

The First-Order Logic of Hyperproperties*

Bernd Finkbeiner¹ and Martin Zimmermann²

- 1 Reactive Systems Group, Saarland University, Saarbrücken, Germany
finkbeiner@react.uni-saarland.de
- 2 Reactive Systems Group, Saarland University, Saarbrücken, Germany
zimmermann@react.uni-saarland.de

Abstract

We investigate the logical foundations of hyperproperties. Hyperproperties generalize trace properties, which are sets of traces, to sets of sets of traces. The most prominent application of hyperproperties is information flow security: information flow policies characterize the secrecy and integrity of a system by comparing two or more execution traces, for example by comparing the observations made by an external observer on execution traces that result from different values of a secret variable.

In this paper, we establish the first connection between temporal logics for hyperproperties and first-order logic. Kamp’s seminal theorem (in the formulation due to Gabbay et al.) states that linear-time temporal logic (LTL) is expressively equivalent to first-order logic over the natural numbers with order. We introduce first-order logic over sets of traces and prove that HyperLTL, the extension of LTL to hyperproperties, is strictly subsumed by this logic. We furthermore exhibit a fragment that is expressively equivalent to HyperLTL, thereby establishing Kamp’s theorem for hyperproperties.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases Hyperproperties, Linear Temporal Logic, First-order Logic

Digital Object Identifier 10.4230/LIPIcs.STACS.2017.30

1 Introduction

Linear-time temporal logic (LTL) [19] is one of the most commonly used logics in model checking [2], monitoring [17], and reactive synthesis [10], and a prime example for the “unusual effectiveness of logic in computer science” [16]. LTL pioneered the idea that the correctness of computer programs should not just be specified in terms of a relation between one-time inputs and outputs, but in terms of the infinite sequences of such interactions captured by the *execution traces* of the program. The fundamental properties of the logic, in particular its ultimately periodic model property [21], and the connection to first-order logic via Kamp’s theorem [18], have been studied extensively and are covered in various handbook articles and textbooks (cf. [7, 22]).

In this paper, we revisit these foundations in light of the recent trend to consider not only the individual traces of a computer program, but properties of *sets* of traces, so-called *hyperproperties* [5]. The motivation for the study of hyperproperties comes from information flow security. Information flow policies characterize the secrecy and integrity of a system by relating two or more execution traces, for example by comparing the observations made by

* This work was partially supported by the German Research Foundation (DFG) under the project “SpAGAT” (FI 936/2-3) in the Priority Program 1496 “Reliably Secure Software Systems” and under the project “TriCS” (ZI 1516/1-1).



an external observer on traces that result from different values of a secret variable. Such a comparison can obviously not be expressed as a property of individual traces, but it can be expressed as a property of the full set of system traces. Beyond security, hyperproperties also occur naturally in many other settings, such as the symmetric access to critical resources in distributed protocols, and Hamming distances between code words in coding theory [12].

HyperLTL [4], the extension of LTL to hyperproperties, uses *trace quantifiers* and *trace variables* to refer to multiple traces at the same time. For example, the formula

$$\forall\pi. \forall\pi'. \mathbf{G} (a_\pi \leftrightarrow a_{\pi'}) \tag{1}$$

expresses that *all* computation traces must *agree* on the value of the atomic proposition a at all times. The extension is useful: it has been shown that most hyperproperties studied in the literature can be expressed in HyperLTL [20]. There has also been some success in extending algorithms for model checking [12], monitoring [1], and satisfiability [11] from LTL to HyperLTL. So far, however, we lack a clear understanding of how deeply the foundations of LTL are affected by the extension. Of particular interest would be a characterization of the models of the logic. Are the models of a satisfiable HyperLTL formula still “simple” in the sense of the ultimately periodic model theorem of LTL?

It turns out that the differences between LTL and HyperLTL are surprisingly profound. Every satisfiable LTL formula has a model that is a (single) ultimately periodic trace. Such models are in particular finite and finitely representable. One might thus conjecture that a satisfiable HyperLTL formula has a model that consists of a finite set of traces, or an ω -regular set of traces, or at least *some* set of ultimately periodic traces. In Section 3, we refute *all* these conjectures. Some HyperLTL formulas have only infinite models, some have only non-regular models, and some have only aperiodic models. We can even encode the prime numbers in HyperLTL!

Is there some way, then, to characterize the expressive power of HyperLTL? For LTL, Kamp’s seminal theorem [18] (in the formulation due to Gabbay et al. [14]) states that LTL is expressively equivalent to first-order logic $FO[<]$ over the natural numbers with order. In order to formulate a corresponding “Kamp’s theorem for HyperLTL,” we have to decide how to encode sets of traces as relational structures, which also induces the signature of the first-order logic we consider. We chose to use relational structures that consist of disjoint copies of the natural numbers with order, one for each trace. To be able to compare positions on different traces, we add the *equal-level predicate* E (cf. [23]), which relates the same time points on different traces. The HyperLTL formula (1), for example, is equivalent to the $FO[<, E]$ formula

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

In Section 4, we show that $FO[<, E]$ is *strictly more expressive* than HyperLTL, i.e., every HyperLTL formula can be translated into an equivalent $FO[<, E]$ formula, but there exist $FO[<, E]$ formulas that cannot be translated to HyperLTL. Intuitively, $FO[<, E]$ can express requirements which relate at some point in time an *unbounded* number of traces, which is not possible in HyperLTL. To obtain a fragment of $FO[<, E]$ that is expressively equivalent to HyperLTL, we must rule out such properties. We consider the fragment where the quantifiers either refer to initial positions or are guarded by a constraint that ensures that the new position is on a trace identified by an initial position chosen earlier. In this way, a formula can only express properties of the bounded number of traces selected by the quantification of initial positions. We call this fragment HyperFO, the *first-order logic of hyperproperties*. Theorem 9, the main result of the paper, then shows that HyperLTL and

HyperFO are indeed expressively equivalent, and thus proves that Kamp's correspondence between temporal logic and first-order logic also holds for hyperproperties.

All proofs omitted due to space restrictions can be found in the full version [13].

2 HyperLTL

Fix a finite set AP of atomic propositions. A trace over AP is a map $t: \mathbb{N} \rightarrow 2^{\text{AP}}$, denoted by $t(0)t(1)t(2)\dots$. The set of all traces over AP is denoted by $(2^{\text{AP}})^\omega$. The projection of t to AP' is the trace $(t(0) \cap \text{AP}')(t(1) \cap \text{AP}')(t(2) \cap \text{AP}')\dots$ over AP'. A trace t is ultimately periodic, if $t = t_0 \cdot t_1^\omega$ for some $t_0, t_1 \in (2^{\text{AP}})^+$, i.e., there are $s, p > 0$ with $t(n) = t(n+p)$ for all $n \geq s$. A set T of traces is ultimately periodic, if every trace in T is ultimately periodic.

The formulas of HyperLTL are given by the grammar

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

where a ranges over atomic propositions in AP and where π ranges over a given countable set \mathcal{V} of *trace variables*. Conjunction, implication, equivalence, and exclusive disjunction \oplus as well as the temporal operators eventually \mathbf{F} and always \mathbf{G} are derived as usual. A sentence is a closed formula, i.e., the formula has no free trace variables.

The semantics of HyperLTL is defined with respect to a trace assignment, a partial mapping $\Pi: \mathcal{V} \rightarrow (2^{\text{AP}})^\omega$. 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 . Furthermore, $\Pi[j, \infty]$ denotes the assignment mapping every π in Π 's domain to $\Pi(\pi)(j)\Pi(\pi)(j+1)\Pi(\pi)(j+2)\dots$.

For sets T of traces and trace-assignments Π we define

- $(T, \Pi) \models a_\pi$, if $a \in \Pi(\pi)(0)$,
- $(T, \Pi) \models \neg \psi$, if $(T, \Pi) \not\models \psi$,
- $(T, \Pi) \models \psi_1 \vee \psi_2$, if $(T, \Pi) \models \psi_1$ or $(T, \Pi) \models \psi_2$,
- $(T, \Pi) \models \mathbf{X} \psi$, if $(T, \Pi[1, \infty]) \models \psi$,
- $(T, \Pi) \models \psi_1 \mathbf{U} \psi_2$, if there is a $j \geq 0$ such that $(T, \Pi[j, \infty]) \models \psi_2$ and for all $0 \leq j' < j$: $(T, \Pi[j', \infty]) \models \psi_1$,
- $(T, \Pi) \models \exists \pi. \varphi$, if there is a trace $t \in T$ such that $(T, \Pi[\pi \rightarrow t]) \models \varphi$, and
- $(T, \Pi) \models \forall \pi. \varphi$, if for all traces $t \in T$: $(T, \Pi[\pi \rightarrow t]) \models \varphi$.

We say that T satisfies a sentence φ , if $(T, \Pi_\emptyset) \models \varphi$. In this case, we write $T \models \varphi$ and say that T is a model of φ . Although HyperLTL sentences are required to be in prenex normal form, they are closed under boolean combinations, which can easily be seen by transforming such formulas into prenex normal form.

3 The Models of HyperLTL

Every satisfiable LTL formula has an ultimately periodic model, i.e., a particularly simple model: It is trivially finite (and finitely represented) and forms an ω -regular language. An obvious question is whether every satisfiable HyperLTL sentence has a simple model, too. Various notions of simplicity could be considered here, e.g., cardinality based ones, being ω -regular, or being ultimately periodic, which all extend the notion of simplicity for the LTL case. In this section, we refute all these possibilities: We show that HyperLTL models have to be in general infinite, might necessarily be non-regular, and may necessarily be aperiodic.

3.1 No Finite Models

Our first result shows that HyperLTL does not have the finite model property (in the sense that every satisfiable sentence is satisfied by a finite set of traces). The proof is a straightforward encoding of an infinite set of traces that appears again in the following proofs.

► **Theorem 1.** *There is a satisfiable HyperLTL sentence that is not satisfied by any finite set of traces.*

Proof. Consider the conjunction φ of the following formulas over $\text{AP} = \{a\}$:

- $\forall\pi. (\neg a_\pi) \mathbf{U} (a_\pi \wedge \mathbf{X} \mathbf{G} \neg a_\pi)$: on every trace there is exactly one occurrence of a .
- $\exists\pi. a_\pi$: there is a trace where a holds true in the first position.
- $\forall\pi. \exists\pi'. \mathbf{F} (a_\pi \wedge \mathbf{X} a_{\pi'})$: for every trace, say where a holds at position n (assuming the first conjunct is satisfied), there is another trace where a holds at position $n + 1$.

It is straightforward to verify that φ is satisfied by the infinite set $T = \{\emptyset^n \cdot \{a\} \cdot \emptyset^\omega \mid n \geq 0\}$ and an induction over n shows that every model has to contain T . Here, one uses the first and second conjunct in the induction start and the first and third conjunct in the induction step. Actually, the first conjunct then implies that T is the only model of φ . ◀

Next, we complement the lower bound with a matching upper bound.

► **Theorem 2.** *Every satisfiable HyperLTL sentence has a countable model.*

Proof. Let φ be a satisfiable HyperLTL sentence and let T be a model. If T is countable, then we are done. Thus, assume T is uncountable and thus in particular non-empty. Furthermore, we assume w.l.o.g. $\varphi = \forall\pi_0. \exists\pi'_0. \dots \forall\pi_k. \exists\pi'_k. \psi$ with quantifier-free ψ .

As T is a model of φ , there is a Skolem function $f_i: T^i \rightarrow T$ for every $i \leq k$ satisfying the following property: $(T, \Pi) \models \psi$ for every trace assignment Π that maps each π_i to some arbitrary $t_i \in T$ and every π'_i to $f_i(t_0, \dots, t_i)$. Note that the relation $(T, \Pi) \models \psi$ does only depend on Π and ψ , but not on T , as ψ is quantifier-free.

Given a subset $S \subseteq T$ and a Skolem function f_i we define

$$f_i(S) = \{f_i(t_0, \dots, t_i) \mid t_0, \dots, t_i \in S\}.$$

Now, fix some $t \in T$. Define $S_0 = \{t\}$ and $S_{n+1} = S_n \cup \bigcup_{i=0}^k f_i(S_n)$ for every n , and $S = \bigcup_{n \geq 0} S_n$. The limit stage S is closed under applying the Skolem functions, i.e., if $t_0, \dots, t_i \in S$, then $f_i(t_0, \dots, t_i) \in S$. Also, every stage S_n is finite by a straightforward induction, hence S is countable. We conclude the proof by showing that S is a model of φ .

Every trace assignment Π mapping π_i to some $t_i \in S$ and every π'_i to $f_i(t_0, \dots, t_i) \in S$ satisfies $(T, \Pi) \models \psi$, as argued above. Also, as argued above, this is independent of T due to ψ being quantifier-free. Hence, we obtain $(S, \Pi) \models \psi$. Finally, a simple induction over the quantifier prefix shows $(S, \Pi_\emptyset) \models \varphi$, i.e., S is indeed a model of φ . ◀

3.2 No Regular Models

The construction presented in the proof of Theorem 1, which pushes a single occurrence of the proposition a through the traces to enforce the set $\{\emptyset^n \cdot \{a\} \cdot \emptyset^\omega \mid n \geq 0\}$ is reused to prove the main result of this subsection. We combine this construction with an inductive swapping construction to show that HyperLTL sentences do not necessarily have ω -regular models. To illustrate the swapping, consider the following finite traces:

$$\begin{array}{ll} t_0 = \{a\} \cdot \emptyset \cdot \{a\} \cdot \emptyset \cdot \{a\} \cdot \emptyset & t_2 = \{a\} \cdot \{a\} \cdot \emptyset \cdot \{a\} \cdot \emptyset \cdot \emptyset \\ t_1 = \{a\} \cdot \{a\} \cdot \emptyset \cdot \emptyset \cdot \{a\} \cdot \emptyset & t_3 = \{a\} \cdot \{a\} \cdot \{a\} \cdot \emptyset \cdot \emptyset \cdot \emptyset \end{array}$$

The trace t_1 is obtained from t_0 by swapping the first occurrence of \emptyset one position to the right (a swap may only occur between adjacent positions, one where a holds and one where it does not). Furthermore, with two more swaps, one turns t_1 into t_2 and t_2 into t_3 .

Our following proof is based on the following three observations: (1) In an alternating sequence of even length such as t_1 , the number of positions where a holds and where a does not hold is equal. Such a sequence is expressible in (Hyper)LTL. (2) A swap does not change this equality and can be formalized in HyperLTL. (3) Thus, if all occurrences of $\{a\}$ are swapped to the beginning, then the trace has the form $\{a\}^n \cdot \emptyset^n$ for some n . Hence, if we start with all alternating sequences as in t_0 , then we end up with the non-regular language $\{\{a\}^n \cdot \emptyset^n \mid n > 0\}$.

► **Theorem 3.** *There is a satisfiable HyperLTL sentence that is not satisfied by any ω -regular set of traces.*

Proof. Consider the conjunction φ of the formulas φ_i , $i \in \{1, \dots, 8\}$ over $\text{AP} = \{a, b, 1, 2, \dagger\}$.

$$\blacksquare \varphi_1 = \forall \pi. (1_\pi \oplus 2_\pi) \wedge \neg \dagger_\pi \wedge \neg \dagger_\pi \quad \mathbf{UG} (\dagger_\pi \wedge \neg a_\pi).$$

Every trace from a set of traces satisfying φ_1 either satisfies 1 or 2 at the first position. Consequently, we speak of traces of type i for $i \in \{1, 2\}$. Also, on every such trace the truth value of \dagger changes exactly once, from false to true, after being false at least at the first position. In the following, we are only interested in the unique maximal prefix of a trace where \dagger does not hold, which we call the *window* of the trace. Note that a may only hold in the window of a trace. Considering windows essentially turns infinite traces into finite ones.

The balance $\text{bal}(t)$ of a trace t is the absolute value of the difference between the number of window positions where a holds and the number of those where a does not hold, i.e.,

$$\text{bal}(t) = |\{n \mid a \in t(n) \text{ and } \dagger \notin t(n)\}| - |\{n \mid a \notin t(n) \text{ and } \dagger \notin t(n)\}|.$$

$$\blacksquare \varphi_2 = \forall \pi. 1_\pi \rightarrow (a_\pi \wedge \mathbf{G}(a_\pi \rightarrow \mathbf{X} \neg a_\pi \wedge \mathbf{X} \neg \dagger_\pi \wedge \mathbf{X} \mathbf{X}(a_\pi \vee \dagger_\pi)))$$

$$\blacksquare \varphi_3 = \exists \pi. 1_\pi \wedge a_\pi \wedge \mathbf{X} \mathbf{X} \dagger_\pi$$

$$\blacksquare \varphi_4 = \forall \pi. \exists \pi'. 1_\pi \rightarrow (1_{\pi'} \wedge \mathbf{F}(\neg \dagger_\pi \wedge \mathbf{X} \dagger_\pi \wedge \mathbf{X} \mathbf{X} \neg \dagger_{\pi'} \wedge \mathbf{X} \mathbf{X} \mathbf{X} \dagger_{\pi'}))$$

If $\varphi_1 \wedge \dots \wedge \varphi_4$ is satisfied by a set of traces, then the projection to $\{a\}$ of the window of every type 1 trace has the form $(\{a\} \cdot \emptyset)^n$ for some $n > 0$, due to φ_2 . In particular, every type 1 trace has balance zero. Furthermore, due to φ_3 and φ_4 , there is a trace with such a window for every $n > 0$.

$$\blacksquare \varphi_5 = \forall \pi. 2_\pi \rightarrow b_\pi \wedge b_\pi \quad \mathbf{UG} \neg b_\pi$$

Finally, φ_5 requires every type 2 trace to have a prefix where b holds true, after which it never holds true again. The length of this prefix is the *rank* of the trace, which is finite.

The next formula implements the swapping process. Each swap has to decrease the rank until a type 1 trace is reached. This rules out models satisfying the formulas by cyclic swaps.

$$\blacksquare \varphi_6 = \forall \pi. \exists \pi'. 2_\pi \rightarrow (\mathbf{F}(\dagger_\pi \wedge \dagger_{\pi'} \wedge \mathbf{X} \neg \dagger_\pi \wedge \mathbf{X} \neg \dagger_{\pi'})) \wedge \varphi_{\text{swp}}(\pi, \pi') \wedge [\\ (1_{\pi'} \wedge b_\pi \wedge \mathbf{X} \neg b_\pi) \vee \\ (2_{\pi'} \wedge \mathbf{F}(b_{\pi'} \wedge \mathbf{X} \neg b_{\pi'} \wedge \mathbf{X} b_\pi \wedge \mathbf{X} \mathbf{X} \neg b_\pi))]$$

where

$$\varphi_{\text{swp}}(\pi, \pi') = (a_\pi \leftrightarrow a_{\pi'}) \mathbf{U} ((a_\pi \oplus \mathbf{X} a_\pi) \wedge (a_{\pi'} \oplus \mathbf{X} a_{\pi'})) \wedge (a_\pi \oplus a_{\pi'}) \wedge \mathbf{X} \mathbf{X} \mathbf{G}(a_\pi \leftrightarrow a_{\pi'}).$$

Intuitively, this formula requires for every trace t of type 2 the existence of a trace t' of the same window length and where the difference in the truth values of a in t and t' is only a single swap at adjacent positions (first line). Furthermore, if t has rank one, then t' has to be of type 1 (line two); otherwise, if t has rank $r > 1$, then t' has to be of type 2 and has to

have rank $r - 1$ (line three). Thus, the rank is an upper bound on the number of swaps that can be executed before a trace of type 1 is reached.

An induction over the rank of type 2 traces shows that every such trace has balance zero, as a swap as formalized by φ_{swp} does not change the balance.

- $\varphi_7 = \exists \pi. 2_\pi \wedge a_\pi$
- $\varphi_8 = \forall \pi. \exists \pi'. 2_\pi \rightarrow (2_{\pi'} \wedge (a_\pi \wedge a_{\pi'})) \mathbf{U} (\mathbf{G} \neg a_\pi \wedge a_{\pi'} \wedge \mathbf{X} \mathbf{G} \neg a_{\pi'})$

The last two formulas imply for every $n > 0$ the existence of a trace of type 2 which has a prefix where a holds true at exactly the first n positions, after which it never holds true again. Due to the balance of type 2 traces being zero (assuming all previous formulas are satisfied), the projection to $\{a\}$ of the window of such a trace has the form $\{a\}^n \cdot \emptyset^n$.

Now, towards a contradiction, assume that $T \models \varphi$ for some ω -regular T . It follows from the observations made above that projecting T to $\{a, \dagger\}$ and intersecting it with the ω -regular language $\{a\}^* \cdot \emptyset^* \cdot \{\dagger\}^\omega$ results in the language $\{\{a\}^n \cdot \emptyset^n \cdot \{\dagger\}^\omega \mid n > 0\}$, which is not ω -regular. This yields the desired contradiction.

To conclude, it suffices to remark that φ is satisfied by taking the union of the set of all required type 1 traces and of the set of all type 2 traces with finite window length, balance zero, and with rank equal to the number of swaps necessary to reach a type 1 trace. ◀

Note that this result can be strengthened by starting with type 1 traces of the form $(\emptyset \cdot \{a\} \cdot \{a'\} \cdot \{a, a'\})^+ \{\dagger\}^\omega$ for some fresh proposition a' and then modify the swap operation to obtain sequences of the form $\emptyset^n \cdot \{a\}^n \cdot \{a'\}^n \cdot \{a, a'\}^n \{\dagger\}^\omega$. These form, when ranging over all n , a non- ω -contextfree language (see [6] for a formal definition of these languages). Thus, not every HyperLTL sentence has an ω -contextfree model.

► **Theorem 4.** *There is a satisfiable HyperLTL sentence that is not satisfied by any ω -contextfree set of traces.*

It is an interesting question to find a non-trivial class of languages that is rich enough for every satisfiable HyperLTL sentence to be satisfied by a model from this class.

3.3 No Periodic Models

Next, we extend the techniques developed in the previous two subsections to show our final result on the complexity of HyperLTL models: although every LTL formula has an ultimately periodic model, one can construct a HyperLTL sentence without ultimately periodic models.

► **Theorem 5.** *There is a satisfiable HyperLTL sentence that is not satisfied by any set of ultimately periodic traces.*

Proof. A trace t is *not* ultimately periodic, if for every $s, p > 0$ there is an $n \geq s$ with $t(n) \neq t(n + p)$. In the following, we construct auxiliary traces that allow us to express this property in HyperLTL. The main difficulty is to construct traces of the form $(\{b\}^p \cdot \emptyset^p)^\omega$ for every p , to implement the quantification of the period length p .

We construct a sentence φ over $\text{AP} = \{a, b, 1, 2, \$\}$ with the desired properties, which is a conjunction of several subformulas. The first conjunct requires every trace in a model of φ to have exactly one occurrence of the proposition a . If it holds at position n , then we refer to $n + 1$ as the *characteristic* of the trace (recall that a trace starts at position 0).

As in the proof of Theorem 3, we have two special types of traces in models of φ , which are identified by either 1 or 2 holding true at the first position of every trace, but there might be other traces as well. Type 1 traces are of the form $\emptyset^c \cdot \{a\} \cdot \emptyset^\omega$ for $c \geq 0$. As in the proof

of Theorem 1, one can construct a conjunct that requires the models of φ to contain a type 1 trace for every such c , but no other traces of type 1.

The projection to $\{b\}$ of a trace t of type 2 is a suffix of $(\{b\}^c \cdot \emptyset^c)^\omega$, where c is the characteristic of t . We claim that one can construct a conjunct of φ that requires all models of φ to contain all these type 2 traces, i.e., all possible suffixes for every $c > 0$. This is achieved by formalizing the following properties in HyperLTL:

1. Every type 2 trace has infinitely many positions where b holds and infinitely many positions where b does not hold. A block of such a trace is a maximal infix whose positions coincide on their truth values of b , i.e., either b holds at every position of the infix, but not at the last one before the infix (if it exists) and not at the first position after the infix or b does not hold at every position of the infix, but at the last one before it (if it exists) and at the first position after it.
2. For every type 1 trace there is at least one type 2 trace of the same characteristic.
3. The length of the first block of every type 2 trace is not larger than its characteristic.
4. If a block ends at the unique position of a type 2 trace where its a holds, then it has to be the first block.
5. For every type 2 trace there is another one of the same characteristic that is obtained by shifting the truth values of b one position to the left.

Assume a set T of traces satisfies all these properties and assume there is a type 2 trace $t \in T$ whose projection to $\{b\}$ is not a suffix of $(\{b\}^c \cdot \emptyset^c)^\omega$, where c is the characteristic of t . The length of its first block is bounded by c , due to the third property. Thus, there has to be a non-first block whose length ℓ is not equal to c . If $\ell > c$, we can use the fifth property to shift this block to the left until we obtain a type 2 trace of characteristic c in T whose first block has the same length ℓ . This trace violates the third property. If $\ell < c$, then we can again shift this block to the left until we obtain a trace in T of characteristic c that has a block of length ℓ that ends at the unique position where a holds. Due to $\ell < c$, this cannot be the first block, i.e., we have derived a contradiction to the fourth property.

On the other hand, for every $c > 0$, there is a some type 2 trace of characteristic c in T . As shown above, its projection to $\{b\}$ is a suffix of $(\{b\}^c \cdot \emptyset^c)^\omega$. Thus, applying the left-shift operation $2c - 1$ times yields all possible suffixes of $(\{b\}^c \cdot \emptyset^c)^\omega$. Thus, T does indeed contain all possible type 2 traces, if it satisfies the formulas described above.

Recall that we have to express the following property: there is a trace t such that for every $s, p > 0$ there is an $n \geq s$ with $t(n) \neq t(n + p)$. To this end, we first existentially quantify a trace π (the supposedly non-ultimately periodic one). Then, we universally quantify two type 1 traces π_s and π_p (thereby fixing s and p as the characteristics of π_s and π_p). Thus, it remains to state that π has two positions n and n' satisfying $s \leq n < n' = n + p$ such that the truth value of $\$$ differs at these positions. To this end, we need another trace π'_p of the same characteristic p as π_p , so that a block of π'_p starts at position n , which allows to determine $n' = n + p$ by just advancing to the end of the block starting at n .

Formally, consider the following statement: there is a trace π such that for all type 1 traces π_s and π_p (here, we quantify over s and p) there is a type 2 trace π'_p that has the same characteristic as π_p such that the following is true: there is a position n no earlier than the one where a holds in π_s such that

- the truth value of b in π'_p differs at positions $n - 1$ and n (i.e., a block begins at n), and
- the atomic proposition $\$$ holds at n in π and not at n' in π or vice versa, where $n' > n$ is the smallest position such that the truth value of b in π'_p differs at $n' - 1$ and n' (i.e., the next block begins at position n'), which implies $n' = n + p$.

The formalization of this statement in HyperLTL is the final conjunct of φ . Hence, φ has no models that contain an ultimately periodic trace.

Finally, φ is satisfied by all models that contain all possible type 1 and all possible type 2 traces as well as at least one trace that is not ultimately periodic when projected to $\{\$$. ◀

Note that the type 1 and type 2 traces above are ultimately periodic, i.e., although we have formalized the existence of a single non-ultimately periodic trace, the model always has ultimately periodic ones as well. By slightly extending the construction, one can even construct a satisfiable sentence whose models contain not a single ultimately periodic trace. To this end, one requires that every trace (in particular the type 1 and type 2 traces) is non-ultimately periodic, witnessed by the proposition $\$$ as above.

► **Theorem 6.** *There is a satisfiable HyperLTL sentence that is not satisfied by any set of traces that contains an ultimately periodic trace.*

As a final note on the expressiveness of HyperLTL we show how to encode the prime numbers. Let type 1 and type 2 traces be axiomatized as in the proof of Theorem 5. Recall projecting a type 2 trace to $\{b\}$ yields a suffix of $(\{b\}^c \cdot \emptyset^c)^\omega$, where $c > 0$ is the trace's characteristic. We say that such a trace is *proper*, if its projection equal to $(\{b\}^c \cdot \emptyset^c)^\omega$. Being proper can be expressed in HyperLTL, say by the formula $\varphi_{\text{prp}}(\pi)$ with a single free variable, relying on the fact that the only occurrence of a induces the characteristic c . Also, we add a new atomic proposition $'$ to AP to encode the prime numbers as follows: the proposition $'$ holds at the first position of a type 1 trace of characteristic c if, and only if, c is a prime number.

Now, consider the following formula, which we add as a new conjunct to the axiomatization of type 1 and type 2 traces:

$$\begin{aligned} & \forall \pi^1. \forall \pi^2. (1_{\pi^1} \wedge '_{\pi^1} \wedge \varphi_{\text{prp}}(\pi^2) \rightarrow \neg \psi(\pi^1, \pi^2)) \wedge \\ & \forall \pi^1. \exists \pi^2. (1_{\pi^1} \wedge \neg '_{\pi^1} \rightarrow \varphi_{\text{prp}}(\pi^2) \wedge \psi(\pi^1, \pi^2)) \end{aligned}$$

Here, the formula $\psi(\pi^1, \pi^2)$ expresses that the single a in π^1 appears at the end of a non-first block in π^2 and that the characteristic of π^2 is strictly greater than one. Thus, $\psi(\pi^1, \pi^2)$ holds if, and only if, the characteristic of π^2 is a non-trivial divisor of the characteristic of π^1 . Thus, the first conjunct expresses that a type 1 trace of characteristic $c > 1$ may only have a $'$ at the first position, if c has only trivial divisors, i.e., if c is prime. Similarly, the second conjunct expresses that a type 1 trace of characteristic $c > 1$ may only not have a $'$ at the first position, if c has a non-trivial divisor, i.e., if c is not prime. Thus, by additionally hardcoding that 1 is not a prime, one obtains a formula φ such that every model T of φ encodes the primes as follows: c is prime if, and only if, there is a type 1 trace of characteristic c in T with $'$ holding true at its first position.

4 First-order Logic for Hyperproperties

Kamp's seminal theorem [18] states that Linear Temporal Logic with the until-operator \mathbf{U} and its dual past-time operator “since” is expressively equivalent to first-order logic over the integers with order, $\text{FO}[\lt]$ for short. Later, Gabbay et al. [14] proved that LTL as introduced here (i.e., exclusively with future-operators) is expressively equivalent to first-order logic over the *natural numbers* with order. More formally, one considers relational structures of the form $(\mathbb{N}, \lt, (P_a)_{a \in \text{AP}})$ where \lt is the natural ordering of \mathbb{N} and each P_a is a subset of \mathbb{N} . There is a bijection mapping a trace t over AP to such a structure \underline{t} . Furthermore, $\text{FO}[\lt]$ is

first-order logic¹ over the signature $\{<\} \cup \{P_a \mid a \in \text{AP}\}$ with equality. The result of Gabbay et al. follows from the existence of the following effective translations: (1) For every LTL formula φ there is an FO[<] sentence φ' such that for all traces t : $t \models \varphi$ if, and only if, $\underline{t} \models \varphi'$. (2) For every FO[<] sentence φ there is an LTL formula φ' such that for all traces t : $\underline{t} \models \varphi$ if, and only if, $t \models \varphi'$.

In this section, we investigate whether there is a first-order logic that is expressively equivalent to HyperLTL. The first decision to take is how to represent a set of traces as a relational structure. The natural approach is to take disjoint copies of the natural numbers, one for each trace and label them accordingly. Positions on these traces can be compared using the order. To be able to compare different traces, we additionally introduce a (commutative) equal-level predicate E, which relates the same time points on different traces.

Formally, given a set $T \subseteq (2^{\text{AP}})^\omega$ of traces over AP, we define the relational structure $\underline{T} = (T \times \mathbb{N}, <^{\underline{T}}, E^{\underline{T}}, (P_a^{\underline{T}})_{a \in \text{AP}})$ with

- $<^{\underline{T}} = \{((t, n), (t, n')) \mid t \in T \text{ and } n < n' \in \mathbb{N}\}$,
- $E^{\underline{T}} = \{((t, n), (t', n)) \mid t, t' \in T \text{ and } n \in \mathbb{N}\}$, and
- $P_a^{\underline{T}} = \{(t, n) \mid a \in t(n)\}$.

We consider first-order logic over the signature $\{<, E\} \cup \{P_a \mid a \in \text{AP}\}$, i.e., with atomic formulas $x = y$, $x < y$, $E(x, y)$, and $P_a(x)$ for $a \in \text{AP}$, and disjunction, conjunction, negation, and existential and universal quantification over elements. We denote this logic by FO[<, E]. We use the shorthand $x \leq y$ for $x < y \vee x = y$ and freely use terms like $x \leq y < z$ with the obvious meaning. A sentence is a closed formula, i.e., every occurrence of a variable is in the scope of a quantifier binding this variable. We write $\varphi(x_0, \dots, x_n)$ to denote that the free variables of the formula φ are among x_0, \dots, x_n .

► **Example 7.**

1. The formula $\text{Succ}(x, y) = x < y \wedge \neg \exists z. x < z < y$ expresses that y is the direct successor of x on some trace.
2. The formula $\text{min}(x) = \neg \exists y. \text{Succ}(y, x)$ expresses that x is the first position of a trace.

Our first result shows that full FO[<, E] is too expressive to be equivalent to HyperLTL. To this end, we apply a much stronger result due to Bozzelli et al. [3] showing that a certain property expressible in KLTL (LTL with the epistemic knowledge operator **K** [9]) is not expressible in HyperCTL*, which subsumes HyperLTL.

► **Theorem 8.** *There is an FO[<, E] sentence φ that has no equivalent HyperLTL sentence: For every HyperLTL sentence φ' there are two sets T_0 and T_1 of traces such that*

1. $\underline{T_0} \not\models \varphi$ and $\underline{T_1} \models \varphi$, but
2. φ' cannot distinguish T_0 and T_1 , i.e., either both $T_0 \models \varphi'$ and $T_1 \models \varphi'$ or both $T_0 \not\models \varphi'$ and $T_1 \not\models \varphi'$.

Proof. Fix $\text{AP} = \{p\}$ and consider the following property of sets T of traces over AP: there is an $n > 0$ such that $p \notin t(n)$ for every $t \in T$. This property is expressible in FO[<, E], but Bozzelli et al. [3] proved that it is not expressible in HyperLTL by constructing sets T_0, T_1 of traces with the desired property.² ◀

¹ We assume familiarity with the syntax and semantics of first-order logic. See, e.g., [8], for an introduction to the topic.

² Actually, they proved a stronger result showing that the property cannot be expressed in HyperCTL*, which subsumes HyperLTL. As the latter logic is a branching-time logic, they actually constructed Kripke structures witnessing their result. However, it is easy to show that taking the languages of traces of these Kripke structures proves our claim.

As already noted by Bozzelli et al., the underlying insight is that HyperLTL cannot express requirements which relate at some point in time an unbounded number of traces. By ruling out such properties, we obtain a fragment of $\text{FO}[\langle, E]$ that is equivalent to HyperLTL. Intuitively, we mimic trace quantification of HyperLTL by quantifying initial positions and then only allow quantification of potentially non-initial positions on the traces already quantified. Thus, such a sentence can only express properties of the bounded number of traces selected by the quantification of initial positions.

To capture this intuition, we have to introduce some notation: $\exists^M x. \varphi$ is shorthand for $\exists x. \min(x) \wedge \varphi$ and $\forall^M x. \varphi$ is shorthand for $\forall x. \min(x) \rightarrow \varphi$, i.e., the quantifiers \exists^M and \forall^M only range over the first positions of a trace in \underline{T} . We use these quantifiers to mimic trace quantification in HyperLTL.

Furthermore, $\exists^G y \geq x. \varphi$ is shorthand for $\exists y. y \geq x \wedge \varphi$ and $\forall^G y \geq x. \varphi$ is shorthand for $\forall y. y \geq x \rightarrow \varphi$, i.e., the quantifiers \exists^G and \forall^G are guarded by a free variable x and range only over greater-or-equal positions on the same trace that x is on. We call the free variable x the *guard* of the quantifier.

We consider sentences of the form

$$\varphi = Q_1^M x_1 \cdots Q_k^M x_k. Q_1^G y_1 \geq x_{g_1} \cdots Q_\ell^G y_\ell \geq x_{g_\ell}. \psi \quad (2)$$

with $Q \in \{\exists, \forall\}$, where we require the sets $\{x_1, \dots, x_k\}$ and $\{y_1, \dots, y_\ell\}$ to be disjoint, every guard x_{g_j} to be in $\{x_1, \dots, x_k\}$, and ψ to be quantifier-free with free variables among the $\{y_1, \dots, y_\ell\}$. We call this fragment HyperFO. Note that the subformula starting with the quantifier Q_1^G being in prenex normal form and ψ only containing the variables y_j simplifies our reasoning later on, but is not a restriction.

► **Theorem 9.** *HyperLTL and HyperFO are equally expressive.*

We prove this result by presenting effective translations between HyperLTL and HyperFO (see Lemma 12 and Lemma 13). We begin with the direction from HyperFO to HyperLTL. Consider a HyperFO sentence φ as in (2). It quantifies k traces with the quantifiers \exists^M and \forall^M . Every other quantification is then on one of these traces. As trace quantification is possible in HyperLTL, we only have to take care of the subformula starting with the guarded quantifiers. After replacing these quantifiers by unguarded ones, we only have to remove the equal-level predicate to obtain an $\text{FO}[\langle]$ sentence. To this end, we merge the k traces under consideration into a single one, which reduces the equal-level predicate to the equality predicate (cf. [23]). The resulting sentence is then translated into LTL using the theorem of Gabbay et al., the merging is undone, and the quantifier prefix is added again. We show that the resulting sentence is equivalent to the original one.

Fix a HyperFO sentence φ as in (2) and consider the subformula

$$\chi = Q_1^G y_1 \geq x_{g_1} \cdots Q_\ell^G y_\ell \geq x_{g_\ell}. \psi$$

obtained by removing the quantification of the guards. We execute the following replacements to obtain the formula χ_m :

1. Replace every guarded existential quantification $\exists^G y_j \geq x_{g_j}$ by $\exists y_j$ and every guarded universal quantification $\forall^G y_j \geq x_{g_j}$ by $\forall y_j$.
2. Replace every atomic formula $P_a(y_j)$ by $P_{(a, g_j)}(y_j)$, where x_{g_j} is the guard of y_j .
3. Replace every atomic formula $E(y_j, y_{j'})$ by $y_j = y_{j'}$.

As we have removed all occurrences of the free guards, the resulting formula χ_m is actually a sentence over the signature $\{\langle\} \cup \{P_a \mid a \in \text{AP} \times \{1, \dots, k\}\}$, i.e., an $\text{FO}[\langle]$ sentence.

Given a list (t_1, \dots, t_k) of traces over AP, define the trace $\text{mrg}(t_1, \dots, t_k) = A_0 A_1 A_2 \cdots$ over $\text{AP} \times \{1, \dots, k\}$ via $A_n = \bigcup_{j=1}^k t_j(n) \times \{j\}$, i.e., we merge the t_j into a single trace.

► **Claim 10.** *Let T be a set of traces and let $\beta_0: \{x_1, \dots, x_k\} \rightarrow T \times \{0\}$ be a variable valuation of the guards x_1, \dots, x_k to elements of \underline{T} . Then, $(\underline{T}, \beta_0) \models \chi$ if, and only if, $\text{mrg}(t_1, \dots, t_k) \models \chi_m$, where t_j is the unique trace satisfying $\beta_0(x_{g_j}) = (t_j, 0)$.*

This claim can be proven by translating a winning strategy for either player in the model checking game [15] for (\underline{T}, χ) (starting with the initial variable valuation β_0) into a winning strategy for the same player in the model checking game for $(\text{mrg}(t_1, \dots, t_k), \chi_m)$.

Now, we apply the theorem of Gabbay et al. [14] to χ_m and obtain an LTL formula χ'_m over $\text{AP} \times \{1, \dots, k\}$ that is equivalent to χ_m . Let χ' be the HyperLTL formula obtained from χ'_m by replacing every atomic proposition (a, j) by a_{π_j} , i.e., we undo the merging. The following claim is proven by a simple structural induction over χ_m .

► **Claim 11.** *Let T be a set of traces and let $\Pi: \{\pi_1, \dots, \pi_k\} \rightarrow T$ be a trace assignment. Then, $\text{mrg}(\Pi(\pi_1), \dots, \Pi(\pi_k)) \models \chi'_m$ if, and only if, $(T, \Pi) \models \chi'$.*

Now, we add the quantifier prefix $Q_1\pi_1 \dots Q_k\pi_k$ to χ' , where $Q_j = \exists$, if $Q_j^M = \exists^M$, and $Q_j = \forall$, if $Q_j^M = \forall^M$. Call the obtained HyperLTL sentence φ' .

► **Lemma 12.** *For every HyperFO sentence φ , there is a HyperLTL sentence φ' such that for every $T \subseteq (2^{\text{AP}})^\omega$: $\underline{T} \models \varphi$ if, and only if, $T \models \varphi'$.*

Proof. Fix a HyperFO sentence φ and let the χ , χ_m , χ'_m , χ' , and φ' be as constructed as above. Let β_0 be a variable valuation as in Claim 10, let the traces $t_1, \dots, t_k \in T$ be defined as in this claim, and let the trace assignment Π map π_j to t_j .

Then, the following equivalences hold:

$$(\underline{T}, \beta_0) \models \chi \stackrel{\text{Claim 10}}{\Leftrightarrow} \text{mrg}(t_1, \dots, t_k) \models \chi_m \stackrel{\text{by def.}}{\Leftrightarrow} \text{mrg}(t_1, \dots, t_k) \models \chi'_m \stackrel{\text{Claim 11}}{\Leftrightarrow} (T, \Pi) \models \chi'.$$

Finally, the equivalence of φ and φ' follows from the fact that one can identify quantification of initial elements of paths in \underline{T} and trace quantification in T , as both φ and φ' have the *same* quantifier prefix. ◀

It remains to consider the translation of HyperLTL into HyperFO, which is straightforward, as usual.

► **Lemma 13.** *For every HyperLTL sentence φ , there is a HyperFO sentence φ' such that for every $T \subseteq (2^{\text{AP}})^\omega$: $T \models \varphi$ if, and only if, $\underline{T} \models \varphi'$.*

Proof. Let π_1, \dots, π_k be the trace variables appearing in φ and fix a set $G = \{x_1, \dots, x_k, x_t\}$ of first-order variables, which we use as guards: the x_j with $j \leq k$ are identified with the trace variables and we use variables guarded by x_t to model the flow of time. We inductively construct a formula $\text{fo}(\varphi)$ satisfying the following invariant: For each subformula ψ of φ , the free variables of the formula $\text{fo}(\psi)$ comprise of a subset of G and one additional (different!) variable, which we call the time-variable of $\text{fo}(\psi)$. We require the time-variables of the subformulas to be fresh unless stated otherwise and also different from the guards in G . Intuitively, the time-variables are used to mimic the flow of time when translating a temporal operator. Formally, we define:

- $\text{fo}(a_{\pi_j}) = \exists^G y \geq x_j. E(y, z) \wedge P_a(y)$, i.e., z is the time-variable of $\text{fo}(a_{\pi_j})$.
- $\text{fo}(\neg\psi_1) = \neg\text{fo}(\psi_1)$, i.e., the time-variable is unchanged.
- $\text{fo}(\psi_1 \vee \psi_2) = \text{fo}(\psi'_1) \vee \text{fo}(\psi_2)$, where we assume w.l.o.g. that $\text{fo}(\psi_1)$ and $\text{fo}(\psi'_2)$ have the same time-variable, which is also the time-variable of the disjunction.
- $\text{fo}(\mathbf{X}\psi_1) = \exists^G z_1 \geq x_t. \text{Succ}(z, z_1) \wedge \text{fo}(\psi_1)$, where z_1 is the time-variable of $\text{fo}(\psi_1)$. Hence, z is the time-variable of $\text{fo}(\mathbf{X}\psi_1)$.

- $\text{fo}(\psi_1 \mathbf{U} \psi_2) = \exists^G z_2 \geq x_t. z \leq z_2 \wedge \text{fo}(\psi_2) \wedge \forall^G z_1 \geq x_t. z \leq z_1 < z_2 \rightarrow \text{fo}(\psi_1)$, where z_i is the time-variable of $\text{fo}(\psi_i)$. Hence, z is the time-variable of $\text{fo}(\psi_1 \mathbf{U} \psi_2)$.
- $\text{fo}(\exists \pi_j. \psi) = \exists^M x_j. \text{fo}(\psi)$, i.e., the time-variable is unchanged.
- $\text{fo}(\forall \pi_j. \psi) = \forall^M x_j. \text{fo}(\psi)$, i.e., the time-variable is unchanged.

Now, we define $\varphi' = \exists^M x_t. \exists^M z. x_t = z \wedge \text{fo}(\varphi)$, where z is the time-variable of $\text{fo}(\varphi)$. It is straightforward to show that φ' is equivalent to φ . Finally, φ' can be rewritten into prenex normal form (with quantifiers Q^M and $Q^G!$) so that the outermost quantifiers bind the guards while the inner ones are guarded. ◀

5 Conclusion and Discussion

The extension from LTL to HyperLTL has fundamentally changed the models of the logic. While a satisfiable LTL formula is guaranteed to have an ultimately periodic model, we have shown that there is no guarantee that a satisfiable HyperLTL formula has a model that is finite, ω -regular, or even just ω -contextfree. Characterizing the expressive power of HyperLTL is thus a formidable challenge. Nevertheless, the results of this paper provide a first such characterization. With the definition of $FO[<, E]$ and HyperFO, and the resulting formulation and proof of Kamp's theorem for hyperproperties, we have established the first connection between temporal logics for hyperproperties and first-order logic. This connection provides a strong basis for a systematic exploration of the models of hyperproperties.

While hyperproperties have recently received a lot of attention from a practical perspective (cf. [1, 4, 12]), their logical and language-theoretic foundations are far less understood, and it is our hope that this paper will attract more research into this exciting area. An important open problem is to find a non-trivial class of languages so that every satisfiable HyperLTL formula is guaranteed to be satisfied by a model from this class. In Section 3, we have ruled out some of the obvious candidates for such a class of languages, such as the ω -regular and ω -contextfree languages. The challenge remains to identify a class of languages that is rich enough for every satisfiable HyperLTL formula.

Another major open problem is to find a temporal logic that is expressively equivalent to $FO[<, E]$. In Section 4, we have shown that HyperLTL is less expressive than $FO[<, E]$, by arguing that HyperLTL cannot express requirements which relate at some point in time an unbounded number of traces. Since KLTL [9] can express such properties, KLTL and related epistemic temporal logics are natural candidates for logics that are expressively equivalent to $FO[<, E]$. Another promising candidate is HyperLTL with past operators, motivated by the results on HyperCTL* with past [3].

Acknowledgements. We thank Markus N. Rabe and Leander Tentrup for fruitful discussions.

References

- 1 Shreya Agrawal and Borzoo Bonakdarpour. Runtime verification of k -safety hyperproperties in HyperLTL. In *CSF 2016*, pages 239–252. IEEE Computer Society, 2016. doi:10.1109/CSF.2016.24.
- 2 Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. The MIT Press, 2008.
- 3 Laura Bozzelli, Bastien Maubert, and Sophie Pinchinat. Unifying hyper and epistemic temporal logics. In A. M. Pitts, editor, *FoSSaCS 2015*, volume 9034 of *Lecture Notes in Computer Science*, pages 167–182. Springer, 2015. doi:10.1007/978-3-662-46678-0_11.

- 4 Michael R. Clarkson, Bernd Finkbeiner, Masoud Koleini, Kristopher K. Micinski, Markus N. Rabe, and César Sánchez. Temporal logics for hyperproperties. In Martín Abadi and Steve Kremer, editors, *POST 2014*, volume 8414 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2014. doi:10.1007/978-3-642-54792-8_15.
- 5 Michael R. Clarkson and Fred B. Schneider. Hyperproperties. *Journal of Computer Security*, 18(6):1157–1210, 2010. doi:10.3233/JCS-2009-0393.
- 6 Rina S. Cohen and Arie Y. Gold. Theory of omega-languages. I. characterizations of omega-context-free languages. *Journal of Computer and System Sciences*, 15(2):169–184, 1977. doi:10.1016/S0022-0000(77)80004-4.
- 7 Stéphane Demri, Valentin Goranko, and Martin Lange. *Temporal Logics in Computer Science*. Cambridge University Press, 2016.
- 8 Heinz-Dieter Ebbinghaus, Jörg Flum, and Wolfgang Thomas. *Mathematical Logic (2. ed.)*. Undergraduate texts in mathematics. Springer, 1994.
- 9 Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- 10 Bernd Finkbeiner. Synthesis of reactive systems. In Javier Esparza, Orna Grumberg, and Salomon Sickert, editors, *Dependable Software Systems Engineering*, volume 45 of *NATO Science for Peace and Security Series – D: Information and Communication Security*, pages 72–98. IOS Press, 2016. doi:10.3233/978-1-61499-627-9-72.
- 11 Bernd Finkbeiner and Christopher Hahn. Deciding hyperproperties. In Josée Desharnais and Radha Jagadeesan, editors, *CONCUR 2016*, volume 59 of *LIPICs*, pages 13:1–13:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICs.CONCUR.2016.13.
- 12 Bernd Finkbeiner, Markus N. Rabe, and César Sánchez. Algorithms for model checking HyperLTL and HyperCTL*. In Daniel Kroening and Corina S. Pasareanu, editors, *CAV 2015 (Part I)*, volume 9206 of *Lecture Notes in Computer Science*, pages 30–48. Springer, 2015. doi:10.1007/978-3-319-21690-4_3.
- 13 Bernd Finkbeiner and Martin Zimmermann. The first-order logic of hyperproperties. *arXiv*, 1610.04388, 2016. URL: <http://arxiv.org/abs/1610.04388>.
- 14 Dov M. Gabbay, Amir Pnueli, Saharon Shelah, and Jonathan Stavi. On the temporal basis of fairness. In Paul W. Abrahams, Richard J. Lipton, and Stephen R. Bourne, editors, *POPL 1980*, pages 163–173. ACM Press, 1980. doi:10.1145/567446.567462.
- 15 Erich Grädel, Phokion G. Kolaitis, Leonid Libkin, Maarten Marx, Joel Spencer, Moshe Y. Vardi, Yde Venema, and Scott Weinstein. *Finite Model Theory and Its Applications*. Springer, 2005.
- 16 Joseph Y. Halpern, Robert Harper, Neil Immerman, Phokion G. Kolaitis, Moshe Y. Vardi, and Victor Vianu. On the unusual effectiveness of logic in computer science. *Bulletin of Symbolic Logic*, 7(2):213–236, 2001. doi:10.2307/2687775.
- 17 Klaus Havelund and Grigore Rosu. Efficient monitoring of safety properties. *International Journal on Software Tools for Technology Transfer*, 6(2):158–173, 2004. doi:10.1007/s10009-003-0117-6.
- 18 Hans W. Kamp. *Tense Logic and the Theory of Linear Order*. PhD thesis, Computer Science Department, University of California at Los Angeles, USA, 1968.
- 19 Amir Pnueli. The Temporal Logic of Programs. In *FOCS 1977*, pages 46–57, 1977.
- 20 Markus N. Rabe. *A Temporal Logic Approach to Information-flow Control*. PhD thesis, Saarland University, 2016.
- 21 A. Prasad Sistla and Edmund M. Clarke. The complexity of propositional linear temporal logics. *Journal of the ACM*, 32(3):733–749, 1985. doi:10.1145/3828.3837.

30:14 The First-Order Logic of Hyperproperties

- 22 Wolfgang Thomas. Languages, automata, and logic. In Grzegorz Rozenberg and Arto Salomaa, editors, *Handbook of Formal Languages, Vol. 3*, pages 389–455. Springer, 1997. doi:10.1007/978-3-642-59126-6.
- 23 Wolfgang Thomas. Path logics with synchronization. In Kamal Lodaya, Madhavan Mukund, and R. Ramanujam, editors, *Perspectives in Concurrency Theory*, pages 469–481. IARCS-Universities, Universities Press, 2009.