The First-Order Logic of Hyperproperties*

Bernd Finkbeiner and Martin Zimmermann
Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany

Linear-time temporal logic (LTL) is one of the most commonly used logics in model checking, monitoring, and reactive synthesis, and a prime example for the "unusal effectiveness of logic in computer science". 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, and the connection to first-order logic via Kamp's theorem [5], have been studied extensively and are covered in handbook articles and textbooks.

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* [2]. 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 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.

HyperLTL [1], 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 $\forall \pi.\ \forall \pi'.\ \mathbf{G}\,(a_\pi \leftrightarrow a_{\pi'})$ 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. There has also been some success in extending algorithms for model checking, monitoring, and satisfiability 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. We refute *all* these conjectures. Some HyperLTL formulas have only infinite models, some have only non-regular models, and

* Full version appeared at STACS 2017 [3].

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 [5] (in the formulation due to Gabbay et al. [4]) 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, which relates the same time points on different traces. The HyperLTL formula from above, for example, is equivalent to the $FO[<, \mathrm{E}]$ formula

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

We show that $FO[<, \mathrm{E}]$ is *strictly more expressive* than HyperLTL, i.e., every HyperLTL formula can be translated into an equivalent $FO[<, \mathrm{E}]$ formula, but there exist $FO[<, \mathrm{E}]$ formulas that cannot be translated to HyperLTL. Intuitively, $FO[<, \mathrm{E}]$ 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 $FO[<, \mathrm{E}]$ 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*. Our main result 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.

References

1. 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.
2. Michael R. Clarkson and Fred B. Schneider. Hyperproperties. *Journal of Computer Security*, 18(6):1157–1210, 2010.
3. Bernd Finkbeiner and Martin Zimmermann. The first-order logic of hyperproperties. In Heribert Vollmer and Brigitte Vallée, editors, *STACS 2017*, volume 66 of *LIPIcs*, pages 30:1–30:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
4. 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.
5. Hans W. Kamp. *Tense Logic and the Theory of Linear Order*. PhD thesis, Computer Science Department, University of California at Los Angeles, USA, 1968.