Automata and Fixpoints for Asynchronous Hyperproperties

JENS OLIVER GUTSFELD, Westfälische Wilhelms Universität Münster, Germany

MARKUS MÜLLER-OLM, Westfälische Wilhelms Universität Münster, Germany

CHRISTOPH OHREM, Westfälische Wilhelms Universität Münster, Germany

Hyperproperties have received increasing attention in the last decade due to their importance e.g. for security analyses. Past approaches have focussed on synchronous analyses, i.e. techniques in which different paths are compared lockstepwise. In this paper, we systematically study asynchronous analyses for hyperproperties by introducing both a novel automata model (Alternating Asynchronous Parity Automata) and the temporal fixpoint calculus H_μ , the first fixpoint calculus that can systematically express hyperproperties in an asynchronous manner and at the same time subsumes the existing logic HyperLTL. We show that the expressive power of both models coincides over fixed path assignments. The high expressive power of both models is evidenced by the fact that decision problems of interest are highly undecidable, i.e. not even arithmetical. As a remedy, we propose approximative analyses for both models that also induce natural decidable fragments.

CCS Concepts: • **Theory of computation** → **Automata over infinite objects; Modal and temporal logics; Verification by model checking.**

Additional Key Words and Phrases: Asynchronicity, Fixpoint calculus, Hyperproperties

ACM Reference Format:

Jens Oliver Gutsfeld, Markus Müller-Olm, and Christoph Ohrem. 2021. Automata and Fixpoints for Asynchronous Hyperproperties. *Proc. ACM Program. Lang.* 5, POPL, Article 38 (January 2021), 29 pages. <https://doi.org/10.1145/3434319>

1 INTRODUCTION

Hyperproperties [Clarkson and Schneider 2010] are a recent innovation in theoretical computer science. While a traditional trace property (like liveness or safety) refers to single traces, a hyperproperty refers to *sets of traces*. Hyperproperties of interest include security properties like non-interference or observational determinism since it can only be inferred from combinations of traces and their relation to each other whether a system fulfills these properties. Analysis methods for hyperproperties have been proposed in many contexts, including abstract interpretation [Mastroeni and Pasqua 2017, 2018], runtime verification [Finkbeiner et al. 2019], synthesis [Finkbeiner et al. 2020] and model checking [Clarkson et al. 2014; Finkbeiner et al. 2015; Gutsfeld et al. 2020a; Rabe 2016]. In model checking, several temporal logics for hyperproperties have been proposed, including *hyperized* variants of LTL [Clarkson et al. 2014; Finkbeiner et al. 2015; Rabe 2016], CTL* [Clarkson et al. 2014; Finkbeiner et al. 2015; Rabe 2016], QPTL [Coenen et al. 2019; Rabe 2016] and PDL- Δ [Gutsfeld et al. 2020a]. In all these logics, specifications are *synchronous*, i.e. the modalities

Authors' addresses: Jens Oliver Gutsfeld, Institut für Informatik, Westfälische Wilhelms Universität Münster, Einsteinstraße 62, Münster, North Rhine-Westphalia, 48149, Germany, jens.gutsfeld@uni-muenster.de; Markus Müller-Olm, Institut für Informatik, Westfälische Wilhelms Universität Münster, Einsteinstraße 62, Münster, North Rhine-Westphalia, 48149, Germany, markus.mueller-olm@uni-muenster.de; Christoph Ohrem, Institut für Informatik, Westfälische Wilhelms Universität Münster, Einsteinstraße 62, Münster, North Rhine-Westphalia, 48149, Germany, christoph.ohrem@uni-muenster.de.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2021 Copyright held by the owner/author(s).

2475-1421/2021/1-ART38

<https://doi.org/10.1145/3434319>

only allow for lockstepwise traversal of different paths. The same is true of the automata-theoretic frameworks underlying the algorithms for these logics. However, the restriction to synchronous traversal of traces is a conceptual limitation of existing approaches [Finkbeiner 2017] that seems rather artificial. Arguably, the ability to accommodate *asynchronous* specifications is an important requirement for hyperproperty verification in various scenarios since many interesting properties require comparing traces at different points of time. For instance, synchronous formulations of information flow security properties such as non-interference are often too strict for system abstractions with varying granularity of steps. Here, a proper security analysis requires asynchronicity in order to match points of interest. Asynchronous hyperproperties also arise naturally in the context of multithreaded programs where each thread is represented by a single trace as the overall system behaviour is an asynchronous interleaving of the individual traces.

In order to investigate the boundaries of automatic analysis of asynchronous hyperproperties induced by undecidability and complexity limitations, we propose both an automata-theoretic framework, Alternating Asynchronous Parity Automata (AAPA), and a temporal fixpoint calculus, H_μ , for asynchronous hyperproperties in this paper. Our contribution is threefold: first of all, we show that both perspectives indeed coincide over fixed sets of paths by providing a direct and intuitive translation between AAPA and H_μ formulas. Secondly, using this correspondence, we highlight the limitations of the analysis of asynchronous hyperproperties by showing that major problems of interest for both models (model checking, satisfiability, automata emptiness) are not even arithmetical. Thus, these problems are not only undecidable, but also exhaustive approximation analyses are impossible as they require recursive enumerability. Finally, we consider natural semantic restrictions – k -synchronicity and k -context-boundedness – of both models that give rise to families of increasingly precise over- and underapproximate analyses. Also, we identify settings where these analyses yield precise results. We provide precise completeness results for all but one of the corresponding decision problems. Our complexity results for restricted classes of AAPA also shed new light on the classical theory of multitape automata over finite words as both the restrictions and the proofs can be directly transferred.

The rest of the paper is structured as follows: in Section 2, we provide some basic notation and recall the definitions of Alternating Parity Automata and Regular Transducers. Then, in Section 3, we introduce Alternating Asynchronous Parity Automata (AAPA) as a model for the asynchronous analysis of multiple input words and study their closure and decidability properties. As the emptiness problem is undecidable, we discuss approximate analyses which lead to decidability for corresponding fragments. We introduce H_μ as a novel fixpoint logic for hyperproperties in Section 4. Section 5 establishes the connection between AAPA and H_μ . In Section 6 and Section 7, this connection is used to transfer the approximate analyses of AAPA to H_μ and to obtain tight complexity bounds for corresponding decision problems. We summarise the paper in Section 8. Due to lack of space, we have transferred some proofs to an extended version [Gutsfeld et al. 2020b].

Related work. Hyperproperties were systematically considered in [Clarkson and Schneider 2010]. The temporal logics HyperLTL and HyperCTL* were introduced in [Clarkson et al. 2014] and efficient algorithms for them were developed in [Finkbeiner et al. 2015]. The polyadic μ -calculus [Andersen 1994] is directly related to hyperproperties. It extends the modal μ -calculus by branching over tuples of states instead of single states and has recently been considered in the context of so-called *incremental hyperproperties* [Milushev and Clarke 2013]. This logic can express properties that are not expressible in HyperLTL and vice versa [Rabe 2016]. The same relation holds between the polyadic μ -calculus and H_μ : On the one hand, HyperLTL can be embedded into H_μ trivially, and on the other hand, H_μ is a linear time logic, while the polyadic μ -calculus is a branching time logic, implying the logics are expressively incomparable. The polyadic μ -calculus was later reinvented [Lange 2015] under the name *higher-dimensional μ -calculus* [Otto 1999] and it was shown that

every bisimulation-invariant property of finite graphs that can be decided in polynomial time can be expressed in it.

A different class of logics with the ability to express hyperproperties are the first- and second-order logics with equal-level predicate $\text{MPL}[E]$, $\text{MSO}[E]$, $\text{FOL}[\prec, E]$ and $\text{S1S}[E]$ [Coenen et al. 2019; Finkbeiner 2017; Spelten et al. 2011]. We believe that H_μ can be embedded into the most powerful of these logics, $\text{S1S}[E]$ and $\text{MSO}[E]$. Since $\text{MPL}[E]$ and $\text{MSO}[E]$ are branching time logics while H_μ is a linear time logic, just like $\text{FOL}[\prec, E]$ and $\text{S1S}[E]$, we restrict our further analysis to the relationship between H_μ and these latter two logics. We believe that the expressive power of $\text{FOL}[\prec, E]$ and H_μ is incomparable: As for HyperLTL [Bozzelli et al. 2015], the property that an atomic proposition does not occur on a certain level in the tree (of traces) – which is directly expressible in $\text{FOL}[\prec, E]$ – likely is not expressible in H_μ . On the other hand, for singleton trace sets, H_μ and $\text{FOL}[\prec, E]$ degenerate to the linear time μ -calculus and $\text{FOL}[\prec]$, respectively, and it is known that $\text{FOL}[\prec]$ – unlike the linear time μ -calculus – cannot express all ω -regular properties. Notwithstanding $\text{S1S}[E]$, we think the study of H_μ is of interest because (i) it is closer to logics traditionally used in model checking and (ii) there is no obvious way to define decidable approximate analyses for $\text{S1S}[E]$ as we do for H_μ . Indeed, all results concerning $\text{S1S}[E]$ we are aware of are undecidability results.

The logic H_μ proposed in the current paper is based on the linear-time μ -calculus [Vardi 1988] and our model checking algorithms use a construction based on alternating parity word automata with holes in the flavour of [Lange 2005] while handling quantifiers via the constructions for HyperCTL* from [Finkbeiner et al. 2015]. AAPA are asynchronous ω -automata with a parity acceptance condition. Asynchronous automata on finite words were already introduced in the seminal paper by Rabin and Scott [Rabin and Scott 1959] on finite automata and later considered in many other contexts [Furia 2014; Geidmanis 1987; Ibarra and Trn 2013]. On infinite words, Bchi automata on multiple input words were considered in the context of recursion theory and descriptive set theory [Finkel 2006, 2016]. As far as we are aware, automata of this type with a parity acceptance condition or alternation have not been studied yet and neither have algorithms for decidable restrictions and their exact complexity. A different line of research discusses variants of asynchronous automata for the analysis of concurrent programs [Muscholl 1996; Peled and Penczek 1996; Zielonka 1987]. However, unlike AAPA, these models are concerned with language recognition for trace languages in the context of concurrent systems. For AAPA (and H_μ), we use two types of restrictions: k -synchronicity and k -context-boundedness. The first restriction has been discussed in the context of multitape automata [Furia 2014; Ibarra and Trn 2013], while the second restriction is inspired by a similar condition used in the analysis of concurrent programs [Atig et al. 2009; Bansal and Demri 2013; Qadeer 2008; Qadeer and Rehof 2005]. In [Krebs et al. 2018], Krebs et. al. consider a team semantics based approach to the verification of hyperproperties using variants of LTL with synchronous and asynchronous semantics. Of course, there is a large body of work on the analysis of asynchronous systems, e.g. [Durand-Gasselin et al. 2015; Esparza et al. 2016; Ganty and Majumdar 2012; Ganty et al. 2009]. However, we are not aware of any such work concerning hyperproperties.

2 PRELIMINARIES

Let AP be a finite set of atomic propositions. A *Kripke Structure* is a tuple $\mathcal{K} := (S, s_0, \delta, L)$ where S is a finite set of states, $s_0 \in S$ is an initial state, $\delta \subseteq S \times S$ is a transition relation and $L : S \rightarrow 2^{\text{AP}}$ is a labeling function. We assume that there are no states without outgoing edges, that is for each $s \in S$, there is an $s' \in S$ with $(s, s') \in \delta$. A path in a Kripke Structure $\mathcal{K}$ is an infinite sequence $s_0 s_1 \dots \in S^\omega$ where s_0 is the initial state of $\mathcal{K}$ and $(s_i, s_{i+1}) \in \delta$ holds for all $i \geq 0$. We denote by $\text{Paths}(\mathcal{K})$ the set of paths in $\mathcal{K}$ starting in s_0 . A trace is an infinite sequence from the set $(2^{\text{AP}})^\omega$. For a path $s_0 s_1 \dots$,

the induced trace is given by $L(s_0)L(s_1)\dots$. We write $\text{Traces}(\mathcal{K})$ to denote the traces induced by paths of a Kripke Structure $\mathcal{K}$ starting in s_0 .

Let Σ be a finite input alphabet. A (nondeterministic) regular transducer over Σ is a tuple $\mathcal{T} = (Q, q_0, \gamma)$ where Q is a finite set of control locations, $q_0 \in Q$ is an initial location and $\gamma : Q \times \Sigma \rightarrow 2^{Q \times \Sigma}$ is a transition function. Given a word $w = w_1 \dots w_n \in \Sigma^*$, a run of $\mathcal{T}$ on w is an alternating sequence $q_0 v_1 q_1 v_2 \dots v_n q_n$ such that $(q_{i+1}, v_{i+1}) \in \gamma(q_i, w_{i+1})$ for all $0 \leq i < n$. We then call $\mathcal{T}(w) := v_1 \dots v_n \in \Sigma^*$ an output of $\mathcal{T}$ on w . Intuitively, a nondeterministic regular transducer can be seen as a nondeterministic finite automaton (NFA) with output.

An Alternating Parity Automaton (APA) over Σ is a tuple $\mathcal{A} = (Q, q_0, \rho, \Omega)$ such that Q is a finite, non-empty set of control locations, $q_0 \in Q$ is an initial control location, $\rho : Q \times \Sigma \rightarrow \mathcal{B}^+(Q)$ is a function that maps control locations and input symbols to positive boolean formulas over control locations and $\Omega : Q \rightarrow \{0, 1, \dots, k\}$ is a function that maps control locations to priorities. We assume that every APA has two distinct states *true* with priority 0 and *false* with priority 1 such that $\rho(\text{true}, \sigma) = \text{true}$ and $\rho(\text{false}, \sigma) = \text{false}$ for all $\sigma \in \Sigma$. If $\rho(q, \sigma)$ only consists of disjunctions for every q and σ , we call an APA a Nondeterministic Parity Automaton (NPA) and denote $\rho(q, \sigma)$ as a set of control locations. Additionally, we allow states X with $\rho(X, \sigma) = \perp$ and $\Omega(X) = \perp$, which we call holes [Lange 2005]. Intuitively, a hole is a state where the construction of an APA is not yet finished. By $\mathcal{A}[X : \mathcal{A}']$, we denote the APA $\mathcal{A}$ where the hole X is replaced by the automaton $\mathcal{A}'$.

A tree T is a subset of $\mathbb{N}^*$ such that for every node $t \in \mathbb{N}^*$ and every positive integer $n \in \mathbb{N}$: $t \cdot n \in T$ implies (i) $t \in T$ (we then call $t \cdot n$ a child of t), and (ii) for every $0 < m < n$, $t \cdot m \in T$. We assume every node has at least one child. A path in a tree T is a sequence of nodes $t_0 t_1 \dots$ such that $t_0 = \varepsilon$ and t_{i+1} is a child of t_i for all $i \in \mathbb{N}_0$. A run of an APA $\mathcal{A}$ on an infinite word $w \in \Sigma^\omega$ is defined as a Q -labeled tree (T, r) where $r : T \rightarrow Q$ is a labelling function such that $r(\varepsilon) = q_0$ and for every node $t \in T$ with children $t_1, \dots, t_k$, we have $1 \leq k \leq |Q|$ and the valuation assigning true to the control locations $r(t_1), \dots, r(t_k)$ and false to all other control locations satisfies $\rho(r(t), w(|t|))$. A run (T, r) is an accepting run iff for every path $t_1 t_2 \dots$ in T , the lowest priority occurring infinitely often is even. A word w is accepted by $\mathcal{A}$ iff there is an accepting run of $\mathcal{A}$ on w . The set of infinite words accepted by $\mathcal{A}$ is denoted by $\mathcal{L}(\mathcal{A})$. Extending the notion of holes, we write $\mathcal{A}[X : \mathcal{L}]$ for a language $\mathcal{L} \subseteq \Sigma^\omega$ to denote $\mathcal{A}[X : \mathcal{A}']$ for some automaton $\mathcal{A}'$ with $\mathcal{L}(\mathcal{A}') = \mathcal{L}$. We call an APA (resp. NPA) an Alternating Büchi Automaton (resp. Nondeterministic Büchi Automaton) iff its priorities are 0 and 1. We abbreviate these automata as ABA and NBA. In the remainder of this paper, we use known theorems about parity and Büchi automata (without holes):

PROPOSITION 2.1 ([DAX AND KLAEDTKE 2008]). *For every APA $\mathcal{A}$ with n states and k priorities, there is a nondeterministic Büchi Automaton with $2^{O(n \cdot k \cdot \log n)}$ states accepting the same language.*

PROPOSITION 2.2. *For every APA $\mathcal{A}$ with n states and k priorities, there is an APA $\overline{\mathcal{A}}$ with n states and k priorities that recognises the complement language.*

PROPOSITION 2.3. *The emptiness problem is PSPACE-complete for APA and NLOGSPACE-complete for NPA and NBA.*

Proposition 2.2 and **Proposition 2.3** can be found e.g. in [Demri et al. 2016]. On multiple occasions in this paper, we use a function for nested exponentials. Specifically, we define $g_{c,p}(0, n) := p(n)$ and $g_{c,p}(d+1, n) := c^{g_{c,p}(d, n)}$ for a constant $c > 1$ and a polynomial p . For $c = 2$ and $p = \text{id}$, i.e. the identity function, we omit the subscripts in order to improve readability. By slight abuse of notation, we say that a function f is in $O(g(d, n))$ if f is in $O(g_{c,p}(d, n))$ for some constant $c > 1$ and polynomial p . We straightforwardly extend this notion to multiple g functions, where different constants $c > 1$ and polynomials p can be used for the various g functions. Also, we write $\text{SPACE}(g(d, n))$ as an abbreviation for $\bigcup_{c,p} \text{SPACE}(g_{c,p}(d, n))$.

3 ALTERNATING ASYNCHRONOUS PARITY AUTOMATA

We introduce a new class of automata for the asynchronous traversal of multiple ω -words.

Definition 3.1 (Alternating Asynchronous Parity Automaton). Let $M = \{1, 2, \dots, n\}$ be a set of directions and Σ an input alphabet. An Alternating Asynchronous Parity Automaton (AAPA) is a tuple $\mathcal{A} = (Q, \rho_0, \rho, \Omega)$ where

- Q and Ω are the same as in an APA,
- $\rho_0 \in \mathcal{B}^+(Q)$ is a positive boolean combination of initial states, and
- $\rho : Q \times \Sigma \times M \rightarrow \mathcal{B}^+(Q)$ maps triples of control locations, input symbols and directions to positive boolean combinations of control locations.

Just as for APA, we call an AAPA where $\rho(q, \sigma, d)$ and ρ_0 only consist of disjunctions a Nondeterministic Asynchronous Parity Automaton (NAPA). Compared to an APA, where a single word over Σ is given as input, an AAPA has access to n input words over Σ and can perform steps on them individually. The M argument of the transition function indicates on which input word to progress. Note that any APA can be seen as an AAPA with $n = 1$. The definition of a run T of an AAPA is similar to the one for a run of an APA, but with the following modifications:

- the run is defined over n input words $w_1, \dots, w_n \in \Sigma^\omega$ instead of a single word w ,
- for each $t \in T$, we have n offset counters $c_1^t, \dots, c_n^t$ starting at $c_i^t = 0$ for all i and t with $|t| \leq 1$,
- we have $\{r(t) \mid t \in T, |t| = 1\} \models \rho_0$, and
- when node $t \in T \setminus \{\varepsilon\}$ has children $t_1, \dots, t_k$, then there is a $d \in M$ such that (i) $c_d^{t_i} = c_d^t + 1$ and $c_{d'}^{t_i} = c_{d'}^t$ for all i and $d' \neq d$, (ii) we have $1 \leq k \leq |Q|$ and (iii) the valuation assigning true to $r(t_1), \dots, r(t_k)$ and false to all other states satisfies $\rho(r(t), w_d(c_d^t), d)$.

These automata are particularly suitable for the analysis of our new logic H_μ , which we introduce in the next section. Indeed, AAPA and H_μ are able to express the same asynchronous restrictions on multiple ω -words, as shown in [Section 5](#). In order to compare AAPA to different automata models over a single input word, we interpret the n input words $w_1, \dots, w_n$ with $w_i = w_i^0 w_i^1 \dots \in \Sigma^\omega$ as a single word $w = (w_1^0, \dots, w_n^0)(w_1^1, \dots, w_n^1) \dots \in (\Sigma^n)^\omega$. We introduce some notation for the asynchronous manipulation of such words. For this purpose, let $v = (v_1, \dots, v_n) \in \mathbb{N}_0^n$ be a vector. Then, we use $w[v] = (w_1^{v_1}, \dots, w_n^{v_n})(w_1^{v_1+1}, \dots, w_n^{v_n+1}) \dots$ to denote w shifted left according to entries in v . We write $v + v'$ and $v + e_i$ for standard vector addition with an arbitrary vector v' and a unit vector $e_i = (0, \dots, 0, 1, 0, \dots, 0)$, respectively.

We establish some properties of AAPA and their nondeterministic counterparts. First, a standard argument using alternation and priority shifts gives us the following result about AAPA:

THEOREM 3.2. *AAPA are closed under union, intersection and complement. The constructions are linear in the size of the input automata.*

However, for their nondeterministic counterpart, the above result does not hold:

THEOREM 3.3. *NAPA are not closed under intersection. In particular, it is undecidable whether there is a NAPA recognizing the intersection $\mathcal{L}(\mathcal{A}_1) \cap \mathcal{L}(\mathcal{A}_2)$ for two NAPA $\mathcal{A}_1, \mathcal{A}_2$.*

PROOF. Let finite state Asynchronous Automata (AA) be NAPA on finite words which accept input words by reaching accepting states. Fix two languages $\mathcal{L}_1, \mathcal{L}_2$ recognizable by AA over an alphabet Σ . Let $\perp$ be a fresh symbol. Then $\hat{\mathcal{L}}_1 = \{w\perp^\omega \mid w \in \mathcal{L}_1\}$ and $\hat{\mathcal{L}}_2 = \{w\perp^\omega \mid w \in \mathcal{L}_2\}$ are recognizable by two NAPA. By a standard argument, from a NAPA recognizing $\hat{\mathcal{L}}_1 \cap \hat{\mathcal{L}}_2$, we obtain an AA recognizing $\mathcal{L}_1 \cap \mathcal{L}_2$, but AA are not closed under intersection and it is undecidable whether there is an AA recognizing the intersection [[Furia 2014](#)]. $\square$

Due to this difference in closure properties, we obtain a gap in expressivity:

COROLLARY 3.4. *There is an AAPA such that no NAPA recognises the same language. It is undecidable whether an AAPA can be translated to a NAPA that recognises the same language.*

A related result regarding the intersection of NAPA is:

THEOREM 3.5. *The emptiness problem for the intersection of NAPA is undecidable.*

PROOF. The proof is by reduction from the Post Correspondence Problem (PCP). Let $I = (w_1, u_1) \dots (w_n, u_n)$ be a PCP instance. We again choose $\perp$ as a fresh symbol. Let $\mathcal{L}_1$ be the language consisting of all possible concatenations of (w_i, u_i) followed by $(\perp, \perp)^\omega$ and $\mathcal{L}_2$ be the language $\{(a_1, a_1)(a_2, a_2) \dots (a_l, a_l)(\perp, \perp)^\omega \mid a_i \in \Sigma, l \in \mathbb{N}_0\}$. $\mathcal{L}_1$ can be recognised by a NAPA with loops on an initial state that first read w_i in direction 1 and then read u_i in direction 2. $\mathcal{L}_2$ can be recognised by inspecting a single symbol of both directions in turns and making sure that the same symbol is being read. Then the PCP instance has a solution iff $\mathcal{L}_1 \cap \mathcal{L}_2$ is non-empty. $\square$

Since the intersection of two NAPA can be recognised by an AAPA, we immediately obtain:

COROLLARY 3.6. *The emptiness problem for AAPA is undecidable.*

Furthermore, since AAPA can be used to decide inclusion problems for two-tape Büchi automata [Finkel and Lecomte 2009], this result can be strengthened drastically. We refer the interested reader to the extended version for a detailed proof of the strengthened claim.

THEOREM 3.7. *The emptiness problem for AAPA is hard for the level Σ_2^1 of the analytical hierarchy.*

When comparing NAPA to synchronous automata, it is not hard to see that the language $\mathcal{L} = \{(a, a)^n(b, a)^n(b, b, a)^n(b, b, b)^\omega \mid n \in \mathbb{N}_0\}$ can be recognised by a NAPA, while there is no synchronous parity automaton recognising $\mathcal{L}$ since it is not an ω -regular language. Despite this increase in expressive power, the emptiness problem for NAPA can still be reduced to an emptiness test on their synchronous counterparts.

THEOREM 3.8. *The emptiness problem for NAPA is PSPACE-complete.*

PROOF. Given a NAPA $\mathcal{A}$ with m states and n input words over the alphabet Σ , we construct a NAPA $\mathcal{A}'$ with $m \cdot |\Sigma|^n$ states over Σ that stores the currently accessible input vector in its state and guesses the next input symbol for each direction. For hardness, we can encode the configurations of a Turing machine whose space is bounded by a polynomial $p(x)$ in a NAPA with $n = p(x)$ directions. Consistency of successive configurations can be checked locally such that there is no need to represent full configurations in the input alphabet or the state space. Therefore, the size of the NAPA stays polynomial despite of the fact that the Turing machine has exponentially many configurations. Thus, we can check for the existence of an accepting run of the Turing machine. $\square$

Since the translation used in the proof is only exponential in n and NAPA subsume synchronous Büchi automata, we obtain the following corollary:

COROLLARY 3.9. *For fixed n , the emptiness problem for NAPA is NLOGSPACE-complete.*

As emptiness of AAPA cannot be decided and as alternation elimination is not possible, we study analyses that consider well-specified subclasses of runs and identify fragments of AAPA for which these analyses are precise.

3.1 k -synchronous Analysis of AAPA

Definition 3.10. We call a run T of an AAPA with n input words k -synchronous for a $k \in \mathbb{N} \cup \{\infty\}$, if in every node t in T , the offset counters $c_1^t, \dots, c_n^t$ satisfy $|c_i^t - c_j^t| \leq k$ for all i and j .

Intuitively, a k -synchronous run has the property that the AAPA can never be *ahead* more than k steps in one direction than in any other. This gives rise to an approximate analysis where only the k -synchronous runs of an AAPA are considered. Since a k -synchronous run is particularly a k' -synchronous run for all $k' \geq k$, the approximation improves with increasing k , capturing the whole semantics at $k = \infty$ for all AAPA. Since an analysis with $k = \infty$ is impossible, we assume that $k < \infty$ in the remainder of this section. We show that k -synchronous runs of an AAPA can be analysed via a reduction to APA:

THEOREM 3.11. *For every AAPA $\mathcal{A}$ over Σ with l priorities, n input words and m states, there is an APA over Σ^n with $O(l \cdot m \cdot |\Sigma|^{k \cdot n})$ states recognizing all words accepted by a k -synchronous run of $\mathcal{A}$.*

PROOF. We read input vectors synchronously and maintain a $k \cdot n$ window of the input words in the state space. Since in a k -synchronous run of an AAPA each direction can be ahead each other direction at most k steps, no direction can leave this window without all rearmost directions performing steps. The content of the window can therefore be stored in an APA's state space.

We simulate steps of the AAPA by moving markers forward in each row of the window. When the last marker leaves the rearmost column of the window, a new input vector is read and added to the front. A step that would leave the window must not be simulated since that would move that direction $k + 1$ steps ahead of the rearmost direction. Therefore, we move to *false* instead. This way, we ensure that each run of the APA simulates a k -synchronous run of the AAPA.

Since an input vector is only read when all directions have left the last column, one step of the APA has to simulate multiple successive steps of the AAPA. This is done by a nondeterministic choice over all sequences of steps resulting in the rearmost column being erased. In order to correctly mirror the priorities, for each simulated sequence of steps, we move to a copy of the reached state annotated with the lowest priority encountered in this sequence.

In order to fill the window, k input vectors are read in an *initialisation* phase of the APA. Overall, the size of the state spaces increases by a factor of $O(l \cdot |\Sigma|^{k \cdot n})$. $\square$

Note that [Theorem 3.11](#) yields an underapproximation of an AAPA's behaviour. By transitioning to *true* instead of *false* at a violation of k -synchronicity, we can instead perform an overapproximation that ignores non k -synchronous branches when determining whether a run is accepting. Obviously, these approximations yield exact results for AAPA where all runs are k -synchronous. We call such AAPA k -synchronous. Synchronicity can be enforced by a syntactic restriction, namely that in all cycles in the transition graph, every direction occurs the same number of times. The parameter k is then induced by the largest difference of the number of occurrences of two directions on any path in the transition graph. We establish a tight complexity bound for this analysis:

THEOREM 3.12. *The problem to decide whether there is a k -synchronous accepting run of an AAPA and thus the emptiness problem for k -synchronous AAPA is EXPSPACE-complete.*

PROOF. Using [Theorem 3.11](#), we can construct an APA whose size is exponential in the size of the input. Since APA can be tested for emptiness in PSPACE, this establishes membership in EXPSPACE. For hardness, we reduce from the acceptance problem for deterministic exponentially space bounded Turing machines (DTMs). Let $\mathcal{M}$ be a DTM with control locations Q and space complexity $f(x) = 2^{p(x)}$ for some polynomial p . For a given input w of $\mathcal{M}$, we construct a k -synchronous AAPA $\mathcal{A}$ with $n := p(|w|) + 1$ and $k := 2$ as follows:

DTM	1	#	0	0	1	(1,q)	0	1	1	...
Bit 1	0	1	0	1	0	1	0	1	0	...
Bit 2	0	1	0	1	0	1	0	1	0	...
Bit 3	0	1	0	1	0	1	0	1	0	...

Fig. 1. A 2-synchronous virtual counter gadget with 3 bits and current value 101. The Turing machine is in state q and its head is on the fourth bit of the current configuration.

In the first direction, we successively read configurations of $\mathcal{M}$ separated by a marker. We encode such configurations by words over $\{0, 1\} \cup \{0, 1\} \times Q$ where the occurrence of a tuple (b, q) indicates that the head points to this position, that this position of the tape contains the bit b , and that the current state is q . Presence of initial and final configuration can easily be checked via alternation.

Additionally, we need to check that successive configurations are constructed in accordance with the transition function of $\mathcal{M}$. For this purpose, we have to count to exponentially large values in order to compare positions with the same (or neighbouring) index in successive configurations with each other. Since we cannot store exponentially large counter values in the state, we instead construct a *virtual counter gadget*: for every additional direction, we can enforce that it consist of the word $(01)^\omega$. As we illustrate in Figure 1, the concatenation of the current values of these directions is interpreted as a counter. Then, we conjunctively move to each position in a synchronous manner, save the value of the first direction in a state of the AAPA and initialise the counter: if the directions have value 1, we advance them by one symbol and otherwise, we keep the position. We then advance the first direction and increase our virtual counter by changing the appropriate bits via single advancements and maintaining the other bits by nondeterministically either maintaining the current symbol or advancing the direction by two positions. Then, one of the nondeterministic choices preserves 2-synchronicity. When the counter has reached $2^{P(|w|)}$, we have found the matching tape cell in the next configuration and can check whether it is admissible by comparing its value with the value saved in the control state of the AAPA.

It is easy to see that $\mathcal{A}$ has a 2-synchronous accepting run iff $\mathcal{M}$ accepts w . $\square$

Note that in the proof of the lower bound detailed above, k can be chosen as a fixed value. This raises the question whether there is a construction for the emptiness test that is exponential only in n , but not in k . Indeed, we can construct an APA that has a non-empty language if and only if the given AAPA has an accepting k -synchronous run. However, unlike the APA in the proof of Theorem 3.11, the constructed APA only accepts a certain encoding of the word accepted by the k -synchronous run instead of the word itself. More specifically, it expects the input word to consist of concatenations of input windows (with size $k \cdot n$) of the original input words, one window for each simulated step of the AAPA. For each direction it maintains a counter indicating its current position in the input window. Using alternation and additional counters, the APA can ensure that the succession of input windows is consistent. A single step of the AAPA can then be simulated by reading the current copy of the input window in order to check that the direction inducing the step has the correct symbol. An upper bound on the number of states of this APA is dominated by the n counters up to k indicating on what position in the window each direction is. This results in a factor of k^n on the number of states but stays polynomial for fixed n . As single steps of the AAPA are simulated separately, this construction also avoids the additional factor l .

Together with the fact that already the emptiness problem for APA is PSPACE-hard, these considerations lead to the following corollary.

COROLLARY 3.13. *For fixed n , the emptiness problem for k -synchronous AAPA is PSPACE-complete.*

Fig. 2. A 3-context-bounded run of an AAPA where green nodes (circles) move in direction 1 and red nodes (squares) move in direction 2. The corresponding guess is $(q_1, \{(q_2, \{q_1, q_6\}), (q_3, \{q_4, q_6\})\})$.

3.2 k -context-bounded Analysis of AAPA

Since different words can only diverge from each other up to k steps, the ability of AAPA to asynchronously traverse words is severely restricted in k -synchronous runs. We thus consider a further class of runs where the positions on different words can diverge unboundedly. For this restriction, we introduce the notion of a context:

Definition 3.14. A context is a (possibly infinite) subpath $p = t_1 t_2 \dots$ in a run of an AAPA over $w_1, \dots, w_n$ such that transitions between successive states all use the same direction, that is there is a $d \in M$ such that for all $i \in \{1, \dots, |p| - 1\}$ we have $c_d^{t_{i+1}} = c_d^{t_i} + 1$. We call a run T of an AAPA k -context-bounded if every path in T consists of at most k contexts.

We propose an approximate analysis which checks only for the existence of a k -context-bounded accepting run of an AAPA. The restriction that AAPA can only consider a single direction during each context is well-chosen in the sense that additionally allowing contexts in which a selection of directions is traversed synchronously leads to undecidability.

Our analysis is based on the observation that a run of an AAPA can be divided into maximal sections only reading a single direction and that in the case of a k -context-bounded run, the interaction between neighbouring sections can be described in a finite way. Given such a description, each direction can be analysed independently from the others. We illustrate this with an example with just two directions and two context switches. Consider the run in Figure 2 where green states move in direction 1 and red states move in direction 2 and inspect the left part of the tree. Instead of moving to q_2 , the automaton responsible for direction 1 can directly move to q_1 and q_6 from q_1 if the automaton responsible for direction 2 ensures existence of an accepting run from q_2 that reaches only these two direction 1 states. Consequently, the automaton responsible for direction 2 can safely cut off the branches between states q_2 and q_1 as well as between q_5 and q_6 since the automaton responsible for direction 1 ensures an accepting continuation from q_1 and q_6 . We furnish states with descriptions of this interplay between assumptions and guarantees and call such descriptions *guesses*. These guesses allow us to split the run from Figure 2 into two runs on single directions as shown in Figure 3. Guesses are nested and the nesting depth corresponds to the number of context switches that can still be performed in a k -context-bounded run. They are constructed inductively: In the most simple case, where no context switch can be made, the guess is empty. If context switches are still possible, the guess is a set of states enriched with guesses with one context switch less.

Formally, when analysing an AAPA $\mathcal{A} = (Q, \rho_0, \rho, \Omega)$ over n words from Σ^ω , we inductively define the set of guesses G_i for i context switches as follows: $G_0 := \{\perp\}$, $G_{i+1} := 2^{Q \times G_i}$ for $i \geq 0$, and $G := \bigcup_{i=0}^{k-1} G_i$. For simplicity, we assume that each state $q \in Q$ moves in a unique direction d (i.e. the transition function yields *false* for other directions); the set of states that move in direction d is

Fig. 3. Translation of the run into multiple synchronous runs with $G_1 = \{(q_2, G_2), (q_3, G_3)\}$, $G_2 = \{q_1, q_6\}$ and $G_3 = \{q_4, q_6\}$. Note that it suffices to include one subtree for each of q_6 and q_4 on the third level of the green tree since all subtrees reached after returning to direction 0 have the same offset for this direction.

called Q_d . This increases the size of the state space by a factor of at most n . As described above, we skip sections belonging to direction $d' \neq d$ when considering a direction d . In order to extract the potential reentry points for directions d from a guess, we use a frontier function $F_d : G \rightarrow G$ defined as follows:

$$F_d(g) = \{(q, g') \in g \mid q \in Q_d\} \cup \bigcup \{F_d(g') \mid (q, g') \in g, q \notin Q_d\}.$$

For the analysis, we define from $\mathcal{A}$ an APA-like structure $\mathcal{S} = (S, \rho_S, \Omega_S)$ over the alphabet Σ with

$$S = Q \times G,$$

$$\begin{aligned} \rho_S((q, g), \sigma) &= \bigvee \left\{ \bigwedge Q'_d \times \{g\} \wedge \bigwedge \{F_d(g'') \mid (q, g'') \in g'\} \mid \right. \\ &\quad \left. Q'_d \subseteq Q_d, g' \subseteq g, Q'_d \cup \{q \mid \exists g'' : (q, g'') \in g'\} \models \rho(q, \sigma, d) \right\} \text{ for } q \in Q_d, \text{ and} \\ \Omega_S((q, g)) &= \Omega(q). \end{aligned}$$

Thus, the successors of a state (q, g) in $\mathcal{S}$ are composed of two sets: a set Q'_d of states that belong to the same section and the set of reentry points extracted from the second component of the guesses in g' annotating states for other directions in g . These sets are chosen such that Q'_d and the states in the first component of the guesses in g' satisfy the corresponding transition function of $\mathcal{A}$.

For this structure, we have:

LEMMA 3.15. *There is a k -context-bounded accepting run in $\mathcal{A}$ iff there is a $g \in G$ such that $\mathcal{S}$ accepts from $\bigwedge \{g' \in F_d(\{(q_0, g)\}) \mid q_0 \in Q_0\}$ for all $d \in \{1, \dots, n\}$ and some Q_0 with $Q_0 \models \rho_0$.*

With a little effort, this structure can be translated to an NPA over Σ^n recognising the words accepted by k -context-bounded runs of the original AAPA $\mathcal{A}$. For this, transitions in each state $q \in Q_d$ of $\mathcal{S}$ are adjusted to the input alphabet Σ^n by allowing arbitrary symbols in other directions than d . The key insight is that the choice of the guess g from Lemma 3.15 can be integrated into an NPA after an alternation removal as a nondeterministic choice. We parameterise the structure $\mathcal{S}$ in the initial guess g (for which we assume w.l.o.g. $g \in G_{k-1}$) and obtain a structure $\mathcal{S}(g)$ of size $O(g(k-2, |\mathcal{A}|))$, where the fixed $g \in G$ indicates the initial guess. The initial state of $\mathcal{S}(g)$ can nondeterministically guess a state set Q_0 with $Q_0 \models \rho_0$ and then move to $\bigwedge_{d \in \{1, \dots, n\}} \bigwedge_{q_0 \in Q_0} \bigwedge F_d(\{(q_0, g)\})$ to obtain the desired test. We eliminate alternation from $\mathcal{S}(g)$ to obtain an NPA $\mathcal{S}'(g)$ of size $O(g(k-1, |\mathcal{A}|))$ for each guess $g \in G_k$. Thus, an NPA which nondeterministically guesses a $g \in G$ and then moves to $\mathcal{S}'(g)$ performs the desired test from Lemma 3.15. Since there are $|G| = O(g(k-1, |\mathcal{A}|))$ possible guesses, this NPA's size asymptotically is $O(g(k-1, |\mathcal{A}|))$ as well. We obtain:

COROLLARY 3.16. *For every AAPA $\mathcal{A}$, there is an NPA with $O(g(k-1, |\mathcal{A}|))$ states recognizing all words accepted by a k -context-bounded run of $\mathcal{A}$.*

Note that similar to the construction for [Theorem 3.11](#), a violation of k -context-boundedness leads to a transition to *false* in our construction (in this case by an empty disjunction in the transition function since there is no $g' \subseteq g$ fulfilling the conditions for $g = \perp$). This again yields an underapproximation of the AAPA's behaviour and can be transformed into an overapproximation by instead transitioning to *true*.

Again, the developed analysis is precise for AAPA having only k -context-bounded runs, which we call k -context-bounded AAPA. A syntactic restriction ensuring that an AAPA is context-bounded is that each strongly connected component (SCC) in its transition graph uses steps of a unique direction only. The parameter k is then induced by the maximum number of switches between directions in the DAG of SCCs.

We show that the size of this construction cannot asymptotically be improved.

THEOREM 3.17. *The problem to decide if there is a k -context-bounded accepting run of an AAPA $\mathcal{A}$ and thus the emptiness problem for k -context-bounded AAPA $\mathcal{A}$ is complete for $(k - 2)\text{EXPSPACE}$.*

For the proof, we establish two helpful lemmas about AAPA with context bounds. We make use of Stockmeyer's nested index encoding [[Stockmeyer 1974](#)]. For a (finite) word $w = w_0 \dots w_n$, the encoding is inductively defined by:

$$\begin{aligned} \text{stock}_0(w) &:= [{}_0 w_0 \dots w_n]_0 \\ \text{stock}_{k+1}(w) &:= [{}_{k+1} \text{stock}_k(\text{bin}(0)) w_0 \dots \text{stock}_k(\text{bin}(n)) w_n]_{k+1} \end{aligned}$$

Here, $\text{bin}(i)$ denotes a binary encoding of $i \in \mathbb{N}_0$. For our proofs, it will be useful to use a least-significant-bit-first encoding for $\text{bin}(i)$. Brackets $[_k$ and $]_k$ allow us to identify the beginning and end of a level k encoding without necessarily knowing the exact length of the encoding. For level $k \geq 1$, every letter of the word is preceded by its index in the word. These indices are written in binary and thus form words of their own, which can be encoded as well. The number of times the encoding of indices is nested is determined by the level of the encoding. Since indices for a word of length n can be encoded in words of length $\log n$, every level of the encoding allows to encode exponentially larger words when starting from a fixed size. Thus, with level 0 encodings of length m , one can encode words of length $g(k, m)$ on level k .

The first lemma shows that by using this encoding, we can perform certain operations on sequences of words of length $g(k, m)$. It is formulated in a generic way using regular transducers and regular languages in order to be applicable to nested index encodings as well as successions of Turing machine configurations. The first is needed for its inductive proof; the second allows us to apply it for the proof of [Theorem 3.17](#).

The operations we need for our result are the following:

- Checking whether the first and the last word of the sequence are contained in two given regular languages and
- checking whether each word (apart from the first one) is obtained by applying a given regular transducer to the previous word.

More concretely, for sequences of indices, we want to check:

- Is the first index $0 \dots 0$?
- Is the last index $1 \dots 1$?
- Is each successor obtained by adding 1 to the previous index?

For Turing machine configurations, the checks are:

- Is the first configuration the starting configuration?
- Is the last configuration an accepting configuration?

- Is each configuration obtained by applying the Turing machine's transition function to its predecessor configuration?

Note that for our reduction to be polynomial, each of the regular language acceptors and regular transducers have to be of polynomial size in the input size, which they indeed are.

With these considerations in mind, it is straightforward to see that an application of the next Lemma can be used to encode the acceptance problem of $g(k - 2, p(n))$ space bounded Turing machines in k -context-bounded AAPA. It is also easy to see that an application to sequences of indices can be used to check the validity of nested index encodings. We now formulate the lemma:

LEMMA 3.18. *Given a size bound m , an AAPA of size polynomial in m with access to directions i and j , as well as $k \geq 2$ context switches starting in a context for direction i can:*

- (1) *enforce that the word in direction i contains a level $k - 2$ nested index encoding of all numbers between 0 and $g(k - 1, m) - 1$ and*
- (2) *enforce that the word in direction j contains a sequence of level $k - 1$ encoded words $w_1 w_2 \dots w_l$, each of length $g(k - 1, m)$, separated by markers, such that*
 - *w_1 is contained in a given regular language with an acceptor of size $O(m)$,*
 - *w_l is contained in another given regular language with an acceptor of size $O(m)$, and*
 - *for $h < l$, w_{h+1} is obtained from w_h by applying a regular transducer of size $O(m)$.*

For the proof, we need the following lemma about AAPA:

LEMMA 3.19. *Given a size bound m , an AAPA of size polynomial in m with $k \geq 1$ context switches can check whether the level $k - 1$ encoded words of length $g(k - 1, m)$ on directions i and j differ.*

PROOF OF LEMMA 3.19. The lemma can be shown by an induction on k .

In the **base case** $k = 1$, the words are level 0 encoded and have length m . Since the number of context switches is restricted, it is not possible to check the two words for differing bits directly by moving the two directions forward one step at a time. Instead, we use alternation to perform this test. For every number of steps s up to m , we disjunctively move s steps forward in direction i , perform a context switch while memorising the number of steps and the last read symbol and then move the same number of steps in direction j and compare the symbols. If the symbols differ, the words cannot be the same and we accept, otherwise we reject. If all disjunctive tests reject, there are no differing bits in the words and thus no accepting runs of the AAPA.

In the **inductive step**, we have $k + 1$ context switches available to compare level $k \geq 1$ encoded words of length $g(k, m)$. Here, we cannot count the index of a symbol in the word in the state space since that would increase the AAPA's size beyond polynomial. We can, however, make use of the fact that the words are encoded such that each symbol in the word is preceded by its index in a level $k - 1$ encoding. Thus, we start on direction i and disjunctively move to the start of some index where we expect the two words to differ. We then perform a context switch and conjunctively test for each position whether (i) the indices differ or (ii) the symbols differ. Test (i) can be done using the induction hypothesis with the k remaining context switches. Test (ii) can be done using a single additional context switch by moving to the symbol, memorising it, then moving to the respective symbol on direction i and comparing the two. $\square$

PROOF OF LEMMA 3.18. The lemma can be shown by an induction on k .

The **base case** is $k = 2$. Item (1) can be enforced by checking for the sequence of words bracketed by $[_0$ and $]_0$ in direction i , that (i) these words have length m , that (ii) the first one is $0\dots 0$, that (iii) the last one is $1\dots 1$ and that (iv) each successor is obtained from the previous word by adding 1. These four checks can be done via alternation and a counter to obtain (1).

For (2), we first need to check that direction j contains level 1 nested index encodings of words of length 2^m . This can be done in the same way as the check in item (1) with the difference that bits belonging to the encoded words themselves have to be skipped. These can, however, easily be found since the indices to be checked are bracketed by $[_0$ and $]_0$. w_1 and w_l belonging to regular languages can be checked by their respective language acceptors, where bits not belonging to the words themselves can be skipped. The main difference between the two tests is that l is not determined and thus the point where l is reached has to be guessed nondeterministically.

For the third item of (2), the challenge is that corresponding positions in w_h and w_{h+1} are $O(2^m)$ steps apart from each other due to the length of the words w_h and thus cannot be matched by counting steps in the state of the AAPA as that would violate its size restriction. However, we can use the two context switches and the fact that we have already checked item (1) to ensure this item. A test for this starts in direction i and conjunctively switches to direction j at the start of each number from item (1). The copy that performs the context switch before number n then has to ensure the correct transduction of the n th bit from each word w_h to w_{h+1} . For this purpose, it conjunctively moves to the start of each word w_h and performs the transduction of bit n . As the state space is not large enough to store the value n , this is done in the following way: for each position in w_h it is checked whether (i) the correct transduction is being performed or (ii) the index does not match n . The latter is done by using [Lemma 3.19](#) with the remaining context switch. Note that in order to ensure that the former check can be performed, the control location of the transducer has to be tracked. This can be done by enriching the input word in direction j with states of the transducer in each bit and checking their correct succession during the transduction tests. Since the transducer's control location is available for each position, the copy can apply the transducer's transition function to the current bit and nondeterministically choose one of the possible tuples of new control location and output bit. It then checks for the new control location at the next position in w_h . It also checks for the output bit in w_{h+1} by again disjunctively testing for each position whether the bit is present or the index is different from n .

In the **inductive step**, we show the claim for $k + 1$ under the induction hypothesis (*IH*) for k . For the proof of (1) we swap the roles of direction i and j and use the induction hypothesis: as argued before, the first index found being $0\dots 0$, the last index being $1\dots 1$ and each index being obtained from the last by adding 1 can be described by regular languages (resp. transducers of constant size). Hence, claim (1) for $k + 1$ follows from claim (2) for k . It remains to argue that claim (2) for $k + 1$, which we will show next, and claim (1) for k , which we have used here, are consistent with each other since they formulate different requirements for the same direction. This is the case since a higher level nested index encoding contains all level $k - 1$ encoded indices from $0\dots 0$ to $1\dots 1$ in ascending order. They can easily be identified by the level $k - 1$ brackets $[_{k-1}$ and $]_{k-1}$.

For claim (2), we first have to show that direction j contains a sequence of words correctly level k encoded. Like in the base case, this is done in a similar way as claim (1) is enforced, namely by checking presence of level $k - 1$ encodings of $0\dots 0$ to $1\dots 1$ in ascending order while skipping bits belonging to the encoded word itself. Note that (*IH*) applies here since we can use the level $k - 1$ indices of the level k encodings that are present in direction i as we have shown already. The main difference to the argument for claim (1) is that we have to start on direction j instead of i , conjunctively move to the start of each word, and then perform a context switch to be able to use (*IH*). This results in $k + 1$ context switches for this case and thus does not violate the context restriction. Also (just like in the base case), item one and two of claim (2) can easily be achieved by applying the language acceptors to the bits on the highest level of the encodings. The third item of (2) can also be shown as in the base case since after moving from direction i to j , there are k context switches left which allows us to apply [Lemma 3.19](#). $\square$

Now, we are finally able to show the desired result.

PROOF OF THEOREM 3.17. Inclusion follows from the construction underlying Corollary 3.16. For hardness, we reduce from the acceptance problem for $g(k-2, p(n))$ space bounded Turing machines using Lemma 3.18 as explained above. $\square$

As mentioned in the introduction, our constructions and completeness proofs can also be applied to alternating asynchronous automata (multitape automata) on finite words: the definitions of k -synchronicity and k -context-boundedness carry over in a direct manner and our algorithms for the emptiness tests can be applied. Furthermore, our hardness proofs can be transferred because our reductions only require reachability of control states and are therefore not dependent on a parity condition. Thus, we believe our results also shed new light on the theory of multitape automata.

4 A μ -CALCULUS FOR HYPERPROPERTIES

We now define our new logic, H_μ . In order to capture asynchronous hyperproperties, we combine ideas from HyperLTL [Finkbeiner et al. 2015], the polyadic μ -calculus [Andersen 1994], and the linear time μ -calculus [Barringer et al. 1986; Vardi 1988] in a novel fashion. From HyperLTL we take the ideas to quantify over path variables and to relativise atomic propositions to path variables. Inspired by the indexed next-operator of the polyadic μ -calculus, we relativise the next-operator to progress only on a single path identified by a path variable. Finally, complex hyperproperties can be specified by fixpoints. In this way, we extend the means provided by the linear time μ -calculus for specifying properties to hyperproperties. Note that none of the logics that inspired the design of H_μ is able to capture asynchronous hyperproperties. We use the following syntax:

Definition 4.1 (Syntax). Let AP be a set of atomic propositions, $N = \{\pi_1, \dots, \pi_n\}$ be a set of path variables and $\chi = \{X_1, \dots, X_m\}$ be a set of predicates. We define H_μ formulas over AP, N and χ by the following grammar:

$$\begin{aligned}\varphi &:= \exists \pi. \varphi | \forall \pi. \varphi | \psi \\ \psi &:= a_\pi | X | \bigcirc_\pi \psi | \psi \vee \psi | \neg \psi | \mu X. \psi\end{aligned}$$

where $a \in \text{AP}$, $\pi \in N$ and $X \in \chi$. We call expressions generated by the nonterminal φ quantified formulas and expressions generated by the nonterminal ψ quantifier-free formulas.

In this paper, we will use two representations of a H_μ formula φ . The first and more common one is its syntax tree which we denote by $\text{rep}_t(\varphi)$. For the second one, we compress the syntax tree into a syntax directed acyclic graph (DAG) where syntactically equivalent subformulas share the same nodes and write $\text{rep}_d(\varphi)$ for this representation. This offers an exponentially more succinct representation for some families of formulas while not increasing the complexity of the algorithms we consider. Then, we use $|\varphi|_t := |\text{rep}_t(\varphi)|$ and $|\varphi|_d := |\text{rep}_d(\varphi)|$ as two measures for the size of a formula. Note that since the DAG is obtained from the syntax tree by compressing it, $|\varphi|_d \leq |\varphi|_t$ holds for all formulas φ . Therefore, all complexity upper bounds in which the size of a formula is measured by $|\cdot|_d$ trivially transfer to complexity upper bounds for the other size measure.

We add common connectives as syntactic sugar: $\text{true} \equiv a_\pi \vee \neg a_\pi$, $\text{false} \equiv \neg \text{true}$, $\psi \wedge \psi' \equiv \neg(\neg\psi \vee \neg\psi')$, $\psi \rightarrow \psi' \equiv \neg\psi \vee \psi'$, $\psi \leftrightarrow \psi' \equiv \psi \rightarrow \psi' \wedge \psi' \rightarrow \psi$ and $\nu X. \psi \equiv \neg \mu X. \neg\psi[\neg X/X]$.

Using these additional connectives, we impose some additional constraints on the syntax of H_μ formulas. First, we assume that in a quantified formula φ all predicates are bound by a fixpoint operator. We also assume a *strictly guarded* form, where predicates X are only allowed in scope of an even number of negations inside $\mu X. \varphi$ and are directly preceded by $\bigcirc_\pi$ for some π . The latter part of this can indeed be required without loss of generality: if one constructs a formula where there is no progress through $\bigcirc_\pi$ between μX and X , the fixpoint can equally be eliminated;

if there is progress through $\bigcirc_\pi$ for some π , the $\bigcirc_\pi$ operator can be moved inwards such that it directly occurs in front of X . Second, a formula is in *positive normal form* when $\neg$ only occurs in front of atomic propositions and all bound predicates and path variables are distinct. Finally, we say that a formula is in *closed form* when all path variables and predicates are bound.

Since our logic extends the linear time μ -calculus with path quantification from HyperLTL, the formula constructs behave in similar ways as they do in those two logics. The main difference to the linear time μ -calculus is that a formula reasons about a set of paths or traces instead of over a single path or trace. Thus, constructs a_π and $\bigcirc_\pi\psi$ are indexed to express that a holds on path π or that ψ holds when path π moves one step forward. Indexing the $\bigcirc$ -operator allows us to express asynchronous behaviour. Path quantification $\exists\pi.\varphi$ or $\forall\pi.\varphi$ allows to require that for one or for all paths π , the set of paths obtained by adding π to the previously considered set fulfills φ . Boolean connectives are interpreted in the standard way. Finally, the constructs X and $\mu X.\psi$ allow us to formulate iterative properties by least fixpoint constructions.

With this logic, it becomes possible to specify asynchronous hyperproperties, i.e. properties that do not rely on traversing different paths lockstepwise. We now sketch a few scenarios in which this is useful. One potential application of hyperlogics is the analysis of multithreaded software. In this scenario, different path variables used in a formula refer to different threads of the system and both the interaction between threads as well as the specification are captured by the formula. For example, let π_0 and π_1 refer to executions of two different threads of a system that synchronise via locking. Then the formula $\mu X.(\psi_{error} \vee (\psi_{move_0} \wedge \bigcirc_{\pi_0} X) \vee (\psi_{move_1} \wedge \bigcirc_{\pi_1} X))$ expresses that the two threads can reach an error state identified by a formula ψ_{error} through a lock-sensitive interleaving. In this example, the atomic proposition $lock_{\pi_i}$ indicates that thread i currently holds the lock and the formula $\psi_{move_i} = \neg lock_{\pi_{1-i}} \vee \neg \bigcirc_{\pi_i} lock_{\pi_i}$ expresses that thread i can perform a step. Such a property clearly requires asynchronous traversal of the different paths. More complex interaction strategies can be handled by modifying the formulas ψ_{move_i} .

Asynchronicity is also useful in applications of hyperlogics in security, e.g. when steps observed by the environment do not correspond to the same number of steps in different paths of the model. This occurs, for instance, in models of software systems that reflect internal computations. In such a situation, the formula $\varphi = \exists\pi.\forall\pi'.\nu X.\mu Y.((a_\pi \leftrightarrow a_{\pi'}) \wedge \bigcirc_\pi \bigcirc_{\pi'} X) \vee \bigcirc_{\pi'} Y$ expresses that there is a path π such that for all paths π' , π and π' repeatedly agree on the atomic proposition a . However, contrary to the HyperLTL formula $\exists\pi.\forall\pi'.\mathcal{G}(a_\pi \leftrightarrow a_{\pi'})$ a step on π can be matched by an arbitrary number of steps on π' . This illustrates how H_μ allows us to relate a path π describing expected observable behaviour with paths π' with additional unobservable steps. A similar technique can be used to specify asynchronous variants of classical hyperproperties from the literature like generalised non-interference or observational determinism [Clarkson and Schneider 2010]. The formula $\varphi_{obs} = \forall\pi.\forall\pi'.((\psi_{eqL} \rightarrow \nu X.\mu Y.\psi_{eqL} \wedge \bigcirc_\pi \bigcirc_{\pi'} X) \vee (\neg obs_\pi \wedge \bigcirc_\pi Y) \vee (\neg obs_{\pi'} \wedge \bigcirc_{\pi'} Y))$ with $\psi_{eqL} = \bigwedge_{a \in L} a_\pi \leftrightarrow a_{\pi'}$, for instance, specifies an asynchronous variant of observational determinism, which intuitively states that a system appears to be deterministic to a low security user who can only see propositions from the set L . More specifically, it states that all pairs of executions which agree on the atomic propositions visible to a low security observer at the beginning of their computation agree on these atomic propositions in all observable situations. Compared to its synchronous counterpart, it skips over unobservable states in both executions.

Later, in Section 6 and Section 7, we adapt the two families of approximate analyses, k -synchronous and k -context-bounded analyses, from AAPA to H_μ . We now illustrate the utility of the resulting analyses. Since any violation of the properties φ and φ_{obs} from above occurs after a finite number of context switches, a k -context-bounded analysis with sufficiently large k can be used to disprove these properties. Another example is the property that one cannot locally distinguish whether fragments of a computation belong to one trace or another, or more precisely that for every point in π_1 there is a

point in π_2 such that the next n steps are indistinguishable. In order to make this property's encoding in H_μ more readable, we introduce some additional syntactic sugar: $\mathcal{G}_\pi \psi \equiv \nu X.(\psi \wedge \bigcirc_\pi X)$ expresses that when progressing on π , some property ψ *generally* holds; $\mathcal{F}_\pi \psi \equiv \mu X.(\psi \vee \bigcirc_\pi X)$ expresses that when progressing on π , some property ψ *finally* holds. With this syntactic sugar, the described property can be encoded as $\mathcal{G}_{\pi_1} \mathcal{F}_{\pi_2} \bigwedge_{i \leq n} \bigwedge_{a \in \text{AP}} (\bigcirc_{\pi_2}^i a_{\pi_2} \leftrightarrow \bigcirc_{\pi_1}^i a_{\pi_1})$. In this example, the k -context-bounded analyses with $k \geq 3$ yield precise results. For the usefulness of the k -synchronous analyses, let us consider the following example: imagine we expect a property described by a formula ψ on n paths $\pi_1, \dots, \pi_n$ in a synchronous manner after an initialisation phase that takes a different number of steps on each path. Then, the formula $\mu X.(\bigvee_{i \leq n} (\text{Init}_{\pi_i} \wedge \bigcirc_{\pi_i} X) \vee (\bigwedge_{i \leq n} (\neg \text{Init}_{\pi_i}) \wedge \psi(\pi_1, \dots, \pi_n)))$ describes this expectation in a natural way and can be precisely analysed by the k -synchronous setup, if the lengths of the initialisation phases differ by at most k . While the two of the above properties can also be expressed using synchronous formulas, this most likely requires exponentially larger and less intuitive formulas.

For the definition of the semantics, we introduce some notation. We call a function $\Pi : N \rightarrow \text{Paths}(\mathcal{K})$ for a set of path variables N a path assignment and $\mathcal{V} : \mathcal{X} \rightarrow \text{PA} \rightarrow 2^{\mathbb{N}_0^n}$ a predicate valuation, where we use PA to denote the set of all path assignments. By $\Pi[\pi \mapsto p]$ (resp. $\mathcal{V}[X \mapsto M]$) we denote a path assignment Π' (resp. predicate valuation $\mathcal{V}'$) with $\Pi'(\pi') = \Pi(\pi')$ for $\pi' \neq \pi$ and $\Pi'(\pi) = p$ (resp. $\mathcal{V}'(X') = \mathcal{V}(X')$ for $X' \neq X$, $\mathcal{V}'(X) = M$). Also, we extend the notion of word shifts according to vectors from ω -words to path assignments: given a path assignment Π binding path variables $\pi_1, \dots, \pi_n$ and a vector $v = (v_1, \dots, v_n) \in \mathbb{N}_0^n$, we use $\Pi[v]$ to denote the path assignment where each path π_i is shifted left according to v_i . For the definition of fixpoints, we define an order $\sqsubseteq$ on functions $\text{PA} \rightarrow 2^{\mathbb{N}_0^n}$ such that $\xi \sqsubseteq \xi'$ iff $\xi(\Pi) \subseteq \xi'(\Pi)$ for all Π . In this way, $(\text{PA} \rightarrow 2^{\mathbb{N}_0^n}, \sqsubseteq)$ forms a complete lattice with $\perp \equiv \lambda \Pi. \emptyset$ as its smallest element.

We now define the semantics of H_μ . The semantics is indexed by a $k \in \mathbb{N}_0 \cup \{\infty\}$ restricting how far different paths can diverge from each other. In the semantics indexed by k , we only consider situations where the foremost path is ahead of the rearmost path at most k steps, or more formally, we restrict ourselves to index tuples from the set $G_k^n := \{(j_1, \dots, j_n) \in \mathbb{N}_0^n \mid \forall i, i'. |j_i - j_{i'}| \leq k\}$. Note that for $k = \infty$, G_k^n contains all index tuples from $\mathbb{N}_0^n$. We distinguish between a quantifier semantics and a path semantics for quantified and quantifier-free formulas, respectively. We write $\Pi \models_k^\mathcal{K} \varphi$ to denote that a path assignment Π in the context of a Kripke Structure $\mathcal{K}$ satisfies a quantified formula φ in the k -semantics. This is extended to Kripke Structures: we write $\mathcal{K} \models_k \varphi$ iff $\{\} \models_k^\mathcal{K} \varphi$ for the empty path assignment $\{\}$. Path semantics on the other hand applies to quantifier-free formulas ψ with possibly free predicates. It is defined in the context of predicate valuations and captures which index combinations from G_k^n fulfill the formula for the given path assignment. We write $(j_1, \dots, j_n) \in \llbracket \varphi \rrbracket_k^\mathcal{V}(\Pi)$ for $(j_1, \dots, j_n) \in G_k^n$ to denote that in the context of a predicate valuation $\mathcal{V}$, when we consider a path assignment Π mapping the variables $\pi_1, \dots, \pi_n$ to paths $p_1, \dots, p_n$, the combination of suffixes $p_1[j_1], \dots, p_n[j_n]$ satisfies the formula φ . Since no restrictions are imposed on index tuples for $k = \infty$, we omit the subscript in this situation and write $\mathcal{K} \models_\infty \varphi$ as $\mathcal{K} \models \varphi$, $\Pi \models_\infty^\mathcal{K} \varphi$ as $\Pi \models^\mathcal{K} \varphi$ and $\llbracket \psi \rrbracket_\infty^\mathcal{V}$ as $\llbracket \psi \rrbracket^\mathcal{V}$, respectively.

Definition 4.2 (Quantifier Semantics).

$$\begin{aligned} \Pi \models_k^\mathcal{K} \exists \pi. \varphi &\text{ iff } \Pi[\pi \mapsto p] \models_k^\mathcal{K} \varphi \text{ for some } p \in \text{Paths}(\mathcal{K}) \\ \Pi \models_k^\mathcal{K} \forall \pi. \varphi &\text{ iff } \Pi[\pi \mapsto p] \models_k^\mathcal{K} \varphi \text{ for all } p \in \text{Paths}(\mathcal{K}) \\ \Pi \models_k^\mathcal{K} \psi &\text{ iff } (0, \dots, 0) \in \llbracket \psi \rrbracket_k^\mathcal{V}(\Pi) \text{ for some } \mathcal{V} \end{aligned}$$

for a quantified formula φ and a quantifier-free formula ψ .

Definition 4.3 (Path Semantics).

$$\begin{aligned}
\llbracket a_{\pi_i} \rrbracket_k^{\mathcal{V}} &:= \lambda \Pi. \{(j_1, \dots, j_n) \in G_k^n \mid a \in L(\Pi(\pi_i)(j_i))\} \\
\llbracket X \rrbracket_k^{\mathcal{V}} &:= \mathcal{V}(X) \\
\llbracket \bigcirc_{\pi_i} \psi \rrbracket_k^{\mathcal{V}} &:= \lambda \Pi. \{(j_1, \dots, j_n) \in G_k^n \mid (j_1, \dots, j_i + 1, \dots, j_n) \in \llbracket \psi \rrbracket_k^{\mathcal{V}}(\Pi)\} \\
\llbracket \psi \vee \psi' \rrbracket_k^{\mathcal{V}} &:= \lambda \Pi. \llbracket \psi \rrbracket_k^{\mathcal{V}}(\Pi) \cup \llbracket \psi' \rrbracket_k^{\mathcal{V}}(\Pi) \\
\llbracket \neg \psi \rrbracket_k^{\mathcal{V}} &:= \lambda \Pi. G_k^n \setminus \llbracket \psi \rrbracket_k^{\mathcal{V}}(\Pi) \\
\llbracket \mu X. \psi \rrbracket_k^{\mathcal{V}} &:= \bigcap \{ \xi : PA \rightarrow 2^{G_k^n} \mid \xi \sqsupseteq \llbracket \psi \rrbracket_k^{\mathcal{V}[X \mapsto \xi]} \}
\end{aligned}$$

We now establish some properties of H_μ 's semantics. The first one is that the semantics of $\mu X. \psi$ indeed characterises a fixpoint.

THEOREM 4.4. $\alpha(\xi) := \llbracket \psi \rrbracket_k^{\mathcal{V}[X \mapsto \xi]}$ is monotone for all $k, \mathcal{V}, X$ and ψ in positive normal form.

The Knaster-Tarski fixpoint theorem [Cousot and Cousot 1979; Tarski 1955] then gives a constructive characterisation of the semantics of fixpoint formulas via ordinal numbers:

COROLLARY 4.5. $\llbracket \mu X. \psi \rrbracket_k^{\mathcal{V}}$ is the least fixpoint of α . It can be characterised by its approximants $\bigsqcup_{\kappa \geq 0} \alpha^\kappa(\perp)$ with $\alpha^0(\xi) = \xi$, $\alpha^{\kappa+1}(\xi) = \alpha(\alpha^\kappa(\xi))$ and $\alpha^\lambda(\xi) = \lambda \Pi. \bigcup_{\kappa < \lambda} \alpha^\kappa(\xi)(\Pi)$ where κ ranges over ordinals and λ over limit ordinals.

The second property is that the k -semantics properly approximates the full semantics.

THEOREM 4.6. $\beta(k) := \llbracket \psi \rrbracket_k^{\mathcal{V}}$ is monotone for all ψ in positive normal form and $\mathcal{V}$.

COROLLARY 4.7. For all Kripke Structures $\mathcal{K}$, formulas φ in positive normal form and $k, k' \in \mathbb{N}_0 \cup \{\infty\}$ with $k \leq k'$, we have: $\mathcal{K} \models_k \varphi$ implies $\mathcal{K} \models_{k'} \varphi$.

For the sake of reasoning about satisfiability, we define a variant of the semantics for H_μ formulas on traces in the straightforward way. Instead of considering path assignments $\Pi : N \rightarrow \text{Paths}(\mathcal{K})$, we instead consider trace assignments $\Pi : N \rightarrow \mathcal{T}$ for a set of traces $\mathcal{T}$. Existential and universal quantifiers then quantify traces $t \in \mathcal{T}$ instead of paths $p \in \text{Paths}(\mathcal{K})$. Also, atomic propositions require the proposition to be included in the respective set of atomic propositions directly instead of in the set obtained by applying the labelling function to a state. We write $\mathcal{T} \models \varphi$ to denote that a set of traces $\mathcal{T}$ fulfills a formula. By choosing $\mathcal{T} = \text{Traces}(\mathcal{K})$, the two semantics coincide. For a more formal specification of this semantics variation, we refer the reader to the extended version.

Given these two variants of H_μ 's semantics, we consider two decision problems:

- *Model Checking*: given a closed H_μ formula φ and a Kripke Structure $\mathcal{K}$, does $\mathcal{K} \models \varphi$ hold?
- *Satisfiability*: given a closed H_μ formula φ , is there a non-empty set of traces $\mathcal{T}$ such that $\mathcal{T} \models \varphi$ holds?

For this purpose we first establish a connection between H_μ and AAPA. This allows us to apply the results on AAPA that we have established already. In order to transfer results for restricted classes of AAPA we define corresponding fragments of the logic H_μ next.

Definition 4.8. We call a H_μ formula k -synchronous for a Kripke Structure $\mathcal{K}$ if the following condition holds: $\mathcal{K} \models \varphi$ implies $\mathcal{K} \models_k \varphi$. A 0-synchronous formula is called synchronous.

As a small technical detail, $\llbracket \bigcirc_{\pi} \psi \rrbracket_0^{\mathcal{V}}$, i.e. the 0-semantics of $\bigcirc_{\pi} \psi$ is always empty. Strictly speaking, this renders the above definition useless for the case $k = 0$. In order to cure this defect, we allow a synchronous next operator $\bigcirc \psi$ that advances all direction simultaneously in synchronous formulas. Its semantics is defined straightforwardly.

Definition 4.9 (Synchronous syntactic fragment). A H_μ formula belongs to the synchronous fragment if it uses the synchronous next operator $\bigcirc$ instead of the indexed next operators $\bigcirc_\pi$.

For the next two fragments, we need a notion of *extended syntax tree* of a formula φ . Thereby, we mean the (infinite) tree obtained by repeatedly substituting fixpoint variables by their respective formula in the syntax tree, i.e. for every fixpoint expression $\mu X.\psi$ (or $\nu X.\psi$), replacing every occurrence of X in ψ with ψ . It is easy to see that membership in the two syntactic fragments defined next is decidable in polynomial time.

Definition 4.10 (k -synchronous syntactic fragment). A H_μ formula belongs to the k -synchronous syntactic fragment if the difference between the number of occurrences of $\bigcirc_\pi$ and $\bigcirc_{\pi'}$ for $\pi \neq \pi'$ on any path starting in the root of the formula's extended syntax tree is at most k .

Definition 4.11 (k -context-bounded syntactic fragment). A H_μ formula is from the k -context-bounded syntactic fragment if on every path starting in the root of the formula's extended syntax tree, the number of switches between directions π for $\bigcirc_\pi$ constructs is at most $k - 1$.

5 EXPRESSIVE EQUALITY

In this section we establish the correspondence between AAPA and H_μ , a relation that is essential for transferring the results from Section 3 to H_μ . Crucial for this correspondence is the fact that a path assignment Π for paths $\pi_1, \dots, \pi_n$ over the set of states S can be encoded into a word over the alphabet S^n . Then, free path variables in a formula correspond to components of the input alphabet and free path predicates correspond to holes in the automaton, where a fitting semantics can be plugged in once the predicate is bound. Because of this correspondence, we use the same name for a predicate and its corresponding hole in an automaton.

Thus, given a path assignment Π with $\Pi(\pi_i) = s_i^0 s_i^1 \dots$, we define its translation into a word over the alphabet $\mathcal{S} = S^n$ as $w_\Pi = (s_1^0, \dots, s_n^0)(s_1^1, \dots, s_n^1) \dots \in \mathcal{S}^\omega$. While viewing path assignments as such words, we handle predicate valuations $\mathcal{V}$ by languages $\mathcal{L}(\mathcal{V}(X_j)(\Pi)) = \{(w_\Pi[v] \in \mathcal{S}^\omega \mid v \in \mathcal{V}(X_j)(\Pi))\}$. For these languages to be well-defined, we need to restrict ourselves to *well-formed* valuations with the property that for all vectors $v, v' \in \mathbb{N}_0^n$, path assignments Π, Π' and predicates X , we have that $\Pi[v] = \Pi'[v']$ implies $v \in \mathcal{V}(X)(\Pi)$ iff $v' \in \mathcal{V}(X)(\Pi')$. However, we can show via induction that only valuations with this property occur during fixpoint iterations of H_μ .

We introduce the notion of $\mathcal{K}$ -equivalence:

Definition 5.1 ($\mathcal{K}$ -equivalence). Given a Kripke Structure $\mathcal{K} = (S, s_0, \delta, L)$, a H_μ formula ψ over $\{\pi_1, \dots, \pi_n\}$ with free predicates $X_1, \dots, X_m$ and an alternating (asynchronous) parity automaton $\mathcal{A}$ with holes $X_1, \dots, X_m$ over the alphabet S^n , we call $\mathcal{A}$ $\mathcal{K}$ -equivalent to ψ , iff the following condition holds: for all path assignments Π , well-formed predicate valuations $\mathcal{V}$ and vectors $v \in \mathbb{N}_0^n$, we have $v \in \llbracket \psi \rrbracket^{\mathcal{V}}(\Pi)$ iff $w_\Pi[v] \in \mathcal{L}(\mathcal{A}[X_1 : \mathcal{L}(\mathcal{V}(X_1)(\Pi)), \dots, X_m : \mathcal{L}(\mathcal{V}(X_m)(\Pi))])$.

The definition of $\mathcal{K}$ -equivalence is straightforwardly extended to quantified formulas φ : $\mathcal{A}$ is called $\mathcal{K}$ -equivalent to φ iff for all Π the statements (i) $\Pi \models^{\mathcal{K}} \varphi$ and (ii) $w_\Pi \in \mathcal{L}(\mathcal{A})$ are equivalent.

This notion allows us to formulate the following theorem:

THEOREM 5.2. *Let $\mathcal{K}$ be a Kripke Structure.*

- (1) *For every quantifier-free H_μ formula ψ in positive normal form, there is an AAPA $\mathcal{A}_\psi$ of linear size in $|\psi|_d$ (and hence also in $|\psi|_t$) such that $\mathcal{A}_\psi$ is $\mathcal{K}$ -equivalent to ψ .*
- (2) *For every AAPA $\mathcal{A}$ over the alphabet Σ , there is a quantifier-free H_μ formula $\psi_{\mathcal{A}}$ with $|\psi_{\mathcal{A}}|_d$ linear and $|\psi_{\mathcal{A}}|_t$ exponential in $|\mathcal{A}|$ such that $\mathcal{A}$ is $\mathcal{K}$ -equivalent to $\psi_{\mathcal{A}}$.*

We dedicate the following two subsections to the constructions underlying the proof of this theorem. A detailed correctness proof for these constructions can be found in the extended version.

5.1 From H_μ to AAPA: Construction for Theorem 5.2, Part 1

Intuitively, the AAPA $\mathcal{A}_\psi$ has a node for each node in ψ 's syntax DAG. For most constructs, the node can straightforwardly check the semantics of the formula, either directly or by transitioning to nodes for subformulas in a suitable manner. The most interesting cases are those for predicates and fixpoint expressions. As in the definition of $\mathcal{K}$ -equivalence, free predicates correspond to holes in the automaton and are thus translated to such. In the same definition, bound predicates do not occur as holes in the automaton; indeed, the construction for fixpoints fills the corresponding holes. This is done by a backwards edge to the start of the automaton. Taking this backwards edge corresponds to one unfolding of a fixpoint iteration. Depending on whether we have a least or greatest fixpoint, the iteration can be performed a finite or infinite number of times. This is captured by assigning the predicate state a priority that is lower than any other in the automaton and that is odd for least and even for greatest fixpoints.

We inductively construct the AAPA $\mathcal{A}_\psi = (Q, \rho_0, \rho, \Omega)$ for a formula ψ . In each step, we assume that automata for all subformulas of ψ are already constructed. The construction is linear in $|\psi|_d$ since in each step at most a constant number of states is added to the already existing partial automata and automata for syntactically (and thus semantically) equivalent subformulas can be shared. We now describe each case of the inductive construction. Recall that we assume that states *true* and *false* are present in every automaton. We therefore do not mention them explicitly.

Case $\psi = a_{\pi_i}$. For atomic propositions, we let the automaton enter an accepting self-loop if the proposition is fulfilled at the current index and a rejecting self-loop otherwise. We set $Q = \{a_{\pi_i}\}$ and $\rho_0 = a_{\pi_i}$. Furthermore, we let $\rho(a_{\pi_i}, s, d) = \text{true}$ if $a \in L(s)$ and $d = i$ and let $\rho(a_{\pi_i}, s, d) = \text{false}$, otherwise. For the priority assignments, we set $\Omega(a_{\pi_i}) = 1$. The case $\psi = \neg a_{\pi_i}$ is analogous.

Case $\psi = X$. A predicate is transformed into a hole in the automaton: formally, we set $Q = \{X\}$, $\rho_0 = X$, $\rho(X, s, d) = \perp$ and $\Omega(X) = \perp$.

Case $\psi = \psi_1 \vee \psi_2$. From the induction hypothesis, we have automata $\mathcal{A}_{\psi_i} = (Q_i, \rho_{0,i}, \rho_i, \Omega_i)$ for the formulas ψ_i . We set $Q = Q_1 \cup Q_2$ and $\rho_0 = \rho_{0,1} \vee \rho_{0,2}$. The transition function is induced by the transition functions of the automata $\mathcal{A}_i$, i.e. $\rho(q, s, d) = \rho_i(q, s, d)$ for $q \in Q_i$. For the priorities, we pick $\Omega(q) = \Omega_i(q)$ for $q \in Q_i$. This automaton accepts iff one of the automata $\mathcal{A}_i$ accepts.

Case $\psi = \bigcirc_i \psi_1$. From the induction hypothesis, we obtain $\mathcal{A}_{\psi_1} = (Q_1, \rho_{0,1}, \rho_1, \Omega_1)$ for ψ_1 . We then set $Q = \{\psi\} \cup Q_1$ and $\rho_0 = \psi$. Furthermore, we set $\rho(\psi, s, d) = \rho_{0,1}$ for $d = i$ and $\rho(\psi, s, d) = \text{false}$ otherwise. For the states $q \in Q_1$ and priorities, we choose $\rho(q, s, d) = \rho_1(q, s, d)$, $\Omega(\psi) = 0$ and $\Omega(q) = \Omega_1(q)$ for $q \in Q_1$. The choice of priority for state ψ does not matter, since every run visiting this node will either visit it only a finite number of times or visit states with lower or equal priority infinitely often. This is ensured by the construction for fixpoint formulas.

Case $\psi = \mu X. \psi_1$ or $\psi = \nu X. \psi_1$. The induction hypothesis gives us an automaton $\mathcal{A}_{\psi_1} = (Q_1, \rho_{0,1}, \rho_1, \Omega_1)$ for ψ_1 . We can assume w.l.o.g. that there is only one hole for the path predicate X , since all holes have the same behaviour. Let $p := \min\{\Omega_1(q_i) \mid q_i \in Q_1\}$. Let $p_{\text{even}} = p$ and $p_{\text{odd}} = p - 1$ if p is even; otherwise, let $p_{\text{even}} = p - 1$ and $p_{\text{odd}} = p$. Intuitively, p_{even} (resp. p_{odd}) is an even (resp. odd) lower bound on the priorities occurring in $\mathcal{A}_{\psi_1}$. If a priority in the automaton below is negative, all priorities are shifted by a multiple of 2 such that the lowest priority is 0 or 1. We set $Q = Q_1$, $\rho_0 = X$, $\rho(X, s, d) = \rho_1(\rho_{0,1}, s, d)$ and $\rho(q, s, d) = \rho_1(q, s, d)$ for $q \in Q_1 \setminus \{X\}$. For the priorities, we choose $\Omega(q) = \Omega_1(q)$ for $q \in Q_1 \setminus \{X\}$. The state X is assigned the priority $\Omega(X) = p_{\text{odd}}$ if ψ is a μ formula and $\Omega(X) = p_{\text{even}}$ otherwise. In this definition, we use $\rho_1(\rho_{0,1}, s, d)$ to denote a variant of $\rho_{0,1}$ in which every occurrence of q is substituted by $\rho_1(q, s, d)$. The automaton evaluates the subformula ψ_1 , but switches to the start in case a predicate X is encountered. The choice of priority for X reflects that an infinite unfolding of a μ -formula should lead to a rejecting run, while an infinite unfolding of a ν -formula should lead to an accepting run.

5.2 From AAPA to H_μ : Construction for Theorem 5.2, Part 2

Given an AAPA $\mathcal{A} = (Q, \rho_0, \rho, \Omega)$, we construct formulas ψ_i^h by induction on i . Our construction is inspired by a construction from [Bozzelli 2007] in the context of a fixpoint logic for visibly pushdown languages.¹ We fix an ordering $q_1, \dots, q_n$ on states of Q such that $\Omega(q_i) \geq \Omega(q_j)$ for $i < j$ for non-hole states and holes are the last m states of the ordering. Intuitively, the formula ψ_0^h describes the local behaviour of a state q_h and the formula ψ_i^h expresses the existence of an accepting run of $\mathcal{A}$ starting in q_h , where only states with a priority higher than $\Omega(q_i)$ are visited infinitely often. For each state q_i , we introduce a predicate X_i which is bound in the i -th step of the inductive construction. In the construction, i ranges from 0 to $n - m$ and h ranges from 1 to n . Therefore, when the construction is finished, only the hole states of $\mathcal{A}$ remain as unbound predicates in the formula and we can choose $\rho_0[q_1/\psi_{n-m}^1] \dots [q_n/\psi_{n-m}^n]$ as the desired formula.

Construction of ψ_0^h . For holes, that is for $h > n - m$, the formula ψ_0^h is given as the predicate X_h . For non-holes, that is for $h \leq n - m$, we first construct $\hat{\rho}(q_h, \sigma, d)$ from $\rho(q_h, \sigma, d)$ by substituting every occurrence of a state q_l with X_l . ψ_0^h is then given as $\bigvee_{\sigma \in \Sigma} \bigvee_{d \in M} (\sigma_{\pi_d} \wedge \bigcirc_d \hat{\rho}(q_h, \sigma, d))$ to describe that some σ is currently being read in direction d and in the next step we are in some combination of states satisfying $\rho(q_h, \sigma, d)$ after moving on in direction d .

Construction of ψ_i^h for $i > 0$. We assume that ψ_{i-1}^h for all h is already constructed. As a first step, we construct the formula ψ_i^i by binding the predicate X_i . We differentiate two cases based on the priority of state q_i . If $\Omega(q_i)$ is odd, then ψ_i^i is given as $\mu X_i. \psi_{i-1}^i$. In the other case, where $\Omega(q_i)$ is even, we construct ψ_i^i as $\nu X_i. \psi_{i-1}^i$. Then we construct ψ_i^h for all $h \neq i$ by substituting ψ_i^i for every occurrence of the predicate X_i in the previous formula, that is: $\psi_i^h = \psi_{i-1}^h[X_i/\psi_i^i]$.

Size of the construction. In the base case of the construction, a number of syntax DAG nodes is created that is linear in the size of ρ . In the inductive step, only a single node is added, that is the node for the fixpoint in ψ_i^i . The other formulas in this step, namely ψ_i^h for $h \neq i$, can be obtained from ψ_{i-1}^h by redirecting edges in the syntax DAG. Since only a single node is added in each step and there is a step for each node of the automaton, $|Q|$ nodes are added in total. Finally, when combining the formulas, at most $|\rho_0|$ nodes are added. Thus, we have $|\psi_{\mathcal{A}}|_d = O(|\rho| + |Q| + |\rho_0|) = O(|\mathcal{A}|)$.

For the second measure of size, $|\psi_{\mathcal{A}}|_t$, the formula must be represented as a syntax tree. Here, the substitution in the construction of ψ_i^h for $h \neq i$ cannot be performed without adding nodes to the syntax tree since the nodes representing ψ_i^i must be duplicated. This results in a linear increase in each of the steps, which leads to an overall exponential size increase in the worst case.

5.3 The Consequences of Theorem 5.2

Since the emptiness problem for AAPA is Σ_2^1 -hard and we can effectively build a H_μ formula for every AAPA and check it against a structure that generates Σ^ω , we obtain:

THEOREM 5.3. *Model checking H_μ against a Kripke model is Σ_2^1 -hard.*

Likewise, the emptiness problem of AAPA can be reduced to the satisfiability problem for H_μ :

THEOREM 5.4. *The satisfiability problem for H_μ is Σ_2^1 -hard.*

6 APPROXIMATE MODEL CHECKING AND DECIDABLE FRAGMENTS

The Σ_2^1 -hardness of model checking and satisfiability imply that these problems can neither be algorithmically solved nor can a complete approximation procedure (e.g. via bounded model

¹More precisely, the construction is found in the appendix of an extended version of that paper that was provided to us by the author. While the paper states that the construction can be performed in linear time and thus with at most linear size increase, it is not discernable how the construction can be performed without an at worst exponential blowup when using the syntax tree of a formula as the basis for measuring its size.

Table 1. Complexity results for model checking.

Fragment	Complexity	
full	UNDECIDABLE	
synchronous	$\text{SPACE}(g(d, \varphi _d))$ -complete	$\text{NSPACE}(g(d-1, \mathcal{K}))$ -complete
k -synchronous	$\text{SPACE}(g(d+1, \varphi _d))$ -complete	$\text{NSPACE}(g(d-1, \mathcal{K}))$ -complete
k -context-bounded	$\text{SPACE}(g(d+k-2, \varphi _d))$	$\text{NSPACE}(g(d-1, \mathcal{K}))$
	$\text{SPACE}(g(\tilde{d}+k-2, \varphi _d))$ -hard	$\tilde{d} = \min(d, 1)$

checking) be developed for the full logic. However, the concepts of k -synchronicity and k -context-boundedness have already proved to lead to decidability for AAPA. We therefore consider similarly restricted analyses for H_μ and explain how they relate to the matching analyses of AAPA.

For model checking H_μ against Kripke structures, two main problems have to be solved: the quantifier-free part of the input formula has to be suitably represented and the representation has to take the quantifiers into account. As representation, we use the AAPA derived from the construction of [subsection 5.1](#). The main idea for handling quantifiers is to take a nondeterministic automaton representing the inner formula and to form a product with the Kripke structure in order to check for the existence of appropriate paths in the structure. Additionally, the automata need to be complemented in the case of universal quantifiers. Both k -synchronous and k -context-bounded AAPA are suitable for this purpose as we can dealternate them to NPA. In this way, we lift our approximate analyses from automata to formulas. Moreover, we even obtain precise results on several fragments of the logic. This stems from the following facts:

- THEOREM 6.1.** • *The automaton $\mathcal{A}_\psi$ for a synchronous H_μ formula ψ can be transformed into an APA of asymptotically the same size.*
 • *The automaton $\mathcal{A}_\psi$ for a k -synchronous H_μ formula ψ is a k -synchronous AAPA.*
 • *The automaton $\mathcal{A}_\psi$ for a k -context-bounded H_μ formula ψ is a k -context-bounded AAPA.*

The following fact allows an application of the k -synchronous analysis to formulas:

THEOREM 6.2. *For a quantifier-free formula ψ , a path assignment Π and a well-formed predicate valuation $\mathcal{V}$, the automaton $\mathcal{A}_\psi$ from [Theorem 5.2](#) has a k -synchronous accepting run over $w_\Pi[v]$ with holes filled according to $\mathcal{V}$ iff $v \in \llbracket \psi \rrbracket_k^{\mathcal{V}}(\Pi)$.*

This generalises the definition of $\mathcal{K}$ -equivalence. The proof relies on the observation that each node in $\mathcal{A}_\psi$'s runs corresponds to a subformula of ψ and that for each such subformula, the offset counters that appear in an accepting run correspond to the tuples in the semantics of ψ . Thus, the restriction induced by the k -semantics and k -synchronous runs deliver the same tuples.

[Theorem 6.2](#) implies that we can use the k -synchronous analysis underlying [Theorem 3.12](#) on the automaton $\mathcal{A}_\psi$ for a formula ψ to determine that formula's k -semantics. Since the approximations given by the k -semantics improve with increasing k for quantifier-free ([Theorem 4.6](#)) as well as quantified ([Corollary 4.7](#)) formulas, this procedure yields increasingly precise analyses of arbitrary H_μ formulas. Indeed, any approximate analysis parameterised in some k on AAPA that supplies us, for increasing k , with an increasing number of offsets v from which there is an accepting run for a fixed w_Π can be used to approximate H_μ as well. This is due to the fact that such an analysis directly corresponds to a monotone semantics approximator as in [Theorem 4.6](#) and is thus applicable in a proof like the one from [Corollary 4.7](#). In particular, this applies to the k -context-bounded analysis underlying [Theorem 3.17](#) using the above argument about the correspondence between offset counters and tuples in the semantics.

6.1 Constructions for Quantifiers

Given a k -synchronous or k -context-bounded AAPA for a quantified formula, the quantifier can be handled in the same way as for HyperLTL in [Finkbeiner et al. 2015].

Construction for existential quantifiers. We perform the construction for a formula $\exists\pi_{n+1}.\varphi$ and a structure $\mathcal{K} = (S, s_0, \delta, L)$. Our construction simulates one input component of $\mathcal{A}_\varphi$ in the state space of $\mathcal{A}_{\exists\pi_{n+1}.\varphi}$. In order to allow us to do so, we assume $\mathcal{A}_\varphi$ to be given as an NPA. For the fragments we consider, a translation to NPA is possible, as seen in Section 3. Therefore, $\mathcal{A}_\varphi$ has the form $(Q_1, q_{0,1}, \rho_1, \Omega_1)$ and input alphabet S^{n+1} . We construct the NPA $\mathcal{A}_{\exists\pi_{n+1}.\varphi} = (Q, q_0, \rho, \Omega)$ with input alphabet S^n , states $Q = Q_1 \times S$, initial state $q_0 = (q_{0,1}, s_0)$, transition function $\rho((q, s), s) = \{(q', s') \in Q_1 \times S \mid q' \in \rho_1(q, s + s), s' \in \delta(s)\}$, and priorities $\Omega((q, s)) = \Omega_1(q)$.

As mentioned, the last component of $\mathcal{S}$ is now simulated in the last component of the state space of $\mathcal{A}_{\exists\pi_{n+1}.\varphi}$. Simultaneously this component makes sure that transitions are taken according to the transition function δ of $\mathcal{K}$. Choosing $(q_{0,1}, s_0)$ as the starting state ensures that the path which is simulated in this way starts in the initial state s_0 of $\mathcal{K}$.

Construction for universal quantifiers. In order to handle universal quantifiers, we complement the automaton $\mathcal{A}_\varphi$ by swapping $\wedge$ and $\vee$ in transitions and increasing all priorities by one. We then construct the automaton $\exists\pi_{n+1}.\neg\varphi$, and complement it again by the same procedure. Note that when several of these constructions are combined, double negations can be cancelled out to avoid unnecessary complementation constructions.

6.2 Complexity Analysis

In order to analyse the size of the resulting automata, we need a notion of alternation depth, which is defined as the number of switches between $\exists$ and $\forall$ quantifiers in the quantifier-prefix of the formula. Since our quantifier construction involves transforming the automaton into an NPA, we formulate our evaluation of $\mathcal{A}_\varphi$'s size for the resulting NPA:

LEMMA 6.3. *The NPA $\mathcal{A}_\varphi$ for a closed formula φ with alternation depth d has size:*

- $O(g(d, |\mathcal{K}|) \cdot g(1, |\varphi|_d))$ in a synchronous analysis,
- $O(g(d, |\mathcal{K}|) \cdot g(2, |\varphi|_d))$ in a k -synchronous analysis, and
- $O(g(d, |\mathcal{K}|) \cdot g(k-1, |\varphi|_d))$ in a k -context-bounded analysis.

For the decision problems corresponding to our analyses, we obtain the following complexity results. The upper bounds also apply for the respective analyses.

THEOREM 6.4. *Model checking a closed synchronous H_μ formula φ with alternation depth d against a Kripke Structure is complete for $\text{SPACE}(g(d, |\varphi|_d))$, $\text{SPACE}(g(d, |\varphi|_t))$ and $\text{NSPACE}(g(d-1, |\mathcal{K}|))$ ².*

PROOF. The upper bounds follow from Lemma 6.3 and the NLOGSPACE complexity of emptiness tests on NPA. For hardness, we can reduce from HyperLTL Model Checking [Rabe 2016] to obtain the desired results. $\square$

THEOREM 6.5. *Model checking a closed k -synchronous H_μ formula φ with alternation depth d against a Kripke Structure is complete for $\text{SPACE}(g(d+1, |\varphi|_d))$ and $\text{NSPACE}(g(d-1, |\mathcal{K}|))$.*

PROOF. The upper bounds follow from Lemma 6.3 and the NLOGSPACE complexity of the emptiness problem for NPA.

²For the case of $d = 0$, we use the definition from [Finkbeiner et al. 2015], where $\text{NSPACE}(g(-1, n))$ was defined as NLOGSPACE. For $d > 0$, we can use Savitch's Theorem to see that the problems are actually complete for the deterministic space classes.

For hardness, we reduce from the acceptance problem for $g(k, p(n))$ space bounded deterministic Turing machines. The reduction is based on Sistla's classical yardstick construction for the satisfiability problem of Quantified Propositional Temporal Logic (QPTL) [Sistla 1983] and the adaptation of that reduction to HyperLTL [Rabe 2016]. More concretely, in the reduction for HyperLTL, a formula $\varphi_{k,m}(P_x, P_y)$ is constructed such that P_x and P_y are true at exactly one point and the points at which both propositions are fulfilled are exactly $N_{k,m}$ steps away where $N_{k,m} \geq g(k, m)$. While in that reduction, an alternation free formula is constructed for a polynomial $N_{k,m}$, we can instead construct a quantifier free formula for an exponential $N_{k,m}$ by building a 2-synchronous AAPA as in the proof of Theorem 3.12 that can check for two propositions whether they are exponentially many indices away from each other. For this AAPA, we can obtain a 2-synchronous H_μ formula in polynomial time. We can then inductively construct the formulas as in the proof for HyperLTL and translate it to H_μ . Note that since HyperLTL progresses on different paths synchronously, the inductive construction can easily be adapted to preserve 2-synchronicity of the formula. Accordingly, we need one quantifier alternation less in our H_μ formula than in the corresponding HyperLTL or QPTL formulas to build a yardstick of length $N_{k,m}$ and obtain the desired lower bound.

The hardness claim for fixed formulas can be obtained by a direct reduction from HyperLTL model checking [Rabe 2016]. $\square$

By the space hierarchy theorem, the completeness for $\text{SPACE}(g(d+1, |\varphi|_d))$ implies that when measuring size of formulas φ by $|\cdot|_t$, at least space $O(g(d, |\varphi|_t))$ is needed. Note that this does not imply hardness for the corresponding class since the implication is based on an exponential time reduction, but hardness requires polynomial time reductions. A similar reasoning applies to later theorems in which completeness results based on $|\cdot|_d$ are given.

THEOREM 6.6. *Model checking a closed k -context-bounded H_μ formula φ with alternation depth d against a Kripke Structure is in $\text{SPACE}(g(d+k-2, |\varphi|_d))$ and $\text{NSPACE}(g(d-1, |\mathcal{K}|))$ and is hard for $\text{SPACE}(g(\tilde{d}+k-2, |\varphi|_d))$ where $\tilde{d} = \min(d, 1)$.*

PROOF. The upper bounds follow from Lemma 6.3 and the NLOGSPACE complexity of emptiness tests on NPA. For hardness, we show the claim for $d = 0$ and $d \geq 1$ separately.

In the first case where $d = 0$, we reduce from the emptiness problem for k -context-bounded AAPA. More specifically, an AAPA $\mathcal{A}$ is non-empty iff $\mathcal{K} \models \exists \pi_1 \dots \exists \pi_n. \psi_{\mathcal{A}}$ for the structure $\mathcal{K}$ whose traces correspond to arbitrary words over $\mathcal{A}$'s input alphabet.

In the second case, where $d \geq 1$, we reduce from the acceptance problem for $g(k-1, p(n))$ space bounded deterministic Turing machines. Similar to the proof of Theorem 3.17, we show a more general result about regular transductions on level $k-1$ Stockmeyer encoded words as in Lemma 3.18. For a regular transducer and two regular language acceptors of size polynomial in n , we construct a Kripke Structure $\mathcal{K}$ and a formula φ such that $\mathcal{K} \models \varphi$ iff there is a sequence of level $k-1$ Stockmeyer encoded words $w_0, w_1, \dots, w_l$ such that

- (a) w_0 is accepted by the first regular language acceptor,
- (b) w_l is accepted by the second regular language acceptor, and
- (c) for all $i < l$, w_{i+1} is obtained from w_i by applying the regular transducer $\mathcal{T} = (Q, q_o, \gamma)$.

As the first step, we construct the Kripke Structure $\mathcal{K}$. Its set of atomic propositions is given as $AP = \{0, 1\} \cup \{[i,]_i | i \leq k\} \cup Q$, where $\{0, 1\}$ is the alphabet for the word itself, $\{[i,]_i | i \leq k\}$ is used for the Stockmeyer encodings and Q is the set of states of the regular transducer. For this set of atomic propositions, we choose $\mathcal{K}$ as the structure that produces arbitrary traces over AP .

To make the formula φ more readable, we introduce some modalities as syntactic sugar that can easily be encoded in H_μ . First, we use the LTL-like modality $\mathcal{G}_\pi$ that was previously defined as well as the Until-modality $\mathcal{U}_\pi$ that can be defined analogously. Also, we introduce two variations of

$\bigcirc_\pi$ namely $\bigcirc_\pi^{symbol}$ and $\bigcirc_\pi^{word}$. In a sequence of Stockmeyer encoded words as inspected here, $\bigcirc_\pi^{symbol}\psi$ expects ψ to hold at the next encoded symbol and $\bigcirc_\pi^{word}\psi$ expects ψ to hold at the start of the next encoded word in the sequence. These modalities can straightforwardly be encoded in H_μ . Finally, we use $\bigcirc_\pi^{word}\psi$ to encode modalities $\mathcal{G}_\pi^{word}\psi$ and $\mathcal{F}_\pi^{word}\psi$ which express that ψ holds for each or one encoded word on π , respectively.

To increase readability of the formula even further, we also introduce some auxiliary formulas:

- $\psi_{stock}^{k-1}(\pi, \pi')$ is obtained from [Lemma 3.18](#). It expresses that the next word of length $g(k-1, p(n))$ on π is a level $k-1$ Stockmeyer encoded word using π' and $k-1$ context switches
- $\psi_{same}^{k-2}(\pi, \pi')$ is obtained from [Lemma 3.19](#) and expresses that the first level $k-2$ Stockmeyer encoded words on π and π' match each other with $k-1$ context switches (starting with π)
- $\psi_{start}(\pi)$ and $\psi_{end}(\pi)$ state that the first encoded word on π is accepted by the first and second regular language acceptor respectively
- $\psi_{enc}(\pi)$ checks that π respects a specific encoding. This means:
 - On every index, exactly one atomic proposition from $\{0, 1\} \cup \{[i,] | i \leq k\}$ is true
 - On every index of an encoded word, exactly one atomic proposition from Q is true
 - The first Q -proposition in every encoded word is q_0

Now we are able to construct the formula φ :

$$\exists \pi_{seq}. \exists \pi_{help}. \forall \pi_{index}. \psi_{stock}^{k-2}(\pi_{index}, \pi_{help}) \rightarrow \quad (1)$$

$$\psi_{enc}(\pi_{seq}) \wedge \mathcal{G}_{\pi_{seq}}^{word} \psi_{stock}^{k-1}(\pi_{seq}, \pi_{help}) \wedge \quad (2)$$

$$\psi_{start}(\pi_{seq}) \wedge \mathcal{F}_{\pi_{seq}}^{word} \psi_{end}(\pi_{seq}) \wedge \quad (3)$$

$$\mathcal{G}_{\pi_{seq}} \psi_{same}^{k-2}(\pi_{seq}, \pi_{index}) \rightarrow \quad (4)$$

$$\bigcirc_{\pi_{seq}}^{symbol} \bigvee_{i \in \{0,1\}, q \in Q} i_{\pi_{seq}} \wedge q_{\pi_{seq}} \wedge \bigvee_{(q', i') \in \gamma(q, i)} \bigcirc_{\pi_{seq}}^{symbol} q'_{\pi_{seq}} \wedge$$

$$\bigcirc_{\pi_{seq}}^{word} \neg \psi_{same}^{k-2}(\pi_{seq}, \pi_{index}) \mathcal{U}_{\pi_{seq}} \psi_{same}^{k-2}(\pi_{seq}, \pi_{index}) \wedge \bigcirc_{\pi_{seq}}^{symbol} i'_{\pi_{seq}}$$

Intuitively, this formula quantifying the three paths π_{seq} , π_{help} and π_{index} can be understood in the following way: on π_{seq} , we check for the existence of the sequence of Stockmeyer encoded words. We use π_{help} as a path yielding indices of lower level encodings to be used for checking the presence of Stockmeyer encodings as in [Lemma 3.18](#) or more specifically the translation of the AAPA referred to in that lemma into H_μ . The path π_{index} is universally quantified and is used to obtain all Stockmeyer indices for a level k encoding. The condition in (1) ensures that it is indeed a level $k-1$ Stockmeyer encoded word and each such word over $\{0, 1\}$ represents an index for a level k encoding. In (2), we ensure that π_{seq} has the required encoding. Part (3) of the formula expresses the two conditions (a) and (b) that the sequence of words must fulfill. Finally, (4) expresses condition (c) of the sequence, i.e. that each word is obtained from the previous one by applying the regular transducer. This is done by checking the transductions for each index of the words separately (which is achieved by the universal quantification over π_{index}): whenever we find the index we are currently checking, we determine the current symbol of the word and the state the transducer is in. We choose one of $\mathcal{T}$'s possible transitions for the current symbol and state to determine its output and next state. The next state is then expected at the next symbol of the current encoded word. The output on the other hand is expected the next time this index occurs in the sequence. $\square$

A few remarks are in order about these last results. First, since the formula constructed for the reduction for [Theorem 6.6](#) works in a similar way as the automaton from [Lemma 3.18](#), one could

Table 2. Complexity results for satisfiability when representing a formula with $rep_d(\cdot)$.

Fragment	Complexity
full	UNDECIDABLE
alternation free synchronous	PSPACE-complete
alternation free k -synchronous	EXPSPACE-complete
alternation free k -context-bounded	$(k - 2)$ EXPSPACE-complete
$\exists^* \forall^*$ synchronous	EXPSPACE-complete
$\exists^* \forall^*$ k -synchronous	EXPSPACE-complete
$\exists^* \forall^*$ k -context-bounded	$(k - 2)$ EXPSPACE-complete

ask whether the hardness estimates can be inductively lifted to $\tilde{d} = d$ in a similar way as was done in [Lemma 3.18](#). This would require the formula ψ_{same} used in our construction to handle longer nested index encodings. As the number of context switches can not be increased further, this seems to require quantification. Indeed, by adding suitable prenex alternating quantifiers, a corresponding formula could be defined. However, this formula could no longer be used inside fixpoints since this would lead to formulas with non-prenex quantifiers and it is unclear whether such alternating quantifiers could be moved to the front of the formula without changing the semantics.

Secondly, while the inability to inductively lift the hardness result to an arbitrary number of quantifier alternations raises the question whether quantifier alternations increase the complexity of a context bounded analysis at all, our result shows that this is at least the case for the first one. A complete answer to this question remains for future work.

Finally, the k -synchronous and k -context-bounded analyses can be combined by interpreting the quantifier free part of the given formula as a boolean combination of subformulas each of which is analysed with one of these analyses. Our setup can easily handle this since the resulting automata for the analyses of these formulas are all synchronous automata which can be combined straightforwardly. In this setting, the hardness proof for the k -context-bounded analysis of automata can be lifted with quantifier alternations in the same way as in the proof of [Theorem 6.5](#) since the subformulas added in the inductive step of the proof are synchronous. This yields a hardness result for the combined analysis that matches the upper bound for the k -context-bounded analysis.

7 APPROXIMATE SATISFIABILITY CHECKING FOR FRAGMENTS OF THE LOGIC

Since certain combinations of path quantifiers lead to undecidability already for synchronous hyperlogics like HyperLTL [[Finkbeiner and Hahn 2016](#)], we consider satisfiability for formulas with restricted quantifier prefixes only. We say that a formula is in the $\exists^*$ fragment of H_μ if it only has $\exists$ quantifiers in its quantifier prefix. Analogously, if a formula's quantifier prefix consists only of $\forall$ quantifiers, it is in the $\forall^*$ fragment. Finally, we say that a formula is in the $\exists^* \forall^*$ fragment if its quantifier prefix has the form $\exists \pi_1 \dots \exists \pi_n \forall \pi'_1 \dots \forall \pi'_m$. Note that the same restrictions have been considered in [[Finkbeiner and Hahn 2016](#)] for HyperLTL in order to obtain fragments for which the satisfiability problem is decidable. However, as our proof of [Theorem 5.4](#) shows, a quantifier prefix with just existential quantifiers suffices for undecidability in the case of H_μ , if the quantifier-free part of the formula is not restricted or approximated as well. Thus, just like for model checking, we make use of the approximate analyses developed in [Section 3](#) for the quantifier-free part of the formula. We now present complexity results for the decision problems corresponding to the resulting approximate analyses. The upper bounds apply to these analyses as well. An overview of our results can be found in [Table 2](#).

THEOREM 7.1. *The satisfiability problem for the $\exists^*$ - and $\forall^*$ -fragment of*

- (1) *synchronous H_μ is PSPACE-complete*
- (2) *k -synchronous H_μ is EXPSPACE-complete*
- (3) *k -context-bounded H_μ is $(k - 2)$ EXPSPACE-complete*

when using $\text{rep}_d(\cdot)$ to represent formulas. The statement in (1) also holds for $\text{rep}_t(\cdot)$.

PROOF. Using [Theorem 5.2](#), the three problems become interreducible to the emptiness problems for the corresponding AAPA restrictions. This is done in the following way:

(i) An $\exists^*$ formula $\exists\pi_1\ldots\exists\pi_n.\psi$ is satisfiable iff $\mathcal{A}_\psi$ is non-empty. (ii) A $\forall^*$ formula $\forall\pi_1\ldots\forall\pi_n.\psi$ is satisfiable iff $\mathcal{A}_{\neg\psi}$ is empty. (iii) An AAPA $\mathcal{A}$ is non-empty iff the formula $\exists\pi_1\ldots\exists\pi_n.\psi_{\mathcal{A}}$ is satisfiable. (iv) An AAPA $\mathcal{A}$ is empty iff the formula $\forall\pi_1\ldots\forall\pi_n.\neg\psi_{\mathcal{A}}$ is satisfiable.

This yields both upper and lower bounds for each of the problems using $\text{rep}_d(\cdot)$ as representation. The hardness result for item (1) and $\text{rep}_t(\cdot)$ can be obtained by reducing from the satisfiability problem for alternation-free HyperLTL instead. $\square$

THEOREM 7.2. *The satisfiability problem for the $\exists^*\forall^*$ -fragment of*

- (1) *synchronous H_μ is EXPSPACE-complete*
- (2) *k -synchronous H_μ is EXPSPACE-complete*
- (3) *k -context-bounded H_μ is $(k - 2)$ EXPSPACE-complete*

when using $\text{rep}_d(\cdot)$ to represent formulas.

PROOF. For the upper bounds, we adapt the idea from [\[Finkbeiner and Hahn 2016\]](#) to test satisfiability of HyperLTL formulas only with a minimal set of traces: the set of traces chosen for the existential quantifiers. In such a set, the universal quantifiers can be instantiated in every possible combination and thus eliminated. More specifically, a HyperLTL formula $\exists\pi_1\ldots\exists\pi_n\forall\pi'_1\ldots\forall\pi'_m.\psi$ is transformed into the equisatisfiable formula $\exists\pi_1\ldots\exists\pi_n.\bigwedge_{j_1=1}^n\ldots\bigwedge_{j_m=1}^n\psi[\pi_{j_1}/\pi'_1]\ldots[\pi_{j_m}/\pi'_m]$, which can then be tested for satisfiability using the same method as for $\exists^*$ formulas.

However, in our setting, a direct instantiation of the universal quantifiers via substitution is possible only for synchronous H_μ formulas since tests for atomic propositions can occur with different offsets otherwise. For k -synchronous and k -context-bounded formulas, we have to incorporate the conjunctive test for each of these *arrangements*, i.e. each set of substitutions, directly into our analysis of the corresponding automata.

In the case of k -synchronous formulas $\exists\pi_1\ldots\exists\pi_n\forall\pi'_1\ldots\forall\pi'_m.\psi$, we first create n^m copies of $\mathcal{A}_\psi$ on $n + m$ input words, one for each arrangement. Then, for each arrangement a , we transform one copy of the AAPA into an APA $\mathcal{A}_\psi(a)$ with n input directions and size $O(|\psi|_d^2 \cdot |\Sigma|^{n \cdot k})$. This is done using a variation of the procedure from [Theorem 3.11](#) which eliminates m input directions and substitutes the corresponding checks according to the arrangement a . When we substitute a path π_i with a path π_j , we have to make sure that all moves that were previously performed on direction j are now performed on direction i without manipulating the moves that were previously made on i . Thus, we introduce a second input marker for π_j on direction i in the $k \cdot n$ window from the proof of [Theorem 3.11](#). It is advanced whenever a symbol from direction j is read and does not affect the other markers on the same direction. These additional m markers do not asymptotically increase the size of the construction, thus we obtain the previously mentioned size. To perform the satisfiability test for the original formula, we now perform an emptiness test on the APA that makes a conjunctive move to $\mathcal{A}_\psi(a)$ for all arrangements a . Since there are n^m different arrangements, this APA has size $O(n^m \cdot |\psi|_d^2 \cdot |\Sigma|^{k \cdot n})$ and the test can be performed in EXPSPACE.

For a k -context-bounded analysis, we also have to incorporate the conjunctive test for all arrangements into the analysis. We first construct the structure $\mathcal{S}(g)$ for each guess g without

identifying any directions. We combine this structure with an arrangement a by replacing the test from Lemma 3.15 with $\bigwedge_{d'/d \in a} \bigwedge \{g' \in F_{d'}(\{q_0, g\}) \mid q_0 \in Q_0\}$ for all d and some Q_0 with $Q_0 \models \rho_0$. Here, we use $d'/d \in a$ to denote that in the arrangement a , d' is substituted by d . When using this notation, we assume that every arrangement contains the substitution d/d such that the original paths are still considered. This integration of a into $\mathcal{S}(g)$ to obtain $\mathcal{S}(a, g)$ does not increase its size beyond $O(g(k-2, |\psi|_d))$ asymptotically. When combining the tests for all arrangements $a_1, \dots, a_{n^m}$, we have to keep in mind that each test can use a different guess. Thus, we construct APA $\mathcal{S}(g_1, \dots, g_{n^m})$ that are parameterised in guesses $g_1, \dots, g_{n^m}$ and conjunctively move into $\mathcal{S}(a_i, g_i)$. Since they consist of n^m APA of size $O(g(k-2, |\psi|_d))$ and an initial state, they asymptotically have size $O(g(k-2, |\psi|_d))$ as well. We remove alternation from these parameterised APA to obtain NPA $\mathcal{S}'(g_1, \dots, g_{n^m})$ of size $O(g(k-1, |\psi|_d))$. The final NPA, which we test for emptiness to solve the satisfiability problem, nondeterministically guesses g_1 to g_{n^m} and moves to $\mathcal{S}'(g_1, \dots, g_{n^m})$. Since there are $|G| = O(g(k-1, |\psi|_d))$ possible guesses, we have $|G|^{n^m} = O(g(k-1, |\psi|_d))$ possible combinations of guesses. Thus, the final NPA also has an asymptotical size of $O(g(k-1, |\psi|_d))$. This yields an emptiness test in $(k-2)\text{EXPSpace}$.

For the lower bound, we need two different reductions. A reduction from $\exists^* \forall^*$ HyperLTL satisfiability yields the first and second lower bound. The third one is obtained from the fact that a k -context-bounded $\exists^*$ formula is especially an $\exists^* \forall^*$ formula. $\square$

8 CONCLUSION

In this paper, we introduced Alternating Asynchronous Parity Automata (AAPA) and the novel fixpoint logic H_μ as tools for the analysis of asynchronous hyperproperties. We showed the most interesting decision problems for both models to be highly undecidable in general, but exhibited families of increasingly precise under- and overapproximations for both AAPA and H_μ and presented asymptotically optimal algorithms for most corresponding decision problems. We also identified syntactic fragments where these analyses yield precise results.

Several questions remain for future work. Firstly, while we have established an equivalence between AAPA and H_μ formulas over fixed path assignments, an interesting question is whether there is a natural model of tree automata possibly extending AAPA and equivalent to the full logic with quantifiers, analogous to the correspondence between the modal μ -calculus and Alternating Parity Tree Automata [Emerson and Jutla 1991]. This could possibly lead to a more direct automata-theoretic approach to H_μ model checking. Secondly, it would be interesting to identify further approximate analyses and corresponding decidable fragments.

ACKNOWLEDGMENTS

This work was partially funded by DFG project MoNaLog (MU 1508/3). We thank the reviewers for their helpful comments and Roland Meyer and Sören van der Wall for valuable discussions. We also thank Laura Bozzelli for providing us with an extended version of [Bozzelli 2007].

REFERENCES

- Henrik Reif Andersen. 1994. *A polyadic modal μ -calculus*. Technical Report ID-TR: 1994-195. Dept. of Computer Science, Technical University of Denmark, Copenhagen.
- Mohamed Faouzi Atig, Ahmed Bouajjani, and Shaz Qadeer. 2009. Context-Bounded Analysis for Concurrent Programs with Dynamic Creation of Threads. In *Tools and Algorithms for the Construction and Analysis of Systems, 15th International Conference, TACAS 2009 (Lecture Notes in Computer Science, Vol. 5505)*, Stefan Kowalewski and Anna Philippou (Eds.). Springer, 107–123. https://doi.org/10.1007/978-3-642-00768-2_11
- Kshitij Bansal and Stéphane Demri. 2013. Model-Checking Bounded Multi-Pushdown Systems. In *Computer Science - Theory and Applications - 8th International Computer Science Symposium in Russia, CSR 2013, Ekaterinburg, Russia, June 25-29, 2013. Proceedings*. 405–417. https://doi.org/10.1007/978-3-642-38536-0_35

- Howard Barringer, Ruurd Kuiper, and Amir Pnueli. 1986. A Really Abstract Concurrent Model and its Temporal Logic. In *Conference Record of the Thirteenth Annual ACM Symposium on Principles of Programming Languages, St. Petersburg Beach, Florida, USA, January 1986*. ACM Press, 173–183. <https://doi.org/10.1145/512644.512660>
- Laura Bozzelli. 2007. Alternating Automata and a Temporal Fixpoint Calculus for Visibly Pushdown Languages. In *CONCUR 2007 – Concurrency Theory*, Luís Caires and Vasco T. Vasconcelos (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 476–491. https://doi.org/10.1007/978-3-540-74407-8_32
- Laura Bozzelli, Bastien Maubert, and Sophie Pinchinat. 2015. Unifying hyper and epistemic temporal logics. In *International Conference on Foundations of Software Science and Computation Structures*. Springer, 167–182. https://doi.org/10.1007/978-3-662-46678-0_11
- Michael R. Clarkson, Bernd Finkbeiner, Masoud Koleini, Kristopher K. Micinski, Markus N. Rabe, and César Sánchez. 2014. Temporal Logics for Hyperproperties. In *Principles of Security and Trust*, Martín Abadi and Steve Kremer (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 265–284. https://doi.org/10.1007/978-3-642-54792-8_15
- Michael R. Clarkson and Fred B. Schneider. 2010. Hyperproperties. *J. Comput. Secur.* 18, 6 (Sept. 2010), 1157–1210. <https://doi.org/10.3233/JCS-2009-0393>
- 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*. 1–13. <https://doi.org/10.1109/LICS.2019.8785713>
- Patrick Cousot and Radhia Cousot. 1979. Constructive versions of Tarski’s fixed point theorems. *Pacific J. Math.* 82, 1 (1979), 43–57. <https://projecteuclid.org/443/euclid.pjm/1102785059>
- Christian Dax and Felix Klaedtke. 2008. Alternation elimination by complementation. In *International Conference on Logic for Programming Artificial Intelligence and Reasoning*. Springer, 214–229. https://doi.org/10.1007/978-3-540-89439-1_16
- Stéphane Demri, Valentin Goranko, and Martin Lange. 2016. *Temporal Logics in Computer Science: Finite-State Systems*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139236119>
- Antoine Durand-Gasselin, Javier Esparza, Pierre Ganty, and Rupak Majumdar. 2015. Model Checking Parameterized Asynchronous Shared-Memory Systems. In *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 9206)*, Daniel Kroening and Corina S. Pasareanu (Eds.). Springer, 67–84. https://doi.org/10.1007/978-3-319-21690-4_5
- E. Allen Emerson and Charanjit S. Jutla. 1991. Tree Automata, Mu-Calculus and Determinacy (Extended Abstract). In *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*. 368–377. <https://doi.org/10.1109/SFCS.1991.185392>
- Javier Esparza, Pierre Ganty, and Rupak Majumdar. 2016. Parameterized Verification of Asynchronous Shared-Memory Systems. *J. ACM* 63, 1 (2016), 10:1–10:48. <https://doi.org/10.1145/2842603>
- Bernd Finkbeiner. 2017. Temporal Hyperproperties. *Bulletin of the EATCS* 123 (2017).
- Bernd Finkbeiner and Christopher Hahn. 2016. Deciding Hyperproperties. In *CONCUR 2016*. 13:1–13:14. <https://doi.org/10.4230/LIPIcs.CONCUR.2016.13>
- Bernd Finkbeiner, Christopher Hahn, Philip Lukert, Marvin Stenger, and Leander Tentrup. 2020. Synthesis from hyperproperties. *Acta Informatica* 57, 1-2 (2020), 137–163. <https://doi.org/10.1007/s00236-019-00358-2>
- Bernd Finkbeiner, Christopher Hahn, Marvin Stenger, and Leander Tentrup. 2019. Monitoring hyperproperties. *Formal Methods Syst. Des.* 54, 3 (2019), 336–363. <https://doi.org/10.1007/s10703-019-00334-z>
- Bernd Finkbeiner, Markus N. Rabe, and César Sánchez. 2015. Algorithms for Model Checking HyperLTL and HyperCTL*. In *CAV 2015*. 30–48. https://doi.org/10.1007/978-3-319-21690-4_3
- Olivier Finkel. 2006. On the Accepting Power of 2-Tape Büchi Automata. In *STACS 2006, 23rd Annual Symposium on Theoretical Aspects of Computer Science, Marseille, France, February 23-25, 2006, Proceedings*. 301–312. https://doi.org/10.1007/11672142_24
- Olivier Finkel. 2016. Infinite games specified by 2-tape automata. *Ann. Pure Appl. Logic* 167, 12 (2016), 1184–1212. <https://doi.org/10.1016/j.apal.2016.05.005>
- Olivier Finkel and Dominique Lecomte. 2009. Decision problems for Turing machines. *Inf. Process. Lett.* 109, 23-24 (2009), 1223–1226. <https://doi.org/10.1016/j.ipl.2009.09.002>
- Carlo A. Furia. 2014. Asynchronous Multi-Tape Automata Intersection: Undecidability and Approximation. *CoRR* abs/1206.4860 (2014). arXiv:1206.4860v5 <http://arxiv.org/abs/1206.4860v5>
- Pierre Ganty and Rupak Majumdar. 2012. Algorithmic verification of asynchronous programs. *ACM Trans. Program. Lang. Syst.* 34, 1 (2012), 6:1–6:48. <https://doi.org/10.1145/2160910.2160915>
- Pierre Ganty, Rupak Majumdar, and Andrey Rybalchenko. 2009. Verifying liveness for asynchronous programs. In *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, Zhong Shao and Benjamin C. Pierce (Eds.). ACM, 102–113. <https://doi.org/10.1145/1480881.1480895>
- Dainis Geidmanis. 1987. On the Capabilities of Alternating and Nondeterministic Multitape Automata. In *Fundamentals of Computation Theory, International Conference FCT’87, Kazan, USSR, June 22-26, 1987, Proceedings*. 150–154. https://doi.org/10.1007/978-3-642-54792-8_15

[//doi.org/10.1007/3-540-18740-5_35](https://doi.org/10.1007/3-540-18740-5_35)

- Jens Oliver Gutsfeld, Markus Müller-Olm, and Christoph Ohrem. 2020a. Propositional Dynamic Logic for Hyperproperties. In *31st International Conference on Concurrency Theory, CONCUR 2020, September 1-4, 2020, Vienna, Austria (Virtual Conference) (LIPIcs, Vol. 171)*, Igor Konnov and Laura Kovács (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 50:1–50:22. <https://doi.org/10.4230/LIPIcs.CONCUR.2020.50>
- Jens Oliver Gutsfeld, Markus Müller-Olm, and Christoph Ohrem. 2020b. Automata and Fixpoints for Asynchronous Hyperproperties. arXiv:2010.11605 [cs.LO]
- Oscar H. Ibarra and Nicholas Q. Trn. 2013. How to synchronize the Heads of a Multitape Automaton. *Int. J. Found. Comput. Sci.* 24, 6 (2013), 799–814. <https://doi.org/10.1142/S0129054113400194>
- Andreas Krebs, Arne Meier, Jonni Virtema, and Martin Zimmermann. 2018. Team Semantics for the Specification and Verification of Hyperproperties. In *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018, August 27-31, 2018, Liverpool, UK (LIPIcs, Vol. 117)*, Igor Potapov, Paul G. Spirakis, and James Worrell (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 10:1–10:16. <https://doi.org/10.4230/LIPIcs.MFCS.2018.10>
- Martin Lange. 2005. Weak Automata for the Linear Time μ -Calculus. In *Verification, Model Checking, and Abstract Interpretation*, Radhia Cousot (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 267–281. https://doi.org/10.1007/978-3-540-30579-8_18
- Martin Lange. 2015. The Arity Hierarchy in the Polyadic μ -Calculus. In *FICS (EPTCS, Vol. 191)*. 105–116. <https://doi.org/10.4204/EPTCS.191.10>
- Isabella Mastroeni and Michele Pasqua. 2017. Hyperhierarchy of Semantics - A Formal Framework for Hyperproperties Verification. In *Static Analysis - 24th International Symposium, SAS 2017, New York, NY, USA, August 30 - September 1, 2017, Proceedings (Lecture Notes in Computer Science, Vol. 10422)*, Francesco Ranzato (Ed.). Springer, 232–252. https://doi.org/10.1007/978-3-319-66706-5_12
- Isabella Mastroeni and Michele Pasqua. 2018. Verifying Bounded Subset-Closed Hyperproperties. In *Static Analysis - 25th International Symposium, SAS 2018, Freiburg, Germany, August 29-31, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 11002)*, Andreas Podelski (Ed.). Springer, 263–283. https://doi.org/10.1007/978-3-319-99725-4_17
- Dimiter Milushev and Dave Clarke. 2013. Incremental Hyperproperty Model Checking via Games. In *Proceedings of the 18th Nordic Conference on Secure IT Systems - Volume 8208 (Ilulissat, Greenland) (NordSec 2013)*. Springer-Verlag New York, Inc., New York, NY, USA, 247–262. https://doi.org/10.1007/978-3-642-41488-6_17
- Anca Muscholl. 1996. On the Complementation of Asynchronous Cellular Bchi Automata. *Theor. Comput. Sci.* 169, 2 (1996), 123–145. [https://doi.org/10.1016/S0304-3975\(96\)00117-X](https://doi.org/10.1016/S0304-3975(96)00117-X)
- Martin Otto. 1999. Bisimulation-invariant PTIME and higher-dimensional μ -calculus. *Theor. Comput. Sci.* 224, 1-2 (1999), 237–265. [https://doi.org/10.1016/S0304-3975\(98\)00314-4](https://doi.org/10.1016/S0304-3975(98)00314-4)
- D. Peled and W. Penczek. 1996. Using Asynchronous Bchi Automata for Efficient Automatic Verification of Concurrent Systems. In *Protocol Specification, Testing and Verification XV: Proceedings of the Fifteenth IFIP WG6.1 International Symposium on Protocol Specification, Testing and Verification, Warsaw, Poland, June 1995*. Springer US, Boston, MA, 315–330. https://doi.org/10.1007/978-0-387-34892-6_20
- Shaz Qadeer. 2008. The Case for Context-Bounded Verification of Concurrent Programs. In *Model Checking Software, 15th International SPIN Workshop, Los Angeles, CA, USA, August 10-12, 2008, Proceedings (Lecture Notes in Computer Science, Vol. 5156)*, Klaus Havelund, Rupak Majumdar, and Jens Palsberg (Eds.). Springer, 3–6. https://doi.org/10.1007/978-3-540-85114-1_2
- Shaz Qadeer and Jakob Rehof. 2005. Context-Bounded Model Checking of Concurrent Software. In *Tools and Algorithms for the Construction and Analysis of Systems, 11th International Conference, TACAS 2005 (Lecture Notes in Computer Science, Vol. 3440)*, Nicolas Halbwachs and Lenore D. Zuck (Eds.). Springer, 93–107. https://doi.org/10.1007/978-3-540-31980-1_7
- Markus N. Rabe. 2016. *A temporal logic approach to Information-flow control*. Ph.D. Dissertation. Saarland University.
- Michael O. Rabin and Dana S. Scott. 1959. Finite Automata and Their Decision Problems. *IBM Journal of Research and Development* 3, 2 (1959), 114–125. <https://doi.org/10.1147/rd.32.0114>
- Aravinda Prasad Sistla. 1983. *Theoretical Issues in the Design and Verification of Distributed Systems*. Ph.D. Dissertation. Carnegie-Mellon University, USA.
- Alex Spelten, Wolfgang Thomas, and Sarah Winter. 2011. Trees over Infinite Structures and Path Logics with Synchronization. In *Proceedings 13th International Workshop on Verification of Infinite-State Systems, INFINITY 2011, Taipei, Taiwan, 10th October 2011 (EPTCS, Vol. 73)*, Fang Yu and Chao Wang (Eds.). 20–34. <https://doi.org/10.4204/EPTCS.73.5>
- Larry Joseph Stockmeyer. 1974. *The complexity of decision problems in automata theory and logic*. Ph.D. Dissertation. MIT.
- Alfred Tarski. 1955. A lattice-theoretical fixpoint theorem and its applications. *Pacific J. Math.* 5, 2 (1955), 285–309. <https://projecteuclid.org/443/euclid.pjm/1103044538>
- Moshe Y. Vardi. 1988. A Temporal Fixpoint Calculus. In *POPL*. ACM Press, 250–259. <https://doi.org/10.1145/73560.735822>
- Wieslaw Zielonka. 1987. Notes on Finite Asynchronous Automata. *ITA* 21, 2 (1987), 99–135.