



Tracy, traces, and transducers: computable counterexamples and explanations for HyperLTL model-checking

Sarah Winter¹ · Martin Zimmermann²

Received: 26 November 2024 / Accepted: 16 July 2025
© The Author(s) 2025

Abstract

HyperLTL model-checking enables the automated verification of information-flow properties for security-critical systems. However, it only provides a binary answer. Here, we consider the problem of computing counterexamples and explanations for HyperLTL model-checking, thereby considerably increasing its usefulness. Based on the maxim “counterexamples/explanations are Skolem functions for the existentially quantified trace variables”, we consider (Turing machine) computable Skolem functions. As not every finite transition system and formula have computable Skolem functions witnessing that the system satisfies the formula, we consider the problem of deciding whether such functions exist. Our main result shows that this problem is decidable by reducing it to solving multiplayer games with hierarchical imperfect information. Furthermore, our algorithm also computes transducers implementing such functions, if they exist.

1 Introduction

Prologue. Tracy sits in her office and needs to print her latest travel reimbursement claim. After hitting the print button, she walks to the printer room only to find out that the document has not been printed. So, she walks back to her office, hits the print button again, walks to the printer and is slightly surprised to find her document. Sometimes Tracy wonders whether the print system is nondeterministic. If only there was a way to find out.

Information-flow properties, which are crucial in the specification of security-critical systems, require the simultaneous reasoning about multiple executions of a system. However, most classical specification languages like LTL and CTL* refer to a single execution trace at a time. Clarkson and Schneider [11] coined the term *hyperproperties* for properties

✉ Sarah Winter
winter@irif.fr

✉ Martin Zimmermann
mzi@cs.aau.dk

¹ IRIF, Université Paris Cité, Paris, France

² Aalborg University, Aalborg, Denmark

that require the reasoning about multiple traces. Just like ordinary trace and branching-time properties, hyperproperties can be specified using temporal logics, e.g., HyperLTL and HyperCTL* [8], expressive, but intuitive specification languages that are able to express typical information-flow properties such as noninterference, noninference, declassification, and input determinism. Due to their practical relevance and theoretical elegance, hyperproperties and their specification languages have received considerable attention during the last decade.

HyperLTL is obtained by extending LTL [29], the most influential specification language for linear-time properties, by trace quantifiers to refer to multiple executions of a system. Hence, a HyperLTL formula is indeed evaluated over a set of traces, which forms the universe for the quantifiers. For example, the HyperLTL formula $\varphi_{\text{id}} = \forall \pi, \pi'. G(i_\pi \leftrightarrow i_{\pi'}) \rightarrow G(o_\pi \leftrightarrow o_{\pi'})$ expresses input determinism, i.e., every pair of traces that always has the same input (represented by the proposition i) also always has the same output (represented by the proposition o). Having learned about HyperLTL, Tracy wonders whether she can formally prove that the print system violates φ_{id} .

In this work, we focus on the model-checking problem for HyperLTL, which intuitively asks whether a given (finite model of a) system satisfies a given HyperLTL specification. This problem is decidable, albeit TOWER-complete [27, 30].

But the model-checking problem as described above is “just” a decision problem, i.e., the user only learns whether the system satisfies the specification or not, but not the *reason* it does or does not. It has been argued that this binary answer is in general not useful [25]: Most real-life systems are too complex to be modelled faithfully by a finite transition system. Hence, one always checks an abstraction, not the actual system. Then, a positive answer to the model-checking problem does not show that the actual system is correct, bugs in it might have been abstracted away when constructing a finite transition system modelling it. The actual killer application of model-checking is the automated generation of counterexamples in case the specification is not satisfied by the abstraction. Given a counterexample in the abstraction one can then check whether this (erroneous) behaviour also exists in the actual system, or whether it was introduced during the abstraction [9]. In the latter case, the abstraction has to be refined and checked again. But if the erroneous behaviour can be found in the actual system, then this bug can be fixed in the actual system.

But what is a counterexample in HyperLTL model-checking? For the formula φ_{id} expressing input determinism this is straightforward: if a transition system does not satisfy the formula, then it has two traces that coincide on their input, but not on their output. However, the situation becomes more interesting in the presence of existentially quantified variables and quantifier alternations. Consider, for example, a formula of the form $\varphi = \exists \pi \forall \pi'. \psi$ with quantifier-free ψ and let \mathfrak{T} be a transition system with set $\text{Tr}(\mathfrak{T})$ of traces. If $\mathfrak{T} \not\models \varphi$, then for every choice of $t \in \text{Tr}(\mathfrak{T})$ there is a $t' \in \text{Tr}(\mathfrak{T})$ such that the variable assignment $\{\pi \mapsto t, \pi' \mapsto t'\}$ does not satisfy ψ . Thus, a counterexample is described by a Skolem function $f: \text{Tr}(\mathfrak{T}) \rightarrow \text{Tr}(\mathfrak{T})$ for the existentially quantified variable π' in the negation $\forall \pi \exists \pi'. \neg \psi$ of φ . It gives, for every choice t for the existentially quantified π in φ a trace $f(t)$ for the universally quantified π' in φ such that $\{\pi \mapsto t, \pi' \mapsto f(t)\} \models \neg \psi$, i.e., $\{\pi \mapsto t, \pi' \mapsto f(t)\} \not\models \psi$, thereby explaining for every choice of t why it is not a good one. The maxim “counterexamples are Skolem functions for existentially quantified variables in the negation of the specification” is true for arbitrary formulas. But before we explore this approach further, let us first consider a second application.

Explainability, the need to explain to, e.g., users, customers, and regulators, what a system does, is an aspect of system design that gains more and more significance. This is in particular true when it comes to systems designed by algorithms, e.g., machine-learning or synthesis. For any nontrivial system of this kind, it is impossible for humans to develop an explanation of their behaviour or a witness for their correctness. This is a major obstacle preventing the wide-spread use of (unexplained) machine-generated software in safety-critical applications [3]. Also here, HyperLTL model-checking can be useful: Assuming the system is supposed to satisfy a HyperLTL specification and indeed does so, then Skolem functions “explain” why the specification is satisfied.

Our Contributions. In this work, we are interested in computing counterexamples/explanations for HyperLTL, which, as argued above, boils down to computing Skolem functions for HyperLTL. Before we explain our contributions, let us remark that counterexamples are just explanations for the negation of the specification, as we have seen above. Hence, in the following we will focus on explanations, as this setting spares us from dealing with a negation. Also, let us remark that for every transition system \mathcal{T} and every HyperLTL formula φ , we either have $\mathcal{T} \models \varphi$ or $\mathcal{T} \models \neg\varphi$. Hence, our framework will either explain why \mathcal{T} satisfies φ or explain why \mathcal{T} satisfies $\neg\varphi$, i.e., explain why \mathcal{T} does not satisfy φ .

In general, we are given a transition system \mathcal{T} and a HyperLTL formula φ such that $\mathcal{T} \models \varphi$, and we want to compute Skolem functions for the existentially quantified variables in φ . Note that the actual explanation-phase employing the Skolem functions is an interactive process between the user (i.e., Tracy) and these functions: Tracy has to specify choices for the universally quantified variables, which are then fed into the Skolem functions, yielding choices for the existentially quantified variables such that the combination of all of these traces satisfies the quantifier-free part of the specification.

To *apply* Skolem functions in that manner, they need to be finitely representable. To be as general as possible, we consider here functions that are computable by Turing machines (in a very natural sense). But even for such a general model, $\mathcal{T} \models \varphi$ may not have a computable explanation. The underlying reason is that such Turing machines can only compute continuous functions: Intuitively, if two inputs coincide on a “long” prefix, then the corresponding outputs also coincide on a “long” prefix. However, it is straightforward to construct a pair $\mathcal{T} \models \varphi$ that does not have continuous Skolem functions (see Theorem 1). Hence, our main focus is on the following question: given \mathcal{T} and φ with $\mathcal{T} \models \varphi$, is $\mathcal{T} \models \varphi$ witnessed by computable Skolem functions? Our main result shows that this problem is decidable. To prove it, we combine techniques developed in the theory of uniformization [15], delay games [24], and multiplayer games with hierarchical imperfect information [5] to express the existence of computable Skolem functions by a multi-player game with hierarchical imperfect information. Intuitively, there is one player for each existentially quantified variable and they form a coalition against a player corresponding to the universally quantified variables. Hierarchical imperfect information then captures the structure of the quantifier prefix, e.g., the Skolem function for π_1 in a formula of the form $\forall\pi_0\exists\pi_1\forall\pi_2\exists\pi_3.\psi$ depends only on π_0 while the one for π_3 depends on π_0 and π_2 (as usual one can assume that Skolem functions only depend on universally quantified variables). Furthermore, delay games are a general approach to deciding the existence of continuous functions in synthesis and uniformization [15, 20, 22, 24].

As a byproduct of our game-theoretic characterization, we show that if $\mathcal{T} \models \varphi$ has (Turing machine) computable Skolem functions, then it also has ones that are computed by

word-to-word (one-way) transducers with bounded delay between input and output, a much more modest machine model. In fact, our algorithm computes transducers implementing Skolem functions whenever computable Skolem functions exist. This allows for the effective computation and simulation of computable Skolem functions as described above.

2 Preliminaries

¹ We denote the set of nonnegative integers by \mathbb{N} . The domain of a partial function $f: A \rightarrow B$ is denoted by $\text{dom}(f) = \{a \in A \mid f(a) \text{ is defined}\}$. More generally, we denote the domain $\{a \in A \mid (a, b) \in R \text{ for some } b \in B\}$ of a relation $R \subseteq A \times B$ by $\text{dom}(R)$.

2.1 Languages, Transition Systems, and Automata

An alphabet is a nonempty finite set. The sets of finite and infinite words over an alphabet Σ are denoted by Σ^* and Σ^ω , respectively. The length of a finite word w is denoted by $|w|$. Given n infinite words w_0, \dots, w_{n-1} , let their merge (also known as zip), which is an infinite word over Σ^n , be defined as

$$\begin{aligned} \text{mrg}(w_0, \dots, w_{n-1}) \\ = (w_0(0), \dots, w_{n-1}(0))(w_0(1), \dots, w_{n-1}(1))(w_0(2), \dots, w_{n-1}(2)) \cdots \end{aligned}$$

We define $\text{mrg}(w_0, \dots, w_{n-1})$ for finite words w_0, \dots, w_{n-1} of the same length analogously.

The set of prefixes of an infinite word $w = w(0)w(1)w(2) \cdots \in \Sigma^\omega$ is $\text{Prfs}(w) = \{w(0) \cdots w(i-1) \mid i \geq 0\}$, which is lifted to languages $L \subseteq \Sigma^\omega$ via $\text{Prfs}(L) = \bigcup_{w \in L} \text{Prfs}(w)$. A language $L \subseteq \Sigma^\omega$ is closed if $\{w \in \Sigma^\omega \mid \text{Prfs}(w) \subseteq \text{Prfs}(L)\} \subseteq L$.

Throughout this paper, we fix a finite set AP of atomic propositions. A transition system $\mathfrak{T} = (V, E, v_I, \lambda)$ consists of a finite set V of vertices, a set $E \subseteq V \times V$ of (directed) edges, an initial vertex $v_I \in V$, and a labelling $\lambda: V \rightarrow 2^{\text{AP}}$ of the vertices by sets of atomic propositions. We assume that every vertex has at least one outgoing edge. A path ρ through \mathfrak{T} is an infinite sequence $\rho = v_0 v_1 v_2 \cdots$ of vertices with $v_0 = v_I$ and $(v_n, v_{n+1}) \in E$ for every $n \geq 0$. The trace of ρ is defined as $\lambda(\rho) = \lambda(v_0)\lambda(v_1)\lambda(v_2) \cdots \in (2^{\text{AP}})^\omega$. The set of traces of \mathfrak{T} is $\text{Tr}(\mathfrak{T}) = \{\lambda(\rho) \mid \rho \text{ is a path of } \mathfrak{T}\}$.

Remark 1 The following facts follow directly from the definition of closed languages.

1. Let \mathfrak{T} be a transition system. Then, $\text{Tr}(\mathfrak{T})$ is closed.
2. If $L_0, \dots, L_{n-1} \subseteq \Sigma^\omega$ are closed, then so is $\{\text{mrg}(w_0, \dots, w_{n-1}) \mid w_i \in L_i \text{ for } 0 \leq i < n\}$.

A Büchi automaton $\mathcal{A} = (Q, \Sigma, q_I, \delta, F)$ consists of a finite set Q of states containing the initial state $q_I \in Q$ and the subset $F \subseteq Q$ of accepting states, an alphabet Σ , and a

¹ We like the twentieth letter of the alphabet. In fact, we like it so much that we use ι to denote traces, T to denote sets of traces, \mathfrak{T} to denote transition systems, and \mathcal{T} to denote transducers. We hope this footnote will help the reader keep track of them.

transition function $\delta: Q \times \Sigma \rightarrow 2^Q$. Let $w = w(0)w(1)w(2) \cdots \in \Sigma^\omega$. A run of \mathcal{A} on w is a sequence $q_0q_1q_2 \cdots$ with $q_0 = q_I$ and $q_{n+1} \in \delta(q_n, w(n))$ for all $n \geq 0$. A run $q_0q_1q_2 \cdots$ is (Büchi) accepting if there are infinitely many $n \in \mathbb{N}$ with $q_n \in F$. The language (Büchi) recognized by \mathcal{A} , denoted by $L(\mathcal{A})$, is the set of infinite words over Σ that have an accepting run of \mathcal{A} on w that is accepting.

Büchi automata recognize exactly the ω -regular languages, but require nondeterminism to do so. In the following, it is sometimes prudent to work with deterministic ω -automata, which also accept exactly the ω -regular languages (see, e.g., [18], for more details). A deterministic parity automaton $\mathcal{P} = (Q, \Sigma, q_I, \delta, \Omega)$ consists of a finite set Q of states containing the initial state $q_I \in Q$, an alphabet Σ , a transition function $\delta: Q \times \Sigma \rightarrow Q$, and a coloring $\Omega: Q \rightarrow \mathbb{N}$ of the states by colors in \mathbb{N} . Let $w = w(0)w(1)w(2) \cdots \in \Sigma^\omega$. Then, \mathcal{P} has a unique run $q_0q_1q_2 \cdots$ on w , defined as $q_0 = q_I$ and $q_{n+1} = \delta(q_n, w(n))$ for all $n \geq 0$. A run $q_0q_1q_2 \cdots$ is (parity) accepting if the maximal color appearing infinitely often in the sequence $\Omega(q_0)\Omega(q_1)\Omega(q_2) \cdots$ is even. The language (parity) recognized by \mathcal{A} , denoted by $L(\mathcal{A})$, is the set of infinite words over Σ such that the run of \mathcal{A} on w is accepting.

2.2 HyperLTL

The formulas of HyperLTL are given by the grammar

$$\varphi ::= \exists \pi. \varphi \mid \forall \pi. \varphi \mid \psi \quad \psi ::= a_\pi \mid \neg \psi \mid \psi \vee \psi \mid X \psi \mid \psi U \psi$$

where a ranges over AP and where π ranges over a fixed countable set \mathcal{V} of (trace) variables. Conjunction (\wedge), exclusive disjunction (\oplus), implication (\rightarrow), and equivalence (\leftrightarrow) are defined as usual, and the temporal operators “eventually” (F) and “always” (G) are derived as $F\psi = \neg \psi U \psi$ and $G\psi = \neg F \neg \psi$. A sentence is a formula without free variables, which are defined as expected.

The semantics of HyperLTL is defined with respect to a trace assignment, a partial mapping $\Pi: \mathcal{V} \rightarrow (2^{\text{AP}})^\omega$. The assignment with empty domain is denoted by Π_\emptyset . Given a trace assignment Π , a variable π , and a trace t we denote by $\Pi[\pi \rightarrow t]$ the assignment that coincides with Π everywhere but at π , which is mapped to t . Furthermore, $\Pi[j, \infty)$ denotes the trace assignment mapping every π in Π ’s domain to $\Pi(\pi)(j)\Pi(\pi)(j+1)\Pi(\pi)(j+2) \cdots$, the suffix of $\Pi(\pi)$ starting at position j .

For sets T of traces and trace assignments Π we define

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

We say that T satisfies a sentence φ if $(T, \Pi_\emptyset) \models \varphi$. In this case, we write $T \models \varphi$ and say that T is a model of φ . A transition system \mathfrak{T} satisfies φ , written $\mathfrak{T} \models \varphi$, if $\text{Tr}(\mathfrak{T}) \models \varphi$.

Although HyperLTL sentences are required to be in prenex normal form, they are closed under Boolean combinations, which can be easily seen by transforming such a formula into an equivalent formula in prenex normal form. In particular, the negation $\neg\varphi$ of a sentence φ satisfies $T \models \neg\varphi$ iff $T \not\models \varphi$. Also, note that the statement $(T, \Pi) \models \psi$ for quantifier-free formulas ψ is independent of T . Hence, we often just write $\Pi \models \psi$ for the sake of readability.

Remark 2 Let ψ be a quantifier-free HyperLTL formula (with k free variables π_0, \dots, π_{k-1}) and let \mathfrak{T} be a transition system. There is an (effectively computable) Büchi automaton $\mathcal{A}_{\psi}^{\mathfrak{T}}$ such that $L(\mathcal{A}_{\psi}^{\mathfrak{T}})$ is equal to

$$\{\text{mrg}(\Pi(\pi_0), \dots, \Pi(\pi_{k-1})) \mid \Pi(\pi_i) \in \text{Tr}(\mathfrak{T}) \text{ for } 0 \leq i \leq k-1 \text{ and } (\text{Tr}(\mathfrak{T}), \Pi) \models \psi\}.$$

It can be obtained by noting that ψ is almost an LTL formula as it is quantifier-free, but its atomic propositions are still labeled by trace variables, i.e., they are all of the form a_{π_i} for some $i \in \{0, \dots, k-1\}$. Let ψ' be the (proper) LTL formula obtained from ψ by replacing each atomic proposition a_{π_j} by the atomic proposition (a, j) , i.e., ψ' is defined over the set $\text{AP} \times \{0, \dots, k-1\}$ of atomic propositions. For ψ' there exists a Büchi automaton $\mathcal{A}_{\psi'}$ with $|\psi| \cdot 2^{|\psi|}$ many states recognizing the language

$$L(\mathcal{A}_{\psi'}) = \{t \in (2^{\text{AP} \times \{0, \dots, k-1\}})^{\omega} \mid t \models \psi'\}$$

(see, e.g., [4]). By replacing each letter $A \in 2^{\text{AP} \times \{0, \dots, k-1\}}$ on a transition of $\mathcal{A}_{\psi'}$ by the letter

$$(\{a \in \text{AP} \mid (a, 0) \in A\}, \dots, \{a \in \text{AP} \mid (a, k-1) \in A\}) \in (2^{\text{AP}})^k$$

we obtain a Büchi automaton \mathcal{A}_{ψ} recognizing

$$\{\text{mrg}(\Pi(\pi_0), \dots, \Pi(\pi_{k-1})) \mid \Pi(\pi_j) \in (2^{\text{AP}})^{\omega} \text{ for all } 0 \leq j \leq k-1 \text{ and } (\text{Tr}(\mathfrak{T}), \Pi) \models \psi\}.$$

Finally, by taking the product of \mathcal{A}_{ψ} with k copies of \mathfrak{T} (where the i -th one restricts the i -th trace of $\text{mrg}(\Pi(\pi_0), \dots, \Pi(\pi_{k-1}))$ to traces of \mathfrak{T}), we obtain the desired Büchi automaton $\mathcal{A}_{\psi}^{\mathfrak{T}}$ recognizing

$$\{\text{mrg}(\Pi(\pi_0), \dots, \Pi(\pi_{k-1})) \mid \Pi(\pi_i) \in \text{Tr}(\mathfrak{T}) \text{ for } 0 \leq i \leq k-1 \text{ and } (\text{Tr}(\mathfrak{T}), \Pi) \models \psi\}.$$

Finally, it has $|\psi| \cdot 2^{|\psi|} \cdot |\mathfrak{T}|^k$ many states.

2.3 Skolem functions for HyperLTL

Let $\varphi = Q_0\pi_0 \cdots Q_{k-1}\pi_{k-1} \cdot \psi$ be a HyperLTL sentence such that ψ is quantifier-free and let T be a set of traces. Moreover, let $i \in \{0, 1, \dots, k-1\}$ be such that $Q_i = \exists$ and let $U_i = \{j < i \mid Q_j = \forall\}$ be the indices of the universal quantifiers preceding Q_i . Furthermore, let $f_i: T^{|U_i|} \rightarrow T$ for each such i (note that f_i is a constant, if U_i is empty). We say that a variable assignment Π with $\text{dom}(\Pi) \supseteq \{\pi_0, \pi_1, \dots, \pi_{k-1}\}$ is consistent with the f_i if $\Pi(\pi_j) \in T$ for all j with $Q_j = \forall$ and $\Pi(\pi_j) = f_j(\Pi(\pi_{i_0}), \Pi(\pi_{i_1}), \dots, \Pi(\pi_{i_{|U_j|-1}}))$ for all

j with $Q_j = \exists$, where $U_j = \{i_0 < i_1 < \dots < i_{|U_j|-1}\}$. If $\Pi \models \psi$ for each Π that is consistent with the f_i , then we say that the f_i are Skolem functions witnessing $T \models \varphi$.

Remark 3 $T \models \varphi$ iff there are Skolem functions for the existentially quantified variables of φ that witness $T \models \varphi$.

Note that only traces for universal variables are inputs for Skolem functions, but not those for existentially quantified variables. As usual, this is not a restriction, as the inputs of a Skolem function for an existentially quantified variable π_i is a superset of the inputs of a Skolem function for another existentially quantified variable π_j with $j < i$.

Example 1 Let $\varphi = \forall \pi \exists \pi' \exists \pi'' . G(a_\pi \leftrightarrow (a_{\pi'} \oplus a_{\pi''}))$. We have $(2^{\{a\}})^\omega \models \varphi$. Now, for every function $f': (2^{\{a\}})^\omega \rightarrow (2^{\{a\}})^\omega$, there is a function $f'': (2^{\{a\}})^\omega \rightarrow (2^{\{a\}})^\omega$ such that f', f'' are Skolem functions witnessing $(2^{\{a\}})^\omega \models \varphi$, i.e., we need to define f'' such that $(f''(t))(n) = (f'(t))(n)$ for all $n \in \mathbb{N}$ such that $t(n) = \emptyset$ and $(f''(t))(n) = \overline{(f'(t))(n)}$ for all $n \in \mathbb{N}$ such that $t(n) = \{a\}$, where $\overline{\{a\}} = \emptyset$ and $\overline{\emptyset} = \{a\}$. Hence, f'' depends on f' , but the value of $f'(t)$ (for the existentially quantified π_1) does not need to be an input to f'' , it can be determined from the input t for the universally quantified π . This is not surprising, but needs to be taken into account in our constructions.

3 Problem statement

Our goal is to construct (Turing machine) computable Skolem functions that serve as algorithmic explanations for the satisfaction of a HyperLTL property. At first glance, this is a very ambitious goal, as it requires working with Turing machines processing infinite inputs and producing infinite outputs. To tackle this issue, we re-formulate our problem as a (rather non-standard, more synthesis-like) uniformization problem.

Here, we consider the variant where one is given a relation $R \subseteq A \times B$ and the goal is to determine whether there is a function uniformizing R (i.e., a partial function $f: A \rightarrow B$ such that $\text{dom}(f) = \text{dom}(R)$ and $\{(a, f(a)) \mid a \in \text{dom}(R)\} \subseteq R$) that is computed by a machine from some fixed class of machines. For the case where A is the Cartesian product of the set of traces of a transition system and B is the set of traces of the same transition system, we have captured the problem of computing a Skolem function as a uniformization problem.

Thus, for a sentence with quantifier prefix $\forall^* \exists^*$ we can obtain computable Skolem functions by interpreting the problem as a uniformization problem. However, for more complex quantifier prefixes, this is no longer straightforward, as the dependencies between the variables have to be considered, e.g., the Skolem function of an existentially quantified variable only has inputs corresponding to outermore universally quantified variables, but may also depend on outermore existentially quantified variables, i.e., outputs are also (implicit) inputs for other functions.

Another issue is that Turing machines are a very expressive model of computation. Filiot and Winter [15] studied synthesis of computable functions from rational specifications (e.g., specifications recognized by a Büchi automaton): they proved that uniformization by Turing machines coincides with uniformization by transducers with bounded delay (if the domain of the specification is closed), a much nicer class of machines computing functions from

infinite words to infinite words. Crucially, the functions computed by such transducers are also continuous in the Cantor topology over infinite words (we refer to [15] for definitions and details). In the setting of Skolem functions for HyperLTL model-checking for $\forall^*\exists^*$ -sentences, that means that if two inputs agree on a “long” prefix, then the corresponding outputs also agree on a “long” prefix. Continuity is a desirable property when using Skolem functions on-the-fly: If Tracy has fixed a long prefix of the inputs for the Skolem functions, then future inputs do not change the output prefixes already produced by the Skolem functions for the fixed prefix.

However, we will show that there is a HyperLTL sentence that does not have continuous Skolem functions (and thus also no computable ones). Hence, it is natural to ask if it is decidable whether a given pair (\mathfrak{T}, φ) has a computable explanation. We prove that this is indeed the case, even for sentences with arbitrary quantifier prefixes.

3.1 Uniformization by computable functions

In the following, ω -regular for languages over the alphabet $\Sigma \times \Gamma$, i.e., $L \subseteq (\Sigma \times \Gamma)^\omega$, we speak about its induced relation $R_L = \{(x, y) \in \Sigma^\omega \times \Gamma^\omega \mid \text{mrg}(x, y) \in L\}$. For the sake of readability, we often do not distinguish between languages $L \subseteq (\Sigma \times \Gamma)^\omega$ and induced relations $R_L \subseteq \Sigma^\omega \times \Gamma^\omega$ and call R_L ω -regular if L is ω -regular. A function $f: \Sigma^\omega \rightarrow \Gamma^\omega$ uniformizes a relation $R \subseteq \Sigma^\omega \times \Gamma^\omega$ if the domain of f is equal to the domain of R and the graph $\{(w, f(w)) \mid w \in \text{dom}(f)\}$ of f is a subset of R . In our setting, the uniformization problem asks whether, for a given relation, there is a computable function that uniformizes it.

To define the computability of a function from Σ^ω to Γ^ω , we consider deterministic three-tape Turing machines \mathcal{M} with the following setup (following [15]): the first tape is a read-only, one-way tape and contains the input in Σ^ω , the second one is a two-way working tape, and the third one is a write-only, one-way tape on which the output in Γ^ω is generated. Formally, we say that \mathcal{M} computes the partial function $f: \Sigma^\omega \rightarrow \Gamma^\omega$ if, when started with input $w \in \text{dom}(f)$ on the first tape, \mathcal{M} produces (in the limit) the output $f(w)$ on the third tape. Note that we do not require the Turing machine to check whether its input is in the domain of f . We just require it to compute the correct output for those inputs in the domain, it may behave arbitrarily on inputs outside of the domain of f . This is done so that the uniformization function only has to capture the complexity of transforming possible inputs into outputs, but does not have to capture the complexity of checking whether an input is in the domain. In our setting, this can be taken care of by deterministic ω -automata that can be effectively computed.

We say that such an \mathcal{M} has bounded delay, if there is a $d \in \mathbb{N}$ such that to compute the first n letters of the output only $n + d$ letters of the input are read (i.e., the other cells of the input tape are not visited before n output letters have been generated).

Lemma 1 *Let f be computable and let $\text{dom}(f)$ be closed. Then f is computable by a Turing machine with bounded delay.*

Proof This follows from the fact that the set Σ^ω is a compact space when equipped with the Cantor distance. A closed subset of a compact space is compact (see, e.g., [2]). Hence, $\text{dom}(f)$ is a compact space. Further, every computable function is continuous (see, e.g., [15]). Now, the Heine-Cantor theorem states that every continuous function between metric

spaces $f:M \rightarrow N$ where M is a compact space is in fact uniformly continuous. Thus, f lies in the intersection of the classes of computable and uniformly continuous functions which implies that f can also be computed with bounded delay:

Intuitively, f being uniformly continuous means there exists some k such that in order to get i output symbols it suffices to consider $i + k$ input symbols. Assume a Turing machine \mathcal{M} computes f . We briefly sketch how to obtain a Turing machine \mathcal{M}' from \mathcal{M} that computes f and has delay at most k .

As long as \mathcal{M} maintains that the i -th output symbol is produced before the $(i + k + 1)$ -th input symbol is read, \mathcal{M}' behaves like \mathcal{M} . However, as soon as \mathcal{M} would violate this, \mathcal{M}' continues to simulate \mathcal{M} on a fixed valid continuation of the input word (regardless of how the actual input word is continued). We refer to this as dummy continuation. We note here that our input words are traces of some transition system. Hence, an input word is a trace and valid continuations of some trace prefix can easily be generated from the transition system. After \mathcal{M} has produced another output symbol (while reading the dummy continuation), \mathcal{M}' produces the same output symbol. The choice of the dummy continuation has no relevance for the produced output symbol, as f is uniformly continuous: Towards a contradiction, assume the output symbol changes for different dummy inputs. That means there exist α and α' in the domain of f that agree on a prefix of length $i + k$, but the i -th symbol of $f(\alpha)$ is different from the i -th symbol of $f(\alpha')$ which contradicts that f is uniformly continuous.

After the output symbol has been produced (using the dummy continuation), the dummy continuation can be discarded and \mathcal{M}' can restart the simulation of \mathcal{M} on the actual input in the configuration that \mathcal{M} was in before the simulation on the dummy continuation started. It is important to note that the next output symbol that is computed by \mathcal{M} (on the actual continuation) does not have to be produced by \mathcal{M}' as it was already determined using the dummy continuation (whereas the concrete dummy continuation does not influence the symbol as argued above). Hence, \mathcal{M}' must keep track of the size of the lead of the output symbols obtained on dummy continuations compared to the number of output symbols computed on the actual input word as not to produce outputs multiple times. \square

3.2 Transducers

Of course, Turing machines are a very expressive model of computation. Filiot and Winter show that for the uniformization of ω -regular relations, much less expressiveness is sufficient, i.e., for such relations, transducers, i.e., finite automata with output, suffice.

Formally, a (one-way deterministic finite) transducer \mathcal{T} is a tuple $(Q, \Sigma, \Gamma, q_I, \delta, \Omega)$ that consists of a finite set Q of states containing the initial state q_I , an input alphabet Σ , an output alphabet Γ , a transition function $\delta: Q \times \Sigma \rightarrow Q \times \Gamma^*$, and a coloring $\Omega: Q \rightarrow \mathbb{N}$. The (unique) run of \mathcal{T} on an input $w = w(0)w(1)w(2)\dots \in \Sigma^\omega$ is the sequence $q_0q_1q_2\dots$ of states defined by $q_0 = q_I$ and q_{i+1} being the unique state with $\delta(q_i, w(i)) = (q_{i+1}, x_i)$ for some $x_i \in \Gamma^*$. The run is accepting if the maximal color appearing infinitely often in $\Omega(q_0)\Omega(q_1)\Omega(q_2)\dots$ is even. With the run $q_0q_1q_2\dots$ on w we associate the output $x_0x_1x_2\dots$, where the x_i are as defined above. As the transducer is deterministic, it induces a map from inputs to outputs. Note that the output may, a priori, be a finite or an infinite word over Γ . In the following, we only consider transducers where the output is infinite for every input with an accepting run. In this case, \mathcal{T} computes a partial func-

tion $f_{\mathcal{T}}: \Sigma^\omega \rightarrow \Gamma^\omega$ defined as follows: the domain of $f_{\mathcal{T}}$ is the set of infinite words $w \in \Sigma^\omega$ such that the run of \mathcal{T} on w is accepting and $f_{\mathcal{T}}(w)$ is the output induced by this (unique) run.

We say that \mathcal{T} has delay $d \in \mathbb{N}$ if for every accepting run and every induced sequence $x_0x_1x_2\cdots$ of outputs (x_i is the output on the i -th transition), we have $i - d \leq |x_0 \cdots x_{i-1}| \leq i$ for all $i \geq 0$, i.e., the output is, at any moment during an accepting run, at most d letters shorter than the input and never longer. We say that a transducer \mathcal{T} is a bounded-delay transducer if there exists a d such that \mathcal{T} has delay d .

Proposition 1 ([15]) *The following are equivalent for a relation R encoded by a Büchi automaton \mathcal{A} and with closed $\text{dom}(R)$:*

1. R is uniformized by a computable function.
2. R is uniformized by a function implemented by a bounded-delay transducer.

As explained above, this covers the case of $\forall^*\exists^*$ formulas. In the remainder, we generalize this result to full HyperLTL, i.e., arbitrary quantifier alternations.

4 Computing Skolem Functions for HyperLTL

Our goal is to determine under which circumstances $\mathfrak{T} \models \varphi$ has a computable explanation, i.e., there are computable Skolem functions witnessing $\mathfrak{T} \models \varphi$, and whether such Skolem functions can be computed by “simpler” models of computation, i.e., bounded-delay transducers.

We start by showing that $\mathfrak{T} \models \varphi$ does not necessarily have a computable explanation.

Theorem 1 *There is a HyperLTL sentence φ and a transition system \mathfrak{T} such that $\mathfrak{T} \models \varphi$ is not witnessed by computable Skolem functions.*

Proof Consider $\varphi = \forall \pi \exists \pi'. (F a_\pi) \leftrightarrow (X a_{\pi'})$ and \mathfrak{T} with $\text{Tr}(\mathfrak{T}) = \emptyset(2^{\{a\}})^\omega$.

Towards a contradiction, assume there is a computable Skolem function for π' . Then, due to Lemma 1, there is also one that is implemented by a bounded-delay Turing machine \mathcal{M} , say with delay d . Now, let \mathcal{M} run on an input with prefix $\emptyset^{d+2} \in \text{Prfs}(\text{Tr}(\mathfrak{T}))$. As \mathcal{M} has bounded delay, it will produce the first two output letters $\emptyset A \in \emptyset 2^{\{a\}}$ after processing the prefix \emptyset^{d+2} (note that all traces of \mathfrak{T} start with \emptyset , the label of the initial state).

If $A = \emptyset$, then the output of \mathcal{M} on the input $\emptyset^{d+2}\{a\}^\omega$ starts with $\emptyset\emptyset$ (as this output only depends on the prefix \emptyset^{d+2}), but the input contains an $\{a\}$. These traces do not satisfy $(F a_\pi) \leftrightarrow (X a_{\pi'})$. On the other hand, if $A = \{a\}$, then the output of \mathcal{M} on the input \emptyset^ω starts with $\emptyset\{a\}$ (again, the output only depends on the prefix \emptyset^{d+2}), but the input contains no $\{a\}$. Again, these traces do not satisfy $(F a_\pi) \leftrightarrow (X a_{\pi'})$. So, in both cases, \mathcal{M} does not implement a Skolem function for π' , i.e., we have the desired contradiction. \square

So, as not every $\mathfrak{T} \models \varphi$ is witnessed by computable Skolem functions, it is natural to ask whether it is decidable, given \mathfrak{T} and φ , if $\mathfrak{T} \models \varphi$ has such a witness. Before we study this problem, we consider another example showing that for some transition system \mathfrak{T} and

sentence φ , even if $\mathfrak{T} \models \varphi$ does have computable Skolem functions, not every (computable) Skolem function is a “good” Skolem function: Fixing a Skolem function for an outermore variable may block innermore variables having computable Skolem functions.

Example 2 Consider the sentence $\exists\pi\forall\pi'\exists\pi''. (X a_\pi) \rightarrow ((F a_{\pi'}) \leftrightarrow (X a_{\pi''}))$ and a transition system \mathfrak{T} with $\text{Tr}(\mathfrak{T}) = \emptyset(2^{\{a\}})^\omega$. Note that every trace of \mathfrak{T} starts with \emptyset . Also, as the quantification of π is not in the scope of any other quantifier we can identify Skolem functions for π with traces that are assigned to π .

Now, if we pick a trace t for π with $a \in t(1)$ then there is no computable Skolem function for π'' (see Theorem 1). However, if we pick a trace t for π with $a \notin t(1)$ then every function is a Skolem function for π'' , as satisfaction is independent of the choices for π' and π'' in this case. In particular, π'' has a computable Skolem function.

Thus, the wrong choice of a (computable) Skolem function for some variable may result in other variables not having computable Skolem functions. By carefully accounting for the dependencies between the Skolem functions we show that the existence of computable Skolem functions is decidable.

Theorem 2 *The following problem is decidable: “Given a transition system \mathfrak{T} and a Hyper-LTL sentence φ with $\mathfrak{T} \models \varphi$, is $\mathfrak{T} \models \varphi$ witnessed by computable Skolem functions?” If the answer is yes, our algorithm computes bounded-delay transducers implementing such Skolem functions.*

The next section is dedicated to presenting a decidable game-theoretic characterization of the existence of computable Skolem functions, thereby proving Theorem 2.

5 A game for computable Skolem functions

Recall that in Subsection 3.1, we have explained that the special case of $\forall^*\exists^*$ sentences can be solved by a reduction to a uniformization problem. We begin this section by giving some intuition for this reduction. To simplify our notation, we consider a sentence of the form $\forall\pi\exists\pi'. \psi$ where ψ is quantifier-free. Here, we need to decide whether there is a computable function $f: \text{Tr}(\mathfrak{T}) \rightarrow \text{Tr}(\mathfrak{T})$ such that $\{\pi \mapsto t, \pi' \mapsto f(t)\} \models \psi$ for all t . Note that $\{\text{mrg}(t, t') \mid t, t' \in \text{Tr}(\mathfrak{T}) \text{ and } \{\pi \mapsto t, \pi' \mapsto t'\} \models \psi\}$ is accepted by a Büchi automaton (see Remark 2). Hence, the problem indeed boils down to a uniformization problem for an ω -regular relation. This problem was first posed (and partially solved) by Hosch and Landweber in 1971 [22] and later completely solved in a series of works [15, 20, 24]. Let us sketch the main ideas underlying the solution, as we will generalize them in the following.

Let $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ be ω -regular. Then, the existence of a function $f: \Sigma_I^\omega \rightarrow \Sigma_O^\omega$ that uniformizes L is captured by a (perfect information) two-player game $\Gamma(L)$ of infinite duration played between Player I (the input player) and Player O (the output player) in rounds $r = 0, 1, 2, \dots$ as follows: In each round r , Player I picks an $a_r \in \Sigma_I$ and then Player O picks a $b_r \in \Sigma_O \cup \{\varepsilon\}$. Thus, the outcome of a play of $\Gamma(L)$ is an infinite word $a_0a_1a_2 \dots \in \Sigma_I^\omega$ picked by Player I and a finite or infinite word $b_0b_1b_2 \dots \in \Sigma_O^* \cup \Sigma_O^\omega$ picked by Player O . The outcome is winning for Player O if $a_0a_1a_2 \dots \in \text{dom}(L)$ implies

$\text{mrg}(a_0a_1a_2\cdots, b_0b_1b_2\cdots) \in L$ (which requires that $b_0b_1b_2\cdots$ is infinite, i.e., Player O has to pick infinitely often a letter from Σ_O). Now, one can show that a winning strategy for Player O can be turned into a function that uniformizes L and every function uniformizing L can be turned into a winning strategy for Player O .

So far, the uniformizing function may be arbitrary, in particular not computable. Also, the delay between input and output in plays that are consistent with a strategy can be unbounded, e.g., if Player O picks ε in every second round. A crucial insight is that this is not necessary: for every ω -regular $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ such that Player O wins $\Gamma(L)$, there is a bound ℓ (that only depends on the size of a (minimal) Büchi automaton accepting L) such that she has a winning strategy that picks ε at most ℓ times [20].

This insight allows to change the rules of $\Gamma(L)$, giving Player O the advantage gained by using ε a bounded number of times from the beginning of a play and grouping moves into blocks of letters of a fixed length. How the block length is obtained is explained below.

The block game $\Gamma_b(L)$ is played in rounds $r = 0, 1, 2, \dots$ as follows: In round 0, Player I picks two blocks $x_0, x_1 \in \Sigma_I^\ell$ and then Player O picks a block $y_0 \in \Sigma_O^\ell$. Then, in every round $r > 0$, Player I picks a block $x_{r+1} \in \Sigma_I^\ell$ and then Player O picks a block $y_r \in \Sigma_O^\ell$. Note that Player I is one block ahead, as he has to pick two blocks in round 0. This accounts for the delay allowed in the definition of computable functions. The outcome of a play of $\Gamma_b(L)$ is an infinite word $x_0x_1x_2\cdots \in \Sigma_I^\omega$ picked by Player I and an infinite word $y_0y_1y_2\cdots \in \Sigma_O^\omega$ picked by Player O . The outcome is winning for Player O if $x_0x_1x_2\cdots \in \text{dom}(L)$ implies $\text{mrg}(x_0x_1x_2\cdots, y_0y_1y_2\cdots) \in L$.

Now, one can show that L is uniformizable iff Player O has a winning strategy for $\Gamma_b(L)$. As $\Gamma_b(L)$ is a finite two-player game with ω -regular winning condition, Player O has a finite-state winning strategy (one implemented by a transducer). Such a finite-state winning strategy can be turned into a computable function uniformizing L , as a transducer can be simulated by a Turing machine. Hence, $\Gamma_b(L)$ does indeed characterize uniformizability of ω -regular relations by computable functions.

Next, let us give some intuition of how to obtain the bound ℓ . To this end, let \mathcal{A} be a Büchi automaton over some alphabet $\Sigma_I \times \Sigma_O$ (we first ignore the acceptance condition in this discussion and later hint at how this is taken into account). As usual, if two finite words $w_0 \in (\Sigma_I \times \Sigma_O)^*$ and $w_1 \in (\Sigma_I \times \Sigma_O)^*$ induce the same state transformations (e.g., for all states p and q , processing w_0 from p leads \mathcal{A} to q iff processing w_1 from p leads \mathcal{A} to q), then these words are indistinguishable for \mathcal{A} (again, we are ignoring acceptance for the time being), i.e., one can replace w_0 by w_1 without changing the possible runs that \mathcal{A} has. This indistinguishability is captured by an equivalence relation over $(\Sigma_I \times \Sigma_O)^*$ with finite index, i.e., the equivalence relation has only finitely many equivalence classes.

However, to capture the interaction described in $\Gamma(L)$ above, we need a more refined approach. Assume Player I picks a sequence $x \in \Sigma_I^*$ of letters. Then, Player O will have to “complete” this block by picking a block $y \in \Sigma_I^{|x|}$ so that $\text{mrg}(x, y)$ is processed by the automaton. In this situation, we can say that x_0 and $x_1 \in \Sigma^*$ are equivalent if they are indistinguishable w.r.t. to their completions to words of the form $\text{mrg}(x_i, y_i)$, e.g., for all states p and q , there is a completion $\text{mrg}(x_0, y_0) \in (\Sigma_I \times \Sigma_O)^*$ of x_0 that leads \mathcal{A} from p to q iff there is a completion $\text{mrg}(x_1, y_1) \in (\Sigma_I \times \Sigma_O)^*$ of x_1 that leads \mathcal{A} from p to q . Intuitively, one does not need to distinguish between x_0 and x_1 because they allow Player O to achieve the same state transformations in \mathcal{A} . This indistinguishability is captured by an equivalence

relation over Σ_I^* of finite index. Now, ℓ can be picked as an upper bound on the length of a minimal word in all equivalence classes.

Thus, the intuition behind the definition of $\Gamma_b(L)$ is that blocks of length ℓ are *rich* enough to capture the full strategic choices for both players in $\Gamma(L)$: every longer word has an equivalent one of length at most ℓ .

Finally, let us briefly mention how to deal with the Büchi acceptance condition we have ignored thus far. As the state transformations are concerned with finite runs of the automaton, we can just keep track of whether an accepting state has been visited or not during this run, all other information is irrelevant. Thus, the equivalence relation will take this (single bit of) information into account as well.

After having sketched the special case of a sentence of the form $\forall\pi\exists\pi'. \psi$, let us now illustrate the challenges we have to address to deal with more quantifier alternations, e.g., for a sentence of the form $\varphi = \forall\pi_0\exists\pi_1 \cdots \forall\pi_{k-2}\exists\pi_{k-1}. \psi$.

- We will consider a multi-player game with one player being in charge of providing traces for the universally quantified variables (generalizing Player *I* above) and one variable player for each existentially quantified variable (generalizing Player *O* above), i.e., altogether we have $\frac{k}{2} + 1$ players. Thus, the player in charge of the universally quantified variables produces traces t_0, t_2, \dots, t_{k-2} while each variable player produces a trace t_i (one for each odd i). These traces are again picked block-wise in rounds.
- The choices by the variable player producing t_i (i.e., i is odd) may only depend on the traces t_0, t_1, \dots, t_{i-1} in order to faithfully capture the semantics of φ . Hence, we need to consider a game of imperfect information, which allows us to hide the traces t_{i+1}, \dots, t_{k-1} from the player in charge of π_i .
- Recall that Player *I* is always one block ahead of Player *O* in $\Gamma_b(L)$, which accounts for the delay allowed in the definition of computable functions. With k traces to be picked (and t_i depending on t_0, t_1, \dots, t_{i-1}), there must be a gap of one block for each even i .

Now, we are able to present the details of our construction. For the remainder of this section, we fix a HyperLTL sentence φ and a transition system \mathfrak{T} with $\mathfrak{T} \models \varphi$. We assume² $\varphi = \forall\pi_0\exists\pi_1 \cdots \forall\pi_{k-2}\exists\pi_{k-1}. \psi$, and use the Büchi automaton $\mathcal{A}_\psi^\mathfrak{T} = (Q, (2^{\text{AP}})^k, q_I, \delta, F)$ constructed in Remark 2 recognizing the language

$$\{\text{mrg}(\Pi(\pi_0), \dots, \Pi(\pi_{k-1})) \mid \Pi(\pi_i) \in \text{Tr}(\mathfrak{T}) \text{ for all } 0 \leq i < k \text{ and } (\text{Tr}(\mathfrak{T}), \Pi) \models \psi\}.$$

In the following, we often need to work with tuples of finite words of the same length. To simplify our notation, from now on we only write $(w_0, w_1, \dots, w_{i-1})$ if each w_j is a word in $(2^{\text{AP}})^*$ such that $|w_0| = |w_1| = \dots = |w_{i-1}|$.

Equivalence Relations.

We begin by defining equivalence relations that capture the concept of indistinguishability discussed in the intuition above.

We write $\mathcal{A} : p \xrightarrow{w} q$ for states p, q of a Büchi automaton \mathcal{A} over an alphabet Σ , if \mathcal{A} has a run from p to q processing the word $w \in \Sigma^*$. Furthermore, we write $\mathcal{A} : p \xrightarrow{w} q$, if \mathcal{A} has a run from p to q processing the word $w \in \Sigma^*$ such that the run visits at least one accepting

²The following reasoning can easily be extended to general sentences with arbitrary quantifier prefixes, albeit at the cost of more complex notation. We substantiate this claim in Remark 4.

state. Finally, we write $\mathfrak{T} : u \xrightarrow{w} v$ for vertices u, v of a transition system \mathfrak{T} , if \mathfrak{T} has a path from u to v labeled by the word $w \in (2^{\text{AP}})^*$.

We continue by defining, for each $1 \leq i \leq k$, an equivalence relation \equiv_i between i -tuples of (finite) words with the intuition that two such tuples are i -equivalent if they do not need to be distinguished. For $i = k$, this means that the two tuples cannot be distinguished by $\mathcal{A}_\psi^\mathfrak{T}$ while for $1 \leq i < k$ this means that both i -tuples can be completed (by adding an $(i + 1)$ -th component) so that the resulting $(i + 1)$ -tuples are \equiv_{i+1} -equivalent.

Formally, we define

$$(w_0, w_1, \dots, w_{k-1}) \equiv_k (\tilde{w}_0, \tilde{w}_1, \dots, \tilde{w}_{k-1})$$

if

- for all states p, q of $\mathcal{A}_\psi^\mathfrak{T}$ we have $\mathcal{A}_\psi^\mathfrak{T} : p \xrightarrow{\text{mrg}(w_0, w_1, \dots, w_{k-1})} q$ if and only if $\mathcal{A}_\psi^\mathfrak{T} : p \xrightarrow{\text{mrg}(\tilde{w}_0, \tilde{w}_1, \dots, \tilde{w}_{k-1})} q$,
- for all states p, q of $\mathcal{A}_\psi^\mathfrak{T}$ we have $\mathcal{A}_\psi^\mathfrak{T} : p \xRightarrow{\text{mrg}(w_0, w_1, \dots, w_{k-1})} q$ if and only if $\mathcal{A}_\psi^\mathfrak{T} : p \xRightarrow{\text{mrg}(\tilde{w}_0, \tilde{w}_1, \dots, \tilde{w}_{k-1})} q$, and
- for all vertices u, v of \mathfrak{T} and all $0 \leq j \leq k - 1$ we have $\mathfrak{T} : u \xrightarrow{w_j} v$ if and only if $\mathfrak{T} : u \xrightarrow{\tilde{w}_j} v$.

Lemma 2 *Let*

$$w = \text{mrg}(w_0^0 w_0^1 w_0^2 \dots, w_1^0 w_1^1 w_1^2 \dots, \dots, w_{k-1}^0 w_{k-1}^1 w_{k-1}^2 \dots)$$

and

$$\tilde{w} = \text{mrg}(\tilde{w}_0^0 \tilde{w}_0^1 \tilde{w}_0^2 \dots, \tilde{w}_1^0 \tilde{w}_1^1 \tilde{w}_1^2 \dots, \dots, \tilde{w}_{k-1}^0 \tilde{w}_{k-1}^1 \tilde{w}_{k-1}^2 \dots)$$

be such that

$$(w_0^n, w_1^n, \dots, w_{k-1}^n) \equiv_k (\tilde{w}_0^n, \tilde{w}_1^n, \dots, \tilde{w}_{k-1}^n)$$

for all $n \in \mathbb{N}$.

1. $w \in L(\mathcal{A}_\psi^\mathfrak{T})$ if and only if $\tilde{w} \in L(\mathcal{A}_\psi^\mathfrak{T})$.
2. For all $0 \leq i \leq k - 1$, we have $w_i^0 w_i^1 w_i^2 \dots \in \text{Tr}(\mathfrak{T})$ if and only if $\tilde{w}_i^0 \tilde{w}_i^1 \tilde{w}_i^2 \dots \in \text{Tr}(\mathfrak{T})$.

Proof 1.) Let $w \in L(\mathcal{A}_\psi^\mathfrak{T})$, Then, there exists an accepting run of $\mathcal{A}_\psi^\mathfrak{T}$ on w . Let q_{n+1} be the state reached by the run after the prefix

$$\text{mrg}(w_0^0 \dots w_0^{n-1}, w_1^0 \dots w_1^{n-1}, \dots, w_{k-1}^0 \dots w_{k-1}^{n-1}),$$

which implies $q_0 = q_I$. Then, we have $\mathcal{A}_\psi^\Sigma : q_{n-1} \xrightarrow{(w_0^n, w_1^n, \dots, w_{k-1}^n)} q_n$ for all n and $\mathcal{A}_\psi^\Sigma : q_{n-1} \xrightarrow{(w_0^n, w_1^n, \dots, w_{k-1}^n)} q_n$ for infinitely many n .

Due to

$$(w_0^n, w_1^n, \dots, w_{k-1}^n) \equiv_k (\tilde{w}_0^n, \tilde{w}_1^n, \dots, \tilde{w}_{k-1}^n)$$

for all n , we also have $\mathcal{A}_\psi^\Sigma : q_{n-1} \xrightarrow{(\tilde{w}_0^n, \tilde{w}_1^n, \dots, \tilde{w}_{k-1}^n)} q_n$ for all n and $\mathcal{A}_\psi^\Sigma : q_{n-1} \xrightarrow{(\tilde{w}_0^n, \tilde{w}_1^n, \dots, \tilde{w}_{k-1}^n)} q_n$ for infinitely many n . This allows us to conclude that there is also an accepting run of \mathcal{A}_ψ^Σ on \tilde{w} .

2.) The proof is very analogous one to the previous one, we just have to argue about paths and vertices of \mathcal{T} instead of runs and states of \mathcal{A}_ψ^Σ (and ignore acceptance) and consider the words w_i^n and \tilde{w}_i^n from the i -th component instead of full k -tuples $(w_0^n, w_1^n, \dots, w_{k-1}^n)$ and $(\tilde{w}_0^n, \tilde{w}_1^n, \dots, \tilde{w}_{k-1}^n)$. \square

Now, for $1 \leq i < k$, we define \equiv_i inductively as follows:

$$(w_0, w_1, \dots, w_{i-1}) \equiv_i (\tilde{w}_0, \tilde{w}_1, \dots, \tilde{w}_{i-1})$$

if

- for all w_i with $|w_i| = |w_0|$ there exists a \tilde{w}_i with $|\tilde{w}_i| = |\tilde{w}_0|$ such that

$$(w_0, w_1, \dots, w_i) \equiv_{i+1} (\tilde{w}_0, \tilde{w}_1, \dots, \tilde{w}_i),$$

and

- for all \tilde{w}_i with $|\tilde{w}_i| = |\tilde{w}_0|$ there exists a w_i with $|w_i| = |w_0|$ such that

$$(w_0, w_1, \dots, w_i) \equiv_{i+1} (\tilde{w}_0, \tilde{w}_1, \dots, \tilde{w}_i).$$

Let us show that each \equiv_i is indeed an equivalence relation and has only finitely many equivalence classes.

Lemma 3 *Every \equiv_i is an equivalence relation of finite index.*

Proof By induction over i from k to 1. The induction start $i = k$ was proven by Büchi [31], so consider $i < k$.

First, it is straightforward to verify that \equiv_i is an equivalence relation, as \equiv_{i+1} is an equivalence relation. Now, we define $\text{ext}((w_0, w_1, \dots, w_{i-1}))$ to be the set of \equiv_{i+1} -equivalence classes containing a (w_0, w_1, \dots, w_i) for some w_i with $|w_i| = |w_0|$. Now, we define $(w_0, w_1, \dots, w_{i-1}) \equiv_i' (\tilde{w}_0, \tilde{w}_1, \dots, \tilde{w}_{i-1})$ if and only if $\text{ext}((w_0, w_1, \dots, w_{i-1})) = \text{ext}((\tilde{w}_0, \tilde{w}_1, \dots, \tilde{w}_{i-1}))$, which is an equivalence relation of finite index: The codomain of ext has at most 2^n elements, where n is the index of \equiv_{i+1} . Finally, \equiv_i' refines \equiv_i , which implies that \equiv_i has finite index as well. \square

Let ℓ be minimal such that each word w with $|w| \geq \ell$ is in an infinite \equiv_1 equivalence class. This is well-defined, as \equiv_1 has finite index, which implies that there are only finitely many words in finite equivalence classes. A block is a word in $(2^{\text{AP}})^\ell$.

Now we have the definitions at hand to define the game $\mathcal{G}(\mathfrak{T}, \varphi)$ that captures the existence of computable Skolem functions. To keep the notation manageable, we describe the game abstractly and defer the concrete definition as a multi-player graph game of imperfect information to Section 5.1.

The Abstract Game. The game $\mathcal{G}(\mathfrak{T}, \varphi)$ is played between Player U who picks traces for the universally quantified variables (by picking blocks) and a coalition of variable players $\{1, 3, \dots, k-1\}$, who pick traces for the existentially quantified variables (Player i for π_i), also by picking blocks. As in the intuition given above for the case of a formula of the form $\forall \pi \exists \pi'$. ψ , the rules of the game $\mathcal{G}(\mathfrak{T}, \varphi)$ need to account for the delay inherent to the definition of computable functions. In the $\forall \exists$ setting, this is covered by the fact that the player in charge of π is one block ahead of the player in charge of π' . With more quantifier alternations, we generalize this as follows for $\varphi = \forall \pi_0 \exists \pi_1 \dots \forall \pi_{k-2} \exists \pi_{k-1}$. ψ :

- The player in charge of π_{k-2} must be one block ahead of the player in charge of π_{k-1} .
- The player in charge of π_{k-3} does not have to be ahead of the player in charge of π_{k-2} , but must also not be behind.
- The player in charge of π_{k-4} must be one block ahead of the player in charge of π_{k-3} . This implies that the player in charge of π_{k-4} must be two blocks ahead of the player in charge of π_{k-1} .
- And so on.

So, the player in charge of π_{k-1} picks one block in round 0, the player in charge of π_{k-2} picks two blocks in round 0 (to be one block ahead), the player in charge of π_{k-3} picks two blocks in round 0, the player in charge of π_{k-4} picks three blocks in round 0 and so on. In general, we define $\Delta_i = \frac{k-(i-1)}{2}$ for (odd) $i \in \{1, 3, \dots, k-1\}$ and $\Delta_i = \Delta_{i+1} + 1$ for (even) $i \in \{0, 2, \dots, k-2\}$, e.g., we have $\Delta_{k-1} = 1$, $\Delta_{k-2} = 2$, $\Delta_{k-3} = 2$, and $\Delta_{k-4} = 3$ capturing the “delay” described above.

Now, we split each round $r = 0, 1, 2, \dots$ into subrounds (r, i) for by $i = 0, 1, \dots, k-1$.

- In subround $(0, i)$ of round 0 for even i , Player U picks Δ_i blocks $t_{i-1}^0, t_{i-1}^1, \dots, t_{i-1}^{\Delta_i-1}$.
- In subround $(0, i)$ of round 0 for odd i , Player i picks Δ_i blocks $t_i^0, t_i^1, \dots, t_i^{\Delta_i-1}$.
- In subround (r, i) of round $r > 0$ for even i , Player U picks a block $t_{i-1}^{\Delta_i+r}$.
- In subround (r, i) of round $r > 0$ for odd i , Player i picks a block $t_i^{\Delta_i-1+r}$.

Figure 1 illustrates the evolution of a play and illustrates the number of blocks picked in round 0 and the resulting “delay” between the selection of blocks for the different variables.

During a play of $\mathcal{G}(\mathfrak{T}, \varphi)$ the players build traces t_0, t_1, \dots, t_{k-1} defined as $t_i = t_i^0 t_i^1 t_i^2 \dots$. We call $(t_0, t_1, \dots, t_{k-1})$ the outcome of the play. The coalition of variable players wins the play if $t_i \notin \text{Tr}(\mathfrak{T})$ for some even i or if $\text{mrg}(t_0, t_1, \dots, t_{k-1}) \in L(\mathcal{A}_{\psi}^{\mathfrak{T}})$, i.e., the variable assignment mapping each π_i to t_i satisfies ψ and each t_i is in $\text{Tr}(\mathfrak{T})$.

As already alluded to above, the game described above must be a game of imperfect information to capture the fact that the Skolem function for an existentially quantified π_i




Fig. 1 The evolution of a play of $\mathcal{G}(\Sigma, \varphi)$ for a sentence φ with six variables. Each gray shape is a subround, consisting of a move of Player U or a move of one variable player. We have $\Delta_5 = 1$, $\Delta_4 = \Delta_3 = 2$, $\Delta_2 = \Delta_1 = 3$, and $\Delta_0 = 4$, which corresponds to the number of blocks picked by the player in charge of variable π_i in round 0

depends only on the universally quantified variables π_j with $j \in \{0, 2, \dots, i-1\}$. Intuitively, we capture this by giving Player i access to all blocks picked in subrounds (r, j) with $j \in \{0, 2, \dots, i-1\}$, but hiding all other picks made by the players in subrounds (r, j) with $j \in \{1, 3, \dots, i-2, i, i+1, i+2, \dots, k-1\}$. Note that Player i not having access to their own moves is not a restriction, as they can always be reconstructed, if necessary.

Formally, a strategy for Player i is a function σ_i mapping sequences of the form³

$$(t_0^0, t_2^0, \dots, t_{i-1}^0)(t_1^1, t_3^1, \dots, t_{i-1}^1) \dots (t_0^{\Delta_i+r}, t_2^{\Delta_i+r}, \dots, t_{i-1}^{\Delta_i+r})(t_0^{\Delta_i+r+1}, t_2^{\Delta_i+r+1}, \dots, t_{i-1}^{\Delta_i+r+1}) \\ (t_0^{\Delta_i+r+2}, t_2^{\Delta_i+r+2}, \dots, t_{i-5}^{\Delta_i+r+2}) \dots (t_0^{\Delta_i+r+\frac{\Delta_i}{2}-1}, t_2^{\Delta_i+r+\frac{\Delta_i}{2}-1}, \dots, t_{i-1}^{\Delta_i+r+\frac{\Delta_i}{2}-1}) \quad (1)$$

for $r \geq 0$ to a block (or a sequence of Δ_i blocks for $r = 0$). The vectors getting shorter at the end is a manifestation of the fact that the players in charge of variables π_j with smaller j are ahead of the players in charge of variables with larger j' (see Figure 1).

A (finite or infinite) play is consistent with σ_i , if the pick of Player i in each round r is the one prescribed by σ_i . A collection $(\sigma_i)_{i \in \{1, 3, \dots, k-1\}}$ of strategies, one for each variable player, is winning, if every play that is consistent with all σ_i is won by the variable players. We say that a strategy σ_i is finite-state, if it is implemented by a transducer that reads inputs as in Equation (1) (over some suitable finite alphabet) and produces an output block (or a sequence of Δ_i blocks in round 0).

The following lemma shows that the existence of a winning collection of strategies characterizes the existence of computable Skolem functions. Note that there is a slight mismatch, as the first implication requires the strategies to be finite-state, while the second implication only yields arbitrary strategies. This gap will be closed later.

³ In this equation, we use column vectors for a more compact presentation.

- Lemma 4 1.** If the coalition of variable players has a winning collection of finite-state strategies then $\mathfrak{T} \models \varphi$ has computable Skolem functions.
2. If $\mathfrak{T} \models \varphi$ has computable Skolem functions, then the coalition of variable players has a winning collection of strategies.

Proof We first show Item 1. So, let $(\sigma_i)_{i \in \{1,3,\dots,k-1\}}$ be a winning collection of finite-state strategies for the variable players. We construct computable Skolem functions $(f_i)_{i \in \{1,3,\dots,k-1\}}$ witnessing $\mathfrak{T} \models \varphi$. So, fix some $i \in \{1,3,\dots,k-1\}$.

The machine \mathcal{M}_i computing f_i works in iterations $n = 0, 1, 2, \dots$ coinciding with the rounds of $\mathcal{G}(\mathfrak{T}, \varphi)$. Its input is $\text{mrg}(t_0, t_2, \dots, t_{i-1})$ (encoding $\frac{i}{2}$ input traces as a single infinite word on the input tape), where we split each t_j into blocks $t_j^0 t_j^1 t_j^2 \dots$. Recall that the block length ℓ is a constant, i.e., \mathcal{M}_i can read its input blockwise.

Now, in iteration 0, \mathcal{M}_i reads the first $\Delta_i + \frac{i}{2}$ blocks of the input, which yields $\Delta_i + \frac{i}{2}$ blocks of each t_j . These blocks can be used to simulate the moves of Player U in subrounds $(0, j)$ for even $j < i$. Note that this does not require all blocks of the t_j for $j > 0$. These have to be stored in the working tape for later use, as the reading tape is one-way. The simulated moves by Player U can be fed into the finite-state implementation of σ_i , yielding blocks $t_i^0, t_i^1, \dots, t_i^{\Delta_i-1}$ as output. The word $t_i^0 t_i^1 \dots t_i^{\Delta_i-1}$ is then written to the output tape of \mathcal{M}_i , which completes iteration 0.

In general, assume \mathcal{M}_i has completed iteration $n - 1$ and now starts iteration $n > 0$. This iteration begins with \mathcal{M}_i reading another block of the input, which yields another block of each t_j . The new block of t_0 , and the oldest stored block for each t_j with $j > 0$ can be used to continue the simulated play (restricted to moves by Player U in subrounds for the variables π_j for $j \in \{0, 2, \dots, i - 1\}$) by feeding them into the finite-state implementation of σ_i , yielding a block t as output. This block is then appended on the output tape. The unused new blocks of t_j with $j > 0$ are again stored on the working tape. This ends iteration n .

To simulate the play, \mathcal{M}_i can just store the whole play prefix on the working tape. To process the play prefix by the finite-state implementation of σ_i , \mathcal{M}_i can just store the whole run prefix on the working tape, although a more economical approach is possible (see the proof of Theorem 2 on Page 29).

Now, we show that the functions f_i computed by the \mathcal{M}_i constructed above are indeed Skolem functions witnessing $\mathfrak{T} \models \varphi$. To this end, let Π with $\text{dom}(\Pi) \supseteq \{\pi_0, \pi_1, \dots, \pi_{k-1}\}$ be a variable assignment that is consistent with the f_i , i.e., each $\Pi(\pi_i)$ with even i is in $\text{Tr}(\mathfrak{T})$ and each $\Pi(\pi_i)$ with odd i is equal to $f_i(\Pi(\pi_0), \Pi(\pi_2), \dots, \Pi(\pi_{i-1}))$. We need to show that each $\Pi(\pi_i)$ for odd i is in $\text{Tr}(\mathfrak{T})$ (i.e., the functions f_i are well-defined) and that $\Pi \models \psi$, i.e., $\text{mrg}(\Pi(\pi_0), \Pi(\pi_1), \dots, \Pi(\pi_{k-1})) \in L(\mathcal{A}_{\psi}^{\mathfrak{T}})$.

By construction, $(\Pi(\pi_0), \Pi(\pi_1), \dots, \Pi(\pi_{k-1}))$ is the outcome of a play of $\mathcal{G}(\mathfrak{T}, \varphi)$ that is consistent with the σ_i and therefore winning for the variable players. As each $\Pi(\pi_i)$ with even i is in $\text{Tr}(\mathfrak{T})$, we conclude $\text{mrg}(\Pi(\pi_0), \Pi(\pi_1), \dots, \Pi(\pi_{k-1})) \in L(\mathcal{A}_{\psi}^{\mathfrak{T}})$, as required. Note that this does also imply $\Pi(\pi_i)$ for odd i is in $\text{Tr}(\mathfrak{T})$, as $L(\mathcal{A}_{\psi}^{\mathfrak{T}})$ only contains tuples of traces from $\text{Tr}(\mathfrak{T})$.

Now, let us consider Item 2. Let $f_i: (\text{Tr}(\mathfrak{T}))^{\frac{i+1}{2}} \rightarrow \text{Tr}(\mathfrak{T})$ for $i \in \{1, 3, \dots, k-1\}$ be computable Skolem functions witnessing $\mathfrak{T} \models \varphi$, say each f_i is implemented by a Turing machine \mathcal{M}_i . By Lemma 1, each \mathcal{M}_i can be assumed to have bounded delay: there is a d_i such that to compute the first n letters of the output only $n + d_i$ letters of the input are

read. Note that we can run such a Turing machine \mathcal{M}_i with delay d_i on a finite input w of length $n + d_i$ and obtain the first n letters of the output $f_i(w')$ of every infinite w' that starts with the prefix w . We will apply this fact to simulate the \mathcal{M}_i on-the-fly on longer and longer prefixes.

Also note that our definition of a function f being computed by a Turing machine \mathcal{M} only requires it to compute the output $f(w)$ for all $w \in \text{dom}(f)$, but it can produce arbitrary (even finite) outputs for $w \notin \text{dom}(f)$. To simplify our construction, we assume here that each \mathcal{M}_i produces an infinite output for every input, even if it is not in the domain of f_i . This can be done w.l.o.g., as the \mathcal{M}_i have bounded delay d_i : as soon as \mathcal{M}_i wants to access input letter $n + d_i + 1$ without having produced $n + 1$ output letters so far (this can be detected, as d_i is a constant), the run does not have delay d_i , which implies that the input cannot be in $\text{dom}(f_i)$. Hence, a designated state can be entered, which produces an arbitrary infinite output while ignoring the remaining input. The resulting machine still has delay d_i , but a complete domain.

Let $d \in \mathbb{N}$ be minimal such that each \mathcal{M}_i has delay at most d . We inductively define a winning collection of strategies for the variable players.

Round 0.

Subrounds (0, 0) and (0, 1). Assume Player U picks $t_0^0, t_0^1, \dots, t_0^{\Delta_0-1}$ in subround (0, 0) to start a play. We fix $\tilde{t}_0^0 = t_0^0$ and fix \tilde{t}_0^n for $n \in \{1, 2, \dots, \Delta_0 - 1\}$ such that $\tilde{t}_0^n \equiv_1 t_0^n$ and $|\tilde{t}_0^0 \tilde{t}_0^1 \dots \tilde{t}_0^n| \geq |\tilde{t}_0^0 \tilde{t}_0^1 \dots \tilde{t}_0^{n-1}| + d$ for all such n . This is always possible, as the \equiv_1 equivalence class of each t_0^n is infinite and therefore contains arbitrarily long words.

Let $\tilde{t}_1^0 \tilde{t}_1^1 \dots \tilde{t}_1^{\Delta_0-2}$ be the output of \mathcal{M}_1 when given the partial input $\tilde{t}_0^0 \tilde{t}_0^1 \dots \tilde{t}_0^{\Delta_0-1}$ such that $|\tilde{t}_1^n| = |\tilde{t}_0^n|$ for all n . This is well-defined by the choice of the length of the \tilde{t}_0^n and the fact that \mathcal{M}_1 has delay d . Note that \mathcal{M}_1 might produce even more output on that input. Any such additional output is ignored in this subround.

As we have $t_0^n \equiv_1 \tilde{t}_0^n$ for all such n , there also exists a t_1^n with $|t_1^n| = |\tilde{t}_0^n|$ such that $(t_0^n, t_1^n) \equiv_2 (\tilde{t}_0^n, \tilde{t}_1^n)$. We define σ_1 such that it picks $t_1^0, t_1^1, \dots, t_1^{\Delta_0-2}$ in subround (0, 1). As $\Delta_0 - 2 = \Delta_1 - 1$, these are Δ_1 many blocks, as required by the definition of $\mathcal{G}(\mathcal{T}, \varphi)$.

Subrounds (0, 2) and (0, 3). Now, assume Player U picks $t_2^0, t_2^1, \dots, t_2^{\Delta_2-1}$ in subround (0, 2). We fix $\tilde{t}_2^0 = t_2^0$ and then fix \tilde{t}_2^n for $n \in \{1, 2, \dots, \Delta_2 - 1\}$ such that $(\tilde{t}_0^n, \tilde{t}_1^n, \tilde{t}_2^n) \equiv_3 (t_0^n, t_1^n, t_2^n)$. This is possible, as we have $(t_0^n, t_1^n) \equiv_2 (\tilde{t}_0^n, \tilde{t}_1^n)$ for all such n . Let $\tilde{t}_3^0 \tilde{t}_3^1 \dots \tilde{t}_3^{\Delta_2-2}$ be the output of \mathcal{M}_3 when given the partial input $\text{mrg}(\tilde{t}_0^0 \tilde{t}_0^1 \dots \tilde{t}_0^{\Delta_0-1}, \tilde{t}_2^0 \tilde{t}_2^1 \dots \tilde{t}_2^{\Delta_2-1})$ such that $|\tilde{t}_3^n| = |\tilde{t}_0^n|$ for all n (again, this is well-defined due to the choice of the length of the \tilde{t}_0^n and \mathcal{M}_3 having delay d , and might require to ignore some output).

As we have $(\tilde{t}_0^n, \tilde{t}_1^n, \tilde{t}_2^n) \equiv_3 (t_0^n, t_1^n, t_2^n)$ for all such n , there also exists a t_3^n with $|t_3^n| = |\tilde{t}_0^n|$ such that $(t_0^n, t_1^n, \dots, t_3^n) \equiv_4 (\tilde{t}_0^n, \tilde{t}_1^n, \dots, \tilde{t}_3^n)$. We define σ_3 such that it picks $t_3^0, t_3^1, \dots, t_3^{\Delta_2-2}$ in subround (0, 3). As $\Delta_2 - 2 = \Delta_3 - 1$, these are Δ_3 many blocks, as required by the definition of $\mathcal{G}(\mathcal{T}, \varphi)$.

Subrounds (0, $i - 1$) and (0, i) for odd $i \in \{5, 7, \dots, k - 1\}$. Assume Player U picks $t_{i-1}^0, t_{i-1}^1, \dots, t_{i-1}^{\Delta_{i-1}-1}$ in subround (0, $i - 1$). As before, we fix $\tilde{t}_{i-1}^0 = t_{i-1}^0$ and then fix \tilde{t}_{i-1}^n for $n \in \{1, 2, \dots, \Delta_{i-1} - 1\}$ such that $(\tilde{t}_0^n, \tilde{t}_1^n, \dots, \tilde{t}_{i-1}^n) \equiv_i (t_0^n, t_1^n, \dots, t_{i-1}^n)$. This is possible, as $(t_0^n, t_1^n, \dots, t_{i-2}^n) \equiv_{i-1} (\tilde{t}_0^n, \tilde{t}_1^n, \dots, \tilde{t}_{i-2}^n)$ for all such n is an invariant of our construction.

Let $\tilde{t}_i^0 \tilde{t}_i^1 \dots \tilde{t}_i^{\Delta_{i-1}-2}$ be the output of \mathcal{M}_i when given the partial input

$$\text{mrg}(\tilde{t}_0^0 \tilde{t}_0^1 \dots \tilde{t}_0^{\Delta_0-1}, \tilde{t}_2^0 \tilde{t}_2^1 \dots \tilde{t}_2^{\Delta_2-1}, \dots, \tilde{t}_{i-1}^0 \tilde{t}_{i-1}^1 \dots \tilde{t}_{i-1}^{\Delta_{i-1}-1})$$

such that $|\tilde{t}_i^n| = |\tilde{t}_0^n|$ for all n . As in the previous cases, this is well-defined.

As we have $(\tilde{t}_0^n, \tilde{t}_1^n, \dots, \tilde{t}_{i-1}^n) \equiv_i (t_0^n, t_1^n, \dots, t_{i-1}^n)$ for all such n , there also exists a t_i^n with $|t_i^n| = |\tilde{t}_0^n|$ such that $(t_0^n, t_1^n, \dots, t_i^n) \equiv_{i+1} (\tilde{t}_0^n, \tilde{t}_1^n, \dots, \tilde{t}_i^n)$ for all n , satisfying the invariant again. We define σ_i such that it picks $t_i^0, t_i^1, \dots, t_i^{\Delta_{i-1}-2}$ in subround $(0, i)$. As $\Delta_{i-1} - 2 = \Delta_i - 1$, these are Δ_i many blocks, as required by the definition of $\mathcal{G}(\mathfrak{T}, \varphi)$.

Round $r > 0$. Now, we consider a round $r > 0$, assuming the σ_i are already defined for all earlier rounds. The construction is very similar to the one for round 0, but simpler as each player (also Player U !) only picks a single block in each subround (r, i) of round r .

Subrounds $(r, 0)$ and $(r, 1)$ Assume Player U picks $t_0^{\Delta_0-1+r}$ in subround $(r, 0)$. We fix $\tilde{t}_0^{\Delta_0-1+r}$ such that $t_0^{\Delta_0-1+r} \equiv_1 \tilde{t}_0^{\Delta_0-1+r}$ and $|\tilde{t}_0^0 \tilde{t}_1^1 \dots \tilde{t}_0^{\Delta_0-1+r}| \geq |\tilde{t}_0^0 \tilde{t}_1^1 \dots \tilde{t}_0^{\Delta_0+r-2}| + d$. This is possible, as the \equiv_1 equivalence class of $t_0^{\Delta_0-1+r}$ is infinite and therefore contains arbitrarily long words.

We run \mathcal{M}_1 on $\tilde{t}_0^0 \tilde{t}_1^1 \dots \tilde{t}_0^{\Delta_0-1+r}$ and obtain another block $\tilde{t}_1^{\Delta_0+r-2}$. There is a $t_1^{\Delta_0+r-2}$ with $|t_1^{\Delta_0+r-2}| = |t_0^{\Delta_0+r-2}|$ such that $(t_0^{\Delta_0+r-2}, t_1^{\Delta_0+r-2}) \equiv_2 (\tilde{t}_0^{\Delta_0+r-2}, \tilde{t}_1^{\Delta_0+r-2})$, as we have $t_0^{\Delta_0+r-2} \equiv_1 \tilde{t}_0^{\Delta_0+r-2}$. We define σ_1 such that it picks the block $t_1^{\Delta_0+r-2}$ in subround $(r, 1)$ (note that $\Delta_0 + r - 2 = \Delta_1 - 1 + r$).

Subrounds $(r, i-1)$ and (r, i) for $i \in \{3, 5, \dots, k-1\}$. Now, assume Player U picks $t_{i-1}^{\Delta_{i-1}-1+r}$ in subround $(r, i-1)$. We fix $\tilde{t}_{i-1}^{\Delta_{i-1}-1+r}$ such that $(\tilde{t}_0^{\Delta_{i-1}-1+r}, \tilde{t}_1^{\Delta_{i-1}-1+r}, \dots, \tilde{t}_{i-1}^{\Delta_{i-1}-1+r}) \equiv_i (t_0^{\Delta_{i-1}-1+r}, t_1^{\Delta_{i-1}-1+r}, \dots, t_{i-1}^{\Delta_{i-1}-1+r})$. This is possible, as $(\tilde{t}_0^{\Delta_{i-1}-1+r}, \tilde{t}_1^{\Delta_{i-1}-1+r}, \dots, \tilde{t}_{i-2}^{\Delta_{i-1}-1+r}) \equiv_{i-1} (t_0^{\Delta_{i-1}-1+r}, t_1^{\Delta_{i-1}-1+r}, \dots, t_{i-2}^{\Delta_{i-1}-1+r})$ is an invariant of our construction.

We run \mathcal{M}_i on

$$\text{mrg}(\tilde{t}_0^0 \tilde{t}_1^1 \dots \tilde{t}_0^{\Delta_{i-1}-1+r}, \tilde{t}_2^0 \tilde{t}_1^1 \dots \tilde{t}_2^{\Delta_{i-1}-1+r}, \dots, \tilde{t}_{i-1}^0 \tilde{t}_{i-1}^1 \dots \tilde{t}_{i-1}^{\Delta_{i-1}-1+r}),$$

yielding another block $\tilde{t}_i^{\Delta_{i-1}+r-2}$. As

$$(\tilde{t}_0^{\Delta_{i-1}+r-2}, \tilde{t}_1^{\Delta_{i-1}+r-2}, \dots, \tilde{t}_{i-1}^{\Delta_{i-1}+r-2}) \equiv_i (t_0^{\Delta_{i-1}+r-2}, t_1^{\Delta_{i-1}+r-2}, \dots, t_{i-1}^{\Delta_{i-1}+r-2}),$$

there also exists a $t_i^{\Delta_{i-1}+r-2}$ with $|t_i^{\Delta_{i-1}+r-2}| = |t_0^{\Delta_{i-1}+r-2}|$ such that $(t_0^{\Delta_{i-1}+r-2}, t_1^{\Delta_{i-1}+r-2}, \dots, t_i^{\Delta_{i-1}+r-2}) \equiv_{i+1} (\tilde{t}_0^{\Delta_{i-1}+r-2}, \tilde{t}_1^{\Delta_{i-1}+r-2}, \dots, \tilde{t}_i^{\Delta_{i-1}+r-2})$. We define σ_i such that it picks $t_i^{\Delta_{i-1}+r-2}$ in subround (r, i) (note that $\Delta_{i-1} + r - 2 = \Delta_i - 1 + r$).

This completes the definition of the σ_i . Note that each σ_i does indeed only depend on the blocks picked in subrounds (r, j) with $j \in \{0, 2, \dots, i-1\}$, i.e., σ_i is indeed a strategy for Player i in $\mathcal{G}(\mathfrak{T}, \varphi)$.

It remains to show that the σ_i are a winning collection of strategies. To this end, let $(t_0, t_1, \dots, t_{k-1})$ be an outcome of a play that is consistent with the σ_i . If a t_i with even i is not in $\text{Tr}(\mathfrak{T})$, then the variable players win immediately. So, assume each t_i with even i is in $\text{Tr}(\mathfrak{T})$. Let $\tilde{t}_0, \tilde{t}_1, \dots, \tilde{t}_{k-1}$ be the traces constructed during the inductive definition of the σ_i . By applying Remark 2.2, we obtain that each \tilde{t}_i with even i is in $\text{Tr}(\mathfrak{T})$ as well. Also, the \tilde{t}_i for odd i satisfy $\tilde{t}_i = f_i(\tilde{t}_0, \tilde{t}_2, \dots, \tilde{t}_{i-1})$ by construction, i.e., they are obtained by

applying the Skolem functions. Hence, the variable assignment mapping π_i to \tilde{t}_i satisfies ψ , which implies that $\text{mrg}(\tilde{t}_0, \tilde{t}_1, \dots, \tilde{t}_{i-1})$ is in $L(\mathcal{A}_{\psi}^{\mathcal{T}})$.

Now, as $(\tilde{t}_0^n, \tilde{t}_1^n, \dots, \tilde{t}_{k-1}^n) \equiv_k (t_0^n, t_1^n, \dots, t_{k-1}^n)$ for all n , applying Lemma 2.1 yields that $\text{mrg}(t_0, t_1, \dots, t_{i-1})$ is in $L(\mathcal{A}_{\psi}^{\mathcal{T}})$ as well, i.e., the variable players do indeed win. \square

Remark 4 Let us substantiate the claim in Footnote 2 that our construction can also be applied to formulas with arbitrary, i.e., not strictly alternating, quantifier prefixes. Then, we have one player in the coalition for each block of existentially quantified variables in the prefix, who selects traces for these variables. The variables of each of these players share their Δ_i values, as these only depend on the number of alternations between quantifiers that are outmore than the currently considered one.

For example, for a formula of the form

$$\forall \pi_0 \exists \pi_1 \exists \pi_2 \forall \pi_3 \exists \pi_4 \exists \pi_5. \psi$$

with quantifier-free ψ , we have just two players (say, Player 12 and Player 45) in the coalition. Player 12 is in charge of the variables π_1 and π_2 , and Player 45 is in charge of the variables π_4 and π_5 . These two players are in a coalition against Player U , who is in charge of the variables π_0 and π_3 . In Round 0, Player U picks three blocks for π_0 , then Player 12 picks two blocks for π_1 and two blocks for π_2 , then Player U picks two blocks for π_3 , and then Player 45 picks one block for π_4 and one block for π_5 . Afterwards, the players pick in alternation (one block for each of their variables), i.e., π_0 is always one block ahead of both π_1 and π_2 and two blocks ahead of both π_4 and π_5 , and π_3 is always one block ahead of both π_4 and π_5 .

Note that just adding universally quantified dummy variables between any pair of consecutive existentially quantified variables is logically equivalent, but makes explanations more complex (as they may depend on the dummy variables now) and increases the complexity of our approach (as that depends on the number of players (see Subsection 5.2)).

Now, one can formalize $\mathcal{G}(\mathcal{T}, \varphi)$ as a multi-player graph game of (hierarchical) imperfect information. The existence of a winning collection of strategies is decidable for such games [5, 28]. Furthermore, if there is a winning collection of strategies, then also a winning collection of finite-state strategies, which closes the gap in the statement of Lemma 4: $\mathcal{T} \models \varphi$ has computable Skolem functions iff the coalition of variable players has a winning collection of finite-state strategies. Furthermore, similarly to the proof of Lemma 4.1, one can show that such finite-state strategies can even be implemented by bounded-delay transducers, thereby completing the proof of Theorem 2.

5.1 The concrete game

After having shown that the abstract game $\mathcal{G}(\mathcal{T}, \varphi)$ characterizes the existence of computable Skolem functions, we now model $\mathcal{G}(\mathcal{T}, \varphi)$ as a multi-player graph game of imperfect information using the notation of Berwanger et al. [5]. In Subsection 5.1.1, we introduce the necessary definitions before we model the game in Subsection 5.1.2. The games considered by Berwanger et al. are concurrent games (i.e., the players make their moves simultane-

ously), while $\mathcal{G}(\mathcal{T}, \varphi)$ is turn-based, i.e., the players make their moves one after the other. Hence, we will also introduce some notation for the special case of turn-based games, which simplifies our modeling.

5.1.1 Distributed games

Game Graphs Fix some set $N = \{1, \dots, n\}$ of players and a distinguished agent called Nature (which is not in N). A profile is a list $p = (p_1, \dots, p_n)$ of elements $p_i \in P_i$ for sets P_i that will be clear from context. For each player $i \in N$ we fix a finite set A^i of actions and a finite set B^i of observations. A game graph $G = (V, E, v_I, (\beta^i)_{i \in N})$ consists of a finite set V of positions, an edge relation $E \subseteq V \times A \times V$ representing simultaneous moves labeled by action profiles (i.e., $A = A^1 \times \dots \times A^n$), an initial position $v_I \in V$, and a profile $(\beta^i)_{i \in N}$ of observation functions $\beta^i: V \rightarrow B^i$ that label, for each player, the positions with observations. We require that E has no dead ends, i.e., for every $v \in V$ and every $a \in A$ there is a $v' \in V$ with $(v, a, v') \in E$.

A game graph $(V, E, v_I, (\beta^i)_{i \in N})$ yields hierarchical information if there exists a total order \preceq over N such that if $i \preceq j$ then for all $v, v' \in V$, $\beta^i(v) = \beta^i(v')$ implies $\beta^j(v) = \beta^j(v')$, i.e., if Player i cannot distinguish v and v' , then neither can Player j for $i \preceq j$.

Plays Intuitively, a play starts at position $v_I \in V$ and proceeds in rounds. In a round at position v , each Player i chooses simultaneously and independently an action $a^i \in A^i$, then Nature chooses a successor position v' such that $(v, a, v') \in E$. Now, each player receives the observation $\beta^i(v')$ and the next round is played at position v' . Thus, a play of G is an infinite sequence $v_0 v_1 v_2 \dots$ of vertices such that $v_0 = v_I$ and for all $r \geq 0$ there is an $a_r \in A$ such that $(v_r, a_r, v_{r+1}) \in E$.

A history is a prefix $v_0 v_1 \dots v_r$ of a play. We denote the set of all histories by $\text{Hist}(G)$ and extend $\beta^i: V \rightarrow B^i$ to plays and histories by defining $\beta^i(v_0 v_1 v_2 \dots) = \beta^i(v_1) \beta^i(v_2) \beta^i(v_3) \dots$. Note that the observation of the initial position is discarded for technical reasons [5]. We say two histories h and h' are indistinguishable to Player i , denoted by $h \sim_i h'$, if $\beta^i(h) = \beta^i(h')$.

Strategies A strategy for Player i is a mapping $s^i: V^* \rightarrow A^i$ that satisfies $s^i(h) = s^i(h')$ for all h, h' with $h \sim_i h'$ (i.e., the action selected by the strategy only depends on the observations of the history). A play $v_0 v_1 v_2 \dots$ is consistent with s^i if for every $r \geq 0$, there is an $a_r = (a^1, \dots, a^n) \in A$ with $(v_r, a_r, v_{r+1}) \in E$ and $a^i = s^i(v_0 v_1 \dots v_r)$. A play is consistent with a strategy profile (s^1, \dots, s^n) if it is consistent with each s^i . The set of possible outcomes of a strategy profile is the set of all plays that are consistent with s .

A distributed game $\mathcal{G} = (G, W)$ consists of a game graph and a winning condition $W \subseteq V^\omega$, where V is the set of positions of G . A play is winning if it is in W and a strategy profile S is winning in \mathcal{G} if all its outcomes are winning.

Finite-state Strategies Next, we define what it means for a strategy for Player i to be finite-state. So far, we used transducers, i.e., automata with output on transitions to implement strategies in a finitary manner. From now on, we follow the definitions used by Berwanger et al. [5] and use Moore machines, i.e., finite automata with output on states. However, each Moore machine can be transformed into a transducer by “moving” the output from a state to all its outgoing transitions.

Let A^i and B^i be the actions and observations of Player i . A Moore machine $\mathcal{S} = (M, m_I, \text{upd}, \text{nxt})$ for Player i consists of a finite set M of memory states containing the initial memory state m_I , a memory update function $\text{upd}: M \times B^i \rightarrow M$, and a next-move function $\text{nxt}: M \rightarrow A^i$. We extend upd to words over B^i by defining $\text{upd}(\varepsilon) = m_I$ and $\text{upd}(b_0 b_1 \cdots b_r) = \text{upd}(\text{upd}(b_0 b_1 \cdots b_{r-1}), b_r)$. We say \mathcal{S} implements the strategy mapping $v_0 v_1 \cdots v_r$ to $\text{nxt}(\text{upd}(\beta^i(v_0 v_1 \cdots v_r)))$. A strategy is finite-state, if it is implemented by some Moore machine.

Proposition 2 ([5, 28])

1. *The following problem is decidable: Given a distributed game with ω -regular winning condition, does it have a winning strategy profile?*
2. *A distributed game with ω -regular winning condition has a winning strategy profile if and only if it has a winning profile of finite-state strategies.*

Turn-based Game Graphs We say that a game graph $G = (V, E, v_I, (\beta^i)_{i \in N})$ is turn-based, if there is a function $o: V \rightarrow N$ such that if $(v, a, v') \in E$, $(v, a', v'') \in E$, and the action profiles a and a' having the same action for Player $o(v)$ implies $v' = v''$. To simplify our notation, we label the edges leaving v only by actions of Player $o(v)$. Thus, in a turn-based game graph, at every position v Player $o(v)$ determines the possible next moves, and Nature selects one of them. Turn-based distributed games are distributed games whose game graphs are turn-based.

5.1.2 Formalization of the abstract game

Now, we are finally ready to formalize the abstract game $\mathcal{G}(\mathfrak{T}, \varphi)$ described in Section 5 as a turn-based distributed game.

We begin by introducing notation for the positions of the game, which intuitively keep track of blocks picked by the players of $\mathcal{G}(\mathfrak{T}, \varphi)$ until they can be processed by an automaton recognizing the winning condition. Due to the delay between the choices by the different players, this requires some notation.

Recall that we have defined Δ_i for even i to be the number of blocks Player U picks in subround $(0, i)$ for variable i and for odd i to be the number of blocks Player i picks in subround $(0, i)$ for variable i . Now, let

$$D = \{(j, x) \mid j \in \{0, 1, \dots, k-1\} \text{ and } x \in \{0, 1, \dots, \Delta_i - 1\}\}.$$

A configuration is a partial function $c: D \rightarrow \mathcal{B}$, where $\mathcal{B} = (2^{\text{AP}})^\ell$ denotes the set of blocks. Let C denote the set of all configurations. The following definitions are visualized in Figure 2.

We only need certain types of configurations c for our construction. We say that c is

- a full configuration if $\text{dom}(c) = D$,
- an initial i -configuration (for $i \in \{0, 1, \dots, k-1\}$) if

$$\text{dom}(c) = \{(j, x) \mid j \in \{0, 1, \dots, i-1\} \text{ and } x \in \{0, 1, \dots, \Delta_j - 1\}\},$$




Fig. 2 Illustrating configurations for a formula with six variables, where a filled circle denotes an element in the domain, and an unfilled circle an element that is not in the domain. The upper row shows initial i -configurations for $i = 0, 1, 2, \dots, 5$ (from left to right), the lower row shows looping i -configurations for $i = 0, 1, 2, \dots, 5$ (from left to right), and the configuration on the right in between the rows is full. Solid arrows show the effect of extending and shifting configurations. The shifting operation is illustrated using the numbers in the circles. Finally, in a play of $\mathcal{G}(\mathcal{I}, \varphi)$ the configurations stored in positions follow the solid arrows, but take the dashed shortcuts avoiding full configurations

and

- a looping i -configuration (again for $i \in \{0, 1, \dots, k-1\}$) if

$$\text{dom}(c) = \{(j, x) \mid j \in \{0, 1, \dots, i-1\} \text{ and } x \in \{0, 1, \dots, \Delta_j - 1\}\} \cup \{(j, x) \mid j \in \{i, i+1, \dots, k-1\} \text{ and } x \in \{0, 1, \dots, \Delta_j - 2\}\}.$$

Note that for $i = k-1$, both the definition of initial and looping i -configuration coincides, as we have $\Delta_{k-1} = 1$. Hence, in the following, we will just speak of $(k-1)$ -configurations whenever convenient.

Given an initial i -configuration c and a sequence $\bar{b} = b_0, b_1, \dots, b_{\Delta_i-1}$ of Δ_i blocks, we define $\text{ext}(c, \bar{b})$ to be the configuration c' defined as

$$c'(j, x) = \begin{cases} c(j, x) & \text{if } (j, x) \in \text{dom}(c), \\ b_x & \text{if } j = i, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Furthermore, given a looping i -configuration c and a block b , we define $\text{ext}(c, b)$ to be the configuration c' defined as

$$c'(j, x) = \begin{cases} c(j, x) & \text{if } (j, x) \in \text{dom}(c), \\ b & \text{if } j = i \text{ and } x = \Delta_i - 1, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Finally, given a full configuration c , we define $\text{shft}(c)$ to be the configuration c' defined as

$$c'(j, x) = \begin{cases} c(j, x + 1) & \text{if } x < \Delta_j - 1, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

The following remark collects how these operations update initial and looping configurations.

Remark 5 1. If c is an initial i -configuration for $i < k - 1$, then $\text{ext}(c, \bar{b})$ is an initial $(i + 1)$ -configuration.

2. If c is a $(k - 1)$ -configuration, then $\text{ext}(c, b)$ is a full configuration.

3. If c is a full configuration, then $\text{shft}(c)$ is a looping 0-configuration.

4. If c is a looping i -configuration for $i < k - 1$, then $\text{ext}(c, b)$ is a looping $(i + 1)$ -configuration.

Finally, let $\mathcal{P} = (Q, (2^{\text{AP}})^k, q_I, \delta, \Omega)$ be a deterministic parity automaton accepting the language of words of the form $\text{mrg}(t_0, t_1, \dots, t_{k-1}) \in ((2^{\text{AP}})^k)^\omega$ such that either $t_i \notin \text{Tr}(\mathfrak{T})$ for some even i or if $\text{mrg}(t_0, t_1, \dots, t_{k-1}) \in L(\mathcal{A}_{\psi}^{\mathfrak{T}})$ (i.e., accepting the winning outcomes of plays of $\mathcal{G}(\mathfrak{T}, \varphi)$).

Now, we are able to formally define $\mathcal{G}(\mathfrak{T}, \varphi)$. It will be played by the players in $N = \{1, 3, \dots, k - 1\}$ (ignoring, for the sake of readability, the fact that N is not of the form $\{1, 2, \dots, n\}$ for some n , as required by the definitions in Subsection 5.1.1). Furthermore, the role of Player U will be played by Nature.

We define the set of positions to contain all tuples (i, c, q) where $i \in \{0, 1, \dots, k - 1\}$, c is an (initial or looping) i -configuration, and q is a state of \mathcal{P} , together with a sink state s_\perp . The initial position is $(0, c_\perp, q_I)$ where c_\perp is the configuration with empty domain (which is the unique initial 0-configuration).

We define the action set for Player i (for odd i) as $\mathcal{B}^{\Delta_i} \cup \mathcal{B}$, where actions in \mathcal{B}^{Δ_i} are intended for round 0 and those in \mathcal{B} are intended for all other rounds. If the wrong action is used, then the sink state will be reached. Next, we define the function o determining which player picks an action to continue a play: we have $o(i, c, q) = i$ for odd i , $o(i, c, q) = 1$ for even i (we will soon explain how Player U moves are simulated even though Player 1 owns the corresponding positions), and $o(s_\perp) = 1$.

The set E of edges is defined as follows (recall that we label edges by actions of a single player, as we define a turn-based game):

- We begin by modelling the moves of Player U . Recall that in a turn-based distributed game, Nature resolves the nondeterminism left after the player who is in charge at that positions has picked an action. Thus, we simulate a move of Player U by giving the position to (say) Player 1. Then, we define the edges such that Player 1's move is irrelevant, but the nondeterminism models the choice of Player U .

Formally, for $i \in \{0, 2, \dots, k - 2\}$, an initial i -configuration c , and a state q of \mathcal{P} , we have the edge $((i, c, q), a, (i + 1, \text{ext}(c, \bar{b}), q))$ for every action a of Player 1 and every $\bar{b} \in \mathcal{B}^{\Delta_i}$: No matter which action a Player 1 picks, Nature can for each possible \bar{b} pick a successor that extends c by \bar{b} . This indeed simulates the move of Player U in sub-round $(0, i)$.

- For $i \in \{1, 3, \dots, k - 3\}$, an initial i -configuration c , and a state q of \mathcal{P} , we have the

- edge $((i, c, q), \bar{b}, (i + 1, \text{ext}(c, \bar{b}), q))$ for each $\bar{b} \in \mathcal{B}^{\Delta_i}$ (this simulates the move of Player i in subround $(0, i)$) as well as the edge $((i, c, q), b, s_{\perp})$ for each $b \in \mathcal{B}$ (Player i may not pick a single block in subround $(0, i)$ if $i < k - 1$). Note that there is no nondeterminism to resolve for Nature, as there is a unique successor position for each action.
- For a $(k - 1)$ -configuration c and state q of \mathcal{P} , we have the edge $((k - 1, c, q), b, (0, \text{shift}(c'), q'))$ for each $b \in \mathcal{B}$ (recall that we have $\Delta_{k-1} = 1$ here), where $c' = \text{ext}(c, b)$ and q' is the state reached by \mathcal{P} when processing $\text{mrg}(c'(0, 0), c'(1, 0), \dots, c'(k - 1, 0))$ from q .
 - For $i \in \{0, 2, \dots, k - 2\}$, a looping i -configuration c , and a state q of \mathcal{P} , we have the edge $((i, c, q), a, (i + 1, \text{ext}(c, b), q))$ for every action a of Player 1 and every $b \in \mathcal{B}$: No matter which action a Player 1 picks, Nature can for each possible b pick a successor that extends c by b . This simulates the move of Player U in a subround $(r, i + 1)$ for $r > 0$ using the same mechanism as described above for moves of Player U in initial configurations.
 - For $i \in \{1, 3, \dots, k - 3\}$, a looping i -configuration c , and a state q of \mathcal{P} , we have the edge $((i, c, q), b, (i + 1, \text{ext}(c, b), q))$ for each $b \in \mathcal{B}$ (this simulates the move of Player i in subround (r, i) for $r > 0$) as well as the edge $((i, c, q), \bar{b}, s_{\perp})$ for each $\bar{b} \in \mathcal{B}^{\Delta_i}$ (Player i may not pick a sequence of blocks in a subround (r, i)). Again, there is no nondeterminism to resolve for Nature, as there is a unique successor position for each action.
 - For completeness, we have the edge $(s_{\perp}, a, s_{\perp})$ for every action a of Player 1.

It remains to define the observation functions β^i as $\beta^i(i, c, q) = c \upharpoonright i$ where $c \upharpoonright i$ is the configuration defined as

$$(c \upharpoonright i)(j, x) = \begin{cases} c(j, x) & \text{if } j \in \{0, 2, \dots, i - 1\}, \\ \text{undefined} & \text{otherwise,} \end{cases}$$

i.e., Player i can only observe the blocks picked for the variables π_j with even $j < i$. For completeness, we also define the observation of the sink state s_{\perp} to be \perp , where $\perp \notin C$. Then, the order $1 \succeq 3 \succeq \dots \succeq k - 1$ witnesses that the game graph yields hierarchical information.

This completes the definition of the game graph. To complete the definition of the game we define the winning condition as follows: Let $(i_0, c_0, q_0)(i_1, c_1, q_1)(i_2, c_2, q_2) \dots$ be a play and let $(i_0, c_0, q_0)(i_k, c_k, q_k)(i_{2k}, c_{2k}, q_{2k}) \dots$ be the subsequence of all (initial and looping) 0-configurations. Recall that the state q_{rk} in such a position (for $r > 0$) is obtained by processing some $\text{mrg}(b_0, \dots, b_{k-1})$ from $q_{(r-1)k}$, where the b_i are blocks picked by the players. We say $(i_0, c_0, q_0)(i_1, c_1, q_1)(i_2, c_2, q_2) \dots$ is in the winning condition, if $\Omega(q_0)\Omega(q_k)\Omega(q_{2k}) \dots$ satisfies the parity condition, which is an ω -regular winning condition. In particular, no winning play may visit the sink s_{\perp} .

Remark 6 The (concrete) distributed game constructed here is a formalization of the abstract game $\mathcal{G}(\mathfrak{T}, \varphi)$ described in Section 5. In particular, a winning collection of (finite-state) strategies for the coalition of players in the abstract game corresponds to a winning (finite-state) strategy profile in the concrete game and vice versa.

Hence, whenever convenient below, we do not distinguish between the concrete and the abstract game.

Lemma 5 *If $\mathcal{G}(\mathfrak{T}, \varphi)$ has a winning profile of finite-state strategies, then $\mathfrak{T} \models \varphi$ is witnessed by Skolem functions computed by bounded-delay transducers which can be effectively computed from the finite-state strategies.*

Proof Let $(s^1, s^3, \dots, s^{k-1})$ be a winning profile of finite-state strategies in (the concrete) $\mathcal{G}(\mathfrak{T}, \varphi)$. We show by induction over i how the finite-state strategies s^i can be turned into bounded-delay transducers \mathcal{T}_i computing Skolem functions witnessing $\mathfrak{T} \models \varphi$. The proof follows closely the analogous results shown in Lemma 4.1 for Turing machine computable Skolem functions.

Let us fix some $i \in \{1, 3, \dots, k-1\}$ and let \mathcal{S}_i be a Moore machine for Player i implementing s^i . Recall that \mathcal{S}_i reads observations of Player i (configurations of the form $c \upharpoonright i$ for initial and looping configurations) and returns actions of Player i .

On the other hand, \mathcal{T}_i reads an input $\text{mrg}(t_0, t_2, \dots, t_{i-1}) \in ((2^{\text{AP}})^{\frac{\ell}{2}})^{\omega}$ where we split each t_j into blocks $t_j = t_j^0 t_j^1 t_j^2 \dots$.

We construct \mathcal{T}_i so that it works in two phases, an initialization phase and a looping phase, that is repeated ad infinitum. We begin by describing the initialization. It begins by \mathcal{T}_i reading Δ_i blocks from each t_j . These blocks can now be assembled into a sequence of observations of Player i corresponding to the play prefix of $\mathcal{G}(\mathfrak{T}, \varphi)$ in which Player U picks these blocks. Note that this requires to process each such observation twice, as Player i 's observation does not get updated, when Player j for odd $j < i$ makes a move. All unused blocks are stored in the state space of \mathcal{T} . Now, \mathcal{T}_i simulates the run of the Moore machine \mathcal{S}_i implementing s^i on this sequence of observations, yielding an action a^i . As s^i is winning, this action is a sequence $\bar{b} = b_0, b_1, \dots, b_{\Delta_i-1}$ of blocks. Then, \mathcal{T}_i outputs $b_0 b_1 \dots b_{\Delta_i-1}$. Now, we process the last observation another $k-i$ times with \mathcal{S}_i , simulating the moves for the remaining variables (which are hidden from Player i which implies the observation is unchanged). This concludes the initialization phase.

The looping phase begins with \mathcal{T}_i reading another block from each t_j . Again, these blocks and the ones stored in the state space can be assembled into observations of Player i corresponding to a continuation of the simulated play prefix in which Player U pick these blocks. Then, the run of the Moore machine \mathcal{S}_i implementing s^i can be continued, yielding an action a^i . This is now a block b , which is output by \mathcal{T}_i . Again, we process the last observation just assembled another $k-i$ times to simulate the moves for the remaining variables, which concludes one looping phase. Note that the delay of \mathcal{T}_i is bounded by $k \cdot \ell$, where ℓ is the block length.

By storing blocks from the input that have been read but not yet used in observations, by discarding blocks no longer needed, and by keeping track of the state the simulated run of \mathcal{S}_i ends in, this behaviour can indeed be implemented using a finite state-space. We leave the tedious, but straightforward, formal definition of \mathcal{T}_i to the reader. \square

Now, our main theorem (Theorem 2) is a direct consequence of Proposition 2, Lemma 4, and Lemma 5. Recall that the theorem states that the problem “Given a transition system \mathfrak{T} and a HyperLTL sentence φ with $\mathfrak{T} \models \varphi$, is $\mathfrak{T} \models \varphi$ witnessed by computable Skolem functions?” is decidable and that, if the answer is yes, our algorithm computes bounded-delay transducers implementing such Skolem functions.

Proof of Theorem 2 Due to Lemma 4 and Proposition 2.2, the following statements are equivalent:

- $\mathfrak{T} \models \varphi$ has computable Skolem functions.
- $\mathcal{G}(\mathfrak{T}, \varphi)$ has a winning strategy profile.
- $\mathcal{G}(\mathfrak{T}, \varphi)$ has a winning profile of finite-state strategies.

The last statement can be decided effectively due to Proposition 2.1. Thus, the existence of computable Skolem functions is decidable. Finally, bounded-delay transducers computing Skolem functions witnessing $\mathfrak{T} \models \varphi$ can be effectively computed due to Lemma 5. \square

5.2 Complexity analysis

In this subsection, we determine the complexity of our algorithm. Our benchmark here is the complexity of the model-checking problem for HyperLTL, which is TOWER-complete. More precisely, checking whether a given transition system \mathfrak{T} satisfies a given formula φ is complete for nondeterministic space bounded by a k -fold exponential in $|\varphi|$, where k is the number of alternations in φ [14, 30].

In the following, we bound the size of $\mathcal{G}(\mathfrak{T}, \varphi)$ constructed in Subsection 5.1.2 and then apply known complexity results for solving distributed games, thereby giving an upper bound on the complexity of deciding whether $\mathfrak{T} \models \varphi$ is witnessed by computable Skolem functions.

Recall that we start the construction of $\mathcal{G}(\mathfrak{T}, \varphi)$ with a transition system \mathfrak{T} and a HyperLTL sentence $\varphi = \forall \pi_0 \exists \pi_1 \cdots \forall \pi_{k-2} \exists \pi_{k-1}. \psi$. First, we construct the Büchi automaton $\mathcal{A}_{\psi}^{\mathfrak{T}}$ recognizing the language

$$\{\text{mrg}(\Pi(\pi_0), \dots, \Pi(\pi_{k-1})) \mid \Pi(\pi_i) \in \text{Tr}(\mathfrak{T}) \text{ for all } 0 \leq i < k \text{ and } (\text{Tr}(\mathfrak{T}), \Pi) \models \psi\}.$$

Due to Remark 2, we can bound the size of $\mathcal{A}_{\psi}^{\mathfrak{T}}$ by

$$|\psi| \cdot 2^{|\psi|} \cdot |\mathfrak{T}|^k \leq 2^{\log(|\psi| \cdot 2^{|\psi|} \cdot |\mathfrak{T}|^k)} = 2^{\log(|\psi|) + |\psi| + k \cdot \log(|\mathfrak{T}|)} \leq 2^{\mathcal{O}(|\psi| \cdot |\mathfrak{T}|)},$$

i.e., exponentially in $|\psi| \cdot |\mathfrak{T}|$.

Now, let n denote the size of $\mathcal{A}_{\psi}^{\mathfrak{T}}$ and s the size of \mathfrak{T} . Then, \equiv_k has index at most $3^{n^2} \cdot 2^{ks^2} \leq 2^{\mathcal{O}(n^2 + ks^2)}$: There are 3 choices for each pair (p, q) of states of $\mathcal{A}_{\psi}^{\mathfrak{T}}$, i.e.,

1. there is a run between p and q that visits at least one accepting state,
2. there is a run between p and q , but none that visits an accepting state, and
3. there is no run between p and q at all,

resulting in the factor 3^{n^2} , and there are two choices for each $j \in \{0, 1, \dots, k-1\}$ and each pair (u, v) of vertices of \mathfrak{T} , i.e.,

1. there is a path of \mathfrak{T} from u to v and
2. there is no path of \mathfrak{T} from u to v ,

resulting in the factor 2^{ks^2} .

Furthermore, every equivalence class of \equiv_k , which is a language of finite words over the alphabet $(2^{\text{AP}})^k$, is recognized by a DFA with at most

$$(3^n)^n \cdot (2^s)^{s \cdot k} = 3^{n^2} \cdot 2^{ks^2} \leq 2^{\mathcal{O}(n^2 + ks^2)}$$

states: For each state q of $\mathcal{A}_{\psi}^{\mathfrak{T}}$ the DFA keeps track of which states are reachable from q (with and without having seen an accepting state) by processing the input word, which requires the product of n modified powerset automata derived from $\mathcal{A}_{\psi}^{\mathfrak{T}}$, each with 3^n states. Furthermore, for each component j and each vertex v of \mathfrak{T} , it also keeps track of which vertices are reachable from v by paths labeled with the j -th component of the input word, again requiring s powerset automata derived from \mathfrak{T} , each with 2^s states. The DFA accepts if exactly the right states and vertices (i.e., those uniquely identifying the equivalence class) are reachable.

Now, let \equiv_{i+1} for $0 < i < k$ have index idx and let each of its equivalence classes be accepted by a DFA of size at most sz . Then, the reasoning used to prove Lemma 3 shows that \equiv_i has index at most 2^{idx} and that each equivalence class of \equiv_i is recognized by a DFA with at most sz^{idx} states, as membership in an \equiv_i equivalence class is determined by membership and non-membership in \equiv_{i+1} equivalence classes. This can be checked by the product of idx many DFA for the equivalence classes (or their complements) of \equiv_{i+1} . Note that we are working with DFA, so complementation is for free.

Now, one can show by an induction that the index of \equiv_i is bounded by a $(k - i + 2)$ -fold exponential in $|\psi| \cdot |\mathfrak{T}|$. Using this, a second induction shows that each equivalence class of \equiv_i is recognized by a DFA whose size is bounded by a $(k - i + 3)$ -fold exponential in $|\psi| \cdot |\mathfrak{T}|$. This relies on the fact that $(2^x)^y$ is equal to 2^{xy} , which implies that sz^{idx} is “just” exponentially larger than sz , even though both idx and sz are in general towers of exponentials.

Now, we can bound ℓ , which is defined so that every w in a finite equivalence class of \equiv_1 satisfies $|w| < \ell$. As a DFA with n states recognizing a finite language can only accept words of length at most $n - 1$, we can bound ℓ by the size of the DFA accepting the equivalence classes of \equiv_1 , i.e., by a $(k + 2)$ -fold exponential in $|\psi| \cdot |\mathfrak{T}|$.

The number of vertices of $\mathcal{G}(\mathfrak{T}, \varphi)$ is $k \cdot |C| \cdot |\mathcal{P}|$ where C is the set of configurations and where \mathcal{P} is a deterministic parity automaton recognizing the language of words of the form $\text{mrg}(t_0, t_1, \dots, t_{k-1}) \in ((2^{\text{AP}})^k)^\omega$ such that either $t_i \notin \text{Tr}(\mathfrak{T})$ for some even i or if $\text{mrg}(t_0, t_1, \dots, t_{k-1}) \in L(\mathcal{A}_{\psi}^{\mathfrak{T}})$.

The number of configurations can be bounded by

$$|C| \leq |\mathcal{B}|^{|D|} \leq ((|2^{\text{AP}}|)^\ell)^{k^2},$$

i.e., by a $(k + 3)$ -fold exponential in $|\psi| \cdot |\mathfrak{T}|$. Furthermore, the size of \mathcal{P} can be bounded doubly-exponentially in $|\varphi| \cdot |\mathfrak{T}|$, as it can be constructed as the product of a deterministic parity automaton that is equivalent to the nondeterministic Büchi automaton $\mathcal{A}_{\psi}^{\mathfrak{T}}$ and $\frac{k}{2}$ deterministic parity automata for the language $(2^{\text{AP}})^\omega \setminus \text{Tr}(\mathfrak{T})$. Altogether, the number of vertices of $\mathcal{G}(\mathfrak{T}, \varphi)$ is bounded by a $(k + 3)$ -fold exponential in $|\psi| \cdot |\mathfrak{T}|$.

As distributed games with hierarchical information can be solved in time that is bounded by an $(n + c)$ -fold exponential in the number of states (where n is the number of players and where c is a small constant that depends on the type of winning condition) [28], we conclude that $\mathcal{G}(\mathcal{T}, \varphi)$ can be solved in time that is bounded by a $(\frac{k}{2} + c)$ -fold exponential in the number of vertices of $\mathcal{G}(\mathcal{T}, \varphi)$.

Combining these estimations, we obtain an upper bound on the runtime of our algorithm for a given transition system \mathcal{T} and formula $\varphi = \forall \pi_0 \exists \pi_1 \dots \forall \pi_{k-2} \exists \pi_{k-1}. \psi$ with quantifier-free ψ .

Lemma 6 *The algorithm described in Section 5 determining whether $\mathcal{T} \models \varphi$ is witnessed by computable Skolem functions whose runtime is bounded by a $(k + \frac{k}{2} + c')$ -fold exponential in $|\psi| \cdot |\mathcal{T}|$ for a small constant c' .*

This results should again be contrasted with HyperLTL model-checking being TOWER-complete [27, 30] (in the number of quantifier-alternations).

6 Related work

As explained in the introduction, computing counterexamples for debugging is the most important application of model-checking. In the framework of LTL model-checking, a counterexample is a single trace of the system that violates the specification. Such a counterexample is typically obtained by running a model-checking algorithm (which are in fact based on searching for such counterexamples). Variations of the problem include bounded model-checking [7], which searches for “short” counterexamples. Also, counterexample-guided abstraction refinement [9] and bounded synthesis [16] rely on counterexample computation.

Counterexamples have not only been studied in the realm of linear-time logics, but also for many other frameworks, e.g., for \forall CTL [10], for CTL [19, 32], for probabilistic temporal logics [13, 23], and for discrete-time Markov models [1].

Furthermore, counterexamples have also been studied in the realm of HyperLTL model-checking. Horak et al. [21] developed HYPERVIS, a webtool which provides interactive visualizations of a given model, specification, and counterexample computed by the HyperLTL model-checker MCHYPER [14]. In complementary work, Coenen et al. [12] present a causality-based explanation method for HyperLTL counterexamples, again computed by MCHYPER. However, the counterexamples computed by MCHYPER are just sets of traces that violate the formula. More specifically, MCHYPER only considers the outermost universal quantifiers and returns a variable assignment to those. This is obviously complete for formulas with quantifier-prefix \forall^* , i.e., without existential quantifiers, but not for more general formulas. In fact, this approach ignores the dynamic dependencies between universally and existentially quantified variables that is captured by Skolem functions, which we analyze here. Finally, let us also mention that counterexamples are the foundation of the bounded synthesis algorithm for \forall^* HyperLTL specifications [17]. In this setting, it is again sufficient to only consider sets of traces and not general Skolem functions.

Beutner and Finkbeiner [6] have presented a reduction from HyperLTL model-checking to nondeterministic multi-agent planning so that if the resulting planning problem admits a plan, then the original model-checking problem is satisfied. Conceptually, their techniques

are similar to ours presented here, i.e., for general formulas they rely on multiple planning agents (as here, one for each existentially quantified variable), nondeterminism in the planning domain (selecting traces for the universally quantified variables as our Player U), and incomplete information (to capture the dependencies of the variables, as here). The plans computed by solving the planning problem can be understood as Skolem functions for the existentially quantified variables. They note that their approach is sound, but incomplete in the sense that there might be instances where $\mathcal{T} \models \varphi$, but the resulting planning problem admits no plan. However, they left open whether their approach is complete for all instances $\mathcal{T} \models \varphi$ being witnessed by computable Skolem functions.

Finally, Beutner and Finkbeiner also presented a sound and complete model-checking algorithm for $\forall^*\exists^*$ -sentences that relies on a *delay-free* game extended with prophecies that require the player in charge of the universal variables to make commitments about future moves (w.r.t. to membership in a finite list of ω -regular properties). Winter and Zimmermann [33] recently extended the prophecy-based approach to full HyperLTL with arbitrary quantifier alternations, thereby making game-based HyperLTL model-checking complete: their (multi-player incomplete information) game characterizes $\mathcal{T} \models \varphi$, even if there are no computable Skolem functions witnessing satisfaction.

7 Conclusion

Based on the maxim “counterexamples/explanations are Skolem functions for the existentially quantified trace variables”, we have shown how to explain why a given transition system satisfies a given HyperLTL formula or, equivalently, to provide counterexamples in case the system does not satisfy the formula. We consider arbitrary computable Skolem functions as explanations. However, this leads to incompleteness as not every $\mathcal{T} \models \varphi$ is witnessed by computable Skolem functions. Nevertheless, we have shown that the existence of computable explanations is decidable, and that they are effectively computable (whenever they exist).

Recall that the runtime of our algorithm is bounded by a $(k + \frac{k}{2} + c)$ -fold exponential while HyperLTL model-checking is complete for nondeterministic space bounded by a k -fold exponential in $|\varphi|$, where k is the number of alternations in φ [14, 30]. In future work we aim to determine whether computing counterexamples/explanations is inherently harder than model-checking.

After having laid the theoretical foundations, it remains to investigate under which circumstances such explanations can be useful in the verification work-flow. For restricted settings, this line of work has been studied as discussed in Section 6. We propose to continue this line of work for more expressive fragments of HyperLTL and explanations. However, for fragments with few quantifier alternations, the situation is clearer. For example, we conjecture that results on delay games with LTL winning conditions [26] and uniformization of LTL definable relations [15] can be adapted to show that deciding the existence of computable Skolem functions for $\forall^*\exists^*$ -sentences is 3EXPTIME-complete.

Epilogue After her lunch break, Tracy walks by the printer room and finds a second copy of her document in the printer tray. It turns out the print system satisfies φ_{id} , but the explanation shows that there is a large delay in the network.

Acknowledgements Supported by DIREC - Digital Research Centre Denmark.

Author Contributions S.W. and M.Z. contributed equally.

Funding Open access funding provided by Aalborg University

Data Availability No datasets were generated or analysed during the current study.

Declarations

Competing interests The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Ábrahám, E., Becker, B., Dehnert, C., Jansen, N., Katoen, J., Wimmer, R.: Counterexample generation for discrete-time Markov models: An introductory survey. In: Bernardo, M., Damiani, F., Hähnle, R., Johnsen, E.B., Schaefer, I. (eds.) SFM 2014. LNCS, vol. 8483, pp. 65–121. Springer, (2014). https://doi.org/10.1007/978-3-319-07317-0_3
2. Arkhangel'skiĭ, A.V., Fedorchuk, V.V.: The basic concepts and constructions of general topology. In: Arkhangel'skiĭ, L.S. A. V. and Pontryagin (ed.) General Topology I: Basic Concepts and Constructions Dimension Theory, pp. 1–90. Springer, (1990). https://doi.org/10.1007/978-3-642-61265-7_1
3. Arrieta, A.B., Rodríguez, N.D., Ser, J.D., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., Herrera, F.: Explainable artificial intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI. Inf. Fusion **58**, 82–115 (2020). <https://doi.org/10.1016/j.inffus.2019.12.012>
4. Baier, C., Katoen, J.: Principles of Model Checking. MIT Press (2008)
5. Berwanger, D., Mathew, A.B., Bogaard, M.: Hierarchical information and the synthesis of distributed strategies. Acta Informatica **55**(8), 669–701 (2018). <https://doi.org/10.1007/S00236-017-0306-5>
6. Beutner, R., Finkbeiner, B.: Non-deterministic planning for hyperproperty verification. In: Bernardini, S., Muise, C. (eds.) ICAPS 2024, pp. 25–30. AAAI Press, (2024). <https://doi.org/10.1609/ICAPS.V34I1.31457>
7. Biere, A., Cimatti, A., Clarke, E.M., Strichman, O., Zhu, Y.: Bounded model checking. Adv. Comput. **58**, 117–148 (2003). [https://doi.org/10.1016/S0065-2458\(03\)58003-2](https://doi.org/10.1016/S0065-2458(03)58003-2)
8. Clarkson, M.R., Finkbeiner, B., Koleini, M., Micinski, K.K., Rabe, M.N., Sánchez, C.: Temporal logics for hyperproperties. In: Abadi, M., Kremer, S. (eds.) POST 2014. LNCS, vol. 8414, pp. 265–284. Springer, (2014). https://doi.org/10.1007/978-3-642-54792-8_15
9. Clarke, E.M., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement for symbolic model checking. J. ACM **50**(5), 752–794 (2003). <https://doi.org/10.1145/876638.876643>
10. Clarke, E.M., Jha, S., Lu, Y., Veith, H.: Tree-like counterexamples in model checking. In: LICS 2002, pp. 19–29. IEEE Computer Society, (2002). <https://doi.org/10.1109/LICS.2002.1029814>
11. Clarkson, M.R., Schneider, F.B.: Hyperproperties. J. Comput. Secur. **18**(6), 1157–1210 (2010). <https://doi.org/10.3233/JCS-2009-0393>
12. Coenen, N., Dachsel, R., Finkbeiner, B., Frenkel, H., Hahn, C., Horak, T., Metzger, N., Siber, J.: Explaining hyperproperty violations. In: Shoham, S., Vizel, Y. (eds.) CAV 2022, Part I. LNCS, vol. 13371, pp. 407–429. Springer, (2022). https://doi.org/10.1007/978-3-031-13185-1_20
13. Damman, B., Han, T., Katoen, J.: Regular expressions for PCTL counterexamples. In: QEST 2008, pp. 179–188. IEEE Computer Society, (2008). <https://doi.org/10.1109/QEST.2008.11>

14. Finkbeiner, B., Rabe, M.N., Sánchez, C.: Algorithms for Model Checking HyperLTL and HyperCTL*. In: Kroening, D., Pasareanu, C.S. (eds.) CAV 2015, Part I. LNCS, vol. 9206, pp. 30–48. Springer, (2015). https://doi.org/10.1007/978-3-319-21690-4_3
15. Filiot, E., Winter, S.: Synthesizing computable functions from rational specifications over infinite words. *Int. J. Found. Comput. Sci.* **35**(1 & 2), 179–214 (2024). <https://doi.org/10.1142/S012905412348009X>
16. Finkbeiner, B., Schewe, S.: Bounded synthesis. *Int. J. Softw. Tools Technol. Transf.* **15**(5–6), 519–539 (2013). <https://doi.org/10.1007/S10009-012-0228-Z>
17. Finkbeiner, B., Hahn, C., Lukert, P., Stenger, M., Tentrup, L.: Synthesis from hyperproperties. *Acta Informatica* **57**(1–2), 137–163 (2020). <https://doi.org/10.1007/s00236-019-00358-2>
18. Grädel, E., Thomas, W., Wilke, T. (eds.): Automata, Logics, and Infinite Games: A Guide to Current Research. LNCS, vol. 2500. Springer, (2002). <https://doi.org/10.1007/3-540-36387-4>
19. Gurfinkel, A., Chechik, M.: Proof-like counter-examples. In: Garavel, H., Hatcliff, J. (eds.) TACAS 2003. LNCS, vol. 2619, pp. 160–175. Springer, (2003). https://doi.org/10.1007/3-540-36577-X_12
20. Holtmann, M., Kaiser, L., Thomas, W.: Degrees of lookahead in regular infinite games. *Log. Methods Comput. Sci.* **8**(3) (2012) [https://doi.org/10.2168/LMCS-8\(3:24\)2012](https://doi.org/10.2168/LMCS-8(3:24)2012)
21. Horak, T., Coenen, N., Metzger, N., Hahn, C., Flemisch, T., Méndez, J., Dimov, D., Finkbeiner, B., Dachselt, R.: Visual analysis of hyperproperties for understanding model checking results. *IEEE Trans. Vis. Comput. Graph.* **28**(1), 357–367 (2022). <https://doi.org/10.1109/TVCG.2021.3114866>
22. Hosch, F.A., Landweber, L.H.: Finite delay solutions for sequential conditions. In: ICALP 1972, pp. 45–60 (1972)
23. Han, T., Katoen, J.: Counterexamples in probabilistic model checking. In: Grumberg, O., Huth, M. (eds.) TACAS 2007. LNCS, vol. 4424, pp. 72–86. Springer, (2007). https://doi.org/10.1007/978-3-540-71209-1_8
24. Klein, F., Zimmermann, M.: How much lookahead is needed to win infinite games? *Log. Methods Comput. Sci.* **12**(3) (2016) [https://doi.org/10.2168/LMCS-12\(3:4\)2016](https://doi.org/10.2168/LMCS-12(3:4)2016)
25. Kurshan, R.P.: Verification technology transfer. In: Grumberg, O., Veith, H. (eds.) 25 Years of Model Checking - History, Achievements, Perspectives. LNCS, vol. 5000, pp. 46–64. Springer, (2008). https://doi.org/10.1007/978-3-540-69850-0_3
26. Klein, F., Zimmermann, M.: Prompt delay. In: Lal, A., Akshay, S., Saurabh, S., Sen, S. (eds.) FSTTCS 2016. LIPIcs, vol. 65, pp. 43–14314. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, (2016). <https://doi.org/10.4230/LIPICS.FSTTCS.2016.43>
27. Mascle, C., Zimmermann, M.: The keys to decidable HyperLTL satisfiability: Small models or very simple formulas. In: Fernández, M., Muscholl, A. (eds.) CSL 2020. LIPIcs, vol. 152, pp. 29–12916. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, (2020). <https://doi.org/10.4230/LIPICS.CSL.2020.29>
28. Pnueli, A., Rosner, R.: Distributed reactive systems are hard to synthesize. In: FOCS 1990, Volume II, pp. 746–757. IEEE Computer Society, (1990). <https://doi.org/10.1109/FSCS.1990.89597>
29. Pnueli, A.: The temporal logic of programs. In: FOCS 1977, pp. 46–57. IEEE, (1977). <https://doi.org/10.1109/SFCS.1977.32>
30. Rabe, M.N.: A temporal logic approach to information-flow control. PhD thesis, Saarland University (2016). <http://scidok.sulb.uni-saarland.de/volltexte/2016/6387/>
31. Richard Büchi, J.: Symposium on decision problems: On a decision method in restricted second order arithmetic. In: Nagel, E., Suppes, P., Tarski, A. (eds.) *Logic, Methodology and Philosophy of Science. Studies in Logic and the Foundations of Mathematics*, vol. 44, pp. 1–11. Elsevier, (1966). [https://doi.org/10.1016/S0049-237X\(09\)70564-6](https://doi.org/10.1016/S0049-237X(09)70564-6)
32. Shoham, S., Grumberg, O.: A game-based framework for CTL counterexamples and 3-valued abstraction-refinement. *ACM Trans. Comput. Log.* **9**(1), 1 (2007). <https://doi.org/10.1145/1297658.1297659>
33. Winter, S., Zimmermann, M.: Prophecies all the way: Game-based model-checking for hyperqptl beyond $\forall^*\exists^*$. *arXiv:2504.08575* (2025) <https://doi.org/10.48550/ARXIV.2504.08575>