

The Complexity of Model Checking Higher Order Fixpoint Logic

Martin Lange¹ and Rafał Somla²

¹ Institut für Informatik, University of Munich, Germany
mlange@informatik.uni-muenchen.de

² IT department, Uppsala University, Sweden
rsomla@it.uu.se

Abstract. This paper analyses the computational complexity of the model checking problem for Higher Order Fixpoint Logic – the modal μ -calculus enriched with a typed λ -calculus. It is hard for every level of the elementary time/space hierarchy and in elementary time/space when restricted to formulas of bounded type order.

1 Introduction

Temporal logics are well-established tools for the specification of correctness properties and their verification in hard- and software design processes. One of the most famous temporal logics is Kozen’s modal μ -calculus \mathcal{L}_μ [16] which extends multi-modal logic with extremal fixpoint quantifiers. \mathcal{L}_μ subsumes many other temporal logics like CTL* [10], and with it CTL [9] and LTL [19], as well as PDL [12]. \mathcal{L}_μ can be embedded into infinitary modal logic. Hence, it does not distinguish between bisimilar models which is a desired property of a logic specifying program behaviour. The model checking problem for \mathcal{L}_μ on finite transition systems is in $\text{NP} \cap \text{co-NP}$, and the best deterministic algorithms are exponential in the alternation depth of the input formula only.

On the other hand, \mathcal{L}_μ is equi-expressive to the bisimulation-invariant fragment of Monadic Second Order Logic over trees or graphs [11, 15]. Since the latter is equi-expressive to Rabin tree automata, properties expressed by formulas of the modal μ -calculus are only regular. There are, however, many interesting correctness properties of programs that are not regular. Examples include *uniform inevitability* [8] which states that a certain event occurs globally at the same time in all possible runs of the system; counting properties like “at any point in a run of a protocol there have never been more *send*- than *receive*-actions”; formulas saying that an unbounded number of data does not lose its order during a transmission process; etc.

When program verification was introduced to computer science, programs as well as their correctness properties were mainly specified in temporal logics. Hence, verification meant to check formulas of the form $\varphi \rightarrow \psi$ for validity, or equally formulas of the form $\varphi \wedge \psi$ for satisfiability. An intrinsic problem for this approach and non-regular properties is undecidability. Note that the intersection problem for context-free languages is already undecidable [1].

One of the earliest attempts at verifying non-regular properties of programs was *Non-Regular PDL* [13] which enriches ordinary PDL by context-free programs. Non-Regular PDL is highly undecidable, hence, the logic did not receive much attention for program verification purposes. Its model checking problem, however, remains decidable on finite transition systems.

Although the theoretical complexity of the model checking problem is normally below that of its satisfiability problem, it often requires a lot more time or space to do model checking. This is simply because the input to a model checker is usually a lot bigger compared to that of a satisfiability checker. Hence, the feasibility of model checking is very much limited by the state space explosion problem: real-world examples result in huge transition systems that are very hard to model check simply because of their sheer size. However, in recent years various clever techniques have been invented that can cope with huge state spaces, starting with *local model checking*, and resulting in symbolic methods like *BDD-based* [4] or *bounded model checking* [6]. They are also a reason for the shift in importance from the satisfiability checking to the model checking problem for program verification.

More expressive power naturally comes with higher complexities. But with good model checking techniques at hand, verifying non-regular properties has become worthwhile again. This is for example reflected in the introduction of *Fixpoint Logic with Chop*, FLC, [18] which extends \mathcal{L}_μ with a sequential composition operator. It is capable of expressing many non-regular – and even non-context-free – properties, and its model checking problem on finite transition systems is decidable in deterministic exponential time [17].

Another logic capable of expressing non-regular properties is the *Modal Iteration Calculus*, MIC, [7] which extends \mathcal{L}_μ with inflationary fixpoint quantifiers. Similar to FLC, the satisfiability checking problem for MIC is undecidable but its model checking problem is decidable in deterministic polynomial space [7].

In order to achieve non-regular effects in FLC, the original \mathcal{L}_μ semantics is lifted to a function from sets of states to sets of states. This idea has been followed consequently in the introduction of *Higher Order Fixpoint Logic*, HFL, [23] which incorporates a typed λ -calculus into the modal μ -calculus. This gives it even more expressive power than FLC. HFL is, for example, capable of expressing *assume-guarantee-properties*. Still, HFL’s model checking problem on finite transition systems remains decidable. This has been stated in its introductory work [23] but no analysis of its computational complexity has been done so far.

Here we set out to answer the open question concerning the complexity of HFL’s model checking problem. We start by recalling the logic and giving a few examples in Section 2. Section 3 presents a reduction from the satisfiability problem for First Order Logic over finite words to HFL’s model checking problem. Consequently, the latter is hard for every level of the elementary time/space hierarchy. I.e. there is no model checking algorithm for HFL that runs in time given by a tower of exponentials whose height does not depend on the input formula. The reduction even yields a non-elementary lower complexity bound for transition systems of fixed size. This lower bound can also be obtained – as

(var) $\Gamma, X^v : \tau \vdash X : \tau$ if $v \in \{0, +\}$	(neg) $\frac{\Gamma^- \vdash \varphi : \mathbf{Prop}}{\Gamma \vdash \neg \varphi : \mathbf{Prop}}$
(or) $\frac{\Gamma \vdash \varphi : \mathbf{Prop} \quad \Gamma \vdash \psi : \mathbf{Prop}}{\Gamma \vdash \varphi \vee \psi : \mathbf{Prop}}$	(mod) $\frac{\Gamma \vdash \varphi : \mathbf{Prop}}{\Gamma \vdash \langle a \rangle \varphi : \mathbf{Prop}}$
(abs) $\frac{\Gamma, X^v : \sigma \vdash \varphi : \tau}{\Gamma \vdash \lambda(X^v : \sigma). \varphi : (\sigma^v \rightarrow \tau)}$	(fix) $\frac{\Gamma, X^+ : \tau \vdash \varphi : \tau}{\Gamma \vdash \mu(X : \tau). \varphi : \tau}$
(app^+) $\frac{\Gamma \vdash \varphi : (\sigma^+ \rightarrow \tau) \quad \Gamma \vdash \psi : \sigma}{\Gamma \vdash (\varphi \psi) : \tau}$	(app^-) $\frac{\Gamma \vdash \varphi : (\sigma^- \rightarrow \tau) \quad \Gamma^- \vdash \psi : \sigma}{\Gamma \vdash (\varphi \psi) : \tau}$
(app^0) $\frac{\Gamma \vdash \varphi : (\sigma^0 \rightarrow \tau) \quad \Gamma \vdash \psi : \sigma \quad \Gamma^- \vdash \psi : \sigma}{\Gamma \vdash (\varphi \psi) : \tau}$	$(prop)$ $\frac{}{\Gamma \vdash p : \mathbf{Prop}}$

Fig. 1. Type inference rules for HFL.

a referee pointed out – from a result by Statman who investigated complexity issues of λ -calculi without fixpoint operators that can be seen as fragments of HFL [21].

Section 4 shows that HFL’s model checking problem is indeed in elementary time when the order of the function types occurring in the input formulas is bound. We conclude in Section 5 with a discussion concerning the use of HFL as a specification formalism for program correctness.

2 Preliminaries

Let $\mathcal{P} = \{p, q, \dots\}$ be a set of atomic propositions, $\mathcal{A} = \{a, b, \dots\}$ be a finite set of action names, and $\mathcal{V} = \{X, Y, X_1, \dots\}$ a set of variable names. For simplicity, we fix \mathcal{P} , \mathcal{A} , and \mathcal{V} for the rest of the paper.

A $v \in \{-, +, 0\}$ is called a variance. The set of HFL-types is the smallest set containing the atomic type \mathbf{Prop} and being closed under function typing with variances, i.e. if σ and τ are HFL-types and v is a variance, then $\sigma^v \rightarrow \tau$ is an HFL-type..

Formulas of HFL are given by the following grammar:

$$\varphi ::= q \mid X \mid \neg \varphi \mid \varphi \vee \varphi \mid \langle a \rangle \varphi \mid \varphi \varphi \mid \lambda(X^v : \tau). \varphi \mid \mu(X : \tau). \varphi$$

where τ is an HFL-type and v is a variance. We use the standard abbreviations: $\mathbf{tt} := q \vee \neg q$ for some $q \in \mathcal{P}$, $\mathbf{ff} := \neg \mathbf{tt}$, $\varphi \wedge \psi := \neg(\neg \varphi \wedge \neg \psi)$, $[a]\psi := \neg \langle a \rangle \neg \psi$, and $\nu X. \varphi := \neg \mu X. \neg \varphi[\neg X/X]$. We will assume that any variable without an explicit type annotation is of the ground type \mathbf{Prop} . Also, if a variance is omitted it is implicitly assumed to be 0.

A sequence Γ of the form $X_1^{v_1} : \tau_1, \dots, X_n^{v_n} : \tau_n$ where X_i are variables, τ_i are types and v_i are variances is called a *context* (we assume all X_i are distinct).

$$\begin{aligned}
\mathcal{T}[\Gamma \vdash q : \mathbf{Prop}] \eta &= \{s \in \mathcal{S} \mid q \in L(s)\} \\
\mathcal{T}[\Gamma \vdash X : \tau] \eta &= \eta(X) \\
\mathcal{T}[\Gamma \vdash \neg \varphi : \mathbf{Prop}] \eta &= \mathcal{S} - \mathcal{T}[\Gamma^- \vdash \varphi : \mathbf{Prop}] \eta \\
\mathcal{T}[\Gamma \vdash \varphi \vee \psi : \mathbf{Prop}] \eta &= \mathcal{T}[\Gamma \vdash \varphi : \mathbf{Prop}] \eta \cup \mathcal{T}[\Gamma \vdash \psi : \mathbf{Prop}] \eta \\
\mathcal{T}[\Gamma \vdash \langle a \rangle \varphi : \mathbf{Prop}] \eta &= \{s \in \mathcal{S} \mid s \xrightarrow{a} t \text{ for some } t \in \mathcal{T}[\Gamma \vdash \varphi : \mathbf{Prop}] \eta\} \\
\mathcal{T}[\Gamma \vdash \lambda(X^v : \tau). \varphi : \tau^v \rightarrow \tau'] \eta &= F \in \mathcal{T}[\tau^v \rightarrow \tau'] \text{ s.t. } \forall d \in \mathcal{T}[\tau] \\
&\quad F(d) = \mathcal{T}[\Gamma, X^v : \tau \vdash \varphi : \tau'] \eta[X \mapsto d] \\
\mathcal{T}[\Gamma \vdash \varphi \psi : \tau'] \eta &= \mathcal{T}[\Gamma \vdash \varphi : \tau^v \rightarrow \tau'] \eta (\mathcal{T}[\Gamma' \vdash \psi : \tau] \eta) \\
\mathcal{T}[\Gamma \vdash \mu(X : \tau) \varphi : \tau] \eta &= \prod_{\mathcal{T}[\tau]} \{d \in \tau \mid \\
&\quad \mathcal{T}[\Gamma, X^+ : \tau \vdash \varphi : \tau] \eta[X \mapsto d] \leq_{\mathcal{T}[\tau]} d\}
\end{aligned}$$

Fig. 2. Semantics of HFL

An HFL-formula φ has type τ in context Γ if the statement $\Gamma \vdash \varphi : \tau$ can be inferred using the rules of Figure 1. We say that φ is *well-formed* if $\Gamma \vdash \varphi : \tau$ for some Γ and τ .

For a variance v , we define its complement v^- as $+$ if $v = -$, as $-$, if $v = +$, and 0 otherwise. For a context $\Gamma = X_1^{v_1} : \tau_1, \dots, X_n^{v_n} : \tau_n$, the complement Γ^- is defined as $X_1^{v_1^-} : \tau_1, \dots, X_n^{v_n^-} : \tau_n$.

A (labelled) transition system is a structure $\mathcal{T} = (\mathcal{S}, \{\xrightarrow{a}\}, L)$ where \mathcal{S} is a finite set of states, \xrightarrow{a} is a binary relation on states for each $a \in \mathcal{A}$, and $L : \mathcal{S} \rightarrow 2^{\mathcal{P}}$ is a labeling function denoting the set of propositional constants that are true in a state.

The semantics of a type w.r.t. a transition system \mathcal{T} is a complete lattice, inductively defined on the type as

$$\begin{aligned}
\mathcal{T}[\mathbf{Prop}] &= (2^{\mathcal{S}}, \subseteq) \\
\mathcal{T}[\sigma^v \rightarrow \tau] &= (\mathcal{T}[\sigma])^{v^-} \rightarrow \mathcal{T}[\tau]
\end{aligned}$$

Here, for two partial orders $\bar{\tau} = (\tau, \leq_{\tau})$ and $\bar{\sigma} = (\sigma, \leq_{\sigma})$, $\bar{\sigma} \rightarrow \bar{\tau}$ denotes the partial order of all monotone functions ordered pointwise, and, $\bar{\tau}^v$ denotes (τ, \leq_{τ}^v) . \leq_{τ}^+ is \leq_{τ} , $a \leq_{\tau}^- b$ iff $b \leq_{\tau} a$, and $\leq_{\tau}^0 = \leq_{\tau} \cap \leq_{\tau}^-$.

An environment η is a possibly partial map on the variable set \mathcal{V} . For a context $\Gamma = X_1^{v_1} : \tau_1, \dots, X_n^{v_n} : \tau_n$, we say that η respects Γ , denoted by $\eta \models \Gamma$, if $\eta(X_i) \in \mathcal{T}[\tau_i]$ for $i \in \{1, \dots, n\}$. We write $\eta[X \mapsto a]$ for the environment that maps X to a and otherwise agrees with η . If $\eta \models \Gamma$ and $a \in \mathcal{T}[\tau]$ then $\eta[X \mapsto a] \models \Gamma, X : \tau$, where X is a variable that does not appear in Γ .

For any well-typed term $\Gamma \vdash \varphi : \tau$ and environment $\eta \models \Gamma$, Figure 2 defines the semantics of φ inductively to be an element of $\mathcal{T}[\tau]$. In the clause for function application ($\varphi \psi$) the context Γ' is Γ if $v \in \{+, 0\}$, and is Γ^- if $v = -$.

We consider fragments of formulas that can be typed using types of restricted order only. Let

$$\text{ord}(\text{Prop}) := 1 \quad \text{ord}(\alpha \rightarrow \beta) := \max(\text{ord}(\alpha) + 1, \text{ord}(\beta))$$

and $\text{HFL}^k := \{ \varphi \in \text{HFL} \mid \varphi \text{ can be typed using types } \tau \text{ with } \text{ord}(\tau) \leq k \text{ for all subformulas only } \}$.

Example 1. The following HFL formula expresses the non-regular property “on any path the number of a ’s seen at any time never exceeds the number of b ’s seen so far.”

$$\nu X.[a]\mathbf{ff} \wedge [b](\nu(Z : \text{Prop} \rightarrow \text{Prop}).\lambda Y.([a]Y \wedge [b](Z(ZY))) X)$$

Note how function composition is used to “remember” in the argument to Z how many times a b -action has been seen along any path. Every b -action gives the potential to do another a -action later on which is remembered in the additional application of Z . And every a -action “uses up” one Z . If there have been as many a ’s as b ’s then the current state must be in the semantics of X again, hence, cannot do another a -action, etc.

Example 2. Let $\mathbf{2}_0^n := n$ and $\mathbf{2}_{m+1}^n := 2^{2^m}$. For any $n \in \mathbb{N}$, there is an HFL formula φ_n expressing the fact that there is a maximal path of length $\mathbf{2}_n^1$ (number of states on this path) through a transition system. It can be constructed using a typed version of the Church numeral 2. Let $P_0 = \text{Prop}$ and $P_{i+1} = P_i \rightarrow P_i$. For $i \geq 1$ define ψ_i of type P_{i+1} as follows

$$\psi_i := \lambda(F : P_i).\lambda(X : P_{i-1}).F(F X)$$

Then

$$\varphi_n := \psi_n \psi_{n-1} \dots \psi_1 (\lambda X.\langle - \rangle X) ([-]\mathbf{ff})$$

where $\langle - \rangle \varphi := \bigvee_{a \in \mathcal{A}} \langle a \rangle \varphi$ and $[-]\varphi := \bigwedge_{a \in \mathcal{A}} [a]\varphi$. Note that for any $n \in \mathbb{N}$, φ_n is of size polynomial in n . This indicates that HFL is able to express computations of Turing Machines of arbitrary elementary complexity. The next section shows that this is indeed the case.

3 The Lower Complexity Bound

Let Σ be a finite alphabet. Formulas of FO in negation normal form over words in Σ^* are given by the following grammar.

$$\varphi ::= x \leq y \mid x < y \mid P_a(x) \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \exists x.\varphi \mid \forall x.\varphi$$

where x, y are variables and $a \in \Sigma$.

A word $w \in \Sigma^*$ of length n is a function of type $\{0, \dots, n-1\} \rightarrow \Sigma$. Thus, $w(i)$ denotes the i -th letter of w . FO formulas are interpreted over words in the

usual way, written $w \models_{\eta} \varphi$ for a word w , a formula φ and an environment η evaluating the free variables in φ by positions in w .

Let Σ_0 and Π_0 be the set of all quantifier-free formulas of FO. Σ_{k+1} is the closure of $\Sigma_k \cup \Pi_k$ under the boolean operators and existential quantification. Similarly, Π_{k+1} is constructed from $\Sigma_k \cup \Pi_k$ using universal quantification instead.

Let $\text{DTime}(f(n))$ and $\text{DSpace}(f(n))$ be classes of languages that can be decided by a deterministic Turing Machine in time, resp. space, $f(n)$ where n measures the length of the input word to the machine. The k -th level of the elementary time/space hierarchy is

$$k\text{EXPTIME} = \bigcup_{c,d \in \mathbb{N}} \text{DTime}(2_k^{c n^d}), \quad k\text{ExpSpace} = \bigcup_{c,d \in \mathbb{N}} \text{DSpace}(2_k^{c n^d})$$

Furthermore, the elementary time/space hierarchy is

$$\text{ELTIME} := \bigcup_{k \in \mathbb{N}} k\text{EXPTIME} = \bigcup_{k \in \mathbb{N}} k\text{ExpSpace}$$

The standard translation of an FO formula into a finite automaton [3] and the encoding of space-bounded Turing Machine computations in FO [22, 20] yield the following results.

Theorem 1. *An FO sentence $\varphi \in \Sigma_{k+1}$ of length n has a model iff it has a model of length 2_k^n .*

Theorem 2. *For all $k \geq 1$: The satisfiability problem for FO formulas in Σ_{k+1} is hard for $k\text{ExpSpace}$.*

We will use these results to prove $k\text{ExpSpace}$ hardness of HFL. Our first step is to translate the problem of deciding whether a given binary word w of length 2_k^n is a model of an FO sentence φ into the HFL model checking problem.

Let us fix n and define \mathcal{T}_n to be a transition system with states $S_n = \{0, 1, \dots, n-1\}$, an empty labeling (we will not use propositional constants) and a cyclic next state relation $\rightarrow \subseteq S_n \times S_n$ given by $0 \rightarrow n-1$ and $i \rightarrow i-1$ for $i = 1, \dots, n-1$.

We represent the Boolean values *false* and *true* by the two elements of $\mathcal{B} = \{\emptyset, S_n\} \subset \mathcal{T}_n[\text{Prop}]$. In the sequel we will implicitly assume that all semantic interpretations are w.r.t \mathcal{T}_n and omit it in front of semantic brackets. Hence we can write $\llbracket \text{ff} \rrbracket$ and $\llbracket \text{tt} \rrbracket$ for the representations of *false* and *true*, respectively.

Let $\mathcal{N}_0^n = \{\{0\}, \{1\}, \dots, \{n-1\}\}$ and let $\mathcal{N}_{k+1}^n = \mathcal{N}_k^n \rightarrow \mathcal{B}$. Clearly, each \mathcal{N}_k^n has exactly 2_k^n elements which we will use to represent numbers in the range $0, \dots, 2_k^n - 1$. We have also $\mathcal{N}_k^n \subseteq \llbracket N_k \rrbracket$ where $N_0 = \text{Prop}$ and $N_{k+1} = N_k \rightarrow \text{Prop}$.

Note that the elements of $\mathcal{N}_{k+1}^n = \mathcal{N}_k^n \rightarrow \mathcal{B}$ can be equivalently viewed as predicates over \mathcal{N}_k^n , subsets of \mathcal{N}_k^n or binary words of length 2_k^n .

For a number $j \in \{0, 1, \dots, 2_k^n - 1\}$ let $\llbracket j \rrbracket^k$ be the element of \mathcal{N}_k^n representing j defined inductively by $\llbracket j \rrbracket^0 = \{j\}$ and

$$\llbracket j \rrbracket^k(x) := \begin{cases} \llbracket \text{tt} \rrbracket & \text{if } x = \llbracket i \rrbracket^{k-1} \text{ and } j_i = 1, \\ \llbracket \text{ff} \rrbracket & \text{o.w.} \end{cases}$$

where j_i is the i -th bit in the binary representation of j . For a binary word w of length 2_k^n let $\llbracket w \rrbracket^k := \llbracket \sum_i w_i \cdot 2^i \rrbracket^{k+1}$.

The possibility of a compact encoding of FO logic in HFL relies on the existence of polynomially sized HFL formulas describing basic operations on numbers represented as elements of \mathcal{N}_k^n . We define

$$\text{inc}_k : N_k \rightarrow N_k \quad , \quad \text{eq}_k : N_k \rightarrow N_k \rightarrow \text{Prop} \quad , \quad \text{search}_k : (N_k \rightarrow \text{Prop}) \rightarrow \text{Prop}$$

adhering to the following specifications.

$$\llbracket \text{inc}_k \rrbracket \llbracket j \rrbracket^k = \llbracket j + 1 \rrbracket^k \quad , \quad \llbracket \text{eq}_k \rrbracket \llbracket j \rrbracket^k \llbracket i \rrbracket^k = \begin{cases} \llbracket \text{tt} \rrbracket & , \text{ iff } j = i \\ \llbracket \text{ff} \rrbracket & , \text{ o.w.} \end{cases}$$

For a predicate $p \in \mathcal{N}_{k+1}^n$,

$$\llbracket \text{search}_k \rrbracket p = \begin{cases} \llbracket \text{tt} \rrbracket & \text{ iff exists } x \in \mathcal{N}_k^n \text{ s.t. } p(x) = \llbracket \text{tt} \rrbracket \\ \llbracket \text{ff} \rrbracket & \text{ o.w.} \end{cases}$$

The search function search_k can be implemented using inc_k and recursion. A helper function $\text{search}'_k P x$ applies P to the successive numbers, starting from x , taking the union of the results.

$$\begin{aligned} \text{search}_k &:= \lambda(P : N_k \rightarrow \text{Prop}). \text{search}'_k P \perp^k \\ \text{search}'_k &:= \lambda(P : N_k \rightarrow \text{Prop}). \mu(Z : N_k \rightarrow \text{Prop}). \\ &\quad \lambda(X : N_k). (P X) \vee (Z (\text{inc}_k X)). \end{aligned}$$

Formula $\perp^k : N_k$ represents 0 and is defined as

$$\perp^0 := [-] \text{ff}, \quad \perp^k := \lambda(X : N_{k-1}). \text{ff} \quad \text{for } k > 0.$$

Functions eq_k and inc_k are defined by induction on k . For $k = 0$ we set

$$\text{eq}_0 := \lambda X. \lambda Y. (X \leftrightarrow Y) \quad \text{inc}_0 := \lambda X. \langle - \rangle X$$

For $k > 0$, function eq_k is implemented by searching for an argument at which two number representations differ:

$$\text{eq}_k := \lambda(X : N_k). \lambda(Y : N_k). \neg(\text{search}_{k-1} \lambda(I : N_{k-1}). \neg(X I \leftrightarrow Y I))$$

Function inc_k is the usual incrementation of a number in binary representation. The helper function $\text{inc}'_k x i$ adds one to the i -th bit of n and possibly the following bits if the carry-over occurs.

$$\text{inc}_k := \lambda(X : N_k). \text{inc}'_k X \perp^{k-1}$$

The value of $\llbracket \text{inc}'_k \rrbracket \llbracket x \rrbracket^k \llbracket i \rrbracket^{k-1}$ is a function which for each j returns the j -th bit of $x + 2^i$ (encoded as $\llbracket \text{tt} \rrbracket$ or $\llbracket \text{ff} \rrbracket$). For $j = i$ the corresponding bit is $\neg x_i$. If there

is no carry-over ($x_i = 0$) then the remaining bits are unchanged. Otherwise the remaining bits are the same as in $x + 2^{i+1}$.

$$\begin{aligned} \text{inc}'_k &:= \lambda(X : N_k). \mu(Z : N_{k-1} \rightarrow N_k). \lambda(I : N_{k-1}). \\ &\quad \lambda(J : N_{k-1}). \text{if } (\text{eq}_{k-1} J I) \\ &\quad \quad (\neg(X I)) \\ &\quad \quad (\text{if } \neg(X I) (X J) (Z (\text{inc}_{k-1} I) J)) \end{aligned}$$

where $\text{if} := \lambda P. \lambda Q. \lambda R. (P \wedge Q) \vee (\neg P \wedge R)$.

Note that the lengths of inc_k , eq_k and search_k as strings can be exponential in k . However, the number of their subformulas is only polynomial in k .

Lemma 1. *For any $k \geq 0$, any $i \in \{0, \dots, 2_k^n - 1\}$, and any $p \in \mathcal{N}_k^n \rightarrow \mathcal{B}$ we have: $\llbracket \text{search}'_k \rrbracket p \llbracket i \rrbracket^k = \llbracket \text{tt} \rrbracket$ iff $p(\llbracket j \rrbracket^k) = \llbracket \text{tt} \rrbracket$ for some $i \leq j < 2_k^n$.*

Proof. Simply because $\llbracket \text{search}'_k \rrbracket p \llbracket i \rrbracket^k \equiv \bigcup_{j=i}^{2_k^n - 1} p(\llbracket j \rrbracket^k)$. □

Lemma 2. *For any k , $\text{eq}_k, \text{inc}_k \in \text{HFL}^{k+1}$ and $\text{search}_k, \text{search}'_k \in \text{HFL}^{k+2}$.*

Proof. It is easy to see that $\text{ord}(N_k) = k$. The type $N_k \rightarrow N_k \rightarrow \text{Prop}$ of eq_k has order $k + 1$ and, by induction hypothesis, eq_k doesn't contain any subformulas of a higher order. Similar, the type of inc_k which is $N_k \rightarrow N_k$ has order $k + 1$, the same as the order of the type $N_k \rightarrow N_{k-1} \rightarrow N_k$ of inc'_k . By induction, neither inc_k nor inc'_k contain a subformula of a higher order. Finally, inspecting the definitions of search_k and search'_k we can see that they contain no subformula whose type has order higher than $k + 2$. □

Let φ be an FO sentence. For given k we translate φ into an HFL^{k+2} formula $\text{tr}_k(\varphi) : N_{k+1} \rightarrow \text{Prop}$ s.t. for any word w of length 2_k^n , w is a model of φ iff $(\mathcal{T}_n \llbracket \text{tr}_k(\varphi) \rrbracket) (\mathcal{T}_n \llbracket w \rrbracket^k) = \mathcal{T}_n \llbracket \text{tt} \rrbracket$.

$$\begin{aligned} \text{tr}_k(x \leq y) &:= \lambda(w : N_{k+1}). \text{search}'_k x (\lambda(u : N_k). \text{eq}_k u y) \\ \text{tr}_k(P_0(x)) &:= \lambda(w : N_{k+1}). \neg(w x) \\ \text{tr}_k(P_1(x)) &:= \lambda(w : N_{k+1}). (w x) \\ \text{tr}_k(\exists x. \varphi) &:= \lambda(w : N_{k+1}). \text{search}_k (\lambda(x : N_k). \text{tr}_k(\varphi) w) \\ \text{tr}_k(\neg \varphi) &:= \lambda(w : N_{k+1}). \neg(\text{tr}_k(\varphi) w) \\ \text{tr}_k(\varphi \vee \psi) &:= \lambda(w : N_{k+1}). (\text{tr}_k(\varphi) w) \vee (\text{tr}_k(\psi) w) \end{aligned}$$

Note that free variables of φ become free variables of type N_k in $\text{tr}_k(\varphi)$.

Lemma 3. *For any FO sentence φ the translation $\text{tr}_k(\varphi)$ is a predicate. That is, $\mathcal{T}_n \llbracket \text{tr}_k(\varphi) \rrbracket$ is an element of \mathcal{N}_{k+2}^n – a function which when applied to an argument from \mathcal{N}_{k+1}^n returns either $\mathcal{T}_n \llbracket \text{tt} \rrbracket$ or $\mathcal{T}_n \llbracket \text{ff} \rrbracket$.*

Proof. This follows from the fact that, by their specifications, search_k and search'_k are predicates. Hence $\text{tr}_k(\varphi)$ is a predicate as a Boolean combination of predicates.

Lemma 4. For any k and φ , $tr_k(\varphi) \in \text{HFL}^{k+2}$.

Proof. Follows immediately from Lemma 2.

Lemma 5. Let φ be an FO formula with variables x_1, \dots, x_l . For any word w of length 2_k^n and FO-environment η we have

$$w \models_{\eta} \varphi \quad \text{iff} \quad (\mathcal{T}_n \llbracket tr_k(\varphi) \rrbracket \rho) (\mathcal{T}_n \llbracket w \rrbracket^k) = \mathcal{T}_n \llbracket \mathbf{tt} \rrbracket$$

where ρ is an HFL environment given by $\rho(x_i) = \mathcal{T}_n \llbracket \eta(x_i) \rrbracket^k$.

Proof. By induction on the structure of φ . We fix n and k and as before omit \mathcal{T}_n in front of semantic brackets.

Case $\varphi = x_i \leq x_j$: Then $\llbracket tr_k(\varphi) \rrbracket \rho = \llbracket \text{search}'_k \rrbracket \llbracket a \rrbracket^k p$ where $a = \eta(x_i)$, $b = \eta(x_j)$ and predicate p is given by $p(x) = \llbracket \text{eq}_k \rrbracket x \llbracket b \rrbracket^k$. We have $w \models_{\eta} \varphi$ iff $a \leq b$ iff exists $a \leq c < 2_k^n$ s.t. $p(\llbracket c \rrbracket^k) = \llbracket \mathbf{tt} \rrbracket$ iff $\llbracket \text{search}'_k \rrbracket \llbracket a \rrbracket^k p = \llbracket \mathbf{tt} \rrbracket$, by Lemma 1.

Case $\varphi = \exists x. \psi$: Then $\llbracket tr_k(\varphi) \rrbracket \rho = \llbracket \text{search}_k \rrbracket p$ where $p(\llbracket i \rrbracket^k) = \llbracket tr_k(\psi) \rrbracket \rho[x \mapsto \llbracket i \rrbracket^k] \llbracket w \rrbracket^k$. By the specification of search_k and induction hypothesis we have $(\llbracket tr_k(\varphi) \rrbracket \rho) \llbracket w \rrbracket^k = \llbracket \mathbf{tt} \rrbracket$ iff $p(\llbracket i \rrbracket^k) = \llbracket \mathbf{tt} \rrbracket$ for some i iff $w \models_{\eta[x \mapsto i]} \psi$ iff $w \models_{\eta} \varphi$.

Case $\varphi = P_0(x_i)$: Then $w \models_{\eta} \varphi$ iff $w_{\eta(x_i)} = 0$ iff $\llbracket w \rrbracket^k \llbracket \eta(x_i) \rrbracket^k = \llbracket \mathbf{ff} \rrbracket$ iff $(\llbracket tr_k(\varphi) \rrbracket \rho) \llbracket w \rrbracket^k = \llbracket \mathbf{tt} \rrbracket$.

The case of $\varphi = P_1(x_i)$ is analogous. The cases of $\varphi = \neg \psi$ and $\varphi = \psi_1 \vee \psi_2$ follow immediately from the induction hypothesis. \square

The model checking problem for HFL^k is the following:

Given an HFL^k sentence $\varphi : \text{Prop}$, a transition system \mathcal{T} and a set of states A decide whether or not $\mathcal{T} \llbracket \varphi \rrbracket = A$.

Lemma 6. The satisfiability problem for FO^k is polynomially reducible to the model checking problem for HFL^{k+2} .

Proof. First note that we can restrict ourselves to the satisfiability problem for FO^k over the binary alphabet $\Sigma = \{0, 1\}$ because any other alphabet can be encoded by Σ at a logarithmic expense only.

Given an FO^k formula φ of length n we can construct in polynomial time and space an instance of the HFL^{k+2} model checking problem consisting of a formula $\varphi' = \text{search}_{k+1} tr_k(\varphi)$, transition system \mathcal{T}_n and the set $\mathcal{T}_n \llbracket \mathbf{tt} \rrbracket = \{0, 1, \dots, n-1\}$. Note that $\varphi' \in \text{HFL}^{k+2}$ by Lemmas 2 and 4. From Lemmas 1, 3 and 5 it follows that $\mathcal{T}_n \llbracket \varphi' \rrbracket = \mathcal{T}_n \llbracket \mathbf{tt} \rrbracket$ iff φ has a model of size 2_k^n which, by Theorem 1, is equivalent to φ having a model. \square

Theorem 2 together with the reduction of Lemma 6 yields the following result.

Theorem 3. The model checking problem for HFL^{k+3} is hard for $k\text{ExpSpace}$ under polynomial time reductions.

Corollary 1. The model checking problem for HFL is not in ELTIME .

Note that the reduction only uses modal formulas $\langle - \rangle \varphi$ and $[-] \varphi$ because 2_k^n is an upper bound on a minimal model for an FO sentence in Σ_{k+1} of length n . However, $2_k^n = 2_{k+\log^* n}^1$. This enables us to use modality-free formulas and transition systems of fixed size 1 in the reduction. The price to pay is that, in order to achieve $k\text{ExpSpace}$ -hardness, one needs formulas with unrestricted types. At least it shows (again) that HFL model checking is not in ELTIME for fixed and arbitrarily small transition systems already.

4 The Upper Complexity Bound

In the following we will identify a type τ and its underlying complete lattice induced by a transition system \mathcal{T} with state set \mathcal{S} . In order to simplify notation we fix \mathcal{T} for the remainder of this section.

Suppose $|\mathcal{S}| = n$ for some $n \in \mathbb{N}$. We define the order $\text{ord}(\tau)$ of a type τ , its size $\#(\tau)$, as well as $\text{rp}(\tau)$ – a space measure for the representation of one of its elements.

$$\begin{aligned} \#(\mathbf{Prop}) &:= 2^n & \#(\alpha \rightarrow \beta) &:= \#(\beta)^{\#(\alpha)} \\ \text{rp}(\mathbf{Prop}) &:= n & \text{rp}(\alpha \rightarrow \beta) &:= \#(\alpha) \cdot \text{rp}(\beta) \end{aligned}$$

Lemma 7. (a) For any HFL type τ over a transition system of size n there are only $\#(\tau)$ many different elements of τ .

(b) An element x of any HFL type τ can be represented using space $\text{rp}(\tau)$.

(c) $\#(\tau) \leq 2_{\text{ord}(\tau)+1}^n$ and $\text{rp}(\tau) \leq 2_{\text{ord}(\tau)}^n$.

Proof. Both can easily be proved by induction on the structure of τ . (a) Note that \mathbf{Prop} represents the power set of a state space. Hence, $\#(\mathbf{Prop}) = 2^n$. Furthermore, for two types τ_1 and τ_2 there are $\#(\tau_2)^{\#(\tau_1)} = \#(\tau_1 \rightarrow \tau_2)$ many functions from τ_1 to τ_2 .

(b) A subset of a state space of size n can be represented for example by a boolean array of size n . In order to completely represent a function of type $\tau_1 \rightarrow \tau_2$ one needs to store for each element of τ_1 a value from τ_2 . Hence, that space is bounded by $\#(\tau_1) \cdot \text{rp}(\tau_2) = \text{rp}(\tau_1 \rightarrow \tau_2)$.

(c) The first inequality trivially holds for $\tau = \mathbf{Prop}$. Assuming that it holds for types α and β with $\text{ord}(\alpha) = k$ and $\text{ord}(\beta) = l$ we have

$$\#(\alpha \rightarrow \beta) = \#(\beta)^{\#(\alpha)} = 2^{2^{l-1} \cdot 2_k^n} \leq 2^{(2_{m-1}^n)^2} \leq 2^{2_m^n} = 2_{m+1}^n$$

where $m = \max(k+1, l) = \text{ord}(\alpha \rightarrow \beta)$. The other inequality follows from the fact that $\text{rp}(\tau) = \log_2 \#(\tau)$ for any τ . \square

Theorem 4. For all $k \in \mathbb{N}$, the model checking problem for HFL^k and finite transition systems is in $(k+1)\text{EXPTIME}$.

Proof. For a finite transition system $\mathcal{T} = (\mathcal{S}, \{\overset{a}{\rightarrow} \mid a \in \mathcal{A}\}, L)$ with $|\mathcal{S}| = n$ and an HFL formula φ of type τ , we describe an alternating procedure for finding

the denotation of φ . The existential player \exists proposes an element of τ as $\llbracket\varphi\rrbracket$ and the universal player \forall challenges her choice. The game proceeds along the structure of φ in the following way.

If $\varphi = \lambda(X : \sigma).\psi$ then \forall chooses an entry in the table written by \exists as a value of φ . The entry consists of the value of X and the corresponding value of φ . \forall can invoke the verification protocol to check that this is the correct value of $\psi(X)$ when X has the value given by the entry.

For $\varphi = \psi_1 \psi_2$, \forall can ask \exists to write a table of ψ_1 . The first entry in the table should contain a previously guessed value of φ as the function result. Now \forall can either verify that the whole table is correct or that the function argument in the first entry corresponds to ψ_2 .

To verify a guess x of the value of $\varphi = \mu(X : \tau).\psi$ first \exists writes a table of a function $\lambda(X : \tau).\psi$ with the first entry (x, x) . Then \forall can either verify correctness of the whole table or indicate another entry (y, y) with y smaller than x . When φ is a variable then it is enough to compare the value chosen by \exists with the one stored in the variable environment.

In all other cases φ is of type **Prop** and its value is a bit vector of length n . Correctness of Boolean operations can be easily verified by \forall using additional $2n$ space for storing the values of the operands.

Clearly the space needed to perform the above procedure is bounded by the maximal $rp(\sigma)$ where σ is a type of a subformula of φ or $3n$, whichever is bigger. Hence, by Lemma 7 the model checking problem is in alternating 2_k^n space thus in $(k + 1)\text{EXPTIME}$ [5]. \square

5 Conclusion

We have shown that the model checking problem for HFL is hard for every $k\text{EXPTIME}$ and consequently of non-elementary complexity. Theoretically, it is tempting to dismiss HFL as a specification formalism that can practically be used. But the same argument would also rule out any practical implementation of a satisfiability checker for Monadic Second Order Logic over words or trees (MSO) since this problem has non-elementary complexity, too. However, the verification tool MONA [14] shows that in many practical cases satisfiability of MSO formulas can be checked efficiently. This is mainly because of the use of efficiently manipulable data structures like BDDs [2], and the fact that formulas used in practical cases do not coincide with those that witness the high complexity.

Thus, the theoretical complexity bounds proved in this paper need to be seen as a high alert warning sign for someone building a model checking tool based on HFL. This will certainly require the use of efficient data structures as well as other clever optimisations. However, only such an implementation will be able to judge the use of HFL as a specification formalism properly.

The exact reason for the high complexity also remains to be analysed from a theoretical point of view. This is begun in Corollary ?? showing that it is only the formula part of the input. Furthermore, the reduction in Section 3 hints that in order to reach high levels in ELTIME , one needs formulas of type $(\dots((\mathcal{P} \rightarrow \mathcal{P}) \rightarrow \mathcal{P}) \rightarrow \dots) \rightarrow \mathcal{P}$ which probably do not occur for natural

correctness properties anyway. It remains to be seen whether this is reflected on the theoretical side, i.e. whether the model checking problem for HFL restricted to formulas of order k type is complete for k' EXPTIME for some k' with a constant distance to k .

Acknowledgments We would like to thank Martin Leucker for discussing the topic at hand.

References

1. Y. Bar-Hillel, M. Perles, and E. Shamir. On formal properties of simple phrase structure grammars. *Zeitschrift für Phonologie, Sprachwissenschaft und Kommunikationsforschung*, 14:113–124, 1961.
2. R. E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers*, 35(8):677–691, August 1986.
3. J. R. Büchi. On a decision method in restricted second order arithmetic. In *Proc. Congress on Logic, Method, and Philosophy of Science*, pages 1–12, Stanford, CA, USA, 1962. Stanford University Press.
4. J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang. Symbolic model checking: 10^{20} states and beyond. *Information and Computation*, 98(2):142–170, June 1992.
5. A. K. Chandra, D. C. Kozen, and L. J. Stockmeyer. Alternation. *Journal of the ACM*, 28(1):114–133, January 1981.
6. E. M. Clarke, A. Biere, R. Raimi, and Y. Zhu. Bounded model checking using satisfiability solving. *Formal Methods in System Design*, 19(1):7–34, 2001.
7. A. Dawar, E. Grädel, and S. Kreutzer. Inflationary fixed points in modal logic. In L. Fribourg, editor, *Proc. 15th Workshop on Computer Science Logic, CSL'01*, LNCS, pages 277–291, Paris, France, September 2001. Springer.
8. E. A. Emerson. Uniform inevitability is tree automaton ineffable. *Information Processing Letters*, 24(2):77–79, January 1987.
9. E. A. Emerson and J. Y. Halpern. Decision procedures and expressiveness in the temporal logic of branching time. *Journal of Computer and System Sciences*, 30:1–24, 1985.
10. E. A. Emerson and J. Y. Halpern. “Sometimes” and “not never” revisited: On branching versus linear time temporal logic. *Journal of the ACM*, 33(1):151–178, January 1986.
11. E. A. Emerson and C. S. Jutla. Tree automata, μ -calculus and determinacy. In *Proc. 32nd Symp. on Foundations of Computer Science*, pages 368–377, San Juan, Puerto Rico, October 1991. IEEE.
12. M. J. Fischer and R. E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18(2):194–211, April 1979.
13. D. Harel, A. Pnueli, and J. Stavi. Propositional dynamic logic of nonregular programs. *Journal of Computer and System Sciences*, 26(2):222–243, April 1983.
14. J. G. Henriksen, J. Jensen, M. Joergensen, and N. Klarlund. MONA: Monadic second-order logic in practice. *LNCS*, 1019:89–110, 1995.
15. D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ -calculus with respect to monadic second order logic. In U. Montanari and V. Sassone, editors, *Proc. 7th Conf. on Concurrency Theory, CONCUR'96*, volume 1119 of *LNCS*, pages 263–277, Pisa, Italy, August 1996. Springer.

16. D. Kozen. Results on the propositional μ -calculus. *TCS*, 27:333–354, December 1983.
17. M. Lange and C. Stirling. Model checking fixed point logic with chop. In M. Nielsen and U. H. Engberg, editors, *Proc. 5th Conf. on Foundations of Software Science and Computation Structures, FOSSACS'02*, volume 2303 of *LNCS*, pages 250–263, Grenoble, France, April 2002. Springer.
18. M. Müller-Olm. A modal fixpoint logic with chop. In C. Meinel and S. Tison, editors, *Proc. 16th Symp. on Theoretical Aspects of Computer Science, STACS'99*, volume 1563 of *LNCS*, pages 510–520, Trier, Germany, 1999. Springer.
19. A. Pnueli. The temporal logic of programs. In *Proc. 18th Symp. on Foundations of Computer Science, FOCS'77*, pages 46–57, Providence, RI, USA, October 1977. IEEE.
20. K. Reinhardt. The complexity of translating logic to finite automata. In E. Grädel, W. Thomas, and Th. Wilke, editors, *Automata, Languages, and Infinite Games*, volume 2500 of *LNCS*, pages 231–238. Springer, 2002.
21. R. Statman. The typed λ -calculus is not elementary recursive. *Theoretical Computer Science*, 9:73–81, 1979.
22. L. Stockmeyer. *The Computational Complexity of Word Problems*. PhD thesis, MIT, 1974.
23. M. Viswanathan and R. Viswanathan. A higher order modal fixed point logic. In Ph. Gardner and N. Yoshida, editors, *Proc. 15th Int. Conf. on Concurrency Theory, CONCUR'04*, volume 3170 of *LNCS*, pages 512–528, London, UK, 2004. Springer.