

Capturing Bisimulation-Invariant Exponential-Time Complexity Classes

Florian Bruse

University of Kassel
Kassel, Germany

florian.bruse@uni-kassel.de

David Kronenberger

University of Kassel
Kassel, Germany

Martin Lange

University of Kassel
Kassel, Germany

martin.lange@uni-kassel.de

Otto's Theorem characterises the bisimulation-invariant PTIME queries over graphs as exactly those that can be formulated in the polyadic μ -calculus, hinging on the Immerman-Vardi Theorem which characterises PTIME (over ordered structures) by First-Order Logic with least fixpoints. This connection has been extended to characterise bisimulation-invariant EXPTIME by an extension of the polyadic μ -calculus with functions on predicates, making use of Immerman's characterisation of EXPTIME by Second-Order Logic with least fixpoints.

In this paper we show that the bisimulation-invariant versions of all classes in the exponential time hierarchy have logical counterparts which arise as extensions of the polyadic μ -calculus by higher-order functions. This makes use of the characterisation of k -EXPTIME by Higher-Order Logic (of order $k + 1$) with least fixpoints, due to Freire and Martins.

1 Introduction

Descriptive complexity theory aims at characterising complexity classes – usually defined via computational resources like time or space consumption – by means of *logical* resources. Its central notion is that of a complexity class \mathcal{C} being *captured* by a logic \mathcal{L} in the sense that the properties which can be checked in complexity \mathcal{C} are exactly those that can be defined in \mathcal{L} . This provides a characterisation of computational complexity that is independent of a machine model. Instead, computational difficulty is characterised by the need for particular logical resources like k -order quantifiers or fixpoints of a particular type. It is widely believed that this provides a more promising line of attack for notoriously difficult problems of separating complexity classes; at least it makes machinery that is traditionally used for measuring the expressive power of logics available for such tasks.

Here we adopt terminology from database theory as this is traditionally close to descriptive complexity, and speak of *queries* being *answered* instead of problems being solved or languages being decided, different names for the same thing.

Ever since Fagin's seminal work showing that the complexity class NP is captured by Existential Second-Order Logic [5], descriptive complexity has provided logical characterisations of many standard complexity classes. The Abiteboul-Vianu Theorem for instance states that PSPACE is captured by First-Order Logic with Partial Fixpoint Operators [1]. The proofs of these results rely on the existence of a total order on the structure at hand. This is not a restriction for characterisations of complexity classes including and above NP as the resources available there are sufficient to construct such an order.

There is no known way to define or construct such a total order in deterministic polynomial time which is a major obstacle for capturing the complexity class P. There is some belief that such a logic should exist, possibly in the form of first-order logic with additional operators like fixpoints, counting and others, cf. [7]. This is grounded in the characterisation of the complexity class of *Ordered Polynomial*

Time – i.e. those queries that can be answered in polynomial time on structures that are equipped with a total order – by First-Order Logic with Least Fixpoints [20, 8].

An interesting result was then found by Otto who considered another restriction of the class P, namely that of *bisimulation-invariant* queries (on graphs, naturally). He showed that this class, denoted P/\sim is captured by the polyadic μ -calculus \mathcal{L}_μ^ω [17], a generalisation of the well-known modal μ -calculus to interpretations of formulas not in states but in tuples of states of fixed arity [2]. This is particularly interesting as it shifts the borderline at which the availability of an order becomes critical, from between P and NP to between P/\sim and $NLOGSPACE/\sim$.

This opens up the question of further capturing results of bisimulation-invariant complexity classes by (modal) logics. Indeed, characterisations have been found for $EXPTIME/\sim$ and $PSPACE/\sim$ in terms of an extension of \mathcal{L}_μ^ω by first-order functions from predicates to predicates [15]. The logic capturing $EXPTIME/\sim$ is coined PHFL¹ – Polyadic Higher-Order Fixpoint Logic of order 1. The higher-order extension is borrowed from HFL which extends the modal μ -calculus with a simply typed λ -calculus [21]. A syntactical restriction called *tail-recursiveness* [4] has been identified that captures $PSPACE/\sim$ [15], and when applying this restriction to \mathcal{L}_μ^ω or, likewise, PHFL⁰, one captures $NLOGSPACE/\sim$ (with the help of a particular partial order only, though) [15].

Considering bisimulation-invariant complexity classes above NP/\sim does not serve the same purpose as it does for smaller ones as one of the key motivations for moving to the bisimulation-invariant world is to avoid the need for a total order. On the other hand, for two complexity classes $\mathcal{C}, \mathcal{C}'$ that have complete and bisimulation-invariant problems we have $\mathcal{C} \neq \mathcal{C}'$ iff $\mathcal{C}/\sim \neq \mathcal{C}'/\sim$. Most of the standard complexity classes possess such problems, for example (1-letter) NFA universality for NP, resp. PSPACE, unbounded tree automaton intersection for EXPTIME, etc. Moreover, separating bisimulation-invariant classes may be easier due to their close connection to modal logics where separating their expressiveness is routinely done (not for the relatively complex higher-order modal fixpoint logics mentioned here, though).

In this paper we extend the characterisation of bisimulation-invariant time complexity classes to all the levels of the exponential time hierarchy. We show that k -EXPTIME/ \sim is captured by PHFL^k, the polyadic version of the higher-order extension of the modal μ -calculus with functions up to type order k . We remark that a similar characterisation of the space complexity classes k -EXSPACE is also possible [13] but needs to be omitted for lack of space and is therefore left for a future publication.

The paper is organised as follows. In Sect. 2 we recall the necessary preliminaries, mainly about the logics studied here. In Sect. 3 we provide the easy half of the capturing result by putting together known results and constructions which witness that any PHFL^k query can be answered in k -fold exponential time. In Sect. 4 we prepare for the more difficult and other half of the capturing result. We rely on a logical characterisation of k -EXPTIME in terms of Higher-Order Predicate Logic with Fixpoints [6], and a key step in expressing such (bisimulation-invariant) queries in Higher-Order Modal Fixpoint is the modelling of higher-order quantification using an enumeration technique. In Sect. 5 we put this to use for showing that any k -EXPTIME/ \sim -query is definable in PHFL^k. In Sect. 6 we conclude with remarks on further work etc.

2 Preliminaries

Let \mathbf{P} and \mathbf{A} be finite sets of *propositions*, resp. *actions*. A labelled transition system (LTS) is a tuple $T = (S, \{\xrightarrow{a}\}_{a \in \mathbf{A}}, \ell)$ where S is a nonempty set of *states*, $\xrightarrow{a} \subseteq S \times S$ is a *transition relation* for each action $a \in \mathbf{A}$, and $\ell: S \rightarrow 2^{\mathbf{P}}$ labels the states with those propositions that are true in them. We write $s \xrightarrow{a} t$ instead of $(s, t) \in \xrightarrow{a}$.

For $d \geq 1$, a d -pointed LTS is a pair $T, (s_1, \dots, s_d)$ of an LTS and a d -tuple of states in it. It is also simply called *pointed* when d is clear from the context. Note that in the case of $d = 1$, this coincides with the usual notion of a pointed LTS. Given some tuple $\bar{s} = (s_1, \dots, s_d)$ and $i \leq d$, we write $\bar{s}[t/i]$ to denote the tuple $(s_1, \dots, s_{i-1}, t, s_{i+1}, \dots, s_d)$.

2.1 Polyadic Higher-Order Fixpoint Logic

We assume familiarity with the modal μ -calculus \mathcal{L}_μ [12]. Polyadic Higher-Order Fixpoint Logic PHFL [15] extends \mathcal{L}_μ in several ways: (I) HFL [21] adds to it a simply typed λ -calculus; (II) the polyadic $\mathcal{L}_\mu, \mathcal{L}_\mu^\omega$ [2, 17] is obtained by lifting the interpretation of formulas in states to tuples of states of fixed arity. Now PHFL merges both extensions. Its introduction requires a few technicalities.

Types. Types are used to govern the syntax of PHFL, especially those parts that denote functions. They are derived from the grammar

$$\tau ::= \bullet \mid \tau^v \rightarrow \tau$$

where \bullet is a ground type for propositions, $v \in \{+, -, 0\}$ are *variances* denoting whether a function is monotonically increasing, monotonically decreasing, or constant in its argument. Variances are only needed to check well-typedness; we often do not display them for the sake of readability.

The *order* ord of a type is defined via $ord(\bullet) = 0$ and $ord(\tau_1 \rightarrow \tau_2) = \max(ord(\tau_1) + 1, ord(\tau_2))$. Types associate to the right whence every type is of the form $\tau_1 \rightarrow \dots \rightarrow \tau_n \rightarrow \bullet$.

Given some LTS $T = (S, \{\xrightarrow{a}\}_{a \in \mathbf{A}}, \ell)$ and some $0 < d \in \mathbb{N}$, the semantics $\llbracket \tau \rrbracket^T$ of a type τ is a complete lattice, defined inductively as follows:

$$\llbracket \bullet \rrbracket^T = (2^{S^d}, \subseteq) \qquad \llbracket \tau_2^v \rightarrow \tau_1 \rrbracket^T = (\llbracket \tau_2 \rrbracket^T \rightarrow \llbracket \tau_1 \rrbracket^T, \sqsubseteq_{\tau_2^v \rightarrow \tau_1})$$

where we tacitly identify a lattice with its domain and where the order $\sqsubseteq_{\tau_2^v \rightarrow \tau_1}$ is given by

$$f \sqsubseteq_{\tau_2^v \rightarrow \tau_1} g \iff \text{for all } x \in \llbracket \tau_2 \rrbracket^T \text{ we have } \begin{cases} f(x) \sqsubseteq_{\tau_1} g(x), & \text{if } v = + \\ g(x) \sqsubseteq_{\tau_1} f(x), & \text{if } v = - \\ f(x) = g(x), & \text{if } v = 0. \end{cases}$$

Hence, the semantics of \bullet is the powerset lattice over S^d and $\tau_2^v \rightarrow \tau_1$ is the lattice of all monotonically increasing, monotonically decreasing or constant functions from $\llbracket \tau_2 \rrbracket^T$ to $\llbracket \tau_1 \rrbracket^T$, depending on v . Such a set together with the above point-wise order forms a complete lattice if $\llbracket \tau_1 \rrbracket^T$ is one. Hence, monotone functions in such a lattice always have a least and greatest fixpoint due to the Knaster-Tarski-Theorem. We write $\bigsqcup_{\tau_2^v \rightarrow \tau_1}$ to denote the join operator in this lattice.

Note that, technically, $\llbracket \tau_1 \rrbracket^T$ is also parameterised in d . However, since d is usually clear from context, we do not display it to avoid clutter.

Syntax. Let \mathbf{P} and \mathbf{A} be as above and let \mathbf{F} and \mathbf{L} be finite sets of *fixpoint variables*, resp. *lambda variables*. We use upper case letters X, Y, \dots for the former and lower case letters x, y, f, g for the latter.

Let $d \geq 1$. By $[d]$ we denote the set $1, \dots, d$. The set of – potentially non-well-formed – formulas of the d -adic fragment of PHFL, called PHFL $_d$, is derived via

$$\varphi ::= p_i \mid \varphi \vee \varphi \mid \neg \varphi \mid \langle a_i \rangle \varphi \mid \{ \sigma \} \varphi \mid \lambda(x: \tau). \varphi \mid x \mid (\varphi \varphi) \mid \mu(X: \tau). \varphi \mid X$$

$$\begin{array}{c}
\frac{}{\Gamma \vdash p_i : \bullet} \quad \frac{\Gamma \vdash \varphi_1 : \bullet \quad \Gamma \vdash \varphi_2 : \bullet}{\Gamma \vdash \varphi_1 \vee \varphi_2 : \bullet} \quad \frac{\bar{\Gamma} \vdash \varphi : \bullet}{\Gamma \vdash \neg \varphi : \bullet} \quad \frac{\Gamma \vdash \varphi : \bullet}{\Gamma \vdash \langle a_i \rangle \varphi : \bullet} \quad \frac{\Gamma \vdash \varphi : \bullet}{\Gamma \vdash \{\sigma\} \varphi : \bullet} \\
\\
\frac{\Gamma, x^v : \tau_1 \vdash \varphi : \tau_2}{\Gamma \vdash \lambda(x^v : \tau_1). \varphi : \tau_1^v \rightarrow \tau_2} \quad \frac{v \in \{+, 0\}}{\Gamma, x^v : \tau \vdash x : \tau} \quad \frac{\Gamma, X^+ : \tau \vdash \varphi : \tau}{\Gamma \vdash \mu(X : \tau). \varphi : \tau} \quad \frac{}{\Gamma, X^+ : \tau \vdash X : \tau} \\
\\
\frac{\Gamma \vdash \varphi_1 : \tau_2^+ \rightarrow \tau_1 \quad \Gamma \vdash \varphi_2 : \tau_2}{\Gamma \vdash (\varphi_1 \varphi_2) : \tau_1} \quad \frac{\Gamma \vdash \varphi_1 : \tau_2^- \rightarrow \tau_1 \quad \bar{\Gamma} \vdash \varphi_2 : \tau_2}{\Gamma \vdash (\varphi_1 \varphi_2) : \tau_1} \\
\\
\frac{\Gamma \vdash \varphi_1 : \tau_2^0 \rightarrow \tau_1 \quad \Gamma \vdash \varphi_2 : \tau_2 \quad \bar{\Gamma} \vdash \varphi_2 : \tau_2}{\Gamma \vdash \varphi_1 \varphi_2 : \tau_1}
\end{array}$$

Figure 1: The PHFL typing system.

where $x \in \mathbf{L}$, $X \in \mathbf{F}$, $p \in \mathbf{P}$, $a \in \mathbf{A}$, $1 \leq i \leq d$ and $\sigma : [d] \rightarrow [d]$ is a mapping on so-called *indices*. We assume that standard connectives such as ff , \wedge , $[a]$ and $v(X : \tau)$ are available via the obvious dualities if needed. We will also allow ourselves to use more convenient notation for the definition of (more complex) functions and their applications. For instance, $(\dots((\varphi \psi_1) \psi_2) \dots) \psi_n$ is simply written as $\varphi(\psi_1, \dots, \psi_n)$, and $\lambda(x_1 : \tau_1). \lambda(x_2 : \tau_2). \dots \lambda(x_n : \tau_n). \psi$ is written as $\lambda(x_1 : \tau_1, \dots, x_n : \tau_n). \psi$, or even $\lambda(x_1, \dots, x_n : \tau). \psi$ if $\tau = \tau_i$ for all $1 \leq i \leq n$.

The notions of free and bound variables in a formula are as usual, with $\lambda(x : \tau). \varphi$ binding x and with $\mu(X : \tau). \varphi$ binding X .

Over an LTS T , a PHFL_d formula intuitively defines a set of d -tuples in T , or a function that transforms such a set into such a set, or into a function etc., depending on the formula's type which will be explained shortly. Before that we briefly introduce the intuitive meaning of the operators in the syntax. The first five listed in the grammar above all define a subset of S^d . For example, p_i denotes all tuples such that p holds at their i th state. The operator $\{\sigma\} \varphi$ rearranges the positions in tuples in such a subset. The modality $\langle a_i \rangle \varphi$ expresses that φ holds on a tuple after replacing the i th state in it by some a -successor. A formula of the form $\lambda(x : \tau^v). \varphi$ defines a function that consumes an argument of type τ and is monotonically increasing, monotonically decreasing, or constant in this argument, depending on v . A formula of the form $(\varphi \psi)$ denotes the application of the semantics of φ to the object defined by ψ . Finally, fixpoints can now also define higher-order functions.

Obviously, not all formulas that can be derived from the above grammar can be given a semantics in a meaningful way; consider e.g. $(p_1 p_2)$. Moreover, as is typical in a situation involving least and greatest fixpoints, the use of negation has to be restricted, cf. the example $\mu(X : \bullet). \neg X$. Hence, PHFL has a type system to filter out formulas that cannot be endowed with a proper semantics.

A finite sequence Γ of *hypotheses* of the form $X^v : \tau$ or $x^v : \tau$ in which each variable occurs at most once is called a *context*. The dual context $\bar{\Gamma}$ is obtained from Γ by replacing all the hypotheses of the form $X^+ : \tau$ by $X^- : \tau$ and vice versa, and doing the same for lambda variables. We say that φ has type τ in the context Γ if the statement $\Gamma \vdash \varphi : \tau$ can be derived from the rules in Fig. 1. A formula without free variables is *well-typed* if the statement $\emptyset \vdash \varphi : \bullet$ can be derived from these rules. We tacitly assume that each fixpoint variable and each lambda variable is bound at most once in a well-typed formula, and that no variable occurs both freely and bound in a formula. Hence, each variable has a unique type in the context of a given, well-formed formula. If the type information is clear from context or not important,

we drop it from binders, simply writing $\lambda x. \varphi$ and $\mu X. \varphi$ for better readability.

A formula is said to be of order k if the maximal order of the type of any subformula in φ is k . By PHFL_d^k we denote the set of well-typed formulas in PHFL_d that are of order k . Note that PHFL does indeed constitute an extension of other known formalisms, namely

- PHFL_1^k is the same as HFL^k for any $k \geq 0$ and, thus $\text{PHFL}_1 = \text{HFL}$, and in particular
- PHFL_1^0 is just the modal μ -calculus \mathcal{L}_μ while
- PHFL_d^0 is the d -dimensional polyadic μ -calculus.

Semantics. Let $T = (S, \{\xrightarrow{a}\}_{a \in \mathbf{A}}, \ell)$ be an LTS and let φ be a well-typed formula. An *environment* is a function η that assigns to each fixpoint variable and each lambda variable of type τ an element of $\llbracket \tau \rrbracket^T$.

Let $d \geq 1$. The semantics $\llbracket \varphi : \tau \rrbracket_\eta^T$ of a PHFL_d formula φ of type τ , relative to an LTS T and an environment η , is an object of $\llbracket \tau \rrbracket^T$, defined recursively as follows.

$$\begin{aligned}
\llbracket \Gamma \vdash p_i : \bullet \rrbracket_\eta^T &= \{(s_1, \dots, s_d \in S^d \mid p \in \ell(s_i))\} \\
\llbracket \Gamma \vdash \varphi_1 \vee \varphi_2 : \bullet \rrbracket_\eta^T &= \llbracket \Gamma \vdash \varphi_1 : \bullet \rrbracket_\eta^T \cup \llbracket \Gamma \vdash \varphi_2 : \bullet \rrbracket_\eta^T \\
\llbracket \Gamma \vdash \neg \varphi : \bullet \rrbracket_\eta^T &= S^d \setminus \llbracket \Gamma \vdash \varphi : \bullet \rrbracket_\eta^T \\
\llbracket \Gamma \vdash \langle a_i \rangle \varphi : \bullet \rrbracket_\eta^T &= \{(s_1, \dots, s_d) \in S^d \mid \text{ex. } t \text{ s.t.} \\
&\quad (s_1, \dots, s_{i-1}, t, s_{i+1}, \dots, s_d) \in \llbracket \Gamma \vdash \varphi : \bullet \rrbracket_\eta^T \text{ and } s_i \xrightarrow{a} t\} \\
\llbracket \Gamma \vdash \{\sigma\} \varphi : \bullet \rrbracket_\eta^T &= \{(s_1, \dots, s_d) \in S^d \mid (s_{\sigma(1)}, \dots, s_{\sigma(d)}) \in \llbracket \Gamma \vdash \varphi : \bullet \rrbracket_\eta^T\} \\
\llbracket \Gamma \vdash \lambda(x^v : \tau_2). \varphi : \tau_1^v \rightarrow \tau_1 \rrbracket_\eta^T &= f \in \llbracket \tau_2^v \rightarrow \tau_1 \rrbracket^T \text{ s.t. f.a. } y \in \llbracket \tau_2 \rrbracket^T. f(y) = \llbracket \Gamma, x^v : \tau_2 \vdash \varphi : \tau_1 \rrbracket_{\eta[x \mapsto y]}^T \\
\llbracket \Gamma \vdash x : \tau \rrbracket_\eta^T &= \eta(x) \\
\llbracket \Gamma \vdash (\varphi_1 \varphi_2) : \tau_1 \rrbracket_\eta^T &= \llbracket \Gamma \vdash \varphi_1 : \tau_2^v \rightarrow \tau_1 \rrbracket_\eta^T (\llbracket \Gamma \vdash \varphi_2 : \tau_2 \rrbracket_\eta^T) \\
\llbracket \Gamma \vdash \mu(X : \tau). \varphi : \tau \rrbracket_\eta^T &= \bigsqcap_{\tau \rightarrow \tau} \{d \in \llbracket \tau \rrbracket_\eta^T \mid \llbracket \Gamma, X : \tau^+ \vdash \varphi : \tau \rrbracket_{\eta[X \mapsto d]}^T \sqsubseteq \tau d\} \\
\llbracket \Gamma \vdash X : \tau \rrbracket_\eta^T &= \eta(X)
\end{aligned}$$

If the type is clear from context, or not important, we simply write $\llbracket \varphi \rrbracket_\eta^T$. We write $T, (s_1, \dots, s_d) \models_\eta \varphi$ if $\varphi : \bullet$ and $(s_1, \dots, s_d) \in \llbracket \varphi \rrbracket_\eta^T$. We say that two formulas φ and ψ are *equivalent*, written $\varphi \equiv \psi$, if for all T and all η we have $\llbracket \varphi \rrbracket_\eta^T = \llbracket \psi \rrbracket_\eta^T$.

It is well-known that the semantics of HFL and, hence, PHFL is invariant under β -reduction and admits the fixpoint unfolding principle, i.e. $\mu X. \varphi \equiv \varphi[\mu X. \varphi/X]$ where substitution is defined as usual.

Bisimilarity. A *bisimulation* R on an LTS $T = (S, \{\xrightarrow{a}\}_{a \in \mathbf{A}}, \ell)$ is a symmetric relation $R \subseteq S \times S$ satisfying the following for all $(s, t) \in R$.

- $\ell(s) = \ell(t)$,
- if there is $a \in \mathbf{A}$ and $s' \in S$ s.t. $s \xrightarrow{a} s'$ then there is $t' \in S$ with $t \xrightarrow{a} t'$ and $(s', t') \in R$,
- if there is $a \in \mathbf{A}$ and $t' \in S$ s.t. $t \xrightarrow{a} t'$ then there is $s' \in S$ with $s \xrightarrow{a} s'$ and $(s', t') \in R$.

Two states s, t are bisimilar, written $s \sim t$, if there is a bisimulation R with $(s, t) \in R$.

Let $d \geq 1$. A set $T \subseteq S$ is called *bisimulation-invariant* if for all $\bar{s} = (s_1, \dots, s_d), \bar{t} = (t_1, \dots, t_d) \in S^d$ such that $s_i \sim t_i$ for all $i \in [d]$, we have $\bar{s} \in T$ iff $\bar{t} \in T$. The notion of bisimulation-invariance can straight-forwardly be lifted to objects of type $\llbracket \tau \rrbracket^T$ for types $\tau \neq \bullet$, cf. [21].

It is well-known that modal logics cannot distinguish bisimilar models, and it is not surprising that the extensions beyond pure modal logic that are available in PHFL do not break this property.

Proposition 1 ([21, 15]). *Let $d \geq 1$, T be an LTS with state set S and φ be a closed PHFL $_d$ formula of type \bullet . Then $\llbracket \varphi \rrbracket^T \subseteq S^d$ is bisimulation-invariant.*

Example 2 ([17]). A standard example shows that bisimilarity itself is definable in PHFL $_2^0$, provided that \mathbf{P} and \mathbf{A} are finite. The PHFL $_2^0$ formula

$$\varphi_{\sim} := \nu(X : \bullet). \left(\bigwedge_{p \in \mathbf{P}} p_1 \leftrightarrow p_2 \right) \wedge \left(\bigwedge_{a \in \mathbf{A}} [a_1] \langle a_2 \rangle X \right) \wedge \{1 \mapsto 2, 2 \mapsto 1\} X$$

is satisfied by a pair (s, t) of some T iff $s \sim t$. The formula essentially states that (s, t) needs to belong to the largest set X of states (s', t') that agree on all propositions (first conjunct) and for which t' can match any a -transition out of s' , for any $a \in \mathbf{A}$, to a pair $(s'', t'') \in X$ (second conjunct). Moreover, X needs to be a symmetric relation (third conjunct).

To exemplify the use of higher-orderness (here: first-order functions) we can use the definability of finite-trace equivalence in PHFL $_2^1$.

Example 3. Two states s, t are finite-trace equivalent if whenever there is a sequence $s \xrightarrow{a_1} s_1 \xrightarrow{a_2} s_2 \xrightarrow{a_3} \dots \xrightarrow{a_n} s_n$ then there are t_1, \dots, t_n s.t. $t \xrightarrow{a_1} t_1 \xrightarrow{a_2} t_2 \xrightarrow{a_3} \dots \xrightarrow{a_n} t_n$ and vice-versa.

$$\varphi_{\text{fte}} := \left(\nu F(x, y). (x \leftrightarrow y) \wedge \bigwedge_{a \in \mathbf{A}} F(\langle a_1 \rangle x, \langle a_2 \rangle y) \right) (\text{tt}, \text{tt})$$

is then satisfied by a pair (s, t) iff s and t are finite-trace equivalent in the sense above. The greatest fixpoint in the formula expresses an infinite conjunction over all finite paths, these can be thought of to be built step-wise via fixpoint unfolding; the n th unfolding expresses that all finite traces of length n are available in the first component of the tuple (and built via $\langle a_1 \rangle$) iff they are available in the second component (built via $\langle a_2 \rangle$). For better readability we have omitted the types in the formula. The types of x, y are \bullet and that of F is, consequently, $\bullet^0 \rightarrow \bullet^0 \rightarrow \bullet$.

We remark that it is easily possible to extend the formula to also check for matching atomic propositions along the traces emerging from s and t . Cf. [16] for further examples of how various other process equivalences and preorders can be expressed in PHFL 1 .

2.2 Higher-Order Logic with Least Fixpoints

We introduce Higher-Order Logic with Least Fixpoints (HO(LFP)) to make use of the characterisation of k -EXPTIME over the class of ordered structures as the queries definable in order- $(k+1)$ HO(LFP), due to Immerman and Vardi [8, 20, 9], resp. Freire and Martins [6].

Types. Types for Higher-Order Logic with Least Fixpoints are constructed, similar to those for PHFL, from a single base type and one constructor: $\tau' ::= \odot \mid (\tau', \dots, \tau')$. Here, however, \odot is the type of *individuals* (like states rather than sets of states), and the tuple type is used to denote (higher-order) relations. We define the order 1 of a type via $\text{ord}(\odot) = 1$ and $\text{ord}(\tau'_1, \dots, \tau'_n) = 1 + \max\{\text{ord}(\tau'_1), \dots, \text{ord}(\tau'_n)\}$.

Let $T = (S, \{\xrightarrow{a}\}_{a \in \mathbf{A}}, \ell)$ be an LTS. This induces a set-theoretic interpretation of types via $\llbracket \odot \rrbracket^T = S$ and $\llbracket (\tau'_1, \dots, \tau'_n) \rrbracket^T = 2^{\llbracket \tau'_1 \rrbracket^T \times \dots \times \llbracket \tau'_n \rrbracket^T}$.

¹Note the discrepancy in the traditional ways to assign numerals to orders in the two logics considered here: order 1 in HO(LFP) refers to, like in ‘‘First-Order Logic’’, individual elements and order 2 is for relations like sets thereof. In PHFL, order 1 refers to the order of a function, i.e. one that takes sets of arguments. This explains why PHFL k corresponds to the fragment of HO(LFP) of order $k+1$, see also the right column in Fig. 2.

Syntax. Let $\mathbf{V} = \{X, \dots\}$ be a countable set of *higher-order variables*, each implicitly equipped with a type. Let \mathbf{P}, \mathbf{A} be sets of propositions, resp. actions. The syntax of HO(LFP) formulas is derived from the following grammar:

$$\varphi ::= p(X) \mid a(X, Y) \mid X(Y_1, \dots, Y_N) \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists(X: \tau'). \varphi \mid (\text{lfp}(X, Y_1, \dots, Y_N). \varphi)(Z_1, \dots, Z_n)$$

where $p \in \mathbf{P}$, $a \in \mathbf{A}$ and $X, Y_1, \dots, Y_n, Z_1, \dots, Z_n \in \mathbf{V}$. Again, dual operators such as \wedge and *gfp* (for greatest fixpoints) are available via the obvious dualities. The variable X is bound in $\exists(X: \tau'). \varphi$ and X, Y_1, \dots, Y_n are bound in $(\text{lfp}(X, Y_1, \dots, Y_N). \varphi)(Z_1, \dots, Z_n)$. A formula is *well-formed* if each variable is bound at most once, and, moreover, variables occur only in a way that matches their type. For example, only a variable of type \odot can occur in a subformula of the form $p(X)$, and if a subformula of the form $X(Y_1, \dots, Y_n)$ occurs, then X has type $(\tau'_1, \dots, \tau'_n)$ for some τ'_1, \dots, τ'_n and Y_i has type τ'_i for all $1 \leq i \leq n$. Moreover, in a subformula of the form $(\text{lfp}(X, Y_1, \dots, Y_N). \varphi)(Z_1, \dots, Z_n)$, the variable X occurs only under an even number of negations in φ . For a more detailed introduction into Higher-Order Logic including formal ways to define well-formedness of formulas, cf. [19]. In a well-formed formula each variable has a unique type. An HO(LFP) formula φ has order k if the order of the highest type of a variable in φ is at most k . We write $\text{HO}^k(\text{LFP})$ for the set of HO(LFP) formulas of order at most k .

Semantics. Let $T = (S, \{\xrightarrow{a}\}_{a \in \mathbf{A}}, \ell)$ be an LTS. A variable assignment α is a function that maps each variable in \mathbf{V} of type τ' into $\llbracket \tau' \rrbracket^T$.

Let φ be an HO(LFP) formula with free variables in X, Y_1, \dots, Y_n such that X occurs only under an even number of negations in φ and let $\tau' = (\tau'_1, \dots, \tau'_n)$ be the type of X in φ while τ'_i is the type of Y_i in φ for $1 \leq i \leq n$. Then, given some variable assignment α , the formula φ defines a monotone function $f: \llbracket \tau' \rrbracket^T \rightarrow \llbracket \tau' \rrbracket^T$ via

$$f(M) \mapsto \{(m_1, \dots, m_n) \in \llbracket \varphi \rrbracket_{\alpha[Y_1 \mapsto m'_1, \dots, Y_n \mapsto m'_n]}^T \mid (m'_1, \dots, m'_n) \in M\}.$$

By the Knaster-Tarski-Theorem [11, 18] this function has a least fixpoint denoted by $LFP(X, Y_1, \dots, Y_n)\varphi$. Note that we suppress T and α here since they will always be clear from context.

The satisfaction relation between an LTS T , a variable assignment α and an HO(LFP) formula φ is defined inductively as follows.

$$\begin{aligned} T, \alpha &\models p(X) \text{ iff } p \in \ell(\alpha(X)) \\ T, \alpha &\models a(X, Y) \text{ iff } \alpha(X) \xrightarrow{a} \alpha(Y) \\ T, \alpha &\models X(Y_1, \dots, Y_n) \text{ iff } (\alpha(Y_1), \dots, \alpha(Y_n)) \in \alpha(X) \\ T, \alpha &\models \neg\varphi \text{ iff } T, \alpha \not\models \varphi \\ T, \alpha &\models \varphi_1 \vee \varphi_2 \text{ iff } T, \alpha \models \varphi_1 \text{ or } T, \alpha \models \varphi_2 \\ T, \alpha &\models \exists(X: \tau') \text{ iff ex. } d \in \llbracket \tau' \rrbracket^T \text{ s.t. } T, \alpha[X \mapsto d] \models \varphi \\ T, \alpha &\models (\text{lfp}(X, Y_1, \dots, Y_N). \varphi)(Z_1, \dots, Z_n) \text{ iff } (\alpha(Z_1), \dots, \alpha(Z_n)) \in LFP(X, Y_1, \dots, Y_n)\varphi. \end{aligned}$$

2.3 Descriptive Complexity

A *query* (of dimension d) is a set Q of pairs $(T, (s_1, \dots, s_d))$ s.t. each T is finite. It is expressed by the HO(LFP) formula φ with free variables X_1, \dots, X_d , if $Q = \{(T, (s_1, \dots, s_d)) \mid T, [X_1 \mapsto s_1, \dots, X_d \mapsto s_d] \models \varphi\}$. Let $k \geq 0$. The complexity class k -EXPTIME is defined as $DTIME(2_k^{m^{\sigma(1)}})$ where $2_0^m := m$ and $2_{k+1}^m := 2^{2_k^m}$. Note that 0-EXPTIME equals P.

For a complexity class \mathcal{C} , a \mathcal{C} -query is one that can be decided within the resource bounds given by \mathcal{C} . We write \mathcal{C}/\sim for the complexity class of \mathcal{C} -queries that are bisimulation-invariant.

We say that a logic \mathcal{L} captures a complexity class \mathcal{C} if the model checking problem for \mathcal{L} is in \mathcal{C} and each \mathcal{C} -query can be expressed in \mathcal{L} .

The Immerman-Vardi Theorem characterises the P-queries over ordered structures as those expressible in $\text{HO}^1(\text{LFP})$, resp. first-order logic with least fixpoints. It's generalisation is the following:

Proposition 4 ([8, 20, 9, 6]). *For each $k \geq 0$, $\text{HO}^{k+1}(\text{LFP})$ captures $k\text{-EXPTIME}$ over the class of ordered structures.*

The ordering is only important for the case of $k = 0$ as such an ordering can be defined in second-order logic, i.e. as soon as $k \geq 1$. It is an open problem whether a logic exists that captures P over the class of all structures (cf. e.g. [7]).

The first capturing result for a bisimulation-invariant class is given by Otto's Theorem.

Proposition 5 ([17]). *The polyadic modal μ -calculus \mathcal{L}_μ^0 , or, equivalently PHFL^0 , captures P/\sim , or equivalently, $0\text{-EXPTIME}/\sim$.*

We give a quick sketch of its proof in order to prepare for the technical developments in Sects. 4 and 5. The model checking problem for PHFL^0 is readily seen to be in P (see also [14]). The interesting part is to show that every bisimulation-invariant P query can be expressed in PHFL^0 . In principle, this could be done by encoding runs of polynomially time-bounded Turing machines, but it is not immediately clear how bisimulation-invariance of the query in question can be used. Instead, the proof for this rests on a key observation: non-bisimilarity of two states can be expressed in PHFL_2^0 by a least-fixpoint formula, since bisimilarity can be defined via the greatest fixpoint formula in Ex. 2. Hence, non-bisimilarity of two states, as a least fixpoint, has a well-founded reason, i.e. one that can be found in finitely many fixpoint unfoldings. Ordering the atomic types in an arbitrary way entails a total order on the bisimulation-equivalence classes, and this order can be defined in PHFL_2^0 . Hence, the LTS in question is ordered, and the Immerman-Vardi Theorem is available, whence the problem reduces to showing that every query defined by a bisimulation-invariant $\text{HO}^1(\text{LFP})$ -query can be expressed equivalently in PHFL^0 .

This latter reduction now follows from a rather straightforward translation from $\text{HO}^1(\text{LFP})$ formulas to PHFL^0 formulas. Variables of type \odot are emulated through polyadicity and variables of type (\odot, \dots, \odot) are represented as order-0 PHFL^0 -variables. Since, in the bisimulation-invariant setting, one can always assume that the LTS in question already is its own bisimulation quotient, bisimilarity and equality coincide. Hence, a subformula of the form $a(X_i, X_j)$ can be replaced by the statement that the i th component of the relation defined has an a -successor that is bisimilar and, hence equal to the j th component. Moreover, since all states in such a bisimulation quotient of a pointed LTS are reachable from a distinguished state, existential quantification can be replaced by reachability of a suitable state. It remains to translate least fixpoints in $\text{HO}^1(\text{LFP})$ into order-0 fixpoints of PHFL^0 .

Using reasoning along similar lines, Otto's Theorem has also been generalised by one order.

Proposition 6 ([15]). *PHFL^1 captures $\text{EXPTIME}/\sim$.*

3 Upper Bounds

Capturing a complexity class \mathcal{C} , defined by some restricted resource consumption, by a logic \mathcal{L} contains two parts: what is commonly seen as the lower bound consists of showing that every query which can be evaluated in complexity \mathcal{C} can also be defined in the logic \mathcal{L} . The upper bound is established by

showing the contrary. This is relatively easy as it suffices to show that queries definable in \mathcal{L} can be evaluated in complexity \mathcal{C} , in other words that the model checking problem for \mathcal{L} belongs to class \mathcal{C} .

Here we do this for the fragments of PHFL of arbitrary but fixed arity d and arbitrary order k , w.r.t. the classes k -EXPTIME of the exponential time hierarchy. We do so by extending the reduction of the model checking problem for a polyadic logic to that of its monadic fragment [14]. Note that PHFL_1^k equals HFL^k – the fragment of (non-polyadic) Higher-Order Fixpoint Logic of formulas of type order at most k . The complexity of model checking such fragments is known:

Proposition 7 ([3]). *Let $k \geq 1$. The model checking problem for HFL^k is k -EXPTIME-complete.*

The corresponding result for PHFL_d^k , first stated without proof in [15], follows via a reduction:

Theorem 8. *Let $k, d \geq 1$. The model checking problem for PHFL_d^k is in k -EXPTIME.*

Proof. By a polynomial reduction to the model checking problem for HFL. Let $k, d \geq 1$ and an LTS $T = (S, \{\xrightarrow{a}\}_{a \in \mathbf{A}}, \ell)$ over \mathbf{A} and \mathbf{P} be given. We construct its d -product T^d over the action set $\mathbf{A}^d := \{a_i \mid a \in \mathbf{A}, i \in [d]\} \cup \{\sigma \mid \sigma : [d] \rightarrow [d]\}$ and atomic propositions $\mathbf{P}^d := \{q_i \mid q \in \mathbf{P}, i \in [d]\}$ as $(S^d, \{\xrightarrow{x}\}_{x \in \mathbf{A}_d^d}, \ell')$ where, for all $q \in \mathbf{P}$, $i \in [d]$, $a \in \mathbf{A}$, $s_1, \dots, s_d, t_1, \dots, t_d \in S$ we have

- $q_i \in \ell'(s_1, \dots, s_d)$ iff $q \in \ell(s_i)$,
- $(s_1, \dots, s_d) \xrightarrow{a_i} (t_1, \dots, t_d)$ iff $s_i \xrightarrow{a} t_i$ and $t_j = s_j$ for all $j \neq i$,
- $(s_1, \dots, s_d) \xrightarrow{\sigma} (t_1, \dots, t_d)$ iff $t_j = \sigma(s_j)$ for all $j \in [d]$.

Next, we translate a PHFL_d^k formula φ inductively into a PHFL_1 formula $\widehat{\varphi}$ as follows. The operation $\widehat{\cdot}$ acts homomorphically on all operators apart from the following three cases.

$$\widehat{\langle \sigma \rangle \psi} := \langle \sigma \rangle \widehat{\psi} \quad , \quad \widehat{\langle a_i \rangle \psi} := \langle a_i \rangle \widehat{\psi} \quad , \quad \widehat{p_i} := p_i .$$

The latter two cases may be confusing as $\widehat{\cdot}$ seems to not change those operators either. However, in e.g. the second case, $\langle a_i \rangle$ on the left side is a polyadic modality combining the action $a \in \mathbf{A}$ with the index $i \in [d]$. On the right side, $\langle a_i \rangle$ is a monadic modality over the action $a_i \in \mathbf{A}^d$. Likewise in the third case.

This equality in syntax for two technically different modal operators, resp. atomic formulas is intended because of the following connection: the d -tuple (s_1, \dots, s_d) of states in T satisfies the PHFL_d formula $\langle a_i \rangle \psi$ iff the state (s_1, \dots, s_d) of T^d satisfies the PHFL_1 formula $\langle a_i \rangle \widehat{\psi}$. A similar statement can be made for atomic propositions and these can easily be generalised to show by induction on the syntax of PHFL that for all PHFL_d formulas φ , all $s_1, \dots, s_d \in S$ and all environments η we have: $(s_1, \dots, s_d) \in \llbracket \varphi \rrbracket_\eta^T$ iff $(s_1, \dots, s_d) \in \llbracket \widehat{\varphi} \rrbracket_\eta^{T^d}$.

This establishes correctness of the reduction. Note that $\widehat{\varphi}$ is a PHFL_1^k , i.e. HFL^k formula whenever $\varphi \in \text{PHFL}_d^k$. Moreover, both T^d and $\widehat{\varphi}$ are easily seen to be constructible in polynomial time (for fixed d). Thus, by Prop. 7 model checking PHFL_d^k is also in k -EXPTIME. \square

4 Higher-Order Quantification

To show that every bisimulation-invariant k -EXPTIME query can be expressed in PHFL^k it suffices, due to Prop. 4, to show that every $\text{HO}^{k+1}(\text{LFP})$ -query can be translated into a PHFL^k formula. The main challenge here is to deal with existential quantification, which has no obvious equivalent in PHFL. In [17], first-order existential quantification is replaced by reachability of a suitable state, which is sufficient in the bisimulation-invariant setting. Higher-order quantification does not have such an obvious

correspondent - there is no notion of a set, or a set of sets, etc. being reachable. Moreover, PHFL does only have a type for sets of tuples of states, not for sets of sets etc. We solve this problem by replacing higher-order types by a variant of their characteristic function, something that fits quite naturally into the PHFL world. We then lift the order on the states inherited from [17] to sets of tuples of states and then to said characteristic functions by ordering them lexicographically. We can enumerate sets, functions and so on alongside this order, and hence, we can mimic existential quantification over some $\text{HO}(\text{LFP})$ type by an enumeration of corresponding characteristic functions.

The increased complexity of this approach compared to simple reachability requires two adaptations: first, since we use the order of sets, functions etc. present in the bisimulation-invariant setting, we often have to compare two such objects w.r.t. this order. Comparing e.g. two sets of d -ary tuples, however, requires a formula containing types of width $2d$. Also, $\text{HO}(\text{LFP})$ -formulas can define a query of some width, yet contain types, resp. quantification over objects of much higher width. Hence, in order to keep the presentation simple, the formulas we develop subsequently will be of some unspecified, yet generally quite high arity, i.e. they will be in PHFL_d^k for some d that is large compared with the width of the original query. The exact value of d will be given towards the end of the translation.

Since we have already agreed to blow up the width used in our translated formulas, we can also make things easier by reserving certain positions in the tuples we work with for special tasks. We generally use the last two positions in our tuples (those with indices d and $d - 1$) to compare individual states w.r.t. the order from [17], and we will use the next r positions from the right, i.e. those with indices $d - r - 1, \dots, d - 2$ for some r , to keep copies of states such that the whole LTS is reachable from at least one of these states, in order to keep the pattern for first-order quantification valid. This will be made formal after we revisit the pattern for existential quantification just below. We then subsequently expand quantification towards sets and characteristic functions. The final translation is then given in Sec. 5.

Reachable States. In [17], existential quantification of first-order logic was replaced by reachability of a suitable state in the LTS in question, using the pattern given subsequently. First, we recall the role of the substitution operator: Let $\sigma^{i \leftarrow j}$ be defined by $\sigma^{i \leftarrow j}(i) = j$ and $\sigma^{i \leftarrow j}(i') = i'$ if $i' \neq i$. Then $(\bar{s} \in \llbracket \{\sigma^{i \leftarrow j}\} \varphi \rrbracket_\eta^T$ iff $\bar{s}[s_j/i] = (s_1, \dots, s_{i-1}, s_j, s_{i+1}, \dots, s_d) \in \llbracket \varphi \rrbracket_\eta^T$ for all T and η .

Now let \mathbf{A} be a set of actions, let $d > 2$, $r \leq d - 2$ and $i \leq d - r - 2$. Recall that, for the time being, we assume that every tuple we work with is such that all states in an LTS are reachable from one of the states at indices $d - r - 1, \dots, d - 2$ of the tuple and that we reserve the last two positions for comparisons (see below). Consider the formula

$$\exists_i \varphi := \bigvee_{j=d-r-1}^{d-2} \{ \sigma^{i \leftarrow j} \} (\mu(X : \bullet). \varphi \vee \bigvee_{a \in \mathbf{A}} \langle a_i \rangle X).$$

for some $\varphi \in \text{PHFL}_d$. We have the following:

Observation 9. Let T be an LTS over \mathbf{A} and let $s = (s_1, \dots, s_d)$ such that all states in T are reachable from at least one state in $s_{d-r-1}, \dots, s_{d-2}$. Then $T, \bar{s} \models \exists_i \varphi$ iff there is a state t in T such that $T, \bar{s}[t/i] \models \varphi$.

As said before, quantification over higher-order types is more complicated, but we can replace reachability by enumeration in lexicographical order for higher-order types. Towards this, note that for all $d \geq 2$ there is some PHFL_2^0 -formula $\varphi_<$ defining a transitive and irreflexive relation $<$ such that, for all LTS T and d -tuples $\bar{s} = (s_1, \dots, s_d)$ we have that $T, \bar{s} \models \varphi_<$ iff $s_{d-1} < s_d$. This formula is defined in [17] using a variant of the negation of the formula from Ex. 2. The actual position of the tuple elements that are compared is not important, we choose to fix it here for consistency.

A crucial ingredient for the correctness of the quantification pattern above is that every state in the LTS is reachable from the states in positions $d - r - 1, \dots, d - 2$. For the first-order case, this can be guaranteed by never manipulating the components with the respective indices. In the higher-order setting, this does not suffice as one deals with arbitrary sets, functions, etc. Hence, for the remainder of the section all formulas are assumed to have a free lambda variable e of type \bullet that is as follows:

Definition 10. Let T be an LTS and let r be fixed. Then an interpretation η is *good* if $\eta(e)$ is a set of the form $M \times \{s_{d-r-1}\} \times \dots \times \{s_{d-2}\} \times S^2$ such that $\emptyset \neq M \subseteq S^{d-2-r}$ and each state of T is reachable from one of the $s_{d-r-1}, \dots, s_{d-2}$.

We do not make this free variable explicit, since it is always assumed to be there. We will see in Sect. 5 how this intended interpretation can be enforced. This stipulation formalises the informal idea given above. Note that in [17], it was assumed that all states of the LTS in question are reachable from a singular state, but this is not a necessary requirement for the argument to work.

Finally, let w and d be such that $2w + r + 2 \leq d$. The intuition here is that, in order to translate from HO(LFP), we will have to deal with sets and higher-order sets of arity at most w .

Quantification for Sets. We now define a similar pattern to that for the first-order case which allows us to iterate over all sets of w -tuples in an LTS. This is not the same as enumerating $\llbracket \bullet \rrbracket^T$ since $w < d$.

Let σ_i be defined via $\sigma_i(d-1) = i$, $\sigma_i(d) = i + w$ and $\sigma_i(j) = j$ if $j < d - 1$. Then $\bar{s} \in \llbracket \{\sigma_i\} \varphi \rrbracket_\eta^T$ iff $\bar{s}[s_i/d-1, s_{i+w}/d] \in \llbracket \varphi \rrbracket_\eta^T$. The intended use for this substitution is to compare the elements at indices i and $i + w$ w.r.t. to the order induced by $\varphi_{<}$. Remember that this formula always compares the last two elements of the tuple. Moreover, let $\sigma_{\rightarrow w}$ be defined by $\sigma_{\rightarrow w}(i) = w + i$ for $i \leq w$ and $\sigma_{\rightarrow w}(j) = j$ for $j > w$. Then $(s_1, \dots, s_d) \in \llbracket \{\sigma_{\rightarrow w}\} \varphi \rrbracket_\eta^T$ iff $(s_1, \dots, s_w, s_1, \dots, s_w, s_{2w+1}, \dots, s_d) \in \llbracket \varphi \rrbracket_\eta^T$. The intended use here is to shift the first w elements of a formula to the right to make room for another tuple at the first w positions such that the two tuples can be compared lexicographically.

Consider the formula $\exists^{(w)}x. \varphi := (\mu(F : \bullet \rightarrow \bullet). \lambda(x : \bullet). \varphi \vee F(\text{next}^{(w)}(x))) \text{ff}$ where

$$\begin{aligned} \varphi_{<}^w &:= \bigvee_{i=1}^w \{\sigma_i\} \varphi_{<} \wedge \bigwedge_{j=1}^{i-1} \{\sigma_j\} \neg \varphi_{<} \\ \varphi_{<}^{(w)}(x, y) &:= \exists_1 \dots \exists_w. y \wedge \neg x \wedge \{\sigma_{\rightarrow w}\} (\forall_1 \dots \forall_w. \varphi_{<}^w \rightarrow x \rightarrow y). \\ \text{next}^{(w)}(x) &:= \lambda(x : \bullet). e \wedge \neg x \wedge \{\sigma_{\rightarrow w}\} (\forall_1 \dots \forall_w. \varphi_{<}^w \rightarrow x) \\ &\quad \vee e \wedge x \wedge \{\sigma_{\rightarrow w}\} (\exists_1 \dots \exists_w. \varphi_{<}^w \wedge \neg x) \end{aligned}$$

The first formula $\varphi_{<}^w$ implements lexicographical comparison of the first w elements in a tuple to the second w elements. The second formula $\varphi_{<}^{(w)}$ lifts the order from tuples to sets of tuples via the lexicographical order induced by the order on the tuples. Finally, the formula $\text{next}^{(w)}$ is a predicate transformer that consumes a set of tuples. It returns a set of tuples that is the lexicographical successor of the input in the order induced by the order on the first w elements of the individual tuples: the output contains a tuple iff either the input does contain it, but not all lexicographically smaller tuples, or if the input does not contain it, but all lexicographically smaller tuples. Also note the role of e that filters out all tuples that do not adhere to our stipulation that the whole LTS be reachable from one of the states at indices $d - r - 1, \dots, d - 2$.

Lemma 11. Let T be an LTS with state set S and let η be good. Then $T, \bar{s} \models_\eta \exists^{(w)}x. \varphi$ iff there is $M \subseteq S^w$ such that $T, \bar{s} \models_{\eta[x \rightarrow M \times S^{d-w} \cap \eta(e)]} \varphi$.

The proof consists of verifying the informal intuition above. We write $\forall^{(w)}x. \varphi$ to denote $\neg \exists^{(w)}x. \neg \varphi$. We write $\exists^{(w)}x_1, \dots, x_n. \varphi$ for $\exists^{(w)}x_1 \dots \exists^{(w)}x_n. \varphi$, and similarly for $\forall^{(w)}x. \varphi$.

Generalised Higher-Order Quantification. We have just seen how existential quantification can be emulated for individual states in an LTS and, with some restrictions, for sets of w -tuples of an LTS. For other types that commonly appear in HO(LFP), i.e. relations of higher order, there is no immediate PHFL equivalent, since all types beyond \bullet are function types. However, we can use these function types to emulate the HO(LFP) types to a sufficient degree. For the sake of simplicity, we only consider types of a special form; we argue in Sect. 5 why this is not a restriction.

Let $\tau_{w,k}$ be inductively defined via $\tau_{w,0} = \bullet$ and $\tau_{w,i+1} = \tau_{w,i} \rightarrow \dots \rightarrow \tau_{w,i} \rightarrow \bullet$ where $\tau_{w,i}$ is repeated w many times. Given T , let $\llbracket \tau_{w,i}^\circ \rrbracket^T = \llbracket \tau_{w,i} \rrbracket^T$ for $i \leq 1$ and let

$$\llbracket \tau_{w,k}^\circ \rrbracket^T = \{f \in \llbracket \tau_{w,k} \rrbracket^T \mid f(f_1, \dots, f_w) = S \text{ or } f(f_1, \dots, f_w) = \emptyset \text{ f.a. } f_1, \dots, f_w \in \llbracket \tau_{w,k-1}^\circ \rrbracket^T\}$$

for $k \geq 2$. The important distinction here is that $\llbracket \tau_{w,k}^\circ \rrbracket^T$ is the restriction of $\llbracket \tau_{w,k} \rrbracket^T$ to those functions that always return either the full set of states or the empty set, at least on inputs from $\llbracket \tau_{w,k-1}^\circ \rrbracket^T$. This is desirable since we want to use functions in $\llbracket \tau_{w,k}^\circ \rrbracket^T$ to emulate higher-order variables of a special form. Given x, x_1, \dots, x_w of the appropriate type, the question whether $\bar{s} \in \llbracket x(x_1, \dots, x_w) \rrbracket_\eta^T$ does not depend on \bar{s} (as in modal logics), but is uniform over the while LTS. However, since $\llbracket \tau_{w,k} \rrbracket^T$ also contains functions that are not uniform starting from PHFL-order 2, we restrict ourselves to functions that are uniform on the necessary inputs (i.e. those that are themselves sufficiently uniform).

Consider the following formulas for $k \geq 1$, where $\bar{x} = x_1, \dots, x_w$ and $\bar{y} = y_1, \dots, y_w$:

$$\begin{aligned} \varphi_{<}^{w,k-1}(x_1, \dots, x_w, y_1, \dots, y_w) &:= \bigvee_{i=1}^w \varphi_{<}^{(w),k-1}(x_i, y_i) \wedge \bigwedge_{j=1}^{i-1} \neg \varphi_{<}^{(w),k-1}(x_j, y_j) \\ \varphi_{<}^{(w),k}(x, y) &:= \exists^{w,k-1} \bar{x}. y(\bar{x}) \wedge \neg x(\bar{x}) \\ &\quad \wedge \forall^{w,k-1} \bar{y}. \varphi_{<}^{w,l}(\bar{y}, \bar{x}) \rightarrow x(\bar{y}) \rightarrow y(\bar{y}) \\ \mathbf{ff}_{(w),k} &:= \lambda(\bar{x}: \tau_{w,k-1}). \mathbf{ff} \\ \mathbf{next}^{w,k}(x) &:= \lambda(x: \tau_{w,k}). \lambda(\bar{x}: \tau_{w,k-1}). \\ &\quad (\neg x(\bar{x}) \wedge \forall^{w,k} \bar{y}. \varphi_{<}^{w,k}(\bar{y}, \bar{x}) \rightarrow x(\bar{y})) \\ &\quad \vee (x(\bar{x}) \wedge \exists^{w,k} \bar{y}. \varphi_{<}^{w,k}(\bar{y}, \bar{x}) \wedge \neg x(\bar{y})) \\ \exists^{w,k}(x). \varphi &:= (\mu(F: \tau_{w,k} \rightarrow \tau_{w,k}). \lambda(x: \tau_{w,k}). \varphi \vee F(\mathbf{next}^{w,k} x)) \mathbf{ff}_{(w),k} \end{aligned}$$

where $\varphi_{<}^{(w),k-1} = \varphi_{<}^{(w)}$ in case $k = 1$ and $\exists^{w,k-1} x. \varphi = \exists^{(w)} x. \varphi$ if $k = 1$.

Similarly as in the definitions given before Lemma 11, these formulas lift quantification up by one level on the type hierarchy. The formula $\varphi_{<}^{w,k-1}$ compares width- w -tuples of functions of type $\tau_{w,k-1}^\circ$ lexicographically using the previously defined formula $\varphi_{<}^{(w),k-1}$ that compares individual such functions. The formula $\varphi_{<}^{(w),k}$ then lifts this to individual functions of the next type, using existential and universal quantification. The formula $\mathbf{next}^{w,k}$ again consumes a function of type $\tau_{w,k}^\circ$ and returns the lexicographically next one using the standard definition of binary incrementation, while $\exists^{w,k}$ implements existential quantification for $\tau_{w,k}^\circ$ by iterating through all possible candidates using $\mathbf{next}^{w,k}$.

Lemma 12. *Let T be an LTS and let η be good. Then $T, \bar{s} \models_\eta \exists^{w,k} x. \varphi$ iff there is $f \in \llbracket \tau_{w,k}^\circ \rrbracket^T$ such that $T, \bar{s} \models_{\eta[x \rightarrow f]} \varphi$.*

The proof follows the same pattern as that of Lemma 11 by verifying that the individual formulas do what is claimed above.

Lemmas 11 and 12 justify the use of HO(LFP)-style quantification symbols for individual states, sets of type $M \times S^{d-w} \cap \eta(e)$ for $M \subseteq S^s$ and for $\tau_{w,k}^\circ$ for all $k \geq 2$. Note that the latter do neither coincide with HO(LFP) types nor with the respective $\tau_{w,k}$.

5 Lower Bounds

Homogeneous Types. In order to simplify the translation from HO(LFP) to PHFL, we restrict the set of types that can be used in HO(LFP) formulas. Let $w \geq 2$ be fixed but arbitrary. Define $\tau'_{w,k}$ as $\tau'_{w,1} = \odot$, $\tau'_{w,i+1} = (\tau'_{w,i}, \dots, \tau'_{w,i})$ with w many repetitions of $\tau'_{w,i}$.

Lemma 13. *If $\varphi \in \text{HO}^k(\text{LFP})$ defines a query Q , then there is $\varphi' \in \text{HO}^k(\text{LFP})$ that defines the same query, but the only types used in φ' are $\tau'_{w,0}, \dots, \tau'_{w,k}$ for some w .*

Proof. There are two principles that are used here: If φ contains a type of the form $\tau'' = (\tau', \dots, \tau')$ with $w' \leq w$ many repetitions of w , we can replace τ'' by (τ', \dots, τ') with exactly w many repetitions of τ' by changing the respective type everywhere in the formula and requiring at quantifiers that, e.g., the last $w - w'$ components are equal to the $w - w' - 1$ st in every tuple contained in a set. Hence, every type can be assumed to have width exactly w .

It remains to deal with inhomogeneous types, e.g. those of the form (τ', τ'', \dots) such that $\text{ord}(\tau') \neq \text{ord}(\tau'')$. This can be remedied by increasing the type of the lower order by one order, and requiring at quantification steps that the only tuple in the set be a singleton of the form (M, \dots, M) of width w . This procedure needs to be chained if the orders of constituent types diverge by more than one. \square

Type Correspondence. Now that we can assume w.l.o.g. that all $\text{HO}^k(\text{LFP})$ -definable queries are defined by an $\text{HO}^k(\text{LFP})$ formula which uses only the types $\tau'_{w,i}$ for $i \leq k$ we can observe that these types are quite similar to the types $\tau_{w,k}$ defined in the previous section. In fact, we want to emulate the type $\tau'_{w,k}$ with $k > 1$ by $\tau_{w,k-2}^\circ$. The type $\tau'_{s,1}$ will be handled by polyadicity as in [17].²

Let T be an LTS with state set S and let $d \geq w \geq 2$. For each $k \geq 2$, we define a translation $\text{tptr}_k^T : \llbracket \tau'_{w,k} \rrbracket^T \rightarrow \llbracket \tau_{w,k-2}^\circ \rrbracket^T$ via $\text{tptr}_2^T(M) = M \times S^{d-w}$ and $\text{tptr}_{i+1}^T(M) = f \in \llbracket \tau_{w,k-2} \rrbracket^T$ s.t.

$$f(f_1, \dots, f_n) = \begin{cases} S, & \text{if } f_j = \text{tptr}_i^T(x_i) \text{ f.a. } 0 \leq j \leq w \text{ and } (x_1, \dots, x_w) \in M \\ \emptyset, & \text{otherwise.} \end{cases}$$

Hence, a variable assignment α such that all variables are of types $\tau'_{w,2}, \dots, \tau'_{w,k}$ induces an environment η_α with definitions for all those variables of order ≥ 2 via $\eta_\alpha(X) = \text{tptr}_i^T(X)$ where i is the order of X .

The Translation. We are now ready to extend the translation given in [17] to $\text{HO}^k(\text{LFP})$ with $k \geq 2$. Towards this, we present a syntactical translation trans from $\text{HO}^k(\text{LFP})$ with $k \geq 2$ into PHFL_{k-1}^d for some $d \geq 2$.

Lemma 14. *Let $\varphi \in \text{HO}^k(\text{LFP})$ be bisimulation invariant and have free first-order variables X_1, \dots, X_r and, hence define a query of width r . Moreover, let $\tau'_{w,0}, \dots, \tau'_{w,k}$ be the only types in φ . W.l.o.g. let $2w \geq r$. Let $d = 2w + r + 2$. Then there is $\varphi' \in \text{PHFL}_{k-1}^d$ such that, for all LTS T that are a bisimulation quotient, we have $T, \alpha \models \varphi$ iff $T, (\alpha(X_1), \dots, \alpha(X_r)) \models_{\eta_\alpha[e \mapsto M]} \varphi'$ where $\emptyset \neq M' \subseteq S^r$ and $M = M' \times S^{d-2r-2} \times \{\alpha(X_1)\} \times \dots \times \{\alpha(X_r)\} \times S^2$.*

²It would also be possible to completely eliminate this type from $\text{HO}^k(\text{LFP})$ for $k \geq 2$; cf. similar constructions in the context of MSO and automata theory.

Proof. Let *trans* be given via

$$\begin{aligned}
\text{trans}(p(X_i)) &:= p_i \\
\text{trans}(a(X_i, X_j)) &:= \langle a_i \rangle \{ \sigma_{i,j} \} \varphi_{\sim} \\
\text{trans}(\varphi_1 \vee \varphi_2) &:= \text{trans}(\varphi_1) \vee \text{trans}(\varphi_2) \\
\text{trans}(\neg \varphi) &:= \neg \text{trans}(\varphi) \\
\text{trans}(\exists(X_i : \odot). \varphi) &:= \exists_i. \text{trans}(\varphi) \\
\text{trans}(\exists(X : \tau'_{w,1}). \varphi) &:= \exists^{(w)}(x). \text{trans}(\varphi) \\
\text{trans}(\exists(X : \tau'_{w,k}). \varphi) &:= \exists^{w,k}(x). \text{trans}(\varphi) \text{ if } k \geq 2 \\
\text{trans}(X(X_{i_1}, \dots, X_{i_w})) &:= \{ \sigma \} x \text{ if } (X : \tau'_{w,1}) \\
\text{trans}(X(X_1, \dots, X_w)) &:= x(x_1, \dots, x_w) \text{ if } (X : \tau'_{w,k}) \text{ and } k \geq 2 \\
\text{trans}(\text{lfp}(X, Y_1, \dots, Y_w). \varphi)(Z_{i_1}, \dots, Z_{i_w}) &:= \{ \sigma \} \mu(F : \bullet \rightarrow \bullet). \text{trans}(\varphi) \text{ if } (X : \tau_{w,2}) \\
\text{trans}(\text{lfp}(X, Y_1, \dots, Y_w). \varphi)(Z_1, \dots, Z_w) &:= (\mu(F : \tau_{w,k-2} \rightarrow \tau_{w,k-2}). \lambda(y_1, \dots, y_w : \tau_{w,k-3}). \\
&\quad \text{trans}(\varphi))(Z_1, \dots, Z_w) \text{ if } (X : \tau_{w,k}) \text{ and } k \geq 3
\end{aligned}$$

where $\sigma_{i,j}$ is defined via $\sigma_{i,j}(d-1) = i$, $\sigma_{i,j}(d) = j$ and $\sigma_{i,j}(k) = k$ for $k < d-1$, and φ_{\sim} is the formula from Ex. 2, resp. [17] that holds iff the states in positions $d-1$ and d are bisimilar, and, finally σ is defined via $\sigma(j) = i_j$ for $0 \leq j \leq w$ and $\sigma(k) = k$ for $k \geq w$. The substitutions $\sigma_{i,j}$, resp. σ serve to swap the elements i and j to positions $d-1$ and d , resp. to reorder the elements according to the variable order used on the left side of the translation.

The claim now follows by induction over the syntax tree of φ . The first five cases are as in [17]. The next four cases are by Lemmas 11 and 12, resp. the definition of η_α . The fixpoint cases are by an induction over the individual stages of the fixpoint iteration; note that the Knaster-Tarski-based semantics of the fixpoints in HO(LFP) and PHFL can be equivalently replaced by semantics based on the Kleene Fixpoint Theorem [10]. This mirrors the proof of the fixpoint case in [17]. \square

Theorem 15. *PHFL^k captures k-EXPTIME/_~ for all $k \geq 0$.*

Proof. The cases of $k = 0$ and $k = 1$ are of course already known [17, 15]. For $k \geq 2$, the upper bound is shown in Thm. 8. For the lower bound, it suffices, due to Prop. 4, to show that for any bisimulation-invariant query defined by an HO^{k+1}(LFP) formula there is an equivalent formula in PHFL^k. Let φ be a formula defining such a query Q of width $r \geq 1$, and w.l.o.g. φ contains only the types $\tau_{w,0}, \dots, \tau_{w,k}$.

Let φ' be the formula obtained from the translation in Lemma 14 and let σ be defined via $\sigma(d-r-2+i) = i$ for $1 \leq i \leq r$, and $\sigma(j) = j$ for $j \leq d-r-2$ or $j \geq d-1$. Note that σ simply copies the first r many states in a tuple to the positions $d-r-1, \dots, d-2$ where they serve to retain goodness (cf. Def. 10). Then $\psi = (\lambda(e : \bullet)\varphi)\{\sigma\}\text{tt}$ defines the query $\{(T, (s_1, \dots, s_d)) \mid (T, (s_1, \dots, s_r)) \in Q\}$. Towards this, note that $\llbracket \{\sigma\}\text{tt} \rrbracket^T = \{(s_1, \dots, s_d) \mid s_i = s_{d-r-2+i} \text{ for } 1 \leq i \leq r\} = M$. Hence, $\llbracket \psi \rrbracket^T = \llbracket \varphi' \rrbracket_\eta^T$ where $\eta = \emptyset[e \mapsto M]$. Since η is good, Lemma 14 is applicable.

What remains is to argue that a query that, on each LTS T , returns all tuples in $M \times S^{d-r}$ where $M \subseteq S^r$, is in fact a query that defines tuples of width r . This can either be done formally by expanding the semantics of the substitution operator $\{\sigma\}$ to return tuples in the width of its co-domain, or by simply considering queries of the above form to be of width r , cf. [17]. \square

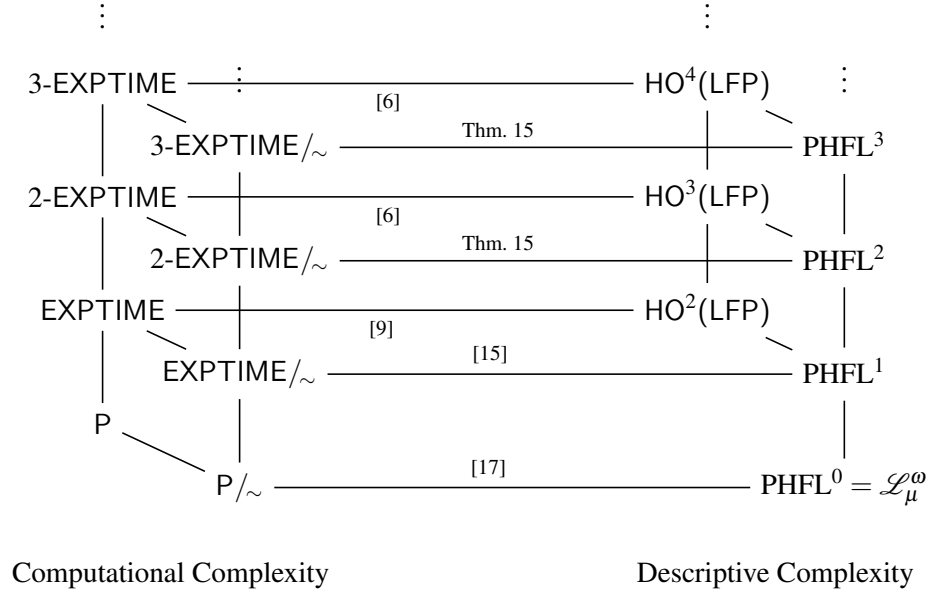


Figure 2: Capturing results for time complexity classes - an overview.

6 Conclusion

We have extended the descriptive complexity of bisimulation-invariant queries, as started by Otto [17] with the capturing of the class P/\sim with the logic \mathcal{L}_μ^ω , to further time complexity classes, namely the classes $k\text{-EXPTIME}/\sim$ of bisimulation-invariant queries that can be answered in k -fold exponential time. These turn out to be exactly those that can be defined in PHFL, the polyadic extension of the higher-order modal fixpoint logic HFL. This is genuinely an extension as \mathcal{L}_μ^ω coincides with PHFL^0 , the fragment with no higher-order constructs. Moreover, the level in the exponential-time hierarchy corresponds to the type order used in formulas: $k\text{-EXPTIME}/\sim = \text{PHFL}^k$.

The resulting picture of descriptive time complexity is shown in Fig. 2. A natural way to extend this is of course to integrate results about the space complexity classes $k\text{-EXPSPACE}/\sim$ sitting in between $k\text{-EXPTIME}/\sim$ and $(k+1)\text{-EXPTIME}/\sim$. Here we can only state that it is possible to capture these in terms of natural fragments of PHFL as well but the proof requires a few more ingredients than the time-complexity case. Here we could use the already known characterisation of $k\text{-EXPTIME}$ by Higher-Order Predicate Logics with Least Fixpoints [6]. For the space complexity classes, we need to first develop such characterisations that extend the Abiteboul-Vianu Theorem. This can be done [13], but due to space restrictions here its detailed presentation is left for a future publication.

A noteworthy observation for the capturing of $k\text{-EXPTIME}/\sim$ with $k \geq 1$ is that, unlike in the case of $k = 0$, the requirement of structures being ordered is not necessary for the invocation of the (generalised) Immerman-Vardi Theorem. However, starting from $k \geq 1$ the order being provided in the bisimulation-invariant setting plays a crucial role to emulate existential quantification from the HO(LFP) side: the LTS in question coming with a PHFL_2^0 -definable order allows us to enumerate sufficiently many sets, functions, etc. to emulate existential quantification. It remains to see in further detail whether this is an artifact of the proof strategy or a genuine and inherent part of the correspondence between complexity classes and logics in the bisimulation-invariant framework.

References

- [1] S. Abiteboul & V. Vianu (1987): *A Transaction Language Complete for Database Update and Specification*. In: *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, San Diego, CA, pp. 260–268, doi:10.1145/28659.28688.
- [2] H. R. Andersen (1994): *A Polyadic Modal μ -Calculus*. Technical Report ID-TR: 1994-195, Dept. of Computer Science, Technical University of Denmark, Copenhagen, doi:10.1.1.42.1859.
- [3] R. Axelsson, M. Lange & R. Somla (2007): *The Complexity of Model Checking Higher-Order Fixpoint Logic*. *Logical Methods in Computer Science* 3, pp. 1–33, doi:10.2168/LMCS-3(2:7)2007.
- [4] F. Bruse, M. Lange & É. Lozes (2017): *Space-Efficient Fragments of Higher-Order Fixpoint Logic*. In: *Proc. 11th Int. Workshop on Reachability Problems, RP'17, LNCS 10506*, pp. 26–41, doi:10.1007/978-3-319-67089-8_3.
- [5] R. Fagin (1974): *Generalized First-Order Spectra and Polynomial-Time Recognizable Sets*. *Complexity and Computation* 7, pp. 43–73.
- [6] C. M. Freire & A. T. Martins (2011): *The Descriptive Complexity of the Deterministic Exponential Time Hierarchy*. In: *Proc. 5th Workshop on Logical and Semantic Frameworks with Applications, LSFA'10*, 269, pp. 71–82, doi:10.1016/j.entcs.2011.03.006.
- [7] M. Grohe (2008): *The Quest for a Logic Capturing PTIME*. In: *Proc. 23rd Annual IEEE Symp. on Logic in Computer Science, LICS'08*, IEEE Computer Society, pp. 267–271, doi:10.1109/LICS.2008.11.
- [8] N. Immerman (1986): *Relational Queries Computable in Polynomial Time*. *Information and Control* 68(1-3), pp. 86–104, doi:10.1016/S0019-9958(86)80029-8.
- [9] N. Immerman (1987): *Languages That Capture Complexity Classes*. *SIAM Journal of Computing* 16(4), pp. 760–778, doi:10.1137/0216051.
- [10] S. C. Kleene (1938): *On Notation for Ordinal Numbers*. *Journal of Symbolic Logic* 3(4), pp. 150–155, doi:10.2307/2267778.
- [11] B. Knaster (1928): *Un théorème sur les fonctions d'ensembles*. *Annals Soc. Pol. Math* 6, pp. 133–134.
- [12] D. Kozen (1983): *Results on the Propositional μ -calculus*. *Theor. Comp. Sci.* 27, pp. 333–354, doi:10.1016/0304-3975(82)90125-6.
- [13] D. Kronenberger (2018): *Capturing Bisimulation-Invariant Complexity Classes by Polyadic Higher-Order Fixpoint Logic*. Master's thesis, University of Kassel.
- [14] M. Lange & É. Lozes (2012): *Model Checking the Higher-Dimensional Modal μ -Calculus*. In: *Proc. 8th Workshop on Fixpoints in Computer Science, FICS'12, Electr. Proc. in Theor. Comp. Sc.* 77, pp. 39–46, doi:10.4204/EPTCS.77.
- [15] M. Lange & É. Lozes (2014): *Capturing Bisimulation-Invariant Complexity Classes with Higher-Order Modal Fixpoint Logic*. In: *Proc. 8th Int. IFIP Conf. on Theoretical Computer Science, TCS'14, LNCS 8705*, Springer, pp. 90–103, doi:10.1007/978-3-662-44602-7.
- [16] M. Lange, É. Lozes & M. Vargas Guzmán (2014): *Model-Checking Process Equivalences*. *Theor. Comp. Sci.* 560, pp. 326–347, doi:10.1016/j.tcs.2014.08.020.
- [17] M. Otto (1999): *Bisimulation-invariant PTIME and higher-dimensional μ -calculus*. *Theor. Comput. Sci.* 224(1-2), pp. 237–265, doi:10.1016/S0304-3975(98)00314-4.
- [18] A. Tarski (1955): *A Lattice-theoretical Fixpoint Theorem and its Application*. *Pacific Journal of Mathematics* 5, pp. 285–309, doi:10.2140/pjm.1955.5.285.
- [19] J. van Benthem & K. Doets (1983): *Higher-Order Logic*. In D. Gabbay & F. Guenther, editors: *Handbook of Philosophical Logic, Volume I: Elements of Classical Logic*, D. Reidel Publishing Co., pp. 275–329, doi:10.1007/978-94-009-7066-3_4.

- [20] M. Y. Vardi (1982): *The Complexity of Relational Query Languages (Extended Abstract)*. In: *Proc. 14th Symp. on Theory of Computing, STOC'82*, ACM, San Francisco, CA, USA, pp. 137–146, doi:10.1145/800070.
- [21] M. Viswanathan & R. Viswanathan (2004): *A Higher Order Modal Fixed Point Logic*. In: *Proc. 15th Int. Conf. on Concurrency Theory, CONCUR'04, LNCS 3170*, Springer, pp. 512–528, doi:10.1007/978-3-540-28644-8_33.