

The Arity Hierarchy in the Polyadic μ -Calculus

Martin Lange

School of Electrical Engineering and Computer Science, University of Kassel, Germany

The polyadic μ -calculus is a modal fixpoint logic whose formulas define relations of nodes rather than just sets in labelled transition systems. It can express exactly the polynomial-time computable and bisimulation-invariant queries on finite graphs. In this paper we show a hierarchy result with respect to expressive power inside the polyadic μ -calculus: for every level of fixpoint alternation, greater arity of relations gives rise to higher expressive power. The proof uses a diagonalisation argument.

1 Introduction

The modal μ -calculus \mathcal{L}_μ is a well-studied logic [14, 4, 5], obtained by adding restricted second-order quantification in the form of least and greatest fixpoints to a multi-modal logic interpreted over labelled transition systems. A formula of the modal μ -calculus is thus interpreted in a state of such transition systems which means that such formulas *define* sets of states in transition systems. For example, $\nu X.\mu Y.\langle a \rangle X \vee \langle b \rangle Y$ defines the set of all states from which there is a path with labels ‘ a ’ and ‘ b ’ that contains infinitely many occurrences of the symbol ‘ a ’.

The polyadic μ -calculus \mathcal{L}_μ^ω is a much less known extension of the modal μ -calculus whose formulas define *relations* rather than sets of states. They are interpreted in a tuple of states rather than a single state, and there are modal operators for each position in this tuple. Thus, one states “the third state has an ‘ a ’-successors” for instance rather than just “there is an ‘ a ’-successors.” Combining such simple modal statements with fixpoint quantifiers yields an expressive logic with interesting applications: the polyadic μ -calculus was first defined by Andersen [2] and used as a logic for defining process equivalences like bisimilarity [15, 16]. Later it was re-invented by Otto under the name *Higher-Dimensional μ -Calculus* [19] and shown to capture the complexity class P over bisimulation-invariant class of finite graphs. I.e. a bisimulation-invariant property of finite graphs can be computed in polynomial time iff it is definable in \mathcal{L}_μ^ω .

There is a natural hierarchy in \mathcal{L}_μ^ω given by fragments of bounded arity. The polyadic μ -calculus itself can be seen as a fragment of FO+LFP, i.e. First-Order Logic extended with fixpoint quantifiers. The translation naturally extends the standard translation of modal logic into first-order formulas with one free variable, seen as the point of reference for the interpretation of the property expressed by the modal formula. Polyadic formulas get interpreted in tuples of states, hence they can be seen as special first-order formulas with several free variables. The arity of a polyadic formula is then the minimal number of free variables needed to express this property in FO+LFP or, equivalently, the length of the tuples used to interpret the formula.

The aim of this article is to show that the hierarchy formed by fragments of bounded arity, denoted $\mathcal{L}_\mu^1, \mathcal{L}_\mu^2, \dots$ is strict. This is not too surprising when taken literally: clearly, *any* satisfiable but non-valid formula in \mathcal{L}_μ^{k+1} is not equivalent to any formula in \mathcal{L}_μ^k since the former get interpreted in $k+1$ -tuples and

*The European Research Council has provided financial support under the European Community’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 259267.

the latter only in k -tuples. We therefore need to employ a convention that allows different fragments to be compared with respect to expressive power and still yields a meaningful hierarchy result. We consider formulas that are interpreted in a single state at the top-level, regardless of their arity. I.e. we show that for every $k \geq 1$ there is a \mathcal{L}_μ^{k+1} -formula Φ_{k+1} such that there is no \mathcal{L}_μ^k -formula ψ which yields

$$\mathcal{T}, \underbrace{(s, \dots, s)}_{k \text{ times}} \models \psi \quad \text{iff} \quad \mathcal{T}, \underbrace{(s, \dots, s)}_{k+1 \text{ times}} \models \Phi_{k+1}$$

for all labelled transition systems \mathcal{T} and all their states s .

Arity hierarchies have been studied before, most notably by Grohe for fixpoint extensions of first-order logic including FO+LFP [10]. Even though each \mathcal{L}_μ^k can be embedded into FO+LFP, the arity hierarchy in \mathcal{L}_μ^ω does not follow immediately from the one in FO+LFP. Grohe constructs formulas of arity $k+1$ in FO+LFP – they belong to the smaller FO+TC already – and shows that they are not equivalent to any formulas of arity k in FO+LFP – not even the much larger FO+sPFP. However, these witnessing formulas are not bisimulation-invariant since they express a relation formed by the transitive closure of a clique relation and being a clique is clearly not bisimulation-invariant. Hence, these witnessing formulas are not expressible in \mathcal{L}_μ^ω and therefore the arity hierarchy is not transferred immediately.

It could of course be checked whether the proof used to show the arity hierarchy in FO+LFP could be adapted to work for \mathcal{L}_μ^ω as well. It would require the search for a similar witnessing property and the adaption of the Ehrenfeucht-Fraïssé argument to the polyadic μ -calculus. Such model comparison games exist for the modal μ -calculus [21] but using them to obtain inexpressibility results has proved to be quite difficult.

Instead we use a simple diagonalisation argument in order to obtain a strictness result regarding arity hierarchies. A k -ary formula φ can be seen syntactically as a labelled transition system \mathcal{T}_φ , roughly based on the syntax-tree representation. We can then define a $k+1$ -ary formula that simulates the evaluation of φ on \mathcal{T}_φ and accepts those \mathcal{T}_φ which are not accepted by φ itself. Hence, we need to find a generic way of dualising the operators in φ . This is no particular problem, for instance, when one sees a disjunction then one needs to check *both* disjuncts, for a conjunction one only needs to check one of them. However, fixpoint formulas may hold or not because of infinite recursive unfoldings through fixpoint operators. This needs to be dualised as well, and the only way that we can see to do this is to equip the simulating formula with a fixpoint structure that is at least as rich as the one of the simulated formula. Consequently, we obtain an arity hierarchy relative to the alternation hierarchy. This does not happen for extensions of First-Order Logic since it is known that there is no alternation hierarchy: every FO+LFP formula can be expressed with a single least fixpoint operator only [12, 23]. The situation for modal logics is different: more fixpoint alternation generally gives higher expressive power, at least so in the modal μ -calculus [6], and presumably then so in \mathcal{L}_μ^ω as well.

The rest of this paper is organised as follows. In Section 2 we recall the polyadic μ -calculus and necessary tools like fixpoint alternation and model checking games. In Section 3 we prove the hierarchy results, and in Section 4 we conclude with a discussion on further work.

2 The Polyadic μ -Calculus

Labelled Transition Systems. Let $\text{Prop} = \{p, q, \dots\}$ and $\text{Act} = \{a, b, \dots\}$ be two fixed, countably infinite sets of atomic propositions and action names. A labeled transition system (LTS) over Prop and Act is a tuple $\mathcal{T} = (S, \rightarrow, \lambda, s_I)$ where S is a set of states, $\rightarrow \subseteq S \times \text{Act} \times S$ is the transition relation, $\lambda : S \rightarrow 2^{\text{Prop}}$ labels the states with atomic propositions, and s_I is some designated starting state. We will write $s \xrightarrow{a} t$ instead of $(s, a, t) \in \rightarrow$.

The Syntax of \mathcal{L}_μ^ω . Let $\text{Var} = \{X, Y, \dots\}$ be an infinite set of second-order variables. The syntax of the polyadic modal μ -calculus \mathcal{L}_μ^ω is similar to that of the ordinary modal μ -calculus. However, modalities and propositions are relativised to a natural number pointing at a position in a tuple of states used to interpret the formula.

A *replacement* is a $\kappa : \mathbb{N} \rightarrow \mathbb{N}$ which acts like the identity function on almost all arguments. We write $\mathbb{N} \dashrightarrow \mathbb{N}$ to denote the space of all replacements. Such a replacement is then written as $\{\kappa(i_1) \leftarrow i_1, \dots, \kappa(i_m) \leftarrow i_m\}$ when $i_1 < \dots < i_m$ are all those indices for which we have $\kappa(i_j) \neq i_j$. We will sometimes allow ourselves to deviate from this and to use some shorter but equally intuitive notation for such functions. For instance $\{1 \leftrightarrow 2\}$ should denote the swap between 1 and 2, i.e. it abbreviates $\{2 \leftarrow 1, 1 \leftarrow 2\}$.

For technical convenience, we define the logic directly in positive normal form. Formulas are then given by the grammar

$$\varphi ::= p(i) \mid \neg p(i) \mid X \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \langle a \rangle_i \varphi \mid [a]_i \varphi \mid \mu X. \varphi \mid \nu X. \varphi \mid \kappa \varphi$$

where $p \in \text{Prop}$, $a \in \text{Act}$, $1 \geq i \in \mathbb{N}$ and κ is a replacement. We require that every second-order variable gets bound by a unique fixpoint quantifier μ or ν . Then for every formula φ there is a function fp_φ which maps each second-order variable X occurring in φ to its unique binding formula $fp_\varphi(X) = \eta X. \psi$.

The set $\text{Sub}(\varphi)$ of subformulas of φ is defined as usual, with $\text{Sub}(\mu X. \varphi) = \{\mu X. \varphi\} \cup \text{Sub}(\varphi)$ for instance.

Later we will use the abbreviation $\ell \rightarrow \varphi$ when ℓ is a literal $q(i)$ or $\neg q(i)$. This behaves like ordinary implication – note that we have defined the logic in positive normal form and can therefore not simply introduce implication via negation – for such formulas when seen as $\bar{\ell} \vee \varphi$ where $\bar{\ell}$ is the usual complementary literal to ℓ .

The *arity* of a formula φ , denoted $ar(\varphi)$ is the largest index i occurring in the operators $p(i)$, $\langle a \rangle_i$, $[a]_i$ and $\{\kappa\}$ in any of its subformulas. The fragment of arity k is $\mathcal{L}_\mu^k := \{\varphi \mid ar(\varphi) \leq k\}$. Hence, $\varphi := \nu X. \langle a \rangle_1 \{2 \leftrightarrow 1\} X$ has arity 2 and it therefore belongs to all fragments \mathcal{L}_μ^2 , \mathcal{L}_μ^3 , etc., because it defines a relation of arity 2 which can also be seen as a relation of higher arity in which the 3rd, 4th, etc. components of its tuples are simply unrestrained.

The Semantics of \mathcal{L}_μ^ω . Formulas of \mathcal{L}_μ^k are interpreted in k -tuples of states of a transition system $\mathcal{T} = (S, \rightarrow, \lambda, s_I)$. An interpretation $\rho : \text{Var} \rightarrow 2^{S^k}$ is needed in order to define this inductively and give a meaning to formulas with free variables. For each \mathcal{L}_μ^k -formula φ , $\llbracket \varphi \rrbracket_\rho^\mathcal{T}$ is a k -ary relation of states in \mathcal{T} , namely the relation defined by φ under the assumption that its free variables are interpreted by ρ .

$$\begin{aligned} \llbracket p(i) \rrbracket_\rho^\mathcal{T} &:= \{(s_1, \dots, s_k) \mid p \in \lambda(s_i)\} \\ \llbracket \neg p(i) \rrbracket_\rho^\mathcal{T} &:= \{(s_1, \dots, s_k) \mid p \notin \lambda(s_i)\} \\ \llbracket X \rrbracket_\rho^\mathcal{T} &:= \rho(X) \\ \llbracket \varphi \vee \psi \rrbracket_\rho^\mathcal{T} &:= \llbracket \varphi \rrbracket_\rho^\mathcal{T} \cup \llbracket \psi \rrbracket_\rho^\mathcal{T} \\ \llbracket \varphi \wedge \psi \rrbracket_\rho^\mathcal{T} &:= \llbracket \varphi \rrbracket_\rho^\mathcal{T} \cap \llbracket \psi \rrbracket_\rho^\mathcal{T} \\ \llbracket \langle a \rangle_i \varphi \rrbracket_\rho^\mathcal{T} &:= \{(s_1, \dots, s_k) \mid \exists t \text{ s.t. } s_i \xrightarrow{a} t \text{ and } (s_1, \dots, s_{i-1}, t, s_{i+1}, \dots, s_k) \in \llbracket \varphi \rrbracket_\rho^\mathcal{T}\} \\ \llbracket [a]_i \varphi \rrbracket_\rho^\mathcal{T} &:= \{(s_1, \dots, s_k) \mid \forall t : \text{ if } s_i \xrightarrow{a} t \text{ then } (s_1, \dots, s_{i-1}, t, s_{i+1}, \dots, s_k) \in \llbracket \varphi \rrbracket_\rho^\mathcal{T}\} \\ \llbracket \mu X. \varphi \rrbracket_\rho^\mathcal{T} &:= \bigcap \{R \subseteq S^k \mid \llbracket \varphi \rrbracket_{\rho[X \mapsto R]}^\mathcal{T} \subseteq R\} \\ \llbracket \nu X. \varphi \rrbracket_\rho^\mathcal{T} &:= \bigcup \{R \subseteq S^k \mid \llbracket \varphi \rrbracket_{\rho[X \mapsto R]}^\mathcal{T} \supseteq R\} \\ \llbracket \kappa \varphi \rrbracket_\rho^\mathcal{T} &:= \{(s_{\kappa(1)}, \dots, s_{\kappa(k)}) \mid (s_1, \dots, s_k) \in \llbracket \varphi \rrbracket_\rho^\mathcal{T}\} \end{aligned}$$

Note that the partial order \subseteq makes S^k a complete lattice with meets and joins given by \cap and \cup , and the semantics of fixpoint formulas is then well-defined according to the Knaster-Tarski Theorem [13, 22].

We also write $\mathcal{T}, s_1, \dots, s_k \models_\rho \varphi$ instead of $(s_1, \dots, s_k) \in \llbracket \varphi \rrbracket_\rho^\mathcal{T}$. If φ has no free second-order variables then we also drop ρ . In Section 3 we will often consider situations with tuples of the form (s, \dots, s) of some length k derivable from the context. We will then simply write $\mathcal{T}, s \models \varphi$ as a short form for $\mathcal{T}, s, \dots, s \models \varphi$.

Two formulas $\varphi, \psi \in \mathcal{L}_\mu^k$ are *equivalent*, written $\varphi \equiv \psi$, if $\llbracket \varphi \rrbracket_\rho^\mathcal{T} = \llbracket \psi \rrbracket_\rho^\mathcal{T}$ for any \mathcal{T} and corresponding variable interpretation ρ . Note that two formulas can be equivalent even if they do not belong to the same arity fragment: if $\varphi \in \mathcal{L}_\mu^k$ and $\psi \in \mathcal{L}_\mu^{k'}$ and $k \neq k'$ then $\varphi, \psi \in \mathcal{L}_\mu^{\max\{k, k'\}}$, i.e. we can interpret the one of smaller arity as a formula of larger arity that simply does not constrain the additional elements in the tuples of the relation it defines.

Examples. The standard example of a \mathcal{L}_μ^ω -formula, indeed a \mathcal{L}_μ^2 -formula, is the one defining *bisimilarity*.

$$\varphi_\sim := \nu X. \left(\bigwedge_{p \in \text{Prop}} p(1) \rightarrow p(2) \right) \wedge \left(\bigwedge_{a \in \text{Act}} [a]_1 \langle a \rangle_2 X \right) \wedge \{1 \leftrightarrow 2\} X$$

It is indeed the case that $\mathcal{T}, s, t \models \varphi_\sim$ iff $s \sim t$, i.e. s and t are bisimilar in \mathcal{T} .

As a second example consider an \mathcal{T} with an edge relation *flight* and two atomic propositions *warm* and *safe*. When seeing the nodes of the LTS as cities (which can or cannot be warm and/or safe and are potentially linked by direct flight connections), then

$$\{3 \leftarrow 1\} \varphi_\sim \wedge \langle \text{flight} \rangle_2 \mu X. \text{warm}(2) \wedge \text{safe}(2) \wedge \langle \text{flight} \rangle_1 \varphi_\sim \wedge (\{3 \leftarrow 1\} \varphi_\sim \vee [\text{flight}]_2 X) \wedge \{2 \leftarrow 3\} X$$

yields all triples (s, t, u) of cities such that there is a roundtrip from t which only traverses through warm and safe cities that can be reached from city s in one step – in case someone in s wants to come and visit – such that the trip can be traversed in both directions. This description of course uses equality (“roundtrip”) on cities which is not available in the logic. Instead we use bisimilarity in the formula, so for instance “roundtrip from t ” is to be understood as a trip starting in t and ending in a city that cannot be distinguished from t with the means of bisimilarity.

Fixpoint Alternation. The proof of the arity hierarchy carried out in Section 3 needs a closer look at the dependencies of fixpoints inside a formula. This phenomenon is well-understood leading to the notion of alternation hierarchy [9, 18]. We give a brief introduction to fixpoint alternation that is sufficient for the purposes of the next section.

Let $k \geq 1$ and $\varphi \in \mathcal{L}_\mu^k$ be fixed. For two variables $X, Y \in \text{Sub}(\varphi)$ we write $X \geq_\varphi Y$ if X has a free occurrence in $\text{fp}_\varphi(Y)$. We use $>_\varphi$ to denote the strict part of its transitive closure. E.g. in

$$\varphi := \mu X. p(2) \vee \langle b \rangle_1 (\nu Y. q(1) \wedge \nu Y'. (\mu Z. Y' \vee \langle a \rangle_1 Z) \wedge [b]_2 Y)$$

we have $X >_\varphi Y >_\varphi Y' >_\varphi Z$ even though there is no free occurrence of X in the fixpoint formula for Z .

Names of variables do not matter but their fixpoint types do. So we abstract this chain of fixpoint dependencies into a chain $\mu >_\varphi \nu >_\varphi \nu >_\varphi \mu$. The *alternation type* of a formula is a maximal descending chain of variables (represented by their fixpoint types) such that adjacent types in this chain are different. The alternation type of φ above is therefore just (μ, ν, μ) . We then define the *alternation hierarchy* as follows: Σ_m^k , respectively Π_m^k consists of all formulas of arity k and alternation type of length at most m such that the m -th last in this chain is μ , respectively ν , if it exists. For instance, the formula φ above belongs to Σ_3^2 and therefore also to Σ_m^2 and Π_m^2 for all $m > 3$. It does not belong to Π_2^2 .

Each variable X occurring in φ is also given an *alternation depth* $ad_\varphi(X)$. It is the index in a maximal chain of dependencies $X_m >_\varphi \dots >_\varphi X_1$ such that adjacent variables have different fixpoint types. E.g. in the example above we have $ad_\varphi(X) = 3$, $ad_\varphi(Y) = ad_\varphi(Y') = 2$ and $ad_\varphi(Z) = 1$.

The next observation is easy to see.

Lemma 1. *Let $\varphi \in \Sigma_m^k$ and $X \in \text{Sub}(\varphi)$ be one of its fixpoint variables. Then the fixpoint type of X is uniquely determined by $ad_\varphi(X)$, namely it is μ if m and i are both odd or both even, otherwise it is ν .*

Model Checking Games. We briefly recall model checking games for the polyadic μ -calculus [15]. They are defined in the same style as the model checking games for the modal μ -calculus [20] as a game played between players VERIFIER and REFUTER on the product space of an LTS and a formula. Such games can be used to reason about the satisfaction of a formula in a structure since both satisfaction and non-satisfaction are reduced to the existence of winning strategies for one of the players in these model checking games.

As with the modal μ -calculus games, the model checking games for the polyadic μ -calculus are nothing more than parity games. However, they are played using k pebbles in the LTS and one pebble on the set of subformulas of the input formula. Hence, a configuration is a $k+1$ -tuple written $s_1, \dots, s_k \vdash \psi$ where the s_i are states of the underlying LTS $\mathcal{T} = (S, \rightarrow, \lambda, s_I)$ and ψ is a subformula of the underlying formula φ .

The rules are as follows.

- In a configuration of the form $s_1, \dots, s_k \vdash \psi_1 \vee \psi_2$, player VERIFIER chooses an $i \in \{1, 2\}$ and the play continues with $s_1, \dots, s_k \vdash \psi_i$. Intuitively, VERIFIER moves the formula pebble to a disjunct from the current disjunction.
- Likewise, in a configuration of the form $s_1, \dots, s_k \vdash \psi_1 \wedge \psi_2$, player REFUTER chooses such an i . Here, this can be seen as refuter moving the formula pebble.
- In a configuration of the form $s_1, \dots, s_k \vdash \langle a \rangle_i \psi$, player VERIFIER chooses a t such that $s_i \xrightarrow{a} t$ and the play continues with $s_1, \dots, s_{i-1}, t, s_{i+1}, \dots, s_k \vdash \psi$. Intuitively, VERIFIER moves the i -th state pebble along an outgoing a -transition. The other $k-1$ pebbles that are on states remain where they are. The formula pebble is also moved to the next subformula.
- Likewise, in a configuration of the form $s_1, \dots, s_k \vdash [a]_i \psi$, player REFUTER chooses such a t .
- In a configuration of the form $s_1, \dots, s_k \vdash \eta X.\psi$ or $s_1, \dots, s_k \vdash X$ such that $fp_\varphi(X) = \eta X.\psi$, the formula pebble is simply moved to ψ , i.e. the play continues with $s_1, \dots, s_k \vdash \psi$.

A player wins a play if the opponent cannot carry out a move anymore. Moreover, VERIFIER wins a play that reaches a configuration of the form $s_1, \dots, s_k \vdash q(i)$ if $q \in \lambda(s_i)$. If, on the other hand, $q \notin \lambda(s_i)$ then player REFUTER wins this play. Finally, there are infinite plays, and the winner is determined by the necessarily unique outermost fixpoint variable (i.e. the largest with respect to $>_\varphi$) that occurs infinitely often in this play. If its fixpoint type is ν , then VERIFIER wins, otherwise it is μ and REFUTER wins.

The main advantage of these model checking games is the characterisation of the satisfaction relation via winning strategies in parity games (which they essentially are).

Proposition 2 ([15]). *Player VERIFIER has a winning strategy in the game in \mathcal{T} and a closed φ , starting in the configuration $s_1, \dots, s_k \vdash \varphi$ iff $\mathcal{T}, s_1, \dots, s_k \models \varphi$.*

3 The Arity Hierarchy

3.1 The Principle Construction

The aim of this section is to show that $\mathcal{L}_\mu^1, \mathcal{L}_\mu^2, \dots$ forms a strict hierarchy with respect to expressive power. The principles underlying the proof are easily explained: first we associate with each \mathcal{L}_μ^k -formula φ an LTS \mathcal{T}_φ with a designated starting state which we also call φ . Then we construct a closed \mathcal{L}_μ^{k+1} -formula that, when given a \mathcal{T}_φ , reads off what φ is from \mathcal{T}_φ and simulates its evaluation on it, checking that it does *not* hold on itself.

We first present the constructions principally, then discuss what results are achieved with the details of these constructions, and finally optimise the constructions such that the desired hierarchy result is achieved. We use a singleton Act which means that we simply write $s \rightarrow t$ instead of $s \xrightarrow{a} t$ for the single action name 'a'. Likewise, we write \diamond_i and \square_i instead of $\langle a \rangle_i$ and $[a]_i$.

Construction of \mathcal{T}_φ . Let $k \geq 1$ be fixed and take an arbitrary closed $\varphi \in \mathcal{L}_\mu^k$. We assume that the set of propositions underlying φ is $\text{Prop} = \{q_0, q_1, q_2, \dots\}$. The construction of \mathcal{T}_φ is largely based on the syntax-tree, respectively syntax-DAG of φ . We have $\mathcal{T}_\varphi = (\text{Sub}(\varphi), \rightarrow, \lambda, \varphi)$ with transitions given as follows.

$$\begin{array}{ll}
 \psi_1 \odot \psi_2 \rightarrow \psi_i & \text{for every } \psi_1 \odot \psi_2 \in \text{Sub}(\varphi), \odot \in \{\wedge, \vee\} \text{ and every } i \in \{1, 2\} \\
 \odot \psi \rightarrow \psi & \text{for every } \odot \psi \in \text{Sub}(\varphi), \odot \in \{\diamond_i, \square_i, \kappa\} \text{ and every } i \in \{1, \dots, k\} \\
 \eta X. \psi \rightarrow \psi & \text{for every } \eta X. \psi \in \text{Sub}(\varphi) \text{ and } \eta \in \{\mu, \nu\} \\
 X \rightarrow \psi & \text{for every } X \in \text{Sub}(\varphi) \text{ such that } fp_\varphi(X) = \eta X. \psi
 \end{array}$$

Thus, the graph structure of \mathcal{T}_φ is indeed almost the one of the syntax-DAG of φ except for additional edges from fixpoint variables to their defining fixpoint formula.

The labelling of the nodes in \mathcal{T}_φ remains to be defined. Remember that the ultimate goal is to construct a formula Φ^{k+1} which simulates the evaluation of φ on \mathcal{T}_φ . We will use k pebbles in order to simulate the k pebbles used in φ , and one additional pebble in order to store the subformula that is currently in question. Note that the satisfaction of a (closed) formula on an LTS only depends on the satisfaction of its subformulas. The position of this additional pebble will determine which subformula is currently evaluated. We therefore need to make the kind of subformula at a node in \mathcal{T}_φ visible to a formula that is interpreted over it. This is what the state labels will be used for. Let

$$\text{Prop}_0 := \{p_{j,i}^+, p_{j,i}^- \mid 1 \leq i \leq k, j \in \mathbb{N}\} \cup \{p^\wedge, p^\vee\} \cup \{p_i^\diamond, p_i^\square \mid 1 \leq i \leq k\} \cup \{p_i^{\text{FP}} \mid 0 \leq i \leq m\} \cup \{p_\kappa^{\text{rp}} \mid \kappa \in \mathbb{N} \dashrightarrow \mathbb{N}\}.$$

The labelling in \mathcal{T}_φ is given as follows. Note that Prop is countably infinite.

$$\begin{array}{ll}
 p_{j,i}^+ \in \lambda(q_j(i)) & \text{for every positive literal } q_j(i) \in \text{Sub}(\varphi) \\
 p_{j,i}^- \in \lambda(q_j(i)) & \text{for every negative literal } \neg q_j(i) \in \text{Sub}(\varphi) \\
 p^\wedge \in \lambda(\psi_1 \wedge \psi_2) & \text{for every } \psi_1 \wedge \psi_2 \in \text{Sub}(\varphi) \\
 p^\vee \in \lambda(\psi_1 \vee \psi_2) & \text{for every } \psi_1 \vee \psi_2 \in \text{Sub}(\varphi) \\
 p_i^\diamond \in \lambda(\diamond_i \psi) & \text{for every } \diamond_i \psi \in \text{Sub}(\varphi), 1 \leq i \leq k \\
 p_i^\square \in \lambda(\square_i \psi) & \text{for every } \square_i \psi \in \text{Sub}(\varphi), 1 \leq i \leq k \\
 p_\kappa^{\text{rp}} \in \lambda(\kappa \psi) & \text{for every } \kappa \psi \in \text{Sub}(\varphi), \kappa : \mathbb{N} \dashrightarrow \mathbb{N}
 \end{array}$$

$$p_i^{\text{FP}} \in \lambda(\eta X.\psi), \lambda(X) \quad \text{for every } \eta X.\psi, X \in \text{Sub}(\varphi), \eta \in \{\mu, \nu\} \text{ with } \text{ad}_\varphi(X) = i$$

With those labels a formula can see what the subformula at a node is that it is interpreted over, for instance whether it is a formula with a replacement modality as the principle operator, etc.

The construction of the simulating formulas. Next we construct formulas that simulate a φ on its own LTs representation \mathcal{T}_φ and check that they do not satisfy themselves. The trick is simple: if $\varphi \in \mathcal{L}_\mu^k$ then we use k pebbles to simulate what φ would do with its k pebbles, and one additional pebble to check which subformula we are currently evaluating. We let this pebble move through the syntax-DAG in a form that is dual to the semantics of the actual operators in the underlying φ ; for instance in a conjunction we look for one conjunct, in a disjunction we continue with both disjuncts. We will use several fixpoint variables to dualise the fixpoint condition similar to the way it is done in the Walukiewicz formulas that express the winning conditions in parity games [24].

Let $m \geq 0$ and $k \geq 1$ be fixed. We construct a formula $\Phi_m^{k+1} \in \mathcal{L}_\mu^{k+1}$ as follows.

$$\begin{aligned} \Phi_m^{k+1} := & \nu X_m.\mu X_{m-1} \dots \eta X_1. \left(\bigwedge_{i=1}^k \bigwedge_{j \in \mathbb{N}} p_{j,i}^+(k+1) \rightarrow \neg q_j(i) \right. \\ & \wedge \bigwedge_{i=1}^k \bigwedge_{j \in \mathbb{N}} p_{j,i}^-(k+1) \rightarrow q_j(i) \\ & \wedge p^\wedge(k+1) \rightarrow \diamond_{k+1} X_1 \\ & \wedge p^\vee(k+1) \rightarrow \square_{k+1} X_1 \\ & \wedge \bigwedge_{i=1}^k p_i^\diamond(k+1) \rightarrow \square_i \square_{k+1} X_1 \\ & \wedge \bigwedge_{i=1}^k p_i^\square(k+1) \rightarrow \diamond_i \square_{k+1} X_1 \\ & \wedge \bigwedge_{\kappa \in \mathbb{N} \rightarrow \mathbb{N}} p_\kappa^{\text{FP}}(k+1) \rightarrow \kappa \square_{k+1} X_1 \\ & \left. \wedge \bigwedge_{i=1}^m p_i^{\text{FP}}(k+1) \rightarrow \square_{k+1} X_i \right) \end{aligned}$$

where $\eta = \nu$ if m is odd and $\eta = \mu$ otherwise.

Remark 1. Of course, Φ_m^{k+1} is not a formula strictly speaking because of the potentially infinite conjunctions in the first two clauses. There is an easy way to fix this: we assume a finite set of atomic propositions $\{p, q, \dots\}$. Then a finite conjunction obviously suffices and Φ_m^{k+1} is indeed a formula. However, we need to address the issue of choice of atomic propositions in Section 3.2 below anyway. So for the moment we simply accept the small flaw about infinite conjunctions as an intermediate step and as a means to separate the principles from the details in this construction.

Note that this problem does not arise in the clause with the p_κ^{FP} since k is fixed, and κ can at most change the first k pebbles. Hence, there are only finitely many such κ .

We need two observations about Φ_m^{k+1} . The first, a syntactic one, is easy to verify.

Lemma 3. *For every $m \geq 0$ and every $k \geq 1$ we have $\Phi_m^{k+1} \in \Pi_m^{k+1}$.*

The second one is of a semantic nature and states that Φ_m^{k+1} does what it is supposed to do.

Lemma 4. *Let $m \geq 0$, $k \geq 1$ and $\varphi \in \Sigma_m^k$. Then we have $\mathcal{T}_{\varphi, \varphi} \models \Phi_m^{k+1}$ iff $\mathcal{T}_{\varphi, \varphi} \not\models \varphi$.*

Proof. We argue using model checking games for \mathcal{L}_μ^ω .

“ \Leftarrow ” Suppose we have $\mathcal{T}_{\varphi, \varphi} \not\models \varphi$, i.e. REFUTER has a winning strategy for the game \mathcal{G} played on \mathcal{T}_φ , k pebbles initially placed on the node φ in it, and the \mathcal{L}_μ^k -formula φ itself. This gives rise to a strategy for player VERIFIER in the game \mathcal{G}' played on \mathcal{T}_φ , now $k+1$ pebbles placed on node φ initially, and the formula Φ_m^{k+1} . The fact that each node in \mathcal{T}_φ satisfies exactly one atomic proposition of the kind p^* and at most one i or at most on κ means that any play which REFUTER does not lose immediately selects a clause in Φ_m^{k+1} , carries out some operation on the pebbles and then loops through some fixpoint variable.

It is not hard to see that VERIFIER can use REFUTER’s strategy from \mathcal{G} to follow the operations carried out on the pebbles prescribed by each clause without losing. For instance, if the third clause demands her to choose a successor for the $k+1$ -st pebble then she takes the one that represents the conjunct that REFUTER would choose in the same situation in \mathcal{G} . This way, every play in \mathcal{G}' that conforms to her strategy has an underlying play in \mathcal{G} that conforms to REFUTER’s strategy there. If that one is won by REFUTER because VERIFIER got stuck at some point then this can only be because the play reached a position of the form $(s_1, \dots, s_k) \vdash \diamond_i \psi$ and s_i has no successor. In the corresponding play in \mathcal{G}' , pebble $k+1$ will be on a node with label p_i^\diamond , and this requires REFUTER to move the i -th pebble to a successor which equally he cannot. Notice that the clause with p_i^\diamond contains the operator \Box and vice-versa. Thus, VERIFIER wins the corresponding play in \mathcal{G}' .

Suppose that the underlying play in \mathcal{G} is won by REFUTER because the largest fixpoint variable X that is seen infinitely often is of type μ . Then we must have $ad_\varphi(X) = i$ for some i , and then the play in \mathcal{G}' will infinitely often go through positions that are labelled with p_i^{FP} , and it will eventually not go through positions that are labelled with $p_{i'}^{\text{FP}}$ with $i' > i$ anymore. All that remains to be seen in this case is that the largest variable seen infinitely often in the play on Φ_m^{k+1} is of type ν . This is a direct consequence of Lemma 1. Hence, VERIFIER wins such plays, too, which shows that her strategy derived from REFUTER’s winning strategy in \mathcal{G} is winning for her in \mathcal{G}' .

“ \Rightarrow ” This is shown by contraposition in the same way now assuming a winning strategy for VERIFIER in the game on \mathcal{T}_φ and φ and turning it into a winning strategy for REFUTER in the game on \mathcal{T}_φ and Φ_m^{k+1} .

Lemma 5. *Let $m \geq 0$ and $k \geq 1$. There is no $\varphi \in \Sigma_m^k$ such that $\varphi \equiv \Phi_m^{k+1}$.*

Proof. Suppose there was such a φ . Then we would have

$$\mathcal{T}_{\varphi, \varphi} \models \varphi \quad \text{iff} \quad \mathcal{T}_{\varphi, \varphi} \models \Phi_m^{k+1} \quad \text{iff} \quad \mathcal{T}_{\varphi, \varphi} \not\models \varphi$$

first because of the assumed equivalence and second because of Lemma 4.

Thus, we could summarise the findings from these lemmas and also uses the observation that the entire construction is equally possible for formula in Π_m^k then yielding a $\Phi_m^{k+1} \in \Sigma_m^{k+1}$. Then we get that for all $m \geq 0$ and $k \geq 1$ we have $\Sigma_m^k \not\subseteq \Pi_m^{k+1}$ and $\Pi_m^k \not\subseteq \Sigma_m^{k+1}$. Consequently, we have $\Sigma_m^k \subsetneq \Sigma_{m+1}^{k+1}$ and $\Pi_m^k \subsetneq \Pi_{m+1}^{k+1}$.

The reason why we do not formally state this as a theorem (yet) is discussed next.

3.2 The Hierarchy over a Fixed Small Signature

Consider what is happening with the set of atomic propositions in the construction of the previous Section 3.1. We have already seen in Remark 1 that the construction does not work for an infinite set of

atomic propositions Prop . Even if this is finite, then the construction does work but it has the following effect: we simulate a formula with k pebbles over Prop by a formula with $k+1$ pebbles over $\text{Prop} \cup \text{Prop}_0$. It is not surprising that we obtain formulas over this extended signature which cannot be expressed over the smaller one. In order to argue that the hierarchy of inexpressibility as laid out in the previous section is truly meaningful we would need $\text{Prop} = \text{Prop} \cup \text{Prop}_0$ or, at least, that the two sets have equal cardinality so that some bijection between them could be used as an encoding.

In the following we will show how the construction can be fixed such that it works over a fixed finite set

$$\text{Prop}_1 := \{p^+, p^-, p^\wedge, p^\vee, p^\diamond, p^\square, p^{\text{FP}}, p^{\text{rp}}, p^{\text{sw}}, p^\bullet\}$$


of atomic propositions. Thus, we do not encode the index of propositions, the level in the fixpoint hierarchy, and the kind of operation on pebbles in those propositions anymore. Instead we will encode this missing information in the graph structure of \mathcal{T}_φ (rather than in its labels). For the replacement modalities κ we need a little preparation.

A replacement κ is called simple if it is of the form $\{i \leftarrow j\}$ or $\{i \leftrightarrow j\}$. A formula is called *normalised* if every replacement in it is simple. The following is a simple consequence of the fact that every function $\kappa : \mathbb{N} \dashrightarrow \mathbb{N}$ that leaves all numbers greater than k untouched, can be constructed by a sequence of swaps between $i, j \leq k$, followed by some simple mappings from some i to a j .

Lemma 6. *Let $m \geq 0, k \geq 1$. Every $\varphi \in \Sigma_m^k$, respectively Π_m^k is equivalent to a normalised $\varphi' \in \Sigma_m^k$, respectively Π_m^k .*

We therefore assume that from now on, all formulas φ to be simulated are normalised. We change the construction of \mathcal{T}_φ as follows.

1. Suppose there is a state s of the form $q_j(i)$ or $\neg q_j(i)$, necessarily labeled with $p_{j,i}^+$ or $p_{j,i}^-$. Replace the proposition by p^+ , respectively p^- , and add a new finite path of length $i+j$ to this node such that the $(i-1)$ -st new state has the label p^\bullet .



Let $\kappa_{\text{cyc}}^\leftarrow := \{k \leftarrow 1, 1 \leftarrow 2, \dots, k-1 \leftarrow k\}$ and consider the formula

$$\text{searchPeb} := (p^\bullet \wedge q_j(1)) \vee \kappa_{\text{cyc}}^\leftarrow \diamond_{k+1} ((p^\bullet \wedge q_j(1)) \vee \kappa_{\text{cyc}}^\leftarrow \diamond_{k+1} ((p^\bullet \wedge q_j(1)) \vee \dots \vee \kappa_{\text{cyc}}^\leftarrow \diamond_{k+1} (p^\bullet \wedge q_j(1)) \dots)) \quad (1)$$

with $k-1$ occurrences of $\kappa_{\text{cyc}}^\leftarrow$. It is true in s at pebble $k+1$ with the additional path iff $q_j(i)$ was true in s with the original construction. The first part of this new path is used to shift the pebbles until the i -th has become the first and then, instead of checking whether the i -th pebble is on a state satisfying q_j , we can now check the first one instead. Note that this formula moves pebble number $k+1$ along this new path but the other k pebbles remain where they are; apart from being cyclically changed around.

We can of course equally construct such a formula that mimicks the checking of $\neg q_j(i)$.

Finally, we also need to use the remaining path to read off the encoding of j . This can easily be done as follows.

$$\text{searchProp}_j := q_0(1) \vee \diamond_{k+1}(q_1(1) \vee \diamond_{k+1}(q_2(1) \vee \dots \vee \diamond_{k+1}(q_{h-2}(1) \vee \diamond_{k+1}q_{h-1}(1)) \dots))$$

This formula is then used instead of $q_j(1)$ in (1), and the resulting formula is used instead of $q_j(i)$ in the clause for $p_{j,i}^-$ in Φ_m^{k+1} . Hence, this clause simply becomes

$$\dots \wedge p_{j,i}^- \rightarrow \text{searchPeb}[\text{searchProp}/q_j(1)]$$

where $\psi[\chi/\chi']$ denotes the formula that results from ψ by replacing every subformula χ' with χ .

2. An edge of the form $\eta X.\psi \rightarrow \psi$ or $X \rightarrow \psi$ is replaced in similar style by a sequence of i edges, marking the last state after them with p^\bullet . Then we can replace the label p_i^{FP} with p^{FP} in the first state, and the corresponding clause in Φ_m^{k+1} with

$$\dots \wedge p^{\text{FP}} \rightarrow \square_{k+1}((p^\bullet \wedge X_1) \vee \square_{k+1}((p^\bullet \wedge X_2) \vee \square_{k+1}(\dots \vee \square_{k+1}(p^\bullet \wedge X_m) \dots)))$$

3. An edge of the form $\diamond_i \psi \rightarrow \psi$ or $\square_i \psi \rightarrow \psi$ is replaced by a sequence of $2i$ edges via new states, and p^\bullet must hold after i and after $2i$ steps. The trick to use here is to cycle the first k pebbles until the i -th one becomes the first, then execute the corresponding action for the i -pebble on the first one instead, and then cycle them back again. Let $\kappa_{\text{cyc}}^{\leftarrow}$ be as above and $\kappa_{\text{cyc}}^{\rightarrow} := \{2 \leftarrow 1, 3 \leftarrow 2, \dots, 1 \leftarrow k\}$. Then we can replace the clause for p_i^\diamond in Φ_m^{k+1} by

$$\dots \wedge p^\diamond \rightarrow (p^\bullet \wedge \square_{k+1}(\square_1 \text{goBack} \vee \kappa_{\text{cyc}}^{\leftarrow}(p^\bullet \wedge \square_{k+1}(\square_1 \text{goBack} \vee \dots (p^\bullet \wedge \square_{k+1} \square_1 \text{goBack}) \dots)))$$

with exactly $k-1$ occurrences of $\kappa_{\text{cyc}}^{\leftarrow}$ and

$$\text{goBack} := \square_{k+1}((p^\bullet \wedge X_1) \vee \kappa_{\text{cyc}}^{\rightarrow} \square_{k+1}((p^\bullet \wedge X_1) \vee \dots \square_{k+1}(p^\bullet \wedge X_1) \dots))$$

with exactly $k-1$ occurrences of $\kappa_{\text{cyc}}^{\rightarrow}$.

Likewise, we can use the same trick to eliminate the dependence on i of the formula Φ_m^{k+1} in the clause for p_i^\square which is equally replaced by p^\square , and those paths of length $2i$ can be used to decode the value i from the graph structure instead of reading it straight off the atomic proposition.

4. Finally, we can use the same trick in a slightly more elaborate fashion to handle replacement modalities of the form $\{i \leftarrow j\}$ and $\{i \leftrightarrow j\}$. We mark nodes in \mathcal{T}_φ that correspond to the form by p^{P} and those that correspond to the latter by p^{sw} . A swap of the form $\{i \leftrightarrow j\}$ can be handled as follows: assume $i < j$.

- (a) Cyclically shift the pebbles $1, \dots, k$ for i positions to the left.
- (b) Cyclically shift the pebbles $2, \dots, k$ for $j-i-1$ positions to the left.
- (c) Swap pebbles 1 and 2.
- (d) Cyclically shift the pebbles $2, \dots, k$ for $j-i-1$ positions to the right.
- (e) Cyclically shift the pebbles $1, \dots, k$ for i positions to the right.

Hence, we replace a transition of the form $\{i \leftarrow j\} \psi \rightarrow \psi$ by a path of length $2j-2$ and mark the states at positions $i, j-1, 2j-2$ and the last one with p^\bullet so that we can, like above, construct a formula that mimicks the five steps above to carry out the swapping of pebbles i and j .

The construction for replacements of the form $\{i \leftrightarrow j\}$ is similar. Again, the trick is to cycle i and j to positions 1 and 2, carry the replacement out on these fixed positions, and cycle the pebbles back again. These eliminates the dependence of Φ_m^{k+1} on propositions which carry such a value.

With this being done, Φ_m^{k+1} becomes a formula that is defined over a fixed set Prop_1 of atomic propositions of size 10, and we can use it to simulate formulas $\varphi \in \Sigma_m^k$ over the same Prop_1 . Then the inexpressibility result of the previous section becomes meaningful. Using standard encoding techniques we can break the result down to \mathcal{L}_μ^ω over two atomic propositions only, using binary encoding, or a single one, using unary encoding. The atomic propositions can also be eliminated entirely by appending certain finite trees to the states in which they hold such that these trees are checkable using fixpoint-free formulas of \mathcal{L}_μ^ω . Hence, we get the following.

Theorem 7. *For all $m \geq 0$ and $k \geq 1$ we have $\Sigma_m^k \not\subseteq \Pi_m^{k+1}$ and $\Pi_m^k \not\subseteq \Sigma_m^{k+1}$. Consequently, we have $\Sigma_m^k \subsetneq \Sigma_{m+1}^{k+1}$ and $\Pi_m^k \subsetneq \Pi_{m+1}^{k+1}$. These results hold independently of the underlying signature Prop and Act .*

4 Conclusion and Further Work

We have shown that the arity hierarchy in the polyadic μ -calculus, a modal fixpoint logic for specifying bisimulation-invariant relational properties of states in transition systems, is strict in the sense that higher arity gives higher expressive power provided that one is allowed to use a little bit more fixpoint alternation ($\Sigma_m^k \subsetneq \Sigma_{m+1}^{k+1}$). If alternation must not increase then higher arity yields not necessarily more but different expressiveness ($\Sigma_m^k \not\subseteq \Pi_m^{k+1}$).

Obviously, the exact effects on expressive power that should be attributed to arity and to fixpoint alternation need to be separated. A first step would be to prove the strictness of the alternation hierarchy within each \mathcal{L}_μ^k . For $k = 1$, i.e. the ordinary μ -calculus, this is known for arbitrary and in particular for finite transition systems [6, 17]. Subsequently, the result could be shown for several other classes of transition systems, for instance binary trees [3, 7], nested words [11] and graphs whose edge relation satisfies certain properties like being transitive for instance [8, 1].

We suspect that not only is the alternation hierarchy within each \mathcal{L}_μ^k also strict, but equally that Arnold's proof [3] using a similar diagonalisation argument for \mathcal{L}_μ can be extended. It relies on the interreducibility between model checking for \mathcal{L}_μ and parity games [20] and in particular the existence of the Walukiewicz formulas defining winning regions in parity games [24]. It is known [15] that the model checking problem for \mathcal{L}_μ^k and any $k \geq 1$ can equally be reduced to a parity game, and it seems feasible to extend the construction of the Walukiewicz formulas to higher arity. This would use similar principles as those underlying the construction of Φ_m^{k+1} in Section 3.

Model checking \mathcal{L}_μ^k can also be reduced to model checking \mathcal{L}_μ directly using k -products of transition systems, i.e. there is a translation of \mathcal{L}_μ^k -formulas to \mathcal{L}_μ -formulas that preserves truth under taking k -fold products of transition systems [19, 15]. Hence, the question of the strictness of the alternation hierarchy in \mathcal{L}_μ^k is equivalent to the question after the strictness of the \mathcal{L}_μ alternation hierarchy over the class of all k -fold products of transition systems.

References

- [1] L. Alberucci & A. Facchini (2009): *The modal μ -calculus over restricted classes of transition systems*. *Journal of Symbolic Logic* 74(4), pp. 1367–1400, doi:10.2178/jsl/1254748696.
- [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] A. Arnold (1999): *The modal μ -calculus alternation hierarchy is strict on binary trees*. *RAIRO - Theoretical Informatics and Applications* 33, pp. 329–339, doi:10.1051/ita:1999121.

- [4] J. Bradfield & C. Stirling (2001): *Modal logics and μ -calculi: an introduction*. In J. Bergstra, A. Ponse & S. Smolka, editors: *Handbook of Process Algebra*, Elsevier, pp. 293–330, doi:10.1016/B978-044482830-9/50022-9.
- [5] J. Bradfield & C. Stirling (2007): *Modal μ -calculi*. In P. Blackburn, J. van Benthem & F. Wolter, editors: *Handbook of Modal Logic: Studies in Logic and Practical Reasoning Volume 3*, Elsevier, pp. 721–756, doi:10.1016/S1570-2464(07)80015-2.
- [6] J. C. Bradfield (1996): *The Modal μ -calculus Alternation Hierarchy Is Strict*. In: *Proc. 7th Conf. on Concurrency Theory, CONCUR'96, LNCS 1119*, Springer, pp. 233–246, doi:10.1007/3-540-61604-7_58.
- [7] J. C. Bradfield (1999): *Fixpoint Alternation: Arithmetic, Transition Systems, and the Binary Tree*. *RAIRO - Theoretical Informatics and Applications* 33(4/5), pp. 341–356, doi:10.1051/ita:1999122.
- [8] G. D'Agostino & Giacomo Lenzi (2010): *On the μ -calculus over transitive and finite transitive frames*. *Theoretical Computer Science* 411(50), pp. 4273–4290, doi:10.1016/j.tcs.2010.09.002.
- [9] E. A. Emerson & C. L. Lei (1986): *Efficient Model Checking in Fragments of the Propositional μ -Calculus*. In: *Symposium on Logic in Computer Science*, IEEE, Washington, D.C., USA, pp. 267–278.
- [10] M. Grohe (1996): *Arity hierarchies*. *Annals of Pure and Applied Logic* 82(2), pp. 103–163, doi:10.1016/0168-0072(95)00072-0.
- [11] J. Gutierrez, F. Klaedtke & M. Lange (2014): *The μ -Calculus Alternation Hierarchy Collapses over Structures with Restricted Connectivity*. *Theoretical Computer Science* 560(3), pp. 292–306, doi:10.1016/j.tcs.2014.03.027.
- [12] 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.
- [13] B. Knaster (1928): *Un théorème sur les fonctions d'ensembles*. *Annals Soc. Pol. Math* 6, pp. 133–134.
- [14] D. Kozen (1982): *Results on the Propositional μ -Calculus*. In: *Proc. 9th Int. Coll. on Automata, Languages and Programming, ICALP'82, LNCS 140*, Springer, pp. 348–359, doi:10.1007/BFb0012782.
- [15] M. Lange & E. 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.
- [16] M. Lange, E. Lozes & M. Vargas Guzmán (2014): *Model-Checking Process Equivalences*. *Theoretical Computer Science* 560, pp. 326–347, doi:10.1016/j.tcs.2014.08.020.
- [17] G. Lenzi (1996): *A Hierarchy Theorem for the μ -Calculus*. In: *Proc. 23rd Int. Coll. on Automata, Languages and Programming, ICALP'96, LNCS 1099*, Springer, pp. 87–97, doi:10.1007/3-540-61440-0_119.
- [18] D. Niwiński (1988): *Fixed Points vs. Infinite Generation*. In: *Proc. 3rd Ann. Symp. on Logic in Computer Science, LICS'88*, IEEE Computer Society, pp. 402–409.
- [19] 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.
- [20] C. Stirling (1995): *Local Model Checking Games*. In: *Proc. 6th Conf. on Concurrency Theory, CONCUR'95, LNCS 962*, Springer, pp. 1–11, doi:10.1007/3-540-60218-6_1.
- [21] C. Stirling (1996): *Games and Modal μ -Calculus*. In T. Margaria & B. Steffen, editors: *Proc. 2nd Int. Workshop on Tools and Algorithms for the Construction and Analysis of Systems, TACAS'96, LNCS 1055*, Springer, pp. 298–312, doi:10.1007/3-540-61042-1_51.
- [22] 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.
- [23] 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.
- [24] I. Walukiewicz (2002): *Monadic second-order logic on tree-like structures*. *Theor. Comput. Sci* 275(1-2), pp. 311–346, doi:10.1016/S0304-3975(01)00185-2.