The variable hierarchy of the μ-calculus is strict*

Dietmar Berwanger[1] and Giacomo Lenzi[2]

[1] Mathematische Grundlagen der Informatik, RWTH Aachen, D-52056 Aachen,
`berwanger@cs.rwth-aachen.de`

[2] Dipartimento di Matematica, Università di Pisa, via Buonarroti 2, I-56127 Pisa,
`lenzi@mail.dm.unipi.it`

Abstract. The modal μ-calculus L_μ attains high expressive power by extending basic modal logic with monadic variables and binding them to extremal fixed points of definable operators. The number of variables occurring in a formula provides a relevant measure of its conceptual complexity. In a previous paper we have shown, for the existential fragment of L_μ, that this conceptual complexity is also reflected in an increase of semantic complexity, by providing examples of existential formulae with k variables that are not equivalent to any existential formula with less than k variables.

In this paper, we prove an existential preservation theorem for the family of L_μ-formulae over at most k variables that define simulation closures of finite strongly connected structures. Since hard formulae for the level k of the existential hierarchy belong to this family, it follows that the bounded variable fragments of the full modal μ-calculus form a hierarchy of strictly increasing expressive power.

Keywords: μ-calculus, structural complexity.

1 Introduction

Among the various formalisms for reasoning about dynamic systems, the modal μ-calculus L_μ enjoys a prominent position due to its high expressive power and model-theoretic robustness, in balance with its fairly manageable computational complexity. As such, L_μ offers a frame of reference for virtually every logic for specifying the operational behaviour of reactive and concurrent systems.

Typically, such systems are modelled as transition structures with elementary states labelled by propositions and binary transition relations labelled by actions. A simple language for speaking about these models is basic modal logic, or Hennessy-Milner logic [10], which extends propositional logic by modalities associated to actions, i.e., existential and universal quantifiers over the successors of a state which are reachable via the specified action.

The μ-calculus of Kozen [13] extends basic modal logic by adding monadic variables bound by least and greatest fixed points of definable operators. This

* This research has been partially supported by the European Community Research Training Network "Games and Automata for Synthesis and Validation" (GAMES)

provides a notion of recursion which invests the logic with high expressive power. On the other hand, the variables also import considerable conceptual complexity.

A well studied source of conceptual complexity is the *alternation depth* of L_μ-formulae, that is, the number of (genuine) alternations between least and greatest fixed points. In [7] Bradfield showed that the alternation hierarchy of the μ-calculus is semantically strict; variants of this result have also been proven by Lenzi [15] and Arnold [1]. Hence, this notion of syntactic complexity of a formula is reflected in its semantic complexity.

Interestingly, most of the formalisms commonly used for process description allow translations into low levels of the L_μ alternation hierarchy. On its first level this hierarchy already captures, for instance, PDL as well as CTL, while their expressive extensions ΔPDL and CTL* do not exceed the second level. Still, the low levels of this hierarchy do not exhaust the significant properties expressible in L_μ. A comprehensive example of formulae distributed over all levels of the alternation hierarchy is provided by parity games. Thus, strictly on level n, there is a formula stating that the first player has a winning strategy in parity games with n priorities.

By reusing fixed point variables several times, it is possible to write many L_μ-formulae, even with highly nested fixed-point definitions, using only very few variables. For any k, let us denote by $L_\mu[k]$ the fragment of L_μ consisting of those formulae that make use of at most k distinct fixed-point variables. It turns out that most specification properties of transition structures can be embedded into $L_\mu[2]$. This is actually the case for all the aforementioned formalisms, CTL, PDL, CTL*, and ΔPDL (see [17]). But the expressive power of the two-variable fragment of L_μ goes well beyond this. As shown in [3], the formulae describing the winning position of a parity game, can also be written with only two variables.

In this context, the question arises, whether a higher number of variables is indeed necessary, or, in other words, whether the number of variables of a formula is reflected as a measure of its semantic complexity.

As a first step towards answering this question, we have proved, together with Grädel in [5], that the variable hierarchy of the existential fragment of L_μ is strict. This syntactic fragment, consisting of the formulae built from atoms and negated atoms by means of boolean connectives, existential modalities, and least and greatest fixed points, is sharply contoured under semantical viewpoint. In [8], D'Agostino and Hollenberg proved that it captures precisely those L_μ-expressible properties ψ that are preserved under extensions, in the sense that whenever $\mathcal{K}, v \models \psi$ and $\mathcal{K} \subseteq \mathcal{K}'$, then also $\mathcal{K}', v \models \psi$. Unfortunately, the technique used in this proof does not comply with a variable-confined setting, and the question whether the variable hierarchy is strict for the full μ-calculus remained open.

To witness the strictness of the variable hierarchy in the restricted existential case considered in [5], we have provided examples of formulae $\psi^k \in L_\mu[k]$, for each level k, that cannot be equivalently expressed by any formula over less than k variables using only existential modalities. Essentially, the formulae ψ^k describe the class of structures extending a clique with k states, where every pair of states i, j is linked by a transition labelled ij.

In the present paper, we prove a preservation theorem stating that every formula defining the extensions of a finite strongly connected structure can be transformed into an existential formula without increasing the number of variables. In particular, this holds for the witnesses ψ^k to the strictness of the existential hierarchy provided in [5]. Consequently, even if the use of universal modalities is allowed, none of the formulae ψ^k can be equivalently written as a formula with less than k variables. In this way, we can answer positively the question concerning the strictness of the variable hierarchy of the full μ-calculus by reducing it to the existential case.

Besides revealing a new aspect of the rich inner structure of the μ-calculus, this result settles an open question formulated in [17] regarding the expressive power of Parikh's Game Logic GL. This logic, introduced in [16] as a generalisation of PDL for reasoning about games, subsumes ΔPDL and intersects nontrivially all the levels of the L_μ alternation hierarchy [3]. When interpreted on transition structures, GL can be translated into $L_\mu[2]$. However it was unknown, up to now, whether the inclusion in L_μ was proper. The strictness of the variable hierarchy implies that already $L_\mu[3]$ is more expressive than GL.

The paper is structured as follows. In Section 2, we introduce the necessary background on the μ-calculus. Section 3 is dedicated to the proof of the Preservation Theorem. We conclude by stating the Hierarchy Theorem in Section 4.

2 The Modal μ-Calculus

Fix a set ACT of actions and a set PROP of atomic propositions. A transition structure for ACT and PROP is a structure $\mathcal{K}$ with universe K (whose elements are called *states*), binary relations $E_a \subseteq K \times K$ for each $a \in$ ACT, and monadic relations $p \subseteq K$ for each atomic proposition $p \in$ PROP.

Syntax. For a set ACT of actions, a set PROP of atomic propositions, and a set VAR of monadic variables, the formulae of L_μ are defined by the grammar

$$\varphi ::= \mathit{false} \mid \mathit{true} \mid p \mid \neg p \mid X \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle a \rangle \varphi \mid [a]\varphi \mid \mu X.\varphi \mid \nu X.\varphi$$

where $p \in$ PROP, $a \in$ ACT, and $X \in$ VAR. An L_μ-formula in which no universal modality $[a]\varphi$ occurs is called *existential*.

The number of distinct variables appearing in an L_μ-formula induces the following syntactic hierarchy. For any $k \in \mathbb{N}$, the *k-variable fragment* $L_\mu[k]$ of the μ-calculus is the set of formulae $\psi \in L_\mu$ that contain at most k distinct variables.

Semantics. Formulae of L_μ are evaluated on transition structures at a particular state. Given a sentence ψ and a structure $\mathcal{K}$ with state v, we write $\mathcal{K}, v \models \psi$ to denote that ψ holds in $\mathcal{K}$ at state v. The set of states $v \in K$ such that $\mathcal{K}, v \models \psi$ is denoted by $[\![\psi]\!]^{\mathcal{K}}$.

Here, we only define $[\![\psi]\!]^{\mathcal{K}}$ for fixed-point formulae ψ. Towards this, note that a formula $\psi(X)$ with a monadic variable X defines on every transition structure $\mathcal{K}$ (providing interpretations for all free variables other than X occurring in

ψ) an operator $\psi^{\mathcal{K}} : \mathcal{P}(K) \to \mathcal{P}(K)$ assigning to every set $X \subseteq K$ the set $\psi^{\mathcal{K}}(X) := [\![\psi]\!]^{\mathcal{K},X} = \{v \in K : (\mathcal{K}, X), v \models \psi\}$. As X occurs only positively in ψ, the operator $\psi^{\mathcal{K}}$ is *monotone* for every $\mathcal{K}$, and therefore, by a well-known theorem due to Knaster and Tarski, has a least fixed point $\mathrm{lfp}(\psi^{\mathcal{K}})$ and a greatest fixed point $\mathrm{gfp}(\psi^{\mathcal{K}})$. Now we put

$$[\![\mu X.\psi]\!]^{\mathcal{K}} := \mathrm{lfp}(\psi^{\mathcal{K}}) \text{ and } [\![\nu X.\psi]\!]^{\mathcal{K}} := \mathrm{gfp}(\psi^{\mathcal{K}}).$$

Least and greatest fixed points can also be constructed inductively. Given a formula $\nu X.\psi$, we define for each ordinal α, the stage X^α of the gfp-induction of $\psi^{\mathcal{K}}$ by $X^0 := K$, $X^{\alpha+1} := [\![\psi]\!]^{(\mathcal{K},X^\alpha)}$, and $X^\alpha := \bigcap_{\beta<\alpha} X^\beta$ if α is a limit ordinal. By monotonicity, the stages of the gfp-induction decrease until a fixed point is reached. By ordinal induction, one easily proves that this inductively constructed fixed point coincides with the greatest fixed point. The *finite approximants* of a formula $\nu X.\varphi$ are defined by $\varphi_0 :=$ *true* and $\varphi_{n+1} = \varphi[X/\varphi_n]$ (the formula obtained by replacing every occurrence of X in φ, by φ_n). Obviously, $\nu X.\varphi$ implies φ_n for all n.

The validity of existential L_μ-formulae is preserved under model extensions and, more generally, under the following notion of simulation.

Definition 1. A *simulation* from a structure $\mathcal{K}$ to a structure $\mathcal{K}'$ is a relation $Z \subseteq K \times K'$ respecting the atomic propositions $p \in \textsc{prop}$ in the sense that $\mathcal{K}, v \models p$ iff $\mathcal{K}', v' \models p$, for $(v, v') \in Z$, which satisfies the following condition. For all $(v, v') \in Z$, $a \in \textsc{act}$, and every w such that $(v, w) \in E_a$, there exists a $w' \in K'$ such that $(v', w') \in E'_a$ and $(w, w') \in Z$. We say that $\mathcal{K}, u$ *simulates* $\mathcal{K}', u'$ and write $\mathcal{K}, u \lesssim \mathcal{K}', u'$, if between the structures there is a simulation containing (u, u').

As a modal logic, the μ-calculus distinguishes between transitions structures only up to behavioural equivalence, captured by the notion of bisimulation.

Definition 2. A *bisimulation* between two transition structures $\mathcal{K}$ and $\mathcal{K}'$ is a simulation Z from $\mathcal{K}$ to $\mathcal{K}'$ so that the inverse relation Z^{-1} is a simulation from $\mathcal{K}'$ to $\mathcal{K}$. Two transition structures $\mathcal{K}, u$ and $\mathcal{K}', u'$ are *bisimilar*, if there is a bisimulation Z between them, with $(u, u') \in Z$.

An important model-theoretic feature of modal logics is the *tree model property* meaning that every satisfiable formula is satisfiable in a tree. This is a straightforward consequence of bisimulation invariance, since $\mathcal{K}, u$ is bisimilar to its *tree unravelling*, i.e., a tree whose nodes correspond to the finite paths in $\mathcal{K}, u$. Every such path π inherits the atomic propositions of its last node v; for every node w reachable from v in $\mathcal{K}$ via an a transition, π is connected to its prolongation by w via an a-transition.

Another significant feature of L_μ is its *finite model property.*

Theorem 3 ([12]). *Every satisfiable* L_μ*-formula has a finite model.*

Since the unravelling of a finite model is a finitely branching tree, we obtain the following corollary.

Corollary 4. *Every satisfiable* L_μ*-formula holds in some finitely branching tree.*

For later use, we state a further consequence of the finite model property.

Corollary 5. *For* $\psi \in \mathrm{L}_\mu$*, let* $\psi[\nu := \nu^n]$ *denote the result of replacing every occurrence* $\nu X.\eta$ *of a* ν*-predicate in* ψ *with its* n*-th approximant* η_n*. Then, a formula* $\varphi \in \mathrm{L}_\mu$ *implies* ψ *if, and only if,* φ *implies* $\psi[\nu := \nu^n]$*, for each* n*.*

Model-checking games. The semantics of L_μ can also be described in terms of *parity games*, in which two players form a path in a given graph with nodes labelled by natural numbers called priorities. If a player cannot move, he loses. If this never occurs, the winner is decided by looking at the (infinite) sequence of priorities occurring in the play. The first player wins if the least priority appearing infinitely often in this sequence is even, otherwise his opponent wins. The Forgetful Determinacy Theorem states that these games are always determined, and the winner has a memoryless winning strategy, that is, a strategy that does not depend on the history of the play but only on the current position.

Theorem 6 (Forgetful Determinacy, [9])**.** *In any parity game, one of the players has a memoryless winning strategy.*

Given a transition structure $\mathcal{K}, v_0$ and a L_μ-sentence ψ, the model-checking game $\mathcal{G}(\mathcal{K}, \psi)$ is a parity game associated with the problem whether $\mathcal{K}, v_0 \models \psi$. Deviating from the more traditional way to define this game with positions associated to subformulae of ψ (see, e.g., [4, 18]), we use a variant more familiar in automata theory which, instead of subformulae, refers to their closure [9, 14].

Definition 7. Let $\psi \in \mathrm{L}_\mu$ be a formula without free variables. For each subformula φ in ψ, we define its *closure* $\mathrm{cl}_\psi(\varphi)$ as the formula obtained by replacing recursively every free occurrence of a variable in φ by its binding definition. By $\mathrm{cl}(\psi)$ we denote the set of closures of all subformulae in ψ.

The positions in the game $\mathcal{G}(\mathcal{K}, \psi)$ are pairs (v, φ) of states $v \in K$ and sentences $\varphi \in \mathrm{cl}(\psi)$. The first player, called Verifier, moves from the positions $(v, \varphi_1 \vee \varphi_2)$, $(v, \langle a \rangle \varphi)$, (v, p) with $v \notin p$, and $(v, \neg p)$ with $v \in p$ and his opponent, called Falsifier, moves from every other position. All plays start at (v_0, ψ) and can proceed as follows:

- no moves are possible from (v, α) where α is atomic or negated atomic;
- from $(v, \varphi_1 \vee \varphi_2)$ or $(v, \varphi_1 \wedge \varphi_2)$ available moves lead to (v, φ_1) and (v, φ_2);
- from $(v, \langle a \rangle \varphi)$ or $(v, [a]\varphi)$ there are available moves to all positions (w, φ) where w is an a-successor of v;
- from $(v, \lambda X.\varphi(X))$ the play moves to $(v, \varphi(\lambda X.\varphi(X))$.

The priority labelling assigns even priorities to positions $(v, \nu X.\varphi)$ and odd priorities to positions $(v, \mu X.\varphi)$, respecting the nesting of greatest and least fixed-point operators. For details (which are not needed in this paper), see [4].

Theorem 8 ([18])**.** *Verifier has a winning strategy in the model-checking game* $\mathcal{G}(\mathcal{K}, \psi)$ *from position* (u, ψ) *iff* $\mathcal{K}, u \models \psi$*.*

Simultaneous Fixed Points. There is a variant of L_μ that admits simultaneous fixed points of several formulae. This does not increase the expressive power but allows more transparent formalisations. The mechanism for building simultaneous fixed-point formulae is the following. Given formulae $\varphi_1, \dots, \varphi_n$ and variables $X_1, \dots, X_n$, we can write an *equational system* $S := \{X_1 = \varphi_1, \dots, X_n = \varphi_n\}$ and build formulae $(\mu X_i : S)$ and $(\nu X_i : S)$. On every structure $\mathcal{K}$, the system S defines an operator $S^{\mathcal{K}}$ mapping an n-tuple $\bar{X} = (X_1, \dots, X_n)$ of sets of states to $S_1^{\mathcal{K}}(\bar{X}), \dots, S_n^{\mathcal{K}}(\bar{X})$ so that, for each i we have: $S_i^{\mathcal{K}}(\bar{X}) := [\![\varphi_i]\!]^{(\mathcal{K},\bar{X})}$. As $S^{\mathcal{K}}$ is monotone, it has extremal fixed points $\mathrm{lfp}(S) = (X_1^\mu, \dots, X_n^\mu)$ respectively $\mathrm{gfp}(S) = (X_1^\nu \dots, X_n^\nu)$, and we set $[\![(\mu X_i : S)]\!]^{\mathcal{K}} := X_i^\mu$ and $[\![(\nu X_i : S)]\!]^{\mathcal{K}} := X_i^\nu$.

It is known that simultaneous least fixed points can be eliminated in favour of nested individual fixed points.

Proposition 9 ([2]). *Every formula in L_μ with simultaneous fixed points can be translated into an equivalent formula in plain L_μ without increasing the number of variables.*

3 The Preservation Theorem

The key argument in our proof of the Hierarchy Theorem consists in the preservation property stated in the current section, which implies that the formulae proposed in [5] to separate the hierarchic levels of the existential fragment also witness the strictness of full μ-calculus variable hierarchy.

This preservation property concerns formulae which define simulation closures of certain structures. The *simulation closure* of a rooted transition structure $\mathcal{K}, v_0$ is the class

$$(\mathcal{K}, v_0)^{\lesssim} := \{\, \mathcal{K}', v_0' \mid \mathcal{K}, v_0 \lesssim \mathcal{K}', v_0' \,\}.$$

Clearly, if $\mathcal{K}$ is finite, this class can be axiomatised by an L_μ-formula. For convenience, we will use simultaneous fixed points. Let S be the system defining, for every node $v \in K$, a proposition X_v via the equation

$$X_v = \bigwedge_{p \mid v \in p} p \wedge \bigwedge_{a \in \mathrm{ACT}, (v,w) \in E_a} \langle a \rangle X_w .$$

It can be easily seen that on any transition structure $\mathcal{K}'$, the greatest solution of this system maps every variable X_v to the set $\{\, v' \in K' \mid \mathcal{K}, v \lesssim \mathcal{K}', v' \,\}$. Hence, for any state $v' \in [\![\nu X_v : S]\!]^{\mathcal{K}'}$, we have $\mathcal{K}, v \lesssim \mathcal{K}', v'$.

For further use, let us denote the L_μ-formula obtained as a translation of the equational expression $\nu X_v : S$ by $\psi_v^{\mathcal{K}}$. For the formula $\psi_{v_0}^{\mathcal{K}}$ associated to the designated root v_0 of $\mathcal{K}$, we write $\psi^{\mathcal{K}}$, and call it the *canonical axiom* of $(\mathcal{K}, v_0)^{\lesssim}$.

Our main technical contribution is stated in the following theorem.

Theorem 10. *Every formula over k variables $\psi \in \mathrm{L}_\mu[k]$ that defines the simulation closure $(\mathcal{K}, v_0)^{\lesssim}$ of a finite strongly connected structure is equivalent to an existential formula $\psi' \in \mathrm{L}_\mu[k]$.*

To prove that universal modalities can be safely eliminated from any formula ψ of the considered kind, we take a detour and first show that they can be eliminated from the formula expressing that a node at which ψ holds is reachable. To refer to this formula, we use a shorthand borrowed from temporal logics:

$$\mathsf{F}\psi := \mu X.\psi \vee \bigvee_{a\in \text{ACT}} \langle a\rangle X.$$

Lemma 14 in the second part of this section then states that from any formula equivalent to $\mathsf{F}\psi$, an existential formula equivalent to ψ can be recovered without increasing the number of variables.

Lemma 11. *Let $\mathcal{K}$ be a finite strongly connected structure and let $\psi^{\mathcal{K}}$ be the canonical axiom of its simulation closure $(\mathcal{K}, v_0)^{\lesssim}$. Then, every formula $\chi \equiv \mathsf{F}\psi^{\mathcal{K}}$ can be transformed, without increasing the number of variables, into an equivalent formula χ' with the following properties:*

(i) *no universal modalities occur in χ';*
(ii) *χ' is of shape $\mathsf{F}\psi$, where ψ contains no μ-operators;*
(iii) *every formula $\varphi \in \mathrm{cl}(\chi')$ holds at some vertex of $\mathcal{K}$.*

Proof. (i) Given an L_μ-formula χ, we say that a subformula $\langle a\rangle\varphi$ starting with a diamond is *vital*, if $\mathrm{cl}_\chi(\varphi)$ implies $\mathsf{F}\psi^{\mathcal{K}}$. Dually, a subformula $[a]\varphi$ starting with a box is vital, if the negation $\neg\,\mathrm{cl}_\chi(\varphi)$ implies $\mathsf{F}\psi^{\mathcal{K}}$.

Eliminating vital boxes. For $\chi \equiv \mathsf{F}\psi^{\mathcal{K}}$, let χ' be the formula obtained by replacing any occurrence of a vital box-subformula $[a]\varphi$ with *true*. Then, χ obviously implies χ'. For the converse, let us consider a tree model $\mathcal{T}$ of χ'. If, at all its nodes, $\mathcal{T}, v \models [a]\,\mathrm{cl}_\chi(\varphi)$ holds, then $\mathcal{T} \models \chi$. Else, there exists a node $v \in T$ with $\mathcal{T}, v \models \langle a\rangle\neg\,\mathrm{cl}_\chi(\varphi)$. But, since $[a]\varphi$ is vital, this means that $\mathcal{T}, v$ and hence $\mathcal{T}$ verifies $\mathsf{F}\psi^{\mathcal{K}}$. Either way, we obtain $\mathcal{T} \models \chi$ and hence $\chi \equiv \chi'$.

Eliminating non-vital modalities. By iterating the above elimination step a finite number of times, we obtain a formula $\chi \equiv \mathsf{F}\psi^{\mathcal{K}}$ without vital box-subformulae. Let now χ' be the formula obtained from χ by substituting simultaneously all remaining (i.e., non-vital) box-subformulae with *false* and all non-vital diamond-subformulae with *true*.

We will first show that the obtained formula χ' implies χ. Let $\mathcal{T}$ be a tree model of χ' and, for every non-vital subformula $\langle a\rangle\varphi$ of χ, let $\mathcal{T}_\varphi$ be a tree model of $\mathrm{cl}_\chi(\varphi) \wedge \neg\mathsf{F}\psi^{\mathcal{K}}$. At hand with the latter models, we construct an extension $\mathcal{T}'$ of $\mathcal{T}$ by introducing for every node $v \in T$ and every non-vital subformula $\langle a\rangle\varphi$ of χ, a fresh copy of $\mathcal{T}_\varphi$ to which we connect v via an a-edge.

Since χ' contains no box-subformulae, it is closed under extensions. Consequently $\mathcal{T}' \models \chi'$ and Verifier has a winning strategy σ in the model-checking game $\mathcal{G}(\mathcal{T}', \chi')$. Also, for every tree $\mathcal{T}_\varphi$, Verifier has a winning strategy σ_φ in the game $\mathcal{G}(\mathcal{T}_\varphi, \mathrm{cl}_\chi(\varphi))$. We can combine these strategies, to obtain a winning strategy for Verifier in the game $\mathcal{G}(\mathcal{T}', \chi)$ as follows. Move according to σ unless a position with a non-vital subformula of χ is met; up to that point, the play cannot leave T, otherwise, since $\mathsf{F}\psi^{\mathcal{K}}$ is falsified at any node $w \in T' \setminus T$, any vital subformula $\langle a \rangle \varphi$ would fail at w. Moreover, no subformula $[a]\varphi$ can occur, as it would correspond to a *false* position in $\mathcal{G}(\mathcal{T}', \chi')$. Consequently, σ leads the play to a position $(v, \langle a \rangle \varphi)$ where $v \in T$ and $\langle a \rangle \varphi$ is non-vital. At that event, let the Verifier choose the a-successor at the root of $\mathcal{T}_\varphi$ and proceed with his memoryless winning strategy σ_φ for the remaining game. In this way, Verifier finally wins any play of $\mathcal{G}(\mathcal{T}', \chi)$. Notice that, for all nodes $w \in T' \setminus T$, we have $\mathcal{T}', w \not\models \mathsf{F}\psi^{\mathcal{K}}$, and hence $\mathcal{T}'$ verifies $\mathsf{F}\psi^{\mathcal{K}}$ (or, equivalently, χ) if, and only if, $\mathcal{T}$ does. Hence, we have the following chain of implications, showing that χ' implies χ:

$$\mathcal{T} \models \chi' \implies \mathcal{T}' \models \chi' \implies \mathcal{T}' \models \chi \implies \mathcal{T} \models \chi.$$

For the converse, consider a tree model $\mathcal{T} \models \chi$ and, for every (non-vital) subformula $[a]\varphi$ of χ, a tree model $\mathcal{T}_{\neg}\varphi \models \neg\, \mathrm{cl}_\chi(\varphi) \wedge \neg \mathsf{F}\psi^{\mathcal{K}}$. As in the previous step, we construct an extension $\mathcal{T}'$ of $\mathcal{T}$ by connecting every node $v \in T$ via an a-edge to a fresh copy of $\mathcal{T}_{\neg}\varphi$, for every subformula $[a]\varphi$ of χ. Since $\chi \equiv \mathsf{F}\psi^{\mathcal{K}}$ is preserved under extensions, $\mathcal{T}'$ is still a model of χ. Let σ be a winning strategy for Verifier in the model-checking game $\mathcal{G}(\mathcal{T}', \chi)$. We will show that σ is also a winning strategy for Verifier in $\mathcal{G}(\mathcal{T}, \chi')$.

Notice that, in $\mathcal{G}(\mathcal{T}', \chi)$ Falsifier has a winning strategy from every position $(v, [a]\varphi)$ with $v \in T$, by moving to the a-successor of v at the root of $\mathcal{T}_{\neg}\varphi$. Consequently, any play according to Verifier's strategy σ will avoid such positions. Besides this, at every position $(v, \langle a \rangle \varphi)$ where $v \in T$ and $\langle a \rangle \varphi$ is a vital subformula of χ, the strategy σ will appoint a successor position (w, φ) with $w \in T$, otherwise, since any a-successor $w' \in T' \setminus T$ falsifies $\mathsf{F}\psi^{\mathcal{K}}$, φ would fail too. Summarising, every play of $\mathcal{G}(\mathcal{T}', \chi)$ according to σ, will avoid universal modalities and meet only nodes $v \in T$, unless a position a non-vital subformula $\langle a \rangle \varphi$ occurs. But under these conditions, we can replicate every play of $\mathcal{G}(\mathcal{T}', \chi)$ according to σ as a play of $\mathcal{G}(\mathcal{T}, \chi')$: in case a non-vital subformula $\langle a \rangle \varphi$ of χ is met in the former game, Verifier immediately wins $\mathcal{G}(\mathcal{T}, \chi')$, since the non-vital diamond-subformulae have been replaced by *true*. Otherwise, the outcome of the play is the same for both games and Verifier wins as well.

This concludes the proof that $\chi \equiv \chi'$.

(ii) By the above result, we can assume without loss that $\chi \equiv \mathsf{F}\psi^{\mathcal{K}}$ contains no box-modalities. For n being the number of states in $\mathcal{K}$, let ψ be the formula obtained by replacing every occurrence of a least fixed-point subformula $\mu X.\varphi$ in χ by it's n-th approximant φ^n. Then, by definition of the μ-operator, ψ implies χ and thus $\mathsf{F}\psi$ implies $\mathsf{F}\chi$, which is equivalent to χ. Conversely, since $\mathcal{K}, v_0 \models \chi$ and $\mathcal{K}$ has n states, we have $\mathcal{K}, v_0 \models \psi$. Since ψ is preserved under extensions, this means that $\psi^{\mathcal{K}}$ implies ψ. Accordingly $\mathsf{F}\psi^{\mathcal{K}}$, which is equivalent to χ, implies $\mathsf{F}\psi$. Hence, $\chi \equiv \mathsf{F}\psi$.

Note that the transformation of χ into $\mathsf{F}\psi$ does not increase the number of variables, as we can pick any of the variables already occurring in χ to expand the F-notation.

(iii) By the previous argument, we can assume that χ is of shape $\mathsf{F}\psi$ where ψ contains no boxes, i.e., $\chi = \mu X.\psi \vee \bigvee_a \langle a \rangle X$. Clearly, χ itself holds at every node of $\mathcal{K}$ and therefore, for every transition a occurring in $\mathcal{K}$, there is a node $v \in K$ where $\langle a \rangle \chi$, and thus $\mathrm{cl}_\chi\big(\langle a \rangle X\big)$, holds. Hence, any subformula φ of χ, with $\mathcal{K}, v \not\models \mathrm{cl}_\chi(\varphi)$ for all v, must actually be a subformula of ψ. Let ψ' be the formula obtained by replacing every such occurrence φ in ψ with *false*. On the one hand, ψ' then obviously implies ψ. On the other hand, as $\mathcal{K}, v_0 \models \mathsf{F}\psi$, there must exists a node v of $\mathcal{K}$ where ψ holds. At that node we also have $\mathcal{K}, v \models \psi'$ and, because ψ' is closed under extensions, this means that $\psi_v^{\mathcal{K}}$ implies ψ'. But then $\mathsf{F}\psi^{\mathcal{K}}$ implies $\mathsf{F}\psi'$ and, by $\mathsf{F}\psi \equiv \mathsf{F}\psi^{\mathcal{K}}$, it follows that $\mathsf{F}\psi$ implies $\mathsf{F}\psi'$. $\square$

Radical formulae and crisp models. Before we proceed towards proving the Preservation Theorem, we will introduce some notions which will be useful in the proof of Lemma 14

Given a formula $\psi \in \mathrm{L}_\mu$, we call a subformula φ *radical*, if it appears directly under a modal quantifier in ψ. We refer to the closure of radicals in ψ by

$$\mathrm{cl}_0(\psi) := \{\psi\} \cup \{\, \varphi \in \mathrm{cl}(\psi) \mid \langle a \rangle \varphi \in \mathrm{cl}(\psi) \text{ or } [a]\varphi \in \mathrm{cl}(\psi) \text{ for some } a \in \textsc{act}\,\}.$$

Radical formulae are the first to be met when a play of the model-checking game reaches a new node of the transition structure. For this reason, we need to care for game positions carrying radical formulae when merging strategies of different games.

Let $\mathcal{M}$ be a model of $\psi \in \mathrm{L}_\mu$ and σ a winning strategy for Verifier in $\mathcal{G}(\mathcal{M}, \psi)$. For any node $v \in M$, we define the *strategic type* of v in $\mathcal{M}$ under σ as follows:

$$\mathrm{tp}_\sigma^{\mathcal{M}}(v) := \{\, \varphi \in \mathrm{cl}_0(\psi) \mid \text{position } (v, \varphi) \text{ is reachable in } \mathcal{G}(\mathcal{M}, \psi) \text{ following } \sigma \,\}.$$

In arbitrary games, the type of a node can be rather complex. However, for existential formulae, Verifier has full control over the moves in the transition structure. In the ideal case, he can foresee for every node, a single radical formula to be proven there.

Given a transition structure $\mathcal{M}$ and a formula ψ, we say that a Verifier strategy σ in the model-checking game $\mathcal{G}(\mathcal{M}, \psi)$ is *crisp*, if the strategic type $\mathrm{tp}_\sigma^{\mathcal{M}}(v)$ of any $v \in M$ consists of not more than one radical. Accordingly, we call a model $\mathcal{M}$ of ψ crisp (under σ), if Verifier has a crisp winning strategy σ in the associated model-checking game.

The subsequent lemmas, that can be easily proved, provide us with sharp tools for manipulating models of existential formulae.

Lemma 12. *Every existential formula $\psi \in \mathrm{L}_\mu$ satisfied in some model $\mathcal{M} \models \psi$ also has a tree model $\mathcal{T}$ bisimilar to $\mathcal{M}$ which is crisp. Moreover, if $\mathcal{M}$ is finitely branching, then $\mathcal{T}$ can be chosen so as well.*

Lemma 13. *Let $\mathcal{T}$ be a crisp tree model of a formula $\psi \in \mathrm{L}_\mu$ under a strategy σ and let $x \in T$ be a node with strategic type $\mathrm{tp}_\sigma^{\mathcal{T}}(x) = \{\varphi\}$. Then, for every crisp tree model $\mathcal{S}$ of φ, the tree $\mathcal{T}[x/\mathcal{S}] \models \psi$, obtained by replacing the subtree of $\mathcal{T}$ rooted at x with $\mathcal{S}$, is still a crisp model of ψ.*

We are now ready for the final step, the elimination of the F-operator.

Lemma 14. *Let $\psi^{\mathcal{K}}$ be the canonical axiom for the simulation closure $(\mathcal{K}, v_0)^{\lesssim}$ of a finite strongly connected structure $\mathcal{K}$. Then, every formula ψ so that $\mathsf{F}\psi \equiv \mathsf{F}\psi^{\mathcal{K}}$ can be transformed, without increasing the number of variables, into a formula ψ' without universal modalities, so that $\psi' \equiv \psi^{\mathcal{K}}$.*

Proof. According to Lemma 11, we can assume that ψ contains no universal modalities or least fixed point operators and that (the closure of) every subformula is true at some node in $\mathcal{K}$.

We will first show that for any node v in $\mathcal{K}$, there is a subformula φ of ψ whose closure $\mathrm{cl}_\psi(\varphi)$ implies $\psi_v^{\mathcal{K}}$. Actually, we always find a radical formula with this property.

Towards a contradiction, let us assume that $\psi_v^{\mathcal{K}}$ is not implied by any radical subformula of ψ. This means that every $\varphi \in \mathrm{cl}_0(\psi)$ has a tree model $\mathcal{T}_\varphi$ which falsifies $\psi_v^{\mathcal{K}}$. According to Corollary 5, we can choose $\mathcal{T}_\varphi$ to be a finitely branching tree that falsifies already an approximant of $\psi_v^{\mathcal{K}}$ to some finite stage m_φ. Observe that this approximant $(\psi_v^{\mathcal{K}})[\nu := \nu^{m_\varphi}]$ is a modal formula. Let us denote its modal depth by n_φ. Further, let us fix a number n which is greater than any n_φ for $\varphi \in \mathrm{cl}_0(\psi)$ and co-prime to every number up to $|K|$.

By Lemma 12, we can assume without loss of generality that each $\mathcal{T}_\varphi$ is a crisp model of φ, this being witnessed by a crisp winning strategy for Verifier in the game $\mathcal{G}(\mathcal{T}_\varphi, \varphi)$. In particular, $\mathcal{T}_\psi$ is a crisp model of ψ. Let σ_ψ be a crisp winning strategy for Verifier in the model-checking game $\mathcal{G}(\mathcal{T}_\psi, \psi)$.

At hand with these, we construct a sequence of trees $(\mathcal{T}_i)_{0 \le i < \omega}$, together with crisp Verifier strategies σ_i witnessing that $\mathcal{T}_i \models \psi$. To start, we set $\mathcal{T}_0 := \mathcal{T}_\psi$ and $\sigma_0 := \sigma_\psi$. In every step $i > 0$, the tree $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by performing the following manipulations at depth $n(i+1)$. For each subtree of $\mathcal{T}_i$ rooted at a node x of this depth, we check whether $\mathcal{T}_i, x \models \psi_v^{\mathcal{K}}$. If this is not the case, the subtree remains unchanged. Else, we look at the strategic type of x under σ_i. If the type is empty, we simply cut all successors of x. Otherwise, $\mathrm{tp}_{\sigma_i}^{\mathcal{T}_i}(x)$ consists of a single radical formula φ, and we replace the subtree $\mathcal{T}_i, x$ with $\mathcal{T}_\varphi$. According to Lemma 13, the resulting tree $\mathcal{T}_{i+1}$ is a model of ψ, and the composition of the strategy σ_i with the crisp strategies σ_φ on the newly appended subtrees $\mathcal{T}_\varphi$ yields a crisp Verifier strategy σ_{i+1} for the model-checking game $\mathcal{G}(\mathcal{T}_{i+1}, \psi)$.

By construction, each of the trees $\mathcal{T}_i$ is finitely branching and the sequence $(\mathcal{T}_i)_{0 \le i < \omega}$ converges in the prefix topology of finitely branching trees (see [11]). Let $\mathcal{T}_\omega$ be the limit of this sequence. Since no μ-operators occur in ψ, its model class is topologically closed on finitely branching trees, according to [11]. Consequently, $\mathcal{T}_\omega$ is still a model of ψ. By our hypothesis, ψ implies $\mathsf{F}\psi^{\mathcal{K}}$. Thus, at some depth d in $\mathcal{T}_\omega$ a node x with $\mathcal{T}_\omega, x \models \psi_v^{\mathcal{K}}$ appears. Since $\mathcal{K}$ is strongly connected, v must lie on a cycle in $\mathcal{K}$. Hence, for $k \le |K|$ being the length of such a cycle, there exist nodes y with $\mathcal{T}_\omega, y \models \psi_v^{\mathcal{K}}$ at every depth $d + jk$. However, our construction eliminated all subtrees carrying the similarity type of v at depths multiple of n. Since n was chosen to be co-prime to any integer up to $|K|$, it follows that $\mathcal{T}_\omega$ cannot satisfy ψ. This is a contradiction which invalidates our assumption that $\psi_v^{\mathcal{K}}$ is not implied by any $\varphi \in \mathrm{cl}_0(\psi)$.

Hence, for every node $v \in K$, there exists a formula $\varphi_v \in \mathrm{cl}_0(\psi)$ implying $\psi_v^{\mathcal{K}}$. We can show that the converse also holds, if v is maximal with respect to the preorder $\lesssim$, in the sense that for every w with $v \lesssim w$ we have $w \lesssim v$. Recall that, by Lemma 11 (iii), the formula φ_v must be verified at some node w in $\mathcal{K}$. Since φ_v is existential and thus preserved under extension, it follows that $\psi_w^{\mathcal{K}}$ implies φ_v, which further implies $\psi_v^{\mathcal{K}}$. But this means that $v \lesssim w$ and, by maximality of v, that w and v are bisimilar. Hence, $\mathcal{K}, v \models \varphi_v$ and consequently $\psi_v^{\mathcal{K}} \equiv \varphi_v$.

This concludes the proof for the case when v_0 is maximal in $\mathcal{K}$ with respect to $\lesssim$. Otherwise, we could not guarantee, of course, that $\varphi_{v_0} \equiv \psi_{v_0}^{\mathcal{K}}$. But in that case, a formula equivalent to $\psi_v^{\mathcal{K}}$ can be recovered from $\mathrm{cl}_0(\psi)$ without great difficulty. □

4 The Hierarchy Theorem

In [5], it was shown that every level k of the variable hierarchy contains existential formulae which are not equivalent to any existential formula from a lower hierarchical level. Examples of such formulae are obtained by considering so-called clique structures $\mathcal{C}^k$ over the set of states $\{0, \ldots, k-1\}$ with transition relations $E_{ij} = \{(i,j)\}$, for $0 \leq i, j < k$. For each k, the canonical axiom ψ^k of the simulation closure of $\mathcal{C}^k$ is an existential L_μ-formula over k variables. The Hierarchy Theorem for the existential fragment states that, if we restrict to formulae using only existential quantification, k variables are indeed necessary.

Theorem 15 ([5]). *For every $k > 0$, the simulation closure of $\mathcal{C}^k$ cannot be defined by any existential formula in $\mathrm{L}_\mu[k-1]$.*

However, this left open the question whether a formula from $\mathrm{L}_\mu[k-1]$ which uses universal quantification may be equivalent to ψ^k. Due to our Preservation Theorem, we are now able to assert that this cannot be the case.

Theorem 16. *For every $k > 0$, the formula $\psi^k \in \mathrm{L}_\mu[k]$ defining the simulation closure of $\mathcal{C}^k$ is not equivalent to any formula in $\mathrm{L}_\mu[k-1]$.*

Proof. Let us assume that there exists a formula $\psi \in \mathrm{L}_\mu[k-1]$ equivalent to ψ^k. Since ψ defines the simulation closure of $\mathcal{C}^k$, a finite strongly connected structure, we can apply Theorem 10 to conclude that there also exists a formula $\psi' \in \mathrm{L}_\mu[k-1]$ using only existential modalities which is equivalent to ψ^k. But this contradicts the Hierarchy Theorem 15 for the existential fragment. □

As a direct consequence, we can separate the expressive power of Parikh's Game Logic [16] and the μ-calculus, thus answering an open question posed by Pauly in [17]. Since Game Logic can be translated into the two variable fragment of L_μ, its expressive power is strictly subsumed already by $\mathrm{L}_\mu[3]$.

Corollary 17. *The modal μ-calculus is strictly more expressive than Game Logic interpreted over transition structures.*

Notice that the examples of strict formulae for $\mathrm{L}_\mu[k]$ given in [5] use a vocabulary consisting of k^2 actions. In a forthcoming paper [6], we provide examples of hard formulae over a fixed alphabet of only two actions for every level k.

Literature

1. A. Arnold, *The mu-calculus alternation-depth is strict on binary trees*, RAIRO Informatique Thorique et Applications, 33 (1999), pp. 329–339.
2. A. Arnold and D. Niwiński, *Rudiments of μ-calculus*, North Holland, 2001.
3. D. Berwanger, *Game logic is strong enough for parity games*, Studia Logica, 75 (2003), pp. 205–219. Special issue on Game Logic and Game Algebra edited by M. Pauly and R. Parikh.
4. D. Berwanger and E. Grädel, *Games and model checking for guarded logics*, in Proceedings of LPAR 2001, Lecture Notes in Computer Science Nr. 2250, Springer, 2001, pp. 70–84.
5. D. Berwanger, E. Grädel, and G. Lenzi, *On the variable hierarchy of the modal mu-calculus*, in Computer Science Logic, CSL 2002, J. Bradfield, ed., vol. 2471 of LNCS, Springer-Verlag, 2002, pp. 352–366.
6. D. Berwanger and G. Lenzi, *Robbers, guardians, and a few diamonds. On the variable hierarchy of the bimodal existential μ-calculus.* Forthcoming.
7. J. Bradfield, *The modal μ-calculus alternation hierarchy is strict*, Theoretical Computer Science, 195 (1998), pp. 133–153.
8. G. d'Agostino and M. Hollenberg, *Logical questions concerning the μ-calculus: interpolation, Lyndon, and łos-Tarski*, Journal of Symbolic Logic, 65 (2000), pp. 310–332.
9. A. Emerson and C. Jutla, *Tree automata, mu-calculus and determinacy*, in Proc. 32nd IEEE Symp. on Foundations of Computer Science, 1991, pp. 368–377.
10. M. C. B. Hennessy and R. Milner, *On observing nondeterminism and concurrency*, in Automata, Languages and Programming, 7th Colloquium, J. W. de Bakker and J. van Leeuwen, eds., vol. 85 of LNCS, Springer-Verlag, 1980.
11. D. Janin and G. Lenzi, *On the logical definability of topologically closed recognizable languages of infinite trees*, Computing and Informatics, 21 (2002), pp. 185–203.
12. D. Kozen, *A finite model theorem for the propositional μ-calculus*, Studia Logica, 47 (1983), pp. 233–241.
13. ———, *Results on the propositional μ-calculus*, Theoretical Computer Science, 27 (1983), pp. 333–354.
14. O. Kupferman, M. Vardi, and P. Wolper, *An automata-theoretic approach to branching-time model checking*, Journal of the ACM, 47 (2000), pp. 312–360.
15. G. Lenzi, *A hierarchy theorem for the mu-calculus*, in Proceedings of the 23rd International Colloquium on Automata, Languages and Programming, ICALP '96, F. Meyer auf der Heide and B. Monien, eds., vol. 1099 of Lecture Notes in Computer Science, Springer-Verlag, July 1996, pp. 87–97.
16. R. Parikh, *The logic of games and its applications*, Annals of Discrete Mathematics, 24 (1985), pp. 111–140.
17. M. Pauly, *Logic for Social Software*, PhD thesis, University of Amsterdam, 2001.
18. C. Stirling, *Bisimulation, modal logic and model checking games*, Logic Journal of the IGPL, 7 (1999), pp. 103–124.