



Modal Separation of Fixpoint Formulae

Jean Christoph Jung  

TU Dortmund University

Jędrzej Kołodziejcki  

TU Dortmund University

Abstract

Modal separability for modal fixpoint formulae is the problem to decide for two given modal fixpoint formulae φ, φ' whether there is a modal formula ψ that separates them, in the sense that $\varphi \models \psi$ and $\psi \models \neg\varphi'$. We study modal separability and its special case modal definability over various classes of models, such as arbitrary models, finite models, trees, and models of bounded outdegree. Our main results are that modal separability is PSPACE-complete over words, that is, models of outdegree ≤ 1 , EXPTIME-complete over unrestricted and over binary models, and 2-EXPTIME-complete over models of outdegree bounded by some $d \geq 3$. Interestingly, this latter case behaves fundamentally different from the other cases also in that modal logic does not enjoy the Craig interpolation property over this class. Motivated by this we study also the induced interpolant existence problem as a special case of modal separability, and show that it is CONEXPTIME-complete and thus harder than validity in the logic. Besides deciding separability, we also investigate the problem of efficient construction of separators. Finally, we consider in a case study the extension of modal fixpoint formulae by graded modalities and investigate separability by modal formulae and graded modal formulae.

2012 ACM Subject Classification Theory of computation \rightarrow Modal and temporal logics

Keywords and phrases Modal Logic, Fixpoint Logic, Separability, Interpolation

Digital Object Identifier 10.4230/LIPIcs.STACS.2025.63

1 Introduction

For given logics $\mathcal{L}, \mathcal{L}^+$, the \mathcal{L} -separability problem for \mathcal{L}^+ is to decide given two \mathcal{L}^+ -formulae φ, φ' whether there is an \mathcal{L} -formula ψ that separates φ and φ' in the sense that $\varphi \models \psi$ and $\psi \models \neg\varphi'$. Obviously, a separator can only exist when φ and φ' are mutually exclusive, and the problem is only meaningful when \mathcal{L} is less expressive than \mathcal{L}^+ . Intuitively, a separator formulated in a “simpler” logic \mathcal{L} explains a given inconsistency in a “complicated” logic \mathcal{L}^+ . Note that, for logics \mathcal{L}^+ closed under negation, \mathcal{L} -separability generalizes the \mathcal{L} -definability problem for \mathcal{L}^+ : decide whether a given \mathcal{L}^+ -formula is equivalent to an \mathcal{L} -formula. Indeed, $\varphi \in \mathcal{L}^+$ is equivalent to an \mathcal{L} -formula iff φ and $\neg\varphi$ are \mathcal{L} -separable. Since separability is more general than definability, solving it requires an even better understanding of the logics under consideration. Both separability and definability are central problems with many applications in computer science. As seminal work let us only mention definability and separability of regular word languages by first-order logic [26, 29, 9].

In this paper we study definability and separability of formulae of the modal μ -calculus μ ML [27, 20] by formulae in propositional modal logic ML. μ ML is the extension of ML with fixpoints that encompasses virtually all specification languages such as PDL [12] and LTL and CTL [3]. Let us consider an example.

► **Example 1.** Consider the following properties P_1, P_2, P_3 of vertex-labelled trees:

- P_1 : there is an infinite path starting in the root on which each point satisfies a ;
- P_2 : on every path there are only finitely many points satisfying a ;
- P_3 : on every path at most two points satisfy a .

The properties are expressible in μ ML but not in ML, and both P_1, P_2 and P_1, P_3 are mutually exclusive. The properties P_1, P_3 are separated by the ML-formula $\psi = a \wedge \diamond(a \wedge \diamond a)$ which



© Jean Christoph Jung and Jędrzej Kołodziejcki;
licensed under Creative Commons License CC-BY 4.0

42nd International Symposium on Theoretical Aspects of Computer Science (STACS 2025).

Editors: Olaf Beyersdorff, Michał Pilipczuk, Elaine Pimentel, and Nguyen Kim Thang; Article No. 63; pp. 63:1–63:32



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

	all models	words	binary trees	d -ary trees, $d \geq 3$
ML-definability	EXPTIME [24]	PSPACE	EXPTIME [24]	EXPTIME [24]
ML-separability	EXPTIME	PSPACE	EXPTIME	2-EXPTIME
separator construction	double exp.	single exp.	double exp.	triple exp.
ML interpolant existence	always	always	always	CONEXPTIME

■ **Table 1** Overview of our results. All complexity results are completeness results.

expresses that there is a path starting with three points satisfying a . On the other hand, no ML-formula separates P_1, P_2 . The intuitive reason for this is that any ML-formula ψ only sees trees up to depth $|\psi|$, and one can find two trees with properties P_1, P_2 which nonetheless look the same up to depth $|\psi|$. ◀

We explore the definability and separability problems over several classes of models relevant for computer science: all models, words, trees of bounded or unbounded outdegree; as well as restrictions of all these classes to finite models. On top of analyzing the decision problems, we also address the problem of constructing efficient definitions and separators whenever they exist. The starting point for our research is the seminal paper of Otto [24], where he solves modal definability over models of bounded and unbounded outdegree. In this paper, we continue this line of research and establish a fairly complete and interesting picture. Table 1 summarizes our results. We now explain its content further.

The first line essentially repeats Otto’s results; we only add the observation that ML-definability over words is PSPACE-complete. Interestingly, separability is substantially more difficult. The case of words is the easiest one, both in terms of computational complexity and required arguments. Next come the cases of binary and of unrestricted trees. These two classes possess some nice structural properties which (although true for different reasons) enable a common algorithmic treatment. Finally, the cases of trees with outdegree bounded by a number $d \geq 3$ enter the stage. These trees lack the good properties essential for previous constructions which results in higher computational complexity. The hardness result for $d \geq 3$ is interesting for two reasons. First, as it is entirely standard to encode trees of higher outdegree into binary ones, one could expect the ternary (and higher) case to have the same complexity as the binary one. And second, even though there are known cases when separation is provably harder than definability (regularity of visibly pushdown languages is decidable [23, Theorem 19] but regular separability thereof is not [19, Theorem 2.4]), to the best of our knowledge our results are the only such case known in logic.

The complexity landscape for deciding separability is also reflected in the maximal sizes of the separators that we construct. Relying on the well-known connection of μ ML to automata, we provide effective constructions for the cases of all models, words, and binary trees. It is worth mentioning that equally effective constructions for definability over all models are given in [22], but they do not work for separability. The ternary case follows from a general argument. Our construction of separators over words is optimal. Under mild assumptions (there are at least two modalities) the constructions over binary and over unrestricted trees are optimal as well, but we leave it open whether these assumptions are needed for the lower bounds. In the case of ternary and higher outdegree trees we only conjecture optimality of the constructed separators.

Finally, we observe that ML lacks the Craig interpolation property over trees of outdegree bounded by $d \geq 3$. Recall that a *Craig interpolant* for $\varphi \models \varphi'$ in some logic \mathcal{L} is a formula $\psi \in \mathcal{L}$ only using the common symbols of φ and φ' and such that $\varphi \models \psi \models \varphi'$. A logic satisfies the *Craig interpolation property (CIP)* if a Craig interpolant of $\varphi \models \varphi'$ always exists.

It is known that ML enjoys CIP over all models and over words [15] and it follows from our techniques that this transfers to binary trees. In contrast and as mentioned above, over ternary and higher-arity trees ML lacks the CIP. It is worth mentioning that modal logic over frames of arity bounded by some d has been studied under the name $\mathbf{K} \oplus \mathit{alt}_d$ [4]. Our results imply that $\mathbf{K} \oplus \mathit{alt}_d$ enjoys CIP iff $d \leq 2$. Motivated by the lack of CIP over higher-arity trees, we study the induced interpolant existence problem – determining whether two given ML-formulae φ, φ' admit a Craig interpolant – as a special case of separability. We show it to be CONEXPTIME-complete over higher arity trees, and thus harder than validity. Interpolant existence has recently been studied for other logics without CIP [18, 1].

As an application of our results for d -ary trees with $d \geq 3$ we additionally present a case study: separability in the *graded* setting in which we allow counting modalities saying “there are at least k children such that [...]” [11]. Counting modalities are a standard extension of modal logic that is especially relevant in applications in knowledge representation for conceptual modeling [2]. We show that ML-separability of graded μML is 2-EXPTIME-complete, while it is EXPTIME-complete if we allow counting modalities also in the separator. The intuitive reason for the hardness in the former case is that trees of bounded arity are definable in graded μML . This former case is also related with a recent study about separating logics supporting counting quantifiers by logics without these [21].

It is worth to mention that ML-definability of μML -formulae generalizes the *boundedness problem* which asks whether a formula with a single fixpoint is equivalent to a modal formula. Boundedness has been studied for other logics such as monadic-second order logic [6], datalog [16], and the guarded fragment of first-order logic [5]. Our paper is an extension of the preliminary paper [17].

The paper is organized as follows. After this introduction 1, we set notation and recall basic facts in the preliminary Section 2. Next, we introduce some topic-specific terminology, discuss a relevant construction of Otto, and solve the case of all models in Section 3. In the following Sections 4 and 5 we deal with unary and binary trees, and in Section 6 we solve the most challenging case of trees of outdegree bounded by $d \geq 3$. Section 7 applies our results to the case with graded modalities. The last Section 8 contains conclusions and final remarks.

2 Preliminaries

We recall the main notions about modal logic ML and the modal μ -calculus μML . For the rest of this paper fix disjoint, countably infinite sets \mathbf{Prop} of *atomic propositions* and \mathbf{Var} of *variables*. The syntax of μML is given by the rule

$$\varphi ::= \tau \mid \neg\tau \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond\varphi \mid \square\varphi \mid x \mid \mu x.\varphi \mid \nu x.\varphi$$

where $\tau \in \mathbf{Prop}$ and $x \in \mathbf{Var}$. We assume that formulae of μML are in a normal form such that every $x \in \mathbf{Var}$ appears at most once in a formula, and if it does appear then its appearance has a unique superformula ψ beginning with μx or νx . Modal logic ML is defined as the fragment of μML with no fixpoint operators μ and ν nor variables. Both in ML and μML , we use abbreviations like \top (for $a \vee \neg a$ for some $a \in \mathbf{Prop}$), $\diamond^n\varphi$ (for a formula $\diamond \dots \diamond\varphi$ with n leading \diamond 's), and $\neg\varphi$. We denote with $\text{sig}(\varphi)$ the set of propositions that occur in φ , and recall that the *modal depth* of an ML formula is the maximal nesting of \diamond, \square . With ML^n we denote the class of all ML-formulae of modal depth at most n , and with ML_σ^n we denote its subclass restricted to signature σ . The *size* $|\varphi|$ of a formula φ is the length of φ represented as a string. This choice of the simplest possible measure of size does not matter for most of our results. We will briefly discuss alternative notions of size in the concluding Section 8.

Both ML and μ ML are interpreted in pointed Kripke structures. More formally, a model \mathcal{M} is a quadruple $\mathcal{M} = (M, v_I, \rightarrow, \text{val})$ consisting of a set M called its *universe*, a distinguished point $v_I \in M$ called the *root*, an accessibility relation $\rightarrow \subseteq M \times M$, and a valuation $\text{val} : M \rightarrow \mathcal{P}(\text{Prop})$.

The semantics of μ ML can be defined in multiple equivalent ways. The one most convenient for us is through parity games (see [32] for an introduction). Given a model \mathcal{M} and a formula $\varphi \in \mu\text{ML}$ we define a *semantic game* $\mathcal{G}(\mathcal{M}, \varphi)$ played between players \exists ve and \forall dam. The positions are $M \times \text{SubFor}(\varphi)$. The moves depend on the topmost connective. From a position of the shape $(v, \psi \vee \psi')$ or $(v, \psi \wedge \psi')$ it is allowed to move to either (v, ψ) or (v, ψ') . From $(v, \diamond\psi)$ and $(v, \square\psi)$ the allowed moves lead to all (w, ψ) such that $v \rightarrow w$. In position (v, τ) or $(v, \neg\tau)$ the game stops and \exists ve wins iff v satisfies the formula component τ or $\neg\tau$, respectively. From $(v, \mu x.\psi)$ and $(v, \nu x.\psi)$ the game moves to (v, ψ) , and from (v, x) to (v, ψ) where ψ is the unique superformula of x beginning with μx or νx . \exists ve owns positions whose formula component has \vee or \diamond as the topmost connective and \forall dam owns all other positions. \exists ve wins an infinite play π if the outermost subformula seen infinitely often in π begins with ν . We say that \mathcal{M}, v *satisfies* φ and write $\mathcal{M}, v \models \varphi$ if \exists ve wins the game $\mathcal{G}(\mathcal{M}, \varphi)$ from position (v, φ) . Since \mathcal{M} is by definition pointed, we abbreviate $\mathcal{M}, v_I \models \varphi$ with $\mathcal{M} \models \varphi$.

The same symbol denotes entailment: $\varphi \models \psi$ means that every model of φ is a model of ψ . In the case only models from some fixed class \mathbf{C} are considered we talk about satisfiability and entailment *over* \mathbf{C} . Let \mathcal{L} be a subset of μML such as ML or ML_σ^n . If two models \mathcal{M} and \mathcal{N} satisfy the same formulae of \mathcal{L} then we call them \mathcal{L} -equivalent and write $\mathcal{M} \equiv_{\mathcal{L}} \mathcal{N}$.

In the paper we will study models of bounded and unbounded outdegree. The *outdegree* of a point $w \in M$ in a model $\mathcal{M} = (M, v_I, \rightarrow, \text{val})$ is the number of successors of w in the underlying directed graph $G_{\mathcal{M}} = (M, \rightarrow)$. We say that \mathcal{M} has *finite outdegree* if every point has finite outdegree and *bounded outdegree* if there is a finite uniform upper bound d on the outdegree of its points. In the latter case, we will call \mathcal{M} *d-ary*, and *binary* or *ternary* if $d = 2$ or $d = 3$. If $d = 1$, then we call \mathcal{M} a *word*. A *d-ary* model is *full* if each of its nodes is either a leaf (i.e. has no children) or has precisely d children. A model \mathcal{M} is a *tree* if $G_{\mathcal{M}}$ is a (directed) tree with root v_I . We denote with \mathbb{T}^d the class of all *d-ary* tree models. Both ML and μML are invariant under bisimulation, and every (*d-ary*) model is bisimilar to a (*d-ary*) tree. Hence, we do not lose generality by only looking at tree models.

A *prefix* of a tree is a subset of its universe closed under taking ancestors. When no confusion arises we identify a prefix $N \subseteq M$ with the induced subtree \mathcal{N} of \mathcal{M} that has N as its universe. The *depth* of a point is the distance from the root. The prefix of depth n (or just *n-prefix*) is the set of all points at depth at most n and is denoted by $M_{|_n}$ (and the corresponding subtree by $\mathcal{M}_{|_n}$).

Bisimulations

We define bisimulations and bisimilarity for trees, assuming for convenience that bisimulations only link points at the same depth. Let $\mathcal{M}, \mathcal{M}'$ be trees and $Z \subseteq M \times M'$ a relation between M and M' that relates only points of the same depth. Then, Z is a *bisimulation between \mathcal{M} and \mathcal{M}'* if it links the roots $v_I Z v'_I$, and for every $w Z w'$ the following conditions are satisfied:

(atom) $\text{val}(w) = \text{val}'(w')$,

(forth) for every $v \in M$ with $w \rightarrow v$ there is a $v' \in M'$ with $w' \rightarrow v'$ and $v Z v'$, and

(back) for every $v' \in M'$ with $w' \rightarrow v'$ there is a $v \in M$ with $w \rightarrow v$ and $v Z v'$.

A *functional bisimulation* (also known as *bounded morphism*) is a function whose graph is a bisimulation. If Z is a functional bisimulation from \mathcal{M} to \mathcal{M}' then we write $Z : \mathcal{M} \xrightarrow{\text{bis}} \mathcal{M}'$ and call \mathcal{M}' a *bisimulation quotient* of \mathcal{M} . The *bisimilarity quotient* of \mathcal{M} is a quotient \mathcal{M}' of \mathcal{M}

such that if $Z' : \mathcal{M}' \rightarrow \mathcal{M}''$ then $\mathcal{M}' = \mathcal{M}''$. It follows from analogous results for arbitrary models that every tree $\mathcal{M} \in \mathbb{T}^d$ has a unique (up to isomorphism) bisimilarity quotient $\mathcal{M}' \in \mathbb{T}^d$ and that two trees are bisimilar iff their bisimilarity quotients are isomorphic.

Further, for every $n \in \mathbb{N}$ and every subset $\sigma \subseteq \mathbf{Prop}$ of the signature we consider a restricted variant of bisimulations called (σ, n) -bisimulations. In a (σ, n) -bisimulation the atom condition is only checked with respect to σ and the back and forth conditions only for points at depth smaller than n . Formally, a relation $Z \subseteq M \times M'$ is a (σ, n) -bisimulation if it is a bisimulation between the n -prefixes of the σ -reducts of $\mathcal{M}, \mathcal{M}'$. We call a (σ, n) -bisimulation between $\mathcal{M}, \mathcal{M}'$ a (σ, n) -isomorphism if it is bijective on the n -prefixes of $\mathcal{M}, \mathcal{M}'$. We write $\mathcal{M} \simeq_{\sigma}^n \mathcal{M}'$ if there exists a (σ, n) -bisimulation between \mathcal{M} and \mathcal{M}' and $\mathcal{M} \cong_{\sigma}^n \mathcal{M}'$ if there is a (σ, n) -isomorphism between them. Crucially, over every class \mathbf{C} of models and for every finite σ the equivalences $\equiv_{\mathbf{ML}_{\sigma}^n}$ and \simeq_{σ}^n coincide, for every n .

Automata

We exploit the well-known connection of $\mu\mathbf{ML}$ and automata that read tree models. A *nondeterministic parity tree automaton (NPTA)* is a tuple $\mathcal{A} = (Q, \Sigma, q_I, \delta, \mathbf{rank})$ where Q is a finite set of states, $q_I \in Q$ is the initial state, $\Sigma = \mathcal{P}(\sigma)$ for some finite set $\sigma \subseteq \mathbf{Prop}$, \mathbf{rank} assigns each state a priority, and δ is a transition function of type:

$$\delta : Q \times \Sigma \rightarrow \mathcal{P}(Q^{\leq d}),$$

where $Q^{\leq d}$ denotes the set of all tuples over Q of length at most d . A run of \mathcal{A} on a tree \mathcal{M} is an assignment $\rho : M \rightarrow Q$ sending the root of the tree to q_I and consistent with δ in the sense that $(\rho(v_1), \dots, \rho(v_k)) \in \delta(\rho(v), \mathbf{val}(v) \cap \sigma)$ for every point v with children v_1, \dots, v_k . On occasion when considering trees of unbounded outdegree we will use automata with transition function of type $\delta : Q \times \Sigma \rightarrow \mathcal{P}(\mathcal{P}(Q))$. Then, consistency of ρ with δ means that $\{\rho(v') \mid v' \in V\} \in \delta(\rho(v), \mathbf{val}(v) \cap \sigma)$ for every v with a set V of children. In either case, we call the run ρ *accepting* if for every infinite path $v_0, v_1 \dots$ in \mathcal{M} the sequence $\mathbf{rank}(\rho(v_0)), \mathbf{rank}(\rho(v_1)), \dots$ satisfies the parity condition. We write $\mathcal{M} \models \mathcal{A}$ in case \mathcal{A} has an accepting run on \mathcal{M} . An automaton that is identical to \mathcal{A} except that the original initial state is replaced with q is denoted $\mathcal{A}[q_I \leftarrow q]$. The *size* of an automaton \mathcal{A} is the number of its states and is denoted by $|\mathcal{A}|$.

An NPTA \mathcal{A} is *equivalent* to a formula $\varphi \in \mu\mathbf{ML}$ over a class \mathbf{C} of trees when $\mathcal{M} \models \varphi$ iff $\mathcal{M} \models \mathcal{A}$ for every tree $\mathcal{M} \in \mathbf{C}$. We rely on the following classical result (see for example the discussion in [31] and the well-presented Dealternation Theorem 5.7 in [7]):

► **Theorem 2.** *For every $\mu\mathbf{ML}$ -formula φ and class \mathbf{C} of trees, we can construct an NPTA with exponentially many states equivalent to φ over \mathbf{C} . The construction takes exponential time when $\mathbf{C} \subseteq \mathbb{T}^d$ for some d , and doubly exponential time in the unrestricted case.*

3 Foundations of Separability

We start with recalling the notion of separability and discuss some of its basic properties.

► **Definition 3.** *Assume a subset \mathcal{L} of all $\mu\mathbf{ML}$ formulae. Given $\varphi, \varphi' \in \mu\mathbf{ML}$, an \mathcal{L} -separator of φ, φ' is a formula $\psi \in \mathcal{L}$ with $\varphi \models \psi$ and $\psi \models \neg\varphi'$. If additionally $\mathbf{sig}(\psi) \subseteq \sigma$ for some signature σ , ψ is called an \mathcal{L}_{σ} -separator.*

The \mathcal{L} -separability problem is to determine, given formulae $\varphi, \varphi' \in \mu\mathbf{ML}$ and a signature σ , if they admit an \mathcal{L}_{σ} -separator ψ . \mathcal{L} -definability is the special case of \mathcal{L} -separability in which

$\varphi' = \neg\varphi$, since an \mathcal{L} -separator of $\varphi, \neg\varphi$ is equivalent to φ . All notions can be relativized to a class \mathbf{C} of models by considering entailment over that class. We investigate ML-separability and ML-definability over different classes of models. The reader may have expected the problems to be defined without restrictions on σ , but in fact such versions of the problems are special instances of our problems with $\sigma = \text{sig}(\varphi) \cup \text{sig}(\varphi')$. Conversely, all lower bounds already hold for such special instances.

We start with observing that, by the tree model property and the finite model property of μML , ψ is an ML_σ -separator of φ, φ' (over all models) iff ψ is an ML_σ -separator of φ, φ' over trees iff ψ is an ML_σ -separator of φ, φ' over finite models. Thus, separability coincides over all these classes. Moreover, with the help of the μML -formula $\theta_\infty = \nu x. \diamond x$ expressing the existence of an infinite path originating in the root, ML-separability over finite trees reduces to ML-separability over all models. More formally:

► **Lemma 4.** *Let $\varphi, \varphi' \in \mu\text{ML}$ and $\psi \in \text{ML}$. Then ψ is an ML_σ -separator of φ, φ' over finite trees iff ψ is an ML_σ -separator of $\varphi \wedge \neg\theta_\infty, \varphi' \wedge \neg\theta_\infty$. This is also true inside \mathbb{T}^d , for $d \in \mathbb{N}$.*

This lemma allows us to transfer all upper bounds obtained in the paper also to the restrictions of the classes to finite models. The lower bounds do not follow from this lemma, but analyzing the proofs yields that they actually work as well. Thus, in the rest of the paper we focus on the classes of all models and \mathbb{T}^d , for $d \in \mathbb{N}$.

The starting point for the technical developments in the paper are model-theoretic characterizations for separability. Similar to what has been done in the context of interpolation, see for example [28], they are given in terms of joint consistency, which we introduce next. Let R be a binary relation on some class of models, such as (σ, n) -isomorphism \cong_σ^n or ML_σ^n -equivalence $\equiv_{\text{ML}_\sigma^n}$. We call two formulae φ, φ' *joint consistent up to R* (in short *joint R -consistent*) if there are models $\mathcal{M} \models \varphi$ and $\mathcal{M}' \models \varphi'$ with $R(\mathcal{M}, \mathcal{M}')$. For technical reasons we will sometimes also talk about joint consistency of automata $\mathcal{A}, \mathcal{A}'$ in place of formulae φ, φ' . Joint R -consistency over a class \mathbf{C} of models is defined by only looking at models from \mathbf{C} . Clearly, if $R' \subseteq R$ and $\mathbf{C}' \subseteq \mathbf{C}$ then joint R' -consistency over \mathbf{C}' implies joint R -consistency over \mathbf{C} . We use the following standard equivalence:

$$\varphi, \varphi' \text{ are not } \text{ML}_\sigma^n\text{-separable over } \mathbf{C} \iff \varphi, \varphi' \text{ are joint } \equiv_\sigma^n\text{-consistent over } \mathbf{C}. \quad (\text{Base})$$

for every $\varphi, \varphi' \in \mu\text{ML}$, $n \in \mathbb{N}$, finite σ , and class \mathbf{C} . The implication from right to left is immediate. The opposite one follows from the observation that for every $n \in \mathbb{N}$ and finite σ there are only finitely many equivalence classes of \equiv_σ^n , and each such class is fully described with a single modal formula.

Let us illustrate how Equivalence (Base) is used to solve ML-separability. Let φ_1 and φ_2 be μML -formulae expressing the respective properties P_1 and P_2 from Example 1. Let \mathcal{M} be an infinite path in which every point satisfies a , and let \mathcal{M}_n be a finite path of length n in which every point satisfies a . Then, for each n the models $\mathcal{M}, \mathcal{M}_n$ witness joint \equiv_σ^n -consistency of φ_1, φ_2 . By Equivalence (Base) this means that φ_1, φ_2 are not ML^n -separable for any n , and thus not ML-separable at all.

Definability is a special case of separability. Since the tools used for solving definability are a starting point for our work, we recall them now.

Modal Definability: A Recap

In his seminal paper [24] Otto showed that ML-definability of μML -formulae is EXPTIME -complete over all models and over \mathbb{T}^d for every $d \geq 2$.

► **Theorem 5.** [24, Main Theorem and Proposition 5] *Over the class of all models, as well as over \mathbb{T}^d for every $d \geq 2$, ML-definability of μML -formulae is EXPTIME-complete.*

We start by recalling and rephrasing Otto's construction and fixing a small mistake in the original proof. The lower bound follows by an immediate reduction from satisfiability of μML -formulae. We look at the upper bound. The first step is the following lemma, which is the heart of [24, Lemma 2].

► **Lemma 6.** *For every $\varphi \in \mu\text{ML}$ and $n, d \in \mathbb{N}$ the following are equivalent:*

1. $\varphi, \neg\varphi$ are joint \cong_σ^n -consistent over \mathbb{T}^d .
2. $\varphi, \neg\varphi$ are joint \cong_σ^n -consistent over \mathbb{T}^d .

The lemma is true, but its proof in [24] is mistaken. The problem there is that the construction duplicates subtrees and hence may turn d -ary models into ones with outdegree greater than d . We present an easy alternative proof.

Proof. Only the implication $1 \Rightarrow 2$ is nontrivial. To prove it assume d -ary $\mathcal{M} \models \varphi, \mathcal{N} \models \neg\varphi$ with $\mathcal{M} \cong_\sigma^n \mathcal{N}$ and assume towards contradiction that $\varphi, \neg\varphi$ are not \cong_σ^n -consistent over \mathbb{T}^d . We have $\mathcal{M} \cong_\sigma^n \mathcal{M}_{|n}^\sigma \cong \mathcal{M}'$ where \mathcal{M}^σ is the σ -reduct of \mathcal{M} , and $\mathcal{M}' \in \mathbb{T}^d$ is the bisimilarity quotient of its n -prefix $\mathcal{M}_{|n}^\sigma$. By the assumption that $\varphi, \neg\varphi$ are not joint \cong_σ^n -consistent, $\mathcal{M} \models \varphi$ implies $\mathcal{M}_{|n}^\sigma \models \varphi$. By invariance of φ under \cong , this in turn implies $\mathcal{M}' \models \varphi$. We construct $\mathcal{N}' \models \neg\varphi$ symmetrically. By definition, $\mathcal{M} \cong_\sigma^n \mathcal{N}$ means that $\mathcal{M}_{|n}^\sigma$ and $\mathcal{N}_{|n}^\sigma$ are bisimilar, which is equivalent to saying that their bisimilarity quotients \mathcal{M}' and \mathcal{N}' are isomorphic, and hence (σ, n) -isomorphic. Thus, $\mathcal{M}', \mathcal{N}'$ witness joint \cong_σ^n -consistency of $\varphi, \neg\varphi$ over \mathbb{T}^d , a contradiction. ◀

Using automata-based techniques we to decide if Item 2 in Lemma 6 holds for all n .

► **Proposition 7.** *For every parity automata $\mathcal{A}, \mathcal{A}'$ and $d \in \mathbb{N}$: $\mathcal{A}, \mathcal{A}'$ are joint \cong_σ^n -consistent over \mathbb{T}^d for all $n \in \mathbb{N}$ iff $\mathcal{A}, \mathcal{A}'$ are joint \cong_σ^m -consistent over \mathbb{T}^d for $m = |\mathcal{A}| + |\mathcal{A}'| + 1$. The latter condition can be checked in time polynomial in $|\mathcal{A}| + |\mathcal{A}'|$.*

Proof. (Sketch) Due to well-known relativization techniques we do not lose generality by only running $\mathcal{A}, \mathcal{A}'$ on full d -ary trees with no leaves. Let L be a language of finite full d -ary trees over σ such that $\mathcal{M} \in L$ iff \mathcal{M} is a prefix of a reduct of a model of \mathcal{A} . Let L' be an analogous language for \mathcal{A}' . The *tallness* of a finite tree is the minimal distance from the root to a leaf. Observe that $\mathcal{A}, \mathcal{A}'$ are \cong_σ^n -consistent over \mathbb{T}^d iff $L \cap L'$ contains a tree of tallness n . Thus, it suffices to check if $L \cap L'$ contains trees of arbitrarily high tallness. To that end construct an automaton \mathcal{B} recognizing $L \cap L'$ of size polynomial in $|\mathcal{A}| + |\mathcal{A}'|$. An easy pumping argument shows that the language $L \cap L'$ of \mathcal{B} contains trees of arbitrarily high tallness iff it contains a tree of tallness $m = |\mathcal{B}| + 1$. To test the latter condition it is enough to inductively compute a sequence $S_1 \supseteq S_2 \supseteq \dots \supseteq S_{|\mathcal{B}|+1}$ of subsets of states of \mathcal{B} , where S_i is the set of all states q such that $\mathcal{B}[q_T \leftarrow q]$ recognizes a tree of tallness at least i . ◀

We are ready to solve ML-definability over \mathbb{T}^d in exponential time. Assume μML -formula φ . For every n , we know by Equivalence (Base) that φ is equivalent over \mathbb{T}^d to some $\psi \in \text{ML}_\sigma^n$ iff $\varphi, \neg\varphi$ are not joint \cong_σ^n -consistent over \mathbb{T}^d . By Lemma 6 this is equivalent to the lack of joint \cong_σ^n -consistency of $\varphi, \neg\varphi$ over \mathbb{T}^d . By Theorem 2 we can compute exponentially-sized automata $\mathcal{A}, \mathcal{A}'$ equivalent to φ and $\neg\varphi$ over \mathbb{T}^d . It follows that φ is not ML_σ -definable over \mathbb{T}^d iff $\mathcal{A}, \mathcal{A}'$ are joint \cong_σ^n -consistent over \mathbb{T}^d for every n . The last condition is decided using Proposition 7. The runtime of our algorithm is polynomial in $|\mathcal{A}| + |\mathcal{A}'|$, and thus exponential in $|\varphi|$. This proves the part of Theorem 5 about \mathbb{T}^d . The remaining part concerning unrestricted models is a special case of Theorem 9, which we will prove next.

Modal Separation: the Unrestricted Case

Over unrestricted models, separability turns out to be only slightly more complicated than definability. Lemma 6 becomes false if $\neg\varphi$ is replaced with arbitrary φ' (which would be the statement relevant for separability). We have the following lemma, however.

► **Lemma 8.** *For every $\varphi, \varphi' \in \mu\text{ML}$ and $n \in \mathbb{N}$ the following are equivalent:*

1. φ, φ' are joint \cong_σ^n -consistent over all models.
2. φ, φ' are joint \cong_σ^n -consistent over \mathbb{T}^d , where $d = |\varphi| + |\varphi'|$.

Proof. The implication (1) \Leftarrow (2) is immediate. To prove the other one (1) \Rightarrow (2) consider an intermediate property:

$$\varphi, \varphi' \text{ are joint } \cong_\sigma^n\text{-consistent over all models.} \quad (1.5)$$

The implication (1) \Rightarrow (1.5) can be read off from Otto's original proof. The remaining one (1.5) \Rightarrow (2) is a special case of a stronger claim which we prove later: the implication (3) \Rightarrow (4) of Lemma 27. ◀

Lemma 8 allows us to solve ML-separability in exponential time.

► **Theorem 9.** *Over all models, ML-separability of μML -formulae is EXPTIME-complete.*

Proof. The proof is almost the same as our proof of Theorem 5. The only difference is that we consider an arbitrary φ' in place of $\neg\varphi$, and hence use Lemma 8 in place of Lemma 6. ◀

Apart from deciding separability we also construct separators when they exist. Given a subset \mathcal{L} of μML formulae, $\varphi \in \mu\text{ML}$, and $\psi \in \mathcal{L}$, we call ψ an \mathcal{L} -uniform consequence of φ if $\psi \models \theta$ for every $\theta \in \mathcal{L}$ such that $\varphi \models \theta$. The notion relativizes to a fixed class \mathbf{C} of models by only considering entailment over that class. Observe that if φ, φ' are \mathcal{L} -separable and ψ is an \mathcal{L} -uniform consequence of φ then ψ is an \mathcal{L} -separator for φ, φ' . The same is true over any class \mathbf{C} .

Note that it follows from the proof of Theorem 9 that if φ, φ' are ML-separable then they admit a separator of modal depth n at most exponential in $|\varphi| + |\varphi'|$. It follows that constructing an ML_σ -separator for φ, φ' boils down to constructing an ML_σ^n -uniform consequence of φ . A naive construction which always works is to take the disjunction of all ML_σ^n -types consistent with φ over \mathbf{C} . Here, by an ML_σ^n -type we mean a maximal consistent subset of ML_σ^n . Since up to equivalence there are only finitely many formulae in ML_σ^n , each ML_σ^n -type can be represented as a single ML_σ^n -formula and the mentioned disjunction ψ is well-defined. This construction is non-elementary in n over all models and doubly exponential in n over models of bounded outdegree.

We present an efficient construction of ML_σ^n -uniform consequences. The construction works over unrestricted models, over \mathbb{T}^1 and over \mathbb{T}^2 but not over \mathbb{T}^d for $d \geq 3$. Since in the following Section 4 we will provide a more efficient construction for \mathbb{T}^1 , now we only look at the unrestricted and binary case. For convenience, we construct ML_σ^n -uniform consequences of automata instead of formulae, with definition adapted in an obvious way.

► **Proposition 10.** *Let \mathbf{C} be the class of all models or \mathbb{T}^2 . Assume an NPTA \mathcal{A} over \mathbf{C} , a signature σ and $n \in \mathbb{N}$. An ML_σ^n -uniform consequence of \mathcal{A} over \mathbf{C} can be constructed in time $|A|^{O(n \cdot |A|)}$ if \mathbf{C} is the class of all models and in time $2^{O(n \cdot |A|)}$ if $\mathbf{C} = \mathbb{T}^2$.*

Proof. Let \mathcal{A} be an NPTA. Let $\mathcal{B} = (Q, \Sigma, q_I, \delta, \text{rank})$ be an automaton of the same size recognizing σ -reducts of models of \mathcal{A} . A formula ψ is an ML_σ^n -uniform consequence of \mathcal{A} over \mathbf{C} iff it is an ML^n -uniform consequence of \mathcal{B} over \mathbf{C} . Thus, it suffices to construct the latter.

We construct $\psi_{n,q}$ for every $q \in Q$ and $n \in \mathbb{N}$ by induction on $n \in \mathbb{N}$. For the base case we put:

$$\psi_{0,q} = \bigvee \{c \in \Sigma \mid \text{there is } \mathcal{N} \in \mathbf{C} \text{ with } \mathcal{N} \models \mathcal{B}[q_I \leftarrow q] \text{ and } \mathcal{N} \models c\}$$

For the induction step define:

$$\psi_{n+1,q} = \bigvee_{c \in \Sigma} \bigvee_{S \in \delta(q,c)} c \wedge \nabla \{\psi_{n,p} \mid p \in S\}$$

where $\nabla \Phi$ is an abbreviation for $\bigwedge_{\theta \in \Phi} \diamond \theta \wedge \square \bigvee_{\theta \in \Phi} \theta$. Assume \mathbf{C} is either the class of all models or \mathbb{T}^2 . The construction preserves the following invariant:

$$\mathcal{M} \models \psi_{n,q} \iff \text{there exists } \mathcal{N} \in \mathbf{C} \text{ with } \mathcal{N} \models \mathcal{B}[q_I \leftarrow q] \text{ and } \mathcal{M} \simeq^n \mathcal{N} \quad (1)$$

for every structure $\mathcal{M} \in \mathbf{C}$. Hence, ψ_{n,q_I} is an ML_σ^n -uniform consequence of \mathcal{A} over \mathbf{C} . It is routine to check that in either case the formula has the right size.

The proof of (1) proceeds by induction, with slightly different arguments in the cases of binary and of unrestricted models. The details of this proof are found in Appendix A.1. It is worth to point out, however, that the implication \Rightarrow from left to right would not be valid over \mathbb{T}^d with $d \geq 3$. \blacktriangleleft

Given the exponential construction of automata from Theorem 2 and the exponential upper bound on modal depth n of separators, Proposition 10 yields an efficient construction of separators.

► **Theorem 11.** *If φ, φ' are ML_σ -separable, then one can compute an ML_σ -separator in time doubly exponential in $|\varphi| + |\varphi'|$.*

It is not difficult to show that, in the presence of at least two accessibility relations \diamond_1, \diamond_2 , the construction is optimal: one can express in μML that the model embeds a full binary tree of depth 2^n and in which each inner node has both a \diamond_1 - and a \diamond_2 -successor. Using standard techniques, one can show that any modal formula expressing this property is of doubly exponential size [13]. Whether having two accessibility relations is necessary for this lower bound is an interesting question which we leave open.

It is interesting to note that the separators we compute are *not* the logically strongest separators and, in fact, strongest separators do not even have to exist.

► **Example 12.** Consider $\varphi = \theta_\infty$ from before and $\varphi' = \square \perp$. For every $n \in \mathbb{N}$, the modal formula $\diamond^n \top$ separates φ from φ' , and $\diamond^m \top \models \diamond^n \top$ whenever $m \geq n$.

The remaining open cases are the problems of ML -separability (and separator construction) over \mathbb{T}^d for $d \geq 1$. We investigate the cases of unary ($d = 1$), binary ($d = 2$), and higher maximal outdegree ($d \geq 3$) in turn. We emphasize that the outdegree d is not a part of the input but rather a property of the considered class of models.

4 Unary Case

We first investigate ML -separability over \mathbb{T}^1 , that is, models that are essentially *words*. Note that satisfiability of μML over words is PSPACE-complete (an upper bound follows, e.g., via the translation to automata and the lower bound is inherited from LTL [30, Theorem 4.1]) which suggests that also definability and separability could be easier. Indeed, we show:

► **Theorem 13.** *ML-definability and ML-separability of μML -formulae is PSPACE-complete over \mathbb{T}^1 .*

Proof. The lower bound is by a reduction from satisfiability, and applies to definability.

Given formulae $\varphi, \varphi' \in \mu\text{ML}$ and a subset of the signature σ , consider the set of finite words $L = \{W \in \mathcal{P}(\sigma)^* \mid W \text{ is a } \sigma\text{-reduct of a prefix } V \text{ of some model } U \text{ of } \varphi\}$. Let L' be a similar language defined for φ' . Two unary models are bisimilar iff they are identical. Hence, by Equivalence (Base) the formulae $\varphi, \varphi' \in \mu\text{ML}$ are not ML_σ -separable over \mathbb{T}^1 iff $L \cap L'$ is infinite. It is standard to define a finite automaton \mathcal{A} recognizing $L \cap L'$ and check if its language is infinite (which is equivalent to checking if $L \cap L'$ contains input longer than $|\mathcal{A}|$). To do it in polynomial space, we nondeterministically guess the long input, letter by letter, and only remember the current state and a binary counter measuring the length of the input guessed so far. ◀

We conclude this section with proving that ML_σ -separators can be constructed in exponential time and are thus of at most exponential size. Note that this is optimal, since over \mathbb{T}^1 , μML is exponentially more succinct than ML . Indeed, it is standard to implement an exponential counter using a polynomially sized μML -formula.

► **Theorem 14.** *If $\varphi, \varphi' \in \mu\text{ML}$ are ML_σ -separable over \mathbb{T}^1 , then one can compute an ML_σ -separator in time exponential in $|\varphi| + |\varphi'|$.*

As argued in the previous section, it suffices to construct an ML^n -uniform consequence of the NPTA equivalent to φ , which we do next.

► **Proposition 15.** *Let \mathcal{A} be an NPTA over \mathbb{T}^1 with ℓ states, $n \in \mathbb{N}$, and σ a signature. An ML_σ^n -uniform consequence of \mathcal{A} over \mathbb{T}^1 can be constructed in time polynomial in n , σ , and ℓ .*

Proof. As argued in the previous section, it suffices to construct an ML^n -uniform consequence of the NPTA \mathcal{B} which recognizes precisely the σ -reducts of models of \mathcal{A} . Let \mathcal{B} have states Q . By construction of \mathcal{B} , we have $|Q| = \ell$. As an auxiliary step, we define for every $p, q \in Q$ and $m \leq n$ a formula $\psi_{p,q}^m \in \text{ML}_\sigma^n$ such that for every $\mathcal{M} \in \mathbb{T}^1$:

$$\mathcal{M} \models \psi_{p,q}^m \iff \text{there is a run of } \mathcal{B} \text{ from } p \text{ to } q \text{ over the } m\text{-prefix of } \mathcal{M}. \quad (2)$$

The ψ_{pq}^m are defined inductively with the base cases ($m \leq 1$) read off from \mathcal{B} , and using *divide and conquer* in the inductive step ($m > 1$), to keep the formulae small. More formally, we define ψ_{pq}^m for $m > 1$ and all $p, q \in Q$ by taking:

$$\psi_{pq}^m = \bigvee_{q' \in Q} (\psi_{pq'}^{\lfloor m/2 \rfloor} \wedge \diamond^{\lfloor m/2 \rfloor} \psi_{q'q}^{\lceil m/2 \rceil})$$

It is routine to verify that ψ_{pq}^m satisfies (2) and is of size $|\psi_{pq}^m| \in O(|Q| \cdot m^2)$. Based on the ψ_{pq}^m , one can define a formula ψ_n that describes all possible prefixes of length $\leq n$ of models of \mathcal{B} , and thus is the sought ML_σ -uniform consequence of \mathcal{B} . One can think of ψ_n as the disjunction of formulae $\psi_{q_0 q}^n$ for q_0 the initial state of \mathcal{B} , but the full construction is slightly more involved since models accepted by \mathcal{B} might be also shorter than n . ◀

5 Binary Case

We next handle the binary case \mathbb{T}^2 . The key observation here is that, between full binary trees, bisimilarity entails isomorphism.

► **Proposition 16.** *Assume full binary trees $\mathcal{M}, \mathcal{M}' \in \mathbb{T}^2$. If \mathcal{M} and \mathcal{M}' are σ -bisimilar then they are σ -isomorphic.*

Proof. By definition a σ -bisimulation between two models is a bisimulation between their reducts to σ , and σ -isomorphism is such a bisimulation which is additionally bijective. It therefore suffices to show that if $\mathcal{M}, \mathcal{M}'$ are full binary trees and Z is a bisimulation between them then there is a bijective bisimulation $Z' \subseteq Z$. We pick such Z' inductively starting with the pair of roots (v_I, v'_I) . The key observation is that if v has children v_1, v_2 and w has children w_1, w_2 and vZw then either (i) v_1Zw_1 and v_2Zw_2 or (ii) v_1Zw_2 and v_2Zw_1 (the cases are not exclusive). The details can be found in Appendix C.1. ◀

Proposition 16 can be used to prove the Craig interpolation property of ML over \mathbb{T}^2 and implies the following separability-variant of Lemma 6 over \mathbb{T}^2 .

► **Lemma 17.** *For every $\varphi, \varphi' \in \mu\text{ML}$ and $n \in \mathbb{N}$ the following are equivalent:*

1. φ, φ' are joint $\stackrel{n}{\equiv}_\sigma$ -consistent over \mathbb{T}^2 .
2. φ, φ' are joint \cong_σ^n -consistent over \mathbb{T}^2 .

Proof. We show only the nontrivial implication $1 \Rightarrow 2$. Assume binary $\mathcal{M} \models \varphi, \mathcal{M}' \models \varphi'$ with $\mathcal{M} \stackrel{n}{\equiv}_\sigma \mathcal{M}'$. Let $\mathcal{N} \models \varphi$ and $\mathcal{N}' \models \varphi'$ be full binary trees obtained from \mathcal{M} and \mathcal{M}' by duplicating subtrees. By Proposition 16, $\mathcal{N} \cong_\sigma^n \mathcal{N}'$ which proves 2. ◀

Similarly to the definability case, Lemma 17 combined with Equivalence (Base) and Proposition 7 immediately give an exponential procedure for separability. Since the lower bound is inherited from definability, we get the following result.

► **Theorem 18.** *ML-separability and ML-definability of μML -formulae is EXPTIME-complete over \mathbb{T}^2 .*

With the same argument as for Theorem 11 we use Proposition 10 to conclude:

► **Theorem 19.** *If φ, φ' are ML_σ -separable over \mathbb{T}^2 , then one can compute an ML_σ -separator in time doubly exponential in $|\varphi| + |\varphi'|$.*


6 Ternary and Beyond

In this section we address the case of models with outdegree bounded by a number $d \geq 3$. We illustrate that this case behaves differently as it lacks the Craig interpolation property.

► **Example 20.** Consider ML-formulae $\varphi = \diamond(a \wedge b) \wedge \diamond(a \wedge \neg b)$ and $\varphi' = \diamond(\neg a \wedge c) \wedge \diamond(\neg a \wedge \neg c)$. Clearly, $\varphi \models \neg \varphi'$ over \mathbb{T}^3 . Observe that models $\mathcal{M}, \mathcal{M}'$ in Figure 1 witness that φ, φ' are joint $\stackrel{\{a\}}{\equiv}$ -consistent and thus joint $\stackrel{n}{\equiv}_{\{a\}}$ -consistent for every $n \in \mathbb{N}$. By Equivalence (Base) there is no $\text{ML}_{\{a\}}$ -separator, which is nothing else than a Craig interpolant. ◀

Motivated by the lack of the Craig interpolation property, we study the ML-*interpolant existence* problem: given $\varphi, \varphi' \in \text{ML}$ and signature σ , decide whether there is an ML_σ -separator of $\varphi, \neg \varphi'$, that is, $\psi \in \text{ML}_\sigma$ with $\varphi \models \psi \models \varphi'$. *Craig ML-*interpolant existence** is the special case in which $\sigma = \text{sig}(\varphi) \cap \text{sig}(\varphi')$. Observe that ML-*interpolant existence* is the special case of ML-separability of μML -formulae in which the input to the separability is restricted to ML-formulae. We show that already ML-*interpolant existence* over \mathbb{T}^3 is harder than ML-separability of μML -formulae over arbitrary models.

► **Theorem 21.** *For $d \geq 3$, ML-*interpolant existence* over \mathbb{T}^d is CONEXPTIME-complete. Hardness already applies to Craig ML-*interpolant existence* over \mathbb{T}^d .*



■ **Figure 1** Witness of joint consistency: dashed lines and colors indicate the $\{a\}$ -bisimulation.

Proof. The upper bound is easy to establish based on the observation that $\varphi, \neg\varphi'$ of modal depth at most m do not admit an ML_σ -separator over \mathbb{T}^d iff they are joint \equiv_σ^m -consistent over \mathbb{T}^d . The witness $\mathcal{M}, \mathcal{M}'$ of joint \equiv_σ^m -consistency of $\varphi, \neg\varphi'$ can be assumed to be of depth m . Such models are of exponential size (they have at most d^m points) and can thus be guessed by a non-deterministic exponential time bounded Turing machine.

The lower bound is more intriguing and relies on an extension of Example 20. Reconsidering the example it is important to note that in *every* witness $\mathcal{M}, \mathcal{M}'$ of joint $\equiv_{\{a\}}^n$ -consistency of φ, φ' , there are two successors of v_I that are bisimilar to the same successor of v'_I . We extend the idea and enforce exponentially many bisimilar points. More precisely, consider families $(\psi_i)_{i \in \mathbb{N}}, (\psi'_i)_{i \in \mathbb{N}}$ of modal formulae inductively defined as follows:

$$\begin{aligned} \psi_0 &= \psi'_0 = \top \\ \psi_{i+1} &= \diamond(a \wedge b_i) \wedge \diamond(a \wedge \neg b_i) \wedge \square(a \rightarrow (\psi_i \wedge (b_i \rightarrow \bigwedge_{j < i} \square^j b_i) \wedge (\neg b_i \rightarrow \bigwedge_{j < i} \square^j \neg b_i))) \\ \psi'_{i+1} &= \diamond(\neg a \wedge c) \wedge \diamond(\neg a \wedge \neg c) \wedge \diamond(a \wedge \psi'_i) \end{aligned}$$

Clearly, the size of ψ_i, ψ'_i is polynomial in i . Moreover, by induction on i , it is readily verified that for every $i \in \mathbb{N}$, for every $\mathcal{M}, \mathcal{M}' \in \mathbb{T}^3$ with $\mathcal{M} \models \psi_i, \mathcal{M}' \models \psi'_i$, and every $(\{a\}, i)$ -bisimulation S witnessing $\mathcal{M} \equiv_{\{a\}}^i \mathcal{M}'$, there are points w_0, \dots, w_{2^i-1} in depth i in \mathcal{M} and a point \hat{w} in depth i in \mathcal{M}' such that $(w_j, \hat{w}) \in S$ for all j and such that distinct w_j, w_k can be distinguished by some proposition in b_0, \dots, b_{i-1} . Intuitively, this means that ψ_i, ψ'_i enforce in joint $\equiv_{\{a\}}^i$ -consistent models $\mathcal{M}, \mathcal{M}'$ that \mathcal{M} contains 2^i points w_0, \dots, w_{2^i-1} which are all linked to the same point \hat{w} in \mathcal{M}' . We exploit this link to synchronize information between the w_j , following a strategy that has recently been used to show CONEXPTIME -hardness for interpolant existence in some description logics [1].

We reduce a NEXPTIME -complete tiling problem [14]: Given a set Δ of tile types and horizontal and vertical compatibility relations $H, V \subseteq \Delta \times \Delta$, and some $n \in \mathbb{N}$ in unary, decide whether one can tile the $2^n \times 2^n$ torus with tiles from Δ complying with H, V . Given Δ, H, V, n , we define formulae $\varphi_n = \psi_{2n} \wedge \square^{2n} \chi_n, \varphi'_n = \psi'_{2n} \wedge \square^{2n} \chi'_n$ of modal depth m and with common signature $\sigma = \text{sig}(\varphi_n) \cap \text{sig}(\varphi'_n)$ such that

$$\Delta, H, V, n \text{ has a solution} \quad \Leftrightarrow \quad \varphi_n, \varphi'_n \text{ are joint } \equiv_\sigma^m\text{-consistent.}$$

To explain the idea, let $\mathcal{M}, \mathcal{M}'$ witness joint \equiv_σ^m -consistency of φ_n, φ'_n . The gadget formulae ψ_{2n}, ψ'_{2n} enforce 2^{2n} points $w_0, \dots, w_{2^{2n}-1}$ in depth $2n$ in \mathcal{M} which are all linked via the bisimulation to a single point \hat{w} in \mathcal{M}' . These 2^{2n} points shall represent the $2^n \times 2^n$ cells of the torus. The intended solution of the tiling problem is represented via propositions $p_d \in \sigma$, for each $d \in \Delta$. To synchronize them we proceed as follows. Using the $2n$ propositions b_0, \dots, b_{2n-1} (which are not in σ), we can associate coordinates $(x_i, y_i) \in \{0, \dots, 2^n - 1\} \times \{0, \dots, 2^n - 1\}$ to each point w_i in the torus. To understand the purpose of χ_n, χ'_n , suppose for a moment that the outdegree of the points \hat{w} and the w_i is at most 2^{2n}

(instead of 3). Then we could proceed by enforcing (via χ_n) below each w_i with coordinates (x_i, y_i) three successors v_i^1, v_i^2, v_i^3 such that

- v_i^1, v_i^2, v_i^3 have coordinates $(x_i, y_i), (x_i, y_i + 1),$ and $(x_i + 1, y_i),$ respectively;
- the coordinates of the v_i^j are made visible using propositions in σ ;
- v_i^1, v_i^2, v_i^3 satisfy $p_{d_1}, p_{d_2}, p_{d_3}$ for $d_1, d_2, d_3 \in \Delta$ such that $(d_1, d_2) \in V$ and $(d_1, d_3) \in H$.

These three successors stipulate bisimilar successors of \hat{w} . Since each point in the torus is stipulated three times as successor of some w_i and since the outdegree of \hat{w} is restricted to 2^{2^n} , the three copies of the same point satisfy the same proposition p_d . By the last item above, the selected propositions comply with V, H and thus represent a solution to the tiling problem. Now, since the outdegree below \hat{w} is at most 3 (and not 2^{2^n} as assumed), the χ_n, χ'_n have to be a bit more complicated, but the idea remains the same. ◀

We show next that the situation for the full separability problem is even worse.

► **Theorem 22.** *For every $d \geq 3$, ML-separability of μ ML-formulae over \mathbb{T}^d is 2-EXPTIME-complete.*

Thus, over \mathbb{T}^d for $d \geq 3$, ML-separability is provably harder than ML-definability, c.f. Theorem 5. Both the upper and the lower bound of Theorem 22 are non-trivial; we provide proof sketches in the following two subsections. Before doing that let us conclude this part with separator construction.

► **Theorem 23.** *If φ, φ' are ML_σ -separable over \mathbb{T}^d , $d \geq 3$, then one can compute an ML_σ -separator in time triply exponential in $|\varphi| + |\varphi'|$.*


Proof. (Sketch) It follows from the upper bound proof of Theorem 22 that, if φ, φ' admit an ML_σ -separator, then they admit one of modal depth bounded doubly exponentially in $|\varphi| + |\varphi'|$. Observe that over the signature of φ and φ' there are only triple exponentially many trees of fixed outdegree d and double exponential depth, and that each such tree is characterized by a modal formula of triply exponential size. The sought separator is then the disjunction of all such formulae consistent with φ . ◀

6.1 Lower Bound for Theorem 22

We reduce the word problem of exponentially space bounded alternating Turing machines (ATMs), which is known to be 2-EXPTIME-complete [8]. Informally, the states of such ATMs are partitioned into *universal states* Q_\forall and *existential states* Q_\exists . Configurations of ATMs are defined as usual, but computations are not sequences of configurations but trees of configurations such that an existential configuration has exactly one successor labeled with a universal configuration and a universal configuration has exactly two successors labeled with existential configurations. A computation tree for an input w is a tree whose root is labeled with the initial configuration and such that successor nodes contain successor configurations. w is accepted if there is a computation tree in which each path is infinite (this acceptance condition is slightly non-standard, but eases the proof).

The reduction relies on the same gadget formulae $(\psi_i)_{i \in \mathbb{N}}, (\psi'_i)_{i \in \mathbb{N}}$ as used in the proof of Theorem 21 and additionally uses ideas for showing 2-EXPTIME-hardness for recently studied interpolant existence problems for description logics [1]. For a given ATM \mathfrak{A} and input w of length n , we construct formulae $\varphi_n = \psi_n \wedge \Box^n \chi, \varphi'_n = \psi'_n \wedge \chi'$ such that

$$\varphi_n, \varphi'_n \text{ are joint } \Leftrightarrow_\sigma^m \text{-consistent for every } m \in \mathbb{N} \quad \text{iff } \mathfrak{A} \text{ accepts } w.$$



■ **Figure 2** Computation tree of \mathfrak{A} below some w_i (drawn horizontally for space constraints).

This suffices by Equivalence (Base). The signature σ will consist of a, z , and propositions c_α for every possible cell content α of \mathfrak{A} , that is, $\alpha \in \Gamma \cup (Q \times \Gamma)$. Additionally, φ_n and φ'_n will use auxiliary propositions, e.g., to encode counters. The only purpose of χ' is to mention the propositions in σ ; the main work is done by ψ_n, ψ'_n, χ .

To explain the idea, let us consider witnesses $\mathcal{M}, \mathcal{M}'$ for joint \Leftrightarrow_σ^m -consistency of φ_n, φ'_n for sufficiently large m . By the properties of ψ_n, ψ'_n , we find 2^n points w_0, \dots, w_{2^n-1} in depth n in \mathcal{M} which are bisimilar to a single point \hat{w} in depth n in \mathcal{M}' . Recall that in every w_i , we have access to its index i via a counter using propositions b_0, \dots, b_n . Now, χ is a μML -formula with the following properties, see also Figure 2 for illustration.

- χ enforces the “skeleton” of a computation tree for \mathfrak{A} , in which each configuration is modeled by a path of length 2^n (using an exponential counter), and in which universal and existential configurations alternate.
- χ also enforces that each point of the skeleton is labeled with some cell content via σ -propositions c_α , but without any synchronization except the initial configuration.
- χ makes sure that below w_i the positions $2^n - i$ of successor configurations are coordinated. The key point is that this enforces (due to bisimilarity) a computation tree below \hat{w} in which, due to the last item above, *all* positions of configurations are coordinated.

We remark that the hardness also holds when σ is not part of the input: one can reduce separability of φ, φ' by ML_σ -formulae to separability of φ, φ' by (arbitrary) ML -formulae.

6.2 Upper Bound for Theorem 22

We show that over models of outdegree at most d , ML -separability of fixpoint formulae can be solved in doubly exponential time. Let us start with establishing a technical but useful fact. For every language of d -ary trees $L \subseteq \mathbb{T}^d$ denote the language:

$$\text{bisQuot}(L) = \{\mathcal{M} \in \mathbb{T}^d \mid \text{there is } \mathcal{N} \in L \text{ and a functional bisimulation } Z : \mathcal{N} \xrightarrow{\text{bis}} \mathcal{M}\}$$

of bisimulation quotients of trees from L .

► **Proposition 24.** *For every NPTA \mathcal{A} , an NPTA \mathcal{B} recognizing $\text{bisQuot}(\mathcal{L}(\mathcal{A}))$ can be computed in time exponential in the size of \mathcal{A} .*

Proof. Fix an NPTA $\mathcal{A} = (Q, \Sigma, q_I, \delta, \text{rank})$. For every $\mathcal{M} \in \mathbb{T}^d$, we characterize existence of d -ary $\mathcal{N} \models \mathcal{A}$ with $\mathcal{N} \xrightarrow{\text{bis}} \mathcal{M}$ with the following parity game $\mathcal{G}_{\text{bisQuot}}(\mathcal{M}, \mathcal{A})$. The game has the set $M \times Q$ as positions. The pair (v_I, q_I) consisting of the root v_I of \mathcal{M} and q_I is the initial position. From a position (v, q) first \exists ve chooses $S \in \delta(q, \text{val}(c))$ and a surjective map $h : S \rightarrow \{v_1, \dots, v_k\}$ where $\{v_1, \dots, v_k\}$ is the set of children of v . Then \forall dam responds with a choice of $p \in S$ and the next round starts in position $(h(p), p)$. The game is a parity

game: the ranks are inherited from \mathcal{A} in the sense that the rank of (v, q) equals $\text{rank}(q)$. In Appendix D.2 we prove:

$$\exists \text{ve wins } \mathcal{G}_{\text{bisQuot}}(\mathcal{M}, \mathcal{A}) \iff \mathcal{M} \in \text{bisQuot}(\mathcal{L}(\mathcal{A})) \quad (3)$$

for every $\mathcal{M} \in \mathbb{T}^d$. Using (3) we prove Proposition 24. It suffices to construct an automaton \mathcal{B} which accepts \mathcal{M} iff $\exists \text{ve wins } \mathcal{G}_{\text{bisQuot}}(\mathcal{M}, \mathcal{A})$. To that end, using standard techniques we encode $\exists \text{ve}$'s positional strategies for $\mathcal{G}_{\text{bisQuot}}(\mathcal{M}, \mathcal{A})$ as colorings of \mathcal{M} with $\mathcal{P}(Q \times Q)$ and construct, in time exponential in $|Q|$, an automaton \mathcal{B}^+ recognizing models labelled with such winning positional strategies. We then obtain \mathcal{B} recognizing $\text{bisQuot}(\mathcal{L}(\mathcal{A}))$ by projecting out the additional colors $\mathcal{P}(Q \times Q)$ from \mathcal{B}^+ . \blacktriangleleft

With the help of Proposition 24 we prove Theorem 22. Fix d , μML -formulae φ and φ' and signature σ . By Equivalence (Base), it suffices to check if φ and φ' are jointly \approx_σ^n -consistent over \mathbb{T}^d for every n . However, unlike with definability or in the binary case, we cannot conclude joint \cong_σ^n -consistency from joint \approx_σ^n -consistency. Instead, we use Proposition 24 to directly decide joint \approx_σ^n -consistency for all n . For a language $L \subseteq \mathbb{T}^d$, define the language:

$$\text{QPL}(L) = \{\mathcal{N} \in \mathbb{T}^d \mid \text{there is } \mathcal{M} \in L, \text{ finite prefix } \mathcal{M}_0 \text{ of } \mathcal{M} \text{ and } Z : \mathcal{M}_0 \xrightarrow{\text{bis}} \mathcal{N}\}$$

of finite d -ary trees which are bisimulation quotients of finite prefixes of models from L . By Proposition 24 and the closure properties of parity automata, for every \mathcal{A} one can construct in exponential time an automaton \mathcal{B} recognizing $\text{QPL}(\mathcal{L}(\mathcal{A}))$.

We prove the upper bound from Theorem 22. Using Theorem 2 compute automata $\mathcal{A}, \mathcal{A}'$ accepting σ -reducts of models of φ, φ' . Compute $\mathcal{B}, \mathcal{B}'$ recognizing $\text{QPL}(\mathcal{L}(\mathcal{A}))$ and $\text{QPL}(\mathcal{L}(\mathcal{A}'))$. Recall that any two trees are bisimilar iff they have isomorphic bisimulation quotients. It follows that φ, φ' admit a ML_σ^n -separator over \mathbb{T}^d iff $\mathcal{A}, \mathcal{A}'$ are joint \approx_σ^n -consistent iff $\mathcal{B}, \mathcal{B}'$ are joint \cong^n consistent. By Proposition 7, the latter condition holds for all $n \in \mathbb{N}$ iff it holds for $n = |\mathcal{B}| + |\mathcal{B}'| + 1$ and this can be tested in time polynomial in $|\mathcal{B}| + |\mathcal{B}'|$. Since $\mathcal{A}, \mathcal{A}'$ are exponential, and $\mathcal{B}, \mathcal{B}'$ are doubly exponential in the size of φ, φ' , this gives the upper bound from Theorem 22.

7 Case Study: Graded Modalities

In this section we apply our techniques and results to the case with *graded modal operators*. Formally, we extend μML with formulae of the shape $\diamond_{\sim g} \psi$ and $\square_{\sim g} \psi$, where $\sim \in \{\leq, \geq\}$ and the *grade* $g \in \mathbb{N}$ is a natural number. Intuitively, $\diamond_{\geq g} \psi$ is true in a point w if w has at least g successors satisfying ψ and dually, $\square_{\leq g} \psi$ is true in w if all but at most g successors satisfy ψ [11, 25]. We denote with grML and $\text{gr}\mu\text{ML}$ the extension of ML and μML , respectively, with such graded modalities. Clearly, for any $d \in \mathbb{N}$, \mathbb{T}^d is $\text{gr}\mu\text{ML}$ -definable by the formula $\theta_d = \nu x. (\diamond_{\leq d} \top \wedge \square x)$, which is an additional motivation to study grML and $\text{gr}\mu\text{ML}$.

Indeed, using the results and techniques from the previous section one can easily prove that ML -separability of $\text{gr}\mu\text{ML}$ -formulae (defined as expected) is 2-EXPTIME-complete.

► **Theorem 25.** *ML-separability of $\text{gr}\mu\text{ML}$ -formulae is 2-EXPTIME-complete.*

Proof. For the lower bound, we reduce ML -separability of μML -formulae over \mathbb{T}^3 in spirit similar to Lemma 4. Since the former problem is 2-EXPTIME-hard by Theorem 22, the latter is as well. Recall the formula θ_3 defining \mathbb{T}^3 . Then, for any μML -formulae φ, φ' and

$\psi \in \text{ML}$, we have that ψ is an ML_σ -separator of φ, φ' over \mathbb{T}^3 iff ψ is an ML_σ -separator of $\varphi \wedge \theta_d, \varphi' \wedge \theta_d$.

Towards the upper bound, suppose $\varphi, \varphi' \in \text{gr}\mu\text{ML}$. Using standard arguments, one can show that φ, φ' are ML -separable over all models iff they are ML -separable over \mathbb{T}^d , where $d = g \times (|\varphi| + |\varphi'|)$ and g is the greatest grade occurring in φ, φ' . We then construct NPTA $\mathcal{A}, \mathcal{A}'$ equivalent to φ, φ' over d -ary trees via (an analogue for $\text{gr}\mu\text{ML}$ of) Theorem 2 and proceed with $\mathcal{A}, \mathcal{A}'$ as described in the upper bound proof of Theorem 22. \blacktriangleleft

Interestingly, the problem becomes easier if we allow grades in the separating formula.

► **Theorem 26.** *grML-separability of gr μ ML-formulae is EXPTIME-complete.*

The lower bound follows by the usual reduction from satisfiability. We thus focus on the upper bound. Similarly to the non-graded case, we establish first a model-theoretic characterization, based on the appropriate notion of bisimilarity that characterizes the expressive power of grML [10]. A relation Z between models is a *graded bisimulation* if it satisfies (atom) and *graded* variants of the (back) and (forth) conditions of bisimulations. The graded (forth) condition says that if vZw then for every $k \in \mathbb{N}$ and pairwise different children v_1, \dots, v_k of v , there are pairwise different children w_1, \dots, w_k of w satisfying v_iZw_i for all $i \leq k$. The graded (back) condition is symmetric. It is a *g-graded bisimulation* if the graded (forth) and (back) conditions need to be satisfied only for $k \leq g$. We denote with $\mathcal{M} \stackrel{\text{grd}}{\simeq} \mathcal{M}'$ (resp., $\mathcal{M} \stackrel{g}{\simeq} \mathcal{M}'$) the fact that there is a graded bisimulation (resp., a g -graded bisimulation) between \mathcal{M} and \mathcal{M}' that relates their roots. Variants with bounded depth n and/or given signature σ are defined and denoted as expected.

► **Lemma 27.** *For every $\varphi, \varphi' \in \text{gr}\mu\text{ML}$ with maximal grade g_{\max} , signature σ , and $n \in \mathbb{N}$, the following are equivalent:*

1. φ, φ' are not grML_σ^n -separable (over all models).
2. φ, φ' are joint $\stackrel{n}{\simeq}_{\text{grd}, \sigma}^n$ -consistent (over all models).
3. φ, φ' are joint $\stackrel{n}{\cong}_\sigma^n$ -consistent (over all models).
4. φ, φ' are joint $\stackrel{n}{\cong}_\sigma^n$ -consistent over \mathbb{T}^d for $d = g_{\max} \times (|\varphi| + |\varphi'|)$.

Using Lemma 27, one can solve grML -separability of $\text{gr}\mu\text{ML}$ formulae in exponential time, following the approach described in Section 3. More precisely, given φ, φ' , we construct NPTA $\mathcal{A}, \mathcal{A}'$ equivalent to φ, φ' over d -ary trees, d as in Lemma 27, and decide whether $\mathcal{A}, \mathcal{A}'$ are joint $\stackrel{n}{\cong}_\sigma^n$ -consistent over \mathbb{T}^d for all n via Proposition 7.

Let us provide some details on the proof of the central Lemma 27.

Proof. We show the implications $1 \Rightarrow 2 \Rightarrow 3 \Rightarrow 4 \Rightarrow 1$ in turn. The implication $4 \Rightarrow 1$ is immediate.

For $1 \Rightarrow 2$, suppose φ, φ' are not grML_σ^n -separable. Hence, for every $g \in \mathbb{N}$ there is a pair of models $\mathcal{M}_g \models \varphi$ and $\mathcal{M}'_g \models \varphi'$ with $\mathcal{M}_g \stackrel{n}{\simeq}_{g, \sigma}^n \mathcal{M}'_g$. One can encode with an FO-sentence θ that two models \mathcal{M} and \mathcal{M}' are depth- n trees, \mathcal{M} is a prefix of some $\mathcal{M}_+ \models \varphi$ and \mathcal{M}' of some $\mathcal{M}'_+ \models \varphi'$. If Z is a fresh binary symbol, then it is also possible to encode with an (infinite) set T of FO-sentences that Z is a graded bisimulation between \mathcal{M} and \mathcal{M}' . Every finite fragment of $\{\theta\} \cup T$ only mentions finitely many grades and hence by assumption is satisfiable. Thus, by compactness of FO, the entire $\{\theta\} \cup T$ is satisfiable. This gives us $\mathcal{M} \stackrel{n}{\simeq}_{\text{grd}, \sigma}^n \mathcal{M}'$ with extensions $\mathcal{M}_+ \models \varphi$ and $\mathcal{M}'_+ \models \varphi'$.

For $2 \Rightarrow 3$, fix witnesses $\mathcal{M}, \mathcal{M}'$ of joint $\stackrel{n}{\simeq}_{\text{grd}, \sigma}^n$ -consistency, that is, $\mathcal{M} \stackrel{n}{\simeq}_{\text{grd}, \sigma}^n \mathcal{M}'$ and there are extensions $\mathcal{M}_+, \mathcal{M}'_+$ of $\mathcal{M}, \mathcal{M}'$ with $\mathcal{M}_+ \models \varphi$ and $\mathcal{M}'_+ \models \varphi'$. By the Löwenheim-Skolem property of FO we may assume that both models are at most countable. It remains to apply the known fact that countable trees \mathcal{N} and \mathcal{N}' satisfy $\mathcal{N} \stackrel{n}{\simeq}_{\text{grd}} \mathcal{N}'$ iff \mathcal{N} and \mathcal{N}' are

isomorphic. For the sake of completeness, we add a brief justification of this latter statement. Assume $w \in \mathcal{N}$ and $w' \in \mathcal{N}'$ with respective children $w_1, w_2, \dots = \bar{w}$ and $w'_1, w'_2, \dots = \bar{w}'$ such that $w \leftrightarrow_{\text{grd}} w'$. For every $\leftrightarrow_{\text{grd}}$ -equivalence class X of \bar{w} the corresponding equivalence class $\{w'_i \mid \exists_{j \leq k}. w_j \leftrightarrow_{\text{grd}} w'_i\} = X'$ has the same cardinality as X . This is immediate for finite X , and for infinite X it follows because in countable models every two infinite subsets have the same cardinality. This allows us to inductively pick a *bijective* subrelation Z of $\leftrightarrow_{\text{grd}}$ between \mathcal{N} and \mathcal{N}' which is still a graded bisimulation.

For $3 \Rightarrow 4$, fix witnesses $\mathcal{M}, \mathcal{M}'$ of joint \cong_{σ}^n -consistency, that is, $\mathcal{M} \cong_{\sigma}^n \mathcal{M}'$ and there are extensions $\mathcal{M}_+, \mathcal{M}'_+$ of $\mathcal{M}, \mathcal{M}'$ with $\mathcal{M}_+ \models \varphi$ and $\mathcal{M}'_+ \models \varphi'$. We trim \mathcal{M}_+ and \mathcal{M}'_+ so that the outdegree becomes at most d . Without losing generality we assume that the prefixes of \mathcal{M}_+ and \mathcal{M}'_+ are not only isomorphic but identical. The semantics of every $\psi \in \mu\text{ML}$ in a model \mathcal{N} is captured by a parity game whose positions are $N \times \text{SubFor}(\psi)$. We extend the definition of the game to $\mu\text{ML}_{\text{grd}}$. The set of positions $N \times \text{SubFor}(\psi)$ and the winning condition are defined as in the classical case, and so are the moves for all the positions with topmost connective other than the graded modalities. In the classical game, from $(v, \diamond\theta)$ \exists ve chooses a child v' of v and the next position is (v', θ) . In $(v, \diamond_{\geq k}\theta)$, first \exists ve chooses a subset v_1, \dots, v_k of size k of children of v , then \forall dam chooses one of these children v_i and the next round starts at (v_i, θ) . Dually, in $(v, \square_{\leq k}\theta)$ first \exists ve picks a subset v_1, \dots, v_k of at most k v 's children, then \forall dam responds with a choice of some v' not in v_1, \dots, v_k and the next position is (v', θ) . It is tedious but straightforward to check that \exists ve wins the game from v, ψ iff ψ is true at v , as in the classical case. Note that if we take a submodel \mathcal{N}_0 of \mathcal{N} which contains at least the root and all \diamond -witnesses (that is, points chosen by a winning strategy ζ in for positions of shape $(v, \diamond_{\geq k}\theta)$) then (the restriction of) ζ to \mathcal{N}_0 is a winning strategy for $\mathcal{G}(\mathcal{N}_0, \psi)$.

Let ζ and ζ' be positional winning strategies for \exists ve in the semantic games $\mathcal{G}(\mathcal{M}_+, \varphi)$ and $\mathcal{G}(\mathcal{M}'_+, \varphi')$. We take submodel $\mathcal{M}_0 \models \varphi$ of \mathcal{M}_+ as follows. In the n -prefix we take the root and all \diamond -witnesses for both ζ and ζ' . In the rest of the model we only take \diamond -witnesses for ζ . A submodel \mathcal{M}'_0 of \mathcal{M}'_+ is defined symmetrically. It follows that $\mathcal{M}_0 \models \varphi$ and $\mathcal{M}'_0 \models \varphi'$.

Recall that g is the maximal grade appearing in φ and φ' . Since the respective sets of positions of $\mathcal{G}(\mathcal{M}_+, \varphi)$ and $\mathcal{G}(\mathcal{M}'_+, \varphi')$ are $M_+ \times \text{SubFor}(\varphi)$ and $M'_+ \times \text{SubFor}(\varphi')$, for every point v there are at most $g \times |\varphi|$ \diamond -witnesses chosen by ζ from a position which has v on the first coordinate. Consequently, the outdegree of \mathcal{M}_0 and \mathcal{M}'_0 is not greater than $d = g \times (|\varphi| + |\varphi'|)$. This proves Lemma 27. \blacktriangleleft

8 Conclusion

We have presented an in-depth study of modal separation of μML -formulae over different classes of structures. For us, the most interesting results are the differences that are obtained over classes of bounded outdegree for different bounds $d = 1$, $d = 2$, $d \geq 3$. Without much effort our results on trees of bounded outdegrees can be transferred to infinite words and to *ranked* trees, via reductions similar to Lemma 4.

Throughout the paper we used the simplest possible measure of formula size: the length of a formula written as a string. Alternative more succinct measures, such as the number of non-isomorphic subformulae (*DAG-size*), are also interesting. Thus, a natural question is to what extent our results depend on the choice of size measure. In principle, using a more succinct measure makes the problems of definability and separability harder. However, all our decision procedures, with an exception of Theorem 21, are automata-based. Consequently, these procedures carry over with unchanged complexity to any size measure for which the

translation from logic to nondeterministic automata has the same complexity as in Theorem 2. In the remaining case of Theorem 21 a sufficient assumption is that the modal depth of a formula is at most polynomial in its size. Both the mentioned assumptions are arguably modest.

A place where the choice of size measure matters a little more is the *construction* of modal definitions and separators. In the cases of unrestricted, unary (\mathbb{T}^1), and high outdegree models (\mathbb{T}^d for $d \geq 3$) the constructed formulae have DAG-size essentially the same as size: doubly, singly, and triply exponential, respectively. Interestingly, however, in the binary case \mathbb{T}^2 our formulae have only singly exponential DAG-size, which is easily seen to be optimal and contrasts with their doubly exponential size. This demonstrates that the lower bounds for size of modal definitions over \mathbb{T}^2 cannot work for DAG-size. The same lower bound construction fails for DAG-size over unrestricted models, although there the exact DAG-size complexity of optimal separators remains unknown.

We mention some interesting open problems. First, the relative succinctness of μML over ML is to the best of our knowledge open in the setting with only one accessibility relation. Second, as we have mentioned in Section 3, the separators we compute are not necessarily the logically strongest ones. The logically strongest separators of φ, φ' are precisely the ML -uniform consequences of φ (if they exist) and are a natural object of study. Clearly, modal definability of φ is a sufficient condition, but not a necessary one. Let $\varphi = \psi \wedge \neg\theta_\infty$ and $\varphi' = \psi$ for some $\psi \in \text{ML}$. Then φ is not equivalent to a modal formula, but ψ is a strongest separator. In the context of $\text{gr}\mu\text{ML}$, open questions are ML -definability (and separability) and μML -definability (and separability) of $\text{gr}\mu\text{ML}$ -formulae. We conjecture them to be easier than 2-EXPTIME. Finally, let us mention that definability of μML -formulae by *safety formulae* has been studied in [22]. It would be natural to investigate separability there as well.

References

- 1 Alessandro Artale, Jean Christoph Jung, Andrea Mazzullo, Ana Ozaki, and Frank Wolter. Living without Beth and Craig: Definitions and interpolants in description and modal logics with nominals and role inclusions. *ACM Trans. Comput. Log.*, 24(4):34:1–34:51, 2023. doi:10.1145/3597301.
- 2 Franz Baader, Ian Horrocks, Carsten Lutz, and Ulrike Sattler. *An Introduction to Description Logic*. Cambridge University Press, 2017.
- 3 Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- 4 Fabio Bellissima. On the lattice of extensions of the modal logics KAlt_n . *Arch. Math. Log.*, 27(2):107–114, 1988. doi:10.1007/BF01620760.
- 5 Michael Benedikt, Balder ten Cate, Thomas Colcombet, and Michael Vanden Boom. The complexity of boundedness for guarded logics. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS*, pages 293–304. IEEE Computer Society, 2015. doi:10.1109/LICS.2015.36.
- 6 Achim Blumensath, Martin Otto, and Mark Weyer. Decidability results for the boundedness problem. *Log. Methods Comput. Sci.*, 10(3), 2014. doi:10.2168/LMCS-10(3:2)2014.
- 7 Mikołaj Bojańczyk and Wojciech Czerwiński. *Automata Toolbox*. University of Warsaw, 2018. URL: <https://www.mimuw.edu.pl/~bojan/papers/toolbox.pdf>.
- 8 Ashok K. Chandra, Dexter Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981. doi:10.1145/322234.322243.
- 9 Sang Cho and Dung T. Huynh. Finite-automaton aperiodicity is PSpace-complete. *Theor. Comput. Sci.*, 88(1):99–116, 1991. doi:10.1016/0304-3975(91)90075-D.
- 10 Maarten de Rijke. A note on graded modal logic. *Stud Logica*, 64(2):271–283, 2000. doi:10.1023/A:1005245900406.

- 11 Kit Fine. In so many possible worlds. *Notre Dame J. Formal Log.*, 13(4):516–520, 1972. doi:10.1305/NDJFL/1093890715.
- 12 Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *J. Comput. Syst. Sci.*, 18(2):194–211, 1979. doi:10.1016/0022-0000(79)90046-1.
- 13 Tim French, Wiebe van der Hoek, Petar Iliev, and Barteld P. Kooi. On the succinctness of some modal logics. *Artif. Intell.*, 197:56–85, 2013. doi:10.1016/j.artint.2013.02.003.
- 14 Martin Fürer. The computational complexity of the unconstrained limited domino problem (with implications for logical decision problems). In *Logic and Machines: Decision Problems and Complexity, Proceedings of the Symposium "Rekursive Kombinatorik"*, volume 171 of *Lecture Notes in Computer Science*, pages 312–319. Springer, 1983. doi:10.1007/3-540-13331-3_48.
- 15 Dov M. Gabbay. Craig’s interpolation theorem for modal logics. In *Conference in Mathematical Logic — London ’70*, pages 111–127, Berlin, Heidelberg, 1972. Springer Berlin Heidelberg.
- 16 Gerd G. Hillebrand, Paris C. Kanellakis, Harry G. Mairson, and Moshe Y. Vardi. Undecidable boundedness problems for datalog programs. *J. Log. Program.*, 25(2):163–190, 1995. doi:10.1016/0743-1066(95)00051-K.
- 17 Jean Christoph Jung and Jędrzej Kołodziejski. Modal separability of fixpoint formulae. In *Proceedings of the 37th International Workshop on Description Logics (DL 2024)*, volume 3739 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2024. URL: <https://ceur-ws.org/Vol-3739/paper-5.pdf>.
- 18 Jean Christoph Jung and Frank Wolter. Living without Beth and Craig: Definitions and interpolants in the guarded and two-variable fragments. In *Proceedings of Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–14. IEEE, 2021. doi:10.1109/LICS52264.2021.9470585.
- 19 Eryk Kopczynski. Invisible pushdown languages. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 867–872. ACM, 2016. doi:10.1145/2933575.2933579.
- 20 Dexter Kozen. Results on the propositional mu-calculus. *Theor. Comput. Sci.*, 27:333–354, 1983. doi:10.1016/0304-3975(82)90125-6.
- 21 Louwe Kuijter, Tony Tan, Frank Wolter, and Michael Zakharyashev. Separating counting from non-counting in fragments of two-variable first-order logic (extended abstract). In *Proc. of DL 2024*, 2024.
- 22 Karoliina Lehtinen and Sandra Quickert. Deciding the first levels of the modal mu alternation hierarchy by formula construction. In *Proceedings of Annual Conference on Computer Science Logic CSL*, volume 41 of *LIPICs*, pages 457–471. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPICs.CSL.2015.457.
- 23 Christof Löding and Christopher Spinrath. Decision problems for subclasses of rational relations over finite and infinite words. *Discrete Mathematics & Theoretical Computer Science*, Vol. 21 no. 3, Jan 2019. doi:10.23638/DMTCS-21-3-4.
- 24 Martin Otto. Eliminating recursion in the μ -calculus. In *Proceedings of 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 531–540. Springer, 1999. doi:10.1007/3-540-49116-3_50.
- 25 Martin Otto. Graded modal logic and counting bisimulation. *CoRR*, abs/1910.00039, 2019. URL: <http://arxiv.org/abs/1910.00039>, arXiv:1910.00039.
- 26 Thomas Place and Marc Zeitoun. Separating regular languages with first-order logic. *Log. Methods Comput. Sci.*, 12(1), 2016. doi:10.2168/LMCS-12(1:5)2016.
- 27 Vaughan R. Pratt. A decidable mu-calculus: Preliminary report. In *Proceedings of 22nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 421–427. IEEE Computer Society, 1981. doi:10.1109/SFCS.1981.4.
- 28 Abraham Robinson. A result on consistency and its application to the theory of definition. *Journal of Symbolic Logic*, 25(2):174–174, 1960. doi:10.2307/2964240.
- 29 Marcel Paul Schützenberger. On finite monoids having only trivial subgroups. *Inf. Control.*, 8(2):190–194, 1965. doi:10.1016/S0019-9958(65)90108-7.

63:20 Modal Separation of Fixpoint Formulae

- 30 A. Prasad Sistla and Edmund M. Clarke. The complexity of propositional linear temporal logics. *J. ACM*, 32(3):733–749, 1985. doi:10.1145/3828.3837.
- 31 Moshe Y. Vardi. Reasoning about the past with two-way automata. In *Proceedings of International Colloquium Automata, Languages and Programming (ICALP)*, volume 1443 of *Lecture Notes in Computer Science*, pages 628–641. Springer, 1998. doi:10.1007/BFB0055090.
- 32 Yde Venema. Lectures on the modal μ -calculus, 2020.

A Proofs for Section 3

A.1 Correctness of the Construction of n -uniform Consequences

We prove (1). The base case is immediate so we focus on the inductive case with:

$$\psi_{n+1,q} = \bigvee_{c \in \Sigma} \bigvee_{S \in \delta(q,c)} c \wedge \nabla \{\psi_{n,p} \mid p \in S\}.$$

Fix \mathcal{M} , denote the color of the root by c and the set of all children of the root by M_0 . If \mathcal{M} satisfies $\psi_{n+1,q}$ then there is $\{p_1, \dots, p_l\} = S \in \delta(q, c)$ such that nabla ∇ of $\Phi = \{\psi_{n,p} \mid p \in S\}$ is satisfied in the root.

We show that without loss of generality the root v_I of \mathcal{M} has sufficiently many children to find a separate witness for each $\psi_{n,p} \in \Phi$. That is, we want to show that there is a surjective assignment $h : M_0 \rightarrow \Phi$ that maps every $v \in M_0$ to some formula $\psi_{n,p}$ true in v . With no restrictions on outdegree this follows from bisimulation-invariance because we can duplicate children of v_I and their subtrees. In the binary case we cannot freely duplicate children so a different argument is required. We want a different child $v_p \models \psi_{n,p}$ in M_0 for each $\psi_{n,p} \in \Phi$. The only scenario in which we do *not* have such unique witnesses is when S has two elements p_1 and p_2 (recall that \mathcal{B} is an automaton over binary trees and so it has no transitions to *more* than two states) and M_0 has only one element v . Let \mathcal{M}' be a model obtained from \mathcal{M} by duplicating v and its subtree. Clearly, $\mathcal{M} \equiv \mathcal{M}'$ and so \mathcal{M}' satisfies the nabla of Φ . Moreover, it has the desired separate witnesses for φ_{n,p_1} and φ_{n,p_2} .

By induction hypothesis, for each $v \in M_0$ with $h(v) = \psi_{n,p}$ there is a model $\mathcal{N}_p \models \mathcal{B}[q_I \leftarrow p]$ n -bisimilar to the subtree of \mathcal{M} rooted in v . Define \mathcal{N} as follows: first take the disjoint union $\{v\} \sqcup \bigsqcup \{\mathcal{N}_p \mid p \in S\}$ of all the \mathcal{N}_p 's and a fresh point v of color c ; then for every \mathcal{N}_p add an edge from v to the root of \mathcal{N}_p and set v as the new root. It is easy to see that $\mathcal{N} \models \mathcal{B}[q_I \leftarrow q]$ and $\mathcal{M} \equiv^{n+1} \mathcal{N}$, as desired.

Conversely, assume $\mathcal{M} \equiv^{n+1} \mathcal{N}$ and $\mathcal{N} \models \mathcal{B}[q_I \leftarrow q]$ witnessed by an $(n+1)$ -bisimulation Z and a run $\rho : N \rightarrow Q$. Denote the children of the root of \mathcal{N} by N_0 . Since ρ is a run, the set $S = \rho[N_0]$ of states assigned to N_0 belongs to $\delta(q, c)$. Every $v \in M_0$ is n -bisimilar to some $w \in N_0$ and hence by the induction hypothesis satisfies $\psi_{n,p}$ for $p = \rho(w) \in S$. Symmetrically, for every $p \in S$ there is $w \in N_0$ accepted by $\mathcal{B}[q_I \leftarrow p]$. Since that w is n -bisimilar to some $v \in M_0$, by induction hypothesis v satisfies $\psi_{n,p}$. It follows that the root of \mathcal{M} satisfies $\nabla \{\psi_{n,p} \mid p \in S\}$ and therefore also $\psi_{n+1,q}$.

B Proofs for Section 4

We provide the full construction of the ML_σ^n -uniform consequence of an NPTA \mathcal{A} over \mathbb{T}^1 .

► **Proposition 15.** *Let \mathcal{A} be an NPTA over \mathbb{T}^1 with ℓ states, $n \in \mathbb{N}$, and σ a signature. An ML_σ^n -uniform consequence of \mathcal{A} over \mathbb{T}^1 can be constructed in time polynomial in n , σ , and ℓ .*

Proof. Let \mathcal{A} be an NPTA and $n \in \mathbb{N}$ and σ a signature. Let $\mathcal{B} = (Q, \Sigma, \delta, q_I, \text{rank})$ with $\Sigma = \mathcal{P}(\sigma)$ be an NPTA of the same size that recognizes σ -reducts of models of \mathcal{A} . A formula ψ is an ML_σ^n -uniform consequence of \mathcal{A} iff it is an ML^n -uniform consequence of \mathcal{B} . We first construct formulae $\psi_{pq}^m \in \text{ML}_\sigma^n$, for $m \leq n$ and $p, q \in Q$ such that for every $\mathcal{M} \in \mathbb{T}^1$:

(*) $\mathcal{M} \models \psi_{pq}^m$ iff there is a run of \mathcal{B} from p to q on the m -prefix of \mathcal{M} .

63:22 Modal Separation of Fixpoint Formulae

The definition is by induction on m :

$$\begin{aligned}\psi_{pq}^0 &= \text{if } p \neq q \text{ then } \perp \text{ else } \top \\ \psi_{pq}^1 &= \bigvee \{c \mid c \in \Sigma, \{q\} \in \delta(p, c)\} \\ \psi_{pq}^m &= \bigvee_{q' \in Q} (\psi_{pq'}^{\lfloor m/2 \rfloor} \wedge \diamond^{\lfloor m/2 \rfloor} \psi_{q'q}^{\lceil m/2 \rceil}) \quad \text{for } 1 < m \leq n\end{aligned}$$

It is routine to verify that ψ_{pq}^m satisfies $(*)$ and is of size $|\psi_{pq}^m| \in O(|Q| \cdot m^2)$.

Before we can construct the desired ML_σ^n -uniform consequence, we need to introduce some more notation. Since we are working over \mathbb{T}^1 , $\delta(q, c)$ contains only sets of cardinality at most 1. The case $\emptyset \in \delta(q, c)$ is of particular interest because this means that the automaton in state q reading color c can “accept” even if it has not finished reading the input; in particular, the automaton can accept finite words as well. Denote with Acc_q the set of $c \in \Sigma$ with $\emptyset \in \delta(q, c)$. Further denote with Cont_q the set of all c such that $\mathcal{B}[q_I \leftrightarrow q]$ accepts a word starting with c . We finish the construction by setting:

$$\psi_n = \bigvee_{q \in Q} \left(\psi_{q_I q}^n \wedge \square^n \bigvee_{c \in \text{Cont}_q} c \right) \vee \bigvee_{m \leq n} \bigvee_{q \in Q} \left(\psi_{q_I q}^m \wedge \square^{m+1} \perp \wedge \square^m \bigvee_{c \in \text{Acc}_q} c \right).$$

It is readily checked that ψ_n satisfies the required size bounds. To verify that $\psi_n \models \theta$ for every $\theta \in \text{ML}^n$ with $\mathcal{B} \models \theta$, we show the following equivalence for all $\mathcal{M} \in \mathbb{T}^1$:

$$\mathcal{M} \models \psi_n \quad \iff \quad \text{there exists } \mathcal{N} \models \mathcal{B} \text{ with } \mathcal{N} \simeq^n \mathcal{M}. \quad (4)$$

For \Rightarrow , fix $\mathcal{M} \in \mathbb{T}^1$ with $\mathcal{M} \models \psi_n$. If $\mathcal{M} \models \psi_{q_I q}^n \wedge \square^n \bigvee_{c \in \text{Cont}_q} c$ for some q , then by $(*)$, there is a run of \mathcal{B} from the initial state q_0 to some state $q \in Q$ when reading the n -prefix of \mathcal{M} , and the last color in the prefix is c . Since $c \in \text{Cont}_q$, we can extend the n -prefix of \mathcal{M} to a $\mathcal{N} \in \mathbb{T}^1$ accepted by \mathcal{B} . If $\mathcal{M} \models \psi_{q_I q}^m \wedge \square^{m+1} \perp \wedge \square^m \bigvee_{c \in \text{Acc}_q} c$, for some $m \leq n$ and $q \in Q$, then \mathcal{M} is a finite tree of depth m that is accepted by the automaton. We can take $\mathcal{N} = \mathcal{M}$ in this case.

For \Leftarrow , let $\mathcal{M} \in \mathbb{T}^1$ a tree such that there is some $\mathcal{N} \models \varphi$ with $\mathcal{N} \simeq^n \mathcal{M}$. The former condition implies that $\mathcal{N} \models \mathcal{B}$ and thus there is an accepting run ρ of \mathcal{B} on \mathcal{N} , and the latter implies that \mathcal{N} and \mathcal{M} coincide on their n -prefixes. We distinguish cases. If the depth of \mathcal{N} is greater than n , then the n -prefix of ρ ending in state q witnesses $\mathcal{M} \models \psi_{q_I q}^n \wedge \square^n \bigvee_{c \in \text{Cont}_q} c$. Otherwise, the depth of \mathcal{N} is $m \leq n$ and the run ρ ending in q witnesses that $\mathcal{M} \models \psi_{q_I q}^m \wedge \square^{m+1} \perp \wedge \square^m \bigvee_{c \in \text{Acc}_q} c$. \blacktriangleleft

C Proofs for Section 5

C.1 Proof of Proposition 16

Assuming full binary $\mathcal{M}, \mathcal{M}'$ and a bisimulation Z between them we pick a bijective bisimulation $Z' \subseteq Z$. To that end, we inductively construct a descending sequence:

$$Z \supseteq Z_1 \supseteq Z_2 \supseteq Z_3 \supseteq \dots$$

of bisimulations such that for each n the restriction of Z_n to the n -prefixes of \mathcal{M} and \mathcal{M}' is bijective. The induction base $n = 1$ is trivial with $Z_1 = Z$. For the induction step $n + 1$ let $Z_n \subseteq Z$ be the bisimulation given by the inductive hypothesis. Z_n bijectively maps the

points v_1, \dots, v_k at depth n in \mathcal{M} to the respective points w_1, \dots, w_k at depth n in \mathcal{M}' . For each i we have $v_i Z w_i$. Hence, either both v_i and w_i are leaves (i), or both have respective children v_i^l, v_i^r and w_i^l, w_i^r . In the latter case either (ii) $v_i^l Z w_i^l$ and $v_i^r Z w_i^r$ or (iii) $v_i^l Z w_i^r$ and $v_i^r Z w_i^l$ (the cases (ii) and (iii) are not exclusive). Consider the bijective relation $K^i \subseteq Z$ between children of v_i and children of w_i :

$$K^i = \begin{cases} \emptyset & \text{if } v_i \text{ and } w_i \text{ are leaves (i),} \\ \{(v_i^l, w_i^l), (v_i^r, w_i^r)\} & \text{if (ii),} \\ \{(v_i^l, w_i^r), (v_i^r, w_i^l)\} & \text{otherwise.} \end{cases}$$

The bisimulation Z_{n+1} is constructed as follows. It is identical to Z_n between points at levels at most n , to $\bigcup_{i \leq k} K^i$ between points at level exactly $n+1$, and to Z between points at strictly greater levels. No points at mismatching levels are linked. It is straightforward to verify that such $Z_{n+1} \subseteq Z_n$ is a bisimulation, and that its restriction to the $n+1$ -prefixes of \mathcal{M} and \mathcal{M}' is bijective.

We conclude the proof by taking the limit $Z' = \bigcap_{n \in \mathbb{N}} Z_n$ as the desired bijective bisimulation between \mathcal{M} and \mathcal{M}' .

C.2 Craig Interpolation over \mathbb{T}^2 from Proposition 16

TODO in the future.

D Proofs for Section 6

The lower bound proofs for Section 6 rely on the families of formulae $(\psi_i)_{i \in \mathbb{N}}, (\psi'_i)_{i \in \mathbb{N}}$ that were already defined in the main part of the paper.

► **Lemma 28.** *There are families of formulae $(\psi_i)_{i \in \mathbb{N}}, (\psi'_i)_{i \in \mathbb{N}}$ such that:*

1. *The size of the formulae ψ_i, ψ'_i is polynomial in i .*
2. *$\text{sig}(\psi_i) = \{a, b_0, \dots, b_{i-1}\}$ and $\text{sig}(\psi'_i) = \{a, c\}$.*
3. *For every $i \in \mathbb{N}$ and arbitrary formulae $\chi, \chi' \in \mu\text{ML}$ the following holds. For every $\mathcal{M}, \mathcal{M}' \in \mathbb{T}^3$ with $\mathcal{M} \models \psi_i \wedge \chi$, $\mathcal{M} \models \psi'_i \wedge \chi'$, for every signature τ that contains a but not c, b_0, \dots, b_{i-1} , and every (τ, ℓ) -bisimulation S witnessing $\mathcal{M} \stackrel{\ell}{\approx}_{\tau} \mathcal{M}'$ for some $\ell \geq i$, there are points w_0, \dots, w_{2^i-1} in depth i in \mathcal{M} and a point \widehat{w} in depth i in \mathcal{M}' such that $(w_j, \widehat{w}) \in S$ for all j and such that distinct w_j, w_k can be distinguished by some proposition in b_0, \dots, b_{i-1} .*

Proof. We take the following formulae inductively defined as in the main part.

$$\begin{aligned} \psi_0 &= \psi'_0 = \top \\ \psi_{i+1} &= \diamond(a \wedge b_i) \wedge \diamond(a \wedge \neg b_i) \wedge \square(\psi_i \wedge (b_i \rightarrow \bigwedge_{j < i} \square^j b_i) \wedge (\neg b_i \rightarrow \bigwedge_{j < i} \square^j \neg b_i)) \\ \psi'_{i+1} &= \diamond(\neg a \wedge c) \wedge \diamond(\neg a \wedge \neg c) \wedge \diamond(a \wedge \psi'_i) \end{aligned}$$

Properties 1–3 can be verified by induction on i . ◀

D.1 Proof of Theorem 21

The proof of Theorem 21 is based on the following (straightforward) consequence of Equivalence (Base).

► **Lemma 29.** *Let $\varphi, \varphi' \in \text{ML}^n$ and σ be a signature. Then the following are equivalent for all classes \mathbb{T}^d :*

1. φ, φ' do not admit an ML_σ -separator over \mathbb{T}^d .
2. φ, φ' are joint \Leftrightarrow_σ^n -consistent over \mathbb{T}^d .

Proof. The implication from the first to the second item is immediate consequence of Equivalence (Base). For the converse implication, suppose φ, φ' are joint \Leftrightarrow_σ^m -consistent and let $\mathcal{M} \models \varphi, \mathcal{M}' \models \varphi'$ with $\mathcal{M} \Leftrightarrow_\sigma^m \mathcal{M}'$ witness this. Since φ, φ' have modal depth at most m , we can assume without loss of generality that $\mathcal{M}, \mathcal{M}'$ have depth at most m . But then, $\mathcal{M}, \mathcal{M}'$ witness joint \Leftrightarrow_σ^n -consistency of φ, φ' for every $n \geq m$. By Equivalence (Base) φ, φ' do not admit an ML_σ -separator. ◀

► **Theorem 21.** *For $d \geq 3$, ML -interpolant existence over \mathbb{T}^d is CONEXPTIME -complete. Hardness already applies to Craig ML -interpolant existence over \mathbb{T}^d .*

Proof. We start with the upper bound. By Lemma 29, interpolant *non*-existence can be decided by a standard “guess-and-check”-procedure:

1. Guess two structures $\mathcal{M}, \mathcal{M}' \in \mathbb{T}^d$ of depth at most the maximal modal depth m of φ, φ' ,
2. Verify that $\mathcal{M} \models \varphi, \mathcal{M}' \models \varphi'$, and $\mathcal{M} \Leftrightarrow_\sigma^m \mathcal{M}'$.

The runtime of the procedure is exponential: the size of any structure $\mathcal{M} \in \mathbb{T}^d$ of depth m is bounded by d^m . Moreover, model checking in modal logic and bisimulation testing is possible in polynomial time in the size of the structure and the given formulae.

We next show the lower bound for $d = 3$, but the proof is analogous for $d \geq 4$. We reduce an appropriate tiling problem. Let Δ be a finite set of *tile types*, and $V, H \subseteq \Delta \times \Delta$ be *vertical* and *horizontal* compatibility relations. Then, a mapping $\tau : [n] \times [n] \rightarrow \Delta$ is called a *torus tiling* for Δ, V, H, n if for all $i, j \in [n]$, we have:

- $(\tau(i, j), \tau(i \oplus 1, j)) \in H$ and
- $(\tau(i, j), \tau(i, j \oplus 1)) \in V$,

where \oplus denotes addition modulo n . The *exponential torus tiling problem* is to decide given Δ, V, H, n (in unary) whether there exists a torus tiling for $\Delta, V, H, 2^n$. It is well-known that the exponential torus tiling problem is NEXPTIME -complete [14].

Let Δ, V, H, n be an input to the exponential torus tiling problem. We will provide formulae φ_n, φ'_n of modal depth $4n$ and with common signature $\sigma = \text{sig}(\varphi_n) \cap \text{sig}(\varphi'_n)$ such that:

φ_n, φ'_n are joint $\Leftrightarrow_\sigma^{4n}$ -consistent iff there is a torus tiling for $\Delta, V, H, 2^n$.

The common signature σ will consist of propositions a, b , and one proposition t_d , for every $d \in \Delta$. Both φ_n and φ'_n will use auxiliary propositions to encode counters. The formulae φ_n, φ'_n are based on the families of formulae $(\psi_i)_{i \in \mathbb{N}}, (\psi'_i)_{i \in \mathbb{N}}$ defined in Lemma 28 and will take the shape:

$$\begin{aligned}\varphi_n &= \psi_{2n} \wedge \Box^{2n} \chi_1 \\ \varphi'_n &= \psi'_{2n} \wedge \Box^{2n} \chi_2\end{aligned}$$

for formulae χ_1, χ_2 to be defined below.

Consider models $\mathcal{M} \models \psi_{2n}$ and $\mathcal{M}' \models \psi'_{2n}$ with $\mathcal{M} \Leftrightarrow_\sigma^{4n} \mathcal{M}'$. Let $w_0, \dots, w_{2^{2n}-1}$ be the points in \mathcal{M} and \hat{w} be the point in \mathcal{M}' that exist due to Lemma 28. Recall that, by the lemma, all w_i are linked to \hat{w} by a $(\sigma, 2n)$ -bisimulation. We associate two numbers x_i, y_i with each point w_i as follows:

- x_i is the number encoded by the valuation of b_0, \dots, b_{n-1} in w_i , and
- y_i is the number encoded by the valuation of b_n, \dots, b_{2n-1} in w_i .

Note that by the properties of $w_0, \dots, w_{2^{2n}-1}$, for every pair x, y with $0 \leq x, y < 2^n$, there is some w_i with $x = x_i$ and $y = y_i$. We denote that point with $w(x, y)$. Hence, the numbers x_i, y_i can serve as addresses of the $2^n \times 2^n$ cells in the intended torus tiling. Intuitively, we will exploit that all w_i are linked to w by a $(\sigma, 2n)$ -bisimulation to synchronize the tile types in each cell. For what follows, it is convenient to denote with b_0^i, \dots, b_{2n-1}^i the values of propositions b_j in point w_i , and with c_0^i, \dots, c_{2n-1}^i and d_0^i, \dots, d_{2n-1}^i the value of the propositions b_j in the encoding of $x_i \oplus 1, y_i$ and $x_i, y_i \oplus 1$, respectively. It is not difficult to write formulae χ_1, χ_2 (of polynomial size) in modal logic (possibly using auxiliary non- σ propositions) that express the following Conditions 1 and 2, respectively. Recall that σ contains a, b and propositions t_d for every $d \in \Delta$.

1. There are three paths p_1, p_2, p_3 of length $2n$ with the following properties:
 - a. each path satisfies a in each point;
 - b. on p_1 , the j -th point satisfies b iff $b_{j-1}^i = 1$, for $1 \leq j \leq 2^n$;
 - c. on p_2 , the j -th point satisfies b iff $c_{j-1}^i = 1$, for $1 \leq j \leq 2^n$;
 - d. on p_3 , the j -th point satisfies b iff $d_{j-1}^i = 1$, for $1 \leq j \leq 2^n$;
 - e. the ends of p_1, p_2, p_3 are labeled with propositions $t_{d_1}, t_{d_2}, t_{d_3}$, respectively, such that $(d_1, d_2) \in H$ and $(d_1, d_3) \in V$.
2. There is a (ternary) tree of depth $2n$ with the following properties:
 - a. each node has three successors: one not satisfying a , one satisfying a and b , and one satisfying a and not b ;
 - b. in the leaves of this tree, we require that at most one proposition t_d is true.

This finishes the definition of the formulae φ_n, φ'_n and we can proceed with showing the correctness of the reduction.

Claim. $\psi_{2n} \wedge \Box^{2n} \chi_1$ and $\psi'_{2n} \wedge \Box^{2n} \chi_2$ are joint $\Leftrightarrow_{\sigma}^{4n}$ -consistent iff there is a torus tiling for $\Delta, H, V, 2^n$.

To prove the claim, suppose first that $\psi_{2n} \wedge \Box^{2n} \chi_1$ and $\psi'_{2n} \wedge \Box^{2n} \chi_2$ are joint $\Leftrightarrow_{\sigma}^{4n}$ -consistent, witnessed by models $\mathcal{M}, \mathcal{M}'$. Let also be $w_0, \dots, w_{2^{2n}-1}$ and \hat{w} be the points that exist in $\mathcal{M}, \mathcal{M}'$ due to Lemma 28.

We define a torus tiling τ as follows. Let x, y be any cell, that is, $0 \leq x, y < 2^n$. Conditions 1b–1d enforce that, three paths are stipulated: one (via 1b) in point $w(x, y)$, one (via 1c) in point $w(x \ominus 1, y)$, and one (via 1d) in point $w(x, y \ominus 1)$, where, similar to \oplus, \ominus denotes subtraction modulo 2^n . Each of these paths is labeled, using $b \in \sigma$ with (the encoding of) x, y along its elements. Due to Conditions 1a and 2a, these three paths (in \mathcal{M}) can only be bisimilar to the path (in \mathcal{M}') in the tree stipulated below \hat{w} that is labeled with (the encoding of) x, y . In particular, the ends of the paths are bisimilar to the same leaf in the tree. Since, by Condition 2b, every leaf in the tree satisfies at most one t_d , and, by Condition 1e, the end of each path satisfies at least one t_d , all ends are labeled with the same t_d . We set $\tau(x, y) = d$. Synchronization is then achieved by Condition 1e. Indeed, consider (x, y) and $(x \oplus 1, y)$. Then Condition 1e ensures that the end of the paths for (x, y) and $(x \oplus 1, y)$ stipulated in $w(x, y)$ are labeled with $t_d, t_{d'}$ such that $(d, d') \in H$. The argument for vertical compatibility is symmetric.

In the other direction, it is not difficult to construct joint $\Leftrightarrow_{\sigma}^{4n}$ -consistent models from a given torus tiling for $\Delta, H, V, 2^n$. ◀

D.2 Upper Bound Proof of Theorem 22

We show adequacy (3) of the game $\mathcal{G}_{\text{bisQuot}}(\mathcal{M}, \mathcal{A})$ defined in the main part.

We first prove the implication “ \Rightarrow ” from left to right. Assume a winning strategy ζ for $\mathcal{G}_{\text{bisQuot}}(\mathcal{M}, \mathcal{A})$. We construct $\mathcal{N} \in \mathbb{T}^d \models \mathcal{A}$ and $Z : \mathcal{N} \xrightarrow{\text{bis}} \mathcal{M}$. The universe N consists of all finite ζ -plays. There is an edge $\pi \rightarrow \pi'$ from π to π' if π' extends π with one move of $\exists\text{ve}$ followed by a response of $\forall\text{dam}$. This means that the outdegree of π equals the size of the transition S chosen by ζ as a response to π . In particular, \mathcal{N} is d -ary. Define $Z : N \rightarrow M$ such that $Z(\pi)$ is the point component from the last configuration in π . We complete the definition of \mathcal{N} by putting $\text{val}^{\mathcal{N}}(\pi) = \text{val}^{\mathcal{M}}(Z(\pi))$ for all $\pi \in N$.

We claim that the function Z is a bisimulation. The **atom** case follows immediately from the definition of $\text{val}^{\mathcal{N}}$. To prove the **back** and **forth** assume $\pi \in \mathcal{N}$. Denote $Z(\pi) = v$, let v_1, \dots, v_k be the children of v and $h : S \rightarrow \{v_1, \dots, v_k\}$ be the move chosen by ζ as a response to π .

To show the **forth** condition assume $\pi \rightarrow \pi'$. Existence of the edge $\pi \rightarrow \pi'$ implies that π' is of shape $\pi(h(p), p)$ for some $p \in S$. Let $v' = h(p)$. Since $v \rightarrow v'$ and $Z(\pi') = v'$ this completes the argument of the **forth** condition. Towards the **back** condition assume $v \rightarrow v'$ for some v' . Since v' is a child of v and h is surjective, there is $p \in S$ such that $h(p) = v'$. Then $\pi' = \pi(v', p)$ is a ζ -play such that $Z(\pi') = v'$ and $v \rightarrow v'$. Hence, v' witnesses the **back** condition.

It remains to construct an accepting run $\rho : N \rightarrow Q$. For each $\pi \in N$ we define $\rho(\pi)$ as the state component of the last configuration in π . This ρ is consistent with δ . To show this assume π with children π_1, \dots, π_k . Denote $\pi_i = \pi(v_i, p_i)$ for each $i \leq k$ and let (v, q) be the last configuration in π . There is $S \in \delta(q, \text{val}^{\mathcal{M}}(v))$ such that $S = \{p_1, \dots, p_k\}$. Since $\rho(\pi) = q$, $\rho(\pi_i) = p_i$ for every $i \leq k$ and $\text{val}^{\mathcal{N}}(\pi) = \text{val}^{\mathcal{M}}(v)$, the transition S is legal in π . To see that ρ is accepting assume an infinite path π_1, π_2, \dots in \mathcal{N} and for each i let (v_i, q_i) be the last configuration of π_i (so in particular $\rho(\pi_i) = q_i$). We need to show that the sequence $q_1 q_2 \dots$ of states satisfies the parity condition. This is true because each π_i is a ζ -play and therefore so is their infinite limit $(v_1, q_1)(v_2, q_2) \dots = \pi$. This completes the proof of the implication “ \Rightarrow ” in (3).

Let us prove the other implication “ \Leftarrow ” in (3). Assume $\mathcal{N} \in \mathbb{T}^d$ such that $\mathcal{N} \models \mathcal{A}$ and $Z : \mathcal{N} \xrightarrow{\text{bis}} \mathcal{M}$, and let $\rho : N \rightarrow Q$ be an accepting run witnessing $\mathcal{N} \models \mathcal{A}$. We construct a winning strategy ζ for $\exists\text{ve}$ in $\mathcal{G}_{\text{bisQuot}}(\mathcal{M}, \mathcal{A})$. The constructed strategy preserves as an invariant that for every ζ -play $\pi = (v_1, q_1) \dots (v_l, q_l)$ there is a path $w_1 \dots w_l$ in \mathcal{N} with $Z(w_i) = v_i$ and $\rho(w_i) = q_i$ for each $i \leq l$. The invariant holds in the initial position (v_I, q_I) . To define moves dictated by ζ assume a play $\pi = (v_1, q_1) \dots (v_l, q_l)$ and a path $w_1 \dots w_l$ from the invariant. Let v'_1, \dots, v'_k be the children of v_l . We define an $\exists\text{ve}$'s move $h : S \rightarrow \{v'_1, \dots, v'_k\}$ dictated by ζ as a response to π .

Let S be the transition chosen by ρ in w_l . By assumption no state of \mathcal{A} appears more than once in a single transition S of δ . This means that ρ is bijective between the set of children of w_l and S : for each $p \in S$ there exists a unique child w^p of w_l such that $\rho(w^p) = p$. We set $h(p) = Z(w^p)$ for every $p \in S$. By the **atom** condition w_l and v_l have the same color so to show that this h is a legal move for $\exists\text{ve}$ it suffices to show that it is surjective. By the **back** condition for every child v' of v_l there is some child w' of w_l with $Z(w') = v'$. Hence, $h(\rho(w')) = v'$ which proves surjectivity of h . Moreover, the invariant is preserved: if π' extends π by $\forall\text{dam}$'s response $(h(p), p)$ to h then $Z(w^p) = h(p)$ and $\rho(w^p) = p$ and so we extend $\bar{w} = w_1 \dots w_k$ with $w_{k+1} = w^p$.

To see that the strategy ζ is winning observe that, thanks to the invariant, for every infinite ζ -play $\pi = (v_1, q_1)(v_2, q_2) \dots$ there exists an infinite path $\bar{w} = w_1 w_2 \dots$ in \mathcal{N} such that

$Z(w_i) = v_i$ and $\rho(w_i) = q_i$ for all i . Thus, Eve wins π because ρ is accepting. This completes the proof of (3).

D.3 Lower Bound Proof of Theorem 22

We establish the following lower bound.

► **Lemma 30.** *ML-separability of μML -formulae over \mathbb{T}^3 is 2-EXPTIME-hard.*

The proof is a reduction of the word problem for languages recognized by exponentially space bounded, alternating Turing machines, which we introduce next.

An *alternating Turing machine (ATM)* is a tuple $\mathfrak{A} = (Q, \Theta, \Gamma, q_0, \Delta)$ where $Q = Q_{\exists} \uplus Q_{\forall}$ is a finite set of states partitioned into *existential states* Q_{\exists} and *universal states* Q_{\forall} . Further, Θ is the input alphabet and Γ is the tape alphabet that contains a *blank symbol* $\text{blank} \notin \Theta$, $q_0 \in Q_{\forall}$ is the *initial state*, and $\Delta \subseteq Q \times \Gamma \times Q \times \Gamma \times \{L, R\}$ is the *transition relation*. We assume without loss of generality that the set $\Delta(q, a) := \{(q', a', M) \mid (q, a, q', a', M) \in \Delta\}$ contains exactly zero or two points for every $q \in Q_{\exists} \cup Q_{\forall}$ and $a \in \Gamma$. Moreover, the state q' (in the transition (q, a, q', a', M)) must be from Q_{\forall} if $q \in Q_{\exists}$ and from Q_{\exists} otherwise, that is, existential and universal states alternate. A *configuration* of an ATM is a word wqw' with $w, w' \in \Gamma^*$ and $q \in Q$. We say that wqw' is *existential* if q is, and likewise for *universal*. *Successor configurations* are defined in the usual way. Note that every configuration has exactly zero or two successor configurations. A *computation tree* of an ATM \mathfrak{A} on input w is a (possibly infinite) tree whose nodes are labeled with configurations of \mathfrak{A} such that

- the root is labeled with the initial configuration q_0w ;
- if a node is labeled with an existential configuration wqw' , then it has a single successor which is labeled with a successor configuration of wqw' ;
- if a node is labeled with a universal configuration wqw' , then it has two successors which are labeled with the two successor configurations of wqw' .

An ATM \mathfrak{A} *accepts* an input w if there is a computation tree of \mathfrak{A} on w in which every branch is infinite.¹ It is well-known that there are 2^n -space bounded ATMs which recognize a 2-EXPTIME-hard language [8], where n is the length of the input w .

Let $\mathfrak{A} = (Q, \Theta, \Gamma, q_0, \Delta)$ be such a 2^n -space bounded ATM and $w = a_0 \dots a_{n-1}$ be an input of length n . We will provide formulae φ_n, φ'_n and signature σ such that


$$\varphi_n, \varphi'_n \text{ are joint } \Leftrightarrow_{\sigma}^m\text{-consistent for every } m \in \mathbb{N} \quad \text{iff} \quad \mathfrak{A} \text{ accepts } w.$$

This suffices by the Equivalence (Base). The signature σ consists of a, z , and propositions c_{α} for every possible cell content α of \mathfrak{A} , that is, $\alpha \in \Gamma \cup (Q \times \Gamma)$. Additionally, φ_n and φ'_n will use auxiliary propositions to encode counters. Both φ_n, φ'_n will use the family of formulae ψ_i, ψ'_i defined in Lemma 28. More precisely, the formulae φ_n, φ'_n will be of the form

$$\begin{aligned} \varphi_n &= \psi_n \wedge \Box^n \chi \\ \varphi'_n &= \psi'_n \wedge \chi' \end{aligned}$$

The only purpose of χ' is to mention the propositions in σ ; the main work is done by ψ_n, ψ'_n, χ . To see what χ achieves consider models $\mathcal{M}, \mathcal{M}'$ witnessing joint $\Leftrightarrow_{\sigma}^m$ consistency of φ_n, φ'_n for some sufficiently large $m \geq n$. By Lemma 28, there are points w_0, \dots, w_{2^n-1} in \mathcal{M} and

¹ This is a slight variation of the more standard acceptance condition in terms of accepting and rejecting states. It is, however, easily seen to be equivalent.



■ **Figure 3** Computation tree of \mathfrak{A} , repeated.

\hat{w} in \mathcal{M}' which are linked by a $(\sigma, m - n)$ -bisimulation. The idea is now that χ enforces below each w_i (the encoding of a) computation tree T_i of \mathfrak{A} on input w . The structure of these computation trees can easily be enforced using standard techniques. The challenge will be to synchronize successor configurations for which we exploit the mentioned bisimilarity between the w_j . We start with detailing the structure of the trees.

- Configurations of \mathfrak{A} are encoded by paths of length 2^n in which each point is labeled with exactly one c_α .
- Since accepting computations are infinite, it is convenient to enforce that every point has a successor. Moreover, along the infinite path we maintain a counter that counts modulo 2^n and thus dissects the path into pieces of length 2^n . Let's call this counter the C-counter.
- According to the structure of computation trees, every other 2^n points, the tree has outdegree two.

As mentioned before, the challenge is to coordinate successor configurations, and we rely on the bisimilarity of all w_i to do so. More precisely, we maintain an additional counter, the B-counter, that counts modulo 2^n as well. We reuse variables b_0, \dots, b_{n-1} for this purpose; recall that they are initialized to i at w_i for all $0 \leq i < 2^n$, in contrast to the C-counter which is initialized to 0 at each w_i . The idea is then to synchronize in the subtree below w_i consecutive points having C-counter 0. These points have distance 2^n and correspond to the same cell in consecutive configurations of \mathfrak{A} . Since all this is done with non- σ -propositions, this coordination is not “visible” across different subtrees below the w_i .

We will now provide χ more concretely. It is a conjunction $\chi = \chi_0 \wedge \Box^* \chi_1$. Here, χ_0 initializes the C-counter to 0 using propositional variables c_0, \dots, c_{n-1} and marks the first configuration as universal (recall that $q_0 \in Q_{\forall}$) using proposition p_{\forall} as follows:

$$\chi_0 = \neg c_0 \wedge \dots \wedge \neg c_{n-1} \wedge p_{\forall}.$$

Formula χ_1 in turn is a conjunction of several formulae. Note that due to the \Box^* in χ all formulae below are intended to hold in the entire subtree below w_i . One conjunct is responsible for incrementing the counters modulo 2^n . This is really standard, so we refrain from detailing it. Another conjunct is $\Diamond \top$ to enforce infinite trees. The most important conjunct is responsible for the synchronization and enforcing the structure of the computation tree. The structure of tree T_i below w_i is depicted in Figure 3. Universal configurations are marked with proposition p_{\forall} while existential configurations are marked with p_{\exists}^i , $i \in \{1, 2\}$, depending on the number of the successor configuration. Points shown as \circ mark the beginning of a configuration, that is, where the C-counter is 0.

Below, we will use $(C = i)$ (or similar expressions) as an abbreviation for the combination of the propositions c_0, \dots, c_{n-1} that encode value i . We next enforce the structure of the

computation tree:

$$\begin{aligned}
(C < 2^n - 1) \wedge p_{\forall} &\rightarrow \Box p_{\forall} && (\dagger) \\
(C < 2^n - 1) \wedge p_{\exists}^i &\rightarrow \Box p_{\exists}^i && i \in \{1, 2\} \\
(C = 2^n - 1) \wedge p_{\forall} &\rightarrow \Box(p_{\exists}^1 \vee p_{\exists}^2) \\
(C = 2^n - 1) \wedge p_{\exists}^i &\rightarrow \Box p_{\forall} && i \in \{1, 2\} \\
(C = 2^n - 1) \wedge p_{\forall} &\rightarrow \Diamond z \wedge \Diamond \neg z
\end{aligned}$$

These implications enforce that all points which represent a configuration satisfy one of $p_{\forall}, p_{\exists}^1, p_{\exists}^2$ indicating the kind of configuration and, if existential, also a choice of the transition function. The symbol $z \in \Sigma$ enforces the branching.

The initial configuration on input $w = a_0 \dots a_{n-1}$ is enforced by

$$c_{q_0, a_0} \wedge \Diamond(c_{a_1} \wedge \Diamond(c_{a_2} \wedge (\dots \wedge \Diamond(c_{a_{n-1}} \wedge \Diamond \chi_{blank}) \dots))),$$

where χ_{blank} enforces label c_{blank} until the end of the configuration (we omit the details).

To coordinate successor configurations, we associate with M functions f_i , $i \in \{1, 2\}$ that map the content of three consecutive cells of a configuration to the content of the middle cell in the i -th successor configuration (assuming an arbitrary order on the set $\Delta(q, a)$, for all q, a). In what follows, we ignore the corner cases that occur at the border of configurations; they can be treated in a similar way. Clearly, for each possible triple $(\alpha_1, \alpha_2, \alpha_3) \in (\Gamma \cup (Q \times \Gamma))^3$, the ML-formula $\varphi_{\alpha_1, \alpha_2, \alpha_3} = c_{\alpha_1} \wedge \Diamond(c_{\alpha_2} \wedge \Diamond c_{\alpha_3})$ is true at an points v of the computation tree iff v is labeled with c_{α_1} , a successors u of v is labeled with c_{α_2} , and a successors t of u is labeled with c_{α_3} . In each configuration, we synchronize points with B-counter 0 by including for every $(\sigma_1, \sigma_2, \sigma_3)$ and $i \in \{1, 2\}$ the following implications:

$$\begin{aligned}
(B = 2^n - 1) \wedge (C < 2^n - 2) \wedge \varphi_{\alpha_1, \alpha_2, \alpha_3} \wedge p_{\forall} &\rightarrow \Box q_{f_1(\alpha_1, \alpha_2, \alpha_3)}^1 \wedge \Box q_{f_2(\alpha_1, \alpha_2, \alpha_3)}^2 \\
(B = 2^n - 1) \wedge (C < 2^n - 2) \wedge \varphi_{\alpha_1, \alpha_2, \alpha_3} \wedge p_{\exists}^i &\rightarrow \Box q_{f_i(\alpha_1, \alpha_2, \alpha_3)}^i
\end{aligned}$$

At this point, the importance of the superscript in p_{\exists}^* becomes apparent: since different cells of a configuration are synchronized in different trees T_k the superscript makes sure that all trees rely on the same choice for existential configurations. Propositions q_{α}^i are used as markers (not in σ) and are propagated for 2^n steps, exploiting the C-counter. The superscript $i \in \{1, 2\}$ determines the successor configuration that the symbol is referring to. After crossing the end of a configuration, the symbol α is propagated using propositions q'_{α} (the superscript is not needed anymore because the branching happens at the end of the configuration, based on z).

$$\begin{aligned}
(C < 2^n - 1) \wedge q_{\alpha}^i &\rightarrow \Box q_{\alpha}^i \\
(C = 2^n - 1) \wedge p_{\forall} \wedge q_{\alpha}^1 &\rightarrow \Box(z \rightarrow A'_{\alpha}) \\
(C = 2^n - 1) \wedge p_{\forall} \wedge q_{\alpha}^2 &\rightarrow \Box(\neg z \rightarrow q'_{\alpha}) \\
(C = 2^n - 1) \wedge p_{\exists}^i \wedge q_{\alpha}^i &\rightarrow \Box q'_{\alpha} && i \in \{1, 2\} \\
(B < 2^n - 1) \wedge q'_{\alpha} &\rightarrow \Box q'_{\alpha} \\
(B = 2^n - 1) \wedge q'_{\alpha} &\rightarrow \Box q_{\alpha}
\end{aligned}$$

For those (q, a) with $\Delta(q, a) = \emptyset$, we add the conjunct

$$\neg c_{q, a}.$$

The following lemma establishes correctness of the reduction.

► **Lemma 31.** *The following conditions are equivalent:*

1. M accepts w ;
2. φ_n, φ'_n are joint \Leftrightarrow_σ^m -consistent for every $m \in \mathbb{N}$.

Proof. “1 \Rightarrow 2”. If M accepts w , there is a computation tree of \mathfrak{A} on w . We construct two models $\mathcal{M} \models \varphi_n$ and $\mathcal{M}' \models \varphi'_n$ such that $\mathcal{M} \Leftrightarrow_\sigma \mathcal{M}'$ (which implies joint \Leftrightarrow_σ^m -consistency for every m). Let $\widehat{\mathcal{M}}$ be the infinite tree-shaped model that represents the computation tree of \mathfrak{A} on w as described above, that is, configurations are represented by sequences of 2^n points and labeled by $p_\forall, p_\exists^1, p_\exists^2$ depending on whether the configuration is universal or existential, and in the latter case the superscript indicates which choice has been made for the existential state. Finally, the first points of the first successor configuration of a universal configuration is labeled with z . Observe that $\widehat{\mathcal{M}}$ interprets only the symbols in σ as non-empty. Now, we obtain models $\mathcal{M}_k, k < 2^n$ from $\widehat{\mathcal{M}}$ by interpreting non- σ -symbols as follows:

- the C -counter starts at 0 at the root and counts modulo 2^n along each path starting in the root;
- the B -counter starts at k at the root and counts modulo 2^n along each path starting in the root;
- the auxiliary concept names of the shape q_α^i and q'_α are interpreted in a minimal way so as to satisfy the implications starting from (\dagger). Note that, by definition of these implications, there is a unique result.

Now, \mathcal{M}' is defined as follows:

- start with a path of length n in which each node satisfies a ,
- add one successor satisfying $\neg a, \neg c$ and one successor satisfying $\neg a, c$ to each node in the path, and
- at the end \widehat{w} of the path, attach a copy of $\widehat{\mathcal{M}}$.

Next, obtain \mathcal{M} from the \mathcal{M}_k as follows:

- Start with a full binary tree \mathcal{M}_0 of depth n ,
- add one successor satisfying $\neg a$ to each node in the tree,
- interpret propositions b_0, \dots, b_{n-1} in a way such that the B -counter values of the 2^n leaves w_0, \dots, w_{2^n-1} (of the original binary tree) range from 0 to $2^n - 1$, and
- attach at each leaf w_k the tree \mathcal{M}_k .

It can be verified that the reflexive, transitive, and symmetric closure of

- all pairs (u, v) for points u satisfying a in \mathcal{M} at level i and points v satisfying a in \mathcal{M}' at level i ,
- all pairs (u, v) for points u satisfying $\neg a$ in \mathcal{M} at level i and points v satisfying $\neg a$ in \mathcal{M}' at level i , and
- all pairs (v, v') , with v in $\widehat{\mathcal{M}}$ and v' a copy of v in some tree \mathcal{M}_k

witnesses $\mathcal{M} \Leftrightarrow_\sigma \mathcal{M}'$.

“2 \Rightarrow 1”. Suppose φ_n, φ'_n are joint \Leftrightarrow_σ^m -consistent for every $m \in \mathbb{N}$. Since we work over models of finite outdegree and due to the form of φ_n, φ'_n (they do not contain any “eventualities” $\diamond^* \psi$ and no other fixpoints), we can construct in a standard way (by ‘skipping bisimulations’) models $\mathcal{M} \models \varphi_n$ and $\mathcal{M}' \models \varphi'_n$ such that $\mathcal{M} \Leftrightarrow_\sigma \mathcal{M}'$. By Lemma 28, there are pairwise σ -bisimilar points w_0, \dots, w_{2^n-1} . That is, the trees starting at w_0, \dots, w_{2^n-1} are bisimilar. Since these trees are additionally models of $\chi_0 \wedge \square^* \chi_1$ it follows that in the tree below w_k , the cell contents of the $(2^n - k)$ -th cell is coordinated, between any two consecutive configurations. Overall, all cell contents are coordinated and thus all trees below some w_k contain a computation tree of \mathfrak{A} on input w (which is solely represented with σ -symbols). Thus \mathfrak{A} accepts w . ◀

We show next that ML_σ -separability reduces to *modal separability* which is the problem whether two given $\varphi, \varphi' \in \text{ML}$ have an ML -separator, so without any signature restriction. This means that the hardness proved in Theorem 22 does not depend on having the signature available in the input.

► **Lemma 32.** *For any $d \geq 3$, there is a polynomial time reduction of ML -separability of μML -formulae over \mathbb{T}^d to modal separability of μML -formulae over \mathbb{T}^d .*

Proof. Let $d \geq 3$ and $\varphi, \varphi' \in \mu\text{ML}$ and σ be a signature. We construct μML -formulae $\widehat{\varphi}, \widehat{\varphi}'$ such that, over \mathbb{T}^d ,

$$\varphi, \varphi' \text{ are } \text{ML}_\sigma\text{-separable} \quad \text{iff} \quad \widehat{\varphi}, \widehat{\varphi}' \text{ are } \text{ML}\text{-separable.} \quad (\dagger)$$

Let $\sigma_\varphi = \text{sig}(\varphi) \setminus \sigma$ and $\sigma_{\varphi'} = \text{sig}(\varphi') \setminus \sigma$, respectively, be the sets of all propositions that occur in φ , respectively φ' , but not in σ . Let o be a fresh proposition (o as in “original”); this means that we extend the set **Prop** by o . Then,

$$\begin{aligned} \widehat{\varphi} &= o \wedge h(\varphi, \sigma_\varphi) \\ \widehat{\varphi}' &= o \wedge h(\varphi', \sigma_{\varphi'}) \end{aligned}$$

where $h(\psi, Q)$ is the formula obtained from ψ by

- first replacing each subformula $\diamond\theta$ with $\diamond o \wedge \square(o \rightarrow \diamond(o \wedge \theta))$ and each subformula $\square\theta$ with $\square(o \rightarrow \square(o \rightarrow \theta))$,
- replacing each subformula of the shape p for $p \in Q$ with $\diamond(\neg o \wedge \diamond^*p)$.

Intuitively, in the first step h relativizes ψ to o and “skips” every second level, and the second step replaces irrelevant propositions by something not modally definable and not visible in the o -relativization. Clearly, the size of $h(\varphi, Q)$ is polynomial in the size of its inputs. We show correctness of the construction (\dagger) .

For “ \Leftarrow ”, suppose that φ, φ' are not ML_σ -separable. By Equivalence (Base), for every $n \in \mathbb{N}$, there are models $\mathcal{M}_n \models \varphi, \mathcal{M}'_n \models \varphi'$ with $\mathcal{M}_n \not\equiv_\sigma^n \mathcal{M}'_n$. We construct models $\mathcal{N}_n \models \widehat{\varphi}, \mathcal{N}'_n \models \widehat{\varphi}'$ with $\mathcal{N}_n \equiv^{2n} \mathcal{N}'_n$ as follows. The model \mathcal{N}_n is obtained from \mathcal{M}_n as follows:

- introduce copies u' for every points u in \mathcal{M}_n ;
- make u' a successor of u and each successor of u (in \mathcal{M}_n) a successor of u' (in \mathcal{N}_n);
- make all points and their copies satisfy o ;
- add an infinite path satisfying $\neg o$ everywhere to every original point u .
- if an original points u satisfied p in \mathcal{M}_n for some $p \in \sigma_\varphi$, then it does not do so anymore in \mathcal{N}_n . Instead, p is made true on the infinite path starting in u at depth $2n + 1$.

The model \mathcal{N}'_n is obtained analogously from \mathcal{M}_n . Using the game theoretic semantics of μML it is not hard to show that:

Claim. $\mathcal{N}_n \models h(\varphi, \sigma_\varphi)$ and $\mathcal{N}'_n \models h(\varphi, \sigma_{\varphi'})$.

Let S be the (σ, n) -bisimulation witnessing $\mathcal{M}_n \equiv_\sigma^n \mathcal{M}'_n$. Based on the fact that, by construction of $\mathcal{N}_n, \mathcal{N}'_n$, no proposition from $\sigma_\varphi \cup \sigma_{\varphi'}$ appears in the first $2n$ levels of $\mathcal{N}_n, \mathcal{N}'_n$, it is not difficult to verify that S' defined as the union of

$$S \cup \{(u', v') \mid (u, v) \in S\}$$

and all pairs (u', v') such that there is an $(u, v) \in S$ and u' and v' both lie in the same distance at most $2n$ to u and v on the outgoing path starting in u and v , witnesses $\mathcal{N}_n \equiv^{2n} \mathcal{N}'_n$.

63:32 Modal Separation of Fixpoint Formulae

For “ \Rightarrow ”, suppose that $\widehat{\varphi}, \widehat{\varphi}'$ are not ML-separable. By Equivalence (Base), for every $n \in \mathbb{N}$, there are models $\mathcal{M}_n \models \widehat{\varphi}, \mathcal{M}'_n \models \widehat{\varphi}'$ with $\mathcal{M}_n \Leftrightarrow^{2n} \mathcal{M}'_n$. Note that we can assume without loss of generality that no proposition from $\sigma_\varphi \cup \sigma_{\varphi'}$ is satisfied in the first $2n$ levels of $\mathcal{M}_n, \mathcal{M}'_n$ (otherwise just push it down arbitrarily far). We can also assume that from a point not satisfying o we never a point satisfying o . Indeed, as we relativized our formulae to o , we can just obtain a new model from \mathcal{M}_n that makes o in every subtree rooted at a point not satisfying o and the new model will still be a model of $\widehat{\varphi}$. Finally, we can assume that, in even levels,² every point satisfying o has at most one successor satisfying o (this can be easily verified based on the replacement rules for $\diamond\psi$ and $\Box\psi$). Given an points u satisfying o at an even level, we denote with u' its unique successor satisfying o (if it exists).

We construct models $\mathcal{N}_n \models \varphi, \mathcal{N}'_n \models \varphi'$ such that $\mathcal{N}_n \Leftrightarrow^n \mathcal{N}'_n$ as follows.

- The domain of \mathcal{N}_n is the smallest set N such that:
 - N contains the root of \mathcal{M}_n , and
 - if $u \in N$ then N contains also all successors of u' (in \mathcal{M}_n) that satisfy o .
 - For $p \in \sigma$ and $u \in N_n$, we have $\mathcal{N}_n, u \models p$ iff $\mathcal{M}_n, u \models p$,
 - For $p \in \sigma_\varphi$ and $u \in N_n$, we have $\mathcal{N}_n, u \models p$ iff $\mathcal{M}_n, u \models \diamond(\neg o \wedge \diamond^* p)$,
- The model \mathcal{N}'_n is constructed analogously from \mathcal{M}'_n . Using the game theoretic semantics of μML it is not hard to show that:

Claim. $\mathcal{N}_n \models \varphi$ and $\mathcal{N}'_n \models \varphi'$.

Let S be any $2n$ -bisimulation witnessing $\mathcal{M}_n \Leftrightarrow^{2n} \mathcal{M}'_n$. It is routine to verify that the restriction of S to $N_n \times N'_n$ witnesses $\mathcal{N}_n \Leftrightarrow^n \mathcal{N}'_n$. ◀

² The root is on even level 0.