



INSTITUTE OF MATHEMATICS

THE CZECH ACADEMY OF SCIENCES

Rules with parameters in modal logic II

Emil Jeřábek

Preprint No. 24-2020

PRAHA 2020

Rules with parameters in modal logic II

Emil Jeřábek

The Czech Academy of Sciences, Institute of Mathematics
Žitná 25, 115 67 Praha 1, Czech Republic, email: jerabek@math.cas.cz

May 3, 2020

Abstract

We analyze the computational complexity of admissibility and unifiability with parameters in transitive modal logics. The class of cluster-extensible (clx) logics was introduced in the first part of this series of papers [8]. We completely classify the complexity of unifiability or inadmissibility in any clx logic as being complete for one of Σ_2^{exp} , NEXP, coNEXP, PSPACE, or Π_2^p . In addition to the main case where arbitrary parameters are allowed, we consider restricted problems with the number of parameters bounded by a constant, and the parameter-free case.

Our upper bounds are specific to clx logics, but we also include similar results for logics of bounded depth and width. In contrast, our lower bounds are very general: they apply each to a class of all transitive logics whose frames allow occurrence of certain finite subframes.

We also discuss the baseline problem of complexity of derivability: it is coNP-complete or PSPACE-complete for each clx logic. In particular, we prove PSPACE-hardness of derivability for a broad class of transitive logics that includes all logics with the disjunction property.

Keywords: modal logic, computational complexity, admissible rule, equational unification

MSC (2020): 03B45, 68Q17

1 Introduction

In this paper, we continue the analysis of admissible rules with parameters (constants) in transitive modal logics satisfying certain extension properties started in Jeřábek [8]. We recall that the first part was devoted to structural results. We introduced the class of *cluster-extensible* (clx) logics, encompassing the most common transitive modal logics such as **K4**, **S4**, **GL**, **S4Grz**, **S4.3**, and many others. We proved that in the setting of rules with parameters, all formulas have projective approximations in any clx logic L , whence L -admissibility is decidable, and we can compute finite complete sets of L -unifiers to any given formula. We provided semantic characterization of L -admissibility in terms of certain classes of frames (called L -extensible), and axiomatizations of L -admissible rules by explicit bases.

The topic of this second part is the *computational complexity* of admissibility with parameters, and of the closely related problem of unifiability, in clx and other transitive logics. We

mention that the complexity of admissibility in transitive modal logics was previously studied in [5]: the main results were that admissibility in certain logics (called extensible) is either coNP-complete or coNEXP-complete depending on if the logic is linear (i.e., of width 1) or not; the lower bound, stating that admissibility is coNEXP-hard, was proved under a weak hypothesis applicable to a larger class of logics. The class of extensible logics of [5] is incomparable with the clx logics of [8]: on the one hand, the condition of extensibility only constraints frames with a one-element root cluster, hence in this sense it is less restrictive; on the other hand, the definition of extensibility in [5] did not accommodate nonlinear logics of bounded branching. However, the principal difference is that [5] only considered admissibility *without parameters*.

As we will see, the introduction of parameters leads to a richer and more complicated landscape: we will encounter several more complexity classes than just coNP and coNEXP, and while linearity of the logic will remain an important dividing line, the complexity of the problem will also be influenced by other factors, the most important being if the logic allows clusters of unbounded size.

On the other hand, the usage of parameters makes our results on complexity more robust, and simpler to formulate. This is most clearly seen for lower bounds: first, they apply already to the special case of unifiability rather than to the full (in)admissibility problem, and second, they only require weak and easily checkable assumptions on the logics, such as being of width ≥ 2 , or having unbounded cluster size. (In contrast, the parameter-free coNEXP lower bound from [5] only applies to admissibility, and relies on a peculiar extensibility condition on the logic.) In general, our lower bounds will have the form that L -unifiability is \mathcal{C} -hard (for a particular complexity class \mathcal{C}) whenever certain finite frames may be embedded into L -frames, or more precisely, when there exist L -frames that subreduce or weakly subreduce to said finite frames; see below for definitions.

Wherever possible, we also include results on the complexity of restricted unifiability or admissibility problems in which only a constant number of parameters are allowed, though in this case the conditions on logics get more complicated, and do not seem optimal.

In contrast to lower bounds, decent upper bounds can be proved only for well-behaved logics, as random transitive modal logics may be quite wild, with already the set of tautologies being of high complexity or even undecidable. We are primarily interested in the case of clx logics, and one of the goals of this paper is to find a complete classification of the complexity of unifiability and admissibility in clx logics, both in the regime with arbitrary parameters, and with only constantly many parameters.

Additionally, we present upper bounds on the complexity of admissibility and unifiability in *logics of bounded depth and width*. To this end, we need first to prove a few structural results on admissibility in logics of bounded depth—in particular, semantical characterizations—since such logics are not covered by the framework from [8].

Similar to clx logics, logics of bounded depth and width are tame and well behaved, which allows us to prove interesting general results about their complexity. However, the class of these logics is structurally quite different from the class of clx logics. Clx logics are few and far between, they are particularly nice logics cherry-picked from the lattice NExt **K4** of all transitive logics, both weak and strong. In contrast, logics of bounded depth and width form an *ideal* in

NExt **K4**: in particular, all extensions of any logic from this class are also in the class. Thus, logics of bounded depth and width may be construed as a toy model that is representative of all well-behaved logics, whereas it may easily happen that there are classes of nice logics whose properties are quite different from clx logics. In fact, as we will see in the sequel, there *are* reasonable classes of nice non-clx logics, including logics with a single top cluster such as **S4.2**, and some of these have different complexity of unifiability and admissibility from what we will encounter in this paper. These cases duly manifest already for logics of bounded depth and width, and for this reason we cannot give a complete classification of complexity of unifiability in these logics in the present paper.

We acknowledge that this guiding principle should not be taken too seriously: for example, we will also see that there is an interesting class of logics for which unifiability and admissibility are PSPACE-complete, which cannot happen for logics of bounded depth and width, and likewise for all nontrivial results concerning the setting with constantly many parameters.

The main focus of the paper is on the complexity of admissibility and unifiability, but as a starting point, we also settle the complexity of *derivability* in clx logics. Generalizing results of Ladner [10], we show that nonlinear clx logics are PSPACE-complete. For the lower bound, we prove PSPACE-hardness for a broad class of transitive logics that includes all logics with the disjunction property, and many other logics such as **S4.2**: similar to our lower bounds on unifiability, the result applies to all transitive logics L such that, roughly speaking, all finite trees may be embedded into L -frames.

In the planned third part of this series of papers, we will adapt the set-up of clx logics to logics with a single top cluster such as **K4.2** and **S4.2**, including both structural results as in [8], and computation complexity results as here. Among other things, this will allow us to complete the classification of the complexity of unifiability and admissibility for logics of bounded depth and width, and to obtain a complete classification of *hereditary* properties of transitive logics that guarantee hardness of unifiability for some complexity class. Both of these suggest that our results on complexity are in a certain sense optimal. We also plan to adapt our results to superintuitionistic logics.

1.1 Overview of results

The linear organization of the paper is as follows. Section 2 contains a few preliminary definitions and facts; its Subsection 2.1 reviews the needed background in complexity theory, including a proof of completeness of certain problems for levels of the exponential hierarchy.

Section 3 deals with the complexity of derivability.

Section 4 is devoted to upper bounds on the complexity of admissibility and unifiability in certain logics: nonlinear and linear clx logics in Subsections 4.1 and 4.2 (respectively), and logics of bounded depth (and, apart from structural results, bounded width) in Subsection 4.3.

Section 5 presents lower bounds on the complexity of unifiability and admissibility. Hardness results for levels of the exponential hierarchy (NEXP, coNEXP, Σ_2^{EXP}) appear in Subsection 5.1, except for results in the setting with a constant number of parameters, which are more complicated and are relegated to Subsection 5.2. Hardness results for PSPACE and levels of the polynomial hierarchy are in Subsection 5.3.

logic		$\not\vdash_L$	unifiability, $\not\vdash_L$			examples
branching	cluster size		parameters:			
			no	$O(1)$	any	
0	$< \infty$	NP			Π_2^p	S5 \oplus Alt_k, Verum
	∞				coNEXP	S5, K4B
1	$< \infty$		PSPACE			GL.3, S4Grz.3
	∞				coNEXP	S4.3, K4.3
≥ 2	$< \infty$	PSPACE	NEXP			GL, K4Grz, S4Grz
	∞				Σ_2^{exp}	K4, S4, K4BB_k

Table 1: Complexity of nonderivability, unifiability, and inadmissibility problems for clx logics. Note: results in the parameter-free column only apply to inadmissibility, not unifiability.

Section 6 concludes the paper.

A summary of results on the complexity of derivability, unifiability, and admissibility in consistent clx logics is given in Table 1. Each entry in the table should be understood so that the stated problem is *complete* for the complexity class indicated, for every clx logic that meets the description. Notice that for a given logic, unifiability and admissibility usually have dual complexity, because unifiability is a special case of inadmissibility; in order to avoid confusing switching between dual complexity classes, we adopt in this overview the convention to indicate the complexity of *unifiability* and *inadmissibility* rather than admissibility. In accordance with this, we also indicate the complexity of *nonderivability* (or equivalently: local satisfiability) rather than derivability. However, detailed statements of theorems later in the paper will often mention both.

The referee expressed surprise that some of these classes are existential and some universal. On the most basic level, unifiability and inadmissibility are existential properties as they amount to the existence of certain substitutions, but this is an unbounded existential quantifier, and even though our results also imply computable bounds, this usually makes a quite inefficient test. We rely instead on semantic criteria characterizing inadmissibility by the existence of certain finite models, and depending on properties of the logic, different quantifiers in this characterization become dominant.

Roughly speaking, in the Σ_2^{exp} case, the outer \exists quantifies over a model, and the inner \forall comes from verification of pseudoextensibility of the model, quantifying over valuations of parameters in a cluster. If cluster size or the number of parameters is bounded, the \forall becomes polynomial-size and disappears, leading to NEXP. If the logic is linear, the model is an upside-down tree of clusters, and rather than fixing it beforehand, we can explore it one branch at a time; this leads to coNEXP as the outer \exists disappears. If both restrictions are met, we are left with only polynomial-size quantifiers: basically, alternating \forall over parameter valuations with \exists over variable valuations. This gives PSPACE; furthermore, if the logic has bounded depth, the number of alternations is likewise bounded, leading to Π_{2d}^p .

Here is a cross-reference of our complexity results sorted according to complexity classes:

- Σ_2^{exp}
 - Upper bounds: inadmissibility in clx logics (Thm. 4.2) and in logics of bounded depth and width (Thm. 4.25).
 - Lower bound: unifiability in logics with frames weakly subreducing to $(\circ + \textcircled{n})^\circ$ (Thm. 5.10).
- NEXP
 - Upper bounds: inadmissibility in clx logics with bounded cluster size or bounded number of parameters (Thm. 4.2), inadmissibility in tabular logics (Thm. 4.20).
 - Lower bounds: unifiability in nonlinear logics (Thm. 5.3), unifiability with $O(1)$ parameters in certain logics (Thms. 5.15, 5.18, 5.20, Cor. 5.16), inadmissibility without parameters in certain logics (Thm. 5.23).
- coNEXP
 - Upper bounds: inadmissibility in linear clx logics (Thm. 4.3) and in linear logics of bounded depth (Thm. 4.21).
 - Lower bound: unifiability in logics of unbounded cluster size (Thm. 5.6).
- PSPACE
 - Upper bounds: (non)derivability in clx logics (Thm. 3.3), (in)admissibility in linear clx logics with bounded cluster size or with $O(1)$ parameters (Thm. 4.3).
 - Lower bounds: (non)derivability in logics subframe-universal for trees (Thm. 3.12, cf. Cor. 3.13), unifiability with $O(1)$ parameters in logics of unbounded depth (Thms. 5.25, 5.31).
- Π_{2d}^p
 - Upper bound: inadmissibility in linear tabular logics (Thm. 4.21, cf. Cor. 4.22).
 - Lower bound: unifiability in logics of depth $\geq d$ (Thm. 5.25).
- NP
 - Upper bounds: nonderivability in consistent linear clx logics (Thm. 3.2), unifiability with $O(1)$ parameters in logics of bounded depth (Thm. 4.6), inadmissibility with $O(1)$ parameters in logics of bounded depth and width (Thm. 4.19), inadmissibility without parameters in consistent linear clx logics [5, Thm. 2.6].
 - Lower bound: unifiability without parameters in consistent logics (Thm. 5.2).

2 Preliminaries

This paper is a continuation of [8], and we assume the reader has access to that paper. We generally follow the same terminology and notational conventions as in [8], which we shall not repeat here, as it would considerably add to the length of the paper. In particular, we assume the reader is familiar with the content of [8, §2], which lays out our basic concepts and terminology. Moreover, we will rely on the definition of clx logics (Def. 4.1) and their main structural properties (§4.3); the definition of tight predecessors and extensible frames (Def. 5.1) along with the ensuing semantical characterization of admissibility in clx logics (Thm. 5.18); and the characterization of admissibility in terms of pseudoextensible models (Def. 5.21, Thm. 5.24).

We stress that unless stated otherwise, admissibility refers to admissibility with arbitrary parameters, and similarly for unification and other related notions.

All logics in this paper are normal modal logics extending **K4**, which we will not always state explicitly.

As a piece of a more obscure notation, we recall from [8] that $\mathbf{2} = \{0, 1\}$; if P is a finite set of formulas (in this paper, typically $P \subseteq \text{Par}$), then $\mathbf{2}^P$ is the set of all assignments $t: P \rightarrow \mathbf{2}$, and for any $t \in \mathbf{2}^P$, we put $P^t = \bigwedge_{\varphi \in P} \varphi^{t(\varphi)}$, where $\varphi^1 = \varphi$, $\varphi^0 = \neg\varphi$.

If $n \in \omega$ is a natural number (which we take to include 0), we will sometimes denote the set $\{0, \dots, n-1\} = \{i : i < n\}$ as just n . (This in fact agrees with the common von Neumann *definition* of natural numbers.) Given an indexed sequence $P = \{\varphi_i : i < n\}$ and $t \in \mathbf{2}^n$, we extend the notation above so that $P^t = \bigwedge_{i < n} \varphi_i^{t(i)}$. Finally, we will use this notation also for subsets $t \subseteq n$, that are identified in this context with their characteristic functions: that is, $P^t = \bigwedge_{i < n} \varphi_i^{[i \in t]}$. Here, $[\dots]$ is the *Iverson bracket*: for any predicate R ,

$$[R] = \begin{cases} 1 & \text{if } R \text{ is true,} \\ 0 & \text{otherwise.} \end{cases}$$

The set of all subsets of n is denoted $\mathcal{P}(n)$.

In accordance with [8, §5.1], we define a logic L to have the *disjunction property* if it admits the rules

$$(\text{DP}_n) \quad \bigvee_{i < n} \Box x_i \ / \ \{x_i : i < n\}$$

for all $n \in \omega$, or equivalently, for $n = 0, 2$ (the admissibility of DP_0 being equivalent to the consistency of L).

If C is a cluster or a point in a model, a cluster type, or an extension condition, let $\text{refl}(C) \in \{\bullet, \circ\}$ denote its *reflexivity*: $\text{refl}(C) = \bullet$ if C is irreflexive, and $\text{refl}(C) = \circ$ if it is reflexive.

Recall that a frame $\langle V, \prec, B \rangle$ is a *generated subframe* of $\langle W, <, A \rangle$ (written as $V \subseteq \cdot W$) if $V \subseteq W$, $\prec = < \cap (V \times W)$ (which implies $V \uparrow \subseteq V$), and $B = \{X \cap V : X \in A\}$. The *disjoint sum* of frames $\langle W_i, <_i, A_i \rangle$ ($i \in I$), denoted $\sum_{i \in I} W_i$, is the frame $\langle W, <, A \rangle$ whose underlying set W is the disjoint union $\bigcup_{i \in I} W_i$, $< = \bigcup_{i \in I} <_i$, and $A = \{X \subseteq W : \forall i \in I (X \cap W_i \in A_i)\}$. A *subreduction* of $\langle W, <, A \rangle$ to $\langle V, \prec, B \rangle$ (see [3, §9.1]) is a partial map $f: W \rightarrow V$ such that

$$(i) \ Y \in B \implies f^{-1}[Y] \in A,$$

(ii) $w < w' \implies f(w) \prec f(w')$, and

(iii) $f(w) \prec v \implies \exists w' > w \ f(w') = v$

for all $w, w' \in W$, $v \in V$, and $Y \subseteq V$. If V is a (Kripke) finite frame, condition (i) simplifies to

(i') $f^{-1}[v] \in A$ for all $v \in V$.

A subreduction is *cofinal* if $\text{dom}(f)\uparrow \subseteq \text{dom}(f)\downarrow$. A subreduction of W onto V is a subreduction whose image is all of V . A frame W (*cofinally*) *subreduces* onto V if there exists a (cofinal, resp.) subreduction of W onto V .

We will often encounter conditions concerning the occurrence of certain patterns as subframes that are oblivious to the reflexivity of individual points. In order to facilitate their formulation, we define a *weak subreduction* of W to V to be a partial map $f: W \rightarrow V$ that satisfies (i), (ii), and the modified condition

(iii') $f(w) \prec v \implies \exists w' \geq w \ f(w') = v$

for all $w \in W$ and $v \in V$. We define other derived notions such as weak subreductions onto V and cofinal weak subreductions similarly as above.

The *reflexivization* of a frame $\langle W, <, A \rangle$, denoted W_\circ , is the frame $\langle W, \leq, A \rangle$. The following is an immediate consequence of the definitions:

Observation 2.1 *Let $\langle W, <, A \rangle$ and $\langle V, \prec, B \rangle$ be frames, and f a partial map from W to V .*

- *If f is a weak subreduction of W to V , it is a subreduction of W_\circ to V_\circ .*

- *If f is a subreduction of W_\circ to V , it is a weak subreduction of W to V .* □

If $L \supseteq \mathbf{K4}$ is a logic with the finite model property, and P a finite set of parameters, let $U_L(P)$ denote the *universal L -frame for parameters P* : it is defined as the upper part of the universal L -frame of rank $|P|$ (denoted $\mathcal{F}_L^{<\infty}(|P|)$ in [3, §8.7]), endowed with the canonical valuation of parameters from P to make it a parametric frame. The double dual of $U_L(P)$ is the canonical frame $C_L(P, \emptyset)$ ([3, Cor. 8.89]); in particular, if $\varphi \in \text{Form}(P, \emptyset)$ is unprovable in L , then $u \not\models \varphi$ for some $u \in U_L(P)$.

Explicitly, we may construct $U_{\mathbf{K4}}(P)$ as the union $\bigcup_{n \in \omega} W_n$ of the following inductively defined chain of finite parametric frames:

- We start with $W_0 = \emptyset$.

- The frame W_{n+1} includes W_n as a generated subframe. Moreover, for every $X \subseteq W_n$, $* \in \{\bullet, \circ\}$, and $\emptyset \neq E \subseteq \mathbf{2}^P$, where $|E| = 1$ if $* = \bullet$: if X has no tight $\langle *, E \rangle$ -predecessor in W_n , we include one in W_{n+1} .

Then $U_L(P) \subseteq U_{\mathbf{K4}}(P)$ consists of all points $u \in U_{\mathbf{K4}}(P)$ such that $u\uparrow$ is an L -frame. We stress that if $X \subseteq W_n$ is a rooted subframe with a reflexive root cluster C such that all assignments from E are realized in C , then no tight $\langle *, E \rangle$ -predecessor of X is added to W_{n+1} , as it already has one (viz., a subset of C).

If L is a clx logic, we may also describe $U_L(P)$ in our terminology as the minimal locally finite $\text{bas}(L)$ -extensible parametric Kripke frame for the parameters P .

It is well known that every point of $U_L(P)$ is definable by a formula (see e.g. [3, Thm. 8.83]). We will need an explicit description of such formulas in the simplest case $P = \emptyset$. Notice that in this case, the universal frame has no proper clusters.

For each $u \in U_{\mathbf{K4}}(\emptyset)$, we define a formula β_u by induction on the depth of u as

$$(1) \quad \beta_u = \begin{cases} \bigwedge_{v>u} \diamond \beta_v \wedge \square \bigvee_{v>u} \beta_v, & \text{refl}(u) = \bullet, \\ \diamond \left(\bigwedge_{v \geq u} \diamond \beta_v \right) \wedge \square \left(\bigvee_{v \geq u} \beta_v \vee \diamond \bigwedge_{v \geq u} \diamond \beta_v \right), & \text{refl}(u) = \circ. \end{cases}$$

Lemma 2.2 *Let $L \supseteq \mathbf{K4}$ have fmp, $u, u' \in U_L(\emptyset)$, and $\varphi \in \text{Form}(\emptyset, \emptyset)$.*

$$(2) \quad u' \models \beta_u \quad \text{iff} \quad u' = u,$$

$$(3) \quad u \models \varphi \quad \text{iff} \quad \vdash_L \beta_u \rightarrow \varphi.$$

Proof: We prove (2) by induction on the depth of u . The right-to-left implication amounts to $u \models \beta_u$, which is straightforward to check.

Assume that $u' \models \beta_u$. If u is irreflexive, the definition of β_u and the induction hypothesis imply $u' \uparrow = u \uparrow$. Notice that there is no $v \in u \uparrow$ such that $v \uparrow = u \uparrow$: this would make $u \uparrow$ rooted with a reflexive root v , hence no additional tight predecessor u would be added to $U_L(\emptyset)$. Thus, $u' \notin u \uparrow$, hence u' is irreflexive, i.e., it is the unique irreflexive tight predecessor of $u \uparrow$: $u' = u$.

If u is reflexive, we have $X := u \uparrow \setminus \{u\} \subseteq u \uparrow$, which implies $u' \notin X$ as above. Assuming $u' \neq u$ for contradiction, let $w \geq u'$ be maximal such that $w \notin u \uparrow$. By maximality, $w \uparrow \setminus \{w\} \subseteq u \uparrow$. Since $u' \models \beta_u$, we have $w \models \diamond \bigwedge_{v \in X} \diamond \beta_v$, i.e., there is $w' > w$ such that $X \subseteq w' \uparrow$. This again implies $w' \notin X$, hence $w' = u$ or $w' = w$ by maximality, i.e., $w < u$ or w is reflexive. If $w < u$, then $w \uparrow \setminus \{w\} = u \uparrow$, i.e., w is a tight predecessor of $u \uparrow$ other than u itself, but no such tight predecessor was added into $U_L(\emptyset)$. Otherwise $w \uparrow \setminus \{w\} = X$ and w is reflexive, i.e., w is a reflexive tight predecessor of X other than u , which is again impossible.

(3): The right-to-left implication follows from $u \models \beta_u$. Left-to-right: assume that $\not\vdash_L \beta_u \rightarrow \varphi$. Since $(\beta_u \rightarrow \varphi) \in \text{Form}(\emptyset, \emptyset)$, this implies $u' \not\models \beta_u \rightarrow \varphi$ for some $u' \in U_L(\emptyset)$ by universality. We must have $u' = u$ by (2), thus $u \not\models \varphi$. \square

2.1 Complexity classes

Since the topic of this paper is computational complexity, we will assume some degree of familiarity with basic computation models and complexity classes. We refer the reader to e.g. Arora and Barak [1] for general background on complexity theory, but for convenience, we review the definitions of classes that appear in this paper.

For any function $t: \omega \rightarrow \omega$, let $\text{DTIME}(t)$ (or $\text{DTIME}(t(n))$ for emphasis) denote the class of all languages $L \subseteq \Sigma^*$ (where Σ is a finite alphabet) computable by a *deterministic (multitape) Turing machine (DTM)* in time at most $t(n)$, where n is the length of input. If F is a family of such functions, such as the family of all polynomially bounded functions $\text{poly}(n) = n^{O(1)}$, we

put $\text{DTIME}(F) = \bigcup_{t \in F} \text{DTIME}(t)$. The *polynomial time* and *exponential time* classes are then defined as

$$\begin{aligned} \text{P} &= \text{DTIME}(\text{poly}(n)), \\ \text{EXP} &= \text{DTIME}(2^{\text{poly}(n)}). \end{aligned}$$

(This is exponential time with polynomial exponent; exponential classes with linear exponent, such as $\text{E} = \text{DTIME}(2^{O(n)})$, will not be used in this paper.) Likewise, $\text{DSPACE}(s(n))$ is the class of languages computable by a DTM using $s(n)$ cells of memory, and

$$\text{PSPACE} = \text{DSPACE}(\text{poly}(n)).$$

For polynomial-space languages it does not matter, but recall that in general, space usage is defined so that it only accounts for the content of work tapes, excluding the input tape (which is assumed read-only), and—if we are computing a function rather than language membership—excluding the output tape (which is write-only).

In particular, our basic notion of reduction will be *many-one logarithmic-space (logspace) reductions*. If X is a complexity class, a language L_0 is *X-hard* if every language $L \in X$ is logspace-reducible to L_0 , and L_0 is *X-complete* if additionally $L_0 \in X$.

Nondeterministic Turing machines (NTM) may have a choice between multiple possible transitions in any configuration; the machine is declared to accept a given input if there *exists* a run of the machine that ends in an accepting state. The class of languages accepted by a NTM in time $t(n)$ is denoted $\text{NTIME}(t)$, and we put

$$\begin{aligned} \text{NP} &= \text{NTIME}(\text{poly}(n)), \\ \text{NEXP} &= \text{NTIME}(2^{\text{poly}(n)}). \end{aligned}$$

An equivalent definition of NP is that it consists of languages L such that membership in L can be *witnessed* by a polynomial-size certificate whose validity can be checked in P.

We could also define nondeterministic space classes, but the only example we are interested in is NPSpace , which equals PSPACE by Savitch's theorem.

For any class X of languages $L \subseteq \Sigma^*$, $\text{co}X$ denotes the dual class $\{\Sigma^* \setminus L : L \in X\}$.

The deterministic and nondeterministic time classes above can be relativized: for any language A , a *Turing machine with oracle A* may query membership of words in A (by writing them on a dedicated oracle query tape) at unit time cost. Then for any class X , NP^X denotes the set of languages computable in polynomial-time by a NTM with oracle $A \in X$, and similarly for other classes. The *polynomial* and *exponential hierarchies* are defined by $\Sigma_0^p = \Delta_0^p = \Pi_0^p = \text{P}$, and for $k > 0$,

$$\begin{aligned} \Sigma_k^p &= \text{NP}^{\Sigma_{k-1}^p}, & \Delta_k^p &= \text{P}^{\Sigma_{k-1}^p}, & \Pi_k^p &= \text{coNP}^{\Sigma_{k-1}^p}, \\ \Sigma_k^{\text{exp}} &= \text{NEXP}^{\Sigma_{k-1}^p}, & \Delta_k^{\text{exp}} &= \text{EXP}^{\Sigma_{k-1}^p}, & \Pi_k^{\text{exp}} &= \text{coNEXP}^{\Sigma_{k-1}^p}. \end{aligned}$$

(Notice that an (N)EXP machine may supply exponentially long queries to the oracle, hence the Σ_{k-1}^p oracle in the definition of Σ_k^{exp} should be thought of actually having the power of $\Sigma_{k-1}^{\text{exp}}$.) In particular, $\Sigma_1^p = \text{NP}$, and $\Sigma_1^{\text{exp}} = \text{NEXP}$.

Many of the classes above can be equivalently defined using *alternating Turing machines* (ATM), which is a view we will favour especially when proving upper bounds (Sections 3–4). An ATM is similar to a NTM in that multiple possible transitions may be defined for any given state. However, the definition of acceptance is different. Each non-final state of an ATM is labelled as either *existential* or *universal*, and we define inductively the set of accepting configurations of the machine as the smallest set satisfying the following conditions:

- A configuration in an accepting final state¹ is accepting.
- A configuration in an existential state is accepting if there exists a transition to an accepting configuration.
- A configuration in a universal state is accepting if all transitions lead to accepting configurations.

(States with exactly one possible transition can be thought of as deterministic; it makes no difference whether they are labelled as existential or universal.)

If AP denotes the class of languages computable by an ATM in polynomial time, we have

$$(4) \quad \text{AP} = \text{PSPACE}.$$

(On a related note, APSPACE = EXP, but we will not need this.)

We are especially interested in classes with bounded alternation. Here, an ATM is said to make an *alternation* when it transitions from an existential state to a universal state, or vice versa. For any $k > 0$ and $t: \omega \rightarrow \omega$, $\Sigma_k\text{-TIME}(t)$ denotes the class of languages computable in time $t(n)$ by an ATM that starts in an existential state, and then makes at most $k - 1$ alternations. The class $\Pi_k\text{-TIME}(t)$ is defined similarly, but starting from a universal state. In particular, $\Sigma_1\text{-TIME}(t) = \text{NTIME}(t)$, and $\Pi_1\text{-TIME}(t) = \text{coNTIME}(t)$. We have the following characterization for any $k > 0$:

$$\begin{aligned} \Sigma_k^p &= \Sigma_k\text{-TIME}(\text{poly}(n)), & \Pi_k^p &= \Pi_k\text{-TIME}(\text{poly}(n)), \\ \Sigma_k^{\text{exp}} &= \Sigma_k\text{-TIME}(2^{\text{poly}(n)}), & \Pi_k^{\text{exp}} &= \Pi_k\text{-TIME}(2^{\text{poly}(n)}). \end{aligned}$$

We will also need convenient complete languages for our classes. The set TAUT of tautologies of **CPC** is the canonical coNP-complete language, and the dual language SAT of satisfiable classical propositional formulas is NP-complete. The standard PSPACE-complete language² is QSAT: the language of all true quantified Boolean sentences

$$(5) \quad Q_0 x_0 Q_1 x_1 \dots Q_{n-1} x_{n-1} \varphi(\vec{x}),$$

where φ is a propositional formula, and each $Q_i \in \{\exists, \forall\}$ quantifies over a truth value $x_i \in \mathbf{2}$.

¹In fact, we could dispense with final states altogether: an accepting (rejecting) final state is equivalent to a universal (existential, resp.) state with no possible transitions out.

²It is more transparent to think of it as an AP-complete language: the existential and universal quantifiers almost directly correspond to existential and universal states of an ATM. The usual textbook proof of PSPACE-completeness of QSAT is for the most part actually a proof of the equality (4).

Let $k > 0$. A quantified Boolean sentence (5) is in Σ_k^q if the quantifier prefix $Q_0 \dots Q_{n-1}$ may be written as at most k alternating blocks, the first block consisting of existential quantifiers, the second of universal quantifiers, and so on; Π_k^q is defined dually (i.e., starting with a universal block). Then the language Σ_k^q -SAT consisting of all true Σ_k^q sentences is Σ_k^p -complete, and the dual language Π_k^q -SAT is Π_k^p -complete. Notice that Σ_1^q -SAT is just a notational variant of SAT: a propositional formula is satisfiable iff the corresponding existentially quantified sentence is true.

Finally, we need complete languages for Σ_k^{exp} (in particular, for $k = 2$). Recall that in descriptive complexity, we encode words $w = w_0 \dots w_{n-1} \in \Sigma^*$ by models M_w with domain $n = \{0, \dots, n-1\}$ endowed with the order relation $<$ (and possibly other arithmetical predicates, which we will not need here), and unary predicates P_a for each symbol $a \in \Sigma$, such that

$$M_w \models P_a(i) \quad \text{iff} \quad w_i = a.$$

By Fagin's theorem, a language L is in NP iff there is a Σ_1^1 (i.e., existential second-order) sentence Φ such that

$$(6) \quad w \in L \quad \text{iff} \quad M_w \models \Phi.$$

More generally, Σ_k^p languages are exactly those that are Σ_k^1 -definable. This correspondence can be generalized to the exponential hierarchy, using *third-order* sentences: a language is in Σ_k^{exp} iff it is Σ_k^2 -definable. (See Kołodziejczyk [9, Prop. 2.6], which also includes a brief historical discussion. A similar statement in Hella and Turull-Torres [4, Thm. 7] suffers from an off-by-one error.)

Since we only need complete problems rather than exact descriptions of the languages, we may simplify the Σ_k^2 sentences to a convenient form. This was already done in [5, L. 3.1] for the special case $k = 1$ (i.e., NEXP); here we generalize it to higher levels of the exponential hierarchy (with a more detailed proof).

Theorem 2.3 *Let $k \geq 1$. Put $Q = \exists$ for k odd, and $Q = \forall$ for k even, and let \bar{Q} be its dual. Then the set of true Σ_k^2 sentences of the form*

$$(7) \quad \exists X_1 \subseteq \mathcal{P}(n) \forall X_2 \subseteq \mathcal{P}(n) \dots Q X_k \subseteq \mathcal{P}(n) \bar{Q} t_0, \dots, t_{m-1} \subseteq n \varphi$$

is a Σ_k^{exp} -complete language, where n is given in unary, and φ is a Boolean combination of atomic formulas of the form $i \in t_\alpha$ or $t_\alpha \in X_j$ for $\alpha < m$, $i < n$, and $j = 1, \dots, k$.

Proof: For ease of notation, we will assume that k is odd, so that $Q = \exists$. The case of k even is dual. We will denote third-order variables by capital letters X, Y, \dots (with indices etc.), second-order variables by lower-case letters t, s, \dots , and first-order variables with Greek letters ξ, η, \dots .

First, any Σ_k^{exp} language reduces to a language $L \in \Sigma_k\text{-TIME}(2^n)$ by a simple padding argument: if $L_0 \in \Sigma_k\text{-TIME}(2^{n^c})$, then $L = \{0^{n^c}1w : w \in L_0, |w| = n\}$ is computable in $\Sigma_k\text{-TIME}(2^n)$, and $w \mapsto 0^{n^c}1w$ is a logspace reduction of L_0 to L . Thus, let us fix a Σ_k^{exp} -complete language $L \in \Sigma_k\text{-TIME}(2^n)$.

By Kołodziejczyk [9] (Thm. 5.6 and a comment below Def. 5.9), there is a $s[\Sigma_k^2]^{\leq 1}$ sentence Φ that defines L as in (6): that is, Φ has the form

$$\exists \vec{X}_1 \forall \vec{X}_2 \dots \exists \vec{X}_k \Psi(\vec{X}_1, \dots, \vec{X}_k),$$

where Ψ is a second-order formula, and all the second-order and third-order variables are unary: i.e., for a model with domain n , the second-order variables range over $\mathcal{P}(n)$, and third-order variables over $\mathcal{P}(\mathcal{P}(n))$. (Second-order variables will remain unary for the rest of the proof, but we will introduce third-order variables of higher arity during subsequent manipulations.) We may assume Ψ only uses $=$ for first-order objects.

We may write Ψ in prenex normal form, and moreover, we may assume that all second-order quantifiers precede all first-order quantifiers: this is easily accomplished by exploiting the equivalences

$$\begin{aligned} \exists \xi Q_0 t_0 \dots Q_l t_l \psi &\Leftrightarrow \exists s (s \neq \emptyset \wedge Q_0 t_0 \dots Q_l t_l \forall \xi (s(\xi) \rightarrow \psi)) \\ &\Leftrightarrow \exists s Q_0 t_0 \dots Q_l t_l (s \neq \emptyset \wedge \forall \xi (s(\xi) \rightarrow \psi)), \end{aligned}$$

where $s \neq \emptyset$ stands for $\exists \xi s(\xi)$; similarly for universal quantifiers. Thus, we may write

$$\Phi = \exists \vec{X}_1 \forall \vec{X}_2 \dots \exists \vec{X}_k \forall t_0 \exists s_1 \forall t_1 \dots \forall t_{l-1} \exists s_l \psi(\vec{X}_1, \dots, \vec{X}_k, \vec{t}, \vec{s}),$$

where ψ is first-order. Next, we get rid of existential second-order quantifiers by introducing Skolem functions: Φ is equivalent to

$$\begin{aligned} \exists \vec{X}_1 \forall \vec{X}_2 \dots \exists \vec{X}_k \exists Y_1 \dots \exists Y_l \forall t_0 \dots \forall t_{l-1} \\ \psi(\vec{X}_1, \dots, \vec{X}_k, \vec{t}, \{\xi : Y_1(t_0, \xi)\}, \dots, \{\xi : Y_l(t_0, \dots, t_{l-1}, \xi)\}), \end{aligned}$$

where each Y_i is a third-order relation variable with i second-order and one first-order argument (i.e., it ranges over $\mathcal{P}(\mathcal{P}(n)^i \times n)$.) Eliminating the comprehension symbols, Φ is equivalent to

$$\exists \vec{X}_1 \forall \vec{X}_2 \dots \exists \vec{X}_k \exists Y_1 \dots \exists Y_l \forall t_0 \dots \forall t_{2l-1} (\gamma(\vec{Y}, \vec{t}) \rightarrow \psi(\vec{X}_1, \dots, \vec{X}_k, \vec{t})),$$

where γ denotes the first-order formula

$$(8) \quad \bigwedge_{i < l} \forall \xi (t_{l+i}(\xi) \leftrightarrow Y_{i+1}(t_0, \dots, t_i, \xi)).$$

By increasing l or adding dummy quantifiers if necessary, we may assume that all the tuples \vec{X}_j also have length l .

In order to simplify each third-order quantifier block to a single unary variable, we pass from M_w to a larger model M_w^* , defined as follows. The domain of M_w^* is $n^* = n \times (l+1)^2$ (i.e., $(l+1)^2$ disjoint copies of the domain of M_w). For each $i < (l+1)^2$, M_w^* includes a unary predicate C_i that selects the i th copy of M_w (i.e., C_i is satisfied by elements of $n \times \{i\}$), and a binary predicate E that defines the equivalence relation

$$\{\langle \langle \xi, i \rangle, \langle \xi, j \rangle \rangle : \xi < n \text{ and } i, j < (l+1)^2\}$$

(i.e., the kernel of the projection $M_w^* \rightarrow M_w$). The original relations of M_w are included on the 0th copy $n \times \{0\}$. Clearly, the mapping $w \mapsto M_w^*$ is still logspace computable. We will translate Φ to a formula Φ^* such that

$$(9) \quad w \in L \quad \text{iff} \quad M_w \models \Phi \quad \text{iff} \quad M_w^* \models \Phi^*.$$

We may represent elements $\xi \in M_w$ by elements of M_w^* that satisfy C_0 , and subsets $t \subseteq M_w$ by subsets of C_0 ; however, we will actually need to quantify over the copies of t in each C_i as well. We represent an l -tuple $\vec{X} = \langle X_i : i < l \rangle$ of third-order objects $X_i \subseteq \mathcal{P}(n)$ by a single third-order object $X^* \subseteq \mathcal{P}(n^*)$ defined as

$$X^* = \bigcup_{i < l} \{t \times \{i\} : t \in X_i\}.$$

Moreover, if we are in addition to \vec{X} given a tuple Y_1, \dots, Y_l , where $Y_i \subseteq \mathcal{P}(n)^i \times n$, we represent \vec{X}, \vec{Y} together by

$$\bigcup_{i < l} \{t \times \{i\} : t \in X_i\} \cup \bigcup_{i=1}^l \left\{ \bigcup_{j < i} (t_j \times \{i(l+1) + j\}) \cup (\{\xi\} \times \{i(l+1) + l\}) : \langle t_0, \dots, t_{i-1}, \xi \rangle \in Y_i \right\}.$$

With this representation in mind, Φ^* will have the form

$$\exists X_1^* \forall X_2^* \dots \exists X_k^* \forall \vec{t}_0 \dots \forall \vec{t}_{2l-1} (\gamma^*(X_k^*, \vec{t}_0, \dots, \vec{t}_{2l-1}) \rightarrow \psi^*(X_1^*, \dots, X_k^*, \vec{t}_0, \dots, \vec{t}_{2l-1})),$$

where each \vec{t}_i is an $(l+1)^2$ -tuple $\langle t_{i,j} : j < (l+1)^2 \rangle$, the formula γ^* is a conjunction of (10) and (11) below, and ψ^* is constructed from ψ as follows. We replace first-order quantifiers $\exists \xi \dots$ and $\forall \xi \dots$ by $\exists \xi (C_0(\xi) \wedge \dots)$ and $\forall \xi (C_0(\xi) \rightarrow \dots)$, respectively. Atomic subformulas of ψ that only mention first-order objects are left intact, atomic formulas $t_i(\xi)$ are replaced with $t_{i,0}(\xi)$, and atomic formulas $X_{j,p}(t_i)$ are replaced with $X_j^*(t_{i,p})$.

The first conjunct of γ^* ,

$$(10) \quad \bigwedge_{\substack{i < 2l \\ j, j' < (l+1)^2}} \forall \xi (t_{i,j'}(\xi) \leftrightarrow C_{j'}(\xi) \wedge \exists \eta (C_j(\eta) \wedge E(\xi, \eta) \wedge t_{i,j}(\eta))),$$

ensures that the sets \vec{t}_i are correctly formed: i.e., $t_{i,j} \subseteq C_j$, and for fixed i , all the $t_{i,j}$ are copies of the same set $t_i \subseteq n$. The second conjunct is a translation of (8), which can be written as

$$\bigwedge_{i=1}^l \forall \xi \left(t_{l+i-1, i(l+1)+l}(\xi) \leftrightarrow X_k^* \left(\bigcup_{j < i} t_{j, i(l+1)+j} \cup \{\xi\} \right) \right).$$

Expanding the definition of \cup , we arrive at

$$(11) \quad \bigwedge_{i=1}^l \forall \xi \exists s \left(\forall \eta \left(s(\eta) \leftrightarrow \eta = \xi \vee \bigvee_{j < i} t_{j, i(l+1)+j}(\eta) \right) \wedge \left(t_{j+i-1, i(l+1)+l}(\xi) \leftrightarrow X_k^*(s) \right) \right).$$

By construction, (9) holds. Notice that γ^* , specifically (11), contains a second-order quantifier. When we bring Φ^* to prenex normal form, we obtain a sentence of the form (dropping the $*$ decoration from variables)

$$\exists X_1 \forall X_2 \dots \exists X_k \forall \vec{t} \exists \xi \forall \vec{s} \varphi(X_1, \dots, X_k, \vec{t}, \vec{s}, \xi),$$

where φ is first-order. (This is still a single, constant-size sentence that only depends on L , not on w .)

As a final step, we transform Φ^* for any given word w to a sentence Φ_w that embeds the structure of M_w^* as follows. Using constants for elements $i \in M_w^*$ (i.e., $i < n^*$), we expand each first-order quantifier $\exists \xi \dots$ to a disjunction $\bigvee_{i < n^*} \dots$, and $\forall \xi \dots$ to $\bigwedge_{i < n^*} \dots$. Then we evaluate each atomic formula that does not involve higher-order variables, and replace it with \top or \perp according to its truth-value.

The resulting formula has size polynomial in n (the exponent being roughly the nesting depth of first-order quantifiers in Φ^*), and it is easy to see that it is logspace computable. It has the form (7), and we have

$$w \in L \quad \text{iff} \quad M_w^* \models \Phi^* \quad \text{iff} \quad n^* \models \Phi_w$$

by construction. □

3 Derivability

Before we embark on our main quest for the complexity of admissibility and unifiability in clx logics, let us first settle a more basic question: what is the complexity of *tautology* or *derivability* in these logics.

Remark 3.1 Notice that (single-conclusion) derivability has the same complexity as tautology for any transitive logic L , as

$$\Gamma \vdash_L \varphi \quad \text{iff} \quad \vdash_L \bigwedge \Gamma \rightarrow \varphi$$

gives a logspace reduction.

Recall that any clx logic L is $\forall\exists$ -definable on finite frames [8, Thm. 4.29], and as a consequence, finite L -frames are recognizable in polynomial time (in fact, in $\Pi_2\text{-TIME}(O(\log n))$, a subclass of uniform AC^0).

Theorem 3.2 *For any consistent linear clx logic L , derivability in L is coNP-complete.*

Proof: Since L is a conservative extension of classical propositional logic, \vdash_L is coNP-hard.

On the other hand, if φ is a formula of size $n = |\varphi|$ such that $\not\vdash_L \varphi$, then there exists an L -model $F \not\models \varphi$ of depth and cluster size at most n , hence size at most n^2 , by [8, Thm. 4.38]. (In fact, it is easy to show that size n is enough.) Since finite L -frames are polynomial-time recognizable by Remark 3.1, this shows that $\not\vdash_L$ is in NP. □

The complexity of nonlinear clx logics is a bit more difficult to establish (although it is just a variant of standard PSPACE-completeness results for modal logics starting with Ladner [10]). For this reason, we state the upper and lower bounds separately. We begin with the former.

Recall from [8, Def. 4.20] that $\text{bas}(L)$ is a finite set of extension conditions that determines the shape of L -frames.

Theorem 3.3 *For any clx logic L , L -derivability is in PSPACE.*

Proof: Recall from [8, Thm. 4.38] that if $\not\vdash_L \varphi$, then φ is falsified in a rooted L -model which is a tree of depth $\leq n$ of clusters of size $\leq n$. Similarly to the standard case of **K4** and other common logics, we may search for such a tree in polynomial space—or as we prefer to think about it, in alternating polynomial time—by exploring one branch at a time.

In more detail, let φ be a formula whose provability in L we want to determine. Since variables and parameters work the same way with respect to derivability, we may assume φ contains no parameters. Put $\Sigma = \text{Sub}(\Box\varphi)$, $B = \{\psi : \Box\psi \in \Sigma\}$, $V = \Sigma \cap \text{Var}$, and $n = |\Sigma|$. For any $v \in \mathbf{2}^V$ and $X \subseteq B$, let $A(v, X)$ denote the Boolean assignment to modal formulas that agrees with v on variables, and that makes $\Box\psi$ true for $\psi \in X$, and false for $\psi \in B \setminus X$. If W is an L -model, let $B^+(W) = \{\psi \in B : W \vDash \psi\}$.

We intend to construct an algorithm to compute the predicate

$$S(X) \quad \text{iff} \quad \exists W \in \text{Mod}_L \ B^+(W) = X$$

for $X \subseteq B$. (I.e., S determines what assignments to the set $\{\Box\psi : \psi \in B\}$ are consistent with L ; recall that Mod_L denotes the class of rooted finite L -models.) We may then express (un)provability of φ as

$$(12) \quad \not\vdash_L \varphi \quad \text{iff} \quad \exists X \subseteq B \setminus \{\varphi\} \ S(X).$$

We will present a recursive description of $S(X)$ expressing how to build a model witnessing $S(X)$ by attaching a new root cluster to a disjoint sum of rooted models witnessing $S(Z)$ for some sets $Z \supseteq X$.

To this end, we consider auxiliary predicates $H_C(X, Y)$ and $U_m(X, Y)$ for $X \subseteq Y \subseteq B$, where $\langle C, m \rangle \in EC^\infty$. Their intended meaning is that $H_C(X, Y)$ holds if, given a (not necessarily rooted) model W' such that $B^+(W') = Y$, we can construct a model W such that $B^+(W) = X$ by attaching to W' a root cluster of type C with a suitable valuation; $U_m(X, Y)$ holds if there are rooted models $\{W_i : i < m\}$ such that $B^+(\sum_{i < m} W_i) = Y$ and $B^+(W_i) \supseteq X$ for each i . We interpret this for infinite extension conditions such that if $C = \otimes$, the root cluster can be reflexive of arbitrary size, and if $m = \infty$, any positive number of W_i is allowed.

Formally, we define H_C and U_m as follows:

$$\begin{aligned} H_\bullet(X, Y) \quad \text{iff} \quad & \exists v \in \mathbf{2}^V \ A(v, Y) \vDash \bigwedge_{\psi \in X} \psi \wedge \bigwedge_{\chi \in Y \setminus X} \neg\chi, \\ H_{\otimes}(X, Y) \quad \text{iff} \quad & \exists \{v_i : i < k\} \subseteq \mathbf{2}^V \left(\forall i < k \ A(v_i, X) \vDash \bigwedge_{\psi \in X} \psi \right. \\ & \left. \& \forall \chi \in Y \setminus X \ \exists i < k \ A(v_i, X) \vDash \neg\chi \right), \end{aligned}$$

$$\begin{aligned}
H_{\otimes}(X, Y) &\text{ iff } \forall \chi \in (Y \setminus X) \cup \{\perp\} \exists v \in \mathbf{2}^V A(v, X) \models \bigwedge_{\psi \in X} \psi \wedge \neg \chi, \\
U_m(X, Y) &\text{ iff } \exists \{Z_j : j < m\} \left(\forall j < m (X \subsetneq Z_j \ \& \ S(Z_j) \ \& \ Y \subseteq Z_j \subseteq B) \right. \\
&\quad \left. \& \ \forall \chi \in B \setminus Y \exists j < m \ \chi \notin Z_j \right), \\
U_{\infty}(X, Y) &\text{ iff } \begin{cases} \forall \chi \in B \setminus Y \exists Z (X \subsetneq Z \ \& \ S(Z) \ \& \ Y \subseteq Z \subseteq B \setminus \{\chi\}), & Y \subsetneq B, \\ X \subsetneq B \ \& \ S(B), & Y = B. \end{cases}
\end{aligned}$$

for $1 \leq k < \omega$ and $m < \omega$.

Claim 3.3.1 *For any $X \subseteq B$,*

$$(13) \quad S(X) \text{ iff } \exists \langle C, m \rangle \in \text{bas}(L) \exists Y (X \subseteq Y \subseteq B \ \& \ H_C(X, Y) \ \& \ U_m(X, Y)).$$

Proof: Left-to-right: Fix a finite rooted L -model W such that $B^+(W) = X$, and such that $B^+(w\uparrow) \supseteq X$ for all $w \in W \setminus \text{rcl}(W)$. Assume that W is of type t , and fix $\langle C, m \rangle \in \text{bas}(L)$ such that $t \preceq \langle C, m \rangle$. Put $Y = B^+(W \setminus \text{rcl}(W))$.

We claim that $H_C(X, Y)$: if $C = \bullet$, let $v \in \mathbf{2}^V$ be the assignment of variables in the (unique) root r , i.e., $v = \text{Sat}_V(r)$. Since $Y = \{\psi \in B : r \models \Box\psi\}$, we have $\text{Sat}_{\Sigma}(r) = A(v, Y)$, and it is readily seen that

$$r \models \bigwedge_{\psi \in X} \psi \wedge \bigwedge_{\chi \in Y \setminus X} \neg \chi,$$

thus v witnesses that $H_{\bullet}(X, Y)$ holds. If C is reflexive, then $X = \{\psi \in B : r \models \Box\psi\}$, thus for any $w \in \text{rcl}(W)$, $\text{Sat}_{\Sigma}(w) = A(v, X)$, where $v = \text{Sat}_V(w)$. If $C = \textcircled{k}$, fix a (not necessarily injective) enumeration $\text{rcl}(W) = \{w_i : i < k\}$; then $H_{\textcircled{k}}(X, Y)$ is witnessed by $v_i = \text{Sat}_V(w_i)$. Likewise, if $C = \textcircled{\infty}$, we may take $v = \text{Sat}_V(w)$ for any $w \in \text{rcl}(W)$ such that $w \models \neg \chi$.

We also claim that $U_m(X, Y)$. If $m < \infty$, let $\{C_j : j < m\}$ be a (not necessarily injective) enumeration of immediate successor clusters of $\text{rcl}(W)$, and put $Z_j = B^+(C_j\uparrow)$. Then $S(Z_j)$ by definition, and $Y = \bigcap_{j < m} Z_j$, which implies that $Y \subseteq Z_j \subseteq B$, and for each $\chi \in B \setminus Y$, $\chi \notin Z_j$ for some j . Moreover, our assumptions on W ensure that $X \subsetneq Z_j$. Thus, $\{Z_j : j < m\}$ witness that $U_m(X, Y)$ holds. If $m = \infty$, then for any $\chi \in B \setminus Y$, there exists a point $x \in W \setminus \text{rcl}(W)$ such that $x \not\models \Box\chi$; then $Z = B^+(x\uparrow)$ has the required properties. If $Y = B$, we just do the same for an arbitrary $x \in W \setminus \text{rcl}(W)$.

Right-to-left: Fix $\langle C, m \rangle \in \text{bas}(L)$ and $X \subseteq Y \subseteq B$ such that $H_C(X, Y)$ and $U_m(X, Y)$. We will construct a finite L -model W whose root witnesses $S(X)$. First, we build the root cluster. If $C = \bullet$, $\text{rcl}(W)$ will be an irreflexive point with valuation of variables v chosen as a witness for $H_{\bullet}(X, Y)$. Likewise, if $C = \textcircled{k}$, let $\{v_i : i < k\}$ be witnesses for $H_{\textcircled{k}}(X, Y)$; then $\text{rcl}(W)$ is a k -element reflexive cluster $\{w_i : i < k\}$ with $\text{Sat}_V(w_i) = v_i$. If $C = \textcircled{\infty}$, then for each $\chi \in Y \setminus X$, fix $v_{\chi} \in \mathbf{2}^V$ such that $A(v_{\chi}, X) \models \bigwedge_X \psi \wedge \neg \chi$; then $\text{rcl}(W)$ is a reflexive cluster consisting of points $\{w_{\chi} : \chi \in Y \setminus X\}$ such that $\text{Sat}_V(w_{\chi}) = v_{\chi}$. This does not work if $X = Y$; in this case, we let $\text{rcl}(W)$ be a reflexive singleton satisfying v where $A(v, X) \models \bigwedge_{\psi \in X} \psi$.

Next, we construct the rest of the model. If $m < \infty$, let $\{Z_j : j < m\}$ be witnesses for $U_m(X, Y)$. Since $S(Z_j)$ for each $j < m$, we may fix a finite rooted L -model W_j such that

$B^+(W_j) = Z_j$. We let $W \setminus \text{rcl}(W) = \sum_{j < m} W_j$. Similarly, if $m = \infty$, then for each $\chi \in B \setminus Y$, we fix Z_χ satisfying the condition from the definition of U_\otimes , and we find a rooted L -model W_χ such that $B^+(W_\chi) = Z_\chi$. We define $W \setminus \text{rcl}(W)$ as the disjoint union of W_χ for all $\chi \in B \setminus Y$, as long as $Y \subsetneq B$. If $Y = B$, let $W \setminus \text{rcl}(W)$ be a rooted L -model satisfying $\bigwedge_{\psi \in B} \Box\psi$.

Since $Y = \bigcap_{j < m} Z_j$ or $Y = \bigcap_{\chi \in B \setminus Y} Z_\chi$ (respectively), we see that

$$W \setminus \text{rcl}(W) \models \bigwedge_{\psi \in Y} \Box\psi,$$

and

$$\text{rcl}(W) \models \bigwedge_{\chi \in B \setminus Y} \neg\Box\chi.$$

Then it is easy to check that $\text{rcl}(W)$ witnesses $S(X)$. □ (Claim 3.3.1)

Notice that if we expand the occurrences of H_C and U_m in (13) using their definitions, we obtain a recursive expression for $S(X)$ in terms of $S(Z)$ or $S(Z_j)$, where $X \subsetneq Z, Z_j$. We may read it as a recipe for an algorithm computing $S(X)$ (and $\vdash_L \varphi$, in view of (12)) on an alternating Turing machine, where we interpret the existential and universal quantifiers (and disjunctions and conjunctions) in the expression as nondeterministic choices in existential or universal states (respectively). As just noted, the argument strictly increases on each recursive call to $S(X)$, hence the recursion depth is at most n . Moreover, each quantifier takes $O(n)$ bits, and the conditions in-between (e.g., the truth of some formulas from Σ under $A(v, X)$) can be checked deterministically in polynomial time. Thus, the algorithm works in alternating polynomial time, i.e., in PSPACE by (4). □

Remark 3.4 The bounds in Theorems 3.2 and 3.3 apply not just to derivability of single-conclusion rules as in Remark 3.1, but also to derivability of multiple-conclusion rules: its definition in terms of single-conclusion derivability

$$\Gamma \vdash_L^m \Delta \quad \text{iff} \quad \exists \varphi \in \Delta \Gamma \vdash_L^1 \varphi$$

shows that \vdash_L^m is in coNP or PSPACE whenever \vdash_L^1 is. (In general, the definition gives a logspace disjunctive truth-table (dtt) reduction of \vdash_L^m to \vdash_L^1 ; both PSPACE and coNP are closed under dtt-reductions.)

We now turn to the lower bound. There is more than one way to approach it. One possibility is to extend Ladner's proof to show that all normal modal logics with the disjunction property are PSPACE-hard. (Even though this is straightforward to carry out, and seems to be a fundamental result, the author is not aware of any published reference. The corresponding result for si logics was proved by Chagrov [2], cf. [3, Thm. 18.30].) We will use another method, namely by reduction from intuitionistic logic (**IPC**) using a series of translations. This route is more useful for our purposes, because the resulting statement is relatively more general in the context of transitive modal logics (it applies to all transitive logics with the disjunction property, and it also applies to various extensions of **K4.2**, which will be relevant in the sequel).

Definition 3.5 Let T denote the *Gödel–McKinsey–Tarski translation* of **IPC** (formulated using connectives $\{\rightarrow, \wedge, \vee, \perp\}$) in **S4**: $\mathsf{T}(\varphi) = \Box\varphi$ if φ is an atom, T commutes with \wedge , \vee , and \perp , and

$$\mathsf{T}(\varphi \rightarrow \psi) = \Box(\mathsf{T}(\varphi) \rightarrow \mathsf{T}(\psi)).$$

The *relativization translation* R of **K4** in itself is defined as follows. Given a formula φ , we fix an atom r (a parameter, in contexts where the distinction between variables and parameters becomes relevant) that does not occur in φ . We define an auxiliary translation ψ^r such that it preserves atoms, commutes with Boolean connectives, and

$$(\Box\psi)^r = \Box(r \rightarrow \psi^r).$$

Then we define $\mathsf{R}(\varphi)$ as $r \rightarrow \varphi^r$.

Let B denote the *boxdot translation* of **S4** in **K4**, which preserves atoms, commutes with Boolean connectives, and satisfies

$$\mathsf{B}(\Box\varphi) = \Box\mathsf{B}(\varphi).$$

Notice that when expanded out, this formula contains two occurrences of $\mathsf{B}(\varphi)$, and as a consequence, the boxdot translation may exponentially enlarge formulas with deeply nested boxes. For this reason, we define the *efficient boxdot translation* $\mathsf{B}'(\varphi)$ as follows. We introduce a new variable z_ψ for each formula ψ such that $\Box\psi \subseteq \varphi$. For any $\chi \subseteq \varphi$, let χ' denote the formula resulting from χ by replacing all topmost occurrences of boxed subformulas $\Box\psi$ with the corresponding z_ψ . Then

$$\mathsf{B}'(\varphi) := \bigwedge_{\Box\psi \subseteq \varphi} \Box(z_\psi \leftrightarrow \Box\psi') \rightarrow \varphi'.$$

Lemma 3.6 *For any formula φ , the formulas $\mathsf{B}(\varphi)$ and $\mathsf{B}'(\varphi)$ are equiderivable. Specifically, we have*

$$(14) \quad \vdash_{\mathbf{K4}} \mathsf{B}(\varphi) \rightarrow \mathsf{B}'(\varphi),$$

$$(15) \quad \vdash_{\mathbf{K4}} \sigma(\mathsf{B}'(\varphi)) \rightarrow \mathsf{B}(\varphi),$$

where σ is the substitution such that $\sigma(z_\psi) = \mathsf{B}(\Box\psi)$.

Proof:

(14): We can prove

$$\vdash_{\mathbf{K4}} \bigwedge_{\Box\psi \subseteq \varphi} \Box(z_\psi \leftrightarrow \Box\psi') \rightarrow (\chi' \leftrightarrow \mathsf{B}(\chi))$$

for each $\chi \subseteq \varphi$ by induction on the complexity of χ .

(15): Since B and σ both commute with Boolean connectives, we have $\sigma(\chi') = \mathsf{B}(\chi)$ for each $\chi \subseteq \varphi$, thus $\sigma(\mathsf{B}'(\varphi))$ is

$$\bigwedge_{\Box\psi \subseteq \varphi} \Box(\Box\mathsf{B}(\psi) \leftrightarrow \Box\mathsf{B}(\psi)) \rightarrow \mathsf{B}(\varphi),$$

which is equivalent to just $\mathsf{B}(\varphi)$. □

Recall that for any frame W , W_\circ denotes its reflexivization. Following [3], if W is reflexive, let ϱW denote the *skeleton* of W , i.e., the quotient of W by the cluster equivalence relation \sim , considered as an intuitionistic frame. The following is standard (see e.g. [3, L. 8.28]):

Lemma 3.7 *For any reflexive frame W , and an intuitionistic formula φ , we have*

$$\varrho W \vDash \varphi \quad \text{iff} \quad W \vDash \mathsf{T}(\varphi). \quad \square$$

We leave the straightforward proof of the corresponding property of the boxdot translation to the reader:

Lemma 3.8 *For any frame W , and a formula φ , we have*

$$W_\circ \vDash \varphi \quad \text{iff} \quad W \vDash \mathsf{B}(\varphi). \quad \square$$

The $\mathsf{R}(\varphi)$ translation is equivalent to φ^{sf} from Chagrov and Zakharyashev [3, Exer. 9.7].

Lemma 3.9 *Let φ be a formula, r an atom, W a frame, W_0 its subframe, and \vDash a valuation in W such that*

$$W, w \vDash r \quad \text{iff} \quad w \in W_0.$$

Then

$$W, w \vDash \varphi^r \quad \text{iff} \quad W_0, w \vDash \varphi$$

for all $w \in W_0$.

Proof: By induction on the complexity of φ . □

Corollary 3.10 *A frame W validates $\mathsf{R}(\varphi)$ iff φ holds in all subframes of W (and thus in all frames that W subreduces onto).*

Consequently, if L is a subframe logic and $\vdash_L \varphi$, then $\vdash_L \mathsf{R}(\varphi)$. □

Definition 3.11 A logic $L \supseteq \mathbf{K4}$ is (cofinally) *subframe-universal for trees* if for each finite tree T , considered as a reflexive frame, there exists an L -frame W that (cofinally, resp.) weakly subreduces onto T .

Notice that in order to verify subframe universality for trees, it is enough to check *binary* trees; on the other hand, if a logic is subframe-universal for trees, it is also subframe-universal for all finite reflexive rooted frames without proper clusters. This follows from the fact that every such frame is a p-morphic image of a binary tree. The same considerations apply, mutatis mutandis, to cofinal subframe universality.

Theorem 3.12 *Let a logic $L \supseteq \mathbf{K4}$ be subframe-universal for trees. Then L -derivability is PSPACE-hard.*

Proof: Derivability in **IPC** is PSPACE-complete by Statman [11], hence it suffices to provide a reduction of **IPC** to L . Now, the translations T , R , and B' increase the size of the formula only linearly, and are readily seen to be logspace-computable, thus we only need to prove

$$\vdash_{\mathbf{IPC}} \varphi \quad \text{iff} \quad \vdash_L \mathsf{B}'(\mathsf{R}(\mathsf{T}(\varphi))).$$

By Lemma 3.6, this is equivalent to

$$\vdash_{\mathbf{IPC}} \varphi \quad \text{iff} \quad \vdash_L \mathsf{B}(\mathsf{R}(\mathsf{T}(\varphi))).$$

For the left-to-right implication, it is well known that T is an interpretation of **IPC** in **S4**, and B is an interpretation of **S4** in **K4**. By Corollary 3.10, R is a self-interpretation of **S4**.

For the right-to-left implication, assume that $\not\vdash_{\mathbf{IPC}} \varphi$, whence there exists a finite tree T such that φ is not valid in ϱT (the intuitionistic version of T). By assumption, there exists a weak subreduction from an L -frame W onto T , hence a subreduction from the reflexivization W_\circ to T by Observation 2.1. Then

$$\varrho T \not\vdash \varphi \implies T \not\vdash \mathsf{T}(\varphi) \implies W_\circ \not\vdash \mathsf{R}(\mathsf{T}(\varphi)) \implies W \not\vdash \mathsf{B}(\mathsf{R}(\mathsf{T}(\varphi)))$$

by Lemma 3.7, Corollary 3.10, and Lemma 3.8. □

The previous theorem applies in particular to all logics with the disjunction property, as they are subframe universal for trees. We will even prove that logics with the disjunction property are *cofinally* subframe-universal for trees, which fact we will need later. We note that the special case for logics $L \supseteq \mathbf{S4}$ follows from [3, Prop. 15.13]; indeed, we could alternatively prove the general result by reducing it to the reflexive case, but the proof below is more elementary (it avoids the machinery of canonical formulas).

Theorem 3.13 *Every logic $L \supseteq \mathbf{K4}$ satisfying the disjunction property is cofinally subframe-universal for trees. Consequently, L -derivability is PSPACE-hard.*

Proof: We will show by induction on the depth of T that for every finite tree T , there exists a descriptive L -frame W cofinally weakly subreducing to T . For the base case, any nonempty L -frame (which exists as L is consistent) weakly subreduces to the one-element tree.

For the induction step, let r be the root of T , let r_i , $i < n$, be the immediate successors of r , and for each $i < n$, let T_i be the subtree of T rooted at r_i . We may assume $n \geq 2$ by duplicating T_0 if necessary. By the induction hypothesis, we may fix a descriptive L -frame W_i , and a cofinal weak subreduction f_i from W_i onto T_i for each $i < n$. We may assume that W_i is rooted, and its root w_i is mapped to r_i by f_i .

By [3, Thm. 15.1], there exists a rooted descriptive L -frame $\langle W, <, A \rangle$ such that the disjoint sum $\sum_{i < n} W_i$ is a generated subframe of W . Notice that f_i is not necessarily a weak subreduction from W to T_i , as W_i may not be an admissible subset of W : it is in general only an intersection of admissible subsets. Nevertheless, assume for now that $\bigcup_{i < n} f_i$ extends to a cofinal weak subreduction g from W to $\bigcup_{i < n} T_i \subseteq T$. We may further extend it to a subreduction $h \supseteq g$ from W onto T by putting

$$h(v) = r \quad \text{iff} \quad v \in \bigcap_{i < n} g^{-1}[r_i] \downarrow.$$

(In particular, the root of W is mapped to r .) The subreduction h is not necessarily cofinal. However, if $v \in W$ sees no point of $\text{dom}(h)$ (i.e., violates cofinality), then points below v cannot be mapped by h to anything else than r , as g is cofinal. Thus, if we fix an arbitrary leaf node $t \in T$, the partial mapping $f: W \rightarrow T$ that extends h by putting

$$f(v) = t \quad \text{if } v \notin \text{dom}(h) \Downarrow$$

is still a weak subreduction from W onto T , and it is clearly cofinal.

Now, it remains to construct g . Let us first fix $i < n$ and $u \in W_i$. Since W is refined, we may find for each $j < n$, $j \neq i$, a set $X_{u,j} \in A$ such that $w_j \in \Box X_{u,j}$, and $u \notin X_{u,j}$. Putting $X_u = W \setminus \bigcup_{j \neq i} X_{u,j}$, we have $u \in X_u \in A$, while $X_u \Downarrow \cap \bigcup_{j \neq i} W_j = \emptyset$. The sets $X_u \cap W_i$ are admissible in W_i ; using the compactness of W_i , there exists a finite set $\{u_k : k < m\} \subseteq W_i$ such that W_i is covered by $Y_i = \bigcup_{k < m} X_{u_k} \in A$.

Let \vDash be an admissible valuation on W such that

$$v \vDash y_i \quad \text{iff } v \in Y_i$$

for each $i < n$ and $v \in W$, and

$$v \vDash x_t \quad \text{iff } f_i(v) = t$$

for each $i < n$, $t \in T_i$, and $v \in W_i$ (the valuation of x_t may be arbitrary outside W_i , which allows \vDash to be admissible in W). We define a partial function $g: W \rightarrow T$ such that for any $i < n$, $t \in T_i$, and $v \in W$,

$$\begin{aligned} g(v) = t \quad \text{iff} \quad & v \vDash x_t \wedge \Box y_i \wedge \Box \bigvee_{s \in T_i} \Diamond x_s \wedge \bigwedge_{\substack{s, s' \in T_i \\ s \neq s'}} \Box \neg(x_s \wedge x_{s'}) \\ & \wedge \bigwedge_{\substack{s, s' \in T_i \\ s \not\leq s'}} \Box (x_s \rightarrow \Box \neg x_{s'}) \wedge \bigwedge_{\substack{s, s' \in T_i \\ s \leq s'}} \Box (x_s \rightarrow \Diamond x_{s'}). \end{aligned}$$

The $\Box y_i$ conjunct ensures that the purported preimages of distinct points are disjoint, hence g is indeed a well-defined partial mapping. It is then easy to check that $g \supseteq f_i$ for each $i < n$, and that g is a cofinal weak subreduction from W onto $\bigcup_{i < n} T_i$. \square

4 Upper bounds

4.1 Nonlinear cluster-extensible logics

In order to prove our basic estimate of the complexity of admissibility in clx logics, we will use the description of admissibility in terms of pseudoextensible models from [8, Thm. 5.24], thus we need to know what it takes to check pseudoextensibility.

Lemma 4.1 *Let us fix an extension condition $\langle C, m \rangle \in EC^\infty$.*

Given a finite set of formulas Σ which is closed under subformulas, and a finite model F , we can check if F is $\langle C, m \rangle$ -pseudoextensible wrt Σ in $\text{coNTIME}((N + 2^{n^2})^{O(1)})$, where $n = |\Sigma|$, and $N = |F|$.

If moreover C is finite, we can even check it in $\text{DTIME}((N + 2^n)^{O(1)})$.

Proof: First, observe that we can compute the valuation of all formulas from Σ in F in about $O(nN^2)$ steps.

Let $P = \text{Par} \cap \Sigma$, $B = \{\varphi : \Box\varphi \in \Sigma\}$, and $* = \text{refl}(C)$. (Recall that $* \in \{\bullet, \circ\}$ denotes the reflexivity of C .)

Given $X \subseteq F$, $\emptyset \neq E \subseteq \mathbf{2}^P$ (with $|E| = 1$ if $* = \bullet$), and $u = \{u_e : e \in E\} \subseteq F$, we can test in $\text{DTIME}((N + 2^n)^{O(1)})$ if u is a tight $\langle *, E \rangle$ -pseudopredecessor of X wrt Σ just using the definition: this amounts to checking $u_e \models P^e$ for all $e \in E$, and for each $\varphi \in B$, to look at the valuation of φ and $\Box\varphi$ in every u_e and $w \in X$.

Better yet, given E and X , we can check the *existence* of a $\langle *, E \rangle$ -tpp of X wrt Σ in $\text{DTIME}((N + 2^{n^2})^{O(1)})$: we compute $B^+ = \bigcap_{w \in X} \{\varphi \in B : w \models \Box\varphi\}$, and $B^- = B \setminus B^+$. If $* = \bullet$, we try every $u \in F$ to see if

$$u \models P^e \wedge \bigwedge_{\varphi \in B^+} \Box\varphi \wedge \bigwedge_{\varphi \in B^-} \neg\Box\varphi.$$

If $* = \circ$, we try every $D \subseteq B^+$ and $f: D \rightarrow E$ to see if any satisfy the condition

$$\forall e \in E \exists u \in F u \models P^e \wedge \bigwedge_{\varphi \in B^+ \setminus D} \Box\varphi \wedge \bigwedge_{\varphi \in B^- \cup D} \neg\Box\varphi \wedge \bigwedge_{\substack{\varphi \in D \\ f(\varphi)=e}} \neg\varphi$$

(cf. [8, Def. 5.21 and L. 5.22]). This condition in turn can be checked by trying all possible e and u . Notice that there are at most 2^n choices for D , and at most $|E|^n \leq 2^{n^2}$ choices for f ; if $|E|$ is bounded by a constant, the latter bound is $2^{O(n)}$.

If m is finite, then F is $\langle C, m \rangle$ -pseudoextensible wrt Σ iff for all $E \subseteq \mathbf{2}^P$ of size $|E| \leq_0 |C|$, and for all $X = \{w_i : i < m\} \subseteq F$, there exists a $\langle *, E \rangle$ -tpp of X wrt Σ . We check this by co-nondeterministically choosing E , and then deterministically trying all N^m possibilities for X , resulting in a $\text{coNTIME}((N + 2^{n^2})^{O(1)})$ algorithm. If moreover C is finite with $|C| = k$, we only need to check $E \subseteq \mathbf{2}^P$ of size $|E| \leq k$; we may do this deterministically by trying all 2^{kn} possibilities, resulting in a $\text{DTIME}((N + 2^n)^{O(1)})$ algorithm.

If $m = \infty$, we could choose $X \subseteq F$ co-nondeterministically as well. However, in order to get a deterministic algorithm for C finite, we cannot afford to try all 2^N possibilities for X . We observe that the existence of a $\langle *, E \rangle$ -tpp of X wrt Σ does not depend on X as such, but only on the set $B^+ \subseteq B$ as defined above. Thus, instead of checking all $X \subseteq F$, we only check all (at most 2^n) possibilities for $B^+ \subseteq B$. Given B^+ , we can test if there exists a corresponding $\langle *, E \rangle$ -tpp by the same method as above; moreover, we can test if there exists a set $X \subseteq F$ such that $B^+ = \bigcap_{w \in X} \{\varphi \in B : w \models \Box\varphi\}$ as this is equivalent to the condition

$$\forall \varphi \in B \setminus B^+ \exists w \in F w \models \neg\Box\varphi \wedge \bigwedge_{\psi \in B^+} \Box\psi,$$

easily verifiable by going through all φ and w . □

Theorem 4.2 *For any clx logic L , L -admissibility is computable in Π_2^{exp} , and L -unifiability is computable in Σ_2^{exp} .*

If L has bounded cluster size, or if the number of parameters is bounded by a constant, L -admissibility is computable in coNEXP , and L -unifiability in NEXP .

Proof: Using [8, Thm. 5.24], a rule Γ / Δ is not L -admissible if and only if it fails in some L -model F of size at most 4^n that is $\text{bas}(L)$ -pseudoextensible wrt Σ , where $\Sigma = \text{Sub}(\Gamma \cup \Delta)$, and $n = |\Sigma|$.

We can check this by nondeterministically choosing such a model F of size $N \leq 4^n$ equipped with an assignment $\models: \Sigma \rightarrow \mathbf{2}$, and verifying it is indeed a counterexample for admissibility of Γ / Δ . Given F , we can check in time $N^{O(1)} = 2^{O(n)}$ whether F is an L -frame (using [8, Thm. 4.29]), whether \models is an honest valuation respecting the connectives, and whether Γ / Δ is satisfied in the model $\langle F, \models \rangle$. By Lemma 4.1, we can check that F is $\langle C, m \rangle$ -pseudoextensible wrt Σ for each $\langle C, m \rangle \in \text{bas}(L)$ by a $\text{coNTIME}(2^{O(n^2)})$ algorithm. Overall, this gives a Σ_2^{exp} algorithm for checking inadmissibility.

If L has bounded cluster size, all conditions $\langle C, m \rangle \in \text{bas}(L)$ have C finite, thus we can check $\text{bas}(L)$ -pseudoextensibility of F in $\text{DTIME}(2^{O(n)})$, and inadmissibility in NEXP (in fact, in NE). The same bound also holds for arbitrary L in case the number of parameters is bounded by a constant k , as we may reduce the size of each C to at most 2^k . \square

4.2 Linear cluster-extensible logics

Observe that linear clx logics are cofinal-subframe logics.

Theorem 4.3 *Let L be a linear clx logic. Then L -admissibility is computable in NEXP , and L -unifiability in coNEXP .*

If L has bounded cluster size, or if the number of parameters is bounded by a constant, then L -admissibility and L -unifiability are in PSPACE .

Proof: Put $\text{bas}^i(L) = \{\langle C, m \rangle \in \text{bas}(L) : m = i\}$ and $\text{bas}_i(L) = \{C : \langle C, i \rangle \in \text{bas}(L)\}$ for $i = 0, 1$. Let us fix a rule Γ / Δ , and put $\Sigma = \text{Sub}(\Gamma \cup \Delta)$, $P = \Sigma \cap \text{Par}$, $B = \{\varphi : \Box\varphi \in \Sigma\}$, and $n = |\Sigma|$.

If W is a finite rooted L -model, we put $B^+(W) = \{\varphi \in B : W \models \varphi\}$, and we define predicates $G_i(W)$ for $i \geq |B^+(W)|$ inductively by

$$(16) \quad G_i(W) \quad \text{iff} \quad \forall C \in \text{bas}_1(L) \forall E \subseteq \mathbf{2}^P, |E| \leq_0 |C| : \\ \exists W' \in \text{Mod}_L \left(W \sqsubseteq \cdot W' \models \Gamma \ \& \ W' \setminus W \text{ is a } \langle \text{refl}(C), E \rangle\text{-tp of } W \right. \\ \left. \ \& \ (B^+(W') \subsetneq B^+(W) \Rightarrow G_{i-1}(W')) \right).$$

Notice that the constraint $B^+(W') \subsetneq B^+(W)$ ensures that the condition $i \geq |B^+(W)|$ is preserved; in particular, $G_0(W)$ never refers to the nonexistent G_{-1} , and it is well defined. Finally, we define $G(W)$ for arbitrary finite L -models W by

$$(17) \quad G(W) \quad \text{iff} \quad W \models \Gamma \ \& \ \forall u \in W \ G_{|B^+(W_u)|}(W_u).$$

Claim 4.3.1 *If W_0 and W_1 are rooted L -models of Γ such that $B^+(W_0) = B^+(W_1)$, and $i \geq |B^+(W_0)|$, then $G_i(W_0)$ iff $G_i(W_1)$.*

Proof: By induction on i . Fix C and E as in (16), and let W'_0 be a witness for $G_i(W_0)$, so that $U = W'_0 \setminus W_0$ is a $\langle \text{refl}(C), E \rangle$ -tp of W_0 . Put $W'_1 = W_1 \cup U$ (with $U < W_1$ so that $W_1 \subseteq \cdot W'_1$). Then W'_1 is an L -model, and $W'_1 \setminus W_1$ is a $\langle \text{refl}(C), E \rangle$ -tp of W_1 in W'_1 . We have $W'_1 \models \Gamma$: Γ holds in $W_1 \subseteq \cdot W'_1$ by assumption; it holds in U since it holds in W'_0 , and satisfaction of formulas from Σ in U only depends (apart from valuation of atoms in U , which did not change) on $B^+(W_1)$, equal to $B^+(W_0)$ by assumption. For the same reason, we obtain $B^+(W'_1) = B^+(W'_0)$, therefore, in case $B^+(W'_1) \subsetneq B^+(W_1)$, we have $G_{i-1}(W'_1)$ by the induction hypothesis. \square (Claim 4.3.1)

Claim 4.3.2 $G(W)$ holds if and only if W is a generated submodel of an L -model $F \models \Gamma$ such that F is $\text{bas}^1(L)$ -pseudoextensible wrt Σ .

Proof:

Left-to-right: Start with W , and by unwinding (17) and (16), attach all the tight predecessors that are asserted to exist. Call the resulting model F . It is clear that F is a finite L -model, $W \subseteq \cdot F$, and $F \models \Gamma$. In order to see that F is $\text{bas}^1(L)$ -pseudoextensible wrt Σ , let $C \in \text{bas}_1(L)$, and $E \subseteq \mathbf{2}^P$ with $|E| \leq_0 |C|$. By construction, all rooted submodels $W' \subseteq \cdot F$ satisfy $G_i(W')$ for an appropriate i , which guarantees the existence of a $\langle \text{refl}(C), E \rangle$ -tp (and a fortiori $\langle \text{refl}(C), E \rangle$ -tpp wrt Σ), with the following exception: no tp's of W' were added when W' resulted by attaching the cluster $\text{rcl}(W')$ below the model $W'' = W' \setminus \text{rcl}(W')$ which satisfied $G_i(W'')$ for some i , and such that $B^+(W') = B^+(W'')$. But then F includes a $\langle \text{refl}(C), E \rangle$ -tp of W'' , and since $B^+(W') = B^+(W'')$, this is also a $\langle \text{refl}(C), E \rangle$ -tpp of W' wrt Σ .

Right-to-left: It suffices to show that $G_i(W)$ if furthermore W is rooted, and $i \geq |B^+(W)|$. Let us proceed by induction on i as in the definition of $G_i(W)$. For any $C \in \text{bas}_1(L)$ and $E \subseteq \mathbf{2}^P$ such that $|E| \leq_0 |C|$, we may fix a $\langle \text{refl}(C), E \rangle$ -tpp $\{u_e : e \in E\} \subseteq F$ of W wrt Σ . We construct W' by attaching to W a new cluster $\{v_e : e \in E\}$, reflexive or not according to $\text{refl}(C)$, such that $v_e \models P^e$ satisfies the same variables as u_e . Clearly, $W \subseteq \cdot W'$, W' is an L -model, and $W' \setminus W$ is a $\langle \text{refl}(C), E \rangle$ -tp of W . Using the fact that $\{u_e : e \in E\}$ is a $\langle \text{refl}(C), E \rangle$ -tpp wrt Σ , we can prove $u_e \models \psi \Leftrightarrow v_e \models \psi$ for all $\psi \in \Sigma$ and $e \in E$ by induction on the complexity of ψ . This implies $W' \models \Gamma$, and $B^+(W') = B^+(F_{u_e})$ for any $e \in E$. Thus, either $B^+(W') = B^+(W)$, or $i > 0$ and $G_{i-1}(F_{u_e})$ by the induction hypothesis, which implies $G_{i-1}(W')$ by Claim 4.3.1. \square (Claim 4.3.2)

Claim 4.3.3 $\Gamma \not\vdash_L \Delta$ if and only if

$$(18) \quad \forall C \in \text{bas}_0(L) \forall E \subseteq \mathbf{2}^P, |E| \leq_0 |C| \exists W \in \text{Mod}_L (G(W) \ \& \ W \text{ is a } \langle \text{refl}(C), E \rangle\text{-tp of } \emptyset) \\ \& \ \forall \varphi \in \Delta \exists W \in \text{Mod}_L (G(W) \ \& \ W \not\models \varphi \ \& \ |W| \leq n + 1).$$

Proof:

Right-to-left: For each $C \in \text{bas}_0(L)$ and $E \subseteq \mathbf{2}^P$ such that $|E| \leq_0 |C|$, there exists an L -model $F_{C,E} \models \Gamma$ which is $\text{bas}^1(L)$ -pseudoextensible wrt Σ by Claim 4.3.2. Likewise, for each $\varphi \in \Delta$, there exists a $\text{bas}^1(L)$ -pseudoextensible wrt Σ L -model $F_\varphi \models \Gamma$ such that $F_\varphi \not\models \varphi$. Let F be the disjoint union of all the $F_{C,E}$'s and F_φ 's. Then F is an L -model refuting Γ / Δ that is $\text{bas}(L)$ -pseudoextensible wrt Σ , hence $\Gamma \not\vdash_L \Delta$ by [8, Thm. 5.24].

Left-to-right: Using [8, Thm. 5.24] again, there exists a locally finite $\text{bas}(L)$ -extensible L -model F such that $F \models \Gamma$, and $F \not\models \varphi$ for each $\varphi \in \Delta$. For any C and E as in (18), let W be a $\langle \text{refl}(C), E \rangle$ -tp of \emptyset inside F ; then $G(W)$ by Claim 4.3.2, as F is in particular $\text{bas}^1(L)$ -pseudoextensible wrt Σ .

A similar argument shows that for any $\varphi \in \Delta$, there is a finite rooted submodel $W \subseteq F$ such that $G(W)$, and $W \not\models \varphi$. Let $W_0 \subseteq W$ be a subset of size at most $n + 1$ that contains a point in a final cluster of W , and for each $\psi \in (B \setminus B^+(W)) \cup \{\varphi\}$, a point $u_\psi \not\models \psi$ such that ψ holds in all clusters strictly above $\text{cl}(u_\psi)$. Then W_0 is an L -model, and since W is a chain of clusters, it is easy to see that W_0 agrees with W on the truth of all formulas from Σ ; in particular, $W_0 \models \Gamma$, and $W_0, u_\varphi \not\models \varphi$. Moreover, we have $G(W_0)$ using Claim 4.3.1: for any $u \in W_0$, $B^+((W_0)_u) = B^+(W_u)$. \square (Claim 4.3.3)

We will now use Claim 4.3.3 and the definition of $G(W)$ to estimate the complexity of $\not\models_L$. Assume first that L has cluster size bounded by k , or that the number of parameters is bounded, in which case we put $k = 2^{|P|}$. Then all the sets E referred to in (16) and (18) have size at most k . We may directly “execute” (18), (17), and (16) by an alternating Turing machine, with existential and universal quantifiers implemented by nondeterministic choices in existential or universal states (respectively).

Since $|B^+(W)| < n$, there are $O(n)$ alternations. For the universal quantifiers, we need $O(1)$ bits to specify C , $O(kn)$ bits to specify E , and $O(\log n)$ bits to specify $\varphi \in \Delta$. For the existential quantifiers, we need $n^{O(1)}$ bits to specify the model W of size $\leq n + 1$ on the second line of (18); the remaining existential quantifiers quantify over a single cluster of size $\leq k$ each, hence they can be specified with $O(kn)$ bits. The other conditions in the definition (e.g., truth of Γ) may be tested deterministically in polynomial time from the data. Thus, we obtain an alternating polynomial-time algorithm for checking $\Gamma \not\models_L \Delta$, showing that it is in $\text{AP} = \text{PSPACE}$.

In the general case where both the cluster size and the number of parameters are unbounded, we need $O(2^n)$ bits to specify the universally quantified E and the existentially quantified W' . The former is expected as we are shooting for coNEXP , but the latter is a problem. Similarly to the proof of Lemma 4.1, we solve it by existentially quantifying only over the sets $B^+(W') \subseteq B$ instead of W' proper, which makes sense because of Claim 4.3.1.

In more detail, let $G'_i(B^+)$ (with $B^+ \subseteq B$ and $i \geq |B^+|$) stand for $G_i(W)$ with input $B^+(W)$ in place of W . Put $V = \Sigma \cap \text{Var}$. For any $v \in \mathbf{2}^V$, $e \in \mathbf{2}^P$, and $B^+ \subseteq B$, let $A(v, e, B^+)$ denote the Boolean assignment that agrees with v and e on variables and parameters (respectively), and that makes $\Box\varphi$ true for $\varphi \in B^+$, and false for $\varphi \in B \setminus B^+$. Then we can express $G'_i(B^+)$ as the conjunction of

$$\begin{aligned} \forall E \subseteq \mathbf{2}^P \setminus \{\emptyset\} \exists D \subseteq B^+ \exists f: D \rightarrow E \\ \left(\forall e \in E \exists v \in \mathbf{2}^V A(v, e, B^+ \setminus D) \models \bigwedge \Gamma \wedge \bigwedge_{\varphi \in B^+ \setminus D} \varphi \wedge \bigwedge_{\substack{\varphi \in D \\ f(\varphi) = e}} \neg \varphi \right. \\ \left. \& (D \neq \emptyset \Rightarrow G'_{i-1}(B^+ \setminus D)) \right) \end{aligned}$$

if $\otimes \in \text{bas}_1(L)$, a similar conjunct restricted to $|E| \leq k$ if $\textcircled{k} \in \text{bas}_1(L)$, and the conjunct

$$\forall e \in \mathbf{2}^P \exists v \in \mathbf{2}^V \exists D \subseteq B^+ \left(D = \{\varphi \in B^+ : A(v, e, B^+) \not\models \varphi\} \right. \\ \left. \& A(v, e, B^+) \models \bigwedge \Gamma \& (D \neq \emptyset \Rightarrow G'_{i-1}(B^+ \setminus D)) \right)$$

if $\bullet \in \text{bas}_1(L)$. We may rewrite the first conjunct in (18) analogously.

In this way, we get an alternating algorithm for $\not\models_L$ using time $2^{n^{O(1)}}$, out of which time $t_{\exists}(n) = n^{O(1)}$ is spent in existential states. We can simulate the existential states by deterministically trying all possible choices; this will multiply the overall running time by a factor of $2^{t_{\exists}(n)}$, which is still $2^{n^{O(1)}}$, hence we obtain a coNEXP algorithm. \square

Remark 4.4 Admissibility in linear clx logics with no parameters at all is even easier than PSPACE, viz. coNP, by [5, Thm. 2.6].

4.3 Logics of bounded depth

Our basic goal is to completely determine the complexity of admissibility and unification for clx logics. As we will see later, the only cases remaining where we can improve upon the upper bounds proved so far is for clx logics of branching 0 (i.e., depth 1) when, moreover, the cluster size or the number of parameters is bounded. This is a tiny class of problems, most of which are not really interesting. For this reason, we broaden the scope of this section beyond clx logics, and we will look at larger classes of logics of bounded depth. We are primarily interested in cases where we can get the complexity inside the polynomial hierarchy, but we will also include for completeness some results of higher complexity, analogous to clx logics.

The added generality comes at a cost: the analysis of admissible rules from [8] does not apply to non-clx logics, hence we need to come up with a relevant theory of admissible rules for logics of bounded depth. Fortunately, this is fairly easy to do. The key factor is that logics of bounded depth are *locally tabular*:

Theorem 4.5 ([3, Thm. 12.21]) *Every logic L of finite depth is locally tabular, i.e., there are only finitely many nonequivalent formulas over any finite set of atoms, or equivalently, the canonical frames $C_L(P, V)$ are finite for any finite $P \subseteq \text{Par}$ and $V \subseteq \text{Var}$ (and more generally, all finitely generated refined L -frames are finite).*

Consequently, all logics L of finite depth, as well as atomic multi-conclusion consequence relations that extend them (such as \sim_L), have the finite model property. \square

The finite model property with respect to admissibility avoids most of the trouble we had to go through in [8].

Let us start with a simple result on unifiability in *all* logics of bounded depth as a teaser.

Theorem 4.6 *If L is a logic of bounded depth, then L -unifiability with a constant number of parameters is in NP.*

In particular, this holds if L is a clx logic of branching 0.

Proof: A formula φ with parameters from a fixed finite set P is L -unifiable iff it has a variable-free unifier iff there is a (definable) valuation in the canonical frame $C_L(P, \emptyset)$ that makes φ true in all points of the model. By Theorem 4.5, $C_L(P, \emptyset)$ is a fixed finite frame, which makes the condition above checkable in NP. \square

Subsequent complexity results will only apply under further restrictions, not to all logics of bounded depth. When the number of parameters is unlimited, we will see in the next section that L -unifiability is NEXP-hard or coNEXP-hard whenever L has width ≥ 2 or unbounded cluster size (Theorems 5.3 and 5.6). Thus, we may hope to get better complexity only for linear logics of bounded cluster size, which in view of bounded depth means that L is tabular.

Contrary to Theorem 4.6, we need some restrictions even if the number of parameters is finite when we discuss admissibility rather than unifiability: we know from [5] that logics such as **S4BD₃** have a coNEXP-hard admissibility problem already with no parameters at all. Even if we just wanted results analogous to the previous sections, placing admissibility on a low level of the exponential hierarchy, this does not seem possible for general logics of bounded depth: cf. the construction of logics of depth d in [3, §18.4] such that some nonderivable formulas require countermodels whose size is a tower of exponentials of height $d - 2$. With these examples in mind, it seems we can only expect reasonably general results if we restrict attention to logics of bounded width (which is equivalent to bounded branching for logics of bounded depth).

But first, we need a rudimentary theory of admissibility in logics of bounded depth.

Definition 4.7 Let L be a logic of bounded depth, and assume that Par is finite. A finite parametric L -frame F has a *loosely separated root* if it is rooted, and points of $\text{rcl}(F)$ have pairwise distinct valuations of parameters. The frame F has a *separated root* if, moreover, it is *not* the case that $F \setminus \text{rcl}(F)$ is a rooted frame with reflexive root cluster C , and for each $w \in \text{rcl}(F)$ there is $x \in C$ such that $w \equiv_{\text{Par}} x$.

A finite parametric L -frame W is *L -extensible* (or *strongly L -extensible*) if for every parametric L -frame F with a separated root (a loosely separated root, resp.) such that $F \setminus \text{rcl}(F) \subseteq W$, $\text{id}_{F \setminus \text{rcl}(F)}$ extends to an isomorphism of F onto a generated subframe of W .

While the concept of strong L -extensibility looks simpler, and it is easier to work with, L -extensibility is a structurally more fundamental property due to the following characterization: L -extensible frames are exactly the finitely generated *injective* L -frames, or duals of finitely generated projective L -algebras.

Lemma 4.8 Let L be a logic of bounded depth, and assume that Par is finite. The following are equivalent for any finite parametric L -frame W .

- (i) W is L -extensible.
- (ii) For every pair of finite parametric L -frames $F_0 \subseteq F$, and a p -morphism $f_0: F_0 \rightarrow W$, there exists a p -morphism $f: F \rightarrow W$ such that $f \supseteq f_0$.

Proof:

(i) \rightarrow (ii): We may extend the p -morphism to one cluster after another in a top-down fashion, hence we may assume that F is rooted, and $F_0 = F \setminus \text{rcl}(F)$. Let F' be $f_0[F_0] \subseteq W$

extended with a copy of $\text{rcl}(F)$ as a new root cluster: then f_0 extends to a p-morphism of F onto F' , hence replacing F with F' , we may assume without loss of generality that $F_0 \subseteq \cdot W$, and $f_0 = \text{id}_{F_0}$. We may shrink the root cluster of F by a p-morphism identifying all points that satisfy the same parameters, hence we may assume that points of $\text{rcl}(F)$ have pairwise distinct valuations of parameters. If F_0 has a reflexive root cluster, and for each $w \in \text{rcl}(F)$ there exists $x_w \in \text{rcl}(F_0)$ such that $w \equiv_{\text{Par}} x_w$, then id_{F_0} extended with the mapping $w \mapsto x_w$ is a p-morphism $F \rightarrow F_0$. Otherwise F has a separated root, hence $f_0 = \text{id}_{F_0}$ extends to a p-morphism $F \rightarrow W$ using (i).

(ii) \rightarrow (i): Assume that F has a separated root, and $F_0 = F \setminus \text{rcl}(F) \subseteq \cdot W$. Using (ii), id_{F_0} extends to a p-morphism $f: F \rightarrow W$. Since F has a separated root, f cannot identify distinct points of $\text{rcl}(F)$, and it cannot map any point of $\text{rcl}(F)$ to F_0 , thus it is an isomorphic embedding in W . \square

Lemma 4.9 *Let L be a logic of bounded depth, and assume that Par is finite.*

- (i) *For any finite $V \subseteq \text{Var}$, the canonical parametric frame $C_L(\text{Par}, V)$ is L -extensible.*
- (ii) *If W is an L -extensible finite parametric L -frame, there exists a finite $V \subseteq \text{Var}$ such that W is a p-morphic image of $C_L(\text{Par}, V)$.*

Proof:

(i): The frame $C_L(\text{Par}, V)$ is finite by Theorem 4.5. If F is a finite L -frame with a separated root such that $F_0 = F \setminus \text{rcl}(F) \subseteq \cdot C_L(\text{Par}, V)$, let \models be a valuation in F that agrees in F_0 with the canonical valuation of $C_L(\text{Par}, V)$, and is arbitrary in $\text{rcl}(F)$. Define a mapping $f: F \rightarrow C_L(\text{Par}, V)$ by $f(x) = \text{Sat}_{\text{Form}(\text{Par}, V)}(F, x)$. Then f is a p-morphism of F to $C_L(\text{Par}, V)$ extending id_{F_0} . As in Lemma 4.8, f is an isomorphic embedding because F has separated root.

(ii): By [8, Lem. 2.3], W is isomorphic to a generated subframe $W' \subseteq \cdot C_L(\text{Par}, V)$ for sufficiently large finite $V \subseteq \text{Var}$. The isomorphism $W' \rightarrow W$ extends to a p-morphism $C_L(\text{Par}, V) \rightarrow W$ by Lemma 4.8. \square

Remark 4.10 In fact, L -extensible frames are exactly the *retracts* of the finitely generated canonical frames $C_L(\text{Par}, V)$.

It is not generally true that p-morphic images of $C_L(\text{Par}, V)$ are L -extensible. For example, let F_0 be the disjoint union of the two-element cluster and the two-element reflexive chain, and let L be the logic of F_0 . If $V \neq \emptyset$, then one of the connected components of $C_L(\text{Par}, V)$ is a two-element cluster where both elements satisfy the same parameters. We may contract this cluster to a single point by a p-morphism, but then the resulting frame is not L -extensible (if $\text{Par} \neq \emptyset$). For much the same reason, there exist inadmissible rules $\Gamma \not\vdash_L \Delta$ with $\Gamma \cup \Delta \subseteq \text{Form}(\text{Par}, V)$ such that there is no projective formula $\theta \in \text{Form}(\text{Par}, V)$ such that $\theta \vdash_L \bigwedge \Gamma$, and $\theta \not\vdash_L \varphi$ for each $\varphi \in \Delta$. An explicit example is the rule $p \vee \diamond x \rightarrow \Box x / x$ for $\text{Par} = \{p\}$ and $V = \{x\}$.

It is easy to show (e.g. using the frame rules from [6]) that p-morphic images of finitely generated canonical frames are exactly the finite L -frames that validate all L -admissible rules.

The canonical frames $C_L(\text{Par}, V)$ are in general not strongly L -extensible. For example, let $L = \mathbf{S4.3} \oplus \mathbf{BD}_2$. Then $C_L(\text{Par}, V)$ contains a cluster C such that every possible valuation of

$\text{Par} \cup V$ is realized by a point of C . There cannot be any cluster below C , as any such could be p-morphically contracted into C while preserving the canonical valuation. Thus, $C_L(\text{Par}, V)$ is not strongly L -extensible.

Lemma 4.11 *Let L be a logic of bounded depth, and assume that Par is finite. Then any finite parametric L -frame is a generated subframe of a strongly L -extensible finite parametric L -frame.*

Proof: Let W_0 be the given frame, and d be the depth of L . We define a sequence of parametric L -frames $W_0 \subseteq \cdot W_1 \subseteq \cdot \dots \subseteq \cdot W_d$ by induction as follows: for every parametric L -frame F with a loosely separated root such that $F \setminus \text{rcl}(F) \subseteq \cdot W_i$, but $\text{id}_{F \setminus \text{rcl}(F)}$ does not extend to an isomorphism of F to a generated subframe of W_i , we add a copy of $\text{rcl}(F)$ to W_{i+1} . (We do not distinguish among isomorphic frames F : they all get a single common copy of $\text{rcl}(F)$.)

Since all the clusters being added have size at most $2^{|\text{Par}|}$, there are only finitely many choices for F up to isomorphism, hence all the frames W_i are finite. Moreover, we see by induction on i that all frames F as above with $F \setminus \text{rcl}(F) \subseteq \cdot W_i$ such that the depth of F is less than i already have a copy in W_i , i.e., only clusters of depth $\geq i + 1$ are added to W_{i+1} . It follows that W_d is strongly L -extensible. \square

Each of the classes of finitely generated canonical frames, p-morphic images of canonical frames, L -extensible frames, or strongly L -extensible frames, provides an adequate semantics for admissible rules. We will mostly work with strongly L -extensible frames, as the other possibilities are not suitable for our purposes for various reasons: canonical frames are astronomically large, and are not flexible enough so that they could be constructed from small pieces as needed; p-morphic images of canonical frames are difficult to algorithmically recognize. There is not much difference between L -extensible and strongly L -extensible frames, however strongly L -extensible frames have a somewhat simpler definition, and are easier to manipulate.

Theorem 4.12 *Let L be a logic of bounded depth, and assume that Par is finite. The following are equivalent for any rule Γ / Δ :*

- (i) $\Gamma \vdash_L \Delta$.
- (ii) Γ / Δ holds in all L -extensible finite parametric L -frames.
- (iii) Γ / Δ holds in all strongly L -extensible finite parametric L -frames.

Proof:

(i) \leftrightarrow (ii): By [8, Lem. 2.2], Γ / Δ is L -admissible iff it holds in $C_L(\text{Par}, V)$ for all finite sets of variables V . Thus, on the one hand, if $\Gamma \vdash_L \Delta$, then Γ / Δ holds in all L -extensible frames by Lemma 4.9, as the validity of rules is preserved by p-morphic images. On the other hand, if Γ / Δ holds in L -extensible frames, it holds in particular in every $C_L(\text{Par}, V)$ by Lemma 4.9, hence $\Gamma \vdash_L \Delta$.

(ii) \rightarrow (iii) is trivial. (iii) \rightarrow (ii): Assume that Γ / Δ holds in all strongly L -extensible frames, and let W be an L -extensible frame. By Lemma 4.11, there exists a strongly L -extensible frame W' such that $W \subseteq \cdot W'$. By Lemma 4.8, there exists a p-morphism $f: W' \rightarrow W$ extending id_W . Thus, $W' \models \Gamma / \Delta$ implies $W \models \Gamma / \Delta$. \square

Remark 4.13 We can, and will, apply the characterization from Theorem 4.12 even if Par is infinite, by considering frames with valuation of a finite set of parameters $P \subseteq \text{Par}$, which will be formally accomplished by shrinking Par to a finite subset for the duration of an argument. We formulated the criteria the way we did for notational convenience.

Lemma 4.14 *Any logic L of finite depth d is axiomatizable over $\mathbf{K4BD}_d$ by a set of frame formulas $\{\alpha^\sharp(F_i, \perp) : i \in I\}$.*

Proof: Let $\{F_i : i \in I\}$ be an enumeration of finite rooted frames of depth d such that $F_i \not\models L$, and put $L' = \mathbf{K4BD}_d \oplus \{\alpha^\sharp(F_i, \perp) : i \in I\}$. On the one hand, $F_i \not\models L$ implies $\vdash_L \alpha^\sharp(F_i, \perp)$, hence $L' \subseteq L$. Assume for contradiction that $\vdash_L \varphi$, but $\not\vdash_{L'} \varphi$. Being a logic of bounded depth, L' has the finite model property, hence there exists a finite rooted L' -frame $F \not\models \varphi$ of depth d . Since $\vdash_L \varphi$, F is not an L -frame. But then $\alpha^\sharp(F, \perp)$ is an axiom of L' , which implies that F is not an L' -frame after all, a contradiction. \square

We now focus on logics of bounded depth and width. All such logics are finitely axiomatizable (see [7, Thm. 5.16]), and therefore their finite frames are coNP-recognizable. In fact, we can improve this to P:

Lemma 4.15 *Let L be a logic of bounded depth and width. Then finite L -frames are recognizable in polynomial time.*

Proof: By [7, Thm. 5.16] and Lemma 4.14, L is axiomatizable by $\mathbf{K4BD}_k \mathbf{BW}_k$ for some constant k , plus finitely many frame formulas $\alpha^\sharp(F_i, \perp)$. We can compute the depth and width of a finite frame in polynomial time. Moreover, given a rooted frame W of width and depth at most k , we can recognize in polynomial time if W reduces to F_i : if F_i has no proper clusters, then for any p-morphism $f: W \rightarrow F_i$, preimages of all points are unions of clusters of W . Thus, there are only $O(1)$ choices for f , which are easy to check. In the general case, let F'_i be F_i with all clusters reduced to size 1. Then a p-morphism $W \rightarrow F_i$ induces a p-morphism $W \rightarrow F'_i$; conversely, a p-morphism $W \rightarrow F'_i$ is induced by a p-morphism $W \rightarrow F_i$ if and only if the maximal clusters in the preimage of each point of F'_i have size at least the size of the corresponding cluster of F_i . This is again easy to check in polynomial time. \square

Corollary 4.16 *For any logic L of bounded depth and width, there exists a constant c with the following property: whenever W is a Kripke L -frame, and C is a cluster of W of size $|C| \geq c$, then the frames obtained from W by enlarging C to arbitrary cardinality are also L -frames.*

Proof: Let c be the maximum of cluster sizes of the frames F_i in the proof of Lemma 4.15. \square

Incidentally, since logics of bounded depth and width have the polynomial-size model property (see e.g. [7, L. 5.13]), Lemma 4.15 implies that all such logics are in coNP. (This will also follow from Theorem 4.19.)

Corollary 4.17 *If L is a logic of bounded depth and width, then L -derivability is in coNP. \square*

Lemma 4.18 *Let L be a logic of bounded depth and width, and assume that Par is finite. There is a polynomial $p(x)$ such that whenever Γ / Δ is a rule of length n , and $\Gamma \not\vdash_L \Delta$, then Γ / Δ is falsified in a strongly L -extensible parametric L -frame of size at most $p(n)$.*

Proof: Put $\Sigma = \text{Sub}(\Gamma \cup \Delta)$, $B = \{\varphi : \Box\varphi \in \Sigma\}$, and $n = |\Sigma|$. Let \models be a valuation falsifying Γ / Δ in a strongly L -extensible parametric finite L -frame W , using Theorem 4.12. For each $\varphi \in \Delta$, fix $x_\varphi \in W$ such that $x_\varphi \not\models \varphi$, and put $W_0 = \bigcup_{\varphi \in \Delta} x_\varphi \uparrow$.

Starting with W_0 , we define a sequence $W_0 \subseteq W_1 \subseteq \dots \subseteq W_d$ of generated subframes of W as follows: $W_{i+1} \supseteq W_i$, and whenever $F \subseteq W$ has a loosely separated root, $F \setminus \text{rcl}(F) \subseteq W_i$, but $\text{id}_{F \setminus \text{rcl}(F)}$ does not extend to a p-morphic embedding $F \rightarrow W_i$, we add a copy of $\text{rcl}(F)$ to W_{i+1} using strong L -extensibility of W . As in Lemma 4.11, we do not distinguish among isomorphic F .

By the same argument as in Lemma 4.11, the frame W_d is strongly L -extensible. Let s_i be the number of clusters in W_i . Since L has bounded depth and width, any rooted L -frame has only $O(1)$ clusters, thus $s_0 = O(|\Delta|) = O(n)$. When constructing W_{i+1} from W_i , each $F \setminus \text{rcl}(F)$ is generated by an antichain in W_i ; since L has bounded width w , there are at most s_i^{w+1} choices for $F \setminus \text{rcl}(F)$. For each of them, there are at most $2^{2^{|\text{Par}|}} + 2^{|\text{Par}|} = O(1)$ choices for $\text{rcl}(F)$ up to isomorphism. Thus, $s_{i+1} = s_i^{O(1)}$. Since a constant number of iterations of a polynomial is still polynomial, $s_d = n^{O(1)}$.

We construct a submodel $W' \subseteq W_d$ as follows:

- The points x_φ , $\varphi \in \Delta$, are included in W' .
- For each cluster C of W_d , and each assignment $e \in \mathbf{2}^{\text{Par}}$ that is realized in C , we include in W' at least one point from C satisfying e .
- For each cluster C of W_d , and each formula $\varphi \in B$ that is falsified in some point of C , we include at least one such point in W' .
- All clusters of W_d of size $\leq c$ are fully included in W' , where c is the constant from Corollary 4.16.
- For each cluster C of W_d of size $\geq c$, we make sure W' includes at least c points of C .

We may choose W' satisfying these requirements so that it has cluster size at most $|B| + 2^{|\text{Par}|} + |\Delta| + c = O(n)$, thus the overall size of W' is polynomial in n . Notice that W' has the same skeleton as W_d : we merely shrunk large clusters to control their size.

The choice of W' ensures that it is an L -model, and that it agrees with W on the satisfaction of formulas from Σ ; in particular, Γ / Δ fails in W' . We claim that W' is strongly L -extensible. Let F' be a parametric L -frame with a loosely separated root such that $F' \setminus \text{rcl}(F') \subseteq W'$. Let F denote $(F' \setminus \text{rcl}(F')) \uparrow \subseteq W_d$ with $\text{rcl}(F')$ attached below (i.e., we expand the clusters of W' included in $F' \setminus \text{rcl}(F')$ back to what they were in W_d). By Corollary 4.16 and the choice of W' , F is still an L -frame with a loosely separated root. Since W_d is strongly L -extensible, $\text{id}_{F \setminus \text{rcl}(F)}$ extends to a p-morphic embedding $f: F \rightarrow W_d$. Since all points of $\text{rcl}(F) = \text{rcl}(F')$ are distinguished by valuation of parameters, all of $\text{rcl}(F)$ is included in W' , thus $f: F \rightarrow W'$. \square

Theorem 4.19 *Let L be a logic of bounded depth and width. Then L -admissibility with a constant number of parameters is in coNP.*

In particular, this holds if L is a clx logic of branching 0.

Proof: By Lemma 4.18, we have $\Gamma \not\vdash_L \Delta$ iff Γ / Δ is falsified in a strongly L -extensible L -model W of polynomial size. Clearly, we can check in polynomial time that $W \not\models \Gamma / \Delta$, and by Lemma 4.15, we can check that W is an L -model.

We can also check in polynomial time that W is strongly L -extensible: by the same argument as in the proof of Lemma 4.18, there are only polynomially many choices for a parametric L -frame F with a loosely separated root such that $F \setminus \text{rcl}(F) \subseteq \cdot W$. We can generate each such frame in polynomial time, check that it is indeed an L -frame with a loosely separated root, and check if we can find a suitable copy of $\text{rcl}(F)$ inside W . \square

The situation for admissibility with unrestricted number of parameters is more complicated. As we already mentioned, cluster size and linearity will come into play. Notice that logics of bounded depth, width, and cluster size are tabular.

Theorem 4.20 *Let L be a tabular logic. Then L -admissibility is decidable in coNEXP, and L -unifiability in NEXP.*

Proof: Let us fix a rule Γ / Δ , and put $\Sigma = \text{Sub}(\Gamma \cup \Delta)$ and $n = |\Sigma|$. In order to apply the criterion from Theorem 4.12, we will proceed as if Par were finite, consisting of only the parameters from Σ (cf. Remark 4.13). However, we bear in mind that $|\text{Par}|$ is bounded by n , not by a constant, for complexity estimates.

We claim that if $\Gamma \not\vdash_L \Delta$, there exists a strongly L -extensible L -model W falsifying Γ / Δ of size $2^{O(n)}$. This can be shown by constructing a sequence of models $W_0 \subseteq \cdot W_1 \subseteq \cdot \dots \subseteq \cdot W_d$ as in the proof of Lemma 4.18. Since L has bounded cluster size, there are only $2^{O(n)}$ possible root clusters in frames with loosely separated roots; with this in mind, it is easy to compute that each W_i has size $2^{O(n)}$. (We do not need to select $W' \subseteq W_d$, as W_d already has bounded cluster size.)

Given a model W of size $2^{O(n)}$, it is easy to check in deterministic exponential time that W is an L -frame, and that $W \not\models \Gamma / \Delta$. We may also check in exponential time if W is strongly L -extensible: by the same counting argument as above, there are $2^{O(n)}$ possible frames F that we need to test, and we can generate and check them in exponential time. \square

Theorem 4.21 *Let L be a linear logic of depth d . Then L -admissibility is decidable in NEXP, and L -unifiability in coNEXP.*

If L has bounded cluster size (i.e., it is tabular), then L -admissibility is in Σ_{2d}^p , and L -unifiability in Π_{2d}^p .

Proof: We use a variant of the algorithm presented in Theorem 4.3. We fix Γ / Δ , and put $\Sigma = \text{Sub}(\Gamma \cup \Delta)$, $B = \{\varphi : \Box\varphi \in \Sigma\}$, and $n = |\Sigma|$. We will again proceed as if Par were finite of size $\leq n$.

Given a finite parametric L -frame F , let $\text{Ext}_L(F)$ denote the set of single-cluster parametric frames C such that F^C is an L -frame with a loosely separated root, where F^C denotes F with C

added as a new root cluster. Using Lemma 4.15, it is polynomial-time decidable if $C \in \text{Ext}_L(F)$. If W is an L -model, let $\text{Fr}(W)$ denote its underlying parametric frame. By abuse of notation, we will write $\text{Ext}_L(W)$ instead of $\text{Ext}_L(\text{Fr}(W))$.

Similarly to the proof of Theorem 4.3, but avoiding the complicated set-up using $B^+(W)$ to guard the search, we define predicates $G_i(W)$ for W a finite rooted L -model, and $i < d$, such that $G_0(W)$ is always true, and

$$(19) \quad G_{i+1}(W) \quad \text{iff} \quad \forall C \in \text{Ext}_L(W) \exists W' (W \subseteq \cdot W' \models \Gamma \ \& \ \text{Fr}(W') = \text{Fr}(W)^C \ \& \ G_i(W')).$$

We define $G(W)$ by

$$(20) \quad G(W) \quad \text{iff} \quad W \models \Gamma \ \& \ \forall u \in W \ G_{d-1}(W_u).$$

Claim 4.21.1 $\Gamma \not\vdash_L \Delta$ if and only if

$$(21) \quad \begin{aligned} \forall C \in \text{Ext}_L(\emptyset) \exists W \in \text{Mod}_L (G(W) \ \& \ \text{Fr}(W) = C) \\ \ \& \ \forall \varphi \in \Delta \exists W \in \text{Mod}_L (G(W) \ \& \ W \not\models \varphi). \end{aligned}$$

Moreover, we may restrict the size of W on the second line of (21) to $n + O(1)$.

Proof:

Left-to-right: Let U be a finite strongly L -extensible L -model falsifying Γ / Δ , using Theorem 4.12. First, we show by induction on i that if $u \in U$ is of depth at least $d - i$, then $G_i(U_u)$ holds. This is trivial for $i = 0$. For the induction step, if $i > 0$, and $C \in \text{Ext}_L(U_u)$, then by strong L -extensibility of U , we can find a copy of C in U so that the model $W = C \uparrow$ satisfies $U_u \subseteq \cdot W$ and $\text{Fr}(W) = \text{Fr}(U_u)^C$. Since $W \subseteq \cdot U$, $W \models \Gamma$. We have $G_{i-1}(W)$ by the induction hypothesis as $\text{rcl}(W)$ has depth at least $d - i + 1$.

By taking $i = 1$, it follows that $G(U_u)$ for all $u \in U$. Then the first line of (21) holds by L -extensibility of U , and the second one as each $\varphi \in \Delta$ is falsified in U .

Concerning the bound on $|W|$, let $W = U_u$ for some $u \neq \varphi$. If L is tabular, $|W| = O(1)$ and we are done. In general, W is a chain of $j \leq d$ clusters $C_j < C_{j-1} < \dots < C_1$. We first modify the model U to a model U' so that for each $1 \leq i < j$, we put a new copy of the submodel $C_{i+1} \downarrow$ below C_i , and likewise, we include a copy of $C_1 \downarrow$ as a new connected component of U' . Then we shrink each original cluster C_i to a subcluster $C'_i \subseteq C_i$ such that:

- $|C'_i| \geq c$, where c is the constant from Corollary 4.16.
- $u \in C'_j$.
- For each formula $\psi \in B$ such that $W \not\models \psi$, let i be minimal such that $C_i \uparrow \not\models \psi$. Then C'_i includes a point falsifying ψ .

Let $W'' = \bigcup_i C'_i \subseteq W$, and $U'' = (U' \setminus W) \cup W'' \subseteq U'$. Clearly, U'' is an L -frame, $W'' \subseteq \cdot U''$, and the construction ensures that $\text{Sat}_\Sigma(U', x) = \text{Sat}_\Sigma(U'', x)$ for each $x \in U''$; in particular, $U'' \models \Gamma$, and $W'' \not\models \varphi$. We have $|W''| \leq cd + n$.

We claim that U'' is strongly L -extensible, which implies $G(U''_x)$ for all $x \in U''$ by the first part of the proof, and, in particular, $G(W'')$. So, let $F'_0 \subseteq \cdot U''$, and $C \in \text{Ext}_L(F'_0)$. Let $F_0 \supseteq F'_0$

be the frame obtained by replacing each C'_i with C_i (if it is included in F'_0). By Corollary 4.16, $C \in \text{Ext}_L(F_0)$. If $F_0 \subseteq U$, strong L -extensibility of U ensures there is an isomorphism f of F_0^C to a generated subframe of U extending F_0 . If $f[C] = C_i$ for some i , we may modify f so that $f[C]$ is instead the new copy of C_i included in U' . Thus, we may assume that $f[C]$ is disjoint from W , hence $f[C] \subseteq U''$. But then f maps F_0^C to a generated subframe of U'' extending F'_0 .

If $F_0 \not\subseteq U$, it is a rooted frame whose root cluster is inside one of the new copies of $C_i \downarrow$ in U' . Then strong L -extensibility of U , applied in the original copy of $C_i \downarrow$ and transferred back, ensures that id_{F_0} extends to an isomorphism f of F_0^C to a generated subframe of U' . Since $f[C]$ is also in the new copy of $C_i \downarrow$, it is included in U'' , hence f restricts to an isomorphism of F_0^C to a generated subframe of U'' .

Right-to-left: For each $C \in \text{Ext}_L(\emptyset)$, we fix a model W_C based on C such that $G(W_C)$, and for each $\varphi \in \Delta$, we fix a model $W_\varphi \not\models \varphi$ such that $G(W_\varphi)$. For each W_C and W_φ , we unwind the definition of G to extend the models downwards with all the “tight predecessors” to an upside-down tree of clusters. We let U be the disjoint sum of the models we obtain. Clearly, U is a finite L -model such that $U \models \Gamma$, and $U \not\models \varphi$ for each $\varphi \in \Delta$. By induction on i , we see that any point $u \in U$ of depth i satisfies $G_j(U_u)$ for some $j \geq d - i$, which was unwinded in the course of the construction of U . It follows that U is strongly L -extensible: let $F_0 \subseteq U$, and $C \in \text{Ext}_L(F_0)$. If $F_0 = \emptyset$, then a copy of C was added to U as W_C . Otherwise F_0 is a rooted L -model of depth at most $d - 1$, hence $G_j(F_0)$ for some $j \geq 1$, hence F_0 was extended with a copy of C while unwinding it. \square (Claim 4.21.1)

As in the proof of Theorem 4.3, (19)–(21) give an algorithm deciding $\Gamma \not\models_L \Delta$ on an alternating Turing machine. If L is tabular, the clusters C and models W or W' appearing in (19)–(21) have bounded size, hence they can be specified using $O(n)$ bits (encoding the relevant valuations of parameters or variables). Thus, the algorithm works in alternating polynomial time. Since it consists of d pairs of universal quantifiers followed by existential quantifiers, it is a Π_{2d}^P algorithm.

If L has unbounded cluster size, the quantifiers over C , W , and W' in (19)–(21) have exponential size, hence as such, we only get a Π_{2d}^{exp} algorithm. We keep the $\forall C$ quantifiers, but we will need something more efficient for the existential quantifiers. We can do it in the same way as in Lemma 4.1 or Theorem 4.3: instead of specifying the full models W or W' , we only specify the set of formulas $\Box\psi$, $\psi \in B$, that they satisfy, while checking separately that this choice of the subset of B is satisfiable in the frame, and validates Γ . Of course, we also need to know the underlying frame of the model, but in (19) and on the first line of (21), this is already fully specified by C . The models W on the second line of (21) may be written down explicitly, as they have polynomial size.

In this way, the existential quantifiers in the algorithm become polynomial-size, hence they can be simulated in deterministic exponential time by trying all possibilities. The universal quantifiers remain, hence we obtain a coNEXP algorithm for $\Gamma \not\models_L \Delta$. \square

Corollary 4.22 *If L is a clx logic of branching 0 and bounded cluster size, then L -admissibility is in Σ_2^P , and L -unifiability in Π_2^P .* \square

For completeness, we will also include an algorithm for admissibility with unlimited pa-

rameters in general logics of bounded depth and width with no bound on cluster size. The argument in Lemma 4.18 and Theorem 4.20 only shows that an inadmissible rule is refuted in a strongly extensible model of size $2^{O(2^n)}$, and this bound cannot be improved. This shows that L -admissibility is decidable in $\text{coNTIME}(2^{O(2^n)})$, i.e., in coNEE . We need a little more elaborate argument to obtain complexity that matches lower bounds from the next section (Theorem 5.10).

Definition 4.23 Assume that Par is finite. Let L be a logic of finite width and depth, and Σ a finite set of formulas closed under subformulas.

A finite L -model W is a Σ -pruned L -extensible model if for every parametric L -frame F with a loosely separated root such that $F \setminus \text{rcl}(F)$ is a generated subframe of (the underlying parametric frame of) W , there exists a subframe $F \setminus \text{rcl}(F) \subsetneq F_1 \subseteq F$ such that $\text{id}_{F \setminus \text{rcl}(F)}$ extends to an isomorphism f of F_1 to a generated parametric subframe of W , and the valuation on F_1 lifted from W via f extends to a valuation on F such that

$$F \models \varphi \quad \text{iff} \quad F_1 \models \varphi$$

for all $\varphi \in \Sigma$.

Lemma 4.24 Assume that Par is finite. Let L be a logic of finite width and depth, Σ a finite set of formulas closed under subformulas of size $n = |\Sigma|$, and $\Gamma \cup \Delta \subseteq \Sigma$.

- (i) If $\Gamma \not\vdash_L \Delta$, then Γ / Δ fails in some Σ -pruned L -extensible model of size $2^{O(n^2)}$.
- (ii) If $\Gamma \vdash_L \Delta$, then Γ / Δ holds in all Σ -pruned L -extensible models.

Proof:

(i): Let c be the constant from Corollary 4.16. We may assume $n \geq c$. Since Γ / Δ is not admissible, we may fix a strongly L -extensible model W where it fails.

We first construct a submodel $W' \subseteq W$ by reducing each cluster $C \subseteq W$ as follows:

- If $|C| \leq c$, then $C \subseteq W'$. If $|C| \geq c$, we make sure $|C \cap W'| \geq c$.
- For each formula $\varphi \in \Sigma$ falsified in some point of C , we include at least one such point in W' .

Clearly W' is an L -model, and it has cluster size at most n . The satisfaction of all formulas from Σ is preserved in all points of W' ; in particular, W' refutes Γ / Δ . We claim that W' is a Σ -pruned L -extensible model, by a similar argument as in the proof of Lemma 4.18: if F' is a parametric L -frame with a loosely separated root such that $F' \setminus \text{rcl}(F') \subseteq W'$, let F be $\text{rcl}(F')$ attached below $(F' \setminus \text{rcl}(F')) \uparrow \subseteq W$. By Corollary 4.16, F is still an L -frame, hence $\text{id}_{F \setminus \text{rcl}(F)}$ extends to an isomorphism f of F to a generated subframe of W . Then $F_1 = f^{-1}[f[F] \cap W']$ has the property required by Definition 4.23.

Next, if C and C' are two clusters of W' such that $G := C \uparrow \setminus C = C' \uparrow \setminus C'$, and id_G extends to an isomorphism of $C \uparrow$ to $C' \uparrow$ (including valuation), let W'' be the model obtained from W' by deleting C' and everything below it. Then W'' is an L -model refuting Γ / Δ , and it is still a

Σ -pruned L -extensible model: for any F as in the definition with $F \setminus \text{rcl}(F) \subseteq \cdot W''$, let us find $F_1 \subseteq F$ and $f: F_1 \rightarrow W'$ using Σ -pruned L -extensibility of W' so that the conclusion of the definition holds. If $\text{rcl}(f[F_1]) = C'$, we modify it so that $\text{rcl}(f[F_1]) = C$. Then $f: F_1 \rightarrow W''$, and it still has the required property.

By continuing in this fashion, we obtain a Σ -pruned L -extensible model U of cluster size $\leq n$ such that no two clusters have the same successors and valuations of parameters and variables. We claim that this implies $|U| = 2^{O(n^2)}$. For $i \leq d$, let U_i denote the set of points of depth at most i in U . We will prove by induction that $|U_i| = 2^{O(n^2)}$. We have $U_0 = \emptyset$. For the induction step, each cluster C of $U_{i+1} \setminus U_i$ has $C \uparrow \setminus C$ generated by an antichain in U_i , for which there are $|U_i|^{O(1)} = 2^{O(n^2)}$ choices. Apart from this, C is uniquely determined by its reflexivity, and its multiset of $\leq n$ valuations to $\text{Par} \cup V$, of which there are $\leq (2^{|\text{Par} \cup V|})^{n+1} = 2^{O(n^2)}$ choices. Thus, $|U_{i+1}| = 2^{O(n^2)}$.

(ii): Assume that W_0 is a Σ -pruned L -extensible model refuting Γ / Δ . Put $f_0 = \text{id}_{W_0}$. We will construct a sequence of models $W_0 \subseteq \cdot W_1 \subseteq \cdot \dots \subseteq \cdot W_d$, and maps $f_i: W_i \rightarrow W_0$ such that $f_0 \subseteq f_1 \subseteq \dots \subseteq f_d$, by induction on i . Each f_i will be a depth-preserving p-morphism of the underlying (non-parametric) frames. It will not preserve the valuation as such, but it will preserve the satisfaction of formulas of the form $\Box\varphi$ for $\varphi \in \Sigma$.

Assume that W_i has been constructed. For each parametric L -frame F with a loosely separated root such that $F_0 = F \setminus \text{rcl}(F) \subseteq \cdot W_i$, let F' denote the subframe $f_i[F_0] \subseteq \cdot W_0$ with $\text{rcl}(F)$ attached below. Since f_i is a frame p-morphism, F' is still an L -frame with a loosely separated root, and we may apply Definition 4.23 to find $F'_1 \subseteq F'$ with $f_i[F_0] \subsetneq F'_1$, and a parametric isomorphism $f \supseteq \text{id}_{f_i[F_0]}$ of F'_1 to a generated subframe $f[F'_1] \subseteq \cdot W_0$. Moreover, we can fix a valuation in $\text{rcl}(F')$ so that the same formulas from Σ hold in F' as in $f[F'_1]$. The same formulas from Σ also hold in F under the valuation obtained using $\text{rcl}(F) = \text{rcl}(F')$, as f_i preserves the truth of formulas $\Box\varphi$, $\varphi \in \Sigma$, between F_0 and $f_i[F_0]$. We include $\text{rcl}(F)$ with this valuation in W_{i+1} , and extend f_i to f_{i+1} so that $\text{rcl}(F)$ is mapped onto $\text{rcl}(f[F'_1])$. We carry out this construction for all frames F simultaneously, so that W_{i+1} is W_i together with all the new clusters $\text{rcl}(F)$.

The construction ensures that W_d is a model based on a strongly L -extensible parametric frame. Since f_d preserves the truth of $\Box\varphi$ for $\varphi \in \Gamma$, we have $W_d \models \Gamma$, while $W_d \not\models \varphi$ for any $\varphi \in \Delta$, as $W_0 \subseteq \cdot W_d$. Thus, $\Gamma \not\vdash_L \Delta$ by Theorem 4.12. \square

Theorem 4.25 *Let L be a logic of bounded depth and width. Then L -admissibility is decidable in Π_2^{exp} , and L -unifiability in Σ_2^{exp} .*

Proof: Let us fix a rule Γ / Δ , and put $\Sigma = \text{Sub}(\Gamma \cup \Delta)$. As in Theorem 4.20, we will proceed as if Par were finite of size $\leq n$.

By Lemma 4.24, we have $\Gamma \not\vdash_L \Delta$ iff there exists a Σ -pruned L -extensible model $W \not\models \Gamma / \Delta$ of size $|W| = 2^{O(n^2)}$. This will give a Σ_2^{exp} algorithm for inadmissibility if we show how to check that W is a Σ -pruned L -extensible model in $\text{coNTIME}(2^{O(n^2)})$. We do this by co-nondeterministically choosing F as in Definition 4.23, and then deterministically testing it satisfies the required condition.

We specify F by selecting one of the $|W|^{O(1)}$ choices for $F \setminus \text{rcl}(F) \subseteq \cdot W$, and by indicating

the reflexivity of $\text{rcl}(F)$ as well as the set of assignments $E \subseteq \mathbf{2}^{\text{Par}}$ that are realized in $\text{rcl}(F)$. This can be done with $O(2^n)$ bits, and we can make sure the choice is sound (i.e., that F is an L -frame, and that $|E| = 1$ if $\text{rcl}(F)$ is irreflexive).

In order to test a given F , we go through the clusters C of W such that $C \uparrow \setminus C = F \setminus \text{rcl}(F)$, $\text{refl}(C) = \text{refl}(\text{rcl}(F))$, C is loosely separated, and the set E_1 of assignment to parameters realized in C is a subset of E . Notice that $E_1 \subsetneq E$ may only happen in the reflexive case. Then the condition on extension of valuation amounts to

$$\forall e \in E \setminus E_1 \exists v \in \mathbf{2}^V A(v, e, X) \models \bigwedge_{\varphi \in X} \varphi,$$

where $V = \Sigma \cap \text{Var}$, $X = \{\varphi \in \Sigma : C \models \Box\varphi\}$, and $A(v, e, X)$ is as in the proof of Theorem 4.3. \square

5 Lower bounds

Our lower bounds generally have the form that unifiability in all logics from a certain class (usually determined by the presence of particular frames as subframes) is hard for some complexity class \mathcal{C} , and we will in fact prove more: the reduction showing \mathcal{C} -hardness, and even the unifiers of positive instances, will be independent of the logic. We introduce terminology so that we can state this more concisely.

Definition 5.1 Let \mathcal{L} be a class of logics, and \mathcal{C} a complexity class. We say that *unifiability in \mathcal{L} is uniformly \mathcal{C} -hard*, if for every $A \in \mathcal{C}$, there exists a logspace-computable function $w \mapsto \xi_w$ such that:

- (i) If $w \in A$, there is a substitution σ such that $\vdash_L \sigma(\xi_w)$ for every $L \in \mathcal{L}$.
- (ii) If $w \notin A$, then ξ_w is not L -unifiable for any $L \in \mathcal{L}$.

Let us illustrate the concept on a simple example:

Theorem 5.2 *Parameter-free unifiability in all consistent logics L is uniformly NP-hard.*

Proof: A propositional formula φ with no modal operators or parameters is classically satisfiable if and only if it is L -unifiable: on the one hand, a satisfying assignment provides a $\{\perp, \top\}$ -substitution that makes the formula a classical tautology; on the other hand, if $\vdash_L \sigma(\varphi)$, then the evaluation of σ in any fixed element of a fixed L -model gives a satisfying assignment for φ . \square

5.1 Exponential hierarchy

In this section, we prove hardness of unifiability for several levels of the exponential hierarchy (NEXP, coNEXP, Σ_2^{exp}) for logics admitting suitable patterns as subframes, complementing the matching upper bounds from Section 4. In all three cases, we proceed by reducing the truth of third-order formulas from Theorem 2.3 to unifiability. Generally speaking, we use variables

to encode existential quantifiers, and parameters to encode universal quantifiers; while second-order objects $t \subseteq n$ may be encoded by n atoms in a straightforward way, third-order objects $X \subseteq \mathcal{P}(n)$ require more complex machinery to ensure they are specified consistently, and it is at this point that the frame patterns are needed.

Theorem 5.3 *Unifiability in nonlinear logics is uniformly NEXP-hard.*

Proof: By Theorem 2.3, it is enough to exhibit a logspace reduction from the problem of determining the truth of sentences of the form

$$(22) \quad \Phi = \exists X \subseteq \mathcal{P}(n) \forall t_0, \dots, t_{m-1} \subseteq n \varphi(\dots, i \in t_\alpha, \dots, t_\alpha \in X, \dots),$$

where n is given in unary and φ is a Boolean formula. Given Φ , we define a formula ξ_Φ in variables x, x_α , and parameters $p_i, p_{\alpha,i}, q, r$ for $i < n$ and $\alpha < m$, as the conjunction of the formulas

$$(23) \quad \bigwedge_{i < n} [\Box((q \rightarrow p_{\alpha,i}) \wedge (r \rightarrow p_i)) \vee \Box((q \rightarrow \neg p_{\alpha,i}) \wedge (r \rightarrow \neg p_i))] \\ \rightarrow \Box((q \rightarrow x_\alpha) \wedge (r \rightarrow x)) \vee \Box((q \rightarrow \neg x_\alpha) \wedge (r \rightarrow \neg x))$$

for $\alpha < m$, and

$$(24) \quad \varphi(\dots, p_{\alpha,i}, \dots, x_\alpha, \dots).$$

Clearly, $|\xi_\Phi| = \text{poly}(n, |\Phi|)$, and the mapping $\Phi \mapsto \xi_\Phi$ is logspace-computable.

Claim 5.3.1 *If Φ is true, then ξ_Φ is **K4**-unifiable.*

Proof: Let $W = U_{\mathbf{K4}}(\{p_i, p_{\alpha,i}, q, r : i < n, \alpha < m\})$, and let us fix a witness X for the existential quantifier in (22). By [8, Thm. 5.18], it suffices to construct a valuation of the variables in W that makes ξ_Φ true in W . For any point $u \in W$, we define $t_\alpha^u, t^u \subseteq n$ for $\alpha < m$ by

$$t_\alpha^u = \{i < n : u \models p_{\alpha,i}\}, \\ t^u = \{i < n : u \models p_i\}.$$

Then we define the valuation of variables in u by

$$u \models x_\alpha \quad \text{iff} \quad t_\alpha^u \in X, \\ u \models x \quad \text{iff} \quad t^u \in X.$$

The choice of X ensures that $\varphi(\dots, t_\alpha^u(i), \dots, t_\alpha^u \in X, \dots)$ is true for any $u \in W$, hence (24) is true in u by the definition of t_α^u and of the valuation of x_α .

Concerning (23), let $\alpha < m$, and assume that

$$u \models \Box((q \rightarrow p_{\alpha,i}) \wedge (r \rightarrow p_i)) \vee \Box((q \rightarrow \neg p_{\alpha,i}) \wedge (r \rightarrow \neg p_i))$$

for all $i < n$. This means there exists $t \subseteq n$ such that $t_\alpha^v = t$ for all $v \geq u$ satisfying q , and $t^v = t$ for all $v \geq u$ satisfying r . It follows that

$$u \models \Box((q \rightarrow x_\alpha) \wedge (r \rightarrow x))$$

if $t \in X$, and

$$u \models \Box((q \rightarrow \neg x_\alpha) \wedge (r \rightarrow \neg x))$$

if $t \notin X$.

□ (Claim 5.3.1)

Claim 5.3.2 *If L is a nonlinear logic, and ξ_Φ is L -unifiable, then Φ is true.*

Proof: Let σ be an L -unifier of ξ_Φ . We may assume that the target formulas of σ include no parameters except those occurring in ξ_Φ , and no variables. By assumption, there is an L -frame W that weakly subreduces onto a two-prong fork; that is, we may assume that W is rooted, and we may fix disjoint nonempty admissible subsets $A, B \subseteq W$ such that $A \cap B\downarrow = B \cap A\downarrow = \emptyset$. Put $P = \{p_i : i < n\}$, and $P_\alpha = \{p_{\alpha,i} : i < n\}$ for $\alpha < m$.

For any $t \subseteq n$, let W_t denote the parametric frame based on W where we make $r \wedge P^t$ true in A , and all parameters false elsewhere. In the root of W_t , the premise of (23) (for arbitrary $\alpha < m$) is true, hence $W_t \models \sigma(\xi_\Phi)$ implies

$$W_t \models r \rightarrow \sigma(x) \quad \text{or} \quad W_t \models r \rightarrow \neg \sigma(x).$$

We define a set $X \subseteq \mathcal{P}(n)$ by

$$X = \{t \subseteq n : W_t \models r \rightarrow \sigma(x)\}.$$

We claim that X is a witness for the truth of Φ . Let $t_0, \dots, t_{m-1} \subseteq n$; we need to show that

$$(25) \quad \varphi(\dots, i \in t_\alpha, \dots, t_\alpha \in X, \dots)$$

is true.

Define a new parametric frame $W^{t_0, \dots, t_{m-1}}$ based on W , where $q \wedge \bigwedge_{\alpha < m} P_\alpha^{t_\alpha}$ holds in B , and all parameters are false elsewhere. Since $W^{\vec{t}} \models \sigma(\xi_\Phi)$, we have

$$W^{\vec{t}} \models \varphi(\dots, p_{\alpha,i}, \dots, \sigma(x_\alpha), \dots),$$

and fixing $u \in B$,

$$W^{\vec{t}}, u \models p_{\alpha,i} \quad \text{iff} \quad i \in t_\alpha$$

by the definition of $W^{\vec{t}}$, hence (25) will follow if we establish

$$(26) \quad W^{\vec{t}}, u \models \sigma(x_\alpha) \quad \text{iff} \quad t_\alpha \in X$$

for each $\alpha < m$.

Consider the mixed parametric frame $W_{t_\alpha}^{\vec{t}}$: that is, make $r \wedge P^{t_\alpha}$ true in A , $q \wedge \bigwedge_{\beta < m} P_\beta^{t_\beta}$ true in B , and all parameters false elsewhere. On the one hand, $W_{t_\alpha}^{\vec{t}}$ coincides with $W^{\vec{t}}$ everywhere except in A , and $B\uparrow \cap A = \emptyset$, hence

$$W^{\vec{t}}, u \models \sigma(x_\alpha) \quad \text{iff} \quad W_{t_\alpha}^{\vec{t}}, u \models \sigma(x_\alpha).$$

On the other hand, $W_{t_\alpha}^{\vec{t}}$ coincides with W_{t_α} everywhere except in B , and $A \uparrow \cap B = \emptyset$, hence

$$W_{t_\alpha}^{\vec{t}}, v \models \sigma(x) \quad \text{iff} \quad W_{t_\alpha}, v \models \sigma(x) \quad \text{iff} \quad t_\alpha \in X$$

for any $v \in A$. To connect these, the premise of (23) for α holds in the root of $W_{t_\alpha}^{\vec{t}}$, hence

$$W_{t_\alpha}^{\vec{t}} \models \Box((q \rightarrow \sigma(x_\alpha)) \wedge (r \rightarrow \sigma(x))) \vee \Box((q \rightarrow \neg\sigma(x_\alpha)) \wedge (r \rightarrow \neg\sigma(x))),$$

which implies

$$W_{t_\alpha}^{\vec{t}}, u \models \sigma(x_\alpha) \quad \text{iff} \quad W_{t_\alpha}^{\vec{t}}, v \models \sigma(x).$$

This shows (26). □ (Claim 5.3.2)

The result now follows from Claims 5.3.1 and 5.3.2. □

Corollary 5.4 *Let L be a nonlinear clx logic of bounded cluster size, or a nonlinear tabular logic. Then L -unifiability is NEXP-complete, and L -admissibility is coNEXP-complete.*

Proof: By Theorems 4.2, 4.20, and 5.3. □

Example 5.5 The previous corollary applies to the logics **GL**, **S4Grz**, and **K4Grz**, as well as to their extensions by the **BB_k** axioms for $k \geq 2$.

Theorem 5.6 *Unifiability in logics of unbounded cluster size is uniformly coNEXP-hard.*

Proof: Using Theorem 2.3 again, we will present a logspace reduction from the truth of sentences of the form

$$(27) \quad \Phi = \forall Y \subseteq \mathcal{P}(n) \exists t_0, \dots, t_{m-1} \subseteq n \varphi(\dots, i \in t_\alpha, \dots, t_\alpha \in Y, \dots),$$

where n is given in unary and φ is a Boolean formula. Given Φ , we define a formula ξ_Φ in parameters p_i, q , and variables $z_{\alpha,i}$ for $i < n$ and $\alpha < m$, as the conjunction of the formulas

$$(28) \quad \Diamond q \rightarrow \Diamond(q \wedge \gamma),$$

$$(29) \quad q \wedge \gamma \rightarrow \varphi(\dots, z_{\alpha,i}, \dots, \eta(z_{\alpha,0}, \dots, z_{\alpha,n-1}), \dots),$$

where

$$\gamma = \bigwedge_{\substack{\alpha < m \\ i < n}} (\Box(\Diamond q \rightarrow z_{\alpha,i}) \vee \Box(\Diamond q \rightarrow \neg z_{\alpha,i})),$$

$$\eta(x_0, \dots, x_{n-1}) = \Diamond \left(\neg q \wedge \Diamond q \wedge \bigwedge_{i < n} (p_i \leftrightarrow x_i) \right).$$

The result follows from the next two claims.

Claim 5.6.1 *If Φ is true, then ξ_Φ is **K4**-unifiable.*

Proof: Putting $W = U_{\mathbf{K4}}(\vec{p}, q)$, we will construct a valuation in W that makes ξ_Φ true. Since Φ holds, we may fix for every $Y \subseteq \mathcal{P}(n)$ some sets $t_0^Y, \dots, t_{m-1}^Y \subseteq n$ that make φ true. Write $P = \{p_i : i < n\}$. For each point $u \in W$, we put

$$Y(u) = \{t \subseteq n : u \models \diamond(\neg q \wedge \diamond q \wedge P^t)\},$$

and define

$$u \models z_{\alpha,i} \quad \text{iff} \quad i \in t_\alpha^{Y(u)}.$$

In order to show that (28) holds, assume $u \models \diamond q$. We can find $u < v \models q$ which is \prec -maximal, i.e., $\neg q$ holds in all points strictly above $\text{cl}(v)$. We claim that this makes $v \models \gamma$: indeed, the only points $v' \geq v$ that satisfy $\diamond q$ are $v' \sim v$, and these have $Y(v') = Y(v)$, hence they agree on the satisfaction of all the variables $z_{\alpha,i}$.

As for (29), assume that $u \models q \wedge \gamma$. This means that for each $\alpha < m$, there is $t_\alpha \subseteq n$ such that

$$(30) \quad u \models \Box(\diamond q \rightarrow z_{\alpha,i}^{[i \in t_\alpha]})$$

for all $i < n$. Since $u \models q$, we must have $t_\alpha = t_\alpha^{Y(u)}$, thus

$$\varphi(\dots, i \in t_\alpha, \dots, t_\alpha \in Y(u), \dots)$$

is true. This will imply

$$u \models \varphi(\dots, z_{\alpha,i}, \dots, \eta(z_{\alpha,0}, \dots, z_{\alpha,n-1}), \dots)$$

if we show

$$u \models \eta(z_{\alpha,0}, \dots, z_{\alpha,n-1}) \quad \text{iff} \quad t_\alpha \in Y(u).$$

But (30) ensures that

$$\begin{aligned} u \models \eta(z_{\alpha,0}, \dots, z_{\alpha,n-1}) & \quad \text{iff} \quad u \models \eta([0 \in t_\alpha], \dots, [n-1 \in t_\alpha]) \\ & \quad \text{iff} \quad u \models \diamond(\neg q \wedge \diamond q \wedge P^{t_\alpha}), \end{aligned}$$

which is the definition of $t_\alpha \in Y(u)$. □ (Claim 5.6.1)

Claim 5.6.2 *If ξ_Φ is unifiable in a logic L of cluster size at least $2^n + 1$, then Φ is true.*

Proof: Let σ be an L -unifier of ξ_Φ such that the target formulas of σ include no parameters except those occurring in ξ_Φ , and no variables. Fix $Y \subseteq \mathcal{P}(n)$: we need to find $t_0, \dots, t_{m-1} \subseteq n$ satisfying (27).

Let $\{s_j : j < l\}$ be an enumeration of Y with $l = |Y| \leq 2^n$. By assumption, there exists an L -frame W that subreduces onto the $(l+1)$ -element cluster: i.e., we may fix disjoint nonempty admissible subsets $A_j \subseteq W$, $j \leq l$, such that $A_j \subseteq A_{j'} \downarrow$ for all $j, j' \leq l$. We may assume w.l.o.g. that

$$A_l \downarrow \subseteq \bigcup_{j \leq l} A_j$$

by adding $A_l \downarrow \setminus \bigcup_{j \leq l} A_j$ to A_l if necessary. (Notice that $A_j \downarrow = A_l \downarrow$ for all $j \leq l$.) With this condition, $\{A_j : j \leq l\}$ is in fact a partition of $A_l \downarrow$.

We consider the parametric frame based on W such that A_j satisfies P^{s_j} for $j < l$, A_l satisfies q , and no parameters hold elsewhere. Since $\vdash_L \sigma(\xi_\Phi)$, there exists $u \in A_l$ such that $u \models \sigma(\gamma)$ by (28), thus there are $t_0, \dots, t_{m-1} \subseteq n$ such that

$$u \models \Box(\Diamond q \rightarrow \sigma(z_{\alpha,i})^{[i \in t_\alpha]})$$

for each $\alpha < m$ and $i < n$. As in the previous claim, this implies

$$u \models \sigma(\eta(z_{\alpha,0}, \dots, z_{\alpha,n-1})) \leftrightarrow \Diamond(\neg q \wedge \Diamond q \wedge P^{t_\alpha}).$$

Now,

$$u \models \Diamond(\neg q \wedge \Diamond q \wedge P^{t_\alpha}) \quad \text{iff} \quad t_\alpha \in Y:$$

on the one hand, for each $j < l$, u sees a point $v \in A_j$, and we have $v \models \neg q \wedge \Diamond q \wedge P^{s_j}$; on the other hand, all points satisfying $\neg q \wedge \Diamond q$ are of this form.

Thus,

$$\varphi(\dots, i \in t_\alpha, \dots, t_\alpha \in Y, \dots)$$

follows from $u \models \sigma(29)$.

□ (Claim 5.6.2)

□

Corollary 5.7 *Let L be a linear clx logic of unbounded cluster size, or a linear logic of bounded depth and unbounded cluster size. Then L -unifiability is coNEXP-complete, and L -admissibility is NEXP-complete.*

Proof: By Theorems 4.3, 4.21, and 5.6. □

Example 5.8 The previous corollary applies to **S5** and to the related logics **K45**, **D45**, and **K4B**. It also applies to the logics **S4.3**, **K4.3**, and **D4.3**, and their extensions by the **BD_k** axioms for $k \geq 1$.

For our most complex logics, we will prove a Σ_2^{exp} lower bound. The argument will combine ideas from both Theorems 5.3 and 5.6. First, let us fix a notation for frames that appear in our Σ_2^{exp} -hardness criterion.

Definition 5.9 Let $n \geq 1$. The symbol $(\circ + \textcircled{n})^\circ$ denotes a finite reflexive Kripke frame consisting of a root and two successor clusters, one of which is an n -cluster, and the other one a single point.

Theorem 5.10 *Unifiability is uniformly Σ_2^{exp} -hard in the class of logics $L \supseteq \mathbf{K4}$ such that for every $n \geq 1$, there exists an L -frame weakly subreducing onto $(\circ + \textcircled{n})^\circ$.*

Proof: We use Theorem 2.3 once again. Given a sentence

$$(31) \quad \Phi = \exists X \subseteq \mathcal{P}(n) \forall Y \subseteq \mathcal{P}(n) \exists t_0, \dots, t_{m-1} \subseteq n \varphi(\dots, i \in t_\alpha, \dots, t_\alpha \in X, \dots, t_\alpha \in Y, \dots),$$

where n is written in unary and φ is a Boolean formula, we define a formula ξ_Φ in parameters p_i, q, r , and variables $x, x_\alpha, z_{\alpha,i}$ for $i < n$ and $\alpha < m$, as the conjunction of the formulas

$$(32) \quad \diamond q \rightarrow \diamond(q \wedge \gamma),$$

$$(33) \quad q \wedge \gamma \rightarrow \varphi(\dots, z_{\alpha,i}, \dots, x_\alpha, \dots, \eta(z_{\alpha,0}, \dots, z_{\alpha,n-1}), \dots),$$

and

$$(34) \quad \bigwedge_{i < n} [\Box((q \rightarrow z_{\alpha,i}) \wedge (r \rightarrow p_i)) \vee \Box((q \rightarrow \neg z_{\alpha,i}) \wedge (r \rightarrow \neg p_i))] \\ \rightarrow \Box((q \rightarrow x_\alpha) \wedge (r \rightarrow x)) \vee \Box((q \rightarrow \neg x_\alpha) \wedge (r \rightarrow \neg x))$$

for $\alpha < m$, where

$$\gamma = \bigwedge_{\substack{\alpha < m \\ i < n}} (\Box(\diamond q \rightarrow z_{\alpha,i}) \vee \Box(\diamond q \rightarrow \neg z_{\alpha,i})), \\ \eta(x_0, \dots, x_{n-1}) = \diamond(\neg q \wedge \diamond q \wedge \bigwedge_{i < n} (p_i \leftrightarrow x_i)).$$

Claim 5.10.1 *If Φ is true, then ξ_Φ is **K4**-unifiable.*

Proof: Let $W = U_{\mathbf{K4}}(\vec{p}, q, r)$; we will construct a valuation in W that makes ξ_Φ true. Let X be a witness for the truth of Φ , and for each $Y \subseteq \mathcal{P}(n)$, let us fix sets $t_0^Y, \dots, t_{m-1}^Y \subseteq n$ such that

$$(35) \quad \varphi(\dots, i \in t_\alpha^Y, \dots, t_\alpha^Y \in X, \dots, t_\alpha^Y \in Y, \dots)$$

is true. Let $P = \{p_i : i < n\}$. For each point $u \in W$, we put

$$Y(u) = \{t \subseteq n : u \models \diamond(\neg q \wedge \diamond q \wedge P^t)\},$$

and we define

$$\begin{aligned} u \models x & \quad \text{iff} \quad \text{Sat}_P(u) \in X, \\ u \models x_\alpha & \quad \text{iff} \quad t_\alpha^{Y(u)} \in X, \\ u \models z_{\alpha,i} & \quad \text{iff} \quad i \in t_\alpha^{Y(u)}. \end{aligned}$$

As in the proof of Claim 5.6.1, if u is a \lesssim -maximal point satisfying q , then $u \models \gamma$, which shows that W satisfies (32). Moreover, if u is any point such that $u \models q \wedge \gamma$, then

$$\begin{aligned} u \models \eta(z_{\alpha,0}, \dots, z_{\alpha,n-1}) & \quad \text{iff} \quad u \models \eta([0 \in t_\alpha^{Y(u)}], \dots, [n-1 \in t_\alpha^{Y(u)}]) \\ & \quad \text{iff} \quad u \models \diamond(\neg q \wedge \diamond q \wedge P^{t_\alpha^{Y(u)}}) \\ & \quad \text{iff} \quad t_\alpha^{Y(u)} \in Y(u), \end{aligned}$$

hence (35) (with $Y = Y(u)$) implies

$$u \models \varphi(\dots, z_{\alpha, i}, \dots, x_{\alpha}, \dots, \eta(z_{\alpha, 0}, \dots, z_{\alpha, n-1}), \dots),$$

establishing the truth of (33).

It is also straightforward to check (34): if its premise holds in u , let $t \subseteq n$ be such that

$$u \models \Box((q \rightarrow z_{\alpha, i}^{[i \in t]}) \wedge (r \rightarrow p_i^{[i \in t]}))$$

for all $i < n$. Then all points $v \geq u$ satisfying q have $t_{\alpha}^{Y(v)} = t$, thus

$$v \models x_{\alpha} \quad \text{iff} \quad t \in X,$$

while all points $w \geq u$ satisfying r have $\text{Sat}_P(w) = t$, thus

$$w \models x \quad \text{iff} \quad t \in X$$

again. □ (Claim 5.10.1)

Claim 5.10.2 *If ξ_{Φ} is L -unifiable, and some L -frame weakly subreduces onto $(\circ + \textcircled{L})^{\circ}$ for $l > 2^n$, then Φ is true.*

Proof: Let σ be an L -unifier of ξ_{Φ} such that the target formulas of σ include no parameters except those occurring in ξ_{Φ} , and no variables. By the assumption, we may fix a rooted L -frame W , and disjoint nonempty admissible subsets $A_j, B \subseteq W$, $j \leq 2^n$, such that $B \cap A_j \downarrow = A_j \cap B \downarrow = \emptyset$, and $A_j \subseteq A_{j'} \downarrow$ for $j, j' \leq 2^n$. This implies $A_j \downarrow = A_{j'} \downarrow$ for all $j, j' \leq 2^n$. Moreover, we may assume

$$A_j \downarrow \cap A_{j'} \uparrow \subseteq A_j \downarrow \setminus B \downarrow \subseteq \bigcup_{i \leq 2^n} A_i$$

by adding the excess $A_j \downarrow \setminus (B \downarrow \cup \bigcup_{i \leq 2^n} A_i)$ to A_{2^n} .

In order to verify that Φ is true, we first construct a purported witness $X \subseteq \mathcal{P}(n)$ to the first existential quantifier. For any $t \subseteq n$, let W_t denote the parametric frame based on W such that $r \wedge P^t$ holds in B , and no parameters elsewhere. Then the root of W satisfies the premise of $\sigma(34)$ (for arbitrary $\alpha < m$), hence

$$W_t \models r \rightarrow \sigma(x) \quad \text{or} \quad W_t \models r \rightarrow \neg \sigma(x).$$

We define

$$X = \{t \subseteq n : W_t \models r \rightarrow \sigma(x)\}.$$

Next, let an arbitrary $Y \subseteq \mathcal{P}(n)$ be given; we need to find $t_0, \dots, t_{m-1} \subseteq n$ satisfying φ . Let $l = |Y| \leq 2^n$, and fix an enumeration $Y = \{s_j : j < l\}$. We define a parametric frame W^Y based on W so that $A_j \models P^{s_j}$ for $j < l$, $A_j \models q$ for $l \leq j \leq 2^n$, and no parameters are true elsewhere. Since the root satisfies $\diamond q$, $W^Y \models \sigma(32)$ ensures we can find $u \models q \wedge \sigma(\gamma)$, thus there are $t_0, \dots, t_{m-1} \subseteq n$ such that

$$(36) \quad W^Y, u \models \Box(\diamond q \rightarrow \sigma(z_{\alpha, i})^{[i \in t_{\alpha}]})$$

for all $\alpha < m$ and $i < n$. It remains to prove that

$$\varphi(\dots, i \in t_\alpha, \dots, t_\alpha \in X, \dots, t_\alpha \in Y, \dots)$$

is true. We have

$$W^Y, u \models \varphi(\dots, \sigma(z_{\alpha,i}), \dots, \sigma(x_\alpha), \dots, \eta(\sigma(z_{\alpha,0}), \dots, \sigma(z_{\alpha,n-1})), \dots)$$

by $W^Y \models \sigma(33)$,

$$W^Y, u \models \sigma(z_{\alpha,i}) \quad \text{iff} \quad i \in t_\alpha$$

for all $\alpha < m$ and $i < n$ by (36), and (36) also implies

$$\begin{aligned} W^Y, u \models \eta(\sigma(z_{\alpha,0}), \dots, \sigma(z_{\alpha,n-1})) & \quad \text{iff} \quad W^Y, u \models \diamond(\neg q \wedge \diamond q \wedge P^{t_\alpha}) \\ & \quad \text{iff} \quad t_\alpha \in Y \end{aligned}$$

for all $\alpha < m$ as in the proof of Claim 5.6.2. Thus, to complete the proof, we only need to show

$$(37) \quad W^Y, u \models \sigma(x_\alpha) \quad \text{iff} \quad t_\alpha \in X$$

for all $\alpha < m$.

We proceed by constructing a parametric frame combining W^Y and W_{t_α} , exploiting \vdash_L $\sigma(34)$, but we will need to be more careful than in the proof of Claim 5.3.2, because there may be other points satisfying $q \wedge \sigma(\gamma)$ in the model than u , and they may disagree on (36). Thus, let us consider the parametric frame W' based on W such that $r \wedge P^{t_\alpha}$ holds in B (as in W_{t_α}), P^{s_j} holds in A_j for $j < l$ (as in W^Y), and q holds in the set

$$A' = \left\{ v : W^Y, v \models q \wedge \bigwedge_{\substack{\alpha < m \\ i < n}} \Box(\diamond q \rightarrow \sigma(z_{\alpha,i})^{[i \in t_\alpha]}) \right\}.$$

Notice that $u \in A'$, and in fact W' agrees with W^Y in $u \uparrow$ (using $u \uparrow \cap B = \emptyset$), thus

$$W^Y, u \models \sigma(x_\alpha) \quad \text{iff} \quad W', u \models \sigma(x_\alpha).$$

The definition of W' ensures that the premise of $\sigma(34)$ for α is satisfied in the root, hence

$$W' \models (q \rightarrow \sigma(x_\alpha)) \wedge (r \rightarrow \sigma(x)) \quad \text{or} \quad W' \models (q \rightarrow \neg \sigma(x_\alpha)) \wedge (r \rightarrow \neg \sigma(x)).$$

But W' agrees with W_{t_α} in $B \uparrow$, hence

$$W' \models r \rightarrow \sigma(x) \quad \text{iff} \quad t_\alpha \in X$$

by the definition of X , which implies

$$W', u \models \sigma(x_\alpha) \quad \text{iff} \quad t_\alpha \in X$$

in view of $W', u \models q$. Thus, (37) holds. □ (Claim 5.10.2)

□

Corollary 5.11 *Let L be a nonlinear clx logic of unbounded cluster size, or a nonlinear logic of bounded depth, bounded width, and unbounded cluster size. Then L -unifiability is Σ_2^{EXP} -complete, and L -admissibility is Π_2^{EXP} -complete.*

Proof: By Theorems 4.2, 4.25, and 5.10. □

Example 5.12 The previous corollary applies to the basic transitive logics **K4**, **D4**, and **S4**, their variants with the McKinsey axiom **K4.1**, **D4.1**, **S4.1**, and extensions of these logics with the bounded branching axioms **BB_k** for $k \geq 2$, possibly combined with the bounded depth axioms **BD_l** for $l \geq 2$.

5.2 NEXP lower bounds with $O(1)$ parameters

The lower bounds in the previous subsection relied on the supply of arbitrarily many parameters, and it is clear from the upper bounds in Section 4 that Theorems 5.6 and 5.10 intrinsically need them. This leaves the possibility of an NEXP lower bound on unifiability or inadmissibility with a constant number of parameters. However, Theorems 4.6 and 4.19 show that we need considerably stronger assumptions than in Theorem 5.3 for that to happen: in particular, no fixed finite frame pattern is enough, even allowing for infinite clusters. It is unclear what should be the optimal condition that, on the one hand, guarantees NEXP-hardness of L -unifiability (or at least L -admissibility) with finitely many parameters, and on the other hand applies to a class of logics as wide as possible.

One such lower bound already appeared in [5]: inadmissibility with no parameters at all is NEXP-complete in logics satisfying a certain ad hoc extensibility condition that operates only with frames of depth ≤ 3 , but requires unbounded width at that depth. As a special case, it applies to all clx logics of *unbounded branching*. While, as we mentioned, it is not clear what is the optimal lower bound to expect, we will present variants of this construction that apply to $\langle \bullet, 2 \rangle$ -extensible or $\langle \circ, 2 \rangle$ -extensible logics, including all *nonlinear* clx logics.

We will actually present several lower bounds, trading the number of parameters for the strength of assumptions on the logic. All the statements will be proved by varying the details of the same generic construction. As in Theorem 5.3, the overall strategy is to provide a reduction from the truth problem for sentences of the form

$$(38) \quad \Phi = \exists X \subseteq \mathcal{P}(n) \forall t_0, \dots, t_{m-1} \subseteq n \varphi(\dots, i \in t_\alpha, \dots, t_\alpha \in X, \dots).$$

Assume we have an efficiently computable list of formulas

$$(39) \quad \beta_0, \dots, \beta_{n-1}, \theta_0, \dots, \theta_{m-1}, \theta$$

that do not contain the variable x . (In fact, they will not contain *any* variables in most cases.) We will use points satisfying β_i as a kind of labels for encoding of the sets $t_\alpha \subseteq n$ in (38): a valid encoding of t_α will be a point that satisfies θ_α , and sees β_i for $i \in t_\alpha$. We will also make use of encoding of $t \subseteq n$ without reference to $\alpha < m$, which will employ θ instead of θ_α . The formulas θ_α and θ in turn will be defined, on a case-by-case basis, in terms of seeing points

satisfying certain other formulas γ_α and δ . (The precise meaning of this vague description will be made clear in the specific proofs.)

Given (39), we define a formula ξ_Φ as the conjunction of the two formulas

$$(40) \quad \bigwedge_{i < n} (\Box(\theta \rightarrow \Diamond\beta_i) \vee \Box(\theta \rightarrow \neg\Diamond\beta_i)) \rightarrow \Box(\theta \rightarrow x) \vee \Box(\theta \rightarrow \neg x)$$

and

$$(41) \quad \bigwedge_{\alpha < m} \left(\Diamond\theta_\alpha \wedge \bigwedge_{i < n} (\Box(\theta_\alpha \rightarrow \Diamond\beta_i) \vee \Box(\theta_\alpha \rightarrow \neg\Diamond\beta_i)) \right) \\ \rightarrow \varphi(\dots, \Diamond(\theta_\alpha \wedge \Diamond\beta_i), \dots, \Diamond(\theta_\alpha \wedge x), \dots).$$

Lemma 5.13 *If Φ is true, then ξ_Φ is $\mathbf{K4}$ -unifiable.*

Proof: Let $W = U_{\mathbf{K4}}(\text{Par} \cap \text{Sub}(\xi_\Phi))$; we will construct a valuation in W that makes ξ_Φ true. If there are any variables other than x , we fix their valuation in an arbitrary way. In order to define the valuation of x , we fix a witness $X \subseteq \mathcal{P}(n)$ for Φ . For any $u \in W$, we define

$$t^u = \{i < n : u \models \Diamond\beta_i\},$$

and we put

$$u \models x \quad \text{iff} \quad t^u \in X.$$

It is easy to see that this makes (40) true, as the value of x in any given point is completely determined by the values of $\Diamond\beta_i$ for $i < n$.

As for (41), if its premise holds in u , let $t_0, \dots, t_{m-1} \subseteq n$ be such that

$$u \models \Box(\theta_\alpha \rightarrow (\Diamond\beta_i)^{[i \in t_\alpha]})$$

for each $\alpha < m$ and $i < n$. Since also $u \models \Diamond\theta_\alpha$, this makes

$$u \models \Diamond(\theta_\alpha \wedge \Diamond\beta_i) \quad \text{iff} \quad i \in t_\alpha.$$

Moreover, $t^v = t_\alpha$ for any $v > u$ such that $v \models \theta_\alpha$, hence the definition of $v \models x$ implies

$$u \models \Diamond(\theta_\alpha \wedge x) \quad \text{iff} \quad t_\alpha \in X.$$

Thus,

$$u \models \varphi(\dots, \Diamond(\theta_\alpha \wedge \Diamond\beta_i), \dots, \Diamond(\theta_\alpha \wedge x), \dots)$$

follows from the truth of $\varphi(\dots, i \in t_\alpha, \dots, t_\alpha \in X, \dots)$. \square

Remark 5.14 We can extract from the proof of Lemma 5.13 a completely explicit unifier σ of ξ_Φ , namely

$$(42) \quad \sigma(x) = \bigvee_{t \in X} \bigwedge_{i < n} (\Diamond\beta_i)^{[i \in t]},$$

with $\sigma(y) = y$ for all variables y other than x .




Figure 1: The frame $F \subseteq U_{\mathbf{S4Grz}}(p)$ from the proof of Theorem 5.15. All points are reflexive; here, \bullet denotes points satisfying p , and \circ points satisfying $\neg p$.

We are ready for our first lower bound. Recall Definition 3.11.

Theorem 5.15 *Unifiability of formulas with 1 parameters and 1 variable is uniformly NEXP-hard in the class of logics cofinally subframe-universal for trees.*

Proof: We need to define the formulas (39) in such a way that ξ_Φ is not L -unifiable if Φ is false. The basic idea is that in order to construct a sufficiently large antichain of formulas in one parameter, we identify suitable antichains in $U_{\mathbf{K4}}(p)$ (actually, in $U_{\mathbf{S4GrzBB}_2}(p)$), and describe their points by formulas in a uniform and algorithmically efficient way. This would be easy enough, but there is an additional serious complication that will force us to tread very carefully: our formulas have to work not only in the original subframe $F \subseteq U_{\mathbf{K4}}(p)$, but in an L -frame that just weakly cofinally subreduces to F .

Formally, we will define a sequence of auxiliary formulas $\{\beta_i^d, \eta_i^d : i < n_d\}$ in parameter p by induction on d , and we will simultaneously build a finite parametric frame $F \subseteq U_{\mathbf{S4}}(p)$ with domain $\{a_i^d : d \in \omega, i < n_d\}$ such that β_i^d holds in a_i^d , and η_i^d in $a_i^d \uparrow$.

For $d = 0$, we put $n_0 = 2$, and

$$\beta_i^0 = \eta_i^0 = \Box p^i$$

for $i < 2$. The points $a_0^0, a_1^0 \in F$ are incomparable, and $a_i^0 \models p^i$.

Assume everything has been defined for $2d$. We put $n_{2d+1} = n_{2d}$, and for each $i < n_{2d}$, let

$$\begin{aligned} \eta_i^{2d+1} &= \Box(p \vee \Box \neg p \rightarrow \eta_i^{2d}), \\ \beta_i^{2d+1} &= \eta_i^{2d+1} \wedge \neg p \wedge \Diamond \beta_i^{2d}. \end{aligned}$$

We introduce new points $\{a_i^{2d+1} : i < n_{2d}\}$ in F such that a_i^{2d+1} is an immediate predecessor of a_i^{2d} , and $a_i^{2d+1} \models \neg p$; however, as an exception, a_0^1 will not be a new point: we put $a_0^1 = a_0^0$. (Notice that $\beta_0^1, \eta_0^1, \beta_0^0$, and η_0^0 are all equivalent.)

At the next level, we put $n_{2d+2} = \binom{n_{2d+1}}{2}$. We will identify $i < n_{2d+2}$ with sets $\{j, k\}$ for $j, k < n_{2d}$ (where j may be equal to k) in some canonical way, and we define

$$\begin{aligned} \eta_{\{j,k\}}^{2d+2} &= \Box(p \rightarrow \eta_j^{2d} \vee \eta_k^{2d} \vee (\Diamond \beta_j^{2d+1} \wedge \Diamond \beta_k^{2d+1})) \wedge \Box(\neg p \rightarrow \Box(p \vee \Box \neg p \rightarrow \eta_j^{2d} \vee \eta_k^{2d})), \\ \beta_{\{j,k\}}^{2d+2} &= \eta_{\{j,k\}}^{2d+2} \wedge p \wedge \Diamond \beta_j^{2d+1} \wedge \Diamond \beta_k^{2d+1}. \end{aligned}$$

For each $j, k < n_{2d}$, the point $a_{\{j,k\}}^{2d+2} \in F$ is an immediate predecessor of a_j^{2d+1} and a_k^{2d+1} , and $a_{\{j,k\}}^{2d+2} \models p$.

Thus, we have $a_i^{2d+e} \models p^e$ for $e = 0, 1$, except that $a_0^0 = a_0^1 \models \neg p$.

The topmost part of F is depicted in Fig. 1.

Claim 5.15.1 *For each $d, d' \in \omega$ and $i < n_d, i' < n_{d'}$,*

$$(43) \quad \vdash_{\mathbf{K4}} \beta_i^d \rightarrow \eta_i^d,$$

$$(44) \quad \vdash_{\mathbf{K4}} \eta_i^d \rightarrow \Box \eta_i^d,$$

$$(45) \quad a_{i'}^{d'} \leq a_i^d \implies \vdash_{\mathbf{K4}} \eta_i^d \rightarrow \eta_{i'}^{d'},$$

$$(46) \quad a_{i'}^{d'} \not\leq a_i^d \implies \vdash_{\mathbf{K4}} \neg(\beta_i^d \wedge \eta_{i'}^{d'}).$$

Proof: (43) and (44) are obvious from the definition. As for (45), apart from the exceptional case $a_0^0 \leq a_0^1$ (where it holds), it suffices to prove it when a_i^d is an immediate successor of $a_{i'}^{d'}$, i.e., $d' = d+1$, and $i' = i$ for d even, or $i' = \{i, j\}$ for d odd. In both cases it follows immediately from the definition of $\eta_{i'}^{d'}$, using (44).

We will prove (46) by induction on $d + d'$. We distinguish several cases.

Assume d is odd. If $a_{i'}^{d'} \not\leq a_i^{d-1}$, then $\vdash \beta_i^d \rightarrow \Diamond \beta_i^{d-1}$ by definition, while $\vdash \eta_{i'}^{d'} \rightarrow \Box \neg \beta_i^{d-1}$ by (44) and the induction hypothesis. Otherwise $a_{i'}^{d'} \leq a_i^{d-1}$ and $a_{i'}^{d'} \not\leq a_i^d$, which is only possible if $\langle d', i' \rangle = \langle d-1, i \rangle \neq \langle 0, 0 \rangle$. In particular, d' is even. If $d' = 0$, we thus have $\eta_{i'}^{d'} = \Box p$, while β_i^d implies $\neg p$. If $d' > 0$, let $i = \{j, k\}$. Since β_i^d implies $\neg p$, the definition of $\eta_{i'}^{d'}$ ensures that $\eta_{i'}^{d'} \wedge \beta_i^d$ implies $\Box(p \rightarrow \eta_j^{d'-2} \vee \eta_k^{d'-2})$. Also, $\vdash \beta_i^d \rightarrow \Diamond \beta_i^{d'}$, where $\vdash \beta_i^{d'} \rightarrow p$, hence

$$\vdash \eta_{i'}^{d'} \wedge \beta_i^d \rightarrow \Diamond(\beta_i^{d'} \wedge (\eta_j^{d'-2} \vee \eta_k^{d'-2})).$$

This implies a contradiction by the induction hypothesis, as $a_j^{d'-2} \not\leq a_i^{d'}$ and $a_k^{d'-2} \not\leq a_i^{d'}$.

If d is even and d' is odd, then $\vdash \beta_i^d \rightarrow p \vee \Box \neg p$, hence the definition of $\eta_{i'}^{d'}$ ensures $\vdash \beta_i^d \wedge \eta_{i'}^{d'} \rightarrow \eta_{i'}^{d'-1}$, which leads to a contradiction by the induction hypothesis as $a_{i'}^{d'-1} \not\leq a_i^d$.

The remaining case is if both d and d' are even. First, if one of the immediate successors a_j^{d-1} of a_i^d satisfies $a_{j'}^{d'} \not\leq a_j^{d-1}$, then $\vdash \beta_i^d \rightarrow \Diamond \beta_j^{d-1}$, while $\vdash \eta_{j'}^{d'} \rightarrow \Box \neg \beta_j^{d-1}$ by (44) and the induction hypothesis. Thus, we may assume that all immediate successors of a_i^d are above $a_{i'}^{d'}$.

If $d' = 0$, this forces $d = 0$ as well; since $i \neq i'$, $\beta_i^d \wedge \eta_{i'}^{d'}$ is $\Box p \wedge \Box \neg p$, which is contradictory.

If $d' > 0$, let $i' = \{j', k'\}$. If $\langle d, i \rangle = \langle 0, 0 \rangle$, i.e., $\beta_i^d = \Box \neg p$, the definition of $\eta_{i'}^{d'}$ gives

$$\vdash \beta_i^d \wedge \eta_{i'}^{d'} \rightarrow \eta_{j'}^{d'-2} \vee \eta_{k'}^{d'-2},$$

and the disjuncts $\eta_{j'}^{d'-2}$ and $\eta_{k'}^{d'-2}$ lead to contradiction by the induction hypothesis, as $a_{j'}^{d'-2} \not\leq a_i^d$ and $a_{k'}^{d'-2} \not\leq a_i^d$. Otherwise $\vdash \beta_i^d \rightarrow p$, hence the definition of $\eta_{i'}^{d'}$ gives

$$\vdash \beta_i^d \wedge \eta_{i'}^{d'} \rightarrow \eta_{j'}^{d'-2} \vee \eta_{k'}^{d'-2} \vee (\Diamond \beta_{j'}^{d'-1} \wedge \Diamond \beta_{k'}^{d'-1}),$$

where $\eta_{j'}^{d'-2}$ and $\eta_{k'}^{d'-2}$ lead to contradiction as above, hence

$$(47) \quad \vdash \beta_i^d \wedge \eta_{i'}^{d'} \rightarrow \Diamond \beta_{j'}^{d'-1} \wedge \Diamond \beta_{k'}^{d'-1}.$$

Notice that $\vdash \beta_{j'/k'}^{d'-1} \rightarrow \neg p$. Thus, (47) directly leads to contradiction if $\langle d, i \rangle = \langle 0, 1 \rangle$, as $\beta_i^d = \Box p$. If $d > 0$ and $i = \{j, k\}$, then the definition of β_i^d ensures

$$\vdash \beta_i^d \rightarrow \Box(\beta_{j'/k'}^{d'-1} \rightarrow \Box(p \vee \Box \neg p \rightarrow \eta_j^{d-2} \vee \eta_k^{d-2})).$$

Since also $\vdash \beta_{j'/k'}^{d'-1} \rightarrow \Diamond \beta_{j'/k'}^{d'-2}$, where $\beta_{j'/k'}^{d'-2}$ implies $p \vee \Box \neg p$, (47) yields

$$(48) \quad \vdash \beta_i^d \wedge \eta_{i'}^{d'} \rightarrow \Diamond(\beta_{j'}^{d'-2} \wedge (\eta_j^{d-2} \vee \eta_k^{d-2})) \wedge \Diamond(\beta_{k'}^{d'-2} \wedge (\eta_j^{d-2} \vee \eta_k^{d-2})).$$

Recall that $a_{i'}^{d'} \leq a_j^{d-1}$ and $a_{i'}^{d'} \leq a_k^{d-1}$. This means

$$(49) \quad a_{j'}^{d'-1} \leq a_j^{d-1} \quad \text{or} \quad a_{k'}^{d'-1} \leq a_j^{d-1},$$

and similarly for a_k^{d-1} ; in particular, $d' \geq d$. On the other hand, using the induction hypothesis, (48) gives the desired $\vdash \neg(\beta_i^d \wedge \eta_{i'}^{d'})$ unless

$$(50) \quad a_j^{d-2} \leq a_{j'}^{d'-2} \quad \text{or} \quad a_k^{d-2} \leq a_{j'}^{d'-2},$$

and similarly for $a_{k'}^{d'-2}$. In particular, this implies $d \geq d'$, thus $d = d'$. Then (49) actually means $j \in \{j', k'\}$, and together with the analogous property of k we obtain $\{j, k\} \subseteq \{j', k'\}$, while (50) and its analogue for k' give $\{j', k'\} \subseteq \{j, k\}$. Thus $i = \{j, k\} = \{j', k'\} = i'$, a contradiction. \square (Claim 5.15.1)

Notice that n_{2d} is strictly increasing. Since n_d roughly squares every other step, it is asymptotically doubly exponential in d : $n_d = 2^{2^{\Omega(d)}}$. On the other hand, the size of the formulas β_i^d , η_i^d is $2^{O(d)}$.

For our application, we take the least even D such that $n_D \geq n + m + 2$, and we define the formulas (39) as follows: $\beta_i = \beta_i^D$ for $i < n$, $\gamma_\alpha = \beta_{n+\alpha}^D$ for $\alpha < m$, $\delta = \beta_{n+m}^D$, and

$$(51) \quad \theta = p \wedge \Diamond \delta \wedge \neg \delta,$$

$$(52) \quad \theta_\alpha = p \wedge \Diamond \delta \wedge \Diamond \gamma_\alpha \wedge \bigwedge_{\alpha' \neq \alpha} \neg \Diamond \gamma_{\alpha'}$$

for each $\alpha < m$. We have $D = O(\log \log(n+m))$, thus $|\beta_i^D| = (\log(n+m))^{O(1)}$. It is not difficult to check that the formulas are logspace-constructible. Claim 5.15.1 shows the key property

$$(53) \quad i \neq i' \implies \vdash_{\mathbf{K4}} \beta_i^D \rightarrow \Box \neg \beta_{i'}^D.$$

In particular, θ_α implies θ .

We also need to construct some finite trees (in the form of parametric reflexive frames) related to the representation of Φ by ξ_Φ . We use $+$ or \sum to denote disjoint sum as usual, and F^p (or F^{-p}) denotes F extended below with a new root, satisfying p (or $\neg p$, respectively).

First, for any $d \in \omega$ and $i < n_d$, let T_i^d be $F_{x_i^d}$ unravelled into a tree: that is, $T_i^0 = \emptyset^{p^i}$ (i.e., the trivial tree satisfying p^i) for $i < 2$; T_i^{2d+1} is $(T_i^{2d})^{-p}$ for $i < n_{2d}$, $\langle d, i \rangle \neq \langle 0, 0 \rangle$, while $T_0^1 = T_0^0$; and $T_{\{j,k\}}^{2d+2} = (T_j^{2d+1} + T_k^{2d+1})^p$ (it does not matter if we simplify this to $(T_j^{2d+1})^p$ in the case $j = k$ or not).

We put $T_{\beta_i} = T_i^D$ for $i < n$, $T_{\gamma_\alpha} = T_{n+\alpha}^D$, and $T_\delta = T_{n+m}^D$ in accordance with the definition of the formulas. We also define a “dummy” tree $T_- = T_{n+m+1}^D$.

For every $t \subseteq n$, let $T_t = (T_- + T_- + \sum_{i \in t} T_{\beta_i})^p$. Next, we put $T_{t,\alpha} = (T_\delta + T_{\gamma_\alpha} + T_t)^p$ for each $\alpha < m$. Then, for every $\vec{t} = \langle t_0, \dots, t_{c-1} \rangle \in (\mathcal{P}(n))^m$, we define $T_{\vec{t}} = (\sum_{\alpha < m} T_{t_\alpha, \alpha})^p$. Finally, we pack all this into one huge tree $T = (\sum_{\vec{t}} T_{\vec{t}})^p$.

If T_{\dots} is any of the trees above, let r_{\dots} denote its root; in particular, r denotes the root of T .

We now assume that L is cofinally subframe-universal for trees, thus there exists an L -frame W and a weak cofinal subreduction g from W onto (the underlying frame of) T . We might assume that W is rooted, and its root is mapped to r by f . We adjust the subreduction by putting

$$f(u) = x \quad \text{iff} \quad u \in g^{-1}[x] \downarrow \setminus \bigcup_{\substack{y \in T \\ x \not\leq y}} g^{-1}[y] \downarrow.$$

for all $u \in W$ and $x \in T$. It is easy to see that the purported f -preimages of points of T are disjoint, hence f is a well-defined partial map $W \rightarrow T$. Moreover, $g^{-1}[x] \subseteq f^{-1}[x]$ and $f^{-1}[x] \downarrow = g^{-1}[x] \downarrow$ for each $x \in T$. This implies that f is still a cofinal weak subreduction from W to T . Moreover, the preimages $f^{-1}[x]$ are *convex*: $f^{-1}[x] \uparrow \cap f^{-1}[x] \downarrow \subseteq f^{-1}[x]$. We now define a valuation of the parameter p in W by

$$u \models p \quad \text{iff} \quad u \in \text{dom}(f) \wedge f(u) \models p.$$

Let h be the p -morphism from T to F mapping the root of each embedded copy of T_i^d to x_i^d . Since $\text{dom}(h)$ is a cofinal generated subframe of T , $h \circ f$ is a cofinal weak subreduction from W to F with convex preimages.

Claim 5.15.2 *For any $u \in \text{dom}(h \circ f)$,*

$$h(f(u)) = a_i^d \implies u \models \beta_i^d.$$

Proof: By induction on d . If $d = 0$, then a_i^d is a leaf of F , thus $(h \circ f)[u \uparrow] = \{a_i^d\}$. By cofinality and convexity of $h \circ f$, we have in fact $u \uparrow \subseteq (h \circ f)^{-1}[a_i^d]$. Thus, $u \models \Box p^i$.

Assume the result holds up to $2d$, and let $h(f(u)) = a_i^{2d+1}$. Clearly, $u \models \neg p$, and $u \models \Diamond \beta_i^{2d}$ by the induction hypothesis. In order to check $u \models \eta_i^{2d+1}$, let $u \leq v \models p \vee \Box \neg p$. If $v \models p$, we have $v \in \text{dom}(f)$, thus $h(f(v)) = a_i^{d'} \geq a_i^{2d+1}$; since $h(f(v)) \models p$, actually $a_i^{d'} \geq a_i^{2d}$. By the induction hypothesis, $v \models \beta_{i'}^{d'}$, hence $v \models \eta_i^{2d}$ by Claim 5.15.1. Otherwise, $v \models \Box \neg p$; by cofinality, there is $w \geq v$ such that $w \in \text{dom}(h \circ f)$, and we may choose it so that $h(f(w))$ is a leaf. By the $d = 0$ case, we cannot have $h(f(w)) = a_1^0$, as $w \models \Box \neg p$. Thus, $h(f(w)) = a_0^0$; since $h \circ f$ is a subreduction, $a_i^{2d} \leq a_0^0$, thus $\vdash \Box \neg p \rightarrow \eta_i^{2d}$ by Claim 5.15.1, thus $v \models \eta_i^{2d}$.

Now, let $h(f(u)) = a_{\{j,k\}}^{2d+2}$. Again, $u \models p$ holds by definition, and $u \models \Diamond \beta_j^{2d+1} \wedge \Diamond \beta_k^{2d+1}$ by the induction hypothesis, hence it suffices to show $u \models \eta_{\{j,k\}}^{2d+2}$.

If $u \leq v \models p$, then $v \in \text{dom}(h \circ f)$, and $a_{i'}^{d'} := h(f(v))$ satisfies $a_{i'}^{d'} = a_{\{j,k\}}^{2d+2}$ or $a_{i'}^{d'} \geq a_j^{2d}$ or $a_{i'}^{d'} \geq a_k^{2d}$. In the first case, we have $v \models \Diamond \beta_j^{2d+1} \wedge \Diamond \beta_k^{2d+1}$ by the induction hypothesis; otherwise, the induction hypothesis gives $v \models \beta_{i'}^{d'}$, hence $v \models \eta_j^{2d}$ or $v \models \eta_k^{2d}$ by Claim 5.15.1.

Assume $u \leq v \leq w$, $v \models \neg p$, and $w \models p \vee \Box \neg p$. If $w \models p$, then again $w \in \text{dom}(h \circ f)$, and $a_{i'}^{d'} := h(f(w))$ satisfies $a_{i'}^{d'} = a_{\{j,k\}}^{2d+2}$ or $a_{i'}^{d'} \geq j$ or $a_{i'}^{d'} \geq a_k^{2d}$. The first case is in fact impossible: if $h(f(u)) = h(f(w)) = a_{\{j,k\}}^{2d+2}$, then $h(f(v)) = a_{\{j,k\}}^{2d+2}$ as well by convexity, contradicting $v \not\models p$. Thus, we have $w \models \eta_j^{2d}$ or $w \models \eta_k^{2d}$ as above³. The remaining case $w \models \Box \neg p$ is left to the reader. \square (Claim 5.15.2)

Claim 5.15.2 and (53) imply that each of the formulas $\beta_0, \dots, \beta_{n-1}, \gamma_0, \dots, \gamma_{m-1}, \delta$ holds in W only in points u such that $f(u)$ is the root of a copy of the corresponding tree $T_{\beta_0}, \dots, T_\delta$. Consequently, θ_α holds only in preimages of copies of $r_{t,\alpha}$ for some $t \subseteq n$.

We are getting ready to bring the proof of Theorem 5.15 to an end. Assume that ξ_Φ is L -unifiable, and let σ be its unifier. We may assume $\sigma(x)$ uses no parameters except for p , and no variables.

We need to define a witness X for Φ . For any $t \subseteq n$, we put $t \in X$ if a preimage of some copy of $r_{t,\alpha}$ for some $\alpha < m$ satisfies $\sigma(x)$. Crucially, this definition does not depend on the choices:

Claim 5.15.3 *If $t \subseteq n$, and $u, v \in W$ are such that $f(u)$ is a copy of $r_{t,\alpha}$ and $f(v)$ a copy of $r_{t,\beta}$ for some $\alpha, \beta < m$, then*

$$u \models \sigma(x) \quad \text{iff} \quad v \models \sigma(x).$$

Proof: Let T' be the frame T with valuation of p modified so that p is false everywhere outside copies of $T_{t,\alpha}$ or $T_{t,\beta}$. Let W' be the corresponding modification of W :

$$W', w \models p \quad \text{iff} \quad w \in \text{dom}(f) \text{ and } T', f(w) \models p.$$

In W' , θ holds only in points w such that $f(w)$ is a copy of $r_{t,\alpha}$ or $r_{t,\beta}$. All these points satisfy the same formulas of the form $\Diamond \beta_i$, namely $w \models \Diamond \beta_i$ iff $i \in t$. Thus, the root r of W' satisfies the premise of $\sigma(40)$. This implies

$$W', r \models \Box(\theta \rightarrow \sigma(x)) \quad \text{or} \quad W', r \models \Box(\theta \rightarrow \neg \sigma(x)),$$

and in particular, $u \models \sigma(x)$ iff $v \models \sigma(x)$. \square (Claim 5.15.3)

Now, let $t_0, \dots, t_{m-1} \subseteq n$; we need to show

$$(54) \quad \varphi(\dots, i \in t_\alpha, \dots, t_\alpha \in X, \dots).$$

Let $u \in \text{dom}(f)$ be such that $f(u)$ is a copy of $r_{\bar{t}}$. The only points above u that satisfy θ_α are preimages of the embedded copy of $T_{t_\alpha, \alpha}$, and these satisfy the same formulas of the form $\Diamond \beta_i$ (viz., such that $i \in t_\alpha$). In other words, u satisfies the premise of $\sigma(41)$, hence

$$u \models \varphi(\dots, \Diamond(\theta_\alpha \wedge \Diamond \beta_i), \dots, \Diamond(\theta_\alpha \wedge \sigma(x)), \dots).$$

³It may happen here that $v \notin \text{dom}(f)$, while $v \in (h \circ f)^{-1}[a_j^{2d}]$ and $v \in (h \circ f)^{-1}[a_k^{2d}]$, even if $j \neq k$. Then v satisfies neither η_j^{2d+1} nor η_k^{2d+1} . It is precisely for this reason that we did not simplify the second conjunct of $\eta_{\{j,k\}}^{2d+2}$ to $\Box(\neg p \rightarrow \eta_j^{2d+1} \vee \eta_k^{2d+1})$, as one might be tempted to do just looking at F .

Moreover,

$$\begin{aligned} u \models \diamond(\theta_\alpha \wedge \diamond\beta_i) & \quad \text{iff} \quad i \in t_\alpha, \\ u \models \diamond(\theta_\alpha \wedge \sigma(x)) & \quad \text{iff} \quad t_\alpha \in X, \end{aligned}$$

hence (54) follows. This completes the proof of Theorem 5.15. \square

Corollary 5.16 *Unifiability of formulas with 1 parameters and 1 variable is NEXP-hard for all logics with the disjunction property.*

Proof: By Theorems 5.15 and 3.13. \square

Example 5.17 The cofinality condition in the statement of Theorem 5.15 cannot be dropped. For example, the logic $L = \mathbf{S4.2Grz}$ is subframe-universal for trees, but L -unifiability of formulas with one parameter is NP-complete. The universal frame $U_L(p)$ consists of two disjoint infinite descending chains alternating between p and $\neg p$, one with p on top, and the other with $\neg p$ on top. It is easy to show that a formula $\varphi \in \text{Form}(\{p\}, \text{Var})$ is L -unifiable iff there exists a valuation \models in $U_L(p)$ that makes $U_L(p) \models \varphi$ iff there exists a valuation that makes φ true in the top $2n$ levels of $U_L(p)$, where $n = |\{\psi : \Box\psi \subseteq \varphi\}| + 1$. This condition can be checked in NP.

Nevertheless, the result holds under somewhat weaker assumptions. In particular, if the single-atom fragment of L is included in a logic cofinally subframe-universal for trees, then L -unifiability of formulas in 1 parameter and 1 variable is NEXP-hard. For example, this applies to bounded top-width logics such as $L = \mathbf{S4BTW}_2$: the single-atom fragment of L is included in $\mathbf{S4Grz}$ (which is cofinally subframe-universal for trees), as $U_{\mathbf{S4Grz}}(p)$ has top width 2.

Unification with 1 parameter is also NEXP-hard for logics that are subframe-universal in a way that respects reflexivity of points, see Theorem 5.20.

Moreover, we can easily generalize the lower bound to all subframe-universal logics at the expense of using one more parameter:

Theorem 5.18 *Unifiability of formulas with 2 parameters and 1 variable is uniformly NEXP-hard in the class of logics subframe-universal for trees.*

Proof: We use almost the same construction as in the proof of Theorem 5.15, except that we take another parameter q , and *relativize* the formulas (39) (recall Definition 3.5): instead of β_i , θ_α , and θ , we use $q \wedge \beta_i^q$, $q \wedge \theta_\alpha^q$, and $q \wedge \theta^q$. The proof proceeds as before, but the subreduction f is not necessarily cofinal, and we put

$$W, u \models q \quad \text{iff} \quad u \in \text{dom}(f).$$

The only place where we used the cofinality of f in the proof of Theorem 5.15 was in Claim 5.15.3, to ensure that $(h \circ f)^{-1}[a_1^0] \models \Box p$, and that certain points satisfying $\Box \neg p$ are in $(h \circ f)^{-1}[a_0^0]$. Relativization side-steps both issues: the former property turns into $(h \circ f)^{-1}[a_1^0] \models \Box(q \rightarrow p)$, which follows immediately from $h \circ f$ being a subreduction and the definition of satisfaction of q ; the latter property holds because the points in question will now satisfy q , hence they will be in the domain of $h \circ f$. \square




Figure 2: The subframe of $U_{\mathbf{K4.2GrzBB}_2}(\emptyset)$ used in the proof of Theorem 5.20.

Going in the opposite direction, we may ask if at least one parameter is needed for NEXP-hardness of unification in, say, nonlinear clx logics. In general, the answer is yes: parameter-free unifiability is in NP for a large class of logics, including all extensions of **D4** or **GL**. However, we can prove NEXP-hardness for logics that allow enough mixing of reflexive and irreflexive points. The theorem below gives a typical and reasonably useful sufficient condition, but similar results also holds for some other logics—basically, what is needed is that we can construct embeddings of arbitrarily large trees into $U_L(\emptyset)$ in a uniform and efficient manner.

Using the relativization technique, we can also obtain additional NEXP-hardness results for unifiability with *one* parameter in this way. Specifically, it applies to logics that satisfy a variant of subframe-universality for trees including both reflexive and irreflexive points. As with Definition 3.11, we may equivalently state the definition for trees, for binary trees, or for arbitrary frames without proper clusters; this time we chose the last one, to emphasize that we treat the objects as frames rather than combinatorial shapes, hence their reflexivity matters.

Definition 5.19 A logic $L \supseteq \mathbf{K4}$ is *subframe-universal for simple frames* if for each finite rooted frame F without proper clusters, there exists an L -frame that subreduces onto F ; equivalently, if $\vdash_L R(\varphi)$ implies $\vdash_{\mathbf{K4Grz}} \varphi$ for all formulas φ .

For example, **D4.2** is subframe-universal for simple frames, but not cofinally subframe-universal for trees, and parameter-free **D4.2**-unifiability is in NP.

Theorem 5.20

- (i) *Unifiability of formulas without parameters and with 1 variable is uniformly NEXP-hard in sublogics of $\mathbf{K4.2GrzBB}_2$.*
- (ii) *Unifiability of formulas with 1 parameters and 1 variable is uniformly NEXP-hard in the class of logics subframe-universal for simple frames.*

Proof:

(i): We follow similar strategy as in the proof of Theorem 5.15. We construct β_i as formulas defining points in a certain antichain in $U_L(\emptyset)$; instead of the parameter p , we will distinguish points by their reflexivity.

We will define antichains $\{a_i^d : i < n_d\} \subseteq U_L(\emptyset)$ by induction on d . There is more than one way to do it. For conciseness of notation, if X is a finite subset of $U_L(\emptyset)$, and $*$ \in $\{\bullet, \circ\}$, let $\text{tp}_*(X)$ denote the unique $*$ -tp of $X \uparrow$ in $U_{\mathbf{K4}}(\emptyset)$. We will only use this notation if it is

guaranteed that actually $\text{tp}_*(X) \in U_L(\emptyset)$ (in particular, X is generated by at most 2 points), and if $X \uparrow$ is not rooted with a reflexive root, so that always $\text{tp}_*(X) \notin X \uparrow$.

We put $n_0 = 1$, $n_1 = n_2 = 2$, $n_3 = 4$, and $n_{d+1} = \binom{n_d+1}{2}$ for $d \geq 3$. Let $a_0^0 = \text{tp}_\bullet(\emptyset)$ be the irreflexive leaf of $U_L(\emptyset)$, $a_0^1 = \text{tp}_\circ(a_0^0)$, and $a_1^1 = \text{tp}_\bullet(a_0^0)$. All remaining points will be irreflexive: $a_0^2 = \text{tp}_\bullet(a_0^1, a_1^1)$, $a_1^2 = \text{tp}_\bullet(a_1^1)$, $a_0^3 = \text{tp}_\bullet(a_0^2)$, $a_1^3 = \text{tp}_\bullet(a_0^2, a_1^2)$, $a_2^3 = \text{tp}_\bullet(a_0^1, a_1^2)$, and $a_3^3 = \text{tp}_\bullet(a_1^2)$. For $d \geq 3$, we identify $i < \binom{n_d+1}{2}$ with sets $\{j, k\}$ for $j, k < n_d$ (not necessarily distinct), and we put $a_{\{j,k\}}^{d+1} = \text{tp}_\bullet(a_j^d, a_k^d)$. (See Fig. 2.)

We define β_i^d as the formulas $\beta_{a_i^d}$ from (1). As in the proof of Theorem 5.15, we have $n_d = 2^{2^{\Omega(d)}}$. In order to see that $|\beta_i^d| = 2^{O(d)}$, it is better to present the formulas in a slightly different way. Let $\eta_u = \bigvee_{v \geq u} \beta_v$; we may define β_u and η_u by simultaneous recursion

$$\begin{aligned}\beta_u &= \bigwedge_{i < m_u} \diamond \beta_{u_i} \wedge \square \bigvee_{i < m_u} \eta_{u_i}, \\ \eta_u &= \beta_u \vee \bigvee_{i < m_u} \eta_{u_i},\end{aligned}$$

where $\{u_i : i < m_u\}$ are the immediate successors of u , and u is irreflexive. (We leave the reflexive case to the reader, as with our choice of antichains, it is irrelevant.) Since in our case $m_u \leq 2$, the formulas β_u and η_u are built from $O(1)$ instances of β_v and η_v for v of lower depth, hence the bound $|\beta_i^d| \leq 2^{O(d)}$ follows by induction on d .

As in the proof of Theorem 5.15, we take D minimal such that $n_D \geq n + m + 2$, and define the formulas (39) by $\beta_i = \beta_i^D$ for $i < n$, $\gamma_\alpha = \beta_{n+\alpha}^D$ for $\alpha < m$, $\delta = \beta_{n+m}^D$, and

$$\begin{aligned}\theta &= \diamond \delta \wedge \neg \delta, \\ \theta_\alpha &= \diamond \delta \wedge \diamond \gamma_\alpha \wedge \bigwedge_{\alpha' \neq \alpha} \neg \diamond \gamma_{\alpha'}.\end{aligned}$$

We have $D = O(\log \log(n + m))$, thus $|\beta_i^D| = (\log(n + m))^{O(1)}$, and it is routine to verify that the formulas are logspace-constructible. We denote $a_{\beta_i} = a_i^D$ for $i < n$, $a_{\gamma_\alpha} = a_{n+\alpha}^D$ for $\alpha < m$, $a_\delta = a_{n+m}^D$, and $a_- = a_{n+m+1}^D$.

For each $t \subseteq n$, we define $a_t \in U_L(\emptyset)$ as a predecessor of $\{a_-\} \cup \{a_{\beta_i} : i \in t\}$ arranged in a binary tree of tp 's: that is, if $t = \{i_j : j < l\}$, we put

$$a_t = \text{tp}_\bullet(a_{\beta_{i_{l-1}}}, \text{tp}_\bullet(\dots, \text{tp}_\bullet(a_{\beta_{i_0}}, a_-) \dots)).$$

For any $\alpha < m$, we define

$$a_{t,\alpha} = \text{tp}_\bullet(a_\delta, \text{tp}_\bullet(a_{\gamma_\alpha}, a_t)).$$

Finally, for each $\vec{t} \in (\mathcal{P}(n))^m$, let

$$a_{\vec{t}} = \text{tp}_\bullet(a_{t_{m-1}, m-1}, \text{tp}_\bullet(\dots, \text{tp}_\bullet(a_{t_{1,1}}, a_{t_{0,0}}) \dots)).$$

We have to show that if $\vdash_L \sigma(\xi_\Phi)$, then Φ is true. We may assume that $\sigma(x) \in \text{Form}(\emptyset, \emptyset)$. We define a set $X \subseteq \mathcal{P}(n)$ by

$$t \in X \quad \text{iff} \quad a_{t,\alpha} \models \sigma(x)$$

for any $t \subseteq n$ and $\alpha < m$. This definition is independent of the choice of α : putting $u = \text{tp}_\bullet(a_{t,\alpha}, a_{t,\alpha'})$, the only element of $u\uparrow$ that satisfies δ is a_δ , hence θ only holds in the three points $u, a_{t,\alpha}, a_{t,\alpha'}$. These all see a_{β_i} for $i \in t$, and for no other $i < n$. Thus, u satisfies the premise of $\sigma(40)$, hence its conclusion, which gives

$$a_{t,\alpha} \models \sigma(x) \quad \text{iff} \quad a_{t,\alpha'} \models \sigma(x).$$

To check that X is a witness for Φ , let $\vec{t} \in (\mathcal{P}(n))^m$. As above, it is easy to see that θ_α is satisfied in $a_{\vec{t}\uparrow}$ only in the point $a_{t_\alpha,\alpha}$, hence $a_{\vec{t}}$ satisfies the premise of $\sigma(41)$, hence

$$a_{\vec{t}} \models \varphi(\dots, \diamond(\theta_\alpha \wedge \diamond\beta_i), \dots, \diamond(\theta_\alpha \wedge \sigma(x)), \dots).$$

Moreover,

$$\begin{aligned} a_{\vec{t}} \models \diamond(\theta_\alpha \wedge \diamond\beta_i) & \quad \text{iff} \quad i \in t_\alpha, \\ a_{\vec{t}} \models \diamond(\theta_\alpha \wedge \sigma(x)) & \quad \text{iff} \quad t_\alpha \in X, \end{aligned}$$

hence (54) is true.

(ii): We relativize the formulas used in (i), i.e., we define $\beta_i = p \wedge (\beta_i^D)^p$, $\gamma_\alpha = p \wedge (\beta_{n+\alpha}^D)^p$, $\delta = p \wedge (\beta_{n+m}^D)^p$, and we define θ, θ_α as in (51) and (52).

Continuing the argument from (i), let $a \in U_{\mathbf{K4}}(\emptyset)$ be a common predecessor of all the points $a_{\vec{t}}$ for $\vec{t} \in (\mathcal{P}(n))^m$. By assumption, there exists an L -frame W , and a subreduction $f: W \rightarrow a\uparrow$. We make W into a parametric frame by putting

$$u \models p \quad \text{iff} \quad u \in \text{dom}(f).$$

Recall that by Lemma 3.9, we have

$$u \models \varphi^p \quad \text{iff} \quad f(u) \models \varphi$$

for all $u \in \text{dom}(f)$ and $\varphi \in \text{Form}(\emptyset, \emptyset)$.

Let σ be an L -unifier of ξ_Φ such that $\sigma(x) \in \text{Form}(p, \emptyset)$. We define $X \subseteq \mathcal{P}(n)$ by

$$(55) \quad t \in X \quad \text{iff} \quad \exists u \in W \exists \alpha < m (f(u) = a_{t,\alpha} \wedge u \models \sigma(x)).$$

Using (40) and a modified valuation that makes p true only in $f^{-1}[a_{t,\alpha}\uparrow \cup a_{t,\alpha'}\uparrow]$, we may check that the definition of X does not depend on the choice of α or u in (55).

Then, for any $\vec{t} \in (\mathcal{P}(n))^m$, we take $u \in f^{-1}[a_{\vec{t}}]$, and using (41) we obtain

$$u \models \varphi(\dots, \diamond(\theta_\alpha \wedge \diamond\beta_i), \dots, \diamond(\theta_\alpha \wedge \sigma(x)), \dots),$$

from which (54) follows as usual. \square

As we already mentioned, parameter-free unification is in NP for a large number of logics of interest. One way to make it harder is to use parameters, but another way is to consider *admissibility* rather than unification: parameter-free admissibility is coNEXP-hard for typical logics of unbounded branching, as shown in [5]. We will generalize this result to a class of logics

that includes all $\langle \circ, 2 \rangle$ -extensible or $\langle \bullet, 2 \rangle$ -extensible logics, and in particular, all nonlinear clx logics.

While we can prove hardness of unification using “static” conditions requiring the presence of suitable subframes, it seems that hardness of admissibility requires “dynamic” conditions expressing closure properties of the class of L -frames. Thus, the condition we employ is a little peculiar to formulate:

Definition 5.21 Let T be a (finite) tree. A logic $L \supseteq \mathbf{K4}$ has the *weak T -extension property* if for every labelling of the leaves of T by finite clusters that are L -frames, there exists an L -frame W with skeleton T such that the final clusters of W are isomorphic to their labels in T .

If \mathcal{T} is a family of trees, L has the *weak extension property wrt \mathcal{T}* if it has the weak T -extension property for each $T \in \mathcal{T}$.

A family of trees \mathcal{T} is *depth-3-universal* if every tree of depth 3 embeds (as a poset) in a tree $T \in \mathcal{T}$.

Example 5.22 The *depth-3 weak extension property*, proved to imply coNEXP-hardness of parameter-free admissibility in [5, Thm. 4.13], is nothing else than the weak extension property wrt the family of all trees of depth 3.

The family of all depth-3 trees is trivially depth-3-universal. The family of all binary trees is also depth-3-universal.

Recall that a *caterpillar* is a tree that becomes a single path if we remove all leaves; a binary caterpillar tree thus consists of a central path $x_0 < x_1 < \dots < x_n$, with at most one additional leaf vertex attached to each x_i . Let us define a *squared caterpillar* to be a tree obtained from a caterpillar by replacing each of its leaves with a caterpillar. Then the family of binary squared caterpillars is depth-3-universal.

Theorem 5.23 *Single-conclusion admissibility without parameters is uniformly coNEXP-hard in the class of all logics satisfying the weak extension property wrt a depth-3-universal family of trees (including all $\langle \circ, 2 \rangle$ -extensible or $\langle \bullet, 2 \rangle$ -extensible logics).*

Proof: Given a sentence Φ as in (38), we define formulas

$$\begin{aligned} \beta_i &= \Box \left(z_i \wedge \bigwedge_{\substack{i' < n+m+4 \\ i' \neq i}} \neg z_{i'} \right), & i < n+m+4, \\ \gamma_\alpha &= \beta_{n+\alpha}, & \alpha < m, \\ \delta_p &= \beta_{n+m+p}, & p < 4, \\ \theta &= \Diamond \delta_1 \wedge \Diamond \delta_2 \wedge \Diamond \delta_3 \wedge \neg \Diamond \delta_0, \\ \theta_\alpha &= \theta \wedge \Diamond \gamma_\alpha \wedge \bigwedge_{\alpha' \neq \alpha} \neg \Diamond \gamma_{\alpha'}, & \alpha < m, \\ \zeta_\Phi &= \bigvee_{i < n+m+4} \Box \neg \beta_i \end{aligned}$$

using variables $\{z_i : i < n + m + 4\}$, and we define ξ_Φ as (40) \wedge (41). We aim to prove

$$(56) \quad \Phi \text{ is true} \quad \text{iff} \quad \xi_\Phi \not\vdash_L \zeta_\Phi.$$

For the left-to-right implication, Lemma 5.13 and Remark 5.14 show that there is a **K4**-unifier σ of ξ_Φ such that $\sigma(z_i) = z_i$ for all $i < n + m + 4$, thus $\sigma(\zeta_\Phi) = \zeta_\Phi$. It remains to observe that $\not\vdash_L \zeta_\Phi$: it suffices to take a finite rooted L -frame with $\geq n + m + 4$ final clusters, where we make z_i (and consequently, β_i) true only in the i th cluster; the existence of such an L -frame follows from the weak extension property.

For the right-to-left implication, let σ be an L -unifier of ξ_Φ such that $\not\vdash_L \sigma(\zeta_\Phi)$. For each $i < n + m + 4$, $\sigma(\beta_i)$ is satisfiable in a finitely generated descriptive L -frame W_i . Since it has the form $\Box(\dots)$, it holds on a cone; using descriptivity, this cone contains a final cluster, which is finite due to W_i being finitely generated. Thus, we may fix a finite single-cluster model $B_i \models \sigma(\beta_i)$ based on an L -frame. We will write $C_\alpha = B_{n+\alpha}$ for $\alpha < m$, and $D_p = B_{n+m+p}$ for $p < 4$.

At least two of the clusters D_1, D_2, D_3 have the same reflexivity; w.l.o.g. $\text{refl}(D_1) = \text{refl}(D_2)$. We may also assume $|D_1| = |D_2|$ by replicating some point from the smaller cluster if necessary.

Let S be the depth-3 tree consisting of a root with $m2^{n+1}$ immediate successors, each of which has $n + 4$ successors. By assumption, S embeds in a tree T such that L has the weak T -extension property. This means that T includes an antichain A of size $m2^{n+1}$ such that each element $a \in A$ sees at least $n+4$ leaves. If an element $a \in A$ has only one immediate successor a' , we may replace a with a' while keeping the number of leaves. Thus, we may assume w.l.o.g. that every $a \in A$ splits. Going in the opposite direction, let $a^- \leq a$ be the minimal element below a such that no element of the chain $[a^-, a]$ splits except a .

We can find a subset $A_0 \subseteq A$ of size $\geq m2^n$ such that for each $a \in A_0$, a^- has a sibling that sees a leaf in $T \setminus A_0 \uparrow$. Indeed, let \sim be the equivalence relation on A defined by $a_0 \sim a_1$ iff a_0^- is a sibling of a_1^- , and let A_0 be constructed from A by removing one element from each \sim -equivalence class of size ≥ 2 . Since we removed at most half of the elements, we have $|A_0| \geq |A|/2 = m2^n$, and it is easy to see that A_0 has the required property. We name $m2^n$ elements of A_0 as $\{r_{t,\alpha} : t \subseteq n, \alpha < m\}$.

We will label each leaf of T with one of the clusters B_i , $i < n + m + 4$, as follows. Let $t \subseteq n$ and $\alpha < m$. There are $\geq n + 4$ leaves in $r_{t,\alpha} \uparrow$. For each $p = 1, 2, 3$, we choose one leaf above $r_{t,\alpha}$ and label it with D_p ; we do it in such a way that the leaves labelled with D_1 and D_2 are above different immediate successors of $r_{t,\alpha}$, thus no element strictly above $r_{t,\alpha}$ sees both. For each $i \in t$, we label one leaf above $r_{t,\alpha}$ with B_i . We label the remaining leaves above $r_{t,\alpha}$ with C_α . Finally, we label every leaf of T that is not above any $r_{t,\alpha}$ with D_0 .

By the weak T -extension property, there exists an L -frame W obtained from T by choosing the reflexivity of each inner node, and replacing each leaf with (the underlying frame of) its label. We will identify the corresponding inner nodes of T and W . We consider W not just as a frame, but also as a model: the valuation of variables in leaf clusters is taken from the corresponding B_i models, and the valuation in inner nodes is arbitrary (say, all variables are false).

Notice that $\vdash_{\mathbf{K4}} \beta_i \rightarrow \Box \neg \beta_{i'}$ for $i \neq i'$. Each leaf cluster of W is labelled with some B_i , and

it then satisfies $\sigma(\beta_i)$, and $\neg\sigma(\beta_{i'})$ for all $i' \neq i$. An inner node of W that sees leaves with two different labels cannot satisfy any $\sigma(\beta_i)$.

It follows that for any $u \in W$,

$$(57) \quad W, u \models \sigma(\theta) \quad \text{iff} \quad \exists t \subseteq n \exists \alpha < m (r_{t,\alpha}^- \leq u \leq r_{t,\alpha}),$$

and for each $\alpha < m$,

$$(58) \quad W, u \models \sigma(\theta_\alpha) \quad \text{iff} \quad \exists t \subseteq n (r_{t,\alpha}^- \leq u \leq r_{t,\alpha}).$$

Indeed, the right-to-left implications follow easily from the definition. On the other hand, if u is not above any $r_{t,\alpha'}^-$, the construction of A_0 ensures that u sees a leaf labelled D_0 , hence $u \models \sigma(\diamond\delta_0)$, which implies $u \not\models \sigma(\theta)$. If $u \not\geq r_{t,\tilde{\alpha}}$ for some $t \subseteq n$ and $\tilde{\alpha} < m$, then $u \not\models \sigma(\diamond\delta_1 \wedge \diamond\delta_2)$, hence $u \not\models \sigma(\theta)$ again. And if $r_{t,\alpha'}^- \leq u \leq r_{t,\alpha'}$ for some t and $\alpha' \neq \alpha$, then $u \models \diamond\gamma_{\alpha'}$, which implies $u \not\models \sigma(\theta_\alpha)$.

Claim 5.23.1 *Let $t \subseteq n$, $\alpha, \alpha' < m$, and $u, u' \in W$ be such that $r_{t,\alpha}^- \leq u \leq r_{t,\alpha}$ and $r_{t,\alpha'}^- \leq u' \leq r_{t,\alpha'}$. Then*

$$(59) \quad W, u \models \sigma(x) \quad \text{iff} \quad W, u' \models \sigma(x).$$

Proof: We define a modified valuation \models^t in W as follows: for any leaf cluster labelled D_1 above some $r_{\tilde{t},\tilde{\alpha}}$, where $\tilde{t} \neq t$, we change the valuation to match D_2 . (This is possible because D_1 and D_2 have isomorphic underlying frames.) This ensures that no point $v \geq r_{\tilde{t},\tilde{\alpha}}^-$ satisfies $\sigma(\diamond\delta_1)$, and consequently, $v \not\models \sigma(\theta)$. Thus, the only points that satisfy $\sigma(\theta)$ in $\langle W, \models^t \rangle$ are the elements of $\bigcup_{\tilde{\alpha} < c} [r_{t,\tilde{\alpha}}^-, r_{t,\tilde{\alpha}}]$. All these points satisfy the same formulas of the form $\sigma(\diamond\beta_i)$, namely those such that $i \in t$.

This means $\langle W, \models^t \rangle$ satisfies the premise of $\sigma(40)$ in the root, hence

$$W \models^t \Box(\sigma(\theta) \rightarrow \sigma(x)) \quad \text{or} \quad W \models^t \Box(\sigma(\theta) \rightarrow \neg\sigma(x)).$$

In particular, (59) follows, as the valuation in $r_{t,\alpha}^- \uparrow \cup r_{t,\alpha'}^- \uparrow$ is unchanged, and $u, u' \models \sigma(\theta)$. □ (Claim 5.23.1)

In view of Claim 5.23.1, we may define a set $X \subseteq \mathbf{2}^n$ such that

$$r_{t,\alpha}^- \leq u \leq r_{t,\alpha} \implies (u \models \sigma(x) \quad \text{iff} \quad t \in X)$$

for all $t \subseteq n$, $\alpha < m$, and $u \in W$. We claim that X is a witness for the truth of Φ . Let $\vec{t} \in (\mathcal{P}(n))^m$, we need to show that

$$(60) \quad \varphi(\dots, i \in t_\alpha, \dots, t_\alpha \in X, \dots)$$

holds. Similarly to the proof of Claim 5.23.1, we consider a modified valuation $\models^{\vec{t}}$ where we “disable” D_1 leaf clusters that are not above $r_{t_\alpha,\alpha}$ for any $\alpha < m$ by turning them into a D_2 .

The new valuation satisfies $\sigma(\theta_\alpha)$ only in the interval $[r_{t_\alpha, \alpha}^-, r_{t_\alpha, \alpha}]$, hence the premise of $\sigma(41)$ holds in the root r of W , hence also

$$W, r \models^{\vec{t}} \varphi(\dots, \sigma(\diamond(\theta_\alpha \wedge \diamond\beta_i)), \dots, \sigma(\diamond(\theta_\alpha \wedge x)), \dots).$$

Moreover, we have

$$\begin{aligned} r \models^{\vec{t}} \sigma(\diamond(\theta_\alpha \wedge \diamond\beta_i)) & \quad \text{iff } i \in t_\alpha, \\ r \models^{\vec{t}} \sigma(\diamond(\theta_\alpha \wedge x)) & \quad \text{iff } t_\alpha \in X, \end{aligned}$$

hence (60) follows. \square

Remark 5.24 We could slightly relax the weak T -extension property: it is enough to ask that there exists an L -frame W , and a weak p-morphism $f: W \rightarrow T$ (i.e., a p-morphism $f: W^\circ \rightarrow T$) such that for each leaf $a \in T$ with label C , there is a p-morphism $f_C: f^{-1}[a] \rightarrow C$. Since the class of L -frames is stable under p-morphic images, this is almost equivalent to Definition 5.21; the only difference arises because a single node of T may have both reflexive and irreflexive preimages in W , hence it may not be possible to choose reflexivity of inner nodes of T in a consistent way so that f becomes a true p-morphism from W . However, we do not know of any interesting application where this would make a difference, hence we prefer the simpler formulation of Definition 5.21.

5.3 PSPACE and below

Theorem 5.25

- (i) For any $d > 0$, unifiability in logics of depth at least d is uniformly Π_{2d}^p -hard.
- (ii) Unifiability of formulas with 2 parameters and 1 variable in logics of unbounded depth is uniformly PSPACE-hard.

Proof: Let us define auxiliary formulas

$$\begin{aligned} \delta_{0,e} &= q^e, & \delta_{i+1,e} &= q^e \wedge \diamond\delta_{i,1-e}, \\ \theta_{i,e} &= \delta_{i,e} \wedge \neg\delta_{i+1,e} = \delta_{i,e} \wedge \Box\neg\delta_{i,1-e}, & \theta_i &= \theta_{i,i \bmod 2} \end{aligned}$$

for $i \in \omega$ and $e \in \mathbf{2}$. The meaning is that $w \models \delta_{i,e}$ iff there is a chain $w = w_i < w_{i-1} < \dots < w_0$ where q is alternately true and false, with it being true (false) in w_i if $e = 1$ (0, resp.), and $w \models \theta_{i,e}$ if, in addition, there is no longer chain. The formula θ_i is normalized so that the chain has $w_0 \not\models q$ on top. (If $w \models q^e$ sees a proper cluster in which both q and $\neg q$ are realized, then $w \models \delta_{i,e}$ and $w \not\models \theta_{i,e}$ for all $i \in \omega$.)

In particular, $w \models \theta_i$ (or $w \models \delta_{i,e}$) implies that w has depth $> i$, and in a model where q holds exactly in points of even depth, we have $w \models \theta_i$ iff the depth of w is $i + 1$. It follows from the description that

$$(61) \quad \vdash_{\mathbf{K4}} \theta_i \rightarrow \neg\theta_j \quad \text{for } i \neq j,$$

$$(62) \quad \vdash_{\mathbf{K4}} \theta_i \rightarrow \diamond\theta_j \quad \text{for } i > j.$$

Let

$$(63) \quad \Phi = \forall \vec{a}_0 \exists \vec{b}_0 \dots \forall \vec{a}_{d-1} \exists \vec{b}_{d-1} \varphi(\vec{a}_0, \dots, \vec{a}_{d-1}, \vec{b}_0, \dots, \vec{b}_{d-1})$$

be a given quantified Boolean sentence with φ quantifier-free, where each \vec{a}_i stands for an m -tuple $a_{i,0}, \dots, a_{i,m-1}$, and similarly for \vec{b}_i . We define a formula ξ_Φ in $m+1$ parameters p_0, \dots, p_{m-1}, q , and m variables x_0, \dots, x_{m-1} , as the conjunction of the formulas

$$(64) \quad \gamma \rightarrow \bigwedge_{i < d} \bigwedge_{j < m} (\Box(\theta_i \rightarrow x_j) \vee \Box(\theta_i \rightarrow \neg x_j)),$$

$$(65) \quad \gamma \wedge \theta_{d-1} \rightarrow \varphi(\dots, a_{i,j} / \Diamond(\theta_i \wedge p_j), \dots, b_{i,j} / \Diamond(\theta_i \wedge x_j), \dots),$$

where

$$\gamma = \bigwedge_{i < d} \bigwedge_{j < m} (\Box(\theta_i \rightarrow p_j) \vee \Box(\theta_i \rightarrow \neg p_j)).$$

Notice that $|\theta_i| = O(i)$, hence $|\xi_\Phi| = O(d^2 m + d|\varphi|) = O(|\Phi|^2)$; it is also easy to see that the mapping $\Phi \mapsto \xi_\Phi$ is logspace-computable.

Claim 5.25.1 *If Φ is true, then ξ_Φ is **K4**-unifiable.*

Proof: For $i \leq d$, denote

$$(66) \quad \Phi_i(\vec{a}_0, \dots, \vec{a}_{i-1}, \vec{b}_0, \dots, \vec{b}_{i-1}) = \forall \vec{a}_i \exists \vec{b}_i \dots \forall \vec{a}_{d-1} \exists \vec{b}_{d-1} \varphi(\vec{a}_0, \dots, \vec{a}_{d-1}, \vec{b}_0, \dots, \vec{b}_{d-1}).$$

Note that $\Phi_0 = \Phi$ and $\Phi_d = \varphi$. For every $i < d$ and $\vec{a}_0, \dots, \vec{a}_i, \vec{b}_0, \dots, \vec{b}_{i-1} \in \mathbf{2}^m$ such that $\exists \vec{b}_i \Phi_{i+1}(\vec{a}_0, \dots, \vec{a}_i, \vec{b}_0, \dots, \vec{b}_i)$ is true, fix a witness $F_i(\vec{a}_0, \dots, \vec{a}_i, \vec{b}_0, \dots, \vec{b}_{i-1}) \in \mathbf{2}^m$ for the $\exists \vec{b}_i$ quantifier block. Also, define

$$G_i(\vec{a}_0, \dots, \vec{a}_i) = F_i(\vec{a}_0, \dots, \vec{a}_i, G_0(\vec{a}_0), \dots, G_{i-1}(\vec{a}_0, \dots, \vec{a}_{i-1}))$$

by induction on $i < d$. Since Φ_0 is true, we can show by induction on i that

$$(67) \quad \Phi_i(\vec{a}_0, \dots, \vec{a}_{i-1}, G_0(\vec{a}_0), \dots, G_{i-1}(\vec{a}_0, \dots, \vec{a}_{i-1}))$$

is true for every $i \leq d$ and every $\vec{a}_0, \dots, \vec{a}_{i-1} \in \mathbf{2}^m$.

Let $W = U_{\mathbf{K4}}(\vec{p}, q)$; by [8, Thm. 5.18], it suffices to construct a valuation of \vec{x} in W that makes $W \models \xi_\Phi$. Let $w \in W$. If $w \models \theta_i \wedge \gamma$ for some $i < d$ (which is unique by (61)), let $\vec{a}_0(w), \dots, \vec{a}_i(w) \in \mathbf{2}^m$ be the unique Boolean vectors such that $v \models p_j^{a_{i',j}(w)}$ for every $i' \leq i$, $j < m$, and $v \geq w$ such that $v \models \theta_{i'}$: their existence and uniqueness follows from (62) and $w \models \gamma$. Putting $\vec{b}_i(w) = G_i(\vec{a}_0(w), \dots, \vec{a}_i(w))$, we define valuation of \vec{x} in w so that $w \models x_j^{b_{i,j}(w)}$ for each $j < m$. If $w \not\models \gamma$, or $w \not\models \theta_i$ for any $i < d$, we define $w \models x_j$ arbitrarily.

Assume that $w \models \gamma$, and fix $i < d$. Then every $v, v' \geq w$ satisfying θ_i also satisfy γ , and we have $\vec{a}_{i'}(v) = \vec{a}_{i'}(v')$ for every $i' \leq i$, hence v and v' agree on the satisfaction of \vec{x} . It follows that $w \models \Box(\theta_i \rightarrow x_j) \vee \Box(\theta_i \rightarrow \neg x_j)$ for each $j < m$. As i was arbitrary, (64) is true in W .

By the same argument, if $w \models \theta_{d-1} \wedge \gamma$, then

$$\begin{aligned} w \models \Diamond(\theta_i \wedge p_j) & \text{ iff } a_{i,j}(w) = 1, \\ w \models \Diamond(\theta_i \wedge x_j) & \text{ iff } b_{i,j}(w) = 1, \end{aligned}$$

for each $i < d$ and $j < m$. Since $\Phi_d = \varphi$, we have $\varphi(\vec{a}_0(w), \dots, \vec{a}_{d-1}(w), \vec{b}_0(w), \dots, \vec{b}_{d-1}(w))$ by (67) and the choice of $\vec{b}_i(w)$, thus (65) holds in W as well. \square (Claim 5.25.1)

Claim 5.25.2 *If a logic L has depth at least d , and ξ_Φ is L -unifiable, then Φ is true.*

Proof: By the assumption, there exists a finite L -frame F of depth d with root r . Let σ be an L -unifier of ξ_Φ , w.l.o.g. chosen so that its target formulas contain no variables, and only the parameters \vec{p}, q . Assume for contradiction that Φ is false. Put $w \models q$ iff $w \in F$ has even depth, so that $w \models \theta_i$ iff w has depth $i + 1$.

By induction on $i \leq d$, we will fix valuation of \vec{p} in points of F of depth at most i , and $\vec{a}_0, \dots, \vec{a}_{i-1}, \vec{b}_0, \dots, \vec{b}_{i-1} \in \mathbf{2}^m$ such that $\Phi_i(\vec{a}_0, \dots, \vec{a}_{i-1}, \vec{b}_0, \dots, \vec{b}_{i-1})$ is false. The valuation will be defined so that $w \models p_j^{a_{i',j}}$ and $w \models \sigma(x_j)^{b_{i',j}}$ for every $w \in F$ of depth $i' + 1$, where $i' < i$; notice that $\sigma(x_j)$ is a formula in the parameters \vec{p}, q , and its value in w is only affected by the valuation of p_j in points of depth $i' + 1 \leq i$ or less, hence it will not change in later steps.

The base step $i = 0$ is trivial—we only observe that Φ_0 is false by assumption. Assume the construction has been carried out up to $i < d$. Since $\Phi_i(\vec{a}_0, \dots, \vec{a}_{i-1}, \vec{b}_0, \dots, \vec{b}_{i-1})$ is false, we can fix $\vec{a}_i \in \mathbf{2}^m$ such that

$$(68) \quad \forall \vec{b}_i \neg \Phi_{i+1}(\vec{a}_0, \dots, \vec{a}_i, \vec{b}_0, \dots, \vec{b}_i).$$

Define $w \models p_j^{a_{i,j}}$ for all $j < m$ and $w \in F$ of depth $i + 1$, and temporarily put (say) $w \not\models p_j$ for all w of depth more than $i + 1$. By the definition and the induction hypothesis, this makes γ true in r . Since $\vdash_L \sigma(\xi_\Phi)$, we have $r \models \Box(\theta_i \rightarrow \sigma(x_j)) \vee \Box(\theta_i \rightarrow \neg\sigma(x_j))$. Thus, there exists a unique $\vec{b}_i \in \mathbf{2}^m$ such that $w \models \sigma(x_j)^{b_{i,j}}$ for every w of depth $i + 1$ and $j < m$. As explained above, this is independent of the temporary definition of valuation of \vec{p} , as the valuation in points of depth at most $i + 1$ has already been fixed. The formula $\Phi_{i+1}(\vec{a}_0, \dots, \vec{a}_i, \vec{b}_0, \dots, \vec{b}_i)$ is false by (68).

When the construction is finished, we have $r \models \theta_{d-1} \wedge \gamma$, and the valuation of $\Diamond(\theta_i \wedge p_j)$ and $\Diamond(\theta_i \wedge \sigma(x_j))$ in r agrees with $a_{i,j}$ and $b_{i,j}$, respectively, for each $i < d$ and $j < m$, hence $\varphi(\vec{a}_0, \dots, \vec{a}_{d-1}, \vec{b}_0, \dots, \vec{b}_{d-1})$ is true due to $\vdash_L \sigma(\xi_\Phi)$. However, this is a contradiction, as $\varphi = \Phi_d$. \square (Claim 5.25.2)

For fixed d , validity of quantified Boolean sentences of the form (63) is a Π_{2d}^P -complete problem, hence Claims 5.25.1 and 5.25.2 imply (i). When d is unrestricted, the validity problem for (63) is PSPACE-complete, and we may arrange $m = 1$ by inserting dummy quantifiers to make universal and existential quantifiers alternate. This shows (ii). \square

We may strengthen Theorem 5.25 by combining it with our other lower bounds, observing that PSPACE (and Π_{2d}^P) is included in both NEXP and coNEXP:

Corollary 5.26

- (i) *Unifiability in all logics except tabular logics of width 1 is uniformly PSPACE-hard.*
 - (ii) *For every $d > 0$, unifiability in all logics except tabular logics of width 1 and depth at most $d - 1$ is uniformly Π_{2d}^P -hard.*
- In particular, unifiability in all consistent logics is uniformly Π_2^P -hard.*

Proof: Let A be a language in PSPACE or Π_{2d}^p as appropriate. Let $w \mapsto \xi_w^0$, $w \mapsto \xi_w^1$, and $w \mapsto \xi_w^2$ be the reductions supplied by Theorems 5.3, 5.6, and 5.25 (respectively). We may assume that the formulas ξ_w^0 , ξ_w^1 , and ξ_w^2 use disjoint sets of variables. Then $\xi_w := \xi_w^0 \wedge \xi_w^1 \wedge \xi_w^2$ is L -unifiable if and only if each ξ_w^i is, thus:

- if $w \in A$, then ξ_w is **K4**-unifiable;
- if $w \notin A$, then ξ_w is not L -unifiable for any logic L that is nonlinear, or of unbounded cluster size, or of sufficiently large depth.

The three classes of logics in the second point together cover the class from the statement of this corollary. \square

Corollary 5.27 *Let L be a consistent linear clx logic of bounded cluster size. If L has branching at least 1 (hence unbounded depth), then L -unifiability and L -admissibility are PSPACE-complete; otherwise, L -unifiability is Π_2^p -complete, and L -admissibility Σ_2^p -complete.*

If L is a linear tabular logic of depth (exactly) $d \geq 1$, then L -unifiability is Π_{2d}^p -complete, and L -admissibility is Σ_{2d}^p -complete.

Proof: By Theorems 4.3, 4.21, and 5.25. \square

Example 5.28 L -unifiability and L -admissibility are PSPACE-complete for the logics $L = \mathbf{GL.3}$, $\mathbf{S4.3Grz}$, and $\mathbf{K4.3Grz}$. Unifiability in $L \oplus \mathbf{BD}_d$ is Π_{2d}^p -complete for $d \geq 1$.

Unifiability is also Π_2^p -complete in the logics $\mathbf{Verum} = \mathbf{K4} \oplus \square \perp$ or $\mathbf{S5} \oplus \mathbf{Alt}_k$ for $k \geq 1$.

In general, two parameters are necessary in Theorem 5.25 (ii):

Example 5.29 The logic $L = \mathbf{S4.2Grz}$ or the clx logic $L = \mathbf{S4Grz.3}$ have unbounded depth, but L -unifiability of formulas with one parameter is NP-complete, as shown in Example 5.17.

Unifiability with one parameter in $\mathbf{S4.3}$ is also NP-complete: the universal frame is only modified by attaching copies of the two-element cluster realizing both p and $\neg p$ below each original node, and including an additional copy as a third connected component. We leave the details to the reader.

However, we will show for completeness that one parameter is sufficient under a stronger hypothesis, namely that the logic admits arbitrarily long *irreflexive* chains (possibly embedded among reflexive points). Note that at least one parameter is needed anyway, as parameter-free unifiability is in NP even for irreflexive logics such as \mathbf{GL} .

Definition 5.30 A logic L has *irreflexive depth at least n* if there exists an L -frame that subreduces onto an n -element irreflexive chain. A logic has *unbounded irreflexive depth* if it has irreflexive depth at least n for all $n \in \omega$.

Theorem 5.31 *Unifiability of formulas with 1 parameter and 1 variable in logics of unbounded irreflexive depth is uniformly PSPACE-hard.*

Proof: As in the proof of Theorem 5.25 (i), we will provide a reduction from validity of quantified Boolean sentences of the form (63) with $m = 1$, but instead of p_0 and q , we will need to reuse a single parameter q for detecting the depth and for specifying a_i . (Since $m = 1$, we will omit the vector signs, and write just a_i, b_i instead of $a_{i,0}, b_{i,0}$.) Roughly, the idea is that each quantifier pair $\forall a_i \exists b_i$ corresponds to three irreflexive points in a chain: the top one will satisfy $\neg q$, the bottom one q , and the middle one will depend on a_i . In this way, we get two alternations in the value of q for each i irrespective of the value of a_i , while we can read off a_i by checking if the i th layer of points satisfying q contains a nontrivial chain $u < v$. (Here we are using the irreflexivity.) The value of b_i will be given using a variable x , one level lower still (so that a_i is reliably fixed).

Formally, let θ_i denote the formulas defined in the proof of Theorem 5.25, and let ξ_Φ be the conjunction of the formulas

$$(69) \quad \gamma \rightarrow \bigwedge_{i < d} (\Box(\theta_{2i+2} \rightarrow x) \vee \Box(\theta_{2i+2} \rightarrow \neg x)),$$

$$(70) \quad \gamma \wedge \theta_{2d} \rightarrow \varphi(\dots, a_i / \Diamond(\theta_{2i+1} \wedge \Diamond\theta_{2i+1}), \dots, b_i / \Diamond(\theta_{2i+2} \wedge x), \dots),$$

where

$$\gamma = \bigwedge_{i < d} (\Box(\theta_{2i+1} \rightarrow \Box\neg\theta_{2i+1}) \vee \Box(\theta_{2i+2} \rightarrow \Diamond(\theta_{2i+1} \wedge \Diamond\theta_{2i+1}))).$$

Again, $\Phi \mapsto \xi_\Phi$ is a logspace-computable mapping.

Claim 5.31.1 *If Φ is true, then ξ_Φ is $\mathbf{K4}$ -unifiable.*

Proof: Define the functions F_i and G_i as in Claim 5.25.1. Let $W = U_{\mathbf{K4}}(q)$; we will construct a valuation of the variable x in W as follows. Let $w \in W$. If $w \models \theta_{2i+2}$ for some $i < d$, let $a_0(w), \dots, a_i(w) \in \mathbf{2}$ be defined by

$$a_{i'}(w) = 1 \quad \text{iff} \quad w \models \Diamond(\theta_{2i'+1} \wedge \Diamond\theta_{2i'+1}).$$

Put $b_i(w) = G_i(a_0(w), \dots, a_i(w))$, and define $w \models x^{b_i(w)}$. For other w , define $w \models x$ in an arbitrary way.

A similar argument as in Claim 5.31.1 shows that under this valuation, all points of W satisfy (69) and (70). Thus, ξ_Φ is $\mathbf{K4}$ -unifiable by [8, Thm. 5.18]. \square (Claim 5.31.1)

Claim 5.31.2 *If L has irreflexive depth at least $3d + 1$, and ξ_Φ is L -unifiable, then Φ is true.*

Proof: Let us fix an L -frame W , and a subreduction f from W onto the irreflexive chain of length $3d + 1$, whose elements we label as $c_0 > c_1 > \dots > c_{3d}$. We may assume that W is rooted, and its root r maps to c_{3d} . For any $i \leq 3d$, put $D_i = f^{-1}[c_i]$. The sets D_i are pairwise disjoint admissible subsets of the frame W ; each D_i is a nonempty irreflexive antichain, and for each $i' < i$, $D_{i'}$ is a cover of D_i (i.e., every element of D_i sees some element of $D_{i'}$). We have $D_{3d} = \{r\}$.

For each $a = \langle a_0, \dots, a_{d-1} \rangle \in \mathbf{2}^d$, let W_a denote the parametric frame for parameter q based on W such that

$$\{w \in W_a : w \models q\} = \bigcup_{i < d} D_{3i+2} \cup \bigcup_{\substack{i < d \\ a_i = 1}} (D_{3i+1} \downarrow \setminus D_{3i+2} \downarrow).$$

It is straightforward to check that in W_a :

- For each $i < d$, points of D_{3i+2} satisfy θ_{2i+1} .
- Points of $W \setminus D_1 \downarrow$ satisfy θ_0 , points of $D_{3d-1} \downarrow$ satisfy θ_{2d} , and for each $0 < i < d$, points of $D_{3i-1} \downarrow \setminus D_{3i+1} \downarrow$ satisfy θ_{2i} . In particular, for each $i \leq d$, points of D_{3i} satisfy θ_{2i} .
- For each $i < d$, points of $D_{3i+1} \downarrow \setminus D_{3i+2} \downarrow$ (which includes D_{3i+1}) satisfy θ_{2i} or θ_{2i+1} according to the value of a_i .

As a consequence, $W_a \models \gamma$: specifically, for any $i < d$, $W_a \models \theta_{2i+1} \rightarrow \Box \neg \theta_{2i+1}$ if $a_i = 0$, and $W_a \models \theta_{2i+2} \rightarrow \Diamond (\theta_{2i+1} \wedge \Diamond \theta_{2i+1})$ if $a_i = 1$.

Now, assume that ξ_Φ is L -unifiable, and let us fix a unifier σ of ξ_Φ whose target formulas contain no variables, and no parameters save q . We define a function $G: \mathbf{2}^d \rightarrow \mathbf{2}^d$ as follows: given $a \in \mathbf{2}^d$, we have $W_a \models \sigma(\xi_\Phi)$ as σ is a unifier, and we know $W_a \models \gamma$, thus using (69), there is for each $i < d$ a unique $b_i \in \mathbf{2}$ such that $W_a \models \theta_{2i+2} \rightarrow \sigma(x)^{b_i}$. We put $G(a) = \langle b_0, \dots, b_{d-1} \rangle$.

By the breakdown above, points $w \in D_{3i+3}$ satisfy θ_{2i+2} in W_a , and the valuation of q in $w \uparrow$ is determined by a_0, \dots, a_i , hence so is the valuation of $\sigma(x)$. That is, b_i only depends on a_0, \dots, a_i : there are functions $G_i: \mathbf{2}^{i+1} \rightarrow \mathbf{2}$, $i < d$, such that

$$G(a_0, \dots, a_{d-1}) = \langle G_0(a_0), G_1(a_0, a_1), \dots, G_{d-1}(a_0, \dots, a_{d-1}) \rangle$$

for each $a \in \mathbf{2}^d$. Since $W_a \models \sigma(\xi_\Phi)$, we have

$$\varphi(a_0, \dots, a_{d-1}, G_0(a_0), \dots, G_{d-1}(a_0, \dots, a_{d-1})) = 1$$

using (70). This makes G_0, \dots, G_{d-1} Skolem functions witnessing the truth of Φ : formally, we may show

$$\forall a_0, \dots, a_{i-1} \in \mathbf{2} \Phi_i(a_0, \dots, a_{i-1}, G_0(a_0), \dots, G_{i-1}(a_0, \dots, a_{i-1})) = 1$$

by reverse induction on $i \leq d$, with Φ_i defined by (66). □ (Claim 5.31.2)

The result now follows from Claims 5.31.1 and 5.31.2. □

Corollary 5.32 *Unifiability of formulas with 1 parameter and 1 variable is PSPACE-hard in all consistent $\langle \bullet, 1 \rangle$ -extensible logics.* □

6 Conclusion

We have undertaken a thorough investigation of the complexity of admissibility and unifiability with parameters in transitive modal logics. In particular, we obtained a classification of these problems for cluster-extensible logics, a number of results for logics of bounded depth and width, and several very general lower bounds on hardness of unifiability in broad classes of transitive logics.

When we reflect on the effect of parameters on the complexity, one recurring observation is that the power of having an unlimited supply of parameters can be fully exploited only for logics that allow arbitrary large clusters, but in that case it blows up the complexity of unifiability to coNEXP (or from NEXP to Σ_2^{exp}). This is perhaps most striking for **S5**, which is at the same time one of the most fundamental and one of the most simple modal logics; the derivability problem for **S5** is coNP-complete, and this also holds for parameter-free admissibility, or even admissibility with a constant number of parameters. However, **S5**-unifiability with arbitrary parameters is coNEXP-complete, i.e., exponentially harder.

A perhaps more intriguing effect is that already the addition of 1–2 parameters may raise the complexity of unifiability from NP to PSPACE in logics of unbounded depth. Even better, we saw that for some logics such as **GL** or **D4**, unifiability with 1 parameter is NEXP-complete, but parameter-free unifiability is NP-complete, hence again we have an exponential blow-up. In these cases, the complexity of parameter-free unifiability is exceptionally low (e.g., parameter-free inadmissibility is already NEXP-complete).

This brings us to one of the rough spots where our analysis is incomplete, even for clx logics: what is the complexity of parameter-free unifiability in nonlinear clx logics? There are some cases where it is NEXP-complete (see Theorem 5.20), and some cases where it is NP-complete (e.g., extensions of **GL** or **D4**, as just mentioned), but we do not know where exactly is the dividing line, or even if there are cases of intermediate complexity.

In the regime with $O(1)$ parameters, our results are enough to settle the complexity for clx logics and some other logics of interest, nevertheless the results are far from satisfactory: the conditions used for the lower bounds have an ad hoc feeling to them, and are unlikely to be exhaustive. We suspect that a more comprehensive classification might be quite hard here.

Finally, in the regime with arbitrarily many parameters, our lower and upper bounds do not match for some interesting classes of logics that are nonlinear and of unbounded cluster size, but do not satisfy the hypothesis of Theorem 5.10. As we already mentioned in the introduction, this gap will be addressed in the sequel.

Acknowledgements

I want to thank the anonymous referee whose suggestions improved the presentation of the paper.

The research was supported by grant 19-05497S of GA ĀR. The Institute of Mathematics of the Czech Academy of Sciences is supported by RVO: 67985840.

References

- [1] Sanjeev Arora and Boaz Barak, *Computational complexity: A modern approach*, Cambridge University Press, 2009.
- [2] Alexander V. Chagrov, *On the complexity of propositional logics*, in: Complexity problems in Mathematical Logic, Kalinin State University, 1985, pp. 80–90 (in Russian).
- [3] Alexander V. Chagrov and Michael Zakharyashev, *Modal logic*, Oxford Logic Guides vol. 35, Oxford University Press, 1997.
- [4] Lauri Hella and José María Turull-Torres, *Computing queries with higher-order logics*, Theoretical Computer Science 355 (2006), no. 2, pp. 197–214.
- [5] Emil Jeřábek, *Complexity of admissible rules*, Archive for Mathematical Logic 46 (2007), no. 2, pp. 73–92.
- [6] —————, *Canonical rules*, Journal of Symbolic Logic 74 (2009), no. 4, pp. 1171–1205.
- [7] —————, *Substitution Frege and extended Frege proof systems in non-classical logics*, Annals of Pure and Applied Logic 159 (2009), no. 1–2, pp. 1–48.
- [8] —————, *Rules with parameters in modal logic I*, Annals of Pure and Applied Logic 166 (2015), no. 9, pp. 881–933.
- [9] Leszek A. Kołodziejczyk, *Truth definitions in finite models*, Journal of Symbolic Logic 69 (2004), no. 1, pp. 183–200.
- [10] Richard E. Ladner, *The computational complexity of provability in systems of modal propositional logic*, SIAM Journal on Computing 6 (1977), no. 3, pp. 467–480.
- [11] Richard Statman, *Intuitionistic propositional logic is polynomial-space complete*, Theoretical Computer Science 9 (1979), no. 1, pp. 67–72.