

The Complexity of Decomposing Modal and First-Order Theories

STEFAN GÖLLER, University of Bremen
 JEAN-CHRISTOPH JUNG, University of Bremen
 MARKUS LOHREY, University of Leipzig

We study the satisfiability problem of the logic $K^2 = K \times K$, i.e., the two-dimensional variant of unimodal logic, where models are restricted to asynchronous products of two Kripke frames. Gabbay and Shehtman proved in 1998 that this problem is decidable in a tower of exponentials. So far the best known lower bound is NEXP-hardness shown by Marx and Mikuláš in 2001.

Our first main result closes this complexity gap: We show that satisfiability in K^2 is nonelementary. More precisely, we prove that it is k -NEXP-complete, where k is the switching depth (the minimal modal rank among the two dimensions) of the input formula, hereby solving a conjecture of Marx and Mikuláš. Using our lower-bound technique allows us to derive also nonelementary lower bounds for the two-dimensional modal logics $K4 \times K$ and $S5_2 \times K$ for which only elementary lower bounds were previously known.

Moreover, we apply our technique to prove nonelementary lower bounds for the sizes of Feferman-Vaught decompositions with respect to product for any decomposable logic that is at least as expressive as unimodal K , generalizing a recent result by the first author and Lin. For the three-variable fragment FO^3 of first-order logic, we obtain the following immediate corollaries: (i) the size of Feferman-Vaught decompositions with respect to disjoint sum are inherently nonelementary and (ii) equivalent formulas in Gaifman normal form are inherently nonelementary.

Our second main result consists in providing effective elementary (more precisely, doubly exponential) upper bounds for the two-variable fragment FO^2 of first-order logic both for Feferman-Vaught decompositions and for equivalent formulas in Gaifman normal form.

Categories and Subject Descriptors: F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic

General Terms: Modal logic, Model theory

Additional Key Words and Phrases: Modal logic, two-dimensional modal logic, Feferman-Vaught decomposition, Gaifman's theorem

ACM Reference Format:

Stefan Göller, Jean-Christoph Jung, and Markus Lohrey, 2014. The complexity of decomposing modal and first-order theories. *ACM Trans. Comput. Logic* 9, 4, Article 39 (March 2010), 43 pages.

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

1.1. Modal logic and many-dimensional modal logic

Modal logic [1; 2] originated in philosophy and for a long time it was known as “the logic of necessity and possibility”. Later, it has been discovered that modal logics are well-suited to talk about relational structures, so called (*Kripke*) frames. Relational

Markus Lohrey was partially supported by the DFG research project GELO (LO 748/7-2). Jean Christoph Jung was supported by the DFG research project ProBDL (LU1417/1-1).

Author's addresses: Stefan Göller, ENS Cachan, France; Jean-Christoph Jung, University of Bremen, Germany; Markus Lohrey, University of Siegen, Germany

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2010 ACM 1529-3785/2010/03-ART39 \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

structures appear in many branches of computer science; consider for example transition systems in verification, semantic networks in knowledge representation, or attribute value structures in linguistics. This has led to various applications of modal logic in areas such as computer science, mathematics, and artificial intelligence.

Depending on the application, a lot of different modal operators have been introduced in the past, each of them tailored towards expressing different features of the domain. For instance there are modalities that talk about time, space, knowledge, beliefs, etc.

However, it turned out that recent application domains require to express properties that *combine* different modalities, e.g., talk about the evolution of knowledge over time. In order to reflect these requirements in theory, *many-dimensional* modal logics have been studied intensively [3; 4]. A particular way of combining two logics \mathcal{L}_1 and \mathcal{L}_2 is building their *product* $\mathcal{L}_1 \times \mathcal{L}_2$ [5]. For products, the semantics is given in terms of structures, whose frames are restricted to be asynchronous products of the (one-dimensional) component frames. The interpretation of the atomic propositions is done in a *uninterpreted* way, i.e., it is independent from the component frames.

An important and well-studied problem in this context is *satisfiability checking*, i.e., to decide whether a given formula admits a model. When considering products of modal logics, it has been shown that the computational complexity of satisfiability checking often increases drastically in comparison to the well-behaved component logics. As an example, consider the basic modal logic \mathbf{K} and its variant $\mathbf{K4}$ for reasoning over the class of transitive frames. Satisfiability is PSPACE-complete for both \mathbf{K} and $\mathbf{K4}$ [6], while for $\mathbf{K} \times \mathbf{K}$ and $\mathbf{K4} \times \mathbf{K}$ only nonelementary upper bounds were known [5]. Even worse, satisfiability becomes undecidable in $\mathbf{K} \times \mathbf{K} \times \mathbf{K}$ [7] and $\mathbf{K4} \times \mathbf{K4}$ [8]. To some extent, this can be explained by the grid-like shape of product structures.

1.2. Logical decomposition

Logical decomposition can concisely be summarized as follows: A logic \mathcal{L} admits decomposition w.r.t. some operation op on structures if all \mathcal{L} -properties that are interpreted on composed (with respect to the operation op) structures, are already determined by the \mathcal{L} -properties of the component structures. Logical decomposition dates back to the work of Mostowski [9] and Feferman and Vaught [10], where it is shown that first-order logic (FO) is decomposable w.r.t. a general product operation, which covers also disjoint union and product. Later, both for more expressive logics and for more sophisticated operations such decomposability results have been proven, see [11] for an excellent survey.

When proving decomposability for a logic \mathcal{L} , one often obtains an effective procedure for computing such decompositions: Given a formula φ from \mathcal{L} evaluated on composed structures, one can effectively compute (i) a finite set of formulas $\{\varphi_1, \dots, \varphi_n\}$, each being evaluated on some specific component, and (ii) a propositional formula β , whose propositions are tests of the form $\mathfrak{S}_i \models \varphi_j$, such that for all composed structures $\mathfrak{S} = \text{op}(\mathfrak{S}_1, \dots, \mathfrak{S}_k)$: $\mathfrak{S} \models \varphi$ if and only if β evaluates to true. The size of the resulting decomposition is typically nonelementary in the size of the original formula. Dawar et al. proved that this is unavoidable if $\mathcal{L} = \text{FO}$ [12].

Decomposition theorems have powerful implications in computer science logic. Let us mention only four of them.

Firstly, assume some decomposable logic \mathcal{L} : Then *decidability of the \mathcal{L} -theory* of some composed structure, for instance a product structure, can be derived from the decidability of the \mathcal{L} -theories of its component structures.

Secondly, let us mention that *model checking* a fixed \mathcal{L} -formula (i.e. the *data complexity*) in a composed structure is not harder than model checking fixed \mathcal{L} -formulas

on the component structures: If the formula is fixed, also the decomposition is fixed (although possibly large).

Moreover, decompositional methods can be applied for showing decidability of *satisfiability checking*: Instead of asking whether a given formula φ is satisfiable in a composed model, one computes a decomposition for φ , translates the decomposition into disjunctive normal form, and finally checks satisfiability of a conjunction of formulas in their corresponding components. Rabinovich proved that basic modal logic \mathbf{K} is decomposable w.r.t. *interpreted* products [13], i.e. where “interpreted” means that an interpretation of the propositions on the respective component structures is applied. It is worth noting that this, however, does not lead to decidability of $\mathbf{K} \times \mathbf{K}$ w.r.t. the classical (uninterpreted) products mentioned above. To the contrary, satisfiability w.r.t. interpreted products is easily reducible to the uninterpreted version.

Finally, an important application of logical decomposition à la Feferman and Vaught is the (original) proof of *Gaifman’s locality theorem* [14] stating that every first-order sentence is equivalent to a boolean combination of basic local sentences, where a basic local sentence admits quantification only relativized to finite neighbourhoods of elements. Gaifman’s locality theorem has important applications such as inexpressibility results for first-order logic. For a further and more recent application of Gaifman’s locality theorem we mention algorithmic meta-theorems for first-order logic [15], stating that first-order properties can be efficiently evaluated on numerous classes of structures.

1.3. Our contributions and related work

As our first main result we show that (even the interpreted variant of) the satisfiability problem of two-dimensional modal logic $\mathbf{K}^2 = \mathbf{K} \times \mathbf{K}$ has nonelementary complexity, hereby solving a fundamental problem that has been open for more than 10 years. Gabbay and Shehtman proved in 1998 that satisfiability in \mathbf{K}^2 is decidable in a tower of exponentials [5]. To the best of the authors’ knowledge, the best known lower bound has been NEXP-hardness shown by Marx and Mikulas in 2001 [16]. In fact, we prove that satisfiability in \mathbf{K}^2 restricted to formulas of switching depth k (the minimal modal rank among the two dimensions) is k -NEXP-complete (where k -NEXP is the set of all problems that can be solved on a nondeterministic Turing machine in k -fold exponential time), hereby confirming a conjecture of Marx and Mikulas [16]. We derive nonelementary lower bounds for the two-dimensional modal logics $\mathbf{K}4 \times \mathbf{K}$ and $\mathbf{S}5_2 \times \mathbf{K}$ for which only elementary lower bounds were known [3].

Our lower bound technique allows us to derive a nonelementary lower bound for the size of Feferman-Vaught decompositions w.r.t. product for \mathbf{K} . Such a result was already shown in [17]. However, in contrast to [17], our proof technique implies that the nonelementary lower bound carries over to all decomposable logics that are at least as expressive as \mathbf{K} . An instance of such a logic is the two-variable fragment FO^2 of first-order logic. Moreover, we prove that the same lower bound holds when relativized to the class of finite trees, answering an open problem formulated in [17].

In the same fashion, we derive the following new results for the three-variable fragment FO^3 of first-order logic: (i) the sizes of Feferman-Vaught decompositions w.r.t. disjoint sum are inherently nonelementary and (ii) equivalent formulas in Gaifman normal form are inherently nonelementary. It is worth mentioning that (i) and (ii) were shown in [12] for full FO. By inspecting the formulas in [12] it turns out that they are in fact FO^4 -formulas. However, it seems to be unclear whether the construction from [12] can be adapted so that it yields FO^3 -formulas.

Finally, we provide effective doubly exponential (and hence elementary) upper bounds for the two-variable fragment FO^2 of first-order logic both for Feferman-Vaught

decompositions and for equivalent formulas in Gaifman normal form. This supports former observations that in many aspects FO^2 is better behaved than FO^3 . For instance, in contrast to FO^3 it has a finite model property and satisfiability is decidable in NEXP [18]. We also prove (non-matching) lower bounds of the form $c^{f(n)}$ (for any constant c and any function $f(n) \in o(\sqrt{n})$) for both Feferman-Vaught decompositions and equivalent formulas in Gaifman normal form for FO^2 .

An extended abstract of this paper appeared as [19].

2. PRELIMINARIES

For $i, j \in \mathbb{Z}$ let $[i, j]$ be the interval $[i, i + 1, \dots, j]$. By $\mathbb{N} \stackrel{\text{def}}{=} \{0, 1, \dots\}$ we denote the non-negative integers. For a set X we denote by $\text{bool}(X)$ the set of boolean formulas with variables ranging over X . Let $u = u_1 \dots u_k \in \Sigma^*$ with $u_i \in \Sigma$ for each $i \in [1, k]$. By $|u| \stackrel{\text{def}}{=} k$ we denote the *length* of u . The *tower function* $\text{Tower} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ is defined as $\text{Tower}(0, n) \stackrel{\text{def}}{=} n$ and $\text{Tower}(\ell + 1, n) \stackrel{\text{def}}{=} 2^{\text{Tower}(\ell, n)}$ for each $\ell, n \in \mathbb{N}$. A function $f : \mathbb{N} \rightarrow \mathbb{N}$ is *elementary* if it can be formed from the successor function, addition, subtraction, and multiplication using compositions, projections, bounded additions and bounded multiplications (of the form $\sum_{z \leq y} g(\bar{x}, z)$ and $\prod_{z \leq y} g(\bar{x}, z)$). We will sometimes use the fact that for each elementary function $f : \mathbb{N} \rightarrow \mathbb{N}$ there exist some $h_0 \in \mathbb{N}$ such that $f(h) < \text{Tower}(h, 2)$ for each $h \geq h_0$.

2.1. Kripke frames and structures

Let us fix a countable set of *action labels* \mathbb{A} and a countable set of *propositional variables* \mathbb{P} . For a finite set $A \subseteq \mathbb{A}$ of action labels, an *A-frame* is a tuple $\mathfrak{F} = (W, \{\overset{a}{\rightarrow} \mid a \in A\})$, where W is set of *worlds* and $\overset{a}{\rightarrow} \subseteq W \times W$ is a binary (accessibility) relation over W for each $a \in A$. Most of the time we write $v \overset{a}{\rightarrow} w$ instead of $(v, w) \in \overset{a}{\rightarrow}$. We say that \mathfrak{F} is a *tree* if

- $W \subseteq U^*$ is a prefix-closed set of words for some set U ,
- $\overset{a}{\rightarrow} \cap \overset{b}{\rightarrow} = \emptyset$ for each $a, b \in A$ with $a \neq b$, and
- for all $v, w \in W$, we have $v \overset{a}{\rightarrow} w$ for some $a \in A$ if and only if there exists $u \in U$ with $w = vu$.

We say that \mathfrak{F} is *finite* if W is finite. In case \mathfrak{F} is a finite tree, the *height* of \mathfrak{F} is defined as $\max\{|w| \mid w \in W\}$.

An (A, \mathbb{P}) -*Kripke structure* (or (A, \mathbb{P}) -*structure* for short), for a finite set $A \subseteq \mathbb{A}$ of action labels and a finite set $\mathbb{P} \subseteq \mathbb{P}$ of propositional variables, is a tuple

$$\mathfrak{G} = (W, \{\overset{a}{\rightarrow} \mid a \in A\}, \{W_p \mid p \in \mathbb{P}\}),$$

where $(W, \{\overset{a}{\rightarrow} \mid a \in A\})$ is an A -frame and $W_p \subseteq W$ is an interpretation for each propositional variable $p \in \mathbb{P}$. By $\mathfrak{F}(\mathfrak{G}) \stackrel{\text{def}}{=} (W, \{\overset{a}{\rightarrow} \mid a \in A\})$ we denote the underlying A -frame of \mathfrak{G} . We say that \mathfrak{G} is a *tree structure* (or simply a *tree*) if $\mathfrak{F}(\mathfrak{G})$ is isomorphic to a tree. By $|\mathfrak{G}| = |W|$ we denote the *size* of \mathfrak{G} . For $s \in W$ let

$$N_{\mathfrak{G}}(s) \stackrel{\text{def}}{=} \{u \in W \mid \exists a \in A : s \overset{a}{\rightarrow} u\}$$

be the set of *successors* of s in \mathfrak{G} . A *pointed (A, \mathbb{P}) -structure* is a pair (\mathfrak{G}, s) where \mathfrak{G} is an (A, \mathbb{P}) -structure and s is a world of \mathfrak{G} . We identify a tree structure \mathfrak{G} with the pointed (A, \mathbb{P}) -structure (\mathfrak{G}, r) , where r is the root of \mathfrak{G} . An $(\{a\}, \mathbb{P})$ -structure is also called *unimodal* and we write $(W, \overset{a}{\rightarrow}, \{W_p \mid p \in \mathbb{P}\})$ instead of $(W, \{\overset{a}{\rightarrow}\}, \{W_p \mid p \in \mathbb{P}\})$.

2.2. Multimodal logic

Formulas of multimodal logic are defined by the following grammar, where a (resp., p) ranges over \mathbb{A} (resp., \mathbb{P}):

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \diamond_a \varphi$$

We introduce the usual abbreviations $\top = p \vee \neg p$ for some $p \in \mathbb{P}$, $\perp = \neg\top$, $\varphi_1 \vee \varphi_2 = \neg(\neg\varphi_1 \wedge \neg\varphi_2)$, and $\Box_a \varphi = \neg \diamond_a \neg\varphi$. We say that φ is *over* (A, P) if the set of action labels (resp., the set of propositional variables) that appear in φ is a subset of A (resp., P). For an (A, P) -structure $\mathfrak{G} = (W, \{\xrightarrow{a} \mid a \in A\}, \{W_p \mid p \in P\})$, $w \in W$, and a multimodal logic formula φ over (A, P) , we define the satisfaction relation $(\mathfrak{G}, w) \models \varphi$ by structural induction on φ , where $a \in A$ and $p \in P$:

$$\begin{aligned} (\mathfrak{G}, w) \models p &\stackrel{\text{def}}{\iff} w \in W_p \\ (\mathfrak{G}, w) \models \neg\varphi &\stackrel{\text{def}}{\iff} (\mathfrak{G}, w) \not\models \varphi \\ (\mathfrak{G}, w) \models \varphi_1 \wedge \varphi_2 &\stackrel{\text{def}}{\iff} (\mathfrak{G}, w) \models \varphi_1 \text{ and } (\mathfrak{G}, w) \models \varphi_2 \\ (\mathfrak{G}, w) \models \diamond_a \varphi &\stackrel{\text{def}}{\iff} \exists w' : w \xrightarrow{a} w' \text{ and } (\mathfrak{G}, w') \models \varphi \end{aligned}$$

Let φ be a multimodal logic formula over (A, P) . An (A, P) -structure \mathfrak{G} is a *model* of φ if $(\mathfrak{G}, w) \models \varphi$ for some world w of \mathfrak{G} . We say that φ is *satisfiable* if φ has a model.

If φ is a boolean formula with propositional variables from $P \subseteq \mathbb{P}$ and $\alpha : P \rightarrow \{0, 1\}$ then we write $\alpha \models \varphi$ if φ evaluates to 1 if every propositional variable $p \in P$ is replaced by the truth value $\alpha(p)$.

2.3. Asynchronous products

Fix non-empty, finite, and pairwise disjoint sets $A_1, \dots, A_d \subseteq \mathbb{A}$ of action labels and non-empty, finite, and pairwise disjoint sets $P_1, \dots, P_d \subseteq \mathbb{P}$ of propositional variables. Let $A = \bigcup_{i \in [1, d]} A_i$ and $P = \bigcup_{i \in [1, d]} P_i$. For A_i -frames $\mathfrak{F}_i = (W_i, \{\xrightarrow{a}_i \mid a \in A_i\})$ ($i \in [1, d]$) we define the *asynchronous product* $\prod_{i \in [1, d]} \mathfrak{F}_i = (W, \{\xrightarrow{a} \mid a \in A\})$ to be the A -frame, where

- $W = \prod_{i \in [1, d]} W_i$, and
- for each $\bar{v} = \langle v_1, \dots, v_d \rangle \in W$ and $\bar{w} = \langle w_1, \dots, w_d \rangle \in W$ we have $\bar{v} \xrightarrow{a} \bar{w}$ if and only if there is some $i \in [1, d]$ such that $a \in A_i$, $v_i \xrightarrow{a}_i w_i$ and $v_j = w_j$ for each $j \in [1, d] \setminus \{i\}$.

An (A, P) -structure $\mathfrak{G} = (W, \{\xrightarrow{a} \mid a \in A\}, \{W_p \mid p \in P\})$ is an *uninterpreted product structure* if $\mathfrak{F}(\mathfrak{G}) = \prod_{i \in [1, d]} \mathfrak{F}_i$, where each \mathfrak{F}_i is some A_i -frame. Thus, we do not make any restrictions on how atomic propositions are interpreted.

Next, let us define interpretations of atomic propositions in products, as introduced in [13]. A (*product*) *interpretation* is a mapping $\sigma : P \rightarrow \text{bool}(P)$, that is, a mapping from the set of propositions to boolean formulas over that set. In our lower bound proofs in Section 3, σ will be the *identity interpretation* id with $\text{id}(p) = p$ for all $p \in P$. Let

$$\mathfrak{G}_i = (W_i, \{\xrightarrow{a}_i \mid a \in A_i\}, \{W_{p,i} \mid p \in P_i\})$$

be an (A_i, P_i) -structure for $i \in [1, d]$. For an interpretation σ , their σ -*product* $\prod_{i \in [1, d]}^\sigma \mathfrak{G}_i$ is defined as the (A, P) -structure $\mathfrak{G} = (W, \{\xrightarrow{a} \mid a \in A\}, \{W_p \mid p \in P\})$ such that

- $\mathfrak{F}(\mathfrak{G}) = \prod_{i \in [1, d]} \mathfrak{F}(\mathfrak{G}_i)$ and
- $\langle w_1, \dots, w_d \rangle \in W_p$ if and only if $\alpha \models \sigma(p)$, where $\alpha(q) = 1$ if and only if $w_i \in W_{q,i}$ for each $i \in [1, d]$ and $q \in P_i$.

If no interpretation is given, we define $\prod_{i \in [1, d]} \mathfrak{S}_i \stackrel{\text{def}}{=} \prod_{i \in [1, d]}^{\text{id}} \mathfrak{S}_i$. Since we mainly deal with the case $d = 2$ and $\sigma = \text{id}$, let us repeat the above definition of the σ -product in this case: $\prod_{i \in [1, 2]}^{\text{id}} \mathfrak{S}_i$ is defined as the (A, P) -structure $\mathfrak{S} = (W, \{\overset{a}{\rightarrow} \mid a \in A\}, \{W_p \mid p \in P\})$ such that

- $\mathfrak{F}(\mathfrak{S}) = \mathfrak{F}(\mathfrak{S}_1) \times \mathfrak{F}(\mathfrak{S}_2)$ and
- $\langle w_1, w_2 \rangle \in W_p$ if and only if either $p \in P_1$ and $w_1 \in W_{p,1}$ or $p \in P_2$ and $w_2 \in W_{p,2}$.

In case (\mathfrak{S}_i, s_i) is a pointed (A_i, P_i) -structure for each $i \in [1, d]$, then by $\prod_{i \in [1, d]}^{\sigma} (\mathfrak{S}_i, s_i)$ we denote the pointed structure $(\prod_{i \in [1, d]}^{\sigma} \mathfrak{S}_i, \langle s_1, \dots, s_d \rangle)$.

2.4. Many-dimensional modal logic

Given (A, P) with $A = \uplus_{i \in [1, d]} A_i$ and $P = \uplus_{i \in [1, d]} P_i$, a $\mathbf{K}^d(A, P)$ -formula is a multimodal formula φ over the signature (A, P) . For a $\mathbf{K}^d(A, P)$ -formula φ and $i \in [1, d]$, we define $\text{rank}_i(\varphi)$ inductively

$$\begin{aligned} \text{rank}_i(p) &\stackrel{\text{def}}{=} 0 && \text{for } p \in P \\ \text{rank}_i(\neg\varphi) &\stackrel{\text{def}}{=} \text{rank}_i(\varphi) \\ \text{rank}_i(\varphi_1 \wedge \varphi_2) &\stackrel{\text{def}}{=} \max\{\text{rank}_i(\varphi_1), \text{rank}_i(\varphi_2)\} \\ \text{rank}_i(\diamond_a \varphi) &\stackrel{\text{def}}{=} \text{rank}_i(\varphi) && \text{if } a \notin A_i \\ \text{rank}_i(\diamond_a \varphi) &\stackrel{\text{def}}{=} \text{rank}_i(\varphi) + 1 && \text{if } a \in A_i \end{aligned}$$

and the *switching depth* of φ as $\min\{\text{rank}_i(\varphi) \mid 1 \leq i \leq d\}$ [16].

An *uninterpreted product model* of φ is an uninterpreted product structure \mathfrak{S} (in the above sense) such that for some world \bar{w} of \mathfrak{S} we have $(\mathfrak{S}, \bar{w}) \models \varphi$. For an interpretation σ , a σ -*model* is a σ -product structure \mathfrak{S} such that $(\mathfrak{S}, \bar{w}) \models \varphi$ for some world \bar{w} of \mathfrak{S} . We say φ is *uninterpreted satisfiable* (resp., σ -*satisfiable*) if φ has an uninterpreted (resp., σ -) product model. Let us introduce the following decision problems.

\mathbf{K}^d -SAT

INPUT: A $\mathbf{K}^d(A, P)$ -formula φ .
QUESTION: Is φ uninterpreted satisfiable?

We introduce the corresponding variant in the presence of an interpretation σ of the atomic propositions.

\mathbf{K}_σ^d -SAT

INPUT: A $\mathbf{K}^d(A, P)$ -formula φ .
QUESTION: Is φ σ -satisfiable?

The first problem (and variants thereof) were intensively studied in the past, for an overview consult [5]. In particular, it is well-known that \mathbf{K}^2 -SAT is decidable in nonelementary time.

The second problem, \mathbf{K}_σ^d -SAT is less standard but is more useful for our purposes. Indeed, it is not hard to show that $\mathbf{K}^d(A, P)$ -SAT is at least as difficult as $\mathbf{K}_\sigma^d(A, P)$ -SAT as stated in the following proposition which be technically useful in Sections 3 and 4. It can be shown for an arbitrary interpretation σ , but we will only need the case $\sigma = \text{id}$.

PROPOSITION 2.1. *There is a polynomial time many-one reduction from \mathbf{K}_{id}^d -SAT to \mathbf{K}^d -SAT which preserves the switching depth. The reduction still holds when restricted to any frame class.*

PROOF. Let $A = \bigcup_{i \in [1, d]} A_i$ be the set of action labels and let $P = \bigcup_{i \in [1, d]} P_i$ be the set of atomic propositions, and let φ be some formula over (A, P) .

The idea is to give a formula χ that admits only models which are id-product structures, in particular, φ is id-satisfiable if and only if $\varphi \wedge \chi$ is satisfiable.

We need the following definition: The set of *modal sequences* $\text{ms}(\psi) \subseteq A^*$ of a formula ψ is inductively defined as follows:

- $\text{ms}(p) \stackrel{\text{def}}{=} \{\varepsilon\}$ for each $p \in P$,
- $\text{ms}(\neg\psi) \stackrel{\text{def}}{=} \text{ms}(\psi)$,
- $\text{ms}(\psi_1 \wedge \psi_2) \stackrel{\text{def}}{=} \text{ms}(\psi_1) \cup \text{ms}(\psi_2)$,
- $\text{ms}(\diamond_a \psi) \stackrel{\text{def}}{=} (\{a\}\text{ms}(\psi)) \cup \{\varepsilon\}$.

We note that $\text{ms}(\psi)$ is prefix closed, that $|\text{ms}(\varphi)| \leq |\varphi|$, and that the maximal length of an element of $\text{ms}(\varphi)$ is the modal rank of φ . If $w = a_1 \dots a_n \in A^*$ we denote with \square_w the sequence of boxes $\square_{a_1} \dots \square_{a_n}$. Particularly, \square_ε is the empty sequence of boxes. Moreover, define the relation $w \xrightarrow{\text{def}} a_1 \circ \xrightarrow{\text{def}} a_2 \circ \dots \circ \xrightarrow{\text{def}} a_n$.

For $i \in [1, d]$ and a word $s \in A^*$ let $s \setminus i \in (A \setminus A_i)^*$ be the word that results from s by removing all occurrences of all symbols from A_i and let $s \upharpoonright i \in A_i^*$ be the word that results from s by removing all occurrences of all symbols from $A \setminus A_i$. We define the following formula χ :

$$\chi \stackrel{\text{def}}{=} \bigwedge_{i \in [1, d]} \bigwedge_{s \in \text{ms}(\varphi)} \bigwedge_{p \in P_i} \square_{s \upharpoonright i} ((p \rightarrow \square_{s \setminus i} p) \wedge (\diamond_{s \setminus i} p \rightarrow p))$$

We define $\varphi' \stackrel{\text{def}}{=} \varphi \wedge \chi$.

Note that φ' has the same switching depth as φ and can be constructed in polynomial time from φ . Therefore it suffices to show that φ is id-satisfiable if and only if φ' is uninterpreted satisfiable.

Since every id-product satisfies χ , it follows that φ' is uninterpreted satisfiable if φ is id-satisfiable. For the other direction let $\mathfrak{S} = (W, \{\xrightarrow{a} \mid a \in A\}, \{W_p \mid p \in P\})$ be an (A, P) -structure such that $\mathfrak{F}(\mathfrak{S}) = \prod_{i \in [1, d]} \mathfrak{F}_i$ where $\mathfrak{F}_i = (W_i, \{\xrightarrow{a} \mid a \in A_i\})$ is an A_i -frame for each $i \in [1, d]$ and assume that $(\mathfrak{S}, \bar{w}_0) \models \varphi \wedge \chi$ for some $\bar{w}_0 \in W$. In particular, we have $W = \prod_{i \in [1, d]} W_i$. Let

$$W_R = \{\bar{w} \in W \mid \exists s \in \text{ms}(\varphi) : \bar{w}_0 \xrightarrow{s} \bar{w}\}.$$

We define for each $i \in [1, d]$ an (A_i, P_i) -structure $\mathfrak{S}_i = (W_i, \{\xrightarrow{a} \mid a \in A_i\}, \{V_p \mid p \in P_i\})$ with $\mathfrak{F}(\mathfrak{S}_i) = \mathfrak{F}_i$ such that $(\prod_{i \in [1, d]} \mathfrak{S}_i, \bar{w}_0) \models \varphi$. For giving the interpretations V_p we need the following statement:

Claim 1. For all $\bar{v}, \bar{w} \in W_R$, $i \in [1, d]$, and $p \in P_i$, if $\bar{v}(i) = \bar{w}(i)$ then $(\bar{v} \in W_p \Leftrightarrow \bar{w} \in W_p)$.

PROOF OF CLAIM 1. Since $\bar{v}, \bar{w} \in W_R$ there exist $s, t \in \text{ms}(\varphi)$ such that

$$\bar{w}_0 \xrightarrow{s} \bar{v} \quad \text{and} \quad \bar{w}_0 \xrightarrow{t} \bar{w}.$$

Since we have a product model and $\bar{v}(i) = \bar{w}(i)$, there exists some \bar{u} with

$$\bar{w}_0 \xrightarrow{t \upharpoonright i} \bar{u}, \quad \bar{w}_0 \xrightarrow{s \upharpoonright i} \bar{u}, \quad \bar{u} \xrightarrow{s \setminus i} \bar{v}, \quad \bar{u} \xrightarrow{t \setminus i} \bar{w}.$$

Since $(\mathfrak{S}, \bar{w}_0) \models \chi$, we get $\bar{v} \in W_p \Rightarrow \bar{u} \in W_p \Rightarrow \bar{w} \in W_p$ and analogously $\bar{w} \in W_p \Rightarrow \bar{v} \in W_p$, which proves Claim 1. \square

Let us define for all $p \in P_i$

$$V_p = \{\bar{w}(i) \in W_i \mid \bar{w} \in W_R \cap W_p\}.$$

With Claim 1, we get for all $\bar{w} \in W_R$:

$$\bar{w} \in W_p \Leftrightarrow \bar{w}(i) \in V_p \quad (1)$$

It remains to show that $(\prod_{i \in [1,d]}^{\text{id}} \mathfrak{S}_i, \bar{w}_0) \models \varphi$. For a sequence $s \in \text{ms}(\varphi)$ let $\text{sub}_s(\varphi)$ be the set of all subformulas ψ of φ such that in the syntax tree for φ there exists a path to an occurrence of ψ such that s is the sequence of modalities along this path. We prove by induction on the structure of a subformula $\psi \in \text{sub}_s(\varphi)$ that for all $\bar{w} \in W_R$ with $\bar{w}_0 \xrightarrow{s} \bar{w}$:

$$(\mathfrak{S}, \bar{w}) \models \psi \Leftrightarrow \left(\prod_{i \in [1,d]}^{\text{id}} \mathfrak{S}_i, \bar{w} \right) \models \psi.$$

For the induction base consider a propositional variable $p \in P$ and assume that $p \in P_i$ and let $\bar{w} \in W_R$. We get:

$$(\mathfrak{S}, \bar{w}) \models p \Leftrightarrow \bar{w} \in W_p \stackrel{(1)}{\Leftrightarrow} \bar{w}(i) \in V_p \Leftrightarrow \left(\prod_{i \in [1,d]}^{\text{id}} \mathfrak{S}_i, \bar{w} \right) \models p$$

For the induction step, the operators \wedge and \neg are straightforward. Finally, let $\psi = \diamond_a \theta \in \text{sub}_s(\varphi)$ and assume that $\bar{w}_0 \xrightarrow{s} \bar{w}$. Hence, $\theta \in \text{sub}_{sa}(\varphi)$. We have:

$$\begin{aligned} (\mathfrak{S}, \bar{w}) \models \diamond_a \theta &\Leftrightarrow \exists \bar{w}' : \bar{w} \xrightarrow{a} \bar{w}' \wedge (\mathfrak{S}, \bar{w}') \models \theta \\ &\stackrel{\text{hyp}}{\Leftrightarrow} \exists \bar{w}' : \bar{w} \xrightarrow{a} \bar{w}' \wedge \left(\prod_{i \in [1,d]}^{\text{id}} \mathfrak{S}_i, \bar{w}' \right) \models \theta \\ &\Leftrightarrow \left(\prod_{i \in [1,d]}^{\text{id}} \mathfrak{S}_i, \bar{w} \right) \models \diamond_a \theta \end{aligned}$$

Since $\varphi \in \text{sub}_\varepsilon(\varphi)$, and $\bar{w}_0 \xrightarrow{\varepsilon} \bar{w}_0$, this shows that $(\prod_{i \in [1,d]}^{\text{id}} \mathfrak{S}_i, \bar{w}_0) \models \varphi$.

Overall, we have given a reduction from \mathbf{K}_{id}^d -SAT to \mathbf{K}^d -SAT. Note that the constructed formula φ' has the same switching depth as φ . \square

Note that the proof of Proposition 2.1 does not make any assumption about the underlying structures. In particular, this implies that Proposition 2.1 holds also for restricted frame classes, e.g., transitive frames.

2.5. Bisimulation equivalence

Let $\mathfrak{S} = (W, \{\xrightarrow{a} \mid a \in A\}, \{W_p \mid p \in P\})$ and $\mathfrak{S}' = (W', \{\xrightarrow{a'} \mid a \in A\}, \{W'_p \mid p \in P\})$ be two (A, P) -structures. A *bisimulation between \mathfrak{S} and \mathfrak{S}'* is a binary relation $R \subseteq W \times W'$ such that for each $(s, s') \in R$ the following holds:

- (1) $s \in W_p$ if and only if $s' \in W'_p$ for all $p \in P$,
- (2) for each $s \xrightarrow{a} t$ there exists $s' \xrightarrow{a'} t'$ such that $(t, t') \in R$, and
- (3) for each $s' \xrightarrow{a'} t'$ there exists $s \xrightarrow{a} t$ such that $(t, t') \in R$.

In case there is a bisimulation R between \mathfrak{S} and \mathfrak{S}' with $(s, s') \in R$ we say that (\mathfrak{S}, s) is *bisimilar to* (\mathfrak{S}', s') and write $(\mathfrak{S}, s) \sim (\mathfrak{S}', s')$ or, in case \mathfrak{S} and \mathfrak{S}' are clear from the context, $s \sim s'$. It is well-known that modal logic cannot distinguish between bisimilar

structures, i.e., if $(\mathfrak{S}, s) \sim (\mathfrak{S}', s')$ then $(\mathfrak{S}, s) \models \varphi$ if and only if $(\mathfrak{S}', s') \models \varphi$ for all modal logic formulas over (A, P) , see for instance [1]. The following proposition (which is straightforward to prove) lifts this statement to many-dimensional modal logics, i.e., modal logic formulas cannot distinguish between interpreted products structures whose components are bisimilar.

PROPOSITION 2.2. *Let $A = \bigcup_{i=1}^d A_i$, $P = \bigcup_{i=1}^d P_i$ and for each $i \in [1, d]$ assume two pointed (A_i, P_i) -structures (\mathfrak{S}_i, w_i) and (\mathfrak{S}'_i, w'_i) with $(\mathfrak{S}_i, w_i) \sim (\mathfrak{S}'_i, w'_i)$. Then for every interpretation $\sigma : P \rightarrow \text{bool}(P)$ we have*

$$\left(\prod_{i \in [1, d]}^{\sigma} \mathfrak{S}_i, \langle w_1, \dots, w_d \rangle \right) \sim \left(\prod_{i \in [1, d]}^{\sigma} \mathfrak{S}'_i, \langle w'_1, \dots, w'_d \rangle \right).$$

In particular $(\prod_{i \in [1, d]}^{\sigma} \mathfrak{S}_i, \langle w_1, \dots, w_d \rangle)$ and $(\prod_{i \in [1, d]}^{\sigma} \mathfrak{S}'_i, \langle w'_1, \dots, w'_d \rangle)$ satisfy the same $\mathbf{K}^d(A, P)$ -formulas.

PROOF. Assume that

$$\begin{aligned} \mathfrak{S}_i &= (W_i, \{\xrightarrow{a}_i \mid a \in A_i\}, \{W_{p,i} \mid p \in P_i\}) \text{ and} \\ \mathfrak{S}'_i &= (W'_i, \{\xrightarrow{a}'_i \mid a \in A_i\}, \{W'_{p,i} \mid p \in P_i\}). \end{aligned}$$

Let $\sigma : P \rightarrow \text{bool}(P)$ be an interpretation, and let $\mathfrak{S} = (W, \{\xrightarrow{a} \mid a \in A\}, \{W_p \mid p \in P\})$ and $\mathfrak{S}' = (W', \{\xrightarrow{a}' \mid a \in A\}, \{W'_p \mid p \in P\})$ be the σ -products of the \mathfrak{S}_i and \mathfrak{S}'_i , respectively. Assume $R_i \subseteq W_i \times W'_i$ is a bisimulation between \mathfrak{S}_i and \mathfrak{S}'_i with $(w_i, w'_i) \in R_i$ for each $i \in [1, d]$. We claim that

$$R = \{(\langle u_1, \dots, u_d \rangle, \langle u'_1, \dots, u'_d \rangle) \mid \forall i \in [1, d] : (u_i, u'_i) \in R_i\}$$

is a bisimulation. For this, let $\bar{u} = \langle u_1, \dots, u_d \rangle \in \prod_{i \in [1, d]} W_i$ and $\bar{u}' = \langle u'_1, \dots, u'_d \rangle \in \prod_{i \in [1, d]} W'_i$ with $(\bar{u}, \bar{u}') \in R$. Hence $u_i \sim u'_i$ and thus $u_i \in W_{p,i} \Leftrightarrow u'_i \in W'_{p,i}$ for each $p \in P_i$ and each $i \in [1, d]$. By definition of interpreted product structures, we get $\bar{u} \in W_p$ if and only if $\bar{u}' \in W'_p$ for each $p \in P$. This establishes point (1) of R being a bisimulation. For proving point (2), let us assume $\bar{u} \xrightarrow{a} \bar{v}$, where $\bar{v} = \langle v_1, \dots, v_d \rangle \in \prod_{i \in [1, d]} W_i$. Then there exists some $i \in [1, d]$ such that $a \in A_i$, $u_i \xrightarrow{a} v_i$ and $v_j = u_j$ for all $j \in [1, d] \setminus \{i\}$. Since $(u_i, u'_i) \in R_i$ and $u_i \xrightarrow{a} v_i$ there exists some $v'_i \in W'_i$ such that $u'_i \xrightarrow{a}' v'_i$ and $(v_i, v'_i) \in R_i$. Choose $\bar{v}' = \langle v'_1, \dots, v'_d \rangle$, where $v'_j = v_j$ for each $j \in [1, d] \setminus \{i\}$ and v'_i as mentioned above. Hence from $(v_j, v'_j) \in R_j$ for each $j \in [1, d]$ it follows $(\bar{v}, \bar{v}') \in R$ by definition of R . Moreover we have $\bar{v} \xrightarrow{a}' \bar{v}'$, thus point (2) follows. Point (3) can be proven analogously. In particular, we have $(\mathfrak{S}, \langle w_1, \dots, w_d \rangle) \sim (\mathfrak{S}', \langle w'_1, \dots, w'_d \rangle)$ and, thus, $(\mathfrak{S}, \langle w_1, \dots, w_d \rangle)$ and $(\mathfrak{S}', \langle w'_1, \dots, w'_d \rangle)$ satisfy the same $\mathbf{K}^d(A, P)$ -formulas. \square

3. \mathbf{K}^2 -SAT IS HARD

The goal of this section is to show a nonelementary lower bound for \mathbf{K}^2 -SAT. As a nonelementary upper bound is already known, we close the complexity gap for this problem. By Proposition 2.1 it suffices to show hardness for \mathbf{K}_{id}^2 -SAT. As a necessary preliminary step we show how to enforce (nonelementarily) big models in \mathbf{K}_{id}^2 . Using this, we prove via a standard reduction from an appropriate tiling problems that \mathbf{K}_{id}^2 -SAT is nonelementary.

In this section, we will only deal with id-products of two structures \mathfrak{S} and $\bar{\mathfrak{S}}$ (over disjoint sets of propositions and actions). To simplify notation we write $\mathfrak{S} \times \bar{\mathfrak{S}}$ for $\mathfrak{S} \times^{\text{id}} \bar{\mathfrak{S}}$.

Recall the tower function $\text{Tower} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ defined as

$$\text{Tower}(0, n) = n \text{ and } \text{Tower}(\ell + 1, n) = 2^{\text{Tower}(\ell, n)}$$

for each $\ell, n \in \mathbb{N}$. In this section, we construct a family $\{\varphi_{\ell, n} \mid \ell, n \geq 1\}$ of $\mathbf{K}^2(\mathbf{A}, \mathbf{P})$ -formulas (\mathbf{A} and \mathbf{P} are specified later) such that for each $\ell, n \in \mathbb{N}$ the following hold:

- (i) $\mathbf{P} = \mathbf{P}_n \uplus \overline{\mathbf{P}_n}$, where \mathbf{P}_n is to be defined and $\overline{\mathbf{P}_n}$ is a disjoint copy of \mathbf{P}_n ,
- (ii) $\mathbf{A} = \{a\} \uplus \{\bar{a}\}$,
- (iii) $|\varphi_{\ell, n}| \leq \exp(\ell) \cdot \text{poly}(n)$, and
- (iv) if $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s, \bar{s} \rangle) \models \varphi_{\ell, n}$, then $|\mathfrak{S}|, |\overline{\mathfrak{S}}| \geq \text{Tower}(\ell, n)$.

Informally speaking, our intention is that if $\ell \geq 1$ and $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s, \bar{s} \rangle) \models \varphi_{\ell, n}$ then (\mathfrak{S}, s) (resp., $(\overline{\mathfrak{S}}, \bar{s})$) will be bisimilar to a particular tree over $(\{a\}, \mathbf{P}_n)$ (resp., $(\{\bar{a}\}, \overline{\mathbf{P}_n})$) that will encode a value in the interval $[0, \text{Tower}(\ell + 1, n) - 1]$ and whose root has $\text{Tower}(\ell, n)$ children. When $\ell = 0$ we will use single-node trees to encode values in the interval $[0, 2^n - 1] = [0, \text{Tower}(1, n) - 1]$.

Let us describe in detail how we encode values in the interval $[0, \text{Tower}(\ell + 1, n) - 1]$ by trees over the signature $(\{a\}, \mathbf{P}_n)$, where

$$\mathbf{P}_n \stackrel{\text{def}}{=} \{b_0, \dots, b_{n-1}, b, \min_b, \min_b^{\leftarrow}, \min_{\neg b}, \min_{\neg b}^{\leftarrow}\}.$$

Definition 3.1. In case $\ell = 0$ we encode a value j from the interval $[0, 2^n - 1] = [0, \text{Tower}(1, n) - 1]$ by using the following single-node tree that we call $\Upsilon_{0, n}^{(0)}(j)$: It consists of a single node v that is (only) labeled with the subset X of the propositions $\{b_0, \dots, b_{n-1}\}$ such that $j = \sum_{b_i \in X} 2^i \in [0, 2^n - 1]$. By $\Upsilon_{0, n}^{(1)}(j)$ we denote the single-node tree that one obtains from $\Upsilon_{0, n}^{(0)}(j)$ by additionally labeling v with b .

For $\ell > 0$ we inductively encode a value j from the interval $[0, \text{Tower}(\ell + 1, n) - 1]$ by the following tree $\Upsilon_{\ell, n}^{(0)}(j)$ that arises as follows:


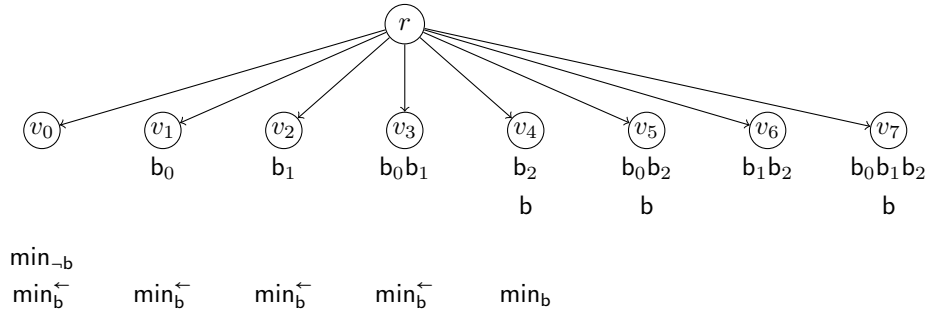
- (i) Let $I \subseteq [0, \text{Tower}(\ell, n) - 1]$ be the unique set such that $j = \sum_{i \in I} 2^i$ and take the disjoint union of a root r and all trees from the set

$$U = \{\Upsilon_{\ell-1, n}^{(1)}(i) \mid i \in I\} \cup \{\Upsilon_{\ell-1, n}^{(0)}(i) \mid i \notin I\}.$$

- (ii) Draw an a -transition from r to the root of each tree in U . Hence r has $\text{Tower}(\ell, n)$ children.
- (iii) For the propositions \min_b and \min_b^{\leftarrow} , proceed as follows: If $I \neq \emptyset$, then define $m_b = \min(I)$; otherwise set $m_b = \text{Tower}(\ell, n)$. If $m_b < \text{Tower}(\ell, n)$, then add the label \min_b to the root of $\Upsilon_{\ell-1, n}^{(1)}(m_b) \in U$. Moreover, add the label \min_b^{\leftarrow} to the root of $\Upsilon_{\ell-1, n}^{(0)}(m') \in U$ for every $m' \in [0, m_b - 1]$.
- (iv) For the propositions $\min_{\neg b}$ and $\min_{\neg b}^{\leftarrow}$, proceed as follows: If $I \neq [0, \text{Tower}(\ell, n) - 1]$, define $m_{\neg b} = \min([0, \text{Tower}(\ell, n) - 1] \setminus I)$; otherwise set $m_{\neg b} = \text{Tower}(\ell, n)$. If $m_{\neg b} < \text{Tower}(\ell, n)$, then add the label $\min_{\neg b}$ to the root of $\Upsilon_{\ell-1, n}^{(0)}(m_{\neg b}) \in U$. Moreover, add the label $\min_{\neg b}^{\leftarrow}$ to the root of $\Upsilon_{\ell-1, n}^{(1)}(m') \in U$ for every $m' \in [0, m_{\neg b} - 1]$.

Let $\Upsilon_{\ell, n}^{(1)}(j)$ denote the tree that one obtains from $\Upsilon_{\ell, n}^{(0)}(j)$ by adding the label b to its root.

Figure 1 shows the tree $\Upsilon_{1, 3}^{(0)}(175)$. First observe that the root r has $8 = 2^3 = \text{Tower}(1, 3)$ successors v_0, \dots, v_7 . Next note that in each v_i the evaluation of the propositions b_j ($j \in [0, 2]$) gives a binary number equal to i . For instance, in v_4 only b_2 is true,


 Fig. 1. The tree $\Upsilon_{1,3}^{(0)}(175)$ with root r .

 Fig. 2. The tree $\Upsilon_{1,3}^{(0)}(176)$ with root r .

hence the corresponding binary number encodes 4. As indicated, the evaluation of b at the children of r gives rise to the binary number 11110101 (least significant bit to the left) which equals 175. By definition, \min_b (resp., \min_{-b}) holds in the minimal position where b holds (resp., does not hold) and \min_b^{\leftarrow} (resp., \min_{-b}^{\leftarrow}) holds in all positions left of \min_b (resp., \min_{-b}). These auxiliary propositions will be useful for enforcing the successor relation. For example, observe that the labelings of b in $\Upsilon_{(1,3)}^{(0)}(175)$ and $\Upsilon_{(1,3)}^{(0)}(176)$, depicted in Figure 2, differ only on positions that were labeled with \min_{-b} and \min_b^{\leftarrow} in Figure 1.

It is worth mentioning that the defined trees $\Upsilon_{\ell,n}(j)$ are similar to the trees $\mathcal{T}(j)$ introduced in [20, Chapter 10] and used for example in [12]. In particular, they both represent the number j and have small depth, but high outdegree. However, there are some differences. Note first that the root of $\mathcal{T}(j)$ has a child for those numbers i such that the i -th bit in j is 1. In contrast, the root of $\Upsilon_{\ell,n}(j)$ has, independent of j , $\text{Tower}(\ell - 1, n)$ children each corresponding to one bit position and the bits set to 1 are marked with the proposition b . Moreover, as we use two-dimensional modal logic instead of first-order logic as in [20] to enforce our trees, we face two problems: First, we cannot enforce them up to isomorphism but only up to bisimulation equivalence. Second, as the logic is much weaker, we need some auxiliary propositional variables (or unary predicates). The particular difficulty is expressing a “less-than” or “successor” predicate and it seems impossible to do this without the auxiliary propositions, already for $\ell = 1$.

Conventions. In the following, let $\overline{\Upsilon}_{\ell,n}^{(\beta)}(j)$ denote the corresponding copy of $\Upsilon_{\ell,n}^{(\beta)}(j)$ over the signature $(\{\overline{a}\}, \overline{P}_n)$, i.e. the tree that one obtains from the tree $\Upsilon_{\ell,n}^{(\beta)}(j)$ by replacing every action label a by \overline{a} and every proposition $p \in P_n$ by $\overline{p} \in \overline{P}_n$. For simplicity, we write $\diamond, \square, \overline{\diamond}, \overline{\square}$ for the modalities $\diamond_a, \square_a, \diamond_{\overline{a}}, \square_{\overline{a}}$, respectively.

We say that a tree \mathfrak{T} is an *extension* of $\Upsilon_{\ell,n}^{(\beta)}(j)$ if \mathfrak{T} evolves from $\Upsilon_{\ell,n}^{(\beta)}(j)$ by adding an arbitrary subset (possibly empty) of $\{b, \min_b, \min_b^{\leftarrow}, \min_{\neg b}, \min_{\neg b}^{\leftarrow}\}$ to the root of \mathfrak{T} . Note that in particular $\Upsilon_{\ell,n}^{(1)}(j)$ is an extension of $\Upsilon_{\ell,n}^{(0)}(j)$. Extensions of $\overline{\Upsilon}_{\ell,n}^{(b)}(j)$ are defined analogously.

Before we define the formulas $\varphi_{\ell,n}$, we introduce auxiliary formulas $\text{eq}_{\ell,n}$, $\text{first}_{\ell,n}$, $\text{last}_{\ell,n}$, and $\text{succ}_{\ell,n}$ whose names indicate their meaning on the asynchronous (id-)product of two trees that are extensions of $\Upsilon_{\ell,n}^{(0)}(j_1)$ and $\overline{\Upsilon}_{\ell,n}^{(0)}(j_2)$, respectively, where $j_1, j_2 \in [0, \text{Tower}(\ell + 1, n) - 1]$.

For $\ell = 0$ they are defined as follows:

$$\begin{aligned} \text{eq}_{0,n} &\stackrel{\text{def}}{=} \bigwedge_{i \in [0, n-1]} b_i \leftrightarrow \overline{b}_i \\ \text{first}_{0,n} &\stackrel{\text{def}}{=} \bigwedge_{i \in [0, n-1]} \neg b_i \wedge \neg \overline{b}_i \\ \text{last}_{0,n} &\stackrel{\text{def}}{=} \bigwedge_{i \in [0, n-1]} b_i \wedge \overline{b}_i \\ \text{succ}_{0,n} &\stackrel{\text{def}}{=} \bigvee_{i \in [0, n-1]} (\neg b_i \wedge \overline{b}_i \wedge \bigwedge_{j \in [0, i-1]} (b_j \wedge \neg \overline{b}_j) \wedge \bigwedge_{j \in [i+1, n-1]} b_j \leftrightarrow \overline{b}_j) \end{aligned}$$

For $\ell > 0$ we define them as follows:

$$\begin{aligned} \text{eq}_{\ell,n} &\stackrel{\text{def}}{=} \square \overline{\square} (\text{eq}_{\ell-1,n} \rightarrow (b \leftrightarrow \overline{b})) \\ \text{first}_{\ell,n} &\stackrel{\text{def}}{=} \square \neg b \wedge \square \neg \overline{b} \\ \text{last}_{\ell,n} &\stackrel{\text{def}}{=} \square b \wedge \square \overline{b} \\ \text{succ}_{\ell,n} &\stackrel{\text{def}}{=} \diamond \neg b \wedge \square \overline{\square} (\text{eq}_{\ell-1,n} \rightarrow ((\min_{\neg b} \leftrightarrow \overline{\min_{\neg b}}) \wedge ((\neg \min_{\neg b}^{\leftarrow} \wedge \neg \min_{\neg b}) \rightarrow (b \leftrightarrow \overline{b})))) \end{aligned}$$

The following lemma shows that the formulas indeed express what they suggest to express. Recall that we identify a tree \mathfrak{T} with the pointed structure (\mathfrak{T}, r) , where r is the root of \mathfrak{T} . Hence, in the following lemma $\mathfrak{T} \times \overline{\mathfrak{T}}$ is identified with the pointed structure $(\mathfrak{T} \times \overline{\mathfrak{T}}, \langle r, \overline{r} \rangle)$, where r (resp., \overline{r}) is the root of \mathfrak{T} (resp., $\overline{\mathfrak{T}}$).

LEMMA 3.2. *Let $\ell, n \geq 0$, let $j_1, j_2 \in [0, \text{Tower}(\ell + 1, n) - 1]$ and let \mathfrak{T} and $\overline{\mathfrak{T}}$ be an extension of $\Upsilon_{\ell,n}^{(0)}(j_1)$ and $\overline{\Upsilon}_{\ell,n}^{(0)}(j_2)$, respectively. Then the following holds:*

- (a) $\mathfrak{T} \times \overline{\mathfrak{T}} \models \text{eq}_{\ell,n}$ if and only if $j_1 = j_2$.
- (b) $\mathfrak{T} \times \overline{\mathfrak{T}} \models \text{first}_{\ell,n}$ if and only if $j_1 = j_2 = 0$.
- (c) $\mathfrak{T} \times \overline{\mathfrak{T}} \models \text{last}_{\ell,n}$ if and only if $j_1 = j_2 = \text{Tower}(\ell + 1, n) - 1$.
- (d) $\mathfrak{T} \times \overline{\mathfrak{T}} \models \text{succ}_{\ell,n}$ if and only if $j_2 = j_1 + 1$.

PROOF. We show the statement by induction on ℓ . Let \mathfrak{T} and $\overline{\mathfrak{T}}$ be as in the lemma. We assume w.l.o.g. that $\mathfrak{T} = \Upsilon_{\ell,n}^{(0)}(j_1)$ and $\overline{\mathfrak{T}} = \overline{\Upsilon}_{\ell,n}^{(0)}(j_2)$: Each of the formulas $\text{eq}_{\ell,n}$, $\text{first}_{\ell,n}$, $\text{last}_{\ell,n}$ and $\text{succ}_{\ell,n}$ holds in $\Upsilon_{\ell,n}^{(0)}(j_1) \times \overline{\Upsilon}_{\ell,n}^{(0)}(j_2)$ if and only if it holds in $\mathfrak{T} \times \overline{\mathfrak{T}}$.

In the following let r and \bar{r} denote the root of $\Upsilon_{\ell,n}^{(0)}(j_1)$ and of $\overline{\Upsilon}_{\ell,n}^{(0)}(j_2)$, respectively.

For the induction base let $\ell = 0$. For (a) we have $(\Upsilon_{\ell,n}^{(0)}(j_1) \times \overline{\Upsilon}_{\ell,n}^{(0)}(j_2)) \models \text{eq}_{0,n}$ if and only if $(b_i \text{ holds in } r \Leftrightarrow \bar{b}_i \text{ holds in } \bar{r})$ for all $i \in [0, n-1]$ if and only if $j_1 = j_2$. Both (b) and (c) can be proven in analogy to (a). For (d) we have $(\Upsilon_{\ell,n}^{(0)}(j_1) \times \overline{\Upsilon}_{\ell,n}^{(0)}(j_2)) \models \text{succ}_{0,n}$ if and only if there is some $i \in [0, n-1]$ such that r does not satisfy b_i , \bar{r} satisfies \bar{b}_i , r satisfies b_j and \bar{r} does not satisfy \bar{b}_j for each $j \in [0, i-1]$ and moreover $(r \text{ satisfies } b_j \Leftrightarrow \bar{r} \text{ satisfies } \bar{b}_j)$ for each $j \in [i+1, n-1]$. The latter is equivalent to $j_2 = j_1 + 1$.

For the induction step let $\ell > 0$. The cases (a), (b) and (c) are straightforward. Let us prove case (d). For $j \in [0, \text{Tower}(\ell, n) - 1]$ let $c_j, d_j \in \{0, 1\}$ such that $\Upsilon_{\ell-1,n}^{(c_j)}(j)$ (resp., $\overline{\Upsilon}_{\ell-1,n}^{(d_j)}(j)$) is a subtree of \mathfrak{T} (resp., $\overline{\mathfrak{T}}$). The formula $\text{succ}_{\ell,n}$ states the following:

- There is a $j \in [0, \text{Tower}(\ell, n) - 1]$ such that $c_j = 0$.
- If j_0 is the minimal $j \in [0, \text{Tower}(\ell, n) - 1]$ such that $c_j = 0$, then j_0 is also the minimal $j \in [0, \text{Tower}(\ell, n) - 1]$ such that $d_j = 1$.
- For all $j_0 < j < \text{Tower}(\ell, n)$ we have $c_j = d_j$.

These conditions are equivalent to

$$1 + \sum_{j=0}^{\text{Tower}(\ell,n)-1} c_j 2^j = \sum_{j=0}^{\text{Tower}(\ell,n)-1} d_j 2^j,$$

i.e., $j_2 = j_1 + 1$. \square

Now we give a family of formulas $\varphi_{\ell,n}$ with the idea that for every model $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s, \bar{s} \rangle)$ of $\varphi_{\ell,n}$ there exists some $j \in [0, \text{Tower}(\ell+1, n) - 1]$ such that $(\mathfrak{S}, s) \sim \mathfrak{T}$ and $(\overline{\mathfrak{S}}, \bar{s}) \sim \overline{\mathfrak{T}}$, where \mathfrak{T} (resp. $\overline{\mathfrak{T}}$) is an extension of $\Upsilon_{\ell,n}^{(0)}(j)$ (resp., $\overline{\Upsilon}_{\ell,n}^{(0)}(j)$).

Definition 3.3. Set $\varphi_{0,n} = \text{eq}_{0,n} \wedge \Box \perp \wedge \overline{\Box} \perp$ and for each $\ell \geq 1$ define $\varphi_{\ell,n}$, by induction on ℓ , as the conjunction of the following formulas:

- (1) $\bigwedge_{i \in [0, n-1]} \neg b_i \wedge \neg \bar{b}_i$
- (2) $\Box \overline{\Diamond} \varphi_{\ell-1,n}$
- (3) $\overline{\Box} \Diamond \varphi_{\ell-1,n}$
- (4) $\Diamond \overline{\Diamond} (\varphi_{\ell-1,n} \wedge \text{first}_{\ell-1,n})$
- (5) $\Box (\overline{\Box} \neg \text{last}_{\ell-1,n} \rightarrow \overline{\Diamond} \text{succ}_{\ell-1,n})$
- (6) $\Box \overline{\Box} (\text{eq}_{\ell-1,n} \rightarrow \bigwedge_{p \in P_n} (p \leftrightarrow \bar{p}))$
- (7) $\Diamond (\min_{\neg b} \vee \min_{\neg \bar{b}}^{\leftarrow}) \wedge \Diamond (\min_b \vee \min_b^{\leftarrow})$
- (8) $\Box (((\min_{\neg b} \vee \min_{\neg \bar{b}}^{\leftarrow}) \rightarrow \neg b) \wedge ((\min_{\neg \bar{b}}^{\leftarrow} \vee \min_b) \rightarrow b))$
- (9) $\Box \overline{\Box} (\text{succ}_{\ell-1,n} \rightarrow \bigwedge_{x \in \{b, \neg b\}} ((\min_x \vee \min_x^{\leftarrow}) \leftrightarrow \min_x^{\leftarrow}))$

We are now ready to present our main theorem. Roughly speaking, it states that $\varphi_{\ell,n}$ enforces structures bisimilar to the product of extensions of $\Upsilon_{\ell,n}^{(0)}(j)$ and $\overline{\Upsilon}_{\ell,n}^{(0)}(j)$ for some j .

THEOREM 3.4. *For every $\ell \geq 0$ the following holds:*

- (a) $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s, \overline{s} \rangle) \models \varphi_{\ell,n}$ if and only if there exists $j \in [0, \text{Tower}(\ell+1, n) - 1]$ such that (\mathfrak{S}, s) is bisimilar to some extension of $\Upsilon_{\ell,n}^{(0)}(j)$ and $(\overline{\mathfrak{S}}, \overline{s})$ is bisimilar to some extension of $\overline{\Upsilon}_{\ell,n}^{(0)}(j)$.
- (b) $|\varphi_{\ell,n}| \leq 3^\ell \cdot \text{poly}(\ell, n)$ and the formula $\varphi_{\ell,n}$ is computable in time $3^\ell \cdot \text{poly}(\ell, n)$.
- (c) The switching depth of $\varphi_{\ell,n}$ is ℓ .

Before giving the complete formal proof of Theorem 3.4, we want to give some intuition. Parts (b) and (c) are straightforward consequences of the definition of $\varphi_{\ell,n}$. For Part (a) observe first that $\varphi_{\ell,n}$ is satisfiable: it is not hard to verify that the product of (any extension of) $\Upsilon_{\ell,n}^{(0)}(j)$ and (any extension of) $\overline{\Upsilon}_{\ell,n}^{(0)}(j)$ satisfies $\varphi_{\ell,n}$. The difficult part is to show that $\varphi_{\ell,n}$ enforces such models, i.e., each model of $\varphi_{\ell,n}$ is of the form $\mathfrak{T} \times \overline{\mathfrak{T}}$, where \mathfrak{T} (resp., $\overline{\mathfrak{T}}$) is bisimilar to an extension of $\Upsilon_{\ell,n}^{(0)}(j)$ (resp., an extension of $\overline{\Upsilon}_{\ell,n}^{(0)}(j)$) for some j . Obviously, this is the case for $\varphi_{0,n}$.

For $\ell > 0$, let $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s, \overline{s} \rangle) \models \varphi_{\ell,n}$. By induction, formula (2) implies that for each successor t of s it holds that (\mathfrak{S}, t) is bisimilar to an extension of $\Upsilon_{\ell-1,n}^{(0)}(i)$ for some i . Formula (3) implies the analogous property for every successor \overline{t} of \overline{s} .

Given this, formulas (3)-(5) together with Lemma 3.2 imply that for every $i \in [0, \text{Tower}(\ell, n) - 1]$ there is a child s_i of s such that (\mathfrak{S}, s_i) is bisimilar to an extension of $\Upsilon_{\ell-1,n}^{(0)}(i)$ and, analogously, there is a child \overline{s}_i of \overline{s} such that $(\overline{\mathfrak{S}}, \overline{s}_i)$ is bisimilar to an extension of $\overline{\Upsilon}_{\ell-1,n}^{(0)}(i)$: By formula (4) and Lemma 3.2, there are such s_0 and \overline{s}_0 . The existence of s_0 and formula (5) imply the existence of \overline{s}_1 . Formula (3) implies that there is some s_1 , and we can repeat the argument using formulas (5) and (3).

Observe now that, in principle, there might be children $s_i \neq s'_i$ of s such that (\mathfrak{S}, s_i) and (\mathfrak{S}, s'_i) are bisimilar to *different* extensions of $\Upsilon_{\ell-1,n}^{(0)}(i)$. This is ruled out by applying formula (6) twice: For any proposition $p \in P_n$ we have: p is satisfied in (\mathfrak{S}, s_i) if and only if \overline{p} is satisfied in $(\overline{\mathfrak{S}}, \overline{s}_i)$ if and only if p is satisfied in (\mathfrak{S}, s'_i) . Hence, we can talk about *the* children s_i and \overline{s}_i , respectively. Of course, children of s and \overline{s} can appear numerous times as copies.

The children s_i and \overline{s}_i encode binary numbers N and \overline{N} , respectively, in the natural way: The i -th bit¹ of N is 1 if and only if (\mathfrak{S}, s_i) satisfies b and analogously for \overline{N} . Note that formula (6) implies that $N = \overline{N}$.

Finally, formulas (7)-(9) ensure that the children s_i and \overline{s}_i are labeled with the propositions $\text{min}_b, \text{min}_{\neg b}, \text{min}_b^{\leftarrow}, \text{min}_{\neg b}^{\leftarrow}$ and $\overline{\text{min}}_b, \overline{\text{min}}_{\neg b}, \overline{\text{min}}_b^{\leftarrow}, \overline{\text{min}}_{\neg b}^{\leftarrow}$, respectively, in a way such that (\mathfrak{S}, s) is bisimilar to an extension of $\Upsilon_{\ell,n}^{(0)}(N)$ and $(\overline{\mathfrak{S}}, \overline{s})$ is bisimilar to an extension of $\overline{\Upsilon}_{\ell,n}^{(0)}(N)$. This is actually the most subtle part of the following proof of Theorem 3.4.

PROOF OF THEOREM 3.4. Part (c) is an immediate consequence of Definition 3.3.

¹The least significant bit is the 0-th bit.

We show Part (b) by induction on ℓ starting with $\ell = 0$. For $\varphi_{0,n} = \bigwedge_{i \in [0, n-1]} b_i \leftrightarrow \bar{b}_i \wedge \square \perp \wedge \bar{\square} \perp$ the statement is trivial. Let now be $\ell > 0$. The formula $\varphi_{\ell-1,n}$ occurs 3 times in $\varphi_{\ell,n}$. The auxiliary formulas $\text{succ}_{\ell-1,n}$, $\text{eq}_{\ell-1,n}$, $\text{last}_{\ell-1,n}$, and $\text{first}_{\ell-1,n}$ are all polynomially sized in ℓ and n . Thus, overall we get $|\varphi_{\ell,n}| = 3 \cdot |\varphi_{\ell-1,n}| + \text{poly}(\ell, n)$. Thus, we obtain by induction hypothesis $|\varphi_{\ell,n}| = 3^\ell \cdot \text{poly}(\ell, n)$.

Let us finally prove Part (a). With (1), (2), \dots , (9) we refer to the formulas from Definition 3.3.

If: We prove the “if”-direction by induction on ℓ .

For the induction base, assume $\ell = 0$. Assume some $j \in [0, \text{Tower}(1, n) - 1] = [0, 2^n - 1]$ such that (\mathfrak{S}, s) is bisimilar to some extension of $\Upsilon_{0,n}^{(0)}(j)$ and $(\bar{\mathfrak{S}}, \bar{s})$ is bisimilar to some extension of $\bar{\Upsilon}_{0,n}^{(0)}(j)$. It is clear that $(\mathfrak{S} \times \bar{\mathfrak{S}}, \langle s, \bar{s} \rangle) \models \square \perp \wedge \bar{\square} \perp$. Moreover, Proposition 2.2 and point (a) of Lemma 3.2 imply that $(\mathfrak{S} \times \bar{\mathfrak{S}}, \langle s, \bar{s} \rangle) \models \text{eq}_{0,n}$. Hence $(\mathfrak{S} \times \bar{\mathfrak{S}}, \langle s, \bar{s} \rangle) \models \varphi_{0,n}$.

For the induction step, assume $\ell \geq 1$. Let j be arbitrary in $[0, \text{Tower}(\ell + 1, n) - 1]$ and assume (\mathfrak{S}, s) is bisimilar to some extension of $\Upsilon_{\ell,n}^{(0)}(j)$ and $(\bar{\mathfrak{S}}, \bar{s})$ is bisimilar to some extension of $\bar{\Upsilon}_{\ell,n}^{(0)}(j)$. By Proposition 2.2, it is sufficient to show that $\Upsilon_{\ell,n}^{(0)}(j) \times \bar{\Upsilon}_{\ell,n}^{(0)}(j) \models \varphi_{\ell,n}$. Let r (resp., \bar{r}) be the root of $\Upsilon_{\ell,n}^{(0)}(j)$ (resp., $\bar{\Upsilon}_{\ell,n}^{(0)}(j)$).

Clearly, formula (1) holds in $\Upsilon_{\ell,n}^{(0)}(j) \times \bar{\Upsilon}_{\ell,n}^{(0)}(j)$ as neither $\Upsilon_{\ell,n}^{(0)}(j)$ satisfies any b_i nor does $\bar{\Upsilon}_{\ell,n}^{(0)}(j)$ satisfy any proposition \bar{b}_i .

For formula (2) let t be any child of r . By Definition 3.1, the subtree \mathfrak{T} rooted in t is an extension of $\Upsilon_{\ell-1,n}^{(0)}(i)$ for some $i \in [0, \text{Tower}(\ell, n) - 1]$. Also by Definition 3.1, there is a child \bar{t} of \bar{r} such that the subtree $\bar{\mathfrak{T}}$ rooted in \bar{t} is an extension of $\bar{\Upsilon}_{\ell-1,n}^{(0)}(i)$. By induction hypothesis, we have $\mathfrak{T} \times \bar{\mathfrak{T}} \models \varphi_{\ell-1,n}$. Formula (3) holds for analogous reasons.

By Definition 3.1, there is a child t (resp., \bar{t}) of r (resp., \bar{r}) such that the subtree \mathfrak{T} (resp., $\bar{\mathfrak{T}}$) rooted in t (resp., \bar{t}) is an extension of $\Upsilon_{\ell-1,n}^{(0)}(0)$ (resp., $\bar{\Upsilon}_{\ell-1,n}^{(0)}(0)$). Point (b) of Lemma 3.2 implies $\mathfrak{T} \times \bar{\mathfrak{T}} \models \text{first}_{\ell-1,n}$. Hence formula (4) is satisfied.

For formula (5) assume that t is an arbitrary child of r . By Definition 3.1, the subtree \mathfrak{T} rooted in t is an extension of $\Upsilon_{\ell-1,n}^{(0)}(i)$ for some $i \in [0, \text{Tower}(\ell, n) - 1]$. We distinguish the following cases on i .

- $i = \text{Tower}(\ell, n) - 1$. By Definition 3.1, there is a child \bar{t} of \bar{r} such that the subtree $\bar{\mathfrak{T}}$ rooted in \bar{t} is an extension of $\bar{\Upsilon}_{\ell-1,n}^{(0)}(i)$. By Point (c) of Lemma 3.2, $\mathfrak{T} \times \bar{\mathfrak{T}} \models \text{last}_{\ell-1,n}$.
- $i < \text{Tower}(\ell, n) - 1$. By Definition 3.1, there is a child \bar{t} of \bar{r} such that the subtree $\bar{\mathfrak{T}}$ rooted in \bar{t} is an extension of $\bar{\Upsilon}_{\ell-1,n}^{(0)}(i+1)$. By Point (d) of Lemma 3.2, $\mathfrak{T} \times \bar{\mathfrak{T}} \models \text{succ}_{\ell-1,n}$.

For formula (6) let t and \bar{t} be arbitrary successors of r and \bar{r} , respectively. There are $k, i \in [0, \text{Tower}(\ell, n) - 1]$ such that the subtree \mathfrak{T} (resp., $\bar{\mathfrak{T}}$) rooted in t (resp., \bar{t}) is an extension of $\Upsilon_{\ell-1,n}^{(0)}(k)$ (resp., $\bar{\Upsilon}_{\ell-1,n}^{(0)}(i)$). Now, assume that $\mathfrak{T} \times \bar{\mathfrak{T}} \models \text{eq}_{\ell-1,n}$. Point (a) of Lemma 3.2 implies that $k = i$. Note that by Definition 3.1 the propositions that hold in t and \bar{t} are uniquely determined by $k = i$. Hence, a proposition $p \in P_n$ holds in t if and only if \bar{p} holds in \bar{t} .

For formula (7) observe that Point (iii) in Definition 3.1 implies that there is either some child of r labeled with min_b or some child of r labeled with min_b^\leftarrow . Thus, $\Upsilon_{\ell,n}^{(0)}(j) \times$

$\overline{\Upsilon}_{\ell,n}^{(0)}(j) \models \diamond(\min_b \vee \min_b^{\leftarrow})$. Similarly, Point (iv) implies that there is either some child of r labeled with \min_{-b} or some child of r is labeled with \min_{-b}^{\leftarrow} . Hence, $\Upsilon_{\ell,n}^{(0)}(j) \times \overline{\Upsilon}_{\ell,n}^{(0)}(j) \models \diamond(\min_{-b} \vee \min_{-b}^{\leftarrow})$.

For formula (8) observe that, by Point (iii) and (iv) in Definition 3.1, every child of r that satisfies \min_{-b} or \min_b^{\leftarrow} does not satisfy b . Analogously, every child of r that satisfies \min_{-b}^{\leftarrow} or \min_b satisfies b .

For formula (9), let t and \bar{t} be arbitrary successors of r and \bar{r} , respectively. There are $i, \bar{i} \in [0, \text{Tower}(\ell, n) - 1]$ such that the subtree \mathfrak{T} (resp., $\overline{\mathfrak{T}}$) rooted in t (resp., \bar{t}) is an extension of $\Upsilon_{\ell-1,n}^{(0)}(i)$ (resp., $\overline{\Upsilon}_{\ell-1,n}^{(0)}(\bar{i})$). Now, assume that $\mathfrak{T} \times \overline{\mathfrak{T}} \models \text{succ}_{\ell-1,n}$. We need to show that $\mathfrak{T} \times \overline{\mathfrak{T}} \models \bigwedge_{x \in \{b, -b\}} ((\overline{\min_x} \vee \min_x^{\leftarrow}) \leftrightarrow \min_x^{\leftarrow})$. We only show it for $x = -b$, because the case $x = b$ can be proven analogously. By Point (d) of Lemma 3.2, we have $\bar{i} = i + 1$.

If $j = \text{Tower}(\ell + 1, n) - 1$, then by Point (iv) of Definition 3.1, every child of r (resp., \bar{r}) is labelled with \min_{-b}^{\leftarrow} (resp., \min_{-b}^{\leftarrow}). Hence, $\mathfrak{T} \times \overline{\mathfrak{T}} \models ((\overline{\min_{-b}} \vee \min_{-b}^{\leftarrow}) \leftrightarrow \min_{-b}^{\leftarrow})$ holds.

Now, assume that $j < \text{Tower}(\ell + 1, n) - 1$. Put $M = \text{Tower}(\ell, n) - 1$ and let t_0, \dots, t_M be the successors of r such that the subtree \mathfrak{T}_k rooted in t_k is an extension of $\Upsilon_{\ell-1,n}^{(0)}(k)$ for all $k \in [0, M]$. Analogously, define $\bar{t}_0, \dots, \bar{t}_M$ to be the successors of \bar{r} such that the subtree $\overline{\mathfrak{T}}_k$ rooted in \bar{t}_k is an extension of $\overline{\Upsilon}_{\ell-1,n}^{(0)}(k)$ for all $k \in [0, M]$. We have that $t = t_i$ and $\bar{t} = \bar{t}_{i+1}$. By Point (iv) of Definition 3.1, there is some $m \in [0, M]$ such that

- $\mathfrak{T}_m \models \min_{-b}$ and $\overline{\mathfrak{T}}_m \models \overline{\min_{-b}}$,
- $\mathfrak{T}_k \models \min_{-b}^{\leftarrow}$ and $\overline{\mathfrak{T}}_k \models \overline{\min_{-b}^{\leftarrow}}$ for all $k < m$, and
- $\mathfrak{T}_k \models \neg \min_{-b} \wedge \neg \min_{-b}^{\leftarrow}$ and $\overline{\mathfrak{T}}_k \models \neg \overline{\min_{-b}} \wedge \neg \overline{\min_{-b}^{\leftarrow}}$ for all $k > m$.

Now, it is easy to verify that $\mathfrak{T} \times \overline{\mathfrak{T}} = \mathfrak{T}_i \times \overline{\mathfrak{T}}_{i+1} \models (\overline{\min_{-b}} \vee \min_{-b}^{\leftarrow}) \leftrightarrow \min_{-b}^{\leftarrow}$.

Only-if: Let us prove the “only-if” direction by induction on ℓ .

For $\ell = 0$ assume $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s, \bar{s} \rangle) \models \varphi_{0,n}$. Thus, both s and \bar{s} do not have any outgoing transitions and due to Point (a) of Lemma 3.2 there exists some $j \in [0, \text{Tower}(1, n) - 1]$ such that (\mathfrak{S}, s) is bisimilar to some extension of $\Upsilon_{0,n}^{(0)}(j)$ and $(\overline{\mathfrak{S}}, \bar{s})$ is bisimilar to some extension of $\overline{\Upsilon}_{0,n}^{(0)}(j)$.

For the induction step, let us assume $\ell \geq 1$ and $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s, \bar{s} \rangle) \models \varphi_{\ell,n}$ and put $m = \text{Tower}(\ell, n) - 1$. For each $i \in [0, m]$ let \mathcal{E}_i denote the set of extensions of $\Upsilon_{\ell-1,n}^{(0)}(i)$ and by $\overline{\mathcal{E}}_i$ the set of extensions of $\overline{\Upsilon}_{\ell-1,n}^{(0)}(i)$.

Claim 1. For each successor t of s we have that (\mathfrak{S}, t) is bisimilar to some element from \mathcal{E}_i for some $i \in [0, m]$ and, conversely, for each $i \in [0, m]$ there is a successor s_i of s such that (\mathfrak{S}, s_i) is bisimilar to some element from \mathcal{E}_i . An analogous property holds for $\overline{\mathfrak{S}}$.

PROOF OF CLAIM 1. Let t be an arbitrary successor of s . By formula (2), there is a successor \bar{t} of \bar{s} such that $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle t, \bar{t} \rangle) \models \varphi_{\ell-1,n}$. By induction hypothesis, there is a $i \in [0, m]$ such that (\mathfrak{S}, t) is bisimilar to some extension of $\Upsilon_{\ell-1,n}^{(0)}(i)$. Analogous reasoning using formula (3) yields that for every successor \bar{t} of \bar{s} there is some $i \in [0, m]$ such that $(\overline{\mathfrak{S}}, \bar{t})$ is bisimilar to some extension of $\overline{\Upsilon}_{\ell-1,n}^{(0)}(i)$.

Moreover, formulas (3),(4), and (5) imply that there are successors s_0, \dots, s_m of s and $\bar{s}_0, \dots, \bar{s}_m$ of \bar{s} such that (\mathfrak{S}, s_i) is bisimilar to some element from \mathcal{E}_i and $(\overline{\mathfrak{S}}, \bar{s}_i)$ is bisimilar to some element from $\overline{\mathcal{E}}_i$: By formula (4) and point (b) of Lemma 3.2, there

are such s_0, \bar{s}_0 . By formula (5) and points (c) and (d) of Lemma 3.2, there is such an \bar{s}_1 . By formula (3) (and reasoning as above), there is such an s_1 . Inductively repeating the argument yields the claimed s_0, \dots, s_m and $\bar{s}_0, \dots, \bar{s}_m$. This proves Claim 1. \square

Claim 2. For each $i \in [0, m]$ there is a unique element $\mathfrak{E}_i \in \mathcal{E}_i$ with the following property: for every successor t of s such that (\mathfrak{S}, t) is bisimilar to some element from \mathcal{E}_i , we have that (\mathfrak{S}, t) is bisimilar to \mathfrak{E}_i . In the same way, there are elements $\bar{\mathfrak{E}}_i \in \bar{\mathcal{E}}_i$.

PROOF OF CLAIM 2. Consider two successors t, t' of s such that both are bisimilar to elements from \mathcal{E}_i . It suffices to prove that (\mathfrak{S}, s_i) and (\mathfrak{S}, s'_i) satisfy the same propositions since both are bisimilar to extensions of $\Upsilon_{\ell-1, n}^{(0)}(i)$. By applying formula (6) twice, we have for each proposition $p \in P_n$ that $(\mathfrak{S}, t) \models p$ iff $(\bar{\mathfrak{S}}, \bar{s}_i) \models \bar{p}$ iff $(\mathfrak{S}, t') \models p$. This shows Claim 2. \square

Fix the elements $\mathfrak{E}_i, \bar{\mathfrak{E}}_i, i \in [0, m]$ that exist due to Claim 2. For each $i \in [0, m]$ let $x_i = 1$ if the tree \mathfrak{E}_i is an extension of $\Upsilon_{\ell-1, n}^{(1)}(i)$ and $x_i = 0$ otherwise. Note that by formula (6) the latter is equivalent to saying that for each $i \in [0, m]$ we set $x_i = 1$ if the tree $\bar{\mathfrak{E}}_i$ is an extension of $\bar{\Upsilon}_{\ell-1, n}^{(1)}(i)$ and $x_i = 0$ otherwise. Let

$$j = \sum_{i \in [0, m]} x_i \cdot 2^i \in [0, \text{Tower}(\ell + 1, n) - 1].$$

Moreover, let X denote the set of atomic propositions that hold in $(\mathfrak{S} \times \bar{\mathfrak{S}}, \langle s, \bar{s} \rangle)$. Let

$$Y = \{\mathfrak{b}, \min_{\mathfrak{b}}, \min_{\mathfrak{b}}^{\leftarrow}, \min_{\neg \mathfrak{b}}, \min_{\neg \mathfrak{b}}^{\leftarrow}\} \quad \text{and} \quad \bar{Y} = \{\bar{\mathfrak{b}}, \overline{\min_{\mathfrak{b}}}, \overline{\min_{\mathfrak{b}}^{\leftarrow}}, \overline{\min_{\neg \mathfrak{b}}}, \overline{\min_{\neg \mathfrak{b}}^{\leftarrow}}\}.$$

By formula (1), we have $X = Q \cup \bar{Q}$ for some $Q \subseteq Y$ and some $\bar{Q} \subseteq \bar{Y}$.

Let us define the tree \mathfrak{T} as the unique extension of $\Upsilon_{\ell, n}^{(0)}(j)$ satisfying precisely the propositions Q in its root t . Analogously, we define the tree $\bar{\mathfrak{T}}$ as the extension of $\bar{\Upsilon}_{\ell, n}^{(0)}(j)$ that satisfies precisely the propositions \bar{Q} in its root \bar{t} . Moreover, for each $i \in [0, m]$ let \mathfrak{T}_i be the unique subtree of \mathfrak{T} that is an extension of $\Upsilon_{\ell-1, n}^{(0)}(i)$ and define $\bar{\mathfrak{T}}_i$ analogously.

Claim 3. $(\mathfrak{T}, t) \sim (\mathfrak{S}, s)$ and $(\bar{\mathfrak{T}}, \bar{t}) \sim (\bar{\mathfrak{S}}, \bar{s})$.

PROOF OF CLAIM 3. We start with $(\mathfrak{T}, t) \sim (\mathfrak{S}, s)$. Note that s and t satisfy the same atomic propositions (namely those from Q) by definition of \mathfrak{T} . Thus, it remains to show the ‘back-and-forth’ condition of bisimulation. By Claim 2, for every $u \in N_{\mathfrak{S}}(s)$ there exists $i \in [0, m]$ such that $(\mathfrak{S}, u) \sim \mathfrak{E}_i$, and every child of t is the root of a tree \mathfrak{T}_j . Hence, it is sufficient to show that $\mathfrak{E}_i \sim \mathfrak{T}_i$ holds for each $i \in [0, m]$. But since both \mathfrak{E}_i and \mathfrak{T}_i are extensions of $\Upsilon_{\ell-1, n}^{(0)}(i)$ it is sufficient to prove that \mathfrak{E}_i and \mathfrak{T}_i satisfy the same set of atomic propositions in their roots. Note that the same arguments apply to $(\bar{\mathfrak{T}}, \bar{t}) \sim (\bar{\mathfrak{S}}, \bar{s})$, that is, it suffices to show that $\bar{\mathfrak{E}}_i$ and $\bar{\mathfrak{T}}_i$ satisfy the same propositions. This is, however, an immediate consequence: $\bar{\mathfrak{E}}_i \models \bar{p}$ iff $\mathfrak{E}_i \models p$ iff $\mathfrak{T}_i \models p$ iff $\bar{\mathfrak{T}}_i \models \bar{p}$, where the first equivalence is due to formula (6) and the last follows from Definition 3.1.

Since \mathfrak{T} is an extension of $\Upsilon_{\ell, n}^{(0)}(j)$, where $j = \sum_{i \in [0, m]} x_i \cdot 2^i \in [0, \text{Tower}(\ell + 1, n) - 1]$, the roots of \mathfrak{E}_i and \mathfrak{T}_i agree on all atomic propositions from $\{\mathfrak{b}_0, \dots, \mathfrak{b}_{n-1}, \mathfrak{b}\}$. Next, let us prove that the roots of \mathfrak{E}_i and \mathfrak{T}_i are labelled with the same atomic propositions from $\{\min_{\neg \mathfrak{b}}, \min_{\neg \mathfrak{b}}^{\leftarrow}\}$. That they are labelled with the same atomic propositions from $\{\min_{\mathfrak{b}}, \min_{\mathfrak{b}}^{\leftarrow}\}$ can be proven analogously.

First, note that we have the following:

- (a) For each $i \in [0, m]$ we have that \mathfrak{E}_i does not satisfy both \min_{-b} and \min_{-b}^{\leftarrow} since otherwise this would imply $\mathfrak{E}_i \models b \wedge \neg b$ by formula (8), a contradiction.
- (b) If there exists some $i_0 \in [0, m]$ such that $\mathfrak{E}_{i_0} \models \min_{-b} \vee \min_{-b}^{\leftarrow}$, then $\mathfrak{E}_i \models \min_{-b}^{\leftarrow}$ for all $i \in [0, i_0 - 1]$: By formula (6), we have $\mathfrak{E}_{i_0} \models \min_{-b} \vee \min_{-b}^{\leftarrow}$. By formula (9), we obtain $\mathfrak{E}_{i_0-1} \models \min_{-b}^{\leftarrow}$. This argument can be continued inductively.

Since $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s, \bar{s} \rangle)$ satisfies formula (7) we can distinguish two cases:

Case 1. There exists some $i_0 \in [0, m]$ such that $\mathfrak{E}_{i_0} \models \min_{-b}$. Then we proceed as follows:

- (c) $\mathfrak{E}_{i_0} \models \neg b$ by formula (8).
- (d) $\mathfrak{E}_i \models \min_{-b}^{\leftarrow}$ for each $i \in [0, i_0 - 1]$ by Point (b) above and hence $\mathfrak{E}_i \models b$ by formula (8) for each $i \in [0, i_0 - 1]$.
- (e) $\mathfrak{E}_i \not\models (\min_{-b}^{\leftarrow} \vee \min_{-b})$ for each $i \in [i_0 + 1, m]$. Assume the contrary, namely $\mathfrak{E}_i \models (\min_{-b}^{\leftarrow} \vee \min_{-b})$ for some $i \in [i_0 + 1, m]$. By Point (b) above, we get in particular $\mathfrak{E}_{i_0} \models \min_{-b}^{\leftarrow}$, which is a contradiction due to Point (a).

By comparing the above Points (a) to (e) with Point (iv) of Definition 3.1, one sees that \mathfrak{E}_i and \mathfrak{T}_i satisfy the same propositions from $\{\min_{-b}, \min_{-b}^{\leftarrow}\}$ for each $i \in [0, m]$.

Case 2. There does not exist any $i \in [0, m]$ such that $\mathfrak{E}_i \models \min_{-b}$. Then we proceed as follows:

- (c') There exists some $i_0 \in [0, m]$ such that $\mathfrak{E}_{i_0} \models \min_{-b}^{\leftarrow}$ by formula (7) and hence $\mathfrak{E}_{i_0} \models b$ by formula (8).
 - (d') $\mathfrak{E}_i \models \min_{-b}^{\leftarrow}$ for each $i \in [0, i_0 - 1]$ by Point (b) above and hence $\mathfrak{E}_i \models b$ by formula (8) for each $i \in [0, i_0 - 1]$.
 - (e') $\mathfrak{E}_i \models \min_{-b}^{\leftarrow}$ and thus $\mathfrak{E}_i \models b$ by formula (8) for each $i \in [i_0 + 1, m]$.
- This is proven by applying the following three steps inductively:
- (i) By formula (9), we have $\mathfrak{E}_{i_0+1} \models (\min_{-b} \vee \min_{-b}^{\leftarrow})$.
 - (ii) Assume $\mathfrak{E}_{i_0+1} \models \min_{-b}$. By formula (6), we have $\mathfrak{E}_{i_0+1} \models \min_{-b}$, which is a contradiction. Thus, $\mathfrak{E}_{i_0+1} \models \min_{-b}^{\leftarrow}$.
 - (iii) By formula (6), we have $\mathfrak{E}_{i_0+1} \models \min_{-b}^{\leftarrow}$.

By comparing the above Points (a), (b), and (c') to (e') to Point (iv) of Definition 3.1 one sees that \mathfrak{E}_i and \mathfrak{T}_i satisfy the same propositions from $\{\min_{-b}, \min_{-b}^{\leftarrow}\}$ for each $i \in [0, m]$. This concludes the proof of Claim 3. \square

Claim 3 shows statement (a) from Theorem 3.4. \square

We are finally ready to proceed to the main result of this section. By making use of the models that are enforced by the formulas $\varphi_{\ell, n}$, we can encode big numbers. In the proof of the following proposition we use these numbers to encode big tiling problems. Let ℓ -NEXP = NTIME(Tower(ℓ , poly(n))) for $\ell \geq 0$.

PROPOSITION 3.5. *The following holds:*

- For each $\ell \geq 1$, \mathbf{K}_{id}^2 -SAT restricted to formulas of switching depth ℓ is ℓ -NEXP-hard under polynomial time many-one reductions.
- In particular, \mathbf{K}_{id}^2 -SAT is nonelementary.

For the proof of Proposition 3.5, we need to introduce tilings and the tiling problem. A *tiling system* is a tuple $S = (\Theta, \mathbb{H}, \mathbb{V})$, where Θ is a finite set of *tile types*, $\mathbb{H} \subseteq \Theta \times \Theta$ is a *horizontal matching relation*, and $\mathbb{V} \subseteq \Theta \times \Theta$ is a *vertical matching relation*. A mapping $\tau : [0, k - 1] \times [0, k - 1] \rightarrow \Theta$ (where $k \geq 0$) is a *k-solution* for S if for all $(x, y) \in [0, k - 1] \times [0, k - 1]$ the following holds:

- if $x < k - 1$, $\tau(x, y) = \theta$, and $\tau(x + 1, y) = \theta'$, then $(\theta, \theta') \in \mathbb{H}$, and
- if $y < k - 1$, $\tau(x, y) = \theta$, and $\tau(x, y + 1) = \theta'$, then $(\theta, \theta') \in \mathbb{V}$.

Let $\text{Sol}_k(S)$ denote the set of all k -solutions for S . Let $w = \theta_0 \cdots \theta_{n-1} \in \Theta^n$ be a word and let $k \geq n$. With $\text{Sol}_k(S, w)$ we denote the set of all $\tau \in \text{Sol}_k(S)$ such that $\tau(x, 0) = \theta_x$ for all $x \in [0, n - 1]$. We denote with $|S|$ the size of a tiling system $S = (\Theta, \mathbb{H}, \mathbb{V})$, that is, the number of elements in $\Theta \uplus \mathbb{H} \uplus \mathbb{V}$. Now, we define tiling problems as follows:

ℓ -EXP-TILING PROBLEM

INPUT: A tiling system $S = (\Theta, \mathbb{H}, \mathbb{V})$, $\theta_0 \in \Theta$.

QUESTION: Does $\text{Sol}_{\text{Tower}(\ell, n)}(S, \theta_0) \neq \emptyset$ hold, where $n = |S|$?

The following result is folklore. It is based on the fact that, in general, from a non-deterministic $t(n)$ -time bounded Turing machine M and an input word w one can construct a (polynomially sized) tiling system S_M and an initial tile type t_0 which simulate M on input w (see e.g. [21; 22]) in the following sense: M accepts w if and only if $\text{Sol}_{t(n)}(S_M, \theta_0) \neq \emptyset$.

THEOREM 3.6 (FOLKLORE). *For each $\ell \geq 1$, the ℓ -EXP-TILING PROBLEM is hard for ℓ -NEXP under polynomial time many-one reductions.*

We can finally prove Proposition 3.5.

PROOF OF PROPOSITION 3.5. Due to technical reasons, we do the proof only for $\ell \geq 3$. The proof is via a polynomial time many-one reduction from the ℓ -EXP-TILING PROBLEM $\text{K}_{\text{id}}^2\text{-SAT}$ restricted to formulas of switching depth ℓ . Let $S_\ell = (\Theta, \mathbb{H}, \mathbb{V})$ be some tiling system of size $n = |S|$ and set $m = \text{Tower}(\ell - 1, n) - 1$.

We add to the set of propositions P_n from the previous section all tile types from Θ and two additional propositions x and y . To \bar{P}_n we add copies $\bar{\theta}$ ($\theta \in \Theta$), \bar{x} , and \bar{y} of these propositions. For $\beta \in \{0, 1\}$ and $z \in \{x, y\}$ we define $\Upsilon_{\ell-1, n}^{(\beta, z)}(j)$ (resp., $\bar{\Upsilon}_{\ell-1, n}^{(\beta, z)}(j)$) as the tree that is obtained from $\Upsilon_{\ell-1, n}^{(\beta)}(j)$ (resp., $\bar{\Upsilon}_{\ell-1, n}^{(\beta)}(j)$) by adding the label z (resp., \bar{z}) to the root.

We first define for all $X, Y \in [0, \text{Tower}(\ell, n) - 1]$ and all $\theta \in \Theta$ the *grid element tree* $\mathfrak{G}(X, Y, \theta)$ as follows: Define the subsets $I_x, I_y \subseteq [0, m]$ uniquely by

$$X = \sum_{i \in I_x} 2^i \quad \text{and} \quad Y = \sum_{i \in I_y} 2^i.$$

Then $\mathfrak{G}(X, Y, \theta)$ is obtained as follows:

- Take the disjoint union of a root node r and all trees from the set

$$U = \{\Upsilon_{\ell-2, n}^{(1, x)}(i) \mid i \in I_x\} \cup \{\Upsilon_{\ell-2, n}^{(0, x)}(i) \mid i \notin I_x\} \cup \\ \{\Upsilon_{\ell-2, n}^{(1, y)}(i) \mid i \in I_y\} \cup \{\Upsilon_{\ell-2, n}^{(0, y)}(i) \mid i \notin I_y\}.$$

- Add an edge from the root r to the root of each tree from U .
- Label the root r with θ .

The tree $\bar{\mathfrak{G}}(X, Y, \theta)$ is obtained from $\mathfrak{G}(X, Y, \theta)$ by replacing every action label a by \bar{a} and every proposition p by \bar{p} .

In order to enforce grid element trees, we need to slightly modify the formulas used in the proof of Theorem 3.4. For this purpose, it is useful to have for $z \in \{x, y\}$ the abbreviations $\diamond_z \psi = \diamond(z \wedge \psi)$, $\square_z \psi = \square(z \rightarrow \psi)$, $\bar{\diamond}_z \psi = \bar{\diamond}(\bar{z} \wedge \psi)$, and $\bar{\square}_z \psi = \bar{\square}(\bar{z} \rightarrow \psi)$. Then, for $z \in \{x, y\}$ we can define relativized formulas $\varphi_{\ell-1, n}^z$, $\text{eq}_{\ell-1, n}^z$, $\text{first}_{\ell-1, n}^z$, $\text{last}_{\ell-1, n}^z$, and $\text{succ}_{\ell-1, n}^z$ by replacing in the definitions of the formulas $\varphi_{\ell-1, n}$, $\text{eq}_{\ell-1, n}$, $\text{first}_{\ell-1, n}$,

$\text{last}_{\ell-1,n}$, and $\text{succ}_{\ell-1,n}$ every modality \diamond (resp., \square , $\overline{\diamond}$, $\overline{\square}$) by \diamond_z (resp., \square_z , $\overline{\diamond}_z$, $\overline{\square}_z$). All occurrences of $\varphi_{\ell-2,n}$, $\text{eq}_{\ell-2,n}$, $\text{first}_{\ell-2,n}$, $\text{last}_{\ell-2,n}$, and $\text{succ}_{\ell-2,n}$ are not changed, i.e., we do not replace modalities within these formulas. The following claim can be verified along the lines of the proof of Lemma 3.2.

Claim 1. Let $N_x, N_y, \overline{N}_x, \overline{N}_y \in [0, \text{Tower}(\ell, n) - 1]$ and $\theta, \theta' \in \Theta$ and let $\mathfrak{T} = \mathfrak{G}(N_x, N_y, \theta)$ and $\overline{\mathfrak{T}} = \overline{\mathfrak{G}}(\overline{N}_x, \overline{N}_y, \theta')$ be grid element trees. Then the following holds for all $z \in \{x, y\}$:

- (a) $\mathfrak{T} \times \overline{\mathfrak{T}} \models \text{eq}_{\ell-1,n}^z$ if and only if $N_z = \overline{N}_z$.
- (b) $\mathfrak{T} \times \overline{\mathfrak{T}} \models \text{first}_{\ell-1,n}^z$ if and only if $N_z = \overline{N}_z = 0$.
- (c) $\mathfrak{T} \times \overline{\mathfrak{T}} \models \text{last}_{\ell-1,n}^z$ if and only if $N_z = \overline{N}_z = \text{Tower}(\ell, n) - 1$.
- (d) $\mathfrak{T} \times \overline{\mathfrak{T}} \models \text{succ}_{\ell-1,n}^z$ if and only if $\overline{N}_z = N_z + 1$.

Using the relativized version of $\varphi_{\ell-1,n}$ we can enforce grid element trees. We define *gridel* as the conjunction of

$$\bigvee_{\theta \in \Theta} (\theta \wedge \overline{\theta} \wedge \bigwedge_{\kappa \in \Theta \setminus \{\theta\}} (-\kappa \wedge \overline{-\kappa})) \wedge \square \square ((x \oplus y) \wedge (\overline{x} \oplus \overline{y})) \wedge \varphi_{\ell-1,n}^x \wedge \varphi_{\ell-1,n}^y,$$

and

$$\bigwedge_{p \in P_n \setminus \Theta} (-p \wedge \overline{-p}) \wedge \bigwedge_{2 \leq i \leq \ell-1} \bigwedge_{p \in \Theta \cup \{x, y\}} \square^i \overline{\square}^i (-p \wedge \overline{-p})$$

where \oplus denotes “exclusive or” and \square^i denotes the sequence of i boxes \square . Intuitively, the first formula expresses that (i) the root is labeled with precisely one symbol $\theta \in \Theta$ and (ii) we can associate precisely two values with the grid element structure: the value enforced by $\varphi_{\ell-1,n}^x$ (in analogy to Theorem 3.4) and the value enforced by $\varphi_{\ell-1,n}^y$. The second formula is just an auxiliary formula restricting the newly introduced propositions appropriately, similar to formula (1) of Definition 3.3. The following claim makes this property of *gridel* explicit.

Claim 2. For all structures (\mathfrak{S}, s) and $(\overline{\mathfrak{S}}, \overline{s})$ we have that $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s, \overline{s} \rangle) \models \text{gridel}$ if and only if there are $X, Y \in [0, \text{Tower}(\ell, n) - 1]$ and $\theta \in \Theta$ such that (\mathfrak{S}, s) and $(\overline{\mathfrak{S}}, \overline{s})$ are bisimilar to grid element structures $\mathfrak{G}(X, Y, \theta)$ and $\overline{\mathfrak{G}}(X, Y, \theta)$, respectively.

Next, let $\tau : [0, \text{Tower}(\ell, n) - 1]^2 \rightarrow \Theta$ be a mapping. We define the *tiling tree* $\mathfrak{T}(\tau)$ as follows:

— Take the disjoint union of a root node r and all trees from the set

$$U = \{\mathfrak{G}(X, Y, \tau(X, Y)) \mid X, Y \in [0, \text{Tower}(\ell, n) - 1]\}.$$

— Add an edge from the root r to the root of each tree from U .

Intuitively, a tiling tree $\mathfrak{T}(\tau)$ represents the mapping τ as follows: for every X, Y in the domain of τ it has a successor that is a grid element tree encoding the triple $(X, Y, \tau(X, Y))$. Again, the copy $\overline{\mathfrak{T}}(\tau)$ is defined as usual. The following claim states the existence of a formula that enforces tiling trees.

Claim 3. There is a formula *tiling* of switching depth ℓ such that for all pointed structures (\mathfrak{S}, s) and $(\overline{\mathfrak{S}}, \overline{s})$ we have $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s, \overline{s} \rangle) \models \text{tiling}$ if and only if there is a $\text{Tower}(\ell, n)$ -solution τ of S_ℓ such that (\mathfrak{S}, s) is bisimilar to the tiling tree $\mathfrak{T}(\tau)$ and $(\overline{\mathfrak{S}}, \overline{s})$ is bisimilar to $\overline{\mathfrak{T}}(\tau)$.

PROOF OF CLAIM 3. We take for tiling the conjunction of the following formulas, where recall that $P_n \cup \overline{P}_n$ is the set of all atomic propositions:

- (1) $\bigwedge_{p \in P_n} (\neg p \wedge \neg \bar{p})$
- (2) $\diamond \bar{\diamond} (\text{gridel} \wedge \text{first}_{\ell-1,n}^x \wedge \text{first}_{\ell-1,n}^y)$
- (3) $\square \bar{\diamond} \text{gridel}$
- (4) $\bar{\square} \diamond \text{gridel}$
- (5) $\bar{\square} \square ((\text{eq}_{\ell-1,n}^x \wedge \text{eq}_{\ell-1,n}^y) \rightarrow \bigwedge_{\theta \in \Theta} (\theta \leftrightarrow \bar{\theta}))$
- (6) $\square (\bar{\square} \neg \text{last}_{\ell-1,n}^x \rightarrow \bar{\diamond} (\text{succ}_{\ell-1,n}^x \wedge \text{eq}_{\ell-1,n}^y \wedge \bigvee_{(\theta_1, \theta_2) \in \mathbb{H}} (\theta_1 \wedge \bar{\theta}_2)))$
- (7) $\square (\bar{\square} \neg \text{last}_{\ell-1,n}^y \rightarrow \bar{\diamond} (\text{succ}_{\ell-1,n}^y \wedge \text{eq}_{\ell-1,n}^x \wedge \bigvee_{(\theta_1, \theta_2) \in \mathbb{V}} (\theta_1 \wedge \bar{\theta}_2)))$

To show that the formula tiling satisfies the statements from the Claim we proceed similarly as in the proof of Theorem 3.4. Observe first that tiling has switching depth ℓ . For the “if”-direction of the statement assume a $\text{Tower}(\ell, n)$ -solution τ and structures (\mathfrak{S}, s) and $(\bar{\mathfrak{S}}, \bar{s})$ that are bisimilar to $\mathfrak{T}(\tau)$ and $\bar{\mathfrak{T}}(\tau)$, respectively. It is routine to verify that $(\mathfrak{S} \times \bar{\mathfrak{S}}, (s, \bar{s}))$ satisfies all the formulas (1)-(7) given above.

For the “only-if”-direction assume that $(\mathfrak{S} \times \bar{\mathfrak{S}}, (s, \bar{s})) \models$ tiling. Formulas (3) and (4) enforce that for all successors t and \bar{t} of s and \bar{s} , respectively, we have that (\mathfrak{S}, t) is bisimilar to some grid element tree and $(\bar{\mathfrak{S}}, \bar{t})$ is bisimilar to some grid element tree. Formula (2) enforces the existence of successors t and \bar{t} such that (\mathfrak{S}, t) is bisimilar to the grid element tree $\mathfrak{G}(0, 0, \theta)$ and $(\bar{\mathfrak{S}}, \bar{t})$ is bisimilar to the grid element tree $\bar{\mathfrak{G}}(0, 0, \theta)$ for some $\theta \in \Theta$. Formulas (6) and (7) inductively enforce the remaining grid element trees as successors: Assume a successor t of s such that (\mathfrak{S}, t) is bisimilar to the grid element tree $\mathfrak{G}(i, j, \theta)$. If $i < \text{Tower}(\ell, n) - 1$, then formula (6) enforces the existence of a successor \bar{t} of \bar{s} such that $(\bar{\mathfrak{S}}, \bar{t})$ is bisimilar to a grid element tree $\bar{\mathfrak{G}}(i+1, j, \theta')$ with $(\theta, \theta') \in \mathbb{H}$. By formula (4), there is some successor s' of s such that (\mathfrak{S}, s') is bisimilar to $\mathfrak{G}(i+1, j, \theta')$. Similarly, if $j < \text{Tower}(\ell, n) - 1$, then formula (7) together with formula (3) enforces the existence of a successor s' of s such that (\mathfrak{S}, s') is bisimilar to a grid element tree $\mathfrak{G}(i, j+1, \theta')$ with $(\theta, \theta') \in \mathbb{V}$. Thus, for each $i, j \in [0, \text{Tower}(\ell, n) - 1]$ there are states t, \bar{t} , and a tile type $\theta, \bar{\theta} \in \Theta$ such that

- t is a successor of s and \bar{t} is a successor of \bar{s} ;
- (\mathfrak{S}, t) is bisimilar to the grid element tree $\mathfrak{G}(i, j, \theta)$;
- $(\bar{\mathfrak{S}}, \bar{t})$ is bisimilar to the grid element tree $\bar{\mathfrak{G}}(i, j, \bar{\theta})$.

Note first that, if two successors t and \bar{t} of s and \bar{s} , respectively, are bisimilar to $\mathfrak{G}(i, j, \theta)$ and $\bar{\mathfrak{G}}(i, j, \bar{\theta})$, then we have $(\mathfrak{S} \times \bar{\mathfrak{S}}, (t, \bar{t})) \models \text{eq}_{\ell-1,n}^x \wedge \text{eq}_{\ell-1,n}^y$ by Claim 1. By formula (5), we obtain $\theta = \bar{\theta}$. Assume now that there are successors t, t' of s such that (\mathfrak{S}, t) is bisimilar to $\mathfrak{G}(i, j, \theta_1)$ and (\mathfrak{S}, t') is bisimilar to $\mathfrak{G}(i, j, \theta_2)$. By the above, there is a successor \bar{t} of \bar{s} such that $(\bar{\mathfrak{S}}, \bar{t})$ is bisimilar to some $\bar{\mathfrak{G}}(i, j, \theta')$. By Claim 1 and formula (5), we obtain $\theta' = \theta_1$. Analogously, we get $\theta' = \theta_2$ and thus $\theta_1 = \theta_2$. Hence, for every $i, j \in [0, \text{Tower}(\ell, n) - 1]$ there is a *unique* θ_{ij} such that for all successors t of s with (\mathfrak{S}, t) bisimilar to $\mathfrak{G}(i, j, \theta')$ we have $\theta' = \theta_{ij}$ (and analogously, for all successors \bar{t} of \bar{s}). Thus, the mapping τ defined by $\tau(i, j) = \theta_{ij}$ is well-defined. By what was said above, it is easy to verify that it is a $\text{Tower}(\ell, n)$ -solution for S_ℓ .

It remains to verify that (\mathfrak{S}, s) and $(\bar{\mathfrak{S}}, \bar{s})$ are bisimilar to the tiling trees $\mathfrak{T}(\tau)$ and $\bar{\mathfrak{T}}(\tau)$, respectively. In order to do this, it suffices to note that the labels in the points s and \bar{s} are the same as in a tiling tree, by formula (1). Together with the properties observed above, this yields the bisimulation and finishes the proof of Claim 3. \square

Finally, define a formula φ_{θ_0} expressing that the tile type at position $(0, 0)$ is precisely the input tile type θ_0 :

$$\varphi_{\theta_0} = \diamond \overline{\diamond} (\text{gridel} \wedge \text{first}_{\ell-1, n}^x \wedge \text{first}_{\ell-1, n}^y \wedge \theta_0).$$

It is not hard to see that $\text{Sol}_{\text{Tower}(\ell, n)}(S_\ell, \theta_0) \neq \emptyset$ if and only if $\text{tiling} \wedge \varphi_{\theta_0}$ is id-satisfiable. It remains to note that the size of the formula $\text{tiling} \wedge \varphi_{\theta_0}$ is polynomial in the size n of S . By Theorem 3.4(b), the size of $\varphi_{\ell-1, n}$ (and thus of $\varphi_{\ell-1, n}^z$) is polynomial in n . Hence, also the size of gridel is polynomial in n . Moreover, all the auxiliary formulas like $\text{first}_{\ell-1, n}^x$ are clearly polynomial in n . Overall, we get that $\text{tiling} \wedge \varphi_{\theta_0}$ is polynomially sized in $n = |S|$. This concludes the proof. \square

The following corollary is an immediate consequence of Proposition 2.1 and Proposition 3.5.

COROLLARY 3.7. *The following holds:*

- For each $\ell \geq 1$, $\text{K}^2\text{-SAT}$ restricted to formulas of switching depth ℓ is $\ell\text{-NEXP-hard}$ under polynomial time many-one reductions.
- In particular, $\text{K}^2\text{-SAT}$ is nonelementary.

4. HARDNESS RESULTS FOR $\text{K}4 \times \text{K}$ AND $\text{S}5_2 \times \text{K}$

In this section, we prove further nonelementary lower bound results for the satisfiability problem of two-dimensional modal logics on restricted classes of frames. We hereby close nonelementary complexity gaps that were stated as open problems in [8]. Although in [8] uninterpreted product models for these logics are of interest, we prove our lower bounds for the id-interpretation only: For each of the logics studied here, the id-interpretation case can be reduced in polynomial time to the uninterpreted case by Proposition 2.1.

We define the following logics:

- $\text{K}4 \times \text{K}$: Two-dimensional logic restricted to product models $\mathfrak{S}_1 \times \mathfrak{S}_2$ where $\mathfrak{F}(\mathfrak{S}_1)$ is a frame (W, \rightarrow_a) such that \rightarrow_a is *transitive*.
- $\text{S}5 \times \text{K}$: Two-dimensional logic restricted to product models $\mathfrak{S}_1 \times \mathfrak{S}_2$ where $\mathfrak{F}(\mathfrak{S}_1)$ is a frame (W, \equiv) with an *equivalence relation* \equiv .
- $\text{S}5_2 \times \text{K}$: Two-dimensional modal logic restricted to models $\mathfrak{S}_1 \times \mathfrak{S}_2$ where $\mathfrak{F}(\mathfrak{S}_1)$ is a frame (W, \equiv, \approx) with equivalence relations \equiv and \approx .

Note that the lower bounds from the last section already hold for formulas having only one action label a_i in every component $i \in \{1, 2\}$. Hence, throughout this section we fix $A = \{a_1, a_2\}$ and some countable set $P = P_1 \cup P_2$ of propositions. As in the previous section we will abbreviate \diamond_{a_1} with \diamond and \diamond_{a_2} with $\overline{\diamond}$.

Let us start with $\text{K}4 \times \text{K}$. We adapt in a straightforward way the reduction from K -satisfiability to $\text{K}4$ -satisfiability to the two-dimensional case. When following a transition in a $\text{K}4$ -frame one has no control over how far one is actually going due to transitivity of the frame. The idea for the reduction is to introduce additional propositions h_0, \dots, h_n and enforce *levels* in the models. Intuitively, h_i is true in w' precisely when w' is in level i seen from the world w where the formula is evaluated. Following a transition is then restricted to increase the level only by 1.

Let φ be a K^2 -formula with $\text{rank}_1(\varphi) = r$ and let h_0, \dots, h_r be fresh propositions. For every $0 \leq k \leq r$, we specify by structural induction a translation function t_k such that

t_k is defined for an input formula ψ whenever $\text{rank}_1(\psi) + k \leq r$. More precisely, we set

$$\begin{aligned} t_k(p) &\stackrel{\text{def}}{=} H_k \wedge p \\ t_k(\neg\psi) &\stackrel{\text{def}}{=} H_k \wedge \neg t_k(\psi) \\ t_k(\psi_1 \wedge \psi_2) &\stackrel{\text{def}}{=} t_k(\psi_1) \wedge t_k(\psi_2) \\ t_k(\overline{\diamond}\psi) &\stackrel{\text{def}}{=} \overline{\diamond} t_k(\psi) \\ t_k(\diamond\psi) &\stackrel{\text{def}}{=} H_k \wedge \diamond(H_{k+1} \wedge t_{k+1}(\psi)), \end{aligned}$$

where $H_k \stackrel{\text{def}}{=} h_k \wedge \bigwedge_{i \neq k} \neg h_i$ and $k < r$ in the definition of $t_k(\diamond\psi)$. We show that the translation is satisfiability preserving. More precisely, we prove the following lemma.

LEMMA 4.1. *For every $\mathbf{K}^2(A, P)$ -formula φ we have: φ is id-satisfiable in \mathbf{K}^2 if and only if $t_0(\varphi)$ is id-satisfiable in $\mathbf{K4} \times \mathbf{K}$.*

PROOF. We assume that φ is defined over $P = P_1 \cup P_2$ for disjoint P_1 and P_2 . Moreover set $r \stackrel{\text{def}}{=} \text{rank}_1(\varphi)$. As in Section 3 we will write $\mathfrak{S}_1 \times \mathfrak{S}_2$ for $\prod_{i \in [1, 2]}^{\text{id}} \mathfrak{S}_i$.

Assume first that φ is id-satisfiable in $\mathbf{K} \times \mathbf{K}$. Thus, there are structures

$$\mathfrak{S}_i = (W_i, \longrightarrow_i, \{W_{i,p} \mid p \in P_i\})$$

($i \in \{1, 2\}$) and $\bar{w} = \langle w_1, w_2 \rangle \in W_1 \times W_2$ such that $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{w}) \models \varphi$. Without loss of generality we assume that \mathfrak{S}_1 is a tree with root w_1 . Define

$$\mathfrak{S}'_1 = (W_1, \longrightarrow_1^+, \{W'_{1,p} \mid p \in P_1 \cup \{h_0, \dots, h_r\}\}),$$

where

- \longrightarrow_1^+ is the transitive closure of \longrightarrow_1 ,
- $W'_{1,p} \stackrel{\text{def}}{=} W_{1,p}$ for all $p \in P_1$, and
- $W'_{1,h_i} \stackrel{\text{def}}{=} V_i$, where V_i is defined to be the set of worlds w' such that the (unique) path in \mathfrak{S}_1 from w_1 to w' has length i , i.e., consists of i transitions.

We prove by induction on the structure of φ that for each subformula ψ it holds: for all $i \in [0, r]$ we have

$$\text{rank}_1(\psi) \leq i \quad \Rightarrow \quad ((\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models \psi \quad \Leftrightarrow \quad (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models t_{r-i}(\psi))$$

for each $\bar{x} \in V_{r-i} \times W_2$.

For the induction base, assume $\psi = p$ for some atomic proposition $p \in P_1 \cup P_2$, i arbitrary in $[0, r]$, and fix an arbitrary $\bar{x} = \langle x_1, x_2 \rangle \in V_{r-i} \times W_2$. By definition of \mathfrak{S}'_1 , we have $\bar{x} \in W'_{1,h_{r-i}} \times W_2$ and $\bar{x} \notin W'_{1,h_j} \times W_2$ for $j \neq r-i$; hence, $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models H_r$. Finally, the following equivalences hold, where we assume that $p \in P_j$ ($j \in \{1, 2\}$): $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models \psi$ if and only if $x_j \in W_{j,p}$ if and only if $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models H_r \wedge p = t_r(\psi)$.

For the induction step, assume ψ is not atomic and $i \in [0, r]$ such that $i \geq \text{rank}_1(\psi)$, and let us fix some $\bar{x} \in V_{r-i} \times W_2$. Note that $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models H_{r-i}$ by definition of \mathfrak{S}'_1 . We make a case distinction on the structure of ψ . For the cases $\neg\chi$, $\chi_1 \wedge \chi_2$, and $\overline{\diamond}\chi$ the equivalence follows straightforwardly from the induction hypothesis since $\text{rank}_1(\psi) = \text{rank}_1(\chi)$ and $\text{rank}_1(\chi_i) \leq \text{rank}_1(\psi)$ for $i \in \{1, 2\}$.

It remains to consider the case $\psi = \diamond\chi$. Then $\text{rank}_1(\chi) = \text{rank}_1(\psi) - 1 \leq i - 1$ and we have

$$\begin{aligned}
(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models \psi &\Leftrightarrow \exists \bar{y} \in V_{r-(i-1)} \times W_2 : \bar{x} \longrightarrow_1 \bar{y} \text{ and } (\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{y}) \models \chi \\
&\stackrel{\text{IH}}{\Leftrightarrow} \exists \bar{y} \in V_{r-i+1} \times W_2 : \bar{x} \longrightarrow_1 \bar{y} \text{ and } (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{y}) \models t_{r-i+1}(\chi) \\
&\stackrel{\bar{y} \in V_{r-i+1} \times W_2}{\Leftrightarrow} \exists \bar{y} \in V_{r-i+1} \times W_2 : \bar{x} \longrightarrow_1^+ \bar{y} \text{ and} \\
&\quad (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{y}) \models H_{r-i+1} \wedge t_{r-i+1}(\chi) \\
&\Leftrightarrow (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models \diamond(H_{r-i+1} \wedge t_{r-i+1}(\chi)) \\
&\stackrel{\bar{x} \in V_{r-i} \times W_2}{\Leftrightarrow} (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models H_{r-i} \wedge \diamond(H_{r-i+1} \wedge t_{r-i+1}(\chi)) \\
&\Leftrightarrow (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models t_{r-i}(\psi).
\end{aligned}$$

Since $\text{rank}_1(\varphi) = r$, $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{w}) \models \varphi$, and $\bar{w} \in V_0 \times W_2$ we get $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{w}) \models t_0(\varphi)$. Hence $t_0(\varphi)$ is id-satisfiable in $\mathbf{K4} \times \mathbf{K}$.

For the other direction assume that $t_0(\varphi)$ is id-satisfiable in $\mathbf{K4} \times \mathbf{K}$. Thus, there are a transitive structure $\mathfrak{S}_1 = (W_1, \longrightarrow_1, \{W_{1,p} \mid p \in P_1 \cup \{h_0, \dots, h_r\}\})$ and a structure $\mathfrak{S}_2 = (W_2, \longrightarrow_2, \{W_{2,p} \mid p \in P_2\})$ and $\bar{w} = \langle w_1, w_2 \rangle \in W_1 \times W_2$ such that $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{w}) \models t_0(\varphi)$. For each $0 \leq i \leq r$ we set

$$T_i \stackrel{\text{def}}{=} W_{1,h_i} \setminus \left(\bigcup_{j \neq i} W_{1,h_j} \right),$$

which corresponds to the formulas H_i in \mathfrak{S}_1 .

Now, define the structure $\mathfrak{S}'_1 = (W'_1, \longrightarrow'_1, \{W'_{1,p} \mid p \in P_1\})$ by taking

- $W'_1 = \bigcup_{0 \leq i \leq r} T_i$,
- $\longrightarrow'_1 = \longrightarrow_1 \cap (\bigcup_{0 \leq i < r} T_i \times T_{i+1})$, and
- $W'_{1,p} = W_{1,p} \cap W'_1$ for all $p \in P_1$.

We prove by structural induction that for each subformula ψ of φ we have: for all $i \in [0, r]$ it holds

$$\text{rank}_1(\psi) \leq i \quad \Rightarrow \quad ((\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models t_{r-i}(\psi)) \quad \Leftrightarrow \quad (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models \psi)$$

for each $\bar{x} \in T_{r-i} \times W_2$.

For the induction base assume $\psi = p$ for some atomic proposition $p \in P_1 \cup P_2$ and $i \geq \text{rank}_1(\psi) = 0$. By definition, we have $t_{r-i}(\psi) = H_{r-i} \wedge p$. Take an arbitrary $\bar{x} \in T_r \times W_2$. By definition of T_{r-i} , we have $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models h_{r-i}$ and $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \not\models h_j$ for each $j \in [0, r] \setminus \{r-i\}$. Thus, $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models H_{r-i}$. Moreover we have $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models p$ if and only if $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models p$ by definition of $W'_{1,p}$. Thus,

$$(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models t_{r-i}(\psi) \quad \Leftrightarrow \quad (\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models H_{r-i} \wedge \psi \quad \Leftrightarrow \quad (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models \psi.$$

For the induction step assume that ψ is not atomic, let $i \in [0, r]$ be such that $\text{rank}_1(\psi) \leq i$ and fix an arbitrary $\bar{x} = \langle x_1, x_2 \rangle \in T_{r-i} \times W_2$. Note that we have $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models H_{r-i}$, by definition of T_{r-i} . We make a case distinction on the structure of ψ . For the cases $\neg\chi$, $\chi_1 \wedge \chi_2$, and $\diamond\chi$ the equivalence follows directly from the induction hypothesis.

For the remaining case $\psi = \diamond\chi$ we have $\text{rank}_1(\chi) = \text{rank}_1(\psi) - 1 \leq i - 1$ and

$$\begin{array}{lcl}
(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models t_{r-i}(\psi) & \Leftrightarrow & (\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models H_{r-i} \wedge \diamond(H_{r-i+1} \wedge t_{r-i+1}(\chi)) \\
& \stackrel{x_1 \in T_{r-i}}{\Leftrightarrow} & (\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{x}) \models \diamond(H_{r-i+1} \wedge t_{r-i+1}(\chi)) \\
& \stackrel{\text{Def. } T_{r-i+1}}{\Leftrightarrow} & \exists \bar{y} \in T_{r-i+1} \times W_2 : \bar{x} \longrightarrow_1 \bar{y} \text{ and} \\
& & (\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{y}) \models H_{r-i+1} \wedge t_{r-i+1}(\chi) \\
\bar{y} \in T_{r-i+1} \times W_2 & \Leftrightarrow & \exists \bar{y} \in T_{r-i+1} \times W_2 : \bar{x} \longrightarrow_1 \bar{y} \text{ and} \\
& & (\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{y}) \models t_{r-i+1}(\chi) \\
& \stackrel{\text{Def. } \longrightarrow'_1}{\Leftrightarrow} & \exists \bar{y} \in T_{r-i+1} \times W_2 : \bar{x} \longrightarrow'_1 \bar{y} \text{ and} \\
& & (\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{y}) \models t_{r-i+1}(\chi) \\
\text{rank}_1(\chi) \leq i-1, \text{hyp} & \Leftrightarrow & \exists \bar{y} \in T_{r-i+1} \times W_2 : \bar{x} \longrightarrow'_1 \bar{y} \text{ and } (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{y}) \models \chi \\
& \Leftrightarrow & (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models \diamond\chi \\
& \Leftrightarrow & (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{x}) \models \psi.
\end{array}$$

By assumption we have $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{w}) \models t_0(\varphi)$, $\bar{w} \in T_0 \times W_2$, and $\text{rank}_1(\varphi) = r$. Thus, the above equivalence implies $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{w}) \models \varphi$ and thus, φ is id-satisfiable in $\mathbf{K} \times \mathbf{K}$. \square

Lemma 4.1 provides a reduction of \mathbf{K}_{id}^2 -SAT to id-satisfiability in $\mathbf{K4} \times \mathbf{K}$. Finally, Proposition 2.1 together with Proposition 3.5 yields the following result.

THEOREM 4.2. *Satisfiability in $\mathbf{K4} \times \mathbf{K}$ is nonelementary.*

Next, we study combinations of \mathbf{K} with $\mathbf{S5}$ and $\mathbf{S5}_2$. It is well-known that the complexity for checking satisfiability jumps from NP for $\mathbf{S5}$ to PSPACE for $\mathbf{S5}_2$. We will show that also the complexity for deciding satisfiability in the product logics $\mathbf{S5} \times \mathbf{K}$ and $\mathbf{S5}_2 \times \mathbf{K}$, respectively, differs. In particular, we will again reduce from \mathbf{K}_{id}^2 -SAT in order to show a nonelementary lower bound for the latter logic, which is in sharp contrast to the following result by Marx [23].

THEOREM 4.3 ([23]). *Satisfiability in $\mathbf{S5} \times \mathbf{K}$ is NEXP-complete.*

PSPACE-hardness for satisfiability in $\mathbf{S5}_2$ can be shown via a straightforward reduction from \mathbf{K} [3]. We adapt this reduction to the two-dimensional case by defining a translation \dagger by

$$\begin{aligned}
q^\dagger &\stackrel{\text{def}}{=} p^* \wedge q \\
(\varphi_1 \wedge \varphi_2)^\dagger &\stackrel{\text{def}}{=} p^* \wedge \varphi_1^\dagger \wedge \varphi_2^\dagger \\
(\neg\varphi)^\dagger &\stackrel{\text{def}}{=} p^* \wedge \neg(\varphi^\dagger) \\
(\overline{\diamond}\varphi)^\dagger &\stackrel{\text{def}}{=} p^* \wedge \overline{\diamond}\varphi^\dagger \\
(\diamond\varphi)^\dagger &\stackrel{\text{def}}{=} p^* \wedge \diamond_{\equiv}(\neg p^* \wedge \diamond_{\approx}(p^* \wedge \varphi^\dagger))
\end{aligned}$$

where \diamond_{\equiv} and \diamond_{\approx} refer to the two modalities in $\mathbf{S5}_2$ and p^* is a fresh propositional variable in the signature of the first component. Intuitively, *one* transition in \mathbf{K} is simulated by *two* transitions in $\mathbf{S5}_2$. This is possible since the composition of two equivalence relations is neither symmetric nor transitive in general and using the fresh variable p^* we can enforce a non-trivial transition, i.e., no loops. It can be proven along the lines of the proof in [3] that \dagger preserves id-satisfiability.

LEMMA 4.4. *For every $\mathbf{K}^2(A, P)$ -formula φ we have: φ is id-satisfiable in \mathbf{K}^2 if and only if φ^\dagger is id-satisfiable in $\mathbf{S5}_2 \times \mathbf{K}$.*

PROOF. We assume that φ is defined over $P = P_1 \cup P_2$ for disjoint P_1 and P_2 with $p^* \notin P_1 \cup P_2$. Again, we write $\mathfrak{S}_1 \times \mathfrak{S}_2$ for $\prod_{i \in [1,2]}^{\text{id}} \mathfrak{S}_i$.

Assume first that φ is id-satisfiable in $\mathbf{K} \times \mathbf{K}$. Thus, there are $\mathfrak{S}_1 = (W_1, \xrightarrow{a}, \{W_{1,p} \mid p \in P_1\})$, $\mathfrak{S}_2 = (W_2, \xrightarrow{b}, \{W_{2,p} \mid p \in P_2\})$, and $\bar{s} \in W_1 \times W_2$ such that $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{s}) \models \varphi$. Define an $\mathbf{S5}_2$ -structure $\mathfrak{S}'_1 = (W'_1, \equiv, \approx, \{W'_{1,p} \mid p \in P_1 \cup \{p^*\}\})$ as follows:

- $W'_1 \stackrel{\text{def}}{=} W_1 \uplus \xrightarrow{a}$,
- \equiv is the reflexive, transitive, and symmetric closure of $\{(w, (w, w')) \mid w \xrightarrow{a} w'\}$,
- \approx is the reflexive, transitive, and symmetric closure of $\{((w, w'), w') \mid w \xrightarrow{a} w'\}$,
- $W'_{1,p} \stackrel{\text{def}}{=} W_{1,p}$ for $p \in P_1$,
- $W'_{1,p^*} \stackrel{\text{def}}{=} W_1$.

Now, one can prove by induction on the structure of a formula ψ that for every world $\bar{w} \in W_1 \times W_2$ we have:

$$(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{w}) \models \psi \quad \Leftrightarrow \quad (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{w}) \models \psi^\dagger.$$

For the induction base, i.e., when ψ is a propositional variable, the statement is immediately true, by definition of the structure \mathfrak{S}'_1 . For the cases $\neg\chi$, $\chi_1 \wedge \chi_2$ and $\diamond\chi$, the statement follows directly from the induction hypothesis.

So assume ψ is of the form $\diamond\chi$. Suppose first that $(\mathfrak{S}_1 \times \mathfrak{S}_2, \langle w_1, w_2 \rangle) \models \diamond\chi$. Thus, there is some world w'_1 such that $w_1 \xrightarrow{a} w'_1$ and $(\mathfrak{S}_1 \times \mathfrak{S}_2, \langle w'_1, w_2 \rangle) \models \chi$. By induction hypothesis, we have $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \langle w'_1, w_2 \rangle) \models \chi^\dagger$. By definition of \mathfrak{S}'_1 , we have $w_1 \equiv (w_1, w'_1)$, $(w_1, w'_1) \approx w'_1$, $w_1, w'_1 \in W'_{1,p^*}$, and $(w_1, w'_1) \notin W'_{1,p^*}$. Obviously, this yields $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \langle w_1, w_2 \rangle) \models p^* \wedge \diamond \equiv (\neg p^* \wedge \diamond \approx (p^* \wedge \chi^\dagger))$. For the other direction suppose $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \langle w_1, w_2 \rangle) \models p^* \wedge \diamond \equiv (\neg p^* \wedge \diamond \approx (p^* \wedge \chi^\dagger))$. Thus, there are worlds $v, w'_1 \in W'_1$ with $w_1 \equiv v$, $v \approx w'_1$, $w_1, w'_1 \in W'_{1,p^*}$, and $v \notin W'_{1,p^*}$ such that $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \langle w'_1, w_2 \rangle) \models \chi^\dagger$. By definition of \mathfrak{S}'_1 , we know that $w'_1 \in W_1$, $v = (w_1, w'_1)$, and $w_1 \xrightarrow{a} w'_1$. As $w'_1 \in W_1$, the induction hypothesis implies $(\mathfrak{S}_1 \times \mathfrak{S}_2, \langle w'_1, w_2 \rangle) \models \chi$. Hence, $(\mathfrak{S}_1 \times \mathfrak{S}_2, \langle w_1, w_2 \rangle) \models \diamond\chi$.

In particular, we obtain $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{s}) \models \varphi^\dagger$, thus, φ^\dagger is id-satisfiable in $\mathbf{S5}_2 \times \mathbf{K}$.

Assume now that φ^\dagger is id-satisfiable in $\mathbf{S5}_2 \times \mathbf{K}$. Hence, there is an $\mathbf{S5}_2$ -structure

$$\mathfrak{S}_1 = (W_1, \equiv, \approx, \{W_{1,p} \mid p \in P_1 \cup \{p^*\}\}),$$

a structure $\mathfrak{S}_2 = (W_2, \xrightarrow{b}, \{W_{2,p} \mid p \in P_2\})$, and $\bar{s} \in W_1 \times W_2$ such that $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{s}) \models \varphi^\dagger$. Define a structure $\mathfrak{S}'_1 = (W'_1, \xrightarrow{a}, \{W'_{1,p} \mid p \in P_1\})$ as follows:

- $W'_1 \stackrel{\text{def}}{=} W_{1,p^*}$
- $\xrightarrow{a} \stackrel{\text{def}}{=} \{(u, v) \mid \exists w \in W_1 \setminus W_{1,p^*} : u \equiv w \approx v\}$
- $W'_{1,p} \stackrel{\text{def}}{=} W_{1,p} \cap W_{1,p^*}$ for all $p \in P_1$

One can prove by induction on the structure of a formula ψ that for every world $\bar{w} \in W'_1 \times W_2$ we have:

$$(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{w}) \models \psi^\dagger \quad \Leftrightarrow \quad (\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{w}) \models \psi$$

Again, the case when ψ is a propositional variable is immediately clear from the definition of \mathfrak{S}'_1 . Also the cases $\neg\chi$, $\chi_1 \wedge \chi_2$, and $\diamond\chi$ are direct consequences from the induction hypothesis.

For the case $\psi = \diamond\chi$ assume first that $(\mathfrak{S}_1 \times \mathfrak{S}_2, \langle w_1, w_2 \rangle) \models (\diamond\chi)^\dagger$. By the semantics, there is some world v with $w_1 \equiv v$ and $v \approx w_2$ such that $w_1, w_2 \in W_{1,p^*}$, $v \notin W_{1,p^*}$, and $(\mathfrak{S}_1 \times \mathfrak{S}_2, \langle w'_1, w_2 \rangle) \models \chi^\dagger$. By induction, we have that $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \langle w'_1, w_2 \rangle) \models \chi$. Moreover, the definition of \mathfrak{S}'_1 yields $w_1 \xrightarrow{a} w'_1$. By the semantics, we get $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \langle w_1, w_2 \rangle) \models \diamond\chi$. For the other direction assume $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \langle w_1, w_2 \rangle) \models \diamond\chi$. Hence, there is some world w'_1 such that $w_1 \xrightarrow{a} w'_1$ and $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \langle w'_1, w_2 \rangle) \models \chi$. By induction, $(\mathfrak{S}_1 \times \mathfrak{S}_2, \langle w'_1, w_2 \rangle) \models \chi^\dagger$. By definition of \xrightarrow{a} , there is some $v \in W_1 \setminus W_{1,p^*}$ such that $w_1 \equiv v \approx w'_1$. By definition of \mathfrak{S}'_1 , we have $w_1, w'_1 \in W_{1,p^*}$. Thus, the semantics yields $(\mathfrak{S}_1 \times \mathfrak{S}_2, \langle w_1, w_2 \rangle) \models (\diamond\chi)^\dagger$.

Observe now that $(\mathfrak{S}_1 \times \mathfrak{S}_2, \bar{s}) \models \varphi^\dagger$ implies $\bar{s} \in W'_1 \times W_2$ by the definition of † . Therefore, we get $(\mathfrak{S}'_1 \times \mathfrak{S}_2, \bar{s}) \models \varphi$ and φ is id-satisfiable in $\mathbf{K} \times \mathbf{K}$. \square

The following theorem is an immediate consequence of Lemma 4.4, Proposition 3.5, and Proposition 2.1.

THEOREM 4.5. *Satisfiability in $\mathbf{S5}_2 \times \mathbf{K}$ is nonelementary.*

5. FEFERMAN-VAUGHT DECOMPOSITIONS FOR PRODUCTS

The Feferman-Vaught decomposition theorem for many-dimensional modal logic \mathbf{K}^d can be formulated as follows, and was proven in [13]. Recall the notion of an interpretation σ from Section 2.3.

THEOREM 5.1 ([13]). *From an interpretation σ and a $\mathbf{K}^d(\mathbf{A}, \mathbf{P})$ -formula φ with $\mathbf{A} = \bigsqcup_{i \in [1, d]} \mathbf{A}_i$, $\mathbf{P} = \bigsqcup_{i \in [1, d]} \mathbf{P}_i$, one can compute a tuple $(\Psi_1, \dots, \Psi_d, \beta)$ with $\Psi_i = \{\psi_i^j \mid j \in J_i\}$ a finite set of multimodal formulas over $(\mathbf{A}_i, \mathbf{P}_i)$ and β a positive boolean formula with variables from $X = \{x_i^j \mid i \in [1, d], j \in J_i\}$ such that for every $(\mathbf{A}_i, \mathbf{P}_i)$ -structure \mathfrak{S}_i and every world w_i of \mathfrak{S}_i ($i \in [1, d]$):*

$$\left(\prod_{i \in [1, d]}^{\sigma} \mathfrak{S}_i, \langle w_1, \dots, w_n \rangle \right) \models \varphi \quad \Leftrightarrow \quad \mu \models \beta$$

Here, $\mu : X \rightarrow \{0, 1\}$ is defined by $\mu(x_i^j) = 1$ if and only if $(\mathfrak{S}_i, w_i) \models \psi_i^j$.

We call $\mathbf{D} \stackrel{\text{def}}{=}} (\Psi_1, \dots, \Psi_d, \beta)$ the *decomposition* of φ and define $|\mathbf{D}| \stackrel{\text{def}}{=} |\beta| + \sum_{i,j} |\psi_i^j|$ to be its *size*.

Note that Theorem 5.1 only holds in the presence of an interpretation σ for the atomic propositions since interpretations establish the connection between the product and component structures. We also mention that Theorem 5.1 has been proven in [13] for much more elaborated notions of interpretations. However, note that not every logic admits decomposition: An example for this is the temporal logic CTL. More precisely, it has been shown in [13, Theorem 11] that decomposition with respect to asynchronous products and a particular interpretation for the atomic proposition (in the sense of Section 2.3) fails for every logic that can express the property EG_p meaning “there is a maximal path (a path is maximal if it is either infinite or ends in a dead-end) on which every world satisfies p ”.

An analogous theorem can be stated for first-order sentences, see [11] for a survey. We assume standard definitions concerning first-order logic. We will consider only relational signatures τ . For a finite set \mathbf{A} of action labels and a finite set of propositions \mathbf{P} we identify the pair (\mathbf{A}, \mathbf{P}) with the signature, where every $a \in \mathbf{A}$ has arity 2 and every $p \in \mathbf{P}$ has arity 1. This allows to consider Kripke structures as ordinary relational structures. In the following we consider decomposition theorems for finite variable fragments FO^k of first-order logic. A formula φ is in FO^k if at most k different variables

occur in φ . Note that a formula, in which every subformula has at most k free variables is equivalent to an FO^k -formula.

THEOREM 5.2 ([10]). *From an interpretation σ and an FO^k -sentence φ over the signature (A, P) with $A = \uplus_{i \in [1, d]} A_i$, $P = \uplus_{i \in [1, d]} P_i$, one can compute a tuple $(\Psi_1, \dots, \Psi_d, \beta)$ with $\Psi_i = \{\psi_i^j \mid j \in J_i\}$ a finite set of FO^k -sentences over the signature (A_i, P_i) and β a positive boolean formula with variables from $X = \{x_i^j \mid i \in [1, d], j \in J_i\}$ such that for every (A_i, P_i) -structure \mathfrak{S}_i ($i \in [1, d]$):*

$$\prod_{i \in [1, d]}^{\sigma} \mathfrak{S}_i \models \varphi \quad \Leftrightarrow \quad \mu \models \beta.$$

Here, $\mu : X \rightarrow \{0, 1\}$ is defined by $\mu(x_i^j) = 1$ if and only if $\mathfrak{S}_i \models \psi_i^j$.

Note that the proofs of both Theorem 5.1 and Theorem 5.2 yield decompositions of nonelementary size. In this section we provide matching lower bounds for Feferman-Vaught decompositions for \mathbf{K}^d and FO^k for $k \geq 2$. Having enforced nonelementarily branching trees with small formulas (Theorem 3.4) allows us to prove a nonelementary lower bound for the sizes of Feferman-Vaught decompositions for 2-dimensional modal logic. Without making this explicit in the statement, our lower bound is more general than the nonelementary lower bound for 2-dimensional modal logic from [17] in the following sense. We provide a family of small formulas which are “inherently hard to decompose”: When assuming, by contradiction, the existence of small decompositions for our formulas, *any* model for them can be used to deduce the desired contradiction, whereas in [17] *appropriately chosen models* had to be defined for this. Our proof strategy is similar to the proof of Theorem 5.1 in [12].

THEOREM 5.3. *Feferman-Vaught decompositions for many-dimensional modal logic w.r.t. asynchronous product are inherently nonelementary. More precisely, for every elementary function $f(n)$ there exists $\ell \geq 1$ such that the \mathbf{K}^2 -formula $\varphi_{\ell, 2}$ from Definition 3.3 has no decomposition D_ℓ in the sense of Theorem 5.1 with $|D_\ell| \leq f(|\varphi_{\ell, 2}|)$. The same lower bound holds when relativized to product structures $\mathfrak{T} \times \mathfrak{T}'$, where $\mathfrak{F}(\mathfrak{T})$ and $\mathfrak{F}(\mathfrak{T}')$ are finite trees.*

PROOF. Assume by contradiction that there were an elementary function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that for each $\ell \geq 1$ there is a decomposition $D_\ell = (\Psi^{(\ell)}, \overline{\Psi}^{(\ell)}, \beta_\ell)$ of $\varphi_{\ell, 2}$ in the sense of Theorem 5.1 with $|D_\ell| \leq f(|\varphi_{\ell, 2}|)$. In particular, $|\beta_\ell| \leq f(|\varphi_{\ell, 2}|)$. Since $|\varphi_{\ell, 2}| \leq \exp(\ell)$ by Theorem 3.4(b), there exists an elementary function g such that $|\beta_\ell| \leq g(\ell)$ for all $\ell \geq 0$. Thus, there exists an $h_0 \geq 0$ with $2^{g(h-1)} < \text{Tower}(h, 2)$ for all $h \geq h_0$; let us fix such an h_0 .

By Theorem 3.4(a), $\varphi_{h_0, 2}$ is id-satisfiable. Assume that $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s, \overline{s} \rangle) \models \varphi_{h_0, 2}$ for some pointed structure (\mathfrak{S}, s) over $(\{a\}, P_n)$ and some pointed structure $(\overline{\mathfrak{S}}, \overline{s})$ over $(\{\overline{a}\}, \overline{P}_n)$. By Theorem 3.4(a) there exists some $k \in [0, \text{Tower}(h_0 + 1, n)]$ such that (\mathfrak{S}, s) is bisimilar to some extension of $\Upsilon_{h_0, n}^{(0)}(k)$ and $(\overline{\mathfrak{S}}, \overline{s})$ is bisimilar to some extension of $\overline{\Upsilon}_{h_0, n}^{(0)}(k)$.

By Definition 3.1, for each $i \in [0, \text{Tower}(h_0, 2) - 1]$ there exist successors s_i of s and \overline{s}_i of \overline{s} such that (\mathfrak{S}, s_i) is bisimilar to some extension of $\Upsilon_{h_0-1, 2}^{(0)}(i)$ and $(\overline{\mathfrak{S}}, \overline{s}_i)$ is bisimilar to some extension of $\overline{\Upsilon}_{h_0-1, 2}^{(0)}(i)$. Also note that

$$(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s_i, \overline{s}_j \rangle) \models \varphi_{h_0-1, 2} \Leftrightarrow i = j \tag{2}$$

for all $i, j \in [0, \text{Tower}(h_0, 2) - 1]$. Consider our decomposition $D_{h_0-1} = (\Psi^{(h_0-1)}, \overline{\Psi}^{(h_0-1)}, \beta_{h_0-1})$ of $\varphi_{h_0-1,2}$ where $\Psi^{(h_0-1)} = \{\psi_j \mid j \in J\}$ and $\overline{\Psi}^{(h_0-1)} = \{\overline{\psi}_j \mid j \in \overline{J}\}$ for some indexed sets J, \overline{J} , and β_{h_0-1} is a positive boolean formula with variables from $X = \{x_j \mid j \in J\} \cup \{\overline{x}_j \mid j \in \overline{J}\}$. By assumption, we have $|\beta_{h_0-1}| \leq g(h_0 - 1)$ and hence, $|X| \leq g(h_0 - 1)$.

For each $r \in [0, \text{Tower}(h_0, 2) - 1]$ we define a truth assignment $\mu_r : X \rightarrow \{0, 1\}$ as follows:

$$\begin{aligned} \mu_r(x_j) = 1 &\Leftrightarrow (\mathfrak{S}, s_r) \models \psi_j \\ \mu_r(\overline{x}_j) = 1 &\Leftrightarrow (\overline{\mathfrak{S}}, \overline{s}_r) \models \overline{\psi}_j \end{aligned}$$

Since for β_{h_0-1} there are $2^{|X|} \leq 2^{g(h_0-1)} < \text{Tower}(h_0, 2)$ many truth assignments, by the pigeonhole principle there exist $0 \leq a < b < \text{Tower}(h_0, 2)$ with $\mu_a = \mu_b$. In other words, $(\mathfrak{S}, s_a) \models \psi_j$ if and only if $(\mathfrak{S}, s_b) \models \psi_j$ and $(\overline{\mathfrak{S}}, \overline{s}_a) \models \overline{\psi}_j$ if and only if $(\overline{\mathfrak{S}}, \overline{s}_b) \models \overline{\psi}_j$. By the definition of a Feferman-Vaught decomposition and the fact that $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s_a, \overline{s}_a \rangle) \models \varphi_{h_0-1,2}$, we obtain $(\mathfrak{S} \times \overline{\mathfrak{S}}, \langle s_a, \overline{s}_b \rangle) \models \varphi_{h_0-1,2}$. But this contradicts (2).

Note that the lower bound also holds when restricting models to products of finite trees, since every pointed structure (\mathfrak{S}, s) (resp., $(\overline{\mathfrak{S}}, \overline{s})$) that is bisimilar to an extension of $\Upsilon_{\ell,n}^{(0)}(j)$ (resp., to an extension of $\overline{\Upsilon}_{\ell,n}^{(0)}(j)$) is bisimilar to a finite tree. \square

Note that the lower bound from Theorem 5.3 would even hold if we defined the size of a decomposition $(\Psi_1, \dots, \Psi_d, \beta)$ as the size of the boolean formula β only (and not accounting for the sizes of the Ψ_i); the same proof works for this variant. In contrast to [17] the proof of Theorem 5.3 allows to derive nonelementary lower bounds on decompositions for any decomposable logic (in the sense of Theorem 5.1) that is at least as expressive as modal logic and only elementarily less succinct than modal logic.

COROLLARY 5.4. *Every logic that is at least as expressive as and at most elementarily less succinct than modal logic does not have elementary sized Feferman-Vaught decompositions with respect to asynchronous product.*

PROOF. We exemplarily provide the proof for FO^2 sentences. The proof for any other logic that satisfies the properties from Corollary 5.4 works analogously. It is well known that a modal logic formula φ can be translated (in polynomial time) into an equivalent FO^2 -formula $\widehat{\varphi}(x)$ with one free variable, see for instance [1, Section 2.4]. The family of FO^2 -sentences that witnesses that there are no elementarily-sized decompositions is simply $\{\exists x : \widehat{\varphi}_{\ell,2}(x) \mid \ell \geq 1\}$. \square

6. FEFERMAN-VAUGHT DECOMPOSITIONS FOR SUM

So far, we only considered Feferman-Vaught decompositions for asynchronous products. Another important and natural operation on structures is the disjoint sum. Let us fix a relational signature τ and for $i \in [1, d]$ let $\mathfrak{S}_i = (D_i, \{P_{i,a} \mid a \in \tau\})$ be a τ -structure such that $D_i \cap D_j = \emptyset$ for $i \neq j$. Let $A_i \notin \tau$ be a fresh unary predicate symbol for each $i \in [1, d]$. The the *disjoint sum* $\sum_{i=1}^d \mathfrak{S}_i$ is the following structure over the signature $\tau \cup \{A_1, \dots, A_d\}$:

$$\sum_{i=1}^d \mathfrak{S}_i \stackrel{\text{def}}{=} \left(\bigcup_{i \in [1, d]} D_i, \left\{ \bigcup_{i \in [1, d]} P_{i,a} \mid a \in \tau \right\} \cup \{D_i \mid i \in [1, d]\} \right).$$

Here, $\bigcup_{i \in [1, d]} P_{i,a}$ is the interpretation for $a \in \tau$ and D_i is the interpretation for the fresh symbol A_i . Note that the fresh symbol A_i allows to recover the component structure

\mathfrak{S}_i . In other words, we can express in FO over $\sum_{i=1}^d \mathfrak{S}_i$ that \mathfrak{S}_i satisfied a certain FO-sentence. The following result is again classical [10; 11].

THEOREM 6.1. *For every FO^k-sentence φ over the signature $\tau \uplus \{A_1, \dots, A_d\}$ one can compute a tuple $(\Psi_1, \dots, \Psi_d, \beta)$, where each $\Psi_i = \{\psi_i^j \mid j \in J_i\}$ is a finite set of FO^k-sentences over the signature τ and where β is a positive boolean formula with variables from $X = \{x_i^j \mid i \in [1, d], j \in J_i\}$ such that for all τ -structures $\mathfrak{S}_1, \dots, \mathfrak{S}_d$:*

$$\sum_{i=1}^d \mathfrak{S}_i \models \varphi \quad \text{if and only if} \quad \mu \models \beta.$$

Here, $\mu: X \rightarrow \{0, 1\}$ is defined by: $\mu(x_i^j) = 1$ if and only if $\mathfrak{S}_i \models \psi_i^j$.

The following result is a simple corollary of Corollary 5.4.

COROLLARY 6.2. *For every $k \geq 3$, there is no elementary function f such that every FO^k-formula φ has a Feferman-Vaught decomposition w.r.t. disjoint sum of size $f(|\varphi|)$.*

PROOF. Recall that $\varphi_{\ell,2}$ is an $\mathbf{K}^2(A, P)$ -formula with $(A, P) \stackrel{\text{def}}{=} (\{a, \bar{a}\}, P_2 \cup \overline{P_2})$. For a structure $\overline{\mathfrak{S}}$ over $(\{\bar{a}\}, \overline{P_2})$ let $\widetilde{\mathfrak{S}}$ be the corresponding structure over $(\{a\}, P_2)$. We translate the formula $\varphi_{\ell,2}$ from above (see the proof of Theorem 5.3) into an FO³-formula $\varphi_{\ell,2}^*(x, x')$ with two free variables over the signature $\tau = \{a\} \cup P_2 \cup \{A_1, A_2\}$ inductively as follows:

- $p^*(x, x') \stackrel{\text{def}}{=} p(x)$ for each $p \in P_2$
- $\bar{p}^*(x, x') \stackrel{\text{def}}{=} p(x')$ for each $\bar{p} \in \overline{P_2}$
- $(\neg\psi)^*(x, x') \stackrel{\text{def}}{=} \neg\psi^*(x, x')$
- $(\psi_1 \wedge \psi_2)^*(x, x') \stackrel{\text{def}}{=} \psi_1^*(x, x') \wedge \psi_2^*(x, x')$
- $(\diamond\psi)^*(x, x') \stackrel{\text{def}}{=} \exists y : (A_1(y) \wedge a(x, y) \wedge \psi^*(y, x'))$
- $(\overline{\diamond}\psi)^*(x, x') \stackrel{\text{def}}{=} \exists y' : (A_2(y') \wedge a(x', y') \wedge \psi^*(x, y'))$

Note that this translation indeed yields an FO³-formula because every subformula has at most three free variables.

The reader can easily verify by induction that for every pointed structure (\mathfrak{S}, s) over $(\{a\}, P_2)$ and every pointed structure $(\overline{\mathfrak{S}}, \bar{s})$ over $(\{\bar{a}\}, \overline{P_2})$ and every $\mathbf{K}^2(A, P)$ -formula φ we have $(\mathfrak{S} \times \overline{\mathfrak{S}}, (s, \bar{s})) \models \varphi$ if and only if $\mathfrak{S} + \overline{\mathfrak{S}} \models \varphi^*(s, \bar{s})$. The family of FO³-sentences that witnesses that there are no elementarily-sized decompositions is thus simply $\{\exists x \exists x' : A_1(x) \wedge A_2(x') \wedge \varphi_{\ell,2}^*(x, x') \mid \ell \geq 1\}$ – the proof is analogous to the proof of Theorem 5.3. \square

Corollary 6.2 raises the question whether even Feferman-Vaught decompositions for FO² w.r.t. disjoint sum become nonelementary. We give a negative answer to this question.

THEOREM 6.3. *The following is computable in doubly exponential time:*
INPUT: An FO²-sentence φ over $\tau \uplus \{A_1, \dots, A_d\}$.
OUTPUT: A decomposition $(\Psi_1, \dots, \Psi_d, \beta)$, where $\Psi_i = \{\psi_i^j \mid j \in J_i\}$ is a finite set of FO²-sentences over τ and β is a positive boolean formula with variables from $X = \{x_i^j \mid$

$i \in [1, d], j \in J_i$ such that for all τ -structures $\mathfrak{S}_1, \dots, \mathfrak{S}_d$:

$$\sum_{i=1}^d \mathfrak{S}_i \models \varphi \quad \text{if and only if} \quad \mu \models \beta.$$

Here, $\mu : X \rightarrow \{0, 1\}$ is defined by: $\mu(x_i^j) = 1$ if and only if $\mathfrak{S}_i \models \psi_i^j$.

We will prove Theorem 6.3 only for the case $d = 2$; the general case can be shown in the same way. Hence, let us fix a signature τ of relational symbols and let $A_1, A_2 \notin \tau$ be two additional unary symbols. Let \mathfrak{S}_1 and \mathfrak{S}_2 be relational structures over the signature τ .

We define a partial order \leq on the set of all first-order formulas by setting $\psi_1 \leq \psi_2$ if ψ_1 is a subformula of ψ_2 . For a formula φ we denote with \mathcal{Q}_φ the set of all subformulas of φ that start with a quantifier. With $\mathcal{Q}_\varphi^{\text{cl}}$ we denote the set of those formulas in \mathcal{Q}_φ that are closed, i.e., do not have free variables. In a formula $\exists x : A_i(x) \wedge \psi$ (resp. $\forall x : A_i(x) \rightarrow \psi$), where $i \in \{1, 2\}$, we say that x is *relativized to A_i* , and for better readability we write $\exists x \in A_i : \psi$ (resp. $\forall x \in A_i : \psi$) for that formula.

A formula φ over the signature $\tau \cup \{A_1, A_2\}$ is called *pure* if φ is a boolean combination of formulas $\varphi_1, \dots, \varphi_n$ such that for every $1 \leq i \leq n$ there exists $j \in \{1, 2\}$ such that for every $(Qx : \psi) \in \mathcal{Q}_{\varphi_i}$ (where $Q \in \{\exists, \forall\}$), x is relativized to A_j in $Qx : \psi$. Equivalently, φ is pure, if the following two conditions hold:

- For all $(Qx : \psi) \in \mathcal{Q}_\varphi$, x is relativized in $(Qx : \psi)$ to either A_1 or A_2 .
- For all $(Q_1x : \psi_1), (Q_2y : \psi_2) \in \mathcal{Q}_\varphi$ with $(Q_1x : \psi_1) \leq (Q_2y : \psi_2)$, x is relativized in $(Q_1x : \psi_1)$ to the same A_i as y in $(Q_2y : \psi_2)$.

To prove Theorem 6.3 (for $d = 2$), it suffices to transform an FO²-sentence over the signature $\tau \cup \{A_1, A_2\}$ in doubly exponential time into an equivalent pure FO²-sentence over the signature $\tau \cup \{A_1, A_2\}$; this will be shown as Theorem 6.6 below.

A formula φ over the signature $\tau \cup \{A_1, A_2\}$ is called *almost pure* if it satisfies the following conditions:

- For all $(Qx : \psi) \in \mathcal{Q}_\varphi$, x is relativized in $(Qx : \psi)$ to either A_1 or A_2 .
- If $(Q_1x : \psi_1), (Q_2y : \psi_2) \in \mathcal{Q}_\varphi$ with $(Q_1x : \psi_1) \leq (Q_2y : \psi_2)$, then x is relativized in $(Q_1x : \psi_1)$ to the same A_i as y in $(Q_2y : \psi_2)$, or there exists $\theta \in \mathcal{Q}_\varphi^{\text{cl}}$ with $(Q_1x : \psi_1) \leq \theta \leq \psi_2$.

In other words, whenever a chain of subformulas $(Q_1x : \psi_1) \leq (Q_2y : \psi_2) \leq \varphi$ does not satisfy the pureness condition, then $(Q_1x : \psi_1)$ occurs within a proper subsentence of $(Q_2y : \psi_2)$ that moreover starts with a quantifier. Clearly, every pure formula is almost pure. Vice versa, we have:

LEMMA 6.4. *From a given almost pure formula φ over the signature $\tau \cup \{A_1, A_2\}$ one can compute a logically equivalent pure formula φ' of size $2^{|\mathcal{Q}_\varphi^{\text{cl}}|} \cdot O(|\varphi|)$. If φ is an FO²-formula then φ' is an FO²-formula as well.*

PROOF. The idea is to replace the topmost occurrences of sentences from the set $\mathcal{Q}_\varphi^{\text{cl}}$ by truth values in all possible ways in a big disjunction over all possible truth assignments. Since sentences from $\mathcal{Q}_\varphi^{\text{cl}}$ may also violate the pureness condition, we have to iterate this replacement step.

Let φ be almost pure and let \mathcal{F} be the set of all mappings from $\mathcal{Q}_\varphi^{\text{cl}} \setminus \{\varphi\}$ to $\{\text{true}, \text{false}\}$. For $f \in \mathcal{F}$ and a formula θ let $\theta[f]$ be the formula that results from θ by replacing every \leq -maximal formula ψ from the set $\mathcal{Q}_\varphi^{\text{cl}} \setminus \{\theta\}$ by the truth value $f(\psi)$. Then, we define φ'

as the disjunction

$$\bigvee_{f \in \mathcal{F}} (\varphi[f] \wedge \bigwedge_{\psi \in \mathcal{Q}_\varphi^{\text{cl}} \setminus \{\varphi\}} (f(\psi) \leftrightarrow \psi[f])).$$

Let us first show that φ' is pure: Since φ is almost pure, in every subformula $(Qx : \psi) \in \mathcal{Q}_\varphi$, x is either relativized to A_1 or A_2 . Hence, the same holds for φ' . Now, let $(Q_1x : \psi'_1), (Q_2y : \psi'_2) \in \mathcal{Q}_{\varphi'}$ with $(Q_1x : \psi'_1) \leq (Q_2y : \psi'_2)$. In order to get a contradiction, assume that x is relativized in $(Q_1x : \psi'_1)$ to A_1 and y is relativized in $(Q_2y : \psi'_2)$ to A_2 . The two subformulas $(Q_1x : \psi'_1)$ and $(Q_2y : \psi'_2)$ must be of the form $(Q_1x : \psi_1[f])$ and $(Q_2y : \psi_2[f])$, respectively, where $(Q_1x : \psi_1), (Q_2y : \psi_2) \in \mathcal{Q}_\varphi$ with $(Q_1x : \psi_1) \leq \theta \leq \psi_2$. Since φ is almost pure, there exists a closed formula $\theta \in \mathcal{Q}_\varphi^{\text{cl}}$ with $(Q_1x : \psi_1) \leq \theta \leq \psi_2$. But then, by construction of φ' we cannot have $(Q_1x : \psi_1[f]) \leq (Q_2y : \psi_2[f])$ in φ' , since the whole subformula θ (which contains $Q_1x : \psi_1$) is replaced by a truth value within $Q_2y : \psi_2[f]$.

Next, let us argue that φ' is equivalent to φ . For this, let us fix a τ -structure \mathfrak{S} . First, assume that $\mathfrak{S} \models \varphi(\bar{a})$ (where \bar{a} is a tuple of values for the free variables in φ ; note that we do not assume that φ is closed). We define the mapping $f : \mathcal{Q}_\varphi^{\text{cl}} \setminus \{\varphi\} \rightarrow \{\text{true}, \text{false}\}$ by $f(\psi) = \text{true}$ if and only if $\mathfrak{S} \models \psi$. Since $\mathfrak{S} \models \varphi(\bar{a})$ we get

$$\mathfrak{S} \models \varphi[f](\bar{a}) \wedge \bigwedge_{\psi \in \mathcal{Q}_\varphi^{\text{cl}} \setminus \{\varphi\}} (f(\psi) \leftrightarrow \psi[f]).$$

On the other hand, if there is a mapping $f : \mathcal{Q}_\varphi^{\text{cl}} \setminus \{\varphi\} \rightarrow \{\text{true}, \text{false}\}$ with

$$\mathfrak{S} \models \varphi[f](\bar{a}) \wedge \bigwedge_{\psi \in \mathcal{Q}_\varphi^{\text{cl}} \setminus \{\varphi\}} (f(\psi) \leftrightarrow \psi[f]),$$

then an induction on the formula size shows that for every $\psi \in \mathcal{Q}_\varphi^{\text{cl}} \setminus \{\varphi\}$, $f(\psi)$ is the truth value of the closed formula ψ in the structure \mathfrak{S} . Since moreover $\mathfrak{S} \models \varphi[f](\bar{a})$, we get $\mathfrak{S} \models \varphi(\bar{a})$.

Finally, the size of the formula

$$\varphi[f] \wedge \bigwedge_{\psi \in \mathcal{Q}_\varphi^{\text{cl}} \setminus \{\varphi\}} (f(\psi) \leftrightarrow \psi[f])$$

is in $O(|\varphi|)$ since the formulas $\varphi[f], \psi[f]$ (for $\psi \in \mathcal{Q}_\varphi^{\text{cl}} \setminus \{\varphi\}$) form a kind of partition of the whole formula φ . Hence, the size of φ' is bounded by $2^{|\mathcal{Q}_\varphi^{\text{cl}}|} \cdot O(|\varphi|)$. \square

LEMMA 6.5. *From a given FO^2 -formula $\varphi(x)$ over the signature $\tau \cup \{A_1, A_2\}$ with at most one free variable x , one can compute FO^2 -formulas $\varphi'(x)$ and $\varphi''(x)$ over the signature $\tau \cup \{A_1, A_2\}$ of size $2^{O(|\varphi|^2)}$ such that the following holds for all structures \mathfrak{S}_1 and \mathfrak{S}_2 over the signature τ .*

- $Qx \in A_1 : \varphi'(x)$ and $Qx \in A_2 : \varphi''(x)$ are almost pure (where $Q \in \{\forall, \exists\}$).
- For all $a \in \mathfrak{S}_1$, $\mathfrak{S}_1 + \mathfrak{S}_2 \models \varphi(a)$ if and only if $\mathfrak{S}_1 + \mathfrak{S}_2 \models \varphi'(a)$.
- For all $a \in \mathfrak{S}_2$, $\mathfrak{S}_1 + \mathfrak{S}_2 \models \varphi(a)$ if and only if $\mathfrak{S}_1 + \mathfrak{S}_2 \models \varphi''(a)$.

Moreover, $|\mathcal{Q}_{\varphi'(x)}^{\text{cl}}| \in 2^{O(|\varphi|)}$ and $|\mathcal{Q}_{\varphi''(x)}^{\text{cl}}| \in 2^{O(|\varphi|)}$.

PROOF. Let us construct the formula $\varphi'(x)$ ($\varphi''(x)$ is constructed analogously) by induction over the structure of the formula $\varphi(x)$. For this, we assume that $\varphi(x)$ is in negation normal form, i.e., negations appear only in front of atomic formulas. The case when $\varphi(x)$ is quantifier-free is easy: simply replace every occurrence of $A_1(x)$ by true and every occurrence of $A_2(x)$ by false.

The case that the top-most operator in $\varphi(x)$ is a boolean operator is clear, e.g., set $(\varphi_1 \wedge \varphi_2)' = \varphi_1' \wedge \varphi_2'$.

Let us now assume that $\varphi(x) = \exists y : \psi(x, y)$. Since $\varphi(x)$ is an FO²-formula, the formula $\psi(x, y)$ can be obtained from a positive boolean formula $B(p_1, \dots, p_k)$ by replacing every propositional variable p_i by

- (a) some $\alpha(x) \in \mathcal{Q}_\varphi$, where only x may occur freely, or by
- (b) some $\beta(y) \in \mathcal{Q}_\varphi$, where only y may occur freely, or by
- (c) a possibly negated atomic formula (i.e., a literal) that involves a subset of the variables $\{x, y\}$.

Let $\psi'(x, y)$ be the formula that results from $\psi(x, y)$ by replacing every subformula $\alpha(x)$ (resp., $\beta(y)$) of type (a) (resp., (b)) by $\alpha'(x)$ (resp., $\beta'(y)$). Since by induction, every formula $\exists x \in A_1 : \alpha'(x)$ and every formula $\exists y \in A_1 : \beta'(y)$ is almost pure, also $\exists x \in A_1 \exists y \in A_1 : \psi'(x, y)$ is almost pure.

By transforming B into DNF, we can write $\psi(x, y)$ as $\bigvee_{i=1}^r \psi_i$, where every ψ_i is a conjunction of formulas of the types (a)–(c). Hence, we can write ψ_i as

$$\psi_i = \alpha_i(x) \wedge \beta_i(y) \wedge \gamma_i(x, y),$$

where α_i is a conjunction of type-(a) formulas, β_i is a conjunction of type-(b) formulas, and $\gamma_i(x, y)$ is a conjunction of type-(c) formulas. Note that $r \leq 2^{|B|} \leq 2^{|\varphi(x)|}$.

Clearly, over a structure $\mathfrak{S}_1 + \mathfrak{S}_2$, the formula $\exists y : \psi(x, y)$ is equivalent to $\exists y \in A_1 : \psi(x, y) \vee \exists y \in A_2 : \psi(x, y)$, i.e., to

$$\exists y \in A_1 : \psi(x, y) \vee \bigvee_{i=1}^r \exists y \in A_2 : (\alpha_i(x) \wedge \beta_i(y) \wedge \gamma_i(x, y)).$$

By induction, for all $x \in \mathfrak{S}_1$, this formula is equivalent to

$$\exists y \in A_1 : \psi'(x, y) \vee \bigvee_{i=1}^r \exists y \in A_2 : (\alpha_i'(x) \wedge \beta_i''(y) \wedge \gamma_i(x, y)).$$

In this formula, every occurrence of a literal in $\gamma_i(x, y)$, in which both x and y occur, can be replaced either by true (if the literal is negative) or false (if the literal is positive). The reason for this is that no atomic relations of $\mathfrak{S}_1 + \mathfrak{S}_2$ involve both elements of \mathfrak{S}_1 and \mathfrak{S}_2 . Clearly, if a literal in $\gamma_i(x, y)$ is replaced by false then we can remove the whole disjunct $\exists y \in A_2 : (\alpha_i'(x) \wedge \beta_i''(y) \wedge \gamma_i(x, y))$; let us assume that this occurs for $q+1 \leq i \leq r$. We therefore obtain an equivalent formula of the form

$$\exists y \in A_1 : \psi'(x, y) \vee \bigvee_{i=1}^q (\alpha_i'(x) \wedge \delta_{i,1}(x) \wedge \exists y \in A_2 : (\beta_i''(y) \wedge \delta_{i,2}(y))).$$

Here $\delta_{i,1}(x)$ (resp., $\delta_{i,2}(y)$) is the conjunction of all literals in $\gamma_i(x, y)$ that only involve the variable x (resp., y). Let $\varphi'(x)$ be the above formula. We have to show that the formula

$$\exists x \in A_1 \left(\exists y \in A_1 : \psi'(x, y) \vee \bigvee_{i=1}^q (\alpha_i'(x) \wedge \delta_{i,1}(x) \wedge \exists y \in A_2 : (\beta_i''(y) \wedge \delta_{i,2}(y))) \right)$$

is almost pure. This follows inductively from the fact that $\exists x \in A_1 \exists y \in A_1 : \psi'(x, y)$, $\exists x \in A_1 : \alpha_i'(x)$, and $\exists y \in A_2 : \beta_i''(y)$ are almost pure, and the fact that $\exists y \in A_2 : (\beta_i''(y) \wedge \delta_{i,2}(y))$ is closed. This concludes the case $\varphi(x) = \exists y : \psi(x, y)$. The case $\varphi(x) = \forall y : \psi(x, y)$ can be treated analogously.

If we allow \wedge 's and \vee 's of arbitrary width, then the depth (i.e., the height of the syntax tree) of $\varphi'(x)$ is bounded by $O(|\varphi|)$. Due to forming CNFs and DNFs, the width of \wedge 's

and \vee 's can be bounded by $2^{|\varphi(x)|}$. Hence, the syntax tree of $\varphi'(x)$ has height $O(|\varphi|)$ and branching degree $2^{|\varphi(x)|}$, and therefore has $2^{O(|\varphi|^2)}$ nodes. Replacing \wedge 's and \vee 's of arbitrary width $\leq 2^{|\varphi(x)|}$ by 2-ary \wedge 's and \vee 's only multiplies the number of nodes by $2^{|\varphi(x)|}$. Hence, $\varphi'(x)$ is of size $2^{O(|\varphi|^2)}$.

For the bound $|\mathcal{Q}_{\varphi'(x)}^{\text{cl}}| \in 2^{O(|\varphi|)}$ note that in the above construction, the number of closed subformulas that start with a quantifier is increased by at most $q+1 \leq r+1$ (due to the formulas $\exists y \in A_2 : (\beta_i''(y) \wedge \delta_{i,2}(y))$ for $i \in [1, q]$ and possibly $\exists y \in A_1 : \psi'(x, y)$). Since r is exponential in the size of the boolean formula B , the bound $|\mathcal{Q}_{\varphi'(x)}^{\text{cl}}| \in 2^{O(|\varphi|)}$ follows. \square

THEOREM 6.6. *From a given closed FO^2 -formula φ over the signature $\tau \cup \{A_1, A_2\}$ one can compute a pure closed FO^2 -formula ψ of size $2^{2^{O(|\varphi|)}}$ such that for all structures \mathfrak{S}_1 and \mathfrak{S}_2 over the signature τ , $\mathfrak{S}_1 + \mathfrak{S}_2 \models \varphi$ if and only if $\mathfrak{S}_1 + \mathfrak{S}_2 \models \psi$.*

PROOF. We first apply Lemma 6.5 to φ and obtain a closed almost pure FO^2 -formula θ such that $\mathfrak{S}_1 + \mathfrak{S}_2 \models \varphi$ if and only if $\mathfrak{S}_1 + \mathfrak{S}_2 \models \theta$. The size of θ is bounded by $2^{O(|\varphi|^2)}$. Finally, we apply Lemma 6.4 to θ and obtain an equivalent pure FO^2 -formula ψ of size $2^{|\mathcal{Q}_{\theta}^{\text{cl}}|} \cdot O(|\theta|)$. Since $|\theta| \in 2^{O(|\varphi|^2)}$ and $|\mathcal{Q}_{\theta}^{\text{cl}}| \in 2^{O(|\varphi|)}$ this yields the upper bound $2^{2^{O(|\varphi|)}}$ for the size of ψ . \square

Let us conclude this section with a (non-matching) lower bound on Feferman-Vaught decompositions for FO^2 .

PROPOSITION 6.7. *There is no function $f(n) \in o(\sqrt{n})$ and $c > 1$ such that every FO^2 -formula φ has a Feferman-Vaught decompositions w.r.t. disjoint sum of size $c^{f(|\varphi|)}$.*

PROOF. Let us define the family of unary predicate symbols $P_n = \{p_0, \dots, p_{n-1}, p_b\}$ and $\tau_n = P_n \cup \{A_1, A_2\}$ for each $n \geq 0$.

One can define a family of FO^2 -sentences $\{\varphi_n \mid n \geq 0\}$, where each formula φ_n is defined over the signature τ_n that has precisely models of the form $\mathfrak{S}_1 + \mathfrak{S}_2$, where

- \mathfrak{S}_1 and \mathfrak{S}_2 are both P_n -structures,
- \mathfrak{S}_1 has precisely 2^n elements u_0, \dots, u_{2^n-1} ,
- $\mathfrak{S}_1 \models p_j(u_i)$ if and only if the j^{th} least significant bit of the binary representation of i is 1 (where $j \in [0, n-1]$),
- \mathfrak{S}_2 has precisely 2^n elements v_0, \dots, v_{2^n-1} ,
- $\mathfrak{S}_2 \models p_j(v_i)$ if and only if the j^{th} least significant bit of the binary representation of i is 1, and
- for every $i \in [0, 2^n - 1]$, $\mathfrak{S}_1 \models p_b(u_i)$ if and only if $\mathfrak{S}_2 \models p_b(v_i)$.

By a standard argument, one can construct FO²-formulas φ_n of size $O(n^2)$ that realize the above-mentioned properties:

$$\begin{aligned} \varphi_n &\stackrel{\text{def}}{=} \forall x : x \in A_1 \leftrightarrow x \notin A_2 \\ &\wedge \bigwedge_{i \in \{1,2\}} \exists x \in A_i : \bigwedge_{j \in [0, n-1]} \neg p_j(x) \\ &\wedge \bigwedge_{i \in \{1,2\}} \forall x, y \in A_i : \bigwedge_{j \in [0, n-1]} (p_j(x) \leftrightarrow p_j(y)) \rightarrow x = y \\ &\wedge \bigwedge_{i \in \{1,2\}} \forall x \in A_i : \bigwedge_{j \in [0, n-1]} \exists y \in A_i : p_j(x) \leftrightarrow \neg p_j(y) \wedge \bigwedge_{k \in [0, n-1] \setminus \{j\}} (p_k(x) \leftrightarrow p_k(y)) \\ &\wedge \forall x \in A_1, y \in A_2 : \bigwedge_{j \in [0, n-1]} (p_j(x) \leftrightarrow p_j(y)) \rightarrow (p_b(x) \leftrightarrow p_b(y)) \end{aligned}$$

We can assign both to \mathfrak{S}_1 and to \mathfrak{S}_2 a number in $[0, 2^{2^n} - 1]$ by simply interpreting the 2^n worlds as positions of a binary string of length 2^n . Formally, let $b_i \in \{0, 1\}$ for $i \in [0, 2^n - 1]$, where $b_i = 1$ if and only if $\mathfrak{S}_1 \models p_b(u_i)$ (resp., $\mathfrak{S}_2 \models p_b(v_i)$). We define

$$\text{val}(\mathfrak{S}_j) \stackrel{\text{def}}{=} \sum_{i=0}^{2^n-1} b_i 2^i \in [0, 2^{2^n} - 1].$$

for each $j \in \{1, 2\}$. Recall that formula φ_n enforces $\text{val}(\mathfrak{S}_1) = \text{val}(\mathfrak{S}_2)$. Also note that conversely for each $i \in [0, 2^{2^n} - 1]$ there is a unique P_n -structure $\mathfrak{S}_{1,i}^{(n)}$ and a unique P_n -structure $\mathfrak{S}_{2,i}^{(n)}$ such that $\mathfrak{S}_{1,i}^{(n)} + \mathfrak{S}_{2,i}^{(n)} \models \varphi_n$ and $\text{val}(\mathfrak{S}_{1,i}^{(n)}) = \text{val}(\mathfrak{S}_{2,i}^{(n)}) = i$. In fact, we have

$$\mathfrak{S}_{1,i}^{(n)} + \mathfrak{S}_{2,j}^{(n)} \models \varphi_n \Leftrightarrow i = j. \quad (3)$$

Assume by contradiction that there were some $c > 1$, a function $f(n) \in o(\sqrt{n})$, and for every $n \geq 1$ a decomposition

$$D_n = (\Psi^{(n)}, \Theta^{(n)}, \beta_n)$$

where

- each $\Psi^{(n)} = \{\psi_j^{(n)} \mid j \in J_n\}$ is a finite set of FO² sentences over the signature P_n ,
- each $\Theta^{(n)} = \{\theta_h^{(n)} \mid h \in H_n\}$ is a finite set of FO² sentences over the signature P_n ,
- β_n is a positive boolean formula with variables $\{x_j^{(n)} \mid j \in J_n\} \cup \{y_h^{(n)} \mid h \in H_n\}$, and
- $|D_n| \leq c^f(|\varphi_n|) \leq c^{f(O(n^2))}$.

such that for every two P_n -structure \mathfrak{S}_1 and \mathfrak{S}_2 we have

$$\mathfrak{S}_1 + \mathfrak{S}_2 \models \varphi_n \quad \Leftrightarrow \quad \mu \models \beta_n.$$

Here, μ assigns variables of β_n as follows:

- $\mu(x_j^{(n)}) = 1$ if and only if $\mathfrak{S}_1 \models \psi_j^{(n)}$ and
- $\mu(y_h^{(n)}) = 1$ if and only if $\mathfrak{S}_2 \models \theta_h^{(n)}$.

Note that the number of variables of β_n is bounded by $c^{f(O(n^2))}$. Since $f(n) \in o(\sqrt{n})$ (and thus $f(dn^2) \in o(n)$ for every constant d) there exists an n such that the number of variables of β_n is strictly smaller than 2^n . Let us fix this n in the following consideration.

For $i \in [0, 2^{2^n} - 1]$ define the truth assignment $\mu_i^{(n)}$ as follows:

- $\mu_i^{(n)}(x_j^{(n)}) \stackrel{\text{def}}{=} 1$ if and only if $\mathfrak{S}_{1,i}^{(n)} \models \psi_j^{(n)}$ and
- $\mu_i^{(n)}(y_h^{(n)}) \stackrel{\text{def}}{=} 1$ if and only if $\mathfrak{S}_{2,i}^{(n)} \models \theta_h^{(n)}$.

Since there are strictly less than 2^{2^n} truth assignments for β_n (by the choice of n), there exist $i < j$ such that $\mu_i^{(n)} = \mu_j^{(n)}$. Since $\mathfrak{S}_{1,i}^{(n)} + \mathfrak{S}_{2,i}^{(n)} \models \varphi_n$ we must have $\mathfrak{S}_{1,i}^{(n)} + \mathfrak{S}_{2,j}^{(n)} \models \varphi_n$ as well. Hence $i = j$ by (3), which is a contradiction. \square

7. GAIFMAN NORMAL FORM

Our technique from the proof of Theorem 6.3 can be used to prove a doubly exponential upper bound on the size (and construction) of Gaifman normal forms [14]. Let us start with a few definitions.

Let $\mathfrak{S} = (D, \{P_a \mid a \in \tau\})$ be a structure over a finite relational signature τ . The *Gaifman graph* of \mathfrak{S} is the undirected graph $G(\mathfrak{S}) = (D, E)$, where the edge relation E contains a pair $(u, v) \in D \times D$ with $u \neq v$ if and only if there exists a relation P_a ($a \in \tau$) of arity say n and a tuple $(u_1, \dots, u_n) \in P_a$ such that $u, v \in \{u_1, \dots, u_n\}$. For $u, v \in D$, the distance $d_{\mathfrak{S}}(u, v)$ is the length (number of edges) of a shortest path from u to v in $G(\mathfrak{S})$. For a tuple $\bar{u} = (u_1, \dots, u_n) \in D^n$ and $v \in D$, let $d_{\mathfrak{S}}(\bar{u}, v) = \min\{d_{\mathfrak{S}}(u_i, v) \mid 1 \leq i \leq n\}$. For $n \in \mathbb{N}$, the *n-sphere* around \bar{u} is $S_{\mathfrak{S},n}(\bar{u}) = \{v \in D \mid d_{\mathfrak{S}}(\bar{u}, v) \leq n\}$. We write $S_n(\bar{u})$ for $S_{\mathfrak{S},n}(\bar{u})$, if \mathfrak{S} is clear from the context.

Note that for every $n \in \mathbb{N}$, there exists a first-order formula $d_n(\bar{x}, y)$ such that for all τ -structures \mathfrak{S} and all elements \bar{u}, v of \mathfrak{S} , $\mathfrak{S} \models d_n(\bar{u}, v)$ if and only if $d_{\mathfrak{S}}(\bar{u}, v) \leq n$. For better readability, we write $d(\bar{x}, y) \leq n$ instead of $d_n(\bar{x}, y)$. The formula $d(\bar{x}, y) > n$ should be understood similarly. In a formula of the form $\exists y : d(\bar{x}, y) \leq r \wedge \psi$ or $\forall y : d(\bar{x}, y) \leq r \rightarrow \psi$, we say that the variable y is *relativized* to $S_r(\bar{x})$. A formula φ is called *r-local around \bar{x}* if for every subformula $(Qy : \psi) \in \mathcal{Q}_{\varphi}$, the variable y is relativized in $(Qy : \psi)$ to a sphere $S_q(\bar{x})$ for some $q \leq r$. A sentence ψ is called an *r-local Gaifman-sentence* if it is of the form

$$\exists x_1, \dots, x_n : \bigwedge_{1 \leq i < j \leq n} d(x_i, x_j) > 2q \wedge \bigwedge_{1 \leq i \leq n} \varphi(x_i),$$

where $\varphi(x_i)$ is q -local around (the single variable) x_i for some $q \leq r$.

THEOREM 7.1 (GAIFMAN'S THEOREM [14]). *Every first-order formula $\varphi(\bar{x})$ is equivalent to a boolean combination $\psi(\bar{x})$ of r -local formulas around \bar{x} and q -local Gaifman-sentences for suitable r and q (that are exponential in the size of $\varphi(\bar{x})$).*

We call the formula $\psi(\bar{x})$ from Theorem 7.1 a *Gaifman normal form* for $\varphi(\bar{x})$. In [12] it was shown that (for FO^4 -formulas already) the size of equivalent formulas in Gaifman normal form cannot be bounded elementarily. By using our formulas $\varphi_{\ell,n}$ from Section 3 and analogous ideas as in [12], we can strengthen the latter result to FO^3 .

PROPOSITION 7.2. *There is no elementary function f such that every FO^3 -formula φ has an equivalent formula in Gaifman normal form of size $f(|\varphi|)$.*

PROOF (SKETCH). We only give a sketch of the proof because the overall proof strategy is very similar to the proof of Theorem 4.2 in [12].

Recall that the K^2 -formula $\varphi_{\ell,2}$ was defined over $(\{a, \bar{a}\}, P_2 \cup \overline{P_2})$. Recall the translation of the K^2 -formula $\varphi_{\ell,2}$ into the FO^3 -formula $\varphi_{\ell,2}^*$ over the signature $\tau = \{a\} \cup P_2 \cup \{A_1, A_2\}$ from the proof of Corollary 6.2.

For each $i \in \{1, 2\}$ and each $j \in [0, \text{Tower}(\ell+1, 2) - 1]$ we define the τ -tree $\mathfrak{T}_{\ell,2}^{(i)}(j)$ with root $r_{\ell,2}^{(i)}(j)$ as the tree that evolves from $\Upsilon_{\ell,2}^{(0)}(j)$ (see Definition 3.1) by adding the unary predicate A_i everywhere.

Consider the structure

$$\mathfrak{F}_\ell \stackrel{\text{def}}{=} \biguplus_{j \in [0, \text{Tower}(\ell+1, 2) - 1]} \mathfrak{T}_{\ell,2}^{(1)}(j) \uplus \mathfrak{T}_{\ell,2}^{(2)}(j),$$

where \uplus denotes disjoint union. For each $\ell \geq 1$ let us define the FO^3 -formula φ_ℓ as the conjunction of the following two formulas:

$$\begin{aligned} & \text{--- } \exists x, x' : A_1(x) \wedge A_2(x') \wedge \varphi_{\ell,2}^*(x, x') \wedge \text{first}_{\ell,2}^*(x, x') \\ & \text{--- } \forall x, x' : \left(A_1(x) \wedge A_2(x') \wedge \varphi_{\ell,2}^*(x, x') \rightarrow \right. \\ & \quad \left. \text{last}_{\ell,2}^*(x, x') \vee \exists x' : (\text{succ}_{\ell,2}^*(x, x') \wedge \exists x : \varphi_{\ell,2}^*(x, x')) \right) \end{aligned}$$

Let us interpret the formula φ_ℓ on the structure \mathfrak{F}_ℓ . The two trees $\Upsilon_{\ell,2}^{(1)}(0)$ and $\Upsilon_{\ell,2}^{(2)}(0)$ witness the first conjunct of the formula φ_ℓ . The second conjunct of φ_ℓ holds in \mathfrak{F}_ℓ since for each $j \in [0, \text{Tower}(\ell+1, 2) - 1]$ either $j = \text{Tower}(\ell+1, 2) - 1$ and hence $\mathfrak{F}_\ell \models \text{last}_{\ell,2}^*(r_{\ell,2}^{(1)}(j), r_{\ell,2}^{(2)}(j))$ or $j < \text{Tower}(\ell+1, 2) - 1$ and hence $\mathfrak{F}_\ell \models \text{succ}^*(r_{\ell,2}^{(1)}(j), r_{\ell,2}^{(2)}(j+1))$ and moreover we have $\mathfrak{F}_\ell \models \varphi_{\ell,2}^*(r_{\ell,2}^{(1)}(j+1), r_{\ell,2}^{(2)}(j+1))$.

For each $j \in [0, \text{Tower}(\ell+1, 2) - 1]$ let \mathfrak{F}_ℓ^{-j} be the τ -structure that one obtains from \mathfrak{F}_ℓ by entirely removing $\mathfrak{T}_{\ell,2}^{(2)}(j)$ from it. Note that $\mathfrak{F}_\ell \models \varphi_\ell$, but $\mathfrak{F}_\ell^{-j} \not\models \varphi_\ell$ for every $j \in [0, \text{Tower}(\ell+1, 2) - 1]$.

Assume now that there were an elementary function f such that every formula φ_ℓ has an equivalent formula ψ_ℓ in Gaifman normal form of size at most $f(|\varphi_\ell|)$. Note that $f(|\varphi_\ell|)$ is elementarily bounded in ℓ . Hence, there exists ℓ such that $f(|\varphi_\ell|) < \text{Tower}(\ell+1, 2)$. Since $\mathfrak{F}_\ell \models \varphi_\ell$, we also have $\mathfrak{F}_\ell \models \psi_\ell$. We can now prove in exactly the same way as in the proof of Theorem 4.2 in [12] that there must exist $j \in [0, \text{Tower}(\ell+1, 2) - 1]$ with $\mathfrak{F}_\ell^{-j} \models \psi_\ell$, i.e., $\mathfrak{F}_\ell^{-j} \models \varphi_\ell$, which is a contradiction. \square

Next, we show that for the fragment FO^2 such an elementary (in fact, doubly exponential) bound is possible: The *quantifier rank* of a first-order formula φ is the maximal nesting depth of quantifiers in φ ; it is denoted by $\text{qr}(\varphi)$.

THEOREM 7.3. *Every FO^2 -formula $\varphi(x)$ is equivalent to a boolean combination $\psi(x)$ of r -local formulas around x and q -local Gaifman-sentences with $r \leq 3\text{qr}(\varphi)$, $q \leq 6\text{qr}(\varphi)$, and $|\psi| \leq 2^{2^{O(|\varphi|)}}$.*

In Theorem 7.3, x is a single variable. This is no restriction, since every FO^2 -formula can be written as a boolean combination of atomic formulas $a(\bar{x})$ (which are r -local around \bar{x} for every r) and formulas that (i) start with a quantifier, and (ii) that have at most one free variable. In the rest of this section, all r -local formulas will be r -local around a single variable x . For the proof of Theorem 7.3 it is useful to define *almost r -local formulas* around x and *almost r -local Gaifman-sentences*. We do this by simultaneous induction:

- Every formula that is built up from atomic formulas and almost p -local Gaifman-sentences (for arbitrary p) using boolean operators and quantifiers relativized to

$S_q(x)$ for arbitrary $q \leq r$ is an almost r -local formula around x (hence, every r -local formula around x is almost r -local around x).

— If the formula $\varphi(x_i)$ is almost q -local around x_i ($1 \leq i \leq n$) for some $q \leq r$, then the sentence

$$\exists x_1, \dots, x_n : \bigwedge_{1 \leq i < j \leq n} d(x_i, x_j) > 2q \wedge \bigwedge_{1 \leq i \leq n} \varphi(x_i) \quad (4)$$

is an almost r -local Gaifman-sentence.

For a formula φ , let $G(\varphi)$ be the set of all almost p -local Gaifman-sentences ψ (for arbitrary p) with $\psi \leq \varphi$.

LEMMA 7.4. *From an almost r -local formula $\varphi(x)$ (around x) one can compute a logically equivalent Boolean combination $\varphi'(x)$ of r -local formulas around x and q -local Gaifman sentences. Here, the size of $\varphi'(x)$ is bounded by $2^{|G(\varphi)|} \cdot O(|\varphi|)$ and q is the maximum of all p such that $G(\varphi)$ contains an almost p -local Gaifman sentence.*

PROOF. Let $\varphi(x)$ be almost r -local around x and let \mathcal{F} be the set of all mappings from $G(\varphi)$ to $\{\text{true}, \text{false}\}$. For $f \in \mathcal{F}$ and a formula θ let $\theta[f]$ be the formula that results from θ by replacing every \leq -maximal formula ψ from the set $G(\varphi) \setminus \{\theta\}$ by the truth value $f(\psi)$. Then, we define φ' as the disjunction

$$\bigvee_{f \in \mathcal{F}} (\varphi[f] \wedge \bigwedge_{\psi \in G(\varphi)} (f(\psi) \leftrightarrow \psi[f])).$$

The proof that φ' is equivalent to φ and that φ' is r -local around x is analogous to the proof of Lemma 6.4. The size of the formula

$$\varphi[f] \wedge \bigwedge_{\psi \in G(\varphi)} (f(\psi) \leftrightarrow \psi[f])$$

is in $O(|\varphi|)$ since the formulas $\varphi[f]$, $\psi[f]$ (for $\psi \in G(\varphi)$) form a kind of partition of the whole formula ψ . Hence, the size of φ' is bounded by $2^{|G(\varphi)|} \cdot O(|\varphi|)$. \square

LEMMA 7.5. *From an FO^2 -formula $\varphi(x)$ with at most one free variable x , one can compute an equivalent almost r -local formula $\varphi^\ell(x)$ of size $2^{O(|\varphi|^2)}$ with $r \leq 3qr(\varphi)$, $|G(\varphi^\ell)| \leq 2^{O(|\varphi|)}$, and every $\psi \in G(\varphi^\ell)$ is an almost $2r$ -local Gaifman sentence.*

PROOF. We prove the lemma by induction over the structure of the formula $\varphi(x)$. The case that the top-most operator in $\varphi(x)$ is a boolean operator is clear, e.g., set $(\varphi_1 \wedge \varphi_2)^\ell = \varphi_1^\ell \wedge \varphi_2^\ell$.

Now, assume that $\varphi(x) = \exists y : \psi(x, y)$. Since $\varphi(x)$ is an FO^2 -formula, the formula $\psi(x, y)$ can be obtained from a positive boolean formula $B(p_1, \dots, p_k)$ by replacing every propositional variable p_i by

- (a) a formula $\alpha(x) \in \mathcal{Q}_\varphi$, which may only contain x freely, or by
- (b) a formula $\beta(y) \in \mathcal{Q}_\varphi$, which may only contain y freely, or by
- (c) a possibly negated atomic formula (i.e., a literal) that involves a subset of the variables $\{x, y\}$.

Inductively, we replace each of the formulas $\alpha(x)$ and $\beta(y)$ in (a) and (b), respectively, by $\alpha^\ell(x)$ and $\beta^\ell(y)$, respectively. These formulas are almost r -local with $r \leq 3qr(\psi)$. Let us denote the resulting formula by $\exists y : \psi^\ell(x, y)$. It is clearly equivalent to

$$\exists y : (d(x, y) \leq 1 \wedge \psi^\ell(x, y)) \vee \exists y : (d(x, y) \geq 2 \wedge \psi^\ell(x, y)).$$

The formula $\exists y : (d(x, y) \leq 1 \wedge \psi^\ell(x, y))$ can be transformed into an almost $(r+1)$ -local formula around x . To see this, note that $d(x, y) \leq 1$ implies that every quantification

that is relativized to $S_r(y)$ can be replaced by a quantification that is relativized to $S_{r+1}(x)$, see also [14]. We will use this argument several times below.

So, let us concentrate on the second formula $\exists y : (d(x, y) \geq 2 \wedge \psi^\ell(x, y))$. We transform the boolean formula $B(p_1, \dots, p_k)$ into disjunctive normal form. Hence, for $\exists y : (d(x, y) \geq 2 \wedge \psi^\ell(x, y))$ we obtain an equivalent formula of the form

$$\exists y : (d(x, y) \geq 2 \wedge \bigvee_{i=1}^p (\alpha_i(x) \wedge \beta_i(y) \wedge \gamma_i(x, y))),$$

where

- $\alpha_i(x)$ is a conjunction of formulas $\alpha^\ell(x)$, where $\alpha(x)$ is of type (a),
- $\beta_i(y)$ is a conjunction of formulas $\beta^\ell(y)$, where $\beta(y)$ is of type (b), and
- $\gamma_i(x, y)$ is a conjunction of literals in the free variables x and y .

Note that all $\alpha_i(x)$ and $\beta_i(y)$ are almost r -local with $r \leq 3qr(\psi)$. Since we assume that $d(x, y) \geq 2$, every occurrence of an atomic formula in $\gamma_i(x, y)$, in which both x and y occur, can be replaced by false. We thus obtain an equivalent formula of the form

$$\exists y : (d(x, y) \geq 2 \wedge \bigvee_{i=1}^p (\alpha_i(x) \wedge \beta_i(y) \wedge \delta_{i,1}(x) \wedge \delta_{i,2}(y))).$$

Here $\delta_{i,1}(x)$ (resp., $\delta_{i,2}(y)$) is the conjunction of all literals in $\gamma_i(x, y)$ that only involve the variable x (resp., y). The above formula is equivalent to

$$\bigvee_{i=1}^p (\alpha_i(x) \wedge \delta_{i,1}(x) \wedge \exists y : (d(x, y) \geq 2 \wedge \beta_i(y) \wedge \delta_{i,2}(y))).$$

The formulas $\alpha_i(x) \wedge \delta_{i,1}(x)$ are almost r -local around x . So, let us concentrate on the formulas $\exists y : (d(x, y) \geq 2 \wedge \beta_i(y) \wedge \delta_{i,2}(y))$. Let us consider a specific such formula and let us just write

$$\exists y : (d(x, y) \geq 2 \wedge \theta(y)) \tag{5}$$

for it, where $\theta(y)$ is almost r -local around y . Consider the sentence

$$\begin{aligned} \rho = & \exists x_1, x_2 : (d(x_1, x_2) > 2r \wedge \theta(x_1) \wedge \theta(x_2)) \wedge \\ & \exists x_1, x_2 : (3 \leq d(x_1, x_2) \leq 2r \wedge \theta(x_1) \wedge \theta(x_2)). \end{aligned}$$

The part in the first line is an almost r -local Gaifman-sentence. The part in the second line can be rewritten as

$$\exists z : (\exists x_1, x_2 \in S_r(z) : d(x_1, x_2) \geq 3 \wedge \theta(x_1) \wedge \theta(x_2)). \tag{6}$$

This sentence is an almost $2r$ -local Gaifman-sentence (with $n = 1$ in (4)): Since $\theta(y)$ is almost r -local around y , the formula $\exists x_1, x_2 \in S_r(z) : d(x_1, x_2) \geq 3 \wedge \theta(x_1) \wedge \theta(x_2)$ is almost $2r$ -local around z .

Hence, ρ is the conjunction of an almost r -local Gaifman-sentence and an almost $2r$ -local Gaifman-sentence, which states that there exist two elements with distance at least 3 that satisfy θ .

We claim that the formula in (5) is equivalent to the following almost $(r + 3)$ -local formula around x :

$$(\exists z : \theta(z)) \wedge \left(\neg(\exists z \in S_1(x) : \theta(z)) \vee \rho \vee (\exists y \in S_3(x) : d(x, y) \geq 2 \wedge \theta(y)) \right).$$

Let us first assume that there exists y with $d(x, y) \geq 2$ and $\theta(y)$. Hence, $\exists z : \theta(z)$ holds. Moreover, assume that $\exists z \in S_1(x) : \theta(z)$ and $\neg\rho$ hold. We have to show that

$\exists y \in S_3(x) : d(x, y) \geq 2 \wedge \theta(y)$ holds. Since we have $\neg\rho$, all elements that satisfy the formula θ have pairwise distance at most 2. Since there is an element of distance at most 1 from x that satisfies θ , the element y with $d(x, y) \geq 2$ and $\theta(y)$ has distance at most 3 from x .

For the other direction, assume that $\exists z : \theta(z)$ and one of $\neg(\exists z \in S_1(x) : \theta(z))$, ρ , or $\exists y \in S_3(x) : d(x, y) \geq 2 \wedge \theta(y)$ holds. We have to show that there exists y with $d(x, y) \geq 2$ and $\theta(y)$. The case that $\neg(\exists z \in S_1(x) : \theta(z))$ or $\exists y \in S_3(x) : d(x, y) \geq 2 \wedge \theta(y)$ holds is clear. If ρ holds, then there exist two elements with distance at least 3 that satisfy θ . Since two elements in $S_1(x)$ have distance at most 2, there must exist $y \notin S_1(x)$ satisfying θ .

We have shown that $\varphi(x) = \exists y : \psi(x, y)$ is equivalent to an almost $(r+3)$ -local formula around x (with $r \leq \text{qr}(\psi)$), which we can take for $\varphi^\ell(x)$. Hence, $\varphi^\ell(x)$ is indeed almost r' -local for some $r' \leq 3\text{qr}(\varphi)$. Moreover, every sentence in $G(\varphi^\ell(x))$ is an almost $2r'$ -local Gaifman sentence (the factor 2 comes from the formula (6)).

In order to bound the size of $\varphi^\ell(x)$, note that the depth of $\varphi^\ell(x)$ is bounded by $O(|\varphi|)$ if we allow \wedge 's and \vee 's of arbitrary width. Since the width can be bounded by $2^{|\varphi(x)|}$, the size of $\varphi^\ell(x)$ can be bounded by $2^{O(|\varphi|^2)}$, see the proof of Lemma 6.5 for an analogous argument.

For the bound $|G(\varphi^\ell)| \leq 2^{O(|\varphi|)}$ note that in the above construction, the number of almost local Gaifman sentences that are introduced is bounded by $O(p)$. Since p is exponential in the size of the boolean formula B , the bound $|G(\varphi^\ell)| \leq 2^{O(|\varphi|)}$ follows. Again, see the proof of Lemma 6.5 for an analogous argument. \square

Let us finally prove Theorem 7.3. We first apply Lemma 7.5 to $\varphi(x)$ and obtain an equivalent almost r -local formula $\theta(x)$ with $|\theta| \leq 2^{O(|\varphi|^2)}$. Moreover $r \leq 3\text{qr}(\varphi)$ and every sentence in $G(\theta)$ is an almost $2r$ -local Gaifman sentence. Finally, we apply Lemma 7.4 to θ and obtain an equivalent Boolean combination $\psi(x)$ of r -local formulas around x and $2r$ -local Gaifman sentences. The size of $\psi(x)$ is bounded by $2^{|G(\theta)|} \cdot O(|\theta|)$. Since $|\theta| \leq 2^{O(|\varphi|^2)}$ and $|G(\theta)| \leq 2^{O(|\varphi|)}$, this yields the upper bound $2^{2^{O(|\varphi|)}}$ for the size of $\psi(x)$. \square

Finally, we give a (non-matching) lower bound on the size of equivalent formulas in Gaifman normal form for FO^2 ; the proof is again based on techniques from [12].

PROPOSITION 7.6. *There is no function $f(n) \in o(\sqrt{n})$ and $c > 1$ such that every FO^2 -formula φ has an equivalent formula in Gaifman normal form of size $c^{f(|\varphi|)}$.*

PROOF. This proof uses a very similar strategy as in [12]. Let us define the signature $\tau_n = \{p_0, \dots, p_{n-1}\}$ of solely unary predicate symbols for each $n \geq 1$. For each $n \geq 1$ it is standard to define an FO^2 -sentence φ_n of size $O(n^2)$ such that there is a unique τ_n -structure \mathfrak{S}_n with $\mathfrak{S}_n \models \varphi_n$, where \mathfrak{S}_n satisfies the following properties:

- \mathfrak{S}_n has 2^n elements u_0, \dots, u_{2^n-1} and
- $\mathfrak{S}_n \models p_j(u_i)$ if and only if the j^{th} least significant bit of the binary representation of i is 1 (where $j \in [0, n-1]$).

Assume by contradiction that there is a function $f(n) \in o(\sqrt{n})$ and $c > 1$ such that for every $n \geq 1$, φ_n has an equivalent sentence ψ_n in Gaifman normal form with $|\psi_n| \leq c^{f(|\varphi_n|)} = c^{f(O(n^2))}$. Hence, we have $\mathfrak{S}_n \models \psi_n$. Since $f(n) \in o(\sqrt{n})$ there is an n such that $|\psi_n| < 2^n$. Let us fix such an n in the following.

The sentence ψ_n is a boolean combination of sentences χ_1, \dots, χ_ℓ , where each sentence

$$\chi_i = \exists x_1, \dots, x_{n_i} : \chi'_i$$

is an r_i -local Gaifman-sentence. Without loss of generality, we can assume that there is some $h \in [1, \ell]$ such that

- (1) $\mathfrak{S}_n \models \chi_i$ for each $i \in [1, h]$ and
- (2) $\mathfrak{S}_n \not\models \chi_i$ for each $i \in [h+1, \ell]$.

Recall that \mathfrak{S}_n consists precisely of the elements u_0, \dots, u_{2^n-1} . By (1) we have that for each $i \in [1, h]$, we can fix elements $w_1^{(i)}, \dots, w_{n_i}^{(i)} \in \{u_0, \dots, u_{2^n-1}\}$ that witness $\mathfrak{S}_n \models \chi_i$, i.e. $\mathfrak{S}_n \models \chi_i'(w_1^{(i)}, \dots, w_{n_i}^{(i)})$. Let $d \stackrel{\text{def}}{=} \sum_{i=1}^h n_i \leq \sum_{i=1}^{\ell} n_i \leq |\psi_n| < 2^n$. For each $j \in [0, 2^n - 1]$, let the τ_n -structure \mathfrak{S}_n^{-j} be obtained from \mathfrak{S}_n by removing the element u_j . Note that obviously $\mathfrak{S}_n^{-j} \not\models \varphi_n$ for every $j \in [0, 2^n - 1]$. Since $d < 2^n$, by the pigeonhole principle, there exists some $j \in [0, 2^n - 1]$ such that

- $u_j \notin \{w_k^{(i)} \mid i \in [1, h], k \in [1, n_i]\}$ and thus
- $\mathfrak{S}_n^{-j} \models \chi_i$ (since χ_i is an r_i -local Gaifman-sentence) for each $i \in [1, h]$.

Recall that by (2) we have $\mathfrak{S}_n \not\models \chi_i$, or equivalently

$$\mathfrak{S}_n \models \neg \exists x_1, \dots, x_{n_i} : \chi_i'$$

for each $i \in [h+1, \ell]$. Also, note that since each formula χ_i is an r_i -local Gaifman-sentence it follows $\mathfrak{S}_n^{-j} \models \neg \chi_i$ for each $i \in [h+1, \ell]$. In total we have

- $\mathfrak{S}_n^{-j} \models \chi_i$ for each $i \in [1, h]$ and
- $\mathfrak{S}_n^{-j} \not\models \chi_i$ for each $i \in [h+1, \ell]$.

and hence $\mathfrak{S}_n^{-j} \models \psi_n$, contradicting $\mathfrak{S}_n^{-j} \not\models \varphi_n$. \square

8. CONCLUSIONS AND OPEN PROBLEMS

We have defined a class of trees that can be enforced in two-dimensional modal logic $\mathbf{K} \times \mathbf{K}$. Using these trees we were able to show nonelementary lower bounds for satisfiability in $\mathbf{K} \times \mathbf{K}$ and, via reductions from this, for $\mathbf{K4} \times \mathbf{K}$ and $\mathbf{S5}_2 \times \mathbf{K}$. As our trees have very large outdegree, we believe that it is not possible to adapt our techniques to the case when the frames in one dimension are restricted to be some kind of ‘linear’, e.g., the frame $(\mathbb{N}, <)$ or transitive, weakly connected frames, so called $\mathbf{K4.3}$ -frames. However, it would be interesting to study the precise complexity of such logics as well. Note that there have been results in this direction. For example, it is known that satisfiability in $\mathbf{PTL} \times \mathbf{K}$ and $\mathbf{PTL} \times \mathbf{S5}_2$ is hard for nonelementary time [3; 24], where \mathbf{PTL} is propositional linear temporal logic including the ‘until’-operator, which is not available in our setting.

Fusions are another way to combine two modal logics; intuitively, they are the independent join of the two participating logics. For example, the fusion of \mathbf{K} and \mathbf{K} is just the bimodal logic \mathbf{K}_2 . While it is well-known that satisfiability in \mathbf{K}_2 is PSPACE-complete [6], nothing is known for combinations that lie between fusions and products. In particular, \mathbf{K}_2 can be axiomatized by taking two copies of the well-known \mathbf{K} -axiom, one for each modal operator; an axiomatization for $\mathbf{K} \times \mathbf{K}$ can be obtained from the axiomatization for the fusion by adding axioms postulating that the two modal operators commute and have the Church-Rosser property. Thus, combinations between fusion and product can be obtained by dropping either Church-Rosser or half of commutativity, and it would be interesting to see which of the additional axioms suffice to prove the nonelementary lower bound.

An open problem from the second part of the paper concerns the size of Feferman-Vaught decompositions (w.r.t. disjoint sum) and equivalent formulas in Gaifman normal form for FO^2 . For both formalisms, we proved a doubly exponential upper bound

and a lower bound of the form $c^{o(\sqrt{n})}$ (for any constant $c > 1$). We conjecture that the upper bound can be improved to a singly exponential bound.

The Łoś-Tarski theorem gives another normal form: the first-order sentences that are preserved under extensions coincide with the ones that are equivalent to an existential sentence. It has been shown in [12, Theorem 6.1] that the size of this normal form is also inherently nonelementary, however the used formulas had at least four free variables. Of course one might ask whether our technique yields formulas in FO^3 .

A further potential question could be to study logics for which Feferman-Vaught decompositions do not exist in general (like CTL) and to find a (possibly decidable) characterization of those formulas which are decomposable.

Acknowledgments

We would like to thank the anonymous reviewers for their helpful comments.

REFERENCES

- P. Blackburn, M. de Rijke, and Y. Venema, *Modal Logic*. Cambridge University Press, 2001.
- P. Blackburn, F. Wolter, and J. van Benthem, Eds., *Handbook of Modal Logic*. Elsevier, 2006.
- D. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyashev, *Many-Dimensional Modal Logics: Theory and Applications*. Elsevier, 2003.
- M. Marx and Y. Venema, *Multi-Dimensional Modal Logic*. Kluwer Academic Press, 1996.
- D. M. Gabbay and V. B. Shehtman, “Products of Modal Logics, Part 1,” *Logic Journal of the IGPL*, vol. 6, no. 1, pp. 73–146, 1998.
- R. E. Ladner, “The Computational Complexity of Provability in Systems of Modal Propositional Logic,” *SIAM J. Comput.*, vol. 6, no. 3, pp. 467–480, 1977.
- R. Hirsch, I. M. Hodkinson, and Á. Kurucz, “On modal logics between $K \times K \times K$ and $S5 \times S5 \times S5$,” *J. Symb. Log.*, vol. 67, no. 1, pp. 221–234, 2002.
- D. Gabelaia, A. Kurucz, F. Wolter, and M. Zakharyashev, “Products of ‘transitive’ modal logics,” *J. Symb. Log.*, vol. 70, no. 3, pp. 993–1021, 2005.
- A. Mostowski, “On direct products of theories,” *J. Symbolic Logic*, vol. 17, pp. 1–31, 1952.
- S. Feferman and R. L. Vaught, “The first order properties of products of algebraic systems,” *Fundamenta Mathematicae*, vol. 47, pp. 57–103, 1959.
- J. A. Makowsky, “Algorithmic uses of the Feferman-Vaught Theorem,” *Ann. Pure Appl. Logic*, vol. 126, no. 1–3, pp. 159–213, 2004.
- A. Dawar, M. Grohe, S. Kreutzer, and N. Schweikardt, “Model Theory Makes Formulas Large,” in *Proc. of ICALP*, ser. Lecture Notes in Computer Science, vol. 4596. Springer, 2007.
- A. Rabinovich, “On compositionality and its limitations,” *ACM Trans. Comput. Log.*, vol. 8, no. 1, 2007.
- H. Gaifman, “On local and nonlocal properties,” in *Logic Colloquium '81*, J. Stern, Ed. North Holland, 1982, pp. 105–135.
- S. Kreutzer, “Algorithmic meta-theorems,” *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 16, p. 147, 2009.
- M. Marx and S. Mikuláš, “Products, or How to Create Modal Logics of High Complexity,” *Logic Journal of the IGPL*, vol. 9, no. 1, pp. 71–82, 2001.
- S. Göller and A. W. Lin, “Concurrency Makes Simple Theories Hard,” in *Proc. of STACS*, ser. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012, pp. 344–355.
- E. Grädel, P. G. Kolaitis, and M. Y. Vardi, “On the decision problem for two-variable first-order logic,” *Bulletin of Symbolic Logic*, vol. 3, no. 1, pp. 53–69, 1997.
- S. Göller, J. C. Jung, and M. Lohrey, “The complexity of decomposing modal and first-order theories,” in *Proceedings of the 27th Annual IEEE Symposium on Logic in Computer Science, LICS 2012*. IEEE, 2012, pp. 325–334.
- J. Flum and M. Grohe, *Parametrized Complexity Theory*. Springer, 2006.
- E. Börger, E. Grädel, and Y. Gurevich, *The classical decision problem*, ser. Universitext. Berlin: Springer-Verlag, 2001.
- B. S. Chlebus, “From domino tilings to a new model of computation,” in *Symposium on Computation Theory*, ser. Lecture Notes in Computer Science, A. Skowron, Ed., vol. 208. Springer, 1984, pp. 24–33.

M. Marx, "Complexity of products of modal logics," *J. Log. Comput.*, vol. 9, no. 2, pp. 197–214, 1999.

J. Y. Halpern and M. Y. Vardi, "The complexity of reasoning about knowledge and time. I. lower bounds," *J. Comput. Syst. Sci.*, vol. 38, no. 1, pp. 195–237, 1989.