

Elementary Modal Logics over Transitive Structures

Jakub Michaliszyn^{1,2} and Jan Otop^{1,3}

- 1 University Of Wrocław
- 2 Imperial College London
- 3 IST Austria

Abstract

We show that modal logic over universally first-order definable classes of transitive frames is decidable. More precisely, let \mathcal{K} be an arbitrary class of transitive Kripke frames definable by a universal first-order sentence. We show that the global and finite global satisfiability problems of modal logic over \mathcal{K} are decidable in NP, regardless of choice of \mathcal{K} . We also show that the local satisfiability and the finite local satisfiability problems of modal logic over \mathcal{K} are decidable in NEXPTIME.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases Modal logic, Transitive frames, Elementary modal logics, Decidability

Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

1 Introduction

Modal logic was first introduced by philosophers as the study of the deductive behaviour of the expressions ‘it is necessary that’ and ‘it is possible that’. Nowadays, it is widely used in several areas of computer science, including formal verification and artificial intelligence.

Syntactically, modal logic extends propositional logic by two unary operators: \Diamond and \Box . The formal semantics is usually given in terms of Kripke structures. Basically, a Kripke structure is a directed graph, called a *frame*, together with a valuation of propositional variables. Vertices of this graph are called *worlds*. For each world truth values of all propositional variables can be defined independently. In this semantics, $\Diamond\varphi$ means *the current world is connected to some world in which φ is true*; and $\Box\varphi$, equivalent to $\neg\Diamond\neg\varphi$, means *φ is true in all worlds to which the current world is connected*.

“Classical” modal logic, defined as above, is very simple, and therefore it has limited applications. For that reason, many modifications of modal logic are studied. One way to enrich modal logic is to add more modalities and obtain so called *multimodal logic*. Another popular modification is to add some constraints on the interpretation of operators, e.g. by requiring that the modal operator represents a relation that is reflexive and transitive (S4). Finally, by combining these two techniques, we may obtain multimodal logics with nonuniform modal operators, like Linear Temporal Logic (LTL), Computation Tree Logic (CTL) or Halpern–Shoham logic (HS).

The first author was generously supported by Polish National Science Center based on the decision number DEC-2011/03/N/ST6/00415. The second author was supported in part by the Austrian Science Fund NFN RiSE (Rigorous Systems Engineering) and by the ERC Advanced Grant QUAREM (Quantitative Reactive Modeling).

© J. Michaliszyn and J. Otop
Conference title on which this volume is based on.
Editors: Billy Editor, Bill Editors; pp. 1–15



Leibniz International Proceedings in Informatics
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Variants of modal logic vary in the complexity and the decidability of the satisfiability problem. While there are some logics, like S5, that are NP-complete, most of them are PSPACE-hard. Modal logic itself is PSPACE-complete, and so is the temporal logic LTL. Logics CTL and CTL* are even harder, EXPTIME-complete and 2-EXPTIME-complete, resp. Finally, the Halpern–Shoham logic is a simple example of a temporal logic that is undecidable, even if we consider some unimodal fragments [3, 16].

There are several ways of adding constraints on the interpretation of operators. The syntactic way goes by adding some additional axioms and considering only modal logics that satisfy these axioms. Another way is by restricting the class of the admissible frames. It also can be done in many ways, but one of the simplest is to define the class of frames by a first order logic sentence that uses a single binary relation R , which is interpreted as the *transition relation*. For example, the sentence $\forall xyz.xRy \wedge yRz \Rightarrow xRz$ defines the class of all transitive frames. Modal logic over a class of frames definable by a first order logic sentence is called *an elementary modal logic*.

The main goal of our work is to classify all elementary modal logics with respect to decidability of their satisfiability problems. In [9], it was shown that there is an universal first-order formula such that the global satisfiability problem over the class of frames that satisfy this formula is undecidable. A slight modification of that formula yields an analogous result for the local satisfiability problem. In [12] it was shown that even a very simple formula with three variables without equality leads to undecidability.

Many modal logics used in automatic verification contain operators that are interpreted as transitive relations. For example, Linear Temporal Logic (LTL) contains transitive operators F and G [2]. In epistemic modal logics, the knowledge operators K_i are interpreted as relations that are not only transitive, but also reflexive and symmetric [4]. Another example is the logic of subintervals [16], which is a fragment of Halpern–Shoham logic with a single modality that can be read as “in some subinterval”.

Main results. In this paper, an logic is called a *subframe logic* if it is an unimodal logic defined by restricting the class of the admissible frames to a class that is closed under subframes. Fine[5] showed that there are undecidable transitive subframe logics. An open question suggested in [9] is whether there is an undecidable transitive subframe logic that is an elementary modal logic. Due to [5], such a logic would not be finitely axiomatizable. We show that such a logic does not exist.

The Łoś–Tarski preservation theorem (see, e.g., [10]) states that a first order definable class of frames is closed under subframes if and only if it is universally first order definable. Therefore, to answer the question discussed above, we may restrict our attention to universal formulae. We prove the following theorem.

► **Theorem 1.** *Let \mathcal{K} be a class of frames defined by a universal first order formula that implies transitivity. The local and the global satisfiability problems for unimodal logic over \mathcal{K} are decidable.*

The finiteness constraint can make the satisfiability problem easier or harder. There are decidable modal logics that are finitely undecidable, and there are undecidable modal logics that are finitely decidable [6, 19]. However, this is not the case here — the finite satisfiability problems are decidable as well.

► **Theorem 2.** *Let \mathcal{K} be a class of frames defined by a universal first order formula that implies transitivity. The local and the global finite satisfiability problems for unimodal logic over \mathcal{K} are decidable.*

We focus on the case where a first-order formula, that defines the class of frames, is a parameter of the problem, and the input consists of a modal formula only. However, our results hold even if we allow a first-order formula to be a part of the input.

Related work. Decidability of modal logic over various classes of frames can be shown by employing the so-called standard translation of modal logic to first-order logic [2]. Indeed, the satisfiability of a modal formula φ in \mathcal{K}_Φ is equivalent to satisfiability of $ST(\varphi) \wedge \Phi$, where $ST(\varphi)$ is the standard translation of φ . In this way, we can show that (multi)modal logic is decidable over any class defined by two-variable logic [20], even extended with a linear order [21], counting quantifiers [22], one transitive relation [24], or equivalence closures of two distinguished binary relations [13]. The same holds for formulae of the guarded fragment [7], even if we allow for some restricted application of transitive relations [14, 23], fixed-points [1, 8] and transitive closures [17]. In many cases, the decidability results hold also when only finite frames are considered. However, the complexity bounds obtained this way are high — usually between EXPTIME and 2NEXPTIME.

Many natural classes of frames, including transitive, reflexive, symmetric and Euclidean, can be defined by first-order sentences even if we further restrict the language to universal Horn formulae. In [18], it was shown that the problems defined by universal Horn formulae are always decidable, and the precise complexity (depending on properties of first-order formulae) was established. In [11], the decidability of elementary modal logics defined by universal Horn formulae was extended to the finite satisfiability case.

Observe however, that the above results do not solve our problem — for example, the class of transitive frames that do not contain cliques of size 3 cannot be defined in any of the languages mentioned above.

Overview of the paper. The proof of Theorem 2 is much easier than the proof of Theorem 1. For a given modal formula φ and its transitive model, we show that we can remove a world from any path containing at least $|\varphi| + 1$ worlds such that the resulting structure is a model. By iterating this procedure, we obtain a model in which there are no paths longer than $|\varphi|$, and the existence of such a model can be verified by an NEXPTIME algorithm. The whole proof of Theorem 2 is in Section 3.

The general satisfiability case is discussed in Section 4. We start with the global satisfiability problem in Section 4.1. We show that each satisfiable formula has a model which is a clique or an infinite path.

In Section 4.2 we present the proof for the local satisfiability problem. This proof is much more complicated than the other proofs. The general shape of the proof is similar to the previous proofs: we show that each satisfiable formula has a “nice” model, i.e., a model whose description is exponential w.r.t. the size of the modal formula, and that description admits efficient algorithm checking whether described model satisfies a given modal (or first order) formula.

The goal of this paper is to prove the decidability. However, we also discuss the complexity that arises from our algorithms. In the case of the global (finite and infinite) satisfiability problems, we provide the optimal complexity proving that these problems is always in NP (see Proposition 10). The corresponding lower bound comes from the trivial reduction of SAT problem. For the local satisfiability problem we show only that the membership is in NEXPTIME (see Propositions 5 and 15), which is not optimal — for example, the satisfiability problem of modal logic over the class of all transitive frames is PSPACE-complete.

Finally, the possible future work is discussed in Section 5.

2 Preliminaries

As we work with both first-order logic and modal logic we help the reader to distinguish them in our notation; we denote first-order formulae with Greek capital letters, and modal formulae with Greek small letters. We assume that the reader is familiar with first-order logic and propositional logic.

Modal logic. Formulae of modal logic are interpreted in Kripke structures, which are triples of the form $\langle M, R, \pi \rangle$, where M is a set of worlds, $\langle M, R \rangle$ is a directed graph called a *frame*, and π is a *labelling*, a function that assigns to each world a set of propositional variables which are true at this world. We say that a structure $\langle M, R, \pi \rangle$ is *based* on the frame $\langle M, R \rangle$. For a given class of frames \mathcal{K} , we say that a structure is \mathcal{K} -based if it is based on some frame from \mathcal{K} . We will use calligraphic letters \mathcal{M}, \mathcal{N} to denote frames and Fraktur letters $\mathfrak{M}, \mathfrak{N}$ to denote structures. Whenever we consider a structure \mathfrak{M} , we assume that its frame is \mathcal{M} and its universe is M (and the same holds for other letters).

The semantics of modal logic is defined recursively. A modal formula φ is (locally) *satisfied* in a world w of a model $\mathfrak{M} = \langle M, R, \pi \rangle$, denoted as $\mathfrak{M}, w \models \varphi$ if

- $\varphi = p$, where p is a variable, and $\varphi \in \pi(w)$,
- $\varphi = \neg\varphi'$ and $\mathfrak{M}, w \not\models \varphi'$,
- $\varphi = \varphi_1 \wedge \varphi_2$ and $\mathfrak{M}, w \models \varphi_1$ and $\mathfrak{M}, w \models \varphi_2$,
- $\varphi = \Diamond\varphi'$ and there is a world $v \in M$ such that $(w, v) \in R$ and $\mathfrak{M}, v \models \varphi'$,

Boolean connectives $\vee, \Rightarrow, \Leftrightarrow$ and constants \top, \perp are introduced in the standard way. We abbreviate $\neg\Diamond\neg\varphi$ by $\Box\varphi$.

We say that a formula φ is *globally* satisfied in \mathfrak{M} , denoted as $\mathfrak{M} \models \varphi$, if for all worlds w of \mathfrak{M} , we have $\mathfrak{M}, w \models \varphi$. By $|\varphi|$ we denote the length of φ .

For a given class of frames \mathcal{K} , we say that a formula φ is *locally* (resp. *globally*) \mathcal{K} -*satisfiable* if there exists a \mathcal{K} -based structure \mathfrak{M} , and a world $w \in W$ such that $\mathfrak{M}, w \models \varphi$ (resp. $\mathfrak{M} \models \varphi$). We study four versions of the satisfiability problem.

For a given formula φ , a Kripke structure \mathfrak{M} , and a world $w \in M$ we define the *type* of w (with respect to φ) in \mathfrak{M} as $tp_{\mathfrak{M}}^{\varphi}(w) = \{\psi : \mathfrak{M}, w \models \psi \text{ and } \psi \text{ is a subformula of } \varphi\}$. We write $tp_{\mathfrak{M}}(w)$ if the formula is clear from the context. Note that $|tp_{\mathfrak{M}}^{\varphi}(w)| \leq |\varphi|$.

First-order logic. The class of (equality-free) universal first order sentences is defined as a subclass of first-order sentences such that each sentence is of the form $\forall \vec{x} \Psi(\vec{x})$, where $\Psi(\vec{x})$ is quantifier-free formula over the language $\{R\}$, where R is a binary relation symbol interpreted as the successor relation in modal logic. As we work only with universally quantified formulae, we often skip the quantifier prefix, e.g., write xRx instead of $\forall x.xRx$.

The set of *transitive formulae*, \forall_T , is defined as the set of those Φ over the language $\{R\}$ which are of the form $(\forall xyz.xRy \wedge yRz \Rightarrow xRz) \wedge \Psi$, where Ψ is an arbitrary universal first-order formula.

Decision problems. The *local* (resp. *global*) *satisfiability problem* \mathcal{K} -SAT (resp. \mathcal{K} -GSAT) as follows. Is a given modal formula locally (resp. globally) \mathcal{K} -satisfiable? The *finite local* (global) *satisfiability problem*, \mathcal{K} -FINSAT (\mathcal{K} -GFINSAT), is defined in the same way, but we are only interested in finite models (the class \mathcal{K} may still contain infinite structures).

In context of decision problems, we use the word “general” as an antonym of “finite”. Furthermore, for a given universal first-order formula Φ , we define \mathcal{K}_{Φ} as the class of frames satisfying Φ .

3 Finite satisfiability

In this section, we show that if a modal formula has a finite model in which the transition relation is transitive, then it has one with the size bounded exponentially in the size of modal formula. This clearly leads to an NEXPTIME algorithm that simply guesses an exponential model and verifies it.

For a given modal formula φ , we say that a world w in a model \mathfrak{M} is *special* w.r.t. φ , if there is a subformula ψ of φ such that w satisfies ψ , no successor of w different than w satisfies ψ . Worlds that are not special w.r.t. φ are *regular* (w.r.t. φ). We omit explicit reference to φ when it is clear from the context.

The decidability of FINSAT and GFINSAT comes from the following lemma.

► **Lemma 3.** *Let Φ be a \forall_T formula, φ be a modal formula, \mathfrak{M} be a model satisfying Φ , w be a regular (w.r.t. φ) world in \mathfrak{M} , and \mathfrak{N} be the result of removing w (and all connected edges) from \mathfrak{M} . Then, for each $v \in N$, $tp_{\mathfrak{M}}^{\varphi}(v) = tp_{\mathfrak{N}}^{\varphi}(v)$.*

Proof. The proof goes by induction on φ . Consider a subformula ψ of φ . The only non-trivial case is when $\psi = \diamond\psi'$, w satisfies ψ' and a world w' is a predecessor of w . Since w is regular, there is a world $v \neq w$ reachable from w that satisfies ψ' in \mathfrak{M} and, by inductive assumptions, in \mathfrak{N} . Since R is transitive, v is a successor of w' , so w' satisfies ψ in \mathfrak{N} . ◀

► **Example 4.** Consider a formula Φ that states that R is transitive, $\varphi = \square\square\perp$ and the structure $\mathfrak{M} = \langle \{a, b, c\}, \{(a, b), (b, c), (a, c)\}, \pi \rangle$ with $\mathfrak{M}, a \not\models \varphi$. It may seem that the worlds b and c are regular, but if we removed any of them, the type of a would not contain φ .

However, \square is an abbreviation, thus an unrolled form of φ is $\neg\diamond\neg\diamond\neg\perp$. Then, only a satisfies $\diamond\neg\diamond\neg\perp$, b (but not c) satisfies $\diamond\neg\perp$, and only c satisfies $\neg\diamond\neg\perp$. Thus, all the worlds in \mathfrak{M} are special w.r.t. φ .

We would like to thank an anonymous reviewer for suggesting the above example. Theorem 2 follows from the following proposition.

► **Proposition 5.** *Let Φ be a \forall_T formula. Then \mathcal{K}_{Φ} -FINSAT and \mathcal{K}_{Φ} -GFINSAT are decidable in NEXPTIME.*

Proof. Let φ be a modal formula with a finite model \mathfrak{M} based on a frame from \mathcal{K}_{Φ} . Since the relation of \mathfrak{M} is transitive, by iterating Lemma 3, we can obtain a model \mathfrak{N} without regular worlds. Since all worlds in \mathfrak{N} are special, the maximal length of a path of different worlds in \mathfrak{N} is bounded by $|\varphi|$.

We obtain a small model by applying standard selection technique [2, 15]. In first stage, we mark an arbitrary world satisfying φ . Then, in consecutive stages, for each world w marked in the previous stage and each subformula ψ of φ , if w is connected to a world satisfying ψ , but not to a marked world satisfying ψ , then we mark one successor of w satisfying ψ . Note that this procedure ends after $|\varphi|$ stages — otherwise there would be a path containing more than $|\varphi|$ different worlds. Finally, we remove all worlds but marked. It is not hard to see that the types of remaining worlds do not change and that the size of obtained model is bounded by $|\varphi|^{|\varphi|}$. We have shown that the considered logic has the exponential model property. The NEXPTIME algorithm simply guesses an exponential model and verifies it. ◀

4 General satisfiability

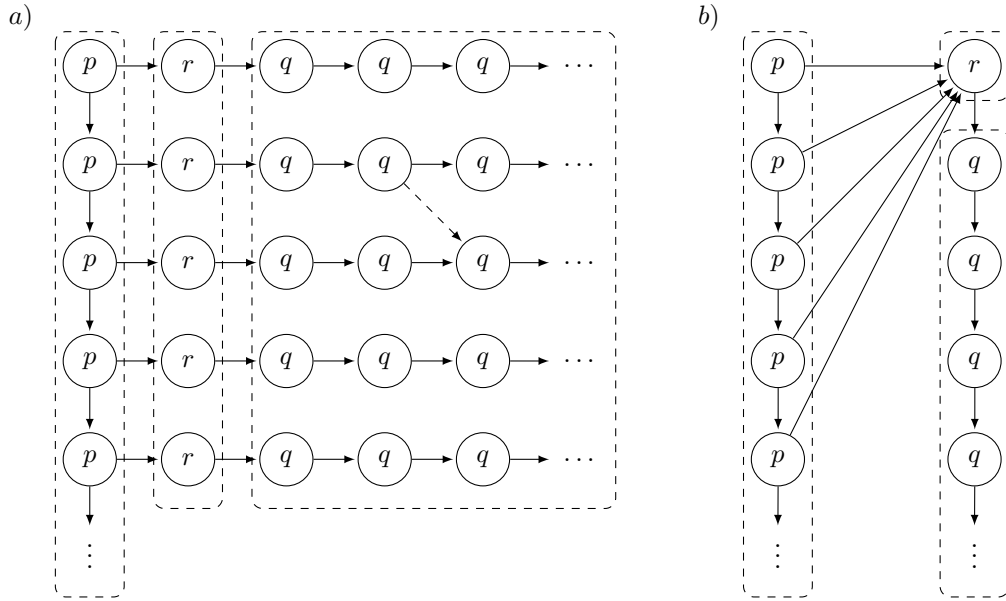
Consider the formula that defines the class of transitive and irreflexive frames:

$$\Gamma = (\forall xyz.xRy \wedge yRz \Rightarrow xRz) \wedge \forall x.\neg xRx$$

A quick check shows that a modal formula $\Diamond\top \wedge \Box\Diamond\top$ has only infinite models satisfying Γ . In order to illustrate some later concepts, we consider a more complex formula ν which is a conjunction of the following properties.

1. $\Box\neg((p \wedge q) \vee (p \wedge r) \vee (q \wedge r)) \wedge \Box(p \vee q \vee r)$ stating that each reachable world satisfies exactly one of variables p, q, r (recall that we consider transitive structures).
2. $\Diamond p \wedge \Box(p \Rightarrow \Diamond p)$ stating that there is a reachable world satisfying p , and each reachable world satisfying p has a successor satisfying p .
3. $\Box(p \Rightarrow \Diamond r) \wedge \Box(r \vee q \Rightarrow \Diamond q \wedge \Box q)$ stating that each reachable world satisfying p has a successor satisfying r , and each reachable world satisfying r or q has a successor satisfying q and all its successors satisfy q .

Fig. 1 contains two models of ν satisfying Γ (the formula is satisfied at the root). It is not hard to see that each such model has to contain two infinite chains: one labelled by p and one labelled by q .



■ **Figure 1** The running example — two models of ν over Γ . The edges that follow from the transitivity are omitted for better readability. Note that the structure a) is a model of ν even if we remove the dashed edge.

4.1 Disentangled structures and global satisfiability

By employing the standard translation [2] and the Löwenheim–Skolem theorem we know that each satisfiable modal formula has a countable model over any class of first-order definable frames. In this section we consider only countable structures.

In order to simplify the exposition in this section, we assume that a modal formula φ is fixed and all notions defined below (a tight structure, a special clique, ...) depend on φ and are defined with respect to φ .

We say that a (sub)structure is a *clique* if it is a single strongly-connected component, i.e., for all two different worlds w, v we have $wRv \wedge vRw$. Note that a single world is a clique, regardless of whether it is reflexive; we will call such a world a *trivial clique*.

Let φ be a modal formula. We say that a structure is *tight* if the size of its cliques is bounded by $|\varphi|$. We omit explicit reference to φ when it is clear from the context.

► **Lemma 6.** *Let Φ be a \forall_T formula, \mathfrak{M} be a structure satisfying Φ and φ be a modal formula. Then there is a tight (w.r.t. φ) substructure \mathfrak{N} of \mathfrak{M} such that for each $v \in N$, $tp_{\mathfrak{M}}^{\varphi}(v) = tp_{\mathfrak{N}}^{\varphi}(v)$.*

Proof. In each clique C in \mathfrak{M} with size greater than $|\varphi|$, we do the following. For each subformula ψ of φ satisfied in some world of C , we mark one world satisfying ψ in C . We remove all worlds from C that are not marked. A straightforward induction w.r.t. the size of subformulae shows that the types (w.r.t. φ) of remaining worlds do not change. We do the operation described above independently for all cliques in the model. ◀

Let M be a model containing a world w . A formula ψ is a *reachable formula* of w if it is a subformula of φ and there is a world v satisfying ψ such that $wRv \wedge \neg vRw$. Let B_w denote the set of all reachable formulae of w . We extend this notation to cliques — for a clique C that contains a world w , we put $B_C = B_w$ (notice that B_C does not depend on the choice of w from C).

We say that a clique is *special* if some world w of this clique satisfies some subformula ψ of φ with $\psi \notin B_w$, and regular otherwise.

For a set B of subformulae of φ , we define (B, s) -zones in a structure \mathfrak{M} as follows. The $(B, 0)$ -zone is the substructure of \mathfrak{M} that consists of worlds from all regular cliques C such that $B_C = B$. The $(B, 1)$ -zone is the substructure of \mathfrak{M} that consists of worlds from all special cliques C such that $B_C = B$. We say that a (B, s) -zone is *special* if $s = 1$ and *regular* otherwise. Clearly, each world is in exactly one zone and the number of (B, s) -zones is bounded by $O(2^{|\varphi|})$.

We define a partial order \preceq on zones as follows. We say that (B, s) -zone $\preceq (B', s')$ -zone if $B \subset B'$ or $B = B'$ and $s \leq s'$.

► **Lemma 7.** *If \mathfrak{M} is transitive, then for any B, B', s, s' , world w from a (B, s) -zone and w' from a (B', s') -zone, if wRw' , then $(B', s') \preceq (B, s)$.*

Proof. Suppose that wRw' . Of course, $B_{w'} \subseteq B_w$. Assume that $(B', s') \not\preceq (B, s)$, i.e., $B_{w'} = B_w$, $s = 0$ and $s' = 1$. Then, w' is a special (trivial) clique and w is not, so there is no edge from w' to w . As w' is special, it satisfies some $\psi \notin B$. But, since wRw' , $\psi \in B_w = B$ — a contradiction. ◀

We say that a frame is a *line* if its reflexive closure is isomorphic with $\langle \mathbb{N}, \leq \rangle$. We extend this notation to structures, saying that a (sub)structure is a *line* if its frame is a line. We jointly call cliques and lines *units*. For a line l , we denote by l^i the *i th world* of l , i.e., the world that have i proper predecessors in l .

Our aim is to show that all modal formulae satisfiable over some class definable by a \forall_T formula Φ have models with a bounded number of units.

We say that a structure \mathfrak{M} is *disentangled* if each special zone of \mathfrak{M} consist of not connected cliques and each regular zone of \mathfrak{M} consists of lines l_1, l_2, \dots such that if there is an edge from a world in l_i to a world in l_j , then $i \leq j$.

► **Example 8.** Consider structures presented in Fig. 1 and a modal formula ν as a reference formula. All worlds satisfying r are special (w.r.t. ν) — none of their successors satisfies r . All other worlds are regular (w.r.t. ν).

Consider any of the structures presented in Fig. 1. Let $\mathfrak{p}, \mathfrak{q}$ and \mathfrak{r} be arbitrary worlds of that structure labelled by p, q and r , respectively. Observe that the structure consists of three zones: the $(B_{\mathfrak{p}}, 0)$ -zone that contains all elements satisfying p , the $(B_{\mathfrak{r}}, 1)$ -zone that contains elements satisfying r and the $B_{\mathfrak{q}}$ -zone that contains all elements satisfying q . Note that $B_{\mathfrak{q}} \subseteq B_{\mathfrak{r}} \subseteq B_{\mathfrak{p}}$ — if there is an edge from one zone to another, then by transitivity the first one contains all the successors of some world in the second one. The second inclusion is strict — case in point, $r \in B_{\mathfrak{p}} \setminus B_{\mathfrak{r}}$. But the first one is not — it is not hard to see that $B_{\mathfrak{r}} = B_{\mathfrak{q}}$.

The model b) is disentangled, and each of its zones contains precisely a single unit — a line or a trivial clique (consisting of a single world). The model a) is not disentangled, but it contains a disentangled submodel — it is enough to remove the dashed edge.

► **Lemma 9.** *Let Φ be a \forall_T formula, φ be a modal formula and \mathfrak{M} be a tight (w.r.t. φ) structure satisfying Φ and φ . Then there is a tight (w.r.t. φ) disentangled substructure \mathfrak{N} of \mathfrak{M} such that for each $v \in N$, $tp_{\mathfrak{M}}^{\varphi}(v) = tp_{\mathfrak{N}}^{\varphi}(v)$.*

Proof. We show that units in a $(B, 1)$ -zone in \mathfrak{M} are not connected. Assume that for some worlds w, v from two different units of this zone we have wRv . World v is in a special clique, therefore there is a world v' in this clique and a subformula $\psi \notin B$ of φ satisfied by v' . Since $\neg v'Rw$ (otherwise v', w would be in the same clique) and wRv' it follows that $\psi \in B_w = B$. This is a contradiction. The structure \mathfrak{N} therefore contains all the $(B, 1)$ -zones of \mathfrak{M} .

Now we consider regular zones; we transform a $(B, 0)$ -zone to a substructure that consists of ordered lines. First, we remove all points with only finite numbers of successors. Such worlds are regular, as they are in a regular zone, therefore Lemma 3 guarantees that the types of the remaining worlds are not changed.

Let w_1, w_2, \dots be an enumeration of those worlds in the $(B, 0)$ -zone such that for all worlds v , if v has a successor in the $(B, 0)$ -zone then it has a successor among w_1, w_2, \dots , and no two worlds among w_1, w_2, \dots are in the same clique. Below we inductively define a sequence l_1, l_2, \dots such that each l_i is a line or it is empty.

For each i , if w_i has a successor in some l_j with $j < i$, then we leave l_i empty. Otherwise, let l_i be an arbitrary path starting in w_i that does not contain non-trivial cliques, i.e., cliques of consisting of more than one world.

We remove all the worlds that are not in any l_j . Clearly, there are no edges from l_i to l_j for any $i > j$. We add the remaining worlds to \mathfrak{N} — these worlds form the $(B, 0)$ -zone of \mathfrak{N} . To obtain a sequence matching the requirements of the definition of disentangled structure, we simply remove empty elements from l_1, l_2, \dots . ◀

► **Proposition 10.** *For any \forall_T formula Φ , \mathcal{K}_{Φ} -GSAT and \mathcal{K}_{Φ} -GFINSAT are in NP.*

Proof. We prove that each satisfiable modal formula has a model which is a single unit. Clearly, the existence of such a model can be verified in NP.

Assume that φ has a model satisfying Φ . By Lemma 9 there is a disentangled model \mathfrak{M} of φ and Φ . Let a (B, s) -zone be a maximal nonempty zone w.r.t. \preceq . It is not hard to see (Lemma 7) that there are no edges from the (B, s) -zone to another zone. Let \mathfrak{N} be any unit of this zone. Of course, \mathfrak{N} satisfies Φ as it is a submodel of \mathfrak{M} . Now we define a single-unit structure \mathfrak{N}' that satisfies φ .

If $s = 1$, then we simply put $\mathfrak{N}' = \mathfrak{N}$. By the definition of disentangled models, there are no edges from \mathfrak{N} to different units in the same zone. Moreover, by choice of the (B, s) -zone

there are no edges in \mathfrak{M} from any world from \mathfrak{N} to any different zone. Therefore, each world in \mathfrak{N} has the same successors in \mathfrak{N} as in \mathfrak{M} , and therefore its type is the same in both models.

If $s = 0$, then we define a line \mathfrak{N}' over a universe $\underline{0}, \underline{1}, \dots$ containing fresh worlds as follows. Let the frame \mathcal{N}' be isomorphic with \mathcal{N} and let $\{t_0, t_1, \dots, t_{k-1}\}$ be a set of labels of worlds of the (B, s) -zone (restricted to the propositional variables of ϕ). For every $i \geq 0$, we label a world \underline{i} as $t_{i \bmod k}$. A quick check shows that the types of worlds in \mathfrak{N}' are the same as the types of worlds in the (B, s) -zone.

This proof can be simply adjusted to the finite global satisfiability — in any finite model, we can find a clique without successors. \blacktriangleleft

4.2 Local satisfiability

First, we define a property of sets of worlds called *homogeneity*. Roughly speaking, all the worlds from a homogeneous set have the same set of predecessors and successors outside of this set.

A world v is called a *proper predecessor (successor)* of w w.r.t. W if v is a predecessor (successor, resp.) of w and $v \notin W$.

A set of worlds W is (P, S) -*homogeneous* (in \mathfrak{M}) if for every world $w \in W$ the set of proper successors of w w.r.t. W is equal to S and the set of proper predecessors of w w.r.t. W is equal to P .

A structure \mathfrak{M}' is the (P, S) -*homogenization* of W in a structure \mathfrak{M} if \mathfrak{M}' is a structure with the same universe as \mathfrak{M} such that W is (P, S) -homogeneous in \mathfrak{M}' and \mathfrak{M}' is obtained by changing the proper predecessors and the proper successors (w.r.t W) of worlds from W to P and S respectively. More precisely, there is an edge from w to v in \mathfrak{M}' if and only if one of the following holds.

- w and v are both in W or both outside W and there is an edge from w to v in \mathfrak{M} .
- w is in P and v is in W or it is a successor of some world of W .
- v is in Q and w is in W or it is a predecessor of some world of W .

A set W is *homogeneous* (in \mathfrak{M}) if there are P, S such that W is (P, S) -homogeneous (in \mathfrak{M}). Figure 2 contains an example of a homogenization. Observe that for all P, S, W and \mathfrak{M} there is a unique a (P, S) -homogenization of W in \mathfrak{M} .

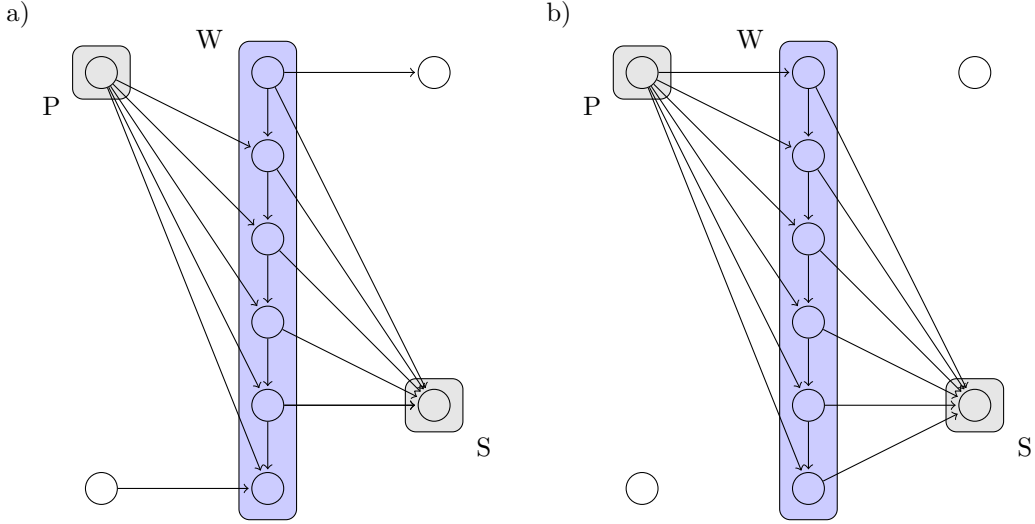
A set of worlds W is a *witness for a set of worlds V* if for each world $v \in V$ and each subformula $\diamond\psi$ of φ , if v has a proper successor w.r.t. V satisfying ψ , then it has one in W . Note that if V is a finite set of worlds within one unit, there is a witness for V of size bounded by $|\varphi|$.

Our aim is to prove the following lemma.

► **Lemma 11.** *Let Φ be a \forall_T formula and φ be a modal formula satisfiable over \mathcal{K}_Φ . Then there is a tight (w.r.t. φ) disentangled model of φ satisfying Φ s.t. the number of units is bounded exponentially in the size of φ and all its units are homogeneous.*

The first step towards the proof of Lemma 11 is enforcing an additional restriction on lines. Namely, we can assume, without loss of generality, that every line consists of either only reflexive worlds or only irreflexive worlds.

► **Lemma 12.** *Let Φ be a \forall_T formula and φ be a modal formula satisfiable over \mathcal{K}_Φ . Then there exists a tight (w.r.t. φ) disentangled model of φ satisfying Φ such that every line is isomorphic with $\langle \mathbb{N}, \leq \rangle$ or $\langle \mathbb{N}, < \rangle$.*



■ **Figure 2** The structure on the picture b) is the result of the (P, S) -homogenization of W in the structure on the picture a). Some edges that follow from the transitivity are omitted for better readability.

Proof. If φ is a modal formula satisfiable over \mathcal{K}_Φ , there exists a tight disentangled model \mathfrak{M} of φ satisfying Φ .

Let l_1, l_2, \dots be an enumeration of lines in \mathfrak{M} . We show that for every $i > 0$, there is a model \mathfrak{M}_i of φ satisfying Φ such that for every $j \in \{1, \dots, i-1\}$, $\mathcal{M}_i \cap l_j$ is isomorphic with $\langle \mathbb{N}, \leq \rangle$ or $\langle \mathbb{N}, < \rangle$.

Let l_i be a line in \mathfrak{M}_i . If l_i contains finitely many reflexive worlds, they can be removed. Let \mathfrak{M}_{i+1} be the resulting structure. Due to Lemma 3, \mathfrak{M}_{i+1} satisfies φ and, as a substructure of \mathfrak{M}_i , it satisfies Φ . Clearly, the line $\mathcal{M}_{i+1} \cap l_i$ is isomorphic with $\langle \mathbb{N}, < \rangle$.

Otherwise, l_i contains infinitely many reflexive worlds. Let \mathfrak{N} be a substructure of \mathfrak{M}_i resulting from removing all irreflexive worlds from l_i . Clearly, $\mathcal{N} \cap l_i$ is isomorphic with $\langle \mathbb{N}, \leq \rangle$. The structure \mathfrak{N} satisfies Φ , but it may violate φ .

Therefore, we define \mathfrak{M}_{i+1} as a structure defined on the frame \mathcal{N} such that the labelling of \mathfrak{M}_{i+1} coincides with the labelling of \mathfrak{M}_i on worlds the do not belong to l_i , i.e., $w \in M_{i+1} \setminus l_i$. Now, on worlds $w \in M_{i+1} \cap l_i$ the labelling is defined in such a way that every labelling that occurs on l_i in \mathfrak{M}_{i+1} , occurs infinitely often in $M_{i+1} \cap l_i$ in \mathfrak{M}_{i+1} . It is easy to verify that the types for all proper predecessors of $M_{i+1} \cap l_i$ in \mathfrak{M}_{i+1} remain the same as in \mathfrak{M}_i . ◀

Before we show Lemma 11, we provide an outline. Define $n = |\varphi|$ and $N = |\Phi|$. We start from an arbitrary tight disentangled model \mathfrak{M} such that every line is isomorphic with $\langle \mathbb{N}, \leq \rangle$ or $\langle \mathbb{N}, < \rangle$ and modify it in three steps.

Step 1. Select recursively worlds from \mathfrak{M} . Initially, select an arbitrary element a of \mathfrak{M} such that $\mathfrak{M}, a \models \varphi$. Then, recursively, for every world w selected in the previous stage do the following. If w is in a non-trivial clique, select a witness for this clique of size bounded by $|\varphi|$.

If w is in a line l , select from l the world l^k with the following property: for every selected world v preceding l , if v has a successor in l , then l^k is a successor of v as well. (There are finitely many selected v preceding l , thus such l^k exists.) Then, remove all worlds from l preceding w' , and select worlds $l^k, l^{k+1}, \dots, l^{k+N \cdot n^{2-n}}$. Finally, choose for them a witness

W and select worlds from $W \setminus l$.

Let \mathfrak{M}_1 be the structure obtained in this step. We show that \mathfrak{M}_1 satisfies Φ and φ and at most $n^{2 \cdot n}$ units contain selected worlds.

Step 2. Remove from \mathfrak{M}_1 those units that do not contain selected worlds, obtaining the structure \mathfrak{M}_2 . Clearly, \mathfrak{M}_2 satisfies Φ as it is a substructure of \mathfrak{M} . However, it may happen that the modal formula φ is not satisfied in \mathfrak{M}_2 .

Consider a model presented at Fig. 1 a) and the unit that consists of all worlds satisfying p . If we select only finitely many units satisfying r and remove all other such units, then the obtained structure would not satisfy the ν . We find a witness for the set of first $N \cdot n^{2 \cdot n}$ worlds from the line and select all its worlds. Later we will show how we can use them as witnesses for all other worlds from this line.

Step 3. In each line l in \mathfrak{M}_2 , starting from the greatest zone w.r.t. \preceq , we find a homogeneous set of N worlds that is before the world $l^{N \cdot n^{2 \cdot n}}$. Since the number of all units in \mathfrak{M}_2 is bounded by $n^{2 \cdot n}$, such a sequence can be easily found. Then, we use this sequence to obtain a line such that all worlds in this line have the same set of successors as the worlds in the sequence and that transformation does not violate Φ (this trick will be presented as Lemma 13).

Let \mathfrak{M}_3 be the structure obtained by applying this procedure to all lines. We conclude by an inductive proof that shows that \mathfrak{M}_3 satisfies φ .

Let us discuss the properties of \mathfrak{M}_1 defined in Step 1. At every stage of the construction, either the selected worlds are from the same unit or they belong to strictly greater zones w.r.t. \preceq . Since chains w.r.t. \preceq are bounded by $2 \cdot n$, there is at most $2 \cdot n$ stages of the construction. Every unit that contains a selected world has at most n selected witnesses. Thus, there are at most $n^{2 \cdot n}$ units with selected worlds. The structure \mathfrak{M}_1 results from \mathfrak{M} by removing finitely many regular worlds, therefore by Lemma 3, for every $v \in \mathfrak{M}_1$, $tp_{\mathfrak{M}_1}(v) = tp_{\mathfrak{M}}(v)$. In particular, \mathfrak{M}_1 satisfies φ . As a substructure of \mathfrak{M} , \mathfrak{M}_1 satisfies Φ .

Observe that \mathfrak{M}_1 satisfies the following property: if a selected world v has a successor in a unit U with a selected world, all worlds of U are successors of v . This is clear if U is a clique, and we remove from lines worlds that violate this. Every unit in \mathfrak{M}_2 contains a selected world, therefore \mathfrak{M}_2 satisfies the following property:

- (1) if a selected world v has a successor in a unit U , all worlds from U are successors of v .

Now, we elaborate on Step 3, i.e., we show a construction of \mathfrak{M}_3 from \mathfrak{M}_2 and we show that it satisfies Φ . We start with the following lemma stating that careful homogenization of a structure satisfying Φ satisfies Φ as well.

► **Lemma 13.** *Let Φ be a \forall_T formula, $N = |\Phi|$, \mathfrak{M} be a structure over \mathcal{K}_Φ , l be a line in \mathfrak{M} isomorphic with $\langle \mathbb{N}, \leq \rangle$ or $\langle \mathbb{N}, < \rangle$. If there is k such that $\{l^k, l^{k+1}, \dots, l^{k+N-1}\}$ is (P, S) -homogeneous, then the $(P \setminus l, S \setminus l)$ -homogenization of l in \mathfrak{M} satisfies Φ .*

Proof. Let \mathfrak{M}' be the $(P \setminus l, S \setminus l)$ -homogenization of l in \mathfrak{M} . Let us assume that l is isomorphic with $\langle \mathbb{N}, < \rangle$. The case where l is isomorphic with $\langle \mathbb{N}, \leq \rangle$ is similar. We show that \mathfrak{M}' satisfies Φ . To avoid confusion with modal logic, we show that \mathcal{M}' , which is a relational structure, satisfies the first order formula Φ .

The formula Φ is a conjunction of the formula $(\forall xyz. xRy \wedge yRz \Rightarrow xRz)$ and a universal first-order formula. Thus, it can be transformed to an equivalent formula of the form $\forall \vec{x}. \Psi(\vec{x})$, where $\Psi(\vec{x})$ is quantifier-free and $|\vec{x}| \leq N$. Suppose that \mathfrak{M}' does not satisfy Φ and let $\vec{u} = \langle u_0, \dots, u_{|\vec{x}|-1} \rangle$ be such that \mathcal{M}' violate $\Psi(\vec{u})$. We can rearrange \vec{u} such that

$\vec{u} = \vec{u}_1 \cup \vec{u}_2$, \vec{u}_1 are from l , \vec{u}_2 are disjoint with l , and \vec{u}_1 are ordered w.r.t. the successors relation of \mathfrak{M}' . Let $f : \vec{u} \mapsto \mathcal{M}'$ be defined as

$$f(u_i) = \begin{cases} l^{k+i} & \text{if } u_i \in \vec{u}_1 \\ u_i & \text{otherwise} \end{cases}$$

Let $f(\vec{u})$ denote image of \vec{u} under f , i.e., $f(\vec{u}) = \langle f(u_0), \dots, f(u_{|\vec{u}|-1}) \rangle$. We show that

- (i) for all $v, v' \in \vec{u}$, \mathcal{M}' satisfies vRv' iff \mathcal{M}' satisfies $f(v)Rf(v')$,
- (ii) for all $v, v' \in \vec{u}$, \mathcal{M}' satisfies $f(v)Rf(v')$ iff \mathcal{M} satisfies $f(v)Rf(v')$.

It follows that if \mathcal{M}' violates $\Psi(\vec{u})$, then \mathcal{M} violates $\Psi(f(\vec{u}))$. Thus, if \mathcal{M} satisfies Φ , then \mathcal{M}' satisfies Φ .

(i) : Clearly, (i) holds when v, v' are both from \vec{u}_1 or both from \vec{u}_2 . If $v \in \vec{u}_1$ and $v' \in \vec{u}_2$, then $v \in l$ and $v' \notin l$. ($P \setminus l, S \setminus l$)-homogeneity implies that \mathcal{M}' satisfies vRv' iff \mathcal{M}' satisfies $f(v)Rv'$. Since $v' \notin l$, $v' = f(v')$. Therefore, \mathcal{M}' satisfies vRv' iff \mathcal{M}' satisfies $f(v)Rf(v')$. The case $v \in \vec{u}_2$ and $v' \in \vec{u}_1$ is similar.

(ii) : ($P \setminus l, S \setminus l$)-homogenization only changes successors and predecessors of worlds from l . This implies that (ii) holds if v, v' are both from \vec{u}_1 or both from \vec{u}_2 .

If $v \in \vec{u}_1$ and $v' \in \vec{u}_2$, then $f(v) \in \{l^k, \dots, l^{k+N-1}\}$ and $v' \notin l$. Thus, \mathcal{M}' satisfies $f(v)Rf(v')$ iff $v' \in S \setminus l$ iff \mathcal{M} satisfies $f(v)Rf(v')$. Similarly, if $v \in \vec{u}_2$ and $v' \in \vec{u}_1$, then $f(v') \in \{l^k, \dots, l^{k+N-1}\}$ and \mathcal{M}' satisfies $f(v)Rf(v')$ iff $v \in P \setminus l$ iff \mathcal{M} satisfies $f(v)Rf(v')$. \blacktriangleleft

Let \sqsubset be a linear order relation on units of \mathfrak{M}_2 compatible with the reversed \preceq , i.e., if U_1 is from a greater zone than U_2 , then $U_1 \sqsubset U_2$. Since \mathfrak{M}_2 has finitely many units, \sqsubset is well-founded. We define structures \mathfrak{N}_U , where U ranges over units of \mathfrak{M}_2 . The structures \mathfrak{N}_U will be defined inductively w.r.t. the order \sqsubset . They will satisfy the following induction assumption: \mathfrak{N}_U satisfies Φ and for every unit U' such that $U' \sqsubseteq U$, U' is homogeneous.

For the induction base we introduce an empty unit \emptyset and we state that it is the least unit w.r.t. \sqsubset . Clearly, $\mathfrak{N}_\emptyset = \mathfrak{M}_2$ satisfies all the assumptions.

For the induction step assume that \mathfrak{N}_V has been defined, it satisfies Φ and every unit $U' \sqsubseteq V$ is homogenous in \mathfrak{N}_V . Let U be the least unit greatest than V (w.r.t. \sqsubset). Since cliques are homogeneous, $\mathfrak{N}_U = \mathfrak{N}_V$ satisfies all the assumptions.

Assume that U is a line l . (1) implies the following dichotomy for selected worlds: for every selected world w and every unit U' either all worlds from U' are successors of w or no world from U' is a successor of w . Therefore, a subset of units determines the set of successors of a given selected world. Due to transitivity, for all w, w' ; if w' is a successor of w , the set of successors of w' is a subset of the set of successors of w . Since there are $n^{2 \cdot n}$ different units, among $N \cdot n^{2 \cdot n}$ consecutive selected worlds there are N consecutive worlds \vec{u} that have the same set of successors outside the line l . Observe (1) implies that all selected worlds in l have the same predecessors. Therefore, \vec{u} have the same sets of predecessors and successors outside l . As they are consecutive on l , predecessors of u_0 are predecessors of \vec{u} and successors of u_{N-1} are successors of \vec{u} . Hence, there are P, S such that \vec{u} is a (P, S) -homogeneous set. Let \mathfrak{N}_U be $(P \setminus l, S \setminus l)$ -homogenization of l in \mathfrak{N}_V . By Lemma 13, \mathfrak{N}_U satisfies Φ . All units homogeneous in \mathfrak{N}_V remain homogeneous in \mathfrak{N}_U and $U = l$ is homogeneous.

We have shown that the construction given in Steps 1, 2 and 3 can be executed and that \mathfrak{M}_3 satisfies all the postulated conditions from Lemma 11 but one. It remains to be verified that \mathfrak{M}_3 satisfies φ .

► **Lemma 14.** *The structure \mathfrak{M}_3 satisfies φ .*

Proof. We show that worlds from \mathfrak{M}_3 have the same types in \mathfrak{M}_3 as in \mathfrak{M}_1 , i.e., for each $v \in M_3$, $tp_{\mathfrak{M}_1}(v) = tp_{\mathfrak{M}_3}(v)$. Observe that \mathfrak{M}_3 contains the root of \mathfrak{M}_1 , thus it satisfies φ . To show that the types do not change, we show the following property:

(2) for every $v \in \mathfrak{M}_3$, $\Diamond\psi \in tp_{\mathfrak{M}_1}(v)$ iff v has a successor u in \mathfrak{M}_3 such that $\psi \in tp_{\mathfrak{M}_1}(u)$.

We show (2) in two steps. First, observe that

(3) for every selected $v \in \mathfrak{M}_2$, $\Diamond\psi \in tp_{\mathfrak{M}_1}(v)$ iff v has a successor u in \mathfrak{M}_2 s.t. $\psi \in tp_{\mathfrak{M}_1}(u)$.

Clearly, \mathfrak{M}_2 is a substructure of \mathfrak{M}_1 , therefore if v has a successor u in \mathfrak{M}_2 such that $\psi \in tp_{\mathfrak{M}_1}(u)$, then $\Diamond\psi \in tp_{\mathfrak{M}_1}(v)$.

Conversely, assume that $\Diamond\psi \in tp_{\mathfrak{M}_1}(v)$. We need to show that v has a successor $u \in \mathfrak{M}_1$ such that $\psi \in tp_{\mathfrak{M}_1}(u)$ and u is in a unit that contains a selected world from \mathfrak{M}_1 . Clearly, such u belongs to \mathfrak{M}_2 , thus (3) holds.

Recall that, for every $w \in \mathfrak{M}_1$, $tp_{\mathfrak{M}}(w) = tp_{\mathfrak{M}_1}(w)$. Let v' be the last selected successor (in \mathfrak{M}) of v that satisfies $\Diamond\psi$ (it can be v itself). The world v' has a selected successor u satisfying ψ . If u belongs to \mathfrak{M}_1 , we are done. We show that $u \notin M_1$ is impossible. Suppose that u belongs to a line l in \mathfrak{M} , but it has been removed at Step 1. Then, due to regularity of l , there is a successor u' of u in l , which is selected and belongs to \mathfrak{M}_1 . Since every world on l is regular, u' has a successor in l satisfying ψ . Thus, u' is a selected strict successor of v' and it satisfies $\Diamond\psi$, which contradicts the assumption that v' is the last selected successor of v that satisfies $\Diamond\psi$.

Observe that (3) implies (2): If v belongs to a clique, then it has exactly the same successors in \mathfrak{M}_3 as in \mathfrak{M}_2 , thus (3) implies (2).

Assume that v belongs to a line l . Let w be a selected world from the same line l which belongs to a chosen homogeneous set. Since w has the same successors in \mathfrak{M}_2 and \mathfrak{M}_3 , (3) implies (2) for w , i.e., $\Diamond\psi \in tp_{\mathfrak{M}_1}(w)$ iff w has a successor u' in \mathfrak{M}_3 such that $\psi \in tp_{\mathfrak{M}_1}(u')$.

All worlds in l have the same set of reachable formulae, therefore $\Diamond\psi \in tp_{\mathfrak{M}_1}(v)$ iff $\Diamond\psi \in tp_{\mathfrak{M}_1}(w)$. Since all worlds in l share the set B , the set of reachable formulae, alike in \mathfrak{M}_1 as in \mathfrak{M}_3 , v has a successor u in l satisfying $\psi \in tp_{\mathfrak{M}_1}(u)$ iff w has a successor u' in l satisfying $\psi \in tp_{\mathfrak{M}_1}(u')$. Also, l is homogeneous in \mathfrak{M}_3 , therefore v and w have the same successors outside l . Thus, w has a successor u' in \mathfrak{M}_3 such that $\psi \in tp_{\mathfrak{M}_1}(u')$ iff u has a successor u'' in \mathfrak{M}_3 such that $\psi \in tp_{\mathfrak{M}_1}(u'')$. This implies (2).

Finally, we show that the property (2) implies that for each $v \in M_3$, $tp_{\mathfrak{M}_1}(v) = tp_{\mathfrak{M}_3}(v)$.

The *modal depth*, denoted by $MD(\varphi)$, of a modal formula φ is defined recursively:

- $MD(p) = 0$ where p is a propositional variable or a logical constant (\top , \perp).
- $MD(\neg\varphi) = MD(\varphi)$.
- $MD(\varphi_1 \wedge \varphi_2) = \max(MD(\varphi_1), MD(\varphi_2))$.
- $MD(\Diamond\varphi) = MD(\varphi) + 1$.

We show by induction on n that for every $v \in M_3$, for every subformula ψ of φ with $MD(\psi) \leq n$, $\psi \in tp_{\mathfrak{M}_1}(v)$ iff $\psi \in tp_{\mathfrak{M}_3}(v)$.

Induction base. The labelling π is not changed at any stage of construction, therefore the types $tp_{\mathfrak{M}_1}(v)$ and $tp_{\mathfrak{M}_2}(v)$ agree on formulae of modal depth zero.

Induction step. Assume that for every $v \in \mathfrak{M}_3$, $tp_{\mathfrak{M}_1}(v)$ and $tp_{\mathfrak{M}_3}(v)$ agree on formulae of modal depth n . Let $w \in \mathfrak{M}_3$. Consider a subformula $\Diamond\psi$ of φ , where modal depth of ψ is n . By (2), $\Diamond\psi \in tp_{\mathfrak{M}_1}(w)$ iff w has a successor u in \mathfrak{M}_3 , such that $\psi \in tp_{\mathfrak{M}_1}(u)$. Since ψ has modal depth n , the induction assumption implies that $\psi \in tp_{\mathfrak{M}_1}(u)$ iff $\psi \in tp_{\mathfrak{M}_3}(u)$. Therefore, $\Diamond\psi \in tp_{\mathfrak{M}_1}(w)$ iff $\Diamond\psi \in tp_{\mathfrak{M}_3}(w)$.

Observe that if $tp_{\mathfrak{M}_1}(w)$ and $tp_{\mathfrak{M}_3}(w)$ agree on formulae of modal depth $n + 1$ of the form $\diamond\psi$, then they agree on all formulae of modal depth $n + 1$. ◀

► **Proposition 15.** *Let Φ be a \forall_T formula. Then \mathcal{K}_Φ -SAT is decidable in NEXPTIME.*

Proof. For a given \forall_T formula Φ , we describe an algorithm that checks whether a given modal formula φ has a model that satisfies Φ , φ and consists of exponentially many homogeneous units. Lemma 11 implies that such an algorithm is complete. The algorithm works as follows:

Guess a finite description of a structure \mathfrak{N} . Observe that if every unit of \mathfrak{N} is homogeneous, then for all units U_1, U_2 , either all worlds from U_1 are successors (predecessors) of all worlds from U_2 or no world from U_1 is a successor (predecessor) of any world from U_2 . Therefore, if \mathfrak{N} has at most M units, it has a description of size $M \cdot (M \cdot 2^{|\varphi|})$. For every unit U it needs to be stated whether U is a line or a clique, which units are successors of U and which types occur in $\{tp_{\mathfrak{N}}(u) : u \in U\}$. The order of types in a line is irrelevant and we can assume that each type occurs infinitely often.

Verify the modal formula. It suffices to check whether every type in every unit has appropriate witnesses and if a formula $\neg\diamond\psi$ belongs to this type, the formula ψ does not belong to any type in the current unit or successor units.

Verify the first-order formula. The first order formula is universal, i.e., $\Phi = \forall\vec{x}.\Psi(\vec{x})$, where Ψ is quantifier-free. Therefore, it has to be checked whether every instantiation of \vec{x} in \mathfrak{N} satisfies Ψ . However, observe that for all \vec{a}, \vec{b} satisfying (a) for every $i \in [1, N]$, a_i is in the same unit as b_i , and (b) for all a_i, a_j from the same unit, a_i is a successor of a_j iff b_i is a successor of b_j , we have \mathcal{M}_3 (the frame of \mathfrak{M}_3) satisfies $\Psi(\vec{a})$ iff \mathcal{M}_3 satisfies $\Psi(\vec{b})$.

It follows that we have to check at most finitely many different types of tuples \vec{a} . Each type of a tuple is determined by the assignment of its elements to the units of \mathfrak{N} (cond. (a)) and relative order of elements in the same unit (cond. (b)). Notice that these two characteristics determine relations among worlds. Hence, it can be easily verified whether tuples of a given type satisfy Ψ .

The number of types of tuples \vec{a} that have to be checked is bounded by $c^N(2N)!$, where c is the number of units. Since N is a parameter of the problem and c is of exponential order, they can be checked in exponential time in $|\varphi|$. Even if Φ was a part of an instance, N is linear in the size of the input and $c^N(2N)!$ is still exponential in the input. ◀

5 Conclusion and future work

We proved that all elementary modal logics over universally defined classes of transitive structures are decidable. In case of the global satisfiability problem, we proved that all satisfiability problems and finite satisfiability problems are in NP and, by a straightforward reduction from SAT, we can conclude that these problems are NP-complete. The case of the local satisfiability problem is more complicated. We proved that all the satisfiability problems and all the finite satisfiability problems are in NEXPTIME, but the precise complexity may vary. For example, the logic K4 is PSPACE-complete, while S5 is NP-complete. Providing the full characterisation with respect to complexity is left as an open question.

Modal logic over a class of transitive frames can be seen as a temporal logic. Indeed, the logic K4 may be defined as a syntactic variant of a fragment of CTL that allows only two modalities — AG and EF. Therefore, Theorem 1 may be treated as a first step in study of *elementary temporal logics*. A natural question that arises is whether the decidability results can be extended for different CTL operators, such as EG and AF.

References

- 1 Vince Bárány and Mikolaj Bojanczyk. Finite satisfiability for guarded fixpoint logic. *Inf. Process. Lett.*, 112(10):371–375, 2012.
- 2 Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Comp. Sc.* Cambridge University Press, Cambridge, 2001.
- 3 Davide Bresolin, Dario Della Monica, Valentin Goranko, Angelo Montanari, and Guido Sciavicco. The dark side of interval temporal logic: Sharpening the undecidability border. In *TIME*, pages 131–138. IEEE, 2011.
- 4 R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- 5 Kit Fine. Logics containing K4. Part II. *J. of Symbolic Logic* 50, pages 619–651, 1985.
- 6 Igor Gorbunov. A decidable modal logic that is finitely undecidable. In *Advances in Modal Logic*, pages 247–258. College Publications, 2006.
- 7 Erich Grädel. On the restraining power of guards. *J. Symbolic Logic*, 64:1719–1742, 1999.
- 8 Erich Grädel and Igor Walukiewicz. Guarded fixed point logic. In *Fourteenth Annual IEEE Symposium on Logic in Computer Science*, pages 45–54, 1999.
- 9 Edith Hemaspaandra and Henning Schnoor. A universally defined undecidable unimodal logic. In *MFCS*, volume 6907 of *LNCS*, pages 364–375. Springer, 2011.
- 10 Wilfrid Hodges. *Model Theory*. Cambridge University Press, 1993.
- 11 Emanuel Kieroński and Jakub Michaliszyn. Two-variable universal logic with transitive closure. In *Proceedings of Computer Science Logic 2012*, 2012.
- 12 Emanuel Kieroński, Jakub Michaliszyn, and Jan Otop. Modal logics definable by universal three-variable formulas. In *FSTTCS*, volume 13 of *LIPICs*, pages 264–275, 2011.
- 13 Emanuel Kieroński, Jakub Michaliszyn, Ian Pratt-Hartmann, and Lidia Tendera. Extending two-variable first-order logic with equivalence closure. In *LICS '12: Proceedings of the 29th IEEE symposium on Logic in Computer Science*, 2012.
- 14 Emanuel Kieroński and Lidia Tendera. On finite satisfiability of the guarded fragment with equivalence or transitive guards. In *LPAR*, volume 4790 of *LNCS*, pages 318–332. Springer, 2007.
- 15 Marcus Kracht. *Tools and Techniques in Modal Logic*. Studies in Logic and the Foundations of Mathematics. Elsevier, Amsterdam, 1999.
- 16 Jerzy Marcinkowski and Jakub Michaliszyn. The ultimate undecidability result for the Halpern-Shoham logic. In *LICS*, pages 377–386. IEEE Computer Society, 2011.
- 17 Jakub Michaliszyn. Decidability of the guarded fragment with the transitive closure. In *ICALP (2)*, volume 5556 of *LNCS*, pages 261–272. Springer, 2009.
- 18 Jakub Michaliszyn and Jan Otop. Decidable elementary modal logics. In *LICS '12: Proceedings of the 29th IEEE symposium on Logic in Computer Science*, 2012.
- 19 Jakub Michaliszyn, Jan Otop, and Piotr Witkowski. Satisfiability vs. finite satisfiability in elementary modal logics. In *Proceedings of GandALF 2012*, pages 141–154, 2012.
- 20 Michael Mortimer. On languages with two variables. *Mathematical Logic Quarterly*, 21(1):135–140, 1975.
- 21 Martin Otto. Two variable first-order logic over ordered domains. *Journal of Symbolic Logic*, 66:685–702, 1998.
- 22 Ian Pratt-Hartmann. Complexity of the two-variable fragment with (binary-coded) counting quantifiers. *Journal of Logic, Language and Information*, 2005.
- 23 Wiesław Szwaś and Lidia Tendera. The guarded fragment with transitive guards. *Annals of Pure and Applied Logic*, 128:227–276, 2004.
- 24 Lidia Tendera and Wiesław Szwaś. FO^2 with one transitive relation is decidable. In *Proc. of STACS 2013*, 2013.