On the Decidability of Elementary Modal Logics

Jakub Michaliszyn, Imperial College London and University of Wrocław
Jan Otop, IST Austria and University of Wrocław
Emanuel Kieroński, University of Wrocław

We consider the satisfiability problem for modal logic over first-order definable classes of frames. We confirm the conjecture from [Hemaspaandra and Schnoor 2008] that modal logic is decidable over classes definable by universal Horn formulae. We provide a full classification of Horn formulae with respect to the complexity of the corresponding satisfiability problem. It turns out, that except for the trivial case of inconsistent formulae, local satisfiability is either NP-complete or PSPACE-complete, and global satisfiability is NP-complete, PSPACE-complete, or EXPTIME-complete. We also show that the finite satisfiability problem for modal logic over Horn definable classes of frames is decidable. On the negative side, we show undecidability of two related problems. First, we exhibit a simple universal three-variable formula defining the class of frames over which modal logic is undecidable. Second, we consider the satisfiability problem of bimodal logic over Horn definable classes of frames, and also present a formula leading to undecidability.

Categories and Subject Descriptors: F.4.1 [**Mathematical Logic**]: Modal logic

General Terms: Theory

Additional Key Words and Phrases: modal logic, elementary logics, computational complexity

ACM Reference Format:
Jakub Michaliszyn, Jan Otop, Emanuel Kieroński, 2013. On the Decidability of Elementary Modal Logics. *ACM Trans. Comput. Logic* 2, 6, Article 1 (August 2015), 45 pages.
DOI:http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

Modal logic for almost hundred years has been an important topic in many academic disciplines, including philosophy, mathematics, linguistics and computer science. Currently, it seems to be most intensively investigated by computer scientists. Among numerous branches in which modal logic, sometimes in disguise, finds applications are hardware and software verification, cryptography and knowledge representation.

Modal logic was introduced by philosophers to study modes of truth. The idea was to extend propositional logic by some new constructions, of which two most important were $\Diamond\varphi$ and $\Box\varphi$, originally read as φ *is possible*, and φ *is necessary*, respectively. A typical question was, given a set of axioms $\mathcal{A}$, corresponding usually to some intuitively acceptable aspects of truth, what is the logic defined by $\mathcal{A}$, i.e., which formulae are provable from $\mathcal{A}$ in a Hilbert-style system.

This paper is an extended version of three conference papers: [Kieroński et al. 2011], [Michaliszyn and Otop 2012], and [Michaliszyn and Kieroński 2012]. The first author was 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 European Research Council (ERC) under grant agreement 267989 (QUAREM) and by the Austrian Science Fund (FWF) NFN projects S11402-N23 (RiSE) and Z211-N23 (Wittgenstein Award). The third author was partially supported by Polish National Science Center grant DEC-2013/09/B/ST6/01535.

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

One of the most important steps in the history of modal logic was the invention of a formal semantics based on the notion of the so-called 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 often called *worlds*. For each world the 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*.

It turned out that there is a beautiful connection between syntactic and semantic approaches to modal logic [Sahlqvist 1975]: logics defined by natural axioms can be equivalently defined by imposing some simple restrictions on the classes of frames. E.g., the axiom $\Diamond\Diamond P \rightarrow \Diamond P$ (*if it is possible that* P *is possible, then* P *is possible*), defining modal logic K4, is valid precisely in the class of transitive frames; the axiom $P \rightarrow \Diamond P$ (*if* P *is true, then* P *is possible*), defining logic T — in the class of reflexive frames, $P \rightarrow \Box\Diamond P$ (*if* P *is true, then it is necessary that* P *is possible*), defining logic B — in the class of symmetric frames, and the axiom $\Diamond P \rightarrow \Box\Diamond P$ (*if* P *is possible, then it is necessary that* P *is possible*), defining logic K5 — in the class of Euclidean frames.

One may think that each modal formula φ defines a class of frames, namely the class of those frames in which φ is valid. Formally, a formula φ is valid in a frame $\mathcal{M}$ if for any possible truth-assignment of propositional variables to the worlds of $\mathcal{M}$, φ is true at every world. This definition has a second-order flavour, since it involves the quantification over sets of elements: for each variable P and each subset V of the set of worlds we have to consider the case in which P is true precisely in the worlds from V. Note however, that many important classes of frames, in particular all the classes we mentioned in the previous paragraph, can be defined by, often very simple, first-order formulae. For a given first-order sentence Φ over the signature consisting of a single binary symbol R we define $\mathcal{K}_\Phi$ to be the set of those frames that satisfy Φ. E.g., if $\Phi = \forall xyz(xRy \wedge yRz \Rightarrow xRz)$ then $\mathcal{K}_\Phi$ is the class of transitive frames. Classes of frames definable in this manner by first-order formulae are called *elementary classes of frames*. A modal logic of an elementary class of frames, i.e., the set of modal formulae valid in this class, is called an *elementary modal logic*.

It turns out that most everyday modal logics are elementary. Thus the decidability of elementary modal logics is an important and active research topic. In this paper we contribute to this area. Instead of validity we rather present our results in terms of satisfiability. Clearly, the question whether a modal formula φ is valid in a class of frames $\mathcal{K}$ is equivalent to the question whether $\neg\varphi$ is unsatisfiable over $\mathcal{K}$. Besides the standard, *local*, notion of modal satisfiability we also consider *global* satisfiability, i.e., the problem of deciding if a given formula is satisfied at every world of some Kripke structure.

As first-order logic is sufficiently strong to axiomatise grid structures, it is not surprising that there are elementary classes of frames over which the satisfiability problem for modal logic is undecidable. It turns out that this is true even if we restrict our attention to formulae in prenex normal form containing only universal quantifiers. The first result of this kind appears in [Hemaspaandra 1996]. A universal first-order formula with equality is exhibited defining the class of frames over which the global satisfiability problem is undecidable. This is improved in [Hemaspaandra and Schnoor 2011], where it is shown that there exists a universal formula without equality, such that even the local satisfiability problem over the class of frames defined by this formula is undecidable. The formula from [Hemaspaandra and Schnoor 2011] uses nine variables, and the proof is fairly complicated. A natural question arises, how many variables are necessary to obtain undecidability. Note that many natural classes of frames, including, classes of transitive, reflexive, symmetric, Euclidean, or equivalence frames are definable by formulae with at most three variables. The satisfiability problem for modal logic over those classes is known to be decidable [Ladner 1977]. It turns out however that there exist universal first-order formulae without equality with only three variables defining the classes of frames over which satisfiability problem for modal logic is undecidable. Exhibiting such formulae is the first contribution of our paper. We believe, that even though we put an effort to use as small number of variables as possible, our undecidability proof is simpler than the one from [Hemaspaandra and Schnoor 2011]. In particular, in the case of local satisfiability our formula is a single clause.

THEOREM 1.1. *There exist three-variable universal formulae* Γ, Γ' *without equality such that the global (finite) satisfiability problem for modal logic over* $\mathcal{K}_\Gamma$ *and the local (finite) satisfiability problem for modal logic over* $\mathcal{K}_{\Gamma'}$ *are undecidable.*

The above result is optimal with respect to the number of variables, since the well known *standard translation* of modal logic to first-order logic fits into the two-variable fragment, and thus satisfiability of modal logic over classes of frames defined by two-variable formulae reduces to satisfiability of the two-variable fragment, which is known to be decidable [Mortimer 1975]. Using the standard translation we can justify decidability of some elementary modal logics. Over the signature containing only the single binary symbol R, two-variable logic itself is probably too weak to express any interesting properties of frames. However, we could use some of its decidable extensions, e.g., by counting quantifiers [Pacholski et al. 1997; Graedel et al. 1997; Pratt-Hartmann 2005], or by the statement about transitivity of R [Szwast and Tendera 2013]. Another idea is to consider some orthogonal generalisations of the image of the standard translation of modal logic, like the guarded fragment [Grädel 1999], or the guarded negation fragment [Bárány et al. 2011].

Inspecting the formulae defining classes of frames corresponding to the best known decidable modal logics, including the mentioned logics T, B, K4, and K5, one easily observes that they share another natural syntactic restriction. Namely, all of them are universal Horn formulae, UHF. Such formulae were considered in [Hemaspaandra and Schnoor 2008], where a dichotomy result was proved, that the satisfiability problem for modal logic over the class of frames defined by an UHF formula (with an arbitrary number of variables) is either in NP or PSPACE-hard. In the same paper the authors proved decidability of a rich class of Horn definable modal logic, including all the cases in which the Horn formula implies reflexivity. However the question was left open, e.g., for formulae involving variants of transitivity. Also, in that paper a conjecture is stated that the problem is decidable for all universal Horn formulae. The main result of our paper is confirming this hypothesis.

THEOREM 1.2. *Let* Φ *be a* UHF *sentence. Then the local and the global satisfiability problems for modal logic over* $\mathcal{K}_\Phi$ *are in* PSPACE *and* EXPTIME*, respectively.*

This theorem reproduces some known decidability results for modal logics. It also works for some interesting classes of frames, for which, up to our knowledge, decidability has not been established so far. An example is the class defined by $\forall xyzv(xRy \wedge yRz \wedge zRv \Rightarrow xRv)$.

Besides showing the decidability we provide a full classification of universal Horn formulae with respect to the complexity of the satisfiability problem of modal logic over classes of frames they define (as mentioned, a division into formulae leading to PSPACE-hardness and to membership in NP is also given in [Hemaspaandra and Schnoor 2008]; our classification is stated in a slightly different manner). It turns out, that except for the trivial case of inconsistent formulae for which the problem is in P, local satisfiability is either NP-complete or PSPACE-complete, and global satisfiability is NP-complete, PSPACE-complete, or EXPTIME-complete

Our next contribution concerns the finite satisfiability problem. In the case of some UHF formulae Φ, we show decidability of the corresponding modal logics by demonstrating the finite model property with respect to $\mathcal{K}_\Phi$, i.e., by proving that every modal formula satisfiable over $\mathcal{K}_\Phi$ has also a finite model in $\mathcal{K}_\Phi$. However, it is not always possible, as it is not hard to construct a UHF formula Φ, such that some modal formulae have only infinite models over $\mathcal{K}_\Phi$. As an example assume that Φ enforces irreflexivity and transitivity, and consider the following modal formula: $\Diamond P \wedge \Box\Diamond P$. This naturally leads to the question, whether for any UHF formula Φ the finite satisfiability problem for modal logic over $\mathcal{K}_\Phi$ is decidable. This question is particularly important, if one considers practical applications, in which the structures (corresponding, e.g., to knowledge bases or descriptions of programs) are usually required to be finite. Decision procedures for the finite satisfiability problem for modal and related logics are very often more complex than procedures for unrestricted satisfiability. In this paper we are however able to positively answer the given question.

THEOREM 1.3. *Let Φ be a universal Horn formula. Then the local finite and the global finite satisfiability problems for modal logic over $\mathcal{K}_\Phi$ are in* NEXPTIME.

Decidability results for modal logics often extend to their *multimodal* variants. In a multimodal logic we have several pairs of operators $\Diamond^i, \Box^i$, each of them corresponding to a different binary relation R_i. As in the case of (uni)modal logics some decidability results can be obtained using the standard translation. Again, to define decidable classes of frames, we may use the two-variable fragment and its extensions, e.g., the variant in which two relations are required to be equivalences [Kieroński and Otto 2005; Kieroński et al. 2012], or the guarded logics.

Our question is whether Theorem 1.2 can be strengthened to cover multimodal logics. The answer turns out to be negative even for *bimodal* logic, i.e., for the case in which there are two binary relations.

THEOREM 1.4. *There exist universal Horn formulae Γ^b, Γ'^b, over the signature consisting with two binary relations R, R' such that the global (finite) satisfiability problem for bimodal logic over $\mathcal{K}_{\Gamma_b}$, and the local (finite) satisfiability problem for bimodal logic over $\mathcal{K}_{\Gamma'_b}$ are undecidable.*

Worth mentioning here is that Γ^b and Γ'^b use a non-trivial interaction between R and R'. We conjecture that if we allow to speak in a single clause only about one of the binary relations the problem becomes decidable.

Plan of the paper. The paper is organized as follows. In Section 2 we introduce basic definitions and recall some useful tools. In Section 3 we prove our undecidability results, i.e., Theorems 1.1 and 1.4. Section 4 establish our main decidability results. To keep the presentation light, all the technicalities are postponed to Section 5. In Section 6 we study the satisfiability problems more carefully to obtain precise complexity bounds. In Section 7 we consider satisfiability over finite models from Horn-definable classes, i.e., we prove Theorem 1.3. Finally, in Section 8 we conclude our results and describe some possible directions of further research.

2. PRELIMINARIES

2.1. Kripke structures and frames

Kripke structures are triples of the form $\langle W, R, \pi\rangle$, where W is a set of worlds, $\langle W, R\rangle$ is a directed graph called a *frame*, and π is a labelling function that assigns to each world a set of propositional variables which are true at this world.

We say that a Kripke structure $\langle W, R, \pi\rangle$ is *based* on the frame $\langle W, R\rangle$. 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. To keep the notation light, we use $\langle \mathcal{M}, \pi\rangle$, where $\mathcal{M} = \langle W, R\rangle$ is a frame, to denote the structure $\langle W, R, \pi\rangle$.

For a frame $\langle W, R\rangle$ and a subset $W' \subseteq W$, we define $R_{\restriction W'} = R \cap (W' \times W')$. Similarly, for a labelling function π, we define $\pi_{\restriction W'}$ to be such that $\pi_{\restriction W'}(w) = \pi(w)$ for all $w \in W'$ and $\pi_{\restriction X}$ to be such that $\pi_{\restriction X}(w) = \pi(w) \cap X$. For any $W' \subseteq W$, we define the restriction of a frame $\langle W, R\rangle_{\restriction W'}$ as $\langle W', R_{\restriction W'}\rangle$, and the restriction of a Kripke structure $\langle W, R, \pi\rangle_{\restriction W'}$ as $\langle W', R_{\restriction W'}, \pi_{\restriction W'}\rangle$.

A world w is said to be *k-followed* (*k-preceded*) in a frame $\mathcal{M}$, if there exists a directed path $(w, u_1, u_2, \ldots, u_k)$ (resp. $(u_1, u_2, \ldots, u_k, w)$) in $\mathcal{M}$. Note that we do not require this path to consist of distinct elements. We say that a world w is *k-inner* in $\mathcal{M}$ if it is k-preceded and k-followed. We use also naturally defined notions of ∞-preceded, ∞-followed, and ∞-inner worlds. In particular, a world on a cycle is ∞-inner.

2.2. Logics and types

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.

Except for Section 3.2, we consider only unimodal logics. Modal logic extends propositional logic with the operator $\Diamond$ and its dual $\Box$. The syntax of modal logic is given by the following BNF:

$$\varphi ::= p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \Diamond\varphi \mid \Box\varphi$$

where p is a propositional variable. Note that all formulae are in the negation normal form.

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

(i) $\varphi = p$ where p is a variable and $\varphi \in \pi(w)$,
(ii) $\varphi = \neg p$ where p is a variable and $\varphi \notin \pi(w)$,
(iii) $\varphi = \varphi_1 \vee \varphi_2$ and $\mathfrak{M}, w \models \varphi_1$ or $\mathfrak{M}, w \models \varphi_2$,
(iv) $\varphi = \varphi_1 \wedge \varphi_2$ and $\mathfrak{M}, w \models \varphi_1$ and $\mathfrak{M}, w \models \varphi_2$,
(v) $\varphi = \Diamond\varphi'$ and there exists a world $v \in W$ such that $(w, v) \in R$ and $\mathfrak{M}, v \models \varphi'$,
(vi) $\varphi = \Box\varphi'$ and for all worlds $v \in W$ such that $(w, v) \in R$ we have $\mathfrak{M}, v \models \varphi'$.

By $|\varphi|$ we denote the length of φ. 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 $\mathfrak{M}, w \models \varphi$.

We say that a structure $\mathfrak{M}$ is a *model* of φ if there is a world w such that $\mathfrak{M}, w \models \varphi$, and that $\mathfrak{M}$ is a *global model* of φ if $\mathfrak{M} \models \varphi$.

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 (resp. global) model of φ, and that φ is *finitely locally* (resp. *globally*) *$\mathcal{K}$-satisfiable* if there exists a finite $\mathcal{K}$-based (resp. global) model of φ.

The set of *universal Horn formulae* (without equality), UHF, is defined as the set of those Φ over the language $\{R\}$ which are of the form $\forall\vec{x}.\Phi_1 \wedge \Phi_2 \wedge \ldots \wedge \Phi_i$, where each Φ_i is a Horn clause. A Horn clause is a disjunction of literals of which at most one is positive. We usually present Horn clauses as implications. For example, the formula $\forall xyz.(xRy \wedge yRz \Rightarrow xRz) \wedge (xRx \Rightarrow \bot)$ defines the set of transitive and irreflexive frames. We often skip the quantifiers and represent such formulae as a set of clauses, e.g.: $\{xRy \wedge yRz \Rightarrow xRz, xRx \Rightarrow \bot\}$. By Φ^p we denote the set of the clauses from Φ containing a positive literal, i.e., all clauses of Φ except those of the form $\Psi \Rightarrow \bot$.

We assume without loss of generality that each Horn clause uses variables from the list x, y, $z_1, z_2, \ldots$, and is of the form $\Psi \Rightarrow \bot$, $\Psi \Rightarrow xRx$, or $\Psi \Rightarrow xRy$. We define $\Phi(v_x, v_y, v_1, \ldots, v_k)$ as the instantiation of Φ with $x = v_x$, $y = v_y$, $z_1 = v_1, \ldots, z_k = v_k$, e.g.

$$(xRz_1 \wedge z_1Rz_2 \wedge z_2Ry \Rightarrow xRy)(a, b, c, d) = aRc \wedge cRd \wedge dRb \Rightarrow aRb$$

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

2.3. Decision problems

For a given $\Phi \in \mathsf{UHF}$, we define $\mathcal{K}_\Phi$ as the class of frames satisfying Φ.

We define the (*unrestricted*) *local* (resp. *global*) *satisfiability problem* Φ-SAT$_L$ (resp. Φ-SAT$_G$) for modal logic *over* $\mathcal{K}_\Phi$ as the question whether a given modal formula is locally (resp. globally) $\mathcal{K}_\Phi$-satisfiable. The *finite local* (resp. *finite global*) *satisfiability problem* Φ-FINSAT$_L$ (reps. Φ-FINSAT$_G$) for modal logic *over* $\mathcal{K}_\Phi$ is the question whether a given modal formula φ is finitely locally (resp. finitely globally) $\mathcal{K}_\Phi$-satisfiable.

We abbreviate $\mathbb{N} \cup \{\infty\}$ by $\mathbb{N}_\infty$ and $\{0, 1, \ldots, k-1\}$ by $\mathbb{Z}_k$. To keep the notation light, we assume that $n \mod \infty = n$ for any n.

A *domino system* is a tuple $\mathcal{D} = (D, D_H, D_V)$, where D is a set of domino pieces and $D_H, D_V \subseteq D \times D$ are binary relations specifying admissible horizontal and vertical adjacencies. We say that $\mathcal{D}$ *tiles* $\mathbb{N} \times \mathbb{N}$ if there exists a function $t : \mathbb{N} \times \mathbb{N} \mapsto D$ such that $\forall i, j \in \mathbb{N}$ we have $(t(i, j), t(i + 1, j)) \in D_H$ and $(t(i, j), t(i, j + 1)) \in D_V$. Similarly, $\mathcal{D}$ tiles $\mathbb{Z}_k \times \mathbb{Z}_l$, for $k, l \in \mathbb{N}$, if there exists $t : \mathbb{Z}_k \times \mathbb{Z}_l \mapsto D$ such that $(t(i, j), t(i + 1 \mod k, j)) \in D_H$ and $(t(i, j), t(i, j + 1 \mod l)) \in D_V$.

The following lemma comes from [Berger 1966; Gurevich and Koryakov 1972].

LEMMA 2.1. *The following problems are undecidable:*
(i) *For a given domino system $\mathcal{D}$ determine if $\mathcal{D}$ tiles $\mathbb{N} \times \mathbb{N}$.*
(ii) *For a given domino system $\mathcal{D}$ determine if there exists $k \in \mathbb{N}$ such that $\mathcal{D}$ tiles $\mathbb{Z}_k \times \mathbb{Z}_k$.*

The *bounded-space domino problem* is defined as follows. For a given tuple $\langle D, D_H, D_V, n\rangle$, where $D_H, D_V \subseteq D \times D$, and n is a natural number given in unary, is there a tiling $t : \mathbb{Z}_n \times \mathbb{N} \to D$ such that for all $k < n$ and $l \in \mathbb{N}$, $(t(k, l), t(k, l + 1)) \in D_V$ and if $k < n - 1$, then $(t(k, l), t(k + 1, l)) \in D_H$? This problem is known to be PSPACE-complete [van Emde Boas 1997].

3. UNDECIDABILITY

In this section we present our undecidability results stated in Theorems 1.1 and 1.4. In both proofs we work with sets of propositional variables including $\mathcal{P} = \{P_{ij} : 0 \le i, j \le 2\}$. As our approach employs reductions from domino systems, we are particularly interested in grid-like structures. The intended meaning of P_{ij} is to mark elements encoding the grid nodes whose coordinates modulo 3 are equal i, j. To simplify notation we assume in this section that the indices in P_{ij} are always taken modulo 3, e.g., if $i = 2, j = 0$, then $P_{i+1,j-1}$ denotes P_{02}.

3.1. Modal logic over classes defined by three variable formulae

First, we prove Theorem 1.1. Let

$$\Gamma = \neg xRy \vee yRx \vee \neg xRz \vee zRx \vee yRz \vee zRy.$$

Note that we skipped the prefix $\forall xyz$ as mentioned in the preliminaries. First, we show that Γ-SAT$_G$ and Γ-FINSAT$_G$ are undecidable. Then we use the trick from [Hemaspaandra and Schnoor 2011] and show that also Γ'-SAT$_L$ and Γ'-FINSAT$_L$ are undecidable, for Γ' being a modification of Γ, using still only three variables.

Note that Γ can be rewritten as $(xRy \wedge \neg yRx \wedge xRz \wedge \neg zRx) \Rightarrow (yRz \vee zRy)$, i.e., it says, that if there are one-way connections from a world x to worlds y, z, then there is also a connection (not necessarily one-way) between y and z. The structure $\mathfrak{G}_\mathbb{N}$, illustrated in Fig. 1, whose universe is $\mathbb{N} \times \mathbb{N}$, is a model of Γ. This model is reflexive (the reflexive arrows are omitted from the picture for clarity). Note that, actually, Γ enforces reflexivity at all worlds having incoming edges. It is also important that some connections are two-way; with one-way arrows, edges between some distant worlds would be required by Γ. Observe that owing to $\mathcal{P}$-labels each world can identify (using the modal language) its horizontal and its vertical successor.

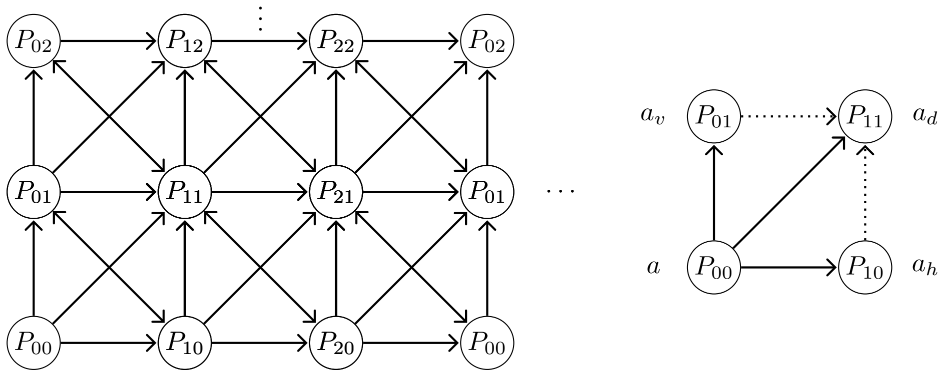

Fig. 1. The structure $\mathfrak{G}_{\mathbb{N}}$. Reflexive arrows at all nodes are omitted for clarity.

To show the undecidability we construct a modal formula ζ, capturing some properties of $\mathfrak{G}_{\mathbb{N}}$, such that any model $\mathfrak{M} \models \zeta$ from $\mathcal{K}_\Gamma$ locally looks like a grid. Namely, ζ says that every element satisfying P_{ij} has three R-successors: one in $P_{i+1,j}$, one in $P_{i,j+1}$, one in $P_{i+1,j+1}$, and forbids connections from $P_{i+1,j+1}$ to $P_{i,j+1}$, $P_{i+1,j}$, and P_{ij}. If we consider now any element a in a model, we see that ζ enforces the existence of its horizontal successor a_h, its vertical successor a_v and its upper-right diagonal successor a_d (see the right part of Fig. 1). By ζ, the connections to these successors are one-way, so we need, by Γ, connections between a_h and a_d, and a_v and a_d. Again, by ζ, these connections has to go from a_h to a_d, and from a_v to a_d, so a_d is indeed a horizontal successor of a_v, and a vertical successor of a_h.

Below we present the details, and cover also the case of finite satisfiability, i.e., satisfiability in the class of finite models. The technique we employ is quite standard. It is similar, e.g., to the technique used in [Otto 2001].

We define ζ as follows:

$$\zeta = \zeta_0 \wedge \bigwedge_{0 \le i,j \le 2} (\zeta_{ij}^{\Diamond} \wedge \zeta_{ij}^{\Box}),$$

where ζ_0 says that each element satisfies precisely one variable from $\mathcal{P}$, $\zeta_{ij}^{\Diamond}$ ensures that all elements have appropriate horizontal, vertical and upper-right diagonal successors, and $\zeta_{ij}^{\Box}$ forbids reversing horizontal, vertical and upper-right diagonal arrows.

$$\zeta_0 = \dot{\bigvee}_{0 \le i,j \le 2} P_{ij},$$

$$\zeta_{ij}^{\Diamond} = P_{ij} \to (\Diamond P_{i+1,j} \wedge \Diamond P_{i,j+1} \wedge \Diamond P_{i+1,j+1}),$$

$$\zeta_{ij}^{\Box} = P_{ij} \to \Box(\neg P_{i-1,j} \wedge \neg P_{i,j-1} \wedge \neg P_{i-1,j-1}).$$

where $\dot{\bigvee}_{0 \le i,j \le 2} P_{ij}$ means that exactly one of $P_{00}, \ldots, P_{22}$ is satisfied. Note that $\zeta_{ij}^{\Box}$ admits reflexive edges.

We encode an instance of the domino problem $\mathcal{D} = (D, D_H, D_V)$ by a modal formula in a standard fashion. For every $d \in D$ we introduce a fresh propositional variable Q_d. Let λ_0 says that each world labelled by a variable from $\mathcal{P}$ contains a domino piece, and let λ_{ij}^H and λ_{ij}^V say that the pairs of elements satisfying horizontal or vertical adjacency relations respect D_H and D_V, respectively. These can be expressed as follows.

$$\lambda_0 = \bigwedge_{0 \le i,j \le 2} (P_{ij} \to \dot{\bigvee}_{d \in D} Q_d),$$

$$\lambda^H_{ij} = \bigwedge_{d \in D} ((Q_d \wedge P_{ij}) \rightarrow \Box(P_{i+1,j} \rightarrow \bigvee_{d':(d,d') \in D_H} Q_{d'})),$$

$$\lambda^V_{ij} = \bigwedge_{d \in D} ((Q_d \wedge P_{ij}) \rightarrow \Box(P_{i,j+1} \rightarrow \bigvee_{d':(d,d') \in D_V} Q_{d'})).$$

We define

$$\lambda^{\mathcal{D}} = \lambda_0 \wedge \bigwedge_{0 \le i,j \le 2} (\lambda^H_{ij} \wedge \lambda^V_{ij}).$$

Undecidability of Γ-SAT$_G$ and Γ-FINSAT$_G$ follows from Lemma 2.1 and the following claim.

CLAIM 3.1. (i) *$\mathcal{D}$ tiles $\mathbb{N} \times \mathbb{N}$ iff there exists $\mathfrak{M} \in \mathcal{K}_\Gamma$ such that $\mathfrak{M} \models \zeta \wedge \lambda^{\mathcal{D}}$.*
(ii) *$\mathcal{D}$ tiles some $\mathbb{Z}_m \times \mathbb{Z}_m$ iff there exists a finite $\mathfrak{M} \in \mathcal{K}_\Gamma$ such that $\mathfrak{M} \models \zeta \wedge \lambda^{\mathcal{D}}$.*

PROOF. As in the case of symbols P_{ij}, when referring to $\zeta^{\Box}_{ij}$ or $\zeta^{\Diamond}_{ij}$ we assume that subscripts are taken modulo 3.

Part (i), $\Rightarrow$ Let t be a tiling of $\mathbb{N} \times \mathbb{N}$. We construct $\mathfrak{M}$ by expanding $\mathfrak{G}_{\mathbb{N}}$ in such a way that for every $i, j \in \mathbb{N}$ the element (i, j) satisfies $Q_{t(i,j)}$. One can check that $\mathfrak{M}$ is as required.

Part (i), $\Leftarrow$ We explain how to construct a homomorphism-like function $f : \mathbb{N} \times \mathbb{N} \mapsto M$, where M denotes the universe of $\mathfrak{M}$, such that for every $i, j \in \mathbb{N}$: (a) $\mathfrak{M} \models P_{ij}(f(i, j))$, (b) $\mathfrak{M} \models f(i, j)Rf(i + 1, j)$, (c) $\mathfrak{M} \models f(i, j)Rf(i, j + 1)$.

First we show how to define f on $\mathbb{N} \times \{0\}$. Let $f(0, 0) = c$ for an arbitrary world c from $\mathfrak{M}$ satisfying P_{00}. Such c exists owing to ζ_0 and $\zeta^{\Diamond}_{ij}$. Assume inductively that for some $i > 0$ we have defined $f(i - 1, 0) = a$, and let a_h be an R-successor of a satisfying P_{i0}. Such a_h exists owing to $\zeta^{\Diamond}_{i-1,0}$. Define $f(i, 0) = a_h$.

Assume now that f is defined for $\mathbb{N} \times [0, \ldots, j-1]$ for some $j > 0$. We extend this definition to $\mathbb{N} \times \{j\}$. Let $f(0, j - 1) = a$. By the inductive assumption a satisfies $P_{0,j-1}$. Choose a_v to be an R-successor of a satisfying P_{0j}. Such a_v exists by $\zeta^{\Diamond}_{0,j-1}$. Set $f(0, j) = a_v$. Assume inductively that we have defined $f(i - 1, j - 1) = a$, $f(i - 1, j) = a_v$, and $f(i, j - 1) = a_h$. By the inductive assumptions $\mathfrak{M} \models P_{i-1,j-1}(a) \wedge P_{i-1,j}(a_v) \wedge P_{i,j-1}(a_h) \wedge aRa_h \wedge aRa_v$. Choose a_d to be an R-successor of a satisfying P_{ij}. Such a_d exists by $\zeta^{\Diamond}_{i-1,j-1}$. By $\zeta^{\Box}_{ij}$, a_h, a_v and a_d cannot be connected to a, so Γ enforces R-connections between a_h and a_d, and between a_v and a_d. Since $\zeta^{\Box}_{ij}$ forbids connection from a_d to a_h, and from a_d to a_v, it has to be that $\mathfrak{M} \models a_hRa_d \wedge a_vRa_d$. This finishes the definition of f with the desired properties.

We define a tiling $t : \mathbb{N} \times \mathbb{N}$ by setting $t(i, j) = d$ for such d that $f(i, j)$ satisfies Q_d (there is precisely one such d owing to λ_0). We argue that this tiling is correct. Let $(i, j) \in \mathbb{N} \times \mathbb{N}$. Let $a = f(i, j), a_h = f(i + 1, j), a_v = f(i, j + 1)$. By property (a) of f we have $\mathfrak{M} \models P_{ij}(a) \wedge P_{i+1,j}(a_h) \wedge P_{i,j+1}(c)$. Properties (b) and (c) imply that $\mathfrak{M} \models aRa_h \wedge aRa_v$. Assume that $t(i, j) = d$, $t(i + 1, j) = d_h$, $t(i, j + 1) = d_v$ and thus $\mathfrak{M} \models Q_d(a) \wedge Q_{d_h}(a_h) \wedge Q_{d_v}(a_v)$. By λ^V_{ij} it follows that $(d, d_h) \in D_H$ and by λ^V_{ij} – that $(d, d_v) \in D_V$.

Part (ii) $\Rightarrow$ Let $l = 3k$ for some $k \in \mathbb{Z}$. We define $\mathfrak{G}_l$ to be the quotient of $\mathfrak{G}_{\mathbb{N}}$ by the relation $\approx$ defined as follows: $(i, j) \approx (i', j')$ iff $i \equiv i' \mod l$ and $j \equiv j' \mod l$. The frame of $\mathfrak{G}_l$ can be seen as a grid on $\mathbb{Z}_l \times \mathbb{Z}_l$ torus with the additional diagonal edges. It is readily checked that for every $k \in \mathbb{N}$ we have $\mathfrak{G}_{3k} \models \Gamma$ and $\mathfrak{G}_{3k} \models \zeta$.

If $\mathcal{D}$ tiles $\mathbb{Z}_m \times \mathbb{Z}_m$ then it also tiles $\mathbb{Z}_{3m} \times \mathbb{Z}_{3m}$. Let t be a tiling of $\mathbb{Z}_{3m} \times \mathbb{Z}_{3m}$. We construct $\mathfrak{M}$ by expanding $\mathfrak{G}_{3m}$ in such a way that for every $i, j \in \mathbb{Z}_{3m}$ the element (i, j) satisfies $Q_{t(i,j)}$. Again, checking that $\mathfrak{M}$ is as required is straightforward.

Part (ii) $\Leftarrow$ We want to define for some $k, l \in \mathbb{N}$ a function $f : \mathbb{Z}_k \times \mathbb{Z}_l \mapsto M$ satisfying: (a) $\mathfrak{M} \models P_{ij}(f(i,j))$, (b) $\mathfrak{M} \models f(i,j)Rf(i+1 \mod k, j)$, (c) $\mathfrak{M} \models f(i,j)Rf(i,j+1 \mod l)$.

We define f as a partial function on $\mathbb{N} \times \mathbb{N}$ and then restrict it to an appropriate domain. We first define f on $\mathbb{N} \times \{0\}$, exactly as in the proof of Part (i), $\Leftarrow$. Since $\mathfrak{M}$ is finite this time, it has to be that $f(k,0) = f(k',0)$ for some $k > k'$. To simplify the presentation we assume $k' = 0$, but this assumption is not relevant. Observe that for $i \in [0,k)$ we have $\mathfrak{M} \models f(i,0)Rf(i+1 \mod k, 0)$. We extend the definition of f to $[0,k) \times \mathbb{N}$ inductively. Assume that f is defined on $[0,k) \times \{0,\ldots,j-1\}$. We define it on $[0,k) \times \{j\}$. For each $i \in [0,k)$ we find an element a_d^i in M such that $\mathfrak{M} \models P_{i+1,j}(a_d^i) \wedge f(i,j-1)Ra_d^i$. Such a_d^i exists owing to $\zeta^{\Diamond}_{i,j-1}$. We set $f(i+1 \mod k, j) = a_d^i$. Now Γ and formulae of the type $\zeta^{\Box}$ enforce for all $i \in [0,k)$ that $\mathfrak{M} \models f(i,j-1)Rf(i,j)$, and $\mathfrak{M} \models f(i,j)Rf(i+1 \mod k, j)$.

Finiteness of $\mathfrak{M}$ implies that for some $l > l'$ we have $f \restriction [0,k) \times \{l\} = f \restriction [0,k) \times \{l'\}$. Again for simplicity we assume that $l' = 0$. Observe that at this moment f is as desired on $\mathbb{Z}_k \times \mathbb{Z}_l$. We define a tiling $t : \mathbb{Z}_k \times \mathbb{Z}_l$ by setting $t(i,j) = d$ for such d that $f(i,j)$ satisfies Q_d (there is precisely one such d owing to λ_0). As in the proof of Part (i), $\Leftarrow$, properties (a), (b), (c) of f and the formulae λ^H_{ij} and λ^V_{ij} imply that t is a correct tiling of $\mathbb{Z}_k \times \mathbb{Z}_l$. This guarantees that there exists also a correct tiling of $\mathbb{Z}_m \times \mathbb{Z}_m$ for $m = gcd(k,l)$. $\Box$

Now we consider the case of local satisfiability. Observe that our proof of the undecidability of global satisfiability over $\mathcal{K}_\Gamma$ works for the subclass $\mathcal{K}^{ref}_\Gamma$ consisting of all reflexive structures from $\mathcal{K}_\Gamma$. This allows us to reproduce the trick from [Hemaspaandra and Schnoor 2011] and reduce global satisfiability over $\mathcal{K}^{ref}_\Gamma$ to local satisfiability over $\mathcal{K}_{\Gamma'}$, for Γ' saying that each world with an incoming edge is reflexive and has an incoming edge from all irreflexive worlds, and enforcing Γ in the substructure consisting of all the reflexive worlds:

$$\Gamma' = (xRy \wedge \neg zRz \Rightarrow yRy \wedge zRy) \wedge (xRx \wedge yRy \wedge zRz \Rightarrow \Gamma).$$

The reduction can be carried out by requiring the existence of an irreflexive *universal* world, i.e., an element connected to all relevant elements in the model.

We recall a classical observation (see, e.g., [Blackburn et al. 2001]), which will be used in our arguments.

LEMMA 3.2. *Let $\mathfrak{M}$ be a Kripke structure. If $\mathfrak{M}'$ is its generated substructure, i.e., it is closed under R, then for every world a from $\mathfrak{M}'$, and every modal formula we have $\mathfrak{M}', a \models \varphi$ iff $\mathfrak{M}, a \models \varphi$.*

We are ready to show the reduction.

LEMMA 3.3. *For every modal formula φ, φ is globally satisfiable in a (finite) model from $\mathcal{K}^{ref}_\Gamma$ iff $\varphi' = Q_U \wedge \Box\neg Q_U \wedge \Diamond\top \wedge \Box\varphi$, where Q_U is a fresh propositional variable, is locally satisfiable in a (finite) model from $\mathcal{K}_{\Gamma'}$.*

PROOF. $\Rightarrow$. Let $\mathfrak{M} \models \varphi$ be a model from $\mathcal{K}^{ref}_\Gamma$. Let $\mathfrak{M}'$ be the structure consisting of $\mathfrak{M}$ and one additional, irreflexive world c, labelled with Q_U, having no incoming edges, and having outgoing edges to the all worlds from $\mathfrak{M}$. Note that $\mathfrak{M}$ is a generated substructure of $\mathfrak{M}'$, so by Lemma 3.2, for its every element a, we have $\mathfrak{M}', a \models \varphi$. It implies that $\mathfrak{M}', c \models \Box\varphi$. Obviously $\mathfrak{M}', c \models Q_U \wedge \Box\neg Q_U \wedge \Diamond\top$, and thus $\mathfrak{M}', c \models \varphi'$. It is straightforward to verify that $\mathfrak{M}' \in \mathcal{K}_{\Gamma'}$.

$\Leftarrow$ Assume that for some $\mathfrak{M} \in \mathcal{K}_{\Gamma'}$ and some world c from $\mathfrak{M}$ we have $\mathfrak{M}, c \models \varphi'$. Let $\mathfrak{M}'$ be the structure obtained by restricting $\mathfrak{M}$ to all worlds with at least one incoming edge. By the first implication in Γ' all worlds from $\mathfrak{M}'$ are reflexive and thus, by the second implication, $\mathfrak{M}' \in \mathcal{K}^{ref}_\Gamma$. As $\mathfrak{M}, c \models Q_U \wedge \Box\neg Q_U$ it follows that c is an irreflexive world, and, thus, by the first implication in Γ' it has edges in $\mathfrak{M}$ to the all worlds of $\mathfrak{M}'$. Consider any world $a \in \mathfrak{M}'$. Since $\mathfrak{M}, c \models \Box\varphi$, we have $\mathfrak{M}, a \models \varphi$. By definition $\mathfrak{M}'$ is closed under R, so by Lemma 3.2, $\mathfrak{M}', a \models \varphi$. As a has been chosen arbitrarily we conclude that $\mathfrak{M}' \models \varphi$. $\Box$

3.2. Bimodal logic over Horn-definable classes of frames

The syntax of bimodal logic extends the syntax of modal logic by the additional modalities $\Diamond'$, $\Box'$. Its formulae are interpreted over bimodal Kripke structures containing an additional binary relation. A bimodal Kripke frame is a triple $\langle W, R, R'\rangle$, where $R, R' \subseteq W^2$. A bimodal Kripke structure is a pair $(\langle W, R, R'\rangle, \pi)$ where $\langle W, R, R'\rangle$ is a bimodal frame and π is defined as in the case of unimodal logic. The semantics of bimodal logic extends the semantics of modal logic by the following rules:

(vii) $\varphi = \Diamond'\varphi'$ and there is a world $v \in W$ such that $(w, v) \in R'$ and $\mathfrak{M}, v \models \varphi'$,
(viii) $\varphi = \Box'\varphi'$ and for all worlds $v \in W$ such that $(w, v) \in R'$ we have $\mathfrak{M}, v \models \varphi'$.

First-order formulae defining classes of bimodal Kripke frames are now allowed to use two binary symbols R, R'. E.g., if $\Phi = R(x, y) \Leftrightarrow R'(y, x)$ then $\mathcal{K}_\Phi$ is the class of all frames in which R' is interpreted as the inverse of R.

In the global satisfiability case the undecidability of bimodal logic can be shown over the class of structures definable by a very simple formula $zRx \wedge zR'y \Rightarrow xRy$. However, as we want to extend the result to cover also the local satisfiability case we use a slightly more complicated one:

$$\Gamma^b = zRx \wedge xRs \wedge zRu \wedge uR'y \Rightarrow xRy.$$

The structure $\mathfrak{G}^b_{\mathbb{N}}$, depicted in Fig. 2, is its model. The universe of $\mathfrak{G}^b_{\mathbb{N}}$ consists of two copies of $\mathbb{N}\times\mathbb{N}$, i.e., it is equal $\{(i, j), \overline{(i, j)} : i, j \in \mathbb{N}\}$. Each (i, j) has an R-edge to $(i{+}1, j)$, $(i, j{+}1)$ and $\overline{(i, j)}$, and it has no outgoing R'-edges. Each $\overline{(i, j)}$ has an R'-edge to $(i{+}1, j{+}1)$, and has no outgoing R-edges. Besides $\mathcal{P}$ its signature includes another set of propositional variables $\overline{\mathcal{P}} = \{\overline{P}_{ij} : 0 \le i, j \le 2\}$. Each element satisfies precisely one propositional variable, namely, (i, j) satisfies P_{ij}, and $\overline{(i, j)}$ satisfies $\overline{P}_{ij}$. Note the role played by the atom xRs in Γ^b – without it, the $\overline{\mathcal{P}}$-labelled worlds would have to be R-connected to some $\mathcal{P}$-labelled worlds.

Our proof strategy is similar to the one we applied in the proof of Claim 3.1. We capture some properties of $\mathfrak{G}^b_{\mathbb{N}}$ by a modal formula ξ:

$$\xi = \xi_0 \wedge \bigwedge_{0 \le i,j \le 2} (\xi^{\Diamond}_{ij} \wedge \wedge \xi'^{\Diamond}_{ij}),$$

where ξ_0 says that each element satisfies precisely one of the variables from the set $\mathcal{P} \cup \overline{\mathcal{P}}$, $\xi^{\Diamond}_{ij}$ ensures that the $\mathcal{P}$-labelled worlds have appropriate successors, and $\xi'^{\Diamond}_{ij}$ ensures that $\overline{\mathcal{P}}'$-labelled worlds have appropriate successors.

$$\xi_0 = \dot{\bigvee}_{0 \le i,j \le 2} (P_{ij} \vee \overline{P}_{ij}) \wedge \bigwedge_{0 \le i,j \le 2} (\neg P_{ij} \vee \neg \overline{P}_{ij})$$

$$\xi^{\Diamond}_{ij} = P_{ij} \rightarrow \Diamond P_{i+1,j} \wedge \Diamond P_{i,j+1} \wedge \Diamond \overline{P}_{ij},$$

$$\xi'^{\Diamond}_{ij} = \overline{P}_{ij} \rightarrow \Diamond' P_{i+1,j+1}$$

For a given domino system $\mathcal{D}$ we construct the formula $\lambda^{\mathcal{D}}$ precisely as in the previous subsection. The undecidability of Γ^b-SAT$_G$ and Γ^b-FINSAT$_G$ follows now from Lemma 2.1 and the following observation, which is a counterpart of Claim 3.1 and can be proved in a similar way.

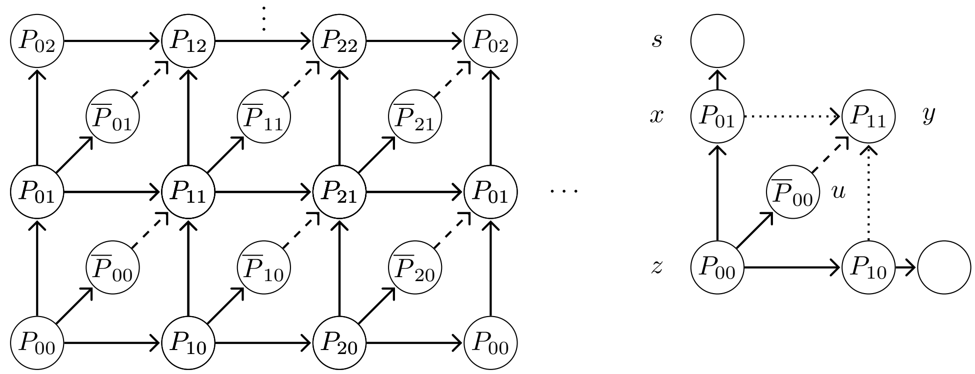

Fig. 2. The structure $\mathfrak{G}^b_{\mathbb{N}}$. Relation R' is represented by dashed arrows. All worlds are irreflexive.

CLAIM 3.4. (i) *$\mathcal{D}$ tiles $\mathbb{N} \times \mathbb{N}$ iff there exists $\mathfrak{M} \in \mathcal{K}_{\Gamma^b}$ such that $\mathfrak{M} \models \xi \wedge \lambda^{\mathcal{D}}$.*
(ii) *$\mathcal{D}$ tiles some $\mathbb{Z}_m \times \mathbb{Z}_m$ iff there exists a finite $\mathfrak{M} \in \mathcal{K}_{\Gamma^b}$ such that $\mathfrak{M} \models \xi \wedge \lambda^{\mathcal{D}}$.*

To show undecidability of the local satisfiability problem, we use this time the fact that the proof for the global satisfiability works for the subclass $\mathcal{K}^s_{\Gamma^b}$ of models containing an element with an R-successor, and in which every element has some R-successors iff it has no R'-successors. It allows us to use again the trick with a universal world (this time with respect to R'-relation). Namely, we require that every element with both R- and R'-successors is connected by R' to every world with at least one R'-predecessor, or at least one R-predecessor. This is obtained by modifying Γ^b to Γ'^b:

$$\Gamma'^b = (xRu_1 \wedge xR'u_2 \wedge vR'y \Rightarrow xR'y) \wedge (xRu_1 \wedge xR'u_2 \wedge vRy \Rightarrow xR'y) \wedge \Gamma^b.$$

We recall the generalisation of Lemma 3.2 to bimodal logic.

LEMMA 3.5. *Let $\mathfrak{M}$ be a bimodal Kripke structure. If $\mathfrak{M}'$ is its generated substructure, i.e., it is closed under R and under R', then for every world a from $\mathfrak{M}'$, and every bimodal formula we have $\mathfrak{M}', a \models \varphi$ iff $\mathfrak{M}, a \models \varphi$.*

In the following claim we reduce global satisfiability over $\mathcal{K}^s_{\Gamma'^b}$ to local satisfiability over $\mathcal{K}_{\Gamma'^b}$, which implies the undecidablity of Γ'^b-SAT$_L$ and Γ'^b-SAT$_G$.

CLAIM 3.6. *For every modal formula φ, φ is globally satisfiable in a (finite) model from $\mathcal{K}^s_{\Gamma'^b}$ iff $\varphi' = \Diamond\top \wedge \Diamond'\top \wedge \Box'\varphi \wedge \Box'(\Diamond\top \leftrightarrow \Diamond'\top) \wedge \Diamond'\Diamond\top$ is locally satisfiable in a (finite) model from $\mathcal{K}_{\Gamma'^b}$.*

PROOF. Let $\mathfrak{M} \models \varphi$ be model from $\mathcal{K}^s_{\Gamma'^b}$ and b its element with an R-successor (and thus no R'-successors). Let $\mathfrak{M}'$ be the structure consisting of $\mathfrak{M}$ and one additional world c with an outgoing R'-edge to each element of $\mathfrak{M}$, and with an outgoing R-edge to b. It is straightforward to verify that $\mathfrak{M}' \in \mathcal{K}_{\Gamma'^b}$. $\mathfrak{M}$ is a generated substructure of $\mathfrak{M}'$, so by Lemma 3.5, for all a from $\mathfrak{M}$ we have $\mathfrak{M}', a \models \varphi$. Thus $\mathfrak{M}', c \models \Box'\varphi$. By definition of $\mathcal{K}^s_{\Gamma'^b}$ it follows that also $\mathfrak{M}', c \models \Box'(\Diamond\top \leftrightarrow \Diamond'\top) \wedge \Diamond'\Diamond\top$. Finally, $\mathfrak{M}', c \models \varphi'$.

In the opposite direction, assume that for some $\mathfrak{M} \in \mathcal{K}_{\Gamma'^b}$ and some world c from $\mathfrak{M}$ we have $\mathfrak{M}, c \models \varphi'$. Let $\mathfrak{M}'$ be the structure obtained by restricting $\mathfrak{M}$ to all worlds with at least one incoming R- or R'-edge. Obviously, as a substructure of $\mathfrak{M}$, $\mathfrak{M}' \in \mathcal{K}_{\Gamma^b}$. By Γ'^b, $\mathfrak{M}'$ contains precisely those worlds of $\mathfrak{M}$ which are R'-successors of c. By φ' this means that $\mathfrak{M}' \in \mathcal{K}^s_{\Gamma'^b}$. Also, by its definition, $\mathfrak{M}'$ is closed under R and R'. Since for every a from $\mathfrak{M}'$ we have $\mathfrak{M}, a \models \varphi$ it follows by Lemma 3.5 that $\mathfrak{M}', a \models \varphi$. Thus $\mathfrak{M}' \models \varphi$. □

4. THE DECIDABILITY

A well-known result says that every satisfiable modal formula is satisfied in a finite tree. This *tree-model property* is crucial for the robust decidability of modal logics. Standard restrictions of classes of frames lead to similar results, stating that some "nice" models exist for all satisfiable formulae. Here we generalize those results for the classes of models that are definable by the Horn formulae. To demonstrate our technique in a more reader-friendly way, in this section we focus on the big picture and postpone most of the proofs to Section 5. For the rest of this section we fix an arbitrary UHF sentence Φ.

First of all, we show that for every modal formula φ, if φ is $\mathcal{K}_\Phi$-satisfiable then it has a "nice" model. This is performed in a few steps. In Subsection 4.1 we define the *closure operator* and we prove that each satisfiable formula has a model based on the closure of a tree. Then, we consider the following structures.

Definition 4.1. The linear structure $\mathcal{L}_\mathbb{Z}$ is defined as $\langle\{\underline{i} : i \in \mathbb{Z}\}, \{(\underline{i}, \underline{i+1}) : i \in \mathbb{Z}\}\rangle$. The infinite binary tree $\mathcal{T}_\infty$ is defined as $\langle\{\underline{s} : s \in \{0,1\}^*\}, \{(\underline{s}, \underline{si}) : s \in \{0,1\}^* \wedge i \in \{0,1\}\}\rangle$.

Notice that the structure $\mathcal{L}_\mathbb{Z}$, unlike typical frames in modal logic, has no world from which all the other world are reachable. In Subsections 4.2 and 4.3 we study properties of the closures of $\mathcal{L}_\mathbb{Z}$, $\mathcal{T}_\infty$, and we reveal the connections between properties of the closures of arbitrary trees and properties of the closures of $\mathcal{L}_\mathbb{Z}$ and $\mathcal{T}_\infty$.

Recall that Φ is not a part of an instance, therefore to prove Theorem 1.2, it is enough to show that for every Φ there is an algorithm solving Φ-SAT$_L$. We describe the algorithms in Sections 4.4 and 4.5.

4.1. Minimal tree-based models

We start from an arbitrary $\mathcal{K}_\Phi$-based model $\mathfrak{M} \models \varphi$ and unravel it (using standard unraveling technique, as in [Sahlqvist 1975] and [Blackburn et al. 2001]) into a model $\mathfrak{M}_0$ whose frame is a tree with the degree of its nodes bounded by $|\varphi|$. Clearly the frame of $\mathfrak{M}_0$ is not necessarily a member of $\mathcal{K}_\Phi$. Next, we add to $\mathfrak{M}_0$ the edges implied by the Horn clauses of Φ. This is performed in countably many steps, until the least fixed point is reached. Observe that the resulting structure, $\mathfrak{M}_\infty$, is still a model of φ, and its frame belongs to $\mathcal{K}_\Phi$.

Formally, we say that an edge (w, w') is a *consequence* of Φ in $\mathcal{M} = \langle W, R\rangle$, if for some worlds $v_1, \ldots, v_k \in W$ and $\Psi_1 \Rightarrow \Psi_2 \in \Phi$ we have $\mathcal{M} \models \Psi_1(w, w', v_1, \ldots, v_k)$, and $\Psi_2(w, w', v_1, \ldots, v_k) = wRw'$. We denote the set of all consequences of Φ in $\mathcal{M}$ by $\mathrm{C}^\Phi_\rightsquigarrow(\mathcal{M})$. We define the *consequence operator* and the *closure operator* as follows.

$$\textsc{Cons}_{\Phi,W}(R) = R \cup \mathrm{C}^\Phi_\rightsquigarrow(\langle W, R\rangle)$$
$$\textsc{Closure}_{\Phi,W}(R) = \bigcup_{i>0} \textsc{Cons}^i_{\Phi,W}(R)$$

Note that the closure operator is the least fixed-point of *Cons*.

Example 4.2. Consider the tree $\langle W, R\rangle$ presented in Fig. 3 and $\Phi = \{xRz \wedge zRy \Rightarrow yRy, xRx \wedge xRy \wedge xRz \Rightarrow yRz\}$. Reflexive edges belong to $\textsc{Cons}_{\Phi,W}(R)$, dashed edges belong to $\textsc{Cons}^2_{\Phi,W}(R)$, and dotted edges belong to $\textsc{Cons}^3_{\Phi,W}(R)$. Quick check shows that $\textsc{Cons}^3_{\Phi,W}(R) = \textsc{Cons}^4_{\Phi,W}(R)$, hence $\textsc{Cons}^3_{\Phi,W}(R)$ equals $\textsc{Closure}_{\Phi,W}(R)$.

For a tree $\mathcal{T} = \langle W, R\rangle$, we define *the closure of* $\mathcal{T}$, as $\mathfrak{C}_\Phi(\mathcal{T}) = \langle W, \textsc{Closure}_{\Phi,W}(R)\rangle$. We denote by Φ^+ the set of the clauses of Φ containing a positive literal, i.e., all clauses of Φ except those of the form $\Psi \Rightarrow \bot$. Note that $\mathfrak{C}_\Phi(\mathcal{T})$ is the smallest (w.r.t. inclusion of the sets of edges) model of Φ^+ containing all edges from R. If $\mathfrak{C}_\Phi(\mathcal{T})$ is a model of Φ, we call it the $\mathcal{T}$*-based model of* Φ.

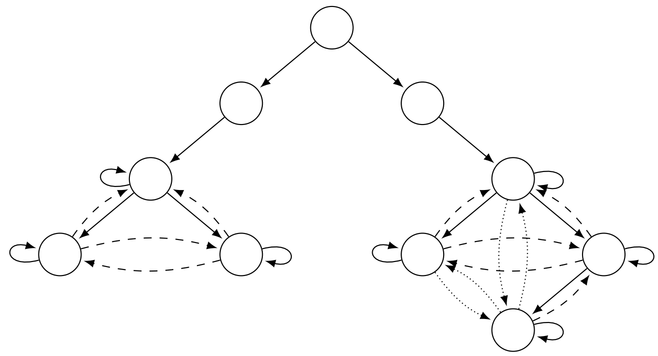

Fig. 3. A closure for $\Phi = \{xRz \wedge zRy \Rightarrow yRy, xRx \wedge xRy \wedge xRz \Rightarrow yRz\}$.

Not all models can be obtained as closures. The following lemma shows, however, that we can restrict our attention to models that are $\mathcal{T}$-*based* for some tree $\mathcal{T}$ with bounded degree.

LEMMA 4.3. *If a modal formula φ has a $\mathcal{K}_\Phi$-based model (global model), then there exists a tree $\mathcal{T}$ with the degree bounded by $|\varphi|$ and a labelling $\pi_\mathcal{T}$, such that*

(i) $\langle \mathcal{T}, \pi_\mathcal{T} \rangle$ is a model (resp. global model) of φ;
(ii) $\langle \mathfrak{C}_\Phi(\mathcal{T}), \pi_\mathcal{T} \rangle$ is a model (resp. global model) of φ that satisfies Φ.

4.2. The closures of the linear structure

We study the possible shapes of $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$. We say that an edge $(\underline{i}, \underline{j})$ is *forward* if $i < j$, *backward* if $i > j$, *short* if $|i - j| < 2$, and *long* if $|i - j| \geq 2$. We say that Φ *forces* long (resp. backward) edges if there is a long (resp. backward) edge in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ and that Φ *forces only long forward edges* if it forces long edges but it does not force backward edges.

Definition 4.4. We say that Φ satisfies

S1 if Φ does not force long edges,
S2 if Φ forces only long forward edges and there is a finite set $\chi \subseteq \mathbb{N}$ s.t. for all $i \geq 0$, $b > 0$, there is an edge from $\underline{i}$ to $\underline{i+b}$ in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ iff $b - 1$ is in the additive closure of χ.
S3 if Φ forces long and backward edges and there exists m such that for all worlds $\underline{i}$, $\underline{i+b}$, there is an edge from $\underline{i}$ to $\underline{i+b}$ in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ iff m divides $|b-1|$.

In the examples below, we abbreviate $xRu_1 \wedge u_1Ru_2 \wedge \cdots \wedge u_{i-2}Ru_{i-1} \wedge u_{i-1}Ry$ by xR^iy.

Example 4.5. Consider a formula $xR^2y \Rightarrow yRx$. Here, Property S3 is satisfied for $m = 3$. For example, $\underline{0}$ is connected to $\underline{1}$, $\underline{4}$, $\underline{7}$ and so on, while $\underline{2}$, $\underline{5}$, $\underline{8}$ and so on are connected to $\underline{0}$ (see Fig. 4a). In general, formula $xR^iy \Rightarrow yRx$ satisfies Property S3 with $m = i + 1$.

Example 4.6. Consider a formula $\varphi_3 \wedge \varphi_4$, where $\varphi_i = xR^iy \Rightarrow xRy$. Here, Property S2 is satisfied for $\chi = \{2, 3\}$. For example, $\underline{0}$ is connected to $\underline{1}$ (as in $\mathcal{L}_\mathbb{Z}$), $\underline{3}$ (because of φ_3), $\underline{4}$ (because of φ_4), $\underline{5}$ (because of φ_3, $\underline{0}R\underline{3}$, $\underline{3}R\underline{4}$, and $\underline{4}R\underline{5}$), and so on (see Fig. 4b). In general, a formula of the form $\varphi_i \wedge \varphi_j$ satisfies Property S2 with $\chi = \{i - 1, j - 1\}$.

It turns out that Properties S1, S2, and S3 cover all possible formulae.

LEMMA 4.7. *Φ satisfies S1, S2, or S3.*

4.3. The closures of trees

The following property tells us whether Φ enforces edges between different branches of $\mathcal{T}_\infty$.

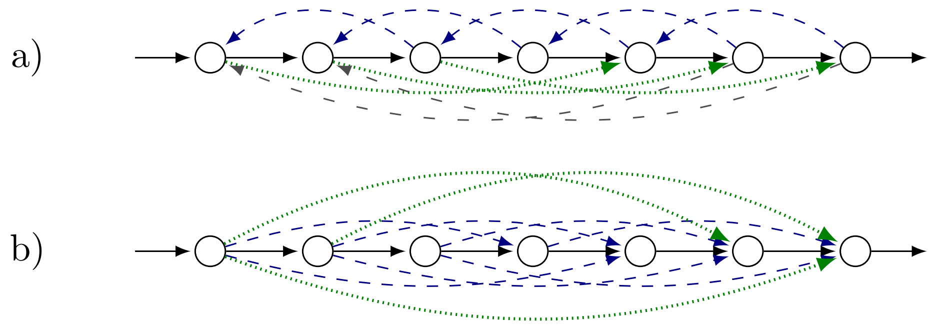

Fig. 4. Two closures of the linear structure.

Definition 4.8. We say that Φ *forks at the level* i if for every $\underline{s} \in \mathcal{T}_\infty$ with $|s| = i$ and for all $t, t' \in \{0,1\}^*$ there are no edges between $\underline{s0t}$ and $\underline{s1t'}$ in $\mathfrak{C}_\Phi(\mathcal{T}_\infty)$. By $\mathfrak{g}^\Phi_{nf}$ we denote the smallest number i such that Φ does not fork at the level i. If Φ forks at all the levels we say that Φ has the *tree-compatible model property* (TCMP) and put $\mathfrak{g}^\Phi_{nf} = \infty$.

The following tool plays a crucial role in our proofs.

Definition 4.9. A function f from $\mathcal{M}_1$ into $\mathcal{M}_2$ is a *morphism* iff for all worlds w, w' if $\mathcal{M}_1 \models wRw'$, then $\mathcal{M}_2 \models f(w)Rf(w')$.

The morphism between structures is also a morphism between their closures.

Observation 4.10. *Let $\mathcal{M}_1, \mathcal{M}_2$ be frames and f be a function from $\mathcal{M}_1$ into $\mathcal{M}_2$. If f is a morphism from $\mathcal{M}_1$ into $\mathcal{M}_2$, then f is a morphism from $\mathfrak{C}_\Phi(\mathcal{M}_1)$ into $\mathfrak{C}_\Phi(\mathcal{M}_2)$.*

It is not hard to see that if Φ has TCMP, then in all tree-based models of Φ there are no edges among the worlds from disjoint subtrees. Indeed, if there is an edge between two different subtrees $\mathcal{S}_1, \mathcal{S}_2$ of a model $\mathcal{M}$, one can define a morphism from $\mathcal{M}$ to $\mathcal{T}_\infty$ which maps $\mathcal{S}_1$ and $\mathcal{S}_2$ into disjoint subtrees of $\mathcal{T}$. This implies that some world above $\mathcal{S}_1$ and $\mathcal{S}_2$ does not fork, and Φ does not have TCMP.

We reveal why linear structures are important. In the tree-compatible case, along each path almost all worlds are connected as in the linear structure. The only exception is for the worlds that are close to the "ends" of the model.

We often use the morphism that maps the worlds of trees to the worlds of $\mathcal{L}_\mathbb{Z}$, $h_\mathcal{T} : \mathcal{T} \to \mathcal{L}_\mathbb{Z}$, defined as $h_\mathcal{T}(v) = \underline{i}$, where i is the length of the path from the root of $\mathcal{T}$ to v.

Lemma 4.11. *There exists a constant $\mathfrak{g}^\Phi_{in} < \infty$ such that for any tree $\mathcal{T}$ and any $\mathfrak{g}^\Phi_{in}$-inner worlds v_i, v_j at the same path there is an edge from v_i to v_j in $\mathfrak{C}_\Phi(\mathcal{T})$ iff there is an edge from $h_\mathcal{T}(v_i)$ to $h_\mathcal{T}(v_j)$ in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$.*

By applying the above lemma we prove the following.

Lemma 4.12. *If Φ satisfies S3, then it does not have the tree-compatible model property.*

Proof. Assume that Φ satisfies S3 for some $m > 0$. Let $k = \mathfrak{g}^\Phi_{in}$. By Lemma 4.11 we see that there are edges from $0^{k+(i+1)(m-1)}$ to $0^{k+i(m-1)}$ in $\mathfrak{C}_\Phi(\mathcal{T}_\infty)$ for every $i \geq 0$. Define $h : \mathcal{L}_\mathbb{Z} \to \mathfrak{C}_\Phi(\mathcal{T}_\infty)$ as $h(\underline{x}) = 0^{k-x(m-1)}$ for $x < 0$ and $h(\underline{x}) = 0^k1^x$ otherwise. Clearly h is a morphism, and by Observation 4.10 it is also a morphism from $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ to $\mathfrak{C}_\Phi(\mathcal{T}_\infty)$. Since in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ there is an edge from $\underline{1}$ to $\underline{1-m+1}$, there is also an edge from 0^k1 to $0^{k+m(m-1)}$ and therefore Φ does not fork on the level k. $\square$

Lemma 4.11 characterizes the edges in the closure between the worlds that are on the same path in a tree. If the formula has TCMP, then it characterizes all the edges in the closures. Below we study formulae that do not fork at some level. We start from two definitions.

Definition 4.13. Worlds w, w' of a frame $\mathcal{M}$ are *equivalent* if for all worlds u we have uRw iff uRw'.

Definition 4.14. We say that structures $\mathfrak{M} = \langle W, R, \pi \rangle$ and $\mathfrak{M}' = \langle W', R', \pi' \rangle$ are *indistinguishable* w.r.t. φ if (1) the sets of types realized in $\mathfrak{M}$ and $\mathfrak{M}'$ are the same, i.e., $\{tp^{\varphi}_{\mathfrak{M}}(w) : w \in W\} = \{tp^{\varphi}_{\mathfrak{M}'}(w) : w \in W'\}$; (2) types of the common worlds are the same, i.e., for all $w \in W \cap W'$ we have $tp^{\varphi}_{\mathfrak{M}}(w) = tp^{\varphi}_{\mathfrak{M}'}(w)$.

If $\mathfrak{M}$ and $\mathfrak{M}'$ are indistinguishable w.r.t. φ, then φ is globally satisfied in $\mathfrak{M}$ iff its globally satisfied in $\mathfrak{M}'$ and φ is locally satisfied in some world of $\mathfrak{M}$ iff its locally satisfied in some world of $\mathfrak{M}'$.

The following observation implies that each satisfiable formula has a model that does not contain a set of equivalent worlds of the cardinality greater than $|\varphi|$.

OBSERVATION 4.15. *Let $\mathfrak{M}$ be a model of φ over universe W and $X \subseteq W$ be a set of pairwise equivalent worlds. Then X can be divided into disjoint sets Y, Z such that $|Z| < |\varphi|$ and $\mathfrak{M}'$ and $\mathfrak{M}_{\restriction W \setminus Y}$ are indistinguishable w.r.t. φ.*

The proof is straightforward – we put in Z one world for each subformula of φ satisfied by some world in X. Clearly $|Z| < |\varphi|$.

We argue that if $\mathfrak{g}^{\Phi}_{nf} < \infty$, then in structures reachable from worlds at the level $\mathfrak{g}^{\Phi}_{nf}$ such equivalence is very common.

LEMMA 4.16. *Assume that Φ does not have TCMP. Then, there is a constant $\mathfrak{g}^{\Phi}_{eq} < \infty$ such that $\mathfrak{g}^{\Phi}_{eq} \geq \mathfrak{g}^{\Phi}_{nf}$ and for each tree $\mathcal{T}$ with a bounded degree, a world w at the level $\mathfrak{g}^{\Phi}_{eq}$ in $\mathfrak{C}_{\Phi}(\mathcal{T})$, and all i, all the $\mathfrak{g}^{\Phi}_{eq}$-followed descendants of w at the level $\mathfrak{g}^{\Phi}_{eq} + i$ are equivalent in the frame $\mathfrak{C}_{\Phi}(\mathcal{T})$.*

Example 4.17. Consider the formula $\Phi = \{\varphi_1, \varphi_2\}$, where $\varphi_1 = xRz \wedge zRy \Rightarrow yRy$ and $\varphi_2 = xRx \wedge xRy \wedge xRz \Rightarrow yRz$, and the tree in Fig. 3. The formula φ_1 enforces the following property: each world that has a predecessor that has a predecessor is reflexive. The formula φ_2 makes the relation R Euclidean except for the non-reflexive worlds. Formula Φ forks at the first two levels.

Let $P_{\mathcal{M}}(w)$ denote the set of predecessors of w in $\mathcal{M}$. We extend this notion to sets by defining $P_{\mathcal{M}}(X) = \bigcup_{w \in X} P_{\mathcal{M}}(w)$. We also extend these notions to Kripke structures in the obvious way. Finally, for $d > 0$ we define $P^d_{\mathcal{M}}(w) = \{v : v \in P_{\mathcal{M}}(w), v \text{ is } d\text{-inner in } \mathcal{M}\}$.

Consider a world w of a frame of the form $\mathfrak{C}_{\Phi}(\mathcal{T})$. The *original path* of w, denoted by $\overline{w}$ is the unique path from the root to w in $\mathcal{T}$. For an original path $\overline{w_i} = w_1 \ldots w_i$, we denote the world w_k by $\overline{w}_k$ and the world w_{i-k} by $\overline{w_{-k}}$.

A sequence is *b-bounded* if it consists of at most b numbers from $[0, b]$.

For a world w with $|\overline{w}| = i$, the *profile of* w, denoted as *prof*(w, b, g), is the following pair: $(bs(w, b), i \mod g)$, where $bs(w, b)$ denotes the b-bounded sequence $p_{\max(1,i-b)}, \ldots, p_i$ such that each p_k is the maximum number not greater than b such that $\overline{w}_k$ is p_k-followed in its model. We say that a pair (p, i) is a *possible* (b, g)*-profile* if p is a b-bounded sequence and $0 \leq i < g$. Note that for a given b and g, there are only finitely many different possible (b, g)-profiles.

We say that a triple (C, D, E) is a *predecessors generator* of a world w if $C, D, E \subseteq \{1, \ldots, |\overline{w}|\}$ and the following holds

$$P_{\mathfrak{M}}(w) = \{\overline{w}_d : d \in D\} \cup \{\overline{w}_{-d} : d \in E\} \cup \bigcup_{c \in C} P^b_{\mathfrak{M}}(\overline{w}_{-c})$$

The following lemma shows that the set of the predecessors of a given world can be computed based on the information of a bounded number of its predecessors.

LEMMA 4.18. *Let Φ be a formula satisfying S2 with χ. There is a constant $\mathfrak{g}^{\Phi}_{bs} < \infty$ such that for each possible $(\mathfrak{g}^{\Phi}_{bs}, gcd(\chi))$-profile l there is a predecessors generator $\mathbb{PG}^{\Phi}_l$ such that for each world w of a model $\mathfrak{C}_{\Phi}(\mathcal{T})$ s.t. $|\overline{w}| > \mathfrak{g}^{\Phi}_{bs}$ and $l = prof(w, \mathfrak{g}^{\Phi}_{bs}, gcd(\chi))$, $\mathbb{PG}^{\Phi}_l$ is a predecessors generator of w.*

Finally, we define $\mathfrak{g}^{\Phi}_{pre}$ as the maximum of $\mathfrak{g}^{\Phi}_{bs}$ and the maximal number that appears in all the predecessors generators of the form $\mathbb{PG}^{\Phi}_l$ for all the possible profiles $prof(w, \mathfrak{g}^{\Phi}_{bs}, gcd(\chi))$. The intuitive meaning of $\mathfrak{g}^{\Phi}_{pre}$ is as follows. If an algorithm keeps in memory the first $\mathfrak{g}^{\Phi}_{pre}$ and the last $\mathfrak{g}^{\Phi}_{pre}$ worlds of $\underline{w}$, then for each predecessor v of w it keeps v or some successor of v different than w. We will use it in the algorithm to deal with the subformulae such as $p \to \Box\neg q$.

To simplify the notation, we define $\mathfrak{g}^{\Phi} = \max\{\mathfrak{g}^{\Phi}_{pre}, \mathfrak{g}^{\Phi}_{bs}, \mathfrak{g}^{\Phi}_{in}, \mathfrak{g}^{\Phi}_{nf}, \mathfrak{g}^{\Phi}_{eq}\}$.

4.4. Formulae with the tree-compatible model property

This subsection contains a high-level description of algorithms for the formulae satisfying the tree-compatible model property. The efficient algorithms are, however, quite complicated. To help the reader understand the key ideas, we presented here simpler (inefficient) algorithms, and we do not discuss some details. The detailed discussion and the efficient algorithms are given in Section 6.

Assume that Φ has the tree-compatible model property and let φ be a modal formula. By Lemma 4.11, if φ has a model, then there is a tree $\mathcal{T}$ such that φ has a model based on $\mathfrak{C}_{\Phi}(\mathcal{T})$. By Lemma 4.12, we know that Φ satisfies S1 or S2, i.e., $\mathfrak{C}_{\Phi}(\mathcal{T})$ does not contain long backward edges.

The algorithms solving Φ-SAT$_L$ and Φ-SAT$_G$ are alternating algorithms, which try to construct such a model in a top-down fashion. Below we describe only the algorithm for the local satisfiability problem.

The alternating algorithm computes a list of formulae that have to be satisfied (initially the list contains φ) by a current world, guesses types of its successors, checks if the world satisfies all the formulae, and then calls itself universally for all the successors.

To construct the list of formulae that have to be satisfied in a world, we need to know the list of the types of predecessors of this world. Recall that the algorithm attempts to construct a model without edges between different branches and long backward edges. A naive solution would be to keep the whole path from the root to the current node along with the guessed types of these worlds. We observe that this can be avoided.

To define the accepting condition, we simply add a counter that counts the number of universal calls (which is equal to the depth of the current node) and simply accept if the value of the counter is large enough to ensure that some configuration has been visited twice. Below we argue that the algorithm can be implemented in such a way that it uses only $O(2^{|\varphi|})$ memory.

If Φ satisfies S1, then it is enough to remember only one world. Below we focus on the formulae that satisfy S2.

Assume that Φ satisfies S2. The algorithm first guesses the first $\mathfrak{g}^{\Phi}$ levels of the constructed model. Then, consider a model $\mathfrak{C}_{\Phi}(\mathcal{T})$ and a path $v_1 \ldots v_s$ in $\mathcal{T}$. The algorithm guesses the type, the number of the successors and the profile of v_s. Next, it computes the set A of types of the predecessors of v_s. By Lemma 4.18 it can be computed given the types of initial worlds $v_1, \ldots, v_a$, recent words $v_{s-b}, \ldots, v_{s-1}$ and the sets of types of $\mathfrak{g}^{\Phi}_{in}$-inner predecessors of recent worlds $v_{s-c}, \ldots, v_{s-1}$, where a, b, c depend only on Φ, and not on s. Finally, the algorithm verifies whether the type of v_s is consistent with A and recursively calls itself on successors of v_s (and verifies that the profile of v_s has been guessed correctly). The described algorithm requires only $O((a+b)\cdot|\varphi| + c \cdot 2^{|\varphi|})$ memory.

4.5. Formulae without the tree-compatible model property

Let Φ be a formula without the tree-compatible model property. Recall that two worlds w, w' of a frame $\mathcal{M}$ are *equivalent* if for each world u we have uRw iff uRw'. We are going to exploit the property guaranteed by Lemma 4.16. We start with the observation that if we have two equivalent worlds v, w with the same types, then we can remove one of them.

OBSERVATION 4.19. *Let $\mathfrak{M} = \langle W, R, \pi\rangle$ be a $\mathcal{K}_\Phi$-based (global) model of a modal formula φ. If for all subformulae ψ of φ satisfied by a world w there is a world $w' \neq w$ of $\mathfrak{M}$ such that w, w' are equivalent and w' satisfies ψ, then for $W' = W \setminus \{w\}$, $\mathfrak{M}_{\restriction W'}$ is a $\mathcal{K}_\Phi$-based (global) model of φ.*

The proof is straightforward — the types of the remaining worlds do not change.

Let $\mathfrak{M}$ be a tree-based model based on the frame $\mathfrak{C}_\Phi(\mathcal{T})$. We denote by *level i of* $\mathfrak{M}$ the set of worlds from $\mathfrak{M}$ such that the length of the path from root to w in $\mathcal{T}$ (notice that $\mathcal{T}$ is a tree) is equal i.

OBSERVATION 4.20. *Let φ be a formula satisfiable over $\mathcal{K}_\Phi$. Then there is a $\mathcal{K}_\Phi$-based model $\mathfrak{M}$ of φ such that the size of each level of $\mathfrak{M}$ is bounded polynomially in $|\varphi|$.*

First, observe that the number of worlds at level i can be bounded by $|\varphi|^i$ because Lemma 4.3 guarantees that the degree of the tree is bounded by $|\varphi|$. It follows from Lemma 4.16 that for all worlds w at the level $\mathfrak{g}^\Phi$ and all $i > 2\mathfrak{g}^\Phi$, all descendants of w at the level i are equivalent. Therefore, by Observation 4.15, we can remove all but $|\varphi|$ of them. Since the number of worlds at the level $\mathfrak{g}^\Phi$ can be bounded by $|\varphi|^{\mathfrak{g}^\Phi}$, the number of worlds at the level $i > 2\mathfrak{g}^\Phi$ can be bounded by $|\varphi|^{\mathfrak{g}^\Phi} \cdot |\varphi|$, so polynomially in $|\varphi|$.

Observation 4.20 says that we can reduce the number of worlds needed at each level and bound them by some polynomial of $|\varphi|$. The existence of such models can be verified by a non-deterministic machine working in polynomial space that first guesses first $2\mathfrak{g}^\Phi$ levels, and then recursively guesses and verifies the consecutive levels, similarly to the tree-compatible case. Since the number of worlds needed at each level can be bounded polynomially in $|\varphi|$, such an algorithm works in NPSPACE=PSPACE [Savitch 1970]. We conclude that the satisfiability problem is in PSPACE. This ends the proof of Theorem 1.2. However, it does not lead to the optimal complexity.

5. PROPERTIES OF CLOSURES

In this section we prove the lemmas. As in the previous section, we fix a formula $\Phi \in \mathsf{UHF}$. By n we denote the number of variables in Φ.

5.1. Proof of Lemma 4.3

LEMMA 4.3. *If a modal formula φ has a $\mathcal{K}_\Phi$-based model (global model), then there exists a tree $\mathcal{T}$ with the degree bounded by $|\varphi|$ and a labelling $\pi_\mathcal{T}$, such that*

(i) $\langle \mathcal{T}, \pi_\mathcal{T}\rangle$ is a model (resp. global model) of φ;
(ii) $\langle \mathfrak{C}_\Phi(\mathcal{T}), \pi_\mathcal{T}\rangle$ is a model (resp. global model) of φ that satisfies Φ.

PROOF. Assume that there exist $\mathfrak{M} = \langle W, R, \pi\rangle$ and $u_0 \in W$ such that $\mathfrak{M} \models \Phi$ and $\mathfrak{M}, u_0 \models \varphi$. We construct $\mathfrak{M}_0 = \langle \mathcal{T}, \pi_\mathcal{T}\rangle$, where $\mathcal{T} = \langle W_0, R_0\rangle$, by unravelling of $\mathfrak{M}$ as follows. The universe W_0 consists of finite sequences of elements of W. We define W_0 and R_0 inductively. Initially, we put $(u_0) \in W_0$. Assume that $(u_0, \ldots, u_k) \in W_0$. Let $\Diamond\psi_1, \ldots, \Diamond\psi_s$ be all formulae of the form $\Diamond\psi$ from $tp^{\varphi}_{\mathfrak{M}}(u_k)$. There exist $u^1_{k+1}, \ldots, u^s_{k+1} \in W$, such that for every $i \in \{1, \ldots, s\}$ we have $\mathfrak{M} \models u_k R u^i_{k+1}$ and $\psi_i \in tp^{\varphi}_{\mathfrak{M}}(u^i_{k+1})$. For each such i we put $(u_0, \ldots, u_k, u^i_{k+1})$ into W_0 and add $((u_0, \ldots, u_k), (u_0, \ldots, u_k, u^i_{k+1}))$ to R_0. We define $\pi_{\mathcal{T}}$ as $\pi_{\mathcal{T}}((u_0, \ldots, u_k)) = \pi(u_k)$. Observe that $\mathcal{T}$ is a tree in which the degree of every node is bounded by $|\varphi|$.

Let $f : W_0 \mapsto W$ be defined as $f((u_0, \ldots, u_k)) = u_k$. Observe that:

(i) for every $\vec{u} = (u_0, \ldots, u_k) \in W_0$ we have $\pi_{M_0}(\vec{u}) = \pi_M(f(\vec{u}))$,
(ii) for all $\vec{u_1}, \vec{u_2} \in W_0$, $\mathfrak{M}_0 \models \vec{u_1} R \vec{u_2}$ implies $\mathfrak{M} \models f(\vec{u_1}) R f(\vec{u_2})$, and
(iii) for all $\vec{u} = (u_0, \ldots, u_k) \in W_0$, $u_{k+1} \in W$ if $\mathfrak{M} \models u_k R u_{k+1}$, then there exists $u'_{k+1} \in W$ such that (1) $tp^{\varphi}_{\mathfrak{M}}(u_{k+1}) = tp^{\varphi}_{\mathfrak{M}}(u'_{k+1})$ and (2) $\mathfrak{M}_0 \models (u_0, \ldots, u_k) R (u_0, \ldots, u_k, u'_{k+1})$.

Having the conditions *(i), (ii)* and *(iii)*, one can show, by induction on the modal depth of a formula, that for every $\vec{u} \in W_0$, $tp^{\varphi}_{\mathfrak{M}_0}(\vec{u}) = tp^{\varphi}_{\mathfrak{M}}(f(\vec{u}))$. This implies that $\mathcal{T}, (u_0) \models \varphi$.

Observe that if $\mathfrak{M}_0$ is substituted by $\mathfrak{C}_\Phi(\mathfrak{M}_0)$, the conditions *(i)* and *(iii)* are preserved. The condition *(ii)* states that f is a morphism from $\mathfrak{M}_0$ to $\mathfrak{M}$. Due to Observation 4.10, f is a morphism of $\mathfrak{C}_\Phi(\mathfrak{M}_0)$ to $\mathfrak{C}_\Phi(\mathfrak{M})$. Since $\mathfrak{M}$ satisfies Φ, $\mathfrak{C}_\Phi(\mathfrak{M}) = \mathfrak{M}$ and f is a morphism of $\mathfrak{C}_\Phi(\mathfrak{M}_0)$ to $\mathfrak{M}$. Therefore, condition *(ii)* is preserved as well. Thus, $\mathfrak{C}_\Phi(\mathfrak{M}_0)$ satisfies φ.

It remains to be shown that $\mathfrak{C}_\Phi(\mathfrak{M}_0)$ satisfies Φ. By definition $\mathfrak{C}_\Phi(\mathcal{M}_0)$ satisfies every $\Psi_1 \Rightarrow \Psi_2$ from Φ^+. Suppose that $\mathfrak{C}_\Phi(\mathcal{M}_0)$ does not satisfy $\Psi \Rightarrow \bot$ from Φ; there are some $\vec{w_1}, \vec{w_2}, \vec{v_1}, \vec{v_2}, \ldots, \vec{v_k}$ such that $\mathfrak{C}_\Phi(\mathcal{M}_0) \models \Psi(\vec{w_1}, \vec{w_2}, \vec{v_1}, \vec{v_2}, \ldots, \vec{v_k})$. But then $\mathfrak{M} \models \Psi(f(\vec{w_1}), f(\vec{w_2}), f(\vec{v_1}), f(\vec{v_2}), \ldots, f(\vec{v_k}))$. This contradicts the assumption that $\mathfrak{M} \models \Phi$. $\square$

5.2. Proof of Lemma 4.7

LEMMA 4.7. *Φ satisfies S1, S2, or S3.*

We use the following fact that can be easily proved by employing Euclidian algorithm.

FACT 5.1. *Assume that X is a (possibly infinite) set of positive numbers that is closed under addition. Then, there exists a finite subset X' of X such that $gcd(X) = gcd(X')$. Moreover, for each $x > lcm(X')$, $gcd(X')$ divides x iff $x \in X$.*

The following definition will prove useful in the sequel. For $s \in \mathbb{Z}$ and frames $\mathcal{M}_1, \mathcal{M}_2$ over the universe $\{\underline{i} : i \in \mathbb{Z}\}$ we define *the shift* $sh_s : \mathcal{M}_1 \mapsto \mathcal{M}_2$ as $sh_s(\underline{i}) = \underline{i+s}$. Let $\mathcal{M}$ be a frame over the universe $\{\underline{i} : i \in \mathbb{Z}\}$ that contains $\mathcal{L}_\mathbb{Z}$. We say that $\mathcal{M}$ is *uniform* if for every $s \in \mathbb{Z}$ the shift sh_s is an automorphism of $\mathcal{M}$. We say that $\mathcal{M}$ is *closed under composition* iff for every world $\underline{i}$, positive k and $a_1, a_2, \ldots, a_k \in \mathbb{Z}$:

(1) if $\mathcal{M} \models \underline{i} R \underline{i+k}$ and $\mathcal{M} \models \underline{i} R \underline{i+a_1} \wedge \ldots \wedge \underline{i+a_{k-1}} R \underline{i+a_k}$ then $\mathcal{M} \models \underline{i} R \underline{i+a_k}$, and
(2) if $\mathcal{M} \models \underline{i} R \underline{i-k}$ and $\mathcal{M} \models \underline{i+a_k} R \underline{i+a_{k-1}} \wedge \ldots \wedge \underline{i+a_1} R \underline{i}$ then $\mathcal{M} \models \underline{i} R \underline{i+a_k}$.

A frame $\mathcal{N}$ is the *composite closure* of $\mathcal{M}$ iff $\mathcal{N}$ is the least (w.r.t. the relation R) closed under composition frame containing $\mathcal{M}$.

For every $s \in \mathbb{Z}$, sh_s is an automorphism of $\mathcal{L}_\mathbb{Z}$ onto itself, therefore by Observation 4.10, sh_s is a morphism of $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ onto itself. Hence, $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ is uniform. As $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ is closed under consequences and it is uniform, one can easily show that $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ is closed under composition. In consequence we have the following:

LEMMA 5.2. *Frame $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ is uniform and closed under composition.*

For any uniform and closed under composition $\mathcal{N}$ we define $X^{\mathcal{N},i}_\Phi = \{a : \mathcal{N} \models \underline{i} R \underline{i+a+1}\}$ and $X^{\mathcal{N}}_\Phi = X^{\mathcal{N},0}_\Phi$. For every i the function sh_i is an automorphism of $\mathcal{N}$, hence $X^{\mathcal{N},i}_\Phi = X^{\mathcal{N}}_\Phi$.

LEMMA 5.3. *Let $\mathcal{N}$ be a uniform frame closed under composition.*

(i) If $x, y \in X_\Phi^\mathcal{N}$ and $x \geq 0$, then $x + y \in X_\Phi^\mathcal{N}$.
(ii) If $x, y \in X_\Phi^\mathcal{N}$ and $s, x \geq 0$, then $x + sy \in X_\Phi^\mathcal{N}$.
(iii) For every $a > 2$, if $-a \in X_\Phi^\mathcal{N}$, then $a \in X_\Phi^\mathcal{N}$.
(iv) If $X_\Phi^\mathcal{N}$ contains a positive number, then for all $a \geq 0$, if $-a \in X_\Phi^\mathcal{N}$, then $a \in X_\Phi^\mathcal{N}$.

PROOF. For the (i) part, observe that if $\mathcal{N} \models \underline{0}R\underline{x+1}$ and $\mathcal{N} \models \underline{0}R\underline{y+1}$, where $x > 0$, then $\mathcal{N} \models (\underline{0}R\underline{y+1}) \wedge (\underline{y+1}R\underline{y+2}) \wedge \ldots \wedge (\underline{i+y+x}R\underline{i+y+x+1})$. Hence, as composite closure, $\mathcal{N}$ satisfies $\underline{i}R\underline{i+x+y+1}$ and $x + y \in X_\Phi^\mathcal{N}$. Property (ii) follows from a straightforward induction based on (i).

For (iii), note that if $\mathcal{N} \models \underline{0}R\underline{-(a-1)}$ where $a > 2$, then due to uniformity we have in $\mathcal{N}$ that $\underline{a+1}R\underline{2}$ and $\underline{2}R\underline{-a+3}$. Therefore, by composition of $\mathcal{N} \models \underline{0}R\underline{-(a-1)}$ and $\mathcal{N} \models \underline{a+1}R\underline{2} \wedge \underline{2}R\underline{-a+3} \wedge \underline{-a+3}R\underline{-a+2} \wedge \underline{-a+2}R\underline{-a+1} \wedge \cdots \wedge \underline{-1}R\underline{0}$ we get $\mathcal{N} \models \underline{0}R\underline{a+1}$.

Finally, for (iv), let $b \in X_\Phi^\mathcal{N}$ be a positive number and $a \geq 0$. If $a = 1$, then by Property (ii) for $x = b$, $y = -1$ and $s = b - 1$ we have $a \in X_\Phi$. If $a = 2$, then consider two cases. If b is odd, we use Property (ii) for $x = b$, $y = -a$ and $s = (b+1)/2$ to show that $1 \in X_\Phi^\mathcal{N}$ and, by (i), that $2 = 1 + 1 \in X_\Phi^\mathcal{N}$. If b is even, use Property (ii) for $x = b$, $y = -a$ and $s = b/2 - 1$ to show that $2 \in X_\Phi^\mathcal{N}$. If $a > 2$, then the statement follows from Property (iii). □

Now we consider the case when $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ contains long backward edges.

LEMMA 5.4. *Suppose that a uniform, closed under composition frame $\mathcal{N}$ contains a backward edge and a long edge. There exists a finite subset X' of $X_\Phi^\mathcal{N}$ such that $gcd(X') = gcd(X_\Phi^\mathcal{N})$ and, for every b, $\mathcal{N} \models \underline{i}R\underline{i+b+1}$ iff $gcd(X')$ divides b.*

PROOF. Due to Lemma 5.3 (iv), for every $b > 0$, $-b \in X_\Phi^\mathcal{N}$ implies $b \in X_\Phi^\mathcal{N}$, therefore $gcd(X_\Phi^\mathcal{N}) = gcd(X^+)$, where X^+ contains all positive elements from $X_\Phi^\mathcal{N}$. By Fact 5.1, there exists finite X' such that for every $b > lcm(X')$ we have $\mathcal{N} \models \underline{i}R\underline{i+b+1}$ iff $gcd(X')$ divides b. We show that for every $b \in \mathbb{Z}$, we have $b \in X_\Phi^\mathcal{N}$ iff $gcd(X')$ divides b.

Of course, all elements of $X_\Phi^\mathcal{N}$ are divisible by $gcd(X')$. Suppose that $b \in \mathbb{Z}$ is divisible by $gcd(X')$. Since $\mathcal{N}$ contains backward edges and long edges, there exists $a_j \in X_\Phi^\mathcal{N}$ such that $a_j < -1$. From Lemma 5.3 we know that $-a_j \in X_\Phi^\mathcal{N}$ and, moreover, for any $s \geq 0$, $b + s \cdot (-a_j) \in X_\Phi^\mathcal{N}$. Let s be such that $b' = b + s \cdot (-a_j) > lcm(X')$. Since $gcd(X')$ divides b and $-a_j$, $b' \in X_\Phi^\mathcal{N}$. Due to additivity of $X_\Phi^\mathcal{N}$, we conclude that $b = (b' + s \cdot a_j) \in X_\phi^\mathcal{N}$. □

We are ready to prove Lemma 4.7. If $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ contains no long edges, then Φ satisfies S1 and we are done. If $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ contains only forward edges, then we define χ to be a minimal set such that $X_\Phi^{\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})}$ is the additive closure of χ (it exists thanks to Fact 5.1). If $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ contains long and backward edges, then Lemma 5.4 implies that there is some m such that for all worlds $\underline{i}$, $\underline{i+b}$, there is an edge from $\underline{i}$ to $\underline{i+b}$ if and only if m divides $|b-1|$.

5.3. Proof of Lemma 4.11

LEMMA 4.11. *There exists a constant $\mathfrak{g}_{in}^\Phi < \infty$ such that for any tree $\mathcal{T}$ and any $\mathfrak{g}_{in}^\Phi$-inner worlds v_i, v_j at the same path there is an edge from v_i to v_j in $\mathfrak{C}_\Phi(\mathcal{T})$ iff there is an edge from $h_\mathcal{T}(v_i)$ to $h_\mathcal{T}(v_j)$ in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$.*

For each $s \in \mathbb{N}_\infty$, we define $\mathcal{I}_s = \mathcal{L}_{\mathbb{Z}\restriction W_s}$, where $W_s = \{\underline{i} : 0 \leq i < s\}$. Lemma 4.11 follows from the following.

LEMMA 5.5. *There exists $\mathfrak{g}_{in}^\Phi$, that depends only on Φ, such that for every $s > 2 \cdot \mathfrak{g}_{in}^\Phi$ if u, v are $\mathfrak{g}_{in}^\Phi$-inner in $\mathfrak{C}_\Phi(\mathcal{I}_s)$, then $\mathfrak{C}_\Phi(\mathcal{I}_s) \models uRv$ iff $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z}) \models uRv$.*

Indeed, the proof of "⇒" is a simple application of Observation 4.10 to the morphism $h_{\mathcal{T}}$. For the "⇐", let $v_0, v_1, \dots$ be a path containing v_i and v_j such that v_i, v_j are $\mathfrak{g}_{in}^{\Phi}$-inner in this path and let $s \in \mathbb{N}_\infty$ be the length of this path. Then this path is isomorphic with $\mathcal{I}_s$. Due to Lemma 5.5 there is an edge from $h_{\mathcal{T}}(v_i)$ to $h_{\mathcal{T}}(v_j)$ in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ if and only if there is such an edge in $\mathfrak{C}_\Phi(\mathcal{I}_s)$, and therefore there is such an edge in $\mathfrak{C}_\Phi(\mathcal{T})$.

The rest of this subsection is devoted to the proof of Lemma 5.5. We prove it in several steps. In Lemma 5.6 we show some basic properties of the closures of frames $\mathfrak{C}_\Phi(\mathcal{I}_s)$. Then, we define the notion of *regularity*. In Lemma 5.8 we prove that any closed under composition, uniform, regular frame $\mathcal{M}$ has the following property: for any finite sequence of worlds, we can find another sequence of worlds such that the subframes defined by these sequences are isomorphic, and the latter sequence consists of worlds $\underline{i_1}, \dots, \underline{i_k}$ such that the numbers $i_1, \dots, i_k$ are contained in an interval bounded by some function depending only on k and Φ. Roughly speaking, it means that each configuration of worlds that appears in $\mathcal{M}$ appears also with worlds fairly close to one another, so an edge $(\underline{i}, \underline{j})$ in the closure can be obtained by applying the closure operator to worlds that are "close" to $\underline{i}$ or $\underline{j}$.

Then we define the *approximations* of the frame $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ — an infinite sequence of frames that converge to $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$. In Lemma 5.9 we show that all the elements of this sequence (except for, perhaps, a few initial ones) are regular, and that the sequence stabilise after a finite index. This allows us to prove Lemma 5.5 by induction on this sequence.

LEMMA 5.6. *Let $s \in \mathbb{N}_\infty$.*

(i) Let k be the maximal number such that v is k-preceded (k-followed) in $\mathcal{I}_s$. If v is $k+1$-preceded (resp. $k+1$-followed) in $\mathfrak{C}_\Phi(\mathcal{I}_s)$, then it is ∞-preceded (resp. ∞-followed) in $\mathfrak{C}_\Phi(\mathcal{I}_s)$.

(ii) If $\underline{i}, \underline{i+j}$ are ∞-inner in $\mathfrak{C}_\Phi(\mathcal{I}_s)$, then there is a morphism f from $\mathcal{L}_\mathbb{Z}$ into $\mathfrak{C}_\Phi(\mathcal{I}_s)$ such that $f(\underline{i}) = \underline{i}$ and $f(\underline{i+j}) = \underline{i+j}$.

For (i) of Lemma 5.6, we observe that increase of k in k-preceded (resp., k-followed) case involves formation of a cycle. For (ii) one can easily construct such a morphism from $\mathcal{L}_\mathbb{Z}$ into $\mathfrak{C}_\Phi(\mathcal{I}_s)$ and extend it to a morphism from $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ into $\mathfrak{C}_\Phi(\mathcal{I}_s)$.

Definition 5.7. A frame $\mathcal{M}$ over the universe $\{\underline{i} : i \in \mathbb{Z}\}$ is t,p-regular for $t,p \in \mathbb{N}$ if for all $b > 0$ (i) if $\mathcal{M} \models \underline{i}R\underline{i+b+1}$ then p divides b; and (ii) if $b > t$ and p divides b, then $\mathcal{M} \models \underline{i}R\underline{i+b+1}$. We say that $\mathcal{M}$ is *regular* if it is t,p-regular for some t,p. The *period* of a regular frame $\mathcal{M}$, denoted by $pi(\mathcal{M})$, and the *threshold* of $\mathcal{M}$, denoted by $tr(\mathcal{M})$, are the numbers such that $\mathcal{M}$ is $tr(\mathcal{M}), pi(\mathcal{M})$-regular and for all p,t such that $\mathcal{M}$ is t,p-regular, $p = pi(\mathcal{M})$ and $t \geq tr(\mathcal{M})$.

LEMMA 5.8. *Let $\mathcal{M}$ be a closed under composition, uniform, regular frame and $u_1, \dots, u_k$ be a sequence of integers. There is a sequence $v_1, \dots, v_k$ such that for every $i,j \in \{1, \dots, k\}$*

(1) $\mathcal{M} \models \underline{u_i}R\underline{u_j}$ iff $\mathcal{M} \models \underline{v_i}R\underline{v_j}$,
(2) $u_i \leq u_j$ iff $v_i \leq v_j$, and
(3) $|v_i - v_j| \leq k \cdot tr(\mathcal{M}) + pi(\mathcal{M})$.
(4) $|v_i - v_j| = |u_i - u_j| \mod pi(\mathcal{M})$ and $|v_i - v_j| \geq tr(\mathcal{M})$ iff $|u_i - u_j| \geq tr(\mathcal{M})$.

Moreover, for any $a,b \in \{1, \dots, k\}$ there is a sequence $v'_1, \dots, v'_k$ such that $v'_a = u_a$, $v'_b = u_b$ and for every $i,j \in \{1, \dots, k\}$ $\mathcal{M} \models \underline{u_i}R\underline{u_j}$ iff $\mathcal{M} \models \underline{v'_i}R\underline{v'_j}$ and $min(|v'_i - v'_a|, |v'_i - v'_b|) < (k-1) \cdot tr(\mathcal{M}) + pi(\mathcal{M})$.

PROOF. Without loss of generality, assume that the sequence $u_1, \dots, u_k$ is ascending. Let j be such $u_{j+1} - u_j > tr(\mathcal{M}) + pi(\mathcal{M})$. Define $u'_{j+i} = u_{j+i} - pi(\mathcal{M})$ for all $1 \leq i \leq k-j$.

Since $\mathcal{M}$ is uniform, the relations among $\underline{u'_{j+1}}, \ldots, \underline{u'_n}$ are the same as the relations among $\underline{u_{j+1}}, \ldots, \underline{u_n}$.

Consider $i \leq j$ and $i' > j$. As $u'_{i'} - u_i - 1 \geq tr(\mathcal{M})$, the following statements are equivalent:

(i) there is an edge between $\underline{u'_{i'}}$ and $\underline{u_i}$;
(ii) $pi(\mathcal{M})$ divides $u'_{i'} - u_i - 1$;
(iii) $pi(\mathcal{M})$ divides $u'_{i'} - u_i - 1 + pi(\mathcal{M}) = u_{i'} - u_i - 1$;
(iv) there is an edge between $\underline{u_{i'}}$ and $\underline{u_i}$.

Clearly, $u_1, \ldots, u_j, u'_{j+1}, \ldots, u'_k$ satisfies (1) and (2), and $u'_{j+1} - u_j < u_{j+1} - u_j$. By iterating the operation defined above finitely many times we obtain sequences that satisfy the required properties. □

Let $W = \{\underline{i} : i \in \mathbb{Z}\}$ be the domain of $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ and R be any set of edges over W. We define R^s as a number such that $\langle \underline{0}, \underline{R^s} \rangle \in Cons_\Phi(R) \setminus R$ and for every b, if $\langle \underline{0}, \underline{b} \rangle \in Cons_\Phi(R) \setminus R$, then $|b| > |R^s|$ or $|b| = |R^s|$ and $b \leq R^s$. If $Cons_\Phi(R) \setminus R = \emptyset$, we put $R^s = 1$.

We define the *sequence of approximations* $\mathcal{N}_0, \mathcal{N}_1, \ldots$ of $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ in the following way:

— $\mathcal{N}_0 = \mathcal{L}_\mathbb{Z} = \langle W, R_0 \rangle$ and
— $\mathcal{N}_{p+1} = \langle W, R_{p+1} \rangle$ is the composite closure of $\langle W, R_p \cup \{(\underline{i}, \underline{i + R^s_p}) : i \in \mathbb{Z}\} \rangle$.

Clearly, the limit of the sequence $\mathcal{N}_0, \mathcal{N}_1, \ldots$ is $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ and all approximations are closed under composition and uniform.

LEMMA 5.9. *Let $\mathcal{N}_0, \mathcal{N}_1, \ldots$ be the sequence of approximations of $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$. For every n, if $\mathcal{N}^\Phi_p$ contains long edges, then $\mathcal{N}^\Phi_p$ is regular. Moreover, the sequence stabilises after some finite index.*

PROOF. We prove the lemma by induction.
The induction base. If $\mathcal{N}_p$ is the first frame containing a long edge, then $\mathcal{N}_{p-1}$ is equal to $\mathcal{L}_\mathbb{Z}$ or a reflexive or symmetric closure (or both) of $\mathcal{L}_\mathbb{Z}$. A quick check shows that in this case both the threshold and period of $\mathcal{N}^\Phi_p$ are bounded by n.
The induction step. Assume that $R^s_p > 0$. First, we show that $R^s_p \in [1, n \cdot tr(\mathcal{N}_p) + pi(\mathcal{N}_p) - 1]$. Suppose that $R^s_p > n \cdot tr(\mathcal{N}_p) + pi(\mathcal{N}_p)$. The edge $(\underline{0}, \underline{R^s_p})$ is implied by the formula Φ applied to some worlds $\underline{u_1}, \ldots, \underline{u_n}$ with $u_s = 0$ and $u_t = R^s_p$ for some s, t. Let $v_1, \ldots, v_n$ be a result of application of Lemma 5.8 to the sequence $u_1, \ldots, u_n$ and the frame $\mathcal{N}_p$. Since the connections among $\underline{v_1}, \ldots, \underline{v_n}$ are the same as among $\underline{u_1}, \ldots, \underline{u_n}$, the edge $(\underline{v_s}, \underline{v_t})$ is a consequence of Φ in $\mathcal{N}_p$. If $(\underline{v_s}, \underline{v_t})$ was an edge in $\mathcal{N}_p$, then, by Property 4 of Lemma 5.8, $(\underline{0}, \underline{R^s_p})$ would be an edge as well. But it is not the case by the definition of R^s_p. The existence of the edge $(\underline{v_s}, \underline{v_t})$ contradicts the minimality of R^s_p. Therefore $R^s_p \leq n \cdot tr(\mathcal{N}_p) + pi(\mathcal{N}_p)$.

Let $j = R^s_p - 1$. If j is divisible by $pi(\mathcal{N}_i)$, then $\mathcal{N}_{p+1}$ is $tr(\mathcal{N}_p), pi(\mathcal{N}_p)$-regular. Assume otherwise, that j is not divisible by $pi(\mathcal{N}_p)$. Let $P = gcd(pi(\mathcal{N}_p), j)$ and $T = tr(\mathcal{N}_p) + j \cdot pi(\mathcal{N}_p)$. Every number divisible by P from the interval $[0, pi(\mathcal{N}_p) - 1]$ is the remainder of some number from $\{j, 2 \cdot j, \ldots, pi(\mathcal{N}_p) \cdot j\}$ from division by $pi(\mathcal{N}_p)$. Hence, every number $b \geq 0$ divisible by P is equal to $\alpha j + \beta pi(\mathcal{N}_p)$ where $\alpha \in \{1, \ldots, pi(\mathcal{N}_p)\}$ and $\beta \in \mathbb{Z}$. If $b > T$, then $\beta pi(\mathcal{N}_p) > tr(\mathcal{N}_p)$ and $\beta pi(\mathcal{N}_p) \in X^{\mathcal{N}_p}_\Phi$. Thus, for every $b \geq T$, if P divides b, then $b \in X^{\mathcal{N}_{p+1}}_\Phi$. On the other hand, the frame $\mathcal{M} = \langle \{\underline{i} : i \in \mathbb{Z}\}, \{(\underline{i}, \underline{i + \alpha P}) : i \in \mathbb{Z}, \alpha \geq 0\} \rangle$ is closed under composition and it contains the frame $\mathcal{N}_p$ and the edges $\{(\underline{i}, \underline{i + j + 1}) : i \in \mathbb{Z}\}$. This implies that $\mathcal{N}_{p+1}$ is contained in $\mathcal{M}$ and if $b \in X^{\mathcal{N}_{p+1}}_\Phi$, then P divides b. Hence, the threshold of $\mathcal{N}_{p+1}$ is bounded by T and the period is equal P.

Now assume that $R^s_p < 0$. Then, $\mathcal{N}_{p+1}$ satisfies the condition of Lemma 5.4, i.e., there exists d such that $\underline{i} R_{p+1} \underline{j}$ iff $d \mid j - 1 - i$. Of course, d is the period of $\mathcal{N}_{p+1}$ and its threshold is equal 0. Moreover, since d divides all $b \in X_\Phi^{\mathcal{N}_p}$ and $X_\Phi^{\mathcal{N}_p}$ contains a number from the interval $[2, n]$ we have $d \leq n$.

Clearly, for all i we have $pi(\mathcal{N}_i) \geq pi(\mathcal{N}_{i+1})$, i.e., the period never increases. Furthermore, if $pi(\mathcal{N}_i) = pi(\mathcal{N}_{i+1})$, then $tr(\mathcal{N}_i) = tr(\mathcal{N}_{i+1})$. Since $|R_p^s| < n \cdot tr(\mathcal{N}_p) + pi(\mathcal{N}_p)$, there is only a finitely many edges that can be added before the period changes. Therefore, the sequence $\mathcal{N}_0, \mathcal{N}_1, \ldots$ stabilises after finitely many steps. □

PROOF OF LEMMA 5.5. Clearly, for every s and all $u, v \in \mathcal{I}_s$, $\mathfrak{C}_\Phi(\mathcal{I}_s) \models uRv$ implies $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z}) \models uRv$. Let $\mathcal{N}_0, \mathcal{N}_1, \ldots$ be the sequence of approximations of $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$. Due to Lemma 5.9, $\mathcal{N}_0, \mathcal{N}_1, \ldots$ stabilises after finitely many steps, i.e. there is p' such that $\mathcal{N}_{p'} = \mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$. Thus, there exists T and P that are the upper bounds on the threshold and the period of the regular frames among $\mathcal{N}_0, \mathcal{N}_1, \ldots$. Let $b(p) = pn(T+n)$. We show by induction on p that if $\underline{i}, \underline{i+j+1}$ are $b(p)$-inner in $\mathcal{I}_s$ and $\mathcal{N}_p \models \underline{i}R\underline{i+j+1}$, then $\mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i}R\underline{i+j+1}$. We conclude that for the stabilising index p' (such that $\mathcal{N}_{p'} = \mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$, which exists by Lemma 5.9), $\mathfrak{g}_{in}^\Phi = b(p')$ satisfies the statement of the lemma.
The induction base. The frame $\mathcal{N}_0 = \mathcal{L}_\mathbb{Z}$ contains only short, forward edges that are also in $\mathcal{I}_s$ and therefore in $\mathfrak{C}_\Phi(\mathcal{I}_s)$.
The induction step. If $\mathcal{N}_{p+1} = \mathcal{N}_p$ then we are done. Otherwise, assume that for all $s > 0, i \in \mathbb{Z}$ if $\underline{i}, \underline{i+j+1}$ are $b(p)$-inner in $\mathfrak{C}_\Phi(\mathcal{I}_s)$ and $\mathcal{N}_p \models \underline{i}R\underline{i+j+1}$, then $\mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i}R\underline{i+j+1}$.

We assume that $\mathcal{N}_p$ contains only forward edges; we argue below that if $\mathcal{N}_p$ contains a backward edge we can stop induction. We proceed in two steps: In Step 1 we show that for all $s > 0, i \in \mathbb{Z}$ if $\underline{i}, \underline{i+R_p^s}$ are $b(p+1)$-inner in $\mathfrak{C}_\Phi(\mathcal{I}_s)$, then $\mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i}R\underline{i+R_p^s}$. Next, in Step 2 we show that Step 1 implies the induction hypothesis.

Step 1: Let $u_1, u_2, \ldots, u_n$, with $u_1 = \underline{i}$ and $u_2 = \underline{i+R_p^s}$, be the worlds that imply the edge $(\underline{i}, \underline{i+R_p^s})$ in $\mathcal{N}_{p+1}$. We consider two cases. If $\mathcal{N}_p = \mathcal{L}_\mathbb{Z}$, then we may assume that $u_1, u_2, \ldots, u_n$ are contained in $\{\underline{i-n}, \ldots, \underline{i+n-1}, \underline{i+n}\}$, so if $\underline{i}$ is n-inner in $\mathcal{I}_s$, then $\mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i}R\underline{i+R_p^s}$.

Otherwise, the frame $\mathcal{N}_p$ contains a long edge and, therefore by Lemma 5.9, it is regular. Clearly, $|R_p^s| < tr(\mathcal{N}_{p+1}) + pi(\mathcal{N}_{p+1}) \leq T+n$. Thus, by Lemma 5.8 applied to $u_1, u_2, \ldots$, with $a = 1$ and $b = 2$, we can assume, without loss of generality, that for every $i \in \{1, \ldots, k\}$, if $x, y \in \mathbb{N}$ are such that $u_i = \underline{x}$, $u_{i+1} = \underline{y}$, then $|x - y| \leq kT + n$. Thus, if u_1, u_2 are $b(p+1)$-inner, then all $u_1, \ldots, u_k$ are $b(p)$-inner. By the inductive hypothesis for all $v, v' \in \{u_1, \ldots, u_n\}$, if $\mathcal{N}_p \models vRv'$, then $\mathfrak{C}_\Phi(\mathcal{I}_s) \models vRv'$. Hence, $\mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i}R\underline{i+R_p^s}$.

Step 2: If $R_p^s \leq 0$, then for any $s > 0$ if some worlds $\underline{i}, \underline{i+R_p^s}$ are $b(p+1)$-inner in $\mathcal{I}_s$, then, by $\mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i}R\underline{i+R_p^s}$, we have $\underline{i}, \underline{i+R_p^s}, \underline{i+R_p^s+1}, \ldots, \underline{i}$ is a cycle, and therefore $\underline{i}$ and $\underline{i+R_p^s}$ are ∞-inner in $\mathfrak{C}_\Phi(\mathcal{I}_s)$. If some world $\underline{i}$ is not $b(p+1)$-preceded in $\mathcal{I}_s$ but it is $b(p+1)$-preceded in $\mathfrak{C}_\Phi(\mathcal{I}_s)$, then by Lemma 5.6 (i) it is ∞-preceded and, since it has a ∞-inner descendant, it is ∞-preceded. Therefore all $b(p+1)$-inner worlds in $\mathfrak{C}_\Phi(\mathcal{I}_s)$ are ∞-inner. For any ∞-inner worlds $\underline{i}, \underline{i+j}$ Lemma 5.6 (ii) shows that if $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z}) \models \underline{i}R\underline{i+j}$ then $\mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i}R\underline{i+j}$. This proves the lemma and therefore we can stop the induction here.

It remains to consider the case when $R_p^s > 1$. Let $r = R_p^s - 1$. We observe that $X_\Phi^{\mathcal{N}_{p+1}} = \{j + \alpha r : \alpha \geq 0, j \in X_\Phi^{\mathcal{N}_p}\}$, i.e., $X_\Phi^{\mathcal{N}_{p+1}}$ is the additive closure of $X_\Phi^{\mathcal{N}_p} \cup \{r\}$. Indeed, the set $X_\Phi^{\mathcal{N}_{p+1}}$ contains $X_\Phi^{\mathcal{N}_p} \cup \{r\}$ and Lemma 5.3 implies that $X_\Phi^{\mathcal{N}_{p+1}}$ is additively closed. Conversely, for every additively closed set Y the frame $\langle \{\underline{i} : i \in Z\}, \{(\underline{i}, \underline{i+j+1}) : i \in \mathbb{Z}, j \in Y\} \rangle$ is closed under composition. Thus, we have to show that for every $s, \alpha, i \in \mathbb{N}$, if $\underline{i}, \underline{i+j+\alpha r+1}$ are $b(p+1)$-inner in $\mathcal{I}_s$ and $j \in X_\Phi^{\mathcal{N}_{p+1}}$, then $\mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i}R\underline{i+j+\alpha r+1}$. We show that by induction on α.

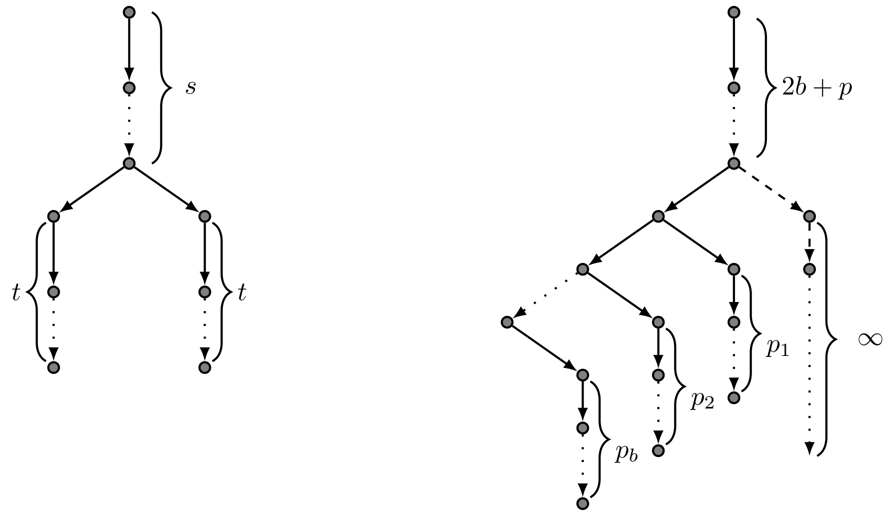

Fig. 5. The frames $\mathcal{C}_t^s$ on the left (see Section 5.4) and $\widehat{\mathcal{V}_l^p}$ on the right (see Section 5.5), where $l = p_1, \ldots, p_b$. The frame $\mathcal{V}_l^p$ results from $\widehat{\mathcal{V}_l^p}$ by removing the infinite path.

The base case, $\alpha = 0$, follows from the inductive hypothesis for p.
The induction step. Let $t = s-(j+r(\alpha-1))$. Assume that $\underline{i}, \underline{i+j+\alpha r+1}$ are $b(p+1)$-inner in $\mathcal{I}_s$.

Consider a function f defined as follows:

$$f(\underline{p}) = \begin{cases} \underline{p} & \text{if } p \le i+r \\ \underline{p+j+(\alpha-1)\cdot r} & \text{if } p > i+r \end{cases}$$

When, $p \neq i+r$, $f(\underline{p}), f(\underline{p+1})$ are successive worlds in $\mathcal{I}_s$, thus $\mathcal{I}_s \models f(\underline{p}) R f(\underline{p+1})$. For $p = i+r$, $f(\underline{i+r}) = \underline{i+r}$ and $f(\underline{i+r+1}) = \underline{i+j+(\alpha-1)r+1}$. By the induction assumption $\mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i+r} R \underline{i+j+(\alpha-1)r+1}$, thus $\mathcal{I}_s \models f(\underline{i+r}) R f(\underline{i+r+1})$. Hence, f is a morphism from $\mathcal{I}_t$ into $\mathfrak{C}_\Phi(\mathcal{I}_s)$. By Observation 4.10 f is a morphism from $\mathfrak{C}_\Phi(\mathcal{I}_t)$ into $\mathfrak{C}_\Phi(\mathcal{I}_s)$.

Since $\underline{i}, \underline{i+j+\alpha r+1}$ are $b(p+1)$-inner in $\mathcal{I}_s$, $\underline{i+r}$ is $b(p+1)$-inner in $\mathcal{I}_t$. Hence, $\mathfrak{C}_\Phi(\mathcal{I}_t) \models \underline{i} R \underline{i+r-1}$. Thus, the morphism f applied to $\underline{i}, \underline{i+r-1}$ implies that $\mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i} R \underline{i+j+\alpha r+1}$. $\square$

5.4. Proof of Lemma 4.16

LEMMA 4.16. *Assume that Φ does not have TCMP. Then, there is a constant $\mathfrak{g}_{eq}^\Phi < \infty$ such that $\mathfrak{g}_{eq}^\Phi \geq \mathfrak{g}_{nf}^\Phi$ and for each tree $\mathcal{T}$ with a bounded degree, a world w at the level $\mathfrak{g}_{eq}^\Phi$ in $\mathfrak{C}_\Phi(\mathcal{T})$, and all i, all the $\mathfrak{g}_{eq}^\Phi$-followed descendants of w at the level $\mathfrak{g}_{eq}^\Phi + i$ are equivalent in the frame $\mathfrak{C}_\Phi(\mathcal{T})$.*

For $k_1 \in \mathbb{N}$ and $k_2 \in \mathbb{N}_\infty$, we define the frame $\mathcal{C}_{k_2}^{k_1}$ as $\mathcal{T}_{\infty \restriction W_\lambda}$, where $W_\lambda = \{\underline{s} : s \sqsubseteq 0^{k_1+k_2} \vee s \sqsubseteq 0^{k_1} 1^{k_2}\}$ ($\sqsubseteq$ denotes the prefix relation).

Frames $\mathcal{C}_{k_2}^{k_1}$ play a similar role to $\mathcal{L}_\mathbb{Z}$ – they are canonical frames for analysing Φ. Indeed, those frames have exactly two branches, and we show in Lemma 5.10, that if Φ does not fork, it is reflected in a finite frame $\mathcal{C}_{k_2}^{k_1}$. We enhance this in Lemma 5.11 by showing that if Φ does not fork, there is an edge in $\mathcal{C}_{k_2}^{k_1}$ that connects consecutive levels on different branches. Then, in Lemma 5.12, we show a connection between frames $\mathcal{C}^{k_1}_{k_2}$ and closures of arbitrary trees. Finally, the proof of Lemma 4.16 follows.

LEMMA 5.10. *Assume that* Φ *does not fork. Then, there exist* $s, t, x, y \in \mathbb{N}$ *such that* $x, y > 0$ *and* $\mathfrak{C}_\Phi(\mathcal{C}^s_t) \models \underline{0^{s+x}}R\underline{0^s1^y}$.

PROOF. Since Φ does not fork, there exist $p, p_0, p_1 \in \mathcal{T}_\infty$ such that $\mathfrak{C}_\Phi(\mathcal{T}_\infty) \models \underline{p0p_0}, \underline{p1p_1}$, i.e., the edge $(\underline{p0p_0}, \underline{p1p_1})$ violates tree-compatibility.

Let $\mathcal{T}_\infty = \langle W, R\rangle$ and i be such that the edge $(\underline{p0p_0}, \underline{p1p_1})$ belongs to $\textsc{Cons}^i_{\Phi,W}(R) \setminus \textsc{Cons}^{i-1}_{\Phi,W}(R)$ (cf. Section 4.1). Observe that this edge is a consequence of Φ and at most $|\Phi|$ edges that belong $\textsc{Cons}^{i-1}_{\Phi,W}(R)$. Similarly, every edge in $\textsc{Cons}^{i-1}_{\Phi,W}(R)$ is a consequence of at most $|\Phi|$ edges from $\textsc{Cons}^{i-2}_{\Phi,W}(R)$, and so on. In consequence, the edge $(\underline{p0p_0}, \underline{p1p_1})$ is a consequence of at most $|\Phi|^i$ edges in $\mathcal{T}_\infty$. Thus, there is a number k such that for a full binary tree $\mathcal{T}_k$ of height k, $\mathfrak{C}_\Phi(\mathcal{T}_k) \models \underline{p0p_0}R\underline{p1p_1}$.

Now, consider a function g from $\mathcal{T}_k$ into $\mathcal{C}^{|p|}_k$ defined as follows:

$$g(\underline{s}) = \begin{cases} \underline{0^{|s|}} & \text{if } |s| \le |p| \\ \underline{0^{|p|}z^{|s|-|p_1|}} & \text{if } |s| \in [|p|+1, k] \text{ and } z \text{ is the } (|p|+1)\text{th letter in } s \end{cases}$$

The function g is a morphism from $\mathcal{T}_k$ into $\mathcal{C}^{|p|}_k$ such that $g(\underline{p}) = \underline{0^{|p|}}$, $g(\underline{p0p_0}) = \underline{0^{|p|}0^{|p_0|+1}}$ and $g(\underline{p1p_1}) = \underline{0^{|p|}1^{|p_1|+1}}$. Due to Observation 4.10, g is a morphism from $\mathfrak{C}_\Phi(\mathcal{T}_k)$ into $\mathfrak{C}_\Phi(\mathcal{C}^{|p|}_k)$. Thus, $\mathfrak{C}_\Phi(\mathcal{T}_k) \models \underline{p0p_0}, \underline{p1p_1}$ and g imply that $\mathfrak{C}_\Phi(\mathcal{C}^{|p|}_k) \models \underline{0^{|p|}0^{|p_0|+1}}R\underline{0^{|p|}1^{|p_1|+1}}$. □

LEMMA 5.11. *Assume that* Φ *does not fork. Then, there exist* $s, t, y \in \mathbb{N}$ *such that* $x > 0$ *and* $\mathfrak{C}_\Phi(\mathcal{C}^s_t) \models \underline{0^{s+y}}R\underline{0^{s+1}1^y}$.

PROOF. Let $s', t', x', y' \in \mathbb{N}$ be such that in $\mathfrak{C}_\Phi(\mathcal{C}^{s'}_{t'})$ there is an edge from $\underline{0^{s'+x'}}$ to $\underline{0^{s'}1^{y'}}$. Since rotation of branches, i.e., the function $r : \mathfrak{C}_\Phi(\mathcal{C}^{s'}_{t'}) \to \mathfrak{C}_\Phi(\mathcal{C}^{s'}_{t'})$ defined as $f(\underline{0^z}) = \underline{0^z}$ for $z \le s'$, $f(\underline{0^{s'+z}}) = \underline{0^{s'}1^z}$ for $z \le t'$ and $f(\underline{0^{s'}1^z}) = \underline{0^{s'+z}}$ for $z \le t'$, is an automorphism of $\mathfrak{C}_\Phi(\mathcal{C}^{s'}_{t'})$, we can assume that $x' \le y'$. Also, the morphism $f : \mathcal{C}^{s'}_{t'} \to \mathcal{C}^{s'}_{t'}$ defined as $f(\underline{i}) = \underline{0^{|i|}}$ implies that $\mathfrak{C}_\Phi(\mathcal{C}^{s'}_{t'})$ contains edge $(\underline{0^{s'+t'}}, \underline{0^{s'+y'}})$.

Let $s = s' + 1$, $t = t' + |x' - y'|$ and $y = max(x', y')$. We define a function $g : \mathcal{C}^{s'}_{t'} \to \mathcal{C}^s_t$ as follows.

$$g(\underline{p}) = \begin{cases} \underline{0^{|p|+1}} & \text{if } p \text{ ends with 0 and } |p| < s' + x' \\ \underline{0^{|p|+y'-x'}} & \text{if } p \text{ ends with 0 and } |p| \ge s' + x' \\ \underline{0p} & \text{if } p \text{ ends with 1} \end{cases}$$

Since $\mathfrak{C}_\Phi(\mathcal{C}^s_t) \models \underline{0^{s'+x'}}R\underline{0^{s'+y'}}$, the function g is a morphism from $\mathcal{C}^{s'}_{t'}$ to $\mathfrak{C}_\Phi(\mathcal{C}^s_t)$ and, by employing Observation 4.10, it is a morphism from $\mathfrak{C}_\Phi(\mathcal{C}^{s'}_{t'})$ to $\mathfrak{C}_\Phi(\mathcal{C}^s_t)$. Hence, $\mathfrak{C}_\Phi(\mathcal{C}^s_t) \models g(\underline{0^{s'+x'}})Rg(\underline{0^{s'}1^{y'}})$. Since $g(\underline{0^{s'+x'}}) = \underline{0^{s+y}}$, $g(\underline{0^{s'}1^{y'}}) = \underline{0^s1^x}$, the result follows. □

LEMMA 5.12. *Assume that* Φ *does not forks. There exist* s, t, y *such that for every tree of bounded degree* $\mathcal{T}$ *and every world* w *at a level* s *in* $\mathcal{T}$*, for every* $i \ge 0$*, if* u_1, u_2 *are* t*-followed descendants of* w *in* $\mathfrak{C}_\Phi(\mathcal{T})$ *at levels* $s + y + i$ *and* $s + y + i + 1$*, then* $\mathfrak{C}_\Phi(\mathcal{T}) \models u_1Ru_2$.

Remark that in contrary to Lemmas 5.10 and 5.11 in Lemma 5.12 the worlds need to be t-followed only in $\mathfrak{C}_\Phi(\mathcal{T})$. Some worlds that are not t-followed in $\mathcal{T}$ may become t-followed in $\mathfrak{C}_\Phi(\mathcal{T})$.

PROOF. Let s, t, y be the constants guaranteed by Lemma 5.11. The proof is by induction on i.

The base case, $i = 0$. Let u_1, u_2 be t-followed descendants of w in $\mathfrak{C}_\Phi(\mathcal{T})$ at levels $s+y$ and $s+y+1$. There is a morphism f from $\mathcal{C}_t^s$ into $\mathfrak{C}_\Phi(\mathcal{T})$ such that $f(\underline{0^s}) = w$, $f(\underline{s_1}) = u_1$ and $f(\underline{s_2}) = u_2$ for $s_1 = 0^{s+y}$ and $s_2 = 0^s 1^{y+1}$. Therefore $\mathfrak{C}_\Phi(\mathcal{T}) \models u_1 R u_2$.
The induction step. Let u_1, u_2 be t-followed descendants of w in $\mathfrak{C}_\Phi(\mathcal{T})$ at levels $s+y+(i+1)$ and $s+y+(i+1)+1$ and let v_1, v_2 be predecessors of u_1, u_2 in $\mathfrak{C}_\Phi(\mathcal{T})$ at levels $s+y+i, s+y+(i+1)$. Clearly, v_1, v_2 are t-followed descendants of w in $\mathfrak{C}_\Phi(\mathcal{T})$. By the induction assumption, $\mathfrak{C}_\Phi(\mathcal{T}) \models v_1 R v_2$.

Notice that there is a morphism g from $\mathcal{C}_t^s$ into $\mathcal{C}_{t-y+1}^{s+y-1}$ such that $g(0^{s+y}) = 0^{s+y}$ and $g(0^s 1^{y+1}) = 0^{s+y-1}11$. Since $\mathfrak{C}_\Phi(\mathcal{C}_t^s) \models \underline{s_1} R \underline{s_2}$, the morphism g, extended to a morphism from $\mathfrak{C}_\Phi(\mathcal{C}_t^s)$ to $\mathfrak{C}_\Phi(\mathcal{C}_{t-y+1}^{s+y-1})$, implies that $\mathfrak{C}_\Phi(\mathcal{C}_{t-y+1}^{s+y-1}) \models \underline{0^{s+y-1}0} R \underline{0^{s+y-1}11}$.

There is a morphism h from $\mathcal{C}_{t-y+1}^{s+y-1}$ into $\mathfrak{C}_\Phi(\mathcal{T})$ such that $h(\underline{0^{s+y-1}}) = v_2$, $h(\underline{0^{s+y-1}0}) = u_1$ and $h(\underline{0^{s+y-1}11}) = u_2$. By Observation 4.10, h is a morphism from $\mathfrak{C}_\Phi(\mathcal{C}_{t-y+1}^{s+y-1})$ to $\mathfrak{C}_\Phi(\mathcal{T})$. Since $\mathfrak{C}_\Phi(\mathcal{C}_{t-y+1}^{s+y-1}) \models \underline{0^{s+y-1}0} R \underline{0^{s+y-1}11}$, we have $\mathfrak{C}_\Phi(\mathcal{T}) \models u_1 R u_2$. □

PROOF OF LEMMA 4.16. Let s, t, y be the constants guaranteed by Lemma 5.11, $\mathcal{T}$ be a tree of bounded degree, w be a world at a level s and $i > 0$. We prove that all $(t+1)$-followed descendants of w in $\mathfrak{C}_\Phi(\mathcal{T}^+)$ at the level $s+y+i$ are equivalent.

Let $i > 0$ and a, b be $(t+1)$-followed (in $\mathfrak{C}_\Phi(\mathcal{T}^+)$) descendants of w at a level $L = s+y+i$. It is sufficient to show that all predecessors of a are predecessors of b. Let c be a predecessor of a in $\mathfrak{C}_\Phi(\mathcal{T})$. Clearly, c is $(t+2)$-followed in $\mathfrak{C}_\Phi(\mathcal{T})$.

Let p_a (p_b) be the predecessor of a (b resp.) in $\mathcal{T}$. Let $S_{a,0}, S_{a,1}$ (resp., $S_{b,0}, S_{b,1}$) be the partition of the successors of a (resp., b) such that worlds from $S_{a,1}$ (resp., $S_{b,1}$) are (t)-followed in $\mathfrak{C}_\Phi(\mathcal{T})$ and worlds from $S_{a,0}$ (resp., $S_{b,0}$) are not. The sets $S_{a,1}, S_{b,1}$ are nonempty, since a, b are $(t+1)$-followed. We define $\mathcal{T}^+$ as the frame, based on $\mathcal{T}$, resulting from removing from $\mathcal{T}$ worlds $S_{a,0}, S_{b,0}$ and their descendants in $\mathcal{T}$, i.e., the successors of a or b which are not (t)-followed in $\mathfrak{C}_\Phi(\mathcal{T})$ and their descendants (in $\mathcal{T}$). Additionally, $\mathcal{T}^+$ has the edges $(p_a, b), (p_b, a)$ and edges (x, y) for $x \in \{a, b\}$ and $y \in S_{a,1} \cup S_{b,1}$. Notice that a, b have the same predecessors and successors in $\mathcal{T}^+$. Since c is $(t+2)$-followed in $\mathfrak{C}_\Phi(\mathcal{T})$, it belongs to $\mathcal{T}^+$.

In the following we show three facts:
(i) $\mathfrak{C}_\Phi(\mathcal{T}^+) \models cRa$ (ii) $\mathfrak{C}_\Phi(\mathcal{T}^+) \models cRb$ (iii) $\mathfrak{C}_\Phi(\mathcal{T}) \models cRb$

In order to show (i), we consider a morphism f from $\mathcal{T}$ into itself, which is identity on $\mathcal{T}^+$ and maps worlds from $S_{a,0}$ and their descendants (in $\mathcal{T}$) into a path of maximal length beginning in a. Similarly, worlds from $S_{b,0}$ and their descendants are mapped into a path of maximal length beginning in b. We notice that the range of f is contained in $\mathcal{T}^+$, since any maximal path beginning in a has to contain a world from $S_{a,1}$. Otherwise, when a maximal path π beginning in a contains a world s_0 from $S_{a,0}$, there is a morphism which maps all the descendants of a on π, and this implies that s_0 is t-followed in $\mathfrak{C}_\Phi(\mathcal{T})$.

Hence, the morphism f is really a morphism from $\mathcal{T}$ into $\mathcal{T}^+$, $f(a) = a$ and $f(c) = c$. This implies that $\mathfrak{C}_\Phi(\mathcal{T}^+) \models cRa$.

For (ii), notice that since a and b have the same predecessors and the same successors in $\mathcal{T}^+$, the function g which is identity on $\mathcal{T}^+$ except that it swaps a with b is an automorphism of $\mathcal{T}^+$. Hence, by Observation 4.10 g is an automorphism of $\mathfrak{C}_\Phi(\mathcal{T}^+)$. This implies that if $a \neq c$, then $\mathfrak{C}_\Phi(\mathcal{T}^+) \models cRb$.

If $a = c$, then a is reflexive in $\mathfrak{C}_\Phi(\mathcal{T}^+)$ and g implies that b is also reflexive in $\mathfrak{C}_\Phi(\mathcal{T}^+)$. Then, the function h from $\mathcal{T}^+$ into $\mathcal{T}^+$ which maps all the descendants of b on b is a morphism from $\mathcal{T}^+$ into $\mathfrak{C}_\Phi(\mathcal{T}^+)$. The composition of h with f is a morphism from $\mathcal{T}$ into $\mathfrak{C}_\Phi(\mathcal{T}^+)$. The function $(h \circ f)$ is also a morphism from $\mathfrak{C}_\Phi(\mathcal{T})$ into $\mathfrak{C}_\Phi(\mathcal{T}^+)$ which maps a to a and the descendants of b into b. Let s be t-followed descendant of b. Due to Lemma 5.12, $\mathfrak{C}_\Phi(\mathcal{T}) \models aRs$. Hence, $\mathfrak{C}_\Phi(\mathcal{T}^+) \models (h \circ f)(a) R (h \circ f)(s)$ where $(h \circ f)(a) = a$ and $(h \circ f)(s) = b$.

For (iii), Lemma 5.12 implies that the identity is a morphism from $\mathcal{T}^+$ into $\mathfrak{C}_\Phi(\mathcal{T})$. Hence, the identity on $\mathcal{T}^+$ is a morphism from $\mathfrak{C}_\Phi(\mathcal{T}^+)$ into $\mathfrak{C}_\Phi(\mathcal{T})$ and $\mathfrak{C}_\Phi(\mathcal{T}^+) \models cRb$ implies $\mathfrak{C}_\Phi(\mathcal{T}) \models cRb$.

We showed that all predecessors of a are predecessors of b and, by symmetry, all predecessors of b are predecessor of a, and therefore a and b are equivalent.

To conclude the proof, we put $\mathfrak{g}^\Phi_{eq} = \max(t+1, s+y)$. Since $s > \mathfrak{g}^\Phi_{nf}$, $\mathfrak{g}^\Phi_{eq} > \mathfrak{g}^\Phi_{nf}$. □

5.5. Proof of Lemma 4.18

LEMMA 4.18. *Let Φ be a formula satisfying S2 with χ. There is a constant $\mathfrak{g}^\Phi_{bs} < \infty$ such that for each possible $(\mathfrak{g}^\Phi_{bs}, gcd(\chi))$-profile l there is a predecessors generator $\mathbb{PG}^\Phi_l$ such that for each world w of a model $\mathfrak{C}_\Phi(\mathcal{T})$ s.t. $|\overline{w}| > \mathfrak{g}^\Phi_{bs}$ and $l = prof(w, \mathfrak{g}^\Phi_{bs}, gcd(\chi))$, $\mathbb{PG}^\Phi_l$ is a predecessors generator of w.*

The above lemma says (roughly) that the set of predecessors of a world w of a frame $\mathfrak{C}_\Phi(\mathcal{T})$ for Φ satisfying S2 can be computed based on the profile of w and the following information about the path p from the root of $\mathcal{T}$ to w: types of worlds in a prefix (of bounded size) of p, types of worlds in a suffix of p, and types of ($\mathfrak{g}^\Phi_{in}$-inner) predecessors of worlds in a suffix of p.

We prove Lemma 4.18 in three steps. (1) In Lemma 5.13 we characterise the set of predecessors of w that are far from the beginning and the end of p. (2) In Lemma 5.14 we study the remaining predecessors of w, i.e., predecessors that are in a prefix (of bounded size) of p or in a suffix (of bounded size) of p. (3) We combine (1) and (2) to show Lemma 4.18.

Let $\underline{i}, \underline{j} \in \mathcal{I}_s$ and let $\mathcal{N}$ be a frame containing $\mathcal{I}_s$. We say that $\underline{i}$ is a b-descendant (resp., b-ancestor) of $\underline{j}$ iff $i - j > b$ (resp., $j - i > b$). We say that a set of worlds $W \subseteq \mathcal{I}_s$ is g-periodic on A iff for every $\underline{i} \in W \cap A$, $\underline{i+g} \in A$ implies $\underline{i+g} \in W$, and $\underline{i-g} \in A$ implies $\underline{i-g} \in W$. Observe that the empty set of worlds is g-periodic on any A.

LEMMA 5.13. *Assume that Φ satisfies S2 with χ and let $g = gcd(\chi)$. Then, there exists a bound $\mathfrak{g}^\Phi_{5.13} \in \mathbb{N}$ such that for every $s > 0$ and every w, the set of successors of w in $\mathfrak{C}_\Phi(\mathcal{I}_s)$ is g-periodic on the set of $\mathfrak{g}^\Phi_{5.13}$-inner $\mathfrak{g}^\Phi_{5.13}$-descendants of w (in $\mathfrak{C}_\Phi(\mathcal{I}_s)$), and the set of predecessors of w in $\mathfrak{C}_\Phi(\mathcal{I}_s)$ is g-periodic on the set of $\mathfrak{g}^\Phi_{5.13}$-inner $\mathfrak{g}^\Phi_{5.13}$-ancestors of w (in $\mathfrak{C}_\Phi(\mathcal{I}_s)$).*

PROOF. We show that for every $i \in [1, \mathfrak{g}^\Phi_{in}]$, then there are $b_i, c_i \geq lcm(\chi)$ such that: for every $s > b_i$, if $\underline{i}$ has long outgoing edges in $\mathfrak{C}_\Phi(\mathcal{I}_s)$ then

$$\text{for every } k \in [b_i - i, s - i - b_i] \text{ we have } \mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i}R\underline{i+k} \text{ iff } g \mid k-1 \tag{1}$$

and for every $s > c_i$ and if $\underline{s-i}$ that has long incoming edges in $\mathfrak{C}_\Phi(\mathcal{I}_s)$ then

$$\text{for every } k \in [c_i, s - i - c_i] \text{ we have } \mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{s-i-k}R\underline{s-i} \text{ iff } g \mid k-1 \tag{2}$$

Then, for $b = max(\mathfrak{g}^\Phi_{in}, b_1, \ldots, b_{\mathfrak{g}^\Phi_{in}}, c_1, \ldots, c_{\mathfrak{g}^\Phi_{in}})$, the following holds: for every $\underline{i} \in \mathfrak{C}_\Phi(\mathcal{I}_s)$:

(i) the set of b-inner b-descendants of $\underline{i}$ is $\{\underline{i+k} : g \mid (k-1), k \in [b, s-i-b]\}$,
(ii) the set of b-inner b-ancestors of $\underline{i}$ is $\{\underline{i-k} : g \mid (k-1), k \in [b, i-b]\}$.

It follows that $\mathfrak{g}^\Phi_{5.13} = b$ satisfies the statement of the lemma.

(i) and (ii), $\subseteq$: The set of b-inner b-descendants of $\underline{i}$ is contained in $\{\underline{i+k} : k \in [b, s-(i+b)]\}$, and the set of b-inner b-ancestors of $\underline{i}$ is contained in $\{\underline{i-k} : k \in [b, i-b]\}$. Since there is a morphism $\mathfrak{C}_\Phi(\mathcal{I}_s)$ into $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$, the property S2 implies that if in $\mathfrak{C}_\Phi(\mathcal{I}_s)$ we have $\underline{i}R\underline{i+k}$, then $k-1$ is in the additive closure of χ. In particular, $g \mid k-1$. Thus, we have $\subseteq$ inclusion for (i) and (ii).

(i), $\supseteq$: First, assume that $\underline{i}$ is $\mathfrak{g}^\Phi_{in}$-inner in $\mathfrak{C}_\Phi(\mathcal{I}_s)$. Then, the property S2 implies that $\mathfrak{C}_\Phi(\mathcal{I}_s) \models \underline{i}R\underline{i+k}$ iff $k-1$ is in the additive closure of χ. But, for $k \geq lcm(\chi)$, $k-1$ belongs to the additive closure of χ iff $g \mid k-1$. Thus, for $b > max(\mathfrak{g}^{\Phi}_{in}, lcm(\chi))$, the set of successors of $\underline{i}$ is g-periodic on the set of b-inner b-descendants of $\underline{i}$.

If $\underline{i}$ is not $\mathfrak{g}^{\Phi}_{in}$-followed, then it has no $\mathfrak{g}^{\Phi}_{in}$-followed successors, and since $b \geq \mathfrak{g}^{\Phi}_{in}$, the statement is satisfied. If $\underline{i}$ is not $\mathfrak{g}^{\Phi}_{in}$-preceded in $\mathfrak{C}_{\Phi}(\mathcal{I}_s)$, then $\underline{i}$ is one of the worlds among $\underline{1}, \dots, \underline{\mathfrak{g}^{\Phi}_{in}}$ and the property (1) implies the statement.

It remains to prove (1): Consider $i \in [1, \mathfrak{g}^{\Phi}_{in}]$. If for every $d > 0$, $\underline{i}$ does not have long outgoing edges, then its set of successors that are b-descendants is empty, thus it is g-periodic over any set. Otherwise, let $d > 0$ be such that $\underline{i}$ has a long edge in $\mathfrak{C}_{\Phi}(\mathcal{I}_d)$, i.e., there is $p > 1$ such that $\mathfrak{C}_{\Phi}(\mathcal{I}_d) \models \underline{i}R\underline{i+p}$. A simple argument shows that this edge can be iterated, so that for some b_i and $r > \mathfrak{g}^{\Phi}_{in} + 2$ we have $\mathfrak{C}_{\Phi}(\mathcal{I}_{b_i}) \models \underline{i}R\underline{r}$. Since there is a morphism $\mathfrak{C}_{\Phi}(\mathcal{I}_{b_i})$ into $\mathfrak{C}_{\Phi}(\mathcal{L}_{\mathbb{Z}})$, the property S2 implies that $g \mid (r-i)-1$.

Let $s > 0$. Consider $k > lcm(\chi)$ that satisfies $i+k \in [b_i, s-b_i]$ (assuming that $b_i + g < s - b_i$ for a given s) and $g \mid k-1$. Then, $g \mid (r-i)-1$ implies $g \mid i+k-r$. Moreover, $i+k-1 = (i+k-r)+(r-1)$ and $i+k-1 \geq b_i > r+c$ implies $i+k-r-1 > lcm(\chi)$. This and the fact that $\underline{r-1}$ is $\mathfrak{g}^{\Phi}_{in}$-inner in $\mathfrak{C}_{\Phi}(\mathcal{I}_s)$ imply $\mathfrak{C}_{\Phi}(\mathcal{I}_s) \models \underline{r-1}R\underline{i+k}$.

Now, $\mathfrak{C}_{\Phi}(\mathcal{I}_s) \models \underline{r-1}R\underline{i+k}$ and $i+k < s-b_i$ imply that the function f_k from $\mathcal{I}_{b_i}$ to $\mathfrak{C}_{\Phi}(\mathcal{I}_s)$ defined as follows:

$$f_k(\underline{j}) = \begin{cases} \underline{j} & \text{if } j \leq r-1 \\ \underline{j+(i+k)-r} & \text{otherwise} \end{cases}$$

is a morphism. Due to Observation 4.10, f_k is a morphism from $\mathfrak{C}_{\Phi}(\mathcal{I}_{b_i})$ to $\mathfrak{C}_{\Phi}(\mathcal{I}_s)$. Observe that $f_k(\underline{i}) = \underline{i}$ and $f_k(\underline{r}) = \underline{i+k}$, therefore $\mathfrak{C}_{\Phi}(\mathcal{I}_s) \models \underline{i}R\underline{i+k}$.

(ii) $\supseteq$: The proof that the set of predecessors of $\underline{i}$ is g-periodic on the set of b-inner b-ancestors of $\underline{i}$ follows the same idea. The proof of the case where $\underline{i}$ is $\mathfrak{g}^{\Phi}_{in}$-followed is straightforward, and for the proof of the case where $\underline{i}$ is not $\mathfrak{g}^{\Phi}_{in}$-followed we utilise (and prove) (2) instead of (1). $\square$

We say that (D, E) is a *boundary predecessors generator* of a world w if $D, E \subseteq \{1, \dots, |\overline{w}|\}$ and for some $c > 0$ we have

$$P_{\mathfrak{M}}(w) \cap (\{\overline{w}_d : d < 2c\} \cup \{\overline{w}_{-d} : d < 3c\}) = \{\overline{w}_d : d \in D\} \cup \{\overline{w}_{-d} : d \in E\}$$

LEMMA 5.14. *Let Φ be a formula satisfying S2 with χ. There are constants $\mathfrak{g}^{\Phi}_{5.14}$, $\mathfrak{g}^{\Phi}_{prof} > \mathfrak{g}^{\Phi}_{5.13}$ such that for each possible $(\mathfrak{g}^{\Phi}_{prof}, gcd(\chi))$-profile l there is a finite boundary predecessors generator $\mathbb{BPG}^{\Phi}_l$ such that for each world w of $\mathfrak{C}_{\Phi}(\mathcal{T})$ such that $|\overline{w}| > \mathfrak{g}^{\Phi}_{5.14}$ and $l = prof(w, \mathfrak{g}^{\Phi}_{prof}, gcd(\chi))$, $\mathbb{BPG}^{\Phi}_l$ is a boundary predecessors generator of w.*

PROOF. Let $g = gcd(\chi)$ and $b = max(2 \cdot \mathfrak{g}^{\Phi}_{5.13}, 4 \cdot |\Phi| \cdot lcm(\chi) + \mathfrak{g}^{\Phi}_{in})$, where $\mathfrak{g}^{\Phi}_{5.13}$ is a constant from Lemma 5.13. Let $p > 0$ and $l = p_1, \dots, p_b$ be a b-bounded sequence such that $p_1 > p_2 > \dots > p_b$.

To simplify the notion, for any $s \in \{0,1\}^*$ we identify $\underline{\overline{s}}$ and s, e.g., we treat 0^3 as a path containing worlds $\underline{\epsilon}, \underline{0}, \underline{00}, \underline{000}$.

We define frames $\mathcal{V}^p_l, \widehat{\mathcal{V}^p_l}$ as follows. The frame $\mathcal{V}^p_l$ is a tree consisting of the following paths: $0^{3b+p}, 0^{2b+p+1}1^{p_1}, 0^{2b+p+2}1^{p_2}, \dots, 0^{2b+p+b}1^{p_b}$, and $\widehat{\mathcal{V}^p_l}$ is an extension of $\mathcal{V}^p_l$ by the infinite path: $0^{2b+p}1^{\omega}$. Let $x = 3b+p$. Observe that the world $\underline{0^x}$ has the profile $(l, x \mod g)$ in $\mathcal{V}^p_l$ and $\widehat{\mathcal{V}^p_l}$.

We show that for every tree $\mathcal{T}$ and every $w \in \mathcal{T}$ with $|\overline{w}| = x+1$ (recall that $\overline{w}$ denotes the path in a tree ($\mathcal{T}$ in this case) from the root to w) and $prof(w, b, g) = (l, x \mod i)$, the paths 0^x in $\mathfrak{C}_{\Phi}(\mathcal{V}^p_l)$ and $\overline{w}$ in $\mathfrak{C}_{\Phi}(\mathcal{T})$ are isomorphic, i.e., for all $\overline{w}_{m_1}, \overline{w}_{m_2} \in \overline{w}$:

$$\mathfrak{C}_{\Phi}(\mathcal{V}^p_l) \models \underline{0^{m_1}}R\underline{0^{m_2}} \iff \mathfrak{C}_{\Phi}(\mathcal{T}) \models \overline{w}_{m_1}R\overline{w}_{m_2} \quad (3)$$

Assuming the equivalence (3), let associate with each p sets D_p, E_p such that

$$d \in D_p \text{ iff } \underline{d} \in P_{\mathfrak{C}_\Phi(\mathcal{V}_l^p)}(\underline{0^x}) \cap (\{\underline{0^d} : 0 \leq d < 2b\}$$
$$d \in E_p \text{ iff } \underline{d} \in P_{\mathfrak{C}_\Phi(\mathcal{V}_l^p)}(\underline{0^x}) \cap \{\underline{0^{x-d}} : 0 \leq d < 3b\}$$

Observe that since Φ satisfies S2, the sets D_p, E_p monotonically increase with p, i.e., for every $y \in \chi$ (from S2), $D_p \subseteq D_{p+y}$ and $E_p \subseteq E_{p+y}$. Indeed, there is a morphism of $\mathcal{V}_l^p$ into $\mathfrak{C}_\Phi(\mathcal{V}_l^{p+y})$ such that $\{\underline{0^d} : d < 2b\}$ is mapped onto itself and $\{\underline{0^{x-d}} : d < 3b\}$ is mapped onto $\{\underline{0^{x+y-d}} : d < 3b\}$. Therefore, there is p_L, such that D_{p_L}, E_{p_L} are saturated, i.e., for every $p' > p_L$, if g divides $p' - p_L$, then $D_{p'} = D_{p_L}$ and $E_{p'} = E_{p_L}$. The equivalence (3) implies that for all $\mathcal{T}$, $w \in \mathcal{T}$, if $|\overline{w}| > 3b + p_L$, g divides $|\overline{w}| - 3b + p_L$ and $prof(w, b, g) = (l, |\overline{w}| \mod i)$, then

$$P_{\mathfrak{C}_\Phi(\mathcal{T})}(w) \cap (\{\overline{w}_d : d < 2b\} \cup \{\overline{w}_{-d} : d < 3b\}) = \{\overline{w}_d : d \in D_{p_L}\} \cup \{\overline{w}_{-d} : d \in E_{p_L}\}$$

Therefore, $\mathfrak{g}^\Phi_{prof} = b$ and $\mathfrak{g}^\Phi_{5.14} = 3b + p_L + lcm(\chi)$ satisfy the statement of the lemma.

Now we shall show the equivalence (3). Consider a tree $\mathcal{T}$. Observe that for every $w \in T$, if $|\overline{w}| = x + 1$ and $prof(w, b, g) = (l, |\overline{w}| \mod i)$, then there is a morphism from $\mathcal{V}_l^p$ into $\mathcal{T}$ such that the path 0^x is mapped on the path $\overline{w}$. Conversely, there is a morphism from $\mathcal{T}$ into $\widehat{\mathcal{V}_l^p}$ such that the path $\overline{w}$ is mapped on the path 0^x. This implies that for all $a, b \leq x$,

C1) $\mathfrak{C}_\Phi(\mathcal{V}_l^p) \models \underline{0^a} R \underline{0^b}$ implies $\mathfrak{C}_\Phi(\mathcal{T}) \models \overline{w}_a R \overline{w}_b$, and
C2) $\mathfrak{C}_\Phi(\mathcal{T}) \models \overline{w}_a R \overline{w}_b$ implies $\mathfrak{C}_\Phi(\widehat{\mathcal{V}_l^p}) \models \underline{0^a} R \underline{0^b}$.

Now, we assume the following: (*) $\mathfrak{C}_\Phi(\widehat{\mathcal{V}_l^p})_{\restriction V_l^p} = \mathfrak{C}_\Phi(\mathcal{V}_l^p)$.

Property (*) together with C1 and C2 imply the equivalence (3). It remains to prove (*). Intuitively, we show that the path $0^{2b+p}1^\omega$ does not help in deriving more consequences.

The inclusion $\mathfrak{C}_\Phi(\widehat{\mathcal{V}_l^p}_{\restriction V_l^p}) \supseteq \mathfrak{C}_\Phi(\mathcal{V}_l^p)$ is immediate. For the converse inclusion consider a sequence of frames $\mathcal{L}_0, \mathcal{L}_1, \ldots$, where $\widehat{\mathcal{V}_l^p} = \mathcal{L}_0$ and for each i $\mathcal{L}_{i+1}$ results from applying a single step closure to $\mathcal{L}_i$, i.e., for $\mathcal{L}_i = \langle \widehat{V_l^p}, R_i \rangle$, $\mathcal{L}_{i+1} = \langle \widehat{V_l^p}, \textsc{Cons}_{\Phi, \widehat{V_l^p}}(R_i) \rangle$. Clearly, the limit of this sequence $\bigcup_{i>0} \mathcal{L}_i$ is $\mathfrak{C}_\Phi(\widehat{\mathcal{V}_l^p})$. We show by induction on i that $(\mathcal{L}_i)_{\restriction V_l^p}$ is a subframe of $\mathfrak{C}_\Phi(\mathcal{V}_l^p)$. The basis of the induction, $i = 0$, follows from the definition of $\widehat{\mathcal{V}_l^p}$.

Assume that $(\mathcal{L}_i)_{\restriction V_l^p}$ is a subframe of $\mathfrak{C}_\Phi(\mathcal{V}_l^p)$ and $u_1, \ldots, u_k$ are worlds of $\mathcal{L}_i$, where $u_1, u_2 \in \mathcal{V}_l^p$. We claim that there are $u'_1, \ldots, u'_n \in \mathcal{V}_l^p$ with $u'_1 = u_1, u'_2 = u_2$ such that for all $m_1, m_2 \in [1, n]$, $\mathcal{L}_i \models u_{m_1} R u_{m_2}$ implies $\mathfrak{C}_\Phi(\mathcal{V}_l^p) \models u'_{m_1} R u'_{m_2}$. Thus, if $\mathcal{L}_{i+1} \models u_1 R u_2$, then $\mathfrak{C}_\Phi(\mathcal{V}_l^p) \models u_1 R u_2$. It follows that $(\mathcal{L}_{i+1})_{\restriction V_l^p}$ is a subframe of $\mathfrak{C}_\Phi(\mathcal{V}_l^p)$.

For every world u_j, if $u_j \in \mathcal{V}_l^p$, we put $u'_j = u_j$. Let $d > 2$ be such that $u_d, \ldots, u_m$ do not belong to $\mathcal{V}_l^p$, i.e., there are $z_d, \ldots, z_m$ such that $u_d, \ldots, u_m$ are of the forms $\underline{0^{2b+p}1^{z_d}}, \ldots, \underline{0^{2b+p}1^{z_m}}$. Due to TCMP $u_d, \ldots, u_m$ are connected only with their predecessors and successors on $0^{2b+p}1^\omega$.

Let $z'_d, \ldots, z'_m$ be a sequence of natural numbers such that

— for every $y \in \{d, \ldots, m\}$, g divides $z_y - z'_y$,
— for all $y_1, y_2 \in \{d, \ldots, m\}$, $z_{y_1} < z_{y_2}$ implies $z'_{y_1} < z'_{y_2}$,
— for all $y_1, y_2 \in \{d, \ldots, m\}$, $|z_{y_1} - z_{y_2}| \in [lcm(\chi), 2 \cdot lcm(\chi)]$.

Observe that $max(z'_d, \ldots, z'_m) < 2(m-d)lcm(\chi) < \frac{1}{2}b - lcm(\chi)$. For every $y \in \{d, \ldots, m\}$, we define u'_y as $\underline{0^{2b+p+lcm(\chi)+z'_y}}$ (recall that $u_y = \underline{0^{2b+p}1^{z_y}}$). We claim that for all $y_1, y_2 \in \{1, \ldots, m\}$, if $\mathcal{L}_i \models u_{y_1} R u_{y_2}$, then $\mathfrak{C}_\Phi(\mathcal{V}_l^p) \models u'_{y_1} R u'_{y_2}$.

Indeed, if $y_1, y_2 < d$, then $u_{y_1} = u'_{y_1}$ and $u_{y_1} = u'_{y_1}$. Assume that $y_2 \in \{d, \ldots, m\}$ and consider two cases. If u_{y_1} is $\mathfrak{g}^{\Phi}_{in}$-inner and $\mathcal{L}_i \models u_{y_1} R u_{y_2}$, the property S2 and TCMP imply that g divides $|\overline{u_{y_2}}| - |\overline{u_{y_1}}|$. Now, by definition of $|\overline{u'_{y_2}}|$, g divides $|\overline{u'_{y_2}}| - |\overline{u_{y_1}}|$, $|\overline{u'_{y_2}}| - |\overline{u_{y_1}}| > lcm(\chi)$ and u_{y_1}, u_{y_2} are $\mathfrak{g}^{\Phi}_{in}$-inner. Thus, Lemma 5.8 implies that $\mathfrak{C}_{\Phi}(\mathcal{V}^p_l) \models u_{y_1} R u'_{y_2}$.

Otherwise, if u_{y_1} in not $\mathfrak{g}^{\Phi}_{in}$-inner, but $\mathcal{L}_i \models u_{y_1} R u_{y_2}$, then u_{y_1} belongs to the path $0^{\mathfrak{g}^{\Phi}_{in}-1}$. By Lemma 5.13 in $\mathfrak{C}_{\Phi}(\widehat{\mathcal{V}^p_l})$, the set of successors of u_{y_1} is g-periodic on $\{\underline{0^{\mathfrak{g}^{\Phi}_{in}+b}}, \ldots, \underline{0^{2b+p}}, \underline{0^{2b+p}1}, \underline{0^{2b+p}1^2}, \ldots\}$. Thus, $\mathcal{L}_i \models u_{y_1} R u_{y_2}$ implies that there exists $z' \in [0, g]$ such that g divides $|u_{y_2}| - (b + p + z' + 1)$ we have $\mathfrak{C}_{\Phi}(\widehat{\mathcal{V}^p_l}) \models u_{y_1} R 0^{b+p+z'}$. But, the set of $\mathfrak{g}^{\Phi}_{5.13}$-descendants of u_{y_1} that are $\mathfrak{g}^{\Phi}_{5.13}$-inner in both $\mathfrak{C}_{\Phi}(\mathcal{V}^p_l)$ and $\mathfrak{C}_{\Phi}(\widehat{\mathcal{V}^p_l})$ contains $\underline{0^{b+p+z'}}$. Thus, by Lemma 5.13 in $\mathfrak{C}_{\Phi}(\mathcal{V}^p_l)$ the set of successors of u_{y_1} is g-periodic on $\{\underline{0^{\mathfrak{g}^{\Phi}_{in}+b}}, \ldots, \underline{0^{2b+p}}\}$, and nonempty by $(\mathcal{L}_i)_{\restriction V^p_l} \subseteq \mathfrak{C}_{\Phi}(\mathcal{V}^p_l)$. Therefore, $\mathfrak{C}_{\Phi}(\mathcal{V}^p_l) \models u_{y_1} R 0^{b+p+z'}$. Finally, observe that g divides $(b+p+z'+1) - |u'_{y_2}|$ and u'_{y_2} is a $\mathfrak{g}^{\Phi}_{5.13}$-inner $\mathfrak{g}^{\Phi}_{5.13}$-descendant of u_{y_1}, therefore $\mathfrak{C}_{\Phi}(\mathcal{V}^p_l) \models u_{y_1} R y'_{y_2}$. □

Finally, we prove Lemma 4.18.

Proof of Lemma 4.18. Let $\mathfrak{M} = \langle W, R, \pi \rangle$ be a model based on $\mathfrak{C}_{\Phi}(\mathcal{T}) \in \mathcal{K}_{\Phi}$ of a modal formula φ, where Φ is tree-compatible and satisfies S2. Recall that $\mathfrak{g}^{\Phi}_{5.14}, \mathfrak{g}^{\Phi}_{prof}$ are constants from Lemma 5.14. Let b be a minimal number divisible by g greater than $\mathfrak{g}^{\Phi}_{prof}$. Then, every path $\overline{w}$ in $\mathcal{T}$ with $|\overline{w}| > \mathfrak{g}^{\Phi}_{5.14} + \mathfrak{g}^{\Phi}_{prof}$ satisfies one of the following:

$$P_{\mathfrak{M}}(w) \cap \{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\} = \emptyset \text{ or} \tag{4}$$

$$P_{\mathfrak{M}}(w) \cap \{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\} = P_{\mathfrak{M}}(\overline{w}_{-b}) \cap \{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\} \tag{5}$$

Observe that by Lemma 5.14, emptiness of $P_{\mathfrak{M}}(w) \cap \{\overline{w}_{-3b}, \overline{w}_{-3b+1}, \ldots, \overline{w}_{-2b-1}\}$ depends only on $prof(w, b, g)$. By Lemma 5.13, the set of successors of w is g-periodic on $\{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\}$, which by $\mathfrak{g}^{\Phi}_{5.13} < b$ implies that $P_{\mathfrak{M}}(w) \cap \{\overline{w}_{-3b}, \overline{w}_{-3b+1}, \ldots, \overline{w}_{-2b-1}\}$ is empty iff $P_{\mathfrak{M}}(w) \cap \{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\}$ is empty. Thus, the above dichotomy and Lemma 5.13 imply the statement of the lemma with $\mathfrak{g}^{\Phi}_{bs} = max(\mathfrak{g}^{\Phi}_{5.14}, b)$. Indeed, let (D, E) is the boundary predecessors generator $\mathbb{BPG}^{\Phi}_{prof(w, \mathfrak{g}^{\Phi}_{bs}, gcd(\chi))}$ which exists by Lemma 5.13. If $P_{\mathfrak{M}}(w) \cap \{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\} = \emptyset$, then the predecessors generator $\mathbb{PG}^{\Phi}_{prof(w, \mathfrak{g}^{\Phi}_{bs}, gcd(\chi))}$ is defined as $(\emptyset, D, E)$. Otherwise, when $P_{\mathfrak{M}}(w) \cap \{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\} = P_{\mathfrak{M}}(\overline{w}_{-b}) \cap \{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\}$, $\mathbb{PG}^{\Phi}_{prof(w, \mathfrak{g}^{\Phi}_{bs}, gcd(\chi))}$ is defined as $(\{b\}, D, E)$. Such $\mathbb{PG}^{\Phi}_{prof(w, \mathfrak{g}^{\Phi}_{bs}, gcd(\chi))}$ clearly satisfies the statement.

We show the above dichotomy. Since $b > \mathfrak{g}^{\Phi}_{5.13}$, all worlds in $\{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\}$ are b-inner b-ancestors of both, w and $\overline{w}_{-b}$. Thus, Lemma 5.13 implies that $P_{\mathfrak{M}}(w)$ and $P_{\mathfrak{M}}(\overline{w}_{-b})$ are g-periodic on $\{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\}$. Since g divides b and there is a morphism from $\mathcal{M}$ into $\mathfrak{C}_{\Phi}(\mathcal{L}_{\mathbb{Z}})$, it follows that if they are both non-empty, then $P_{\mathfrak{M}}(w) \cap \{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\} = P_{\mathfrak{M}}(\overline{w}_{-b}) \cap \{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\}$.

Observe that if $P_{\mathfrak{M}}(w) \cap \{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\}$ is non-empty then $P_{\mathfrak{M}}(\overline{w}_{-b}) \cap \{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\}$ is non-empty as well. Indeed, by $|\overline{w}| > \mathfrak{g}^{\Phi}_{5.14} + b$, $\overline{w}_{-b}$ is $\mathfrak{g}^{\Phi}_{in}$-inner, thus the property S2 implies that $\overline{w}_{-b}$ has predecessors among $\{\overline{w}_b, \overline{w}_{b+1}, \ldots, \overline{w}_{-2b-1}\}$. □

6. THE COMPLEXITY

In this section we study the satisfiability problems more carefully to obtain the complexity bounds summarized in Table I.

Table I. A summary of a complexity of a satisfiability problem for modal logic defined by consistent Horn formulae.

Properties of Φ	Φ-SAT$_G$	Φ-SAT$_L$
Φ is bounded	NP-c (6.2)	
Φ is unbounded, …		
…has the TCMP and satisfies S1	ExpTime-c (6.3)	PSpace-c (6.3)
…has the TCMP and satisfies S2	NP-c (6.5)	PSpace-c (6.4, 6.6)
…has the TCMP and satisfies S3	impossible (Lemma 4.12)	
…lacks the TCMP and satisfies S1	PSpace-c (6.8)	NP-c (6.7)
…lacks the TCMP and satisfies S2	NP-c (6.9)	NP-c (6.9)
…lacks the TCMP and satisfies S3	NP-c (6.12)	NP-c (6.12)

6.1. Boundedness

We say that a formula Φ is *bounded* if $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ is not a model of Φ, and *unbounded* otherwise. If the formula is bounded, then there is a number k such that the length of each path in each model of Φ is bounded by k, and the value of k depends only on Φ. As the formula Φ is a parameter of the problem, not a part of an input, the exact value of k is irrelevant from the complexity point of view as it is regarded as a constant.

Example 6.1. Consider any $n \in \mathbb{N}$ and the formula $\Phi_n = x_1Rx_2 \wedge x_2Rx_3 \wedge \cdots \wedge x_{n-1}Rx_n \Rightarrow \bot$. It is not hard to see that no model of Φ_n contains a path of length n.

Inconsistent formulae are a special case of bounded formulae. For an inconsistent formula Φ, problems Φ-SAT$_L$ and Φ-SAT$_G$ can be solved in constant time by an algorithm that always returns "No". For consistent formulae, the satisfiability problems are NP-complete.

Proposition 6.2. *If Φ is a consistent, bounded* UHF *formula, then Φ-SAT$_L$ and Φ-SAT$_G$ are* NP*-complete.*

Proof. The lower bound comes from the trivial reduction from the SAT problem. Below we show the polynomial model property that leads to a straightforward nondeterministic algorithm that guesses a model and verifies it.

Let Φ be a bounded formula and φ be a modal formula. Then for any model $\mathfrak{M} = \langle W, R, \pi\rangle$ of φ and Φ, we can find a $W' \subseteq W$ such that $\mathfrak{M}_{\restriction W'}$ is a model of φ and $|W'|$ is polynomial in $|\varphi|$. First, we add an arbitrary world from $\mathfrak{M}$ that satisfies φ to W'. Then, recursively, for each world w in W' and each subformula $\Diamond\psi$ of φ, if w has a witness for $\Diamond\psi$ in W but not in W', we add one such witness to W'. We proceed until a fixed-point is reached. Observe that since the length of each path is bounded by k, then this procedure takes at most k recursive steps, and in each of them, it adds at most $|\varphi|$ worlds for each element of W'. Therefore, we have $|W'| \leq \sum_{i=0}^{k} |\varphi|^i < |\varphi|^{k+1}$ and $\mathfrak{M}_{\restriction W'}$ is a model of φ, so indeed we find a polynomial model of φ. Of course, since Φ is universal, $\langle W', R_{\restriction W'}\rangle$ satisfies Φ. □

6.2. Formulae with the tree compatible model property

Proposition 6.3. *For a given unbounded* UHF *formula Φ, if Φ has the tree-compatible model property and satisfies S1, then Φ-SAT$_L$ is* PSpace*–complete and Φ-SAT$_G$ is* ExpTime*–complete.*

Proof. Assume that all edges in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ are short. We use standard approaches to satisfiability of modal logic over the class of all models. For local satisfiability we can bound the depth of tree-models and the degree of their worlds linearly in φ and then check the existence of such models in a depth-first search manner in PSPACE (see e.g. [Ladner 1977]; please note that while the cited result does not consider reflexivity and symmetry, there are only some minor changes needed to cover these cases).

For the global satisfiability, we can enforce models whose depth is exponential with respect to the length of the modal formula φ. The existence of models can be checked by an alternating procedure, which first guesses the type of the root and then guesses types of its children and universally repeats the procedure for the children. This algorithm works in alternating polynomial space, and thus the problem is in EXPTIME. The corresponding lower bound can be obtained by encoding the halting problem for alternating Turing machine with polynomial space. $\square$

We use the following notion. For a given world w, *universal requirements* of w, denoted by $UR(w)$, is the subset of the type of w that consists of formulae of the form $\Box\varphi$. We extends this notion to sets by putting $UR(X) = \bigcup_{w \in X} UR(w)$. Moreover, we define *predecessors requirements* of w, denoted by $PR(w)$, as the set of the universal requirements of the predecessors of w, i.e., $UR(\{v : v \text{ is a predecessor of } w\})$. As before, we extend this notion to sets defining $PR(X) = \bigcup_{w \in X} PR(w)$. Finally, we define *$d$-inner predecessors requirements* $PR^d(w)$ as $UR(\{v : v \text{ is a } d\text{-inner predecessor of } w\})$.

PROPOSITION 6.4. *For a given unbounded* UHF *formula Φ, if Φ has the tree-compatible model property and satisfies S2, then Φ-SAT_L and Φ-SAT_G are in* PSPACE.

PROOF. Assume that the condition S2 holds for some χ. This case can be treated similarly to the case of satisfiability over the class of transitive models, i.e., the case of logic K4 (see [Ladner 1977] or Section 6.7 in [Blackburn et al. 2001]). Let A be the additive closure of χ and c be the product of all positive elements of χ.

Lemma 4.18 implies that for all $i \geq \mathfrak{g}^\Phi$ there is a predecessors generator (C, D, E) that depends only on $prof(\underline{i}, \mathfrak{g}^\Phi, gcd(\chi))$ such that

$$PR(\underline{i}) = \bigcup \{UR(\overline{w}_d) : d \in D\} \cup \bigcup \{UR(\overline{w}_{-d}) : d \in E\} \cup \bigcup \{PR^{\mathfrak{g}^\Phi}(\overline{w}_{-c}) : c \in C\} \quad (6)$$

Let $k \in \mathbb{N}$ be greater than $\mathfrak{g}^\Phi$ such that k bounds numbers in the sets C, D, E in all predecessors generators.

We are ready to design an alternating algorithm that guesses a tree-based structure in a top-down manner. For an input φ, it first guesses and verifies first $\mathfrak{g}^\Phi$ levels. Then, the algorithm recursively calls procedure $verify(head, PRs^{\mathfrak{g}^\Phi}, URs, CR, \Diamond\psi)$ where

- $head$ contains information about the first k levels of structure;
- $PRs^{\mathfrak{g}^\Phi}$ is a list of $\mathfrak{g}^\Phi$-inner predecessors requirements of previous k k-inner worlds;
- URs is a list of universal requirements of previous k k-inner worlds;
- CR is a set of predecessors requirements for the current world;
- $\Diamond\psi$ is a subformula of φ.

The procedure guesses a type t that satisfies ψ and all requirements. Then it guesses a subset of subformulae of φ in order to provide all witnesses for the current world, and for each of them guesses their predecessor profile. Next, for all witnesses, it universally calls itself for this subformula with $PRs^{\mathfrak{g}^\Phi}$ and CR updated using Equation (6). Concurrently, it checks whether the predecessor profile has been guessed correctly.

The algorithm described above verifies if φ has a model, but it may run forever. Therefore we add one more parameter to procedure $verify$: a list of visited configurations (i.e. quadruple $(PRs^{\mathfrak{g}^\Phi}, URs, CR, \Diamond\psi)$), and additional condition: return "Yes" if the same configuration is visited second time.

It is not hard to see that if this algorithm returns “Yes”, then it is possible to build a model. Also, thanks to the property (i) of Lemma 4.3, if φ has a model, then it has a tree-based model such that all witnesses for the world at the level k are realized at the level $k+1$. In such tree-based model, worlds are connected only if they are on the same path in tree and, moreover, k-inner worlds v, w are connected if and only if $h_{\mathcal{T}}(v)$ and $h_{\mathcal{T}}(w)$ are. Such a canonical model can be guessed and verified by the algorithm. What remain to be explained is that this algorithm works in polynomial time.

The key observation here is that predecessors requirements cannot shrink, i.e., if we have two configurations $(PRs_1^{\mathfrak{g}^\Phi}, URs_1, CR_1, \Diamond\psi_1)$ and $(PRs_2^{\mathfrak{g}^\Phi}, URs_2, CR_2, \Diamond\psi_2)$ such that the algorithm visits the second one after the first one, then for each $r \in PRs_1^{\mathfrak{g}^\Phi} \cup \{CR_1\}$ (we abuse a notation here since no confusion will result) there is $r' \in PRs_2^{\mathfrak{g}^\Phi} \cup \{CR_2\}$ such that $r \subseteq r'$. It means that the number of possible $PRs^{\mathfrak{g}^\Phi}$ lists can be bounded by $|\varphi|^{2c} \cdot (2c)!$. Also, since universal requirements result from $\Diamond\psi$, and $\Diamond\psi_2$ is a subformula of $\Diamond\psi_1$, there are at most $|\varphi|$ different universal requirements along each path. Thus, the number of all configurations can be bounded by $(|\varphi|^b \cdot (2c)!) \cdot |\varphi|^b \cdot |\varphi|$, which is polynomial in $|\varphi|$. Therefore, after a polynomial number of steps some configuration must occur twice. Since APTime = PSpace, it leads to the membership is PSpace in both global and local case. □

For the global satisfiability problem, we obtain a better algorithm.

Proposition 6.5. *For a given unbounded* UHF *formula* Φ*, if* Φ *has the tree-compatible model property and satisfies S2, then* Φ*-SAT*$_G$ *is* NP*–complete.*

Proof. The lower bound comes from the trivial reduction from the SAT problem. Below we discuss only the upper bound.

Let Φ satisfy S2 for some χ, let c be the product of all positive elements of χ, and $\mathfrak{M}$ be a $\mathcal{T}$-based model of φ from $\mathcal{K}_\Phi$. Observe that if $\mathcal{T}$ has a finite path, then the last world on that path (the leaf) is a model of Φ and φ. Otherwise, if each path is infinite, we prove that φ has a $\mathcal{K}_\Phi$-based model with the number of types bounded by $|\varphi| \cdot c$.

We say that a world w at the level i (of $\mathcal{T}$) is *saturated* if for all k and every successors w' of w at the levels $i + kc$, $PR(w) = PR(w')$.

Observe that in $\mathfrak{M}$ there is a world w such that the subtree rooted in w contains only saturated worlds. Indeed, Lemma 4.11 implies that for all k and every successors w' of w at the levels $i+kc$, $PR(w) \subseteq PR(w')$. Since $PR(w)$ is finite, below a certain level all worlds are saturated. Let $\mathfrak{M}'$ be this subtree. Of course, $\mathfrak{M}'$ is a $\mathcal{K}_\Phi$-model of φ. For each subformula $\Diamond\psi$ of φ and each $i < c$, if there is a world in $\mathfrak{M}'$ at a level $jc + i$ for some j that satisfies ψ, then we take a 1-type of one such world and call it $t_{\psi,j}$. It is not hard to see that there exists a model $\mathfrak{M}''$ that contains only w and worlds of these types — we can construct such a model starting from w, and then recursively constructing new levels that contain all needed witnesses for the previous level.

The non-deterministic algorithm proceeds as follows. First, it guesses sets of requirements PR_0, PR_1, …, PR_{c-1}, and a subset of types of the form $t_{\psi,j}$. If this types are consistent with requirements and for each $t_{\psi,i}$ we can find $t_{\psi_1,i+1 \mod c}$, …, $t_{\psi_s,i+1 \mod c}$ such that these types provide all needed witnesses for a world of type $t_{\psi,i}$, then it returns “Yes”, otherwise it returns “No”. The algorithm works in polynomial time and solves Φ-SAT$_G$. □

Proposition 6.6. *For a given unbounded* UHF *formula* Φ*, if* Φ *has the tree-compatible model property, then* Φ*-SAT*$_L$ *is* PSpace*-hard.*

Proof. We encode the QBF problem, adjusting the usual technique (see e.g. [Ladner 1977]). Let $P = \vartheta_1 p_1 \vartheta_2 p_2 \ldots \vartheta_n p_n.\rho$ be an instance of QBF problem, where $\vartheta_i \in \{\forall, \exists\}$ and ρ is quantifier-free. We define a modal formula φ such that P is true if and only if φ has a $\mathcal{K}_\Phi$-based model.

We define an operator $\Box_i \psi = \psi \wedge \Box\psi \wedge \Box\Box\psi \wedge \cdots \wedge \Box^i \psi$. Formula φ contains the variables $l_0, l_1, \ldots, l_n$ and $p_1, \ldots, p_n$ and is a conjunction of the following formulae.

(1) l_0
(2) $\Box_n \bigvee_{0 \le i \le n} l_i \wedge \Box_n \bigwedge_{j \neq i} \neg(l_i \wedge l_j)$
(3) $\Box_n(l_i \to \Diamond l_{i+1})$ for each $i < n$ such that $\vartheta_{i+1} = \exists$
(4) $\Box_n(l_i \to \Diamond(l_{i+1} \wedge p_{i+1}) \wedge \Diamond(l_{i+1} \wedge \neg p_{i+1}))$ for each $i < n$ such that $\vartheta_{i+1} = \forall$
(5) $\Box_n((l_i \wedge p_i \to \Box_{n-i} p_i) \wedge (l_i \wedge \neg p_i \to \Box_{n-i} \neg p_i))$ for each $i < n$
(6) $\Box_n(l_n \to \psi)$

Consider a tree $\mathcal{T}$ that consists of $n+1$ levels, and each world at ith level has one successor if $\vartheta_i = \exists$ and two successors otherwise.

Assume that P is true. We define a labelling π of $\mathcal{T}$ inductively, starting from the root. Let the root satisfy only l_0. Let w be at a level i. Define $t_{i+1} = \pi(w) \setminus \{l_i\} \cup \{l_{i+1}\}$. If $\vartheta_i = \forall$, then w has two successors and we set their labellings to be t_{i+1} and $t_{i+1} \cup \{p_{i+1}\}$. Otherwise, set the labelling of the successor of w to t_{i+1}, if formula $\vartheta_{i+2} p_{i+2} \ldots \vartheta n p_n.\rho$ is satisfied for a valuation that makes p_{i+1} false, and to $t_{i+1} \cup \{p_{i+1}\}$ otherwise. Then, $\langle \mathfrak{C}_\Phi(\mathcal{T}), \pi \rangle$ is a $\mathcal{K}_\Phi$-based model of φ.

On the other hand, if φ has a model, then we can show that $\mathcal{T}$ can be homomorphically embedded in this model and the image of this embedding is a justification that P is true. $\Box$

6.3. Formulae without the tree compatible model property

Proposition 6.7. *For a given unbounded* UHF *formula Φ, if Φ does not have the tree-compatible model property and satisfies S1, then it has a polynomial model property for the local satisfiability problem.*

Proof. First, by Observation 4.20 we assume that the size of each level is bounded by $p(|\varphi|)$, where p is a polynomial. Proposition 6.7 follows from the fact that in the local satisfiability case for any tree-based model $\mathfrak{M}$ based of $\mathfrak{C}_\Phi(\mathcal{T})$ such that of $\mathfrak{M}_0, \underline{0} \models \varphi$, we can simply remove all worlds w that are at the levels greater than the quantifier depth of φ. Indeed, S1 says that there are only short edges in closures and therefore the removed worlds were not reachable by φ. The resulting model contains at most $p(|\varphi|) \cdot |\varphi|$ worlds, which is polynomial in $|\varphi|$. $\Box$

Proposition 6.8. *For a given unbounded* UHF *formula Φ, if Φ does not have the tree-compatible model property and satisfies S1, then Φ-SAT_G is* PSpace*-complete.*

Proof. The upper bound comes from Observation 4.20, as explained in Section 4.5. Below we discuss only the lower bound. To ease reading, we consider only the formula $\Phi = \{sRt \wedge tRy \wedge sRx \Rightarrow xRy\}$. Proofs for other cases are similar.

Let $\langle D, D_H, D_V, n \rangle$ be an instance of the bounded-space domino problem. We define a formula $\varphi = \psi_c \wedge \psi_v \wedge \psi_h \wedge \psi_e$ over variables $\{t_0, \ldots, t_{n-1}\} \cup D$ where:

— $\psi_c = \bigvee_{d \in D} d \wedge \bigwedge_{d,d' \in D, d \neq d'} (\neg d \vee \neg d')$;
— $\psi_e = \bigwedge_{i<n} \Diamond t_i$;
— $\psi_v = \bigwedge_{i<n} \bigwedge_{d \in D} (t_i \wedge d \to (\bigvee_{(d,d') \in D_V} \Box(t_i \to d')))$;
— $\psi_h = \bigwedge_{i<n-1} \bigwedge_{d \in D} (\Box(t_i \wedge d) \to (\bigvee_{(d,d') \in D_H} \Box(t_{i+1} \to d')))$.

The reduction is polynomial. Suppose that $\mathfrak{M}$ is a model of Φ and φ and v_0 is any world of $\mathfrak{M}$. We define the tiling t by repeating the following procedure. For a given i, we define $v_{j,i}$ as a successor of v_i that satisfies t_j and we put $t(j,i) = d$, where d is satisfied in $v_{j,i}$. Note that ψ_e guarantees that such a successor exists, ψ_v guarantees that if there is more than one such successor, then all of them satisfy the same d, and ψ_c guarantees that all worlds satisfy precisely one d. Finally, we set v_{i+1} as any successor of v_i that satisfies t_0.

It is not hard to see that for all $k < n-1$ and $l \in \mathbb{N}$ property $(t(k,l), t(k+1,l)) \in D_H$ is guaranteed by ψ_h since both $v_{k,l}$ and $v_{k+1,l}$ are successors of v_l. To check the other property, consider any $l \in \mathbb{N}$ and $k < n$. Since $v_l R v_{l+1}$, $v_{l+1} R v_{k,l+1}$, and $v_l R v_{k,l}$, Φ guarantees that we have $v_{k,l} R v_{k,l+1}$ and therefore ψ_v guarantees that $(t(k,l), t(k,l+1)) \in D_V$.

We showed that if φ has a model that satisfies Φ, then the domino problem has a solution. It is easy to see that the converse is also true. □

PROPOSITION 6.9. *For a given unbounded* UHF *formula Φ, if Φ does not have the tree-compatible model property and satisfies S2, then Φ-SAT_G and ϕ-SAT_L are* NP*-complete.*

PROOF. Let Φ be an unbounded UHF formula that does not have the tree-compatible model property and satisfies S2 for some χ, φ be a modal formula and $\mathfrak{M}$ be a tree-based model of Φ and φ. Let c be a product of all positive elements of χ and for a world w at the level $\mathfrak{g}^\Phi$ and $i > \mathfrak{g}^\Phi$, set C_i^w be the set of all descendants of w at the level i. According to Observation 4.20, we may assume that the size of each such set is polynomial in $|\varphi|$. Our goal is to show that for any w, it is enough to consider only polynomially many non-isomorphic sets C_i^w, which will make the algorithm described in the previous section polynomial.

The technique used in the previous section is not sufficient to prove containment in NP — now, it is not enough just to satisfy one formula of the form $\Diamond\psi$ at each level. We solve this problem in the following way: in each C_i^w, we put as many witnesses as possible. Note that since all worlds in C_i^w are equivalent, for any $v \in C_i^w$ we have $PR(v) = PR(C_i^w)$.

OBSERVATION 6.10. *Let w, v be worlds such that $v \in C_j^w$ for some j and let i be such that $\mathfrak{g}^\Phi < i < j$ and c divides $j - i$. If $UR(v) \subseteq PR(C_{i+1}^w)$ and $PR(v) = PR(C_i^w)$, then model obtained by adding a copy v' of v to C_i^w satisfies both Φ and φ.*

Note that the set of successors of v is a subset of the set of successors of v', and therefore v has all the needed witnesses. Moreover, the set of predecessors of v' is a subset of the set of predecessors of v, so v' does not violate any predecessor requirements. Finally, since v' does not add any new requirements, it should be clear that new model satisfies φ. Therefore the new model satisfies φ and, in an obvious way, Φ.

OBSERVATION 6.11. *Let w be a world at the level $\mathfrak{g}^\Phi$, $i > \mathfrak{g}^\Phi$, and $A = \{0, 1, \dots\}$ be a (possibly finite) set of consecutive numbers. Let $\mathcal{C} = \{C_{i+ac}^w : a \in A\}$ be such that for all $j, j' \in A$, $PR(C_j^w) = PR(C_{j'}^w)$ and $PR(C_{j+1}^w) = PR(C_{j'+1}^w)$. Then, we can define a set C' with $|C'| \leq |\varphi|$ such that each element of $\bigcup\mathcal{C}$ can be replaced by a copy of an element from C' in a way such that the obtained model is still a model of φ and Φ.*

Let $C = \bigcup\mathcal{C}$. We define a $C' \subseteq C$ in the following way. For every subformula of φ of the form $\Diamond\psi$, if there is a type t satisfying ψ such that t is realized in infinitely many elements of $\mathcal{C}$, then we take one world of this type and add it to C'. If there is no such type, but there is a world in C that satisfies ψ, then we find a maximal $j \in A$ such that there is such a world $v \in C_{i+jc}^w$ and we add v to C'. Clearly, $|C'| \leq |\varphi|$. Then, we define $C'\underline{^{i+jc}} = C' \cap \bigcup_{a \in A, a \geq j} C_{i+ac}^w$ and replace each C_{i+jc}^w by $C'\underline{^{i+jc}}$. The obtained model satisfies both φ and Φ.

Having Observation 6.11 we can prove the statement of the lemma. Let w be a world at the level $\mathfrak{g}^\Phi$ and i be such that $\mathfrak{g}^\Phi \leq i < \mathfrak{g}^\Phi + c$. The sequence $PR(C_i^w), PR(C_{i+c}^w), PR(C_{i+2c}^w) \dots$ never shrinks, and the same holds for $PR(C_{i+1}^w), PR(C_{i+c+1}^w), PR(C_{i+2c+1}^w), \dots$. Therefore, the sequence $C_i^w, C_{i+c}^w, C_{i+2c}^w$ can be split into at most $|\varphi|^2$ subsequences that satisfy the requirements of Observation 6.11, so the number of different sets of the form C_i^w can be bounded by $|\varphi|^3$. Taking into account all possible w and i, we can bound the number of possible sets C_i^w by $|\varphi|^{\mathfrak{g}^\Phi} \cdot c \cdot |\varphi|^3$, which is polynomial in φ. □

Table II. A summary of a results for the finite satisfiability problems for modal logic defined by consistent Horn formulae.

Properties of Φ	Φ-FINSAT$_G$	Φ-FINSAT$_L$
Φ is bounded	FMP, NP-c (6.2)	
Φ is unbounded and satisfies S3	FMP, NP-c (4.12, 6.12)	
Φ is unbounded and satisfies S2	NExpTime (7.12)	
Φ is unbounded, satisfies S1, and...		
has the TCMP and merges at some level	no FMP (7.3) PSpace-c (7.7, 7.8)	FMP (7.1) PSpace-c (6.3)
has the TCMP and does not merge at any level	FMP (7.9) ExpTime-c (6.3)	FMP (7.1) PSpace-c (6.3)
lacks the TCMP	FMP (7.10) PSpace-c (6.8)	FMP (7.1) NP-c (6.7)

Proposition 6.12. *For a given unbounded* UHF *formula Φ, if Φ does not have the tree-compatible model property and satisfies S3, then it has the polynomial model property.*

Proof. Suppose that Φ does not have the tree-compatible model property and satisfies S3 for some m. Observe that in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ for all i and $l \geq 0$, worlds $\underline{i}$ and $\underline{i+lm}$ are equivalent. Let $\mathfrak{M}$ be a model of φ. It follows from Lemma 4.16 that for all w at the level $\mathfrak{g}^\Phi$ and all i, all descendants of w at the levels $2\mathfrak{g}^\Phi+i$, $2\mathfrak{g}^\Phi+i+m$, $2\mathfrak{g}^\Phi+i+2m$, ... are equivalent. So we can remove all but polynomially many of them and obtain a smaller model that satisfies φ. By repeating this procedure for all such w we obtain a model of polynomial size in $|\varphi|$. □

7. FINITE MODELS

In practical applications of modal logic (as well as any other logic) one is often interested in finite structures because most structures about which we may want to reason, like knowledge bases or models of computations, are essentially finite. Thus, with respect to the decision problem, we are specially interested not only in the question, whether the formula is just satisfiable but if it is finitely satisfiable, i.e. if it has a finite model.

We say that modal logic has the *finite model property* (resp. *finite global model property*), FMP, *with respect to a class of frames* $\mathcal{K}$, if any formula that is locally (resp. globally) satisfiable over $\mathcal{K}$ is also finitely locally (resp. globally) satisfiable over $\mathcal{K}$.

We have already seen that with respect to classes defined by many UHF formulae modal logic has the finite and global finite model properties. Namely, it holds for all bounded formulae and all unbounded formulae satisfying S3. In such cases, the question about the existence of a finite model is equivalent to the question about the existence of any model, and therefore finite and unrestricted satisfiability problems coincide. In this section we show FMP in some other cases (for which it was not demonstrated in the previous sections, since it was not necessary for establishing decidability and precise complexity).

We show below that in some cases FMP fails. Nevertheless, we prove decidability of Φ-FINSAT$_L$ and Φ-FINSAT$_G$ for any UHF formula Φ. Table II summarises our results.

7.1. Formulae that do not force long edges

In this subsection we consider unbounded formulae $\Phi \in$ UHF that satisfy S1, i.e. do not force long edges. First, we show the (local) finite model property. We prove it essentially by an application of the standard selection argument.

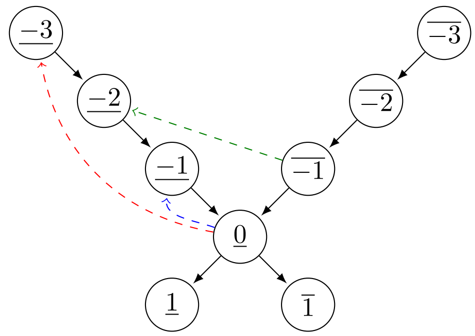

Fig. 6. A fragment of the frame $\mathcal{X}$ (circles and solid arrows). Consider a formula $\Phi = yRw \wedge wRv \wedge xRv \Rightarrow xRy$ that forces edge $(\overline{-1}, \underline{-2})$. When applied to $x = w = \underline{0}$, $y = \underline{-1}$ and $v = \underline{1}$, it forces edge $(\underline{0}, \underline{-1})$. Then, applied to $x = \underline{0}$, $v = \underline{-1}$, $w = \underline{-2}$ and $y = \underline{-3}$ it forces long edge $(\underline{0}, \underline{-3})$.

Proposition 7.1. *Each unbounded* UHF *formula* Φ *that does not force long edges has the finite model property in the local satisfiability case.*

Proof. Assume that φ is locally $\mathcal{K}_\Phi$-satisfiable. Let $\mathcal{T}$ be a tree of bounded degree guaranteed by Lemma 4.3. Thus, there exists a model $\mathfrak{M}$ based on the frame $\mathfrak{C}_\Phi(\mathcal{T}) \in \mathcal{K}_\Phi$, such that $\mathfrak{M}, w \models \varphi$, where w is the root of $\mathcal{T}$. Recall the morphism $h_\mathcal{T} : \mathcal{T} \to \mathcal{L}_\mathbb{Z}$. By observation 4.10, $h_\mathcal{T}$ is also a morphism from $\mathfrak{C}_\Phi(\mathcal{T})$ to $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$. Since Φ does not force long edges, $\mathfrak{C}_\Phi(\mathcal{T})$ has only edges between nodes on the same level or on two consecutive levels.

In order to obtain a finite model, we simply remove from $\mathfrak{M}$ all worlds from levels greater than $|\varphi|$. Since the truth of φ depends only on the worlds that are reachable from the root w by a path whose length is bounded by $|\varphi|$ (more precisely: by the modal depth of φ), the resulting model is a finite model of φ and, of course, it satisfies Φ since Φ is universal. $\square$

We showed that φ has a $\mathcal{K}_\Phi$-based model if and only if it has a finite $\mathcal{K}_\Phi$-based model, so Φ-SAT$_L$ is equal to Φ-FINSAT$_L$. Finite global satisfiability problem is much more complicated. In particular for some formulae Φ we will have to deal with infinite models.

In previous sections, where we dealt with unrestricted satisfiability we analised the behavior of the first-order formula on $\mathcal{T}_\infty$ and $\mathcal{L}_\mathbb{Z}$. For our current purposes we need one more frame, $\mathcal{X}$, that contains a world with in-degree 2. Formally, $\mathcal{X}$ is of the form $\langle W_X, R_X \rangle$, where $W_X = \{\underline{i} : i \in \mathbb{Z}\} \cup \{\overline{i} : i \in \mathbb{Z} \setminus \{0\}\}$ and $R_X = \{(\underline{i}, \underline{i+1}) : i \in \mathbb{Z}\} \cup \{(\overline{i}, \overline{i+1}) : i \in \mathbb{Z} \setminus \{-1, 0\}\} \cup \{(\overline{-1}, \underline{0}), (\underline{0}, \overline{1})\}$. Fig. 6 shows a fragment of $\mathcal{X}$.

We say that a formula Φ *merges* at a level $k < 0$ if in $\mathfrak{C}_\Phi(\mathcal{X})$ there is an edge from $\underline{k-1}$ to $\overline{k}$. For example, the formula $\Phi = xRz \wedge zRv \wedge yRv \Rightarrow xRy$ merges at the level -1. Note that $\mathcal{T}_\infty$ and $\mathcal{L}_\mathbb{Z}$ satisfy Φ.

We distinguish now three cases, corresponding to the last three rows of Table II. Recall that we say that Φ has TCMP if it forks at all levels. Each of the cases is treated in a separate subsection.

7.1.1. Formulae that have TCMP and merge.

The following lemma shows an important regularity in models of formulae that merge.

Lemma 7.2. *Let* Φ *be an unbounded* UHF *formula that does not force long edges, and merges at a level* k*,* $\mathfrak{M}$ *be a model of* Φ*,* $v_1, v_2, \ldots, v_l$ *be a walk (i.e. a path, but not necessarily simple) in* $\mathfrak{M}$ *such that all* v_i *are* ∞*-inner.*

(i) If $v_l R v_{l-c}$ *for some* $c > 0$*, then for all* $i > c$*,* $v_i R v_{i-c}$*.*
(ii) If $v_{l-c} R v_l$ *for some* $c > 0$*, then for all* $i > c$*,* $v_{i-c} R v_i$*.*

PROOF. Let $\dots, v_{-2}, v_{-1}, v_0, v_1$ and $v_l, v_{l+1}, \dots$ be infinite walks in $\mathfrak{M}$. Such walks exist since v_1 and v_l are ∞-inner.

We prove (i) by induction. Assume that for some $i > 0$, for all $j > i$ we have $v_j R v_{j-c}$. We define a morphism h from $\mathcal{X}$ into $\mathfrak{M}$ as follows

$$h(w) = \begin{cases} v_{i+s+1} & \text{if } w = \underline{k+s} \text{ for some } s \leq 0 \\ v_{i-c+s} & \text{if } w = \underline{k+s} \text{ for some } s > 0 \\ v_{i-c+s} & \text{if } w = \overline{k+s} \text{ for some } s \in \mathbb{Z} \end{cases}$$

A quick check shows that h is indeed a morphism and since $\mathfrak{C}_\Phi(\mathcal{X})$ contains an edge from $\underline{k-1}$ to $\overline{k}$, $\mathfrak{M}$ contains an edge from v_i to v_{i-c}. The proof of (ii) is similar. □

We use the above lemma to show the lack of the finite model property.

PROPOSITION 7.3. *Let Φ be an unbounded* UHF *formula that does not force long edges, merges at a level $k < 0$ and forks at all levels. Then modal logic lacks the finite global model property with respect to $\mathcal{K}_\Phi$.*

PROOF. Let $\lambda = \lambda_0 \wedge \lambda_1 \wedge \lambda_2 \wedge \lambda_3 \wedge \lambda_4$, where:

$$\lambda_0 = \bigvee_{i \in \{1,2,3,4\}} p_i \quad \wedge \bigwedge_{i,j \in \{1,2,3,4\}, i \neq j} \neg(p_i \wedge p_j)$$

$$\lambda_1 = p_1 \rightarrow (\Diamond p_2 \wedge \Box p_2) \qquad \lambda_2 = p_2 \rightarrow (\Diamond p_3 \wedge \Box p_3)$$

$$\lambda_3 = p_3 \rightarrow (\Diamond p_2 \wedge \Diamond p_4 \wedge \Box(p_2 \vee p_4)) \qquad \lambda_4 = p_4 \rightarrow (\Diamond p_1 \wedge \Box p_1)$$

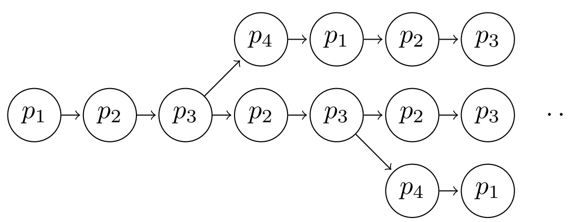

Fig. 7. An infinite model of τ.

An infinite model of λ is presented in Fig. 7. It is not hard to see that its frame belongs to $\mathcal{K}_\Phi$ for any Φ meeting the assumptions. Assume to the contrary that $\mathfrak{M}$ is a finite $\mathcal{K}_\Phi$-based model of λ. Take an arbitrary world w_1 satisfying p_3. Quick check shows that such a world must exist. Let $w_1, w_2, w_3, \dots$ be an infinite path in $\mathfrak{M}$ such that for odd i, w_i satisfies p_3, and for even i, w_i satisfies p_2. Such a path is guaranteed by λ_2 and λ_3. Since $\mathfrak{M}$ is finite, it must be the case that for some $0 < s < t$ we have $\mathfrak{M} \models w_t R w_s$. If s is odd (and thus w_s satisfies p_3) then let v_0 be a successor of w_s satisfying p_4. Otherwise let v_0 be a successor of w_{s+1} satisfying p_4. Such a world v_0 is guaranteed by λ_3. Note that v_0 if ∞-preceded. Let v_1 be a successor of v_0 satisfying p_1 (guaranteed by λ_4), let v_2 be a successor of v_1 satisfying p_2 (guaranteed by λ_1), and let $v_3, v_4, v_5, \dots$ be an infinite path such that v_i satisfies p_3 for odd i, and it satisfies p_2 for even i (existing due to λ_2 and λ_3). Again, since $\mathfrak{M}$ is finite, it must be the case that for some $0 < k < l$ we have $\mathfrak{M} \models v_l R v_k$. Please note that all of the elements $v_1, \dots, v_l$ are ∞-inner. It follows from Lemma 7.2 (i) that $\mathfrak{M} \models v_{l-k+1} R v_1$. But v_{l-k+1} satisfies p_2 or p_3, v_1 satisfies p_1, and thus λ_2 or λ_3 forbid this connection. Therefore there is no finite model of λ based on a frame from $\mathcal{K}_\Phi$. □

Before we show the algorithm that solves the finite satisfiability problem, we argue that we may restrict our attention to strongly connected models.

LEMMA 7.4. *Let $\Phi \in$ UHF, φ be a modal formula and $\mathfrak{M}$ be a finite $\mathcal{K}_\Phi$-based model such that $\mathfrak{M} \models \varphi$. Then there is a $\mathcal{K}_\Phi$-based submodel $\mathfrak{N}$ of $\mathfrak{M}$ such that $\mathfrak{N} \models \varphi$ and the frame of $\mathfrak{N}$ is strongly connected.*

PROOF. Consider a partial order on the set of strongly connected components of $\mathfrak{M}$, defined in such a way that $\mathfrak{N} \leq \mathfrak{N}'$ iff there is a path from an element of $\mathfrak{N}'$ to an element of $\mathfrak{N}$, or if $\mathfrak{N} = \mathfrak{N}'$. Since $\mathfrak{M}$ is finite, there must be a component $\mathfrak{N}_{min}$ which is minimal with respect to $\leq$. As Φ is universal, $\mathfrak{N}_{min}$ satisfies Φ. Moreover, since each world from $\mathfrak{N}_{min}$ has all its successors in $\mathfrak{N}_{min}$ (there are no paths to worlds in other connected components) it is a generated substructure of $\mathfrak{M}$, and thus, by Lemma 3.2, $\mathfrak{N}_{min} \models \varphi$. $\square$

We say that a frame $\mathcal{M}$ is *k-periodic* if its universe can be divided into pairwise disjoint, non-empty sets of worlds $W_1, W_2, \ldots, W_k$ such that for each v, w from $\mathcal{M}$ there is an edge from v to w if and only if for some $i \leq k$, $v \in W_i$ and $w \in W_{(i \mod k)+1}$. Notice that a 1-periodic frame is a clique. For each $k \in \mathbb{N}$ we define the cycle $\mathcal{C}_k$ as $\mathcal{I}_k$ with one additional edge, namely $(\underline{k-1}, \underline{0})$. Clearly, each $\mathcal{C}_k$ is k-periodic.

We are going to prove decidability by showing that each satisfiable formula has a model that is k-periodic for some k. In order to do so, we introduce two technical lemmas.

LEMMA 7.5. *Let $\Phi \in$ UHF.*

(a) If Φ has a k-periodic model $\mathcal{M}$, then $\mathcal{C}_k$ is a model of Φ.
(b) If $\mathcal{C}_k$ is a model of Φ, then any k-periodic frame is a model of Φ.
(c) If $\mathcal{L}_\mathbb{Z}$ is a model of Φ, then for all $c > |\Phi|$, $\mathcal{C}_c$ is a model of Φ.
(d) If for some $k > |\Phi|$ the frame $\mathcal{C}_k$ is a model of Φ, then $\mathcal{L}_\mathbb{Z}$ is a model of Φ.

PROOF. For (a), observe that if a periodic model $\mathcal{M}$ that consists of sets $W_1, W_2, \ldots, W_k$ is a model of Φ, then $\mathcal{C}_k$ is isomorphic with an induced substructure of $\mathfrak{M}$ that contains one world from every W_i.

We say that a morphism $h : \mathcal{M} \to \mathcal{M}'$ is *complete* if for all v, v' we have $h(v)Rh(v')$ if and only if vRv'. Note that if there is a complete morphism $h : \mathcal{M} \to \mathcal{M}'$ and Φ does not hold in $\mathcal{M}$, then it does not hold in $\mathcal{M}'$.

For (b), assume that there is a periodic frame $\mathcal{M}$ that consists of sets $W_1, W_2, \ldots, W_k$ and is not a model of Φ, but $\mathcal{C}_k$ is a model of Φ. We define a complete morphism $f : \mathcal{M} \to \mathcal{C}_k$ as $f(v) = \underline{i}$ for $v \in W_i$. Since Φ does not hold in $\mathcal{M}$ and f is a complete morphism, Φ does not hold in $\mathcal{C}_k$ — a contradiction.

We prove (c) as follows. Let $c > |\Phi|$. Assume that there is a clause Ψ satisfied in $\mathcal{L}_\mathbb{Z}$ but not in $\mathcal{C}_c$, and let $v_1, v_2, \ldots, v_n$ be worlds of $\mathcal{C}_c$ such that $\Psi(v_1, \ldots, v_n)$ is false. Let k be such that no world among $v_1, \ldots, v_n$ is equal $\underline{\mathrm{k}}$.

Consider the function $f : \mathcal{C}_{c\restriction\{v_1,\ldots,v_n\}} \to \mathcal{L}_\mathbb{Z}$ defined as $f(\underline{s}) = \underline{s}$ for $s > k$ and $f(\underline{s}) = \underline{c+s}$ for $s < k$. A quick check shows that the function f is a complete morphism. Since $\Psi(v_1, \ldots, v_n)$ does not hold in $\mathcal{C}_c$, it follows that $\Psi(f(v_1), \ldots, f(v_n))$ does not hold in $\mathcal{L}_\mathbb{Z}$. But $\mathcal{L}_\mathbb{Z} \models \Psi$, a contradiction.

For the proof of (d), let $k > |\Phi|$, $\Psi \Rightarrow \Psi'$ be satisfied in $\mathcal{C}_k$ but not in $\mathcal{L}_\mathbb{Z}$. Let $v_1 = \underline{s}, v_1 = \underline{t}, v_3 \ldots, v_n$ be worlds of $\mathcal{L}_\mathbb{Z}$ such that $\Psi(v_1, \ldots, v_n)$ is true, $\Psi'(v_1, \ldots, v_n)$ is not, and $|s-t|$ is minimal. Let $f(\underline{i}) = \underline{i \mod k}$ be a morphism from $\mathcal{L}_\mathbb{Z}$ onto $\mathcal{C}_k$. If $t - s \mod k \neq 1$, then $\Psi \Rightarrow \Psi'(f(v_1), \ldots, f(v_n))$ does not hold and we have a contradiction. Otherwise, $|s-t| \geq k-1$ so there is a world $\underline{l}$ such that l is between s and t and $\underline{l}$ is different from all of $\underline{s}, \underline{t}, v_3, \ldots, v_n$. But then, morphism $g : \mathcal{L}_{\mathbb{Z}\restriction\{v_1,\ldots,v_n\}} \to \mathcal{L}_\mathbb{Z}$ defined as $g(\underline{s}) = \underline{s}$ for $s < l$ and $g(\underline{s}) = \underline{s-1}$ leads to the contradiction with the minimality of $|s-t|$. $\square$

LEMMA 7.6. *Let Φ be an unbounded* UHF *formula that does not force long edges and such that in $\mathfrak{C}_\Phi(\mathcal{X})$ for some $i, j < 0$ we have $\overline{i}R\underline{j}$ or $\underline{i}R\overline{j}$. Then $j - i = 1$ and $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z}) = \mathcal{L}_\mathbb{Z}$.*

PROOF. As $\mathcal{X}$ has a symmetric shape, $\mathfrak{C}_\Phi(\mathcal{X}) \models \underline{i}R\overline{j}$ implies $\mathfrak{C}_\Phi(\mathcal{X}) \models \overline{i}R\underline{j}$. So, wlog. we assume that $\mathfrak{C}_\Phi(\mathcal{X}) \models \underline{i}R\overline{j}$. Let us consider the morphism f from $\mathcal{X}$ into $\mathcal{L}_\mathbb{Z}$ defined as

$$f(\underline{k}) = f(\overline{k}) = \underline{k}$$

If $|j - i| > 1$, then there is a long edge in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ and it contradicts the assumption that Φ does not force long edges. If $j - i = -1$, then the morphism f implies that there is an edge $(\underline{j}, \underline{j-1})$ in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ and, since $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ is uniform, for all k there are edges $(\underline{k}, \underline{k-1})$ in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$. We define another morphism g to show that then $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ contains a long edge. Let g be the morphism from $\mathcal{X}$ into $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ defined as

$$g(w) = \begin{cases} \underline{|k|} & \text{if } w = \underline{k} \text{ for some } k \\ \underline{-|k|} & \text{if } w = \overline{k} \text{ for some } k \end{cases}$$

It is not hard to see that g is indeed a morphism and therefore that $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ contains a long edge $(\underline{|i|}, \underline{-|j|})$. An example is presented in Fig. 6.

If $j = i$, then the morphism f implies that there is a reflexive world in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$, and therefore all worlds are reflexive. Consider a morphism h from $\mathcal{X}$ into $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ defined as

$$h(w) = \begin{cases} \underline{1} & \text{if } w = \overline{k} \text{ for some } k \leq i \\ \underline{0} & \text{otherwise} \end{cases}$$

Since all worlds in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ are reflexive, h is indeed a morphism, so in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ there is edge $(\underline{1}, \underline{0})$ and, as in the previous case, all edges in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ are symmetric and therefore $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ contains a long edge.

For the proof of $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z}) = \mathcal{L}_\mathbb{Z}$, recall that if $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ contains a symmetric or reflexive edge, then it contains long edges. But Φ does not force long edges, and therefore $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z}) = \mathcal{L}_\mathbb{Z}$. □

For a given model $\mathfrak{M}$, we define a *characteristic cycle* of $\mathfrak{M}$ as a walk $v_0, v_1, \ldots, v_{l-1}$ that contains all worlds from $\mathfrak{M}$ and, moreover, in $\mathfrak{M}$ there is an edge from v_{l-1} to v_0. Note that every strongly connected model with at least two worlds contains a characteristic cycle.

PROPOSITION 7.7. *Let Φ be an unbounded* UHF *formula that does not force long edges, merges at a level $k < 0$ and forks at all levels. Then Φ-FINSAT$_G$ is in* PSPACE.

PROOF. Let φ be a modal formula and $\mathfrak{M}$ be a strongly connected model of φ from $\mathcal{K}_\Phi$. Such a model exists due to Lemma 7.4. Assume that $\mathfrak{M}$ contains at least two worlds and let $v_0, v_1, \ldots, v_{l-1}$ be a characteristic cycle of $\mathfrak{M}$. For better readability, below we omit " mod l" in subscripts of vs.

Our aim is to show that $\mathfrak{M}$ is s-periodic for some s.

Let $\chi^{\mathfrak{M}} \subseteq \mathbb{N}$ be such that $k \in \chi^{\mathfrak{M}}$ if and only if there is v_i such that $\mathfrak{M} \models v_i R v_{i+k+1}$. Lemma 7.2 implies that for all v_i and $k \in \chi^{\mathfrak{M}}$, $\mathfrak{M} \models v_i R v_{i+k+1}$.

We show that $\chi^{\mathfrak{M}}$ is additively closed. Assume that $x, y \in \chi^{\mathfrak{M}}$. It means that $\mathfrak{M}$ contains edges (v_{x+y+1}, v_{x+y+2}) , (v_{x+1}, v_{x+y+2}) and (v_0, v_{x+1}). We define a morphism h from $\mathcal{X}$ to $\mathcal{M}$, the frame of $\mathfrak{M}$, as

$$h(w) = \begin{cases} v_s & \text{if } w = \underline{k-1+s} \text{ for all } s \leq 0 \\ v_{x+1} & \text{if } w = \underline{k} \\ v_{x+y+1+s} & \text{if } w = \underline{k+s} \text{ for all } s > 0 \\ v_{x+y+1+s} & \text{if } w = \overline{k+s} \text{ for all } s \in \mathbb{Z} \end{cases}$$

We see that $h(\underline{k-1}) = v_0$ and $h(\overline{k}) = v_{x+y+1}$, and since in $\mathcal{X}$ there is an edge from $\underline{k-1}$ to $\overline{k}$ it implies that in $\mathcal{M}$ there is an edge from v_0 to v_{x+y+1} and thus $x + y \in \chi^{\mathfrak{M}}$.

Let $\chi_l^{\mathfrak{M}} = \{i \mod l : i \in \chi^{\mathfrak{M}}\}$. By Fact 5.1, $\chi_l^{\mathfrak{M}}$ can be represented as $\{i \cdot gcd(\chi^{\mathfrak{M}}) \mod l : i \in \mathbb{N}\}$. Define $W_i = \{v_{i+j\cdot gcd(\chi^{\mathfrak{M}})} : j \in \mathbb{N}\}$. It follows that all elements of W_i have all successors in W_{i+1}, and therefore $\mathfrak{M}$ is $gcd(\chi^{\mathfrak{M}})$-periodic.

We are now going to compress the sets W_i. For each i and each subformula ψ of φ, if there is a world in W_i that satisfies ψ, we mark one such world. Then we remove all unmarked worlds. It is easy to see that the types of worlds remain the same.

We have proved that all models of φ are s-periodic and that their sets can be compressed to a size bounded by $|\varphi|$, but the value of s can be arbitrary large. We show that there is an NPSPACE (=PSPACE) procedure that checks, for a given modal formula φ, if φ has a Φ-based finite global periodic model.

Our NPSPACE algorithm works as follows. First, it checks if there is a single world or a single clique (1-periodic set) with size bounded by $|\varphi|$, that satisfies both φ and Φ. If it is the case the algorithm returns "Yes". Otherwise, it guesses a set W_1 of size bounded by $|\varphi|$ and then, recursively, guesses the successive sets of size similarly bounded, checking if guessed worlds are consistent with their predecessor, and returns "no" otherwise. The algorithm stops after $\binom{2^{|\varphi|}}{|\varphi|} + 1$ steps and returns "yes".

If there is a model of φ, then the algorithm returns "yes". Indeed, we showed that φ has a single world model or an s-periodic model with size of sets bounded by $|\varphi|$, and the algorithm can simply guess this world or successively guess consecutive sets of this model.

If the algorithm returns "yes", then it visited two sets satisfying the same subformulae, so there is a sequence of sets $V_1, V_2, \ldots, V_k, V_1$ with $k \leq 2^{|\varphi|}$ such that each set contains all witnesses needed by its predecessors. We build an s-periodic model that contains sets $V_1, \ldots, V_k$ repeated $\lceil |\Phi|/k \rceil + 1$ times. Clearly, the obtained model satisfies φ. By Lemma 7.6, $\mathcal{L}_{\mathbb{Z}} = \mathfrak{C}_\Phi(\mathcal{L}_{\mathbb{Z}})$, and by Lemma 7.5 it is also a model of Φ. □

We show the corresponding lower bound by encoding of the bounded-space domino problem.

PROPOSITION 7.8. *Let Φ be an unbounded* UHF *formula that does not force long edges, merges at a level $k < 0$ and forks at all levels. Then Φ-FINSAT$_G$ is* PSPACE*-hard.*

PROOF. Let $\mathcal{D} = (D, D_H, D_V, n)$ be an instance of the bounded-space domino problem. We will construct a modal formula η which is globally, finitely $\mathcal{K}_\Phi$-satisfiable iff $\mathcal{D}$ has a solution. In our intended model a single world represents a whole row of a solution.

We employ propositional variables p_i^d for $i < n$ and $d \in D$. The intended meaning of p_i^d is that the point in column i is tiled by d.

We put $\eta = \eta^l \wedge \eta^h \wedge \eta^v$, where η^l that guarantees that each point is tiled by exactly one element of D, η^h ensures that the tilling respects D_H, and η^v ensures that each world has a successor that describes the row which is consistent with the current one with respect to the relation D_V.

$$\eta^l = \bigwedge_{i<n} \left(\bigvee_{d \in D} p_i^d \wedge \bigwedge_{d,d' \in D, d \neq d'} \neg(p_i^d \wedge p_i^{d'}) \right)$$

$$\eta^h = \bigwedge_{i<n-1} \bigvee_{(d,d') \in D_H} (p_i^d \wedge p_{i+1}^{d'})$$

$$\eta^v = \Diamond\top \wedge \bigwedge_{i<n} \bigvee_{(d,d') \in D_V} (p_i^d \wedge \Box p_i^{d'})$$

Assume that $\langle D, D_H, D_V, n\rangle$ has a solution that consists of rows $r_1, r_2, \ldots$. Then among first n^n+1 of them some rows r_i, r_j with $i<j$ are tiled identically. Let $l = c(j-i)$, for some $c > |\Phi|$. We encode the solution on $\mathcal{C}_l$ in such a way that $\underline{s}$ represents row $i + (s \mod l)$. Note that by Lemma 7.6 and Lemma 7.5(c) it follows that $\mathcal{C}_l$ belongs to $\mathcal{K}_\Phi$. Conversely, if η has a model $\mathfrak{M}$ then we can construct a solution by starting from an arbitrary world of $\mathfrak{M}$, translating it to the initial row of a solution in a natural way, and recursively building successive rows as translations of the worlds guaranteed by η^v. $\square$

7.1.2. Formulae that do not merge and have TCMP. We prove that in the case of formulae Φ that do not force long edges, fork at all levels and do not merge at any level, modal logic has the finite global model property with respect to $\mathcal{K}_\Phi$. In the proof we start from a possibly infinite tree–based model $\mathfrak{M}$, and construct a very large structure that locally looks like a part of $\mathfrak{M}$, but is finite. We need to do it carefully in order not to violate the first-order formula Φ.

PROPOSITION 7.9. *Let Φ be an unbounded* UHF *formula that does not force long edges, has TCMP and does not merge at any level $k<0$. Then modal logic has the finite global model property with respect to $\mathcal{K}_\Phi$.*

PROOF. Let $\mathfrak{M}^b$ be a model of φ and Φ based on a tree $\mathcal{T}^b$ guaranteed by Lemma 4.3. Let $n = |\varphi|$ and $N = |\Phi|$. If there is a world in $\mathfrak{M}^b$ without a proper successor, then the structure that contains only this world is a model of φ and Φ. Otherwise, all worlds are ∞-followed. We assume that every world has degree n – if a world has a smaller degree, then we can replicate any of its subtrees.

Let w be any $\mathfrak{g}^\Phi$-inner world in $\mathcal{T}^b$, $\mathcal{T}$ be a subtree of $\mathcal{T}^b$ rooted at w, and $\mathfrak{M}$ be a substructure of $\mathfrak{M}^b$ that consists of the worlds from $\mathcal{T}$. Clearly, $\mathfrak{M}$ satisfies Φ and φ.

Let M be the universe of $\mathfrak{M}$. For each $w \in M$, we define $\mathcal{S}'_w$ to be the subtree of $\mathcal{T}$ rooted at w, $\mathcal{S}_w$ to be the frame that contains first $2N$ levels of $\mathcal{S}'_w$, and $\mathfrak{S}_w$ to be the substructure of $\mathfrak{M}$ that contains the worlds from $\mathcal{S}_w$. Let $tp(\mathfrak{M})$ be a set of all types realized in $\mathfrak{M}$. For each type $t \in tp(\mathfrak{M})$, we pick one world w_t of this type and define $\mathfrak{S}_t = \mathfrak{S}_{w_t}$ and $\mathcal{S}_t = \mathcal{S}_{w_t}$.

For each $\mathcal{S}_t$, we label the leaves of $\mathcal{S}_t$ in a consecutive way, e.g. from left to right, such that the leaves labeled with 1, 2, …, n have the same parent and so on.

For each $s \in \{0,1\}$, $p \in \{1,\ldots,n\}$ and $t \in tp(\mathfrak{M})$, we define $\mathfrak{T}^s_{t,p}$ as a copy of $\mathfrak{S}_t$. We define the finite structure $\mathfrak{M}_s$ as a disjoint union of all possible $\mathfrak{T}^s_{t,p}$. We say that a world w is *at the level k* in $\mathfrak{T}^s_{t,p}$ if it is a copy of a world that is at the level k in $\mathcal{S}_t$ and that it is *at the level k* in $\mathfrak{M}_s$ if it is at the level k in some tree of $\mathfrak{M}_s$. We say that a world v is a *parent* of v' in $\mathfrak{M}_k$ if wRv, v is at the level k and v' is at the level $k+1$ for some k. For any two worlds v, v' that are in the same tree, we define $lca(v,v')$ as the lowest common ancestor of v and v' (w.r.t. the relation parent). We define $llca(v,v')$ as the level of $lca(v,v')$ if such a world exists and $llca(v,v') = -1$ otherwise.

We define $\mathfrak{M}'$ as the disjoint union of $\mathfrak{M}_0$ and $\mathfrak{M}_1$ with additional edges defined as follows. Consider tree $\mathfrak{T}^0_{t,p}$ and its leaf v labeled by p. Let w be a world in $\mathfrak{M}$ with the same type and $t_1,\ldots,t_k$ be types of successors of w in $\mathcal{T}$. For each $j \le k$ we add an edge from v to the root of $\mathfrak{T}^1_{t_j,p}$ and, if some connection between w and its successors is symmetric, we make this edge symmetric as well. We do the same for the leaves from $\mathfrak{M}_1$, but we connect them with the roots from $\mathfrak{M}_0$.

It is not hard to see that all worlds in $\mathfrak{M}'$ satisfy φ. We prove that $\mathfrak{M}'$ satisfies Φ. Assume to the contrary that this is not the case. Let $\Psi \Rightarrow \Psi'$ be a formula which is not satisfied in $\mathfrak{M}'$. Then there are worlds $v_1,\ldots,v_n$ such that $\Psi(v_1,\ldots,v_n)$ holds but $\Psi'(v_1,\ldots,v_n)$ does not. Define the function $\nu_k : \mathfrak{M}' \to \{0, \dots, 4N-1\}$ as

$$\nu_k(v) = \begin{cases} s-k & \text{for each } v \text{ at a level } s \geq k \text{ in } M_0 \\ s+2N-k & \text{for each } v \text{ at a level } s \text{ in } M_1 \\ s+4N-k & \text{for each } v \text{ at a level } s < k \text{ in } M_0 \end{cases}$$

Let k be such that no world among $v_1, \dots, v_n$ is at the level k in $\mathfrak{M}_0$ and $\mathfrak{M}_1$. The function $f : \mathfrak{M}'{\restriction_{\{v_1,\dots,v_n\}}} \to \mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ defined as $f(v) = \underline{\nu_k(v)}$ is a morphism.

The case $\Psi' = \bot$ is impossible. Indeed, since Φ is unbounded, $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ is a model of Φ and Ψ'. Similarly, if $\Psi' = xRx$, then some world in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ would be reflexive and, since all worlds in $\mathfrak{M}$ are $\mathfrak{g}^\Phi$-inner in $\mathfrak{M}^b$, $\Psi'(v_1, \dots, v_n)$ would be satisfied.

The only remaining case is $\Psi' = xRy$. Let v_1 be at a level l_1 in $\mathfrak{M}_{s_1}$ and v_2 be at a level l_2 is $\mathfrak{M}_{s_2}$. There are two cases: either $s_1 = s_2$ and $|l_1 - l_2| \leq 1$, or $s_1 \neq s_2$ and one of v_1, v_2 is a root and the other one is a leaf. Otherwise, Φ would force long edges.

Assume that $s_1 < s_2$ and let k be such that no world among $v_1, \dots, v_n$ is at a level k in $\mathfrak{M}_0$. Consider a morphism $g : \mathfrak{M}'{\restriction_{\{v_1,\dots,v_n\}}} \to \mathfrak{M}'$ defined as

$$g(v) = \begin{cases} v' & \text{if } v \text{ is at a level } i \geq k \text{ in } \mathfrak{M}_0 \text{ and } v' \text{ is a parent of } v \\ v & \text{otherwise} \end{cases}$$

It implies that Φ requires also an edge from some world that is not a leaf to some root, and so by the morphism f we can show that Φ forces long edges. The case $s_1 > s_2$ is symmetric.

Assume that $s_1 = s_2 = 0$. If $v_1 = v_2$, then, by the morphism f, all worlds of $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ are reflexive and Ψ' would be satisfied, as before. If v_2 is a parent of v_1, then, by the morphism f, all edges in $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z})$ are symmetric and Ψ' would be satisfied. So we can assume that v_1 and v_2 are not on the same path in $\mathfrak{M}_0$.

Assume that $l_1 \leq N$ and $l_2 \leq N$ and let $k > N$ be such that no world among $v_3, \dots, v_N$ is at the level k in $\mathfrak{M}_0$. We define a morphism $h_1 : \mathfrak{M}'{\restriction_{\{v_1,\dots,v_n\}}} \to \mathcal{T}_\infty$ as follows.

$$h_1(v) = \begin{cases} \underline{0^{\nu_k(v)}} & \text{if } \nu_k(v) < 4N-k \\ \underline{0^{4N-k+llca(v,v_1)}1^{s-llca(v,v_1)}} & \text{if } v \text{ at the level } s \text{ and } \nu_k(v) \geq 4N-k \end{cases}$$

Let $m = llca(v_1, v_2)$. Since v_1 and v_2 are not on the same path, $m < min(l_1, l_2)$. Since $h_1(v_1) = \underline{0^{4N-k+l_1}}$ and $h_1(v_2) = \underline{0^{4N-k+m}1^{l_2-m}}$ and h_1 is a morphism, it implies that Φ does not fork a the level $\underline{0^{4N-k+m}}$ — a contradiction.

Consider the case when $l_1 \geq N$ and $l_2 \geq N$. Let $k < N$ be such that no world among $v_3, \dots, v_N$ is at the level k in $\mathfrak{M}_0$.

If $llca(v_1, v_2) \leq k$, then Φ merges at some level. We prove it using the following morphism $h_2 : \mathfrak{M}'{\restriction_{\{v_1,\dots,v_n\}}} \to \mathcal{X}$. Let $\mathfrak{T}^0_{t,p}$ be the tree that contains v_1.

$$h_2(v) = \begin{cases} \underline{s-2N} & \text{if } v \text{ at a level } s \geq k \text{ in } \mathfrak{M}_0 \text{ and } llca(v_1, v) > k \\ \overline{\underline{s-2N}} & \text{if } v \text{ at a level } s \geq k \text{ in } \mathfrak{M}_0 \text{ and } llca(v_1, v) \leq k \\ \underline{s} & \text{if } v \text{ at a level } s \text{ in } \mathfrak{M}_1 \\ \underline{2N+s} & \text{if } v \text{ at a level } s \text{ in } \mathfrak{M}_0 \end{cases}$$

It is readily checkable that h_2 is a morphism and it implies that Φ merges at some level.

Let $llca(v_1, v_2) > k$. We prove that Φ does not fork at some level. To this end, let k' be such that no world among $v_3, \dots, v_N$ is at the level k' in $\mathfrak{M}_1$. We define $V_1 = V_{M0} \cup V_{M1}$ as follows. Set $v \in V_{M0}$ if and only if v is at a level $s > k$ in $\mathfrak{M}_0$ and $lcm(v_1, v) \in \{v_1, v\}$ (in other worlds, v is an ancestor or descendant of v_1 in $\mathfrak{M}_0$). Finally, for each leaf w from V_{M0} labeled by m and each $t \in tp(\mathfrak{M})$, V_{M1} contains all worlds from levels less than k' in $\mathfrak{T}^1_{t,m}$.

Let $t = llca(v_1, v_2) - k$, We define a morphism $h_3 : \mathfrak{M}'{\restriction_{\{v_1,\ldots,v_n\}}} \to \mathcal{T}_\infty$.

$$h_3(v) = \begin{cases} \underline{0^{\nu_k(v)}} & \text{if } v \in V_1 \text{ or } \nu_k(v) < t \\ \underline{0^t 1^{\nu_k(v)-t}} & \text{otherwise} \end{cases}$$

It is readily checkable that h_3 is a morphism and it implies that Φ does not fork at the level t. The case when $s_1 = s_2 = 1$ is symmetric. □

7.1.3. Formulae that do not have TCMP. In the case of formulae that do not force long edges and do not fork at some level, the finite model property follows from the fact that each satisfiable formula has a k-periodic model for some k.

PROPOSITION 7.10. *Let Φ be an unbounded* UHF *formula that does not force long edges and does not fork at some level $k > 0$. Then modal logic has the finite global model property with respect to $\mathcal{K}_\Phi$.*

PROOF. Let $\mathfrak{M}$ be defined as in the proof of Proposition 7.9. First, observe that $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z}) = \mathcal{L}_\mathbb{Z}$ and, since Φ is unbounded, $\mathcal{L}_\mathbb{Z}$ is a model of Φ. Let v be a world at level $\mathfrak{g}^\Phi$ and let $\mathfrak{M}'$ be the model that consists of all descendants of v from levels greater than $2\mathfrak{g}^\Phi$. By Lemma 4.16, all worlds in $\mathfrak{M}'$ at the same level are equivalent. Since the number of types is bounded, there exist two levels k, l in $\mathfrak{M}'$ such that $k - l > |\Phi| + 1$ and the sets of types realized at levels k and l are equal. We create model $\mathfrak{M}''$ by removing all worlds from levels greater than or equal to k, and connecting all worlds from level $k-1$ to worlds from level l. Finally, we define $\mathfrak{M}'''$ by taking for each level one world of each type realized at this level. A quick check shows that models $\mathfrak{M}'$, $\mathfrak{M}''$, and $\mathfrak{M}'''$ satisfy φ and that $\mathfrak{M}'''$ is finite.

We justify that $\mathfrak{M}'''$ is a model of Φ. Since $\mathcal{L}_\mathbb{Z}$ is a model of Φ, Lemma 7.5 shows that C_{k-l} is a model of Φ, and the same lemma shows that therefore any $k-l$–periodic model is a model of Φ. Model $\mathfrak{M}'''$ is obviously $k-l$-periodic. □

7.2. Formulae that force long edges

As mentioned earlier, for formulae that satisfy S3, the polynomial model property follows from the previous results. Thus, the rest of this section is devoted to unbounded formulae $\Phi \in$ UHF that satisfy S2. First, observe that for some such formulae Φ modal logic may lack the finite model property (local and global) with respect to $\mathcal{K}_\Phi$. Consider e.g. the formula $(xRz_1 \wedge z_1Ry \Rightarrow xRy) \wedge (xRx \Rightarrow \bot)$ (defining the class transitive and irreflexive frames) and the modal formula $\Diamond\top \wedge \Box\Diamond\top$. A quick check shows that all models of these formulae are infinite (in local and global cases). On the other hand, it is known that modal logic has the finite model property with respect to the class defined by $(xRz_1 \wedge z_1Ry \Rightarrow xRy) \wedge xRx$, i.e. the class of transitive and reflexive frames. We are not going to determine all first-order formulae Φ leading to the finite model property. However, for each Φ satisfying S2 we establish decidability of the corresponding finite satisfiability problem by proving that if a modal formula φ has a finite model (in local or global case) in $\mathcal{K}_\Phi$, then it has a model in $\mathcal{K}_\Phi$ of size bounded by $|\varphi|^{O(|\varphi|)}$. Clearly, it leads to a NEXPTIME algorithm that simply guesses such a small model and verifies it.

Consider a modal formula φ and its $\mathcal{K}_\Phi$-based model $\mathfrak{M}$ with universe M. We say that a world w is *redundant* for φ and $\mathfrak{M}$ if $\mathfrak{M}{\restriction_{M\setminus\{w\}}}$ is a model of φ. We prove the following lemma by showing that a model that is large enough has to contain a redundant world.

LEMMA 7.11. *Let Φ be an unbounded* UHF *formula satisfying S2. If φ has a finite $\mathcal{K}_\Phi$-based model, then it has a $\mathcal{K}_\Phi$-based model of size bounded by $|\varphi|^{O(|\varphi|)}$.*

PROOF. Let Φ be an unbounded UHF formula that satisfies S2 for some χ (as in Def. 4.4), and φ be a modal formula with a $\mathcal{K}_\Phi$-based model (local or global) $\mathfrak{M}$. Let c be any positive element of χ. Observe that for all $i \in \mathbb{Z}$ and $k \geq 0$ we have $\mathfrak{C}_\Phi(\mathcal{L}_\mathbb{Z}) \models \underline{i}R\underline{i + kc + 1}$.

We start from bounding the number of worlds that are not $\mathfrak{g}^{\Phi}$-preceded. We use the standard selection technique [Blackburn et al. 2001] — we start from an arbitrary world that satisfies φ, and then recursively for each world added in the previous stage we pick at most $|\varphi|$ witnesses. Let $\mathfrak{M}'$ be a model obtained this way. We define *the royal part* of $\mathfrak{M}'$ as the set of worlds that contain all worlds that are not $\mathfrak{g}^{\Phi}$-preceded and *the court* as the set of the $\mathfrak{g}^{\Phi}$-preceded worlds that were added as witnesses for some worlds from the royal part. Clearly, the total size of the royal part and the court are bounded by $|\varphi|^{\mathfrak{g}^{\Phi}+1}$.

Let w be a $\mathfrak{g}^{\Phi}$-inner world not from the court such that for each subformula $\Diamond\psi$ of φ such that ψ is satisfied in w there exists a $\mathfrak{g}^{\Phi}$-inner world $w_{\psi} \neq w$ that satisfies ψ and that there is a path from w to w_{ψ} with the length cj for some j. We show that w is redundant.

Consider any predecessor w' of w. If w' is not $\mathfrak{g}^{\Phi}$-preceded, then it has all the required witnesses in the court and the royal part. Otherwise, let ψ be a subformula of φ such that w satisfies ψ. We show that there is an edge from w' to w_{ψ}. To this end, consider a path $v_1, v_2, \ldots, v_{\mathfrak{g}^{\Phi}},\ w', w,\ v'_1, v'_2, \ldots, v'_{cj},\ w_{\psi}, v''_1, v''_2, \ldots, v''_{\mathfrak{g}^{\Phi}}$. Such a path exists since w' is $\mathfrak{g}^{\Phi}$-preceded and w_{ψ} is $\mathfrak{g}^{\Phi}$-inner, and there is a straightforward morphism from $\mathcal{I}_{2\mathfrak{g}^{\Phi}+2+cj}$ into this path. So it is enough to show that there is an edge from $\underline{\mathfrak{g}^{\Phi}+1}$ to $\underline{\mathfrak{g}^{\Phi}+1+cj+1}$ in $\mathfrak{C}_{\Phi}(\mathcal{I}_{2\mathfrak{g}^{\Phi}+2+j})$. By earlier observations, $\mathfrak{C}_{\Phi}(\mathcal{L}_{\mathbb{Z}})$ contains an edge from $\underline{\mathfrak{g}^{\Phi}+1}$ to $\underline{\mathfrak{g}^{\Phi}+1+cj+1}$, and Lemma 4.11 implies that there is an edge from $\underline{\mathfrak{g}^{\Phi}+1}$ to $\underline{\mathfrak{g}^{\Phi}+1+cj+1}$ in $\mathfrak{C}_{\Phi}(\mathcal{I}_{2\mathfrak{g}^{\Phi}+2+cj})$.

By iterating the above argument we can remove all $\mathfrak{g}^{\Phi}$—inner worlds except for at most $|\varphi|^{c\cdot|\varphi|}$ worlds. Finally, we again use the selection technique to bound the number of worlds that are not $\mathfrak{g}^{\Phi}$-followed by $|\varphi|^{c\cdot|\varphi|}\cdot|\varphi|^{\mathfrak{g}^{\Phi}}$. Since Φ is not a part of an instance, we reduced the number of worlds to $|\varphi|^{O(|\varphi|)}$. □

PROPOSITION 7.12. *If Φ is an unbounded* UHF *formula that forces long edges, then Φ-FINSAT$_L$ and Φ-FINSAT$_G$ are in* NEXPTIME.

Establishing precixe complexity in the case of formulae satisfying S2 remains open.

8. CONCLUSION AND FUTURE WORK

The main result of this paper states that the satisfiability problem of modal logic is decidable over any class of frames definable by universal first-order Horn formulae. The result works for four classes of the decidability problem: local satisfiability, global satisfiability, local finite satisfiability, and global finite satisfiability. We also provided the complexity results summarized in Tables I and II.

We proved that the result is optimal in two ways. First, we showed that there is a very simple universal first-order formula with only three variables that defines the class of frames such that modal logic over this class is undecidable. Furthermore, we proved a similar result for the class of frames definable by universal Horn formula in the case of bimodal logic.

All the results here are proved for the first-order formulae without equality. In [Michaliszyn and Otop 2012], we described how the results can be extended for the formulae containing equality. However, adding the equality makes the construction more complicated, and we decided to keep it as light as possible.

In this paper, we focused on the case when the first-order formula is fixed. However, the question about the precise complexity of the satisfiability problem when both formulae are parts of instances is also interesting. Is this problem in PSPACE for the case of local satisfiability?

The ultimate aim of this study is to classify *all* (universally) first-order definable modal logics with respect to the decidability of their satisfiability problem. One interesting question in this context is whether the following "metaproblem" is decidable. Given a (universal) first-order formula Φ, is Φ-SAT_L decidable?

References

Vince Bárány, Balder ten Cate, and Luc Segoufin. 2011. Guarded Negation. In *Proceedings of 38th International Colloquium on Automata, Languages and Programming (ICALP'11)*. Springer-Verlag, Berlin, Heidelberg, 356–367.

Robert Berger. 1966. The undecidability of the domino problem. *Mem. AMS* 66 (1966).

Patrick Blackburn, Maarten de Rijke, and Yde Venema. 2001. *Modal Logic*. Cambridge Tracts in Theoretical Comp. Sc., Vol. 53. Cambridge University Press, Cambridge.

Erich Grädel. 1999. On The Restraining Power of Guards. *Journal of Symbolic Logic* 64, 4 (1999), 1719–1742.

Erich Graedel, Martin Otto, and Eric Rosen. 1997. Two-variable logic with counting is decidable. In *Proceedings of the 12th Annual IEEE Symposium on Logic in Computer Science (LICS '97)*. IEEE Computer Society, Washington, DC, USA, 306–317.

Yu. Sh. Gurevich and I. O. Koryakov. 1972. Remarks on Berger's paper on the domino problem. *Siberian Mathematical Journal* 13 (1972), 319–321.

Edith Hemaspaandra. 1996. The Price of Universality. *Notre Dame Journal of Formal Logic* 37, 2 (1996), 174–203.

Edith Hemaspaandra and Henning Schnoor. 2008. On the Complexity of Elementary Modal Logics. In *Proceedings of 25th Annual Symposium on Theoretical Aspects of Computer Science (STACS'08)*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 349–360.

Edith Hemaspaandra and Henning Schnoor. 2011. A Universally Defined Undecidable Unimodal Logic. In *MFCS (LNCS)*, Vol. 6907. Springer, Berlin, Heidelberg, 364–375.

Emanuel Kieroński, Jakub Michaliszyn, and Jan Otop. 2011. Modal Logics Definable by Universal Three-Variable Formulas. In *FSTTCS (LIPIcs)*, Vol. 13. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 264–275.

Emanuel Kieroński, Jakub Michaliszyn, Ian Pratt-Hartmann, and Lidia Tendera. 2012. Two-Variable First-Order Logic with Equivalence Closure. In *Proceedings of the 2012 27th Annual IEEE/ACM Symposium on Logic in Computer Science (LICS '12)*. IEEE Computer Society, Washington, DC, USA, 431–440.

Emanuel Kieroński and Martin Otto. 2005. Small Substructures and Decidability Issues for First-Order Logic with Two Variables. In *Proceedings of the 20th Annual IEEE Symposium on Logic in Computer Science (LICS '05)*. IEEE Computer Society, Washington, DC, USA, 448–457.

Richard E. Ladner. 1977. The Computational Complexity of Provability in Systems of Modal Propositional Logic. *Siam Journal on Computing* 6 (1977), 467–480.

Jakub Michaliszyn and Emanuel Kieroński. 2012. Finite Satisfiability of Modal Logic over Horn-Definable Classes of Frames. In *Advances in Modal Logic*. College Publications, London, UK, 464–482.

Jakub Michaliszyn and Jan Otop. 2012. Decidable Elementary Modal Logics. In *Proceedings of the 27th Annual IEEE/ACM Symposium on Logic in Computer Science (LICS'12)*. IEEE Computer Society, Washington, DC, USA, 491–500.

Michael Mortimer. 1975. On languages with two variables. *Mathematical Logic Quarterly* 21, 1 (1975), 135–140.

Martin Otto. 2001. Two Variable First-Order Logic over Ordered Domains. *Journal of Symbolic Logic* 66, 2 (2001), 685–702.

Leszek Pacholski, Wieslaw Szwast, and Lidia Tendera. 1997. Complexity of Two-Variable Logic with Counting. In *LICS*. IEEE Computer Society, Washington, DC, USA, 318–327.

Ian Pratt-Hartmann. 2005. Complexity of the Two-Variable Fragment with Counting Quantifiers. *Journal of Logic, Language and Information* 14, 3 (2005), 369–395.

Henrik Sahlqvist. 1975. Completeness and correspondence in the first and second order semantics for modal logic. *Studies in Logic and the Foundations of Mathematics* 82 (1975), 110–143.

Walter J. Savitch. 1970. Relationships between nondeterministic and deterministic tape complexities. *J. Comput. Syst. Sci.* 4 (April 1970), 177–192. Issue 2.

Wieslaw Szwast and Lidia Tendera. 2013. FO^2 with one transitive relation is decidable. In *Proceedings of 30th International Symposium on Theoretical Aspects of Computer Science (STACS'13)*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 317–328.

Peter van Emde Boas. 1997. The Convenience of Tilings. In *Complexity, Logic, and Recursion Theory*. Marcel Dekker Inc, New York, USA, 331–363.