Symposium on Theoretical Aspects of Computer Science 2008 (Bordeaux), pp. 349-360
www.stacs-conf.org

ON THE COMPLEXITY OF ELEMENTARY MODAL LOGICS

EDITH HEMASPAANDRA AND HENNING SCHNOOR

Department of Computer Science, Rochester Institute of Technology, Rochester, NY 14623, U.S.A.
E-mail address, Edith Hemaspaandra: eh@cs.rit.edu
E-mail address, Henning Schnoor: hs@cs.rit.edu

Abstract. Modal logics are widely used in computer science. The complexity of modal satisfiability problems has been investigated since the 1970s, usually proving results on a case-by-case basis. We prove a very general classification for a wide class of relevant logics: Many important subclasses of modal logics can be obtained by restricting the allowed models with first-order Horn formulas. We show that the satisfiability problem for each of these logics is either NP-complete or PSPACE-hard, and exhibit a simple classification criterion. Further, we prove matching PSPACE upper bounds for many of the PSPACE-hard logics.

1. Introduction

Modal logics have proven to be a valuable tool in mathematics and computer science. The traditional uni-modal logic enriches the propositional language with the operator $\Diamond$, where $\Diamond\varphi$ is interpreted as φ *possibly holds*. The usual semantics interpret modal formulas over graphs, where $\Diamond\varphi$ means "there is a successor world where φ is true." In addition to their mathematical interest, modal logics are widely used in practical applications: In artificial intelligence, modal logic is used to model the knowledge and beliefs of an agent, see e.g. [BZ05]. Modal logics also can be applied in cryptographic and other protocols [FHJ02, CDF03, HMT88, LR86]. For many specific applications, there exist tailor-made variants of modal logics [BG04].

Due to the vast number of applications, complexity issues for modal logics are very relevant, and have been examined since Ladner's seminal work [Lad77]. Depending on the application, modal logics with different properties are studied. For example, one might want the formula $\varphi \rightarrow \Diamond\varphi$ to be an axiom—if something is true, then it should be considered possible. Or $\Diamond\Diamond\varphi \rightarrow \Diamond\varphi$—if it is possible that φ is possible, then φ itself should be possible. Classical results [Sah73] show that there is a close correspondence between modal logics defined by axioms and logics obtained by restricting the class of considered graphs. Requiring the axioms mentioned above corresponds to restricting the classes of graphs to those which are reflexive or transitive, respectively. Determining the complexity of a given

Supported in part by NSF grants CCR-0311021 and IIS-0713061, a Friedrich Wilhelm Bessel Research Award, and the DAAD postdoc program.

STACS 2008
Symposium on Theoretical Aspects of Computer Science 2008 (Bordeaux), pp. 349-360
http://drops.dagstuhl.de/opus/volltexte/2008/1356

modal logic, defined either by the class of considered graphs or via a modal axiom system, has been an active line of research since Ladner's results. In particular, the complexity classes NP and PSPACE have been at the center of attention.

Most complexity results have been on a case-by-case basis, proving results for individual logics both for standard modal logics and variations like temporal or hybrid logics [HM92, Ngu05, SC85]. Examples of more general results include Halpern and Rêgo's proof that logics including the *negative introspection* axiom, which corresponds to the Euclidean graph property, have an NP-complete satisfiability problem [HR07]. In [SP06], Schröder and Pattinson show a way to prove PSPACE upper bounds for modal logics defined by modal axioms of modal depth 1. In [Lad77], Ladner proved PSPACE-hardness for all logics for which reflexive and transitive graphs are admissible models. In [Spa93], Hemaspaandra showed that all normal logics extending S4.3 have an NP-complete satisfiability problem, and work on the *Guarded Fragment* has shown that some classes of modal logics can be seen as a decidable fragment of first-order logic [AvBN98].

While these results give hardness or upper bounds for classes of logics, they do not provide a full case distinction identifying *all* "easy" or "hard" cases in the considered class. We achieve such a result: For a large class of modal logics containing many important representatives, we identify *all* cases which have an NP-complete satisfiability problem, and show that the satisfiability problem for *all* other non-trivial logics in that class is PSPACE-hard. Hence these problems avoid the infinitely many complexity classes between NP and PSPACE, many of which have natural complete problems arising from logical questions. To our knowledge, such a general result has not been achieved before.

To describe the considered class of modal logics, note that many relevant properties of modal models can be expressed by first-order formulas: A graph is transitive if its edge-relation R satisfies the clause $\forall xyz\,(xRy \wedge yRz \rightarrow xRz)$ and symmetric if it satisfies $\forall xy\,(xRy \rightarrow yRx)$. Many other graph properties can be defined using similar formulas, where the presence of a certain pattern of edges in the graph forces the existence of another. Analogously to propositional logic, we call conjunctions of such clauses *universal Horn formulas*. Many relevant logics can be defined in this way: All examples form [Lad77] fall into this category, as well as logics over Euclidean graphs.

We study the following problem: Given a universal Horn formula $\hat{\psi}$, what is the complexity of the modal satisfiability problem over the class of graphs defined by $\hat{\psi}$?

The main results of this paper are the following: First, we identify all cases which give a satisfiability problem solvable in NP (which then for every nontrivial logic is NP-complete), and show that all other cases are PSPACE-hard. Second, we prove a generalization of a "tree-like model property," and use it to obtain PSPACE upper bounds for a large class of logics. As a corollary, we prove that Ladner's classic hardness result is "optimal" in the class of logics defined by universal Horn formulas. A further corollary is that in the universal Horn class, all logics whose satisfiability problem is not PSPACE-hard already have the "polynomial-size model property," which is only one of several known ways to prove NP upper bounds for modal logics.

Various work was done on restricting the syntax of the modal formulas by restricting the propositional operators [BHSS06], the nesting degree and number of variables [Hal95] or considering modal formulas in Horn form [CL94]. While these results are about restricting the *syntax* of the modal formulas, the current work studies different *semantics* of modal logics, where the semantics are specified by Horn formulas.

logic name	**graph property**	**formula definition**
K	All graphs	$\mathsf{K}(\hat{\varphi}_{\text{taut}})$
T	reflexive	$\mathsf{K}(\hat{\varphi}_{\text{refl}})$
B	symmetric	$\mathsf{K}(\hat{\varphi}_{\text{symm}})$
K4	transitive graphs	$\mathsf{K}(\hat{\varphi}_{\text{trans}})$
S4	transitive and reflexive	$\mathsf{K}(\hat{\varphi}_{\text{trans}} \wedge \hat{\varphi}_{\text{refl}})$
S5	equivalence relations	$\mathsf{K}(\hat{\varphi}_{\text{trans}} \wedge \hat{\varphi}_{\text{refl}} \wedge \hat{\varphi}_{\text{symm}})$

Table 1: Common modal logics

In Section 2, we introduce terminology and generalize classic complexity results. Section 3 introduces universal Horn formulas and states our main result. In Section 4, we explain the main ideas of the proof, before we obtain matching PSPACE upper bounds for many of the involved logics in Section 5, generalizing many previously known results. Proofs and precise definitions can be found in the full version of the paper.

2. Preliminaries

2.1. Basic Concepts and Notation

Modal logic is an extension of propositional logic. A modal formula is a propositional formula using variables, the usual logical symbols $\wedge, \vee, \neg$, and a unary operator $\Diamond$. (A dual operator $\Box$ is often considered as well, this can be regarded as abbreviation for $\neg\Diamond\neg$.) A model for a modal formula is a set of connected "worlds" with individual propositional assignments. To be precise, a *frame* is a directed graph $G = (W, R)$, where the vertices in W are called "worlds," and an edge $(u, v) \in R$ is interpreted as v is "considered possible" from u. A *model* $M = (G, X, \pi)$ consists of a frame $G = (W, R)$, a set X of propositional variables and a function π assigning to each variable $x \in X$ a subset of W, the set of worlds in which x is true. We say the model M is *based* on the frame (W, R). If $\mathcal{F}$ is a class of frames, then a model is an $\mathcal{F}$-model if it is based on a frame in $\mathcal{F}$. With $|M|$ we denote the number of worlds in M. For a world $w \in W$, we define when a modal formula ϕ is *satisfied* at w in M (written $M, w \models \phi$). If ϕ is a variable x, then $M, w \models \phi$ if and only if $w \in \pi(x)$. As usual, $M, w \models \phi_1 \wedge \phi_2$ if and only if $M, w \models \phi_1$ and $M, w \models \phi_2$, and $M, w \models \neg\phi$ iff $M, w \not\models \phi$. For the modal operator, $M, w \models \Diamond\phi$ if and only if there is a world $w' \in W$ such that $(w, w') \in R$ and $M, w' \models \phi$.

We describe a way to define classes of frames with propositional formulas. The *frame language* is the first-order language containing (in addition to the operators $\wedge, \vee$, and $\neg$) the binary relation R, which is interpreted as the edge relation in a graph. Semantics are defined in the obvious way, e.g., a graph satisfies the formula $\hat{\varphi}_{\text{trans}} := \forall xyz(xRy \wedge yRz \rightarrow xRz)$ if and only if it is transitive. In order to separate modal formulas from first-order formulas, we use $\hat{}$ to denote the latter, e.g., $\hat{\varphi}$ is a first-order formula, while ϕ is a modal formula.

A modal logic usually is defined as the set of the formulas provable in it. Since a formula is satisfiable iff its negation is not provable, we can define a logic by the set of formulas satisfiable in it. For a first-order formula $\hat{\psi}$ over the frame language, we define the logic $\mathsf{K}(\hat{\psi})$ as the logic in which a modal formula ϕ is satisfiable if and only if there is a model M and a world $w \in M$ such that the frame which M is based on satisfies $\hat{\psi}$ (we simply write $M \models \hat{\psi}$ for this), and $M, w \models \phi$. Many relevant modal logics can be expressed in this way: In addition to the formula $\hat{\varphi}_{\text{trans}}$ defined above, let $\hat{\varphi}_{\text{refl}} := \forall w(wRw)$, and let $\hat{\varphi}_{\text{symm}} := \forall xy(xRy \rightarrow yRx)$. Finally, let $\hat{\varphi}_{\text{taut}}$ be some tautology over the frame language. Table 1 introduces common modal logics and how they can be expressed in our framework. For a formula $\hat{\psi}$ over the frame language, we consider the following problem:

Problem: $\mathsf{K}(\hat{\psi})$-SAT
Input: A modal formula ϕ
Question: Is ϕ satisfiable in a model based on a frame satisfying $\hat{\psi}$?

As an example, the problem $\mathsf{K}(\hat{\varphi}_{\text{trans}})$-$\mathsf{SAT}$ is the problem of deciding if a given modal formula can be satisfied in a transitive frame, and is the same as the satisfiability problem for the logic $\mathsf{K4}$. In the problem $\mathsf{K}(\hat{\psi})$-SAT, the formula $\hat{\psi}$ is fixed. It is also interesting to study the *uniform* version of the problem, where we are given a formula $\hat{\psi}$ over the frame language and a modal formula ϕ, and the goal is to determine whether there exists a model satisfying both. This problem obviously is PSPACE-hard (this easily follows from Theorem 2.2). In this paper, we study the complexity behavior of *fixed* modal logics, and focus on the NP-PSPACE-gap in complexity. The property of having "small models" is often crucial, as these often lead to a satisfiability problem in NP. A modal logic KL has the *polynomial-size model property*, if there is a polynomial p, such that for every KL-satisfiable formula ϕ, there is a KL-model M and a world $w \in M$ such that $M, w \models \phi$, and $|M| \leq p(|\phi|)$. Since modal logic is an extension of propositional logic, the satisfiability problem for every non-trivial modal logic is NP-hard. Hence, showing that a modal logic has a satisfiability problem in NP is an optimal complexity bound. The following standard observation is the basis of our NP containment proofs:

Proposition 2.1. *Let $\hat{\psi}$ be a first-order formula over the frame language, such that $\mathsf{K}(\hat{\psi})$ has the polynomial-size model property. Then $\mathsf{K}(\hat{\psi})$-$\mathsf{SAT} \in$ NP.*

2.2. Ladner's Theorem and Applications

In [Lad77], Ladner proved complexity results for a variety of modal logics. An *extension* of a modal logic KL is a modal logic KL' such that every formula which is valid in KL is also valid in KL', or equivalently such that every KL'-satisfiable formula is KL-satisfiable. For example, every logic that we consider is an extension of K, and $\mathsf{S4}$ is an extension of $\mathsf{K4}$. It is easy to see that if $\hat{\psi}_1$ and $\hat{\psi}_2$ are formulas over the frame language such that $\hat{\psi}_1$ implies $\hat{\psi}_2$, then $\mathsf{K}(\hat{\psi}_1)$ is an extension of $\mathsf{K}(\hat{\psi}_2)$. Ladner's main result can be stated as follows.

Theorem 2.2 ([Lad77]). *(1) The satisfiability problems for* $\mathsf{K}, \mathsf{K4}$, *and* $\mathsf{S4}$ *are* PSPACE-*complete, and* $\mathsf{S5}$-SAT *is* NP-*complete.*
(2) If $\mathsf{S4}$ *is an extension of* KL, *and* KL *extends* K, *then* KL-SAT *is* PSPACE-*hard.*

The ideas from Ladner's proof for Theorem 2.2 can be extended to show similar results. We define a generalization of transitivity. For a number k, we say that a graph G is k-transitive if for every pair of vertices u, v in G such that there is a k-step path from u to v in G, there is an edge (u, v). Note that a graph is transitive if and only if it is 2-transitive. For a set $S \subseteq \mathbb{N}$, a graph is S-transitive if it is k-transitive for every $k \in S$. A *strict tree* is a directed connected graph which has a root w from which all other vertices can be reached via a unique path. A reflexive/S-transitive/symmetric tree is the reflexive/S-transitive/symmetric closure of a strict tree.

Theorem 2.3. *Let $\hat{\psi}$ be a first-order formula over the frame language such that one of the following cases applies:*

- *$\hat{\psi}$ is satisfied in every strict tree,*
- *$\hat{\psi}$ is satisfied in every reflexive tree,*
- *there is a set $S \subseteq \mathbb{N}$ such that $\hat{\psi}$ is satisfied in every S-transitive tree,*
- *$\hat{\psi}$ is satisfied in every symmetric tree,*
- *$\hat{\psi}$ is satisfied in every tree which is both reflexive and symmetric,*
- *there is a set $S \subseteq \mathbb{N}$ such that $\hat{\psi}$ is satisfied in every tree which is both reflexive and S-transitive.*

Then $\mathsf{K}(\hat{\psi})$-SAT *is* PSPACE-*hard and* $\mathsf{K}(\hat{\psi})$ *does not have the polynomial-size model property.*

This theorem follows with straightforward observations from Ladner's proof, by applying the well-known tree-model property for the involved logics, and using similar ideas for the case of symmetric models. In the next section, we will present the main result of this paper: In the class of modal logics defined by universal Horn formulas, all non-trivial cases that are not covered by Theorem 2.3 have an NP-complete satisfiability problem.

3. Universal Horn Formulas and the Main Result

We now consider a syntactically restricted case of universal first-order formulas, namely Horn formulas. Many well-known modal logics can be expressed in this way. Usually, a Horn clause is defined as a disjunction of literals of which at most one is positive. If a positive literal occurs, then the clause $\overline{x_1} \vee \cdots \vee \overline{x_n} \vee y$ can be written as the implication $x_1 \wedge \cdots \wedge x_n \rightarrow y$. Since in the context of the frame language, an atomic proposition is of the form xRy, the following is the natural version of Horn clauses for our purposes:

Definition 3.1. A *universal Horn clause* over the basic frame language is a formula of the form $x_1^1 R x_2^1 \wedge \cdots \wedge x_1^k R x_2^k \rightarrow C$, where C is of the form xRy or $C = \mathsf{false}$, where all (not necessarily distinct) variables are implicitly universally quantified.

A *universal Horn formula* is a conjunction of universal Horn clauses. Due to space reasons, in this paper we only consider Horn clauses of the first form (known as positive Horn clauses), our results hold for the second form analogously. With universal Horn formulas, many of the frame properties usually considered can be expressed, like transitivity, symmetry, euclidicity, etc. We now state the classification theorem:

Theorem 3.2. *Let $\hat{\psi}$ be a universal Horn formula. Then either $\hat{\psi}$ satisfies the condition from Theorem 2.3, (and thus* $\mathsf{K}(\hat{\psi})$-SAT *is* PSPACE-*hard) or* $\mathsf{K}(\hat{\psi})$ *has the polynomial-size model property and* $\mathsf{K}(\hat{\psi})$-$\mathsf{SAT} \in$ NP.

4. A Proof Sketch

We now give an overview of the ideas used to prove our main result. We first show that universal Horn clauses can be represented as graphs in such a way that the properties of the involved logics can be characterized with homomorphic images of the defining Horn clauses. In Section 4.2, we then demonstrate at an example how NP results can be shown for universal Horn clauses, before outlining the strategy for the proof of Theorem 3.2.

4.1. Universal Horn clauses and homomorphisms

For a universal Horn clause $\hat{\varphi} = x_1^1 R x_2^1 \wedge \cdots \wedge x_1^k R x_2^k \rightarrow x R y$, the *prerequisite graph of* $\hat{\varphi}$, denoted with $\text{prereq}(\hat{\varphi})$, consists of the variables appearing on the left-hand side of the implication $\hat{\varphi}$ as vertices, where for variables $u, v \in \{x_1^1, \ldots, x_2^k\}$, there is an edge (u, v) if the clause uRy appears on the left-hand side of the formula. The *conclusion edge of* $\hat{\varphi}$, denoted with $\text{conc}(\hat{\varphi})$, is the edge (x, y). As usual, a *homomorphism* between graphs is a function on the vertices preserving the edge relation. The definition of the prerequisite graph and the conclusion edge of a universal Horn formula establishes a one-to-one correspondence between universal Horn clauses and their representation as graphs. These definitions allow us to relate truth of a Horn clause to homomorphic images of the involved graphs. In the following, $\text{prereq}(\hat{\varphi}) \cup \{x, y\}$ is the graph obtained from $\text{prereq}(\hat{\varphi})$ by adding (if not already present) the vertices x and y, but no additional edges.

Proposition 4.1. *Let $\hat{\varphi}$ be a universal Horn clause with $\text{conc}(\hat{\varphi}) = (x, y)$. A graph G satisfies $\hat{\varphi}$ if and only if for every homomorphism $\alpha\colon \text{prereq}(\hat{\varphi}) \cup \{x, y\} \rightarrow G$, there is an edge $(\alpha(x), \alpha(y))$ in G.*

This observation is central, as it shows that properties of a logic $\mathsf{K}(\hat{\varphi})$ depend on the types of homomorphic images of $\text{prereq}(\hat{\varphi})$.

4.2. Example of a Case in NP

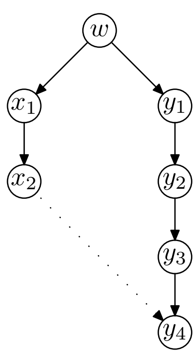

Figure 1: $\hat{\varphi}^{2\rightarrow 4}$

We now give an example of the proof of NP membership. Let $\hat{\varphi}^{k\rightarrow l}$ be the formula $wRx_1 \wedge x_1 R x_2 \wedge \cdots \wedge x_{k-1} R x_k \wedge wRy_1 \wedge y_1 R y_2 \wedge \cdots \wedge y_{l-1} R y_l \rightarrow x_k R y_l$, where all variables are universally quantified (and in the case that $k = 0$ or $l = 0$, we replace x_0 or y_0 with w). In Figure 1, we present the graph representation of the formula $\hat{\varphi}^{2\rightarrow 4}$. A graph G satisfies $\hat{\varphi}^{k\rightarrow l}$ if and only if for any nodes $w, x_k, y_l \in G$, if there is a k-step path from w to x_k and an l-step path from w to y_l, then (x_k, y_l) is an edge in G. This definition generalizes several well-known examples: A graph is reflexive if and only if it satisfies $\hat{\varphi}^{0\rightarrow 0}$, symmetry is expressed with $\hat{\varphi}^{1\rightarrow 0}$, and k-transitivity with $\hat{\varphi}^{0\rightarrow k}$. A graph is Euclidean iff it satisfies $\hat{\varphi}^{1\rightarrow 1}$. Thus, this notation allows us to capture many interesting graph properties. We give the proof idea for a relatively easy special case. Generalizations of this idea are the main ingredients for our polynomial-size model proofs.

Theorem 4.2. *Let $k \geq 1$, and let $\hat{\psi}$ be a universal formula over the frame language such that $\hat{\psi}$ implies $\hat{\varphi}^{k\rightarrow k}$. Then $\mathsf{K}(\hat{\psi})$ has the polynomial-size model property, and $\mathsf{K}(\hat{\psi})$-$\mathsf{SAT} \in \text{NP}$.*

Proof Sketch. It suffices to prove the polynomial-size model property. We need to show that every $\mathsf{K}(\hat{\psi})$-satisfiable formula ϕ has a "small" model. Let M be a model and w a world such that $M, w \models \phi$, and M also satisfies the first-order formula $\hat{\psi}$, in particular it then satisfies $\hat{\varphi}^{k\rightarrow k}$, which means that for t, u, v in M, if there is a k-step path from t to u and also from t to v, then (u, v) is an edge in M. Note that since $\hat{\psi}$ is universal, every restriction of M still satisfies $\hat{\psi}$. With a standard construction, we can restrict the number of vertices in M which do not have a k-step predecessor to polynomial size.

We show that the edge relation restricted to the set C of nodes having k-step predecessors is an equivalence relation, then with a standard argument (similar to the proof for the logic $\mathsf{S5}$), we can pick one world for every satisfied subformula of ϕ from polynomially many "clusters" in C, and obtain a small submodel M' of M satisfying $M', w \models \phi$.

We first show that every node v in C is reflexive: It has a k-step predecessor t, and hence by the $\hat{\varphi}^{k \to k}$-property, there is an edge (v, v). We show that C is symmetric: Let (u, v) be an edge in C. From the above, we know that both of these nodes are reflexive in C. Hence there is a k-step path from u to both u and v. From the $\hat{\varphi}^{k \to k}$-property, it follows that (v, u) is an edge. Finally, we prove that C is transitive. Let (t, u) and (u, v) be edges in C. Since we already showed that C is symmetric, we know that (u, t) is an edge as well. Since all of the involved nodes are also reflexive, there is a k-step path from u to both t and v, implying that (t, v) is an edge as well. □

The analog of Theorem 4.2 can be shown to hold for many other cases, for example showing that all logics of the form $\mathsf{K}(\hat{\varphi}^{k \to l})$ for $1 \leq k, l$ or $k \geq 2, l = 0$ have satisfiability problems in NP. Note that the case $k = 1$ of Theorem 4.2 follows from the main result of [HR07]. Our results and theirs are incomparable: They achieve the NP result for *all* modal logics extending what in our notation is $\mathsf{K}(\hat{\varphi}^{1 \to 1})$, where our results only hold for logics defined by universal formulas. On the other hand, Theorem 3.2 gives NP membership for a large class of logics which do not follow from their result, namely all logics defined by universal Horn formulas not satisfied on the various forms of trees mentioned in Theorem 2.3.

Similarly to Proposition 4.1, it can be shown that there is a close relationship between implications of Horn clauses and homomorphisms between their prerequisite graphs. This relationship and Theorem 4.2 imply that every universal Horn clause which can be homomorphically mapped into the clause $\hat{\varphi}^{k \to k}$ in an appropriate way defines a logic having the polynomial-size model property and a satisfiability problem in NP. Therefore we can collect "homomorphism properties" of Horn clauses that guarantee this property of the corresponding logics (we need to be careful with what a homomorphism exactly is in this context—we will not go into the technical details here). We can define a class $\mathcal{H}_{\mathrm{NP}}$ of graphs, containing among others representations of generalizations of $\hat{\varphi}^{k \to l}$ for those k, l leading to the polynomial-size model property, which has the following properties:

(1) For all universal Horn clauses $\hat{\varphi}$ which can be homomorphically mapped into an element of $\mathcal{H}_{\mathrm{NP}}$, $\mathsf{K}(\hat{\varphi})$ has the polynomial-size model property, and $\mathsf{K}(\hat{\varphi})$-$\mathsf{SAT} \in \mathrm{NP}$.
(2) $\mathcal{H}_{\mathrm{NP}}$ is large enough to form the basic building blocks of almost all of our NP results.

An example for a graph in $\mathcal{H}_{\mathrm{NP}}$ is the following: If a clause $\hat{\varphi}$ with $\mathrm{conc}\,(\hat{\varphi}) = (x, y)$ can be mapped into a "generalized line" $(l_0, \ldots, l_n)$ with a homomorphism α satisfying $\alpha(y) = l_i$ and $\alpha(x) = l_{i+k}$ for $k \geq 2$, then the NP conditions are met. This is satisfied, for example, for clauses of the form $\hat{\varphi}^{k \to 0}$ for $k \geq 2$. While the set $\mathcal{H}_{\mathrm{NP}}$ contains a large class of graphs, on its own it is not sufficient to prove all NP results for logics defined by universal Horn formulas. The main reason is that it does not take into account interference between clauses in Horn formulas. For example, it is well known that reflexivity, symmetry, and transitivity alone lead to PSPACE-complete logics. But the combination of these requirements defines the logic $\mathsf{S5}$, which has a satisfiability problem in NP. Hence, interference between clauses in a Horn formulas is of crucial importance for the complexity of the defined logic.

4.3. An Alternative Formulation and Proof Idea of the Main Theorem

We give an idea of the main strategy used to prove the classification Theorem 3.2. For this, we restate the complexity classification as the algorithm Horn-Classification (see Figure 3) which, given a universal Horn formula $\hat{\psi}$ as input, determines the complexity of $\mathsf{K}(\hat{\psi})$-SAT. The output of the algorithm and the statement of Theorem 3.2 agree in all cases. The purpose of the algorithm is not to be implemented, but to serve as a case distinction used to prove Theorem 3.2, and several later corollaries.

The idea of the algorithm is the following: In *types-list*, it maintains a list of implications of the formula $\hat{\psi}$. For example, Horn-Classification puts `refl` into *types-list* if it detects the formula $\hat{\psi}$ to require any graph G satisfying it to be "near reflexive" (every vertex having long enough paths ending in and originating in it must be reflexive). Similarly, $\texttt{symm} \in$ *types-list* ($\texttt{trans}^k \in$ *types-list*) means that $\hat{\psi}$ requires a graph to be "near symmetric" ("near k-transitive"). The occurring class $\mathcal{H}_{\mathrm{NP}}^{\textit{types-list}}$ is a generalization of the set $\mathcal{H}_{\mathrm{NP}}$ introduced earlier. Since we are not considering only individual clauses anymore, this class is not constant, but dynamically grows corresponding to collected implications of the formula kept in *types-list*. For example, $\mathcal{H}_{\mathrm{NP}}^{\{\texttt{refl}\}}$ contains reflexive closures of elements in $\mathcal{H}_{\mathrm{NP}}$.

If Horn-Classification detects that a clause $\hat{\varphi}$ can be mapped into an element of $\mathcal{H}_{\mathrm{NP}}^{\textit{types-list}}$, then the clause $\hat{\varphi}$, in addition with the requirements kept in *types-list*, implies the polynomial-size model property. Another condition leading to NP containment of the logic is the following: It is well known that the modal logic over the class of frames which are both transitive and symmetric has a satisfiability problem in NP. Generalizing this, when Horn-Classification detects that $\hat{\psi}$ implies "near symmetry" and "near k-transitivity," this also leads to NP solvability of the satisfiability problem.

For a set *types-list* $\subseteq \left\{\texttt{refl}, \texttt{symm}, \texttt{trans}^k \mid k \in \mathbb{N}\right\}$, a graph G satisfies the conditions of *types-list* if it has the corresponding properties, i.e., if $\texttt{refl} \in$ *types-list* ($\texttt{symm} \in$ *types-list*, $\texttt{trans}^k \in$ *types-list*), then G is required to be reflexive (symmetric, k-transitive). A *types-list* tree is obtained from a strict tree by adding exactly those edges required to make it satisfy the conditions of *types-list* (this is a natural closure operator). For a universal Horn clause $\hat{\varphi}$, let *types-list*-$T_{\hat{\varphi}}^{\mathrm{hom}}$ denote the pairs (α, T) such that T is a *types-list* tree, and $\alpha\colon \mathrm{prereq}\,(\hat{\varphi}) \rightarrow T$ is a homomorphism. Due to Proposition 4.1, this is the set of *types-list* trees about which the clause $\hat{\varphi}$ "makes a statement," along with the corresponding homomorphisms.

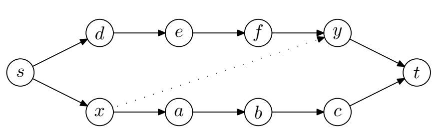

Figure 2: Example Formula

We now define the properties of Horn clauses which do not lead to NP containment of the satisfiability problems on their own. Recalling Section 2.2, it is natural that clauses which are satisfied in every reflexive, transitive, or symmetric tree are among these. This is captured by the following definitions: If $\hat{\varphi}$ is a universal Horn clause with $\mathrm{conc}\,(\hat{\varphi}) = (x, y)$, we say that ($\hat{\varphi}$, *types-list*) satisfies the *reflexive case*, if for every $(\alpha, T) \in$ *types-list*-$T_{\hat{\varphi}}^{\mathrm{hom}}$ it holds that $\alpha(x) = \alpha(y)$. ($\hat{\varphi}$, *types-list*) satisfies the *transitive case for $k \in \mathbb{N}$*, if for every $(\alpha, T) \in$ *types-list*-$T_{\hat{\varphi}}^{\mathrm{hom}}$ there is a path from $\alpha(x)$ to $\alpha(y)$ in T, and there is some $(\alpha, T) \in$ *types-list*-$T_{\hat{\varphi}}^{\mathrm{hom}}$ such that $\alpha(y)$ is exactly k levels below $\alpha(x)$ in T. Finally, ($\hat{\varphi}$, *types-list*) satisfies the *symmetric case* if for every $(\alpha, T) \in$ *types-list*-$T_{\hat{\varphi}}^{\mathrm{hom}}$, there

Input: Universal Horn formula $\hat{\psi}$
  types-list $:= \emptyset$
  while not done **do**
    if every clause in $\hat{\psi}$ is satisfied on every *types-list* tree **then**
      $\mathsf{K}(\hat{\psi})$-SAT is PSPACE-hard
    end if
    Let $\hat{\varphi}$ be a clause in $\hat{\psi}$ not satisfied on every *types-list* tree
    if $\hat{\varphi}$ can homomorphically be mapped into a graph from $\mathcal{H}_{\mathrm{NP}}^{\textit{types-list}}$ **then**
      $\mathsf{K}(\hat{\psi})$ has the polynomial-size model property, and $\mathsf{K}(\hat{\psi})$-$\mathsf{SAT} \in \mathrm{NP}$.
    else
      if ($\hat{\varphi}$, *types-list*) satisfies the reflexive case **then**
        types-list $:=$ *types-list* $\cup \{\texttt{refl}\}$
      else if ($\hat{\varphi}$, *types-list*) satisfies the transitive case for $k \geq 2$ **then**
        types-list $:=$ *types-list* $\cup \{\texttt{trans}^k\}$
      else if ($\hat{\varphi}$, *types-list*) satisfies the symmetric case **then**
        types-list $:=$ *types-list* $\cup \{\texttt{symm}\}$
      end if
    end if
    if for some k, $\{\texttt{symm}, \texttt{trans}^k\} \subseteq$ *types-list* **then**
      $\mathsf{K}(\hat{\psi})$ has the polynomial-size model property, and $\mathsf{K}(\hat{\psi})$-$\mathsf{SAT} \in \mathrm{NP}$.
    end if
  end while

Figure 3: The Algorithm Horn-Classification

is an edge $(\alpha(y), \alpha(x))$ in T. If one of these cases applies, $\hat{\varphi}$ is satisfied in every reflexive, transitive, or symmetric tree, and hence recalling Theorem 2.3, it is not surprising that these conditions do not lead to NP membership on their own—but in combination with others, they very well might. In the variable *types-list*, the algorithm keeps a list of these conditions encountered.

For the correctness proof, we first need to show that the choices that the algorithm has to make always can be made: In the relevant situations, at least one of the "reflexive," "transitive," or "symmetric" conditions occurs. We also need to prove that it actually comes to a halt—the main argument is to show that only finitely many $\texttt{trans}^k$-conditions are added to *types-list*, and no element is added twice. Building on Ladner's results, proving correctness of the PSPACE hard cases is trivial. The interesting statement of the classification is that all logics not covered by these cases have an NP-complete satisfiability problem, which is as low a complexity bound as we can hope for. We give an example for a logic which Horn-Classification determines to have a satisfiability problem in NP.

As an example, let $\hat{\varphi}$ be the universal Horn clause with prerequisite graph as shown in Figure 2, with $\mathrm{conc}\,(\hat{\varphi}) = (x, y)$, and let $\hat{\psi}$ be the Horn formula having $\hat{\varphi}$ as its only clause. Horn-Classification starts with *types-list* $= \emptyset$, and in its first iteration checks if $\hat{\varphi}$ is satisfied in every strict tree. This is not the case, as Figure 4 shows (here, we simply marked each node with the names of the vertices which are preimages of the homomorphism): This is a homomorphic image of $\mathrm{prereq}\,(\hat{\varphi})$ as a line (in particular a strict tree), in which the images of x and y are not connected with an edge. Thus $\hat{\varphi}$ is not satisfied in this strict tree. However, if we map prereq $(\hat{\varphi})$ homomorphically into a strict tree, then the vertices between s and t must be "pairwise identified" like in Figure 4.

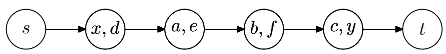

Figure 4: Image as line

Therefore the transitive case is satisfied for $k = 3$, and HORN-CLASSIFICATION adds $\mathtt{trans}^3$ to *types-list*. Next it checks if $\hat{\varphi}$ is satisfied in every $\{\mathtt{trans}^3\}$ tree. This is not the case, and the homomorphic image of prereq $(\hat{\varphi})$ as a $\{\mathtt{trans}^3\}$ line in Figure 5 (here we only included those edges added by the $\mathtt{trans}^3$ closure required for the homomorphism) shows that $\hat{\varphi}$ satisfies the homomorphic property in $\mathcal{H}_{\mathrm{NP}}$ mentioned at the end of Section 4.2. This also clarifies what a "generalized line" is: At this point, the variable *types-list* only contains the condition $\mathtt{trans}^3$, and hence a "generalized line" is the 3 transitive closure of a line. Therefore, $\mathsf{K}(\hat{\varphi})$-$\mathsf{SAT} \in \mathrm{NP}$.

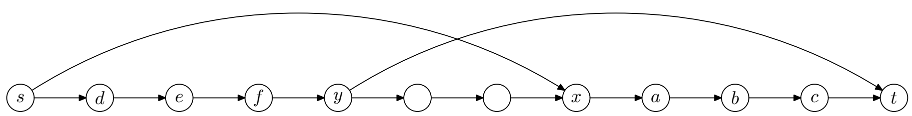

Figure 5: Homomorphic image as $\{\mathtt{trans}^3\}$-line

This example demonstrates that in the run of the algorithm, a clause $\hat{\varphi}$ can meet different cases depending on the content of the variable *types-list* : In the situation that *types-list* $= \emptyset$, the clause $\hat{\varphi}$ satisfies the transitive case, but when *types-list* $= \{\mathtt{trans}^3\}$, this is no longer true.

5. Tree-like Models and PSPACE Upper Bounds

We now prove PSPACE upper bounds for many of our logics. The first step is to prove a tree-like model property for the PSPACE-hard cases. This generalizes many standard results, like the fact that $\mathsf{K4}$-satisfiable formulas are always satisfiable on a transitive tree. For a model M, let $\mathsf{edges}\,(M)$ denote the edges of the frame that M is based on.

Theorem 5.1. *Let $\hat{\psi}$ be a universal Horn formula satisfied on every types-list tree, where there is no k such that types-list contains both* $\mathtt{trans}^k$ *and* $\mathtt{symm}$. *Then every $\mathsf{K}(\hat{\psi})$-satisfiable modal formula is satisfiable in a model M such that M satisfies $\hat{\psi}$ and there is a strict tree T_0 and a types-list tree T_1 such that T_0, T_1, and M have the same set of worlds, and* $\mathsf{edges}\,(T_0) \subseteq \mathsf{edges}\,(M) \subseteq \mathsf{edges}\,(T_1)$.

The theorem shows that in a certain way, the "converse" of the prerequisite that $\hat{\psi}$ is satisfied in every *types-list* tree is true as well: Not only is every *types-list* tree a model in the logic $\mathsf{K}(\hat{\psi})$, but we can restrict ourselves to models which are "close" to *types-list* trees. This is not a real "converse:" For example, if *types-list* contains only the symmetric condition, there might be formulas satisfiable in $\mathsf{K}(\hat{\psi})$, but not in a symmetric tree. The main idea of the proof is to start with a tree model for a satisfiable formula ϕ, and step by step add enough edges to the model in a way which again gives a model for ϕ. Using this theorem, we can construct a PSPACE decision algorithm for the involved logics similar in spirit to Ladner's decision procedure, with the additional difficulty that we need to ensure that the properties demanded by the first-order formula are also met by the model. This is a major obstacle: Recalling Proposition 4.1, we need to consider all homomorphisms from the prerequisite graphs of the clauses in the formula $\hat{\psi}$ into the (potentially exponential size) model that we construct, at the same time may only keep a polynomial fragment of the model in memory. This can be done since we can restrict ourselves to connected components of small diameter in a strict tree. The theorem gives a unified proof for showing that the logics K, T, and B, among others, have satisfiability problems in PSPACE.

Theorem 5.2. *Let $\hat{\psi}$ be a universal Horn formula such that* Horn-Classification *does not add any* $\texttt{trans}^k$ *to types-list on input $\hat{\psi}$. Then* $\mathsf{K}(\hat{\psi})$-$\mathsf{SAT} \in$ PSPACE.

There are some interesting applications of our results: Ladner proved that all normal modal logics KL such that $\mathsf{S4}$ is an extension of KL have a PSPACE-hard satisfiability problem. Using Theorems 3.2 and 5.1, we show that his result is optimal in the sense that every universal Horn logic which is a "proper extension" of $\mathsf{S4}$ already has an NP-solvable satisfiability problem.

Theorem 5.3. *Let $\hat{\psi}$ be a universal Horn formula such that $\hat{\psi}$ implies $\hat{\varphi}_{\text{refl}} \wedge \hat{\varphi}_{\text{trans}}$. Then either $\mathsf{K}(\hat{\psi}) = \mathsf{S4}$, or $\mathsf{K}(\hat{\psi})$ has the polynomial-size model property and $\mathsf{K}(\hat{\psi})$-$\mathsf{SAT} \in$* NP.

We further can show a PSPACE upper bound for all universal Horn logics which are extensions of the logic T, and hence, from Theorem 3.2, conclude that these are all either solvable in NP (and thus NP-complete if they are consistent), or PSPACE-complete.

Theorem 5.4. *Let $\hat{\psi}$ be a universal Horn formula such that $\hat{\psi}$ implies $\hat{\varphi}_{\text{refl}}$. Then $\mathsf{K}(\hat{\psi})$-$\mathsf{SAT} \in$* PSPACE.

In a similar way, we can prove that all universal Horn logics which imply a variant of symmetry give rise to a satisfiability problem in PSPACE. A noteworthy difference in the prerequisites of Theorem 5.4 and Corollary 5.5 is that the former requires the reflexivity condition to be implied by the formula $\hat{\psi}$, while the latter only needs a "near symmetry"-condition as detected by Horn-Classification.

Corollary 5.5. *Let $\hat{\psi}$ be a universal Horn formula such that* Horn-Classification *adds* $\texttt{symm}$ *to types-list on input $\hat{\psi}$. Then* $\mathsf{K}(\hat{\psi})$-$\mathsf{SAT} \in$ PSPACE. *In particular, any universal Horn logic which is an extension of* B *has a satisfiability problem solvable in* PSPACE.

6. Conclusion and Future Research

We analyzed the complexity of modal logics defined by universal Horn formulas, covering many well-known logics. We showed that the non-trivial satisfiability problems for these logics are either NP-complete or PSPACE-hard, and gave an easy criterion to recognize these cases. Our results directly imply that (unless NP = PSPACE) such a logic has a satisfiability problem in NP if and only if it has the polynomial-size model property. We also demonstrated that a wide class of the considered logics has a satisfiability problem solvable in PSPACE.

Open questions include determining complexity upper bounds for the satisfiability problems for all modal logics defined by universal Horn formulas. We strongly conjecture that all of these are decidable, and consider it possible that all of these problems are in PSPACE. A successful way to establish upper complexity bounds is the guarded fragment [AvBN98, Grä99]. This does not seem to be applicable to our logics, since it cannot be used for transitive logics, and we obtain PSPACE-upper bounds for all of our logics except those involving a variant of transitivity.

The next major open challenges are generalizing our results to formulas not in the Horn class, and allowing arbitrary quantification. Initial results show that even when considering only universal formulas over the frame language, undecidable logics appear. An interesting enrichment of Horn clauses is to allow the equality relation. Preliminary results indicate that Theorem 3.2 holds for this more general case as well.

Acknowledgments: The second author thanks Thomas Schneider for helpful hints. We also thank the anonymous referees for many helpful suggestions.

References

[AvBN98] H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.

[BG04] B. Bennett and A. Galton. A unifying semantics for time and events. *Artificial Intelligence*, 153(1-2):13–48, 2004.

[BHSS06] M. Bauland, E. Hemaspaandra, H. Schnoor, and I. Schnoor. Generalized modal satisfiability. In *Proc. of the 23nd Symposium on Theoretical Aspects of Computer Science*, pp. 500–511, 2006.

[BZ05] C. Baral and Y. Zhang. Knowledge updates: Semantics and complexity issues. *Artificial Intelligence*, 164(1-2):209–243, 2005.

[CDF03] T. Coffey, R. Dojen, and T. Flanagan. On the automated implementation of modal logics used to verify security protocols. In *Proceedings of the 1st international symposium on Information and communication technologies*, pages 329–334. Trinity College Dublin, 2003.

[CL94] C. Chen and I. Lin. The computational complexity of the satisfiability of modal horn clauses for modal propositional logics. *Theoretical Computer Science*, 129(1):95–121, 1994.

[FHJ02] U. Frendrup, Hüttel, and J. Jensen. Modal logics for cryptographic processes. In *Proceedings of EXPRESS*, 2002.

[Grä99] E. Grädel. Why are modal logics so robustly decidable? *Bulletin of the European Association for Theoretical Computer Science*, 68:90103, 1999.

[Hal95] J. Halpern. The effect of bounding the number of primitive propositions and the depth of nesting on the complexity of modal logic. *Artificial Intelligence*, 75(2):361–372, 1995.

[HM92] J. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54(2):319–379, 1992.

[HMT88] J. Halpern, Y. Moses, and M. Tuttle. A knowledge-based analysis of zero knowledge. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, pages 132–147, New York, NY, USA, 1988. ACM Press.

[HR07] J. Halpern and L. Rêgo. Characterizing the NP-PSPACE gap in the satisfiability problem for modal logic. In *International Joint Conferences on Artificial Intelligence*, pages 2306–2311, 2007.

[Lad77] R. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal on Computing*, 6(3):467–480, 1977.

[LR86] R. Ladner and J. Reif. The logic of distributed protocols: Preliminary report. In *Proceedings of the 1986 Conference on Theoretical Aspects of Reasoning About Knowledge*, pages 207–222, San Francisco, CA, USA, 1986. Morgan Kaufmann Publishers Inc.

[Ngu05] L. Nguyen. On the complexity of fragments of modal logics. In *Advances in Modal Logic - Volume 5*, pages 249–268. King's College Publications, 2005.

[Sah73] H. Sahlqvist. Completeness and correspondence in the first and second order semantics for modal logic. In *Proceedings of the Third Scandinavian Logic Symposium*, 1973.

[SC85] A. Sistla and E. Clarke. The complexity of propositional linear temporal logics. *Journal of the ACM*, 32(3):733–749, 1985.

[SP06] L. Schröder and D. Pattinson. PSPACE bounds for rank-1 modal logics. In *LICS*, pages 231–242, 2006.

[Spa93] E. Spaan. *Complexity of Modal Logics*. PhD thesis, Department of Mathematics and Computer Science, University of Amsterdam, 1993.