

LTCS–Report

Aachen University of Technology
Research group for
Theoretical Computer Science

Modal Logic and the two-variable fragment (Revised Version)

Carsten Lutz, Ulrike Sattler, and Frank Wolter

LTCS-Report 01-04

RWTH Aachen
LuFg Theoretische Informatik
<http://www-lti.informatik.rwth-aachen.de>

Ahornstr. 55
52074 Aachen
Germany

Modal Logic and the two-variable fragment (Revised Version)

Carsten Lutz¹, Ulrike Sattler¹, and Frank Wolter²

¹ LuFG Theoretical Computer Science, RWTH Aachen,
Ahornstr. 55, 52074 Aachen, Germany

² Institut für Informatik, Universität Leipzig,
Augustus-Platz 10-11, 04109 Leipzig, Germany

Abstract. We introduce a modal language L which is obtained from standard modal logic by adding the Boolean operators on accessibility relations, the identity relation, and the converse of relations. It is proved that L has the same expressive power as the two-variable fragment FO^2 of first-order logic, but speaks less succinctly about relational structures: if the number of relations is bounded, then L -satisfiability is EXPTIME-complete but FO^2 satisfiability is NEXPTIME-complete. We indicate that the relation between L and FO^2 provides a general framework for comparing modal and temporal languages with first-order languages.

1 Introduction

Ever since it was observed that many modal logics can be regarded as fragments of first-order logic, exploring the connection between these two families of languages has been a major research issue. The starting point was Kamp's result [18] stating that modal logic with binary operators *Since* and *Until* has the same expressive power as monadic first-order logic over structures such as $\langle \mathbb{N}, < \rangle$ and $\langle \mathbb{R}, < \rangle$. Van Benthem [26, 27] provided a systematic model theoretic analysis of the relation between families of modal logics and predicate logics and Gabbay [10, 9] extended Kamp's result to a systematic investigation of expressively complete modal logics. As part of his investigation, Gabbay made the basic observation that modal languages are often contained in *finite variable fragments* of first-order logic. For example, the basic modal language with unary operators can be embedded in the two-variable fragment FO^2 of first-order logic. In the early 1990s, this observation was regarded as an explanation for the decidability of many modal logics: the decidability of FO^2 (cf. [21, 23, 14]) explains the decidability of standard modal logics simply because they are contained in it.³ The situation is different as soon as our concern is computational complexity: while most standard modal logics are decidable in EXPTIME, in PSPACE,

³ More recently it has been argued that some "modal phenomena" are better explained by their *tree-model-property* [28] (i.e., they are determined by tree-like structures) and/or by embedding them in bounded (or guarded) fragments of first-order logic [1, 13]. The logics we consider here do not have those properties.

or in NP (see e.g. [19, 2, 25]), the two-variable fragment FO^2 is NEXPTIME-complete [14]. Thus, the question arises why modal logics are often of a lower complexity than the two-variable fragment. There are two possible explanations for this phenomenon:

- The “standard modal logics” contained in FO^2 have strictly less expressive power than FO^2 itself;
- Although the expressive power of some standard modal logics coincides with the expressive power of FO^2 , the way in which FO^2 speaks about relational structures is strictly more succinct than the way in which modal languages do.

In temporal logic, it follows from considerations of expressiveness and computational complexity that the second explanation is the correct one, see for example the paper [7] by Etessami et al. The main contribution of this paper is to show that, also on non-linear structures, the second explanation is the correct one: we define a natural modal logic L and prove the following:

1. L has the *same expressive power* as FO^2 ;
2. as soon as we allow only for a bounded number of relation symbols, L -satisfiability is only EXPTIME-complete (whereas it is NEXPTIME-complete for an unbounded number of relation symbols).

The logic L extends the basic multi-modal logic \mathbf{K}_m by (i) Boolean combinations of accessibility relations, (ii) the converse of accessibility relations, and (iii) the identity relation. All those ingredients have been investigated and applied intensively: see [11, 17, 20] for (i), [5, 12, 30] for (ii), and [6] for (iii). Hence, L can certainly be regarded as a standard member of the modal family and Property 1 above rules out the first explanation for the good computational behaviour of modal logic. To further support the second explanation, we use a simple argument to prove that L is exponentially more succinct than FO^2 .

The usefulness of our result that L is expressively complete for FO^2 is demonstrated by showing that it provides a general framework for comparing the expressive power and complexity of modal logic and first-order logic. For example, in “weak” temporal logic (where the only temporal operators are “always in the future” and “always in the past”) interpreted over strict linear orderings, the Boolean operations and the identity relation are definable. Thus, weak temporal logic has the same expressive power as FO^2 over strict linear orderings (and without further binary relation symbols). For the strict linear ordering $\langle \mathbb{N}, < \rangle$, this was first proved by Etessami et al. in [7]. In this case, the complexity-gap is even wider: Over $\langle \mathbb{N}, < \rangle$, weak temporal logic is NP-complete [24] while FO^2 is NEXPTIME-complete [7, 15]. In the present paper, we show that this holds for $\langle \mathbb{Q}, < \rangle$ and $\langle \mathbb{R}, < \rangle$ as well.

2 Expressivity

We start with defining the languages under consideration. FO^2 comprises exactly those first-order formulas without constants and function symbols but with

equality whose only variables are x and y and whose relation symbols have arity ≤ 2 . The unary predicates are denoted by P_1, \dots while the binary ones are R_1, \dots . For $m \leq \omega$ we denote by FO_m^2 the fragment of FO^2 consisting of formulas containing only the first m binary relations. FO^2 is interpreted in the standard manner in structures of the form $\langle W, \mathcal{P}_1, \dots, \mathcal{R}_1, \dots \rangle$ in which the \mathcal{P}_i interpret the P_i and the \mathcal{R}_i interpret the R_i .

The modal language $\mathcal{ML}^{\neg, \cap, \cup, -, id}$ is Boolean modal logic [11, 20] enriched with a converse constructor and the identity relation.

Definition 1. A complex modal parameter is an expression built up from atomic modal parameters R_1, \dots , the identity parameter id , and the operators \neg, \cap, \cup , and \cdot^- . For $m \leq \omega$ we denote by $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ the modal language defined inductively as follows:

- all propositional variables p_1, p_2, \dots belong to $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$;
- if $\varphi, \psi \in \mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ and S is a complex modal parameter built from the first m atomic modal parameters R_1, \dots, R_m and id , then $\neg\psi$, $\varphi \wedge \psi$, and $\langle S \rangle \varphi$ belong to $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$.

We abbreviate $\top = p_1 \vee \neg p_1$ and $\perp = \neg \top$. The box operator $[S]\varphi$ and other Boolean connectives are defined as abbreviations in the standard manner.

A Kripke-model is a structure $\mathcal{M} = \langle W, \pi, \mathcal{R}_1, \dots \rangle$ in which π associates with every variable p a subset $\pi(p)$ of W . Let S be a (possibly complex) modal parameter. Then the extension $\mathcal{E}(S)$ is inductively defined as follows:

- | | |
|------------------------------------|---|
| if $S = R_i$ (i.e., S is atomic) | then $\mathcal{E}(S) = \mathcal{R}_i$ |
| if $S = id$ | then $\mathcal{E}(S) = \{(w, w) \mid w \in W\}$ |
| if $S = \neg S'$ | then $\mathcal{E}(S) = (W \times W) \setminus \mathcal{E}(S')$ |
| if $S = S_1 \cap S_2$ | then $\mathcal{E}(S) = \mathcal{E}(S_1) \cap \mathcal{E}(S_2)$ |
| if $S = S_1^-$ | then $\mathcal{E}(S) = \{(w, w') \mid (w', w) \in \mathcal{E}(S_1)\}$ |

The semantics of formulas is defined inductively in the standard way, e.g. for the diamond operator we have

$$\mathcal{M}, w \models \langle S \rangle \varphi \text{ iff } \exists w' \in W \text{ with } (w, w') \in \mathcal{E}(S) \text{ and } \mathcal{M}, w' \models \varphi$$

Given a Kripke-model $\mathcal{M} = \langle W, \pi, \mathcal{R}_1, \dots \rangle$, define a corresponding first-order model $\mathcal{M}_\sigma = \langle W, \mathcal{P}_1, \dots, \mathcal{R}_1, \dots \rangle$ by setting $\mathcal{P}_i = \pi(p_i)$.

We start our investigation of the relationship between FO_m^2 and $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ by showing that these logics are equally expressive. If we write $\varphi(x), \varphi(y)$ for formulas, we assume that at most the displayed variable occurs free in φ .

Theorem 1 (Expressive completeness for 2-variable-logic). For every $\varphi \in \mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ there exists a formula $\varphi^\sharp(x) \in FO_m^2$ whose length is linear in the length of φ such that the following holds for all Kripke-models \mathcal{M} and all $a \in W$:

$$\mathcal{M}, a \models \varphi \Leftrightarrow \mathcal{M}_\sigma \models \varphi^\sharp(a).$$

Conversely, given $\varphi(x) \in FO_m^2$ there exists a formula $\varphi^{\sigma_x} \in \mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ whose length is exponential in the length of φ such that the following holds for all Kripke-models \mathcal{M} and all $a \in W$:

$$\mathcal{M}, a \models \varphi^{\sigma_x} \Leftrightarrow \mathcal{M}_\sigma \models \varphi(a).$$

Proof. The proof of the first claim is standard [26, 27], so we concentrate on the second one, whose proof is rather similar to the proof provided in [7] for temporal logics.

An FO^2 -formula $\rho(x, y)$ is called a *binary atom* if it is an atom of the form $R_i(x, y)$, $R_i(y, x)$, or $x = y$. A *binary type* t for a formula ψ is a set of FO^2 -formulas containing (i) either χ or $\neg\chi$ for each binary atom χ occurring in ψ , (ii) either $x = y$ or $x \neq y$, and (iii) no other formulas than these. The set of binary types for ψ is denoted by \mathcal{R}_ψ . A formula ξ is called a *unary atom* if it is of the form $R_i(x, x)$, $R_i(y, y)$, $A_i(x)$, or $A_i(y)$.

Let $\varphi(x) \in FO_m^2$. We assume $\varphi(x)$ is built using \exists , \wedge , and \neg only. We inductively define two mappings \cdot^{σ_x} and \cdot^{σ_y} where the former one takes each FO_m^2 -formula $\varphi(x)$ to the corresponding $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ -formula φ^{σ_x} and the latter does the same for FO_m^2 -formulas $\varphi(y)$. We only give the details of \cdot^{σ_x} since \cdot^{σ_y} is defined analogously by switching the roles of x and y .

Case 1. If $\varphi(x) = P_i(x)$, then put $(\varphi(x))^{\sigma_x} = p_i$.

Case 2. If $\varphi(x) = R_i(x, x)$, then put $(\varphi(x))^{\sigma_x} = \langle id \cap R_i \rangle \top$.

Case 3. If $\varphi(x) = \chi_1 \wedge \chi_2$, then put, recursively, $(\varphi(x))^{\sigma_x} = \chi_1^{\sigma_x} \wedge \chi_2^{\sigma_x}$.

Case 4. If $\varphi(x) = \neg\chi$, then put, recursively, $(\varphi(x))^{\sigma_x} = \neg(\chi)^{\sigma_x}$.

Case 5. If $\varphi(x) = \exists y\chi(x, y)$, then $\chi(x, y)$ can clearly be written as

$$\chi(x, y) = \gamma[\rho_1, \dots, \rho_r, \gamma_1(x), \dots, \gamma_l(x), \xi_1(y), \dots, \xi_s(y)],$$

i.e., as a Boolean combination γ of ρ_i , $\gamma_i(x)$, and $\xi_i(y)$; the ρ_i are binary atoms; the $\gamma_i(x)$ are unary atoms or of the form $\exists y\gamma'_i$; and the $\xi_i(y)$ are unary atoms or of the form $\exists x\xi'_i$. We may assume that x occurs free in $\varphi(x)$. Our first step is to move all formulas without a free variable y out of the scope of \exists : obviously, $\varphi(x)$ is equivalent to

$$\bigvee_{\langle w_1, \dots, w_\ell \rangle \in \{\top, \perp\}^\ell} \left(\bigwedge_{1 \leq i \leq \ell} (\gamma_i \leftrightarrow w_i) \wedge \exists y \gamma(\rho_1, \dots, \rho_r, w_1, \dots, w_l, \xi_1, \dots, \xi_s) \right). \quad (1)$$

For every binary type $t \in \mathcal{R}_\varphi$ and binary atom α from φ , we have $t \models \alpha$ or $t \models \neg\alpha$ —hence we can “guess” a binary type t and then replace all binary atoms by either true or false. For $t \in \mathcal{R}_\varphi$, let $\rho_i^t = \top$ if $t \models \rho_i$, and $\rho_i^t = \perp$, otherwise. Then $\varphi(x)$ is equivalent to

$$\bigvee_{\langle w_1, \dots, w_\ell \rangle \in \{\top, \perp\}^\ell} \left(\bigwedge_{1 \leq i \leq \ell} (\gamma_i \leftrightarrow w_i) \wedge \bigvee_{t \in \mathcal{R}_\varphi} \exists y \left(\left(\bigwedge_{\alpha \in t} \alpha \right) \wedge \gamma(\rho_1^t, \dots, \rho_r^t, w_1, \dots, w_l, \xi_1, \dots, \xi_s) \right) \right). \quad (2)$$

Define, for every negated and unnegated binary atom α , a complex modal parameter α^{σ_x} as follows:

$$\begin{aligned} (x = y)^{\sigma_x} &= id & (\neg(x = y))^{\sigma_x} &= \neg id \\ (R_i(x, y))^{\sigma_x} &= R_i & (\neg R_i(x, y))^{\sigma_x} &= \neg R_i \\ (R_i(y, x))^{\sigma_x} &= R_i^- & (\neg R_i(y, x))^{\sigma_x} &= \neg R_i^-. \end{aligned}$$

Put, for every binary type $t \in \mathcal{R}_\varphi$, $t^{\sigma_x} = \bigcap_{\alpha \in t} \alpha^{\sigma_x}$. Now compute, recursively, $\gamma_i^{\sigma_x}$ and $\xi_i^{\sigma_y}$, and define $\varphi(x)^\sigma$ as

$$\bigvee_{\langle w_1, \dots, w_\ell \rangle \in \{\top, \perp\}^\ell} \left(\bigwedge_{1 \leq i \leq \ell} (\gamma_i^{\sigma_x} \leftrightarrow w_i) \wedge \bigvee_{t \in \mathcal{R}_\varphi} \langle t^{\sigma_x} \rangle \gamma(\rho_1^t, \dots, \rho_r^t, w_1, \dots, w_l, \xi_1^{\sigma_y}, \dots, \xi_s^{\sigma_y}) \right).$$

□

Note that φ^{σ_x} can be computed in polynomial time in the length of φ^{σ_x} . We should like to stress that the existence of formalisms with some ‘modal flavour’ and the same expressive power as FO^2 is known [4, 10]. However, these formalisms have a number of purely technical constructs which did not find applications in modal or description logic. In [4], for example, Borgida constructs a counterpart L' of FO^2 in which accessibility relations \mathcal{R} can be defined as products of extensions of formulas: for any two formulas φ_1, φ_2 one can form $\mathcal{R} = \{w \in W : w \models \varphi_1\} \times \{w \in W : w \models \varphi_2\}$. The expressive completeness result for L' becomes rather straightforward. In fact, the translation provided by Borgida is polynomial so that L' is speaking about relational structures as succinctly as FO^2 does.

Regarding the succinctness of $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$, we show that it is exponentially less succinct than FO^2 .

Theorem 2. *For $n \geq 1$, let φ_n be the following formula of FO^2 :*

$$\begin{aligned} \forall x \exists y \left(\bigwedge_{k=0..n-1} \left(\bigwedge_{j=0..k-1} P_j(x) \right) \rightarrow (P_k(x) \leftrightarrow P_k(y)) \right) \\ \wedge \bigwedge_{k=0..n-1} \left(\bigvee_{j=0..k-1} \neg P_j(x) \right) \rightarrow (P_k(x) \leftrightarrow P_k(y)) \end{aligned}$$

Every formula of $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ that is equivalent to φ_n is of length at least $2^n/2$, for all $n \geq 1$.

Proof. We prove Theorem 2 for the class of *universal structures*, i.e., structures in which all binary relations symbols are interpreted as the cross-product of the domain. Clearly, this implies the general case.

It is easily seen that a first-order structure $\mathcal{M} = \langle W, \mathcal{P}_1, \dots \rangle$ is a model of the formula φ_n iff there exist (distinct) elements $w_0, \dots, w_{2^n-1} \in W$ such that for $i < 2^n$, the truth values of P_0, \dots, P_{n-1} at w_i are the binary encoding of the number i . Moreover, on the class of universal structures, every $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ -formula can be converted into an equivalent formula that uses only the modal

operators $[a \cup \neg a]\varphi$ and $[\neg id]\varphi$ and whose size is bounded polynomially in that of the original formula. Similar to the standard proof of the linear-size model property of the modal logic **S5**, it is easy to show that, if a formula φ of length k in this language is satisfiable in a universal structure, then it is satisfiable in a universal structure of cardinality at most $2k$. Since every formula equivalent to φ_n enforces a domain of cardinality at least 2^n , every equivalent $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ -formula must be of size at least $2^n/2$.

3 Complexity

We show that, for $0 < m < \omega$, $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ -satisfiability is EXPTIME-complete and hence in a lower complexity class than FO_m^2 -satisfiability which is known to be NEXPTIME-complete [14]. The EXPTIME lower bound for $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ -satisfiability is an immediate consequence of the fact that \mathcal{ML}_m^{\neg} is EXPTIME-hard even if $m = 1$ [20]. Hence, we concentrate on the upper bound. It is established by first (polynomially) reducing $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ -satisfiability to a certain variant of $\mathcal{ML}_k^{\neg id}$ -satisfiability (where $\mathcal{ML}_k^{\neg id}$ is multi-modal **K** enriched with the difference modality [6]) and then showing that this variant of $\mathcal{ML}_k^{\neg id}$ -satisfiability can be decided in EXPTIME.

3.1 Reducing $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ to $\mathcal{ML}_{s,t,n}^{\neg id}$

In this section, we generally assume that $0 < m < \omega$. The following languages are used in the reduction:

Definition 2 (Languages). (1) By $\mathcal{ML}_{s,t,n}^{\neg id}$ we denote the modal language $\mathcal{ML}_k^{\neg id}$ with $k = 2s + t + n$ modal parameters

$$\mathfrak{P} = \{K_1, \dots, K_s, I_1, \dots, I_s, X_1, \dots, X_t, Y_1, \dots, Y_n\}$$

and the difference modality $\langle d \rangle$, where d is an abbreviation for $\neg id$.

(2) $\mathcal{ML}_m^{(\neg), \cap, -, id}$ is $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ with negation of modal parameters restricted to atomic modal parameters and without union of modal parameters.

(3) By $\mathcal{ML}_{s,t,n}^-$ we denote the modal language \mathcal{ML}_k^- with converse and $k = s + t + n$ modal parameters $K_1, \dots, K_s, X_1, \dots, X_t, Y_1, \dots, Y_n$.

Definition 3 (Semantics). A structure

$$\mathcal{M} = \langle W, \mathcal{K}_1, \dots, \mathcal{K}_s, \mathcal{X}_1, \dots, \mathcal{X}_t, \mathcal{Y}_1, \dots, \mathcal{Y}_n \rangle$$

is called a c-frame iff

1. the relations \mathcal{K}_i are irreflexive and antisymmetric,
2. the relations \mathcal{X}_i are irreflexive and symmetric,
3. the relations \mathcal{Y}_i are subsets of $\{(w, w) \mid w \in W\}$,

4. for all $w, w' \in W$ with $w \neq w'$, there exists a unique

$$\mathcal{S} \in \{\mathcal{K}_1, \dots, \mathcal{K}_s, \mathcal{K}_1^{-1}, \dots, \mathcal{K}_s^{-1}, \mathcal{X}_1, \dots, \mathcal{X}_t\}$$

such that $(w, w') \in \mathcal{S}$, and

5. for each $w \in W$, there exists a unique i with $1 \leq i \leq n$ such that $(w, w) \in \mathcal{Y}_i$,

where \mathcal{R}_i^{-1} is used to denote the converse of a binary relation \mathcal{R}_i . An $\mathcal{ML}_{s,t,n}^-$ -formula is called *c-satisfiable* iff it has a model which is based on a *c-frame*. Such a model is called a *c-model*.

A structure $\mathcal{M} = \langle W, \mathcal{K}_1, \dots, \mathcal{K}_s, \mathcal{I}_1, \dots, \mathcal{I}_s, \mathcal{X}_1, \dots, \mathcal{X}_t, \mathcal{Y}_1, \dots, \mathcal{Y}_n \rangle$ is called an *s-frame* iff there exists a *c-frame*

$$\mathcal{M}' = \langle W, \mathcal{K}'_1, \dots, \mathcal{K}'_s, \mathcal{X}'_1, \dots, \mathcal{X}'_t, \mathcal{Y}'_1, \dots, \mathcal{Y}'_n \rangle$$

such that $\mathcal{K}_i \subseteq \mathcal{K}'_i$, $\mathcal{I}_i \subseteq \mathcal{K}_i^{-1}$, $\mathcal{X}_i \subseteq \mathcal{X}'_i$, and $\mathcal{Y}_i \subseteq \mathcal{Y}'_i$. An $\mathcal{ML}_{s,t,n}^{-id}$ -formula is called *s-satisfiable* iff it is satisfiable in a model based on an *s-frame*. Such a model is called an *s-model*.

A *literal* is a modal parameter that matches one of the following descriptions:

- an atomic parameter or the negation thereof,
- the inverse of an atomic parameter or the negation thereof,
- the identity parameter or the negation of the identity parameter.

The reduction is comprised of a series of polynomial reduction steps. Let φ be a $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ -formula.

Step 1. Exhaustively apply the following rewrite rules to modal parameters in φ :

$$\begin{array}{lll} (\neg S)^- \rightsquigarrow \neg(S^-) & (S_1 \cup S_2)^- \rightsquigarrow S_1^- \cup S_2^- & id^- \rightsquigarrow id \\ S^{--} \rightsquigarrow S & (S_1 \cap S_2)^- \rightsquigarrow S_1^- \cap S_2^- & \neg id^- \rightsquigarrow \neg id \end{array}$$

In the resulting formula φ_1 , all modal parameters are Boolean combinations of literals.

Step 2. Convert all modal parameters in φ_1 to disjunctive normal form over literals using a truth table (as, e.g., described in [22], page 20). If the “empty disjunction” is obtained when converting a modal parameter S , then replace every occurrence of $\langle S \rangle \psi$ with \perp . Call the result of the conversion φ_2 . The conversion can be done in linear time since the number m of atomic modal and we use a truth table for the conversion (instead of applying equivalences). It is easy to see that φ_2 is satisfiable iff φ_1 is satisfiable. Since the conversion to DNF was done using a truth table, each disjunct occurring in a modal parameter in φ_2 is a *relational type*, i.e., of the form $S_0 \cap S_1 \cap \dots \cap S_m \cap S'_1 \cap \dots \cap S'_m$, where

1. $S_0 = id$ or $S_0 = \neg id$,
2. $S_i = R_i$ or $S_i = \neg R_i$ for $1 \leq i \leq m$, and
3. $S'_i = R_i^-$ or $S'_i = \neg(R_i^-)$ for $1 \leq i \leq m$.

Let $\Gamma_{=}$ be the set of all relational types with $S_0 = id$, Γ_{\neq} be the set of all relational types with $S_0 = \neg id$, and $\Gamma = \Gamma_{=} \cup \Gamma_{\neq}$.

Step 3. We reduce satisfiability of $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ -formulas of the form of φ_2 (i.e, the modal parameters are disjunctions of relational types) to the satisfiability of $\mathcal{ML}_m^{(\neg), \cap, -, id}$ -formulas in which all modal parameters are relational types. As the first step, recursively apply the following substitution to φ_2 from the inside to the outside (i.e., no union on modal parameters occurs in φ)

$$\langle S_1 \cup \dots \cup S_k \rangle \varphi \rightsquigarrow \langle S_1 \rangle p_\varphi \vee \dots \vee \langle S_k \rangle p_\varphi$$

where p_φ is a new propositional variable. Call the result of these substitutions φ'_2 . Secondly, define

$$\varphi_3 := \varphi'_2 \wedge \bigwedge_{p_\varphi \text{ occurs in } \varphi'_2} \bigwedge_{S \in \Gamma} [S](p_\varphi \leftrightarrow \varphi).$$

φ_3 is an $\mathcal{ML}_m^{(\neg), \cap, -, id}$ -formula as required.⁴ Furthermore, φ_2 is satisfiable iff φ_3 is satisfiable, and the reduction is linear.

Step 4. It is not hard to see that the set Γ_{\neq} (from Step 3) can be partitioned into three sets Γ_{\neq}^s , Γ_{\neq}^1 , and Γ_{\neq}^2 such that there exists a bijection F from Γ_{\neq}^1 onto Γ_{\neq}^2 and, for every Kripke structure \mathcal{M} with set of worlds W , and $w, w' \in W$, the following holds:

1. for all $S \in \Gamma_{\neq}^s$: $\mathcal{M}, (w, w') \models S$ iff $\mathcal{M}, (w', w) \models S$ and
2. for all $S \in \Gamma_{\neq}^1$: $\mathcal{M}, (w, w') \models S$ iff $\mathcal{M}, (w', w) \models F(S)$.

Given this, it is easy to reduce satisfiability of $\mathcal{ML}_m^{(\neg), \cap, -, id}$ -formulas of the form of φ_3 to c-satisfiability of $\mathcal{ML}_{s,t,n}^-$ -formulas, where $s = |\Gamma_{\neq}^1|$, $t = |\Gamma_{\neq}^s|$, and $n = |\Gamma_{=} |$. Let r be some bijection between Γ_{\neq}^1 and the set $\{K_1, \dots, K_s\}$, r' some bijection between Γ_{\neq}^s and the set $\{X_1, \dots, X_t\}$, and r'' some bijection between $\Gamma_{=}$ and the set $\{Y_1, \dots, Y_n\}$. The formula φ_4 is obtained from φ_3 by replacing (1.) each element S of Γ_{\neq}^1 that appears in φ_3 with $r(S)$, (2.) each element S of Γ_{\neq}^2 with $r(S)^-$, (3.) each element S of Γ_{\neq}^s with $r'(S)$, and (4.) each element S of $\Gamma_{=}$ with $r''(S)$. It can be proved that φ_3 is satisfiable iff φ_4 is c-satisfiable. Furthermore, the reduction is obviously linear.

Step 5. We reduce c-satisfiability of $\mathcal{ML}_{s,t,n}^-$ -formulas to s-satisfiability of $\mathcal{ML}_{s,t,n}^{-id}$ -formulas. W.l.o.g., we assume that φ_3 does not contain modal parameters of the form X_i^- and Y_i^- : since these parameters are interpreted by symmetric relations, X_i^- (resp. Y_i^-) can be replaced by X_i (resp. Y_i). For $\chi \in \mathcal{ML}_{s,t,n}^-$ (without X_i^- and Y_i^-), denote by $\chi^* \in \mathcal{ML}_{s,t,n}^{-id}$ the formula obtained from χ by replacing all occurrences of K_i^- with I_i .

⁴ We use $\bigwedge_{S \in \Gamma} [S](p_\varphi \leftrightarrow \varphi)$ instead of the more natural $[R](p_\varphi \leftrightarrow \varphi) \wedge [\neg R](p_\varphi \leftrightarrow \varphi)$ (for some atomic R) to ensure that all modal parameters in φ_3 are still relational types after the application of Step 3.

For each $S \in \mathfrak{P}$ (see Definition 2), we use S^\sim to denote (i) I_i if $S = K_i$, (ii) K_i if $S = I_i$, (iii) X_i if $S = X_i$, and (iv) Y_i if $S = Y_i$. For convenience, we define two more sets

$$\mathfrak{P}_1 = \{K_1, \dots, K_s, I_1, \dots, I_s, X_1, \dots, X_t\} \text{ and } \mathfrak{P}_2 = \{Y_1, \dots, Y_n\}$$

. Define φ_5 as the conjunction of φ_4^* with all formulas $\vartheta \wedge [d]\vartheta$, where ϑ can be obtained from the following formulas by replacing ψ and all ψ_S with subformulas of φ_4^* .

$$\begin{aligned} \chi_1 &:= \left(\bigwedge_{S \in \mathfrak{P}_2} [S]\psi_S \right) \rightarrow \left(\bigvee_{S \in \mathfrak{P}_2} \psi_S \right) \\ \chi_2 &:= \bigwedge_{\mathcal{P} \subseteq \mathfrak{P}_1} \left[\left(\bigwedge_{S \in \mathcal{P}} [S]\neg\psi_S \wedge \langle d \rangle \left(\bigwedge_{S \in \mathcal{P}} \psi_S \wedge \bigwedge_{S \in \mathfrak{P}_1 \setminus \mathcal{P}} [S^\sim]\neg\psi_S \right) \right) \rightarrow \bigvee_{S \in \mathfrak{P}_1 \setminus \mathcal{P}} \neg\psi_S \right] \\ \chi_3 &:= \bigwedge_{S \in \mathfrak{P}} \psi \rightarrow [S]\langle S^\sim \rangle \psi \end{aligned}$$

Obviously, φ_5 is an $\mathcal{ML}_{s,t,n}^{-id}$ -formula. The formula χ_1 deals with Item 5 from the definition of c-frames, χ_2 with Item 4, and χ_3 with symmetry from Item 2 and with the semantics of the converse operator. Note that the length of φ_5 is polynomial in the length $|\varphi_4|$ of φ_4 since the set of modal parameters is fixed.

Lemma 1. φ_4 is c-satisfiable iff φ_5 is s-satisfiable.

Proof: The “only if” direction is straightforward: Let

$$\mathcal{M} = \langle W, \pi, \mathcal{K}_1, \dots, \mathcal{K}_s, \mathcal{X}_1, \dots, \mathcal{X}_t, \mathcal{Y}_1, \dots, \mathcal{Y}_n \rangle$$

be a c-model for φ_4 . It is readily checked that

$$\mathcal{M}' = \langle W, \pi, \mathcal{K}_1, \dots, \mathcal{K}_s, \mathcal{K}_1^{-1}, \dots, \mathcal{K}_s^{-1}, \mathcal{X}_1, \dots, \mathcal{X}_t, \mathcal{Y}_1, \dots, \mathcal{Y}_n \rangle$$

is an s-frame and that the ϑ formulas from above are true in \mathcal{M}' . Hence, by the semantics of converse, \mathcal{M}' is obviously a model for φ_5 .

It remains to prove the “if” direction. Let

$$\mathcal{M} = \langle W, \pi, \mathcal{K}_1, \dots, \mathcal{K}_s, \mathcal{I}_1, \dots, \mathcal{I}_s, \mathcal{X}_1, \dots, \mathcal{X}_t, \mathcal{Y}_1, \dots, \mathcal{Y}_n \rangle$$

be an s-model for φ_5 . In particular, this implies that all formulas derived from χ_1 to χ_3 are true in \mathcal{M} . Before we construct the c-model for φ_4 , we prove two claims:

Claim 1. For each $w, w' \in W$ with $w \neq w'$, there exists an $S \in \mathfrak{P}_1$ such that, for all subformulas ψ of φ_4^* , we have that $\mathcal{M}, w \not\models \langle S \rangle \psi$ implies $\mathcal{M}, w' \not\models \psi$ and $\mathcal{M}, w' \not\models \langle S^\sim \rangle \psi$ implies $\mathcal{M}, w \not\models \psi$.

Proof: Assume that the claim does not hold. Fix $w, w' \in W$ with $w \neq w'$ that do not have the property from the claim. This means that, for each $S \in \mathfrak{P}_1$,

- (i) there is a subformula ψ_S^1 of φ_4^* such that $\mathcal{M}, w \models [S]\neg\psi_S^1$ and $\mathcal{M}, w' \models \psi_S^1$
or
(ii) there is a subformula ψ_S^2 of φ_4^* such that $\mathcal{M}, w' \models [S^\sim]\neg\psi_S^2$ and $\mathcal{M}, w \models \psi_S^2$.

Let \mathcal{P} be the subset of \mathfrak{P}_1 such that $S \in \mathcal{P}$ iff S satisfies (i) and let $\psi_S = \psi_S^1$ if $S \in \mathcal{P}$ and $\psi_S = \psi_S^2$ otherwise. Let ϑ be the instantiation of χ_2 with \mathcal{P} and the ψ_S .⁵ Since all formulas derived from χ_2 are true in \mathcal{M} , we have $\mathcal{M}, w \models \vartheta$. It is straightforward to verify that this is a contradiction to the properties of the ψ_S as stated under (i) and (ii). \square

Claim 2. For each $w \in W$, there exists an $S \in \mathfrak{P}_2$ such that, for all subformulas ψ of φ_4^* , we have that $\mathcal{M}, w \models [S]\psi$ implies $\mathcal{M}, w \models \psi$.

Proof: Similar to the previous claim, only simpler using χ_1 in place of χ_2 . \square

Construct a Kripke model $\mathcal{M}' = \langle W, \pi, \mathcal{K}'_1, \dots, \mathcal{K}'_s, \mathcal{X}'_1, \dots, \mathcal{X}'_t, \mathcal{Y}'_1, \dots, \mathcal{Y}'_n \rangle$ as follows: Initially, set $\mathcal{K}'_i := \mathcal{K}_i \cup \mathcal{I}_i^{-1}$, $\mathcal{X}'_i := \mathcal{X}_i \cup \mathcal{X}_i^{-1}$, and $\mathcal{Y}'_i := \mathcal{Y}_i$. Then, augment the relations as follows:

1. For each $w, w' \in W$ with $w \neq w'$, if

$$(w, w') \notin \bigcup_{1 \leq i \leq s} \mathcal{K}'_i \cup \bigcup_{1 \leq i \leq s} (\mathcal{K}'_i)^{-1} \cup \bigcup_{1 \leq i \leq t} \mathcal{X}'_i$$

then choose an $S \in \mathfrak{P}_1$ as in Claim 1 and set

- $\mathcal{K}'_i := \mathcal{K}'_i \cup \{(w, w')\}$ if $S = K_i$,
 - $\mathcal{K}'_i := \mathcal{K}'_i \cup \{(w', w)\}$ if $S = I_i$, and
 - $\mathcal{X}'_i := \mathcal{X}'_i \cup \{(w, w'), (w', w)\}$ if $S = X_i$.
2. For each $w \in W$, if $(w, w) \notin \bigcup_{1 \leq i \leq n} \mathcal{Y}'_i$ then choose a $Y_i \in \mathfrak{P}_2$ as in Claim 2 and set $\mathcal{Y}'_i := \mathcal{Y}'_i \cup \{(w, w)\}$.

It is not hard to check that \mathcal{M}' is a c-model, i.e., that the properties from Definition 2 are satisfied. It hence remains to prove that

$$\mathcal{M}, w \models \psi^* \text{ iff } \mathcal{M}', w \models \psi$$

for all subformulas ψ of φ_4 . The proof is by induction over the structure of ψ . The only interesting case is the following:

$\psi = \langle T \rangle \rho$. Let

- $S = K_i, \mathcal{S} = \mathcal{K}_i, \mathcal{T} = \mathcal{K}'_i$ if $T = K_i$,
- $S = I_i, \mathcal{S} = \mathcal{I}_i, \mathcal{T} = (\mathcal{K}'_i)^{-1}$ if $T = K_i^-$,
- $S = X_i, \mathcal{S} = \mathcal{X}_i, \mathcal{T} = \mathcal{X}'_i$ if $T = X_i$, and
- $S = Y_i, \mathcal{S} = \mathcal{Y}_i, \mathcal{T} = \mathcal{Y}'_i$ if $T = Y_i$.

⁵ For the cases $\mathcal{P} = \emptyset$ and $\mathcal{P} = \mathfrak{P}_1$, we assume that the “empty” conjunction is equivalent to \top and the “empty” disjunction equivalent to \perp .

Obviously, $\psi^* = \langle S \rangle \rho^*$. The “only if” direction is trivial since $\mathcal{S} \subseteq \mathcal{T}$ by construction of \mathcal{M}' . Hence, let us concentrate on the “if” direction. We prove the contrapositive, i.e.,

$$\mathcal{M}, w \not\models \langle S \rangle \rho^* \text{ implies } \mathcal{M}', w \not\models \langle T \rangle \rho.$$

Hence, assume $\mathcal{M}, w \not\models \langle S \rangle \rho^*$. We show that $(w, w') \in \mathcal{T}$ implies $\mathcal{M}', w' \not\models \rho$. Several cases need to be distinguished:

- (w, w') has been added to \mathcal{T} in the initial step due to $(w, w') \in \mathcal{S}$.⁶ Since $\mathcal{M}, w \not\models \langle S \rangle \rho^*$, this implies $\mathcal{M}, w' \not\models \rho^*$. By induction, we have $\mathcal{M}', w' \not\models \rho$ and are done.
- (w, w') has been added to \mathcal{T} in the initial step, but $(w, w') \notin \mathcal{S}$.⁵ Assume to the contrary that $\mathcal{M}', w' \models \rho$. By induction, we have $\mathcal{M}, w' \models \rho^*$. Let $\mathcal{S}^\sim = \mathcal{I}_i$ if $\mathcal{S} = \mathcal{K}_i$, and $\mathcal{S}^\sim = \mathcal{S}^{-1}$ otherwise. By definition of the initial step and since $(w, w') \notin \mathcal{S}$, we have $(w', w) \in \mathcal{S}^\sim$. Since $\mathcal{M}, w' \models \rho^*$ and $\mathcal{M}, w' \models \chi'_3$ for all instantiations χ'_3 of χ_3 , we have $\mathcal{M}, w' \models [S^\sim] \langle S \rangle \rho^*$ and, since $(w', w) \in \mathcal{S}^{-1}$, we have also $\mathcal{M}, w \models \langle S \rangle \rho^*$. This is obviously in contradiction to our initial assumption $\mathcal{M}, w \not\models \langle S \rangle \rho^*$ and we conclude that $\mathcal{M}, w' \not\models \rho$.
- (w, w') has been added to \mathcal{T} during the first part of the augmentation step. We make a case distinction as follows:
 - $\mathcal{T} = \mathcal{K}'_i$. Then $S = K_i$ was chosen in the augmentation step leading to the addition of (w, w') to \mathcal{K}'_i . Hence, S is as in Claim 1 for (w, w') . Since, clearly, ρ^* is a subformula of φ_4^* and $\mathcal{M}, w \not\models \langle S \rangle \rho^*$, this implies $\mathcal{M}, w \not\models \rho^*$. It remains to apply the induction hypothesis.
 - $\mathcal{T} = (\mathcal{K}'_i)^{-1}$. Then $S = I_i$ was chosen in the augmentation step leading to the addition of (w', w) to \mathcal{K}'_i . Hence, S is as in Claim 1 for (w, w') . Again, this implies $\mathcal{M}, w \not\models \rho^*$ and it remains to apply the induction hypothesis.
 - $\mathcal{T} = \mathcal{X}_i$. Then $S = X_i$ was chosen in the augmentation step leading to the addition of (w, w') to \mathcal{K}_i . Hence, by definition of the augmentation step, S either S is as in Claim 1 for (w, w') or S is as in Claim 1 for (w', w) . Since $X_i^\sim = X_i$, the latter case implies the further case and we conclude that S is as in Claim 1 for (w, w') . We continue as in the previous cases.
- (w, w') has been added during the second part of the augmentation step. Analogous to the subcase $\mathcal{T} = \mathcal{K}'_i$ in the previous case.

Since \mathcal{M} is a model for φ_4^* , we have that \mathcal{M}' is a model for φ_4 . □

⁶ This also includes the case that $\mathcal{T} = (\mathcal{K}'_i)^{-1}$ and (w', w) have been added to \mathcal{K}_i .

3.2 An EXP_{TIME} upper bound for $\mathcal{ML}_{s,t,n}^{\neg id}$

We show that s -satisfiability of $\mathcal{ML}_{s,t,n}^{\neg id}$ -formulas can be decided in deterministic exponential time. Consider an $\mathcal{ML}_{s,t,n}^{\neg id}$ -formula φ with modal parameters from $\{d\} \cup \mathfrak{P}_1 \cup \mathfrak{P}_2$ as defined above. Denote by $\text{cl}(\varphi)$ the closure under single negation of the set of all subformulas of φ . In what follows we identify $\neg\neg\psi$ with ψ . A φ -type \mathbf{t} is a subset of $\text{cl}(\varphi)$ with

- $\neg\chi \in \mathbf{t}$ iff $\chi \notin \mathbf{t}$, for all $\neg\chi \in \text{cl}(\varphi)$;
- $\chi_1 \wedge \chi_2 \in \mathbf{t}$ iff $\chi_1, \chi_2 \in \mathbf{t}$, for all $\chi_1 \wedge \chi_2 \in \text{cl}(\varphi)$.

Given a world w in a model, the set of formulas in $\text{cl}(\varphi)$ which are realized in w is a (φ) -type. We use the following notation:

- for $R \in \mathfrak{P}$ we write $\mathbf{t}_1 \rightarrow_R \mathbf{t}_2$ iff $\{\neg\chi \mid \neg\langle R \rangle \chi \in \mathbf{t}_1\} \subseteq \mathbf{t}_2$;
- a φ -type \mathbf{t} is called a χ -singleton type if $\{\chi, \neg\langle d \rangle \chi\} \subseteq \mathbf{t}$.

Intuitively, singleton types are types which cannot be realized by two different worlds in a model. A *candidate* for φ is a maximal set (w.r.t. \subseteq) \mathcal{T} of φ -types with the following properties:

- (C1)** for all $\mathbf{t} \in \mathcal{T}$: if $\langle Y_i \rangle \chi_1 \in \mathbf{t}$ and $\langle Y_j \rangle \chi_2 \in \mathbf{t}$, then $i = j$;
- (C2)** for all $\mathbf{t} \in \mathcal{T}$: if for some $i \leq n$, $\langle Y_i \rangle \chi \in \mathbf{t}$, then $\mathbf{t} \rightarrow_{Y_i} \mathbf{t}$ and $\{\chi \mid \langle Y_i \rangle \chi \in \mathbf{t}\} \subseteq \mathbf{t}$;
- (C3)** if \mathcal{T} contains a χ -singleton type \mathbf{t} , then $\neg\chi \in \mathbf{t}'$, for all $\mathbf{t}' \in \mathcal{T} - \{\mathbf{t}\}$,
- (C4)** for every $\langle d \rangle \chi \in \text{cl}(\varphi)$ and $\mathbf{t}, \mathbf{t}' \in \mathcal{T}$: $\neg\chi, \neg\langle d \rangle \chi \in \mathbf{t}$ iff $\neg\chi, \neg\langle d \rangle \chi \in \mathbf{t}'$.

Intuitively, (C1) says that it suffices to add at most a single reflexive edge Y_i to each world of type \mathbf{t} which is necessary since we are heading for s -models. By (C2), for each $\langle Y_i \rangle$ -formula in \mathbf{t} we find a witness in \mathbf{t} itself. (C3) states that, for every $\langle d \rangle \chi \in \text{cl}(\varphi)$, \mathcal{T} does not contain more than one χ -singleton type (C3). (C4) should be obvious by the semantics of $\langle d \rangle$. The number of candidates for φ is exponential in the length of φ .

1. for each $\langle d \rangle \chi \in \text{cl}(\varphi)$, each candidate \mathcal{T} contains at most one χ -singleton type,
2. the number of distinct $\langle d \rangle$ -formulas in $\text{cl}(\varphi)$ is linear in the length of φ , and
3. if \mathcal{T} and \mathcal{T}' agree on the singleton types and $\mathcal{T} \cap \mathcal{T}' \neq \emptyset$, then $\mathcal{T} = \mathcal{T}'$ since candidates are maximal sets.

This yields an upper bound of $n_\varphi = 2^{(|\text{cl}(\varphi)|+1)^2}$ for the number of candidates.

A *relational candidate* is a triple $\langle \mathcal{T}, \mathcal{F}, \mathcal{I} \rangle$ consisting of

- a candidate \mathcal{T} for φ ;
- a function $\mathcal{F} : \{1, \dots, k\} \rightarrow \mathcal{T}_N$ with $k \leq |\text{cl}(\varphi)|^2$ (in what follows we often use $\mathcal{F} = \{(1, \mathcal{F}(1)), \dots, (k, \mathcal{F}(k))\}$);
- and a function \mathcal{I} mapping each modal parameter $R \in \mathfrak{P}$ to a relation $R^{\mathcal{I}} \subseteq (\mathcal{T}_S \cup \mathcal{F}) \times (\mathcal{T}_S \cup \mathcal{F})$ such that

- (R1)** $\langle \mathcal{T}_S \cup \mathcal{F}, (R^{\mathcal{I}} : R \in \mathfrak{P}) \rangle$ is an s-frame;
- (R2)** for all $R \in \mathfrak{P}$, $m, m' \leq k$ and types $\mathfrak{t}, \mathfrak{t}'$: if $\mathfrak{t}R^{\mathcal{I}}\mathfrak{t}'$, $\mathfrak{t}R^{\mathcal{I}}(m, \mathfrak{t}')$, $(m, \mathfrak{t})R^{\mathcal{I}}\mathfrak{t}'$, or $(m, \mathfrak{t})R^{\mathcal{I}}(m', \mathfrak{t}')$, then $\mathfrak{t} \rightarrow_R \mathfrak{t}'$;
- (R3)** for all $R \in \mathfrak{P}$, $\langle R \rangle \chi \in \text{cl}(\varphi)$, and $\mathfrak{t} \in \mathcal{T}_S$ with $\langle R \rangle \chi \in \mathfrak{t}$ we find $\mathfrak{t}' \in \mathcal{T}_S$ with $\mathfrak{t}R^{\mathcal{I}}\mathfrak{t}'$ and $\chi \in \mathfrak{t}'$ or we find $(m, \mathfrak{t}') \in \mathcal{F}$ with $\mathfrak{t}R^{\mathcal{I}}(m, \mathfrak{t}')$ and $\chi \in \mathfrak{t}'$;
- (R4)** for all $R \in \mathfrak{P}_2$ and $(m, \mathfrak{t}) \in \mathcal{F}$, if $\langle R \rangle \chi \in \mathfrak{t}$, then $(m, \mathfrak{t})R^{\mathcal{I}}(m, \mathfrak{t})$.

Intuitively, \mathcal{T}_S is the set of worlds realizing singleton types, \mathcal{F} is the set of worlds providing witnesses for diamond formulas in singleton types, and \mathcal{I} fixes the extension of the modal parameters on $\mathcal{T}_S \cup \mathcal{F}$. Note that \mathcal{F} need not contain more than $|\text{cl}(\varphi)|^2$ worlds since each candidate contains at most $|\text{cl}(\varphi)|$ singleton types (one for each $\langle d \rangle \chi \in \text{cl}(\varphi)$, see above) and each type may contain at most $|\text{cl}(\varphi)|$ diamond formulas. (R2) ensures that the relations fixed satisfy all box formulas. (R3) guarantees that diamond-formulas in $\mathfrak{t} \in \mathcal{T}_S$ with parameters $R \in \mathfrak{P}$ have witnesses in $\mathcal{T}_S \cup \mathcal{F}$. And (R4) says that relations from \mathfrak{P}_2 are interpreted by \mathcal{I} as enforced by the diamond formulas. We need not consider types from \mathcal{T}_S in (R4) since the corresponding claim already follows from (R1) and (R3). The number of relational candidates is bounded by $n_\varphi \cdot 2^{|\text{cl}(\varphi)|^3} \cdot |\text{cl}(\varphi)|^{6 \cdot |\mathfrak{P}|+2}$ [?].

The number of relational candidates is bounded by $n_\varphi \cdot 2^{|\text{cl}(\varphi)|^3} \cdot |\text{cl}(\varphi)|^{6 \cdot |\mathfrak{P}|+2}$: there exist at most n_φ candidates and, for each candidate, there are at most $|\text{cl}(\varphi)|^2 \cdot 2^{|\text{cl}(\varphi)|^3}$ choices for \mathcal{F} and, for given \mathcal{T} and \mathcal{F} , at most $|\text{cl}(\varphi)|^{6 \cdot |\mathfrak{P}|}$ choices for \mathcal{I} .

Our algorithm enumerates all (exponentially many) relational candidates and performs, for each such candidate, an elimination procedure that checks whether the candidate under consideration induces a model or not. Concerning the enumeration of relational candidates, note that it can be checked in polynomial time whether some \mathcal{I} defines an s-frame as required by (R1) above: It is tedious but straightforward to write down explicit conditions that determine s-frames. We now describe the elimination procedure. Intuitively, we remove those non-singleton types whose diamond formulas are not witnessed: for a given relational candidate $\langle \mathcal{T}, \mathcal{F}, \mathcal{I} \rangle$ we can form a sequence $\mathcal{T} = \mathcal{T}_0 \supseteq \mathcal{T}_1 \supseteq \dots$ inductively as follows: put $\mathcal{T}_0 = \mathcal{T}$. Suppose \mathcal{T}_i is defined. Then delete non-singleton types $\mathfrak{t} \in \mathcal{T}_i$ which are not in the range of \mathcal{F} whenever

- (E1)** there are no pairwise disjoint relations $R^{\mathcal{I}} \subseteq \{\mathfrak{t}\} \times \mathcal{T}_S$ for all $R \in \mathfrak{P}_1$, such that (i) $\mathfrak{t} \rightarrow_R \mathfrak{t}'$ whenever $\mathfrak{t}R^{\mathcal{I}}\mathfrak{t}'$, and (ii) for all $\langle R \rangle \chi \in \mathfrak{t}$, $R \in \mathfrak{P}_1$, there exists \mathfrak{t}' with $\mathfrak{t}R^{\mathcal{I}}\mathfrak{t}'$ and $\chi \in \mathfrak{t}'$ or there exists $\mathfrak{t}' \in \mathcal{T}_i - \mathcal{T}_S$ with $\mathfrak{t} \rightarrow_R \mathfrak{t}'$ and $\chi \in \mathfrak{t}'$, or
- (E2)** there is $\langle d \rangle \chi \in \mathfrak{t}$ but no $\mathfrak{t}' \in \mathcal{T}_i$ with $\chi \in \mathfrak{t}'$

and denote the result by \mathcal{T}_{i+1} . Clearly, $\mathcal{T}_i = \mathcal{T}_{i+1}$ after at most $2^{|\text{cl}(\varphi)|}$ rounds. We denote the result of the elimination procedure started on \mathcal{T} with $\widehat{\mathcal{T}}$. Obviously, for each non-singleton type \mathfrak{t} in $\widehat{\mathcal{T}}$ which is not in the range of \mathcal{F} , each diamond formula in \mathfrak{t} is witnessed by some type in $\widehat{\mathcal{T}}$ such that at most one “edge” from \mathfrak{t} to any $\mathfrak{t}' \in \mathcal{T}_S$ is required (this is crucial for building s-models). Together with (R3), (C2), and (R4), this implies that the only diamond formulas not witnessed

in $\widehat{\mathcal{T}}$ are either $\langle R \rangle$ -formulas in types from the range of \mathcal{F} with $R \in \{d\} \cup \mathfrak{P}_1$, or are $\langle d \rangle$ -formulas in types from \mathcal{T}_S . Since we are building s-models, we must be careful choosing singleton types as witnesses for these formulas:

Lemma 2. φ is s-satisfiable iff there exists a relational candidate $\langle \mathcal{T}, \mathcal{F}, \mathcal{I} \rangle$ such that $\langle \widehat{\mathcal{T}}, \mathcal{F}, \mathcal{I} \rangle$ has the following properties:

- there exists $\mathfrak{t} \in \widehat{\mathcal{T}}$ with $\varphi \in \mathfrak{t}$,
- for every $(m, \mathfrak{t}) \in \mathcal{F}$ and all $\langle R \rangle \chi \in \mathfrak{t}$ with $R \in \mathfrak{P}_1$: (i) there exists $\mathfrak{t}' \in \widehat{\mathcal{T}} - \mathcal{T}_S$ with $\mathfrak{t} \rightarrow_R \mathfrak{t}'$ and $\chi \in \mathfrak{t}'$ or (ii) there exists $\mathfrak{t}' \in \mathcal{T}_S$ with $(m, \mathfrak{t})R^{\mathcal{I}}\mathfrak{t}'$ and $\chi \in \mathfrak{t}'$.
- for every $(m, \mathfrak{t}) \in \mathcal{F}$ and $\langle d \rangle \chi \in \mathfrak{t}$ we find $\mathfrak{t}' \in \widehat{\mathcal{T}}$ with $\chi \in \mathfrak{t}'$.
- for every $\mathfrak{t} \in \mathcal{T}_S$ and $\langle d \rangle \chi \in \mathfrak{t}$ we find a $\mathfrak{t}' \in \widehat{\mathcal{T}}$ with $\mathfrak{t} \neq \mathfrak{t}'$ such that $\chi \in \mathfrak{t}'$.

Proof. Suppose φ is s-satisfiable. Take a witness $\langle W, \pi, \mathcal{R}_1, \dots, \mathcal{R}_k \rangle$. Let, for $w \in W$,

$$\mathfrak{t}(w) = \{\chi \in \text{cl}(\varphi) \mid w \models \chi\},$$

and $\mathcal{T} = \{\mathfrak{t}(w) \mid w \in W\}$. Due to the semantics of the modal operator $\langle d \rangle$, for each singleton type $\mathfrak{t} \in \mathcal{T}$, we find precisely one $w_{\mathfrak{t}}$ with $\mathfrak{t}(w_{\mathfrak{t}}) = \mathfrak{t}$. Select, for each singleton type $\mathfrak{t} \in \mathcal{T}$ and each $\langle R \rangle \chi \in \mathfrak{t}$, $R \in \mathfrak{P}_1$, a world $v_{\mathfrak{t}, \langle R \rangle \chi} \in W$ such that $w_{\mathfrak{t}}\mathcal{R}v_{\mathfrak{t}, \langle R \rangle \chi}$ and $v_{\mathfrak{t}, \langle R \rangle \chi} \models \chi$. Let v_1, \dots, v_r be an enumeration of those $v_{\mathfrak{t}, \langle R \rangle \chi}$ for which $\mathfrak{t}(v_{\mathfrak{t}, \langle R \rangle \chi})$ is not a singleton-type and put

$$\mathcal{F} = \{(i, \mathfrak{t}(v_i)) \mid 1 \leq i \leq r\}.$$

Note that $r \leq |\text{cl}(\varphi)|^2$. Let, for $x, y \in \mathcal{T}_S \cup \mathcal{F}$ and $R \in \mathfrak{P}$:

$$xR^{\mathcal{I}}y \Leftrightarrow \begin{cases} w_{\mathfrak{t}}\mathcal{R}v_m & : x = \mathfrak{t}, y = (m, \mathfrak{t}'), \\ v_m\mathcal{R}w_{\mathfrak{t}} & : x = (m, \mathfrak{t}'), y = \mathfrak{t}, \\ w_{\mathfrak{t}}\mathcal{R}w_{\mathfrak{t}'} & : x = \mathfrak{t}, y = \mathfrak{t}', \\ v_n\mathcal{R}v_m & : x = (n, \mathfrak{t}), y = (m, \mathfrak{t}'). \end{cases}$$

Now take a candidate $\mathcal{S} \supseteq \mathcal{T}$ (\mathcal{T} itself may violate the maximality condition) containing precisely the singleton types from \mathcal{T} , and \mathcal{I} and \mathcal{F} as defined above. It is easy to see that the elimination procedure applied to $\langle \mathcal{S}, \mathcal{F}, \mathcal{I} \rangle$ terminates with a structure satisfying the four properties in Lemma 2.

Conversely, suppose the elimination procedure terminates with $\langle \mathcal{T}, \mathcal{F}, \mathcal{I} \rangle$ that satisfies all four conditions in Lemma 2. We define an s-model satisfying φ as follows: W consists of $\mathcal{T}_S \cup \mathcal{F}$ and the set W_S of finite sequences

$$(\mathfrak{t}_{i_0}, R_{i_0}, \mathfrak{t}_{i_1}, R_{i_1}, \dots, \mathfrak{t}_{i_k}),$$

where $\mathfrak{t}_{i_j} \in \mathcal{T}_N$, $R_{i_j} \in \mathfrak{P}_1$, and $k \geq 0$.

Note that adding the *paths* from W_S instead of elements of \mathcal{T}_N to $\mathcal{T}_S \cup \mathcal{F}$ allows us to make sure that the same type reached via different paths yields different worlds. Like in standard unravelling, a path represents its last element,

and will therefore be interpreted according to its last type. So, define a valuation π into W as follows:

$$x \in \pi(p) \Leftrightarrow \begin{cases} p \in x & : x \in \mathcal{T}_S, \\ p \in \mathbf{t} & : x = (m, \mathbf{t}) \in \mathcal{F}, \\ p \in \mathbf{t}_{i_k} & : x = (\mathbf{t}_{i_0}, R_{i_0}, \mathbf{t}_{i_1}, R_{i_1}, \dots, \mathbf{t}_{i_k}) \in W_S \end{cases}$$

It remains to define the relational structure of our s-model. Intuitively, we start with the relational structure provided by \mathcal{I} and then, for $R \in \mathfrak{P}_1$ and each non-singleton type $\mathbf{t} \in \mathcal{T}_N$ which is not in the range $\text{ran}(\mathcal{F})$ of \mathcal{F} , take $R^{\mathcal{I}} \subseteq \{\mathbf{t}\} \times \mathcal{T}_S$ supplied by (E1). For every non-singleton type $\mathbf{t} \in \text{ran}(\mathcal{F})$, take an $m_{\mathbf{t}}$ with $\mathcal{F}(m_{\mathbf{t}}) = \mathbf{t}$. Define, for $x, y \in W$ and $R \in \mathfrak{P}_1$:

$$x\mathcal{R}y \Leftrightarrow \begin{cases} xR^{\mathcal{I}}y & : x, y \in \mathcal{T}_S \cup \mathcal{F}, \\ \mathbf{t} \rightarrow_R \mathbf{t}_{i_k} & : x = (m, \mathbf{t}), y = (\mathbf{t}_{i_0}, \dots, R_{i_{k-1}}, \mathbf{t}_{i_k}), R = R_{i_{k-1}}, \\ \mathbf{t}_{i_n} \rightarrow_R \mathbf{t}_{i_{n+1}}, & : x = (\dots, \mathbf{t}_{i_n}), y = (\dots, \mathbf{t}_{i_n}, R_{i_n}, \mathbf{t}_{i_{n+1}}), R = R_{i_n}, \\ \mathbf{t}_{i_k} R^{\mathcal{I}} \mathbf{t} & : x = (\mathbf{t}_{i_0}, \dots, R_{i_{k-1}}, \mathbf{t}_{i_k}), y = \mathbf{t} \in \mathcal{T}_S, \mathbf{t}_{i_k} \in \text{ran}(\mathcal{F}), \\ (m_{\mathbf{t}_{i_k}}, \mathbf{t}_{i_k}) R^{\mathcal{I}} \mathbf{t} & : x = (\mathbf{t}_{i_0}, \dots, R_{i_{k-1}}, \mathbf{t}_{i_k}), y = \mathbf{t} \in \mathcal{T}_S, \mathbf{t}_{i_k} \notin \text{ran}(\mathcal{F}). \end{cases}$$

Define, for $x, y \in W$ and $R \in \mathfrak{P}_2$:

$$x\mathcal{R}y \Leftrightarrow \begin{cases} xR^{\mathcal{I}}y & : x, y \in \mathcal{T}_S \cup \mathcal{F}, \\ \exists \psi. \langle R \rangle \psi \in \mathbf{t}_{i_k} & : x = y = (\mathbf{t}_{i_0}, \dots, R_{i_{k-1}}, \mathbf{t}_{i_k}) \in W_S. \end{cases}$$

It is left to the reader to check that $\langle W, \pi, (\mathcal{R} : R \in \mathfrak{P}) \rangle$ is an s-model satisfying φ . \square

Obviously, the conditions listed in Lemma 2 can be checked in exponential time and we have obtained an EXPTIME upper bound for $\mathcal{ML}_{s,t,n}^{-id}$ -satisfiability. The reduction of $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ to $\mathcal{ML}_{s,t,n}^{-id}$ given in Section 3.1 immediately yields an EXPTIME upper bound for the satisfiability of $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ -formulas.

Theorem 3. *For $0 < m < \omega$, satisfiability of $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ -formulas is EXPTIME-complete.*

Note that, for $m = \omega$, satisfiability of $\mathcal{ML}_m^{\neg, \cap, \cup, -, id}$ -formulas is NEXPTIME-complete: In [20], it is proved that satisfiability in $\mathcal{ML}_\omega^{\neg, \cap, \cup}$ is NEXPTIME-hard and the upper bound follows from Theorem 1 and the NEXPTIME upper bound for FO_ω^2 . So, in the modal language, the complexity depends on whether we have a bounded number of accessibility relations or not, while FO^2 does not “feel” this difference.

4 The temporal case

We briefly indicate that the expressive completeness result presented in this paper provides a general framework for comparing the expressivity of modal languages with first-order languages.

Fix a class \mathcal{K} of frames of the form $\mathfrak{F} = \langle W, \mathcal{R}_1, \dots, \mathcal{R}_m \rangle$. Denote by $\mathcal{E}_{\mathfrak{F}}$ the mapping which determines the extension of any complex modal parameter in \mathfrak{F} . A set \mathcal{S} of complex modal parameters over $\{R_1, \dots, R_m, id\}$ is called *exhaustive* for \mathcal{K} if for every complex modal parameter S , such that there exists $\mathfrak{F} \in \mathcal{K}$ with $\mathcal{E}_{\mathfrak{F}}(S) \neq \emptyset$, we find $S_1, \dots, S_k \in \mathcal{S}$ such that $\mathcal{E}_{\mathfrak{F}}(S) = \mathcal{E}_{\mathfrak{F}}(S_1 \cup \dots \cup S_k)$ for all $\mathfrak{F} \in \mathcal{K}$. Denote by $\mathcal{ML}(\mathcal{S})$ the modal language with operators $\langle S \rangle$, $S \in \mathcal{S}$.

Theorem 4. *Let \mathcal{K} be a class of frames and \mathcal{S} a set of complex modal parameters which is exhaustive for \mathcal{K} . Then $\mathcal{ML}(\mathcal{S})$ is expressively complete for the two-variable fragment over \mathcal{K} ; i.e., for every $\varphi \in FO^2$ we find a $\varphi^\sigma \in \mathcal{ML}(\mathcal{S})$ such that for all $\mathcal{M} = \langle W, \pi, \mathcal{R}_1, \dots, \mathcal{R}_m \rangle$ with $\langle W, \mathcal{R}_1, \dots, \mathcal{R}_m \rangle \in \mathcal{K}$ and all $a \in W$:*

$$\mathcal{M}, a \models \varphi^\sigma \Leftrightarrow \mathcal{M}_\sigma \models \varphi(a).$$

Moreover, given φ the formula φ^σ is exponential in the size of φ and can be computed in polynomial time in the size of φ^σ .

The proof of this theorem is similar to the proof of Theorem 1. We provide two examples from temporal logic:

(i) Let \mathcal{K} be a class of strict linear orderings $\langle W, \mathcal{R} \rangle$. Then $\mathcal{S} = \{R, R^-, id\}$ is exhaustive for \mathcal{K} . Hence, $\mathcal{ML}(\mathcal{S})$ is expressively complete for the two-variable fragment over \mathcal{K} . It is not hard to see that any $\mathcal{ML}(\mathcal{S})$ -formula ψ can be translated into an equivalent $\mathcal{ML}(\{R, R^-\})$ -formula ψ' whose length is linear in the length of ψ . In other words, the language of temporal logic with operators ‘always in the future’ and ‘always in the past’ [5, 12] is expressively complete for the two-variable fragment over any class of strict linear orderings.

(ii) Consider again a class \mathcal{K} of strict linear orderings $\langle W, \mathcal{R} \rangle$. Let, for every $\mathfrak{F} = \langle W, \mathcal{R} \rangle \in \mathcal{K}$, $\mathcal{Int}(\mathfrak{F}) = \langle \mathcal{I}(\mathfrak{F}), \mathcal{R}_1, \dots, \mathcal{R}_{13} \rangle$, where $\mathcal{I}(\mathfrak{F})$ is the set of intervals in \mathfrak{F} and $\mathcal{R}_1, \dots, \mathcal{R}_{13}$ is the list of Allen’s relations over $\mathcal{I}(\mathfrak{F})$. $\mathcal{S} = \{R_1, \dots, R_{13}\}$ is exhaustive for $\{\mathcal{Int}(\mathfrak{F}) \mid \mathfrak{F} \in \mathcal{K}\}$ and so $\mathcal{ML}(\mathcal{S})$ is expressively complete for $\{\mathcal{Int}(\mathfrak{F}) \mid \mathfrak{F} \in \mathcal{K}\}$. This interval-based temporal logic was introduced in [16].

Using (i), we obtain the following complexity result for the two-variable fragment interpreted in strict linear orderings:

Theorem 5. *Suppose \mathcal{K} is a class of strict linear orderings such that satisfiability of temporal propositional formulas with operators ‘always in the future’ and ‘always in the past’ in \mathcal{K} is in NP and \mathcal{K} contains an infinite ordering. Then satisfiability of FO^2 with one binary relation interpreted by the strict linear ordering is NEXPTIME-complete.*

Proof. NEXPTIME-hardness follows from the condition that \mathcal{K} contains an infinite structure and that FO^2 without binary relation symbols is NEXPTIME-hard already. Conversely, the following algorithm is in NEXPTIME: given φ , compute φ^σ (in exponential time) and check whether φ^σ is satisfiable in \mathcal{K} . \square

This Theorem applies to e.g. (i) the class of all strict linear orderings, (ii) $\{\langle\mathbb{N}, <\rangle\}$, (iii) $\{\langle\mathbb{Q}, <\rangle\}$, and (iv) $\{\langle\mathbb{R}, <\rangle\}$, see [24, 29].

References

1. H. Andréka, I. Németi, and J. van Benthem. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.
2. P. Blackburn, M. de Rijke, and Y. Venema. Modal logic. In print.
3. E. Börger, E. Grädel, and Yu. Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer, 1997.
4. A. Borgida. On the relative expressiveness of description logics and predicate logics. *Artificial Intelligence*, 82(1 - 2):353–367, 1996.
5. J.P. Burgess. Basic tense logic. In D.M. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume 2, pages 89–133. Reidel, Dordrecht, 1984.
6. Maarten de Rijke. The modal logic of inequality. *The Journal of Symbolic Logic*, 57(2):566–584, June 1992.
7. K. Etesami, M. Vardi, and T. Wilke. First-order logic with two variables and unary temporal logic. In *Proceedings of 12th. IEEE Symp. Logic in Computer Science*, pages 228–235, 1997.
8. M. Fürer. The computational complexity of the unconstrained limited domino problem (with implications for logical decision problems). In *Logic and Machines: Decision problems and complexity*, pages 312–319. Springer, 1984.
9. D. Gabbay, I. Hodkinson, and M. Reynolds. *Temporal Logic: Mathematical Foundations and Computational Aspects, Volume 1*. Oxford University Press, 1994.
10. D.M. Gabbay. Expressive functional completeness in tense logic. In U. Mönnich, editor, *Aspects of Philosophical Logic*, pages 91–117. Reidel, Dordrecht, 1981.
11. G. Gargov and S. Passy. A note on boolean modal logic. In D. Skordev, editor, *Mathematical Logic and Applications*, pages 253–263, New York, 1987. Plenum Press.
12. R.I. Goldblatt. *Logics of Time and Computation*. Number 7 in CSLI Lecture Notes, Stanford. CSLI, 1987.
13. E. Grädel. Why are modal logics so robustly decidable? *Bulletin of the European Association for Theoretical Computer Science*, 68:90–103, 1999.
14. E. Grädel, P. Kolaitis, and M. Vardi. On the Decision Problem for Two-Variable First-Order Logic. *Bulletin of Symbolic Logic*, 3:53–69, 1997.
15. E. Grädel and M. Otto. On Logics with Two Variables. *Theoretical Computer Science*, 224:73–113, 1999.
16. J. Halpern and Y. Shoham. A propositional modal logic of time intervals. *Journal of the ACM*, 38:935–962, 1991.
17. I. L. Humberstone. Inaccessible worlds. *Notre Dame Journal of Formal Logic*, 24(3):346–352, 1983.
18. H. Kamp. *Tense Logic and the Theory of Linear Order*. Ph. D. Thesis, University of California, Los Angeles, 1968.
19. R.E. Ladner. The computational complexity of provability in systems of modal logic. *SIAM Journal on Computing*, 6:467–480, 1977.
20. C. Lutz and U. Sattler. The complexity of reasoning with boolean modal logics. In Frank Wolter, Heinrich Wansing, Maarten de Rijke, and Michael Zakharyashev, editors, *Advances in Modal Logics Volume 3*. CSLI Publications, Stanford, 2001.

21. M. Mortimer. On languages with two variables. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 21:135–140, 1975.
22. Anil Nerode and Richard A. Shore. *Logic for Applications*. Springer Verlag, New York, 1997.
23. D. Scott. A decision method for validity of sentences in two variables. *Journal of Symbolic Logic*, 27(377), 1962.
24. A. Sistla and E. Clarke. The complexity of propositional linear temporal logics. *Journal of the Association for Computing Machinery*, 32:733–749, 1985.
25. E. Spaan. *Complexity of Modal Logics*. PhD thesis, Department of Mathematics and Computer Science, University of Amsterdam, 1993.
26. J. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Napoli, 1983.
27. J. van Benthem. Correspondence theory. In D.M. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume 2, pages 167–247. Reidel, Dordrecht, 1984.
28. M. Vardi. Why is modal logic so robustly decidable? In *Descriptive Complexity and Finite Models*, pages 149–184. AMS, 1997.
29. F. Wolter. Tense logics without tense operators. *Mathematical Logic Quarterly*, 42:145–171, 1996.
30. F. Wolter. Completeness and decidability of tense logics closely related to logics containing $K4$. *Journal of Symbolic Logic*, 62:131–158, 1997.