One-dimensional Fragment of First-order Logic

Lauri Hella [1]

School of Information Sciences
University of Tampere
Finland

Antti Kuusisto [2]

Institute of Computer Science
University of Wrocław
Poland

Abstract

We introduce a novel decidable fragment of first-order logic. The fragment is *one-dimensional* in the sense that quantification is limited to applications of blocks of existential (universal) quantifiers such that at most one variable remains free in the quantified formula. The fragment is closed under Boolean operations, but additional restrictions (called *uniformity conditions*) apply to combinations of atomic formulae with two or more variables. We argue that the notions of *one-dimensionality* and *uniformity* together offer a novel perspective on the *robust decidability* of modal logics. We also show that the one-dimensional fragment is expressively equivalent to a polyadic modal logic with the capacity of permuting and forming Boolean combinations of accessibility relations. Furthermore, we establish that minor modifications to the restrictions of the syntax of the one-dimensional fragment lead to undecidable formalisms. Namely, the *two-dimensional* and *non-uniform one-dimensional* fragments are shown undecidable. Finally, we prove that with regard to expressivity, the one-dimensional fragment is incomparable with both the guarded negation fragment and two-variable logic with counting. Our proof of the decidability of the one-dimensional fragment is based on a technique involving a direct reduction to the monadic class of first-order logic. The novel technique is itself of an independent mathematical interest, and one of the principal contributions of the paper.

Keywords: Extensions of modal logic, fragments of first-order logic, Boolean modal logic, decidability.

[1] The research of Lauri Hella was partially funded by a Professor Pool grant awarded by the Finnish Cultural Foundation.

[2] Antti Kuusisto acknowledges that this work was carried out during a tenure of the ERCIM "Alain Bensoussan" Fellowship Programme, and that the research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement number 246016.

1 Introduction

Decidability questions constitute one of the core themes in computer science logic. Decidability properties of several fragments of first-order logic have been investigated after the completion of the program concerning the classical decision problem. Currently perhaps the most important two frameworks studied in this context are the *guarded fragment* [1] and *two-variable logics*.

Two-variable logic FO^2 was introduced by Henkin in [10] and showed decidable in [14] by Mortimer. The satisfiability and finite satisfiability problems of two-variable logic were proved to be NEXPTIME-complete in [8]. The extension of two-variable logic with counting quantifiers, FOC^2, was shown decidable in [9], [15]. It was subsequently proved to be NEXPTIME-complete in [16].

Research concerning decidability of variants of two-variable logic has been very active in recent years. Recent articles in the field include for example [3] [5], [11], [17], and several others. The recent research efforts have mainly concerned decidability and complexity issues in restriction to particular classes of structures, and also questions related to different built-in features and operators that increase the expressivity of the base language.

Guarded fragment GF was originally conceived in [1]. It is a restriction of first-order logic that only allows quantification of "guarded" new variables—a restriction that makes the logic rather similar to modal logic.

The guarded fragment has generated a vast literature, and several related decidability questions have been studied. The fragment has recently been significantly generalized in [2]. The article introduces the *guarded negation first-order logic* GNFO. This logic only allows negations of formulae that are guarded in the sense of the guarded fragment. The guarded negation fragment has been shown complete for 2NEXPTIME in [2].

Two-variable logic and guarded-fragment are examples of decidable fragments of first-order logic that are not based on restricting the quantifier alternation patterns of formulae, unlike the prefix classes studied in the context of the classical decision problem. Surprisingly, not many such frameworks have been investigated in the literature.

In this paper we introduce a novel decidable fragment that essentially allows arbitrary quantifier alternation patterns. The *uniform one-dimensional fragment* UF_1 of first-order logic is obtained by restricting quantification to blocks of existential (universal) quantifiers that *leave at most one free variable* in the resulting formula. Additionally, a *uniformity condition* applies to the use of atomic formulae: if $n, k \geq 2$, then a Boolean combination of atoms $R(x_1, ..., x_k)$ and $S(y_1, ..., y_n)$ is allowed only if $\{x_1, ..., x_k\} = \{y_1, ..., y_n\}$. Boolean combinations of formulae with at most one free variable can be formed freely.

We establish decidability of the satisfiability and finite satisfiability problems of UF_1. We also show that if the uniformity condition is lifted, we obtain an undecidable logic. Furthermore, if we keep uniformity but go two-dimensional by allowing existential (universal) quantifier blocks that leave two variables free, we again obtain an undecidable formalism. Therefore, *if we lift either of the two restrictions that our fragment is based on, we obtain an* *undecidable logic.*

In addition to studying decidability, we also show that UF_1 is incomparable in expressive power with both FOC^2 and GNFO.

In [18], Vardi initiated an intriguing research effort that aims to understand phenomena behind the *robust decidability* of different variants of modal logic. In addition to [18], see also for example [7] and the introduction of [2]. Modal logic indeed has several features related to what is known about decidability. In particular, modal logic embeds into both FO^2 and GF.

However, there exist several important and widely applied decidable extensions of modal logic that do not embed into *both* FO^2 and GF. Such extensions include *Boolean modal logic* (see [6], [13]) and basic *polyadic modal logic*, i.e, modal logic containing accessibility relations of arities higher than two (see [4]). Boolean modal logic allows Boolean combinations of accessibility relations and therefore can express for example the formula $\exists y\big(\neg R(x,y) \wedge P(y)\big)$. Polyadic modal logic can express the formula $\exists x_2 ... \exists x_k \big(R(x_1, ..., x_k) \wedge P(x_2) \wedge ... \wedge P(x_k)\big)$. Boolean modal logic and polyadic modal logic are both inherently one-dimensional, and furthermore, satisfy the uniformity condition of UF_1. Both logics embed into UF_1. The notions of *one-dimensionality* and *uniformity* can be seen as novel features that can help, in part, *explain* decidability phenomena concerning modal logics.

Importantly, also the equality-free fragment of FO^2 embeds into UF_1. In fact, when attention is restricted to vocabularies with relations of arities at most two, the expressivities of UF_1 and the equality-free fragment of FO^2 coincide. Instead of seeing this as a weakness of UF_1, we in fact regard UF_1 as *a canonical generalisation* of (equality-free) FO^2 into contexts with arbitrary relational vocabularies. The fragment UF_1 can be regarded as a *vectorisation* of FO^2 that offers new possibilities for extending research efforts concerning two-variable logics. *It is worth noting that for example in database theory contexts, two-variable logics as such are not always directly applicable due to the arity-related limitations.* Thus we believe that the one-dimensional fragment is indeed a worthy discovery that *extends the scope of research on two-variable logics to the realm involving relations of arbitrary arities.*

Instead of extending basic techniques from the field of two-variable logic, our decidability proof is based on a direct satisfiability preserving translation of UF_1 into monadic first-order logic. The novel proof technique is mathematically interesting in its own right, and is in fact a central contribution of this article; the proof technique is clearly robust and can be modified and extended to give other decidability results. Furthermore, as a by-product of our proof, we identify a natural polyadic modal logic MUF_1, which is expressively equivalent to the one-dimensional fragment. This *modal normal form* for the one-dimensional fragments is also—we believe—a nice contribution.

2 Preliminaries

Let $\mathbb{Z}_+$ denote the set of positive integers. Let $\mathcal{T}$ denote a *complete relational vocabulary*, i.e., $\mathcal{T} := \bigcup_{k \in \mathbb{Z}_+} \tau_k$, where τ_k denotes a countably infinite set of

k-ary relation symbols. Each vocabulary τ we consider below is assumed to be a subset of $\mathcal{T}$. A τ-formula of first-order logic is a formula whose set of non-logical symbols is a subset of τ. A τ-model is a model whose set of interpreted non-logical symbols is τ.

Let VAR denote the countably infinite set $\{ x_i \mid i \in \mathbb{Z}_+ \}$ of *variable symbols*. We define the set of $\mathcal{T}$-formulae of first-order logic in the usual way, assuming that all variable symbols are from VAR. Below we use *meta-variables* x, y, z in order to denote variables in VAR. Also symbols of the type y_i and z_i, where $i \in \mathbb{Z}_+$, will be used as meta-variables. In addition to meta-variables, we also need to directly use the variables $x_i \in \text{VAR}$ below. Note that for example the meta-variables y_1 and y_2 may denote the same variable in VAR, while the variables $x_1, x_2 \in \text{VAR}$ of course simply *are* different variables.

Let R be a k-ary relation symbol, $k \in \mathbb{Z}_+$. An atomic formula $R(y_1, ..., y_k)$ is called *m-ary* if there are exactly m distinct variables in the set $\{y_1, ..., y_k\}$. For example, if x, y are distinct variables, then $S(x, y)$ and $T(y, x, y, y)$ are binary, and $U(x_1, x_6, x_3, x_2, x_1, x_6)$ is 4-ary. An *m-ary τ-atom* is an atomic formula that is m-ary, and the relation symbol of the formula is in τ.

Let $\tau \subseteq \mathcal{T}$. Let $\mathfrak{M}$ a τ-model with the domain M. A function f that maps some subset of VAR into M is an *assignment*. Let φ be a τ-formula with the free variables $y_1, ..., y_k$. Let f be an assignment that interprets the free variables of φ in M. We write $\mathfrak{M}, f \models \varphi$ if $\mathfrak{M}$ satisfies φ when the free variables of φ are interpreted according to f. Let $u_1, ..., u_k \in M$. Let φ be a τ-formula whose free variables are among $y_1, ..., y_k$. We write $\mathfrak{M}, \frac{(u_1, ..., u_k)}{(y_1, ..., y_k)} \models \varphi$ if $\mathfrak{M}, f \models \varphi$ for some assignment f such that $f(y_i) = u_i$ for each $i \in \{1, .., k\}$.

By a *non-empty conjunction* we mean a finite conjunction with at least one conjunct; for example $R(x, y) \wedge \exists y P(y)$ and $\top$ are non-empty conjunctions.

By *monadic first-order logic*, or MFO, we mean the fragment of first-order logic *without equality*, where formulae contain only unary relation symbols.

Let $k \in \mathbb{Z}_+$. A *k-permutation* is a bijection $\sigma : \{1, ..., k\} \to \{1, ..., k\}$. When k is irrelevant or clear from the context, we simply talk about permutations.

Let $k \in \mathbb{Z}_+$. We let $(u, ..., u)_k$ and u^k denote the k-tuple containing k copies of the object u. When $k = 1$, this tuple is identified with the object u.

Let l and $k \leq l$ be positive integers. Let K be a set, and let $(s_1, ..., s_l) \in K^l$ be a tuple. We let $(s_1, ..., s_l) \upharpoonright k$ denote the tuple $(s_1, ..., s_k)$. Let $R \subseteq K^l$ be an l-ary relation. We let $R \upharpoonright k$ denote the k-ary relation $R' \subseteq K^k$ defined such that for each $(s_1, ..., s_k) \in K^k$, we have $(s_1, ..., s_k) \in R'$ iff $(s_1, ..., s_k) = (u_1, ..., u_l) \upharpoonright k$ for some tuple $(u_1, ..., u_l) \in R$.

Recall that $\bigwedge \emptyset$ is assumed to be always true, while $\bigvee \emptyset$ is always false.

3 The one-dimensional fragment

We shall next define the *uniform one-dimensional* fragment UF_1 of first-order logic. Let $Y = \{y_1, ..., y_n\}$ be a set of variable symbols, and let R be a k-ary relation symbol. An atomic formula $R(y_{i_1}, ..., y_{i_k})$ is called a *Y-atom* if $\{y_{i_1}, ..., y_{i_k}\} = Y$. A finite set of Y-atoms is called a *Y-uniform set*. When Y is irrelevant or known from the context, we may simply talk about a *uniform set*.

For example, assuming that x, y, z are distinct variables, $\{T(x,y), S(y,x)\}$ and $\{R(x,x,y), R(y,y,x), S(y,x)\}$ are uniform sets, while $\{R(x,y,z), R(x,y,y)\}$ is not. The empty set is a $\emptyset$-uniform set.

Let $\tau \subseteq \mathcal{T}$. The set $\mathrm{UF}_1(\tau)$, or the set of τ-formulae of the one-dimensional fragment, is the smallest set $\mathcal{F}$ satisfying the following conditions.

(i) Every unary τ-atom is in $\mathcal{F}$, and $\bot, \top \in \mathcal{F}$.

(ii) If $\varphi \in \mathcal{F}$, then $\neg\varphi \in \mathcal{F}$. If $\varphi_1, \varphi_2 \in \mathcal{F}$, then $(\varphi_1 \wedge \varphi_2) \in \mathcal{F}$.

(iii) Let $Y = \{y_1, ..., y_k\}$ be a set of variable symbols. Let U be a finite set of formulae $\psi \in \mathcal{F}$ whose free variables are in Y. Let $V \subseteq Y$. Let F be a V-uniform set of τ-atoms. Let φ be any Boolean combination of formulae in $U \cup F$. Then $\exists y_2 ... \exists y_k\, \varphi \in \mathcal{F}$.

(iv) If $\varphi \in \mathcal{F}$, then $\exists y\, \varphi \in \mathcal{F}$.

Notice that there is no equality symbol in the language. Notice also that the formation rule (iv) is strictly speaking not needed since the rule (iii) covers it. Concerning the rule (i), notice that also atoms of the type $S(x, ..., x)_k$, where $k \neq 1$, are legitimate formulae. Let UF_1 denote the set $\mathrm{UF}_1(\mathcal{T})$.

3.1 Intuitions underlying the decidability proof

We show decidability of the satisfiability and finite satisfiability problems of UF_1 by translating UF_1-formulae into equisatisfiable MFO-formulae. We first translate UF_1 into a logic DUF_1. This logic is a normal form for UF_1 such that all literals of arities higher than one appear in simple conjunctions, as for example in the formula $\exists y \exists z \big(R(x,z,y,z) \wedge \neg S(y,x,z) \wedge \varphi(y)\big)$. The logic DUF_1 is then translated into a modal logic MUF_1, which is an essentially variable-free formalism for DUF_1. In Section 4 we show how formulae of the logic MUF_1 are translated into equisatisfiable formulae of MFO, which is well-known to have the finite model property.

The semantics of MUF_1 is defined (see Section 3.4) with respect to pointed models $(\mathfrak{M}, u)$, where $u \in M = Dom(\mathfrak{M})$. If φ is a formula of MUF_1, we let $\|\varphi\|^{\mathfrak{M}}$ denote the set $\{\, v \in M \mid (\mathfrak{M}, v) \models \varphi \,\}$. In Section 4 we fix a MUF_1-formula ψ and translate it to an MFO-formula $\psi^*(x)$. We prove that if $(\mathfrak{M}, v) \models \psi$, then $\psi^*(x)$ is satisfied in a model $\mathfrak{T}$, whose domain is $M \times T$, where T is the domain of an *m-dimensional hypertorus of arity l*. Such a hypertorus is a structure $(T, R_1, ..., R_m)$, where the m different relations R_i are all l-ary. Intuitively, the domain of $\mathfrak{T}$ consists of several copies of M, one copy for each point of the hypertorus. Let SUB_ψ denote the set of subformulae of ψ. The vocabulary of $\mathfrak{T}$ consists of monadic predicates P_α and P_t, where $\alpha \in \mathrm{SUB}_\psi$ and $t \in T$. The predicates are interpreted such that $P_\alpha^{\mathfrak{T}} := \|\alpha\|^{\mathfrak{M}} \times T$ and $P_t^{\mathfrak{T}} := M \times \{t\}$.

We will give a *rigorous and self-contained* proof of the decidability of UF_1, but to get an (admittedly very rough) initial idea of some of the related background intuitions, consider the following construction. (*It may also help to refer back to this section while internalizing the proof.*)

Consider a formula of ordinary unimodal logic φ and a Kripke model $\mathfrak{N}$.

We can *maximize* the accessibility relation R of $\mathfrak{N}$ by defining a new relation $S \subseteq N \times N$ such that $(u, v) \in S$ iff for all formulae $\Diamond\beta \in \mathrm{SUB}_\varphi$, we have

$$(\mathfrak{N}, v) \models \beta \Rightarrow (\mathfrak{N}, u) \models \Diamond\beta. \tag{1}$$

If we replace R by S in $\mathfrak{N}$, then each point w in the new model will satisfy exactly the same subformulae of φ as w satisfied in the old model. Thus we can *encode information concerning* R by using the (so-called filtration) condition given by Equation 1. The equation talks about the *sets* $\|\beta\|^{\mathfrak{N}}$ and $\|\Diamond\beta\|^{\mathfrak{N}}$, and thus it turns out that we can encode the information given by the equation by *monadic predicates* P_β and $P_{\Diamond\beta}$ corresponding to the sets $\|\beta\|^{\mathfrak{N}}$ and $\|\Diamond\beta\|^{\mathfrak{N}}$ (cf. the formulae $PreCons_\delta$ and $Cons_\delta$ in Section 4.1). *This way we can encode information concerning accessibility relations by using formulae of* MFO.

This construction does not work if one tries to maximize *both* a binary relation R *and its complement* $\overline{R}$ at the same time: the problem is that the maximized relations S and $\overline{S}$ will not necessarily be complements of each other. For this reason we need to *make enough room* for maximizing accessibility relations. Below we will simultaneously maximize several types of accessibility relations that cannot be allowed to intersect. Thus we need to use an *n-dimensional* hypertorus (rather than a usual 2D torus). Each k-ary accessibility relation type δ of the translated MUF_1-formula will be reserved a sequence $\overline{r} := (M \times \{t_1\}, ..., M \times \{t_k\})$ of copies of M from the domain of $\mathfrak{T}$. Information concerning δ will be encoded into this sequence $\overline{r}$ of models.

3.2 Diagrams

Let $\tau \subseteq \mathcal{T}$ be a *finite* vocabulary. Let $k \geq 2$ be an integer, and let $Y = \{y_1, ..., y_k\}$ be a set of *distinct* variable symbols. A *uniform k-ary τ-diagram* is a maximal satisfiable set of Y-atoms and negated Y-atoms of the vocabulary τ. (The empty set is *not* considered to be a uniform k-ary τ-diagram; this case is relevant when τ contains no relation symbols of the arity k or higher.)

For example, let $\tau = \{P, R, S\}$, where the arities of P, R, S are 1, 2, 3, respectively. Now $\{R(x, y), \neg R(y, x), S(y, x, x), S(x, y, x), \neg S(x, x, y), S(x, y, y), \neg S(y, x, y), S(y, y, x)\}$ is a uniform binary τ-diagram. Here we assume that x and y are distinct variables.

Let $\tau \subseteq \mathcal{T}$ be a *finite* vocabulary. The set $\mathrm{DUF}_1(\tau)$ is the smallest set $\mathcal{F}$ satisfying the following conditions.

(i) Every unary τ-atom is in $\mathcal{F}$. Also $\bot, \top \in \mathcal{F}$.

(ii) If $\varphi \in \mathcal{F}$, then $\neg\varphi \in \mathcal{F}$. If $\varphi_1, \varphi_2 \in \mathcal{F}$, then $(\varphi_1 \wedge \varphi_2) \in \mathcal{F}$.

(iii) Let δ be a uniform k-ary τ-diagram in the variables y_1,...,y_k, where $k \geq 2$. Let φ be a non-empty conjunction of a finite set U of formulae in $\mathcal{F}$ whose free variables are among y_1,...,y_k. Then $\exists y_2 ... \exists y_k \left(\bigwedge \delta \wedge \varphi \right) \in \mathcal{F}$.

(iv) If $\varphi \in \mathcal{F}$ has at most one free variable, y, then $\exists y\, \varphi \in \mathcal{F}$.

Let DUF_1 denote the set of formulae φ such that for some finite $\tau \subseteq \mathcal{T}$, we have $\varphi \in \mathrm{DUF}_1(\tau)$. UF_1 translates effectively into DUF_1; see the appendix for the proof. Here we briefly *sketch* the principal idea behind the translation.

Consider a UF_1-formula $\exists \overline{y}\, \psi$, where $\overline{y}$ is a tuple of variables. Put ψ into disjunctive normal form $\psi_1 \vee ... \vee \psi_k$. Thus $\exists \overline{y}\, \psi$ translates into the formula $\exists \overline{y}\, \psi_1 \vee ... \vee \exists \overline{y}\, \psi_k$, where each ψ_i is a conjunction. Each ψ_i is equivalent to a disjunction $\psi_{i,1} \vee ... \vee \psi_{i,m}$, where $\psi_{i,j}$ is of the desired type $\big(\bigwedge \delta \wedge \varphi\big)$.

3.3 Hypertori

We next define a class of hypertori. It may help to have a look at Lemma 3.1 before internalizing the definition. Let $l \geq 2$ and $n \geq 2$ be integers. Define $T := \{1, ..., n\} \times \{1, ..., l\} \times \{0, 1, 2\}$. Let $(t_1, ..., t_l) \in T^l$ be a tuple. Let $t_1 = (m, m', m'')$. Let $j \in \{1, ..., n\}$. A tuple $(t_1, ..., t_l) \in T^l$, where each $t_i = (p, q_i, r)$, is the *j-th good l-ary sequence originating from t_1*, if for each $i \in \{2, ..., l\}$, the following conditions hold.

(i) $p - m \equiv j - 1 \mod n$.

(ii) $q_i - m' \equiv i - 1 \mod l$.

(iii) $r - m'' \equiv 1 \mod 3$.

Define the relation $R_j \subseteq T^l$ such that $(s_1, ..., s_l) \in R_j$ iff $(s_1, ..., s_l)$ is the j-th good l-ary sequence originating from s_1. The structure $\big(T, R_1, ..., R_n\big)$ is the *n-dimensional hypertorus of the arity l*. The following lemma is easy to prove.

Lemma 3.1 *Let $\big(T, R_1, ..., R_n\big)$ be an n-dimensional hypertorus of the arity l. Let $j \in \{1, ..., n\}$ and $k \in \{2, ..., l\}$. Then the following conditions hold.*

(i) *For each $t \in T$, there exists exactly one tuple $(s_1, ..., s_k) \in R_j \upharpoonright k$ such that $t = s_1$. We have $s_i \neq s_j$ for all $i, j \in \{1, ..., k\}$ such that $i \neq j$.*

(ii) *Let $(s_1, ..., s_k) \in R_j \upharpoonright k$. Let σ be a k-permutation, and let $i \in \{1, ..., n\} \setminus \{j\}$. Then $(s_{\sigma(1)}, ..., s_{\sigma(k)}) \notin R_i \upharpoonright k$.*

(iii) *Let $(s_1, ..., s_k) \in R_j \upharpoonright k$. Let μ be any k-permutation other than the identity permutation. Then $(s_{\mu(1)}, ..., s_{\mu(k)}) \notin R_j \upharpoonright k$.*

In the rest of the article, we let $\mathfrak{T}(n, l)$ denote the n-dimensional hypertorus of the arity l. We let $T(n, l)$ and $R_j(n, l)$ denote, respectively, the domain and the relation R_j of $\mathfrak{T}(n, l)$.

3.4 Translation into a modal logic

Let $\tau \subseteq \mathcal{T}$ be a *finite* vocabulary, and let $k \geq 2$ be an integer. Let $\mathfrak{M}$ be a τ-model with the domain M. Let δ be a uniform k-ary τ-diagram in the variables $x_1, ..., x_k$. Notice that here we use the standard variables $x_1, ..., x_k$ from VAR. The diagram δ is a *standard uniform k-ary τ-diagram.* We define $\|\delta\|^{\mathfrak{M}}$ to be the relation $\{\, (u_1, ..., u_k) \in M^k \mid \mathfrak{M}, \frac{(u_1, ..., u_k)}{(x_1, ..., x_k)} \models \bigwedge \delta \,\}$. Standard variables are needed in order to uniquely specify the *order of elements* in tuples of $\|\delta\|^{\mathfrak{M}}$.

Let δ be a standard uniform k-ary τ-diagram. Let $q \leq k$ be a positive integer. Let $t : \{1, ..., k\} \to \{1, ..., q\}$ be a surjection. We let δ/t denote the set obtained from δ by replacing each variable x_i by $x_{t(i)}$.

Let k and q be positive integers such that $2 \leq q \leq k$. Let η and δ be standard uniform q-ary and k-ary τ-diagrams, respectively. Let $f : \{1, ..., k\} \to \{1, ..., q\}$

be a surjection. Assume that $\bigwedge \eta \models \bigwedge \delta/f$, i.e., the implication $\mathfrak{M}, h \models \eta \Rightarrow \mathfrak{M}, h \models \delta/f$ holds for each τ-model $\mathfrak{M}$ and each assignment h interpreting the variables $x_1, ..., x_q$ in the domain of $\mathfrak{M}$. Then we write $\eta \leq_f \delta$.

We then define a modal logic that provides an essentially variable-free representation of UF_1. Define the set $\mathrm{MUF}_1(\tau)$ to be the smallest set $\mathcal{F}$ such that the following conditions are satisfied.

(i) If $S \in \tau$ is a relation symbol of any arity, then $S \in \mathcal{F}$. Also $\bot, \top \in \mathcal{F}$.

(ii) If $\varphi \in \mathcal{F}$, then $\neg\varphi \in \mathcal{F}$. If $\varphi_1, \varphi_2 \in \mathcal{F}$, then $(\varphi_1 \wedge \varphi_2) \in \mathcal{F}$.

(iii) If $\varphi_1, ..., \varphi_k \in \mathcal{F}$ and δ is a standard uniform k-ary τ-diagram, then $\langle\delta\rangle(\varphi_1, ..., \varphi_k) \in \mathcal{F}$.

(iv) If $\varphi \in \mathcal{F}$, then $\langle E\rangle\varphi \in \mathcal{F}$. (Here $\langle E\rangle$ denotes the *universal modality*; see below for the semantics.)

The semantics of $\mathrm{MUF}_1(\tau)$ is defined with respect to *pointed σ-models* $(\mathfrak{M}, w)$, where $\mathfrak{M}$ is an ordinary σ-model of predicate logic for some vocabulary $\sigma \supseteq \tau$, and w is an element of the domain M of $\mathfrak{M}$. Obviously we define that $(\mathfrak{M}, w) \models \top$ always holds, and that $(\mathfrak{M}, w) \models \bot$ never holds. Let $S \in \tau$ be an n-ary relation symbol. We define $(\mathfrak{M}, w) \models S \Leftrightarrow w^n \in S^{\mathfrak{M}}$, where $S^{\mathfrak{M}}$ is the interpretation of the relation symbol S in the model $\mathfrak{M}$. The Boolean connectives $\neg$ and $\wedge$ have their usual meaning. For formulae of the type $\langle\delta\rangle(\varphi_1, ..., \varphi_k)$, we define that $(\mathfrak{M}, w) \models \langle\delta\rangle(\varphi_1, ..., \varphi_k)$ if and only if there exists a tuple $(u_1, ..., u_k) \in \|\delta\|^{\mathfrak{M}}$ such that $u_1 = w$ and $(\mathfrak{M}, u_i) \models \varphi_i$ for each $i \in \{1, ..., k\}$. For formulae $\langle E\rangle\varphi$, we define $(\mathfrak{M}, w) \models \langle E\rangle\varphi$ if and only if there exists some $u \in M$ such that $(\mathfrak{M}, u) \models \varphi$.

When φ is a $\mathrm{MUF}_1(\tau)$-formula and $\mathfrak{M}$ a σ-model with the domain M, we let $\|\varphi\|^{\mathfrak{M}}$ denote the set $\{\, u \in M \mid (\mathfrak{M}, u) \models \varphi \,\}$. We let MUF_1 denote the union of all sets $\mathrm{MUF}_1(\tau)$, where τ is a finite subset of $\mathcal{T}$.

It is very easy to show that there is an effective translation that turns any formula $\gamma(x) \in \mathrm{DUF}_1$ into a formula $\chi \in \mathrm{MUF}_1$ such that $(\mathfrak{M}, w) \models \chi \Leftrightarrow \mathfrak{M}, \frac{w}{x} \models \gamma(x)$ for all τ-models $\mathfrak{M}$, where τ is the set of non-logical symbols in $\gamma(x)$. (The set of non-logical symbols in χ is contained in τ, and the formula $\gamma(x)$ can either be a sentence or have the free variable x.)

4 UF_1 is decidable

Let us *fix* a formula ψ of MUF_1. We will first define a translation of ψ to an MFO-formula $\psi^*(x)$ in Section 4.1. We will then show in Sections 4.2 and 4.3 that the translation indeed preserves equivalence of satisfiability over finite models as well as over all models. Due to the above effective translations from UF_1 to DUF_1 and from DUF_1 to MUF_1, this implies that the satisfiability and finite satisfiability problems of UF_1 are decidable.

4.1 Translating MUF_1 into monadic first-order logic

We assume, w.l.o.g., that ψ contains at least one subformula of the type $\langle\delta\rangle(\chi_1, \chi_2)$. If not, we redefine ψ. The vocabulary of ψ may of course grow. We also assume, w.l.o.g., that ψ does not contain occurrences of the symbols $\top$,

$\bot$. Furthermore, we assume, w.l.o.g., that if R is a relation symbol occurring in some diagram of ψ, then $\neg R$ also occurs in ψ as a subformula: we can of course always add the conjunct $R \vee \neg R$ to ψ.

Let V_ψ be the set of all relation symbols in ψ, whether they occur in diagrams or as atomic subformulae; in fact, due to our assumptions above, the set of atomic formulae in ψ is equal to V_ψ. Let D_ψ be the set of relation symbols occurring in the diagrams of ψ. Let $V_\psi(k)$ denote the set of k-ary relation symbols in V_ψ. Define $D_\psi(k)$ analogously. Due to the assumption that ψ contains a subformula $\langle\delta\rangle(\chi_1, \chi_2)$, each relation symbol of some arity $m \geq 2$ that occurs as an atom in ψ, also occurs in the diagram δ. (This is due to the definition of MUF_1.) Thus $V_\psi(n) = D_\psi(n)$ for all $n > 1$.

Let $\mathcal{M}$ denote the maximum arity of all diagrams in ψ. For each $k \in \{2, ..., \mathcal{M}\}$, let Δ_k denote the set of exactly all standard uniform k-ary V_ψ-diagrams. Let Δ denote the union of the sets Δ_k, where $k \in \{2, ..., \mathcal{M}\}$. Let $\mathcal{N} := max\{\ |\Delta_k| \mid k \in \{2, ..., \mathcal{M}\}\ \}$. Recall that $T(\mathcal{N}, \mathcal{M})$ denotes the domain of the $\mathcal{N}$-dimensional hypertorus of the arity $\mathcal{M}$. For each $k \in \{2, ..., \mathcal{M}\}$, define an injection $b_k : \Delta_k \longrightarrow \{\ R_1(\mathcal{N}, \mathcal{M}), ..., R_{\mathcal{N}}(\mathcal{N}, \mathcal{M})\ \}$. For a k-ary diagram $\delta \in \Delta_k$, let T_δ denote the k-ary relation $\big(b_k(\delta)\big) \restriction k$.

Let SUB_ψ denote the set of subformulae of the formula ψ. Fix fresh unary relation symbols P_α and P_t for each formula $\alpha \in \mathrm{SUB}_\psi$ and torus point $t \in T(\mathcal{N}, \mathcal{M})$. The vocabulary of the translation $\psi^*(x)$ of ψ will be the set $\{\ P_\alpha \mid \alpha \in \mathrm{SUB}_\psi\ \} \cup \{\ P_t \mid t \in T(\mathcal{N}, \mathcal{M})\ \}$. We let V^* denote this set.

We shall next define a collection of auxiliary formulae needed in order to define $\psi^*(x)$. If a pointed model $(\mathfrak{M}, u)$ satisfies ψ, then $\psi^*(x)$ will be satisfied in a larger model; the related model construction is defined in the beginning of Section 4.2. The predicates of the type P_α will be used to encode information about sets $\|\alpha\|^{\mathfrak{M}}$, while the predicates P_t encode information about the *diagrams* of ψ. The predicates P_t are crucial when defining a V_ψ-model $\mathfrak{B}$ that satisfies ψ based on a V^*-model $\mathfrak{A}$ of $\psi^*(x)$ in Section 4.3.

Let $\delta \in \Delta_k$. Define $PreCons_\delta(x_1, ..., x_k)$ to be the formula

$$\bigwedge_{\langle\delta\rangle(\chi_1,...,\chi_k)\ \in\ \mathrm{SUB}_\psi} \Big(P_{\chi_1}(x_1) \wedge ... \wedge P_{\chi_k}(x_k) \rightarrow\ P_{\langle\delta\rangle(\chi_1\,,\,...\,,\,\chi_k)}(x_1)\Big).$$

Let $\Delta(\delta)$ be the set of pairs (η, f), where $\eta \in \Delta$ is a p-ary diagram for some $p \geq k$, and $f : \{1, ..., p\} \rightarrow \{1, ..., k\}$ is a surjection such that we have $\delta \leq_f \eta$. The set $\Delta(\delta)$ is the set of *inverse projections of* δ *in* Δ. Define

$$Cons_\delta(x_1, ..., x_k)\ :=\ \bigwedge_{(\eta,f)\,\in\,\Delta(\delta)} PreCons_\eta(x_{f(1)}, ..., x_{f(p)}).$$

The following formula is the principal formula that encodes information about diagrams of δ (cf. Lemma 4.1).

$$Diag_\delta(x_1, ..., x_k) := \bigvee_{(t_1\,,\,...\,,\,t_k)\ \in\ T_\delta} P_{t_1}(x_1) \wedge ... \wedge P_{t_k}(x_k)\ \wedge\ Cons_\delta(x_1, ..., x_k).$$

Let $+(\delta)$ denote the set of relation symbols R that occur positively in δ, i.e., there exists some atom $R(y_1, ..., y_n) \in \delta$, where n is the arity of R. Let $-(\delta)$ be the relation symbols R that occur negatively in δ, i.e., $\neg R(y_1, ..., y_n) \in \delta$ for some atom $R(y_1, ..., y_n)$. The following three formulae encode information about atomic formulae in ψ. Define

$$\begin{array}{rcl} \mathit{Local}_\delta(x) & := & \bigwedge\limits_{R \,\in\, +(\delta)} P_R(x) \;\wedge \bigwedge\limits_{R \,\in\, -(\delta)} P_{\neg R}(x), \\ \mathit{LocalDiag}_\delta(x) & := & \mathit{Local}_\delta(x) \;\rightarrow\; \mathit{PreCons}_\delta(x, ..., x)_k, \\ \psi_{local} & := & \bigwedge\limits_{\delta \,\in\, \Delta} \forall x\, \mathit{LocalDiag}_\delta(x). \end{array}$$

The next formula is essential in the construction of a V_ψ-model of ψ from a V^*-model of $\psi^*(x)$ in Section 4.3. The two models have the same domain. The formula states that each tuple can be interpreted to satisfy *some* diagram δ such that information concerning the unary predicates in V^* is consistent with δ. See the way $\mathfrak{B}$ is defined based on $\mathfrak{A}$ in Section 4.3 for further details. Define

$$\psi_{total} := \bigwedge_{k \,\in\, \{2,...,\mathcal{M}\}} \forall x_1 ... \forall x_k \bigvee_{\delta \,\in\, \Delta_k} \mathit{Cons}_\delta(x_1, ..., x_k).$$

Also the following formula is crucial for the definition of $\mathfrak{B}$.

$$\psi_{uniq} := \bigwedge_{t,\, s \;\in\; T(\mathcal{N},\mathcal{M}),\; t \neq s} \neg\exists x\big(\, P_t(x) \;\wedge\; P_s(x)\,\big).$$

Let $\neg\alpha$, $(\beta \wedge \gamma)$, $\langle E \rangle \chi$, and $\langle \delta \rangle(\chi_1, ..., \chi_k)$ be formulae in SUB_ψ. The following formulae recursively encode information concerning subformulae of ψ. Define

$$\begin{array}{rcl} \psi_{\neg\alpha} & := & \forall x\Big(P_{\neg\alpha}(x) \leftrightarrow \neg P_\alpha(x)\Big), \\ \psi_{(\beta\wedge\gamma)} & := & \forall x\Big(P_{(\beta\wedge\gamma)}(x) \leftrightarrow \big(P_\beta(x) \wedge P_\gamma(x)\big)\Big), \\ \psi_{\langle E \rangle \chi} & := & \forall x\Big(P_{\langle E \rangle \chi}(x) \leftrightarrow \exists y P_\chi(y)\Big), \\ \psi_{\langle \delta \rangle(\chi_1,...,\chi_k)} & := & \forall x_1\Big(P_{\langle \delta \rangle(\chi_1,...,\chi_k)}(x_1) \\ & & \quad \leftrightarrow \exists x_2 ... x_k\big(\mathit{Diag}_\delta(x_1, ..., x_k) \\ & & \qquad \wedge P_{\chi_1}(x_1) \wedge ... \wedge P_{\chi_k}(x_k)\,\big)\Big). \end{array}$$

Let $\psi_{sub} := \bigwedge\limits_{\alpha \,\in\, \mathrm{SUB}_\psi} \psi_\alpha$. Finally, we define

$$\psi^*(x) \;:=\; \psi_{total} \;\wedge\; \psi_{uniq} \;\wedge\; \psi_{local} \;\wedge\; \psi_{sub} \;\;\wedge\; P_\psi(x).$$

4.2 Satisfiability of ψ implies satisfiability of $\psi^*(x)$

Fix an arbitrary model V_ψ-model $\mathfrak{M}$ with the domain M. Fix a point $w \in M$. Assume $(\mathfrak{M}, w) \models \psi$. We shall next construct a model $\mathfrak{T}$ with the domain $M \times T(\mathcal{N}, \mathcal{M})$. We then show that $\mathfrak{T}, \frac{(w,t)}{x} \models \psi^*(x)$, where t is a torus point. If $\mathfrak{M}$ is a finite model, then so is $\mathfrak{T}$.

The domain $M \times T(\mathcal{N}, \mathcal{M})$ of the V^*-model $\mathfrak{T}$ consists of copies of M, one copy for each torus point $t \in T(\mathcal{N}, \mathcal{M})$. Let us define interpretations of the symbols in V^*. Consider a symbol P_α, where $\alpha \in \mathrm{SUB}_\psi$. If $(u,t) \in Dom(\mathfrak{T})$, then $(u,t) \in P_\alpha^{\mathfrak{T}} \Leftrightarrow u \in \|\alpha\|^{\mathfrak{M}}$. Consider then a symbol P_t, where $t \in T(\mathcal{N}, \mathcal{M})$. If $(u,t') \in Dom(\mathfrak{T})$, then $(u,t') \in P_t^{\mathfrak{T}} \Leftrightarrow t' = t$.

Lemma 4.1 *Let $\langle\delta\rangle(\chi_1, ..., \chi_k) \in \mathrm{SUB}_\psi$ and $(u,t) \in Dom(\mathfrak{T})$. Then $(\mathfrak{M}, u) \models \langle\delta\rangle(\chi_1, ..., \chi_k)$ iff $\mathfrak{T}, \frac{(u,t)}{x_1} \models \exists x_2 ... \exists x_k \big(Diag_\delta(x_1, ..., x_k) \wedge P_{\chi_1}(x_1) \wedge ... \wedge P_{\chi_k}(x_k) \big)$.*

Proof. Define $u_1 := u$ and $t_1 := t$. Assume $(\mathfrak{M}, u_1) \models \langle\delta\rangle(\chi_1, ..., \chi_k)$. Thus $(u_1, ..., u_k) \in \|\delta\|^{\mathfrak{M}}$ for some tuple $(u_1, ..., u_k)$ such that $u_i \in \|\chi_i\|^{\mathfrak{M}}$ for each i. Hence $(u_i, s) \in P_{\chi_i}^{\mathfrak{T}}$ for each i and each torus point s. To conclude the first direction of the proof, it suffices to prove that $\mathfrak{T}, \frac{((u_1,t_1),...,(u_k,t_k))}{(x_1,...,x_k)} \models Diag_\delta(x_1, ..., x_k)$ for some torus points $t_2, ..., t_k$.

Let $t_2, ..., t_k$ be the torus points such that $(t_1, ..., t_k) \in T_\delta$. In order to establish that $\mathfrak{T}, \frac{((u_1,t_1),...,(u_k,t_k))}{(x_1,...,x_k)} \models Cons_\delta(x_1, ..., x_k)$, assume that $\delta \leq_f \eta$, where $\eta \in \Delta_p$ and $p \geq k$. Assume that $\langle\eta\rangle(\gamma_1, ..., \gamma_p) \in \mathrm{SUB}_\psi$, and that $\mathfrak{T}, \frac{((u_1,t_1),...,(u_k,t_k))}{(x_1,...,x_k)} \models P_{\gamma_1}(x_{f(1)}) \wedge ... \wedge P_{\gamma_p}(x_{f(p)})$. We must show that $(u_{f(1)}, t_{f(1)}) \in P^{\mathfrak{T}}_{\langle\eta\rangle(\gamma_1,...,\gamma_p)}$.

For each $i \in \{1, ..., p\}$, as $(u_{f(i)}, t_{f(i)}) \in P_{\gamma_i}^{\mathfrak{T}}$, we have $u_{f(i)} \in \|\gamma_i\|^{\mathfrak{M}}$ by the definition of $P_{\gamma_i}^{\mathfrak{T}}$. As $(u_1, ..., u_k) \in \|\delta\|^{\mathfrak{M}}$ and $\delta \leq_f \eta$, we have $(u_{f(1)}, ..., u_{f(p)}) \in \|\eta\|^{\mathfrak{M}}$. Therefore we have $u_{f(1)} \in \|\langle\eta\rangle(\gamma_1, ..., \gamma_p)\|^{\mathfrak{M}}$. Thus $(u_{f(1)}, t_{f(1)}) \in P^{\mathfrak{T}}_{\langle\eta\rangle(\gamma_1,...,\gamma_p)}$ by the definition of $P^{\mathfrak{T}}_{\langle\eta\rangle(\gamma_1,...,\gamma_p)}$.

We then deal with the converse implication of the lemma. Define $s_1 := t$ and $v_1 := u$. Assume $\mathfrak{T}, \frac{(v_1,s_1)}{x_1} \models \exists x_2 ... \exists x_k \big(Diag_\delta(x_1, ..., x_k) \wedge P_{\chi_1}(x_1) \wedge ... \wedge P_{\chi_k}(x_k) \big)$. Hence $\mathfrak{T}, \frac{((v_1,s_1),...,(v_k,s_k))}{(x_1,...,x_k)} \models Diag_\delta(x_1, ..., x_k)$ for some tuple $\big((v_1, s_1), ..., (v_k, s_k)\big)$ such that $(v_i, s_i) \in P_{\chi_i}^{\mathfrak{T}}$ for each i. As now $\mathfrak{T}, \frac{((v_1,s_1),...,(v_k,s_k))}{(x_1,...,x_k)} \models PreCons_\delta(x_1, ..., x_k)$, we infer that $(v_1, s_1) \in P^{\mathfrak{T}}_{\langle\delta\rangle(\chi_1,...,\chi_k)}$. By the definition of $P^{\mathfrak{T}}_{\langle\delta\rangle(\chi_i,...,\chi_k)}$, we have $(\mathfrak{M}, v_1) \models \langle\delta\rangle(\chi_1, ..., \chi_k)$. □

Lemma 4.2 *Let t be any torus point. Under the assumption $(\mathfrak{M}, w) \models \psi$, we have $\mathfrak{T}, \frac{(w,t)}{x} \models \psi^*(x)$.*

Proof. See the appendix. □

4.3 Satisfiability of $\psi^*(x)$ implies satisfiability of ψ

Let $\mathfrak{A}$ be a V^*-model with the domain A. Assume that $\mathfrak{A}, \frac{w}{x} \models \psi^*(x)$. We next define a V_ψ-model $\mathfrak{B}$ with the same domain A, and then show that $(\mathfrak{B}, w) \models \psi$.

Let U be a non-empty set, and let $p \in \mathbb{Z}_+$. Let $(u_1, ..., u_p) \in U^p$ be a tuple. We say that the tuple $(u_1, ..., u_p)$ *spans the set* $\{u_1, ..., u_p\}$.

Let $k \in \mathbb{Z}_+$, and let $S \in V_\psi$ be a k-ary symbol. We define $(u, ..., u)_k \in S^{\mathfrak{B}}$ iff $u \in P_S^{\mathfrak{A}}$. This settles the interpretation of the symbols $S \in V_\psi$ on tuples that span sets of size one. Interpretation of the symbols on tuples that span larger sets is more complicated. We begin with the following lemma, whose proof is straightforward by Lemma 3.1.

Lemma 4.3 *Let $u_1, ..., u_k \in A$. Assume $\mathfrak{A}, \frac{(u_1,...,u_k)}{(x_1,...,x_k)} \models Diag_\delta(x_1, ..., x_k)$. Then $\mathfrak{A}, \frac{(u_{\sigma(1)},...,u_{\sigma(k)})}{(x_1,...,x_k)} \not\models Diag_\eta(x_1, ..., x_k)$ holds for all all k-permutations σ and all $\eta \in \Delta_k \setminus \{\delta\}$. Also $\mathfrak{A}, \frac{(u_{\mu(1)},...,u_{\mu(k)})}{(x_1,...,x_k)} \not\models Diag_\delta(x_1, ..., x_k)$ holds for all k-permutations μ other than the identity permutation.*

Let $q \in \{2, ..., \mathcal{M}\}$. Consider subsets of A that have exactly $q \geq 2$ elements. Let us divide such sets into two classes. Let $U = \{u_1, ..., u_q\}$ be a set with q distinct elements. Assume first that there exists some q-permutation σ and some $\eta \in \Delta_q$ such that $\mathfrak{A}, \frac{(u_{\sigma(1)},...,u_{\sigma(q)})}{(x_1,...,x_q)} \models Diag_\eta(x_1, ..., x_q)$. Define $tuple(U) := (u_{\sigma(1)}, ..., u_{\sigma(q)})$ and $diagram(U) := \eta$. Define also $type(U) = 1$. Assume then that $\mathfrak{A}, \frac{(u_{\sigma(1)},...,u_{\sigma(q)})}{(x_1,...,x_q)} \not\models Diag_\eta(x_1, ..., x_q)$ holds for all $\eta \in \Delta_q$ and all q-permutations σ. As $\mathfrak{A} \models \psi_{total}$, there exists some diagram $\delta \in \Delta_q$ such that $\mathfrak{A}, \frac{(u_1,...,u_q)}{(x_1,...x_q)} \models Cons_\delta(x_1, ..., x_q)$. Define $tuple(U) = (u_1, ..., u_q)$ and $diagram(U) := \delta$. Define also $type(U) = 2$.

Notice that by our assumptions in Section 4.1, there are no relation symbols $S \in V_\psi \setminus D_\psi$ of any arity higher than one. Recall that $\mathcal{M}$ is the maximum arity of diagrams in Δ. We next define the relations $S^{\mathfrak{B}}$, where $S \in D_\psi$, on tuples of elements of A that span sets with $q \in \{2, ..., \mathcal{M}\}$ elements. The definition has the property—as Lemma 4.5 below establishes—that if $(u_1, ..., u_k) \in \|\delta\|^{\mathfrak{B}}$, where $\delta \in \Delta_k$, then $\mathfrak{A}, \frac{(u_1,...,u_k)}{(x_1,...,x_k)} \models PreCons_\delta(x_1, ..., x_k)$. In fact this holds also for tuples that span a singleton set, see Lemma 4.5.

Let $q \in \{2, ..., \mathcal{M}\}$, and let $U \subseteq A$ be a set of the size q. Assume first that $type(U) = 1$. Let $diagram(U) = \eta \in \Delta_q$ and $tuple(U) = (u_1, ..., u_q)$. We have $\mathfrak{A}, \frac{(u_1,...,u_q)}{(x_1,...,x_q)} \models Diag_\eta(x_1, ..., x_q)$. Let $k \geq q$ be an integer. Interpret each k-ary symbol $S \in D_\psi$ such that $\mathfrak{B}, \frac{(u_1,...,u_q)}{(x_1,...,x_q)} \models \eta$. This definition uniquely specifies the interpretation of S on each k-ary tuple that spans the set $\{u_1, ..., u_q\}$. To see this, let $f : \{1, ..., k\} \to \{1, ..., q\}$ be a surjection. Now we have $(u_{f(1)}, ..., u_{f(k)}) \in S^{\mathfrak{B}}$ iff $S(x_{f(1)}, ..., x_{f(k)}) \in \eta$.

Assume then that $type(U) = 2$. Let $diagram(U) = \delta \in \Delta_q$ and $tuple(U) = (v_1, ..., v_q)$. We have $\mathfrak{A}, \frac{(v_1,...,v_q)}{(x_1,...,x_q)} \models Cons_\delta(x_1, ..., x_p)$. Let $k \geq q$ be an integer. Interpret each k-ary symbol $S \in D_\psi$ such that $\mathfrak{B}, \frac{(v_1,...,v_q)}{(x_1,...,x_q)} \models \delta$.

We investigate each $q \in \{2, ..., \mathcal{M}\}$, and thereby obtain a definition of $\mathfrak{B}$; if there are symbols of arity $r > \mathcal{M}$ in D_ψ, we arbitrarily define the interpretations of such symbols on tuples that span sets with more than $\mathcal{M}$ elements.

Lemma 4.4 *If $(u_1, ..., u_k) \in A^k$ and $\mathfrak{A}, \frac{(u_1,...,u_k)}{(x_1,...,x_k)} \models Diag_\delta(x_1, ..., x_k)$ for some $\delta \in \Delta_k$, then $(u_1, ..., u_k) \in \|\delta\|^{\mathfrak{B}}$.*

Proof. Assume $\mathfrak{A}, \frac{(u_1,...,u_k)}{(x_1,...,x_k)} \models Diag_\delta(x_1,...,x_k)$. Notice that $k \geq 2$, since diagrams have by definition an arity at least two. As $\mathfrak{A} \models \psi_{uniq}$, the set $U = \{u_1,...,u_k\}$ has exactly k elements. We have *type*$(U) = 1$, and by Lemma 4.3, *tuple*$(U) = (u_1,...,u_k)$. Thus $\mathfrak{B}, \frac{(u_1,...,u_k)}{(x_1,...,x_k)} \models \delta$. □

Lemma 4.5 *Let* $k \in \{1,...,\mathcal{M}\}$. *If* $(u_1,...,u_k) \in \|\delta\|^{\mathfrak{B}}$, *where* $\delta \in \Delta_k$, *then* $\mathfrak{A}, \frac{(u_1,...,u_k)}{(x_1,...,x_k)} \models PreCons_\delta(x_1,...,x_k)$.

Proof. The case where $(u_1,..,u_k)$ spans a singleton set follows since $\mathfrak{A} \models \psi_{local}$. Let us consider the cases where $(u_1,..,u_k)$ spans a set of the size two or larger.

Assume that $(u_1,...,u_k) \in \|\delta\|^{\mathfrak{B}}$ is a tuple such that $U = \{u_1,...,u_k\}$ contains exactly $q \geq 2$ elements. Let $m : \{1,...,q\} \to \{1,...,k\}$ be an injection such that the tuple $(u_{m(1)},...,u_{m(q)})$ spans the set $\{u_1,...,u_k\}$.

Assume first that we have $\mathfrak{A}, \frac{(u_{m(\sigma(1))},...,u_{m(\sigma(q))})}{(x_1,...,x_q)} \models Diag_\eta(x_1,...,x_q)$ for some $\eta \in \Delta_q$ and some q-permutation σ. Thus *type*$(U) = 1$. By Lemma 4.3, we have *tuple*$(U) = (u_{m(\sigma(1))},...,u_{m(\sigma(q))})$ and *diagram*$(U) = \eta$. Let $s : \{1,...,q\} \to \{1,...,k\}$ be the injection such that $s(i) = m(\sigma(i))$ for each $i \in \{1,...,q\}$. As *tuple*$(U) = (u_{s(1)},...,u_{s(q)})$, we have $\mathfrak{B}, \frac{(u_{s(1)},...,u_{s(q)})}{(x_1,...,x_q)} \models \eta$. As $\mathfrak{A}, \frac{(u_{s(1)},...,u_{s(q)})}{(x_1,...,x_q)} \models Diag_\eta(x_1,...,x_q)$, we have $\mathfrak{A}, \frac{(u_{s(1)},...,u_{s(q)})}{(x_1,...,x_q)} \models Cons_\eta(x_1,...,x_q)$.

The rest or the argument for the case where *type*$(U) = 1$, will be dealt with below. Let us next elaborate some details related to the case where *type*$(U) = 2$. So, assume *type*$(U) = 2$. Let $t : \{1,...,q\} \to \{1,...,k\}$ be an injection such that *tuple*$(U) = (u_{t(1)},...,u_{t(q)})$. Let *diagram*$(U) = \rho \in \Delta_q$. Thus $\mathfrak{A}, \frac{(u_{t(1)},...,u_{t(q)})}{(x_1,...,x_q)} \models Cons_\rho(x_1,...,x_q)$ and $\mathfrak{B}, \frac{(u_{t(1)},...,u_{t(q)})}{(x_1,...,x_q)} \models \rho$.

We then complete the arguments for both cases *type*$(U) = 1$ and *type*$(U) = 2$. Let $(h,\nu) \in \{(s,\eta),(t,\rho)\}$, where s and t are the injections defined above, and of course η and ρ are the related diagrams.

Let $g : \{1,...,k\} \to \{1,...,q\}$ be the surjection such that $g(i) = j$ iff $u_i = u_{h(j)}$. Notice that $(u_{h(1)},...,u_{h(q)}) \in \|\nu\|^{\mathfrak{B}}$ and $(u_1,...,u_k) \in \|\delta\|^{\mathfrak{B}}$, and these two tuples span the same set with q elements. Thus we have $\nu \leq_g \delta$.

We have $\mathfrak{A}, \frac{(u_{h(1)},...,u_{h(q)})}{(x_1,...,x_q)} \models Cons_\nu(x_1,...,x_q)$. As $\nu \leq_g \delta$, we have $\mathfrak{A}, \frac{(u_{h(1)},...,u_{h(q)})}{(x_1,...,x_q)} \models PreCons_\delta(x_{g(1)},...,x_{g(k)})$. Recalling that $g(i) = j$ iff $u_i = u_{h(j)}$, we conclude that $\mathfrak{A}, \frac{(u_1,...,u_k)}{(x_1,...,x_k)} \models PreCons_\delta(x_1,...,x_k)$, as required. □

Lemma 4.6 *Let* $\alpha \in \mathrm{SUB}_\psi$ *and* $u \in A$. *We have* $(\mathfrak{B}, u) \models \alpha$ *iff* $\mathfrak{A}, \frac{u}{x} \models P_\alpha(x)$.

Proof. See the appendix. □

Due to Lemma 4.6, we observe that since $\mathfrak{A}, \frac{w}{x} \models P_\psi(x)$, we must have $(\mathfrak{B}, w) \models \psi$. Together with Lemma 4.2, this establishes the following theorems.

Theorem 4.7 *The one dimensional fragment has the finite model property.*

Corollary 4.8 *The satisfiability and finite satisfiability problems of the one dimensional fragment are decidable.*

5 Undecidable extensions

The *general one-dimensional fragment* GF_1 of first-order logic is defined in the same way as UF_1, except that the uniformity condition is relaxed. The set of τ-formulae of GF_1 is the smallest set $\mathcal{F}$ satisfying the following conditions.

(i) If φ is a unary τ-atom, then $\varphi \in \mathcal{F}$. Also $\top, \bot \in \mathcal{F}$.

(ii) If $\varphi \in \mathcal{F}$, then $\neg\varphi \in \mathcal{F}$. If $\varphi_1, \varphi_2 \in \mathcal{F}$, then $(\varphi_1 \wedge \varphi_2) \in \mathcal{F}$.

(iii) Let $Y = \{y_1, ..., y_k\}$ be a set of variable symbols. Let U be a finite set of formulae $\psi \in \mathcal{F}$ with free variables in Y. Let F be a set of τ-atoms with free variables in Y. Let φ be any Boolean combination of formulae in $F \cup U$. Then $\exists y_2 ... \exists y_k\, \varphi \in \mathcal{F}$ and $\exists y_1 ... \exists y_k\, \varphi \in \mathcal{F}$.

There are different natural ways of generalizing UF_1 so that a two-dimensional logic is obtained. Here we consider a formalism which we call the *strongly uniform two-dimensional fragment* SUF_2 of first-order logic. The set of τ-formulae of SUF_2 is the smallest set $\mathcal{F}$ satisfying the following conditions.

(i) If φ is a unary or a binary τ-atom, then $\varphi \in \mathcal{F}$. Also $\top, \bot \in \mathcal{F}$.

(ii) If $\varphi \in \mathcal{F}$, then $\neg\varphi \in \mathcal{F}$. If $\varphi_1, \varphi_2 \in \mathcal{F}$, then $(\varphi_1 \wedge \varphi_2) \in \mathcal{F}$.

(iii) Let y_1 and y_2 be variable symbols. Let U be a finite set of formulae $\psi \in \mathcal{F}$ whose free variables are in $\{y_1, y_2\}$. Let φ be any Boolean combination of formulae in U. Then $\exists y_2\, \varphi \in \mathcal{F}$ and $\exists y_1 \exists y_2\, \varphi \in \mathcal{F}$.

(iv) Let $Y = \{y_1, ..., y_k\}$, $k \geq 3$, be a set of variable symbols. Let U be a finite set of formulae $\psi \in \mathcal{F}$ such that each ψ has at most one free variable, and the variable is in Y. Let F be a V-uniform set, $V \subseteq Y$, of τ-atoms. Let φ be any Boolean combination of formulae in $F \cup U$. Then $\exists y_3 ... \exists y_k\, \varphi \in \mathcal{F}$, $\exists y_2 ... \exists y_k\, \varphi \in \mathcal{F}$ and $\exists y_1 ... \exists y_k\, \varphi \in \mathcal{F}$.

Both of these extensions of UF_1 are Π^0_1-complete; see the appendix for the proofs. This shows that if we lift either of the two principal syntactic restrictions of UF_1, we obtain an undecidable formalism.

6 Expressivity

Guarded negation first-order logic GNFO is a novel fragment of first-order logic introduced in [2]. GNFO subsumes the guarded fragment GF. It turns out that UF_1 is incomparable in expressivity with both GNFO and the two-variable fragment with counting quantifiers FOC^2. This is proved in the appendix.

7 Conclusion

The main contribution of this paper is the discovery of the fragment UF_1 via the introduction of the notions of *uniformity* and *one-dimensionality.* The notions offer a new perspective on why modal logics are robustly decidable. Also, UF_1 extends equality-free FO^2 in a natural way, and thus provides a possible novel direction in the currently very active research on two-variable logics. Also, we believe that our satisfiability preserving translation of UF_1 into the monadic class is of independent mathematical interest. The translation is robust and can be altered and extended to give other decidability proofs.

In the future we intend to study variants of UF_1 with identity. It was observed in [2] that adding the formula $\forall x \forall y\big(Rxy \leftrightarrow x \neq y\big)$ to GNFO leads to an undecidable formalism. It is not immediately clear whether the extension of UF_1 with the free use of equality and inequality results in undecidability. We are currently working on related decidability and complexity questions.

We conjecture that our decidability result can be carried out by an alternative method combining a generalization of Scott normal form and the *dual Maslov class*. The alternative method does not involve a new proof technique, unlike the work above.

Appendix

A Translation $\mathrm{UF}_1 \rightarrow \mathrm{DUF}_1$

Proposition A.1 *There is an effective translation that transforms each formula in* UF_1 *to an equivalent formula in* DUF_1.

Proof. Let $\chi := \exists y_2 ... \exists y_k\, \varphi$ be a formula of UF_1 formed using the formation rule (iii) in the definition of UF_1. We may assume, w.l.o.g., that that the variables $y_1, ..., y_k$ are distinct, and that $k \geq 2$. Define $Y := \{y_1, ..., y_k\}$. Let τ_χ be the set of relation symbols in χ of the arity two and higher.

Put φ into disjunctive normal form. We obtain a formula $\exists y_2 ... \exists y_k \big(\varphi_1 \vee ... \vee \varphi_n\big)$. Now distribute the existential quantifier prefix $\exists y_2 ... \exists y_k$ over the disjunctions, obtaining the formula $\exists y_2 ... \exists y_k\, \varphi_1 \vee ... \vee \exists y_2 ... \exists y_k\, \varphi_n$.

Now consider the formula φ_j. Assume first that φ_j is of the type $\alpha \wedge \psi$, where α is a non-empty conjunction of atoms and negated atoms of the arity $m \geq 2$, and ψ is a non-empty conjunction of formulae that have at most one free variable. Let $z_2, ..., z_p \in Y$ denote the variables in $Y \setminus \{y_1\}$ that occur in α. Notice that $p = m$ if and only if y_1 occurs in α. Let z_1 denote y_1.

Let $z_{p+1}, ..., z_k \in Y$ be the variables in $Y \setminus \{z_1, ..., z_p\}$. Notice that the formula ψ is equivalent to the conjunction $\psi_1(z_1) \wedge ... \wedge \psi_k(z_k) \wedge \beta$, where each formula $\psi_i(z_i)$ is the conjunction of exactly all conjuncts of ψ with the free variable z_i, in the case such conjuncts exist, and $\psi_i(z_i)$ is the formula $\top$ otherwise; the formula β is the conjunction of the conjuncts of ψ without free variables. The formula φ_j is equivalent to the formula $\exists z_2 ... \exists z_p\big(\, \alpha \wedge \psi_1(z_1) \wedge ... \wedge \psi_p(z_p)\,\big) \wedge \exists z_{p+1} \psi_{p+1}(z_{p+1}) \wedge ... \wedge \exists z_k \psi_k(z_k) \wedge \beta$. Notice that for each i, the formula $\exists z_i \psi_i(z_i)$ is a DUF_1-formula if $\psi_i(z_i)$ is.

Consider the formula $\gamma := \exists z_2 ... \exists z_p\big(\, \alpha \wedge \psi_1(z_1) \wedge ... \wedge \psi_p(z_p)\,\big)$. The formula α is either equivalent to $\perp$, or equivalent to a non-empty disjunction $\delta_1 \vee ... \vee \delta_l$, where each δ_i denotes a conjunction over some uniform m-ary τ_χ-diagram. (Notice that since α is quantifier-free, the equivalence checking can be done effectively.) Assume first that α is equivalent to $\delta_1 \vee ... \vee \delta_l$. Therefore the formula γ is equivalent to the disjunction $\exists z_2 ... \exists z_p\big(\, \delta_1 \wedge \psi_1(z_1) \wedge ... \wedge \psi_p(z_p)\,\big) \vee ... \vee \exists z_2 ... \exists z_p\big(\, \delta_l \wedge \psi_1(z_1) \wedge ... \wedge \psi_p(z_p)\,\big)$. Notice that the disjunct $\exists z_2 ... \exists z_p\big(\, \delta_i \wedge \psi_1(z_1) \wedge ... \wedge \psi_p(z_p)\,\big)$ is a DUF_1-formula if the formulae $\psi_1(z_1), ..., \psi_p(z_p)$ are;

we may need to use the formation rule (iv) of DUF_1 in addition to rule (iii) if $\exists z_2 ... \exists z_p\big(\,\delta_i \wedge \psi_1(z_1) \wedge ... \wedge \psi_p(z_p)\big)$ does not contain the free variable z_1. In the case α is equivalent to $\bot$, then γ is equivalent to $\bot$.

We have now discussed the case where φ_j is of the type $\exists y_2 ... \exists y_k\big(\,\alpha \wedge \psi\,\big)$, where α is a non-empty conjunction of atoms and negated atoms of some arity higher than one, and ψ is a non-empty conjunction of formulae with at most one free variable. The case where φ_j is $\exists y_2 ... \exists y_k\, \alpha$, can be reduced to the case already discussed by considering the formula $\exists y_2 ... \exists y_k\big(\,\alpha \wedge \top\big)$. Assume thus that φ_j is the formula $\exists y_2 ... \exists y_k\, \psi$, where ψ is some conjunction $\psi_1(y_1) \wedge ... \wedge \psi_k(y_k) \wedge \beta$, where the formulae $\psi_i(y_i)$ have at most one free variable, and β has no free variables. Now φ_j is equivalent to the formula $\psi_1(y_1) \wedge \exists y_2 \psi_2(y_2) \wedge ... \wedge \exists y_k \psi_k(y_k) \wedge \beta$. Each conjunct $\exists y_i \psi_1(y_i)$ is a DUF_1-formula if $\psi_i(y_i)$ is.

All other cases the translation from UF_1 to DUF_1 are straightforward. □

B Proofs for Section 4

Proof of Lemma 4.2. We establish the claim of the lemma by showing that

$$\mathfrak{T}, \frac{(w,t)}{x} \models \psi_{total} \wedge \psi_{uniq} \wedge \psi_{local} \wedge \psi_{sub} \wedge P_{\psi}(x).$$

To show that $\mathfrak{T} \models \psi_{total}$, let $\big((u_1,t_1),...,(u_k,t_k)\big) \in (Dom(\mathfrak{T}))^k$, where $k \in \{2,...,\mathcal{M}\}$. We need to show that $\big((u_1,t_1),...,(u_k,t_k)\big)$ satisfies $Cons_{\delta}(x_1,...,x_k)$ for some $\delta \in \Delta_k$. Consider the tuple $(u_1,...,u_k) \in M^k$. Let η be the unique standard uniform k-ary V_{ψ}-diagram η such that $(u_1,...,u_k) \in \|\eta\|^{\mathfrak{M}}$. Let $p \in \{2,...,\mathcal{M}\}$, $p \geq k$. Let $\rho \in \Delta_p$. Let $f : \{1,...,p\} \rightarrow \{1,...,k\}$ be a surjection, and assume that $\eta \leq_f \rho$. Thus $(u_{f(1)},...,u_{f(p)}) \in \|\rho\|^{\mathfrak{M}}$. In order to conclude that $\mathfrak{T} \models \psi_{total}$, we need to show that $\mathfrak{T}, \frac{\left((u_1,t_1),...,(u_k,t_k)\right)}{(x_1,...,x_k)} \models PreCons_{\rho}(x_{f(1)},...,x_{f(p)})$. Therefore we assume that $\mathfrak{T}, \frac{\left((u_1,t_1),...,(u_k,t_k)\right)}{(x_1,...,x_k)} \models P_{\chi_1}(x_{f(1)}) \wedge ... \wedge P_{\chi_p}(x_{f(p)})$. Thus we have $(\mathfrak{M}, u_{f(i)}) \models \chi_i$ for each $i \in \{1,...,p\}$. As $(u_{f(1)},...,u_{f(p)}) \in \|\rho\|^{\mathfrak{M}}$, we therefore have $u_{f(1)} \in \|\langle\rho\rangle(\chi_1,...,\chi_p)\|^{\mathfrak{M}}$. Thus $(u_{f(1)}, t_{f(1)}) \in P^{\mathfrak{T}}_{\langle\rho\rangle(\chi_1,...,\chi_p)}$, whence $\mathfrak{T}, \frac{\left((u_1,t_1),...,(u_k,t_k)\right)}{(x_1,...,x_k)} \models P_{\langle\rho\rangle(\chi_1,...,\chi_p)}(x_{f(1)})$. Therefore $\mathfrak{T} \models \psi_{total}$.

It is immediate by the definition of the domain of $\mathfrak{T}$ and the predicates $P^{\mathfrak{T}}_t$, where t is a torus point, that $\mathfrak{T} \models \psi_{uniq}$.

To show that $\mathfrak{T} \models \psi_{local}$, assume $\mathfrak{T}, \frac{(u,t)}{x} \models Local_{\delta}(x)$ for some k-ary diagram $\delta \in \Delta$. Thus $(u,...,u)_k \in \|\delta\|^{\mathfrak{M}}$. To show that $\mathfrak{T}, \frac{(u,t)}{x} \models PreCons_{\delta}(x,...,x)_k$, let $\langle\delta\rangle(\chi_1,...,\chi_k) \in \mathrm{SUB}_{\psi}$ and assume that $\mathfrak{T}, \frac{(u,t)}{x} \models P_{\chi_1}(x) \wedge ... \wedge P_{\chi_k}(x)$. Therefore $u \in \|\chi_i\|^{\mathfrak{M}}$ for each $i \in \{1,...,k\}$, whence $u \in \|\langle\delta\rangle(\chi_1,...,\chi_k)\|^{\mathfrak{M}}$. Thus $(u,t) \in P^{\mathfrak{T}}_{\langle\delta\rangle(\chi_1,...,\chi_k)}$, as required.

The non-trivial part in proving that $\mathfrak{T} \models \psi_{sub}$ involves showing that $\mathfrak{T} \models \psi_{\langle\delta\rangle(\chi_1,...,\chi_k)}$ for formulae of the type $\langle\delta\rangle(\chi_1,...,\chi_k)$. This follows directly by Lemma 4.1, since $P^{\mathfrak{T}}_{\langle\delta\rangle(\chi_1,...,\chi_k)} = \|\langle\delta\rangle(\chi_1,...,\chi_k)\|^{\mathfrak{M}} \times Dom(\mathfrak{T})$.

Since $(\mathfrak{M}, w) \models \psi$ and $P_\psi^{\mathfrak{T}} = \|\psi\|^{\mathfrak{M}} \times Dom(\mathfrak{T})$, we have $\mathfrak{T}, \frac{(w,t)}{x} \models P_\psi(x)$.□

Proof of Lemma 4.6. We establish the claim by induction on the structure of α. For all atomic formulae $S \in \mathrm{SUB}_\psi$, the claim follows directly from the definition of the relations $S^{\mathfrak{B}}$ on tuples that span a singleton set. The cases where α is of form $\neg\beta$ or $(\beta \wedge \gamma)$ are straightforward since $\mathfrak{A} \models \psi_{sub}$.

Define $u_1 := u$ and $x_1 := x$. Assume that $\mathfrak{B}, \frac{u_1}{x_1} \models \langle\delta\rangle(\chi_1, ..., \chi_k)$, where $\langle\delta\rangle(\chi_1, ..., \chi_k) \in \mathrm{SUB}_\psi$. Thus $(u_1, ..., u_k) \in \|\delta\|^{\mathfrak{B}}$ for some tuple $(u_1, ..., u_k)$ such that $u_i \in \|\chi_i\|^{\mathfrak{B}}$ for each $i \in \{1, ..., k\}$. Now, for each $i \in \{1, ..., k\}$, we have $P_{\chi_i}^{\mathfrak{A}} = \|\chi_i\|^{\mathfrak{B}}$ by the induction hypothesis, and therefore $u_i \in P_{\chi_i}^{\mathfrak{A}}$. By Lemma 4.5, we have $\mathfrak{A}, \frac{(u_1,...,u_k)}{(x_1,...,x_k)} \models PreCons_\delta(x_1, ..., x_k)$. By the definition of the formula $PreCons_\delta(x_1, ..., x_k)$, we conclude that $\mathfrak{A}, \frac{u_1}{x_1} \models P_{\langle\delta\rangle(\chi_1,...,\chi_k)}(x_1)$.

For the converse, assume $\mathfrak{A}, \frac{u_1}{x_1} \models P_{\langle\delta\rangle(\chi_1,...,\chi_k)}(x_1)$. As $\mathfrak{A} \models \psi_{\langle\delta\rangle(\chi_1,...,\chi_k)}$, we have $\mathfrak{A}, \frac{u_1}{x_1} \models \exists x_2 ... \exists x_k \big(Diag_\delta(x_1, ..., x_k) \wedge P_{\chi_1}(x_1) \wedge ... \wedge P_{\chi_k}(x_k)\big)$. Hence there exists some tuple $(u_1, ..., u_k)$ such that $u_i \in P_{\chi_i}^{\mathfrak{A}}$ for each i and $\mathfrak{A}, \frac{(u_1,...,u_k)}{(x_1,...,x_k)} \models Diag_\delta(x_1, ..., x_k)$. By Lemma 4.4, we have $(u_i, ..., u_k) \in \|\delta\|^{\mathfrak{B}}$. As $\|\chi_i\|^{\mathfrak{B}} = P_{\chi_i}^{\mathfrak{A}}$ for each i by the induction hypothesis, we conclude that $(\mathfrak{B}, u_1) \models \langle\delta\rangle(\chi_1, ..., \chi_k)$.

Assume first that $(\mathfrak{B}, u) \models \langle E\rangle\chi$, where $\langle E\rangle\chi \in \mathrm{SUB}_\psi$. Thus $(\mathfrak{B}, v) \models \chi$ for some v, whence $\mathfrak{A}, \frac{v}{y} \models P_\chi(y)$ by the induction hypothesis. Thus $\mathfrak{A} \models \exists y P_\chi(y)$. As $\mathfrak{A} \models \psi_{sub}$, we have $\mathfrak{A}, \frac{u}{x} \models P_{\langle E\rangle\chi}(x)$. Assume then that $\mathfrak{A}, \frac{u}{x} \models P_{\langle E\rangle\chi}(x)$. As $\mathfrak{A} \models \psi_{sub}$, we have $\mathfrak{A} \models \exists y P_\chi(y)$, whence $\mathfrak{A}, \frac{v}{y} \models P_\chi(y)$ for some v. By the induction hypothesis, we have $(\mathfrak{B}, v) \models \chi$, whence $(\mathfrak{B}, u) \models \langle E\rangle\chi$. □

C Arguments concerning undecidable extensions

We recall the tiling problem of the infinite grid $\mathbb{N} \times \mathbb{N}$. A tile is a map $t : \{R, L, T, B\} \to C$, where C is a countably infinite set of colours. We use the notation $t_X := t(X)$ for $X \in \{R, L, T, B\}$. Intuitively, t_R, t_L, t_T and t_B are the colours of the right edge, left edge, top edge and bottom edge of the tile t.

Let $\mathbb{T}$ be a finite set of tiles. A $\mathbb{T}$-tiling of $\mathbb{N} \times \mathbb{N}$ is a function $f : \mathbb{N} \times \mathbb{N} \to \mathbb{T}$ that satisfies the following horizontal and vertical tiling conditions:

(T_H) For all $i, j \in \mathbb{N}$, if $f(i, j) = t$ and $f(i+1, j) = t'$, then $t_R = t'_L$.

(T_V) For all $i, j \in \mathbb{N}$, if $f(i, j) = t$ and $f(i, j+1) = t'$, then $t_T = t'_B$.

Thus, f is a proper tiling iff the colors on the matching edges of any two adjacent tiles coincide. The tiling problem for $\mathbb{N} \times \mathbb{N}$ asks whether for a finite set $\mathbb{T}$ of tiles, there is a $\mathbb{T}$-tiling of $\mathbb{N} \times \mathbb{N}$. It is well known that this problem is undecidable (Π_1^0-complete). Using the problem, it is easy to prove the following.

Proposition C.1 *The satisfiability problem of* GF_1 *is* Π_1^0*-complete.*

Proof. Let $\tau = \{H, V\}$ be a vocabulary, where H and V are binary relation symbols. The infinite grid $\mathbb{N} \times \mathbb{N}$ can be represented by a τ-structure $\mathfrak{G} := (\mathbb{N} \times \mathbb{N}, H^{\mathfrak{G}}, V^{\mathfrak{G}})$, where $H^{\mathfrak{G}} := \{((i, j), (i+1, j)) \mid i, j \in \mathbb{N}\}$ and $V^{\mathfrak{G}} := \{((i, j), (i, j+1)) \mid i, j \in \mathbb{N}\}$. Let Γ be the conjunction of the three τ-sentences $\eta_H := \forall x \exists y\, H(x,y)$, $\eta_V := \forall x \exists y\, V(x,y)$, and $\eta_{Com} := \forall x \forall y \forall z \forall w\, \big((H(x,y) \wedge V(x,z) \wedge H(z,w)) \rightarrow V(y,w)\big)$. It is easy to see that η_H, η_V and η_{Com} are in GF_1.

It is straightforward to show that if $\mathfrak{M}$ is a τ-model such that $\mathfrak{M} \models \Gamma$, then there exists a homomorphism $h : \mathfrak{G} \rightarrow \mathfrak{M}$.

Let $\mathbb{T}$ be a set of tiles. We simulate tiles by unary relation symbols P_t for each $t \in \mathbb{T}$. We denote the corresponding vocabulary $\tau \cup \{P_t \,|\, t \in \mathbb{T}\}$ by $\sigma_{\mathbb{T}}$. The tiling conditions (T_H) and (T_V) can be expressed by the $\sigma_{\mathbb{T}}$-sentences $\psi_H := \forall x \forall y \bigwedge_{t,t' \in \mathbb{T},\, t_R \neq t'_L} (P_t(x) \wedge P_{t'}(y)) \rightarrow \neg H(x,y)$ and $\psi_V := \forall x \forall y \bigwedge_{t,t' \in \mathbb{T},\, t_T \neq t'_B} (P_t(x) \wedge P_{t'}(y)) \rightarrow \neg V(x,y)$. Let $\Psi_{\mathbb{T}} := \psi_H \wedge \psi_V \wedge \psi_{part}$, where ψ_{part} is a sentence saying that every element is in exactly one of the relations P_t, $t \in \mathbb{T}$. Clearly ψ_{part} can be expressed in GF_1.

It is straightforward to show that the sentence $\Gamma \wedge \Psi_{\mathbb{T}}$ is satisfiable if and only if $\mathbb{N} \times \mathbb{N}$ is $\mathbb{T}$-tilable. Since the sentence $\Gamma \wedge \Psi_{\mathbb{T}}$ is in GF_1 for each finite set $\mathbb{T}$ of tiles, the tiling problem is effectively reducible to the satisfiability problem of GF_1. Hence the satisfiability problem is Π^0_1-hard. On the other hand, GF_1 is a fragment of first-order logic, whence its satisfiability problem is in Π^0_1. □

Let $\tau_+ = \{H_+, V_+, S\}$ be a vocabulary, where H_+ and V_+ are ternary relation symbols and S is a binary relation symbol. We will represent the infinite grid $\mathbb{N} \times \mathbb{N}$ as a τ_+-structure $\mathfrak{G}_+ := (\mathbb{N}, H_+^{\mathfrak{G}_+}, V_+^{\mathfrak{G}_+}, S^{\mathfrak{G}_+})$, where $H_+^{\mathfrak{G}_+} := \{(i, i+1, j) \,|\, i,j \in \mathbb{N}\}$, $V_+^{\mathfrak{G}_+} := \{(i,j,j+1) \,|\, i,j \in \mathbb{N}\}$, and $S_+^{\mathfrak{G}_+} := \{(i, i+1) \,|\, i \in \mathbb{N}\}$. Notice that $(u,v,w) \in V_+^{\mathfrak{G}_+}$ iff (u,v) connects to (u,w) via the vertical successor $V^{\mathfrak{G}}$ of the standard Cartesian grid $\mathfrak{G}$ defined in the proof of Proposition C.1. On the other hand, $(u,v,w) \in H_+^{\mathfrak{G}_+}$ iff $\big((u,w),(v,w)\big) \in H^{\mathfrak{G}}$. We shall next form a τ_+-sentence Γ_+ of SUF_2 such that $\mathfrak{G}_+ \models \Gamma_+$, and there is a homomorphism from $\mathfrak{G}_+$ to any model of Γ_+. Define Γ_+ to be the conjunction of the formulae $\theta_S := \forall x \exists y\, S(x,y)$, $\theta_H := \forall x_1 \forall x_2\, (S(x_1,x_2) \rightarrow \forall y\, H_+(x_1,x_2,y))$, and $\theta_V := \forall y_1 \forall y_2\, (S(y_1,y_2) \rightarrow \forall x\, V_+(x,y_1,y_2))$.

Lemma C.2 *If $\mathfrak{M}$ is a τ_+-model such that $\mathfrak{M} \models \Gamma_+$, then there exists a homomorphism $h : \mathfrak{G}_+ \rightarrow \mathfrak{M}$.*

Proof. We define a function $h : \mathbb{N} \rightarrow M$ by recursion as follows. Choose an arbitrary point $a_0 \in M$, and set $h(0) := a_0$. Assume that $h(i) = a$ has been defined. Since $\mathfrak{M} \models \theta_S$, there is $b \in M$ such that $(a,b) \in S^{\mathfrak{M}}$. Define $h(i+1) := b$. Observe first that $(h(i), h(i+1)) \in S^{\mathfrak{M}}$ for each $i \in \mathbb{N}$. Furthermore, since $\mathfrak{M} \models \theta_H \wedge \theta_V$, we have $(h(i), h(i+1), h(j)) \in H_+^{\mathfrak{M}}$ and $(h(i), h(j), h(j+1)) \in V_+^{\mathfrak{M}}$ for all $i,j \in \mathbb{N}$. Thus h is a homomorphism $\mathfrak{G}_+ \rightarrow \mathfrak{A}$. □

Theorem C.3 *The satisfiability problem of* SUF_2 *is* Π^0_1*-complete.*

Proof. By Lemma C.2, we know that if $\mathfrak{M}$ is a τ_+-model such that $\mathfrak{M} \models \Gamma_+$, then there exists a homomorphism $h : \mathfrak{G}_+ \rightarrow \mathfrak{M}$. (We also have $\mathfrak{G}_+ \models \Gamma_+$.)

Let $\mathbb{T}$ be a set of tiles. This time we simulate tiles by fresh ternary relation symbols $P_{X,t}$, where $X \in \{R, L, T, B\}$ and $t \in \mathbb{T}$. Let $\rho_{\mathbb{T}} := \tau_+ \cup \{P_{X,t} \,|\, X \in \{R,L,T,B\}, t \in \mathbb{T}\}$ be the corresponding vocabulary.

The idea here is that if $(a, b, c) \in P_{R,t}$ and $(a, b, c) \in P_{L,t'}$, then the right edge of (a, c) is coloured with t_R and the left edge of (b, c) is coloured with t'_L; recall that $(a, b, c) \in H_+^{\mathfrak{G}+}$ means that $\big((a, c), (b, c)\big) \in H^{\mathfrak{G}}$. Similarly, if $(a, b, c) \in P_{T,t}$ and $(a, b, c) \in P_{B,t'}$, then the top edge of (a, b) is coloured with t_T and the bottom edge of (a, c) is coloured with t'_B. Thus, we can express the tiling conditions (T_H) and (T_V) by the following SUF_2-sentences:

$$\varphi_H := \forall x_1 \forall x_2 \forall y \bigwedge_{t,t' \in \mathbb{T},\ t_R \neq t'_L} \Big(\big(P_{R,t}(x_1, x_2, y) \wedge P_{L,t'}(x_1, x_2, y)\big) \to \neg H_+(x_1, x_2, y)\Big),$$

$$\varphi_V := \forall x \forall y_1 \forall y_2 \bigwedge_{t,t' \in \mathbb{T},\ t_T \neq t'_B} \Big(\big(P_{T,t}(x, y_1, y_2) \wedge P_{B,t'}(x, y_1, y_2)\big) \to \neg V_+(x, y_1, y_2)\Big).$$

We also need a sentence φ_{prop} stating that each pair (a, b) is tiled by exactly one $t \in \mathbb{T}$. This amounts to stating, firstly, that the interpretation of each symbol $P_{R,t}$ depends only on the first and the last variable: $\bigwedge_{t \in \mathbb{T}} \forall x_1 \forall y\, (\exists x_2\, P_{R,t}(x_1, x_2, y) \to \forall x_2\, P_{R,t}(x_1, x_2, y))$, and analogously for $P_{L,t}$, $P_{T,t}$ and $P_{B,t}$. Secondly, the four colors of each pair correspond to the same tile, meaning that $\bigwedge_{t \in \mathbb{T}} \forall x_1 \forall y\, (\exists x_2\, P_{R,t}(x_1, x_2, y) \leftrightarrow \exists x_2\, P_{L,t}(x_2, x_1, y))$ holds, and similar conditions for the other pairs $(P_{X,t}, P_{Y,t})$ hold. Thirdly, for each $X \in \{L, R, B, T\}$, every triple is in exactly one of the relations $P_{X,t}$, $t \in \mathbb{T}$.

Clearly there is such a sentence φ_{prop} in SUF_2. Let $\Phi_{\mathbb{T}}$ be the conjunction of the sentences φ_H, φ_V and φ_{prop}. Thus the sentence $\Gamma_+ \wedge \Phi_{\mathbb{T}}$ is satisfiable if and only if $\mathbb{N} \times \mathbb{N}$ is $\mathbb{T}$-tilable. Hence we conclude that SUF_2 is Π^0_1-complete.□

D Expressivity

Theorem D.1 UF_1 *is incomparable in expressivity with both two-variable logic with counting* (FOC^2) *and guarded negation fragment* (GNFO)*.*

Proof. The expressivity of FOC^2 is seriously limited when it comes to properties of relations of arities greater than two. It is easy to show that for example the UF_1-sentence $\exists x \exists y \exists z\, R(x, y, z)$ is not expressible in FOC^2. Thus UF_1 is not contained in FOC^2.

It is straightforward to show by using the bisimulation for GNFO, provided in [2], that the UF_1-sentence $\exists x \exists y\, \neg R(x, y)$ is not expressible in GNFO. This follows from the fact that structures $\big(\{a\}, \{(a, a)\}\big)$ and $\big(\{a, b\}, \{(a, a), (b, b)\}\big)$ are bisimilar in the sense of GNFO. Thus UF_1 is not contained in GNFO.

The FO^2-sentence $\forall x \forall y (x = y)$ cannot be expressed in UF_1. This can be seen (for example) by observing that the two directions of our decidability proof together entail that satisfiable sentences of the equality-free logic UF_1 can always be satisfied in a larger model. Thus UF_1 does not contain FO^2.

It follows immediately from the definition of UF_1 that the equality-free fragment of FO^2 is contained in UF_1. In fact, it is easy to prove that in restriction to models with relation symbols of arities at most two, the expressivities of UF_1 and the identity-free fragment of FO^2 coincide. (Consider for example the translation from UF_1 to MUF_1 in the case of such vocabularies.)

To see that UF_1 does not contain GNFO, consider the GNFO-sentence $\exists x \exists y \exists z(Rxy \wedge Ryz \wedge Rzx)$. It is easy to show (by a pebble game argument, see [12]), that this property is not expressible in FO^2. As UF_1 is contained in FO^2 when attention is restricted to models with only binary relations, UF_1 does not contain GNFO. □

References

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

[2] Bárány, V., B. ten Cate and L. Segoufin, *Guarded negation*, in: *Proceedings of Automata, Languages and Programming - 38th International Colloquium, (ICALP), Part II* (2011), pp. 356–367.

[3] Benaim, S., M. Benedikt, W. Charatonik, E. Kieroński, R. Lenhardt, F. Mazowiecki and J. Worrell, *Complexity of two-variable logic on finite trees*, in: *Proceedings of Automata, Languages and Programming - 40th International Colloquium, (ICALP), Part II* (2013), pp. 74–88.

[4] Blackburn, P., M. de Rijke and Y. Venema, "Modal Logic," Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 2001.

[5] Charatonik, W. and P. Witkowski, *Two-variable logic with counting and trees*, in: *Proceedings of the 28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2013, pp. 73–82.

[6] Gargov, G. and S. Passy, *Modal environment for boolean speculations*, in: *Mathematical logic and its applications. Proc. of the Summer School and Conference dedicated to the 80th Anniversary of Kurt Gödel* (1987), pp. 253–263.

[7] Grädel, E., *Why are modal logics so robustly decidable?*, in: *Current Trends in Theoretical Computer Science*, World Scientific, 2001 pp. 393–408.

[8] Grädel, E., P. G. Kolaitis and M. Y. Vardi, *On the decision problem for two-variable first-order logic*, Bulletin of Symbolic Logic **3** (1997), pp. 53–69.

[9] Grädel, E., M. Otto and E. Rosen, *Two-variable logic with counting is decidable*, in: *12th Annual IEEE Symposium on Logic in Computer Science (LICS)*, 1997, pp. 306–317.

[10] Henkin, L., *Logical systems containing only a finite number of symbols.*, Presses De l'Université De Montréal (1967).

[11] Kieroński, E., J. Michaliszyn, I. Pratt-Hartmann and L. Tendera, *Two-variable first-order logic with equivalence closure*, in: *Proceedings of the 27th Annual IEEE Symposium on Logic in Computer Science (LICS)*, 2012, pp. 431–440.

[12] Libkin, L., "Elements of Finite Model Theory," Texts in Theoretical Computer Science. An EATCS Series, Springer, 2004.

[13] Lutz, C. and U. Sattler, *The complexity of reasoning with boolean modal logics*, in: *Proceedings of Advances in Modal Logic (AiML)*, 2000, pp. 329–348.

[14] Mortimer, M., *On languages with two variables.*, Mathematical Logic Quarterly **21** (1975), pp. 135–140.

[15] Pacholski, L., W. Szwast and L. Tendera, *Complexity of two-variable logic with counting*, in: *Proceedings of the 12th Annual IEEE Symposium on Logic in Computer Science (LICS)*, 1997, pp. 318–327.

[16] Pratt-Hartmann, I., *Complexity of the two-variable fragment with counting quantifiers*, Journal of Logic, Language and Information **14** (2005), pp. 369–395.

[17] Szwast, W. and L. Tendera, FO^2 *with one transitive relation is decidable*, in: *Proceedings of 30th International Symposium on Theoretical Aspects of Computer Science (STACS)*, 2013, pp. 317–328.

[18] Vardi, M. Y., *Why is modal logic so robustly decidable?*, in: *Descriptive Complexity and Finite Models, Proceedings of a DIMACS Workshop* (1996), pp. 149–184.