

Complexity and Expressivity of Uniform One-Dimensional Fragment with Equality

Emanuel Kieroński* and Antti Kuusisto†

April 24, 2014

Abstract

Uniform one-dimensional fragment $UF_1^=$ is a formalism obtained from first-order logic by limiting quantification 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 formulas with two or more variables. The fragment can be seen as a canonical generalization of two-variable logic, defined in order to be able to deal with relations of arbitrary arities. The fragment was introduced recently, and it was shown that the satisfiability problem of the equality-free fragment UF_1 of $UF_1^=$ is decidable. In this article we establish that the satisfiability and finite satisfiability problems of $UF_1^=$ are NEXPTIME-complete. We also show that the corresponding problems for the extension of $UF_1^=$ with counting quantifiers are undecidable. In addition to decidability questions, we compare the expressivities of $UF_1^=$ and two-variable logic with counting quantifiers FOC^2 . We show that while the logics are incomparable in general, $UF_1^=$ is strictly contained in FOC^2 when attention is restricted to vocabularies with the arity bound two.

1 Introduction

Two-variable logic FO^2 was introduced by Henkin in [9] and proved decidable in [12] by Mortimer. The satisfiability and finite satisfiability problems of FO^2 were shown to be NEXPTIME-complete in [6]. The extension of two-variable logic with counting quantifiers, FOC^2 , was proved decidable in [7], [13]. It was subsequently shown to be NEXPTIME-complete in [14]. Research on extensions and variants of two-variable logic is *currently very*

*Institute of Computer Science, University of Wrocław, Poland

†University of Tampere Logic Group and Institute of Computer Science, University of Wrocław, Poland

active. 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. Recent articles in the field include for example [3], [4], [10], [16], and several others.

Typical systems of modal logic are contained in two-variable logic, or some variant of it, and hence investigations on two-variable logics have direct implications on various fields of computer science, including verification of software and hardware, distributed systems, knowledge representation and artificial intelligence. However, two-variable logics do not cope well with relations of arities greater than two, and therefore the *scope of related research is significantly restricted.* In database theory contexts, for example, two-variable logics as such are usually not directly applicable due to the *severe arity-related limitations.*

The recent article [8] introduces the *uniform one-dimensional fragment*, UF_1^- , which is a natural generalization of FO^2 to contexts with relations of arbitrary arities. The logic UF_1^- is a fragment of first-order logic 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 formulas: if $n, k \geq 2$, then a Boolean combination of atoms $R(x_1, \dots, x_k)$ and $S(y_1, \dots, y_n)$ is allowed only if $\{x_1, \dots, x_k\} = \{y_1, \dots, y_n\}$. Boolean combinations of formulas with at most one free variable can be formed freely, and the use of equality is unrestricted.

It is established in [8] that already the equality-free fragment UF_1 of UF_1^- can define properties not expressible in FOC^2 and also properties not expressible in the recently introduced *guarded negation fragment* [2], which significantly generalizes the *guarded fragment* [1]. The article [8] also shows, inter alia, that the equality-free logic UF_1 is decidable, and furthermore, that minor modifications to the syntax of UF_1 lead to undecidable formalisms. Namely, the *uniform two-dimensional* and *non-uniform one-dimensional* fragments are shown undecidable.

In this article we establish that the satisfiability and finite satisfiability problems of the uniform one-dimensional fragment with equality (UF_1^-) are NEXPTIME-complete. These results are obtained by appropriately generalizing and modifying the construction in [6] that provides small models for satisfiable FO^2 -formulas in *Scott normal form*. The NEXPTIME-completeness of FOC^2 raises the natural question whether the extension UFC_1^- of UF_1^- with *counting quantifiers* remains decidable. We answer this question in the negative by showing that the satisfiability and finite satisfiability problems of UFC_1^- are complete for Π_1^0 and Σ_1^0 , respectively. These results are established by tiling arguments that make an appropriate use of a *ternary* relation, together with the usual unary relations commonly employed in similar undecidability proofs.

We also study the expressivity of UF_1^- . We establish that while UF_1^- and FOC^2 are incomparable in expressivity in general, in restriction to vocabularies with the arity bound two, we have $\text{UF}_1^- < \text{FOC}^2$.

The uniform one-dimensional fragment UF_1^- canonically extends FO^2 , and in fact the equality-free fragments of UF_1^- and FO^2 coincide when attention is limited to binary vocabularies. We believe that UF_1^- is an interesting fragment that can be used in order to *extend the scope of research on two-variable logics to the realm involving relations of arbitrary arities*.

2 Preliminaries

Let m and $n \geq m$ be integers. We let $[m, n]$ denote the set of integers i such that $m \leq i \leq n$. If φ and ψ are first-order formulas, then $\varphi \equiv \psi$ indicates that the formulas are equivalent. If \mathcal{L} and \mathcal{L}' are fragments of first-order logic, we write $\mathcal{L} \leq \mathcal{L}'$ to indicate that for every sentence of \mathcal{L} , there exists an equivalent sentence of \mathcal{L}' . We let $\text{VAR} := \{v_i \mid i \in \mathbb{N}\}$ denote the set of first-order variable symbols. We mostly use *metavariables* x, y, z, x_i, y_i, z_i , etc., in order to refer to symbols in VAR . Notice that for example x_1 and x_2 may denote the same variable in VAR , while v_1 and v_2 are necessarily different variables. The set of free variables of a formula ψ is denoted by $\text{free}(\psi)$.

Let $X = \{x_1, \dots, x_n\}$ be a finite set of variable symbols. Let R be a k -ary relation symbol. An atomic formula $R(x_{i_1}, \dots, x_{i_k})$ is called an X -atom if $\{x_{i_1}, \dots, x_{i_k}\} = X$. A finite set of X -atoms is an X -uniform set. When X is irrelevant or known from the context, we may simply talk about a *uniform set*. For example, if x, y, z are distinct variables, then $\{T(x, y), S(y, x)\}$ and $\{R(x, x, y), R(y, y, x), S(y, x)\}$ are uniform sets, while $\{R(x, y, z), R(x, y, y)\}$ and $\{S(x, y), x = y\}$ are not (uniform sets are not allowed to contain equality atoms). The empty set is an X -uniform set for every finite subset of VAR , including \emptyset .

Let \mathbb{Z}_+ denote the set of positive integers. Let \mathcal{V} denote a *complete relational vocabulary*, i.e., $\mathcal{V} := \bigcup_{k \in \mathbb{Z}_+} \tau_k$, where τ_k denotes a countably infinite set of k -ary relation symbols. Every vocabulary τ we consider below is assumed to be a subset of \mathcal{V} . A k -ary τ -atom is an atomic τ -formula ψ such that $|\text{free}(\psi)| = k$. For example, if $P \in \tau$ is a unary and $R \in \tau$ a binary symbol, then $P(x), x = x, R(x, x)$ are unary τ -atoms, and $R(v_1, v_2), v_1 = v_2$ are binary τ -atoms. If τ is known from the context or irrelevant, we may simply talk about k -ary atoms.

Let $\tau \subseteq \mathcal{V}$. The set $\text{UF}_1^-(\tau)$, or the set of τ -formulas of the uniform one-dimensional fragment, is the smallest set \mathcal{F} satisfying the following conditions.

1. Every unary τ -atom is in \mathcal{F} . Also $\perp, \top \in \mathcal{F}$.

2. Every identity atom $x = y$ is in \mathcal{F} .
3. If $\varphi \in \mathcal{F}$, then $\neg\varphi \in \mathcal{F}$. If $\varphi_1, \varphi_2 \in \mathcal{F}$, then $(\varphi_1 \wedge \varphi_2) \in \mathcal{F}$.
4. Let $X = \{x_0, \dots, x_k\} \subseteq \text{VAR}$. Let U be a finite set of formulas $\psi \in \mathcal{F}$ whose free variables are in X . Let $V \subseteq X$. Let F be a V -uniform set of τ -atoms. Let φ be any Boolean combination of formulas in $U \cup F$. Then $\exists x_1 \dots \exists x_k \varphi \in \mathcal{F}$ and $\exists x_0 \dots \exists x_k \varphi \in \mathcal{F}$.

Let UF_1^- denote the set $\text{UF}_1^-(\mathcal{V})$.

Let \bar{x} denote a tuple of variables, and let $\chi := \exists \bar{x} \varphi$ be a UF_1^- -formula formed by using the rule 4 above. Assume that φ is quantifier-free. Then we call φ a UF_1^- -matrix. If φ does not contain k -ary atoms for any $k \geq 2$, with the possible exception of equality atoms $x = y$, then we define $S_\varphi := \emptyset$. Otherwise we define S_φ to be the set V used in the construction of χ (see rule 4). The set S_φ is the set of *live variables* of φ .

Let $\psi(x_0, \dots, x_k)$ be a UF_1^- -matrix, where (x_0, \dots, x_k) enumerates the variables of ψ . Let \mathfrak{A} be a structure. Let $a_0, \dots, a_k \in A$, where A is the domain of \mathfrak{A} . We let $\text{live}(\psi(x_0, \dots, x_k)[a_0, \dots, a_k])$ denote the set $T \subseteq \{a_0, \dots, a_k\}$ such that $a_i \in T$ iff x_i is a live variable of $\psi(x_0, \dots, x_k)$. We may write $\text{live}(\psi[a_0, \dots, a_k])$ instead of $\text{live}(\psi(x_0, \dots, x_k)[a_0, \dots, a_k])$ when no confusion can arise. Notice that since the elements a_i are not required to be distinct, it is possible that $|\text{live}(\psi[a_0, \dots, a_k])|$ is smaller than the number of live variables in ψ .

Let $\tau \subseteq \mathcal{V}$ be a finite vocabulary. A *1-type* over the vocabulary τ is a maximal satisfiable set of literals (atoms and negated atoms) over τ with the variable v_1 . The set of all 1-types over τ is denoted by $\alpha[\tau]$, or just by α when τ is clear from the context. We identify 1-types α and conjunctions $\bigwedge \alpha$. A *k-table* over τ is a maximal satisfiable set of $\{v_1, \dots, v_k\}$ -atoms and negated $\{v_1, \dots, v_k\}$ -atoms over τ . Recall that a $\{v_1, \dots, v_k\}$ -atom must contain exactly all the variables in $\{v_1, \dots, v_k\}$, and note that a 2-table does not contain equality formulas or negated equality formulas. We identify k -tables β and conjunctions $\bigwedge \beta$.

Let \mathfrak{A} be a τ -structure, and let $a \in A$. Let α be a 1-type over τ . We say that a *realizes* α if α is the unique 1-type such that $\mathfrak{A} \models \alpha[a]$. We let $\text{tp}_{\mathfrak{A}}(a)$ denote the 1-type realized by a . Similarly, for *distinct* elements $a_1, \dots, a_k \in A$, we let $\text{tb}_{\mathfrak{A}}(a_1, \dots, a_k)$ denote the unique k -table *realized* by the tuple (a_1, \dots, a_k) , i.e., the k -table $\beta(v_1, \dots, v_k)$ such that $\mathfrak{A} \models \beta[a_1, \dots, a_k]$. Note that we have $\text{tp}_{\mathfrak{A}}(a) \equiv \text{tb}_{\mathfrak{A}}(a)$ for every $a \in A$.

Let m be the maximum arity of symbols in τ . We observe that to fully define a τ -structure \mathfrak{A} over a known domain A , it is sufficient to consider each set $B \subseteq A$, $|B| \leq m$, and first choose an enumeration $(b_1, \dots, b_{|B|})$ of the elements of B , and then specify $\text{tb}_{\mathfrak{A}}(b_1, \dots, b_{|B|})$.

Observation 2.1. Let $\psi(x_1, \dots, x_k)$ be a UF_1^- -matrix, where (x_1, \dots, x_k) enumerates the variables in ψ . Let \mathfrak{A} be a τ -structure, where τ is the set of

relation symbols in ψ . Let $a_1, \dots, a_k \in A$ be a sequence of (not necessarily distinct) elements. Whether or not $\mathfrak{A} \models \psi[a_1, \dots, a_k]$ holds, depends only on (i) the 1-types of the elements a_i , (ii) the list of pairs (a_i, a_j) such that $a_i = a_j$, and (iii) the table $\text{tb}_{\mathfrak{A}}(b_1, \dots, b_l)$, where (b_1, \dots, b_l) is an arbitrary enumeration of $\text{live}(\psi[a_1, \dots, a_k])$.

3 Complexity of UF_1^-

We now introduce a normal form for UF_1^- inspired by the *Scott normal form* for FO^2 [15]. We say that a UF_1^- -formula φ is in *generalized Scott normal form* if

$$\varphi = \bigwedge_{1 \leq i \leq m_{\exists}} \forall x \exists y_1 \dots y_{k_i} \varphi_i^{\exists}(x, y_1, \dots, y_{k_i}) \wedge \bigwedge_{1 \leq i \leq m_{\forall}} \forall x_1 \dots x_{l_i} \varphi_i^{\forall}(x_1, \dots, x_{l_i}), \quad (1)$$

where formulas φ_i^{\exists} and φ_i^{\forall} are quantifier-free UF_1^- -matrices.

Proposition 3.1. Each UF_1^- -formula φ translates in polynomial time to a UF_1^- -formula φ' in generalized Scott normal form (over a signature extended by some fresh unary symbols) such that φ and φ' are satisfiable over the same domains.

Proof. A simple adaptation of the well-known translation given, e.g., in [5]. \square

Let φ be a UF_1^- -formula in generalized Scott-normal form given in Equation 1. Assume $\mathfrak{A} \models \varphi$. We will build a small τ -model \mathfrak{A}' of φ , where τ is the set of relation symbols in φ . Our construction modifies and generalizes the construction of a small model for a satisfiable FO^2 -formula in Scott normal form from [6]. Let $a \in A$ and $b_1, \dots, b_{k_i} \in A$ be elements such that $\mathfrak{A} \models \varphi_i^{\exists}[a, b_1, \dots, b_{k_i}]$. We say that the structure $\mathfrak{B} := \mathfrak{A} \upharpoonright \{a, b_1, \dots, b_{k_i}\}$ is a *witness structure* for a and φ_i^{\exists} . The substructure of \mathfrak{B} restricted to the elements in $\text{live}(\varphi_i^{\exists}[a, b_1, \dots, b_{k_i}])$ is called the *live part* of \mathfrak{B} . If the live part of \mathfrak{B} does not contain a , then the live part is called *free*. Note that $|B|$ may be smaller than $k_i + 1$ (this may be even imposed by the use of equalities). Also, a may be a member of the live part of \mathfrak{B} even if the variable x is not a live variable of φ_i^{\exists} .

The court. Let n be the width of φ , i.e., $n = \max(\{k_i + 1\}_{1 \leq i \leq m_{\exists}} \cup \{l_i\}_{1 \leq i \leq m_{\forall}})$. We assume, w.l.o.g., that $n \geq 2$. A 1-type α realized in \mathfrak{A} is *royal* if it is realized at most $n - 1$ times in \mathfrak{A} . The points in A that realize a royal 1-type are called *kings*. Let K be the set of all kings in \mathfrak{A} . Clearly $|K| \leq (n - 1)|\alpha|$

We then define a set $D \subseteq A$. For each pair $(\alpha, \varphi_i^{\exists})$, where α is a 1-type realized in \mathfrak{A} , if it is possible, select an element $a \in A$ that realizes

the 1-type α such that there exists a witness structure $\mathfrak{B}_{\alpha,i}$ for a and φ_i^{\exists} whose live part $\bar{\mathfrak{B}}_{\alpha,i}$ is free. Add the elements of $\bar{B}_{\alpha,i}$ to D . Since we add at most $n - 1$ elements for each pair $(\alpha, \varphi_i^{\exists})$, the total size of D is bounded by $(n - 1)m_{\exists}|\alpha|$.

For each $a \in K \cup D$ and each φ_i^{\exists} , select a witness structure in \mathfrak{A} and let $C_{a,i}$ denote its universe. Define $C := K \cup D \cup \bigcup_{a \in K \cup D, 1 \leq i \leq m_{\exists}} C_{a,i}$. We call $\mathfrak{C} := \mathfrak{A} \upharpoonright C$ the *court* of \mathfrak{A} . Note that $|C| \leq n|K \cup D| \leq n((n - 1) + (n - 1)m_{\exists})|\alpha|$. We have $|C| \leq 2|\varphi|^3 2^{|\varphi|}$.

Universe. The court \mathfrak{C} of \mathfrak{A} will be a substructure of \mathfrak{A}' . The remaining part of the universe of \mathfrak{A}' consists of three fresh disjoint sets E, F, G . Each of them contains $m_{\exists} + n$ elements of type α for each non-royal α realized in \mathfrak{A} . The i -th element of type α ($1 \leq i \leq m_{\exists} + n$) in E (resp. F, G) is denoted $e_{\alpha,i}$ (resp. $f_{\alpha,i}, g_{\alpha,i}$). The size of each set E, F, G is bounded by $(n + m_{\exists})|\alpha| \leq 2|\varphi|2^{|\varphi|}$. Thus the total size of $|A'|$ is bounded by $8|\varphi|^3 2^{|\varphi|}$, which is exponential in $|\varphi|$.

Witnesses. Our next aim is to provide witness structures for each element of $a \in A' \setminus (K \cup D)$ and each φ_j^{\exists} . We will choose elements in A' which will form the universe (say, of size s) of the live part of a witness structure for a and φ_j^{\exists} and define the s -table on these elements. The remaining elements of the witness structure (elements not in the live part) will then be very easily found in A' .

Let $a' \in A' \setminus (K \cup D)$. We find a *pattern element* $a \in A$ of a' as follows. If $a' \in C$, then the pattern element is a' itself. If $a' \in E \cup F \cup G$, then we let an arbitrary $a \in A$ such that $\text{tp}_{\mathfrak{A}}(a) = \text{tp}_{\mathfrak{A}'}(a')$ be the pattern element of a' . For each $1 \leq j \leq m_{\exists}$, we find a witness structure $\mathfrak{B}_{a,j}$ for a and φ_j^{\exists} , and let $\bar{\mathfrak{B}}_{a,j}$ be its live part. If this live part is free, then there is nothing to do; an appropriate live part for the witness structure of a' and φ_j^{\exists} already exists in $\mathfrak{D} := \mathfrak{C} \upharpoonright D$. Otherwise, let r_1, \dots, r_k be the kings included in $\bar{\mathfrak{B}}_{a,j}$ (possibly $k = 0$), and let a, b_1, \dots, b_l be the non-royal elements of $\bar{\mathfrak{B}}_{a,j}$ (possibly $l = 0$). Let α_i be the 1-type of b_i ($1 \leq i \leq l$). We consider the following cases.

Case 1. If $l = 0$ and $a' \in C$, then there is nothing to do; a' forms the live part of the desired witness structure together with some elements in K .

Case 2. If $l = 0$ and $a' \notin C$, then we set $\text{tb}_{\mathfrak{A}'}(a', r_1, \dots, r_k) := \text{tb}_{\mathfrak{A}}(a, r_1, \dots, r_k)$.

Case 3. If $l > 0$ and $a' \in E$, then we define $b'_1 := f_{\alpha_1,j}$ and choose b'_2, \dots, b'_l to be distinct elements of types $\alpha_2, \dots, \alpha_l$ from $S := \{f_{\alpha,s} : m_{\exists} + 1 \leq s \leq m_{\exists} + n, \alpha \text{ non-royal}\}$. This is possible since $l < n$ and S contains n realizations of each non-royal 1-type. We set

$$\text{tb}_{\mathfrak{A}'}(a', r_1, \dots, r_k, b'_1, \dots, b'_l) := \text{tb}_{\mathfrak{A}}(a, r_1, \dots, r_k, b_1, \dots, b_l).$$

Case 4. If $l > 0$ and $a' \in F$ (resp. $a' \in G \cup (C \setminus (K \cup D))$), then we proceed as in the previous case, but we take the elements b'_i from G (resp. E).

Claim 3.2. The described procedure of providing live parts of witness structures can be executed without conflicts.

Proof. Consider two tuples of elements $t' = a', r'_1, \dots, r'_{k'}, b'_1, \dots, b'_{l'}$ and $t'' = a'', r''_1, \dots, r''_{k''}, b''_1, \dots, b''_{l''}$ that form in the described process the live parts of witness structures for, respectively, a', φ_i^{\exists} and a'', φ_j^{\exists} .

If $a' \neq a''$, then the two tuples differ, since our strategy of using the sets E, F, G ensures that a' cannot be a member of t'' and a'' cannot be a member of t' .

If $a' = a'', i \neq j$ and $l' > 0$, then the two tuples differ since b'_1 cannot be a member of t'' , due to the following reason. Since b'_1 is selected to be the i -th realization of some α from E, F , or G , it must differ from b''_1 , which is the j -th realization of some α . Furthermore, b'_i must differ from each b''_m , $m > 1$, since each of them is the s -th realization of some α for some $s > m_{\exists}$, while $i \leq m_{\exists}$. The case where $a' = a'', i \neq j$ and $l'' > 0$ is symmetric.

If $a' = a'', i \neq j$ and $l' = l'' = 0$, then it is possible that the two tuples are identical, but in this case they contain only a and some kings $r'_1, \dots, r'_{k'}$, and even though $\text{tb}_{\mathfrak{A}'}(a', r'_1, \dots, r'_{k'})$ is defined twice, it is done precisely the same way both times (setting $\text{tb}_{\mathfrak{A}'}(a', r'_1, \dots, r'_{k'}) := \text{tb}_{\mathfrak{A}'}(c, r'_1, \dots, r'_{k'})$ both times, where c is the pattern element of a'). \square

It is probably worth commenting why we prepared free live parts of witness structures in \mathfrak{D} instead of building them using elements of $E \cup F \cup G$ in a “regular” way. One of the problematic situations arises, e.g., when an element a' from, say, E builds the live part of its witness structure for some φ_i^{\exists} using an element $b' \in F$ and some kings r_1, \dots, r_k . In this case $\text{tb}_{\mathfrak{A}'}(b', r_1, \dots, r_k)$ is defined. However, it may happen that b' needs to form the live part of its witness structure for some φ_j^{\exists} using precisely the elements b', r_1, \dots, r_k , which can lead to a conflict.

Completion. Let a'_1, \dots, a'_k ($a'_i \neq a'_j$ for $i \neq j$, $1 < k \leq n$) be elements in A' such that the table $\text{tb}_{\mathfrak{A}'}(a'_1, \dots, a'_k)$ has not yet been defined. Select distinct elements a_1, \dots, a_k of \mathfrak{A} such that $\text{tp}_{\mathfrak{A}}(a_i) = \text{tp}_{\mathfrak{A}'}(a'_i)$ ($1 \leq i \leq k$). This is always possible due to our strategy of not introducing extra kings. Set $\text{tb}_{\mathfrak{A}'}(a'_1, \dots, a'_k) := \text{tb}_{\mathfrak{A}}(a_1, \dots, a_k)$.

Claim 3.3. $\mathfrak{A}' \models \varphi$.

Proof. Let us first argue that all existential requirements are fulfilled. Consider an arbitrary element $a' \in A$ and a conjunct $\forall x \exists y_1 \dots y_{k_i} \varphi_i^{\exists}(x, y_1, \dots, y_{k_i})$ of φ . If $a' \in K \cup D$, then the witness structure for a' and φ_i^{\exists} is provided in \mathfrak{C} by the elements of $C_{a', i}$. If $a' \in (C \setminus (K \cup D)) \cup E \cup F \cup G$, then the live part of a witness structure is either secured in \mathfrak{D} or according to one of the four cases in the step *Witnesses*. We consider the case where the live part is not in \mathfrak{D} (the arguments involving \mathfrak{D} are similar). Assume that the pattern

element for a' is a and the witness structure for a and φ_i^{\exists} is $\mathfrak{B}_{a,i}$, where $B_{a,i} = \{a, r_1, \dots, r_{k'}, b_1, \dots, b_{l'}\}$. Here the elements r_i are kings and the elements b_i non-kings. Assume that the universe of the live part of $\mathfrak{B}_{a,i}$ is $\bar{B}_{a,i} = \{a, r_1, \dots, r_k, b_1, \dots, b_l\}$, $l \leq l'$, $k \leq k'$. (Note that in the case we are considering, a is necessarily a member of $\bar{B}_{a,i}$). A witness structure for a' can now be formed using the elements $a', r_1, \dots, r_{k'}$, the elements b'_1, \dots, b'_l (if any) chosen according to *Case 3* or *Case 4*, and elements $b'_{l+1}, \dots, b'_{l'}$ from, say, the set E , whose 1-types are equal to the 1-types of $b_{l+1}, \dots, b_{l'}$. Note that we always have sufficiently many elements with appropriate 1-types in E . Due to Observation 2.1, the substructure $\mathfrak{A}' \upharpoonright \{a', r_1, \dots, r_{k'}, b'_1, \dots, b'_{l'}\}$ is indeed an appropriate witness structure.

Let us then argue that the universal conjuncts of φ are satisfied. Consider a conjunct $\forall x_1 \dots x_k \varphi_i^{\forall}(x_1, \dots, x_k)$ and a tuple of (not necessarily distinct) elements a'_1, \dots, a'_k from A' . We must show that $\mathfrak{A}' \models \varphi_i^{\forall}[a'_1, \dots, a'_k]$. Let b'_1, \dots, b'_l be the elements of $\text{live}(\varphi_i^{\forall}[a'_1, \dots, a'_k])$. There are three different ways in which $\text{tb}_{\mathfrak{A}'}(b'_1, \dots, b'_l)$ can become defined: if $l = 1$, then $\text{tb}_{\mathfrak{A}'}(b'_1, \dots, b'_l)$ is just a 1-type, and the table is fixed in step *Universe*, and if $l > 1$, then the table is fixed either according to step *Witnesses* or step *Completion*. In each of the cases there are distinct elements $b_1, \dots, b_l \in A$ such that $\text{tp}_{\mathfrak{A}}(b_i) = \text{tp}_{\mathfrak{A}'}(b'_i)$ for $1 \leq i \leq l$ and $\text{tb}_{\mathfrak{A}'}(b'_1, \dots, b'_l) = \text{tb}_{\mathfrak{A}}(b_1, \dots, b_l)$. Let $b'_1, \dots, b'_l, c'_1, \dots, c'_s$ be a list of all distinct elements of $\{a'_1, \dots, a'_k\}$. It is readily verified that the tuple b_1, \dots, b_l can be extended by elements c_1, \dots, c_s to a tuple of distinct elements of A with $\text{tp}_{\mathfrak{A}}(c_i) = \text{tp}_{\mathfrak{A}'}(c'_i)$ for $1 \leq i \leq s$. Thus we can form a sequence of (not necessarily distinct) elements a_1, \dots, a_k of A corresponding to a'_1, \dots, a'_k in such a way that $a'_i = b'_j \Leftrightarrow a_i = b_j$ and $a'_i = c'_j \Leftrightarrow a_i = c_j$. Obviously $\mathfrak{A} \models \varphi_i^{\forall}[a_1, \dots, a_k]$. By Observation 2.1, it follows that $\mathfrak{A}' \models \varphi_i^{\forall}[a'_1, \dots, a'_k]$. \square

We have proved the following theorem.

Theorem 3.4. UF_1^- has the finite model property. Moreover, every satisfiable UF_1^- formula φ has a model whose size is bounded exponentially in $|\varphi|$.

We can now prove the following theorem.

Theorem 3.5. The satisfiability problem (= finite satisfiability problem) for UF_1^- is NEXPTIME-complete.

Proof. The lower bound follows from the NEXPTIME lower bound for FO^2 from [11]. For the upper bound, we translate the input formula φ to an equisatisfiable formula φ' in generalized Scott normal form. Suppose φ' is satisfiable. We guess an exponentially bounded model \mathfrak{A} of φ' (note that not only the universe of \mathfrak{A} is bounded exponentially, but also the description of \mathfrak{A} , since we are dealing only with at most $|\varphi'|$ relations of arity at most $|\varphi'|$), and verify that it is indeed a model. The last task can be carried out

in an exhaustive way: for each $a \in A$ and each φ_i^{\exists} , guess which elements form a witness structure and check that they indeed form a required witness structure; for each conjunct $\forall x_1, \dots, x_k \varphi_i^{\forall}(x_1, \dots, x_k)$, enumerate all tuples a_1, \dots, a_k of A and check that $\mathfrak{A} \models \varphi_i^{\forall}[a_1, \dots, a_k]$. \square

4 Expressivity

In this section we compare the expressivity of $\text{UF}_1^=$ with the expressivities of FO^2 and FOC^2 . Clearly $\text{UF}_1^=$ contains FO^2 , and it is not hard to see that the inclusion is strict; equalities can be used freely in $\text{UF}_1^=$, and for example the property that there are precisely two elements in a unary relation P is expressible in $\text{UF}_1^=$ but not in FO^2 . The expressivities of $\text{UF}_1^=$ and FOC^2 are related as follows.

Theorem 4.1. $\text{UF}_1^=$ and FOC^2 are incomparable in expressivity.

Proof. It is straightforward to establish that FOC^2 cannot express the $\text{UF}_1^=$ -sentence $\exists x \exists y \exists z R(x, y, z)$, and therefore $\text{UF}_1^= \not\leq \text{FOC}^2$. To show that $\text{FOC}^2 \not\leq \text{UF}_1^=$, let R be a binary relation symbol and consider models over the signature $\{R\}$. We claim that $\text{UF}_1^=$ cannot express the FOC^2 -definable condition that the in-degree (w.r.t. the relation R) at every node is at most one. Assume $\varphi(R)$ is a $\text{UF}_1^=$ -formula that defines the condition. Consider the conjunction $\varphi(R) \wedge \forall x \exists y R(x, y) \wedge \exists x \forall y \neg R(y, x)$. It is easy to see that this formula does not have a finite model, and thereby the assumption that $\text{UF}_1^=$ can express $\varphi(R)$ is false. \square

The rest of this section is devoted to the scenario in which the signature contains only unary and binary relation symbols. We will show that in such a case the expressivity of $\text{UF}_1^=$ lies strictly between FO^2 and FOC^2 .

Let τ be a finite relational vocabulary. Let β be a 2-table over τ , and let x and y be *distinct* variables. Let S be the set of atoms obtained from β by replacing all occurrences of the variables v_1 and v_2 in β by x and y , respectively. We call S a *binary τ -diagram* in the variables (x, y) , and denote it by $\beta(x, y)$. We identify binary diagrams and conjunctions over them. A *binary τ -arrow* in the variables x, y is an atomic formula $R(x, y)$ (or $R(y, x)$), where $R \in \tau$. Notice that neither equality statements nor atoms of the form $R(x, x), R(y, y)$ are binary τ -arrows. It is straightforward to show that if φ is a Boolean combination of binary τ -arrows in the variables x, y , then φ is equivalent to the disjunction of τ -diagrams $\beta(x, y)$ that entail φ , i.e., $\beta(x, y) \models \varphi$. Notice that $\bigvee \emptyset = \perp$ is a legitimate disjunction of diagrams.

Let $\{x_0, \dots, x_k\}$ be a (possibly empty) set of distinct variables. An *identity literal* over $\{x_0, \dots, x_k\}$ is a formula of the type $x_i = x_j$ or $\neg x_i = x_j$, where $i, j \in [0, k]$. An identity literal is *non-trivial* if the variables in it are different. An *identity profile* over $\{x_0, \dots, x_k\}$, or a $\{x_0, \dots, x_k\}$ -*profile*, is a maximal satisfiable set of non-trivial identity literals over $\{x_0, \dots, x_k\}$. We

identify identity profiles and conjunctions over them. We let $\text{diff}(x_0, \dots, x_k)$ denote the conjunction of inequalities $x_i \neq x_j$, where $i, j \in [0, k]$, $i \neq j$. An identity profile is a *discriminate profile* if it is the formula $\text{diff}(x_0, \dots, x_k)$ for some set $\{x_0, \dots, x_k\}$ of distinct variables. Let I be a set of identity literals over $\{x_0, \dots, x_k\}$. Let φ be a $\{x_0, \dots, x_k\}$ -profile. We say that φ is *consistent with I* if $\varphi \models \bigwedge I$.

A UF_1^- -formula φ is a *block formula* if φ is of the type $\exists \bar{x} \psi$ or $\neg \exists \bar{x} \psi$. Here $\exists \bar{x}$ denotes a vector of one or more existentially quantified variables. Formulas $\exists \bar{x} \psi$ are called *positive blocks*, while formulas $\neg \exists \bar{x} \psi$ are *negative blocks*. A UF_1^- -formula is *simple* if it is a literal or a block formula.

Let τ be a finite relational vocabulary with the arity bound two. Let x_0, \dots, x_k be distinct variable symbols. Let $\varphi := \exists x_1 \dots x_k \psi$ be a UF_1^- -formula over τ . Let $x, y \in \{x_0, \dots, x_k\}$ be distinct variables. We call φ a τ -*diagram block* if the formula ψ is a conjunction

$$\beta(x, y) \wedge \text{diff}(x_0, \dots, x_k) \wedge \psi_0(x_0) \wedge \dots \wedge \psi_k(x_k),$$

where $\beta(x, y)$ is a binary τ -diagram in the variables (x, y) , and every formula $\psi_i(x_i)$ is a conjunction of simple formulas ψ' such that $\text{free}(\psi') \subseteq \{x_i\}$. Furthermore, if $\chi(x)$ is a τ -diagram block with the free variable x , then also the formula $\exists x \chi(x)$ is a τ -diagram block. A UF_1^- -formula φ is said to be in *diagram normal form* if for every positive block formula φ' that occurs as a subformula in φ , there is a τ such that φ' is a τ -diagram block.

Lemma 4.2. Every positive block formula is equivalent to a disjunction of diagram blocks.

Proof. Consider a block formula $\varphi(x_0) := \exists x_1 \dots \exists x_k \psi$ with the free variable x_0 . Let τ be the set of relation symbols that occur in ψ . First put ψ into disjunctive normal form so that we obtain the formula $\chi_1 \vee \dots \vee \chi_l$, where each disjunct χ_i is a conjunction of simple formulas. We have

$$\varphi(x_0) \equiv \varphi'(x_0) := \exists x_1 \dots \exists x_k (\chi_1 \vee \dots \vee \chi_l).$$

Now distribute the existential quantifier prefix of $\varphi'(x_0)$ over the disjunctions. We have

$$\varphi(x_0) \equiv \varphi''(x_0) := \bigvee_{i \in \{1, \dots, l\}} \exists x_1 \dots \exists x_k \chi_i. \quad (2)$$

Consider an arbitrary disjunct $\exists x_1 \dots \exists x_k \chi_i$ of $\varphi''(x_0)$. Recall that χ_i is a conjunction simple formulas. Let I be the set of conjuncts of χ_i that are non-trivial identity literals. Let Π be the set of all $\{x_0, \dots, x_k\}$ -profiles consistent with I . Let J be the set of conjuncts of χ_i that are not in I . Thus

$$\chi_i \equiv \chi'_i := \bigvee_{\pi \in \Pi} (\pi \wedge \bigwedge J). \quad (3)$$

Write the formula $\bigwedge J$ in the form $\alpha \wedge \beta$, where α is a conjunction of formulas with at most one free variable and β is a conjunction of literals with two free variables. Notice that due to the uniformity condition, β has exactly two free variables, or, alternatively, β is the formula $\bigwedge \emptyset = \top$. Let $\beta_1 \vee \dots \vee \beta_m$ be the disjunction of τ -diagrams that entail β . Thus

$$\bigwedge J \equiv \alpha \wedge \bigvee_{j \in [1, m]} \beta_j. \quad (4)$$

Combining Equations 3 and 4, we infer that

$$\chi_i \equiv \bigvee_{\substack{\pi \in \Pi \\ j \in [1, m]}} (\pi \wedge \alpha \wedge \beta_j).$$

Therefore we have

$$\exists x_1 \dots \exists x_k \chi_i \equiv \chi_i'' := \bigvee_{\substack{\pi \in \Pi \\ j \in [1, m]}} \exists x_1 \dots \exists x_k (\pi \wedge \alpha \wedge \beta_j). \quad (5)$$

Consider an arbitrary disjunct $\exists x_1 \dots \exists x_k (\pi \wedge \alpha \wedge \beta_j)$ of χ_i'' . The profile formula π may contain unnegated literals. We want to get rid of them. If π contains an unnegated identity $x = y$ as a conjunct, we get rid of the variable y altogether by erasing the conjunct $x = y$ and renaming variables in the remaining conjuncts of $\pi \wedge \alpha \wedge \beta_j$. The renaming is done such that x_0 is never erased. By renaming variables in this fashion, we obtain a formula $\pi' \wedge \alpha' \wedge \beta'_j$ such that

$$\exists x_1 \dots \exists x_k (\pi \wedge \alpha \wedge \beta_j) \equiv \exists y_1 \dots \exists y_n (\pi' \wedge \alpha' \wedge \beta'_j), \quad (6)$$

where $\{y_1, \dots, y_n\} \subseteq \{x_1, \dots, x_k\}$, π' is a discriminate $\{y_1, \dots, y_n\}$ -profile, α' is a conjunction of simple formulas with at most one free variable, and β'_j is a binary τ -diagram or the formula \top . By combining Equations 2, 5 and 6, it is easy to see that the original formula $\varphi(x_0)$ is indeed equivalent to some disjunction of diagram blocks.

We still need to show that positive block formulas *without* free variables are equivalent to disjunctions of diagram blocks. Consider a formula $\exists x_0 \dots \exists x_k \chi$. We assume, w.l.o.g., that $k > 0$ and that the variable x_0 occurs free in χ . We translate the formula $\varphi'(x_0) := \exists x_1 \dots \exists x_k \chi$ in the way described above to a disjunction $\varphi'_1 \vee \dots \vee \varphi'_n$ of diagram blocks. We then observe that $\exists x_0 \varphi'_1 \vee \dots \vee \exists x_0 \varphi'_n$ is equivalent to the original formula $\exists x_0 \dots \exists x_k \chi$. \square

Corollary 4.3. Each UF_1^- -formula is equivalent to a formula in diagram normal form.

Proof. By induction on the structure of formulas, using Lemma 4.2. \square

Let τ be a finite relational vocabulary with the arity bound 2. Let $k \in \mathbb{Z}_+$, and let x_0, \dots, x_k be distinct variable symbols. Let $\varphi(x_0, \dots, x_k) := \text{diff}(x_0, \dots, x_k) \wedge \psi$, where ψ is conjunction of τ -literals such that the following conditions hold.

1. The variables of each conjunct of ψ are in $\{x_0, \dots, x_k\}$.
2. If ψ has $R(x, y)$ or $\neg R(x, y)$ as a conjunct, where $R \in \tau$ is a binary relation symbol and x, y distinct variables, then $x_0 \in \{x, y\}$.

Then we call $\varphi(x_0, \dots, x_k)$ a τ -star formula in the variables (x_0, \dots, x_k) . The variable x_0 is called the *centre variable* of φ . Consider then a quantifier-free τ -formula $\psi(x_0, \dots, x_k) := \text{diff}(x_0, \dots, x_k) \wedge \beta \wedge \alpha$ such that the following conditions are satisfied.

1. The formula β is a conjunction $\beta_1(x_0, x_1) \wedge \dots \wedge \beta_k(x_0, x_k)$, where each $\beta_i(x_0, x_i)$ is a binary τ -diagram in the variables (x_0, x_i) .
2. The formula α is a conjunction $\alpha_0(x_0) \wedge \dots \wedge \alpha_k(x_k)$, where each $\alpha_i(x_i)$ is a 1-type over τ and in the variable x_i . (The variable v_1 is replaced by x_i .)

The formula $\psi(x_0, \dots, x_k)$ is called a τ -star type in the variables (x_0, \dots, x_k) . The variable x_0 is the *centre variable* of the τ -star type. It is straightforward to show that every τ -star formula in the variables (x_0, \dots, x_k) is equivalent to a disjunction of τ -star types in the variables (x_0, \dots, x_k) .

Let $\varphi(x_0, \dots, x_k)$ be a τ -star formula in the variables (x_0, \dots, x_k) . Then the formula $\exists x_1 \dots \exists x_k \varphi(x_0, \dots, x_k)$ is called a τ -star centre formula of the width k . Let $\psi(x_0, \dots, x_k)$ be a τ -star type in the variables (x_0, \dots, x_k) . Then the formula $\exists x_1 \dots \exists x_k \psi(x_0, \dots, x_k)$ is called a τ -star centre type of the width k . The following lemma follows immediately from the fact that every τ -star formula is equivalent to a disjunction of τ -star types.

Lemma 4.4. Every τ -star centre formula of the width k is equivalent to a disjunction of τ -star centre types of the width k .

A *2-type* over τ is a maximal satisfiable set of τ -literals in the variables v_1 and v_2 (equalities and negated equalities are considered to be τ -literals). If \mathcal{T} is a 2-type over τ and x, y distinct variables, we let $\mathcal{T}(x, y)$ denote the set obtained from \mathcal{T} by replacing all occurrences of v_1 and v_2 in \mathcal{T} by x and y , respectively. Below we identify sets $\mathcal{T}(x, y)$ and conjunctions over them.

Let ψ be a τ -star type in the variables (x_0, \dots, x_k) . Each pair (x_0, x_i) , where $i \in [1, k]$, is called a *ray* of ψ . Let \mathcal{T} be a 2-type over τ . We say that the ray (x_0, x_i) of ψ *realizes* \mathcal{T} if $\psi \models \mathcal{T}(x_0, x_i)$.

Theorem 4.5. Let τ be a relational vocabulary with the arity bound 2. Then $\text{UF}_1^-(\tau) \leq \text{FOC}^2(\tau)$. The inclusion is strict if τ contains a binary symbol.

Proof. We have above shown that $\text{FOC}^2(\tau) \not\leq \text{UF}_1^-(\tau)$ if τ contains a binary relation symbol. Therefore it suffices to show that $\text{UF}_1^-(\tau) \leq \text{FOC}^2(\tau)$. The claim is established by induction on the structure of UF_1^- -formulas in diagram normal form. We discuss the case involving quantifiers.

Let $\sigma \subseteq \tau$ and consider a σ -diagram block $\varphi(x_0) := \exists x_1 \dots \exists x_k \psi$, where $\psi := \beta(x_0, x_1) \wedge \text{diff}(x_0, \dots, x_k) \wedge \psi_0(x_0) \wedge \dots \wedge \psi_k(x_k)$. Note that we assume that the free variable x_0 of $\varphi(x_0)$ occurs in the σ -diagram $\beta(x_0, x_1)$ —unless σ does not contain binary relation symbols and thus $\beta(x_0, x_1) = \top$. The case where σ contains binary relation symbols and x_0 does not occur in the binary σ -diagram of ψ , is discussed later. Write each formula $\psi_i(x_i)$ in a form $\delta_i(x_i) \wedge \delta'_i(x_i)$, where $\delta_i(x_i)$ is a conjunction of the literals that occur as conjuncts in $\psi_i(x_i)$ and $\delta'_i(x_i)$ is the conjunction of the block formulas of $\psi_i(x_i)$. We have

$$\psi \equiv \beta(x_0, x_1) \wedge \text{diff}(x_0, \dots, x_k) \wedge (\delta_0(x_0) \wedge \delta'_0(x_0)) \wedge \dots \wedge (\delta_k(x_k) \wedge \delta'_k(x_k)).$$

Let P_0, \dots, P_k be fresh unary relation symbols. Consider the formula

$$\psi' := \beta(x_0, x_1) \wedge \text{diff}(x_0, \dots, x_k) \wedge (\delta_0(x_0) \wedge P_0(x_0)) \wedge \dots \wedge (\delta_k(x_k) \wedge P_k(x_k)).$$

Let σ' be the set of relation symbols in ψ' . Let us consider the formula $\chi(x_0) := \exists x_1 \dots \exists x_k \psi'$. By Lemma 4.4, we have $\chi(x_0) \equiv \chi'(x_0) := \theta_0(x_0) \vee \dots \vee \theta_m(x_0)$, where each $\theta_i(x_0)$ is a σ' -star centre type. We shall next show that each σ' -star centre type can be expressed in FOC^2 . This will conclude the argument concerning the formula $\varphi(x_0)$, as the disjuncts $\theta_i(x_0)$ of $\chi'(x_0)$ can first be replaced by equivalent FOC^2 -formulas, and after that, each subformula $P_i(z)$ ($0 \leq i \leq k$) in the resulting formula can be replaced by an FOC^2 -formula $\delta''_i(z) \equiv \delta'_i(z)$ obtained by the induction hypothesis. Here z is either of the variables in the two-variable formula we are constructing. If necessary, variables in δ'_i can be circulated to avoid variable capture. This way we obtain an FOC^2 -formula equivalent to $\varphi(x_0)$.

The notion of a star centre type was of course designed to be expressible in FOC^2 . Consider the σ' -star centre type $\exists x_1 \dots \exists x_k \gamma$, where γ is the σ' -star type

$$\gamma := \text{diff}(x_0, \dots, x_k) \bigwedge_{i \in \{1, \dots, k\}} \beta_i(x_0, x_k) \wedge \bigwedge_{i \in \{0, \dots, k\}} \alpha_i(x_i).$$

For each 2-type \mathcal{T} over σ' , let $\#\mathcal{T}$ denote the number of rays of γ that realize \mathcal{T} . Let T denote the set of all 2-types over σ' . Define

$$\gamma'(x_0) := \bigwedge_{\mathcal{T} \in T} \exists^{\geq \#\mathcal{T}} y \mathcal{T}(x_0, y).$$

It is easy to see that the FOC^2 -formula $\gamma'(x_0)$ is equivalent to the σ' -star centre type $\exists x_1 \dots \exists x_k \gamma$.

Let us then consider a σ -diagram block formula $\theta(x_0) := \exists x_1 \dots \exists x_k \eta$, where $\eta := \beta(x_i, x_j) \wedge \text{diff}(x_0, \dots, x_k) \wedge \psi_0(x_0) \wedge \dots \wedge \psi_k(x_k)$, and

$x_0 \notin \{x_i, x_j\}$. Let \bar{x} denote a tuple containing exactly the variables in $\{x_0, \dots, x_k\} \setminus \{x_i\}$. Consider the block formula $\theta'(x_i) := \exists \bar{x} \eta$. In this formula, the free variable x_i occurs in the part $\beta(x_i, x_j)$ of η . Thus, by our argument above, $\theta'(x_i)$ is equivalent to a formula $\theta''(x_i)$ of FOC^2 .

By the induction hypothesis, there are FOC^2 -formulas $\psi'_0(x_0) \equiv \psi_0(x_0)$ and $\psi'_j(x_j) \equiv \psi_j(x_j)$. Let $\psi'_j(x_0)$ denote the FOC^2 -formula obtained from $\psi'_j(x_j)$ by changing the free variable x_j to x_0 , and circulating variables, if necessary. Let the variables used in $\psi'_0(x_0)$ and $\psi'_j(x_0)$ be x_0 and x_i . Define $\mathcal{B}(x_i, x_0) := \beta(x_i, x_0) \wedge \psi'_j(x_0)$. The original formula $\theta(x_0)$ is equivalent to the FOC^2 -formula

$$\psi'_0(x_0) \wedge \exists x_i \left(x_i \neq x_0 \wedge \theta''(x_i) \wedge (\neg \mathcal{B}(x_i, x_0) \vee \exists^{\geq 2} x_0 (x_i \neq x_0 \wedge \mathcal{B}(x_i, x_0))) \right).$$

To conclude the proof, we need to discuss the case involving a block formula $\chi := \exists x_0 \dots \exists x_k \chi'$ that does not contain a free variable. Assume, w.l.o.g., that $k \geq 1$. Convert the block formula $\exists x_1 \dots \exists x_k \chi'$ to an FOC^2 -formula $\pi(x_0)$. Thus the original formula χ is equivalent to the FOC^2 -formula $\exists x_0 \pi(x_0)$. \square

5 Undecidability of UFC_1^-

Since FOC^2 and UF_1^- are decidable, it is natural to ask whether the extension of UF_1^- by counting quantifiers, UFC_1^- , remains decidable. Formally, UFC_1^- is obtained from UF_1^- by allowing the free substitution of quantifiers \exists by quantifiers $\exists^{\geq k}, \exists^{\leq k}, \exists^=k$. We next show that both the general and the finite satisfiability problems of UFC_1^- are undecidable.

For the proofs, we use the standard tiling and periodic tiling arguments. A *tile* is a mapping $t : \{R, L, T, B\} \rightarrow C$, where C is a countably infinite set of colours. We use the subscript notation $t_X := t(X)$ for $X \in \{R, L, T, B\}$. Intuitively, t_R, t_L, t_T and t_B are the colors of the right edge, left edge, top edge and bottom edge of the tile t , respectively.

Let $\mathfrak{S} := (S, H, V)$ be a structure with domain S and binary relations H and V . Let \mathbb{T} be a finite nonempty set of tiles. A \mathbb{T} -tiling of \mathfrak{S} is a function $f : S \rightarrow \mathbb{T}$ that satisfies the following conditions.

(T_H) For all $a, b \in S$, if $f(a) = t, f(b) = t'$ and $(a, b) \in H$, then $t_R = t'_L$.

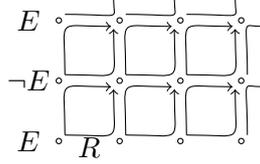
(T_V) For all $a, b \in S$, if $f(a) = t, f(b) = t'$ and $(a, b) \in V$, then $t_T = t'_B$.

The *tiling problem* for \mathfrak{S} asks, given a finite nonempty set \mathbb{T} of tiles, whether there exists a \mathbb{T} -tiling of \mathfrak{S} .

The standard *grid* is the structure $\mathfrak{G} := (\mathbb{N} \times \mathbb{N}, H, V)$, where $H = \{((i, j), (i+1, j)) \mid i, j \in \mathbb{N}\}$ and $V = \{((i, j), (i, j+1)) \mid i, j \in \mathbb{N}\}$ are binary relations. It is well known that the tiling problem for \mathfrak{G} is Π_1^0 -complete. Let n be a positive integer. Let $T := [0, n-1] \times [0, n-1]$. An $(n \times n)$ -torus is the structure (T, H, V) such that $H = \{((i, j), (i+1, j)) \mid (i, j) \in T\}$ and

$V = \{((i, j), (i, j + 1)) \mid (i, j) \in T\}$, where the sum is taken modulo n . The *periodic tiling problem* asks, given a finite nonempty set \mathbb{T} of tiles, whether there exist an $n \in \mathbb{Z}_+$ such the $(n \times n)$ -torus is \mathbb{T} -tilable. It is well known that the periodic tiling problem is Σ_1^0 -complete.

We shall below define a UFC_1^- -formula η which axiomatizes a sufficiently rich class of grid-like structures. In order to encode grids *with* UFC_1^- -formulas, we employ a *ternary* predicate R and a unary predicate E . Intuitively, E labels elements that represent the nodes of the even rows of a grid, and R contains triples (a, b, c) such that b is the horizontal successor of a and c is the vertical successor of b , or b is the vertical successor of a and c is the horizontal successor of b . The following figure depicts an initial portion of an infinite structure $\mathfrak{A}_{\mathfrak{G}}$ (over the signature $\{R, E\}$) which is our intended encoding of the standard infinite grid \mathfrak{G} . An arrow from a node a via b to c means that $R(a, b, c)$ holds.



Define the formulas $\varphi_H(x, y) := \exists z((R(x, y, z) \vee R(z, x, y)) \wedge (E(x) \leftrightarrow E(y)))$ and $\varphi_V(x, y) := \exists z((R(x, y, z) \vee R(z, x, y)) \wedge (E(x) \leftrightarrow \neg E(y)))$. Note that these are not UF_1^- -formulas. Let \mathfrak{A} be a structure over the vocabulary $\{R, E\}$. We let \mathfrak{A}^* be the structure over the vocabulary $\{H, V\}$ such that the \mathfrak{A}^* has the same domain A as \mathfrak{A} , and the relation $H^{\mathfrak{A}^*}$ ($V^{\mathfrak{A}^*}$) is the set of pairs $(a, a') \in A$ such that $\mathfrak{A} \models \varphi_H[a, a']$ ($\mathfrak{A} \models \varphi_V[a, a']$). Note that $\mathfrak{A}_{\mathfrak{G}}^*$ is the standard grid \mathfrak{G} .

We next define a UFC_1^- -formula η that captures some essential properties of $\mathfrak{A}_{\mathfrak{G}}$. Let η be the conjunction of the formulas (7) – (13) below. Note that the syntactic restrictions of UFC_1^- are indeed met.

$$\exists x E(x) \tag{7}$$

$$\forall x \exists^{=1} y \exists z (R(x, y, z) \wedge (E(x) \leftrightarrow E(y))) \tag{8}$$

$$\forall x \exists^{=1} y \exists z (R(x, y, z) \wedge (E(x) \leftrightarrow \neg E(y))) \tag{9}$$

$$\forall x \exists^{=1} z \exists y R(x, y, z) \tag{10}$$

$$\forall x \forall y \forall z (R(x, y, z) \rightarrow (E(x) \leftrightarrow \neg E(z))) \tag{11}$$

$$\forall x \exists^{=1} y \exists z (((E(x) \leftrightarrow E(y)) \wedge (R(z, x, y) \vee R(x, y, z)))) \tag{12}$$

$$\forall x \exists^{=1} y \exists z (((E(x) \leftrightarrow \neg E(y)) \wedge (R(z, x, y) \vee R(x, y, z)))) \tag{13}$$

We claim that η has the following properties.

- (i) There exists a model $\mathfrak{A} \models \eta$ such that $\mathfrak{A}^* = \mathfrak{G}$.

(ii) For every model $\mathfrak{A} \models \eta$, there is a homomorphism from \mathfrak{G} to \mathfrak{A}^* .

Assume we can show that η indeed has the above properties. Let \mathbb{T} be an arbitrary input to the tiling problem, and let $\mathcal{P}_{\mathbb{T}} := \{P_t \mid t \in \mathbb{T}\}$ be a set of fresh unary predicate symbols. Construct a UFC_1^- -formula $\varphi_{\mathbb{T}} := \psi_0 \wedge \psi_H \wedge \psi_V$ over the vocabulary $\{R, E\} \cup \mathcal{P}_{\mathbb{T}}$ as follows.

1. ψ_0 states that each point of the model is in the interpretation of exactly one predicate symbol P_t , $t \in \mathbb{T}$.
2. $\psi_H \equiv \forall x \forall y \bigwedge_{t, t' \in \mathbb{T}, t_R \neq t'_L} \neg(\varphi_H(x, y) \wedge P_t(x) \wedge P_{t'}(y))$. Note that the right hand side here is not a UFC_1^- -formula, but it can easily be modified so that the resulting formula is.
3. $\psi_V \equiv \forall x \forall y \bigwedge_{t, t' \in \mathbb{T}, t_T \neq t'_B} \neg(\varphi_V(x, y) \wedge P_t(x) \wedge P_{t'}(y))$.

It is easy to see that $\eta \wedge \varphi_{\mathbb{T}}$ has a model iff there exists a \mathbb{T} -tiling of \mathfrak{G} .

Claim 5.1. The formula η has the properties (i) and (ii).

Proof. For (i), we take as \mathfrak{A} the structure $\mathfrak{A}_{\mathfrak{G}}$. Let us consider the property (ii). Let \mathfrak{A} be a model such that $\mathfrak{A} \models \eta$. We will show how to construct a homomorphism h from the standard grid \mathfrak{G} to \mathfrak{A}^* . Let us first define the embedding of the first two rows of the grid. Let $a' \in A$ be an element witnessing the conjunct (7). Let $h(0, 0) = a'$. Let $b', c' \in A$ be elements such that $\mathfrak{A} \models R[a', b', c'] \wedge (E[a'] \leftrightarrow \neg E[b'])$, guaranteed to exist by (9). Let $h(0, 1) = b'$. We observe that $V[h(0, 0), h(1, 0)]$ holds in \mathfrak{A}^* , as required. We also have $\mathfrak{A} \models E[a'] \wedge \neg E[b']$.

Assume, for the sake of induction, that we have defined $h(k, 0) = a$ and $h(k, 1) = b$ for some $k \geq 0$. Assume that $\mathfrak{A}^* \models H[h(m', 0), h(m' + 1, 0)]$ and $\mathfrak{A}^* \models V[h(m, 0), h(m, 1)]$ for all $m' \in \{0, \dots, k - 1\}$ and $m \in \{0, \dots, k\}$. Assume also that $\mathfrak{A} \models E[h(m, 0)] \wedge \neg E[h(m, 1)]$ for all $m \in \{0, \dots, k\}$.

Let $c, d \in A$ be elements such that $\mathfrak{A} \models R[a, c, d] \wedge (E[a] \leftrightarrow E[c])$, guaranteed to exist by (8). Choose $h(k + 1, 0) = c$ and $h(k + 1, 1) = d$. We have $H[h(k, 0), h(k + 1, 0)]$ in \mathfrak{A}^* .

Let $c', d' \in A$ be elements such that $\mathfrak{A} \models R[a, c', d'] \wedge (E[a] \leftrightarrow \neg E[c'])$, guaranteed to exist by (9). We have $d' = d$ by (10). Thus $\mathfrak{A} \models R[a, c', d] \wedge (E[a] \leftrightarrow \neg E[c'])$. By (13), we conclude that $c' = b$. By (11), we have $\mathfrak{A} \models (E[a] \leftrightarrow \neg E[d])$, and thus $\mathfrak{A} \models (E[d] \leftrightarrow E[c'])$. Hence we have $\mathfrak{A} \models R[a, b, d] \wedge (E[b] \leftrightarrow E[d])$, whence $H[h(k, 1), h(k + 1, 1)]$ holds in \mathfrak{A}^* .

We still need to show that $\mathfrak{A}^* \models V[h(k + 1, 0), h(k + 1, 1)]$, i.e., $\mathfrak{A}^* \models V[c, d]$. We already know that \mathfrak{A} satisfies $E(a) \leftrightarrow E(c)$ and $E(a) \leftrightarrow \neg E(d)$. Thus $\mathfrak{A} \models E(c) \leftrightarrow \neg E(d)$. Since we also know that $\mathfrak{A} \models R[a, c, d]$, we conclude that $\mathfrak{A}^* \models V[c, d]$. Since $\mathfrak{A} \models E[a] \wedge (E[a] \leftrightarrow E[c]) \wedge (E[a] \leftrightarrow \neg E[d])$, we have $\mathfrak{A} \models E[h(k + 1, 0)] \wedge \neg E[h(k + 1, 1)]$.

We have defined h for the first two rows of \mathfrak{G} . Assume, for the sake of induction, that we have defined h for the first $l \geq 2$ rows of the grid, and

that the homomorphism conditions are satisfied. Assume also that for all $m \in \mathbb{N}$ and $m' \in \{0, \dots, l\}$, we have $\mathfrak{A} \models E[h(m, m')]$ iff m' is even.

We extend the definition of h to the $(l+1)$ -st row. As above, we proceed by induction on the columns. Let $h(0, l) = a$. Let $b, c \in A$ be elements such that $\mathfrak{A} \models R[a, b, c] \wedge (E[a] \leftrightarrow \neg E[b])$, guaranteed to exist by (9). Let $h(0, l+1) = b$. We may assume, by symmetry, that $\mathfrak{A} \models E[h(0, l)]$. We have $\mathfrak{A}^* \models V[h(0, l), h(0, l+1)]$ and $\mathfrak{A} \models \neg E[h(0, l+1)]$.

Assume, for the sake of induction, that there exists a k such that we have defined $h(m, l+1)$ for all $m \leq k$. Assume that $\mathfrak{A}^* \models H[h(m', l+1), h(m'+1, l+1)]$ for all $m' \in \{0, \dots, k-1\}$ and $\mathfrak{A}^* \models V[h(m, l), h(m, l+1)]$ for all $m \in \{0, \dots, k\}$. Assume also that $\mathfrak{A} \models \neg E[h(m, l+1)]$ for all $m \in \{0, \dots, k\}$.

Let $b, c \in A$ be elements such that $\mathfrak{A} \models R[a, b, c] \wedge (E[a] \leftrightarrow E[b])$, guaranteed to exist by (8). Define $h(k+1, l+1) = c$. The arguments concerning the homomorphism conditions (and the predicate E) are similar to the arguments concerning the first two rows, but involve also the condition (12). The straightforward details are left to the reader. \square

Thus the satisfiability problem of UFC_1^- is Π_1^0 -hard. Since UFC_1^- is a fragment of first-order logic, the following theorem holds.

Theorem 5.2. The satisfiability problem of UFC_1^- is Π_1^0 -complete.

The above argument leading to Theorem 5.2 can be used with minor modifications in order to show Σ_1^0 -completeness of the finite satisfiability problem of UFC_1^- using the periodic tiling problem.

Theorem 5.3. The finite satisfiability problem of UFC_1^- is Σ_1^0 -complete.

Proof. We use the same formulas η and $\varphi_{\mathbb{T}}$ as in the proof of Theorem 5.2. We claim that η has the following properties.

- (iii) For every positive integer n , there is a structure $\mathfrak{A} \models \eta$ such that \mathfrak{A}^* is the $(2n \times 2n)$ -torus. (We use the factor 2 here in order to deal with the predicate E appropriately.)
- (iv) For every finite model $\mathfrak{A} \models \eta$, there exists some $(n \times n)$ -torus \mathfrak{T} such that there is a homomorphism from \mathfrak{T} to \mathfrak{A}^* .

Assume we can show that η satisfies (iii) and (iv). It is then easy to show that for a nonempty finite set \mathbb{T} of tiles, the formula $\eta \wedge \varphi_{\mathbb{T}}$ has a finite model iff there exists some $n \in \mathbb{Z}_+$ such that the $(n \times n)$ -torus is \mathbb{T} -tilable. Thus the finite satisfiability problem for UFC_1^- is Σ_1^0 -hard. Since the finite satisfiability problem of first-order logic is in Σ_1^0 , the theorem holds.

We then sketch a proof that η indeed has the properties (iii) and (iv). For each positive integer n , we let \mathfrak{T}_n denote the $(n \times n)$ -torus. For (iii), we define \mathfrak{A} to be the natural quotient of $\mathfrak{A}_{\mathfrak{G}}$ with $2n \times 2n$ elements such that $\mathfrak{A}^* = \mathfrak{T}_{2n}$. To prove (iv), let \mathfrak{A} be a finite structure such that $\mathfrak{A} \models \eta$. We

first construct a homomorphism h from the standard grid \mathfrak{G} to \mathfrak{A}^* precisely as in the proof of Theorem 5.2. Since \mathfrak{A} is finite, there exist some $i, j, i < j$, such that we have $h(i, 0) = h(j, 0)$. We can show by induction, using formula (13), that $h(i, s) = h(j, s)$ for all $s \in \mathbb{N}$. Now, for each $s \in \mathbb{N}$, let π_s denote the tuple

$$(h(i, s), h(i + 1, s), \dots, h(j - 1, s)).$$

Finiteness of \mathfrak{A} guarantees that there exist $k, l, k < l$ such that $\pi_k = \pi_l$. Let $h' : [0, j - i - 1] \times [0, l - k - 1] \rightarrow A$ be the function defined so that $h'(x, y) = h(i + x, k + y)$. The arguments given above imply that h' is a homomorphism from the (not necessarily square) $(j - i) \times (l - k)$ torus to \mathfrak{A}^* . Define $t := \text{lcm}(j - i, l - k, 2)$, where lcm denotes the least common multiple operation. Let $m := \frac{t}{2}$. Using h' , it is easy to define an embedding from the square $(2m \times 2m)$ -torus to \mathfrak{A}^* . \square

References

- [1] H. Andr eka, J. van Benthem, and I. N emeti. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27(3):217–274, 1998.
- [2] V. B arany, B. ten Cate, and L. Segoufin. Guarded negation. In *ICALP (2)*, pages 356–367, 2011.
- [3] S. Benaim, M. Benedikt, W. Charatonik, E. Kieroński, R. Lenhardt, F. Mazowiecki, and J. Worrell. Complexity of two-variable logic on finite trees. In *ICALP (2)*, pages 74–88, 2013.
- [4] W. Charatonik and P. Witkowski. Two-variable logic with counting and trees. In *LICS*, pages 73–82, 2013.
- [5] H.-D. Ebbinghaus and J. Flum. *Finite model theory*. Perspectives in Mathematical Logic. Springer, 1995.
- [6] E. Gr adel, P. Kolaitis, and M. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.
- [7] E. Gr adel, M. Otto, and E. Rosen. Two-variable logic with counting is decidable. In *LICS*, pages 306–317, 1997.
- [8] L. Hella and A. Kuusisto. One-dimensional fragment of first-order logic. *arXiv:1404.4004*, 2014.
- [9] L. Henkin. Logical systems containing only a finite number of symbols. *Presses De l’Universit e De Montr eal*, 1967.

- [10] E. Kieroński, J. Michaliszyn, I. Pratt-Hartmann, and L. Tendera. Two-variable first-order logic with equivalence closure. In *LICS*, pages 431–440, 2012.
- [11] H. R. Lewis. Complexity results for classes of quantificational formulas. *Journal of Computer and System Sciences*, 21(3):317 – 353, 1980.
- [12] M. Mortimer. On languages with two variables. *Mathematical Logic Quarterly*, 21(1):135–140, 1975.
- [13] L. Pacholski, W. Szwast, and L. Tendera. Complexity of two-variable logic with counting. In *LICS*, pages 318–327. IEEE, 1997.
- [14] I. Pratt-Hartmann. Complexity of the two-variable fragment with counting quantifiers. *Journal of Logic, Language and Information*, 14(3):369–395, 2005.
- [15] D. Scott. A decision method for validity of sentences in two variables. *Journal Symbolic Logic*, 27:477, 1962.
- [16] W. Szwast and L. Tendera. FO^2 with one transitive relation is decidable. In *STACS*, pages 317–328, 2013.