

One-Dimensional Logic over Words*

Emanuel Kieroński

University of Wrocław, Poland
kiero@cs.uni.wroc.pl

Abstract

One-dimensional fragment of first-order logic is obtained by restricting quantification to blocks of existential quantifiers that leave at most one variable free. We investigate one-dimensional fragment over words and over ω -words. We show that it is expressively equivalent to the two-variable fragment of first-order logic. We also show that its satisfiability problem is NEXPTIME-complete. Further, we show undecidability of some extensions, whose two-variable counterparts remain decidable.

1998 ACM Subject Classification F.4 Mathematical Logic and Formal Languages

Keywords and phrases satisfiability, expressivity, words, fragments of first-order logic

Digital Object Identifier 10.4230/LIPIcs.CSL.2016.38

1 Introduction

One-dimensional fragment of first-order logic, F_1 , is obtained by restricting quantification to blocks of existential quantifiers that leave at most one variable free. It is not difficult to show that over general relational structures the satisfiability problem for F_1 is undecidable [8]. Its *uniform* variant, UF_1 , was introduced by Hella and Kuusisto in [8] as a generalization of the two-variable fragment of first-order logic, FO^2 , to contexts with relations of arity higher than two. In that paper the decidability and the finite model property for UF_1 without equality was proved. Roughly speaking, the uniformity restrictions allow for Boolean combination of atoms Rx_1, \dots, x_m and Sy_1, \dots, y_n of arity greater than one only if the sets of variables $\{x_1, \dots, x_m\}$ and $\{y_1, \dots, y_n\}$ are equal; Boolean combinations of atoms of arity one can be formed freely. In [9] the finite model property was extended to UF_1 with free (i.e. not necessarily uniform) use of equality. It was also shown that the satisfiability problem is NEXPTIME-complete. Both results were obtained, to some extent, by a generalisation of the classical techniques used by Grädel, Kolaitis and Vardi in [7] in context of FO^2 (whose satisfiability problem is also NEXPTIME-complete). A nice survey of the results on UF_1 can be found in a recent paper by Kuusisto [13], which also reveals some connections between UF_1 and description logics.

The uniformity restriction is indeed crucial for the decidability of UF_1 . Unfortunately, it also limits the possible scenarios in which this logic can be used. A question is if there are any ways of weakening it without losing decidability. Please note that actually the variant of UF_1 considered in [9] has a non-uniform ingredient, namely the equality predicate. This non-uniformity indeed gives an additional power: it turns out that UF_1 restricted to signatures with only unary and binary relational symbols is expressively equivalent to FO^2 if the uniform use of equality is imposed, while it is more expressive if there is no such restriction, offering in particular some sort of counting. Even a stronger decidable

* Partially supported by the Polish National Science Centre grant DEC-2013/09/B/ST6/01535.



© Emanuel Kieroński;

licensed under Creative Commons License CC-BY

25th EACSL Annual Conference on Computer Science Logic (CSL 2016).

Editors: Jean-Marc Talbot and Laurent Regnier; Article No. 38; pp. 38:1–38:15



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

logic was identified in [10]. It is UF_1 with a free use of one equivalence relation, shown to be 2-NEXPTIME-complete (or even NEXPTIME-complete when some natural variation is considered). That paper demonstrates however that this decidability result is fragile. E.g., adding a second equivalence relation leads to undecidability. This contrasts with the case of FO^2 which is decidable in the presence of two equivalences and becomes undecidable only after adding the third one [11]. UF_1 becomes also undecidable when extended with a non-uniform use of a single transitive relation. Again, this reveals a difference in comparison with FO^2 whose satisfiability in models with one transitive relation was shown decidable by Szwast and Tendera [19].

Instead of extending UF_1 with some non-uniform ingredients we may also try to investigate full F_1 over some classes of structures, in which the meaning of non-unary symbols is fixed. One of the simplest, but very important such class is the class of words. Many formalisms over words have been investigated so far. It is known that the satisfiability problem for full first-order logic is decidable, but with non-elementary complexity, as shown by Stockmeyer [18]. In fact, already the fragment with three variables is non-elementary. On the other hand a reasonable complexity appears when the number of variables is restricted to two. The satisfiability problem for FO^2 over words and ω -words was shown to be NEXPTIME-complete by Etessami, Vardi and Wilke [6]. In the same paper it is observed that the expressive power of FO^2 over words is equal to the expressive power of unary temporal logic, UTL, i.e., temporal logic with four navigational operators: *next state*, *somewhere in the future*, *previous state*, *somewhere in the past*. FO^2 , however, turns out to be exponentially more succinct than UTL. An extension of FO^2 with counting quantifiers, C^2 , is shown to be NEXPTIME-complete over words by Charatonik and Witkowski, [4]. In fact, it is not difficult to observe that over words C^2 has also the same expressive power as plain FO^2 . Another interesting extension of FO^2 , which significantly increases its expressive power is an extension with the *between* predicate recently studied by Krebs et al. [12].

In this paper we study the expressive power and the complexity of the satisfiability problem of the full one-dimensional fragment of first-order logic over words, $F_1[<, +1]$. First, we show that its expressive power is the same as the expressive power of FO^2 , and thus also of UTL, and, as mentioned, of C^2 .

The advantage of F_1 over those other formalisms is that it allows to specify some properties in a more natural and elegant way. If we want to say that a word contains some (especially not fully specified) pattern, consisting of more than two elements, we can just quantify an appropriate number of positions, and say how they should be labelled and related to each other. Expressing the same in FO^2 will usually require some heavy recycling of the two available variables. Let us look at two simple examples. Consider a system whose behaviour we model as a word or an ω -word, in which one or more of n atomic propositions out of P_1, \dots, P_n can hold in a given point of time. To say that there are m non-overlapping time intervals (sets of consecutive positions of the word) in each of which each of P_i holds at least once, we can use the following $F_1[<, +1]$ sentence:

$$\exists y_0 y_1 \dots y_n x_{11} \dots x_{1n} \dots x_{m1} \dots x_{mn} \left(\bigwedge_{i=1}^m \bigwedge_{j=1}^n y_{i-1} \leq x_{ij} \wedge x_{ij} < y_i \wedge P_j x \right).$$

As another example¹ take the property saying that it is possible to choose no more than m positions satisfying together all of P_i :

¹ Pointed out to the author by Jakub Michaliszyn.

$$\exists x_1 \dots x_m \left(\bigwedge_{i=1}^n \bigvee_{j=1}^m P_i x_j \right).$$

The reader is asked to check that expressing the above properties in $\text{FO}^2[<, +1]$ is indeed not straightforward and leads to complicated formulas.

In fact, what is worth mentioning here, our translation of $F_1[<, +1]$ to $\text{FO}^2[<, +1]$ has an exponential blow-up, which seems to be hard to avoid, and which thus suggests that $F_1[<, +1]$ may be able to express some properties more succinctly than $\text{FO}^2[<, +1]$, and possibly, even C^2 .

Further, we turn our attention to the satisfiability problem for F_1 , managing to show that, in spite of the exponential blow-up in the translation, $F_1[<, +1]$ retains the complexity of $\text{FO}^2[<, +1]$, i.e., is NEXPTIME -complete. While our proof has some similarities to the proof of Etesami, Vardi and Wilke [6] for $\text{FO}^2[<, +1]$, it is technically more involved, due to the combinatorically more complicated nature of the objects involved. Nevertheless, the basic idea in the proof is rather straightforward and is based on an appropriately tuned contraction procedure.

We conclude the paper examining some possible extensions of $F_1[<, +1]$. Probably, the most significant of them is the extension of $F_1[<, +1]$ with an equivalence relation, inspired by an analogous extension of $\text{FO}^2[<, +1]$ (FO^2 over *data words*), studied by Bojańczyk et al. [2]. The satisfiability problem for FO^2 over data words, even though very hard, is decidable. We show that $F_1[<, +1]$ over data words becomes undecidable.

Finally, we suggest some related open problems concerning F_1 over some specific classes of structures.

2 Preliminaries

We assume that the reader is familiar with basic concepts in mathematical logic and computational complexity theory. Throughout this paper we mostly use standard terminology and notation.

By *one-dimensional fragment* of first-order logic, F_1 , we mean the relational fragment in which quantification is restricted to blocks of existential quantifiers that leave at most one variable free. Formally, F_1 over relational signature τ and some countably infinite set of variables Var is the smallest set such that:

- $R\bar{x} \in F_1$ for all $R \in \tau$ and all tuples \bar{x} of variables from Var ,
- $x = y \in F_1$ for all variables $x, y \in Var$,
- F_1 is closed under \vee and \neg ,
- if φ is an F_1 formula with free variables x_0, \dots, x_k then formulas $\exists x_0, \dots, x_k \varphi$ and $\exists x_1, \dots, x_k \varphi$ belong to F_1 .

As usually, we can use standard abbreviations for other Boolean operations, like \wedge , \rightarrow , \top , etc., as well as for universal quantification. The length of a formula φ is measured in a natural way, and denoted $\|\varphi\|$. The *width* of a formula is the maximum of the numbers of free variables in its subformulas.

We will be primarily interested in signatures consisting of a possibly infinite set of unary symbols and two binary symbols $+1$ and $<$. The obtained logic is then denoted $F_1[<, +1]$. Sometimes we will use another binary symbol \ll , which is an abbreviation: $x \ll y \equiv x < y \wedge \neg(+1(x, y))$. In Section 5 we will consider also some other binary symbols, whose meaning will be then explained.

We denote structures with Gothic capital letters, possibly with decorations: $\mathfrak{M}, \mathfrak{M}', \mathfrak{M}_1$, etc., and their universes with the corresponding Roman capital letters M, M', M_1 , etc. We are interested in structures in which $<$ is interpreted as a linear order and $+1$ as its induced successor relation. Such a structure is called a *finite word*, or just a *word*, if its universe is finite, and ω -*word* if after dropping the interpretation of unary relations it is isomorphic to $(\mathbb{N}, <, +1)$. If \mathfrak{M} is a word and \mathfrak{M}' a word or ω -word we denote by $\mathfrak{M}\mathfrak{M}'$ the word obtained by the concatenation of \mathfrak{M} and \mathfrak{M}' ; if \mathfrak{M}' consists of just one element a then this concatenation is written as $\mathfrak{M}a$. An ω -word built out of a finite word \mathfrak{M}_0 followed by infinitely many copies of a finite word \mathfrak{M}_1 is denoted by $\mathfrak{M}_0\mathfrak{M}_1^\omega$. Such an ω -word is called *periodic*. When referring to the elements of a model \mathfrak{M} we will sometimes denote by $a + i$, for $a \in M$ and $i \in \mathbb{Z}$, the element located i positions to the right from a if $i > 0$, $-i$ positions to the left from a if $i < 0$, and the element a if $i = 0$. The *satisfiability problem* over words (ω -words) for a logic \mathcal{L} is to check if for a given sentence $\varphi \in \mathcal{L}$ there exists a word (ω -word) \mathfrak{M} such that $\mathfrak{M} \models \varphi$.

3 Expressivity

It is known that $\text{FO}^2[<, +1]$ is expressively equivalent over words and ω -words to unary temporal logic, UTL, i.e., temporal logic with four navigational operators: *next state*, *somewhere in the future*, *previous state*, *somewhere in the past* [6]. Here we show that $\text{F}_1[<, +1]$ shares their expressivity:

► **Theorem 1.** $\text{F}_1[<, +1]$ and $\text{FO}^2[<, +1]$ are expressively equivalent over words and ω -words.

Obviously, $\text{FO}^2[<, +1]$ can be seen as a fragment of $\text{F}_1[<, +1]$. Here we present a translation from $\text{F}_1[<, +1]$ to $\text{FO}^2[<, +1]$ which justifies Thm. 1. More specifically, we show that for any $\text{F}_1[<, +1]$ sentence there is an FO^2 sentence satisfied in precisely the same models, and that for any $\text{F}_1[<, +1]$ formula with one free variable there is an FO^2 formula with one free variable such that they are satisfied at the same positions of any model. The crux is to show how to translate formulas starting with a block of quantifiers.

► **Lemma 2.** For any $\text{F}_1[<, +1]$ formula $\psi = \exists y_1 \dots, y_k \psi_0(y_0, y_1, \dots, y_k)$ with free variable y_0 there exists an FO^2 formula ψ' with one free variable such that for every word or ω -word \mathfrak{M} and every $a \in M$ we have $\mathfrak{M} \models \psi[a]$ iff $\mathfrak{M} \models \psi'[a]$. Similarly, for any $\text{F}_1[<, +1]$ sentence $\psi = \exists y_1 \dots, y_k \psi_0(y_1, \dots, y_k)$ there exists an FO^2 sentence ψ' such that for any word or ω -word \mathfrak{M} we have $\mathfrak{M} \models \psi$ iff $\mathfrak{M} \models \psi'$.

Proof. We prove this lemma by induction over the quantifier depth of ψ , measured as the maximal nesting depth of blocks of quantifiers rather than of individual quantifiers. We explicitly consider the case of a subformula with a free variable (the case of sentences can be treated similarly). Let us take any

$$\psi = \exists y_1 \dots, y_k \psi_0(y_0, y_1, \dots, y_k), \quad (1)$$

convert ψ_0 into disjunctive form and distribute existential quantifiers over disjunctions, obtaining

$$\psi = \bigvee_{i=1}^l \exists y_1 \dots, y_k \psi_i(y_0, y_1, \dots, y_k), \quad (2)$$

for some $l \in \mathbb{N}$, where each ψ_i is a conjunction of literals, subformulas with one free variable of the form $\exists z_1, \dots, z_k \psi_0(y_j, z_1, \dots, z_k)$, subsentences of the form $\exists z_1, \dots, z_k \psi_0(z_1, \dots, z_k)$, and negations of such formulas.

Recall that possible atoms are Ay_i for a unary symbol A , $y_i < y_j$, $+1(y_i, y_j)$ and $y_i = y_j$, for some i, j .

An *ordering scheme* over variables y_0, \dots, y_k is a formula of the form $\eta_0(y_{i_0}, y_{i_1}) \wedge \eta_1(y_{i_1}, y_{i_2}) \wedge \dots \wedge \eta_{k-1}(y_{i_{k-1}}, y_{i_k})$, where $\eta_i(v, w)$ is one of the following formulas: $v = w$, $+1(v, w)$ or $v \ll w$, and $i_0, i_1, i_2, \dots, i_k$ is a permutation of $0, 1, \dots, k$.

Consider now a single disjunct $\exists y_1 \dots, y_k \psi_i(y_0, y_1, \dots, y_k)$ of (2) (again assuming that it is not a sentence and has free variable y_0), and replace it by the following disjunction over all possible ordering schemes π over y_0, \dots, y_k :

$$\bigvee_{\pi} \exists y_1 \dots, y_k (\pi(y_0, \dots, y_k) \wedge \psi_i^{\pi}(y_0, y_1, \dots, y_k)), \quad (3)$$

where ψ_i^{π} is obtained from ψ_i by replacing all atoms $y_i < y_j$, $+1(y_i, y_j)$ and $y_i = y_j$, which are not bounded by the quantifiers of the subformulas of ψ_i by \top or \perp , according to the information recorded in π . Let us rearrange ψ_i^{π} into $\psi_{i,*}^{\pi} \wedge \bigwedge_{j=0}^k \psi_{i,j}^{\pi}(y_j)$, where $\psi_{i,*}^{\pi}$ consists of the conjuncts without free variables and $\psi_{i,j}^{\pi}(y_j)$ consists of the conjuncts with free variable y_j . We now explain how to translate a single disjunct

$$\exists y_1 \dots, y_k (\pi(y_0, \dots, y_k) \wedge \psi_{i,*}^{\pi} \wedge \bigwedge_{j=0}^k \psi_{i,j}^{\pi}(y_j)) \quad (4)$$

of (3). Let i_0, i_1, \dots, i_k be the permutation used to generate π , and let s be such that $i_s = 0$. By the inductive assumption we can replace $\psi_{i,*}^{\pi}$ by an equivalent FO^2 sentence $\psi'_{i,*}$, with two variables. We can also replace in each $\psi_{i,j}^{\pi}(y_j)$ any conjunct of the form $\exists z_1, \dots, z_k \chi(y_j, z_1, \dots, z_k)$ by an equivalent two-variable conjunct with one free variable. Thus, in turn, $\psi_{i,j}^{\pi}(y_j)$ can be replaced by an equivalent FO^2 formula $\psi'_{i,j}$ with one free variable.

We finally replace (4) by the conjunction of

$$\psi'_{i,*} \wedge \psi'_{i,i_s}(y_0), \quad (5)$$

$$\exists y (\eta_{s-1}(y, y_0) \wedge \psi'_{i,i_{s-1}}(y) \wedge \exists y_0 (\eta_{s-2}(y_0, y) \wedge \psi'_{i,i_{s-2}}(y_0) \wedge \dots)), \quad (6)$$

$$\exists y (\eta_{s+1}(y, y_0) \wedge \psi'_{i,i_{s+1}}(y) \wedge \exists y_0 (\eta_{s+2}(y_0, y) \wedge \psi'_{i,i_{s+2}}(y_0) \wedge \dots)), \quad (7)$$

in which (5) enforces the satisfaction of the proper subsentences, (6) takes care of witnesses smaller than (or equal) to y_0 , passing the word from y_0 to the left, and (7) takes care of witnesses greater than (or equal to) y_0 , passing the word from y_0 to the right. Of course, in all the above formulas we appropriately rename the variables if necessary, so that only y_0 and y are used. \blacktriangleleft

Having translated formulas starting with blocks of quantifiers, we can easily translate other formulas with at most one free variable, since they are just boolean combinations of the former. This gives a translation from $\text{F}_1[<, +1]$ to $\text{FO}^2[<, +1]$.

Observe that starting from an $\text{F}_1[<, +1]$ formula this translation may produce a formula in $\text{FO}^2[<, +1]$ which is exponentially longer. Essentially, there are two sources of this exponential blow-up. The first is the transformation to disjunctive form, and the second is considering all possible permutations of variables quantified in a single block of quantifiers. The question whether this blow-up is necessary is left open.

Let us mention here that $\text{C}^2[<, +1]$, the two-variable logic with counting quantifiers easily translates to $\text{F}_1[<, +1]$. For example, to express $\exists^{\geq k} y \psi(x, y)$ we can just write $\exists y_1, \dots, y_k (\bigwedge_{1 \leq i < j \leq k} y_i \neq y_j \wedge \bigwedge_{i=1}^k \psi(x, y_i))$. Thus all the logics from the following list: UTL , $\text{FO}^2[<, +1]$, $\text{C}^2[<, +1]$, $\text{F}_1[<, +1]$ are expressively equivalent over words and ω -words.

4 Satisfiability

We next turn our attention to satisfiability. We prove that the satisfiability problem for $F_1[<, +1]$ both over words and ω -words is NEXPTIME-complete. To this end we start with introducing a convenient normal form, inspired by Scott normal form for FO^2 (a similar normal form is used also in [9]). Then we develop a contraction method involving a careful analysis of certain similarities between elements in a model, and explain how to use it in order to obtain small model properties for $F_1[<, +1]$ both over words and ω -words. Then the complexity result will easily follow.

4.1 Normal form

We adapt here the well known Scott normal form for FO^2 [17] to our purposes. We say that an $F_1[<, +1]$ formula φ is in *normal form* if φ has the following shape:

$$\bigwedge_{1 \leq i \leq m_{\exists}} \forall y_0 \exists y_1 \dots y_{k_i} \varphi_i^{\exists} \wedge \bigwedge_{1 \leq i \leq m_{\forall}} \forall x_1 \dots x_{l_i} \varphi_i^{\forall}, \quad (8)$$

where $\varphi_i^{\exists} = \varphi_i^{\exists}(y_0, y_1, \dots, y_{k_i})$ and $\varphi_i^{\forall} = \varphi_i^{\forall}(x_1, \dots, x_{l_i})$ are quantifier-free. Please note that the width of φ is the maximum of the set $\{k_i + 1\}_{1 \leq i \leq m_{\exists}} \cup \{l_j\}_{1 \leq j \leq m_{\forall}}$. The following fact can be proved in a standard fashion, see, e.g., [5] for a more detailed exposition of the technique.

► **Lemma 3.** *For every $F_1[<, +1]$ formula φ , one can compute in polynomial time an $F_1[<, +1]$ formula φ' in normal form (over the signature extended by some fresh unary symbols) such that: (i) any model of φ can be expanded to a model of φ' by appropriately interpreting new unary symbols; (ii) any model of φ' restricted to the signature of φ is a model of φ .*

Proof (Sketch). We successively replace innermost subformulas ψ of φ of the form $\exists y_1, \dots, y_k \varphi(y_0, y_1, \dots, y_k)$ by atoms $P_{\psi}(y_0)$, where P_{ψ} is a fresh unary symbol, and axiomatize P_{ψ} using two normal form conjuncts: $\forall y_0 \exists y_1, \dots, y_k (P_{\psi}(y_0) \rightarrow \varphi(y_0, y_1, \dots, y_k))$ and $\forall y_0, y_1, \dots, y_k (\neg \varphi(y_0, y_1, \dots, y_k) \vee P_{\psi}(y_0))$. ◀

The above lemma allows us, when dealing with satisfiability or when analysing the size and shape of models, to restrict attention to normal form formulas.

4.2 Contraction

Let τ be a finite unary signature. A *1-type* over τ is a subset of τ . We say that an element $a \in M$ realizes a 1-type α in a word or ω -word \mathfrak{M} , and write $\mathfrak{M} \models \alpha[a]$, if for each $A \in \tau$ we have $\mathfrak{M} \models A[a]$ iff $A \in \alpha$. The 1-type realized by a in \mathfrak{M} is denoted by $\text{type}^{\mathfrak{M}}(a)$.

Let us introduce some new, more sophisticated concepts which will turn out to be helpful in our constructions.

► **Definition 4.**

- An *ordered k -type* is a tuple of the form $(\alpha_1, \eta_1, \alpha_2, \dots, \alpha_{k-1}, \eta_{k-1}, \alpha_k)$, where α_i is an atomic 1-type ($1 \leq i \leq k$) and $\eta_i(x, y)$ is either $+1(x, y)$ or $x \ll y$ ($1 \leq i \leq k-1$).
- For a given word or ω -word \mathfrak{M} we say that a tuple of its distinct elements $a_1, \dots, a_k \in M$ realizes an ordered k -type $(\alpha_1, \eta_1, \alpha_2, \dots, \alpha_{k-1}, \eta_{k-1}, \alpha_k)$ if $\mathfrak{M} \models \alpha_i[a_i]$ for $1 \leq i \leq k$ and $\mathfrak{M} \models \eta_i[a_i, a_{i+1}]$ for $1 \leq i < k$.

Thus, an ordered k -type stores some basic information about k distinct elements in a model and about their relative location. When k is clear from the context we sometimes talk just about the *ordered type* of a tuple. The ordered type realized by a tuple a_1, \dots, a_k in \mathfrak{M} is denoted as $\text{ordtype}^{\mathfrak{M}}(a_1, \dots, a_k)$. We are going to abstract the information about a single element of a model using the following notion.

► **Definition 5.** For a given word or ω -word \mathfrak{M} , $a \in M$, and a natural number $n > 0$, we say that the n -profile of a is the tuple $(\alpha_{-n}, \dots, \alpha_{-1}, \alpha_0, \alpha_1, \dots, \alpha_n, L_{-1}, \dots, L_{-n}, R_1, \dots, R_n)$, where

- α_i is the 1-type of element $a + i$ of \mathfrak{M} if $a + i$ is defined, or $\alpha_i = \emptyset$ otherwise,
- L_{-i} is the set of the ordered k -types ($1 \leq k \leq n$) realized in the prefix of \mathfrak{M} ending at $a - i - 1$,
- R_i is the set of the ordered k -types ($1 \leq k \leq n$) realized in the suffix of \mathfrak{M} starting at $a + i + 1$.

The n -profile of an element $a \in M$ is denoted as $\text{prof}_n^{\mathfrak{M}}(a)$. We sometimes say that an element *realizes* its profile.

Let us observe that the number of possible n -profiles realized in a model is not very large.

► **Lemma 6.** *If \mathfrak{M} is a word or ω -word over a signature τ then the number of different n -profiles of elements of \mathfrak{M} is bounded exponentially in $|\tau|$ and in n .*

Proof. The number of atomic 1-types realized in \mathfrak{M} is at most $2^{|\tau|}$. The number of ordered k -types is bounded by $2^{k|\tau|} \cdot 2^{(k-1)} \leq 2^{k(|\tau|+1)}$. Thus the size of each of the sets L_i and R_i in an n -profile is bounded exponentially by $n \cdot 2^{n(|\tau|+1)}$. Moreover, these sets behave monotonically in \mathfrak{M} : If $(\alpha_{-n}, \dots, \alpha_{-1}, \alpha_0, \alpha_1, \dots, \alpha_n, L_{-1}, \dots, L_{-n}, R_1, \dots, R_n)$ and $(\alpha'_{-n}, \dots, \alpha'_{-1}, \alpha'_0, \alpha'_1, \dots, \alpha'_n, L'_{-1}, \dots, L'_{-n}, R'_1, \dots, R'_n)$ are n -profiles of elements $a, a' \in M$, respectively, such that $\mathfrak{M} \models a < a'$, then $L_i \subseteq L'_i$ and $R_i \supseteq R'_i$ ($-n \leq i \leq n$). Hence, when moving from the first position of \mathfrak{M} to the right, each of the sets L_i and R_i may change at most exponentially many times. Since the number of n -profiles with all L_i and R_i fixed is bounded by $2^{(2n+1)|\tau|}$ the claim follows. ◀

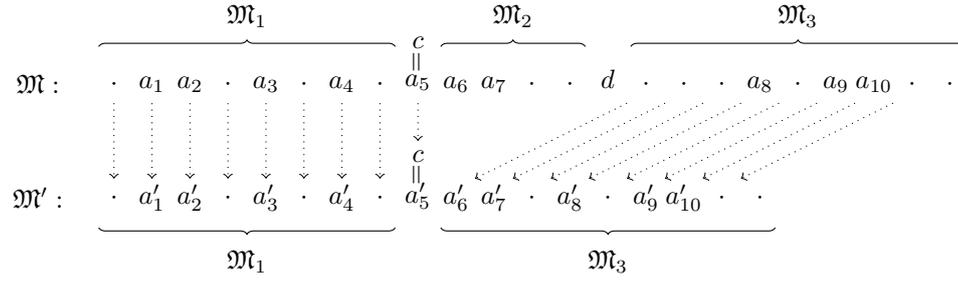
We are ready to prove the crucial contraction lemma. Namely, we observe that removing a fragment of a word between two realizations of the same profile, does not change the profiles of the surviving elements. As we will see later such surgery also does not affect the satisfaction of certain normal form formulas.

► **Lemma 7.** *Let $\mathfrak{M} = \mathfrak{M}_1 c \mathfrak{M}_2 d \mathfrak{M}_3$ be a word or ω -word and $n > 0$ a natural number. Assume that $\text{prof}_n^{\mathfrak{M}}(c) = \text{prof}_n^{\mathfrak{M}}(d)$ and $\mathfrak{M}' = \mathfrak{M}_1 c \mathfrak{M}_3$. Then for each $a' \in M'$ we have $\text{prof}_n^{\mathfrak{M}'}(a') = \text{prof}_n^{\mathfrak{M}}(a')$.*

Proof. Consider the case when $a' \in M_1 \cup \{c\}$. Note that the prefix of \mathfrak{M} ending in a' is equal to the prefix of \mathfrak{M}' ending in a' . It follows that for all $-n \leq i < 0$ we have $\text{prof}_n^{\mathfrak{M}'}(a').\alpha_i = \text{prof}_n^{\mathfrak{M}}(a').\alpha_i$ and $\text{prof}_n^{\mathfrak{M}'}(a').L_i = \text{prof}_n^{\mathfrak{M}}(a').L_i$, since these components refer only to positions located in this initial prefix.

Take such $l \geq 0$ that in \mathfrak{M}' we have $c = a' + l$. For $0 \leq i \leq l$, it is the case that $a' + i$ in \mathfrak{M} equals $a' + i$ in \mathfrak{M}' and thus $\text{prof}_n^{\mathfrak{M}'}(a').\alpha_i = \text{prof}_n^{\mathfrak{M}}(a').\alpha_i$; if $l < i \leq n$ then $a' + i = c + (i - l)$ in \mathfrak{M}' and it equals $d + (i - l)$ in \mathfrak{M} , and since $\text{prof}_n^{\mathfrak{M}'}(c).\alpha_{i-l} = \text{prof}_n^{\mathfrak{M}}(d).\alpha_{i-l}$ it also gives $\text{prof}_n^{\mathfrak{M}'}(a').\alpha_i = \text{prof}_n^{\mathfrak{M}}(a').\alpha_i$.

It remains to see that $\text{prof}_n^{\mathfrak{M}'}(a').R_i = \text{prof}_n^{\mathfrak{M}}(a').R_i$ for $0 < i \leq n$. To show that $\text{prof}_n^{\mathfrak{M}'}(a').R_i \subseteq \text{prof}_n^{\mathfrak{M}}(a').R_i$ take any k -ordered type γ belonging to $\text{prof}_n^{\mathfrak{M}'}(a').R_i$ and let a'_1, \dots, a'_k be a realization of γ in \mathfrak{M}' such that $\mathfrak{M}' \models a' + i < a'_1 < \dots < a'_k$. See Fig. 1. Let us divide the sequence a'_1, \dots, a'_k into three (possibly empty) fragments:



■ **Figure 1** Contraction of \mathfrak{M} into \mathfrak{M}' . The tuple a_1, \dots, a_{10} has the same ordered type as a'_1, \dots, a'_{10} .

- a'_1, \dots, a'_s containing all elements a'_j such that $\mathfrak{M}' \models a'_j \leq c$,
- $a'_{s+1}, \dots, a'_{s+t}$ which is empty if $a'_s \neq c$; and is the maximal fragment of consecutive elements located just to the right from c , i.e., such that in \mathfrak{M}' we have $a'_{s+i} = c + i$ for $1 \leq i \leq t$, and $a'_{s+t+1} \neq c + t + 1$,
- a'_{t+1}, \dots, a'_k being the remaining fragment.

In the example from Fig. 1 we have $s = 5$, and $t = 2$. We show that there is a realization a_1, \dots, a_k of γ in \mathfrak{M} such that $a' + i < a_1 < \dots < a_k$. For $j \leq s$ it suffices to take $a_j := a'_j$. For $s < j \leq s+t$ we take $a_j := c + (j - s)$ (as computed in \mathfrak{M}). Finally, for $s+t < j \leq k$, we again take $a_j := a'_j$. It is readily verified that $\text{ordtype}^{\mathfrak{M}}(a_1, \dots, a_k) = \text{ordtype}^{\mathfrak{M}'}(a'_1, \dots, a'_k) = \gamma$, and thus $\gamma \in \text{prof}_n^{\mathfrak{M}}(a').R_i$.

To show that $\text{prof}_n^{\mathfrak{M}}(a').R_i \subseteq \text{prof}_n^{\mathfrak{M}'}(a').R_i$ we take any k -ordered type γ belonging to $\text{prof}_n^{\mathfrak{M}}(a').R_i$ and let a_1, \dots, a_k be a realization of γ in \mathfrak{M} such that $\mathfrak{M} \models a' + i < a_1 < \dots < a_k$. We again split the sequence a_1, \dots, a_k into three (possibly empty) fragments:

- a_1, \dots, a_s containing all elements a_j such that $\mathfrak{M} \models a_j \leq c$,
- a_{s+1}, \dots, a_{s+t} which is empty if $a_s \neq c$; and is the maximal fragment of consecutive elements located just to the right from c , i.e., such that in \mathfrak{M} we have $a_{s+i} = c + i$ for $1 \leq i \leq t$, and $a_{s+t+1} \neq c + t + 1$,
- a_{t+1}, \dots, a_k being the remaining fragment.

Note that this time the second fragment belongs to $\mathfrak{M}_2 d \mathfrak{M}_3$. We show that there is a realization a'_1, \dots, a'_k of γ in \mathfrak{M}' such that $\mathfrak{M}' \models a' + i < a'_1 < \dots < a'_k$. For $j \leq s$ it suffices to take $a'_j := a_j$. For $s < j \leq s+t$ we take $a'_j := c + (j - s)$ (note that $c + (j - s)$ is computed in \mathfrak{M}' so it belongs to M_3). Finally, let $\gamma' = \text{ordtype}^{\mathfrak{M}}(a_{s+t+1}, \dots, a_k)$. Observe that $\gamma' \in \text{prof}_n^{\mathfrak{M}}(c).R_{t+1}$. Thus $\gamma' \in \text{prof}_n^{\mathfrak{M}}(d).R_{t+1}$, and this means that there is a realization of γ in \mathfrak{M}_3 , to the right from $d + t + 1$. In \mathfrak{M}' the same realization is located to the right from $c + t + 1$ and we can take its elements as a'_{s+t+1}, \dots, a'_k . Again, it is readily verified that $\text{ordtype}^{\mathfrak{M}}(a_1, \dots, a_k) = \text{ordtype}^{\mathfrak{M}'}(a'_1, \dots, a'_k) = \gamma$, and thus $\gamma \in \text{prof}_n^{\mathfrak{M}'}(a').R_i$.

The case when $a' \in M_3$ can be treated symmetrically: this time we get equality of the R_i components of profiles for free and to show equality of the L_i components we use the equality of the L_i components of the profiles of c and d . ◀

Let us now observe how the notion of an n -profile is closely related to the satisfaction of normal form $F_1[<, +1]$ formulas.

► **Lemma 8.** *Let φ be a normal form $F_1[<, +1]$ formula of width n and let \mathfrak{M} be a word or ω -word such that $\mathfrak{M} \models \varphi$. Let \mathfrak{M}' be a word or ω -word such that for each $a' \in M'$ there is $a \in M$ such that $\text{prof}_n^{\mathfrak{M}'}(a') = \text{prof}_n^{\mathfrak{M}}(a)$. Then $\mathfrak{M}' \models \varphi$.*

Proof. This observation follows from the definition of the notion of a profile and the definition of the shape of normal form formulas. Below we present some details.

Existential conjuncts. Let us first check that all elements of \mathfrak{M}' have appropriate witnesses for the existential conjuncts. Consider any conjunct $\forall y_0 \exists y_1 \dots y_{k_i} \varphi_i^{\exists}$ of φ and any element $a'_0 \in M'$. Let $a_0 \in M$ be such that $\text{prof}_n^{\mathfrak{M}}(a_0) = \text{prof}_n^{\mathfrak{M}'}(a'_0)$. Let $a_1, \dots, a_{k_i} \in M$ be a tuple of (not necessarily distinct) elements of M forming a witness tuple for a and φ_i^{\exists} , i.e., such that $\mathfrak{M} \models \varphi_i^{\exists}[a_0, a_1, \dots, a_{k_i}]$. Let $b_0, \dots, b_{k'_i}$ be the sequence of distinct elements of M such that $\mathfrak{M} \models b_0 < b_1 < \dots < b_{k'_i}$ and $\{b_0, \dots, b_{k'_i}\} = \{a_0, a_1, \dots, a_{k_i}\}$. (Note that k'_i may be smaller than k_i because of potential equalities among a'_i 's.) Assume that $a_0 = b_u$, and that $b_s, \dots, b_u, \dots, b_t$ is the maximal subsequence of $b_0, \dots, b_{k'_i}$ consisting of consecutive elements of \mathfrak{M} , containing $b_u = a_0$. Our aim now is to demonstrate that there exists a sequence $b'_0, \dots, b'_{k'_i}$ in \mathfrak{M}' such that $\text{ordtype}^{\mathfrak{M}}(b_0, \dots, b_{k'_i}) = \text{ordtype}^{\mathfrak{M}'}(b'_0, \dots, b'_{k'_i})$ and $b'_u = a'_0$. Such a sequence can be defined in the following way:

- Let $\gamma = \text{ordtype}^{\mathfrak{M}}(b_0, \dots, b_{s-1})$. Observe that $\gamma \in \text{prof}_n^{\mathfrak{M}}(a_0).L_{-(u-s+1)}$ and thus also $\gamma \in \text{prof}_n^{\mathfrak{M}'}(a'_0).L_{-(u-s+1)}$. This guarantees that there is a realization of γ in \mathfrak{M}' to the left from $a'_0 - (u - s + 1)$. We take the elements of this realization as b'_0, \dots, b'_{s-1} .
- for $s \leq j \leq t$ we take $b'_s = a'_0 + (j - u)$.
- Let $\gamma' = \text{ordtype}^{\mathfrak{M}}(b_{t+1}, \dots, b_{k'_i})$. Observe that $\gamma' \in \text{prof}_n^{\mathfrak{M}}(a_0).R_{t-u+1}$ and thus also $\gamma' \in \text{prof}_n^{\mathfrak{M}'}(a'_0).R_{t-u+1}$. This guarantees that there is a realization of γ' in \mathfrak{M}' to the right from $a' + (t - s + 1)$. We take the elements of this realization as $b'_{t+1}, \dots, b'_{k'_i}$.

It is readily verified that $\text{ordtype}^{\mathfrak{M}'}(b'_0, \dots, b'_{k'_i})$ is as desired. For $1 \leq j \leq k_i$ we take $a'_j := b'_k$ for such that k that $a_j = b_k$. It follows that $\mathfrak{M}' \models \varphi_i^{\exists}[a'_0, a'_1, \dots, a'_{k_i}]$.

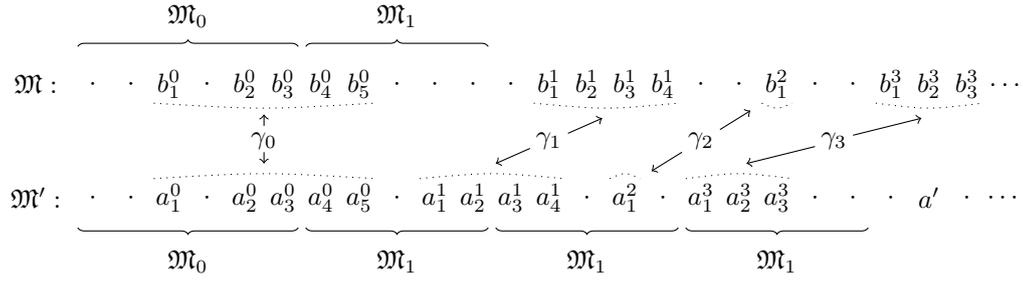
Universal conjuncts. Consider now a conjunct $\forall x_1 \dots x_{l_i} \varphi_i^{\forall}(x_1, \dots, x_{l_i})$ of φ . Let a'_1, \dots, a'_{l_i} be any sequence of (not necessarily distinct) elements of M' . We want to see that $\mathfrak{M}' \models \varphi_i^{\forall}[a'_1, \dots, a'_{l_i}]$. Let $b'_1, \dots, b'_{l'_i}$ be the sequence of elements such that $\mathfrak{M}' \models b'_1 < \dots < b'_{l'_i}$ and $\{b'_1, \dots, b'_{l'_i}\} = \{a'_1, \dots, a'_{l_i}\}$. It is sufficient to show that there is a sequence $b_1, \dots, b_{l'_i}$ of elements of \mathfrak{M} such that $\text{ordtype}^{\mathfrak{M}}(b_1, \dots, b_{l'_i}) = \text{ordtype}^{\mathfrak{M}'}(b'_1, \dots, b'_{l'_i})$. Assume that b'_1, \dots, b'_s is the maximal prefix of $b'_1, \dots, b'_{l'_i}$ consisting of consecutive elements of \mathfrak{M}' . By assumption there is $b_1 \in M$ such that $\text{prof}_n^{\mathfrak{M}}(b_1) = \text{prof}_n^{\mathfrak{M}'}(b'_1)$. We now take $b_j := b_1 + (j - 1)$ for $1 \leq j \leq s$. Let $\gamma = \text{ordtype}^{\mathfrak{M}'}(b'_{s+1}, \dots, b'_{l'_i})$. Observe that $\gamma \in \text{prof}_n^{\mathfrak{M}'}(b'_1).R_s$. Thus also $\gamma \in \text{prof}_n^{\mathfrak{M}}(b_1).R_s$, which implies that there is a realization of γ in \mathfrak{M} to the right from $b_1 + s$. The elements of this realization are taken respectively as $b_{s+1}, \dots, b_{l'_i}$. It is readily verified that $\text{ordtype}^{\mathfrak{M}}(b_1, \dots, b_{l'_i}) = \text{ordtype}^{\mathfrak{M}'}(b'_1, \dots, b'_{l'_i})$ as desired. ◀

4.3 Surgery on ω -words

In this subsection we work over ω -words. Namely we show how to transform a given ω -word into a periodic one without introducing any new profiles.

► **Lemma 9.** *Let \mathfrak{M} be an ω -word and $n > 0$ a natural number. Let \mathfrak{M}_0 be the shortest prefix of \mathfrak{M} such that it contains all the elements having the n -profiles which are realized finitely many times in \mathfrak{M} . Let a be the first element not belonging to M_0 , and π its n -profile. Let \mathfrak{M}_1 be the shortest fragment of \mathfrak{M} such that*

- *it starts at a ,*
- *contains the realizations of all n -profiles realized in \mathfrak{M} infinitely many times,*



■ **Figure 2** Building a periodic model \mathfrak{M}' from \mathfrak{M} .

- ends at an element whose successor has n -profile π ,
- and its length is greater than n .

Consider the ω -word $\mathfrak{M}' = \mathfrak{M}_0 \mathfrak{M}_1^\omega$. For each $a' \in M'$, if $a' \in M_0$ then $\text{prof}_n^{\mathfrak{M}'}(a') = \text{prof}_n^{\mathfrak{M}}(a')$, and if a' belongs to a copy of \mathfrak{M}_1 then $\text{prof}_n^{\mathfrak{M}'}(a') = \text{prof}_n^{\mathfrak{M}}(a)$ for $a \in M_1$ being the element whose a' is a copy of.

Proof. We consider explicitly the case when a' is a copy of an element $a \in M_1$. (The case of $a' \in M_0$ is similar and simpler.) It should be clear that for $-n \leq i \leq n$ we have $\text{prof}_n^{\mathfrak{M}'}(a') = \text{prof}_n^{\mathfrak{M}}(a)$. Since the prefix of \mathfrak{M} ending at the predecessor of a is also a prefix of \mathfrak{M}' , and a' is located to the right from it in \mathfrak{M}' , it follows that for all $-n \leq i < 0$ we have $\text{prof}_n^{\mathfrak{M}}(a).L_i \subseteq \text{prof}_n^{\mathfrak{M}'}(a').L_i$. Let us show the opposite containment. Take any $\gamma \in \text{prof}_n^{\mathfrak{M}'}(a').L_i$ and its realization \bar{a}'_π located to the left from $a' + i$. Let us write the elements of \bar{a}'_π , in the increasing order, as $a_1^0, \dots, a_u^0, \dots, a_{s_0}^0, a_1^1, \dots, a_{s_1}^1, \dots, a_1^k, \dots, a_{s_k}^k$, where (cf. Fig. 2):

- a_1^0, \dots, a_u^0 are all members of \bar{a}'_π from M_0 ; in the example from Fig. 2 we have $u = 3$,
- if a_u^0 is not the rightmost element of \mathfrak{M}_0 then $u = s_0$; otherwise s_0 is chosen so that $\mathfrak{M}' \models \bigwedge_{u \leq j < s_0} +1(a_j^0, a_{j+1}^0)$ and $\mathfrak{M}' \models a_{s_0}^0 \ll a_1^1$; note that since $|M_1| > n$ it follows that all $a_1^0, \dots, a_u^0, \dots, a_{s_0}^0$ belong to $M_0 \cup M_1$; in the example from Fig. 2 we have $s_0 = 5$,
- for $i > 0$ the fragments $a_1^i, \dots, a_{s_i}^i$ are maximal fragments of \bar{a}'_π such that $\mathfrak{M}' \models \bigwedge_{1 \leq j < s_i} +1(a_j^i, a_{j+1}^i)$.

For $i \geq 0$ let $\gamma_i := \text{ordtype}^{\mathfrak{M}'}(a_1^i, \dots, a_{s_i}^i)$. We want to find a realization of γ in \mathfrak{M} . Let us set $b_1^0, \dots, b_u^0, \dots, b_{s_0}^0$ to be a realization of γ_0 from $M_0 \cup M_1$ consisting just of elements $a_1^0, \dots, a_u^0, \dots, a_{s_0}^0$. Observe that for $i > 0$ each of γ_i is realized in \mathfrak{M} infinitely many times. To see this let us denote by $\alpha_1, \dots, \alpha_{s_i}$ the 1-types of $a_1^i, \dots, a_{s_i}^i$, respectively, in \mathfrak{M}' . Recall that a_1^i is a copy of an element b from M_1 whose n -profile is realized in \mathfrak{M} infinitely many times. Due to our construction the elements $b + j$ from \mathfrak{M} , for $1 \leq j \leq s_i$ have 1-types $\alpha_1, \dots, \alpha_{s_i}$, respectively. Thus there are infinitely many tuples of consecutive elements of 1-types $\alpha_1, \dots, \alpha_{s_i}$ in \mathfrak{M} ; any such tuple obviously realizes γ_i . Now we can compose a realization of γ in \mathfrak{M} starting from the chosen realization of γ_0 and iteratively finding a realization of γ_i separated by at least one element from the realization of γ_{i-1} . Moreover, since $\text{prof}_n^{\mathfrak{M}}(a).L_i$ is realized infinitely many times in \mathfrak{M} we can find its realization starting at least n positions to the right from the last element of the chosen realization of γ . This implies that $\gamma \in \text{prof}_n^{\mathfrak{M}}(a).L_i$.

The equality $\text{prof}_n^{\mathfrak{M}}(a).R_i = \text{prof}_n^{\mathfrak{M}'}(a').R_i$ can be proved in a similar fashion. ◀

4.4 Complexity

Lemmas 7, 8 immediately lead to the following small model property.

► **Lemma 10.** *Every normal form $F_1[<, +1]$ formula φ satisfiable over a finite word has a model of size bounded exponentially in $\|\varphi\|$.*

Proof. It suffices to take any finite model of φ and perform the contraction procedure from Lemma 7 as long as possible, i.e., as long as the model contains a pair of elements with the same n -profile (with n – the width of φ). By Lemma 6 the number of elements in the resulting model is bounded exponentially in $\|\varphi\|$. By Lemmas 7, 8 it indeed satisfies φ . ◀

Similarly, using Lemmas 9, 8 we get:

► **Lemma 11.** *Every normal form $F_1[<, +1]$ formula φ satisfiable over an ω -word has a model \mathfrak{M} of the form $\mathfrak{M} = \mathfrak{M}_1\mathfrak{M}_2^\omega$ where both $|M_1|$ and $|M_2|$ are bounded exponentially in $\|\varphi\|$.*

Proof. We start from an arbitrary model and build a periodic model as guaranteed by Lemma 9. Then we shorten its initial and periodic part using Lemma 7, obtaining a model as required. ◀

Finally, we are ready to state our complexity result.

► **Theorem 12.** *The satisfiability problems for $F_1[<, +1]$ over words (ω -words) is NEXPTIME-complete.*

Proof. For a given $F_1[<, +1]$ formula φ convert it into its normal form φ' . Then guess a finite model of φ' of size bounded exponentially as guaranteed by Lemma 10 (exponentially bounded initial and periodic parts of a regular ω -model as guaranteed by Lemma 11) and verify that it is indeed a model of φ' (they generate a model of φ). ◀

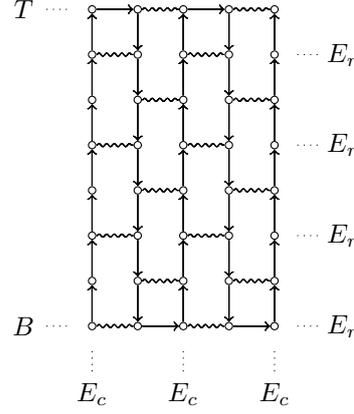
5 Undecidable extensions

5.1 Data words

A *data word* (ω -data word) is a word (ω -word) with an additional binary relation \sim which is required to be interpreted as an equivalence relation, and which is intended to model data equality tests. Data words are motivated by their connections to XML. FO^2 over data words becomes at least as hard as reachability in Petri nets [2]. Nevertheless, the satisfiability problem remains decidable. We show that $F_1[<, +1]$ over data words becomes undecidable. Undecidability can be even shown in the absence of $<$.

► **Theorem 13.** *The satisfiability problem for $F_1[+1, \sim]$ over finite data words and over ω -data-words is undecidable.*

Proof. We employ the standard apparatus of tiling systems. A *tiling system* is a quadruple $\mathcal{T} = \langle C, c_0, Hor, Ver \rangle$, where C is a non-empty, finite set of *colours*, c_0 is an element of C , and Hor, Ver are binary relations on C called the *horizontal* and *vertical* constraints, respectively. We say that \mathcal{T} *tiles* the $m \times n$ grid if there is a function $f : \{0, 1, \dots, m-1\} \times \{0, 1, \dots, n-1\} \rightarrow C$ such that $f(0, 0) = c_0$, for all $0 \leq i < m-1, 0 \leq j \leq n-1$ we have $\langle f(i, j), f(i+1, j) \rangle$ is in Hor , and for all $0 \leq i < m, 0 \leq j < n-1$ we have $\langle f(i, j), f(i, j+1) \rangle$ is in Ver . It is well known that the problem of checking if for a given tiling system \mathcal{T} there



■ **Figure 3** The grid-like structure used to show undecidability of $F_1[+1, \sim]$.

are m, n such that \mathcal{T} tiles the $m \times n$ grid is undecidable. The problem remains undecidable if we require m to be even and n odd.

To show undecidability of the satisfiability problem for $F_1[+1, \sim]$ over finite words we construct a formula $\Phi_{\mathcal{T}}$ which is satisfied in a finite word iff \mathcal{T} tiles the $m \times n$ grid for some even m and odd n . We begin the construction of $\Phi_{\mathcal{T}}$ with enforcing that its model is a finite grid-like structure, in which the relation $+1$ forms a snake-like path from its lower-left corner to the upper-right corner, and the equivalence relation connects some elements from neighbouring columns. See Fig. 3. As mentioned, we assume that the number of columns is odd and the number of rows is even. We employ the following unary predicates: B, T, E_c, E_r , whose intended purpose is to mark elements in the bottom row, top row, even columns, and even rows, respectively.

The first two formulas say that the lower left and upper right corners of the grid exist:

$$\exists x(Bx \wedge \neg Tx \wedge E_c x \wedge E_r x \wedge \neg \exists y(+1(y, x))) \quad (9)$$

$$\exists x(Tx \wedge \neg Bx \wedge E_c x \wedge \neg E_r x \wedge \neg \exists y(+1(x, y))) \quad (10)$$

Next we take care of $+1$ relation, ensuring that it respects the intended meaning of the unary predicates:

$$\begin{aligned} \forall xy \quad (+1(x, y) \rightarrow & \quad (11) \\ & (E_c x \wedge E_c y \rightarrow (\neg B y \wedge \neg T x \wedge (E_r x \leftrightarrow \neg E_r y))) \wedge \\ & (E_c x \wedge \neg E_c y \rightarrow (T x \wedge T y \wedge \neg B x \wedge \neg B y \wedge \neg E_r x \wedge \neg E_r y)) \wedge \\ & (\neg E_c x \wedge E_c y \rightarrow (B x \wedge B y \wedge \neg T x \wedge \neg T y \wedge E_r x \wedge E_r y)) \wedge \\ & (\neg E_c x \wedge \neg E_c y \rightarrow (\neg B x \wedge \neg T y \wedge (E_r \leftrightarrow \neg E_r y)))) \end{aligned}$$

Further, we enforce the appropriate \sim -connections. (We abbreviate a formula guaranteeing that x_1, \dots, x_k agree on E_c -predicate by $SameColumn(x_1, \dots, x_k)$.)

$$\forall xyz t(+1(x, y) \wedge +1(y, z) \wedge +1(z, t) \wedge T y \wedge T z \rightarrow x \sim t) \quad (12)$$

$$\forall xyz t(+1(x, y) \wedge +1(y, z) \wedge +1(z, t) \wedge B y \wedge B z \rightarrow x \sim t) \quad (13)$$

$$\begin{aligned} \forall xyz t u w (SameColumn(x, y, z) \wedge SameColumn(t, u, w) \wedge \\ +1(x, y) \wedge +1(y, z) \wedge z \sim t \wedge +1(t, u) \wedge +1(u, w) \rightarrow x \sim w) \quad (14) \end{aligned}$$

And finally, we say that T and B are appropriately propagated.

$$\forall xy(x \sim y \rightarrow (Tx \leftrightarrow Ty) \wedge (Bx \leftrightarrow By)) \quad (15)$$

$$\begin{aligned} \forall xyzt(\text{SameColumn}(x, y) \wedge \text{SameColumn}(z, t) \wedge \\ +1(x, y) \wedge y \sim z \wedge +1(z, t) \rightarrow (Tx \leftrightarrow Tt) \wedge (Bx \leftrightarrow Bt)) \end{aligned} \quad (16)$$

Formulas (9)-(16) ensure that all the vertical segments of the snake-like path are of the same length and thus that any model indeed looks as in Fig. 3. It remains to encode the tiling problem. We use a unary predicate P_c for each $c \in C$. We say that each node of the grid is coloured by precisely one colour from C and that $(0, 0)$ is coloured by c_0 :

$$\forall x(\bigvee_{c \in C} P_c(x) \wedge \bigwedge_{c \neq d} \neg(P_c(x) \wedge P_d(x))), \quad (17)$$

$$\forall x((\neg \exists y +1(y, x)) \rightarrow P_{c_0}(x)). \quad (18)$$

Let us abbreviate by $\Theta_H(x, y)$ the formula $\bigwedge_{\langle c, d \rangle \notin \text{Hor}} (\neg P_c(x) \wedge \neg P_d(y))$ stating that x, y respect the horizontal constraints of \mathcal{T} and by $\Theta_V(x, y)$ the analogous formula for vertical constraints. We take care of vertical adjacencies:

$$\forall xy(E_c(x) \wedge E_c(y) \wedge +1(x, y) \vee \neg E_c(x) \wedge \neg E_c(y) \wedge +1(y, x) \rightarrow \Theta_V(x, y)), \quad (19)$$

and of horizontal adjacencies:

$$\forall xyzt(+1(x, y) \wedge +1(y, z) \wedge +1(z, t) \wedge Ty \wedge Tz \rightarrow \Theta_H(y, z)), \quad (20)$$

$$\forall xyzt(+1(x, y) \wedge +1(y, z) \wedge +1(z, t) \wedge By \wedge Bz \rightarrow \Theta_H(y, z)), \quad (21)$$

$$\begin{aligned} \forall xyztuw(\text{SameColumn}(x, y, z) \wedge \text{SameColumn}(t, u, w) \wedge z \sim t \wedge \\ +1(x, y) \wedge +1(y, z) \wedge +1(t, u) \wedge +1(u, w) \rightarrow \Theta_H(x, w) \wedge \Theta_H(y, u) \wedge \Theta_H(z, t)). \end{aligned} \quad (22)$$

Let $\Phi_{\mathcal{T}}$ be the conjunction of (9)-(22). From any model of $\Phi_{\mathcal{T}}$, we can read off a tiling of an $m \times n$ grid by inspecting the colours assigned to the elements of the model. On the other hand, given any tiling for \mathcal{T} , we can construct a finite model of $\Phi_{\mathcal{T}}$ in the obvious way. We leave the detailed arguments to the reader.

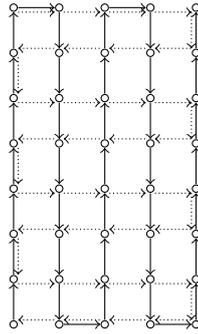
The case of ω -words can be treated essentially in the same way. We just mark one element in a model, corresponding to the upper-right corner of the grid, with a special unary symbol, and relativize all our formulas to positions smaller than this element (marked with another fresh unary symbol). In effect, it is irrelevant what happens in the infinite fragment of a model starting in this marked element.

What is probably worth commenting is that in our undecidability proof we use the equivalence relation \sim in a very limited way, actually not benefiting from its transitivity or symmetry. In fact, the transitivity of \sim does not help, being rather an obstacle in our construction. \blacktriangleleft

5.2 Other variants

Both $\text{FO}^2[<]$ and $\text{FO}^2[+1]$ remain decidable when, besides $<$ or $+1$, the signature may contain other binary symbols, whose interpretation is not restricted ([15], [3]). We can easily see that this is not the case for F_1 .

► **Theorem 14.** *The satisfiability problem for $F_1[+1]$ and $F_1[<]$ is undecidable when an additional uninterpreted binary relation is available.*



■ **Figure 4** The grid-like structure used to show undecidability of $F_1[+1_a, +1_b]$.

Actually, the proof is trivial and can be obtained even without using the linear order (just a simple grid axiomatization using a single binary predicate and some unary coordinate predicates). See [8] for a proof using two binary symbols.

There is a huge list of other, more specialized variations, which was studied in the context of FO^2 , and which thus may also be considered here. One of the options is to have two linear orders rather than just one. The second linear order may be interpreted, e.g., as a comparison relation on data values. $FO^2[+1_a, +1_b]$, the two-variable fragment accessing the linear orders through their successor relations only, is decidable in $NEXPTIME$ [3]. Showing that a corresponding variant of F_1 is undecidable is again easy. We can define a grid-like structure using the first linear order to form a snake-like path as in the proof of Thm. 12 and the second to form another snake-like path, starting in the upper-left corner, ending in the lower-left corner and going horizontally through our grid, with steps down only on the borders. See Fig. 4. The required structure can be defined with help of some additional unary predicates. Since the details of the construction do not differ significantly from the details of the proof of Thm. 12 we omit them here.

► **Theorem 15.** *The satisfiability problem for $F_1[+1_a, +1_b]$ is undecidable.*

6 Conclusion

This article is a starting point for investigations of the one-dimensional fragment of first-order logic over restricted classes of structures. We proved that $F_1[<, +1]$ is expressively equivalent over words and ω -words to $FO^2[<, +1]$, and that it also retains the complexity of the satisfiability problem of the latter. We argued on the other hand that some natural extensions of $F_1[<, +1]$, whose two-variable counterparts remain decidable, become undecidable. One of such extensions are data words. Regarding the case of words we leave some open questions, the most important of which is if F_1 is more succinct than FO^2 . Our working hypothesis is that it is true. A positive answer would be a good motivation for $F_1[<, +1]$, showing that it allows to describe some properties of words not only in a more elegant, but simply also in a shorter way.

Our plan is also to check if the techniques developed in this paper can be used to show that the satisfiability problem of F_1 over trees retains the $EXSPACE$ -complexity of FO^2 [1]. We suspect that it is true, even though this time the expressive power of both formalisms seems to differ: e.g., over unordered words, having access only to the descendant predicate, one can say in F_1 that there are at least three nodes satisfying a unary predicate P , such that none of them is a descendant of other, which is inexpressible in FO^2 .

As mentioned in Section 5.2 there are a couple of specialized variations whose satisfiability may be considered. We suggest two of them here:

1. $F_1[<_a, <_b]$ (two linear orders accessed by *less then* predicates); in the case of FO^2 this variant is decidable at least over finite models [16].
2. $F_1[<, \sim]$ (data words, without access to the successor relation); this variant is NEXPTIME-complete for FO^2 [2]; moreover, in the case of FO^2 it remains decidable even if additional uninterpreted binary symbols are allowed [14]; of course, due to Thm. 14 we surely will not be able to lift the latter result to F_1 .

References

- 1 Saguy Benaim, Michael Benedikt, Witold Charatonik, Emanuel Kieroński, Rastislav Lenhardt, Filip Mazowiecki, and James Worrell. Complexity of two-variable logic on finite trees. In *ICALP (2)*, pages 74–88, 2013. doi:10.1007/978-3-642-39212-2_10.
- 2 Mikołaj Bojańczyk, Claire David, Anca Muscholl, Thomas Schwentick, and Luc Segoufin. Two-variable logic on data words. *ACM Trans. Comput. Log.*, 12(4):27, 2011.
- 3 Witold Charatonik and Piotr Witkowski. Two-variable logic with counting and trees. In *LICS*, pages 73–82, 2013.
- 4 Witold Charatonik and Piotr Witkowski. Two-variable logic with counting and a linear order. In *Computer Science Logic*, volume 41 of *LIPICs*, pages 631–647, 2015.
- 5 Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite model theory*. Perspectives in Mathematical Logic. Springer, 1995.
- 6 Kousha Etessami, Moshe Y. Vardi, and Thomas Wilke. First-order logic with two variables and unary temporal logic. *Inf. Comput.*, 179(2):279–295, 2002. doi:10.1006/inco.2001.2953.
- 7 Erich Grädel, Phokion Kolaitis, and Moshe Y. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.
- 8 Lauri Hella and Antti Kuusisto. One-dimensional fragment of first-order logic. In *Advances in Modal Logic 10*, pages 274–293, 2014.
- 9 Emanuel Kieronski and Antti Kuusisto. Complexity and expressivity of uniform one-dimensional fragment with equality. In *Mathematical Foundations of Computer Science, Part I*, pages 365–376, 2014.
- 10 Emanuel Kieronski and Antti Kuusisto. Uniform one-dimensional fragments with one equivalence relation. In *Computer Science Logic*, volume 41 of *LIPICs*, pages 597–615, 2015.
- 11 Emanuel Kieroński and Martin Otto. Small substructures and decidability issues for first-order logic with two variables. *Journal of Symbolic Logic*, 77:729–765, 2012.
- 12 Andreas Krebs, Kamal Lodaya, Paritosh Pandya, and Howard Straubing. Two-variable logic with a between predicate. In *LICS*, 2016.
- 13 Antti Kuusisto. On the uniform one-dimensional fragment. In *Proceedings of International Workshop on Description Logic*, 2016.
- 14 Angelo Montanari, Marco Pazzaglia, and Pietro Sala. Metric propositional neighborhood logic with an equivalence relation. In *TIME*, pages 49–58, 2014.
- 15 Martin Otto. Two-variable first-order logic over ordered domains. *Journal of Symbolic Logic*, 66:685–702, 2001.
- 16 Thomas Schwentick and Thomas Zeume. Two-variable logic with two order relations. *Logical Methods in Computer Science*, 8(1), 2012. doi:10.2168/LMCS-8(1:15)2012.
- 17 Dana Scott. A decision method for validity of sentences in two variables. *Journal Symbolic Logic*, 27:477, 1962.
- 18 Larry J. Stockmeyer. *The Complexity of Decision Problems in Automata Theory and Logic*. PhD thesis, MIT, Cambridge, Massasuchets, USA, 1974.
- 19 Wiesław Szwał and Lidia Tendera. FO^2 with one transitive relation is decidable. In *STACS*, pages 317–328, 2013.