Satisfiability for two-variable logic with two successor relations on finite linear orders[☆]

Diego Figueira

University of Edinburgh

Abstract

We study the finite satisfiability problem for first order logic with two variables and two binary relations, corresponding to the induced successor relations of two finite linear orders. We show that the problem is decidable in NExpTime.

1. Introduction

First-order logic with two variables (henceforth denoted by FO^2) is of importance in computer science due to its decidable satisfiability problem (contrary to fragments of FO with 3 or more variables), and since it has connections with many formalisms, such as modal, temporal or description logics. Many fragments of FO^2 have been studied because of this, especially in the presence of linear orders or equivalence relations. There are, still, a few relevant basic problems that remain open, and our work aims at expanding the classification of FO^2 in the presence of linear orders. In this setting, linear orders are related with temporal logics, but it is also applicable in other scenarios, like in databases or description logics.

We study the two variable fragment of first-order logic with two variables and two successor relations on two finite linear orders. We show that the problem is decidable in NExpTime. This bound is optimal, since the problem is NExpTime-hard [4]. This logic has been previously claimed to be decidable in 2NExpTime in [12], but the proof was flawed.[1] As a corollary of the results from the report [14], this logic is shown to be decidable with a non-primitive-recursive algorithm.[2] Our result also trivially extends to the satisfiability of existential monadic second order logic with two variables ($EMSO^2$) and two successor relations on finite linear orders.

This work focuses on the *finite* satisfiability problem and hence all the results discussed next are relative to finite structures. FO^2 is a well-known decidable fragment of first-order logic. Over arbitrary relations, it is known to be decidable [15], NExpTime-complete [5]. FO^2 over words (*i.e.*, with two relations: a successor relation over a finite linear order, and its transitive closure) is NExpTime-complete [4]. The satisfiability problem was shown to be undecidable: in the presence of two transitive relations (even without equality) the satisfiability problem is undecidable [8]; in the presence of one transitive relation and one equivalence relation [11]; in the presence of three linear orders [9]; or in the presence of three equivalence relations [10]. However, if it has only two equivalence relations it is decidable [11]. Over words with one equivalence relation it is decidable [1]. If it only has a transitive closure over a finite linear order and an equivalence relation, then it is NExpTime-complete [1]. If it only has a successor relation over a finite linear order and an equivalence relation, it is in 2NExpTime [1]. On trees, with only successor relations (*i.e.*, the *child* and *next sibling* relations) and an equivalence relation it is decidable in 3NExpTime [2].

There have also been works in the presence of

[☆]Work supported by the FET-Open grant agreement FOX, number FP7-ICT-233599.

[1]In fact, according to its author, the proof of Lemma 4 in [12] is wrong [13], and there does not appear to be an easy way of fixing it. This is a key lemma employed to obtain the decidability results contained in [12]. (Of course, this does not affect the undecidability results contained in [12].) Here we adopt a different strategy to prove decidability.

[2]By this we mean an algorithm whose time or space is not bounded by any primitive-recursive function.

a linear order and a linear preorder [17, 14]. In the presence of two finite linear orders, if there is a successor and its transitive closure over one linear order, and a successor over another linear order, it is decidable, and as hard as reachability of VAS according to the report [14]. If there are only two successors, it is known to be NExpTime-hard. Indeed, it is already NExpTime-hard even when no binary relations are present [4]. Here, we show that it is indeed NExpTime-complete. In fact, it sits in the same complexity class as FO^2 with just one successor relation on a linear order.

2. Preliminaries

Let $\mathbb{N} = \{1, 2, \dots\}$, $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$, and for every $n, m \in \mathbb{N}$, $[n, m] = \{i \in \mathbb{N} \mid n \leq i \leq m\}$, $[n] = [1, n]$. Given a function $f : A \to B$ and a set $A' \subseteq A$, by $f|_{A'}$ we denote f restricted to the elements of A', and by $f[a \mapsto b] : A \cup \{a\} \to B \cup \{b\}$ we denote the function where $f[a \mapsto b](a) = b$ and $f[a \mapsto b](a') = f(a')$ for all $a' \in A \setminus \{a\}$. We write $[a \mapsto b]$ to denote $\iota\,[a \mapsto b]$, where ι is the identity function, and $[a \mapsto b, a' \mapsto b']$ to denote $[a \mapsto b][a' \mapsto b']$. For any number n, we denote by $|n|$ its absolute value. We write $\#S$ to denote the number of elements of a set S. Given a string $w \in A^*$, we write $|w|$ to denote the length of w, $w[i, j]$ to denote the subword of w restricted to positions $[i, j]$, and $w[i]$ to denote $w[i, i]$, for any $1 \leq i \leq j \leq |w|$.

2.1. Permutations

First-order structures with n elements and two linear orders can be naturally represented as a permutation on $[n]$ with n valuations.[3] This representation will prove useful in the proofs that follow. A permutation of $[n]$ is represented as a set $\pi \subseteq [n] \times [n]$ so that for every $i \in [n]$, π has exactly one pair with i in the first component, and exactly one pair with i in the second component. We will normally use the symbols $(r, c), (s, d)$ to denote elements of a permutation.

Formally, given $n \in \mathbb{N}$, we say that $\pi \subseteq [n] \times [n]$ is an ***n*-permutation** if for every $k \in [n]$ we have $\#\{c \mid (k, c) \in \pi\} = \#\{r \mid (r, k) \in \pi\} = 1$. We say that π is a permutation if it is an n-permutation for some n. Given a permutation π and $(r, c), (r', c') \in \pi$ we say that the **neighborhood type** of $(r, c), (r', c')$ in π is an element $t \in \{\bullet, \nearrow, \uparrow, \nwarrow, \leftarrow, \swarrow, \downarrow, \searrow, \rightarrow, \infty\}$ so that: $t \in \{\searrow, \downarrow, \swarrow\}$ iff $r' - r = 1$, $t \in \{\nearrow, \uparrow, \nwarrow\}$ iff $r' - r = -1$, $t \in \{\nearrow, \rightarrow, \searrow\}$ iff $c' - c = 1$, $t \in \{\nwarrow, \leftarrow, \swarrow\}$ iff $r' - r = -1$, and $t = \bullet$ iff $(r, c) = (r', c')$. We denote it with $[(r, c), (r', c')]_\pi$. Figure 1-*a* contains a graphical representation of a 4-permutation, where $[(2, 2), (3, 4)]_\pi = \downarrow$, $[(2, 2), (1, 3)]_\pi = \nearrow$, $[(1, 3), (4, 1)]_\pi = \infty$.

Let us fix $\mathbb{V}$ to be an enumerable set of propositional letters. A **valued permutation** is a pair (π, σ) consisting of a permutation π and a function $\sigma : \pi \to 2^{\mathbb{V}}$ that assigns a set of propositional letters to each element of π. We say that $\sigma(r, c) \subseteq \mathbb{V}$ is the **valuation** of (r, c) in (π, σ). Since for every r [resp. for every c] there is only one c [resp. only one r] such that $(r, c) \in \pi$ we use the notation $\sigma(r, \_)$ [resp. $\sigma(\_, c)$] to denote $\sigma(r, c)$ for the only c [resp. only r] such that $(r, c) \in \pi$. For example, the valued permutation of Figure 1-*b* is such that $\sigma(3, \_) = \sigma(3, 4) = \{r, q\}$. A valued permutation can be seen as a finite first-order structure with two linear orders, where the permutation element (r, c) represents the r-th element in the first linear order, and the c-th element in the second linear order, and its valuation is $\sigma(r, c)$. Likewise, any finite first-order structure with two linear orders can be represented by a valued permutation.

For convenience in our proofs, we also represent an n-permutations with any set $\pi' \subseteq S \times T$ with $S, T \subseteq \mathbb{N}$, $\#S = \#T = n$ such that for every $(c, r) \in S \times T$, $\#\{c' \mid (r, c') \in \pi'\} = \#\{r' \mid (r', c) \in \pi'\} = 1$. In this case we say that π' is a **permutation over** $S \times T$. Note that for every n-permutation π' over $S \times T$ there is an n-permutation π and a bijection $f : \pi' \to \pi$ that preserves the order of the elements: for all $f(r, c) = (r', c')$, $f(s, d) = (s', d')$, we have that $r \leq s$ iff $r' \leq s'$ and $c \leq d$ iff $c' \leq d'$. For any $(r, c), (r', c') \in \pi'$ we define $[(r, c), (r', c')]_{\pi'} = [f(r, c), f(r', c')]_\pi$.

Valued permutations represent, precisely, first-order finite structures with two linear orders. We define then the semantics of $\mathrm{FO}^2(\rightarrow, \downarrow)$ over valued permutations.

2.2. FO^2 with two linear orders

We define FO^2 on finite structures with the induced successor relations of two finite linear orders, that we denote by $\mathrm{FO}^2(\rightarrow, \downarrow)$. The atoms of $\mathrm{FO}^2(\rightarrow, \downarrow)$ are: $p(a)$, $a{\rightarrow}b$, and $a{\downarrow}b$, for every $a, b \in \{x, y\}$ and every propositional letter $p \in \mathbb{V}$.

[3] We do not claim that we are the first to use this encoding, it may have been used before.

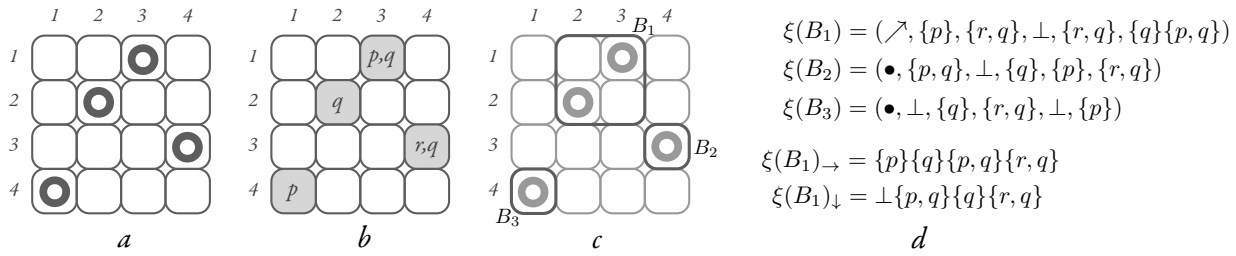

Figure 1: Representation of: (*a*) a 4-permutation $\pi = \{(4,1),(2,2),(1,3),(3,4)\}$; (*b*) a valued permutation (π,σ) where, *e.g.*, $\sigma(1,3) = \{p,q\}$; (*c*) the maximal blocks of π; (*d*) the fingerprints of (π,σ).

If φ, ψ are formulas of FO^2, so are $\exists a.\varphi$, $\forall a.\varphi$, $\neg\varphi$, $\varphi \wedge \psi$, $\varphi \vee \psi$, where $a \in \{x, y\}$. For any $\varphi \in \mathrm{FO}^2(\rightarrow,\downarrow)$, let $\mathcal{V}_\varphi \subseteq \mathbb{V}$ be the set of all propositional variables occurring in φ.

Semantics. We define $\mathrm{FO}^2(\rightarrow,\downarrow)$ on valued permutations. The semantics are as expected, we give only some cases to fix notation. Here, (π,σ) is a valued permutation and μ is a partial function $\mu : \{x, y\} \to \pi$.

$(\pi,\sigma) \models_\mu p(x)$	if $p \in \sigma(\mu(x))$
$(\pi,\sigma) \models_\mu \exists x.\varphi$	if for some $(r,c) \in \pi$ we have $(\pi,\sigma) \models_{\mu[x\mapsto(r,c)]} \varphi$
$(\pi,\sigma) \models_\mu (x{\rightarrow}y)$	if for some $(r,c),(r',c+1) \in \pi$ we have $\mu(x) = (r,c)$, $\mu(y) = (r',c+1)$
$(\pi,\sigma) \models_\mu (x{\downarrow}y)$	if for some $(r,c),(r+1,c') \in \pi$ we have $\mu(x) = (r,c)$, $\mu(y) = (r+1,c')$

For any closed formula φ, we define $(\pi,\sigma) \models \varphi$ if $(\pi,\sigma) \models_{v_\emptyset} \varphi$, where $v_\emptyset(x) = v_\emptyset(y) = \bot$. In this case we say that (π,σ) *satisfies* φ. For example, the valued permutation of Figure 1-*b* satisfies the formula $\forall x \forall y.\neg(x{\rightarrow}y \wedge y{\downarrow}x \wedge p(x))$. The **satisfiability problem** for $\mathrm{FO}^2(\rightarrow,\downarrow)$ is then, given a closed formula $\varphi \in \mathrm{FO}^2(\rightarrow,\downarrow)$, whether $(\pi,\sigma) \models \varphi$ for some (π,σ).

Scott normal form. Any formula $\varphi \in \mathrm{FO}^2(\rightarrow,\downarrow)$ can be converted into a satisfiability equivalent formula in Scott normal form, which is of the form

$$\forall x \forall y\ \chi \quad \wedge \quad \bigwedge_i \forall x \exists y\ \psi_i,$$

where χ and all the ψ_i's are quantifier-free formulas of $\mathrm{FO}^2(\rightarrow,\downarrow)$. The resulting formula is linear in terms of the size of the original formula. Further, this reduction is polynomial-time (see, *e.g.*, [6]). Henceforward we assume that all the formulas we work with are in Scott normal form, unless otherwise stated.

3. Results

Theorem 1. The satisfiability problem for $\mathrm{FO}^2(\rightarrow,\downarrow)$ is NExpTime-complete.

As an immediate corollary we have that the same bound holds for $\mathrm{EMSO}^2(\rightarrow,\downarrow)$, where $\mathrm{EMSO}^2(\rightarrow,\downarrow)$ stands for formulas of $\mathrm{FO}^2(\rightarrow,\downarrow)$ prefixed by existential quantification over sets of permutation elements.

Corollary 1. The satisfiability problem for $\mathrm{EMSO}^2(\rightarrow,\downarrow)$ is NExpTime-complete.

Proof sketch. First, in Section 4 we show a property of the blocks of a valued permutation. A block of a permutation can be seen as a set of positions $\{(r,c),(r+1,c+1),\ldots,(r+k,c+k)\}$ (or $\{(r,c),(r+1,c-1),\ldots,(r+k,c-k)\}$) of the permutation.[4] We prove that if a $\mathrm{FO}^2(\rightarrow,\downarrow)$ formula φ is satisfiable, then it is satisfiable in a valued model where every block is of size bounded exponentially in the size of φ. Moreover, the number of different types of blocks that can appear in the valued permutation is also bounded exponentially in the size of φ.

Second, in Section 5 we show a combinatorial proposition. This involves what we call n-permutation *constraints*, which are sets of positions of $[n] \times [n]$ where a permutation satisfying this constraint is not allowed to have an element. We give a sufficient condition on how large n must be to ensure that there exists a permutation satisfying any constraint with a certain property—namely that it has at most 4 elements in any row or column.

Finally, in Section 6 we introduce a problem called the Restricted Labeled Permutation problem (RLP), which we show to be decidable in NP using the result of Section 5. We then show that satisfiability for $\mathrm{FO}^2(\rightarrow,\downarrow)$ can be reduced in NExpTime to the

[4] A similar notion of *block* is also used in [14].

RLP problem, by using the results on the size of the blocks of Section 4. Thus, decidability of the satisfiability problem for $\mathrm{FO}^2(\rightarrow,\downarrow)$ follows, with a tight upper bound of NExpTime.

4. Few and small blocks properties

Definition 1 (Block)**.** Given an n-permutation π we say that $B \subseteq [n] \times [n]$ is a **block** of π if $B = [i, i+k] \times [j, j+k]$ for some $k \in \mathbb{N}_0$ so that $i, j, i+k, j+k \in [n]$, and either

- $\#B = 1$, and in this case we say that B has type '$\bullet$', or, otherwise,
- for every $(r,c),(r+1,c') \in B \cap \pi$, we have $c' = c+1$, and in this case we say that B has type '$\searrow$', or
- for every $(r,c),(r+1,c') \in B \cap \pi$, we have $c' = c-1$, and in this case we say that B has type '$\nearrow$'.

We say that k is the **size** of the block B. A block B is **maximal** if there is no block B' of π with $B \subsetneq B'$. Figure 1-c shows the three maximal blocks of a permutation, one with type $\nearrow$ and two with type $\bullet$.

Proposition 1. Any minimal valued permutation satisfying $\varphi \in \mathrm{FO}^2(\rightarrow,\downarrow)$ is such that every block is of size at most exponential in φ.

In fact, note that $\mathrm{FO}^2(\rightarrow,\downarrow)$ on blocks is basically like $\mathrm{FO}^2(\rightarrow)$ (first order logic with a successor relation on a linear order), where we have the exponential length model property [4]. However, note that a block is within a context of other blocks, and special care must be taken in order to preserve all the elements that may be needed outside the block.

Definition 2 (Fingerprint)**.** Given a maximal block $B = [k, k+n] \times [l, l+n]$ of a valued permutation (π,σ), we define $\xi(B) = (t, b^{\rightarrow}_{-1}, b^{\rightarrow}_{+1}, b^{\downarrow}_{-1}, b^{\downarrow}_{+1}, a_0 \cdots a_n)$, where

- $t \in \{\nearrow, \searrow, \bullet\}$ is the type of B,
- $b^{\rightarrow}_{+1} = \sigma(k+n+1, \_)$ if it exists, or $b^{\rightarrow}_{+1} = \perp$ otherwise; $b^{\downarrow}_{+1} = \sigma(\_, l+n+1)$ if it exists, or $b^{\downarrow}_{+1} = \perp$ otherwise,
- $b^{\rightarrow}_{-1} = \sigma(k-1, \_)$ if it exists, or $b^{\rightarrow}_{-1} = \perp$ otherwise; $b^{\downarrow}_{-1} = \sigma(\_, l-1)$ if it exists, or $b^{\downarrow}_{-1} = \perp$ otherwise,
- $a_i = \sigma(k+i, \_)$ for all $0 \leq i \leq n$.

$\xi(B)$ is the **fingerprint** of B, and t is the **type** of $\xi(B)$ (notation: $\mathrm{type}(\tau) = t$, where $\xi(B) = \tau$). For $\xi(B) = \tau$, we also define $\tau_{\rightarrow} = b^{\rightarrow}_{-1} a_0 \cdots a_n b^{\rightarrow}_{+1}$; and $\tau_{\downarrow} = b^{\rightarrow}_{-1} a_0 \cdots a_n b^{\rightarrow}_{+1}$ if $t = \searrow$, or $\tau_{\downarrow} = b^{\rightarrow}_{-1} a_n \cdots a_0 b^{\rightarrow}_{+1}$ otherwise. The set of fingerprints of a valued permutation is the set of the fingerprints of all its maximal blocks. Figure 1-d contains an example of the fingerprints of a valued permutation.

Proposition 2. If φ is satisfiable, then it is satisfiable in a minimal model with at most an exponential number of fingerprints.

By Propositions 1 and 2, we can restrict our attention to permutations with labels, over the exponential alphabet of fingerprints of maximal blocks. However, to do this we need restrict the possible permutations. For example, there cannot be two elements $(r,c),(r+1,c+1)$ in the permutation where both its labels contain fingeprints with type $\searrow$. Indeed, this would imply that the blocks to which these fingerprint correspond were not actually maximal. This suggests that we need to deal with some sort of *constraints* defining valid permutations. This is the theme of the following section.

5. Permutations under constraints

We define constraints that restrict where permutations may or may not contain elements. A constraint specify some positions in which a permutation satisfying it is not allowed to have an element. More precisely, a (n,k) constraint contains not more than k forbidden positions in an n-permutation.

Definition 3 ((n,k)-constraint)**.** Given $n, k \in \mathbb{N}$, and $S, T \subseteq \mathbb{N}$ with $\#S = \#T = n$, we say that $\zeta \subseteq S \times T$ is a (n,k)-constraint over $S \times T$ if for every $(r,c) \in S \times T$ we have $\#\{c' \mid (r,c') \in \zeta\} \leq k$ and $\#\{r' \mid (r',c) \in \zeta\} \leq k$. An n-permutation π over $S \times T$ satisfies a (n,k)-constraint ζ over $S \times T$ if $\pi \cap \zeta = \emptyset$. If $S = T = [n]$ we say that ζ is just a (n,k)-constraint.

Remark 1. As with the permutations, any n-permutation π over S satisfying a (n,k)-constraint ζ can be equivalently seen as a n-permutation π' satisfying a (n,k)-constraint ζ' and vice-versa.

Proposition 3. For every (n,k)-constraint ζ with $n > 2k$, there is an n-permutation π satisfying ζ.

Proof. This can be shown by a simple application of Hall's Marriage Theorem [7] (see also [3, p.36]). Remember that Hall's theorem—in its finite, graph theoretic formulation—states that for any bipartite

graph $G = (V_1 \cup V_2, E)$ with bipartite sets V_1 and V_2 of equal size, G has a perfect matching if, and only if, every subset $S \subseteq V_1$ verifies $|N_G(S)| \geq |S|$. In the formulation, $N_G(S) \subseteq V_2$ is the neighbourhood of S in G (*i.e.*, the set of vertices adjacent to some vertex of S).

Let ζ be a (n, k) constraint where $n > 2k$. Consider a bipartite graph $G = (V_r \cup V_c, E)$, where $V_r = \{r\} \times [n]$, $V_c = \{c\} \times [n]$. Vertices from V_r represent rows and vertices from V_c represent a columns. The set of edges $E \subseteq V_r \times V_c$ is defined as all pairs $((r, i), (c, j))$ so that $(i, j) \notin \zeta$ (*i.e.*, they represent permutation positions that do not interefere any constraint). Hence, there is a perfect matching between V_r and V_c if, and only if, there is an n-permutation satisfying ζ. To prove that there is such a matching, by Halls' theorem it suffices to verify $|N_G(S)| \geq |S|$ for every $S \subseteq V_r$. We show this by case distinction.

- Suppose first $|S| \leq n - k$. Note that every vertex of V_r has at least $n - k$ edges because there are only k constraints in ζ. Then, $N_G(S)$ has $n - k$ vertices because every single vertex in S has already $n - k$ edges.
- Suppose now $|S| > n - k$. Since $n > 2k$, we have $n - k > k$. Then, every vertex from V_c already has one neighbor in S, as every vertex from V_c has at least $n - k > k$ neighbors.

Hence, there is a perfect matching, and thus there exists a permutation satisfying ζ. □

We then have the following corollary.

Corollary 2. For every $(n, 4)$-constraint ζ with $n \geq 9$, there is an n-permutation π that satisfies ζ.

6. Labeled permutations

In this section we prove the NExpTime upper bound of the satisfiability problem for $\mathrm{FO}^2(\rightarrow, \downarrow)$, using the developments of the two previous sections. The idea is to guess the (exponentially many) blocks (of exponential size) of a minimal model that satisfies φ, and use them as letters of our alphabet. Using this guessing, we reduce the satisfiability problem into a problem we introduce next, the *Restricted Labeled Permutation problem* (RLP).

6.1. Restricted labeled permutation problem

Definition 4. A **labeled permutation** over a (finite) alphabet $\mathbb{A}$ is a pair (π, λ) where π is a permutation and $\lambda : \pi \to \mathbb{A}$. Note that (π, λ) is nothing else than a valued permutation where exactly one propositional variable holds at any position.

Definition 5 ($(\pi, \lambda)_{\rightarrow}$, $(\pi, \lambda)_{\downarrow}$)**.** Given a labeled n-permutation (π, λ) we define $(\pi, \lambda)_{\rightarrow}$ and $(\pi, \lambda)_{\downarrow}$ as follows: $(\pi, \lambda)_{\rightarrow} = \lambda(1, c_1) \cdots \lambda(n, c_n)$ and $(\pi, \lambda)_{\downarrow} = \lambda(r_1, 1) \cdots \lambda(r_n, n)$, where $\{(1, c_1), \ldots, (n, c_n)\} = \{(r_1, 1), \ldots, (r_n, n)\} = \pi$.

Definition 6. A **label restriction** over an alphabet $\mathbb{A}$ is a triple (a, t, b), where $a, b \in \mathbb{A}$ and $t \in \{\nearrow, \searrow\}$. We say that a labeled permutation (π, λ) *satisfies* (a, t, b) if for every $(r, c), (s, d) \in \pi$ such that $\lambda(r, c) = a$, $\lambda(s, d) = b$ we have $[(r, c), (s, d)]_\pi \neq t$. We say that (π, λ) satisfies a *set* of label restrictions if it satisfies all of its members.

We define the main problem of this section. Using the result of Section 5 on permutations under constraints, we show that this problem is in NP.

Problem:	The Restricted Labeled Permutation problem (RLP)
Input:	A finite alphabet $\mathbb{A}$, a set of label restrictions R, and two regular languages $\mathcal{L}_1, \mathcal{L}_2 \subseteq \mathbb{A}^*$ (given as NFA).
Question:	Is there a labeled permutation (π, λ) satisfying R such that $(\pi, \lambda)_{\rightarrow} \in \mathcal{L}_1$ and $(\pi, \lambda)_{\downarrow} \in \mathcal{L}_2$?

Proposition 4. The RLP problem is in NP.

Proof. Let $\mathcal{A}_1, \mathcal{A}_2$ be two NFA over the alphabet $\mathbb{A}$ corresponding to the regular languages $\mathcal{L}_1, \mathcal{L}_2$ respectively. Let R be a set of restrictions.

The algorithm first guesses some properties of the labeled permutation (π, λ) that satisfies R and is such that $(\pi, \lambda)_{\rightarrow} \in \mathcal{L}_1$ and $(\pi, \lambda)_{\downarrow} \in \mathcal{L}_2$. (We cannot simply guess (π, λ) because it may be too big.) For each letter $a \in \mathbb{A}$ we guess if it appears exactly k times in (π, λ) for some $k \leq 17$, or if it appears more than 17 times. Let $g : \mathbb{A} \to \{0, 1, \ldots, 17, \infty\}$ be this guessing; and let us define $\mathbb{A}_{\leq} = \{a \in \mathbb{A} \mid g(a) \neq \infty\}$ and $\mathbb{A}_{>} = \{a \in \mathbb{A} \mid g(a) = \infty\}$.

Let us call *zone* to any set $S = [i, i + l] \times [j, j + l]$ for $i, j, l \in \mathbb{N}$. Next, we guess a *small* labeled permutation (π', λ') over the alphabet $\mathbb{A}_{\leq} \cup \{\square\}$. This labeled permutation is such that:

1. All the letters $a \in \mathbb{A}_{\leq}$ in (π', λ') appear exactly $g(a)$ times.
2. There is no zone S so that
   - $\pi' \cap S$ has at least two elements, and
   - $\lambda'(r, c) = \square$ for all $(r, c) \in \pi' \cap S$.
3. (π', λ') satisfies the restrictions R.

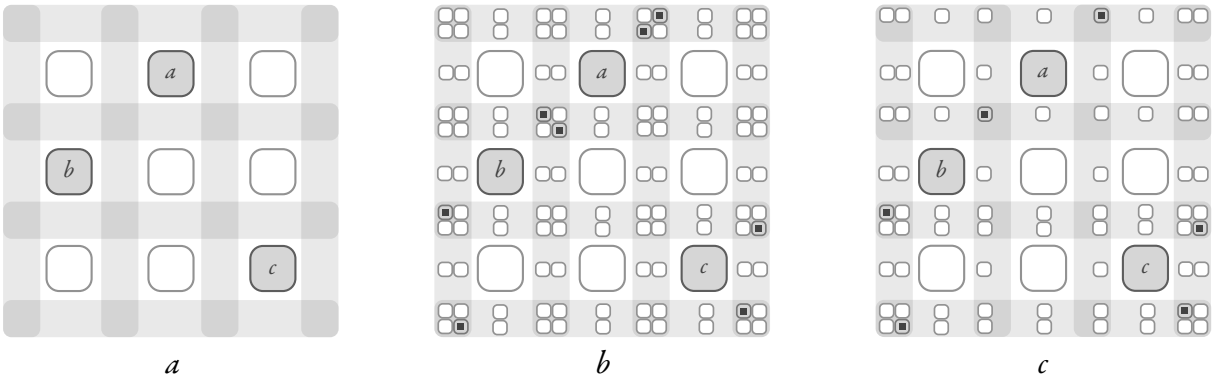

Figure 2: Example of the labeled permutation (π', λ') over the alphabet $\{a, b, c, \square\}$, showing the zones (depicted as) where the label $\square$ (depicted as) can appear. Here, $g(a) = g(b) = g(c) = 1$.

Claim 1. All possible (π', λ') satisfying the conditions above are labeled permutations of size polynomially bounded by $\#\mathbb{A}$.

Proof. Note that, once we fix g, there are not more than $N = (\ 1 + \sum_{\substack{a \in \mathbb{A} \\ g(a) \neq \infty}} g(a)\ )^2$ different zones containing only $\square$ labels, that cover all positions where the label $\square$ can occur in (π', λ'). These are the zones defined in between the elements of $\{a \in \mathbb{A} \mid g(a) \neq \infty\}$. Then, there cannot be more than N elements with label $\square$, since otherwise there would be at least one zone with more than one element, contradicting condition (2). Since $N \leq (1+17 \cdot \#\mathbb{A})^2$, the claim follows. For example, Figure 2-*a* depicts the $16 = (1 + g(a) + g(b) + g(c))^2$ possible zones where the label $\square$ can appear as the dark gray areas. In Figure 2-*b* we see that there is one zone (in fact, two) that contains more than one element $\square$, and therefore condition (2) is falsified (for instance, when $S = [4, 5] \times [4, 5]$). Finally, Figure 2-*c* shows a labeled permutation (π', λ') satisfying condition (2). $\square$

Let $\mathcal{A}_>$ be an NFA over $\mathbb{A}$ that accepts all words $w \in \mathbb{A}^*$ such that every $a \in \mathbb{A}_>$ appears more than 17 times in w. Let e_1 [resp. e_2] be the regular expression resulting from replacing every appearance of $\square$ in $(\pi', \lambda')_\rightarrow$ [resp. in $(\pi', \lambda')_\downarrow$] with the expression $(\mathbb{A}_>)^+$. Notice that, for every $i \in \{1, 2\}$, any word w of $L(e_i) \cap L(\mathcal{A}_>)$ is such that the number of appearances of $a \in \mathbb{A}_\leq$ in w is exactly $g(a)$, and every other letter $a \in \mathbb{A}_>$ appears more than 17 times. Let $\mathcal{A}'_i$ denote the NFA corresponding to $L(\mathcal{A}_i) \cap L(e_i) \cap L(\mathcal{A}_>)$, for every $i \in \{1, 2\}$. Observe that $\mathcal{A}_>$, $\mathcal{A}'_1$ and $\mathcal{A}'_2$ can be built in polynomial time. Given a language $\mathcal{L} \subseteq \mathbb{A}^*$, let $pk(\mathcal{L})$ denote the Parikh image of $\mathcal{L}$. We finally check whether

$$pk(L(\mathcal{A}'_1)) \ \cap \ pk(L(\mathcal{A}'_2)) \ \neq \ \emptyset.$$

This can be verified in NP by computing the existential Presburger formulas for both automata in polynomial time [18] and checking for emptiness of its intersection in NP [16].

Claim 2. $pk(L(\mathcal{A}'_1)) \ \cap \ pk(L(\mathcal{A}'_2)) \ \neq \ \emptyset$ if, and only if, there is a labeled permutation (π, λ) that satisfies R, such that $(\pi, \lambda)_\rightarrow \in \mathcal{L}_1$ and $(\pi, \lambda)_\downarrow \in \mathcal{L}_2$.

The rest of the proof is dedicated to prove the statement above.

[$\Rightarrow$] If $pk(L(\mathcal{A}'_1)) \cap pk(L(\mathcal{A}'_2)) \neq \emptyset$, we show that there is a labeled permutation (π, λ) that satisfies R and such that $(\pi, \lambda)_\rightarrow \in \mathcal{L}_1$ and $(\pi, \lambda)_\downarrow \in \mathcal{L}_2$.

Let $w_1 \in L(\mathcal{A}'_1)$, $w_2 \in L(\mathcal{A}'_2)$ such that $pk(w_1) = pk(w_2)$, and let $m = |w_1| = |w_2|$. For every $a \in \mathbb{A}$, let $X_a \subseteq [m] \times [m]$ be defined as all $(r, c) \in [m] \times [m]$ such that $w_1[r] = a$ or $w_2[c] = a$ (*i.e.*, X_a is the set of possible permutation elements labeled with a). For any $A \subseteq \mathbb{A}$, let $X_A = \bigcup_{a \in A} X_a$.

We define the labeled permutation $(\pi_{\mathbb{A}_\leq}, \lambda_{\mathbb{A}_\leq})$ as all the elements $(r + r', c + c')$ such that

- $(r, c) \in \pi'$, $\lambda'(r, c) \neq \square$, and
- r' [resp. c'] is $k - \ell$, where
  - k is the number of occurrences of letters from $\mathbb{A}_>$ in w_1 [resp. in w_2] before the r-th [resp. c-th] appearance of a letter from $\mathbb{A}_\leq$, and
  - ℓ is the number of letters $\square$ in $(\pi', \lambda')_\rightarrow[1, r]$ [resp. in $(\pi', \lambda')_\downarrow[1, c]$].

  Note that $k \geq \ell$.

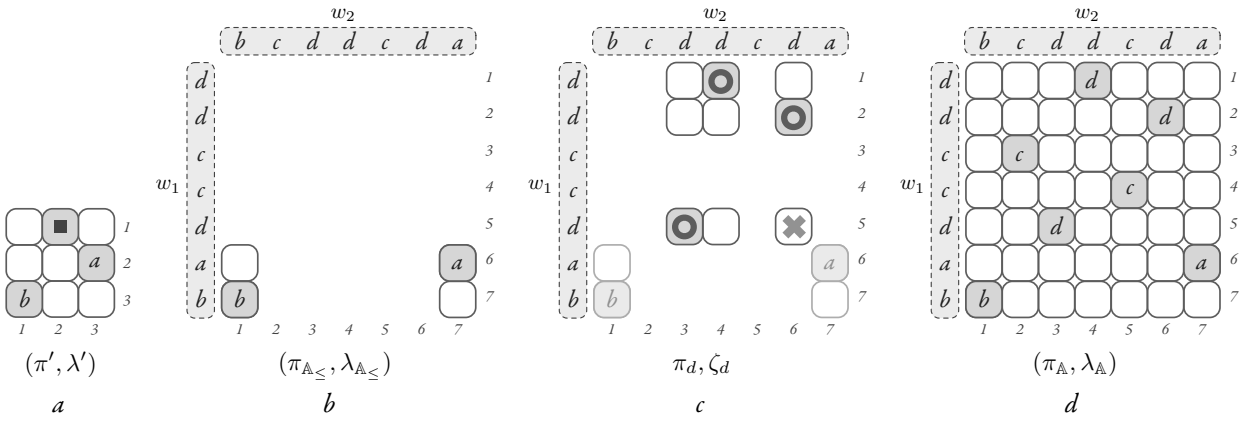

Figure 3: Example where $\mathbb{A}_\leq = \{a, b\}$, $\mathbb{A}_> = \{c, d\}$, $w_1 = ddccdab$, $w_2 = bcddcda$, $m = 7$. For illustration purposes, we let the threshold defining $\mathbb{A}_\leq$ and $\mathbb{A}_>$ to be 1 instead of 17.

We define $\lambda_{\mathbb{A}_\leq}(r+r', c+c') = \lambda'(r, c)$. We have that $(\pi_{\mathbb{A}_\leq}, \lambda_{\mathbb{A}_\leq})$ is a labeled permutation over $X_{\mathbb{A}_\leq}$ satisfying R, as (π', λ') satisfies R by (3). In fact, it is equivalent to (π', λ') when restricted to elements with labels in $\mathbb{A}_\leq$ (*cf.* Figures 3-*a*, 3-*b*). We build, for every $H \subseteq \mathbb{A}_>$, a labeled permutation $(\pi_{\mathbb{A}_\leq \cup H}, \lambda_{\mathbb{A}_\leq \cup H})$ over $X_{\mathbb{A}_\leq \cup H}$ satisfying R, so that $(\pi_{\mathbb{A}_\leq \cup H}, \lambda_{\mathbb{A}_\leq \cup H})_{\rightarrow}$ [resp. $(\pi_{\mathbb{A}_\leq \cup H}, \lambda_{\mathbb{A}_\leq \cup H})_{\downarrow}$] is w_1 [resp. w_2] projected onto $\mathbb{A}_\leq \cup H$. Note that when $H = \mathbb{A}_>$, $(\pi_{\mathbb{A}_\leq \cup H}, \lambda_{\mathbb{A}_\leq \cup H})$ is the labeled permutation over $X_{\mathbb{A}} = [m] \times [m]$ we are looking for. We build these inductively.

The base case is when $H = \emptyset$ and we then have $(\pi_{\mathbb{A}_\leq}, \lambda_{\mathbb{A}_\leq})$, which clearly satisfies R. Suppose now we have constructed $(\pi_{\mathbb{A}_\leq \cup H}, \lambda_{\mathbb{A}_\leq \cup H})$ for some H, and let $a \in \mathbb{A}_>$ be such that $a \notin H$. Let ζ_a be the set of all $(r, c) \in X_a$ such that there is some $(r', c') \in \pi_{\mathbb{A}_\leq \cup H}$ with $|r - r'| = |c - c'| = 1$. Using Corollary 2 one can show that there is always a permutation over X_a satisfying ζ_a, so that it does not have any two elements one next to the other. In the example of Figure 3, we see in item *c* an illustration of a possible such permutation over X_d, and the constraints originating in this case from $(\pi_{\mathbb{A}_\leq}, \lambda_{\mathbb{A}_\leq})$.)

Claim 3. There is a permutation π_a over X_a satisfying the constraints ζ_a such that there are no two $(r, c), (r', c') \in \pi_a$ with $|r - r'| = |c - c'| = 1$.

Proof. We first partition X_a into two sets X'_a, X''_a so that there are no two elements in X'_a (resp. X''_a) with neighboring rows or columns. Let X'_a be the set of all $(r, c) \in X_a$ such that $\#\{r' \mid r' \leq r \wedge (r', c) \in X_a\}$ and $\#\{c' \mid c' \leq c \wedge (r, c') \in X_a\}$ are *odd*; and let X''_a be the set of all $(r, c) \in X_a$ such that $\#\{r' \mid r' \leq r \wedge (r', c) \in X_a\}$ and $\#\{c' \mid c' \leq c \wedge (r, c') \in X_a\}$ are *even*. Since $\#X_a \geq 18^2$, we have that $\#X'_a \geq 9^2$ and $\#X''_a \geq 9^2$. Note that ζ_a does not have more than 4 restrictions on each row and on each column. Then, $\zeta_a \cap X'_a$ is an $(\ell, 4)$-constraint with $\ell \geq 9$. Hence, by Corollary 2, there exists a permutation π'_a over X'_a that satisfies $\zeta_a \cap X'_a$. Now let ζ'_a be the set of all $(r, c) \in X''_a$ such that there is some $(r', c') \in \pi'_a \cup \pi_{\mathbb{A}_\leq \cup H}$ where $|r - r'| = |c - c'| = 1$. Notice that $\zeta_a \cap X''_a \subseteq \zeta'_a$. Remember that $\#X''_a \geq 9^2$ and note that ζ'_a does not have more than 4 restrictions on each row and on each column. Then, applying again Corollary 2, there is a permutation π''_a over X''_a that satisfies $\zeta'_a \cap X''_a$. By definition of X'_a and X''_a we have that $\pi_a = \pi'_a \cup \pi''_a$ is a permutation over X_a that satisfies both $\zeta_a \cap X'_a$ and ζ'_a. Further, since $\zeta_a \cap X''_a \subseteq \zeta'_a$, we have that π_a satisfies ζ_a. Also, by definition of ζ'_a, we further have that there are no two $(r, c), (r', c') \in \pi_a$ with $|r - r'| = |c - c'| = 1$. □

We therefore define the new permutation $\pi_{\mathbb{A}_\leq \cup H \cup \{a\}} = \pi_a \cup \pi_{\mathbb{A}_\leq \cup H}$; and $\lambda_{\mathbb{A}_\leq \cup H \cup \{a\}}(r, c) = a$ for all $(r, c) \in X_a$, and $\lambda_{\mathbb{A}_\leq \cup H \cup \{a\}}(r, c) = \lambda_{\mathbb{A}_\leq \cup H}(r, c)$ otherwise. Note that it satisfies R since for all new positions (r, c) added there is no other position (r', c') such that $|r - r'| = |c - c'| = 1$.

Finally, the desired labeled permutation is $(\pi_{\mathbb{A}_\leq \cup \mathbb{A}_>}, \lambda_{\mathbb{A}_\leq \cup \mathbb{A}_>})$.

[⇐] Suppose that $pk(L(\mathcal{A}'_1)) \cap pk(L(\mathcal{A}'_2)) = \emptyset$; we show that there is no labeled permutation (π, λ) verifying the restrictions imposed by the problem.

First, we show that there cannot be a solution whose every label appears at most 17 times. By

means of contradiction, suppose (π, λ) is such a solution. Then, the algorithm can guess (π, λ), and we would then have that, by construction, $pk(L(\mathcal{A}'_1)) = pk(L(\mathcal{A}'_2)) \neq \emptyset$, which is in contradiction of our hypothesis.

Now suppose that (π, λ) is a solution, where $\mathbb{A}_> \neq \emptyset$ is the set of letters that appear more than 17 times in (π, λ). Let (π'', λ'') be the replacement in (π, λ) of every label from $\mathbb{A}_>$ with $\square$. We can now build (π', λ') from (π'', λ''), by replacing each block containing only elements $\square$ (and is maximal in size with respect to this property) with only one element $\square$. For example, if (π'', λ'') is as depicted in Figure 2-*b*, we produce (π', λ') by removing two rows and two columns, ending up with the labeled permutation of Figure 2-*c*.

Let $e_1, e_2, \mathcal{A}_>, \mathcal{A}'_1, \mathcal{A}'_2$ be defined as before from (π', λ'). It is clear that $(\pi, \lambda)_\rightarrow$ is in $L(e_1)$ and $(\pi, \lambda)_\rightarrow$ in $L(e_2)$. It is therefore true that $(\pi, \lambda)_\rightarrow \in L(\mathcal{A}'_1)$ and $(\pi, \lambda)_\downarrow \in L(\mathcal{A}'_2)$. Since (π, λ) is a labeled permutation we have that $pk((\pi, \lambda)_\rightarrow) = pk((\pi, \lambda)_\downarrow)$. Then, $pk(L(\mathcal{A}'_1)) \cap pk(L(\mathcal{A}'_2)) \neq \emptyset$, which is in contradiction with our hypothesis. Therefore, if $pk(L(\mathcal{A}'_1)) \cap pk(L(\mathcal{A}'_2)) = \emptyset$, there cannot be a solution (π, λ) to the RLP instance. $\square$

6.2. Satisfiability for FO^2

We now show that there is a NExpTime reduction from the satisfiability problem for $\mathrm{FO}^2(\rightarrow, \downarrow)$ into the RLP problem. This, combined with the fact that RLP is in NP (Proposition 4), concludes the proof of Theorem 1, showing that the satisfiability problem for $\mathrm{FO}^2(\rightarrow, \downarrow)$ is in NExpTime; hence it is NExpTime-complete [4].

Before going into the reduction, we show that the satisfaction of a formula in a valued permutation depends solely on its sets of fingerprints, plus some summary information. This summary information says, for every possible valuation S, how many times S appears in (π, σ) (counting up to a threshold of 3). In the reduction from $\mathrm{FO}^2(\rightarrow, \downarrow)$ into RLP we guess the summary information and set of fingerprints of a minimal model (bounded by Propositions 1 and 2), translating the formula into a RLP instance.

Definition 7. Given a set of fingerprints X, let $\hat{X}$ be the set of all the valuations in X, that is $\hat{X} = \{\tau_\rightarrow[i] \mid \tau \in X, 1 \leq i \leq |\tau_\rightarrow| \wedge \tau_\rightarrow[i] \neq \bot\}$. Given a set of fingerprints X and disjoint sets of valuations $V_1, V_2, V_3 \subseteq \hat{X}$, **a valued permutation** (π, σ) **over** X, V_1, V_2, V_3 is any valued permutation such that X is the set of fingerprints of (π, σ), and for every $1 \leq i \leq 3$, V_i is the set of valuations that appear exactly i times in (π, σ). We also say that X, V_1, V_2, V_3 is the **summary** of (π, σ).

Given a formula $\forall y.\psi$ where ψ is a quantifier-free formula, and a valued permutation (π, σ) with a block B, whether all the elements from B verify $\forall y.\psi$ or not, depends only on: the summary of (π, σ), and the fingerprint of B. Similarly for formulas $\exists y.\psi$. Moreover, we can test this in polynomial time.

We introduce the concept of a fingerprint being **consistent** with a formula $\forall x \forall y.\chi$ [resp. $\forall x \exists y.\psi$] and a summary.[5] These are the necessary and sufficient conditions to ensure that every element of a block with such fingerprint in a valued permutation over such summary satisfies $\forall y.\chi$ [resp. $\exists y.\psi$].

First, note that for any quantifier-free formula ψ of $\mathrm{FO}^2(\rightarrow, \downarrow)$, the validity of $(\pi, \sigma) \models_\mu \psi$ only depends on: S, S' and t, where: $S = \sigma(\mu(x))$, $S' = \sigma(\mu(y))$, and $t = [\mu(x), \mu(y)]_\pi$. We will then write $(S)\, t\, (S') \models \psi$, to denote that ψ holds in any model that assigns S to x, S' to y and so that the neighborhood type between x and y is t. Notice that we can decide $(S)\, t\, (S') \models \psi$ in polynomial time. For example, if $\psi = x{\rightarrow}y \wedge (a(x) \vee \neg b(y))$, we have $(\{b\}) \nearrow (\{a, c\}) \models \psi$ but $(\{a\}) \downarrow (\{a, b\}) \not\models \psi$.

The formal definition of consistency is given next.

Definition 8. r Let τ be a fingerprint, and $m = |\tau_\rightarrow|$. Given a set of valuations $Y \subseteq 2^{\mathcal{V}_\varphi}$ and three disjoint sets $V_1, V_2, V_3 \subseteq Y$, we say that τ is **consistent with a universal formula** $\forall y.\chi$ and V_1, V_2, V_3, Y if all of the following conditions hold:

1. For every $1 < i < m$, we have $(\tau_\rightarrow[i]) \bullet (\tau_\rightarrow[i]) \models \chi$.
2. If $\tau_\rightarrow[1] \neq \bot$, $(\tau_\rightarrow[2]) \leftarrow (\tau_\rightarrow[1]) \models \chi$.
3. If $\tau_\rightarrow[m] \neq \bot$, $(\tau_\rightarrow[m-1]) \rightarrow (\tau_\rightarrow[m]) \models \chi$.
4. For every $1 < i, i+1 < m$, if $\mathrm{type}(\tau) = \searrow$, $(\tau_\rightarrow[i]) \searrow (\tau_\rightarrow[i{+}1]) \models \chi$, $(\tau_\rightarrow[i{+}1]) \nwarrow (\tau_\rightarrow[i]) \models \chi$, otherwise, if $\mathrm{type}(\tau) = \nearrow$, $(\tau_\rightarrow[i]) \nearrow (\tau_\rightarrow[i{+}1]) \models \chi$, $(\tau_\rightarrow[i+1]) \swarrow (\tau_\rightarrow[i]) \models \chi$.
5. For every $1 \leq i-1, i, i+1 \leq m$,
   - for every $g \in Y \setminus \{\tau_\rightarrow[i-1], \tau_\rightarrow[i], \tau_\rightarrow[i+1]\}$ then $(\tau_\rightarrow[i]) \infty (g) \models \chi$,
   - for every $g \in \{\tau_\rightarrow[i-1], \tau_\rightarrow[i], \tau_\rightarrow[i+1]\}$, if $g \notin V_j$ for $j = \#\{t \in \{1, 0, -1\} \mid \tau_\rightarrow[i+t] = g\}$ then $(\tau_\rightarrow[i]) \infty (g) \models \chi$.
6. Idem to items 2 and 3, but replacing $\rightarrow$ with $\downarrow$, $\leftarrow$ with $\uparrow$, and $\tau_\rightarrow$ with $\tau_\downarrow$.

[5] In fact, we do not need the set of fingerprints X to define this notion but just the set of valuations $\hat{X}$, and it is therefore defined over $\hat{X}$.

We say that τ is **consistent with an existential formula** $\exists y.\psi$ if for every $1 < i < m$ either

1. (a) $(\tau_{\rightarrow}[i]) \bullet (\tau_{\rightarrow}[i]) \models \psi$,
   (b) $i+1 = m$ and $\tau_{\rightarrow}[i+1] \neq \bot$ and $(\tau_{\rightarrow}[i]) \rightarrow (\tau_{\rightarrow}[i+1]) \models \psi$, or
   (c) $i = 2$, $\tau_{\rightarrow}[1] \neq \bot$ and $(\tau_{\rightarrow}[2]) \leftarrow (\tau_{\rightarrow}[1]) \models \psi$,
2. $\text{type}(\tau) = \searrow$, and either
   (a) $1 < i, i+1 < m$ and $(\tau_{\rightarrow}[i]) \searrow (\tau_{\rightarrow}[i+1]) \models \psi$,
   (b) $1 < i-1, i < m$ and $(\tau_{\rightarrow}[i]) \nwarrow (\tau_{\rightarrow}[i-1]) \models \psi$,
   (c) $1 \leq i-1, i, i+1 \leq m$ and there is some $g \in Y \setminus \{\tau_{\rightarrow}[i-1], \tau_{\rightarrow}[i], \tau_{\rightarrow}[i+1]\}$, such that $(\tau_{\rightarrow}[i]) \infty (g) \models \psi$, or
   (d) $1 \leq i-1, i, i+1 \leq m$ and there is $g \in \{\tau_{\rightarrow}[i-1], \tau_{\rightarrow}[i], \tau_{\rightarrow}[i+1]\}$ with $g \notin V_j$ for $j = \#\{t \in \{1, 0, -1\} \mid \tau_{\rightarrow}[i+t] = g\}$ such that $(\tau_{\rightarrow}[i]) \infty (g) \models \psi$,
3. $\text{type}(\tau) = \nearrow$, and some condition as the ones in item 2 holds, where $\searrow$ and $\nwarrow$ are replaced with $\nearrow$ and $\swarrow$, or
4. Idem as condition 1, replacing $\rightarrow$ with $\downarrow$, $\leftarrow$ with $\uparrow$, and $\tau_{\rightarrow}$ with $\tau_{\downarrow}$.

Finally, we say that a fingerprint τ is consistent with a formula in Scott normal form $\varphi = \forall x \forall y.\chi \wedge \bigwedge_i \forall x \exists y.\psi_i$ and sets Y, V_1, V_2, V_3 if it is consistent with $\forall y.\chi$ and with $\exists y.\psi_i$ for all i.

The following Lemmas follow straightforward from the previous definitions.

Lemma 1. For every valued permutation (π, σ) over X, V_1, V_2, V_3, with a maximal block B, and for every formula $\forall y.\chi$ where χ is quantifier-free, we have that all the elements from B verify $\forall y.\chi$ if and only if $\xi(B)$ is consistent with $\forall y.\chi$ and $\hat{X}, V_1, V_2, V_3$.

Lemma 2. For every valued permutation (π, σ) over X, V_1, V_2, V_3, with a maximal block B, for every formula $\exists y.\psi$ where ψ is quantifier-free, we have that all the elements from B verify $\exists y.\psi$ if and only if $\xi(B)$ is consistent with $\exists y.\psi$ and $\hat{X}, V_1, V_2, V_3$.

Remark 2. Note that the property of consistency of Definition 8 can be checked in polynomial time.

Lemma 3. There is a NExpTime reduction from the satisfiability problem for $\text{FO}^2(\rightarrow, \downarrow)$ into the RLP problem.

Proof. Let $\varphi \in \text{FO}^2(\rightarrow, \downarrow)$ be in Scott normal form, $\varphi = \forall x \forall y.\chi \wedge \bigwedge_i \forall x \exists y.\psi_i$. By Proposition 2, there are at most an exponential number of different fingerprints, and by Proposition 1 each one of them is at most of exponential size. The algorithm guesses the set X of all fingerprints needed in a minimal valued permutation that satisfies φ, and summary sets $V_1, V_2, V_3 \subseteq \hat{X}$. The algorithm checks that for every $\tau \in X$, τ is consistent with φ and $\hat{X}, V_1, V_2, V_3$.

We define the language $\mathcal{L}_1 \subseteq X^*$ [resp. $\mathcal{L}_2 \subseteq X^*$] of all words $w \in X^*$ such that

- the first element of $(w[1])_{\rightarrow}$ [resp. of $(w[1])_{\downarrow}$] is $\bot$, and the last element of $(w[|w|])_{\rightarrow}$ [resp. of $(w[|w|])_{\downarrow}$] is $\bot$,
- for every $1 \leq i < |w|$ such that $(w[i])_{\rightarrow} = a_1 \cdots a_{n-1} a_n$ [resp. $(w[i])_{\downarrow} = a_1 \cdots a_{n-1} a_n$] and $(w[i+1])_{\rightarrow} = b_1 b_2 \cdots b_m$ [resp. $(w[i+1])_{\downarrow} = b_1 b_2 \cdots b_m$] we have $a_{n-1} a_n = b_1 b_2$,
- all the elements of X appear in w,
- for every $1 \leq i \leq 3$, V_i is the set of valuations that appear exactly i times in $\hat{w}$, where $\hat{w} = \hat{w}_1 \cdots \hat{w}_{|w|}$ and $\hat{w}_i = (w[i])_{\rightarrow}[2] \cdots (w[i])_{\rightarrow}[m_i - 1]$ for $m_i = |(w[i])_{\rightarrow}|$.

It is immediate that $\mathcal{L}_1$ and $\mathcal{L}_2$ are regular languages, and that they can be defined by two NFA that can be built in polynomial time in the size of X.

Finally, we define the label restrictions, avoiding having two blocks that actually define a bigger block (because these blocks are supposed to be maximal). Let R be the set of all triples (τ, d, τ') such that $\tau, \tau' \in X$ and either

- $d = \searrow$, $\text{type}(\tau) \neq \nearrow$, $\text{type}(\tau') \neq \nearrow$, or
- $d = \nearrow$, $\text{type}(\tau) \neq \searrow$, $\text{type}(\tau') \neq \searrow$.

We reduced the satisfiability problem into the RLP problem for $(X, R, \mathcal{L}_1, \mathcal{L}_2)$, concluding the proof.

Claim 4. The RLP instance $(X, R, \mathcal{L}_1, \mathcal{L}_2)$ has a positive solution iff the formula φ is satisfiable. □

7. Conclusion

Our work shows that the following combinatorial problem is at the core of the satisfiability for $\text{FO}^2(\rightarrow, \downarrow)$ and of the RLP problem, and is decidable in NP. Given two regular languages $\mathcal{L}, \mathcal{L}' \subseteq \mathbb{A}^*$, is there a word $a_1 \cdots a_n \in \mathcal{L}$ and a permutation $p : [n] \rightarrow [n]$ so that $a_{p(1)} \cdots a_{p(n)} \in \mathcal{L}'$ and $|p(i+1) - p(i)| > 1$ for all $i \in [n]$?

A natural question left open is whether this decidability result can be extended to FO^2 with k successor relations over finite linear orders is decidable, for arbitrary k (or at least for $k = 3$).

References

[1] Mikołaj Bojańczyk, Claire David, Anca Muscholl, Thomas Schwentick, and Luc Segoufin. Two-variable logic on data words. *ACM Trans. Comput. Log.*, 2010.

[2] Mikołaj Bojańczyk, Anca Muscholl, Thomas Schwentick, and Luc Segoufin. Two-variable logic on data trees and XML reasoning. *Journal of the ACM*, 56(3):1–48, 2009.
[3] Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Heidelberg, third edition, 2005.
[4] 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.
[5] Erich Grädel, Phokion G. Kolaitis, and Moshe Y. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.
[6] Erich Grädel and Martin Otto. On logics with two variables. *Theoretical Computer Science*, 224(1-2):73–113, 1999.
[7] Philip Hall. On representatives of subsets. *Journal of the London Mathematical Society*, 10:26–30, 1935.
[8] Emanuel Kieroński. Results on the guarded fragment with equivalence or transitive relations. In *CSL*, volume 3634 of *Lecture Notes in Computer Science*, pages 309–324. Springer, 2005.
[9] Emanuel Kieroński. Decidability issues for two-variable logics with several linear orders. In *EACSL Annual Conference on Computer Science Logic (CSL'11)*, volume 12 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 337–351. Schloss Dagstuhl, 2011.
[10] Emanuel Kieroński and Martin Otto. Small substructures and decidability issues for first-order logic with two variables. In *LICS*, pages 448–457. IEEE Computer Society, 2005.
[11] Emanuel Kieroński and Lidia Tendera. On finite satisfiability of two-variable first-order logic with equivalence relations. In *LICS*, pages 123–132. IEEE Computer Society, 2009.
[12] Amaldev Manuel. Two orders and two variables. In *Int. Symp. on Mathematical Foundations of Comp. Sci. (MFCS'10)*, Lecture Notes in Computer Science. Springer, 2010.
[13] Amaldev Manuel. Personal communication, 2012.
[14] Amaldev Manuel and Thomas Zeume. Two-variable logic with a linear successor and a preorder. Unpublished manuscript available at `http://www.imsc.res.in/~amal/manuelzeume.pdf`, 2011.
[15] Michael Mortimer. On languages with two variables. *Mathematical Logic Quarterly*, 21:135–140, 1975.
[16] Christos H. Papadimitriou. On the complexity of integer programming. *Journal of the ACM*, 28(4):765–768, 1981.
[17] Thomas Schwentick and Thomas Zeume. Two-variable logic with two order relations. *Logical Methods in Computer Science*, 8(1:15):1–27, 2012.
[18] Kumar Neeraj Verma, Helmut Seidl, and Thomas Schwentick. On the complexity of equational horn clauses. In *International Conference on Automated Deduction (CADE'05)*, volume 3632 of *Lecture Notes in Computer Science*, pages 337–352. Springer, 2005.

Appendix A. Missing proofs

Proof of Lemma 1.

[⇒] Suppose that every element $(r,c) \in B \cap \pi$ verifies $(\pi,\sigma) \models_{[x\mapsto(r,c)]} \forall y.\chi$. We show that $\tau = \xi(B)$ is consistent with $\forall y.\chi$ and $\hat{X}, V_1, V_2, V_3$.

Condition 1 is met, since in particular $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r,c)]} \chi$ for every $(r,c) \in B \cap \pi$; this means that $(\sigma(r,c)) \bullet (\sigma(r,c)) \models \chi$, which implies condition 1.

If i is the smallest column element of B and $i > 1$, it means that $\tau_{\rightarrow}[1] \neq \bot$, in fact $\tau_{\rightarrow}[1] = \sigma(\_, i-1)$. Since for some $(r,i),(r',i-1) \in \pi$, $(\pi,\sigma) \models_{[x\mapsto(r,i),y\mapsto(r',i-1)]} \chi$, and since B is maximal, we have that $|r - r'| > 1$ and hence that $(\sigma(r,i)) \leftarrow (\sigma(r',i-1)) \models \chi$. As $\sigma(r',i-1) = \tau_{\rightarrow}[1]$ and $\sigma(r,i) = \tau_{\rightarrow}[2]$, condition 2 is met. Condition 3 is similar.

Suppose $\text{type}(\tau) = \searrow$ and $1 < i, i+1 < |\tau_{\rightarrow}|$. This means that there are $(r,c),(r+1,c+1) \in B \cap \pi$ with $\sigma(r,c) = \tau_{\rightarrow}[i]$, $\sigma(r+1,c+1) = \tau_{\rightarrow}[i+1]$. Since $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r+1,c+1)]} \chi$ and $(\pi,\sigma) \models_{[y\mapsto(r,c),x\mapsto(r+1,c+1)]} \chi$ we have that $(\sigma(r,c)) \searrow (\sigma(r+1,c+1)) \models \chi$ and $(\sigma(r+1,c+1)) \nwarrow (\sigma(r,c)) \models \chi$. Hence, $(\tau_{\rightarrow}[i]) \searrow (\tau_{\rightarrow}[i+1]) \models \chi$ and $(\tau_{\rightarrow}[i+1]) \nwarrow (\tau_{\rightarrow}[i]) \models \chi$ and thus condition 4 holds. The proof for $\text{type}(\tau) = \nwarrow$ is similar.

Suppose now that $1 \leq i-1, i, i+1 \leq |\tau_{\rightarrow}|$. Let $g \in \hat{X} \setminus \{\tau_{\rightarrow}[i-1], \tau_{\rightarrow}[i], \tau_{\rightarrow}[i+1]\}$. This means that there is some $(r,c) \in B \cap \pi$ with $\sigma(r,c) = \tau_{\rightarrow}[i]$ and some $(r',c') \in \pi$ with $\sigma(r',c') = g$, $[(r,c),(r',c')]_\pi = \infty$. Since $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r',c')]} \chi$, we have that $(\sigma(r,c)) \infty (\sigma(r',c')) \models \chi$ and hence the first part of condition 5 is met. On the other hand, if $g \in \{\tau_{\rightarrow}[i-1], \tau_{\rightarrow}[i], \tau_{\rightarrow}[i+1]\}$ and $g \notin V_j$ for $j = \#\{t \in \{1,0,-1\} \mid \tau_{\rightarrow}[i+t] = g\}$, this means that there must be necessarily some other position with valuation g. That is, there is some $(r',c') \in \pi$ with $\sigma(r',c') = g$, $[(r,c),(r',c')]_\pi = \infty$. Then, the same reasoning as before applies to show that the second part of condition 5 is met.

[⇐] Suppose that $\tau = \xi(B)$ is consistent with $\forall y.\chi$ and $\hat{X}, V_1, V_2, V_3$. Let us show that every element $(r,c) \in B \cap \pi$ verifies $(\pi,\sigma) \models_{[x\mapsto(r,c)]} \forall y.\chi$.

Let $(r,c) \in B \cap \pi$ and $(r',c') \in \pi$. If $[(r,c),(r',c')]_\pi = \bullet$ then $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r',c')]} \chi$ by condition 1. If $[(r,c),(r',c')]_\pi = \leftarrow$, then $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r',c')]} \chi$ by condition 2. If $[(r,c),(r',c')]_\pi = \rightarrow$, then $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r',c')]} \chi$ by condition 3. If $[(r,c),(r',c')]_\pi \in \{\nearrow, \nwarrow, \searrow, \swarrow\}$, then $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r',c')]} \chi$ by condition 4. And if $[(r,c),(r',c')]_\pi = \infty$, then $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r',c')]} \chi$ by condition 5. Hence, every element $(r,c) \in B \cap \pi$ verifies $(\pi,\sigma) \models_{[x\mapsto(r,c)]} \forall y.\chi$. □

Proof of Lemma 2. [⇒] Suppose that every element $(r,c) \in B \cap \pi$ verifies $(\pi,\sigma) \models_{[x\mapsto(r,c)]} \exists y.\psi$. We

show that $\tau = \xi(B)$ is consistent with $\exists y.\psi$ and $\hat{X}, V_1, V_2, V_3$.

Let $(r,c) \in B \cap \pi$, then there must be $(r',c') \in \pi$ such that $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r',c')]} \psi$. If $[(c,r),(r',c')]_\pi = \bullet$ then condition 1a is met. If $[(c,r),(r',c')]_\pi = \rightarrow$, then condition 1b is met. If $[(c,r),(r',c')]_\pi = \leftarrow$, then condition 1c is met. If $[(c,r),(r',c')]_\pi = \searrow$, then condition 2a is met. If $[(c,r),(r',c')]_\pi = \nwarrow$, then condition 2b is met. If $[(c,r),(r',c')]_\pi = \infty$, then either condition 2c or 2d is met. If $[(c,r),(r',c')]_\pi \in \{\nearrow, \swarrow\}$, then condition 3 is met. If $[(c,r),(r',c')]_\pi \in \{\uparrow, \downarrow\}$, then condition 4 is met. Thus, $\tau = \xi(B)$ is consistent with $\exists y.\psi$ and $\hat{X}, V_1, V_2, V_3$.

[$\Leftarrow$] Suppose that $\tau = \xi(B)$ is consistent with $\exists y.\psi$ and $\hat{X}, V_1, V_2, V_3$. Let us show that every element $(r,c) \in B \cap \pi$ verifies $(\pi,\sigma) \models_{[x\mapsto(r,c)]} \exists y.\psi$.

Let $(r,c) \in B \cap \pi$. There must be some $1 < i < |\tau_\rightarrow|$ such that $\tau_\rightarrow[i] = \sigma(r,c)$. If condition 1a holds, then $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r,c)]} \psi$. If condition 1b holds, then $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r',c+1)]} \psi$ for some $(r',c+1) \in \pi$ with $|r'-r| > 1$. If condition 1c holds, then $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r',c-1)]} \psi$ for some $(r',c-1) \in \pi$ with $|r'-r| > 1$. If condition 2a holds, then $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r+1,c+1)]} \psi$ for $(r+1,c+1) \in \pi$. If condition 2b holds, then $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r-1,c-1)]} \psi$ for $(r-1,c-1) \in \pi$. If condition 2c or 2d holds, then $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r-1,c-1)]} \psi$ for some $(r',c') \in \pi$ with $[(r,c),(r',c')]_\pi = \infty$. A similar reasoning applies if 3 or 4 hold. □

Proof of Proposition 1. We show the following statement: Any minimal valued permutation satisfying $\varphi \in \mathrm{FO}^2(\rightarrow,\downarrow)$ is such that every block is of size less or equal to

$$3 \cdot 2^{4\#\mathcal{V}_\varphi+3} + 2^{3(\#\mathcal{V}_\varphi+1)} + 3 \cdot 2^{\#\mathcal{V}_\varphi}.$$

Let (π,σ) be a minimal valued permutation such that $(\pi,\sigma) \models \varphi$. Without loss of generality assume $\sigma : \pi \to 2^{\mathcal{V}_\varphi}$.

Suppose, by means of contradiction, that $B = [i,i+k] \times [i',i'+k]$ is a block of π of size $k > 3 \cdot 2^{4\#\mathcal{V}_\varphi+3} + 2^{3(\#\mathcal{V}_\varphi+1)} + 3 \cdot 2^{\#\mathcal{V}_\varphi}$. Let us assume that B is of type $\searrow$ (if it has type $\nearrow$ a symmetrical reasoning applies). For every $S \subseteq \mathcal{V}_\varphi$, choose any three elements from $\{(r,c) \in \pi \cap B \mid \sigma(r,c) = S\}$, or, if there are less than three, all the elements; let S_π be the set of these three (or less) elements. Now consider all the sub-blocks $B' \subsetneq B$ defined strictly between elements of $\bigcup_{S\subseteq\mathcal{V}_\varphi} S_\pi$, that is, such that $B' \cap \bigcup_{S\subseteq\mathcal{V}_\varphi} S_\pi = \emptyset$, consider the ones that are maximal with respect to inclusion. Since $\#\bigcup_{S\subseteq\mathcal{V}_\varphi} S_\pi \leq 3 \cdot 2^{\#\mathcal{V}_\varphi}$, there are at most $3 \cdot 2^{\#\mathcal{V}_\varphi} + 1$ such sub-blocks having a total of at least $k - 3 \cdot 2^{\#\mathcal{V}_\varphi}$ elements. By the Pigeonhole Principle, since $k - 3 \cdot 2^{\#\mathcal{V}_\varphi} > 3 \cdot 2^{4\#\mathcal{V}_\varphi+3} + 2^{3(\#\mathcal{V}_\varphi+1)} = (3 \cdot 2^{\#\mathcal{V}_\varphi} + 1) \cdot 2^{3(\#\mathcal{V}_\varphi+1)}$, this means that there must be a sub-block with at least $2^{3(\#\mathcal{V}_\varphi+1)} + 1$ elements; suppose it is

$$B' = ([i+\ell, i+\ell+k'] \times [i'+\ell, i'+\ell+k']) \subseteq B,$$

where $k' > 2^{3(\#\mathcal{V}_\varphi+1)}$ and $\ell + k' \leq k$. There must be two distinct elements $(r,c),(r',c') \in \pi \cap B'$ with

$$\begin{aligned} \sigma(r,c) &= \sigma(r',c'), \\ \sigma(r-1,\_) &= \sigma(r'-1,\_), \text{ and} \\ \sigma(r+1,\_) &= \sigma(r'+1,\_). \end{aligned} \qquad (\dagger)$$

Without any loss of generality, suppose that $r < r'$ and $c < c'$ (the same argument works if $r < r'$ and $c' < c$). Let $\pi' = \pi \setminus ([r+1,r'] \times [c+1,c'])$. It is easy to verify that $(\pi',\sigma') \models \varphi$, where $\sigma' = \sigma|_{\pi'}$.

Claim 5. $(\pi',\sigma') \models \varphi$

Proof. Let $\varphi = \forall x \forall y.\chi \wedge \bigwedge_j \forall x \exists y.\psi_j$. For any $(s,d),(s',d') \in \pi$ such that $(s,d) \in \pi'$ and $(\pi,\sigma) \models_{[x\mapsto(s,d),y\mapsto(s',d')]} \psi_j$, we can find $(s'',d'') \in \pi'$ such that $\sigma(s',d') = \sigma'(s'',d'')$ and such that $[(s'',d''),(s,d)]_{\pi'} = [(s',d'),(s,d)]_\pi$. This is because, if the neighborhood type is not ∞ then it is taken care of by (†). Otherwise, it is witnessed by one of the elements from S_π, for $\sigma(s',d') = S$. Therefore, (π',σ') verifies $\forall x \exists y.\psi_j$ for every j.

On the other hand, if $(s,d),(s',d') \in \pi'$ where $(\pi,\sigma) \models_{[x\mapsto(s,d),y\mapsto(s',d')]} \chi$, we have that, if $[(s,d),(s',d')]_{\pi'} = [(s,d),(s',d')]_\pi$, then $(\pi',\sigma') \models_{[x\mapsto(s,d),y\mapsto(s',d')]} \chi$. Otherwise, we necessarily have that $(s,d) = (r,c)$, $(s',d') = (r'+1,c'+1)$ (or vice-versa). But since $\sigma(r+1,c+1) = \sigma(r'+1,c'+1)$ and $(\pi,\sigma) \models_{[x\mapsto(r,c),y\mapsto(r+1,c+1)]} \chi$, then $(\pi',\sigma') \models_{[x\mapsto(s,d),y\mapsto(s',d')]} \chi$, because $[(s,d),(s',d')]_{\pi'} = [(r,c),(r+1,c+1)]_\pi = \searrow$. □

Since π' is smaller than π, and (π',σ') verifies φ, it cannot be that π is minimal in size, which is in contradiction with our hypothesis. Therefore, π does not have blocks of size bigger than $3 \cdot 2^{4\#\mathcal{V}_\varphi+3} + 2^{3(\#\mathcal{V}_\varphi+1)} + 3 \cdot 2^{\#\mathcal{V}_\varphi}$. □

Proof of Proposition 2. Let $(\pi,\sigma) \models \varphi$ be such that (π,σ) is minimal in size. Without any loss of generality we assume that $\sigma : \pi \to 2^{\mathcal{V}_\varphi}$. We show

how to build another minimal valued permutation from (π, σ) with a number block fingerprints that is bounded by an exponential function on $|\varphi|$.

Let $\pi = B_1 \cup \cdots \cup B_N$, where each B_i is a maximal block. We are going to mark some blocks (only exponentially many), and use only these to build a new valued permutation satisfying φ. For each $S \subseteq \mathcal{V}_\varphi$, let $S_\pi \subseteq \pi$ be a set of 4 elements of $\{(r,c) \in \pi \mid \sigma(r,c) = S\}$ —if the set has less than 4 elements, then let S_π be all of them. For every $S \subseteq \mathcal{V}_\varphi$ and every $(r,c) \in S_\pi$, mark the block B_k such that $(r,c) \in B_k$ with a color red. It follows that we end up with at most $4 \cdot 2^{\#\mathcal{V}_\varphi}$ blocks marked with red. On the other hand, for every $i \in [N]$, let us define $h_i = (d, b^{\rightarrow}_{-1}, b^{\rightarrow}_{+1}, b^{\downarrow}_{-1}, b^{\downarrow}_{+1})$, where $\xi(B_i) = (d, b^{\rightarrow}_{-1}, b^{\rightarrow}_{+1}, b^{\downarrow}_{-1}, b^{\downarrow}_{+1}, a_0 \cdots a_t)$. For each $h \in \{h_i \mid i \in [N] \text{ s.t. } B_i \text{ is not marked red}\}$ let us mark with color green one block B_i such that $h_i = h$ and such that B_i is not marked red. Since there are no more than $3 \cdot 2^{4\#\mathcal{V}_\varphi}$ different h_i's, there are only an exponential number of blocks marked with colors red or green.

Now, let B_m be a block with a fingerprint different from all the fingerprints of the marked blocks. There must be a block B_n marked with green, so that $h_n = h_m$. Let (π', σ') be the result of replacing of B_m with B_n in (π, σ), in the expected way. We then have that $(\pi', \sigma') \models \varphi$.

Claim 6. $(\pi', \sigma') \models \varphi$.

Proof. Let $\varphi = \forall x \forall y.\chi \wedge \bigwedge_i \forall x \exists y.\psi_i$.

We first show that $\forall x \forall y.\chi$ holds in (π', σ'). Let $(r,c), (r',c') \in \pi'$.

- If (r,c) is in the edited block of the permutation (*i.e.*, in the new copy of the fingerprint of B_n).
  - If $[(r,c),(r',c')]_\pi \neq \infty$, then both elements are described inside the fingerprint $\xi(B_n)$. Hence, there are two elements in the block B_n of (π, σ) with the same neighborhood type and with the same valuations. Since for those two we have that χ holds, then it holds for $(r,c), (r',c')$ as well. Hence, $(\pi', \sigma') \models_{[x \mapsto (r,c), y \mapsto (r',c')]} \chi$.
  - Otherwise, suppose $[(r,c),(r',c')]_\pi = \infty$. Note that in the red colored blocks of (π, σ) there must be at least 4 elements with value $\sigma'(r,c)$. By the same reason there must be 4 elements with value $\sigma'(r',c')$. Since marked blocks are preserved, there must be $(s,l), (s',d') \in \pi$ such that $[(s,d),(s',d')]_\pi = \infty$, and $\sigma(s,d) = \sigma'(r,c)$, $\sigma(s',d') = \sigma'(r',c')$. Since $(\pi, \sigma) \models_{[x \mapsto (s,d), y \mapsto (s',d')]} \chi$, we have $(\pi', \sigma') \models_{[x \mapsto (r,c), y \mapsto (r',c')]} \chi$.
- If both (r,c) and (r',c') are in the part that was not modified, then there must clearly be $(s,d), (s',d') \in \pi$ with $[(s,d),(s',d')]_\pi = [(r,c),(r',c')]_{\pi'}$, so that $\sigma(s,d) = \sigma'(r,c)$, $\sigma(s',d') = \sigma'(r',c')$. Hence, $(\pi', \sigma') \models_{[x \mapsto (r,c), y \mapsto (r',c')]} \chi$.

We now show that for every i and $(r,c) \in \pi'$, $\exists y.\varphi_i$ holds.

- If (r,c) is in the edited block of the permutation (*i.e.*, in the new copy of the fingerprint of B_n), then there must be some (r',c') in B_n such that $\sigma(r',c') = \sigma'(r,c)$. Since $(\pi, \sigma) \models_{[x \mapsto (r',c')]} \exists y.\varphi_i$, there must be some $(s',d') \in \pi$ such that $(\pi, \sigma) \models_{[x \mapsto (r',c'), y \mapsto (s',d')]} \varphi_i$.
  - If $[(s',d'),(r',c')]_\pi \neq \infty$, then it is described in the fingerprint of B_n and therefore there must be some (s,d) such that $[(s,d),(r,c)]_{\pi'} = [(s',d'),(r',c')]_\pi$, and $\sigma'(s,d) = \sigma(s',d')$, $\sigma'(r,c) = \sigma(r',c')$. Then, $(\pi', \sigma') \models_{[x \mapsto (r,c), y \mapsto (s,d)]} \varphi_i$ since $(\pi, \sigma) \models_{[x \mapsto (r',c'), y \mapsto (s',d')]} \varphi_i$.
  - Otherwise, suppose that $[(s',d'),(r',c')]_\pi = \infty$. There must be 4 elements of (π, σ) with the same value $\sigma(s',d')$ inside blocks marked red. Then, there must be a block marked red in π' containing some (s,d) with $\sigma'(s,d) = \sigma(s',d')$ and such that $[(s,d),(r,c)]_{\pi'} = \infty$. Then, $(\pi', \sigma') \models_{[x \mapsto (r,c), y \mapsto (s,d)]} \varphi_i$ since $(\pi, \sigma) \models_{[x \mapsto (r',c'), y \mapsto (s',d')]} \varphi_i$.
- If (r,c) is in the part that was not modified, then there is some $(r',c') \in \pi$ with $\sigma(r',c') = \sigma'(r,c)$ and $(r',c') \notin B_m$. Since $(\pi, \sigma) \models_{[x \mapsto (r',c')]} \exists y.\varphi_i$, there must be some $(s',d') \in \pi$ such that $(\pi, \sigma) \models_{[x \mapsto (r',c'), y \mapsto (s',d')]} \varphi_i$.
  - If $[(s',d'),(r',c')]_\pi \neq \infty$, then (s',d') is described in the fingerprint of the block where (r',c') belongs. Since this fingerprint is preserved (because it is not B_m), there must be some (s,d) with $\sigma'(s,d) = \sigma(s',d')$ and such that $[(r,c),(s,d)]_{\pi'} = [(r',c'),(s',d')]_\pi$. Then, $(\pi, \sigma) \models_{[x \mapsto (r',c'), y \mapsto (s',d')]} \varphi_i$.
  - Otherwise, if $[(s',d'),(r',c')]_\pi = \infty$, there must be 4 other positions with the same value as (s',d') in (π, σ), in blocks marked red. Then, there must necessarily be some $(s,d) \in \pi'$ such that $\sigma'(s,d) = \sigma(s',d')$

and such that $[(s,d),(r,c)]_{\pi'} = \infty$. Then, $(\pi,\sigma) \models_{[x\mapsto(r',c'),y\mapsto(s',d')]} \varphi_i$.

Therefore, $(\pi',\sigma') \models_{[x\mapsto(r,c)]} \exists y.\varphi_i$. □

Note that since B_m and B_n have the same type, then the number of maximal blocks in π' is the same as in π, and moreover the number of maximal blocks with fingerprints different from the marked fingerprints is decremented by one. Note that in particular this means that $\#B_m = \#B_n$. Otherwise, if $\#B_m > \#B_n$, we would have that (π',σ') is smaller than (π,σ) which cannot be since (π,σ) is minimal in size. Conversely, if $\#B_m < \#B_n$, we could have marked B_m with green instead of B_n and we arrive to the same contradiction.

We can repeat this operation with all the unmarked blocks ending up with a valued permutation whose every block has the fingerprint of a marked block. Since there are only exponentially many marked blocks, there are exponentially many fingerprints. □