On the expressive power of invariant logics over sparse classes of structures

Julien Grange

THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PSL

Préparée à l'Ecole Normale Supérieure

On the Expressive Power of Invariant Logics over Sparse Classes of Structures

Soutenue par
Julien GRANGE
Le 29 juin 2020

Ecole doctorale n° 386
Sciences Mathématiques de Paris-Centre

Spécialité
Informatique

Composition du jury :

Cristina SIRANGELO Professeur, Université Paris Diderot	*Présidente*
Anuj DAWAR Professeur, University of Cambridge	*Rapporteur*
Etienne GRANDJEAN Professeur émérite, Université de Caen	*Rapporteur*
Mamadou KANTE Maître de conférences, Université Clermont Auvergne	*Examinateur*
Leonid LIBKIN Professeur, Ecole Normale Supérieure	*Examinateur*
Luc SEGOUFIN Directeur de recherches, Ecole Normale Supérieure	*Directeur de thèse*

Je voudrais remercier sincèrement mon directeur de thèse, Luc Segoufin, pour son encadrement particulièrement formateur et sa disponibilité sans faille au cours de ces trois années.

J'adresse un grand merci, sans me lancer dans un inventaire qui risquerait à tort d'être perçu comme non order-invariant, à tous mes proches.

Cette thèse est dédiée à la mémoire de Jacquot l'Anglais.

Contents

Chapter 1

Introduction

Finite model theory offers to develop formalisms for expressing properties of finite structures, and focuses on proving relations between, and bounds on, these languages.

The most basic of these formalisms is probably *first-order logic*, FO, in which one can quantify over elements of the structures. Its expressive power has been studied extensively and is known to be very restrained, as it can only express properties that are local, which roughly means that it can only talk about the immediate surroundings of a small number of elements, and it is unable to count.

There are several ways to define logics extending FO in order to enhance its expressive power. For instance, *monadic second-order logic* MSO takes FO as a building block, and allows quantification not only over elements, but also over sets of elements. Alternatively, one could add a fixpoint operator to FO and get *least fixpoint first-order logic* LFP. Both these additions break the local character of the logic.

Another way to define logics from FO is through the addition, in an invariant way, of arithmetic predicates on the structure that are exterior to the vocabulary. This amounts to arbitrarily identifying the domain of the structure with an initial segment of the integers, and allowing some arithmetic on them. However, we want these extensions to define properties of the structures, and not to depend on a particular ordering on their elements: thus we focus on invariant extensions of FO.

If the only predicate allowed is the order, we get *order-invariant first-order logic*, $<$-inv FO. Restricting a bit the additional relation, we get *successor-invariant first-order*, Succ-inv FO. In this formalism, we only grant an access to the successor relation derived from the order, provided that the evaluation of a sentence using this successor relation is independent of the choice of a particular successor. The focus of this thesis will be on those two logics. Note that we use here the term "logic" somewhat liberally, since $<$-inv FO and Succ-inv FO do not have a recursive syntax, which is a usual requirement for a logic.

A strong motivation for the study of these two formalisms comes from database theory. On top of that, some profound relations between complexity theory and invariant logics exist, reinforcing the interest in these notions.

As databases are commonly stored on disks that implicitly order their memory segments, when one wishes to express a query in FO, one has access to an additional order on the elements of the database. However, making use of this order without care could result in queries that evaluate differently on two implementations of the same database, which is clearly an undesirable behavior breaking the physical data independence principle from [10]. We want to use this order only in an invariant manner; that way, the result of a query depends only on the database it is run on, and not on the way the data is stored on disk. This amounts exactly to the definition of $<$-inv FO, or Succ-inv FO if we restrict the way this order can be accessed.

It is straightforward that $<$-inv FO is at least as expressive as Succ-inv FO, which in turn can express any FO-definable property. Gurevich constructed a class of finite structures that can be defined by an $<$-inv FO sentence, but which is not FO-definable. Though this construction wasn't published by Gurevich, it can be found e.g. in Section 5.2 of [27]. Rossman extended this result, and proved in [33] that on finite structures, Succ-inv FO is strictly more expressive than FO.

Grohe and Schwentick [24] proved that these logics are Gaifman-local, giving an upper bound to their power of expression. Other upper bounds were given by Benedikt and Segoufin [5], who proved that $<$-inv FO, and hence Succ-inv FO, are included in MSO on classes of bounded treewidth and on classes of bounded degree. Elberfeld, Frickenschmidt and Grohe [16] extended the first inclusion to a broader setting, that of decomposable structures. Whether these logics are included in MSO in general is still an open question.

The classes of structures involved in the separating examples by Gurevich and Rossman are dense, and no other example is known on classes that are sparse. Far from it, $<$-inv FO and *a fortiori* Succ-inv FO are known to collapse to FO on several sparse classes, meaning that $<$-inv FO and FO can define the same properties on those classes of structures. Benedikt and Segoufin [5] proved the collapse on trees, while Eickmeyer, Elberfeld and Harwarth [14] obtained an analogous result on graphs of bounded tree-depth.

As for model-checking considerations, the result from Grohe, Kreutzer and Siebertz [23] stating that the model checking problem for FO is fixed-parameter tractable on nowhere dense classes of graphs has been extended to Succ-inv FO by Van den Heuvel, Kreutzer, Pilipczuk, Quiroz, Rabinovich and Siebertz [38], as long as we restrict ourselves to classes of bounded expansion.

Contribution: In this thesis, we improve the understanding of the expressive power of Succ-inv FO and $<$-inv FO by broadening the setting in which they are known to collapse to FO.

In Chapter 3, we prove that Succ-inv FO collapses to FO on classes of structures of bounded degree.

To do this, we show how to construct successors on two FO-similar structures (that is, structures that satisfy the same FO-sentences which quantification nesting is below some threshold) of small degree, such that the two structures remain FO-similar when considering the additional successor relation.

A sparsity notion orthogonal to degree boundedness is that of treewidth, which measures in some sense the distance from a structure to a tree. It is an open question whether $<$-inv FO or Succ-inv FO collapse to FO on classes of graphs of bounded treewidth, or even bounded pathwidth.

In Chapter 4, which is based on a joint work with Luc Segoufin [22], we take a step in that direction, by showing that $<$-inv FO collapses to FO over the class of hollow trees.

Hollow trees are a generalization of binary trees, and are closely related to structures of pathwidth 2.

More precisely, the vocabulary of hollow trees contains two binary relations. They are interpreted so that the resulting structure is a tree with the following features: each element has at most four neighbors - its first child, its last child and up to two siblings. One of the binary relation is symmetrical and defines the sibling relation while the other one is oriented and defines the partial parent-child relation. A parent may have an arbitrary number of children, but it is directly only related to two of them. Note that while this can be done in MSO, FO cannot reconstruct the complete parent-child relation of every node within a hollow tree.

The proof that $<$-inv FO is no more expressive that FO on hollow trees follows a strategy similar to the that used for binary trees in [5]: we first exhibit a set of operations over hollow trees (or, more precisely, over structures FO-similar to hollow trees) that preserve order-invariance similarity. We then show that if two hollow trees are FO similar then one of them can be transformed using our set of operations into the other, lifting FO similarity to $<$-inv FO similarity. The first part is standard, though it adds a new kind of operation to previously used set of operations, and makes use of the locality of $<$-inv FO [24]. The second part is more combinatorial and forms the main technical contribution of this Chapter.

Chapter 2

Preliminaries

In this chapter, we recall some of the basic definitions, examples and properties from finite model theory, as well as some notions relating to tree-decompositions. We also review the state of the art with respect to order-invariant and successor-invariant logics.

For a complete course in finite model theory, we refer the reader to [27].

We start by setting in Section 2.1 the general framework of finite model theory, by introducing relational structures and defining first-order and second-order logics on such structures. We then define in Section 2.2 several extensions of these logics by means of invariant relations. Classical methods for establishing expressivity results for those logics are described in Section 2.3, and used in Section 2.4 where we recall the main results known about the expressive power of invariant logics.

In Section 2.5, we turn to structures of bounded treewidth and pathwidth. After recalling the classical definitions for graphs, we extend these notions to all structures and establish a logical framework for them.

Section 2.6 describes the strategies that will later be used in Chapter 3 and 4 for proving a collapse between two logics. We then give an overview of the complexity of the model checking problem in the case of successor-invariant first-order logic in Section 2.7, before concluding the chapter by listing in Section 2.8 all the expressivity results spread throughout the previous sections.

2.1 Structures and logics

We define the basic notions of finite model theory in three steps.

First, we explain in Section 2.1.1 how mathematical objects can be seen as relational structures. Sections 2.1.2 and 2.1.3 then provide two usual formalisms to express properties about these objects. Those languages, or logics, are respectively first-order logic, FO, and monadic second-order logic, MSO. While FO only allows quantification over elements of the structures (e.g. edges of a graph), MSO broadens the quantification to sets of elements, thus allowing to express more properties such as connectivity.

2.1.1 Relational structures

By **relational vocabulary**, we mean a set $\Sigma := \{R_1, \cdots, R_n, c_1, \cdots, c_m\}$ of symbols together with an **arity** function $\mathrm{Ar} : \{R_1, \cdots, R_n\} \to \mathbb{N}$. The R_i are called $Ar(R_i)$-ary relation symbols, and the c_j are called the constant symbols.

For a finite alphabet σ of symbols, we define the relational vocabulary

$$P_\sigma := \{P_s : s \in \sigma\},$$

where every P_s is a unary relation symbol, i.e. such that $\mathrm{Ar}(P_s) = 1$.

For the remainder of this chapter, we consider a relational vocabulary

$$\Sigma := \{R_1, \cdots, R_n, c_1, \cdots, c_m\}.$$

A **Σ-structure** $\mathcal{A}$ is a tuple $(A, R_1^{\mathcal{A}}, \cdots, R_n^{\mathcal{A}}, c_1^{\mathcal{A}}, \cdots, c_m^{\mathcal{A}})$ where:

- A is a nonempty set
- for every $i \in \{1, \cdots, n\}$, $R_i^{\mathcal{A}} \subseteq A^{\mathrm{Ar}(R_i)}$
- for every $j \in \{1, \cdots, m\}$, $c_j^{\mathcal{A}} \in A$

A is called the **domain** of $\mathcal{A}$. We say that each relation symbol R_i is **interpreted** as the $\mathrm{Ar}(R_i)$-ary relation $R_i^{\mathcal{A}}$, and that each constant symbol c_j is **interpreted** as the element $c_j^{\mathcal{A}}$

We say that $\mathcal{A}$ is a **finite Σ-structure** if its domain is finite.

Two structures $\mathcal{A}$ and $\mathcal{B}$ are **isomorphic** if there exists a bijection f between their domains such that f and f^{-1} preserve relations and constants. This is denoted $\mathcal{A} \simeq \mathcal{B}$. Every class of structures under consideration will be closed under isomorphism.

Let's illustrate the definition of structure with several examples:

Example 2.1.1. *Let σ be a finite alphabet. There are several ways to represent a word $w \in \sigma^*$ as a relational structure.*

First, let's consider an additional relation symbol $<$ of arity 2. Every word $w \in \sigma^*$ can be seen as a $P_\sigma \cup \{<\}$-structure $\mathcal{A}_w$, where:

- *A_w is the set $\{0, \cdots, |w| - 1\}$ (where $|w|$ denotes the length of w)*
- *$<$ is interpreted as the natural order on A_w*
- *for every letter $s \in \sigma$, P_s is interpreted as the set of positions at which the letter s occurs in w.*

*By extension, we will call such a structure $\mathcal{A}_w$ a **word** over σ, as it is the most common way in the literature to represent a word of σ^* as a logical structure.*

Now, let's consider the relational vocabulary $P_\sigma \cup \{S\}$ instead of $P_\sigma \cup \{<\}$, where S is also a binary relation symbol. We can define $\mathcal{B}_w$ as for $\mathcal{A}_w$, but where S is interpreted as the successor relation on $\{0, \cdots, |w| - 1\}$.

*Such a structure will be referred to as a **dipath** over σ.*

*If we now symmetrize the interpretation of S in $\mathcal{B}_w$, we get a **path** over σ.*

Depending on whether we choose to represent $w \in \sigma^*$ as a word, a dipath or a path, the properties that we are able to express may vary.

For instance on the alphabet $\sigma := \{a, b, c\}$, using first-order logic as defined in Section 2.1.2, one is able to state that there exists an occurrence of the letter a that comes before every occurrence of the letter b in $\mathcal{A}_w$; however this cannot be expressed in $\mathcal{B}_w$. In other words, whether the language

$$c^* a \sigma^*$$

is definable in first-order logic depends on the representation we chose.

Likewise, the language

$$c^* abc^* abc^*$$

is definable in first-order on words and dipaths, but not on paths.

As we will see in Section 2.1.3, the more expressive logic MSO *is able to define transitive closure. In that case, every property definable on words is also definable on dipaths.*

Example 2.1.2. *Similarly, one can easily consider any tree t over an alphabet σ as a relational structure over $P_\sigma \cup \{S\}$, where S is a binary symbol. Depending on the setting, we can interpret S as the ancestor-descendant relation, the parent-child relation, or their symmetric closure.*

As for words, the choice of the model has an impact on the properties that are expressible in weak logics, such as first-order logic. When considering words or trees as relational structures, we will always make clear which representation we choose.

Example 2.1.3. *More generally, any colored finite graph with a set σ of colors can be seen as a $P_\sigma \cup \{E\}$-structure, whose domain is the set of nodes of the graph, and which interprets the binary relation symbol E as the edge relation of the graph.*

It will be convenient to identify the graph and the corresponding structure (or better still, to define a graph as such a structure).

Proviso. *Unless stated otherwise, we will only consider finite structures.*

For clarity purposes, structures will always be denoted through calligraphic upper-case letters and their domain through the corresponding standard upper-case letter.

Given a Σ-structure $\mathcal{A}$, we will often identify the symbols $R_i, c_j \in \Sigma$ with their interpretations in $\mathcal{A}$.

The **Gaifman graph** $\mathcal{G}_\mathcal{A}$ of a Σ-structure $\mathcal{A}$ is defined as (A, V) where $(x, y) \in V$ iff x and y appear in the same tuple of a relation of $\mathcal{A}$. In particular, if a graph is seen as a relational structure on the vocabulary $\{E\}$, its Gaifman graph is the unoriented version of itself.

The **degree** of $\mathcal{A}$ is the degree of its Gaifman graph, and a class $\mathcal{C}$ of Σ-structures is said to be of **bounded degree** if there exists some $d \in \mathbb{N}$ such that the degree of every $\mathcal{A} \in \mathcal{C}$ is at most d.

Note 2.1.4. *There exist other definitions for the degree of a structure. For instance, one could consider the degree of an element $x \in A$ to be the total number of tuples in which it appears, i.e. the cardinal of the set*

$$\{\bar{y} \in R^{\mathcal{A}} : x \in \bar{y}, R \in \Sigma\}.$$

This definition is equivalent to the one we have adopted, in the sense that classes of bounded degree are the same for both notions, as long as the vocabulary is finite.

By $\mathrm{dist}_{\mathcal{A}}(x, y)$, we denote the distance between x and y in $\mathcal{G}_{\mathcal{A}}$. Given two sets S and T of elements of A and $m \in \mathbb{N}$, we say that S and T are **m-distant** in $\mathcal{A}$, if $\mathrm{dist}_{\mathcal{A}}(x, y) \geq m$ for all $x \in S$ and all $y \in T$.

We now give the definition of the neighborhood type of an element.

Let c be a constant symbol that doesn't appear in Σ.

For $k \in \mathbb{N}$ and $x \in A$, the **k-neighborhood** $\mathcal{N}^k_{\mathcal{A}}(x)$ of x is the $(\Sigma \cup \{c\})$-structure whose Σ-restriction is the substructure of $\mathcal{A}$ induced by $\{y \in A : \mathrm{dist}_{\mathcal{A}}(x, y) \leq k\}$, and where c is interpreted as x.

The **k-neighborhood type** $\tau = \mathrm{tp}^k_{\mathcal{A}}(x)$ is the isomorphism class of its k-neighborhood. We say that τ is a neighborhood type over Σ, and that x is an **occurrence** of τ in $\mathcal{A}$. $|\mathcal{A}|_\tau$ denotes the number of occurrences of τ in $\mathcal{A}$, and for $t \in \mathbb{N} \cup \{\infty\}$, we write $[\![\mathcal{A}]\!]_k =^t [\![\mathcal{B}]\!]_k$ to mean that for every k-neighborhood type τ, $|\mathcal{A}|_\tau$ and $|\mathcal{B}|_\tau$ are either equal, or both larger than t. If $t = \infty$, it means that $|\mathcal{A}|_\tau$ and $|\mathcal{B}|_\tau$ are equal for every k-neighborhood type τ.

We extend those definitions to tuples of elements by considering several new constant symbols, fixing the tuples pointwise.

This notion of neighborhood type is not to be confused with the other use of the word "type" in logic, which refers to the set of formulas with one free variable satsfied by an element. When speaking of types, we always refer to the former meaning.

2.1.2 First-order logic

Let Var be a countably infinite set of fresh symbols, called the **variable symbols**. We will usually denote the variable symbols through the letters $x, y, z, \cdots$

A **Σ-atom** is a token of the form $R_i(t_1, \cdots, t_{\mathrm{Ar}(R_i)})$ for $i \in \{1, \cdots, n\}$ or $t_1 = t_2$, where the t_k are either one of the constant symbols $c_1, \cdots, c_m$ or a variable symbol of Var.

The set of **first-order formulas over** Σ is defined as the closure of the set of atoms under the unary symbols $\{\neg\} \cup \{\exists x : x \in \mathrm{Var}\}$ and the binary symbols $\{\vee\}$.

The set $\mathrm{FO}(\Sigma)$ of **first-order sentences over** Σ contains every first-order formula φ over Σ having no **free variable**, i.e. such that for every $x \in \mathrm{Var}$ occurring in an atom a of φ, there exists an ancestor of a labelled $\exists x$ or $\forall x$ in the syntactic tree of φ.

Formulas and sentences will be denoted through the letters $\varphi, \phi, \psi, \cdots$

Σ will be omitted when it is clear from the context.

The **quantifier rank** of an FO-formula is the maximal number of nodes labelled $\exists x$, $x \in \mathrm{Var}$ on a branch of the syntactic tree of φ.

We now give a semantics *à la Tarksi* to FO(Σ).

Let $\mathcal{A}$ be a Σ-structure $\mathcal{A}$ together with a partial function v : Var $\rightarrow A$ covering all the free variables of a FO-formula φ. It will be convenient to extend v with $v : c_j \mapsto c_j^{\mathcal{A}}$ for $j \in \{1, \cdots, m\}$.

We define inductively the relation $\mathcal{A}, v \models \varphi$ as follows:

- φ is the atom $R_i(t_1, \cdots, t_n)$ and $(v(t_1), \cdots, v(t_n)) \in R_i^{\mathcal{A}}$
- φ is the atom $t_1 = t_2$, and $v(t_1) = v(t_2)$ (the first '=' is just a symbol, while the second one refers to the equality of elements of A)
- φ is of the form $\exists x.\psi$, and there exists an element $a \in A$ such that, defining v' as v together with $x \mapsto a$ (where this mapping takes precedence over $x \mapsto v(x)$ if $v(x)$ was defined), $\mathcal{A}, v' \models \psi$
- φ is of the form $\neg\psi$ (resp. $\psi \vee \phi$), and $\mathcal{A}, v \not\models \psi$ (resp. $\mathcal{A}, v \models \psi$ or $\mathcal{A}, v \models \phi$)

If φ is a FO-sentence, then we say that $\mathcal{A}$ **satisfies** φ if $\mathcal{A}, v \models \varphi$ where v is the empty valuation. We then write $\mathcal{A} \models \varphi$.

Similarly, if $x_1, \cdots, x_n \in$ Var are the free variables of φ, and $a_1, \cdots, a_n$ are elements of A, then for any partial function v : Var $\rightarrow A$,

$$\mathcal{A}, v \uplus \{x_1 \mapsto a_1, \cdots, x_n \mapsto a_n\} \models \varphi$$

will be abbreviated as

$$\mathcal{A} \models \varphi(a_1, \cdots, a_n)\,.$$

We have only defined the logical connectors $\exists x$, $\neg$ and $\vee$, but we will freely use the shortcuts $\forall x$, $\wedge$, $\rightarrow$ and $\leftrightarrow$, which can be derived from the first one accordingly to the well-known De Morgan laws.

Each sentence φ of FO(Σ) (and actually, of any logic $\mathcal{L}$) defines a property of Σ-structures; namely, the set $\{\mathcal{A} : \mathcal{A} \models \varphi\}$. When we mention the expressivity of a logic, we refer to the range of properties that can be defined with sentences of that logic.

We will define precisely what it means for a logic to be more expressive than another one on a class of structures in Section 2.3.

2.1.3 Monadic second-order logic

The expressive power of FO is very limited. We will see in Section 2.4.2 that it is local, namely that FO can only define properties that concern a small radius around a limited number of elements.

Even simple properties such as the connectivity of a graph cannot be defined in FO. We now turn to a fundamental extension of FO: monadic second-order logic, MSO. Not only the transitive closure of a relation is definable in MSO (and thus the connectivity property), but it has the very nice property of capturing exactly the regular languages on words and trees.

In FO, one can only quantify over elements of the structure, whereas in MSO, we allow quantification on sets of elements. For that, we need a new set of variables: let VAR be a countably infinite set of fresh symbols, called the **set-variable symbols**. We will usually refer to those variables as $X, Y, Z, \cdots$

We define the set of **MSO-formulas over** Σ as we did the set of FO-formulas over Σ, with the following additions:

- for every $X \in \mathrm{VAR}$ and $t \in \mathrm{Var} \cup \{c_1, \cdots, c_m\}$, $X(t)$ is an atom
- we consider the new unary operations $\{\exists X : X \in \mathrm{VAR}\}$ for the closure

As for FO, an **MSO-sentence** is an MSO-formula which has no free variable and no free set-variable.

The **quantifier rank** of an MSO-formula is the maximal number of nodes labelled $\exists x$ or $\exists X$, $x \in \mathrm{Var}, X \in \mathrm{VAR}$ on a branch of the syntactic tree of φ.

To define the semantics of an MSO-formula φ on a Σ-structure $\mathcal{A}$, in addition to the partial function $v : \mathrm{Var} \to A$ (extended to $\{c_1, \cdots, c_m\}$), we consider a partial function $V : \mathrm{VAR} \to \mathcal{P}(A)$ ($\mathcal{P}(A)$ being the power set of A) covering the set of free variables of φ.

In addition to the semantics rules for FO, we need the following rules:

- $\mathcal{A}, v, V \models X(c_i)$ if $v(c_i) \in V(X)$
- $\mathcal{A}, v, V \models \exists X.\psi$ if there exists $E \subseteq A$ such that $\mathcal{A}, v, V' \models \psi$, where V' is defined as V together with $X \mapsto E$ (this mapping taking once again precedence over $X \mapsto V(X)$ if it was defined)

If φ is an MSO-sentence, v and V can be omitted, and we write $\mathcal{A} \models \varphi$.
Once again, we will use the shortcut $\forall X.\psi$ for $\neg\exists X.\neg\psi$.

Example 2.1.5. *Let's prove that the reflexive transitive closure of a binary relation is* MSO*-definable. Given a binary relation $R \in \Sigma$, consider the following* MSO*-formula with two free variables:*

$$\phi_{R^*}(x, y) := \forall X.\ [X(x) \wedge (\forall u.\forall v.X(u) \wedge R(u, v) \to X(v))] \ \to \ X(y)\,.$$

For any Σ-structure $\mathcal{A}$ and elements $a, b \in A$, $\mathcal{A}, \{x \mapsto a, y \mapsto b\} \models \phi_{R^*}(x, y)$ if and only if (a, b) belongs to the reflexive transitive closure of $R^{\mathcal{A}}$.

Indeed, if $\mathcal{A}, \{x \mapsto a, y \mapsto b\} \models \phi_{R^*}(x, y)$, then mapping X to the set $\{e \in A : (a, e) \in R^*\}$ of all the elements reachable by R from a, we get that b is reachable from a.

Conversely, suppose that $(a, b) \in R^*$, and that we map X to a set E. If E satisfies the premise, then $a \in E$ and E is stable by R, which entails that E contains all the elements reachable by R from a. In particular, $b \in E$. Thus, $\mathcal{A}, \{x \mapsto a, y \mapsto b\} \models \phi_{R^*}(x, y)$.

We say that $\phi_{R^*}(x, y)$ defines the reflexive transitive closure of R.

Once again, to each MSO-sentence corresponds the class of structures that satisfy this sentence. Remember that we showed in Example 2.1.1 how a word on σ could be thought of as a $P_\sigma \cup \{<\}$-structure. We now state a fundamental result due to Büchi [9]:

Proposition 2.1.6. *The* MSO*-definable classes of words are exactly the regular languages.*

This means that for every regular language L, there exists an MSO-sentence φ_L which is satisfied by a word w (or more precisely, by its associated structure $\mathcal{A}_w$) iff $w \in L$. Conversely, the language of words satisfying any MSO-sentence is regular.

Note that when considering MSO, it doesn't matter whether we choose to include the full order $<$ or just the successor relation S in the translation from words to structures. Indeed, given a successor relation S, the MSO-formula ϕ_{S^*} from Example 2.1.5 defines exactly $<$. Conversely, it is not hard to write a formula defining S from $<$ (to define that restriction, FO is even expressive enough). We will encounter again such a notion of two-way definability, albeit only for FO, when we formally define bi-FO-interpretations in Section 2.3.2. This is a convenient tool to state that two models are equivalent.

A similar result for ranked trees was proven by Thatcher and Wright [36]:

Proposition 2.1.7. *The* MSO*-definable classes of trees are exactly the regular languages of trees.*

The previous proposition also holds on unranked trees.

Let's now add modulo quantification $\exists^{(m)}x$ for $m \in \mathbb{N}$, to MSO. The semantics of these new quantifiers is

$$\mathcal{A}, v, V \models \exists^{(m)}x.\varphi \quad \text{if} \quad |\{a \in A : \mathcal{A}, v \uplus \{x \mapsto a\}, V \models \varphi\}| \equiv 0 \bmod m\,.$$

In other words, we are now allowed to count, modulo some integer, the number of witnesses to a formula.

The logic thus defined is **monadic second-order logic with counting**, CMSO.

2.2 Invariant logics

In Section 2.1 were introduced two logical framework to express properties on relational structures. Let's now extend these logics by introducing some additional information on the structures at hand, which can only be used in an invariant way. For instance, a structure stored on a disk is enriched with an order on its elements: what happens if we let FO and MSO access this order, but require that the evaluation of a sentence on an ordered structure does not depend on the choice of the order?

We start by defining the general notion of invariance in Section 2.2.1, before specifying it to the case where the additional information is an order relation in Section 2.2.2, or a successor relation in Section 2.2.3).

2.2.1 General setting

Let Σ and Σ' be two disjoint relational vocabularies, and consider a class $\mathcal{C}'$ of Σ'-structures. Let $\mathcal{L} \in \{\mathrm{FO}, \mathrm{MSO}\}$ be a logic.

The notion of $\mathcal{C}'$-invariance allows us to express that an $\mathcal{L}$-formula possibly uses the structure of $\mathcal{C}'$ without depending on it.

Let $\mathcal{A}$ be a Σ-structure and let $\mathcal{A}'$ be a Σ'-structure with the same domain A. We define $(\mathcal{A}, \mathcal{A}')$ as the $(\Sigma \cup \Sigma')$-structure with domain A, where every symbol of Σ (resp. Σ') is interpreted as in $\mathcal{A}$ (resp. $\mathcal{A}'$).

We say that a sentence φ of $\mathcal{L}(\Sigma \cup \Sigma')$ is **$\mathcal{C}'$-invariant over $\mathcal{A}$** if, for every Σ'-structures $\mathcal{A}', \mathcal{A}'' \in \mathcal{C}'$ with domain A,

$$(\mathcal{A}, \mathcal{A}') \models \varphi \quad \leftrightarrow \quad (\mathcal{A}, \mathcal{A}'') \models \varphi\,.$$

In that case, the choice of the overlaying Σ'-structure of C' doesn't matter, and we write $\mathcal{A} \models \varphi$ if $(\mathcal{A}, \mathcal{A}') \models \varphi$ for any (or equivalently, for every) $\mathcal{A}' \in \mathcal{C}'$.

We say that φ is **$\mathcal{C}'$-invariant** if it is $\mathcal{C}'$-invariant over every finite Σ-structure. The set of $\mathcal{C}'$-invariant $\mathcal{L}(\Sigma \cup \Sigma')$-sentences is denoted **$\mathcal{C}'$-inv** $\mathcal{L}(\Sigma)$.

Note 2.2.1. *The reason for requiring the invariance only over finite structures will be apparent in Note 2.2.4. Indeed, if an* $\mathrm{FO}(\Sigma \cup \Sigma')$*-sentence is* $\mathcal{C}'$*-invariant over all finite and infinite structures, Beth Definability Theorem [6] ensures that it is equivalent to an* $\mathrm{FO}(\Sigma)$*-sentence as long as* $\mathcal{C}'$ *is* $\mathrm{FO}(\Sigma')$*-definable.*

We now define **$\mathcal{C}'$-inv**$/_{\mathcal{C}}$ $\mathcal{L}(\Sigma)$, where $\mathcal{C}$ is a class of Σ-structures, as the set of $\mathcal{L}(\Sigma \cup \Sigma')$-sentences that are $\mathcal{C}'$-invariant over every structure of $\mathcal{C}$.

When we prove that some invariant logic collapses to FO, we will mention whether the result still holds when we restrict the class of structures over which our sentences are required to be invariant.

Note that $\mathcal{L}(\Sigma) \subseteq \mathcal{C}'\text{-inv}/_{\mathcal{C}}\ \mathcal{L}(\Sigma)$ for every $\mathcal{C}, \mathcal{C}'$: a sentence of $\mathcal{L}(\Sigma)$ makes no use of symbols of Σ', and is thus necessarily $\mathcal{C}'$-invariant over $\mathcal{C}$.

In particular, $\mathcal{L}(\Sigma) \subseteq \mathcal{C}'\text{-inv}\ \mathcal{L}(\Sigma)$.

Furthermore, if $\mathcal{C}_1 \subseteq \mathcal{C}_2$, then

$$\mathcal{C}'\text{-inv}/_{\mathcal{C}_2}\ \mathcal{L}(\Sigma) \subseteq \mathcal{C}'\text{-inv}/_{\mathcal{C}_1}\ \mathcal{L}(\Sigma)\,.$$

The converse is not true. We will see later on, e.g. with $\mathcal{C}'$ being the class of linear orderings, $\mathcal{C}_2$ being the class of all finite and infinite structures and $\mathcal{C}_1$ being the class of all finite structures, that

$$\mathcal{C}'\text{-inv}/_{\mathcal{C}_1}\ \mathcal{L}(\Sigma) \not\subseteq \mathcal{C}'\text{-inv}/_{\mathcal{C}_2}\ \mathcal{L}(\Sigma)\,.$$

Let's now focus on two such classes $\mathcal{C}'$, which will be at the center of this thesis: the class of ordered structures, and the class of structures with a successor.

2.2.2 Order-invariance

Consider the vocabulary Σ' consisting of a single binary relation symbol $<$, and the class $\mathcal{C}'$ of linear orders, i.e. the class of Σ-structures $\mathcal{A}'$ such that $<^{\mathcal{A}'}$ is a total order on the domain A.

For the sake of simplicity, $\mathcal{C}'$-inv $\mathcal{L}(\Sigma)$ will be abbreviated as $<$-inv $\mathcal{L}(\Sigma)$, or even $<$-inv $\mathcal{L}$ when Σ is clear from the context.

If $\mathcal{L} = \mathrm{FO}$, we get $<$-inv FO, **order-invariant FO**. If $\mathcal{L} = \mathrm{MSO}$, we get $<$-inv MSO, **order-invariant MSO**.

To reformulate, an $\mathrm{FO}(\Sigma \cup \{<\})$-sentence φ is in $<$-inv FO iff for any Σ-structure $\mathcal{A}$ and any linear orders $<_1$ and $<_2$ on A,

$$(\mathcal{A}, <_1) \models \varphi \quad \text{iff} \quad (\mathcal{A}, <_2) \models \varphi\,.$$

Let's look at a non-trivial property definable in $<$-inv FO, based on a example from Potthoff [31].

Example 2.2.2. *An unordered binary tree with descendant is a* $\{S, D\}$*-structure whose restriction to* $\{S\}$ *is a binary tree (where* $S(x, y)$ *holds iff* x *is the parent of* y*, and a node has either* 0 *or* 2 *children), and whose interpretation of* D *is the transitive closure of* S*. Note that no distinction is made between the two children of an internal node.*

It is not hard to see that the class of unordered binary trees with descendant is definable in $\mathrm{FO}(\{S, D\})$*, through an* FO*-sentence* Φ*.*

Let $\mathcal{C}$ *be the class of unordered binary trees with descendant whose branches all agree on the parity of their length, i.e. such that all the paths from the root to the leaves have an odd length, or all of them have an even length. Let's show that* $\mathcal{C}$ *is definable in* $<$-inv $\mathrm{FO}(\{S, D\})$*.*

Suppose that we are given an order $<$ *on an unordered binary tree with descendant. Such an order induces an order among the children of any internal node, which can be used to state that the zig-zag branch going out of a node* x*, i.e. the path from* x *to a leaf alternating between first and second children, has even length. Indeed, consider the* $\mathrm{FO}(\{S, D, <\})$ *formula* $\varphi_{even\_zigzag}(x)$ *stating that either* x *is a leaf, or*

- *there exists a leaf* l *that is a descendant of* x
- *for every successive nodes* u_1, u_2 *and* u_3 *belonging to the branch between* x *and* l*,* u_2 *is the first child of* u_1 *iff* u_3 *is the second child of* u_2
- *the first node of the branch is the first child of* x
- l *is the second child of its parent.*

Let's emphasize that such a property is FO*-definable only because we considered the descendant relation; if a tree only has access to the parent-child relation, one wouldn't be able to define branches in* FO*.*

Now, note that if some unordered binary tree with descendant $\mathcal{T}$ *doesn't belong to* $\mathcal{C}$*, it must contain an internal node* t *whose two children* u *and* v *are such that all the branches from* u *are of even length, and all the branches from* v *are of odd length. An illustration is given in Figure 2.1.*

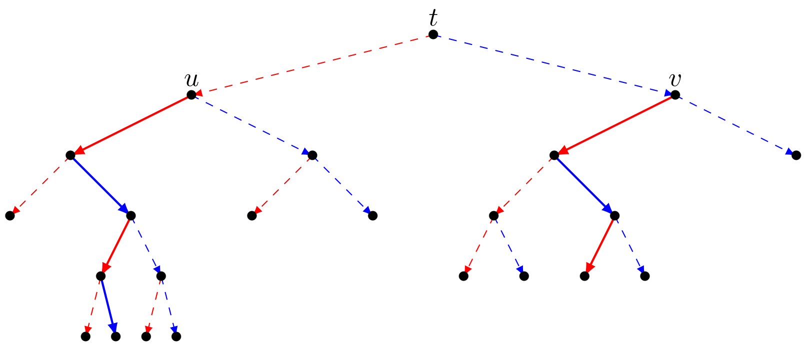

Figure 2.1: A node t witnessing that $\mathcal{T} \notin \mathcal{C}$. Red arrows denote the relation parent/first (wrt. $<$) child, and the blue ones the relation parent/second child. The zigzag branches of different parity are the drawn in plain arrows

Then a $\{S, D, <\}$-structure, which interprets $<$ as a linear order, satisfies

$$\varphi_{\mathcal{C}} \;:=\; \Phi \;\wedge\; \neg\exists t, u, v.\ S(t,u) \wedge S(t,v) \wedge \varphi_{even\_zigzag}(u) \wedge \neg\varphi_{even\_zigzag}(v)$$

iff it belongs to $\mathcal{C}$. Since belonging to $\mathcal{C}$ has nothing to do with the interpretation of $<$, it becomes apparent that $\varphi_{\mathcal{C}}$ is $<$-invariant, and thus $\mathcal{C}$ is definable in $<$-inv FO($\{S, D\}$).

It is important to note that two different orders on the same tree may give rise to different witnessing branches in φ_{even_zigzag}; however, the formula will evaluate similarly on both ordered structures.

As will be developed in Section 2.4.1, this example illustrates that $<$-inv FO *is strictly more expressive than* FO*, as it turns out that $\mathcal{C}$ is not definable in* FO($\{S, D\}$) *(this can easily be proven with the tools developed in Section 2.3.1).*

The problem of determining whether an FO-sentence using an order or a successor relation is invariant wrt. this relation is undecidable, by reduction from Trakhtenbrot's theorem. Hence we use here the term logic somewhat liberally, since having a recursive syntax is a usual requirement for a logic.

Proposition 2.2.3. *If Σ contains a relation symbol R with* $\mathrm{Ar}(R) \geq 2$*, then* $<$-inv FO(Σ) *doesn't have a recursive syntax.*

Proof. Let R be a relation symbol of Σ of arity 2. It is straightforward to adapt this proof to any relation symbol of arity at least 3.

Let INV be the problem at hand:

input: $\varphi \in \mathrm{FO}(\Sigma \cup \{<\})$

question: is φ order-invariant?

Let's consider the modified finite satisfiability problem MSAT for FO(Σ):

input: $\varphi \in \mathrm{FO}(\Sigma)$

question: does there exist a finite Σ-structure $\mathcal{A}$ such that:

$$\begin{cases} \mathcal{A} & \models \quad \exists x.\exists y.R(x,y) \;\wedge\; \exists x.\exists y.\neg R(x,y) \\ \mathcal{A} & \models \quad \varphi \end{cases}$$

The well-known Trakhtenbrot's theorem [37] states that MSAT in undecidable.

We reduce MSAT to the complement of INV: from a sentence $\varphi \in \mathrm{FO}(\Sigma)$, we compute the $\mathrm{FO}(\Sigma \cup \{<\})$-sentence

$$\bar{\varphi} := \varphi \;\wedge\; \forall \min .\forall \max .(\forall x.\neg(x < \min) \wedge \neg(\max < x)) \to R(\min, \max)\,.$$

If φ is a positive instance of MSAT, with a witnessing structure $\mathcal{A}$, then $\bar{\varphi}$ is not order-invariant over $\mathcal{A}$. Indeed, depending on the order, the second part of $\bar{\varphi}$ will or won't be satisfied.

Conversely, if $\bar{\varphi}$ is a negative instance of INV, then there is a struture $\mathcal{A}$ on which $\bar{\varphi}$ holds or not depending on the order. Such an $\mathcal{A}$ must satisfy φ, and $R^{\mathcal{A}}$ can be neither the complete nor the empty relation.

□

Note that determining whether a sentence is order-invariant is already undecidable on the class of dipaths [5].

If however each relation symbol of Σ is at most unary, then INV becomes decidable. Indeed, order-invariance then amounts to the commutativity of a language definable in FO (thus a regular language), which can be decided by checking whether its syntactic monoid is commutative.

Order-invariance is also decidable on the fragment of FO that uses only two variables, as shown by Harwath and Zeume [40].

Note 2.2.4. *In the definition, we required for any sentence φ of* $<$-inv FO *to be order-invariant over all finite structures, and not all finite and infinite structures. With the latter definition,* $<$-inv FO *would immediately collapse to* FO*. This can be proved either with Beth Definability Theorem [6], or with Craig's Interpolation Theorem [12] from classical model theory; we do the latter. Let's first state Craig's Interpolation Theorem:*

Proposition 2.2.5. *Let Σ_1, Σ_2 be two relational vocabularies and $\Sigma := \Sigma_1 \cap \Sigma_2$.*

Let Φ_1 and Φ_2 be respectively an $\mathrm{FO}(\Sigma_1)$*-sentence and an* $\mathrm{FO}(\Sigma_2)$*-sentence.*

If

$$\Phi_1 \to \Phi_2$$

holds in every (finite or infinite) $(\Sigma_1 \cup \Sigma_2)$-structure, there there exists an $\mathrm{FO}(\Sigma)$*-sentence Φ such that both*

$$\Phi_1 \to \Phi \quad \text{and} \quad \Phi \to \Phi_2$$

hold in every (finite or infinite) $(\Sigma_1 \cup \Sigma_2)$-structure.

Suppose that an $\mathrm{FO}(\Sigma \cup \{<\})$*-sentence φ is order-invariant over all finite and infinite structures. Using Proposition 2.2.5, let's prove that φ is equivalent to some* $\mathrm{FO}(\Sigma)$*-sentence ψ.*

Let $<_1$ and $<_2$ be two distinct order symbols, $\Sigma_1 := \Sigma \cup \{<_1\}$, $\Sigma_2 := \Sigma \cup \{<_2\}$, and φ_1 (resp. φ_2) be the result of the replacement in φ of all the occurrences of the symbol $<$ by the symbol $<_1$ (resp. $<_2$).

Now, let Ψ_1 (resp. Ψ_2) be an $\mathrm{FO}(\Sigma_1)$*-sentence (resp.* $\mathrm{FO}(\Sigma_2)$*-sentence) stating that $<_1$ (resp. $<_2$) is a linear order.*

The formula

$$\Psi_1 \wedge \varphi_1 \quad \to \quad (\Psi_2 \to \varphi_2)$$

holds in every finite or infinite $(\Sigma_1 \cup \Sigma_2)$-structure, by order-invariance of φ.

Proposition 2.2.5 guarantees the existence of some Σ-sentence ψ such that both

$$\Psi_1 \wedge \varphi_1 \quad \to \quad \psi \tag{2.1}$$

and

$$\psi \quad \to \quad (\Psi_2 \to \varphi_2) \tag{2.2}$$

hold in every finite or infinite $\Sigma \cup \{<_1, <_2\}$-structure.

Let's now prove that ψ is equivalent to φ on every Σ-structure $\mathcal{A}$:

- *Suppose that $\mathcal{A} \models \varphi$. Then for any linear order $<_1^{\mathcal{A}}$ on A,*

$$(\mathcal{A}, <_1^{\mathcal{A}}) \models \Psi_1 \wedge \varphi_1 \,,$$

  which together with (2.1) entails $(\mathcal{A}, <_1^{\mathcal{A}}) \models \psi$, hence $\mathcal{A} \models \psi$.

- *Suppose now that $\mathcal{A} \models \psi$. For any linear order $<_2^{\mathcal{A}}$ on A, (2.2) gives $(\mathcal{A}, <_2^{\mathcal{A}}) \models \varphi_2$, i.e. $\mathcal{A} \models \varphi$.*

This reasoning explains why order-invariance is only interesting on finite structures.

Furthermore, combined with the fact that FO $\subsetneq$ $<$-inv FO *(we've begun to see this with Example 2.2.2, and this will be developed in Section 2.4.1) , it shows that there is no equivalent to Craig's Interpolation Theorem in the finite.*

Let's now give an example of a property definable in $<$-inv MSO over the empty vocabulary:

Example 2.2.6. *One can write an* MSO($\{<\}$)*-formula Φ_{alt} with one free set variable X, stating that X contains every second element wrt. $<$:*

$$\Phi_{alt}(X) \quad := \quad \forall x.\forall y.\ (x < y \wedge \neg\exists z.\ x < z \wedge z < y) \rightarrow (X(x) \leftrightarrow \neg X(y))\,.$$

Using $\Phi_{alt}(X)$, it is now easy to express that the domain of a structure has even size, with

$$\begin{aligned} \varphi_{even} \quad := \quad \exists \min.\exists \max.\exists X. \quad & \forall x.\ \neg(x < \min) \wedge \neg(\max < x) \\ \wedge \quad & X(\min) \wedge \Phi_{alt}(X) \wedge \neg X(\max)\,. \end{aligned}$$

Indeed, an ordered structure satisfies φ_{even} iff the set containing the minimal element and every second element doesn't contain the last element. Note that this property depends only on the domain of the structure, and not on the order: thus $\varphi_{even} \in$ $<$-inv MSO($\emptyset$).

2.2.3 Successor-invariance

A order-invariant sentence can make use of an order as long as the result of its evaluation doesn't depend on the choice of this order. Let's now weaken the additional structure our sentences can access, by restricting the order to the successor relation.

We say that a binary relation on a finite set A is a **successor relation on** A if it is the graph of a circular permutation of A, i.e. a bijective function from A to A with a single orbit. This differs from the standard notion of successor in that there is neither minimal nor maximal element. However, this doesn't have any impact on our results, as we will prove in Proposition 2.2.7.

Let Σ' be the relational vocabulary containing a single binary relation symbol S, and let $\mathcal{C}'$ be the class of Σ'-structures $\mathcal{A}'$ such that $S^{\mathcal{A}'}$ is a successor relation on the domain A.

Then $\mathcal{C}'$-inv FO is abbreviated as Succ-inv FO, **successor-invariant FO**.

To reformulate, an $\mathrm{FO}(\Sigma \cup \{S\})$-sentence φ is in Succ-inv FO iff for any Σ-structure $\mathcal{A}$ and any successors S_1 and S_2 on A,

$$(\mathcal{A}, S_1) \models \varphi \quad \text{iff} \quad (\mathcal{A}, S_2) \models \varphi\,.$$

Let's now prove that our decision to consider circular successors instead of the more traditional linear ones (with a minimal and a maximal element) bears no consequence on the expressivity of the logic defined. If we define LinSucc-inv FO in the same way that Succ-inv FO, but where the invariant relation is a linear successor $\bar{S}$, we get:

Proposition 2.2.7. *For every vocabulary* Σ*,*

$$\text{Succ-inv FO}(\Sigma) = \text{LinSucc-inv FO}(\Sigma)\,,$$

i.e. Succ-inv $\mathrm{FO}(\Sigma)$ *and* LinSucc-inv $\mathrm{FO}(\Sigma)$ *define the same properties of* Σ*-structures.*

Proof. Given $\varphi \in$ Succ-inv FO, let's prove that there exists a formula

$$\psi \in \text{LinSucc-inv FO}$$

such that ψ is equivalent to φ, i.e. for every Σ-structure $\mathcal{A}$, $\mathcal{A} \models \varphi$ iff $\mathcal{A} \models \psi$.

Let ψ be defined as φ in which every atom $S(x, y)$ has been replaced with $\bar{S}(x, y) \vee \neg\exists z.(\bar{S}(x, z) \vee \bar{S}(z, y))$.

Let $\mathcal{A}$ be a Σ-structure and $\bar{S}^{\mathcal{A}}$ be a linear successor on A. Then

$$(\mathcal{A}, \bar{S}^{\mathcal{A}}) \models \psi \quad \text{iff} \quad (\mathcal{A}, S^{\mathcal{A}}) \models \varphi\,,$$

where $S^{\mathcal{A}}$ is the circular successor obtained from $\bar{S}^{\mathcal{A}}$ by adding an edge from the maximal element to the minimal one.

This guarantees that $\psi \in$ LinSucc-inv FO, and that ψ and φ are equivalent.

Conversely, let $\psi \in$ LinSucc-inv FO and let φ be the formula $\exists \min . \mathrm{Cut}(\psi)$, where $\mathrm{Cut}(\psi)$ is obtained by replacing in ψ every $\bar{S}(x, y)$ with $S(x, y) \wedge \neg y = \min$.

Let $\mathcal{A}$ be a Σ-structure, let $S^{\mathcal{A}}$ be a circular successor on A, and let $\min \in A$. Then

$$(\mathcal{A}, S^{\mathcal{A}}, \min) \models \mathrm{Cut}(\psi) \quad \text{iff} \quad (\mathcal{A}, \bar{S}^{\mathcal{A}}) \models \psi\,,$$

where $\bar{S}^{\mathcal{A}}$ is the linear successor obtained from $S^{\mathcal{A}}$ by removing the edge pointing to min. Hence $(\mathcal{A}, S^{\mathcal{A}}) \models \varphi$ iff there exists a linear successor $\bar{S}^{\mathcal{A}}$ obtained from $S^{\mathcal{A}}$ by an edge removal such that $(\mathcal{A}, \bar{S}^{\mathcal{A}}) \models \psi$, that is iff $\mathcal{A} \models \psi$.

This ensures that $\varphi \in$ Succ-inv FO and that φ and ψ are equivalent. □

Note that although LinSucc-inv FO and Succ-inv FO have the same expressive power, one may be more concise than the other. However, the succinctness (both in terms of quantifier rank and size of the formulas) of one with respect to the other cannot be more than linear.

Note 2.2.8. *There is no need to define* Succ-inv MSO*. Indeed, recall from Example 2.1.5 that the transitive closure is definable in* MSO*.*

Mimicking the proof techniques from Proposition 2.2.7, one can easily show that $<$-inv MSO *and* Succ-inv MSO *define the same properties.*

Let's give an example of Succ-inv FO sentence. For this first example, it will be more convenient to use the LinSucc-inv FO framework.

Example 2.2.9. *Let's consider the empty vocabulary. Note that* $\emptyset$*-structures are bare sets.*

We define $\mathrm{FO}(\{\bar{S}\})$*-formulas* $(\psi_k(x,y))_{k\in\mathbb{N}}$ *such that, if* $\bar{S}$ *is a linear successor on* A *and if* $a,b \in A$*, then* $\mathcal{A} \models \psi_k(a,b)$ *iff* a *and* b *are at distance* 2^k *according to* $\bar{S}$ *(not taking the orientation of* $\bar{S}$ *into consideration).*

The $(\psi_k(x,y))_k$ *are defined by induction on* k *as follows:*

- $\psi_0(x,y) := \bar{S}(x,y) \vee \bar{S}(y,x)$
- $\psi_{k+1}(x,y) := x \neq y \wedge \exists m.\forall z.(z = x \vee z = y) \to \psi_k(z,m)\,.$

The slight complexification in the definition of ψ_{k+1} *from* ψ_k *allows the* $(\psi_k(x,y))_k$ *to be linear in* k*, both in terms of quantifier rank and of size.*

For $n \in \mathbb{N}$*, we define* ϕ_n *as*

$$\exists \min . \exists \max . \ (\forall z.\neg\bar{S}(z,\min) \wedge \neg\bar{S}(\max,z)) \ \wedge \ \psi_n(\min,\max)\,.$$

Provided that $\bar{S}^{\mathcal{A}}$ *is indeed a linear successor relation on* A*, we get that* $\mathcal{A} \models \phi_n$ *iff* $|A| = 2^n + 1$*.*

Furthermore, it is clear that $\phi_n \in$ LinSucc-inv $\mathrm{FO}(\emptyset)$*, as the this property depends only on the size of* A*, and not on the particular interpretation of* $\bar{S}$*.*

It is well known, and can easily be shown using Ehrenfeucht-Fraïssé games, that any sentence of $\mathrm{FO}(\emptyset)$ *defining this property must have quantifier rank (and thus, size) at least* $2^n + 2$*.*

This example proves the following proposition.

Proposition 2.2.10. Succ-inv $\mathrm{FO}(\emptyset)$ *is at least exponentially more succinct than* $\mathrm{FO}(\emptyset)$*, both in terms of quantifier rank and size.*

For the same reasons as $<$-inv FO, Succ-inv FO doesn't have a recursive syntax in general:

Proposition 2.2.11. *If* Σ *contains a relation symbol* R *with* $\mathrm{Ar}(R) \geq 2$*, then* Succ-inv $\mathrm{FO}(\Sigma)$ *doesn't have a recursive syntax.*

2.3 Tools for proving expressivity results

We have defined a number of different logics in Sections 2.1 and 2.2, and seen that some are more expressive than others when it comes to the range of properties they can define. In this section, we develop the framework as well as some useful tools to prove such results about the expressivity of the previously defined logics. The first results that follow from these techniques are presented in Section 2.4.

The notion of similarity between two structures, and thus their indistinguishability by a given logic, is defined in Section 2.3.1. Ehrenfeucht-Fraïssé games, in which quantification is seen as the choice of a move in a two-player game, and their link to FO-similirity, are also developed in that section.

We then introduce in Section 2.3.2 the notion of FO-interpretation as a way to define structures from other structures, and show how they can be used to lift expressivity results from one class of structures to another.

Let's first establish some vocabulary to express relations between the expressive power of two logics.

Given two classes of Σ-structures $\mathcal{P} \subseteq \mathcal{C}$ and a logic $\mathcal{L}$, we say that $\mathcal{P}$ is **$\mathcal{L}$-definable on** $\mathcal{C}$ if there exists an $\mathcal{L}$-sentence φ such that

$$\forall \mathcal{A} \in \mathcal{C}, \quad \mathcal{A} \in \mathcal{P} \quad \text{iff} \quad \mathcal{A} \models \varphi \,.$$

If $\mathcal{C}$ is not mentioned, it is understood to be the class of all finite Σ-structures.

Note that the notion of $\mathcal{L}$-definability on $\mathcal{C}$ differs from the notion of $\mathcal{L}$-definability as soon as $\mathcal{C}$ itself is not $\mathcal{L}$-definable.

Given a class of structures $\mathcal{C}$ and two logics $\mathcal{L}, \mathcal{L}'$, we say that $\mathcal{L}'$ **is at least as expressive as** $\mathcal{L}$ **on** $\mathcal{C}$, abbreviated $\mathcal{L} \subseteq \mathcal{L}'$ on $\mathcal{C}$, if every property $\mathcal{P} \subseteq \mathcal{C}$ that is $\mathcal{L}$-definable on $\mathcal{C}$ is also $\mathcal{L}'$-definable on $\mathcal{C}$.

We write $\mathcal{L} = \mathcal{L}'$ on $\mathcal{C}$ if $\mathcal{L}$ and $\mathcal{L}$ are equally expressive on $\mathcal{C}$, i.e. if $\mathcal{L} \subseteq \mathcal{L}'$ on $\mathcal{C}$ and $\mathcal{L}' \subseteq \mathcal{L}$ on $\mathcal{C}$.

Once again, if $\mathcal{C}$ is not specified, we understand it to be the class of all finite Σ-structures.

Note 2.3.1. *One has to be careful with the class of structures over which invariance is assumed. In this thesis, by*

$$<\text{-inv FO} = \text{FO} \;\; \textit{on } \mathcal{C}$$

we mean that for every $\mathrm{FO}(\Sigma \cup \{<\})$*-sentence* φ *that is* <-invariant over all finite structures*, there exists an* $\mathrm{FO}(\Sigma)$*-sentence that is equivalent on* $\mathcal{C}$*.*

This is not to be confused with the stronger property requiring that, for every $\mathrm{FO}(\Sigma \cup \{<\})$*-sentence* ψ *that is* <-invariant over $\mathcal{C}$*, there exists an* $\mathrm{FO}(\Sigma)$*-sentence that is equivalent on* $\mathcal{C}$*. This is denoted*

$$<\text{-inv}/_{\mathcal{C}} \text{ FO} = \text{FO} \;\; \textit{on } \mathcal{C} \,.$$

The relations between such invariant logics with restricted invariance support do not relativize. For instance, while

$$<\text{-inv FO} = \text{FO} \;\; \textit{on } \mathcal{C}$$

entails

$$<\text{-inv FO} = \text{FO} \;\; \textit{on } \mathcal{C}'$$

as soon as $\mathcal{C}' \subseteq \mathcal{C}$*, there is no reason for*

$$<\text{-inv}/_{\mathcal{C}} \text{ FO} = \text{FO} \;\; \textit{on } \mathcal{C}$$

to entail

$$<\text{-inv}/_{\mathcal{C}'} \text{ FO} = \text{FO} \;\; \textit{on } \mathcal{C}' \,.$$

Indeed, with the weakening of the invariance requirement from $\mathcal{C}$ *to* $\mathcal{C}'$*, a sentence of* $<\text{-inv}/_{\mathcal{C}'}$ FO *could fall outside of* $<\text{-inv}/_{\mathcal{C}}$ FO*.*

In this section, we give a few standard tools that will be convenient when delimiting the expressive power of a logic.

2.3.1 Similarity and Ehrenfeucht-Fraïssé games

It will prove useful to measure to which extent two structures are alike, from the point of view of a specific logic.

Given a logic

$$\mathcal{L} \in \{\text{FO}, \text{MSO}, <\text{-inv FO}, \text{Succ-inv FO}\}$$

and two Σ-structures $\mathcal{A}, \mathcal{B}$, we say that $\mathcal{A}$ **is $\mathcal{L}$-similar to $\mathcal{B}$ at depth** k, and we note $\mathcal{A} \equiv_k^{\mathcal{L}} \mathcal{B}$, if $\mathcal{A}$ and $\mathcal{B}$ agree on every $\mathcal{L}$-sentence φ of quantifier rank at most k; that is, if $\mathcal{A} \models \varphi$ iff $\mathcal{B} \models \varphi$ for every such φ.

For every $k \in \mathbb{N}$ and $\mathcal{L}$, the relation $\equiv_k^{\mathcal{L}}$ is obviously an equivalence relation on the class of Σ-structures. Furthermore, it has finite index, i.e. the number of equivalence classes for this relation is finite (recall that we assumed the finiteness of Σ). On top of that, every equivalence class is definable by an $\mathcal{L}$-sentence of quantifier rank k.

The main technique for proving an inclusion between two logics follows from this remark:

Proposition 2.3.2. *Let Σ be a relational vocabulary, let $\mathcal{C}$ be a class of Σ-structures and let $\mathcal{L}, \mathcal{L}' \in \{\text{FO}(\Sigma), \text{MSO}(\Sigma), <\text{-inv FO}(\Sigma), \text{Succ-inv FO}(\Sigma)\}$.*

If there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that

$$\forall k \in \mathbb{N},\ \forall \mathcal{A}, \mathcal{B} \in \mathcal{C} \quad \mathcal{A} \equiv_{f(k)}^{\mathcal{L}'} \mathcal{B} \quad \to \quad \mathcal{A} \equiv_k^{\mathcal{L}} \mathcal{B}\,,$$

then

$$\mathcal{L} \subseteq \mathcal{L}' \text{ on } \mathcal{C}\,.$$

Proof. Let $\mathcal{P} \subseteq \mathcal{C}$ be a property $\mathcal{L}$-definable on $\mathcal{C}$: there exists an $\mathcal{L}$-sentence φ of quantifier rank k such that

$$\forall \mathcal{A} \in \mathcal{C}, \quad \mathcal{A} \in \mathcal{P} \quad \text{iff} \quad \mathcal{A} \models \varphi\,.$$

By definition, all the Σ-structures in any equivalence class for $\equiv_k^{\mathcal{L}}$ agree on φ. The class of Σ-structures satisfying φ is thus a union of equivalence classes for $\equiv_k^{\mathcal{L}}$; let E_φ be the (finite) set of equivalence classes of that union. We have

$$\mathcal{P} = \bigcup_{c \in E_\varphi} c \cap \mathcal{C}\,.$$

By hypothesis, for each $c \in E_\varphi$, there exists a (finite) set E'_c of equivalence classes for $\equiv_{f(k)}^{\mathcal{L}'}$ such that

$$c \cap \mathcal{C} = \bigcup_{c' \in E'_c} c' \cap \mathcal{C}$$

Combining these, we get

$$\mathcal{P} = \bigcup_{c \in E_\varphi} \bigcup_{c' \in E'_c} c' \cap \mathcal{C}\,.$$

All the unions being finite, and all the c' being definable via an $\mathcal{L}'$-sentence of quantifier rank $f(k)$, we get that $\mathcal{P}$ is $\mathcal{L}'$-definable on $\mathcal{C}$ by the disjunction of these sentences, thus proving

$$\mathcal{L} \subseteq \mathcal{L}' \text{ on } \mathcal{C}\,.$$

□

We now define Ehrenfeucht-Fraïssé games for FO, which capture exactly FO-similarity. There exist similar games for MSO, but we won't need them in this thesis.

Given two Σ-structures $\mathcal{A}$ and $\mathcal{B}$, the k-round Ehrenfeucht-Fraïssé game for FO(Σ) is played between two players: the Spoiler and the Duplicator. At round $i \in \{1, \cdots, r\}$, the Spoiler chooses an element $p_i^{\mathcal{A}}$ in A or $p_i^{\mathcal{B}}$ in B. The Duplicator then chooses an element $p_i^{\mathcal{A}}$ in A (if the Spoiler played in B) or $p_i^{\mathcal{B}}$ in B (if the Spoiler played in A).

After k rounds, one can think of the $(p_i^{\mathcal{A}})_i$ (resp. $(p_i^{\mathcal{B}})_i$) as the interpretations in $\mathcal{A}$ (resp. in $\mathcal{B}$) of some fresh constant symbols (p_i). Consider the new vocabulary $\Sigma' := \Sigma \uplus \{p_i : i \in \{1, \cdots, k\}\}$, and the two thus defined Σ'-structures $\mathcal{A}'$ and $\mathcal{B}'$.

The Duplicator has won the k-round Ehrenfeucht-Fraïssé game if for every relation symbol $R \in \Sigma \cup \{=\}$ of arity r ($=$ being of arity 2), and every constant symbols $c_1, \cdots, c_r \in \Sigma'$ (thus including the true constants of Σ, plus the elements chosen during the game),

$$\mathcal{A} \models R(c_1, \cdots, c_r) \quad \text{iff} \quad \mathcal{B} \models R(c_1, \cdots, c_r)\,.$$

We say that the Duplicator has a wining strategy for this game if he can win no matter the choices of the Spoiler.

It turns out the k-round Ehrenfeucht-Fraïssé game for FO captures exactly the FO-similarity at depth k:

Proposition 2.3.3. *For every $k \in \mathbb{N}$, the Duplicator has a winning strategy in the k-round Ehrenfeucht-Fraïssé game between $\mathcal{A}$ and $\mathcal{B}$ iff $\mathcal{A} \equiv_k^{\mathrm{FO}} \mathcal{B}$.*

The following result, due to Fagin, Stockmeyer and Vardi [18], is a key component of our proofs. It states that when the degree is bounded, an equivalence class for $\equiv_k^{\mathrm{FO}}$ is characterized by the number of occurrences of neighborhood types of a large enough radius, up to a large enough threshold.

Recall that $[\![\mathcal{A}]\!]_r =^t [\![\mathcal{B}]\!]_r$ means that $\mathcal{A}$ and $\mathcal{B}$ contain the same number of occurrences of each r-neighborhood type up to a threshold t.

Proposition 2.3.4. *For every vocabulary Σ, and for every $k, d \in \mathbb{N}$, there exist $r, t \in \mathbb{N}$ such that for every Σ-structures $\mathcal{A}$ and $\mathcal{B}$ of degree at most d,*

$$[\![\mathcal{A}]\!]_r =^t [\![\mathcal{B}]\!]_r \quad \to \quad \mathcal{A} \equiv_k^{\mathrm{FO}} \mathcal{B}\,.$$

Describing a winning strategy for the Duplicator in the k-round Ehrenfeucht-Fraïssé game between two structures gets out of hand as soon as the structures are slightly non-trivial, and is tedious task even on the simplest structures.

When proving that $\mathcal{A} \equiv_k^{\mathrm{FO}} \mathcal{B}$, Proposition 2.3.4 allows us to replace the description of such a strategy with a simple counting of every neighborhood type, which is often much more convenient.

Note that the converse obviously holds: when the degree is bounded, one can state in FO that every fixed-radius neighborhood type has a given number of occurrences, up to some threshold.

2.3.2 FO-interpretations and bi-FO-interpretations

At numerous points in this thesis, it will be convenient to define new structures from existing ones. MSO-transductions and FO-interpretations are probably the two most standard tools for doing so in a purely logical fashion. As our interest goes to weakly-expressive logics, we will only make use of the latter; being less permissive, it allows greater control on the thus defined structures.

Let Σ and Σ' be two relational vocabularies. It will be convenient to assume that $\Sigma' = \{R_1, \cdots, R_n\}$ is purely relational, i.e. doesn't contain any constant symbol.

An **interpretation** $\mathcal{I}$ **from** Σ **to** Σ' is a tuple $(\phi, \varphi_1, \cdots, \varphi_n)$ of FO(Σ)-formulas where

- ϕ has r free variables
- for $i \in \{1, \cdots, n\}$, φ_i has $\text{Ar}(R_i) \cdot r$ free variables.

We call **arity** of $\mathcal{I}$ the number r of free variables in ϕ, and **depth** of $\mathcal{I}$ the maximum among the quantifier ranks of $\phi, \varphi_1, \cdots, \varphi_n$.

Given a Σ-structure $\mathcal{A}$, the Σ'-structure $\mathcal{I}(\mathcal{A})$ is defined as follows:

- its domain is $\{(a_1, \cdots, a_r) \in A^r : \mathcal{A} \models \phi(a_1, \cdots, a_r)\}$
- for every $i \in \{1, \cdots, n\}$, the $\text{Ar}(R_i)$ r-tuples

$$(a_1^1, \cdots, a_r^1), \cdots, (a_1^{\text{Ar}(R_i)}, \cdots, a_r^{\text{Ar}(R_i)})$$

belong to the interpretation of R_i iff

$$\mathcal{A} \models \varphi_i(a_1^1, \cdots, a_r^1, \cdots, a_1^{\text{Ar}(R_i)}, \cdots, a_r^{\text{Ar}(R_i)}) .$$

Proposition 2.3.5. *Let* $\mathcal{L} \in \{\text{FO}, <\text{-inv FO}, \text{Succ-inv FO}\}$ *and let* $\mathcal{I}$ *be an* FO*-interpretation from* Σ *to* Σ'*, of arity* r *and depth* d*.*

For every Σ*-structures* $\mathcal{A}, \mathcal{B}$ *and for every* $k \in \mathbb{N}$*,*

$$\mathcal{A} \equiv_{rk+d}^{\mathcal{L}(\Sigma)} \mathcal{B} \quad \to \quad \mathcal{I}(\mathcal{A}) \equiv_k^{\mathcal{L}(\Sigma')} \mathcal{I}(\mathcal{B}) .$$

Proof. This is a well-known result when $\mathcal{L}$ is FO; let's start by proving the proposition in that case.

The interpretation $\mathcal{I}$ translates into a transformation $\mathcal{I}^{-1}$ from FO(Σ') to FO(Σ). $\mathcal{I}^{-1}$ is defined by induction as follows:

- for every $i \in \{1, \cdots, n\}$, and $x^1, \cdots, x^{\text{Ar}(R_i)} \in \text{Var}$,

$$\mathcal{I}^{-1}(R_i(x^1, \cdots, x^{\text{Ar}(R_i)})) := \varphi_i(x_1^1, \cdots, x_r^1, \cdots, x_1^{\text{Ar}(R_i)}, \cdots, x_r^{\text{Ar}(R_i)})$$

- for every $x, y \in \text{Var}$,

$$\mathcal{I}^{-1}(x = y) := \bigwedge_{j=1}^{r} x_j = y_j$$

- for every $x \in \mathrm{Var}$, $\mathcal{I}^{-1}(\exists x.\psi) := \exists x_1.\cdots.\exists x_r.\phi(x_1,\cdots,x_r) \wedge \mathcal{I}^{-1}(\psi)$
- $\mathcal{I}^{-1}(\neg\psi) := \neg\mathcal{I}^{-1}(\psi)$ and $\mathcal{I}^{-1}(\psi \vee \psi') := \mathcal{I}^{-1}(\psi) \vee \mathcal{I}^{-1}(\psi')$.

If $\psi \in \mathrm{FO}(\Sigma')$ is of quantifier rank k, then the quantifier rank of $\mathcal{I}^{-1}(\psi)$ is $rk + d$. Furthermore, for every Σ-structure $\mathcal{A}$,

$$\mathcal{A} \models \mathcal{I}^{-1}(\psi) \quad \leftrightarrow \quad \mathcal{I}(\mathcal{A}) \models \psi\,. \tag{2.3}$$

If $\mathcal{A} \equiv^{\mathrm{FO}}_{rk+d} \mathcal{B}$, then for any such ψ, both structures agree on $\mathcal{I}^{-1}(\psi)$, hence $\mathcal{I}(\mathcal{A})$ and $\mathcal{I}(\mathcal{B})$ agree on ψ by (2.3). This entails $\mathcal{I}(\mathcal{A}) \equiv^{\mathrm{FO}}_k \mathcal{I}(\mathcal{B})$.

Let's now prove that an analogous result holds for $<$-inv FO, the proof for Succ-inv FO being similar.

We expand the given interpretation $\mathcal{I}$ from Σ to Σ' into an interpretation $\mathcal{I}_<$ from $\Sigma\cup\{<\}$ to $\Sigma'\cup\{<\}$, with the new order being defined as the lexicographical order on r-tuples based on the first order. Note that $\mathcal{I}_<$ has the same depth as $\mathcal{I}$, since the lexicographical order is definable by an FO-formula of quantifier rank 0.

The proof relies on the following remark: given a sentence $\psi \in {<}\text{-inv FO}(\Sigma')$ of quantifier k, $(\mathcal{I}_<)^{-1}(\psi)$ is order-invariant. Indeed, if $<_1$ and $<_2$ are two orders on the domain of a Σ-structure $\mathcal{A}$, then

$$\begin{array}{llcll} (\mathcal{A},<_1) \models (\mathcal{I}_<)^{-1}(\psi) & \text{iff} \quad \mathcal{I}_<(\mathcal{A},<_1) & \models & \psi & \qquad \text{(by (2.3))} \\ & \text{iff} \quad \mathcal{I}(\mathcal{A}) & \models & \psi & \\ & \text{iff} \quad \mathcal{I}_<(\mathcal{A},<_2) & \models & \psi & \\ & \text{iff} \quad (\mathcal{A},<_2) & \models & (\mathcal{I}_<)^{-1}(\psi)\,. & \end{array}$$

Now, if $\mathcal{A} \equiv^{<\text{-inv FO}}_{rk+d} \mathcal{B}$, then for every $\psi \in {<}\text{-inv FO}(\Sigma')$ of quantifier rank at most k, $\mathcal{A}$ and $\mathcal{B}$ agree on the order-invariant sentence $(\mathcal{I}_<)^{-1}(\psi)$. Hence for any orders $<_A$ and $<_B$ on their respective domains, $\mathcal{I}_<(\mathcal{A},<_A)$ and $\mathcal{I}_<(\mathcal{B},<_B)$ agree on ψ. This entails, ψ being order-invariant, that $\mathcal{I}(\mathcal{A})$ and $\mathcal{I}(\mathcal{B})$ agree on ψ.

In the end, we get that $\mathcal{I}(\mathcal{A}) \equiv^{<\text{-inv FO}}_k \mathcal{I}(\mathcal{B})$.

□

We now define a tool for reducing the collapse of $<$-inv FO (or similarly, of Succ-inv FO) to FO from one class of structures to another.

Let $\mathcal{C}_1, \mathcal{C}_2$ be two classes of structures over the respective vocabularies Σ_1 and Σ_2.

We say that $\mathcal{C}_1$ is **bi-FO-interpretable** through $\mathcal{C}_2$ if there exist two FO-interpretations $\mathcal{I}_{12}$ and $\mathcal{I}_{21}$, respectively from Σ_1 to Σ_2, and from Σ_2 to Σ_1, such that for every $\mathcal{A} \in \mathcal{C}_1$, $\mathcal{I}_{12}(\mathcal{A}) \in \mathcal{C}_2$ and $\mathcal{I}_{21}(\mathcal{I}_{12}(\mathcal{A})) \simeq \mathcal{A}$. The following result is rather straightforward:

Proposition 2.3.6. *Let $\mathcal{L} \in \{<\text{-inv FO}, \text{Succ-inv FO}\}$.*

If $\mathcal{C}_1$ is bi-FO-interpretable through $\mathcal{C}_2$ and $\mathcal{L} = \mathrm{FO}$ over $\mathcal{C}_2$, then $\mathcal{L} = \mathrm{FO}$ over $\mathcal{C}_1$.

Proof. We show that there exists some function f such that for every $k \in \mathbb{N}$ and $\mathcal{A}, \mathcal{B} \in \mathcal{C}_1$, if $\mathcal{A} \equiv^{\mathrm{FO}}_{f(k)} \mathcal{B}$, then $\mathcal{A} \equiv^{\mathcal{L}}_k \mathcal{B}$. Proposition 2.3.2 then allows us to conclude.

As $\mathcal{L} = \text{FO}$ over $\mathcal{C}_2$ we know that there is a function g such that for all $k \in \mathbb{N}$ and $\mathcal{A}, \mathcal{B} \in \mathcal{C}_2$, if $\mathcal{A} \equiv^{\text{FO}}_{g(k)} \mathcal{B}$, then $\mathcal{A} \equiv^{\mathcal{L}}_{k} \mathcal{B}$: to any $\mathcal{L}$-sentence with a quantifier rank less than k, choose arbitrarily a FO-sentence equivalent to it over $\mathcal{C}_2$ and take $g(k)$ as the max of their quantifier rank.

Let r_{12}, d_{12} (resp. r_{21}, d_{21}) be the arity and depth of $\mathcal{I}_{12}$ (resp. $\mathcal{I}_{21}$), and set $f(k) := r_{12}g(r_{21}k + d_{21}) + d_{12}$.

Assume now that $\mathcal{A}, \mathcal{B} \in \mathcal{C}_1$ are such that $\mathcal{A} \equiv^{\text{FO}}_{f(k)} \mathcal{B}$.

Applying $\mathcal{I}_{12}$ to both structures gives us $\mathcal{I}_{12}(\mathcal{A}) \equiv^{\text{FO}}_{g(r_{21}k+d_{21})} \mathcal{I}_{12}(\mathcal{B})$, in virtue of Proposition 2.3.5.

Hence by the choice of g, $\mathcal{I}_{12}(\mathcal{A}) \equiv^{\mathcal{L}}_{r_{21}k+d_{21}} \mathcal{I}_{12}(\mathcal{B})$, which yields $\mathcal{A} \equiv^{\mathcal{L}}_{k} \mathcal{B}$ after applying $\mathcal{I}_{21}$ (Proposition 2.3.5). □

Note 2.3.7. *Proposition 2.3.6 offers a convenient way to lift a collapse result from one class of structures to another. However, it is not necessary to rely on a bi-*FO*-interpretation to prove such a result: it is enough to prove that there exists functions $f_{12}, f_{21} : \mathbb{N} \to \mathbb{N}$ such that, for every $\mathcal{A}, \mathcal{B} \in \mathcal{C}_1$, there exist $P_{\mathcal{B}}(\mathcal{A}), P_{\mathcal{A}}(\mathcal{B}) \in \mathcal{C}_2$ such that*

$$\forall n \in \mathbb{N}, \quad \mathcal{A} \equiv^{\text{FO}}_{f_{12}(n)} \mathcal{B} \quad \to \quad P_{\mathcal{B}}(\mathcal{A}) \equiv^{\text{FO}}_{n} P_{\mathcal{A}}(\mathcal{B}) \tag{2.4}$$

and such that

$$\forall n \in \mathbb{N}, \quad P_{\mathcal{B}}(\mathcal{A}) \equiv^{\mathcal{L}}_{f_{21}(n)} P_{\mathcal{A}}(\mathcal{B}) \quad \to \quad \mathcal{A} \equiv^{\mathcal{L}}_{n} \mathcal{B}\,. \tag{2.5}$$

Under these conditions, the proof of Proposition 2.3.6 can trivially be adapted to get that the collapse of $\mathcal{L}$ to FO *on $\mathcal{C}_2$ entails the collapse of $\mathcal{L}$ to* FO *on $\mathcal{C}_1$.*

*Bi-*FO*-interpretations have the advantage of being easy to manipulate and to guarantee (2.4) and (2.5). They can be thought of as uniform version of the previous statement: indeed, $P_{\mathcal{B}}(\mathcal{A})$ now only depends on $\mathcal{A}$, and not anymore on $\mathcal{B}$. In that sense, bi-*FO*-interpretations are reminiscent of bicontinuous mappings from $\mathcal{C}_1$ in $\mathcal{C}_2$: if two structures are close together in $\mathcal{C}_1$, so are their images in $\mathcal{C}_2$, and conversely.*

Note 2.3.8. *Though we will not concern ourselves with this matter throughout this thesis, note that this bi-*FO*-interpretation method can be adapted to lift other kinds of upper-bound results.*

For instance, if one knows that $<$-inv FO $\subseteq$ MSO *on a class $\mathcal{C}_2$ and wants to prove that this inclusion holds on another class $\mathcal{C}_1$, it is enough to show the existence of an* MSO*-transduction (which is basically the equivalent of an* FO*-interpretation, but with* MSO*-formulas) from $\mathcal{C}_1$ to $\mathcal{C}_2$ and of an* FO*-interpretation from $\mathcal{C}_2$ to $\mathcal{C}_1$.*

2.4 First expressivity results

We now turn to the main focus of this thesis: the study of the expressive power of order-invariant and successor-invariant logics.

We list in this section some known results about the expressivity of the logics defined in Sections 2.1 and 2.2. We start by showing in Section 2.4.1 that order and successor-invariant logics are more expressive (at least when the vocabulary is rich enough) than their classical counterparts.

On the other hand, the following sections establish upper bounds on the expressivity of invariant logics. The locality of invariant first-order logics, which amounts to their inability to distinguish between elements that are locally similar, is discussed in Section 2.4.2. We then present in Section 2.4.3 a series of sparse classes of structures on which order and successor-invariant first-order logics are known to have a limited expressive power: may that be an inclusion in MSO, or even a collapse to FO.

2.4.1 Separating examples

Let's first establish a few straighforward inclusions between plain and invariant first-order logics.

As a sentence of FO(Σ) makes use of neither the successor relation nor the order, it is obvious that FO $\subseteq$ Succ-inv FO and FO $\subseteq$ $<$-inv FO on any class of structures.

Furthermore, the FO-formula

$$\varphi_S(x, y) := x < y \ \wedge \ \neg\exists z.\ x < z \wedge z < y$$

defines the linear successor relation associated to an order. With this formula, we can define an interpretation $\mathcal{I}_S$ from $\Sigma \cup \{<\}$ to $\Sigma \cup \{S\}$ which transforms the order into the corresponding successor relation. Using the notations from the proof of Proposition 2.3.5, we get a transformation $(\mathcal{I}_S)^{-1}$ from FO($\Sigma \cup \{S\}$) to FO($\Sigma \cup \{<\}$). It is easy to notice that if ψ is successor-invariant, then $(\mathcal{I}_S)^{-1}(\psi)$ is order-invariant. Furthermore, $\mathcal{A} \models \psi$ iff $\mathcal{A} \models (\mathcal{I}_S)^{-1}(\psi)$, thus entailing LinSucc-inv FO $\subseteq$ $<$-inv FO.

Combining this inclusion with Proposition 2.2.7, we get the first sequence of inclusions

$$\text{FO} \subseteq \text{Succ-inv FO} \subseteq <\text{-inv FO}\,.$$

We can now ask ourselves whether there inclusions are strict. The second one is conjectured to be strict [33], but we know no proof of that fact.

The first example separating order-invariant FO from plain FO dates back to the beginning of the 1990s and is due to Gurevich, though he did not publish it. He exhibited a class of graphs that is definable in $<$-inv FO, but not in FO.

Proposition 2.4.1. *As long as Σ contains a relation of arity at least 2,*

$$\text{FO}(\Sigma) \subsetneq <\text{-inv FO}(\Sigma)\,.$$

Proof. We describe the class exhibited by Gurevich and give a sentence of $<$ -inv FO defining it. Full details of the proof that this class is not FO-definable can be found in Section 5.2 of [27].

To understand the genesis of this example, recall Example 2.2.6 where we defined an $<$-inv MSO-sentence stating that the domain of a structure has even size, by asserting the existence of a set containing the first element wrt. $<$, every second element, and not containing the last element.

Of course, one cannot directly quantify over sets in FO, but the idea of the Gurevich separating example is to simulate in FO quantification over sets, by having elements in the structure that behave like those sets. Then quantifying over those elements amounts to quantifying over the set they stand for, and it becomes possible to simulate the previous separating example in $<$-inv FO. The most natural way to treat subsets of X as elements is to take the powerset $\mathcal{P}(X)$ as working structure.

Let's consider the most general setting, where Σ only contains a binary relation symbol; it will be convenient to denote it $\subset$.

Given $n \in \mathbb{N}$, let $\mathcal{B}_n$ denote the boolean algebra over $\{1, \cdots, n\}$, that is the Σ-structure having the powerset $\mathcal{P}(\{1, \cdots, n\})$ as domain, and where $\subset$ is interpreted as the strict inclusion over the subsets of $\{1, \cdots, n\}$.

One can construct an FO-sentence ϕ defining the class of all Σ-structures that are isomorphic to some $\mathcal{B}_n$. Start by stating that $\subset$ is antireflexive, antisymmetric and transitive, and that there is exactly one element of in-degree 0 (corresponding to the empty set) included in all the elements. Then the elements of in-degree 1 are meant to be the singletons. From there, one can state in FO all the properties of boolean algebras (two sets are equal if they contain the same elements, i.e. if they are supersets of the same singletons; the union of two sets exists; etc.)

Recall the example separating $<$-inv MSO from MSO. In a structure satisfying ϕ, we mimick this idea by stating in a sentence ψ_{even} of $\text{FO}(\Sigma \cup \{<\})$ the existence of an element (of the powerset) containing the minimal singleton with respect to $<$, containing every second singleton wrt. $<$, and not containing the last singleton wrt. $<$.

Now, $\phi \wedge \psi_{\text{even}}$ defines exactly the Σ-structures that are isomorphic to $\mathcal{B}_n$ for some even n. Since this property does not depend on the order, we have that $\phi \wedge \psi_{\text{even}} \in <$-inv $\text{FO}(\Sigma)$.

It only remains to prove that for any $k \in \mathbb{N}$, there exists a large enough $n \in \mathbb{N}$ such that the Duplicator has a winning strategy in the k-round Ehrenfeucht-Fraïssé game between $\mathcal{B}_{2n}$ and $\mathcal{B}_{2n+1}$. Though this result seems quite obvious, describing precisely such a strategy is rather tedious; we refer the reader to Section 5.2 of [27], where one is described through a compositional argument.

Recall from Proposition 2.3.3 that Ehrenfeucht-Fraïssé games capture FO-similarity: this entails that FO cannot discriminate between boolean algebras over odd and even sets, as long as they are large enough. Thus there is no FO-sentence equivalent to $\phi \wedge \psi_{\text{even}}$, and FO $\subsetneq$ $<$-inv FO on Σ. ☐

To the best of our knowledge, there exist two other examples separating $<$-inv FO from FO. Otto [30] formulated a property relating to connectivity, definable in $<$-inv FO (but not in FO) in a specific setting involving once again a boolean algebra, used to simulate quantification over sets when only allowed to quantify over elements. The third example, very different from the above two, is due to Potthoff [31] and explicited in Example 2.2.2.

If, on the other hand, Σ doesn't contain relation of arity at least 2, $<$-inv FO collapses to FO:

Proposition 2.4.2. *If Σ contains only relations of arity at most 1, then*

$$\text{FO}(\Sigma) = <\text{-inv FO}(\Sigma)\,.$$

Proof. Let's assume for simplicity's sake that all the relations of Σ are of arity 1. Relations of arity 0 are not an issue, since they are nothing but propositional variables.

We show that for every $\alpha \in \mathbb{N}$, there exists $f(\alpha) \in \mathbb{N}$ such that for every Σ-structures $\mathcal{G}_1, \mathcal{G}_2$,

$$\mathcal{G}_1 \equiv^{\mathrm{FO}}_{f(\alpha)} \mathcal{G}_2 \quad \to \quad \mathcal{G}_1 \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{G}_2\,;$$

Proposition 2.3.2 then concludes the proof.

For that, we prove the existence, under this hypothesis, of orders $<_1$ and $<_2$ such that

$$(\mathcal{G}_1, <_1) \equiv^{\mathrm{FO}}_{\alpha} (\mathcal{G}_2, <_2)\,.$$

Let $f(\alpha) := 2^\alpha + 1$. For every $\mathcal{G}_1, \mathcal{G}_2$ such that $\mathcal{G}_1 \equiv^{\mathrm{FO}}_{f(\alpha)} \mathcal{G}_2$, we have that $[\![\mathcal{G}_1]\!]_0 =^{2^\alpha} [\![\mathcal{G}_2]\!]_0$.

For every subset S of Σ, let G_1^S (resp. G_2^S) be the set of elements of G_1 (resp. G_2) appearing exactly in the interpretations of the unary relations in S. Then $[\![\mathcal{G}_1]\!]_0 =^{2^\alpha} [\![\mathcal{G}_2]\!]_0$ means that for every $S \subseteq \Sigma$, G_1^S and G_2^S have the same cardinality up to a threshold 2^α.

Let's now fix an arbitrary order $\prec$ on $\mathcal{P}(\Sigma)$. This order yields preorders $\prec_1$ on G_1 and $\prec_2$ on G_2, such that $x \prec_i y$ iff $x \in G_i^S$ and $y \in G_i^{S'}$ with $S \prec_i S'$.

Let $<_1$ (resp. $<_2$) be any order on G_1 (resp. G_2) refining $\prec_1$ (resp. $\prec_2$). We claim that $(\mathcal{G}_1, <_1) \equiv^{\mathrm{FO}}_{\alpha} (\mathcal{G}_2, <_2)$: it is possible for the Duplicator to win the α-round Ehrenfeucht-Fraïssé game by applying, for each $S \subseteq \Sigma$, either the identity strategy (if $|G_1^S| = |G_2^S|$) or the well-known strategy guaranteeing that two linear orders of length at least 2^α are α-similar. Such a strategy is described, for instance, in the proof of Theorem 3.6 of [27].

This entails the collapse $\mathrm{FO}(\Sigma) = <$-inv $\mathrm{FO}(\Sigma)$. □

The proof of the following proposition, due to Rossman [33], is much more involved than the previous one.

Proposition 2.4.3. *As long as Σ contains a relation of arity at least 2,*

$$\mathrm{FO}(\Sigma) \subsetneq \text{ Succ-inv } \mathrm{FO}(\Sigma)\,.$$

We now turn to the expressive power of $<$-inv MSO. Example 2.2.6 ensures that MSO $\subsetneq$ $<$-inv MSO, as the parity of the domain is definable in $<$-inv MSO but not in MSO. A straightforward generalization of this example shows that $<$-inv MSO can count modulo any integer, yielding the inclusion CMSO $\subseteq$ $<$-inv MSO, which turns out to be strict as well, according to a result by Ganzow and Rubin [20]:

Proposition 2.4.4. CMSO $\subsetneq$ $<$-inv MSO

2.4.2 Locality

All the properties that one can define in first-order logic take into account only the immediate neighborhoods of a fixed number of elements. There is no hope, for instance, to define in FO the reachability property in a graph: a path joining two elements could be arbitrarily long and venture arbitrarily far from them.

The notion of locality is one way to formalize this idea. It comes in two versions: Gaifman locality, and Hanf locality.

A formula φ of any logic with free (element) variables $x_1, \cdots, x_n$ is said to be **Gaifman local with radius** r if for any structure $\mathcal{A}$, and any elements $a_1, \cdots, a_n$ and $b_1, \cdots, b_n$, if

$$\mathrm{tp}^r_{\mathcal{A}}(a_1, \cdots, a_n) = \mathrm{tp}^r_{\mathcal{A}}(b_1, \cdots, b_n)$$

then

$$\mathcal{A} \models \varphi(a_1, \cdots, a_n) \quad \text{iff} \quad \mathcal{A} \models \varphi(b_1, \cdots, b_n)\,.$$

In other words, the set of n-uple of elements of A defined by φ only depends on the r-neighborhood type of these tuples.

We say that a logic is **Gaifman local** if each formula φ of that logic is Gaifman local with some radius r_φ.

A formula φ of any logic with free variables $x_1, \cdots, x_n$ is said to be **Hanf local with radius** r if for any structures $\mathcal{A}$ and $\mathcal{B}$ and any elements $a_1, \cdots a_n \in A$ and $b_1, \cdots, b_n \in B$, if

$$[\![(\mathcal{A}, a_1, \cdots, a_n)]\!]_r =^\infty [\![(\mathcal{B}, b_1, \cdots, b_n)]\!]_r$$

(where it is understood that the a_i and b_i are the interpretations of new constants symbols c_i), then

$$\mathcal{A} \models \varphi(a_1, \cdots, a_n) \quad \text{iff} \quad \mathcal{B} \models \varphi(b_1, \cdots, b_n)\,.$$

In the case where φ is a sentence, stating that φ is Hanf local with radius r amounts to saying that φ cannot discriminate between two structures that have the same number of occurrences of every r-neighborhood type.

We say that a logic is **Hanf local** if each formula φ of that logic is Hanf local with some radius r_φ.

Hella, Libkin and Nurmonen [25] proved that Hanf locality is a stronger notion than Gaifman locality:

Proposition 2.4.5. *If a formula φ is Hanf local with radius r, then it is Gaifman local with radius $3r + 1$.*

The following result, due to Grohe and Schwentick [24], allows to prove many inexpressibility results for $<$-inv FO. The core of their proof is a construction of orders that maintain FO-similarity between two very close structures: we will use their construction in Chapter 4 to prove the order-invariance of some operations.

Proposition 2.4.6. $<$-inv FO *is Gaifman local.*

However, whether $<$-inv FO is Hanf local is still an open question.

For $k \in \mathbb{N}$, let's define FO+MOD$_k$ as the extension of FO with modulo-counting quantifiers of the form $\exists^{i \bmod k} x$ with $0 \le i < k$, such that $\mathcal{A}, v \models \exists^{i \bmod k} x.\psi$ iff

$$|\{a \in A : \mathcal{A}, v' \models \psi\}| = i \bmod k\,,$$

where v' is defined as v together with $x \mapsto a$.

In other words, FO+MOD$_k$ allows to count modulo k. A consequence of the results of [25] is that for every $k \in \mathbb{N}$, FO+MOD$_k$ is Gaifman local.

It is intersting to note that unlike for plain FO, the addition of an order to FO+MOD$_k$ breaks the local property of the logic. Indeed, Niemistö [29] showed that for any $k \geq 2$, $<$-inv (FO+MOD$_k$) is not Gaifman local. In particular, his example separates $<$-inv (FO+MOD$_k$) from FO+MOD$_k$. He however proved that $<$-inv (FO+MOD$_k$) retains some weaker notion of locality, which he called *alternating Gaifman locality.*

Let's now consider the infinite vocabulary Σ' consisting of a relation symbol $\tilde{P}$ for every numerical predicate $P \subseteq \mathbb{N}^r$. Let Arb be the class containing all the Σ'-structures whose domain can be identified with an initial segment of $\mathbb{N}$ so that every $\tilde{P}$ is interpreted as P.

In accordance with the definition of invariant logics from Section 2.2, let Arb-inv FO(Σ) denote the set of FO($\Sigma \cup \Sigma'$)-sentences that are invariant wrt Arb. Arb-inv FO is obviously more expressive than $<$-inv FO, as the order is a particular numerical predicate.

Anderson, van Melkebeek, Schweikardt and Segoufin [1] studied the locality of Arb-inv FO, and proved that it was not Gaifman local. However, Arb-inv FO is Gaifman local with polylogarithmic radius in the following sense: for every formula $\varphi \in$ Arb-inv FO, there exists a polylogarithmic function f such that if two tuples $\bar{a}$ and $\bar{b}$ of elements of some structure $\mathcal{A}$ have the same $f(|A|)$-neighborhood type, then

$$\mathcal{A} \models \varphi(\bar{a}) \quad \text{iff} \quad \mathcal{A} \models \varphi(\bar{b})\,.$$

Whether a similar result holds when considering Hanf locality is not known.

2.4.3 Upper bounds and collapses of invariant logics

Let's define informally ∃SO (resp. ∀SO) as second-oder logic where, in contrast with MSO which can only quantify over unary relations, one can quantify over relations of arbitrary arity, provided that this is done in a purely existential (resp. universal) fashion. In those logics, it is possible to introduce an order relation, which is either quantified existentially or universally (which is equivalent by order-invariance). It is thus apparent that

$$<\text{-inv FO} \subseteq \exists\text{SO} \cap \forall\text{SO}\,.$$

Using the fact that an order is MSO-definable modulo set variables on classes of structures of bounded degree, Benedikt and Segoufin [5] proved the following upper bound for $<$-inv FO.

Proposition 2.4.7. *Let $d \in \mathbb{N}$ and let $\mathcal{C}_d$ denote the class of graphs of degree at most d. Then*

$$<\text{-inv}/_{\mathcal{C}_d}\ \text{FO} \subseteq \text{MSO} \textit{ on } \mathcal{C}_d\,.$$

Recall the discussion from Note 2.3.1 about the class on which sentences are required to be invariant. Proposition 2.4.7 states that for any FO-sentence φ

that is order-invariant over all the graphs of $\mathcal{C}_d$, there exists an MSO-sentence that is equivalent to φ on $\mathcal{C}_d$. In particular, this result is stronger than

$$<\text{-inv FO} \subseteq \text{MSO on } \mathcal{C}_d\,,$$

which only asserts that for every FO-sentence φ that is *order-invariant over all finite graphs*, there exists an MSO-sentence that is equivalent to φ on $\mathcal{C}_d$.

In the same article, Benedikt and Segoufin proved the following collapse on trees. Here, trees do not have access to the descendant relation (recall Example 2.2.2). Trees can either be n-ranked (meaning that the degree is at most n and the children of any node are ordered through an other binary relation) or unranked and unordered (in which case the degree is not bounded, but there is no distinction between elements of a siblinghood).

Proposition 2.4.8. *Let $\mathcal{T}ree$ be either the class of n-ranked trees for some $n \in \mathbb{N}$ or the class of unranked trees. Then*

$$<\text{-inv}/_{\mathcal{T}ree}\ \text{FO} = \text{FO} \ \textit{on}\ \mathcal{T}ree\,.$$

Once again this entails that $<$-inv FO = FO on $\mathcal{T}ree$.

It is not known whether $<$-inv FO collapses to FO on classes of trees of unbounded degree when the siblinghoods are ordered. However, in these conditions, the lexicographical order is definable in MSO, which guarantees that $<$-inv FO $\subseteq$ MSO.

Many techniques used in Chapter 4 are closely inspired from those presented in this article. We discuss them more in depth in Section 2.6.

The collapse on trees was proved independently by Niemistö [29].

The treedepth of a graph $\mathcal{G}$ is defined as the minimal height of a forest $\mathcal{F}$ with the descendant relation (recall Example 2.2.2), such that $\mathcal{G}$ is a subgraph of $\mathcal{F}$ when forgetting about the orientation of edges.

Eickmeyer, Elberfeld and Harwath [14] proved the following result:

Proposition 2.4.9. *Let $\mathcal{T}ree\mathcal{D}epth_n$ be the class of graphs of treedepth at most n. Then*

$$<\text{-inv}/_{\mathcal{T}ree\mathcal{D}epth_n}\ \text{FO} = \text{FO} \ \textit{on}\ \mathcal{T}ree\mathcal{D}epth_n\,.$$

2.5 Treewidth and pathwidth

When one takes a look at the separating examples from Section 2.4.1, it seems that dense classes of graphs are required for order and successor-invariant logics to be more expressive than plain first-order logic. On the other hand, results from Section 2.4.3 lead us to believe that on sparse classes, these invariant logics collapse to FO.

Trying to extend these collapse results, it is natural to study the behavior of invariant logics on classes of bounded treewidth and bounded pathwidth. These notions were introduced in 1986 by Robertson and Seymour [32] to measure how far a given graph is from being a tree or a path.

We first recall those definitions in Section 2.5.1, and extend them in Section 2.5.2 to all purely relational structures. We then consider a special kind of decompositions, more fitting for our purpose, in Sections 2.5.3 and 2.5.4.

2.5.1 Treewidth and pathwidth of a graph

Let $\mathcal{G}$ be a graph. A tree-decomposition of $\mathcal{G}$ is a tree $\mathcal{T}$ together with a bag function $\beta : T \to \mathcal{P}(G)$ such that

- every element $x \in G$ appears in some bag: $\exists t \in T, x \in \beta(t)$
- for every edge (x, y) of $\mathcal{G}$, $\exists t \in T. x \in \beta(t) \wedge y \in \beta(t)$
- the set of elements of $\mathcal{T}$ containing a given element of $\mathcal{G}$ is connected:

$$\forall x \in G,\ \{t \in T : x \in \beta(t)\} \text{ is connected in } \mathcal{T}\,.$$

It will be convenient to assume that a bag is always non-empty.

A tree-decomposition $(\mathcal{T}, \beta)$ is called path-decomposition if $\mathcal{T}$ is a dipath.

The width of the tree or path-decomposition $(\mathcal{T}, \beta)$ is $\max\{|\beta(t)| : t \in T\} - 1$. The first condition is only required to ensure that the notion of width makes sense when there are only isolated points in the graph. We say that a class $\mathcal{C}$ has treewidth (resp. pathwidth) $k \in \mathbb{N}$ if every graph of $\mathcal{C}$ admits a tree-decomposition (resp. path-decomposition) of width at most k, and $\mathcal{C}$ is said to be a class of bounded treewidth (resp. bounded pathwidth) if it has treewidth (resp. pathwidth) at most k for some $k \in \mathbb{N}$.

Example 2.5.1. *Let $\mathcal{T}$ be a tree. It is not hard to see that $\mathcal{T}$ together with*

$$\beta :\ t \in T\ \mapsto\ \{t, u\}\,,\ \textit{where } u \textit{ is the parent of } t \textit{ in } \mathcal{T} \textit{ if it exists,}$$

is a tree-decomposition of $\mathcal{T}$.

Hence the class of trees has treewidth 1*. This is the reason we subtracted* 1 *to* $\max\{|\beta(t)| : t \in T\}$ *in the definition of treewidth.*

Note 2.5.2. *Any class of graphs of treedepth d has pathwidth at most d: visiting in a lexicographical order the leaves of a forest witnessing the fact that a graph has treedepth d yields a path-decomposition of width d.*

Benedikt and Segoufin [5] gave an upper bound to $<$ -inv FO when the treewidth of a class is bounded:

Proposition 2.5.3. *Let $\mathcal{T}ree\mathcal{W}idth_n$ be the class of graphs of treewidth at most $n \in \mathbb{N}$. Then*

$$<\text{-inv}/_{\mathcal{T}ree\mathcal{W}idth_n}\ \text{FO} \subseteq \text{MSO}\ \textit{on}\ \mathcal{T}ree\mathcal{W}idth_n\,.$$

Elberfeld, Frickenschmidt and Grohe [16] extended this result to what they called *decomposable structures*, which covers graphs of bounded treewidth as well as planar graphs.

Proposition 2.5.4. *On the class of planar graphs, and on any class of bounded treewidth:*

- $<$-inv FO $\subseteq$ MSO
- $<$-inv MSO $=$ CMSO .

2.5.2 Logical perspective on decompositions

Our aim is to combine FO-interpretations and tree or path-decompositions in an attempt to lift the collapse $<$-inv FO = FO on trees from Proposition 2.4.8 to classes of bounded treewidth or pathwidth.

Recall the discussion from Section 2.3.2: the first step towards lifting this collapse is to show that two FO-similar graphs admit FO-similar decompositions.

To give a meaning to the notion of FO-similarity between two decompositions, we need to fix a logical framework for these decompositions. In the process, we extend these definitions to the context of purely relational structures. We omit the constant symbols for simplicity's sake, but one could consider vocabulary with constant with an extra bit of care.

Intuitively, a tree-decomposition is a colored tree, where the color of a node describes the content of the corresponding bag (i.e. the isomorphism class of the substructure induced by the bag) as well as how to glue that bag to the bags of the parent and the children of the node.

Let Σ be a relation vocabulary without constant symbols. For $k \in \mathbb{N}$, let's consider the vocabulary $\Lambda_k := \Sigma \cup \{I_0, \cdots, I_k, O_0, \cdots, O_k\}$ where all the I_i, O_i are fresh unary relation symbols. Bags of a decomposition of width k are seen as Λ_k-structures, where all the relations in Σ are inherited from the underlying structure, and the I_i (resp. O_i) contain the elements that also appear in the parent's bag (resp. some child's bag).

More precisely, we say that a Λ_k-structure $\mathcal{B}$ is a k-bag if it satisfies the following conditions:

- $|B| \leq k+1$
- for every $0 \leq i \leq k$, $|I_i^{\mathcal{B}}| \leq 1$ and $|O_i^{\mathcal{B}}| \leq 1$
- for every $0 \leq i < j \leq k$, $I_i^{\mathcal{B}} \cap I_j^{\mathcal{B}} = O_i^{\mathcal{B}} \cap O_j^{\mathcal{B}} = \emptyset$.

Let σ_k be the (finite) set of isomorphism classes of k-bags, and $\Sigma_k := P_{\sigma_k} \cup \{S\}$ where P_{σ_k} is the vocabulary defined in Section 2.1.1 and S is a new binary relation symbol. For $c \in \sigma_k$, we will write $c \models \varphi$ if any (or equivalently, every) $\mathcal{B} \in c$ satisfies φ.

We say that a Σ_k-structure $\mathcal{T}$ is a **valid tree-decomposition of width** k if

- its restriction to $\{S\}$ is an unranked tree, meaning that an element may have an arbitrary number of children, and makes no difference among them
- the interpretations of the unary relations of P_{σ_k} partition its domain T
- $\forall c \in \sigma_k$, $\forall t \in T$ such that $\mathcal{T} \models P_c(t)$ and $\forall i \in \{0, \cdots, k\}$,

  $c \models \exists x, O_i(x)$ iff there exists a child u of t in $\mathcal{T}$ such that $\mathcal{T} \models P_d(u)$ with some $d \in \sigma_k$ such that $d \models \exists y, I_i(y)$
- if $\mathcal{T} \models P_c(r)$ where r is the root of $\mathcal{T}$, then $c \models \bigwedge_{0 \leq i \leq k} \neg \exists x, I_i(x)$.

Basically, these conditions amount to saying that an *output* element O_i is always interfaced with a corresponding *input* element I_i in some child's bag, and conversely. Those elements will be merged in the construction to come.

Let TD_k be the class of all valid tree-decompositions of width k, and PD_k the subclass of TD_k containing all the tree-decompositions whose restriction to $\{S\}$ is a dipath.

Every $\mathcal{T} \in \mathrm{TD}_k$ generates a Σ-structure $\mathcal{E}\mathrm{xt}(\mathcal{T})$ in the natural way: we take the disjoint union of the Σ-structures which correspond to the colors of all the elements of $\mathcal{T}$, and we identify in this union an element colored with I_i with the element colored with the corresponding O_i in the parent's bag. Note that all of this can be done via an MSO-transduction.

More precisely, $\mathcal{E}\mathrm{xt}(\mathcal{T})$ is constructed as follows. For every $c \in \sigma_k$ and every element $t \in \mathcal{T}$ belonging to $\mathcal{P}_c^{\mathcal{T}}$, we consider a k-bag $\mathcal{B}_t$ belonging to c, such that the $(\mathcal{B}_t)_{t\in T}$ have pairwise disjoint domains. Let $A := \bigcup_{t\in T} B_t$. We define the interface relation I on A, such that $I(x, y)$ holds iff $x \in B_t$ and $y \in B_u$ for some $t, u \in T$, such that

- $\mathcal{T} \models S(t, u)$, i.e. t is the parent of u in $\mathcal{T}$
- $\mathcal{B}_t \models O_i(x)$ for some $0 \le i \le k$
- $\mathcal{B}_u \models I_i(x)$ for the same i .

Now, let $\sim$ be the most coarse-grained equivalence relation extending I. $\mathcal{E}\mathrm{xt}(\mathcal{T})$ is defined as follows:

- its domain is $A/\!\sim$
- for every r-ary relation symbol $R \in \Sigma$,

  for every equivalence classes $\bar{x}_1, \cdots, \bar{x}_r$ for $\sim$,

  $\mathcal{E}\mathrm{xt}(\mathcal{T}) \models R(\bar{x}_1, \cdots, \bar{x}_r)$ iff there exist $t \in T$, and $x_1, \cdots, x_n \in B_t$ such that $x_i \in \bar{x}_i$ $(1 \le i \le r)$ and $\mathcal{B}_t \models E(x_1, \cdots, x_r)$.

Roughly speaking, a tuple belongs to the interpretation of R in $\mathcal{E}\mathrm{xt}(\mathcal{T})$ iff it appears in some bag of $\mathcal{T}$.

An illustration of this process is given in Figure 2.2. In this example, $k = 2$ and Σ is the vocabulary of graphs colored with four colors, i.e. $\Sigma = \{E, P_{\mathrm{red}}, P_{\mathrm{yellow}}, P_{\mathrm{green}}, P_{\mathrm{blue}}\}$, represented in the figure as colored circles. The colored squares represent the equivalence classes of 2-bags, i.e. the unary predicates of Σ_2.

We are now ready to define the logical notion of tree-decompositions and treewidth in the case of Σ-structures. For the remainder of this thesis, let's consider that these definition override the classical ones given in Section 2.5.1 for graphs. This is not an issue, since there is a correspondence between decompositions in both senses. In particular, a graph (seen as a $\{E\}$-structure) admits a tree-decomposition in TD_k in the following sense iff it has treewidth k in the classical sense.

A **tree-decomposition** (resp. **path-decomposition**) **of width** k of a Σ-structure $\mathcal{A}$ is a Σ_k-structure $\mathcal{T} \in \mathrm{TD}_k$ (resp. $\mathcal{T} \in \mathrm{PD}_k$) such that $\mathcal{E}\mathrm{xt}(\mathcal{T}) \simeq \mathcal{A}$.

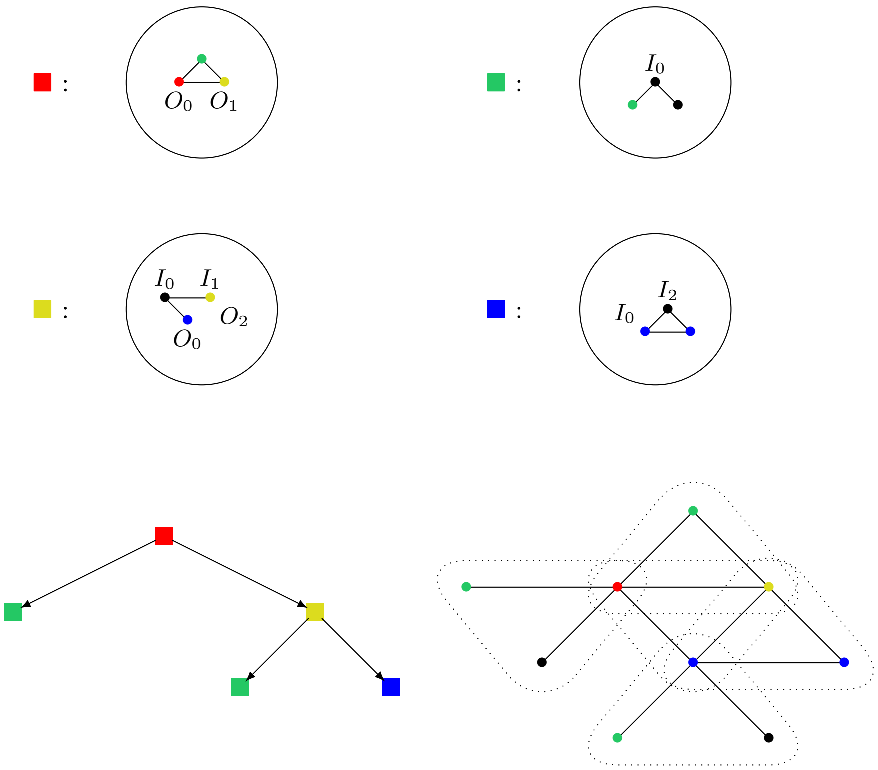

Figure 2.2: On colored graphs, $\mathcal{T} \in \mathrm{TD}_2$ (on the left) and $\mathcal{E}\mathrm{xt}(\mathcal{T})$ (on the right)

Let $\mathcal{TW}_k$ (resp. $\mathcal{PW}_k$) denote the class of all Σ-structures of **treewidth** (resp. **pathwidth**) **at most** k, that is structures that admit a tree-decomposition (resp. path-decomposition) of width at most k.

We say that a class $\mathcal{C}$ has **treewidth** (resp. **pathwidth**) **at most** k if $\mathcal{C} \subseteq \mathcal{TW}_k$ (resp. $\mathcal{C} \subseteq \mathcal{PW}_k$), and that $\mathcal{C}$ is a class of **bounded treewidth** (resp. **bounded pathwidth**) if it has treewidth (resp. pathwidth) at most k for some $k \in \mathbb{N}$.

2.5.3 Domino decompositions

In the classical definition of a tree-decomposition (see Section 2.5.1) of a graph, there is no restriction to the number of bags, and their distance, containing a given element of the graph. This carries over to the definition we gave of a tree-decomposition of a Σ-structure, in that the equivalence classes for the relation $\sim$ may be composed of elements coming from k-bags that are arbitrarily far apart in the tree-decomposition.

For us to be able to combine FO and tree and path-decompositions in an attempt to lift the collapse of Proposition 2.4.8 to classes of structures of bounded treewidth or pathwidth, it will be necessary to consider particular decompositions. Indeed, if elements appear in bags that are far apart in the decomposition, then it becomes impossible to interpret in FO the structure in its decomposition.

To see that, consider the following example on graphs for $k = 1$. Let's define the path-decomposition $\mathcal{T}_n$ as in Figure 2.3, where n is the number of occurrences of the middle bags.

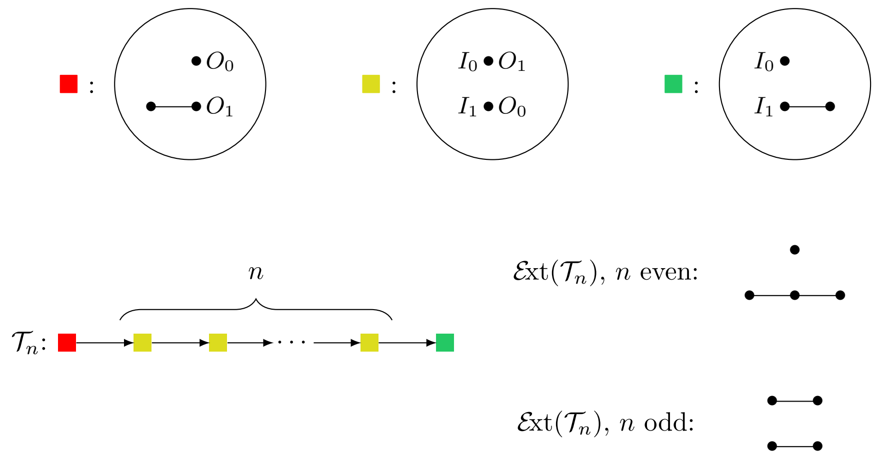

Figure 2.3: A path-decomposition (on the bottom left) $\mathcal{T}_n \in \mathrm{PD}_1$, and $\mathcal{E}\mathrm{xt}(\mathcal{T}_n)$ (on the bottom right), depending on the parity of n

FO is unable to count modulo two, hence

$$\forall m \in \mathbb{N},\ \exists n \in \mathbb{N}, \quad \mathcal{T}_{2n} \equiv_m^{\mathrm{FO}} \mathcal{T}_{2n+1}\,,$$

yet

$$\mathcal{E}\mathrm{xt}(\mathcal{T}_{2n}) \not\equiv_2^{\mathrm{FO}} \mathcal{E}\mathrm{xt}(\mathcal{T}_{2n+1})\,.$$

This shows that there is no hope to use the bi-FO-interpretation method from Proposition 2.3.6 or any similar method to lift the collapse of $<$-inv FO to FO on trees (Proposition 2.4.8) to structures of bounded treewidth, as long as we allow such decompositions.

To avoid this issue, we consider a more restrictive notion of decomposition, based on the definition introduced independently by Ding and Oporowski [13] and by Bodlaender and Engelfriet [8]. Let's first give the definition from [8].

A tree or path-decomposition (in the classical sense) $(\mathcal{T}, \beta)$ of a graph $\mathcal{G}$ is said to be domino if for every $x \in G$, there are at most two $t \in T$ such that x belongs to $\beta(t)$.

If a (simple and loopless) graph admits a domino tree-decomposition of width k, then its degree is necessarily bounded by $2k$, as any element can be linked only to an element appearing in the same bag. The converse question of finding a domino tree-decomposition of small width for any graph of bounded degree and bounded treewidth was answered in [13] and [8]. Bodlaender later improved the bound in [7]:

Proposition 2.5.5. *Let $k \in \mathbb{N}$ and $d \in \mathbb{N}$.*

Any graph of treewidth at most k and degree at most d admits a domino tree-decomposition of width at most $(9k + 7)d(d + 1) - 1$.

Let's show that Proposition 2.5.5 doesn't have an equivalent for domino path-decomposition. For that, we exhibit a family of structures (namely, graphs) $(\mathcal{A}_n)_{n>1}$ of degree 3 and pathwidth 2 such that for every $n \in \mathbb{N}$, $\mathcal{A}_n$ admits no domino path-decomposition of width at most n.

For $n > 1$, let $\mathcal{A}_n$ be the structure from Figure 2.4 with

$$m := \left\lfloor \frac{n^2(n+1)}{(n^2-2n-1)} \right\rfloor + 1$$

and where every dashed path has length n^2.

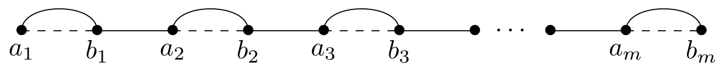

Figure 2.4: The graph $\mathcal{A}_n$: all the plain arcs are edges, and dashed arcs are paths of length n^2.

First, note that the family $(\mathcal{A}_n)_n$ is indeed of pathwidth 2: going through these graphs from left to right (and adding b_i to the bags from the moment we visit a_i until we reach b_i along the dashed path) yields a simple path-decomposition of width 2.

Now, let's show that $\mathcal{A}_n$ doesn't admit any domino path-decomposition of width at most n.

We will need the observation that if x and y are two elements of a graph $\mathcal{G}$, then in any domino decomposition of $\mathcal{G}$, a bag containing x and a bag containing y cannot be at distance more than $\text{dist}_{\mathcal{G}}(x,y)+1$. The proof of this fact is a straightforward induction on $\text{dist}_{\mathcal{G}}(x,y)$.

Suppose that there exists a domino path-decomposition $\mathcal{T}_n$ of width at most n of $\mathcal{A}_n$. The diameter of $\mathcal{A}_n$ is less than n^2+2m-2, hence according to the previous remark, the length of $\mathcal{T}_n$ is at most n^2+2m-1.

There are at most n^2+2m bags in $\mathcal{T}_n$, and each one contains at most $n+1$ elements. Yet, the size of A_n is $m(n^2+1)$; this is absurd, since we chose m to ensure $m(n^2+1) > (n+1)(n^2+2m)$.

This concludes the proof that the family $(\mathcal{A}_n)_{n>1}$ doesn't admit domino path-decompositions of any fixed width.

2.5.4 Decompositions of bounded diameter

Let's now extend the definition of domino decompositions to the general setting of Σ-structures. Our motivation is to be able to interpret in FO a structure in its decompositions, i.e. make sure that the MSO-transduction that defines a structure in any of its decompositions is already in FO. For that, we may weaken the condition that an element can appear only in two (or a bounded number of bags). Indeed, as long as all the bags containing a given element are in a bounded radius, FO will be able to recreate its neighborhood in the structure; it is not necessary for their number to be bounded.

The **diameter** of a tree-decomposition $\mathcal{T} \in \mathcal{TW}_k$ is defined as the maximum over the equivalence classes c of $\sim$ (recall the definition of $\sim$ from Section 2.5.2) of the maximal $\text{dist}_{\mathcal{T}}(t,u)$, where $t,u \in T$ each contain an element of c. In other word, the diameter of $\mathcal{T}$ is the maximal distance between two bags containing elements that are merged in $\mathcal{E}\text{xt}(\mathcal{T})$.

Let TD_k^δ (resp. PD_k^δ) denote the class of tree-decompositions (resp. path-decompositions) of width at most k and of diameter at most δ, and $\mathcal{TW}_k^\delta$ (resp. $\mathcal{PW}_k^\delta$) denote the class of Σ-structures admitting a decomposition in TD_k^δ (resp. PD_k^δ).

Note 2.5.6. *In the case of graphs, TD_k^1 corresponds to all the domino tree-decompositions of width k.*

However, as soon as $\delta \geq 2$, $\mathcal{TW}_k^\delta$ contains graphs of arbitrarily large degree. For instance, $\mathcal{TW}_1^2$ contains all the star graphs.

It is straightforward to construct, for any $k, \delta \in \mathbb{N}$, an FO-interpretation $\mathcal{I}_k^\delta$ such that for every $\mathcal{T} \in \mathrm{TD}_k^\delta$, $\mathcal{I}_k^\delta(\mathcal{T}) \simeq \mathcal{E}\mathrm{xt}(\mathcal{T})$.

To be able to use Proposition 2.3.6, one would need the converse, namely that it is possible to interpret in FO a tree-decomposition of bounded diameter in structures of $\mathcal{TW}_k^\delta$. In other words, one would need FO-interpretations $\mathcal{J}_k^\delta$ such that, for every $\mathcal{A} \in \mathcal{TW}_k^\delta$, $\mathcal{E}\mathrm{xt}(\mathcal{J}_k^\delta(\mathcal{A})) \simeq \mathcal{A}$.

We show in Section 2.5.6 that such FO-interpretations do not exist. Before that, we show in Section 2.5.5 that it isn't possible either to interpret in FO a path-decomposition of bounded diameter in structures of $\mathcal{PW}_k^\delta$. We actually prove stronger results, that deny the hope to use the more general method from Note 2.3.7 as well.

2.5.5 Interpretation of path-decompositions

Let's prove that it isn't possible to interpret in FO a path-decomposition of bounded diameter in structures of $\mathcal{PW}_k^\delta$. We in fact prove a stronger result, which entails this one.

Let's consider the vocabulary of colored graphs $\Sigma := \{E, P_0, P_1\}$ where E is binary and P_0, P_1 are unary.

Let $k, \delta \in \mathbb{N}$. We exhibit two families $(\mathcal{G}_\beta)_{\beta\in\mathbb{N}}$ and $(\mathcal{H}_\beta)_{\beta\in\mathbb{N}}$ of Σ-structures of $\mathcal{PW}_2^2$ such that

- $\forall \beta \in \mathbb{N}, \quad \mathcal{G}_\beta \equiv_\beta^{\mathrm{FO}} \mathcal{H}_\beta$
- $\exists \alpha \in \mathbb{N}, \forall \beta \in \mathbb{N}$, for all decompositions $\mathcal{D}(\mathcal{G}_\beta), \mathcal{D}(\mathcal{H}_\beta) \in \mathrm{PD}_k^\delta$,
$$\mathcal{D}(\mathcal{G}_\beta) \not\equiv_\alpha^{\mathrm{FO}} \mathcal{D}(\mathcal{H}_\beta)\,.$$

In other words, we show that no matter how large, both in term of width and diameter, we allow our decompositions to be, there are similar structures of $\mathcal{PW}_2^2$ that do not have similar path-decompositions of this width and diameter. In particular, this entails that there is no hope to find an FO-interpretation of a path-decomposition in structures of $\mathcal{PW}_2^2$, even by blowing the width and diameter.

Let $k, \delta \in \mathbb{N}$. We set $\alpha := \delta(n+1)$, where the value of n will be apparent later on, and depends only on k. Let $\beta \in \mathbb{N}$.

The structures $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$ will be based on a series of gadgets. The first one, $\mathcal{A}_n^p$, is defined in Figure 2.5. The value of p will be specified later on.

For any integers $n_1, n_2 \leq n$, we define $\mathcal{A}_n^p[n_1, n_2]$ as $\mathcal{A}_n^p$, where the path (of length $pn-1$, hence having pn nodes) from a_0 to a_1 is colored with P_0, P_1 as

$$(0^{n-n_1}1^{n_1})^p$$

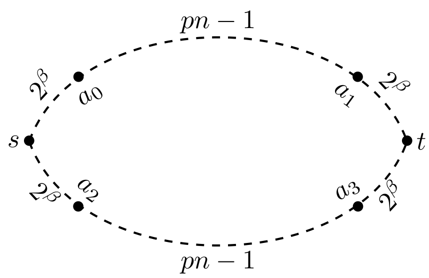

Figure 2.5: The gadget $\mathcal{A}_n^p$, where the lengths of the different unoriented paths depend on $\beta, p, n \in \mathbb{N}$.

and the path from a_2 to a_3 is colored with P_0, P_1 as

$$(0^{n-n_2}1^{n_2})^p\,.$$

Σ being fixed before k and δ are known, these colorings are a way to encode a number of colors which can depend on k and δ. The integer α, which can depend on k and δ, will help us decode these colorings.

Note 2.5.7. *If x and y are two elements of a structure $\mathcal{G}$, then in any decomposition of $\mathcal{G}$ of diameter at most δ, a bag containing x and a bag containing y must at distance at most $\delta(dist_{\mathcal{G}}(x,y)+1)$ from each other.*

The proof of this fact is a straightforward induction on $dist_{\mathcal{G}}(x,y)$.

Consider a path-decomposition $\mathcal{P} \in \mathrm{PD}_k^\delta$ of $\mathcal{A}_n^p[n_1,n_2]$.

Suppose that any two bags containing respectively s and t are at distance at least

$$2\delta(2^\beta+1)+1$$

from one another in $\mathcal{P}$.

Then there must exist a bag containing some node of the path $[a_0,a_1]$ as well as some node of the path $[a_2,a_3]$. Indeed, according to Note 2.5.7 any bag containing a_0 must be at distance at most $\delta(2^\beta+1)$ from any bag containing s and similarly for a_2 and s, a_1 and t and a_3 and t. By assumption, there must exists at least one bag in $\mathcal{P}$ that separates all the bags containing a_0 or a_2 from all the bags containing a_1 or a_3. Such a bag must contain both a node of $[a_0,a_1]$ and a node of $[a_2,a_3]$.

In that case, $\mathcal{P}$ satisfies the property $P_n(n_1,n_2)$: "there exists a bag containing both a node that is part of a path $0^{n-n_1}1^{n_1}$, and a node that is part of a path $0^{n-n_2}1^{n_2}$".

Note that any path-decomposition $\mathcal{P}' \in \mathrm{PD}_k^\delta$ such that $\mathcal{P}' \equiv_\alpha^{\mathrm{FO}} \mathcal{P}$ must also satisfy $P_n(n_1,n_2)$, by choice of α.

We now define $\mathcal{A}_n^p[n_1,n_2]^m$ as the concatenation of m copies of $\mathcal{A}_n^p[n_1,n_2]$, as illustrated in Figure 2.6.

Now, consider a path-decomposition $\mathcal{P} \in \mathrm{PD}_k^\delta$ of $\mathcal{A}_n^p[n_1,n_2]^m$.

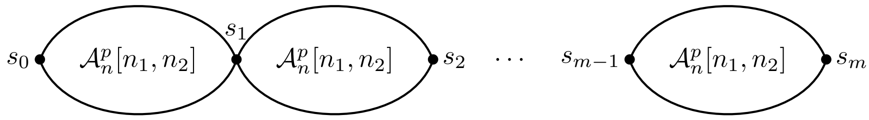

Figure 2.6: The gadget $\mathcal{A}_n^p[n_1, n_2]^m$.

Suppose that for every $0 \leq i < m$, there exist bags containing respectively s_i and s_{i+1} that are at distance at most

$$2\delta(2^\beta + 1)$$

from one another. Then the length of $\mathcal{P}$ cannot be too large. Indeed, any element of $\mathcal{A}_n^p[n_1, n_2]^m$ is at distance at most

$$2^\beta + \frac{pn-1}{2}$$

from the nearest s_i. This means, according to Note 2.5.7, that the length of $\mathcal{P}$ can be bounded by

$$\delta(2^\beta + \frac{pn-1}{2}) + m \cdot (2\delta(2^\beta + 1) + \delta) + \delta(2^\beta + \frac{pn-1}{2})$$

(joining a bag containing the nearest s_i, then jump at most m times to bags containing other s_j until we reach the closest to the destination point, then reach the destination).

This expression can be coarsely bounded by

$$\delta(2^{\beta+4}m + pn)\,.$$

This means that the size of $\mathcal{A}_n^p[n_1, n_2]^m$ is at most

$$(k+1)[\delta(2^{\beta+4}m + pn) + 1]\,. \tag{2.6}$$

However, $\mathcal{A}_n^p$ has size

$$2^{\beta+2} + 2pn - 2\,,$$

hence $\mathcal{A}_n^p[n_1, n_2]^m$ has size

$$m(2^{\beta+2} + 2pn - 2) - (m-1)\,,$$

which is greater than

$$m(2^{\beta+2} + 2pn - 3)\,. \tag{2.7}$$

As in the example from Section 2.5.3, p and m can be chosen such that the expression (2.7) is bigger than (2.6), which is absurd.

Indeed, showing that (2.7) is bigger that (2.6) amounts to finding p and m such that

$$m[2^{\beta+2} + 2pn - 3 - (k+1)\delta 2^{\beta+4}] > (k+1)(\delta pn + 1)\,.$$

This can be done by first choosing p so that

$$2^{\beta+2} + 2pn - 3 - (k+1)\delta 2^{\beta+4} > 0\,,$$

and choosing m in consequence.

Hence, there must exist $0 \le i < m$ such that no pair of bags containing respectively s_i and s_{i+1} at distance at most $\delta \cdot 2^{\beta+1}$ from each other. As we've seen, this means that $\mathcal{P}$ satisfies $P_n(n_1, n_2)$.

We now have constructed a gadget $\mathcal{A}^p_n[n_1, n_2]^m$ which is such that any path-decomposition $\mathcal{P} \in \mathrm{PD}^\delta_k$ of $\mathcal{A}^p_n[n_1, n_2]^m$ satisfies $P_n(n_1, n_2)$.

Since $P_n(n_1, n_2)$ is a property that is preserved by $\equiv^{\mathrm{FO}}_\alpha$, we will use it as a lever to prove that there is no path-decompositions in PD^δ_k of $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$ that are FO-similar at depth α.

For that, we define $\mathcal{G}_\beta$ as in Figure 2.7. The value of n and l will be fixed in the remainder of the proof. Recall that while l will depend on β, n must not, since α is a function of n.

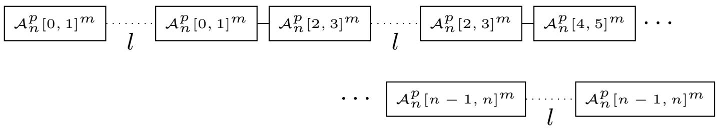

Figure 2.7: The structure $\mathcal{G}_\beta$, where p, m are chosen as above, and n is an odd integer to be fixed. The dotted edges are unoriented paths of length l, which will also be fixed later on.

Using the same values for n and l, $\mathcal{H}_\beta$ is defined as in Figure 2.8.

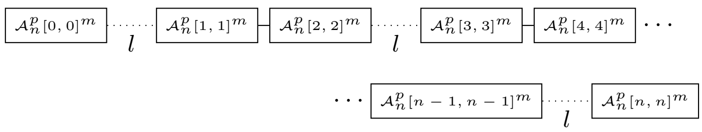

Figure 2.8: The structure $\mathcal{H}_\beta$.

One can easily see that $\mathcal{G}_\beta \equiv^{\mathrm{FO}}_\beta \mathcal{H}_\beta$. This result from the observation that, for every $0 \le i \le \frac{n-1}{2}$, two copies of $\mathcal{A}^p_n[2i, 2i+1]^m$ are FO-similar at depth β to the union of $\mathcal{A}^p_n[2i, 2i]^m$ and $\mathcal{A}^p_n[2i+1, 2i+1]^m$. Indeed, the paths of length 2^β in all the $\mathcal{A}^p_n[n_1, n_2]$ prevent the Spoiler in the β-round Ehrenfeucht-Fraïssé game from spotting which integers n_1, n_2 appear in the same $\mathcal{A}^p_n[n_1, n_2]$.

Furthermore, $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$ belong to $\mathcal{PW}^2_2$. Indeed, they each admit a path-decomposition of width 2 that goes from their left to their right, which moves in each $\mathcal{A}^p_n[n_1, n_2]$ one step in the top path, then one step in the bottom path.

It remains to prove that $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$ do not admit decompositions in PD^δ_k that are FO-similar at depth α.

Suppose that $\mathcal{D}(\mathcal{G}_\beta), \mathcal{D}(\mathcal{H}_\beta) \in \mathrm{PD}^\delta_k$ are respective decompositions of $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$ such that

$$\mathcal{D}(\mathcal{G}_\beta) \equiv^{\mathrm{FO}}_\alpha \mathcal{D}(\mathcal{H}_\beta)\,. \tag{2.8}$$

As we've seen above, $\mathcal{D}(\mathcal{G}_\beta)$ must satisfy all the properties

$$P_n(0,1), P_n(2,3), \cdots, P_n(n-1,n)\,.$$

By (2.8), $\mathcal{D}(\mathcal{H}_\beta)$ must satisfy them too. By construction of $\mathcal{H}_\beta$, this means that for every $0 \leq i \leq \frac{n-1}{2}$, there exists a bag of $\mathcal{D}(\mathcal{H}_\beta)$ containing both a node of $\mathcal{A}_n^p[2i,2i]^m$ and a node of $\mathcal{A}_n^p[2i+1,2i+1]^m$. This prevents $\mathcal{D}(\mathcal{H}_\beta)$ from being to long.

More precisely, since any $\mathcal{A}_n^p[n_1,n_2]^m$ has diameter bounded by

$$m(pn-1+2^{\beta+1})\,.$$

Note 2.5.7 entails that any of its decompositions in PD_k^δ has length at most

$$\delta[m(pn-1+2^{\beta+1})+1]\,.$$

With the requirement that for every i, $\mathcal{A}_n^p[2i,2i]^m$ and $\mathcal{A}_n^p[2i+1,2i+1]^m$ overlap in $\mathcal{D}(\mathcal{H}_\beta)$, the length of $\mathcal{D}(\mathcal{H}_\beta)$ is at most

$$\frac{l}{2}+n\delta[m(pn-1+2^{\beta+1})+2]+\frac{l}{2}\,,$$

which can be bounded by

$$n\delta m(pn+2^{\beta+1})+l-1\,.$$

This implies that the size of $\mathcal{H}_\beta$ cannot exceed

$$(k+1)[n\delta m(pn+2^{\beta+1})+l]\,. \tag{2.9}$$

However, recall from (2.7) that any $\mathcal{A}_n^p[n_1,n_2]^m$ has size at least

$$m(2^{\beta+2}+2pn-3)\,,$$

hence the size of $\mathcal{H}_\beta$ is at least

$$(n+1)m(2^{\beta+2}+2pn-3)+n(l-1)\,. \tag{2.10}$$

The right choice of n and l make (2.10) bigger than (2.9), which is absurd.

Indeed, (2.10) is bigger that (2.9) iff

$$l[n-(k+1)] > (k+1)[n\delta m(pn+2^{\beta+1})+l]-(n+1)m(2^{\beta+2}+2pn-3)+n\,.$$

Choosing $n := k+2$ allows us to set l so that this inequality holds.

It follows that $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$ do not have decompositions in PD_k^δ that are FO-similar at depth α, thus concluding the proof.

2.5.6 Interpretation of tree-decompositions

We have seen proved in Section 2.5.5 that path-decompositions of bounded diameter are not FO-interpretable in structures of $\mathcal{PW}_k^\delta$, even when increasing the width and the diameter.

Let's now show an equivalent result for tree-decomposisions. For that, we proceed in a similar way.

Let's consider the vocabulary of colored graphs $\Sigma := \{E, P_0, P_1\}$ where E is binary and P_0, P_1 are unary.

Let $k, \delta \in \mathbb{N}$. We exhibit two families $(\mathcal{G}_\beta)_{\beta\in\mathbb{N}}$ and $(\mathcal{H}_\beta)_{\beta\in\mathbb{N}}$ of Σ-structures of $\mathcal{TW}_2$ and of degree 5 such that

- $\forall \beta \in \mathbb{N}, \quad \mathcal{G}_\beta \equiv^{\mathrm{FO}}_\beta \mathcal{H}_\beta$
- $\exists \alpha \in \mathbb{N}, \forall \beta \in \mathbb{N}$, for all decompositions $\mathcal{D}(\mathcal{G}_\beta), \mathcal{D}(\mathcal{H}_\beta) \in \mathrm{TD}^\delta_k$,
$$\mathcal{D}(\mathcal{G}_\beta) \not\equiv^{\mathrm{FO}}_\alpha \mathcal{D}(\mathcal{H}_\beta)\,.$$

We introduce in Figure 2.9 the gadget $\mathcal{L}^l_p$, which is composed of two complete binary trees of height p, whose leaves are pairwise linked by a path of length l.

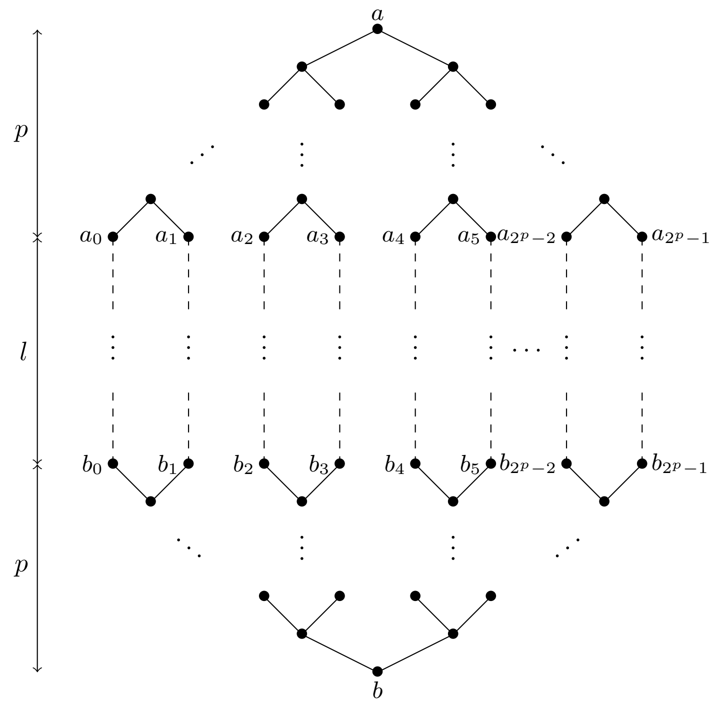

Figure 2.9: The gadget $\mathcal{L}^l_p$, where l is the length of the dotted paths. All the edges are undirected. The nodes a and b are the **sources** of $\mathcal{L}^l_p$.

The interest of $\mathcal{L}^l_p$ resides in the fact that, provided that p is large enough with respect to the width of a tree-decomposition, the sources (that is, the elements a and b) cannot be far apart in this decomposition, while they can be made arbitrarily distant, by choice of l, in the original graph.

More precisely, let $\mathcal{T}$ be a tree-decomposition of $\mathcal{L}^l_p$ in TD^δ_k, where
$$p = \lceil \log(k+2) \rceil\,.$$
We claim that in $\mathcal{T}$, any bag containing a and any bag containing b are at distance at most $d := 2\delta(p+2) + 2\delta$ from one another.

Suppose otherwise, and consider the bags t_a and $t_b \in T$ which minimize $\mathrm{dist}_\mathcal{T}(t_a, t_b)$ among bags containing respectively a and b. Since $\mathcal{T}$ has diameter δ, we must have that
$$\mathrm{dist}_\mathcal{T}(t_a, t_b) \geq 2\delta(p+1) + 1\,.$$

Let $t \in T$ be a bag on the path between t_a and t_b at distance at least $\delta(p+1)+1$ from each other. By virtue of $\mathcal{T}$ being a tree and in view of Note 2.5.7, t disconnects all the bags containing the a_i from all the bags containing the b_i.

This means that all the 2^p disjoint paths from a_i to b_i must intersect t, which is absurd since $2^p \geq k+2$.

This proves that any bag containing a and any bag containing b are at distance at most $2\delta(p+2)$ from one another in $\mathcal{T}$.

Once k and δ are given, we set $p := \lceil \log(k+2) \rceil$.

Our construction will depend on an integer n whose value depends only on k and δ, and will be set later on. The integer α will also be chosen later on.

Let $\beta \in \mathbb{N}$, and set $l := 2^\beta$.

Let's now construct the structures $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$. Both of them will amount to a concatenation of many instances of $\mathcal{L}_p^l$ (by concatenation of two $\mathcal{L}_p^l$, we mean the disjoint union of those structures, where we merge one of their sources). On top of that, all the sources will have a label. As in Section 2.5.6, labels are encoded with the unary relations P_0, P_1 on a path of length at most n.

Let's deal with $\mathcal{G}_\beta$ first. We start by considering some nodes $(s_w)_{w \in \{0,1\}^{\leq n}}$, i.e. one element for each one of the $2^{n+1}-1$ sequences of bits of length at most n. All those s_w belong to G_β.

We now add the aforementionned labels: for every $w \in \{0,1\}^{\leq n}$ we attach to s_w a path of length $|w|$, and color it with P_0 and P_1 in order to code w.

On top of that, for every $w \in \{0,1\}^{\leq n-1}$, we link s_w to s_{w0} with an copy of $\mathcal{L}_p^l$; in other words, s_w and s_{w0} are the sources of this gadget. All those copies are disjoint. Note that up until now, each s_{w1} with $|w| = n-1$ is alone with its label in its connected component of $\mathcal{G}_\beta$. This stage of the construction is depicted in Figure 2.10.

As of now, $\mathcal{G}_\beta$ is a union of concatenations of copies of $\mathcal{L}_p^l$, together with some isolated nodes. However, we will specify later how to agglomerate all those connected components so that $\mathcal{G}_\beta$ becomes a single concatenantion of copies of $\mathcal{L}_p^l$. Basically, we will add copies of $\mathcal{L}_p^l$ between the connected components to group them into a single sequence of $\mathcal{L}_p^l$, as illustrated in Figure 2.11.

Similarly, $\mathcal{H}_\beta$ is obtained by linking each s_w to the corresponding s_{w1} with a copy of $\mathcal{L}_p^l$. In the end $\mathcal{H}_\beta$ will also be a single concatenation of copies of $\mathcal{L}_p^l$.

The idea behind those constructions is that, in the complete binary tree where the nodes are the words of $\{0,1\}^{\leq n}$, each edge (w, w') is taken into account as a copy of $\mathcal{L}_p^l$ either in G_β (if $w' = w0$) or in H_β (if $w' = w1$).

The labels will help us identify s_w in $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$, with the help of the FO-similarity at depth α of their tree-decompositions.

We can now start establishing some results. Suppose that $\mathcal{D}(\mathcal{G}_\beta)$ and $\mathcal{D}(\mathcal{H}_\beta)$ are respective tree-decompositions of $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$ in TD_k^δ such that

$$\mathcal{D}(\mathcal{G}_\beta) \equiv_\alpha^{\mathrm{FO}} \mathcal{D}(\mathcal{H}_\beta)\,. \tag{2.11}$$

We've seen that in $\mathcal{D}(\mathcal{G}_\beta)$, for any $w \in \{0,1\}^{\leq n-1}$, any bags containing respectively s_w and s_{w0} are at distance at most $d := 2\delta(p+2)$ from one another.

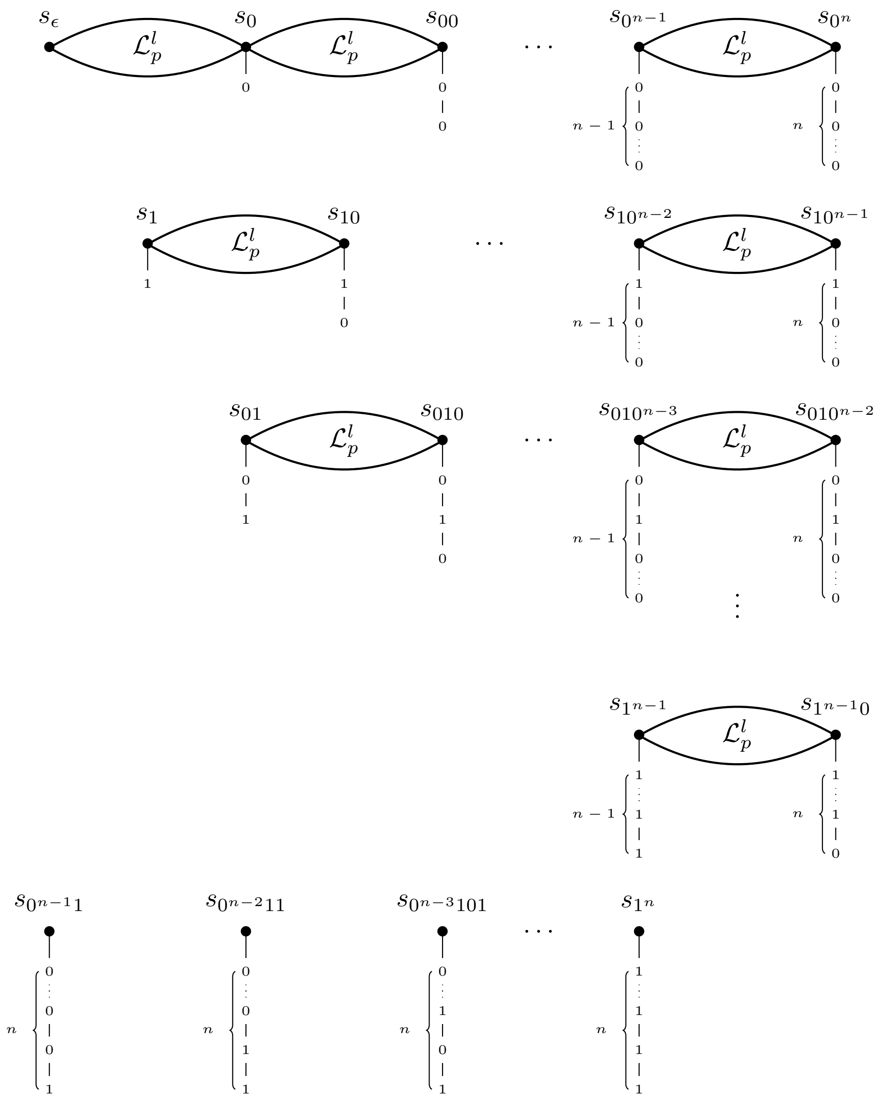

Figure 2.10: $\mathcal{G}_\beta$, at the beginning of the construction.

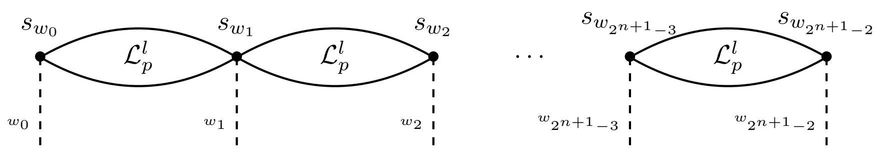

Figure 2.11: $\mathcal{G}_\beta$, at the end of the construction, for some ordering $w_0, \cdots, w_{2^{n+1}-2}$ of the words of $\{0,1\}^{\leq n}$ to be fixed later on. It extends the construction from Figure 2.10.

Let's choose a large enough α with respect to n, d and δ so that the property

"any bags containing nodes with labels w and $w0$ are at distance at most d"

is expressible as an FO-sentence of quantifier rank α. By (2.11), any bags of $\mathcal{D}(\mathcal{H}_\beta)$ containing respectively s_w and s_{w0} must be at distance at most d.

Similarly, both in $\mathcal{D}(\mathcal{G}_\beta)$ and $\mathcal{D}(\mathcal{H}_\beta)$, any bags containing respectively s_w and s_{w1} must be at distance at most $d := 2\delta(p+2)$ from one another.

In the end, for $w, w' \in \{0,1\}^{\leq n}$, in $\mathcal{D}(\mathcal{G}_\beta)$ as well as in $\mathcal{D}(\mathcal{H}_\beta)$, any bags containing respectively s_w and $s_{w'}$ must be at distance at most $2nd$ from one another. This is because the complete binary tree of height n as diameter $2n$.

Thus in both decompositions, assuming (2.11), the $2^{n+1}-1$ nodes s_w are contained in a subtree whose diameter is linear in n. This compactness is key to this counter-example, and is the reason for this tree-like construction.

The next step is to show that we can identify the parts of $\mathcal{D}(\mathcal{G}_\beta)$ and $\mathcal{D}(\mathcal{H}_\beta)$ which contain the elements s_w.

More precisely, we show that there exist subtrees $\mathcal{S}_G$ and $\mathcal{S}_H$ of $\mathcal{D}(\mathcal{G}_\beta)$ and $\mathcal{D}(\mathcal{H}_\beta)$ such that

- $\mathcal{S}_G \simeq \mathcal{S}_H$
- the diameter of $\mathcal{S}_G, \mathcal{S}_H$ is at most $2nd$
- $\mathcal{S}_G$ and $\mathcal{S}_H$ have degree at most $2k+3$
- every s_w belongs to at least one bag of $\mathcal{S}_G$, and one bag of $\mathcal{S}_H$.

For that, we start by considering the minimal subtree $\mathcal{S}_G$ of $\mathcal{D}(\mathcal{G}_\beta)$ which contains all the bags containing any s_w. As we've seen, this subtree has diameter at most $2nd$.

There is however no restriction on the degree of $\mathcal{S}_G$. To get the desired properties, we trim $\mathcal{S}_G$ in the following way.

While there exists at least one, pick a bag t of $\mathcal{S}_G$ of degree greater than $2k+3$. Let $\mathcal{S}_1, \cdots, \mathcal{S}_r$ be the connected components of $\mathcal{S}_G \setminus \{t\}$.

We claim that at most $2k+3$ of the $\mathcal{S}_i$ contain some s_w which does not appear in t. Recall, although at this point we have only partially constructed $\mathcal{G}_\beta$, that in the end it will be a concatenation of $\mathcal{L}_p^l$ (i.e. will consist of all the s_w, arranged in some sequence, and pairwise linked with a copy of $\mathcal{L}_p^l$).

Let $w_0, \cdots, w_{2^{n+1}-2}$ be the sequence of words of $\{0,1\}^{\leq n}$ appearing in the same order as the s_w in $\mathcal{G}_\beta$, c.f. Figure 2.11.

Suppose that there are at least $2k+4$ connected components $\mathcal{S}_i$ containing some s_w which does not appear in t.

For each such $\mathcal{S}_i$, let $m(i)$ be the maximal index such that $s_{w_{m(i)}}$ belongs to $\mathcal{S}_i$ and not to t. For at most one $\mathcal{S}_i$ we can have $m(i) = 2^{n+1}-2$. For all the others (i.e. for at least $2k+3$ of them), $s_{w_{m(i)+1}}$ does not belong to $\mathcal{S}_i$.

Given that t is a bag of size at most $k+1$, there must exist at least $k+2$ indexes i such that

- $s_{w_{m(i)}}$ belongs to $\mathcal{S}_i$ and not to t,
- $s_{w_{m(i)+1}}$ belongs neither to $\mathcal{S}_i$ neither to t.

By construction, $s_{w_{m(i)}}$ and $s_{w_{m(i)+1}}$ are linked with a copy of $\mathcal{L}_p^l$ in $\mathcal{G}_\beta$. Thus for each of these $k+2$ couples, there exists a path from s_{w_j} to $s_{w_{j+1}}$, which must intersect t. All such path being disjoint, t must intersect $k+2$ distinct paths, which is absurd.

We trim out of $\mathcal{S}_G$ all the $\mathcal{S}_i$ which contain no s_w which does not appear in t. In the new subtree, t has degree at most $2k+3$, and each s_w still belongs to at least one of its bags.

In the end, $\mathcal{S}_G$ has degree at most $2k+3$, and its diameter is a most $2nd$.

We find $\mathcal{S}_H$ using (2.11): by setting α big enough wrt. n, k and δ, the Spoiler in the α-round Ehrenfeucht-Fraïssé game between $\mathcal{D}(\mathcal{G}_\beta)$ and $\mathcal{D}(\mathcal{H}_\beta)$ can cover $\mathcal{S}_G$ (which has bounded diameter and degree) as well as enough bags of $\mathcal{D}(\mathcal{G}_\beta)$ to cover the labels of each of the s_w. The corresponding moves of the Duplicator in $\mathcal{D}(\mathcal{H}_\beta)$ yield $\mathcal{S}_H$.

In the remainder of the proof, we let $\mathcal{S} := \mathcal{S}_G \simeq \mathcal{S}_H$.

To conclude the proof, we are now going to show that these decompositions are too compact to exist.

In $\mathcal{S}$, we pick a node t. Let $\mathcal{S}_1, \cdots, \mathcal{S}_r$ be the connected components of $\mathcal{S} \setminus \{t\}$, with $r \leq 2k+3$.

Let's establish some vocabulary. We say that s_w **only occurs in** $\mathcal{S}_i$, which we denote $s_w \in^! \mathcal{S}_i$, if s_w is contained in some bag of $\mathcal{S}_i$, but not in t. Note that this implies, by nature of tree-decompositions, that s_w doesn't belong to any bag of any $\mathcal{S}_j, j \neq i$.

If $s_w \in^! \mathcal{S}_i$, we say that s_w is an $\mathcal{S}_i$**-inode**. If $s_{w'} \in^! \mathcal{S}_j$ for some $j \neq i$, we say that s_w is an $\mathcal{S}_i$**-onode**.

We say that s_w and $s_{w'}$ are **adjacent in** $\mathcal{G}_\beta$ (resp. **in** $\mathcal{H}_\beta$) if they are linked by a copy of $\mathcal{L}_p^l$ in $\mathcal{G}_\beta$ (resp. $\mathcal{H}_\beta$).

We say that they are **adjacent** if they are adjacent in $\mathcal{G}_\beta$ or in $\mathcal{H}_\beta$. For now, s_w and $s_{w'}$ are adjacent iff $w' = w0$, $w' = w1$, $w = w'0$ or $w = w'1$, but as said previously, we are going to add some copies of $\mathcal{L}_p^l$ in both structures.

We say that $\{s_w, s_{w'}\}$ is an $\mathcal{S}_i$**-bridge** if

- $s_w \in^! \mathcal{S}_i$
- $s_{w'} \in^! \mathcal{S}_j$ for some $j \neq i$
- s_w and $s_{w'}$ are adjacent.

Let $\mathcal{S}_i$ be a connected component of $\mathcal{S} \setminus \{t\}$. Then there are at most $2k+2$ $\mathcal{S}_i$-bridges. Otherwise there would exist at least $k+2$ $\mathcal{S}_i$-inode adjacent in $\mathcal{G}_\beta$ (without loss of generality) to some $\mathcal{S}_i$-onode. Thus t would intersect with at least $k+2$ disjoint paths, which is absurd.

Let $h < n$ whose value will be apparent later on. With this remark in mind, let's now show that there cannot exist i such that

$$N \leq |\{s_w \in^! \mathcal{S}_i\}| \leq (2^{n+1}-1) - (k+1) - N\,, \tag{2.12}$$

where

$$N := (2^h - 1) + (4k+4)(2^{n-h+1}-1) + 1\,.$$

This amount to saying that the s_w cannot be spread evenly across the $\mathcal{S}_1, \cdots, \mathcal{S}_r$: as soon as at least N s_w only occur in some $\mathcal{S}_i$, then most of the s_w must only occur in that $\mathcal{S}_i$. We will then see that this is absurd.

Let's now show that there doesn't exist any $\mathcal{S}_i$ satisfying (2.12). Suppose that there does exist such an $\mathcal{S}_i$.

For any w of length h, let $\mathcal{T}_w$ be the set of $s_{w'}$ such that w is a prefix of w', together with a copy of $\mathcal{L}_p^l$ joining every $s_{w'}$ and $s_{w''}$ where $w'' = w'0$ or $w'' = w'1$. Basically, $\mathcal{T}_w$ corresponds to the subtree rooted in w in the complete binary trees of universe $\{0,1\}^{\leq n}$, where the edges are replaced with copies of $\mathcal{L}_p^l$. There are 2^h such $\mathcal{T}_w$. Let's call them **h-trees**.

Note that if some h-tree $\mathcal{T}_w$ contains

- an $\mathcal{S}_i$-inode,
- an $\mathcal{S}_i$-onode,
- no s_w appearing in t,

then there must exist an $\mathcal{S}_i$-bridge included in $\mathcal{T}_w$. Since, as seen earlier, there cannot exist more than $2k+2$ $\mathcal{S}_i$-bridges and t contains at most $k+1$ s_w, this entails that of the 2^h disjoint h-trees, at most $3k+3$ can contain at the same time some $\mathcal{S}_i$-inode and some $\mathcal{S}_i$-onode.

Recall that we supposed in (2.12) that there were at least N $\mathcal{S}_i$-inodes. Since there are only $2^h - 1$ s_w not belonging to any h-tree (those for which $|w| < h$), and since each h-tree contains $2^{n-h+1} - 1$ s_w, by choice of N there must exist at least $4k+5$ h-trees containing some $\mathcal{S}_i$-inode.

Similarly, since the second inequality of (2.12) entails that there exist at least N $\mathcal{S}_i$-onodes, there must exist at least $4k+5$ h-trees contining some $\mathcal{S}_i$-onode.

We've seen that at most $3k+3$ of the h-trees can contain at the same time an $\mathcal{S}_i$-inode and an $\mathcal{S}_i$-onode. Thus, there must exist at least $k+2$ h-trees $\mathcal{T}$ containing only $\mathcal{S}_i$-inodes and s_w appearing in t; hence there exists some h-tree $\mathcal{T}$ containing only $\mathcal{S}_i$-inodes. Similarly, there must exist some h-tree $\mathcal{T}'$ containing only $\mathcal{S}_i$-onodes.

Let's now finish the construction of $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$ so that the existence of $\mathcal{T}$ and $\mathcal{T}'$ yields at least $2k+3$ $\mathcal{S}_i$-bridges. For that, we add copies of $\mathcal{L}_p^l$ between the s_w, $|w| = n$ in order to make sure that for every h-trees $\mathcal{T}_w \neq \mathcal{T}_{w'}$, there are $2k+3$ leaves of $\mathcal{T}_w$ which are adjacent to leaves of $\mathcal{T}_{w'}$. Once we've shown how to do this, we get $2k+3$ $\mathcal{S}_i$-bridges involving leaves of $\mathcal{T}$ (which are $\mathcal{S}_i$-inodes) and leaves of $\mathcal{T}'$ (which are $\mathcal{S}_i$-onodes).

Consider an h-tree $\mathcal{T}_w$. In $\mathcal{T}_w$, there are 2^{n-h-1} leaves $s_{w'}$ with w' ending with a 1. Those are isolated (not taking into account their label) in $\mathcal{G}_\beta$. Since there are $2^h - 1$ other h-trees, as long as $2^{n-h-1} \geq (2k+3)(2^h - 1)$ i.e. in particular for

$$h := \lfloor \frac{n-1-\log(2k+3)}{2} \rfloor , \tag{2.13}$$

we can arbitrarily put $2k+3$ copies of $\mathcal{L}_p^l$ between the leaves ending with 1 of every pair of h-tree. This leads to an impossibility, which comes from the assumption of the existence of an $\mathcal{S}_i$ satisfying (2.12).

At this point, $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$ are still unions of concatenations of copies of $\mathcal{L}_p^l$ (indeed, in the previous step, copies of $\mathcal{L}_p^l$ were only added between isolated s_w of $\mathcal{G}_\beta$). As promised, we now arbitrarily add copies of $\mathcal{L}_p^l$ in $\mathcal{G}_\beta$ and $\mathcal{H}_\beta$ so that both of them are a concatenation of copies of $\mathcal{L}_p^l$.

As desired, both structures have degree 5, and treewidth at most 2 (indeed, they are series-parralel graphs). Setting $l := 2^\beta$ ensures that

$$\mathcal{G}_\beta \equiv_\beta^{\mathrm{FO}} \mathcal{H}_\beta \,,$$

since the Spoiler in the β-round Ehrenfeucht-Fraïssé game has no way of determining in which order the s_w are linked in $\mathcal{G}_\beta$ and $\mathcal{H}_k$.

We've seen that for every bag $t \in S$, there are at most $2k+3$ connected components $\mathcal{S}_1, \cdots, \mathcal{S}_r$ in $\mathcal{S} \setminus \{t\}$. Furthermore, no $\mathcal{S}_i$ can satisfy (2.12).

It is not possible for all the connected components of $\mathcal{S} \setminus \{t\}$ to have less than N inodes, as long as n is large enough, since

$$(2k+3)(N-1) + (k+1) < 2^{n+1} - 1 \,.$$

Hence there must exist some $\mathcal{S}_i$ with at least N $\mathcal{S}_i$-inodes. Since $\mathcal{S}_i$ cannot satisfy (2.12), there must exist more than $(2^{n+1}-1) - (k+1) - N$ $\mathcal{S}_i$-inodes.

Each other connected component $\mathcal{S}_j, j \neq i$ must then have less than N $\mathcal{S}_j$-inode.

This unique $\mathcal{S}_i$ is called the **large connected component of** $\mathcal{S} \setminus \{t\}$.

We are now ready to conclude the proof. For that, consider Algorithm 1.

Algorithm 1

Arbitrarily pick a bag $t \in S$
while true **do**
    Print t
    $t \leftarrow$ the neighbor of t in the large connected component of $\mathcal{S} \setminus \{t\}$.
end while

Let's look at the infinite sequence output by Algorithm 1.

$\mathcal{S}$ being acyclic and finite, some sequence t_1, t_2, t_1, with $t_1, t_2 \in S$, must occur at some point in the output string.

Let $\mathcal{S}_1$ (resp. $\mathcal{S}_2$) be the connected component of $\mathcal{S} \setminus \{(t_1, t_2)\}$ containing t_1 (resp. t_2). The apparition of the sequence t_1, t_2, t_1 in the output means that

- $\mathcal{S}_2$ is the large connected component of $\mathcal{S} \setminus \{t_1\}$
- $\mathcal{S}_1$ is the large connected component of $\mathcal{S} \setminus \{t_2\}$.

In particular, this means that there exist more that $(2^{n+1}-1)-(k+1)-N$ $\mathcal{S}_1$-inodes, and more than $(2^{n+1}-1)-(k+1)-N$ $\mathcal{S}_2$-inodes. However, those sets of inodes are disjoint, and the sum of their number exceeds $2^{n+1}-1$ as long as n is chosen large enough wrt. k and δ (recall that $N = O(2^{\frac{n}{2}})$).

We have thus reached an impossibility, proving that there cannot exist tree-decompositions $\mathcal{D}(\mathcal{G}_\beta)$ and $\mathcal{D}(\mathcal{H}_\beta)$ in TD_k^δ such that

$$\mathcal{D}(\mathcal{G}_\beta) \equiv_\alpha^{\mathrm{FO}} \mathcal{D}(\mathcal{H}_\beta) \,.$$

As a corollary, we get the impossibility to FO-interpret, in structures of bounded treewidth and degree, one of their tree-decompositions of bounded diameter.

2.6 Methodology for proving a collapse

When attempting to prove that $\mathcal{L} \in \{<\text{-inv FO}, \text{Succ-inv FO}\}$ collapses to FO on a class $\mathcal{C}$ of structures, one can follow several strategies. We detail in this section three such strategies.

2.6.1 Lifting a pre-established collapse

The first idea is to find another class of structures $\mathcal{C}'$ on which it is known that $\mathcal{L} = \text{FO}$, and such that $\mathcal{C}$ is bi-FO-interpretable through $\mathcal{C}'$. Proposition 2.3.6 then allows us to lift the collapse from $\mathcal{C}'$ to $\mathcal{C}$.

This method is first used in Corollary 3.1.2, in which we lift the collapse of Succ-inv FO to FO from graph classes of bounded degree to near-uniform graph classes (see [19]), which are exaclty the classes of graphs FO-interpretable in graph classes of bounded degree. It turns out that each one of these near-uniform classes can be FO-intepreted in a class of graphs of bounded degree, with the property that the FO-interpretation can be reversed: its converse is not a proper FO-interpretation (hence we fall outside the scope of bi-FO-interpretations), but it still preserves FO-similarity, allowing the use of Note 2.3.7.

We also use this method in Corollary 4.2.1, to lift the collapse $<$-inv FO = FO on a class of graphs to a class containing directed version of those graphs. For that, it is enough to provide FO-definable encoding and decoding of the orientation in the undirected setting.

Although this seems to be the natural way to lift the collapse of $<$-inv FO to FO from trees to structures of bounded treewidth and degree, we've seen in Section 2.5.6 that this cannot be done. It might be possible, when restricting to structures of bounded pathwidth and degree, to find an FO-interpretation which defines in each such structure one of its tree-decompositions of bounded diameter. Although not excluding this possibility, the result from Section 2.5.5 leads us to believe that there is not much hope in that direction, by showing that path-decompositions of bounded diameter are not definable in this way.

2.6.2 The direct method

If the bi-FO-interpretation approach fails, one could try the direct method. For example, in the case of $<$-inv FO, it amounts to proving that there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that, for every structures $\mathcal{A}, \mathcal{A}' \in \mathcal{C}$, there exist orders $<$ on A and $<'$ on A' such that

$$\forall n \in \mathbb{N}, \quad \mathcal{A} \equiv^{\text{FO}}_{f(n)} \mathcal{A}' \quad \rightarrow \quad (\mathcal{A}, <) \equiv^{\text{FO}}_{n} (\mathcal{A}', <') \,.$$

This is indeed enough to prove that $<$-inv/$_{\mathcal{C}}$ FO = FO on $\mathcal{C}$ in virtue of Proposition 2.3.2.

Similarly, to prove that Succ-inv/$_{\mathcal{C}}$ FO = FO, one would need, given a pair of structures of $\mathcal{C}$, to construct two successors that maintain FO-similarity to some degree between those two structures. We will use this method in Chapter 3 to prove that Succ-inv FO is no more expressive than plain FO on classes of structures of bounded degree.

While it appears to be possible to construct such successors on well-behaved graphs, such a task seems ambitious in the case of orders. In fact, it is not even clear that such a construction is possible on the class of dipaths, on which $<$-inv FO is known to collapse to FO (Proposition 2.4.8). Even if such orders existed, the fact that they are hard to find already on paths suggests that there is not much hope to exhibit such orders in the case, for instance, of structures of bounded treewidth.

Note that the technique used in [14] to prove that $<$-inv FO collapses to FO when the treedepth is bounded is related to the direct method. Instead of constructing orders maintaining FO-similarity between two structures, the authors exhibit canonical orders on structures of bounded treedepth such that the equivalence classes of the ordered structures are FO-definable. Once again, it seems difficult to extend this technique to a broader setting.

2.6.3 The chaining method

As we've seen, already on paths, it appears difficult to construct orders on FO-similar structures and preserve similarity. This is mainly because the FO-similarity constraint on the initial structures enforces only local likeness (recall Section 2.4.2), while FO-similarity in the presence of an order requires a global likeness, as the diameter of the Gaifman graph of an ordered structure is 1. There is no guarantee that such orders even exist.

The techniques developed in [4] and [5] to prove that $<$-inv FO = FO on trees, although not presented through that lens in those papers, can be seen as a more flexible construction. It generalizes the previous method insofar as, instead of exhibiting orders $<$ and $<'$ respectively on $\mathcal{A}$ and $\mathcal{A}'$ that directly maintain FO-similarity, it goes through a chain of intermediate structures and orders.

More precisely, the key is to show the existence of a function $f : \mathbb{N} \to \mathbb{N}$ such that, for every $n \in \mathbb{N}$ and $\mathcal{A}, \mathcal{A}' \in \mathcal{C}$ satisfying $\mathcal{A} \equiv^{\text{FO}}_{f(n)} \mathcal{A}'$, there exist structures $(\mathcal{A}_i)_{1\leq i\leq p}$ and corresponding orders $(<_i, <'_i)_{1\leq i\leq p}$ as well as orders $<$ on A and $<'$ on A' such that

$$\begin{array}{llllllll}
(\mathcal{A},<) & \equiv^{\text{FO}}_n & (\mathcal{A}_1,<'_1) & & & & & \\
& & (\mathcal{A}_1,<_1) & \equiv^{\text{FO}}_n & (\mathcal{A}_2,<'_2) & & & \\
& & & & \ddots & & & \\
& & & & & (\mathcal{A}_p,<_p) & \equiv^{\text{FO}}_n & (\mathcal{A}',<')
\end{array}$$

This progression in the equivalences entails that all the $\mathcal{A}_i$, as well as $\mathcal{A}$ and $\mathcal{A}'$, belong to the same equivalence class for $\equiv^{<\text{-inv FO}}_n$. We can conclude with Proposition 2.3.2.

This chaining method corresponds to the *n-flip* defined by Barceló and Libkin in [3].

It is shown in [16] (for $<$-inv MSO, but the proof also holds for $<$-inv FO) that as soon as $<$-inv FO = FO on $\mathcal{C}$, such a chain exists for every pair of structures of $\mathcal{C}$ for some f. In other words, this method is complete; of course, finding the intermediate structures may not be an easy task.

In [4], each progression in the chaining corresponds to some operation on the tree. The orders are derived from the proof given in [24] of Proposition 2.4.6.

Finding orders maintaining FO-similarity is a manageable task in that context because there are few modifications between a structure of the chain and the following one.

A similar technique is used in Chapter 4, although some complications arise. Indeed, while in [4], all the intermediate structures of the chain belong to the class under consideration (namely, the class of trees), in Chapter 4, we consider structures that fall out of the class at hand.

If for a class $\mathcal{C}$ of structures one manage to prove that such a chain always exists, and that all intermediate structures belong to a class $\mathcal{C}' \supseteq \mathcal{C}$, we get not only that $<$-inv FO = FO on $\mathcal{C}$, but also that

$$<\text{-inv}/_{\mathcal{C}'}\ \text{FO} = \text{FO on } \mathcal{C}\,.$$

In other words, an FO($\Sigma \cup \{<\}$)-sentence need only be order-invariant on $\mathcal{C}'$ to have an FO-sentence equivalent on $\mathcal{C}$.

2.7 Model checking for successor-invariance

A related question to that of the expressivity of invariant logics is the matter of the complexity of the model checking for these logics, i.e. the difficulty one has to determine whether a model satisfies a given invariant sentence.

This question will not play a direct role in this thesis, but we mention the state of the art in this domain for completeness.

The **model checking problem for a logic** $\mathcal{L}$, MC($\mathcal{L}$), is given a structure $\mathcal{A}$ and an $\mathcal{L}$-sentence φ, to determine whether $\mathcal{A} \models \varphi$. There are several ways to measure the complexity of the model checking problem for a logic, depending on which of $\mathcal{A}$ and φ are treated as input or as parameter.

We speak of **combined complexity** of the model checking problem considers if both $\mathcal{A}$ and φ are part of the input. Stockmeyer [35] proved that the combined complexity of MC(FO) is PSPACE-complete. The hardness comes from the formula part: even on a model with two elements, MC(FO) is PSPACE-complete.

To see to what extent the structure weighs in the complexity of the model checking problem, Vardi [39] defined the **data complexity**, where only $\mathcal{A}$ is considered as part of the input, while φ is treated as a parameter of the problem. The data complexity of MC(FO) is in LOGSPACE (and in fact, even in AC^0).

To get a finer-grained understanding of the complexity of the model checking problem, it is useful to consider parametrized complexity. MC($\mathcal{L}$) is said to be **fixed-parameter tractable** if it is solvable in time $f(|\varphi|) \cdot |\mathcal{A}|^c$ for some constant c, where it is only asked of f to be a computable function (here, $|\varphi|$ denotes the size of φ). If $c = 1$, we say that this problem is **fixed-parameter linear**.

The fundamental result that MC(MSO) is fixed-parameter linear on classes of graphs of bounded treewidth was proved by Courcelle [11].

The classes of graphs on which the model checking for FO is fixed-parameter tractable has been widely studied. It has originally been proven by Seese [34] that MC(FO) is fixed-parameter linear on any class of bounded degree. After a series of improvements on this result, Grohe, Kreutzer and Siebertz [23] showed that MC(FO) is fixed-parameter tractable on any nowhere dense class of graph. Without getting in the technical definitions, let's just state that the notion of nowhere dense classes of graphs captures some idea of sparsity. In particular, it generalizes the notions of classes excluding a minor (hence, of classes of bounded treewidth) and of classes of bounded degree.

We have seen in Proposition 2.4.3 that Succ-inv FO is more expressive than plain FO. Whether this improvement in expressivity comes at the cost of the complexity of the model checking problem is an interesting question, which has been investigated e.g. in [17, 15, 38].

Van den Heuvel, Kreutzer, Pilipczuk, Quiroz, Rabinovich and Siebertz[38] proved that MC(Succ-inv FO) is fixed-parameter tractable on any class of bounded expansion (which is less general than the nowhere dense setting, but also includes any class of bounded degree or treewidth).

Since there is no indication that Succ-inv FO is more expressive than FO on classes of bounded expansion, this could possibly be due to a collapse of Succ-inv FO to FO on those classes. Our restult from Chapter 3 showing that Succ-inv FO collapses to FO on classes of bounded degree, together with the aforementioned result from [34], gives an alternative proof of the fact that MC(Succ-inv FO) is fixed-parameter linear when the degree is bounded.

We know of no result concerning the complexity of MC($<$-inv FO) other than the ones ensuing from bounds on $<$-inv FO, either in the form of a collapse to FO or of an inclusion in MSO.

2.8 Review of expressivity results

Before we list the expressivity results on invariant logics exposed throughout this chapter, let us mention the fundamental result proved independently by Immerman, Vardi and Livchak (see, for instance, [26]), linking order-invariance to complexity theory.

Proposition 2.8.1. $<$-inv LFP *captures* PTIME.

Here, LFP is the expansion of FO with a least fixpoint operator; see [27] for a formal definition of that logic.

Finding a logic with recursive syntax capturing PTIME is one of the long lasting quests of descriptive complexity. A profound understanding of the expressive power of order-invariance is necessary if one is to find an equivalent logic to LFP, which hopefully can be acquired through the study of weaker logics, such as $<$-inv FO and Succ-inv FO. This result, as well as the following proposition by Vardi [39], underline the importance of invariant logics.

Proposition 2.8.2. $<$-inv PFP *captures* PSPACE.

Here, PFP is the expansion of FO with a partial fixpoint operator.

Let's now list the few inclusions between invariant and plain logics that we have mentioned throughout this chapter. Concerning order and successor-invariance, we have seen that

$$\text{FO} \subsetneq \text{Succ-inv FO} \subseteq <\text{-inv FO}$$

and that

$$\text{CMSO} \subsetneq <\text{-inv MSO}\,.$$

Recall that all the known separating examples are dense. On the other hand, we have seen some upper-bounds for invariant logics on sparse structure. Namely, the collapse of $<$ -inv FO (hence, of Succ-inv FO) to FO has been proved on

- dipaths,
- unordered trees without the descendant relation,
- ranked trees without the descendant relation (as well as with the descendant relation, trivially),
- classes of bounded treedepth,

while the inclusion of $<$ -inv FO (hence, of Succ-inv FO) in MSO has been proved on

- planar graphs,
- classes of bounded treewidth,
- classes of bounded degree.

As for $<$-inv MSO, it is known to collapses to CMSO on

- planar graphs,
- classes of bounded treewidth.

This thesis contributes to this field by adding the following results:

- Succ-inv FO collapses to FO on classes of bounded degree (Theorem 3.1.1)
- $<$-inv FO collapses to FO on paths (Corollary 4.2.2)
- $<$-inv FO collapses to FO on hollow trees (Theorem 4.2.1), as defined in Section 4.1.2

Chapter 3

Successor-Invariant FO on Classes of Bounded Degree

As we've seen in Proposition 2.4.3, Succ-inv FO is in general more expressive than plain FO. However, the separating example by Rossman [33], which is to the best of our knowledge the only one, involves a dense class of structures.

This leaves open the possibility of a collapse of Succ-inv FO to FO on sparse classes of structures. In this chapter, we show that this is the case when the degree is bounded.

We start by stating the collapse result and giving an overview of the proof in Section 3.1.

We then define fractal types in Section 3.2, which is a key notion in our proof. We eventually move to the core of the proof, which is detailed in the several subparts of Section 3.3.

This chapter reformulates the content of [21].

The remainder in the division of $n \in \mathbb{N}$ by $m > 0$ is denoted $n[m]$.

3.1 Overview of the result

The main theorem of this section is the following collapse result:

Theorem 3.1.1. *Let Σ be a relational vocabulary, let $d \in \mathbb{N}$ and let $\mathcal{C}_d$ denote the class of all Σ-structures of degree at most d. Then*

$$\text{Succ-inv}/_{\mathcal{C}_d}\ \text{FO} = \text{FO} \ \textit{on}\ \mathcal{C}_d\,.$$

Let's first state a corollary of this result. Gajarský, Hliněný, Obdržálek, Lokshtanov and Ramanujan [19] characterized the graph classes which are FO-interpretable in a class of graphs of bounded degree as the near-uniform graph classes.

Let $\mathcal{D}$ be such a class of graphs, which is FO-interpretable in a class of graphs of bounded degree $\mathcal{C}'$. The construction from the aforementioned paper exhibits a class $\mathcal{C}$ of graphs of bounded degree (which may be larger that the degree of $\mathcal{C}'$) and an FO-interpretation $\mathcal{I}$ such that every $\mathcal{H} \in \mathcal{D}$ is isomorphic to $\mathcal{I}(\mathcal{G})$ for some $\mathcal{G} \in \mathcal{C}$, and such that $\mathcal{I}$ admits a converse which preserves FO-similarity (this is strongly reminiscent of a bi-FO-interpretation, although the converse in not an FO-interpretation itself); namely, for every $k \in \mathbb{N}$, there exists $k' \in \mathbb{N}$ such that for every $\mathcal{H}, \mathcal{H}' \in \mathcal{D}$ such that

$$\mathcal{H} \equiv_{k'}^{\mathrm{FO}} \mathcal{H}'$$

then there exist $\mathcal{G}, \mathcal{G}' \in \mathcal{C}$ such that

$$\mathcal{H} \simeq \mathcal{I}(\mathcal{G}) \quad \wedge \quad \mathcal{H}' \simeq \mathcal{I}(\mathcal{G}') \quad \wedge \quad \mathcal{G} \equiv_{k}^{\mathrm{FO}} \mathcal{G}' .$$

Note 2.3.7 together with Theorem 3.1.1 entail the following corollary:

Corollary 3.1.2. *Let $\mathcal{D}$ be a near-uniform class of graphs, i.e. a class of graphs* FO*-interpretable in a class of graphs of bounded degree. Then*

$$\text{Succ-inv}/_{\mathcal{D}}\ \mathrm{FO} = \mathrm{FO}\ \textit{on}\ \mathcal{D} .$$

The proof of Theorem 3.1.1 is given in Section 3.3, and constitutes the core of this chapter. We give here a sketch of this proof; this will motivate the definitions given in Section 3.2.

To prove this collapse, we follow the *direct method* as developed in Section 2.6.2. Namely, given two structures $\mathcal{G}_1$ and $\mathcal{G}_2$ of degree at most d that are FO-similar (that is, such that $\mathcal{G}_1 \equiv_n^{\mathrm{FO}} \mathcal{G}_2$ for a large enough n), our goal is to construct a successor relation S_1 on $\mathcal{G}_1$ and S_2 on $\mathcal{G}_2$ such that $(\mathcal{G}_1, S_1)$ and $(\mathcal{G}_2, S_2)$ stay FO-similar. We can then use Proposition 2.3.2 to conclude the proof.

It thus remains to construct suitable successor relations S_1 and S_2. Recall that the neighborhood type of an element is the description, up to isomorphism, of its neighborhood. When the degree is bounded, there are a finite number of neighborhood types, and if some type has many occurrences, then some occurrences must be far apart from each other. With this in mind, we separate the neighborhood types occurring in $\mathcal{G}_1$ and $\mathcal{G}_2$ into two categories:

- on the one hand, the rare types, which have few occurrences in $\mathcal{G}_1$ and $\mathcal{G}_2$ (and thus, that have the same number of occurrences in both structures, by FO-similarity)
- on the other hand, the frequent types, which have many occurrences both in $\mathcal{G}_1$ and $\mathcal{G}_2$.

In order to make the proof of FO-similarity of $(\mathcal{G}_1, S_1)$ and $(\mathcal{G}_2, S_2)$ as simple as possible, we want an element of $\mathcal{G}_1$ (and similarly for $\mathcal{G}_2$ and S_2) and its successor by S_1 to have the same neighborhood type in $\mathcal{G}_1$ as much as possible, and to be far enough in $\mathcal{G}_1$, in order for the neighborhood types occurring in $(\mathcal{G}_1, S_1)$ to be as "regular" as possible. As long as there are at least two different neighborhood types, the first constraint obviously cannot be satisfied, but we will construct S_1 as close as possible to satisfying it.

For instance, suppose that $\mathcal{G}_1$ contains three frequent neighborhood types τ_0, τ_1 and τ_2, and one rare neighborhood type χ with two occurrences. At the end of the construction, S_1 will (mostly) look like in Figure 3.1, where the relations of $\mathcal{G}_1$ have been omitted and the arrows represent S_1, which is indeed a circular successor.

Note that all the elements of type τ_1 form a segment wrt. S_1, as well as all the elements of type τ_2. τ_0, the first frequent type, has a special role in that it is used to embed all the elements of rare type (here, χ). Furthermore, and this is not apparent in the figure, two successive elements for S_1 are always distant in $\mathcal{G}_1$.

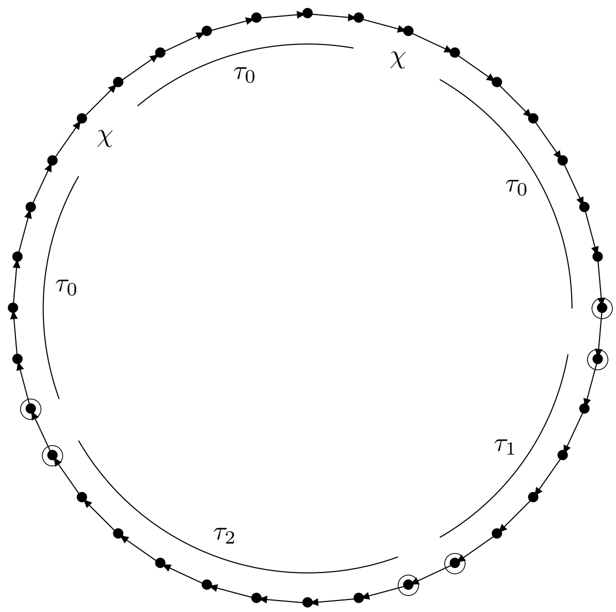

Figure 3.1: Illustration of S_1 when there are three frequent neighborhood types (τ_0, τ_1, τ_3) and one rare type (χ) in $\mathcal{G}_1$. The elements of rare type are surrounded by occurrences of the first frequent type, τ_0. Junction elements are circled.

Keeping this idea in mind, S_1 (and similarly, S_2) is constructed iteratively, by adding S-edges to the initial structures one at a time. For practical reasons, we will start the construction of S_1 around occurrences of rare types: for each element x of rare type, we find two elements of neighborhood type τ_0 that are far apart in $\mathcal{G}_1$, and far from x. Then we add two S-edges in order for those two elements to become the S_1-predecessor and the S_1-successor of x. We repeat this process for every element of rare type (and actually, for every element that belongs to the neighborhood of a rare element) until each one is protected by a ball of elements of frequent type. This is possible because there are few elements of rare type, and many elements of any frequent type; since the degree is bounded, those elements of frequent type are spread across the structure, and can be found far from the current construction.

Once this is done, we apply a similar construction around elements of frequent types that will, in the end, be the S_1-predecessor or S_1-successor of an element of another frequent type - that is, elements that will be at the border of the segments (for S_1) of a given frequent type. Such elements are circled in

Figure 3.1. We must choose only a small number of such elements (two for each frequent type, of which there are few due to the degree boundedness hypothesis), hence we can find enough far-apart elements of frequent type to embed them. Once again, degree boundedness is crucial.

After these two steps, S_1 has been constructed around all the singular points. It only remains to complete S_1 by adding edges between the remaining elements (all of which are occurrences of frequent types), in such a way that elements of a same frequent type end up forming a segment for S_1, and such that S_1 brings together elements that were far apart in the initial structure $\mathcal{G}_1$. Once again, the high number of occurrences of each frequent type allows us to do so.

Applying the same construction to $\mathcal{G}_2$, we end up with two structures $(\mathcal{G}_1, S_1)$ and $(\mathcal{G}_2, S_2)$ that cannot be distinguished by FO-formulas of small (wrt. the initial FO-similarity index between $\mathcal{G}_1$ and $\mathcal{G}_2$) quantifier rank, which concludes the proof.

We have given a global overview of the construction process of S_1; however, there are technical difficulties to take care of, which are dealt with in Section 3.3. For that, we need the definitions given in Section 3.2, which formalize the notion of regularity of a neighborhood type in $(\mathcal{G}_1, S_1)$ and $(\mathcal{G}_2, S_2)$.

3.2 Fractal types and layering

To prove Theorem 3.1.1, we will start from two structures $\mathcal{G}_1$ and $\mathcal{G}_2$ that are FO-similar, and construct successor relations S_1 and S_2 on their domains so that the structures remain FO-similar when we take into account these additional successor relations.

We want to construct S_ϵ, for $\epsilon \in \{1, 2\}$, in a way that makes $\text{tp}^k_{(\mathcal{G}_\epsilon, S_\epsilon)}(a)$ as regular as possible for every $a \in G_\epsilon$, in order to ease the proof of FO-similarity of $(\mathcal{G}_1, S_1)$ and $(\mathcal{G}_2, S_2)$.

Ideally, the S_ϵ-successors and S_ϵ-predecessors of any element should have the same k-neighborhood type in $\mathcal{G}_\epsilon$ as this element. On top of that, there should not be any overlap between the k-neighborhoods in $\mathcal{G}_\epsilon$ of elements that are brought closer by S_ϵ (this "independence" is captured by the layering property, introduced in Definition 3.2.3).

If we now try to visualize what $\text{tp}^k_{(\mathcal{G}_\epsilon, S_\epsilon)}(a)$ would look like in those perfect conditions, we realize that it reminds of a fractal (although the patterns - that is, the neighborhood types - are obviously repeated only a finite number of times).

This is why we introduce in Definition 3.2.1 the **fractal type** $[\tau]_k$.

Aside from a small number of exceptions (namely, for neighborhood types that don't occur frequently enough, and around the transitions between frequent types), every element of k-neighborhood type τ in $\mathcal{G}_\epsilon$ will have the fractal type $[\tau]_k$ in $(\mathcal{G}_\epsilon, S_\epsilon)$.

If $\mathcal{N}$ is a representative of a neighborhood type τ, $c^{\mathcal{N}}$ is called the **center of $\mathcal{N}$**. Recall from Section 2.1.1 that c is the constant symbol added to Σ when considering neighborhood types to pinpoint the central element of a neighborhood.

Definition 3.2.1 (Fractal types). *We define by induction on $k \in \mathbb{N}$, for every k-neighborhood type τ over Σ, the k-neighborhood types $[\tau]_k$, $[\tau]_k^+$ and $[\tau]_k^-$ over $\Sigma \cup \{S\}$.*

For $k = 0$, $[\tau]_0 = [\tau]_0^+ = [\tau]_0^- = \tau$ (meaning that S is interpreted as the empty relation in $[\tau]_0$, $[\tau]_0^+$ and $[\tau]_0^-$).

Starting from a representative $\mathcal{N}$ of center a of the isomorphism class τ, we construct $\mathcal{N}'$ as follows.

For every $x \in N$ at distance $d \leq k-1$ from a, let $\mathcal{M}_x^+$ and $\mathcal{M}_x^-$ be structures of respective isomorphism type $[\chi]_{k-d-1}^+$ and $[\chi]_{k-d-1}^-$, where χ is the $(k-d-1)$-neighborhood type of x in $\mathcal{N}$, and of respective center x^+ and x^-.

$\mathcal{N}'$ is defined as the disjoint union of $\mathcal{N}$ and all the $\mathcal{M}_x^+$ and the $\mathcal{M}_x^-$, for $x \neq a$, together with all the edges $S(x, x^+)$ and $S(x^-, x)$.

From there, $\mathcal{N}^+$ (resp. $\mathcal{N}^-$) is defined as the the disjoint union of $\mathcal{N}'$ and $\mathcal{M}_a^+$ (resp. $\mathcal{M}_a^-$) together with the edge $S(a, a^+)$ (resp. $S(a^-, a)$). Likewise, $\mathcal{N}^{+/-}$ is defined as the disjoint union of $\mathcal{N}'$, $\mathcal{M}_a^+$ and $\mathcal{M}_a^-$ together with the edges $S(a, a^+)$ and $S(a^-, a)$. In each case, a is taken as the center.

Now, $[\tau]_k$, $[\tau]_k^+$ and $[\tau]_k^-$ are defined respectively as the isomorphism type of $\mathcal{N}^{+/-}$, $\mathcal{N}^+$ and $\mathcal{N}^-$.

An illustration of this definition is given in Figure 3.2.

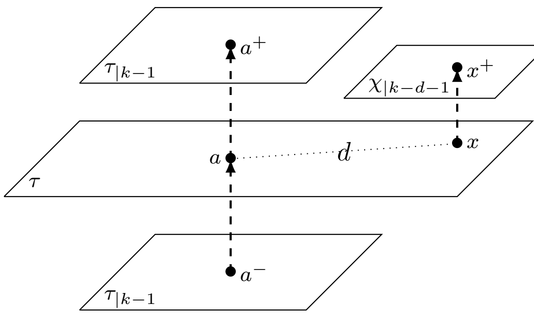

Figure 3.2: Partial representation of $\mathcal{N}^{+/-}$, of type $[\tau]_k$. Here, χ is the $(k-d)$-neighborhood type of the element x, at distance d from a in τ. The dashed arrows represent S-edges.

Definition 3.2.2 (Path and cycles). *A* ***cycle*** *of length $l \geq 2$ in the $\Sigma \cup \{S\}$-structure $\mathcal{A}$ is a sequence $(x_0, \ldots, x_{l-1})$ of distinct vertices of A such that for every $0 \leq i < l$, x_i and $x_{i+1[l]}$ appear in the same tuple of some relation of $\mathcal{A}$ (in other words, it is a cycle in $\mathcal{G}_\mathcal{A}$). If furthermore $(x_i, x_{i+1[l]}) \in S$ for every i, then we say that it is an S-cycle. If for some i, $(x_i, x_{i+1[l]}) \in S$ or $(x_{i+1[l]}, x_i) \in S$, then we say that the cycle goes through an S-edge. A* ***path*** *is defined similarly, without the requirement on x_{l-1} and x_0, and its length is $l-1$ instead of l.*

Definition 3.2.3 (Layering). *We say that an r-neighborhood $\mathcal{N}$ over $\Sigma \cup \{S, c\}$ is* ***layered*** *if it doesn't contain any cycle going through an S-edge. Every $[\tau]_r$ is obviously layered by construction.*

We say that a structure over $\Sigma \cup \{S\}$ satisfies the property (`Layer[r]`) iff all the r-neighborhoods of this structure are layered.

It turns out (`Layer[`r`]`) can be reformulated in a way that doesn't involve the r-neighborhoods of the structure:

Lemma 3.2.4. *A structure $\mathcal{G}$ over $\Sigma \cup \{S\}$ satisfies* (`Layer[r]`) *if and only if it contains no cycle of length at most $2r+1$ going through an S-edge.*

Proof. If $\mathcal{G}$ contains a cycle of length at most $2r+1$ going through an S-edge, then the r-neighborhood of any vertex of this cycle contains the whole cycle, thus (`Layer[r]`) doesn't hold in $\mathcal{G}$.

Conversely, suppose that there exists $x \in G$ such that $\mathcal{N}^r_{\mathcal{G}}(x)$ contains a cycle going through an S-edge, and let $S(y,z)$ be such an edge.

For any $u \in \mathcal{N}^r_{\mathcal{G}}(x)$, we define the **cone** C_u **at** u as the set of elements $v \in \mathcal{N}^r_{\mathcal{G}}(x)$ such that every shortest path from x to v in $\mathcal{N}^r_{\mathcal{G}}(x)$ goes through u.

There are two cases, depending on the relative position of y, z and their cones:

- If $z \notin C_y$ and $y \notin C_z$, let $p_{y \to x}$ (resp. $p_{x \to z}$) be a path of minimal length from y to x, not going through z (resp. from x to z, not going through y).

  Let X be the set of nodes appearing both in $p_{y \to x}$ and $p_{x \to z}$. X is not empty, as $x \in X$, and $y, z \notin X$. Let $v \in X$ such that $\mathrm{dist}_{\mathcal{G}}(x,v)$ is maximal among the nodes of X, and let $p_{y \to v}$ (resp. $p_{v \to z}$) be the segment of $p_{y \to x}$ (resp. of $p_{x \to z}$) from y to v (resp. from v to z).

  Then $p_{v \to z} \cdot (z,y) \cdot p_{y \to v}$ is a cycle going through an S-edge, and is of length at most $2r+1$. This is illustrated in Figure 3.3.

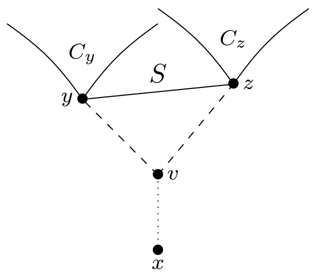

Figure 3.3: Existence of a short cycle joining y, z and v.

- Otherwise, suppose without loss of generality that $z \in C_y$. This entails that $y \notin C_z$ and $\mathrm{dist}_{\mathcal{G}}(x,z) = d+1$ where $d := \mathrm{dist}_{\mathcal{G}}(x,y)$.

  Let the initial cycle be $(z, v_1, \cdots, v_{m-1}, y)$, with the notation $v_0 = z$ and $v_m = y$.

  Let i be the minimal integer such that $v_i \notin C_z$. Let $p_{x \to v_i}$ be a shortest path from x to v_i: by definition, it doesn't intersect C_z, and has length at most r. Thus, there exists a path $p_{y \to v_i} = p_{y \to x} \cdot p_{x \to v_i}$ from y to v_i of length at most $r+d$ going only through nodes outside of C_z.

  Since $v_{i-1} \in C_z$, there exists a path $p_{v_{i-1} \to z}$ from v_{i-1} to z of length at most $r-(d+1)$ going only through nodes of C_z.

  Hence $p_{y \to v_i} \cdot (v_i, v_{i-1}) \cdot p_{v_{i-1} \to z} \cdot (z,y)$ is a cycle going though an S-edge, and its length is at most $2r+1$. This is depicted in Figure 3.4.

□

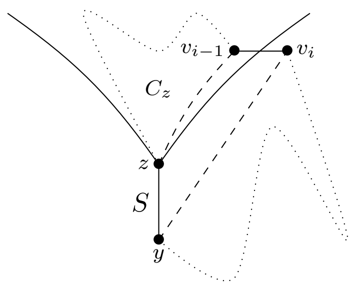

Figure 3.4: Existence of a short cycle joining y, z, v_{i-1} and v_i.

This characterization of (`Layer`[r]) allows us to state the following lemma, which is now straightforward. It gives a method to add an S-edge without breaking the property (`Layer`[r]).

Lemma 3.2.5. *Let $r \in \mathbb{N}$, and $(\mathcal{G}, S)$ be a structure satisfying (`Layer`[r]).*

If $x, y \in G$ are such that $dist_{(\mathcal{G},S)}(x, y) > 2r + 1$, then (`Layer`[$r$]) holds in $(\mathcal{G}, S \cup \{(x, y)\})$.

3.3 Proof of the collapse

We are now ready to prove Theorem 3.1.1. Recall the sketch of proof from Section 3.1. We proceed in several steps:

Section 3.3.1 details the general framework of the proof. In Section 3.3.2, we divide the neighborhood types into rare ones and frequent ones.

We then begin the construction of S_1: Section 3.3.3 is dedicated to the construction of S_1 around the occurrences in $\mathcal{G}_1$ of rare types. Then, in Section 3.3.4, we keep constructing S_1 around the occurrences (two for each neighborhood type) of frequent types that are designed to make, when the construction is complete, the S_1-junction between two frequent types.

At this point, S_1 will be fully built around the singular points of $\mathcal{G}_1$. Section 3.3.5 deals with the transfer of this partial successor relation S_1 over to $\mathcal{G}_2$: this will result in a partial S_2, built in a similar way around the singular points of $\mathcal{G}_2$.

In Section 3.3.6, S_1 and S_2 are completed independently, to cover G_1 and G_2. These expansions do not need to be coordinated, since at this point, the elements that are not already covered by S_1 and S_2 are occurrences of frequent types and their resulting types will be regular (i.e. fractal) both in $(\mathcal{G}_1, S_1)$ and $(\mathcal{G}_2, S_2)$.

We then give some simple examples in Section 3.3.7, before establishing properties of S_1 and S_2 in Section 3.3.8, and concluding the proof in Section 3.3.9.

3.3.1 General method

Let $\mathcal{C}_d$ be the class of Σ-structures of degree at most d. We show the following: for every $\alpha \in \mathbb{N}$, there exists some $f(\alpha) \in \mathbb{N}$ such that, given $\mathcal{G}_1, \mathcal{G}_2 \in \mathcal{C}_d$, if $\mathcal{G}_1 \equiv^{\mathrm{FO}}_{f(\alpha)} \mathcal{G}_2$ then $\mathcal{G}_1 \equiv^{\text{Succ-inv}/\mathcal{C}_d\ \mathrm{FO}}_{\alpha} \mathcal{G}_2$. For that, we will exhibit successor relations S_1 and S_2 such that $(\mathcal{G}_1, S_1) \equiv^{\mathrm{FO}}_{\alpha} (\mathcal{G}_2, S_2)$.

To prove that $(\mathcal{G}_1, S_1) \equiv_\alpha^{\mathrm{FO}} (\mathcal{G}_2, S_2)$, we will use Proposition 2.3.4: there exist integers r and t depending on Σ, α and d such that

$$[\![(\mathcal{G}_1, S_1)]\!]_r =^t [\![(\mathcal{G}_2, S_2)]\!]_r \quad \rightarrow \quad (\mathcal{G}_1, S_1) \equiv_\alpha^{\mathrm{FO}} (\mathcal{G}_2, S_2)\,.$$

We will construct S_1 and S_2 iteratively in a way that ensures, at each step, that the property `(Layer[`r`])` holds in $(\mathcal{G}_1, S_1)$ and in $(\mathcal{G}_2, S_2)$. `(Layer[`r`])` is obviously satisfied in $(\mathcal{G}_1, \emptyset)$. Each time we add an S_1-edge or an S_2-edge, we will make sure that we are in the right conditions to call upon Lemma 3.2.5, so that `(Layer[`r`])` is preserved.

Note 3.3.1. *Note that the size of any r-neighborhood of degree at most d is bounded by a function N of d and r; namely by $N(d, r) := d \cdot \frac{(d-1)^r - 1}{d-2} + 1$ if $d \neq 2$, and by $N(2, r) := 2r + 1$ if $d = 2$.*

3.3.2 Separation between rare and frequent types

Knowing the values of r and t as defined in Section 3.3.1, we are now able to divide the r-neighborhood types of degree at most d (that is, occurring in structures of degree at most d) into two categories: the **rare types** and the **frequent types**. The intent is that the two structures have the same number of occurrences of every rare type, and that frequent types have many occurrences (wrt. the total number of occurrences of rare types) in both structures. This "many occurrences wrt." is formalized through a function g which is to be specified later on.

More precisely,

Lemma 3.3.2. *Given $d, r \in \mathbb{N}$ and an increasing function $g : \mathbb{N} \rightarrow \mathbb{N}$, there exists $p \in \mathbb{N}$ such that for every Σ-structures $\mathcal{G}_1, \mathcal{G}_2 \in \mathcal{C}_d$ satisfying $\mathcal{G}_1 \equiv_p^{\mathrm{FO}} \mathcal{G}_2$, we can divide the r-neighborhood types over Σ of degree at most d into rare types and frequent types, such that*

- *every rare type has the same number of occurrences in $\mathcal{G}_1$ and in $\mathcal{G}_2$,*
- *both in $\mathcal{G}_1$ and in $\mathcal{G}_2$, every frequent type has at least $g(\beta)$ occurrences, where β is the number of occurrences of all the rare types in the structure,*
- *if there is no frequent type, then $\mathcal{G}_1$ and $\mathcal{G}_2$ are isomorphic.*

Proof. Let $\chi_1, \cdots, \chi_n$ be an enumeration of all the r-neighborhood types over Σ of degree at most d, ordered in such a way that $\forall i < j, |\mathcal{G}_1|_{\chi_i} \leq |\mathcal{G}_1|_{\chi_j}$. Note that n is a function of d and r.

The classification of neighborhood types between rare ones and frequent ones is done through Algorithm 2.

At the end of Algorithm 2, we call $\chi_1, \cdots, \chi_{i-1}$ the rare types, and $\chi_i, \cdots, \chi_n$ the frequent ones.

Note that β indeed counts the total number of occurrences of rare types in $\mathcal{G}_1$.

We now define the integers $(a_i)_{1 \leq i \leq n}$ as $a_1 := g(0)$ and

$$a_{i+1} := \max\{a_i, g(i a_i)\}\,.$$

Algorithm 2 Separation between rare and frequent types

```
β ← 0
i ← 1
while i ≤ n and |G_1|_{χ_i} < g(β) do
    β ← β + |G_1|_{χ_i}
    i++
end while
                          ▷ If i ≤ n, χ_i is the frequent type with the least
                            occurrences in G_1.
                            If i = n + 1, all the neighborhood types are rare.
```

As g is monotone, it is easy to show by induction that for each rare type χ_j with $j < i$, $|\mathcal{G}_1|_{\chi_j} < a_j$.

As long as p is chosen large enough so that $\mathcal{G}_1 \equiv_p^{\mathrm{FO}} \mathcal{G}_2$ entails $[\![\mathcal{G}_1]\!]_r =^{a_n} [\![\mathcal{G}_2]\!]_r$, we have by construction that every rare type has the same number of occurrences (which is smaller that a_n) in $\mathcal{G}_1$ and in $\mathcal{G}_2$. Furthermore, in $\mathcal{G}_1$ as in $\mathcal{G}_2$, if β denotes the total number of occurrences of rare types, every frequent type has at least $g(\beta)$ occurrences.

We just need to make sure that the two structures are isomorphic when all the neighborhood types are rare. If this is the case, then $|\mathcal{G}_1| = |\mathcal{G}_2| \leq n(a_n - 1)$. Hence, as long as $p \geq n(a_n - 1)$, $\mathcal{G}_1 \equiv_p^{\mathrm{FO}} \mathcal{G}_2$ implies $\mathcal{G}_1 \simeq \mathcal{G}_2$ when all the types are rare. □

Let $\tau_0, \cdots, \tau_{m-1}$ be the frequent types. From now on, we suppose that $m \geq 1$: there is nothing to do if $m = 0$, since $\mathcal{G}_1$ and $\mathcal{G}_2$ are isomorphic. Let β be the total number of occurrences of rare types in $\mathcal{G}_1$.

3.3.3 Construction of S_1 around elements of rare type

To begin with, let's focus on $\mathcal{G}_1$, and start the construction of S_1 around occurrences of rare types. Algorithm 3, on page 69, deals with this construction.

For a given occurrence x of some rare type, we choose as its S_1-successor and S_1-predecessor two occurrences of type τ_0 (the first frequent type), far apart from one another and from x. The existence of those elements relies on the bounded degree hypothesis. This is done on lines 8 and 11.

When line 14 is reached, every occurrence of rare type has an S_1-predecessor and an S_1-successor of neighborhood type τ_0.

It is not enough, however, only to deal with the occurrences of rare types. We need to "protect" them up to distance r in $(\mathcal{G}_1, S_1)$. For that purpose, we construct the subsets R_k of G_1, for $0 \leq k \leq r$. In the following, $R_{\leq k}$ denotes $\bigcup\limits_{0 \leq j \leq k} R_j$.

For each k, the subset R_k is constructed in order to be the set of elements at distance exactly k in $(\mathcal{G}_1, S_1)$ from the set of occurrences of rare types. Until we have reached $k = r$ (that is, distance r from occurrences of rare types), every element of R_k is given an S_1-successor (line 21) and/or an S_1-predecessor (line 26) of its neighborhood type, if it doesn't already have one. Once again, those elements are required to be far (i.e. at distance greater than $2r + 1$) from what already has been constructed.

Provided that g is large enough, it is always possible to find x^+ and x^- on lines 8, 11, 21 and 26. Indeed, all the neighborhood types considered are frequent ones, and the size of the $(2r+1)$-neighborhood of $R_{\leq k+1}$ is bounded by a function of d, r and β (the total number of occurrences of rare types in $\mathcal{G}_1$). More precisely, at any point of the construction, $(\mathcal{G}_1, S_1)$ has degree at most $d+2$. Hence, the $(2r+1)$-neighborhood of R_r has size at most

$$\beta N(d+2, 3r+1)$$

(recall the definition of N from Note 3.3.1), and it is enough to make sure that

$$g(\beta) \geq \beta N(d+2, 3r+1) + 1 \,.$$

3.3.4 Construction of S_1 around the junctions between two frequent types

Recall that there is a second kind of singular elements: those which will be at the junction between two successive frequent types. That is, elements of neighborhood type τ_i that will, in the final structure $(\mathcal{G}_1, S_1)$, have an S_1-successor of neighborhood type $\tau_{i+1[m]}$, or an S_1-predecessor of type $\tau_{i-1[m]}$.

Those junction elements need to be treated in a similar way as the occurrences of rare types in Section 3.3.3. This construction is done throughout Algorithm 4 on page 70.

The idea of Algorithm 4 is very similar to that of Algorithm 3. We start by picking two elements x_i^+ and x_i^- for every frequent type τ_i (for loop line 2), that are far from each other and from the previous construction.

Then we build m S_1-edges between those elements on line 9: these edges are intended to be at the junction between the frequent types in the final structure.

The set P_0 of those $2m$ elements will have the same role as the set R_0 of occurrences of rare types for Algorithm 3: we build S_1-edges at depth r around it. This is done through the subsets P_k of G_1, for $0 \leq k \leq r$, P_k being the set of elements at distance k from P_0 in $(\mathcal{G}_1, S_1)$. Once again, $P_{\leq k}$ denotes $\bigcup_{0 \leq j \leq k} P_j$.

For the same reason as for Algorithm 3, it is always possible to find elements x^+ and x^- on lines 17 and 22.

Note that if $m = 1$, there is obviously no transition elements: we simply construct an S_1-edge between x_0^+ and x_0^-.

3.3.5 Carrying S_1 over to $\mathcal{G}_2$

In Sections 3.3.3 and 3.3.4, S_1 has been constructed around the singular points of $\mathcal{G}_1$, i.e. occurrences of rare types and elements that are to make the junction between two S_1-segments of frequent types.

Before we extend S_1 to the remaining elements (all of them being occurrences of frequent types) of $\mathcal{G}_1$, we carry it over to $\mathcal{G}_2$. This transfer is possible under the starting hypothesis that $\mathcal{G}_1$ and $\mathcal{G}_2$ are FO-similar.

Let

$$\begin{cases} A_1 := R_{\leq r} \cup P_{\leq r} \\ B := \{x \in G_1 : \mathrm{dist}_{(\mathcal{G}_1, S_1)}(x, A_1) \leq r\} \end{cases}$$

Algorithm 3 Construction of S_1 around elements of rare type

$S_1 \leftarrow \emptyset$
$R_0 \leftarrow \{x \in G_1 : \mathrm{tp}^r_{\mathcal{G}_1}(x) \text{ is rare}\}$
$R_1, \cdots, R_r \leftarrow \emptyset$
for all $x \in R_0$ **do**
  for all neighbor $y \notin R_{\leq 1}$ of x in $\mathcal{G}_1$ **do**
    $R_1 \leftarrow R_1 \cup \{y\}$
  end for
  find x^+ such that
    $\mathrm{tp}^r_{\mathcal{G}_1}(x^+) = \tau_0$ and
    $\mathrm{dist}_{(\mathcal{G}_1,S_1)}(x^+, R_{\leq 1}) > 2r + 1$
    ▷ We pick a node at distance greater than $2r + 1$ in compliance with Lemma 3.2.5, so that neighborhoods stay layered. Recall that τ_0 is the first frequent type.
  $R_1 \leftarrow R_1 \cup \{x^+\}$
  $S_1 \leftarrow S_1 \cup \{(x, x^+)\}$
  find x^- such that
    $\mathrm{tp}^r_{\mathcal{G}_1}(x^-) = \tau_0$ and
    $\mathrm{dist}_{(\mathcal{G}_1,S_1)}(x^-, R_{\leq 1}) > 2r + 1$
  $R_1 \leftarrow R_1 \cup \{x^-\}$
  $S_1 \leftarrow S_1 \cup \{(x^-, x)\}$
end for
  ▷ At this point, every element of rare type has an S_1-predecessor and an S_1-successor of neighborhood type τ_0
for k from 1 to $r - 1$ **do**
  for all $x \in R_k$ **do** ▷ $\mathrm{tp}^k_{\mathcal{G}_1}(x)$ is a frequent type
    for all neighbor $y \notin R_{\leq k+1}$ of x in $\mathcal{G}_1$ **do**
      $R_{k+1} \leftarrow R_{k+1} \cup \{y\}$
    end for
    if x doesn't have a successor by S_1 **then**
      find x^+ such that
        $\mathrm{tp}^r_{\mathcal{G}_1}(x^+) = \mathrm{tp}^r_{\mathcal{G}_1}(x)$ and
        $\mathrm{dist}_{(\mathcal{G}_1,S_1)}(x^+, R_{\leq k+1}) > 2r + 1$
      $R_{k+1} \leftarrow R_{k+1} \cup \{x^+\}$
      $S_1 \leftarrow S_1 \cup \{(x, x^+)\}$
    end if
    if x doesn't have a predecessor by S_1 **then**
      find x^- such that
        $\mathrm{tp}^r_{\mathcal{G}_1}(x^-) = \mathrm{tp}^r_{\mathcal{G}_1}(x)$ and
        $\mathrm{dist}_{(\mathcal{G}_1,S_1)}(x^-, R_{\leq k+1}) > 2r + 1$
      $R_{k+1} \leftarrow R_{k+1} \cup \{x^-\}$
      $S_1 \leftarrow S_1 \cup \{(x^-, x)\}$
    end if
  end for
end for

If we let t^d_r be the number of r-neighborhood types of degree at most d over Σ, we must have that $m \leq n$ thus $|A_1|$ can be bounded by

$$(\beta + 2t^d_r)N(d + 2, r)\,.$$

Algorithm 4 Construction of S_1 around the junctions between two frequent types

$P_0, \cdots, P_r \leftarrow \emptyset$
for i from 0 to $m-1$ **do**
  find x_i^+ such that
    $\mathrm{tp}^r_{\mathcal{G}_1}(x_i^+) = \tau_i$ and
    $\mathrm{dist}_{(\mathcal{G}_1, S_1)}(x_i^+, R_{\leq r} \cup P_0) > 2r+1$
  $P_0 \leftarrow P_0 \cup \{x_i^+\}$
  find x_i^- such that
    $\mathrm{tp}^r_{\mathcal{G}_1}(x_i^-) = \tau_i$ and
    $\mathrm{dist}_{(\mathcal{G}_1, S_1)}(x_i^-, R_{\leq r} \cup P_0) > 2r+1$
  $P_0 \leftarrow P_0 \cup \{x_i^-\}$
end for
for i from 0 to $m-1$ **do**
  $S_1 \leftarrow S_1 \cup \{(x_i^-, x^+_{i+1[m]})\}$
end for
for k from 0 to $r-1$ **do**
  for all $x \in P_k$ **do**
    for all neighbor $y \notin P_{\leq k+1}$ of x in $\mathcal{G}_1$ **do**
      $P_{k+1} \leftarrow P_{k+1} \cup \{y\}$
    end for
    if x doesn't have a successor by S_1 **then**
      find x^+ such that
        $\mathrm{tp}^r_{\mathcal{G}_1}(x^+) = \mathrm{tp}^r_{\mathcal{G}_1}(x)$ and
        $\mathrm{dist}_{(\mathcal{G}_1, S_1)}(x^+, R_{\leq r} \cup P_{\leq k+1}) > 2r+1$
      $P_{k+1} \leftarrow P_{k+1} \cup \{x^+\}$
      $S_1 \leftarrow S_1 \cup \{(x, x^+)\}$
    end if
    if x doesn't have a predecessor by S_1 **then**
      find x^- such that
        $\mathrm{tp}^r_{\mathcal{G}_1}(x^-) = \mathrm{tp}^r_{\mathcal{G}_1}(x)$ and
        $\mathrm{dist}_{(\mathcal{G}_1, S_1)}(x^-, R_{\leq r} \cup P_{\leq k+1}) > 2r+1$
      $P_{k+1} \leftarrow P_{k+1} \cup \{x^-\}$
      $S_1 \leftarrow S_1 \cup \{(x^-, x)\}$
    end if
  end for
end for

Similarly, the size of B can be bounded by

$$(\beta + 2t_r^d)N(d+2, 2r)\,,$$

which is a function of β, r and d. Hence as long as $f(\alpha)$ is larger than that number, the Duplicator has a winning strategy in the Ehrenfeucht-Fraïssé game between $\mathcal{G}_1$ and $\mathcal{G}_2$ in which the Spoiler chooses every element of B. Let $h : B \to G_2$ be the function resulting from such a strategy.

h defines an isomorphism from $\mathcal{G}_1|_B$ to $\mathcal{G}_2|_{\mathrm{Im}(h)}$. Let $A_2 := h(A_1)$. By taking $f(\alpha)$ one higher than required, to make sure that $\mathrm{Im}(h)$ covers the r-neighborhood in $\mathcal{G}_2$ of every element of A_2, we have that for every $x \in A_1$, $\mathrm{tp}^r_{\mathcal{G}_2}(h(x)) = \mathrm{tp}^r_{\mathcal{G}_1}(x)$.

We set $S_2 := \{(h(x), h(y)) : (x, y) \in S_1\}$. h now defines an isomorphism from $(\mathcal{G}_1, S_1)|_B$ to $(\mathcal{G}_2, S_2)|_{\mathrm{Im}(h)}$, and for every $x \in A_1$, $\mathrm{tp}^r_{(\mathcal{G}_2,S_2)}(h(x)) = \mathrm{tp}^r_{(\mathcal{G}_1,S_1)}(x)$.

3.3.6 Completion of S_1 and S_2

Now that S_1 and S_2 are constructed around all the singular points both in $\mathcal{G}_1$ and $\mathcal{G}_2$, it remains to extend their construction to all the other elements of the structures. Recall that all the remaining elements are occurrences of frequent types.

From $(\mathcal{G}_\epsilon, S_\epsilon)$ ($\epsilon \in \{1,2\}$) at any point in the construction, let's define the partial function $S_\epsilon^* : G_\epsilon \to G_\epsilon$ that maps $x \in G_\epsilon$ to the (unique) y that is S_ϵ-reachable (while taking the orientation into account) from x and that doesn't have an S_ϵ-successor. This function is defined on every element that doesn't belong to an S_ϵ-cycle (and in particular, on every element without an S_ϵ-predecessor).

Likewise, we define S_ϵ^{-*} by reversing the arrows of S_ϵ.

At this point, for every $x \notin A_1$, $S_1^*(x) = S_1^{-*}(x) = x$, and for every $x \notin A_2$, $S_2^*(x) = S_2^{-*}(x) = x$.

We now run Algorithm 5, on page 72. We first treat $\mathcal{G}_1$, and then apply a similar method to $\mathcal{G}_2$, replacing x_i^+ and x_i^- by $h(x_i^+)$ and $h(x_i^-)$. The idea is, for every frequent type τ_i, to insert all its remaining occurrences between (in the sense of S_1) x_i^+ and x_i^-.

The first approach (the loop at line 2) is greedy: while constructing S_ϵ on nodes of neighborhood type τ_i, we choose as the successor of the current node any occurrence of τ_i that is at distance greater than $2r+1$ from the current node s and the closing node of neighborhood type τ_i, $S_\epsilon^{-*}(x_i^-)$. This, together with Lemma 3.2.5, ensures that `(Layer[`r`])` holds after every addition. The conditions line 11 also ensure that the final edge addition, line 15, doesn't break `(Layer[`r`])`.

Once we cannot apply this greedy approach anymore, we know that only a small number (which can be bounded by $2N(d+2, 2r+1)$) of nodes of neighborhood type τ_i remain without S_1-predecessor. The loop at line 17 considers one such node x at a time. As long as g is large enough, we have constructed S_1 around enough elements of type τ_i in the greedy approach to ensure the existence of some $S_1(y,z)$, with y, z of type τ_i and at distance greater than $2r+1$ from x; x is inserted between y and z (line 20). For that, it is enough to have constructed at least

$$2N(d+2, 2r+1) + 1$$

S_1-edges in the greedy phase. This is the case in particular when there are at least

$$4N(d+2, 2r+1) + 1$$

elements of neighborhood type τ_i without S_1-predecessor at the beginning of Algorithm 5, which can be ensured by having

$$g(\beta) \geq |A_1| + 4N(d+2, 2r+1) + 1\,.$$

This holds in particular when

$$g(\beta) \geq (\beta + 2t_r^d)N(d+2, r) + 4N(d+2, 2r+1) + 1\,.$$

We will prove in Lemma 3.3.6 that all these insertions preserve `(Layer[`r`])`.

Algorithm 5 Completion of S_ϵ

for ϵ from 1 to 2 **do**
    for i from 0 to $m-1$ **do**
        if $\epsilon = 1$ **then**
            $s \leftarrow S_1^*(x_i^+)$
            $t \leftarrow S_1^{-*}(x_i^-)$
        else
            $s \leftarrow S_2^*(h(x_i^+))$
            $t \leftarrow S_2^{-*}(h(x_i^-))$
        end if
        while such an x exists **do**
            find x with no S_ϵ-predecessor, such that $\mathrm{tp}^r_{\mathcal{G}_\epsilon}(x) = \tau_i$,
                $\mathrm{dist}_{(\mathcal{G}_\epsilon, S_\epsilon)}(s, x) > 2r+1$,
                $\mathrm{dist}_{(\mathcal{G}_\epsilon, S_\epsilon)}(x, t) > 2r+1$ and
                $\mathrm{dist}_{(\mathcal{G}_\epsilon, S_\epsilon)}(S_\epsilon^*(x), t) > 2r+1$
            $S_\epsilon \leftarrow S_\epsilon \cup \{(s, x)\}$
            $s \leftarrow S_\epsilon^*(x)$
        end while
        ▷ At this point, only a bounded number of elements of type τ_i are left without an S_ϵ-predecessor
        $S_\epsilon \leftarrow S_\epsilon \cup \{(s, t)\}$
    end for
    for i from 0 to $m-1$ **do**
        for all x without S_ϵ-predecessor, s.t. $\mathrm{tp}^r_{\mathcal{G}_\epsilon}(x) = \tau_i$ **do**
            find $y, z \notin A_\epsilon$ such that
                $\mathrm{tp}^r_{\mathcal{G}_\epsilon}(y) = \mathrm{tp}^r_{\mathcal{G}_\epsilon}(z) = \tau_i$,
                $(y, z) \in S_\epsilon$,
                $\mathrm{dist}_{(\mathcal{G}_\epsilon, S_\epsilon)}(y, x) > 2r+1$ and
                $\mathrm{dist}_{(\mathcal{G}_\epsilon, S_\epsilon)}(S_\epsilon^*(x), z) > 2r+1$
            $S_\epsilon \leftarrow S_\epsilon \setminus \{(y, z)\} \cup \{(y, x), (S_\epsilon^*(x), z)\}$
        end for
    end for
end for

3.3.7 Examples of construction

Before we give the proof of correctness of these algorithms, let us see how they apply in some simple cases.

Example 3.3.3. *Suppose that there are no occurrences of rare types, and only one frequent type τ_0, and assume $r = 2$.*

In this case, Algorithm 3 is irrelevant, and all Algorithm 4 does is pick x_0^- and x_0^+ far from each other, and start building S_1 around those nodes in order to construct their complete r-neighborhood in $(\mathcal{G}_1, S_1)$. In order to make the figure more readable, let us consider that x_0^- and x_1^+ have only one neighbor. In Figure 3.5, the plain lines represent edges in $\mathcal{G}_1$, and the dashed arrows represent S_1.

We now apply Algorithm 5. The first step is to add elements between (in the sense of S_1) $S_1^*(x_0^+)$ and $S_1^{-*}(x_0^-)$ in order to join them, in a greedy fashion. Once this is done, there only remain a few elements that haven't been assigned an S_1-predecessor. This is depicted in Figure 3.6.

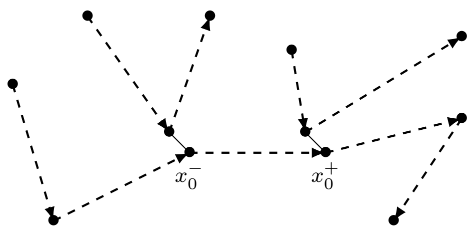

Figure 3.5: After Algorithm 4, with one frequent type.

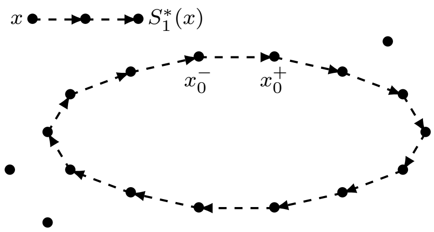

Figure 3.6: After the greedy part of Algorithm 5, with one frequent type.

Now we consider one by one each of the elements that don't have an S_1-predecessor: let's start with x in Figure 3.6. Our goal is to insert it in the S_1-cycle while still respecting (`Layer[r]`)*. For that, we find two successive elements y, z of the cycle that are far from x and $S_1^*(x)$, and we insert x between them, as shown in Figure 3.7.*

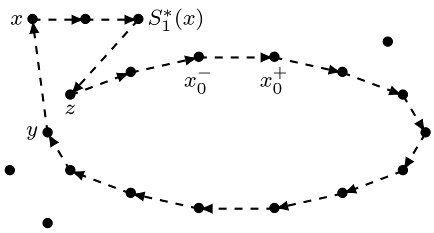

Figure 3.7: Inserting x in the S_1-cycle, as in the second part of Algorithm 5, with one frequent type.

We treat all the elements without an S_1-predecessor in the same way, until S_1 is fully built.

Example 3.3.4. *Suppose now that there are two frequent types τ_0 and τ_1, and still no occurrences of rare types.*

The procedure is very similar: in Algorithm 4, we build the r-neighborhood in $(\mathcal{G}_1, S_1)$ of the four nodes x_0^-, x_0^+, x_1^- and x_1^+.

After the greedy part of Algorithm 5, S_1 looks like in Figure 3.8, where occurrences of τ_0 are represented as $\bullet$ and occurrences of τ_1 as $\circ$. The remaining of Algorithm 5 is as unchanged.

Note that if there existed some occurrences of rare types, they would be embedded in the τ_0 part of the S_1-cycle.

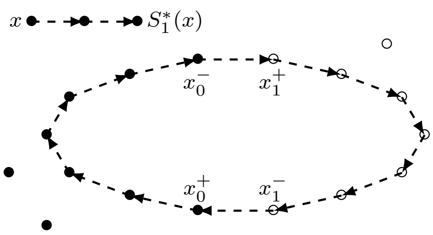

Figure 3.8: After the greedy part of Algorithm 5, with two frequent types.

3.3.8 Properties of S_1 and S_2

We are now ready to show that, after the successive run of Algorithms 3, 4 and 5,

- S_1 and S_2 are indeed successor relations (Lemma 3.3.5),
- $(\mathcal{G}_1, S_1)$ and $(\mathcal{G}_2, S_2)$ satisfy `(Layer[`r`])` (Lemma 3.3.6),
- any singular element (around a rare or a junction element) of $(\mathcal{G}_1, S_1)$ and its corresponding element via h in $(\mathcal{G}_2, S_2)$ have the same r-neighborhood type (Lemma 3.3.9), while any other element in both structures has a regular (i.e. fractal) r-neighborhood type (Lemma 3.3.8).

These properties will allow us to prove in Section 3.3.9 that $(\mathcal{G}_1, S_1)$ and $(\mathcal{G}_2, S_2)$ have the same number of occurrences of every r-neighborhood type, up to a threshold t.

Lemma 3.3.5. *S_1 (resp. S_2) is a successor relation on G_1 (resp. G_2).*

Proof. This result is rather transparent, but a rigorous proof requires the usage of a somewhat cumbersome invariant.

Let us focus on $\mathcal{G}_1$; the proof is the same for $\mathcal{G}_2$, replacing every x_i^+ and x_i^- with $h(x_i^+)$ and $h(x_i^-)$.

Let $a \in G_1$ be defined as $S_1^{-*}(x_{m-1}^-)$ at the beginning of Algorithm 5. By construction, $\mathrm{tp}_{\mathcal{G}}^r(a) = \tau_{m-1}$ and a has no S_1-predecessor as of now.

We show that at any point before line 15 of the loop iteration $i = m - 1$ of Algorithm 5,

(i). $S_1^{-*}(s) = a$

(ii). $S_1^{-*}(x_{m-1}^-) = a$

(iii). let $y, z \notin A_1$ such that $(y, z) \in S_1$ and $\mathrm{tp}_{\mathcal{G}_1}^r(y) = \mathrm{tp}_{\mathcal{G}_1}^r(z) = \tau_j$ for some j; then $S_1^{-*}(y) = a$

(iv). for every i, $(x_i^-, x_{i+1[m]}^+) \in S_1$

(v). there is no S_1-cycle

(vi). for every $j > i$, $\mathrm{tp}_{\mathcal{G}_1}^r(S_1^{-*}(x_j^-)) = \tau_j$.

This is obviously satisfied at the beginning of Algorithm 5: there are not yet such y, z as in (iii), and $s = S_1^*(x_0^+)$ is S_1-reachable from x_{m-1}^- (since $(x_{m+1}^-, x_0^+) \in S_1$) hence (i) holds.

Line 4 preserves the invariant. Indeed, the new value of s is S_1-reachable from its previous value (this is guaranteed by (iv)), which means that they have the same image through S_1^{-*}, namely a.

Let's prove that line 12 preserves the invariant. (i) and (ii) still hold since $x \neq a$: indeed, for $i < m-1$, x and a don't share the same neighborhood type, while for $i = m-1$, $a = t$ (because of (ii)) and the distance condition prohibits $x = a$. (iii) still holds, as the only new possibility for such a couple (y, z) is (s, x), which is such that $S_1^{-*}(y) = a$ (because of (i)). (iv) obviously holds, as does (v), since the only way for an S_1-cycle to have been created is if $x = S_1^{-*}(s)$, that is $x = a$. We have seen that this is absurd. (vi) is satisfied, as the only way for it to fail is for x to be some $S_1^{-*}(x_j^-)$, for $j > i$, which is impossible due to type requirements.

Now, let's move to line 13. Only (i) needs verification, and the argument is the same as for line 4.

Finally, let's look at line 15, for $i < m-1$. $t \neq a$ since their neighborhood types are different, hence (i), (ii) and (v) still hold. (iii) still holds, as the only new possibility for such a couple (y, z) is (s, t), which is such that $S_1^{-*}(y) = a$ because of (i) (actually, (s, t) doesn't even fit the condition, since $t \in A_1$). (iv) is still satisfied. (vi) holds, as the only way for it to fail is for t to be some $S_1^{-*}(x_j^-)$, for $j > i$, which is impossible due to type requirements.

We now prove that from line 17 until the end of Algorithm 5, there is exactly one S_1-cycle, which contains every $y, z \notin A_1$ such that $(y, z) \in S_1$ and $\mathrm{tp}^r_{\mathcal{G}_1}(y) = \mathrm{tp}^r_{\mathcal{G}_1}(z) = \tau_j$ for some j.

This is true after line 15 of the loop iteration $m-1$, which creates the first S_1-cycle, as (i) and (ii) ensure $t = a = S_1^{-*}(s)$. (iii) guarantees that this newly created S_1-cycle contains all the couple (y, z) satisfying the condition.

It remains to show that line 20 preserves this property: by hypothesis, y and z belong to the S_1-cycle. After line 20, there is still exactly one S_1-cycle, which corresponds to the previous one where the S_1-edge has been replaced by the S_1-segment $[x, S_1^*(x)]$. The only S_1-edges that have been added belong to the S_1-cycle, hence the second part of the property still holds.

In the end, every element of G_1 has a predecessor by S_1, hence S_1 is a permutation of G_1. We've shown that it has a single orbit. □

Lemma 3.3.6. *(*`Layer[r]`*) holds in $(\mathcal{G}_\epsilon, S_\epsilon)$, for $\epsilon \in \{1, 2\}$.*

Proof. This property is guaranteed by the distance conditions of the form

$$\mathrm{dist}_{(\mathcal{G}_\epsilon, S_\epsilon)}(.,.) > 2r + 1$$

imposed throughout Algorithms 3, 4 and 5, and by Lemma 3.2.5.

One can very easily verify that (`Layer[r]`) is guaranteed by Lemma 3.2.5 to hold in $(\mathcal{G}_\epsilon, S_\epsilon)$ prior to the run of Algorithm 5.

We focus on Algorithm 5, and we use Lemma 3.2.5 to prove that (`Layer[r]`) remains valid in $(\mathcal{G}_\epsilon, S_\epsilon)$ throughout its run. There are three edge additions we have to prove correct:

- For the edge addition of line 12, this follows directly from Lemma 3.2.5.
- For the edge addition of line 15, we show that the invariant

  $$\text{dist}_{(\mathcal{G}_\epsilon, S_\epsilon)}(s,t) > 2r+1$$

  is satisfied at the beginning and at the end of the while line 10. This invariant, together with Lemma 3.2.5, will be enough to conclude.

  The invariant holds before the first execution of the while loop (except for $m = 1$, where it only bootstraps after two executions of the loop).

  Working towards a contradiction, suppose that the invariant is broken during an execution of the loop. We use the pre notations. There must exists a path from $S_\epsilon^*(x)$ (which is to become the new value of s at the end of the loop) to t in $(\mathcal{G}_\epsilon, S_\epsilon \cup \{(s,x)\})$ of length at most $2r+1$; consider a shortest one. As is cannot be valid in $(\mathcal{G}_\epsilon, S_\epsilon)$ by choice of x, it must go through the newly added edge (s,x). This means that in $(\mathcal{G}_\epsilon, S_\epsilon)$, either there exist paths of length at most $2r+1$ from $S_\epsilon^*(x)$ to s and from x to t, or paths of length at most $2r+1$ from $S_\epsilon^*(x)$ to x and from s to t. The former is absurd considering the way x was chosen, and the latter contradicts the previous invariant.

- Let's prove that the addition of the two S_ϵ-edges of line 20 doesn't break `(Layer[`r`])`. By choice of y, we know that $\text{dist}_{(\mathcal{G}_\epsilon, S_\epsilon)}(y,x) > 2r+1$. *A fortiori*, we must have $\text{dist}_{(\mathcal{G}_\epsilon, S_\epsilon \setminus \{(y,z)\})}(y,x) > 2r+1$, and Lemma 3.2.5 ensures that $(\mathcal{G}_\epsilon, S_\epsilon \setminus \{(y,z)\} \cup \{(y,x)\})$ satisfies the property `(Layer[`r`])`.

  Now, to the second addition: let's prove that, at the beginning of line 20, $\text{dist}_{(\mathcal{G}_\epsilon, S_\epsilon \setminus \{(y,z)\} \cup \{(y,x)\})}(S_\epsilon^*(x), z) > 2r+1$. We are then able to conclude with Lemma 3.2.5.

  Suppose it's not the case and consider a shortest path from $S_\epsilon^*(x)$ to z, which must be of length at most $2r+1$. This path cannot be valid in $(\mathcal{G}_\epsilon, S_\epsilon)$, thus it has to go through the new edge (y,x). Since there cannot exist a path of length at most $2r$ from $S_\epsilon^*(x)$ to y in $(\mathcal{G}_\epsilon, S_\epsilon)$ (as this would contradict $\text{dist}_{(\mathcal{G}_\epsilon, S_\epsilon)}(S_\epsilon^*(x), z) > 2r+1$), it has to borrow the edge from x to y.

  Then in $(\mathcal{G}_\epsilon, S_\epsilon \setminus \{(y,z)\})$, there is a path of length at most $2r$ from y to z, which contradicts `(Layer[`r`])` in $(\mathcal{G}_\epsilon, S_\epsilon)$.

□

The following lemma states that the only time S_ϵ joins two nodes that have different r-neighborhood types in $\mathcal{G}_\epsilon$ is when one of them is an occurrence of a rare type (in which case its S_ϵ-predecessor and S_ϵ-successor are of neighborhood type τ_0) or when they are the elements which make the transition between two frequent types (that is, one is x_i^- and the other is $x_{i+1[m]}^+$, for some $i < m$).

Lemma 3.3.7. *$\forall x, y \in G_1$ such that $(x,y) \in S_1$ and ($x \notin R_0$ and $y \notin R_0$) and ($x \notin P_0$ or $y \notin P_0$), then $tp^r_{\mathcal{G}_1}(x) = tp^r_{\mathcal{G}_1}(y)$.*

$\forall x, y \in G_2$ such that $(x,y) \in S_2$ and ($x \notin h(R_0)$ and $y \notin h(R_0)$) and ($x \notin h(P_0)$ or $y \notin h(P_0)$), then $tp^r_{\mathcal{G}_2}(x) = tp^r_{\mathcal{G}_2}(y)$.

Proof. The property clearly holds at the end of Algorithm 3 and Algorithm 4.

For any i from 0 to $m-1$, the only S_1-edges (resp. S_2-edges) that are added during the i-th loop are between two nodes of neighborhood type τ_i. □

Recall the discussion at the beginning of Section 3.2. We now prove that, as long as an element is far from any occurrence of a rare type and from the elements that make the transition between two frequent types, its neighborhood type in $(\mathcal{G}_\epsilon, S_\epsilon)$ is the fractal of its neighborhood type in $\mathcal{G}_\epsilon$.

Lemma 3.3.8. *For $\epsilon \in \{1,2\}$ and for every $0 \leq k \leq r$ and $x \notin R_{\leq k} \cup P_{\leq k}$ (if $\epsilon = 1$) or $x \notin h(R_{\leq k} \cup P_{\leq k})$ (if $\epsilon = 2$),*

$$tp^k_{(\mathcal{G}_\epsilon, S_\epsilon)}(x) = [tp^k_{\mathcal{G}_\epsilon}(x)]_k\,.$$

Proof. We prove the result by induction on k. For $k = 0$, there is nothing to do but note that no edge $S_\epsilon(x,x)$ has been created.

Suppose that we've proven the result for some $k < r$, and let $x \notin R_{\leq k+1} \cup P_{\leq k+1}$, or $x \notin h(R_{\leq k+1} \cup P_{\leq k+1})$.

Let y be such that $\mathrm{dist}_{\mathcal{G}_\epsilon}(x,y) = d$, for some $1 \leq d \leq k+1$. By construction of the R_i and P_i, and of h, we have that $y \notin R_{\leq k+1-d} \cup P_{\leq k+1-d}$, or $y \notin h(R_{\leq k+1-d} \cup P_{\leq k+1-d})$ (this is easily shown by induction on d). Hence, $\mathrm{tp}^{k+1-d}_{(\mathcal{G}_\epsilon, S_\epsilon)}(y) = [\mathrm{tp}^{k+1-d}_{\mathcal{G}_\epsilon}(y)]_{k+1-d}$.

Because Lemma 3.3.6 ensures that the $(k+1)$-neighborhood of x in $(\mathcal{G}_\epsilon, S_\epsilon)$ is layered, it only remains to show that the S_ϵ-successor x^+ and predecessor x^- of x are such that $\mathrm{tp}^k_{(\mathcal{G}_\epsilon, S_\epsilon)}(x^+) = \mathrm{tp}^k_{(\mathcal{G}_\epsilon, S_\epsilon)}(x^-) = [\mathrm{tp}^k_{\mathcal{G}_\epsilon}(x)]_k$. Let's show this for x^+.

Lemma 3.3.7 ensures $\mathrm{tp}^r_{\mathcal{G}_\epsilon}(x^+) = \mathrm{tp}^r_{\mathcal{G}_\epsilon}(x)$. It only remains to note that $x^+ \notin R_{\leq k} \cup P_{\leq k}$, or $x^+ \notin h(R_{\leq k} \cup P_{\leq k})$, and the induction hypothesis allows us to conclude. □

When we first defined h, it preserved r-neighborhood types by construction. The last step before we are able to conclude the proof of Theorem 3.1.1 is to make sure that h still preserves r-neighborhood types, taking into account the S_ϵ-edges added during the run of Algorithm 5.

Lemma 3.3.9. $\forall x \in A_1, tp^r_{(\mathcal{G}_2, S_2)}(h(x)) = tp^r_{(\mathcal{G}_1, S_1)}(x)$.

Proof. We prove by induction on $0 \leq k \leq r$ that

$$\forall x \in A_1, \mathrm{tp}^k_{(\mathcal{G}_2, S_2)}(h(x)) = \mathrm{tp}^k_{(\mathcal{G}_1, S_1)}(x)\,.$$

There is nothing to prove for $k = 0$.

Moving from k to $k+1$, let $x \in A_1$ and let y be such that $\mathrm{dist}_{\mathcal{G}_1}(x,y) = d$, for some $1 \leq d \leq k+1$. Note that $y \in B$, hence it has an image by h.

If $y \in A_1$, the induction hypothesis allows us to conclude that

$$\mathrm{tp}^{k+1-d}_{(\mathcal{G}_2, S_2)}(h(y)) = \mathrm{tp}^{k+1-d}_{(\mathcal{G}_1, S_1)}(y)\,.$$

Else, Lemma 3.3.8 ensures that:

$$\mathrm{tp}^r_{(\mathcal{G}_2, S_2)}(h(y)) = [\mathrm{tp}^r_{\mathcal{G}_2}(h(y))]_r = [\mathrm{tp}^r_{\mathcal{G}_1}(y)]_r = \mathrm{tp}^r_{(\mathcal{G}_1, S_1)}(y)\,.$$

In both cases, $\mathrm{tp}^{k+1-d}_{(\mathcal{G}_2, S_2)}(h(y)) = \mathrm{tp}^{k+1-d}_{(\mathcal{G}_1, S_1)}(y)$.

Because of (`Layer[`r`]`), it only remains to show that the S_ϵ-successors of x and $h(x)$, as well as their S_ϵ-predecessors, have the same k-neighborhood type in $(\mathcal{G}_\epsilon, S_\epsilon)$. Let's prove this for the successors, respectively named x^+ and $h(x)^+$.

If $x^+ \in A_1$, then by construction $h(x)^+ = h(x^+)$, and the induction hypothesis allows us to conclude.

Otherwise, neither x^+ nor $h(x)^+$ belongs to A_1. Under this hypothesis, Lemma 3.3.7 ensures that

$$\mathrm{tp}^r_{\mathcal{G}_2}(h(x)^+) = \mathrm{tp}^r_{\mathcal{G}_2}(h(x)) = \mathrm{tp}^r_{\mathcal{G}_1}(x) = \mathrm{tp}^r_{\mathcal{G}_1}(x^+)\,.$$

Now, Lemma 3.3.8 ensures that

$$\begin{aligned}\mathrm{tp}^r_{(\mathcal{G}_2,S_2)}(h(x)^+) &= [\mathrm{tp}^r_{\mathcal{G}_2}(h(x)^+)]_r\\ &= [\mathrm{tp}^r_{\mathcal{G}_1}(x^+)]_r\\ &= \mathrm{tp}^r_{(\mathcal{G}_1,S_1)}(x^+)\,.\end{aligned}$$

□

3.3.9 Conclusion of the proof

We are now able to conclude the proof. Recall that we want to prove that

$$[\![(\mathcal{G}_1, S_1)]\!]_r =^t [\![(\mathcal{G}_2, S_2)]\!]_r\,.$$

Let τ be an r-neighborhood type over $\Sigma \cup \{S\}$ which occurs in $(\mathcal{G}_1, S_1)$. There are two cases to consider:

- if τ occurs outside of A_1, then Lemma 3.3.8 ensures that $\tau = [\chi]_r$ for some frequent r-neighborhood type χ. We can choose g so that χ is guaranteed to have at least t occurrences in $\mathcal{G}_1$ outside of A_1, and in $\mathcal{G}_2$ outside of A_2. This is ensured as long as

  $$g(\beta) \geq |A_1| + t\,,$$

  and in particular when

  $$g(\beta) \geq (\beta + 2t^d_r)N(d+2, r) + t\,.$$

  Lemma 3.3.8 then ensures that τ occurs at least t times both in $(\mathcal{G}_1, S_1)$ and in $(\mathcal{G}_2, S_2)$.

- if τ occurs only in A_1, then it cannot occur in $(\mathcal{G}_2, S_2)$ outside of A_2 (for the same reasons as above).

  Lemma 3.3.9 guarantees that τ has the same number of occurrences in A_1 and in A_2, hence in $(\mathcal{G}_1, S_1)$ and in $(\mathcal{G}_2, S_2)$.

Recall that r and t where chosen accordingly to 2.3.4 for Σ and α: we have that $[\![(\mathcal{G}_1, S_1)]\!]_r =^t [\![(\mathcal{G}_2, S_2)]\!]_r$ entails

$$(\mathcal{G}_1, S_1) \equiv^{\mathrm{FO}}_\alpha (\mathcal{G}_2, S_2)\,.$$

We have thus shown that for every $\alpha \in \mathbb{N}$, there exists some $f(\alpha) \in \mathbb{N}$ such that for any Σ-structures $\mathcal{G}_1, \mathcal{G}_2$ of degree at most d,

$$\mathcal{G}_1 \equiv^{\text{FO}}_{f(\alpha)} \mathcal{G}_2 \quad \to \quad \mathcal{G}_1 \equiv^{\text{Succ-inv}/_{\mathcal{C}_d}\ \text{FO}}_{\alpha} \mathcal{G}_2 \,.$$

This, together with Proposition 2.3.2, proves that

$$\text{Succ-inv}/_{\mathcal{C}_d}\ \text{FO} \subseteq \text{FO on } \mathcal{C}_d \,.$$

3.4 Conclusion

We have shown in this section that Succ-inv FO collapses to FO on any class of bounded degree, as well as on classes of graphs which are FO-interpretable in graph classes of bounded degree, namely near-uniform graph classes as defined in [19].

Our proof gives a constructive translation from Succ-inv FO to FO on classes of bounded degree. The quantifier rank of the translated sentence is triple-exponential in the quantifier rank of the original formula. As seen in Proposition 2.2.10, the blowup is at least exponential, but we do not know if an exponential translation is at all possible.

Similar considerations arise when we take into account the length of the sentences instead of their quantifier rank - in this regard, our construction is even non-elementary, and all we know is that the blowup is at least exponential.

An interesting task would be to improve the succinctness of the translation, or to give tighter lower bounds on such constructions.

Apart from these considerations, there are two main directions in which one could look to extend the present result. One possibility would be to keep looking at classes of bounded degree while climbing up in the ladder of expressivity, and ask whether $<$-inv FO collapses to FO as well on these classes of structures. New techniques would be needed, as contrary to what was the case with a successor, the addition of an order doesn't preserve the bounded degree property. Furthermore, even if $<$-inv FO = FO in this setting, it is not clear whether such orders can be directly constructed. It may be necessary to construct, as in [5], a chain of intermediate structures and orders.

Alternatively, we could change the setting, and study the expressivity of Succ-inv FO on other sparse classes of structures, e.g. on classes of bounded treewidth. If showing the collapse of Succ-inv FO to FO on these classes proved itself to be out of reach, a possibility would be to aim at proving that Succ-inv FO is Hanf local (which would be stronger than the known Gaifman locality). In that case, the starting hypothesis on the structures $\mathcal{G}_1$ and $\mathcal{G}_2$ would be stronger, as the existence of a k-neighborhood type-preserving bijection between the two structures would be assumed.

These tasks are much harder without any bound on the degree, which was what guaranteed that we could find elements of a given frequent type far from each other.

Chapter 4

Order-Invariant FO on Hollow Trees

After proving the collapse of Succ-inv FO to FO on classes of bounded degree in Chapter 3, we turn to $<$-inv FO.

Contrary to a successor relation, adding an order to a structure brings all the elements close together, as the Gaifman graph of an ordered structure has diameter 1. With this in mind, using a method similar to that of Chapter 3 (i.e. what we called the *direct method* in Section 2.6) to prove a collapse of $<$-inv FO to FO on any non-trivial class of structures seems ambitious.

Such a task seems difficult, if at all possible, already for dipaths. Recall from Proposition 2.4.8 that $<$ -inv FO = FO on the class of dipaths. However, it is not clear whether, given two FO-similar dipaths, there exist orders on those dipaths that preserve FO-similarity.

It thus seems that in order to prove collapses of $<$ -inv FO to FO, the *chaining method* from Section 2.6.3 is the most appropriate. For instance, it was the strategy used in [4] and [5] to prove that $<$-inv FO = FO on trees.

In this chapter, we prove the collapse of $<$-inv FO to FO on hollow trees. We believe that this result is a step towards proving the collapse on classes of pathwidth 2, as will be explained in Section 4.1.2. As a corollary, we get that $<$-inv FO = FO on paths.

We start by defining hollow trees in Section 4.1, and we explain in which way they relate to graphs of pathwidth 2. In Section 4.2, we state the main result of this chapter, as well as some corollaries.

Sections 4.3 through 4.6 detail the techniques used in the proof of this result, which is given in Section 4.7.

This chapter, which is a slightly altered version of [22], results from a joint work with Luc Segoufin.

4.1 Definitions

4.1.1 General notations

In the course of this chapter, we will need the following definitions.

For $k \in \mathbb{N}$, we define the **k-enrichment** $\mathcal{E}_k(\mathcal{A})$ of a Σ-structure $\mathcal{A}$ as $\mathcal{A}$ itself where each element has been recolored with its k-neighborhood type. $\mathcal{E}_k(\mathcal{A})$ is a structure over the vocabulary Σ augmented with a unary predicate for every k-neighborhood type over Σ: there is a finite number of them as long as we consider classes of structures of bounded degree.

Recall that, for any k-neighborhood type τ, we denote by $|\mathcal{A}|_\tau$ the number of elements of A whose k-neighborhood type is τ.

In this chapter, most of our structures will be enriched by recoloring each element by its k-neighborhood type, hence we will essentially use 0-neighborhood types, which capture k-neighborhood types in the original structure. In view of this we denote by $[\![\mathcal{A}]\!]$ the function $\tau \mapsto |\mathcal{A}|_\tau$ whose domain is the set of 0-neighborhood types over the considered vocabulary. Note that $[\![\mathcal{E}_k(\mathcal{A})]\!]$ is the function $\tau \mapsto |\mathcal{A}|_\tau$ whose domain is the set of all k-neighborhood types over Σ. In particular, $[\![\mathcal{E}_k(\mathcal{A})]\!] = [\![\mathcal{E}_k(\mathcal{B})]\!]$ iff $[\![\mathcal{A}]\!]_k =^\infty [\![\mathcal{B}]\!]_k$.

Let $d, D \in \mathbb{N}$, and f, g be functions from a same domain to $\mathbb{N}$; in practice, f and g will be $[\![\mathcal{E}_k(\mathcal{A})]\!]$ for some k and some structure $\mathcal{A}$. We say that $f \leq_d^D g$ if, for every x in the domain,

- if $f(x) \leq d$, then $f(x) = g(x)$
- if $f(x) \neq g(x)$, then $g(x) \geq f(x) + D$.

By $f < g$, we mean that $\forall x, (f(x) < g(x) \text{ or } f(x) = g(x) = 0)$.

Let $\mathcal{A}$ be a structure over a vocabulary containing the binary relation symbol R. We say that $U \subseteq A$ is **R-stable** if

$$\forall x \in U, \forall y \in A, (R(x,y) \vee R(y,x)) \rightarrow y \in U .$$

4.1.2 Hollow trees

Barát, Hajnal, Lin and Yang [2] proved that any graph of pathwidth at most 2 can be decomposed in a series of what they called *tracks*. Thus, a first step towards proving the collapse of $<$-inv FO to FO on classes of pathwidth at most 2 is to show that $<$-inv FO = FO on the class of tracks.

A typical example of track of degree 3 is depicted in Figure 4.1, where the dashed arcs are colored paths, and all the chords are single edges. Each chord could actually be a single edge or the juxtaposition of two edges with a single vertex in the middle; however Proposition 2.3.6 allows us to ignore that case, since there exists a simple bi-FO-interpretation that gets rid of the middle vertices by coloring the chords according to whether they are simple or double edges.

We show in Figure 4.2 how such a track can be turned into a structure resembling a tree. We add color and number identifiers to clarify the translation.

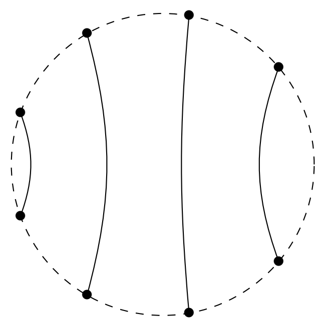

Figure 4.1: Example of track of degree 3.

Note that those two transformations, as well as their inverse, are definable as FO-interpretations as soon as the square edge is part of the track. Hence, Note 2.3.7 guarantees that the collapse of $<$-inv FO to FO on one of those classes of structures amounts to the collapse on the others.

This remark motivates the definition of hollow trees. Informally, hollow trees resemble the aforementioned tree-like structure with the key difference that the vertical edges (i.e. the parent-child edges) and the horizontal one are distinguishable. In return for that specification, we do not restrict the complexity of the underlying tree, while the tree-like structures resulting from the transformation of a track are very constrained. In particular, the class of hollow trees has unbounded pathwidth.

Let's now formally define hollow trees.

An unranked ordered tree is a tree with a successor relation among the children of any node. We see unranked ordered trees as structures over the signature composed of two binary relation symbols S and S', where S is interpreted as the parent-child relation, and S' as the horizontal successor. A set of nodes that share the same parent is called a siblinghood.

We define a mapping H from the set of unranked ordered trees to structures over two binary predicates S and E. Given an unranked ordered tree $\mathcal{T}$, $H(\mathcal{T})$ is defined as follows:

- its domain is T,
- $H(\mathcal{T}) \models S(x, y)$ iff $\mathcal{T} \models S(x, y)$ and y is either the first or the last of its siblings,
- E is interpreted as the symmetrical closure of S'.

The image of H is the set of **hollow trees**, denoted $\mathbb{H}$. If $\mathcal{P} = H(\mathcal{T})$ then $\mathcal{T}$ is the underlying tree structure of $\mathcal{P}$.

In other words, within a hollow tree, only the two children at the endpoints of a siblinghood know their parent. Notice that we do not distinguish between the first and last child, nor do we between the left and right sibling. This makes the model more general, as explained in Section 4.2.2.

An example of hollow tree is given in Figure 4.3.

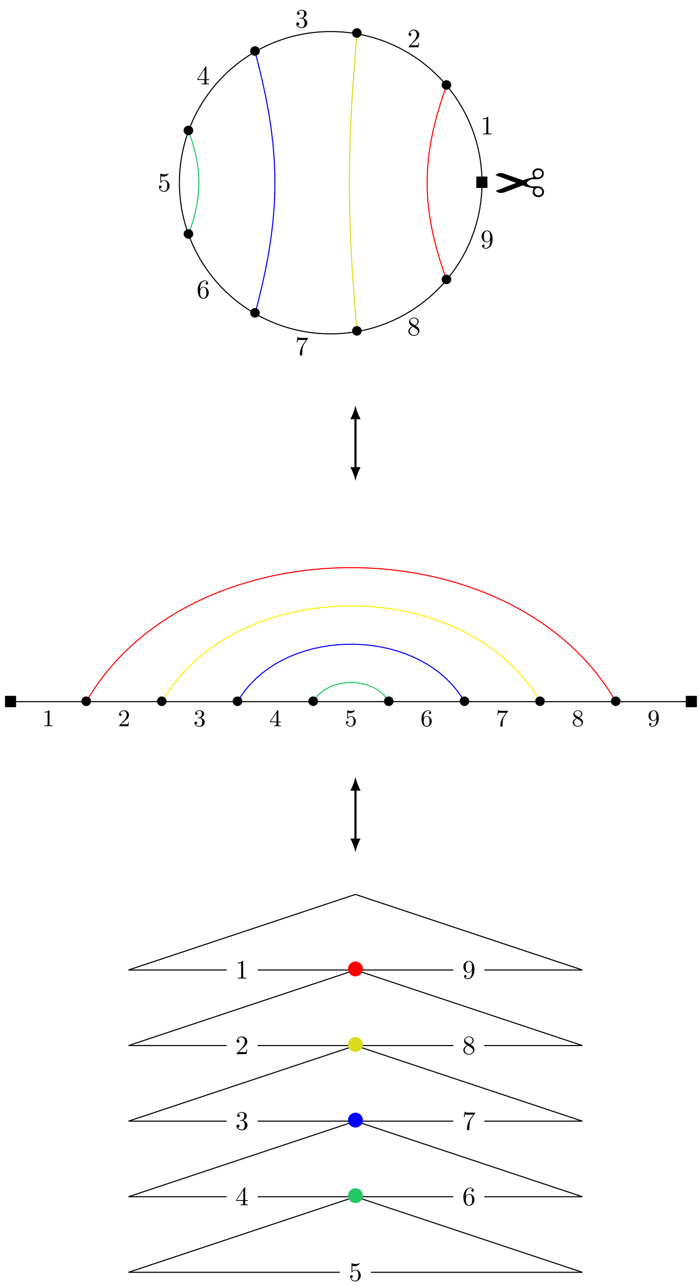

Figure 4.2: Turning a track of degree 3 into a tree-like structure. One goes from the track to the intermediate structure by cutting the edge represented as a square, and from the intermediate structure to the tree-like structure by contracting each chord into a vertex.

Given a finite alphabet σ, we define $\mathbb{H}_\sigma$, the set of **hollow trees over** σ, as the set of colored extensions of hollow trees using the vocabulary P_σ, where the interpretations of the predicates of P_σ partition the domain.

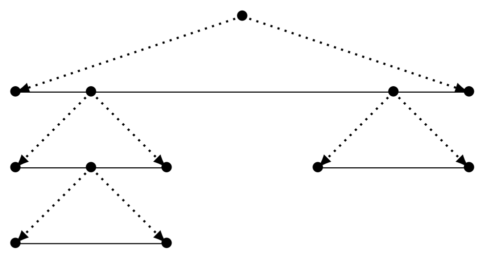

Figure 4.3: An example of hollow tree. The dotted arrows represent S and the plain (symmetrical) lines represent E.

Note that for every σ, $\mathbb{H}_\sigma$ is a class of structures of treewidth 2. Therefore, according to Proposition 2.5.3, $<$-inv FO $\subseteq$ MSO on $\mathbb{H}_\sigma$.

4.2 Overview of the results

4.2.1 Main result

The main result we prove in this chapter is the collapse of $<$-inv FO to FO on hollow trees.

Theorem 4.2.1. *For every alphabet σ,*

$$<\text{-inv FO} = \text{FO} \textit{ on } \mathbb{H}_\sigma\,.$$

We sketch here the proof of this result, before developping in Sections 4.3 through 4.6 the tools used in that proof. It will then be given in full details Section 4.7.

This proof follows the chaining method introduced in Section 2.6.3. Recall that our goal is to find some function f such that

$$\forall \alpha \in \mathbb{N}, \forall \mathcal{P}, \mathcal{Q} \in \mathbb{H}_\sigma, \quad \mathcal{P} \equiv^{\text{FO}}_{f(\alpha)} \mathcal{Q} \quad \rightarrow \quad \mathcal{P} \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{Q}\,.$$

Proposition 2.3.2 then allows us to conclude.

To show this we fix some $\alpha \in \mathbb{N}$ and consider two hollow trees $\mathcal{P}$ and $\mathcal{Q}$, such that $\mathcal{P} \equiv^{\text{FO}}_{f(\alpha)} \mathcal{Q}$ for a large enough $f(\alpha)$. As explained in Section 2.6.3, the general idea is to exhibit a sequence of intermediate ordered structures between $\mathcal{P}$ and $\mathcal{Q}$, which are pairwise FO-similar at depth α.

The hard part is to find such intermediate structures and orders. Each intermediate structure results from the application of some operation to the previous one, starting from $\mathcal{Q}$, until we reach $\mathcal{P}$. All of those operations are designed to be invisible to all formulas of $<$-inv FO of quantifier rank less than α.

In the end, we get that $\mathcal{P} \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{Q}$.

We will use two kinds of operations, as described in Section 4.3: “swap operations”, which preserve $<$-inv FO, and one which preserves MSO (and *a fortiori* $<$-inv FO as $<$-inv FO $\subseteq$ MSO on $\mathbb{H}_\sigma$ by Proposition 2.5.3).

The MSO-preserving operation will be used in Section 4.3.3, in order to pump $\mathcal{Q}$ to make sure that every neighborhood type is present at least as many times in $\mathcal{Q}$ as in $\mathcal{P}$.

Once this is done, we explain in Section 4.4 how to transform $\mathcal{Q}$ with swap operations in order to include $\mathcal{P}$ into it. Since $\mathcal{Q}$ may be larger than $\mathcal{P}$, there could be some extra material in $\mathcal{Q}$ that we call “loops”. The last step is to remove those loops and this is the goal of Section 4.6.

When performing the swap operations, there will be a constant need for reorganizing the S-edges (in particular to make sure that the loops are S-stable). Section 4.5 and Section 4.6.3 compile the results that allow us to do so.

4.2.2 Corollaries

Recall the results from Section 2.3.2: Proposition 2.3.6 states that the collapse $<$-inv FO = FO can be lifted from a class to another as long as there exists a bi-FO-interpretation from the latter through the former.

The first consequence of this result is that we can assume a normal form on hollow trees without loss of generality. Namely, we can assume that each S-parent has exactly two S-children, and that no element is at the same time an S-parent and an S-child. Indeed, there is a simple bi-FO-interpretation that transforms a general hollow tree into one having the desired properties (by duplicating nodes that are only child, and those that are simultaneously S-parent and S-child, and marking them with new unary predicates), and back to the initial one.

Recall that, in the definition of hollow trees, the relation E is symmetric. This turns out to be more general than choosing E as an arbitrary directed binary relation, as shown in the following result, where **directed hollow trees** are defined as hollow trees but with a directed binary relation E instead of a symmetric one. Note that we do not assume that E is a successor relation among siblings: the direction of E could be arbitrary, but the result below is valid in particular when E is a successor relation on siblinghoods.

Via a simple bi-FO-interpretation which uses extra colors to encode the direction of the edges, we get the following result:

Corollary 4.2.1. *For every alphabet σ,*

$$<\text{-inv FO} = \text{FO} \textit{ on the class of directed hollow trees over } \sigma\,.$$

Proof. We use Lemma 2.3.6 and exhibit a bi-FO-interpretation from directed hollow trees over σ through hollow trees over $\sigma \cup \{—, |\}$.

We give the first FO-interpretation $\mathcal{I}$ (from directed hollow trees to hollow trees), and leave the reverse one to the reader.

To avoid confusion in the notations, let's rename the directed binary relation E as F in the vocabulary of directed hollow trees: hence $\mathcal{I}$ goes from the vocabulary $\{F, S\} \cup P_\sigma$ to $\{E, S\} \cup P_{\sigma \cup \{—,|\}}$.

Given a σ directed hollow tree $\mathcal{T}$, $\mathcal{I}(\mathcal{T})$ is defined as follows. An illustration of this interpretation is given in Figure 4.4.

- Its domain is T, plus two new elements v_{xy} and v_{yx} for every $x, y \in T$ such that $\mathcal{T} \models F(x, y)$,

- the interpretation of S is unchanged,
- E is interpreted as the union of

$$\{(x, v_{xy}), (v_{xy}, x), (v_{xy}, v_{yx}), (v_{yx}, v_{xy}), (v_{yx}, y), (y, v_{yx})\}$$

  for every $x, y \in T$ such that $\mathcal{T} \models F(x, y)$,
- the interpretation of every $P \in P_\sigma$ is unchanged,
- $P_{—}$ is interpreted as $\{v_{xy} : x, y \in T, \mathcal{T} \models F(x, y)\}$,
- $P_{|}$ is interpreted as $\{v_{yx} : x, y \in T, \mathcal{T} \models F(x, y)\}$.

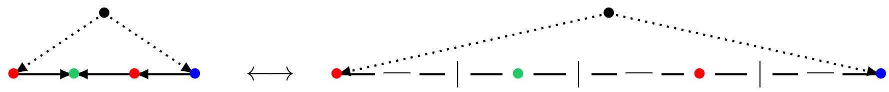

Figure 4.4: The encoding of a directed hollow tree $\mathcal{T}$ as the (symmetrical) hollow tree $\mathcal{I}(\mathcal{T})$.

This encoding allows the converse FO-interpretation to recover the orientation of the F-edges of $\mathcal{T}$ from $\mathcal{I}(\mathcal{T})$ in a straightforward way. □

Recall from Example 2.1.1 the definition of a path over σ. The class of paths over σ is obviously bi-FO-interpretable through $\mathbb{H}_\sigma$: from paths to hollow trees, just add an S-parent to the endpoints of the path, and from hollow trees to paths, forget this element. Thus we get the following result:

Corollary 4.2.2. *For every alphabet σ,*

$$<\text{-inv FO} = \text{FO} \textit{ on the class of paths over } \sigma$$

Similarly, a straightforward bi-FO-interpretation in conjunction with Theorem 4.2.1 give us back the result from [5] that $<$ -inv FO = FO on ranked trees.

4.3 Swaps and pumping

In this section we provide a few operations, which we call *swaps*, that preserve $\equiv_\alpha^{<\text{-inv FO}}$. Although the k-neighborhood type of every element will be left unchanged, applying these operations may break the somewhat rigid structure of hollow trees. In order to work with the intermediate structures, we loosen the definition of hollow trees and define hollow quasitrees as follows.

Definition 4.3.1. *For $k > 0$ and σ a set of colors, we define the set of* ***hollow k-quasitrees on σ****, quasi-$\mathbb{H}_\sigma^k$, as the set of all finite structures over $\{E, S\} \cup P_\sigma$ such that the k-neighborhood type of any of their elements is the k-neighborhood type of some element in some hollow tree in $\mathbb{H}_\sigma$, and which are such that their relation E is acyclic.*

In other words a hollow quasitree differs from a hollow tree by its relation S which may not induce a tree structure: a node may have its S-children in two distinct siblinghoods and a hollow quasitree may have cycles using the relation S (but not using only the relation E). All these properties are obviously not definable in FO: locally, a hollow quasitree looks like a hollow tree.

Note that by definition $\mathbb{H}_\sigma \subseteq$ quasi-$\mathbb{H}_\sigma^k$ for every k. An example of what a hollow quasitree could look like is given in Figure 4.5.

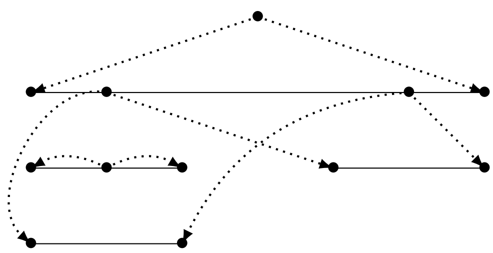

Figure 4.5: An example of hollow quasitree. The dotted arrows represent S and the plain (symmetrical) lines represent E.

Let $\mathcal{T} \in$ quasi-$\mathbb{H}_\sigma^k$. We define the **support** of $\mathcal{T}$ as its restriction to the vocabulary $P_\sigma \cup \{E\}$. The n**-enriched support** of $\mathcal{T}$, denoted $\mathcal{S}\mathrm{upp}_n(\mathcal{T})$, is the support of its n-enrichment (and not the other way around). Hence, it keeps in memory the local behavior within $\mathcal{T}$. The set $\mathrm{End}(\mathcal{T})$ of **endpoints** of $\mathcal{T}$ is the set of elements of the support having degree one. A connected component of the support of $\mathcal{T}$ is called a **thread**[1]. Note that by E-acyclicity of $\mathcal{T}$, each of its threads is a path, hence contains exactly two endpoints, but those endpoints may not have the same S-parent. We say that a hollow k-quasitree has the **matching endpoints property** if the two endpoints of each thread have the same S-parent. Note that a hollow tree has the matching endpoints property. Notice also that in a hollow k-quasitree, any thread of length less than $2k+1$ has matching endpoints, for otherwise the k-neighborhood of the midway element of the thread would witness that the matching endpoints property is not satisfied, and its k-neighborhood type would thus never appear in a hollow tree. For $x, y \in T$ belonging to the same thread, $[x, y]$ denotes the set of elements that lie between them (formally, those who disconnect x from y in $\mathcal{S}\mathrm{upp}_0(\mathcal{T})$), including x and y. We naturally define $[x, y[$ as $[x, y] \setminus \{y\}$.

The following lemma, implicit in the proof of Proposition 2.4.6 by Grohe and Schwentick [24], will allow us to prove that our operations preserve order-invariance equivalence.

Lemma 4.3.2. *Let Σ be a relational vocabulary and let $p, \alpha \in \mathbb{N}$. There exists $o_p^\Sigma(\alpha) \in \mathbb{N}$ such that for every structure $\mathcal{A}$ over Σ, and for every p-tuples of elements $\bar{a}, \bar{b} \in A^p$ that have the same $o_p^\Sigma(\alpha)$-neighborhood type in $\mathcal{A}$, there are two orders $<_{\bar{a}\bar{b}}$ and $<_{\bar{b}\bar{a}}$ on A such that*

- $(\mathcal{A}, <_{\bar{a}\bar{b}}) \equiv_\alpha^{\mathrm{FO}} (\mathcal{A}, <_{\bar{b}\bar{a}})$,

[1] A thread is nothing other than a siblinghood when the quasitree is a tree.

- $\bar{a}\bar{b}$ *is an initial segment of* $<_{\bar{a}\bar{b}}$,
- $\bar{b}\bar{a}$ *is an initial segment of* $<_{\bar{b}\bar{a}}$.

Our operations are divided into three families depending on whether we modify the relation S, the relation E, or whether we do a global pumping,

In the following, $\mathcal{R}$ is a hollow $(m+1)$-quasitree on σ.

4.3.1 Crossing-S-swaps

Let $a, a', a'', b, b', b'' \in R$ be such that $S(a,a'), S(a,a''), S(b,b'), S(b,b'')$ and such that $\text{tp}^m_{\mathcal{R}}(a,a',a'') = \text{tp}^m_{\mathcal{R}}(b,b',b'')$.

Let $\mathcal{R}^- := \mathcal{R} \setminus \{S(a,a'), S(a,a''), S(b,b'), S(b,b'')\}$ and assume that the sets $\{a',a''\}, \{b',b''\}$ and $\{a,b\}$ are pairwise $(2m+3)$-distant in $\mathcal{R}^-$.

$\mathcal{R}' := \mathcal{R}^- \cup \{S(a,b'), S(a,b''), S(b,a'), S(b,a'')\}$ is called the **m-guarded crossing-S-swap between a and b in $\mathcal{R}$** (see Figure 4.6).

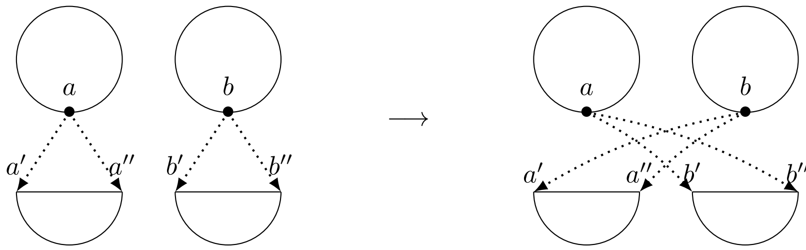

Figure 4.6: The crossing-S-swap between a and b.

Note 4.3.3. *A particular case where the distance condition is met is when* $dist_{\mathcal{R}}(a,b) \geq 2m+5$.

Lemma 4.3.4. *For all* $\alpha \in \mathbb{N}$ *there exists* $s(\alpha) \in \mathbb{N}$ *such that for all* $m \geq s(\alpha)$*, and every hollow* $(m+1)$*-quasitree* $\mathcal{R}$*,*

if $\mathcal{R}'$ *is the* m*-guarded crossing-*S*-swap between* a *and* b *in* $\mathcal{R}$*,*

then $\mathcal{R}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$*, and* $\forall x \in R$*,* $tp^{m+1}_{\mathcal{R}'}(x) = tp^{m+1}_{\mathcal{R}}(x)$ *.*

Moreover $\mathcal{R}' \in$ *quasi-*$\mathbb{H}^{m+1}_\sigma$ *and* $\mathcal{S}upp_{m+1}(\mathcal{R}') = \mathcal{S}upp_{m+1}(\mathcal{R})$ *.*

In order to prove that $\mathcal{R}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$ we need to exhibit a linear order over $\mathcal{R}$ and one over $\mathcal{R}'$ such that we can play an α-round Ehrenfeucht-Fraïssé game between the resulting ordered structures. The linear orders are constructed using Lemma 4.3.2 applied to (a',a'') and (b',b'') and the structure $\mathcal{R}^-$. A simple FO-interpretation is then used to transfer the corresponding orders onto $\mathcal{R}$ and $\mathcal{R}'$.

Proving that the neighborhood type of an element is unchanged is rather straightforward. It then follows that $\mathcal{R}' \in$ quasi-$\mathbb{H}^{m+1}_\sigma$, since $\mathcal{R}$ was itself a hollow $(m+1)$-quasitree, and the same $(m+1)$-neighborhood types occur in both structures. Similarly, it entails that $\mathcal{S}\text{upp}_{m+1}(\mathcal{R}') = \mathcal{S}\text{upp}_{m+1}(\mathcal{R})$, since only S-edges has changed between $\mathcal{R}$ and $\mathcal{R}'$, which in $\mathcal{S}\text{upp}_{m+1}(\mathcal{R})$ and $\mathcal{S}\text{upp}_{m+1}(\mathcal{R}')$ are only accounted for in the neighborhood type of every element (which is unchanged).

Proof. We first show that $\mathcal{R}' \equiv_{\alpha}^{<\text{inv FO}} \mathcal{R}$. This is essentially a reduction to Lemma 4.3.2 using FO-interpretations. Recall the function o_p^{Σ} given by Lemma 4.3.2. We use it with $p = 2$ and $\Sigma := P_\sigma \cup \{E, S, P_{1/2}, P_{3/4}\}$ where $P_{1/2}$ and $P_{3/4}$ are unary. Assume now that $m \geq o_2^{\Sigma}(\alpha + c)$ where c is a constant to be chosen later on.

Consider the extension $\mathcal{R}^\star$ of $\mathcal{R}^-$ to Σ where the interpretation of $P_{1/2}$ is $\{a\}$ and that of $P_{3/4}$ is $\{b\}$. Since $P_{1/2}^{\mathcal{R}^\star}$ and $P_{3/4}^{\mathcal{R}^\star}$ are at distance $> m$ from a', a'', b' and b'' in $\mathcal{R}^\star$, we have that $\text{tp}^m_{\mathcal{R}^\star}(a', a'') = \text{tp}^m_{\mathcal{R}^\star}(b', b'')$.

We can therefore apply Lemma 4.3.2, and get two orders $<_{a'a''b'b''}$ and $<_{b'b''a'a''}$ such that $(\mathcal{R}^\star, <_{a'a''b'b''}) \equiv^{\text{FO}}_{\alpha+c} (\mathcal{R}^\star, <_{b'b''a'a''})$. Now, consider the FO-interpretation that adds an S-edge between u and v if either

- $P_{1/2}(u)$ and v is the first or the second element of $<$
- or $P_{3/4}(u)$ and v is the third or the fourth element of $<$,

and then forgets about $P_{1/2}$ and $P_{3/4}$.

Take c to be the depth of this FO-interpretation (which has arity 1).

Note that the result of this FO-interpretation on $(\mathcal{R}^\star, <_{a'a''b'b''})$ is an ordered extension of $\mathcal{R}$ and its result on $(\mathcal{R}^\star, <_{b'b''a'a''})$ is an ordered extension of $\mathcal{R}'$.

This entails $\mathcal{R}' \equiv_{\alpha}^{<\text{inv FO}} \mathcal{R}$.

Now, let $x \in R$, and let's show that $\text{tp}^{m+1}_{\mathcal{R}'}(x) = \text{tp}^{m+1}_{\mathcal{R}}(x)$.

First, if x is at distance $> m + 1$ of $\{a, a', a'', b, b', b''\}$ in $\mathcal{R}$, there isn't any change in its $(m + 1)$-neighborhood.

Otherwise, there are several cases to consider, according to whether x belongs to $\mathcal{N}^{m+1}_{\mathcal{R}^-}(a)$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(a')$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(a'')$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(b)$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(b')$, or $\mathcal{N}^{m+1}_{\mathcal{R}^-}(b')$; we treat the first one of them, the others being similar.

Set $d_a := \text{dist}_{\mathcal{R}^-}(x, a)$ and $d_b := \text{dist}_{\mathcal{R}^-}(x, b)$. By hypothesis, $d_a \leq m + 1$.

We distinguish two cases:

- if $d_b > m + 1$: because of the distance constraint, we can partition $\mathcal{N}^{m+1}_{\mathcal{R}}(x)$ into $\mathcal{N}^{m+1}_{\mathcal{R}^-}(x)$ and $\mathcal{N}^{m-d_a}_{\mathcal{R}^-}(a', a'')$, with two S-edges joining a in the first and a', a'' in the second. These two parts are at distance ≥ 2 in R^-, hence they are fully independent (no overlap, and no edge between the two except for $S(a, a')$ and $S(a, a'')$).

  Likewise, we can partition $\mathcal{N}^{m+1}_{\mathcal{R}'}(x)$ into $\mathcal{N}^{m+1}_{\mathcal{R}^-}(x)$ and $\mathcal{N}^{m-d_a}_{\mathcal{R}^-}(b', b'')$.

  $\mathcal{N}^{m-d_a}_{\mathcal{R}^-}(a', a'') \simeq \mathcal{N}^{m-d_a}_{\mathcal{R}^-}(b', b'')$, hence $\mathcal{N}^{m+1}_{\mathcal{R}}(x) \simeq \mathcal{N}^{m+1}_{\mathcal{R}'}(x)$.

- otherwise, $d_b \leq m + 1$: now, we can partition $\mathcal{N}^{m+1}_{\mathcal{R}}(x)$ into $\mathcal{N}^{m+1}_{\mathcal{R}^-}(x)$, $\mathcal{N}^{m-d_a}_{\mathcal{R}^-}(a', a'')$ and $\mathcal{N}^{m-d_b}_{\mathcal{R}^-}(b', b'')$, with two S-edges joining a in the first to a', a'' in the second and two S-edges joining b in the first to b', b'' in the third, as depicted in Figure 4.7. These three parts are at distance ≥ 2 in $\mathcal{R}^-$, hence they are fully independent (no overlap, and no edge between the two except for $S(a, a'), S(a, a''), S(b, b')$ and $S(b, b'')$).

  Likewise, we can partition $\mathcal{N}^{m+1}_{\mathcal{R}'}(x)$ into $\mathcal{N}^{m+1}_{\mathcal{R}^-}(x)$, $\mathcal{N}^{m-d_a}_{\mathcal{R}^-}(b', b'')$ and $\mathcal{N}^{m-d_b}_{\mathcal{R}^-}(a', a'')$, as shown in Figure 4.7. We have that

  $$\mathcal{N}^{m-d_a}_{\mathcal{R}^-}(a', a'') \simeq \mathcal{N}^{m-d_a}_{\mathcal{R}^-}(b', b'')$$

and

$$\mathcal{N}_{\mathcal{R}^-}^{m-d_b}(a', a'') \simeq \mathcal{N}_{\mathcal{R}^-}^{m-d_b}(b', b'')\,,$$

hence

$$\mathcal{N}_{\mathcal{R}}^{m+1}(x) \simeq \mathcal{N}_{\mathcal{R}'}^{m+1}(x)\,.$$

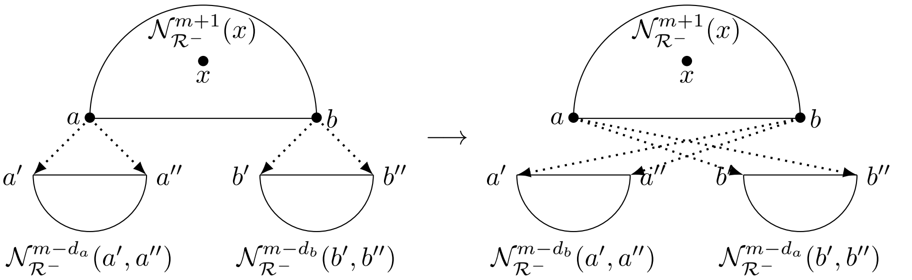

Figure 4.7: Evolution of the neighborhood of x before and after a crossing-S-swap in the proof of Lemma 4.3.4. We see $\mathcal{N}_{\mathcal{R}}^{m+1}(x)$ on the left and $\mathcal{N}_{\mathcal{R}'}^{m+1}(x)$ on the right.

□

4.3.2 E-swaps

We start by defining four different kinds of E-swaps: crossing-E-swaps, mirror-E-swaps, segment-E-swaps and contiguous-segment-E-swaps.

We then prove some results holding for every kind of E-swap.

Crossing-E-swaps

Let $a, b, a', b' \in R$ be such that $E(a, b), E(a', b')$, a, b and a', b' appear in two different threads of $\mathcal{R}$ and such that $\{a, b, a', b'\}$ and $\mathrm{End}(\mathcal{R})$ are $(2m+3)$-distant in $\mathcal{S}\mathrm{upp}_0(\mathcal{R})$. Furthermore, assume that $\mathrm{tp}_{\mathcal{R}}^m(a, b) = \mathrm{tp}_{\mathcal{R}}^m(a', b')$.

Let $\mathcal{R}' := \mathcal{R} \setminus \{E(a, b), E(a', b')\} \cup \{E(a, b'), E(a', b)\}$.

Then $\mathcal{R}'$ is called the **m-guarded crossing-E-swap between ab and $a'b'$ in** $\mathcal{R}$ (c.f. Figure 4.8).

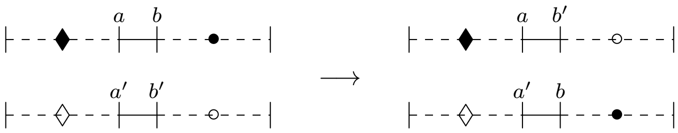

Figure 4.8: Illustration of the m-guarded crossing-E-swap between ab and $a'b'$ in $\mathcal{R}$.

Mirror-E-swaps

Let $a, b, b', a' \in R$ appear in that order in a single thread of $\mathcal{R}$, such that $E(a, b), E(a', b')$, and such that $\{a, b, a', b'\}$ and $\mathrm{End}(\mathcal{R})$ are $(2m+3)$-distant in $\mathcal{S}\mathrm{upp}_0(\mathcal{R})$. Furthermore, assume that $\mathrm{tp}_{\mathcal{R}}^m(a, b) = \mathrm{tp}_{\mathcal{R}}^m(a', b')$.

Let $\mathcal{R}' := \mathcal{R} \setminus \{E(a,b), E(a',b')\} \cup \{E(a,b'), E(a',b)\}$.

Then $\mathcal{R}'$ is called the **m-guarded mirror-E-swap at** $[b,b']$ **in** $\mathcal{R}$ (c.f. Figure 4.9).

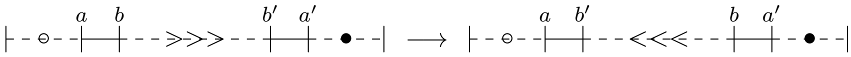

Figure 4.9: Illustration of the m-guarded mirror-E-swap at $[b,b']$ in $\mathcal{R}$.

Segment-E-swaps

Consider now $a,b,c,d,a',b',c',d' \in R$ appearing in that order in a single thread of $\mathcal{R}$, satisfying

- $E(a,b), E(c,d), E(a',b')$ and $E(c',d')$,
- $\{a,b,c,d,a',b',c',d'\}$ and $\mathrm{End}(\mathcal{R})$ are $(2m+3)$-distant in $\mathcal{S}\mathrm{upp}_0(\mathcal{R})$,
- $\mathrm{tp}^m_{\mathcal{R}}(a,b) = \mathrm{tp}^m_{\mathcal{R}}(a',b')$,
- $\mathrm{tp}^m_{\mathcal{R}}(c,d) = \mathrm{tp}^m_{\mathcal{R}}(c',d')$.

$$\begin{aligned}\text{Let } \mathcal{R}' \;:=\; \mathcal{R} \;&\setminus\; \{E(a,b), E(a',b'), E(c,d), E(c',d')\} \\ &\cup\; \{E(a,b'), E(a',b), E(c,d'), E(c',d)\}\,.\end{aligned}$$

Then $\mathcal{R}'$ is called the **m-guarded segment-E-swap between** $[b,c]$ **and** $[b',c']$ **in** $\mathcal{R}$ (c.f. Figure 4.10).

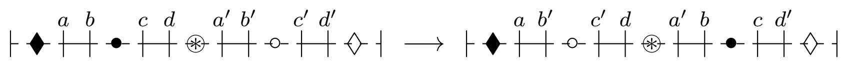

Figure 4.10: Illustration of the m-guarded segment-E-swap between $[b,c]$ and $[b',c']$ in $\mathcal{R}$.

Contiguous-segment-E-swaps

Finally, let a,b,a',b',a'',b'' be elements of R appearing in that order in a single thread of $\mathcal{R}$, such that

- $E(a,b), E(a',b')$ and $E(a'',b'')$,
- $\{a,b,a',b',a'',b''\}$ and $\mathrm{End}(\mathcal{R})$ are $(2m+3)$-distant in $\mathcal{S}\mathrm{upp}_0(\mathcal{R})$,
- $\mathrm{tp}^m_{\mathcal{R}}(a,b) = \mathrm{tp}^m_{\mathcal{R}}(a',b') = \mathrm{tp}^m_{\mathcal{R}}(a'',b'')$.

$$\begin{aligned}\text{Let } \mathcal{R}' \;:=\; \mathcal{R} \;&\setminus\; \{E(a,b), E(a',b'), E(a'',b'')\} \\ &\cup\; \{E(a,b'), E(a',b''), E(a'',b)\}\,.\end{aligned}$$

Then $\mathcal{R}'$ is called the **m-guarded contiguous-segment-E-swap between** $[b,a']$ **and** $[b',a'']$ **in** $\mathcal{R}$ (c.f. Figure 4.11).

As long as m is large enough, all the m-guarded E-swaps preserve $\equiv_{\alpha}^{<\text{-inv FO}}$ and the $(m+1)$-neighborhood type of every element.

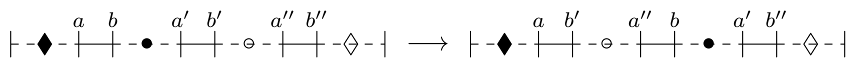

Figure 4.11: Illustration of the m-guarded contiguous-segment-E-swap between $[b, a']$ and $[b', a'']$ in $\mathcal{R}$.

Lemma 4.3.5. *For all $\alpha \in \mathbb{N}$ there exists $s(\alpha) \in \mathbb{N}$ such that for every $m \geq s(\alpha)$ and every hollow $(m+1)$-quasitree $\mathcal{R}$, if $\mathcal{R}'$ is either*

- *the m-guarded crossing-E-swap between ab and $a'b'$ in $\mathcal{R}$*
- *or the m-guarded mirror-E-swap at $[b, b']$ in $\mathcal{R}$*
- *or the m-guarded contiguous-segment-E-swap between $[b, a']$ and $[b', a'']$ in $\mathcal{R}$*
- *or the m-guarded segment-E-swap between $[b, c]$ and $[b', c']$ in $\mathcal{R}$,*

then $\mathcal{R}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$, $\forall x \in R, tp_{\mathcal{R}'}^{m+1}(x) = tp_{\mathcal{R}}^{m+1}(x)$ and $\mathcal{R}' \in$ quasi-$\mathbb{H}_\sigma^{m+1}$.

Note 4.3.6. *In each of these cases, the swap doesn't introduce any E-loop. Hence, once we've shown that every elements keeps its $(m+1)$-neighborhood type, we immediately get that $\mathcal{R}' \in$ quasi-$\mathbb{H}_\sigma^{m+1}$.*

However, these operations do not preserve hollowtreeness. This is the reason why we introduced the notion of quasitree.

The proof is a tedious case analysis. Basically it amounts to the following idea: if the elements involved in the swap are far away from each other then we can use Lemma 4.3.2 in the structure $\mathcal{R}$ minus the E-edges of interest, and get orders on $\mathcal{R}$ and $\mathcal{R}'$ which make these structures similar as in the proof of Lemma 4.3.4.

On the other hand, if the elements are close to each other, then the fact that they share the same neighborhood type induces some periodicity on their neighborhoods. These neighborhoods can therefore be decomposed into several consecutive similar pieces. We can then apply Lemma 4.3.2 to these smaller components to conclude.

We prove Lemma 4.3.5 separately for every type of E-swap. We will need the following lemmas.

Lemma 4.3.7. *Let $\mathcal{Q} \in$ quasi-$\mathbb{H}_\sigma^{m+1}$ and $x, y \in Q$ such that the sets $\{x\}, \{y\}$ and End$(\mathcal{Q})$ are pairwise $(2m+3)$-distant in $\mathcal{S}upp_0(\mathcal{Q})$.*

Then $dist_{\mathcal{Q}}(x, y) \geq 2m+3$.

Proof. This is a consequence of the fact that in a hollow $(m+1)$-quasitree, a thread of length less than $2m+1$ must have matching endpoints.

Suppose that there exists path of length $\leq 2m+2$ in $\mathcal{Q}$ from x to y, and let p be such a path of shortest length.

The path p may use either an E-edge or an S-edges. We divide p into segments between two consecutive S-edges. Let $t_1, \cdots, t_r$ be the corresponding threads in that order (with possible repetitions).

We know that $x \in t_1$ and $y \in t_r$. As x and y are at distance at least $2m+3$ when using only E-edges, we must have $r \geq 2$.

There are two ways for p to go from t_i to t_{i+1}: either (1) using and S-edge $S(a, e_{i+1})$ with $a \in t_i$ and $e_{i+1} \in \text{End}(t_{i+1})$, or (2) using an S-edge $S(b, e_i)$ with $b \in t_{i+1}$ and $e_i \in \text{End}(t_i)$.

As x is far from the endpoints of t_1, p must go from t_1 to t_2 using case (1). Similarly, p must go from t_{r-1} to t_r using case (2). Hence, there must exist some $1 < i < r$ such that p moves from t_{i-1} to t_i following (1), and from t_i to t_{i+1} following (2).

Since p is a shortest path, the two endpoints of t_i involved in (1) and (2) cannot be the same; hence p goes from one endpoint of t_i to the other and the length of t_i must be $\leq 2m+1$. Since $Q \in$ quasi-$\mathbb{H}_\sigma^{m+1}$, branches of length $\leq 2m+1$ must have matching endpoints. This contradicts the minimality of p, since p could have avoided t_i completely. □

Lemma 4.3.8. *Let $\mathcal{Q} \in$ quasi-$\mathbb{H}_\sigma^{m+1}$ and $x \neq y \in Q$ belonging to the same thread, such that $\{x, y\}$ and End($\mathcal{Q}$) are $(2m+3)$-distant in $\mathcal{S}upp_0(\mathcal{Q})$.*

Every path of length $< 2m+3$ between x and y goes through every E-edge of $[x, y]$.

In other words, if $\mathcal{Q}^-$ is $\mathcal{Q}$ minus any E-edge of $[x, y]$, $\text{dist}_{\mathcal{Q}^-}(x, y) \geq 2m+3$.

Proof. The proof is identical as the one of Lemma 4.3.7, by considering a shortest path of length $\leq 2m+2$ from x to y that doesn't go through every E-edge of $[x, y]$: we arrive at the same contradiction. □

The following lemma will only be needed in Section 4.5; however, we state it here as its proof is very similar to the previous ones.

Lemma 4.3.9. *Let $\mathcal{Q} \in$ quasi-$\mathbb{H}_\sigma^{m+1}$, and $x, y \in$ End($\mathcal{Q}$) such that*

$$dist_{\mathcal{S}upp_0(\mathcal{Q})}(x, y) \geq 2m+3\,.$$

Then any path of length $\leq 2m+3$ from x to y goes through at least one of their S-parents.

Proof. We proceed similarly as in Lemma 4.3.7; let's use the same notations.

Suppose that p doesn't go through x's parent neither y's. Let's show that p goes from t_1 to t_2 using case (1): if not, it uses case (2) through the other endpoint of t_1. Hence, t_1 would be of length $\leq 2m+1$ (since it takes at least 2 to reach y from there), and would have matching endpoints; this is absurd, since p would go through x's parent.

Similarly, we show that p moves from t_{r-1} to t_r following (2).

We can conclude exactly as in Lemma 4.3.7. □

We are now ready to prove Lemma 4.3.5 in each of the four cases.

Proof of Lemma 4.3.5 for crossing-E-swaps

We let $\mathcal{R}^- := \mathcal{R} \setminus \{E(a, b), E(a', b')\}$.

It follows from Lemma 4.3.7 and Lemma 4.3.8 that a, b, a' and b' are at distance at least $2m+3$ from each other in $\mathcal{R}^-$.

First, we show that $\mathcal{R}' \equiv_\alpha^{\text{<-inv FO}} \mathcal{R}$. This is essentially a reduction to Lemma 4.3.2 using FO-interpretations.

Recall the function o_p^Σ given by Lemma 4.3.2. We use it with $p = 1$ and $\Sigma := P_\sigma \cup \{E, S, P_1, P_2\}$ where P_1 and P_2 are unary. Let's now assume that $m \geq o_1^\Sigma(\alpha + c)$, where c is a constant to be chosen later on.

Consider the extension $\mathcal{R}^\star$ of $\mathcal{R}^-$ to Σ where the interpretation of P_1 is $\{b\}$ and that of $\mathcal{P}_2$ is $\{b'\}$. Since $P_1^{\mathcal{R}^\star}$ and $P_2^{\mathcal{R}^\star}$ are at distance $> m$ from a and a' in $\mathcal{R}^\star$, we have that $\mathrm{tp}^m_{\mathcal{R}^\star}(a) = \mathrm{tp}^m_{\mathcal{R}^\star}(a')$.

We can therefore apply Lemma 4.3.2, and get two orders $<_{aa'}$ and $<_{a'a}$ such that $(\mathcal{R}^\star, <_{aa'}) \equiv^{\mathrm{FO}}_{\alpha+c} (\mathcal{R}^\star, <_{a'a})$. Now, consider the FO-interpretation that adds a (symmetrical) E-edge between u and v if either

- $P_1(u)$ and v is the first element of $<$
- or $P_2(u)$ and v is the second element of $<$,

and then forgets about P_1 and P_2.

Take c to be the depth of this FO-interpretation (which has arity 1). Note that the result of this FO-interpretation on $(\mathcal{R}^\star, <_{aa'})$ is an ordered extension of $\mathcal{R}$ and that its result on $(\mathcal{R}^\star, <_{a'a})$ is an ordered extension of $\mathcal{R}'$. This entails $\mathcal{R}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$.

Now, let $x \in R$, and let's show that $\mathrm{tp}^{m+1}_{\mathcal{R}'}(x) = \mathrm{tp}^{m+1}_{\mathcal{R}}(x)$.

First, if x is at distance $> m + 1$ of $\{a, b, a', b'\}$ in $\mathcal{R}$, there isn't any change in its $(m + 1)$-neighborhood.

Otherwise, there are several cases to consider, according to whether x belongs to $\mathcal{N}^{m+1}_{\mathcal{R}^-}(a)$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(a')$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(b)$ or $\mathcal{N}^{m+1}_{\mathcal{R}^-}(b')$; we treat the first one, the others being similar.

Set $d := \mathrm{dist}_{\mathcal{R}}(x, a)$.

We can partition $\mathcal{N}^{m+1}_{\mathcal{R}}(x)$ into $\mathcal{N}^{m+1}_{\mathcal{R}^-}(x)$ and $\mathcal{N}^{m-d}_{\mathcal{R}^-}(b)$, with an E-edge joining a in the first and b in the second.

Because of the distance condition, these two parts are at distance ≥ 2 in $\mathcal{R}^-$, hence they are fully independent (no overlap, and no edge between the two except for $E(a, b)$).

Likewise, we can partition $\mathcal{N}^{m+1}_{\mathcal{R}'}(x)$ into $\mathcal{N}^{m+1}_{\mathcal{R}^-}(x)$ and $\mathcal{N}^{m-d}_{\mathcal{R}^-}(b')$.

$\mathcal{N}^{m-d}_{\mathcal{R}^-}(b) \simeq \mathcal{N}^{m-d}_{\mathcal{R}^-}(b')$, hence $\mathcal{N}^{m+1}_{\mathcal{R}}(x) \simeq \mathcal{N}^{m+1}_{\mathcal{R}'}(x)$.

Proof of Lemma 4.3.5 for mirror-E-swaps

Let $\mathcal{R}^- := \mathcal{R} \setminus \{E(a, b), E(a', b')\}$.

It follows from Lemma 4.3.8 that the three sets $\{a\}, \{a'\}$ and $\{b, b'\}$ are $(2m + 3)$-distant in $\mathcal{R}^-$.

The proof that $\mathcal{R}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$ is done exactly as in the case of crossing-E-swaps.

Now, let $x \in R$, and let's show that $\mathrm{tp}^{m+1}_{\mathcal{R}'}(x) = \mathrm{tp}^{m+1}_{\mathcal{R}}(x)$.

First, if x is at distance $> m + 1$ of $\{a, b, a', b'\}$ in $\mathcal{R}$, there isn't any change in its $(m + 1)$-neighborhood.

Otherwise, there are several cases to consider, according to whether x belongs to $\mathcal{N}^{m+1}_{\mathcal{R}^-}(a)$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(a')$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(b)$ or $\mathcal{N}^{m+1}_{\mathcal{R}^-}(b')$; the first two are similar to the cases appearing in the proof for crossing-E-swaps.

We treat the third one, the fourth being symmetrical.

Set $d := \mathrm{dist}_{\mathcal{R}^-}(x, b)$ and $d' := \mathrm{dist}_{\mathcal{R}^-}(x, b')$. By hypothesis, $d \leq m + 1$.

$\mathcal{N}^{m+1}_{\mathcal{R}}(x)$ can partitioned into $\mathcal{N}^{m+1}_{\mathcal{R}^-}(x)$, $\mathcal{N}^{m-d}_{\mathcal{R}^-}(a)$ and the possibly empty $\mathcal{N}^{m-d'}_{\mathcal{R}^-}(a')$, with an E-edge joining b in the first to a in the second, and an E-edge joining b' in the first to a' in the third (if it is nonempty).

We claim that any two of these three neighborhoods are at distance ≥ 2 in $\mathcal{R}^-$, hence they are fully independent: no overlap, and no edge between any two of them, except (possibly) for $E(a, b)$ and (possibly) $E(a', b')$.

Indeed, suppose (the other pairs of neighborhoods are treated similarly) that $\mathcal{N}^{m+1}_{\mathcal{R}^-}(x)$ and $\mathcal{N}^{m-d}_{\mathcal{R}^-}(a)$ are at distance ≤ 1. Then $\mathrm{dist}_{\mathcal{R}^-}(a, x) \leq 2m + 2 - d$, hence $\mathrm{dist}_{\mathcal{R}^-}(a, b) \leq 2m + 2$, which contradicts Lemma 4.3.8 for a and b (recall that $\{a, b\}$ and $\mathrm{End}(\mathcal{R})$ are $(2m + 3)$-distant in $\mathcal{S}\mathrm{upp}_0(\mathcal{R})$).

Likewise we can partition $\mathcal{N}^{m+1}_{\mathcal{R}'}(x)$ into $\mathcal{N}^{m+1}_{\mathcal{R}^-}(x)$, $\mathcal{N}^{m-d}_{\mathcal{R}^-}(a')$ and $\mathcal{N}^{m-d'}_{\mathcal{R}^-}(a)$.

We have that $\mathcal{N}^{m-d}_{\mathcal{R}^-}(a) \simeq \mathcal{N}^{m-d}_{\mathcal{R}^-}(a')$ and $\mathcal{N}^{m-d'}_{\mathcal{R}^-}(a') \simeq \mathcal{N}^{m-d'}_{\mathcal{R}^-}(a)$, hence $\mathcal{N}^{m+1}_{\mathcal{R}}(x) \simeq \mathcal{N}^{m+1}_{\mathcal{R}'}(x)$.

Proof of Lemma 4.3.5 for contiguous-segment-E-swaps

Let $\mathcal{R}^- := \mathcal{R} \setminus \{E(a, b), E(a', b'), E(a'', b'')\}$.

Let x, y be non-endpoint elements belonging to the same thread of some $\mathcal{Q} \in \text{quasi-}\mathbb{H}^{n+1}_{\sigma}$. Let $\mathcal{Q}^- := \mathcal{Q} \setminus \{E(x', x), E(y, y')\}$, where x' (resp. y') is the E-neighbor of x (resp. y) that doesn't belong to $[x, y]$.

We denote by $[x, y]^{\mathcal{Q}}_n$ the substructure of $\mathcal{Q}^-$ induced by the set of nodes at distance $\leq n$ in $\mathcal{Q}^-$ from $[x, y]$, together with a new color marking x as the left endpoint.

We define concatenation as follows: if x, x_1, y_1, y appear in the same thread in that order, and $E(x_1, y_1)$, then we write $[x, y]^{\mathcal{Q}}_n =: [x, x_1]^{\mathcal{Q}}_n \cdot [y_1, y]^{\mathcal{Q}}_n$.

Let us abbreviate $\mathrm{dist}_{\mathcal{S}\mathrm{upp}_0(\mathcal{Q})}(x, y)$ as $|[x, y]|$ (that is, the distance from x to y if we are only allowed E-edges).

We first prove that $(m + 1)$-neighborhood types are unchanged by an m-guarded contiguous-segment-E-swap.

Let $x \in R$, and let's show that $\mathrm{tp}^{m+1}_{\mathcal{R}'}(x) = \mathrm{tp}^{m+1}_{\mathcal{R}}(x)$

If x is at distance $> m+1$ of $\{a, b, a', b', a'', b''\}$ in $\mathcal{R}$, there isn't any change in its $(m + 1)$-neighborhood.

Otherwise, there are several cases to consider, according to whether x belongs to $\mathcal{N}^{m+1}_{\mathcal{R}^-}(a)$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(b)$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(a')$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(b')$, $\mathcal{N}^{m+1}_{\mathcal{R}^-}(a'')$ or $\mathcal{N}^{m+1}_{\mathcal{R}^-}(b'')$. We treat the second one, the other ones being similar or simpler.

Set $d_b := \mathrm{dist}_{\mathcal{R}^-}(x, b)$, $d_{a'} := \mathrm{dist}_{\mathcal{R}^-}(x, a')$, and $d_{b'a''} := \mathrm{dist}_{\mathcal{R}^-}(b', a'')$. By hypothesis, $d_b \leq m + 1$.

We can partition $\mathcal{N}^{m+1}_{\mathcal{R}}(x)$ into $\mathcal{N}^{m+1}_{\mathcal{R}^-}(x)$, $\mathcal{N}^{m-d_b}_{\mathcal{R}^-}(a)$, and the possibly empty $\mathcal{N}^{m-d_{a'}}_{\mathcal{R}^-}(b')$ and $\mathcal{N}^{m-1-d_{a'}-d_{b'a''}}_{\mathcal{R}^-}(b'')$, with an E-edge joining b in the first to a in the second, an E-edge joining a' in the first to b' in the third, and an E-edge joining a'' in the third to b'' in the fourth (in the case they are non-empty).

We claim that any two of these four neighborhoods are at distance ≥ 2 in $\mathcal{R}^-$, hence they are fully independent: no overlap, and no edge between any two of them except (possibly) for $E(a, b)$, $E(a', b')$ and $E(a'', b'')$.

Indeed, suppose (the other pairs of neighborhoods are treated similarly) that $\mathcal{N}^{m+1}_{\mathcal{R}^-}(x)$ and $\mathcal{N}^{m-d_b}_{\mathcal{R}^-}(a)$ are at distance ≤ 1. Then $\mathrm{dist}_{\mathcal{R}^-}(a, x) \leq 2m + 2 - d_b$, hence $\mathrm{dist}_{\mathcal{R}^-}(a,b) \leq 2m+2$, which contradicts Lemma 4.3.8 for a and b (recall that $\{a,b\}$ and $\mathrm{End}(\mathcal{R})$ are $(2m+3)$-distant in $\mathcal{S}\mathrm{upp}_0(\mathcal{R})$).

Likewise, we can partition $\mathcal{N}^{m+1}_{\mathcal{R}'}(x)$ into

$$\mathcal{N}^{m+1}_{\mathcal{R}^-}(x),\ \mathcal{N}^{m-d_b}_{\mathcal{R}^-}(a''),\ \mathcal{N}^{m-1-d_b-d_{b'a''}}_{\mathcal{R}^-}(a) \text{ and } \mathcal{N}^{m-d_{a'}}_{\mathcal{R}^-}(b'')\,.$$

Because $\mathrm{tp}^m_{\mathcal{R}}(a,b) = \mathrm{tp}^m_{\mathcal{R}}(a',b') = \mathrm{tp}^m_{\mathcal{R}}(a'',b'')$, $\mathcal{N}^{m-d_b}_{\mathcal{R}^-}(a)$ is isomorphic to the union of $\mathcal{N}^{m-d_b}_{\mathcal{R}^-}(a'')$ and $\mathcal{N}^{m-1-d_b-d_{b'a''}}_{\mathcal{R}^-}(a)$ with an E-edge joining b' in the first and a in the second (if they are both nonempty).

Similarly, the union of $\mathcal{N}^{m-d_{a'}}_{\mathcal{R}^-}(b')$ and $\mathcal{N}^{m-1-d_{a'}-d_{b'a''}}_{\mathcal{R}^-}(b'')$ with an E-edge joining a'' in the first and b'' in the second (if they are both nonempty) is isomorphic to $\mathcal{N}^{m-d_{a'}}_{\mathcal{R}^-}(b'')$.

Hence $\mathcal{N}^{m+1}_{\mathcal{R}}(x) \simeq \mathcal{N}^{m+1}_{\mathcal{R}'}(x)$.

We now exhibit a $s(\alpha)$ such that for every $m \geq s(\alpha)$, m-guarded contiguous-segment-E-swaps preserve $\equiv^{<\text{-inv FO}}_{\alpha}$.

We will first set $N \in \mathbb{N}$ instead of $s(\alpha)$, that will be sufficient for most cases. Then, we will define $s(\alpha) \geq N$ which will work for all cases.

Recall the function o^{Σ}_p needed for Lemma 4.3.2, and consider $n := o^{\Sigma}_2(\alpha+c)$ where c is to be chosen later on, and $\Sigma := P_\sigma \cup \{E, S, P_1, P_4\}$ where P_1 and P_4 are unary. We distinguish between several cases depending on whether a, a' and a'' are close or not, where "close" is relative to n.

1. Assume first that $\mathrm{tp}^n_{\mathcal{R}^-}(b,a') = \mathrm{tp}^n_{\mathcal{R}^-}(b',a'')$.

   This case covers the instances where $[b,a']^{\mathcal{R}}_n \simeq [b',a'']^{\mathcal{R}}_n$, as well as those where $|[a,a']|$ and $|[a',a'']|$ are both $> 2n+2$.

   Consider the extension $\mathcal{R}^\star$ of $\mathcal{R}^-$ to Σ where $P_1^{\mathcal{R}^\star} := \{a\}$ and $P_4^{\mathcal{R}^\star} := \{b''\}$. Since $P_1^{\mathcal{R}^\star}$ and $P_4^{\mathcal{R}^\star}$ are at distance $> n$ from $\{b,a',b',a''\}$ (this is guaranteed by Lemma 4.3.8, because we will make sure that $s(\alpha) \geq n$), $\mathrm{tp}^n_{\mathcal{R}^\star}(b,a') = \mathrm{tp}^n_{\mathcal{R}^\star}(b',a'')$.

   Hence, we can apply Lemma 4.3.2, and get two orders $<_{ba'b'a''}$ and $<_{b'a''ba'}$ such that $(\mathcal{R}^\star, <_{ba'b'a''}) \equiv^{\mathrm{FO}}_{\alpha+c} (\mathcal{R}^\star, <_{b'a''ba'})$.

   Now, consider the FO-interpretation that adds a symmetrical E-edge between u and v if either

   - $P_1(u)$ and v is the first element of $<$
   - or u is the second element of $<$ and v is either its third one
   - or u is the fourth element of $<$ and $P_4(v)$,

   and then forgets about P_1 and P_4.

   Take c to be the depth of this FO-interpretation (which has arity 1).

   Note that the result of this FO-interpretation on $(\mathcal{R}^\star, <_{ba'b'a''})$ is an ordered extension of $\mathcal{R}$ and that its result on $(\mathcal{R}^\star, <_{b'a''ba'})$ is an ordered extension of $\mathcal{R}'$.

   This entails $\mathcal{R}' \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{R}$.

2. Assume next that $[b', a'']_n^{\mathcal{R}}$ can be decomposed as

$$[b', a_1]_n^{\mathcal{R}} \cdot [b_1, a_2]_n^{\mathcal{R}} \cdots [b_k, a'']_n^{\mathcal{R}},$$

where each of these enriched segments is isomorphic to $[b, a']_n^{\mathcal{R}}$.

We can then apply $k+1$ times Case 1 and obtain $\mathcal{R}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$ as desired.

3. From now on, $N \geq l(2n+2)+n$ for a large enough l to be chosen later on.

As we are not in Case 1, we can restrict our study to the cases where $|[a, a']| \leq 2n+2$ (the cases where $|[a', a'']| \leq 2n+2$ can be treated similarly).

We need the following claim, which relies on the Lyndon-Schützenberger Theorem [28].

Claim 4.3.10. *Let $n \in \mathbb{N}$ and $\mathcal{R} \in$ quasi-$\mathbb{H}_\sigma^{n+1}$.*

Let $a_1, b_1, a_2, b_2, a_3, b_3, a_4, b_4$ appear in that order in a single thread of $\mathcal{R}$, such that $E(a_1, b_1), E(a_2, b_2), E(a_3, b_3)$ and $E(a_4, b_4)$.

Suppose that $[b_1, a_2]_n^{\mathcal{R}} \simeq [b_3, a_4]_n^{\mathcal{R}}$ and $[b_1, a_3]_n^{\mathcal{R}} \simeq [b_2, a_4]_n^{\mathcal{R}}$.

Then there exist decompositions $[b_1, a_2]_n^{\mathcal{R}} = \mathcal{U}_1 \cdots \mathcal{U}_p$, $[b_2, a_3]_n^{\mathcal{R}} = \mathcal{V}_1 \cdots \mathcal{V}_q$, and $[b_3, a_4]_n^{\mathcal{R}} = \mathcal{W}_1 \cdots \mathcal{W}_p$, where all the $\mathcal{U}_i, \mathcal{V}_i$ and $\mathcal{W}_i$ are isomorphic.

Proof. Consider Θ_n which maps $[x, y]_n^{\mathcal{R}}$ to the word $[x, y]$ where each element is colored with its n-neighborhood type in $[x, y]_n^{\mathcal{R}}$.

Let $u := \Theta_n([b_1, a_2]_n^{\mathcal{R}})$, $v := \Theta_n([b_2, a_3]_n^{\mathcal{R}})$ and $w := \Theta_n([b_3, a_4]_n^{\mathcal{R}})$.

The hypothesis guarantee $u = w$ and $uv = vw$. Hence $uv = vu$.

By Lyndon-Schtzenberger Theorem, there must exist a word a and integers p, q such that $u = w = a^p$ and $v = a^q$ [28].

We can decompose $[b_1, a_2]_n^{\mathcal{R}}, [b_2, a_3]_n^{\mathcal{R}}$ and $[b_3, a_4]_n^{\mathcal{R}}$ alongside those decompositions of u, v and w, to get $[b_1, a_2]_n^{\mathcal{R}} = \mathcal{U}_1 \cdots \mathcal{U}_p$, $[b_2, a_3]_n^{\mathcal{R}} = \mathcal{V}_1 \cdots \mathcal{V}_q$, and $[b_3, a_4]_n^{\mathcal{R}} = \mathcal{W}_1 \cdots \mathcal{W}_p$, where all the $\mathcal{U}_i, \mathcal{V}_i$ and $\mathcal{W}_i$ are mapped to a by Θ_n, hence are isomorphic. □

Let ϕ be an isomorphism between the N-neighborhood of (a, b) and that of (a', b').

As $|[a, a']| \leq 2n+2$, a' and b' are in the N-neighborhood of (a, b): set $x_0 := a'$ and $y_0 := b'$. Construct by induction $x_{i+1} := \phi(x_i)$ and $y_{i+1} := \phi(y_i)$ until $i > l$. Our choice of N ensures that x_i and y_i are well defined as x_{i-1} and y_{i-1} remain in the N-neighborhood of (a, b). For all $j \leq l$, $\mathcal{X}_j := [y_{j-1}, x_j]_n^{\mathcal{R}}$ is isomorphic to $[b, a']_n^{\mathcal{R}}$.

Likewise, starting from (a'', b'') instead of (a', b'), we show that there exist $x'_1, y'_1, \cdots, x'_l, y'_l$ such that for $j \in [1, l]$ (and with the convention that $x'_0 = a''$), $\mathcal{X}'_j := [y'_j, x'_{j-1}]_n^{\mathcal{R}}$ is isomorphic to $[b, a']_n^{\mathcal{R}}$.

We distinguish several cases.

(a) Suppose that $|[a', a'']| \geq 2N$. This ensures that all the $(x_i)_{i\geq 1}$, $(y_i)_{i\geq 0}$, $(x'_i)_{i\geq 0}$ and $(y'_i)_{i\geq 1}$ belong to $[b', a'']$.
We can decompose $[b, a'']^{\mathcal{R}}_n$ as

$$[b,a']^{\mathcal{R}}_n \cdot \underbrace{[y_0,x_1]^{\mathcal{R}}_n}_{\mathcal{X}_1} \cdots \underbrace{[y_{l-1},x_l]^{\mathcal{R}}_n}_{\mathcal{X}_l} \cdot [y_l,x'_l]^{\mathcal{R}}_n \cdot \underbrace{[y'_l,x'_{l-1}]^{\mathcal{R}}_n}_{\mathcal{X}'_l} \cdot \underbrace{[y'_{l-1},x'_{l-2}]^{\mathcal{R}}_n}_{\mathcal{X}'_{l-1}} \cdots \underbrace{[y'_1,x'_0]^{\mathcal{R}}_n}_{\mathcal{X}'_1} .$$

Let now $\mathcal{R}_1$ be the n-guarded contiguous-segment-E-swap between $[b, x'_{l-1}]$ and $[y'_{l-1}, a'']$ in $\mathcal{R}$ (recall that $a'' = x'_0$).
If l is chosen large enough, namely $l \geq 2n + 4$, this swap falls in Case 1 of this Lemma and therefore $\mathcal{R}_1 \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{R}$.
In $\mathcal{R}_1$, $[y'_{l-1}, x'_{l-1}]^{\mathcal{R}_1}_n$ (that is, the segment strictly between a and b'') is decomposed as

$$\underbrace{[y'_{l-1},x'_{l-2}]^{\mathcal{R}}_n}_{\mathcal{X}'_{l-1}} \cdots \underbrace{[y'_1,x'_0]^{\mathcal{R}}_n}_{\mathcal{X}'_1} \cdot [b,a']^{\mathcal{R}}_n \cdot \underbrace{[y_0,x_1]^{\mathcal{R}}_n}_{\mathcal{X}_1} \cdots \underbrace{[y_{l-1},x_l]^{\mathcal{R}}_n}_{\mathcal{X}_l} \cdot [y_l,x'_l]^{\mathcal{R}}_n \cdot \underbrace{[y'_l,x'_{l-1}]^{\mathcal{R}}_n}_{\mathcal{X}'_l} .$$

Now, let $\mathcal{R}_2$ be the n-guarded contiguous-segment-E-swap between $[y'_{l-1}, a']$ and $[y_0, x'_{l-1}]$ in $\mathcal{R}_1$.
By choice of l, this swap falls again in Case 1 of this Lemma. Thus, $\mathcal{R}_2 \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{R}_1$.
Observe now that in $\mathcal{R}_2$, $[b, a'']^{\mathcal{R}_2}_n$ (the segment strictly between a and b'') is decomposed as

$$\underbrace{[y_0,x_1]^{\mathcal{R}}_n}_{\mathcal{X}_1} \cdots \underbrace{[y_{l-1},x_l]^{\mathcal{R}}_n}_{\mathcal{X}_l} \cdot [y_l,x'_l]^{\mathcal{R}}_n \cdot \underbrace{[y'_l,x'_{l-1}]^{\mathcal{R}}_n}_{\mathcal{X}'_l} \cdot \underbrace{[y'_{l-1},x'_{l-2}]^{\mathcal{R}}_n}_{\mathcal{X}'_{l-1}} \cdots \underbrace{[y'_1,x'_0]^{\mathcal{R}}_n}_{\mathcal{X}'_1} \cdot [b,a']^{\mathcal{R}}_n ,$$

that is, $[b', a'']^{\mathcal{R}}_n \cdot [b, a']^{\mathcal{R}}_n$. Hence $\mathcal{R}_2 = \mathcal{R}'$, and we get $\mathcal{R}' \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{R}$.

(b) Suppose now that $|[a', a'']| < 2N$. Set $s(\alpha) := (2N + 1)(2n + 2) + n$. Just as before (by replacing l with $2N + 1$), we define $x_0, y_0, \cdots, x_{2N+1}, y_{2N+1}$ and $x'_0, y'_0, \cdots, x'_{2N+1}, y'_{2N+1}$, and accordingly, $\mathcal{X}_1, \cdots, \mathcal{X}_{2N+1}$ and $\mathcal{X}'_1, \cdots, \mathcal{X}'_{2N+1}$ that all are isomorphic to $[b, a']^{\mathcal{R}}_n$.
Not all of the $(x'_i)_{0\leq i\leq 2N}$ can be in $[b', a'']$. Let k be the smallest index such that $x'_i \notin [b', a'']$ (we know that $1 \leq k \leq 2N$).
If $x'_k = a'$, we can conclude using Case 2.
Otherwise, $a, b, x'_k, y'_k, a', b', x'_{k-1}, y'_{k-1}$ must appear in that order in the thread.
$[b, a']^{\mathcal{R}}_n \simeq [y'_k, x'_{k-1}]^{\mathcal{R}}_n$ by definition.
To see that $[b, x'_k]^{\mathcal{R}}_n \simeq [b', x'_{k-1}]^{\mathcal{R}}_n$, consider the restriction of an isomorphism between $\mathcal{X}'_k$ and $\mathcal{X}'_{k+1}$ to the final segments of length $|[b', x'_{k-1}]| = |[b, x'_k]|$.
We can now apply Claim 4.3.10, and get decompositions $[b, x'_k]^{\mathcal{R}}_n = \mathcal{U}_1 \cdots \mathcal{U}_p$, $[y'_k, a']^{\mathcal{R}}_n = \mathcal{V}_1 \cdots \mathcal{V}_q$, and $[b', x'_{k-1}]^{\mathcal{R}}_n = \mathcal{W}_1 \cdots \mathcal{W}_p$, where all the $\mathcal{U}_i, \mathcal{V}_i$ and $\mathcal{W}_i$ are isomorphic.

Hence, $[b,a']_n^{\mathcal{R}}$ can be decomposed as $\mathcal{U}_1 \cdots \mathcal{U}_p \cdot \mathcal{V}_1 \cdots \mathcal{V}_q$, and such a decomposition can be transposed onto each $\mathcal{X}'_i, 0 < i < k$, as $\mathcal{X}'_i = \mathcal{Y}_1^i \cdots \mathcal{Y}_{p+q}^i$, where all the $\mathcal{Y}_j^i$, the $\mathcal{U}_i$, the $\mathcal{V}_i$ and the $\mathcal{W}_i$ are isomorphic.
We can now decompose $[b,a'']_n^{\mathcal{R}}$ as

$$\underbrace{\mathcal{U}_1 \cdots \mathcal{U}_p \cdot \mathcal{V}_1 \cdots \mathcal{V}_q}_{[b,a']_n^{\mathcal{R}}} \cdot \underbrace{\underbrace{\mathcal{W}_1 \cdots \mathcal{W}_p}_{[b',x'_{k-1}]_n^{\mathcal{R}}} \cdot \underbrace{\mathcal{Y}_1^{k-1} \cdots \mathcal{Y}_{p+q}^{k-1}}_{\mathcal{X}'_{k-1}} \cdots \underbrace{\mathcal{Y}_1^{1} \cdots \mathcal{Y}_{p+q}^{1}}_{\mathcal{X}'_1}}_{[b',a'']_n^{\mathcal{R}}} .$$

Now, we can use Case 2 to swap $\mathcal{V}_q$ with $[b',a'']_n^{\mathcal{R}}$: $\mathcal{R}_1 \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$, where in $\mathcal{R}_1$, the segment strictly between a and b'' is

$$\mathcal{U}_1 \cdots \mathcal{U}_p \cdot \mathcal{V}_1 \cdots \mathcal{V}_{q-1} \cdot \mathcal{W}_1 \cdots \mathcal{W}_p \cdot \mathcal{Y}_1^{k-1} \cdots \mathcal{Y}_{p+q}^{k-1} \cdots \mathcal{Y}_1^{1} \cdots \mathcal{Y}_{p+q}^{1} \cdot \mathcal{V}_q .$$

Repeating this operation $p+q-1$ times allows us to conclude that $\mathcal{R}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$.

Proof of Lemma 4.3.5 for segment-E-swaps

The proof that $\forall x \in R, \text{tp}_{\mathcal{R}'}^{m+1}(x) = \text{tp}_{\mathcal{R}'}^{m+1}(x)$ is done as for the contiguous-segment-E-swaps.

Let's now find $s(\alpha) \in \mathbb{N}$ that guarantees the $\equiv_\alpha^{<\text{-inv FO}}$ invariance of any m-guarded segment-E-swaps, for $m \geq s(\alpha)$.

Let $n := o_2^\Sigma(\alpha + c)$ where c is the depth of some FO-interpretation to be specified later on, and $\Sigma := P_\sigma \cup \{E, S, P_1, P_2, P_3, P_4\}$ where P_1, P_2, P_3 and P_4 are unary.

Let $\mathcal{R}^- := \mathcal{R} \setminus \{E(a,b), E(a',b'), E(c,d), E(c',d')\}$.

1. Assume first that $\text{tp}_{\mathcal{R}^-}^n(b,c) = \text{tp}_{\mathcal{R}^-}^n(b',c')$.

   This case covers the instances where $[b,c]_n^{\mathcal{R}} \simeq [b',c']_n^{\mathcal{R}}$, as well as those where $|[a,c]|$ and $|[a',c']|$ both are $> 2n+2$.

   Consider the extension $\mathcal{R}^\star$ of $\mathcal{R}^-$ to Σ where $P_1^{\mathcal{R}^\star} := \{a\}$, $P_2^{\mathcal{R}^\star} := \{d\}$, $P_3^{\mathcal{R}^\star} := \{a'\}$ and $P_4^{\mathcal{R}^\star} := \{d'\}$. Since $P_1^{\mathcal{R}^\star}, P_2^{\mathcal{R}^\star}, P_3^{\mathcal{R}^\star}$ and $P_4^{\mathcal{R}^\star}$ are at distance $> n$ from $\{b, c, b', c'\}$ (this is guaranteed by Lemma 4.3.8, because we will make sure that $s(\alpha) \geq n$), $\text{tp}_{\mathcal{R}^\star}^n(b,c) = \text{tp}_{\mathcal{R}^\star}^n(b',c')$.

   Hence, we can apply Lemma 4.3.2, and get two orders $<_{bcb'c'}$ and $<_{b'c'bc}$ such that $(\mathcal{R}^\star, <_{bcb'c'}) \equiv_{\alpha+c}^{\text{FO}} (\mathcal{R}^\star, <_{b'c'bc})$.

   Now, consider the FO-interpretation that adds a symmetrical E-edge between u and v if either

   - $P_1(u)$ and v is the first element of $<$
   - or u is the second element of $<$ and $P_2(v)$
   - or $P_3(u)$ and v is the third element of $<$

- or u is the fourth element of $<$ and $P_4(v)$,

and then forgets about P_1, P_2, P_3 and P_4.

Take c to be the depth of this FO-interpretation (which has arity 1).

Note that the result of this FO-interpretation on $(\mathcal{R}^\star, <_{bcb'c'})$ is an ordered extension of $\mathcal{R}$ and that its result on $(\mathcal{R}^\star, <_{b'c'bc})$ is an ordered extension of $\mathcal{R}'$.

This entails $\mathcal{R}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$.

2. We can now, without loss of generality, assume that $|[a,c]| \leq 2n+2$.

   Let s be the threshold $s(\alpha)$ from the proof of contiguous-segment-E-swaps. Let us increase that threshold for it to account for segment-E-swaps: set $s(\alpha) := (2n+2) + M$ with $M := \max(s, n)$.

   Consider an isomorphism φ from $\mathcal{N}_{\mathcal{R}}^{s(\alpha)}(a,b)$ to $\mathcal{N}_{\mathcal{R}}^{s(\alpha)}(a',b')$.

   By choice of $s(\alpha)$, $\text{tp}_{\mathcal{R}}^M(c,d) = \text{tp}_{\mathcal{R}}^M(\varphi(c), \varphi(d))$.

   Since $(\varphi(c), \varphi(d)) \neq (c', d')$ (for otherwise, we would be in the Case 1), there are only two subcases to consider:

   - if $a', b', \varphi(c), \varphi(d), c', d'$ appear in that order, i.e. the segment strictly between a and d' can be decomposed as

     $$[b,c]_n^{\mathcal{R}} \cdot [d,a']_n^{\mathcal{R}} \cdot [b', \varphi(c)]_n^{\mathcal{R}} \cdot [\varphi(d), c']_n^{\mathcal{R}} .$$

     Let $\mathcal{R}_1$ be the M-guarded segment-E-swap between $[b,c]$ and $[b', \varphi(c)]$ in $\mathcal{R}$.

     This swap falls under the scope of Case 1 since $M \geq n$, hence $\mathcal{R}_1 \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$ and $\forall z \in R$, $\text{tp}_{\mathcal{R}_1}^{M+1}(z) = \text{tp}_{\mathcal{R}}^{M+1}(z)$. In $\mathcal{R}_1$, the segment strictly between a and d' can be decomposed as

     $$[b', \varphi(c)]_n^{\mathcal{R}} \cdot [d,a']_n^{\mathcal{R}} \cdot [b,c]_n^{\mathcal{R}} \cdot [\varphi(d), c']_n^{\mathcal{R}} .$$

     Hence we are in the conditions (since $M \geq s$) to apply Lemma 4.3.5 in the case of the M-guarded contiguous-segment-E-swap between $[d,c]$ and $[\varphi(d), c']$ in $\mathcal{R}_1$.

     We get $\mathcal{R}_2 \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}_1$. In $\mathcal{R}_2$, the segment strictly between a and d' can be decomposed as

     $$[b', \varphi(c)]_n^{\mathcal{R}} \cdot [\varphi(d), c']_n^{\mathcal{R}} \cdot [d,a']_n^{\mathcal{R}} \cdot [b,c]_n^{\mathcal{R}} = [b', c']_n^{\mathcal{R}} \cdot [d,a']_n^{\mathcal{R}} \cdot [b,c]_n^{\mathcal{R}} .$$

     That is, $\mathcal{R}_2 = \mathcal{R}'$, and we get $\mathcal{R}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$ as desired.

   - if $a', b', c', d', \varphi(c), \varphi(d)$ appear in that order, i.e. the segment strictly between a and $\varphi(d)$ can be decomposed as

     $$[b,c]_n^{\mathcal{R}} \cdot [d,a']_n^{\mathcal{R}} \cdot [b', c']_n^{\mathcal{R}} \cdot [d', \varphi(c)]_n^{\mathcal{R}} .$$

     Let $\mathcal{R}_1$ be the M-guarded segment-E-swap between $[b,c]$ and $[b', \varphi(c)]$ in $\mathcal{R}$.

This swap falls under the scope of Case 1 since $M \geq n$, hence $\mathcal{R}_1 \equiv_{\alpha}^{<\text{-inv FO}} \mathcal{R}$ and $\forall z \in R$, $\mathrm{tp}_{\mathcal{R}_1}^{M+1}(z) = \mathrm{tp}_{\mathcal{R}}^{M+1}(z)$. In $\mathcal{R}_1$, the segment strictly between a and $\varphi(d)$ can be decomposed as

$$[b', c']_n^{\mathcal{R}} \cdot [d', \varphi(c)]_n^{\mathcal{R}} \cdot [d, a']_n^{\mathcal{R}} \cdot [b, c]_n^{\mathcal{R}} .$$

Hence we are in the conditions (since $M \geq s$) to apply Lemma 4.3.5 in the case of the M-guarded contiguous-segment-E-swap between $[d', \varphi(c)]$ and $[d, c]$ in $\mathcal{R}_1$.
We get $\mathcal{R}_2 \equiv_{\alpha}^{<\text{-inv FO}} \mathcal{R}_1$. In $\mathcal{R}_2$, the segment strictly between a and $\varphi(d)$ can be decomposed as

$$[b', c']_n^{\mathcal{R}} \cdot [d, a']_n^{\mathcal{R}} \cdot [b, c]_n^{\mathcal{R}} \cdot [d', \varphi(c)]_n^{\mathcal{R}} .$$

That is, the segment strictly between a and d' is

$$[b', c']_n^{\mathcal{R}} \cdot [d, a']_n^{\mathcal{R}} \cdot [b, c]_n^{\mathcal{R}} .$$

Hence $\mathcal{R}_2 = \mathcal{R}'$, and we get $\mathcal{R}' \equiv_{\alpha}^{<\text{-inv FO}} \mathcal{R}$ as desired.

4.3.3 Pumping

The next operation makes use of the fact that $<$-inv FO $\subseteq$ MSO over hollow trees. Hence our hollow trees can be "pumped" in order to duplicate some of their parts.

Proposition 4.3.11. $\forall \alpha, n, d \in \mathbb{N}, \exists M \in \mathbb{N}, \forall D \in \mathbb{N}$,
for every $\mathcal{P}, \mathcal{Q} \in \mathbb{H}_\sigma$ such that $\mathcal{P} \equiv_M^{\text{FO}} \mathcal{Q}$,
there exists $\mathcal{Q}' \in \mathbb{H}_\sigma$ such that

$$\mathcal{Q}' \equiv_{\alpha}^{<\text{-inv FO}} \mathcal{Q} \text{ and } [\![\mathcal{E}_{n+1}(\mathcal{P})]\!] \leq_d^D [\![\mathcal{E}_{n+1}(\mathcal{Q}')]\!] .$$

The proof is a pumping argument: by setting M large enough, we make sure in FO that if a $(n+1)$-neighborhood type has more occurrences in $\mathcal{P}$ than in $\mathcal{Q}$, then it has enough occurrences in $\mathcal{Q}$ so that we can find a context in $\mathcal{Q}$ containing at least one occurrence, and no occurrence of a rare type, such that we can duplicate this context inside $\mathcal{Q}$ without changing its MSO-type.

Proof. Hollow trees have treewidth at most 2, hence Proposition 2.5.3 guarantees that $<$-inv FO $\subseteq$ MSO on $\mathbb{H}_\sigma$.

In particular, there is a $\beta \in \mathbb{N}$ such that $\equiv_\beta^{\text{MSO}}$ subsumes $\equiv_{\alpha}^{<\text{-inv FO}}$. We will construct $\mathcal{Q}'$ such that $\mathcal{Q}' \equiv_\beta^{\text{MSO}} \mathcal{Q}$.

Let $d' > d$ be a number that will be specified during the proof. We choose M large enough to make sure that every $(n+1)$-neighborhood type has the same number of occurrences in $\mathcal{P}$ and in $\mathcal{Q}$ up to a threshold d' (this can be expressed in FO).

We prove the proposition by induction on κ, where κ is the number of $(n+1)$-neighborhood types τ such that $|\mathcal{P}|_\tau \neq |\mathcal{Q}|_\tau$ and $|\mathcal{Q}|_\tau < |\mathcal{P}|_\tau + D$.

If $\kappa = 0$, there is nothing to do as $\mathcal{Q}' := \mathcal{Q}$ fits. Otherwise, let τ be such a type. Notice that because $\mathcal{P} \equiv_M^{\text{FO}} \mathcal{Q}$ we must have $|\mathcal{Q}|_\tau > d'$.

There are two cases to consider.

- Suppose that there exists a thread in $\mathcal{Q}$ which contains at least l nodes $x_1, x_2, \cdots, x_l$ (in that order) having the same $(n+1)$-neighborhood type, whose subtrees each contains at least one node of neighborhood type τ in $\mathcal{Q}$, and such that for every $i < l$, duplicating within the thread the forest below $[x_i, x_{i+1}[$ does not affect the $\equiv_\beta^{\text{MSO}}$ of $\mathcal{Q}$, where l is chosen large enough so that there exists $i < l$ such that the forest below $[x_i, x_{i+1}[$ doesn't contain any occurrence of a $(n+1)$-neighborhood type τ' such that $|\mathcal{Q}|_{\tau'} \leq d$.

  Then we construct $\mathcal{Q}'$ from $\mathcal{Q}$ by duplicating the forest below $[x_i, x_{i+1}[$ as many times as necessary to have enough nodes of type τ. This decreases κ and guarantees that $\mathcal{Q}' \equiv_\beta^{\text{MSO}} \mathcal{Q}$ and we can conclude by induction.

- Assume now that there is a chain for the ancestor relation $x_1, x_2, \cdots, x_l$ having the same $(n+1)$-neighborhood type such that each of the contexts $\mathcal{C}_\mathcal{P}(x_i, x_{i+1})$ (we use here the notations introduced for Lemma 4.6.4 to denote the context between x_i and x_{i+1}, but they should be transparent enough) contains at least one node of type τ and $\forall i, j, \mathcal{S}_\mathcal{P}(x_i) \equiv_\beta^{\text{MSO}} \mathcal{S}_\mathcal{P}(x_j)$ (the subtrees at x_i and x_j), where l is large enough to guarantee the existence of some $i < l$ such that $\mathcal{C}_\mathcal{P}(x_i, x_{i+1})$ contains no node of any neighborhood type τ' such that $|\mathcal{Q}|_{\tau'} \leq d$.

  Let $\mathcal{Q}' := \mathcal{P}_\mathcal{P}(x_i) \cdot \mathcal{C}_\mathcal{P}(x_i, x_{i+1})^k \cdot \mathcal{S}_\mathcal{P}(x_{i+1})$ (that is, we've duplicated k times the context between x_i and x_{i+1}) with k large enough so we have enough nodes of neighborhood type τ. We have $\mathcal{Q}' \equiv_\beta^{\text{MSO}} \mathcal{Q}$ and κ has decreased by 1: we can conclude by induction.

It remains to fix d' large enough so that one of the two cases above must hold. □

4.4 Inclusion and pseudo-inclusion

Recall that our ultimate goal is to show that if two hollow trees agree on the same FO sentences of quantifier rank $f(\alpha)$ then they agree on all $<$-inv FO sentences of quantifier rank α. For this, we will show that if $\mathcal{P}$ and $\mathcal{Q}$ are hollow trees that agree on all FO sentences of quantifier rank $f(\alpha)$ then we can use operations such as the swap operations described in Section 4.3 to transform $\mathcal{Q}$ into $\mathcal{P}$. As these operations preserve $<$-inv FO we get the desired result.

In this section we perform the first step towards transforming $\mathcal{Q}$ into $\mathcal{P}$. We show that using the swap operations we can transform $\mathcal{Q}$ into $\mathcal{Q}'$ so that $\mathcal{Q}'$ "includes" $\mathcal{P}$. The resulting structure $\mathcal{Q}'$ will be a hollow quasitree. In the next sections we will continue the transformation and remove from $\mathcal{Q}'$ all the extra material it contains, deriving $\mathcal{P}$.

In order to define what we mean by "inclusion" we need the notion of an n-abstract context of a hollow quasitree. Intuitively this is an S-stable n-enriched substructure. More formally, given a hollow quasitree $\mathcal{T} \in$ quasi-$\mathbb{H}_\sigma^n$ and a set U of its domain that is S-stable, then $\mathcal{C} := \mathcal{T}_{|U}$, together with the function $\text{tp}^n(.)$ that maps $x \in U$ to its n-neighborhood type in $\mathcal{T}$, is called an **n-abstract context** denoted $\mathcal{C} = \text{Ctxt}_n(\mathcal{T}|_U)$. The set of n-abstract contexts is denoted Ctxt_σ^n. Note that $\text{tp}^n(x)$ denotes $\text{tp}_\mathcal{T}^n(x)$ and not $\text{tp}_\mathcal{C}^n(x)$. We need to remember, at least locally, how $\mathcal{C}$ was glued to the rest of $\mathcal{T}$ in order to preserve n-neighborhood types when moving $\mathcal{C}$ to some other place.

We are now ready to define the notion of "inclusion". We actually define both "inclusions" and "pseudo-inclusions". We will need to pseudo-include a hollow quasitree into another (Proposition 4.4.2), and then to include an abstract context into a hollow quasitree (Proposition 4.4.3). Since a hollow k-quasitree $\mathcal{T} \in$ quasi-$\mathbb{H}_\sigma^k$ can be seen as a k-abstract context ($\mathcal{T} = \mathrm{Ctxt}_k(\mathcal{T}|_T)$), we only need to define (pseudo-)inclusions from an abstract context into a hollow quasitree.

Definition 4.4.1. *Let $k \in \mathbb{N}$, $\mathcal{U} \in Ctxt_\sigma^k$ and $\mathcal{Q} \in$ quasi-$\mathbb{H}_\sigma^k$. We say that $h : U \to Q$ is a k-**pseudo-inclusion** if h is injective and for all $x, y, z \in U$ the following is verified:*

1. *$tp_\mathcal{Q}^k(h(x)) = tp^k(x)$,*
2. *if x and y are in the same thread of $\mathcal{U}$ then $h(x)$ and $h(y)$ are also on the same thread of $\mathcal{Q}$ and if moreover $z \in [x, y]$ then $h(z) \in [h(x), h(y)]$,*
3. *if $\mathcal{U} \models E(x, y)$ and t is the E-neighbor of $h(x)$ in $[h(x), h(y)]$ then t is the image of y by an isomorphism (induced by the fact that they share the same k-neighborhood type) between the k-neighborhood of x and that of $h(x)$.*

*If $\mathcal{U} \models E(x, y)$ and $\mathcal{Q} \not\models E(h(x), h(y))$ then $\{x, y\}$ is said to be a **jumping pair** for h, and $tp_Q^{k-1}(h(x), t)$, where t is the E-neighbor of $h(x)$ in $[h(x), h(y)]$, is called its type.*[2]

*A k-pseudo-inclusion is said to be **reduced** if there is at most one jumping pair of a given type.*

*A k-pseudo-inclusion is called a k-**inclusion** if it has no jumping pairs, that is if it preserves E.*

The last condition of pseudo-inclusion is a complication induced by the fact that E is not oriented and that we thus cannot distinguish between the two siblings of a node. It ensures that h preserves the neighborhoods in the right order. We can now state the main result of this section. Note that the precondition that $\mathcal{Q}$ has more realizations for each neighborhood type than $\mathcal{U}$ or $\mathcal{P}$ will not be a problem in view of Proposition 4.3.11. The second proposition is stronger than the first one as it derives inclusion instead of pseudo-inclusion, but it requires the stronger hypothesis that every occurring neighborhood type has strictly more realizations in $\mathcal{Q}$ than in $\mathcal{U}$.

Proposition 4.4.2. *For every $\alpha, m \in \mathbb{N}$, there exists $N \in \mathbb{N}$ such that for every $\mathcal{P}, \mathcal{Q} \in$ quasi-$\mathbb{H}_\sigma^{N+1}$ such that*

$$[\![\mathcal{E}_{N+1}(\mathcal{P})]\!] \leq [\![\mathcal{E}_{N+1}(\mathcal{Q})]\!],$$

there exists $\mathcal{Q}' \in$ quasi-$\mathbb{H}_\sigma^{m+1}$ such that

$$\mathcal{Q}' \equiv_\alpha^{<\text{inv FO}} \mathcal{Q} \text{ and } [\![\mathcal{E}_{m+1}(\mathcal{Q}')]\!] = [\![\mathcal{E}_{m+1}(\mathcal{Q})]\!]$$

and there exists an $(m+1)$-pseudo-inclusion h from $\mathcal{P}$ into $\mathcal{Q}'$.

[2]This is an ease of notation; to be more precise, we should make the type of a jumping pair symmetrical.

Proposition 4.4.3. *For every $\alpha, m \in \mathbb{N}$, there exists $N \in \mathbb{N}$ such that for every $\mathcal{U} \in Ctxt_\sigma^{N+1}$ and every $\mathcal{Q} \in quasi\text{-}\mathbb{H}_\sigma^{N+1}$ such that*

$$[\![\mathcal{E}_{N+1}(\mathcal{U})]\!] < [\![\mathcal{E}_{N+1}(\mathcal{Q})]\!],$$

there exists $\mathcal{Q}' \in quasi\text{-}\mathbb{H}_\sigma^{m+1}$ such that

$$\mathcal{Q}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q} \text{ and } [\![\mathcal{E}_{m+1}(\mathcal{Q}')]\!] = [\![\mathcal{E}_{m+1}(\mathcal{Q})]\!]$$

and such that $\mathcal{U}$ is $(m+1)$-included in $\mathcal{Q}'$.

Recall that

$$[\![\mathcal{E}_{N+1}(\mathcal{U})]\!] < [\![\mathcal{E}_{N+1}(\mathcal{Q})]\!]$$

means that *every* $(N+1)$-neighborhood type has more occurrences in $\mathcal{Q}$ than in $\mathcal{U}$.

Proof. We mainly focus on the proof of Proposition 4.4.2. We will then explain how to modify the proof using the extra hypothesis in order to get inclusion instead of pseudo-inclusion thus proving Proposition 4.4.3.

We modify $\mathcal{Q}$ using E-swaps in order to construct a pseudo-inclusion h from $\mathcal{P}$. This is done step by step, extending the domain of h thread by thread and, inside each thread, from one of its endpoint to the other.

We distinguish between two kinds of threads of $\mathcal{P}$. The short ones will be easily taken care of as they can be completely described in first-order. The long ones will require more work.

In view of Lemma 4.3.5, we assume that $m \geq s(\alpha)$. We set $n := 3m+3$. We will only perform swaps involving nodes at distance (along E) $\geq n-m$ from the endpoints; hence, the "distant from endpoints" conditions of m-guarded E-swaps will always be satisfied.

A thread is short if its length (the distance along E between its two endpoints) is at most $2(n-m)$. By taking N large enough, our hypothesis

$$[\![\mathcal{E}_{N+1}(\mathcal{P})]\!] \leq [\![\mathcal{E}_{N+1}(\mathcal{Q})]\!]$$

guarantees that we can find an injective mapping from the short threads of $\mathcal{P}$ to that of $\mathcal{Q}$, which sends each short thread to one having an isomorphic $(m+1)$-enrichment. We initialize h according to this mapping. It is clear that h is a partial $(m+1)$-pseudo-inclusion mapping.

It remains to extend the domain of h to the long threads.

Let a be an endpoint of a long thread of $\mathcal{P}$: $\text{segtype}_{m+1,\mathcal{P}}^{n-m}(a)$ denotes the isomorphism type of the segment $[a,b]$, where b is the element at distance $n-m$ of a in its thread, and every element is colored with its $(m+1)$-neighborhood type in $\mathcal{P}$. By $\text{End}_{m+1}^{n-m}(\mathcal{P}) \leq \text{End}_{m+1}^{n-m}(\mathcal{Q})$, we mean that every $\text{segtype}_{m+1,.}^{n-m}(.)$ has at least as many occurrences in $\mathcal{Q}$ as in $\mathcal{P}$.

Let $\mathring{\mathcal{S}}_{m+1}^{n-m}(\mathcal{P})$ be the restriction of $\mathcal{S}\text{upp}_{m+1}(\mathcal{P})$ to elements that are at distance $> n-m$ from $\text{End}(\mathcal{P})$.

Every intermediate structure $\mathcal{Q}'$ will verify the following invariant:

$$\begin{cases} \mathcal{Q}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q} \\ [\![\mathcal{E}_{m+1}(\mathcal{Q}')]\!] = [\![\mathcal{E}_{m+1}(\mathcal{Q})]\!] \\ \text{End}_{m+1}^{n-m}(\mathcal{P}) \leq \text{End}_{m+1}^{n-m}(\mathcal{Q}') \text{ and } h \text{ preserves } \text{segtype}_{m+1,.}^{n-m}(.) \\ [\![\mathring{\mathcal{S}}_{m+1}^{n-m}(\mathcal{P})]\!] \leq [\![\mathring{\mathcal{S}}_{m+1}^{n-m}(\mathcal{Q}')]\!] \end{cases} \tag{4.1}$$

As long as N is large enough, the hypothesis guarantees that $\mathcal{Q}$ verifies (4.1).

Assume we have already constructed a partial $(m+1)$-pseudo-inclusion h from $\mathcal{P}$ to $\mathcal{Q}'$ where $\mathcal{Q}'$ verifies (4.1). We show that we can construct a new hollow quasitree $\mathcal{Q}'' \in \text{quasi-}\mathbb{H}_\sigma^{m+1}$ verifying (4.1), using a sequence of E-swaps applied to $\mathcal{Q}'$ in such a way that h can be extended by at least one element of P.

To this end, assume first that the domain of h is the union of a number of threads, each contained in its entirety (this is the case at the beginning). Let $\text{Im}(h)$ denote the image of h. Let t be any thread of $\mathcal{P}$ not in the domain of h and let x be an endpoint of t. We want to extend h in order for its domain to contain x.

(4.1) ensures that there exists an endpoint $y \notin \text{Im}(h)$ of a long thread of $\mathcal{Q}'$ such that $\text{segtype}^{n-m}_{m+1,\mathcal{Q}'}(y) = \text{segtype}^{n-m}_{m+1,\mathcal{P}}(x)$. We don't modify $\mathcal{Q}'$ and extend h by sending every $z \in [x, x']$ to the corresponding $h(z) \in [y, y']$ (where x', y' are the elements at distance $n-m$ of x, y in their threads). Every z and $h(z)$ have the same $(m+1)$-neighborhood type, and h preserves E on $[x, x']$.

By construction, h is a partial $(m+1)$-pseudo-inclusion mapping as desired.

Suppose now that the domain of h contains a set of (entire) threads and the initial segment of a thread t of $\mathcal{P}$, that includes at least the points of t at distance $\leq n-m$ from its endpoint in the domain of h. Let x' be the last element of t in the domain of h and x be the first element of t not in the domain of h. In particular we have $E(x', x)$. Assume furthermore that x is at distance greater than $n-m$ from the other endpoint of t.

Since $[\![\mathring{\mathcal{S}}^{n-m}_{m+1}(\mathcal{P})]\!] \leq [\![\mathring{\mathcal{S}}^{n-m}_{m+1}(\mathcal{Q}')]\!]$, there must exist an element

$$y \in \mathring{\mathcal{S}}^{n-m}_{m+1}(\mathcal{Q}') \setminus \text{Im}(h)$$

having the same $(m+1)$-neighborhood type as x. Let y' be the image of x' by an isomorphism between $\mathcal{N}^{m+1}_{\mathcal{P}}(x)$ and $\mathcal{N}^{m+1}_{\mathcal{Q}'}(y)$, and let $\hat{x}$ be the image of x by an isomorphism between $\mathcal{N}^{m+1}_{\mathcal{P}}(x')$ and $\mathcal{N}^{m+1}_{\mathcal{Q}'}(h(x'))$.

By definition, $\text{tp}^m_{\mathcal{Q}'}(h(x'), \hat{x}) = \text{tp}^m_{\mathcal{Q}'}(y', y)$.

If $y = \hat{x}$, leave $\mathcal{Q}'$ unchanged and let h map x to $\hat{x}$. Otherwise, there are several cases to consider depending on the positions of y and y'.

1. If y is on a thread that does not intersect $\text{Im}(h)$.

   Let $\mathcal{Q}''$ be the m-guarded crossing-E-swap between $h(x')\hat{x}$ and $y'y$ in $\mathcal{Q}'$. Extend h by setting $h(x)$ to y (c.f. Figure 4.12, in which $\text{Im}(h)$ is represented as double lines).

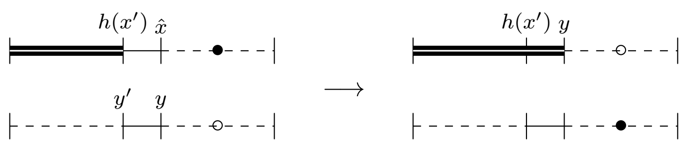

Figure 4.12: An illustration of the progression in the (pseudo-)inclusion (case 1).

2. If y is between $h(z)$ and $h(z')$ where z and z' are consecutive node of the current thread t already in the domain of h and such that y' is between

$h(z)$ and y (that is, they are in the right order for a segment-E-swap), c.f. Figure 4.13.

Let u' and u be the respective E-neighbors of $h(z)$ and $h(z')$ in $[h(z), h(z')]$. h being a pseudo-inclusion, $\mathrm{tp}^m_{\mathcal{Q}'}(h(z), u') = \mathrm{tp}^m_{\mathcal{Q}'}(u, h(z'))$.

Let now $\mathcal{Q}''$ be the m-guarded segment-E-swap between $[u', y']$ and $[h(z'), h(x')]$ in $\mathcal{Q}'$, and extend h by setting $h(x)$ to y.

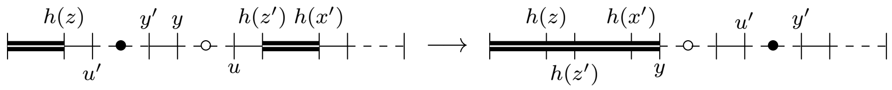

Figure 4.13: An illustration of the progression in the (pseudo-)inclusion (case 2).

3. If y is between $h(z)$ and $h(z')$ where z and z' are consecutive nodes of the current thread t already in the domain of h and such that y' is between y and $h(z')$ (that is, they are not in the right order for a segment-E-swap), c.f. Figure 4.14. This means that $y, y', h(x'), \hat{x}$ appear in that order.

   Let $\mathcal{R}$ be the m-guarded mirror-E-swap at $[y', h(x')]$ in $\mathcal{Q}'$.

   In $\mathcal{R}$, $h(z), u', h(z'), u$ now appear in that order.

   Let $\mathcal{Q}''$ be the m-guarded mirror-E-swap at $[u', h(z')]$ in $\mathcal{R}$ and extend h by setting $h(x)$ to y.

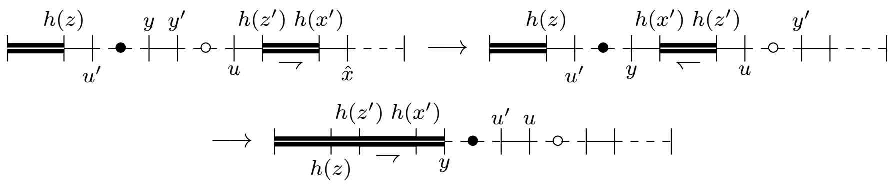

Figure 4.14: An illustration of the progression in the (pseudo-)inclusion (case 3).

4. If y is between $h(z)$ and $h(z')$ where z and z' are consecutive node in some thread different from t already in the domain of h (c.f. Figure 4.15).

   Let $\mathcal{R}$ be the m-guarded crossing-E-swap between $y'y$ and $h(x')\hat{x}$ in $\mathcal{Q}'$.

   Let $\mathcal{Q}''$ be the m-guarded crossing-E-swap between $h(z)u'$ and $uh(z')$ in $\mathcal{R}$, and extend h by setting $h(x)$ to y.

5. If y is on the same thread as $h(x')$, such that $h(x'), \hat{x}, y, y'$ appear in that order (c.f. Figure 4.16).

   Then let $\mathcal{Q}''$ be the m-guarded mirror-E-swap at $[\hat{x}, y]$ in $\mathcal{Q}'$ and extend h by setting $h(x)$ to y.

6. Finally if y is on the same thread as $h(x')$ but $h(x'), \hat{x}, y', y$ appear in that order.

   This is the case where we cannot achieve inclusion without extra hypothesis. For Proposition 4.4.2, we simply allow a "jump" and set $h(x)$ to y without changing $\mathcal{Q}'$.

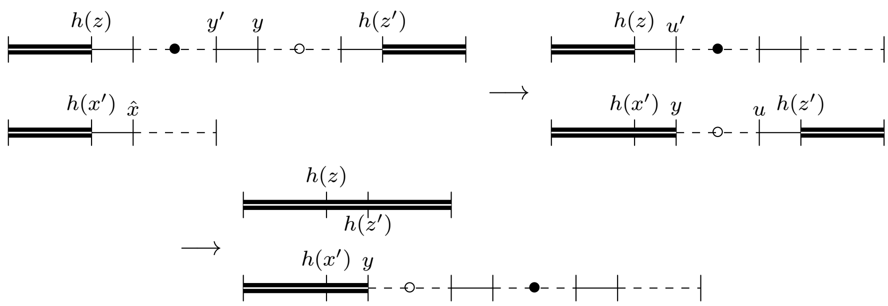

Figure 4.15: An illustration of the progression in the (pseudo-)inclusion (case 4).

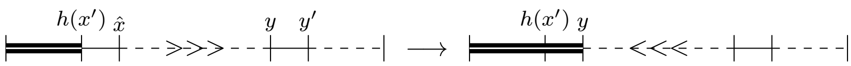

Figure 4.16: An illustration of the progression in the (pseudo-)inclusion (case 5).

In the previous case analysis, in order to perform E-swaps, it was important for x (and therefore y) to be far away from the endpoint e of t that is not already in the domain of h. In order to conclude the proof of Proposition 4.4.3, it remains to consider the case where x is at distance $n-m$ from e.

By hypothesis, there exists an endpoint a outside of $\mathrm{Im}(h)$ of a long thread such that $\mathrm{segtype}^{n-m}_{m+1,\mathcal{Q}'}(a) = \mathrm{segtype}^{n-m}_{m+1,\mathcal{P}}(e)$. Let ξ be the isomorphism between $[e,x]$ and $[a,y]$, where y is the element at distance $n-m$ of a in its thread.

If a is not on the same thread as $h(x')$, let y' be the E-neighbor of y not in $[y,a]$. We let $\mathcal{Q}''$ be the m-guarded crossing-E-swap between $h(x')\hat{x}$ and $y'y$ in $\mathcal{Q}'$ and extend h by setting $h(u)$ to $\xi(u)$ for all u in $[x,e]$.

Otherwise, we don't modify $\mathcal{Q}'$ and simply extend h by setting $h(u)$ to $\xi(u)$ for every u in $[x,e]$. Notice that there may be a jump between $h(x')$ and $h(x)$.

This concludes the proof of Proposition 4.4.2. We now move to the proof of Proposition 4.4.3.

We decompose U as $P \uplus V$, where P is the union of the threads of $\mathcal{U}$ whose endpoints were endpoints in the structure from which $\mathcal{U}$ is derived (that is, their neighborhood type is a type of endpoint). We let $\mathcal{P}$ be $(U)_{|P}$ and proceed as above with the threads of $\mathcal{P}$.

It all works as above except for the two cases where we introduced a jump. Consider again the situation of Case 6. Our extra cardinality hypothesis ensures that there is a $z \neq y$ verifying the same conditions as y (otherwise we would be in a previous case). Assume without loss of generality that $h(x'), y, z$ appear in that order (c.f. Figure 4.17).. Set $\mathcal{Q}''$ to be the m-guarded contiguous-segment-E-swap between $[\hat{x}, y']$ and $[y, z']$ in $\mathcal{Q}'$, and extend h by setting $h(x)$ to y. h is now an inclusion.

We also introduced a jump when extending h to the endpoint of some thread. But the cardinality condition ensures that we have two endpoints $a_1 \neq a_2$

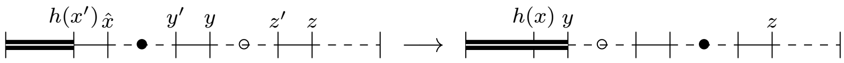

Figure 4.17: An illustration of the progression in the inclusion (case 6, for Proposition 4.4.3).

outside of $\mathrm{Im}(h)$ such that

$$\mathrm{segtype}^{n-m}_{m+1,\mathcal{Q}'}(a_1) = \mathrm{segtype}^{n-m}_{m+1,\mathcal{Q}'}(a_2) = \mathrm{segtype}^{n-m}_{m+1,\mathcal{P}}(e)\,.$$

Hence at least one of them is on a different thread than $h(x')$ and the procedure described above yields an inclusion.

In order to conclude the proof of Proposition 4.4.3 it remains to extend the domain of h to V. This done in the exact same way but, as the threads of V may not include the endpoints, it gives rise to new cases. We use the same notations. Let t be the thread under investigation and let u be its first element in V. Note that u doesn't have to be an endpoint of t.

The first difference is in Case 6: it may be the case that there is no z verifying the same conditions as y. In this case, and if no previous case is applicable, it must be the case that such a z appear "before" $h(u)$: that is, $z, h(u), h(x'), y$ appear in that order. There are now two possibilities:

- as described in Figure 4.18, $z', z, h(x')$ are in that order, where z' is the image of x' by an isomorphism mapping the neighborhood of x to that of z. Then set $\mathcal{Q}''$ to be the m-guarded contiguous-segment-E-swap between $[z, h(x')]$ and $[\hat{x}, y']$ in $\mathcal{Q}'$, and extend h by setting $h(x)$ to y.

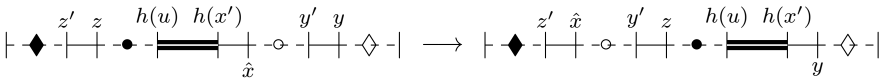

Figure 4.18: An illustration of the progression in the inclusion of $\mathcal{V}$, first completion of case 6.

- otherwise, $z, z', h(x')$ appear in that order (c.f. Figure 4.19).

  Set $\mathcal{Q}''$ to be the m-guarded mirror-E-swap at $[z', h(x')]$ in $\mathcal{Q}'$, and extend h by setting $h(x)$ to z. Notice that we have "reversed" the direction on the inclusion of the current thread but this isn't an issue since E is not oriented.

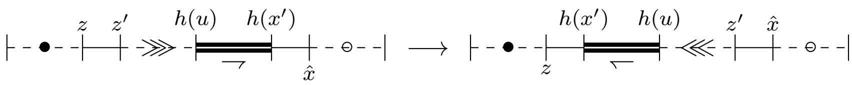

Figure 4.19: An illustration of the progression in the inclusion of $\mathcal{V}$, second completion of case 6.

The second difference is that it is now possible that none of the cases described above are applicable. In that situation, there must exist two nodes y

and z "before" $h(u)$ having the same neighborhood type as x. If at least one of them (say z) is in reverse order (i.e. $z, z', h(x')$ appear in that order, c.f. Figure 4.19) we proceed exactly as before.

Otherwise, it means that we can set $\mathcal{Q}''$ to be (assuming without loss of generality that $y, z, h(x')$ appear in that order, c.f. Figure 4.20) the m-guarded contiguous-segment-E-swap between $[y, z']$ and $[z, h(x')]$ in $\mathcal{Q}'$ and extend h by setting $h(x)$ to y.

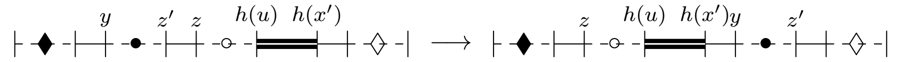

Figure 4.20: An illustration of the progression in the inclusion of $\mathcal{V}$, if no previous case is applicable.

□

4.5 Tools for reorganizing S-edges

In the previous section, we have seen how to "rewrite" $\mathcal{Q}$ using E-swap operations in order to pseudo-include $\mathcal{P}$ into the resulting quasitree. By definition, the pseudo-inclusion h of $\mathcal{P}$ into $\mathcal{Q}$ respects the enriched support but can be completely wild relatively to the S-edges. For instance, in $\mathcal{Q}$, the endpoints of a thread may not have the same S-parent. In this section we show how to use S-swaps in order to ensure that our pseudo-inclusion mapping takes into account (to various degrees) the S-edges. We say that two nodes of a quasitree are S-siblings if they share the same S-parent.

In Section 4.5.1, we show how to make sure that the pseudo-inclusion respects the S-siblings relation. In Section 4.5.2 we show how to ensure that the image of a pseudo-inclusion is S-stable. S-stability is required to define and operate on the loops, as will be established in Section 4.6.

4.5.1 S-siblings re-association

The following lemma shows how to modify a pseudo-inclusion in order for it to preserve the S-siblings relation. Note that it doesn't necessarily mean that the image structure has the matching endpoint property because the initial structure itself may not have this property as it is derived from a quasitree.

Lemma 4.5.1. $\forall \alpha, m \in \mathbb{N}, \exists N \in \mathbb{N}$ *such that*

for every $\mathcal{W} \in Ctxt_\sigma^N$ *and* $\mathcal{Q} \in$ *quasi-*$\mathbb{H}_\sigma^N$, *if* $h : W \to Q$ *is an* N*-pseudo-inclusion,*

then there exists some $\mathcal{Q}' \in$ *quasi-*$\mathbb{H}_\sigma^{m+1}$ *and some* $(m+1)$*-pseudo-inclusion* $h' : W \to Q'$ *such that*

$$\mathcal{Q}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q} \text{ and } \mathcal{S}upp_{m+1}(\mathcal{Q}') \simeq \mathcal{S}upp_{m+1}(\mathcal{Q})$$

and such that if x *and* y *are* S*-siblings in* $\mathcal{W}$, *then so are* $h'(x)$ *and* $h'(y)$ *in* $\mathcal{Q}'$.

Proof. We can assume that $m \geq s(\alpha)$. N is to be fixed later, and will chosen such that $2N \geq 2(2m+3)+1$.

Let $(x_1, y_1), \cdots, (x_r, y_r)$ denote all the pairs of endpoints of threads of length $\leq 2N-1$ of $\mathcal{W}$ (they must be S-siblings), and let $(x_{r+1}, y_{r+1}), \cdots, (x_s, y_s)$ denote the other pairs of S-siblings of $\mathcal{W}$, in an arbitrary order.

We are going to construct a sequence of structures

$$\mathcal{Q} = \mathcal{Q}_r \equiv_\alpha^{<\text{-inv FO}} \cdots \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q}_s$$

of same $(m+1)$-enriched support, and functions $f_r, \cdots, f_s$ such that:

- f_i $(m+1)$-pseudo-includes $\mathcal{W}$ in $\mathcal{Q}_i$
- $\forall j \leq i$, $f_i(x_j)$ and $f_i(y_j)$ are S-siblings in $\mathcal{Q}_i$
- $\forall j > i$, let z be the S-sibling of $f_i(x_j)$ in $\mathcal{Q}_i$. Let Z (resp. Y) be the element at distance $2m+3$ of z (resp. $f_i(y_j)$) in $\mathcal{S}\text{upp}_0(\mathcal{Q}_i)$ (Z and Y exist since their threads are of length $\geq 2N$).

  Then $\mathcal{S}\text{upp}_{m+1}(\mathcal{Q}_i)|_{[z,Z]} \simeq \mathcal{S}\text{upp}_{m+1}(\mathcal{Q}_i)|_{[f_i(y_j),Y]}$.

For $i = r$, set $\mathcal{Q}_r := \mathcal{Q}$ and $f_r := h$. Note that threads of $\mathcal{Q}$ of length at most $2N-1$ must have matching endpoints. N is chosen large enough so that the last property holds - $N := 2 + (2m+3) + (m+1)$ is enough.

Assume now that we have constructed $\mathcal{Q}_i$ and f_i as required. If $f_i(x_{i+1})$ and $f_i(y_{i+1})$ are S-siblings in $\mathcal{Q}_i$, set $\mathcal{Q}_{i+1} := \mathcal{Q}_i$ and $f_{i+1} := f_i$.

Otherwise, let z be the S-sibling of $f_i(x_{i+1})$, Z (resp. Y) be the element at distance $2m+3$ of z (resp. $f_i(y_{i+1})$) in $\mathcal{S}\text{upp}_0(\mathcal{Q}_i)$, and Z' (resp. Y') be the element at distance $2m+4$ of z (resp. $f_i(y_{i+1})$) in $\mathcal{S}\text{upp}_0(\mathcal{Q}_i)$. We know that $\mathcal{S}\text{upp}_{m+1}(\mathcal{Q}_i)|_{[z,Z]} \simeq \mathcal{S}\text{upp}_{m+1}(\mathcal{Q}_i)|_{[f_i(y_{i+1}),Y]}$ (witnessed by an isomorphism ϕ).

In particular, $\text{tp}^m_{\mathcal{Q}_i}(Z, Z') = \text{tp}^m_{\mathcal{Q}_i}(Y, Y')$, and $\{Y, Y', Z, Z'\}$ and $\text{End}(\mathcal{Q}_i)$ are $(2m+3)$-distant in $\mathcal{S}\text{upp}_0(\mathcal{Q}_i)$ by choice of N.

We distinguish between two cases:

- if Y, Y' and Z, Z' are in different threads.

  Let $\mathcal{Q}_{i+1}$ be the m-guarded crossing-E-swap between ZZ' and YY' in $\mathcal{Q}_i$ (c.f. Figure 4.21).

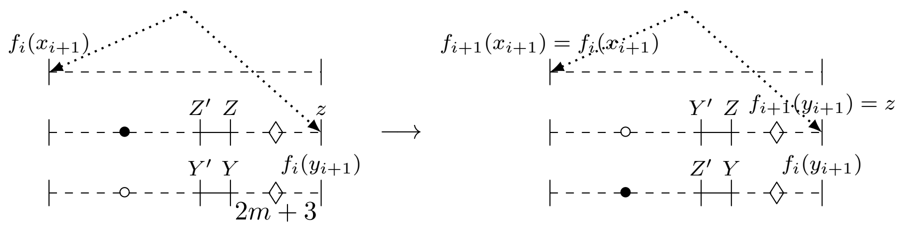

Figure 4.21: After the m-guarded crossing-E-swap between ZZ' and YY' in $\mathcal{Q}_i$, $f_{i+1}(x_{i+1})$ and $f_{i+1}(y_{i+1})$ are S-siblings.

- if Y, Y' and Z, Z' are in the same thread.

  Let $\mathcal{Q}_{i+1}$ be the m-guarded mirror-E-swap at $[Z', Y']$ in $\mathcal{Q}_i$ (c.f. Figure 4.22).

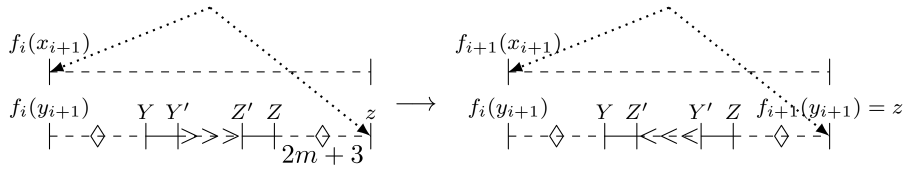

Figure 4.22: After the m-guarded mirror-E-swap at $[Z', Y']$ in $\mathcal{Q}_i$, $f_{i+1}(x_{i+1})$ and $f_{i+1}(y_{i+1})$ are S-siblings.

In both cases, we define f_{i+1} as $\Phi \circ f_i$ where Φ the permutation of Q_i defined as ϕ on $[z, Z]$, ϕ^{-1} on $[f_i(y_{i+1}), Y]$, and the identity elsewhere.

Lemma 4.3.5 guarantees that in both cases, $\mathcal{Q}_{i+1} \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q}_i$ and, together with the invariant $\mathcal{S}\text{upp}_{m+1}(\mathcal{Q}_i)|_{[z,Z]} \simeq \mathcal{S}\text{upp}_{m+1}(\mathcal{Q}_i)|_{[f_i(y_{i+1}),Y]}$, that $\mathcal{S}\text{upp}_{m+1}(\mathcal{Q}_{i+1}) \simeq \mathcal{S}\text{upp}_{m+1}(\mathcal{Q}_i)$.

Furthermore, $f_{i+1}(x_{i+1})$ and $f_{i+1}(y_{i+1})$ are S-siblings, and it is straightforward to see that f_{i+1} is a $(m+1)$-pseudo-inclusion, and that the other conditions are still respected.

In the end, $\mathcal{Q}' := \mathcal{Q}_s$ and $h' := f_s$ fit. □

A particular case of the previous lemma is when $\mathcal{W}$ is a hollow tree and h is surjective: then $\mathcal{Q}'$ has the matching endpoints property. This result will be useful in the proof of Proposition 4.6.8.

4.5.2 S-stabilization

We will often need to state that several sets are far from each other. To this end we introduce the notion of scattering, which is a compact way of saying that. For a subset A of a structure $\mathcal{R}$ whose vocabulary contains the binary relation S, define $\mathcal{R} \setminus S(A)$ to be $\mathcal{R}$ minus all the S-edges adjacent to any element of A. If $A = \{z\}$, we note $\mathcal{R} \setminus S(z)$ instead of $\mathcal{R} \setminus S(\{z\})$.

Definition 4.5.2. *Let $A_1, \cdots, A_k, B$ be subsets of R, and $\delta \in \mathbb{N}$.*

*We say that $A_1, \cdots, A_k$ **are δ-scattered wrt.** B if $A_1, \cdots, A_k, B$ are pairwise δ-distant in $\mathcal{R} \setminus S(A_1 \cup \cdots \cup A_k)$.*

The following lemma will be useful in a couple of proofs. It gives a setting in which we can apply simultaneous crossing-S-swaps.

Lemma 4.5.3. *Let $\alpha, s \in \mathbb{N}$, $m \geq s(\alpha)$, $\mathcal{R} \in$ quasi-$\mathbb{H}_\sigma^{m+1}$, and let*

$$a_1, a_1', a_1'', b_1, b_1', b_1'', \cdots, a_s, a_s', a_s'', b_s, b_s', b_s'' \in R$$

be distinct elements such that, for every i,

- *a_i', a_i'' (resp. b_i', b_i'') are the S-children of a_i (resp. b_i)*
- *$tp_{\mathcal{R}}^m(a_i, a_i', a_i'') = tp_{\mathcal{R}}^m(b_i, b_i', b_i'')$*

and let

$$B \supseteq \{b_1, b_1', b_1'', \cdots, b_s, b_s', b_s''\}$$

be such that $\{a_1\}, \cdots, \{a_s\}, B$ are pairwise $(2m+5)$-distant.

Let $\mathcal{R}'$ be $\mathcal{R}$ where all the $S(a_i, a'_i)$, $S(a_i, a''_i)$, $S(b_i, b'_i)$ and $S(b_i, b''_i)$ have been replaced by $S(a_i, b'_i)$, $S(a_i, b''_i)$, $S(b_i, a'_i)$ and $S(b_i, a''_i)$. In other words, $\mathcal{R}'$ is the simultaneous crossing-S-swap between every a_i and b_i in $\mathcal{R}$.

Then

- $\mathcal{R}' \equiv_{\alpha}^{\text{<-inv FO}} \mathcal{R}$,
- $\mathcal{S}upp_{m+1}(\mathcal{R}') = \mathcal{S}upp_{m+1}(\mathcal{R})$ *(in particular, $\mathcal{R}' \in$ quasi-$\mathbb{H}_{\sigma}^{m+1}$),*
- *and $\{a_1\}, \{a'_1, a''_1\}, \cdots, \{a_s\}, \{a'_s, a''_s\}$ are m-scattered wrt. B in $\mathcal{R}'$.*

Proof. Recall from Note 4.3.3 that $2m{+}5$ provides a sufficient distance condition to apply an m-guarded crossing-S-swap.

We construct a sequence of structures

$$\mathcal{R} = \mathcal{R}_0 \equiv_{\alpha}^{\text{<-inv FO}} \cdots \equiv_{\alpha}^{\text{<-inv FO}} \mathcal{R}_s$$

having the same $(m{+}1)$-enriched support, where $\mathcal{R}_i$ is the simultaneous crossing-S-swap between a_j and b_j in $\mathcal{R}$ for every $j \leq i$.

Let's show by induction that for every i, $\mathcal{R}_i$ verifies (P_i):

1. $\{a_{i+1}\}, \cdots, \{a_s\}, B$ are pairwise $(2m+5)$-distant
2. $\{a_1\}, \{a'_1, a''_1\}, \cdots, \{a_i\}, \{a'_i, a''_i\}$ are m-scattered wrt. B.

$\mathcal{R}$ verifies (P_0). Suppose that we have constructed $\mathcal{R}_i$ and let $\mathcal{R}_{i+1}$ be the m-guarded crossing-S-swap between a_{i+1} and b_{i+1} in $\mathcal{R}_i$: (P_i).1 ensures that $\text{dist}_{\mathcal{R}_i}(a_{i+1}, b_{i+1}) \geq 2m + 5$. Lemma 4.3.4 gives $\mathcal{R}_{i+1} \equiv_{\alpha}^{\text{<-inv FO}} \mathcal{R}_i$ and $\mathcal{S}\text{upp}_{m+1}(\mathcal{R}_{i+1}) = \mathcal{S}\text{upp}_{m+1}(\mathcal{R}_i)$. Let's show that $\mathcal{R}_{i+1}$ verifies (P_{i+1}):

(P_{i+1}).1 holds since the $(2m+4)$-neighborhoods of the $(a_j)_{j>i+1}$ haven't seen any change, because of (P_i).1.

(P_{i+1}).2 : Let $\mathcal{R}_i^-$ denote $\mathcal{R}_i \setminus S(\{a_1, a'_1, a''_1, \cdots, a_i, a'_i, a''_i\})$.
Let $\mathcal{R}_{i+1}^-$ denote $\mathcal{R}_{i+1} \setminus S(\{a_1, a'_1, a''_1, \cdots, a_{i+1}, a'_{i+1}, a''_{i+1}\})$.

Let $x, y \in \{a_1\} \cup \{a'_1, a''_1\} \cup \cdots \cup \{a_i\} \cup \{a'_i, a''_i\}$ be elements of two different sets.

(P_i).2 entails that x and y are each at distance $\geq m$ in $\mathcal{R}_i^-$ from each other and from B, and (P_i).1 implies (since x and y are at distance 1 of B) that they are at distance $\geq m$ in $\mathcal{R}_i$ (hence in $\mathcal{R}_i^-$) from $\{a_{i+1}, a'_{i+1}, a''_{i+1}\}$.

Hence, the swap doesn't affect their m-neighborhoods in $\mathcal{R}_i^-$, and they are still at distance at least m from each other and from B in

$$\mathcal{R}_{i+1} \setminus S(\{a_1, a'_1, a''_1, \cdots, a_i, a'_i, a''_i\})$$

hence in $\mathcal{R}_{i+1}^-$.

Let $a \in \{a_{i+1}, a'_{i+1}, a''_{i+1}\}$ and $b \in B$. A path in $\mathcal{R}_{i+1}^-$ of length $\leq m - 1$ from a to b or from a to x doesn't go through the new S-edges, hence is valid in $\mathcal{R}_i$ and contradicts (P_i).1 (in the second case, because x is at distance 1 from B in $\mathcal{R}_i$).

This entails $\mathrm{dist}_{\mathcal{R}^-_{i+1}}(a,b) \geq m$ and $\mathrm{dist}_{\mathcal{R}^-_{i+1}}(a,x) \geq m$.

It remains to show that $\mathrm{dist}_{\mathcal{R}^-_{i+1}}(a_{i+1}, a'_{i+1}) \geq m$ (and similarly for a_{i+1} and a''_{i+1}).

Suppose that there is a path of length $\leq m-1$ in $\mathcal{R}^-_{i+1}$ from a_{i+1} to a'_{i+1}. This path is valid in $\mathcal{R}^-_i$. Hence, there would be a "vertical loop" in $\mathcal{N}^{\mathcal{R}_i}_{m+1}(a_{i+1})$, contradicting $\mathcal{R}_i \in$ quasi-$\mathbb{H}^{m+1}_\sigma$.

We set $\mathcal{R}' := \mathcal{R}_s$, which has the desired properties. □

The image of a pseudo-inclusion has no reason to be S-stable, thus neither has its complement. However, this is a crucial requirement to apply the results presented in the next section, Section 4.6, in order to remove the extra material not in the image of the pseudo-inclusion.

The next result provides a method to ensure that the image (and its complement) of a pseudo-inclusion is S-stable.

Recall that a pseudo-inclusion is said to be reduced if there is at most one jumping pair of a given type. At the end of this process, we get a reduced pseudo-inclusion, which will allow us to minimize the complement of its image in Section 4.6.1.

Proposition 4.5.4. *For every $\alpha, m \in \mathbb{N}$, there exist $N, d, D \in \mathbb{N}$ such that for every $\mathcal{P} \in \mathbb{H}_\sigma$ and $\mathcal{Q} \in$ quasi-$\mathbb{H}^{N+1}_\sigma$ such that*

$$[\![\mathcal{E}_{N+1}(\mathcal{P})]\!] \leq^D_d [\![\mathcal{E}_{N+1}(\mathcal{Q})]\!]$$

and such that $\mathcal{P}$ is $(N+1)$-pseudo-included in $\mathcal{Q}$ through some h, there exist h' and $\mathcal{Q}' \in$ quasi-$\mathbb{H}^{m+1}_\sigma$ where

$$\mathcal{Q}' \equiv^{<\text{-inv FO}}_\alpha \mathcal{Q} \text{ and } \mathcal{S}upp_{m+1}(\mathcal{Q}') \simeq \mathcal{S}upp_{m+1}(\mathcal{Q})$$

and where h' is a reduced $(m+1)$-pseudo-inclusion of $\mathcal{P}$ in $\mathcal{Q}'$ such that

$$Q' \setminus \mathrm{Im}(h') \text{ is } S\text{-stable in } \mathcal{Q}' .$$

Proof. The idea of the proof is to consider all the pairs of elements x, y which break the S-stability of $\mathrm{Im}(h)$, i.e. such that $S(x,y)$, $x \in \mathrm{Im}(h)$ and $y \notin \mathrm{Im}(h)$. If there are many of them, then at least two of them are far from each other and we can apply a crossing-S-swap to correct the mapping h. We end up with a bounded number of problematic pairs that can be corrected separately.

We can assume that $m \geq s(\alpha)$. We will first provide a non-necessarily reduced $(m'+1)$-pseudo-inclusion verifying those conditions, with $m' := 2m+3$, and then modify it as well as the underlying structure to get a fitting reduced $(m+1)$-pseudo-inclusion.

For every $n \in \mathbb{N}$ (we will assign a value to n later on), there is an N such that, under the hypothesis, Lemma 4.5.1 yields $\mathcal{R} \equiv^{<\text{-inv FO}}_\alpha \mathcal{Q}$ such that $\mathcal{S}\mathrm{upp}_{n+2}(\mathcal{R}) \simeq \mathcal{S}\mathrm{upp}_{n+2}(\mathcal{Q})$ and g which $(n+2)$-pseudo-includes $\mathcal{P}$ in $\mathcal{R}$ and respects the S-siblings relation; we denote by V the complement of $\mathrm{Im}(g)$ in Q.

This implies that two nodes having the same S-parent are both either in $\mathrm{Im}(g)$ (and are the two endpoints of the same thread) or in V.

We say that $z \in V$ is **misassociated** if its S-children are in $\mathrm{Im}(g)$. Likewise, we say that $g(x)$ is misassociated if its S-children are in V. The $(n+2)$-type of this element is called the type of the misassociation. Note that the number of misassociations in V and in $\mathrm{Im}(g)$ is the same.

First, we deal with all but a bounded number of misassociations. There exists an M (which depends only on n) such that, if there are more than $2M$ misassociations, then we can find a misassociated element of V and one of $\mathrm{Im}(g)$ that have the same type, and are at distance $\geq 2(n+1)+5$ from one another: this is because a hollow quasitree has degree at most 4. We can solve these misassociations by a $(n+1)$-guarded crossing-S-swap, according to Lemma 4.3.4 and Note 4.3.3, which preserves $\mathcal{S}\mathrm{upp}_{n+2}(\mathcal{R})$.

Once we've done that, we're left with at most M misassociations in V, and the same number in $\mathrm{Im}(g)$. Let $(z_1, g(x'_1), g(x''_1)), \cdots, (z_r, g(x'_r), g(x''_r))$ be an arbitrary enumeration of the misassociated elements of V, together with their S-children (recall that x'_i and x''_i are S-siblings in $\mathcal{P}$, and let x_i be their S-parent).

Fix i between 1 and r. There exists a sequence $x_i = x_i^1, \cdots, x_i^{s_i}$ of elements of P, such that, if we name $x_i'^j$ and $x_i''^j$ the S-children of x_i^j in $\mathcal{P}$, for every j, $g(x_i^j)$ is the S-parent of $g(x_i'^{j+1})$ and $g(x_i''^{j+1})$, and $g(x_i^{s_i})$ is misassociated ; let $z'_i, z''_i \in V$ be its S-children, and rename $y_i := x_i^{s_i}$ for ease. This sequence is represented in Figure 4.23.

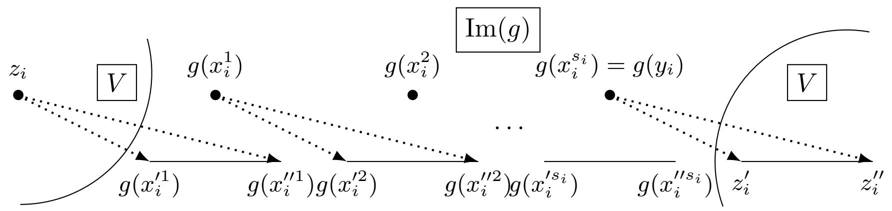

Figure 4.23: From z_i to (z'_i, z''_i).

For every j,

$$\mathrm{tp}_{\mathcal{P}}^{n+2}(x_i'^{j+1}) = \mathrm{tp}_{\mathcal{R}}^{n+2}(g(x_i'^{j+1}))\,,$$

hence

$$\mathrm{tp}_{\mathcal{P}}^{n+1}(x_i^{j+1}) = \mathrm{tp}_{\mathcal{R}}^{n+1}(g(x_i^{j}))\,,$$

which in turn implies that

$$\mathrm{tp}_{\mathcal{R}}^{n+1}(g(x_i^{j+1})) = \mathrm{tp}_{\mathcal{R}}^{n+1}(g(x_i^{j}))\,.$$

For the same reason, we have that $\mathrm{tp}_{\mathcal{R}}^{n+1}(z_i) = \mathrm{tp}_{\mathcal{R}}^{n+1}(g(x_i^1))$. In the end, we get

$$\tau_i := \mathrm{tp}_{\mathcal{R}}^{n+1}(z_i) = \mathrm{tp}_{\mathcal{R}}^{n+1}(g(y_i))\,.$$

Let B be the set containing the z_i, the $g(y_i)$, for $1 \leq i \leq r$, and their S-children.

Since we've bounded r by M and $\mathcal{R}$ is of degree 4, we can choose d and D large enough so that we are able to find $t_1, \cdots, t_r \in V$ and $u_1, \cdots, u_r \in \mathrm{Im}(g)$, with respective S-children t_i', t_i'' and u_i', u_i'', such that t_i and u_i are of neighborhood type τ_i (since $z_i \in V$ is of type τ_i, there must be at least d elements of this type in $\mathrm{Im}(g)$ and D in V), and such that $\{t_1\}, \cdots, \{t_r\}, \{u_1\}, \cdots, \{u_r\}, B$ are pairwise $(2n+5)$-distant in $\mathcal{R}$.

We can apply Lemma 4.5.3 with

- $s := 2r$
- $(a_1, \cdots, a_s) := (t_1, \cdots, t_r, u_1, \cdots, u_r)$
- $(b_1, \cdots, b_s) := (z_1, \cdots, z_r, g(y_1), \cdots, g(y_r))$.

This ensures that $\mathcal{R}'$ (which is the simultaneous crossing-S-swaps between z_i and t_i and crossing-S-swaps between $g(y_i)$ and u_i) is such that

$$\mathcal{R}' \equiv_\alpha^{<\text{inv FO}} \mathcal{R} \quad \text{and} \quad \mathcal{S}\mathrm{upp}_{n+1}(\mathcal{R}') = \mathcal{S}\mathrm{upp}_{n+1}(\mathcal{R})\,.$$

Furthermore, $\{t_1\}, \{t_1', t_2'\}, \{u_1\}, \{u_1', u_1''\}, \cdots, \{t_r\}, \{t_r', t_r''\}, \{u_r\}, \{u_r', u_r''\}$ are n-scattered wrt. B in $\mathcal{R}'$.

Note that we haven't added any new misassociated element in the process: the only misassociated elements in $\mathcal{R}'$ are now the t_i and the u_i.

Choose retrospectively $n := 2m' + 5$.

Let's show that $\{t_1\}, \cdots, \{t_r\}, \{u_1\}, \cdots, \{u_r\}$ are pairwise $(2m'+5)$-distant in $\mathcal{R}'$.

Let x, y be distinct elements among them, and let's prove that

$$\mathrm{dist}_{\mathcal{R}'}(x, y) \geq 2m' + 5\,.$$

Suppose that's false, and consider a shortest path from x to y. It cannot be valid in $\mathcal{R}$, hence it must go though at least one new S-edge, and the first one must be $S(x, x')$, with x' being either z_i' or z_i'' (if $x = u_i$) or $g(x_i')$ or $g(x_i'')$ (if $x = t_i$) for some i.

The only way to reach y from x' in less than $2m'+4$ is through an S-children y' of y.

Now, $x \neq y$, hence $\mathrm{dist}_{\mathcal{S}\mathrm{upp}_0(\mathcal{R}')}(x', y') \geq 2m' + 3$ (either they are endpoints of two different threads, either the thread they're both in doesn't have the matching endpoint property, which ensures that it is of length $> 2n+1$). We can thus apply Lemma 4.3.9, which states that the path of length $\leq 2m' + 3$ from x' to y' must go through either x or y. This contradicts the minimality hypothesis.

We can proceed to the sequence of m'-guarded crossing-S-swap between u_i and t_i in $\mathcal{R}'$ for every i.

After the r swaps, we end up with $\mathcal{R}'' \equiv_\alpha^{<\text{inv FO}} \mathcal{Q}$, such that

$$\mathcal{S}\mathrm{upp}_{m'+1}(\mathcal{R}'') \simeq \mathcal{S}\mathrm{upp}_{m'+1}(\mathcal{Q})$$

and with no misassociation left wrt. g.

If the pseudo-inclusion g isn't reduced, we reduce it by eliminating one by one its redundant jumping pairs.

Seeing g as an m-pseudo-inclusion, we get that nodes involved in a jumping pair are at distance $\geq 2m+3$ from the endpoints in $\mathcal{S}\mathrm{upp}_0(\mathcal{R}'')$.

Let $\{x, x'\}$ and $\{y, y'\}$ be two jumping pairs with the same type, and u', u (resp v', v) be the E-neighbors of $g(x), g(x')$ (resp. $g(y), g(y')$) in $[g(x), g(x')]$ (resp. $[g(y), g(y')]$). We have that

$$\mathrm{tp}^m_{\mathcal{R}''}(g(x), u') = \mathrm{tp}^m_{\mathcal{R}''}(u, g(x')) = \mathrm{tp}^m_{\mathcal{R}''}(g(y), v') = \mathrm{tp}^m_{\mathcal{R}''}(v, g(y')) .$$

If their images are on two different threads (c.f. Figure 4.24), we can perform two m-guarded crossing-E-swaps: first, the m-guarded crossing-E-swap between $g(x)u'$ and $vg(y')$ in $\mathcal{R}''$, and then the m-guarded crossing-E-swap between $g(x)g(y')$ and $ug(x')$ in the previous swap, after which $\{x, x'\}$ is no longer a jumping pair.

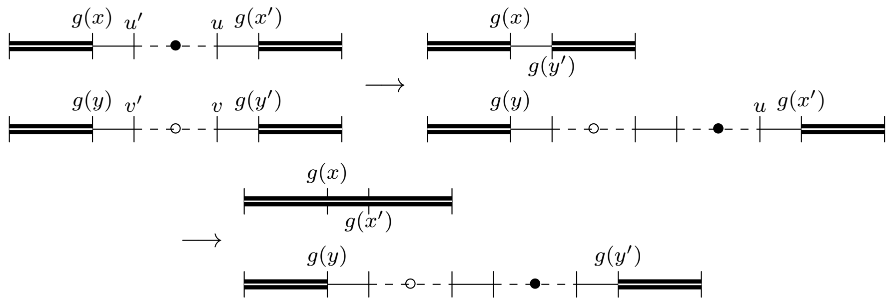

Figure 4.24: Elimination of a jumping pair among two, in different threads.

If their images appear on the same thread in the order $g(y), g(y'), g(x), g(x')$ (c.f. Figure 4.25), we can perform the m-guarded contiguous-segment-E-swap between $[v', g(x)]$ and $[u', u]$ in $\mathcal{R}''$, after which $\{x, x'\}$ is no longer a jumping pair.

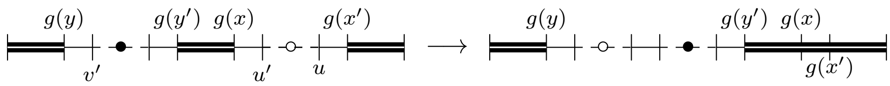

Figure 4.25: Elimination of a jumping pair among two, with a contiguous-segment-E-swap.

Otherwise, we can assume as in Figure 4.26 that the images appear in the order $g(y), g(y'), g(x'), g(x)$. We can perform two consecutive m-guarded mirror-E-swaps: first the m-guarded mirror-E-swap at $[v', g(x')]$ in $\mathcal{R}''$, and then (in order to reverse again the segment $[g(y'), g(x')]$ into the initial direction) the m-guarded mirror-E-swap at $[g(x'), u']$ in the previous swap, after which $\{x, x'\}$ is no longer a jumping pair.

In the end, we get $\mathcal{Q}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q}$, such that $\mathcal{S}\mathrm{upp}_{m+1}(\mathcal{Q}') \simeq \mathcal{S}\mathrm{upp}_{m+1}(\mathcal{Q})$, and a reduced h' that $(m+1)$-pseudo-includes $\mathcal{P}$ in $\mathcal{Q}'$. Notice that during the transformation from g (which was misassociation-free) to h, we never created any misassociation. Hence, $Q' \setminus \mathrm{Im}(h')$ is S-stable.

□

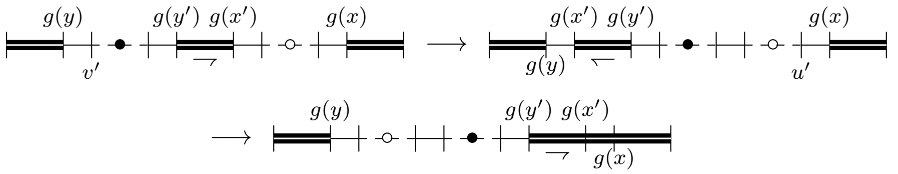

Figure 4.26: Elimination of a jumping pair among two, with two mirror-E-swaps.

4.6 Removing unnecessary material

In this section we show how to remove the material in $\mathcal{Q}$ that is not present in the image of the pseudo-inclusion of $\mathcal{P}$. From the previous section we can assume that the pseudo-inclusion mapping preserves the S-siblings relation and that its image is S-stable. The remaining part of $\mathcal{Q}$ is then a union of "loops" in the sense that they connect nodes that have the same neighborhood type. After defining properly the notion of loop, we will use in Section 4.6.1 a pumping argument in order to reduce the size of the loop to some constant while preserving $\equiv_\alpha^{\text{<-inv FO}}$. In Section 4.6.2 we then show how to remove small loops without affecting the order-invariant equivalence class. Finally, in Section 4.6.3 we show that if a hollow tree and a hollow quasitree have the same enriched support, then they are $\equiv_\alpha^{\text{<-inv FO}}$: this concludes the proof of Theorem 4.2.1.

We start with the definition of an abstract loop.

Let $n \in \mathbb{N}$. Let $\text{Type}_\sigma^n[2]$ denote the set of $(n-1)$-neighborhood types for pairs over the vocabulary $P_\sigma \cup \{E, S\}$, of degree ≤ 4. Let Σ_n be the vocabulary enriching $P_\sigma \cup \{E, S\}$ with two unary symbols J_τ^1 and J_τ^2 for every $\tau \in \text{Type}_\sigma^n[2]$.

Let h be a reduced n-pseudo-inclusion from $\mathcal{P} \in \mathbb{H}_\sigma$ to $\mathcal{Q} \in \text{quasi-}\mathbb{H}_\sigma^n$, such that $V := Q \setminus \text{Im}(h)$ is S-stable.

Let $\mathcal{Q}_+$ be an extension of $\mathcal{Q}$ to Σ_n obtained in the following way. Since h is reduced, for every $\tau \in \text{Type}_\sigma^n[2]$, there is at most one jumping pair of type τ. If there isn't, J_τ^1 and J_τ^2 are interpreted as the empty set. Else, let $\{x, x'\}$ be this pair, and u' (resp. u) be the E-neighbor of $h(x)$ (resp. $h(x')$) in $[h(x), h(x')]$. Interpret J_τ^1 as $\{h(x), u'\}$ and J_τ^2 as $\{h(x'), u\}$ (the assignments $x \mapsto 1$ and $x' \mapsto 2$ are arbitrary). This is illustrated on the left part of Figure 4.27, where the double line represents $\text{Im}(h)$. We say that $\mathcal{Q}_+$ is a **h-jump-extension of $\mathcal{Q}$**.

We define $\mathcal{V}_+ = \text{Ctxt}_n(\mathcal{Q}_+|_V)$ as the extension of $\text{Ctxt}_n(\mathcal{Q}|_V)$ to Σ_n where every J_τ^i is defined consistently with $\mathcal{Q}_+$ (i.e. $\forall x \in V, \mathcal{V}_+ \models J_\tau^i(x)$ iff $\mathcal{Q}_+ \models J_\tau^i(x)$). This process is illustrated in Figure 4.27. $\mathcal{V}_+$ is called an **n-abstract loop**. Let $\mathbb{L}_\sigma^n$ be the set of n-abstract loops.

Every Σ_n-structure will have a '+' symbol in its name. When we omit it, we mean the reduction of the structure to $P_\sigma \cup \{E, S\}$ (for instance, from $\mathcal{V}_+ \in \mathbb{L}_\sigma^n$, we get $\mathcal{V} := \text{Ctxt}_n(\mathcal{Q}|_V) \in \text{Ctxt}_\sigma^n$).

Let $\mathcal{W}_+ \in \mathbb{L}_\sigma^n$, and g be an n-inclusion from $\mathcal{W}$ to some $\mathcal{R} \in \text{quasi-}\mathbb{H}_\sigma^n$.

Let $\mathcal{R}_+$ be an extension of $\mathcal{R}$ to Σ_n obtained in the following way. For every $\tau \in \text{Type}_\sigma^n[2]$, and $i \in \{1, 2\}$, such that there exists (a unique) $x_\tau^i \in W$ such that $\mathcal{W}_+ \models J_\tau^i(x_\tau^i)$, J_τ^i is interpreted in $\mathcal{R}_+$ as $\{g(x_\tau^i), y_\tau^i\}$, where $y_\tau^i \notin \text{Im}(g)$ and $E(g(x_\tau^i), y_\tau^i)$. The existence and unicity of such y_τ^i is guaranteed. This process

Figure 4.27: Example of a h-jump-extension $\mathcal{Q}_+$ of $\mathcal{Q}$ (on the left), and its associated abstract loop $\mathcal{V}_+$ of support $V := Q \setminus \mathrm{Im}(h)$ (on the right).

is depicted in Figure 4.28. Every other J_τ^i is interpreted as the empty set. We say that $\mathcal{R}_+$ is the **g-border-extension of $\mathcal{R}$**.

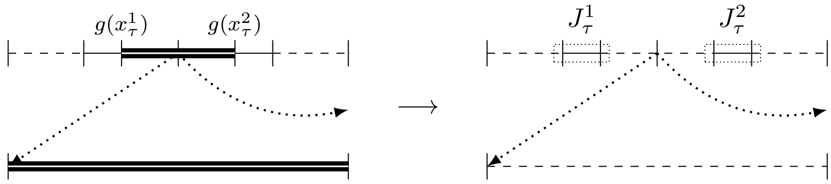

Figure 4.28: From an inclusion g (double line) of the previous $\mathcal{V}$ in $\mathcal{R}$ to the g-border-extension $\mathcal{R}_+$.

Let $\mathcal{I}_n$ be the FO-interpretation from the vocabulary Σ_n to $P_\sigma \cup \{E, S\}$, which adds an E-edge between a and b if $a \neq b$, $J_\tau^i(a)$ and $J_\tau^i(b)$ for some $i \in \{1, 2\}$ and τ, and then forgets about the $(J_\tau^i)_{(i,\tau)}$. Every $\mathcal{I}_n$ has arity 1 and depth 0. Hence for every $k \in \mathbb{N}$ and Σ_n-structures $\mathcal{A}$ and $\mathcal{B}$, $\mathcal{A} \equiv_k^{<\text{-inv FO}} \mathcal{B}$ entails $\mathcal{I}_n(\mathcal{A}) \equiv_k^{<\text{-inv FO}} \mathcal{I}_n(\mathcal{B})$.

4.6.1 Loop minimization

It will be crucial to bound the size of the loops left by a pseudo-inclusion. The following result does this using a simple pumping argument.

Proposition 4.6.1. *For every $\alpha, n \in \mathbb{N}$, there exists $N \in \mathbb{N}$ such that*

for every $\mathcal{P} \in \mathbb{H}_\sigma$, $\mathcal{Q} \in$ quasi-$\mathbb{H}_\sigma^n$ and for every reduced n-pseudo-inclusion $h : P \to Q$, such that $V := Q \setminus \mathrm{Im}(h)$ is S-stable,

there exists $\mathcal{Q}' \in$ quasi-$\mathbb{H}_\sigma^n$ and a reduced n-pseudo-inclusion $h' : P \to Q'$ such that

$$\mathcal{Q}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q}$$

and such that $U := Q' \setminus \mathrm{Im}(h')$ is S-stable and $|U| \leq N$.

Proof. For every equivalence class $\mathcal{C}$ of $\equiv_\alpha^{<\text{-inv FO}}$ on $\mathbb{L}_\sigma^n$, pick a representative $\mathcal{U}_+^{\mathcal{C}}$. Now, set $N := \max\{|U_+^{\mathcal{C}}| : \mathcal{C}$ equivalence class for $\equiv_\alpha^{<\text{-inv FO}}\}$. N is well defined since $\equiv_\alpha^{<\text{-inv FO}}$ is of finite index.

Let $\mathcal{Q}_+$ be a h-jump-extension of $\mathcal{Q}$.

Let $\mathcal{U}_+$ be the representative of the class of $\mathcal{V}_+ := \mathrm{Ctxt}_n(\mathcal{Q}_+|_V)$.

Since V is S-stable in $\mathcal{Q}$,

$$\mathcal{Q}^+ \setminus \{E(h(x), u'), E(u, h(x')) : \{x, x'\} \text{ jumping pair}\}$$

can be decomposed as $\mathcal{V}_+ \uplus \mathcal{R}_+$ for some Σ_n-structure $\mathcal{R}_+$.

Note that $\mathcal{Q} = \mathcal{I}_n(\mathcal{V}_+ \uplus \mathcal{R}_+)$. We set $\mathcal{Q}' := \mathcal{I}_n(\mathcal{U}_+ \uplus \mathcal{R}_+)$ and $h' := h$ (this makes sense since $R = \mathrm{Im}(h)$).

By definition of $\mathcal{U}_+$,

$$\mathcal{U}_+ \uplus \mathcal{R}_+ \equiv_\alpha^{<\text{-inv FO}} \mathcal{V}_+ \uplus \mathcal{R}_+ \,.$$

Applying $\mathcal{I}_n$ yields

$$\mathcal{Q}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q} \,.$$

It remains to show that h' is an n-pseudo-inclusion. Every thread of $\mathcal{P}$ is still sent on a single thread: indeed, for every jumping pair $\{x, x'\}$ for h, $h'(x)$ and $h'(x')$ lie on the same thread. This is because in $\mathcal{U}_+ \in \mathbb{L}_\sigma^n$.

All that's left to prove is that for every $a \in R, \mathrm{tp}_{\mathcal{Q}'}^n(a) = \mathrm{tp}_{\mathcal{Q}}^n(a)$. This follows from the fact that for every τ and $i \in \{1, 2\}$, the element of $\mathcal{U}_+$ colored with J_τ^i and the element of $\mathcal{V}_+$ coloured with J_τ^i (if they exist) have the same n-neighborhood type, once again because $\mathcal{U}_+ \in \mathbb{L}_\sigma^n$. □

4.6.2 Loop elimination

It now remains to get rid of the small loops. This is a consequence of the "aperiodicity" of $<$-inv FO: we cannot distinguish in $<$-inv FO between k and $k+1$ copies of the same object if k is sufficiently large. Starting from a small loop, we can use the inclusion results of Section 4.4 to recreate many copies of the loop within $\mathcal{Q}$, then, according to the following proposition, get rid of one copy using aperiodicity.

We now turn to loop elimination.

Our goal is to get rid of the extra material found outside of the image of the pseudo-inclusion. For that, we make sure it is S-stable (Proposition 4.5.4), we minimize it (Proposition 4.6.1), then we include (Proposition 4.4.3) a great number a times this loop in $\mathcal{Q}$. However, to be able to remove a copy while staying in the same $\equiv_\alpha^{<\text{-inv FO}}$-class, we need to recreate every of these loops to the original cape we included: recall indeed that the inclusion preserves the E-edges, but not necessarily the S-edges.

The following lemma gives a method to modify the including structure so that the pseudo-inclusion respects S-edges.

Lemma 4.6.2. $\forall \alpha, n, \exists N, \forall M, \exists D \in \mathbb{N}$ *such that*

for every $\forall \mathcal{Q} \in$ *quasi-*$\mathbb{H}_\sigma^{N+1}$ *and* $\forall \mathcal{W} \in \mathit{Ctxt}_\sigma^{N+1}$ *such that* $|W| \leq M$, *and*

for every $(N+1)$*-inclusion* $h : W \to Q$ *such that for every* $(N+1)$*-neighborhood type* τ *that occurs in* $\mathcal{W}$, *there are at least* D *elements of type* τ *in* $Q \setminus \mathrm{Im}(h)$,

there exist some $\mathcal{Q}' \in$ *quasi-*$\mathbb{H}_\sigma^n$ *and some* g *that* n*-includes* $\mathcal{W}$ *into* $\mathcal{Q}'$ *such that*

$$\mathcal{Q}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q} \text{ and } \mathcal{S}upp_n(\mathcal{Q}') \simeq \mathcal{S}upp_n(\mathcal{Q})$$

and

$$\mathcal{W} \models S(x, y) \quad \to \quad \mathcal{Q}' \models S(g(x), g(y)) \,.$$

Proof. We can assume that $n \geq s(\alpha)$. We'll assign values to m and N later on, in that order.

Keep in mind from Note 4.3.3 that a crossing-S-swap is guarded as long as it happens between elements of same neighborhood type that are distant enough. First, we re-associate the S-edges going in/out of the images of every S-parent and S-child. The hypothesis on the number of excess occurrences of every neighborhood type allows us to scatter their S-neighbors across the including structure. Recall that we introduced the notion of scattering in Definition 4.5.2.

Let's enumerate arbitrarily as $(x_1, x'_1, x''_1), \cdots, (x_r, x'_r, x''_r)$ the elements of $\mathcal{W}$ such that $S(x_i, x'_i) \wedge S(x_i, x''_i) \wedge x'_i \neq x''_i$.

First, we use Lemma 4.5.1 to find $\mathcal{R} \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q}$ such that

$$\mathcal{S}\text{upp}_{m+1}(\mathcal{R}) \simeq \mathcal{S}\text{upp}_{m+1}(\mathcal{Q})$$

and g such that $g : W \to R$ is an $(m+1)$-pseudo-inclusion that respects the S-siblings relation, for some m to be specified later on. This sets the value for N.

Let $B := \{b_1, b'_1, b''_1, \cdots, b_s, b'_s, b''_s\}$ be such that b'_i, b''_i are the S-children of b_i and $\forall i \leq r, \exists j, k \leq s$, $g(x_i) = b_j, g(x'_i) = b'_k$ and $g(x''_i) = b''_k$ (the existence of a k comes from the fact that g respects the S-siblings relation). Note that the minimal such B is $\text{Im}(g)$ plus the S-children of every $g(x_i)$ (if they are not already in $\text{Im}(g)$), plus the S-parent of every $g(x'_i), g(x''_i)$ (if it's not already in $\text{Im}(g)$). This guarantees that $s \leq 2r$.

Every hollow 1-quasitree has degree at most 4. In $\mathcal{R}$, $|\text{Im}(g)| \leq M$; hence as long as D is large enough, there must exist elements $(a_i)_{1 \leq i \leq s} \in R$, such that for every i, a'_i, a''_i being the S-children of a_i in $\mathcal{R}$, $\text{tp}^m_{\mathcal{R}_0}(a_i, a'_i, a''_i) = \text{tp}^m_{\mathcal{R}}(b_i, b'_i, b''_i)$, and $\{a_1\}, \cdots, \{a_s\}, B$ are pairwise $(2m+5)$-distant in $\mathcal{R}$, where $m := 2n + 5$.

We are in the right conditions to apply Lemma 4.5.3, and get $\mathcal{R}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{R}$, with $\mathcal{S}\text{upp}_{2n+6}(\mathcal{R}') = \mathcal{S}\text{upp}_{2n+6}(\mathcal{R})$ and $\{a_1\}, \{a'_1, a''_1\}, \cdots, \{a_s\}, \{a'_s, a''_s\}$ are $(2n+5)$-scattered wrt. B in $\mathcal{R}'$. Note that $g : W \to R'$ still preserves the S-siblings relation.

Not all of the a_i, a'_i, a''_i are of interest. We re-index them, and focus on $u_1, u'_1, u''_1, \cdots, u_r, u'_r, u''_r$, where u_i is the S-parent of $g(x'_i), g(x''_i)$ and u'_i, u''_i are the S-children of $g(x_i)$.

The scattering of the a_i, a'_i, a''_i entails that $\{u_1\}, \{u'_1, u''_1\}, \cdots, \{u_r\}, \{u'_r, u''_r\}$ are $(2n+5)$-scattered wrt. $\text{Im}(g)$ in $\mathcal{R}'$.

Set $\mathcal{W}_i := \mathcal{W} \uplus \{\bar{x}_{i+1}, \cdots, \bar{x}_r\}$ where, for every $j > i$, $S(x_j, x'_j)$ and $S(x_j, x''_j)$ have been replaced by $S(\bar{x}_j, x'_j)$ and $S(\bar{x}_j, x'_j)$. There cannot be a path of length $\leq 2n + 5$ from x'_j (or x''_j) to x_j, as long as $N + 1 \geq 2n + 5$, for otherwise there would be a vertical loop in $\text{tp}^{N+1}_{\mathcal{W}}(x_j)$.

Now, let's re-associate the S-edges back so that g respects S. We construct a sequence of structures

$$\mathcal{T}_0 \equiv_\alpha^{<\text{-inv FO}} \cdots \equiv_\alpha^{<\text{-inv FO}} \mathcal{T}_r$$

having the same $(n+1)$-enriched support, where $\mathcal{T}_i$ is the simultaneous crossing-S-swap between $g(x_j)$ and u_j in $\mathcal{R}'$ for $j \leq i$.

Let's prove that for every i, $\mathcal{T}_i$ verifies (Q_i):

1. $\{u_{i+1}\}, \{u'_{i+1}, u''_{i+1}\}, \cdots, \{u_r\}, \{u'_r, u''_r\}$ are $(2n+5)$-scattered wrt. $\mathrm{Im}(g)$
2. $\forall j, k > i$, let $a_j \in \{x'_j, x''_j\}$. Then

$$\mathrm{dist}_{\mathcal{T}_i}(g(a_j), g(x_k)) \geq \min(\mathrm{dist}_{\mathcal{W}_i}(a_j, x_k), 2n+6)$$

3. $\forall j \neq k > i, a_j \in \{x'_j, x''_j\}$ and $a_k \in \{x'_k, x''_k\}$, $\mathrm{dist}_{\mathcal{T}_i}(g(a_j), g(a_k)) > 2n+5$

Set $\mathcal{T}_0 := \mathcal{R}'$.

We check that (Q_0).2 holds, for x'_j and x_k (it is similar for x''_j). Let them be such that

$$\mathrm{dist}_{\mathcal{T}_0}(g(x'_j), g(x_k)) \leq 2n+5$$

and let's prove that

$$\mathrm{dist}_{\mathcal{W}_0}(g(x'_j), g(x_k)) \leq \mathrm{dist}_{\mathcal{T}_0}(g(x'_j), g(x_k))\,.$$

Consider a shortest path from $g(x'_j)$ to $g(x_k)$ in $\mathcal{T}_0$.

Suppose it goes through at least one S-edge: the first time it does, it must be one that goes out of the thread containing $g(x'_j)$, which is contained (because g is an inclusion) in $\mathrm{Im}(g)$. (Q_0).1 rules out the possibility for this S-edge to be of the form $S(g(x_l), u'_l)$ (or $S(g(x_l), u''_l)$): from u'_l, the only way to reach $g(x_k)$ in $\leq 2n+4$ is through $S(g(x_l), u''_l)$, which contradicts the minimality of this path.

Moreover, it cannot be the S-edge landing on the other endpoint of the thread, since this would mean that the thread is of length $\leq 2n+4$, and since $\mathcal{R}' \in$ quasi-$\mathbb{H}_\sigma^{2n+6}$, the other endpoint is guaranteed to be $g(x''_j)$. In this case, there would be a shortest path from $g(x'_j)$ to u_j, which would directly borrow $S(u_j, g(x'_j))$.

Hence, the first S-edge can only be $S(u_j, g(x'_j))$, and (Q_0).1 ensures that the only way this would result in a path of length $\leq 2n+5$ is if the second edge it goes through is $S(u_j, g(x''_j))$, from which we can repeat the same reasoning to prove that from there, the path doesn't go through any S-edge.

The other possibility is that the path doesn't go through any S-edge. In either case, it means that $g(x'_j)$ or $g(x''_j)$ and $g(x_k)$ are on the same thread, and the shortest path follows the E-edges of this thread. Hence, a path as short exists in $\mathcal{W}_0$ between x'_j and x_k.

We now check that (Q_0).3 holds: let x'_j and x'_k (and similarly for x''_j and for x''_k) be such that

$$\mathrm{dist}_{\mathcal{T}_0}(g(x'_j), g(x'_k)) \leq 2n+5$$

and consider a shortest path from $g(x'_j)$ to $g(x'_k)$.

The same reasoning as before ensures that $g(x'_k)$ is on the same thread as $g(x'_j)$ or $g(x''_j)$, and that the shortest path follows the E-edges of that thread, which must then be of length $\leq 2n+5$ which in turn implies that $j = k$.

Now suppose that we have constructed $\mathcal{T}_i$ and let $\mathcal{T}_{i+1}$ be the n-guarded crossing-S-swap between $g(x_{i+1})$ and u_{i+1} in $\mathcal{T}_i$. Suppose that

$$\mathrm{dist}_{\mathcal{T}_i}(g(x_{i+1}), u_{i+1}) < 2n+5\,.$$

Then

$$\mathrm{dist}_{\mathcal{T}_i}(g(x_{i+1}), g(x'_{i+1})) \leq 2n+5\,,$$

and $(Q_i).2$ ensures that

$$\mathrm{dist}_{\mathcal{W}_i}(x_{i+1}, x'_{i+1}) \leq 2n+5$$

which, as seen above, is absurd.

Lemma 4.3.4 ensures that

$$\mathcal{T}_{i+1} \equiv_{\alpha}^{<\text{inv FO}} \mathcal{T}_i \text{ and } \mathcal{S}\mathrm{upp}_{n+1}(\mathcal{T}_{i+1}) = \mathcal{S}\mathrm{upp}_{n+1}(\mathcal{T}_i)\,.$$

Let's show that $\mathcal{T}_{i+1}$ verifies (Q_{i+1}).

$(Q_{i+1}).1$ is straightforward: we only need to note that the new S-edges appeared at distance at least $2n+4$ from every

$$A \in \{\{u_{i+2}\}, \{u'_{i+2}, u''_{i+2}\}, \cdots, \{u_r\}, \{u'_r, u''_r\}\}$$

in $\mathcal{T}_i \setminus S(A)$.

$(Q_{i+1}).2$: let $j, k > i+1$ and suppose that there is a path of length $l \leq 2n+5$ between $g(x'_j)$ (and similarly for $g(x''_j)$) and $g(x_k)$ in $\mathcal{T}_{i+1}$, and consider a shortest such path. Let's show that $\mathrm{dist}_{\mathcal{W}_{i+1}}(x'_j, x_k) \leq l$. If this path doesn't go through any of the new S-edges, $(Q_i).2$ allows us to conclude (any path going through $\bar{x}_{i+1}$ in $\mathcal{W}_i$ can now through x_{i+1} instead).

Otherwise, $(Q_i).1$ ensures that it doesn't go through $S(u_{i+1}, u'_{i+1})$ or $S(u_{i+1}, u''_{i+1})$. Thus we can decompose this path in a sequence of two (since it's a shortest path) paths valid in $\mathcal{T}_i$, joined either by the edge $S(g(x_{i+1}), g(x'_{i+1}))$, or $S(g(x_{i+1}), g(x''_{i+1}))$, or one then the other.

It is not possible for the path to be decomposable as

$$g(x'_j) \overset{p_1}{\rightsquigarrow} g(a_{i+1}) S g(x_{i+1}) \overset{p_2}{\rightsquigarrow} g(x_k)$$

(for $a_{i+1} \in \{x'_{i+1}, x''_{i+1}\}$), because p_1 would be a path of length at most $2n+5$ in $\mathcal{T}_i$ from $g(x'_j)$ to $g(a_{i+1})$, which contradicts $(Q_i).3$.

Hence the path can be decomposed as

$$g(x'_j) \overset{p_1}{\rightsquigarrow} g(x_{i+1}) S g(a_{i+1}) \overset{p_2}{\rightsquigarrow} g(x_k)$$

with p_1 and p_2, of respective length l_1 and l_2 (with $l = l_1 + l_2 + 1$) being valid in $\mathcal{T}_i$.

$(Q_i).2$ allows us to reflect p_1 as a path from x'_j to x_{i+1} in $\mathcal{W}_i$ of length at most l_1, and p_2 as a path from a_{i+1} to x_k in $\mathcal{W}_i$ of length at most l_2. Replacing $\bar{x}_{i+1}$ by x_{i+1} in those paths gives us paths at least as short valid in $\mathcal{W}_{i+1}$. We then link them with $S(x_{i+1}, a_{i+1}) \in \mathcal{W}_{i+1}$, and get

$$\mathrm{dist}_{\mathcal{W}_{i+1}}(x'_j, x_k) \leq l_1 + l_2 + 1 = l\,.$$

$(Q_{i+1}).3$: let $j \neq k > i+1$, $a_j \in \{x'_j, x''_j\}$ and $a_k \in \{x'_k, x''_k\}$.

Suppose that there is a path (take a shortest witness) p of length at most $2n+5$ between $g(a_j)$ and $g(a_k)$ in $\mathcal{T}_{i+1}$. Because of $(Q_i).3$, p cannot be valid in $\mathcal{T}_i$. Because of $(Q_i).1$, it cannot go through $S(u_{i+1}, u'_{i+1})$ or $S(u_{i+1}, u''_{i+1})$.

Hence, it must go through $S(g(x_{i+1}), g(x'_{i+1}))$ or $S(g(x_{i+1}), g(x''_{i+1}))$. It cannot go through both, for otherwise we could replace

$$g(x'_{i+1})Sg(x_{i+1})Sg(x''_{i+1})$$

in p by

$$g(x'_{i+1})Su_{i+1}Sg(x''_{i+1})$$

and get a path as short in $\mathcal{T}_i$.

We can decompose p either as

$$g(a_j) \overset{p_1}{\rightsquigarrow} g(a_{i+1})Sg(x_{i+1}) \overset{p_2}{\rightsquigarrow} g(a_k)$$

or, if it goes through the S-edge in the other direction, as

$$g(a_j) \overset{p_1}{\rightsquigarrow} g(x_{i+1})Sg(a_{i+1}) \overset{p_2}{\rightsquigarrow} g(a_k)$$

with $a_{i+1} \in \{x'_{i+1}, x''_{i+1}\}$ and p_1, p_2 valid in $\mathcal{T}_i$, and of length at most $2n+5$.

This is absurd since either p_1 or p_2 breaks $(Q_i).3$.

We set $\mathcal{Q}' := \mathcal{T}_r$ together with g, which have the desired properties. □

Proposition 4.6.3. $\forall \alpha \in \mathbb{N}, \exists l \in \mathbb{N}, \forall m \in \mathbb{N}, \exists n \in \mathbb{N}, \forall M \in \mathbb{N}, \exists K \in \mathbb{N}$ *such that*

for every abstract loop $\mathcal{U}_+ \in \mathbb{L}_\sigma^{n+1}$ *and every* $\mathcal{Q} \in$ *quasi-*$\mathbb{H}_\sigma^{n+1}$ *such that*

$$|U| \leq M \text{ and } (l+1) \cdot [\![\mathcal{E}_{n+1}(\mathcal{U})]\!] < [\![\mathcal{E}_{n+1}(\mathcal{Q})]\!]$$

and such that for every $(n+1)$*-neighborhood type* χ *that occurs in* $\mathcal{U}$, $|\mathcal{Q}|_\chi \geq K$,

there exists $\mathcal{Q}' \in$ *quasi-*$\mathbb{H}_\sigma^m$ *such that*

$$\mathcal{Q}' \equiv_\alpha^{<\text{inv FO}} \mathcal{Q} \text{ and } [\![\mathcal{E}_m(\mathcal{Q})]\!] = [\![\mathcal{E}_m(\mathcal{Q}')]\!] + [\![\mathcal{E}_m(\mathcal{U})]\!] .$$

Proof. The proof is based on the well known result that first-order formulas of quantifier-rank k cannot distinguish between a linear order of length 2^k and a linear order of length 2^k+1 (see, for instance, [27]). Hence if a loop is repeated at least 2^k+1 times, we can eliminate one instance without changing the $\equiv_k^{<\text{inv FO}}$ class of the structure.

First, we include many copies of the loop in $\mathcal{Q}$. The inclusion may not preserve S-edges: the next step is to re-associate these S-edges with crossing-S-swaps in order for these copies to be isomorphic. This is made possible by the hypothesis on the number of occurrences of neighborhood types appearing in $\mathcal{U}$: it gives us room to make sure the crossing-S-swaps are guarded.

Once this is done, we can remove one copy in a $<$-inv FO-indistinguishable way.

We can assume that $m \geq s(\alpha)$. Let m_1 be given by Lemma 4.6.2 from m, D be given by Lemma 4.6.2 from m and M and n be given by Lemma 4.4.3 from m_1. Set $K := D + (l+1)M$.

We construct $\mathcal{U}^l_+, \mathcal{U}^{l+1}_+ \in \mathbb{L}^{n+1}_\sigma$, such that

- $[\![\mathcal{E}_{n+1}(\mathcal{U}^l)]\!] = l \cdot [\![\mathcal{E}_{n+1}(\mathcal{U})]\!]$,
- $[\![\mathcal{E}_{n+1}(\mathcal{U}^{l+1})]\!] = (l+1) \cdot [\![\mathcal{E}_{n+1}(\mathcal{U})]\!]$,
- $\mathcal{U}^l_+ \equiv^{\text{<-inv FO}}_\alpha \mathcal{U}^{l+1}_+$.

Consider the FO-interpretation $\mathcal{J}$ (of arity 2 and depth d, independent of n) from the vocabulary $\Sigma_{n+1} \cup \{N, <\}$ (where N is a unary relational symbol) to $\Sigma_{n+1} \cup \{<\}$, which, given a structure $\mathcal{V}_+$, returns $\mathcal{J}(\mathcal{V}_+)$ as follows. For the sake of simplicity, we will name $1, \cdots, r$ the elements of $N^{\mathcal{V}_+}$ accordingly to $<^{\mathcal{V}_+}$.

- its domain is $\{1, \cdots, r\} \times (V \setminus N^{\mathcal{V}_+})$
- $\mathcal{J}(\mathcal{V}_+) \models S((i,x),(j,y))$ iff $i = j$ and $\mathcal{V}_+ \models S(x,y)$
- $\mathcal{J}(\mathcal{V}_+) \models E((i,x),(j,y)) \wedge E((j,y),(i,x))$ iff $i = j$ and $\mathcal{V} \models E(x,y)$, or $j = i+1$ and $\mathcal{V}_+ \models J^2_\tau(x)$ and $\mathcal{V}_+ \models J^1_\tau(y)$ for some τ
- for every τ, $\mathcal{J}(\mathcal{V}_+) \models J^1_\tau(i,x)$ iff $i = 1$ and $\mathcal{V}_+ \models J^1_\tau(x)$
- for every τ, $\mathcal{J}(\mathcal{V}_+) \models J^2_\tau(i,x)$ iff $i = r$ and $\mathcal{V}_+ \models J^2_\tau(x)$
- $<^{\mathcal{J}(\mathcal{V}_+)}$ is the lexicographical order

In other words, if we add r elements to the abstract loop $\mathcal{U}_+$, color them with N and add an order, its image by $\mathcal{J}$ is the r-fold concatenation of $\mathcal{U}_+$ to itself (in the same direction each time), with an order.

Fix an arbitrary order $<^U$ on U. For $r \in \mathbb{N}$, let $\mathcal{U}^{[r]}_+$ be the $\Sigma_{n+1} \cup \{N, <\}$-structure obtained by adding $\{1, \cdots, r\}$ to the domain of $\mathcal{U}_+$, interpreting N as $\{1, \cdots, r\}$ and ordering the elements as $1, \cdots, r$ and then accordingly to $<^U$.

Now let $(\mathcal{U}^r_+, <^r) := \mathcal{J}(\mathcal{U}^{[r]}_+)$. See Figure 4.29 for an example.

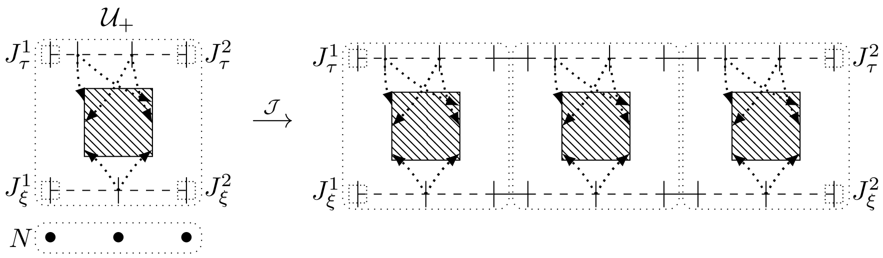

Figure 4.29: Application of $\mathcal{J}$ to $\mathcal{U}^{[3]}_+$. In this illustration, two jumping pair types τ and ξ are relevant in $\mathcal{U}_+$. The new order $<^3$ is the concatenation of the old ones.

If we choose $l := 2^{2\alpha+d}$, we have

$$\mathcal{U}_+^{[l]} \equiv^{\mathrm{FO}}_{2\alpha+d} \mathcal{U}_+^{[l+1]},$$

hence

$$(\mathcal{U}_+^l, <^l) \equiv^{\mathrm{FO}}_\alpha (\mathcal{U}_+^{l+1}, <^{l+1}),$$

which entails

$$\mathcal{U}_+^l \equiv^{<\text{-inv FO}}_\alpha \mathcal{U}_+^{l+1}.$$

By construction,

$$[\![\mathcal{E}_{n+1}(\mathcal{U}^l)]\!] = l \cdot [\![\mathcal{E}_{n+1}(\mathcal{U})]\!]$$

and

$$[\![\mathcal{E}_{n+1}(\mathcal{U}^{l+1})]\!] = (l+1) \cdot [\![\mathcal{E}_{n+1}(\mathcal{U})]\!].$$

By hypothesis, $[\![\mathcal{E}_{n+1}(\mathcal{U}^{l+1})]\!] < [\![\mathcal{E}_{n+1}(\mathcal{Q})]\!]$, thus we can apply Proposition 4.4.3 to get $\mathcal{R} \equiv^{<\text{-inv FO}}_\alpha \mathcal{Q}$ such that $[\![\mathcal{E}_{m_1+1}(\mathcal{R})]\!] = [\![\mathcal{E}_{m_1+1}(\mathcal{Q})]\!]$ and a (m_1+1)-inclusion h from $\mathcal{U}^{l+1}$ to $\mathcal{R}$.

Now, for every (m_1+1)-neighborhood type ξ occurring in $\mathcal{U}^{l+1}$, $|\mathcal{R}|_\xi = |\mathcal{Q}|_\xi \geq K$, hence $|\mathcal{R}|_{R\setminus \mathrm{Im}(h)}|_\xi \geq D$ by choice of K.

We can apply Lemma 4.6.2, which yields some $\mathcal{R}^{l+1} \equiv^{<\text{-inv FO}}_\alpha \mathcal{R}$ such that $\mathcal{S}\mathrm{upp}_m(\mathcal{R}^{l+1}) \simeq \mathcal{S}\mathrm{upp}_m(\mathcal{R})$, and g m-includes $\mathcal{U}^{l+1}$ in $\mathcal{R}^{l+1}$, and respects S.

Let $\mathcal{R}_+^{l+1}$ be the g-border-extension of $\mathcal{R}^{l+1}$.

Since $g(\mathcal{U}^{l+1})$ is S-stable in $\mathcal{R}^{l+1}$, we can decompose

$$\mathcal{R}_+^{l+1} \setminus \{E(x,y) : x,y \in R^{l+1}, i \in \{1,2\}, J^i_\tau(x) \wedge J^i_\tau(y)\}$$

as $g(\mathcal{U}_+^{l+1}) \uplus \mathcal{R}'_+$ for some Σ_n structure $\mathcal{R}'_+$, where $g(\mathcal{U}_+^{l+1})$ is the abstract loop based upon $g(\mathcal{U}^{l+1})$ such that g respects every J^i_τ.

Note that $\mathcal{R}^{l+1} = \mathcal{I}_m(g(\mathcal{U}_+^{l+1}) \uplus \mathcal{R}'_+)$, and let $\mathcal{R}^l := \mathcal{I}_m(g(\mathcal{U}_+^l) \uplus \mathcal{R}'_+)$.

$g(\mathcal{U}_+^l) \uplus \mathcal{R}'_+ \equiv^{<\text{-inv FO}}_\alpha g(\mathcal{U}_+^{l+1}) \uplus \mathcal{R}'_+$, hence $\mathcal{R}^l \equiv^{<\text{-inv FO}}_\alpha \mathcal{R}^{l+1}$.

Now, set $\mathcal{Q}' := \mathcal{R}^l$. We have that $\mathcal{Q}' \equiv^{<\text{-inv FO}}_\alpha \mathcal{Q}$, and, by construction,

$$\begin{aligned}
[\![\mathcal{E}_m(\mathcal{Q})]\!] &= [\![\mathcal{E}_m(\mathcal{U}^{l+1})]\!] + [\![\mathcal{E}_m(\mathcal{R}')]\!] \\
&= [\![\mathcal{E}_m(\mathcal{U})]\!] + [\![\mathcal{E}_m(\mathcal{U}^l)]\!] + [\![\mathcal{E}_m(\mathcal{R}')]\!] \\
&= [\![\mathcal{E}_m(\mathcal{U})]\!] + [\![\mathcal{E}_m(\mathcal{Q}')]\!].
\end{aligned}$$

□

4.6.3 S-parents re-association

We now turn to the last step of the proof of Theorem 4.2.1.

After the removal of the extra material in $\mathcal{Q}$, we have transformed our initial hollow tree $\mathcal{Q}$ into a hollow quasitree having the same number of occurrences of any neighborhood type as the initial $\mathcal{P}$. They both have the same threads but may differ with their S-edges. The following proposition states that they are $\equiv^{<\text{-inv FO}}_\alpha$, thus ending the proof of Theorem 4.2.1.

The techniques used in the proof of the following proposition are strongly reminiscent of those used in [4]; it requires a notion of vertical-S-swaps adapted to hollow trees.

The notion of *vertical swap* in a tree has been introduced in [5] and is a crucial operation in their proof. We need here a version of these vertical swaps adapted to hollow trees. Unlike the other swaps, vertical-S-swap preserve hollow trees. In the following, $\mathcal{T}$ is a hollow tree on σ.

We start by defining classical notions making use of the tree structure of $\mathcal{T}$.

The (strict) ancestor relation within a hollow tree is inherited from the original tree and is denoted by $x \preceq y$ (resp. $x \prec y$). Note that this relation is not part of the schema and not expressible in FO from E and S.

Let x, y be two nodes of T such that $x \prec y$. We define the **context** $\mathcal{C}_\mathcal{T}(x, y)$ **at** x **and** y **in** $\mathcal{T}$ (referred using the simplified notation $\mathcal{C}$ in the following) as the substructure of $\mathcal{T}$ induced by the set $\{z \in T : x \prec z \ \wedge \ y \not\prec z\}$, with three distinguished nodes colored by two new unary predicates $\top$ and $\bot$: the S-children x' and x'' of x are $\mathcal{C}$'s **top-anchors** ($\top^\mathcal{C} = \{x', x''\}$), and y its **bottom-anchor** ($\bot^\mathcal{C} = \{y\}$). The set $V(\mathcal{C}) := \{z \in C : z \preceq y\}$ is the set of **vertebræ** of $\mathcal{C}$. The **height** height$(\mathcal{C})$ is $|V(\mathcal{C})|$ and correspond to the difference of depth between y and x. Given $n \in \mathbb{N}$, $\mathcal{C}$'s n**-skeleton**, denoted $\mathrm{Sk}_n(\mathcal{C})$, is the substructure of $\mathcal{C}$ induced by the nodes at distance at most n of $V(\mathcal{C})$, of S-children of nodes of $V(\mathcal{C})$, or of $\mathcal{C}$'s top-anchors. Additionally, $\mathrm{Sk}_n(\mathcal{C})$ inherits the restriction of $\prec$ to $V(\mathcal{C})$. Two contexts are said to be n**-similar** if their n-skeletons are isomorphic. Given two contexts $\mathcal{C}$ and $\mathcal{D}$, we denote by $\mathcal{C} \cdot \mathcal{D}$ the context obtained as the disjoint union of $\mathcal{C}$ and $\mathcal{D}$, with an S-edge from $\mathcal{C}$'s bottom-anchor to each of $\mathcal{D}$'s top-anchor, and where the anchors are redefined in the natural way: $\top^{\mathcal{C}\cdot\mathcal{D}} := \top^\mathcal{C}$ and $\bot^{\mathcal{C}\cdot\mathcal{D}} := \bot^\mathcal{D}$. Similarly, we define the **prefix** $\mathcal{P}_\mathcal{T}(y)$ **at** y **in** $\mathcal{T}$ as the substructure of $\mathcal{T}$ induced by $\{z \in T : y \not\prec z\}$ (the only additional relation being $\bot$), and the **suffix** $\mathcal{S}_\mathcal{T}(x)$ **at** x **in** $\mathcal{T}$ as the substructure of $\mathcal{T}$ induced by $\{z \in T : x \prec z\}$ (here, the only additional relation is $\top$). The concatenation between a prefix and a context, a prefix and a suffix, and a context and a suffix are defined in the natural way (and results respectively in a prefix, a hollow tree, and a suffix). Concatenation is associative.

Let

$$x \prec x_A \prec x_B \prec x_C \in T$$

and

$$(x', x''), (x'_A, x''_A), (x'_B, x''_B), (x'_C, x''_C)$$

be their respective S-children.

Suppose that

$$\mathrm{tp}^k_\mathcal{T}(x, x', x'') = \mathrm{tp}^k_\mathcal{T}(x_B, x'_B, x''_B)$$

and

$$\mathrm{tp}^k_\mathcal{T}(x_A, x'_A, x''_A) = \mathrm{tp}^k_\mathcal{T}(x_C, x'_C, x''_C)\,.$$

Let us define $\mathcal{P} := \mathcal{P}_\mathcal{T}(x)$, $\mathcal{A} := \mathcal{C}_\mathcal{T}(x, x_A)$, $\mathcal{B} := \mathcal{C}_\mathcal{T}(x_A, x_B)$, $\mathcal{C} := \mathcal{C}_\mathcal{T}(x_B, x_C)$ and $\mathcal{S} := \mathcal{S}_\mathcal{T}(x_C)$. With these definitions, $\mathcal{T} = \mathcal{P} \cdot \mathcal{A} \cdot \mathcal{B} \cdot \mathcal{C} \cdot \mathcal{S}$.

In this case, $\mathcal{T}' := \mathcal{P} \cdot \mathcal{C} \cdot \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{S}$ is called the k**-guarded vertical-**S**-swap between** $[x, x_A]$ **and** $[x_B, x_c]$ **in** $\mathcal{T}$, c.f. Figure 4.30.

We wish to show the following lemma:

Lemma 4.6.4. *For all* $\alpha \in \mathbb{N}$*, there exists* $N \in \mathbb{N}$ *such that*

for every hollow tree $\mathcal{T}$ *on* σ*, if* $\mathcal{T}'$ *is the* N*-guarded vertical-*S*-swap between* $[x, x_A]$ *and* $[x_B, x_C]$ *in* $\mathcal{T}$*,*

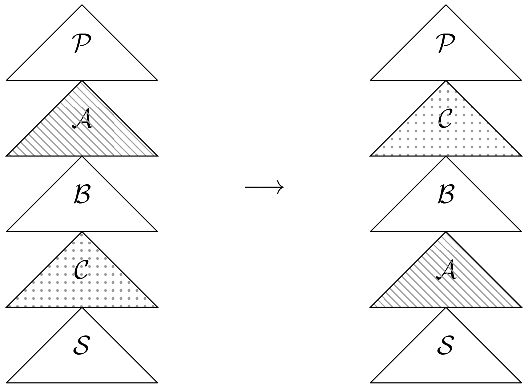

Figure 4.30: A vertical-S-swap from $\mathcal{T}$ to $\mathcal{T}'$.

then every node in T has the same $(N+1)$-neighborhood type in $\mathcal{T}$ and in $\mathcal{T}'$, and $\mathcal{T}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{T}$.

It is immediate to check that an N-guarded vertical-S-swap preserves $(N+1)$-neighborhood types. The following is devoted to the proof that it also preserves $\equiv_\alpha^{<\text{-inv FO}}$, concluding the proof of Lemma 4.6.4.

We start by proving a special case of Lemma 4.6.4. We will reduce the general case to it. This case is illustrated in Figure 4.31.

Lemma 4.6.5. *For all $\alpha \in \mathbb{N}$, there exists $M \in \mathbb{N}$ such that the following holds.*

Let $\mathcal{T} \in \mathbb{H}_\sigma$ and

$$x \prec x_A \prec x_B \in T$$

having for respective S-children

$$(x', x''), (x'_A, x''_A) \text{ and } (x'_B, x''_B)$$

Suppose that

$$tp^M_{\mathcal{T}}(x, x', x'') = tp^M_{\mathcal{T}}(x_A, x'_A, x''_A) = tp^M_{\mathcal{T}}(x_B, x'_B, x''_B)\,.$$

Let $\mathcal{P} := \mathcal{P}_{\mathcal{T}}(x)$, $\mathcal{A} := \mathcal{C}_{\mathcal{T}}(x, x_A)$, $\mathcal{B} := \mathcal{C}_{\mathcal{T}}(x_A, x_B)$ and $\mathcal{S} := \mathcal{S}_{\mathcal{T}}(x_B)$.

Then

$$\mathcal{P} \cdot \mathcal{A} \cdot \mathcal{B} \cdot \mathcal{S} \equiv_\alpha^{<\text{-inv FO}} \mathcal{P} \cdot \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{S}\,.$$

Proof. We will first set $N \in \mathbb{N}$ instead of M, that will be sufficient for most cases. Then, we will define $M \geq N$ which will work for all cases.

Recall the function o_p^Σ introduced in Lemma 4.3.2, and consider

$$n := o_3^\Sigma(\alpha + c)\,,$$

where c is to be chosen later on, and $\Sigma := P_\sigma \cup \{E, S, P_{1/2}, P_6\}$ where $P_{1/2}$ and P_6 are new unary symbols. We distinguish between several cases depending on whether x, x_A and x_B are close or not, where "close" is relative to n.

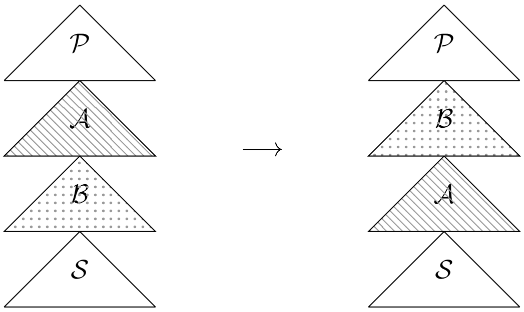

Figure 4.31: vertical-S-swap (special case) from $\mathcal{P}\cdot\mathcal{A}\cdot\mathcal{B}\cdot\mathcal{S}$ to $\mathcal{P}\cdot\mathcal{B}\cdot\mathcal{A}\cdot\mathcal{S}$.

1. Assume first that $\mathrm{tp}^n_{\mathcal{A}}(x', x'', x_A) = \mathrm{tp}^n_{\mathcal{B}}(x'_A, x''_A, x_B)$.

   This case covers the instances where $\mathcal{A}$ and $\mathcal{B}$ are n-similar, as well as those where $\mathrm{dist}_{\mathcal{T}}(x, x_A)$ and $\mathrm{dist}_{\mathcal{T}}(x_A, x_B)$ are $> 2n+2$.

   Consider the extension $\mathcal{T}^-$ of $\mathcal{P}\uplus\mathcal{A}\uplus\mathcal{B}\uplus\mathcal{S}$ to Σ where the interpretation of $P_{1/2}$ only contains the bottom-anchor of $\mathcal{P}$, and that of P_6 contains the top-anchors of $\mathcal{S}$. Since $P^{\mathcal{T}^-}_{1/2}$ and $P^{\mathcal{T}^-}_6$ are at distance $+\infty$ from $\mathcal{A}$ and $\mathcal{B}$, $\mathrm{tp}^n_{\mathcal{T}^-}(x', x'', x_A) = \mathrm{tp}^n_{\mathcal{T}^-}(x'_A, x''_A, x_B)$.

   Hence, we can apply Lemma 4.3.2, and get two orders $<_{AB}$ (whose first elements are $x', x'', x_A, x'_A, x''_A, x_B$) and $<_{BA}$ (whose first elements are $x'_A, x''_A, x_B, x', x'', x_A$) such that $(\mathcal{T}^-, <_{AB}) \equiv^{\mathrm{FO}}_{\alpha+c} (\mathcal{T}^-, <_{BA})$.

   Now, consider the FO-interpretation that adds an S-edge between u and v if either

   - $P_{1/2}(u)$ and v is either the first or the second element of $<$
   - or u is the third element of $<$ and v is either its fourth or fifth one
   - or u is the sixth element of $<$ and $P_6(v)$,

   and then forgets about $P_{1/2}$ and P_6.

   Take c to be the depth of this FO-interpretation (which has arity 1).

   Note that the result of this FO-interpretation of $(\mathcal{T}^-, <_{AB})$ is an ordered extension of $\mathcal{P}\cdot\mathcal{A}\cdot\mathcal{B}\cdot\mathcal{S}$ and that its result on $(\mathcal{T}^-, <_{BA})$ is an ordered extension of $\mathcal{P}\cdot\mathcal{B}\cdot\mathcal{A}\cdot\mathcal{S}$.

   This entails $\mathcal{P}\cdot\mathcal{A}\cdot\mathcal{B}\cdot\mathcal{S} \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{P}\cdot\mathcal{B}\cdot\mathcal{A}\cdot\mathcal{S}$.

2. Assume next that $\mathcal{B}$ can be decomposed as $\mathcal{B}_1\cdots\mathcal{B}_k$, where each $\mathcal{B}_i$ is n-similar to $\mathcal{A}$.

   We can then apply k times Case 1 and obtain

   $$\mathcal{P}\cdot\mathcal{A}\cdot\mathcal{B}\cdot\mathcal{S} \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{P}\cdot\mathcal{B}\cdot\mathcal{A}\cdot\mathcal{S}$$

   as desired.

3. From now on, $N \geq l(2n+2)+n$ for a large enough l to be chosen later on.

As we are not in Case 1, we can restrict our study to the cases where $\mathrm{dist}_{\mathcal{T}}(x, x_A) \leq 2n + 2$ (the cases where $\mathrm{dist}_{\mathcal{T}}(x_A, x_B) \leq 2n + 2$ can be treated similarly).

We will need the following claims. The first one is just a simple observation.

Claim 4.6.6. *Let $\mathcal{U}, \mathcal{V}$ be two n-similar contexts of $\mathcal{T}$.*
For every decomposition

$$\mathcal{U} = \mathcal{U}_1 \cdots \mathcal{U}_p \,,$$

there exists a decomposition

$$\mathcal{V} = \mathcal{V}_1 \cdots \mathcal{V}_p$$

such that for every i, $\mathcal{U}_i$ and $\mathcal{V}_i$ are n-similar.

Proof. Let φ be an isomorphism from $\mathrm{Sk}_n(\mathcal{U})$ to $\mathrm{Sk}_n(\mathcal{V})$, and

$$x_0 \prec x_1 \prec \cdots \prec x_p \in T$$

be such that $\mathcal{U}_i = \mathcal{C}_{\mathcal{T}}(x_{i-1}, x_i)$. Since φ is $\prec$-monotonous on $V(\mathcal{U})$, the $\mathcal{V}_i = \mathcal{C}_{\mathcal{T}}(\varphi(x_{i-1}), \varphi(x_i))$ are well-defined.

We have that $\mathcal{V} = \mathcal{V}_1 \cdots \mathcal{V}_p$; it remains to show that for every i, $\mathcal{U}_i$ and $\mathcal{V}_i$ are n-similar. Again, the $\prec$-monotonicity of φ entails that

$$\varphi(\mathrm{Sk}_n(\mathcal{U}_i)) \subseteq \mathrm{Sk}_n(\mathcal{V}_i) \,,$$

which allows us to conclude. ☐

The next one is a variant of Lyndon-Schtzenberger Theorem stated for contexts of hollow trees instead of words.

Claim 4.6.7. *Let $n \in \mathbb{N}$, let $\mathcal{T} \in \mathbb{H}_\sigma$, let $x \prec y \prec z \prec t$ be nodes of T, and let $\mathcal{U} := \mathcal{C}_{\mathcal{T}}(x, y)$, $\mathcal{V} := \mathcal{C}_{\mathcal{T}}(y, z)$ and $\mathcal{W} := \mathcal{C}_{\mathcal{T}}(z, t)$ such that $\mathcal{U}$ and $\mathcal{W}$ are n-similar, $\mathcal{U} \cdot \mathcal{V}$ and $\mathcal{V} \cdot \mathcal{W}$ are n-similar.*
Then there exist decompositions

$$\mathcal{U} = \mathcal{U}_1 \cdots \mathcal{U}_p \,,$$

$$\mathcal{V} = \mathcal{V}_1 \cdots \mathcal{V}_q$$

and

$$\mathcal{W} = \mathcal{W}_1 \cdots \mathcal{W}_p$$

where all the $\mathcal{U}_i, \mathcal{V}_i$ and $\mathcal{W}_i$ are n-similar.

Proof. We define θ_n which maps two successive vertebræ $x_i \prec x_{i+1}$ of a context $\mathcal{U}$ to the neighborhood type $\mathrm{tp}^n_{\mathcal{C}_{\mathcal{U}}(x_i, x_{i+1})}(x'_i, x''_i, x_{i+1})$, where x'_i and x''_i are the S-children of x_i.

Now, let Θ_n be the monoid morphism from contexts to words extending θ_n; that is, if $x_0 \prec \cdots \prec x_d$ are all the vertebræ of $\mathcal{U}$, then

$$|\Theta_n(\mathcal{U})| = \mathrm{height}(\mathcal{U}) = d \,,$$

and the ith letter of $\Theta_n(\mathcal{U})$ is $\theta_n(x_i, x_{i+1})$.

Let $u := \Theta_n(\mathcal{U})$, $v := \Theta_n(\mathcal{V})$ and $w := \Theta_n(\mathcal{W})$.

By n-similarity of $\mathcal{U}$ and $\mathcal{W}$, $u = w$, and by n-similarity of $\mathcal{U}\cdot\mathcal{V}$ and $\mathcal{V}\cdot\mathcal{W}$, $uv = vw$. Hence $uv = vu$.

By Lyndon-Schtzenberger Theorem, there must exist a word a and integers p, q such that $u = w = a^p$ and $v = a^q$ [28].

We can decompose $\mathcal{U}, \mathcal{V}$ and $\mathcal{W}$ alongside those decompositions of u, v and w, to get $\mathcal{U} = \mathcal{U}_1 \cdots \mathcal{U}_p$, $\mathcal{V} = \mathcal{V}_1 \cdots \mathcal{V}_q$, and $\mathcal{W} = \mathcal{W}_1 \cdots \mathcal{W}_p$, where all the $\mathcal{U}_i, \mathcal{V}_i$ and $\mathcal{W}_i$ are mapped to a by Θ_n, hence are n-similar. □

Let ϕ be an isomorphism between the N-neighborhood of x and that of x_A.

As $\text{dist}_{\mathcal{T}}(x, x_A) \leq 2n + 2$, x_A is in the N-neighborhood of x and set $x_0 := x_A$ and $x_1 := \phi(x_A)$. Construct by induction $x_{i+1} := \phi(x_i)$ until $i > l$. Our choice of N ensures that x_i is well defined as x_{i-1} remains in the N-neighborhood of x. We claim that for all $j \leq l$,

$$X_j := \mathcal{C}_{\mathcal{T}}(x_{j-1}, x_j)$$

is n-similar to $\mathcal{A}$. This is a simple consequence of the fact that the n-skeleton of of X_j is included into the N-neighborhood of x.

Likewise, starting from x_B instead of x_A, we show that there exist

$$y_1, \cdots, y_l \in T$$

such that for $j \in [1, l]$ (and with the convention that $y_0 := x_B$),

$$Y_j := \mathcal{C}_{\mathcal{T}}(y_j, y_{j-1})$$

is n-similar to $\mathcal{A}$.

We distinguish several cases.

(a) Suppose that $\text{dist}_{\mathcal{T}}(x_A, x_B) \geq 2N$. This ensures that all the $(x_i)_{i\geq 1}$ and $(y_i)_{i\geq 0}$ belong to B.

- Suppose $x_{l-1} \prec y_l$.
  If we let $\mathcal{C} := \mathcal{C}_{\mathcal{T}}(x_{l-1}, y_l)$, we can decompose $\mathcal{B}$ as

  $$\mathcal{X}_1 \cdots \mathcal{X}_{l-1} \cdot \mathcal{C} \cdot \mathcal{Y}_l \cdots \mathcal{Y}_1 .$$

  If l is chosen large enough, namely $l \geq 2n + 4$, the following decomposition of $\mathcal{T}$:

  $$\mathcal{P} \cdot \underbrace{\mathcal{A} \cdot \mathcal{X}_1 \cdots \mathcal{X}_{l-1} \cdot \mathcal{C} \cdot \mathcal{Y}_l} \cdot \underbrace{\mathcal{Y}_{l-1} \cdots \mathcal{Y}_1} \cdot \mathcal{S}$$

  falls in Case 1 of this Lemma and therefore the following equation holds:

  $$\begin{aligned} & \mathcal{P} \cdot \underbrace{\mathcal{A} \cdot \mathcal{X}_1 \cdots \mathcal{X}_{l-1} \cdot \mathcal{C} \cdot \mathcal{Y}_l} \cdot \underbrace{\mathcal{Y}_{l-1} \cdots \mathcal{Y}_1} \cdot \mathcal{S} \\ \equiv_{\alpha}^{<\text{-inv FO}} \quad & \mathcal{P} \cdot \underbrace{\mathcal{Y}_{l-1} \cdots \mathcal{Y}_1} \cdot \underbrace{\mathcal{A} \cdot \mathcal{X}_1 \cdots \mathcal{X}_{l-1} \cdot \mathcal{C} \cdot \mathcal{Y}_l} \cdot \mathcal{S} \end{aligned}$$

Likewise, we get:

$$\begin{aligned} & \mathcal{P} \cdot \underbrace{\mathcal{Y}_{l-1} \cdots \mathcal{Y}_1 \cdot \mathcal{A}} \cdot \underbrace{\mathcal{X}_1 \cdots \mathcal{X}_{l-1} \cdot \mathcal{C} \cdot \mathcal{Y}_l} \cdot \mathcal{S} \\ \equiv_\alpha^{<\text{-inv FO}} \quad & \mathcal{P} \cdot \underbrace{\mathcal{X}_1 \cdots \mathcal{X}_{l-1} \cdot \mathcal{C} \cdot \mathcal{Y}_l} \cdot \underbrace{\mathcal{Y}_{l-1} \cdots \mathcal{Y}_1 \cdot \mathcal{A}} \cdot \mathcal{S} \end{aligned}$$

Hence, we have

$$\mathcal{P} \cdot \mathcal{A} \cdot \mathcal{B} \cdot \mathcal{S} \equiv_\alpha^{<\text{-inv FO}} \mathcal{P} \cdot \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{S}$$

as desired.

- Suppose now that $x_{l-1} \not\prec y_l$.
  Because $\text{dist}_\mathcal{T}(x_A, x_B) \geq 2N$, we know that $y_l \not\prec x_{l-1}$: let $\mathcal{T}_1$ be the n-guarded crossing-S-swap between x_{l-1} and y_l in $\mathcal{T}$. Lemma 4.3.4 (we can always assume that $n \geq s(\alpha)$) ensures that $\mathcal{T}_1 \equiv_\alpha^{<\text{-inv FO}} \mathcal{T}$ and $\forall z \in T$, $\text{tp}_{\mathcal{T}_1}^{n+1}(z) = \text{tp}_{\mathcal{T}}^{n+1}(z)$. We are now in the situation to apply Case 2. It only remains to do again the n-guarded crossing-S-swap between x_{l-1} and y_l afterwards to derive the desired $\mathcal{T}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{T}$.

(b) Suppose now that $\text{dist}_\mathcal{T}(x_A, x_B) < 2N$. Set

$$M := (2N+1)(2n+2) + n\,.$$

Just as before (by replacing l with $2N+1$), we define

$$x_0, \cdots, x_{2N+1} \text{ and } y_0, \cdots, y_{2N+1}$$

and, accordingly,

$$\mathcal{X}_1, \cdots, \mathcal{X}_{2N+1} \text{ and } \mathcal{Y}_1, \cdots, \mathcal{Y}_{2N+1}$$

that all are n-similar to $\mathcal{A}$.

There are at most $2N$ vertebræ in $\mathcal{B}$, hence not all of the $(y_i)_{0 \leq i \leq 2N}$ can be in $\mathcal{B}$. Let k be the smallest index such that y_k is not a vertebrate of $\mathcal{B}$ (we know that $1 \leq k \leq 2N$). Since $x_A \prec y_{k-1}$ and $y_k \prec y_{k-1}$, x_A and y_k must be related by $\preceq$; by definition of k, we must have $y_k \preceq x_A$. If $y_k = x_A$, we can conclude using Case 2. Otherwise, $y_k \prec x_A \prec y_{k-1}$.

Likewise, either $x \prec y_k$ or $y_k \preceq x$. By n-similarity of $\mathcal{A}$ and $\mathcal{Y}_k$, and because $\text{dist}_\mathcal{T}(x, x_A) \leq 2n+2$, we know that $\text{height}(\mathcal{A}) = \text{height}(\mathcal{Y}_k)$. Hence, it cannot be the case that $y_k \preceq x$.

We now have $x \prec y_k \prec x_A \prec y_{k-1}$. Let

$$\mathcal{U} := C_\mathcal{T}(x, y_k), \quad \mathcal{V} := C_\mathcal{T}(y_k, x_A) \quad \text{and} \quad \mathcal{W} := C_\mathcal{T}(x_A, y_{k-1})\,.$$

We have that $\mathcal{U} \cdot \mathcal{V} = \mathcal{A}$ and $\mathcal{V} \cdot \mathcal{W} = \mathcal{Y}_k$ are n-similar.

To see that $\mathcal{U}$ and $\mathcal{W}$ are n-similar, look at $\mathcal{Y}_{k+1}$: there is an isomorphism φ from $\text{Sk}_n(\mathcal{Y}_k)$ to $\text{Sk}_n(\mathcal{Y}_{k+1})$, which is by definition $\prec$-monotonous on $V(\mathcal{Y}_k)$. Hence φ sends any vertebrate of $\mathcal{Y}_k$ to the vertebrate of $\mathcal{Y}_{k-1}$ whose depth is $\text{height}(\mathcal{A})$ smaller. This entails $\varphi(x_A) = x$, and by restricting φ, $\mathcal{W}$ and $\mathcal{U}$ are n-similar.

We can now apply Claim 4.6.7, and get decompositions $\mathcal{U} = \mathcal{U}_1 \cdots \mathcal{U}_p$, $\mathcal{V} = \mathcal{V}_1 \cdots \mathcal{V}_q$, and $\mathcal{W} = \mathcal{W}_1 \cdots \mathcal{W}_p$, where all the $\mathcal{U}_i, \mathcal{V}_i$ and $\mathcal{W}_i$ are n-similar.

Hence, $\mathcal{A}$ can be decomposed as $\mathcal{U}_1 \cdots \mathcal{U}_p \cdot \mathcal{V}_1 \cdots \mathcal{V}_q$, and such a decomposition can be transposed as in Claim 4.6.6 onto each $\mathcal{Y}_i, 0 < i < k$, as $\mathcal{Y}_i = \mathcal{Y}_1^i \cdots \mathcal{Y}_{p+q}^i$, where all the $\mathcal{Y}_j^i$, the $\mathcal{U}_i$, the $\mathcal{V}_i$ and the $\mathcal{W}_i$ are n-similar. $\mathcal{T}$ can thus be decomposed as

$$\mathcal{P} \cdot \underbrace{\mathcal{U}_1 \cdots \mathcal{U}_p \cdot \mathcal{V}_1 \cdots \mathcal{V}_q}_{\mathcal{A}} \cdot \underbrace{\underbrace{\mathcal{W}_1 \cdots \mathcal{W}_p}_{\mathcal{W}} \cdot \underbrace{\mathcal{Y}_1^{k-1} \cdots \mathcal{Y}_{p+q}^{k-1}}_{\mathcal{Y}_{k-1}} \cdots \underbrace{\mathcal{Y}_1^1 \cdots \mathcal{Y}_{p+q}^1}_{\mathcal{Y}_1}}_{\mathcal{B}} \cdot \mathcal{S}\,.$$

Now, we can use Case 2 with $\mathcal{A} := \mathcal{V}_q$ and derive that $\mathcal{T}$ is $\equiv_\alpha^{<\text{-inv FO}}$ to

$$\mathcal{P} \cdot \mathcal{U}_1 \cdots \mathcal{U}_p \cdot \mathcal{V}_1 \cdots \mathcal{V}_{q-1} \cdot \mathcal{W}_1 \cdots \mathcal{W}_p \cdot \mathcal{Y}_1^{k-1} \cdots \mathcal{Y}_{p+q}^{k-1} \cdots \mathcal{Y}_1^1 \cdots \mathcal{Y}_{p+q}^1 \cdot \mathcal{V}_q \cdot \mathcal{S}\,.$$

Repeating this operation $p + q - 1$ times allows us to conclude that

$$\mathcal{P} \cdot \mathcal{A} \cdot \mathcal{B} \cdot \mathcal{S} \equiv_\alpha^{<\text{-inv FO}} \mathcal{P} \cdot \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{S}\,.$$

□

We are now ready to conclude the proof of Lemma 4.6.4. As in the proof of Lemma 4.6.5, we distinguish between two cases. Let $n := o_3^\Sigma(\alpha + c)$ where c is the depth of some FO-interpretation to be specified later on, and

$$\Sigma := P_\sigma \cup \{E, S, P_{1/2}, P_3, P_{4/5}, P_6\}$$

where $P_{1/2}$, P_3, $P_{4/5}$ and P_6 are new unary symbols.

1. Assume first that $\text{tp}_\mathcal{A}^n(x', x'', x_A) = \text{tp}_\mathcal{C}^n(x'_B, x''_B, x_C)$.

   This case covers the instances where $\mathcal{A}$ and $\mathcal{C}$ are n-similar, as well as those where $\text{dist}_\mathcal{T}(x, x_A)$ and $\text{dist}_\mathcal{T}(x_B, x_C)$ are larger than $2n + 2$.

   Consider the extension $\mathcal{T}^-$ of $\mathcal{P} \uplus \mathcal{A} \uplus \mathcal{B} \uplus \mathcal{C} \uplus \mathcal{S}$ to Σ where $P_{1/2}^{\mathcal{T}^-} := \{x\}$, $P_3^{\mathcal{T}^-} := \{x'_A, x''_A\}$, $P_{4/5}^{\mathcal{T}^-} := \{x_B\}$ and $P_6^{\mathcal{T}^-} := \{x'_C, x''_C\}$.

   Since $P_{1/2}^{\mathcal{T}^-}$, $P_3^{\mathcal{T}^-}$, $P_{4/5}^{\mathcal{T}^-}$ and $P_6^{\mathcal{T}^-}$ are at distance $+\infty$ from $\mathcal{A}$ and $\mathcal{C}$, we have

   $$\text{tp}_{\mathcal{T}^-}^n(x', x'', x_A) = \text{tp}_{\mathcal{T}^-}^n(x'_B, x''_B, x_C)\,.$$

   Hence, we can apply Lemma 4.3.2, and get two orders $<_{AC}$ (whose first elements are $x', x'', x_A, x'_B, x''_B, x_C$) and $<_{CA}$ (whose first elements are $x'_B, x''_B, x_C, x', x'', x_A$) such that $(\mathcal{T}^-, <_{AC}) \equiv_{\alpha+c}^{\text{FO}} (\mathcal{T}^-, <_{CA})$.

   Now, consider the FO-interpretation that adds an S-edge between u and v if either

   - $P_{1/2}(u)$ and v is either the first or the second element of $<$
   - or u is the third element of $<$ and $P_3(v)$

- or $P_{4/5}(u)$ and v is either the fourth or the fifth element of $<$
- or u is the sixth element of $<$ and $P_6(v)$,

and then forgets about $P_{1/2}$, P_3, $P_{4/5}$ and P_6.

Take c to be the depth of this FO-interpretation (which has arity 1).

Note that the result of this FO-interpretation on $(\mathcal{T}^-, <_{AC})$ is an ordered extension of $\mathcal{P} \cdot \mathcal{A} \cdot \mathcal{B} \cdot \mathcal{C} \cdot \mathcal{S}$ and that its result on $(\mathcal{T}^-, <_{CA})$ is an ordered extension of $\mathcal{P} \cdot \mathcal{C} \cdot \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{S}$.

This entails $\mathcal{P} \cdot \mathcal{A} \cdot \mathcal{B} \cdot \mathcal{C} \cdot \mathcal{S} \equiv_\alpha^{<\text{-inv FO}} \mathcal{P} \cdot \mathcal{C} \cdot \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{S}$.

2. We can now, without loss of generality, assume that $\text{dist}_\mathcal{T}(x, x_A) \leq 2n+2$.

   Set $N := (2n+2) + m$ with $m := \max(M, n, s(\alpha))$, where M is given in Lemma 4.6.5.

   Consider an isomorphism φ from $\mathcal{N}_\mathcal{T}^N(x, x', x'')$ to $\mathcal{N}_\mathcal{T}^N(x_B, x'_B, x''_B)$.

   By choice of N,

$$\text{tp}_\mathcal{T}^m(x_A, x'_A, x''_A) = \text{tp}_\mathcal{T}^m(\varphi(x_A), y', y''),$$

   where y' and y'' are the S-children of $\varphi(x_A)$.

   Since $x_B \prec \varphi(x_A)$ and $\varphi(x_A) \neq x_C$ (for otherwise we would be in Case 1), there are only three subcases to consider:

   - if $x_B \prec \varphi(x_A) \prec x_C$, set

$$\mathcal{C}' := \mathcal{C}_\mathcal{T}(x_B, \varphi(x_A))$$

     and

$$\mathcal{X} := \mathcal{C}_\mathcal{T}(\varphi(x_A), x_C).$$

     We then have

$$\mathcal{T} = \mathcal{P} \cdot \mathcal{A} \cdot \mathcal{B} \cdot \mathcal{C}' \cdot \mathcal{X} \cdot \mathcal{S}.$$

     Let

$$\mathcal{T}_1 = \mathcal{P} \cdot \mathcal{C}' \cdot \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{X} \cdot \mathcal{S}$$

     be the m-guarded vertical-S-swap between $[x, x_A]$ and $[x_B, \varphi(x_A)]$ in $\mathcal{T}$. This swap falls under the scope of Case 1 since $m \geq n$, hence

$$\mathcal{T}_1 \equiv_\alpha^{<\text{-inv FO}} \mathcal{T} \text{ and } \forall z \in T, \text{tp}_{\mathcal{T}_1}^{m+1}(z) = \text{tp}_\mathcal{T}^{m+1}(z).$$

     Hence we are in the conditions (since $m \geq M$) to apply Lemma 4.6.5 on $\varphi(x_A) \prec x_A \prec x_C$ in $\mathcal{T}_1$ and get

$$\mathcal{T}_2 := \mathcal{P} \cdot \mathcal{C}' \cdot \mathcal{X} \cdot \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{S} \equiv_\alpha^{<\text{-inv FO}} \mathcal{T}_1.$$

     Note that $\mathcal{T}' = \mathcal{T}_2$, which implies that $\mathcal{T}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{T}$.
     These sequence of operations is depicted in Figure 4.32.

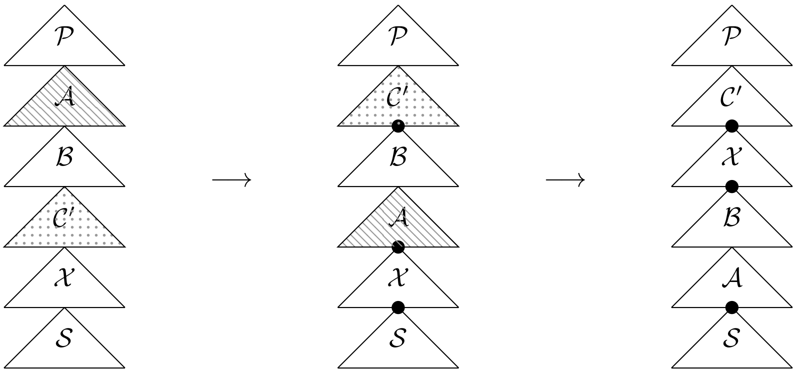

Figure 4.32: The swaps solving the case $x_B \prec \varphi(x_A) \prec x_C$. The second operation swaps the segments between the dark nodes using Lemma 4.6.5.

- if $x_B \prec x_C \prec \varphi(x_A)$, set

$$\mathcal{C}' := \mathcal{C}_{\mathcal{T}}(x_B, \varphi(x_A))$$

and

$$\mathcal{X} := \mathcal{C}_{\mathcal{T}}(\varphi(x_A), x_C)\,.$$

We then have

$$\mathcal{T} = \mathcal{P} \cdot \mathcal{A} \cdot \mathcal{B} \cdot \mathcal{C} \cdot \mathcal{X} \cdot \mathcal{S}'\,.$$

Let

$$\mathcal{T}_1 = \mathcal{P} \cdot \mathcal{C} \cdot \mathcal{X} \cdot \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{S}'$$

be the m-guarded vertical-S-swap between $[x, x_A]$ and $[x_B, \varphi(x_A)]$ in $\mathcal{T}$. This swap falls under the scope of Case 1 since $m \geq n$, hence

$$\mathcal{T}_1 \equiv_{\alpha}^{<\text{-inv FO}} \mathcal{T} \text{ and } \forall z \in T, \text{tp}_{\mathcal{T}_1}^{m+1}(z) = \text{tp}_{\mathcal{T}}^{m+1}(z)\,.$$

Hence we are in the conditions (since $m \geq M$) to apply Lemma 4.6.5 on $x_C \prec \varphi(x_A) \prec x_A$ in $\mathcal{T}_1$ and get

$$\mathcal{T}_2 := \mathcal{P} \cdot \mathcal{C} \cdot \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{X} \cdot \mathcal{S}' \equiv_{\alpha}^{<\text{-inv FO}} \mathcal{T}_1\,.$$

Note that $\mathcal{T}' = \mathcal{T}_2$, which implies that $\mathcal{T}' \equiv_{\alpha}^{<\text{-inv FO}} \mathcal{T}$.
These operations are depicted in Figure 4.33.

- otherwise, we have $x_B \prec x_C$ and $x_B \prec \varphi(x_A)$ but x_C and $\varphi(x_A)$ are $\prec$-incomparable.
Let us decompose $\mathcal{C} \cdot \mathcal{S}$ as $\mathcal{C}'[\mathcal{S}, \mathcal{S}']$ (this notation extends in the natural way that of context, with two bottom anchors), where

$$\mathcal{S}' := \mathcal{S}_{\mathcal{T}}(\varphi(x_A))\,,$$

that is

$$\mathcal{T} = \mathcal{P} \cdot \mathcal{A} \cdot \mathcal{B} \cdot \mathcal{C}'[\mathcal{S}, \mathcal{S}']\,.$$

First, let

$$\mathcal{T}_1 = \mathcal{P} \cdot \mathcal{A} \cdot \mathcal{B} \cdot \mathcal{C}'[\mathcal{S}', \mathcal{S}]$$

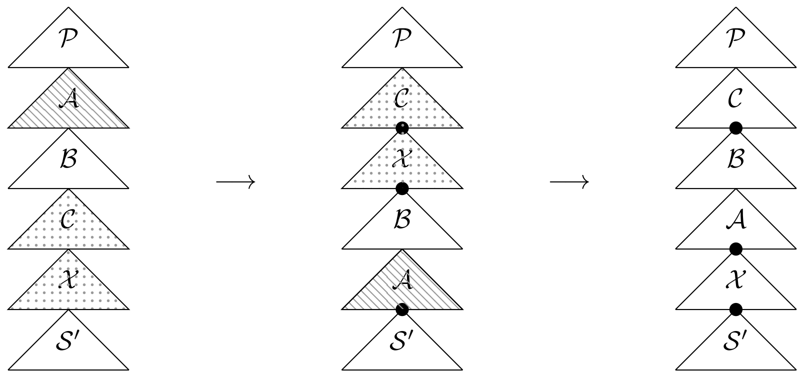

Figure 4.33: The case where $x_B \prec x_C \prec \varphi(x_A)$. The second operation swaps the segments between the dark nodes using Lemma 4.6.5.

be the m-guarded crossing-S-swap between x_C and $\varphi(x_A)$ in $\mathcal{T}$. Lemma 4.3.4 ensures (since $m \geq s(\alpha)$) that

$$\mathcal{T}_1 \equiv_\alpha^{<\text{-inv FO}} \mathcal{T} \text{ and } \forall z \in T, \mathrm{tp}_{\mathcal{T}_1}^{m+1}(z) = \mathrm{tp}_{\mathcal{T}}^{m+1}(z)\,.$$

The distance precondition in Lemma 4.3.4 holds because $\mathcal{T}$ is a hollow tree.

Let

$$\mathcal{T}_2 = \mathcal{P} \cdot \mathcal{C}'[\mathcal{S}', \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{S}]$$

be the m-guarded vertical-S-swap between $[x, x_A]$ and $[x_B, \varphi(x_A)]$ in $\mathcal{T}_1$. This swap falls under Case 1, hence ($m \geq n+1$) we get

$$\mathcal{T}_2 \equiv_\alpha^{<\text{-inv FO}} \mathcal{T}_1 \text{ and } \forall z \in T, \mathrm{tp}_{\mathcal{T}_2}^{m+1}(z) = \mathrm{tp}_{\mathcal{T}_1}^{m+1}(z)\,.$$

Now, let

$$\mathcal{T}_3 = \mathcal{P} \cdot \mathcal{C}'[\mathcal{B} \cdot \mathcal{A} \cdot \mathcal{S}, S']$$

be the m-guarded crossing-S-swap between x_A and x_C in $\mathcal{T}_2$. Lemma 4.3.4 ensures (since $m \geq s(\alpha)$) that $\mathcal{T}_3 \equiv_\alpha^{<\text{-inv FO}} \mathcal{T}_2$.

Note that

$$\mathcal{T}_3 = \mathcal{P} \cdot \mathcal{C} \cdot \mathcal{B} \cdot \mathcal{A} \cdot \mathcal{S}$$

is nothing but $\mathcal{T}'$. Hence $\mathcal{T}' \equiv_\alpha^{<\text{-inv FO}} \mathcal{T}$.

This process is illustrated in Figure 4.34.

Proposition 4.6.8. $\forall \alpha \in \mathbb{N}, \exists n_1 \in \mathbb{N}$ *such that for every* $\mathcal{P} \in \mathbb{H}_\sigma$ *and* $\forall \mathcal{Q} \in$ *quasi-*$\mathbb{H}_\sigma^{n_1}$*, if*

$$\mathcal{S}upp_{n_1}(\mathcal{P}) \simeq \mathcal{S}upp_{n_1}(\mathcal{Q})\,,$$

then

$$\mathcal{P} \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q}\,.$$

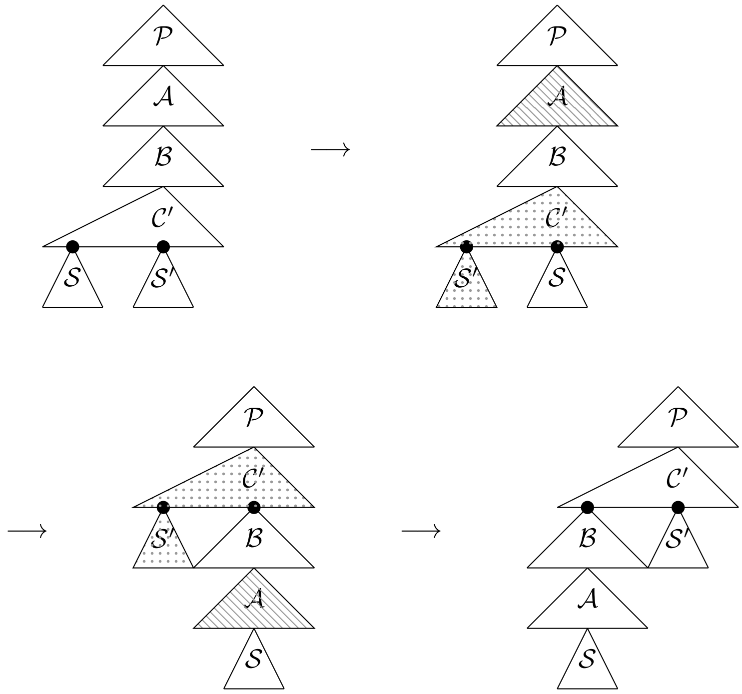

Figure 4.34: The case where x_C and $\varphi(x_A)$ are $\prec$-unrelated. The first operation is an m-guarded crossing-S-swap between $\varphi(x_A)$ and x_C. The second operation uses Case 1. The last operation is the dual of the first one.

Proof. Let n_0 be the maximum between the integers given by Lemma 4.6.4 and Lemma 4.6.5, and $s(\alpha)$.

Let n_1 be the integer given by Lemma 4.5.1 for n_0.

Because of the isomorphism between the n_1-enriched supports, there is a trivial n_1-pseudo-inclusion of $\mathcal{P}$ in $\mathcal{Q}$. Thus, Lemma 4.5.1 yields some $\mathcal{Q}' \in$ quasi-$\mathbb{H}_\sigma^{n_0+1}$ such that $\mathcal{Q}' \equiv_\alpha^{\text{<-inv FO}} \mathcal{Q}$, $\mathcal{S}\text{upp}_{n_0+1}(\mathcal{Q}') \simeq \mathcal{S}\text{upp}_{n_0+1}(\mathcal{Q})$ and some h' which $(n_0{+}1)$-pseudo-includes $\mathcal{P}$ in $\mathcal{Q}'$ and which respects S-siblings relation.

Now, $\mathcal{P}$ is a hollow tree, hence has the matching endpoints property, and h' must be surjective: this entails that $\mathcal{Q}'$ has the matching endpoints property.

For the remainder of this proof, we will need to apply vertical-S-swaps to $\mathcal{Q}'$ (and subsequent hollow quasitrees), even though it is not necessarily a hollow tree. However, the matching endpoints property ensures that the connected component $\mathcal{R}$ containing its root is a hollow tree.

We will only apply vertical-S-swaps in $\mathcal{R}$; when we talk of the vertical-S-swap in $\mathcal{Q}'$, we mean the disjoint union of the vertical-S-swap in $\mathcal{R}$ and of the other connected components of $\mathcal{Q}'$.

A **tree-prefix** of $\mathcal{P}$ or $\mathcal{Q}'$ is a substructure $\mathcal{T}$ which contains the root, is E-stable and such that if $S(x, y)$ and $y \in T$, then $x \in T$.

Let t be a thread with matching endpoints whose parent is y and an element x. We say that $x \prec t$ if $x \preceq y$. If u is a thread we write $t \prec u$ if $y \prec z$, where z is the S-parent of both of u's endpoints.

Let $\mathcal{T}_0, \cdots, \mathcal{T}_r$ be a sequence of tree-prefixes of $\mathcal{P}$ such that T_0 contains only the root of $\mathcal{P}$, $T_r = P$, and we go from T_i to T_{i+1} by adding a single thread.

We construct a sequence of structures $\mathcal{Q}' = \mathcal{Q}_0, \cdots, \mathcal{Q}_r$ with the following properties:

- $\mathcal{Q}_{i+1} \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q}_i$
- $\mathcal{S}\text{upp}_{n_0+1}(\mathcal{Q}_{i+1}) = \mathcal{S}\text{upp}_{n_0+1}(\mathcal{Q}_i)$
- $\mathcal{T}_i$ is **vertically-pseudo-included** in $\mathcal{Q}_i$, that is for every node x and thread t of $\mathcal{T}_i$, if x is the parent of t in $\mathcal{T}_i$ then $x \prec t$ in $\mathcal{Q}_i$. The smallest tree-prefix of $\mathcal{Q}_i$ containing all the threads of $\mathcal{T}_i$ is called the **$\mathcal{T}_i$-pseudo-tree**.

For $i = 0$, there is nothing to do: the root of $\mathcal{P}$ is vertically-pseudo-included in $\mathcal{Q}_0 = \mathcal{Q}'$.

From $\mathcal{T}_i$ to $\mathcal{T}_{i+1}$: we let t be the thread in $T_{i+1} \setminus T_i$ and let x be the parent of t.

- If t is in the $\mathcal{T}_i$-pseudo-tree of $\mathcal{Q}_i$, then there exists some element y and some thread u in $\mathcal{T}_i$ such that in $\mathcal{T}_i$, y is the parent of u, and in $\mathcal{Q}_i$, $y \prec t \prec u$. In $\mathcal{Q}_i$, let's call y' the parent of u, u' the thread whose parent is y, x' the parent of t and t' the thread whose parent is x.

  There are two cases to consider.

  - If $y \prec x$ in $\mathcal{P}$ (c.f. Figure 4.35). Then in $\mathcal{Q}_i$, we must have $y' \prec x$, and we can apply Lemma 4.6.4. Let $\mathcal{Q}_{i+1}$ be the n_0-guarded vertical-S-swap between $[y, x']$ and $[y', x]$ in $\mathcal{Q}_i$. Note that in the limit case where y is the parent of t (that is, $x' = y$), we apply Lemma 4.6.5 instead of Lemma 4.6.4.

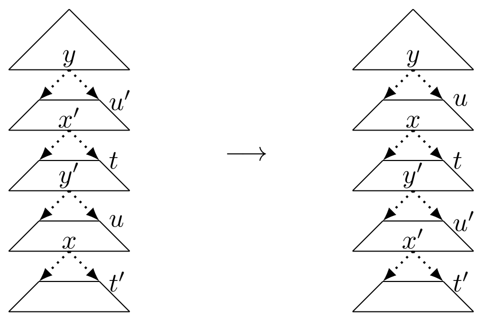

Figure 4.35: We re-associate y to u and x to t with a vertical-S-swap.

  - Otherwise, x and y must be $\prec$-unrelated in $\mathcal{Q}_i$ (c.f. Figure 4.36). Note that because we work in a hollow tree, the conditions to apply Lemma 4.3.4 are met. Set $\mathcal{Q}'$ to be the n_0-guarded crossing-S-swap between x and x' in $\mathcal{Q}_i$. Then, set $\mathcal{Q}_{i+1}$ to be the n_0-guarded crossing-S-swap between y and y' in $\mathcal{Q}'$.

- Otherwise, if $x \prec t$ in $\mathcal{Q}_i$, we set $\mathcal{Q}_{i+1} := \mathcal{Q}_i$.

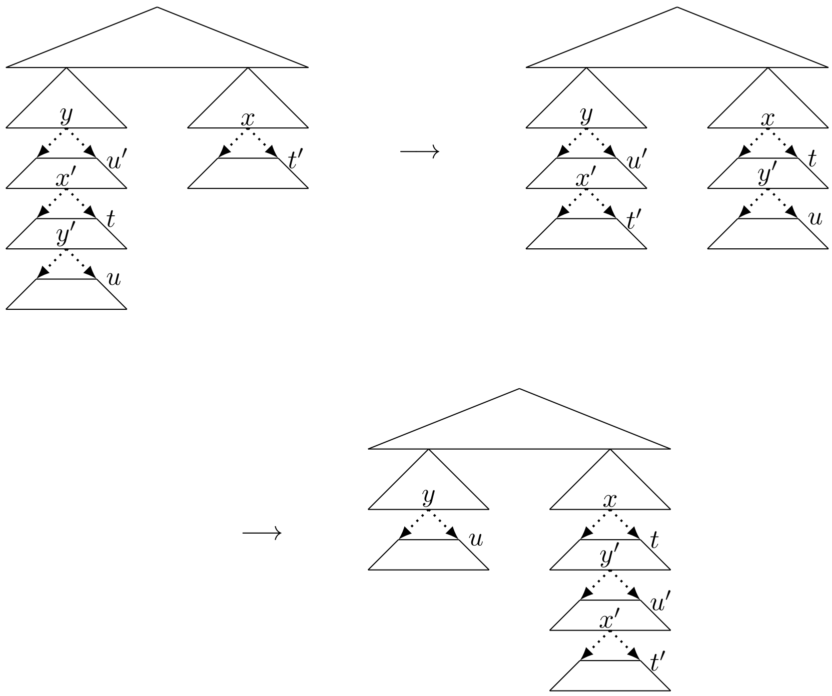

Figure 4.36: We re-associate y to u and x to t with two crossing-S-swaps.

- Otherwise, x and x' are $\prec$-unrelated in $\mathcal{Q}_i$. Once again, we are in the right setting to apply Lemma 4.3.4 because we work in a hollow tree. Let $\mathcal{Q}_{i+1}$ be the n_0-guarded crossing-S-swap between x and x' in $\mathcal{Q}_i$.

In the end, we have vertically-pseudo-included $\mathcal{P}$ into $\mathcal{Q}_r$. Since they have the same support, the vertical-pseudo-inclusion is an isomorphism. Hence,

$$\mathcal{P} \equiv_{\alpha}^{<\text{-inv FO}} \mathcal{Q}\,.$$

□

4.7 Proof of the main result

Let's recall the main result of this chapter:

Theorem 4.2.1. *For every alphabet σ,*

$$<\text{-inv FO} = \text{FO} \textit{ on } \mathbb{H}_\sigma\,.$$

We now have the tools to prove Theorem 4.2.1.

Let $\alpha \in \mathbb{N}$. Recall that we want to find $f(\alpha)$ such that $\forall \mathcal{P}, \mathcal{Q} \in \mathbb{H}_\sigma$, if $\mathcal{P} \equiv_{f(\alpha)}^{\text{FO}} \mathcal{Q}$ then $\mathcal{P} \equiv_{\alpha}^{<\text{-inv FO}} \mathcal{Q}$. Proposition 2.3.2 allows us to conclude from there.

We set, in that order:

- l as in Proposition 4.6.3 (loop elimination), that is such that FO[α] cannot distinguish the linear order on $\{1, \cdots, l\}$ from the linear order on $\{1, \cdots, l+1\}$
- n_1 as in Proposition 4.6.8 (S-parents re-association)
- n_2 as in Proposition 4.4.2 (pseudo-inclusion) for $n_1 - 1$
- n_3 as in Proposition 4.6.3 (loop elimination) for n_2
- M as in Proposition 4.6.1 (loop minimization) for $n_3 + 1$
- K as in Proposition 4.6.3 (loop minimization) for n_2 and M
- n_4, d_1, D as in Proposition 4.5.4 (S-stabilization of the image of a pseudo-inclusion) for n_3
- n_5 as in Proposition 4.4.2 (pseudo-inclusion) for n_4
- $f(\alpha)$ as in Proposition 4.3.11 (pumping) for n_5 and $d := \max(d_1, K)$.

Starting from $\mathcal{P} \equiv^{\text{FO}}_{f(\alpha)} \mathcal{Q}$, we unfold the previously set indexes to apply the corresponding propositions in the reverse order: we transform $\mathcal{Q}$ into $\mathcal{P}$ along a sequence of $\equiv^{<\text{-inv FO}}_{\alpha}$ hollow quasitrees (with smaller and smaller radius) $\mathcal{Q}_i$ as follows.

According to Proposition 4.3.11, we can pump inside $\mathcal{Q}$ to get

$$\mathcal{Q}_0 \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{Q}$$

such that

$$[\![\mathcal{E}_{n_5+1}(\mathcal{P})]\!] \leq^D_d [\![\mathcal{E}_{n_5+1}(\mathcal{Q}_0)]\!] \,. \tag{4.2}$$

Now that we've made sure there were at least as many occurrences of every neighborhood type in $\mathcal{Q}_0$ as in $\mathcal{P}$, Proposition 4.4.2 yields

$$\mathcal{Q}_1 \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{Q}_0$$

such that

$$[\![\mathcal{E}_{n_4+1}(\mathcal{Q}_1)]\!] = [\![\mathcal{E}_{n_4+1}(\mathcal{Q}_0)]\!]$$

and such that $\mathcal{P}$ is $(n_4 + 1)$-pseudo-included in $\mathcal{Q}_1$ by some h.

Since

$$[\![\mathcal{E}_{n_4+1}(\mathcal{P})]\!] \leq^D_{d_1} [\![\mathcal{E}_{n_4+1}(\mathcal{Q}_1)]\!]$$

by (4.2), Proposition 4.5.4 gives

$$\mathcal{Q}_2 \equiv^{<\text{-inv FO}}_{\alpha} \mathcal{Q}_1$$

such that

$$\mathcal{S}\text{upp}_{n_3+1}(\mathcal{Q}_2) \simeq \mathcal{S}\text{upp}_{n_3+1}(\mathcal{Q}_1)$$

as well as some reduced h' which $(n_3 + 1)$-pseudo-includes $\mathcal{P}$ in $\mathcal{Q}_2$, where $V := Q_2 \setminus \text{Im}(h')$ is S-stable in $\mathcal{Q}_2$.

Proposition 4.6.1 gives us some $\mathcal{Q}_3$ and $\mathcal{U}_+ \in \mathbb{L}_\sigma^{n_3+1}$ such that

$$\mathcal{Q}_3 \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q}_2, \qquad |U| \leq M$$

and

$$[\![\mathcal{E}_{n_3+1}(\mathcal{Q}_3)]\!] = [\![\mathcal{E}_{n_3+1}(\mathcal{P})]\!] + [\![\mathcal{E}_{n_3+1}(\mathcal{U})]\!] .$$

Since $|U| \leq M$ and for every (n_3+1)-neighborhood type ξ that occurs in $\mathcal{U}$, $|\mathcal{Q}_3|_\xi \geq K$ (indeed: since ξ occurs in $\mathcal{U}$, $|\mathcal{Q}_0|_\xi \neq |\mathcal{P}|_\xi$ and $|\mathcal{Q}_3|_\xi > |\mathcal{P}|_\xi > d \geq K$), we can remove the extra elements by applying Proposition 4.6.3, which gives

$$\mathcal{Q}_4 \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q}_3$$

such that

$$[\![\mathcal{E}_{n_2+1}(\mathcal{Q}_4)]\!] = [\![\mathcal{E}_{n_2+1}(\mathcal{Q}_3)]\!] - [\![\mathcal{E}_{n_2+1}(\mathcal{U})]\!] = [\![\mathcal{E}_{n_2+1}(\mathcal{P})]\!] .$$

Since $\mathcal{P}$ and $\mathcal{Q}_4$ have the same number of occurrences of every (n_2+1)-neighborhood type, we can pseudo-include one into another according to Proposition 4.4.2, which yields

$$\mathcal{Q}_5 \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q}_4$$

such that

$$\mathcal{S}\text{upp}_{n_1}(\mathcal{Q}_5) \simeq \mathcal{S}\text{upp}_{n_1}(\mathcal{P}) .$$

Indeed, the pseudo-inclusion cannot have any jumping pair.

Finally, Proposition 4.6.8 allows us to conclude that $\mathcal{Q}_5 \equiv_\alpha^{<\text{-inv FO}} \mathcal{P}$, completing our sequence of transformations.

This concludes the proof that $\mathcal{P} \equiv_\alpha^{<\text{-inv FO}} \mathcal{Q}$, which proves Theorem 4.2.1.

Note 4.7.1. *Recall the discussion from Section 2.6.3: not only did we prove that for every alphabet σ,*

$$<\text{-inv FO} = \text{FO} \;\; \textit{on} \; \mathbb{H}_\sigma ,$$

but also that for every σ and every $k \in \mathbb{N}$,

$$<\text{-inv}/_{quasi\text{-}\mathbb{H}_\sigma^k} \; \text{FO} = \text{FO} \;\; \textit{on} \; \mathbb{H}_\sigma .$$

Indeed, as long as $\alpha > k$, all the structures involved in the proof belong to quasi-$\mathbb{H}_\sigma^k$. We not know whether this collapse still hold when k goes to infinity, namely whether

$$<\text{-inv}/_{\mathbb{H}_\sigma} \; \text{FO} = \text{FO} \;\; \textit{on} \; \mathbb{H}_\sigma .$$

It is not clear to us what kind of sentence could be order-invariant on hollow trees but not on hollow quasitrees with arbitrarily large radius.

In our proof, we had to break the structure of hollow trees in the chain of intermediate structures. It is an interesting question whether this can be avoided by finding an alternative strategy staying within the realm of hollow trees.

4.8 Conclusion

We have shown in this chapter that $<$-inv FO = FO on the class of hollow trees. Now recall from Section 4.1.2 that one of the motivations for the study of hollow trees comes from pathwidth 2: in order to lift this collapse to tracks of bounded degree, it suffices to show that $<$-inv FO = FO over structures that have the same underlying graph as hollow trees, but without the possibility to distinguish a sibling from a child. In other words, structures obtained from hollow trees by subsuming the union of E and S in a single binary relation.

Unfortunately our proof does not extend to this class of structures, since being able to tell both types of edges apart is crucial in our proof in order to distinguish between E-swaps and S-swaps. We leave this generalization as an open problem.

We also have no idea on the path to take when the degree is not assumed to be bounded. Indeed, in that setting, there is no hope to transform a track into a tree-like structure by the mean of a bi-FO-interpretation.

Chapter 5

Conclusion

In this thesis, we improved the understanding of the expressive power of invariant logics by broadening the settings in which order and successor-invariant logics are known to collapse to FO.

Our main contributions are the proofs that Succ-inv FO collapses to FO when the degree is bounded, as well as on any class of graphs which is FO-interpretable in a class of graphs of bounded degree (i.e. near-uniform classes of graphs), and that $<$-inv FO collapses to FO on hollow trees.

Besides those results, we believe that the techniques developed throughout this thesis can be useful in future endeavors. Similarly, we hope that newly introduced tools such as fractal types and mirror swaps could prove useful outside of the framework of this thesis.

A natural follow-up to the results presented here would be to extend the setting, and to show that $<$-inv FO or Succ-inv FO collapses to FO on broader classes of structures; for instance on structures of bounded treewidth, or pathwidth.

When the degree is not supposed to be bounded, proving such a collapse seems particularly difficult, and we have no idea about a strategy one could implement to reach that result. It is not even clear to us that $<$-inv FO = FO in that setting.

However, on classes of structures of bounded treewidth and of bounded degree, it seems reasonnable to conjecture that $<$-inv FO collapses to FO. On classes of pathwidth 2 and of bounded degree for instance, which motivated the definition of hollow tree, it may be possible, though technically difficult, to extend the techniques developed in Chapter 4.

In Sections 2.5.5 and Sections 2.5.6, we showed that there was no hope to interpret in FO path-decompositions (resp. tree-decompositions) of bounded diameter in structures of bounded pathwdith (resp. bounded treewidth) and bounded degree. However, these counter-examples do not exclude the possibility to interpret in FO tree-decompositions of bounded diameter in structures of bounded pathwidth and bounded degree. Although it seems unlikely to us that such FO-interpretations exist in the light of our counter-examples, their existence would entail a lift of the collapse of $<$-inv FO to FO from trees to classes of bounded pathwidth and bounded degree.

It is also conceivable, given the trouble one already goes through when trying to prove that Succ-inv FO = FO on classes of pathwidth 2, that the notions of treewidth or pathwidth are not the right measures of sparsity when dealing with invariant logics. One could then try to find a measure that is more adapted to such a context. We have seen that the degree boundedness hypothesis is a lever that can successfully be used in the context of a successor relation; relaxing the hypothesis on the degree in a careful way may allows to extend the collapse of Succ-inv FO to FO to broader sparse classes of structures.

Bibliography

[1] Matthew Anderson, Dieter van Melkebeek, Nicole Schweikardt, and Luc Segoufin. Locality from circuit lower bounds. *SIAM J. Comput.*, 2012.

[2] János Barát, Péter Hajnal, Yixun Lin, and Aifeng Yang. On the structure of graphs with path-width at most two. *Studia Scientiarum Mathematicarum Hungarica*, 2012.

[3] Pablo Barceló and Leonid Libkin. Order-invariant types and their applications. *Log. Methods Comput. Sci.*, 2016.

[4] Michael Benedikt and Luc Segoufin. Regular tree languages definable in FO and in FO_{mod}. *ACM Trans. Comput. Log.*, 2009.

[5] Michael Benedikt and Luc Segoufin. Towards a characterization of order-invariant queries over tame graphs. *J. Symb. Log.*, 2009.

[6] Evert W Beth. On padoa's method in the theory of definition. 1956.

[7] Hans L. Bodlaender. A note on domino treewidth. *Discrete Mathematics & Theoretical Computer Science*, 1999.

[8] Hans L. Bodlaender and Joost Engelfriet. Domino treewidth. *J. Algorithms*, 1997.

[9] J Richard Büchi. Weak second-order arithmetic and finite automata. *Mathematical Logic Quarterly*, 1960.

[10] E. F. Codd. *The Relational Model for Database Management, Version 2.* Addison-Wesley, 1990.

[11] Bruno Courcelle. Graph rewriting: An algebraic and logic approach. Elsevier and MIT Press, 1990.

[12] William Craig. Linear reasoning. A new form of the herbrand-gentzen theorem. *J. Symb. Log.*, 1957.

[13] Guoli Ding and Bogdan Oporowski. Some results on tree decomposition of graphs. *Journal of Graph Theory*, 1995.

[14] Kord Eickmeyer, Michael Elberfeld, and Frederik Harwath. Expressivity and succinctness of order-invariant logics on depth-bounded structures. In *Mathematical Foundations of Computer Science (MFCS)*, 2014.

[15] Kord Eickmeyer, Ken-ichi Kawarabayashi, and Stephan Kreutzer. Model checking for successor-invariant first-order logic on minor-closed graph classes. In *Logic in Computer Science, LICS*, 2013.

[16] Michael Elberfeld, Marlin Frickenschmidt, and Martin Grohe. Order invariance on decomposable structures. In *Logic in Computer Science, LICS*, 2016.

[17] Viktor Engelmann, Stephan Kreutzer, and Sebastian Siebertz. First-order and monadic second-order model-checking on ordered structures. In *Logic in Computer Science, LICS*, 2012.

[18] Ronald Fagin, Larry J. Stockmeyer, and Moshe Y. Vardi. On monadic NP vs. monadic co-np. *Inf. Comput.*, 1995.

[19] Jakub Gajarský, Petr Hlinený, Daniel Lokshtanov, Jan Obdrzálek, and M. S. Ramanujan. A new perspective on FO model checking of dense graph classes. *CoRR*, 2018.

[20] Tobias Ganzow and Sasha Rubin. Order-invariant MSO is stronger than counting MSO in the finite. *CoRR*, 2007.

[21] Julien Grange. Successor-invariant first-order logic on classes of bounded degree. In *LICS, Logic in Computer Science*, 2020.

[22] Julien Grange and Luc Segoufin. Order-Invariant First-Order Logic over Hollow Trees. In *Computer Science Logic, CSL*, 2020.

[23] Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. In *Symposium on Theory of Computing, STOC*. ACM, 2014.

[24] Martin Grohe and Thomas Schwentick. Locality of order-invariant first-order formulas. *ACM Trans. Comput. Log.*, 2000.

[25] Lauri Hella, Leonid Libkin, and Juha Nurmonen. Notions of locality and their logical characterizations over finite models. *J. Symb. Log.*, 1999.

[26] Neil Immerman. Relational queries computable in polynomial time. *Information and Control*, 1986.

[27] Leonid Libkin. *Elements of Finite Model Theory.* Texts in Theoretical Computer Science. An EATCS Series. Springer, 2004.

[28] Roger C Lyndon, Marcel-Paul Schützenberger, et al. The equation $a^M = b^N c^P$ in a free group. *The Michigan Mathematical Journal*, 1962.

[29] Hannu Niemistö. *Locality and order-invariant logics.* PhD thesis, University of Helsinki, 2007.

[30] Martin Otto. Epsilon-logic is more expressive than first-order logic over finite structures. *J. Symb. Log.*, 2000.

[31] Andreas Potthoff. *Logische klassifizierung regulärer baumsprachen.* Universität Kiel. Institut für Informatik und Praktische Mathematik, 1994.

[32] Neil Robertson and Paul D. Seymour. Graph minors. II. algorithmic aspects of tree-width. *J. Algorithms*, 1986.

[33] Benjamin Rossman. Successor-invariant first-order logic on finite structures. *J. Symb. Log.*, 2007.

[34] Detlef Seese. Linear time computable problems and first-order descriptions. *Mathematical Structures in Computer Science*, 1996.

[35] Larry Joseph Stockmeyer. *The complexity of decision problems in automata theory and logic*. PhD thesis, Massachusetts Institute of Technology, 1974.

[36] James W. Thatcher and Jesse B. Wright. Generalized finite automata theory with an application to a decision problem of second-order logic. *Mathematical systems theory*, 1968.

[37] Boris A Trakhtenbrot. Impossibility of an algorithm for the decision problem in finite classes. *Doklady Akademii Nauk SSSR*, 1950.

[38] Jan van den Heuvel, Stephan Kreutzer, Michal Pilipczuk, Daniel A. Quiroz, Roman Rabinovich, and Sebastian Siebertz. Model-checking for successor-invariant first-order formulas on graph classes of bounded expansion. *CoRR*, 2017.

[39] Moshe Y. Vardi. The complexity of relational query languages. In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing*, 1982.

[40] Thomas Zeume and Frederik Harwath. Order-invariance of two-variable logic is decidable. *CoRR*, 2016.

RÉSUMÉ

Cette thèse s'attache à l'étude du pouvoir d'expression de deux logiques définies par invariance : successor-invariant first order logic, Succ-inv FO, et order-invariant first order logic, <-inv FO. Ces formalismes étendent la logique du premier ordre, FO, en autorisant l'accès à une relation de successeur (pour Succ-inv FO) ou à une relation d'ordre (pour <-inv FO) sur le domaine des structures considérées, à la condition que l'évaluation des formules ne dépende pas du choix d'une telle relation.

Il est établi que dans le cadre général, Succ-inv FO et <-inv FO sont plus expressives que la simple logique du premier ordre. Cependant, si l'on se restreint au cas des arbres, ces deux logiques ont le même pouvoir d'expression que FO. Les deux résultats centraux de cette thèse étendent les classes de structures sur lesquelles le pouvoir d'expression de ces logiques définies par invariance est réduit à celui de FO.

Tout d'abord, on montrera que Succ-inv FO n'est pas plus expressif que FO sur les classes de structures dont le degré est borné. La preuve de ce résultat repose sur la construction de successeurs préservant la similarité des structures au premier ordre.

Dans un second temps, on définira une nouvelle classe de structures : celle des hollow trees. Un hollow tree est essentiellement un arbre de rang non borné dans lequel un parent est uniquement relié à son enfant le plus à gauche et à celui le plus à droite. Les nœuds d'une fratrie sont liés par le biais d'une relation binaire symétrique. La notion de hollow tree est une généralisation de celle d'arbre de rang borné, et se présente comme une première étape vers les classes de structures de largeur de chemin au plus 2.

On montrera que sur la classe des hollow trees, <-inv FO n'est pas plus expressive que la logique du premier ordre.

MOTS CLÉS

Théorie des modèles finis – Invariance – Logique du premier ordre

ABSTRACT

This thesis focuses on the expressive power of two invariant logics: successor-invariant first-order logic, Succ-inv FO, and order-invariant first-order logic, <-inv FO. In these formalisms, on top of plain first-order logic, FO, an access to an additional successor relation (for Succ-inv FO) or linear order relation (for <-inv FO) on structures is granted, provided that the evaluation of sentences does not depend on the choice of a particular relation.

It is well known that both Succ-inv FO and <-inv FO are more expressive than FO in general. However, if one considers only trees, these logics are no more expressive than plain FO. The two main results of this thesis extend the classes of structures on which these invariant logics collapse to FO.

First, we prove that Succ-inv FO is no more expressive than FO on classes of bounded degree. For that, we show how successor relations preserving FO-similarity can be constructed.

Second, we define a new class of structures, that of hollow trees. A hollow tree can be seen as an unranked tree, where a parent is only linked to its leftmost and rightmost children. Elements of a siblinghood are linearly related through another binary relation, which is symmetric. The notion of hollow trees is a generalization of ranked trees, and we believe it to be a gateway to structures of pathwidth 2.

We show that <-inv FO collapses to FO on the class of hollow trees.

KEYWORDS

Finite model theory – Order-invariance – Successor-invariance – First-order logic