



HAL
open science

The Treewidth Boundedness Problem for an Inductive Separation Logic of Relations

Marius Bozga, Lucas Bueri, Radu Iosif, Florian Zuleger

► **To cite this version:**

Marius Bozga, Lucas Bueri, Radu Iosif, Florian Zuleger. The Treewidth Boundedness Problem for an Inductive Separation Logic of Relations. 2023. hal-04264291v2

HAL Id: hal-04264291

<https://hal.science/hal-04264291v2>

Preprint submitted on 23 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

THE TREEWIDTH BOUNDEDNESS PROBLEM FOR AN INDUCTIVE SEPARATION LOGIC OF RELATIONS

MARIUS BOZGA ^a, LUCAS BUERI ^a, RADU IOSIF ^a, AND FLORIAN ZULEGER ^b

^a Univ. Grenoble Alpes, CNRS, Grenoble INP, VERIMAG, 38000, France

^b Institute of Logic and Computation, Technische Universität Wien, Austria

ABSTRACT. The treewidth boundedness problem for a logic asks for the existence of an upper bound on the treewidth of the models of a given formula in that logic. This problem is found to be undecidable for first order logic. We consider a generalization of Separation Logic over relational signatures, interpreted over standard relational structures, and describe an algorithm that decides the treewidth boundedness problem for this logic. Furthermore, our algorithm can give an estimate of the bound of the models of a given formula, in case there is a finite such bound.

1. INTRODUCTION

The treewidth of a graph is a positive integer measuring, informally speaking, how far a graph is from a tree. For instance, trees have treewidth one, series-parallel graphs (i.e., circuits with one input and one output that can be either cascaded or overlaid) have treewidth two, whereas $k \times k$ square grids have treewidth k , for any $k \geq 1$. The treewidth parameter is a cornerstone of algorithmic tractability. For instance, many NP-complete graph problems such as Hamiltonicity and 3-Colorability become polynomial-time, when restricted to inputs whose treewidth is bounded by a constant (see, e.g., [FG06, Chapter 11] for a survey of classical treewidth-parameterized tractable problems).

Structures are interpretations of relation symbols that define the standard semantics of first and second order logic [vD94]. They provide a unifying framework for reasoning about a multitude of graph types e.g., graphs with multiple edges, labeled graphs, colored graphs, hypergraphs, etc. The notion of treewidth is straightforwardly generalized from graphs to structures. In this context, bounding the treewidth by a constant sets the frontier between the decidability and undecidability of monadic second order (MSO) logical theories. A result of Courcelle [Cou90] proves that MSO is decidable over bounded treewidth structures, by reduction to the emptiness problem of tree automata. A dual result of Seese [See91] proves that each class of structures with a decidable MSO theory necessarily has bounded treewidth. Since MSO is the yardstick of graph specification logics [CE12], these results show that *treewidth bounded* classes of structures are tantamount to the existence of decision procedures for important classes of properties, in those areas of computing where graphs are relevant such as, e.g., static analysis [JM82], databases [ABS00] and concurrency [DNM08].

Key words and phrases: Model Theory, Treewidth, Separation Logic.

This paper considers the *treewidth boundedness problem*, which asks for the existence of a bound on the treewidths of the models of a formula given in input. We show that for first-order logic (and implicitly MSO) the problem is already undecidable. This negative result for classical logics motivates our focus on *substructural logics* that have, in addition to boolean conjunction, a conjunction-like connective, for which Gentzen’s natural deduction rules of weakening and contraction do not hold. We prove the decidability of this problem for a generalization of Separation Logic to relational signatures, interpreted over structures.

Separation Logic (SL) [IO01, Rey02, CGG02] is a first order substructural logic with a *separating conjunction* $*$ that decomposes structures. For reasons related to its applications to the deductive verification of pointer-manipulating programs, the models of SL are finite partial functions, called *heaps*. In SL, the separating conjunction stands for the union of heaps with disjoint domains.

When combined with *inductive definitions* [Acz77], SL gives concise descriptions of the recursive data structures (singly- and doubly-linked lists, trees, etc.) used in imperative programming (e.g., C, C++, Java, etc.). The shape of these structures can be described using only existentially quantified separating conjunctions of (dis-)equalities and points-to atoms. This subset of SL is referred to as the *symbolic heap* fragment.

SL is a powerful tool for reasoning about low-level pointer updates. It allows to describe actions *locally*, i.e., only with respect to the resources (e.g., memory cells, network nodes) involved, while framing out the part of the state that is irrelevant for the action. This principle of describing mutations, known as *local reasoning* [COY07], is at the heart of scalable compositional proof techniques for pointer programs [CDOY11, CDNQ12, DPJ08, BCO06].

The *Separation Logic of Relations* (SLR) is the generalization of SL to relational signatures, interpreted over structures. This logic has been first considered for relational databases and object-oriented languages [KR04]. Here the separating conjunction splits the interpretation of each relation symbol from the signature into disjoint parts. For instance, the formula $r(x_1, \dots, x_n)$ describes a structure in which all relations are empty and r consists of a single tuple of values x_1, \dots, x_n , whereas $r(x_1, \dots, x_n) * r(y_1, \dots, y_n)$ says that r consists of two distinct tuples, i.e., the values of x_i and y_i differ for at least one index $1 \leq i \leq n$. Moreover, when encoding graphs by structures, SLR allows to specify edges that have no connected vertices, isolated vertices, or both. The same style of composition is found in other spatial logics interpreted over graphs, such as the GL logic of Cardelli et al [CGG02].

Our motivation for studying the models of SLR arose from recent work on deductive verification of self-adapting distributed systems, where Hoare-style local reasoning is applied to write correctness proofs for systems with dynamically reconfigurable network architectures [ABIK22, BBI22a, BBI22b]. The assertion language of these proofs is SLR, with unary relation symbols used to model nodes (processes) of the network and relation symbols of arity two or more used to model links (communication channels) between nodes. Just as user-defined inductive predicates are used in SL to describe data structures (lists, trees, etc.), SLR inductive predicates are used to describe common architectural styles (e.g., pipelines, rings, stars, etc.) that ensure correct and optimal behavior of many distributed applications.

The decidability result from this paper defines the class of SLR formulæ whose models are treewidth bounded and provides a reasonable estimate on the bound, in case one exists. On one hand, this algorithm answers the question *does the set of structures defined by a given system of inductive definitions have a decidable MSO theory?* If this is the case, problems such as, e.g., Hamiltonicity, k -Colorability, Planarity, etc. are decidable on this set of structures. Another application is the decidability of the *entailment problem* $\llbracket \phi \rrbracket_{\Delta} \subseteq \llbracket \psi \rrbracket_{\Delta}$

asking if each model of a formula ϕ is also a model of another formula ψ , when the predicate symbols in ϕ and ψ are interpreted by a set of inductive definitions Δ . In principle, the decidability of this problem depends on (i) ϕ having only treewidth bounded models, for a computable upper bound, and (ii) both ϕ and ψ being MSO-definable [IRS13]. The algorithm described in the paper provides a key ingredient for defining fragments of SLR with a decidable entailment problem, which is tantamount to automating proof generation in Hoare logic [BBI22a, BBI22b].

1.1. Related work. One of the first fragments of SL with a decidable entailment problem relied on an ad-hoc translation into equivalent MSO formulæ, together with a static guarantee of treewidth boundedness, called *establishment* [IRS13]. More recently, the entailment problem in this fragment of SL has been the focus of an impressive body of work [CHO⁺11, KZ20, EIP21a, EIP21b, LL23]. In particular, the *establishment problem* “is a given set of inductive definitions established” has been found to be co-NP-complete [JKM⁺17]. Moreover, lifting the establishment condition leads to the undecidability of entailments, as showed in [EIP22]. The establishment problem can, in fact, be seen as the precursor of the treewidth boundedness problem studied in the present paper.

The treewidth parameter showed also in a recent comparison between the expressivity of SLR with inductive definitions and that of MSO [IZ23]. When restricting the interpretation of the logics to treewidth bounded graphs, SLR strictly subsumes MSO, i.e., for each MSO formula ϕ and integer $k \geq 1$, there exists a formula ψ of SLR that defines the models of ϕ of treewidth at most k . Moreover, the logics are incomparable for classes of graphs of unbounded treewidth.

1.2. Motivating examples. We introduce the reader to SLR and the treewidth boundedness problem by means of examples. Figure 1 (a) shows a chain, defined by an unfolding of the inductive predicate $A(x_1, x_2)$. The chain starts at x_1 and ends at x_2 . The elements of the chain are labeled by a unary relation symbol a and the neighbours are linked by a binary relation r . Each unfolding of the inductive definition $A(x_1, x_2) \leftarrow \exists y . a(x_1) * r(x_1, y) * A(y, x_2)$ instantiates the existential quantifier to an element distinct from the existing ones. This is because every instantiation of an existential quantifier is placed into a set labeled by a and the semantics of the separating conjunction requires that these sets must be disjoint in each decomposition of a model of $a(x_1) * r(x_1, y) * A(y, x_2)$ into models of $a(x_1) * r(x_1, y)$ and $A(y, x_2)$. Then, each model of $\exists x_1 \exists x_2 . A(x_1, x_2)$ is a possibly cyclic chain, because nothing is enforced on x_2 , which can be mapped back to a previous instantiation of y . Hence each model of this sentence has treewidth two at most.

Figure 1 (b) shows a family of models for a slightly modified definition of the chain from Figure 1 (a), given by the recursive rule $A(x_1, x_2) \leftarrow \exists y . r(x_1, y) * A(y, x_2)$, where the instantiations of the existential quantifiers are not placed into any particular set. In this case, one can fold a sufficiently large chain onto itself and creating a square grid, by using the same element of the structure more than once to instantiate a quantifier. Then, the sentence $\exists x_1 \exists x_2 . A(x_1, x_2)$ has an infinite set of models containing larger and larger square grid minors, thus having unbounded treewidth.

Since placing every quantifier instance into the same set guarantees treewidth boundedness, as in, e.g., Figure 1 (a), a natural question that arises is what happens when these instances are placed into two (not necessarily disjoint) sets? The inductive definition of the

predicate A in Figure 1 (c) creates an unbounded number of disconnected r -edges whose endpoints are arbitrarily labeled with a and b , respectively. In this case, one can instantiate a a -labeled (resp. b -labeled) variable with a new element or a previous b (resp. a) element and build chains (or sets of disconnected chains), of treewidth two at most two. Again, this is because a simple cycle with more than two elements has treewidth two.

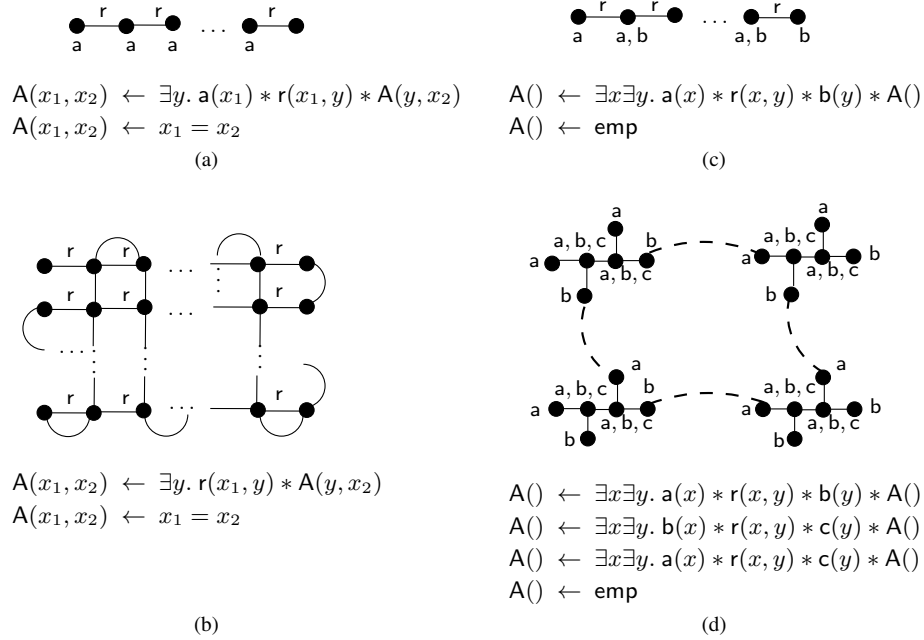


Figure 1: Examples of bounded and unbounded treewidth models

Let us now consider three unary relation symbols a , b and c and three types of disconnected r -edges (according to the labels of their endpoints) created by three recursive definitions of Figure 1 (d), namely a - b , b - c and a - c edges. In this case, the sentence $A()$, where A is a predicate symbol of zero arity, has models with unboundedly large square grid minors, obtained by “glueing” these edges (i.e., instantiating several quantifiers with the same element from different sets). The glued pairs are connected with dotted lines in Figure 1 (d). Hence, these structures form a set of unbounded treewidth.

These examples highlight the main ideas behind an algorithm that decides the existence of a bound on the treewidths of the models of a given formula, with predicates interpreted by set of inductive definitions. First, one needs to identify the definitions that can iterate any number of times producing building blocks of unboundedly large grids (modulo edge contractions). Second, these structures must connect elements from different sets, e.g., a , b or c in Figure 1. A complication is that these sets can be defined not only by monadic relation symbols, but also by n -ary relation atoms where all but one variable have the same values for any occurrence. For instance, the variable x_2 in Figure 1 (a) has the same value in an arbitrarily long unfolding of $A(x_1, x_2)$ and we could have written $r(x_1, x_2)$ instead of $a(x_1)$ in the first rule, with the same effect, while avoid using ‘ a ’ altogether. Last, the interplay between the connectivity and labeling of the building blocks is important. For instance, in Figure 1 (d), the building blocks of the grid are structures consisting of six elements,

that connect three ‘a’ with three ‘b’ elements. In fact, as we shall prove, a necessary and sufficient condition for treewidth-boundedness is that any two such “iterable” substructures that connect at least three elements must also place these elements in at least one common set (e.g., a, b or c in our example).

2. THE SEPARATION LOGIC OF RELATIONS

This section defines formally the Separation Logic of Relations (SLR) and its corresponding treewidth boundedness problem. It also introduces most of the technical notions used throughout the paper.

Let \mathbb{N} be the set of positive integers, zero included and $\mathbb{N}_+ \stackrel{\text{def}}{=} \mathbb{N} \setminus \{0\}$. Given integers i and j , we write $[i..j]$ for the set $\{i, i+1, \dots, j\}$, assumed to be empty if $i > j$. For a set A , we denote by $\text{pow}(A)$ its powerset. The cardinality of a finite set A is $\text{card}(A)$. By writing $S = S_1 \uplus S_2$, we mean that S_1 and S_2 partition S , i.e., that $S = S_1 \cup S_2$ and $S_1 \cap S_2 = \emptyset$.

Multisets are denoted as $\{\!\{a, b, \dots\}\!\}$ and all set operations (union, intersection, etc.) are used with multisets as well. In particular, a binary operation involving a set and a multiset implicitly lifts the set to a multiset and returns a multiset. The multi-powerset (i.e., the set of multisets) of A is denoted as $\text{mpow}(A)$.

For a binary relation $R \subseteq A \times A$, we denote by R^* its reflexive and transitive closure and by R^\equiv the smallest equivalence relation that contains R , i.e., the closure of R^* by symmetry. For a set $S \subseteq A$, we denote by $R|_S$ the relation obtained by removing from R all pairs with an element not in S . A binary relation $R \subseteq A \times B$ is an *A-B matching* iff $\{a, b\} \cap \{a', b'\} = \emptyset$, for all distinct pairs $(a, b), (a', b') \in R$.

2.1. Structures. Let \mathbb{R} be a finite and fixed set of *relation symbols*, of arities $\#r \geq 1$, for all $r \in \mathbb{R}$. A relation symbol of arity one (resp. two) is called *unary* (resp. *binary*).

A *structure* is a pair $\mathbf{S} = (\mathbf{U}, \sigma)$, where \mathbf{U} is an *infinite* set called the *universe* and $\sigma : \mathbb{R} \rightarrow \text{pow}(\mathbf{U}^+)$ is an *interpretation* mapping each relation symbol r into a *finite* subset of $\mathbf{U}^{\#r}$. We consider only structures with finite interpretations, because the logic under consideration (defined below) can only describe sets of finite structures. The *support* $\text{supp}(\sigma) \stackrel{\text{def}}{=} \{u_i \mid \langle u_1, \dots, u_{\#r} \rangle \in \sigma(r), i \in [1..\#r]\}$ of an interpretation is the (necessarily finite) set of elements that occur in a tuple from the interpretation of a relation symbol. The support of a structure is the support of its interpretation.

Two structures (\mathbf{U}_1, σ_1) and (\mathbf{U}_2, σ_2) are *locally disjoint* iff $\sigma_1(r) \cap \sigma_2(r) = \emptyset$, for all $r \in \mathbb{R}$ and *disjoint* iff $\text{supp}(\sigma_1) \cap \text{supp}(\sigma_2) = \emptyset$. Two structures are *isomorphic* iff they differ only by a renaming of their elements (see, e.g., [EF95, Section A3] for a formal definition of isomorphism between structures).

We consider the *composition* as a partial binary operation between structures, defined as pointwise disjoint union of the interpretations of relation symbols:

Definition 2.1. The composition of two locally disjoint structures (U_1, σ_1) and (U_2, σ_2) is $(U_1, \sigma_1) \bullet (U_2, \sigma_2) \stackrel{\text{def}}{=} (U_1 \cup U_2, \sigma_1 \uplus \sigma_2)$, where $(\sigma_1 \uplus \sigma_2)(r) \stackrel{\text{def}}{=} \sigma_1(r) \uplus \sigma_2(r)$, for all $r \in \mathbb{R}$. The composition is undefined if (U_1, σ_1) and (U_2, σ_2) are not locally disjoint.

For example, Figure 2 shows the composition of two structures \mathbf{S}_1 and \mathbf{S}_2 , whose interpretations are represented as hyper-graphs with edges denoting tuples from the interpretation of relation symbols a, b and c, of arities 3, 2 and 2, respectively. Note that \mathbf{S}_1 and \mathbf{S}_2 are locally

disjoint but not disjoint, for instance the elements u_2 and u_3 are present in the support of both structures.




Figure 2: Composition of structures

2.2. Treewidth. A graph is a pair $G = (\mathcal{N}, \mathcal{E})$, such that \mathcal{N} is a finite set of *nodes* and $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$ is a set of *edges*. A (simple) *path* in G is a sequence of (pairwise distinct) nodes v_1, \dots, v_n , such that $(v_i, v_{i+1}) \in \mathcal{E}$, for all $i \in [1..n-1]$. We say that v_1, \dots, v_n is an *undirected path* if $\{(v_i, v_{i+1}), (v_{i+1}, v_i)\} \cap \mathcal{E} \neq \emptyset$ instead, for all $i \in [1..n-1]$. A set of nodes $S \subseteq \mathcal{N}$ is *connected in G* iff between any two nodes in S there is an undirected path in G that involves only nodes from S . A graph G is *connected* iff \mathcal{N} is connected in G .

Given a set Ω of labels, a Ω -*labeled unranked tree* is a tuple $T = (\mathcal{N}, \mathcal{E}, r, \lambda)$, where $(\mathcal{N}, \mathcal{E})$ is a graph, $r \in \mathcal{N}$ is a designated node called the *root*, such that there exists a unique simple path from r to any other node $n \in \mathcal{N} \setminus \{r\}$ and no path from r to r in $(\mathcal{N}, \mathcal{E})$. The mapping $\lambda : \mathcal{N} \rightarrow \Omega$ associates each node of the tree a label from Ω .

Definition 2.2. A *tree decomposition* of a structure $S = (U, \sigma)$ is a pow(U)-labeled unranked tree $T = (\mathcal{N}, \mathcal{E}, r, \lambda)$, such that the following hold:

- (1) for each relation symbol $r \in \mathbb{R}$ and each tuple $\langle u_1, \dots, u_{\#r} \rangle \in \sigma(r)$ there exists a node $n \in \mathcal{N}$, such that $\{u_1, \dots, u_{\#r}\} \subseteq \lambda(n)$,
- (2) for each element $u \in \text{supp}(\sigma)$, the set of nodes $\{n \in \mathcal{N} \mid u \in \lambda(n)\}$ is nonempty and connected in $(\mathcal{N}, \mathcal{E})$.

The *width* of the tree decomposition is $\text{wd}(T) \stackrel{\text{def}}{=} \max_{n \in \mathcal{N}} \text{card}(\lambda(n)) - 1$. The *treewidth* of the structure σ is $\text{tw}(\sigma) \stackrel{\text{def}}{=} \min\{\text{wd}(T) \mid T \text{ is a tree decomposition of } \sigma\}$.

Note that, since we consider only structures with finite support, tree decompositions are finite trees with finite sets as labels, hence the treewidth of a structure is a well-defined integer. A set of structures is *treewidth-bounded* iff the set of corresponding treewidths is finite and *treewidth-unbounded* otherwise. We assume basic acquaintance with the notions of grid and minor. It is known that a set of structures having infinitely many minors isomorphic to some $n \times n$ grid is treewidth-unbounded [Bod98].

2.3. Separation Logic of Relations. The *Separation Logic of Relations* (SLR) uses a set of *variables* $\mathbb{V} = \{x, y, \dots\}$ and a set of *predicates* $\mathbb{P} = \{A, B, \dots\}$ with given arities $\#A \geq 0$. A predicate of zero arity is called *nullary*.

The formulæ of SLR are defined by the syntax in Figure 3 (a). A variable is *free* if it does not occur within the scope of an existential quantifier and $\text{fv}(\phi)$ denotes the set of free variables of ϕ . A *sentence* is a formula with no free variables. For a formula ϕ , we denote by

ϕ^\exists the sentence obtained by existentially quantifying its free variables. A formula without quantifiers is called *quantifier-free*.

Instead of the standard boolean conjunction, SLR has a *separating conjunction* $*$. The formulæ $x \neq y$ and $A(x_1, \dots, x_{\#A})$ are called *disequalities* and *predicate atoms*, respectively. To alleviate notation, we denote by A the predicate atom $A()$, whenever A is nullary. A formula without predicate atoms is called *predicate-free*. A *qpf* formula is both quantifier- and predicate-free.

$$\phi := \text{emp} \mid x = y \mid x \neq y \mid r(x_1, \dots, x_{\#r}) \mid A(x_1, \dots, x_{\#A}) \mid \phi * \phi \mid \exists x . \phi$$

(a)

$$\begin{array}{ll} (\mathbf{U}, \sigma) \models_{\Delta}^{\mathfrak{s}} \text{emp} & \stackrel{\text{def}}{\iff} \sigma(r) = \emptyset, \text{ for all } r \in \mathbb{R} \\ (\mathbf{U}, \sigma) \models_{\Delta}^{\mathfrak{s}} x \sim y & \stackrel{\text{def}}{\iff} (\mathbf{U}, \sigma) \models_{\Delta}^{\mathfrak{s}} \text{emp} \text{ and } \mathfrak{s}(x) \sim \mathfrak{s}(y), \text{ for } \sim \in \{=, \neq\} \\ (\mathbf{U}, \sigma) \models_{\Delta}^{\mathfrak{s}} r(x_1, \dots, x_k) & \stackrel{\text{def}}{\iff} \sigma(r) = \{\langle \mathfrak{s}(x_1), \dots, \mathfrak{s}(x_k) \rangle\} \text{ and } \sigma(r') = \emptyset, \text{ for } r' \in \mathbb{R} \setminus \{r\} \\ (\mathbf{U}, \sigma) \models_{\Delta}^{\mathfrak{s}} A(y_1, \dots, y_n) & \stackrel{\text{def}}{\iff} \sigma \models_{\Delta}^{\mathfrak{s}} \phi[x_1/y_1, \dots, x_n/y_n], \text{ for some } A(x_1, \dots, x_n) \leftarrow \phi \in \Delta \\ (\mathbf{U}, \sigma) \models_{\Delta}^{\mathfrak{s}} \phi_1 * \phi_2 & \stackrel{\text{def}}{\iff} \text{exist structures } (\mathbf{U}_i, \sigma_i), \text{ where } (\mathbf{U}, \sigma) = (\mathbf{U}_1, \sigma_1) \bullet (\mathbf{U}_2, \sigma_2) \\ & \text{and } (\mathbf{U}_i, \sigma_i) \models_{\Delta}^{\mathfrak{s}} \phi_i, \text{ for both } i = 1, 2 \\ (\mathbf{U}, \sigma) \models_{\Delta}^{\mathfrak{s}} \exists x . \phi & \stackrel{\text{def}}{\iff} \sigma \models_{\Delta}^{\mathfrak{s}[x \leftarrow u]} \phi, \text{ for some } u \in \mathbf{U} \end{array}$$

(b)

Figure 3: The syntax (a) and semantics (b) of the Separation Logic of Relations

Definition 2.3. A *set of inductive definitions (SID)* is a finite set Δ of rules of the form $A(x_1, \dots, x_{\#A}) \leftarrow \phi$, where $x_1, \dots, x_{\#A}$ are pairwise distinct variables, called *parameters*, such that $\text{fv}(\phi) \subseteq \{x_1, \dots, x_{\#A}\}$.

The semantics of SLR is given by the satisfaction relation $(\mathbf{U}, \sigma) \models_{\Delta}^{\mathfrak{s}} \phi$ between structures and formulæ, parameterized by a store \mathfrak{s} and a SID Δ . We write $\mathfrak{s}[x \leftarrow u]$ for the store that maps x into u and agrees with \mathfrak{s} on all variables other than x . By $[x_1/y_1, \dots, x_n/y_n]$ we denote the substitution that replaces each free variable x_i by y_i in a formula ϕ . The result of applying the substitution $[x_1/y_1, \dots, x_n/y_n]$ to the formula ϕ is denoted as $\phi[x_1/y_1, \dots, x_n/y_n]$, where, by convention, the existentially quantified variables from ϕ are renamed to avoid clashes with y_1, \dots, y_n . Then $\models_{\Delta}^{\mathfrak{s}}$ is the least relation that satisfies the constraints in Figure 3 (b).

Note that the interpretation of equalities and relation atoms differs in SLR from first-order logic, namely $x = y$ requires that the structure is empty and $r(x_1, \dots, x_{\#r})$ denotes the structure in which all relations symbols are interpreted by empty sets, except for r , which contains the tuple of store values of $x_1, \dots, x_{\#r}$ only. Moreover, every structure (\mathbf{U}, σ) , such that $(\mathbf{U}, \sigma) \models_{\Delta}^{\mathfrak{s}} \phi$, interprets each relation symbol as a finite set of tuples, defined by a finite least fixpoint iteration over the rules from Δ . The assumption that each structure has an infinite universe excludes the cases in which a formula becomes unsatisfiable because there are not enough elements to instantiate the quantifiers introduced by the unfolding of the rules, thus simplifying the definitions.

If ϕ is a sentence (resp. a predicate-free formula), we omit the store \mathfrak{s} (resp. the SID Δ) from $\mathbf{S} \models_{\Delta}^{\mathfrak{s}} \phi$. For a SLR sentence ϕ , let $\llbracket \phi \rrbracket_{\Delta} \stackrel{\text{def}}{=} \{\mathbf{S} \mid \mathbf{S} \models_{\Delta} \phi\}$ be the set of Δ -models of ϕ . If ϕ is, moreover, predicate-free we say that ϕ is *satisfiable* iff $\llbracket \phi \rrbracket \neq \emptyset$.

For a qpf formula ϕ , we write $x \approx_\phi y$ (resp. $x \not\approx_\phi y$) iff $x = y$ is (resp. is not) a logical consequence of ϕ , i.e., $\mathfrak{s}(x) = \mathfrak{s}(y)$ for each store \mathfrak{s} and structure \mathbb{S} , such that $\mathbb{S} \models^{\mathfrak{s}} \phi$. Note that $x \not\approx_\phi y$ is the negation of $x \approx_\phi y$, which is different from that $x \neq y$ is implied by ϕ . We define several quantitative measures relative to SIDs:

Definition 2.4. Let Δ be a SID. We denote by:

- ▷ $\text{maxVars}(\Delta)$ the maximum number of variables that occur, either free or existentially quantified, in a rule from Δ ,
- ▷ $\text{maxRelAtoms}(\Delta)$ the maximum number of relation atoms that occur in a rule from Δ ,
- ▷ $\text{maxPredAtoms}(\Delta)$ the maximum number of predicates that occur in a rule from Δ ,
- ▷ $\text{maxRelArity}(\Delta)$ the maximum arity of relation symbols occurring in Δ ,
- ▷ $\text{predNo}(\Delta)$ the number of predicate symbols occurring in Δ ,
- ▷ $\text{relNo}(\Delta)$ the number of relation symbols occurring in Δ .

2.4. Simplifying assumptions. In the rest of this paper, we simplify the technical development by two assumptions, that lose no generality. The first assumption is that no equalities occur in the given SID (Lemma 2.6).

Definition 2.5. A formula is *equality-free* iff it contains no equalities nor predicate atoms in which the same variable occurs twice. A rule $\mathbf{A}(x_1, \dots, x_n) \leftarrow \phi$ is equality-free iff ϕ is equality-free. A SID is equality-free iff it consists of equality-free rules.

Lemma 2.6. *Given a SID Δ , one can build an equality-free SID Δ' , such that $\llbracket \mathbf{A} \rrbracket_\Delta = \llbracket \mathbf{A} \rrbracket_{\Delta'}$, for each nullary predicate \mathbf{A} . Moreover, all quantitative measures (Definition 2.4) of Δ' are the same as for Δ , except for $\text{predNo}(\Delta') \leq \text{predNo}(\Delta) \cdot \text{maxVars}(\Delta)^{\text{maxVars}(\Delta)}$.*

Proof. See [IZ23, Lemma 9]. The construction of Δ' considers predicates $\mathbf{A}_{I_1, \dots, I_n}$, where \mathbf{A} is a predicate symbol that occurs in Δ and $I_1 \uplus \dots \uplus I_n = [1.. \#\mathbf{A}]$ is a partition. Since the number of partitions of $[1.. \#\mathbf{A}]$ is asymptotically bounded by $\#\mathbf{A}^{\#\mathbf{A}} \leq \text{maxVars}(\Delta)^{\text{maxVars}(\Delta)}$, we obtain the bound on $\text{predNo}(\Delta')$. \square

The following notion of *unfolding* is used to define several technical notions and state the second simplifying assumption. Let ϕ and ψ be formulæ and Δ be a SID. We denote by $\phi \Rightarrow_\Delta \psi$ the fact that ψ is obtained by replacing a predicate atom $\mathbf{A}(y_1, \dots, y_n)$ in ϕ by a formula $\rho[x_1/y_1, \dots, x_n/y_n]$, where $\mathbf{A}(x_1, \dots, x_n) \leftarrow \rho$ is a rule from Δ . A Δ -*unfolding* is a sequence of formulæ $\phi_1 \Rightarrow_\Delta \dots \Rightarrow_\Delta \phi_n$. The Δ -unfolding is *complete* if the last formula is predicate-free. The following statement is a direct consequence of the semantics of SLR:

Proposition 2.7. *Let ϕ be a sentence, Δ a SID and \mathbb{S} a structure. Then $\mathbb{S} \in \llbracket \phi \rrbracket_\Delta$ iff $\mathbb{S} \models^{\mathfrak{s}} \psi$, for a store \mathfrak{s} and complete Δ -unfolding $\phi \Rightarrow_\Delta^* \exists x_1 \dots \exists x_n . \psi$, where ψ is a qpf formula.*

Proof. “ \Leftarrow ” By induction on the definition of the satisfaction relation $\models_\Delta^{\mathfrak{s}}$. “ \Rightarrow ” By induction on the length of the Δ -unfolding. \square

The second assumption is that any Δ -unfolding of a nullary predicate by the given SID yields a predicate-free formula that is satisfiable. Again, this assumption loses no generality (Lemma 2.9).

Definition 2.8. A SID Δ is *all-satisfiable* for a nullary predicate \mathbf{A} iff each predicate-free formula ϕ which is the outcome of a complete Δ -unfolding $\mathbf{A} \Rightarrow_\Delta^* \phi$ is satisfiable.

Lemma 2.9. *Given a SID Δ and a nullary predicate A , one can build a SID $\bar{\Delta}$ all-satisfiable for A , such that $\llbracket A \rrbracket_{\Delta} = \llbracket A \rrbracket_{\bar{\Delta}}$. Moreover, all quantitative measures (Definition 2.4) of Δ' are the same as for Δ , except for $\text{predNo}(\bar{\Delta}) \leq \text{predNo}(\Delta) \cdot \text{relNo}(\Delta) \cdot \text{maxVars}(\Delta)^{\text{maxRelArity}(\Delta)}$.*

Proof. For space reasons, this proof is given in subsection A.1. \square

2.5. The treewidth boundedness problem. We are ready to state the main problem addressed in this paper. Before, we state two technical lemmas that state several relations between qpf formulæ and the upper bounds on the treewidth of their models:

Lemma 2.10. *Let ϕ, ψ be qpf formulæ, $x_0, x_1, x_2, \dots, x_k$ variables and r a relation symbol of arity k . Then, the following hold:*

- (1) $\text{tw}(\llbracket (\phi * \bigstar_{i=1}^k x_0 = x_i) \rrbracket) \leq \text{tw}(\llbracket \phi \rrbracket)$,
- (2) $\text{tw}(\llbracket \phi \rrbracket) - 1 \leq \text{tw}(\llbracket (\phi * \bigstar_{i=1}^k x_0 \neq x_i) \rrbracket) \leq \text{tw}(\llbracket \phi \rrbracket)$ if $\phi * \bigstar_{i=1}^k x_0 \neq x_i$ satisfiable,
- (3) $\text{tw}(\llbracket \phi \rrbracket) - 1 \leq \text{tw}(\llbracket (\phi * r(x_1, \dots, x_k)) \rrbracket) \leq \text{tw}(\llbracket \phi \rrbracket) + k$ if $\phi * r(x_1, \dots, x_k)$ satisfiable.
- (4) $\text{tw}(\llbracket (\phi * \psi) \rrbracket) \leq \text{tw}(\llbracket \phi \rrbracket) + \text{card}(\text{fv}(\psi))$ if ψ contains only relation atoms.

Proof. For space reasons, this proof is given in subsection A.2. \square

Lemma 2.11. *Let ϕ and ψ be qpf formulæ and $F \stackrel{\text{def}}{=} \text{fv}(\phi) \cap \text{fv}(\psi)$, such that $\phi * \psi$ is satisfiable and $x \not\approx_{\phi} y$, for all $x, y \in F$. Let $\eta \stackrel{\text{def}}{=} \bigstar_{x, y \in F, x \approx_{\psi} y} x = y$. Then, $\text{tw}(\llbracket (\phi * \eta) \rrbracket) \leq \text{tw}(\llbracket (\phi * \psi) \rrbracket) + \text{card}(F)$.*

Proof. For space reasons, this proof is given in subsection A.3. \square

The main result of this paper is a decidability proof for the following decision problem:

Definition 2.12. The TWB^{SLR} problem asks whether the set $\llbracket \phi \rrbracket_{\Delta}$ is treewidth-bounded, for an SID Δ and SLR sentence ϕ given as input.

This result is tightened by a proof of the undecidability of the treewidth-boundedness problem for first-order logic (section 5), that further improves our understanding of the relation between the expressivity of classical and substructural logics. The decidability proof proceeds in two steps. First, we show the decidability of the problem for sentences of the form A , where A is a nullary predicate symbol, and a class of SIDs having a particular property, called *expandability*, defined below (section 3). Second, we show how to reduce the treewidth-boundedness problem for arbitrary sentences and SIDs to the problem for nullary predicate atoms and expandable SIDs (section 4).

3. EXPANDABLE SETS OF INDUCTIVE DEFINITIONS

This section introduces the formal definitions of canonical models and expandable SIDs, needed for the first part of the proof of decidability of the TWB^{SLR} problem. The main result of this section is that the treewidth boundedness problem is decidable for the sets of models of a nullary predicate defined by an expandable SID.

For simplicity, in the rest of this paper we shall represent sentences ϕ by nullary predicate atoms A . This loses no generality since $\llbracket \phi \rrbracket_{\Delta} = \llbracket A \rrbracket_{\Delta \cup \{A \leftarrow \phi\}}$ provided that A is not defined by any other rule in Δ . In the rest of this section we fix an arbitrary SID Δ and nullary predicate A .

3.1. Canonical models. Intuitively, a Δ -model of A is *canonical* if it can be defined using a store that matches only those variables that are equated in the outcome of the complete Δ -unfolding of A that “produced” the model, in the sense of Proposition 2.7. A *rich canonical model* records, moreover, the disequalities introduced during the unfolding.

Definition 3.1. A store \mathfrak{s} is *canonical for ϕ* iff $\mathfrak{s}(x) = \mathfrak{s}(y)$ only if $x \approx_\phi y$, for all $x, y \in \text{fv}(\phi)$. A *rich canonical Δ -model* of a sentence ϕ is a pair (S, \mathfrak{d}) , where $S = (U, \sigma)$ is a structure and $\mathfrak{d} \subseteq U \times U$ is a symmetric relation, such that there exists a complete Δ -unfolding $\phi \Rightarrow_\Delta^* \exists x_1 \dots \exists x_n . \psi$, where ψ is qpf, and a store \mathfrak{s} canonical for ψ , such that $S \models^{\mathfrak{s}} \psi$ and $\mathfrak{d}(u, v)$ iff there exist variables $x \in \mathfrak{s}^{-1}(u)$, $y \in \mathfrak{s}^{-1}(v)$ and the disequality $x \neq y$ occurs in ψ . We denote by $[[\phi]]_\Delta^r$ the set of rich canonical Δ -models of ϕ and $[[\phi]]_\Delta^c \stackrel{\text{def}}{=} \{S \mid (S, \mathfrak{d}) \in [[\phi]]_\Delta^r\}$ the set of *canonical Δ -models* of ϕ . If ϕ is predicate-free, we write $[[\phi]]^c$ (resp. $[[\phi]]^r$) instead of $[[\phi]]_\Delta^c$ (resp. $[[\phi]]_\Delta^r$).

A store \mathfrak{s} is *injective* over a set of variables x_1, \dots, x_n iff $\mathfrak{s}(x_i) = \mathfrak{s}(x_j)$ implies $i = j$, for all $i, j \in [1..n]$. Note that the canonical Δ -models of an equality free SID Δ can be defined considering injective, instead of canonical stores. Nevertheless, this more general definition of canonical models using canonical stores will become useful later on, when predicate-free formulæ with equalities will be considered.

Canonical models are important for two reasons. First, their treewidth is bounded:

Lemma 3.2. $\text{tw}(S) \leq \max\text{Vars}(\Delta) - 1$, for any $S \in [[A]]_\Delta^c$.

Proof. Let $S = (U, \sigma) \in [[A]]_\Delta^c$ be a canonical Δ -model of A . We define a tree decomposition $T = (\mathcal{N}, \mathcal{E}, r, \lambda)$ of S as follows. The graph of T is any derivation tree of Δ whose outcome is S . This is a tree labeled with rules from Δ , whose parent-child relation is defined as follows: if n is a node labeled with a rule ρ , for each predicate atom $B(z_1, \dots, z_{\#B})$ that occurs in ρ , there is exactly one child m of n whose label is a rule that defines B . The bag $\lambda(n)$ contains exactly those elements that are the store values of the variables occurring free or bound in ρ . We check that T is a tree decomposition of S by proving the two points of Definition 2.2:

- (1) each tuple $\langle u_1, \dots, u_{\#r} \rangle \in \sigma(r)$ occurs in S because of a relation atom $r(z_1, \dots, z_{\#r})$ that occurs in the label of a node n from the parse tree. Then $u_1, \dots, u_{\#r} \in \lambda(n)$, by the definition of T .
- (2) let $n, m \in \mathcal{N}$ be nodes of T and $u \in \lambda(n) \cap \lambda(m)$ be an element. Then the label of each node on the path between n and m in the parse tree contains a variable whose store value is u , hence the set $\{p \in \mathcal{N} \mid u \in \lambda(p)\}$ is non-empty and connected in T . \square

Second, any model is obtained via an *internal fusion* of a rich canonical model. The internal fusion is a unary operation that takes as input a structure and outputs a set of structures obtained by joining certain elements from its support. This operation is formally defined as quotienting with respect to certain equivalence relations:

Definition 3.3. Let $S = (U, \sigma)$ be a structure and $\approx \subseteq U \times U$ be an equivalence relation, where $[u]_\approx$ is the equivalence class of $u \in U$. The *quotient* $S_{/\approx} = (U_{/\approx}, \sigma_{/\approx})$ is $U_{/\approx} \stackrel{\text{def}}{=} \{[u]_\approx \mid u \in U\}$ and $\sigma_{/\approx}(r) \stackrel{\text{def}}{=} \{\langle [u_1]_\approx, \dots, [u_{\#r}]_\approx \rangle \mid \langle u_1, \dots, u_{\#r} \rangle \in \sigma(r)\}$, for all $r \in \mathbb{R}$.

For example, Figure 4 (a) shows the outcome of quotienting a structure with respect to an equivalence relation, whose equivalence classes are encircled with dashed lines.

A fusion operation glues elements without losing tuples from the interpretation of a relation symbol. For this reason, we consider only equivalence relations that are *compatible* with a given structure and define internal fusion as the following unary operation:

Definition 3.4. An equivalence relation $\approx \subseteq \mathbf{U} \times \mathbf{U}$ is *compatible* with a structure $\mathbf{S} = (\mathbf{U}, \sigma)$ iff for all $r \in \mathbb{R}$ and any two tuples $\langle u_1, \dots, u_{\#r} \rangle, \langle v_1, \dots, v_{\#r} \rangle \in \sigma(r)$, there exists $i \in [1.. \#r]$ such that $u_i \not\approx v_i$. An *internal fusion* of \mathbf{S} is a structure isomorphic to $\mathbf{S}_{/\approx}$, for an equivalence relation \approx compatible with \mathbf{S} . Let $\text{IF}(\mathbf{S})$ be the set of internal fusions of \mathbf{S} and $\text{IF}(\mathcal{S}) \stackrel{\text{def}}{=} \bigcup_{\mathbf{S} \in \mathcal{S}} \text{IF}(\mathbf{S})$, for a set \mathcal{S} of structures.

For example, Figure 4 (b) shows a possible internal fusion of a structure. Note that the equivalence relation from Figure 4 (a) is not compatible with the structure and cannot be used in a fusion.




Figure 4: Quotient (a) and internal fusion (b)

For technical reasons, we introduce also the internal fusion of a rich canonical model, as quotienting with respect to an equivalence relation that does not violate the disequality relation:

Definition 3.5. An equivalence relation $\approx \subseteq \mathbf{U} \times \mathbf{U}$ is *compatible* with a rich canonical model $(\mathbf{S}, \mathfrak{d})$ iff it is compatible with $\mathbf{S} = (\mathbf{U}, \sigma)$ and $\mathfrak{d}(u, v)$ only if $u \not\approx v$. We denote by $\widetilde{\text{IF}}(\mathbf{S}, \mathfrak{d})$ the set of structures isomorphic to $\mathbf{S}_{/\approx}$, where \approx is some equivalence relation compatible with $(\mathbf{S}, \mathfrak{d})$.

The following lemma relates the sets of models, canonical and rich canonical models of a sentence, via the two types of internal fusion:

Lemma 3.6. $\llbracket \mathbf{A} \rrbracket_{\Delta} = \widetilde{\text{IF}}(\llbracket \mathbf{A} \rrbracket'_{\Delta}) \subseteq \text{IF}(\llbracket \mathbf{A} \rrbracket^c_{\Delta})$.

Proof. It is sufficient to prove $\llbracket \mathbf{A} \rrbracket_{\Delta} = \widetilde{\text{IF}}(\llbracket \mathbf{A} \rrbracket'_{\Delta})$, since $\widetilde{\text{IF}}(\llbracket \mathbf{A} \rrbracket'_{\Delta}) \subseteq \text{IF}(\llbracket \mathbf{A} \rrbracket^c_{\Delta})$ is immediate, by Definition 3.5, because any equivalence relation that is compatible with a rich canonical Δ -model $(\mathbf{S}, \mathfrak{d})$ is also compatible with the canonical Δ -model \mathbf{S} .

“ \subseteq ” Let $\mathbf{S} \in \llbracket \mathbf{A} \rrbracket_{\Delta}$ be a structure. By Proposition 2.7, we have $\mathbf{S} \models \exists y_1 \dots \exists y_m . \psi$, where ψ is a qpf formula, such that $\text{fv}(\psi) = \{y_1, \dots, y_m\}$ and $\mathbf{A} \Rightarrow_{\Delta}^* \exists y_1 \dots \exists y_m . \psi$ is a complete Δ -unfolding. Then there exists a store \mathfrak{s} , such that $\mathbf{S} \models^{\mathfrak{s}} \psi$. Let $\overline{\mathbf{S}} = (\overline{\mathbf{U}}, \overline{\sigma})$ be a structure and $\overline{\mathfrak{s}}$ be an injective store over y_1, \dots, y_m . Since Δ is equality-free, there are no equality atoms in ψ , hence such a structure and injective store exist. We consider $\approx \subseteq \overline{\mathbf{U}} \times \overline{\mathbf{U}}$ to be the least equivalence relation such that $\overline{\mathfrak{s}}(y_i) \approx \overline{\mathfrak{s}}(y_j) \stackrel{\text{def}}{\iff} \mathfrak{s}(y_i) = \mathfrak{s}(y_j)$, for all $1 \leq i < j \leq m$. To prove that \approx is compatible with $\overline{\mathbf{S}}$, consider two tuples $\langle \overline{\mathfrak{s}}(z_1), \dots, \overline{\mathfrak{s}}(z_{\#r}) \rangle, \langle \overline{\mathfrak{s}}(z'_1), \dots, \overline{\mathfrak{s}}(z'_{\#r}) \rangle \in \overline{\sigma}(r)$, for some $r \in \mathbb{R}$ and suppose, for a contradiction, that $\overline{\mathfrak{s}}(z_i) \approx \overline{\mathfrak{s}}(z'_i)$, for all $i \in [1.. \#r]$. Then $r(z_1, \dots, z_{\#r}) * r(z'_1, \dots, z'_{\#r})$ is a subformula of ψ , modulo a reordering of atoms. By the definition of \approx , we have $\mathfrak{s}(z_i) = \mathfrak{s}(z'_i)$, for all $i \in [1.. \#r]$, in contradiction with $\mathbf{S} \models^{\mathfrak{s}} \psi$ and the semantics of the separating conjunction. Since $\overline{\mathbf{S}} \models^{\overline{\mathfrak{s}}} \psi$ and $\overline{\mathfrak{s}}$ is injective over y_1, \dots, y_m ,

we obtain that $\bar{\mathfrak{s}}|_{\{y_1, \dots, y_m\}}$ is a bijection between $\{y_1, \dots, y_m\}$ and $\text{supp}(\bar{\sigma})$ hence $\bar{\mathfrak{s}}^{-1}(u)$ is a singleton, for each $u \in \text{supp}(\bar{\sigma})$. Let $\mathfrak{d} \subseteq \bar{U} \times \bar{U}$ be the relation defined as $\mathfrak{d}(u, v)$ iff the disequality $\bar{\mathfrak{s}}^{-1}(u) \neq \bar{\mathfrak{s}}^{-1}(v)$ occurs in ψ . Then \approx is compatible with $(\bar{S}, \mathfrak{d}) \in \llbracket \mathbf{A} \rrbracket_{\Delta}^r$ hence $\bar{S}_{/\approx} \in \widetilde{\text{IF}}(\llbracket \mathbf{A} \rrbracket_{\Delta}^r)$. Finally, the mapping $h : \text{supp}(\sigma) \rightarrow \text{supp}(\bar{\sigma})$ defined as $h(\mathfrak{s}(y_i)) \stackrel{\text{def}}{=} [\bar{\mathfrak{s}}(y_i)]_{\approx}$, for all $i \in [1..m]$ is shown to be an isomorphism between S and $\bar{S}_{/\approx}$, leading to $S \in \widetilde{\text{IF}}(\llbracket \mathbf{A} \rrbracket_{\Delta}^r)$, by the fact that the set $\widetilde{\text{IF}}(\llbracket \mathbf{A} \rrbracket_{\Delta}^r)$ is closed under isomorphism (Definition 3.5).

“ \supseteq ” Let $S \in \widetilde{\text{IF}}(\llbracket \mathbf{A} \rrbracket_{\Delta}^r)$ be a structure. Then there exists a rich canonical Δ -model $(\bar{S}, \mathfrak{d}) \in \llbracket \mathbf{A} \rrbracket_{\Delta}^r$, where $\bar{S} = (\bar{U}, \bar{\sigma})$ and an equivalence relation $\approx \subseteq \bar{U} \times \bar{U}$ such that \approx is compatible with (\bar{S}, \mathfrak{d}) and S is isomorphic to $\bar{S}_{/\approx}$. Since $(\bar{S}, \mathfrak{d}) \in \llbracket \mathbf{A} \rrbracket_{\Delta}^r$, there exists a complete Δ -unfolding $\mathbf{A} \Rightarrow_{\Delta}^* \exists y_1 \dots \exists y_m . \psi$, such that ψ is qpf and a store \mathfrak{s} , injective over y_1, \dots, y_m , such that $S \models^{\mathfrak{s}} \psi$ and $\mathfrak{d}(\mathfrak{s}(z), \mathfrak{s}(z'))$ for each disequality $z \neq z'$ from ψ . Let $\bar{\mathfrak{s}}$ be the store defined as $\bar{\mathfrak{s}}(y_i) = [\mathfrak{s}(y_i)]_{\approx}$, for all $i \in [1..m]$. We prove $\bar{S}_{/\approx} \models^{\bar{\mathfrak{s}}} \psi$ by induction on the structure of ψ , considering the following cases:

- ▷ $\psi = y_i \neq y_j$: because \approx is compatible with (\bar{S}, \mathfrak{d}) , we have $[\mathfrak{s}(y_i)]_{\approx} \neq [\mathfrak{s}(y_j)]_{\approx}$, hence $\bar{\mathfrak{s}}(y_i) \neq \bar{\mathfrak{s}}(y_j)$.
- ▷ $\psi = r(y_{i_1}, \dots, y_{i_{\#r}})$: because $\bar{S} \models^{\mathfrak{s}} r(y_{i_1}, \dots, y_{i_{\#r}})$, we have $\bar{\sigma}(r) = \{\langle \mathfrak{s}(y_{i_1}), \dots, \mathfrak{s}(y_{i_{\#r}}) \rangle\}$ and $\bar{\sigma}_{/\approx}(r) = \{\langle [\mathfrak{s}(y_{i_1})]_{\approx}, \dots, [\mathfrak{s}(y_{i_{\#r}})]_{\approx} \rangle\}$, by Definition 3.3.
- ▷ $\psi = \psi_1 * \psi_2$: because $\bar{S} \models^{\mathfrak{s}} \psi_1 * \psi_2$, there exist locally disjoint structures $\bar{S}_1 \bullet \bar{S}_2 = \bar{S}$, such that $\bar{S}_i \models^{\mathfrak{s}} \psi_i$, for $i = 1, 2$. Since \approx is compatible with S , the structures $\bar{S}_{1/\approx}$ and $\bar{S}_{2/\approx}$ are locally disjoint, by Definition 3.4. Then their composition is defined and we have $\bar{S}_{/\approx} = \bar{S}_{1/\approx} \bullet \bar{S}_{2/\approx}$. By the inductive hypothesis, we have $\bar{S}_{i/\approx} \models^{\bar{\sigma}} \psi_i$, for $i = 1, 2$, thus $\bar{S}_{/\approx} \models^{\bar{\sigma}} \psi_1 * \psi_2$.

Hence $\bar{S}_{/\approx} \in \llbracket \mathbf{A} \rrbracket_{\Delta}$ and $S \in \llbracket \mathbf{A} \rrbracket_{\Delta}$ follows, since the set $\llbracket \mathbf{A} \rrbracket_{\Delta}$ is closed under isomorphism, see, e.g., [IZ23, Proposition 7] for a proof. \square

3.2. Expandable sets of inductive definitions. We introduce the notion of *expandable* SID, a key ingredient of our proof of decidability for the TWB^{SLR} problem. A structure is a *substructure* of another if the former is obtained from the latter by removing elements from its support:

Definition 3.7. Let $S_i = (U_i, \sigma_i)$ be structures, for $i = 1, 2$. S_1 is *included* in S_2 iff $U_1 \subseteq U_2$ and $\sigma_1(r) \subseteq \sigma_2(r)$, for all $r \in \mathbb{R}$. S_1 is a *substructure* of S_2 , denoted $S_1 \sqsubseteq S_2$, iff $S_1 \subseteq S_2$ and $\sigma_1(r) = \{\langle u_1, \dots, u_{\#r} \rangle \in \sigma_2(r) \mid u_1, \dots, u_{\#r} \in \text{supp}(\sigma_1)\}$, for all $r \in \mathbb{R}$.

A SID is *expandable* if any set of canonical models of a sentence are all substructures of the same canonical model of that sentence, that can be, moreover, placed “sufficiently far away” one from another.

Definition 3.8. A SID Γ is *expandable* for a nullary predicate \mathbf{A} iff for each sequence of pairwise disjoint canonical models $S_1 = (U_1, \sigma_1), \dots, S_n = (U_n, \sigma_n) \in \llbracket \mathbf{A} \rrbracket_{\Gamma}^c$, there exists a rich canonical model $(S, \mathfrak{d}) \in \llbracket \mathbf{A} \rrbracket_{\Gamma}^r$, where $S = (U, \sigma)$, such that:

- (1) $S_1 \bullet \dots \bullet S_n \sqsubseteq S$,
- (2) $\mathfrak{d}(u, v)$ holds for no $u \in \text{supp}(\sigma_i)$ and $v \in \text{supp}(\sigma_j)$, where $1 \leq i < j \leq n$, and
- (3) for no relation symbol $r \in \mathbb{R}$ and tuples $\langle u_1, \dots, u_{\#r} \rangle, \langle v_1, \dots, v_{\#r} \rangle \in \sigma(r)$ there exist $1 \leq i < j \leq n$, such that $\{u_1, \dots, u_{\#r}\} \cap \text{supp}(\sigma_i) \neq \emptyset$, $\{v_1, \dots, v_{\#r}\} \cap \text{supp}(\sigma_j) \neq \emptyset$ and $\{u_1, \dots, u_{\#r}\} \cap \{v_1, \dots, v_{\#r}\} \neq \emptyset$.

Example 3.9. The SID Δ from Figure 5 is expandable for A , because any choice of pairwise disjoint canonical models of A (here S_1 , S_2 and S_3) can be embedded in a canonical model of A , such that there are no pairs, in the interpretation of the binary relation symbol e , that stretch from one such model to another (the pairs not entirely inside the support of either S_1 , S_2 or S_3 are depicted in dashed lines).




Figure 5: An expandable SID

Example 3.10. (continued from Example 3.9) Consider the SID Δ' obtained from Δ by changing its last rule into $B(x_1) \leftarrow \text{emp}$. Δ' is not expandable for A , because the canonical models of A are acyclic chains of elements, whose neighbours are related by e , such that all but the last element is labeled by a . Then, two such structures S_1 and S_2 cannot be embedded in a third structure S as substructures, because of the last non-labeled element of S_1 that occurs in the middle of S and must be labeled by a . This violates the definition of substructures (see Definition 3.7), in which the labeling of an element is the same in a substructure and in its enclosing structure.

The *external fusion* is a binary operation that glues elements from disjoint structures:

Definition 3.11. An *external fusion* of the structures $S_1 = (U_1, \sigma_1)$ and $S_2 = (U_2, \sigma_2)$ is a structure isomorphic to $(S'_1 \bullet S'_2)_{/\approx}$, where $S'_i = (U'_i, \sigma'_i)$ are disjoint isomorphic copies of S_i and $\approx \subseteq U'_1 \times U'_2$ is the smallest equivalence relation containing a nonempty $\text{supp}(\sigma'_1) - \text{supp}(\sigma'_2)$ matching that is compatible with $S'_1 \bullet S'_2$. Let $\text{EF}(S_1, S_2)$ be the set of external fusions of S_1 and S_2 . For a set of structures \mathcal{S} , let $\text{EF}^*(\mathcal{S})$ (resp. $\text{IEF}^*(\mathcal{S})$) be the closure of \mathcal{S} under taking external (resp. both internal and external) fusions.

For example, Figure 6 shows the external fusion of two disjoint structures via a matching relation (the equivalence classes of the matching relation are encircled with dashed lines). Note that the conditions (2) and (3) of Definition 3.8 ensure that the external fusion of these substructures is not hindered by their position inside the larger structure. For instance, any matching relation between the supports of the substructures S_1 , S_2 and S_3 from Figure 5 can be used to define an external fusion of these structures. This is because there are no pairs, from the interpretation of the e relation symbol in S , that have an element in common and the other non-common elements in the support of two different substructures. If such pairs existed, the non-common elements could not be fused by an equivalence relation compatible with $S_i \bullet S_j$, for any $1 \leq i < j \leq 3$.

The following lemma proves one direction of the equivalence between the treewidth boundedness of $\llbracket A \rrbracket_\Delta$ and that of the set of structures obtained by applying both internal and external fusion to the canonical models from $\llbracket A \rrbracket_\Delta^c$.

Lemma 3.12. *Let Δ be an expandable SID for a nullary predicate A . Then, (1) $\text{IEF}^*(\llbracket A \rrbracket_\Delta^c)$ is treewidth-bounded only if (2) $\llbracket A \rrbracket_\Delta$ is treewidth-bounded only if (3) $\text{EF}^*(\llbracket A \rrbracket_\Delta^c)$ is treewidth-bounded.*



Figure 6: External fusion

Proof. “(1) \Rightarrow (2)” $\text{IF}(\llbracket \mathbf{A} \rrbracket_{\Delta}^c) \subseteq \text{IEF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$ holds trivially, by Definition 3.11, leading to $\llbracket \mathbf{A} \rrbracket_{\Delta} \subseteq \text{IEF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$, by Lemma 3.6. “(2) \Rightarrow (3)” Let $S = (U, \sigma) \in \text{EF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$ be a structure. It is sufficient to prove that $S \sqsubseteq S'$ for another structure $S' \in \llbracket \mathbf{A} \rrbracket_{\Delta}$, because $\text{tw}(S) \leq \text{tw}(S')$, in this case. Then there exist pairwise disjoint structures $S_1 = (U_1, \sigma_1), \dots, S_n = (U_n, \sigma_n) \in \llbracket \mathbf{A} \rrbracket_{\Delta}^c$ and an equivalence relation $\approx \subseteq (\bigcup_{i=1}^n U_i) \times (\bigcup_{i=1}^n U_i)$, that is compatible with $S_1 \bullet \dots \bullet S_n$, matches only elements from different structures and is not the identity, such that S is isomorphic to $(S_1 \bullet \dots \bullet S_n)_{/\approx}$. By Definition 3.8, there exists a rich canonical model $(S'', \mathfrak{d}) \in \llbracket \mathbf{A} \rrbracket'_{\Delta}$, such that (1) $S \sqsubseteq S''$, (2) $\mathfrak{d}(u, v)$ holds for no $u \in \text{supp}(\sigma_i)$ and $v \in \text{supp}(\sigma_j)$, where $1 \leq i < j \leq n$, and (3) for no relation symbol $r \in \mathbb{R}$ and tuples $\langle u_1, \dots, u_{\#r} \rangle, \langle v_1, \dots, v_{\#r} \rangle \in \sigma(r)$, there exist $1 \leq i < j \leq n$, such that $\{u_1, \dots, u_{\#r}\} \cap \text{supp}(\sigma_i) \neq \emptyset$, $\{v_1, \dots, v_{\#r}\} \cap \text{supp}(\sigma_j) \neq \emptyset$ and $\{u_1, \dots, u_{\#r}\} \cap \{v_1, \dots, v_{\#r}\} \neq \emptyset$. By the last two conditions, \approx is compatible with (S'', \mathfrak{d}) , leading to $S''_{/\approx} \in \widetilde{\text{IF}}(\llbracket \mathbf{A} \rrbracket'_{\Delta}) = \llbracket \mathbf{A} \rrbracket_{\Delta}$ by Lemma 3.6. We conclude by taking $S' = S''_{/\approx}$. \square

The missing direction $\text{EF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c) \Rightarrow \text{IEF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$, that allows to establish the equivalence of the three points of Lemma 3.12, requires the introduction of further technical notions. The proof of the main result of this section relies on an algorithm for the treewidth boundedness of sets $\text{EF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$, obtained by external fusion of disjoint canonical Δ -models of \mathbf{A} . By the equivalence of the treewidth boundedness of the sets $\llbracket \mathbf{A} \rrbracket_{\Delta}$ and $\text{EF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$ (Lemma 3.12 and Lemma 3.26), this is also an algorithm for the TWB^{SLR} problem for expandable SIDs .

3.3. Color schemes. Intuitively, the *color* of an element from the support of a structure is the set of relation symbols labeling solely that element. For the given set \mathbb{R} of relation symbols, we define the set of *colors* as $\mathbb{C} \stackrel{\text{def}}{=} \text{pow}(\mathbb{R})$. The elements of a structure are labeled with colors as follows:

Definition 3.13. The *coloring* of a structure $S = (U, \sigma)$ is the mapping $\mathcal{C}_S : U \rightarrow \mathbb{C}$ defined as $\mathcal{C}_S(u) \stackrel{\text{def}}{=} \{r \in \mathbb{R} \mid \langle u, \dots, u \rangle \in \sigma(r)\}$.

Moreover, we define an abstraction of structures as finite multisets of colors:

Definition 3.14. The *multiset color abstraction* $S^{\#} \in \text{mpow}(\mathbb{C})$ of a structure $S = (U, \sigma)$ is $S^{\#} \stackrel{\text{def}}{=} \{\mathcal{C}_S(u) \mid u \in \text{supp}(\sigma)\}$. For an integer $k \geq 0$, the *k-multiset color abstraction* $S^{\#k} \subseteq \text{mpow}(\mathbb{C})$ is $S^{\#k} \stackrel{\text{def}}{=} \{M \subseteq S^{\#} \mid \text{card}(M) \leq k\}$. These abstractions are lifted to sets \mathcal{S} of structures, yielding the sets of multisets $\mathcal{S}^{\#} \stackrel{\text{def}}{=} \{S^{\#} \mid S \in \mathcal{S}\}$ and $\mathcal{S}^{\#k} \stackrel{\text{def}}{=} \bigcup_{S \in \mathcal{S}} S^{\#k}$.

Colors are organized in *RGB color schemes*, defined below:

Definition 3.15. A partition $(\mathbb{C}^{\text{red}}, \mathbb{C}^{\text{green}}, \mathbb{C}^{\text{blue}})$ of \mathbb{C} is an *RGB-color scheme* iff:

- (1) $\mathcal{C}_1 \cap \mathcal{C}_2 \neq \emptyset$, for all $\mathcal{C}_1, \mathcal{C}_2 \in \mathbb{C}^{\text{blue}}$,
- (2) $\mathcal{C}_1 \cap \mathcal{C}_2 \neq \emptyset$, for all $\mathcal{C}_1 \in \mathbb{C}^{\text{green}}$ and $\mathcal{C}_2 \in \mathbb{C}^{\text{blue}}$,
- (3) for all $\mathcal{C}_1 \in \mathbb{C}^{\text{red}}$ there exists $\mathcal{C}_2 \in \mathbb{C}^{\text{blue}}$ such that $\mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$.

Note that an RGB-color scheme is fully specified by the set \mathbb{C}^{blue} . Indeed, any color not in \mathbb{C}^{blue} is unambiguously placed within \mathbb{C}^{red} or \mathbb{C}^{green} , depending on whether it is disjoint from some color in \mathbb{C}^{blue} . For example, Figure 7 shows several RGB-color schemes for the relational signature $\mathbb{R} = \{a, b, c\}$.




Figure 7: Examples of RGB color schemes

Because a fusion operation only joins element with disjoint colors, blue elements can only be joined with red elements, green elements can be joined with green or red elements, whereas red elements can be joined with elements of any other color, provided that they are disjoint subsets of \mathbb{R} . Moreover, a fusion operation can always join a pair of elements with disjoint colors:

Lemma 3.16. *Let $S_1 = (U_1, \sigma_1)$ and $S_2 = (U_2, \sigma_2)$ be disjoint structures. Let $u_1 \in \text{supp}(\sigma_1)$, $u_2 \in \text{supp}(\sigma_2)$ be elements such that $\mathcal{C}_{S_1}(u_1) \cap \mathcal{C}_{S_2}(u_2) = \emptyset$. Then, the equivalence relation on $U_1 \cup U_2$ generated by (u_1, u_2) is compatible with $S_1 \bullet S_2$.*

Proof. We denote by \approx the relation $\{(u_1, u_2)\}^=$ in the following. Let $r \in \mathbb{R}$ be a relation and let $\langle u_{1,1}, u_{1,2}, \dots, u_{1,\#r} \rangle \in \sigma_1(r)$, $\langle u_{2,2}, u_{2,2}, \dots, u_{2,\#r} \rangle \in \sigma_2(r)$ be distinct tuples. If for some index $i \in [1..\#r]$ either $u_{1,i} \neq u_1$ or $u_{2,i} \neq u_2$ then $u_{1,i} \not\approx u_{2,i}$, by the definition of \approx . Otherwise, if for all indices $i \in [1..\#r]$ both $u_{1,i} = u_1$ and $u_{2,i} = u_2$ then $r \in \mathcal{C}_{S_1}(u_1)$ and $r \in \mathcal{C}_{S_2}(u_2)$. This implies $\mathcal{C}_{S_1}(u_1) \cap \mathcal{C}_{S_2}(u_2) \neq \emptyset$ and contradicts the hypothesis about the choice of u_1, u_2 . Therefore, no tuples from S_1 and S_2 respectively are merged by the fusion. Finally, it is also an easy check that no tuples from S_1 (resp. S_2) are merged, because when restricted to S_1 (resp. S_2) the equivalence \approx becomes the identity. \square

The first ingredient of a decidable condition, equivalent to the treewidth boundedness of a set $\text{EF}^*(\llbracket A \rrbracket_{\Delta}^c)$, is conformance with an RGB color scheme, defined below:

Definition 3.17. A set \mathcal{S} of structures *conforms to* $(\mathbb{C}^{red}, \mathbb{C}^{green}, \mathbb{C}^{blue})$ if and only if:

- (1) for all structures $S = (U, \sigma) \in \mathcal{S}$, if $\mathcal{C}_S(u) \in \mathbb{C}^{red}$, for some element $u \in \text{supp}(\sigma)$, then $\mathcal{C}_S(u') \in \mathbb{C}^{blue}$, for all other elements $u' \in \text{supp}(\sigma) \setminus \{u\}$, and
- (2) $S^\# \cap \mathbb{C}^{green} \subseteq \{\mathcal{C}, \mathcal{C} \mid \mathcal{C} \in \mathbb{C}^{green}\}$, for all structures $S \in \text{EF}^*(\mathcal{S})$.

Moreover, a structure $S \in \mathcal{S}$ is said to be of either type:

- \triangleright **R** if $S^\# \in \text{mpow}(\mathbb{C}^{blue} \cup \mathbb{C}^{red})$ and $\text{card}(S^\# \cap \mathbb{C}^{red}) = 1$,
- \triangleright **G** if $S^\# \in \text{mpow}(\mathbb{C}^{blue} \cup \mathbb{C}^{green})$ and $\text{card}(S^\# \cap \mathbb{C}^{green}) > 0$, and
- \triangleright **B** if $S^\# \in \text{mpow}(\mathbb{C}^{blue})$.

Conformance to some RGB color scheme is the key to bounding the treewidth of the sets of structures obtained by external fusion of a treewidth bounded set of structures.

An equivalence relation \approx is said to be *generated* by a set of pairs $(u_1, v_1), \dots, (u_k, v_k)$ if it is the least equivalence relation, such that $u_i \approx v_i$, for all $i \in [1..k]$. Furthermore, we say that \approx is k -generated if k is the minimal cardinality of a set of pairs that generates \approx .

Lemma 3.18. *Let \mathcal{S} be a treewidth bounded set of structures conforming to an RGB color scheme. Then, for any structure $S \in \text{EF}^*(\mathcal{S})$, the following hold:*

- (1) S is of type either R, G or B,
- (2) if $S = (S_1 \bullet S_2)_{/\approx}$ for some $S_1, S_2 \in \text{EF}^*(\mathcal{S})$ then exactly one of the following hold:
 - (a) \approx is 1-generated, or
 - (b) \approx is 2-generated and either (i) S_1, S_2 are of type R, or (ii) S_1, S_2 are of type G and $\text{card}(S_1^\# \cap \mathbb{C}^{\text{green}}) = \text{card}(S_2^\# \cap \mathbb{C}^{\text{green}}) = 2$
- (3) $\text{tw}(S) \leq \text{tw}(\mathcal{S}) + 1$.

Proof. (1) By induction on the derivation of $S \in \text{EF}^*(\mathcal{S})$ from \mathcal{S} . Table 1 summarizes the possible types of $\text{EF}(S_1, S_2)$ on structures S_1 and S_2 of types R, G or B, respectively.

$\text{EF}(S_1, S_2)$	S_2 of R type	S_2 of G type	S_2 of B type
S_1 of R type	R, G, B	G, B	B
S_1 of G type	G, B	G, B	\perp
S_1 of B type	B	\perp	\perp

Table 1: The types of structures obtained by external fusion (\perp means none)

(2) We distinguish two cases:

- ▷ S_1 is of type R: If S_2 is of type B or G then S_1 and S_2 can be fused only by equivalences \approx generated by a single pair, that contains the element from the support of S_1 with color in \mathbb{C}^{red} , thus matching the case (1) from the statement. Else, if S_2 is of type R then S_1 and S_2 can be fused by equivalences generated by at most two pairs, each containing an element with color from \mathbb{C}^{red} , from either S_1 or S_2 , thus matching the case (2(b)i) from the statement.
- ▷ S_1, S_2 are both of type G: By contradiction, assume they can be fused by an equivalence \approx generated by three pairs of elements $(u_{1i}, u_{2i})_{i=1,2,3}$. Let $G_{1i} = \mathcal{C}_{S_1}(u_{1i})$, $G_{2i} = \mathcal{C}_{S_2}(u_{2i})$ be the colors from $\mathbb{C}^{\text{green}}$ of the matching elements in the two structures, for $i = 1, 2, 3$. Then, we can construct structures using S_1 and S_2 where any of these colors repeat strictly more than twice, henceforth, contradicting the conformance property to the RGB color scheme. The principle of the construction is depicted in Figure 8. Finally, note that the construction depicted in Figure 8 fuse actually only pairs of colors (G_{1i}, G_{2i}) for $i = 1, 2$. Henceforth, the conformance property is also contradicted if S_1 and S_2 can be fused by a 2-generated equivalence relation \approx , such that the support of either S_1 or S_2 contains more than three elements with colors in $\mathbb{C}^{\text{green}}$.

(3) The previous point shows that, under the hypotheses of the lemma, every structure in $\text{EF}^*(\mathcal{S})$ is constructed by external fusion with matchings generated by one or two pairs of elements. This result can be actually refined, i.e., we can consider only external fusions where the 1-generated matchings are applied before the 2-generated matchings. That is, assume $S = ((S_1 \bullet S_2)_{/\approx_2} \bullet S_3)_{/\approx_1}$ where \approx_1, \approx_2 are 1-, resp. 2-generated and $S_1, S_2, S_3 \in \text{EF}^*(\mathcal{S})$. Without loss of generality, assume moreover, \approx_1 is matching some element




Figure 8: External fusion of G structures by 3-generated matchings

of S_3 with an element of S_1 (the other case is symmetric). Then, we can find 1-, resp. 2-generated matchings \approx'_1, \approx'_2 such that $S = ((S_1 \bullet S_3)_{/\approx'_1} \bullet S_2)_{/\approx'_2}$. That is, first fuse S_1 and S_3 by a single matching pair, then fuse the result with S_2 by two matching pairs. Therefore, we can w.l.o.g. assume in the following that $\mathbf{EF}^*(S) = \mathbf{EF}_2^*(\mathbf{EF}_1^*(S))$, where \mathbf{EF}_k^* denotes the external fusion using only k -generated matchings.

Given a tree decomposition T for a structure $S = (U, \sigma)$ and an equivalence relation $\approx \subseteq U \times U$, we denote by $T_{/\approx}$ the tree decomposition of the quotient structure $S_{/\approx}$, obtained by the relabeling of elements u in the bags of T by their representatives $[u]_{\approx}$. We prove the following facts:

Fact 3.19. $\text{tw}(\mathbf{EF}_1^*(S)) \leq \text{tw}(S)$.

Proof. By induction on the derivation of $S \in \mathbf{EF}_1^*(S)$ from S .

Base case: Immediate, as for any $S \in S$ we have $\text{tw}(S) \leq \text{tw}(S)$.

Induction step: Consider $S = (S_1 \bullet S_2)_{/\approx}$ where $\approx = \{(u_1, u_2)\}^=$. Let T_1, T_2 be tree decompositions of respectively S_1, S_2 such that $\text{wd}(T_1), \text{wd}(T_2) \leq \text{tw}(S)$. We first build a tree decomposition T_{12} of $S_1 \bullet S_2$ by

- ▷ transforming T_2 into T'_2 by reversing edges such that a node n_2 containing u_2 in T_2 becomes the root of T'_2 and
- ▷ linking the root of T'_2 to a node n_1 of T_1 containing u_1 .

This ensures that $T_{12/\approx}$ is a valid tree decomposition for S , and moreover $\text{wd}(T_{12/\approx}) = \text{wd}(T_{12}) = \max(\text{wd}(T_1), \text{wd}(T_2)) \leq \text{tw}(S)$. \square

Fact 3.20. Let $S = (U, \sigma) \in \mathbf{EF}_2^*(\mathbf{EF}_1^*(S))$ be a structure. Then, one of the following holds:

- (A) $\text{tw}(S) \leq \text{tw}(S)$,
- (B) S is of type B, $\text{tw}(S) \leq \text{tw}(S) + 1$,
- (C) S is of type G and there exists $T = (\mathcal{N}, \mathcal{E}, r, \lambda)$ a tree decomposition of S such that $\text{wd}(T) \leq \text{tw}(S) + 1$ and $\mathcal{C}_S(u) \in \mathbb{C}^{\text{green}}$ implies $u \in \lambda(r)$, for all $u \in \text{supp}(\sigma)$.

Proof. By induction on the derivation of $S \in \mathbf{EF}_2^*(\mathbf{EF}_1^*(S))$ from $\mathbf{EF}_1^*(S)$.

Base case: Immediate, as we already shown $\text{tw}(S) \leq \text{tw}(S)$, that is, (A) for any $S \in \mathbf{EF}_1^*(S)$.

Induction step: Consider $S = (S_1 \bullet S_2)_{/\approx}$ where $\approx = \{(u_{11}, u_{21}), (u_{12}, u_{22})\}^=$. Since \approx is 2-generated, we know from the previous point (2) that either:

- ▷ S_1, S_2 are of type R: From the induction hypothesis on S_1 and S_2 it follows that both must satisfy (A), hence $\text{tw}(S_1) \leq \text{tw}(S), \text{tw}(S_2) \leq \text{tw}(S)$ respectively. We are therefore in the situation of composing two structures of type R by a 2-generated matching, hence obtaining the structure S of type B. Without loss of generality consider $\mathcal{C}_{S_1}(u_{11}) \in \mathbb{C}^{\text{red}}$,

$\mathcal{C}_{S_1}(u_{12}) \in \mathbb{C}^{blue}$, $\mathcal{C}_{S_2}(u_{21}) \in \mathbb{C}^{blue}$, $\mathcal{C}_{S_2}(u_{22}) \in \mathbb{C}^{red}$. Let T_1, T_2 be tree decompositions of S_1, S_2 respectively. First, we construct a tree decomposition T_{12} for $S_1 \bullet S_2$ as follows:

- construct T'_1 from T_1 by propagating the element u_{11} to all the nodes,
- construct T'_2 from T_2 by propagating the element u_{22} to all the nodes and then reverting the edges such that a node n_2 containing the element u_{21} becomes the root,
- link the root node of T'_2 to a node n_1 of T'_1 containing the element u_{12} .

This ensures that $T_{12/\approx}$ is a valid tree decomposition for S , and moreover $\text{wd}(T_{12/\approx}) = \text{wd}(T_{12}) = \max(\text{wd}(T'_1), \text{wd}(T'_2)) = \max(\text{wd}(T_1) + 1, \text{wd}(T_2) + 1) \leq \text{tw}(\mathcal{S}) + 1$. This completes the proof that S satisfies condition (B).

- ▷ S_1, S_2 are of type G and $\text{card}(S_1^\# \sqcap \mathbb{C}^{green}) = \text{card}(S_2^\# \sqcap \mathbb{C}^{green}) = 2$: First, let us observe that the structure S is either of type B or G. According to the induction hypothesis, we consider the following two cases:

- both S_1 and S_2 satisfy the condition (C), namely there exist the tree decompositions T_1, T_2 of width at most $\text{tw}(\mathcal{S}) + 1$ such that moreover all the elements with colors in \mathbb{C}^{green} are located at their root nodes. Then, we can construct a tree decomposition T_{12} for $S_1 \bullet S_2$ by simply linking the root of T_2 as a child to the root of T_1 . We obtain that $T_{12/\approx}$ is a valid decomposition for S and satisfies $\text{wd}(T_{12/\approx}) = \text{wd}(T_{12}) = \max(\text{wd}(T_1), \text{wd}(T_2)) \leq \text{tw}(\mathcal{S}) + 1$. Moreover, if S is of type G observe that all the elements with colors in \mathbb{C}^{green} are located at the root node of $T_{12/\approx}$. Therefore, in any case, the structure S satisfies either (B) or (C).
- either one or both of S_1 or S_2 satisfy the condition (A). Without loss of generality consider $\text{tw}(S_1) \leq \text{tw}(\mathcal{S})$. We know however that S_1 satisfies $\text{card}(S_1^\# \sqcap \mathbb{C}^{green}) = 2$, that is, it has exactly two elements with colors in \mathbb{C}^{green} . But then, we can show that S_1 satisfies the condition (C) as well. That is, consider a tree decomposition T_1 for S_1 such that $\text{wd}(T_1) \leq \text{tw}(\mathcal{S})$. Let u_{11}, u_{12} be the two elements with colors in \mathbb{C}^{green} . We can build a tree decomposition T'_1 fulfilling (C) by first propagating the element u_{11} to all the nodes of T_1 and then reverting the edges such that some node containing the element u_{12} becomes the root. Obviously, $\text{wd}(T'_1) = \text{wd}(T_1) + 1 \leq \text{tw}(\mathcal{S}) + 1$ and all the elements with color in \mathbb{C}^{green} are located at the root node. We can proceed similarly with S_2 . Then, the proof is completed as in the first case. \square

This completes the proof of the point (3) from the statement. \square

Moreover, conformance with RGB schemes allow us to infer a bound on the set obtained by applying external fusion to a treewidth bounded set of structures:

Lemma 3.21. *Let \mathcal{S} be a treewidth bounded set of structures, that conforms to an RGB color scheme. Then, we have $\text{tw}(\text{IF}(\text{EF}^*(\mathcal{S}))) \leq \text{tw}(\mathcal{S}) + 1$.*

Proof. $\text{EF}^*(\mathcal{S})$ is treewidth bounded as a direct consequence of Lemma 3.18, that establishes the bounds for every type of structure from $\text{EF}^*(\mathcal{S})$. Moreover, $\text{IF}(\text{EF}^*(\mathcal{S}))$ is treewidth-bounded because, using the tree decompositions T constructed for structures S in $\text{EF}^*(\mathcal{S})$ one obtains tree decomposition T' and treewidth bounds for any structures $S' = S_{/\approx}$ obtained by internal fusion, as follows:

- ▷ if S is of type R then by internal fusion one glues the unique element u_1 with color in \mathbb{C}^{red} to some other element in the structure. We know, from the inductive property used in the proof of Lemma 3.18 (3) that if S is of type R then it must satisfies condition (A) of Fact 3.20 that is, $\text{tw}(S) \leq \text{tw}(\mathcal{S})$. Therefore, one can construct T' from T by

- replicating u_1 in all nodes and then $T'_{/\approx}$ is a valid tree decomposition for S' . Obviously $\text{wd}(T') \leq \text{wd}(T) + 1 \leq \text{tw}(\mathcal{S}) + 1$.
- ▷ if S is of type **G** then by internal fusion one glue elements with color in \mathbb{C}^{green} . As before, using the same inductive property, we know that S satisfies either condition (A) or (C) of Fact 3.20. If S satisfies condition (C), as all elements that could be glued are already present in the root node, $T'_{/\approx}$ is a valid tree decomposition for S' . Obviously, the treewidth bound remains unchanged, that is, at most $\text{tw}(\mathcal{S}) + 1$. If S satisfies condition (A) and two of its nodes u_1, u_2 with colors in \mathbb{C}^{green} can be fused, then two copies of S , respectively S', S'' can also be fused via the two-pair matching $(u'_1, u''_2), (u''_1, u'_2)$. Henceforth, u_1 and u_2 must be the unique elements with colors in \mathbb{C}^{green} , otherwise contradicting the point (2) of Lemma 3.18. But then, from the tree decomposition T of S such that $\text{wd}(T) \leq \text{tw}(\mathcal{S})$ we can construct the tree decomposition T' by simply propagating u_2 to all other nodes in T and hence, preserving the bound of $\text{tw}(\mathcal{S}) + 1$.
 - ▷ if S is of type **B** then no non-trivial internal fusion exists, and obviously, the treewidth bound remains unchanged. \square

3.4. Connected structures. We shall check conformance with RGB schemes for sets of *maximally connected* structures, defined below:

Definition 3.22. A *path* from u to v in a structure $S = (U, \sigma)$ is a finite sequence of tuples:

$$\langle u_{1,1}, \dots, u_{1,n_1} \rangle \in \sigma(r_1), \dots, \langle u_{k,1}, \dots, u_{k,n_k} \rangle \in \sigma(r_k), \text{ for some } r_1, \dots, r_k \in \mathbb{R}$$

where $u \in \{u_{1,1}, \dots, u_{1,n_1}\}, v \in \{u_{k,1}, \dots, u_{k,n_k}\}$ and $\{u_{i,1}, \dots, u_{i,n_i}\} \cap \{u_{i+1,1}, \dots, u_{i+1,n_{i+1}}\} \neq \emptyset$, for all $i \in [1..k-1]$. The structure S is *connected* iff there exists a path from u to v , for all $u, v \in \text{supp}(\sigma)$.

Definition 3.23. A structure S_1 is a *maximal connected substructure* of another structure S_2 , denoted $S_1 \sqsubseteq^{mc} S_2$, iff (i) $S_1 \sqsubseteq S_2$ (see Definition 3.7), (ii) S_1 is connected, and (iii) for any connected substructure $S'_1 \sqsubseteq S_2$, we have $S_1 \sqsubseteq S'_1$ only if $S_1 = S'_1$. For a structure S we denote by $\text{split}(S) \stackrel{\text{def}}{=} \{S' \mid S' \sqsubseteq^{mc} S\}$ the set of maximally connected substructures, lifted to sets of structures \mathcal{S} as $\text{split}(\mathcal{S}) \stackrel{\text{def}}{=} \cup_{S \in \mathcal{S}} \text{split}(S)$.

Note that $\text{tw}(\mathcal{S}) = \text{tw}(\text{split}(\mathcal{S}))$ for any set of structures \mathcal{S} . The next lemma shows that both internal and external fusions preserve maximally connected substructures:

Lemma 3.24. *For each set \mathcal{S} of structures, the following hold:*

- (1) $\text{split}(\text{EF}^*(\mathcal{S})) = \text{EF}^*(\text{split}(\mathcal{S}))$, and
- (2) $\text{split}(\text{IEF}^*(\mathcal{S})) = \text{IEF}^*(\text{split}(\mathcal{S}))$.

Proof. For space reasons, this proof is given in subsection B.1. \square

The core of our algorithm is a decidable equivalent condition for the treewidth boundedness of a set obtained by applying external fusion to a set of connected structures. This condition is that, in any of the structures produced by external fusion, there is no way of connecting six elements u_1, v_1, w_1 and u_2, v_2, w_2 , labeled with non-disjoint colors \mathcal{C}_1 and \mathcal{C}_2 , respectively. Assume that this condition is violated by some structures S_1 and S_2 with elements u_1, v_1, w_1 and u_2, v_2, w_2 , such that $\mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$. In this case, Figure 9 depicts the construction of a structure with an $n \times n$ square grid minor, of treewidth at least n , for any $n \geq 1$. Intuitively, $\mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ allows to glue the elements u_1 with u_2 , v_1 with v_2 and w_1 with w_2 , respectively.

Lemma 3.25. *The following are equivalent, for any treewidth-bounded set \mathcal{S} of structures:*

- (1) $\mathbf{EF}^*(\mathcal{S})$ is treewidth bounded,
- (2) $\{\{\mathcal{C}_1, \mathcal{C}_1, \mathcal{C}_1\}, \{\mathcal{C}_2, \mathcal{C}_2, \mathcal{C}_2\}\} \in (\mathbf{EF}^*(\mathbf{split}(\mathcal{S})))^{\#3}$ implies $\mathcal{C}_1 \cap \mathcal{C}_2 \neq \emptyset$, for all $\mathcal{C}_1, \mathcal{C}_2 \in \mathbb{C}$,
- (3) $\mathbf{split}(\mathcal{S})$ conforms to some RGB color scheme.

Proof. “(1) \Rightarrow (2)” If $\mathbf{EF}^*(\mathcal{S})$ is treewidth-bounded then $\mathbf{split}(\mathbf{EF}^*(\mathcal{S}))$ is treewidth-bounded. Using Lemma 3.24 the later set is equal to $\mathbf{EF}^*(\mathbf{split}(\mathcal{S}))$ and henceforth treewidth bounded as well. By contradiction, assume that (2) does not hold. Then, there exist colors $\mathcal{C}_1, \mathcal{C}_2 \in \mathbb{C}$, connected structures $S_1, S_2 \in \mathbf{EF}^*(\mathbf{split}(\mathcal{S}))$ such that $\{\{\mathcal{C}_1, \mathcal{C}_1, \mathcal{C}_1\} \in S_1^\sharp, \{\mathcal{C}_2, \mathcal{C}_2, \mathcal{C}_2\} \in S_2^\sharp$ and moreover $\mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$. We shall use S_1 and S_2 to build infinitely many connected structures containing arbitrarily large square grid minors. First, construct the connected structure $S_{12} \in \mathbf{EF}^*(\mathbf{split}(\mathcal{S}))$ by fusing one pair (u_1, u_2) with colors $\mathcal{C}_1, \mathcal{C}_2$. Let v_1, w_1 resp. v_2, w_2 be the remaining distinct elements of S_{12} with color $\mathcal{C}_1, \mathcal{C}_2$ from respectively S_1, S_2 . For arbitrarily positive n , consider $n \times n$ disjoint copies $(S_{12}^{i,j})_{i,j=1,n}$ of S_{12} . Let $\approx^{1,j}$ be $\{(v_1^{1,j}, v_2^{1,j-1})\} =$, $\approx^{i,1}$ be $\{(w_2^{i,1}, w_1^{i-1,1})\} =$, $\approx^{i,j}$ be $\{(v_1^{i,j}, v_2^{i,j-1}), (w_2^{i,j}, w_1^{i-1,j})\} =$ for all $i, j = 2, n$. Second, construct the grid-like connected structure $X^{n,n} \in \mathbf{EF}^*(\mathbf{split}(\mathcal{S}))$:

$$X^{n,n} = (\dots((S_{12}^{1,1} \bullet S_{12}^{1,2})_{/\approx^{1,2}} \bullet S_{12}^{2,1})_{/\approx^{2,1}} \bullet \dots \bullet S_{12}^{i,j})_{/\approx^{i,j}} \bullet \dots \bullet S_{12}^{n,n})_{/\approx^{n,n}}$$

where structures $S_{12}^{i,j}$ are added to the fusion in increasing order of $i + j$. The construction is illustrated in Figure 9. We can show that $X^{n,n}$ contains an $n \times n$ square grid minor. Finally, as n can be taken arbitrarily large, we conclude that $\mathbf{EF}^*(\mathbf{split}(\mathcal{S}))$ contains structures with arbitrarily large square grid minors, it is not treewidth-bounded, contradicting (1).




Figure 9: The principle of grid construction

“(2) \Rightarrow (3)” We define a RGB color scheme by selecting:

$$\mathbb{C}^{blue} = \{\mathcal{C} \in \mathbb{C} \mid \{\{\mathcal{C}, \mathcal{C}, \mathcal{C}\} \in (\mathbf{EF}^*(\mathbf{split}(\mathcal{S})))^{\#3}\}$$

Since (2) holds, this is a valid definition for \mathbb{C}^{blue} , which induces a partitioning of the remaining colors into \mathbb{C}^{green} and \mathbb{C}^{red} . We show that $\mathbf{split}(\mathcal{S})$ is conforming to this RGB partitioning, by checking the two points of Definition 3.17:

- (1) Let $S \in \mathbf{split}(\mathcal{S})$ and prove that for any two colors $\mathcal{C}_1, \mathcal{C}_2 \in \mathbb{C}$, if $\{\{\mathcal{C}_1, \mathcal{C}_2\} \subseteq S^\sharp$ and $\mathcal{C}_1 \in \mathbb{C}^{red}$ then $\mathcal{C}_2 \in \mathbb{C}^{blue}$. Since $\mathcal{C}_1 \in \mathbb{C}^{red}$, there must exists a color $\mathcal{C}'_1 \in \mathbb{C}^{blue}$, such that $\mathcal{C}_1 \cap \mathcal{C}'_1 = \emptyset$, by Definition 3.15. By the definition of \mathbb{C}^{blue} in our RGB-color

scheme, this further implies $\{\{C'_1, C'_1, C'_1\} \in (\mathbf{EF}^*(\mathbf{split}(\mathcal{S})))^{\#3}$. Henceforth, there exists a structure $S' \in \mathbf{EF}^*(\mathbf{split}(\mathcal{S}))$ such that $\{\{C'_1, C'_1, C'_1\} \subseteq S'^{\#}$. We can now use S' and three disjoint copies of S to build a new structure S'' by gluing progressively, each one of the three elements of color C'_1 in S' to the element of color C_1 of S . Then, by construction, the structure S'' will also contain three elements of color C_2 , one from each disjoint copy of S . Therefore, $\{\{C_2, C_2, C_2\} \in S''^{\#}$ and because $S'' \in \mathbf{EF}^*(\mathbf{split}(\mathcal{S}))$ this implies $\{\{C_2, C_2, C_2\} \in (\mathbf{EF}^*(\mathbf{split}(\mathcal{S})))^{\#3}$ and therefore $C_2 \in \mathbb{C}^{blue}$.

- (2) By contradiction, let $S \in \mathbf{EF}^*(\mathbf{split}(\mathcal{S}))$ be such that $S^{\#} \cap \mathbb{C}^{green} \not\subseteq \{\{C, C \mid C \in \mathbb{C}^{green}\}\}$. Then there exists $C' \in (S^{\#} \cap \mathbb{C}^{green}) \setminus \{\{C, C \mid C \in \mathbb{C}^{green}\}\}$, i.e., $C' \in \mathbb{C}^{green}$ and $\{\{C', C', C'\} \subseteq S^{\#}$. The latter implies $\{\{C', C', C'\} \in S^{\#3} \subseteq (\mathbf{EF}^*(\mathbf{split}(\mathcal{S})))^{\#3}$. But this implies $C' \in \mathbb{C}^{blue}$ according to the definition of the RGB color scheme, contradicting $C' \in \mathbb{C}^{green}$.

“(3) \Rightarrow (1)” By Lemma 3.21, $\mathbf{EF}^*(\mathbf{split}(\mathcal{S}))$ is treewidth bounded. Then, by Lemma 3.24, $\mathbf{split}(\mathbf{EF}^*(\mathcal{S}))$ is treewidth bounded, thus $\mathbf{EF}^*(\mathcal{S})$ is treewidth bounded. \square

A first consequence of this result is the equivalence between the treewidth boundedness of the sets $\llbracket \mathbf{A} \rrbracket_{\Delta}$ and $\mathbf{EF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$. The following lemma establishes this equivalence, by providing the missing direction to Lemma 3.12:

Lemma 3.26. *Given a SID Δ and a nullary predicate symbol \mathbf{A} , $\mathbf{EF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$ is treewidth bounded only if $\mathbf{IEF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$ is treewidth bounded.*

Proof. For any set of structures \mathcal{S} , we have $\mathbf{IEF}^*(\mathcal{S}) = \mathbf{IF}(\mathbf{EF}^*(\mathcal{S}))$, because the operations of internal and external fusion commute, namely $\mathbf{EF}(\mathbf{IF}(S_1), S_2) \subseteq \mathbf{IF}(\mathbf{EF}(S_1), S_2)$, for any structures S_1, S_2 . By Lemma 3.25, $\mathbf{EF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$ is treewidth-bounded only if $\mathbf{split}(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$ conforms to an RGB color scheme. Then, $\mathbf{IF}(\mathbf{EF}^*(\mathbf{split}(\llbracket \mathbf{A} \rrbracket_{\Delta}^c))) = \mathbf{IEF}^*(\mathbf{split}(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)) = \mathbf{split}(\mathbf{IEF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c))$ is treewidth-bounded, by Lemma 3.21 and Lemma 3.24. Thus, $\mathbf{IEF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$ is treewidth bounded. \square

Our algorithm that decides the treewidth boundedness of a set $\mathbf{EF}^*(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$ checks whether the set $(\mathbf{EF}^*(\mathbf{split}(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)))^{\#3}$ meets condition (2) of Lemma 3.25. For this check to be effective, the latter set must be constructed in finite time from the description of Δ and \mathbf{A} , provided as input. This construction proceeds in three consecutive stages. First, we show that, for any set \mathcal{S} of structures, the k -color abstraction $(\#^k \mathbf{EF}^*(\mathcal{S}))$ can be built from $\mathcal{S}^{\#k}$ by an effectively computable abstract operator. Second, we build a SID Γ and a nullary predicate \mathbf{P} , such that $\llbracket \mathbf{P} \rrbracket_{\Gamma}^c = \mathbf{split}(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$, i.e., it encodes the set of maximally connected substructures from some canonical Δ -model of \mathbf{A} (subsection 3.6). Finally, we compute the k -multiset color abstraction of $\llbracket \mathbf{P} \rrbracket_{\Gamma}^c$ (subsection 3.7). We end this section with a proof of decidability for the treewidth boundedness problem for expandable SIDs (Theorem 3.33).

3.5. Color abstractions of externally fused sets. We describe now the effective construction of a k -multiset abstraction $(\mathbf{EF}^*(\mathcal{S}))^{\#k}$ from the abstraction $\mathcal{S}^{\#k}$ of a set \mathcal{S} of structures, for a given integer $k \geq 1$. First, as we are interested only in k -multisets color abstractions, we can restrict external fusion to 1-generated matchings, with no loss of generality.

Definition 3.27. The *single-pair external fusion* of disjoint structures $S_1 = (U_1, \sigma_1)$ and $S_2 = (U_2, \sigma_2)$ is the external fusion (Definition 3.11) induced by 1-generated matchings. We denote by $\mathbf{EF}_1(S_1, S_2)$ the set of structures obtained by single-pair external fusion of S_1 and S_2 . For a set of structures \mathcal{S} , we denote by $\mathbf{EF}_1^*(\mathcal{S})$ the closure of \mathcal{S} under single-pair external fusions.

In general, the single-pair external fusion is strictly less expressive than external fusion, yet it produces the same k -multiset color abstractions:

Lemma 3.28. $(\mathbf{EF}^*(\mathcal{S}))^{\sharp k} = (\mathbf{EF}_1^*(\mathcal{S}))^{\sharp k}$ for any set \mathcal{S} of structures and integer $k \geq 1$.

Proof. “ $(\mathbf{EF}_1^*(\mathcal{S}))^{\sharp k} \subseteq (\mathbf{EF}^*(\mathcal{S}))^{\sharp k}$ ” This direction follows directly from $\mathbf{EF}_1^*(\mathcal{S}) \subseteq \mathbf{EF}^*(\mathcal{S})$. “ $(\mathbf{EF}^*(\mathcal{S}))^{\sharp k} \subseteq (\mathbf{EF}_1^*(\mathcal{S}))^{\sharp k}$ ” We prove the stronger property:

$$\forall \mathbf{S} \in \mathbf{EF}^*(\mathcal{S}). \exists \mathbf{S}' \in \mathbf{EF}_1^*(\mathcal{S}). \mathbf{S}^\sharp \subseteq \mathbf{S}'^\sharp$$

By induction on the derivation of $\mathbf{S} \in \mathbf{EF}^*(\mathcal{S})$ from \mathcal{S} .

Base case: Assume $\mathbf{S} \in \mathcal{S}$. Then $\mathbf{S}' = \mathbf{S}$ satisfies the property.

Induction step: Assume $\mathbf{S} = (\mathbf{S}_1 \bullet \mathbf{S}_2)_{/\approx}$ for some $\mathbf{S}_1, \mathbf{S}_2 \in \mathbf{EF}^*(\mathcal{S})$ and some equivalence relation \approx , defined as $\{(u_{1i}, u_{2i}) \mid i \in [1..n]\}^\#$, that conforms to the requirements of external fusion for $\mathbf{S}_1, \mathbf{S}_2$. Let $\mathcal{C}_{1i} = \mathcal{C}_{\mathbf{S}_1}(u_{1i})$, $\mathcal{C}_{2i} = \mathcal{C}_{\mathbf{S}_2}(u_{2i})$, for all $i \in [1..n]$. According to the definition of external fusion, $\mathbf{S} = (\mathbf{S}_1 \bullet \mathbf{S}_2)_{/\approx}$ implies $\mathcal{C}_{1i} \cap \mathcal{C}_{2i} = \emptyset$ and moreover:

$$\mathbf{S}^\sharp = \{\{\mathcal{C}_{1i} \cup \mathcal{C}_{2i} \mid i \in [1..n]\}\} \cup (\mathbf{S}_1^\sharp \setminus \{\{\mathcal{C}_{1i} \mid i \in [1..n]\}\}) \cup (\mathbf{S}_2^\sharp \setminus \{\{\mathcal{C}_{2i} \mid i \in [1..n]\}\})$$

By induction hypothesis, for $\mathbf{S}_1, \mathbf{S}_2$ there exists $\mathbf{S}'_1, \mathbf{S}'_2 \in \mathbf{EF}_1^*(\mathcal{S})$ such that $\mathbf{S}_1^\sharp \subseteq \mathbf{S}'_1{}^\sharp$, $\mathbf{S}_2^\sharp \subseteq \mathbf{S}'_2{}^\sharp$. We use \mathbf{S}'_1 and n disjoint copies $\mathbf{S}'_{2,1}, \dots, \mathbf{S}'_{2,n}$ of \mathbf{S}'_2 to construct \mathbf{S}' with the required property. The idea is that, for every pair $u_{1i} \approx u_{2i}$, we fuse some element u'_{1i} with color \mathcal{C}_{1i} from \mathbf{S}'_1 with some element u'_{2i} with color \mathcal{C}_{2i} from $\mathbf{S}'_{2,i}$. Such elements always exist, because $\mathbf{S}_1^\sharp \subseteq \mathbf{S}'_1{}^\sharp$, $\mathbf{S}_2^\sharp \subseteq \mathbf{S}'_2{}^\sharp$. Therefore, consider the equivalence relations $\approx'_i = (u'_{1i}, u'_{2i})^\#$ for some pair of elements as above, for all $i \in [1..n]$ and define:

$$\mathbf{S}' = (\dots ((\mathbf{S}'_1 \bullet \mathbf{S}'_{2,1})_{/\approx'_1} \bullet \mathbf{S}'_{2,2})_{/\approx'_2} \bullet \dots \bullet \mathbf{S}'_{2,n})_{/\approx'_n}$$

Then $\mathbf{S}' \in \mathbf{EF}_1^*(\mathcal{S})$ and, moreover, we have $\mathbf{S}^\sharp \subseteq \mathbf{S}'^\sharp$, because:

$$\mathbf{S}'^\sharp = \{\{\mathcal{C}_{1i} \cup \mathcal{C}_{2i} \mid i \in [1..n]\}\} \cup (\mathbf{S}'_1{}^\sharp \setminus \{\{\mathcal{C}_{1i} \mid i \in [1..n]\}\}) \cup \bigcup_{i \in [1..n]} (\mathbf{S}'_{2,i}{}^\sharp \setminus \{\{\mathcal{C}_{2i}\}\}) \quad \square$$

Second, the closure $(\mathbf{EF}_1^*(\mathcal{S}))^{\sharp k}$ can be computed by a least fixpoint iteration of an abstract operation on the domain of k -multiset color abstractions. As the later domain is finite, this fixpoint computation is guaranteed to terminate.

Definition 3.29. The *single-pair multiset fusion* is defined below, for $M_1, M_2 \in \text{mpow}(\mathbb{C})$:

$$\mathbf{ef}_1^{\sharp k}(M_1, M_2) \stackrel{\text{def}}{=} \{M \in \text{mpow}(\mathbb{C}) \mid \exists \mathcal{C}_1 \in M_1. \exists \mathcal{C}_2 \in M_2. \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset, \\ M = \{\{\mathcal{C}_1 \cup \mathcal{C}_2\}\} \cup \bigcup_{i=1,2} (M_i \setminus \{\{\mathcal{C}_i\}\})\}$$

Given an integer $k \geq 1$, the *single-pair k -multiset fusion* is defined for $M_1, M_2 \in \text{mpow}(\mathbb{C})$, such that $\text{card}(M_1) \leq k$ and $\text{card}(M_2) \leq k$:

$$\mathbf{ef}_1^{\sharp k}(M_1, M_2) \stackrel{\text{def}}{=} \{M \mid \exists M' \in \mathbf{ef}_1^{\sharp k}(M_1, M_2). M \subseteq M', \text{card}(M) \leq k\}$$

For a set \mathcal{M} of multisets (resp. k -multisets) of colors, let $\mathbf{ef}_1^{\sharp k}(\mathcal{M})$ (resp. $\mathbf{ef}_1^{\sharp k*}(\mathcal{M})$) be the closure of \mathcal{M} under taking single-pair fusion on multisets (resp. k -multisets).

Lemma 3.30. $(\mathbf{EF}_1^*(\mathcal{S}))^{\sharp k} = \mathbf{ef}_1^{\sharp k*}(\mathcal{S}^{\sharp k})$, for any set \mathcal{S} of structures and integer $k \geq 1$.

Proof. Abusing notation, we write $M^{\sharp k} \stackrel{\text{def}}{=} \{M' \mid M' \subseteq M, \text{card}(M') \leq k\}$. Then, we have $(\mathbf{EF}_1^*(\mathcal{S}))^{\sharp k} = ((\mathbf{EF}_1^*(\mathcal{S}))^\sharp)^{\sharp k}$, by Definition 3.14. Using Definition 3.27 of single pair external fusion and Definition 3.29 of single pair fusion of multisets, we can prove that for all

structures S_1, S_2 it holds $(\mathbf{EF}_1(S_1, S_2))^\sharp = \mathbf{ef}_1^\sharp(S_1^\sharp, S_2^\sharp)$. This immediately extends to their respective closure, henceforth, $(\mathbf{EF}_1^*(S))^\sharp = \mathbf{ef}_1^{*\sharp}(S^\sharp)$. Henceforth, we are left with proving that $(\mathbf{ef}_1^{*\sharp}(S^\sharp))^{\sharp k} = \mathbf{ef}_1^{\sharp k*}(S^{\sharp k})$.

“($\mathbf{ef}_1^{*\sharp}(S^\sharp)$) $^{\sharp k} \subseteq \mathbf{ef}_1^{\sharp k*}(S^{\sharp k})$ ” We prove that, for all $M \in \mathbf{ef}_1^{*\sharp}(S^\sharp)$, we have $M^{\sharp k} \subseteq \mathbf{ef}_1^{\sharp k*}(S^{\sharp k})$. The proof goes by induction on the derivation of $M \in \mathbf{ef}_1^{*\sharp}(S^\sharp)$ from S^\sharp .

Base case: Assume $M \in S^\sharp$. Then $M^{\sharp k} \subseteq (S^\sharp)^{\sharp k} = S^{\sharp k} \subseteq \mathbf{ef}_1^{\sharp k*}(S^{\sharp k})$.

Induction step: Assume $M \in \mathbf{ef}_1^\sharp(M_1, M_2)$ for some multisets of colors M_1, M_2 such that $M_1^{\sharp k}, M_2^{\sharp k} \subseteq \mathbf{ef}_1^{\sharp k*}(S^{\sharp k})$. Then, there exists $C_1 \in M_1, C_2 \in M_2$ such that $C_1 \cap C_2 = \emptyset$ and $M = (M_1 \setminus \{\{C_1\}\}) \cup (M_2 \setminus \{\{C_2\}\}) \cup \{\{C_1 \cup C_2\}\}$, by Definition 3.29. Let $M' \in M^{\sharp k}$, that is, $M' \subseteq M$, $\text{card}(M') \leq k$. We distinguish several cases:

- ▷ $M' \subseteq M_1$ (the case $M' \subseteq M_2$ is symmetric): $M' \in M_1^{\sharp k}$, thus $M' \in \mathbf{ef}_1^{\sharp k*}(S^{\sharp k})$.
- ▷ $M' \not\subseteq M_i$, for $i = 1, 2$ and $C_1 \cup C_2 \notin M'$: M' can be partitioned in two nonempty parts $M'_1 \subseteq M_1, M'_2 \subseteq M_2$ such that $M = M'_1 \uplus M'_2$. As both parts are not empty, we have $M'_1 \in M_1^{\sharp k-1}, M'_2 \in M_2^{\sharp k-1}$, thus $(M'_1 \cup \{\{C_1\}\}) \in M_1^{\sharp k}, (M'_2 \cup \{\{C_2\}\}) \in M_2^{\sharp k}$. It is an easy check that $M' \in \mathbf{ef}_1^{\sharp k}((M'_1 \cup \{\{C_1\}\}), (M'_2 \cup \{\{C_2\}\}))$. This implies $M' \in \mathbf{ef}_1^{\sharp k*}(S^{\sharp k})$ as both subterms belong to $\mathbf{ef}_1^{\sharp k*}(S^{\sharp k})$.
- ▷ $M' \not\subseteq M_i$, for $i = 1, 2$ and $C_1 \cup C_2 \in M'$: we proceed as in the previous case but considering a partitioning of $M' \setminus \{\{C_1 \cup C_2\}\}$. We obtain $M' \in \mathbf{ef}_1^{\sharp k*}(S^{\sharp k})$, as well.

“($\mathbf{ef}_1^{\sharp k*}(S^{\sharp k}) \subseteq (\mathbf{ef}_1^{*\sharp}(S^\sharp))^{\sharp k}$ ” We prove that, for all k -multiset $M' \in \mathbf{ef}_1^{\sharp k*}(S^{\sharp k})$, there exists $M \in \mathbf{ef}_1^{*\sharp}(S^\sharp)$, such that $M' \subseteq M$, by induction on the derivation of M' from $S^{\sharp k}$.

Base case: Assume $M' \in S^{\sharp k} = (S^\sharp)^{\sharp k}$. Then, there exists $M \in S^\sharp$ such that $M' \subseteq M$. Obviously, $M \in \mathbf{ef}_1^{*\sharp}(S^\sharp)$.

Induction step: Assume $M' \in \mathbf{ef}_1^{\sharp k}(M'_1, M'_2)$ for some k -multisets of colors $M'_1, M'_2 \in \mathbf{ef}_1^{\sharp k*}(S^{\sharp k})$. By the inductive hypothesis, there exists multisets $M_1, M_2 \in \mathbf{ef}_1^{*\sharp}(S^\sharp)$ such that $M'_1 \subseteq M_1, M'_2 \subseteq M_2$. Since M'_1, M'_2 can be composed such that to obtain (a superset of) the multiset M' , one can use precisely the same pairs of colors to compose M_1, M_2 and henceforth to obtain the multiset $M \in \mathbf{ef}_1^{*\sharp}(S^\sharp)$, which is the superset of M' . \square

3.6. Maximally connected substructures. Since we consider canonical models, we can assume w.l.o.g. that the given SID Δ contains no disequalities (such atoms are trivially unsatisfiable or valid). We represent the set of maximally connected structures $\text{split}(\llbracket \mathbf{A} \rrbracket_\Delta^c)$ as a set of canonical models $\llbracket \mathbf{P} \rrbracket_\Gamma^c$, for a fresh nullary predicate \mathbf{P} and a SID Γ , whose construction is described next.

Given a qpf formula ψ , we define $\text{conn}(\psi) \subseteq \text{fv}(\psi) \times \text{fv}(\psi)$ to be the least equivalence relation such that $(y, z) \in \text{conn}(\psi)$ if $r(x_1, \dots, x_{\#r})$ occurs in ψ and $y, z \in \{x_1, \dots, x_{\#r}\}$, for some $r \in \mathbb{R}$. Intuitively, $\text{conn}(\psi)$ consists of the pairs of free variables of ψ that are connected by a path (see Definition 3.22) in each canonical model of ψ .

Let $\mathbf{B}(y_1, \dots, y_{\#\mathbf{B}})$ be a predicate atom, $J = \{j_1, \dots, j_p\} \subseteq [1.. \#\mathbf{B}]$ be a set of indices ordered as $j_1 \leq \dots \leq j_p$, $\xi \subseteq J \times J$ be an equivalence relation and \mathbf{B}^ξ be a fresh predicate of arity p . In particular, $\#\mathbf{B}^\xi = 0$ if $\xi = \emptyset$ is the empty relation. We define the shorthands:

$$\begin{aligned} \text{fv}_J(\mathbf{B}(y_1, \dots, y_{\#\mathbf{B}})) &\stackrel{\text{def}}{=} \{y_j \mid j \in J\} & \xi(\mathbf{B}(y_1, \dots, y_{\#\mathbf{B}})) &\stackrel{\text{def}}{=} \{(y_j, y_k) \mid (j, k) \in \xi\} \\ & & \mathbf{B}(y_1, \dots, y_{\#\mathbf{B}}) /_\xi &\stackrel{\text{def}}{=} \mathbf{B}^\xi(y_{j_1}, \dots, y_{j_p}) \end{aligned}$$

We build definitions for the predicate atoms $B^\xi(x_{j_1}, \dots, x_{j_p})$, by “narrowing” the definitions of $B(x_1, \dots, x_{\#B})$, respectively. More precisely, every structure $S \in \llbracket B^\xi(y_{j_1}, \dots, y_{j_p}) \rrbracket_\Delta^c$ will correspond to a set of maximally connected substructures in a structure $S' \in \llbracket B(y_1, \dots, y_{\#B}) \rrbracket_\Delta^c$ such that, moreover (i) for every such substructure, there exists an element associated with y_j , for some $j \in J$, and (ii) y_j and y_k are mapped to elements from the same connected substructure of S' if and only if $(j, k) \in \xi$. In other words, the equivalence relation ξ is used to summarize the information about the maximally connected structures from any canonical model of B^ξ .

We describe the construction of Γ next. Consider a rule of Δ of the form:

$$B_0(x_1, \dots, x_{\#B_0}) \leftarrow \exists y_1 \dots \exists y_m \cdot \psi * \bigstar_{i=1}^\ell B_i(z_{i,1}, \dots, z_{i,\#B_i}) \quad (3.1)$$

formulæ ψ' , ψ'' , sets $J_i \uplus \bar{J}_i = [1..\#B_i]$, equivalence relations $\xi_i \subseteq J_i \times J_i$, for all $i \in [1..\ell]$, an equivalence relation $\Xi \subseteq (\{x_1, \dots, x_{\#B_0}\} \cup \{y_1, \dots, y_m\}) \times (\{x_1, \dots, x_{\#B_0}\} \cup \{y_1, \dots, y_m\})$, such that the following hold:

- (1) $\psi = \psi' * \psi''$ modulo a reordering of atoms, such that $\text{fv}(\psi') \cap \text{fv}(\psi'') = \emptyset$,
- (2) $\text{fv}_{J_i}(B_i(z_{i,1}, \dots, z_{i,\#B_i})) \cap \text{fv}(\psi'') = \emptyset$ and
 $\text{fv}_{\bar{J}_i}(B_i(z_{i,1}, \dots, z_{i,\#B_i})) \cap \text{fv}(\psi') = \emptyset$, for all $i \in [1..\ell]$,
- (3) $\Xi = (\text{conn}(\psi') \cup \bigcup_{i=1}^\ell \xi_i(B_i(z_{i,1}, \dots, z_{i,\#B_i}))) =$.

Intuitively, the conditions (1)–(3) above guarantee that the models of $B_i(z_{i,1}, \dots, z_{i,\#B_i})/\xi_i$ (recall, these are sets of maximally connected structures) compose with a model of ψ' without losing neither connectivity nor maximality, in the context of the rule (3.1). At this point, we distinguish two cases:

▷ If there exist sets $J_0 \uplus \bar{J}_0 = [1..\#B_0]$, $J_0 \neq \emptyset$ and an equivalence relation $\xi_0 \subseteq J_0 \times J_0$, such that:

- (4) $\text{fv}_{J_0}(B_0(x_1, \dots, x_{\#B_0})) \cap \text{fv}(\psi'') = \emptyset$ and $\text{fv}_{\bar{J}_0}(B_0(x_1, \dots, x_{\#B_0})) \cap \text{fv}(\psi') = \emptyset$,
- (5) $\text{fv}_{J_0}(B_0(x_1, \dots, x_{\#B_0})) \cap \text{fv}_{\bar{J}_i}(B_i(z_{i,1}, \dots, z_{i,\#B_i})) = \emptyset$ and
 $\text{fv}_{\bar{J}_0}(B_0(x_1, \dots, x_{\#B_0})) \cap \text{fv}_{J_i}(B_i(z_{i,1}, \dots, z_{i,\#B_i})) = \emptyset$, for all $i \in [1..\ell]$,
- (6) for all $y \in (\text{fv}(\psi') \cup \bigcup_{i=1}^\ell \text{fv}_{J_i}(B_i(z_{i,1}, \dots, z_{i,\#B_i}))) \cap \{y_1, \dots, y_m\}$
there exists $x \in \text{fv}_{J_0}(B_0(x_1, \dots, x_{\#B_0}))$, such that $(x, y) \in \Xi$,
- (7) $\xi_0(B_0(x_1, \dots, x_{\#B_0})) = \Xi \upharpoonright_{\{x_1, \dots, x_{\#B_0}\}} \cup \{(x, x) \mid x \in \text{fv}_{J_0}(B_0(x_1, \dots, x_{\#B_0}))\}$

then we add to Γ the following rule:

$$B_0(x_1, \dots, x_{\#B_0})/\xi_0 \leftarrow \exists y_1 \dots \exists y_m \cdot \psi' * \bigstar_{i \in [1..\ell], J_i \neq \emptyset} B_i(z_{i,1}, \dots, z_{i,\#B_i})/\xi_i \quad (3.2)$$

Intuitively, the conditions (4)–(7) identify the set J_0 and the equivalence relation ξ_0 for which the result of the composition becomes a model of $B_0(x_1, \dots, x_{\#B_0})/\xi_0$. Altogether, these lead to the definition of the rules of the form 3.2 which propagate the construction of maximally connected structures in Γ .

▷ If (8) Ξ defines an unique equivalence class, and (9) $(x, x) \notin \Xi$ for all $x \in \{x_1, \dots, x_{\#B_0}\}$ then we add to Γ the following rule:

$$P \leftarrow \exists y_1 \dots \exists y_m \cdot \psi' * \bigstar_{i \in [1..\ell], J_i \neq \emptyset} B_i(z_{i,1}, \dots, z_{i,\#B_i})/\xi_i \quad (3.3)$$

Intuitively, conditions (8)–(9) in addition to (1)–(3), ensure that by composing the models of $B_i(z_{i,1}, \dots, z_{i,\#B_i})/\xi_i$ with a model of ψ' we obtain a single maximally connected structure in the context of the rule (3.1), which is moreover not referred by any of the parameters of B_0 . Henceforth, the result of this composition is actually a model of $\text{split}(\llbracket A \rrbracket_\Delta^c)$ and consequently is added as a model of P by the rules of the form 3.3.

Recall that Δ was assumed to be equality-free (Definition 2.5) and all-satisfiable for \mathbf{A} (Definition 2.8). Moreover, we assume that every predicate defined by a rule of Δ occurs on some complete Δ -unfolding of \mathbf{A} . Obviously, the rules that do not meet this requirement can be removed from Δ without changing $\llbracket \mathbf{A} \rrbracket_{\Delta}^c$. The following lemma shows that the set of canonical Γ -models of \mathbf{P} is a correct representation of the set of canonical Δ -models of \mathbf{A} :

Lemma 3.31. *For each equality-free SID Δ , which is all-satisfiable for a nullary predicate symbol \mathbf{A} , one can effectively build a SID Γ and a nullary predicate \mathbf{P} , such that $\text{split}(\llbracket \mathbf{A} \rrbracket_{\Delta}^c) = \llbracket \mathbf{P} \rrbracket_{\Gamma}^c$.*

Proof. For space reasons, this proof is given in subsection B.2. \square

3.7. Color abstractions of canonical models. We compute the k -multiset color abstraction $(\llbracket \mathbf{P} \rrbracket_{\Gamma}^c)^{\#k}$ by a least fixpoint iteration in a finite abstract domain, defined directly from the rules in the SID. The elements of the domain are composed of the colors of parameter values and the k -multiset color abstraction of the elements not referenced by parameters.

A k -bounded color triple $\langle X, c, M \rangle$ consists of a finite set of variables $X \subseteq \mathbb{V}$, a mapping $c : X \rightarrow \mathbb{C}$, and a multiset $M \in \text{mpow}(\mathbb{C})$, such that $\text{card}(M) \leq k$. Since X and \mathbb{R} are finite, there are finitely many color triples. The following operations on color triples are lifted to sets, as usual:

k -composition: $\langle X_1, c_1, M_1 \rangle \bullet^{\#k} \langle X_2, c_2, M_2 \rangle \stackrel{\text{def}}{=}$

$$\begin{aligned} & \{ \langle X_1 \cup X_2, c_{12}, M_{12} \rangle \mid c_{12}(x) = c_1(x) \uplus c_2(x), \text{ for all } x \in X_1 \cap X_2, \\ & \quad c_{12}(x) = c_i(x) \text{ for all } x \in X_i \setminus X_{3-i}, \text{ for all } i \in \{1, 2\}, \\ & \quad M_{12} \subseteq M_1 \cup M_2, \text{ card}(M_{12}) \leq k \} \end{aligned}$$

This operation is undefined, if $c_1(x) \cap c_2(x) \neq \emptyset$, for some $x \in X_1 \cap X_2$.

substitution: $\langle X, c, M \rangle [s] \stackrel{\text{def}}{=} \langle Y, c \circ s, M \rangle$, for any bijection $s : Y \rightarrow X$

k -projection: $\langle X, c, M \rangle \downarrow_Y^{\#k} \stackrel{\text{def}}{=} \{ \langle Y, c|_Y, M' \rangle \mid M' \subseteq M \cup \{c(x) \mid x \in X \setminus Y\}, \text{ card}(M') \leq k \}$, for $Y \subseteq X$.

For a qpf formula ψ , let $\gamma(\psi) \stackrel{\text{def}}{=} \langle \text{fv}(\psi), \lambda x \in \text{fv}(\psi) . \{r \in \mathbb{R} \mid r(x, \dots, x) \text{ occurs in } \psi\}, \emptyset \rangle$. Given a predicate \mathbf{B} , we denote by $\langle \mathbf{B} \rangle_{\Gamma}^{\#k}$ the least sets of k -bounded color triples over the variables $x_1, \dots, x_{\#\mathbf{B}}$, the satisfies the following constraints:

$$\langle \mathbf{B}_0 \rangle_{\Gamma}^{\#k} \supseteq (\gamma(\psi) \bullet^{\#k} \bullet_{i \in [1..l]}^{\#k} \langle \mathbf{B}_i \rangle_{\Gamma}^{\#k} [x_1/z_{i,1}, \dots, x_{\#\mathbf{B}_i}/z_{i,\#\mathbf{B}_i}]) \downarrow_{\{x_1, \dots, x_{\#\mathbf{B}_0}\}}^{\#k} \quad (3.4)$$

one for each rule of Γ of the form (3.1). Note that the operations on sets of color triples are monotonic and the sets thereof are finite, since the arity of predicates is finite and k is fixed. Henceforth, the least solution can be computed in finite time by an ascending Kleene iteration. For a n -ary relation R , we denote by $\pi_k(R)$ the set of elements that occur on the k -th position in a tuple from R .

Lemma 3.32. $(\llbracket \mathbf{P} \rrbracket_{\Gamma}^c)^{\#k} = \pi_3(\langle \mathbf{P} \rangle_{\Gamma}^{\#k})$, for any $k \geq 1$, SID Γ and nullary predicate \mathbf{P} .

Proof. For space reasons, this proof is given in subsection B.3. \square

3.8. The expandable treewidth boundedness problem. We end this section with a proof of decidability for the treewidth boundedness problem of the sets of Δ -models of a nullary predicate A , provided that Δ is an expandable SID (see Definition 3.8). In case such a bound exists, we provide the optimal upper bound, in terms of the input SID Δ . We recall that $\max\text{Vars}(\Delta)$ is the maximum number of variables that occur, either free or bound by an existential quantifier, in some rule from Δ (Lemma 3.2).

Theorem 3.33. *There exists an algorithm that decides, for each expandable SID Δ and nullary predicate A , whether the set $\llbracket A \rrbracket_\Delta$ has bounded treewidth. If, moreover, this is the case then $\text{tw}(\llbracket A \rrbracket_\Delta) \leq \max\text{Vars}(\Delta)$.*

Proof. Because Δ is expandable, by Lemma 3.12 and Lemma 3.26, $\llbracket A \rrbracket_\Delta$ is treewidth bounded iff $\text{EF}^*(\llbracket A \rrbracket_\Delta^c)$ is treewidth bounded. By Lemma 3.25, $\text{EF}^*(\llbracket A \rrbracket_\Delta^c)$ is treewidth bounded iff $\{\mathcal{C}_1, \mathcal{C}_1, \mathcal{C}_1\}, \{\mathcal{C}_2, \mathcal{C}_2, \mathcal{C}_2\} \in (\text{EF}^*(\text{split}(\llbracket A \rrbracket_\Delta^c)))^{\#3} \Rightarrow \mathcal{C}_1 \cap \mathcal{C}_2 \neq \emptyset$, for all $\mathcal{C}_1, \mathcal{C}_2 \in \mathcal{C}$. The latter condition can be effectively checked by computing the finite set $(\text{EF}^*(\text{split}(\llbracket A \rrbracket_\Delta^c)))^{\#3}$. By Lemma 3.28 and Lemma 3.30, we have $(\text{EF}^*(\text{split}(\llbracket A \rrbracket_\Delta^c)))^{\#3} = \text{ef}_1^{\#k*}(\text{split}(\llbracket A \rrbracket_\Delta^c))^{\#k}$. By Lemma 3.31, one can effectively build a SID Γ and a nullary predicate P , such that $\text{split}(\llbracket A \rrbracket_\Delta^c) = \llbracket P \rrbracket_\Gamma^c$. Moreover, by Lemma 3.32, we obtain $(\llbracket P \rrbracket_\Gamma^c)^{\#3} = \pi_3(\langle P \rangle_\Gamma^{\#3})$, hence $(\text{EF}^*(\text{split}(\llbracket A \rrbracket_\Delta^c)))^{\#3} = \pi_3(\langle P \rangle_\Gamma^{\#3})$, which is effectively computable by a ascending Kleene iteration in the finite domain of 3-bounded color triples.

For the upper bound, since $\llbracket A \rrbracket_\Delta \subseteq \text{IEF}^*(\llbracket A \rrbracket_\Delta^c)$ (Lemma 3.6), we have $\text{tw}(\llbracket A \rrbracket_\Delta) \leq \text{tw}(\text{IEF}^*(\llbracket A \rrbracket_\Delta^c))$. By Lemma 3.21, $\text{tw}(\llbracket A \rrbracket_\Delta) \leq \text{tw}(\llbracket A \rrbracket_\Delta^c) + 1$ and, by Lemma 3.2, we obtain $\text{tw}(\llbracket A \rrbracket_\Delta) \leq \max\text{Vars}(\Delta)$. \square

The bound given by Theorem 3.33 is optimal, as shown by the following example:

Example 3.34. Let us consider the following SID:

$$\Delta = \begin{cases} A \leftarrow \exists y_1 \exists y_2 . a(y_1) * e(y_1, y_2) * A \\ A \leftarrow \text{emp} \end{cases}$$

This SID is expandable for A , because any canonical Δ -model $S = (U, \sigma)$ of A consists of a set of pairs $(u_1, u_2) \in \sigma(e)$, such that $u_1 \in \sigma(a)$ and $u_2 \notin \sigma(a)$. Hence, any sequence of canonical Δ -models of A can be embedded as substructures in a canonical Δ -model of A . Moreover, $\max\text{Vars}(\Delta) = 2$ and any cyclic list of e -related adjacent elements labeled by a is a Δ -model of A of treewidth 2.

4. THE REDUCTION TO EXPANDABLE SETS OF INDUCTIVE DEFINITIONS

This section completes the proof of decidability of the treewidth boundedness problem TWB^{SLR} , by showing a reduction to the decidable treewidth boundedness problem for expandable SIDs (Theorem 3.33). Moreover, an analysis of this reduction allows to compute upper bounds on the treewidth of the set of models of an SLR sentence, provided that such a bound exists. The core of the reduction is the following lemma:

Lemma 4.1. *Let Δ be a SID and A be a nullary predicate. Then, one can build finitely many SIDs $\Gamma_1, \dots, \Gamma_n$, that are expandable for a nullary predicate B , such that $\llbracket A \rrbracket_\Delta$ is treewidth bounded iff each $\llbracket B \rrbracket_{\Gamma_i}$ is treewidth bounded, for $i \in [1..n]$.*

The rest of this section is concerned with the proof of this lemma. For technical reasons, the construction of expandable SIDs with an equivalent treewidth boundedness problem uses a representation of the SID as a tree automaton (subsection 4.1). This representation allows to distinguish the purely structural aspects, related to the dependencies between rules, from details related to the flow of parameters. An important class of automata distinguish between the so-called *1-transitions*, that occur exactly once, from the ∞ -*transitions*, that may occur any number of times on an accepting run. These automata are called *choice-free* (Definition 4.5). Each automaton admits a finite choice-free decomposition that preserves its language (Lemma 4.9).

Next, we consider tree automata whose alphabets are finite sets of qpf formulæ (subsection 4.2). The trees recognized by these automata are representations of the predicate-free formulæ produced by the complete Δ -unfoldings of A . Since we assume that Δ is all-satisfiable for A (Definition 2.8), each accepting run of an automaton “produces” a canonical Δ -model of A (Lemma 4.13).

Furthermore, we define *persistent* variables, whose values are carried along each sequence of ∞ -transitions of a choice-free automaton (Definition 4.14). Identifying and removing the persistent variables from a choice-free automaton with an alphabet of qpf formulæ constitutes an important ingredient of the construction, because of point (3) of Definition 3.8, that requires the embedded canonical models of an expandable SID to be placed sufficiently far away one from another. In particular, this guarantees that the coloring (Definition 3.13) of an element from an embedded substructure does not change in the larger structure. The effective transformation of a choice-free automaton over an alphabet of qpf formulæ into an automaton without persistent variables is described in subsection 4.4. This transformation does not preserve the language, nor the set of models corresponding to the trees recognized by the automaton, but is shown to preserve the existence of a (computable) bound on the treewidths of these models.

4.1. Tree Automata. Let \mathbb{A} be a *ranked alphabet*, each symbol $a \in \mathbb{A}$ having an associated integer *rank* $\rho(a) \geq 0$. The elements of \mathbb{N}_+^* are finite sequences of strictly positive natural numbers, called *positions*. We write pq for the concatenation of $p, q \in \mathbb{N}^*$ and $q \cdot P \stackrel{\text{def}}{=} \{qp \mid p \in P\}$, for $P \subseteq \mathbb{N}^*$. A *ranked tree* is a finite partial function $t : \mathbb{N}^* \rightarrow \mathbb{A}$, such that the set $\text{dom}(t)$ is *prefix-closed*, i.e., for each $p \in \text{dom}(t)$, if q is a prefix of p , then $q \in \text{dom}(t)$, and *sibling-closed*, i.e., $\{i \in \mathbb{N} \mid pi \in \text{dom}(t)\} = \{1, \dots, \rho(t(p))\}$, for all $p \in \text{dom}(t)$. The *frontier* of t is the set $\text{fr}(t) \stackrel{\text{def}}{=} \{p \in \text{dom}(t) \mid p1 \notin \text{dom}(t)\}$. We denote by $t|_p$ the subtree of t at position $p \in \text{dom}(t)$ i.e., $t|_p$ is the tree such that $\text{dom}(t|_p) = \{q \in \mathbb{N}^* \mid pq \in \text{dom}(t)\}$ and $t|_p(q) = t(pq)$, for each $q \in \text{dom}(t|_p)$. A tree u is *embedded in t at position $p \in \text{dom}(t)$* iff $pq \in \text{dom}(t)$ and $u(q) = t(pq)$, for each $q \in \text{dom}(u)$.

Definition 4.2. An (\mathbb{A} -labeled tree) *automaton* is $\mathcal{A} = (\mathbb{A}, \mathcal{Q}, \mathcal{I}, \delta)$, where \mathcal{Q} is a finite set of states, $\mathcal{I} \subseteq \mathcal{Q}$ is a set of initial states (if \mathcal{I} is a singleton, we denote it by $\iota \in \mathcal{Q}$), δ is a finite set of transitions $\tau : q_0 \xrightarrow{a} (q_1, \dots, q_{\rho(a)})$. For a transition $\tau : q_0 \xrightarrow{a} (q_1, \dots, q_\ell) \in \delta$, let $\bullet\tau \stackrel{\text{def}}{=} q_0$ be the source and $\tau^\bullet \stackrel{\text{def}}{=} \{\{q_1, \dots, q_\ell\}\}$ the multiset of targets of τ . For a set of transitions $T \subseteq \delta$, let $\bullet T \stackrel{\text{def}}{=} \{\bullet\tau \mid \tau \in T\}$ and $T^\bullet \stackrel{\text{def}}{=} \bigcup_{\tau \in T} \tau^\bullet$. For a set of states $S \subseteq \mathcal{Q}$, let $\bullet S \stackrel{\text{def}}{=} \{\tau \mid \bullet\tau \notin S, \tau^\bullet \cap S \neq \emptyset\}$, $S^\bullet \stackrel{\text{def}}{=} \{\tau \mid \bullet\tau \in S, \tau^\bullet \cap S = \emptyset\}$ and $\bullet S^\bullet \stackrel{\text{def}}{=} \{\tau \mid \bullet\tau \in S, \tau^\bullet \cap S \neq \emptyset\}$.

The following notions concern the structure of automata. The relation $\rightsquigarrow \subseteq \mathcal{Q} \times \mathcal{Q}$ is defined as $q \rightsquigarrow q'$ iff there exists $\tau \in \delta$ such that $q = \bullet\tau$ and $q' \in \tau^\bullet$. A *strongly connected*

component (SCC) is a maximal set $S \subseteq \mathcal{Q}$, such that $q \rightsquigarrow^* q'$, for all $q, q' \in S$. An SCC S is *nonlinear* iff there exists a transition $\tau \in \bullet S^\bullet$ such that $\text{card}(\tau^\bullet \cap S) \geq 2$ and *linear* otherwise. The *SCC graph* of \mathcal{A} is the directed graph $\mathfrak{G}_{\mathcal{A}} \stackrel{\text{def}}{=} (\mathcal{N}, \mathcal{E})$, where \mathcal{N} is the set of SCCs of \mathcal{A} and $(S, S') \in \mathcal{E}$ iff $S \neq S'$ and there exists $q \in S$ and $q' \in S'$, such that $q \rightsquigarrow q'$, for all $S, S' \in \mathcal{N}$. We write $\mathfrak{G}_{\mathcal{A}} = (\mathcal{N}, \mathcal{E}, S)$ if $\mathfrak{G}_{\mathcal{A}}$ is a tree with root $S \in \mathcal{N}$.

The execution of automata is defined next. A *run* θ of \mathcal{A} over a ranked tree t is a tree $\theta : \text{dom}(t) \rightarrow \mathcal{Q}$ such that $\theta(p) \xrightarrow{t(p)} (\theta(p1), \dots, \theta(p\ell)) \in \delta$, for all $p \in \text{dom}(t)$, where $\ell = \rho(t(p))$. Note that the frontier of a run is labeled by states q such that there exists a transition $q \xrightarrow{\alpha} () \in \delta$, in analogy to the final states of a word automaton. A weaker notion is that of *partial* runs, where the previous condition holds for $\text{dom}(t) \setminus \text{fr}(t)$, instead of the entire $\text{dom}(t)$. A run θ is *accepting* if $\theta(\epsilon) \in \mathcal{I}$. The *language* of \mathcal{A} is $\mathcal{L}(\mathcal{A}) \stackrel{\text{def}}{=} \bigcup_{q \in \mathcal{I}} \mathcal{L}_q(\mathcal{A})$, where $\mathcal{L}_q(\mathcal{A}) \stackrel{\text{def}}{=} \{t \mid \mathcal{A} \text{ has a run } \theta \text{ over } t \text{ and } \theta(\epsilon) = q\}$.

An automaton is *rooted* iff $\mathcal{I} = \{\iota\}$ and $\iota \notin \delta^\bullet$. For an automaton \mathcal{A} one can build finitely many rooted automata $\mathcal{A}_1, \dots, \mathcal{A}_n$ such that $\mathcal{L}(\mathcal{A}) = \bigcup_{i=1}^n \mathcal{L}(\mathcal{A}_i)$. A rooted automaton \mathcal{A} is *trim* iff $\iota \rightsquigarrow^* q$ and $\mathcal{L}_q(\mathcal{A}) \neq \emptyset$, for each state $q \in \mathcal{Q}$. Each automaton with non-empty language can be transformed into a trim one with the same language, by a simple marking algorithm. We use the following notions of *simulation* and *refinement* between automata:

Definition 4.3. Let $\mathcal{A} = (\mathbb{A}, \mathcal{Q}_{\mathcal{A}}, \iota_{\mathcal{A}}, \delta_{\mathcal{A}})$ and $\mathcal{B} = (\mathbb{A}, \mathcal{Q}_{\mathcal{B}}, \iota_{\mathcal{B}}, \delta_{\mathcal{B}})$ be automata. A mapping $h : \mathcal{Q}_{\mathcal{A}} \rightarrow \mathcal{Q}_{\mathcal{B}}$ is a *simulation* if and only if the following hold:

- (1) $h(\iota_{\mathcal{A}}) = \iota_{\mathcal{B}}$ and
- (2) $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta_{\mathcal{A}}$ only if $h(q_0) \xrightarrow{\alpha} (h(q_1), \dots, h(q_\ell)) \in \delta_{\mathcal{B}}$, for all $q_0, \dots, q_\ell \in \mathcal{Q}_{\mathcal{A}}$.

A simulation h is a *refinement* if and only if, moreover:

- (3) $q'_0 \xrightarrow{\alpha} (q'_1, \dots, q'_\ell) \in \delta_{\mathcal{B}}$ only if there exist $q_0 \in h^{-1}(q'_0), \dots, q_\ell \in h^{-1}(q'_\ell)$, such that $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta_{\mathcal{A}}$, for all $q_0 \in \mathcal{Q}_{\mathcal{A}}$ and $q'_1, \dots, q'_\ell \in \mathcal{Q}_{\mathcal{B}}$.

If $h : \mathcal{Q}_{\mathcal{A}} \rightarrow \mathcal{Q}_{\mathcal{B}}$ is a simulation then \mathcal{B} *simulates* \mathcal{A} . If h is a refinement then \mathcal{A} *refines* \mathcal{B} .

The key properties of simulations and refinements are stated and proved below:

Lemma 4.4. *If \mathcal{B} simulates \mathcal{A} then $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{B})$. If \mathcal{A} is a refinement of \mathcal{B} then $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{B})$.*

Proof. Let $\mathcal{A} = (\mathbb{A}, \mathcal{Q}_{\mathcal{A}}, \iota_{\mathcal{A}}, \delta_{\mathcal{A}})$, $\mathcal{B} = (\mathbb{A}, \mathcal{Q}_{\mathcal{B}}, \iota_{\mathcal{B}}, \delta_{\mathcal{B}})$ and $h : \mathcal{Q}_{\mathcal{A}} \rightarrow \mathcal{Q}_{\mathcal{B}}$ be a mapping. “ \subseteq ” Assume that h is a simulation. Let $t \in \mathcal{L}(\mathcal{A})$ be a tree and θ be an accepting run of \mathcal{A} over t . Then one shows that $h \circ \theta$ is an accepting run of \mathcal{B} over t , by induction on t , using points (1) and (2) of Definition 4.3. “ \supseteq ” Assume that h is a refinement. Let $t \in \mathcal{L}(\mathcal{B})$ be a tree and θ be an accepting run of \mathcal{B} over t . We build an accepting run of \mathcal{A} over t by induction on t , using points (1) and (3) of Definition 4.3. \square

The following structural property of automata is key for building expandable SIDs:

Definition 4.5. An automaton $\mathcal{A} = (\mathbb{A}, \mathcal{Q}, \iota, \delta)$ is *choice-free* iff the following hold:

- (1) the SCC graph of \mathcal{A} is a tree $\mathfrak{G}_{\mathcal{A}} = (\mathcal{N}, \mathcal{E}, S_0)$, where $\bullet S = \{\tau\}$ and $\text{card}(\tau^\bullet \cap S) = 1$, for all $S \in \mathcal{N} \setminus \{S_0\}$, i.e., any non-root SCC is entered by one branch of a single transition,
- (2) there exists a mapping $\Lambda : \mathcal{N} \cup \delta \rightarrow \{1, \infty\}$ such that:
 - (a) for all $S \in \mathcal{N}$, if S is linear and $\Lambda(S) = 1$ then $\text{card}(S^\bullet) = 1$.
 - (b) for all $\tau \in \delta$, $\Lambda(\tau) = 1$ iff $\tau \in S^\bullet$, for some linear $S \in \mathcal{N}$ such that $\Lambda(S) = 1$,
 - (c) for all $S \in \mathcal{N}$, $\Lambda(S) = 1$ iff $S = S_0$ or $\bullet S = \{\tau\}$, for some $\tau \in \delta$ such that $\Lambda(\tau) = 1$.




Figure 10: A choice-free tree automaton and its non-choice-free extension

Let $\delta = \delta^1 \uplus \delta^\infty$, where $\delta^k \stackrel{\text{def}}{=} \{\tau \in \delta \mid \Lambda(\tau) = k\}$ and $k \in \{1, \infty\}$, be the partition of the set of transitions induced by the mapping Λ . A state $q \in (\delta^1)^\bullet \cap \bullet(\delta^\infty)$ is called a *pivot state*. Let $\mathcal{R}_q^\infty(\mathcal{A})$ denote the set of partial runs θ of \mathcal{A} , such that $\theta(\epsilon) = q$ and for all $p \in \text{dom}(\theta) \setminus \text{fr}(\theta)$, there exists $a \in \mathbb{A}$ such that $\theta(p) \xrightarrow{a} (\theta(p1), \dots, \theta(pn)) \in \delta^\infty$.

Intuitively, the structure of choice-free automata allows them to traverse a unique sequence of linear SCCs, before entering a non-linear SCC. Note that the labeling Λ of SCCs and transitions from Definition 4.5 is unique, because the SCC graph of a choice-free automaton is a tree whose root is a 1-SCC, and Λ is determined by the linearity of the SCCs in this tree. Hence, the partition of the transitions of a choice-free automaton into 1- and ∞ -transitions is unambiguous.

Example 4.6. The automaton from Figure 10 is choice free. The linear (resp. nonlinear) SCCs are labeled by L (resp. NL). The labeling of transitions and SCCs, resp. the SCC graph by Λ are represented in Figure 10 (resp. right). The choice-freeness is violated by adding the transitions $q_0 \xrightarrow{b} (q_5)$ and $q_5 \xrightarrow{c} ()$ (in red), because the linear SCC $\{q_0\}$ is labeled with 1 and has two outgoing transitions, thus contradicting point (2a) of Definition 4.5 (the additional transitions are not labeled).

The transitions from δ^1 , called *1-transitions*, are used to move from one linear SCC to another, hence all of them occur exactly once on each accepting run:

Lemma 4.7. *Let $\mathcal{A} = (\mathbb{A}, \mathcal{Q}, \iota, \delta)$ be a choice-free automaton, such that $\delta = \delta^1 \uplus \delta^\infty$ (Definition 4.5) and let θ be an accepting run of \mathcal{A} over a tree t . Then, for each 1-transition $q_0 \xrightarrow{a} (q_1, \dots, q_\ell) \in \delta^1$ there exists exactly one position $p \in \text{dom}(\theta)$, such that $\theta(p) = q_0$, $t(p) = a$ and $\theta(pi) = q_i$, for all $i \in [1..\ell]$.*

Proof. For space reasons, the proof of this lemma is given in subsection C.1. □

The transitions from δ^∞ , called *∞ -transitions*, can be applied any number of times on some accepting run. This fact occurs as an easy consequence of the lemma below:

Lemma 4.8. *Let $\mathcal{A} = (\mathbb{A}, \mathcal{Q}, \iota, \delta)$ be a choice-free automaton, where $\delta = \delta^1 \uplus \delta^\infty$ (Definition 4.5). Then, for any state $q \in \bullet(\delta^\infty)$ there exists a pivot state $q_0 \in (\delta^1)^\bullet \cap \bullet(\delta^\infty)$ and a partial run $\theta_0 \in \mathcal{R}_{q_0}^\infty(\mathcal{A})$ consisting only of ∞ -transitions, such that $\theta_0(p) = q$ for some $p \in \text{fr}(\theta_0)$ and either:*

- (1) $\{\{q, q_0\}\} \subseteq \{\{\theta_0(p) \mid p \in \text{fr}(\theta_0)\}\}$, i.e., if $q = q_0$ then q occurs twice on $\text{fr}(\theta_0)$, or
- (2) each partial run $\theta \in \mathcal{R}_q^\infty(\mathcal{A})$ can be extended to a partial run $\theta' \in \mathcal{R}_q^\infty(\mathcal{A})$ such that q_0 occurs on the frontier of θ' .

Proof. For space reasons, the proof of this lemma is given in subsection C.2. □

Moreover, any automaton can be decomposed into finitely many choice-free automata:

Lemma 4.9. *Given an automaton $\mathcal{A} = (\mathbb{A}, \mathcal{Q}, \iota, \delta)$, one can build finitely many choice-free automata $\mathcal{A}_i = (\mathbb{A}, \mathcal{Q}_i, \iota_i, \delta_i)$, for $i \in [1..n]$, such that $\mathcal{L}(\mathcal{A}) = \bigcup_{i=1}^n \mathcal{L}(\mathcal{A}_i)$ and, moreover, $\text{card}(\delta_i^1) \leq \max(\text{card}(\mathcal{Q}), \max\{\rho(a) \mid a \in \mathbb{A}\}^{\text{card}(\mathcal{Q})})$, for all $i \in [1..n]$.*

Proof. For space reasons, the proof of this lemma is given in subsection C.3. \square

Example 4.10. The choice-free decomposition of the automaton \mathcal{A} from Figure 10 consists of the following choice-free automata, with transitions labeled according to Definition 4.5:

$$\mathcal{A}_1 = \begin{cases} q_0 \xrightarrow{a} (q_0, q_1) & (\infty) \\ q_1 \xrightarrow{b} (q_1) & (\infty) \\ q_1 \xrightarrow{c} () & (\infty) \\ q_0 \xrightarrow{b} (q_2) & (1) \\ q_2 \xrightarrow{a} (q_2, q_2) & (\infty) \\ q_2 \xrightarrow{b} (q_3) & (\infty) \\ q_2 \xrightarrow{b} (q_4) & (\infty) \\ q_3 \xrightarrow{c} () & (\infty) \\ q_4 \xrightarrow{c} () & (\infty) \end{cases} \quad \mathcal{A}_2 = \begin{cases} q_0 \xrightarrow{a} (q_0, q_1) & (\infty) \\ q_1 \xrightarrow{b} (q_1) & (\infty) \\ q_1 \xrightarrow{c} () & (\infty) \\ q_0 \xrightarrow{b} (q_5) & (1) \\ q_5 \xrightarrow{c} () & (1) \end{cases}$$

It can be seen that $\mathcal{L}_{q_0}(\mathcal{A}) = \mathcal{L}_{q_0}(\mathcal{A}_1) \cup \mathcal{L}_{q_0}(\mathcal{A}_2)$ because \mathcal{A} has the choice in q_0 between taking the transition $q_0 \xrightarrow{b} (q_2)$ or $q_0 \xrightarrow{b} (q_5)$. Since these transitions occur at most once in each accepting run of \mathcal{A} , the choice-free decomposition of \mathcal{A} produces two automata in which each such transition occurs exactly once on each accepting run.

4.2. Automata with Alphabets of Formulæ. The construction of the expandable SIDs from Lemma 4.1 uses automata that recognize trees labeled with qpf formulæ taken from a finite set. We recall that every model of a sentence is defined by a complete unfolding that replaces the predicate atoms with corresponding definitions, recursively. The steps of these unfoldings can be placed into a tree labeled with predicate-free formulæ from an alphabet Σ , reflecting the partial order in which the rules from the SID are applied. These unfolding trees form the language of an automaton defined directly from the syntax of the SID. Dually, from any Σ -labeled automaton one can build a SID whose unfolding trees form the language of the automaton.

Definition 4.11. Let Σ be the set of qpf formulæ α of rank $\rho(\alpha) = \ell$, such that:

- (1) $\text{fv}(\alpha) \subseteq \{x_1^{[\epsilon]}, \dots, x_{n_0}^{[\epsilon]}\} \cup \{y_1^{[\epsilon]}, \dots, y_m^{[\epsilon]}\} \cup \bigcup_{i=1}^{\ell} \{x_1^{[i]}, \dots, x_{n_i}^{[i]}\}$, for some $m, n_0, \dots, n_\ell \in \mathbb{N}$;
a variable $x_j^{[i]}$ is called a *i-variable*, for all $i \in \{\epsilon\} \cup [1..\ell]$,
- (2) $x_j^{[i]} \not\approx_\alpha x_k^{[i]}$, for all $i \in [1..\ell]$ and $1 \leq j < k \leq n_i$.

The *characteristic formula* of a Σ -labeled tree t is the qpf formula $\Theta(t) \stackrel{\text{def}}{=} \bigstar_{p \in \text{dom}(t)} t(p)^{[p]}$, where the formulæ $t(p)^{[p]}$ are obtained from $t(p) \in \Sigma$ by replacing each occurrence of a variable $x^{[q]}$ by $x^{[pq]}$, for all $p \in \text{dom}(t)$.

Given a SID Δ , the Σ -labeled automaton $\mathcal{A}_{\Delta, \mathbb{A}} \stackrel{\text{def}}{=} (\Sigma, \mathcal{Q}_\Delta, q_\mathbb{A}, \delta_\Delta)$ is defined as follows:
 $\triangleright \mathcal{Q}_\Delta$ contains states $q_\mathbb{B}$, where \mathbb{B} is a predicate occurring in Δ ; each state has an associated arity $\#q_\mathbb{B} \stackrel{\text{def}}{=} \#\mathbb{B}$,

$\triangleright \delta_\Delta$ contains a transition $q_{A_0} \xrightarrow{\alpha_\rho} (q_{A_1}, \dots, q_{A_\ell})$, where α_ρ is the symbol:

$$\alpha_\rho \stackrel{\text{def}}{=} \psi[x_1/x_1^{[\epsilon]}, \dots, x_{n_0}/x_{n_0}^{[\epsilon]}, y_1/y_1^{[\epsilon]}, \dots, y_m/y_m^{[\epsilon]}] * \bigstar_{i=1}^\ell \bigstar_{j=1}^{n_i} z_{i,j}^{[\epsilon]} = x_j^{[\epsilon]}$$

of rank $\rho(\alpha_\rho) = \ell$ that corresponds to the rule $\rho \in \Delta$, where ψ is a qpf formula:

$$\rho : \mathbf{A}_0(x_1, \dots, x_{n_0}) \leftarrow \exists y_1 \dots \exists y_m . \psi * \bigstar_{i=1}^\ell \mathbf{A}_i(z_{i,1}, \dots, z_{i,n_i})$$

Example 4.12. Let us consider the following SID:

$$\Delta = \begin{cases} \mathbf{A} & \leftarrow \exists y_1 \exists y_2 \exists y_3 . \mathbf{B}(y_1, y_2, y_3) \\ \mathbf{B}(x_1, x_2, x_3) & \leftarrow \exists y_4 . \mathbf{e}(x_1, x_3) * \mathbf{e}(x_1, y_4) * \mathbf{B}(y_4, x_2, x_3) \\ \mathbf{B}(x_1, x_2, x_3) & \leftarrow \mathbf{e}(x_1, x_3) * \mathbf{e}(x_1, x_2) * \mathbf{e}(x_2, x_3) \end{cases}$$

The automaton $\mathcal{A}_{\Delta, \mathbf{A}}$ has the following transitions:

$$\mathcal{A}_{\Delta, \mathbf{A}} = \begin{cases} q_{\mathbf{A}} & \xrightarrow{y_1^{[\epsilon]}=x_1^{[1]} * y_2^{[\epsilon]}=x_2^{[1]} * y_3^{[\epsilon]}=x_3^{[1]}} (q_{\mathbf{B}}) \\ q_{\mathbf{B}} & \xrightarrow{\mathbf{e}(x_1^{[\epsilon]}, x_3^{[\epsilon]}) * \mathbf{e}(x_1^{[\epsilon]}, y_4^{[\epsilon]}) * y_4^{[\epsilon]}=x_1^{[1]} * x_2^{[\epsilon]}=x_2^{[1]} * x_3^{[\epsilon]}=x_3^{[1]}} (q_{\mathbf{B}}) \\ q_{\mathbf{B}} & \xrightarrow{\mathbf{e}(x_1^{[\epsilon]}, x_3^{[\epsilon]}) * \mathbf{e}(x_1^{[\epsilon]}, x_2^{[\epsilon]}) * \mathbf{e}(x_2^{[\epsilon]}, x_3^{[\epsilon]})} () \end{cases}$$

Dually, given an automaton $\mathcal{A} = (\Sigma, \mathcal{Q}, \iota, \delta)$, the SID $\Delta_{\mathcal{A}}$ consists of the following rules, one for each transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta$:

$$\mathbf{A}_{q_0}(x_1, \dots, x_{\#q_0}) \leftarrow \exists y_1 \dots \exists y_m . \alpha[x_1^{[\epsilon]}/x_1, \dots, x_{\#q_0}^{[\epsilon]}/x_{\#q_0}] * \bigstar_{j=1}^\ell \mathbf{A}_{q_j}(x_1^{[j]}, \dots, x_{\#q_j}^{[j]}) \quad (4.1)$$

$$\text{where } \{y_1, \dots, y_m\} \stackrel{\text{def}}{=} \text{fv}(\alpha) \setminus (\{x_1^{[\epsilon]}, \dots, x_{\#q_0}^{[\epsilon]}\} \cup \bigcup_{j=1}^\ell \{x_1^{[j]}, \dots, x_{\#q_j}^{[j]}\})$$

The similarity between SIDs and Σ -labeled automata (Lemma 4.13) motivates the use of similar terminology. For a Σ -labeled automaton \mathcal{A} , we define $\llbracket \mathcal{A} \rrbracket \stackrel{\text{def}}{=} \bigcup_{t \in \mathcal{L}(\mathcal{A})} \llbracket \Theta(t) \rrbracket^\exists$, $\llbracket \mathcal{A} \rrbracket^c \stackrel{\text{def}}{=} \bigcup_{t \in \mathcal{L}(\mathcal{A})} \llbracket \Theta(t) \rrbracket^{\exists^c}$ and $\llbracket \mathcal{A} \rrbracket^r \stackrel{\text{def}}{=} \bigcup_{t \in \mathcal{L}(\mathcal{A})} \llbracket \Theta(t) \rrbracket^{\exists^r}$. Moreover, a Σ -labeled automaton \mathcal{A} is all-satisfiable if the formula $\Theta(t)$ is satisfiable, for all $t \in \mathcal{L}(\mathcal{A})$. The relation between SIDs and Σ -labeled automata is formally stated below:

Lemma 4.13.

- (1) Given a SID Δ and a nullary predicate \mathbf{A} , one can build an automaton $\mathcal{A}_{\Delta, \mathbf{A}}$ such that $\llbracket \mathbf{A} \rrbracket_\Delta = \llbracket \mathcal{A}_{\Delta, \mathbf{A}} \rrbracket$. Moreover, if Δ is all-satisfiable for \mathbf{A} , then $\mathcal{A}_{\Delta, \mathbf{A}}$ is all-satisfiable.
- (2) Given an automaton $\mathcal{A} = (\Sigma, \mathcal{Q}, \iota, \delta)$, one can build a SID $\Delta_{\mathcal{A}}$, such that $\llbracket \mathcal{A} \rrbracket = \llbracket \mathbf{A}_\iota \rrbracket_{\Delta_{\mathcal{A}}}$ and $\llbracket \mathcal{A} \rrbracket^r = \llbracket \mathbf{A}_\iota \rrbracket_{\Delta_{\mathcal{A}}}^r$, for a nullary predicate \mathbf{A} .

Proof. For space reasons, the proof of this lemma is given in subsection C.4. \square

4.3. Persistent Variables. The second ingredient of the construction of the expandable SIDs from Lemma 4.1 are the *persistent* variables of a Σ -labeled choice-free automaton. These are variables introduced by the 1-transitions of the automaton, whose values propagate via equalities throughout each run of the choice-free automaton. We define persistent variables using the notion of *profile*:

Definition 4.14. Let $\mathcal{A} = (\Sigma, \mathcal{Q}, \iota, \delta)$ be a choice-free automaton, where $\delta = \delta^1 \uplus \delta^\infty$ (Definition 4.5). A *positional function* $\mathfrak{P} : \mathcal{Q} \rightarrow \text{pow}(\mathbb{N})$ associates each state q with a set $\mathfrak{P}(q) \subseteq [1.. \#q]$. The *profile of \mathcal{A}* is the pointwise largest positional function $\mathfrak{P}_{\mathcal{A}}$ such that, for each transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta^\infty$, each $k \in [1.. \ell]$ and each $r \in \mathfrak{P}_{\mathcal{A}}(q_k)$, there exists

$s \in \mathfrak{P}_{\mathcal{A}}(q_0)$, such that $x_s^{[\epsilon]} \approx_{\alpha} x_r^{[k]}$. A variable $x_j^{[i]}$ that occurs within the label of a transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_{\ell}) \in \delta$ is said to be *persistent* iff $j \in \mathfrak{P}_{\mathcal{A}}(q_i)$, for all $i \in [0..\ell]$.

Intuitively, $\mathfrak{P}_{\mathcal{A}}(q)$ is the set of indices of those variables, associated with a state, that will be equated, through a chain of equalities in the characteristic formula $\Theta(t)$, to the same variable associated with the corresponding pivot state (Lemma 4.8) in every run of ∞ -transitions of \mathcal{A} over t . Note that the profile is computable by a finite greatest fixpoint Kleene iteration over each SCC of the automaton (see the proof of Lemma 4.18 for an explicit statement of that fixed point).

Without loss of generality, we assume that a Σ -labeled automaton does not have trivial SCCs, i.e., consisting of a single state, with no transitions that are both outgoing and incoming to that state. The profile associated with such state q would be the interval $[1..\#q]$ and any variable in the label of an incoming or outgoing transition would be unnecessarily considered persistent. The trivial SCCs of a Σ -labeled automaton can be eliminated by a pre-processing step which combines the labels of the incoming and outgoing transitions and renames the variables according to the convention (Definition 4.11).

Example 4.15. (continued from Example 4.12) The profile of the automaton $\mathcal{A}_{\Delta, \mathcal{A}}$, that corresponds to the SID Δ from Example 4.12 associates $q_{\mathcal{A}}$ with the empty set and $q_{\mathcal{B}}$ with the set $\{2, 3\}$. Note that $\#q_{\mathcal{A}} = 0$ and $\#q_{\mathcal{B}} = 3$. The first and third transitions in Example 4.12 are 1-transitions, whereas the second transition is an ∞ -transition. The variables $x_2^{[1]}$ and $x_3^{[1]}$ from the label of the second transition (∞) are persistent.

A *context* $\theta_{p \leftarrow q}$ is a partial run over a tree t such that $p \in \text{fr}(\theta_{p \leftarrow q})$, $\theta_{p \leftarrow q}(p) = q$ and $\theta_{p \leftarrow q}(r) \xrightarrow{t(r)} ()$, for all $r \in \text{fr}(\theta_{p \leftarrow q}) \setminus p$, i.e., the partial run has exactly one “open” frontier position p that is labeled with a state q . A key property of automata is that equalities between non-persistent variables vanish in contexts consisting of ∞ -transitions only (Lemma 4.18). These contexts, called *resets*, are formally defined below:

Definition 4.16. A context $\theta_{p \leftarrow q} \in \mathcal{R}_q^{\infty}(\mathcal{A})$ over a tree t is a q -*reset* iff (1) $x_j^{[\epsilon]} \approx_{\Theta(t)} x_j^{[p]}$, for all $j \in \mathfrak{P}_{\mathcal{A}}(q)$, and (2) $x_j^{[\epsilon]} \not\approx_{\Theta(t)} x_k^{[p]}$, for all $j, k \in [1..\#q]$, such that $k \notin \mathfrak{P}_{\mathcal{A}}(q)$. The path between ϵ and p in $\theta_{p \leftarrow q}$ is called a *reset path*.

Example 4.17. (continued from Example 4.12) For instance, in the context $\theta_{1 \leftarrow q_{\mathcal{B}}}$, that consists of the second transition $q_{\mathcal{B}} \xrightarrow{\alpha} q_{\mathcal{B}}$ of the automaton in Example 4.12, we have $x_1^{[\epsilon]} \not\approx_{\alpha} x_1^{[1]}$. Then, the value of $x_1^{[\epsilon]}$ is “forgotten” along any run that iterates this transition at least twice.

Lemma 4.18. *Let $\mathcal{A} = (\Sigma, \mathcal{Q}, \iota, \delta)$ be a trim automaton. Then, there exists a q -reset for (1) each pivot state $q \in (\delta^1)^{\bullet} \cap \bullet(\delta^{\infty})$, and (2) each state $q \in \bullet(\delta^1) \cap \bullet(\delta^{\infty})$, i.e., that is the origin of both a 1-transition and a ∞ -transition of \mathcal{A} .*

Proof. By Definition 4.14, $\mathfrak{P}_{\mathcal{A}}$ is the greatest fixpoint of the monotone function \mathcal{F} on the domain of positional functions $\mathfrak{P} : \mathcal{Q} \rightarrow \text{pow}(\mathbb{N})$, defined below:

$$\mathcal{F}(\mathfrak{P}) \stackrel{\text{def}}{=} \lambda q . \bigcap_{\substack{q_0 \xrightarrow{\alpha} (q_1, \dots, q_{\ell}) \in \delta^{\infty} \\ q = q_k \in \{q_1, \dots, q_{\ell}\}}} \{r \in [1..\#q_k] \mid \exists s \in \mathfrak{P}(q_0) . x_s^{[\epsilon]} \approx_{\alpha} x_r^{[k]}\}$$

Namely, we have $\mathfrak{P}_{\mathcal{A}} = \mathcal{F}^i(\top) = \mathcal{F}^j(\top)$, for a sufficiently large $i \geq 1$ and any $j \geq i$, where \top is the positional function $\lambda q . [1..\#q]$. Now consider the following “big-step” function \mathcal{G}

on the domain of positional functions:

$$\mathcal{G}(\mathfrak{P}) \stackrel{\text{def}}{=} \lambda q . \bigcap_{\substack{\theta \in \mathcal{R}_q^\infty(\mathcal{A}) \text{ partial run over } t \\ p \in \text{fr}(\theta), \text{ such that } \theta(p)=q}} \{r \in [1..\#q] \mid \exists s \in \mathfrak{P}(q) . x_s^{[\epsilon]} \approx_{\Theta(t)} x_r^{[p]}\}$$

We prove the following:

Fact 4.19. $\text{gfp}(\mathcal{F})(q) = \text{gfp}(\mathcal{G})(q)$, for any pivot state q of \mathcal{A} .

Proof. “ \subseteq ” Each partial run $\theta \in \mathcal{R}_q^\infty(\mathcal{A})$ such that $\theta(p) = q$, for some $p \in \text{fr}(\theta)$ corresponds to a finite sequence of transitions from δ^∞ . “ \supseteq ” Since q is a pivot state we have $q \in \bullet(\delta^\infty)$, thus necessarily $q = q_0$, where $q_0 \in (\delta^1)^\bullet$ is the state whose existence is stated by Lemma 4.8. Then every ∞ -transition incoming to q belongs to a partial run $\theta \in \mathcal{R}_q^\infty(\mathcal{A})$, such that q occurs on the frontier of θ . \square

Back to the proof, we prove the two points of the statement below:

(1) Let q be a pivot state of \mathcal{A} . By Fact 4.19, we have $\mathfrak{P}_{\mathcal{A}}(q) = \mathcal{G}^i(\top)(q)$ for a sufficiently large finite integer $i \geq 0$. We show that the latter condition is equivalent to the existence of a q -reset $\theta_{p \leftarrow q} \in \mathcal{R}_q^\infty(\mathcal{A})$.

“ \Leftarrow ” Assume that there exists a q -reset $\theta \in \mathcal{R}_q^\infty(\mathcal{A})$ over some tree t . Then $\mathcal{G}(\top)(q) = \{j \in [1..\#q] \mid \exists k \in [1..\#q] . x_j^{[\epsilon]} \approx_{\Theta(t)} x_k^{[p]}\} = \mathfrak{P}_{\mathcal{A}}(q)$.

“ \Rightarrow ” Assume there exists $i \geq 0$, such that $\mathcal{G}^i(\top)(q) = \mathfrak{P}_{\mathcal{A}}(q)$ and let i be the smallest such integer. Then $\top, \mathcal{G}(\top), \mathcal{G}^2(\top), \dots, \mathcal{G}^i(\top)$ is a strictly decreasing sequence hence, for each $j \in [1..i]$, there exists a partial run $\theta_j \in \mathcal{R}_q^\infty(\mathcal{A})$ over some tree t_j and a position $p_j \in \text{fr}(t_j)$, such that $\theta_j(p_j) = q$ and $\{r \in [1..\#q] \mid \exists s \in \mathcal{G}^{j-1}(\top)(q) . x_s^{[\epsilon]} \approx_{\Theta(t_j)} x_r^{[p_j]}\} \subsetneq \mathcal{G}^{j-1}(\top)(q)$. We compose these partial runs $\theta_1, \dots, \theta_i$ by appending each θ_j to θ_{j-1} at position $p_{j-1} \in \text{fr}(\theta_{j-1})$, for all $j \in [2..i]$ into a partial run $\theta'' \in \mathcal{R}_q^\infty(\mathcal{A})$. We define a context $\bar{\theta}_{p_i \leftarrow q}$ by appending to each position $r \in \text{fr}(\theta) \setminus \{p_i\}$ a complete run starting in $\theta(r)$. By the fact that \mathcal{A} is trim, such a run exists. The context $\bar{\theta}_{p_i \leftarrow q}$ satisfies condition (2) of Definition 4.16, but not necessarily (1). Let $\pi : \mathfrak{P}_{\mathcal{A}}(q) \rightarrow \mathfrak{P}_{\mathcal{A}}(q)$ be a permutation defined as $\pi(i) = j$ iff $x_i \approx_{\Theta(\bar{t})} x_j$, where \bar{t} is the tree recognized by the partial run $\bar{\theta}_{p_i \leftarrow q}$ of \mathcal{A} . Note that the choice of j is not unique, but one exists, by Definition 4.14. Then we define the q -reset $\theta_{p_i \leftarrow q}$ by appending $\bar{\theta}_{p_i \leftarrow q}$ to itself at position p_i a number of times equal to the order of π . Then, one can check that $\theta_{p_i \leftarrow q}$ satisfies both conditions of Definition 4.16.

(2) Let S be the SCC of q in \mathcal{A} . Since $q = \bullet\tau$, for some transition $\tau \in \delta^1$, it must be the case that S is a linear SCC, by Definition 4.5. Also $q \in \bullet(\delta^\infty)$ thus, by Lemma 4.8, there exists a pivot state q_0 in S and let $\theta^0 \in \mathcal{R}_q^\infty(\mathcal{A})$ be a partial run from q to q_0 with transitions from $\bullet S \subseteq \delta^\infty$. From point (1) above we obtain a q_0 -reset $\theta_{p_1 \leftarrow q_0}^1 \in \mathcal{R}_{q_0}^\infty(\mathcal{A})$ such that $x_j^{[\epsilon]} \approx_{\Theta(t)} y_k^{[p_1]}$, for all $j, k \in \mathfrak{P}_{\mathcal{A}}(q_0)$ and $x_j^{[\epsilon]} \not\approx_{\Theta(t)} y_k^{[p_1]}$, for all $j, k \in [1..\#q_0]$, $k \notin \mathfrak{P}_{\mathcal{A}}(q_0)$. Moreover, there exists another context $\theta_{p_2 \leftarrow q}^2 \in \mathcal{R}_{q_0}^\infty(\mathcal{A})$. By the choice of the pivot state q_0 , there exists a position $p_0 \in \text{fr}(\theta^0)$ such that $\theta^0(p_0) = q_0$. Let $p \stackrel{\text{def}}{=} p_0 p_1 p_2$ and $\bar{\theta}_{p \leftarrow q}$ be the context consisting of θ^0 to which we append, in this order:

- $\triangleright \theta^1$ on position p_0 ,
- $\triangleright \theta^2$ on the position $p_0 p_2$,
- \triangleright to any other position $r \in (\text{fr}(\theta^0) \setminus \{p_0\}) \cup (p_0 \cdot \text{fr}(\theta^1) \setminus \{p_1\}) \cup (p_0 p_1 \cdot \text{fr}(\theta^2) \setminus \{p_2\})$ a complete run starting in:
 - $- \theta^0(r)$ if $r \in \text{fr}(\theta^0) \setminus \{p_0\}$,

- $\theta^1(r')$ if $r = p_0 r'$ and $r' \in \text{fr}(\theta^1) \setminus \{p_1\}$, and
- $\theta^2(r'')$ if $r = p_0 p_1 r''$ and $r'' \in \text{fr}(\theta^2) \setminus \{p_2\}$.

Such runs exist by the assumption that \mathcal{A} is trim. Moreover, these runs use only ∞ -transitions, because their states are from ∞ -SCCs (Definition 4.5).

It is easy to check that $\bar{\theta}$ satisfies condition (2) of Definition 4.16. In order to satisfy condition (1), in addition to (2), we append $\bar{\theta}$ to itself at position p , using the same idea as in the construction at point (1). \square

The purpose of introducing resets is proving that any sequence of partial runs consisting of ∞ -transitions can be embedded in a complete run, such that each two such partial runs are separated by any number of resets. This is a key ingredient for the proof of the “embedding” property of the canonical models of expandables SIDs (Definition 3.8).

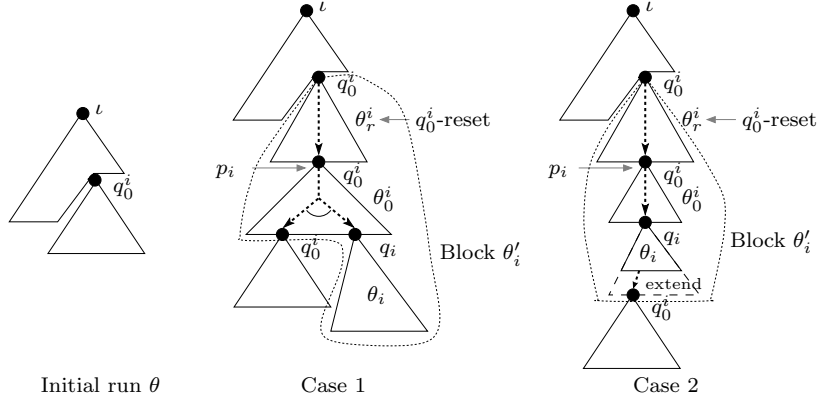


Figure 11: Embedding of a partial run θ_i in θ

Lemma 4.20. *Let \mathcal{A} be a trim choice-free automaton. Given partial runs $\theta_1 \in \mathcal{R}_{q_1}^\infty(\mathcal{A})$, \dots , $\theta_n \in \mathcal{R}_{q_n}^\infty(\mathcal{A})$ and an integer $k \geq 1$, there exists an accepting run θ of \mathcal{A} such that:*

- (1) θ_i is embedded in θ at some position $p_i \in \text{dom}(\theta)$, for each $i \in [1..n]$,
- (2) $p_i \cdot \text{dom}(\theta_i) \cap p_j \cdot \text{dom}(\theta_j) = \emptyset$, for all $1 \leq i < j \leq n$,
- (3) the path between p_i and p_j in θ traverses k times some reset path disjoint from $\bigcup_{\ell=1}^n p_\ell \cdot \text{dom}(\theta_\ell)$, for all $1 \leq i < j \leq n$.

Proof. Let θ be an arbitrary accepting run of \mathcal{A} . By Lemma 4.7, each 1-transition occurs exactly once on θ , hence θ visits each pivot state at least once. The partial runs $\theta_1, \dots, \theta_n$ will be inserted into θ one by one, as described next. First, for each $\theta_i \in \mathcal{R}_{q_i}^\infty(\mathcal{A})$, we have a pivot state q_i^0 and a partial run $\theta_i^0 \in \mathcal{R}_{q_i^0}^\infty(\mathcal{A})$, satisfying condition 1 or 2 of Lemma 4.8. Since q_i^0 occurs on θ , we can insert in θ a new partial run $\theta_i' \in \mathcal{R}_{q_0^i}^\infty(\mathcal{A})$ defined next. By Lemma 4.18 (1), there exists a q_0^i -reset sequence $\theta_r^i \in \mathcal{R}_{q_0^i}^\infty(\mathcal{A})$. The partial run θ_i' is obtained by composing θ_r^i with itself k times, followed by θ_0^i . These compositions are possible, because q_0^i occurs at the root of θ_r^i and θ_0^i , as well as the frontier of θ_r^i . Depending on which condition of Lemma 4.8 is satisfied by q_0^i and θ_0^i , we distinguish the following cases (see Figure 11 for an illustration):

- ▷ condition (1) of Lemma 4.8 holds: in this case q_i and q_0^i occur on different positions on the frontier of θ_0^i , thus we append θ_i on the position where q_i occurs and the rest of θ on the position where q_0^i occurs.
- ▷ condition (2) of Lemma 4.8 holds: in this case only q_0^i occurs on the frontier of θ_0^i , thus we continue with θ_i , which can be extended to reach q_0^i again, by Lemma 4.8. From this second occurrence of q_0^i , we continue with θ .

We prove the points from the statement of the Lemma below:

- (1) The runs $\theta_1, \dots, \theta_n$ are inserted into θ at positions p_1, \dots, p_n , respectively.
- (2) Since $\theta_1, \dots, \theta_n$ are inserted one after the other (the order is not important), we have $p_i \cdot \text{dom}(\theta_i) \cap p_j \cdot \text{dom}(\theta_j) = \emptyset$, for all $1 \leq i < j \leq n$.
- (3) By the definition of $\theta'_1, \dots, \theta'_n$, the path between p_i and p_j traverses k times the θ_r^i or θ_r^j reset sequences that are, moreover, disjoint from each $p_k \cdot \text{dom}(\theta_k)$, for $k \in [1..n]$. \square

4.4. Eliminating Persistent Variables. For the rest of this section, let $\mathcal{A}_{\Delta, \mathbf{A}} = (\Sigma, \mathcal{Q}_{\Delta}, q_{\mathbf{A}}, \delta_{\Delta})$ be the automaton built for the given SID Δ and nullary predicate \mathbf{A} , by the construction of Lemma 4.13. Since Δ was assumed to be all-satisfiable, the same can be assumed about $\mathcal{A}_{\Delta, \mathbf{A}}$, by Lemma 4.13 (1).

Moreover, we can assume, without loss of generality, that $\mathcal{A}_{\Delta, \mathbf{A}}$ is choice-free and let $\delta_{\Delta} = \delta_{\Delta}^1 \uplus \delta_{\Delta}^{\infty}$ be the partition of the transitions of $\mathcal{A}_{\Delta, \mathbf{A}}$ (Definition 4.5). If this is not the case, we consider one of the finitely many automata in the language-preserving choice-free decomposition of $\mathcal{A}_{\Delta, \mathbf{A}}$ (Lemma 4.9).

The transformation proceeds in three stages, denoted (I), (II) and (III) below. The result of each stage is one or more choice-free automata that are treewidth bounded if and only if the set $\llbracket \mathcal{A}_{\Delta, \mathbf{A}} \rrbracket$ is treewidth bounded.

Example 4.21. We shall illustrate each stage of the construction on the following SID:

$$\Delta \left\{ \begin{array}{l} \mathbf{A} \leftarrow \exists y_1 \exists y_2 \exists y_3 \cdot \mathbf{C}_1(y_1, y_2, y_3) \\ \mathbf{C}_1(x_1, x_2, x_3) \leftarrow \exists y_4 \cdot \mathbf{e}(x_1, y_4) * \mathbf{e}(x_3, y_4) * \mathbf{C}_1(y_4, x_2, x_3) \\ \mathbf{C}_1(x_1, x_2, x_3) \leftarrow \exists y_5 \cdot \mathbf{e}(x_1, x_2) * \mathbf{C}_2(x_2, y_5, x_3) \\ \mathbf{C}_2(x_1, x_2, x_3) \leftarrow \exists y_6 \cdot \mathbf{e}(x_1, y_6) * \mathbf{e}(x_3, y_6) * \mathbf{C}_2(y_6, x_2, x_3) \\ \mathbf{C}_2(x_1, x_2, x_3) \leftarrow \mathbf{e}(x_1, x_2) \end{array} \right.$$

For simplicity, the existentially quantified variables are given pairwise distinct names. The automaton $\mathcal{A}_{\Delta, \mathbf{A}}$ has the following transitions:

$$\mathcal{A}_{\Delta, \mathbf{A}} \left\{ \begin{array}{l} \tau_1 : q_{\mathbf{A}} \xrightarrow{y_1^{[\epsilon]=x_1^{[1]} * y_2^{[\epsilon]=x_2^{[1]} * y_3^{[\epsilon]=x_3^{[1]}}} (q_{\mathbf{C}_1}) \quad (1) \\ \tau_2 : q_{\mathbf{C}_1} \xrightarrow{\mathbf{e}(x_1^{[\epsilon]}, y_4^{[\epsilon]}) * \mathbf{e}(x_3^{[\epsilon]}, y_4^{[\epsilon]}) * y_4^{[\epsilon]=x_1^{[1]} * x_2^{[\epsilon]=x_2^{[1]} * x_3^{[\epsilon]=x_3^{[1]}}} (q_{\mathbf{C}_1}) \quad (\infty) \\ \tau_3 : q_{\mathbf{C}_1} \xrightarrow{\mathbf{e}(x_1^{[\epsilon]}, x_2^{[\epsilon]}) * x_2^{[\epsilon]=x_1^{[1]} * y_5^{[\epsilon]=x_2^{[1]} * x_3^{[\epsilon]=x_3^{[1]}}} (q_{\mathbf{C}_2}) \quad (1) \\ \tau_4 : q_{\mathbf{C}_2} \xrightarrow{\mathbf{e}(x_1^{[\epsilon]}, y_6^{[\epsilon]}) * \mathbf{e}(x_3^{[\epsilon]}, y_6^{[\epsilon]}) * y_6^{[\epsilon]=x_1^{[1]} * x_2^{[\epsilon]=x_2^{[1]} * x_3^{[\epsilon]=x_3^{[1]}}} (q_{\mathbf{C}_2}) \quad (\infty) \\ \tau_5 : q_{\mathbf{C}_2} \xrightarrow{\mathbf{e}(x_1^{[\epsilon]}, x_2^{[\epsilon]})} () \quad (1) \end{array} \right.$$

The Δ -models of \mathbf{A} have the structure depicted in Figure 12 (a), with elements that are values of persistent variables annotated by the name of the first occurrence of the persistent variable during the run.




Figure 12: Models of the Running Example for the Elimination of Persistent Variables

I. Removing relation and disequality atoms from 1-transitions. This step replaces each symbol α that labels a 1-transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell)$ of $\mathcal{A}_{\Delta, \mathcal{A}}^I$ with the symbol obtained by removing all relation and disequality atoms from α . The outcome of this transformation is denoted $\mathcal{A}_{\Delta, \mathcal{A}}^I \stackrel{\text{def}}{=} (\Sigma, \mathcal{Q}_{\Delta}, q_{\mathcal{A}}, \delta_{\Delta}^I)$.

Example 4.22. (continued from Example 4.21) The result of removing the relation and disequality atoms from the 1-transitions of the choice-free automaton given in Example 4.21 is shown below:

$$\mathcal{A}_{\Delta, \mathcal{A}}^I \left\{ \begin{array}{l} \tau_1 : q_{\mathcal{A}} \xrightarrow{y_1^{[\epsilon]=x_1^{[1]} * y_2^{[\epsilon]=x_2^{[1]} * y_3^{[\epsilon]=x_3^{[1]}}} (q_{C_1}) \quad (1) \\ \tau_2 : q_{C_1} \xrightarrow{e(x_1^{[\epsilon]}, y_4^{[\epsilon]}) * e(x_3^{[\epsilon]}, y_4^{[\epsilon]}) * y_4^{[\epsilon]=x_1^{[1]} * x_2^{[\epsilon]=x_2^{[1]} * x_3^{[\epsilon]=x_3^{[1]}}} (q_{C_1}) \quad (\infty) \\ \tau_3 : q_{C_1} \xrightarrow{x_2^{[\epsilon]=x_1^{[1]} * y_5^{[\epsilon]=x_2^{[1]} * x_3^{[\epsilon]=x_3^{[1]}}} (q_{C_2}) \quad (1) \\ \tau_4 : q_{C_2} \xrightarrow{e(x_1^{[\epsilon]}, y_6^{[\epsilon]}) * e(x_3^{[\epsilon]}, y_6^{[\epsilon]}) * y_6^{[\epsilon]=x_1^{[1]} * x_2^{[\epsilon]=x_2^{[1]} * x_3^{[\epsilon]=x_3^{[1]}}} (q_{C_2}) \quad (\infty) \\ \tau_5 : q_{C_2} \xrightarrow{\text{emp}} () \quad (1) \end{array} \right.$$

The formal properties of $\mathcal{A}_{\Delta, \mathcal{A}}^I$ are stated and proved below. Note that $\mathcal{A}_{\Delta, \mathcal{A}}^I$ is choice-free, because $\mathcal{A}_{\Delta, \mathcal{A}}$ is choice-free, i.e., the re-labeling of the transitions of $\mathcal{A}_{\Delta, \mathcal{A}}$ does not change the structure of its SCC graph. In the statement of this and the upcoming lemmas, we assume addition and order within the set of natural numbers with infinity, i.e., $n + \infty = \infty + n = \infty$, $n \leq \infty$ and $\infty \leq \infty$, for all $n \in \mathbb{N}$.

Lemma 4.23.

- (1) $\mathcal{A}_{\Delta, \mathcal{A}}^I$ is all-satisfiable
- (2) $\text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}}^I \rrbracket) \leq \text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}} \rrbracket) + \text{card}(\delta_{\Delta}^1) \cdot (\max\text{RelAtoms}(\Delta) + \max\text{Vars}(\Delta))$
- (3) $\text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}} \rrbracket) \leq \text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}}^I \rrbracket) + \text{card}(\delta_{\Delta}^1) \cdot \max\text{Vars}(\Delta)$

Proof. For space reasons, the proof of the lemma is given in subsection C.5. \square

II. Removing equalities involving non-persistent variables. At this point, the labels of the 1-transitions of $\mathcal{A}_{\Delta, \mathcal{A}}^I$ consist of equalities only. We now remove the equalities that would be lost when adding resets before and after 1-transitions that is, we forget equalities involving non-persistent variables while keeping equalities between persistent ones. To this end, we modify the label of each 1-transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell)$ of $\mathcal{A}_{\Delta, \mathcal{A}}^I$ in two steps:

- (1) for each non-persistent ϵ -variable $x_j^{[\epsilon]}$, i.e., $j \in [1.. \#q_0] \setminus \mathfrak{P}_{\mathcal{A}_{\Delta, \mathcal{A}}^I}(q_0)$, occurring in α in some equality with a persistent i -variable, i.e., $x_j^{[\epsilon]} = x_k^{[i]}$, $k \in \mathfrak{P}_{\mathcal{A}_{\Delta, \mathcal{A}}^I}(q_i)$, we substitute $x_j^{[\epsilon]}$ in α with a fresh variable $y_j^{[\epsilon]} \notin \text{fv}(\alpha)$,

- (2) remove each equality involving a non-persistent variable $x_j^{[i]}$, for $i = \epsilon$ and $j \in [1..\#q_0]$, or $i \in [1..\ell]$ and $j \in [1..\#q_i]$.

Example 4.24. (continued from Example 4.22) The profile of the automaton $\mathcal{A}_{\Delta, \mathcal{A}}^I$ from Example 4.22 is $\mathfrak{P}_{\mathcal{A}_{\Delta, \mathcal{A}}^I}(q_{\mathcal{A}}) = \emptyset$ and $\mathfrak{P}_{\mathcal{A}_{\Delta, \mathcal{A}}^I}(q_{C_1}) = \mathfrak{P}_{\mathcal{A}_{\Delta, \mathcal{A}}^I}(q_{C_2}) = \{2, 3\}$. By removing the equalities involving non-persistent variables, we obtain $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ from $\mathcal{A}_{\Delta, \mathcal{A}}^I$ by transforming only the 1-transition τ_3 into:

$$\tau'_3 : q_{C_1} \xrightarrow{y_5^{[\epsilon]=x_2^{[1]} * x_3^{[\epsilon]=x_3^{[1]}}} (q_{C_2}) \quad (1)$$

Since $x_1^{[1]}$ is non-persistent, the equality $x_2^{[\epsilon]} = x_1^{[1]}$ is removed from the label of the original transition. Note that the equalities $y_5^{[\epsilon]} = x_2^{[1]}$ and $x_3^{[\epsilon]} = x_3^{[1]}$ are kept, because $x_2^{[1]}$, $x_3^{[\epsilon]}$ and $x_3^{[1]}$ are persistent variables. Note that $y_5^{[\epsilon]}$ is not persistent, because it is not associated with a state in τ'_3 .

The result is the choice-free automaton $\mathcal{A}_{\Delta, \mathcal{A}}^{II} \stackrel{\text{def}}{=} (\Sigma, \mathcal{Q}_{\Delta}, q_{\mathcal{A}}, \delta_{\Delta}^{II})$, whose properties are stated and proved below:

Lemma 4.25.

- (1) $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ is all-satisfiable,
- (2) $\text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}}^{II} \rrbracket) \leq \text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}}^I \rrbracket) + \text{card}((\delta_{\Delta}^I)^1) \cdot \text{maxVars}(\Delta)$,
- (3) $\text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}}^I \rrbracket) \leq \text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}}^{II} \rrbracket)$.

Proof. For space reasons, the proof of this lemma is given in subsection C.6. \square

III. Removing persistent variables. We build from the choice-free automaton $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ a set of choice-free automata $\tilde{\mathcal{B}}_1, \dots, \tilde{\mathcal{B}}_m$, having no persistent variables within the transition labels, such that $\llbracket \mathcal{A}_{\Delta, \mathcal{A}}^{II} \rrbracket$ is treewidth bounded if and only if $\llbracket \tilde{\mathcal{B}}_i \rrbracket$ is treewidth bounded, for each $i \in [1..m]$.

We recall that each 1-transition of a choice free automata occurs exactly once in each accepting run over a Σ -labeled tree t and each such occurrence corresponds to one subformula $t(p)^p$ of $\Theta(t)$, for a position $p \in \text{dom}(t)$ (Lemma 4.7). Using renaming, if necessary, we can assume that the ϵ -variables that are not associated with the states of the transition have distinct names between the labels of the 1-transitions of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ and let $\mathcal{Y} \stackrel{\text{def}}{=} \{y_1^{[\epsilon]}, \dots, y_{\mathcal{M}}^{[\epsilon]}\}$ denote their set, in the following. For instance, we have $\mathcal{Y} = \{y_1^{[\epsilon]}, y_2^{[\epsilon]}, y_3^{[\epsilon]}, y_5^{[\epsilon]}\}$ in Example 4.22. The transformation is done in three steps:

- (A) We annotate each state q of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ with an injective partial function $a : [1..\#q] \rightarrow [1..\mathcal{M}]$ that maps each persistent variable $x_j^{[i]}$, associated with q , to a variable $y_{a(j)}^{[\epsilon]} \in \mathcal{Y}$, such that $x_j^{[i]} \approx_{\Theta(t)} y_{a(j)}^{[\epsilon]}$ holds for each tree $t \in \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^{II})$.
- (B) We split the automaton obtained from the annotation of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ into several choice-free automata $\tilde{\mathcal{B}}_1, \dots, \tilde{\mathcal{B}}_m$ such that $\mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^{II}) = \bigcup_{i=1}^m \mathcal{L}(\tilde{\mathcal{B}}_i)$.
- (C) The annotation of the states in each $\tilde{\mathcal{B}}_i$ is used to replace each occurrence of a relation atom $r(z_1^{[\epsilon]}, \dots, z_{\#r}^{[\epsilon]})$, occurring within the label of an annotated transition $(q_0, a_0) \xrightarrow{\epsilon} ((q_1, a_1), \dots, (q_\ell, a_\ell))$, with a fresh relation atom $r_g(z_{i_1}^{[\epsilon]}, \dots, z_{i_k}^{[\epsilon]})$, where $g : [1..\#r] \rightarrow [1..\mathcal{M}]$ maps each persistent variable from the set $\{z_1^{[\epsilon]}, \dots, z_{\#r}^{[\epsilon]}\}$ to its corresponding

variable from \mathcal{V} and $\{z_{i_1}^{[\epsilon]}, \dots, z_{i_k}^{[\epsilon]}\}$ are the remaining non-persistent variables. The persistent variables are subsequently removed from α and the remaining variables are renamed according to the conventions from Definition 4.11. Consequently, the arities of the states (q_i, a_i) , $i \in [0..\ell]$ are changed as well.

Example 4.26. (continued from Example 4.24) Let us consider the automaton $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$ from Example 4.24. The ϵ -variables from the labels of 1-transitions, that are not associated with states thereof, are $y_1^{[\epsilon]}, y_2^{[\epsilon]}, y_3^{[\epsilon]}, y_5^{[\epsilon]}$, renamed as $\mathcal{V} = \{y_1^{[\epsilon]}, y_2^{[\epsilon]}, y_3^{[\epsilon]}, y_4^{[\epsilon]}\}$, respectively. We recall that $\mathfrak{P}_{\mathcal{A}_{\Delta, \mathbf{A}}^{II}}(q_{\mathbf{A}}) = \emptyset$ and $\mathfrak{P}_{\mathcal{A}_{\Delta, \mathbf{A}}^{II}}(q_{\mathbf{C}_1}) = \mathfrak{P}_{\mathcal{A}_{\Delta, \mathbf{A}}^{II}}(q_{\mathbf{C}_2}) = \{2, 3\}$ is the profile of $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$. The automaton obtained by annotating the states of $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$ with assignments is already choice-free and the result of the elimination of persistent variables is shown below:

$$\tilde{\mathcal{B}} \left\{ \begin{array}{l} \tau_1 : (q_{\mathbf{A}}, \emptyset) \xrightarrow{\text{emp}} ((q_{\mathbf{C}_1}, a)) \quad (1) \\ \tau_2 : (q_{\mathbf{C}_1}, a) \xrightarrow{e(x_1^{[\epsilon]}, y_4^{[\epsilon]}) * e_g(y_4^{[\epsilon]}) * y_4^{[\epsilon]} = x_1^{[1]}} ((q_{\mathbf{C}_1}, a)) \quad (\infty) \\ \tau_3 : (q_{\mathbf{C}_1}, a) \xrightarrow{\text{emp}} ((q_{\mathbf{C}_2}, a)) \quad (1) \\ \tau_4 : (q_{\mathbf{C}_2}, a) \xrightarrow{e(x_1^{[\epsilon]}, y_6^{[\epsilon]}) * e_g(y_6^{[\epsilon]}) * y_6^{[\epsilon]} = x_1^{[1]}} ((q_{\mathbf{C}_2}, a)) \quad (\infty) \\ \tau_5 : (q_{\mathbf{C}_2}, a) \xrightarrow{\text{emp}} () \quad (1) \end{array} \right.$$

where:

- ▷ $a : [1..3] \rightarrow [1..4]$ is the partial mapping defined as $a(2) = 2$, $a(3) = 3$ and undefined at 1,
- ▷ $g : [1..2] \rightarrow [1..4]$ is the partial mapping defined as $g(1) = 3$ and undefined at 2.

Note that, because the equality between the persistent variables $x_3^{[\epsilon]}$ and $x_3^{[1]}$ has been kept in $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$ (see Example 4.24), both variables are mapped by g to the same variable $y_3^{[\epsilon]}$, hence the same relation symbol e_g replaces both $e(x_3^{[\epsilon]}, y_4^{[\epsilon]})$ in τ_2 and $e(x_3^{[\epsilon]}, y_6^{[\epsilon]})$ in τ_4 . Figure 12 (b) shows the shape of the structures from $[[\tilde{\mathcal{B}}]]$. Since all but the first elements in both the \mathbf{C}_1 and \mathbf{C}_2 chains are now labeled with the same unary relation symbol e_g , these structures are of treewidth at most two.

We recall that $\mathcal{A}_{\Delta, \mathbf{A}}^{II} = (\Sigma, \mathcal{Q}_{\Delta}^I, q_{\mathbf{A}}, \delta_{\Delta}^{II})$ and describe the transformation formally:

(A) Let $\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II} \stackrel{\text{def}}{=} (\Sigma, \tilde{\mathcal{Q}}_{\Delta}^I, (q_{\mathbf{A}}, \emptyset), \tilde{\delta}_{\Delta}^{II})$ be the automaton, whose set of states is:

$$\tilde{\mathcal{Q}}_{\Delta}^I \stackrel{\text{def}}{=} \{(q, a) \mid q \in \mathcal{Q}_{\Delta}^I, a : [1..\#q] \rightarrow [1..\mathcal{M}] \text{ is a partial injective mapping}\}$$

The initial state of $\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}$ consists of the initial state $q_{\mathbf{A}}$ of $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$ annotated with the empty mapping, because we have considered $\#q_{\mathbf{A}} = 0$. The set $\tilde{\delta}_{\Delta}^{II}$ contains a transition $(q_0, a_0) \xrightarrow{\alpha} ((q_1, a_1), \dots, (q_{\ell}, a_{\ell}))$ if and only if either one of the following holds (by Definition 4.5, these conditions are exclusive):

- ▷ $q_0 \xrightarrow{\alpha} (q_1, \dots, q_{\ell}) \in (\delta_{\Delta}^{II})^1$: in this case, for all $k \in [1..\ell]$ and $i \in \mathfrak{P}_{\mathcal{A}_{\Delta, \mathbf{A}}^{II}}(q_k)$, we define:

$$a_k(i) \stackrel{\text{def}}{=} \begin{cases} a_0(j) & \text{if there exists } j \text{ such that } x_i^{[k]} \approx_{\alpha} x_j^{[\epsilon]}, \\ m & \text{else, if } m \text{ is such that } x_i^{[k]} \approx_{\alpha} y_m^{[\epsilon]} \end{cases}$$

Note that a_k is well defined, because each i -variable is equated to a unique ϵ -variable in the definition of $\mathcal{A}_{\Delta, \mathbf{A}}$ and this fact is unchanged by the constructions of $\mathcal{A}_{\Delta, \mathbf{A}}^I$ and $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$.

- ▷ $q_0 \xrightarrow{\alpha} (q_1, \dots, q_{\ell}) \in (\delta_{\Delta}^{II})^{\infty}$: in this case, for all $k \in [1..\ell]$ and $i \in [1..\#q_k]$, we define $a_k(i) \stackrel{\text{def}}{=} a_0(j)$ if there exists j , such that $x_i^{[k]} \approx_{\alpha} x_j^{[\epsilon]}$ and undefined, otherwise.

The property of the first step of persistent variable elimination is summarized below:

Lemma 4.27. $\mathcal{L}(\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}) = \mathcal{L}(\mathcal{A}_{\Delta, \mathbf{A}}^{II})$.

Proof. Let $h : \widetilde{\mathcal{Q}}_{\Delta}^I \rightarrow \mathcal{Q}_{\Delta}^I$ be the function defined as $h((q, a)) \stackrel{\text{def}}{=} q$, for all $(q, a) \in \widetilde{\mathcal{Q}}_{\Delta}^I$. We show that h is a refinement. By Lemma 4.4, we obtain $\mathcal{L}(\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}) = \mathcal{L}(\mathcal{A}_{\Delta, \mathbf{A}}^{II})$. We prove the three points of Definition 4.3:

- (1) $h((q_{\mathbf{A}}, \emptyset)) = q_{\mathbf{A}}$, by the definition of h .
- (2) $(q_0, a_0) \xrightarrow{\alpha} ((q_1, a_1), \dots, (q_{\ell}, a_{\ell})) \in \widetilde{\delta}_{\Delta}^{II}$ only if $q_0 \xrightarrow{\alpha} (q_1, \dots, q_{\ell}) \in \delta_{\Delta}^{II}$, by the definition of $\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}$.
- (3) Let $q_0 \xrightarrow{\alpha} (q_1, \dots, q_{\ell}) \in \delta_{\Delta}^{II}$ be a transition and assume w.l.o.g. that $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$ is trim (if this is not the case, a trim automaton with the same language can be considered instead). Then, there exists an accepting run θ of $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$, such that the transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_{\ell})$ occurs on θ . Each state q on θ can be annotated with an injective partial mapping $a : [1..\#q] \rightarrow [1..\mathcal{M}]$ and the result is an accepting run of $\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}$. Hence there exist injective partial functions $a_i : [1..\#q_i] \rightarrow [1..\mathcal{M}]$, for $i \in [0..\ell]$ such that $(q_0, a_0) \xrightarrow{\alpha} ((q_1, a_1), \dots, (q_{\ell}, a_{\ell})) \in \widetilde{\delta}_{\Delta}^{II}$. Moreover, $(q_i, a_i) \in h^{-1}(q_i)$, for each $i \in [0..\ell]$. \square

(B) The problem, at this point, is that $\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}$ is not necessarily choice-free, because annotating the states of $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$ may cause several transitions to occur between different linear SCCs. These transitions originate from the same 1-transition of $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$ and differ only in the annotations added at step (A). We circumvent this problem by decomposing $\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}$ into choice-free automata $\widetilde{\mathcal{B}}_1, \dots, \widetilde{\mathcal{B}}_m$, such that $\mathcal{L}(\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}) = \bigcup_{i=1}^m \mathcal{L}(\widetilde{\mathcal{B}}_i)$. To this end, we choose sets $\widetilde{\delta}_1^1, \dots, \widetilde{\delta}_m^1$, such that:

- $\triangleright (\widetilde{\delta}^{II})^1 = \bigcup_{i=1}^m \widetilde{\delta}_i^1$, and
- \triangleright for each $i \in [1..m]$, the set $\widetilde{\delta}_i^1$ contains exactly one transition $(q_0, a_0) \xrightarrow{\alpha} ((q_1, a_1), \dots, (q_{\ell}, a_{\ell}))$ from $\widetilde{\delta}_{\Delta}^{II}$, for each transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_{\ell})$ from $(\delta_{\Delta}^{II})^1$.

Moreover, we define the set:

$$(\widetilde{\delta})^{\infty} \stackrel{\text{def}}{=} \{(q_0, a_0) \xrightarrow{\alpha} ((q_1, a_1), \dots, (q_{\ell}, a_{\ell})) \in \widetilde{\delta}_{\Delta}^{II} \mid q_0 \xrightarrow{\alpha} (q_1, \dots, q_{\ell}) \in (\delta_{\Delta}^{II})^{\infty}\}$$

For each $i \in [1..m]$, let $\widetilde{\mathcal{B}}_i \stackrel{\text{def}}{=} (\Sigma, \widetilde{\mathcal{Q}}_{\Delta}^I, (q_{\mathbf{A}}, \emptyset), \widetilde{\delta}_i^1 \uplus \widetilde{\delta}^{\infty})$. We prove below that $\widetilde{\mathcal{B}}_1, \dots, \widetilde{\mathcal{B}}_m$ is indeed a choice-free decomposition of $\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}$:

Lemma 4.28.

- (1) $\widetilde{\mathcal{B}}_i$ is all-satisfiable and choice-free, for $i \in [1..m]$.
- (2) $\mathcal{L}(\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}) = \bigcup_{i=1}^m \mathcal{L}(\widetilde{\mathcal{B}}_i)$.

Proof. For space reasons, the proof of this lemma is given in subsection C.7. \square

(C) Let us fix a choice-free automaton $\widetilde{\mathcal{B}} = (\Sigma, \widetilde{\mathcal{Q}}_{\Delta}^I, (q_{\mathbf{A}}, \emptyset), \widetilde{\delta}^1 \uplus \widetilde{\delta}^{\infty})$ among $\widetilde{\mathcal{B}}_1, \dots, \widetilde{\mathcal{B}}_m$. Consider an arbitrary transition $\tau : (q_0, a_0) \xrightarrow{\alpha} ((q_1, a_1), \dots, (q_{\ell}, a_{\ell})) \in \widetilde{\delta}^1 \uplus \widetilde{\delta}^{\infty}$. We denote by $P_0 \stackrel{\text{def}}{=} \{x_j^{[c]} \mid j \in \text{dom}(a_0)\} \cup \{y_j^{[c]} \in \text{fv}(\alpha) \mid \tau \in \widetilde{\delta}^1\}$, $P_i \stackrel{\text{def}}{=} \{x_j^{[i]} \mid j \in \text{dom}(a_i)\}$, for $i \in [1..\ell]$ and $P \stackrel{\text{def}}{=} \bigcup_{i=0}^{\ell} P_i$ the set of persistent variables occurring in α . The goal of this step is

to remove from α the variables from P . The outcome of this transformation of α will be denoted as $\bar{\alpha}$. In order to guarantee the preservation of the naming conventions from Definition 4.11, the remaining (non-persistent) variables from $\text{fv}(\alpha) \setminus P$ are renamed using the injective mapping $\zeta_\tau : \text{fv}(\alpha) \setminus P \rightarrow \text{fv}(\alpha)$ defined as follows:

- ▷ $\zeta_\tau(x_k^{[i]}) \stackrel{\text{def}}{=} x_{k-p}^{[i]}$ where $p \stackrel{\text{def}}{=} \text{card}(\{x_j^{[i]} \mid j < k\} \cap P_i)$ for all $i \in [0..\ell]$, $k \in [1..n_i]$, $x_k^{[i]} \notin P_i$,
- ▷ $\zeta_\tau(y_k^{[\epsilon]}) = y_k^{[\epsilon]}$, for all $k \in [1..m]$, $y_k^{[\epsilon]} \notin P_0$

where m, n_0, \dots, n_ℓ are the numbers of $y_j^{[\epsilon]}$, $x_j^{[\epsilon]}$, \dots , $x_j^{[i]}$, respectively (see Definition 4.11). Intuitively, the renaming *shifts to the left* the j indices of the $x_j^{[i]}$ variables so that the persistent variables indexed according to the assignments a_i are ignored. The transformation of α to $\bar{\alpha}$ is described for each atom of α , as follows:

- ▷ every relation atom $r(z_1^{[\epsilon]}, \dots, z_{\#r}^{[\epsilon]})$ is replaced by $r_g(\zeta_\tau(z_{i_1}^{[\epsilon]}), \dots, \zeta_\tau(z_{i_k}^{[\epsilon]}))$, where r_g is a fresh relation symbol of arity k such that $\{z_{i_1}, \dots, z_{i_k}\} \stackrel{\text{def}}{=} \{z_1, \dots, z_{\#r}\} \setminus P_0$ and $g : [1..\#r] \rightarrow [1..\mathcal{M}]$ is the partial function:

$$g(j) \stackrel{\text{def}}{=} \begin{cases} a_0(j), & \text{if } z_j^{[\epsilon]} \in P_0 \\ \text{undefined}, & \text{otherwise} \end{cases}$$

Note that the arity of r_g is at least one, for the following reason. Since relation atoms occur only in ∞ -transitions, they can repeat arbitrary many times in characteristic formulæ over runs. Hence, if such an atom has only persistent variables, these characteristic formulæ will become unsatisfiable if the atom repeat more than twice. This contradicts, however, the hypothesis that $\tilde{\mathcal{B}}_i$ is all-satisfiable.

- ▷ every (dis-)equality atom $x \sim y$ is replaced by $\zeta_\tau(x) \sim \zeta_\tau(y)$ for $\sim \in \{=, \neq\}$ if $\{x, y\} \cap P = \emptyset$ and removed otherwise. In particular, note that there is no equality in α between a variable in P and another one not in P due to elimination of equalities with non-persistent variables in 1-transitions and to the rule of propagation through ∞ -transitions. Moreover, this transformation turns the label of each 1-transition into **emp**, because the labels of 1-transitions contain only equalities involving persistent variables.

The result of this transformation of $\tilde{\mathcal{B}}_i$ is denoted $\bar{\mathcal{B}}_i$, for each $i \in [1..m]$. Let $\bar{\mathcal{B}}$ be any of $\bar{\mathcal{B}}_1, \dots, \bar{\mathcal{B}}_m$ and $\tilde{\mathcal{B}}$ be the corresponding automaton before the removal of persistent variables. The properties of this transformation are stated and proved below:

Lemma 4.29.

- (1) $\bar{\mathcal{B}}$ is all-satisfiable and choice-free,
- (2) $\text{tw}(\llbracket \bar{\mathcal{B}} \rrbracket) \leq \text{tw}(\llbracket \tilde{\mathcal{B}} \rrbracket)$,
- (3) $\text{tw}(\llbracket \tilde{\mathcal{B}} \rrbracket) \leq \text{tw}(\llbracket \bar{\mathcal{B}} \rrbracket) + \text{card}(\tilde{\delta}^1) \cdot \text{maxVars}(\Delta)$.

Proof. For space reasons, the proof of this lemma is given in subsection C.8. □

4.5. Wrapping 1-transitions. In subsection 4.4 we have transformed a given automaton $\mathcal{A}_{\Delta, A}$ into choice-free automata $\bar{\mathcal{B}}_1, \dots, \bar{\mathcal{B}}_m$ without persistent variables. We can assume w.l.o.g. that the 1-transitions of these automata are labeled by **emp**, because the only remaining equalities are between variables $y^{[\epsilon]}$ not associated with states and non-persistent variables $x_j^{[i]}$. In the following, let $\bar{\mathcal{B}} = (\Sigma, \mathcal{Q}, \iota, \bar{\delta})$ be any of $\bar{\mathcal{B}}_1, \dots, \bar{\mathcal{B}}_m$ and $\bar{\delta} = \bar{\delta}^1 \uplus \bar{\delta}^\infty$ be the partition of its transitions into 1- and ∞ -transitions (Definition 4.5). In order to obtain an expandable SID from $\bar{\mathcal{B}}$, i.e., using Lemma 4.13 (2), we must be able to embed

any sequence of runs of $\overline{\mathcal{B}}$ into a single run of $\overline{\mathcal{B}}$. However, this is currently not possible, because the labels of the 1-transitions cannot be viewed as labels of ∞ -transitions. The problem is shown by the following example:

Example 4.30. Consider two disjoint structures $\mathcal{S}_1 = (\mathcal{U}_1, \sigma_1), \mathcal{S}_2 = (\mathcal{U}_2, \sigma_2) \in \llbracket \overline{\mathcal{B}} \rrbracket^c$, having the shape depicted in Figure 12 (b). Each \mathcal{S}_i has all but two elements, call them e_i^1 and e_i^2 , labeled with a unary relation symbol e_g . Then, the composition $\mathcal{S}_1 \bullet \mathcal{S}_2$ will have four unlabeled elements e_1^1, e_1^2, e_2^1 and e_2^2 . Since any structure $\mathcal{S} = (\mathcal{U}, \sigma) \in \llbracket \overline{\mathcal{B}} \rrbracket^c$ has exactly two unlabeled elements, the structure $\mathcal{S}_1 \bullet \mathcal{S}_2$ is not a substructure of \mathcal{S} , according to Definition 3.7. This is because $e_1^1, e_1^2, e_2^1, e_2^2 \in (\text{supp}(\sigma_1) \cup \text{supp}(\sigma_2)) \setminus (\sigma_1(e_g) \uplus \sigma_2(e_g))$ are pairwise distinct, hence $\text{card}((\text{supp}(\sigma_1) \cup \text{supp}(\sigma_2)) \setminus (\sigma_1(e_g) \uplus \sigma_2(e_g))) = 4$, whereas $\text{card}(\text{supp}(\sigma) \setminus \sigma(e_g)) = 2$, independently of the choice of \mathcal{S} .

For a Σ -labeled tree t , two positions p and s , such that $p \in \text{dom}(t)$ (nothing is required about s), and a sequence of variables x_1, \dots, x_k , we define the formula:

$$\Omega_t^{p/s}(x_1, \dots, x_k) \stackrel{\text{def}}{=} * \{r(x_1^{[s]}, \dots, x_k^{[s]}) \mid r(x_1^{[p]}, \dots, x_k^{[p]}) \text{ occurs in } \Theta(t), x_i^{[p]} \approx_{\Theta(t)} x_i^{[s]}, \forall i \in [1..k]\}$$

For simplicity, assume, for each transition $q_0 \xrightarrow{\text{emp}} (q_1, \dots, q_\ell) \in \overline{\delta}^1$ of $\overline{\mathcal{B}}$, that q_0, \dots, q_ℓ belong to non-trivial SCCs. If some q_i belongs to a trivial SCC, i.e., with no transitions from q_i to itself or other states in the same SCC, then $\#q_i$ must be zero, because all parameters from its profile have been removed by the previous transformation. In this case, the construction below is adapted by replacing the formulæ $\Omega_{t_0}^{\epsilon/\epsilon}$ and $\Omega_{t_j}^{p_j/j}$ by **emp**, for the corresponding states that belong to trivial SCCs. The automaton \mathcal{B} is obtained from $\overline{\mathcal{B}}$ by replacing the label of each such 1-transition with the following formula, for some trees t_i corresponding to resets $\theta_{p_i \leftarrow q_i}^i$ of $\overline{\mathcal{B}}$, for $i \in [0..\ell]$:

$$q_0 \xrightarrow{*_{i_1, \dots, i_k \in [1..\#q_0]} \Omega_{t_0}^{\epsilon/\epsilon}(x_{i_1}, \dots, x_{i_k}) * *_{j \in [1..\ell], i_1, \dots, i_k \in [1..\#q_j]} \Omega_{t_j}^{p_j/j}(x_{i_1}, \dots, x_{i_k})} (q_1, \dots, q_\ell)$$

Note that the existence of such resets is guaranteed by Lemma 4.18 and the previous assumption. This construction is illustrated by the following example:

Example 4.31. (continued from Example 4.26) The outcome of applying the above transformation to the automaton $\overline{\mathcal{B}}$ from Example 4.26 is the automaton \mathcal{B} given below:

$$\mathcal{B} = \begin{cases} \tau_1 : (q_A, \emptyset) \xrightarrow{e_g(x_1^{[1]})} ((q_{C_1}, a)) & (1) \\ \tau_2 : (q_{C_1}, a) \xrightarrow{e(x_1^{[\epsilon]}, y_4^{[\epsilon]}) * e_g(y_4^{[\epsilon]}) * y_4^{[\epsilon]} = x_1^{[1]}} ((q_{C_1}, a)) & (\infty) \\ \tau_3 : (q_{C_1}, a) \xrightarrow{e_g(x_1^{[1]})} ((q_{C_2}, a)) & (1) \\ \tau_4 : (q_{C_2}, a) \xrightarrow{e(x_1^{[\epsilon]}, y_6^{[\epsilon]}) * e_g(y_6^{[\epsilon]}) * y_6^{[\epsilon]} = x_1^{[1]}} ((q_{C_2}, a)) & (\infty) \\ \tau_5 : (q_{C_2}, a) \xrightarrow{\text{emp}} () & (1) \end{cases}$$

The transformation adds the relation atoms $e_g(x_1^{[1]})$ to the labels of τ_1 and τ_3 , respectively. Note that τ_2 and τ_4 are (q_{C_1}, a) - and (q_{C_2}, a) -resets, respectively. The added relation atoms correspond to the relation atoms that label the $x_1^{[1]}$ variable within the labels of these resets, taken backwards. All elements from a structure $\mathcal{S} \in \llbracket \mathcal{B} \rrbracket^c$ are now labeled by e_g , hence the composition of any sequence of structures from $\llbracket \mathcal{B} \rrbracket^c$ can be embedded as a substructure of a structure from the same set. In other words, the SID obtained from \mathcal{B} using Lemma 4.13 (2) is expandable (Definition 3.8).

The formal properties of \mathcal{B} are stated and proved below:

Lemma 4.32.

- (1) \mathcal{B} is all-satisfiable,
- (2) $\text{tw}(\llbracket \mathcal{B} \rrbracket) \leq \text{tw}(\llbracket \overline{\mathcal{B}} \rrbracket) + \text{card}(\overline{\delta}^1) \cdot \text{maxVars}(\Delta)$,
- (3) $\text{tw}(\llbracket \overline{\mathcal{B}} \rrbracket) \leq \text{tw}(\llbracket \mathcal{B} \rrbracket) + \text{card}(\overline{\delta}^1) \cdot (1 + \text{maxPredAtoms}(\Delta)) \cdot \text{relNo}(\Delta) \cdot \text{maxVars}(\Delta)^{\text{maxRelArity}(\Delta)}$.

Proof. For space reasons, the proof of this lemma is given in subsection C.9. □

4.6. The Proof of Lemma 4.1. We have collected all the ingredients needed to prove the decidability of the treewidth boundedness problem, for SLR formulæ interpreted over general SIDs. The key point is the proof of Lemma 4.1, the main result being an immediate consequence of this lemma and Theorem 3.33. To prove Lemma 4.1, we rely on Lemma 4.20, which states that any sequence of ∞ -runs of a choice-free automaton can be disjointly embedded into an accepting run of the same automaton, such that any two of the embedded runs are separated by a given number of reset paths in the enclosing run.

Before proving Lemma 4.1, we state a property of the outcome of the above transformation of automata (Figure 14). This property uses the following notion:

Definition 4.33. Let \mathcal{A} be a choice-free automaton. A *view* for \mathcal{A} is a tuple $\langle \theta, t, \mathfrak{s}, \mathbf{S} \rangle$, such that $\theta \in \mathcal{R}_q^\infty(\mathcal{A})$ is a partial ∞ -run over a Σ -labeled tree t , \mathfrak{s} is a canonical store for $\Theta(t)$ and \mathbf{S} is a structure, such that $\mathbf{S} \models^{\mathfrak{s}} \Theta(t)$. A structure $\mathbf{S}' = (\mathbf{U}', \sigma')$ is *encapsulated* by the view $\langle \theta, t, \mathfrak{s}, \mathbf{S} \rangle$ if (i) $\mathbf{S}' \sqsubseteq \mathbf{S}$, (ii) $\text{supp}(\sigma') \cap \mathfrak{s}(\{x_1^{[\epsilon]}, \dots, x_{\#\theta(\epsilon)}^{[\epsilon]}\}) = \emptyset$, and (iii) $\text{supp}(\sigma') \cap \mathfrak{s}(\{x_1^{[p]}, \dots, x_{\#\theta}^{[p]}\}) = \emptyset$ if, moreover, the partial run θ is a $\theta_{p \leftarrow q}$ context.

Intuitively, a structure is encapsulated by a view if it is a substructure of the structure in the view and it does not “touch” the values of the variables from the root (resp. frontier) point of the partial run in the view.

Let \mathcal{B} be any choice-free automaton resulting from the transformation in Figure 14. The following lemma shows that each canonical model from $\llbracket \mathcal{B} \rrbracket^c$ can be decomposed into pairwise disjoint structures, each of which being encapsulated by a separate view for \mathcal{B} :

Lemma 4.34. *For each structure $\mathbf{S} = (\mathbf{U}, \sigma) \in \llbracket \mathcal{B} \rrbracket^c$, there exist pairwise disjoint structures $\mathbf{S}_1, \dots, \mathbf{S}_n$ and views $\langle \theta_1, t_1, \mathfrak{s}_1, \mathbf{S}'_1 \rangle, \dots, \langle \theta_n, t_n, \mathfrak{s}_n, \mathbf{S}'_n \rangle$ for \mathcal{B} such that $\mathbf{S}'_1, \dots, \mathbf{S}'_n$ are pairwise disjoint, $\mathbf{S} = \mathbf{S}_1 \bullet \dots \bullet \mathbf{S}_n$ and \mathbf{S}_i is encapsulated by $\langle \theta_i, t_i, \mathfrak{s}_i, \mathbf{S}'_i \rangle$, for each $i \in [1..n]$.*

Proof. Let $\mathcal{B} \stackrel{\text{def}}{=} (\Sigma, \mathcal{Q}, \iota, \delta)$. Since $\mathbf{S} = (\mathbf{U}, \sigma) \in \llbracket \mathcal{B} \rrbracket^c$, there exists a tree $t_0 \in \mathcal{L}(\mathcal{B})$ such that $\mathbf{S} \models^{\mathfrak{s}_0} \Theta(t_0)$ for a store \mathfrak{s}_0 , that is canonical for $\Theta(t_0)$. Let θ_0 be an accepting run of \mathcal{B} over t_0 . Because \mathcal{B} is choice-free, θ_0 can be decomposed into (i) maximal partial runs $\theta_{01} \in \mathcal{R}_{r_1}^\infty(\mathcal{B}), \dots, \theta_{0n} \in \mathcal{R}_{r_n}^\infty(\mathcal{B})$ consisting of (arbitrarily many) connected ∞ -transitions, and (ii) partial runs τ_1, \dots, τ_m consisting of a single 1-transition each, such that r_i is the state at the root of θ_{0i} and q_i is the left-hand side of a 1-transition, for each $i \in [1..n]$. We refer to Figure 13 (a) for an illustration of the decomposition.

For every $i \in [1..n]$ we define \mathbf{S}_i as the substructure of \mathbf{S} constructed along the maximal partial run θ_{0i} . That is, \mathbf{S}_i contains all the relation atoms defined on ∞ -transitions in θ_{0i} and the relation atoms defined on the entering (and possibly exiting) 1-transition(s) involving common variables for entering (resp. exiting) state(s). Intuitively, all these relation atoms occur in the gray part in Figure 13 (a). Note that, since the 1-transitions of \mathcal{B} do not equate variables $x_j^{[i]}$, for $i \in \mathbb{N} \cup \{\epsilon\}$, the structures \mathbf{S}_i are pairwise disjoint, and $\mathbf{S} = \mathbf{S}_1 \bullet \dots \bullet \mathbf{S}_n$.




Figure 13: Run decomposition

We build the views $\{\langle \theta_i, t_i, \mathfrak{s}_i, S'_i \rangle\}_{i \in [1..n]}$ as follows. For every $i \in [1..n]$, we define the partial run θ_i by extending the partial run θ_{0i} by the same r_i - and q_i -resets used in the transformation of $\bar{\mathcal{B}}$ into \mathcal{B} (note that the q_i -reset is needed only if θ_{0i} reaches a 1-transition). This construction is illustrated in Figure 13 (b). Let t_i be the Σ -labeled tree corresponding to the run θ_i . Let \mathfrak{s}_i be a canonical store for $\Theta(t_i)$ constructed by extending \mathfrak{s}_0 such that $\mathfrak{s}_0(x^{[p]}) = \mathfrak{s}_i(x^{[p']})$ whenever p and p' correspond to the same relative position within θ_{0i} . In other words, \mathfrak{s}_i preserves the elements occurring in S_i and associates new distinct elements to all other variables introduced by the resets (in particular, all variables that are equated in $\Theta(t_i)$ are mapped to the same element). The store \mathfrak{s}_i defines a structure $S'_i = (\mathbf{U}, \sigma'_i)$, such that $S'_i \models^{\mathfrak{s}_i} \Theta(t_i)$, as $\sigma'_i(r) \stackrel{\text{def}}{=} \{\langle \mathfrak{s}_i(z_1^{[p_1]}), \dots, \mathfrak{s}_i(z_k^{[p_k]}) \rangle \mid r(z_1^{[p_1]}, \dots, z_k^{[p_k]}) \text{ occurs in } \Theta(t_i)\}$, for each relation symbol $r \in \mathbb{R}$. We prove below that S_i is encapsulated by the view $\langle \theta_i, t_i, \mathfrak{s}_i, S'_i \rangle$:

- ▷ By the choice of \mathfrak{s}_i , that extends \mathfrak{s}_0 as explained above, we have $S_i \subseteq S'_i$. Moreover, no tuples are added to S_i by the inserted resets, because the labeling of 1-transitions in \mathcal{B} is constructed precisely from the resets that guarantee this property, i.e., the construction of \mathcal{B} guarantees that the set of relation atoms occurring in the reset was used to label the 1-transitions. Thus, we have $S_i \sqsubseteq S'_i$.
- ▷ By the construction of the partial run θ_i , the set of variables $\{x_1^{[e]}, \dots, x_{\#r_i}^{[e]}\}$ in θ_i are not related by equalities to any of the variables at the root of θ_{0i} in $\Theta(t_i)$. In particular, this is ensured by the fact that \mathcal{B} has no persistent variables. Since, \mathfrak{s}_i is canonical for $\Theta(t_i)$, we obtain that $\text{supp}(\sigma_i) \cap \{\mathfrak{s}_i(\{x_1^{[e]}, \dots, x_{\#r_i}^{[e]}\})\} = \emptyset$. The same argument applies to the set of variables occurring at the frontier position of θ_i , i.e., when θ_{0i} reaches the left-hand side q_i of a 1-transition, as in Figure 13 (b). \square

Dually, the following lemma gathers pairwise disjoint structures into a single rich canonical model, that meets the conditions of expandability (Definition 3.8):

Lemma 4.35. *Let S_1, \dots, S_n be pairwise disjoint structures encapsulated by the views $\langle \theta_1, t_1, \mathfrak{s}_1, S'_1 \rangle, \dots, \langle \theta_n, t_n, \mathfrak{s}_n, S'_n \rangle$ for \mathcal{B} , where S'_1, \dots, S'_n are also pairwise disjoint. Then, there exists a rich canonical model $(S, \mathfrak{d}) \in \llbracket \mathcal{B} \rrbracket^r$, such that the conditions (1), (2) and (3) from Definition 3.8 hold for S_1, \dots, S_n and (S, \mathfrak{d}) .*

Proof. Since \mathcal{B} is a choice-free automaton, by Lemma 4.20, there exists an accepting run θ of \mathcal{B} over a tree t , such that the following hold:

- (a) each partial run θ_i is embedded in θ at some position $r_i \in \text{dom}(\theta)$, for all $i \in [1..n]$,
- (b) $r_i \cdot \text{dom}(\theta_i) \cap r_j \cdot \text{dom}(\theta_j) = \emptyset$, for all $1 \leq i < j \leq n$,
- (c) each path between the positions r_i and r_j traverses at least once some reset path, disjoint from $\bigcup_{k=1}^n r_k \cdot \text{dom}(\theta_k)$, for all $1 \leq i < j \leq n$.

For each $i \in [1..n]$, since $\langle \theta_i, t_i, \mathfrak{s}_i, S'_i \rangle$ is a view, we have $S'_i \models^{s_i} \Theta(t_i)$, hence $\mathfrak{s}_i(\text{fv}(\Theta(t_i))) \subseteq \text{supp}(\sigma'_i)$, where we assume w.l.o.g. that $S'_i \stackrel{\text{def}}{=} (\mathbf{U}, \sigma'_i)$. By point (b) above, the subformulae corresponding to the subtrees of t with domains $r_i \cdot \text{dom}(\theta_i)$, for $i \in [1..n]$, have disjoint sets of free variables. We define a store \mathfrak{s} as follows, for each variable $x^{[p]} \in \text{fv}(\Theta(t))$:

- ▷ if $p = r_i p'$ and $p' \in \text{dom}(t_i)$, for some $i \in [1..n]$, then we set $\mathfrak{s}(x^{[p]}) \stackrel{\text{def}}{=} \mathfrak{s}_i(x^{[p']})$,
- ▷ otherwise, we chose a fresh value $\mathfrak{s}(x^{[p]})$, such that $\mathfrak{s}(x^{[p]}) \notin \bigcup_{i=1}^n \text{supp}(\sigma'_i)$ and $\mathfrak{s}(x^{[p]}) \neq \mathfrak{s}(z^{[r]})$, for each variable $z^{[r]}$, such that $x^{[p]} \not\approx_{\Theta(t)} z^{[r]}$.

By the fact that $\text{supp}(\sigma'_i) \cap \text{supp}(\sigma'_j) = \emptyset$, i.e., $\mathfrak{s}_i(\text{fv}(\Theta(t_i))) \cap \mathfrak{s}_j(\text{fv}(\Theta(t_j))) = \emptyset$, for all $1 \leq i < j \leq n$, and the construction of \mathfrak{s} , we obtain that \mathfrak{s} is canonical for $\Theta(t)$. The store \mathfrak{s} defines the structure $\mathbf{S} = (\mathbf{U}, \sigma)$, as follows:

$$\sigma(r) \stackrel{\text{def}}{=} \{ \langle \mathfrak{s}(z_1), \dots, \mathfrak{s}(z_{\#r}) \rangle \mid r(z_1, \dots, z_{\#r}) \text{ occurs in } \Theta(t) \}, \text{ for all } r \in \mathbb{R}$$

By the definition of \mathbf{S} , we have $\mathbf{S} \models^{\mathfrak{s}} \Theta(t)$. Moreover, we define the relation:

$$\mathfrak{d} \stackrel{\text{def}}{=} \{ \langle \mathfrak{s}(x), \mathfrak{s}(y) \rangle \mid x \neq y \text{ or } y \neq x \text{ occurs in } \Theta(t) \}$$

We have obtained a rich canonical model $(\mathbf{S}, \mathfrak{d}) \in \llbracket \Theta(t) \rrbracket^r$ and, since $t \in \mathcal{L}(\mathcal{B})$, we have $(\mathbf{S}, \mathfrak{d}) \in \llbracket \mathcal{B} \rrbracket^r$. We prove below the three conditions from Definition 3.8:

- (1) By the construction of \mathbf{S} , we have $S'_i \subseteq \mathbf{S}$, for all $i \in [1..n]$. Since $\mathbf{S}_1, \dots, \mathbf{S}_n$ are pairwise disjoint, their composition is defined, hence $\mathbf{S}_1 \bullet \dots \bullet \mathbf{S}_n \subseteq \mathbf{S}'_1 \bullet \dots \bullet \mathbf{S}'_n \subseteq \mathbf{S}$. W.l.o.g., let $\mathbf{S}_i = (\mathbf{U}, \sigma_i)$, for all $i \in [1..n]$. To prove $\mathbf{S}_1 \bullet \dots \bullet \mathbf{S}_n \subseteq \mathbf{S}$, by Definition 3.7, we must prove that:

$$\sigma_1(r) \uplus \dots \uplus \sigma_n(r) = \{ \langle u_1, \dots, u_{\#r} \rangle \in \sigma(r) \mid u_1, \dots, u_{\#r} \in \text{supp}(\sigma_1) \cup \dots \cup \text{supp}(\sigma_n) \} \quad (4.2)$$

The “ \subseteq ” direction follows from $\mathbf{S}_1 \bullet \dots \bullet \mathbf{S}_n \subseteq \mathbf{S}$, hence we are left with proving the dual “ \supseteq ” direction. Let $\langle u_1, \dots, u_{\#r} \rangle \in \sigma(r)$ be a tuple, such that $u_1, \dots, u_{\#r} \in \bigcup_{i=1}^n \text{supp}(\sigma_i)$. By the definition of σ , there exists a relation atom $r(z_1^{[p]}, \dots, z_{\#r}^{[p]})$ in $\Theta(t)$, such that $\mathfrak{s}(z_i) = u_i$, for all $i \in [1..\#r]$. To simplify matters, we assume that the position of each variable in the relation atom is the same, the case where these positions are either p and pi , or pi and pj , for some $p \in \mathbb{N}^*$ and $i \neq j \in \mathbb{N}$ is treated in a similar way and left to the reader. Moreover, for each $i \in [1..\#r]$, there exists a unique $k_i \in [1..n]$, such that $u_i \in \text{supp}(\sigma_{k_i})$. Suppose, for a contradiction, that $k_i \neq k_j$, for some $1 \leq i < j \leq n$. Then, there exist paths between p and some positions $s_i \in r_i \cdot \text{dom}(t_i)$ and $s_j \in r_j \cdot \text{dom}(t_j)$, such that $z_i^{[p]} \approx_{\Theta(t)} \xi_i^{[s_i]}$ and $z_j^{[p]} \approx_{\Theta(t)} \xi_j^{[s_j]}$. Consider the case where θ_i and θ_j are runs (the case where one of them is a context uses a similar argument and is left to the reader). Since t_i and t_j are embedded in t at positions r_i and r_j , respectively, at least one of these paths, say the one from p to s_i , contains the position r_i . Then, there exists a variable $x_{\ell_i}^{[r_i]}$, for $\ell_i \in [1..\#\theta_i(\epsilon)]$, such that $\mathfrak{s}(x_{\ell_i}^{[r_i]}) = u_i$. Hence, $\text{supp}(\sigma_i) \cap \mathfrak{s}(\{x_1^{[r_i]}, \dots, x_{\#\theta_i(\epsilon)}^{[r_i]}\}) = \text{supp}(\sigma_i) \cap \mathfrak{s}_i(\{x_1^{[\epsilon]}, \dots, x_{\#\theta_i(\epsilon)}^{[\epsilon]}\}) \neq \emptyset$, in contradiction with the fact that \mathbf{S}_i is encapsulated by $\langle \theta_i, t_i, \mathfrak{s}_i, S'_i \rangle$, by condition (ii) of Definition 4.33. We obtained that $k_1 = \dots = k_{\#r}$, hence $u_1, \dots, u_{\#r} \in \text{supp}(\sigma_k)$, leading to $\langle u_1, \dots, u_{\#r} \rangle \in \sigma_k(r)$, for some index $k \in [1..n]$.

This proves (4.2). Since the choice of r was arbitrary, we obtain that $S_1 \bullet \dots \bullet S_n \sqsubseteq S$, by Definition 3.7.

- (2) Suppose, for a contradiction, that there exists a pair $(u, v) \in \mathfrak{D}$, such that $u \in \text{supp}(\sigma_i)$ and $v \in \text{supp}(\sigma_j)$, for some indices $1 \leq i < j \leq n$. Then, there exists a disequality $x^{[p]} \neq y^{[p]}$ (or $y \neq x$, this case being symmetric) in $\Theta(t)$, such that $\mathfrak{s}(x^{[p]}) = u$ and $\mathfrak{s}(y^{[p]}) = v$. Since $u \in \text{supp}(\sigma_i)$, there exists a variable $\xi_i^{[p_i]}$ such that $p_i \in r_i \cdot \text{dom}(t_i)$ and $\mathfrak{s}(\xi_i^{[p_i]}) = u$. Since \mathfrak{s} is canonical for $\Theta(t)$, we have $x \approx_{\Theta(t)} \xi_i^{[p_i]}$. Suppose, for a contradiction, that $p \notin r_i \cdot \text{dom}(t_i)$. Then, by a similar argument as the one used in the proof of point (1), we obtain a contradiction with condition (ii) of Definition 4.33, hence $p \in r_i \cdot \text{dom}(t_i)$. Symmetrically, we obtain $p \in r_j \cdot \text{dom}(t_j)$, hence $r_i \cdot \text{dom}(t_i) \cap r_j \cdot \text{dom}(t_j) \neq \emptyset$, which contradicts point (b) above.
- (3) Suppose, for a contradiction, that there exists a relation symbol $r \in \mathbb{R}$ and tuples $\langle u_1, \dots, u_{\#r} \rangle, \langle v_1, \dots, v_{\#r} \rangle \in \sigma(r)$, such that $\{u_1, \dots, u_{\#r}\} \cap \text{supp}(\sigma_i) \neq \emptyset$, $\{v_1, \dots, v_{\#r}\} \cap \text{supp}(\sigma_j) \neq \emptyset$ and $\{u_1, \dots, u_{\#r}\} \cap \{v_1, \dots, v_{\#r}\} \neq \emptyset$, for some indices $1 \leq i < j \leq n$. Then, there exists two distinct relation atoms $r(z_1^{[p_1]}, \dots, z_{\#r}^{[p_1]})$ and $r(z_1^{[p_2]}, \dots, z_{\#r}^{[p_2]})$ in $\Theta(t)$ and variables $\xi_i^{[s_i]}$, $s_i \in r_i \cdot \text{dom}(t_i)$ and $\xi_j^{[s_j]}$, $s_j \in r_j \cdot \text{dom}(t_j)$, such that $z_k^{[p_1]} \approx_{\Theta(t)} \xi_i^{[s_i]}$, $z_\ell^{[p_2]} \approx_{\Theta(t)} \xi_j^{[s_j]}$, for some indices $k, \ell \in [1.. \#r]$. For simplicity, we consider that the position of the variables is the same in the above relation symbols, i.e., p_1 and p_2 , respectively. By an argument similar to the one used in the proof of point (2), we obtain that $p_1 \in r_i \cdot \text{dom}(t_i)$ and $p_2 \in r_j \cdot \text{dom}(t_j)$. However, since $\{u_1, \dots, u_{\#r}\} \cap \{v_1, \dots, v_{\#r}\} \neq \emptyset$, there exist indices $g, h \in [1.. \#r]$ such that $z_g^{[p_1]} \approx_{\Theta(t)} z_h^{[p_2]}$. For simplicity, we consider the case where θ_i and θ_j are runs, the case where at least one of them is a context uses a similar argument being left to the reader. Then the path between r_i and r_j is contained with the path between p_1 and p_2 . By point (c) above, this path contains a reset path disjoint from $\bigcup_{k=1}^n r_k \cdot \text{dom}(\theta_k)$. Since, moreover, \mathcal{B} has no persistent variables, by construction, we obtain $z_g^{[p_1]} \not\approx_{\Theta(t)} z_h^{[p_2]}$, contradiction. \square




Figure 14: The Chain of Automata Transformations

We are now ready to prove Lemma 4.1:

Proof. Let $\mathcal{A}_{\Delta, A}$ be the Σ -labeled automaton corresponding to Δ and A , such that $\llbracket A \rrbracket_{\Delta} = \llbracket \mathcal{A}_{\Delta, A} \rrbracket$, by Lemma 4.13 (1). The chain of transformations depicted in Figure 14 produces the set $\{\mathcal{B}_j^i\}_{i \in [1..n], j \in [1..m_i]}$ of choice-free and all-satisfiable automata, such that $\llbracket \mathcal{A}_{\Delta, A} \rrbracket$ is treewidth bounded if and only if $\llbracket \mathcal{B}_j^i \rrbracket$ is treewidth bounded, for each $i \in [1..n]$ and $j \in [1..m_i]$. In particular, $\mathcal{A}_{\Delta, A}^1, \dots, \mathcal{A}_{\Delta, A}^n$ is the language-preserving choice-free decomposition of $\mathcal{A}_{\Delta, A}$ (Lemma 4.9) and $\bar{B}_1^i, \dots, \bar{B}_{m_i}^i$ are obtained by removing the persistent variables

from each automaton $\widetilde{\mathcal{B}}^i$ in the language-preserving choice-free decomposition of $(\widetilde{\mathcal{A}}_{\Delta, \mathcal{A}})^{II}$ (Lemma 4.28). We assume, without loss of generality, that the initial state of each \mathcal{B}_j^i is $q_{\mathcal{B}}$ and let Γ_j^i be the SID, such that $\llbracket \mathcal{B}_j^i \rrbracket = \llbracket \mathcal{B} \rrbracket_{\Gamma_j^i}$ and $\llbracket \mathcal{B}_j^i \rrbracket^r = \llbracket \mathcal{B} \rrbracket_{\Gamma_j^i}^r$, by Lemma 4.13 (2). We are left with proving that Γ_j^i is expandable for \mathcal{B} , for each $i \in [1..n]$ and $j \in [1..m_i]$. Let Γ be any of the SIDs $\Gamma_1^1, \dots, \Gamma_{m_n}^n$ and \mathcal{B} be the automaton such that $\llbracket \mathcal{B} \rrbracket_{\Gamma}^c = \llbracket \mathcal{B} \rrbracket^c$, by Lemma 4.13 (2). Let $\mathcal{S}_1, \dots, \mathcal{S}_k \in \llbracket \mathcal{B} \rrbracket_{\Gamma}^c$ be pairwise disjoint structures. By Lemma 4.34, for each $i \in [1..k]$ there exists a decomposition $\mathcal{S}_i = \mathcal{S}_1^i \bullet \dots \bullet \mathcal{S}_{\ell_i}^i$ into pairwise disjoint structures, such that the structure \mathcal{S}_j^i is encapsulated by a view $\langle \theta_j^i, t_j^i, \mathfrak{s}_j^i, \bar{\mathcal{S}}_j^i \rangle$ for \mathcal{B} , for each $i \in [1..k]$ and $j \in [1..\ell_i]$. Without loss of generality, we assume that the structures $\{\bar{\mathcal{S}}_j^i\}_{i \in [1..k], j \in [1..\ell_i]}$ are pairwise disjoint, hence the composition of $\mathcal{S}_1^1, \dots, \mathcal{S}_{\ell_n}^n$ is defined, thus we obtain $\mathcal{S}_1 \bullet \dots \bullet \mathcal{S}_k = \mathcal{S}_1^1 \bullet \dots \bullet \mathcal{S}_{\ell_n}^n$. By Lemma 4.35, there exists a rich canonical model $(\mathcal{S}, \mathfrak{d}) \in \llbracket \mathcal{B} \rrbracket^r$, such that conditions (1), (2) and (3) from Definition 3.8 hold for $\mathcal{S}_1, \dots, \mathcal{S}_{\ell_k}^k$ and $(\mathcal{S}, \mathfrak{d})$. Since $\llbracket \mathcal{B} \rrbracket^r = \llbracket \mathcal{B} \rrbracket_{\Gamma}^r$, by Lemma 4.13 (2), Γ is expandable for \mathcal{B} . \square

4.7. The Decidability of the Treewidth Boundedness Problem for SLR. The proof of the main result (Theorem 4.36) uses a reduction of an arbitrary SID to a set of expandable SIDs, having an equivalent treewidth boundedness status (Lemma 4.1). Since treewidth boundedness is decidable for expandable SIDs (Theorem 3.33), this proves the decidability of the problem, in the general case. We conclude this section with the proof of the main result, that is the decidability of the TWB^{SLR} problem, for unrestricted SIDs:

Theorem 4.36. *There exists an algorithm that decides, for each SID Δ and nullary predicate \mathcal{A} , whether the set $\llbracket \mathcal{A} \rrbracket_{\Delta}$ has bounded treewidth. If, moreover, this is the case, then $\text{tw}(\llbracket \mathcal{A} \rrbracket_{\Delta}) \leq \max\text{Vars}(\Delta) + N \cdot M$, where:*

$$\begin{aligned} M &\stackrel{\text{def}}{=} 2 \cdot \max\text{Vars}(\Delta) + (1 + \max\text{PredAtoms}(\Delta)) \cdot \text{relNo}(\Delta) \cdot \max\text{Vars}(\Delta)^{\max\text{RelArity}(\Delta)} \\ N &\stackrel{\text{def}}{=} \max(K, \max\text{PredAtoms}(\Delta)^K) \\ K &\stackrel{\text{def}}{=} \text{predNo}(\Delta) \cdot \text{relNo}(\Delta) \cdot \max\text{Vars}(\Delta)^{\max\text{Vars}(\Delta) + \max\text{RelArity}(\Delta)} \end{aligned}$$

Proof. By Lemma 4.1, the treewidth boundedness problem for the set $\llbracket \mathcal{A} \rrbracket_{\Delta}$ can be effectively reduced to finitely many treewidth boundedness problems for sets $\llbracket \mathcal{B} \rrbracket_{\Gamma_1}, \dots, \llbracket \mathcal{B} \rrbracket_{\Gamma_k}$, where $\Gamma_1, \dots, \Gamma_k$ are expandable for \mathcal{B} . The latter problem is decidable, by Theorem 3.33. The upper bound follows from the sequence of transformations given in Figure 14. Let Γ be any of the expandable SIDs $\Gamma_1, \dots, \Gamma_k$. By Theorem 3.33, we have $\text{tw}(\llbracket \mathcal{B} \rrbracket_{\Gamma}) \leq \max\text{Vars}(\Delta)$. Note that the maximum number of variables occurring in a rule is not increased by the construction of Γ from Δ . By Lemma 4.23 (3), Lemma 4.25 (3), Lemma 4.29 (3) and Lemma 4.32 (3), we obtain $\text{tw}(\llbracket \mathcal{A} \rrbracket_{\Delta}) \leq \text{tw}(\llbracket \mathcal{B} \rrbracket_{\Gamma}) + \text{card}(\delta^1) \cdot M$, where δ^1 is the set of 1-transitions of the choice-free automaton used to define Γ (point (2) of Lemma 4.13) and $M = 2 \cdot \max\text{Vars}(\Delta) + (1 + \max\text{PredAtoms}(\Delta)) \cdot \text{relNo}(\Delta) \cdot \max\text{Vars}(\Delta)^{\max\text{RelArity}(\Delta)}$. Note that $\text{card}(\delta^1) \cdot M$ is the sum of the increases in the upper bounds of the treewidth along the transformation. By Lemma 4.9, we obtain $\text{card}(\delta^1) \leq \max(\text{card}(\mathcal{Q}_{\Delta}), \max\text{PredAtoms}(\Delta)^{\text{card}(\mathcal{Q}_{\Delta})})$, where \mathcal{Q}_{Δ} denotes the set of states of $\mathcal{A}_{\Delta, \mathcal{A}}$ (Lemma 4.13). Since the entire construction was done assuming that Δ is equality-free and all-satisfiable, by lifting these assumptions, we obtain $\text{card}(\mathcal{Q}_{\Delta}) \leq \text{predNo}(\Delta) \cdot \text{relNo}(\Delta) \cdot \max\text{Vars}(\Delta)^{\max\text{Vars}(\Delta) + \max\text{RelArity}(\Delta)}$ (Lemma 2.6 and Lemma 2.9). \square

Note that, unlike Theorem 3.33, that gives an optimal upper bound for the treewidths of the structures described by an expandable SID, Theorem 4.36 does not provide such an optimal upper bound. In particular, the upper bound of Theorem 4.36 grows doubly exponential in the $\max\text{Vars}(\Delta)$ and $\max\text{RelArity}(\Delta)$ parameters and simply exponential in the $\text{predNo}(\Delta)$ and $\text{relNo}(\Delta)$ parameters. Finding an optimal bound for the general case is considered as subject for future work.

5. THE TREewidth BOUNDEDNESS PROBLEM FOR FIRST ORDER LOGIC

This section proves the undecidability of the treewidth boundedness problem for first-order logic. This result places the frontier of decidability for this problem between the classical first-order logic and substructural SLR with simple inductive definitions, i.e., using only existentially quantified separating conjunctions of atoms.

We recall that the *first-order logic* (FO) is the set of formulæ consisting of equalities and relation atoms, connected by boolean conjunction, negation and existential quantification. The semantics of first order logic is given by the satisfaction relation $(\mathbf{U}, \sigma) \Vdash^s \phi$ between structures and formulæ, parameterized by a *store* $\mathfrak{s} : \mathbb{V} \rightarrow \mathbf{U}$ such that $(\mathbf{U}, \sigma) \Vdash^s r(x_1, \dots, x_{\#r})$ iff $\langle \mathfrak{s}(x_1), \dots, \mathfrak{s}(x_{\#r}) \rangle \in \sigma(r)$. If ϕ is a sentence the store is not important, thus we omit the superscript and write $\mathbf{S} \Vdash \phi$ instead. The set of *models* of a FO sentence ϕ is denoted as $\llbracket \phi \rrbracket \stackrel{\text{def}}{=} \{ \mathbf{S} \mid \mathbf{S} \Vdash \phi \}$. Although we use the same notation for the sets of models of FO and SLR formulæ, the underlying logic is clear from the context.

This section is concerned with the proof of the following theorem:

Theorem 5.1. *The problem is $\llbracket \phi \rrbracket$ treewidth-bounded, for a given FO sentence ϕ with at least two binary relation symbols and several unary relation symbols, is undecidable.*

Proof. We will reduce from the undecidability of the *Tiling Problem* [Ber66]. We first recall its definition. Given a finite set of tiles $S = \{t_1, \dots, t_n\}$ is there a tiling of the plane such that the colors of neighbouring tiles match? (We note that rotating or reflecting the tiles is not allowed.) In more detail: We assume the plane is given by integer coordinates (x, y) with $x, y \in \mathbb{Z}$. We want to put a copy of a tile at every coordinate. We will require that neighbouring tiles match. For this we assume to be given a relation $H \subseteq S \times S$ – specifying which tiles match can be placed next to each other horizontally – and $V \subseteq S \times S$ – specifying which tiles match can be placed next to each other vertically. We now require for every tiling that $(t_i, t_j) \in H$, for all tiles t_i and t_j placed at coordinates (x, y) and $(x + 1, y)$, and $(t_i, t_j) \in V$, for all tiles t_i and t_j placed at coordinates (x, y) and $(x, y + 1)$. It is well-known that it is undecidable whether such a tiling exists [Ber66]. In fact, it is known that is already undecidable whether such a tiling exists for the upper-right quadrant of the plane, i.e., when coordinates (x, y) are restricted to $x, y \in \mathbb{N}$.

We will now reduce the tiling problem to deciding whether a given first-order formula has infinitely many non-isomorphic models of unbounded treewidth. We consider some instance of the tiling problem. For encoding this problem, we define the signature $\mathbb{R} = \{\text{up}, \text{right}, \mathbf{N}, \mathbf{S}, \mathbf{E}, \mathbf{W}, \mathbf{I}, \mathbf{T}_1, \dots, \mathbf{T}_n\}$ to consist of the binary relations **up** and **right**, the unary relations **S**, **E**, **W**, **N**, **I** and the unary relations $\mathbf{T}_1, \dots, \mathbf{T}_n$ (one for each tile in the tiling instance). We then consider the following formula:

$$\phi \stackrel{\text{def}}{=} \bigwedge_{i=1}^{13} \psi_i \wedge \bigwedge_{j=1}^4 \phi_j$$

where:

- ▷ $\psi_1 \stackrel{\text{def}}{=} \forall x \exists^{\leq 1} y . \mathbf{right}(x, y) \wedge \exists^{\leq 1} y . \mathbf{right}(y, x)$ states that \mathbf{right} and \mathbf{right}^{-1} are partial functions, ψ_2 states that that \mathbf{up} and \mathbf{up}^{-1} are partial functions,
- ▷ $\psi_3 \stackrel{\text{def}}{=} \forall x, y, z . \mathbf{up}(x, y) \wedge \mathbf{right}(x, z) \rightarrow \exists w . \mathbf{up}(z, w) \wedge \mathbf{right}(y, w)$ states that \mathbf{right} and \mathbf{up} commute, ψ_4 states that \mathbf{right}^{-1} and \mathbf{up} commute, ψ_5 states that \mathbf{right} and \mathbf{up}^{-1} commute and ψ_6 states that \mathbf{right}^{-1} and \mathbf{up}^{-1} commute,
- ▷ $\psi_7 \stackrel{\text{def}}{=} \forall x . \mathbf{S}(x) \leftrightarrow \neg \exists y . \mathbf{up}(y, x)$ states that south-labelled nodes are exactly the ones that do not have incoming \mathbf{up} edges, and $\psi_8, \psi_9, \psi_{10}$ define the analogous property for the west, east, and north labels,
- ▷ $\psi_{11} \stackrel{\text{def}}{=} \forall x . \mathbf{I}(x) \leftrightarrow \neg(\mathbf{S}(x) \vee \mathbf{E}(x) \vee \mathbf{N}(x) \vee \mathbf{W}(x))$ states that internal nodes are exactly the ones not labelled by south, east, north or west,
- ▷ $\psi_{12} \stackrel{\text{def}}{=} \forall x . \mathbf{I}(x) \rightarrow \exists y . \mathbf{right}(x, y) \wedge \exists y . \mathbf{right}(y, x)$ states that internal nodes have exactly one outgoing and exactly one incoming \mathbf{right} edge, ψ_{13} states that internal nodes have exactly one outgoing and exactly one incoming \mathbf{up} edge,
- ▷ $\phi_1 \stackrel{\text{def}}{=} \forall x . \bigvee_{i=1}^n \mathbf{T}_i(x)$ states that every coordinate holds at least one tile,
- ▷ $\phi_2 \stackrel{\text{def}}{=} \forall x . \bigwedge_{i \neq j} \neg \mathbf{T}_i(x) \vee \neg \mathbf{T}_j(x)$ states that every coordinate holds at most one tile,
- ▷ $\phi_3 \stackrel{\text{def}}{=} \forall x \forall y . \mathbf{right}(x, y) \rightarrow \bigvee_{(t_i, t_j) \in H} \mathbf{T}_i(x) \wedge \mathbf{T}_j(y)$ states that tiles, that are next to each other horizontally, satisfy the horizontal matching constraint, and
- ▷ $\phi_4 \stackrel{\text{def}}{=} \forall x \forall y . \mathbf{up}(x, y) \rightarrow \bigvee_{(t_i, t_j) \in V} \mathbf{T}_i(x) \wedge \mathbf{T}_j(y)$ states that tiles, that are next to each other vertically, satisfy the vertical matching constraint.

As usual, the formulæ $\exists^{\leq 1} x . \varphi$ stand for $\exists x . \varphi \rightarrow \forall y . \varphi[x/y] \rightarrow y = x$. The proof will make use of the fact that ϕ encodes grids and non-standard models of grids, which are (disjoint unions of) grid-like structures. We will argue the following:

Fact 5.2. Each model of ϕ can be decomposed into (disjoint unions of) grids, cylinders, and toruses, where grids have $\mathbf{S}, \mathbf{E}, \mathbf{W}, \mathbf{N}$ borders, cylinders have either \mathbf{S}, \mathbf{N} or \mathbf{E}, \mathbf{W} borders, and toruses do not have any borders and only consist of internal nodes.

Proof. We recall that we are only interested in finite models of first-order formulæ. We note that ϕ specifies that \mathbf{up} and \mathbf{right} are (partial) functions, and, hence, we will use functional notation in the following. We now fix a model (U, σ) of ϕ – as usual, we require that $U \neq \emptyset$. We decompose (U, σ) into its maximally connected components, connected via $\mathbf{up}, \mathbf{right}$. We choose a representative u_C , for each component C . We observe that either there are $j \leq 0 \leq i$ such that $\mathbf{right}^j(u_C)$ is \mathbf{W} -labelled and $\mathbf{right}^i(u_C)$ is \mathbf{E} -labelled, or $\mathbf{right}^i(u_C) = u_C$, for some $i \geq 0$ (because the universe is finite and the functionality of \mathbf{right} and \mathbf{right}^{-1} ensures that the only possible loop returns to u_C). An analogous statement holds for \mathbf{up} as well as the \mathbf{N} and \mathbf{S} labels. We now call a component C a *grid*, if u_C reaches $\mathbf{S}, \mathbf{E}, \mathbf{W}, \mathbf{N}$ via $\mathbf{up}, \mathbf{right}$ and their inverses, a *cylinder* if u_C reaches \mathbf{S}, \mathbf{N} or \mathbf{E}, \mathbf{W} via \mathbf{up} and its inverse resp. \mathbf{right} and its inverse, or a *torus*, otherwise. \square

We now justify the naming of these components:

Fact 5.3. Consider a grid component C with representative u_C such that $\mathbf{right}^j(u_C)$ is \mathbf{W} -labelled, $\mathbf{right}^i(u_C)$ is \mathbf{E} -labelled, $\mathbf{up}^k(u_C)$ is \mathbf{S} -labelled and $\mathbf{up}^l(u_C)$ is \mathbf{N} -labelled, for some $j \leq 0 \leq i$ and $k \leq 0 \leq l$. We claim that:

- (1) the elements $\mathbf{right}^a(\mathbf{up}^b(u_C))$ exist, for all $j \leq a \leq i$ and all $k \leq b \leq l$,
- (2) an element $\mathbf{right}^a(\mathbf{up}^b(u_C))$, for $j \leq a \leq i$ and $k \leq b \leq l$, is \mathbf{E} -labelled iff $a = i$; analogous claims hold for the labels $\mathbf{S}, \mathbf{W}, \mathbf{N}$,
- (3) the elements $\mathbf{right}^a(\mathbf{up}^b(u_C))$ are internal nodes, for all $j < a < i$ and all $k < b < l$,

- (4) the elements $\mathbf{right}^a(\mathbf{up}^b(u_C))$ are pairwise different, for all $j \leq a \leq i$ and all $k \leq b \leq l$,
- (5) all elements of the component can be represented as $\mathbf{right}^a(\mathbf{up}^b(u_C))$, for some $j \leq a \leq i$ and $k \leq b \leq l$, and
- (6) the component is isomorphic to a grid.

Proof. Items (1), (2) and (3) directly follow from the commutativity requirements. For (4), we consider some $j \leq a_1, a_2 \leq i$ and $k \leq b_1, b_2 \leq l$. We will show that $a_1 \neq a_2$ resp. $b_1 \neq b_2$ imply that $\mathbf{right}^{a_1}(\mathbf{up}^{b_1}(u_C)) \neq \mathbf{right}^{a_2}(\mathbf{up}^{b_2}(u_C))$. We will assume that $\mathbf{right}^{a_1}(\mathbf{up}^{b_1}(u_C)) = \mathbf{right}^{a_2}(\mathbf{up}^{b_2}(u_C))$ and derive a contradiction. Let us assume that $a_1 > a_2$ (the other cases are analogous). Then, we have $\mathbf{right}^{a_1+s}(\mathbf{up}^{b_1}(u_C)) = \mathbf{right}^{a_2+s}(\mathbf{up}^{b_2}(u_C))$ for $s = i - a_1 \geq 0$. However, $\mathbf{right}^{a_1+s}(\mathbf{up}^{b_1}(u_C))$ is E-labelled, while $\mathbf{right}^{a_2+s}(\mathbf{up}^{b_2}(u_C))$ is not, by (2), contradiction. For (5), we observe that every node reachable from u_C connected via \mathbf{up} , \mathbf{right} and their inverses can be represented as $\mathbf{right}^a(\mathbf{up}^b(u_C))$, because of the commutativity requirements; further, we must have $j \leq a \leq i$ and $k \leq b \leq l$ because the S, E, W, N borders do not have outgoing edges. For (6), we observe that the component is isomorphic to the structure with domain $\{(x, y) \mid x, y \in [j, i], y \in [k, l]\}$, where \mathbf{right} is interpreted as $\{((x, y), (x + 1, y)) \mid x \in [j, i - 1], y \in [k, l]\}$ and \mathbf{up} as $\{((x, y), (x, y + 1)) \mid x \in [j, i], y \in [k, l - 1]\}$. \square

Fact 5.4. Consider a cylinder component C with representative u_C such that $\mathbf{right}^j(u_C)$ is W-labelled, $\mathbf{right}^i(u_C)$ is E-labelled, and $\mathbf{up}^k(u_C) = u_C$ for some $j \leq 0 \leq i$ and $k \leq 0$, where k is the smallest number with this property. (The properties stated below hold analogously for N, S cylinders). We claim that:

- (1) the elements $\mathbf{right}^a(\mathbf{up}^b(u_C))$ exist, for all $j \leq a \leq i$ and all $0 \leq b < k$,
- (2) an element $\mathbf{right}^a(\mathbf{up}^b(u_C))$, for $j \leq a \leq i$ and $0 \leq b < k$, is E-labelled iff $a = i$; an analogous claim hold for the label W,
- (3) the elements $\mathbf{right}^a(\mathbf{up}^b(u_C))$ are internal nodes, for all $j < a < i$ and all $0 \leq b < k$,
- (4) all the elements $\mathbf{right}^a(\mathbf{up}^b(u_C))$ are pairwise different, for all $j \leq a \leq i$ and all $0 \leq b < k$,
- (5) all nodes of the component can be represented as $\mathbf{right}^a(\mathbf{up}^b(u_C))$, for some $j \leq a \leq i$ and $0 \leq b < k$, and
- (6) the component is isomorphic to a cylinder, i.e., a grid for which the north-border connects to south-border.

Proof. Items (1), (2) and (3) directly follow from the commutativity requirements. For (4), we consider some $j \leq a_1, a_2 \leq i$ and $0 \leq b_1, b_2 < k$. We will show that $a_1 \neq a_2$ resp. $b_1 \neq b_2$ imply that $\mathbf{right}^{a_1}(\mathbf{up}^{b_1}(u_C)) \neq \mathbf{right}^{a_2}(\mathbf{up}^{b_2}(u_C))$. We will assume that $\mathbf{right}^{a_1}(\mathbf{up}^{b_1}(u_C)) = \mathbf{right}^{a_2}(\mathbf{up}^{b_2}(u_C))$ and derive a contradiction. Let us first assume that $a_1 > a_2$ (the case $a_1 < a_2$ is symmetric). Then, we have $\mathbf{right}^{a_1+s}(\mathbf{up}^{b_1}(u_C)) = \mathbf{right}^{a_2+s}(\mathbf{up}^{b_2}(u_C))$ for $s = i - a_1$. However, $\mathbf{right}^{a_1+s}(\mathbf{up}^{b_1}(u_C))$ is E-labelled, while $\mathbf{right}^{a_2+s}(\mathbf{up}^{b_2}(u_C))$ is not, by (2), contradiction. Now we assume $a_1 = a_2$ and $b_1 > b_2$ (the case $b_1 < b_2$ is symmetric). Then, $\mathbf{right}^{a_1}(\mathbf{up}^{b_1}(u_C)) = \mathbf{right}^{a_2}(\mathbf{up}^{b_2}(u_C))$ implies that $\mathbf{up}^{b_1-b_2}(u_C) = u_C$ with $0 \leq b_1 - b_2 < k$. However, this contradicts that k is the smallest number with this property. For (5), we observe that every node reachable from u_C connected via \mathbf{up} , \mathbf{right} and their inverses can be represented as $\mathbf{right}^a(\mathbf{up}^b(u_C))$, because of the commutativity requirements; further, we can in fact choose $0 \leq b < k$ because of commutativity and the assumption that $\mathbf{up}^k(u_C) = u_C$. Moreover, we must have $j \leq a \leq i$ because the E, W borders do not have outgoing edges. For (6), we observe that

the component is isomorphic to the structure with domain $\{(x, y) \mid x, y \in [j, i], y \in [k, l]\}$, where **right** is interpreted as $\{((x, y), (x + 1, y)) \mid x \in [j, i - 1], y \in [0, k - 1]\}$ and **up** as $\{((x, y), (x, y + 1)) \mid x \in [j, i], y \in [0, k - 2]\} \cup \{((x, k), (x, 1)) \mid x \in [j, i]\}$. \square

Fact 5.5. Consider a torus component C , with representative u_C such that $\mathbf{up}^k(u_C) = u_C$ and $\mathbf{right}^l(u_C) = u_C$ for some $k \leq 0$ and $l \leq 0$, where k and l are the smallest numbers with this property. We claim that:

- (1) the elements exists $\mathbf{right}^a(\mathbf{up}^b(u_C))$, for all $0 \leq a < k$ and all $0 \leq b < l$,
- (2) the elements $\mathbf{right}^a(\mathbf{up}^b(u_C))$ are internal nodes, for all $0 \leq a < k$ and all $0 \leq b < l$,
- (3) all nodes of the component can be represented as $\mathbf{right}^a(\mathbf{up}^b(u_C))$, for some $0 \leq a < k$ and $0 \leq b < l$.

Proof. Items (1) and (2) directly follow from the commutativity requirements. For (3), we observe that every node reachable from u_C connected via **up**, **right** and their inverses can be represented as some $\mathbf{right}^a(\mathbf{up}^b(u_C))$ because of the commutativity requirements; further, we can in fact choose $0 \leq a < k$ and $0 \leq b < l$ because of commutativity and the assumptions that $\mathbf{up}^k(u_C) = u_C$ and $\mathbf{right}^l(u_C) = u_C$. We note that the $\mathbf{right}^a(\mathbf{up}^b(u_C))$ are in general not pairwise different (e.g., we might have $\mathbf{right}^a(u_C) = \mathbf{up}^b(u_C)$ for some $0 \leq a < k$ and $0 \leq b < l$). However, in our below argument we do not need to distinguish whether all the elements $\mathbf{right}^a(\mathbf{up}^b(u_C))$ of a torus component are pairwise different. \square

The following claim reduces the treewidth boundedness problem for first-order logic to the tiling problem for the first quadrant of the plane:

Fact 5.6. ϕ has models of unbounded treewidth iff there is a tiling of the upper-right quadrant of the plane.

Proof. “ \Leftarrow ” Let us assume that there is a tiling of the upper-right quadrant of the plane. Then, for every $n \in \mathbb{N}$, this tiling induces a square grid G_n of size $n \times n$ with $G_n \models \phi$: simply take the tiles at positions (x, y) , with $x, y \in [1, n]$, from the tiling of the upper-right quadrant, and verify that in this way we obtain a model of the formula ϕ .

“ \Rightarrow ” We now assume that ϕ has models of unbounded treewidth, i.e., for every $i \geq 1$ there is a finite model (U, σ) with $\text{tw}((U, \sigma)) \geq i$. If any model (U, σ) contains a torus component C , we immediately obtain a tiling of the upper-right quadrant by unrolling the torus: we define the tiling of the upper right quadrant by placing the tile of the element $\mathbf{right}^i(\mathbf{up}^j(u_C))$ at position (i, j) . It is then routine to verify that the subformula $\bigwedge_{j=1}^4 \phi_j$ of ϕ ensures that the matching requirements of a tiling are satisfied. Hence, we are left with the case that no model of ϕ contains a torus component.

We now observe that an $n \times m$ grid has treewidth $\min\{n, m\}$ and an $n \times m$ cylinder has treewidth $\min\{2n, m\}$ resp. $\min\{n, 2m\}$ for **E, W** resp. **S, N** cylinders. For the $n \times m$ grid, this follows from the k -cops and robber game, defined as follows. A position in the game is a pair (γ, r) , where $\gamma \subseteq [1..n] \times [1..m]$, $\text{card}(\gamma) = k$ and $r \in [1..n] \times [1..m] \setminus \gamma$. The game can move from (γ_i, r_i) to (γ_{i+1}, r_{i+1}) iff there exists a path between r_i and r_{i+1} in the restriction of the grid to $[1..n] \times [1..m] \setminus (\gamma_i \cap \gamma_{i+1})$. We say that k cops catch the robber iff every sequence of moves in the game is finite. It is known that, if the treewidth of the graph is greater or equal to k , then $k + 1$ cops catch the robber on a graph G [ST93]. Since $\min\{n, m\} - 1$ cops do not catch the robber (which can always move to the intersection of a cop-free row and a cop-free column) it follows that the treewidth of the grid is greater than $\min\{n, m\} - 1$. At the same time, there exists a tree decomposition of width $\min\{n, m\}$. For the $n \times m$ N-S

cylinder (the case of the E-W cylinder is analogous), we need extra n cops to prevent the robber escaping wrapping around the E-W axis, thus the treewidth is $\min\{2n, m\}$.

We now consider some $i \geq 0$ and some model (U, σ) with $\text{tw}((U, \sigma)) \geq 2i$ that does not contain torus components. Then, N decomposes into grid components and cylinder components. Because of our assumption $\text{tw}((U, \sigma)) \geq 2i$ there must be some component C of N with $\text{tw}(C) \geq 2i$. Now, we can deduce that C contains some square grid M of size $i \times i$ as a substructure (this follows from $2i \leq \min\{n, m\}$ for grids and from $2i \leq \min\{n, 2m\}$ resp. $2i \leq \min\{2n, m\}$ for cylinders). Hence, we can restrict our attention to models of ϕ that are square grids. Let M_1, M_2, \dots be a sequence of models with $M_n \models \phi$, where each M_n is a square grid of size $n \times n$. We are now going to construct a sequence of models G_1, G_2, \dots such that each G_n is a square grid of size n with $G_n \models \phi$, and each G_n is included in G_{n+1} , where we say a model I of ϕ is included in a model J of ϕ if I resp. J are square grids of size $n \times n$ resp. $m \times m$, and we have that $n \leq m$ and all tiles at positions (x, y) , with $x, y \in [1, n]$, are the same in both models. We construct the sequence G_1, G_2, \dots inductively, maintaining an infinite sequence of models M_1^n, M_2^n, \dots , for each $n \in \mathbb{N}$, such that G_n is included in all M_i^n : Take G_1 to be a model that consists of a single tile, which appears infinitely often at position $(1, 1)$ in the models M_1, M_2, \dots ; then we obtain the sequence M_1^1, M_2^1, \dots as the restriction of M_1, M_2, \dots to the models that include G_1 . Assume we have already defined G_n . Choose some square grid G_{n+1} of size $n+1$ that is included infinitely often in models of the sequence M_1^n, M_2^n, \dots (note that such a square grid must exist by the pigeonhole principle); then obtain the sequence $M_1^{n+1}, M_2^{n+1}, \dots$ by restricting the sequence to the M_1^n, M_2^n, \dots to the models that include G_{n+1} . With the sequence G_1, G_2, \dots at hand we now obtain a tiling of the plane: For position (i, j) , with $i, j \in \mathbb{N}$, simply take the tile at this position in $G_{\max\{i, j\}}$. We now verify that the horizontal resp. vertical requirements of a tiling are satisfied. We verify only the horizontal requirement (the vertical one is symmetric). Consider tiles at positions (i, j) and $(i+1, j)$. If $i \neq j$, then both tiles have been defined by $G_{\max\{i, j\}}$, and the matching requirement is satisfied because $G_{\max\{i, j\}}$ is a model of ϕ . If $i = j$ then the tile at position (i, i) is defined by G_i and the tile at position $(i+1, i)$ is defined by G_{i+1} . Now we observe that the tile at position (i, i) in G_i is the same as the tile at position (i, i) in G_{i+1} , because G_i is included in G_{i+1} , and the matching requirement is satisfied because G_{i+1} is a model of ϕ . □

This concludes the proof of the theorem. □

6. CONCLUSIONS

We have presented a decision procedure for the treewidth boundedness problem in the context of SLR, a generalization of Separation Logic over relational signatures, interpreted over structures. This procedure allows to define the precise fragment of SLR in which every formula has a bound on the treewidth of its models. This fragment is the right candidate for the definition of a fragment of SLR with a decidable entailment problem. Another application is checking that each graph defined by a treewidth-bounded SLR formula satisfies MSO-definable properties such as, e.g., Hamiltonicity, or 3-Colorability.

REFERENCES

- [ABIK22] Emma Ahrens, Marius Bozga, Radu Iosif, and Joost-Pieter Katoen. Reasoning about distributed reconfigurable systems. *Proc. ACM Program. Lang.*, 6(OOPSLA2):145–174, 2022. doi:10.1145/3563293.
- [ABS00] Serge Abiteboul, Peter Buneman, and Dan Suciu. *Data on the Web: From Relations to Semistructured Data and XML*. Morgan Kaufmann, 2000.
- [Acz77] Peter Aczel. An introduction to inductive definitions. In *Handbook of Mathematical Logic*, volume 90 of *Studies in Logic and the Foundations of Mathematics*, pages 739–782. Elsevier, 1977. doi:10.1016/S0049-237X(08)71120-0.
- [BBI22a] Marius Bozga, Lucas Bueri, and Radu Iosif. Decision problems in a logic for reasoning about reconfigurable distributed systems. In *Automated Reasoning - 11th International Joint Conference, IJCAR 2022, Haifa, Israel, August 8-10, 2022, Proceedings*, volume 13385 of *Lecture Notes in Computer Science*, pages 691–711. Springer, 2022. doi:10.1007/978-3-031-10769-6_40.
- [BBI22b] Marius Bozga, Lucas Bueri, and Radu Iosif. On an invariance problem for parameterized concurrent systems. In Bartek Klin, Slawomir Lasota, and Anca Muscholl, editors, *33rd International Conference on Concurrency Theory, CONCUR 2022, September 12-16, 2022, Warsaw, Poland*, volume 243 of *LIPICs*, pages 24:1–24:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.CONCUR.2022.24.
- [BCO06] Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem-Paul de Roever, editors, *Formal Methods for Components and Objects*, pages 115–137, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- [Ber66] Robert Berger. *The undecidability of the domino problem*. American Mathematical Soc., 1966.
- [BFPG14] James Brotherston, Carsten Fuhs, Juan Antonio Navarro Pérez, and Nikos Gorogiannis. A decision procedure for satisfiability in separation logic with inductive predicates. In Thomas A. Henzinger and Dale Miller, editors, *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS ’14, Vienna, Austria, July 14 - 18, 2014*, pages 25:1–25:10. ACM, 2014. doi:10.1145/2603088.2603091.
- [Bod98] Hans L. Bodlaender. A partial k -arboretum of graphs with bounded treewidth. *Theor. Comput. Sci.*, 209(1-2):1–45, 1998. doi:10.1016/S0304-3975(97)00228-4.
- [CDNQ12] Wei-Ngan Chin, Cristina David, Huu Hai Nguyen, and Shengchao Qin. Automated verification of shape, size and bag properties via user-defined predicates in separation logic. *Science of Computer Programming*, 77(9):1006–1036, 2012. The Programming Languages track at the 24th ACM Symposium on Applied Computing (SAC’09). doi:10.1016/j.scico.2010.07.004.
- [CDOY11] Cristiano Calcagno, Dino Distefano, Peter W. O’Hearn, and Hongseok Yang. Compositional shape analysis by means of bi-abduction. *J. ACM*, 58(6), December 2011. doi:10.1145/2049697.2049700.
- [CE12] Bruno Courcelle and Joost Engelfriet. *Graph Structure and Monadic Second-Order Logic: A Language-Theoretic Approach*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2012. doi:10.1017/CB09780511977619.
- [CGG02] Luca Cardelli, Philippa Gardner, and Giorgio Ghelli. A Spatial Logic for Querying Graphs. In Peter Widmayer, Francisco Triguero Ruiz, Rafael Morales Bueno, Matthew Hennessy, Stephan Eidenbenz, and Ricardo Conejo, editors, *Proceedings of the 29th International Colloquium on Automata, Languages and Programming (ICALP’02)*, volume 2380 of *Lecture Notes in Computer Science*, pages 597–610. Springer, July 2002.
- [CHO⁺11] Byron Cook, Christoph Haase, Joël Ouaknine, Matthew J. Parkinson, and James Worrell. Tractable reasoning in a fragment of separation logic. In Joost-Pieter Katoen and Barbara König, editors, *CONCUR 2011 - Concurrency Theory - 22nd International Conference, CONCUR 2011, Aachen, Germany, September 6-9, 2011. Proceedings*, volume 6901 of *Lecture Notes in Computer Science*, pages 235–249. Springer, 2011. doi:10.1007/978-3-642-23217-6_16.
- [Cou90] Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, 1990. doi:10.1016/0890-5401(90)90043-H.
- [COY07] Cristiano Calcagno, Peter W. O’Hearn, and Hongseok Yang. Local action and abstract separation logic. In *22nd IEEE Symposium on Logic in Computer Science (LICS 2007), 10-12 July 2007*,

- Wrocław, Poland, *Proceedings*, pages 366–378. IEEE Computer Society, 2007. doi:10.1109/LICS.2007.30.
- [DNM08] Pierpaolo Degano, Rocco De Nicola, and José Meseguer, editors. *Concurrency, Graphs and Models, Essays Dedicated to Ugo Montanari on the Occasion of His 65th Birthday*, volume 5065 of *Lecture Notes in Computer Science*. Springer, 2008. doi:10.1007/978-3-540-68679-8.
- [DPJ08] Dino Distefano and Matthew J. Parkinson J. jStar: towards practical verification for java. In *Proceedings of the 23rd ACM SIGPLAN Conference on Object-Oriented Programming Systems Languages and Applications, OOPSLA '08*, page 213–226, New York, NY, USA, 2008. Association for Computing Machinery. doi:10.1145/1449764.1449782.
- [EF95] Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite model theory*. Perspectives in Mathematical Logic. Springer, 1995.
- [EIP21a] Mnacho Echenim, Radu Iosif, and Nicolas Peltier. Decidable entailments in separation logic with inductive definitions: Beyond establishment. In Christel Baier and Jean Goubault-Larrecq, editors, *29th EACSL Annual Conference on Computer Science Logic (CSL 2021)*, volume 183 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:18, Dagstuhl, Germany, 2021. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CSL.2021.20.
- [EIP21b] Mnacho Echenim, Radu Iosif, and Nicolas Peltier. Unifying decidable entailments in separation logic with inductive definitions. In André Platzer and Geoff Sutcliffe, editors, *Automated Deduction - CADE 28 - 28th International Conference on Automated Deduction, Virtual Event, July 12-15, 2021, Proceedings*, volume 12699 of *Lecture Notes in Computer Science*, pages 183–199. Springer, 2021. doi:10.1007/978-3-030-79876-5_11.
- [EIP22] Mnacho Echenim, Radu Iosif, and Nicolas Peltier. Entailment is undecidable for symbolic heap separation logic formulæ with non-established inductive rules. *Inf. Process. Lett.*, 173:106169, 2022. doi:10.1016/J.IPL.2021.106169.
- [FG06] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006. doi:10.1007/3-540-29953-X.
- [IO01] Samin S. Ishtiaq and Peter W. O’Hearn. BI as an assertion language for mutable data structures. In Chris Hankin and Dave Schmidt, editors, *Conference Record of POPL 2001: The 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, London, UK, January 17-19, 2001*, pages 14–26. ACM, 2001.
- [IRS13] Radu Iosif, Adam Rogalewicz, and Jiri Simáček. The tree width of separation logic with recursive definitions. In Maria Paola Bonacina, editor, *Automated Deduction - CADE-24 - 24th International Conference on Automated Deduction, Lake Placid, NY, USA, June 9-14, 2013. Proceedings*, volume 7898 of *Lecture Notes in Computer Science*, pages 21–38. Springer, 2013. doi:10.1007/978-3-642-38574-2_2.
- [IZ23] Radu Iosif and Florian Zuleger. Expressiveness results for an inductive logic of separated relations. In Guillermo A. Pérez and Jean-François Raskin, editors, *34th International Conference on Concurrency Theory (CONCUR 2023)*, volume 279 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:20, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CONCUR.2023.20.
- [JKM⁺17] Christina Jansen, Jens Katelaan, Christoph Matheja, Thomas Noll, and Florian Zuleger. Unified reasoning about robustness properties of symbolic-heap separation logic. In Hongseok Yang, editor, *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, volume 10201 of *Lecture Notes in Computer Science*, pages 611–638. Springer, 2017. doi:10.1007/978-3-662-54434-1_23.
- [JM82] Neil D. Jones and Steven S. Muchnick. A flexible approach to interprocedural data flow analysis and programs with recursive data structures. In *Proceedings of the 9th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '82, page 66–74, New York, NY, USA, 1982. Association for Computing Machinery. doi:10.1145/582153.582161.
- [KR04] Viktor Kuncak and Martin Rinard. Generalized records and spatial conjunction in role logic. In *Static Analysis*, pages 361–376, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.
- [KZ20] Jens Katelaan and Florian Zuleger. Beyond symbolic heaps: Deciding separation logic with inductive definitions. In Elvira Albert and Laura Kovács, editors, *LPAR 2020: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Alicante, Spain*,

- May 22–27, 2020, volume 73 of *EPiC Series in Computing*, pages 390–408. EasyChair, 2020. doi:10.29007/vkmj.
- [LL23] Quang Loc Le and Xuan-Bach Dinh Le. An efficient cyclic entailment procedure in a fragment of separation logic. In Orna Kupferman and Pawel Sobocinski, editors, *Foundations of Software Science and Computation Structures - 26th International Conference, FoSSaCS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22–27, 2023, Proceedings*, volume 13992 of *Lecture Notes in Computer Science*, pages 477–497. Springer, 2023. doi:10.1007/978-3-031-30829-1_23.
- [Rey02] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22–25 July 2002, Copenhagen, Denmark, Proceedings*, pages 55–74. IEEE Computer Society, 2002. doi:10.1109/LICS.2002.1029817.
- [See91] D. Seese. The structure of the models of decidable monadic theories of graphs. *Annals of Pure and Applied Logic*, 53(2):169–195, 1991. doi:10.1016/0168-0072(91)90054-P.
- [ST93] P.D. Seymour and R. Thomas. Graph searching and a min-max theorem for tree-width. *Journal of Combinatorial Theory, Series B*, 58(1):22–33, 1993. doi:10.1006/jctb.1993.1027.
- [vD94] Dirk van Dalen. *Logic and structure (3. ed.)*. Universitext. Springer, 1994.

APPENDIX A. PROOFS FROM SECTION 2

A.1. Proof of Lemma 2.9. Without loss of generality, we consider that Δ is equality-free (Lemma 2.6). We propose a construction using an idea of Brotherston et al [BFPG14], that characterizes the satisfiability of a predicate by an abstraction consisting of tuples of parameters occurring in the interpretation of relation symbols. A similar abstraction has been used to check satisfiability of SLR formulæ [BBI22a].

Definition A.1. A *base* σ^\sharp is a mapping $\sigma^\sharp : \mathbb{R} \rightarrow \text{mpow}(\mathbb{V}^+)$ of relation symbols r into multisets of tuples of variables of length $\#r$ each. A base is *satisfiable* iff $\sigma^\sharp(r)$ is a set, for all $r \in \mathbb{R}$. Given a set of variables $X \subseteq \mathbb{V}$, let $\text{SatBase}(X)$ denote the set of satisfiable bases whose images contain only variables from X and let $\text{SatBase} \stackrel{\text{def}}{=} \text{SatBase}(\mathbb{V})$.

We consider three partial operations on SatBase . First, the *composition* is $\sigma_1^\sharp \otimes \sigma_2^\sharp \stackrel{\text{def}}{=} \sigma_1^\sharp \cup \sigma_2^\sharp$ if $\sigma_1^\sharp \cup \sigma_2^\sharp$ is satisfiable, and undefined, otherwise. Second, the *substitution* $\sigma^\sharp[x_1/y_1, \dots, x_n/y_n]$ replaces simultaneously each occurrence of x_j by y_j in σ^\sharp , for all $j \in [1..n]$. Third, given a set $X \subseteq \mathbb{V}$ of variables, the *projection* is:

$$\sigma^\sharp \downarrow_X \stackrel{\text{def}}{=} \lambda r . \{ \langle x_1, \dots, x_n \rangle \in \sigma^\sharp(r) \mid x_1, \dots, x_n \in X \}$$

Finally, for a qpf formula ϕ , we define:

$$\text{Base}(\phi) \stackrel{\text{def}}{=} \lambda r . \{ \langle x_1, \dots, x_n \rangle \mid r(x_1, \dots, x_n) \text{ occurs in } \phi \}$$

The predicates $\mathbf{B}^{\sigma^\sharp}$ of the SID $\overline{\Delta}$ are obtained by annotating each predicate \mathbf{B} that occurs in Δ with a satisfiable base σ^\sharp . The arity of each predicate $\mathbf{B}^{\sigma^\sharp}$ is the arity of \mathbf{B} . Then $\overline{\Delta}$ contains the rules:

$$\mathbf{B}_0^{\sigma_0^\sharp}(x_1, \dots, x_{\#B_0}) \leftarrow \exists y_1 \dots \exists y_m . \psi * \bigstar_{i=1}^{\ell} \mathbf{B}_i^{\sigma_i^\sharp}(z_{i,1}, \dots, z_{i,\#B_i})$$

where $\mathbf{B}_0(x_1, \dots, x_{\#B_0}) \leftarrow \exists y_1 \dots \exists y_m . \psi * \bigstar_{i=1}^{\ell} \mathbf{B}_i(z_{i,1}, \dots, z_{i,\#B_i})$ is a rule of Δ , ψ is the largest qpf formula from its right-hand side and, moreover, the following condition holds:

$$\sigma_0^\sharp = \left(\text{Base}(\psi) \otimes \bigotimes_{i=1}^{\ell} \sigma_i^\sharp[x_1/z_{i,1}, \dots, x_{\#B_i}/z_{i,\#B_i}] \right) \downarrow_{\{x_1, \dots, x_{\#B_0}\}}$$

In addition, $\overline{\Delta}$ contains a rule $A \leftarrow A^{\sigma^\sharp}$, for each satisfiable base $\sigma^\sharp \in \text{SatBase}$. Note that $\overline{\Delta}$ is finite, because Δ and SatBase are finite. We now prove the points from the statement of Lemma 2.9:

“ $\overline{\Delta}$ is all-satisfiable for A” We prove the following, more general, fact:

Fact A.2. For each predicate B^{σ^\sharp} that occurs in $\overline{\Delta}$, each predicate-free formula ϕ such that $B^{\sigma^\sharp} \Rightarrow_{\overline{\Delta}}^* \phi$ and each injective store \mathfrak{s} over $\{x_1, \dots, x_{\#B}\}$, there exists a structure $S = (U, \sigma)$ such that $S \models^{\mathfrak{s}} \phi$ and, for all $r \in \mathbb{R}$ and tuples of variables $\langle x_{j_1}, \dots, x_{j_{\#r}} \rangle$ with $j_k \in [1.. \#B]$ for all $k \in [1.. \#r]$, we have $\langle x_{j_1}, \dots, x_{j_{\#r}} \rangle \in \sigma^\sharp(r)$ if and only if $\langle \mathfrak{s}(x_{j_1}), \dots, \mathfrak{s}(x_{j_{\#r}}) \rangle \in \sigma(r)$.

Proof. We proceed by induction on the length of the unfolding $B^{\sigma^\sharp} \Rightarrow_{\overline{\Delta}}^* \phi$. The first step of the unfolding uses a rule:

$$B^{\sigma^\sharp}(x_1, \dots, x_{\#B}) \leftarrow \exists y_1 \dots \exists y_m . \psi * \bigstar_{i=1}^{\ell} B_i^{\sigma_i^\sharp}(z_{i,1}, \dots, z_{i, \#B_i})$$

and for all $i \in [1.. \ell]$, we have $B_i^{\sigma_i^\sharp} \Rightarrow_{\overline{\Delta}}^* \phi_i$, such that

$$\phi = \exists y_1 \dots \exists y_m . \psi * \bigstar_{i=1}^{\ell} \phi_i[x_1/z_{i,1}, \dots, x_{\#B_i}/z_{i, \#B_i}]$$

modulo a reordering of atoms. Let \mathfrak{s}' be an injective store that extends \mathfrak{s} over $\{y_1, \dots, y_m\}$. For all $i \in [1.. \ell]$, we define \mathfrak{s}_i over $\{x_1, \dots, x_{\#B_i}\}$ by $\mathfrak{s}_i(x_j) \stackrel{\text{def}}{=} \mathfrak{s}'(z_{i,j})$ for all $j \in [1.. \#B_i]$. By the inductive hypothesis applied to the unfolding $B_i^{\sigma_i^\sharp} \Rightarrow_{\overline{\Delta}}^* \phi_i$, there exists a structure $S_i = (U_i, \sigma_i)$ such that $S_i \models^{\mathfrak{s}_i} \phi_i$ and $\langle x_{j_1}, \dots, x_{j_{\#r}} \rangle \in \sigma_i^\sharp(r)$ if and only if $\langle \mathfrak{s}_i(x_{j_1}), \dots, \mathfrak{s}_i(x_{j_{\#r}}) \rangle \in \sigma_i(r)$. We can furthermore assume for all $i_1 \neq i_2 \in [1.. \ell]$ that $\text{supp}(\sigma_{i_1}) \cap \text{supp}(\sigma_{i_2}) = \{\mathfrak{s}'(x_1), \dots, \mathfrak{s}'(x_{\#B}), \mathfrak{s}'(y_1), \dots, \mathfrak{s}'(y_m)\}$. Since ψ is equality-free, there exists a structure S_ψ such that $S_\psi \models^{\mathfrak{s}'} \psi$, and $\text{Base}(\psi)$ is a satisfiable base. We then prove that the structure $S = (U, \sigma) \stackrel{\text{def}}{=} S_\psi \bullet S_1 \bullet \dots \bullet S_\ell$ is defined and meets the requirements from the statement:

▷ Let $r \in \mathbb{R}$, and suppose, for a contradiction, that there exists a tuple $\langle u_1, \dots, u_{\#r} \rangle \in \sigma_{i_1}(r) \cap \sigma_{i_2}(r)$ for $i_1 \neq i_2 \in [1.. \ell]$. From the induction hypothesis and because \mathfrak{s}' is injective, there exists $\langle z_1, \dots, z_{\#r} \rangle \in \sigma_{i_1}^\sharp[x_1/z_{i_1,1}, \dots, x_{\#B_{i_1}}/z_{i_1, \#B_{i_1}}](r) \cap \sigma_{i_2}^\sharp[x_1/z_{i_2,1}, \dots, x_{\#B_{i_2}}/z_{i_2, \#B_{i_2}}](r)$ with $\mathfrak{s}'(z_k) = u_k$ for all $k \in [1.. \#r]$. This contradicts that the following base composition:

$$\sigma_{i_1}^\sharp[x_1/z_{i_1,1}, \dots, x_{\#B_{i_1}}/z_{i_1, \#B_{i_1}}] \otimes \sigma_{i_2}^\sharp[x_1/z_{i_2,1}, \dots, x_{\#B_{i_2}}/z_{i_2, \#B_{i_2}}]$$

is defined. However, this must be the case, for the above derivation rule to exist in $\overline{\Delta}$. Therefore, the composition $S_1 \bullet \dots \bullet S_\ell$ is defined. The same type of argument can be used if the tuple occurs in the intersection between the interpretation of a relation symbol in S_ψ and S_i , for some $i \in [1.. \ell]$, thus the composition $S_\psi \bullet S_1 \bullet \dots \bullet S_\ell$ is defined.

▷ $S_\psi \bullet S_1 \bullet \dots \bullet S_\ell \models^{\mathfrak{s}'} \psi * \bigstar_{i=1}^{\ell} B_i^{\sigma_i^\sharp}(z_{i,1}, \dots, z_{i, \#B_i})$ by construction, thus $S \models^{\mathfrak{s}} \phi$.
 ▷ Let $r \in \mathbb{R}$ and $j_k \in [1.. \#B]$ for all $k \in [1.. \#r]$. We omit relations that occur in ψ since this is a simple case. Then $\langle x_{j_1}, \dots, x_{j_{\#r}} \rangle \in \sigma^\sharp(r)$ if and only if $\langle x_{j'_1}, \dots, x_{j'_{\#r}} \rangle \in \sigma_i^\sharp(r)$ for some $i \in [1.. \ell]$ and $x_{j_k} = z_{i, j'_k}$ for all $k \in [1.. \#r]$, if and only if $\langle \mathfrak{s}_i(x_{j'_1}), \dots, \mathfrak{s}_i(x_{j'_{\#r}}) \rangle \in \sigma_i(r)$ by induction hypothesis, if and only if $\langle \mathfrak{s}(x_{j_1}), \dots, \mathfrak{s}(x_{j_{\#r}}) \rangle \in \sigma(r)$. \square

By taking any injective store \mathfrak{s} over $\{x_1, \dots, x_{\#A}\}$, for every predicate-free formula ϕ such that $A^{\sigma^\sharp} \Rightarrow_{\overline{\Delta}}^* \phi$, we find a structure S such that $S \models^{\mathfrak{s}} \phi$, by Fact A.2. Therefore, $\overline{\Delta}$ is all-satisfiable for A.

“ $\llbracket \mathbf{A} \rrbracket_{\Delta} = \llbracket \mathbf{A} \rrbracket_{\overline{\Delta}}$ ” The inclusion $\llbracket \mathbf{A} \rrbracket_{\overline{\Delta}} \subseteq \llbracket \mathbf{A} \rrbracket_{\Delta}$ is immediate by simply removing the base annotations from any derivation run on $\overline{\Delta}$. For the converse, we prove the following fact:

Fact A.3. For every predicate \mathbf{B} of Δ , for every structure $\mathbf{S} \in \llbracket \mathbf{B} \rrbracket_{\Delta}$ and store \mathfrak{s} such that $\mathbf{S} \models^{\mathfrak{s}} \mathbf{B}(x_1, \dots, x_{\#\mathbf{B}})$, there exists a satisfiable base $\sigma^{\#}$, such that $\mathbf{S} \models^{\mathfrak{s}} \mathbf{B}^{\sigma^{\#}}(x_1, \dots, x_{\#\mathbf{B}})$, and $\langle x_{j_1}, \dots, x_{j_{\#r}} \rangle \in \sigma^{\#}(r)$ if and only if $\langle \mathfrak{s}(x_{j_1}), \dots, \mathfrak{s}(x_{j_{\#r}}) \rangle \in \sigma(r)$ for all $r \in \mathbb{R}$ and $j_k \in [1..\#\mathbf{B}]$ for all $k \in [1..\#r]$.

Proof. Let $\mathbf{S} \in \llbracket \mathbf{B} \rrbracket_{\Delta}$ and \mathfrak{s} be such that $\mathbf{S} \models^{\mathfrak{s}} \mathbf{B}(x_1, \dots, x_{\#\mathbf{B}})$. The proof is by induction on the derivation of $\mathbf{S} \models^{\mathfrak{s}} \mathbf{B}(x_1, \dots, x_{\#\mathbf{B}})$. Assume that the first step of this derivation uses a rule $\mathbf{B}(x_1, \dots, x_{\#\mathbf{B}}) \leftarrow \exists y_1 \dots \exists y_m \cdot \psi * \bigstar_{i=1}^{\ell} \mathbf{B}_i(z_{i,1}, \dots, z_{i,\#\mathbf{B}_i})$. Then, one can split $\mathbf{S} = \mathbf{S}_{\psi} \bullet \mathbf{S}_1 \bullet \dots \bullet \mathbf{S}_{\ell}$, such that $\mathbf{S}_{\psi} \models^{\mathfrak{s}'}$ ψ and $\mathbf{S}_i \models^{\mathfrak{s}_i}$ $\mathbf{B}_i(x_1, \dots, x_{\#\mathbf{B}_i})$, where \mathfrak{s}' extends \mathfrak{s}_0 and $\mathfrak{s}_i(x_j) = \mathfrak{s}'(z_{i,j})$ for all $i \in [1..\ell]$ and $j \in [1..\#\mathbf{B}_i]$. By the induction hypothesis, for each $i \in [1..\ell]$, there exists a satisfiable base $\sigma_i^{\#}$, such that $\mathbf{S}_i \models^{\mathfrak{s}_i} \mathbf{B}_i^{\sigma_i^{\#}}(x_1, \dots, x_{\#\mathbf{B}_i})$, and $\langle x_{j_1}, \dots, x_{j_{\#r}} \rangle \in \sigma_i^{\#}(r)$ if and only if $\langle \mathfrak{s}_i(x_{j_1}), \dots, \mathfrak{s}_i(x_{j_{\#r}}) \rangle \in \sigma(r)$. We consider:

$$\sigma^{\#} \stackrel{\text{def}}{=} \left(\text{Base}(\psi) \otimes \bigotimes_{i=1}^{\ell} \sigma_i^{\#}[x_1/z_{i,1}, \dots, x_{\#\mathbf{B}_i}/z_{i,\#\mathbf{B}_i}] \right) \downarrow_{\{x_1, \dots, x_{\#\mathbf{B}}\}}$$

and prove that $\sigma^{\#}$ is properly defined above:

- ▷ ψ is satisfiable, thus so is $\text{Base}(\psi) = \sigma_{\psi}^{\#}$. Projections over base pairs do not change the satisfiability, nor substitutions because $\overline{\Delta}$ is equality-free.
- ▷ We check the satisfiability of the composition in the definition of $\sigma^{\#}$. Suppose, for a contradiction, that there exists a tuple $\langle z_1, \dots, z_{\#r} \rangle \in \sigma_{\psi}^{\#}(r) \cap \sigma_i^{\#}[x_1/z_{i,1}, \dots, x_{\#\mathbf{B}_i}/z_{i,\#\mathbf{B}_i}](r)$ for some $i \in [1..\ell]$ and $z_k \in \{x_1, \dots, x_{\#\mathbf{B}}, y_1, \dots, y_m\}$ for all $k \in [1..\#r]$. Then there exist $j_1, \dots, j_{\#r} \in [1..\#\mathbf{B}_i]$ such that $\langle x_{j_1}, \dots, x_{j_{\#r}} \rangle \in \sigma_i^{\#}(r)$ and $\mathfrak{s}_i(x_{j_k}) = \mathfrak{s}'(z_k)$ for all $k \in [1..\#r]$. Thus we obtain $\langle \mathfrak{s}'(z_1), \dots, \mathfrak{s}'(z_{\#r}) \rangle \in \sigma_{\psi}(r) \cap \sigma_i(r)$, which contradicts the composition $\mathbf{S}_{\psi} \bullet \mathbf{S}_i$. A similar argument ensures that no collisions occur between \mathbf{S}_{i_1} and \mathbf{S}_{i_2} , for any $i_1 \neq i_2 \in [1..\ell]$.

The last part of Fact A.3 is similar to the proof of Fact A.2. □

Therefore, any model \mathbf{S} of $\llbracket \mathbf{A} \rrbracket_{\Delta}$ is a model of $\llbracket \mathbf{A}^{\sigma^{\#}} \rrbracket_{\overline{\Delta}}$ with an appropriate base $\sigma^{\#}$. The upper bound on $\text{predNo}(\overline{\Delta})$ is obtained by noticing that, for each predicate symbol \mathbf{A} that occurs in Δ , we introduce at most $\text{relNo}(\Delta) \cdot \#\mathbf{A}^{\max\text{RelArity}(\Delta)} \leq \text{relNo}(\Delta) \cdot \max\text{Vars}(\Delta)^{\max\text{RelArity}(\Delta)}$ new predicate symbols. □

A.2. Proof of Lemma 2.10. The proof follows a generic guideline. First, recall that for any set of structures \mathcal{S} we have

$$\text{tw}(\mathcal{S}) = \max_{\mathbf{S} \in \mathcal{S}} \text{tw}(\mathbf{S}) = \max_{\mathbf{S} \in \mathcal{S}} \min \{ \text{wd}(T) \mid T \text{ is a tree decomposition of } \mathbf{S} \}.$$

Therefore, in order to prove an inequality of the form $\text{tw}(\llbracket \phi^{\exists} \rrbracket) \leq \text{tw}(\llbracket \psi^{\exists} \rrbracket) + k$ for ϕ, ψ two qpf formulæ, we make use of the alternating max and min by proving the following:

for every structure \mathbf{S} and store \mathfrak{s} such that $\mathbf{S} \models^{\mathfrak{s}} \phi$

there exists a structure \mathbf{S}' and a store \mathfrak{s}' such that $\mathbf{S}' \models^{\mathfrak{s}'} \psi$ and

for every tree decomposition T' of \mathbf{S}'

there exists a tree decomposition T of \mathbf{S} such that $\text{wd}(T) \leq \text{wd}(T') + k$.

- (1) The first point is immediate since $\llbracket (\phi * \bigstar_{i=1}^k x_0 = x_i) \rrbracket^{\exists} \subseteq \llbracket \phi \rrbracket^{\exists}$.
- (2) Since $\llbracket (\phi * \bigstar_{i=1}^k x_0 \neq x_i) \rrbracket^{\exists} \subseteq \llbracket \phi \rrbracket^{\exists}$, we obtain immediately $\text{tw}(\llbracket (\phi * \bigstar_{i=1}^k x_0 \neq x_i) \rrbracket^{\exists}) \leq \text{tw}(\llbracket \phi \rrbracket^{\exists})$. For the reverse inequality, recall that $\phi * \bigstar_{i=1}^k x_0 \neq x_i$ is satisfiable. Then ϕ must also be satisfiable, so let $\mathbf{S} = (\mathbf{U}, \sigma)$ be a model and \mathfrak{s} a store such that $\mathbf{S} \models^{\mathfrak{s}} \phi$. We distinguish two cases:
- ▷ if $\mathfrak{s}(x_0) \neq \mathfrak{s}(x_i)$ for all $i \in [1..k]$ then let $\mathbf{S}' = \mathbf{S}$, $\mathfrak{s}' = \mathfrak{s}$ hence $\mathbf{S}' \models^{\mathfrak{s}'} \phi * \bigstar_{i=1}^k x_0 \neq x_i$ and $\text{tw}(\mathbf{S}) = \text{tw}(\mathbf{S}')$.
 - ▷ if $\mathfrak{s}(x_0) = \mathfrak{s}(x_i)$ for some $i \in [1..k]$ then let us consider a new fresh element $e \in \mathbf{U}$ and define a new store \mathfrak{s}' by $\mathfrak{s}'(y) \stackrel{\text{def}}{=} e$ if $y \approx_{\phi} x_0$, and $\mathfrak{s}'(y) \stackrel{\text{def}}{=} \mathfrak{s}(y)$ otherwise. We define the $\mathbf{S}' = (\mathbf{U} \cup \{e\}, \sigma')$ as follows. For every $r \in \mathbb{R}$, for every tuple $\langle u_1, \dots, u_{\#r} \rangle \in \sigma(r)$, there exists a unique relation atom $r(y_1, \dots, y_{\#r})$ occurring in ϕ such that $\mathfrak{s}(y_j) = u_j$ for all $j \in [1..\#r]$. Then, add the tuple $(\mathfrak{s}'(y_1), \dots, \mathfrak{s}'(y_{\#r}))$ to $\sigma'(r)$. By construction $\mathbf{S}' \models^{\mathfrak{s}'} \phi * \bigstar_{i=1}^k x_0 \neq x_i$. Let T' be a tree decomposition of \mathbf{S}' . We define T by removing the element e from T' and adding $\mathfrak{s}(x_0)$ in every node of T' . T is a tree decomposition of \mathbf{S} of width at most $\text{wd}(T') + 1$. Therefore $\text{wd}(T) \leq \text{wd}(T') + 1$, hence the result.

In both cases we obtain the expected result $\text{tw}(\llbracket \phi \rrbracket^{\exists}) \leq \text{tw}(\llbracket (\phi * \bigstar_{i=1}^k x_0 \neq x_i) \rrbracket^{\exists}) + 1$.

- (3) We prove $\text{tw}(\llbracket (\phi * r(x_1, \dots, x_k)) \rrbracket^{\exists}) \leq \text{tw}(\llbracket \phi \rrbracket^{\exists}) + k$. Let $\mathbf{S} = (\mathbf{U}, \sigma)$ and \mathfrak{s} such that $\mathbf{S} \models^{\mathfrak{s}} \phi * r(x_1, \dots, x_k)$. We define $\mathbf{S}' = (\mathbf{U}, \sigma')$ from \mathbf{S} by removing the tuple $\langle \mathfrak{s}(x_1), \dots, \mathfrak{s}(x_k) \rangle$ from $\sigma(r)$. Let T' be a tree decomposition of \mathbf{S}' . We define T by adding the elements $\mathfrak{s}(x_1), \dots, \mathfrak{s}(x_k)$ to every node in T' . This construction does not break connectedness of the subtree of T containing any element, T still contains a node with all elements of any relation in σ' , and moreover (since T' is not empty) T contains a node (in fact all nodes) with elements $\mathfrak{s}(x_1), \dots, \mathfrak{s}(x_k)$ simultaneously. Therefore T is a tree decomposition of \mathbf{S} of width at most $\text{wd}(T') + k$.

We now prove that $\text{tw}(\llbracket \phi \rrbracket^{\exists}) \leq \text{tw}(\llbracket (\phi * r(x_1, \dots, x_k)) \rrbracket^{\exists}) + 1$. Recall $\phi * r(x_1, \dots, x_k)$ is satisfiable from the hypothesis. Let $\mathbf{S} = (\mathbf{U}, \sigma) \models^{\mathfrak{s}} \phi$, and we distinguish two cases:

- ▷ If $\langle \mathfrak{s}(x_1), \dots, \mathfrak{s}(x_k) \rangle \notin \sigma(r)$, then consider $\mathfrak{s}' = \mathfrak{s}$ and \mathbf{S}' obtained by adding the above tuple to $\sigma(r)$. Then $\mathbf{S}' \models^{\mathfrak{s}'} \phi * r(x_1, \dots, x_k)$ and for any tree decomposition T' of \mathbf{S}' we have $T = T'$ is also a tree decomposition for \mathbf{S} hence ensuring $\text{wd}(T) = \text{wd}(T')$.
- ▷ If $\langle \mathfrak{s}(x_1), \dots, \mathfrak{s}(x_k) \rangle \in \sigma(r)$ then, because $\phi * r(x_1, \dots, x_k)$ is satisfiable, there must exist variables x'_1, \dots, x'_k and $j \in [1..k]$ such that $r(x'_1, \dots, x'_k)$ occurs in ϕ , $\mathfrak{s}(x_i) = \mathfrak{s}(x'_i)$, for every $i \in [1..k]$, and moreover $x_j \not\approx_{\phi} x'_j$. Let us consider a new fresh element $e \in \mathbf{U}$ and define a new store \mathfrak{s}' by $\mathfrak{s}'(y) = e$ if $y \approx_{\phi} x'_j$, and $\mathfrak{s}'(y) = \mathfrak{s}(y)$ otherwise. Let $\mathbf{S}' = (\mathbf{U} \cup \{e\}, \sigma')$ defined as follows. For every $r' \in \mathbb{R}$, for every tuple $\langle u_1, \dots, u_{\#r'} \rangle \in \sigma(r')$, there exists a unique relation atom $r'(y_1, \dots, y_{\#r'})$ occurring in ϕ such that $\mathfrak{s}(y_j) = u_j$ for all $j \in [1..\#r']$. Then, add the tuple $(\mathfrak{s}'(y_1), \dots, \mathfrak{s}'(y_{\#r'}))$ to $\sigma'(r')$. Finally, add the tuple $\langle \mathfrak{s}'(x_1), \dots, \mathfrak{s}'(x_k) \rangle$ to $\sigma'(r)$. By construction $\mathbf{S}' \models^{\mathfrak{s}'} \phi * r(x_1, \dots, x_k)$. Let T' be a tree decomposition of \mathbf{S}' . We define T by removing the element e from T' and adding $\mathfrak{s}(x'_j)$ in every node of T' . T is a tree decomposition of \mathbf{S} of width at most $\text{wd}(T') + 1$, that is, $\text{wd}(T) \leq \text{wd}(T') + 1$.

In both cases we obtain the expected result, that is, $\text{tw}(\llbracket \phi \rrbracket^{\exists}) \leq \text{tw}(\llbracket (\phi * r(x_1, \dots, x_k)) \rrbracket^{\exists}) + 1$.

- (4) Similar to point (3), it generalizes the (right) inequality from formulæ consisting of a single relation atom to formulæ consisting of arbitrarily many relation atoms. \square

A.3. Proof of Lemma 2.11. Note that $\text{fv}(\eta) \subseteq F \subseteq \text{fv}(\phi)$ and henceforth, $\text{fv}(\phi * \eta) = \text{fv}(\phi)$. We follow a similar strategy as in the proof of Lemma 2.10. Consider a structure $\mathbf{S} = (\mathbf{U}, \sigma)$ and a store \mathfrak{s} , such that $\mathbf{S} \models^{\mathfrak{s}} \phi * \eta$. We build a store \mathfrak{s}' and a structure \mathbf{S}' such that $\mathbf{S}' \models^{\mathfrak{s}'} \phi * \psi$ and $\text{tw}(\mathbf{S}) \leq \text{tw}(\mathbf{S}') + \text{card}(F)$. First, consider a store \mathfrak{s}'' , that is canonical for ψ (Definition 3.1), and a structure $\mathbf{S}'' = (\mathbf{U}'', \sigma'')$, such that $\mathbf{S}'' \models^{\mathfrak{s}''} \psi$. Such store and structure exist, because ψ is satisfiable. Assume without loss of generality that \mathbf{S}'' and \mathbf{S} are disjoint structures, that is, $\mathfrak{s}''(\text{fv}(\psi)) \cap \mathfrak{s}(\text{fv}(\phi)) = \emptyset$. Second, we define the store \mathfrak{s}' , as follows:

$$\mathfrak{s}'(y) \stackrel{\text{def}}{=} \begin{cases} \mathfrak{s}''(y) & \text{if } y \in \text{fv}(\psi) \\ \mathfrak{s}''(y') & \text{if } y \in \text{fv}(\phi) \setminus \text{fv}(\psi) \text{ and there exists } y' \in F \text{ such that } y \approx_{\phi} y' \\ \mathfrak{s}(y) & \text{otherwise} \end{cases}$$

Note that the definition of \mathfrak{s}' is consistent. In particular, for any $y \in \text{fv}(\phi) \setminus \text{fv}(\psi)$ there exists at most one variable $y' \in F$, such that $y \approx_{\phi} y'$, because otherwise, the hypothesis $x \not\approx_{\phi} y$ for all $x, y \in F$ would not hold. We build now the structure $\mathbf{S}' = (\mathbf{U}', \sigma')$ where $\mathbf{U}' \stackrel{\text{def}}{=} \mathbf{U} \cup \mathbf{U}''$ and $\sigma'(r)$ is defined for every relation symbol $r \in \mathbb{R}$ as follows:

- ▷ add each tuple $\langle u_1, \dots, u_{\#r} \rangle \in \sigma''(r)$ to $\sigma'(r)$,
- ▷ for every tuple $\langle u_1, \dots, u_{\#r} \rangle \in \sigma(r)$, there exists a unique relation atom $r(y_1, \dots, y_{\#r})$ occurring in ϕ , such that $\mathfrak{s}(y_i) = u_i$ for all $i \in [1..\#r]$; we add the tuple $\langle \mathfrak{s}'(y_1), \dots, \mathfrak{s}'(y_{\#r}) \rangle$ to $\sigma'(r)$.
- ▷ nothing else belongs to $\sigma'(r)$.

This construction guarantees that $\mathbf{S}' \models^{\mathfrak{s}'} \phi * \psi$. Equality and disequality atoms in $\phi * \psi$ are satisfied, by the definition of \mathfrak{s}' . With regard to relation atoms, notice that no tuple is added twice to $\sigma'(r)$ in the definition above. That is, if some $\langle \mathfrak{s}'(y_1), \dots, \mathfrak{s}'(y_{\#r}) \rangle$ obtained from $\sigma(r)$ exists also in $\sigma''(r)$, then $\phi * \psi$ would not be satisfiable. Let T' be a tree decomposition of \mathbf{S}' . We define a tree decomposition T by:

- ▷ removing $\{\mathfrak{s}'(y) \mid y \in \text{fv}(\psi)\}$ from every bag of T' ,
- ▷ adding $\{\mathfrak{s}(y) \mid y \in F\}$ to every bag of T' .

The result is a tree decomposition T of \mathbf{S} of width $\text{wd}(T) \leq \text{wd}(T') + \text{card}(F)$. Since the choice of T' was arbitrary, we obtain $\text{tw}(\mathbf{S}) \leq \text{tw}(\mathbf{S}') + \text{card}(F)$. Since the choice of \mathbf{S} was arbitrary, we obtain $\text{tw}(\llbracket (\phi * \eta)^{\exists} \rrbracket) \leq \text{tw}(\llbracket (\phi * \psi)^{\exists} \rrbracket) + \text{card}(F)$. \square

APPENDIX B. PROOFS FROM SECTION 3

B.1. Proof of Lemma 3.24. (1) “ $\mathbf{EF}^*(\text{split}(\mathcal{S})) \subseteq \text{split}(\mathbf{EF}^*(\mathcal{S}))$ ” By induction on the derivation of $\mathbf{S} \in \mathbf{EF}^*(\text{split}(\mathcal{S}))$ from $\text{split}(\mathcal{S})$.

Base case: Let $\mathbf{S} \in \text{split}(\mathcal{S})$. We have:

$$\exists \mathbf{S}' \in \mathcal{S}. \mathbf{S} \in \text{split}(\mathbf{S}') \Rightarrow \exists \mathbf{S}' \in \mathbf{EF}^*(\mathcal{S}). \mathbf{S} \in \text{split}(\mathbf{S}') \Rightarrow \mathbf{S} \in \text{split}(\mathbf{EF}^*(\mathcal{S}))$$

Induction step: Assume $\mathbf{S} = (\mathbf{S}_1 \bullet \mathbf{S}_2)_{/\approx}$ for some $\mathbf{S}_1, \mathbf{S}_2 \in \mathbf{EF}^*(\text{split}(\mathcal{S}))$ and \approx satisfying the conditions of Definition 3.11 for external fusion of \mathbf{S}_1 and \mathbf{S}_2 . Moreover, assume the induction hypothesis $\mathbf{S}_1, \mathbf{S}_2 \in \text{split}(\mathbf{EF}^*(\mathcal{S}))$. Then

$$\begin{aligned} & \mathbf{S}_1, \mathbf{S}_2 \in \text{split}(\mathbf{EF}^*(\mathcal{S})), \mathbf{S} = (\mathbf{S}_1 \bullet \mathbf{S}_2)_{/\approx} \Rightarrow \\ & (\exists \mathbf{S}'_1, \mathbf{S}'_2 \in \mathbf{EF}^*(\mathcal{S}). \mathbf{S}_1 \sqsubseteq^{mc} \mathbf{S}'_1, \mathbf{S}_2 \sqsubseteq^{mc} \mathbf{S}'_2), \mathbf{S} = (\mathbf{S}_1 \bullet \mathbf{S}_2)_{/\approx} \Rightarrow \\ & \exists \mathbf{S}'_1, \mathbf{S}'_2 \in \mathbf{EF}^*(\mathcal{S}). (\exists \approx' . \mathbf{S} \sqsubseteq^{mc} (\mathbf{S}'_1 \bullet \mathbf{S}'_2)_{/\approx'}) \Rightarrow \end{aligned}$$

$$\exists S' \in \text{EF}^*(\mathcal{S}). S \sqsubseteq^{mc} S' \Rightarrow S \in \text{split}(\text{EF}^*(\mathcal{S}))$$

In the above, the equivalence \approx' is taken as the extension by equality of \approx and henceforth it satisfies the conditions of Definition 3.11 for external fusion of S'_1, S'_2 .

“ $\text{split}(\text{EF}^*(\mathcal{S})) \subseteq \text{EF}^*(\text{split}(\mathcal{S}))$ ” By induction on the derivation of $S' \in \text{EF}^*(\mathcal{S})$ from \mathcal{S} .

Base case: Let $S \in \text{split}(S')$ for some $S' \in \mathcal{S}$. We have:

$$S \in \text{split}(S') \Rightarrow S \in \text{EF}^*(\text{split}(S')) \Rightarrow S \in \text{EF}^*(\text{split}(\mathcal{S}))$$

Induction step: Assume $S \in \text{split}(S')$ for some $S' = (S'_1 \bullet S'_2)_{/\approx'}$ for some $S'_1, S'_2 \in \text{EF}^*(\mathcal{S})$ and equivalence \approx' satisfying the conditions of Definition 3.11 for external fusion of S'_1, S'_2 . Assume the induction hypothesis, that is, $\text{split}(S'_1), \text{split}(S'_2) \subseteq \text{EF}^*(\text{split}(\mathcal{S}))$. Then:

$$\begin{aligned} \text{split}(S'_1), \text{split}(S'_2) \subseteq \text{EF}^*(\text{split}(\mathcal{S})), S' = (S'_1 \bullet S'_2)_{/\approx'}, S \in \text{split}(S') \Rightarrow \\ \text{split}(S'_1), \text{split}(S'_2) \subseteq \text{EF}^*(\text{split}(\mathcal{S})), S' = (S'_1 \bullet S'_2)_{/\approx'}, S \sqsubseteq^{mc} S' \end{aligned}$$

We distinguish two sub-cases:

- S is a maximally connected substructure of S'_1 (the case of S'_2 is symmetric) not affected by the external fusion defined by \approx' :

$$\begin{aligned} \text{split}(S'_1) \subseteq \text{EF}^*(\text{split}(\mathcal{S})), (\exists S_1. S_1 \sqsubseteq^{mc} S'_1, S = S_1) \Rightarrow \\ \text{split}(S'_1) \subseteq \text{EF}^*(\text{split}(\mathcal{S})), S \in \text{split}(S'_1) \Rightarrow S \in \text{EF}^*(\text{split}(\mathcal{S})) \end{aligned}$$

- S is a connected structure including several maximally connected substructures, at least one from each S'_i , for $i = 1, 2$:

$$\begin{aligned} \text{split}(S'_1), \text{split}(S'_2) \subseteq \text{EF}^*(\text{split}(\mathcal{S})), \\ \exists k_1 \geq 1. \exists S_{1,1} \dots \exists S_{1,k_1}. S_{1,i} \sqsubseteq^{mc} S'_1 \text{ for all } i \in [1..k_1], \\ \exists k_2 \geq 1. \exists S_{2,1} \dots \exists S_{2,k_2}. S_{2,j} \sqsubseteq^{mc} S'_2 \text{ for all } j \in [1..k_2], \\ (\exists \approx . S = (S_{1,1} \bullet \dots \bullet S_{1,k_1} \bullet S_{2,1} \bullet \dots \bullet S_{2,k_2})_{/\approx}), S \text{ connected} \Rightarrow \\ \text{split}(S'_1), \text{split}(S'_2) \subseteq \text{EF}^*(\text{split}(\mathcal{S})), \\ \exists k_1 \geq 1. \exists S_{1,1} \dots \exists S_{1,k_1}. S_{1,i} \in \text{split}(S'_1) \text{ for all } i \in [1..k_1], \\ \exists k_2 \geq 1. \exists S_{2,1} \dots \exists S_{2,k_2}. S_{2,j} \in \text{split}(S'_2) \text{ for all } j \in [1..k_2], \\ (\exists \approx . S = (S_{1,1} \bullet \dots \bullet S_{1,k_1} \bullet S_{2,1} \bullet \dots \bullet S_{2,k_2})_{/\approx}), S \text{ connected} \Rightarrow \\ \exists k_1 \geq 1. \exists S_{1,1} \dots \exists S_{1,k_1}. S_{1,i} \in \text{EF}^*(\text{split}(\mathcal{S})) \text{ for all } i \in [1..k_1], \\ \exists k_2 \geq 1. \exists S_{2,1} \dots \exists S_{2,k_2}. S_{2,j} \in \text{EF}^*(\text{split}(\mathcal{S})) \text{ for all } j \in [1..k_2], \\ (\exists \approx . S = (S_{1,1} \bullet \dots \bullet S_{1,k_1} \bullet S_{2,1} \bullet \dots \bullet S_{2,k_2})_{/\approx}), S \text{ connected} \Rightarrow S \in \text{EF}^*(\text{split}(\mathcal{S})) \end{aligned}$$

In the above, the equivalence \approx is the restriction of \approx' to the substructures included in the composition. As \approx' is conforming for external fusion of S'_1, S'_2 and since the resulting structure S is connected, it is always possible to obtain S as a sequence of external fusions conforming to Definition 3.11 from the respective substructures.

(2) “ $\text{IEF}^*(\text{split}(\mathcal{S})) \subseteq \text{split}(\text{IEF}^*(\mathcal{S}))$ ” By induction on the derivation of $S \in \text{IEF}^*(\text{split}(\mathcal{S}))$ from $\text{split}(\mathcal{S})$. The induction proceeds as for (1), with one additional case in the induction step.

Induction step: Let $S = (S_1)_{/\approx}$ for some $S_1 \in \text{IEF}^*(\text{split}(\mathcal{S}))$ and equivalence relation \approx conforming to internal fusion of S_1 . Moreover, assume the induction hypothesis $S_1 \in$

$\text{split}(\text{IEF}^*(\mathcal{S}))$. Then

$$\begin{aligned} & S_1 \in \text{split}(\text{IEF}^*(\mathcal{S})), S = (S_1)_{/\approx} \Rightarrow \\ & (\exists S'_1 \in \text{IEF}^*(\mathcal{S}). S_1 \sqsubseteq^{mc} S'_1), S = (S_1)_{/\approx} \Rightarrow \\ & \exists S'_1 \in \text{IEF}^*(\mathcal{S}). (\exists \approx' . S \sqsubseteq^{mc} (S'_1)_{/\approx'}) \Rightarrow \\ & \exists S' \in \text{IEF}^*(\mathcal{S}), S \sqsubseteq^{mc} S' \Rightarrow S \in \text{split}(\text{IEF}^*(\mathcal{S})) \end{aligned}$$

In the above, the equivalence \approx' is taken as the extension by equality of \approx and hence conforming for internal fusion of structure S'_1 .

” $\text{split}(\text{IEF}^*(\mathcal{S})) \subseteq \text{IEF}^*(\text{split}(\mathcal{S}))$ ” By induction on the derivation of $S' \in \text{IEF}^*(\mathcal{S})$ from \mathcal{S} . The induction proceeds as for (1), with one additional case in the induction step.

Induction step: Let $S \in \text{split}(S')$ for some $S' = (S'_1)_{/\approx'}$ for some $S'_1 \in \text{IEF}^*(\mathcal{S})$ and equivalence \approx' conforming for internal fusion of S'_1 . Moreover, assume the induction hypothesis $\text{split}(S'_1) \subseteq \text{IEF}^*(\text{split}(\mathcal{S}))$. Then

$$\begin{aligned} & \text{split}(S'_1) \subseteq \text{IEF}^*(\text{split}(\mathcal{S})), S' = (S'_1)_{/\approx'}, S \in \text{split}(S') \Rightarrow \\ & \text{split}(S'_1) \subseteq \text{IEF}^*(\text{split}(\mathcal{S})), S' = (S'_1)_{/\approx'}, S \sqsubseteq^{mc} S' \Rightarrow \\ & \quad \text{split}(S'_1) \subseteq \text{IEF}^*(\text{split}(\mathcal{S})), \\ & \quad \exists k \geq 1. \exists S_{1,1} \dots \exists S_{1,k}. S_{1,i} \sqsubseteq^{mc} S'_1 \text{ for all } i \in [1..k], \\ & \quad (\exists \approx . S = (S_{1,1} \bullet \dots \bullet S_{1,k})_{/\approx}, S \text{ connected}) \Rightarrow \\ & \quad \text{split}(S'_1) \subseteq \text{IEF}^*(\text{split}(\mathcal{S})), \\ & \quad \exists k \geq 1. \exists S_{1,1} \dots \exists S_{1,k}. S_{1,i} \in \text{split}(S'_1) \text{ for all } i \in [1..k], \\ & \quad (\exists \approx . S = (S_{1,1} \bullet \dots \bullet S_{1,k})_{/\approx}, S \text{ connected}) \Rightarrow \\ & \quad \exists k \geq 1. \exists S_{1,1} \dots \exists S_{1,k}. S_{1,i} \in \text{IEF}^*(\text{split}(\mathcal{S})) \text{ for all } i \in [1..k], \\ & \quad (\exists \approx . S = (S_{1,1} \bullet \dots \bullet S_{1,k})_{/\approx}, S \text{ connected}) \Rightarrow S \in \text{IEF}^*(\text{split}(\mathcal{S})) \end{aligned}$$

In the above, the equivalence \approx is taken as the restriction of \approx' to the maximal connected substructures included in the construction of connected S . Henceforth, \approx is conforming for internal fusion as well. As the resulting structure S is connected, it is always possible to construct it in $\text{IEF}^*(\text{split}(\mathcal{S}))$ i.e., first by using external fusion conforming to Definition 3.11 to connect all the included substructures and second, by using internal fusion to further restrict the result if needed. \square

B.2. Proof of Lemma 3.31. “ $\text{split}(\llbracket \Delta \rrbracket_A^c) \subseteq \llbracket \Gamma \rrbracket_P^c$ ” We prove first the following fact:

Fact B.1. Let $B_0(x_1, \dots, x_{\#B_0}) \Rightarrow_{\Delta}^* \exists y_1 \dots \exists y_n . \phi$ be a complete Δ -unfolding, where ϕ is a qpf formula, \mathfrak{s} a store injective over $\{x_1, \dots, x_{\#B_0}\} \cup \{y_1, \dots, y_n\}$, $S = (U, \sigma)$ a structure such that $S \models^{\mathfrak{s}} \phi$ and $S' = (U', \sigma')$ a structure, such that $S' \sqsubseteq^{mc} S$ and $\text{supp}(\sigma') \cap \{\mathfrak{s}(x_1), \dots, \mathfrak{s}(x_{\#B_0})\} \neq \emptyset$. Then, there exist a nonempty set $J_0 \subseteq [1..\#B_0]$, equivalence relation $\xi_0 \subseteq J_0 \times J_0$ and complete Γ -unfolding $B_0(x_1, \dots, x_{\#B_0})_{/\xi_0} \Rightarrow_{\Gamma}^* \exists y_1 \dots \exists y_n . \phi'$, where ϕ' is a qpf formula, such that $S' \models^{\mathfrak{s}} \phi'$.

Proof. By induction on the length of the Δ -unfolding. Assume the first rule in this unfolding to be of the form (3.1), for a qpf formula ψ_0 . Then, there exist:

- ▷ unfoldings $B_i(x_1, \dots, x_{\#B_i}) \Rightarrow_{\Delta}^* \exists y_{j_{i,1}} \dots \exists y_{j_{i,k_i}} . \phi_i$, where $y_{j_{i,1}}, \dots, y_{j_{i,k_i}} \in \{y_1, \dots, y_n\}$ and ϕ_i are qpf formulae, for all $i \in [1..\ell]$, and

▷ structures $S_0 = (U_0, \sigma_0), \dots, S_\ell = (U_\ell, \sigma_\ell)$, such that $S_0 \bullet \dots \bullet S_\ell = S$, $S_0 \models^s \psi_0$ and $S_i \models^s \phi_i \theta_i$, where $\theta_i \stackrel{\text{def}}{=} [x_1/z_{i,1}, \dots, x_{\#B_i}/z_{i,\#B_i}]$, for all $i \in [1..\ell]$.

Since $S' \sqsubseteq^{mc} S_0 \bullet \dots \bullet S_\ell$ is a maximally connected structure, there must exist structures $S'_0 = (U_0, \sigma'_0)$, $S'_{i_1} = (U_{i_1}, \sigma'_{i_1}), \dots, S'_{i_k} = (U_{i_k}, \sigma'_{i_k})$, for $i_1, \dots, i_k \in [1..\ell]$, such that:

▷ $S' = S'_0 \bullet S'_{i_1} \bullet \dots \bullet S'_{i_k}$,

▷ $\text{supp}(\sigma'_{i_h}) \cap \{\mathfrak{s}(z_{i_h,1}), \dots, \mathfrak{s}(z_{i_h,\#B_{i_h}})\} \neq \emptyset$ and $S'_{i_h} \sqsubseteq^{mc} S_{i_h}$, for all $h \in [1..k]$,

▷ $\text{supp}(\sigma') \cap \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\} = \emptyset$, for all $i \in [1..\ell] \setminus \{i_1, \dots, i_k\}$.

Since $S_{i_h} \models^s \phi_{i_h} \theta_{i_h}$, we have $S_{i_h} \models^{s \circ \theta_{i_h}^{-1}} \phi_{i_h}$, for all $h \in [1..k]$. By the inductive hypothesis, there exist nonempty sets $J_{i_1} \subseteq [1..\#B_{i_1}], \dots, J_{i_k} \subseteq [1..\#B_{i_k}]$, equivalence relations $\xi_{i_1} \subseteq J_{i_1} \times J_{i_1}, \dots, \xi_{i_k} \subseteq J_{i_k} \times J_{i_k}$ and complete Γ -unfoldings:

$$B_{i_h}(x_1, \dots, x_{\#B_{i_h}}) / \xi_{i_h} \Rightarrow_{\Gamma}^* \exists y_{j_{i_h,1}} \dots \exists y_{j_{i_h,k_{i_h}}} \cdot \phi'_{i_h}$$

such that $S'_{i_h} \models^{s \circ \theta_{i_h}^{-1}} \phi'_{i_h}$, for all $h \in [1..k]$. Then, we define:

▷ sets $\bar{J}_{i_h} \stackrel{\text{def}}{=} [1..\#B_{i_h}] \setminus J_{i_h}$, for all $h \in [1..k]$,

▷ qpf formulæ ψ'_0 and ψ''_0 satisfying points (1) and (2) from the construction of Γ ,

▷ an equivalence relation $\Xi \stackrel{\text{def}}{=} (\text{conn}(\psi'_0) \cup \bigcup_{h=1}^k \xi_{i_h}(B_{i_h}(z_{i_h,1}, \dots, z_{i_h,\#B_{i_h}})))^{\bar{}}$.

We argue that the construction of the formulæ ψ'_0 and ψ''_0 is effective. There are no (dis-)equalities in Δ , i.e., ψ_0 consists of relation atoms only. Each atom α of ψ_0 , such that $\text{fv}(\alpha) \cap \text{fv}_{J_{i_h}}(B_{i_h}(z_{i_h,1}, \dots, z_{i_h,\#B_{i_h}})) \neq \emptyset$, for some $h \in [1..k]$, is added to ψ'_0 . Moreover, each atom α of ψ_0 , such that $\text{fv}(\alpha) \cap \text{fv}_{\bar{J}_{i_h}}(B_{i_h}(z_{i_h,1}, \dots, z_{i_h,\#B_{i_h}})) = \emptyset$, for all $h \in [1..k]$ is added to ψ''_0 . Note that, each atom can only be added either to ψ'_0 or ψ''_0 but not to both, because $S'_{i_h} \sqsubseteq^{mc} S_{i_h}$ and $S'_{i_h} \models^s \phi'_{i_h}$ imply that no further element can be added to S'_{i_h} , for all $h \in [1..k]$. The rest of the atoms α from ψ_0 , i.e., such that $\text{fv}(\alpha) \cap \text{fv}_{J_{i_h}}(B_{i_h}(z_{i_h,1}, \dots, z_{i_h,\#B_{i_h}})) = \emptyset$ and $\text{fv}(\alpha) \cap \text{fv}_{\bar{J}_{i_h}}(B_{i_h}(z_{i_h,1}, \dots, z_{i_h,\#B_{i_h}})) = \emptyset$, for all $h \in [1..k]$, are split between ψ'_0 and ψ''_0 , by repeating the following steps until a fixpoint is reached:

▷ if $\text{fv}(\alpha) \cap \text{fv}(\psi'_0) \neq \emptyset$, then update ψ'_0 as $\psi'_0 * \alpha$,

▷ else, update ψ''_0 as $\psi''_0 * \alpha$.

By construction, we obtain that $\text{fv}(\psi'_0) \cap \text{fv}(\psi''_0) = \emptyset$, as required at point (1).

Let $J_0 \stackrel{\text{def}}{=} \{j \in [1..\#B_0] \mid \mathfrak{s}(x_j) \in \text{supp}(\sigma')\}$. Note that $J_0 \neq \emptyset$ because $\text{supp}(\sigma') \cap \{\mathfrak{s}(x_1), \dots, \mathfrak{s}(x_{\#B_0})\} \neq \emptyset$. We define the equivalence relation $\xi_0 \subseteq J_0 \times J_0$ as follows:

$$(i, j) \in \xi_0 \stackrel{\text{def}}{\iff} (x_i, x_j) \in \Xi$$

Moreover, one can show that $S'_0 \models^s \psi'_0$, by the construction of ψ'_0 and ψ''_0 , hence ξ_0 satisfies the conditions (4)–(7) from the definition of Γ , hence Γ contains a rule of the form (3.2), with qpf formula ψ'_0 . Since $S' \models^s \psi'_0 * \bigstar_{h=1}^k \phi'_{i_h} \theta_h$ and $\phi' = \psi'_0 * \bigstar_{h=1}^k \phi'_{i_h} \theta_h$ modulo a reordering of atoms, we obtain that $S' \models^s \phi'$. \square

The proof is completed as follows. Let $S' \in \text{split}(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$ be a maximally connected substructure of a canonical model $S \in \llbracket \mathbf{A} \rrbracket_{\Delta}^c$. Then, there exists a complete Δ -unfolding $A \Rightarrow_{\Delta}^* \exists y_1 \dots \exists y_m \cdot \bar{\phi}$ and a store \mathfrak{s} , injective over y_1, \dots, y_m , such that $S \models^s \bar{\phi}$. Because S' is connected, there exists a unique (i) predicate atom $B_0(z_1, \dots, z_{\#B_0})$, (ii) subformula ϕ of $\bar{\phi}$ and (iii) structure $S'' \sqsubseteq S$, such that $B_0(z_1, \dots, z_{\#B_0}) \Rightarrow_{\Delta}^* \exists y_{i_1} \dots \exists y_{i_n} \cdot \phi$ is a complete unfolding, $\text{supp}(S'') \cap \{\mathfrak{s}(z_1), \dots, \mathfrak{s}(z_{\#B_0})\} = \emptyset$, $S'' \models^s \phi$ and $S' \sqsubseteq^{mc} S''$. Without losing generality, we assume that the above is the smallest Δ -unfolding with these properties

and assume that the first rule of the Δ -unfolding is of the form (3.1), with a qpf formula ψ . Then, because S' is connected, the right-hand side of this rule contains zero or more predicate atoms $B_{i_h}(z_{i_h,1}, \dots, z_{i_h, \#B_{i_h}})$, for $i_h \in \{i_1, \dots, i_k\} \subseteq [1..\ell]$, such that $\text{supp}(\sigma') \cap \{\mathfrak{s}(z_{i_h,1}), \dots, \mathfrak{s}(z_{i_h, \#B_{i_h}})\} \neq \emptyset$. Accordingly, we decompose $S'' = S'_0 \bullet \dots \bullet S'_k$ such that $S'_0 \models^s \psi_0$ and $S'_h \models^s \phi_h$, where $B_{i_h}(z_{i_h,1}, \dots, z_{i_h, \#B_{i_h}}) \Rightarrow_{\Delta}^* \exists y_{j_h,1} \dots \exists y_{j_h, m_h} \cdot \phi_h$ are complete unfoldings, for all $h \in [1..k]$. This decomposition of S'' induces a decomposition of $S' \sqsubseteq^{mc} S''$ as $S'_0 \bullet \dots \bullet S'_k = S'$ such that $S'_h \sqsubseteq^{mc} S''_h$ and $\text{supp}(\sigma'_h) \cap \{\mathfrak{s}(z_{i_h,1}), \dots, \mathfrak{s}(z_{i_h, \#B_{i_h}})\} \neq \emptyset$, for all $h \in [1..k]$. Applying Fact B.1, we find nonempty subsets $J_h \subseteq [1..\#B_{i_h}]$, equivalence relations $\xi_h \subseteq J_h \times J_h$ and complete Γ -unfoldings $B_{i_h}(z_{i_h,1}, \dots, z_{i_h, \#B_{i_h}}) / \xi_h \Rightarrow_{\Gamma}^* \exists y_{p_h,1} \dots \exists y_{p_h, n_h} \cdot \phi'_h$, such that $S'_h \models^s \phi'_h$, for all $h \in [1..k]$. We define the sets $\bar{J}_h = [1..\#B_{i_h}] \setminus J_h$ and the formulæ ψ'_0 and ψ''_0 such that conditions (1) and (2) are met. Let $\Xi \stackrel{\text{def}}{=} (\text{conn}(\psi'_0) \cup \bigcup_{h=1}^k \xi_{i_h}(B_{i_h}(z_{i_h,1}, \dots, z_{i_h, \#B_{i_h}})))^{\bar{=}}$ be an equivalence relation. Since S' is connected, we argue that Ξ has a single equivalence class. Moreover, $(z, z) \notin \Xi$, for all $z \in \{z_1, \dots, z_{\#B_0}\}$, since $\text{supp}(S'') \cap \{\mathfrak{s}(z_1), \dots, \mathfrak{s}(z_{\#B_0})\} = \emptyset$ and \mathfrak{s} is injective. Then, by definition, Γ contains a rule of the form (3.3). This rule and the complete unfoldings $B_{i_h}(z_{i_h,1}, \dots, z_{i_h, \#B_{i_h}}) \Rightarrow_{\Gamma}^* \exists y_{j_h,1} \dots \exists y_{j_h, m_h} \cdot \phi'_h$, for all $h \in [1..k]$, are composed to make up a complete Γ -unfolding $P \Rightarrow_{\Gamma} \exists y_{q_1} \dots \exists y_{q_r} \cdot \psi'_0 * \bigstar_{h=1}^k \phi'_h$, such that $S' \models^s \psi'_0 * \bigstar_{h=1}^k \phi'_h$. Since $y_{q_1}, \dots, y_{q_r} \in \{y_1, \dots, y_m\}$ and \mathfrak{s} is injective over y_1, \dots, y_m , we obtain that $S' \in \llbracket P \rrbracket_{\Gamma}^c$.

“ $\llbracket \Gamma \rrbracket_P^c \subseteq \text{split}(\llbracket \Delta \rrbracket_A^c)$ ” We prove first two related facts. First, let $B_0(x_1, \dots, x_{\#B_0}) / \xi_0 \Rightarrow_{\Gamma}^* \exists y_1 \dots \exists y_n \cdot \phi$ be a complete Γ -unfolding, where ϕ is a qpf formula, $J_0 \subseteq [1..\#B_0]$ a nonempty set and $\xi_0 \subseteq J_0 \times J_0$ an equivalence relation, \mathfrak{s} be a store injective over $\{x_1, \dots, x_{\#B_0}\} \cup \{y_1, \dots, y_n\}$ and $S = (U, \sigma)$ be a structure such that $S \models^s \phi$. Given an equivalence class $I \subseteq J_0$ of ξ_0 , we define the structure:

$$\text{reach}_{\xi}^s(I) \stackrel{\text{def}}{=} (U, \lambda r \cdot \{ \langle u_1, \dots, u_{\#r} \rangle \in \sigma(r) \mid \forall j \in [1..\#r] \exists i \in I \cdot \mathfrak{s}(x_i) \text{ connected to } u_j \text{ in } S \})$$

Fact B.2. $\text{reach}_{\xi}^s(I) \sqsubseteq^{mc} S$.

Proof. By induction on the length of the Γ -unfolding. Assume that the first rule of the unfolding is of the form (3.2), with a qpf formula ψ_0 . Then, there exist nonempty sets $J_{i_h} \subseteq [1..\#B_{i_h}]$ and equivalence relations $\xi_{i_h} \subseteq J_{i_h} \times J_{i_h}$, for some $i_1, \dots, i_k \in [1..\ell]$ and all $h \in [1..k]$ and an equivalence relation $\Xi \subseteq (\{x_1, \dots, x_{\#B_0}\} \cup \{y_1, \dots, y_m\}) \times (\{x_1, \dots, x_{\#B_0}\} \cup \{y_1, \dots, y_m\})$, that satisfy points (1)–(7) from the definition of Γ . By point (7), $\{x_i \mid i \in I\}$ is an equivalence class of $\Xi \upharpoonright_{x_1, \dots, x_{\#B_0}}$ and let $X \subseteq \text{fv}(\psi'_0) \cup \bigcup_{h=1}^k \{z_{i_h,1}, \dots, z_{i_h, \#B_{i_h}}\}$ be the unique equivalence class of Ξ that contains it. For each $h \in [1..k]$, let $I_h \subseteq [1..\#B_{i_h}]$ be the equivalence class of ξ_{i_h} used to define X (3).

Let $B_{i_h}(x_1, \dots, x_{\#B_{i_h}}) \Rightarrow_{\Gamma}^* \exists y_{j_h,1} \dots \exists y_{j_h, m_h} \cdot \phi_h$ be complete Γ -unfoldings, such that $\phi = \psi'_0 * \bigstar_{h=1}^k \phi_h \theta_h$, where $\theta_h \stackrel{\text{def}}{=} [x_1 / z_{i_h,1}, \dots, x_{\#B_{i_h}} / z_{i_h, \#B_{i_h}}]$, for each $h \in [1..k]$. Since $S \models^s \phi$, there exist structures $S_0 = (U_0, \sigma_0), \dots, S_k = (U_k, \sigma_k)$ such that $S_0 \models^s \psi'_0$ and $S_h \models^s \phi_h \theta_h$, or equivalently, $S_h \models^{s \circ \theta_h^{-1}} \phi_h$ for all $h \in [1..k]$. By the inductive hypothesis, we have $\text{reach}_{S_h}^{s \circ \theta_h^{-1}}(I_h) \sqsubseteq^{mc} S_h$, for all $h \in [1..k]$. Since X is an equivalence class of Ξ , by point (3) of the definition of Γ , we obtain that $\text{reach}_{\xi}^s(I) \sqsubseteq^{mc} S$. \square

Second, let $B_0(x_1, \dots, x_{\#B_0}) \Rightarrow_{\Delta}^* \exists y_1 \dots \exists y_m \cdot \phi'$ be the complete Δ -unfolding obtained by replacing each rule of the form (3.2) with its corresponding rule (3.1) in the above

Γ -unfolding, such that ϕ' is a qpf formula and $y_1, \dots, y_n \in \{y_1, \dots, y_m\}$. Note that the latter can be assumed w.l.o.g., if necessary, by a renaming of the quantified variables.

Fact B.3. There exists a store \mathfrak{s}' , that is injective over y_1, \dots, y_m and agrees with \mathfrak{s} over y_1, \dots, y_n , and a structure $S' = (\mathbf{U}, \sigma')$, such that $S' \models^{\mathfrak{s}'} \phi'$ and $\text{reach}_{\xi}^{\mathfrak{s}}(I) \sqsubseteq^{mc} S'$, for each equivalence class $I \subseteq J_0$ of ξ_0 .

Proof. The store \mathfrak{s}' is defined as:

- ▷ $\mathfrak{s}'(y_i) = \mathfrak{s}(y_i)$, for each $i \in [1..n]$,
- ▷ $\mathfrak{s}'(y_i)$ is chosen from $\mathbf{U} \setminus \{\mathfrak{s}(y_1), \dots, \mathfrak{s}(y_n)\}$ such that, moreover, $\mathfrak{s}'(y_i) \neq \mathfrak{s}'(y_j)$, for all $n+1 \leq i < j \leq m$.

Note that \mathfrak{s}' can be build, because \mathbf{U} is infinite. Because of the assumption that each predicate defined by a rule from Δ occurs on some complete Δ -unfolding of \mathbf{A} , there exists a complete Δ -unfolding:

$$\mathbf{A} \Rightarrow_{\Delta} \dots \Rightarrow_{\Delta} \mathbf{B}_0(z_1, \dots, z_{\#\mathbf{B}_0}) * \phi'' \Rightarrow_{\Delta} (\exists y_1 \dots \exists y_m . \phi')[x_1/z_1, \dots, x_{\#\mathbf{B}_0}/z_{\#\mathbf{B}_0}] * \phi''$$

where ϕ'' is a predicate-free formula, possibly containing existential quantifiers. Since every complete Δ -unfolding of \mathbf{A} yields a satisfiable formula, there exists a store \mathfrak{s}'' that agrees with $\mathfrak{s}' \circ [z_1/x_1, \dots, z_{\#\mathbf{B}_0}/x_{\#\mathbf{B}_0}]$ over y_1, \dots, y_m and a structure S'' , such that:

$$S'' \models^{\mathfrak{s}''} (\exists y_1 \dots \exists y_m . \phi')[x_1/z_1, \dots, x_{\#\mathbf{B}_0}/z_{\#\mathbf{B}_0}] * \phi''$$

Note that, in the above construction, we have taken \mathfrak{s}'' to agree with \mathfrak{s}' over y_1, \dots, y_m . This is possible because there are no (dis-)equalities in Δ and the set of models of a qpf formula is closed under isomorphism-preserving renaming of elements.

Let S' and S''' be structures such that $S'' = S' \bullet S'''$, $S' \models^{\mathfrak{s}''} \phi'[x_1/z_1, \dots, x_{\#\mathbf{B}_0}/z_{\#\mathbf{B}_0}]$, or equivalently $S' \models^{\mathfrak{s}'} \phi'$, and $S''' \models^{\mathfrak{s}''} \phi''$. By induction on the length of the Δ -unfolding, relying on by point (1) of the definition of Γ , one can prove that $\phi' = \phi * \bar{\phi}$, where $\bar{\phi}$ is a qpf formula, such that $\text{fv}(\phi) \cap \text{fv}(\bar{\phi}) = \emptyset$. Since $S' \models^{\mathfrak{s}'} \phi'$, there exists a structure $\bar{S} = (\mathbf{U}, \bar{\sigma})$, such that $S' = S \bullet \bar{S}$. Moreover, since \mathfrak{s}' is injective over y_1, \dots, y_m , by construction, we obtain $\text{supp}(\sigma) \cap \text{supp}(\bar{\sigma}) = \emptyset$. Let $I \subseteq J_0$ be an equivalence class of ξ_0 . By Fact B.2, we have $\text{reach}_{\xi}^{\mathfrak{s}}(I) \sqsubseteq^{mc} S$. Since $S' = S \bullet \bar{S}$ and $\text{supp}(\sigma) \cap \text{supp}(\bar{\sigma}) = \emptyset$, we obtain $\text{reach}_{\xi}^{\mathfrak{s}}(I) \sqsubseteq^{mc} S'$. \square

The proof is completed as follows. Let $S \in [\mathbf{P}]_{\Gamma}^c$ be a canonical Γ -model of \mathbf{P} , i.e., there exists a complete Γ -unfolding $\mathbf{P} \Rightarrow_{\Gamma}^* \exists y_1 \dots \exists y_n . \phi$, where ϕ is a qpf formula, and a store \mathfrak{s} injective over y_1, \dots, y_n such that $S \models^{\mathfrak{s}} \phi$. By the definition of Γ , the first rule of this unfolding is of the form (3.3), with a qpf formula ψ_0 . Then there exist Γ -unfoldings $\mathbf{B}_i(x_1, \dots, x_{\#\mathbf{B}_i})/\xi_i \Rightarrow_{\Gamma}^* \exists y_{j_{i,1}} \dots \exists y_{j_{i,m_i}} . \phi_i$, for some sets $J_i \subseteq [1..\#\mathbf{B}_i]$ and equivalence relations $\xi_i \subseteq J_i \times J_i$, for $i \in [1..\ell]$, such that:

$$\phi = \psi_0 * \bigstar_{i \in [1..\ell], J_i \neq \emptyset} \phi_i \theta_i$$

where $\theta_i \stackrel{\text{def}}{=} [x_1/z_{i,1}, \dots, x_{\#\mathbf{B}_i}/z_{i,\#\mathbf{B}_i}]$, $i \in [1..\ell]$. Let $\{i_1, \dots, i_k\} \stackrel{\text{def}}{=} \{i \in [1..\ell] \mid J_i \neq \emptyset\}$. Then, there exist structures $S_0 \bullet \dots \bullet S_k = S$, such that $S_0 \models^{\mathfrak{s}} \psi_0$ and $S_j \models^{\mathfrak{s}} \phi_{i_j} \theta_{i_j}$, for $i \in [1..\ell]$. By the definition of Γ , there exists a complete Δ -unfolding:

$$\begin{aligned} \mathbf{B}_0(x_1, \dots, x_{\#\mathbf{B}_0}) &\Rightarrow_{\Delta} \exists y_1 \dots \exists y_m . \psi'_0 * \bigstar_{i=1}^{\ell} \mathbf{B}_i(z_{i,1}, \dots, z_{i,\#\mathbf{B}_i}) \Rightarrow_{\Delta} \dots \\ &\Rightarrow_{\Delta}^* \exists y_1 \dots \exists y_p . \psi'_0 * \bigstar_{j=1}^k \phi'_{i_j} \theta_{i_j} * \eta \end{aligned}$$

for qpf formulæ $\psi'_0, \phi'_{i_1}, \dots, \phi'_{i_k}$ and predicate-free formula η . Consider the equivalence relation Ξ over $\{x_1, \dots, x_{\#\mathbf{B}_0}\} \cup \{y_1, \dots, y_m\}$ defined as:

$$\Xi \stackrel{\text{def}}{=} \left(\text{conn}(\psi_0) \cup \bigcup_{j=1}^k \xi_{i_j}(\mathbf{B}_{i_j}(z_{i_j,1}, \dots, z_{i_j, \#\mathbf{B}_{i_j}})) \right) =$$

By point (8), Ξ has a single equivalence class X such that:

- ▷ $X \cap \{x_1, \dots, x_{\#\mathbf{B}_0}\} = \emptyset$,
- ▷ the sets $I_j \stackrel{\text{def}}{=} \{h \in [1.. \#\mathbf{B}_{i_j}] \mid z_{i_j, h} \in X\}$ are unions of equivalence classes of ξ_{i_j} , namely $I_j = I_{j,1} \uplus \dots \uplus I_{j,q_j}$, where $I_{j,h}$ are equivalence classes of ξ_{i_j} , for all $j \in [1..k]$.

By Fact B.3, there exist a store \mathfrak{s}' , that is injective over y_1, \dots, y_p and agrees with \mathfrak{s} over y_1, \dots, y_n , and structures S'_1, \dots, S'_k , such that $S'_j \models^{\mathfrak{s}'} \phi_j$ and $\text{reach}_{S'_j}^{\mathfrak{s}'}(I_{j,h}) \sqsubseteq^{mc} S'_j$, for all $j \in [1..k]$ and $h \in [1..q_j]$. We argue that $S_j = \bullet_{h=1}^{q_j} \text{reach}_{S'_j}^{\mathfrak{s}'}(I_{j,h}) \sqsubseteq^{mc} S'_j$. Moreover, since \mathfrak{s}' is injective over y_1, \dots, y_n , one can build a structure S'_0 , such that $S'_0 \models^{\mathfrak{s}'} \psi'_0$. We define $S' \stackrel{\text{def}}{=} S'_0 \bullet \bullet_{j=1}^k S'_j$. Thus, we have $S \sqsubseteq^{mc} S'$ and we are left with showing that S' can be embedded in a canonical Δ -model of \mathbf{A} .

By the assumption that each predicate defined by Δ occurs on some complete Δ -unfolding of \mathbf{A} , there exists another complete Δ -unfolding:

$$\begin{aligned} \mathbf{A} &\Rightarrow_{\Delta} \dots \Rightarrow_{\Delta} \mathbf{B}_0(z_{0,1}, \dots, z_{0, \#\mathbf{B}_0}) * \zeta \\ &\Rightarrow_{\Delta} (\exists y_1 \dots \exists y_m \cdot \psi_0 * \star_{i=1}^{\ell} \mathbf{B}_i(z_{i,1}, \dots, z_{i, \#\mathbf{B}_i})) [x_1/z_{0,1}, \dots, x_{\#\mathbf{B}_0}/z_{0, \#\mathbf{B}_0}] * \zeta \\ &\Rightarrow_{\Delta} \dots \Rightarrow_{\Delta} (\exists y_1 \dots \exists y_n \cdot \psi'_0 * \star_{i=1}^k \phi'_i \theta_i * \eta) [x_1/z_{0,1}, \dots, x_{\#\mathbf{B}_0}/z_{0, \#\mathbf{B}_0}] * \zeta \end{aligned}$$

for some predicate-free formula $\zeta \stackrel{\text{def}}{=} \exists y_{n+1} \dots \exists y_p \cdot \eta$, for some variables y_{n+1}, \dots, y_p , such that $\{y_{n+1}, \dots, y_p\} \cap \{y_1, \dots, y_n\} = \emptyset$ and a qpf formula η . Since this latter Δ -unfolding yields a satisfiable formula, there exists a structure S'' and a store \mathfrak{s}'' , injective over y_1, \dots, y_p , that agrees with \mathfrak{s}' over y_1, \dots, y_n , such that $S'' \models^{\mathfrak{s}''} \eta$. Then, $S' \bullet S'' \in \llbracket \mathbf{A} \rrbracket_{\Delta}^c$ and, since $S \sqsubseteq^{mc} S'$, we obtain $S \sqsubseteq^{mc} S' \bullet S''$, leading to $S \in \text{split}(\llbracket \mathbf{A} \rrbracket_{\Delta}^c)$. \square

B.3. Proof of Lemma 3.32. Without loss of generality, we can consider that Γ is equality-free (Lemma 2.6) and all-satisfiable for \mathbf{P} (Lemma 2.9).

" $(\llbracket \mathbf{P} \rrbracket_{\Gamma}^c)^{\sharp k} \subseteq \pi_3(\langle \mathbf{P} \rangle_{\Gamma}^{\sharp k})$ " We prove the following, more general, property:

Let $\mathbf{B}_0(x_1, \dots, x_{\#\mathbf{B}_0}) \Rightarrow_{\Gamma}^ \exists y_1 \dots \exists y_n \cdot \phi$ be a complete Γ -unfolding such that ϕ is a qpf formula, \mathfrak{s} be a store injective over $\{x_1, \dots, x_{\#\mathbf{B}_0}\} \cup \{y_1, \dots, y_n\}$, $\mathbf{S} = (\mathbf{U}, \sigma)$ be a structure such that $\mathbf{S} \models^{\mathfrak{s}} \phi$ and $D \subseteq \text{supp}(\sigma) \setminus \{\mathfrak{s}(x_1), \dots, \mathfrak{s}(x_{\#\mathbf{B}_0})\}$ be a set such that $\text{card}(D) \leq k$. Then there exists $\langle \{x_1, \dots, x_{\#\mathbf{B}_0}\}, c, M \rangle \in \langle \mathbf{B}_0 \rangle_{\Gamma}^{\sharp k}$ such that $\mathcal{C}_{\mathbf{S}}(\mathfrak{s}(x_i)) = c(x_i)$, for all $i \in [1.. \#\mathbf{B}_0]$ and $M = \{\{\mathcal{C}_{\mathbf{S}}(u) \mid u \in D\}\}$.*

The proof is by induction on the length of the complete Γ -unfolding. Assume w.l.o.g. that the first rule applied in the unfolding is of the form (3.1), with a qpf formula ψ_0 . Then, there exist structures $\mathbf{S}_0 = (\mathbf{U}_0, \sigma_0), \dots, \mathbf{S}_{\ell} = (\mathbf{U}_{\ell}, \sigma_{\ell})$, such that:

- ▷ $\mathbf{S} = \mathbf{S}_0 \bullet \dots \bullet \mathbf{S}_{\ell}$,
- ▷ $\mathbf{S}_0 \models^{\mathfrak{s}} \psi_0$,
- ▷ there exists a complete Γ -unfolding $\mathbf{B}_i(z_{i,1}, \dots, z_{i, \#\mathbf{B}_i}) \Rightarrow_{\Gamma}^* \exists y_{j_{i,1}} \dots \exists y_{j_{i,k_i}} \cdot \phi_i$, where $j_{i,1}, \dots, j_{i,k_i} \in [1..n]$ and ϕ_i is a subformula of ϕ , such that $\mathbf{S}_i \models^{\mathfrak{s}} \phi_i$, and the indices $j_{i,m}$ are pairwise distinct, for all $m \in [1..k_i]$, for all $i \in [1..\ell]$.

Let \mathfrak{s}_i be the store such that $\mathfrak{s}_i(x_j) = \mathfrak{s}(z_{i,j})$ for all $j \in [1..\#B_i]$ and \mathfrak{s}_i agrees with \mathfrak{s} everywhere else, for all $i \in [1..\ell]$. Then, there exists a complete Γ -unfolding:

$$B_i(x_1, \dots, x_{\#B_i}) \Rightarrow_{\Gamma}^* (\exists y_{j_1,1} \dots \exists y_{j_i,k_i} \cdot \phi_i)[z_{i,1}/x_1, \dots, z_{i,\#B_i}/x_{\#B_i}] = \exists y_{j_1,1} \dots \exists y_{j_i,k_i} \cdot \psi_i$$

such that $S_i \models^{\mathfrak{s}_i} \psi_i$, where ψ_i is a qpf formula, for all $i \in [1..\ell]$. We define the sets:

$$\begin{aligned} D_0 &\stackrel{\text{def}}{=} D \cap (\text{supp}(\sigma_0) \cup \bigcup_{i=1}^{\ell} \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\}) \\ D_i &\stackrel{\text{def}}{=} (\text{supp}(\sigma_i) \cap D) \setminus \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\}, \text{ for each } i \in [1..\ell] \end{aligned}$$

and prove the following fact:

Fact B.4. $D = D_0 \uplus D_1 \uplus \dots \uplus D_{\ell}$

Proof. The sets D_0, \dots, D_{ℓ} are pairwise disjoint, since:

$$\begin{aligned} D_0 &\subseteq \text{supp}(\sigma_0) \cup \bigcup_{i=1}^{\ell} \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\} \\ D_i &\subseteq \text{supp}(\sigma_i) \setminus \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\}, \text{ for all } i \in [1..\ell] \end{aligned}$$

and, moreover for all $1 \leq i < j \leq \ell$:

$$\begin{aligned} \triangleright & (\text{supp}(\sigma_0) \cup \bigcup_{i=1}^{\ell} \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\}) \cap \text{supp}(\sigma_i) \subseteq \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\} \\ \triangleright & \text{supp}(\sigma_i) \cap \text{supp}(\sigma_j) \subseteq \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\} \cap \{\mathfrak{s}(z_{j,1}), \dots, \mathfrak{s}(z_{j,\#B_j})\} \end{aligned}$$

because \mathfrak{s} is injective over y_1, \dots, y_n . “ \supseteq ” We have $D_0 \uplus D_1 \uplus \dots \uplus D_{\ell} \subseteq D$ because $D_i \subseteq D$, for all $i \in [0..\ell]$. “ \subseteq ” Let $u \in D$ be an element. By the choice of D , we have $u \in \text{supp}(\sigma) \setminus \{\mathfrak{s}(x_1), \dots, \mathfrak{s}(x_{\#B})\}$. Since $S = S_0 \bullet \dots \bullet S_{\ell}$, we have $\text{supp}(\sigma) = \bigcup_{i=0}^{\ell} \text{supp}(\sigma_i)$, hence $u \in \text{supp}(\sigma_i)$, for some $i \in [0..\ell]$. If $u \in D_0$ we are done. Otherwise, $u \notin D_0$, hence $u \notin \text{supp}(\sigma_0)$ and $u \in \text{supp}(\sigma_i)$, for some $i \in [1..\ell]$. Moreover, $u \notin \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\}$, for all $j \in [1..\ell]$, hence $u \in D_i$. \square

Back to the proof, since $\text{card}(D_i) \leq \text{card}(D) \leq k$, for all $i \in [1..\ell]$, by the inductive hypothesis, there exist $\langle \{x_1, \dots, x_{\#B_i}\}, c_i, M_i \rangle \in \langle B_i \rangle_{\Gamma}^{\#k}$, for $i \in [1..\ell]$, such that:

$$\begin{aligned} \triangleright & \mathcal{C}_S(\mathfrak{s}_i(x_j)) = c(x_j), \text{ for all } j \in [1..\#B_i], \\ \triangleright & M_i = \{\{\mathcal{C}_{S_i}(u) \mid u \in D_i\}\}. \end{aligned}$$

Let $\langle \text{fv}(\psi_0), c_0, \emptyset \rangle \stackrel{\text{def}}{=} \gamma(\psi_0)$ be a color triple. Since $S_0 \models^{\mathfrak{s}} \psi_0$, we have $\mathcal{C}_{S_0}(\mathfrak{s}(x_j)) = c(x_j)$, for all $j \in [1..\#B_0]$. By definition, there exists a constraint of the form (3.4) for the above rule (3.1). We prove that the $\bullet^{\#k}$ -composition from the right-hand side of the constraint is defined. Suppose, for a contradiction, that $c_i(x) \cap c_j(x) \neq \emptyset$, for some $x \in \{z_{i,1}, \dots, z_{i,\#B_i}\} \cap \{z_{j,1}, \dots, z_{j,\#B_j}\}$ and $1 \leq i < j \leq \ell$. Then $\mathcal{C}_{S_i}(x) \cap \mathcal{C}_{S_j}(x) \neq \emptyset$, contradicting the fact that $S_i \bullet S_j$ is defined. The same reasoning applies if $c_0(x) \cap c_i(x) \neq \emptyset$, for some $x \in \text{fv}(\psi) \cap \{z_{i,1}, \dots, z_{i,\#B_i}\}$ and $i \in [1..\ell]$. Then, there exists a color triple:

$$\langle X', c', M' \rangle \in \langle \text{fv}(\psi_0), c_0, \emptyset \rangle \bullet^{\#k} \bullet^{\#k}_{i \in [1..\ell]} \langle \{x_1, \dots, x_{\#B_i}\}, c_i, M_i \rangle [x_1/z_{i,1}, \dots, x_{\#B_i}/z_{i,\#B_i}]$$

W.l.o.g. we can chose the tuple such that $M' = M_1 \cup \dots \cup M_{\ell}$. This choice is possible since $\text{card}(M') = \sum_{i=1}^{\ell} \text{card}(M_i) = \sum_{i=1}^{\ell} \text{card}(D_i) \leq \text{card}(D) \leq k$. Let $\langle X, c, M \rangle \in \langle X', c', M' \rangle \downarrow_{\{x_1, \dots, x_{\#B_0}\}}^{\#k}$ be such that $M = M' \cup \{\{\mathcal{C}_{S_0}(u) \mid u \in D_0\}\}$. This choice is possible, since $\text{card}(M) \leq \sum_{i=0}^{\ell} \text{card}(D_i) \leq k$. We prove the points of the statement:

\triangleright Let $i \in [1..\#B_0]$ be an index. By the definition of the $\bullet^{\#k}$ -composition, we have:

$$\mathcal{C}_S(\mathfrak{s}(x_i)) = \mathcal{C}_{S_0 \bullet \dots \bullet S_{\ell}}(\mathfrak{s}(x_i)) = \bigcup_{j=0}^{\ell} \mathcal{C}_{S_j}(\mathfrak{s}(x_i)) = \biguplus_{j=0}^{\ell} c_j(x_i) = c'(x_i) = c(x_i)$$

$$\triangleright M = \bigcup_{i=0}^{\ell} \{\{\mathcal{C}_{S_i}(u) \mid u \in D_i\}\} = \{\{\mathcal{C}_{S_0 \bullet \dots \bullet S_{\ell}}(u) \mid u \in \biguplus_{i=0}^{\ell} D_i\}\} = \{\{\mathcal{C}_S(u) \mid u \in D\}\}$$

" $\pi_3(\langle \mathbf{P} \rangle_\Gamma^{\#k}) \subseteq (\llbracket \mathbf{P} \rrbracket_\Gamma)_{\#k}$ " We prove the following, more general, property:

Let $\langle \{x_1, \dots, x_{\#B_0}\}, c, M \rangle \in \langle \mathbf{B}_0 \rangle_\Gamma^{\#k}$ be a color triple. Then there exists a complete Γ -unfolding $\mathbf{B}_0(x_1, \dots, x_{\#B_0}) \Rightarrow_\Gamma^* \exists y_1 \dots \exists y_n \cdot \phi$, whose steps belong to a complete Γ -unfolding of \mathbf{P} , such that ϕ is a qpf formula, a store \mathfrak{s} injective over $\{x_1, \dots, x_{\#B_0}\} \cup \{y_1, \dots, y_n\}$, a structure $\mathbf{S} = (\mathbf{U}, \sigma)$ such that $\mathbf{S} \models^{\mathfrak{s}} \phi$ and $D \subseteq \text{supp}(\sigma) \setminus \{\mathfrak{s}(x_1), \dots, \mathfrak{s}(x_{\#B_0})\}$, $\text{card}(D) \leq k$, such that $\mathcal{C}_S(\mathfrak{s}(x_i)) = c(x_i)$, for all $i \in [1..\#B_0]$ and $M = \{\{\mathcal{C}_S(u) \mid u \in D\}\}$.

The proof is by induction on the length of the finite fixpoint iteration that produced $\langle \{x_1, \dots, x_{\#B_0}\}, c, M \rangle$. Assume that the last step of the iteration corresponds to a constraint of the form (3.4), with a qpf formula ψ_0 . By definition, there exists a rule of the form (3.1) in Γ , with the same qpf formula ψ_0 . Then ψ_0 is satisfiable, because each Γ -unfolding of \mathbf{P} yields a satisfiable formula. Then there exists a color triple:

$$\langle X', c', M' \rangle \in \gamma(\psi_0) \bullet^{\#k} \bullet^{\#k}_{i \in [1..\ell]} \langle \mathbf{B}_i \rangle_\Gamma^{\#k} [x_1/z_{i,1}, \dots, x_{\#B_i}/z_{i,\#B_i}]$$

such that

$$\begin{aligned} \langle \{x_1, \dots, x_{\#B_0}\}, c, M \rangle &\in \langle X', c', M' \rangle \downarrow_{\{x_1, \dots, x_{\#B_0}\}}^{\#k} \\ M &\subseteq M' \cup \{c'(x) \mid x \in X' \setminus \{x_1, \dots, x_{\#B_0}\}\} \end{aligned}$$

Then, there exist

$$\begin{aligned} \langle \text{fv}(\psi_0), c_0, \emptyset \rangle &\stackrel{\text{def}}{=} \gamma(\psi_0) \\ \langle \{z_{i,1}, \dots, z_{i,\#B_i}\}, c'_i, M_i \rangle &\in \langle \mathbf{B}_i \rangle_\Gamma^{\#k} [x_1/z_{i,1}, \dots, x_{\#B_i}/z_{i,\#B_i}], \text{ for all } i \in [1..\ell] \end{aligned}$$

such that

$$\langle X', c', M' \rangle \in \langle \text{fv}(\psi_0), c_0, \emptyset \rangle \bullet^{\#k} \bullet^{\#k}_{i \in [1..\ell]} \langle \{z_{i,1}, \dots, z_{i,\#B_i}\}, c'_i, M_i \rangle$$

Hence, there exist $\langle \{x_1, \dots, x_{\#B_i}\}, c_i, M_i \rangle \in \langle \mathbf{B}_i \rangle_\Gamma^{\#k}$ such that $c'_i = c_i \circ [x_1/z_{i,1}, \dots, x_{\#B_i}/z_{i,\#B_i}]$, for all $i \in [1..\ell]$. By the inductive hypothesis, for all $i \in [1..\ell]$, there exist:

- ▷ a complete unfolding $\mathbf{B}_i(x_1, \dots, x_{\#B_i}) \Rightarrow_\Gamma^* \exists y_{j_i,1} \dots \exists y_{j_i,k_i} \cdot \psi_i$ such that ψ_i is a qpf formula. By applying an α -renaming, if necessary, we assume w.l.o.g. that the variables $y_{j_i,1} \dots y_{j_i,k_i}$ are pairwise distinct and, moreover, distinct from $x_1, \dots, x_{\#B_0}$.
- ▷ a store \mathfrak{s}_i that is injective over $\{x_1, \dots, x_{\#B_i}\} \cup \{y_{j_i,1}, \dots, y_{j_i,k_i}\}$. We assume w.l.o.g. that $\mathfrak{s}_i(x_j) = \mathfrak{s}_i(x_m)$ iff $z_{i,j}$ and $z_{i,m}$ are the same variable in the rule (3.1), for all $1 \leq i < k \leq \ell$, $j \in [1..\#B_i]$ and $m \in [1..\#B_k]$. Note that this assumption does not contradict the fact that \mathfrak{s}_i is injective over $\{x_1, \dots, x_{\#B_i}\} \cup \{y_{j_i,1}, \dots, y_{j_i,k_i}\}$.
- ▷ a structure $\mathbf{S}_i = (\mathbf{U}_i, \sigma_i)$ such that $\mathbf{S}_i \models^{\mathfrak{s}_i} \psi_i$. We assume w.l.o.g. that $\text{supp}(\sigma_i) \cap \text{supp}(\sigma_j) \subseteq \{\mathfrak{s}_i(x_1), \dots, \mathfrak{s}_i(x_{\#B_i})\} \cap \{\mathfrak{s}_j(x_1), \dots, \mathfrak{s}_j(x_{\#B_j})\}$. Note that this is possible by the assumption that Γ is equality-free.
- ▷ a set $D_i \subseteq \text{supp}(\sigma_i) \setminus \{\mathfrak{s}_i(x_1), \dots, \mathfrak{s}_i(x_{\#B_i})\}$, such that $\text{card}(D_i) \leq k$, $\mathcal{C}_{\mathbf{S}_i}(\mathfrak{s}_i(x_j)) = c_i(x_j)$, for all $j \in [1..\#B_i]$ and $M_i = \{\{\mathcal{C}_{\mathbf{S}_i}(u) \mid u \in D_i\}\}$.

We prove the points of the statement. Let θ_i be the substitution $[x_1/z_{i,1}, \dots, x_{\#B_i}/z_{i,\#B_i}]$, for each $i \in [1..\ell]$, where $\mathbf{B}_i(z_{i,1}, \dots, z_{i,\#B_i})$ is a predicate atom that occurs on the right-hand side of the rule (3.1). A complete Γ -unfolding $\mathbf{B}_0(x_1, \dots, x_{\#B_0}) \Rightarrow_\Gamma^* \exists y_1 \dots \exists y_n \cdot \psi$ is built from the rule (3.1) above, with qpf formula ψ_0 , followed by $\mathbf{B}_i(x_1, \dots, x_{\#B_i})\theta_i \Rightarrow_\Gamma^* \exists y_{j_i,1} \dots \exists y_{j_i,k_i} \cdot \psi_i\theta_i$, for all $i \in [1..\ell]$. Hence $\psi = \psi_0 * \bigstar_{i=1}^\ell \psi_i\theta_i$ modulo a reordering of atoms. Let $\mathfrak{s}'_i \stackrel{\text{def}}{=} \mathfrak{s}_i \circ \theta_i$ and define the store \mathfrak{s} as follows:

- ▷ $\mathfrak{s}(z) \stackrel{\text{def}}{=} \mathfrak{s}'_i(z)$, for each $z \in \text{fv}(\psi_i\theta_i)$,

▷ $\mathfrak{s}(z) \notin \bigcup_{i=1}^{\ell} \mathfrak{s}'_i(\text{fv}(\psi_i\theta_i))$, for each variable $z \in \text{fv}(\psi_0) \setminus \bigcup_{i=1}^{\ell} \text{fv}(\psi_i\theta_i)$ such that, moreover, \mathfrak{s} is injective over $\text{fv}(\psi_0)$. Note that this is possible because we assumed Γ to be equality-free.

Then, we consider a structure $S_0 = (\mathbf{U}_0, \sigma_0)$ such that:

▷ $S_0 \models^{\mathfrak{s}} \psi_0$, and

▷ $\text{supp}(\sigma_0) \cap \text{supp}(\sigma_i) \subseteq \mathfrak{s}(\text{fv}(\psi_0)) \cap \mathfrak{s}(\text{fv}(\psi_i\theta_i))$, for all $i \in [1..\ell]$.

Since ψ_0 is satisfiable, such a structure exists and we can consider w.l.o.g. that it satisfies the above conditions, because Γ is equality-free. It is easy to check that the structures S_0, \dots, S_{ℓ} are pairwise locally disjoint, hence $S = (\mathbf{U}, \sigma) \stackrel{\text{def}}{=} S_0 \bullet \dots \bullet S_{\ell}$ is defined. Moreover, we have $S \models^{\mathfrak{s}} \psi$, because $\psi = \psi_0 * \bigstar_{i=1}^{\ell} \psi_i\theta_i$, $S_0 \models^{\mathfrak{s}} \psi_0$ and $S_i \models^{\mathfrak{s}} \psi_i\theta_i$, for all $i \in [1..\ell]$. Further, for all $j \in [1..\#B_0]$, we have:

$$\mathcal{C}_S(\mathfrak{s}(x_j)) = \mathcal{C}_{S_0 \bullet \dots \bullet S_{\ell}}(\mathfrak{s}(x_i)) = \bigsqcup_{i=0}^{\ell} c_i(x_j) = c'(x_j) = c(x_j)$$

We consider the set $D \stackrel{\text{def}}{=} \{u \in \text{supp}(\sigma) \mid \mathcal{C}_S(u) \in M\}$. Suppose, for a contradiction, that $\mathfrak{s}(x_i) \in D$, for some $i \in [1..\#B_0]$. Then $\mathcal{C}_S(\mathfrak{s}(x_i)) \in M$, hence $c(x_i) \in M$. Since $M \subseteq M' \cup \{c'(x) \mid x \in X' \setminus \{x_1, \dots, x_{\#B_0}\}\}$, we must have $c(x_i) \in M' \subseteq \bigcup_{j=1}^{\ell} M_j$ and let $j \in [1..\ell]$ be such that $c(x_i) \in M_j = \{\mathcal{C}_{S_j} \mid u \in D_j\}$, by the inductive hypothesis. Then there exists $k \in [1..\#B_j]$ such that $c(x_i) = c_j(x_k) = \mathcal{C}_{S_j}(\mathfrak{s}_j(x_k))$, thus $\mathfrak{s}_j(x_k) \in D_j \subseteq \text{supp}(\sigma_j) \setminus \{\mathfrak{s}_j(x_1), \dots, \mathfrak{s}_j(x_{\#B_j})\}$, contradiction. We obtained $D \subseteq \text{supp}(\sigma) \setminus \{\mathfrak{s}(x_1), \dots, \mathfrak{s}(x_{\#B_0})\}$ and are left with proving that $M = \{\mathcal{C}_S(u) \mid u \in D\}$. “ \supseteq ” Immediate, by the definition of D . “ \subseteq ” Let $C \in M$ be a color. Then either one of the following holds:

▷ $C = \{r \in \mathbb{R} \mid r(z, \dots, z) \text{ occurs in } \psi_0\} \cup \bigcup_{i=1}^{\ell} c_i(z)$, for some $z \in \text{fv}(\psi_0)$: in this case, $C = \mathcal{C}_S(\mathfrak{s}(z))$ and $\mathfrak{s}(z) \in \text{supp}(\sigma)$, hence $\mathfrak{s}(z) \in D$.

▷ $C \in M_i$, for some $i \in [1..\ell]$: in this case, $C = \mathcal{C}_S(u) = \mathcal{C}_{S_i}(u)$, for some $u \in D_i$, by the inductive hypothesis. Then $u \in \text{supp}(\sigma_i) \subseteq \text{supp}(\sigma)$, hence $u \in D$. \square

APPENDIX C. PROOFS FROM SECTION 4

C.1. Proof of Lemma 4.7. Assume w.l.o.g. that \mathcal{A} is rooted and let $\mathfrak{G}_{\mathcal{A}} = (\mathcal{N}, \mathcal{E}, S_0)$ be the SCC graph of \mathcal{A} . By Definition 4.5, $\mathfrak{G}_{\mathcal{A}}$ is a tree and, moreover, $S_0 = \{\iota\}$, because \mathcal{A} is rooted. Let $\Lambda : \mathcal{N} \cup \delta \rightarrow \{1, \infty\}$ be the labeling from Definition 4.5. For every SCC $S \in \mathcal{N} \setminus \{S_0\}$, let $\text{entry}(S)$ be the unique state q such that $\{q\} = \tau^{\bullet} \cap S$, where $\{\tau\} = \bullet S$, by point (1) of Definition 4.5, and $\text{entry}(S_0) \stackrel{\text{def}}{=} \iota$. Moreover, each linear SCC $S \in \mathcal{N}$ such that $\Lambda(S) = 1$ has a unique transition τ , such that $S^{\bullet} = \{\tau\}$, by point (2a) of Definition 4.5. We prove first an invariant of 1-labeled linear SCCs:

Fact C.1. Let $p \in \text{dom}(\theta)$ be a position, such that $\theta(p) \in S$, for a linear SCC $S \in \mathcal{N}$, such that $\Lambda(S) = 1$. Then there exists a descendant $p' \in \text{dom}(\theta)$ of p , such that $\theta(p') = s_0$, $t(p') = \beta$ and $\theta(p'_i) = s_i$, for all $i \in [1..k]$, where $S^{\bullet} = \{s_0 \xrightarrow{\beta} (s_1, \dots, s_k)\}$.

Proof. Suppose, for a contradiction, that $s_0 \xrightarrow{\beta} (s_1, \dots, s_k)$ never occurs below p in θ . Then every transition that occurs at some position below p in θ must be from $\bullet S^{\bullet}$. This, however, cannot be the case for a transition $\theta(p') \xrightarrow{t(p')} ()$, such that $p' \in \text{fr}(\theta)$. Since, moreover, θ is an accepting run, such a transition must occur on the frontier of θ . \square

The following facts prove the existence and uniqueness of a position labeled with the entry state of each 1-labeled linear SCC:

Fact C.2. For each SCC $S \in \mathcal{N}$, such that $\Lambda(S) = 1$, there exists a position $p \in \text{dom}(\theta)$, such that $\theta(p) = \text{entry}(S)$.

Proof. Because $\mathfrak{G}_{\mathcal{A}}$ is a tree with root S_0 , we have that S is reachable from S_0 in $\mathfrak{G}_{\mathcal{A}}$ by a path of pairs from \mathcal{E} . The proof goes by induction on the length $n \geq 0$ of this path. For the base case $n = 0$ (i.e., $S = S_0$) we take $p = \epsilon$. For the inductive step, let S' be the parent of S in $\mathfrak{G}_{\mathcal{A}}$. By points (1) and (2c) of Definition 4.5, $\bullet S = \{\tau\}$ for some $\tau \in S'^{\bullet} \cap \delta^1$, such that $\{\text{entry}(S)\} = \tau^{\bullet} \cap S$. By the inductive hypothesis, there exists a position $p' \in \text{dom}(\theta)$, such that $\theta(p') = \text{entry}(S')$. By Fact C.1, there exists a descendant p of p' , such that $\theta(p) = \text{entry}(S)$. \square

Fact C.3. For each SCC $S \in \mathcal{N}$, such that $\Lambda(S) = 1$, there exists at most one position $p \in \text{dom}(\theta)$, such that $\theta(p) = \text{entry}(S)$.

Proof. Suppose, for a contradiction, that there exist two positions $p_1, p_2 \in \text{dom}(\theta)$, such that $\theta(p_1) = \theta(p_2) = \text{entry}(S)$. By induction of the length of p_i , we prove the existence of a sequence $S_{i,k_i}, \tau_{i,k_i}, \dots, S_{i,1}, \tau_{i,1}, S_{i,0} = S_0$ such that $\theta(p_i) \in S_{i,k_i}$, $\bullet S_{i,j} = \tau_{i,j}$ and $\{\tau_{i,j}\} = S_{i,j-1}^{\bullet}$, for all $j \in [1..k_i]$ and $i = 1, 2$. Since $p_1 \neq p_2$, there exists an SCC $S_{1,j_1} = S_{2,j_2}$ that violates condition (1) of Definition 4.5. \square

Let $\tau : q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell)$ be a transition, such that $\Lambda(\tau) = 1$. By point (2b) of Definition 4.5, we have $\tau \in S^{\bullet}$ for some linear SCC $S \in \mathcal{N}$, such that $\Lambda(S) = 1$. By Fact C.2 and Fact C.3, there exists a unique position $p \in \text{dom}(\theta)$, such that $\theta(p) = \text{entry}(S)$. By Fact C.1, there exists a position $p' \in \text{dom}(\theta)$, such that $\theta(p') = q_0$, $t(p') = \alpha$ and $\theta(p'i) = q_i$, for all $i \in [1..\ell]$. Suppose, for a contradiction, that this position is not unique, hence there exists another position $p'' \in \text{dom}(\theta)$, such that $\theta(p'') = q_0$, $t(p'') = \alpha$ and $\theta(p''i) = q_i$, for all $i \in [1..\ell]$. Since $\theta(p') = \theta(p'') = q_0 \in S$, there exists a transition τ' with $\text{card}(\tau'^{\bullet} \cap S) \geq 2$, in contradiction with the fact that S is linear. This concludes the proof \square .

C.2. Proof of Lemma 4.8. Let us consider the SCC graph $\mathfrak{G}_{\mathcal{A}} = (\mathcal{N}, \mathcal{E}, S_0)$ and the mapping $\Lambda : \mathcal{N} \cup \delta \rightarrow \{1, \infty\}$ with the properties stated in Definition 4.5 and let $q \in \bullet(\delta^\infty)$ be a state. W.l.o.g., we consider that $\mathcal{L}(\mathcal{A}) \neq \emptyset$ and that \mathcal{A} is trim. Then q is reachable from $S_0 = \{\iota\}$, i.e., there exists a partial run θ_1 on \mathcal{A} and a position p_1 such that $\theta_1(\epsilon) = \iota$ and $\theta_1(p_1) = q$. Let p_2 be the longest strict prefix of p_1 such that the transition $\tau : \theta_1(p_2) \xrightarrow{a} \langle \theta_1(p_21), \dots, \theta_1(p_2\ell) \rangle$ is in δ^1 for some $a \in \mathbb{A}$ and index ℓ . This position p_2 exists, by Definition 4.5, because $S_0 = \{\iota\}$ is linear, $\Lambda(S_0) = 1$ by condition (2c), $\text{card}(S_0^{\bullet}) = 1$ by condition (2a), and the only transition $\tau_0 \in \{\iota\}^{\bullet}$ is in δ^1 by condition (2b). This shows that $\iota \notin \bullet\delta^\infty$ hence $q \neq \iota$ and τ_0 is a transition in δ^1 on the path from ι to q in θ_1 , with τ being the last one.

We decompose $p_1 = p_2 r p_3$ for some index $r \in [1..\ell]$ and position p_3 and define the partial run θ_2 as $\theta_2(\epsilon) \stackrel{\text{def}}{=} \theta_1(p_2 r)$ and, for each $u \in \mathbb{N}^*$ and $i \in \mathbb{N}$ such that $p_2 r u i \in \text{dom}(\theta_1)$ and $p_2 r u$ is a strict prefix of p_1 , by $\theta_2(ui) \stackrel{\text{def}}{=} \theta_1(p_2 r u)$. Then θ_2 starts from the state $q_0 \stackrel{\text{def}}{=} \theta_1(p_2 r) \in \tau^{\bullet}$ and $p_3 \in \text{fr}(\theta_2)$ gives the state $\theta_2(p_3) = q$. Let $S \in \mathcal{N}$ be the SCC in $\mathfrak{G}_{\mathcal{A}}$ such that $q_0 \in S$. Then $\Lambda(\tau) = 1$ and $\tau \in \bullet S$ (hence $\bullet S = \{\tau\}$ by condition (1) of Definition 4.5), thus $\Lambda(S) = 1$ by condition (2c) of Definition 4.5.

We distinguish three cases (see Figure 15 for an illustration):

- ▷ If S is not linear, there exists a transition $\tau' \in \bullet S^{\bullet}$ such that $\text{card}(\tau'^{\bullet} \cap S) \geq 2$. Let $q' \stackrel{\text{def}}{=} \tau'^{\bullet}$ and q'', q''' be the states such that $\{\{q'', q'''\}\} \subseteq \tau'^{\bullet} \cap S$. Since $q_0, q', q'', q''' \in S$, we can

construct a partial run $\theta_3 \in \mathcal{R}_{q_0}^\infty(\mathcal{A})$ with transitions taken from $\bullet S^\bullet$, which reaches q' from q_0 , then applies τ' and reaches q_0 from both q'' and q''' . This gives $\theta_3(p_4) = \theta_3(p_5) = q_0$, for two distinct positions $p_4, p_5 \in \text{fr}(\theta_3)$. We define $\theta_0 \in \mathcal{R}_{q_0}^\infty(\mathcal{A})$ as the partial run with domain $\text{dom}(\theta_3) \cup \{p_5u \mid u \in \text{dom}(\theta_2)\}$, that extends θ_3 by $\theta_0(p_5u) \stackrel{\text{def}}{=} \theta_2(u)$, for all $u \in \text{dom}(\theta_2)$. Then θ_0 satisfies point 1 of the lemma because $\theta_0(p_4) = q_0$ and $\theta_0(p_5p_3) = q$, with $p_4 \neq p_5p_3 \in \text{fr}(\theta_0)$.

- ▷ If S is linear and $q \notin S$, there exists a unique position p_6 and a transition:

$$\tau' : \theta_2(p_6) \xrightarrow{\alpha} \langle \theta_2(p_61), \dots, \theta_2(p_6k) \rangle \in \delta^\infty$$

for some alphabet symbol $\alpha \in \mathbb{A}$ and some index $k \in \mathbb{N}$, such that $\theta_2(p_6) \in S$. Moreover, there exists an index $r \in [1..k]$ such that $\theta_2(p_6r) \notin S$ and p_6r is a prefix of p_3 . Then $\Lambda(\tau') = \infty$ and, by condition (2b) of Definition 4.5, we have $\tau' \notin \bullet S^\bullet$, hence $\tau' \in \bullet S^\bullet$ and $q' \stackrel{\text{def}}{=} \theta_2(p_6r') \in S$ for another index $r' \in [1..k] \setminus \{r\}$. Then there exists a partial run $\theta_4 \in \mathcal{R}_{q'}^\infty(\mathcal{A})$ such that $\theta_4(p_7) = q_0$ for some position $p_7 \in \text{fr}(\theta_4)$. We define the partial run $\theta_0 \in \mathcal{R}_{q_0}^\infty(\mathcal{A})$ with domain $\text{dom}(\theta_2) \cup \{p_6r'u \mid u \in \text{dom}(\theta_4)\}$, by extending θ_2 with $\theta_0(p_6r'u) \stackrel{\text{def}}{=} \theta_4(u)$, for all $u \in \text{dom}(\theta_4)$. Then θ_0 satisfies point 1 of the lemma because $\theta_0(p_6r'p_7) = q_0$ and $\theta_0(p_3) = q$, with $p_6r'p_7 \neq p_3 \in \text{fr}(\theta_0)$.

- ▷ If S is linear and $q \in S$, let $\theta \in \mathcal{R}_q^\infty(\mathcal{A})$ be a partial run. Then $\bullet S^\bullet$ contains only one transition in δ^1 , thus for every position $u \in \text{dom}(\theta) \setminus \text{fr}(\theta)$, such that $\theta(u) \in S$, the transition $\theta(u) \xrightarrow{\alpha} \langle \theta(u1), \dots, \theta(uk) \rangle$ belongs to $\bullet S^\bullet$. Then, there exists an index $i \in [1..k]$ such that $\theta(ui) \in S$, and we can find a path in θ which stays in S and reaches the frontier, that is $q' \stackrel{\text{def}}{=} \theta(p_8) \in S$, for some $p_8 \in \text{fr}(\theta)$. Hence, there exists a partial run $\theta_5 \in \mathcal{R}_{q'}^\infty(\mathcal{A})$ such that $\theta_5(p_9) = q_0$, for some position $p_9 \in \text{fr}(\theta_5)$. We now can extend θ to some partial run $\theta' \in \mathcal{R}_q^\infty(\mathcal{A})$ with domain $\text{dom}(\theta) \cup \{p_8u \mid u \in \text{dom}(\theta_5)\}$, as $\theta'(p_8u) \stackrel{\text{def}}{=} \theta_5(u)$ for all $u \in \text{dom}(\theta_5)$. The partial run θ' satisfies point 2 of the lemma, because $\theta'(p_8p_9) = q_0$, with $p_8p_9 \in \text{fr}(\theta')$. \square




Figure 15: The cases from the proof of Lemma 4.8

C.3. Proof of Lemma 4.9. We assume w.l.o.g. that $\mathcal{A} = (\mathbb{A}, \mathcal{Q}, \iota, \delta)$ is rooted. Let $\mathfrak{G}_{\mathcal{A}} = (\mathcal{N}, \mathcal{E})$ be the SCC graph of \mathcal{A} , where $\mathcal{N} = \{S_1, \dots, S_M\}$ is a topological ordering of the SCCs i.e., if $(S_i, S_j) \in \mathcal{E}$ then $i < j$, for all $i, j \in [1..M]$. For each $i = 1, \dots, M$, we iterate the following transformation of \mathcal{A} :

- ▷ let $\mathcal{S}_i \stackrel{\text{def}}{=} \bigcup_{(S_i, S_j) \in \mathcal{E}^*} S_j$ be the set of states from any SCC reachable from S_i in $\mathfrak{G}_{\mathcal{A}}$,
- ▷ let $k_i \stackrel{\text{def}}{=} \sum_{\tau \in \bullet S_i} \text{card}(\tau^\bullet \cap S_i)$ be the number of edges of $\mathfrak{G}_{\mathcal{A}}$ incoming to S_i ,
- ▷ create k_i copies of the transitions $q_0 \xrightarrow{a} (q_1, \dots, q_\ell) \in \delta$ such that $\{q_0, q_{i_1}, \dots, q_{i_j}\} = \{q_0, \dots, q_\ell\} \cap \mathcal{S}_i$ i.e., add a transition $(q_0, h) \xrightarrow{a} (q_1, \dots, (q_{i_1}, h), \dots, (q_{i_j}, h), \dots, q_\ell)$ for each $h \in [1..k_i]$,
- ▷ connect these new transitions to the rest of the automaton by adequately changing the states $q \in \tau^\bullet \cap \mathcal{S}_i$ for $\tau \in \bullet S_i$ to their corresponding copies (q, h) , for all $h \in [1..k_i]$.

It is easy to check that the resulting automaton fulfills condition (1) of Definition 4.5 and has the same language as \mathcal{A} , using Lemma 4.4. We can thus assume w.l.o.g. in the following that $\mathfrak{G}_{\mathcal{A}} = (\mathcal{N}, \mathcal{E}, S_1)$ is a tree and let S_1, \dots, S_N be a topological ordering of its nodes. We associate a variable x_i (resp. y_τ) ranging over $\{0, 1, \infty\}$ with each SCC $S_i \in \mathcal{N}$ (resp. transition $\tau \in \delta$). Initially, the values of these variables are all zero. We iterate over the finite sequence S_1, \dots, S_N as follows. For each $i \in [1..N]$, we perform the following assignments in this order:

- (i) let $x_i \stackrel{\text{def}}{=} \begin{cases} 1, & \text{if } i = 1 \\ \sum_{\tau \in \bullet S_i} y_\tau \cdot \text{card}(\tau^\bullet \cap S_i), & \text{otherwise} \end{cases}$
- (ii) for each $\tau \in \bullet S_i^\bullet$, let $y_\tau \stackrel{\text{def}}{=} \begin{cases} \infty, & \text{if } x_i > 0 \\ 0, & \text{otherwise} \end{cases}$
- (iii) if $x_i \in \{0, \infty\}$ or S_i is nonlinear, for each $\tau \in S_i^\bullet$, let $y_\tau \stackrel{\text{def}}{=} \begin{cases} \infty, & \text{if } x_i > 0 \\ 0, & \text{otherwise} \end{cases}$
- (iv) else (i.e., $x_i = 1$ and S_i is linear) chose for all $\{y_\tau\}_{\tau \in S_i^\bullet}$ some values from $\{0, 1\}$, such that $x_i = \sum_{\tau \in S_i^\bullet} y_\tau$.

Since, for each SCC $S_i \in \mathcal{N}$, there is at most one transition $\tau \in \bullet S_i$ and $\text{card}(\tau^\bullet \cap S_i) \leq 1$, each variable x_i is assigned either 0, 1 or ∞ at (i). Note that no variable is assigned twice in the above iteration sequence, because every x_i is assigned exactly once, every y_τ , for $\tau \in \bullet S_i$ is assigned before x_i and every y_τ , for $\tau \in \bullet S_i^\bullet \cup S_i^\bullet$ is assigned after x_i . Furthermore, we have $\bullet S_i^\bullet \cap \bullet S_j^\bullet = \emptyset$ and $S_i^\bullet \cap S_j^\bullet = \emptyset$, for all $1 \leq i < j \leq N$, so that each y_τ , for $\tau \in \bullet S_i^\bullet \cup S_i^\bullet$, is assigned exactly once. Moreover, since the choice at (iv) is finite, there are finitely many outcomes of the above nondeterministic iteration, say $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_\ell, \mathbf{y}_\ell)$, where $\mathbf{x}_i = \langle \bar{x}_{i,j} \rangle_{j \in [1..N]}$ and $\mathbf{y}_i = \langle \bar{y}_{i,\tau} \rangle_{\tau \in \delta}$. For each $i \in [1..\ell]$, we define the automaton $\mathcal{A}_i = (\mathbb{A}, \mathcal{Q}_i, \iota, \delta_i)$, where $\mathcal{Q}_i \stackrel{\text{def}}{=} \bigcup \{S_j \mid \bar{x}_{i,j} > 0, j \in [1..N]\}$ and $\delta_i \stackrel{\text{def}}{=} \{\tau \in \delta \mid \bar{y}_{i,\tau} > 0\}$. We are left with proving the following facts:

Fact C.4. Each automaton \mathcal{A}_i is choice-free, for $i \in [1..N]$.

Proof. We prove below the points of Definition 4.5:

- (1) Let S_{j_0} be an SCC of \mathcal{A} , such that $\bar{x}_{i,j_0} > 0$ i.e., S_{j_0} is a vertex in the SCC graph $\mathfrak{G}_{\mathcal{A}_i}$. Since the variable x_{j_0} received its value \bar{x}_{i,j_0} at (i), either $j_0 = 1$ (in which case $\bar{x}_{i,j_0} = 1$) or there exists an incoming transition $\tau \in \bullet S_{j_0}$ such that $\bar{y}_{i,\tau} > 0$. Let S_{j_1} , for some $j_1 < j_0$ be the SCC such that $\bullet \tau \in S_{j_1}$. Then $\bar{x}_{i,j_1} > 0$. Repeating the same argument for j_1 , we discover a maximal finite sequence j_0, \dots, j_k such that $(S_{j_{i+1}}, S_{j_i}) \in \mathcal{E}$, for all $i \in [0..k-1]$. Moreover, it must be the case that $j_k = 1$, or else the sequence could be extended, contradicting its

maximality. Since $\mathfrak{G}_{\mathcal{A}}$ is a tree, the path from S_1 to S_{j_0} must be unique and, since the choice of S_{j_0} was arbitrary, $\mathfrak{G}_{\mathcal{A}_i} = (\mathcal{N}_i, \mathcal{E}_i, S_1)$ is a tree as well. The second point from condition (1) holds already for $\mathfrak{G}_{\mathcal{A}}$, hence it must hold for $\mathfrak{G}_{\mathcal{A}_i}$.

(2) The mapping $\Lambda_i : \mathcal{N}_i \cup \delta_i \rightarrow \{1, \infty\}$ is defined as $\Lambda_i(S_j) = \bar{x}_{i,j}$ for each $S_j \in \mathcal{N}_i$ and $\Lambda_i(\tau) = \bar{y}_{i,\tau}$ for each $\tau \in \delta_i$. We check that Λ_i verifies the conditions (2) from Definition 4.5:

(2a) if $S_j \in \mathcal{N}_i$ is linear and $\bar{x}_{i,j} = 1$ then the choice at step (iv) was $\bar{y}_{i,\tau} = 1$, for exactly one transition $\tau \in S_j^\bullet$.

(2b) $\bar{y}_{i,\tau} = 1$ iff the value of y_τ was set at step (iv) and $\tau \in S_j^\bullet$ is the unique outgoing transition for which a nonzero value was assigned to y_τ , for a linear SCC S_j with $\bar{x}_{i,j} = 1$.

(2c) $\bar{x}_{i,j} = 1$ iff the value of x_j was set at step (i) and either $j = 1$ or for all but one transitions $\tau \in \bullet S_j$ we have $\bar{y}_{i,\tau} = 1$. \square

Fact C.5. $\mathcal{L}(\mathcal{A}) = \bigcup_{i=1}^{\ell} \mathcal{L}(\mathcal{A}_i)$.

Proof. “ \supseteq ” Since $\delta_i \subseteq \delta$, we have $\mathcal{L}(\mathcal{A}_i) \subseteq \mathcal{L}(\mathcal{A})$, for all $i \in [1..\ell]$.

“ \subseteq ” Let $t \in \mathcal{L}(\mathcal{A})$ and θ be an accepting run of \mathcal{A} over t . We show that there exists an iteration (i)–(iv) leading to the values $(\mathbf{x}_i, \mathbf{y}_i)$ such that, for each transition τ occurring on θ at some position $p \in \text{dom}(\theta)$ i.e., $\theta(p) = \bullet \tau$, we have:

- $\triangleright \bar{x}_{i,j} > 0$, where S_j is the unique SCC of \mathcal{A} such that $\bullet \tau \in S_j$, and
- $\triangleright \bar{y}_{i,\tau} > 0$.

By the second point above we obtain that θ is an accepting run of \mathcal{A}_i . The proof is by reverse induction on the size of the subtree of θ rooted at p .

Base case: If $p = \epsilon$, the variable x_1 is always assigned the value 1 at step (i). We chose the values for all $\{y_{\tau'}\}_{\tau' \in S_1^\bullet}$, such that $y_{\tau'}$ is assigned 1 and $y_{\tau'}$ is assigned 0, for all $\tau' \in S_1^\bullet \setminus \{\tau\}$ at step (iv).

Induction step: If $p \in \text{dom}(\theta) \setminus \{\epsilon\}$, since $j \neq 1$, by the inductive hypothesis, the variable $y_{\tau'}$ is assigned non-zero values, for at least one $\tau' \in \bullet S_j$, thus we assign x_j the value $\sum_{\tau' \in \bullet S_j} y_{\tau'} \cdot \text{card}(\tau^\bullet \cap S_j) > 0$ at step (i). If $\tau \in \bullet S_j^\bullet$, then y_τ is assigned ∞ at step (ii). Otherwise, it must be the case that $\tau \in S_j^\bullet$ and we distinguish two cases. If S_j is nonlinear, then y_τ is assigned ∞ at (iii). Else, S_j is linear and we can chose the value 1 for y_τ at step (iv), because x_j has been already assigned to 1. \square

Let $i \in [1..\ell]$. We prove the upper bound on $\text{card}(\delta_i^1)$, as follows. Since the SCC graph of \mathcal{A}_i is a tree, the number of 1-transitions in \mathcal{A}_i equals the number of SCCs in \mathcal{A}_i . Due to the expansion of the first step of the proof, we have $\text{card}(\delta_i^1) \leq \max\{\rho(a) \mid a \in \mathbb{A}\}^s \leq \max\{\rho(a) \mid a \in \mathbb{A}\}^{\text{card}(\mathcal{Q})}$, where $s \leq \text{card}(\mathcal{Q})$ denotes the number of SCCs in \mathcal{A} . \square

C.4. Proof of Lemma 4.13. (1) For the first part, we prove the two directions of the following equivalence: for all structures \mathbf{S} and predicates \mathbf{B} of arity n , there exists a store \mathfrak{s} such that $\mathbf{S} \models_{\Delta}^{\mathfrak{s}} \mathbf{B}(x_1, \dots, x_n)$ iff there exists a tree $t \in \mathcal{L}_{q_{\mathbf{B}}}(\mathcal{A}_{\Delta, \mathbf{A}})$ and a store $\bar{\mathfrak{s}}$ such that $\mathbf{S} \models^{\bar{\mathfrak{s}}} \Theta(t)$ and $\bar{\mathfrak{s}}(x_j^{[e]}) = \mathfrak{s}(x_j)$ for all $j \in [1..n]$.

“ \Rightarrow ” We proceed by induction on the definition of $\mathbf{S} \models_{\Delta}^{\mathfrak{s}} \mathbf{B}(x_1, \dots, x_n)$. Then Δ contains a rule ρ of the form $\mathbf{B}(x_1, \dots, x_n) \leftarrow \exists y_1 \dots \exists y_m . \psi * \bigstar_{i=1}^{\ell} \mathbf{B}_i(z_{i,1}, \dots, z_{i,n_i})$ and we can decompose the structure $\mathbf{S} = \mathbf{S}_0 \bullet \dots \bullet \mathbf{S}_{\ell}$, such that $\mathbf{S}_0 \models^{\mathfrak{s}'} \psi$ and $\mathbf{S}_i \models_{\Delta}^{\mathfrak{s}'_i} \mathbf{B}_i(z_{i,1}, \dots, z_{i,n_i})$ for all $i \in [1..\ell]$, for a store \mathfrak{s}' that agrees with \mathfrak{s} over $\{x_1, \dots, x_n\}$. For all $i \in [1..\ell]$, we consider a store \mathfrak{s}_i such that $\mathfrak{s}_i(x_j) = \mathfrak{s}'(z_{i,j})$, for all $j \in [1..n_i]$. We have $\mathbf{S}_i \models_{\Delta}^{\mathfrak{s}_i} \mathbf{B}_i(x_1, \dots, x_{n_i})$

and, by induction hypothesis, there exists a tree $t_i \in \mathcal{L}_{q_{B_i}}(\mathcal{A}_{\Delta, \mathbf{A}})$ and a store $\bar{\mathfrak{s}}_i$ such that $S_i \models^{\bar{\mathfrak{s}}_i} \Theta(t_i)$ and $\bar{\mathfrak{s}}_i(x_j^{[\epsilon]}) = \mathfrak{s}_i(x_j)$ for all $j \in [1..n_i]$. Let $\bar{\mathfrak{s}}$ be a store such that:

- ▷ $\bar{\mathfrak{s}}(x_j^{[\epsilon]}) = \mathfrak{s}(x_j)$, for all $j \in [1..n]$,
- ▷ $\bar{\mathfrak{s}}(y_j^{[\epsilon]}) = \mathfrak{s}'(y_j)$ for all $j \in [1..m]$,
- ▷ $\bar{\mathfrak{s}}(z^{[ip]}) \stackrel{\text{def}}{=} \bar{\mathfrak{s}}_i(z^{[ip]})$, for all $i \in [1..\ell]$ and $z^{[ip]} \in \text{fv}(\Theta(t_i))$.

Note that $\bar{\mathfrak{s}}$ is well defined because $\bar{\mathfrak{s}}_i(z^{[ip]}) = \bar{\mathfrak{s}}_j(z^{[ip]}) = \mathfrak{s}'(z^{[ip]})$, for all $z^{[ip]} \in \text{fv}(\Theta(t_i)) \cap \text{fv}(\Theta(t_j))$. We have $S_i \models^{\bar{\mathfrak{s}}} \Theta(t_i)^{[i]}$ for all $i \in [1..\ell]$ and $S_0 \models^{\bar{\mathfrak{s}}} \alpha_\rho$, thus $S \models^{\bar{\mathfrak{s}}} \Theta(t)$ where t is the tree consisting of a root labelled by α_ρ and ℓ children t_i for $i \in [1..\ell]$. Since $t \in \mathcal{L}_{q_B}(\mathcal{A}_{\Delta, \mathbf{A}})$ and $\bar{\mathfrak{s}}(x_j^{[\epsilon]}) = \mathfrak{s}(x_j)$, for all $j \in [1..n_0]$, by definition, we obtain the result.

“ \Leftarrow ” The reverse implication is proven by induction on the structure of the tree $t \in \mathcal{L}_{q_B}(\mathcal{A}_{\Delta, \mathbf{A}})$ such that $S \models^{\bar{\mathfrak{s}}} \Theta(t)$. Since $t \in \mathcal{L}_{q_B}(\mathcal{A}_{\Delta, \mathbf{A}})$, there is a transition $q_B \xrightarrow{t(\epsilon)} (q_{B_1}, \dots, q_{B_\ell}) \in \delta_\Delta$ such that $t(\epsilon) = \alpha_\rho$ for some rule ρ of the form above and $t|_i \in \mathcal{L}_{q_{B_i}}(\mathcal{A}_{\Delta, \mathbf{B}})$ for all $i \in [1..\ell]$.

Meanwhile $S \models^{\bar{\mathfrak{s}}} t(\epsilon) * \bigstar_{i=1}^\ell \Theta(t|_i)^{[i]}$, thus we can decompose the structure as $S = S_0 \bullet \dots \bullet S_\ell$, such that $S_0 \models^{\bar{\mathfrak{s}}} \alpha_\rho$ and $S_i \models^{\bar{\mathfrak{s}}} \Theta(t|_i)^{[i]} * \bigstar_{j=1}^{n_i} z_{i,j}^{[\epsilon]} = x_j^{[i]}$ for all $i \in [1..\ell]$. Note that the additional equalities from α_ρ are necessary to remember the links between the variables from ρ . Let $\bar{\mathfrak{s}}_i$ be a store, such that $\bar{\mathfrak{s}}_i(z^{[ip]}) = \bar{\mathfrak{s}}(z^{[ip]})$, for all $z^{[ip]} \in \text{fv}(\Theta(t|_i))$ and all $i \in [1..\ell]$. By the inductive hypothesis on $t|_i$, there exists a store \mathfrak{s}_i such that $S_i \models_{\Delta}^{\bar{\mathfrak{s}}_i} B_i(x_1, \dots, x_{n_i})$ and $\mathfrak{s}_i(x_j) = \bar{\mathfrak{s}}_i(x_j^{[\epsilon]})$, for all $j \in [1..n_i]$. We consider a store \mathfrak{s}' such that $\mathfrak{s}'(x_j) = \bar{\mathfrak{s}}(x_j^{[\epsilon]})$, for all $j \in [1..n]$, and $\mathfrak{s}'(y_j) \stackrel{\text{def}}{=} \bar{\mathfrak{s}}(y_j^{[\epsilon]})$, for all $j \in [1..m]$. For all $i \in [1..\ell]$ and $j \in [1..n_i]$ we have $\mathfrak{s}_i(x_j) = \bar{\mathfrak{s}}_i(x_j^{[\epsilon]}) = \bar{\mathfrak{s}}(x_j^{[i]}) = \bar{\mathfrak{s}}(z_{i,j}^{[\epsilon]}) = \mathfrak{s}'(z_{i,j})$, because $z_{i,j}^{[\epsilon]} = x_j^{[i]}$ holds for $\bar{\mathfrak{s}}$ in the empty structure. Therefore $S_i \models_{\Delta}^{\mathfrak{s}'} B_i(z_{i,1}, \dots, z_{i,n_i})$, for all $i \in [1..\ell]$. Moreover $S_0 \models^{\mathfrak{s}'} \psi$, and by composing the structures and using ρ , we obtain $S \models_{\Delta}^{\mathfrak{s}'} B(x_1, \dots, x_n)$.

(2) To show $\llbracket \mathcal{A} \rrbracket = \llbracket \mathbf{A}_i \rrbracket_{\Delta_{\mathcal{A}}}$ we prove the following equivalence: for all structures S and states $q_0 \in Q$, there exists a tree $t \in \mathcal{L}_{q_0}(\mathcal{A})$ and a store $\bar{\mathfrak{s}}$ such that $S \models^{\bar{\mathfrak{s}}} \Theta(t)$ iff there exists a store \mathfrak{s} such that $S \models_{\Delta_{\mathcal{A}}}^{\mathfrak{s}} A_{q_0}(x_1, \dots, x_{\#q_0})$ and $\mathfrak{s}(x_j) = \bar{\mathfrak{s}}(x_j^{[\epsilon]})$, for all $j \in [1..\#q_0]$.

“ \Rightarrow ” We reason by induction on the structure of the tree $t \in \mathcal{L}_{q_0}(\mathcal{A})$, such that $S \models^{\bar{\mathfrak{s}}} \Theta(t)$.

Since $t \in \mathcal{L}_{q_0}(\mathcal{A})$, there is a transition $q_0 \xrightarrow{t(\epsilon)} (q_1, \dots, q_\ell) \in \delta$ such that $t|_i \in \mathcal{L}_{q_i}(\mathcal{A})$, for all $i \in [1..\ell]$. Meanwhile $S \models^{\bar{\mathfrak{s}}} t(\epsilon) * \bigstar_{i=1}^\ell \Theta(t|_i)^{[i]}$ thus, we can decompose the structure $S = S_0 \bullet \dots \bullet S_\ell$, such that $S_0 \models^{\bar{\mathfrak{s}}} t(\epsilon)$ and $S_i \models^{\bar{\mathfrak{s}}} \Theta(t|_i)^{[i]}$. Let $\bar{\mathfrak{s}}_i$ be a store, such that $\bar{\mathfrak{s}}_i(z^{[ip]}) = \bar{\mathfrak{s}}(z^{[ip]})$, for all $z^{[ip]} \in \text{fv}(\Theta(t|_i))$ and all $i \in [1..\ell]$. By the inductive hypothesis on $t|_i$, there exists a store \mathfrak{s}_i such that $S_i \models_{\Delta_{\mathcal{A}}}^{\bar{\mathfrak{s}}_i} A_{q_i}(x_1, \dots, x_{\#q_i})$ and $\mathfrak{s}_i(x_j) = \bar{\mathfrak{s}}_i(x_j^{[\epsilon]})$ for all $j \in [1..\#q_i]$. We consider a store \mathfrak{s}' , such that $\mathfrak{s}'(x_j) = \bar{\mathfrak{s}}(x_j^{[\epsilon]})$, for all $j \in [1..\#q_0]$, $\mathfrak{s}'(x_j^{[i]}) = \mathfrak{s}_i(x_j)$, for all $i \in [1..\ell]$ and all $j \in [1..\#q_i]$, and $\mathfrak{s}'(z) = \bar{\mathfrak{s}}(z)$, for all other variables $z \in \text{fv}(t(\epsilon))$. Then $S_i \models_{\Delta_{\mathcal{A}}}^{\mathfrak{s}'} A_{q_i}(x_1^{[i]}, \dots, x_{\#q_i}^{[i]})$ and $S_0 \models^{\mathfrak{s}'} t(\epsilon)[x_1^{[\epsilon]}/x_1, \dots, x_{\#q_0}^{[\epsilon]}/x_{\#q_0}]$ thus, by composing the structures, we obtain $S \models_{\Delta_{\mathcal{A}}}^{\mathfrak{s}'} A_{q_0}(x_1, \dots, x_{\#q_0})$.

“ \Leftarrow ” The reverse is shown by induction on the definition of $S \models_{\Delta_{\mathcal{A}}}^{\mathfrak{s}} A_{q_0}(x_1, \dots, x_{\#q_0})$. Then there exists a rule in $\Delta_{\mathcal{A}}$ of the form (4.1) and we can decompose the structure $S = S_0 \bullet \dots \bullet S_\ell$ such that $S_0 \models^{\mathfrak{s}'} \alpha[x_1^{[\epsilon]}/x_1, \dots, x_{\#q_0}^{[\epsilon]}/x_{\#q_0}]$ and $S_i \models_{\Delta_{\mathcal{A}}}^{\mathfrak{s}'} A_{q_i}(x_1^{[i]}, \dots, x_{\#q_i}^{[i]})$, for all $i \in [1..\ell]$, where \mathfrak{s}' is a store that agrees with \mathfrak{s} over $\{x_1, \dots, x_{\#q_i}\}$. For all $i \in [1..\ell]$, we consider a store \mathfrak{s}_i such that $\mathfrak{s}_i(x_j) = \mathfrak{s}'(x_j^{[i]})$. We have $S_i \models_{\Delta_{\mathcal{A}}}^{\mathfrak{s}_i} A_{q_i}(x_1, \dots, x_{\#q_0})$ and, by the inductive hypothesis, there exists a tree $t_i \in \mathcal{L}_{q_i}(\mathcal{A})$ and a store $\bar{\mathfrak{s}}_i$ such that $\sigma_i \models^{\bar{\mathfrak{s}}_i} \Theta(t_i)$ and $\bar{\mathfrak{s}}_i(x_j^{[\epsilon]}) = \mathfrak{s}_i(x_j)$, for all $j \in [1..\#q_i]$. Let $\bar{\mathfrak{s}}$ be a store such that:

- ▷ $\bar{\mathfrak{s}}(x_j^{[e]}) \stackrel{\text{def}}{=} \mathfrak{s}(x_j)$, for all $j \in [1..\#q_0]$,
- ▷ $\bar{\mathfrak{s}}(y_j^{[e]}) \stackrel{\text{def}}{=} \mathfrak{s}'(y_j^{[e]})$, for all $j \in [1..m]$,
- ▷ $\bar{\mathfrak{s}}(z^{[p]}) \stackrel{\text{def}}{=} \bar{\mathfrak{s}}_i(z^{[p]})$, for all $i \in [1..\ell]$ and all $z^{[p]} \in \text{fv}(\Theta(t_i))$.

Note that $\bar{\mathfrak{s}}$ is well defined because $\bar{\mathfrak{s}}_i(z^{[p]}) = \bar{\mathfrak{s}}_j(z^{[p]}) = \mathfrak{s}'(z^{[p]})$, for all $z^{[p]} \in \text{fv}(\Theta(t_i)) \cap \text{fv}(\Theta(t_j))$. We have $S_i \models^{\bar{\mathfrak{s}}} \Theta(t_i)^{[i]}$, for all $i \in [1..\ell]$ and $S_0 \models^{\bar{\mathfrak{s}}} \alpha$, thus $S \models^{\bar{\mathfrak{s}}} \Theta(t)$, where t is the tree consisting of a root labelled by α and children t_i , for $i \in [1..\ell]$. Since $t \in \mathcal{L}_{q_0}(\mathcal{A})$ and $\bar{\mathfrak{s}}(x_j^{[e]}) = \mathfrak{s}(x_j)$ for all $j \in [1..\#q_0]$ by definition, we obtain the result.

To show $\llbracket \mathcal{A} \rrbracket^r \subseteq \llbracket \mathcal{A}_\ell \rrbracket_{\Delta, \mathcal{A}}^r$, let $(S, \mathfrak{d}) \in \llbracket \mathcal{A} \rrbracket^c$, where $S = (U, \sigma)$ is a structure and $\mathfrak{d} \subseteq U \times U$ is a symmetric relation. Then there exists a tree $t \in \mathcal{L}(\mathcal{A})$ and a store \mathfrak{s} canonical for $\Theta(t)$, such that $S \models^{\mathfrak{s}} \Theta(t)$ and, for each $(u, v) \in \mathfrak{d}$, there exist variables $x \in \mathfrak{s}^{-1}(u)$ and $y \in \mathfrak{s}^{-1}(v)$ such that the disequality $x \neq y$ occurs in $\Theta(t)$. Let θ be the accepting run of \mathcal{A} over t . By a depth-first traversal of θ , we build a complete unfolding $A_\ell \Rightarrow_{\Delta, \mathcal{A}}^* \Theta(t)^{\exists}$. Since $S \models^{\mathfrak{s}} \Theta(t)$, we obtain $(S, \mathfrak{d}) \in \llbracket \mathcal{A}_\ell \rrbracket_{\Delta, \mathcal{A}}^r$, by Definition 3.1. Conversely, to show $\llbracket \mathcal{A} \rrbracket^r \supseteq \llbracket \mathcal{A}_\ell \rrbracket_{\Delta, \mathcal{A}}^r$, let $(S, \mathfrak{d}) \in \llbracket \mathcal{A}_\ell \rrbracket_{\Delta, \mathcal{A}}^r$, where $S = (U, \sigma)$ is a structure and $\mathfrak{d} \subseteq U \times U$ is a symmetric relation. Then, there exists a complete unfolding $A_\ell \Rightarrow_{\Delta, \mathcal{A}}^* \exists x_1 \dots \exists x_n . \psi$, where ψ is a qpf formula, and a store \mathfrak{s} canonical for ψ , such that $S \models^{\mathfrak{s}} \psi$ and, for all $(u, v) \in \mathfrak{d}$ there exist variables $x \in \mathfrak{s}^{-1}(u)$ and $y \in \mathfrak{s}^{-1}(v)$, such that the disequality $x \neq y$ occurs in ψ . By induction on the length of the unfolding, one can build an accepting run θ of \mathcal{A} , that recognizes a tree $t \in \mathcal{L}(\mathcal{A})$, such that $\Theta(t)$ differs from ψ by an α -renaming and permutation of atoms via commutativity and associativity of the separating conjunction. Hence $(S, \mathfrak{d}) \in \llbracket \Theta(t)^{\exists} \rrbracket^r$, thus $(S, \mathfrak{d}) \in \llbracket \mathcal{A} \rrbracket^r$. \square

C.5. Proof of Lemma 4.23. (1) Let $t' \in \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^I)$ be a tree. Since $\mathcal{A}_{\Delta, \mathcal{A}}^I$ was obtained from $\mathcal{A}_{\Delta, \mathcal{A}}$ by removing relation and disequality atoms from the labels of its 1-transitions, there exists a tree $t \in \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}})$, such that $\text{dom}(t) = \text{dom}(t')$ and $\Theta(t) = \Theta(t') * \bigstar_{i=1}^n r_i(z_{i,1}, \dots, z_{i,k_i}) * \bigstar_{j=1}^m y_{j,1} \neq y_{j,2}$ modulo reordering of atoms, for some relation symbols r_i and variables $z_{i,1}, \dots, z_{i,k_i}, y_{j,1}, y_{j,2}$. By Lemma 4.13, $\Theta(t)$ is satisfiable, hence there exists a structure (U, σ) and a store \mathfrak{s} , such that $(U, \sigma) \models^{\mathfrak{s}} \Theta(t)$. We define the interpretation $\sigma'(r) = \sigma(r) \setminus \{ \langle \mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,k_i}) \rangle \mid i \in [1..n], r_i = r \}$, for all $r \in \mathbb{R}$. It is easy to check that $(U, \sigma') \models^{\mathfrak{s}} \Theta(t')$, hence $\Theta(t')$ is satisfiable. Since the choice of t' was arbitrary, we obtain that $\mathcal{A}_{\Delta, \mathcal{A}}^I$ is all-satisfiable.

(2) Let $t' \in \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^I)$ be a tree. By the construction of $\mathcal{A}_{\Delta, \mathcal{A}}^I$ from $\mathcal{A}_{\Delta, \mathcal{A}}$, there exists a tree $t \in \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}})$, such that $\text{dom}(t) = \text{dom}(t')$ and $\Theta(t) = \Theta(t') * \bigstar_{i=1}^n r_i(z_{i,1}, \dots, z_{i,k_i}) * \bigstar_{j=1}^m y_{j,1} \neq y_{j,2}$ modulo reordering of atoms, for some relation symbols r_i and variables $z_{i,1}, \dots, z_{i,k_i}, y_{j,1}, y_{j,2}$. Let $Y \stackrel{\text{def}}{=} \bigcup_{j=1}^m \{y_{j,1}, y_{j,2}\}$ be the set of variables occurring in disequality atoms. By Lemma 4.7, each 1-transition of $\mathcal{A}_{\Delta, \mathcal{A}}$ occurs exactly once in each accepting run, hence $n \leq \text{card}(\delta_{\Delta}^1) \cdot \text{maxRelAtoms}(\Delta)$, $\text{card}(Y) \leq \text{card}(\delta_{\Delta}^1) \cdot \text{maxVars}(\Delta)$. By Lemma 2.10 (2, 3), $\text{tw}(\Theta(t')^{\exists}) \leq \text{tw}(\Theta(t)^{\exists}) + n + \text{card}(Y) \leq \text{tw}(\Theta(t)^{\exists}) + \text{card}(\delta_{\Delta}^1) \cdot (\text{maxRelAtoms}(\Delta) + \text{maxVars}(\Delta))$. Since the choice of t' was arbitrary, we obtain that $\llbracket \mathcal{A}_{\Delta, \mathcal{A}}^I \rrbracket$ is treewidth bounded, more precisely $\text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}}^I \rrbracket) \leq \text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}} \rrbracket) + \text{card}(\delta_{\Delta}^1) \cdot (\text{maxRelAtoms}(\Delta) + \text{maxVars}(\Delta))$.

(3) Let $t \in \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}})$ be a tree. By the construction of $\mathcal{A}_{\Delta, \mathcal{A}}^I$ from $\mathcal{A}_{\Delta, \mathcal{A}}$, there exists a tree $t' \in \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^I)$, such that $\text{dom}(t') = \text{dom}(t)$ and $\Theta(t) = \Theta(t') * \bigstar_{i=1}^n r_i(z_{i,1}, \dots, z_{i,k_i}) * \bigstar_{j=1}^m y_{j,1} \neq y_{j,2}$ modulo reordering of atoms, for some relation symbols r_i and variables $z_{i,1}, \dots, z_{i,k_i}, y_{j,1}, y_{j,2}$. Let $Y \stackrel{\text{def}}{=} \bigcup_{i=1}^n \{z_{i,1}, \dots, z_{i,k_i}\}$ be the set of variables occurring in relation atoms.

By Lemma 4.7, each 1-transition of $\mathcal{A}_{\Delta, \mathcal{A}}$ occurs exactly once in each accepting run, hence $\text{card}(Y) \leq \text{card}(\delta_{\Delta}^1) \cdot \text{maxVars}(\Delta)$. By Lemma 2.10 (4) $\text{tw}(\Theta(t)^{\exists}) \leq \text{tw}(\Theta(t')^{\exists}) + \text{card}(Y) \leq \text{tw}(\Theta(t')^{\exists}) + \text{card}(\delta_{\Delta}^1) \cdot \text{maxVars}(\Delta)$. Since the choice of t was arbitrary, we obtain that $\llbracket \mathcal{A}_{\Delta, \mathcal{A}} \rrbracket$ is treewidth bounded, more precisely $\text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}} \rrbracket) \leq \text{tw}(\llbracket \mathcal{A}_{\Delta, \mathcal{A}}^I \rrbracket) + \text{card}(\delta_{\Delta}^1) \cdot \text{maxVars}(\Delta)$. \square

C.6. Proof of Lemma 4.25. (1) Let $t' \in \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^{II})$ be a tree. Since $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ was obtained from $\mathcal{A}_{\Delta, \mathcal{A}}^I$ by removing equality atoms from the labels of its 1-transitions (2), there exists a tree $t \in \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^I)$, such that $\text{dom}(t) = \text{dom}(t')$ and $\Theta(t) = \Theta(t') * \psi$ where ψ is a conjunction of equality atoms. By Lemma 4.23 (1), $\Theta(t)$ is satisfiable, hence there exists a structure \mathbf{S} and a store \mathfrak{s} such that $\mathbf{S} \models^{\mathfrak{s}} \Theta(t)$. We immediately obtain $\mathbf{S} \models^{\mathfrak{s}} \Theta(t')$, hence $\Theta(t')$ is satisfiable. Since the choice of t' was arbitrary, we obtain that $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ is all-satisfiable.

(2) Let $t' \in \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^{II})$ be a tree and θ' be an accepting run over t' . We shall build a tree $t \in \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^I)$ related to t' and θ' and show that $\text{tw}(\llbracket \Theta(t)^{\exists} \rrbracket) \leq \text{tw}(\llbracket \Theta(t')^{\exists} \rrbracket) + K$, where K is constant that does not depend on the choice of t' . The idea of the construction of t from t' is to add resets before and after each 1-transition in the run θ' , so that the equalities removed by the transformation from δ_{Δ}^I to δ_{Δ}^{II} can be added back, without changing the set of models of $\Theta(t')^{\exists}$. To avoid unnecessary complications, we consider each 1-transition separately (recall that there are finitely many 1-transitions in $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$).

We refer to Figure 16 for an illustration of this construction. For a given position $p \in \text{dom}(t')$ such that $\theta'(p) \xrightarrow{t'(p)} (\theta'(p1), \dots, \theta'(p\ell)) \in (\delta_{\Delta}^{II})^1$, we separate the run θ' into a context $\theta_{p \leftarrow \theta(p)}^{\text{init}}$, before the 1-transition, and $\ell \stackrel{\text{def}}{=} \rho(t'(p))$ runs $\theta_1, \dots, \theta_{\ell}$, after the 1-transition, i.e., $\theta_{p \leftarrow \theta(p)}^{\text{init}}(r) \stackrel{\text{def}}{=} \theta'(r)$ for every $r \in \text{dom}(t')$ that is not a suffix of p , and $\theta_i(r) \stackrel{\text{def}}{=} \theta'(pir)$ for every $i \in [1..\ell]$ and position r with $pir \in \text{dom}(t')$.




Figure 16: Construction of t from t' (Lemma 4.25)

Then, we build t and the associated run θ by combining successive partial runs:

- ▷ start θ with $\theta_{p \leftarrow \theta(p)}^{\text{init}}$, i.e., no change from θ' above position p ,
- ▷ add a $\theta'(p)$ -reset $\theta_{u \leftarrow \theta(p)} \in \mathcal{R}_{\theta'(p)}^{\infty}(\mathcal{A}_{\Delta, \mathcal{A}}^I)$ at position p ; such a reset exists by Lemma 4.18 because $\theta'(p) \in \bullet((\delta_{\Delta}^{II})^1) \cap \bullet((\delta_{\Delta}^{II})^{\infty})$, since $\theta'(p)$ belong to a non-trivial SCC,

- ▷ pursue at position pu with the transition $\theta'(p) \xrightarrow{\alpha p} (\theta'(p1), \dots, \theta'(p\ell)) \in (\delta_{\Delta}^I)^1$ corresponding to the original 1-transition from $\mathcal{A}_{\Delta, \mathbf{A}}^I$, after adding back the equalities removed by the transformation,
- ▷ for every $i \in [1..\ell]$, introduce at position pu_i a $\theta'(pi)$ -reset $\theta_{v_i \leftarrow \theta'(pi)} \in \mathcal{R}_{\theta'(pi)}^{\infty}(\mathcal{A}_{\Delta, \mathbf{A}}^{II})$; such a reset exists by Lemma 4.18 since $\theta'(pi) \in ((\delta_{\Delta}^{II})^1)^{\bullet} \cap ((\delta_{\Delta}^{II})^{\infty})$ is a pivot state,
- ▷ continue with θ_i at position $puiv_i$, for every $i \in [1..\ell]$.

Fact C.6. $\text{tw}(\llbracket \Theta(t')^{\exists} \rrbracket) \leq \text{tw}(\llbracket \Theta(t)^{\exists} \rrbracket) + K_1$, where $K_1 \leq \max\text{Vars}(\Delta)$ is the maximal number of i -variables, $i \in \mathbb{N} \cup \{\epsilon\}$, in any 1-transition of $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$.

Proof. Let $\phi \stackrel{\text{def}}{=} \mathbf{*}_p$ not a prefix of r $t'(r)^{[r]} * \mathbf{*}_{pir \in \text{dom}(t')}$ $t'(pir)^{[puiv_i r]}$, which corresponds to the characteristic formula $\Theta(t')$ without the 1-transition at position p , and with the new position labels in t . Now $\Theta(t) = \phi * \psi$ where ψ is the separating conjunction of all $t(pr)^{[pr]}$ with $pr \in \text{dom}(t)$ and uiv_i not the prefix of r for any $i \in [1..\ell]$ (position pr is either the start of the 1-transition or part of one of the reset contexts). Let $F \stackrel{\text{def}}{=} \text{fv}(\phi) \cap \text{fv}(\psi)$ and $\eta \stackrel{\text{def}}{=} \mathbf{*} \{x = y \mid x, y \in V_f, x \approx_{\psi} y\}$. Then, F contains the variables at the extremity of the added part from t' to t , i.e., $F = \{x_k^{[p]} \mid k \in [1..\#t(p)]\} \cup \{x_j^{[puiv_i]} \mid i \in [1..\ell], j \in [1..\#t(puiv_i)]\}$. These variables exactly correspond to the parameters appearing in the 1-transition at position p of t' , i.e., $F = \{x_k^{[p]} \mid k \in [1..\#t'(p)]\} \cup \{x_j^{[pi]} \mid i \in [1..\ell], j \in [1..\#t'(pi)]\}$, hence $\text{card}(F) \leq K_1 \leq \max\text{Vars}(\Delta)$. By Definition 4.11 (2), ϕ does not induce equalities between variables of F . On the other hand, ψ only induces equalities between persistent variables of F , thanks to the introduced reset paths. These equalities η correspond exactly to those occurring in $t'(p)$, thus $\phi * \eta$ is equal to $\Theta(t')$, modulo a renaming of the variables. Now $\phi * \psi$ is satisfiable (since $\mathcal{A}_{\Delta, \mathbf{A}}^I$ is all-satisfiable), thus by Lemma 2.11, we obtain $\text{tw}(\llbracket \Theta(t')^{\exists} \rrbracket) = \text{tw}(\llbracket (\phi * \eta)^{\exists} \rrbracket) \leq \text{tw}(\llbracket (\phi * \psi)^{\exists} \rrbracket) + K_1 = \text{tw}(\llbracket \Theta(t)^{\exists} \rrbracket) + K_1$. \square

After doing this transformation (t' to t) for all 1-transitions, the final tree t satisfies $t \in \mathcal{L}(\mathcal{A}_{\Delta, \mathbf{A}}^I)$ since all the added transitions (1-transition or reset) appear in δ_{Δ}^I . With the inequality at each step (Fact C.6) and since each 1-transition of δ_{Δ}^{II} occurs exactly once in the initial tree t' (Lemma 4.7), we get $\text{tw}(\llbracket \Theta(t')^{\exists} \rrbracket) \leq \text{tw}(\llbracket \Theta(t)^{\exists} \rrbracket) + K$, where $K = \text{card}((\delta_{\Delta}^{II})^1) \cdot K_1 \leq \text{card}((\delta_{\Delta}^{II})^1) \cdot \max\text{Vars}(\Delta)$. As t' has been chosen arbitrary and $\text{card}((\delta_{\Delta}^{II})^1) = \text{card}((\delta_{\Delta}^I)^1)$, we conclude that $\text{tw}(\llbracket \mathcal{A}_{\Delta, \mathbf{A}}^{II} \rrbracket) \leq \text{tw}(\llbracket \mathcal{A}_{\Delta, \mathbf{A}}^I \rrbracket) + \text{card}((\delta_{\Delta}^I)^1) \cdot \max\text{Vars}(\Delta)$.

(3) Let $t \in \mathcal{L}(\mathcal{A}_{\Delta, \mathbf{A}}^I)$. By the construction of $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$ from $\mathcal{A}_{\Delta, \mathbf{A}}^I$, there exists a tree $t' \in \mathcal{L}(\mathcal{A}_{\Delta, \mathbf{A}}^{II})$ such that $\text{dom}(t) = \text{dom}(t')$ and $\Theta(t) = \Theta(t') * \psi$ where ψ is a conjunction of equalities. $\Theta(t')$ is a qpf formula thus by Lemma 2.10 (1), we get $\text{tw}(\llbracket \Theta(t)^{\exists} \rrbracket) = \text{tw}(\llbracket (\Theta(t') * \psi)^{\exists} \rrbracket) \leq \text{tw}(\llbracket \Theta(t')^{\exists} \rrbracket)$. As t has been chosen arbitrary, we obtain that $\text{tw}(\llbracket \mathcal{A}_{\Delta, \mathbf{A}}^I \rrbracket) \leq \text{tw}(\llbracket \mathcal{A}_{\Delta, \mathbf{A}}^{II} \rrbracket)$.

C.7. Proof of Lemma 4.28. (1) By Lemma 4.27, $\mathcal{L}(\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}) = \mathcal{L}(\mathcal{A}_{\Delta, \mathbf{A}}^{II})$, hence $\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}$ is all-satisfiable, because $\mathcal{A}_{\Delta, \mathbf{A}}^{II}$ is all-satisfiable, by Lemma 4.25 (1). Since $\widetilde{\mathcal{B}}_i$ is obtained by removing zero or more transitions from $\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II}$, we have $\mathcal{L}(\widetilde{\mathcal{B}}_i) \subseteq \mathcal{L}(\widetilde{\mathcal{A}}_{\Delta, \mathbf{A}}^{II})$, hence \mathcal{B}_i is all-satisfiable.

Let $\tilde{\delta}^1$ be any of the sets $\tilde{\delta}_1^1, \dots, \tilde{\delta}_m^1$. We prove that $\tilde{\mathcal{B}} \stackrel{\text{def}}{=} (\Sigma, \tilde{\mathcal{Q}}_\Delta^I, (q_A, \emptyset), \tilde{\delta}^1 \uplus \tilde{\delta}^\infty)$ is choice-free. To this end, we define the following labeling of transitions:

$$\Lambda(\tau) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } \tau \in \tilde{\delta}^1 \\ \infty & \text{if } \tau \in \tilde{\delta}^\infty \end{cases}$$

Let $h : \tilde{\mathcal{Q}}_\Delta^I \rightarrow \mathcal{Q}_\Delta^I$ be the function defined as $h((q, a)) \stackrel{\text{def}}{=} q$, for all $(q, a) \in \tilde{\mathcal{Q}}_\Delta^I$. For each SCC \tilde{S} of $\tilde{\mathcal{B}}$, we have $h(\tilde{S}) \subseteq S$, for an SCC S of $\mathcal{A}_{\Delta, A}^{II}$, by the definitions of $\widetilde{\mathcal{A}}_{\Delta, A}^{II}$ and $\tilde{\mathcal{B}}$. Since $\mathcal{A}_{\Delta, A}^{II}$ is choice-free, by Lemma 4.25 (1), we can extend the labeling Λ to the SCCs of $\tilde{\mathcal{B}}$, as follows:

$$\Lambda(\tilde{S}) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } h(\tilde{S}) \subseteq S, \text{ for some 1-SCC } S \text{ of } \mathcal{A}_{\Delta, A}^{II} \\ \infty & \text{if } h(\tilde{S}) \subseteq S, \text{ for some } \infty\text{-SCC } S \text{ of } \mathcal{A}_{\Delta, A}^{II} \end{cases}$$

We prove first the following fact:

Fact C.7. Let \tilde{S} be an SCC of $\tilde{\mathcal{B}}$, such that $\Lambda(\tilde{S}) = 1$. Then $h(\tilde{S})$ is a 1-SCC of $\mathcal{A}_{\Delta, A}^{II}$. Moreover, $h(\tilde{S})$ is linear if \tilde{S} is linear.

Proof. By the definition of Λ , we have that $h(\tilde{S}) \subseteq S$, for a 1-SCC of $\mathcal{A}_{\Delta, A}^{II}$, hence, for the first point, it suffices to prove that $h(\tilde{S}) = S$. Suppose, for a contradiction, that there exists a state $q \in S \setminus h(\tilde{S})$. Let $q' \in h(\tilde{S})$ be a state. Since S is an SCC of $\mathcal{A}_{\Delta, A}^{II}$, we have $q \rightsquigarrow^* q' \rightsquigarrow^* q$ in $\mathcal{A}_{\Delta, A}^{II}$. Since $q' \in h(\tilde{S})$, there exists an injective partial mapping $a' : [1.. \#q'] \rightarrow [1.. \mathcal{M}]$ such that $(q', a') \in \tilde{S}$. By the construction of $\widetilde{\mathcal{A}}_{\Delta, A}^{II}$, we obtain $(q', a') \rightsquigarrow^* (q, a)$ in $\widetilde{\mathcal{A}}_{\Delta, A}^{II}$, for some injective partial mapping $a : [1.. \#q] \rightarrow [1.. \mathcal{M}]$. To show a contradiction, we prove that $(q, a) \rightsquigarrow^* (q', a')$ in $\widetilde{\mathcal{A}}_{\Delta, A}^{II}$, hence $q' \in h(\tilde{S})$. By the construction of $\widetilde{\mathcal{A}}_{\Delta, A}^{II}$, there exist injective partial mappings $a_1, a_2, \dots : [1.. \#q'] \rightarrow [1.. \mathcal{M}]$, such that $(q, a) \rightsquigarrow^* (q', a_1) \rightsquigarrow^* (q', a_2) \rightsquigarrow^* \dots$. Because the underlying transitions of $\mathcal{A}_{\Delta, A}^{II}$ along this path are ∞ -transitions, the mappings a_1, a_2, \dots are undefined everywhere except for $\mathfrak{P}_{\mathcal{A}_{\Delta, A}^{II}}(q')$, where they are defined. Because these transitions are from $\bullet S \bullet$, each of these mappings can be obtained from the previous one by composition with a permutation over $\mathfrak{P}_{\mathcal{A}_{\Delta, A}^{II}}(q')$. Since all such permutations can be enumerated in this way, this leads to $(q', a'_1) \rightsquigarrow^* (q', a')$ in $\widetilde{\mathcal{A}}_{\Delta, A}^{II}$, i.e., contradiction.

For the second point, suppose, for a contradiction, that $h(\tilde{S})$ is not linear. Then, there exists a transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell)$ such that $q_0, q_i, q_j \in h(\tilde{S})$, for some indices $1 \leq i < j \leq \ell$. Since $h(\tilde{S})$ is a connected component of $\mathcal{A}_{\Delta, A}^{II}$, we have $q_i \rightsquigarrow^* q_0$ and $q_j \rightsquigarrow^* q_0$ in $\mathcal{A}_{\Delta, A}^{II}$. Let a_0, \dots, a_ℓ be injective partial mappings, such that $(q_0, a_0) \xrightarrow{\alpha} ((q_1, a_1), \dots, (q_\ell, a_\ell))$ is a transition of $\tilde{\mathcal{B}}$. By the above argument, we obtain $(q_i, a_i) \rightsquigarrow^* (q_0, a_0)$ and $(q_j, a_j) \rightsquigarrow^* (q_0, a_0)$, hence \tilde{S} is non-linear, contradiction. \square

Back to the main proof, we check the points of Definition 4.5 below:

(1) The SCC graph of $\tilde{\mathcal{B}}$ is a tree, by the choice of the set $\tilde{\delta}^1$ corresponding to $\tilde{\mathcal{B}}$. Moreover, each non-root SCC of $\tilde{\mathcal{B}}$ is entered by exactly one branch, since this is already the case for the choice-free automaton $\mathcal{A}_{\Delta, A}^{II}$.

(2a) Let \tilde{S} be a linear SCC of $\tilde{\mathcal{B}}$, such that $\Lambda(\tilde{S}) = 1$. By Fact C.7, $h(\tilde{S})$ is a linear 1-SCC of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$, hence $\text{card}(h(\tilde{S})^\bullet) = 1$. Then, we obtain $\text{card}(\tilde{S}^\bullet) = 1$.

(2b) Let $\tilde{\tau} = (q_0, a_0) \xrightarrow{\alpha} ((q_1, a_1), \dots, (q_\ell, a_\ell))$ be a transition of $\tilde{\mathcal{B}}$. Then $\Lambda(\tilde{\tau}) = 1$ iff $\tilde{\tau} \in \tilde{\delta}^1$, by the definition of Λ . “ \Rightarrow ” Assume that $\tilde{\tau} \in \tilde{\delta}^1$. Then, the underlying transition $\tau = q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell)$ is a 1-transition of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$. Since $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ is choice-free, we have that $\tau \in S^\bullet$, for some linear 1-SCC S of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$. Let \tilde{S} be an SCC of $\tilde{\mathcal{B}}$ such that $h(\tilde{S}) \subseteq S$. Then $\Lambda(\tilde{S}) = 1$, by the definition of Λ . Suppose, for a contradiction that \tilde{S} is non-linear. Then, $h(\tilde{S})$ is non-linear, hence S is non-linear, contradiction. Finally, $\tilde{\tau} \in \tilde{S}^\bullet$ follows from $\tau \in S^\bullet$. “ \Leftarrow ” Assume that $\tilde{\tau} \in \tilde{S}^\bullet$, for a linear SCC \tilde{S} of $\tilde{\mathcal{B}}$, such that $\Lambda(\tilde{S}) = 1$. Then $h(\tilde{S})$ is a linear 1-SCC of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$, by Fact C.7. Moreover, $\tau \in h(\tilde{S})^\bullet$, where τ is the underlying transition of $\tilde{\tau}$ from $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$. Hence τ is a 1-transition of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ and $\tilde{\tau} \in \tilde{\delta}^1$, by the choice of $\tilde{\delta}^1$.

(2c) By Fact C.7, $\Lambda(\tilde{S}) = 1$ iff $h(\tilde{S})$ is a 1-SCC of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$, for each SCC \tilde{S} of $\tilde{\mathcal{B}}$. “ \Rightarrow ” Assume that $h(\tilde{S})$ is a 1-SCC of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$. Since $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ is choice-free, then $h(\tilde{S})$ is the root of the SCC tree of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$ or $\tilde{S} = \{\tilde{\tau}\}$, for some transition $\tilde{\tau}$ of $\tilde{\mathcal{B}}$, such that $\Lambda(\tilde{\tau}) = 1$. In the first case, $h(\tilde{S})$ has no incoming transition, hence \tilde{S} is the root of the SCC tree of $\tilde{\mathcal{B}}$. In the second case, $h(\tilde{S}) = \{\tau\}$, for a 1-transition τ of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$. Then, $\tilde{S} = \{\tilde{\tau}\}$, for a transition $\tilde{\tau}$, such that $\Lambda(\tilde{\tau}) = 1$, by the definition of Λ . “ \Leftarrow ” If $h(\tilde{S})$ is the root of the SCC tree of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$, then \tilde{S} has no incoming transitions, hence it is the root of the SCC tree of $\tilde{\mathcal{B}}$. If $h(\tilde{S}) = \{\tau\}$, for a 1-transition τ of $\mathcal{A}_{\Delta, \mathcal{A}}^{II}$, we obtain $\tilde{S} = \{\tilde{\tau}\}$, for a transition $\tilde{\tau}$ of $\tilde{\mathcal{B}}$, such that $\Lambda(\tilde{\tau}) = 1$, by the definition of Λ .

(2) Since $\mathcal{L}(\tilde{\mathcal{B}}_i) \subseteq \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^{II})$, for each $i \in [1..m]$, we have $\bigcup_{i=1}^m \mathcal{L}(\tilde{\mathcal{B}}_i) \subseteq \mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^{II})$. By the definition of $\tilde{\mathcal{B}}_1, \dots, \tilde{\mathcal{B}}_m$, for each accepting run θ of $\widetilde{\mathcal{A}}_{\Delta, \mathcal{A}}^{II}$, there exists $i \in [1..m]$, such that θ is an accepting run of $\tilde{\mathcal{B}}_i$, thus $\mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^{II}) \subseteq \bigcup_{i=1}^m \mathcal{L}(\tilde{\mathcal{B}}_i)$. \square

C.8. Proof of Lemma 4.29. First, we prove the following fact:

Fact C.8. Let $t \in \mathcal{L}(\tilde{\mathcal{B}})$ be a tree, θ be an accepting run of $\tilde{\mathcal{B}}$ over t and $p \in \text{dom}(t)$ be a position such that $\theta(p) = (q, a)$. Then, $x_i^{[p]} \approx_{\Theta(t)} y_{a(i)}^{[r_i]}$, for each $i \in \text{dom}(a)$, such that r_i is the unique position where a variable $y_{a(i)}^{[e]} \in \mathcal{Y}$ occurs in $\Theta(t)$.

Proof. By induction on the structure of t . \square

(1) Let $\bar{t} \in \mathcal{L}(\bar{\mathcal{B}})$ be a tree. By the construction of $\bar{\mathcal{B}}$, the accepting run $\bar{\theta}$ of $\bar{\mathcal{B}}$ over \bar{t} can be transformed into an accepting run θ of $\tilde{\mathcal{B}}$ over a tree t , such that $\text{dom}(t) = \text{dom}(\bar{t})$, by changing the labels $\bar{\alpha}$ back to the original labels α . Since $\tilde{\mathcal{B}}$ is all-satisfiable, by Lemma 4.28 (1), the formula $\Theta(t)$ is satisfiable, and let (\mathbf{U}, σ) be a structure and \mathfrak{s} be a store such that $(\mathbf{U}, \sigma) \models^{\mathfrak{s}} \Theta(t)$. However, $\Theta(\bar{t})$ is obtained from $\Theta(t)$ by removing several (dis-)equalities from the labels of 1-transitions and by changing each relation atom $r(z_1^{[p]}, \dots, z_{\#r}^{[p]})$ into a relation atom $r_g(\zeta_\tau(z_{i_1}^{[p]}), \dots, \zeta_\tau(z_{i_k}^{[p]}))$, according to the construction of $\bar{\mathcal{B}}$. Suppose, for a contradiction, that there exist two distinct positions $p_1, p_2 \in \text{dom}(\bar{t})$ such that the relation atoms $r_g(\zeta_{\tau_1}(z_{i_1}^{[p_1]}), \dots, \zeta_{\tau_1}(z_{i_k}^{[p_1]}))$ and $r_g(\zeta_{\tau_2}(z_{i_1}^{[p_2]}), \dots, \zeta_{\tau_2}(z_{i_k}^{[p_2]}))$ occur in $\Theta(\bar{t})$ and $\zeta_{\tau_1}(z_{i_j}^{[p_1]}) \approx_{\Theta(\bar{t})} \zeta_{\tau_2}(z_{i_j}^{[p_2]})$, for all $j \in [1..k]$. Then, there exist relation atoms $r(z_1^{[p_1]}, \dots, z_{\#r}^{[p_1]})$

and $r(z_1^{[p_2]}, \dots, z_{\#r}^{[p_2]})$ that occur in $\Theta(t)$, such that $i_1, \dots, i_k \in [1..\#r]$ and, by Fact C.8 and properties of the renaming, we obtain that $z_j^{[p_1]} \approx_{\Theta(t)} z_j^{[p_2]}$, for all $j \in [1..\#r]$. This however contradicts the satisfiability of $\Theta(t)$, hence such relation atoms cannot occur in $\Theta(\bar{t})$. Similarly, we obtain a contradiction if we suppose a disequality atom $\zeta_\tau(x) \neq \zeta_\tau(y)$ occurs in $\Theta(\bar{t})$ and $\zeta_\tau(x) \approx_{\Theta(\bar{t})} \zeta_\tau(y)$. Then, $\Theta(\bar{t})$ is satisfiable and, since the choice of \bar{t} was arbitrary, $\bar{\mathcal{B}}$ is all-satisfiable.

(2) Let $\bar{S} = (\mathbf{U}, \bar{\sigma}) \in \llbracket \bar{\mathcal{B}} \rrbracket$ be a structure. Then, there exists a tree $\bar{t} \in \mathcal{L}(\bar{\mathcal{B}})$ and a store $\bar{\mathfrak{s}}$ such that $\bar{S} \models^{\bar{\mathfrak{s}}} \Theta(\bar{t})$. Let \tilde{t} be the tree obtained from \bar{t} by replacing back each label $\bar{\alpha}$ with the original label α , according to the construction of $\bar{\mathcal{B}}$ from $\tilde{\mathcal{B}}$. Note that $\text{dom}(\tilde{t}) = \text{dom}(\bar{t})$ and let $\bar{\theta}, \tilde{\theta}$ be two related accepting runs of $\bar{\mathcal{B}}, \tilde{\mathcal{B}}$ over \bar{t}, \tilde{t} respectively. We define a store $\tilde{\mathfrak{s}}$ over $\text{fv}(\Theta(\tilde{t}))$ as follows. Let p be an arbitrary position in $\text{dom}(\tilde{t})$. Let (q_0, a_0) be the state assigned to p and $\tau : (q_0, a_0) \xrightarrow{\alpha} ((q_1, a_1), \dots, (q_\ell, a_\ell))$ be the transition taken at p in the run $\bar{\theta}$. Let ζ_τ be the renaming of variables as defined in the transformation. The store $\tilde{\mathfrak{s}}$ is then defined by:

$$\begin{aligned} \triangleright \tilde{\mathfrak{s}}(y_i^{[p]}) &\stackrel{\text{def}}{=} \begin{cases} \bar{\mathfrak{s}}(y_i^{[p]}) & \text{if } \tau \notin \tilde{\delta}^1 \\ u_i & \text{if } \tau \in \tilde{\delta}^1 \text{ for some distinct values } u_i \notin \bar{\mathfrak{s}}(\text{fv}(\Theta(\tilde{t}))), i \in [1..\mathcal{M}] \end{cases} \\ \triangleright \tilde{\mathfrak{s}}(x_i^{[p]}) &\stackrel{\text{def}}{=} \begin{cases} \bar{\mathfrak{s}}(x_j^{[p]}) & \text{if } i \notin \text{dom}(a_0) \text{ and } \zeta_\tau(x_i^{[\epsilon]}) = x_j^{[\epsilon]} \\ u_{a_0(i)} & \text{if } i \in \text{dom}(a_0) \end{cases} \end{aligned}$$

That is, the store $\tilde{\mathfrak{s}}$ allocates fresh distinct values for all the persistent variables and re-uses the already given values for the non-persistent ones in $\bar{\mathfrak{s}}$, while taking into account their renaming by ζ_τ . We now consider the structure $\tilde{S} = (\mathbf{U}, \tilde{\sigma})$, where $\tilde{\sigma}$ is the interpretation that assigns to each relation symbol r the set of tuples $\langle \tilde{\mathfrak{s}}(\xi_1), \dots, \tilde{\mathfrak{s}}(\xi_{\#r}) \rangle$ for every relation atom $r(\xi_1, \dots, \xi_{\#r})$ in $\Theta(\tilde{t})$. We can now check that $(\mathbf{U}, \tilde{\sigma}) \models^{\tilde{\mathfrak{s}}} \Theta(\tilde{t})$. Any two tuples from $\tilde{\sigma}(r)$ defined as above are necessarily distinct, unless the corresponding tuples restricted to non-persistent variables from $\bar{\sigma}(r_g)$ are not distinct. Similarly, equalities on disequality atoms on non-persistent variables only hold in $\Theta(\tilde{t})$ as they already hold in the renamed form in $\Theta(\bar{t})$. Finally, equalities and disequalities in $\Theta(\tilde{t})$ involving persistent variables hold by the construction of $\tilde{\mathfrak{s}}$. But now, every tree decomposition of \tilde{S} is a tree decomposition for \bar{S} . Hence, $\text{tw}(\bar{S}) \leq \text{tw}(\tilde{S})$ and consequently, since the choice of \bar{S} was arbitrary we have $\text{tw}(\llbracket \bar{\mathcal{B}} \rrbracket) \leq \text{tw}(\llbracket \tilde{\mathcal{B}} \rrbracket)$.

(3) Let $\tilde{S} = (\mathbf{U}, \tilde{\sigma}) \in \llbracket \tilde{\mathcal{B}} \rrbracket$ be a structure. Then, there exists a tree $\tilde{t} \in \mathcal{L}(\tilde{\mathcal{B}})$ and a store $\tilde{\mathfrak{s}}$, such that $\tilde{S} \models^{\tilde{\mathfrak{s}}} \Theta(\tilde{t})$. Let $\bar{t} \in \mathcal{L}(\bar{\mathcal{B}})$ be the tree obtained by changing each label α of \tilde{t} into $\bar{\alpha}$, according to the construction of $\bar{\mathcal{B}}$. We consider the structure $\bar{S} \stackrel{\text{def}}{=} (\mathbf{U}, \bar{\sigma})$, where $\bar{\sigma}$ interprets each relation symbol r_g by the set of tuples $\langle \bar{\mathfrak{s}}(\zeta_\tau(z_{i_1}^{[p]})), \dots, \bar{\mathfrak{s}}(\zeta_\tau(z_{i_k}^{[p]})) \rangle$, such that $r_g(\zeta_\tau(z_{i_1}^{[p]}), \dots, \zeta_\tau(z_{i_k}^{[p]}))$ occurs in $\Theta(\bar{t})$. Let \bar{T} be a tree decomposition of \bar{S} . We consider the tree decomposition \tilde{T} obtained by adding the values $\tilde{\mathfrak{s}}(y_i^{[r_i]})$, where r_i is the unique position where a ϵ -variable $y_i^{[\epsilon]} \in \mathcal{Y}$ occurs in $\Theta(\bar{t})$, to the label of each node in \bar{T} . Then, \tilde{T} is a tree decomposition of \tilde{S} and moreover $\text{wd}(\tilde{T}) \leq \text{wd}(\bar{T}) + \mathcal{M}$. Since the choice of \bar{T} was arbitrary we obtain $\text{tw}(\tilde{S}) \leq \text{tw}(\bar{S}) + \mathcal{M}$. Consequently, since the choice of \tilde{S} was arbitrary and $\mathcal{M} \leq \text{card}(\tilde{\delta}^1) \cdot \max\text{Vars}(\Delta)$, we obtain that $\text{tw}(\llbracket \tilde{\mathcal{B}} \rrbracket) \leq \text{tw}(\llbracket \bar{\mathcal{B}} \rrbracket) + \text{card}(\tilde{\delta}^1) \cdot \max\text{Vars}(\Delta)$. \square

C.9. Proof of Lemma 4.32. By hypothesis, $\overline{\mathcal{B}}$ has no persistent variables and since the relabeling of 1-transitions introduces no existential quantifiers, then \mathcal{B} has no persistent variables either.

(1) Let $t \in \mathcal{L}(\mathcal{B})$ be a tree and let θ be an accepting run of \mathcal{B} over t . Let us consider the run θ' of $\overline{\mathcal{B}}$ obtained by replacing with **emp** the labels of the 1-transitions $\tau^{\text{exp}} : q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta^1$. We then define the run θ'' of $\overline{\mathcal{B}}$ by replacing each occurrence of a 1-transition $\tau : q_0 \xrightarrow{\text{emp}} (q_1, \dots, q_\ell)$ in θ' by the partial run θ_τ constructed by extending τ with the resets used in the definition of the transformation, namely:

- ▷ a pre- q_0 -reset $\theta_{p_0 \leftarrow q_0}^0 \in \mathcal{R}_{q_0}^\infty(\overline{\mathcal{B}})$ if $q_0 \in \bullet(\delta^\infty)$ and,
- ▷ a post- q_i -reset $\theta_{p_i \leftarrow q_i}^i \in \mathcal{R}_{q_i}^\infty(\overline{\mathcal{B}})$ for each $i \in [1..\ell]$ such that $q_i \in \bullet(\delta^\infty)$.




Figure 17: Expansion of 1-transitions

The relation between τ from θ_τ is illustrated in Figure 17 (a,b). Note that the resets exist according to Lemma 4.18. Then, it is easy to check that θ'' is indeed a run of $\overline{\mathcal{B}}$. Let $t'' \in \mathcal{L}(\overline{\mathcal{B}})$ be the tree accepted by θ'' . Since $\overline{\mathcal{B}}$ is all-satisfiable, there exists a store \mathfrak{s} and a structure S'' such that $S'' \models^{\mathfrak{s}} \Theta(t'')$. Now, by the definition of \mathcal{B} , the label α of each transition $\tau^{\text{exp}} : q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta^1$ is actually $\omega_0 * \bigstar_{j=1}^{\ell} \omega_j$, where:

$$\begin{aligned} \omega_0 &= \bigstar_{i_1, \dots, i_k \in [1..\#q_0]} \Omega_{t_0}^{\epsilon/\epsilon}(x_{i_1}, \dots, x_{i_k}) \\ \omega_j &= \bigstar_{i_1, \dots, i_k \in [1..\#q_j]} \Omega_{t_j}^{p_j/j}(x_{i_1}, \dots, x_{i_k}), \text{ for } j \in [1..\ell] \end{aligned}$$

and t_0, t_1, \dots, t_ℓ are the Σ -labelled trees of the corresponding resets, respectively $\theta_0, \theta_1, \dots, \theta_\ell$. This construction is illustrated in Figure 17(c). Hence, there exists a substructure $S \sqsubseteq S''$ such that $S \models^{\mathfrak{s}} \Theta(t)$. Actually, S is simply obtained from S'' by removing all but the tuples introduced by the relation atoms of $\omega_0, \dots, \omega_\ell$, defined along the partial runs θ_τ of θ'' . Since the choice of t was arbitrary, \mathcal{B} is all-satisfiable.

(2) Let $t' \in \mathcal{L}(\mathcal{B})$ be a tree. Let $t \in \mathcal{L}(\overline{\mathcal{B}})$ be the tree obtained by t' be removing relational atoms from the labels of 1-transitions, that is, reversing the transformation. As both \mathcal{B} and $\overline{\mathcal{B}}$ are all-satisfiable, the characteristic formulæ $\Theta(t')$ and $\Theta(t)$ are satisfiable. They differ, moreover, only by finitely many relational atoms. Henceforth, by using Lemma 2.10 (4), we obtain that the difference between the treewidth of their models is bounded by the number of free variables occurring on these atoms, that is, at most $\text{card}(\overline{\delta}^1) \cdot \max\text{Vars}(\Delta)$. Then $\text{tw}(\llbracket \Theta(t')^{\exists} \rrbracket) \leq \text{tw}(\llbracket \Theta(t)^{\exists} \rrbracket) + \text{card}(\overline{\delta}^1) \cdot \max\text{Vars}(\Delta)$ and as the choice of t' was arbitrary, we obtain $\text{tw}(\llbracket \mathcal{B} \rrbracket) \leq \text{tw}(\llbracket \overline{\mathcal{B}} \rrbracket) + \text{card}(\overline{\delta}^1) \cdot \max\text{Vars}(\Delta)$.

(3) Let $t \in \mathcal{L}(\overline{\mathcal{B}})$ be a tree. Let $t' \in \mathcal{L}(\mathcal{B})$ be the tree obtained from t by replacing the **emp**-labels of 1-transitions by their corresponding formulæ according to the transformation.

As both $\overline{\mathcal{B}}$ and \mathcal{B} are all-satisfiable, the characteristic formulæ $\Theta(t)$ and $\Theta(t')$ are satisfiable. These formulæ differ only by a finite number of relation atoms, that is, the ones inserted in the labels of the modified 1-transitions. Henceforth, by Lemma 2.10 (3), we obtain that the difference between the treewidth of their models is bounded by the number K of these relation atoms. As the choice of t was arbitrary, we obtain that $\text{tw}(\overline{\mathcal{B}}) \leq \text{tw}(\mathcal{B}) + K$.

We can further obtain an upper bound on K as follows. First, note that for a relation symbol $r \in \mathbb{R}$ and m distinct variables there exists at most distinct $m^{\#r}$ relation atoms. Note that, in our context $m \leq \text{maxVars}(\Delta)$ and $\#r \leq \text{maxRelArity}(\Delta)$. Second, such atoms could be added for every 1-transition, for every one of its states q_i . That is, overall we obtain $K \leq \text{card}(\overline{\delta}^1) \cdot (1 + \text{maxPredAtoms}(\Delta)) \cdot \text{relNo}(\Delta) \cdot \text{maxVars}(\Delta)^{\text{maxRelArity}(\Delta)}$. \square