



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-04264291v1

HAL Id: hal-04264291

<https://hal.science/hal-04264291v1>

Preprint submitted on 30 Oct 2023 (v1), last revised 23 May 2024 (v2)

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 Bozga¹, Lucas Bueri¹, Radu Iosif¹, and Florian Zuleger²

¹ Univ. Grenoble Alpes, CNRS, Grenoble INP, VERIMAG, 38000, France

² 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 for the treewidth boundedness problem in the context of this logic.

1 Introduction

The treewidth of a graph is a positive integer measuring the distance between the graph and 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 $n \times n$ square grids have treewidth n , for any $n \geq 1$. The treewidth is a cornerstone of algorithmic tractability. For instance, many NP-complete graph problems such as Hamiltonicity and 3-Coloring become P, when restricted to inputs whose treewidth is bounded by a constant, see, e.g., [23, Chapter 11].

Structures are interpretations of relation symbols that define the standard semantics of first and second order logic [18]. 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 [16] proves that MSO is decidable over bounded treewidth structures, by reduction to the emptiness problem of tree automata. A dual result of Seese [34] proves that each class of structures with a decidable MSO theory necessarily has bounded treewidth. Since MSO is the yardstick of graph specification logics [17], 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 [28], databases [1] and concurrency [19].

This paper considers the *treewidth boundedness problem* asking for the existence of an upper bound on the treewidths of the models of a given input formula. For first order logic (and implicitly MSO) the problem is already undecidable (Theorem 1), hence we focus on non-classical substructural logics³. We prove the decidability of this prob-

³ Substructural logics have, in addition to boolean conjunction, a conjunction-like connective, for which Gentzen's natural deduction rules of weakening and contraction do not hold.

lem for a generalization of Separation Logic to relational signatures, interpreted over structures (Theorem 2).

Separation Logic (SL) [27,33,14] 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 this context, the separating conjunction is interpreted as the union of heaps with disjoint domains. SL interpreted over heaps 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* [12], is at the heart of scalable compositional proof techniques for pointer programs [11].

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 [30]. 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 (hyper-)graphs by structures, SLR allows to specify (hyper-)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 [14].

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 [2,6,7]. 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 datastructures (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.

A key ingredient of automated proof generation in Hoare logic is the availability of a decision procedure for 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 (1) ϕ having only treewidth-bounded models, and (2) both ϕ and ψ being MSO-definable [25]. The decidability result from this paper (Theorem 2) defines precisely those formulæ of SLR whose models form a treewidth-bounded set.

Motivating examples We introduce the reader to SLR and the treewidth boundedness problem by means of examples. Fig. 1 (a) shows a chain $A(x_1, x_2)$ starting at x_1 and ending at x_2 , whose elements are labeled by a monadic relation symbol a and 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 the 'a' set and the semantics of the separating conjunction requires that these sets

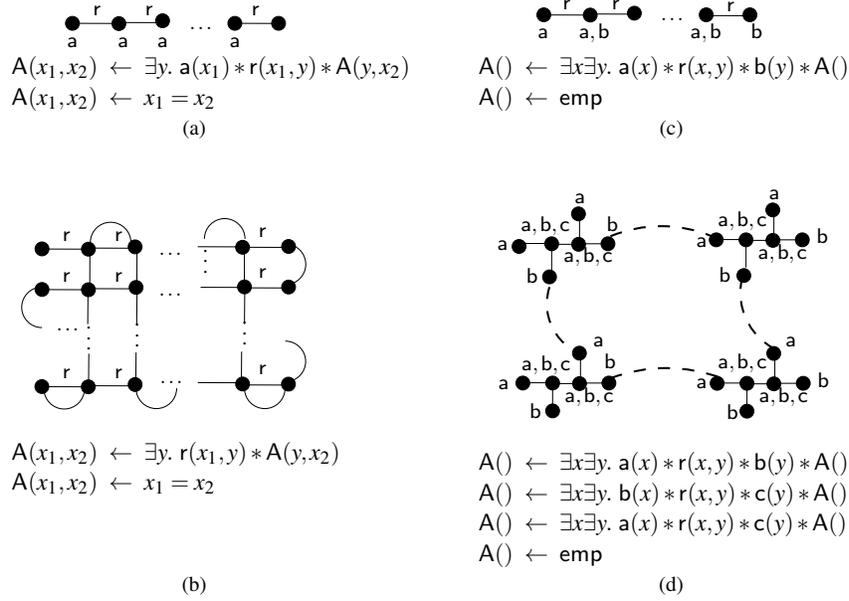


Fig. 1. Examples of Bounded and Unbounded Treewidth Models

must be disjoint in the models of ϕ_1 and ϕ_2 , that compose into a model of $\phi_1 * \phi_2$. Then, any model of $\exists x \exists y. A(x, y)$ is a (possibly cyclic) chain, of treewidth at most two.

Fig. 1 (b) shows a family of models for a slightly modified definition of a chain, 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 a 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 formula $\exists x \exists y. A(x, y)$ 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., Fig. 1 (a), a natural question is what happens when these instances are placed into two (not necessarily disjoint) sets? The inductive definition of the predicate A in Fig. 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 at most two⁴.

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 Fig. 1 (d), namely a - b , b - c and a - c edges. In this case, the formula $A()$, 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 Fig. 1 (d). Consequently, the models of $A()$ form a set of unbounded treewidth.

These examples highlight the ideas behind an algorithm that decides the existence of a bound on the treewidths of the models of a given formula, with predicates interpreted

⁴ For instance, a simple cycle with more than two elements has treewidth two.

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 Fig. 1. A complication is that such sets could 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 Fig. 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 Fig. 1 (d), the building blocks of the grid are structures consisting of six elements, that connect two ‘a’ with two ‘b’ elements.

For space reasons, additional technical material relative to Sections §2, §3, §4, §5 and §6 is given in Appendix §A, §B, §C, §D and §E, respectively.

2 The Treewidth Boundedness Problem

This section defines formally the treewidth boundedness problem and introduces most of the technical definitions.

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., $S = S_1 \cup S_2$ and $S_1 \cap S_2 = \emptyset$.

Multisets are denoted as $[a, b, \dots]$ and all set operations are used with multisets as well. In particular, a binary operation involving a set and a multiset considers the set to be a multiset and yields 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 $(a, b), (a', b') \in R$.

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

A *structure* is a pair $S = (U, \sigma)$, where U is an *infinite* set called the *universe* and $\sigma : \mathbb{R} \rightarrow \text{pow}(U^+)$ is an *interpretation* mapping each relation symbol r into a *finite* subset of $U^{\#r}$. We consider only structures with finite interpretations, because SLR (defined below) can only describe such 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. Two structures (U_1, σ_1) and (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 (a formal definition is given in [20, §A3]).

We consider several operations on structures. The first operation is *composition*, defined as pointwise disjoint union of the interpretations of relation symbols:

Definition 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.*

We define two fusion operations, that glue together elements from the same structure (*internal fusion*) or from distinct structures (*external fusion*). Fusion operations are formally defined via quotienting with respect to certain equivalence relations:

Definition 2. 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 $\sigma_{/\approx}(r) \stackrel{\text{def}}{=} \{([u_1]_{\approx}, \dots, [u_{\#r}]_{\approx}) \mid \langle u_1, \dots, u_{\#r} \rangle \in \sigma(r)\}$, for all $r \in \mathbb{R}$.

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. An equivalence relation $\approx \subseteq U \times U$ is compatible with a structure $S = (U, \sigma)$ 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 S is a structure isomorphic to $S_{/\approx}$, for an equivalence relation \approx compatible with S . Let $\text{IF}(S)$ be the set of internal fusions of S and $\text{IF}(S) \stackrel{\text{def}}{=} \bigcup_{S \in \mathcal{S}} \text{IF}(S)$, for a set \mathcal{S} of structures.

External fusion is a binary operation that glues elements taken from different structures:

Definition 4. 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.

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 there is an undirected path in G between any two nodes in 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 5. A tree decomposition of a structure $S = (U, \sigma)$ is a $\text{pow}(U)$ -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 $\{u_1, \dots, u_{\#r}\} \subseteq \lambda(n)$, and
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

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

(a)

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

(b)

Fig. 2. The Separation Logic of Relations

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 [5].

Logics Let $\mathbb{V} = \{x, y, \dots\}$ be a set of variables. *First order logic* (FO) is the set of formulae consisting of *equalities* $x = y$ and *relation atoms* $r(x_1, \dots, x_{\#r})$ connected by boolean conjunction, negation and existential quantification. 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*. The semantics of first order logic is given by a satisfaction relation $(\mathbb{U}, \sigma) \models^{\mathfrak{s}} \phi$ between structures and formulae, parameterized by a *store* $\mathfrak{s} : \mathbb{V} \rightarrow \mathbb{U}$ such that $(\mathbb{U}, \sigma) \models^{\mathfrak{s}} r(x_1, \dots, x_{\#r})$ 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 $S \models \phi$ instead. The set of *models* of a FO sentence ϕ is denoted as $\llbracket \phi \rrbracket \stackrel{\text{def}}{=} \{S \mid S \models \phi\}$.

The *Separation Logic of Relations* (SLR) uses a set of *predicates* $\mathbb{P} = \{A, B, \dots\}$ with given arities $\#A \geq 0$. A predicate of zero arity is called *nullary*. The formulae of SLR are defined by the syntax in Fig. 2 (a). Instead of the boolean conjunction, SLR has a *separating conjunction* $*$. The formulae $x \neq y$ 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.

Definition 6. A set of inductive definitions (SID) is a finite set 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 $(\mathbb{U}, \sigma) \models_{\Delta}^{\mathfrak{s}} \phi$ between structures and formulae, 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 being denoted as $\phi[x_1/y_1, \dots, x_n/y_n]$. As a 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 Fig. 2 (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 (U, σ) , such that $(U, \sigma) \models_{\Delta}^s \phi$, 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 s (resp. the SID Δ) from $S \models_{\Delta}^s \phi$. For a SLR sentence ϕ , let $[[\phi]]_{\Delta} \stackrel{\text{def}}{=} \{S \mid S \models_{\Delta} \phi\}$ be the set of Δ -models of ϕ . If ϕ is, moreover, predicate-free we omit the subscript and say that ϕ is *satisfiable* iff $[[\phi]] \neq \emptyset$. Although we use the same notation for the sets of models of FO and SLR formulae, the underlying logic will always be clear from the context.

Given predicate-free formulae ϕ and ψ , we say that ψ is a *logical consequence* of ϕ iff for each structure (U, σ) and store s , such that $(U, \sigma) \models^s \phi$, there exists a structure (U, σ') such that $\sigma'(r) \subseteq \sigma(r)$, for all $r \in \mathbb{R}$, and $(U, \sigma') \models^s \psi$.

Decision Problems We consider the following decision problems:

Definition 7. *The TWB^{FO} problem asks, given a FO sentence ϕ , if the set $[[\phi]]$ is treewidth-bounded. The TWB^{SLR} problem asks, given a SID Δ and a SLR sentence ϕ , if the set $[[\phi]]_{\Delta}$ is treewidth-bounded.*

First, we prove that TWB^{FO} is undecidable:

Theorem 1. *The problem is $[[\phi]]$ treewidth-bounded, for a given FO sentence ϕ with at least two binary relation symbols and several unary relations, is undecidable.*

See proof on page 29.

The rest of this paper is concerned with the proof of the following theorem:

Theorem 2. *The TWB^{SLR} problem is decidable.*

Proof overview Let (Δ, ϕ) be an instance of the TWB^{SLR} problem, where Δ is a SID and ϕ is a SLR sentence. Without loss of generality, we assume that ϕ consists of a single nullary predicate atom A , since $[[\phi]]_{\Delta} = [[A]]_{\Delta \cup \{A \leftarrow \phi\}}$ provided that A is not defined by any other rule in Δ . We use a notion of *canonical model*, intuitively, structures obtained by instantiating all existential quantifiers introduced by the unfolding of the rules in Δ to pairwise distinct variables. The set of canonical Δ -models of A is denoted $[[A]]_{\Delta}^c$. A formal definition is given in §3 (Def. 8). We reduce the treewidth boundedness problem for $[[A]]_{\Delta}$ to the same problem for a set generated by external fusion of structures that may occur any number of times embedded in some canonical Δ -model of A .

The first part of the proof builds finitely many SIDs $\Gamma_1, \dots, \Gamma_n$ having a semantic property, called *expandability*. Intuitively, a SID Γ is expandable for a nullary predicate B iff any sequence of canonical Γ -models of B can be disjointly embedded in some canonical Γ -model of B . A formal definition is given in §3 (Def. 12) and the proof of the lemma below is given in §5:

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

The above result allows to assume w.l.o.g. that Δ is expandable for A .

In the second part of the proof we reduce the treewidth boundedness of $\llbracket A \rrbracket_\Delta$ to the treewidth boundedness of sets of structures obtained by applying both internal and external fusion to canonical models. The proof of the lemma below is given in §5:

Lemma 2. *Let Δ be an expandable SID for a nullary predicate A . Then, (1) only if (2) only if (3), where:*

- (1) $\text{IEF}^*(\llbracket A \rrbracket_\Delta^c)$ is treewidth-bounded,
- (2) $\llbracket A \rrbracket_\Delta$ is treewidth-bounded,
- (3) $\text{EF}^*(\llbracket A \rrbracket_\Delta^c)$ is treewidth-bounded.

In the third part of the proof we establish the equivalence of the points (1-3) of Lemma 2, by proving the missing direction (3) only if (1). This reduces the treewidth boundedness of $\llbracket A \rrbracket_\Delta$ to the treewidth boundedness of $\text{EF}^*(\llbracket A \rrbracket_\Delta^c)$. The proofs of the following two lemmas are given in §6:

Lemma 3. *Given a SID Δ and a nullary predicate symbol A , $\text{EF}^*(\llbracket A \rrbracket_\Delta^c)$ is treewidth-bounded only if $\text{IEF}^*(\llbracket A \rrbracket_\Delta^c)$ is treewidth-bounded.*

The above lemma is a consequence of the argument used to show the decidability of the treewidth boundedness problem for sets of the form $\text{EF}^*(\llbracket A \rrbracket_\Delta^c)$:

Lemma 4. *The following problem is decidable: given a SID Δ and a nullary predicate A , is $\text{EF}^*(\llbracket A \rrbracket_\Delta^c)$ treewidth-bounded?*

Finally, the treewidth boundedness for $\text{EF}^*(\llbracket A \rrbracket_\Delta^c)$ is shown to be equivalent to the treewidth boundedness of a set generated by external fusion of a set \mathcal{S} of *connected* structures, i.e., in which there is a path of tuples between any two elements from the support. We prove that (1) $\text{EF}^*(\mathcal{S})$ is treewidth-unbounded iff (2) $\text{EF}^*(\mathcal{S})$ contains infinitely many grid minors iff (3) there exist two disjoint structures $(U_i, \sigma_i) \in \text{EF}^*(\mathcal{S})$ and distinct elements $u_i, v_i, w_i \in \text{supp}(\sigma_i)$ labeled with disjoint sets of relation symbols C_i , for $i = 1, 2$. For the latter condition, Fig. 3 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, the condition $C_1 \cap C_2 = \emptyset$ allows to glue the elements u_1 with u_2 , v_1 with v_2 and w_1 with w_2 , respectively.

The existence of structures satisfying condition (3) above is checked by computing an ascending Kleene sequence in a domain of multisets of relation symbols in which each symbol occurs at most three times. Since this domain is finite, the least fixpoint is attained in a finite number of steps, yielding an algorithm that decides the treewidth-boundedness of the set $\llbracket A \rrbracket_\Delta$.

3 Expandable Sets of Inductive Definitions

This section introduces the formal definitions of canonical models and expandable SIDs, thus completing the overview of the proof of Theorem 2.

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 $A(y_1, \dots, y_n)$ in ϕ by a formula $\rho[x_1/y_1, \dots, x_n/y_n]$, where $A(x_1, \dots, x_n) \leftarrow \rho$ is a rule from Δ . A Δ -*unfolding* is a sequence of formulæ such that $\phi_1 \Rightarrow_\Delta \dots \Rightarrow_\Delta \phi_n$. The Δ -unfolding is *complete* iff ϕ_n is a predicate-free formula. The following is a direct consequence of the semantics of SLR:

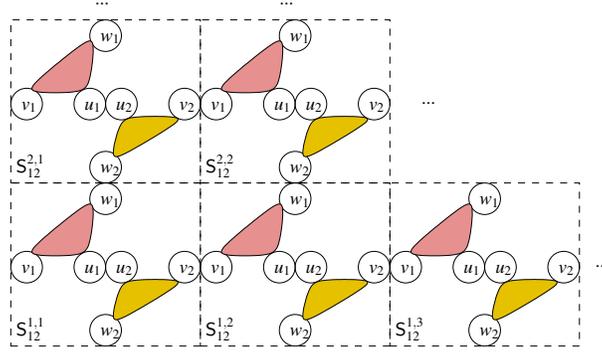


Fig. 3. The principle of grid construction

Proposition 1. Let ϕ be a sentence, Δ a SID and S a structure. Then $S \in \llbracket \phi \rrbracket_{\Delta}$ iff $S \models^s \psi$, for a store s and complete Δ -unfolding $\phi \Rightarrow_{\Delta}^* \exists x_1 \dots \exists x_n \cdot \psi$, where ψ is a qpf formula.

Intuitively, a model of a sentence is *canonical* if it can be defined using a store that matches only those variables that are equated in the result of the unfolding. 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 . A store s is *canonical for ϕ* iff $s(x) = s(y)$ only if $x \approx_{\phi} y$, for all $x, y \in \text{fv}(\phi)$. Moreover, a *rich canonical model* stores information about the disequalities introduced by the unfolding.

Definition 8. Let Δ be a SID and ϕ a sentence. A rich canonical Δ -model of 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 $\phi \Rightarrow_{\Delta}^* \exists x_1 \dots \exists x_n \cdot \psi$, where ψ is qpf, and a store s canonical for ψ , such that $S \models^s \psi$ and $\mathfrak{d}(u, v)$ iff there exist variables $x \in s^{-1}(u)$ and $y \in s^{-1}(v)$ such that $x \neq y$ occurs in ψ . We denote by $\llbracket \phi \rrbracket_{\Delta}^r$ the set of rich canonical Δ -models of ϕ and $\llbracket \phi \rrbracket_{\Delta}^c \stackrel{\text{def}}{=} \{S \mid (S, \mathfrak{d}) \in \llbracket \phi \rrbracket_{\Delta}^r\}$ the set of canonical Δ -models of ϕ . If ϕ is predicate-free, we write $\llbracket \phi \rrbracket^c$ (resp. $\llbracket \phi \rrbracket^r$) instead of $\llbracket \phi \rrbracket_{\Delta}^c$ (resp. $\llbracket \phi \rrbracket_{\Delta}^r$).

In the rest of this section we fix the SID Δ . We simplify the technical development assuming w.l.o.g. that no equalities occur in Δ (Lemma 5).

Definition 9. A formula is *equality-free* iff it contains no equalities nor predicate atoms in which the same variable occurs twice. A rule $A(x_1, \dots, x_n) \leftarrow \phi$ is *equality-free* iff ϕ is equality-free. A SID is *equality-free* iff it consists of equality-free rules.

Lemma 5. Given a SID Δ , one can build an equality-free SID Δ' , such that $\llbracket A \rrbracket_{\Delta} = \llbracket A \rrbracket_{\Delta'}$ for any nullary predicate A .

See proof on page 33.

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

Lemma 6. Given a nullary predicate A , there exists a constant $W \geq 1$, depending only on Δ , such that $\text{tw}(S) \leq W$, for any $S \in \llbracket A \rrbracket_{\Delta}^c$.

See proof on page 33.

Second, any model is obtained via internal fusion of a rich canonical model:

Definition 10. An equivalence relation $\approx \subseteq \mathbf{U} \times \mathbf{U}$ is compatible with a rich canonical model (S, \mathfrak{d}) iff it is compatible with $S = (\mathbf{U}, \sigma)$ and $\mathfrak{d}(u, v)$ only if $u \not\approx v$. We denote by $\widetilde{\text{IF}}(S, \mathfrak{d})$ the set of structures isomorphic to $S_{/\approx}$, where \approx is some equivalence relation compatible with (S, \mathfrak{d}) .

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 stores in Def. 8.

Lemma 7. Let A be a nullary predicate. Then, $\llbracket A \rrbracket_{\Delta} = \widetilde{\text{IF}}(\llbracket A \rrbracket_{\Delta}^c) \subseteq \text{IF}(\llbracket A \rrbracket_{\Delta}^c)$.

See proof on page 33.

A structure is a *substructure* of another if the former is obtained from the latter by removing elements from its support:

Definition 11. Let $S_i = (\mathbf{U}_i, \sigma_i)$ be structures, for $i = 1, 2$. S_1 is included in S_2 iff $\mathbf{U}_1 \subseteq \mathbf{U}_2$ 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 (equivalently, a nullary predicate) are all substructures of the same canonical model of that sentence, that can be, moreover, placed “sufficiently far away” one from another.

Definition 12. A SID Γ is expandable for a nullary predicate B iff for each sequence of pairwise disjoint canonical models $S_1 = (\mathbf{U}_1, \sigma_1), \dots, S_n = (\mathbf{U}_n, \sigma_n) \in \llbracket B \rrbracket_{\Gamma}^c$, there exists a rich canonical model $(S, \mathfrak{d}) \in \llbracket B \rrbracket_{\Gamma}^c$, where $S = (\mathbf{U}, \sigma)$, 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 $\{u_1, \dots, u_{\#r}\} \cap \{v_1, \dots, v_{\#r}\} \neq \emptyset$.

The conditions (2) and (3) of Def. 12 ensure that the external fusion (Def. 4) of these substructures is not hindered by how they are placed inside the larger structure. This definition completes the formalization of the statements of Lemmas 1 and 2 (§5) on which the proof of Theorem 2 rests. We proceed with a proof of Lemma 2.

Proof of Lemma 2 “(1) \Rightarrow (2)” $\text{IF}(\llbracket A \rrbracket_{\Delta}^c) \subseteq \text{IEF}^*(\llbracket A \rrbracket_{\Delta}^c)$ holds trivially, by Def. 4, leading to $\llbracket A \rrbracket_{\Delta} \subseteq \text{IEF}^*(\llbracket A \rrbracket_{\Delta}^c)$, by Lemma 7. “(2) \Rightarrow (3)” Let $S = (\mathbf{U}, \sigma) \in \text{EF}^*(\llbracket A \rrbracket_{\Delta}^c)$ be a structure. It is sufficient to prove that $S \sqsubseteq S'$ for another structure $S' \in \llbracket A \rrbracket_{\Delta}$, because $\text{tw}(S) \leq \text{tw}(S')$, in this case. Then there exist pairwise disjoint structures $S_1 = (\mathbf{U}_1, \sigma_1), \dots, S_n = (\mathbf{U}_n, \sigma_n) \in \llbracket A \rrbracket_{\Delta}^c$ and an equivalence relation $\approx \subseteq (\bigcup_{i=1}^n \mathbf{U}_i) \times (\bigcup_{i=1}^n \mathbf{U}_i)$, 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 Def. 12, there exists a rich canonical model $(S'', \mathfrak{d}) \in \llbracket A \rrbracket_{\Delta}^c$, 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 $S''_{/\approx} \in \widetilde{\text{IF}}(\llbracket A \rrbracket_{\Delta}^c) = \llbracket A \rrbracket_{\Delta}$ by Lemma 7. We conclude by taking $S' = S''_{/\approx}$. \square

4 Encoding Sets of Inductive Definitions by Tree Automata

For technical reasons, the construction of expandable SIDs with an equivalent treewidth boundness problem (Lemma 1) uses a representation of the SID as a tree automaton. This representation allows to distinguish the purely structural aspects, related to the dependencies between rules, from details related to the flow of parameters.

Let \mathbb{A} be a *ranked alphabet*, each symbol $a \in \mathbb{A}$ having an associated integer $\rho(a) \geq 0$, called the *rank* of a . The elements of \mathbb{N}_+^* and 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 13. An $(\mathbb{A}$ -labeled tree) automaton is $\mathcal{A} = (\mathbb{A}, Q, I, \delta)$, where Q is a finite set of states, $I \subseteq Q$ is a set of initial states (if I is a singleton, we denote it by $\iota \in Q$), 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 $\bullet\tau \stackrel{\text{def}}{=} q_0$ and $\tau^\bullet \stackrel{\text{def}}{=} [q_1, \dots, q_\ell]$, i.e., a multiset. 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 $S \subseteq Q$, 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 $\rightsquigarrow \subseteq Q \times Q$ 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 $S \subseteq Q$, 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 $\theta : \text{dom}(t) \rightarrow Q$ 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 $q \xrightarrow{a} () \in \delta$, 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 $\theta(\varepsilon) \in I$. The *language* of \mathcal{A} is $\mathcal{L}(\mathcal{A}) \stackrel{\text{def}}{=} \bigcup_{q \in I} \mathcal{L}_q(\mathcal{A})$, where $\mathcal{L}_q(\mathcal{A}) \stackrel{\text{def}}{=} \{t \mid \mathcal{A} \text{ has a run } \theta \text{ over } t \text{ and } \theta(\varepsilon) = q\}$.

An automaton is *rooted* iff $I = \{\iota\}$ 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 $q \in Q$. Each automaton with non-empty language can be transformed into a trim one with the same language, by a simple marking algorithm.

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

Definition 14. An automaton $\mathcal{A} = (\mathbb{A}, Q, \iota, \delta)$ 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 $S = \{\tau\}$, for some $\tau \in \delta$ such that $\Lambda(\tau) = 1$.

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 $q \in (\delta^1)^\bullet \cap (\delta^\infty)^\bullet$ is called a pivot state. Let $\mathcal{R}_q^\infty(\mathcal{A})$ denote the set of partial runs θ of \mathcal{A} , such that $\theta(\varepsilon) = q$ 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. The transitions from δ^1 , called 1-transitions in the following, are used to move from one linear SCC to another, hence all of them occur exactly once on each accepting run:

Lemma 8. Let $\mathcal{A} = (\mathbb{A}, Q, \mathfrak{1}, \delta)$ be a choice-free automaton, such that $\delta = \delta^1 \uplus \delta^\infty$ (Def. 14) 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 $i \in [1..\ell]$.

See proof on page 34.

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

Lemma 9. Let $\mathcal{A} = (\mathbb{A}, Q, \mathfrak{1}, \delta)$ be a choice-free automaton, where $\delta = \delta^1 \uplus \delta^\infty$ (Def. 14). Then, for any state $q \in (\delta^\infty)^\bullet$ there exists a pivot state $q_0 \in (\delta^1)^\bullet \cap (\delta^\infty)^\bullet$ 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 $\theta' \in \mathcal{R}_{q_0}^\infty(\mathcal{A})$ such that q_0 occurs on the frontier of θ' .

See proof on page 35.

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

Lemma 10. Given an automaton \mathcal{A} , one can build finitely many choice-free automata $\mathcal{A}_1, \dots, \mathcal{A}_n$, such that $\mathcal{L}(\mathcal{A}) = \bigcup_{i=1}^n \mathcal{L}(\mathcal{A}_i)$.

See proof on page 37.

Alphabets of formulæ The construction of expandable SIDs (Lemma 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 15. Let Σ be the set of qpf formulae α of rank $\rho(\alpha) = \ell$, such that:

1. $\text{fv}(\alpha) \subseteq \{x_1^{[e]}, \dots, x_{n_0}^{[e]}\} \cup \{y_1^{[e]}, \dots, y_m^{[e]}\} \cup \bigcup_{i=1}^{\ell} \{x_1^{[i]}, \dots, x_{n_i}^{[i]}\}$, for some $m, n_0, \dots, n_\ell \in \mathbb{N}$; a variable $x_j^{[i]}$ is called a i -variable, for all $i \in \{\varepsilon\} \cup [1..\ell]$,
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 formulae $t(p)^{[p]}$ 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 $\mathcal{A}_{\Delta, A} \stackrel{\text{def}}{=} (\Sigma, Q_\Delta, q_A, \delta_\Delta)$ is defined as follows:

- Q_Δ contains states q_B , where B is a predicate occurring in Δ ; each state has an associated arity $\#q_B \stackrel{\text{def}}{=} \#B$,
- δ_Δ contains a transition $q_{A_0} \xrightarrow{\alpha_p} (q_{A_1}, \dots, q_{A_\ell})$, where α_p is the symbol:

$$\alpha_p \stackrel{\text{def}}{=} \Psi[x_1/x_1^{[e]}, \dots, x_{n_0}/x_{n_0}^{[e]}, y_1/y_1^{[e]}, \dots, y_m/y_m^{[e]}] * \bigstar_{i=1}^{\ell} \bigstar_{j=1}^{n_i} z_{i,j}^{[e]} = x_j^{[i]} \quad (1)$$

of rank $\rho(\alpha_p) = \ell$ that corresponds to the (equality-free) rule $\rho \in \Delta$, where Ψ is qpf:

$$\rho : A_0(x_1, \dots, x_{n_0}) \leftarrow \exists y_1 \dots \exists y_m \cdot \Psi * \bigstar_{i=1}^{\ell} A_i(z_{i,1}, \dots, z_{i,n_i}) \quad (2)$$

The following property of $\mathcal{A}_{\Delta, A}$ is immediate from its definition:

Lemma 11. Let $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta_\Delta$ be a transition of $\mathcal{A}_{\Delta, A}$. Then, for each $i \in [1..\ell]$ and each $j \in [1..\#q_i]$, there exists an ε -variable z , such that $x_j^{[i]} \approx_\alpha z$.

See proof on page 39.

Dually, given an automaton $\mathcal{A} = (\Sigma, Q, \mathbf{1}, \delta)$, the SID $\Delta_{\mathcal{A}}$ consists of the following rules, one for each transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta$:

$$A_{q_0}(x_1, \dots, x_{\#q_0}) \leftarrow \exists y_1 \dots \exists y_m \cdot \alpha[x_1^{[e]}/x_1, \dots, x_{\#q_0}^{[e]}/x_{\#q_0}] * \bigstar_{j=1}^{\ell} A_{q_j}(x_1^{[j]}, \dots, x_{\#q_j}^{[j]}) \quad (3)$$

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

Note that $\Delta_{\mathcal{A}}$ is not equality-free, but can be transformed into an equality-free SID (Lemma 5).

Definition 16. For a Σ -labeled automaton \mathcal{A} , let $\llbracket \mathcal{A} \rrbracket \stackrel{\text{def}}{=} \bigcup_{t \in \mathcal{L}(\mathcal{A})} \llbracket \Theta(t)^{\exists} \rrbracket$ (resp. $\llbracket \mathcal{A} \rrbracket^c \stackrel{\text{def}}{=} \bigcup_{t \in \mathcal{L}(\mathcal{A})} \llbracket \Theta(t)^{\exists} \rrbracket^c$ and $\llbracket \mathcal{A} \rrbracket^r \stackrel{\text{def}}{=} \bigcup_{t \in \mathcal{L}(\mathcal{A})} \llbracket \Theta(t)^{\exists} \rrbracket^r$). The automaton \mathcal{A} is all-satisfiable iff the formula $\Theta(t)$ is satisfiable, for all $t \in \mathcal{L}(\mathcal{A})$.

The relation between SIDs and Σ -labeled automata is stated below:

Lemma 12. (1) Given a SID Δ and a nullary predicate A , one can build an automaton $\mathcal{A}_{\Delta, A}$ such that $\llbracket A \rrbracket_\Delta = \llbracket \mathcal{A}_{\Delta, A} \rrbracket$. (2) Given an automaton $\mathcal{A} = (\Sigma, Q, \mathbf{1}, \delta)$, one can build a SID $\Delta_{\mathcal{A}}$, such that $\llbracket \mathcal{A} \rrbracket = \llbracket A_{\mathbf{1}} \rrbracket_{\Delta_{\mathcal{A}}}$ and $\llbracket \mathcal{A} \rrbracket^r = \llbracket A_{\mathbf{1}} \rrbracket_{\Delta_{\mathcal{A}}}^r$, for a nullary predicate A .

See proof on page 39.

Persistent variables The second ingredient of the construction of expandable SIDs (Lemma 1) are the *persistent* variables introduced by 1-transitions, whose values propagate via equalities throughout each run of the choice-free automaton. We introduce these variables formally using the notion of *profile*:

Definition 17. Let $\mathcal{A} = (\Sigma, Q, \iota, \delta)$ be a choice-free automaton, where $\delta = \delta^1 \uplus \delta^\infty$ (Def. 14). A positional function $\mathfrak{P} : Q \rightarrow \text{pow}(\mathbb{N})$ 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 $s \in \mathfrak{P}_{\mathcal{A}}(q_\varepsilon)$, such that $x_s^{[e]} \approx_\alpha x_r^{[k]}$. 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 $i \in [1..\ell] \cup \{\varepsilon\}$.

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 entry state in every run of ∞ -transitions of \mathcal{A} over t . Note that the profile is computable by a separate finite greatest fixpoint Kleene iteration over sets of SCCs in the automaton interconnected by ∞ -transitions.

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:

Definition 18. A context $\theta_{p \leftarrow q} \in \mathcal{R}_q^\infty(\mathcal{A})$ over a tree t is a q -reset iff (1) $x_j^{[e]} \approx_{\Theta(t)} x_j^{[p]}$, for all $j \in \mathfrak{P}_{\mathcal{A}}(q)$, and (2) $x_j^{[e]} \not\approx_{\Theta(t)} x_k^{[p]}$, 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 a reset path.

Lemma 13. Let $\mathcal{A} = (\Sigma, Q, \iota, \delta)$ be a trim automaton. Then, there exists a q -reset for (1) each pivot state $q \in (\delta^1)^\bullet \cap \bullet(\delta^\infty)$ of \mathcal{A} 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.

See proof on page 41.

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:

Lemma 14. Let \mathcal{A} be a trim 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 $1 \leq i < j \leq n$.

See proof on page 42.

5 A Decomposition into Expandable Sets of Inductive Definitions

This section describes the technical development leading to the proof of Lemma 1. In the rest of this section, let Δ be a given equality-free SID and A be a nullary predicate. The automaton $\mathcal{A}_{\Delta,A}$ recognizes the set of Δ -models of A , by Lemma 12 (1). We shall build from $\mathcal{A}_{\Delta,A}$ finitely many automata $\mathcal{B}_1, \dots, \mathcal{B}_m$, such that $\llbracket \mathcal{A}_{\Delta,A} \rrbracket$ is treewidth-bounded iff $\llbracket \mathcal{B}_i \rrbracket$ is treewidth-bounded, for each $i \in [1..m]$ and, moreover, each SID $\Gamma_i \stackrel{\text{def}}{=} \Delta_{\mathcal{B}_i}$, i.e., obtained from \mathcal{B}_i using Lemma 12 (2) is shown to be expandable for a nullary predicate B , for all $i \in [1..m]$. The construction of $\mathcal{B}_1, \dots, \mathcal{B}_m$ proceeds in five steps, denoted **(I-V)** in what follows.

The automata built in the following will be *simulations* and *refinements* of $\mathcal{A}_{\Delta,A}$:

Definition 19. Let $\mathcal{A} = (\mathbb{A}, Q_{\mathcal{A}}, \iota_{\mathcal{A}}, \delta_{\mathcal{A}})$ and $\mathcal{B} = (\mathbb{A}, Q_{\mathcal{B}}, \iota_{\mathcal{B}}, \delta_{\mathcal{B}})$ be automata. A mapping $h : Q_{\mathcal{A}} \rightarrow Q_{\mathcal{B}}$ is a *simulation* iff (1) $h(\iota_{\mathcal{A}}) = \iota_{\mathcal{B}}$ and (2) $q_0 \xrightarrow{a} (q_1, \dots, q_\ell) \in \delta_{\mathcal{A}}$ only if $h(q_0) \xrightarrow{a} (h(q_1), \dots, h(q_\ell)) \in \delta_{\mathcal{B}}$, for all $q_0, \dots, q_\ell \in Q_{\mathcal{A}}$. A *refinement* h is a simulation such that, moreover (3) $h(q_0) \xrightarrow{a} (q'_1, \dots, q'_\ell) \in \delta_{\mathcal{B}}$ only if there exist states $q_1 \in h^{-1}(q'_1), \dots, q_\ell \in h^{-1}(q'_\ell)$, such that $q_0 \xrightarrow{a} (q_1, \dots, q_\ell) \in \delta_{\mathcal{A}}$, for all $q_0 \in Q_{\mathcal{A}}$ and $q'_1, \dots, q'_\ell \in Q_{\mathcal{B}}$. If a simulation (refinement) $h : Q_{\mathcal{A}} \rightarrow Q_{\mathcal{B}}$ exists then \mathcal{A} *simulates* (refines) \mathcal{B} .

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

Lemma 15. If \mathcal{A} simulates (resp. refines) \mathcal{B} then $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{B})$ (resp. $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{B})$).

See proof on page 43.

We shall also make use of the following relations between qpf formulæ and the upper bounds on the treewidth of their models:

Lemma 16. Let ϕ be a qpf formula, x_1, x_2, \dots, x_k variables and r a relation symbol of arity k , such that $\phi * x_1 \neq x_2$ and $\phi * r(x_1, \dots, x_k)$ are satisfiable. Then, we have:

1. $\text{tw}(\llbracket (\phi * x_1 = x_2)^{\exists} \rrbracket) \leq \text{tw}(\llbracket \phi^{\exists} \rrbracket)$,
2. $\text{tw}(\llbracket \phi^{\exists} \rrbracket) - 1 \leq \text{tw}(\llbracket \phi * x_1 \neq x_2^{\exists} \rrbracket) \leq \text{tw}(\llbracket \phi^{\exists} \rrbracket)$,
3. $\text{tw}(\llbracket \phi^{\exists} \rrbracket) - 1 \leq \text{tw}(\llbracket \phi * r(x_1, \dots, x_k)^{\exists} \rrbracket) \leq \text{tw}(\llbracket \phi^{\exists} \rrbracket) + k$

See proof on page 44.

I. Satisfiability The first step is the construction of an automaton $\mathcal{A}'_{\Delta,A} = (\Sigma, Q'_A, q_A, \delta'_A)$ recognizing the set of trees from the language of $\mathcal{A}_{\Delta,A} = (\Sigma, Q_A, q_A, \delta_A)$ that have, moreover, a satisfiable characteristic formula. This construction uses an idea of Brotherston et al [9], 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 [6].

The states of $\mathcal{A}'_{\Delta,A}$ are *base pairs*, defined below:

Definition 20. A *base pair* (σ^\sharp, π) consists of 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, and a conjunction of disequalities π . A *base pair* is said to be *satisfiable* iff π is satisfiable and the multiplicity of any tuple $\langle x_1, \dots, x_{\#r} \rangle \in \sigma^\sharp(r)$ is one, for all $r \in \mathbb{R}$. Given a set of variables $X \subseteq \mathbb{V}$, let $\text{SatBase}(X)$ denote the set of satisfiable base pairs involving 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 $(\sigma_1^\sharp, \pi_1) \otimes (\sigma_2^\sharp, \pi_2) \stackrel{\text{def}}{=} (\sigma_1^\sharp \cup \sigma_2^\sharp, \pi_1 * \pi_2)$ if $(\sigma_1^\sharp \cup \sigma_2^\sharp, \pi_1 * \pi_2)$ is satisfiable, and undefined, otherwise. Second, the *substitution* $(\sigma^\sharp, \pi)[x_1/y_1, \dots, x_n/y_n]$ replaces simultaneously each occurrence of x_j by y_j in σ^\sharp and π , for all $j \in [1..n]$. Third, given a set $X \subseteq \mathbb{V}$ of variables, the *projection* is $(\sigma^\sharp, \pi) \downarrow_X \stackrel{\text{def}}{=} (\lambda r. \{ \langle x_1, \dots, x_n \rangle \in \sigma^\sharp(r) \mid x_1, \dots, x_n \in X \}, \pi \downarrow_X)$ where, for a formula ϕ , the operation $\phi \downarrow_X$ removes from ϕ all atoms involving variables not from X . Finally, for a qpf formula $\phi = \psi * \pi$, where ψ is a separated conjunction of relation atoms and π is a pure formula, we define:

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

We define the automaton $\mathcal{A}_{\Delta, A}^I = (\Sigma, Q_\Delta^I, \iota_A^I, \delta_\Delta^I)$, where:

- $Q_\Delta^I \stackrel{\text{def}}{=} \{(q, (\sigma^\sharp, \pi)) \mid q \in Q_\Delta, (\sigma^\sharp, \pi) \in \text{SatBase}(x_1^{[e]}, \dots, x_{\#q}^{[e]})\}$,
- $\iota_A^I \stackrel{\text{def}}{=} \{(\iota_A, (\sigma_0^\sharp, \text{emp}))\}$, where σ_0^\sharp interprets each relation symbol as the empty set; recall that, since we assumed $\# \iota_A = 0$, there are no tuples associated with ι_A ,
- δ_Δ^I is the set of transitions $(q_0, (\sigma_0^\sharp, \pi_0)) \xrightarrow{\alpha} \langle (q_1, (\sigma_1^\sharp, \pi_1)), \dots, (q_\ell, (\sigma_\ell^\sharp, \pi_\ell)) \rangle$, such that $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta_\Delta$ and the following condition holds:

$$(\sigma_0^\sharp, \pi_0) = \left(\text{Base}(\alpha) \otimes \bigotimes_{i=1}^{\ell} (\sigma_i^\sharp, \pi_i)[x_1/x_1^{[i]}, \dots, x_{\#q_i}/x_{\#q_i}^{[i]}] \right) \downarrow_{\{x_1^{[e]}, \dots, x_{\#q_0}^{[e]}\}}$$

The formal properties of the automaton $\mathcal{A}_{\Delta, A}^I$ are stated and proved below:

Lemma 17. (1) $\mathcal{A}_{\Delta, A}^I$ is all-satisfiable. (2) $\llbracket \mathcal{A}_{\Delta, A} \rrbracket = \llbracket \mathcal{A}_{\Delta, A}^I \rrbracket$.

See proof on page 45.

II. Removing relation atoms from 1-transitions Without loss of generality, we assume that $\mathcal{A}_{\Delta, A}^I$ is choice-free and let $\delta_\Delta = \delta_\Delta^1 \uplus \delta_\Delta^\infty$ be the partition of the transitions of $\mathcal{A}_{\Delta, A}^I$ (Def. 14). If this is not the case, we decompose $\mathcal{A}_{\Delta, A}^I$ into finitely many choice-free automata whose union has the same language as $\mathcal{A}_{\Delta, A}^I$ (Lemma 10).

This step replaces each symbol α that labels a 1-transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell)$ of $\mathcal{A}_{\Delta, A}^I$ with the symbol obtained by removing all relation atoms from α . Let $\mathcal{A}_{\Delta, A}^{II}$ be the outcome of this transformation, whose properties are stated and proved below:

Lemma 18. (1) $\mathcal{A}_{\Delta, A}^{II}$ is all-satisfiable. (2) $\llbracket \mathcal{A}_{\Delta, A}^{II} \rrbracket$ is treewidth-bounded iff $\llbracket \mathcal{A}_{\Delta, A}^I \rrbracket$ is treewidth-bounded.

See proof on page 47. At this point, the labels of the 1-transitions of $\mathcal{A}_{\Delta, A}^{II}$ consist of (dis-)equalities only. Moreover, since $\mathcal{A}_{\Delta, A}^{II}$ is obtained from the choice-free automaton $\mathcal{A}_{\Delta, A}^I$ by a structure-preserving re-labeling of transitions, it is choice-free as well.

III. Removing equalities involving non-persistent variables We modify the label of each 1-transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell)$ of $\mathcal{A}_{\Delta, A}^{II}$ in two successive steps:

1. for each non-persistent ε -variable $x_j^{[e]}$ (Def. 17), for $j \in [1.. \#q_0]$, that occurs in α in some equality with a persistent i -variable, replace α with $\alpha[x_j^{[e]}/y_k^{[e]}] * x_j^{[e]} = y_k^{[e]}$, where $y_k^{[e]} \notin \text{fv}(\alpha)$ is a fresh variable,

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

The idea is to remove the equalities that would be lost when adding resets before and after every 1-transitions that is, forget equalities involving non-persistent variables (2) while keeping equalities between persistent ones (1). The result is the choice-free automaton $\mathcal{A}_{\Delta, A}^{III}$, whose properties are stated and proved below:

Lemma 19. *Let $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell)$ be a 1-transition of $\mathcal{A}_{\Delta, A}^{III}$. Then, for each $i \in [1..\ell]$ and each $j \in \mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q_i)$, there exists a ε -variable z , such that $x_j^{[i]} \approx_\alpha z$.*

See proof on page 48.

Lemma 20. *Let ϕ and ψ be qpf formulae, such that $\phi * \psi$ is satisfiable and $x \not\approx_\phi y$, for all $x, y \in \text{fv}(\phi) \cap \text{fv}(\psi)$. Let $\psi_{eq} = \star \{x = y \mid x, y \in \text{fv}(\phi) \cap \text{fv}(\psi), x \approx_\psi y\}$. Then, $\text{tw}(\llbracket (\phi * \psi_{eq})^\exists \rrbracket) \leq \text{tw}(\llbracket (\phi * \psi)^\exists \rrbracket) + \text{card}(\text{fv}(\phi) \cap \text{fv}(\psi))$.*

See proof on page 48.

Lemma 21. (1) $\mathcal{A}_{\Delta, A}^{III}$ is all-satisfiable. (2) $\llbracket \mathcal{A}_{\Delta, A}^{III} \rrbracket$ is treewidth-bounded iff $\llbracket \mathcal{A}_{\Delta, A}^{II} \rrbracket$ is treewidth-bounded.

See proof on page 49. Again, $\mathcal{A}_{\Delta, A}^{III}$ is choice-free, because it is obtained by a structure-preserving re-labeling of the choice-free automaton $\mathcal{A}_{\Delta, A}^{II}$.

IV. Removing persistent variables We build from $\mathcal{A}_{\Delta, A}^{III} = (\Sigma, Q_\Delta^I, \iota_A^I, \delta_\Delta^{III})$ an automaton $\mathcal{A}_{\Delta, A}^{IV}$ having no persistent variables at all. We recall that, by Lemma 8, 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 $p \in \text{dom}(t)$. Using a renaming, if necessary, we can assume that the ε -variables $y_i^{[e]}$, i.e., not associated with the states of the transition (Def. 15), have distinct names between the 1-transitions of $\mathcal{A}_{\Delta, A}^{III}$ and let $\mathcal{Y} \stackrel{\text{def}}{=} \{y_1^{[e]}, \dots, y_{\mathcal{M}}^{[e]}\}$ denote their set.

First, we decorate each state q of $\mathcal{A}_{\Delta, A}^{III}$ with an injective partial mapping $a : [1..\#q] \rightarrow [1..\mathcal{M}]$ that refer each parameter from the profile of $\mathcal{A}_{\Delta, A}^{III}$ at q to the unique persistent variable in \mathcal{Y} that must equal that parameter in each run. Second, we use this information to replace each occurrence of a relation atom $r(z_1^{[e]}, \dots, z_{\#r}^{[e]})$ on the label of a decorated transition $(q_0, a_0) \xrightarrow{\alpha} ((q_1, a_1), \dots, (q_\ell, a_\ell))$ with a fresh relation atom $r_g(z_{i_1}^{[e]}, \dots, z_{i_k}^{[e]})$, where $g : [1..\#r] \rightarrow \text{dom}(a_0)$ is the partial mapping of variables to persistent variables and $\{i_1, \dots, i_k\}$ are the remaining non-persistent variables.

Formally, we define $\mathcal{A}_{\Delta, A}^{IV} = (\Sigma, Q_\Delta^{IV}, \iota_A^{IV}, \delta_\Delta^{IV})$, where:

- $Q_\Delta^{IV} \stackrel{\text{def}}{=} \{(q, a) \mid q \in Q_\Delta^I, a : \mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q) \rightarrow [1..\mathcal{M}] \text{ is an injective mapping}\}$, where $\#(q, a) \stackrel{\text{def}}{=} \#q - \text{card}(\mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q))$; we recall that $\mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}$ is the profile of $\mathcal{A}_{\Delta, A}^{III}$ (Def. 17).
- $(q_0, a_0) \xrightarrow{\alpha} ((q_1, a_1), \dots, (q_\ell, a_\ell)) \in \delta_\Delta^{IV}$ iff either one of the following holds:
 - $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in (\delta_\Delta^{III})^1$ and, for all $k \in [1..\ell]$ and $i \in \mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q_k)$, we have:

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

By Lemma 19, one of the cases above must hold, hence a_k is well defined.

- $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in (\delta_\Delta^{III})^\infty$ and, for all $k \in [1..\ell]$, we have $a_k \stackrel{\text{def}}{=} a_0 \circ b_k$, where $b_k(i) = j \iff x_i^{[k]} \approx_\alpha x_j^{[e]}$, for all $i \in \mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q_k)$ and $j \in \mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q_0)$. Note that, by Def. 17, b_k is well defined.

The goal of this transformation is to remove, from the transition label α , the persistent variables associated to one of the states q_0, \dots, q_ℓ . In order to preserve the naming conventions from Def. 15, we rename the remaining (non-persistent) variables using an injective mapping $\eta : \text{fv}(\alpha) \rightarrow \text{fv}(\alpha)$, such that:

- $\eta(\{x_i^{[e]} \mid i \in \mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q_0)\}) = \{x_1^{[e]}, \dots, x_{k_0}^{[e]}\}$, $k_0 \stackrel{\text{def}}{=} n_0 - \text{card}(\mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q_0))$,
- $\eta(\{x_i^{[j]} \mid i \in \mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q_j)\}) = \{x_1^{[j]}, \dots, x_{k_j}^{[j]}\}$, $k_j \stackrel{\text{def}}{=} n_j - \text{card}(\mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q_j))$, for $j \in [1..\ell]$,
- $\eta(y_i^{[e]}) = y_i^{[e]}$, for $i \in [1..m]$,

where m, n_0, \dots, n_ℓ are as in Def. 15. Note that, by the definition (1) of the transition labels of $\mathcal{A}_{\Delta, A}$, each relation atom from α is of the form $r(z_1^{[e]}, \dots, z_{\#r}^{[e]})$ (i.e., these atoms are not changed by the transformations **(I-III)**, with the exception of **(II)**, which removes relation atoms from the 1-transitions).

We distinguish two cases. If α is the label of a 1-transition of $\mathcal{A}_{\Delta, A}^{III}$, we define $\bar{\alpha} \stackrel{\text{def}}{=} \text{emp}$. Otherwise (α labels an ∞ -transition), $\bar{\alpha}$ is obtained by replacing each relation atom $r(z_1^{[e]}, \dots, z_{\#r}^{[e]})$ from α with a relation atom $r_g(\eta(z_{i_1}^{[e]}), \dots, \eta(z_{i_k}^{[e]}))$, where:

- r_g is a new relation symbol of arity k and $g : [1..\#r] \rightarrow [1..\mathcal{M}] \cup \{\perp\}$ is:

$$g(i) \stackrel{\text{def}}{=} \begin{cases} a_0(j), & \text{if } z_i^{[e]} \text{ and } x_j^{[e]} \text{ are the same variable, such that } j \in \mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q_0) \\ \perp & , \text{ otherwise} \end{cases}$$

- $\{i_1, \dots, i_k\} \stackrel{\text{def}}{=} [1..\#r] \setminus \mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}}(q_0)$.

The properties of $\mathcal{A}_{\Delta, A}^{IV}$ are stated and proved below:

Lemma 22. (1) $\mathcal{A}_{\Delta, A}^{IV}$ is all-satisfiable. (2) $\llbracket \mathcal{A}_{\Delta, A}^{IV} \rrbracket$ is treewidth-bounded iff $\llbracket \mathcal{A}_{\Delta, A}^{III} \rrbracket$ is treewidth-bounded.

See proof on page 51.

The construction of $\mathcal{A}_{\Delta, A}^{IV}$ incurs a slight complication. In fact, decorating the states of $\mathcal{A}_{\Delta, A}^{III}$ with assignments (tracking the values of persistent variables) may cause several transitions to occur between different linear SCCs. These transitions originate from the same 1-transition of $\mathcal{A}_{\Delta, A}^{III}$ and differ only in the decoration of their states. We fix this problem by splitting $\mathcal{A}_{\Delta, A}^{IV}$ into finitely many choice-free automata $\bar{B}_1, \dots, \bar{B}_m$, one for each such choice, such that $\mathcal{L}(\mathcal{A}_{\Delta, A}^{IV}) = \bigcup_{i=1}^m \mathcal{L}(\bar{B}_i)$ (each \bar{B}_i has no persistent variables).

Formally, we define $\bar{B}_i = (\Sigma, Q_{\Delta}^{IV}, (q_A, \theta), \bar{\delta}_i)$, where $\bar{\delta}_i \stackrel{\text{def}}{=} (\bar{\delta}_i)^1 \uplus (\bar{\delta})^\infty$ is the following partition of transitions:

- each $(\bar{\delta}_i)^1$ contains exactly one transition $(q_0, a_0) \xrightarrow{\bar{\alpha}} ((q_1, a_1), \dots, (q_\ell, a_\ell)) \in \delta_\Delta^{IV}$, for each transition $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta_\Delta^{III}$,
- $(\bar{\delta})^\infty \stackrel{\text{def}}{=} \{(q_0, a_0) \xrightarrow{\bar{\alpha}} ((q_1, a_1), \dots, (q_\ell, a_\ell)) \in \delta_\Delta^{IV} \mid q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in (\delta_\Delta^{III})^\infty\}$; note that the set $(\bar{\delta})^\infty$ is the same in all \bar{B}_i .

We prove that $\overline{B}_1, \dots, \overline{B}_m$ is indeed a choice-free decomposition of $\mathcal{A}_{\Delta, A}^{IV}$. The proof relies on a stronger notion of automata refinement:

Definition 21. An automaton $\mathcal{A} = (\mathbb{A}, Q_{\mathcal{A}}, \iota_{\mathcal{A}}, \delta_{\mathcal{A}})$ is a strong refinement of $\mathcal{B} = (\mathbb{A}, Q_{\mathcal{B}}, \iota_{\mathcal{B}}, \delta_{\mathcal{B}})$ iff there exists a refinement $h : Q_{\mathcal{A}} \rightarrow Q_{\mathcal{B}}$ such that the following hold:

1. $h^{-1}(S)$ is an SCC of \mathcal{A} , for each SCC S of \mathcal{B} .
 2. for each SCC S of \mathcal{B} and each transition $q_0 \xrightarrow{a} (q_1, \dots, q_{\ell}) \in S^{\bullet}$ there exists exactly one transition $q'_0 \xrightarrow{a} (q'_1, \dots, q'_{\ell}) \in \delta_{\mathcal{A}}$, such that $q'_i \in h^{-1}(q_i)$, for all $i \in [0..{\ell}]$.
- If a strong refinement $h : Q_{\mathcal{A}} \rightarrow Q_{\mathcal{B}}$ exists then \mathcal{A} strongly refines \mathcal{B} .

Lemma 23. If \mathcal{A} strongly refines \mathcal{B} and \mathcal{B} is choice-free, then \mathcal{A} is choice-free.

See proof on page 52.

Lemma 24. Each $\overline{B}_1, \dots, \overline{B}_m$ is all-satisfiable, choice-free and $\mathcal{L}(\mathcal{A}_{\Delta, A}^{IV}) = \bigcup_{i=1}^m \mathcal{L}(\overline{B}_i)$.

See proof on page 54.

V. Wrapping 1-transitions into partial runs of ∞ -transitions At this point, we have finitely many all-satisfiable choice-free automata $\overline{B}_1, \dots, \overline{B}_m$ without persistent variables. In order to obtain expandable SIDs from these automata, using Lemma 12 (2), any sequence of accepting runs of \overline{B}_i must be embedded in an accepting run of the same automaton, for all $i \in [1..m]$. Since all 1-transitions of \overline{B}_i must occur on any accepting run (Lemma 8), we need to “wrap” the labels of 1-transitions into characteristic formulæ of trees recognized by partial runs consisting of ∞ -transitions only. This will enable the use of Lemma 14 to embed several runs consisting of ∞ -transitions into one accepting run. The outcome of this transformation is denoted \mathcal{B}_i , for all $i \in [1..m]$.

Let \overline{B} be any of $\overline{B}_1, \dots, \overline{B}_m$. For a Σ -labeled tree t , two positions p and s , such that only $p \in \text{dom}(t)$ (i.e., nothing is required about s), and a sequence of variables x_1, \dots, x_k , we define the formula:

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

The automaton \mathcal{B} is obtained from \overline{B} by replacing the label of each 1-transition $q_0 \xrightarrow{\text{emp}} (q_1, \dots, q_{\ell})$ of \overline{B} with the following formula, for some trees t_i corresponding to some reset $\theta_{p_i \leftarrow q_i}^i$ of \overline{B} , for $i \in [0..{\ell}]$:

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

Note that the existence of such resets is guaranteed by Lemma 13.

Lemma 25. (1) \mathcal{B} is all-satisfiable. (2) $\llbracket \mathcal{B} \rrbracket$ is treewidth-bounded iff $\llbracket \overline{B} \rrbracket$ is treewidth-bounded.

See proof on page 54.

A view for an automaton \mathcal{A} is a tuple $\langle \theta, t, \mathfrak{s}, \mathfrak{S} \rangle$, such that $\theta \in \mathcal{R}_q^{\infty}(\mathcal{A})$ is a partial run over the Σ -labeled tree t , \mathfrak{s} is a canonical store for $\Theta(t)$ and \mathfrak{S} is a structure, such that $\mathfrak{S} \models^{\mathfrak{s}} \Theta(t)$. The key property of this transformation is given by the following lemma:

Lemma 26. For each structure $S = (U, \sigma) \in \llbracket \mathcal{B} \rrbracket^c$, there exist pairwise disjoint structures S_1, \dots, S_n and 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} such that S'_1, \dots, S'_n are pairwise disjoint, $S = S_1 \bullet \dots \bullet S_n$ and, for all $i \in [1..n]$: (1) $S_i \sqsubseteq S'_i$, (2) $\text{supp}(\sigma) \cap \mathfrak{s}_i(\{x_1^{[e]}, \dots, x_{\#\theta_i(\epsilon)}^{[e]}\}) = \emptyset$, and $\text{supp}(\sigma) \cap \mathfrak{s}_i(\{x_1^{[p_i]}, \dots, x_{\#q_i}^{[p_i]}\}) = \emptyset$ if θ_i is a context $\theta_{p_i \leftarrow q_i}$.

See proof on page 55.

Dually, the following lemma “gathers” pairwise disjoint structures into a single rich canonical model, that meets the conditions of expandability (Def. 12):

Lemma 27. Given pairwise disjoint structures S_1, \dots, S_n and 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} , satisfying the conditions (1) and (2) of Lemma 26, there exists a rich canonical model $(S, \mathfrak{d}) \in \llbracket \mathcal{B} \rrbracket^r$, such that the conditions (1), (2) and (3) from Def. 12 hold for S_1, \dots, S_n and (S, \mathfrak{d}) .

See proof on page 56.

Proof of Lemma 1 For each $i \in [1..m]$, let Γ_i be the SID corresponding to \mathcal{B}_i , as given by Lemma 12 (2). We assume w.l.o.g. that the automata \mathcal{B}_i have the same initial state and B is the predicate corresponding to it. First, we prove that $\llbracket A \rrbracket_\Delta$ is treewidth-bounded iff $\llbracket B \rrbracket_{\Gamma_i}$ is treewidth-bounded, for all $i \in [1..m]$. By Lemmas 12 (2) and 17 (2), we have $\llbracket A \rrbracket_\Delta = \llbracket \mathcal{A}_{\Delta, A} \rrbracket = \llbracket \mathcal{A}'_{\Delta, A} \rrbracket$. By Lemmas 21 (2) and 22 (2), we have that $\llbracket A \rrbracket_\Delta$ is treewidth-bounded iff $\llbracket \mathcal{A}^{IV}_{\Delta, A} \rrbracket$ is treewidth-bounded. By Lemmas 24 and 25 (2), $\llbracket A \rrbracket_\Delta$ is treewidth-bounded iff \mathcal{B}_i is treewidth-bounded, for all $i \in [1..m]$. By Lemma 12 (2), $\llbracket A \rrbracket_\Delta$ is treewidth-bounded iff $\llbracket B \rrbracket_{\Gamma_i}$ is treewidth-bounded, for all $i \in [1..m]$.

Let Γ be any of $\Gamma_1, \dots, \Gamma_m$ and we prove that it is expandable for B . Let \mathcal{B} be its corresponding automaton, such that $\llbracket B \rrbracket_\Gamma^c = \llbracket \mathcal{B} \rrbracket^c$, by Lemma 12 (2). Let $S_1 = (U_1, \sigma_1), \dots, S_n = (U_n, \sigma_n) \in \llbracket \mathcal{B} \rrbracket^c$ be pairwise disjoint structures. By Lemma 26, for each $i \in [1..n]$ there exist a decomposition $S_i = S_{i,1} \bullet \dots \bullet S_{i,m_i}$ and views $\langle \theta_{i,j}, t_{i,j}, \mathfrak{s}_{i,j}, S_{i,j} \rangle$, for $j \in [1..m_i]$, such that conditions (1) and (2) hold. We assume w.l.o.g. that $S_{1,1}, \dots, S_{n,m_n}$ are pairwise disjoint, hence their composition is defined, hence $S_1 \bullet \dots \bullet S_n = S_{1,1} \bullet \dots \bullet S_{n,m_n}$. By Lemma 27, there exists a rich canonical model $(S, \mathfrak{d}) \in \llbracket \mathcal{B} \rrbracket^r = \llbracket B \rrbracket_\Gamma^r$ (by Lemma 12 (2)) that meets the conditions of Def. 12. \square

6 External Fusion

This section gives the definitions and results needed for the proofs of Lemmas 3 and 4. We start by defining connected structures and maximal connected substructures:

Definition 22. A path from u to v in a structure $S = (U, \sigma)$ is a finite sequence of tuples $\{\langle u_{i,1}, \dots, u_{i,n_i} \rangle \in \sigma(r_i)\}_{i=1}^k$ such that $r_1, \dots, r_k \in \mathbb{R}$, $u \in \{u_{1,1}, \dots, u_{1,n_1}\}$, $v \in \{u_{k,1}, \dots, u_{k,n_k}\}$ 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 23. A structure S_1 is a maximal connected substructure of another structure S_2 , denoted $S_1 \sqsubseteq^{mc} S_2$, iff (i) $S_1 \sqsubseteq S_2$, (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 $\text{split}(\mathcal{S}) \stackrel{\text{def}}{=} \cup_{S \in \mathcal{S}} \text{split}(S)$.

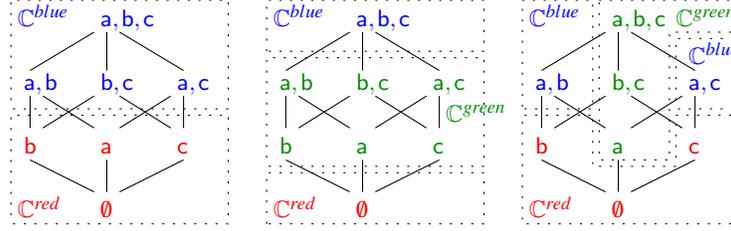


Fig. 4. Examples of RGB color schemes

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 28. *Given a set \mathcal{S} of structures, we have (1) $\text{split}(\text{EF}^*(\mathcal{S})) = \text{EF}^*(\text{split}(\mathcal{S}))$, and (2) $\text{split}(\text{IEF}^*(\mathcal{S})) = \text{IEF}^*(\text{split}(\mathcal{S}))$.*

See proof on page 58.

For a 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 24. *The coloring of a structure $S = (U, \sigma)$ is the mapping $C_S : U \rightarrow \mathbb{C}$ defined as $C_S(u) \stackrel{\text{def}}{=} \{r \in \mathbb{R} \mid \langle u, \dots, u \rangle \in \sigma(r)\}$.*

Intuitively, a fusion operation can only join pairs of elements with disjoint colors. Moreover, we define an abstraction of structures as finite multisets of colors:

Definition 25. *The multiset color abstraction $S^\sharp \in \text{mpow}(\mathbb{C})$ of a structure $S = (U, \sigma)$ is $S^\sharp \stackrel{\text{def}}{=} [C_S(u) \mid u \in \text{supp}(\sigma)]$. For an integer $k \geq 0$, the k -multiset color abstraction $S^{*k} \subseteq \text{mpow}(\mathbb{C})$ is $S^{*k} \stackrel{\text{def}}{=} \{M \subseteq S^\sharp \mid \text{card}(M) \leq k\}$. These abstractions are lifted to sets \mathcal{S} of structures as usual $\mathcal{S}^\sharp \stackrel{\text{def}}{=} \{S^\sharp \mid S \in \mathcal{S}\}$ and $\mathcal{S}^{*k} \stackrel{\text{def}}{=} \bigcup_{S \in \mathcal{S}} S^{*k}$.*

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

Definition 26. *A partition $(\mathbb{C}^{\text{red}}, \mathbb{C}^{\text{green}}, \mathbb{C}^{\text{blue}})$ of \mathbb{C} is an RGB-color scheme iff:*

- (1) $C_1 \cap C_2 \neq \emptyset$, for all $C_1, C_2 \in \mathbb{C}^{\text{blue}}$,
- (2) $C_1 \cap C_2 \neq \emptyset$, for all $C_1 \in \mathbb{C}^{\text{green}}$ and $C_2 \in \mathbb{C}^{\text{blue}}$,
- (3) for all $C_1 \in \mathbb{C}^{\text{red}}$ there exists $C_2 \in \mathbb{C}^{\text{blue}}$ such that $C_1 \cap C_2 = \emptyset$.

Note that an RGB-color scheme is fully specified by the set \mathbb{C}^{blue} ; any other color is unambiguously placed within \mathbb{C}^{red} or $\mathbb{C}^{\text{green}}$, depending on its disjointness from some color in \mathbb{C}^{blue} . Fig. 4 illustrates several RGB-color schemes for a relational signature $\mathbb{R} = \{a, b, c\}$. 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} .

Lemma 29. *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 $C_{S_1}(u_1) \cap C_{S_2}(u_2) = \emptyset$. Then, the equivalence relation $\{(u_1, u_2)\}^\# \subseteq (U_1 \cup U_2) \times (U_1 \cup U_2)$ is compatible with $S_1 \bullet S_2$.*

See proof on page 60.

Definition 27. A set S of structures conforms to $(\mathbb{C}^{red}, \mathbb{C}^{green}, \mathbb{C}^{blue})$ if and only if:

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

If S conforms to $(\mathbb{C}^{red}, \mathbb{C}^{green}, \mathbb{C}^{blue})$, any structure $S \in S$ is of type either:

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

Conformance to an RGB color scheme is key to bounding the treewidth of structures obtained by external fusion of a treewidth-bounded set of connected structures:

Lemma 30. Let S be a treewidth-bounded set of connected structures conforming to an RGB color scheme. Then, for any structure $S \in \text{EF}^*(S)$, the following hold:

1. S is connected and of type either R, G or B,
2. $\text{tw}(S) \leq \begin{cases} \text{tw}(S) & , \text{ if } S \text{ of type R} \\ \max(\text{tw}(S) + 2 \cdot \text{card}(\mathbb{C}^{green}), 3 \cdot \text{card}(\mathbb{C}^{green})) & , \text{ if } S \text{ of type G} \\ \max(\text{tw}(S) + 2 \cdot \text{card}(\mathbb{C}^{green}), 3 \cdot \text{card}(\mathbb{C}^{green}), \text{tw}(S) + 1) & , \text{ if } S \text{ is of type B} \end{cases}$

See proof on page 60.

Lemma 31. Let S be a treewidth-bounded set of connected structures, conforming to an RGB color scheme. Then, $\text{tw}(\text{EF}^*(S))$ and $\text{tw}(\text{IF}(\text{EF}^*(S)))$ are both bounded by $\max(\text{tw}(S) + 2 \cdot \text{card}(\mathbb{C}^{green}), 3 \cdot \text{card}(\mathbb{C}^{green}), \text{tw}(S) + 1)$.

See proof on page 63.

The core of the treewidth boundedness algorithm is a decidable equivalent condition for the treewidth boundedness of a set obtained by external fusion of a treewidth-bounded set of connected structures. Essentially, a set generated by external fusion is treewidth-bounded iff there is no way of connecting six elements u_i, v_i and w_i , labeled with C_i , for $i = 1, 2$, respectively, where $C_1 \cap C_2 \neq \emptyset$ (condition (2) of Lemma 32).

Lemma 32. The following are equivalent, for any treewidth-bounded set S of connected structures:

1. $\text{EF}^*(S)$ is treewidth bounded,
2. $[C_1, C_1, C_1], [C_2, C_2, C_2] \in (\text{EF}^*(\text{split}(S)))^{\sharp 3}$ implies $C_1 \cap C_2 \neq \emptyset$ for all C_1, C_2 ,
3. $\text{split}(S)$ conforms to some RGB color scheme.

See proof on page 63. The algorithm checks if the set $(\text{EF}^*(\text{split}(S)))^{\sharp 3}$ meets condition (2) above. The check is effective, provided that this set can be built in finitely many steps.

In order to decide whether $\llbracket A \rrbracket_\Delta$ is treewidth-bounded, for a given SID Δ and nullary predicate A , we must compute the set $(\text{EF}^*(\text{split}(\llbracket A \rrbracket_\Delta)))^{\sharp k}$ and check condition (2) of Lemma 32. To this end, we first define an abstract operator $\text{ef}_1^{\sharp k^*}(\cdot)$ having the property $(\text{EF}^*(\text{split}(S)))^{\sharp k} = \text{ef}_1^{\sharp k^*}(S^{\sharp k})$, for any set S of structures. Since the set $S^{\sharp k}$ is finite, $\text{ef}_1^{\sharp k^*}(S^{\sharp k})$ is computable by a finite least fixpoint iteration. Second, we use another finite least fixpoint iteration to compute the k -multiset color abstraction of a set of canonical models $(\llbracket B \rrbracket_\Gamma^c)^{\sharp k}$, for a given SID Γ and nullary predicate B . Finally, we define the set of maximally connected components $\text{split}(\llbracket A \rrbracket_\Delta^c)$ as the set of canonical models $\llbracket B \rrbracket_\Gamma^c$, where the SID Γ and the nullary predicate B are built from Δ and A , respectively. Consequently, the set $(\text{EF}^*(\text{split}(\llbracket A \rrbracket_\Delta)))^{\sharp k} = \text{ef}_1^{\sharp k^*}((\llbracket B \rrbracket_\Gamma^c)^{\sharp k})$ can be computed effectively.

6.1 Computing k -Multiset Color Abstraction for External Fusion

We describe now the effective construction of a k -multiset abstraction $(\text{EF}^*(S))^{\sharp k}$ from the abstraction $S^{\sharp k}$ of a set S of structures, for a given integer $k \geq 1$. First, as we are interested only on k -multisets color abstractions, we can restrict external fusion to bipartite equivalence relations generated by a single pair, without loss of precision.

Definition 28. *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 (Def. 4) induced by equivalence relations $\{(u_1, u_2)\}^=$, where $u_i \in \text{supp}(\sigma_i)$, for $i = 1, 2$. We denote by $\text{EF}_1(S_1, S_2)$ the set of structures obtained by single-pair external fusion of S_1 and S_2 . For a set of structures S , we denote by $\text{EF}_1^*(S)$ the closure of 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 33. $(\text{EF}^*(S))^{\sharp k} = (\text{EF}_1^*(S))^{\sharp k}$ for any set S of structures and integer $k \geq 1$.

See proof on page 64. Second, the closure $(\text{EF}_1^*(S))^{\sharp k}$ 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 29. *The single-pair multiset fusion is defined below, for $M_1, M_2 \in \text{mpow}(\mathbb{C})$:*

$$\text{ef}_1^{\sharp(M_1, M_2)} \stackrel{\text{def}}{=} \{M \in \text{mpow}(\mathbb{C}) \mid \exists C_1 \in M_1. \exists C_2 \in M_2. C_1 \cap C_2 = \emptyset, \\ M = [C_1 \cup C_2] \cup \bigcup_{i=1,2} (M_i \setminus [C_i])\}$$

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 $\text{card}(M_2) \leq k$:

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

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

Lemma 34. $(\text{EF}_1^*(S))^{\sharp k} = \text{ef}_1^{\sharp k^*}(S^{\sharp k})$ for any set S of structures, for any integer $k \geq 1$.

See proof on page 65.

6.2 Computing k -Multiset Color Abstraction for SID Canonical Models

We shall apply Lemma 34 to compute the k -multiset color abstraction of a set $[[A]]_{\Delta}^{\mathbb{C}}$ of canonical models, for a given nullary predicate A and SID Δ . To this end, we must first compute its k -multiset color abstraction $([[A]]_{\Delta}^{\mathbb{C}})^{\sharp k}$. This is done by a least fixpoint computation in an abstract domain, defined directly from the rules in the SID, that tracks 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 set of variables $X \subseteq \mathbb{V}$, a mapping $c : X \rightarrow \mathbb{C}$, and a multiset $M \subseteq \text{mpow}(\mathbb{C})$, such that $\text{card}(M) \leq k$. Note that there exists $\text{card}(\mathbb{C})^{\text{card}(X)+k}$ distinct color triples, for given X and k . We consider the following operations on color triples, lifted to sets as usual:

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

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

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 B , we denote by $\llbracket B \rrbracket_{\Delta}^{\#k}$ the least sets of k -bounded color triples over the variables $x_1, \dots, x_{\#B}$, the satisfies the following constraints:

$$\llbracket B_0 \rrbracket_{\Delta}^{\#k} \supseteq (\gamma(\psi) \bullet^{\#k} \bullet^{\#k} \llbracket B_i \rrbracket_{\Delta}^{\#k} [x_1/z_{i,1}, \dots, x_{\#B_i}/z_{i,\#B_i}]) \downarrow_{\{x_1, \dots, x_{\#B_0}\}}^{\#k} \quad (4)$$

one for each rule of Δ of the form:

$$B_0(x_1, \dots, x_{\#B_0}) \leftarrow \exists y_1 \dots \exists y_m . \psi * \bigstar_{i=1}^{\ell} B_i(z_{i,1}, \dots, z_{i,\#B_i}) \quad (5)$$

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 35. ($\llbracket [A]_{\Delta}^c \rrbracket^{\#k} = \pi_3(\llbracket A \rrbracket_{\Delta}^{\#k})$), for any $k \geq 1$, any equality-free SID Δ and nullary predicate A , such that each complete Δ -unfolding of A yields a satisfiable formula.

See proof on page 66.

As a remark, if the given SID is not equality-free, one can build an equivalent equality-free SID (Lemma 5). Moreover, if there are complete unfoldings that yield unsatisfiable formulæ, the SID can be modified to remove such unfoldings (Lemmas 12 and 17).

6.3 Computing Maximally Connected Structures for SID Canonical Models

The final step required to check the condition (2) of Lemma 32 is the computation of the set $\text{split}(\llbracket A \rrbracket_{\Delta}^c)$ of maximally connected substructures from any canonical Δ -model of A . 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 $\text{split}(\llbracket A \rrbracket_{\Delta}^c)$ as $\llbracket [P]_{\Gamma}^c \rrbracket$, for a fresh nullary predicate P and a SID Γ whose construction is described next.

Given a qpf formula ψ , we define $\zeta(\psi) \subseteq \text{fv}(\psi) \times \text{fv}(\psi)$ as the least equivalence relation such that $(y, z) \in \zeta(\psi)$ iff $r(x_1, \dots, x_{\#r})$ occurs in ψ and $y, z \in \{x_1, \dots, x_{\#r}\}$, for some $r \in \mathbb{R}$. Let $B(y_1, \dots, y_{\#B})$ be a predicate atom, $J \subseteq [1.. \#B]$ be a set of indices

$j_1 \leq \dots \leq j_p$ and $\xi \subseteq J \times J$ be an equivalence relation and B^ξ be a fresh predicate of arity p . In particular, $\#B^\xi = 0$ if $\xi = \emptyset$ is the empty relation. We define the shorthands:

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

Consider a rule of Δ of the form (5), formulae 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:

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

We distinguish two cases. If (1) 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:

- (iv) $\text{fv}_{J_0}(B_0(x_1, \dots, x_{\#B_0})) \cap \text{fv}(\psi'') = \emptyset$ and $\text{fv}_{\bar{J}_0}(B_0(x_1, \dots, x_{\#B_0})) \cap \text{fv}(\psi') = \emptyset$,
- (v) $\text{fv}_{J_0}(B_0(x_1, \dots, x_{\#B_0})) \cap \text{fv}_{\bar{J}_i}(B_i(z_{i,1}, \dots, z_{i,\#B_i})) = \emptyset$ 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 $i \in [1..\ell]$,
- (vi) 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 $(x, y) \in \Xi$,
- (vii) $\xi_0(B_0(x_1, \dots, x_{\#B_0})) = \Xi \upharpoonright_{\{x_1, \dots, x_{\#B_0}\}} \cup \{(x, x) \mid x \in \text{fv}_{J_0}(B_0(x_1, \dots, x_{\#B_0}))\}$

then we add to Γ the following rule:

$$B_0(x_1, \dots, x_{\#B_0})/\xi_0 \leftarrow \exists y_1 \dots \exists y_m . \psi' * \bigstar_{i \in [1..\ell], J_i \neq \emptyset} B_i(z_{i,1}, \dots, z_{i,\#B_i})/\xi_i \quad (6)$$

Else, (2) if (viii) Ξ defines an unique equivalence class, and (ix) $(x, x) \notin \Xi$ for all $x \in \{x_1, \dots, x_{\#B_0}\}$ then we add to Γ the following rule:

$$P \leftarrow \exists y_1 \dots \exists y_m . \psi' * \bigstar_{i \in [1..\ell], J_i \neq \emptyset} B_i(z_{i,1}, \dots, z_{i,\#B_i})/\xi_i \quad (7)$$

Same as before (see Lemma 35 in §6.2), we can assume without loss of generality that Δ is equality-free and produces only satisfiable formulae starting from A . Moreover, we assume that every predicate defined by a rule of Δ occurs on some complete Δ -unfolding of A . Obviously, the rules that do not meet this requirement can be removed from Δ without changing $\llbracket A \rrbracket_\Delta^c$.

Lemma 36. *For any equality-free SID Δ and a nullary predicate symbol A , such that each complete Δ -unfolding of A yields a satisfiable formula, one can effectively build a SID Γ and a nullary predicate P such that $\text{split}(\llbracket A \rrbracket_\Delta^c) = \llbracket P \rrbracket_\Gamma^c$.*

See proof on page 69.

Proof of Lemma 3 For any set of structures \mathcal{S} , we have $\text{IEF}^*(\mathcal{S}) = \text{IF}(\text{EF}^*(\mathcal{S}))$, because the operations of internal and external fusion commute, namely $\text{EF}(\text{IF}(S_1), S_2) \subseteq \text{IF}(\text{EF}(S_1, S_2))$, for any structures S_1, S_2 . By Lemma 32, $\text{EF}^*(\llbracket A \rrbracket_\Delta^c)$ is treewidth-bounded only if $\text{split}(\llbracket A \rrbracket_\Delta^c)$ conforms to an RGB color scheme. Then, $\text{IF}(\text{EF}^*(\text{split}(\llbracket A \rrbracket_\Delta^c))) =$

$\text{IEF}^*(\text{split}(\llbracket A \rrbracket_{\Delta}^c)) = \text{split}(\text{IEF}^*(\llbracket A \rrbracket_{\Delta}^c))$ is treewidth-bounded, by Lemma 31 and Lemma 28. Thus, $\text{IEF}^*(\llbracket A \rrbracket_{\Delta}^c)$ is treewidth-bounded. \square

Proof of Lemma 4 By Lemma 36 one builds a SID Γ and a nullary predicate B , such that $\text{split}(\llbracket A \rrbracket_{\Delta}^c) = \llbracket B \rrbracket_{\Gamma}^c$. By Lemma 35, one effectively computes $(\llbracket B \rrbracket_{\Gamma}^c)^{\#3}$ and, by Lemmas 34 and 33, one effectively computes $(\text{EF}^*(\llbracket B \rrbracket_{\Gamma}^c))^{\#3} = \text{ef}_1^{\#3*}(\llbracket B \rrbracket_{\Gamma}^c)^{\#3}$. Therefore, $(\text{EF}^*(\text{split}(\llbracket A \rrbracket_{\Delta}^c)))^{\#3} = \text{ef}_1^{\#3*}(\llbracket B \rrbracket_{\Gamma}^c)^{\#3}$ can be effectively computed. We decide the treewidth-boundedness of $\text{EF}^*(\llbracket A \rrbracket_{\Delta}^c)$, by an application of Lemma 32, because, moreover, $\llbracket A \rrbracket_{\Delta}^c$ is a treewidth-bounded set of structures (Lemma 6). $\text{EF}^*(\llbracket A \rrbracket_{\Delta}^c)$ is treewidth-bounded iff condition (2) of Lemma 32 holds for $(\text{EF}^*(\text{split}(\llbracket A \rrbracket_{\Delta}^c)))^{\#3}$. \square

7 Conclusions and Future Work

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, e.g., hamiltonicity, or 3-colorability.

References

1. Abiteboul, S., Buneman, P., Suciu, D.: Data on the Web: From Relations to Semistructured Data and XML. Morgan Kaufmann (2000)
2. Ahrens, E., Bozga, M., Iosif, R., Katoen, J.: Reasoning about distributed reconfigurable systems. In: Object-Oriented Programming, Systems, Languages and Applications (OOPSLA). vol. To appear (2022), <https://arxiv.org/abs/2107.05253>
3. Arad, C.I.: Programming Model and Protocols for Reconfigurable Distributed Systems. Ph.D. thesis, KTH Royal Institute of Technology (2013)
4. Berger, R.: The undecidability of the domino problem. American Mathematical Soc. (1966)
5. Bodlaender, H.L.: A partial k -arboretum of graphs with bounded treewidth. Theor. Comput. Sci. **209**(1-2), 1–45 (1998). [https://doi.org/10.1016/S0304-3975\(97\)00228-4](https://doi.org/10.1016/S0304-3975(97)00228-4), [https://doi.org/10.1016/S0304-3975\(97\)00228-4](https://doi.org/10.1016/S0304-3975(97)00228-4)
6. Bozga, M., Bueri, L., Iosif, R.: 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. Lecture Notes in Computer Science, vol. 13385, pp. 691–711. Springer (2022). https://doi.org/10.1007/978-3-031-10769-6_40, https://doi.org/10.1007/978-3-031-10769-6_40
7. Bozga, M., Bueri, L., Iosif, R.: On an invariance problem for parameterized concurrent systems. In: Klin, B., Lasota, S., Muscholl, A. (eds.) 33rd International Conference on Concurrency Theory, CONCUR 2022, September 12-16, 2022, Warsaw, Poland. LIPIcs, vol. 243, pp. 24:1–24:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2022). <https://doi.org/10.4230/LIPIcs.CONCUR.2022.24>, <https://doi.org/10.4230/LIPIcs.CONCUR.2022.24>
8. Bradbury, J., Cordy, J., Dingel, J., Wermelinger, M.: A survey of self-management in dynamic software architecture specifications. In: Proceedings of the 1st ACM SIGSOFT workshop on Self-managed systems. pp. 28–33. ACM (2004)
9. Brotherston, J., Fuhs, C., Pérez, J.A.N., Gorogiannis, N.: A decision procedure for satisfiability in separation logic with inductive predicates. In: Henzinger, T.A., Miller, D. (eds.) 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. pp. 25:1–25:10. ACM (2014). <https://doi.org/10.1145/2603088.2603091>, <https://doi.org/10.1145/2603088.2603091>

10. Butting, A., Heim, R., Kautz, O., Ringert, J.O., Rumpe, B., Wortmann, A.: A classification of dynamic reconfiguration in component and connector architecture description. In: Proceedings of MODELS 2017 Satellite Event: Workshops (ModComp). CEUR Workshop Proceedings, vol. 2019, pp. 10–16. CEUR-WS.org (2017)
11. Calcagno, C., Distefano, D., O’Hearn, P.W., Yang, H.: Compositional shape analysis by means of bi-abduction. *J. ACM* **58**(6) (Dec 2011). <https://doi.org/10.1145/2049697.2049700>, <https://doi.org/10.1145/2049697.2049700>
12. Calcagno, C., O’Hearn, P.W., Yang, H.: Local action and abstract separation logic. In: 22nd IEEE Symposium on Logic in Computer Science (LICS 2007), 10-12 July 2007, Wrocław, Poland, Proceedings. pp. 366–378. IEEE Computer Society (2007). <https://doi.org/10.1109/LICS.2007.30>, <https://doi.org/10.1109/LICS.2007.30>
13. Cao, J., Chan, A.T.S., Sun, Y.: GOP: A Graph-Oriented Programming Model for Parallel and Distributed Systems, pp. 21–36. Springer US, Boston, MA (2005), https://doi.org/10.1007/0-387-28967-4_2
14. Cardelli, L., Gardner, P., Ghelli, G.: A Spatial Logic for Querying Graphs. In: Widmayer, P., Ruiz, F.T., Bueno, R.M., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) Proceedings of the 29th International Colloquium on Automata, Languages and Programming (ICALP’02). Lecture Notes in Computer Science, vol. 2380, pp. 597–610. Springer (Jul 2002)
15. Cavalcante, E., Batista, T.V., Oquendo, F.: Supporting dynamic software architectures: From architectural description to implementation. In: Bass, L., Lago, P., Kruchten, P. (eds.) 12th Working IEEE/IFIP Conference on Software Architecture, WICSA 2015. pp. 31–40. IEEE Computer Society (2015)
16. Courcelle, B.: The monadic second-order logic of graphs. i. recognizable sets of finite graphs. *Information and Computation* **85**(1), 12–75 (1990). [https://doi.org/https://doi.org/10.1016/0890-5401\(90\)90043-H](https://doi.org/https://doi.org/10.1016/0890-5401(90)90043-H), <https://www.sciencedirect.com/science/article/pii/089054019090043H>
17. Courcelle, B., Engelfriet, J.: Graph Structure and Monadic Second-Order Logic: A Language-Theoretic Approach. *Encyclopedia of Mathematics and its Applications*, Cambridge University Press (2012). <https://doi.org/10.1017/CBO9780511977619>
18. van Dalen, D.: Logic and structure (3. ed.). Universitext, Springer (1994)
19. Degano, P., Nicola, R.D., Meseguer, J. (eds.): Concurrency, Graphs and Models, Essays Dedicated to Ugo Montanari on the Occasion of His 65th Birthday, Lecture Notes in Computer Science, vol. 5065. Springer (2008). <https://doi.org/10.1007/978-3-540-68679-8>, <https://doi.org/10.1007/978-3-540-68679-8>
20. Ebbinghaus, H., Flum, J.: Finite model theory. *Perspectives in Mathematical Logic*, Springer (1995)
21. Echenim, M., Iosif, R., Peltier, N.: Decidable Entailments in Separation Logic with Inductive Definitions: Beyond Establishment. In: Baier, C., Goubault-Larrecq, J. (eds.) 29th EACSL Annual Conference on Computer Science Logic (CSL 2021). *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 183, pp. 20:1–20:18. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2021). <https://doi.org/10.4230/LIPIcs.CSL.2021.20>, <https://drops.dagstuhl.de/opus/volltexte/2021/13454>
22. Echenim, M., Iosif, R., Peltier, N.: Unifying decidable entailments in separation logic with inductive definitions. In: Platzer, A., Sutcliffe, G. (eds.) Automated Deduction - CADE 28 - 28th International Conference on Automated Deduction, Virtual Event, July 12-15, 2021, Proceedings. *Lecture Notes in Computer Science*, vol. 12699, pp. 183–199. Springer (2021). https://doi.org/10.1007/978-3-030-79876-5_11, https://doi.org/10.1007/978-3-030-79876-5_11
23. Flum, J., Grohe, M.: Parameterized Complexity Theory. *Texts in Theoretical Computer Science. An EATCS Series*, Springer (2006). <https://doi.org/10.1007/3-540-29953-X>, <https://doi.org/10.1007/3-540-29953-X>
24. Greibach, S.: A note on undecidable properties of formal languages. *Math. Systems Theory* **2**, 1–6 (1968). <https://doi.org/https://doi.org/10.1007/BF01691341>

25. Iosif, R., Rogalewicz, A., Simáček, J.: The tree width of separation logic with recursive definitions. In: Bonacina, M.P. (ed.) *Automated Deduction - CADE-24 - 24th International Conference on Automated Deduction*, Lake Placid, NY, USA, June 9-14, 2013. *Proceedings. Lecture Notes in Computer Science*, vol. 7898, pp. 21–38. Springer (2013). https://doi.org/10.1007/978-3-642-38574-2_2, https://doi.org/10.1007/978-3-642-38574-2_2
26. Iosif, R., Zuleger, F.: Expressiveness results for an inductive logic of separated relations. In: *34th International Conference on Concurrency Theory*. vol. 279, p. to appear. *LIPIcs, Dagstuhl* (2023), <https://doi.org/10.48550/arXiv.2208.01520>
27. Ishtiaq, S.S., O’Hearn, P.W.: BI as an assertion language for mutable data structures. In: Hankin, C., Schmidt, D. (eds.) *Conference Record of POPL 2001: The 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, London, UK, January 17-19, 2001. pp. 14–26. ACM (2001)
28. Jones, N.D., Muchnick, S.S.: 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*. p. 66–74. *POPL ’82*, Association for Computing Machinery, New York, NY, USA (1982). <https://doi.org/10.1145/582153.582161>, <https://doi.org/10.1145/582153.582161>
29. Katelaan, J., Zuleger, F.: Beyond symbolic heaps: Deciding separation logic with inductive definitions. In: Albert, E., Kovács, L. (eds.) *LPAR 2020: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning*, Alicante, Spain, May 22-27, 2020. *EPiC Series in Computing*, vol. 73, pp. 390–408. EasyChair (2020). <https://doi.org/10.29007/vkmj>, <https://doi.org/10.29007/vkmj>
30. Kuncak, V., Rinard, M.: Generalized records and spatial conjunction in role logic. In: *Static Analysis*. pp. 361–376. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
31. Le Metayer, D.: Describing software architecture styles using graph grammars. *IEEE Transactions on Software Engineering* **24**(7), 521–533 (1998). <https://doi.org/10.1109/32.708567>
32. Matheja, C., Pagel, J., Zuleger, F.: A decision procedure for guarded separation logic. *ACM Trans. Comput. Logic* (apr 2022). <https://doi.org/10.1145/3534927>, <https://doi.org/10.1145/3534927>, just accepted, available online
33. Reynolds, J.C.: 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*. pp. 55–74. IEEE Computer Society (2002). <https://doi.org/10.1109/LICS.2002.1029817>, <https://doi.org/10.1109/LICS.2002.1029817>
34. Seese, D.: The structure of the models of decidable monadic theories of graphs. *Annals of Pure and Applied Logic* **53**(2), 169–195 (1991). [https://doi.org/https://doi.org/10.1016/0168-0072\(91\)90054-P](https://doi.org/https://doi.org/10.1016/0168-0072(91)90054-P), <https://www.sciencedirect.com/science/article/pii/016800729190054P>
35. Seymour, P., Thomas, R.: Graph searching and a min-max theorem for tree-width. *Journal of Combinatorial Theory, Series B* **58**(1), 22–33 (1993). <https://doi.org/https://doi.org/10.1006/jctb.1993.1027>, <https://www.sciencedirect.com/science/article/pii/S0095895683710270>
36. Taentzer, G., Goedicke, M., Meyer, T.: Dynamic change management by distributed graph transformation: Towards configurable distributed systems. In: *International Workshop on Theory and Application of Graph Transformations*. pp. 179–193. Springer (1998)
37. Wermelinger, M.: Towards a chemical model for software architecture reconfiguration. *IEE Proceedings-Software* **145**(5), 130–136 (1998)
38. Wermelinger, M., Fiadeiro, J.L.: A graph transformation approach to software architecture reconfiguration. *Sci. Comput. Program.* **44**(2), 133–155 (2002)

A Material from Section 2

Theorem 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 relations, is undecidable.*

Proof (Proof of Theorem 1). We will reduce from the undecidability of the *Tiling Problem* [4]. 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 [4]. 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 bounded tree-width. We consider some instance of the tiling problem. For encoding this problem, we define the signature $\mathbb{R} = \{\text{up}, \text{right}, \text{N}, \text{S}, \text{E}, \text{W}, \text{I}, \text{T}_1, \dots, \text{T}_n\}$ to consist of the binary relations up and right , the unary relations $\text{S}, \text{E}, \text{W}, \text{N}, \text{I}$ and the unary relations $\text{T}_1, \dots, \text{T}_n$ (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

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

- $\phi_4 \stackrel{\text{def}}{=} \forall x \forall y . \text{up}(x, y) \rightarrow \bigvee_{(t_i, t_j) \in V} T_i(x) \wedge T_j(y)$ states that tiles, that are next to each other vertically, satisfy the vertical matching constraint.

As usual, the formulæ $\exists^{\leq 1} x . \phi$ stand for $\exists x . \phi \rightarrow \forall y . \phi[x/y] \rightarrow y = x$. 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:

Each model of ϕ can be decomposed into (disjoint unions of) grids, cylinders, and cyclic grids, where grids have S, E, W, N borders, cylinders have either S, N or E, W borders, and cyclic grids do not have any borders and only consist of internal nodes.

We recall that we are only interested in finite models of first-order formulae. We note that ϕ specifies that `up` and `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 `up`, `right`. We choose a representative u_C , for each component C . We observe that either there are $j \leq 0 \leq i$ such that $\text{right}^j(u_C)$ is W-labelled and $\text{right}^i(u_C)$ is E-labelled, or $\text{right}^i(u_C) = u_C$, for some $i \geq 0$ (because the universe is finite and the functionality of `right` and right^{-1} ensures that the only possible loop returns to u_C). An analogous statement holds for `up` as well as the N and S labels. We now call a component C a *grid*, if u_C reaches S, E, W, N via `up`, `right` and their inverses, a *cylinder* if u_C reaches S, N or E, W via `up` and its inverse resp. `right` and its inverse, or a *cyclic grid*, otherwise. We now justify the naming of these components.

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

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

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 $\text{right}^{a_1}(\text{up}^{b_1}(u_C)) \neq \text{right}^{a_2}(\text{up}^{b_2}(u_C))$. We will assume that $\text{right}^{a_1}(\text{up}^{b_1}(u_C)) = \text{right}^{a_2}(\text{up}^{b_2}(u_C))$ and derive a contradiction. Let us assume that $a_1 > a_2$ (the other cases are analogous). Then, we have $\text{right}^{a_1+s}(\text{up}^{b_1}(u_C)) = \text{right}^{a_2+s}(\text{up}^{b_2}(u_C))$ for $s = i - a_1 \geq 0$. However, $\text{right}^{a_1+s}(\text{up}^{b_1}(u_C))$ is E-labelled, while $\text{right}^{a_2+s}(\text{up}^{b_2}(u_C))$ is not, by (2), contradiction. For 5., we observe that every node reachable from u_C connected via `up`, `right` and their inverses can be represented as $\text{right}^a(\text{up}^b(u_C))$, 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 **right** is interpreted as $\{(x, y), (x+1, y) \mid x \in [j, i-1], y \in [k, l]\}$ and **up** as $\{(x, y), (x, y+1) \mid x \in [j, i], y \in [k, l-1]\}$.

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.

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 **up**, **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 $\{(x, y), (x, y+1) \mid x \in [j, i], y \in [0, k-2]\} \cup \{(x, k), (x, 1) \mid x \in [j, i]\}$.

Consider a cyclic grid 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 $0 \leq b < l$.

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 $\text{right}^a(\text{up}^b(u_C))$ 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 $\text{up}^k(u_C) = u_C$ and $\text{right}^l(u_C) = u_C$. We note that the $\text{right}^a(\text{up}^b(u_C))$ are in general not pairwise different (e.g., we might have $\text{right}^a(u_C) = \text{up}^b(u_C)$ 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 $\text{right}^a(\text{up}^b(u_C))$ of a cyclic grid component are pairwise different.

We claim that ϕ has models of unbounded treewidth iff there is a tiling of the upper-right quadrant of the plane.

“ \Rightarrow ” 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 . “ \Leftarrow ” 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 cyclic grid component C , we immediately obtain a tiling of the upper-right quadrant by unrolling the cyclic grid: we define the tiling of the upper right quadrant by placing the tile of the element $\text{right}^i(\text{up}^j(u_C))$ 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 cyclic grid 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 [35]. 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 no cyclic grid 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_i^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 the horizontal requirement (the vertical one is symmetric). Consider some 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 . \square

B Material from Section 3

Lemma 5. *Given a SID Δ , one can build an equality-free SID Δ' , such that $\llbracket A \rrbracket_\Delta = \llbracket A \rrbracket_{\Delta'}$ for any nullary predicate A .*

Proof (Proof of Lemma 5). See [26, Lemma 9].

Lemma 6. *Given a nullary predicate A , there exists a constant $W \geq 1$, depending only on Δ , such that $\text{tw}(S) \leq W$, for any $S \in \llbracket A \rrbracket_\Delta^c$.*

Proof (Proof of Lemma 6). See [26, Lemma 11].

Lemma 7. *Let A be a nullary predicate. Then, $\llbracket A \rrbracket_\Delta = \widetilde{\text{IF}}(\llbracket A \rrbracket_\Delta^c) \subseteq \text{IF}(\llbracket A \rrbracket_\Delta^c)$.*

Proof (Proof of Lemma 7). It is sufficient to prove $\llbracket A \rrbracket_\Delta = \widetilde{\text{IF}}(\llbracket A \rrbracket_\Delta^c)$, since $\widetilde{\text{IF}}(\llbracket A \rrbracket_\Delta^c) \subseteq \text{IF}(\llbracket A \rrbracket_\Delta^c)$ is immediate, by Def. 10 i.e., any equivalence relation that is compatible with a rich canonical Δ -model (S, δ) is also compatible with S .

“ \subseteq ” Let $S \in \llbracket A \rrbracket_\Delta$ be a structure. By Prop. 1, we have $S \models \exists y_1 \dots \exists y_m . \psi$, where ψ is qpf, such that $\text{fv}(\psi) = \{y_1, \dots, y_m\}$ and $A \Rightarrow_\Delta^* \exists y_1 \dots \exists y_m . \psi$ is a complete Δ -unfolding. Then there exists a store \mathfrak{s} , such that $S \models^{\mathfrak{s}} \psi$. Let $\bar{S} = (\bar{U}, \bar{\sigma})$ be a structure and $\bar{\mathfrak{s}}$ be a store such that $\bar{S} \models^{\bar{\mathfrak{s}}} \psi$ and $\bar{\mathfrak{s}}(y_i) \neq \bar{\mathfrak{s}}(y_j)$, for all $i \neq j \in [1..m]$. Since Δ is equality-free, there are no equality atoms in ψ , hence such a structure and a store exist. We consider $\approx \subseteq \bar{U} \times \bar{U}$ to be the least equivalence relation such that $\bar{\mathfrak{s}}(y_i) \approx \bar{\mathfrak{s}}(y_j) \stackrel{\text{def}}{\iff} \mathfrak{s}(y_i) = \mathfrak{s}(y_j)$. To prove that \approx is compatible with \bar{S} , consider two tuples $\langle \bar{\mathfrak{s}}(z_1), \dots, \bar{\mathfrak{s}}(z_{\#r}) \rangle, \langle \bar{\mathfrak{s}}(z'_1), \dots, \bar{\mathfrak{s}}(z'_{\#r}) \rangle \in \bar{\sigma}(r)$, for some $r \in \mathbb{R}$ and suppose, for a contradiction, that $\bar{\mathfrak{s}}(z_i) \approx \bar{\mathfrak{s}}(z'_i)$, for all $i \in [1..\#r]$. Then $r(z_1, \dots, z_{\#r}) * r(z'_1, \dots, z'_{\#r})$ is a subformula of ψ , possibly 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 $S \models^{\mathfrak{s}} \psi$. Since $\bar{S} \models^{\bar{\mathfrak{s}}} \psi$ and $\bar{\mathfrak{s}}$ 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 $\bar{s}^{-1}(u)$ is a singleton, for each $u \in \text{supp}(\bar{\sigma})$. Let $\bar{\delta} \subseteq \bar{U} \times \bar{U}$ be the relation defined as $\bar{\delta}(u, v)$ iff the disequality $\bar{s}^{-1}(u) \neq \bar{s}^{-1}(v)$ occurs in ψ . Then \approx is compatible with $(\bar{S}, \bar{\delta}) \in \llbracket A \rrbracket'_\Delta$ hence $\bar{S}_{/\approx} \in \widetilde{\text{IF}}(\llbracket A \rrbracket'_\Delta)$. Finally, the mapping $h : \text{supp}(\sigma) \rightarrow \text{supp}(\bar{\sigma})$ defined as $h(s(y_i)) \stackrel{\text{def}}{=} [\bar{s}(y_i)]_\approx$, for all $i \in [1..m]$ is shown to be an isomorphism between S and $\bar{S}_{/\approx}$, leading to $S \in \widetilde{\text{IF}}(\llbracket A \rrbracket'_\Delta)$.

“ \supseteq ” Let $S \in \widetilde{\text{IF}}(\llbracket A \rrbracket'_\Delta)$ be a structure. Then there exists a rich canonical Δ -model $(\bar{S}, \bar{\delta}) \in \llbracket A \rrbracket'_\Delta$, where $\bar{S} = (\bar{U}, \bar{\sigma})$ and an equivalence relation $\approx \subseteq \bar{U} \times \bar{U}$ such that \approx is compatible with $(\bar{S}, \bar{\delta})$ and S is isomorphic to $\bar{S}_{/\approx}$. Since $(\bar{S}, \bar{\delta}) \in \llbracket A \rrbracket'_\Delta$, there exists a complete Δ -unfolding $A \Rightarrow^*_\Delta \exists y_1 \dots \exists y_m . \psi$, such that ψ is qpf and a store \mathfrak{s} , injective over y_1, \dots, y_m , such that $\bar{S} \models^{\mathfrak{s}} \psi$ and $\bar{\delta}(s(z), s(z'))$ 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 $(\bar{S}, \bar{\delta})$, 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 Def. 2.
- $\psi = \psi_1 * \psi_2$: 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 Def. 3. 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 $\bar{S}_{i/\approx} \models^{\bar{\mathfrak{s}}} \psi_i$, for $i = 1, 2$, thus $\bar{S}_{/\approx} \models^{\bar{\mathfrak{s}}} \psi_1 * \psi_2$.

Hence $\bar{S}_{/\approx} \in \llbracket A \rrbracket_\Delta$ and $S \in \llbracket A \rrbracket_\Delta$ follows, since $\llbracket A \rrbracket_\Delta$ is closed under isomorphism [26, Prop. 7].

C Material from Section 4

Lemma 8. *Let $\mathcal{A} = (\mathbb{A}, Q, \mathfrak{t}, \delta)$ be a choice-free automaton, such that $\delta = \delta^1 \uplus \delta^\infty$ (Def. 14) 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 (Proof of Lemma 8). 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 Def. 14, $\mathfrak{G}_\mathcal{A}$ is a tree and, moreover, $S_0 = \{\mathfrak{t}\}$, because \mathcal{A} is rooted. Let $\Lambda : \mathcal{N} \cup \delta \rightarrow \{1, \infty\}$ be the labeling from Def. 14. 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 Def. 14, and $\text{entry}(S_0) \stackrel{\text{def}}{=} \mathfrak{t}$. 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 Def. 14.

We prove first an invariant of 1-labeled linear SCCs:

Fact 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 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 $p = \varepsilon$. For the inductive step, let S' be the parent of S in $\mathfrak{G}_{\mathcal{A}}$. By points (1) and (2c) of Def. 14, $\bullet S = \{\tau\}$ 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 1, there exists a descendant p of p' , such that $\theta(p) = \text{entry}(S)$. \square

Fact 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 Def. 14. \square

Let $\tau : q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell)$ be a transition, such that $\Lambda(\tau) = 1$. By point (2b) of Def. 14, we have $\tau \in S \bullet$ for some linear SCC $S \in \mathcal{N}$, such that $\Lambda(S) = 1$. By Facts 2 and 3, there exists a unique position $p \in \text{dom}(\theta)$, such that $\theta(p) = \text{entry}(S)$. By Fact 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.

Lemma 9. *Let $\mathcal{A} = (\mathbb{A}, Q, \iota, \delta)$ be a choice-free automaton, where $\delta = \delta^1 \uplus \delta^\infty$ (Def. 14). 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 $\theta' \in \mathcal{R}_{q_0}^\infty(\mathcal{A})$ such that q_0 occurs on the frontier of θ' .

Proof (Proof of Lemma 9). 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 Def. 14 and let $q \in \bullet(\delta^\infty)$ be a state. W.l.o.g., we consider that $\mathcal{L}(\mathcal{A}) \neq \emptyset$ and \mathcal{A} to be 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 $\theta_1(\varepsilon) = \iota$ and $\theta_1(p_1) = q$. Let p_2 be the longest strict prefix of p_1 such that the transition $\tau : \theta_1(p_2) \xrightarrow{a} \langle \theta_1(p_2 1), \dots, \theta_1(p_2 \ell) \rangle$ is in δ^1 for some $a \in \mathbb{A}$ and index ℓ . This position p_2 exists thanks to Def. 14: $S_0 = \{\iota\}$ is linear, $\Lambda(S_0) = 1$ by cond.2c, $\text{card}(S_0 \bullet) = 1$ by cond.2a, and the only transition $\tau_0 \in \{\iota\} \bullet$ is in δ^1 by cond.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 $r \in [1..l]$ and position p_3 and define the partial run θ_2 as $\theta_2(\varepsilon) \stackrel{\text{def}}{=} \theta_1(p_2 r)$ 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 $S \in \mathcal{X}$ 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 cond.1 of Def. 14), thus $\Lambda(S) = 1$ by cond.2c of Def. 14.

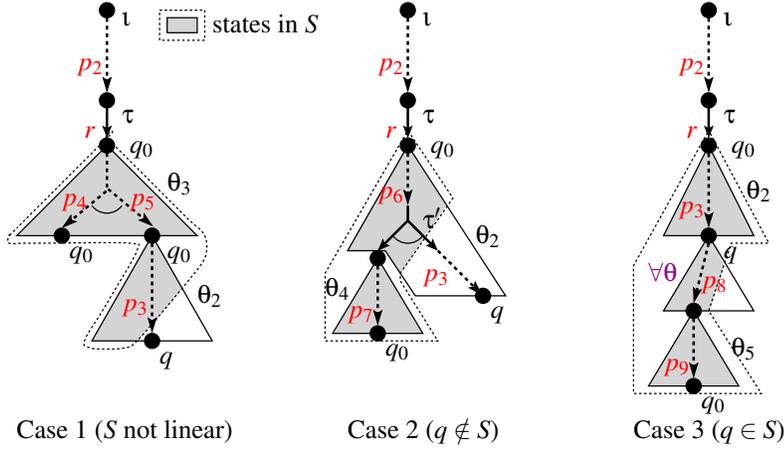


Fig. 5. The cases from the proof of Lemma 9.

We distinguish three cases (see Fig. 5 for an illustration):

1. If S is not linear, there exists a transition $\tau' \in \bullet S$ such that $\text{card}(\tau'^\bullet \cap S) \geq 2$. Let $q' \stackrel{\text{def}}{=} \tau'$ 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 $\bullet S$, 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 $\text{dom}(\theta_3) \cup \{p_5 u \mid u \in \text{dom}(\theta_2)\}$, that extends θ_3 by $\theta_0(p_5 u) \stackrel{\text{def}}{=} \theta_2(u)$, 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 $\theta_0(p_5 p_3) = q$, with $p_4 \neq p_5 p_3 \in \text{fr}(\theta_0)$.
2. If S is linear and $q \notin S$, there exists a unique position p_6 and a transition:

$$\tau' : \theta_2(p_6) \xrightarrow{\alpha} \langle \theta_2(p_6 1), \dots, \theta_2(p_6 k) \rangle \in \delta^\infty$$

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 $\theta_2(p_6 r) \notin S$ and $p_6 r$ is a prefix of p_3 . Then $\Lambda(\tau') = \infty$ and, by cond.2b of Def. 14, we have $\tau' \notin \bullet S$, hence $\tau' \in \bullet S$ and $q' \stackrel{\text{def}}{=} \theta_2(p_6 r') \in S$ 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 $\text{dom}(\theta_2) \cup \{p_6 r' u \mid u \in \text{dom}(\theta_4)\}$,

- by extending θ_2 with $\theta_0(p_6 r' u) \stackrel{\text{def}}{=} \theta_4(u)$, for all $u \in \text{dom}(\theta_4)$. Then θ_0 satisfies point 1 of the lemma because $\theta_0(p_6 r' p_7) = q_0$ and $\theta_0(p_3) = q$, with $p_6 r' p_7 \neq p_3 \in \text{fr}(\theta_0)$.
3. If S is linear and $q \in S$, let $\theta \in \mathcal{R}_q^\infty(\mathcal{A})$ be a partial run. Then S^\bullet contains only one transition in δ^1 , thus for every position $u \in \text{dom}(\theta) \setminus \text{fr}(\theta)$, such that $\theta(u) \in S$ and 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 $\text{dom}(\theta) \cup \{p_8 u \mid u \in \text{dom}(\theta_5)\}$, as $\theta'(p_8 u) \stackrel{\text{def}}{=} \theta_5(u)$ for all $u \in \text{dom}(\theta_5)$. The partial run θ' satisfies point 2 of the lemma, because $\theta'(p_8 p_9) = q_0$, with $p_8 p_9 \in \text{fr}(\theta')$.

Lemma 10. *Given an automaton \mathcal{A} , one can build finitely many choice-free automata $\mathcal{A}_1, \dots, \mathcal{A}_n$, such that $\mathcal{L}(\mathcal{A}) = \bigcup_{i=1}^n \mathcal{L}(\mathcal{A}_i)$.*

Proof (Proof of Lemma 10). We assume w.l.o.g. that $\mathcal{A} = (\mathbb{A}, Q, \iota, \delta)$ 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 $S_i \stackrel{\text{def}}{=} \bigcup_{(S_i, S_j) \in \mathcal{E}^*} S_j$ 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 $\{q_0, q_{i_1}, \dots, q_{i_j}\} = \{q_0, \dots, q_\ell\} \cap S_i$ 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 $q \in \tau^\bullet \cap S_i$ 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 fullfils condition (1) of Def. 14 and has the same language as \mathcal{A} , using Lemma 15. 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 $\tau \in \bullet S_i \cup S_i^\bullet$, 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 $\mathcal{A}_i = (\mathbb{A}, Q_i, \iota, \delta_i)$, where $Q_i \stackrel{\text{def}}{=} \bigcup \{S_j \mid \bar{x}_{i,j} > 0, j \in [1..N]\}$ and $\delta_i \stackrel{\text{def}}{=} \{\tau \in \delta \mid \bar{y}_{i,\tau} > 0\}$. We are left with proving the following facts:

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

Proof. We prove below the points of Def. 14:

- (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 Def. 14:
- (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 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:

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

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 $p = \varepsilon$, 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_τ is assigned 1 and $y_{\tau'}$ is assigned 0, for all $\tau' \in S_1^\bullet \setminus \{\tau\}$ at step (iv).

Inductive step If $p \in \text{dom}(\theta) \setminus \{\varepsilon\}$, 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 $\sum_{\tau' \in \bullet S_j} y_{\tau'} \cdot \text{card}(\tau' \cap S_j) > 0$ 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 non-linear, 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

This concludes the proof of the Lemma.

Lemma 11. Let $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta_\Delta$ be a transition of $\mathcal{A}_{\Delta, A}$. Then, for each $i \in [1.. \ell]$ and each $j \in [1.. \#q_i]$, there exists an ε -variable z , such that $x_j^{[i]} \approx_\alpha z$.

Proof (Proof of Lemma 11). Directly from the definition (1) of the labels on the transitions of $\mathcal{A}_{\Delta, A}$.

Lemma 12. (1) Given a SID Δ and a nullary predicate A , one can build an automaton $\mathcal{A}_{\Delta, A}$ such that $\llbracket A \rrbracket_\Delta = \llbracket \mathcal{A}_{\Delta, A} \rrbracket$. (2) Given an automaton $\mathcal{A} = (\Sigma, Q, \iota, \delta)$, one can build a SID $\Delta_{\mathcal{A}}$, such that $\llbracket \mathcal{A} \rrbracket = \llbracket A_\iota \rrbracket_{\Delta_{\mathcal{A}}}$ and $\llbracket \mathcal{A} \rrbracket^r = \llbracket A_\iota \rrbracket_{\Delta_{\mathcal{A}}}^r$, for a nullary predicate A .

Proof (Proof of Lemma 12). (1) For the first part, we prove the two directions of the following equivalence: for all structures S and predicates B of arity n , there exists a store \mathfrak{s} such that $S \models_\Delta^{\mathfrak{s}} B(x_1, \dots, x_n)$ iff there exists a tree $t \in \mathcal{L}_{q_B}(\mathcal{A}_{\Delta, A})$ and a store $\bar{\mathfrak{s}}$ such that $S \models^{\bar{\mathfrak{s}}} \Theta(t)$ and $\bar{\mathfrak{s}}(x_j^{[i]}) = \mathfrak{s}(x_j)$ for all $j \in [1..n]$.

“ \Rightarrow ” We proceed by induction on the definition of $S \models_\Delta^{\mathfrak{s}} B(x_1, \dots, x_n)$. Then Δ contains a rule ρ of the form $B(x_1, \dots, x_n) \leftarrow \exists y_1 \dots \exists y_m \cdot \Psi * \bigstar_{i=1}^\ell B_i(z_{i,1}, \dots, z_{i,n_i})$ and we can decompose the structure $S = S_0 \bullet \dots \bullet S_\ell$, such that $S_0 \models^{\mathfrak{s}'}$ Ψ and $S_i \models_\Delta^{\mathfrak{s}'}$ $B_i(z_{i,1}, \dots, z_{i,n_i})$ 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 $\sigma_i \models_\Delta^{\mathfrak{s}_i}$ $B_i(x_1, \dots, x_{n_i})$ and, by induction hypothesis, there exists a tree $t_i \in \mathcal{L}_{q_{B_i}}(\mathcal{A}_{\Delta, A})$ and a store $\bar{\mathfrak{s}}_i$ such that $\sigma_i \models^{\bar{\mathfrak{s}}_i} \Theta(t_i)$ and $\bar{\mathfrak{s}}_i(x_j^{[i]}) = \mathfrak{s}_i(x_j)$ for all $j \in [1..n_i]$. Let $\bar{\mathfrak{s}}$ be a store such that:

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

Note that $\bar{\mathfrak{s}}$ is well defined because $\bar{\mathfrak{s}}_i(z^{[i]p}) = \bar{\mathfrak{s}}_j(z^{[i]p}) = \mathfrak{s}'(z^{[i]p})$, for all $z^{[i]p} \in \text{fv}(\Theta(t_i)) \cap \text{fv}(\Theta(t_j))$. We have $S_i \models^{\bar{\mathfrak{s}}} \Theta(t_i)^{[i]}$ for all $i \in [1.. \ell]$ and $S_0 \models^{\bar{\mathfrak{s}}} \alpha_\rho$, 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 $t \in \mathcal{L}_{q_B}(\mathcal{A}_{\Delta, A})$ and $\bar{\mathfrak{s}}(x_j^{[i]}) = \mathfrak{s}(x_j)$, 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 $t \in \mathcal{L}_{q_B}(\mathcal{A}_{\Delta, A})$ such that $S \models^{\bar{\mathfrak{s}}} \Theta(t)$. Since $t \in \mathcal{L}_{q_B}(\mathcal{A}_{\Delta, A})$, there is a transition $q_B \xrightarrow{t(\varepsilon)} (q_{B_1}, \dots, q_{B_\ell}) \in \delta_\Delta$ such that $t(\varepsilon) = \alpha_\rho$ for some rule ρ of the form above and $t|_i \in \mathcal{L}_{q_{B_i}}(\mathcal{A}_{\Delta, B})$ for all $i \in [1.. \ell]$. Meanwhile $S \models^{\bar{\mathfrak{s}}} t(\varepsilon) * \bigstar_{i=1}^\ell \Theta(t|_i)^{[i]}$, 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 $S_i \models^{\bar{\mathfrak{s}}} \Theta(t|_i)^{[i]} * \bigstar_{j=1}^{n_i} z_{i,j}^{[i]} = x_j^{[i]}$ 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 $\bar{\mathfrak{s}}_i(z^{[i]p}) = \bar{\mathfrak{s}}(z^{[i]p})$, for all $z^{[i]p} \in \text{fv}(\Theta(t|_i))$ and all $i \in [1.. \ell]$. By the inductive hypothesis on $t|_i$, there exists a store \mathfrak{s}_i such that $S_i \models_\Delta^{\mathfrak{s}_i} B_i(x_1, \dots, x_{n_i})$ and $\mathfrak{s}_i(x_j) = \bar{\mathfrak{s}}_i(x_j^{[i]})$, for all $j \in [1..n_i]$. We consider a store \mathfrak{s}' such that $\mathfrak{s}'(x_j) = \bar{\mathfrak{s}}(x_j^{[i]})$, for all $j \in [1..n]$, and $\mathfrak{s}'(y_j) \stackrel{\text{def}}{=} \bar{\mathfrak{s}}(y_j^{[i]})$, for all $j \in [1..m]$. For all $i \in [1.. \ell]$ and $j \in [1..n_i]$ we have $\mathfrak{s}_i(x_j) = \bar{\mathfrak{s}}_i(x_j^{[i]}) = \bar{\mathfrak{s}}(x_j^{[i]}) = \bar{\mathfrak{s}}(x_j^{[i]}) = \mathfrak{s}'(z_{i,j})$, because $z_{i,j}^{[i]} = x_j^{[i]}$ 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 $S_0 \models^{s'} \psi$, and by composing the structures and using ρ , we obtain $S \models_{\Delta}^{s'} B(x_1, \dots, x_n)$.

(2) To show $\llbracket \mathcal{A} \rrbracket = \llbracket A_1 \rrbracket_{\Delta_{\mathcal{A}}}$ 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{s} such that $S \models^{\bar{s}} \Theta(t)$ 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 $\mathfrak{s}(x_j) = \bar{s}(x_j^{[e]})$, 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 $S \models^{\bar{s}} \Theta(t)$. Since $t \in \mathcal{L}_{q_0}(\mathcal{A})$, there is a transition $q_0 \xrightarrow{t(\varepsilon)} (q_1, \dots, q_\ell) \in \delta$ such that $t|_i \in \mathcal{L}_{q_i}(\mathcal{A})$, for all $i \in [1..\ell]$. Meanwhile $S \models^{\bar{s}} t(\varepsilon) * \star_{i=1}^{\ell} \Theta(t|_i)^{[i]}$ thus, we can decompose the structure $S = S_0 \bullet \dots \bullet S_\ell$, such that $S_0 \models^{\bar{s}} t(\varepsilon)$ and $S_i \models^{\bar{s}} \Theta(t|_i)^{[i]}$. Let \bar{s}_i be a store, such that $\bar{s}_i(z^{[p]}) = \bar{s}(z^{[p]})$, for all $z^{[p]} \in \text{fv}(\Theta(t|_i))$ and all $i \in [1..\ell]$. By the inductive hypothesis on $t|_i$, there exists a store \mathfrak{s}_i such that $S_i \models_{\Delta_{\mathcal{A}}}^{\mathfrak{s}_i} A_{q_i}(x_1, \dots, x_{\#q_i})$ and $\mathfrak{s}_i(x_j) = \bar{s}_i(x_j^{[e]})$ for all $j \in [1..\#q_i]$. We consider a store \mathfrak{s}' , such that $\mathfrak{s}'(x_j) = \bar{s}(x_j^{[e]})$, 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 $\mathfrak{s}'(z) = \bar{s}(z)$, for all other variables $z \in \text{fv}(t(\varepsilon))$. Then $S_i \models_{\Delta_{\mathcal{A}}}^{\mathfrak{s}'} A_{q_i}(x_1^{[i]}, \dots, x_{\#q_i}^{[i]})$ and $S_0 \models^{\mathfrak{s}'} t(\varepsilon)[x_1^{[e]}/x_1, \dots, x_{\#q_0}^{[e]}/x_{\#q_0}]$ 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 (3) and we can decompose the structure $S = S_0 \bullet \dots \bullet S_\ell$ such that $S_0 \models^{\mathfrak{s}'} \alpha[x_1^{[e]}/x_1, \dots, x_{\#q_0}^{[e]}/x_{\#q_0}]$ 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 $S_i \models_{\Delta_{\mathcal{A}}}^{\mathfrak{s}_i} A_{q_i}(x_1, \dots, x_{\#q_i})$ and, by the inductive hypothesis, there exists a tree $t_i \in \mathcal{L}_{q_i}(\mathcal{A})$ and a store \bar{s}_i such that $\sigma_i \models^{\bar{s}_i} \Theta(t_i)$ and $\bar{s}_i(x_j^{[e]}) = \mathfrak{s}_i(x_j)$, for all $j \in [1..\#q_i]$. Let \bar{s} be a store such that:

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

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

(2) To show $\llbracket \mathcal{A} \rrbracket^r \subseteq \llbracket A_1 \rrbracket_{\Delta_{\mathcal{A}}}^r$, 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 $A_1 \Rightarrow_{\Delta_{\mathcal{A}}}^* \Theta(t)^{\exists}$. Since $S \models^{\mathfrak{s}} \Theta(t)$, we obtain $(S, \mathfrak{d}) \in \llbracket A_1 \rrbracket_{\Delta_{\mathcal{A}}}^r$, by Def. 8. Conversely, to show $\llbracket \mathcal{A} \rrbracket^r \supseteq \llbracket A_1 \rrbracket_{\Delta_{\mathcal{A}}}^r$, let $(S, \mathfrak{d}) \in \llbracket A_1 \rrbracket_{\Delta_{\mathcal{A}}}^r$, where $S = (U, \sigma)$ is a structure and $\mathfrak{d} \subseteq U \times U$ is a symmetric relation. Then, there exists a complete unfolding $A_1 \Rightarrow_{\Delta_{\mathcal{A}}}^* \exists x_1 \dots \exists x_n . \psi$, 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 $(S, \mathfrak{d}) \in \llbracket \Theta(t) \rrbracket^r$, thus $(S, \mathfrak{d}) \in \llbracket \mathcal{A} \rrbracket^r$.

Lemma 13. *Let $\mathcal{A} = (\Sigma, Q, \mathfrak{l}, \delta)$ be a trim automaton. Then, there exists a q -reset for (1) each pivot state $q \in (\delta^1)^\bullet \cap (\delta^\infty)^\bullet$ of \mathcal{A} and (2) each state $q \in (\delta^1)^\bullet \cap (\delta^\infty)^\bullet$, i.e., that is the origin of both a 1-transition and a ∞ -transition.*

Proof (Proof of Lemma 13). By Def. 17, $\mathfrak{P}_{\mathcal{A}}$ is the greatest fixpoint of the monotone function \mathcal{F} on the domain of positional functions $\mathfrak{P} : Q \rightarrow \text{pow}(\mathbb{N})$, defined below:

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

Namely, we have $\mathfrak{P}_{\mathcal{A}} = \mathcal{F}^i(\top)$ for a sufficiently large $i \geq 1$ and $\mathfrak{P}_{\mathcal{A}} = \mathcal{F}^j(\top)$ for 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:

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

We prove the following:

Fact 6 $\text{gfp}(\mathcal{F})(q) = \text{gfp}(\mathcal{G})(q)$, 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 $q \in (\delta^\infty)^\bullet$, thus necessarily $q = q_0$, where $q_0 \in (\delta^1)^\bullet$ is the state whose existence is stated by Lemma 10. 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 6, 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 $\{j \in [1.. \#q] \mid \exists k \in [1.. \#q] \setminus \mathfrak{P}_{\mathcal{A}}(q) . x_j^{[e]} \approx_{\theta(t)} x_k^{[p]}\} = \mathfrak{P}_{\mathcal{A}}(q)$, thus $\mathcal{G}(\top)(q) = \mathfrak{P}_{\mathcal{A}}(q)$. “ \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:

$$\{r \in [1.. \#q] \mid \exists s \in \mathcal{G}^{j-1}(\top)(q) . x_s^{[e]} \approx_{\theta_j(t_j)} x_r^{[p_j]}\} \subsetneq \mathcal{G}^{j-1}(\top)(q)$$

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 $\theta'' \in \mathcal{R}_q^\infty(\mathcal{A})$. 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 Def. 18, 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 $x_i \approx_{\theta(\bar{i})} x_j$, there \bar{i} 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 Def. 17. 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 Def. 18.

(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 Def. 14. Also $q \in \bullet(\delta^\infty)$ thus, by Lemma 9, 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 $x_j^{[e]} \approx_{\theta^0} y_k^{[p_1]}$, for all $j, k \in \mathfrak{P}_{\mathcal{A}}(q_0)$ and $x_j^{[e]} \not\approx_{\theta^0} y_k^{[p_1]}$, for all $j, k \in [1.. \#q_0]$, $k \notin \mathfrak{P}_{\mathcal{A}}(q_0)$. Moreover, there exists another context $\theta_{p_2 \leftarrow q}^2 \in \mathcal{R}_{q_0}^\infty(\mathcal{A})$. 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:

- θ^1 on some position $p_0 \in \text{fr}(\theta^0)$ such that $\theta^0(p_0) = q_0$; such a position exists by the choice of θ^0 ,
- θ^2 on the position $p_0 p_2$, and
- 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 (Def. 14).

It is easy to check that $\bar{\theta}$ satisfies condition (2) of Def. 18. 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).

Lemma 14. *Let \mathcal{A} be a trim 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 (Proof of Lemma 14). Let θ be an arbitrary accepting run of \mathcal{A} . By Lemma 8, 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 9. Since q_i^0 occurs on θ , we can insert in θ a new partial run $\theta_i^1 \in \mathcal{R}_{q_i^0}^\infty(\mathcal{A})$ defined next. By Lemma 13 (1), there exists a q_i^0 -reset sequence $\theta_i^1 \in \mathcal{R}_{q_i^0}^\infty(\mathcal{A})$. The partial run θ_i^1 is obtained by composing θ_i^0 with itself k times, followed by θ_i^0 . These compositions are possible, because q_i^0 occurs at the root of θ_i^0 and θ_i^1 , as well as the frontier of θ_i^0 . Depending on which condition of Lemma 9 is satisfied by q_i^0 and θ_i^0 , we distinguish the following cases (see Fig. 6 for an illustration):

- Condition (1) In this case q_i and q_0^i occur on different positions on the frontier of θ_0^i , thus we place θ_i on the position of q_i and the rest of θ on the position of q_0^i .
- Condition (2) 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 9. From this second occurrence of q_0^i , we continue with θ .

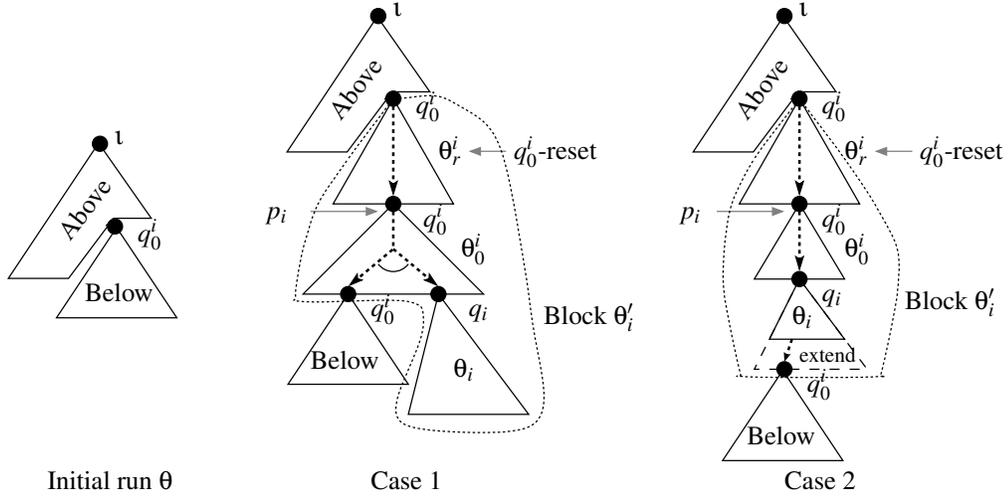


Fig. 6. Embedding of a partial run θ_i in θ .

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

- (1) All runs $\theta_1, \dots, \theta_n$ are inserted into θ at positions p_1, \dots, p_n , respectively, by construction.
- (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 $k \in [1..n]$.

D Material from Section 5

Lemma 15. *If \mathcal{A} simulates (resp. refines) \mathcal{B} then $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{B})$ (resp. $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{B})$).*

Proof (Proof of Lemma 15). Let $\mathcal{A} = (\mathbb{A}, Q_{\mathcal{A}}, \iota_{\mathcal{A}}, \delta_{\mathcal{A}})$, $\mathcal{B} = (\mathbb{A}, Q_{\mathcal{B}}, \iota_{\mathcal{B}}, \delta_{\mathcal{B}})$ and $h: Q_{\mathcal{A}} \rightarrow Q_{\mathcal{B}}$ 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 Def. 19.

“ \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 structural induction on t , using points (1) and (3) of Def. 19. \square

Lemma 16. *Let ϕ be a qpf formula, x_1, x_2, \dots, x_k variables and r a relation symbol of arity k , such that $\phi * x_1 \neq x_2$ and $\phi * r(x_1, \dots, x_k)$ are satisfiable. Then, we have:*

1. $\text{tw}(\llbracket (\phi * x_1 = x_2)^{\exists} \rrbracket) \leq \text{tw}(\llbracket \phi^{\exists} \rrbracket)$,
2. $\text{tw}(\llbracket \phi^{\exists} \rrbracket) - 1 \leq \text{tw}(\llbracket \phi * x_1 \neq x_2^{\exists} \rrbracket) \leq \text{tw}(\llbracket \phi^{\exists} \rrbracket)$,
3. $\text{tw}(\llbracket \phi^{\exists} \rrbracket) - 1 \leq \text{tw}(\llbracket \phi * r(x_1, \dots, x_k)^{\exists} \rrbracket) \leq \text{tw}(\llbracket \phi^{\exists} \rrbracket) + k$

Proof (Proof of Lemma 16). The proof follows a generic guideline. First, recall that for any set of structures \mathcal{S} we have $\text{tw}(\mathcal{S}) = \max_{S \in \mathcal{S}} \text{tw}(S) = \max_{S \in \mathcal{S}} \min\{\text{wd}(T) \mid T \text{ is a tree decomposition of } S\}$. 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 S and store \mathfrak{s} with $S \models^{\mathfrak{s}} \phi$

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

$\text{tw}(S) \leq \text{tw}(S') + k$ or

for all tree decomposition T' of S'

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

(1) The first point is immediate since $\llbracket \phi * x_1 = x_2^{\exists} \rrbracket \subseteq \llbracket \phi^{\exists} \rrbracket$.

(2) Since $\llbracket \phi * x_1 \neq x_2^{\exists} \rrbracket \subseteq \llbracket \phi^{\exists} \rrbracket$, we immediately obtain $\text{tw}(\llbracket \phi * x_1 \neq x_2^{\exists} \rrbracket) \leq \text{tw}(\llbracket \phi^{\exists} \rrbracket)$. For the other inequality, recall that $\phi * x_1 \neq x_2$ is satisfiable. Let $S = (U, \sigma)$ be a model and \mathfrak{s} a store such that $S \models^{\mathfrak{s}} \phi$. We distinguish two cases:

- if $\mathfrak{s}(x_1) \neq \mathfrak{s}(x_2)$ then let $S' = S$, $\mathfrak{s}' = \mathfrak{s}$ hence $S' \models^{\mathfrak{s}'} \phi * x_1 \neq x_2$ and $\text{tw}(S) = \text{tw}(S')$.
- if $\mathfrak{s}(x_1) = \mathfrak{s}(x_2)$ then let consider a new fresh element $e \in U$ and define a new store \mathfrak{s}' by $\mathfrak{s}'(y) \stackrel{\text{def}}{=} e$ if ϕ implies $y = x_2$, and $\mathfrak{s}'(y) \stackrel{\text{def}}{=} \mathfrak{s}(y)$ otherwise. Let moreover define the new structure $S' = (U \cup \{e\}, \sigma')$ 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 $S' \models^{\mathfrak{s}'} \phi * x_1 \neq x_2$. Let T' be a tree decomposition of S' . We define T by removing the element e from T' and adding $\mathfrak{s}(x_1)$ in every node of T' . T is a tree decomposition of S of width at most $\text{wd}(T') + 1$. Therefore $\text{wd}(T) \leq \text{wd}(T') + 1$, hence the result.

(3) Recall $\phi * r(x_1, \dots, x_k)$ is satisfiable from the hypothesis. We first prove that $\text{tw}(\llbracket \phi * r(x_1, \dots, x_k)^{\exists} \rrbracket) \leq \text{tw}(\llbracket \phi^{\exists} \rrbracket) + k$. Let $S = (U, \sigma)$ and \mathfrak{s} such that $S \models^{\mathfrak{s}} \phi * r(x_1, \dots, x_k)$. We define $S' = (U, \sigma')$ from 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 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 components 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 S of width at most $\text{wd}(T') + k$. T' and S were chosen arbitrary, hence the result.

We now prove that $\text{tw}(\llbracket \phi^{\exists} \rrbracket) \leq \text{tw}(\llbracket \phi * r(x_1, \dots, x_k)^{\exists} \rrbracket) + 1$. Let $S = (U, \sigma) \models^{\mathfrak{s}} \phi$, 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 S' obtained by adding the above tuple to $\sigma(r)$. Then $S' \models^{\mathfrak{s}'} \phi * r(x_1, \dots, x_k)$ and for any tree decomposition T' of S' we have $T = T'$ is also a tree decomposition for 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 thanks to the satisfiability hypothesis, this tuple is not “forced” by the formula ϕ , that is, 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 the equality $x_j = x'_j$ is not implied by ϕ .

We define a new store \mathfrak{s}' by $\mathfrak{s}'(y) = e$ if ϕ implies $y = x'_j$, and $\mathfrak{s}'(y) = \mathfrak{s}(y)$ otherwise, with $e \in \mathbb{U}$ a fresh new element. Let $S' = (\mathbb{U} \cup \{e\}, \sigma')$ with $\sigma'(r') \stackrel{\text{def}}{=} \sigma(r')$ for any $r' \neq r$, and $\sigma'(r)$ defined as follows. For every tuple $\langle u_1, \dots, u_k \rangle \in \sigma(r)$, there exists a unique relation atom $r(y_1, \dots, y_k)$ occurring in ϕ such that $\mathfrak{s}(y_i) = u_i$ for all $i \in [1..k]$. Then $\sigma'(r)$ contains the tuple $\langle \mathfrak{s}'(y_1), \dots, \mathfrak{s}'(y_k) \rangle$. Finally, add the tuple $\langle \mathfrak{s}'(x_1), \dots, \mathfrak{s}'(x_k) \rangle$ to $\sigma'(r)$. By construction $S' \models^{\mathfrak{s}'} \phi * r(x_1, \dots, x_k)$.

Let T' a tree decomposition of 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 S of width at most $\text{wd}(T') + 1$, that is, $\text{wd}(T) \leq \text{wd}(T') + 1$.

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

Lemma 17. (1) $\mathcal{A}_{\Delta, A}^I$ is all-satisfiable. (2) $\llbracket \mathcal{A}_{\Delta, A} \rrbracket = \llbracket \mathcal{A}_{\Delta, A}^I \rrbracket$.

Proof (Proof of Lemma 17). (1) To show that $\mathcal{A}_{\Delta, A}^I$ is all-satisfiable, we prove the following fact:

For all states $q_0 \in Q$, base tuples $(\sigma^\#, \pi) \in \text{SatBase}(x_1^{[e]}, \dots, x_{\#q_0}^{[e]})$, trees $t \in \mathcal{L}_{(q_0, (\sigma^\#, \pi))}(\mathcal{A}_{\Delta, A}^I)$, stores \mathfrak{s} injective over $x_1^{[e]}, \dots, x_{\#q_0}^{[e]}$, there exists a store $\bar{\mathfrak{s}}$, that agrees with \mathfrak{s} over $x_1^{[e]}, \dots, x_{\#q_0}^{[e]}$, and a structure $S = (\mathbb{U}, \sigma)$, such that $S \models^{\bar{\mathfrak{s}}} \Theta(t)$ and $\langle x_{j_1}^{[e]}, \dots, x_{j_{\#r}}^{[e]} \rangle \in \sigma^\#(r) \iff \langle \mathfrak{s}(x_{j_1}^{[e]}), \dots, \mathfrak{s}(x_{j_{\#r}}^{[e]}) \rangle \in \sigma(r)$, for all $r \in \mathbb{R}$ and $j_1, \dots, j_{\#r} \in [1..\#q_0]$.

The proof goes by induction on the structure of t . Since $t \in \mathcal{L}_{(q_0, (\sigma^\#, \pi))}(\mathcal{A}_{\Delta, A}^I)$, there exists a transition:

$$(q_0, (\sigma^\#, \pi)) \xrightarrow{t(\varepsilon)} \langle (q_1, (\sigma_1^\#, \pi_1)), \dots, (q_\ell, (\sigma_\ell^\#, \pi_\ell)) \rangle \in \delta'$$

such that $t|_i \in \mathcal{L}_{(q_i, (\sigma_i^\#, \pi_i))}(\mathcal{A}_{\Delta, A}^I)$, for all $i \in [1..\ell]$. By the definition of $\mathcal{A}_{\Delta, A}^I$, the formula $t(\varepsilon)$ is satisfiable, hence there exists a store \mathfrak{s}' and a structure $S_0 = (\mathbb{U}_0, \sigma_0)$ such that $S_0 \models^{\mathfrak{s}'} t(\varepsilon)$. Because the symbols in Σ contain no equalities between the variables $x_i^{[e]} \in \text{fv}(\Theta(t))$, we can assume w.l.o.g. the following:

- \mathfrak{s}' agrees with \mathfrak{s} over $x_1^{[e]}, \dots, x_{\#q_0}^{[e]}$, and
- $\mathfrak{s}'(z_1^{[p_1]}) \neq \mathfrak{s}'(z_2^{[p_2]})$, for all $z_i^{[p_i]} \in \text{fv}(\Theta(t))$, $i = 1, 2$, such that $z_1^{[p_1]} \not\approx_{\Theta(t)} z_2^{[p_2]}$ (\dagger).

For each $i \in [1..\ell]$, we consider a store \mathfrak{s}_i such that $\mathfrak{s}_i(x_j^{[e]}) = \mathfrak{s}'(x_j^{[e]})$, for all $j \in [1..\#q_i]$. Then \mathfrak{s}_i is injective over $x_1^{[e]}, \dots, x_{\#q_i}^{[e]}$, by the (\dagger) assumption on \mathfrak{s}' . Since $t|_i \in \mathcal{L}_{(q_i, (\sigma_i^\#, \pi_i))}(\mathcal{A}_{\Delta, A}^I)$, by the inductive hypothesis, there exists a store $\bar{\mathfrak{s}}_i$, that agrees with \mathfrak{s}_i over $x_1^{[e]}, \dots, x_{\#q_i}^{[e]}$,

and a structure $S_i = (U_i, \sigma_i)$, such that $S_i \models^{\bar{s}_i} \Theta(t|_i)$ and $\langle x_{j_1}^{[e]}, \dots, x_{j_{\#r}}^{[e]} \rangle \in \sigma_i^\#(r) \iff \langle \mathfrak{s}_i(x_{j_1}^{[e]}), \dots, \mathfrak{s}_i(x_{j_{\#r}}^{[e]}) \rangle \in \sigma(r)$, for all $r \in \mathbb{R}$ and $j_1, \dots, j_{\#r} \in [1.. \#q_i]$. Because U_1, \dots, U_ℓ are infinite, we can chose $\bar{s}_1, \dots, \bar{s}_\ell$ such that $\text{img}(\bar{s}_i) \cap \text{img}(\bar{s}_j) \subseteq \text{img}(\mathfrak{s}')$, for all $1 \leq i < j \leq \ell$. We prove the following:

Fact 7 *The structures S_0, \dots, S_ℓ are locally disjoint.*

Proof. Suppose, for a contradiction, that there exists a tuple $\langle u_1, \dots, u_{\#r} \rangle \in \sigma_i(r) \cap \sigma_j(r)$, for some $r \in \mathbb{R}$ and $1 \leq i < j \leq \ell$. Then $u_1, \dots, u_{\#r} \in \text{img}(\mathfrak{s}')$, hence there exist indices $k_{1,1}, \dots, k_{1,\#r} \in [1.. \#q_i]$ and $k_{2,1}, \dots, k_{2,\#r} \in [1.. \#q_j]$, such that $u_h = \mathfrak{s}'(x_{k_{1,h}}^{[i]}) = \mathfrak{s}'(x_{k_{2,h}}^{[j]})$, for all $h \in [1.. \#r]$. By the assumption (\dagger) on \mathfrak{s}' , we obtain $x_{k_{1,h}}^{[i]} \approx_{\Theta(r)} x_{k_{2,h}}^{[j]}$, for all $h \in [1.. \#r]$. Since $\langle u_1, \dots, u_{\#r} \rangle \in \sigma_i(r) \cap \sigma_j(r)$, we obtain $\langle x_{k_{1,1}}^{[i]}, \dots, x_{k_{1,\#r}}^{[i]} \rangle \in \sigma_i^\#(r)$ and $\langle x_{k_{2,1}}^{[j]}, \dots, x_{k_{2,\#r}}^{[j]} \rangle \in \sigma_j^\#(r)$. Then $(\sigma^\#, \pi)$ is not satisfiable, by the definition of $\mathcal{A}_{\Delta, A}^I$, contradiction. A similar contradiction is obtained by considering $i = 0$. \square

Back to the proof, the structure $S = S_0 \bullet \dots \bullet S_\ell$ is well defined. Let \bar{s} be the store:

- $\bar{s}(z^{[e]}) \stackrel{\text{def}}{=} \mathfrak{s}'(z^{[e]}(t))$, for all $z^{[e]} \in \text{fv}(\Theta(t))$, and
- $\bar{s}(z^{[i]p}) \stackrel{\text{def}}{=} \bar{s}_i(z^{[i]p})$, for all $i \in [1.. \ell]$ and $z^{[i]p} \in \Theta(t|_i)$.

Hence we obtain $S \models^{\bar{s}} \Theta(t)$. Moreover, for all $r \in \mathbb{R}$ and indices $j_1, \dots, j_{\#r}$, we have:

$$\begin{aligned} \langle x_{j_1}, \dots, x_{j_{\#r}} \rangle \in \sigma^\#(r) &\iff \langle x_{j_1}, \dots, x_{j_k} \rangle \in \sigma_0^\#(r) \uplus \biguplus_{i=1}^{\ell} \sigma_i^\#[x_1/z_{i,1}, \dots, x_{\#q_i}/z_{i,\#q_i}](r) \\ &\iff \langle \mathfrak{s}'(x_{j_1}), \dots, \mathfrak{s}'(x_{j_{\#r}}) \rangle \in \sigma_0(r) \uplus \biguplus_{i=1}^{\ell} \sigma_i(r) \\ &\iff \langle \mathfrak{s}(x_{j_1}), \dots, \mathfrak{s}(x_{j_{\#r}}) \rangle \in \sigma(r) \end{aligned}$$

(2) We prove that $\llbracket \mathcal{A}_{\Delta, A} \rrbracket = \llbracket \mathcal{A}_{\Delta, A}^I \rrbracket$. The mapping $h : Q_\Delta^I \rightarrow Q_\Delta$ defined as $h((q, (\sigma^\#, \pi))) = q$ is a simulation, by the definition of $\mathcal{A}_{\Delta, A}^I$. Hence $\mathcal{L}(\mathcal{A}_{\Delta, A}^I) \subseteq \mathcal{L}(\mathcal{A}_{\Delta, A})$, by Lemma 15, leading to $\llbracket \mathcal{A}_{\Delta, A}^I \rrbracket \subseteq \llbracket \mathcal{A}_{\Delta, A} \rrbracket$. For the $\llbracket \mathcal{A}_{\Delta, A}^I \rrbracket \supseteq \llbracket \mathcal{A}_{\Delta, A} \rrbracket$ direction, we show the following:

For all states $q_0 \in Q$, trees $t \in \mathcal{L}_{q_0}(\mathcal{A}_{\Delta, A})$, structures $S = (U, \sigma)$ and stores \mathfrak{s} , such that $S \models^{\mathfrak{s}} \Theta(t)$, there exists a base tuple $(\sigma^\#, \pi)$ such that $t \in \mathcal{L}_{(q_0, (\sigma^\#, \pi))}(\mathcal{A}_{\Delta, A}^I)$ and, for all relations $r \in \mathbb{R}$ and indices $j_1, \dots, j_k \in [1.. \#q_0]$, if $\langle x_{j_1}^{[e]}, \dots, x_{j_{\#r}}^{[e]} \rangle \in \sigma^\#(r)$ then $\langle \mathfrak{s}(x_{j_1}^{[e]}), \dots, \mathfrak{s}(x_{j_{\#r}}^{[e]}) \rangle \in \sigma(r)$.

The proof is by induction on the structure of t . Since $t \in \mathcal{L}_{q_0}(\mathcal{A}_{\Delta, A})$, there exists a transition $q_0 \xrightarrow{t(\varepsilon)} (q_1, \dots, q_\ell) \in \delta$ such that $t|_i \in \mathcal{L}_{q_i}(\mathcal{A}_{\Delta, A})$ for all $i \in [1.. \ell]$. Because $S \models^{\mathfrak{s}} \Theta(t)$, we can decompose the structure as $S = S_0 \bullet \dots \bullet S_\ell$, such that $S_0 \models^{\mathfrak{s}} t(\varepsilon)$ and $S_i \models^{\mathfrak{s}} \Theta(t|_i)^{[i]}$ for all $i \in [1.. \ell]$. We can assume w.l.o.g. that $S_i = (U, \sigma_i)$, for all $i \in [0.. \ell]$. Let us consider the stores \mathfrak{s}_i , such that $\mathfrak{s}_i(z^{[i]p}) \stackrel{\text{def}}{=} \mathfrak{s}(z^{[i]p})$, for all $z^{[i]p} \in \text{fv}(\Theta(t|_i))$ and $i \in [1.. \ell]$. Then $S_i \models^{\mathfrak{s}_i} \Theta(t|_i)$ and, by the inductive hypothesis on $t|_i$, there exists a satisfiable base tuple $(\sigma_i^\#, \pi_i)$, such that $t|_i \in \mathcal{L}_{(q_i, (\sigma_i^\#, \pi_i))}(\mathcal{A}_{\Delta, A}^I)$ and $\langle x_{j_1}^{[e]}, \dots, x_{j_k}^{[e]} \rangle \in \sigma_i^\#(r)$

implies $\langle \mathfrak{s}_i(x_{j_1}^{[e]}), \dots, \mathfrak{s}_i(x_{j_k}^{[e]}) \rangle \in \sigma_i(r)$, for all $r \in \mathbb{R}$ and $j_1, \dots, j_k \in [1.. \#q_i]$. Let us denote $(\sigma_0^\#, \pi_0) \stackrel{\text{def}}{=} \text{Base}(t(\varepsilon))$ and define the base tuple:

$$(\sigma^\#, \pi) \stackrel{\text{def}}{=} \left((\sigma_0^\#, \pi_0) \otimes \bigotimes_{i=1}^{\ell} (\sigma_i^\#, \pi_i) [x_1^{[e]}/x_1^{[i]}, \dots, x_{\#q_i}^{[e]}/x_{\#q_i}^{[i]}] \right) \downarrow_{\{x_1^{[e]}, \dots, x_{\#q_0}^{[e]}\}}$$

First we prove that $(\sigma^\#, \pi)$ is satisfiable. The substitutions do not change the satisfiability of the base tuples because they are injective. The projection preserves satisfiability as well. Hence the only possibilities for $(\sigma^\#, \pi)$ being unsatisfiable are the following:

- there are two tuples $\langle z_{1,1}, \dots, z_{1,\#r} \rangle, \langle z_{1,1}, \dots, z_{1,\#r} \rangle \in \sigma^\#(r)$, such that $z_{1,j} \approx_{\pi_0} z_{2,j}$, for all $j \in [1.. \#r]$. Suppose, for a contradiction, that $\langle z_{1,1}, \dots, z_{1,\#r} \rangle, \langle z_{1,1}, \dots, z_{1,\#r} \rangle \in \sigma_i^\#(r)$, for some $i \in [0.. \ell]$. Then $(\sigma_i^\#, \pi_i)$ is not satisfiable. If $i \in [1.. \ell]$, this contradicts the inductive hypothesis. Else $i = 0$, which contradicts the satisfiability of $t(\varepsilon)$, i.e., $S_0 \models^s t(\varepsilon)$. Then it must be the case that $\langle z_{k,1}, \dots, z_{k,\#r} \rangle \in \sigma_{i_k}^\#(r)$, for $k = 1, 2$, for some $i_1 \neq i_2 \in [0.. \ell]$. Consider the case $i_1 \in [1.. \ell]$ and $i_2 \in [0.. \ell]$, the other case being symmetric. By the inductive hypothesis, we have $\langle \mathfrak{s}(z_{1,1}), \dots, \mathfrak{s}(z_{1,\#r}) \rangle \in \sigma_{i_1}(r)$. If $i_2 = 0$, we have $\langle \mathfrak{s}(z_{2,1}), \dots, \mathfrak{s}(z_{2,\#r}) \rangle \in \sigma_0(r)$, by the definition of $(\sigma_0^\#, \pi_0)$, otherwise $i_2 \in [1.. \ell]$ and $\langle \mathfrak{s}(z_{2,1}), \dots, \mathfrak{s}(z_{2,\#r}) \rangle \in \sigma_{i_2}(r)$, by the inductive hypothesis. In both cases, we obtain a contradiction with the satisfiability of $\Theta(t)$, i.e., $S \models^s \Theta(t)$.
- π is unsatisfiable, because there are variables $z_1, z_2 \in \bigcup_{i=0}^{\ell} \text{fv}(\pi_i)$, such that $z_1 \approx_{\pi_{i_1}} z_2$ and $z_1 \neq z_2$ is asserted in π_{i_2} , for some $i_1, i_2 \in [0.. \ell]$. A similar reasoning as above leads to a contradiction with the satisfiability of $\Theta(t)$, i.e., $S \models^s \Theta(t)$.

Since $(\sigma^\#, \pi)$ is satisfiable, by the definition of $\mathcal{A}_{\Delta, A}^I$, there exists a transition:

$$(q_0, (\sigma^\#, \pi)) \xrightarrow{t(\varepsilon)} ((q_1, (\sigma_1^\#, \pi_1)), \dots, (q_\ell, (\sigma_\ell^\#, \pi_\ell))) \in \delta^{\text{sat}}$$

which proves $t \in \mathcal{L}_{(q_0, (\sigma^\#, \pi))}(\mathcal{A}_{\Delta, A}^I)$. To establish the second point of the inductive invariant, let $\langle x_{j_1}^{[e]}, \dots, x_{j_k}^{[e]} \rangle \in \sigma^\#(r)$ be a tuple. If $\langle x_{j_1}^{[e]}, \dots, x_{j_{\#r}}^{[e]} \rangle \in \sigma_0^\#(r)$ then $\langle \mathfrak{s}(x_{j_1}^{[e]}), \dots, \mathfrak{s}(x_{j_{\#r}}^{[e]}) \rangle \in \sigma_0(r) \subseteq \sigma(r)$. Else, there exists $k \in [1.. \ell]$ and a tuple $\langle x_{h_1}^{[e]}, \dots, x_{h_{\#r}}^{[e]} \rangle \in \sigma_k^\#(r)$, such that $x_{j_i}^{[e]} \approx_{\pi} x_{h_i}^{[k]}$, for all $i \in [1.. \#r]$. By the inductive hypothesis, we have $\langle \mathfrak{s}(x_{h_1}^{[k]}), \dots, \mathfrak{s}(x_{h_{\#r}}^{[k]}) \rangle \in \sigma_k(r) \subseteq \sigma(r)$, thus we obtain $\langle \mathfrak{s}(x_{j_1}^{[e]}), \dots, \mathfrak{s}(x_{j_{\#r}}^{[e]}) \rangle \in \sigma(r)$.

Lemma 18. (1) $\mathcal{A}_{\Delta, A}^{II}$ is all-satisfiable. (2) $\llbracket \mathcal{A}_{\Delta, A}^{II} \rrbracket$ is treewidth-bounded iff $\llbracket \mathcal{A}_{\Delta, A}^I \rrbracket$ is treewidth-bounded.

Proof (Proof of Lemma 18). (1) Let $t \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{II})$ be a tree. Since $\mathcal{A}_{\Delta, A}^{II}$ was obtained from $\mathcal{A}_{\Delta, A}^I$ by removing relation atoms from the labels of its 1-transitions, there exists a tree $t' \in \mathcal{L}(\mathcal{A}_{\Delta, A}^I)$, such that $\text{dom}(t) = \text{dom}(t')$ and $\Theta(t) = \Theta(t') * \bigstar_{i=1}^n r_i(z_{i,1}, \dots, z_{i,k_i})$ modulo reordering of atoms, for some relation symbols r_i and variables $z_{i,1}, \dots, z_{i,k_i}$. By Lemma 17, $\Theta(t')$ is satisfiable, hence there exists a structure (U, σ') and a store \mathfrak{s} , such that $(U, \sigma') \models^s \Theta(t')$. 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 $(U, \sigma') \models^s \Theta(t)$, hence $\Theta(t)$ is satisfiable. Since the choice of t was arbitrary, we obtain that $\mathcal{A}_{\Delta, A}^{II}$ is all-satisfiable.

(2) “ \Leftarrow ” Let $W \geq 1$ be an integer such that $\text{tw}(S) \leq W$ for all structures S and all trees $t \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{II})$, such that $S \models \Theta(t)^\exists$. Since $\llbracket \mathcal{A}_{\Delta, A}^{II} \rrbracket$ is treewidth-bounded, such an

integer exists. Let $t \in \mathcal{L}(\mathcal{A}_{\Delta, A}^I)$ be an arbitrary tree. By the construction of $\mathcal{A}_{\Delta, A}^{II}$ from $\mathcal{A}_{\Delta, A}^I$, there exists a tree $t' \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{II})$, such that $\text{dom}(t') = \text{dom}(t)$ and $\Theta(t) = \Theta(t') * \bigstar_{i=1}^n r_i(z_{i,1}, \dots, z_{i,k_i})$ modulo reordering of atoms, for some relation symbols r_i and variables $z_{i,1}, \dots, z_{i,k_i}$. By Lemma 8, each 1-transition of $\mathcal{A}_{\Delta, A}^I$ occurs exactly once in each accepting run, hence $n \leq \text{card}(\delta_{\Delta}^1) \cdot M$, where $M \stackrel{\text{def}}{=} \max\{\text{size}(\alpha) \mid q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell) \in \delta_{\Delta}^1\}$. Moreover, $k_i \leq M$, for all $i \in [1..n]$. By Lemma 16 (3), $\text{tw}(\Theta(t)^{\exists}) \leq \text{tw}(\Theta(t')^{\exists}) + \sum_{i=1}^n k_i \leq \text{card}(\delta_{\Delta}^1) \cdot M^2$. Since the choice of t was arbitrary, we obtain that $\llbracket \mathcal{A}_{\Delta, A}^I \rrbracket$ is treewidth-bounded. “ \Rightarrow ” This direction uses a symmetric argument.

Lemma 19. *Let $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell)$ be a 1-transition of $\mathcal{A}_{\Delta, A}^{III}$. Then, for each $i \in [1..\ell]$ and each $j \in \mathfrak{F}_{\mathcal{A}_{\Delta, A}^{III}}(q_i)$, there exists a ε -variable z , such that $x_j^{[i]} \approx_{\alpha} z$.*

Proof (Proof of Lemma 19). By Lemma 11 and the fact that equalities involving persistent i -variables $x_j^{[i]}$, for $i \in [1..\ell]$ and $j \in [1..\#q_i]$ are either the same in the labels of the 1-transitions $q_0 \xrightarrow{\alpha} (q_1, \dots, q_\ell)$ of $\mathcal{A}_{\Delta, A}^{III}$ and $\mathcal{A}_{\Delta, A}^I$, or replaced by $x_j^{[i]} = y_k^{[i]}$ for some $k \leq 1$.

Lemma 20. *Let ϕ and ψ be qpf formulae, such that $\phi * \psi$ is satisfiable and $x \not\approx_{\phi} y$, for all $x, y \in \text{fv}(\phi) \cap \text{fv}(\psi)$. Let $\psi_{eq} = \bigstar \{x = y \mid x, y \in \text{fv}(\phi) \cap \text{fv}(\psi), x \approx_{\psi} y\}$. Then, $\text{tw}(\llbracket \phi * \psi_{eq} \rrbracket^{\exists}) \leq \text{tw}(\llbracket \phi * \psi \rrbracket^{\exists}) + \text{card}(\text{fv}(\phi) \cap \text{fv}(\psi))$.*

Proof (Proof of Lemma 20). Let denote by $V_f \stackrel{\text{def}}{=} \text{fv}(\phi) \cap \text{fv}(\psi)$, that is, the set of “frontier” variables between ϕ and ψ . Let observe that $\text{fv}(\psi_{eq}) \subseteq V_f \subseteq \text{fv}(\phi)$ and henceforth, $\text{fv}(\phi * \psi_{eq}) = \text{fv}(\phi)$. We follow a similar idea as in the proof of lemma 16. Consider a structure $S = (U, \sigma)$ and a store \mathfrak{s} on $\text{fv}(\phi)$ such that $S \models^{\mathfrak{s}} \phi * \psi_{eq}$. We build a store \mathfrak{s}' and a structure S' such that $S' \models^{\mathfrak{s}'} \phi * \psi$. First, consider a store \mathfrak{s}_{ψ} canonical for ψ and the structure $S_{\psi} = (U_{\psi}, \sigma_{\psi})$ such that $S_{\psi} \models^{\mathfrak{s}_{\psi}} \psi$. Assume without loss of generality that S_{ψ} and S are disjoint structures, that is, $\mathfrak{s}_{\psi}(\text{fv}(\psi)) \cap \mathfrak{s}(\text{fv}(\phi)) = \emptyset$. Second, define the store \mathfrak{s}' on $\text{fv}(\phi) \cup \text{fv}(\psi)$ by taking:

$$\mathfrak{s}'(y) \stackrel{\text{def}}{=} \begin{cases} \mathfrak{s}_{\psi}(y) & \text{if } y \in \text{fv}(\psi) \\ \mathfrak{s}_{\psi}(y') & \text{if } y \in \text{fv}(\phi) \setminus \text{fv}(\psi), y' \in V_f, y \approx_{\phi} y' \\ \mathfrak{s}(y) & \text{otherwise} \end{cases}$$

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 $y' \in V_f$ such that $y \approx_{\phi} y'$ (otherwise, the hypothesis $x \not\approx_{\phi} y$ for all $x, y \in V_f$ does not hold). We build now the structure $S' = (U', \sigma')$ where $U' \stackrel{\text{def}}{=} U \cup U_{\psi}$ and $\sigma'(r)$ is defined for every relation $r \in \mathbb{R}$ as follows:

- for every tuple $\langle u_1, \dots, u_{\#r} \rangle \in \sigma_{\psi}(r)$ add it 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 $\phi * \psi_{eq}$ such that $\mathfrak{s}(y_i) = u_i$ for all $i \in [1..\#r]$; in this case, add the tuple $\langle \mathfrak{s}'(y_1), \dots, \mathfrak{s}'(y_{\#r}) \rangle$ to $\sigma'(r)$.

This construction guarantees that $S' \models^{\mathfrak{s}'} \phi * \psi$. With regard to 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 $\sigma_{\psi}(r)$ then $\phi * \psi$ cannot be satisfiable. Let T' be an optimal tree decomposition of S' , i.e., $\text{tw}(S') = \text{wd}(T')$. We define a tree decomposition T by:

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

The result is a tree decomposition T of S of width $\text{wd}(T) \leq \text{wd}(T') + \text{card}(V_f) = \text{tw}(S') + \text{card}(V_f)$, thus $\text{tw}(S) \leq \text{tw}(S') + \text{card}(V_f)$. Since the choice of S was arbitrary, we obtain $\text{tw}(\llbracket(\phi * \psi_{eq})^\exists\rrbracket) \leq \text{tw}(\llbracket(\phi * \psi)^\exists\rrbracket) + \text{card}(V_f)$ as required.

Lemma 21. (1) $\mathcal{A}_{\Delta, A}^{III}$ is all-satisfiable. (2) $\llbracket\mathcal{A}_{\Delta, A}^{III}\rrbracket$ is treewidth-bounded iff $\llbracket\mathcal{A}_{\Delta, A}^{II}\rrbracket$ is treewidth-bounded.

Proof (Proof of Lemma 21). We assume the necessary changes in step 1 of the transformation have already been done, as they do not change the semantic of $\mathcal{A}_{\Delta, A}^{II}$.

(1) Let $t' \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{III})$ be a tree. Since $\mathcal{A}_{\Delta, A}^{III}$ was obtained from $\mathcal{A}_{\Delta, A}^{II}$ by removing equality atoms from the labels of its 1-transitions (step 2 of the transformation), there exists a tree $t \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{II})$, such that $\text{dom}(t) = \text{dom}(t')$ and $\Theta(t) = \Theta(t') * \psi$ where ψ is a conjunction of equality atoms. By Lemma 18 (1), $\Theta(t)$ is satisfiable, hence there exists a structure S and a store s such that $S \models^s \Theta(t)$. We immediately obtain $S \models^s \Theta(t')$, hence $\Theta(t')$ is satisfiable. Since the choice of t' was arbitrary, we obtain that $\mathcal{A}_{\Delta, A}^{III}$ is all-satisfiable.

(2) We compare $\max_{t \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{III})} \text{tw}(\llbracket\Theta(t)^\exists\rrbracket)$ with $\max_{t \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{II})} \text{tw}(\llbracket\Theta(t)^\exists\rrbracket)$.

“ \Leftarrow ” Assume that $\llbracket\mathcal{A}_{\Delta, A}^{III}\rrbracket$ is treewidth-bounded. Let $t \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{II})$. By the construction of $\mathcal{A}_{\Delta, A}^{III}$ from $\mathcal{A}_{\Delta, A}^{II}$, there exists a tree $t' \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{III})$ 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 16 (1), we get $\text{tw}(\llbracket\Theta(t)^\exists\rrbracket) = \text{tw}(\llbracket(\Theta(t') * \psi)^\exists\rrbracket) \leq \text{tw}(\llbracket\Theta(t')^\exists\rrbracket)$. Since t has been chosen arbitrary, we obtain that $\llbracket\mathcal{A}_{\Delta, A}^{II}\rrbracket$ is treewidth-bounded.

“ \Rightarrow ” Suppose $\llbracket\mathcal{A}_{\Delta, A}^{II}\rrbracket$ is treewidth-bounded by some constant $W \geq 1$. Let $t' \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{III})$ and θ' be an accepting run over t' . We will show that $\text{tw}(\Theta(t')) \leq W + K$, where K does not depend on the choice of t' .

We build a tree $t \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{II})$ related to t' as follows. The idea 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}^{II} to δ_{Δ}^{III} are not impactful and can be added back. To avoid unnecessary complications, we consider one 1-transition at a time (recall that there are finitely many 1-transitions in $\mathcal{A}_{\Delta, A}^{III}$). For one position $p \in \text{dom}(t')$ such that $\theta'(p) \xrightarrow{t'(p)}$ $(\theta'(p1), \dots, \theta'(p\ell)) \in (\delta_{\Delta}^{III})^1$, 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), such that $\theta_{p \leftarrow \theta(p)}^{\text{init}}(r) \stackrel{\text{def}}{=} \theta'(r)$ 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 $pir \in \text{dom}(t')$. We build t and the associated run θ by combining successive partial runs:

- Start θ with $\theta_{p \leftarrow \theta(p)}^{\text{init}}$ (no change from θ' above position p),
- Add a $\theta'(p)$ -reset $\theta_{u \leftarrow \theta'(p)} \in \mathcal{R}_{\theta'(p)}^{\infty}(\mathcal{A}_{\Delta, A}^{II})$ at position p ,
- Pursue at position pu with the 1-transition $\theta'(p) \xrightarrow{\alpha_p} (\theta'(p1), \dots, \theta'(p\ell)) \in (\delta_{\Delta}^{II})^1$ corresponding to the old 1-transition in $\mathcal{A}_{\Delta, A}^{III}$ after adding back the equalities removed by the transformation,

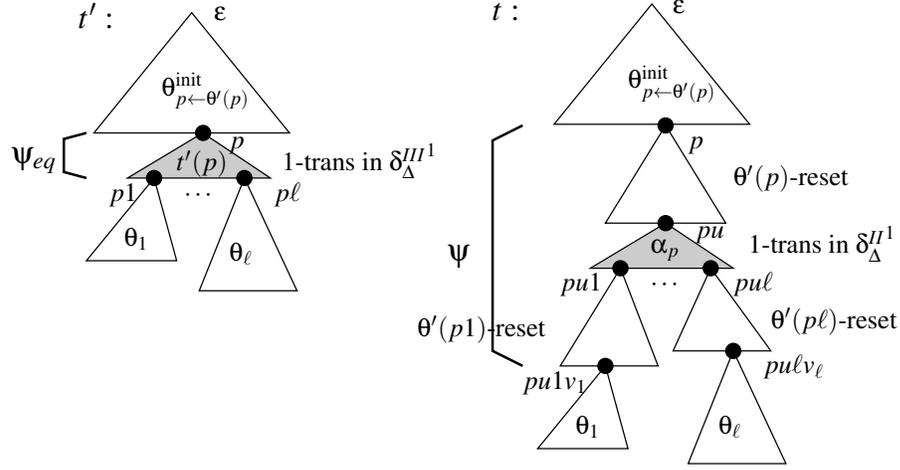


Fig. 7. Construction of t from t' (Lemma 21)

- For every $i \in [1..\ell]$, introduce at position pui a $\theta'(pi)$ -reset $\theta_{v_i \leftarrow \theta'(pi)} \in \mathcal{R}_{\theta'(pi)}^\infty(\mathcal{A}_{\Delta, A}^{II})$,
- Continue with θ_i at position $puiv_i$ for every $i \in [1..\ell]$.

Figure 7 illustrates this construction. We now prove that the transformation from t' to t only increases the treewidth of t' by a constant :

Fact 8 $\text{tw}(\llbracket \Theta(t')^{\exists} \rrbracket) \leq \text{tw}(\llbracket \Theta(t)^{\exists} \rrbracket) + k$ where k is the maximal number of i -variables, $i \in \mathbb{N} \cup \{\varepsilon\}$, in any 1-transition of $\mathcal{A}_{\Delta, A}^{III}$.

Proof. Let $\phi \stackrel{\text{def}}{=} \star_p$ not a prefix of r $t'(r)^{[r]} * \star_{pir \in \text{dom}(t')} t'(pir)^{[puiv_i; r]}$ 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 $V_f \stackrel{\text{def}}{=} \text{fv}(\phi) \cap \text{fv}(\psi)$ and $\Psi_{eq} \stackrel{\text{def}}{=} \star \{x = y \mid x, y \in V_f, x \approx_\psi y\}$. V_f contains parameter variables at the extremity of the added part from t' to t : $V_f = \{x_k^{[p]} \mid k \in [1..\#t(p)]\} \cup \{x_j^{[puiv_i]} \mid i \in [1..\ell], j \in [1..\#t(puiv_i)]\}$. These variables exactly correspond to the parameters $\{x_k^{[p]} \mid k \in [1..\#t'(p)]\} \cup \{x_j^{[pi]} \mid i \in [1..\ell], j \in [1..\#t'(pi)]\}$ appearing in the 1-transition at position p of t' , hence $\text{card}(V_f) \leq k$. By Def. 15 (2), ϕ does not induce equalities between variables of V_f . On the other hand, ψ only induce equalities between persistent variables of V_f , thanks to the reset paths. These equalities Ψ_{eq} exactly correspond to those occurring in $t'(p)$, thus $\phi * \Psi_{eq}$ is equal to $\Theta(t')$ after renaming the variables. Now $\phi * \psi$ is satisfiable (since $\mathcal{A}_{\Delta, A}^{II}$ is all-satisfiable), thus by Lemma 20, we obtain $\text{tw}(\llbracket \Theta(t')^{\exists} \rrbracket) = \text{tw}(\llbracket (\phi * \Psi_{eq})^{\exists} \rrbracket) \leq \text{tw}(\llbracket (\phi * \psi)^{\exists} \rrbracket) + k = \text{tw}(\llbracket \Theta(t)^{\exists} \rrbracket) + k$. \square

After doing this transformation (t' to t) for all 1-transitions, the final tree t satisfies $t \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{III})$ since all the added transitions (1-transition or reset) appear in δ_{Δ}^{III} . With the inequality at each step (Fact 8) and since each 1-transition of δ_{Δ}^{III} occurs exactly once

in the initial tree t' (Def. 8), we get $\text{tw}(\llbracket \Theta(t')^{\exists} \rrbracket) \leq \text{tw}(\llbracket \Theta(t)^{\exists} \rrbracket) + K \leq W + K$, where $K = k \times \text{card}((\delta_{\Delta}^{III})^1)$. We conclude that $\llbracket \mathcal{A}_{\Delta, A}^{III} \rrbracket$ is treewidth-bounded, as t' has been chosen arbitrary.

Lemma 22. (1) $\mathcal{A}_{\Delta, A}^{IV}$ is all-satisfiable. (2) $\llbracket \mathcal{A}_{\Delta, A}^{IV} \rrbracket$ is treewidth-bounded iff $\llbracket \mathcal{A}_{\Delta, A}^{III} \rrbracket$ is treewidth-bounded.

Proof (Proof of Lemma 22). Let \mathcal{A}° be the automaton obtained from $\mathcal{A}_{\Delta, A}^{III}$ by decorating the states $q \in Q_{\Delta}^{III}$ with mappings $a : \mathfrak{P}_{\mathcal{A}_{\Delta, A}^{III}} \rightarrow [1..M]$ and keeping the original labels on the transitions and the original arities of the states, i.e. $\#(q, a) = \#q$ in \mathcal{A}° . It is easy to check that \mathcal{A}° is a refinement of $\mathcal{A}_{\Delta, A}^{III}$ hence, by Lemma 15, we have $\mathcal{L}(\mathcal{A}^{\circ}) = \mathcal{L}(\mathcal{A}_{\Delta, A}^{III})$. Consequently, \mathcal{A}° is all-satisfiable, because $\mathcal{A}_{\Delta, A}^{III}$ is all-satisfiable. Moreover, $\llbracket \mathcal{A}^{\circ} \rrbracket$ is treewidth-bounded iff $\llbracket \mathcal{A}_{\Delta, A}^{III} \rrbracket$ is treewidth-bounded.

Fact 9 Let $t \in \mathcal{L}(\mathcal{A}^{\circ})$, θ be an accepting run of \mathcal{A}° 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 $i \in \mathfrak{P}_{\mathcal{A}^{\circ}}(q)$, such that r_i is the unique position where a variable $y_{a(i)}^{[e]} \in \mathcal{V}$ occurs in $\Theta(t)$.

Proof. By induction on the structure of t , using Lemma 19. □

(1) Let $\bar{t} \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{IV})$ be a tree. By the construction of $\mathcal{A}_{\Delta, A}^{IV}$, the accepting run $\bar{\theta}$ of $\mathcal{A}_{\Delta, A}^{IV}$ over \bar{t} can be transformed into an accepting run θ of \mathcal{A}° 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 , as in the construction of $\mathcal{A}_{\Delta, A}^{IV}$. Since \mathcal{A}° is all-satisfiable, the formula $\Theta(t)$ is satisfiable, and let (U, σ) be a structure and \mathfrak{s} be a store such that $(U, \sigma) \models^{\mathfrak{s}} \Theta(t)$. However, $\Theta(\bar{t})$ is obtained from $\Theta(t)$ by removing (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 $r_g(\eta(z_{i_1}^{[p]}), \dots, \eta(z_{i_k}^{[p]}))$, according to the construction of $\mathcal{A}_{\Delta, A}^{IV}$. Suppose, for a contradiction, that there exist two distinct positions $p_1, p_2 \in \text{dom}(\bar{t})$ such that the relation atoms $r_g(\eta(z_{i_1}^{[p_1]}), \dots, \eta(z_{i_k}^{[p_1]}))$ and $r_g(\eta(z_{i_1}^{[p_2]}), \dots, \eta(z_{i_k}^{[p_2]}))$ occur in $\Theta(\bar{t})$ and $z_{i_j}^{[p_1]} \approx_{\theta(\bar{t})} z_{i_j}^{[p_2]}$, for all $j \in [1..k]$. Let $\theta(p_i) = (q_i, a_i)$, for $i = 1, 2$. 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 9, 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 exist in $\Theta(\bar{t})$. Then, $\Theta(\bar{t})$ is satisfiable and, since the choice of \bar{t} was arbitrary, $\mathcal{A}_{\Delta, A}^{IV}$ is all-satisfiable.

(2) We prove that $\llbracket \mathcal{A}_{\Delta, A}^{IV} \rrbracket$ is treewidth-bounded iff $\llbracket \mathcal{A}^{\circ} \rrbracket$ is treewidth-bounded.

“ \Rightarrow ” Assume that $\llbracket \mathcal{A}_{\Delta, A}^{IV} \rrbracket$ is treewidth-bounded and let $S = (U, \sigma) \in \llbracket \mathcal{A}^{\circ} \rrbracket$ be a structure. Then, there exists a tree $t \in \mathcal{L}(\mathcal{A}^{\circ})$ and a store \mathfrak{s} , such that $S \models^{\mathfrak{s}} \Theta(t)$. Let $\bar{t} \in \mathcal{L}(\mathcal{A}_{\Delta, A}^{IV})$ be the tree obtained by changing each label α of t into $\bar{\alpha}$, according to the construction of $\mathcal{A}_{\Delta, A}^{IV}$. We consider the structure $\bar{S} \stackrel{\text{def}}{=} (U, \bar{\sigma})$, where $\bar{\sigma}$ interprets each relation symbol r_g by the set of tuples $\langle \mathfrak{s}(\eta(z_{i_1}^{[p]})), \dots, \mathfrak{s}(\eta(z_{i_k}^{[p]})) \rangle$, such that $r_g(\eta(z_{i_1}^{[p]}), \dots, \eta(z_{i_k}^{[p]}))$ occurs in $\Theta(\bar{t})$. Let \bar{T} be an optimal tree decomposition of \bar{S} , hence $\text{tw}(\bar{T}) \leq \text{tw}(\llbracket \mathcal{A}_{\Delta, A}^{IV} \rrbracket)$

(by the assumption, $\text{tw}(\llbracket \mathcal{A}_{\Delta, \mathbb{A}}^{IV} \rrbracket)$ is a finite integer). We consider the tree decomposition T obtained by adding the values $\mathfrak{s}(y_i^{[r_i]})$, where r_i is the unique position where a ε -variable $y_i^{[e]} \in \mathcal{Y}$ occurs in $\Theta(\bar{t})$, to the label of each node in \bar{T} . Then, T is a tree decomposition of S of width $\text{wd}(T) \leq \text{wd}(\bar{T}) + \mathcal{M} \leq \text{tw}(\llbracket \mathcal{A}_{\Delta, \mathbb{A}}^{IV} \rrbracket) + \mathcal{M}$. Consequently, $\text{tw}(S) \leq \text{tw}(\llbracket \mathcal{A}_{\Delta, \mathbb{A}}^{IV} \rrbracket) + \mathcal{M}$ and, since the choice of S was arbitrary, we obtain that $\llbracket \mathcal{A}^\circ \rrbracket$ is treewidth-bounded.

“ \Leftarrow ” Let $\bar{S} = (U, \bar{\sigma}) \in \llbracket \mathcal{A}_{\Delta, \mathbb{A}}^{IV} \rrbracket$ be a structure. Then, there exists a tree $\bar{t} \in \mathcal{L}(\mathcal{A}_{\Delta, \mathbb{A}}^{IV})$, hence also an accepting run $\bar{\theta}$ of $\mathcal{A}_{\Delta, \mathbb{A}}^{IV}$ over \bar{t} , and a store \mathfrak{s} such that $\bar{S} \models^{\mathfrak{s}} \Theta(\bar{t})$. Let t' be the tree obtained from \bar{t} by replacing each relation atom $r_g(\eta(z_{i_1}^{[p]}), \dots, \eta(z_{i_k}^{[p]}))$ with the original atom $r(z_1^{[p]}, \dots, z_{\#r}^{[p]})$ and adding the original equalities between persistent variables on the labels corresponding to the occurrences of the transitions of $\mathcal{A}_{\Delta, \mathbb{A}}^{III}$ in the construction of $\mathcal{A}_{\Delta, \mathbb{A}}^{IV}$. Note that $\text{dom}(t') = \text{dom}(\bar{t})$. We consider the store \mathfrak{s}' that agrees with \mathfrak{s} over $\text{fv}(\Theta(\bar{t}))$ and maps each variable from the set $\{y_i^{[p]} \mid y_i^{[e]} \in \mathcal{Y}, p \in \text{dom}(t')\}$ to a distinct value, not in $\mathfrak{s}(\text{fv}(\Theta(\bar{t})))$. This is possible because no equality between these variables occurs as logical consequence of $\Theta(t')$ (Def. 15). We consider the structure $S' = (U, \sigma')$, where σ' is the interpretation that assigns each relation symbol r the set of tuples $(\mathfrak{s}'(\xi_1), \dots, \mathfrak{s}'(\xi_{\#r}))$, where ξ_i is either:

- $y_{a(i)}^{[p]}$, if $i \in \mathfrak{P}_A(q)$ and r is the unique position where $y_{a(i)}^{[p]}$ occurs in $\Theta(t')$, or
- $z_i^{[p]}$, otherwise (i.e., $i \notin \mathfrak{P}_A(q)$),

where the relation symbol $r_g(z_{i_1}^{[p]}, \dots, z_{i_k}^{[p]})$ occurs in $\Theta(\bar{t})$ and $\bar{\theta}(p) = (q, a)$ is the state assigned to the position p by the accepting run $\bar{\theta}$ of $\mathcal{A}_{\Delta, \mathbb{A}}^{IV}$ over \bar{t} . By Fact 9, we obtain $S' \models^{\mathfrak{s}'} \Theta(t')$, hence $S' \in \llbracket \Theta(t')^\exists \rrbracket$ and $\text{tw}(S') \leq \text{tw}(\llbracket \Theta(t')^\exists \rrbracket)$. Since any tree decomposition of S' is also a tree decomposition of \bar{S} , we have $\text{tw}(\bar{S}) \leq \text{tw}(\llbracket \Theta(t')^\exists \rrbracket)$.

Consider the tree t obtained by replacing the labels of t' corresponding to the positions where the 1-transitions of $\mathcal{A}_{\Delta, \mathbb{A}}^{III}$ occur (in the construction of $\mathcal{A}_{\Delta, \mathbb{A}}^{IV}$) with the original labels from these transitions. Note that there are no relation atoms and no equalities involving non-persistent variables in these labels, only equalities between persistent variables and disequalities, possibly involving non-persistent variables.

Then, we have $t \in \mathcal{L}(\mathcal{A}^\circ)$, hence $\text{tw}(\llbracket \Theta(t) \rrbracket^\exists) \leq \text{tw}(\llbracket \mathcal{A}^\circ \rrbracket)$. By Lemma 16 (2), we obtain $\text{tw}(\llbracket \Theta(t')^\exists \rrbracket) \leq \text{tw}(\llbracket \Theta(t) \rrbracket^\exists) + K$, where K is a constant depending exclusively on $\mathcal{A}_{\Delta, \mathbb{A}}^{III}$ (i.e., the number of disequalities added to $\Theta(t')$ to obtain $\Theta(t)$ is bounded by the product between the number of 1-transitions in $\mathcal{A}_{\Delta, \mathbb{A}}^{III}$ and the maximum size among the labels of these transitions). Thus, $\text{tw}(\bar{S}) \leq \text{tw}(\llbracket \mathcal{A}^\circ \rrbracket) + K$. Since the choice of \bar{S} was arbitrary, we conclude that $\llbracket \mathcal{A}_{\Delta, \mathbb{A}}^{IV} \rrbracket$ is treewidth-bounded.

Lemma 23. *If \mathcal{A} strongly refines \mathcal{B} and \mathcal{B} is choice-free, then \mathcal{A} is choice-free.*

Proof (Proof of Lemma 23). Let $\mathcal{A} = (\mathbb{A}, Q_{\mathcal{A}}, \iota_{\mathcal{A}}, \delta_{\mathcal{A}})$ and $\mathcal{B} = (\mathbb{A}, Q_{\mathcal{B}}, \iota_{\mathcal{B}}, \delta_{\mathcal{B}})$, with SCC graphs $\mathfrak{G}_{\mathcal{A}} = (\mathcal{N}_{\mathcal{A}}, \mathcal{E}_{\mathcal{A}})$ and $\mathfrak{G}_{\mathcal{B}} = (\mathcal{N}_{\mathcal{B}}, \mathcal{E}_{\mathcal{B}})$, respectively. Because \mathcal{A} strongly refines \mathcal{B} , there exists a strong refinement $h : Q_{\mathcal{A}} \rightarrow Q_{\mathcal{B}}$.

Fact 10 *For any set $S \subseteq Q_{\mathcal{A}}$, S is a (linear) SCC iff $h(S)$ is a (linear) SCC.*

Proof. “ \Rightarrow ” If S is an SCC, then $q \rightsquigarrow^* q'$, for any $q, q' \in S$. By point (2) of Def. 19, also $h(q) \rightsquigarrow^* h(q')$, thus $h(S)$ is an SCC. Assume, moreover, that S is linear and suppose, for a contradiction, that $h(S)$ is non-linear. Then, there exists a transition $q_0 \xrightarrow{a} (q_1, \dots, q_\ell) \in \delta_{\mathcal{B}}$ such that $q_0 \in h(S)$ and $q_i, q_j \in h(S)$ for $1 \leq i < j \leq \ell$. Because h is a strong refinement, we obtain $h^{-1}(q_0) \cup h^{-1}(q_i) \cup h^{-1}(q_j) \subseteq S$ and, moreover, there exists a transition $q'_0 \xrightarrow{a} (q'_1, \dots, q'_\ell) \in \delta_{\mathcal{A}}$, such that $q'_i \in h^{-1}(q_i)$, for all $i \in [0..\ell]$. Then $q'_0, q'_i, q'_j \in S$, thus S is non-linear, contradiction.

“ \Leftarrow ” If $h(S)$ is an SCC then $S = h^{-1}(h(S))$ is an SCC, because h is a strong refinement. Assume, moreover, that $h(S)$ is linear and suppose, for a contradiction, that S is non-linear. Then there exists a transition $q_0 \xrightarrow{a} (q_1, \dots, q_\ell) \in \delta_{\mathcal{A}}$ such that $q_0, q_i, q_j \in S$, for some $1 \leq i < j \leq \ell$. Because h is a refinement, $h(q_0) \xrightarrow{a} (h(q_1), \dots, h(q_\ell)) \in \delta_{\mathcal{B}}$ and $h(q_0), h(q_i), h(q_j) \in h(S)$, thus $h(S)$ is non-linear, contradiction. \square

We prove the points of Def. 14 for \mathcal{A} :

- (1) Since \mathcal{B} is choice-free, $\mathfrak{G}_{\mathcal{B}}$ is a tree with root node, say $S_0^{\mathcal{B}}$. Since h is an isomorphism between $\mathfrak{G}_{\mathcal{B}}$ and $\mathfrak{G}_{\mathcal{A}}$, by Fact 10, $\mathfrak{G}_{\mathcal{A}}$ is a tree with root node $S_0^{\mathcal{A}} \stackrel{\text{def}}{=} h^{-1}(S_0^{\mathcal{B}})$. Let $S \in \mathcal{N}_{\mathcal{A}} \setminus \{S_0^{\mathcal{A}}\}$ and suppose, for a contradiction, that $\text{card}(\bullet S) \neq 1$. Since $S \neq S_0^{\mathcal{A}}$, we must have $\text{card}(\bullet S) > 1$ and let S' be the parent of S in $\mathfrak{G}_{\mathcal{A}}$. Since $\mathfrak{G}_{\mathcal{A}}$ is a tree, there exist at least two transitions $q_{i,0} \xrightarrow{a_i} (q_{i,1}, \dots, q_{i,\ell_i}) \in (S')^\bullet \cap \bullet S$, for $i = 1, 2$. Since h is a simulation, we obtain $h(q_{i,0}) \xrightarrow{a_i} (h(q_{i,1}), \dots, h(q_{i,\ell_i})) \in (h(S'))^\bullet \cap \bullet (h(S))$, for $i = 1, 2$, contradicting the fact that \mathcal{B} is choice-free. Then let $\bullet S \stackrel{\text{def}}{=} \{q_0 \xrightarrow{a} (q_1, \dots, q_\ell)\}$ and suppose, for a contradiction, that $q_i, q_j \in S$, for some $1 \leq i < j \leq \ell$. In this case, $\bullet h(S) = \{h(q_0) \xrightarrow{a} (h(q_1), \dots, h(q_\ell))\}$ and $h(q_i), h(q_j) \in h(S)$, contradicting the fact that \mathcal{B} is choice-free.
- (2) Let $\Lambda : \mathcal{N}_{\mathcal{B}} \cup \delta_{\mathcal{B}} \rightarrow \{1, \infty\}$ be the labeling from point (2) of Def. 14. We prove the following points, for the labeling $\Lambda \circ h$ of $\mathcal{N}_{\mathcal{A}} \cup \delta_{\mathcal{A}}$, where $h(\tau) \stackrel{\text{def}}{=} h(q_0) \xrightarrow{a} (h(q_1), \dots, h(q_\ell))$ for any transition $\tau : q_0 \xrightarrow{a} (q_1, \dots, q_\ell) \in \delta_{\mathcal{A}}$:
 - (2a) Let $S \in \mathcal{N}_{\mathcal{A}}$ be a linear SCC, such that $\Lambda(h(S)) = 1$. By Fact 10, $h(S)$ is linear, hence $\text{card}(h(S)^\bullet) = 1$. By point (2) of Def. 21, we obtain $\text{card}(\bullet S) = 1$.
 - (2b) Let $\tau \in \delta_{\mathcal{A}}$ be a transition. Since \mathcal{B} is choice-free, we have $\Lambda(h(\tau)) = 1$ iff $h(\tau) \in S^\bullet$ for some linear SCC $S \in \mathcal{N}_{\mathcal{B}}$, such that $\Lambda(S) = 1$. By Fact 10, S is a linear SCC of \mathcal{B} iff $h^{-1}(S)$ is a linear SCC of \mathcal{A} . We obtain that $\Lambda(h(\tau)) = 1$ iff $\tau \in (h^{-1}(S))^\bullet$ for some linear SCC $h^{-1}(S)$ such that $(\Lambda \circ h)(h^{-1}(S)) = 1$.
 - (2c) Let $S \in \mathcal{N}_{\mathcal{A}}$ be a SCC. Then $h(S)$ is an SCC of \mathcal{B} , by Fact 10. Since \mathcal{B} is choice-free, we have $\Lambda(h(S)) = 1$ iff $h(S) = S_0$ or $\bullet(h(S)) = \{\tau\}$, for a transition $\tau \in \delta_{\mathcal{B}}$ such that $\Lambda(\tau) = 1$. By Fact 10, $h(S) = S_0$ iff $S = h^{-1}(S_0)$ and we established already that $h^{-1}(S_0)$ is the root node of the tree-like SCC graph $\mathfrak{G}_{\mathcal{A}}$. By point (2b) of Def. 14, $\Lambda(\tau) = 1$ iff there exists a linear SCC S' of \mathcal{B} such that $\tau \in (S')^\bullet$ and $\Lambda(S') = 1$. Then, there exists exactly one transition $\tau' \in \delta_{\mathcal{A}}$, such that $h(\tau') = \tau$, by point 2 of Def. 21. It follows that $\bullet(h(S)) = \{\tau\}$ iff $\bullet S = \{\tau'\}$, for a transition $\tau' \in \delta_{\mathcal{A}}$ such that $\Lambda(h(\tau')) = \Lambda(\tau) = 1$.

Lemma 24. Each $\overline{\mathcal{B}}_1, \dots, \overline{\mathcal{B}}_m$ is all-satisfiable, choice-free and $\mathcal{L}(\mathcal{A}_{\Delta, \mathcal{A}}^{\text{IV}}) = \bigcup_{i=1}^m \mathcal{L}(\overline{\mathcal{B}}_i)$.

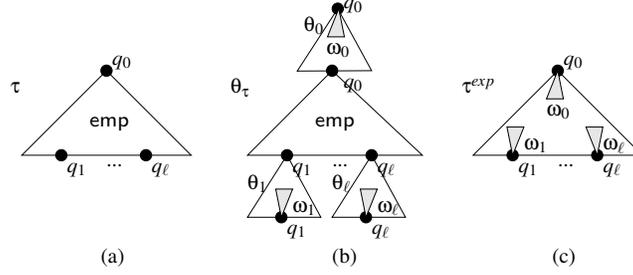


Fig. 8. Expansion of 1-transitions

Proof (Proof of Lemma 24). Each $\overline{\mathcal{B}}_i$ is a strong refinement of $\mathcal{A}_{\Delta, A}^{IV}$, hence $\mathcal{L}(\overline{\mathcal{B}}_i) = \mathcal{L}(\mathcal{A}_{\Delta, A}^{IV})$, by Lemma 15. Since $\mathcal{A}_{\Delta, A}^{IV}$ is all-satisfiable, we obtain that $\overline{\mathcal{B}}_i$ is all-satisfiable. Moreover, $\overline{\mathcal{B}}_i$ is choice-free because $\mathcal{A}_{\Delta, A}^{IV}$ is choice-free, by Lemma 23.

Lemma 25. (1) \mathcal{B} is all-satisfiable. (2) $\llbracket \mathcal{B} \rrbracket$ is treewidth-bounded iff $\llbracket \overline{\mathcal{B}} \rrbracket$ is treewidth-bounded.

Proof (Proof of Lemma 25). 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 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 $q_i \in \bullet(\delta^\infty)$.

The relation between τ from θ_τ is illustrated in Fig. 8(a,b). Note that the resets exist according to Lemma 13. 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 s and a structure S'' such that $S'' \models^s \Theta(t'')$. 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 $\omega_0 * *_{j=1}^\ell \omega_j$ where

$$\omega_0 = * \{ \Omega_{t_0}^{\varepsilon/\varepsilon}(x_{i_1}, \dots, x_{i_k}) \mid i_1, \dots, i_k \in [1.. \#q_0] \}$$

$$\omega_j = * \{ \Omega_{t_j}^{p_j/j}(x_{i_1}, \dots, x_{i_k}) \mid i_1, \dots, i_k \in [1.. \#q_j] \} \text{ for } j \in [1.. \ell]$$

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 Fig. 8(c). Hence, there exists a substructure $S \sqsubseteq S''$ such that $S \models^s \Theta(t)$. Actually, S is simply obtained from S'' by removing all but the relation atoms part of $\omega_0, \dots, \omega_\ell$ defined along the partial runs θ_τ part of θ'' .

2) “ \Rightarrow ” 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 α formula according to the

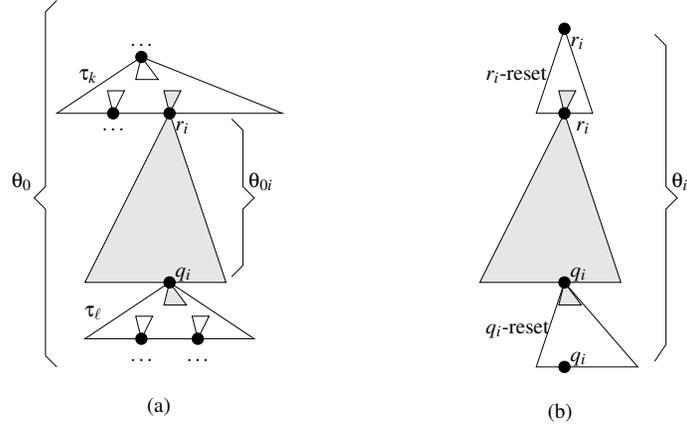


Fig. 9. Principle of run decomposition

transformation. As both $\overline{\mathcal{B}}$ and \mathcal{B} are all-satisfiable, the two characteristic formula $\Theta(t)$ and $\Theta(t')$ are satisfiable. These formula differ only by a finite number of relation atoms, that is, the ones inserted in the α labels of modified 1-transitions. Henceforth, using Lemma 16, point 3, the difference between the treewidth of their models is bounded by some constant. As the choice of t was arbitrary, the result extends to the set of all models of $\overline{\mathcal{B}}$. “ \Leftarrow ” This direction uses a symmetric argument.

Lemma 26. *For each structure $S = (U, \sigma) \in \llbracket \mathcal{B} \rrbracket^c$, there exist pairwise disjoint structures S_1, \dots, S_n and views $\langle \theta_1, t_1, s_1, S'_1 \rangle, \dots, \langle \theta_n, t_n, s_n, S'_n \rangle$ for \mathcal{B} such that S'_1, \dots, S'_n are pairwise disjoint, $S = S_1 \bullet \dots \bullet S_n$ and, for all $i \in [1..n]$: (1) $S_i \sqsubseteq S'_i$, (2) $\text{supp}(\sigma) \cap s_i(\{x_1^{[e]}, \dots, x_{\#\theta_i(e)}^{[e]}\}) = \emptyset$, and $\text{supp}(\sigma) \cap s_i(\{x_1^{[p]}, \dots, x_{\#q_i}^{[p]}\}) = \emptyset$ if θ_i is a context $\theta_{p_i \leftarrow q_i}$.*

Proof (Proof of Lemma 26). Let $\mathcal{B} \stackrel{\text{def}}{=} (\Sigma, Q, \iota, \delta)$. Since $S = (U, \sigma) \in \llbracket \mathcal{B} \rrbracket^c$ there exists a tree $t_0 \in \mathcal{L}(\mathcal{B})$ such that $S \models^{s_0} \Theta(t_0)$ for a store s_0 canonical for $\Theta(t_0)$. Let θ_0 be an accepting run of \mathcal{B} over t_0 . Consider the unique finite decomposition of θ_0 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 (see Fig. 9(a)).

For every $i \in [1..n]$ we define S_i as the substructure of S constructed along the maximal partial run θ_{0i} . That is, 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 Fig. 9(a). Note that, according to the definition of 1-transitions in our transformation the structures S_i as defined above are pairwise disjoint and therefore we have $S = S_1 \bullet \dots \bullet S_n$.

Now we build the views as follows. For every $i \in [1..n]$, we define the partial run θ_i by extending the partial run θ_{0i} by the chosen resets used for the definition of transformation (note that the post-reset is needed only if θ_{0i} reaches a 1-transition). This construction is illustrated in Fig 9(b). Let t_i be the Σ -labelled tree corresponding to the run θ_i . Let s_i be a canonical store for $\Theta(t_i)$ constructed by extending 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 allocates new distinct elements for all other non-equal variables newly introduced along the resets. The store \mathfrak{s}_i defines the structure S'_i , such that $S'_i \models^{\mathfrak{s}_i} \Theta(t_i)$. We prove the points from the statement:

- (1) By the choice of \mathfrak{s}_i , that extends \mathfrak{s}_0 as explained above, we have $S_i \sqsubseteq S'_i$. Actually, no extra relation atoms are added into S_i as part of the newly inserted resets because of the labeling of 1-transitions in \mathcal{B} is constructed precisely from the resets to guarantee this property. That is, our transformation guarantees that the exactly same set of relation atoms occurring in the reset was used to label the 1-transitions.
- (2) By the construction of the partial run θ_i the set of variables occurring at the entry position (i.e., $\{x_1^{[e]}, \dots, x_{\#r_i}^{[e]}\}$) are not related by equalities to any variables at the entry of θ_{0i} . This is ensured by any reset of an automaton without persistent variables, and we note that \mathcal{B} is such an automaton. r_i -reset for automata without persistent variables). Since, moreover, \mathfrak{s}_i is canonical for $\Theta(t_i)$, we obtain that $\text{supp}(\sigma)$ and $\mathfrak{s}_i(\{x_1^{[e]}, \dots, x_{\#r_i}^{[e]}\})$ are disjoint. The same reasoning applies to the set of variables occurring at the exit position of θ_i , that is, when θ_{0i} reaches a 1-transition.

Lemma 27. *Given pairwise disjoint structures S_1, \dots, S_n and 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} , satisfying the conditions (1) and (2) of Lemma 26, there exists a rich canonical model $(S, \mathfrak{d}) \in \llbracket \mathcal{B} \rrbracket^r$, such that the conditions (1), (2) and (3) from Def. 12 hold for S_1, \dots, S_n and (S, \mathfrak{d}) .*

Proof (Proof of Lemma 27). By Lemma 14, there exists an accepting run θ of \mathcal{B} over a tree t , such that the following hold:

1. each partial run θ_i is embedded in θ at some position $r_i \in \text{dom}(\theta)$, for all $i \in [1..n]$,
2. $r_i \cdot \text{dom}(\theta_i) \cap r_j \cdot \text{dom}(\theta_j) = \emptyset$, for all $1 \leq i < j \leq n$,
3. 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 $S'_i \models^{\mathfrak{s}_i} \Theta(t_i)$, hence $\mathfrak{s}_i(\text{fv}(\Theta(t_i))) \subseteq \text{supp}(\sigma'_i)$, where we assume w.l.o.g. that $S'_i \stackrel{\text{def}}{=} (U, \sigma'_i)$. By point (2) 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 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 $x^{[p]} \approx_{\Theta(t)} z^{[r]}$.

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 $S = (U, \sigma)$, 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 S , we have $S \models^{\mathfrak{s}} \Theta(t)$. 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 obtain a rich canonical model (S, \mathfrak{d}) of $\Theta(t)$ and, since $t \in \mathcal{L}(\mathcal{B})$, we have $(S, \mathfrak{d}) \in \llbracket \mathcal{B} \rrbracket^r$. We prove below the three conditions from Def. 12:

- (1) By the construction of S , we have $S'_i \subseteq S$, for all $i \in [1..n]$. Since S_1, \dots, S_n are pairwise disjoint, their composition is defined, hence $S_1 \bullet \dots \bullet S_n \subseteq S$. W.l.o.g., let

$S_i = (U, \sigma_i)$, for all $i \in [1..n]$. To prove $S_1 \bullet \dots \bullet S_n \sqsubseteq S$, 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 $s(z_i) = u_i$, 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 p_i , or p_i and p_j , 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 $\ell_i \in [1..\#\theta_i(\varepsilon)]$, such that $s(x_{\ell_i}^{[r_i]}) = u_i$. Hence, $\text{supp}(\sigma_i) \cap s(\{x_1^{[r_i]}, \dots, x_{\#\theta_i(\varepsilon)}^{[r_i]}\}) = \text{supp}(\sigma_i) \cap s_i(\{x_1^{[\varepsilon]}, \dots, x_{\#\theta_i(\varepsilon)}^{[\varepsilon]}\}) \neq \emptyset$, in contradiction with condition (2) of Lemma 26. 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 $k \in [1..n]$, i.e., $\sigma_1(r) \uplus \dots \uplus \sigma_n(r) = \{\langle u_1, \dots, u_{\#r} \rangle \mid u_1, \dots, u_{\#r} \in \bigcup_{i=1}^n \text{supp}(\sigma_i)\}$. Since the choice of r was arbitrary, we obtain that $S_1 \bullet \dots \bullet S_n \sqsubseteq S$.

- (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 $s(x^{[p]}) = u$ and $s(y^{[p]}) = v$. 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 $s(\xi_i^{[p_i]}) = u$. Since 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 Lemma 26 (2), 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 condition (2) 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 (3) above, this path contains a reset path disjoint from $\bigcup_{k=1}^n r_k \cdot \text{dom}(\theta_k)$, hence $z_g^{[p_1]} \not\approx_{\Theta(t)} z_h^{[p_2]}$, contradiction.

E Material from Section 6

Lemma 28. *Given a set \mathcal{S} of structures, we have (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 (Proof of Lemma 28). (1) “ $\text{split}(\text{EF}^*(\mathcal{S})) \supseteq \text{EF}^*(\text{split}(\mathcal{S}))$ ” By induction on the derivation of $S \in \text{EF}^*(\text{split}(\mathcal{S}))$ from $\text{split}(\mathcal{S})$. Base case: Let $S \in \text{split}(\mathcal{S})$. We have:

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

Induction step: Assume $S = (S_1 \bullet S_2)_{/\approx}$ for some $S_1, S_2 \in \text{EF}^*(\text{split}(\mathcal{S}))$ and \approx satisfying the conditions of Def. 4 for external fusion of S_1 and S_2 . Moreover, assume the induction hypothesis $S_1, S_2 \in \text{split}(\text{EF}^*(\mathcal{S}))$. Then

$$\begin{aligned} S_1, S_2 \in \text{split}(\text{EF}^*(\mathcal{S})), S = (S_1 \bullet S_2)_{/\approx} &\Rightarrow \\ (\exists S'_1, S'_2 \in \text{EF}^*(\mathcal{S}). S_1 \sqsubseteq^{mc} S'_1, S_2 \sqsubseteq^{mc} S'_2), S = (S_1 \bullet S_2)_{/\approx} &\Rightarrow \\ \exists S'_1, S'_2 \in \text{EF}^*(\mathcal{S}). (\exists \approx' . S \sqsubseteq^{mc} (S'_1 \bullet S'_2)_{/\approx'}) &\Rightarrow \\ \exists S' \in \text{EF}^*(\mathcal{S}). S \sqsubseteq^{mc} S' \Rightarrow S \in \text{split}(\text{EF}^*(\mathcal{S})) \end{aligned}$$

In the above, the equivalence \approx' is taken as the extension by equality of \approx and henceforth it satisfies the conditions of Def. 4 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 Def. 4 for external fusion of S'_1, S'_2 . Moreover, 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 $i = 1, 2$:

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

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 Def. 4 from the respective substructures.

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

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

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

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 .

” $\text{split}(\text{IEF}^*(S)) \subseteq \text{IEF}^*(\text{split}(S))$ ” By induction on the derivation of $S' \in \text{IEF}^*(S)$ from 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 $S'_1 \in \text{IEF}^*(S)$ and equivalence \approx' conforming for internal fusion of S'_1 . Moreover, assume the induction

hypothesis $\text{split}(S'_1) \subseteq \text{IEF}^*(\text{split}(S))$. Then

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

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 $\text{IEF}^*(\text{split}(S))$ i.e., first by using external fusion conforming to Def. 4 to connect all the included substructures and second, by using internal fusion to further restrict the result if needed. \square

Lemma 29. *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 $\{(u_1, u_2)\}^\approx \subseteq (U_1 \cup U_2) \times (U_1 \cup U_2)$ is compatible with $S_1 \bullet S_2$.*

Proof (Proof of Lemma 29). We denote by \approx the relation $\{(u_1, u_2)\}^\approx$ in the following. Let $r \in \mathbb{R}$ be a relation and let $\langle u_{1,1}, u_{1,2}, \dots, u_{1,\#r} \rangle \in \sigma_1(r)$, $\langle u_{2,1}, u_{2,2}, \dots, u_{2,\#r} \rangle \in \sigma_2(r)$ 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 indexes $i \in [1..\#r]$ 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

Lemma 30. *Let S be a treewidth-bounded set of connected structures conforming to an RGB color scheme. Then, for any structure $S \in \text{EF}^*(S)$, the following hold:*

1. S is connected and of type either R, G or B,
2. $\text{tw}(S) \leq \begin{cases} \text{tw}(S) & , \text{ if } S \text{ of type R} \\ \max(\text{tw}(S) + 2 \cdot \text{card}(\mathbb{C}^{\text{green}}), 3 \cdot \text{card}(\mathbb{C}^{\text{green}})) & , \text{ if } S \text{ of type G} \\ \max(\text{tw}(S) + 2 \cdot \text{card}(\mathbb{C}^{\text{green}}), 3 \cdot \text{card}(\mathbb{C}^{\text{green}}), \text{tw}(S) + 1) & , \text{ if } S \text{ is of type B} \end{cases}$

Proof (Proof of Lemma 30). (1) By induction on the derivation of $S \in \text{EF}^*(S)$ from S . Table 1 summarizes the possible types of $\text{EF}(S_1, S_2)$ on connected structures S_1 and S_2 of types R, G or B, respectively.

(2) 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 $S_{/\approx}$,

Table 1. Possible resulting type of $\text{EF}(S_1, S_2)$ for S_1, S_2 of R, G, B types.

$\text{EF}(S_1, S_2)$	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	-
S_1 of B type	B	-	-

obtained by the relabeling of elements u in the bags of T by their representatives $[u]_{\approx}$. Let $S \in \text{EF}^*(\mathcal{S})$ be a structure. We consider the following three cases:

S is of type R: We consider the stronger property, that is, S admits a tree decomposition T where (i) the unique element with color in \mathbb{C}^{red} is present in the bag of the root r of T and (ii) the treewidth is bounded by $\text{tw}(\mathcal{S})$. Obviously, this stronger property implies the expected result, namely $\text{tw}(S) \leq \text{tw}(\mathcal{S})$. The proof goes by induction on the derivation of $S \in \text{EF}^*(\mathcal{S})$ from \mathcal{S} .

Base case: Assume that $S \in \mathcal{S}$. Then, obviously a tree decomposition T exists, such that $\text{wd}(T) \leq \text{tw}(\mathcal{S})$. Moreover, T can be progressively transformed by reversing tree edges (hence preserving its width) until one of the nodes containing the element with color in \mathbb{C}^{red} becomes the root.

Induction step: according to Table 1 the only way to construct a structure S of R type is by fusion of two structures S_1, S_2 both of R type. The external fusion between S_1 and S_2 is possible only by using the elements u_1, u_2 with colors in \mathbb{C}^{red} , from respectively S_1, S_2 . Wlog, consider that u_2 is glued to an element $u'_1 \in S_1$ and let \approx be the equivalence relation $\{(u_2, u'_1)\}^=$. Then, from the tree decompositions T_1, T_2 of respectively S_1 and S_2 satisfying the induction hypothesis, we build the tree decomposition T_{12} of $S_1 \bullet S_2$ by linking the root r_2 of T_2 as a child of some node n_1 of T_1 which contains u'_1 . This ensures that $T_{12/\approx}$ is a valid tree decomposition of S . By construction, the unique element with color in \mathbb{C}^{red} after fusion is eventually $[u_1]_{\approx}$ and remains present in the bag of the tree root r_1 . Also, $\text{tw}(T_{12/\approx}) \leq \max(\text{tw}(T_1), \text{tw}(T_2))$, hence we obtain $\text{tw}(T_{12/\approx}) \leq \text{tw}(\mathcal{S})$, by applying the induction hypothesis.

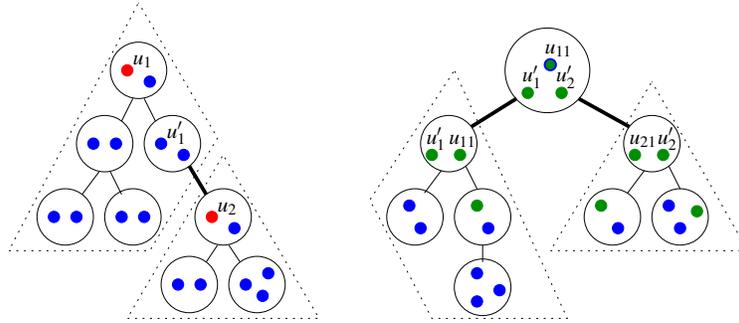


Fig. 10. Tree decompositions for structures of R type (left) and G type (right)

S is of type G: We consider the stronger property, that is, S admits a tree decomposition T where (i) all elements with colors in \mathbb{C}^{green} are present in the bag of the root r of

T and (ii) the treewidth is bounded by $\max(\text{tw}(\mathcal{S}) + 2 \cdot \text{card}(\mathbb{C}^{\text{green}}), 3 \cdot \text{card}(\mathbb{C}^{\text{green}}))$. Obviously, this stronger property implies the expected result, namely the upper bound on $\text{tw}(\mathcal{S})$. We prove it by induction on the derivation of $S \in \text{EF}^*(\mathcal{S})$ from \mathcal{S} .

Base case: Assume that $S \in \mathcal{S}$. Then, obviously a tree decomposition T exists, such that $\text{wd}(T) \leq \text{tw}(\mathcal{S})$. Furthermore, T can be transformed by replicating all elements with colors in $\mathbb{C}^{\text{green}}$ in all nodes (and hence to the root as well). As there are at most $2 \cdot \text{card}(\mathbb{C}^{\text{green}})$ elements in S (point (2) of Def. 27) it follows that the obtained tree decomposition has width bounded by $\text{tw}(\mathcal{S}) + 2 \cdot \text{card}(\mathbb{C}^{\text{green}})$.

Induction step: according to Table 1 the only way to construct a structure S of G is by fusion of two structures S_1, S_2 of type either R or G. In any case, we build a tree decomposition $T_{12/\approx}$ for S from the tree decompositions T_1, T_2 of respectively S_1, S_2 satisfying their respective inductive hypothesis, and the (fusion) equivalence relation \approx . If at least one of the structures is of type R, the external fusion eventually uses the unique element with color in \mathbb{C}^{red} . Then, we can use the same construction as in previous case for the fusion of structures of type R. The tree decomposition $T_{12/\approx}$ will satisfy by construction the condition related to the placement of elements with colors in $\mathbb{C}^{\text{green}}$. Moreover, as previously, it holds $\text{tw}(T_{12/\approx}) \leq \max(\text{tw}(T_1), \text{tw}(T_2))$ which implies the expected bound on the treewidth. If both structures S_1, S_2 are of type G, the construction of T_{12} is a bit different. In this case, the external fusion is necessarily due to the gluing of a non-empty set of pairs $\{(u_{1i}, u_{2i}) \mid i \in I\}$ from respectively S_1, S_2 conforming to conditions stated in Def. 4. Let \approx be the equivalence relation $(\{(u_{1i}, u_{2i}) \mid i \in I\})^-$. We build the tree decomposition T_{12} of $S_1 \bullet S_2$ by adding a fresh root node r_{12} with two children, respectively the roots r_1, r_2 of T_1, T_2 . Moreover, in the bag of r_{12} we copy (i) all elements u'_1, u'_2 from the bags of r_1, r_2 with colors in $\mathbb{C}^{\text{green}}$ that were not glued, (ii) all elements that are glued from one of the sides i.e., u_{1i} for all $i \in I$. Then, it is an easy check that $T_{12/\approx}$ is a valid tree decomposition of $(S_1 \bullet S_2)_{/\approx}$. By construction, all elements with colors in $\mathbb{C}^{\text{green}}$ are located at the root r_{12} . Also $\text{tw}(T_{12/\approx}) \leq \max(\text{tw}(T_1), \text{tw}(T_2), \text{card}(\lambda_{12}(r_{12})))$ where $\lambda_{12}(r_{12})$ denotes the bag of elements at the root r_{12} of T_{12} . We will show that $\text{card}(\lambda_{12}(r_{12})) \leq 3 \cdot \text{card}(\mathbb{C}^{\text{green}})$, that is, a sufficient condition to establish the bound of $\text{tw}(T_{12/\approx})$ knowing moreover the bounds of the widths of T_1, T_2 given by the induction hypothesis. Consider that N_1, N_2 elements from respectively r_1, r_2 were copied to r_{12} without being glued. Consider that N_{12}, N'_{12} pairs of elements were glued and become elements with colors from $\mathbb{C}^{\text{green}}, \mathbb{C}^{\text{blue}}$, respectively. Then, as each one of S_1, S_2, S has at most $2 \cdot \text{card}(\mathbb{C}^{\text{green}})$ elements in $\mathbb{C}^{\text{green}}$ (point (2) of Def. 27), we obtain:

$$N_1 + N_{12} + N'_{12} \leq 2 \cdot \text{card}(\mathbb{C}^{\text{green}})$$

$$N_2 + N_{12} + N'_{12} \leq 2 \cdot \text{card}(\mathbb{C}^{\text{green}})$$

$$N_1 + N_2 + N_{12} \leq 2 \cdot \text{card}(\mathbb{C}^{\text{green}})$$

By summing up, we obtain the approximation of $\text{card}(\lambda_{12}(r_{12}))$ as follows:

$$\begin{aligned} 2 \cdot \text{card}(\lambda_{12}(r_{12})) &= 2 \cdot (N_1 + N_2 + N_{12} + N'_{12}) \\ &\leq 2 \cdot N_1 + 2 \cdot N_2 + 3 \cdot N_{12} + 2 \cdot N'_{12} \\ &\leq 6 \cdot \text{card}(\mathbb{C}^{\text{green}}) \end{aligned}$$

S is of B type: The expected property is established by induction on the derivation of $S \in \text{EF}^*(\mathcal{S})$ from \mathcal{S} . **Base case:** Let $S \in \mathcal{S}$ be a structure. Then, a tree decomposition T exists, such that $\text{wd}(T) \leq \text{tw}(\mathcal{S})$.

Induction step: according to Table 1 there are several ways to construct a structure S of \overline{G} is by fusion of two structures S_1, S_2 of either R, G or B type. In any case, we shall build a tree decomposition $T_{12/\approx}$ for S from the tree decompositions T_1, T_2 of respectively S_1, S_2 satisfying their respective inductive hypothesis and the (fusion) equivalence relation \approx . For the fusion of G-G type structures, the construction of $T_{12/\approx}$ as given in the previous case guarantees the bound on $\text{tw}(T_{12/\approx})$. For the fusion of G-R or B-R type structures, the same construction of $T_{12/\approx}$ as given in the first case guarantees the bound on $\text{tw}(T_{12/\approx})$. For the fusion of R-R structures producing B type structures one can reuse the construction given in the first case, with a slightly adaptation. Consider that elements u_1, u_2 (from S_1, S_2 , respectively) with colors in \mathbb{C}^{red} are glued to respectively u'_2, u'_1 (from S_2, S_1 , respectively). First, we modify the tree decomposition T_1 by replicating u_1 in all nodes and the tree decomposition T_2 by replicating u'_2 in all nodes. Given the induction hypothesis, after modification T_1, T_2 have the width at most $\text{tw}(S) + 1$. Second, we proceed as previously, i.e., construct the tree decomposition T_{12} of $S_1 \bullet S_2$ by making the root r_2 of T_2 the child to some node n_1 of T_1 containing the element u'_1 . It is an easy check that $T_{12/\approx}$ is a valid tree decomposition of S and its width is at most $\text{tw}(S) + 1$. \square

Lemma 31. *Let S be a treewidth-bounded set of connected structures, conforming to an RGB color scheme. Then, $\text{tw}(\text{EF}^*(S))$ and $\text{tw}(\text{IF}(\text{EF}^*(S)))$ are both bounded by $\max(\text{tw}(S) + 2 \cdot \text{card}(\mathbb{C}^{green}), 3 \cdot \text{card}(\mathbb{C}^{green}), \text{tw}(S) + 1)$.*

Proof (Proof of Lemma 31). $\text{EF}^*(S)$ is treewidth bounded as a direct consequence of Lemma 30, that establishes the bounds for every type of structure from $\text{EF}^*(S)$. Moreover, $\text{IF}(\text{EF}^*(S))$ is treewidth-bounded because, using the tree decompositions T constructed for structures S in $\text{EF}^*(S)$ 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 eventually glues the unique element u_1 with color in \mathbb{C}^{red} to some other element in the structure. 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 $\text{wd}(T') \leq \text{wd}(T) + 1 \leq \text{tw}(S) + 1$.
- if S is of type G then by internal fusion ones eventually glue elements with color in \mathbb{C}^{green} . As all such elements are already present in the root node, $T_{/\approx}$ is a valid tree decomposition for S' . Obviously, the treewidth bound remains unchanged.
- if S is of type B then no non-trivial internal fusion exists, and obviously, the treewidth bound remains unchanged. \square

Lemma 32. *The following are equivalent, for any treewidth-bounded set S of connected structures:*

1. $\text{EF}^*(S)$ is treewidth bounded,
2. $[C_1, C_1, C_1], [C_2, C_2, C_2] \in (\text{EF}^*(\text{split}(S)))^{\sharp 3}$ implies $C_1 \cap C_2 \neq \emptyset$ for all C_1, C_2 ,
3. $\text{split}(S)$ conforms to some RGB color scheme.

Proof (Proof of Lemma 32). “(1) \Rightarrow (2)” If $\text{EF}^*(S)$ is treewidth-bounded then $\text{split}(\text{EF}^*(S))$ is treewidth-bounded. Using Lemma 28 the later set is equal to $\text{EF}^*(\text{split}(S))$ and henceforth treewidth bounded as well. By contradiction, assume that (2) does not hold. Then, there exist colors $C_1, C_2 \in \mathbb{C}$, connected structures $S_1, S_2 \in \text{EF}^*(\text{split}(S))$ such that $[C_1, C_1, C_1] \in S_1^{\sharp 3}$, $[C_2, C_2, C_2] \in S_2^{\sharp 3}$ and moreover $C_1 \cap C_2 = \emptyset$. 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 $S_{12} \in \text{EF}^*(\text{split}(S))$ by fusing one pair (u_1, u_2) with colors C_1, C_2 . Let v_1, w_1 resp. v_2, w_2 be the remaining distinct elements of S_{12} with color C_1, 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 $X^{n,n} \in \text{EF}^*(\text{split}(S))$:

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

where structures $S_{12}^{i,j}$ are added to the fusion in increasing order of $i + j$. The construction is illustrated in Fig. 3. We can show that $\text{MIN}(X^{n,n})$ contains an $n \times n$ square grid. Finally, as n can be taken arbitrarily large, we conclude that $\text{MIN}(\text{EF}^*(\text{split}(S)))$ contains arbitrarily large square grids, it is not treewidth-bounded, contradicting (1).

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

$$\mathbb{C}^{blue} = \{C \in \mathbb{C} \mid [C, C, C] \in (\text{EF}^*(\text{split}(S)))^{\sharp 3}\}$$

Since (2) holds, this is a valid definition for \mathbb{C}^{blue} , and henceforth one defined the partitioning of the remaining colors between \mathbb{C}^{green} and \mathbb{C}^{red} . We show that $\text{split}(S)$ is conforming to this RGB partitioning:

- checking point (1) of Def. 27: Let $S \in \text{split}(S)$ and prove that for any two colors $C_1, C_2, [C_1, C_2] \subseteq S^\sharp$ and $C_1 \in \mathbb{C}^{red}$ only if $C_2 \in \mathbb{C}^{blue}$. Since $C_1 \in \mathbb{C}^{red}$, there must exists a color $C'_1 \in \mathbb{C}^{blue}$, such that $C_1 \cap C'_1 = \emptyset$, by Def. 26. Now, given the definition of \mathbb{C}^{blue} in our RGB-color scheme, this further implies $[C'_1, C'_1, C'_1] \in (\text{EF}^*(\text{split}(S)))^{\sharp 3}$. Henceforth, there exists a structure $S' \in \text{EF}^*(\text{split}(S))$ such that $[C'_1, C'_1, C'_1] \subseteq S'^\sharp$. 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 contains three elements of color C_2 , one from each disjoint copy of S . Therefore, $[C_2, C_2, C_2] \in S''^\sharp$ and because $S'' \in \text{EF}^*(\text{split}(S))$ this implies $[C_2, C_2, C_2] \in (\text{EF}^*(\text{split}(S)))^{\sharp 3}$ and therefore $C_2 \in \mathbb{C}^{blue}$.
- checking point (2) of Def. 27: By contradiction, let $S \in \text{EF}^*(\text{split}(S))$ be such that $S^\sharp \cap \mathbb{C}^{green} \not\subseteq [C, C \mid C \in \mathbb{C}^{green}]$. Then there exists a color $C' \in (S^\sharp \cap \mathbb{C}^{green}) \setminus [C, C \mid C \in \mathbb{C}^{green}]$. We have $C' \in \mathbb{C}^{green}$ and also $[C', C', C'] \subseteq S^\sharp$. The later implies $[C', C', C'] \in S^{\sharp 3} \subseteq (\text{EF}^*(\text{split}(S)))^{\sharp 3}$. 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 31, $\text{EF}^*(\text{split}(S))$ is treewidth-bounded. Then, by Lemma 28, $\text{split}(\text{EF}^*(S))$ is treewidth-bounded, thus $\text{EF}^*(S)$ is treewidth bounded. \square

Lemma 33. $(\text{EF}^*(S))^{\sharp k} = (\text{EF}_1^*(S))^{\sharp k}$ for any set S of structures and integer $k \geq 1$.

Proof (Proof of Lemma 33). “ $(\text{EF}_1^*(S))^{\sharp k} \subseteq (\text{EF}^*(S))^{\sharp k}$ ” This direction follows directly from $\text{EF}_1^*(S) \subseteq \text{EF}^*(S)$. “ $(\text{EF}^*(S))^{\sharp k} \subseteq (\text{EF}_1^*(S))^{\sharp k}$ ” We prove the stronger property:

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

By induction on the derivation of $S \in \text{EF}^*(S)$ from \mathcal{S} . Base case: Assume $S \in \mathcal{S}$. Then $S' = S$ satisfies the property.

Induction step: Assume $S = (S_1 \bullet S_2)_{/\approx}$ for some $S_1, S_2 \in \text{EF}^*(S)$ 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 S_1, S_2 . Let $C_{1i} = C_{S_1}(u_{1i})$, $C_{2i} = C_{S_2}(u_{2i})$, for all $i \in [1..n]$. According to the definition of external fusion, $S = (S_1 \bullet S_2)_{/\approx}$ implies $C_{1i} \cap C_{2i} = \emptyset$ and moreover:

$$S^\sharp = [(C_{1i} \cup C_{2i}) \mid i \in [1..n]] \cup (S_1^\sharp \setminus [(C_{1i}) \mid i \in [1..n]]) \cup (S_2^\sharp \setminus [(C_{2i}) \mid i \in [1..n]])$$

By induction hypothesis, for S_1, S_2 there exists $S'_1, S'_2 \in \text{EF}_1^*(S)$ such that $S_1^\sharp \subseteq S'^{\sharp}_1$, $S_2^\sharp \subseteq S'^{\sharp}_2$. We use S'_1 and n disjoint copies $S'_{2,1}, \dots, S'_{2,n}$ of S'_2 to construct 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 C_{1i} from S'_1 with some element u'_{2i} with color C_{2i} from $S'_{2,i}$. Such elements always exist, because $S_1^\sharp \subseteq S'^{\sharp}_1$, $S_2^\sharp \subseteq S'^{\sharp}_2$. 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:

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

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

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

□

Lemma 34. $(\text{EF}_1^*(S))^{\sharp k} = \text{ef}_1^{\sharp k}(S^{\sharp k})$ for any set S of structures, for any integer $k \geq 1$.

Proof (Proof of Lemma 34). Abusing notation, we write $M^{\sharp k} \stackrel{\text{def}}{=} \{M' \mid M' \subseteq M, \text{card}(M') \leq k\}$.

Then, we have $(\text{EF}_1^*(S))^{\sharp k} = ((\text{EF}_1^*(S))^\sharp)^{\sharp k}$, by Def. 25. Using Def. 28 of single pair external fusion and Def. 29 of single pair fusion of multisets, we can prove that for all structures S_1, S_2 it holds $(\text{EF}_1(S_1, S_2))^\sharp = \text{ef}_1^{\sharp(S_1^\sharp, S_2^\sharp)}$. This immediately extends to their respective closure, henceforth, $(\text{EF}_1^*(S))^\sharp = \text{ef}_1^{\sharp}(S^\sharp)$. Henceforth, we are left with proving that $(\text{ef}_1^{\sharp}(S^\sharp))^{\sharp k} = \text{ef}_1^{\sharp k}(S^{\sharp k})$.

'' $(\text{ef}_1^{\sharp}(S^\sharp))^{\sharp k} \subseteq \text{ef}_1^{\sharp k}(S^{\sharp k})$ '' We prove that, for all $M \in \text{ef}_1^{\sharp}(S^\sharp)$, we have $M^{\sharp k} \subseteq \text{ef}_1^{\sharp k}(S^{\sharp k})$.

The proof goes by induction on the derivation of $M \in \text{ef}_1^{\sharp}(S^\sharp)$ from S^\sharp . Base case: Assume $M \in S^\sharp$. Then $M^{\sharp k} \subseteq (S^\sharp)^{\sharp k} = S^{\sharp k} \subseteq \text{ef}_1^{\sharp k}(S^{\sharp k})$.

Induction step: Assume $M \in \text{ef}_1^{\sharp(M_1, M_2)}$ for some multisets of colors M_1, M_2 such that $M_1^{\sharp k}, M_2^{\sharp k} \subseteq \text{ef}_1^{\sharp k}(S^{\sharp k})$. 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 Def. 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 $M' \in \text{ef}_1^{\sharp k}(S^{\sharp k})$.
- $M' \not\subseteq M_i$, 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 $M' \in \text{ef}_1^{\sharp k}((M'_1 \cup [C_1]), (M'_2 \cup [C_2]))$. This implies $M' \in \text{ef}_1^{\sharp k}(S^{\sharp k})$ as both subterms belong to $\text{ef}_1^{\sharp k}(S^{\sharp k})$.
- $M' \not\subseteq M_i$, 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 $M' \in \text{ef}_1^{\sharp k}(S^{\sharp k})$, as well.

“ $\text{ef}_1^{\sharp k^*}(\mathcal{S}^{\sharp k}) \subseteq (\text{ef}_1^{\sharp*}(\mathcal{S}^{\sharp}))^{\sharp k}$ ” We prove that, for all k -multiset $M' \in \text{ef}_1^{\sharp k^*}(\mathcal{S}^{\sharp k})$, there exists $M \in \text{ef}_1^{\sharp*}(\mathcal{S}^{\sharp})$, such that $M' \subseteq M$, by induction on the derivation of M' from $\mathcal{S}^{\sharp k}$.
Base case: Assume $M' \in \mathcal{S}^{\sharp k} = (\mathcal{S}^{\sharp})^{\sharp k}$. Then, there exists $M \in \mathcal{S}^{\sharp}$ such that $M' \subseteq M$. Obviously, $M \in \text{ef}_1^{\sharp*}(\mathcal{S}^{\sharp})$.

Induction step: Assume $M' \in \text{ef}_1^{\sharp k}(M'_1, M'_2)$ for some k -multisets of colors $M'_1, M'_2 \in \text{ef}_1^{\sharp k^*}(\mathcal{S}^{\sharp k})$. By the inductive hypothesis, there exists multisets $M_1, M_2 \in \text{ef}_1^{\sharp*}(\mathcal{S}^{\sharp})$ 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 $M \in \text{ef}_1^{\sharp*}(\mathcal{S}^{\sharp})$, which is the superset of M' . \square

Lemma 35. ($\llbracket A \rrbracket_{\Delta}^c \stackrel{\sharp k}{=} \pi_3(\llbracket A \rrbracket_{\Delta}^{\sharp k})$, for any $k \geq 1$, any equality-free SID Δ and nullary predicate A , such that each complete Δ -unfolding of A yields a satisfiable formula.

Proof (Proof of Lemma 35). “ \subseteq ” We prove the following, more general, property:

Let $B_0(x_1, \dots, x_{\#B_0}) \Rightarrow_{\Delta}^* \exists y_1 \dots \exists y_n . \phi$ be a complete Δ -unfolding such that ϕ is a qpf formula, \mathfrak{s} be a store injective over $\{x_1, \dots, x_{\#B_0}\} \cup \{y_1, \dots, y_n\}$, $S = (U, \sigma)$ be a structure such that $S \models^{\mathfrak{s}} \phi$ and $D \subseteq \text{supp}(\sigma) \setminus \{\mathfrak{s}(x_1), \dots, \mathfrak{s}(x_{\#B_0})\}$ be a set such that $\text{card}(D) \leq k$. Then there exists $\langle \{x_1, \dots, x_{\#B_0}\}, c, M \rangle \in \llbracket B_0 \rrbracket_{\Delta}^{\sharp k}$ such that $C_S(\mathfrak{s}(x_i)) = c(x_i)$, for all $i \in [1.. \#B_0]$ and $M = [C_S(u) \mid u \in D]$.

The proof is by induction on the length of the complete Δ -unfolding. Assume w.l.o.g. that the first rule applied in the unfolding is of the form (5), with a qpf formula ψ_0 . Then, there exist structures $S_0 = (U_0, \sigma_0), \dots, S_{\ell} = (U_{\ell}, \sigma_{\ell})$, such that:

- $S = S_0 \bullet \dots \bullet S_{\ell}$,
- $S_0 \models^{\mathfrak{s}} \psi_0$, and
- there exists a complete unfolding $B_i(z_{i,1}, \dots, z_{i,\#B_i}) \Rightarrow_{\Delta}^* \exists y_{j_{i,1}} \dots \exists y_{j_{i,k_i}} . \phi_i$, where $j_{i,1}, \dots, j_{i,k_i} \in [1..n]$ and ϕ_i is a subformula of ϕ , such that $S_i \models^{\mathfrak{s}} \phi_i$, and the indices $j_{i,m}$ are pairwise distinct, for all $i \in [1..\ell]$ and $m \in [1..k_i]$.

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 unfolding:

$$B_i(x_1, \dots, x_{\#B_i}) \Rightarrow_{\Delta}^* (\exists y_{j_1} \dots \exists y_{j_k} . \phi_i)[z_{i,1}/x_1, \dots, z_{i,\#B_i}/x_{\#B_i}] = \exists y_{j_1} \dots \exists y_{j_k} . \psi_i$$

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:

$$D_0 \stackrel{\text{def}}{=} D \cap (\text{supp}(\sigma_0) \cup \bigcup_{i=1}^{\ell} \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\})$$

$$D_i \stackrel{\text{def}}{=} (\text{supp}(\sigma_i) \cap D) \setminus \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\}, \text{ for each } i \in [1..\ell]$$

and prove the following fact:

Fact 11 $D = D_0 \uplus D_1 \uplus \dots \uplus D_{\ell}$

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

$$D_0 \subseteq \text{supp}(\sigma_0) \cup \bigcup_{i=1}^{\ell} \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\}$$

$$D_i \subseteq \text{supp}(\sigma_i) \setminus \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\}, \text{ for all } i \in [1..\ell]$$

and, moreover:

$$\begin{aligned} & (\text{supp}(\sigma) \cup \bigcup_{i=1}^{\ell} \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\}) \cap \text{supp}(\sigma_i) \subseteq \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\}, \text{ for all } i \in [1..\ell] \\ & \text{supp}(\sigma_i) \cap \text{supp}(\sigma_j) \subseteq \{\mathfrak{s}(z_{i,1}), \dots, \mathfrak{s}(z_{i,\#B_i})\} \cap \{\mathfrak{s}(z_{j,1}), \dots, \mathfrak{s}(z_{j,\#B_j})\}, \text{ for all } 1 \leq i < j \leq \ell \end{aligned}$$

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 $\langle \{x_1, \dots, x_{\#B_i}\}, c_i, M_i \rangle \in \langle \langle B_i \rangle \rangle_{\Delta}^{\#k}$, for $i \in [1..\ell]$, such that:

- $C_S(\mathfrak{s}(x_j)) = c(x_j)$, for all $j \in [1..\#B_i]$,
- $M_i = [C_{S_i}(u) \mid u \in D_i]$.

Let $\langle \text{fv}(\psi_0), c_0, \emptyset \rangle \stackrel{\text{def}}{=} \gamma(\psi_0)$ be a color triple. Since $S_0 \models^s \psi_0$, we have $C_{S_0}(\mathfrak{s}(x_j)) = c(x_j)$, for all $j \in [1..\#B_0]$. By definition, there exists a constraint of the form (4) for the above rule (5). 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 $C_{S_i}(x) \cap C_{S_j}(x) \neq \emptyset$, 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 triple:

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

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 $M = M' \cup [C_{S_0}(u) \mid u \in D_0]$. 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:

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

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

- $M = \bigcup_{i=0}^{\ell} [C_{S_i}(u) \mid u \in D_i] = [C_{S_0 \bullet \dots \bullet S_\ell}(u) \mid u \in \biguplus_{i=0}^{\ell} D_i] = [C_S(u) \mid u \in D]$

” \supseteq ” We prove the following, more general, property:

Let $\langle \{x_1, \dots, x_{\#B_0}\}, c, M \rangle \in \langle \langle B_0 \rangle \rangle_{\Delta}^{\#k}$ be a color triple. Then there exists a complete Δ -unfolding $B_0(x_1, \dots, x_{\#B_0}) \Rightarrow_{\Delta}^ \exists y_1 \dots \exists y_n \cdot \phi$, whose steps belong to a complete Δ -unfolding of A , 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 $S = (U, \sigma)$ such that $S \models^s \phi$ and $D \subseteq \text{supp}(\sigma) \setminus \{\mathfrak{s}(x_1), \dots, \mathfrak{s}(x_{\#B_0})\}$, $\text{card}(D) \leq k$, such that $C_S(\mathfrak{s}(x_i)) = c(x_i)$, for all $i \in [1..\#B_0]$ and $M = [C_S(u) \mid u \in D]$.*

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 (4), with a qpf formula ψ_0 . By definition, there exists a rule of the form (5) in Δ , with the same qpf formula ψ_0 . Then ψ_0 is satisfiable, because each Δ -unfolding of A yields a satisfiable formula. Then there exists a color triple:

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

such that $\langle \{x_1, \dots, x_{\#B_0}\}, c, M \rangle \in \langle X', c', M' \rangle \downarrow_{\{x_1, \dots, x_{\#B_0}\}}^{\#k}$, i.e., $c = c' \downarrow_{\{x_1, \dots, x_{\#B_0}\}}$ and $M \subseteq M' \cup \{c'(x) \mid x \in X' \setminus \{x_1, \dots, x_{\#B_0}\}\}$. Then, there exist $\langle \text{fv}(\psi_0), c_0, \emptyset \rangle \stackrel{\text{def}}{=} \gamma(\psi_0)$ and $\langle \{z_{i,1}, \dots, z_{i,\#B_i}\}, c'_i, M_i \rangle \in \langle \mathbb{B}_i \rangle_{\Delta}^{\#k} [x_1/z_{i,1}, \dots, x_{\#B_i}/z_{i,\#B_i}]$, for all $i \in [1..\ell]$, such that $\langle X', c', M' \rangle \in \langle \text{fv}(\psi_0), c_0, \emptyset \rangle \bullet^{\#k} \bullet^{\#k}_{i=1}^{\ell} \langle \{z_{i,1}, \dots, z_{i,\#B_i}\}, c'_i, M_i \rangle$, hence there exist $\langle \{x_1, \dots, x_{\#B_i}\}, c_i, M_i \rangle \in \langle \mathbb{B}_i \rangle_{\Delta}^{\#k}$ 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 $\mathbb{B}_i(x_1, \dots, x_{\#B_i}) \Rightarrow_{\Delta}^* \exists y_{j_{i,1}} \dots \exists y_{j_{i,k_i}} \cdot \psi_i$ 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 $\mathfrak{s}_i(x_j) = \mathfrak{s}_k(x_m)$ iff $z_{i,j}$ and $z_{k,m}$ are the same variable in the rule (5), 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 $S_i = (\mathbb{U}_i, \sigma_i)$ such that $S_i \models^{\mathfrak{s}_i} \psi_i$. 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 Δ 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 $\text{card}(D_i) \leq k$, $C_{S_i}(\mathfrak{s}_i(x_j)) = c_i(x_j)$, for all $j \in [1..\#B_i]$ and $M_i = [C_{S_i}(u) \mid u \in D_i]$.

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 $\mathbb{B}_i(z_{i,1}, \dots, z_{i,\#B_i})$ is a predicate atom that occurs on the right-hand side of the rule (5). A complete Δ -unfolding $\mathbb{B}_0(x_1, \dots, x_{\#B_0}) \Rightarrow_{\Delta}^* \exists y_1 \dots \exists y_n \cdot \psi$ is built from the rule (5) above, with qpf formula ψ_0 , followed by $\mathbb{B}_i(x_1, \dots, x_{\#B_i})\theta_i \Rightarrow_{\Delta}^* \exists y_{j_{i,1}} \dots \exists y_{j_{i,k_i}} \cdot \psi_i\theta_i$, for all $i \in [1..\ell]$. Hence $\psi = \psi_0 * \ast_{i=1}^{\ell} \psi_i\theta_i$ 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 Δ to be equality-free.

Then, we consider a structure $S_0 = (\mathbb{U}_0, \sigma_0)$ 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 Δ is equality-free. It is easy to check that the structures S_0, \dots, S_{ℓ} are pairwise locally disjoint, hence $S = (\mathbb{U}, \sigma) \stackrel{\text{def}}{=} S_0 \bullet \dots \bullet S_{\ell}$ is defined. Moreover, we have $S \models^{\mathfrak{s}} \psi$, because $\psi = \psi_0 * \ast_{i=1}^{\ell} \psi_i\theta_i$, $S_0 \models^{\mathfrak{s}} \psi_0$ 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:

$$C_S(\mathfrak{s}(x_j)) = C_{S_0 \bullet \dots \bullet S_\ell}(\mathfrak{s}(x_j)) = \bigoplus_{i=0}^{\ell} c_i(x_j) = c'(x_j) = c(x_j)$$

We consider the set $D \stackrel{\text{def}}{=} \{u \in \text{supp}(\sigma) \mid C_S(u) \in M\}$. Suppose, for a contradiction, that $\mathfrak{s}(x_i) \in D$, for some $i \in [1..\#B_0]$. Then $C_S(\mathfrak{s}(x_i)) \in M$, 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 $c(x_i) \in M_j = [C_{S_j} \mid u \in D_j]$, by the inductive hypothesis. Then there exists $k \in [1..\#B_j]$ such that $c(x_i) = c_j(x_k) = C_{S_j}(\mathfrak{s}_j(x_k))$, 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 $M = [C_S(u) \mid u \in D]$. “ \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, $C = C_S(\mathfrak{s}(z))$ 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, $C = C_S(u) = C_{S_i}(u)$, 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

Lemma 36. *For any equality-free SID Δ and a nullary predicate symbol A , such that each complete Δ -unfolding of A yields a satisfiable formula, one can effectively build a SID Γ and a nullary predicate P such that $\text{split}(\llbracket A \rrbracket_{\Delta}^c) = \llbracket P \rrbracket_{\Gamma}^c$.*

Proof (Proof of Lemma 36). “ \subseteq ” We prove first the following fact:

Fact 12 *Let $B_0(x_1, \dots, x_{\#B_0}) \Rightarrow_{\Delta}^* \exists y_1 \dots \exists y_n \cdot \phi$ 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 $B_0(x_1, \dots, x_{\#B_0}) / \xi_0 \Rightarrow_{\Gamma}^* \exists y_1 \dots \exists y_n \cdot \phi'$, 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 (5), for a qpf formula Ψ_0 . Then, there exist:

- unfoldings $B_i(x_1, \dots, x_{\#B_i}) \Rightarrow_{\Delta}^* \exists y_{j_{i,1}} \dots \exists y_{j_{i,k_i}} \cdot \phi_i$, where $y_{j_{i,1}}, \dots, y_{j_{i,k_i}} \in \{y_1, \dots, y_n\}$ and ϕ_i are qpf formulæ, for all $i \in [1..\ell]$, and
- structures $S_0 = (U_0, \sigma_0), \dots, S_\ell = (U_\ell, \sigma_\ell)$, such that $S_0 \bullet \dots \bullet S_\ell = S$, $S_0 \models^{\mathfrak{s}} \Psi_0$ and $S_i \models^{\mathfrak{s}} \phi_i \theta_i$, 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 $S_{i_h} \models^{\mathfrak{s}} \phi_{i_h} \theta_{i_h}$, we have $S_{i_h} \models^{\mathfrak{s} \circ \theta_{i_h}^{-1}} \phi_{i_h}$, 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 $\bar{J}_{i_h} \stackrel{\text{def}}{=} [1..\#B_i] \setminus J_{i_h}$, for all $h \in [1..k]$,
- qpf formulæ ψ'_0 and ψ''_0 satisfying points (i) and (ii) from the construction of Γ ,
- an equivalence relation $\Xi \stackrel{\text{def}}{=} (\zeta(\psi'_0) \cup \bigcup_{h=1}^k \xi_{i_h}(B_{i_h}(z_{i_h,1}, \dots, z_{i_h, \#B_{i_h}})))^{\bar{\cdot}}$.

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 (i).

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:

$$(i, j) \in \xi_0 \iff (x_i, x_j) \in \Xi$$

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 (iv-vii) from the definition of Γ , hence Γ contains a rule of the form (6), 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 $S' \in \text{split}(\llbracket A \rrbracket_{\Delta}^c)$ be a maximally connected substructure of a canonical model $S \in \llbracket A \rrbracket_{\Delta}^c$. 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 (5), 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 12, 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

(i) and (ii) are met. Let $\Xi \stackrel{\text{def}}{=} (\zeta(\Psi'_0) \cup \bigcup_{h=1}^k \xi_{i_h}(\mathbb{B}_{i_h}(z_{i_h,1}, \dots, z_{i_h, \#B_{i_h}})))^=$ 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 (7). This rule and the complete unfoldings $\mathbb{B}_{i_h}(z_{i_h,1}, \dots, z_{i_h, \#B_{i_h}}) \Rightarrow_{\Gamma}^* \exists y_{j_{h,1}} \dots \exists y_{j_{h,m_h}} \cdot \phi'_h$, 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 $S' \models^{\mathfrak{s}} \Psi'_0 * \bigstar_{h=1}^k \phi'_h$. 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 $S' \in \llbracket P \rrbracket_{\Gamma}^c$.

“ \supseteq ” We prove first two related facts. First, let $\mathbb{B}_0(x_1, \dots, x_{\#B_0}) / \xi_0 \Rightarrow_{\Gamma}^* \exists y_1 \dots \exists y_n \cdot \phi$ 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 $S \models^{\mathfrak{s}} \phi$. Given an equivalence class $I \subseteq J_0$ of ξ_0 , we define the structure:

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

Fact 13 $\text{reach}_{\xi}^{\mathfrak{s}}(I) \sqsubseteq^{mc} S$.

Proof. By induction on the length of the Γ -unfolding. Assume that the first rule of the unfolding is of the form (6), 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 (i-vii) from the definition of Γ . By point (vii), $\{x_i \mid i \in I\}$ is an equivalence class of $\Xi_{x_1, \dots, x_{\#B_0}}$ 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 (iii).

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

Second, let $\mathbb{B}_0(x_1, \dots, x_{\#B_0}) \Rightarrow_{\Delta}^* \exists y_1 \dots \exists y_m \cdot \phi'$ be the complete Δ -unfolding obtained by replacing each rule of the form (6) with its corresponding rule (5) 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 14 *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 $S' = (U, \sigma')$, 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 $U \setminus \{\mathfrak{s}(y_1), \dots, \mathfrak{s}(y_n)\}$ 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 U is infinite. Because of the assumption that each predicate defined by a rule from Δ occurs on some complete Δ -unfolding of A , there exists a complete Δ -unfolding:

$$A \Rightarrow_{\Delta} \dots \Rightarrow_{\Delta} B_0(z_1, \dots, z_{\#B_0}) * \phi'' \Rightarrow_{\Delta} (\exists y_1 \dots \exists y_m \cdot \phi') [x_1/z_1, \dots, x_{\#B_0}/z_{\#B_0}] * \phi''$$

where ϕ'' is a predicate-free formula, possibly containing existential quantifiers. Since every complete Δ -unfolding of A yields a satisfiable formula, there exists a store \mathfrak{s}'' that agrees with $\mathfrak{s}' \circ [z_1/x_1, \dots, z_{\#B_0}/x_{\#B_0}]$ over y_1, \dots, y_m and a structure S'' , such that:

$$S'' \models^{\mathfrak{s}''} (\exists y_1 \dots \exists y_m \cdot \phi') [x_1/z_1, \dots, x_{\#B_0}/z_{\#B_0}] * \phi''$$

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 $S'' = S' \bullet S'''$, $S' \models^{\mathfrak{s}''} \phi' [x_1/z_1, \dots, x_{\#B_0}/z_{\#B_0}]$, 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 (i) 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 $\bar{S} = (\mathsf{U}, \bar{\sigma})$, 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 13, 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 $S \in [[\mathsf{P}]]_{\Gamma}^{\mathfrak{s}}$ be a canonical Γ -model of P , i.e., there exists a complete Γ -unfolding $\mathsf{P} \Rightarrow_{\Gamma}^* \exists y_1 \dots \exists y_n \cdot \phi$, 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 (7), with a qpf formula ψ_0 . Then there exist Γ -unfoldings $B_i(x_1, \dots, x_{\#B_i}) / \xi_i \Rightarrow_{\Gamma}^* \exists y_{j_1,1} \dots \exists y_{j_i,m_i} \cdot \phi_i$, for some sets $J_i \subseteq [1..\#B_i]$ and equivalence relations $\xi_i \subseteq J_i \times J_i$, for $i \in [1..\ell]$, such that:

$$\phi = \psi_0 * \bigast_{i \in [1..\ell], J_i \neq \emptyset} \phi_i \theta_i \quad (8)$$

where $\theta_i \stackrel{\text{def}}{=} [x_1/z_{i,1}, \dots, x_{\#B_i}/z_{i,\#B_i}]$, $i \in [1..\ell]$. 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:

$$\begin{aligned} B_0(x_1, \dots, x_{\#B_0}) &\Rightarrow_{\Delta} \exists y_1 \dots \exists y_m \cdot \psi'_0 * \bigast_{i=1}^{\ell} B_i(z_{i,1}, \dots, z_{i,\#B_i}) \Rightarrow_{\Delta} \dots \\ &\Rightarrow_{\Delta}^* \exists y_1 \dots \exists y_p \cdot \psi'_0 * \bigast_{j=1}^k \phi'_{i_j} \theta_{i_j} * \eta \end{aligned} \quad (9)$$

for qpf formulae $\psi'_0, \phi'_{i_1}, \dots, \phi'_{i_k}$ and predicate-free formula η . Consider the equivalence relation:

$$\Xi \stackrel{\text{def}}{=} (\zeta(\psi_0) \cup \bigcup_{j=1}^k \xi_{i_j}(B_{i_j}(z_{i_j,1}, \dots, z_{i_j,\#B_{i_j}}))) =$$

over $\{x_1, \dots, x_{\#B_0}\} \cup \{y_1, \dots, y_m\}$. By point (viii), Ξ has a single equivalence class X such that:

- $X \cap \{x_1, \dots, x_{\#B_0}\} = \emptyset$, and
- the sets $I_j \stackrel{\text{def}}{=} \{h \in [1..\#B_{i_j}] \mid z_{i_j,h} \in X\}$ 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 14, 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 $\text{reach}_{\xi_{i_j}}^{\mathfrak{s}'}(I_{j,h}) \sqsubseteq^{mc} S'_j$, for all $j \in [1..k]$ and $h \in [1..q_j]$. We argue that $S_j = \bullet_{h=1}^{q_j} \text{reach}_{\xi_{i_j}}^{\mathfrak{s}'}(I_{j,h}) \sqsubseteq^{mc} S'_j$. 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 A .

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

$$\begin{aligned}
A &\Rightarrow_{\Delta} \dots \Rightarrow_{\Delta} B_0(z_{0,1}, \dots, z_{0,\#B_0}) * \zeta \\
&\Rightarrow_{\Delta} (\exists y_1 \dots \exists y_m \cdot \psi_0 * \bigstar_{i=1}^{\ell} B_i(z_{i,1}, \dots, z_{i,\#B_i})) [x_1/z_{0,1}, \dots, x_{\#B_0}/z_{0,\#B_0}] * \zeta \\
&\Rightarrow_{\Delta} \dots \Rightarrow_{\Delta} (\exists y_1 \dots \exists y_n \cdot \psi'_0 * \bigstar_{i=1}^k \phi'_i \theta_i * \eta) [x_1/z_{0,1}, \dots, x_{\#B_0}/z_{0,\#B_0}] * \zeta
\end{aligned}$$

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, $S' \bullet S'' \in \llbracket A \rrbracket_{\Delta}^c$ and, since $S \sqsubseteq^{mc} S'$, we obtain $S \sqsubseteq^{mc} S' \bullet S''$, leading to $S \in \text{split}(\llbracket A \rrbracket_{\Delta}^c)$. \square