



HAL
open science

The theory of concatenation over finite models

Dominik D Freydenberger, Liat Peterfreund

► **To cite this version:**

Dominik D Freydenberger, Liat Peterfreund. The theory of concatenation over finite models. 2021.
hal-03104159

HAL Id: hal-03104159


<https://inria.hal.science/hal-03104159>

Preprint submitted on 8 Jan 2021

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.

The theory of concatenation over finite models

Dominik D. Freydenberger 

Loughborough University, Loughborough, United Kingdom

Liat Peterfreund

IRIF - CNRS and Université de Paris, Paris, France

Edinburgh University, Edinburgh, United Kingdom

Abstract

We propose FC, a logic on words that combines the previous approaches of finite-model theory and the theory of concatenation. It has immediate applications to spanners, a formalism for extracting structured data from text that has recently received considerable attention in database theory. In fact, FC is designed to be to spanners what FO is to relational databases.

Like the theory of concatenation, FC is built around word equations; in contrast to it, its semantics are defined to only allow finite models, by limiting the universe to a word and all its subwords. As a consequence of this, FC has many of the desirable properties of FO[<], while being far more expressive. Most noteworthy among these desirable properties are sufficient criteria for efficient model checking and capturing various complexity classes by extending the logic with appropriate closure or iteration operators.

These results allow us to obtain new insights into and techniques for the expressive power and efficient evaluation of spanners. More importantly, FC provides us with a general framework for logic on words that has potential applications far beyond spanners.

2012 ACM Subject Classification Theory of computation → Database query languages (principles); Theory of computation → Finite Model Theory; Theory of computation → Logic and databases; Mathematics of computing → Combinatorics on words

Keywords and phrases document spanners, finite-model theory, theory of concatenation, word equations, descriptive complexity, information extraction, database theory, model checking

Related Version This full version is available at <https://arxiv.org/abs/1912.06110>.

Acknowledgements The authors thank Joel D. Day, Manfred Kufleitner, Leonid Libkin, and Sam M. Thompson for their helpful comments.

1 Introduction

Document spanners (or just *spanners*) are a rule-based framework for information extraction that was proposed by Fagin, Kimelfeld, Reiss, and Vansummeren [25] to study the formal properties of the query language AQL of IBM's SystemT for information extraction. On an intuitive level, the main idea of document spanners can be understood as querying a word like one would query a relational database. More specifically, *extractors* turn text into tables of position intervals, which are then combined using a relational algebra.

Hence, after the extraction step, spanners act like relational databases; and the relational algebra that is common to both maps directly into first-order logic (FO). In fact, relational databases were conceived as an application of FO over relations. This connection has been maintained over the decades, leading to deep insights that touch basically every aspect of query languages that one might want to examine (see e.g. Abiteboul, Hull, and Vianu [1] and the conferences ICDT and PODS). This close connection between relational databases and FO does not directly translate to spanners: These do not have explicit tables but build them via extractors; and this step makes many canonical approaches for relational databases nonviable in a spanner setting (see [32]). This raises the question whether there is a logic on words that can be to spanners what FO is to relational databases.

One of the most common approaches to logic on words is the one that is found in finite-model theory: A word is viewed as a finite linear order (a sequence of positions); for every letter a in the terminal alphabet Σ , one uses a unary predicate $P_a(x)$ to express that a occurs in the x -th position of the word. We refer to this logic as $\text{FO}[\langle]$. On finite universes, $\text{FO}[\langle]$ can express exactly the star-free languages, a subclass of the regular languages (see e.g. Straubing [62]). But even the most basic document spanners from [25] can express all regular languages; and the *core spanners* that correspond to AQL can use word equality selections that allow the definition of non-regular languages like the language of all words of the form ww . A different approach to logic on words is the *theory of concatenation*. First defined by Quine [59], this logic is built on *word equations*, that is, equations of the form $xx \doteq yyy$, where variables like x and y stand for words from Σ^* (instead of representing positions). While less prominent than $\text{FO}[\langle]$, the theory of concatenation has been studied extensively since the 1970s, with particular emphasis on word equations (see related work).

A connection between core spanners and the existential-positive theory of concatenation was first observed in [31]; building on this, [30] introduced SpLog , a variant of that fragment that has the same expressive power as core spanners (with efficient conversions in both directions between these formalisms). Furthermore, [30] also established that the situation is the same when comparing SpLog^\neg (SpLog with negation) to *generalized core spanners*, which are extended with a difference operator (see [25, 58]). Most literature on the theory of concatenation does not use negation. This is usually justified by the fact that, as shown by Quine [59], negation leads to an undecidable theory (i. e., satisfiability is undecidable).

Contrast this to FO over finite models: By Trakhtenbrot’s theorem, satisfiability is undecidable; but the model checking problem is decidable (see e.g. [24, 46]). We can observe the same situation for SpLog^\neg and generalized core spanners: While satisfiability is undecidable (see [25]), model checking and evaluation are in PSPACE (see [30]). This is not surprising, as both models operate on an input word, which is essentially a finite model. But from the perspective of the theory of concatenation (as commonly used in literature), the notion of finite models does not really make sense – there, the universe is Σ^* .

In [30], the “finite model properties” of SpLog and SpLog^\neg are ensured through the syntax: One special variable represents the input word, and the structure of the word equations in the formulas ensures that all variables can only be subwords of that word. While this approach allows one to capture spanners succinctly, the resulting formulas can quickly become unwieldy and require a lot of syntactic sugar. After working with SpLog for a while, the authors decided that it is not as natural to use as FO is for databases. For example, [58] extended core spanners with recursion (inspired by *Datalog*), thus capturing those spanners that can be computed in polynomial time. But mirroring this in SpLog with least fixed-points (analogous to the relation of $\text{FO}[\langle]$ and *Datalog*) turned out to be very inconvenient.

The present paper proposes a different approach: Instead of ensuring finite models by restricting the syntax of the theory of concatenation, we change the semantics. The resulting logic FC , the *finite model theory of concatenation*, treats k -tuples of words as structures, and the set of their subwords as the universe¹. This logic may use various *constraints* (that is, *predicates*); in particular regular constraints that allow to check whether a word belongs to some regular language (these are needed for spanners). Although spanners are an important motivation of FC , we keep the constraints as general as reasonably possible. This is partly due to the fact that the theory of concatenation has recently been linked to string solvers,

¹ In most cases, we will only be interested in the case of $k = 1$, mirroring how spanners operate on a single word. But the more general case is useful when studying the expressive power of FC , see Section 3.2.

where length constraints are relevant as well; but also because our results do not require any artificial restrictions. Furthermore, we also examine $FC[]$ (FC without constraints).

To connect FC to $FO[<]$, Section 3.3 introduces a variant of the latter that is extended with a word equality predicate Eq . Our first main result is that the resulting logic $FO[Eq]$ has the same expressive power as FC ; and there are efficient conversion between both logics. Hence, we can pick whichever logic is most suitable for the task at hand. In Section 3.4, the next set of main results is that $FC[REG]$ (FC with regular constraints) can be used as a logic for spanners; and this applies to various fragments of both models. Compared to $SpLog$, the syntax of $FC[REG]$ is less cumbersome, but the conversions are as efficient. Section 4.1 shows that model checking for FC has the same complexity as for FO . In particular, the width of a formula is a parameter that allows us to make the problem tractable. Building combinatorics on words, we give a sufficient criterion for of word equations that can be decomposed into formulas of bounded width. In Section 4.2, we see that FC that is extended with various iteration operators captures L , NL , P , and $PSPACE$; analogously to classical results for $FO[<]$. This allows us to define $DataSpLog$, a variant of $Datalog$ on words that uses word equations and captures P . Section 4.3 uses the Feferman-Vaught theorem to show an inexpressibility result for $FC[]$ that also provides us with the first inexpressibility proof for generalized core spanners over non-unary terminal alphabets. Finally, Section 4.4 deals with the static analysis and relative succinctness of formulas.

Related work *Document spanners* were introduced by Fagin, Kimelfeld, Reiss, and Vansummeren [25] and have received considerable attention in the last few years. The two main areas of interest are expressive power [25, 30, 31, 53, 55, 58] and efficient evaluation [3, 27, 32, 53, 54, 57]; further topics include updates [3, 34, 48], cleaning [26], distributed query planning [21], and a weighted variant [22].

The *theory of concatenation* was shown to be undecidable by Quine [59]. Later, Büchi and Senger [11] showed that its existential fragment is decidable; building on Makanin’s algorithm that decides the satisfiability of *word equations* [50] (this is also discussed in [41]). A fairly recent survey on the satisfiability of word equations is [20]. More current research on word equations and the theory of concatenation can be found in e.g. [14, 18, 19, 56]. Word equations have recently attracted attention in the context of string solvers (see e.g. [18] for further details and references) and in database theory (see [5, 30]). In particular, [47] presents a fragment of the theory of concatenation that is decidable, see [12], but much more restricted than the logics in the present paper. In contrast to this, the areas of *finite-model theory* in general and *database theory* in particular are much larger. We refer to [24, 46] and [1], respectively. The relative succinctness of formulas that is examined in Section 4.4 was examined in an FO -setting by Berkholtz and Chen [7], with very different results.

There has been a significant amount of work on query languages *string databases*, see [6, 8, 36, 37]. These treat words as entries of the database. Unlike the present paper, they do not operate on a word (or tuple of words), which means that they lack the finite-model property that we discuss in Section 3.1. Furthermore, they offer transformation operations that greatly increase the expressive power of the model and usually allow the query language to express Turing-complete functions from words to words. On a first glance, the present paper might seem more related to [38] than it is: Although that paper also uses a “subword” relation and represents it with \sqsubseteq , this refers to scattered (i.e., non-continuous) subwords. What we call the theory of concatenation C is called $FO(A^*, \cdot)$ in [38].

2 Preliminaries

For two logic fragments \mathcal{F}_1 and \mathcal{F}_2 , we write $\mathcal{F}_1 \equiv \mathcal{F}_2$ to denote that for every formula in one fragment, there is an equivalent formula in the other. We write $\mathcal{F}_1 \equiv_{\text{poly}} \mathcal{F}_2$ if these formulas can be constructed in polynomial time.

Let ε denote the *empty word*. We use $|x|$ for the *length* of a word, a formula, or a regular expression x , or the *number of elements* of a finite set x . We treat tuples as sets, except that we define $|\vec{x}| = k$ for every k -tuple \vec{x} , even if elements are repeated. For a tuple \vec{w} of words, let $\|\vec{w}\| := \sum_{w \in \vec{w}} |w|$. A word v is a *subword* of a word w , written $v \sqsubseteq w$, if there exists (possibly empty) words p, s with $w = pvs$. For words x and y , let $x \sqsubseteq_p y$ (x is a *prefix* of y) if $y = xs$ for some s , and $x \sqsubset_p y$ if $x \sqsubseteq_p y$ and $x \neq y$. For alphabets A, B , a *morphism* is a function $h: A^* \rightarrow B^*$ with $h(u \cdot v) = h(u) \cdot h(v)$ for all $u, v \in A^*$. To define h , it suffices to define $h(a)$ for all $a \in A$. For k -tuples $\vec{x} = (x_1, \dots, x_k)$ over A^* , let $h(\vec{x}) := (h(x_1), \dots, h(x_k))$.

Let Σ be a finite *terminal alphabet*, and let Ξ be an infinite *variable alphabet* that is disjoint from Σ . Unless explicitly stated, we assume that $|\Sigma| \geq 2$.

2.1 Patterns and the theory of concatenation

A *pattern* is a word from $(\Sigma \cup \Xi)^*$. For every pattern $\eta \in (\Sigma \cup \Xi)^*$, let $\text{var}(\eta)$ denote the set of variables that occur in η . A *pattern substitution* (or just *substitution*) is a partial morphism $\sigma: (\Sigma \cup \Xi)^* \rightarrow \Sigma^*$ that satisfies $\sigma(a) = a$ for all $a \in \Sigma$. When applying a substitution σ to a pattern η , we assume $\text{dom}(\sigma) \supseteq \text{var}(\eta)$, where $\text{dom}(\sigma)$ denotes the domain of σ .

A *word equation* is a pair of patterns, that is, a pair (η_L, η_R) with $\eta_L, \eta_R \in (\Sigma \cup \Xi)^*$. We also write $\eta_L \doteq \eta_R$, and refer to η_L and η_R as the *left side* and the *right side* of the word equation (respectively). A *solution* of $\eta_L \doteq \eta_R$ is a substitution σ with $\sigma(\eta_L) = \sigma(\eta_R)$.

The theory of concatenation combines word equations with first-order logic. First the syntax: Let \mathcal{K} be a set of predicates on words (the *constraints*), each of which is identified with a *constraint symbol* $\kappa \in \mathcal{K}$ of some arity $\text{ar}(\kappa)$. The set $\mathbf{C}[\mathcal{K}]$ of *theory of concatenation formulas with \mathcal{K} -constraints* uses word equations $(\eta_L \doteq \eta_R)$ with $\eta_L, \eta_R \in (\Sigma \cup \Xi)^*$ and constraints $\kappa(\vec{x})$ with $\kappa \in \mathcal{K}$ and $\vec{x} \in \Xi^{\text{ar}(\kappa)}$ as atoms. These are combined with conjunction, disjunction, negation, and quantifiers with variables from Ξ . We use \mathbf{C} for the union of all $\mathbf{C}[\mathcal{K}]$. For every $\varphi \in \mathbf{C}$, we define its set of *free variables* $\text{free}(\varphi)$ by $\text{free}(\eta_L \doteq \eta_R) := \text{var}(\eta_L) \cup \text{var}(\eta_R)$ and $\text{free}(\kappa(\vec{x})) = \vec{x}$; extending this canonically.

The semantics build on solutions of word equations and treat constraints as predicates: For all $\varphi \in \mathbf{C}[\mathcal{K}]$ and all pattern substitutions σ with $\text{dom}(\sigma) \supseteq \text{free}(\varphi)$, we define $\sigma \models \varphi$ as follows: Let $\sigma \models (\eta_L \doteq \eta_R)$ if $\sigma(\eta_L) = \sigma(\eta_R)$ and $\sigma \models \kappa(\vec{x})$ if $\sigma(\vec{x}) \in \kappa$. For the existential (or universal) quantifier, we say $\sigma \models \exists x: \varphi$ (or $\sigma \models \forall x: \varphi$) if $\sigma_{x \mapsto w} \models \varphi$ holds for an (or all) $w \in \Sigma^*$, where $\sigma_{x \mapsto w}$ is defined by $\sigma_{x \mapsto w}(x) := w$ and $\sigma_{x \mapsto w}(y) := \sigma(y)$ for all $y \in (\Sigma \cup \Xi) - \{x\}$. The connectives' semantics are defined canonically.

To avoid complexity issues, we assume that for each $\kappa \in \mathcal{K}$, we can evaluate $\kappa(\vec{w})$ in space that is logarithmic in $\|\vec{w}\|$ (and, hence, in polynomial time). When this paper is concerned with specific choices of \mathcal{K} , it mostly focuses on the *constraint-free* formulas from $\mathbf{C}[\emptyset] = \mathbf{C}[\emptyset]$ and on $\mathbf{C}[\text{REG}]$, where REG are the *regular constraints* which we write “ $x \dot{\in} \alpha$ ” for $x \in \Xi$ and any regular expression² α , with $\sigma \models x \dot{\in} \alpha$ if $\sigma(x) \in \mathcal{L}(\alpha)$, where $\mathcal{L}(\alpha)$ is the language of α .

► **Example 2.1.** Let $\varphi := (x \doteq yy) \wedge y \dot{\in} (\text{ab})^*$. Then $\sigma \models \varphi$ if and only if $\sigma(y) = w$ and $\sigma(x) = ww$ for some $w \in (\text{ab})^*$.

² We could allow finite automata instead; for the purpose of this paper, the distinction does not matter.

We freely add and omit parentheses as long as the meaning stays clear and contract quantifiers $\exists x_1: \dots \exists x_k: \varphi$ to $\exists x_1, \dots, x_k: \varphi$. We define EC, the *existential fragment* of C, as those formulas that do not use universal quantifiers and that apply negation only to word equations or constraints. The *existential-positive fragment* EPC allows neither universal quantifiers, nor negations. We shall use the same notation for other logics that we define.

2.2 Spans and document spanners

This section briefly introduces spans and spanners. These definitions are kept at a bare minimum, as the paper will mostly approach spanners through the equivalent logics. Full definitions can be found in [25, 30], which also include numerous examples.

Let $w := a_1 \dots a_n$ with $n \geq 1$ and $a_1, \dots, a_n \in \Sigma$. A *span* of w is an interval $[i, j)$ with $1 \leq i \leq j \leq n + 1$. It describes the word $w_{[i, j)} = a_i \dots a_{j-1}$. For finite $V \subset \Xi$ and $w \in \Sigma^*$, a (V, w) -*tuple* is a function μ that maps each variable in V to a span of w . A *spanner* with variables V is a function P that maps every $w \in \Sigma^*$ to a set $P(w)$ of (V, w) -tuples. We use $\text{SVars}(P)$ for the variables of a spanner P . Like [25], we base spanners on *regex formulas*; regular expressions with variable bindings of the form $x\{\alpha\}$. This matches the same words as the sub-expression α and binds x to the corresponding span of w . A regex formula is *functional* if on every word, every match binds every variable exactly once. The set of functional regex formulas is RGX. For every $\alpha \in \text{RGX}$, we define a spanner $\llbracket \alpha \rrbracket$ as follows: Every match on a word w defines an $(\text{SVars}(\alpha), w)$ -tuple μ , where each $\mu(x)$ is the span where x was bound. Then $\llbracket \alpha \rrbracket(w)$ contains all these μ from matches of α on w .

We use the spanner operations union \cup , natural join \bowtie , projection π , set difference $-$, and equality selection $\zeta^=$, where $\zeta_{x,y}^= P(w)$ is the set of all $\mu \in P(w)$ with $w_{\mu(x)} = w_{\mu(y)}$. The class of *generalized core spanner representations* RGX^{core} consists of combinations of RGX and any of the five operators; the *core spanner representations* RGX^{core} exclude set difference.

In [30], the C[REG]-fragments SpLog and SpLog^- were introduced as alternatives to RGX^{core} and $\text{RGX}^{\text{gcore}}$, respectively. We discuss this further in Section 3.4.

3 A finite-model variant of the theory of concatenation

3.1 The logic FC

By Quine [59], the satisfiability problem for C[] (given $\varphi \in \text{C}[\]$, is there a σ with $\sigma \models \varphi$) is undecidable. As shown by Durnev [23], this holds even for the prefix class $\forall \exists^3$. Hence, given φ and σ , it is also undecidable whether $\sigma \models \varphi$. Compare this to relational first-order logic on finite models (see e. g. [24, 46]): There, satisfiability is also undecidable, but the model checking problem – given a finite structure \mathcal{A} and a formula φ , does $\mathcal{A} \models \varphi$ hold – is decidable. In contrast to this, variables in C can be mapped to any element of Σ^* . Thusly, the universe of C is always assumed to be Σ^* , making all models infinite.

We propose FC, the *finite model version of the theory of concatenation*. As this logic is still built around word equations, we distinguished variables to represent the structure (a word or a tuple of words) in word equations. Accordingly, these are called *structure variables*. As we shall see in Section 3.2, allowing multiple structure variables is convenient. For each instance of a query evaluation or model checking problem, we fix the content of the structure variables to represent the input. As C-variables represent words (as opposed to positions in words or graphs like for FO), the universe for FC is the set of subwords of the structure variables. We define FC following these key ideas, starting with the syntax:

► **Definition 3.1.** Choose a tuple \vec{s} of pairwise distinct variables from Ξ , which we call the structure variables. The set $\text{FC}(\vec{s})$ of FC-formulas with structure variables \vec{s} is the set of all $\varphi \in \mathcal{C}$ that satisfy the following two conditions:

1. every word equation in φ has exactly one variable from Ξ on its left side, and
2. no quantifier in φ binds a structure variable.

We use $\text{struc}(\varphi)$ to denote the set of structure variables of φ . In contrast to \mathcal{C} , we modify $\text{free}(\varphi)$ by declaring that structure variables are not free. A formula is Boolean if it has no free variables. We use FC for the union of all $\text{FC}(\vec{s})$, and we define $\text{FC}[\mathcal{K}]$ as for $\mathcal{C}[\mathcal{K}]$.

Hence, structure variables do not have to appear in φ at all; but if they appear, they may not be bound by quantifiers. For variable tuples \vec{s} and \vec{x} , we write $\varphi(\vec{s})(\vec{x})$ to indicate that $\text{struc}(\varphi) = \vec{s}$, $\text{free}(\varphi) = \vec{x}$, and that the elements in each of \vec{s} and \vec{x} are pairwise different (also, note that \vec{s} and \vec{x} are disjoint by definition). We also write $\varphi(\vec{s})$ (or $\varphi(\vec{x})$) if we are only interested in the structure (or the free) variables and the context is clear.

Although FC-formulas are technically \mathcal{C} -formulas, we modify their semantics to ensure that the universe is the set of subwords of structure variables:

► **Definition 3.2.** For every $\varphi \in \text{FC}$ and every pattern substitution σ with $\text{dom}(\sigma) \supseteq \text{struc}(\varphi) \cup \text{free}(\varphi)$, we define $\sigma \models_{\text{struc}(\varphi)} \varphi$ by extending each step of the recursive definition of $\sigma \models \varphi$ in the semantics of \mathcal{C} with the additional condition that for every $x \in \text{dom}(\sigma)$, there is some $\mathfrak{s} \in \text{struc}(\varphi)$ with $\sigma(x) \sqsubseteq \sigma(\mathfrak{s})$. If the context is clear, in particular after stating $\varphi \in \text{FC}$, we abuse notation and write $\sigma \models \varphi$ to denote $\sigma \models_{\text{struc}(\varphi)} \varphi$.

An immediate consequence of this definition is that $\sigma \models_{\text{struc}(\varphi)} \varphi$ implies that for every $x \in \text{free}(\varphi)$, there is an $\mathfrak{s} \in \text{struc}(\varphi)$ with $\sigma(x) \sqsubseteq \sigma(\mathfrak{s})$. Hence, if we fix the images of the structure variables, the universe is restricted to subwords of these images. In contrast to \mathcal{C} , where we can understand Σ^* as universe, this restriction ensures that all models of an FC-formula are finite. We say $\varphi, \psi \in \text{FC}$ are *equivalent*, written $\varphi \equiv \psi$, if $\text{free}(\varphi) = \text{free}(\psi)$, $\text{struc}(\varphi) = \text{struc}(\psi)$, and $\sigma \models_{\text{struc}(\varphi)} \varphi$ holds if and only if $\sigma \models_{\text{struc}(\psi)} \psi$.

Before we look at examples, we introduce definitions for querying and model checking. For $\varphi(\vec{x}) \in \text{FC}(\vec{s})$ and $\vec{w} \in (\Sigma^*)^{|\vec{s}|}$, let $\llbracket \varphi \rrbracket \langle \vec{w} \rangle$ denote the set of all $\sigma(\vec{x})$ such that $\sigma \models_{\vec{s}} \varphi$ and $\sigma(\vec{s}) = \vec{w}$. If context allows, we also write $\sigma \in \llbracket \varphi \rrbracket \langle \vec{w} \rangle$ for $\sigma \models_{\vec{s}} \varphi$ and $\sigma(\vec{s}) = \vec{w}$. If φ is Boolean, $\llbracket \varphi \rrbracket \langle \vec{w} \rangle$ is either \emptyset (“false”) or the set that contains the empty tuple $()$ (“true”).

► **Definition 3.3.** A Boolean formula $\varphi(\vec{s})() \in \text{FC}$ defines the $|\vec{s}|$ -ary relation $\mathcal{R}(\varphi)$ of all \vec{w} with $\llbracket \varphi \rrbracket \langle \vec{w} \rangle = \{()\}$. If $|\vec{s}| = 1$, we write $\mathcal{L}(\varphi)$ instead of $\mathcal{R}(\varphi)$.

► **Example 3.4.** Let $\varphi_1(\mathfrak{s})(x) := (x \doteq x)$ and $\varphi_2(\mathfrak{s})(x) := \exists p, s: (\mathfrak{s} \doteq p x s)$. Then for every $w \in \Sigma^*$, we have $\llbracket \varphi_1 \rrbracket \langle w \rangle = \llbracket \varphi_2 \rrbracket \langle w \rangle$, which is the set of all σ with $\sigma(x) \sqsubseteq \sigma(\mathfrak{s}) = w$. Let $\varphi_3(\mathfrak{s}_1, \mathfrak{s}_2)() := \exists p, s: (\mathfrak{s}_1 \doteq p \mathfrak{s}_2 s)$. Then $\mathcal{R}(\varphi_3) = \{(u, v) \mid u \in \Sigma^*, u \sqsupseteq v\}$. The formula $\varphi_4(\mathfrak{s})() := \exists x: (\mathfrak{s} \doteq x x)$ defines the language $\mathcal{L}(\varphi_4) = \{ww \mid w \in \Sigma^*\}$. Finally, let $\varphi_5(\mathfrak{s}_1, \mathfrak{s}_2, \mathfrak{s}_3)() := \exists x: (x \doteq \varepsilon)$ and $\varphi_6(\mathfrak{s}_1, \mathfrak{s}_2, \mathfrak{s}_3) := (\mathfrak{s}_1 \doteq \mathfrak{s}_1)$. Then $\mathcal{R}(\varphi_5) = \mathcal{R}(\varphi_6) = (\Sigma^*)^3$.

► **Example 3.5.** A language is called *star-free* if it is defined by a regular expression α that is constructed from the empty set \emptyset , terminals $a \in \Sigma$, concatenation \cdot , union \cup , and complement α^C . Given such an α , we define $\varphi_\alpha(\mathfrak{s})() := \exists x: (\mathfrak{s} \doteq x \wedge \psi_\alpha(x))$, where $\psi_\alpha(x)$ is defined recursively by $\psi_\emptyset(x) := \neg(x \doteq x)$, $\psi_a(x) := (x \doteq a)$, $\psi_{(\alpha_1 \cdot \alpha_2)}(x) := \exists x_1, x_2: (x \doteq x_1 x_2 \wedge \psi_{\alpha_1}(x_1) \wedge \psi_{\alpha_2}(x_2))$, $\psi_{(\alpha_1 \cup \alpha_2)}(x) := \psi_{\alpha_1}(x) \vee \psi_{\alpha_2}(x)$, and $\psi_{\alpha^C}(x) := \neg \psi_\alpha(x)$. Then $\sigma \models \varphi_\alpha$ if and only if $\sigma(x) \in \mathcal{L}(\alpha)$ and $\sigma(x) \sqsubseteq \sigma(\mathfrak{s})$. Thus, $\mathcal{L}(\varphi_\alpha) = \mathcal{L}(\alpha)$.

As star-free languages are exactly those languages that can be expressed by relational first-order logic, this raises the question how those two approaches are related. We explore this in Section 3.3. But before that, we extend our formal toolkit for working with FC.

3.2 FC and relations

We next examine how FC can be used to define relations, and how this can be used to compose formulas in a way that is similar to FO. One side-effect of the finite universe of FC is that some relations are harder to express: For example, in C, we can use word equations like $xy \doteq yx$. Restricting left sides of equations to one variable is not an actual restriction for C, as we can rewrite this to $\exists z: (z \doteq xy \wedge z \doteq yx)$ instead. This does not work in FC, as z (and hence the whole equation) needs to “fit” into one of the structure variables. Thus, it is not obvious that a relation that is definable with $k \geq 2$ structure variables can be expressed with one structure variable. We address this through the following definition, which adapts SpLog-selectability from [30] which, in turn, is based on spanner selectability from [25].

► **Definition 3.6.** *For every $\mathcal{F} \subseteq \text{FC}$ and $k \geq 1$, a relation $R \subseteq (\Sigma^*)^k$ is called \mathcal{F} -selectable if for every $\varphi \in \mathcal{F}$ and all k -tuples \vec{x} over $\text{free}(\varphi)$, there exists $\varphi_{\vec{x}}^R \in \mathcal{F}$ with $\text{free}(\varphi) = \text{free}(\varphi_{\vec{x}}^R)$ and $\text{struc}(\varphi) = \text{struc}(\varphi_{\vec{x}}^R)$ such that $\sigma \models \varphi_{\vec{x}}^R$ if and only if $\sigma \models \varphi$ and $\sigma(\vec{x}) \in R$.*

Intuitively speaking, the fact that a relation R is selectable in an FC-fragment \mathcal{F} means that adding constraint symbols that represent R to \mathcal{F} does not increase the expressive power.

For $k \geq 1$, we use $\text{FC}_{(k)}$ to denote the union of all $\text{FC}(\vec{s})$ with $|\vec{s}| = k$.

► **Lemma 3.7.** *Let $\mathcal{F} \in \{\text{EPFC}[\mathcal{K}], \text{EFC}[\mathcal{K}], \text{FC}[\mathcal{K}]\}$ for some \mathcal{K} . A relation $R \subseteq (\Sigma^*)^k$ is \mathcal{F} -selectable if and only if $R = \mathcal{R}(\varphi)$ for some $\varphi \in \mathcal{F} \cap \text{FC}_{(k)}[\mathcal{K}]$.*

The proof also implies, firstly, that R is $\text{FC}[\mathcal{K}]$ -definable if and only if $\text{FC}[\mathcal{K}] \equiv \text{FC}[\mathcal{K} \cup \{R\}]$ and, secondly, that when constructing $\varphi_{\vec{x}}^R$, we can just use the formula that defines R as a subformula and assume that its structure variables are “normal” variables. The authors find this much handier than the corresponding results for SpLog and spanners (Lemma 5.1 in [30], Proposition 4.15 in [25]). This is the main reason why we allow multiple structure variables. We use this to introduce some relations that can be used in the definition of formulas in all fragments of FC where Lemma 3.7 applies; e. g., φ_2 in Example 3.4 defines the subword relation, which allows us to use \sqsubseteq as shorthand. The relations in Example 3.8 were discussed for SpLog in [30]; but as we use these throughout the paper, we revisit them for FC.

► **Example 3.8.** Let $\varphi^{\neq \varepsilon} := \bigvee_{a \in \Sigma} \exists y: x \doteq a y$. If we choose x as structure variable of φ , we have $\mathcal{L}(\varphi) = \Sigma^+$. Instead, we use $\varphi^{\neq \varepsilon}$ as subformula of $\varphi^{\sqsubset p} := \exists x: ((y_2 \doteq y_1 x) \wedge \varphi^{\neq \varepsilon}(x))$. For $\varphi^{\sqsubset p}(y_1, y_2)$, we have that $\mathcal{R}(\varphi^{\sqsubset p})$ is the set of all (u, v) with $u \sqsubset_p v$. We use $\varphi^{\sqsubset p}(y_1, y_2)$ in

$$\varphi^{\neq}(z_1, z_2) := \varphi^{\sqsubset p}(z_1, z_2) \vee \varphi^{\sqsubset p}(z_2, z_1) \vee \exists x, y_1, y_2: \bigvee_{a \in \Sigma} \bigvee_{b \in \Sigma - \{a\}} (z_1 \doteq x a y_1 \wedge z_2 \doteq x b y_2).$$

We can then define $\varphi(\mathfrak{s})(x) := \exists p_1, p_2, s_1, s_2: (\mathfrak{s} \doteq p_1 x s_1 \wedge \mathfrak{s} \doteq p_2 x s_2 \wedge \varphi^{\neq}(p_1, p_2))$, a query that returns all x that have at least two (potentially overlapping) occurrences in \mathfrak{s} .

The construction for \neq was used in [13] to show $\text{EC}[\neq] \equiv \text{EPC}[\neq]$. We use it analogously:

► **Lemma 3.9.** $\text{EFC}[\neq] \equiv_{\text{poly}} \text{EPFC}[\neq]$ and $\text{EFC}[\text{REG}] \equiv \text{EPFC}[\text{REG}]$.

As the transformation for regular constraints involves complementing regular expressions, it is not polynomial. We leave open whether no polynomial transformation exists.

3.3 FO[<] with word equality

In this section, we establish connections between FC and “classical” relational first-order logic. It is probably safe to say that in finite model theory, the most common way of applying

first-order logic to words is the logic $\text{FO}[\prec]$. This uses the equality \doteq , and a vocabulary that consists of a binary relation symbol \prec and unary relation symbols P_a for each $a \in \Sigma$. Every word $w = a_1 \cdots a_n \in \Sigma^+$ with $n \geq 1$ is represented by a structure \mathcal{A}_w with universe $\{1, \dots, n\}$. For every $a \in \Sigma$, the relation P_a consists of those i that have $a_i = a$. To simplify dealing with ε , we slightly deviate from this standard structure. For every $w \in \Sigma^*$, we extend \mathcal{A}_w to \mathcal{A}'_w by adding an additional “letter-less” node $|w| + 1$ that occurs in no P_a . Then we have a one-to-one correspondence between pairs (i, j) with $i \leq j$ from the universe of \mathcal{A}'_w and the spans $[i, j]$ of w (see Section 2.2), and $w = \varepsilon$ does not require a special case.

► **Definition 3.10.** $\text{FO}[\text{Eq}]$ extends $\text{FO}[\prec]$ with constants min and max , the binary relation symbol succ , and the 4-ary relation symbol Eq . For every $w \in \Sigma^*$ and the corresponding structure \mathcal{A}'_w , these symbols express $\text{min} = 1$, $\text{max} = |w| + 1$, $\text{succ} = \{(i, i + 1) \mid 1 \leq i \leq |w|\}$, and Eq contains those (i_1, j_1, i_2, j_2) with $i_1 \leq j_1$ and $i_2 \leq j_2$ such that $w_{[i_1, j_1]} = w_{[i_2, j_2]}$.

$\text{FO}[\text{Eq}, \mathcal{K}]$ extends $\text{FO}[\text{Eq}]$ by representing each $\kappa \in \mathcal{K}$ with a $2k$ -ary relation symbol $\hat{\kappa}$, where $k := \text{ar}(\kappa)$, for the relation of all $(i_1, j_1, \dots, i_k, j_k)$ such that $(w_{[i_1, j_1]}, \dots, w_{[i_k, j_k]}) \in \kappa$. We write $\alpha \in \llbracket \varphi \rrbracket(w)$ to denote that α is a satisfying assignment for φ on \mathcal{A}'_w .

► **Example 3.11.** The $\text{FO}[\text{Eq}]$ -formula $\exists x: \text{Eq}(\text{min}, x, x, \text{max})$ defines $\{ww \mid w \in \Sigma^*\}$.

When comparing the expressive power of FC and $\text{FO}[\text{Eq}]$, we need to address that for the former, variables range over subwords, while for the latter, variables range over positions. This is similar to comparing concatenation logic to spanners (see Section 3.4), and we address it analogously through the notion of a formula from one logic *realizing* a formula from the other. We start with the direction from relational to concatenation logic.

► **Definition 3.12.** Let $\varphi \in \text{FO}[\text{Eq}, \mathcal{K}]$. For every assignment α for φ on some structure \mathcal{A}'_w , we define its corresponding substitution σ by $\sigma(x) := w_{[1, \alpha(x)]}$ for all $x \in \text{free}(\varphi)$.

A formula $\psi \in \text{FC}_{(1)}[\mathcal{K}]$ realizes φ if $\text{free}(\psi) = \text{free}(\varphi)$ and for all $w \in \Sigma^*$, we have $\sigma \in \llbracket \psi \rrbracket(w)$ if and only if σ is the corresponding substitution of some $\alpha \in \llbracket \varphi \rrbracket(w)$.

Thus, ψ represents node $i \in \{1, \dots, |w| + 1\}$ through the prefix of w that has length $i - 1$.

► **Lemma 3.13.** Given $\varphi \in \text{FO}[\text{Eq}, \mathcal{K}]$, we can compute in polynomial time $\psi \in \text{FC}[\mathcal{K}]$ that realizes φ . This preserves the properties existential and existential-positive.

The direction from FC to FO is less straightforward. We have to increase the number of variables, due to a counting argument: On a word w , the number of possible assignments can be quadratic in $|w|$ for an FC-variable; but there are only $|w| + 1$ possible choices per FO-variable. Accordingly, we shall represent each variable x with two variables x^o and x^c ; and the goal is to express a substitution σ in an assignment α by $\sigma(x) = w_{[\alpha(x^o), \alpha(x^c)]}$.

► **Definition 3.14.** Let $\varphi(\mathfrak{s}) \in \text{FC}[\mathcal{K}]$ and let σ be a substitution for φ . Let $\psi \in \text{FO}[\text{Eq}, \mathcal{K}]$ with $\text{free}(\psi) := \{x^o, x^c \mid x \in \text{free}(\varphi)\}$. An assignment α for ψ on $\mathcal{A}'_{\sigma(\mathfrak{s})}$ expresses σ if $\sigma(x) = \sigma(\mathfrak{s})_{[\alpha(x^o), \alpha(x^c)]}$ for all $x \in \text{free}(\varphi)$. We say ψ realizes φ if, for all $w \in \Sigma^*$, we have

1. if $\alpha \in \llbracket \psi \rrbracket(w)$, then α expresses some $\sigma \in \llbracket \varphi \rrbracket(w)$, and
2. if $\sigma \in \llbracket \varphi \rrbracket(w)$, then $\alpha \in \llbracket \psi \rrbracket(w)$ holds for all α that express σ .

► **Lemma 3.15.** Given $\varphi \in \text{FC}[\mathcal{K}]$, we can compute $\psi \in \text{FO}[\text{Eq}, \mathcal{K}]$ in polynomial time that realizes φ . This preserves the properties existential and existential-positive.

We extend \equiv and \equiv_{poly} to the comparison of FC and $\text{FO}[\text{Eq}]$, using it to express that formulas in one logic are realized by formulas in the other. Combining Lemma 3.13 and 3.15, we get:

► **Theorem 3.16.** $FC_{(1)}[\mathcal{K}] \equiv_{\text{poly}} FO[\text{Eq}, \mathcal{K}]$ and $EFC_{(1)}[\mathcal{K}] \equiv_{\text{poly}} EFO[\text{Eq}, \mathcal{K}]$.

Hence, not only have $FC_{(1)}[\mathcal{K}]$ and $FO[\text{Eq}, \mathcal{K}]$ the same expressive power, the conversions between them can be performed in polynomial time. This only refers to formulas that operate on single words; to extend this to tuples, one would need to generalize $FO[<]$ accordingly.

3.4 Connecting FC[REG] to spanners

We next introduce a fragment of \mathcal{C} that uses guards to simulate the semantics of FC syntactically. This shall allow us to connect $FC[\text{REG}]$ to SpLog^\neg and document spanners.

► **Definition 3.17.** Let $\varphi \in \mathcal{C}$. Then $x \in \Xi$ is a main variable if it appears on the left side of a word equation in φ . Let $\text{main}(\varphi)$ be the set of main variables of φ , and define the set of its auxiliary variables $\text{aux}(\varphi) := \text{free}(\varphi) - \text{main}(\varphi)$. The guarded fragment GC is the subset of \mathcal{C} where all subformulas satisfy the following conditions:

1. Every word equation has exactly one variable on its left side,
2. every disjunction $\psi_1 \vee \psi_2$ has $\text{main}(\psi_1) = \text{main}(\psi_2)$ and $\text{aux}(\psi_1) = \text{aux}(\psi_2)$,
3. quantifiers do not bind main variables,
4. constraints are guarded $\psi \wedge \kappa(\vec{x})$ with $\vec{x} \subseteq \text{free}(\psi)$, and
5. negations are guarded $\psi_1 \wedge \neg\psi_2$ with $\text{main}(\psi_1) = \text{main}(\psi_2)$ and $\text{aux}(\psi_1) = \text{aux}(\psi_2)$.

► **Example 3.18.** The formula $\varphi := \exists y, z: (x \doteq yy \wedge y \doteq zz)$ is not a GC-formula. As y appears on the left side of a word equation, it is a main variable. Hence, Definition 3.17 does not allow binding y with an existential quantifier.

GC[REG] can be seen as a generalization of SpLog (see [30]): For $w \in \Xi$, $\text{SpLog}(w)$ is the set of all $\varphi \in \text{EPGC}[\text{REG}]$ with $\text{main}(\varphi) = w$. The same holds for SpLog^\neg and GC[REG]. Although GC is defined as a fragment of \mathcal{C} , we can treat its formulas as FC-formulas:

► **Lemma 3.19.** If $\sigma \models \varphi$ for $\varphi \in \text{GC}$, then for every $x \in \text{aux}(\varphi)$, there is some $y \in \text{main}(\varphi)$ with $\sigma(y) \supseteq \sigma(x)$.

Hence, for $\varphi \in \text{GC}$, we have $\sigma \models \varphi$ iff. $\sigma \models_{\text{main}(\varphi)} \varphi$. Thus, we can treat φ as an FC-formula by interpreting main as structure variables and auxiliary as free variables. We extend the notion of equivalence accordingly (within GC, and between GC and FC).

► **Lemma 3.20.** Given $\varphi \in \text{FC}$, we can compute in polynomial time an equivalent $\psi \in \text{GC}$. This preserves the properties existential and existential-positive.

Hence, if one prefers the infinite model semantics of the theory of concatenation over the finite model semantics of FC, one could also view GC as the “real logic”, and treat FC as syntactic sugar (this is the essentially how [30] uses SpLog). But the authors think that the safe semantics of FC are simpler to work with and lead to cleaner formulas, in particular for non-unary relations (compare Lemma 3.7 to Lemma 5.1 in [30]).

As spanners reason over positions in a word, and FC and SpLog over words, we cannot directly consider them equivalent. Instead, we use the notion of one *realizing* the other (see [30, 31] or Definition A.4), analogously to Section 3.3. Likewise, we extend \equiv (and \equiv_{poly}) to the comparison of formulas and spanner representations. As shown in [30], we have $\text{SpLog} \equiv_{\text{poly}} \text{RGX}^{\text{core}}$ and $\text{SpLog}^\neg \equiv_{\text{poly}} \text{RGX}^{\text{gcore}}$. Together with Lemma 3.20, we conclude:

► **Theorem 3.21.** $FC_{(1)}[\text{REG}] \equiv_{\text{poly}} \text{RGX}^{\text{gcore}}$ and $\text{EPFC}_{(1)}[\text{REG}] \equiv_{\text{poly}} \text{RGX}^{\text{core}}$.

Hence, like SpLog and SpLog^\neg , $\text{FC}[\text{REG}]$ and its fragments can be used to replace document spanners or to define relations for them. There is a class of spanners that corresponds to $\text{FC}[\]$: We call a regex formula *simple* if the operator $*$ is only applied to Σ or terminal words. Thus, $\Sigma^*x\{\text{foo}\}\Sigma^*$ and $(\text{ab})^*c$ are simple, but $(\text{a} \cup \text{b})^*c$ is not. Let sRGX be the set of all simple regex formulas in RGX , and define $\text{sRGX}^{\text{core}}$ and $\text{sRGX}^{\text{gcore}}$ accordingly.

► **Theorem 3.22.** $\text{FC}_{(1)}[\] \equiv_{\text{poly}} \text{sRGX}^{\text{gcore}}$ and $\text{EFC}_{(1)}[\] \equiv_{\text{poly}} \text{sRGX}^{\text{core}}$.

4 Complexity and expressive power

4.1 Model checking FC

The *model checking problem* for FC is, given a Boolean formula φ and $\vec{w} \in (\Sigma^*)^k$ with $k = |\text{struc}(\varphi)|$, decide whether $\llbracket \varphi \rrbracket \langle \vec{w} \rangle = \{()\}$. The *recognition problem* is, given a formula φ and a pattern substitution σ , decide whether $\sigma \models \varphi$. As common in literature (see e.g. [1, 46]), we distinguish *data complexity*, where the formula is fixed and only \vec{w} or σ are considered input, and *combined complexity*, where the formula is also part of the input. Apart from the PSPACE lower bound, the following results were shown in [30, 31, 32] for SpLog and SpLog^\neg (and, hence, $\text{EPFC}[\text{REG}]$ and $\text{FC}[\text{REG}]$). We provide a matching lower bound and translate all these results into our more general framework:

► **Theorem 4.1.** *The data complexity of the $\text{FC}[\mathcal{K}]$ -recognition problem is in L . The combined complexity is PSPACE-complete for $\text{FC}[\mathcal{K}]$ and NP-complete for $\text{EFC}[\mathcal{K}]$, even if restricted to Boolean $\text{FC}_{(1)}[\]$ - and $\text{EPFC}_{(1)}[\]$ -formulas over single letter words.*

The combined complexity is the same as for relational first-order logic (see e.g. [46]), and the proofs are equally straightforward. We define the *width* $\text{wd}(\varphi)$ of a formula φ as the maximum number of free variables in any of its subformulas. Another result that translates from FO is that bounding the width reduces the complexity of model checking:

► **Theorem 4.2.** *Given $\varphi \in \text{FC}[\]$ and a tuple \vec{w} over Σ^* , we can decide whether $\llbracket \varphi \rrbracket \langle \vec{w} \rangle = \emptyset$ in time $O(k|\varphi|n^{2k})$, where $k := \text{wd}(\varphi)$ and $n := \|\vec{w}\|$.*

As pointed out in the proof, this is only a rough upper bound; taking properties of variables into account lowers the exponent. Moreover, the same approach also works for $\text{FC}[\mathcal{K}]$; and the total time depends on the cost of deciding the constraints.

4.1.1 Patterns and variable-bounded formulas

In principle, we can apply various structure parameters for first-order formulas (see e.g. Adler and Weyer [2]) to FC. But this assumes that we can treat word equations as atomic formulas (as FO treats relation symbols as atomic formulas), which is not the case. First, note that [32] applied the notion of acyclic conjunctive queries (see [1]) to spanner representations, treating the atomic regex formulas like relational atoms, with the conclusion that in this setting, acyclicity does not guarantee tractable evaluation. But if one converts a non-trivial regex formula to an $\text{FC}[\text{REG}]$ -formula, one usually obtains a non-atomic formula with multiple word equations. Hence, from an FC point of view, one can argue that regex formulas should not be considered atoms when defining notions like acyclicity for spanners.

In fact, when defining acyclicity, even word equations should not be considered atomic: The patterns that form the right sides of word equations can be seen as shorthand for terms with binary concatenation as a function. This raises the question which patterns can be

decomposed into simpler formulas; for example, into formulas of bounded width. We develop a sufficient criterion; building on a result from combinatorics on words and formal languages.

In addition to being the fundamental building blocks of word equations, patterns also received significant attention as language generators. The language $\mathcal{L}(\alpha)$ of a pattern α is the set of all $\sigma(\alpha)$, where σ is a pattern substitution. This definition is from Angluin [4], who also showed that the membership problem for pattern languages is NP-complete. Moreover, given two patterns α and β , it is undecidable whether $\mathcal{L}(\alpha) \subseteq \mathcal{L}(\beta)$ holds (see [10]). Hence, many canonical problems for FC are already difficult for pattern languages. Reidenbach and Schmid [60] started a series of articles on classes of pattern languages with a polynomial time membership problem (surveyed in Manea and Schmid [52]), most of which are special cases of the following definition (see [16] or Appendix B.3 for the definition of treewidth):

► **Definition 4.3.** *The standard graph of a pattern $\alpha = \alpha_1 \cdots \alpha_n$ with $n \geq 1$ and $\alpha_i \in (\Sigma \cup \Xi)$ is $\mathcal{G}_\alpha := (V_\alpha, E_\alpha)$ with $V_\alpha := \{1, \dots, n\}$ and $E_\alpha := E_\alpha^< \cup E_\alpha^{\text{eq}}$, where $E^<$ is the set of all $\{i, i+1\}$ with $1 \leq i < n$, and E_α^{eq} is the set of all $\{i, j\}$ such that α_i is some $x \in \Xi$, and α_j is the next occurrence of x in α . Then $\text{tw}(\alpha)$, the treewidth of α , is the treewidth of \mathcal{G}_α .*

► **Theorem 4.4.** *For every pattern α , one can construct in polynomial time $\varphi \in \text{EPFC}[\]$ with $\mathcal{L}(\alpha) = \mathcal{L}(\varphi)$ and $\text{wd}(\varphi) = 2\text{tw}(\alpha) + 3$.*

This allows us to rewrite every formula $\exists \vec{x}: y \doteq \alpha$ into an equivalent EPFC[$\]$ -formula with width $2\text{tw}(\alpha) + k + 3$, where $k = |\text{free}(y \doteq \alpha) - \vec{x}|$. Recall that we treat tuples as sets. Combining Theorems 4.2 and 4.4 yields a (slightly) different proof of the polynomial time decidability of the membership problem for classes of patterns with bounded treewidth from [60]. As pointed out in [17], bounded treewidth does not cover all pattern languages with a polynomial time membership problem, like e. g. patterns α^k where $\text{tw}(\alpha)$ is bounded. But these languages can be expressed by $\exists x: (\mathfrak{s} \doteq x^k \wedge \varphi_\alpha(x))$, where $\varphi_\alpha(x)$ is a formula that selects $\mathcal{L}(\alpha)$, thus increasing the width only by 1. In addition to lowering the width, we can define FC-formulas that are exponentially shorter than the equivalent pattern:

► **Example 4.5.** Let $\alpha_k := x^{2^k}$ for $k \geq 1$. Then $\mathcal{L}(\alpha_k) = \mathcal{L}(\varphi)$ for $\varphi(\mathfrak{s}) := \exists x_1, \dots, x_k: (\mathfrak{s} \doteq x_1 x_1 \wedge \bigwedge_{i=1}^{k-1} x_i \doteq x_{i+1} x_{i+1})$. We can achieve width 2 by pulling quantifiers inwards.

There are more general definitions of the treewidth of a pattern (see [52, 60]). While patterns can be understood as concatenation terms, this translation is ambiguous. Hence, even the defining acyclic patterns is not straightforward. We leave this and related issues for future work; but we briefly consider aspects of succinctness in Section 4.4.

4.2 Iteration and recursion

Like FO, we can extend FC with operators for transitive closure and fixed points (see e. g. [46]). We denote the respective extensions of FC with *deterministic transitive closure*, *transitive closure*, *least fixed-points*, and *partial fixed point* by FC^{dte} , FC^{tc} , FC^{lfp} , and FC^{pfp} (see Appendix B.4 for definitions). We say that such a logic \mathcal{F} captures a complexity class \mathbb{C} if the languages that are definable in \mathcal{F} are exactly the languages in \mathbb{C} .

► **Theorem 4.6.** FC^{dte} , FC^{tc} , FC^{lfp} , FC^{pfp} capture L, NL, P, PSPACE, respectively.

This holds even for formulas that are existential-positive and constraint-free. Thus, FC and even EPFC[$\]$ behave under fixed-points and transitive closures like FO[$\]$. As FO with least-fixed points can be used to define Datalog (see Part D of [1]). To mirror this for FC, we define DataSpLog, a version of Datalog that is based on word equations. For a set of

constraints \mathcal{K} , a $\text{DataSpLog}[\mathcal{K}]$ -program is a tuple $P := (\vec{s}, \mathcal{R}, \Phi, \text{OUT})$, where \vec{s} are the structure variables, \mathcal{R} is a set of relation symbols $\mathcal{K} \cap \mathcal{R} = \emptyset$, each $R \in \mathcal{R}$ has an arity $\text{ar}(R)$, $\text{OUT} \in \mathcal{R}$, and Φ is a finite set of rules $R(\vec{x}) \leftarrow \varphi_1(\vec{y}_1), \dots, \varphi_m(\vec{y}_m)$ with $R \in \mathcal{R}$, $m \geq 1$, each φ_i is an atomic $\text{FC}[\mathcal{K} \cup \mathcal{R}](\vec{s})$ -formula, and each $x \in \vec{x}$ appears in some \vec{y}_i .

We define $\llbracket P \rrbracket(\vec{w})$ incrementally, initializing the relations of all $R \in \mathcal{R}$ to \emptyset . For each rule $R(\vec{x}) \leftarrow \varphi_1(\vec{y}_1), \dots, \varphi_m(\vec{y}_m)$, we enumerate all σ with $\sigma(\vec{s}) := \vec{w}$ and check if $\sigma \models \exists \vec{y}: \bigwedge_{i=1}^m \varphi_i$, where $\vec{y} := (\bigcup_{i=1}^m \vec{y}_i) - \vec{x}$. If this holds, we add $\sigma(\vec{x})$ to R . This is repeated until all relations have stabilized. Then $\llbracket P \rrbracket(\vec{w})$ is the content of the relation OUT .

► **Example 4.7.** Consider the $\text{DataSpLog}[\]$ -program $(\mathfrak{s}, \{O, E\}, \Phi, O)$, with $\text{ar}(O) = 0$, $\text{ar}(E) = 3$, and Φ contains $O() \leftarrow \mathfrak{s} \doteq xyz$, $E(x, y, z) \leftarrow x \doteq \varepsilon, y \doteq \varepsilon, z \doteq \varepsilon$ and $E(x, y, z) \leftarrow x \doteq \hat{x}\mathfrak{a}, y \doteq \hat{y}\mathfrak{b}, z \doteq \hat{z}\mathfrak{c}, E(\hat{x}, \hat{y}, \hat{z})$. This defines the language of all $\mathfrak{a}^n \mathfrak{b}^n \mathfrak{c}^n$, $n \geq 0$.

► **Theorem 4.8.** $\text{DataSpLog}[\]$ captures P .

By our assumptions on \mathcal{K} , this holds for all $\text{DataSpLog}[\mathcal{K}]$. This is unsurprising, considering Datalog on ordered structures captures P , see [1, 24, 46] and the analogous result for spanners with recursion [58]. But it allows using word equations as a basis for Datalog on words and provides potential application for future insights into acyclicity for patterns. Moreover, Datalog can be seen as a generalization of *range concatenation grammars (RCGs)*, see [9, 40], to outputs and relations. As such, DataSpLog might help to develop deeper connections between Datalog on words and RCGs.

4.3 Inexpressibility for FC

There are currently only few inexpressibility methods for $\text{FC}[\]$ and $\text{FC}[\text{REG}]$. Over unary alphabets, $\text{FC}[\text{REG}]$ is as expressible as Presburger arithmetic, as observed in [31, 58] for $\text{EFC}[\text{REG}]$ and $\text{FC}[\text{REG}]$, respectively. This also applies to $\text{FC}[\]$. Larger alphabets are more complicated: For the existential fragment, Lemma 3.20 allows us to treat $\text{EFC}[\]$ as $\text{EC}[\]$ and to use some of the inexpressibility results of Karhumäki, Mignosi, and Plandowski [41]. This translates to $\text{EFC}[\text{REG}]$ under certain conditions (see [30]). For $\text{EFC}[\text{REG}]$, there is also the core-simplification-lemma of Fagin et al. [25].

For FO , a standard inexpressibility tool are Ehrenfeucht–Fraïssé games (e.g. [46]). But equalities make using these for FC or $\text{FO}[\text{Eq}]$ far from straightforward. For sufficiently restricted languages, the Feferman-Vaught theorem (see [51]) can be used:

► **Lemma 4.9.** *There is no $\text{FC}[\]$ -formula that defines the language $\{\mathfrak{a}^n \mathfrak{b}^n \mid n \geq 1\}$.*

Moreover, we can show that regular constraints offer no help for defining this language. This allows us to generalize this inexpressibility to $\text{FC}[\text{REG}]$ and, hence, to $\text{RGX}^{\text{gcore}}$. To the authors' knowledge, this is the first inexpressible result for $\text{RGX}^{\text{gcore}}$ on non-unary alphabets:

► **Theorem 4.10.** *If $|\Sigma| \geq 2$, then the equal length relation is not $\text{FC}[\text{REG}]$ -selectable.*

In this specific case, the limited structure of the language $\{\mathfrak{a}^n \mathfrak{b}^n \mid n \geq 1\}$ allows the use of the Feferman-Vaught theorem. But a more general inexpressibility method for $\text{FC}[\]$ would probably need to combine this with techniques like those in [41].

4.4 Static analysis and satisfiability

Section 4.1.1 discusses how FC (and even EPFC) can express patterns succinctly. Two big open questions are how big this advantage is, and whether we can compute such minimizations. The undecidability results for SpLog (and, hence, for $\text{EPFC}[\text{REG}]$) from Freydenberger and Holldack [31] can be adapted to $\text{EPFC}[\]$. Let FC^k or denote the set of formulas with width $\leq k$.

► **Theorem 4.11.** *Even for Boolean $\text{EPFC}_{\langle 1 \rangle}^4[\Sigma]$, we can decide neither containment, nor equivalence, nor whether the formula defines a language that is Σ^* , regular, a pattern language, or expressible in $\text{FC}_{\langle 1 \rangle}^0[\Sigma]$. Furthermore, given $\varphi \in \text{EPFC}_{\langle 1 \rangle}^4[\Sigma]$, we cannot compute an equivalent ψ such that $|\psi|$ is minimal.*

One might think that fragments like $\text{EPFC}_{\langle 1 \rangle}^1[\Sigma]$ are too limited, but these already allow expressing 1-variable word equations, for which the solutions are far from trivial (see Nowotka and Saarela [56]). Note that this leaves open the decidability of, given $\varphi \in \text{FC}[\Sigma]$ (or $\varphi \in \text{EPFC}[\Sigma]$) and $k > 0$, is there an equivalent $\psi \in \text{FC}^k[\Sigma]$. But without suitable inexpressibility methods (recall Section 4.3), we cannot even show that a language is inexpressible in $\text{FC}^k[\Sigma]$ for some $k > 0$, which complicates tackling this problem.

Via Hartmanis' [39] meta theorem, certain undecidability results provide insights into the relative succinctness of models (see [45] or e.g. [31] for details). For two logics \mathcal{F}_1 and \mathcal{F}_2 , the *tradeoff* from \mathcal{F}_1 to \mathcal{F}_2 is *non-recursive* if, for every computable $f: \mathbb{N} \rightarrow \mathbb{N}$, there exists some $\varphi \in \mathcal{F}_1$ that is expressible in \mathcal{F}_2 , but $|\psi| \geq f(|\varphi|)$ holds for every $\psi \in \mathcal{F}_2$ with $\psi \equiv \varphi$.

► **Theorem 4.12.** *There are non-recursive tradeoffs from $\text{EPFC}_{\langle 1 \rangle}^4[\Sigma]$ to regular expressions and to $\text{FC}_{\langle 1 \rangle}^0[\Sigma]$, and from $\text{FC}_{\langle 1 \rangle}^4[\Sigma]$ to $\text{EPFC}[\Sigma]$, to patterns, and to singleton sets $\{w\}$.*

Another consequence of the proof of Theorem 4.12 is that satisfiability is undecidable for $\text{FC}_{\langle 1 \rangle}^4[\Sigma]$ and trivial for $\text{FC}_{\langle 1 \rangle}^0[\Sigma]$. Leaving the cases of one and two variables open, we show:

► **Proposition 4.13.** *Satisfiability for $\text{FC}_{\langle 1 \rangle}^3[\Sigma]$ is undecidable if $|\Sigma| \geq 2$.*

5 Conclusions and further directions

By defining FC, we have introduced a logic that can be understood as a finite-model approach to the theory of concatenation. Like SpLog and its variants, FC can be used as a logic for document spanners; but in contrast to these logics, FC is less cumbersome. We have also connected FC to FO[Eq], which extends FO[<] with the equality predicate Eq. This allowed us to use the Feferman-Vaught theorem in an inexpressibility proof; but we could also adapt many classical results from FO to the FC-setting.

Many fundamental questions remain open, in particular regarding inexpressibility techniques and efficient model checking (and, later, enumeration of results). For inexpressibility, very little is known: Lemma 4.9 heavily relies on the limited structure of the language. This is the same situation as discussed in Section 6.1 of [30], which describes an inexpressibility technique for $\text{EPFC}[\text{REG}]$. While these two approaches provide us with some means of proving inexpressibility for FC or its fragments, this only touches some special cases, and much remains to be done. It seems likely that a more general method should combine approaches from finite-model theory with techniques as those in [41].

Regarding model checking and evaluation of formulas, Section 4.1.1 already describes that treating word equations as atoms might be not enough to capture all cases of formulas for which tractable model checking. One possible direction is obvious: Develop a theory of acyclic patterns, extend this to conjunctive queries over word equations, and then generalize this to formulas. This could then be used as stepping stone towards a notion of acyclic regex formulas (and other spanner representations) or be adapted to DataSpLog. A potential advantage of DataSpLog over Datalog on words is that it should be possible to extend natural criteria for tractable word equations directly into natural criteria for tractable DataSpLog.

There are many other possible directions. For example, one could easily define a second-order version of FC and adapt various results from SO. Moreover, FC could be examined from an algebra point of view, or related to rational and regular relations.

References

- 1 Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases*. Addison-Wesley, 1995. URL: <http://webdam.inria.fr/Alice/>.
- 2 Isolde Adler and Mark Weyer. Tree-width for first order formulae. *Logical Methods in Computer Science*, 8(1), 2012.
- 3 Antoine Amarilli, Pierre Bourhis, Stefan Mengel, and Matthias Niewerth. Constant-delay enumeration for nondeterministic document spanners. In *Proc. ICDT 2019*, pages 22:1–22:19, 2019.
- 4 Dana Angluin. Finding patterns common to a set of strings. *Journal of Computer and System Sciences*, 21:46–62, 1980.
- 5 Pablo Barceló and Pablo Muñoz. Graph logics with rational relations: The role of word combinatorics. *ACM Transactions on Computational Logic*, 18(2):10:1–10:41, 2017.
- 6 Michael Benedikt, Leonid Libkin, Thomas Schwentick, and Luc Segoufin. Definable relations and first-order query languages over strings. *Journal of the ACM*, 50(5):694–751, 2003.
- 7 Christoph Berkholz and Hubie Chen. Compiling existential positive queries to bounded-variable fragments. In *Proc. PODS 2019*, pages 353–364, 2019.
- 8 Anthony J. Bonner and Giansalvatore Mecca. Sequences, datalog, and transducers. *Journal of Computer and System Sciences*, 57(3):234–259, 1998.
- 9 Pierre Boullier. Range concatenation grammars. In *New developments in parsing technology*, pages 269–289. Springer, 2004.
- 10 Joachim Bremer and Dominik D. Freydenberger. Inclusion problems for patterns with a bounded number of variables. *Information and Computation*, 220:15–43, 2012.
- 11 J. Richard Büchi and Steven Senger. Coding in the existential theory of concatenation. *Archiv für mathematische Logik und Grundlagenforschung*, 26(1):101–106, 1987.
- 12 Taolue Chen, Yan Chen, Matthew Hague, Anthony W. Lin, and Zhilin Wu. What is decidable about string constraints with the replaceall function. *PACMPL*, 2(POPL):3:1–3:29, 2018.
- 13 Christian Choffrut and Juhani Karhumäki. Combinatorics of words. In Grzegorz Rozenberg and Arto Salomaa, editors, *Handbook of Formal Languages*, volume 1: Word, Language, Grammar, chapter 6, pages 329–438. Springer, 1997.
- 14 Laura Ciobanu, Volker Diekert, and Murray Elder. Solution sets for equations over free groups are EDTOL languages. *IJAC*, 26(5):843–886, 2016.
- 15 John H. Conway. *FRACTRAN: A simple universal computing language for arithmetic*, 1987.
- 16 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshantov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- 17 Joel D. Day, Pamela Fleischmann, Florin Manea, Dirk Nowotka, and Markus L. Schmid. On matching generalised repetitive patterns. In *Proc. DLT 2018*, pages 269–281, 2018.
- 18 Joel D. Day, Vijay Ganesh, Paul He, Florin Manea, and Dirk Nowotka. The satisfiability of word equations: Decidable and undecidable theories. In *Proc. RP 2018*, pages 15–29, 2018.
- 19 Joel D. Day, Florin Manea, and Dirk Nowotka. The hardness of solving simple word equations. In *Proc. MFCS 2017*, pages 18:1–18:14, 2017.
- 20 Volker Diekert. More than 1700 years of word equations. In *Proc. CAI 2015*, 2015.
- 21 Johannes Doleschal, Benny Kimelfeld, Wim Martens, Yoav Nahshon, and Frank Neven. Split-correctness in information extraction. In *Proc. PODS 2019*, pages 149–163, 2019.
- 22 Johannes Doleschal, Benny Kimelfeld, Wim Martens, and Liat Peterfreund. Weight annotation in information extraction. In *Proc. ICDT 2020*, 2020. Accepted.
- 23 V. G. Durnev. Undecidability of the positive $\forall\exists^3$ -theory of a free semigroup. *Siberian Mathematical Journal*, 36(5):917–929, 1995.
- 24 Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite Model Theory*. Springer, 2nd edition, 1999.
- 25 Ronald Fagin, Benny Kimelfeld, Frederick Reiss, and Stijn Vansummeren. Document spanners: A formal approach to information extraction. *Journal of the ACM*, 62(2):12:1–12:51, 2015.

- 26 Ronald Fagin, Benny Kimelfeld, Frederick Reiss, and Stijn Vansummeren. Declarative cleaning of inconsistencies in information extraction. *ACM Transactions on Database Systems*, 41(1):6:1–6:44, 2016.
- 27 Fernando Florenzano, Cristian Riveros, Martín Ugarte, Stijn Vansummeren, and Domagoj Vrgoc. Constant delay algorithms for regular document spanners. In *Proc. PODS 2018*, pages 165–177, 2018.
- 28 Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006.
- 29 Dominik D. Freydenberger. Extended regular expressions: Succinctness and decidability. *Theory of Computing Systems*, 53(2):159–193, 2013.
- 30 Dominik D. Freydenberger. A logic for document spanners. *Theory of Computing Systems*, 63(7):1679–1754, 2019.
- 31 Dominik D. Freydenberger and Mario Holldack. Document spanners: From expressive power to decision problems. *Theory of Computing Systems*, 62:854–898, 2018.
- 32 Dominik D. Freydenberger, Benny Kimelfeld, and Liat Peterfreund. Joining extractions of regular expressions. In *Proc. PODS 2018*, pages 137–149, 2018.
- 33 Dominik D. Freydenberger and Markus L. Schmid. Deterministic regular expressions with back-references. *Journal of Computer and System Sciences*, 105:1 – 39, 2019.
- 34 Dominik D. Freydenberger and Sam M. Thompson. Dynamic complexity of document spanners. In *Proc. ICDT 2020*, 2020. Accepted.
- 35 Michael R. Garey and David S. Johnson. *Computers and Intractability*. W. H. Freeman and Company, 1979.
- 36 Seymour Ginsburg and Xiaoyang Sean Wang. Regular sequence operations and their use in database queries. *Journal of Computer and System Sciences*, 56(1):1–26, 1998.
- 37 Gösta Grahne, Matti Nykänen, and Esko Ukkonen. Reasoning about strings in databases. *Journal of Computer and System Sciences*, 59(1):116–162, 1999.
- 38 Simon Halfon, Philippe Schnoebelen, and Georg Zetsche. Decidability, complexity, and expressiveness of first-order logic over the subword ordering. In *Proc. LICS 2017*, pages 1–12, 2017.
- 39 Juris Hartmanis. On Gödel speed-up and succinctness of language representations. *Theoretical Computer Science*, 26:335–342, 1983.
- 40 Laura Kallmeyer. *Parsing Beyond Context-Free Grammars*. Cognitive Technologies. Springer, 2010.
- 41 Juhani Karhumäki, Filippo Mignosi, and Wojciech Plandowski. The expressibility of languages and relations by word equations. *Journal of the ACM*, 47(3):483–505, 2000.
- 42 Phokion G. Kolaitis and Moshe Y. Vardi. Conjunctive-query containment and constraint satisfaction. *Journal of Computer and System Sciences*, 61(2):302–332, 2000.
- 43 Dexter C. Kozen. *Theory of Computation*. Springer, 1st edition, 2010.
- 44 Stuart A. Kurtz and Janos Simon. The undecidability of the generalized Collatz problem. In *Proc. TAMC 2007*, pages 542–553, 2007.
- 45 Martin Kutrib. The phenomenon of non-recursive trade-offs. *International Journal of Foundations of Computer Science*, 16(5):957–973, 2005.
- 46 Leonid Libkin. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. Springer, 2004.
- 47 Anthony Widjaja Lin and Pablo Barceló. String solving with word equations and transducers: towards a logic for analysing mutation XSS. In *Proc. POPL 2016*, pages 123–136, 2016.
- 48 Katja Losemann. *Foundations of Regular Languages for Processing RDF and XML*. PhD thesis, University of Bayreuth, 2015. URL: <https://epub.uni-bayreuth.de/2536/>.
- 49 M. Lothaire. *Combinatorics on Words*. Addison-Wesley, Reading, MA, 1983.
- 50 G. S. Makanin. The problem of solvability of equations in a free semigroup. *Matematicheskii Sbornik*, 103:147–236, 1977.

- 51 Johann A. Makowsky. Algorithmic uses of the Feferman-Vaught theorem. *Annals of Pure and Applied Logic*, 126(1-3):159–213, 2004.
- 52 Florin Manea and Markus L. Schmid. Matching patterns with variables. In *Proc. WORDS 2019*, pages 1–27, 2019.
- 53 Francisco Maturana, Cristian Riveros, and Domagoj Vrgoc. Document spanners for extracting incomplete information: Expressiveness and complexity. In *Proc. PODS 2018*, pages 125–136, 2018.
- 54 Andrea Morciano, Martin Ugarte, and Stijn Vansummeren. Automata-based evaluation of AQL queries. Technical report, Université Libre de Bruxelles, 2016.
- 55 Yoav Nahshon, Liat Peterfreund, and Stijn Vansummeren. Incorporating information extraction in the relational database model. In *Proc. WebDB 2016*, page 6, 2016.
- 56 Dirk Nowotka and Aleksa Saarela. An optimal bound on the solution sets of one-variable word equations and its consequences. In *Proc. ICALP 2018*, pages 136:1–136:13, 2018.
- 57 Liat Peterfreund, Dominik D. Freydenberger, Benny Kimelfeld, and Markus Kröll. Complexity bounds for relational algebra over document spanners. In *Proc. PODS 2019*, pages 320–334, 2019.
- 58 Liat Peterfreund, Balder ten Cate, Ronald Fagin, and Benny Kimelfeld. Recursive programs for document spanners. In *Proc. ICDT 2019*, 2019.
- 59 W. V. Quine. Concatenation as a basis for arithmetic. *The Journal of Symbolic Logic*, 11(4):105–114, 1946.
- 60 Daniel Reidenbach and Markus L. Schmid. Patterns with bounded treewidth. *Information and Computation*, 239:87–99, 2014.
- 61 Markus L. Schmid. Characterising REGEX languages by regular languages equipped with factor-referencing. *Information and Computation*, 249:1–17, 2016.
- 62 Howard Straubing. *Finite Automata, Formal Logic, and Circuit Complexity*. Birkhäuser, 1994.

A Appendix for Section 3

A.1 Proof of Lemma 3.7

A.1.1 Preparations for the proof of Lemma 3.7

For convenience, we define the set of *public variables* of φ as $\text{pub}(\varphi) := \text{struc}(\varphi) \cup \text{free}(\varphi)$. We begin by introducing a definition that shall allow us to compare formulas with the same public variables but different number of structure variables.

► **Definition A.1.** *We say that $\varphi, \psi \in \text{FC}$ are weakly equivalent if $\text{pub}(\varphi) = \text{pub}(\psi)$ and we have that $\sigma \models_{\text{struc}(\varphi)} \varphi$ holds if and only if $\sigma \models_{\text{struc}(\psi)} \psi$.*

The notion of weak equivalence allows us to make two observations on the interplay of free and structure variables. The first is the insight that arbitrarily many free variables can be “promoted” to structure variable status:

► **Lemma A.2.** *Let $\mathcal{F} \in \{\text{EPFC}[\mathcal{K}], \text{EFC}[\mathcal{K}], \text{FC}[\mathcal{K}]\}$. For every $\varphi \in \mathcal{F}$ and every $X \subseteq \text{free}(\varphi)$, there is a weakly equivalent $\psi \in \mathcal{F}$ with $\text{struc}(\psi) = \text{struc}(\varphi) \cup X$ and $\text{free}(\psi) = \text{free}(\varphi) - X$.*

Proof. Choose a fragment \mathcal{F} ; let $\varphi \in \mathcal{F}$ and $X \subseteq \text{free}(\varphi)$. We cannot simply declare the variables from X to be structure variables, as these might result in a formula that accepts substitutions that were not accepted in φ (namely, if the new structure variable has an image that is not subword of the image of some “old” structure variable). To this end, we define

$$\psi := \varphi \wedge \bigwedge_{x \in X} \bigvee_{\mathfrak{s} \in \text{struc}(\varphi)} \exists p, s: (\mathfrak{s} \doteq p x s)$$

with $\text{struc}(\psi) := \text{struc}(\varphi) \cup X$. Then $\sigma \models_{\text{struc}(\varphi)} \varphi$ if and only if $\sigma \models_{\text{struc}(\psi)} \psi$. ◀

The second insight concerns the other direction: A structure variable that always refers to a subword of another structure variable can be “demoted” to free variable status. This is an essential step in the proof of Lemma 3.7.

► **Lemma A.3.** *Let $\varphi \in \text{FC}$. If, for some structure variables $\vec{\mathfrak{s}}$, there is an $\mathfrak{s} \in \vec{\mathfrak{s}}$ such that for every σ with $\sigma \models_{\vec{\mathfrak{s}}} \varphi$, there exists some $\hat{\mathfrak{s}} \in \text{struc}(\varphi) - \{\mathfrak{s}\}$ with $\sigma(\mathfrak{s}) \sqsubseteq \sigma(\hat{\mathfrak{s}})$, then we have $\sigma \models_{\vec{\mathfrak{s}}} \varphi$ if and only if $\sigma \models_{\vec{\mathfrak{s}} - \{\mathfrak{s}\}} \varphi$.*

Proof. Let $\varphi \in \text{FC}$ and $\mathfrak{s} \in \text{struc}(\varphi)$ such that for every σ with $\sigma \models \varphi$, there is some $\hat{\mathfrak{s}} \in \text{struc}(\varphi) - \{\mathfrak{s}\}$ with $\sigma(\mathfrak{s}) \sqsubseteq \sigma(\hat{\mathfrak{s}})$. Our goal is to show that $\sigma \models_{\text{struc}(\varphi)} \varphi$ if and only if $\sigma \models_{\text{struc}(\varphi) - \{\mathfrak{s}\}} \varphi$.

For the *if*-direction, assume $\sigma \models_{\text{struc}(\varphi) - \{\mathfrak{s}\}} \varphi$. Then $\sigma(\mathfrak{s}) \in \bigcup_{x \in \vec{\mathfrak{s}} - \{\mathfrak{s}\}} \{u \mid u \sqsubseteq \sigma(x)\}$ holds by definition. Thus, we do not gain additional possible choices for the quantifiers if we allow them to also range over the subwords of $\sigma(\mathfrak{s})$. Hence, $\sigma \models_{\text{struc}(\varphi)} \varphi$ must hold.

For the *only-if*-direction, assume that $\sigma \models_{\text{struc}(\varphi)} \varphi$. By our assumption, this means that $\sigma(\mathfrak{s}) \sqsubseteq \sigma(\hat{\mathfrak{s}})$ for some other structure variable $\hat{\mathfrak{s}}$. Hence,

$$\bigcup_{x \in \vec{\mathfrak{s}}} \{u \mid u \sqsubseteq \sigma(x)\} = \bigcup_{x \in \vec{\mathfrak{s}} - \{\mathfrak{s}\}} \{u \mid u \sqsubseteq \sigma(x)\}$$

holds, which means that removing \mathfrak{s} from the structure variables does not change the available choices for the quantifiers. Thus, we have $\sigma \models_{\text{struc}(\varphi) - \{\mathfrak{s}\}} \varphi$. ◀

In other words, Lemma A.3 states that if a structure variable \mathfrak{s} of some FC-formula φ is always subword of other structure variables, then we can interpret φ as a weakly equivalent formula $\hat{\varphi}$ with $\text{struc}(\hat{\varphi}) = \text{struc}(\varphi) - \{\mathfrak{s}\}$ and $\text{free}(\hat{\varphi}) = \text{free}(\varphi) \cup \{\mathfrak{s}\}$.

A.1.2 Main part of the proof of Lemma 3.7

Proof. Let $\mathcal{F} \in \{\text{EPFC}[\mathcal{K}], \text{EFC}[\mathcal{K}], \text{FC}[\mathcal{K}]\}$ and choose a relation $R \subseteq (\Sigma^*)^k$ with $k \geq 1$. We want to show that R is \mathcal{F} -selectable if and only if it is definable in $\mathcal{F} \cap \text{FC}_{\langle k \rangle}$.

Only-if-direction: Assume that R is \mathcal{F} -selectable. Let $\vec{\mathfrak{s}} := (\mathfrak{s}_1 \dots, \mathfrak{s}_k)$ and $\vec{x} := (x_1, \dots, x_k)$ consist of pairwise distinct variables, and define $\varphi(\vec{\mathfrak{s}})(\vec{x}) := \bigwedge_{i=1}^k (\mathfrak{s}_i \doteq x_i)$.

Then $\varphi \in \mathcal{F} \cap \text{FC}_{\langle k \rangle}$, and we have that $\sigma \models \varphi$ if and only if $\sigma(\vec{\mathfrak{s}}) = \sigma(\vec{x})$. As we assume R is \mathcal{F} -selectable, there exists $\varphi_{\vec{x}}^R \in \mathcal{F}$ such that $\sigma \models \varphi_{\vec{x}}^R$ if and only if $\sigma \models \varphi$ and $\sigma(\vec{x}) \in R$. Hence, $\sigma \models \varphi_{\vec{x}}^R$ if and only if $\sigma(\vec{\mathfrak{s}}) = \sigma(\vec{x}) \in R$. Thus, $\mathcal{R}(\psi) = R$ for $\psi(\vec{\mathfrak{s}})() := \exists \vec{x}: \varphi_{\vec{x}}^R(\vec{\mathfrak{s}})(\vec{x})$. This proves that R is definable in $\mathcal{F} \cap \text{FC}_{\langle k \rangle}$.

If-direction: Assume that R is definable in $\mathcal{F} \cap \text{FC}_{\langle k \rangle}$. Then there exists some $\psi() \in \mathcal{F}(\vec{\mathfrak{s}})$ with $\vec{\mathfrak{s}} := (\mathfrak{s}_1 \dots, \mathfrak{s}_k)$ and $\sigma \models \psi$ if and only if $\sigma(\vec{\mathfrak{s}}) \in R$. Consider any $\varphi \in \mathcal{F}$ such that $\text{pub}(\varphi)$ and $\text{pub}(\psi)$ are disjoint (renaming variables if necessary). Choose any k -tuple $\vec{x} = (x_1, \dots, x_k)$ over $\text{free}(\varphi)$. Note that these x_i are not required to be distinct, and define $\psi_{\vec{x}}^R := \varphi \wedge \bigwedge_{i=1}^k (x_i \doteq \mathfrak{s}_i) \wedge \psi$ with $\text{struc}(\psi_{\vec{x}}^R) = \text{struc}(\varphi) \cup \text{struc}(\psi)$. Observe that $\sigma \models \psi_{\vec{x}}^R$ holds if and only if $\sigma \models \varphi$, $\sigma(\vec{x}) = \sigma(\vec{\mathfrak{s}})$, and $\sigma(\vec{\mathfrak{s}}) \in R$.

Now observe that for every $\sigma \models \varphi$ and every $1 \leq i \leq k$, there exists some $y \in \text{struc}(\varphi)$ such that $\sigma(x_i) \sqsubseteq \sigma(y)$ and, hence, $\sigma(\mathfrak{s}_i) \sqsubseteq \sigma(y)$. This allows us to apply Lemma A.3 for each \mathfrak{s}_i , and interpret $\psi_{\vec{x}}^R$ as a weakly equivalent formula $\psi'_{\vec{x}}^R$ with $\text{struc}(\psi'_{\vec{x}}^R) = \text{struc}(\varphi)$. We define $\varphi_{\vec{x}}^R := \exists \vec{\mathfrak{s}}: \psi'_{\vec{x}}^R$, and observe that $\sigma \models \varphi_{\vec{x}}^R$ holds if and only if $\sigma \models \varphi$ and $\sigma(\vec{x}) \in R$. As $\varphi_{\vec{x}}^R \in \mathcal{F}$, we conclude that R is \mathcal{F} -selectable. \blacktriangleleft

A direct consequence of the proof of Lemma 3.7 is that for every φ and every definable relation R , we have $|\varphi_{\vec{x}}^R| \in O(|\varphi| + |\psi_R|)$, where ψ_R is a formula that defines R .

A.2 Proof of Lemma 3.9

Proof. In existential formulas, negations can only be applied to word equations and regular constraints.

Every negation of a word equations is replaced with an existential-positive formula, using the construction from Example 3.8. This does not introduce regular constraints; so the resulting formula is constraint-free if and only if the original formula is. The length of the resulting formula is $|\Sigma|^2$ times the length of the original formula and, hence, polynomial. In fact, as we assume $|\Sigma|$ to be fixed, it even is linear. Hence, $\text{EFC} \equiv_{\text{poly}} \text{EPFC}$.

For negations of regular constraints $\neg(x \dot{\in} \alpha)$, we use the fact that the class of regular languages is closed under complement. Thus, there is a regular expression $\hat{\alpha}$ with $\mathcal{L}(\hat{\alpha}) = \Sigma^* - \mathcal{L}(\alpha)$. We replace the negated constraint with $x \dot{\in} \hat{\alpha}$.

Note that the length of $\hat{\alpha}$ can be double-exponential in the length of α , which makes this part of the construction inefficient. Thus, the construction gives us $\text{EFC}[\text{REG}] \equiv \text{EPFC}[\text{REG}]$, but not $\text{EFC}[\text{REG}] \equiv_{\text{poly}} \text{EPFC}[\text{REG}]$. \blacktriangleleft

A.3 Proof of Lemma 3.13

For FC and FO, the *width* $\text{wd}(\varphi)$ of a formula φ is the maximum number of free variables in any of its subformulas. We prove the following, stronger result:

Given $\varphi \in \text{FO}[\text{Eq}, \mathcal{K}]$ with $k := \text{wd}(\varphi)$, we can compute $\psi \in \text{FC}[\mathcal{K}]$ in time $O(k|\varphi|)$ that realizes φ . This preserves the properties existential and existential-positive. Furthermore, we have $\text{wd}(\psi) = k + 2 + \max\{\text{ar}(\kappa) \mid \kappa \in \mathcal{K}\}$.

Proof. We use $(x \sqsubseteq_p \mathfrak{s})$ as shorthand for the formula $\exists z: (\mathfrak{s} \doteq xz)$. This formula is frequently used as a guard to ensure that our construction has the “prefix invariant”, by which we mean that $\sigma \models \psi$ implies $\sigma(x) \sqsubseteq_p \sigma(\mathfrak{s})$ for all constructed ψ and all $x \in \text{free}(\psi)$. Usually, we do not point this out. The reader can safely assume that every occurrence of $(x \sqsubseteq_p \mathfrak{s})$ serves this purpose. Note the use of \sqsubseteq_p can increase the width of the formula by 1; we discuss this in each case.

The main part of the proof is a structural induction along the definition of $\text{FO}[\text{Eq}, \mathcal{K}]$.

Base cases: We begin the construction with the base cases; the length of the constructed formula is discussed at the end of the whole construction.

- $x \doteq y$ where neither x nor y is min or max is realized by

$$(x \doteq y) \wedge (x \sqsubseteq_p \mathfrak{s}) \wedge (y \sqsubseteq_p \mathfrak{s}).$$

Simply using $x \doteq y$ is not enough, as we need to ensure the “prefix invariant”. This can increase the width of the formula by 1. If either of x or y is a constant, we simply replace any occurrence of min with ε and of max with \mathfrak{s} .

- $x < y$ where neither x nor y is a constant is realized by

$$(y \sqsubseteq_p \mathfrak{s}) \wedge \bigvee_{a \in \Sigma} \exists z: (y \doteq x a z).$$

We do not need to include $(x \sqsubseteq_p \mathfrak{s})$, as this is already implicitly ensured by the equations $y \doteq x a z$ in the disjunction. The new variable z increases the width by one (and we can also use this z for \sqsubseteq_p).

Now for the constants: If $y = \text{max}$, we consider three cases for x . If $x = \text{max}$, the formula is not satisfiable, and we realize it with the contradiction $(\mathfrak{s} \doteq a) \wedge (\mathfrak{s} \doteq aa)$ for some $a \in \Sigma$. If $x = \text{min}$, the formula is realized by $\exists z: \bigvee_{a \in \Sigma} (\mathfrak{s} \doteq az)$. If x is a variable, we construct the formula as in the general case and replace y with \mathfrak{s} .

If $y = \text{min}$, the formula is not satisfiable, and we realize it a contradiction (see above). Finally, if y is a variable, we only need to consider $x = \text{max}$ and $x = \text{min}$. In the first case, we have a contradiction (see above); the second is realized by $\exists z: \bigvee_{a \in \Sigma} (x \doteq az)$. Neither of the constructions increases the width by more than one.

- $P_a(x)$ is realized by $\exists z: (\mathfrak{s} \doteq x a z)$ if x is a variable, $P_a(\text{min})$ is realized by $\exists z: (\mathfrak{s} \doteq a z)$, and $P_a(\text{max})$ is realized by $\exists z: (\mathfrak{s} \doteq z a)$. In each case, the width is increased by one.
- $\text{succ}(x, y)$ for variables x and y is realized by

$$(y \sqsubseteq_p \mathfrak{s}) \wedge \bigvee_{a \in \Sigma} (y \doteq x a).$$

If $x = \text{max}$ or $y = \text{min}$, any contradiction realizes $\text{succ}(x, y)$. Moreover, $\text{succ}(\text{min}, \text{max})$ is realized by $\bigvee_{a \in \Sigma} \mathfrak{s} \doteq a$. Finally, for variables x or y , we realize $\text{succ}(x, \text{max})$ $\bigvee_{a \in \Sigma} \mathfrak{s} \doteq xa$ and $\text{succ}(\text{min}, y)$ by $(y \sqsubseteq_p \mathfrak{s}) \wedge \bigvee_{a \in \Sigma} y \doteq a$. Neither of these constructions increases the width by more than one.

- $\text{Eq}(x_1, y_1, x_2, y_2)$ is realized by

$$(y_1 \sqsubseteq_p \mathfrak{s}) \wedge (y_2 \sqsubseteq_p \mathfrak{s}) \wedge \exists z: (y_1 \doteq x_1 z \wedge y_2 \doteq x_2 z)$$

if all four parameters are variables. For constants, we adapt the construction as follows: If $y_i = \text{min}$, we replace $y_i \doteq x_i z$ with $(x_i \doteq \varepsilon) \wedge (z \doteq \varepsilon)$ and omit $(y_i \sqsubseteq_p \mathfrak{s})$. If $x_i = \text{min}$, we replace x_i in the constructed formulas with ε (removing the tautology $\varepsilon \doteq \varepsilon$ if it is created by a combination of this and the previous case occurring together). Every $x_i = \text{max}$ or $y_i = \text{max}$ is replaced with \mathfrak{s} . Again, all cases increase the width by at most one.

- $\hat{\kappa}(x_1, y_1, \dots, x_l, y_l)$ with $l := \text{ar}(\kappa)$ is realized by

$$\bigwedge_{i=1}^l (x_i \sqsubseteq_{\mathfrak{p}} \mathfrak{s}) \wedge \exists z_1, \dots, z_l : (\kappa(z_1, \dots, z_l) \wedge \bigwedge_{i=1}^l y_i \doteq x_i z_i)$$

if all x_i and y_i are variables. If constants are used, replace these with new variables x_{\min} and x_{\max} , add these to the existential quantifiers, and add the formulas $(x_{\min} \doteq \varepsilon)$ and $(x_{\max} \doteq \mathfrak{s})$ to the conjunction under the quantifiers. Hence, this case can increase the width by up to $2 + l$.

Recursive steps: For the recursive steps, let $\varphi, \varphi_1, \varphi_2 \in \text{FO}[\text{Eq}, \mathcal{K}]$ be formulas that are realized by $\psi, \psi_1, \psi_2 \in \text{FC}[\mathcal{K}]$, respectively.

- $\varphi_1 \wedge \varphi_2$ is realized by $\psi_1 \wedge \psi_2$.
- $\exists x : \varphi$ is realized by $\exists x : \psi$.
- $\forall x : \varphi$ is realized by

$$\forall x : (\neg(x \sqsubseteq_{\mathfrak{p}} \mathfrak{s}) \vee \psi),$$

which expresses $\forall x : ((x \sqsubseteq_{\mathfrak{p}} \mathfrak{s}) \rightarrow \psi)$. This guard is necessary, as the FC-quantifier ranges over all subwords of $\sigma(\mathfrak{s})$, but only prefixes of $\sigma(\mathfrak{s})$ are relevant for the FO-quantifier.

In cases where we prefer using a second additional variable over introducing a negation, we could instead use the formula

$$\forall x : \left(\psi \vee \bigvee_{a \in \Sigma} \bigvee_{b \in \Sigma - \{a\}} \exists z_1 : (\exists z_2 : \mathfrak{s} \doteq z_1 a z_2 \wedge \exists z_2 : \mathfrak{s} \doteq z_1 b z_2) \right).$$

- $\neg \varphi$ is realized by $\neg \psi \wedge \bigwedge_{x \in \text{free}(\varphi)} x \sqsubseteq_{\mathfrak{p}} \mathfrak{s}$.
- $\varphi_1 \vee \varphi_2$ is realized by

$$\left(\psi_1 \wedge \bigwedge_{x \in \text{free}(\varphi_2) - \text{free}(\varphi_1)} x \sqsubseteq_{\mathfrak{p}} \mathfrak{s} \right) \vee \left(\psi_2 \wedge \bigwedge_{x \in \text{free}(\varphi_1) - \text{free}(\varphi_2)} x \sqsubseteq_{\mathfrak{p}} \mathfrak{s} \right).$$

Complexity: Regarding the length of the constructed formula, note that the formulas for $<$ and succ depend on Σ . But as we assume Σ to be fixed, this is only a constant factor.

The only formulas that is affected by the width are the constraints and the disjunction: this leads to a factor of k and brings the length of the final formula to $k|\varphi|$. If no disjunctions occur and if the arity of the constraints is bounded (e. g., as we have for $\text{FC}[\text{REG}]$), this factor is not needed, and we get a length of $O(|\varphi|)$.

As all steps are straightforward, we can construct ψ in time $O(|\psi|)$.

The width of the formula is dominated by the construction for constraints, which can increase it by up to $2 + \max\{\text{ar}(\kappa) \mid \kappa \in \mathcal{K}\}$. If $\mathcal{K} = \emptyset$, the total width of the resulting formula is at most $k + 1$ instead (instead of $k + 2$). ◀

A.4 Proof of Lemma 3.15

For FC and FO, the *width* $\text{wd}(\varphi)$ of a formula φ is the maximum number of free variables in any of its subformulas. We prove the following stronger result:

Given $\varphi \in \text{FC}[\mathcal{K}]$ with $k := \text{wd}(\varphi)$, we can compute $\psi \in \text{FO}[\text{Eq}, \mathcal{K}]$ in time $O(k|\varphi|)$ that realizes φ and has $\text{wd}(\psi) = 2k + 3$. This preserves the properties existential and existential-positive.

Proof. We show this with a structural induction along the definition of $\text{FC}[\mathcal{K}]$. Recall that it is our goal to represent each FC -variable x through the two FO -variables x^o and x^c . We shall construct ψ in such a way that $\alpha(x^o) \leq \alpha(x^c)$ holds for all assignments α that satisfy ψ .

As we shall see in the case for word equations, the total number of variables can be lowered to $2|\text{free}(\varphi)| + 2$ if all word equations in φ have \mathfrak{s} on their left side. Our constructions use $x \leq y$ as shorthand for $x < y \vee x \doteq y$.

Word equations: Assume that $\varphi = (x_L \doteq \eta_R)$, with $x_L \in \Xi$ and $\eta_R \in (\Xi \cup \Sigma)^*$. We first handle a few special cases before proceeding to the main construction for word equations.

Word equations, special cases: We first handle the rather straightforward case of $\eta_R = \varepsilon$. Here, we distinguish two cases, namely $x_L = \mathfrak{s}$ and $x_L \neq \mathfrak{s}$. The first means that we are dealing with the equation $\mathfrak{s} \doteq \varepsilon$. This is true if and only if \mathcal{A}'_w contains only a single node. We express this with

$$\psi := (\min \doteq \max).$$

For $x_L \neq \mathfrak{s}$, we can directly define

$$\psi := (x_L^o \doteq x_L^c).$$

Recall that the spans of empty words in some word w are exactly the spans $[j, j]$ with $1 \leq j \leq |w| + 1$. Now for the more interesting case of $\eta_R \neq \varepsilon$. Here, we need to take care of one more special case; namely, that \mathfrak{s} appears in η_R . If η_R contains one or more occurrences of \mathfrak{s} , we distinguish the following sub-cases:

1. η_R contains at least one terminal,
2. η_R contains no terminals.

In the first case, we can conclude that there is no σ with $\sigma \models_{\mathfrak{s}} \varphi$. This is for the following reason: Assume $\sigma(x_L) = \sigma(\eta_R)$. This implies $|\sigma(x_L)| = |\sigma(\eta_R)|$. By definition, we also have $\sigma(x_L) \sqsubseteq \sigma(\mathfrak{s})$ and hence $|\sigma(x_L)| \leq |\sigma(\mathfrak{s})|$. As η_R contains \mathfrak{s} and at least one terminal (which is constant under σ), we have $|\sigma(\eta_R)| \geq |\sigma(\mathfrak{s})| + 1$. Thus, $|\sigma(\eta_R)| > |\sigma(\mathfrak{s})| \geq |\sigma(x_L)|$. Contradiction. As φ is not satisfiable, we choose the unsatisfiable formula

$$\psi := \exists x: (\text{P}_a(x) \wedge (x \doteq \max)).$$

Recall that we assume that we defined the node $|w| + 1$ in \mathcal{A}'_w to be letter-less, which also ensure that this formula is indeed unsatisfiable. This allows us to construct an unsatisfiable $\text{EPFO}[\text{Eq}]$ -formula that also works on \mathcal{A}'_ε and does not assume that $|\Sigma| \geq 2$.

In the second case, we know that $\eta_R \in \Xi^+$ and that it contains \mathfrak{s} at least once. If η_R contains \mathfrak{s} twice, then $\sigma \models_{\mathfrak{s}} \varphi$ can only hold if $\sigma(x) = \varepsilon$ holds for all $x \in \text{var}(\eta_R) \cup \{x_L\}$. This is due to a straightforward length argument: If $\sigma \models_{\mathfrak{s}} \varphi$, then $\sigma(x_L) = \sigma(\eta_R)$ and $\sigma(x_L) \sqsubseteq \sigma(\mathfrak{s})$. The first part implies $|\sigma(x_L)| = |\sigma(\eta_R)|$. As \mathfrak{s} appears at least twice in η_R , we have $|\sigma(\eta_R)| \geq 2|\sigma(\mathfrak{s})|$. Putting this together gives

$$|\sigma(\mathfrak{s})| \geq |\sigma(x_L)| \geq |\sigma(\eta_R)| \geq 2|\sigma(\mathfrak{s})|,$$

which implies $|\sigma(\mathfrak{s})| = 0$. This proves the claim. In this case, we define

$$\psi := \exists x: ((\min \doteq \max) \wedge \bigwedge_{y \in \text{free}(\varphi)} y^o \doteq y^c).$$

The big disjunction only serves to ensure that ψ has the correct free variables; as there are no other possible assignments in \mathcal{A}'_ε , we do not need to make the equality explicit.

Hence, we can safely assume that $\eta_R \in \Xi^+$ and that it contains \mathfrak{s} exactly once. Again we distinguish two cases, namely $|\eta_R| = 1$ and $|\eta_R| \geq 2$.

If $|\eta_R| = 1$, we have $\varphi = (x_L \doteq \mathfrak{s})$. If $x_L = \mathfrak{s}$, we are dealing with the trivial formula $\mathfrak{s} \doteq \mathfrak{s}$, and can just define

$$\psi := \exists x: (x \doteq x),$$

or some other trivially satisfiable formula. If $x_L \neq \mathfrak{s}$, we define

$$\psi := (x_L^o \doteq \min) \wedge (x_L^c \doteq \max)$$

to express this equality. It is convenient not to use Eq here, as x_L must encompass the whole structure.

Now for $|\eta_R| \geq 2$, where η_R contains exactly one occurrence of \mathfrak{s} . If $x_L = \mathfrak{s}$, we can see from a straightforward length argument that $\sigma \models_{\mathfrak{s}} \varphi$ if and only if $\sigma(y) = \varepsilon$ for all $y \in \text{var}(\eta_R) - \{\mathfrak{s}\}$. We express this with the formula

$$\psi := \bigwedge_{y \in \text{var}(\eta_R) - \{\mathfrak{s}\}} y^o \doteq y^c.$$

If $x_L \neq \mathfrak{s}$, we also need to ensure that $\sigma(x_L) = \sigma(\mathfrak{s})$ holds, as we have $\sigma(x_L) \sqsubseteq \sigma(\mathfrak{s})$ by definition and $\sigma(x_L) \supseteq \sigma(\mathfrak{s})$ from the fact that \mathfrak{s} occurs in η_R . We define

$$\psi := (x_L^o \doteq \min) \wedge (x_L^c \doteq \max) \wedge \bigwedge_{y \in \text{var}(\eta_R) - \{\mathfrak{s}\}} y^o \doteq y^c.$$

This also takes care of the case where x_L occurs in η_R . Then, we must have $\sigma(x_L) = \varepsilon$ in addition to $\sigma(x_L) = \sigma(\mathfrak{s})$.

Word equations, main construction: After covering these special cases, we can proceed with the main part of the construction. Let $\eta_R = \eta_1 \cdots \eta_n$, $n \geq 1$, with $\eta_i \in (\Xi \cup \Sigma)^+$ and $\eta_i \neq \mathfrak{s}$ for all $1 \leq i \leq n$. Note that $x_L = \mathfrak{s}$ might hold.

We shall first discuss how to construct an FO[Eq]-formula with $n + 1$ variables in addition to the $2|\text{free}(\varphi)|$ free variables from $\{x^o, x^c \mid x \in \text{var}(\eta_R)\}$ that are required by definition. After that, we shall describe how to reduce this to 3 additional variables (by reordering quantifiers and re-using variables, as commonly done for FO with a bounded number of variables).

These $n + 2$ additional variables are the variables y_1, \dots, y_{n+1} . The idea behind the construction is that each pair (y_i, y_{i+1}) shall represent the part of $\sigma(\eta_R)$ that is created by η_i . If we do not want to keep the number of variables low, we define

$$\hat{\psi} := \exists y_1, \dots, y_{n+1}: \begin{cases} (y_1 \doteq \min) \wedge \bigwedge_{i=1}^n \psi_i(y_i, y_{i+1}) \wedge (y_{n+1} \doteq \max) & \text{if } x_L = \mathfrak{s}, \\ \text{Eq}(x_L^o, x_L^c, y_1, y_{n+1}) \wedge \bigwedge_{i=1}^n \psi_i(y_i, y_{i+1}) & \text{if } x_L \neq \mathfrak{s}, \end{cases}$$

where the formulas ψ_i are defined as follows for all $1 \leq i \leq n$:

$$\psi_i(y_i, y_{i+1}) := \begin{cases} P_a(y_i) \wedge \text{succ}(y_i, y_{i+1}) & \text{if } \eta_i = a \in \Sigma, \\ \text{Eq}(x^o, x^c, y_i, y_{i+1}) & \text{if } \eta_i = x \in X. \end{cases}$$

Although $\hat{\psi}$ is directly obtained from the pattern η_R , some explanations are warranted. Firstly, note that \mathfrak{s} only plays a role if we have $x_L = \mathfrak{s}$. In this case, the use of min and max ensures that η_R encompasses all of \mathfrak{s} .

Moreover, observe that the construction ensures that $\text{free}(\hat{\psi})$ is the set of all x^o and x^c such that $x \in \text{free}(\hat{\varphi})$. If $x_L = \mathfrak{s}$, then $\text{free}(\hat{\varphi}) = \text{var}(\eta_R)$, and the variables x^o and x^c are “introduced” in the ψ_i where $\eta_i = x$ holds. But if $x_L \neq \mathfrak{s}$ and $x_L \notin \text{var}(\eta_R)$, then $\text{Eq}(x_L^o, x_L^c, y_1, y_{n+1})$ not only ensures that x_L and η_R are mapped to the same word, but also that $\text{free}(\hat{\psi})$ contains x_L^o and x_L^c .

Finally, we observe that the construction does not need to specify that $y_i \leq y_{i+1}$ or $x^o \leq x^c$ holds. By definition, succ and Eq guarantee this property and can act as guards.

Keeping this in mind, one can now prove by induction that for every $w \in \Sigma^*$, we have $\alpha \in \llbracket \hat{\psi} \rrbracket(w)$ if and only if α expresses some pattern substitution σ with $\sigma \in \llbracket \varphi \rrbracket(w)$. In other words, $\hat{\psi}$ realizes φ . All that remains is to reduce the number of variables through a standard re-ordering and renaming process.

We first discuss the case of $x_L = \mathfrak{s}$, where we need only two variables. Observe that for $2 \leq i \leq n$, the variable y_i is only used in the sub-formulas ψ_{i-1} and ψ_{i+1} . Similarly, y_1 is only used in ψ_1 and in $(y_1 \doteq \min)$, and y_{n+1} is only used in ψ_n and $(y_{n+1} \doteq \max)$. This allows us to use shift the quantifiers into the conjunction, which leads to the following formula:

$$\begin{aligned} \psi' := \exists y_1, y_2: & \left((y_1 \doteq \min) \wedge \psi_1(y_1, y_2) \right. \\ & \wedge \exists y_3: (\psi_2(y_2, y_3)) \\ & \wedge \exists y_4: (\psi_3(y_3, y_4)) \\ & \quad \vdots \\ & \left. \wedge \exists y_{n+1}: (\psi_n(y_n, y_{n+1}) \wedge (y_{n+1} \doteq \max)) \dots \right), \end{aligned}$$

for which $\psi' \equiv \hat{\psi}$ holds. As observed above, each variable y_i is only used together with y_{i+1} . Accordingly, we now obtain ψ from ψ' by replacing every variable y_i where i is odd with z_1 , and every y_i where i is even with z_2 . Then ψ has only $|\text{free}(\varphi)| + 2$ variables. Moreover, $\psi \equiv \psi' \equiv \hat{\psi}$ holds; and as we already established that $\hat{\psi}$ realizes φ , we conclude that ψ realizes φ .

For the case of $x_L \neq \mathfrak{s}$, observe that $\hat{\psi}$ contains $\text{Eq}(x_L^o, x_L^c, y_1, y_{n+1})$. Hence, we cannot move the quantifier for y_{n+1} to the “bottom” of the formula. Instead, we define

$$\begin{aligned} \psi' := \exists y_1, y_2, y_{n+1}: & \left(\text{Eq}(x_L^o, x_L^c, y_1, y_{n+1}) \wedge \psi_1(y_1, y_2) \right. \\ & \left. \wedge \exists y_3: (\psi_2(y_2, y_3) \wedge \exists y_4: (\psi_3(y_3, y_4) \dots \wedge \exists y_n: (\psi_n(y_n, y_{n+1})) \dots)) \right). \end{aligned}$$

We now obtain ψ by renaming the y_i with $1 \leq i \leq n$ as in the previous case. Hence, the only difference is that y_{n+1} remains unchanged, which leads to a total of $2|\text{free}(\varphi)| + 3$ variables.

Constraints: If $\varphi = \kappa(x_1, \dots, x_l)$ with $l := \text{ar}(\kappa)$, we first consider the straightforward case if $\vec{\mathfrak{s}} \notin \vec{x}$. Then, we can simply define

$$\psi := \hat{\kappa}(x_1^o, x_1^c, \dots, x_l^o, x_l^c).$$

If $\vec{\mathfrak{s}} \in \vec{x}$, let S be the set of all $i \in [l]$ with $x_i = \mathfrak{s}$ and let $X := [l] - S$. We now define

$$\psi := \exists z^o, z^c: ((z^o \doteq \min) \wedge (z^c \doteq \max) \wedge \hat{\kappa}(\vec{y})),$$

where $\vec{y} := (y_1^o, y_1^c, \dots, y_l^o, y_l^c)$ is defined by

$$(y_i^o, y_i^c) := \begin{cases} (z^o, z^c) & \text{if } i \in S, \\ (x_i^o, x_i^c) & \text{if } i \in X. \end{cases}$$

Note that in both cases, the definition of the semantics of $\hat{\kappa}$ also guarantees $x_i^o \leq x_i^c$ for all $i \in X$.

Conjunctions: If $\varphi = \varphi_1 \wedge \varphi_2$, we define $\psi := \psi_1 \wedge \psi_2$, where ψ_1 and ψ_2 realize φ_1 and φ_2 , respectively. The correctness of this construction follows directly from the induction assumption and Definition 3.14.

Disjunctions: If $\varphi = \varphi_1 \vee \varphi_2$, we first construct the FO[Eq, \mathcal{K}]-formulas ψ_1 and ψ_2 that realize φ_1 and φ_2 , respectively. We cannot just define ψ as $\psi_1 \vee \psi_2$. Unless $\text{free}(\varphi_1) = \text{free}(\varphi_2)$ holds, this definition would accept assignments that do not realize any pattern substitution. For example, if we have $\alpha \in \llbracket \psi_1 \rrbracket \langle w \rangle$ with $\alpha(x^o) > \alpha(x^c)$ for some variable $x \in \text{free}(\varphi_2) - \text{free}(\varphi_1)$, then $\alpha \in \llbracket \psi_1 \vee \psi_2 \rrbracket \langle w \rangle$ holds.

We address this problem by guarding variables that are only free in exactly one formula, and define

$$\psi := \left(\psi_1 \wedge \bigwedge_{x \in \text{free}(\varphi_2) - \text{free}(\varphi_1)} (x^o \leq x^c) \right) \vee \left(\psi_2 \wedge \bigwedge_{x \in \text{free}(\varphi_1) - \text{free}(\varphi_2)} (x^o \leq x^c) \right).$$

For all $w \in \Sigma^*$, we now have $\alpha \in \llbracket \psi \rrbracket \langle w \rangle$ if and only if $\alpha \in \llbracket \psi_i \rrbracket \langle w \rangle$ for $i \in \{1, 2\}$, and $\alpha(x^o) \leq \alpha(x^c)$ for all $x \in \text{free}(\varphi)$, which holds if and only if α expresses some $\sigma \in \llbracket \varphi_i \rrbracket \langle w \rangle$.

Negations: If $\varphi = \neg \hat{\varphi}$, we first construct $\hat{\psi}$ that realizes $\hat{\varphi}$, and then define

$$\psi := \neg \hat{\psi} \wedge \left(\bigwedge_{x \in \text{free}(\hat{\varphi})} x^o \leq x^c \right).$$

We face an issue that is analogous to the one for disjunction; defining $\neg \hat{\psi}$ would lead to a formula that accepts assignments that do not express a pattern substitution. Again, the solution is guarding the free variables of $\hat{\varphi}$.

Existential quantifiers: If $\varphi = \exists x: \hat{\varphi}$, construct a formula $\hat{\psi}$ that realizes $\hat{\varphi}$, and define $\psi := \exists x^o, x^c: \hat{\psi}$. As $x^o \leq x^c$ is guaranteed as an induction invariant, we do not need to guard the two variables.

Universal quantifiers: If $\varphi = \forall x: \hat{\varphi}$, construct $\hat{\psi}$ that realizes $\hat{\varphi}$, and define

$$\psi := \forall x^o, x^c: ((x^o > x^c) \vee \hat{\psi}),$$

which amounts to defining $\forall x^o, x^c: ((x^o \leq x^c) \rightarrow \hat{\psi})$. Again, we need to deal with the induction invariant: If we simply defined $\forall x^o, x^c: \hat{\psi}$, then the formula would be invalid on all non-empty w .

Complexity considerations: The special cases for word equations can be checked in time $O(|\varphi|)$ and create formulas of constant length. The main construction for word equations creates a formula of length $O(|\eta_R|)$ and takes proportional time.

The only recursive cases that create formulas of a length more than linear are negation and disjunctions. Here, the guards increase the formula length to $O(k|\varphi|)$, which dominate the final formula length and the total running time. Hence, if φ contains neither negations nor disjunctions, we have $|\psi| \in O(|\varphi|)$; and the same holds for the run time. \blacktriangleleft

A.5 Proof of Lemma 3.19

Proof. This follows from a straightforward induction along the definition of GC.

- *Base case:* If φ is a word equation $y \doteq \eta$, then $\text{main}(\varphi) = \{y\}$, and $x \in \text{var}(\eta)$ holds every auxiliary variable x of φ . Then $\sigma \models \varphi$ if and only if $\sigma(y) = \sigma(\eta)$, which immediately implies $\sigma(y) \sqsupseteq \sigma(x)$.
- *Conjunctions:* If $\varphi = (\varphi_1 \wedge \varphi_2)$ and $x \in \text{aux}(\varphi)$, then $x \in \text{aux}(\varphi_1)$ or $x \in \text{aux}(\varphi_2)$. Let $x \in \text{aux}(\varphi_1)$, the other case proceeds analogously. If $\sigma \models \varphi$, then $\sigma \models \varphi_1$. By the induction hypothesis, there exists $y \in \text{main}(\varphi_1)$ with $\sigma(y) \sqsupseteq \sigma(x)$. As $y \in \text{main}(\varphi)$, this proves the claim.
- *Disjunctions:* As for conjunctions (and even more straightforward, as $\text{aux}(\varphi_1) = \text{aux}(\varphi_2)$).
- *Quantifiers:* If $\varphi = \exists x: \psi$ or $\varphi = \forall x: \psi$, then the claim follows directly from the induction hypothesis, due to $\text{main}(\varphi) = \text{main}(\psi)$ and $\text{aux}(\varphi) = \text{aux}(\psi) - \{x\}$.
- *Guarded constraints and guarded negations:* If $\varphi = (\varphi_1 \wedge \kappa(\vec{x}))$ or $\varphi = (\varphi_1 \wedge \neg\varphi_2)$, then $\sigma \models \varphi$ implies $\sigma \models \varphi_1$. As $\text{main}(\varphi) = \text{main}(\varphi_1)$ and $\text{aux}(\varphi) = \text{aux}(\varphi_1)$, the claim follows directly from the induction hypothesis. ◀

A.6 Proof of Lemma 3.20

Proof. Let \vec{s} be a tuple of structure variables. Our goal is to show that each $\varphi \in \text{FC}[\mathcal{K}](\vec{s})$ can be converted into an equivalent $\psi \in \text{GC}[\mathcal{K}]$. Recall that this implies that $\text{struc}(\varphi) = \text{main}(\psi)$ and $\text{free}(\varphi) = \text{aux}(\psi)$.

We show this using an induction along Definition 3.1. The *width* $\text{wd}(\varphi)$ of some $\varphi \in \text{FC}$ is defined as the maximum number of free variables over all its subformulas. For $\psi \in \text{GC}$, we use the number of auxiliary variables to define $\text{wd}(\psi)$.

In each of the steps, it shall be easy to see that $|\psi| \in O(|\varphi| |\vec{s}|^2 \text{wd}(\varphi))$, and that ψ can be constructed in time that is proportional to its length. The construction also introduces neither new universal quantifiers, nor new negations. Hence, the resulting formula is existential or existential-positive if and only if the original formula had this property.

Word equations: If φ is of the form $x \doteq \eta$, we distinguish two cases. Firstly, consider the case $x \in \vec{s}$. In other words, x is a structure variable. As $y \in \text{var}(\eta)$ must hold for all $y \in \text{aux}(\varphi)$, we know that $\sigma \models \varphi$ implies $\sigma(x) \sqsupseteq \sigma(y)$ for all $y \in \text{aux}(\varphi)$. Hence, we only need to ensure that all structure variables from \vec{s} actually appear in ψ as the left side of a word equation, as GC requires this. To this end, we simply define $\psi := (x \doteq \eta) \wedge \bigwedge_{s \in \vec{s}} s \doteq s$.

Technically, the last conjunction does not need to include x , but $s \in \vec{s}$ is easier to read than $s \in \vec{s} - \{x\}$. While this construction is not optimal, it does not significantly affect the complexity of the construction, so we opt for the slightly less efficient but more readable approach. This also applies to most of the other cases in the proof, but we only point this out in this specific case.

We need to spend a little more effort in the second case, namely that $x \notin \vec{s}$. By definition of the semantics of FC, we know that $\sigma \models \varphi$ implies that there is some $\mathfrak{s} \in \vec{s}$ with $\sigma(\mathfrak{s}) \sqsupseteq \sigma(x)$. Note that \mathfrak{s} might not be unique, and different choices of σ might require different choices of \mathfrak{s} . Accordingly, we define

$$\psi := \bigvee_{\mathfrak{s} \in \vec{s}} \left((\exists p, s: (\mathfrak{s} \doteq p x s \wedge \mathfrak{s} \doteq p \eta s)) \wedge \bigwedge_{\hat{\mathfrak{s}} \in \vec{s}} \hat{\mathfrak{s}} \doteq \hat{\mathfrak{s}} \right).$$

The last part of ψ , the conjunction of the trivial equations $\hat{\mathfrak{s}} \doteq \hat{\mathfrak{s}}$, simply serves to ensure the GC-requirement that all parts of a disjunction have the same free variables. The other part of the formula is more important: for all σ , we have $\sigma \models \exists p, s: (\mathfrak{s} \doteq p x s \wedge \mathfrak{s} \doteq p \eta s)$ if and only if $\sigma(x) \sqsubseteq \sigma(\mathfrak{s})$ and that $\sigma(x) = \sigma(\eta)$.

In both cases, we have $\text{main}(\psi) = \vec{\mathfrak{s}} = \text{struc}(\varphi)$ and $\text{aux}(\psi) = \text{var}(\eta) - \vec{\mathfrak{s}} = \text{free}(\varphi)$.

Constraint symbols: For $\varphi = \kappa(\vec{x})$ with $\kappa \in \mathcal{K}$, we distinguish similar cases as for word equations. If $\vec{x} \subseteq \vec{\mathfrak{s}}$, then we do not need to guard the elements of \vec{x} , and can simply define

$$\psi := \left(\bigwedge_{\mathfrak{s} \in \vec{\mathfrak{s}}} \mathfrak{s} \doteq \mathfrak{s} \right) \wedge \kappa(\vec{x}).$$

On the other hand, every $x \in \vec{x}$ that is not a structure variable needs to be guarded, which we achieve with a similar construction as for word equations. We define

$$\psi := \bigwedge_{x \in (\vec{x} - \vec{\mathfrak{s}})} \bigvee_{\mathfrak{s} \in \vec{\mathfrak{s}}} \left(\exists p, s: \mathfrak{s} \doteq p x s \wedge \bigwedge_{\hat{\mathfrak{s}} \in \vec{\mathfrak{s}}} \hat{\mathfrak{s}} \doteq \hat{\mathfrak{s}} \right) \wedge \kappa(\vec{x})$$

Again, we do not know in which structure variable \mathfrak{s} the free variable x can be embedded; hence, ψ accounts for all possibilities. This disjunction also acts as the guarding formula that GC requires for the use of constraints.

In both cases, we have $\text{main}(\psi) = \vec{\mathfrak{s}} = \text{struc}(\varphi)$ and $\text{aux}(\psi) = \vec{x} = \text{free}(\varphi)$. Take note that $|\psi| \in O(\text{wd}(\kappa(\vec{x}))|\vec{\mathfrak{s}}|^2)$. Recall that according to our definition of free variables for FC, structure variables do not count as free variables. In particular, if $\vec{x} \subseteq \vec{\mathfrak{s}}$, we have $\text{wd}(\kappa(\vec{x})) = 0$.

Conjunctions: For $\varphi = (\varphi_1 \wedge \varphi_2)$ with $\varphi_1, \varphi_2 \in \text{FC}(\vec{\mathfrak{s}})$, we first construct equivalent $\psi_1, \psi_2 \in \text{GC}$, and then define $\psi := \psi_1 \wedge \psi_2$. As GC has no specific requirements for conjunction, this suffices. We have $\text{main}(\psi) = \text{main}(\psi_1) \cup \text{main}(\psi_2) = \vec{\mathfrak{s}} = \text{struc}(\varphi)$ and $\text{aux}(\psi) = \text{aux}(\psi_1) \cup \text{aux}(\psi_2) = \text{free}(\varphi_1) \cup \text{free}(\varphi_2) = \text{free}(\varphi)$.

Disjunctions: For $\varphi = \varphi_1 \vee \varphi_2$ with $\varphi_1, \varphi_2 \in \text{FC}(\vec{\mathfrak{s}})$, we also first construct equivalent $\psi_1, \psi_2 \in \text{GC}$. While these have the same main variables (namely, $\vec{\mathfrak{s}}$), they might have different auxiliary variables, which means that we cannot simply define ψ as their disjunction. Instead, we artificially “inflate” the sets of auxiliary variables in each of the formulas, using the following construction:

$$\psi := \left(\psi_1 \wedge \bigwedge_{x \in \text{aux}(\psi_1) - \text{aux}(\psi_2)} \bigvee_{\mathfrak{s} \in \vec{\mathfrak{s}}} \left((\exists p, s: \mathfrak{s} \doteq p x s) \wedge \bigwedge_{\hat{\mathfrak{s}} \in \vec{\mathfrak{s}}} \hat{\mathfrak{s}} \doteq \hat{\mathfrak{s}} \right) \right) \vee \left(\psi_2 \wedge \bigwedge_{x \in \text{aux}(\psi_2) - \text{aux}(\psi_1)} \bigvee_{\mathfrak{s} \in \vec{\mathfrak{s}}} \left((\exists p, s: \mathfrak{s} \doteq p x s) \wedge \bigwedge_{\hat{\mathfrak{s}} \in \vec{\mathfrak{s}}} \hat{\mathfrak{s}} \doteq \hat{\mathfrak{s}} \right) \right).$$

Again, we ensure that the inner disjunctions range over all structure variables, and that the outer disjunctions ranges over the same auxiliary variables. The new parts of the formula (i. e., everything that is not ψ_1 or ψ_2), simply express that x is subword of some $\mathfrak{s} \in \vec{\mathfrak{s}}$. We observe $\text{main}(\psi) = \vec{\mathfrak{s}} = \text{struc}(\varphi)$ and $\text{aux}(\psi) = \text{aux}(\psi_1) \cup \text{aux}(\psi_2) = \text{free}(\varphi_1) \cup \text{free}(\varphi_2) = \text{free}(\varphi)$. This is one of the cases where the width affects the length of ψ , as we have $|\psi| \in O(|\varphi| |\vec{\mathfrak{s}}|^2 \text{wd}(\varphi))$.

Existential and universal quantifiers: If $\varphi = \exists x: \hat{\varphi}$ or $\varphi = \forall x: \hat{\varphi}$, then by definition $x \notin \vec{\mathfrak{s}}$. Hence, we can construct an equivalent $\hat{\psi} \in \text{GC}$ from $\hat{\varphi}$, and define $\psi := \exists x: \hat{\psi}$ or $\psi := \forall x: \hat{\psi}$. In both cases, we have $\text{main}(\psi) = \text{main}(\hat{\psi}) = \vec{\mathfrak{s}} = \text{struc}(\varphi)$ and $\text{aux}(\psi) := \text{aux}(\hat{\psi}) - \{x\} = \text{free}(\hat{\varphi}) - \{x\} = \text{free}(\varphi)$.

Negations: For $\varphi = \neg\hat{\varphi}$, we first take $\hat{\varphi}$ and construct an equivalent $\hat{\psi} \in \text{GC}$. We then define $\psi := \Upsilon \wedge \neg\hat{\psi}$, where

$$\Upsilon := \bigwedge_{x \in \text{aux}(\hat{\psi})} \bigvee_{\mathfrak{s} \in \vec{\mathfrak{s}}} \left((\exists p, s: \mathfrak{s} \doteq p x s) \wedge \bigwedge_{\hat{\mathfrak{s}} \in \vec{\mathfrak{s}}} \hat{\mathfrak{s}} \doteq \hat{\mathfrak{s}} \right).$$

Clearly, $\sigma \models \Upsilon$ if and only if for every $x \in \text{aux}(\hat{\psi})$, there is an $\mathfrak{s} \in \vec{\mathfrak{s}}$ with $\sigma(x) \sqsubseteq \sigma(\mathfrak{s})$. This means that it can act as a guard for the negation of $\hat{\psi}$, and that ψ is equivalent to φ . Note that $\text{aux}(v) = \text{aux}(\hat{\psi})$ and, hence, $\text{aux}(\psi) = \text{aux}(\hat{\psi}) = \text{free}(\hat{\varphi}) = \text{free}(\varphi)$. This is the other case where the width affects $|\psi|$.

Complexity of the construction: As mentioned above, it is easily seen that $|\psi|$ is in $O(|\varphi| |\vec{\mathfrak{s}}|^2 \text{wd}(\varphi))$. Note that the quadratic blowup comes from ensuring that some auxiliary variable x is a subword of some structure variable $\vec{\mathfrak{s}}$, and we do not know which structure variable needs to be used.

If each x is guaranteed to be subword of one specific structure variable \mathfrak{s} , then the disjunction over all possible structure variables and the associated conjunction over all other structure variables can be avoided. This lowers $|\psi|$ to $O(|\varphi| |\vec{\mathfrak{s}}| \text{wd}(\varphi))$.

Similarly, if all disjunctions and negations are already guarded by other formulas with the same free variables, we can avoid this conjunctions over auxiliary variables, and shave off the factor $|\vec{\mathfrak{s}}|^2 \text{wd}(\varphi)$ for these cases (although we would still have to deal with the $|\vec{\mathfrak{s}}|^2$ factor from the base cases). If both conditions apply, we can lower $|\psi|$ to $O(|\varphi| |\vec{\mathfrak{s}}|)$. ◀

A.7 Definition A.4

▶ **Definition A.4.** Let $\varphi \in \text{SpLog}^\neg(\mathbf{w})$. A spanner P with $\text{SVars}(P) = \text{free}(\varphi) - \{\mathbf{w}\}$ realizes φ if, for all $w \in \Sigma^*$, we have $\mu \in P(w)$ if and only if $\sigma \in \llbracket \varphi \rrbracket(w)$ holds for the substitution that is defined by $\sigma(x) := w_{\mu(x)}$ for all $x \in \text{SVars}(P)$.

A substitution σ expresses a (V, w) -tuple μ if for all $x \in V$, we have $\sigma(x^P) = w_{[1, i]}$ and $\sigma(x^C) = w_{[i, j]}$ for $[i, j] = \mu(x)$. Let P be a spanner and let $\varphi \in \text{SpLog}^\neg(\mathbf{w})$ with $\text{free}(\varphi) = \{\mathbf{w}\} \cup \{x^P, x^C \mid x \in \text{SVars}(P)\}$. Then φ realizes ψ if, for all $w \in \Sigma^*$, we have $\sigma \in \llbracket \varphi \rrbracket(w)$ if and only if σ expresses some $\mu \in P(w)$.

In other words, x^C is $w_{\mu(x)}$ (the content of x), and x^P is the prefix of w before $w_{\mu(x)}$. The main variable \mathbf{w} of a SpLog^\neg -formula has the same role as the input text of document spanner.

We extend this definition from comparing SpLog^\neg -formulas and spanners to comparing $\text{FC}_{(1)}$ -formulas and spanners.

A.8 Proof of Theorem 3.21

Proof. By Lemma 3.20, there are polynomial-time conversions between $\text{FC}[\text{REG}]\langle \mathfrak{s} \rangle$ -formulas and $\text{GC}[\text{REG}]$ -formula with one main variable (i.e., the SpLog^\neg -formulas). Combining these with the polynomial-time conversions between SpLog^\neg and $\text{RGX}^{\text{gcore}}$ from [30], we get $\text{FC}[\text{REG}]\langle \mathfrak{s} \rangle \equiv_{\text{poly}} \text{RGX}^{\text{gcore}}$.

The situation is analogous for the steps between $\text{EPFC}[\text{REG}]\langle \mathfrak{s} \rangle$ and EPGC with one main variable (aka SpLog), and between SpLog and RGX^{core} . ◀

A.9 Proof of Theorem 3.22

Proof. Large parts of this proof follow directly from the proof of Theorem ?? (Theorem 4.9 in [30]) and the results that we already established in the present paper.

From FC to spanners: For the conversion from $\text{FC}[\langle \mathfrak{s} \rangle]$ to $\text{sRGX}^{\text{gcore}}$, we observe that Lemma 3.20 allows us to convert constraint-free FC-formula into an equivalent SpLog^- -formula of polynomial size. The proof of Theorem 4.9 in [30] then converts this into a spanner representation from $\text{RGX}^{\text{gcore}}$. But as the formula is constraint-free, the construction from that proof (Section 4.2.1 of [30]) only creates regex formulas that contain the $*$ operator only in the form of Σ^* . Hence, the regex formulas are simple. Hence, we have a conversion from $\text{FC}[\langle \mathfrak{s} \rangle]$ to $\text{sRGX}^{\text{gcore}}$.

For EFC, we first use Lemma 3.9 to obtain an EPFC-formula. We then use Lemma 3.20 to obtain a constraint-free SpLog -formula, which is converted to a spanner representation from $\text{sRGX}^{\text{core}}$, using the same reasoning as above.

From spanners to FC: We only need to show how to construct a formula $\varphi \in \text{EFC}[\langle \mathfrak{s} \rangle]$ that realizes a given $\alpha \in \text{sRGX}$. The other parts of the construction are then handled by the proof from [30].

Section 4.2.2 of [30] contains a conversion of functional regex formulas to SpLog -formulas. The only reason that why cannot directly use this construction is that it converts regex formulas without variables (that are subformulas of formulas with variables) to regular constraints.

Hence, to adapt this construction, all we need to do is to show that given an $\alpha \in \text{sRGX}$ that has no variables, we can construct in polynomial time a formula $\varphi \in \text{EPFC}[\langle \mathfrak{s} \rangle]$ with $\mathcal{L}(\varphi) = \mathcal{L}(\alpha)$. We define $\varphi(\mathfrak{s}) := \exists x: (\mathfrak{s} \dot{=} x \wedge \varphi_\alpha(x))$, where φ_α is defined by recursively. The cases for all operators except the star are straightforward: We define $\varphi_\emptyset := (x \dot{=} ax)$ for some $a \in \Sigma$ (or any other positive formula that is not satisfiable), $\varphi_\varepsilon := (x \dot{=} \varepsilon)$, $\varphi_a := (x \dot{=} a)$ for each $a \in \Sigma$, $\varphi_{(\alpha_1 \cdot \alpha_2)} := \exists x_1, x_2: (x \dot{=} x_1 x_2 \wedge \varphi_{\alpha_1}(x_1) \wedge \varphi_{\alpha_2}(x_2))$, and $\varphi_{(\alpha_1 \cup \alpha_2)} := \varphi_{\alpha_1}(x) \vee \varphi_{\alpha_2}(x)$.

For the star operator, simple regex-formulas allow only two choices, namely $\alpha = \Sigma^*$ or $\alpha = w^*$ for some $w \in \Sigma^*$. The first case is also straightforward; we define $\varphi_{\Sigma^*} := (x \dot{=} x)$.

For the second case, we exclude the case $w = \varepsilon$, for which the formula $x \dot{=} \varepsilon$ suffices. Our construction adapts the construction for the respective result for EC (Theorem 5 in [41]) to EPFC and uses the following well-known fact from combinatorics on words: For every $w \in \Sigma^*$, let $\varrho(w)$ denote the *root* of w ; that is, the shortest word r such that w can be written as $w = r^k$ for some $k \geq 0$. For all $u, v \in \Sigma^+$, we have $uv = vu$ if and only if $\varrho(u) = \varrho(v)$ (see e. g. Lothaire [49]).

This allows us to express w^* in the following way: Let $p \geq 1$ be the unique value for which $w = \varrho(w)^p$ holds. We now define

$$\varphi_{w^*} := (x \dot{=} \varepsilon) \vee (x \dot{=} w) \vee \psi,$$

where

$$\psi := \begin{cases} \exists y: (x \dot{=} yw \wedge x \dot{=} wy) & \text{if } p = 1, \\ \exists y, z: (x \dot{=} y^p \wedge z \dot{=} y \cdot \varrho(w) \wedge z \dot{=} \varrho(w) \cdot y) & \text{if } p \geq 2. \end{cases}$$

Next, we show that $\sigma \models \varphi_{w^*}$ if and only if $\sigma(x) \in w^*$ and $\sigma(x) \sqsubseteq \sigma(\mathfrak{s})$.

We begin with the *only-if-direction*. Let $\sigma(x) = w^i$ for $i \geq 0$ and $\sigma(\mathfrak{s}) \sqsupseteq \sigma(x)$. If $i = 0$ or $i = 1$, we have $\sigma \models (x \dot{=} \varepsilon)$ or $\sigma \models (x \dot{=} w)$. Hence, we can assume $i \geq 2$. We first consider the case $p = 1$. Let $\tau := \sigma_{y \mapsto w^{i-1}}$. Then $\tau(y) \sqsubseteq \tau(x) \sqsubseteq \tau(\mathfrak{s})$ holds by definition. Furthermore, we have $\tau \models (x \dot{=} yw)$ due to

$$\tau(x) = w^i = w^{i-1}w = \tau(y)w = \tau(yw)$$

and $\tau \models (x \doteq yw)$ for analogous reasons. Hence, $\sigma \models \psi$ and, thereby $\sigma \models \varphi_{w^*}$. This concludes the case $p = 1$.

For the case $p \geq 2$, note that $w^i = \varrho(w)^{ip}$. We define the pattern substitution τ by $\tau(y) := \varrho(w)^i$, $\tau(z) := \varrho(w)^{i+1}$, and $\tau(u) = \sigma(u)$ for all other $u \in \Xi$; and claim $\sigma \models \psi$.

First, note that as $p \geq 2$ and $i \geq 2$, we have $i + 1 \leq ip$. This implies $\tau(z) = \varrho(w)^{i+1} \sqsubseteq \varrho(w)^{ip} = \tau(x)$ and, hence, $\tau(y) \sqsubseteq \tau(z) \sqsubseteq \tau(\mathfrak{s})$. Now we have

$$\begin{aligned} \tau(x) &= \varrho(w)^{ip} = \tau(y^p), \\ \tau(z) &= \varrho(w)^{i+1} = \tau(y \cdot \varrho(w)), \\ \tau(z) &= \varrho(w)^{i+1} = \tau(\varrho(w) \cdot y). \end{aligned}$$

Hence, $\sigma \models \psi$, and thereby $\sigma \models \varphi_{w^*}$. This concludes the case of $p \geq 2$ and this direction of the proof.

For the *if-direction*, assume $\sigma \models \varphi_{w^*}$. Then $\sigma(x) = \varepsilon$, $\sigma(x) = w$, or $\sigma \models \psi$. There is nothing to argue in the first two cases, so assume the third holds. Again, we need to distinguish $p = 1$ and $p \geq 2$.

We begin with $p = 1$, and consider any $v \in \Sigma^*$ such that $\tau \models (x \doteq yw \wedge x \doteq wy)$ for $\tau := \sigma_{y \rightarrow v}$. Then we have $\tau(x) = \tau(yw) = \tau(wy)$ and, hence, $vw = wv$. This holds if and only if $u = \varepsilon$ or, due to the fact mentioned above, $\varrho(u) = \varrho(w)$. In either case, we know that there exists some $i \geq 0$ with $v = w^i$. Hence, $\tau(x) = uw = w^{i+1}$. As $\sigma(x) = \tau(x)$, we have $\sigma(x) \in w^+$.

For $p \geq 2$, consider $u, v \in \Sigma^*$ such that $\tau \models (x \doteq y^p \wedge z \doteq y \cdot \varrho(w) \wedge z \doteq \varrho(w) \cdot y)$ for $\tau := \sigma_{y \rightarrow u, z \rightarrow v}$. Due to the last two equations, we have $u\varrho(w) = \varrho(w)u$. Again, we invoke the fact, and observe there is some $i \geq 0$ with $u = \varrho(w)^i$. Hence, $\tau(x) = \varrho(w)^{ip} = w^i$ and therefore, $\sigma(x) \in w^*$. This concludes this direction and the whole correctness proof.

Regarding the complexity of the construction, note that the length of φ_α is linear in $|\alpha|$. We conclude that φ can be constructed in polynomial time. \blacktriangleleft

B Appendix for Section 4

B.1 Proof of Theorem 4.1

Proof. We first consider the upper bounds for data complexity and combined complexity, and then the lower bounds for EFC and for FC.

Data complexity: We could show this directly, by defining a finite automaton with a read-only input tape and 2 pointers for each variable in the formula. This is straightforward, but somewhat tedious.

Instead, we refer to Lemma 3.15, and observe that ever $\varphi \in \text{FC}[\mathcal{K}]$ can be converted into an $\text{FO}[\text{Eq}, \mathcal{K}]$ -formula ψ that realizes φ . Observe that Eq and, by our complexity assumption on constraints, all $\kappa \in \mathcal{K}$ can be decided in logarithmic space.

Using standard methods for FO and in particular $\text{FO}[\prec]$ (see e. g. [24, 46]), we can then convert ψ into a L-Turing machine that uses machines for Eq and the constraint symbols as sub-programs.

Combined complexity, upper bounds: This is straightforward: For EPFC, we only need to deal with existential quantifiers, and as every quantified variable has to be a subword of a structure variable, these can be guessed. For every substitution and every word equation, $\sigma \models x \doteq \alpha$ can be verified in linear time, and constraints can be checked in polynomial time by

our assumption. This results in an NP-algorithm. For FC, we can represent all quantified variables in polynomial space, and enumerating all possible choices for these still results in a PSPACE-algorithm. Analogous arguments were made for SpLog and SpLog[∇] in [30, 31].

Combined complexity, NP lower bound: This can be directly derived from Theorem 3.1 in [32], which states that there is a subclass of Boolean core spanners for which the evaluation problem is NP-complete, even on words of length 1. A close look reveals that the spanner representations that are constructed in that proof are from sRGX^{core}. Thus, Theorem 3.22 allows us to convert them in polynomial time into EPFC₍₁₎[\cdot]-formulas. Hence, the problem is NP-complete.

Combined complexity, PSPACE lower bound: The idea is very similar to Theorem 6.16 in Libkin [46].

We prove the PSPACE lower bound with a reduction from the QBF-3SAT problem, which is stated as follows: Given a well-formed quantified Boolean formula $\psi = Q_1 v_1 : \dots : Q_k v_k : \psi_C$, where $k \geq 1$, $Q_i \in \{\exists, \forall\}$, and ψ_C is a propositional formula in 3-CNF, decide whether ψ is true. This problem is PSPACE-complete, see e.g. Garey and Johnson [35].

Let $\psi = Q_1 \hat{x}_1 : \dots : Q_k \hat{x}_k : \bigwedge_{i=1}^m (\ell_{i,1} \wedge \ell_{i,2} \wedge \ell_{i,3})$ with $\ell_{i,j} \in \{\hat{x}_l, \neg \hat{x}_l \mid 1 \leq l \leq k\}$. Choose $\mathbf{a} \in \Sigma$. The FC₍₁₎[\cdot]-formula that we construct shall represent each propositional variable \hat{x}_l with a variable x_l , and shall use $x_l = \mathbf{a}$ and $x_l = \varepsilon$ to represent $\hat{x}_l = 1$ and $\hat{x}_l = 0$, respectively.

We define a formula $\varphi(\mathfrak{s})() \in \text{FC}_{(1)}[\cdot]$ by $\varphi := \varphi_1^Q$, where

$$\varphi_i^Q := \begin{cases} \exists x_i : \varphi_{i+1}^Q & \text{if } Q_i = \exists, \\ \forall x_i : ((x_i \doteq \mathbf{a} \vee x_i \doteq \varepsilon) \rightarrow \varphi_{i+1}^Q) & \text{if } Q_i = \forall \end{cases}$$

for $1 \leq i \leq k$, and

$$\varphi_{k+1}^Q := \bigwedge_{j=1}^m (\varphi_{j,1}^\ell \vee \varphi_{j,2}^\ell \vee \varphi_{j,3}^\ell),$$

$$\varphi_{j,l}^\ell := \begin{cases} x_v \doteq \mathbf{a} & \text{if } \ell_{i,j} = \hat{x}_v \\ x_v \doteq \varepsilon & \text{if } \ell_{i,j} = \neg \hat{x}_v \end{cases}$$

for $1 \leq j \leq m$ and $l \in \{1, 2, 3\}$. Clearly, $\tau \models \varphi_{k+1}^Q$ if and only if τ encodes a satisfying assignment of the propositional formula. Moreover, the universal quantifiers $\forall x_i : \cdot$ are used such that the only interesting substitutions for x_i are those that map x_i to \mathbf{a} or to ε .

Let $\sigma(\mathfrak{s}) := \mathbf{a}$. Then $\sigma \models \varphi$ if and only if ψ is true. As φ can be constructed in polynomial time and QBF-3SAT is PSPACE-complete, this means that the recognition problem for Boolean FC₍₁₎[\cdot]-formulas is PSPACE-hard. Note that by choosing only existential quantifiers, this proof could also be used to as a reduction from 3SAT, which is an alternative way of obtaining the NP lower bound. ◀

B.2 Proof of Theorem 4.2

Proof. The proof is an extension of the bottom-up evaluation for the FO-case (see e.g. Theorem 4.24 in Flum and Grohe [28]). Let $\varphi \in \text{FC}[\cdot]$ and \vec{w} over Σ^* . For convenience, let $k := \text{wd}(\varphi)$ and $n := \|\vec{w}\|$. As every variable in φ must be mapped to a subword of some $w \in \vec{w}$, this means that we have $O(n^2)$ possible assignments for each variable.

For every word equation $x \doteq \alpha$, we know that $\text{wd}(x \doteq \alpha) \leq k$. This means that there are $O(n^{2k})$ different τ that could satisfy $\tau \models (x \doteq \alpha)$ and $\tau(\mathfrak{s}) = \sigma(\mathfrak{s})$ for all $\mathfrak{s} \in \text{struc}(\varphi)$. We can create a list of all these τ in time $O(n^{2k+1})$ by enumerating the $O(n^{2k})$ many possible choices and checking each choice in time $O(n)$.

We can lower the complexity to $O(n^{2k})$ by representing each assignment $\tau(x)$ as three pointers (l, i, j) , where l points a word $w \in \vec{w}$ such that $\tau(x) \sqsubseteq w$, and (i, j) determine the beginning and end of $\tau(x)$ in w . We then pre-compute a table of all pairs of such triples that determine the same word (like the Eq-relation for FO[Eq]). This table contains at most $O(n^2)$ entries and can be computed in time $O(n^3)$. By starting with $\tau(x)$, this table can then be used to check $\tau \models (x \doteq \alpha)$ in time $O(k)$, bringing the complexity of generating the list down to $O(kn^{2k})$.

Of course, this is still only a rough upper bound. For example, if a variable y is the first or last variable of α , there are only $O(n)$ possible assignments for y ; and if α starts or ends with terminals, this restricts the possible choices for x .

The lists of results can then be combined as in the relational case, requiring time $O(kn^{2k})$ in each inner node of the parse tree of φ . After computing all these sets, we check whether the list for the root is non-empty, and return the corresponding result. As the number of these is bounded by $|\varphi|$, we arrive at a total running time of $O(k|\varphi|n^{2k})$. ◀

B.3 Proof of Theorem 4.4

Proof. We first give a short summary of the definition of tree decompositions, treewidth, and nice tree decompositions (based on Chapter 7 of [16]). Readers who are familiar with these are invited to skip over to the actual construction.

Tree decompositions: A *tree decomposition* of a graph $G = (V, E)$ is a tree T with a function B that maps every node t of T to a subset of V such that:

1. for every $i \in V$, there is at least one node t of T such that $i \in B(t)$,
2. for every edge $(i, j) \in E$, there is at least one node t of T such that $i, j \in B(t)$,
3. for every $i \in V$, the set of nodes t of T such that $i \in B(t)$ induces a connected subtree of T .

The *width* of a tree decomposition (T, B) is the size of the largest $B(t)$ minus one. The *treewidth* $\text{tw}(G)$ of G is the minimal possible treewidth over all tree decompositions of G . A tree decomposition (T, B) of G is called *nice* if, firstly, T has a root r such that $B(r) = \emptyset$ and $B(l) = \emptyset$ for every leaf l of T , and secondly, every non-leaf node is of one of the following types:

- *introduce node (for i):* a node t with exactly one child t' such that $B(t) = B(t') \cup \{i\}$ with $i \notin B(t')$,
- *forget node (for i):* a node t with exactly one child t' such that $B(t) = B(t') - \{i\}$ with $i \in B(t')$,
- *join node:* a node t with exactly two children t_1 and t_2 such that $B(t) = B(t_1) = B(t_2)$.

Recall that if a graph has a tree decomposition of width at most k , it also has a nice tree decomposition of width at most k . Moreover, for every $i \in V$, there is exactly one forget node.

Construction: Given a pattern $\alpha = \alpha_1 \cdots \alpha_n$ with $n \geq 1$ such that $\text{tw}(\mathcal{G}_\alpha) = k$ for some $k \geq 1$, we first construct an EPFO $^{2k+2}$ [Eq]-formula φ with $\mathcal{L}(\varphi) = \mathcal{L}(\alpha)$. By Lemma 3.13, we then know that $\mathcal{L}(\alpha)$ is expressible in EPFC $^{2k+3}$ [].

The key idea of the construction is basically the same as in the proof Kolaitis and Vardi [42] for variable bounded FO (namely, read the formula directly from the tree decomposition). We only need to take some specifics of patterns into account.

We shall use two variables x_i^o and x_i^c for every $i \in [n]$, such that for every $w \in \mathcal{L}(\alpha)$, with the goal that $w_{[x_i^o, x_i^c]}$ describes the part of w that is generated by α_i (ensuring among other things that multiple occurrences of the same variable are mapped to the same word)

Let (T, B) be a nice tree decomposition (T, B) with width k . For every node t of T , we define a formula φ_t as follows:

1. If t is an introduce node with child t' , we define $\varphi_t := \varphi_{t'}$.
2. If t is a join node with children t_1 and t_2 , we define φ_t as the conjunction of all φ_{t_i} for which the subtree that is rooted at t_i contains at least one forget node.
3. If t is a forget node for $i \in [n]$, we define $\varphi_t := \exists x_i^o, x_i^c : \psi_i$, where ψ_i is a conjunction of the following formulas:
 - a. $(x_1^o \doteq \min)$ if $i = 1$,
 - b. $(x_n^c \doteq \max)$ if $i = n$,
 - c. $\text{succ}(x_i^o, x_i^c) \wedge P_a(x_i^o)$ if $\alpha_i = a$ for some $a \in \Sigma$,
 - d. $x_i^o \leq x_i^c$ if $\alpha_i \in \Xi$,
 - e. $\text{Eq}(x_i^o, x_i^c, x_j^o, x_j^c)$ if $i, j \in E_\alpha^{\text{eq}}$,
 - f. $x_{i-1}^c \doteq x_i^o$ if $(i-1) \in B(t)$,
 - g. $x_i^c \doteq x_{i+1}^o$ if $(i+1) \in B(t)$,
 - h. $\varphi_{t'}$ if the child t' of t contains at least one forget node.

We then define φ as φ_r for the root r of T . As the width of T is at most k , we have $|B(t)| \leq k+1$ for all nodes t of T . As every $i \in [n]$ is represented by two variables, we have $\text{wd}(\varphi) \leq 2k+2$.

As every $i \in [n]$ is forgotten exactly once, it is convenient to use forget nodes to process the variables (hence, they are the only nodes that actually generate interesting parts of the formula). Recall that each (x_i^o, x_i^c) is supposed to represent $\sigma(\alpha_i)$ for each $w = \sigma(\alpha) \in \mathcal{L}(\alpha)$. The cases 3a and 3b ensure that the left- and rightmost positions of α are mapped to the respective ends of w . For $\alpha_i = a$, case 3c ensures that α_i generates exactly this a , and that (x_i^o, x_i^c) describes exactly one letter. If α_i is a variable, the case 3d only states that x_i^o is not to the right of x_i^c . Other occurrences of this variable are handled by case 3e, which sets all occurrences of variables that are connected with E_α^{eq} have to generate the same image. Finally, case 3f and 3g ensure that neighboring parts share the same borders.

As $E_\alpha^< \subseteq E_\alpha$, we have that for every $i \in [n-1]$, there is some $B(t)$ that contains i and $i+1$. On the way from t to the root r , one of i and $i+1$ is forgotten first; this is the place where we ensure that $x_i^c = x_{i+1}^o$ holds. Analogously, $E_\alpha^{\text{eq}} \subseteq E_\alpha$ ensures that all occurrences of a variable have the same image. ◀

B.4 Definitions and results for Section 4.2

B.4.1 Fixed points

Our first step towards defining FC with fixed points is interpreting FC-formulas as functions that map relations on words to relations on words. To this end, we extend FC with a relation symbol \dot{R} that represents the input relation is. Unlike for constraints, the relation R for \dot{R} is not assumed to be fixed. Instead, we define the notion of a *generalized pattern substitution* σ , that also maps \dot{R} to a relation $\sigma(\dot{R}) \subseteq (\Sigma^*)^{\text{ar}(\dot{R})}$. For some structure variables \vec{s} and an $\text{ar}(\dot{R})$ -tuple of variables \vec{x} , we then have $\sigma \models_{\vec{s}} \dot{R}(\vec{x})$ if $\sigma(R) \subseteq (\text{Sub}(\sigma(\vec{s})))^{\text{ar}(\dot{R})}$ and

$\sigma(\vec{x}) \in \sigma(R)$. We call the formulas that are extended in this way $\text{FC}[\dot{R}]$ -formulas. This naturally generalizes to $\text{FC}[\{\dot{R}\} \cup \mathcal{K}]$ -formulas for every set of constraints \mathcal{K} .

► **Definition B.1.** For every word w , we use $\text{Sub}(w)$ to denote the set of all u with $u \sqsubseteq w$, and we extend this to tuples \vec{w} by $\text{Sub}(\vec{w}) = \bigcup_{w \in \vec{w}} \text{Sub}(w)$. Let $\varphi(\vec{s})$ be an $\text{FC}[\{\dot{R}\} \cup \mathcal{K}]$ -formula and $k := \text{ar}(\dot{R})$. For every tuple of words \vec{w} with $|\vec{w}| = |\vec{s}|$, and every k -tuple \vec{x} over $\text{free}(\varphi)$, we define the function from k -ary relations over $\text{Sub}(\vec{w})$ to k -ary relations over $\text{Sub}(\vec{w})$ by

$$F_{\vec{x}, \vec{w}}^\varphi(R) := \{\sigma(\vec{x}) \mid \sigma \models \varphi, \sigma(\dot{R}) = R, \sigma(\vec{s}) = \vec{w}\}$$

for every $R \subseteq \text{Sub}(\vec{w})^k$. We use this to define a sequence of relations by $R_0 := \emptyset$ and $R_{i+1} := F_{\vec{x}, \vec{w}}^\varphi(R^i)$ for all $i \geq 0$.

► **Example B.2.** Let $\text{ar}(\dot{R}) = 2$, and define the $\text{FC}[\dot{R}]$ -formula

$$\varphi(x, y) := (x \doteq \varepsilon \wedge y \doteq \varepsilon) \vee \exists \hat{x}, \hat{y}: \left(\bigvee_{a \in \Sigma} \bigvee_{b \in \Sigma} (x \doteq a \cdot \hat{x} \wedge y \doteq b \cdot \hat{y} \wedge \dot{R}(\hat{x}, \hat{y})) \right).$$

Using a straightforward induction, one can prove that for every $w \in \Sigma^*$, we have that $F_{(x,y), w}^\varphi$ defines a sequence of relations, where each R_i contains the pairs (u, v) where $u, v \in \text{Sub}(w)$ and $|u| = |v| < i$. In other words, for $i > |w|$, we have that R_i expresses the equal length relation on $\text{Sub}(w)$.

For every set A and every function $f: \mathcal{P}(A) \rightarrow \mathcal{P}(A)$, we say that $S \subseteq A$ is a *fixed point* of f if $f(S) = S$. A fixed point S of f is the *least fixed point* if $S \subseteq T$ holds for every fixed point T of f . We denote the least fixed point of f by $\text{lfp}(f)$. Using basic fixed point theory, see e. g. Ebbinghaus and Flum [24], we can prove the FC-version of a basic result for FO:

► **Lemma B.3.** Let $\varphi(\vec{s}) \in \text{EPFC}[\dot{R}]$, let $\vec{w} \in (\Sigma^*)^{|\vec{s}|}$, and let \vec{x} be a k -tuple over $\text{free}(\varphi)$. Then there exists $c \leq \|\vec{w}\|^{2k}$ such that $R_c = \text{lfp}(F_{\vec{x}, \vec{w}}^\varphi)$.

Proof. First, observe that $F_{\vec{x}, \vec{w}}^\varphi$ is a function $F_{\vec{x}, \vec{w}}^\varphi: \mathcal{P}(A) \rightarrow \mathcal{P}(A)$ for $A := (\text{Sub}(\vec{w}))^k$. Furthermore, note that A is a finite set with $|A| \leq \|\vec{w}\|^{2k}$.

To prove the claim, we use two further notions from fixed point theory: $F_{\vec{x}, \vec{w}}^\varphi$ is called *monotone* if $S \subseteq T$ implies $F_{\vec{x}, \vec{w}}^\varphi(S) \subseteq F_{\vec{x}, \vec{w}}^\varphi(T)$ for all $S, T \subseteq A$. It is *inductive* if $R_i \subseteq R_{i+1}$ for all $i \geq 0$.

As we are dealing with the existential-positive fragment of $\text{FC}[\dot{R}]$, the function $F_{\vec{x}, \vec{w}}^\varphi$ is monotone (this can be proven with a straightforward induction). But every monotone function from $\mathcal{P}(A)$ to $\mathcal{P}(A)$ is also inductive (see e. g. Lemma 8.1.2 in [24]). Hence, for $c := |A|$, the relation R_c is the least fixed point of $F_{\vec{x}, \vec{w}}^\varphi$ (this holds for every inductive function $\mathcal{P}(A) \rightarrow \mathcal{P}(A)$, see e. g. Lemma 8.1.1 in [24]). Hence, c is polynomial in $\|\vec{w}\|$. ◀

In other words, least fixed points for sequences of relations that are defined by FC-formulas behave in the same way as for FO-formulas. Accordingly, we can extend FC with least fixed points in the same way that FO can be extended with least fixed points:

► **Definition B.4.** Let \dot{R} be a relation symbol, $k := \text{ar}(\dot{R})$, and $\varphi(\vec{s}) \in \text{EPFC}[\{\dot{R}\} \cup \mathcal{K}]$. For all k -tuples \vec{x} and \vec{y} over $\Xi - \vec{s}$, we define $[\text{lfp } \vec{x}, \dot{R}: \varphi](\vec{y})$ as an LFP-formula that has structure variables \vec{s} and free variables $(\text{free}(\varphi) - \vec{x}) \cup \vec{y}$.

For every pattern substitution σ , we define $\sigma \models_{\vec{s}} [\text{lfp } \vec{x}, \dot{R}: \varphi](\vec{y})$ if there exists an extended pattern substitution τ with

1. $\tau \models \varphi$,

2. $\tau(\vec{x}) = \sigma(\vec{y})$,
3. $\tau(z) = \sigma(z)$ for all $z \in (\text{pub}(\varphi) - \vec{x})$, and
4. $\tau(\dot{R}) = \text{lfp}(F_{\vec{x}, \sigma(\vec{s})}^\varphi)$.

We generalize this multiple relation symbols and to nested fixed point operators, and we use $\text{FC}^{\text{lfp}}[\mathcal{K}]$ to denote the logic that is obtained by adding these LFP-formulas as base cases to the definition of $\text{FC}[\mathcal{K}]$. The union of all these $\text{FC}^{\text{lfp}}[\mathcal{K}]$ is denoted by FC^{lfp} .

Recall that $\text{pub}(\varphi) = \text{struc}(\varphi) \cup \text{free}(\varphi)$.

► **Example B.5.** Recall the formula $\varphi(x, y)$ from Example B.2 such that $\text{lfp}(F_{(x,y),w}^\varphi)$ is the equal length relation on every $w \in \Sigma^*$. We use this to define the $\text{FC}^{\text{lfp}}[\text{REG}]$ -formula

$$\psi(\vec{s})() := \exists x, y: (s \doteq xy \wedge x \in \mathbf{a}^* \wedge y \in \mathbf{b}^* \wedge [\text{lfp}(x, y), \dot{R}: \varphi]((x, y))),$$

which defines the language of all words $\mathbf{a}^n \mathbf{b}^n$ with $n \geq 0$.

Recall that we assume that all constraints in \mathcal{K} can be decided in polynomial time.

► **Lemma B.6.** *The data complexity of the recognition problem for FC^{lfp} is in P.*

Proof. Let $\varphi \in \text{FC}^{\text{lfp}}[\mathcal{K}]$ for some constraint set \mathcal{K} . We want to show that for every pattern substitution σ , we can decide in polynomial time whether $\sigma \models \varphi$. To do that, we extend the proof of Theorem 4.1 (see B.1) to include LFP-formulas.

To check whether $\sigma \models [\text{lfp } \vec{x}, \dot{R}: \varphi](\vec{y})$ for some $\varphi(\vec{s})(\vec{x})$, we compute $\text{lfp}(F_{\vec{x}, \sigma(\vec{s})}^\varphi)$. As shown in Lemma B.3, this is equivalent to computing $R_{|A|}$, for $A := (\text{Sub}(\vec{w}))^k$, where $k := |\vec{x}|$ and $\vec{w} := \sigma(\vec{s})$.

This can be done inductively by computing each R_{i+1} from R_i with $R_0 = \emptyset$. In each of these induction steps, we determine R_{i+1} by enumerating all extended substitutions τ that have $\tau(\dot{R}) = R_i$ and satisfy firstly, $\tau(\vec{y}) = \sigma(\vec{y})$ and, secondly, for every $x \in \vec{x}$, there is some $y \in \vec{y}$ with $\tau(x) \sqsubseteq \tau(y) = \sigma(y)$. For each such τ , we check whether $\tau \models \varphi$.

This check can be done in polynomial time, according to our induction assumption (relation predicates can be evaluated with a lookup if the relation has been computed, and constraints are assumed to be decidable in polynomial time). As $|\vec{x}| = |\text{free}(\varphi)|$, and as there are at most $\|\vec{w}\|^2$ different choices for $\tau(x)$, there are at most $\|\vec{w}\|^{2|\text{free}(\varphi)|}$ different τ . Hence, each level R_{i+1} can be computed using polynomially many checks that each take polynomial time.

We only need to compute polynomially many levels until reaching the least fixed point $R_{|A|}$. Hence, $\text{lfp}(F_{\vec{x}, \sigma(\vec{y})}^\varphi)$ can be computed in time that is polynomial in $\|\vec{w}\|$; and by the induction assumption, $\sigma \models [\text{lfp } \vec{x}, \dot{R}: \varphi](\vec{y})$ can then be decided in polynomial time.

Apart from that, the proof proceeds as for $\text{FC}[\mathcal{K}]$ in the proof Theorem 4.1; substituting P for L. ◀

The function $F_{\vec{x}, \vec{w}}^\varphi$ from Definition B.1 can also be used to define partial fixed points. We define the *partial fixed point* $\text{pfp}(F_{\vec{x}, \vec{w}}^\varphi)$ by $\text{pfp}(F_{\vec{x}, \vec{w}}^\varphi) := R_i$ if $R_i = R_{i+1}$ holds for some $i \geq 0$, and $\text{pfp}(F_{\vec{x}, \vec{w}}^\varphi) := \emptyset$ if $R_i \neq R_{i+1}$ holds for all $i \geq 0$.

We then define PFP-formulas $[\text{pfp } \vec{x}, \dot{R}: \varphi](\vec{y})$ analogously to LFP-formulas, the only difference being that φ can be any $\text{FC}[\mathcal{K}]$ -formula and is not restricted to the existential-positive fragment:

► **Definition B.7.** *Let \dot{R} be a relation symbol, $k := \text{ar}(\dot{R})$, and $\varphi(\vec{s}) \in \text{FC}[\{\dot{R}\} \cup \mathcal{K}]$. For all k -tuples \vec{x} and \vec{y} over $\Xi - \vec{s}$, we define $[\text{pfp } \vec{x}, \dot{R}: \varphi](\vec{y})$ as a PFP-formula that has structure variables \vec{s} and free variables $(\text{free}(\varphi) - \vec{x}) \cup \vec{y}$.*

For every pattern substitution σ , we define $\sigma \models_{\vec{s}} [\text{pfp } \vec{x}, \dot{R}: \varphi](\vec{y})$ if there exists an extended pattern substitution τ with

1. $\tau \models \varphi$,
2. $\tau(\vec{x}) = \sigma(\vec{y})$,
3. $\tau(z) = \sigma(z)$ for all $z \in (\text{pub}(\varphi) - \vec{x})$, and
4. $\tau(\dot{R}) = \text{pfp}(F_{\vec{x}, \sigma(\vec{s})}^{\varphi})$.

We generalize this multiple relation symbols and to nested fixed point operators, and we define $\text{FC}^{\text{pfp}}[\mathcal{K}]$ and FC^{pfp} analogously to Definition B.4.

► **Lemma B.8.** *The data complexity of the recognition problem for FC^{pfp} is in PSPACE.*

Proof. This proof proceeds similar to the one of Lemma B.6, the only difference is the bound on the number of R_i that need to be checked. To test if $\sigma \models [\text{pfp } \vec{x}, \dot{R}: \varphi](\vec{y})$, we need to compute $\text{pfp}(F_{\vec{x}, \sigma(\vec{s})}^{\varphi})$. As the underlying universe $\text{Sub}(\sigma(\vec{s}))$ is finite, we only need to enumerate up to $2^{|\sigma(\vec{s})|^{2k}}$ different R_i , where $k := \text{ar}(\dot{R})$. Moreover, we only need to keep each R_i in memory until R_{i+1} has been constructed; after that, R_i can be overwritten with R_{i+2} . Each current R_i can be represented as $\tau(\vec{x}) \in \text{Sub}(\sigma(\vec{s}))^k$, which means that the whole procedure can run in PSPACE. Apart from this, the proof proceeds as in Lemma B.6 (and then as in Theorem 4.1), but using PSPACE instead of P or L. ◀

We revisit FC^{lfp} and FC^{pfp} in Section B.5 for the proof of Theorem 4.6.

B.4.2 Transitive closures

For every relation $R \subseteq (\Sigma^*)^k \times (\Sigma^*)^k$ with $k \geq 1$, we define its *transitive closure* $\text{tc}(R)$ as the set of all $(r, \hat{r}) \subseteq (\Sigma^*)^k \times (\Sigma^*)^k$ for which there exists a sequence $r_1, \dots, r_n \in (\Sigma^*)^k$ with $n \geq 1$, $r_1 = r$, $r_n = \hat{r}$, and $(r_i, r_{i+1}) \in R$ for $1 \leq i < n$.

The *deterministic transitive closure* of R , written $\text{dtc}(R)$, is defined by adding the additional restriction that for every $1 \leq i < n$, there is no $(r_i, s) \in R$ with $s \neq r_{i+1}$.

► **Definition B.9.** *Let $\varphi \in \text{FC}[\mathcal{K}]$ and for $k \geq 1$, choose two k -tuples \vec{x} and \vec{y} over $\text{free}(\varphi)$, and two k -tuples \vec{s} and \vec{t} over $\Xi - \text{struc}(\varphi)$. Then $[\text{tc } \vec{x}, \vec{y}: \varphi](\vec{s}, \vec{t})$ is a TC-formula and $[\text{dtc } \vec{x}, \vec{y}: \varphi](\vec{s}, \vec{t})$ is a DTC-formula. Both have the same structure variables as φ , and as free variables the set $(\text{free}(\varphi) - (\vec{x} \cup \vec{y})) \cup (\vec{s} \cup \vec{t})$.*

For every pattern substitution σ , we define $\sigma \models [\text{tc } \vec{x}, \vec{y}: \varphi](\vec{s}, \vec{t})$ if $(\sigma(\vec{s}), \sigma(\vec{t})) \in \text{tc}(R_\sigma)$, where R_σ is the set of all $(\tau(\vec{x}), \tau(\vec{y}))$ such that

1. $\tau \models \varphi$, and
2. $\tau(z) = \sigma(z)$ for all $z \in \text{free}(\varphi) - (\vec{x} \cup \vec{y})$.

The analogous definition applies to DTC-formulas, substituting $\text{dtc}(R_\sigma)$ for $\text{tc}(R)_\sigma$.

We generalize this to multiple and nested applications of the closure operators; and we use $\text{FC}^{\text{tc}}[\mathcal{K}]$ or $\text{FC}^{\text{dtc}}[\mathcal{K}]$ to denote the logics that are obtained by adding these TC- or DTC-formulas as base cases to the definition of $\text{FC}[\mathcal{K}]$. FC^{tc} and FC^{dtc} are defined analogously.

As we do not require that \vec{x} and \vec{s} (or \vec{y} and \vec{t}) are distinct, we use $[\text{tc } \vec{x}, \vec{y}: \varphi]$ as shorthand for $[\text{tc } \vec{x}, \vec{y}: \varphi](\vec{x}, \vec{y})$. We now consider some examples.

► **Example B.10.** We define the $\text{EPFC}^{\text{dtc}}[\]$ -formula

$$\varphi(\mathfrak{s})() := \exists x, y: \left((\mathfrak{s} \doteq y) \wedge [\text{dtc } x, y: \psi] \wedge (x \doteq \varepsilon \vee \bigvee_{a \in \Sigma} y \doteq a) \right),$$

$$\psi(x, y) := \bigvee_{a \in \Sigma} (x \doteq a \cdot y \cdot a).$$

Then $\sigma \models \varphi$ if and only if $\sigma(w)$ is a palindrome over Σ . The formula φ expresses that x can be obtained from y by concatenating some one occurrence of some letter a to the left and one to the right of y . By applying the transitive closure, we obtain the relation of all (x, y) such that $x = u \cdot y \cdot u^R$, where $u \in \Sigma^*$ and u^R is the reversal of u .

Note that ψ selects the relation of all (x, y) with $x = aya$ for some $a \in \Sigma$. Hence, each word has exactly one successor in this relation, which means that we can indeed use $[\text{dtc } x, y : \psi]$. But if we wrote $[\text{dte } y, x : \psi]$ instead, then there could be multiple successors for some x (depending on the content of \mathfrak{s}), which means that dte would fail.

► **Example B.11.** Consider a directed graph $G = (V, E)$ with $V = \{v_1, \dots, v_n\} \subseteq \{0, 1\}^+$ and $n \geq 1$. Define $\text{enc}(E)$ as an encoding of E over $\{0, 1, \#, \$\}$ such that $\text{enc}(E)$ contains the subword $\$v_i\#v_j\$$ if and only if $(v_i, v_j) \in E$. We define the $\text{EPFC}^{\text{tc}}[\]$ -formula

$$\varphi(x, y) := [\text{tc } x, y : \exists z : (z \doteq \$x\#y\$ \wedge x \dot{\in} \{0, 1\}^+ \wedge y \dot{\in} \{0, 1\}^+)].$$

Then $\llbracket \varphi \rrbracket(\text{enc}(E))$ is the set of all (v_i, v_j) such that v_j can be reached from v_i in one or more steps.

Next, we examine the data complexity of model-checking FC^{tc} and FC^{dte} . Recall that we assume that constraints can be evaluated in L .

► **Lemma B.12.** *The data complexity of the recognition problem is in NL for FC^{tc} and in L for FC^{dte} .*

Proof. Again, we extend the proof of Theorem 4.1 (see B.1) by describing how we evaluate DTC- and TC-formulas. Although some modifications are required in our setting, the basic idea is the same as for FO-formulas with dte - or tc -operators (see, e. g., Theorem 7.4.1. in [24]). Given σ and a TC-formula $[\text{tc } \vec{x}, \vec{y} : \varphi](\vec{s}, \vec{t})$ (or a DTC-formula like this), first note that on any given structure \vec{w} , the underlying universe can have up to n^2 elements for $n := \|\vec{w}\|$. This means that the closures can create paths up to length n^{2k} .

We then construct a logspace-Turing machine M_0 for φ that will be used as a sub-routine. Using one counter per main variable, we can implement a counter from 1 to n . Combining $2k$ of these, we can create a counter that counts up to n^{2k} . Now we can progress as in the relational case: we invoke M as a subroutine at most n^{2k} times to checking whether there is a path from $\sigma(\vec{x})$ to $\sigma(\vec{y})$. For tc , this involves guessing the next step; for dte , it involves checking that the successor is unique. Hence, all this can be done in NL and L , respectively.

We then integrate this into larger formulas via Theorem 4.1; using the fact that L and NL are both closed under complement. Of course, we could reprove this for NL by imitating the proof of the Immerman-Szelepcsényi theorem by means of FC^{dte} ; but this would not provide us with any new insights. ◀

B.5 Proof of Theorem 4.6

Proof. We want to show that a language is definable in a logic \mathcal{F} if and only if it belongs to the complexity class \mathbb{C} , where \mathcal{F} ranges over FC^{dte} , FC^{tc} , FC^{lfp} , or FC^{pfp} , and \mathbb{C} over L , NL , P , and PSPACE , respectively.

We have already established the direction from the logics to the complexity classes, namely in Lemma B.12 for FC^{dte} and FC^{tc} , in Lemma B.6 for FC^{lfp} , and in Lemma B.8 for FC^{pfp} . These results rely on Theorem 4.1 and, thus, on Lemma 3.15 (which allows us to convert FC -formulas into FO-formulas).

For the other direction, one might ask whether it is possible to use Lemma 3.13 (the other direction of Theorem 3.16). In particular, we have that for each of the extensions of FC, the correspondingly extended version of $\text{FO}[\prec]$ captures the complexity class.

But we have the additional goal of showing that an $\text{EPFC}[\prec]$ -formula is enough; and just applying Lemma 3.13 to the proofs that the authors found in literature would have required considerable hand-waving.

Capturing L and NL with dtc and tc : As explained by e.g. Kozen [43] (Lecture 5), a language L is in L (or in NL) if and only if there is some $k \geq 1$ such that L is accepted by a deterministic (or on-deterministic) finite automaton A that has k -many two-way input heads that are read-only and cannot move beyond the input. We assume without loss of generality that A does not read the left end-marker (this can be realized in the finite control).

Let n denote the number of states of A . We assume that the state set is $[n]$, that the starting state is 1 and that the accepting state is n . Given such an automaton A , our goal is to construct a Boolean formula φ such that $\mathcal{L}(\varphi) = \mathcal{L}(A) \cap \Sigma^{\geq n}$. The finitely missing words can then be added with a straightforward disjunction.

In the construction, the structure variable \mathfrak{s} represents the input w of A . Each head number $i \in [k]$ is modeled by a variable x_i , where its current position $j \in [|w|]$ is represented as $w_{[1,j]}$ (that is, the prefix of w that has length $j - 1$). Likewise, the current state $q \in [n]$ is represented by $w_{[1,q]}$.

Our goal is to define a formula ψ that encodes the successor relation R for A . Using dtc or tc , we can then use to simulate the behavior of A on w . To this end, we define two types of helper formulas. Firstly, for $q \in [n]$, we define a formula $\psi_q^Q(x)$ that expresses “ x represents state q ” by having $\sigma \models \psi_q^Q(x)$ if and only if $\sigma(x) = \sigma(\mathfrak{s})_{[1,q]}$. Let $\psi_1^Q(x) := (x \doteq \varepsilon)$ and

$$\psi_{q+1}^Q(x) := \exists \hat{x}, z: \bigvee_{a \in \Sigma} (x \doteq \hat{x}a \wedge \mathfrak{s} \doteq xz \wedge \psi_q^Q(z))$$

for all $1 \leq q < n$. Next, for each $a \in \Sigma$, we define

$$\psi_a^{\text{read}}(x) := \exists z: \mathfrak{s} \doteq xaz,$$

which expresses that “the letter after the prefix x is a ”. We also define

$$\psi_{\dashv}^{\text{read}}(x) := \exists z: \mathfrak{s} \doteq x$$

We shall use these two types of formulas to check the content of the input heads i (namely, whether head i reads $a \in \Sigma$ or the right end marker \dashv). Finally, we define

$$\psi_{\text{succ}}(x, y) := \bigvee_{a \in \Sigma} y \doteq xa$$

to express that “ y is one letter longer than x ”, which we shall use for the head movements. Now we are ready to put the pieces together. For $\vec{a} = (a_1, \dots, a_k) \in (\Sigma \cup \{\dashv\})^k$ and $q \in [n]$, define

$$\psi_q^{\vec{a}} := \psi_q^Q(x_0) \wedge \bigwedge_{i=1}^k \psi_{a_i}^{\text{read}}(x_i) \wedge \psi_{q, \vec{a}, i}^{\text{mov}}(x_i, y_i),$$

where the head movements are simulated by

$$\psi_{q, \vec{a}, i}^{\text{mov}}(x_i, y_i) := \begin{cases} \psi_{\text{succ}}(x_i, y_i) & \text{if } A, \text{ when reading } \vec{a} \text{ in state } q, \text{ moves head } i \text{ to the right,} \\ \psi_{\text{succ}}(y_i, x_i) & \text{if } A, \text{ when reading } \vec{a} \text{ in state } q, \text{ moves head } i \text{ to the left,} \\ x_i \doteq y_i & \text{if } A, \text{ when reading } \vec{a} \text{ in state } q, \text{ does not move head } i. \end{cases}$$

This gives us the immediate successor for each combination \vec{a} of input letters (including the end-marker) and each state q . To get all possible successors, we combine these into

$$\psi(\vec{x}, \vec{y}) := \bigvee_{q \in [n]} \bigvee_{\vec{a} \in (\Sigma \cup \{-\})^k} \psi_q^{\vec{a}}(\vec{x}, \vec{y}),$$

where $\vec{x} = (x_0, \dots, x_k)$ and $\vec{y} = (y_0, \dots, y_k)$. We now define

$$\begin{aligned} \varphi(\mathfrak{s})() &:= \exists \vec{x}, \vec{y}: \left(\bigwedge_{i=0}^k (x_i \doteq \varepsilon) \wedge \psi_n^Q(y_0) \wedge [\text{dte } \vec{x}, \vec{y}: \psi] \right) && \text{if } A \text{ is deterministic, and} \\ \varphi(\mathfrak{s})() &:= \exists \vec{x}, \vec{y}: \left(\bigwedge_{i=0}^k (x_i \doteq \varepsilon) \wedge \psi_n^Q(y_0) \wedge [\text{tc } \vec{x}, \vec{y}: \psi] \right) && \text{if } A \text{ is non-deterministic.} \end{aligned}$$

Outside the closure operators, the formula expresses that all \vec{x} encodes the initial position (x_0 encodes the starting state 1 and all tapes are at the very left), and that \vec{y} encodes a halting position. If A is non-deterministic, then we can use the tc -operator to obtain the transitive closure of the successor relation on the configurations of A on w . If A is deterministic, then every configuration has at most one successor, meaning that the dte -operator also compute the transitive closure, having the same intended effect.

Thus, for all $w \in \Sigma^*$, we have $\llbracket \varphi \rrbracket(w) \neq \emptyset$ if and only if $w \in \mathcal{L}(A)$ and $|w| \geq n$. As mentioned above, the “missing words” from the set $W := \mathcal{L}(A) - \mathcal{L}(\varphi)$ can be added now by defining a formula $\varphi \vee \bigvee_{w \in W} \mathfrak{s} \doteq w$.

We conclude that $\text{FC}^{\text{tc}}[\]$ captures NL and that $\text{FC}^{\text{dte}}[\]$ captures L. Moreover, note that we used only a single closure operator, and that the formulas are existential-positive (inside and outside of the closure operator).

Capturing P with lfp: Hence, we give an outline of the full proof, which takes key-ideas from the proof of Theorem 7.3.4 in [24]. Again, the main challenge is ensuring that the formula is existential-positive.

For every language $L \in \text{P}$, there is a Turing machine M that decides L in polynomial time. We assume that M has one read-only input tape over Σ and a read-write work tape that extends to the right and has a tape alphabet $\Gamma = \{0, \dots, m\}$ for some $m \geq 1$. For the state set Q , we assume $Q = \{0, \dots, n\}$ for some $n \geq 1$, where 0 is the initial and n the single accepting state. When starting, each head is on the left of its tape (position 0), the machine is in state 1, and each cell of the work tape contains 0.

As M decides L in polynomial time, there is a natural number d such that on each input $w \in \Sigma$, we have that M terminates after at most $|w|^d$ steps. During this run, M will not visit more than $|w|^d$ tape positions.

For each i with $0 \leq i \leq |w|$, let w_i be the prefix of w that has length i . For $k \geq 1$, we identify each k -tuple $\vec{v} = (v_1, \dots, v_k)$ with the number $N(\vec{v}) := \sum_{i=1}^k (|v_i| |w|^{i-1})$. Hence, we can use two d -tuples of variables, a tuple $\vec{t} = (t_1, \dots, t_d)$ that to encode time stamps and a tuple $\vec{p} = (p_1, \dots, p_d)$ that encodes positions on work tape (where 0 is the leftmost position). The construction will ensure that both tuples will only take on prefixes of w as values.

Our simulation of M will be able to run for $(|w| + 1)^d - 1$ steps; but this does not affect the outcome.

Our goal is to define a relation R that starts with the initial configuration of M on w , and then uses the lfp -operator to compute each successor configuration. As the time (and, hence, the space) of M are bounded, this is enough.

The relation R with arity $2d + 2$ shall represent the configuration of M on w in step $N(\vec{t})$ as follows:

- $(\vec{t}, w_0, w_q, \vec{\varepsilon})$ to denote that M is in state $q \in Q$,
 - $(\vec{t}, w_1, w_{i-1}, \vec{\varepsilon})$ to denote that the input head is at position i with $0 \leq i < |w|$,
 - $(\vec{t}, w_2, \varepsilon, \vec{p})$ to denote that the working head is at position $N(\vec{p})$,
 - $(\vec{t}, w_3, w_\gamma, \vec{p})$ to denote that the working tape contains $\gamma \in \Gamma$ at position i ,
- where $\vec{\varepsilon}$ is shorthand for the d -tuple that has ε on all positions.

Like in the case for transitive closures, the constructed formula will only be correct for w of sufficient length; but again, the finitely many exceptions can be added later. In particular, we assume that $|w| \geq c$ for $c := \max\{3, |\Gamma|, |Q|\}$. The only w_i that we refer to explicitly through their number i are w_0 to w_3 for the first component of R , w_q with $q \in Q$, and w_γ with $\gamma \in \Gamma$. For each one of these, $|w| \geq c$ guarantees that the input w is large enough to encode them.

As we encode various things in these prefixes w_i , it is helpful to define a successor relation

$$\psi_{succ}(x, y) := \bigvee_{a \in \Sigma} (y \dot{=} xa) \wedge \exists z: (\mathfrak{s} \dot{=} yz)$$

which expresses that x and y are prefixes of \mathfrak{s} , and y is one letter longer than x . Our first use for tgus is in the shorthand formulas $\psi_i^{pre}(\mathfrak{s})(x)$ for $0 \leq i \leq c$ to express the prefixes w_i . Let $\psi_0^{pre}(x) := (x \dot{=} \varepsilon)$ and $\psi_{i+1}^{pre}(x) := \exists y: \psi_{succ}(y, x)$ for $0 \leq i < c$.

We are now ready to define the formula $\psi_{init}(\mathfrak{s})(x_1, \dots, x_{2d+2})$, which expresses the initial configuration:

$$\begin{aligned} \psi_{init}(\mathfrak{s})(x_1, \dots, x_{2d+2}) := & \\ & \left(\bigwedge_{i=1}^d x_i \dot{=} \varepsilon \wedge (\psi_0^{pre}(x_{d+1}) \vee \psi_1^{pre}(x_{d+1}) \vee \psi_2^{pre}(x_{d+1})) \wedge \bigwedge_{i=d+2}^{2d+2} x_i \dot{=} \varepsilon \right) \\ & \vee \left(\bigwedge_{i=1}^d x_i \dot{=} \varepsilon \wedge \psi_3^{pre}(x_{d+1}) \wedge (x_{d+2} \dot{=} \varepsilon) \wedge \bigwedge_{j=d+3}^{2d+2} \exists y: \mathfrak{s} \dot{=} x_j y \right). \end{aligned}$$

The first part of the formula ensures that the machine starts in state 0, that each head is at position 0. The second part ensures that all tape cells are set to the blank symbol 0. Note that the tape position \vec{p} is stored in the last d components of the tuple (i. e., from $d+3$ to $2d+2$). To get all possible \vec{p} , these variables can take on any prefix w_i of w .

To describe the successor of a time stamp or the movement of the working head, we also define a relation $\psi_d^{succ}(\vec{x}, \vec{y})$ for d -tuples \vec{x} and \vec{y} such that $\sigma \models \psi_d^{succ}$ if and only if every component of $\sigma(\vec{x})$ and $\sigma(\vec{y})$ is a prefix of w , and $N(\sigma(\vec{y})) = N(\sigma(\vec{x})) + 1$. The basic idea is as for ψ_{succ} , but extending it to multiple digits by taking into account all cases where a carry might happen (this is straightforward, but rather tedious). Using this idea and the proper prefix relation \sqsubset_p from Example 3.8, we also construct an existential-positive formula $\psi_d^{\neq}(\vec{x}, \vec{y})$ that expresses $N(\sigma(\vec{x})) \neq N(\sigma(\vec{y}))$, similar to how we expressed \neq in that example.

We also define formulas $\psi_a(x) := \exists y: \mathfrak{s} \dot{=} x a y$ for every $a \in \Sigma$. If x represents the position of the input head, $\psi_a(x)$ expresses that this head is reading the letter a .

This is now all that we need to describe the successor relation R on configurations of M . We define an LFP-formula

$$\psi := [\text{lfp } \vec{x}, \dot{R}: (\psi_{init} \vee \psi_{next})](\vec{x}),$$

where $\vec{x} := (x_1, \dots, x_{2d+2})$ and the EPFC[\dot{R}]-formula ψ_{next} is constructed as follows:

- Using existential quantifiers, we retrieve a time stamp \vec{t} from R , and the uniquely defined state q , input head position i , working head position \vec{p} , and working head content γ for \vec{p} for this time stamp \vec{t} .

- If $\vec{t} = w^d$, nothing needs to be done. Hence, we can assume that this is not the case.
- As M is a deterministic Turing machine, the combination of state, current input symbol, and current tape symbol uniquely determine a combination of head movements and working tape action. Which of these applies can be determined by a big disjunction over all combinations of applying ψ^{pre} to the state and the working tape symbol, and ψ_a to the input symbol. For each of these cases, we create a sub-formula that describes head movements and the tape action in the time stamp \vec{t} with $N(\vec{t}) = N(\vec{t}) + 1$. We shall store \vec{t} in the free variables x_1 to x_d of ψ_{next} .
- The sub-formula then has a disjunction over the four possible choices for x_{d+1} (namely, for $\psi_1^{pre}(x_{d+1})$ to $\psi_4^{pre}(x_{d+1})$).
- For ψ_1^{pre} , the next state, we simply ensure that the correct successor state is stored in x_{d+2} , and set all remaining variables to ε .
- For ψ_2^{pre} , the input head position, we use ψ_{succ} to pick position $i + 1$ or $i - 1$ if the head moves, or just use the same position.
- For ψ_3^{pre} , the working head position, we use ψ_d^{succ} analogously.
- For ψ_4^{pre} , the working tape contents, we distinguish whether the cell is affected by the tape operation or not; that is, whether the cell is at position \vec{p} or not. If it is at that position, we return the new cell content. If not (which can be tested with ψ_d^{\neq}), we retrieve the cell content for time stamp \vec{t} from \vec{R} using existential quantifiers and return it unchanged.

Now, ψ computes the relation of all encodings of configurations that M reaches on input w . All that remains is checking for the existence of an accepting configuration. We define

$$\varphi := \exists \vec{x}: (\psi_0^{pre}(x_{d+1}) \wedge \psi_n^{pre}(x_{d+2}) \wedge \psi(\vec{x}))$$

for $\vec{x} = (x_1, \dots, x_{2d+2})$. Then we have $\llbracket \varphi \rrbracket(w) = \{()\}$ if and only if $\llbracket \psi \rrbracket(w)$ contains the encoding of a configuration that reaches the accepting state n . Hence, $\mathcal{L}(\varphi) = L$.

Capturing PSPACE with lfp: We can show this by modifying the lfp-construction: Instead of using time stamps, each stage of the relation in \vec{R} only keeps the most recent configuration and uses it to construct the next. As L is decidable in PSPACE, this can be done using the tuple \vec{p} . As we have already established that the lfp-construction is possible with an existential-positive formula, this modification is straightforward. ◀

B.6 Proof of Theorem 4.8

Proof. We can directly rewrite every $\text{DataSpLog}[\mathcal{K}]$ -program into an equivalent $\text{EPFC}^{\text{lfp}}[\mathcal{K}]$ -formula. By Theorem 4.6, these are in P.

For the other direction, we know from the proof of the lfp-case of Theorem 4.6 that every language in P is recognized by a formula from $\text{EPFC}^{\text{lfp}}[\]$ that consist of existential quantifiers over a single lfp-operator. After transforming the underlying formula into a union of conjunctive queries (using the same rules as for relational logic), we immediately obtain an equivalent $\text{DataSpLog}[\]$ -program. ◀

B.7 Proof of Lemma 4.9

We start with some preliminaries; the actual proof is in Section B.7.2

B.7.1 Decomposing the structure

In the proof, we use \sqcup to denote the union of two disjoint sets. The key part of the argument is the following formulation of the Feferman-Vaught theorem:

► **Feferman-Vaught theorem** (Theorem 1.6 in [51]). *For every $q \in \mathbb{N}$ and for every first order formula φ of quantifier rank q over a finite vocabulary, one can compute effectively a reduction sequence*

$$\psi_1^A, \dots, \psi_k^A, \psi_1^B, \dots, \psi_k^B$$

of first order formulas the same vocabulary and a Boolean function $B_\varphi : \{0, 1\}^{2k} \rightarrow \{0, 1\}$ such that

$$\mathcal{A} \sqcup \mathcal{B} \models \varphi$$

if and only if $B_\varphi(b_1^A, \dots, b_k^A, b_1^B, \dots, b_k^B) = 1$ where $b_j^A = 1$ iff. $\mathcal{A} \models \psi_j^A$ and $b_j^B = 1$ iff. $\mathcal{B} \models \psi_j^B$.

This proof uses FO[Eq]-formulas instead of FC-formulas (due to Theorem 3.16). Intuitively, we show that any formula φ and structure \mathcal{A}'_w for some word $w \in \mathbf{a}^* \mathbf{b}^*$ can be translated into a formula ψ that operates on the union of two disjoint structures $\tilde{\mathcal{A}}_w^{\mathbf{a}}$ and $\tilde{\mathcal{A}}_w^{\mathbf{b}}$ such that $\mathcal{A}'_w \models \varphi$ if and only if $\tilde{\mathcal{A}}_w^{\mathbf{a}} \sqcup \tilde{\mathcal{A}}_w^{\mathbf{b}}$ satisfies ψ . We then apply the Feferman-Vaught theorem on ψ and obtain some kind of separation of it. Finally, we use the pigeonhole principle to compose a word that it is outside of the language.

Formally, let $\varphi' \in \text{FO}[\text{Eq}]$. Recall that φ' is evaluated on the structures \mathcal{A}'_w for $w \in \Sigma^*$, with the universe $\{1, \dots, |w| + 1\}$, where the node $|w| + 1$ is not marked with any letter. Also recall that the vocabulary of FO[Eq] contains the two unary letter predicates $P_{\mathbf{a}}$ and $P_{\mathbf{b}}$, the binary relations $<$ and succ , the 4-ary relation Eq, and the constant symbols \min and \max .

To apply the Feferman-Vaught theorem, we need to split \mathcal{A}'_w into two disjoint structures, and to rewrite φ into a suitable formula ψ . In this case, “suitable” means that φ and ψ are equivalent on words of the form $w = \mathbf{a}^m \mathbf{b}^n$ with $m, n \geq 1$. On these words, \mathcal{A}'_w contains nodes $1, \dots, m$ that are marked \mathbf{a} , nodes $m + 1, \dots, m + n$ that are marked \mathbf{b} , and the unmarked node $m + n + 1$. Our goal is to split all non-unary relations in \mathcal{A}'_w into a structure $\tilde{\mathcal{A}}_w^{\mathbf{a}}$ for the \mathbf{a} -part and structure $\tilde{\mathcal{A}}_w^{\mathbf{b}}$ for the \mathbf{b} -part. The only technical issue that we need to deal with is that the Eq-relation contains tuples (i_1, j_1, i_2, j_2) with $w_{[i_1, j_1]} = w_{[i_2, j_2]}$. In these tuples, j_1 is the first position that is *not* in $w_{[i_1, j_1]}$, and likewise for j_2 . While this relation is more convenient when expressing spanners and converting from and to FC-formulas, it creates issues when splitting the structures. Hence, we first define a 4-ary relation NEEq that contains those (i_1, j_1, i_2, j_2) with $i_1 \leq j_1$ and $i_2 \leq j_2$ that have $w_{i_1} \dots w_{j_1} = w_{i_2} \dots w_{j_2}$. Hence, unlike Eq, this relation only describes the equality of non-empty words.

We now can directly split the universe into $\{1, \dots, m\}$ and $\{m + 1, \dots, m + n + 1\}$, and define

- $\tilde{\mathcal{A}}_w^{\mathbf{a}}$ over the universe $\{1, \dots, m\}$,
- $\tilde{\mathcal{A}}_w^{\mathbf{b}}$ over the universe $\{m + 1, \dots, m + n + 1\}$,

where each structure $\tilde{\mathcal{A}}_w^c$ with $c \in \{\mathbf{a}, \mathbf{b}\}$ has the relation P_c and the relations $<_c$, succ_c , and NEEq_c , which are restrictions of the corresponding relations in \mathcal{A}'_w or of NEEq to the universe of $\tilde{\mathcal{A}}_w^c$. Likewise, we use \min_c and \max_c , where $\min_{\mathbf{a}}$ and $\max_{\mathbf{a}}$ refer to 1 and m , and $\min_{\mathbf{b}}$ and $\max_{\mathbf{b}}$ refer to $m + 1$ and $m + n + 1$, respectively. We use FO[NEEq] to denote the set of all formulas over this modified vocabulary and observe the following:

► **Lemma B.13.** *For every $\varphi \in \text{FO}[\text{Eq}]$, we can construct $\psi \in \text{FO}[\text{NEEq}]$ such that for every $w = \mathbf{a}^m \mathbf{b}^n$ with $m, n \geq 1$, we have $\mathcal{A}'_w \models \varphi$ if and only if $\tilde{\mathcal{A}}_w^{\mathbf{a}} \sqcup \tilde{\mathcal{A}}_w^{\mathbf{b}} \models \psi$.*

Proof. We obtain ψ from φ by rewriting all parts that use the constant symbols \min and \max or any of the relation symbols $<$, succ , or Eq . As we are only interested in words of the form $\mathbf{a}^m \mathbf{b}^n$ with $m, n \geq 1$, we can replace \min with $\min_{\mathbf{a}}$ and \max with $\max_{\mathbf{b}}$. We now replace the relations as follows:

- Every occurrence of $x < y$ is replaced with the formula

$$(x <_{\mathbf{a}} y) \vee (x <_{\mathbf{b}} y) \vee \left((x <_{\mathbf{a}} \max_{\mathbf{a}} \vee x \doteq \max_{\mathbf{a}}) \wedge (\min_{\mathbf{b}} <_{\mathbf{b}} y \vee \min_{\mathbf{b}} \doteq y) \right),$$

which covers the cases that both variables are in the \mathbf{a} -part, both are in the \mathbf{b} -part, or the remaining case that x is in the \mathbf{a} -part and y in the \mathbf{b} -part.

- Every occurrence of $\text{succ}(x, y)$ is replaced with

$$\text{succ}_{\mathbf{a}}(x, y) \vee \text{succ}_{\mathbf{b}}(x, y) \vee (x \doteq \max_{\mathbf{a}} \wedge y \doteq \min_{\mathbf{b}}),$$

where first two cases have both variables in the same part (as above), and the last describes that x is the last \mathbf{a} and y the first \mathbf{b} .

- To simplify the explanation of the last case, we describe it in two steps. Every $\text{Eq}(x_1, y_1, x_2, y_2)$ is first replaced with

$$\begin{aligned} (x_1 \doteq y_1 \wedge x_2 \doteq y_2) \vee \exists z_1, z_2: & (\text{succ}(z_1, y_1) \wedge \text{succ}(z_2, y_2) \wedge \text{NEEq}_{\mathbf{a}}(x_1, z_1, x_2, z_2)) \\ & \vee \exists z_1, z_2: (\text{succ}_{\mathbf{b}}(z_1, y_1) \wedge \text{succ}_{\mathbf{b}}(z_2, y_2) \wedge \text{NEEq}_{\mathbf{b}}(x_1, z_1, x_2, z_2)) \end{aligned}$$

The first conjunct describes the case where we have two occurrences of the empty word (which can be in any part of w). In all other cases, equal words must both be in the \mathbf{a} -part or the \mathbf{b} -part, which means that they are covered by the respective NEEq_c . As these relations bound words with their last position (unlike Eq), we use the z_i to obtain these positions. After this, we replace each of the two $\text{succ}(z_i, y_i)$ with

$$\text{succ}_{\mathbf{a}}(z_i, y_i) \vee (z_i \doteq \max_{\mathbf{a}} \wedge y_i \doteq \min_{\mathbf{b}}),$$

to account for cases where x_i is in the \mathbf{a} -part and y_i in the \mathbf{b} -part.

Cases where these subformulas involve constants are handled accordingly. On words from $\mathbf{a}^+ \mathbf{b}^+$, each of the new subformulas acts exactly like the one it replaces. Hence, for every $w = \mathbf{a}^m \mathbf{b}^n$ with $m, n \geq 1$, we have $\mathcal{A}'_w \models \varphi$ if and only if $\tilde{\mathcal{A}}_w^{\mathbf{a}} \sqcup \tilde{\mathcal{A}}_w^{\mathbf{b}} \models \psi$. ◀

B.7.2 Actual proof of Lemma 4.9

Proof. Assume that there is some $\varphi \in \text{FO}[\text{Eq}]$ with $\mathcal{L}(\varphi) = \{\mathbf{a}^n \mathbf{b}^n \mid n \geq 1\}$ (by Theorem 3.16, this is the same as assuming that this is an $\text{FC}[\]$ -language). We apply Lemma B.13 to φ and obtain a formula $\psi \in \text{FO}[\text{NEEq}]$ such that for all $w = \mathbf{a}^m \mathbf{b}^n$, we have $\tilde{\mathcal{A}}_w^{\mathbf{a}} \sqcup \tilde{\mathcal{A}}_w^{\mathbf{b}} \models \psi$ if and only if $m = n$. We now invoke the Feferman-Vaught theorem for ψ and obtain a sequence of first order formulas

$$\psi_1^A, \dots, \psi_k^A, \psi_1^B, \dots, \psi_k^B$$

and a Boolean function $B_\varphi : \{0, 1\}^{2k} \rightarrow \{0, 1\}$ such that for any word $w \in \{\mathbf{a}^m \mathbf{b}^n \mid m, n \geq 1\}$, we have $\tilde{\mathcal{A}}_w^{\mathbf{a}} \sqcup \tilde{\mathcal{A}}_w^{\mathbf{b}} \models \psi$ if and only if $B_\varphi(b_1^A, \dots, b_k^A, b_1^B, \dots, b_k^B) = 1$; where $b_j^A = 1$ iff $\tilde{\mathcal{A}}_w^{\mathbf{a}} \models \psi_j^A$ and $b_j^B = 1$ iff $\tilde{\mathcal{A}}_w^{\mathbf{b}} \models \psi_j^B$.

Since $\{\mathbf{a}^n \mathbf{b}^n \mid n \geq 1\}$ contains infinitely many words, we can use the pigeonhole principle to conclude that there exists $m \neq n$ such that for $w_m := \mathbf{a}^m \mathbf{b}^m$ and $w_n := \mathbf{a}^n \mathbf{b}^n$, we have for every j

- $\tilde{\mathcal{A}}_{w_m, \mathbf{a}} \models \psi_j^A$ iff $\tilde{\mathcal{A}}_{w_n, \mathbf{a}} \models \psi_j^A$, and

■ $\tilde{\mathcal{A}}_{w_m, \mathbf{b}} \models \psi_j^B$ iff. $\tilde{\mathcal{A}}_{w_n, \mathbf{b}} \models \psi_j^B$.

In other words, w_m and w_n produce the same vectors of $2k$ bits, which means that the Boolean function B has the same result. Therefore, we can conclude that $\tilde{\mathcal{A}}_{w_m, \mathbf{a}} \sqcup \tilde{\mathcal{A}}_{w_n, \mathbf{b}} \models \psi$. Together with Lemma B.13, this gives us $\mathbf{a}^m \mathbf{b}^n \models \varphi$, which is a contradiction. ◀

B.8 Proof of Theorem 4.10

Proof. We first show that we can extend Lemma 4.9 to show that there is no FC[REG]-formula that defines the language $L_{el} := \{\mathbf{a}^n \mathbf{b}^n \mid n \geq 1\}$. Assume for the sake of a contradiction that there is a formula $\varphi(\mathfrak{s}) \in \text{FC}[\text{REG}]$ such that $\mathcal{L}(\varphi) = L_{el}$. Our goal is now to prove that there exists a formula $\psi(\mathfrak{s}) \in \text{FC}[\]$ such that $\llbracket \varphi \rrbracket(w) = \llbracket \psi \rrbracket(w)$ for all $w \in \mathbf{a}^* \mathbf{b}^*$. To construct ψ , we first obtain $\psi' \in \text{FC}[\text{REG}]$ by replacing every constraint $x \dot{\in} \alpha$ in φ with a constraint for the language $\mathcal{L}(\alpha) \cap \mathbf{a}^* \mathbf{b}^*$. This is possible, as each of these languages $\mathcal{L}(\alpha) \cap \mathbf{a}^* \mathbf{b}^*$ is regular, due to the fact that the class of regular languages is closed under intersection.

Then $\llbracket \psi' \rrbracket(w) = \llbracket \varphi \rrbracket(w)$ holds for all $w \in \mathbf{a}^* \mathbf{b}^*$; as on these words, all variables in φ can only be mapped to elements of $\mathbf{a}^* \mathbf{b}^*$ as well. Next, we use Lemma 6.1 from [30], which states that every bounded regular language is an EPFC[]-language; where a language L is bounded if it is subset of a language $w_1^* \cdots w_k^*$ with $k \geq 1$ and $w_1, \dots, w_k \in \Sigma^*$. Clearly, $\mathbf{a}^* \mathbf{b}^*$ is bounded, which means that all constraints in ψ' use bounded regular languages.

Thus, we can obtain $\psi \in \text{FC}[\]$ from ψ' by replacing every constraint in ψ' with an equivalent EPFC[]-formula. Then we have $\psi \equiv \psi'$, which gives us $\llbracket \psi \rrbracket(w) = \llbracket \varphi \rrbracket(w)$ for all $w \in \mathbf{a}^* \mathbf{b}^*$. We conclude that

$$\varphi \equiv \mathfrak{s} \dot{\in} \mathbf{a}^* \mathbf{b}^* \wedge \psi.$$

Now ψ is an FC[]-formula; and as $\mathbf{a}^* \mathbf{b}^*$ is a simple regular expression, we can rewrite it into an equivalent FC[]-formula, using the construction from the proof of Theorem 3.22. Hence, φ is equivalent to an FC[]-formula; which means that $\mathcal{L}(\varphi) = L_{el}$ is an FC[]-language. This contradicts Lemma 4.9; hence, φ cannot exist.

Now assume that there the equal length relation is selectable in FC[REG], that is, assume there is some $\varphi_{el}(\mathfrak{s})(x, y)$ such that $\sigma \models \varphi_{el}$ if and only if $|\sigma(x)| = |\sigma(y)|$. Then we have $\mathcal{L}(\varphi_{el}) = L_{el}$ for

$$\varphi_{el}(\mathfrak{s})(\cdot) := \mathfrak{s} \dot{=} xy \wedge \varphi_{el}(x, y) \wedge x \dot{\in} \mathbf{a}^* \wedge y \dot{\in} \mathbf{b}^*.$$

This contradicts the previous paragraph, which means that the equal length relation is not selectable in FC[REG]. ◀

B.9 Proof of Theorem 4.11

Before we proceed to the actual proof in Appendix B.9.2, we use Appendix B.9.1 to introduce a new notational shorthand that simplifies our reasoning for that proof. But before that, we observe the following result on the expressive power of $\text{FC}_{(1)}^0[\]$:

► **Lemma B.14.** *A language $L \subseteq \Sigma^*$ is definable in $\text{FC}_{(1)}^0[\]$ if and only if L is finite or co-finite.*

Proof. The *if*-direction is straightforward. If L is finite, we have $L = \{w_1, \dots, w_n\}$ for some $n \geq 0$ and can simply define $\varphi := \bigvee_{i \in [n]} (\mathfrak{s} \dot{=} w_i)$. Likewise, we can define every co-finite language using negation.

For the *only-if*-direction, the proof of Lemma 3.15 (see Appendix A.4, “special cases”) allows us to exclude all cases where the structure variable \mathfrak{s} appears on the right side of a

word equation. Hence, we can assume that $\varphi \in \text{FC}_{\{1\}}^0[\]$ is defined using only equations of the form $\mathfrak{s} \doteq w$ with $w \in \Sigma^*$, conjunctions, disjunctions, and negations (without other variables, quantifiers play no role). Hence, $\mathcal{L}(\varphi)$ is obtained by combining singleton languages $\{w\}$ with intersection, union, and complement. Singleton languages are finite, negation turns a finite into a co-finite language (and vice versa), and unions and intersections preserve the property “finite or co-finite”. Thus, $\mathcal{L}(\varphi)$ is finite or co-finite. \blacktriangleleft

B.9.1 Regex patterns and regex equations

A *regex pattern* is a tuple $\alpha = (\alpha_1, \dots, \alpha_n)$ with $n \geq 0$, where each α_i is either a variable $x \in \Xi$, or a regular expression. If all regular expressions in α are simple (recall Section 3.4), we say that α is a *simple regex pattern*.

We use regex patterns instead of patterns to extend word equations to *regex equations*: every $x \in \Xi$ and every regex pattern α can be combined into a regex equation $x \doteq \alpha$, which is simple if α is simple.

To define the semantics of $x \doteq \alpha$ with $\alpha = (\alpha_1, \dots, \alpha_n)$ and $n \geq 0$, let $R \subseteq [n]$ be the set of all i such that α_i is regular expression, and let $V := [n] - R$ be the set of all i such that α_i is a variable. We now define the EPFC[REG]-formula

$$\varphi_{x,\alpha} := \exists y_1, \dots, y_n : (x \doteq y_1 \cdots y_n \wedge \bigwedge_{i \in R} y_i \dot{\in} \alpha_i \wedge \bigwedge_{i \in V} y_i \doteq \alpha_i).$$

For all structure variables $\vec{\mathfrak{s}}$, we then define that $\sigma \models_{\vec{\mathfrak{s}}} (x \doteq \alpha)$ if and only if $\sigma \models_{\vec{\mathfrak{s}}} \varphi_{x,\alpha}$. We define the free and structure variables of regex equations as for the rest of FC. Hence, if we fix some structure variables $\vec{\mathfrak{s}}$, then $\text{free}(x \doteq \alpha)$ is the set of all non-structure variables in x and α . We write regex patterns like patterns that contain regular expressions; see the following example.

► **Example B.15.** Let $\varphi(\mathfrak{s})() := \exists x, y : (\mathfrak{s} \doteq x \mathbf{ab}^* \mathbf{a} y)$. Then $\mathcal{L}(\varphi)$ is the set of all $w \in \Sigma^*$ that contain a subword $\mathbf{ab}^n \mathbf{a}$ with $n \geq 0$. Note that the word equation in φ is simple.

The following directly follows from techniques from the proofs of Theorem 3.22 and Lemma 3.15.

► **Lemma B.16.** *For every regex equation $(x \doteq \alpha)$, we can construct $\psi \in \text{EPFC}[\text{REG}]$ with $\psi \equiv \varphi_{x,\alpha}$ and $\text{wd}(\psi) = |\text{free}(\varphi_{x,\alpha})| + 3$. If α is simple, then we can have $\psi \in \text{EPFC}[\]$.*

Proof. The first part can be achieved by reordering the quantifiers as in the word equation case in the proof of Lemma 3.15 (see Appendix A.4). For example, if all α_i are regular expressions and n is even, we can define

$$\begin{aligned} \psi := \exists y, z_1 : & (x \doteq y z_1 \wedge y \dot{\in} \alpha_1 \wedge \\ & \exists y, z_2 : (z_1 \doteq y z_2 \wedge y \dot{\in} \alpha_2 \wedge \\ & \exists y, z_1 : (z_2 \doteq y z_1 \wedge y \dot{\in} \alpha_3 \wedge \\ & \vdots \\ & \exists y, z_2 : (z_1 \doteq y z_2 \wedge y \dot{\in} \alpha_n) \cdots))). \end{aligned}$$

If $\alpha_i = x_i$ for some variable $x_i \in \Xi$, we can avoid using y in this case and write $z_1 \doteq x_i z_2$ instead of $(z_1 \doteq x_i z_2) \wedge (y \dot{\in} \alpha_i)$ if i is even, or the respective other case if i is odd. In addition to the free variables from $\text{free}(\varphi_{x,\alpha})$, the width is only increased by the three additional variables y, z_1 , and z_2 .

If the regex pattern α is simple, then all its regular expressions are simple. As shown in the proof of Theorem 3.22 (see Appendix A.9), we can then replace every constraint $y \in \alpha_i$ with an equivalent EPFC[\cdot]-formula $\psi_{\alpha_i}(y)$. Moreover, we can see in that proof that $\text{wd}(\psi_{\alpha_i}) = 3$; and as we can reuse the variables z_1 and z_2 in ψ_{α_i} , replacing the constraints in ψ does not increase its width. \blacktriangleleft

B.9.2 Main part of the proof of Theorem 4.11

Proof. The undecidability results were proven in Theorem 4.6 of [31] for RGX^{core} , which is equivalent to SpLog and to $\text{EPFC}[\text{REG}]$.

But the heavy lifting of that proof was actually done in the proof of Theorem 14 of [29]. That paper examines *extended regular expressions* with one variable, also called *xregex* with one variable in [25] and (due to one anonymous reviewer's strong encouragement) in [31].

These xregex extend classical regular expressions with a variable binding operator $(\alpha)\%x$ and a variable recall operator x for a single variable x . For example, the xregex $((\mathbf{a} \cup \mathbf{b})^*)\%x x$ creates the language of all ww with $w \in \{\mathbf{a}, \mathbf{b}\}^*$, and $(\mathbf{a}^*)\%x \mathbf{b} x \mathbf{b} x$ the language of all $\mathbf{a}^n \mathbf{b} \mathbf{a}^n \mathbf{b} \mathbf{a}^n$ with $n \geq 0$. This is as much understanding of syntax and semantics as we need for purpose of this paper³.

The main idea of the proof of Theorem 14 in [29] is as follows: Given a so-called *extended Turing machine* M , define a language $\text{VALC}(M) \subseteq \{0, \#\}^*$ that contains exactly one word for every *valid* computation of M (i. e., an accepting computation on some input). In other words, there is a one-to-one correspondence between $\text{VALC}(M)$ and each word that is accepted by M . Note that the details of these extended Turing machines do not matter to our proof, as our translations function on a purely syntactical level.

Now define $\text{INVALC}(M) := \{0, \#\}^* - \text{VALC}(M)$. Our main goal is now to show that α can be converted into a formula $\varphi \in \text{EPFC}_{(1)}^4[\cdot]$ with $\mathcal{L}(\varphi) = \mathcal{L}(\alpha) = \text{INVALC}(M)$. We first assume that $\Sigma = \{0, \#\}$ and discuss larger alphabets later. After that, we discuss which undecidability results follow from this construction.

Creating the formula (binary alphabet): As shown in [29], given M , one can construct a one-variable xregex α with $\mathcal{L}(\alpha) = \text{INVALC}(M)$. The construction is rather lengthy; but it is described in a way that allows us to only consider the necessary modifications.

As one might expect, α is obtained by enumerating all possible types of errors that can cause a word to be an element of $\text{INVALC}(M)$. The proof in [29] distinguishes two different types of error: *structural errors*, where a word cannot be interpreted as the result of encoding a sequences of configurations of M , the first configuration is not initial, or the last configuration is not accepting; and *behavioral errors*, where we assume that it is an encoding of a sequence of configuration, but at least one configuration in the sequence does not have the right successor.

While structural errors can be handled with a classical regular expression, behavioral errors require the use of variables to handle the tape contents correctly. This makes makes expressing the structural errors straightforward for xregex, but requires considerable effort for $\text{FC}[\cdot]$. We first deal with structural errors.

³ The interested reader can find much more on xregex can be found in [29] and, more recently, [33]. In particular, the latter uses a much nicer form of semantics that is due to [61] which was also used in [30] to simplify the semantics of regex formulas.

Structural errors: In the encoding that is defined in [29], every configuration of M is encoded as a word from the language

$$L_C := \{00^{t_1}\#00^{t_2}\#00^a\#0^q \mid t_1, t_2 \geq 0, a \in \{0, 1\}, q \in [n]\}$$

where n is the number of states of M (hence, q encodes the current state). Here, 0^{t_1} and 0^{t_2} are unary encodings of the tape contents to the left and right of the head, and a is the head symbol under the head. The sequence of configurations of M is then encoded as a word from the language

$$L_{seq} := \{\#\#c_1\#\#c_2\#\#\dots\#\#c_n\#\# \mid n \geq 1, c_i \in L_C \text{ for all } i \in [n]\}.$$

Now define L_S as the subset of L_{seq} where c_1 has state $q = 1$, $t_1 = 0$, and $t_2 > 0$ (meaning initial state and head starting on the left of a non-empty input), and c_n has symbol a under the head and is in a state q such that M halts. We now say that $w \in \Sigma^*$ has a *structural error* if $w \notin L_V$. As $\text{VALC}(M) \subseteq L_S$ must hold, having $w \notin L_S$ is sufficient for $w \in \text{INVALC}(M)$.

We first define a formula φ_{seq} for the complement of L_{seq} . We define φ_{seq} using we use simple regex equations. As these have no free variables, we can use Lemma B.16 to interpret φ_S as formula from $\text{EPFC}_{\langle 1 \rangle}^3[\]$. We begin with

$$\varphi_{seq,1} := (\mathfrak{s} \doteq \varepsilon) \vee (\mathfrak{s} \doteq 0\Sigma^*) \vee (\mathfrak{s} \doteq \Sigma^*0) \vee (\mathfrak{s} \doteq \#0\Sigma^*) \vee (\mathfrak{s} \doteq \Sigma^*0\#) \vee (\mathfrak{s} \doteq \#) \vee (\mathfrak{s} \doteq \#\#).$$

Then we have $w \notin \mathcal{L}(\varphi_{seq,1})$ if and only if w is of the form $\#\#\Sigma^*\#\#$. Building on this, let

$$\begin{aligned} \varphi_{seq,2} := & \varphi_{S,1} \vee (\mathfrak{s} \doteq \Sigma^*\#\#\#\Sigma^*) \vee (\mathfrak{s} \doteq \Sigma^*\#\#0^+\#\#\Sigma^*) \vee (\mathfrak{s} \doteq \Sigma^*\#\#0^+\#0^+\#\#\Sigma^*) \\ & \vee (\mathfrak{s} \doteq \Sigma^*\#\#0^+\#0^+\#0^+\#\#\Sigma^*) \vee (\mathfrak{s} \doteq \Sigma^*0\#0^+\#0^+\#0^+\#0\Sigma^*). \end{aligned}$$

Observe that L_{seq} uses double hashes $\#\#$ to separate encodings of configurations, and single hashes $\#$ to separate the components within an encoded configuration. Now we have $w \notin \mathcal{L}(\varphi_{seq,2})$ if and only if w is of the form $\#\#(0^+\#0^+\#0^+\#0^+\#)\#$. Next, let

$$\varphi_{seq} := \varphi_{seq,2} \vee (\mathfrak{s} \doteq \Sigma^*00^n\#\#\Sigma^*) \vee (\mathfrak{s} \doteq \Sigma^*00\#0^+\#\#\Sigma^*).$$

In the encoding, each block of 0s to the left of a double hash encodes a state. Hence, the first part of φ_{seq} (after $\varphi_{seq,2}$) expresses that there is an encoding of a state q that is not in the state set $[n]$ of M . Likewise, the second part expresses that there is a tape symbol a that is not 0 or 1. Consequently, we have $w \notin \mathcal{L}(\varphi_{seq})$ if and only if $w \in L_{seq}$. In other words, φ_S defines the complement of L_{seq} . To extend this into a φ_V for the complement of L_V , we need to define two types of errors; namely, that the first encoded configuration is not initial, and that the last configuration is not halting. The first is handled by

$$\varphi_{S,1} := \varphi_{seq} \vee (\mathfrak{s} \doteq \#\#00\Sigma^*) \vee (\mathfrak{s} \doteq \#\#0^+\#0\#\Sigma^*) \vee (\mathfrak{s} \doteq \#\#0^+\#0^+\#0^+\#00\Sigma^*)$$

which has cases where the first configuration has $t_1 \neq 0$, $t_2 = 0$, or a $q \neq 1$ (in this order). Finally, let $\overline{H} \subseteq \{0, 1\} \times [n]$ be the set of all (a, q) such that M does not halt when reading symbol a in state q , and define

$$\varphi_S := \varphi_{S,1} \vee \bigvee_{(a,q) \in \overline{H}} (\mathfrak{s} \doteq \Sigma^*\#00^a\#0^q\#\#),$$

which expresses that M would not halt on the last configuration in the sequence. Now we have $w \in \mathcal{L}(\varphi_S)$ if and only if $w \notin L_S$; which means that φ_S describes exactly the words that have a structural error. Recall that we can interpret φ_S as a formula from $\text{EPFC}_{\langle 1 \rangle}^3[\]$.

Behavioral errors and combining the parts: For these behavioral errors, first note that Section 3.3 of [31] explains that the xregex for $\text{INVALC}(M)$ from the proof in [29] have no stars over the variable operators. Moreover, they can be rewritten into a union of xregexes that have no disjunctions over the variable operators (these are called *regex paths* in [31]), simply by factoring out the disjunctions. But in our terminology, these regex paths can be viewed as Boolean formulas of the form

$$\psi = \exists x: (\mathfrak{s} \doteq \alpha \wedge x \dot{\in} \beta),$$

where α is a regex pattern that has x as only variable and β is a regular expression. Moreover, one can verify by going through all the cases in the definitions of the behavioral errors in [29] that in every case, both the regex pattern α and the regular expression β are simple. Hence, we can apply Lemma B.16 and interpret each ψ as a formula from $\text{EPFC}_{\langle 1 \rangle}^4[\cdot]$. Then we define the Boolean formula φ_B as the disjunction of all these ψ , thus describing all behavioral errors. We then define $\varphi := \varphi_S \vee \varphi_B$ and have $\mathcal{L}(\varphi) = \text{INVALC}(M)$ with $\varphi \in \text{EPFC}_{\langle 1 \rangle}^4[\cdot]$.

Adapting the formula to larger alphabets: For larger alphabets, we need to address the problem that simple regular expressions can only express Σ^* , but not A^* for $A \subset \Sigma$ with $|A| \geq 2$ (recall ??). Hence, while 0^* is not problematic, $\{0, \#\}^*$ is not expressible. Luckily, any word that contains some letter from $\Sigma - \{0, \#\}$ is invalid anyway. The errors that were described by formulas with regex patterns that contain Σ^* still describe the errors they described before; and they also describe new ones. We extend φ_S with an additional disjunction $\bigvee_{a \in \Sigma - \{0, \#\}} (\mathfrak{s} \doteq \Sigma^* a \Sigma^*)$ to catch all words that consist only of the new letters. But no other changes are required.

Undecidable problems as consequences of the construction: As shown in Lemma 10 of [29], the peculiarities of extended Turing machines that are used in the construction do not affect the “usual” undecidability properties that one expects from Turing machines. In particular, we have that, given an extended Turing machine M , the question whether

1. M accepts *at least one* input is semi-decidable but not co-semi-decidable, and
2. M accepts *finitely many* inputs is neither semi-decidable, nor co-semi-decidable,

Given M , we can construct $\varphi \in \text{EPFC}_{\langle 1 \rangle}^4[\cdot]$ with $\mathcal{L}(\varphi) = \text{INVALC}(M)$. Hence, the following questions are undecidable:

- $\mathcal{L}(\varphi) \stackrel{?}{=} \Sigma^*$ is not semi-decidable, as we have $\text{INVALC}(M) = \Sigma^*$ if and only if $\text{VALC}(M) = \emptyset$. This also gives us undecidability of containment and equivalence.
- “Is $\mathcal{L}(\varphi)$ regular?” is neither semi-decidable, nor co-semidecidable. As shown in Lemma 13 of [29], we have that $\text{INVALC}(M)$ is regular if and only if it is co-finite, which holds if and only if $\text{VALC}(M)$ is finite.
- “Is there a pattern α with $\mathcal{L}(\alpha) = \mathcal{L}(\varphi)$?” is not semi-decidable. We shall prove this by showing that such an α exists if and only if $\mathcal{L}(\varphi) = \Sigma^*$. Assume there is an α with $\mathcal{L}(\alpha) = \text{INVALC}(M)$. As $\text{INVALC}(M)$ contains the words 0 and $\#$, we know that α cannot contain any terminals (as these would occur in all words in the pattern language). This means that there must be a variable x that occurs exactly once in α (otherwise, we could generate neither 0 nor $\#$). Hence, $\mathcal{L}(\alpha) = \Sigma^*$, as we can generate every $w \in \Sigma^*$ by defining $\sigma(x) := w$ and $\sigma(y) := \varepsilon$ for all other variables.
- “Is $\mathcal{L}(\varphi)$ expressible in $\text{FC}_{\langle 1 \rangle}^0[\cdot]$ ” is neither semi-decidable, nor co-semi-decidable. By Lemma B.14, the languages that are expressible in $\text{FC}_{\langle 1 \rangle}^0[\cdot]$ are finite or co-finite. $\text{INVALC}(M)$ cannot be finite, and it is co-finite if and only if $\text{VALC}(M)$ is finite.

The non-existence of a computable minimization function also follows from the undecidability of the question whether $\mathcal{L}(\varphi) = \Sigma^*$, using the same argument as for Theorem 4.9 in [31]: Every reasonable definition of the length of the formula will ensure that there are only finitely many φ such that $|\varphi|$ is minimal and $\mathcal{L}(\varphi) = \Sigma^*$. Thus, the set of these minimal representations is finite and thereby decidable. We could then decide $\mathcal{L}(\varphi) \stackrel{?}{=} \Sigma^*$ by applying the minimization algorithm to φ and checking whether the result is in the finite set. ◀

B.10 Proof of Theorem 4.12

Proof. Most of our reasoning relies on the undecidabilities that we established in Theorem 4.11. Like [31], we use a meta-theorem by Hartmanis [39] that basically states that for two systems of representations A and B such that given a representation $r \in B$, it is not co-semi-decidable whether r has an equivalent representation in A , there is a non-recursive tradeoff from B to A . See Kutrib [45] for details and background, and the proof of Theorem 4.10 for a detailed execution of the reasoning behind that meta-theorem.

Hence, Theorem 4.11 gives us non-recursive tradeoffs from $\text{EPFC}_{\langle 1 \rangle}^4[]$ to $\text{FC}_{\langle 1 \rangle}^0[]$ and all representations of regular languages (regular expressions, DFAs, NFAs, etc). Note that the lower bound for the trade-off to patterns remains open, as we have only established that the corresponding problem is not semi-decidable.

Regarding the tradeoffs from $\text{FC}_{\langle 1 \rangle}^4[]$, we first observe that the non-recursive tradeoff to $\text{EPFC}[]$ follows directly from the proof of Theorem 4.11 in [31], which demonstrates a non-recursive tradeoff from RGX^{core} to RGX^{core} . That proof relies on the same construction for $\text{INVALC}(M)$ as Theorem 4.11; and we have established that regular constraint are required for that.

For the remaining tradeoffs, we make use of the fact that we can now use negations. This allows us to adapt the proof of Theorem 4.11 to obtain more undecidability results. Given M , the proof of 4.11 allows us to construct $\varphi \in \text{EPFC}_{\langle 1 \rangle}^4[]$ with $\mathcal{L}(\varphi) = \text{INVALC}(M)$. Hence, we have $\neg\varphi \in \text{FC}_{\langle 1 \rangle}^4[]$ and $\mathcal{L}(\neg\varphi) = \text{VALC}(M)$.

Next, observe that although is not directly shown in [29], it follows directly by using the same methods that given M , it is neither semi-decidable, nor co-semi-decidable whether M accepts *exactly one* input.

Hence, given $\psi \in \text{FC}_{\langle 1 \rangle}^4[]$, the question whether there is a word $w \in \Sigma^*$ with $\mathcal{L}(\psi) = \{w\}$ is neither semi-decidable, nor co-semi-decidable. By invoking Hartmanis' meta-theorem, we obtain the non-recursive tradeoffs from $\text{FC}_{\langle 1 \rangle}^4[]$ to pattern languages.

Furthermore, observe that every pattern language $\mathcal{L}(\alpha)$ is either an infinite language (if α contains at least one variable) or a singleton language $\{w\}$ (if α contains no variables; i. e., $\alpha = w$ for some $w \in \Sigma^*$). Hence, this gives us non-recursive tradeoffs to pattern languages as well.

This raises the question whether the non-recursive tradeoff also exists if we only consider pattern languages with variables (after all, focusing on the special case of singleton languages might be considered close to cheating).

Although qwe leave the case for $\text{FC}_{\langle 1 \rangle}^4[]$ open, we can show non-recursive tradeoff from $\text{FC}_{\langle 1 \rangle}^5[]$ to patterns with variables. Given M , we can define $\psi \in \text{FC}_{\langle 1 \rangle}^5[]$ with

$$\mathcal{L}(\psi) := \text{VALC}(M) 0\#^3 0 \Sigma^*,$$

by defining

$$\psi := \exists x, y: ((s \doteq x 0\#^3 0 y) \wedge \neg\hat{\varphi}(x)),$$

where $\hat{\varphi}$ is obtained from the φ that is constructed from M as in the proof of Theorem 4.11 by replacing all occurrences of \mathfrak{s} with a new variable x .

Now we claim that a pattern α with $\mathcal{L}(\alpha) = \mathcal{L}(\varphi)$ exists if and only if $\text{VALC}(M)$ contains exactly one element. The *if*-direction is clear. Hence, assume such an α exists. As pattern languages are always either singleton languages or infinite, we know that $\text{VALC}(M) \neq \emptyset$. By definition of $\mathcal{L}(\psi)$, this means that $\mathcal{L}(\alpha)$ is infinite, which means that α contains at least one variable.

Moreover, as no word in $\text{VALC}(M)$ contains $\#^3$ as a subword, we know that every $w \in \mathcal{L}(\alpha)$ has a unique factorization $w = u \cdot 0\#^3 0v$ with $u \in \text{VALC}(M)$ and $v \in \Sigma^*$. We now consider the uniquely defined factorization

$$\alpha = u_0 x_1 u_1 \cdots x_n u_n$$

for some $n \geq 1$, with $u_0, \dots, u_n \in \Sigma^*$ and $x_1, \dots, x_n \in \Xi$. Now assume that u_0 does not have a prefix from the language $\text{VALC}(M) \cdot 0\#^3 0$, and define a pattern substitution σ with $\sigma(x_1) := \#^4$. Then $\sigma(\alpha)$ has a prefix of the form $\sigma(u_0 \cdot x_1) = u_0 \cdot \#^4$. But as $\#^4$ is not subword of any word in $\text{VALC}(M)$, this means that $\sigma(\alpha)$ does not have a factorization $w = u \cdot 0\#^3 0v$ with $u \in \text{VALC}(M)$ and $v \in \Sigma^*$, as the $\#^4$ would need to occur in the v , which would lead us to the conclusion that u_0 has a prefix from $\text{VALC}(M) \cdot 0\#^3 0$ and contradict our assumption that this is not the case.

Hence, we now consider the case that u_0 has a prefix from the language $\text{VALC}(M) \cdot 0\#^3 0$. As M cannot continue its computation after stopping, we have $|\text{VALC}(M)| = 1$.

Hence, $\mathcal{L}(\psi)$ can be expressed with a pattern with variables if and only if M accepts exactly one input. This means that this expressibility is neither semi-decidable nor co-semi-decidable; the latter allows us to use Hartmanis' meta-theorem to conclude non-recursive tradeoffs from $\text{FC}_{\langle 1 \rangle}^5[\]$ to patterns with at least one variable. \blacktriangleleft

B.11 Proof of Proposition 4.13

Recall that we discussed that Durnev [23] shows undecidability of satisfiability for EPC[$\]$ -formulas of the form $\forall s: \exists x, y, z: \varphi$, where φ is quantifier-free. But note that the formula that is constructed there is not an FC-formula, as it contains equations of the form $x0 \doteq 0x$. In principle, one could prove Proposition 4.13 by rewriting the proof from [23] or sketching which changes need to be made. But as the following proof is short enough (and as it is an opportunity to use FRACTRAN), we give an original proof instead.

Proof. We show the undecidability by providing a reduction from the halting problem for FRACTRAN-programs (introduced by Conway [15]). A FRACTRAN-program is a finite sequence $P := (\frac{n_1}{d_1}, \frac{d_2}{n_2}, \dots, \frac{d_k}{n_k})$ with $k \geq 1$ and $n_i, d_i \geq 1$. The input (and only memory) is a natural number $m \geq 1$.

This is interpreted as follows: In each step, we search the list of fractions in the program P from left to right until we find the first fraction $\frac{n_i}{d_i}$ such that the product $m \frac{n_i}{d_i}$ is a natural number. If no such fraction can be found, P terminates. Otherwise, we update m to $m \frac{n_i}{d_i}$ and proceed to the next step.

By reducing the fractions, we can ensure that all n_i and d_i are co-prime. Furthermore, as we are interested in termination, we can exclude cases where $d_i = 1$. The halting problem for FRACTRAN (deciding whether a program P terminates on an input number $n \geq 1$) is undecidable (see Kurtz and Simon [44]).

Given P and m , our goal is to construct a Boolean formula $\varphi \in \text{FC}_3$ that is satisfiable if and only if P terminates on input m . Assume that $\Sigma \supseteq \{0, 1\}$. The construction shall ensure

that $\llbracket \varphi \rrbracket \langle w \rangle \neq \emptyset$ holds if and only if $w \in 0(1^+0)^+$ is an encoding of an accepting run of P on m . More formally, we will have

$$w = 01^{c_1} 01^{c_2} 01^{c_3} 0 \dots 01^{c_{t-1}} 01^{c_t} 0,$$

with $c_j \geq 1$ for all j , where $c_1 = m$, each c_{i+1} is the number that succeeds c_i after applying one step of P , and c_t is a number on which P terminates (i. e., c_t is divided by no d_i). We first define φ_{cod} to be the following Boolean FC_3 -formula:

$$\varphi_{cod} := \exists x: \mathfrak{s} \doteq 01^m 0 x \tag{1}$$

$$\wedge \exists x: \mathfrak{s} \doteq x 0 \tag{2}$$

$$\wedge \neg \exists x, y: \mathfrak{s} \doteq x 0 0 y \tag{3}$$

$$\wedge \bigvee_{a \in \Sigma - \{0,1\}} \neg \exists x, y: \mathfrak{s} \doteq x a y \tag{4}$$

The parts of the conjunction have the following roles: (1) expresses that w starts with $01^m 0$, (2) states that it ends on 0, (3) requires that it w does not contain 00 , and (4) forbids all letters other than 0 and 1. Hence, these four parts together ensure that $w \in 01^m(01^+0)^*$ holds. Hence, if $w \in \mathcal{L}(\varphi_{cod})$, we know that w encodes a sequence $c_1, \dots, c_t \geq 1$ for some $t \geq 1$ with $c_1 = m$. The next step is defining the following Boolean formula:

$$\varphi_{term} := \forall x, y: \left((\mathfrak{s} \doteq x 0 y 0 \wedge \neg \exists x, z: y \doteq x 0 z) \rightarrow \neg \bigvee_{i=1}^k \exists x: y \doteq x^{d_i} \right)$$

The left side of the implication states that y contains the last block of 1s in w , the right side that the length of y is not divided by any d_i . In other words, φ_{term} expresses that P terminates on c_t . All that remains is defining a formula that expresses that c_{i+1} is the successor of c_i when one step of P is applied. This is the job of the following Boolean formula:

$$\begin{aligned} \varphi_{step} := \forall x, y, z: & \left((x \doteq 0 y 0 z 0 \wedge \neg \exists x, z: y \doteq x 0 z \wedge \neg \exists x, y: z \doteq x 0 y) \right. \\ & \left. \rightarrow \bigvee_{j=1}^k \left(\exists x: (y \doteq x^{d_j} \wedge z = x^{n_j}) \wedge \bigwedge_{l < j} \neg \exists x: y \doteq x^{d_l} \right) \right) \end{aligned}$$

This formula expresses that, if w contains $01^{c_i} 01^{c_{i+1}} 0$, then $c_{i+1} = \frac{n_j}{d_j} c_i$ holds for some $1 \leq j \leq k$, and c_i is not divided by any d_l with $l < j$.

We now put the parts together and define $\varphi := \varphi_{cod} \wedge \varphi_{step} \wedge \varphi_{term}$. Then $\llbracket \varphi \rrbracket \langle w \rangle \neq \emptyset$ if and only if w encodes a terminating run of the FRACTRAN -program P on the input m . In other words, P terminates on m if and only if the constructed $\varphi \in \text{FC}_3$ is satisfiable. As the halting problem for FRACTRAN is undecidable, we conclude that satisfiability for FC_3 is undecidable. \blacktriangleleft