The Complexity of Width Minimization for Existential Positive Queries

Simone Bova
Institut für Informationssysteme
Technische Universität Wien
Favoritenstraße 9-11
A-1040 Wien, Austria
simone.bova@tuwien.ac.at

Hubie Chen
Departamento LSI, Facultad de Informática
Universidad del País Vasco
E-20018 San Sebastián, Spain
hubie.chen@ehu.es
and
IKERBASQUE, Basque Foundation for Science
E-48011 Bilbao, Spain

ABSTRACT

Existential positive queries are logical sentences built from conjunction, disjunction, and existential quantification, and are also known as select-project-join-union queries in database theory, where they are recognized as a basic and fundamental class of queries. It is known that the number of variables needed to express an existential positive query is the crucial parameter determining the complexity of evaluating it on a database, and is hence a natural measure from the perspective of query optimization and rewriting. In this article, we study the complexity of the natural decision problem associated to this measure, which we call the expressibility problem: Given an existential positive query and a number k, can the query be expressed using k (or fewer) variables? We precisely determine the complexity of the expressibility problem, showing that it is complete for the level Π_2^p of the polynomial hierarchy. Moreover, we prove that the expressibility problem is undecidable in positive logic (that is, existential positive logic plus universal quantification), thus establishing existential positive logic as a maximal syntactic fragment where expressibility is decidable.

Categories and Subject Descriptors

F.1.3 [**Complexity Measures and Classes**]: [Reducibility and completeness, Relations among complexity classes]; F.4.1 [**Mathematical Logic**]: [Logic and constraint programming]; H.2.3 [**Languages**]: [Query languages]

Keywords

Existential positive queries, positive queries, finite-variable logics, computational complexity, query optimization

1. INTRODUCTION

The *width* of a first-order sentence ϕ is the maximum number of free variables in a subformula of ϕ, and is considered as a fundamental measure of logical sentences in database theory, descriptive complexity, and finite model theory [15, 12].

There has long been interest in reducing a given sentence to a logically equivalent sentence of *minimum* width. A primary motivation for this interest is that the natural algorithm for model checking a relational first-order sentence on a finite relational structure has an exponential dependence on the width of the sentence, as observed by Vardi [18]. This model checking is a fundamental computational problem that translates, in the database theory parlance, to the evaluation of a (Boolean) query over a relational database.

Indeed, equivalence preserving rewriting rules are routinely used in query optimization to transform a query into an equivalent query easier to evaluate [17]. For instance, the width-four query $\exists x \exists y \exists z \exists w (Exy \wedge Eyz \wedge Ezw)$ can be transformed to the width-two equivalent query $\exists x \exists y (Exy \wedge \exists z (Eyz \wedge \exists w Ezw))$; in this example, the transformation is justified by syntactic replacements, namely, the logical equivalence of $\exists x(\phi \wedge \psi)$ and $\phi \wedge \exists x \psi$ under the condition that x does not occur free in ϕ.

By renaming bound variables, a sentence of width at most k can be rewritten as a logically equivalent sentence using at most k distinct variable symbols. Consequently, the following natural decision problem, which we call the *expressibility* problem, captures the question of to what extent the width of a sentence can be minimized:

> Given a sentence ϕ in relational first-order logic and a number k, is ϕ logically equivalent to a sentence using k (or fewer) variable symbols?

Expressibility is undecidable in general, and computationally intractable even under severe restrictions on the quantifiers and connectives permitted in the sentences; the reader is referred to the discussion part of this introduction for more details. At the same time, as has been articulated in the literature [13], the typical situation in the database setting is the posing of a relatively short query to relatively large database, or in logical parlance, the evaluation of a short formula on a large relational structure. It has consequently been argued that, in measuring the time complexity of this task, one could reasonably allow a slow (that is, possibly non-polynomial-time) computable preprocessing of the

10.5441/002/icdt.2014.25

formula, so long as the desired evaluation can be performed in polynomial time following this preprocessing. Relaxing polynomial-time computation so that an arbitrary dependence in a *parameter* is tolerated yields, in essence, the notion of *fixed-parameter tractability*. This notion of tractability is the base of *parameterized complexity theory*, which provides a taxonomy for reasoning about and classifying problems where each instance has an associated parameter [11].

A recent study on the complexity of model checking in existential positive logic [6] revealed that the number of variables needed to express a sentence is the crucial parameter determining complexity. Specifically, it was shown that, on a set of sentences having bounded arity (that is, there is a constant upper bounding the arity of all relation symbols appearing in the sentences), model checking is fixed-parameter tractable if and only if there exists a constant k such that each sentence in the set is logically equivalent to a sentence using at most k variable symbols (the negative part holds, as usual, under standard assumptions in parameterized complexity theory). From the perspective of this result, the computational problem of determining exactly how many variables are needed to express a sentence is very well-motivated.

Results. Although undecidable in general, the expressibility problem is decidable in syntactic fragments of relational first-order logic that are semantically equivalent to fundamental classes of database queries, notably *unions of conjunctive queries* (or *select-project-join-union* queries), corresponding to *existential positive* logic (the syntactic fragment obtained by restricting the logical signature to existential quantification, $\exists$, conjunction, $\wedge$, and disjunction, $\vee$). The existential positive fragment has received much attention in database theory in general [1], and in query optimization in particular since the classical work of Sagiv and Yannakakis [16], where the authors showed that the problem of deciding whether two existential positive formulas are logically equivalent (whether two queries have the same evaluation in all databases) is complete for the the level Π_2^p of the polynomial hierarchy.

A pertinent case to introduce our work is that of *conjunctive queries*, corresponding to existential positive logic without disjunction; this fragment is also known as *primitive positive* logic. The equivalence problem for conjunctive queries has been proved NP-complete in the seminal article by Chandra and Merlin [4], whereas the complexity of the expressibility problem for conjunctive queries has been proved NP-complete more recently [8], for all $k \geq 2$ and all signatures containing at least one binary relation symbol.

In this paper, we study the computational complexity of the expressibility problem for existential positive queries, proving the following two results:

- The expressibility problem for existential positive queries is Π_2^p-complete, for all $k \geq 6$ and all signatures containing at least one binary relation symbol (Theorem 6).
- The expressibility problem for positive queries (that is, existential positive queries plus universal quantification) is undecidable, for all $k \geq 3$ and all signatures containing countably many unary relation symbols and three binary relation symbols (Theorem 19).

The first result characterizes exactly the computational complexity of the expressibility problem for existential positive queries, and complements the previous results by Sagiv and Yannakakis by providing an existential positive analogue of the result on expressibility for conjunctive queries. The second result shows that existential positive logic is a maximal syntactic fragment of relational first-order logic where the expressibility problem is decidable (note that existential positive logic plus negation is all first-order logic), in this sense showing the optimality of the previous complexity result. (This second result also tightens the folklore fact that the expressibility problem is undecidable in relational first-order logic).

Discussion. A standard method to define syntactic fragments of relational equality-free first-order logic is by restricting the logical symbols in $\{\forall, \exists, \wedge, \vee, \neg\}$ allowed, subject to having at least one quantifier and at least one binary connective; for the sake of conciseness, we freely identify a fragment with the defining subset of logical symbols.

It is readily verified that if the sentences in one such fragment are logically equivalent to the negations of the sentences in another such fragment (for instance, the sentences in the fragment $\{\forall, \vee\}$ are logically equivalent to the negations of the sentences in $\{\exists, \wedge\}$), then the expressibility problem has the same complexity on both fragments.

Thus the classification of the expressibility problem, over the different syntactic fragments of relational equality-free first-order logic, indeed reduces to the classification of the following fragments: $\{\forall, \exists, \wedge, \vee, \neg\}$, $\{\forall, \exists, \wedge, \vee\}$, $\{\exists, \wedge, \vee\}$, $\{\forall, \exists, \wedge\}$, $\{\exists, \wedge\}$, $\{\exists, \vee\}$.

Note that a sentence in the fragment $\{\exists, \vee\}$ is expressible using at most k variables if and only if no atom contains more than k free variables. In view of this observation and the results in this work and in the literature, the following is known about the classification of the expressibility problem over syntactic fragments:

- $\{\forall, \exists, \wedge, \vee\}$ and $\{\forall, \exists, \wedge, \vee, \neg\}$ are undecidable by Theorem 19;
- $\{\exists, \wedge, \vee\}$ is Π_2^p-complete by Theorem 6;
- $\{\exists, \wedge\}$ is NP-complete by [8, Theorem 6];
- $\{\exists, \vee\}$ is polynomial-time tractable.

Thus understanding the computability (and complexity) status of the expressibility problem in the fragment defined by $\{\forall, \exists, \wedge\}$, also known as *(quantified) conjunctive positive* logic, would complete the classification, which we suggest as a problem for future research. We insist that this problem is, in principle, highly nonobvious, by noting that the entailment and equivalence problems for conjunctive positive logic have been proved decidable only relatively recently; their exact complexity is still quite open [7].

Organization. The paper is organized as follows. In Section 2, we introduce the basic definitions and the fundamental theory of primitive positive logic needed in the rest of the paper. In Section 3, we develop the theory of existential positive logic needed to settle the complexity upper bound for the expressibility problem. In Section 4, we prove the complexity lower bound for the expressibility problem in existential positive logic. In Section 5, we prove undecidability of the expressibility problem in positive logic.

2. PRELIMINARIES

For an integer $k \geq 0$, we use $\underline{k}$ to denote the set $\{1, \ldots, k\}$, with the convention that $\underline{0} = \emptyset$.

In this paper, we focus on relational first-order logic. A *signature* σ is a set of *relation symbols*, each of which has an associated natural number called its *arity*.

2.1 Structures

A *structure* $\mathbf{A}$ (over signature σ) is given by a nonempty set A called the *universe* of the structure and denoted by the corresponding non-bold letter, and a relation $R^{\mathbf{A}} \subseteq A^r$ for each arity r relation symbol $R \in \sigma$. A structure is *finite* if its universe is finite.

A collection of structures is said to be *similar* if they share the same signature. Let $\mathbf{A}, \mathbf{B}$ be similar structures on the signature σ. The *union* of $\mathbf{A}$ and $\mathbf{B}$ is the structure $\mathbf{A} \cup \mathbf{B}$ with universe $A \cup B$ and with $R^{\mathbf{A} \cup \mathbf{B}} = R^{\mathbf{A}} \cup R^{\mathbf{B}}$ for each arity r relation symbol $R \in \sigma$. A *homomorphism* from $\mathbf{A}$ to $\mathbf{B}$ is a mapping $h : A \to B$ such that for each symbol $R \in \sigma$, it holds that $h(R^{\mathbf{A}}) \subseteq R^{\mathbf{B}}$, by which is meant that for each tuple $(a_1, \ldots, a_k) \in R^{\mathbf{A}}$, one has $(h(a_1), \ldots, h(a_k)) \in R^{\mathbf{B}}$. We will sometimes simply write $\mathbf{A} \to \mathbf{B}$ to indicate that there exists a homomorphism from $\mathbf{A}$ to $\mathbf{B}$. We say that $\mathbf{A}$ and $\mathbf{B}$ are *homomorphically equivalent* if $\mathbf{A} \to \mathbf{B}$ and $\mathbf{B} \to \mathbf{A}$ both hold.

The structure $\mathbf{B}$ is a *substructure* of the structure $\mathbf{A}$ if $B \subseteq A$ and $R^{\mathbf{B}} \subseteq R^{\mathbf{A}}$ for all relation symbols R; if $R^{\mathbf{B}} = R^{\mathbf{A}} \cap B^r$ for each arity r relation symbol $R \in \sigma$, the structure $\mathbf{B}$ is called the substructure of $\mathbf{A}$ *generated* by B, also denoted by $\mathbf{A}|_B$. When $\mathbf{B}$ is a substructure of $\mathbf{A}$, there exists a homomorphism h from $\mathbf{A}$ to $\mathbf{B}$, and h fixes each element $b \in B$, the mapping h is said to be a *retraction* from $\mathbf{A}$ to $\mathbf{B}$; when there exists a retraction from $\mathbf{A}$ to $\mathbf{B}$, it is said that $\mathbf{A}$ *retracts* to $\mathbf{B}$. A *core* of the structure $\mathbf{A}$ is a structure $\mathbf{C}$ such that $\mathbf{A}$ retracts to $\mathbf{C}$, but $\mathbf{A}$ does not retract to any proper substructure of $\mathbf{C}$. We will make use of the following well-known facts on cores [14]: (1) each finite structure has a core; (2) all cores of a finite structure are isomorphic. From these facts, it is reasonable to speak of *the* core of a finite structure, which we do, and we use $\mathsf{core}(\mathbf{A})$ to denote a representative from the set of all cores of a finite structure $\mathbf{A}$.

We define the *Gaifman graph* of a structure $\mathbf{B}$ to be the undirected graph $G(\mathbf{B})$ with vertex set B and having an edge $\{b, b'\}$ if and only if b and b' co-occur in a tuple of $\mathbf{B}$.

A *tree decomposition* of an undirected graph G with vertex set B is a pair (T, β) consisting of a tree T and a map $\beta : V^T \to \wp(B)$ defined on the vertex set V^T of T such that, for each vertex $t \in V^T$, it holds that $\beta(t)$ is a non-empty subset of B, called the *bag* of t, and the following conditions hold:

- For each $b \in B$, the vertices $\{t \mid b \in \beta(t)\}$ form a connected subtree of T.
- For each edge $\{b, b'\}$ of G, there exists a vertex $t \in V^T$ such that $\{b, b'\} \subseteq \beta(t)$.

The *width* of a tree decomposition (T, β) is defined as

$$(\max_{t \in V^T} |\beta(t)|) - 1.$$

The *treewidth* of an undirected graph G, denoted by $\mathsf{tw}(G)$, is the minimum width over all tree decompositions of G; the treewidth of a structure $\mathbf{B}$, denoted by $\mathsf{tw}(\mathbf{B})$, is defined as $\mathsf{tw}(G(\mathbf{B}))$.

2.2 Formulas

An *atom* (over signature σ) is an equality of variables ($x = y$) or is a predicate application $R(x_1, \ldots, x_r)$, which we also write $Rx_1 \ldots x_r$, where $x_1, \ldots, x_r$ are variables, and $R \in \sigma$ is an arity r relation symbol. A *formula* (over signature σ) is built from atoms (over σ), conjunction ($\wedge$), disjunction ($\vee$), universal quantification ($\forall$), and existential quantification ($\exists$). A *sentence* is a formula having no free variables. In the sequel, we let FO denote the set of relational first-order formulas. For each set L of first-order formulas and each integer $k \geq 1$, we let L^k denote the subset of L containing formulas that use at most k variables, and L_σ denote the subset of L containing formulas over signature σ.

An *positive* formula (over signature σ) is a formula built from atoms (over σ) using conjunction, disjunction, universal quantification, and existential quantification; we let PFO denote the set of positive formulas. An *existential positive* formula (over signature σ) is a formula built from atoms (over σ) using conjunction, disjunction, and existential quantification; we let EP denote the set of existential positive formulas. A *primitive positive* formula (over signature σ) is a formula built from atoms (over σ) using conjunction and existential quantification; we let PP denote the set of primitive positive formulas.

We use the following standard terminology and notation from logic. For a structure $\mathbf{A}$ and a sentence ϕ over the same signature, we write $\mathbf{A} \models \phi$ if the sentence ϕ is *true* in the structure $\mathbf{A}$. When $\mathbf{A}$ is a structure, f is a mapping from variables to the universe of $\mathbf{A}$, and ψ is a formula over the signature of $\mathbf{A}$, we write $\mathbf{A}, f \models \psi$ to indicate that ψ is satisfied by $\mathbf{A}$ and f. Let ϕ and ψ be sentences over the same signature σ. Then, ϕ *entails* ψ (denoted $\phi \models \psi$) if, for all structures $\mathbf{A}$ over σ, it holds that $\mathbf{A} \models \phi$ implies $\mathbf{A} \models \psi$; also, ϕ and ψ are *logically equivalent* (denoted $\phi \equiv \psi$) if $\phi \models \psi$ and $\psi \models \phi$.

We use the following terminology and notation. Let σ be a signature, let ϕ be a primitive positive formula over σ, and let $\mathbf{A}$ be a finite structure over σ. By the *existential closure* of a formula, we mean the sentence obtained by existentially quantifying the free variables of the formula.

- $\mathbf{C}[\phi]$ denotes the *canonical structure* induced by the existential closure of ϕ, as follows. Let ϕ^c be the existential closure of the prenex form of ϕ. Let $\mathsf{elim}_=(\phi^c)$ be obtained by eliminating equalities from ϕ^c using the following syntactic transformations: for each equality $x = y$ on distinct variables, replace all instances of y with x in the quantifier free part, and remove the quantifier $\exists y$ from the prefix; remove equalities of the form $x = x$.

  Define $\mathbf{C}[\phi]$ to be the structure having a universe element for each existentially quantified variable in the formula $\mathsf{elim}_=(\phi^c)$, and where, for each $R \in \sigma$, the relation $R^{\mathbf{C}[\phi]}$ contains $(x_1, \ldots, x_k)$ if and only if it holds that $R(x_1, \ldots, x_k)$ appears in the quantifier free part of $\mathsf{elim}_=(\phi^c)$.

- $Q[\mathbf{A}]$ denotes the *canonical query* of $\mathbf{A}$, defined as follows. If $A = \{a_1, \ldots, a_n\}$, then $Q[\mathbf{A}]$ is equal to

  $$\exists a_1 \ldots \exists a_n \bigwedge_{R \in \sigma} \bigwedge_{(a_{i_1}, \ldots, a_{i_r}) \in R^{\mathbf{A}}} R(a_{i_1}, \ldots, a_{i_r}),$$

  where r is the arity of $R \in \sigma$.

We will use the following known fact.

Proposition 1. *(Chandra-Merlin [4]) Let ϕ be a sentence and let $\mathbf{A}$ be a finite structure, such that $\phi = Q[\mathbf{A}]$ or $\mathbf{A} = \mathbf{C}[\phi]$. Then, for any structure $\mathbf{B}$, it holds that $\mathbf{A} \to \mathbf{B}$ if and only if $\mathbf{B} \models \phi$.*

It is straightforward to verify that the existential closure of any primitive positive formula ϕ is logically equivalent to $Q[\mathbf{C}[\phi]]$, and that every finite structure $\mathbf{A}$ is homomorphically equivalent to $\mathbf{C}[Q[\mathbf{A}]]$.

3. EXISTENTIAL POSITIVE LOGIC

In this section, we define the problem under study, and then establish some basic facts on existential positive logic.

Definition 2. *Let $\mathrm{L} \in \{\mathrm{PFO}, \mathrm{EP}\}$. The computational problem* L-Expr *is defined as follows:*

> *Given a sentence $\phi \in \mathrm{L}$ and an integer $k \geq 1$, decide whether ϕ is logically equivalent to a sentence in L^k.*

Moreover, for every signature σ and every integer $m \geq 1$, L^m_σ-Expr is the restriction of L-Expr *to instances where $\phi \in \mathrm{L}_\sigma$ and $k = m$.*

Note that throughout this paper, the only notion of reduction that we use is many-one polynomial-time reduction.

Definition 3. *A sentence ϕ in* EP *is in* disjunctive *form if $\phi = \bigvee_{i \in \underline{n}} \phi_i$, where, for all $i \in \underline{n}$, ϕ_i is a sentence in* PP*; such a disjunctive form is* irredundant *if there do not exist distinct $i, j \in \underline{n}$ such that $\phi_i \models \phi_j$.*

We will make use of the following syntactic transformations, which preserve logical equivalence:

$$\exists x(\theta \vee \theta') \equiv \exists x\theta \vee \exists x\theta'; \tag{E1}$$

$$\theta \wedge (\theta' \vee \theta'') \equiv (\theta \wedge \theta') \vee (\theta \wedge \theta''); \tag{E2}$$

$$\exists x\theta \equiv \theta, \text{ if } x \text{ not free in } \theta; \tag{E3}$$

$$\theta \vee \theta' \equiv \theta', \text{ if } \theta \models \theta'. \tag{E4}$$

Given an arbitrary existential positive sentence, an equivalent existential positive sentence in disjunctive form is computable by iterated syntactic replacements exploiting the facts (E1) and (E2) above; also, given an existential positive sentence in disjunctive form, an equivalent existential positive sentence in irredundant disjunctive form is computable by iterated syntactic replacements exploiting the fact (E4) above.

The proof that our computational problem is contained in the complexity class Π^p_2 relies on the following lemma.

Lemma 4. *Let ϕ and ψ be sentences in* EP*. Let $\bigvee_{i \in \underline{m}} \phi_i$ and $\bigvee_{j \in \underline{n}} \psi_j$ be disjunctive forms in* EP *logically equivalent to ϕ and ψ, respectively. The following hold.*

1. *$\phi \models \psi$ if and only if, for all $i \in \underline{m}$, there exists $j \in \underline{n}$ such that $\phi_i \models \psi_j$.*
2. *If the above disjunctive forms are irredundant and $\phi \equiv \psi$, then $m = n$ and there exists a bijection $\pi \colon \underline{m} \to \underline{m}$ such that for all $i \in \underline{m}$ it holds that $\phi_i \equiv \psi_{\pi(i)}$.*
3. *Let $k \geq 1$ be an integer. Then, ϕ is logically equivalent to a sentence in EP^k if and only if, for all $i \in \underline{m}$, there exists $i' \in \underline{m}$ such that $\phi_i \models \phi_{i'}$ and $\phi_{i'}$ is logically equivalent to a sentence in PP^k.*

Proof. For (1), the backwards direction is clear. For the forwards direction, let $i \in \underline{m}$. We have $\mathbf{C}[\phi_i] \models \phi$, from which it follows that $\mathbf{C}[\phi_i] \models \psi$. We must then have that there exists $j \in \underline{n}$ such that $\mathbf{C}[\phi_i] \models \psi_j$, from which the result follows from Proposition 1.

For (2), let $i \in \underline{m}$. By (1), there exists $j \in \underline{n}$ such that $\phi_i \models \psi_j$. We claim that $\phi_i \equiv \psi_j$. This is because there exists $i' \in \underline{m}$ such that $\psi_j \models \phi_{i'}$; if $i \neq i'$, then this implies that the disjunctive form for ϕ is not irredundant, a contradiction. Since the disjunctive form for ψ is irredundant, there is a unique $j \in \underline{n}$ satisfying the condition $\phi_i \equiv \psi_j$, and we thus obtain an injection $\pi : \underline{m} \to \underline{n}$, as well as that $m \leq n$. By symmetric reasoning, we obtain that $n \leq m$ and so $m = n$ and the injection π is a bijection.

For (3), first let $\phi \in \mathrm{EP}$. If ϕ is logically equivalent to a sentence in EP^k, say ϕ', then the disjunctive form of ϕ' obtained using the above transformations (E1), (E2) and (E3) is such that each disjunct is a primitive positive sentence in PP^k. This implies that there is an irredundant disjunctive form $\bigvee_{j \in \underline{n}} \psi_j$ logically equivalent to ϕ where each disjunct is in PP^k. By (1), for any $i \in \underline{m}$, there exists $j \in \underline{n}$ such that $\phi_i \models \psi_j$. Since there is a sub-disjunction of $\bigvee_{i \in \underline{m}} \phi_i$ that is irredundant, by (2) there exists $i' \in \underline{m}$ such that $\phi_{i'}$ and ψ_j are logically equivalent. We then have $\phi_i \models \phi_{i'}$, as desired.

Now suppose that $\rho : \underline{m} \to \underline{m}$ is a mapping such that for each $i \in \underline{m}$, it holds that $\phi_i \models \phi_{\rho(i)}$ and each $\phi_{\rho(i)}$ is logically equivalent to a sentence in PP^k. Clearly, ϕ is logically equivalent to $\bigvee_{i \in \underline{m}} \phi_{\rho(i)}$. $\square$

The condition in Lemma 4(3) allows to establish containment in Π^p_2 for the problem under consideration.

Proposition 5. *The problem* EP-Expr *is in the complexity class Π^p_2.*

4. COMPLEXITY RESULT

We state our main complexity result.

Theorem 6. *Let σ be a signature that contains a relation symbol of at least binary arity. For each $k \geq 6$, the problem EP^k_σ-Expr is Π^p_2-complete.*

Proof. Proposition 5 proves containment in Π^p_2. For Π^p_2-hardness, the case where σ contains a binary relation symbol is proved in Theorem 16; the higher-arity case is easily verified. $\square$

Note that if σ is a signature that contains only unary relation symbols, then each sentence in EP_σ is logically equivalent to a sentence in EP^1_σ, so the problem EP^k_σ-Expr is trivial for all $k \geq 1$. Hence the statement of Theorem 6 also establishes a complete complexity classification of the expressibility problem with respect to all relational signatures, in the spirit of the purely syntactic classification theory in the book by Börger, Grädel, and Gurevich [2].

We devote the rest of this section to the proof of the complexity lower bound.

4.1 Source Problem

When $\mathbf{B}$ is a structure, define Π_k-$\mathsf{QCSP}(\mathbf{B})$ to be the problem of deciding, given a Π_k prenex sentence Φ whose quantifier-free part is a conjunction of atoms without equality, whether or not $\mathbf{B} \models \Phi$; define Σ_k-$\mathsf{QCSP}(\mathbf{B})$ similarly, with respect to Σ_k sentences. For $q \geq 2$, we define the structure $\mathbf{K}_q$, the clique on q vertices, to be the structure with universe $\underline{q}$ and that interprets the binary relation symbol E by $E^{\mathbf{K}_q} = \{(i,j) \in \underline{q}^2 \mid i \neq j\}$. Our Π_2^p hardness results will be proved by showing reductions from the problems Π_2-$\mathsf{QCSP}(\mathbf{K}_q)$, where $q \geq 3$.

Proposition 7. *(follows from [3]) Let $q \geq 3$. For each even $k \geq 2$, the problem Π_k-$\mathsf{QCSP}(\mathbf{K}_q)$ is Π_k^p-complete; and, for each odd $k \geq 3$, the problem Σ_k-$\mathsf{QCSP}(\mathbf{K}_q)$ is Σ_k^p-complete.*

Proof. Let $\mathbf{B}$ be the structure with universe $\{0,1\}$ and with a single relation, $R^{\mathbf{B}} = \{0,1\}^3 \setminus \{(0,0,0),(1,1,1)\}$. Under the bounds on k given in the proposition statement, one has that Π_k-$\mathsf{QCSP}(\mathbf{B})$ and Σ_k-$\mathsf{QCSP}(\mathbf{B})$ are Π_k^p-complete and Σ_k^p-complete, respectively; this follows from [5, Theorem 7.2]. We use the construction of [3, Proposition 5.1] to give a reduction from those problems to the present problems. The only modification needed is the following. Each universally quantified variable in an instance of Π_k-$\mathsf{QCSP}(\mathbf{B})$ or Σ_k-$\mathsf{QCSP}(\mathbf{B})$ is translated to a universally quantified variable followed by two existentially quantified variables. Such existentially quantified variables can be shifted right without changing the truth-value of the sentence. By the assumed bounds on k, each block of universally quantified variables has a block of existentially quantified variables to its right, so we indeed obtain a reduction that preserves the quantifier prefix (in the sense of being Π_k or Σ_k). $\square$

Remark 8. *An inspection of the proof of Proposition 7 yields that the hardness results hold on instances Φ where the quantifier-free part Φ_G has the property that $E^{\mathbf{C}[\Phi_G]}$ is symmetric and irreflexive. In the sequel, we will assume that Φ_G has this property. Indeed, one can always replace $E^{\mathbf{C}[\Phi_G]}$ with its symmetric closure, without affecting the truth-value of Φ on a structure $\mathbf{K}_q$; and note that any instance where this relation is not irreflexive is false on a structure $\mathbf{K}_q$.*

4.2 Auxiliary Structures

We call a structure a *labelled digraph* if it is over a signature that consists of a binary relation symbol E and zero or more unary relation symbols; we call a structure a *digraph* if it is over a signature consisting of just a binary relation symbol E. A digraph or labelled digraph is *symmetric* if it interprets E as a symmetric relation. In previous work [6], a way to encode a given labelled digraph $\mathbf{B}$ as a digraph $\mathbf{B}^*$ was given, and is as follows. Let $L_1, \ldots, L_n$ denote the unary symbols of the signature of $\mathbf{B}$. For each $b \in B$, define a gadget digraph $\mathbf{G}_b$ which has universe

$$G_b = \{b^s, b^c, b^d, b^{s1}, b^{t1}, b^{s2}, b^{t2}, \ldots, b^{sn}, b^{tn}, b^t\} \\ \cup \{b^{ui} \mid b \in L_i^{\mathbf{B}}\} \cup \{b^{vi} \mid b \in L_i^{\mathbf{B}}\}$$

and edge relation

$$\begin{aligned} E^{\mathbf{G}_b} = & \{(b^c, b^s), (b^c, b^d), (b^s, b^d), (b^d, b^{s1})\} \\ & \cup \{(b^{si}, b^{ti}) \mid i \in \{1, \ldots, n\}\} \\ & \cup \{(b^{ti}, b^{s(i+1)}) \mid i \in \{1, \ldots, n-1\}\} \\ & \cup \{(b^{tn}, b^t)\} \\ & \cup \{(b^{ui}, b^{si}), (b^{vi}, b^{ti}), (b^{vi}, b^{ui}) \mid b \in L_i^{\mathbf{B}}\}. \end{aligned}$$

For a subset $C \subseteq B$, we define $C^* = \bigcup_{b \in C} G_b$; the digraph $\mathbf{B}^*$ has universe B^* and edge relation

$$E^{\mathbf{B}^*} = (\bigcup_{b \in B} E^{\mathbf{G}_b}) \cup \{(b^t, b'^s) \mid (b, b') \in E^{\mathbf{B}}\}.$$

The key feature of this construction is that it preserves homomorphisms.

Lemma 9. *(follows from [6, Lemma 17]) Let $\mathbf{A}, \mathbf{B}$ be labelled digraphs over the same signature. There exists a homomorphism $\mathbf{A} \to \mathbf{B}$ if and only if there exists a homomorphism $h : \mathbf{A}^* \to \mathbf{B}^*$; moreover, when the latter condition holds, the image of h is of the form C^* where $C \subseteq B$.*

Tools for understanding the treewidth of structures of the form $\mathbf{B}^*$ are provided in the following lemmas, which relate the treewidth of such a structure to the treewidth of the structure $\mathbf{B}^+$, defined as follows. When $\mathbf{B}$ is a labelled digraph, the structure $\mathbf{B}^+$ has universe $B^+ = \{b^s, b^t \mid b \in B\}$ and edge relation

$$E^{\mathbf{B}^+} = \{(b^s, b^t) \mid b \in B\} \cup \{(b^t, b'^s) \mid (b, b') \in E^{\mathbf{B}}\}.$$

Lemma 10. *([6, Lemma 19]) Let $\mathbf{B}$ be a labelled digraph. It holds that $\mathsf{tw}(\mathbf{B}^*) \leq \max(\mathsf{tw}(\mathbf{B}^+), 5)$.*

Lemma 11. *Let $\mathbf{B}$ be a labelled digraph. It holds that $\mathsf{tw}(\mathbf{B}^+) \leq \mathsf{tw}(\mathbf{B}^*)$.*

Proof. Given a tree decomposition (T, β) of $G(\mathbf{B}^*)$, a tree decomposition (T, β') of $G(\mathbf{B}^+)$ having lower or equal width can be obtained by defining $\beta'(t) = f(\beta(t))$, where, for each $b \in B$, the mapping f sends $G_b \setminus \{b^t\}$ to b^s, and sends b^t to b^t. Clearly, it holds for each vertex t of T that $|\beta'(t)| \leq |\beta(t)|$. It is straightforward to verify that (T, β') is a tree decomposition of $G(\mathbf{B}^+)$; note that the connectivity condition is satisfied because $G_b \setminus \{b^t\}$ is connected in $G(\mathbf{B}^*)$, for each $b \in B$. $\square$

Lemma 12. *Let $\mathbf{B}$ be a symmetric labelled digraph. It holds that $\mathsf{tw}(\mathbf{B}) < \mathsf{tw}(\mathbf{B}^+)$.*

Consider an undirected graph on vertex set V. We say that two subsets C, C' of V *touch* if they have a vertex in common or there is an edge between them. A set of mutually touching connected vertex sets is a *bramble*. We say that a subset S of V *covers* a bramble $\mathcal{M}$ if it nontrivially intersects each set in $\mathcal{M}$. The *order* of a bramble $\mathcal{M}$ is the least number of vertices that covers it. We will use the *treewidth duality theorem* [9], which says that, for $k \geq 0$, a graph has tree-width $\geq k$ if and only if it has a bramble of order $> k$.

Proof. We prove that for each bramble $\mathcal{M}$ of $G(\mathbf{B})$, there exists a bramble $\mathcal{M}^+$ of $G(\mathbf{B}^+)$ of strictly higher order, which suffices by the tree-width duality theorem.

When C is a subset of B, we use C^s to denote the set $\{c^s \mid c \in C\}$, and C^t to denote the set $\{c^t \mid c \in C\}$. Let $\mathcal{M} = \{C_1, \dots, C_n\}$ be a bramble of $G(\mathbf{B})$. Define $\mathcal{M}^+$ to be the set system $\{C_1^s, C_1^t\} \cup \bigcup_{i \geq 2, i \in \underline{n}} \{C_i^s \cup (C_i \setminus C_1)^t, (C_i \setminus C_1)^s \cup C_i^t\}$. We claim that $\mathcal{M}^+$ is a bramble of $G(\mathbf{B}^+)$. We demonstrate this by verifying that each pair of distinct sets in $\mathcal{M}^+$ touch. The following cases are exhaustive, up to symmetry; here, i denotes an element of $\underline{n}$ with $i \geq 2$.

- C_1^s, C_1^t. These touch since for any $c_1 \in C_1$, we have $(c_1^s, c_1^t) \in E^{\mathbf{B}^+}$, and so $\{c_1^s, c_1^t\}$ is an edge in $G(\mathbf{B}^+)$.
- $C_i^s \cup (C_i \setminus C_1)^t, (C_i \setminus C_1)^s \cup C_i^t$. These touch since for any $c_i \in C_i$, we have $(c_i^s, c_i^t) \in E^{\mathbf{B}^+}$, and so $\{c_i^s, c_i^t\}$ is an edge in $G(\mathbf{B}^+)$.
- $C_1^s, C_i^s \cup (C_i \setminus C_1)^t$. If $C_1 \cap C_i$ is non-empty, then so is $C_1^s \cap C_i^s$. Otherwise, there is an edge in $G(\mathbf{B})$ between a vertex $c_1 \in C_1$ and a vertex $c_i \in C_i \setminus C_1$, and so $(c_1^t, c_i^s) \in E^{\mathbf{B}^+}$, implying that the two given sets touch in $G(\mathbf{B}^+)$.
- $C_1^s, (C_i \setminus C_1)^s \cup C_i^t$. If $C_1 \cap C_i$ is non-empty, then let $c \in C_1 \cap C_i$; we have $c^s \in C_1^s$, $c^t \in C_i^t$, and as $(c^s, c^t) \in E^{\mathbf{B}^+}$, the edge $\{c^s, c^t\}$ is present in $G(\mathbf{B}^+)$. Otherwise, there exist vertices $c_1 \in C_1$ and $c_i \in C_i \setminus C_1$ that are adjacent in $G(\mathbf{B})$, and so $(c_i^t, c_1^s) \in E^{\mathbf{B}^+}$, implying that $\{c_1^s, c_i^t\}$ is an edge in $G(\mathbf{B}^+)$.

It remains to show that the order of $\mathcal{M}^+$ is strictly higher than that of $\mathcal{M}$. To show this, we prove that for any cover S^+ of $\mathcal{M}^+$, there exists a cover S of $\mathcal{M}$ with $|S| < |S^+|$. Let S^+ be a cover of $\mathcal{M}^+$, and define S to be the subset of B obtained from removing the s, t superscripts from $S^+ \setminus C_1^t$. Since $C_1^t \in \mathcal{M}^+$, the cover S^+ must contain an element of C_1^t, from which it follows that $|S| < |S^+|$. We now verify that S covers $\mathcal{M}$. We have that S covers C_1, since S^+ covers C_1^s. When $i \geq 2$, we have that S covers C_i, since S^+ covers $C_i^s \cup (C_i \setminus C_1)^t$, which implies that S covers $C_i \cup (C_i \setminus C_1) = C_i$. $\square$

4.3 Reduction

Let

$$\forall y_1 \dots \forall y_m \exists x_1 \dots \exists x_n \phi_G$$

be an instance of Π_2-$\mathsf{QCSP}(\mathbf{K}_q)$. Relative to this instance, we define the following objects.

- Let τ be the signature
$$\{E\} \cup \{U_{y_1}, \dots, U_{y_m}\} \cup \{U_1, \dots, U_q\}$$
where the U_{y_i} and the U_j are unary relation symbols.
- We define the following formulas of signature τ.
$$\phi_K = (\bigwedge_{i \in \underline{q}} U_i(i)) \wedge (\bigwedge_{i,j \in \underline{q}, i \neq j} E(i,j))$$
For each $i \in \underline{m}$, $j \in \underline{q}$,
$$\lambda_{y_i \to j} = U_{y_i}(j)$$
$$\phi_{y_i \to j} = \lambda_{y_i \to j} \wedge \bigwedge_{k \in \underline{q}, k \neq j} (E(y_i, k) \wedge E(k, y_i))$$
For each $f : \{y_1, \dots, y_m\} \to \underline{q}$,
$$\lambda_f = \bigwedge_{i \in \underline{m}} \lambda_{y_i \to f(y_i)}$$
$$\phi_f = \bigwedge_{i \in \underline{m}} \phi_{y_i \to f(y_i)}$$

Observe that, for each mapping $f : \{y_1, \dots, y_m\} \to \underline{q}$,

$$\phi_f = \lambda_f \wedge \bigwedge_{i \in \underline{m}, k \in \underline{q}, f(y_i) \neq k} (E(y_i, k) \wedge E(k, y_i)), \qquad (\dagger)$$

up to a permutation of the conjuncts. In the sequel, we formally view ϕ_G as a formula of signature τ, so that, for instance, $\phi_G \wedge \phi_K \wedge \phi_f$ is a formula of signature τ, and $\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]$ is a structure of signature τ.

LEMMA 13. *Let $\forall y_1 \dots \forall y_m \exists x_1 \dots \exists x_n \phi_G$ be an instance of Π_2-$\mathsf{QCSP}(\mathbf{K}_q)$. If a mapping $f : \{y_1, \dots, y_m\} \to \underline{q}$ has an extension $f' : \{y_1, \dots, y_m, x_1, \dots, x_n\} \to \underline{q}$ such that $\mathbf{K}_q, f' \models \phi_G$, then the following hold.*

1. *$\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]^*$ maps homomorphically to $\mathbf{C}[\phi_K \wedge \lambda_f]^*$.*
2. *If $q \geq 5$, then $\mathsf{tw}(\mathbf{C}[\phi_K \wedge \lambda_f]^*) \leq q$.*

PROOF. For the first part, it is sufficient to prove that $\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]$ maps homomorphically to $\mathbf{C}[\phi_K \wedge \lambda_f]$; the statement then follows by Lemma 9. Note that the universes of the structures are $C[\phi_K \wedge \lambda_f] = \underline{q}$ and $C[\phi_G \wedge \phi_K \wedge \phi_f] = \{y_1, \dots, y_m, x_1, \dots, x_n\} \cup \underline{q}$.

Let $f' : \{y_1, \dots, y_m, x_1, \dots, x_n\} \to \underline{q}$ be an extension of $f : \{y_1, \dots, y_m\} \to \underline{q}$ such that $\mathbf{K}_q, f' \models \phi_G$. Let

$$h \colon \{y_1, \dots, y_m, x_1, \dots, x_n\} \cup \underline{q} \to \underline{q}$$

be the extension of f' defined by $h(j) = j$ for all $j \in \underline{q}$. We claim that h is a homomorphism from $\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]$ to $\mathbf{C}[\phi_K \wedge \lambda_f]$. By hypothesis, h maps homomorphically $\mathbf{C}[\phi_G]$ into $\mathbf{C}[\phi_K]$. Clearly, h maps homomorphically $\mathbf{C}[\phi_K \wedge \lambda_f]$ into $\mathbf{C}[\phi_K \wedge \lambda_f]$. By ($\dagger$), it suffices to show that h is a homomorphism from $\mathbf{C}[\bigwedge_{i \in \underline{m}, k \in \underline{q}, f(y_i) \neq k} (E(y_i, k) \wedge E(k, y_i))]$ to $\mathbf{C}[\phi_K]$. Suppose that the tuples $(y_i, k), (k, y_i)$ occur in the first structure. Then, by definition, $f(y_i) \neq k$. Therefore $(h(y_i), h(k)) = (f(y_i), k) \in E^{\mathbf{C}[\phi_K]}$, and $(h(k), h(y_i)) = (k, f(y_i)) \in E^{\mathbf{C}[\phi_K]}$.

For the second part, assume $q \geq 5$. Then, by Lemma 10, it is sufficient to prove that $\mathsf{tw}(\mathbf{C}[\phi_K \wedge \lambda_f]^+) \leq q$. We establish $\mathsf{tw}(\mathbf{C}[\phi_K \wedge \lambda_f]^+) \leq q$ by providing a tree decomposition of width q of $\mathbf{C}[\phi_K \wedge \lambda_f]^+$. It is straightforward to check that a path of q vertices $v_1, \dots, v_q$, where the bag on v_j is $\{k^s \mid k \in \underline{q}\} \cup \{j^t\}$ for all $j \in \underline{q}$, gives the required tree decomposition. $\square$

LEMMA 14. *There exists an algorithm running in polynomial time that, given an instance*

$$\phi = \forall y_1 \dots \forall y_m \exists x_1 \dots \exists x_n \phi_G$$

of Π_2-$\mathsf{QCSP}(\mathbf{K}_q)$, computes two sentences $\phi', \phi'' \in \mathrm{EP}_{\{E\}}$, where E is a binary relation symbol, such that ϕ' is logically equivalent to the disjunctive form

$$\bigvee_{f : \{y_1, \dots, y_m\} \to \underline{q}} Q[\mathbf{C}[\phi_K \wedge \lambda_f]^*], \qquad \text{(F1)}$$

ϕ'' is logically equivalent to the disjunctive form

$$\bigvee_{f:\{y_1,\ldots,y_m\}\to \underline{q}} Q[\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]^*], \qquad \text{(F2)}$$

and the following hold:

1. *The disjunctive forms (F1) and (F2) are irredundant.*
2. *For all $f : \{y_1,\ldots,y_m\} \to \underline{q}$, $\mathbf{C}[\phi_K \wedge \lambda_f]^*$ maps homomorphically to $\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]^*$; and consequently, $\phi'' \models \phi'$.*
3. *If $\mathbf{K}_q \models \phi$, then ϕ' and ϕ'' are logically equivalent.*

Proof. Let $\phi = \forall y_1 \ldots \forall y_m \exists x_1 \ldots \exists x_n \phi_G$ be an instance of Π_2-$\mathsf{QCSP}(\mathbf{K}_q)$. The algorithm, given ϕ, constructs in polynomial-time the existential positive sentences

$$\phi' = Q[\mathbf{C}[\phi_K]^*] \wedge \bigwedge_{i\in\underline{m}} \bigvee_{j\in\underline{q}} Q[\mathbf{C}[\lambda_{y_i\to j}]^*];$$

$$\phi'' = Q[\mathbf{C}[\phi_G]^*] \wedge Q[\mathbf{C}[\phi_K]^*] \wedge \bigwedge_{i\in\underline{m}} \bigvee_{j\in\underline{q}} Q[\mathbf{C}[\phi_{y_i\to j}]^*].$$

We claim that:

$$\phi' \equiv \bigvee_{f:\{y_1,\ldots,y_m\}\to \underline{q}} Q[\mathbf{C}[\phi_K \wedge \lambda_f]^*]]; \qquad \text{(G1)}$$

$$\phi'' \equiv \bigvee_{f:\{y_1,\ldots,y_m\}\to \underline{q}} Q[\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]^*]. \qquad \text{(G2)}$$

It is sufficient to observe the following logical equivalences. For (G1),

$$\begin{aligned}
\phi' &= Q[\mathbf{C}[\phi_K]^*] \wedge \bigwedge_{i\in\underline{m}} \bigvee_{j\in\underline{q}} Q[\mathbf{C}[\lambda_{y_i\to j}]^*] \\
&\equiv Q[\mathbf{C}[\phi_K]^*] \wedge \bigvee_{f:\{y_1,\ldots,y_m\}\to \underline{q}} Q[\mathbf{C}[\lambda_f]^*] \\
&\equiv \bigvee_{f:\{y_1,\ldots,y_m\}\to \underline{q}} (Q[\mathbf{C}[\phi_K]^*] \wedge Q[\mathbf{C}[\lambda_f]^*]) \\
&\equiv \bigvee_{f:\{y_1,\ldots,y_m\}\to \underline{q}} Q[(\mathbf{C}[\phi_K] \cup \mathbf{C}[\lambda_f])^*] \\
&\equiv \bigvee_{f:\{y_1,\ldots,y_m\}\to \underline{q}} Q[\mathbf{C}[\phi_K \wedge \lambda_f]^*].
\end{aligned}$$

For (G2), we similarly have

$$\begin{aligned}
\phi'' &= Q[\mathbf{C}[\phi_G]^*] \wedge Q[\mathbf{C}[\phi_K]^*] \wedge \bigwedge_{i\in\underline{m}} \bigvee_{j\in\underline{q}} Q[\mathbf{C}[\phi_{y_i\to j}]^*] \\
&\equiv Q[\mathbf{C}[\phi_G]^*] \wedge Q[\mathbf{C}[\phi_K]^*] \wedge \bigvee_{f:\{y_1,\ldots,y_m\}\to \underline{q}} Q[\mathbf{C}[\phi_f]^*] \\
&\equiv \bigvee_{f:\{y_1,\ldots,y_m\}\to \underline{q}} (Q[\mathbf{C}[\phi_G]^*] \wedge Q[\mathbf{C}[\phi_K]^*] \wedge Q[\mathbf{C}[\phi_f]^*]) \\
&\equiv \bigvee_{f:\{y_1,\ldots,y_m\}\to \underline{q}} Q[(\mathbf{C}[\phi_G] \cup \mathbf{C}[\phi_K] \cup \mathbf{C}[\phi_f])^*] \\
&\equiv \bigvee_{f:\{y_1,\ldots,y_m\}\to \underline{q}} Q[\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]^*].
\end{aligned}$$

To prove the stated properties, we observe preliminarily that ϕ_f contains all conjuncts of λ_f by (†), thus $\mathbf{C}[\phi_K \wedge \lambda_f]$ is a substructure of $\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]$.

We prove the first property. By Proposition 1, it is sufficient to check that if f and g are distinct mappings from $\{y_1,\ldots,y_m\}$ to $\underline{q}$, then $\mathbf{C}[\phi_K \wedge \lambda_f]$ does not map homomorphically to $\mathbf{C}[\phi_K \wedge \lambda_g]$. By Lemma 9, this implies that $\mathbf{C}[\phi_K \wedge \lambda_f]^*$ does not map homomorphically to $\mathbf{C}[\phi_K \wedge \lambda_g]^*$, which settles irredundancy of (F1); in turn it follows, by the substructure observation, that $\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]^*$ does not map homomorphically to $\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_g]^*$, which settles irredundancy of (F2). Assume for a contradiction that h maps $\mathbf{C}[\phi_K \wedge \lambda_f]$ homomorphically to $\mathbf{C}[\phi_K \wedge \lambda_g]$. By definition of ϕ_K, it holds that $U_i^{\mathbf{C}[\phi_K]} = \{i\}$ for all $i \in \underline{q}$, therefore h acts identically on $\underline{q}$. Let $j \in \underline{m}$ be such that $f(y_j) = k \neq k' = g(y_j)$. By definition of λ_f and λ_g, it holds that $U_{y_j}^{\mathbf{C}[\phi_f]} = \{k\}$ and $U_{y_j}^{\mathbf{C}[\phi_g]} = \{k'\}$. Therefore, $h(k) = k'$, a contradiction.

We prove the second property. It suffices to prove the first part; that $\phi'' \models \phi'$ is then a consequence by appeal to Proposition 1. Let f be any mapping from $\{y_1,\ldots,y_m\}$ to $\underline{q}$. By the observation that $\mathbf{C}[\phi_K \wedge \lambda_f]$ is a substructure of $\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]$, we have that $\mathbf{C}[\phi_K \wedge \lambda_f]$ maps homomorphically to $\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]$; the statement then follows by Lemma 9.

We prove the third property. Assume $\mathbf{K}_q \models \phi$. Let f be any mapping of $\{y_1,\ldots,y_m\}$ to $\underline{q}$. Then, there exists an extension

$$f' : \{y_1,\ldots,y_m,x_1,\ldots,x_n\} \to \underline{q}$$

of f such that $\mathbf{K}_q, f' \models \phi_G$. Then, by Lemma 13(1), $\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]^*$ maps homomorphically to $\mathbf{C}[\phi_K \wedge \lambda_f]^*$, which implies that $Q[\mathbf{C}[\phi_K \wedge \lambda_f]^*] \models Q[\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]^*]$. Therefore, by Lemma 4(1), $\phi' \models \phi''$. Then, by the second property proved above, $\phi' \equiv \phi''$. $\Box$

Lemma 15. *Let $\forall y_1 \ldots \forall y_m \exists x_1 \ldots \exists x_n \phi_G$ be an instance of Π_2-$\mathsf{QCSP}(\mathbf{K}_q)$. Let $f : \{y_1,\ldots,y_m\} \to \underline{q}$ be a mapping, and suppose that $\mathsf{tw}(\mathsf{core}(\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]^*)) \leq q$. Then, f has an extension $f' : \{y_1,\ldots,y_m,x_1,\ldots,x_n\} \to \underline{q}$ such that $\mathbf{K}_q, f' \models \phi_G$.*

In the proof, we will use the following notation: when $\mathbf{B}$ is a structure on signature σ, and $\sigma' \subseteq \sigma$, use $\mathsf{red}_{\sigma'}(\mathbf{B})$ to denote the *reduct* of $\mathbf{B}$ on σ', that is, the structure on σ' naturally obtained from $\mathbf{B}$ by forgetting the interpretations of the symbols not in σ'.

Proof. Set $\mathbf{A} = \mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]$. Since each core of $\mathbf{A}^*$ is the image of an endomorphism of $\mathbf{A}^*$, then by Lemma 9, each core of $\mathbf{A}^*$ has universe of the form S^* where $S \subseteq A$. Let $S \subseteq A$ be a subset with this property, and let $\mathbf{S}$ be the substructure of $\mathbf{A}$ induced on S. By assumption, $\mathsf{tw}(\mathbf{S}^*) \leq q$. By Lemma 11, $\mathsf{tw}(\mathbf{S}^+) \leq q$. By Lemma 12, $\mathsf{tw}(\mathbf{S}) < q$. It follows that $\mathsf{red}_E(\mathbf{S})$ has a homomorphism to $\mathbf{K}_q$ because, by Remark 8 and construction, $\mathsf{red}_E(\mathbf{S})$ is irreflexive. Since $\mathbf{A}$ has a homomorphism to $\mathbf{S}$, we have that $\mathsf{red}_E(\mathbf{A})$ has a homomorphism to $\mathsf{red}_E(\mathbf{S})$, and by transitivity of the homomorphism relation, we have that $\mathsf{red}_E(\mathbf{A})$ has a homomorphism h to $\mathbf{K}_q$. Observe that $\mathsf{red}_E(\mathbf{A}) = \mathbf{C}[\phi_G \wedge (\bigwedge_{i,j\in\underline{q},i\neq j} E(i,j)) \wedge \bigwedge_{i\in\underline{m},k\in\underline{q},f(y_i)\neq k}(E(y_i,k) \wedge E(k,y_i))]$.

By relabelling the elements of $\mathbf{K}_q$ if necessary, it can be assumed that h is the identity map on $\underline{q}$. Therefore, since

$\mathbf{K}_q, h \models \bigwedge_{i\in \underline{m}, k\in \underline{q}, f(y_i)\neq k}(E(y_i,k)\wedge E(k,y_i))$, we have that h is an extension of f. Indeed, for all $i \in \underline{m}$ and $k \in \underline{q}$ such that $f(y_i) \neq k$, we have $(h(y_i), h(k)) = (h(y_i), k) \in E^{\mathbf{K}_q}$, which implies $h(y_i) \neq k$.

Since $\mathbf{K}_q, h \models \phi_G$, we obtain the result. □

4.4 Hardness Result

THEOREM 16. *Let σ be a signature that contains a relation symbol E of binary arity. For each $k \geq 6$, the problem EP^k_σ-EXPR is Π^p_2-hard.*

PROOF. Assume $q \geq 5$. We show that there is a reduction from Π_2-$\mathsf{QCSP}(\mathbf{K}_q)$ to EP^{q+1}_σ-EXPR, where $\sigma = \{E\}$ and E is a binary relation symbol; this suffices by Proposition 7.

Let $\phi = \forall y_1 \ldots \forall y_m \exists x_1 \ldots \exists x_n \phi_G$ be an instance of the problem Π_2-$\mathsf{QCSP}(\mathbf{K}_q)$. The reduction uses the algorithm in Lemma 14 to compute in polynomial-time the sentence $\phi'' \in \mathrm{EP}_\sigma$ defined there. We prove that $\mathbf{K}_q \models \phi$ if and only if ϕ'' is logically equivalent to a sentence in EP^{q+1}_σ.

Assume that $\mathbf{K}_q \models \phi$. By Lemma 14(3), we have that ϕ'' is logically equivalent to ϕ'. Now look at the formula shown to be logically equivalent to ϕ' in that lemma (Lemma 14). For each $f : \{y_1, \ldots, y_m\} \to \underline{q}$, by Lemma 13(2) and the assumption that $q \geq 5$, we have $\mathsf{tw}(\mathbf{C}[\phi_K \wedge \lambda_f]^*) \leq q$, and therefore by [8, Theorem 5], $Q[\mathbf{C}[\phi_K \wedge \lambda_f]^*]$ is logically equivalent to a primitive positive sentence in PP^{q+1}. Therefore, ϕ'' is logically equivalent to a sentence in EP^{q+1}_σ.

Assume that $\mathbf{K}_q \not\models \phi$. Suppose that

$$f\colon \{y_1, \ldots, y_m\} \to \underline{q}$$

is a mapping such that for all mappings

$$f'\colon \{y_1, \ldots, y_m, x_1, \ldots, x_n\} \to \underline{q}$$

extending f it holds that $\mathbf{K}_q, f' \not\models \phi_G$. Then, by Lemma 15, $\mathsf{tw}(\mathsf{core}(\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]^*)) > q$. Therefore, by [8, Theorem 5], $Q[\mathbf{C}[\phi_G \wedge \phi_K \wedge \phi_f]^*]$ is not logically equivalent to a primitive positive sentence in PP^{q+1}. Since we have, by Lemma 14(1), the disjunctive form in (F2) is irredundant, so by Lemma 4(3), ϕ'' is a "No" instance of EP^{q+1}_σ-EXPR. □

5. POSITIVE LOGIC

In this section, we prove that expressibility is undecidable in positive logic. Recall the *satisfiability problem* for first-order logic: Given a first-order sentence ϕ in prefix normal form, is there a structure $\mathbf{A}$ such that $\mathbf{A} \models \phi$?

In the rest of this section, τ is fixed and denotes a signature that contains countably many unary relation symbols, and a relation symbol of binary arity, say R. The *Kahr class* is the class of equality-free first-order τ-sentences in prefix normal form with quantifier prefix equal to $\forall x \exists y \forall z$. We need the following classical result [2, Theorem 3.1.1].

THEOREM 17. *The Kahr class has an undecidable satisfiability problem.*

The following lemma is an adaptation of the reduction in [7, Theorem 7] proving that positive logic has an undecidable entailment problem.

LEMMA 18. *Let $\rho = \tau \cup \{R_c\}$ where R_c is a relation symbol of binary arity not in τ. The following problem is undecidable: Given a pair (ψ, ψ') of sentences in PFO^3_ρ, does ψ entail ψ'?*

PROOF. We reduce from the satisfiability problem for the Kahr class, undecidable by Theorem 17. Let ϕ be a sentence in the Kahr class. Note that, by definition, $\phi \in \mathrm{FO}^3_\tau$.

We reduce ϕ to a pair $(\phi' \wedge \chi, \chi')$ of sentences in PFO^3_ρ such that ϕ is satisfiable if and only if $\phi' \wedge \chi \not\models \chi'$, which implies the stated undecidability result. The sentences ϕ', χ, and χ' are defined as follows.

Let $U^1, \ldots, U^r$ be the unary symbols occurring in ϕ, and let $U^1_c, \ldots, U^r_c$ be unary symbols in τ not occurring in ϕ. We define the following sentences:

- ϕ' denotes the sentence obtained by taking the negation normal form of ϕ (where the negation connective is applied only to relation symbols), and replacing literals of the form $\neg U^i x$ by $U^i_c x$ for all $i \in \underline{r}$, and literals of the form $\neg Rxy$ by $R_c xy$;
- $\chi = \forall xy(Rxy \vee R_c xy) \wedge \bigwedge_{i\in\underline{r}} \forall x(U^i x \vee U^i_c x)$;
- $\chi' = \exists xy(Rxy \wedge R_c xy) \vee \bigvee_{i\in\underline{r}} \exists x(U^i x \wedge U^i_c x)$.

We claim that ϕ is satisfiable if and only if $(\phi' \wedge \chi) \wedge \neg\chi'$ is satisfiable, which is in turn equivalent to $\phi' \wedge \chi \not\models \chi'$. Noticing that ϕ', χ, and χ' are in PFO^3_ρ by construction, the statement is settled. To prove the claim, it is sufficient to observe that $\chi \wedge \neg\chi'$ is logically equivalent to

$$\forall xy(R_c xy \leftrightarrow \neg Rxy) \wedge \bigwedge_{i\in\underline{r}} \forall x(U^i_c x \leftrightarrow \neg U^i x),$$

from which, by the construction of ϕ', it follows that ϕ is satisfiable if and only if $(\chi \wedge \neg\chi') \wedge \phi'$ is satisfiable. □

To prove the main result of this section, we prepare the following terminology and notation. Let ϕ be a FO-sentence, and let U be a unary relation symbol. The *U-relativization* ϕ^U of ϕ is defined inductively as follows: if ϕ is an atom, then $\phi^U = \phi$; if $\phi = \neg\psi$, then $\phi^U = \neg(\psi^U)$; if $\phi = \psi_1 \wedge \psi_2$, then $\phi^U = \psi_1^U \wedge \psi_2^U$; if $\phi = \psi_1 \vee \psi_2$, then $\phi^U = \psi_1^U \vee \psi_2^U$; if $\phi = \forall x\psi$, then $\phi^U = \forall x(Ux \to \psi)$; if $\phi = \exists x\psi$, then $\phi^U = \exists x(Ux \wedge \psi)$. By the relativization lemma, for all structures $\mathbf{A}$ on a signature including U such that $U^{\mathbf{A}}$ is nonempty, it holds that $\mathbf{A}|_{U^{\mathbf{A}}} \models \phi$ if and only if $\mathbf{A} \models \phi^U$ [10, Lemma 2.4].

Let U and U_c be two distinct unary symbols. Let ϕ be a FO-sentence. Let ϕ^{U,U_c} denote the FO-sentence defined inductively as follows: if ϕ is an atom, then $\phi^{U,U_c} = \phi$; if $\phi = \neg\psi$, then $\phi^{U,U_c} = \neg(\psi^{U,U_c})$; if $\phi = \psi_1 \wedge \psi_2$, then $\phi^{U,U_c} = \psi_1^{U,U_c} \wedge \psi_2^{U,U_c}$; if $\phi = \psi_1 \vee \psi_2$, then $\phi^{U,U_c} = \psi_1^{U,U_c} \vee \psi_2^{U,U_c}$; if $\phi = \forall x\psi$, then $\phi^{U,U_c} = \forall x(U_c x \vee \psi)$; if $\phi = \exists x\psi$, then $\phi^{U,U_c} = \exists x(Ux \wedge \psi)$.

We preliminarily observe the following easy fact. If $\mathbf{A}$ is a structure on a signature including U and U_c and $\mathbf{A} \models \forall x(\neg Ux \leftrightarrow U_c x)$, then $\mathbf{A} \models \phi^U$ if and only if $\mathbf{A} \models \phi^{U,U_c}$. Intuitively, we simulate in the positive fragment the relativization of ϕ to U, which is not positive in general.

THEOREM 19. *Let σ be a signature that contains countably many unary relation symbols and three relation symbols of binary arity. For each $k \geq 3$, the problem PFO^k_σ-EXPR is undecidable.*

PROOF. We reduce from the undecidable problem that is given in Lemma 18, asking whether for a pair (ϕ, ψ) of PFO^3_ρ-sentences, $\phi \models \psi$; here, ρ is a signature that contains countably many unary relation symbols, and two relation symbols of binary arity.

Let E be a binary relation symbol not in ρ (thus neither in ϕ nor in ψ), and let $\sigma = \rho \cup \{E\}$. The reduction maps (ϕ, ψ) to the instance χ of PFO^3_σ-EXPR, defined as follows.

Let U and U_c be unary relation symbols in σ not occurring in ϕ or ψ. We define the following PFO_σ-sentences:

- $\theta = \exists x_1 x_2 x_3 x_4 (\bigwedge_{i,j \in \underline{4}, i \neq j} E x_i x_j)$.
- $\alpha = \exists x U x \wedge \forall x (U_c x \vee U x) \wedge \phi^{U,U_c}$.
- $\beta = \exists x (U_c x \wedge U x) \vee \psi^{U,U_c}$.
- $\chi = \alpha \wedge (\beta \vee \theta^{U_c})$.

We claim that $\phi \models \psi$ if and only if χ is logically equivalent to a PFO^3-sentence.

($\Rightarrow$) Assume $\phi \models \psi$. We claim that $\alpha \models \beta$. By the claim, $\chi \equiv \alpha$. Since $\phi \in \text{PFO}^3$, by construction $\phi^{U,U_c} \in \text{PFO}^3$. By possibly renaming x if it does not occur in ϕ, we conclude that α is logically equivalent to a PFO^3-sentence.

We prove the claim. Let $\mathbf{A}$ be any σ-structure. We distinguish two cases.

Assume that $U^{\mathbf{A}} = \emptyset$, or $U^{\mathbf{A}} \cup U_c^{\mathbf{A}} \neq A$, or $U^{\mathbf{A}} \cap U_c^{\mathbf{A}} \neq \emptyset$. If $U^{\mathbf{A}} = \emptyset$ or $U^{\mathbf{A}} \cup U_c^{\mathbf{A}} \neq A$, then $\mathbf{A} \not\models \alpha$. If $U^{\mathbf{A}} \cap U_c^{\mathbf{A}} \neq \emptyset$, then $\mathbf{A} \models \beta$.

Otherwise, $U^{\mathbf{A}} \neq \emptyset$, $U^{\mathbf{A}} \cup U_c^{\mathbf{A}} = A$, and $U^{\mathbf{A}} \cap U_c^{\mathbf{A}} = \emptyset$. In this case, by the preliminarily observed fact, we have that $\mathbf{A} \models \phi^U$ if and only if $\mathbf{A} \models \phi^{U,U_c}$, and $\mathbf{A} \models \psi^U$ if and only if $\mathbf{A} \models \psi^{U,U_c}$. Assume that $\mathbf{A} \models \alpha$. Then, $\mathbf{A} \models \phi^{U,U_c}$, and by the above, $\mathbf{A} \models \phi^U$. By the relativization lemma, $\mathbf{A}|_{U^{\mathbf{A}}} \models \phi$. Then, by hypothesis, $\mathbf{A}|_{U^{\mathbf{A}}} \models \psi$, which implies that $\mathbf{A} \models \psi^U$ by the relativization lemma, so by the above, $\mathbf{A} \models \psi^{U,U_c}$. Therefore, $\mathbf{A} \models \beta$.

($\Leftarrow$) Assume $\phi \not\models \psi$. Let $\mathbf{A}$ be a ρ-structure such that $\mathbf{A} \models \phi$ and $\mathbf{A} \not\models \psi$. Assume without loss of generality that $A \cap \underline{4} = \emptyset$. Define two σ-structures $\mathbf{A}_0$ and $\mathbf{A}_1$ as follows.

- $A_0 = A \cup \underline{3}$, $U^{\mathbf{A}_0} = A$, $U_c^{\mathbf{A}_0} = \underline{3}$, $E^{\mathbf{A}_0} = \{(i,j) \mid i,j \in \underline{3}, i \neq j\}$, and $R^{\mathbf{A}_0} = R^{\mathbf{A}}$ for all $R \in \sigma$.
- $A_1 = A \cup \underline{4}$, $U^{\mathbf{A}_1} = A$, $U_c^{\mathbf{A}_1} = \underline{4}$, $E^{\mathbf{A}_1} = \{(i,j) \mid i,j \in \underline{4}, i \neq j\}$, and $R^{\mathbf{A}_1} = R^{\mathbf{A}}$ for all $R \in \sigma$.

We observe that, for all r, the Duplicator wins the r-round k-pebble Ehrenfeucht game on $\mathbf{A}_0$ and $\mathbf{A}_1$ by playing the identity on A and maintaining a partial isomorphism between $\underline{3}$ and $\underline{4}$, which is possible because the players have only 3 pebbles. Thus, $\mathbf{A}_0$ and $\mathbf{A}_1$ model the same FO^3-sentences [15, Theorem 6.10]. We claim that $\mathbf{A}_0 \not\models \chi$ but $\mathbf{A}_1 \models \chi$. Therefore, by the claim and the observation, χ is not logically equivalent to a FO^3-sentence.

To prove the claim, we first prove $\mathbf{A}_0 \not\models \beta \vee \theta^{U_c}$, so that $\mathbf{A}_0 \not\models \chi$. This follows from the following three facts:

- $\mathbf{A}_0 \not\models \exists x (U_c x \wedge U x)$, by construction.
- $\mathbf{A}_0 \not\models \psi^{U,U_c}$. Indeed, by construction of $\mathbf{A}_0$, it holds that ψ^{U,U_c} behaves in $\mathbf{A}_0$ as ψ^U. Moreover, by the relativization lemma, $\mathbf{A}_0 \models \psi^U$ if and only if $\mathbf{A}_0|_{U^{\mathbf{A}_0}} \models \psi$ that is, if and only if $\mathbf{A} \models \psi$, which is false by hypothesis.
- $\mathbf{A}_0 \not\models \theta^{U_c}$. We have that $\mathbf{A}_0 \models \theta$ if and only if $\mathbf{A}_0|_{U_c^{\mathbf{A}_0}} \models \theta$, which is false by construction.

We next prove $\mathbf{A}_1 \models \alpha \wedge \theta^{U_c}$, so that $\mathbf{A}_1 \models \chi$. This follows from the following three facts:

- $\mathbf{A}_1 \models \exists x U x \wedge \forall x (U_c x \vee U x)$, by construction.
- $\mathbf{A}_1 \models \phi^{U,U_c}$. Similarly, $\mathbf{A}_1 \models \phi^{U,U_c}$ if and only if $\mathbf{A}_1 \models \phi^U$ if and only if $\mathbf{A}_1|_{U^{\mathbf{A}_1}} \models \phi$, that is, if and only if $\mathbf{A} \models \phi$, which is true by hypothesis.
- $\mathbf{A}_1 \models \theta^{U_c}$. In a similar fashion, $\mathbf{A}_1 \models \theta^{U_c}$ if and only if $\mathbf{A}_1|_{U_c^{\mathbf{A}_1}} \models \theta$, which is true by construction.

The statement is proved. $\square$

6. ACKNOWLEDGEMENTS

The first author was supported by ERC Starting Grant (Complex Reason, 239962) and FWF Austrian Science Fund (Parameterized Compilation, P26200).

The second author was supported by the Spanish Project FORMALISM (TIN2007-66523), by the Basque Government Project S-PE12UN050(SAI12/219), and by the University of the Basque Country under grant UFI11/45.

7. REFERENCES

[1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases.* Addison-Wesley, 1995.

[2] E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem.* Springer, 1997.

[3] F. Börner, A. A. Bulatov, H. Chen, P. Jeavons, and A. A. Krokhin. The Complexity of Constraint Satisfaction Games and QCSP. *Inform. Comput.*, 207(9):923–944, 2009.

[4] A. K. Chandra and P. M. Merlin. Optimal Implementation of Conjunctive Queries in Relational Data Bases. In *Proceedings of STOC'77*, pages 77–90, 1977.

[5] H. Chen. A Rendezvous of Logic, Complexity, and Algebra. *ACM Computing Surveys*, 42(1), 2009.

[6] H. Chen. On the Complexity of Existential Positive Queries. *CoRR*, abs/1206.3902, 2012.

[7] H. Chen, F. Madelaine, and B. Martin. Quantified Constraints and Containment Problems. In *Proceedings of LICS'08*, 2008.

[8] V. Dalmau, P. G. Kolaitis, and M. Y. Vardi. Constraint Satisfaction, Bounded Treewidth, and Finite-Variable Logics. In *Proceedings of CP'02*, 2002.

[9] R. Diestel. *Graph Theory, 4th Edition.* Springer, 2012.

[10] H.-D. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic.* Springer, 1984.

[11] J. Flum and M. Grohe. *Parameterized Complexity Theory.* Springer, 2006.

[12] E. Grädel. Finite model theory and descriptive complexity. In *Finite Model Theory and Its Applications*, pages 125–230. Springer, 2007.

[13] M. Grohe, T. Schwentick, and L. Segoufin. When is the Evaluation of Conjunctive Queries Tractable? In *Proceedings of STOC'2001*, 2001.

[14] P. Hell and J. Nešetřil. The Core of a Graph. *Discrete Math.*, 109:117–126, 1992.

[15] N. Immerman. *Descriptive complexity.* Springer, 1999.

[16] Y. Sagiv and M. Yannakakis. Equivalences Among Relational Expressions with the Union and Difference Operators. *J. ACM*, 27(4):633–655, 1980.

[17] J. D. Ullman. *Database and Knowledge-Base Systems*, Volumes I and II. Computer Science Press, 1989.

[18] M. Y. Vardi. On the Complexity of Bounded-Variable Queries. In *Proceedings of PODS'95*, pages 266–276, 1995.