

Block-Sorted Quantified Conjunctive Queries

Hubie Chen^{1*} and Dániel Marx^{2**}

¹ Universidad del País Vasco and IKERBASQUE, E-20018 San Sebastián, Spain

² Computer and Automation Research Institute, Hungarian Academy of Sciences (MTA SZTAKI), Budapest, Hungary

Abstract. We study the complexity of model checking in quantified conjunctive logic, that is, the fragment of first-order logic where both quantifiers may be used, but conjunction is the only permitted connective. In particular, we study block-sorted queries, which we define to be prenex sentences in multi-sorted relational first-order logic where two variables having the same sort must appear in the same quantifier block. We establish a complexity classification theorem that describes precisely the sets of block-sorted queries of bounded arity on which model checking is fixed-parameter tractable. This theorem strictly generalizes, for the first time, the corresponding classification for existential conjunctive logic (which is known and due to Grohe) to a logic in which both quantifiers are present.

1 Introduction

Model checking, the problem of deciding if a logical sentence holds on a structure, is a fundamental computational task that appears in many guises throughout computer science. Witness its appearance in areas such as computational logic, verification, artificial intelligence, constraint satisfaction, and computational complexity. The case where one wishes to evaluate a first-order sentence on a finite structure is a problem of principal interest in database theory and is the topic of this article. This problem is well-known to be quite intractable in general: it is PSPACE-complete.

As has been articulated in the literature [7], 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 formula, so long as

* Research 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.

** Research supported by the European Research Council (ERC) grant “PARAMTIGHT: Parameterized complexity and the search for tight complexity results,” reference 280152.

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. We follow this paradigm, and focus the discussion on this form of tractability.

First-order model checking is intractable even if one restricts the connectives and quantifiers permitted; for instance, model checking of *existential conjunctive queries*, by which we mean sentences formed using atoms, conjunction (\wedge), and existential quantification (\exists), is well-known to be intractable (it is NP-complete). Thus, a typical way to gain insight into which sentences exhibit tractable behavior is to consider model checking relative to a set Φ of sentences. In the context of existential conjunctive logic, there is a mature understanding of sentence sets. It was proved by Grohe [6] that when Φ is a set of existential conjunctive queries having bounded arity, model checking on Φ is fixed-parameter tractable if there is a constant $k \geq 1$ such that each sentence in Φ is logically equivalent to one whose treewidth is bounded above by k , and is intractable otherwise (under a standard assumption from parameterized complexity). The treewidth of a conjunctive sentence (in prenex form) is measured here via the graph on the sentence’s variables wherein two variables are adjacent if they co-occur in an atom.

An important precursor to Grohe’s theorem was the complexity classification of graph sets for existential conjunctive logic. Grohe, Schwentick, and Segoufin [7] defined model checking relative to a graph set \mathcal{G} as the problem of deciding, given a structure and an existential conjunctive query whose graph is in \mathcal{G} , whether or not the query is true on the structure; they showed that the problem is fixed-parameter tractable when \mathcal{G} has bounded treewidth, and intractable otherwise. In this paper, we restrict our attention to queries of bounded arity (the case of unbounded arity leads to a different theory, where complexity may depend on the choice of representation of relations [3, 8]). For bounded-arity structures, this result is *coarser* than Grohe’s theorem, as it can be taken as a classification of sentence sets Φ that obey the closure property that if a sentence is in Φ , then all sentences having the same graph are also in Φ ; in contrast, Grohe’s theorem classifies arbitrary sentence sets.

This graph classification was recently generalized to quantified conjunctive logic, wherein both quantifiers (\forall, \exists) are permitted in addition to conjunction (\wedge). Define a *prefixed graph* to be a quantifier prefix $Q_1v_1 \dots Q_nv_n$ paired with a graph on the variables $\{v_1, \dots, v_n\}$; each quantified conjunctive query in prenex form can naturally be mapped to a prefixed graph, by simply taking the quantifier prefix of the query along with the graph of the quantifier-free, conjunctive portion of the query. Chen and Dalmau [2] defined a width measure for prefixed graphs, which generalizes treewidth, and proved that model checking on a set of prefixed graphs is fixed-parameter tractable if the set has bounded width, and intractable otherwise. This result generalizes the graph classification by Grohe, Schwentick, and Segoufin, and provides a unified view of this classification as

well as earlier complexity results [5] on quantified conjunctive logic. Note, however, that the present result is incomparable to Grohe’s result: Grohe’s result is on arbitrary sentence sets in a less expressive logic, while the result of Chen and Dalmau considers sentences in more expressive logic, but considers them from the coarser graph-based viewpoint, that is, it classifies sentence sets obeying the (analog of the) described closure property.

In this article, we present a veritable generalization of Grohe’s theorem in quantified conjunctive logic. In the bounded-arity case, our theorem naturally unifies together both Grohe’s theorem and the classification of prefixed graphs in quantified conjunctive logic. The sentences studied by our theorem are of the following type. Define a *block-sorted query* to be a quantified conjunctive sentence in multi-sorted, relational first-order logic where two variables having the same sort must occur in the same quantifier block. This class of sentences includes each sentence having a sort for each quantifier block. As an example, consider the sentence

$$\exists x_1, x_2 \forall y_1, y_2, y_3 \exists z_1, z_2 \\ R(x_1, y_1) \wedge R(x_2, y_3) \wedge S(x_2, y_2, y_3, z_1) \wedge S(x_1, y_1, y_2, z_2) \wedge T(x_1, x_2, y_2),$$

where the variables x_i have the same sort e , the variables y_i have the same sort u , and the variables z_i have the same sort e' ; the arities of the relation symbols R , S , and T are eu , $euue'$, and eeu , respectively. The definitions impose that a structure \mathbf{B} on which such a sentence can be evaluated needs to provide a domain B_s (which is a set) for each sort; quantifying a variable of sort s is performed over the domain B_s . (See the next section for the precise formalization that is studied.)

Our main theorem is the classification of block-sorted queries. We show how to computably derive from each query a second logically equivalent query, and demonstrate that, for a bounded-arity set of block-sorted queries, model checking is fixed-parameter tractable if the width of the derived queries is bounded (with respect to the mentioned width measure [2]), and is intractable otherwise. This studied class of queries encompasses existential conjunctive queries, which can be viewed as block-sorted queries in which there is one existential quantifier block, and all variables have the same sort. Observe that, given any sentence in quantified conjunctive logic (either one-sorted or multi-sorted) and any structure on which the sentence is to be evaluated, one can view the sentence as a block-sorted query. (This is done as follows: for each sort s that appears in more than one quantifier block, introduce a new sort s^b for each block b where it appears; correspondingly, introduce new relation symbols.) Our theorem can thus be read as providing a general tractability result which is applicable to all of quantified conjunctive logic, and a matching intractability result that proves optimality of this tractability result for the class of block-sorted queries.

Our theorem is the first generalization of Grohe’s theorem to a logic where both quantifiers are present. The previous work suggests that we should proceed the following way: take the width measure of Chen and Dalmau [2], and apply it to some analog of the logically equivalent core of Grohe [6]. However,

the execution of these ideas are not at all obvious and we have to overcome a number of technical barriers. For instance, Grohe’s theorem statement (in the formulation given here) makes reference to logical equivalence. While there is a classical and simple characterization of logical equivalence in existential conjunctive logic [1], logical equivalence for first-order logic is of course well-known to be an undecidable property; logical equivalence for quantified conjunctive logic is now known (in the one-sorted case) to be decidable [4], but is perhaps still not well-understood (for instance, its exact complexity is quite open). Despite this situation, we succeed in identifying, for each block-sorted sentence, a logically equivalent sentence whose width characterizes the original sentence’s complexity, obtaining a statement parallel to that of Grohe’s theorem; the definition of this equivalent sentence is a primary contribution of this article. In carrying out this identification, we present a notion of *core* for block-sorted sentences and develop its basic theory; the core of an existential conjunctive sentence (an established notion) is, intuitively, a minimal equivalent sentence, and Grohe’s theorem can be stated in terms of the treewidth of the cores of a sentence set. Another technical contribution of the article is to develop a graph-theoretic understanding of variable interactions (see Section 4), which understanding is sufficiently strong so as to allow for the delicate embedding of hard sentences from the previous work [2] into the sentences under consideration, to obtain the intractability result. Overall, we believe that the notions, concepts, and techniques that we introduce in this article will play a basic role in the investigation of model checking in logics that are more expressive than the one considered here.

2 Preliminaries

2.1 Terminology and setup

We will work with the following formalization of multi-sorted relational first-order logic. A *signature* is a pair (σ, \mathcal{S}) where \mathcal{S} is a set of *sorts* and σ is a set of relation symbols; each relation symbol $R \in \sigma$ has associated with it an element of \mathcal{S}^* , called the *arity* of R and denoted $\text{ar}(R)$. In formulas over signature (σ, \mathcal{S}) , each variable v has associated with it a sort $s(v)$ from \mathcal{S} ; we use *atom* to refer to an atomic formula $R(v_1, \dots, v_k)$ where $R \in \sigma$ and $s(v_1) \dots s(v_k) = \text{ar}(R)$. A *structure* \mathbf{B} on signature (σ, \mathcal{S}) consists of an \mathcal{S} -sorted family $\{B_s \mid s \in \mathcal{S}\}$ of sets called the *universe* of \mathbf{B} , and, for each symbol $R \in \sigma$, an interpretation $R^{\mathbf{B}} \subseteq B_{\text{ar}(R)}$. Here, for a word $w = w_1 \dots w_k \in \mathcal{S}^*$, we use B_w to denote the product $B_{w_1} \times \dots \times B_{w_k}$. We say that two structures are *similar* if they are defined on the same signature. Let \mathbf{B} and \mathbf{C} be two similar structures defined on the same signature (σ, \mathcal{S}) . We say that \mathbf{B} is a *substructure* of \mathbf{C} if for each $s \in \mathcal{S}$, it holds that $B_s \subseteq C_s$, and for each $R \in \sigma$, it holds that $R^{\mathbf{B}} \subseteq R^{\mathbf{C}}$. We say that \mathbf{B} is an *induced substructure* of \mathbf{C} if, in addition, for each $R \in \sigma$ one has that $R^{\mathbf{B}} = R^{\mathbf{C}} \cap B_{\text{ar}(R)}$.

A *quantified conjunctive query* is a sentence built from atoms, conjunction, existential quantification, and universal quantification. It is well-known that such sentences can be efficiently translated into prenex normal form, that is, of the

form $Q_1v_1 \dots Q_nv_n\phi$ where each Q_i is a quantifier and where ϕ is a conjunction of atoms. For such a sentence, it is well-known that the conjunction ϕ can be encoded as a structure \mathbf{A} where A_s contains the variables of sort s that appear in ϕ and, for each relation symbol R , the relation $R^{\mathbf{A}}$ consists of all tuples (v_1, \dots, v_k) such that $R(v_1, \dots, v_k)$ appears in ϕ . In the other direction, any structure \mathbf{A} can be viewed as encoding the conjunction $\bigwedge_{(v_1, \dots, v_k) \in R^{\mathbf{B}}} R(v_1, \dots, v_k)$. We will typically denote a quantified conjunctive query $Q_1v_1 \dots Q_nv_n\phi$ as a pair (P, \mathbf{A}) consisting of the quantifier prefix $P = Q_1v_1 \dots Q_nv_n$ and a structure \mathbf{A} that encodes the quantifier-free part ϕ . Note that when discussing the evaluation of a sentence (P, \mathbf{A}) on a structure, we can and often will assume that all variables appearing in P are elements of \mathbf{A} .

We define a *block-sorted query* to be a quantified conjunctive query in prenex normal form where for all variables v, v' , if $s(v) = s(v')$ then v, v' occur in the same quantifier block. By a quantifier block, we mean a subsequence $Q_iv_i \dots Q_jv_j$ of the quantifier prefix (with $i \leq j$) having maximal length such that $Q_i = \dots = Q_j$. We number the quantifier blocks from left to right (that is, the outermost quantifier block is considered the first). For each sort s having a variable that appears in such a query, either all variables of sort s are universal, in which case we call s a *universal sort* or a \forall -sort, or all variables of sort s are existential, in which case we call s a *existential sort* or a \exists -sort.

2.2 Conventions

In general, when \mathbf{A} is a structure with universe $\{A_s \mid s \in \mathcal{S}\}$, we assume that the sets A_s are pairwise disjoint, and use A to denote $\cup_{s \in \mathcal{S}} A_s$. Correspondingly, we assume that in forming formulas over signature (σ, \mathcal{S}) , the sets of permitted variables for different sorts are pairwise disjoint. Relative to a quantified conjunctive query (P, \mathbf{A}) , we use A_{\exists} to denote the set $\{a \in A \mid s(a) \text{ is an } \exists\text{-sort}\}$; likewise, we use A_{\forall} to denote the set $\{a \in A \mid s(a) \text{ is an } \forall\text{-sort}\}$. In dealing with sets such as these, for a variable v we use a subscript $< v$ to restrict to variables coming before v in the quantifier prefix P ; for instance, we will use $A_{\forall, < v}$ to denote the set of all universally quantified variables that occur before v . When discussing a function from a set whose elements are sorted to another such set, we assume tacitly that the function preserves sort, that is, for each sort s , each element of sort s in the first set is mapped to an element of sort s in the second set.

Let (P, \mathbf{A}) be a block-sorted query and let \mathbf{B} be a structure similar to \mathbf{A} ; we say that a homomorphism $\phi : \mathbf{A} \rightarrow \mathbf{B}$ is *universal-injective* if ϕ is injective on A_{\forall} .

2.3 Basic facts

Intuitively, evaluating the query (P, \mathbf{A}) on the structure \mathbf{B} can be interpreted as game with two players “universal” and “existential.” In the order given by the prefix P , the two players assign values to the variables; existential and universal

sets the values of the existential and the universal variables, respectively. The aim of existential is to ensure that the resulting assignment satisfies the formula, that is, gives a homomorphism from \mathbf{A} to \mathbf{B} , while universal tries to prevent this. The query (P, \mathbf{A}) is true on \mathbf{B} if existential has a winning strategy. We formalize this intuition by the following definition:

Definition 1. Let (P, \mathbf{A}) be a quantified conjunctive query, and let \mathbf{B} be a structure similar to \mathbf{A} . An existential strategy for (P, \mathbf{A}) on \mathbf{B} is a set of mappings $(f_x : (A_{\forall, < x} \rightarrow B) \rightarrow B_{s(x)})_{x \in A_{\exists}}$ such that the following holds: for any $h : A_{\forall} \rightarrow B$, a homomorphism from \mathbf{A} to \mathbf{B} is given by the map $(f, h) : A \rightarrow B$ defined by $(f, h)(x) = f_x(h \upharpoonright A_{\forall, < x})$ for each existential variable x , and $(f, h)(y) = h(y)$ for each universal variable y .

Proposition 2. Let (P, \mathbf{A}) be a quantified conjunctive query, and let \mathbf{B} be a structure similar to \mathbf{A} . Then $\mathbf{B} \models (P, \mathbf{A})$ if and only if there is an existential strategy.

The transitivity of homomorphisms allows us to quickly deduce consequences of the existence of a homomorphism $\mathbf{A} \rightarrow \mathbf{B}$. For example, we know that there is also a homomorphism $\mathbf{A}' \rightarrow \mathbf{B}$ whenever there is a homomorphism $\mathbf{A}' \rightarrow \mathbf{A}$; and there is also homomorphism $\mathbf{A} \rightarrow \mathbf{B}'$ whenever there is a homomorphism $\mathbf{B} \rightarrow \mathbf{B}'$. These quick observations are very useful in the study of the homomorphism problem, where they allow us to restrict our attention to specific type of structures. In our setting, however, the quantified nature of the problem makes such consequences less obvious. In the following, we find analogs of these observations in our setting, that is, assuming that $\mathbf{B} \models (P_A, \mathbf{A})$ holds, we explore under what conditions the structure \mathbf{B} or the query (P_A, \mathbf{A}) can be replaced to obtain another true statement.

First we give a sufficient condition under which the query can be replaced. Let us say that two similar block-sorted queries (P_A, \mathbf{A}) and (P_C, \mathbf{C}) having the same number of quantifier blocks are *mutually respecting* if for each sort s and for each $i \geq 1$, it holds that s is used in the i th quantifier block of P_A if and only if it is used only in the i th quantifier block of P_C .

Proposition 3 Let (P_A, \mathbf{A}) and (P_C, \mathbf{C}) be similar block-sorted queries that are mutually respecting. Suppose that $i : \mathbf{A} \rightarrow \mathbf{C}$ is a universal-injective homomorphism. Then it holds that (P_C, \mathbf{C}) entails (P_A, \mathbf{A}) .

The following proposition gives a sufficient condition for replacing the structure \mathbf{B} on which the query is evaluated:

Proposition 4 Let σ be a signature, let (P, \mathbf{A}) be a block-sorted query over σ , and let \mathbf{B}, \mathbf{B}' be structures over σ . Suppose that $\mathbf{B} \models (P, \mathbf{A})$ and that there exists a homomorphism $g : \mathbf{B} \rightarrow \mathbf{B}'$ that is universal-surjective in the sense that $f(B_{\forall}) = B'_{\forall}$. Then, it holds that $\mathbf{B}' \models (P, \mathbf{A})$.

Note that this proposition can be viewed as a variant of the known fact that, in standard (one-sorted) first-order logic, if a quantified conjunctive query Φ holds on a structure \mathbf{B} and \mathbf{B} admits a surjective homomorphism to \mathbf{B}' , then Φ also holds on \mathbf{B}' (see for example [4, Lemma 1]).

3 The selfish core

Let (P, \mathbf{C}) be a block-sorted query on signature (σ, \mathcal{S}) . When \mathbf{A} is similar to \mathbf{C} , we say that \mathbf{A} is an \exists -substructure of \mathbf{C} if \mathbf{A} is a substructure of \mathbf{C} ; for each \forall -sort u it holds that $A_u = C_u$; and, for each \exists -sort e it holds that $A_e \subseteq C_e$. We say that \mathbf{A} is a *proper* \exists -substructure of \mathbf{C} if, in addition, there exists an \exists -sort e such that the containment $A_e \subseteq C_e$ is proper.

We say that a block-sorted query (P, \mathbf{C}) is *selfish* if $\mathbf{C} \models (P, \mathbf{C})$. We say that a block-sorted query (P, \mathbf{C}) is a *selfish core* if it is selfish and for any proper \exists -substructure \mathbf{A} of \mathbf{C} , either (P, \mathbf{A}) is not selfish or the queries (P, \mathbf{A}) and (P, \mathbf{C}) are not logically equivalent.

We give characterizations of the notion of selfish core in the following proposition; afterwards, we show that each block-sorted query has (in a sense made precise) a selfish core. Let us say that an endomorphism $h : \mathbf{C} \rightarrow \mathbf{C}$ of a structure \mathbf{C} is *proper* if its image is proper, that is, if there exists a sort $s \in \mathcal{S}$ such that $h(C_s) \subsetneq C_s$.

Proposition 5 *Let (P, \mathbf{C}) be a selfish block-sorted query. The following are equivalent.*

1. (P, \mathbf{C}) is a selfish core.
2. There does not exist a proper endomorphism of \mathbf{C} that fixes each universal variable.
3. There does not exist a proper endomorphism of \mathbf{C} that, for each universal sort u , is injective on C_u .

Define a *selfish core of a block-sorted query (P, \mathbf{A})* to be a block-sorted query that is a selfish core and that is logically equivalent to (P, \mathbf{A}) . We now show that each block-sorted query has a selfish core which is computable (from the query).

Definition 6. *Let (P, \mathbf{A}) be a block-sorted query; we define the block-sorted query (P^*, \mathbf{A}^*) the following way.*

- For each \forall -sort u , define $A_u^* = A_u$.
- For each \exists -sort e , define $A_e^* = \{x^g \mid x \in A_e, g : A_{\forall, < x} \rightarrow A_{\forall, < x}\}$.
- P^* is obtained from P by replacing each quantification $\exists x$ with $\exists x^{g_1} \dots \exists x^{g_m}$ where g_1, \dots, g_m is a list of all the mappings from $A_{\forall, < x}$ to $A_{\forall, < x}$.
- $R^{\mathbf{A}^*} = \{(g'(a_1), \dots, g'(a_k)) \mid (a_1, \dots, a_k) \in R^{\mathbf{A}}, g : A_{\forall} \rightarrow A_{\forall}\}$ where g' is the extension of g that maps a value $x \in A_{\exists}$ to $x^{g|A_{\forall, < x}}$.

Example 7. Consider the query $(P, \mathbf{A}) = \forall y_1, y_2 \exists x : R_1(x, y_1) \wedge R_2(x, y_2)$. For $i, j \in \{1, 2\}$, let g_{ij} be the mapping defined by $g_{ij}(y_1) = y_i$ and $g_{ij}(y_2) = y_j$. Then the query (P^*, \mathbf{A}^*) can be defined as

$$\begin{aligned} & \forall y_1, y_2 \exists x^{g^{11}}, x^{g^{12}}, x^{g^{21}}, x^{g^{22}} : \\ & [R_1(x^{g^{11}}, y_1) \wedge R_2(x^{g^{11}}, y_1) \wedge R_1(x^{g^{12}}, y_1) \wedge R_2(x^{g^{12}}, y_2) \\ & \quad \wedge R_1(x^{g^{21}}, y_2) \wedge R_2(x^{g^{21}}, y_1) \wedge R_1(x^{g^{22}}, y_2) \wedge R_2(x^{g^{22}}, y_2)]. \end{aligned}$$

Proposition 8 *Let (P, \mathbf{A}) be a block-sorted query. The following statements concerning (P^*, \mathbf{A}^*) hold.*

1. (P^*, \mathbf{A}^*) and (P, \mathbf{A}) are logically equivalent.
2. (P^*, \mathbf{A}^*) is selfish.
3. The structure \mathbf{A}^* contains an induced substructure \mathbf{C} such that (P^*, \mathbf{C}) is a selfish core of (P, \mathbf{A}) ; moreover, (P^*, \mathbf{C}) is computable from (P, \mathbf{A}) .

4 Strong and weak elements

Throughout this section, we assume that (P, \mathbf{A}) is a block-sorted query; the definitions and claims are all relative to this query. We use $G_{\mathbf{A}}$ to denote the Gaifman graph of the structure \mathbf{A} , that is, the graph with vertex set A and containing an edge $\{a, a'\}$ if and only if a and a' are distinct and co-occur in a tuple of a relation of \mathbf{A} . Relative to (P, \mathbf{A}) , when i is the number of a quantifier block, we will use notation such as $A_{\geq i}$ to denote the set of variables occurring in block i or later, and define for example $A_{< i}$ analogously.

Definition 9. *A level i component is a maximal connected set of \exists -variables in $G_{\mathbf{A}}[A_{\geq i}]$.*

Definition 10. *Let $x \in A_{\exists}$ be an \exists -variable in the i th quantifier block.*

- For $j \leq i$, use $C_{\mathbf{A}}(x, j)$ to denote the level j component containing x .
- Define $N_{\mathbf{A}}(x, j)$, the neighborhood of $C_{\mathbf{A}}(x, j)$, to be the set of all universal variables in $A_{< j}$ adjacent to $C_{\mathbf{A}}(x, j)$ in $G_{\mathbf{A}}$.
- Define $U_{\mathbf{A}}(x)$ to be $\bigcup_{j \leq i} N_{\mathbf{A}}(x, j)$.

In other words, a universal variable y on level j is in $U_{\mathbf{A}}(x)$ if and only if y can be reached from x on a path in $G_{\mathbf{A}}$ such that all the vertices of the path other than y are existential variables on levels greater than j . We remark that the definition of $C_{\mathbf{A}}(x, j)$, as well as that of the other sets, depends on \mathbf{A} as well as the quantifier prefix P ; however, this prefix will be clear from context, and we omit it from the notation.

Definition 11. *We say that an \exists -variable $x^g \in A_{\exists}^*$ is degenerate if g is non-injective on $U_{\mathbf{A}}(x)$.*

Definition 12. *An \exists -variable $x \in A_{\exists}$ is weak if there exists a universal-injective homomorphism $\psi : \mathbf{A} \rightarrow \mathbf{A}^*$ where $\psi(x)$ is a degenerate element of \mathbf{A}^* ; the \exists -variable x is strong otherwise.*

Example 13. Consider the following query (P, \mathbf{A}) :

$$\begin{aligned} \forall y_1, y_2, y_3 \exists x_1 x_2, x_3, x_4, x_5 \\ R_1(x_1, y_1) \wedge R_2(x_2, y_2) \wedge R_3(x_1, x_2) \wedge R_1(x_3, y_3) \wedge R_1(x_5, y_3) \\ \wedge R_2(x_4, y_3) \wedge R_3(x_3, x_4) \wedge R_3(x_5, x_4). \end{aligned}$$

If g is the mapping with $g(y_1) = g(y_2) = g(y_3) = y_3$, then there is a homomorphism ψ from \mathbf{A} to \mathbf{A}^* that is identity on y_1, y_2, y_3, x_1, x_2 and $\psi(x_3) = \psi(x_5) = x_1^g$ and $\psi(x_4) = x_2^g$. Hence x_3, x_4, x_5 are weak elements.

Definition 14. Define the strong substructure of \mathbf{A} to be the substructure of \mathbf{A} induced by the union of A_{\forall} with the strong elements of \mathbf{A} .

The main result of the section is showing that removing the weak elements does not change the sentence. In the proof of the classification theorem, this will allow us to consider the width of the strong substructure as the classification criteria.

Theorem 15. Let \mathbf{S} be the strong substructure of \mathbf{A} . The queries (P, \mathbf{S}) and (P, \mathbf{A}) are logically equivalent.

We conclude this section with a simple lemma that will be of help in establishing the complexity hardness result.

Lemma 16. Suppose that ϕ is a universal-injective endomorphism of \mathbf{A} and that $x \in A_{\exists}$ is a strong variable. Then $\phi(x)$ is strong as well.

Proof. Assume for contradiction that $\phi(x)$ is not strong: there is a universal-injective homomorphism $\psi : \mathbf{A} \rightarrow \mathbf{A}^*$ where $\psi(\phi(x))$ is a degenerate element. Now $\psi(\phi)$ is a universal-injective homomorphism $\mathbf{A} \rightarrow \mathbf{A}^*$ that maps x to a degenerate element of \mathbf{A}^* , contradicting the assumption that x is strong. \square

5 Classification theorem

When Φ is a set of (possibly multi-sorted) first-order sentences, define Φ -MC to be the model checking problem of deciding, given a sentence $\phi \in \Phi$ and a finite structure \mathbf{B} over the same signature, whether or not $\mathbf{B} \models \phi$. We study this problem using parameterized complexity; we use the terminology and conventions for parameterized complexity defined in [2], and take ϕ to be the parameter of an instance (ϕ, \mathbf{B}) .

As defined in [2], a *prefixed graph* consists of a quantifier prefix P paired with an undirected graph whose vertices are the variables appearing in P . In [2], a width measure is defined that associates a natural number with each prefixed graph; we refer the reader to that article for the precise definition. As we use both the algorithmic and hardness results of [2] as black box, the exact definition does not matter for the purposes of this paper. Fix a computable mapping M that, given a block-sorted query ϕ , computes a selfish core (P, \mathbf{A}) of ϕ , and then computes the strong substructure \mathbf{S} of (P, \mathbf{A}) , and outputs (P, \mathbf{S}) .

Theorem 17. Let Φ be a set of block-sorted queries of bounded arity. If the set of prefixed graphs $\{(P, G_{\mathbf{S}}) \mid (P, \mathbf{S}) \in M(\Phi)\}$ has bounded width, then the problem Φ -MC is in FPT; otherwise, the problem Φ -MC is not in FPT, unless $W[1] \subseteq \text{nuFPT}$.

The remainder of this section is devoted to the proof of Theorem 17.

The positive FPT result is obtained as follows. Given an instance (ϕ, \mathbf{B}) of the problem Φ -MC, the algorithm is to evaluate $\mathbf{B} \models M(\phi)$ using the algorithm

of [2]; this evaluation can be performed in polynomial time given $M(\phi)$, and since the computation of $M(\phi)$ depends only on the parameter of the instance (ϕ, \mathbf{B}) , the whole computation is in FPT.

We now give the hardness result. For a block-sorted query (P, \mathbf{S}) over signature (σ, \mathcal{S}) , we define the relativization $(P, \mathbf{S})_{\text{rel}}$ of (P, \mathbf{S}) in the following way. Denote P by $Q_1 v_1 \dots Q_n v_n$, and let θ be the conjunction of atoms corresponding to \mathbf{S} . Define $(P, \mathbf{S})_{\text{rel}}$ to be the one-sorted sentence $Q_1 v_1 \in W_{v_1} \dots Q_n v_n \in W_{v_n} \theta$ over signature $\sigma \cup \{W_{v_1}, \dots, W_{v_n}\}$ where each W_{v_i} is a fresh unary relation symbol and the arity of a symbol $R \in \sigma$ is the *length* of $\text{ar}_{(\sigma, \mathcal{S})}(R)$. Here, $\exists v \in W\psi$ is syntactic shorthand for $\exists v(W(v) \wedge \psi)$; and, $\forall v \in W\psi$ is syntactic shorthand for $\forall v(W(v) \rightarrow \psi)$. Assuming that the set of prefixed graphs given in the theorem statement has unbounded width, the hardness result of [2, Section 6] implies that $\Phi_{\text{rel}} = \{(M(\phi))_{\text{rel}} \mid \phi \in \Phi\}$ is W[1]-hard or coW[1]-hard under nuFPT reductions. It thus suffices to give an nuFPT reduction from $\Phi_{\text{rel}}\text{-MC}$ to $\Phi\text{-MC}$, which we now do. Let $((P, \mathbf{S})_{\text{rel}}, \mathbf{B})$ be an instance of $\Phi_{\text{rel}}\text{-MC}$, and let $\phi \in \Phi$ be such that $(P, \mathbf{S}) = M(\phi)$; let (P, \mathbf{A}) denote the selfish core of ϕ computed by $M(\phi)$ (note that \mathbf{S} is a substructure of \mathbf{A}).

We will work with the structure \mathbf{A}^* . Let A^{id} denote the subuniverse of \mathbf{A}^* containing all universal variables of \mathbf{A}^* and each existential variable of \mathbf{A}^* of the form a^{id} , where id is the identity mapping. (We use id generically to denote the identity mapping, but note that this is defined on $A_{\forall, < a}$ for an existential variable a .) Observe that A^{id} induces in \mathbf{A}^* a copy of the structure \mathbf{A} . With this correspondence, let S^{id} denote the union of A_{\forall} and the strong elements of A^{id} , and let \mathbf{S}^{id} denote the induced substructure of \mathbf{A}^* on S^{id} . Let D denote the subuniverse of \mathbf{A}^* containing all degenerate elements of \mathbf{A}^* . We will sometimes drop the id superscript when it is clear from context.

Define a structure \mathbf{B}' over signature (σ, \mathcal{S}) as follows. The universe is denoted by $\{B'_s \mid s \in \mathcal{S}\}$ and is defined by $B'_s = \{(a, b) \in (A_s^{\text{id}} \cup D_s) \times (B \cup \{\perp\}) \mid (a \in S_s^{\text{id}} \rightarrow b \in U_a^{\mathbf{B}}) \text{ and } (a \notin S_s^{\text{id}} \rightarrow b = \perp)\}$. Here, B denotes the universe of the one-sorted structure \mathbf{B} . Now, for each $R \in \sigma$, define $R^{\mathbf{B}'}$ to be the relation $\{((a_1, b_1), \dots, (a_k, b_k)) \in B'_{\text{ar}(R)} \mid (a_1, \dots, a_k) \in R^{\mathbf{A}^*} \text{ and } ((a_1, \dots, a_k) \in R^{\mathbf{S}^{\text{id}}} \rightarrow (b_1, \dots, b_k) \in R^{\mathbf{B}})\}$. We will use π_i to denote the mapping that projects a tuple onto the i th coordinate.

We claim that $\mathbf{B} \models (P, \mathbf{S})_{\text{rel}}$ if and only if $\mathbf{B}' \models (P, \mathbf{A})$.

We first prove the backwards direction. We will use the following lemma.

Lemma 18. *Suppose that $h : A \rightarrow (A^{\text{id}} \cup D)$ is a homomorphism from \mathbf{A} to \mathbf{A}^* that is identity on A_{\forall} . Then $h(S) = S^{\text{id}}$.*

Proof. Let h_0 be the homomorphism $\mathbf{A} \rightarrow \mathbf{A}^*$ that is identity on universals and maps x to x^{id} . As \mathbf{A} is selfish and (P, \mathbf{A}) and (P, \mathbf{A}^*) are logically equivalent by Proposition 8(1), we have that $\mathbf{A} \models (P, \mathbf{A}^*)$, implying that there is a homomorphism h^* from \mathbf{A}^* to \mathbf{A} that is identity on the universals.

We claim that h^* is injective on A^{id} . Indeed, otherwise $h^*(h_0)$ is noninjective (as A^{id} is the image of h_0), hence it is a proper endomorphism of \mathbf{A} that is

identity on the universals. By Proposition 5, this contradicts the assumption that \mathbf{A} is a selfish core.

Next we claim that h^* maps S^{id} to S . Otherwise, the endomorphism $h^*(h_0)$ of \mathbf{A} maps a strong element to a weak element, contradicting Lemma 16. Together with the fact that h^* is injective on A^{id} , it follows that h^* maps $A^{\text{id}} \setminus S^{\text{id}}$ to $A \setminus S$.

The homomorphism h cannot map a strong element $x \in S$ to D by definition of strong elements. If h maps a strong element $x \in S$ to $A^{\text{id}} \setminus S^{\text{id}}$, then (as shown in the previous paragraph) endomorphism $h^*(h)$ of \mathbf{A} maps x to $A \setminus S$, that is, to a weak element. As $h^*(h)$ is an endomorphism fixing the universals, this contradicts Lemma 16. Thus we have proved that h maps every strong element to S^{id} . \square

Let $(f'_x)_{x \in A_\exists}$ be an existential strategy witnessing $\mathbf{B}' \models (P, \mathbf{A})$. Let $x \in S$ be an existential variable of $(P, \mathbf{S})_{\text{rel}}$, and let s be the sort of x . For any mapping $h : A_{\forall, < x} \rightarrow B$, define $h' : A_{\forall, < x} \rightarrow B'$ by $h'(y) = (y, h(y))$. Observe that under any such mapping h , we have $\pi_1(\{f'_x(h') \mid x \in S_s\}) = S_s$, since we can extend h' to a mapping $h'' : A_\forall \rightarrow B'$ and then the homomorphism (f', h'') given by Definition 1 is from A to $A^{\text{id}} \cup D$, and Lemma 18 can be applied.

We can thus define a strategy (f_x) for $(P, \mathbf{S})_{\text{rel}}$ on \mathbf{B} as follows: for an existential variable $x \in S_\exists$ and a map $h : A_{\forall, < x} \rightarrow B$, define $f_x(h) = b$ if and only if (x, b) is in $\{f'_x(h') \mid x \in S_s\}$. This mapping is well-defined by the observation of the previous paragraph, and for any $h : A_\forall \rightarrow B$ obeying $h(a) \in U_a^{\mathbf{B}}$, we obtain from the definition of \mathbf{B}' that the homomorphism (f, h) given by Definition 1 is from \mathbf{S} to \mathbf{B} with $(f, h)(x) \in U_x^{\mathbf{B}}$ for each \exists -variable x .

We now prove the forwards direction. We will make use of the following lemma.

Lemma 19. *There is an existential strategy (f'_x) for (P, \mathbf{A}) on the substructure of \mathbf{A}^* induced by $(A^{\text{id}} \cup D)$ where $f'_x(h)$ is a degenerate element in D whenever $h : A_{\forall, x} \rightarrow A_{\forall, x}$ is not injective on $U(x)$.*

Let (f_t) witness $\mathbf{B} \models (P, \mathbf{S})_{\text{rel}}$. We will define a strategy (F_x) to witness $\mathbf{B}' \models (P, \mathbf{A})$.

For each partial map $H : A_\forall \rightarrow B'$ and subset $Y \subseteq A_\forall$ containing the domain of H , fix $e(H, Y)$ to be an extension of H defined on Y such that

- if for some universal sort s it holds that $H_1|_{A_s} = H_2|_{A_s}$, then $e(H_1, Y)|_{A_s} = e(H_2, Y)|_{A_s}$; and,
- if for some universal sort s the map $\pi_1(H)$ is injective on A_s , then $\pi_1(e(H, Y))$ is as well.

It is straightforward to verify that such a mapping e exists; note that when using this mapping, Y will be of the form $A_{\forall, < x}$ for an existential variable x .

We define the strategy (F_x) as follows. Let x be an existential variable of (P, \mathbf{A}) , and let $H : A_{\forall, < x} \rightarrow B'$ be a map. Define $H[x]$ as $e(H|U(x), A_{\forall, < x})$. Set $F_x(H)$ to be the pair (c_1, c_2) where

- $c_1 = f'_x(\pi_1(H[x]))$, where (f'_x) is the strategy from Lemma 19; and,

- c_2 is \perp if c_1 is not in S , and otherwise is equal to $f_{c_1}((H[x])(A_{\forall, < x}))$. Note that $(H[x])(A_{\forall, < x})$ is a set of pairs that should be viewed as a function, in passing it to f_{c_1} .

Observe that if c_1 is not degenerate, then by the just-given lemma, the mapping $\pi_1(H[x])$ is injective on $U(x)$; it follows that $\pi_1(H[x])$ is injective on $A_{\forall, < x}$ by the definition of $H[x]$ and the second condition in the definition of e . This implies that $(H[x])(A_{\forall, < x})$ is the graph of a mapping defined on $A_{\forall, < x}$, and c_2 as described above is well-defined.

By the definition of c_1 , we have that (F_x) has the property that for any $H : A_{\forall} \rightarrow B'$, it holds that $\pi_1(F, H)$ is a homomorphism from \mathbf{A} to \mathbf{A}^* . It remains to verify that if $(a_1, \dots, a_k) \in R^{\mathbf{A}}$, then the image $((t_1, b_1), \dots, (t_k, b_k))$ of (a_1, \dots, a_k) under (F, H) has the property that $(t_1, \dots, t_k) \in R^{\mathbf{S}}$ implies $(b_1, \dots, b_k) \in R^{\mathbf{B}}$. For each existential variable x occurring in (a_1, \dots, a_k) , observe that for any universal variable y coming before it in the quantifier prefix, one has $y \in U(x)$ and thus $H(y) = (H[x])(y)$. It thus suffices to show that if x and x' are existential variables in this tuple where x occurs before x' , $H : A_{\forall, < x} \rightarrow B'$ and $H' : A_{\forall, < x'} \rightarrow B'$ are mappings where H' extends H , then $(H'[x'])(A_{\forall, < x'})$ extends $(H[x])(A_{\forall, < x})$. It suffices to show that $H'[x']$ and $H[x]$ agree on $A_{\forall, < x}$. It follows by definition of U that $U(x)|_{A_{\forall, < x}} = U(x')|_{A_{\forall, < x}}$. Thus, for an \forall -sort s occurring before x , we have $(H'|U(x'))|_{A_s} = (H|U(x))|_{A_s}$. So thus by the first condition in the definition of e , it holds that $e(H'|U(x'), A_{\forall, < x'})|_{A_s} = e(H|U(x), A_{\forall, < x})|_{A_s}$ from which we obtain the desired agreement.

References

1. 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.
2. H. Chen and V. Dalmau. Decomposing quantified conjunctive (or disjunctive) formulas. In *LICS*, 2012.
3. H. Chen and M. Grohe. Constraint satisfaction with succinctly specified relations. *Journal of Computer and System Sciences*, 76(8):847–860, 2010.
4. H. Chen, F. Madelaine, and B. Martin. Quantified constraints and containment problems. In *Twenty-Third Annual IEEE Symposium on Logic in Computer Science (LICS)*, 2008.
5. G. Gottlob, G. Greco, and F. Scarcello. The complexity of quantified constraint satisfaction problems under structural restrictions. In *IJCAI 2005*, 2005.
6. M. Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *Journal of the ACM*, 54(1), 2007.
7. M. Grohe, T. Schwentick, and L. Segoufin. When is the evaluation of conjunctive queries tractable? In *STOC 2001*, 2001.
8. D. Marx. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 735–744, 2010.