

The Tractability Frontier of Graph-Like First-Order Query Sets

Hubie Chen

Departamento LSI, Universidad del País Vasco
and IKERBASQUE, Basque Foundation for Science
E-20018 San Sebastián, Spain
hubie.chen@ehu.es

Abstract

We study first-order model checking, by which we refer to the problem of deciding whether or not a given first-order sentence is satisfied by a given finite structure. In particular, we aim to understand on which sets of sentences this problem is tractable, in the sense of parameterized complexity theory. To this end, we define the notion of a graph-like sentence set, which definition is inspired by previous work on first-order model checking wherein the permitted connectives and quantifiers were restricted. Our main theorem is the complete tractability classification of such graph-like sentence sets, which is (to our knowledge) the first complexity classification theorem concerning a class of sentences that has no restriction on the connectives and quantifiers. To present and prove our classification, we introduce and develop a novel complexity-theoretic framework which is built on parameterized complexity and includes new notions of reduction.

Categories and Subject Descriptors F.4.1 [Mathematical logic and formal languages]: Mathematical logic—Computational logic; H.2.3 [Database management]: Languages—Query languages

General Terms Theory

Keywords first-order queries, query evaluation, parameterized complexity, treewidth

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. In this article, we study *first-order model checking*, by which we refer to the case of this problem where one wishes to evaluate a first-order sentence on a finite structure. This case is of principal interest in database theory, where first-order sentences form a basic, heavily studied class of *database queries*, and where it is well-recognized that the problem of evaluating such a query on a database can be taken as a formulation of first-order model checking. Indeed, the investigation of model checking in first-order logic entails an examination of

one of the simplest, most basic logics, and it can be expected that understanding of the first-order case should provide a well-founded basis for studying model checking in other logics, such as those typically considered in verification and database theory. First-order model checking is well-known to be intractable in general: it is PSPACE-complete.

As has been articulated [11, 16], the typical model-checking situation in the database and verification settings is the evaluation of a relatively short sentence on a relatively large structure. Consequently, it has been argued that, in measuring the time complexity of model checking, one could reasonably allow a slow (non-polynomial-time) preprocessing of the sentence, so long as the desired evaluation can be performed in polynomial time following the preprocessing. Relaxing polynomial-time computation to allow arbitrary preprocessing of a *parameter* of a problem instance 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 utilize this paradigm, and focus the discussion on this form of tractability (here, the sentence is the parameter).

A typical way to understand which types of sentences are well-behaved and exhibit desirable, tractable behavior is to simply consider model checking relative to a set Φ of sentences, and to attempt to understand on which sets one has tractable model checking. We restrict attention to sets of sentences having bounded arity.¹ Here, there have been successes in understanding which sets of sentences are tractable (and which are not) in fragments of first-order logic described by restricting the connectives and quantifiers that may be used: there are systematic classification results for so-called conjunctive queries [8, 12, 13] (formed using the connectives and quantifiers in $\{\wedge, \exists\}$) existential positive queries [4] ($\{\wedge, \vee, \exists\}$), and quantified conjunctive queries [5, 7] ($\{\wedge, \exists, \forall\}$). However, to the best of our knowledge, there has been no classification theorem for general first-order logic, without any restriction on the connectives and quantifiers. In this article, we present the first such classification.

Our approach. In the fragments of first-order logic where the only connective permitted is one of the binary connectives ($\{\wedge, \vee\}$)—such as those of conjunctive queries and quantified conjunctive queries—a heavily studied approach to describing sets Φ of sentences is a *graphical approach*. In this graphical approach, one studies a sentence set Φ if it is *graphical* in the following sense: if one prenex sentence ϕ is contained in Φ and a second prenex sentence ψ has the same prefix as ϕ and also has the same graph as

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CSL-LICS 2014, July 14–18, 2014, Vienna, Austria.
Copyright © 2014 ACM 978-1-4503-2886-9...\$15.00.
<http://dx.doi.org/10.1145/2603088.2603119>

¹Note that in the case of unbounded arity, complexity may depend on the choice of representation of relations [6].

ϕ , then ψ is also in Φ . (By the graph of a prenex sentence ϕ , we mean the graph whose vertices are the variables of ϕ and where two vertices are adjacent if they occur together in an atomic formula.) In the fragments where it was considered, this approach of studying graphical sentence sets is not only a natural way to coarsen the project of classifying all sets Φ of sentences, but in fact, can be used cleanly as a key module in obtaining general classifications of sentence sets: such general classifications have recently been proved by using the respective graphical classifications as black boxes [7, 8].

In this article, we adapt this graphical approach to the full first-order setting. To explain how this is done, consider that a graphical set Φ of sentences satisfies certain syntactic closure properties. For instance, if one takes a sentence from such a graphical set Φ and replaces the relation symbol of an atomic formula, the resulting sentence will have the same graph, and will hence continue to be in Φ ; we refer to this property of Φ as *replacement closure*. As another example, if one rewrites a sentence in Φ by invoking associativity or commutativity of the connective $\wedge$, the resulting sentence will likewise still have the same graph and will hence be contained in Φ also. Inspired by these observations, we define a sentence set Φ to be *graph-like* if it is replacement closed and also closed under certain well-known syntactic transformations, such as associativity and commutativity of the binary connectives (see Section 2 for the full definition).

Our principal result is the complete tractability characterization of graph-like sentence sets (see Theorem 6.2 and the corollaries that follow). In particular, we introduce a measure on first-order formulas which we call *thickness*, and show that a set of graph-like sentences is tractable if and only if it has bounded thickness (under standard complexity-theoretic assumptions). In studying unrestricted first-order logic, we believe that our building on the syntactic, graphical approach—which has an established, fruitful tradition—will facilitate the formulation and obtention of future, more general results.

As evidence of our result’s generality and of its faithfulness to the graphical approach, we note that the graphical classification of quantified conjunctive queries [5] can be readily derived from our main classification result (this will be discussed in the full version of this paper); it follows readily that the dual graphical classification of quantified disjunctive queries (also previously derived [5]) can be dually derived from our main classification result. We therefore give a single classification theorem that naturally unifies together these two previous classifications. Indeed, we believe that the technology that we introduce to derive our classification yields a cleaner, deeper and more general understanding of these previous classifications. Observe that, for each of those classifications, since only one binary connective is present, the two quantifiers behave asymmetrically; this is in contrast to the present situation, where in building formulas, wherever a formula may be constructed, its dual may be as well.

Parameterized complexity. An increasing literature investigates the following general situation:

Given a parameterized problem P
whose instances consist of two parts, and a set S ,
define $P[S]$ to be the restricted version of P where
(x, y) is admitted as an instance iff $x \in S$.

Then, attempt to classify and understand, over all sets S ,
the complexity of the problem $P[S]$.

Examples of such classifications and studies include [5, 6, 8–10, 12–14]. It is our view that this literature suffers from the defect that there is no complexity-theoretic framework for discussing the families of problems obtained thusly. As a consequence, different authors and different articles used divergent language and notions

to present hardness results on and reductions between such problems, and applied different computability assumptions on the sets S considered. We attempt to make a foundational contribution and to ameliorate this state of affairs by presenting a complexity-theoretic framework for handling and classifying problems of the described form. In particular, we introduce notions such as reductions and complexity classes for problems of the above type, which we formalize as *case problems* (Section 4). Although we do not carry out this exercise here, we believe that most of the results in the mentioned literature can be shown to be naturally and transparently expressible within our framework.

In order to derive our classification, we present (within our complexity-theoretic framework) a new notion of reduction which we call *accordion reduction* and which is crucial for the proof of our hardness result. (See Sections 7 and 8 for further discussion.) We believe that this notion of reduction may play a basic role in future classification projects of the form undertaken here.

Let us emphasize that, while the establishment of our main classification theorem makes use of the complexity-theoretic framework and accompanying machinery that was just discussed (and is presented in Sections 4 and 7), this framework and machinery is fully generic in that it does not make any reference to and is not specialized to the model checking problem. We believe and hope that the future will find this framework to be a suitable basis for presenting, developing and discussing complexity classification results.

2. Preliminaries

When $g : A \rightarrow B$ and $h : B \rightarrow C$ are mappings, we will typically use $h(g)$ to denote their composition. When f is a partial mapping, we use $\text{dom}(f)$ to denote its domain, and we use $f \upharpoonright S$ to denote its restriction to the set S . We will use π_i to denote the i th projection, that is, the mapping that, given a tuple, returns the value in the tuple’s i th coordinate. For a natural number k , we use $\underline{k}$ to denote the set $\{1, \dots, k\}$.

2.1 First-order logic

We use the syntax and semantics of first-order logic as given by a standard treatment of the subject. In this article, we restrict to relational first-order logic, so the only symbols in signatures are relation symbols. We use letters such as $\mathbf{A}, \mathbf{B}$ to denote structures and A, B to denote their respective universes. The reader may assume for concreteness that relations of structures are represented using lists of tuples, although in general, we will deal with the setting of bounded arity, and natural representations of relations will be (for the complexity questions at hand) equivalent to this one. We assume that equality is not built-in to first-order logic, so a *formula* is created from *atoms*, the usual connectives ($\neg, \wedge, \vee$) and quantification ($\exists, \forall$); by an *atom*, we mean a formula $R(v_1, \dots, v_k)$ where a relation symbol is applied to a tuple of variables (of the arity of the symbol). A formula is *positive* if it does not contain negation ($\neg$). We use $\text{free}(\phi)$ to denote the set of free variables of a formula ϕ . The *width* of a formula ϕ , denoted by $\text{width}(\phi)$, is defined as the maximum of $|\text{free}(\psi)|$ over all subformulas ψ of ϕ . The *arity* of a formula is the maximum arity over all relation symbols that occur in the formula.

We use $\phi_1 \wedge \dots \wedge \phi_n$ as notation for $(\dots((\phi_1 \wedge \phi_2) \wedge \phi_3) \dots)$, and $\phi_1 \vee \dots \vee \phi_n$ is defined dually. We refer to a formula of the shape $\phi_1 \wedge \dots \wedge \phi_n$ as a *conjunction*, and respectively refer to a formula of the shape $\phi_1 \vee \dots \vee \phi_n$ as a *disjunction*. Let Φ be a set of formulas. A *positive combination* of formulas from Φ is a formula in the closure of Φ under conjunction and disjunction. A *CNF* of formulas from Φ is a conjunction of disjunctions of formulas from Φ , and a *DNF* of formulas from Φ is a disjunction of conjunctions of formulas from Φ .

We will use the following terminology which is particular to this article. A subformula ψ of a formula ϕ is a *positively combined subformula* if, in viewing ϕ as a tree, all nodes on the unique path from the root of ϕ to the parent of the root of ψ (inclusive) are conjunctions or disjunctions. We say that a formula ϕ is *variable-loose* if no variable is quantified twice, and no variable is both quantified and a free variable of ϕ . We say that a formula is *symbol-loose* if no relation symbol appears more than once in the formula. We say that a formula is *loose* if it is both variable-loose and symbol-loose.

We now present a number of syntactic transformations; that each preserves logical equivalence is well-known.²

- (α) Associativity and commutativity of $\wedge$ and $\vee$
- (β) $\exists x(\bigvee_{i=1}^n \phi_i) \equiv \bigvee_{i=1}^n (\exists x \phi_i)$,
 $\forall y(\bigwedge_{i=1}^n \phi_i) \equiv \bigwedge_{i=1}^n (\forall y \phi_i)$
- (γ) $\exists x(\phi \wedge \psi) \equiv (\exists x \phi) \wedge \psi$ if $x \notin \text{free}(\psi)$,
 $\forall y(\phi \vee \psi) \equiv (\forall y \phi) \vee \psi$ if $y \notin \text{free}(\psi)$
- (δ) (Distributivity for $\wedge$ and $\vee$)
 $\phi \wedge (\psi \vee \psi') \equiv (\phi \wedge \psi) \vee (\phi \wedge \psi')$,
 $\phi \vee (\psi \wedge \psi') \equiv (\phi \vee \psi) \wedge (\phi \vee \psi')$
- (ϵ) (DeMorgan's laws)
 $\neg \exists v \phi \equiv \forall v \neg \phi$, $\neg \forall v \phi \equiv \exists v \neg \phi$
 $\neg(\phi \wedge \psi) \equiv \neg \phi \vee \neg \psi$, $\neg(\phi \vee \psi) \equiv \neg \phi \wedge \neg \psi$

We say that a set Φ of formulas is *syntactically closed* if, for each $\phi \in \Phi$, when a formula ϕ' can be obtained from ϕ by applying one of the syntactic transformations (α), (β), (γ), (δ), (ϵ) to a subformula of ϕ , it holds that $\phi' \in \Phi$. The *syntactic closure* of a formula ϕ is the intersection of all syntactically closed sets that contain ϕ .

Let us say that a formula ϕ' on signature σ' is obtainable from a formula ϕ on signature σ by *replacement* if ϕ' can be obtained from ϕ by replacing instances of relation symbols in ϕ with instances of relation symbols from σ' (without making any other changes to ϕ).

Example 2.1. Let ϕ be the formula $\forall y \exists x \exists x' (E(y, x) \wedge E(x, x'))$. Let τ be the signature $\{E, F\}$ where E and F are relation symbols of binary arity. Each of the four formulas ϕ , $\forall y \exists x \exists x' (F(y, x) \wedge E(x, x'))$, $\forall y \exists x \exists x' (E(y, x) \wedge F(x, x'))$, and $\forall y \exists x \exists x' (F(y, x) \wedge F(x, x'))$ is obtainable from ϕ by replacement; moreover, these are the only four formulas over signature τ that are obtainable from ϕ by replacement.

Let us say that a set of formulas Φ is *replacement closed* if, for each $\phi \in \Phi$, when ϕ' is obtainable from ϕ by replacement, it holds that $\phi' \in \Phi$.

Definition 2.2. A set of formulas Φ is *graph-like* if it is syntactically closed and replacement closed.

2.2 Graphs and hypergraphs

When S is a set, we use $K(S)$ to denote the set containing all size 2 subsets of S , that is, $K(S) = \{\{s, s'\} \mid s, s' \in S, s \neq s'\}$. For us, a *graph* is a pair (V, E) where V is a set and $E \subseteq K(V)$.

Here, a *hypergraph* H is a pair $(V(H), E(H))$ consisting of a vertex set $V(H)$ and an edge set $E(H)$ which is a subset of the power set $\wp(V(H))$. We will sometimes specify a hypergraph just by specifying the edge set E , in which case the vertex set is understood to be $\bigcup_{e \in E} e$. We associate a hypergraph $(V(H), E(H))$ with the graph $(V(H), \bigcup_{e \in E(H)} K(e))$, and thereby refer to (for

example) the treewidth of or an elimination ordering of a hypergraph.

An *elimination ordering* of a graph (V, E) is a pair

$$((v_1, \dots, v_n), E')$$

that consists of a superset E' of E and an ordering $v_1, \dots, v_n$ of the elements of V such that the following property holds: for each vertex v_k , any two distinct lower neighbors v, v' of v_k are adjacent in E' , that is, $\{v, v'\} \in E'$; here, a *lower neighbor* of a vertex v_k is a vertex v_i such that $i < k$ and $\{v_i, v_k\} \in E'$. Relative to an elimination ordering e , we define the *lower degree* of a vertex v , denoted by $\text{lower-deg}(e, v)$, to be the number of lower neighbors that it has; we define $\text{lower-deg}(e)$ to be the maximum of $\text{lower-deg}(e, v)$ over all vertices v . We assume basic familiarity with the theory of treewidth [2]. The following is a key property of treewidth that we will utilize; here, we use $\text{tw}(H)$ to denote the treewidth of H .

Proposition 2.3. For each $k \geq 2$, there exists a polynomial-time algorithm that, given as input a hypergraph H with a distinguished edge f , will return the following whenever $\text{tw}(H) < k$: an elimination ordering $e = ((v_1, \dots, v_m), E)$ with $\text{lower-deg}(e) = \text{tw}(H)$ and where $\{v_1, \dots, v_{|f|}\} = f$.

The treewidth of (for example) a set of graphs $\mathcal{G}$ is the set $\{\text{tw}(G) \mid G \in \mathcal{G}\}$; it is said to be *unbounded* if this set is infinite, and *bounded* otherwise. We employ similar terminology, in general, when dealing with a complexity measure defined on a class of objects.

3. Parameterized complexity

In this section, we specify the framework of parameterized complexity to be used in this article.

Throughout, we use Σ to denote an alphabet over which languages are defined. As is standard, we will sometimes view elements of $\Sigma^* \times \Sigma^*$ as elements of Σ^* . A *parameterization* is a mapping from Σ^* to Σ^* . A *parameterized problem* is a pair (Q, κ) consisting of a language $Q \subseteq \Sigma^*$ and a parameterization $\kappa : \Sigma^* \rightarrow \Sigma^*$. A *parameterized class* is a set of parameterized problems.

Assumption 3.1. We assume that each parameterized problem (Q, κ) has a non-trivial language Q , that is, that neither $Q = \Sigma^*$ nor $Q = \emptyset$.

Remark 3.2. Let us remark on a difference between our setup and that of other treatments. Elsewhere, a *parameterization* is often defined to be a mapping from Σ^* to $\mathbb{N}$, and in the context of query evaluation, the parameterization studied is typically the size of the query. In contrast, this article takes the parameterization to be the query itself. Since there are finitely many queries of any fixed size, model checking on a set of queries will be fixed-parameter tractable (that is, in the class FPT, defined below) under one of these parameterizations if and only if it is under the other. However, we find that—as concerns the theory in this article—taking the query itself to be the parameter allows for a significantly cleaner presentation. One example reason is that the reductions we present will generally be “slice-to-slice”, that is, they will send all instances with the same query to instances that share another query. Indeed, to understand the complexity of a set of queries, we will apply a closure operator to pass to a larger set of queries having the same complexity, using the notion of accordion reduction (see Sections 7 and 8); we believe that the theory justifying this passage is most cleanly expressed under the used parameterization. $\square$

We now define what it means for a partial mapping r to be *FPT-computable*; we actually first define a non-uniform version of this

²Note that in the transformation (γ), we permit that ϕ is not present.

notion. This definition coincides with typical definitions in the case that r is a total mapping; in the case that r is a partial mapping, we use a “promise” convention, that is, we do not impose any mandate on the behavior of the respective algorithm in the case that the input x is not in the domain of r .

Definition 3.3. Let $\kappa : \Sigma^* \rightarrow \Sigma^*$ be a parameterization.

A partial mapping $r : \Sigma^* \rightarrow \Sigma^*$ is *nu-FPT-computable with respect to κ* if there exist a function $f : \Sigma^* \rightarrow \mathbb{N}$ and a polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that for each $k \in \Sigma^*$, there exists an algorithm A_k satisfying the following condition: on each string $x \in \text{dom}(r)$ such that $\kappa(x) = k$, the algorithm A_k computes $r(x)$ within time $f(\kappa(x))p(|x|)$.

A partial mapping $r : \Sigma^* \rightarrow \Sigma^*$ is *FPT-computable with respect to κ* if there exists a single algorithm A that can play the role of each algorithm A_k in the just-given definition; formally, if there exist a function $f : \Sigma^* \rightarrow \mathbb{N}$, a polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$, and an algorithm A such that for each $k \in \Sigma^*$, the above condition is satisfied when A_k is set equal to A . $\square$

Definition 3.4. We define FPT to be the class that contains a parameterized problem (Q, κ) if and only if the characteristic function of Q is FPT-computable with respect to κ .

Let us remark that we do not put in effect some of the computability assumptions that are often present in other treatments, for instance, we do not require a computable upper bound on the function f in our definition of *FPT-computable*. This is for simplicity; it can be verified that the theory that we present and carry out can also be done so in the case when such assumptions are present.

We now introduce the notion of a reduction between parameterized problems.

Definition 3.5. Let (Q, κ) and (Q', κ') be parameterized problems. A *FPT-reduction* (respectively, *nu-FPT-reduction*) from (Q, κ) to (Q', κ') is a total mapping g that is FPT-computable (respectively, nu-FPT-computable) with respect to κ such that: (1) for each $x \in \Sigma^*$, it holds that $x \in Q$ if and only if $g(x) \in Q'$; and (2) for each $k \in \Sigma^*$, the set $\kappa'(g(\{x \in \Sigma^* \mid \kappa(x) = k\}))$ is finite.

Assumption 3.6. We assume that each parameterized class C is closed under FPT-reductions, that is, if (Q', κ') is in C and (Q, κ) FPT-reduces to (Q', κ') , then (Q, κ) is in C .

Proposition 3.7. Let (Q, κ) , (Q', κ') , and (Q'', κ'') be parameterized problems.

- If g is a FPT-reduction from (Q, κ) to (Q', κ') and h is a FPT-reduction from (Q', κ') to (Q'', κ'') , then their composition $h(g)$ is a FPT-reduction from (Q, κ) to (Q'', κ'') .
- Similarly, if g is a nu-FPT-reduction from (Q, κ) to (Q', κ') and h is a nu-FPT-reduction from (Q', κ') to (Q'', κ'') , then their composition $h(g)$ is a nu-FPT-reduction from (Q, κ) to (Q'', κ'') .

Also, each FPT-reduction from (Q, κ) to (Q', κ') is an nu-FPT-reduction from (Q, κ) to (Q', κ') .

We may now define, for each parameterized class C , a non-uniform version of the class, denoted by *nu- C* .

Definition 3.8. (non-uniform classes) When C is a parameterized class, we define *nu- C* to be the set that contains each parameterized problem that has a nu-FPT-reduction to a problem in C .

Remark 3.9. It is straightforwardly verified that a parameterized problem (Q, κ) is in the class *nu-FPT* (under the above definitions) if and only if the characteristic function of Q is nu-FPT-computable with respect to κ .

We next present two notions of hardness for parameterized classes.

Definition 3.10. (hardness) Let C be a parameterized class. We say that a problem (Q, κ) is *C -hard* if every problem in C has a FPT-reduction to (Q, κ) . We say that a problem (Q, κ) is *non-uniformly C -hard* if every problem in C has a nu-FPT-reduction to (Q, κ) .

We end this section by observing some basic closure properties of what we call *degree-bounded functions*, which, roughly speaking, are the functions which can serve as the running time of algorithms in the definition of *FPT-computable* (Definition 3.3). Let $\kappa : \Sigma^* \rightarrow \Sigma^*$ be a parameterization. A partial function $T : \Sigma^* \rightarrow \mathbb{N}$ is *degree-bounded* with respect to κ if there exist a function $f : \Sigma^* \rightarrow \mathbb{N}$ and a polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that, for each $x \in \text{dom}(T)$, it holds that $T(x) \leq f(\kappa(x))p(|x|)$. We will make use of the observation that a partial mapping $h : \Sigma^* \rightarrow \Sigma^*$ is FPT-computable with respect to κ if and only if there is a degree-bounded function T (with $\text{dom}(T) \supseteq \text{dom}(h)$) and an algorithm that, for all $x \in \text{dom}(h)$, computes $h(x)$ within time $T(x)$.

The following proposition will be of use in establishing that functions are degree-bounded.

Proposition 3.11. Let κ be a parameterization, and let $T_1, \dots, T_m : \Sigma^* \rightarrow \mathbb{N}$ be partial functions sharing the same domain.

1. If each of $T_1, \dots, T_m$ is degree-bounded with respect to κ , then $T_1 + \dots + T_m$ is as well.
2. If each of $T_1, \dots, T_m$ is degree-bounded with respect to κ , then the product $T_1 \cdot \dots \cdot T_m$ is as well.
3. Let $q : \mathbb{N} \rightarrow \mathbb{N}$ be a polynomial; if a partial function $T : \Sigma^* \rightarrow \mathbb{N}$ is degree-bounded with respect to κ , then $q(T)$ is as well.

4. Case complexity

A number of previous works focus on a decision problem Q where each instance consists of two parts, and obtain restricted versions of the problem by taking sets $S \subseteq \Sigma^*$ and, for each such set, considering the restricted version where one allows only instances where (say) the first of the parts falls into S . This is precisely the type of restriction that we will consider here. We are interested in first-order model checking, which we view as the problem of deciding, given a first-order sentence and a structure, whether or not the sentence holds on the structure; our particular interest is to study restricted versions of this problem where the allowed sentences come from a set S .

It has been useful (see for instance the articles [5, 8]) and is useful in the present article to present reductions between such restricted versions of problems. In order to facilitate our doing this, we present a framework wherein we formalize this type of restricted version of problem as a *case problem*, and then present a notion of reduction for comparing case problems. We believe that our notion of reduction, called *slice reduction*, faithfully abstracts out precisely the key useful properties that are typically present in such reductions in the literature. Note that, in the existing literature, different articles imposed different computability assumptions on the sets S considered (assumptions used include that of computable enumerability, of computability, and of no computability assumption). One feature of our framework is that such reductions can be carried out and discussed independently of whether or not any such computability assumption is placed on the sets S ; a general theorem (Theorem 4.5) allows one to derive normal parameterized reductions from slice reductions, where the exact computability of the reduction derivable depends on the computability assumption placed on the sets S .

We now introduce our framework. Suppose that $Q \subseteq \Sigma^* \times \Sigma^*$ is a language of pairs; for a set $T \subseteq \Sigma^*$, we use Q_T to denote the language $Q \cap (T \times \Sigma^*)$ and for a single string $t \in \Sigma^*$, we use Q_t to denote the language $Q \cap (\{t\} \times \Sigma^*)$.

Definition 4.1. A *case problem* consists of a language of pairs $Q \subseteq \Sigma^* \times \Sigma^*$ and a subset $S \subseteq \Sigma^*$, and is denoted $Q[S]$. When $Q[S]$ is a case problem, we use $\text{param-}Q[S]$ to denote the parameterized problem (Q_S, π_1) .

Ultimately, our purpose in discussing a case problem $Q[S]$ is to understand the complexity of the associated parameterized problem $\text{param-}Q[S]$. As mentioned, formalizing the notion of a case problem allows us to cleanly present reductions between such problems.

Remark 4.2. Let (Q, κ) be a parameterized problem. Under the assumption that κ is FPT-computable with respect to itself, the parameterized problem (Q, κ) is straightforwardly verified to be equivalent, under FPT-reduction, to the parameterized problem (Q', π_1) where $Q' = \{(\kappa(x), x) \mid x \in \Sigma^*\}$. Hence, any such given parameterized problem (Q, κ) may be canonically associated to the case problem $Q'[\Sigma^*]$, as one has $\text{param-}Q'[\Sigma^*] = (Q', \pi_1)$.

We define *case-CLIQUE* to be the case problem $Q[\Sigma^*]$ where Q contains a pair (k, G) if and only if G is a graph that contains a clique of size k (we assume that both G and k are encoded as strings over Σ); we define *case-co-CLIQUE* to be the problem $\bar{Q}[\Sigma^*]$.

We now present the notion of *slice reduction*, which allows us to compare case problems.

Definition 4.3. A case problem $Q[S]$ *slice reduces* to a second case problem $Q'[S']$ if there exist:

- a computably enumerable language $U \subseteq \Sigma^* \times \Sigma^*$ and
- a partial function $r : \Sigma^* \times \Sigma^* \times \Sigma^* \rightarrow \Sigma^*$ that has $\text{dom}(r) = U \times \Sigma^*$ and is FPT-computable with respect to the parameterization (π_1, π_2)

such that the following conditions hold:

- (*coverage*) for each $s \in S$, there exists $s' \in S'$ such that $(s, s') \in U$, and
- (*correctness*) for each $(t, t') \in U$, it holds (for each $y \in \Sigma^*$) that

$$(t, y) \in Q \Leftrightarrow (t', r(t, t', y)) \in Q'.$$

We call the pair (U, r) a *slice reduction* from $Q[S]$ to $Q'[S']$.

In this definition, we understand the parameterization (π_1, π_2) to be the mapping that, for all $(s, s') \in U$ and $y \in \Sigma^*$, returns the pair (s, s') given the triple (s, s', y) .

Theorem 4.4. (Transitivity of slice reducibility) Suppose that $Q_1[S_1]$ slice reduces to $Q_2[S_2]$ and that $Q_2[S_2]$ slice reduces to $Q_3[S_3]$. Then $Q_1[S_1]$ slice reduces to $Q_3[S_3]$.

The following theorem allows one to derive an FPT-reduction or an nu-FPT-reduction from a slice reduction.

Theorem 4.5. Suppose that a case problem $Q[S]$ slice reduces to another case problem $Q'[S']$. Then, it holds that $\text{param-}Q[S]$ nu-FPT-reduces to $\text{param-}Q'[S']$; if in addition S and S' are computable, then $\text{param-}Q[S]$ FPT-reduces to $\text{param-}Q'[S']$.

For each parameterized class C , we define *case- C* to be the set of case problems that contains a case problem $Q[S]$ if and only if there exists a case problem $Q'[S']$ such that

- $\text{param-}Q'[S']$ is in C ,
- $Q[S]$ slice reduces to $Q'[S']$, and
- S' is computable.

Proposition 4.6. Suppose that C is a parameterized class, and that $Q[S]$ is a case problem in *case- C* . Then, the problem $\text{param-}Q[S]$ is in *nu- C* ; if it is assumed additionally that S is computable, then the problem $\text{param-}Q[S]$ is in C .

Let $Q[S]$ be a case problem and let C be a parameterized class. We say that $Q[S]$ is *case- C -hard* if there exists a case problem $Q^-[S^-]$ such that

- $\text{param-}Q^-[S^-]$ is C -hard,
- $Q^-[S^-]$ slice reduces to $Q[S]$, and
- S^- is computable.

Proposition 4.7. Suppose that C is a parameterized class, and that $Q[S]$ is a case problem that is *case- C -hard*. Then, the problem $\text{param-}Q[S]$ is non-uniformly C -hard; if it is assumed additionally that S is computable, then the problem $\text{param-}Q[S]$ is C -hard.

5. Thickness

In this section, we define a measure of first-order formulas that we call *thickness*, which we will show is the crucial measure that determines whether or not a graph-like set of sentences is tractable. From a high-level viewpoint, the measure is defined in the following way. We first define a notion of *organized formula* and show that each formula is logically equivalent to a positive combination of organized formulas; we then define a notion of *layered formula* and show that each organized formula is logically equivalent to a layered formula. We then, for each layered formula ϕ , define its thickness (denoted by $\text{thick}_l(\phi)$), and then naturally extend this definition to positive combinations of layered formulas, and hence to all formulas. A key property of thickness, which we prove (Theorem 5.9), is that there exists an algorithm that, given a formula ϕ , outputs an equivalent formula that uses at most $\text{thick}(\phi)$ many variables.

Definition 5.1. We define the set of organized formulas inductively, as follows.

- Each atom and each negated atom is an organized formula.
- If each of $\phi_1, \dots, \phi_n$ is an organized formula and $v \in \text{free}(\phi_1) \cap \dots \cap \text{free}(\phi_n)$, then $\exists v(\bigwedge_{i=1}^n \phi_i)$ and $\forall v(\bigvee_{i=1}^n \phi_i)$ are organized formulas.

Theorem 5.2. There exists an algorithm org^+ that, given a formula ϕ as input, outputs a positive combination $\text{org}^+(\phi)$ of organized formulas that is logically equivalent to ϕ and that is in the syntactic closure of ϕ .

The algorithm of this theorem is defined recursively with respect to formula structure.

In what follows, when V is a set of variables and $Q \in \{\exists, \forall\}$ is a quantifier, we will use QV as shorthand for $Qv_1 \dots Qv_n$, where $v_1, \dots, v_n$ is a list of the elements of V . Our discussion will always be independent of the particular ordering chosen. Relative to a hypergraph H and a subset $S \subseteq V(H)$, we consider a set of edges $\{e_1, \dots, e_k\}$ to be *S-connected* if one has connectedness of the graph with vertices $\{e_1, \dots, e_k\}$ and having an edge between e_i and e_j if and only if $S \cap e_i \cap e_j \neq \emptyset$; we say that the hypergraph H is itself *S-connected* if $E(H)$ is *S-connected*.

Definition 5.3. We define the sets of $\exists$ -layered formulas and of $\forall$ -layered formulas to be the variable-loose formulas that can be constructed inductively, as follows:

- Each atom and each negated atom is both an $\exists$ -layered formula and a $\forall$ -layered formula.
- If each of $\phi_1, \dots, \phi_n$ is a $\forall$ -layered formula, and $X \subseteq \text{free}(\phi_1) \cup \dots \cup \text{free}(\phi_n)$ is such that the hypergraph $\{\text{free}(\phi_1), \dots, \text{free}(\phi_n)\}$ is X -connected, then $\exists X(\bigwedge_{i=1}^n \phi_i)$ is an $\exists$ -layered formula.
- If each of $\phi_1, \dots, \phi_n$ is an $\exists$ -layered formula, and $Y \subseteq \text{free}(\phi_1) \cup \dots \cup \text{free}(\phi_n)$ is such that the hypergraph

$\{\text{free}(\phi_1), \dots, \text{free}(\phi_n)\}$ is Y -connected, then $\forall Y(\bigvee_{i=1}^n \phi_i)$ is a $\forall$ -layered formula.

Theorem 5.4. There exists an algorithm that, given an organized formula ϕ as input, outputs a layered formula that is logically equivalent to ϕ .

The intuitive idea behind the algorithm of Theorem 5.4 is to combine together (into a set quantification QV) quantifiers of the same type that occur adjacently in ϕ .

Theorem 5.5. There exists an algorithm lay^+ that, given a formula ϕ as input, outputs a positive combination $\text{lay}^+(\phi)$ of layered formulas that is logically equivalent to ϕ .

Proof. Immediate from Theorems 5.2 and 5.4. $\square$

Definition 5.6. We define the following measures on layered formulas.

- When ϕ is an atom or a negated atom, we define $\text{thick}_l(\phi) = |\text{free}(\phi)|$.
- Suppose that ϕ is a layered formula of the form $\exists U(\bigwedge_{i=1}^n \phi_i)$ or $\forall U(\bigvee_{i=1}^n \phi_i)$.

We define the *local thickness* of ϕ as

$$\text{local-thick}_l(\phi) = 1 + \text{tw}(\{\text{free}(\phi_i) \mid i \in \underline{n}\} \cup \{\text{free}(\phi)\})$$

where the object to which the treewidth is applied is a hypergraph specified by its edge set, which hypergraph has vertex set $\bigcup_{i=1}^n \text{free}(\phi_i)$.

We define the *thickness* of ϕ inductively as

$$\text{thick}_l(\phi) = \max(\{\text{local-thick}_l(\phi)\} \cup \{\text{thick}_l(\phi_i) \mid i \in \underline{n}\}).$$

We define the *quantified thickness* of ϕ as

$$\text{quant-thick}_l(\phi) = 1 + \text{tw}(\{\text{free}(\phi_i) \cap U \mid i \in \underline{n}\}).$$

Definition 5.7. The *thickness* of an arbitrary formula ϕ is defined as follows: let Ψ be the set of layered formulas that are positively combined subformulas of $\text{lay}^+(\phi)$ such that $\text{lay}^+(\phi)$ is a positive combination over Ψ ; then,

$$\text{thick}(\phi) = \max_{\psi \in \Psi} \text{thick}_l(\psi).$$

Proposition 5.8. The function $\text{thick}(\cdot)$ is computable.

Proof. This follows from the computability of $\text{lay}^+(\phi)$ from ϕ (Theorem 5.5) and the definition of thickness. $\square$

We indeed now demonstrate a principal property of thickness, namely, that this measure provides an upper bound on the number of variables needed to express a formula; this upper bound is effective in that there is an algorithm that computes an equivalent formula using a bounded number of variables. When $k \geq 1$, let us say that a formula *uses k many variables* if the set containing all variables that occur in the formula has size less than or equal to k .

Theorem 5.9. There exists an algorithm that, given as input a formula ϕ , outputs an equivalent formula that uses $\text{thick}(\phi)$ many variables.

Definition 5.10. Let ϕ be a formula of the form $\exists V(\bigwedge_{i=1}^n \phi_i)$ or $\forall V(\bigvee_{i=1}^n \phi_i)$. Let us say that a pair $e = ((v_1, \dots, v_m), E)$ consisting of an ordering $v_1, \dots, v_m$ of the elements in $\bigcup_{i=1}^n \text{free}(\phi_i)$ and a subset $E \subseteq K(\{v_1, \dots, v_m\})$ is an *elimination ordering* of ϕ if:

- the variables in $\text{free}(\phi)$ occur first in the ordering, that is, $\{v_1, \dots, v_{|\text{free}(\phi)|}\} = \text{free}(\phi)$; and,
- e is an elimination ordering of the hypergraph $\{\text{free}(\phi_i)\} \cup \{\text{free}(\phi_i) \mid i \in \underline{n}\}$.

The following lemma can be taken as a variation of a lemma of Kolaitis and Vardi [15, Lemma 5.2].

Lemma 5.11. There exists a polynomial-time algorithm that, given a formula ϕ of the form $\exists V(\bigwedge_{i=1}^n \phi_i)$ or $\forall V(\bigvee_{i=1}^n \phi_i)$ and an elimination ordering e of ϕ , outputs a formula ϕ' that is logically equivalent to ϕ such that $\text{width}(\phi') \leq \max(\{1 + \text{lower-deg}(e)\} \cup \{\text{width}(\phi_i) \mid i \in \underline{n}\})$.

Proof. (Theorem 5.9) By definition of $\text{thick}(\phi)$ and by the computability of $\text{lay}^+(\phi)$ from ϕ (Theorem 5.5), it suffices to prove that there is an algorithm that, given a layered formula ϕ , returns an equivalent formula that uses $\text{thick}_l(\phi)$ many variables.

Consider the following algorithm A defined recursively on layered formulas. When ϕ is an atom or a negated atom, $A(\phi) = \phi$. When ϕ is of the form $\exists V(\bigwedge_{i=1}^n \phi_i)$ or $\forall V(\bigvee_{i=1}^n \phi_i)$, the algorithm proceeds as follows. Set H to be the hypergraph

$$\{\text{free}(\phi_i) \mid i \in \underline{n}\} \cup \{\text{free}(\phi)\}.$$

We have $1 + \text{tw}(H) = \text{local-thick}_l(\phi) \leq \text{thick}_l(\phi)$. By calling the algorithm of Proposition 2.3 on the hypergraph H and the distinguished edge $\text{free}(\phi)$, we can obtain an elimination ordering e of ϕ having $1 + \text{lower-deg}(e) = 1 + \text{tw}(H) \leq \text{thick}_l(\phi)$. Let ϕ' be the formula obtained from ϕ by replacing each formula ϕ_i by $A(\phi_i)$; for each $i \in \underline{n}$, we have $\text{width}(\phi_i) \leq \text{thick}_l(\phi_i)$. By applying the algorithm of Lemma 5.11 to ϕ' and e , we obtain a formula ϕ'' having $\text{width}(\phi'') \leq \text{thick}_l(\phi)$. By renaming variables in ϕ'' , we can obtain an equivalent formula where the number of variables used is equal to $\text{width}(\phi'')$. The output $A(\phi)$ of the algorithm is this equivalent formula. $\square$

6. Graph-like queries

In this section, we state our main theorem and two corollaries thereof, which describe the tractable graph-like sentence sets. In the following two sections, we prove the hardness portion of the main theorem.

Definition 6.1. Define MC to be the language of pairs

$$\{(\phi, \mathbf{B}) \mid \mathbf{B} \models \phi\}$$

where ϕ denotes a first-order sentence, and $\mathbf{B}$ denotes a relational structure.

Theorem 6.2. (Main theorem) Let Φ be a graph-like set of sentences having bounded arity. If Φ has bounded thickness, then the case problem $\text{MC}[\Phi]$ is in case-FPT; otherwise, the case problem $\text{MC}[\Phi]$ is case-W[1]-hard or case-co-W[1]-hard.

We give a proof of this theorem that makes a single forward reference to the main theorem of Section 8.

Proof. Suppose Φ has thickness bounded above by k . Define S' to be the set of sentences having thickness less than or equal to k . The set S' is computable by Proposition 5.8. We have that $\text{param-MC}[S']$ is in FPT via the algorithm that, given an instance $(\phi, \mathbf{B})$, first checks if $\phi \in S'$, and if so, invokes Theorem 5.9 to obtain ϕ' , and then performs the natural bottom-up, polynomial-time evaluation of ϕ' on $\mathbf{B}$ (à la Vardi [17]). The case problem $\text{MC}[\Phi]$ slice reduces to $\text{MC}[S']$ via the slice reduction $((\{s', s'\} \mid s' \in S'), \pi_3)$.

Suppose that Φ has unbounded thickness. Theorem 8.1 yields that either case-CLIQUE or case-co-CLIQUE slice reduces to $\text{MC}[\Phi]$. It then follows by definition that $\text{MC}[\Phi]$ is case-W[1]-hard or case-co-W[1]-hard, since the problems CLIQUE and co-CLIQUE are W[1]-hard and co-W[1]-hard, respectively. $\square$

We now provide two corollaries that describe the complexity of the problems $\text{param-MC}[\Phi]$ addressed by the main theorem; the first corollary assumes that Φ is computable, while the second corollary makes no computability assumption on Φ . Both of these corollaries follow directly from Theorem 6.2 via use of Propositions 4.6 and 4.7.

Corollary 6.3. Let Φ be a computable, graph-like set of sentences having bounded arity. If Φ has bounded thickness, then the problem $\text{param-MC}[\Phi]$ is in FPT; otherwise, the problem $\text{param-MC}[\Phi]$ is not in FPT, unless $W[1] \subseteq \text{FPT}$.

Corollary 6.4. Let Φ be a graph-like set of sentences having bounded arity. If Φ has bounded thickness, then the problem $\text{param-MC}[\Phi]$ is in nu-FPT ; otherwise, the problem $\text{param-MC}[\Phi]$ is not in nu-FPT , unless $W[1] \subseteq \text{nu-FPT}$.

In the full version of this article, we provide a discussion of how the main theorem of the present paper can be used to readily derive the dichotomy theorem of graphical sets of quantified conjunctive queries [5]; a dual argument yields the corresponding theorem on graphical sets of quantified disjunctive queries. Note that it follows immediately from this discussion and [5, Example 3.5] that there exists a set of graph-like sentences having bounded arity that is tractable, but not contained in one of the tractable classes identified by Adler and Weyer [1].

Let us also note here that, as pointed out by Adler and Weyer, it is undecidable, given a first-order sentence ϕ and a value $k \geq 1$, whether or not ϕ is logically equivalent to a k -variable sentence, and hence one cannot expect an algorithm that takes a first-order sentence and outputs an equivalent one that minimizes the number of variables. (This undecidability result holds even for positive first-order logic [3].)

7. Accordion reductions

In this section, we introduce a notion that we call *accordion reduction*. When $C \subseteq \Sigma^* \times \Sigma^*$ is a set of string pairs and $S \subseteq \Sigma^*$ is a set of strings, we use $\text{closure}_C(S)$ to denote the intersection of all sets T containing S and having the closure property that, if $(u, u') \in C$ and $u' \in T$, then $u \in T$. When an accordion reduction exists for such a C (with respect to a language Q), the theorem of this section (Theorem 7.2) yields that the problem $Q[\text{closure}_C(S)]$ slice reduces to $Q[S]$. Hence, an accordion reduction is not itself a slice reduction, but its existence provides a sufficient condition for the existence of a class of slice reductions.

How is this section's theorem proved? One component of an accordion reduction is an FPT-computable mapping r that, for each $(u, u') \in C$, maps a Q -instance (u, y) to a Q -instance (u', y') . Intuitively, to give a slice reduction from $Q[\text{closure}_C(S)]$ to $Q[S]$, one needs to reduce, for any $s \in S$ and $s_1 \in \text{closure}_C(S)$, instances of the form $(s_1, \cdot)$ to instances of the form $(s, \cdot)$. The containment $s_1 \in \text{closure}_C(S)$ implies the existence of a sequence $s_1, \dots, s_k = s$ such that every pair (s_i, s_{i+1}) is in C . This naturally suggests applying the map r repeatedly, but note that there is no constant bound on the length k of the sequence. Hence, r needs to be sufficiently well-behaved so that, when composed with itself arbitrarily many times (in the described way), the end effect is that of an FPT-computable function that may serve as the map in the definition of slice reduction. (The author is mentally reminded of the closing of an accordion in thinking that this potentially long sequence of compositions yields a single well-behaved map.) To ensure this well-behavedness, we impose a condition that we call *measure-linearity*.

In the context of accordion reductions, a *measure* is a mapping $m : \Sigma^* \times \Sigma^* \rightarrow \mathbb{N}$ such that there exist a function $f : \Sigma^* \rightarrow \mathbb{N}$ and

a polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ whereby, for all pairs $(s, y) \in \Sigma^* \times \Sigma^*$, it holds that $m(s, y) \leq |y| \leq f(s)p(m(s, y))$.

Definition 7.1. Let $Q \subseteq \Sigma^* \times \Sigma^*$ be a language of pairs. With respect to Q , an *accordion reduction* consists of:

- a computably enumerable language $C \subseteq \Sigma^* \times \Sigma^*$,
- a measure $m : \Sigma^* \times \Sigma^* \rightarrow \mathbb{N}$, and
- a mapping $r : \Sigma^* \times \Sigma^* \times \Sigma^* \rightarrow \Sigma^*$ that has $\text{dom}(r) = C \times \Sigma^*$, that is FPT-computable with respect to the parameterization (π_1, π_2) , and that is *measure-linear* in that, for each pair $(u, u') \in C$, there exists a constant $B_{(u, u')}$ such that, for each $y \in \Sigma^*$, it holds that $m(u', r(u, u', y)) \leq B_{(u, u')} m(u, y)$,

such that the following condition holds:

- (*correctness*) for each $(u, u') \in C$, it holds (for each $y \in \Sigma^*$) that

$$(u, y) \in Q \Leftrightarrow (u', r(u, u', y)) \in Q.$$

Theorem 7.2. Suppose that $Q[S]$ is a case problem, and that (C, m, r) is an accordion reduction with respect to Q . Then, the case problem $Q[\text{closure}_C(S)]$ slice reduces to the case problem $Q[S]$.

8. Hardness

In this section, we establish the main intractability result of the paper, namely, that the case problem $\text{MC}[\Phi]$ is hard when Φ is graph-like and has unbounded thickness.

Theorem 8.1. Suppose that Φ is a set of graph-like sentences of bounded arity such that $\text{thick}(\Phi)$ is unbounded. Then, either case-CLIQUE or case-co-CLIQUE slice reduces to $\text{MC}[\Phi]$.

Proof. Immediate from Lemmas 8.2 and 8.3, and Theorem 8.5. $\square$

This intractability result is obtained by composing three slice reductions. Define a formula to be *friendly* if it is loose, positive, and layered. We first show (Lemma 8.2) that there exists a set of friendly sentences Ψ , with $\text{thick}_l(\Psi)$ unbounded, such that $\text{MC}[\Psi]$ slice reduces to $\text{MC}[\Phi]$. We next show (Lemma 8.3) that a multi-sorted version $\text{full}(\Psi)$ of Ψ has the property that $\text{MC}_s[\text{full}(\Psi)]$ slice reduces to $\text{MC}[\Psi]$; here, MC_s denotes the multi-sorted generalization of MC. Finally, we directly slice reduce either case-CLIQUE or case-co-CLIQUE to $\text{MC}_s[\Psi]$ (Theorem 8.5); this third reduction is obtained via an accordion reduction.

Lemma 8.2. Suppose that Φ is a set of graph-like sentences of bounded arity such that $\text{thick}(\Phi)$ is unbounded. There exists a set of friendly sentences Ψ such that $\text{MC}[\Psi]$ slice reduces to $\text{MC}[\Phi]$ and such that $\text{thick}_l(\Psi)$ is unbounded.

In what follows, we will work with multi-sorted relational first-order logic, formalized as follows. (For differentiation, we will refer to formulas in the usual first-order logic considered thus far as *one-sorted*.) A *signature* is a pair (σ, S) where S is a set of *sorts* and σ is a set of relation symbols; each relation symbol $R \in \sigma$ has an associated arity $\text{ar}(R)$ which is an element of S^* . In a formula over signature (σ, S) , each variable v has associated with it a sort $s(v)$ from S ; an atom is a 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, S) consists of an S -sorted family $\{B_s \mid s \in S\}$ of sets called the *universe* of $\mathbf{B}$, and for each symbol $R \in \sigma$, an interpretation $R^{\mathbf{B}} \subseteq B_{\text{ar}(R)}$, where for a word $w = w_1 \dots w_k \in S^*$, we use B_w to denote the product $B_{w_1} \times \dots \times B_{w_k}$. We use MC_s to denote the multi-sorted version of MC, that is, it is the language of pairs $(\phi, \mathbf{B})$ where ϕ is a sentence and $\mathbf{B}$ is a structure both having the same signature (σ, S) , and $\mathbf{B} \models \phi$.

Suppose that ϕ is a multi-sorted friendly formula on signature $(\sigma, \mathcal{S})$, and let V be the set of variables occurring in ϕ ; we say that ϕ is *fully-sorted* if $V \subseteq \mathcal{S}$ and for each $v \in V$, the sort of v is v itself (that is, $s(v) = v$). When ψ is a one-sorted friendly formula, and V is the set of variables that occur in ψ , we use $\text{full}(\psi)$ to denote the natural fully-sorted formula induced by ψ , namely, the formula on signature (σ, V) where σ contains those symbols occurring in ψ , and, if $R(v_1, \dots, v_k)$ appears in ψ , then $\text{ar}(R) = v_1 \dots v_k$ (this is well-defined since ϕ is symbol-loose).

Lemma 8.3. Let Ψ be a set of one-sorted friendly sentences. The set $\text{full}(\Psi)$ of fully-sorted friendly sentences has the property that $\text{MC}_s[\text{full}(\Psi)]$ slice reduces to $\text{MC}[\Psi]$.

Proof. We give a slice reduction. Define U to be the set that contains each pair of the form $(\text{full}(\psi), \psi)$. Suppose that ψ and $\mathbf{B}$ are over signature $(\sigma, \mathcal{S})$. Define $r(\text{full}(\psi), \psi, \mathbf{B}) = \mathbf{B}'$, where $\mathbf{B}'$ is defined as follows. Let B' be a set whose cardinality is $\max_{s \in \mathcal{S}} |B_s|$. For each sort $s \in \mathcal{S}$, fix a map $f_s : B' \rightarrow B_s$ that is surjective. For each relation symbol $R \in \sigma$ of arity $s_1 \dots s_k$, define $R^{\mathbf{B}'} = \{(b'_1, \dots, b'_k) \in B'^k \mid (f_{s_1}(b'_1), \dots, f_{s_k}(b'_k)) \in R^{\mathbf{B}}\}$. It is straightforward to prove by induction that, for all subformulas ϕ of ψ , and each assignment g from the set of variables to B' , that $\mathbf{B}', g \models \phi$ if and only if $\mathbf{B}, f \circ^s g \models \text{full}(\phi)$. Here, $f \circ^s g$ denotes the mapping that sends each variable v to $f_{s(v)}(g(v))$ and we view $\text{full}(\phi)$ as a formula over the signature $(\sigma, \mathcal{S})$ of ψ . The correctness of the reduction follows. $\square$

In the remainder of this section, we assume that all formulas and structures under discussion are multi-sorted.

Let us say that a friendly formula is a *simple formula* if it is of the form $\exists X(\bigwedge_{i=1}^n \alpha_i)$ or of the form $\forall Y(\bigwedge_{i=1}^n \alpha_i)$ where each α_i is an atom. When ϕ is a simple formula where the variables V are those that are quantified initially, we say that ψ is a *sentence based on ϕ* if ψ is a simple sentence derivable from ϕ by replacing each atom $R(w_1, \dots, w_k)$ with an atom whose variables are the elements in $\{w_1, \dots, w_k\} \cap V$.

We now present an accordion reduction that will be used to derive our hardness result. When Ψ is a set of friendly sentences, this accordion reduction will allow a simple subformula $\phi' = \exists V \chi$ to simulate a disjunction of atoms on $\text{free}(\phi')$, and likewise for a simple subformula $\phi' = \forall V \chi$ to simulate a corresponding conjunction; this is made precise as follows. The set C is defined to contain a pair (ψ, ϕ) of friendly sentences if there exists a simple subformula $\phi' = QV \chi$ of ϕ such that one of the following conditions holds:

- (1) ψ is a sentence based on ϕ' .
- (2) $Q = \exists$ and ψ is a friendly sentence obtained from ϕ by replacing ϕ' with $\bigvee_{i=1}^m E_i(v_{i1}, v_{i2})$, where the tuples (v_{i1}, v_{i2}) are such that $\{v_{11}, v_{12}\}, \dots, \{v_{m1}, v_{m2}\}$ is a list of the elements in $K(\text{free}(\phi'))$ and the E_i are relation symbols (each of which is fresh in that it does not appear elsewhere in ψ).
- (3) $Q = \forall$ and ψ is a friendly sentence obtained from ϕ by replacing ϕ' with $\bigwedge_{i=1}^m E_i(v_{i1}, v_{i2})$ where the tuples (v_{i1}, v_{i2}) and the symbols E_i are as described in the previous case.

Theorem 8.4. Let M be the measure such that $M(\phi, \mathbf{B})$ is equal to $\max_{s \in \mathcal{S}} |B_s|$ when $\mathbf{B}$ is a multi-sorted structure defined on the signature $(\sigma, \mathcal{S})$ of ϕ , and such that $M(\phi, \mathbf{B})$ is equal to 0 otherwise. There exists a mapping $r : \Sigma^* \times \Sigma^* \times \Sigma^* \rightarrow \Sigma^*$ such that the triple (C, M, r) is an accordion reduction with respect to MC_s . (Here, C is the set defined above.)

Proof. Suppose $(\psi, \phi, \mathbf{A})$ is a triple with $(\psi, \phi) \in C$ and where $\mathbf{A}$ is a structure over the signature of ψ . We define $r(\psi, \phi, \mathbf{A})$ to be the structure $\mathbf{B}$, defined as follows.

It is straightforward to treat the case where there exists a simple subformula ϕ' of ϕ such that ψ is the sentence based on ϕ' (here, we omit discussion of this case).

In the remainder of this proof, we consider the case that there exist a simple subformula $\phi' = \exists V \chi$ of ϕ such that ψ is a friendly sentence obtained from ϕ by replacing ϕ' with a disjunction ψ' as in the definition of C . (Dual to this case is the remaining case where $\phi' = \forall V \chi$ and ψ is obtained from ϕ by replacing ϕ' with a conjunction.) We will denote the disjunction ψ' by $E_1(w_{11}, w_{12}) \vee \dots \vee E_m(w_{m1}, w_{m2})$. We define the universe of $\mathbf{B}$ as follows.

- For each sort u of $\mathbf{A}$, we define $B_u = A_u$.
- For each $v \in V$, define $B_v = \{(E_\ell, u, a) \mid \ell \in \underline{m}, u \in \{w_{\ell 1}, w_{\ell 2}\}, a \in A_u\}$

As $m \leq |V|^2$ and a variable u appearing as the second coordinate in an element of a set B_v must be an element of $\text{free}(\phi')$, we have $M(\phi, \mathbf{B}) \leq |V|^2 \cdot |\text{free}(\phi')| \cdot M(\psi, \mathbf{A})$; this confirms measure-linearity of M , since $|V|$ and $|\text{free}(\phi')|$ are functions of the pair (ψ, ϕ) .

For each atom $R(u_1, \dots, u_k)$ in ϕ that occurs outside of ϕ' (equivalently, that also appears in ψ), define $R^{\mathbf{B}} = R^{\mathbf{A}}$.

For each atom $R(u_1, \dots, u_k)$ that occurs in ϕ' , define $R^{\mathbf{B}}$ to contain a tuple $(b_1, \dots, b_k) \in B_{u_1 \dots u_k}$ if and only if the following two conditions hold:

- For all $i, j \in \underline{k}$, if $u_i \in V$ and $u_j \in V$, then $b_i = b_j$.
- For all $i, j \in \underline{k}$, if $u_i \in V$, $b_i = (E_\ell, u, a)$, and $u_j \notin V$ (equivalently, $u_j \in \text{free}(\phi')$), then
 - $u_j = u$ implies $b_j = a$, and
 - $\{u_j, u\} = \{w_{\ell 1}, w_{\ell 2}\}$ implies $\mathbf{A}, \{(u, a), (u_j, b_j)\} \models E_\ell(w_{\ell 1}, w_{\ell 2})$. (Here, we use a set of pairs to denote a partial map.)

To verify that $\mathbf{A} \models \psi$ if and only if $\mathbf{B} \models \phi$, it suffices to verify that, for any assignment f defined on $\text{free}(\phi') = \text{free}(\psi')$ taking each variable u to an element of A_u , that

$$\mathbf{A}, f \models \psi' \Leftrightarrow \mathbf{B}, f \models \phi'.$$

($\Rightarrow$): There exists $\ell \in \underline{m}$ such that $\mathbf{A}, f \models E_\ell(w_{\ell 1}, w_{\ell 2})$. Pick w to be a variable in $\{w_{\ell 1}, w_{\ell 2}\}$, and consider the extension f^+ of f that sends each variable in V to $(E_\ell, w, f(w))$. It is straightforward to verify that $\mathbf{B}, f^+$ satisfies the conjunction of ϕ' ; we do so as follows. Suppose that $R(u_1, \dots, u_k)$ is an atom in this conjunction. We claim that $(f^+(u_1), \dots, f^+(u_k)) \in R^{\mathbf{B}}$. This tuple clearly satisfies the first condition in the definition of $R^{\mathbf{B}}$; to check the second condition, suppose that $i, j \in \underline{k}$ are such that $u_i \in V$ and $u_j \notin V$. We have $f^+(u_i) = (E_\ell, w, f(w))$. If $u_j = w$, then indeed $f^+(u_j) = f(w)$. If $\{u_j, w\} = \{w_{\ell 1}, w_{\ell 2}\}$, then $\{(w, f(w)), (u_j, f^+(u_j))\}$ is equal to $f \upharpoonright \{w_{\ell 1}, w_{\ell 2}\}$, and we have $\mathbf{A}, f \upharpoonright \{w_{\ell 1}, w_{\ell 2}\} \models E_\ell(w_{\ell 1}, w_{\ell 2})$ by our choice of ℓ .

($\Leftarrow$): Suppose that $\mathbf{B}, f^+$ satisfies the conjunction of ϕ' . By the definition of $\mathbf{B}$ and since ϕ' is a layered formula, f^+ maps all variables in V to the same value (E_ℓ, u, a) . There exists an atom $R(u_1, \dots, u_k)$ in ϕ' such that one of its variables u_j has the property that $\{u_j, u\} = \{w_{\ell 1}, w_{\ell 2}\}$. It follows, from the definition of $R^{\mathbf{B}}$, that $\mathbf{A}, f^+ \upharpoonright \{u, u_j\} \models E_\ell(w_{\ell 1}, w_{\ell 2})$. $\square$

Theorem 8.5. Let Ψ be a set of fully-sorted, friendly sentences such that $\text{thick}_l(\Psi)$ is unbounded. Then, either case-CLIQUE or case-co-CLIQUE slice reduces to $\text{MC}_s[\Psi]$.

Suppose that ψ is a simple friendly formula that occurs as a subformula of a formula ϕ . We say that ψ is an *existential k -clique* if it is of the form $\exists X(\bigwedge_{i=1}^n \alpha_i)$ and there exists a set V of variables, with $|V| \geq k$, that are existentially quantified (in ϕ) such that for each set $\{v, v'\} \in K(V)$, there exists an atom α_i with $\{v, v'\} \subseteq \text{free}(\alpha_i)$. We define a *universal k -clique* dually.

Proof. Let C be as defined above in the discussion. By appeal to Theorem 8.4, it suffices to prove that either case-CLIQUE or case-co-CLIQUE slice reduces to $\text{MC}_s[\text{closure}_C(\Psi)]$. When Θ is a set of layered sentences, let us use the term Θ -*relevant subformula* to refer to a layered subformula ϕ of a sentence θ in Θ . For each Ψ -relevant subformula ϕ , it holds that $\text{local-thick}_l(\phi) \leq \text{quant-thick}_l(\phi) + |\text{free}(\phi)|$ (this is since the hypergraph from which $\text{quant-thick}_l(\phi)$ is defined is the hypergraph from which $\text{local-thick}_l(\phi)$ is defined, but with the vertices in $\text{free}(\phi)$ removed). By the definition of $\text{thick}_l(\cdot)$, the quantity $\text{local-thick}_l(\phi)$ over Ψ -relevant subformulas ϕ is unbounded. Thus, over Ψ -relevant subformulas ϕ , either the quantity $\text{quant-thick}_l(\phi)$ or the quantity $|\text{free}(\phi)|$ is unbounded. We may therefore consider two cases.

Assume that $|\text{free}(\phi)|$ is unbounded over Ψ -relevant subformulas ϕ . By conditions (2) and (3) in the definition of C , we obtain that $|\text{free}(\phi)|$ is unbounded over simple $\text{closure}_C(\Psi)$ -relevant subformulas ϕ . We assume that $|\text{free}(\phi)|$ is unbounded over such simple subformulas ϕ that use existential quantification (if not, then $|\text{free}(\phi)|$ is unbounded over such simple subformulas ϕ that use universal quantification, and the argumentation is dual). We now consider two cases. If the number of universally quantified variables in $\text{free}(\phi)$ is unbounded over such subformulas ϕ , then by condition (2) of the definition of C applied to each such subformula (and the sentence in which it appears), we obtain that, for each $k \geq 1$, the set $\text{closure}_C(\Psi)$ contains a sentence that contains a universal k -clique. Otherwise, the number of existentially quantified variables in $\text{free}(\phi)$ is unbounded over such subformulas ϕ ; in this case, by an application of condition (3) followed by an application of condition (2) (again, to each such subformula ϕ and the sentence in which it appears), we obtain that, for each $k \geq 1$, the set $\text{closure}_C(\Psi)$ contains a sentence that contains an existential k -clique. In this latter case, let us explain how to exhibit a slice reduction from case-CLIQUE to $\text{MC}_s[\text{closure}_C(\Psi)]$ (in the former case, one dually obtains a reduction from case-co-CLIQUE to $\text{MC}_s[\text{closure}_C(\Psi)]$). We define U to be the set of pairs (k, θ) such that $k \geq 1$ and θ is a sentence that contains an existential k -clique as a subformula. Let us define $r(k, \theta, (V, E))$ to be the structure $\mathbf{B}$ described as follows. Let W be the set of variables that witnesses the existential k -clique. For each relation symbol R of an atom in θ that does not witness the existential k -clique, we may set $R^{\mathbf{B}}$ to either $\emptyset$ or B^r (here, r is the arity of R) in such a way that $\mathbf{B} \models \theta$ if and only if $\mathbf{B} \models \exists W(\bigwedge_j \beta_j)$, where the β_j are the atoms witnessing the existential k -clique. By defining, for each remaining relation symbol F (that is, for each relation symbol F of an atom β_j), the relation $F^{\mathbf{B}}$ to be E , we obtain that (V, E) contains a k -clique if and only if $\mathbf{B} \models \theta$.

Now assume that $\text{quant-thick}_l(\phi)$ is unbounded over Ψ -relevant subformulas ϕ . By conditions (2) and (3) in the definition of C , we obtain that $\text{quant-thick}_l(\phi)$ is unbounded over simple $\text{closure}_C(\Psi)$ -relevant subformulas ϕ . By considering the sentences based on these subformulas (and by invoking condition (1) in the definition of C), we obtain that $\text{closure}_C(\Psi)$ contains simple sentences $\exists X(\bigwedge \alpha_i)$ or contains simple sentences $\exists Y(\bigvee \alpha_i)$ such that, over these sentences, $\text{quant-thick}_l(\cdot)$ is unbounded. In the former case, the intractability result of Grohe, Schwentick and Segoufin [13] directly yields a slice reduction from case-CLIQUE to $\text{MC}_s[\text{closure}_C(\Psi)]$ where the set U contains

a pair (k, θ) when $k \geq 1$ and θ is a sentence $\exists X(\bigwedge \alpha_i)$ whose corresponding graph admits a k -by- k grid as a minor. The latter case is dual (one obtains a slice reduction from case-co-CLIQUE to $\text{MC}_s[\text{closure}_C(\Psi)]$). $\square$

Acknowledgments

The author was supported by the Spanish Project FORMALISM (TIN2007-66523), by the Basque Government Project SPE12UN050(SAI12/219), and by the University of the Basque Country under grant UFI11/45. The author thanks Montserrat Hermo and Simone Bova for useful comments.

References

- [1] I. Adler and M. Weyer. Tree-width for first order formulae. *Logical Methods in Computer Science*, 8(1), 2012.
- [2] H. L. Bodlaender. A partial k -arborescence of graphs with bounded treewidth. *Theoretical Computer Science*, 209:1–45, 1998.
- [3] S. Bova and H. Chen. The complexity of width minimization for existential positive queries. In *ICDT*, pages 235–244, 2014.
- [4] H. Chen. On the complexity of existential positive queries. *ACM Trans. Comput. Log.*, 15(1), 2014.
- [5] H. Chen and V. Dalmau. Decomposing quantified conjunctive (or disjunctive) formulas. In *LICS*, 2012.
- [6] H. Chen and M. Grohe. Constraint satisfaction with succinctly specified relations. *Journal of Computer and System Sciences*, 76(8):847–860, 2010.
- [7] H. Chen and D. Marx. Block-sorted quantified conjunctive queries. In *ICALP*, 2013.
- [8] H. Chen and M. Müller. The fine classification of conjunctive queries and parameterized logarithmic space complexity. In *PODS*, 2013.
- [9] Y. Chen, M. Thurley, and M. Weyer. Understanding the complexity of induced subgraph isomorphisms. In *ICALP 2008*, pages 587–596, 2008.
- [10] A. Durand and S. Mengel. Structural tractability of counting of solutions to conjunctive queries. In *Proceedings of the 16th International Conference on Database Theory (ICDT 2013)*, 2013.
- [11] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.
- [12] M. Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *Journal of the ACM*, 54(1), 2007.
- [13] M. Grohe, T. Schwentick, and L. Segoufin. When is the evaluation of conjunctive queries tractable? In *STOC 2001*, 2001.
- [14] P. Jonsson, V. Lagerkvist, and G. Nordh. Blowing holes in various aspects of computational problems, with applications to constraint satisfaction. In *CP*, pages 398–414, 2013.
- [15] P. Kolaitis and M. Vardi. Conjunctive-Query Containment and Constraint Satisfaction. *Journal of Computer and System Sciences*, 61:302–332, 2000.
- [16] C. Papadimitriou and M. Yannakakis. On the Complexity of Database Queries. *Journal of Computer and System Sciences*, 58(3):407–427, 1999.
- [17] M. Y. Vardi. On the complexity of bounded-variable queries. In *PODS'95*, pages 266–276, 1995.