

A Trichotomy in the Complexity of Counting Answers to Conjunctive Queries

Hubie Chen¹ and Stefan Mengel^{*2}

- 1 Universidad del País Vasco, E-20018 San Sebastián, Spain, *and* IKERBASQUE, Basque Foundation for Science, E-48011 Bilbao, Spain
- 2 LIX UMR 7161, École Polytechnique, France

Abstract

Conjunctive queries are basic and heavily studied database queries; in relational algebra, they are the select-project-join queries. In this article, we study the fundamental problem of counting, given a conjunctive query and a relational database, the number of answers to the query on the database. In particular, we study the complexity of this problem relative to sets of conjunctive queries. We present a trichotomy theorem, which shows essentially that this problem on a set of conjunctive queries is either tractable, equivalent to the parameterized CLIQUE problem, or as hard as the parameterized counting CLIQUE problem; the criteria describing which of these situations occurs is simply stated, in terms of graph-theoretic conditions.

1998 ACM Subject Classification H.2.3 [Database Management]: Languages – Query languages, F.1.3 [Computation by Abstract Devices]: Complexity measures and classes, G.2.2 [Discrete Mathematics]: Graph Theory – Hypergraphs

Keywords and phrases database theory, query answering, conjunctive queries, counting complexity

Digital Object Identifier 10.4230/LIPIcs.ICDT.2015.110

1 Introduction

Conjunctive queries are the most basic and most heavily studied database queries. They can be formalized logically as formulas consisting of a sequence of existentially quantified variables, followed by a conjunction of atomic formulas; in relational algebra, they are the *select-project-join* queries (see e.g. [1]). Ever since the landmark 1977 article of Chandra and Merlin [2], complexity-theoretic aspects of conjunctive queries have been a research subject of persistent and enduring interest which continues to the present day (as a sampling, we point to the works [13, 16, 10, 12, 18, 14, 17, 5]; see the discussions and references therein for more information). The problem of evaluating a Boolean (closed) conjunctive query on a relational database is equivalent to a number of well-known problems, including conjunctive query containment, the homomorphism problem on relational structures, and the constraint satisfaction problem [2, 13]. That this evaluation problem appears in many equivalent guises attests to the foundational and primal nature of this problem, and it has correspondingly been approached and studied from a wide variety of perspectives and motivations.

In this article, we study the fundamental problem of counting, given a conjunctive query and a relational database, the number of query answers, that is, the number of assignments that make the query true with respect to the database; we denote this problem by #CQ.

* Partially supported by a Qualcomm grant administered by École Polytechnique.



In addition to being a natural problem in its own right, let us remark that all practical query languages supported by database management systems have a counting operator. We study the complexity of $\#CQ$ relative to sets of conjunctive queries, that is, we study a problem family: each set of conjunctive queries gives rise to a restricted version of the general problem. Our objective is to determine on which sets of conjunctive queries $\#CQ$ is tractable, and more broadly, to understand the complexity behavior of the problem family at hand. Throughout, we assume that each considered set of conjunctive queries is of *bounded arity*, by which we mean that there is a constant bounding the arity of all relation symbols in all queries in the set.¹

Surprisingly, despite the natural and basic character of the counting problem $\#CQ$, the project of understanding its complexity behavior over varying sets of queries has been carried out in previous work for only two particular types of conjunctive queries. In the case of Boolean conjunctive queries, the problem $\#CQ$ specializes to the problem of deciding whether or not such a query evaluates to *true* or *false* on a database. A classification of sets of Boolean conjunctive queries was given by Grohe [12], showing essentially that this decision problem is either polynomial-time tractable, or is hard under a typical complexity-theoretic assumption from parameterized complexity, namely, that the parameterized $CLIQUE$ problem is not fixed-parameter tractable (see Theorem 9).² It is well-known that a conjunctive query can be naturally mapped to a relational structure (see Definition 7); the tractable sets of Grohe's classification are such queries whose corresponding structures have *cores* of *bounded treewidth*. The *core* of a structure \mathbf{A} is, intuitively, the smallest structure that is (in a certain sense) equivalent to the structure \mathbf{A} , and *bounded treewidth* is a graph-theoretical condition that can intuitively be taken as a notion of tree similitude. Following Grohe's work, Dalmau and Jonsson [6] studied the case of quantifier-free conjunctive queries (which they phrase as the problem of counting homomorphisms between a given pair of relational structures). They proved that bounded treewidth is the property that determines polynomial-time tractability for this case; in contrast to Grohe's theorem, the statement of their classification (Theorem 10) does not refer to the notion of core, and is proved under the assumption that the *counting* version $\#CLIQUE$ of the parameterized $CLIQUE$ problem is not fixed-parameter tractable.

In this article, we present a trichotomy theorem describing the complexity of the counting problem $\#CQ$ on each possible set of conjunctive queries. Our trichotomy theorem unifies, generalizes, and directly implies the two discussed prior classifications. This trichotomy yields that counting on a set of conjunctive queries is polynomial-time tractable, is interreducible with the $CLIQUE$ problem, or admits a reduction from and is thus as hard as the counting problem $\#CLIQUE$. In order to prove and state our trichotomy theorem, we work with a notion of *core* of a (structure associated with a) conjunctive query whose definition crucially takes into account which variables of the conjunctive query are free (Definition 12). We also use a notion of hypergraph of a conjunctive query whose vertices are the free variables of the query (Definition 21). The properties of a query set that determine which case of our trichotomy theorem applies are whether or not the cores have bounded treewidth, and whether or not the just-mentioned hypergraphs have bounded treewidth; these two conditions correspond, respectively, to the conditions that describe the dichotomies for the Boolean case

¹ It is known that when no such bound on the arity is assumed, the complexity of query evaluation can be highly sensitive to the representation of database relations [4]. In contrast, natural representations are equivalent under polynomial-time translation; thus, the study of bounded arity queries can be viewed as the investigation of the representation-independent case.

² This is equivalent to the assumption that the parameterized complexity class $W[1]$ is not contained in the parameterized complexity class FPT , which is the phrasing that Grohe employs.

and the quantifier-free case. The proof of our trichotomy draws on recent work of Durand and Mengel [8], who presented a classification for the problem $\#CQ$ based on hypergraphs (see Theorem 11).

Note that it is readily verified that the classes of queries for which we show that $\#CQ$ is tractable are equivalent to those described in recent work by Greco and Scarcello [11]. In contrast to their work, we show that the promise version of $\#CQ$ is not only fixed-parameter tractable on these classes of queries but can even be solved in polynomial time. Moreover, and more importantly, we show that these classes are the *only* classes of queries for which $\#CQ$ can be solved efficiently because all other classes of queries are intractable under standard complexity assumptions.

In order to prove and present our trichotomy, we introduce a version of the *case complexity* framework [3] which is suitable for dealing with counting problems. Among other features, this framework facilitates the presentation of reductions between parameterized problems which are restricted in terms of the permitted parameters (or *slices*); this is the type of restriction we deal with here, as the parameter of an instance of $\#CQ$ is taken to be the query, and we consider $\#CQ$ with respect to various query sets. This framework also allows the straightforward derivation of complexity consequences as a function of the computability assumption placed on the query sets; witness the derivation of Theorems 23 and 24 from Theorem 22.

2 Preliminaries

For an integer $i \geq 1$, we define π_i to be the operator that, given a tuple, returns the value in the i th coordinate.

We assume that the reader is familiar with basic graph theoretic notions. In particular, we will use some very basic properties of treewidth, which can be found, for example, in [9, Chapter 11] or in [15, Section 2.3].

2.1 Structures, homomorphisms and cores

A relational vocabulary is defined to be a set of relation symbols $\tau := \{R_1, R_2, \dots, R_\ell\}$ where each R_i has an arity r_i . A relational structure \mathbf{A} over τ is a tuple $(A, R_1^{\mathbf{A}}, \dots, R_\ell^{\mathbf{A}})$ where A is a set called the *domain* of \mathbf{A} and $R_i^{\mathbf{A}} \subseteq A^{r_i}$ is a relation of arity r_i . All vocabularies and structures in this article are assumed to be relational. We assume each structure in this article to be *finite* in that it has a finite domain. We denote structures by the bold letters, $\mathbf{A}, \mathbf{B}, \dots$, and their corresponding domains by A, B, \dots

We assume each class of structures in this article to be of bounded arity, that is, for each such class we assume there exists a constant $c \geq 1$ such that the arity of each relation of a structure in the class is at most c . Since in the bounded arity setting the sizes of all reasonable encodings of a structure are polynomially related, we do not fix a specific encoding but assume that all structures are encoded in any reasonable way.

► **Definition 1.** Let \mathbf{A} and \mathbf{B} be two structures over the same vocabulary τ . A *homomorphism* from \mathbf{A} to \mathbf{B} is a function $h : A \rightarrow B$ such that for each relation symbol $R \in \tau$ and each $t = (t_1, \dots, t_\ell) \in R^{\mathbf{A}}$ we have $(h(t_1), \dots, h(t_\ell)) \in R^{\mathbf{B}}$. A homomorphism h from \mathbf{A} to \mathbf{B} is called an *isomorphism* if h is bijective and h^{-1} is a homomorphism from \mathbf{B} to \mathbf{A} ; when such an isomorphism exists, we say that \mathbf{A} and \mathbf{B} are *isomorphic*. An isomorphism from a structure to itself is called an *automorphism*.

► **Definition 2.** Two structures \mathbf{A} and \mathbf{B} are *homomorphically equivalent* if there are homomorphisms from \mathbf{A} to \mathbf{B} and from \mathbf{B} to \mathbf{A} .

A structure is a *core* if it is not homomorphically equivalent to a proper substructure of itself. A structure \mathbf{B} is a *core of a structure* a structure \mathbf{A} if \mathbf{B} is a substructure of \mathbf{A} , \mathbf{B} is homomorphically equivalent to \mathbf{A} , and \mathbf{B} is a core.

We state two basic properties of cores of structures; due to the first, we will speak of *the* core of a structure instead of *a* core. The second seems to be folklore; a proof can be found, for example, in [9].

► **Lemma 3.** *Every structure \mathbf{A} has at least one core. Furthermore, every two cores \mathbf{B}_1 and \mathbf{B}_2 of \mathbf{A} are isomorphic.*

► **Lemma 4.** *Let \mathbf{A} and \mathbf{B} be two homomorphically equivalent structures, and let \mathbf{A}' and \mathbf{B}' be cores of \mathbf{A} and \mathbf{B} , respectively. Then \mathbf{A}' and \mathbf{B}' are isomorphic.*

2.2 Complexity theory background

Throughout, we use Σ to denote an alphabet over which strings are formed. All problems to be considered are viewed as counting problems. So, a *problem* is a mapping $Q : \Sigma^* \rightarrow \mathbb{N}$. We view decision problems as problems where, for each $x \in \Sigma^*$, it holds that $Q(x)$ is equal to 0 or 1. We use FP (as usual) to denote the class of problems (which, again, are mappings $\Sigma^* \rightarrow \mathbb{N}$) that can be computed in polynomial time. A *parameterization* is a mapping $\kappa : \Sigma^* \rightarrow \Sigma^*$. A parameterized problem is a pair (Q, κ) consisting of a problem Q and a parameterization κ . A partial function $T : \Sigma^* \rightarrow \mathbb{N}$ is *polynomial-multiplied* with respect to a parameterization κ if there exists a computable 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 $T(x) \leq f(\kappa(x))p(|x|)$.

► **Definition 5.** Let $\kappa : \Sigma^* \rightarrow \Sigma^*$ be a parameterization. A partial mapping $r : \Sigma^* \rightarrow \Sigma^*$ is *FPT-computable* with respect to κ if there exist a polynomial-multiplied function $T : \Sigma^* \rightarrow \mathbb{N}$ (with respect to κ) with $\text{dom}(T) = \text{dom}(r)$ and an algorithm A such that, for each string $x \in \text{dom}(r)$, the algorithm A computes $r(x)$ within time $T(x)$; when this holds, we also say that r is *FPT-computable* with respect to κ *via* A .

As is standard, we may and do freely interchange among elements of Σ^* , $\Sigma^* \times \Sigma^*$, and \mathbb{N} . We define FPT to be the class that contains a parameterized problem (Q, κ) if and only if Q is FPT-computable with respect to κ .

We now introduce a notion of reduction for counting problems, which is a form of Turing reduction. We use $\wp_{\text{fin}}(A)$ to denote the set containing all finite subsets of A .

► **Definition 6.** A *counting FPT-reduction* from a parameterized problem (Q, κ) to a second parameterized problem (Q', κ') consists of a computable function $h : \Sigma^* \rightarrow \wp_{\text{fin}}(\Sigma^*)$, and an algorithm A such that:

- on an input x , A may make oracle queries of the form $Q'(y)$ with $\kappa'(y) \in h(\kappa(x))$, and
- Q is FPT-computable with respect to κ via A .

We use CLIQUE to denote the decision problem where (k, G) is a yes-instance when G is a graph that contains a clique of size $k \in \mathbb{N}$. By #CLIQUE we denote the problem of counting, given (k, G) , the number of k -cliques in the graph G . The parameterized versions of these problems, denoted by p -CLIQUE and p -#CLIQUE, are defined via the parameterization π_1 . We will make tacit use of the following well-known facts: FPT is closed under counting FPT-reduction; p -CLIQUE is complete for $W[1]$ under counting FPT-reduction; and, p -#CLIQUE complete for # $W[1]$ under counting FPT-reduction.

A *promise problem* is a pair $\langle Q, I \rangle$ where Q is a problem and $I \subseteq \Sigma^*$. When C is a complexity class, define $\text{prom-}C$ to contain a promise problem $\langle Q, I \rangle$ when there exists $P \in C$ such that, for all $x \in I$, it holds that $P(x) = Q(x)$. A *parameterized promise problem* is a pair $\langle \langle Q, I \rangle, \kappa \rangle$ consisting of a promise problem $\langle Q, I \rangle$ and a parameterization κ ; such a problem will also be notated by $\langle \langle Q, \kappa \rangle, I \rangle$. When C is a parameterized complexity class, define $\text{prom-}C$ to contain a promise problem $\langle \langle Q, I \rangle, \kappa \rangle$ when there exists a problem P such that $(P, \kappa) \in C$ and for all $x \in I$, it holds that $P(x) = Q(x)$.

3 Conjunctive Queries and Computational Problems

A *conjunctive query* is a relational first-order formula (possibly with free variables) of the form $\exists v_1 \dots \exists v_n \bigwedge_{i=1}^m \alpha_i$ where each α_i is a predicate application, that is, an atomic formula of the form $R(\vec{u})$ where R is a relation symbol and \vec{u} is a tuple of variables. Since the only type of queries that we are concerned with in this article are conjunctive queries, we will sometimes simply use *query* to refer to a conjunctive query.

► **Definition 7.** To a conjunctive query ϕ over the vocabulary τ , we assign a structure $\mathbf{A} = \mathbf{A}_\phi$, called the *natural model*, as follows: the domain of \mathbf{A} is $\text{var}(\phi)$; \mathbf{A} is over the vocabulary τ ; and for each relation symbol $R \in \tau$, we set $R^{\mathbf{A}} := \{\vec{a} \mid R(\vec{a}) \text{ is an atom of } \phi\}$.

To each conjunctive query ϕ we assign the pair (\mathbf{A}, S) where \mathbf{A} is the natural model of ϕ and S the set of its free variables. From such a pair (\mathbf{A}, S) , it is easy to reconstruct the corresponding query ϕ : each tuple of a relation of \mathbf{A} is made into an atom, and then, one existentially quantifies the elements of A not in S to obtain ϕ . Because of this easy correspondence between queries and pairs (\mathbf{A}, S) , in a slight abuse of notation, we do not differentiate between pairs (\mathbf{A}, S) and queries throughout. In particular, we will call a pair (\mathbf{A}, S) a query, and we will use \mathcal{C} interchangeably for classes of queries and of pairs (\mathbf{A}, S) .

Let ϕ be a conjunctive query with assigned pair (\mathbf{A}, S) and let \mathbf{B} be a structure. Then a function $h : S \rightarrow B$ is a satisfying assignment of ϕ if and only if it can be extended to a homomorphism from \mathbf{A} to \mathbf{B} ; we denote the set of such functions by $\text{hom}(\mathbf{A}, \mathbf{B}, S)$. In this article, we are interested in the following counting problem.

#CQ

Input: A query (\mathbf{A}, S) and a structure \mathbf{B} .

Problem: Compute $|\text{hom}(\mathbf{A}, \mathbf{B}, S)|$.

In the case of a conjunctive query (\mathbf{A}, \emptyset) without free variables, the problem #CQ amounts to deciding whether or not there exists a homomorphism from \mathbf{A} to \mathbf{B} . We define this case as the problem CQ.

CQ

Input: A query (\mathbf{A}, \emptyset) and a structure \mathbf{B} .

Problem: Decide if there exists a homomorphism from \mathbf{A} to \mathbf{B} .

We define $p\text{-\#CQ}$ to be the parameterized problem $(\text{\#CQ}, \pi_1)$, that is, we take the parameter of each instance $((\mathbf{A}, S), \mathbf{B})$ to be (\mathbf{A}, S) . (Formally, we view each instance of $p\text{-\#CQ}$ as a pair of strings, where the first component encodes the query, and the second component encodes the structure.) In analogy to #CQ, we define $p\text{-CQ}$ to be the parameterized problem (CQ, π_1) .

We define the hypergraph of a query (\mathbf{A}, S) to be the hypergraph $\mathcal{H} = (V, E)$ where V is the domain of \mathbf{A} and $E := \{\text{dom}(t) \mid t \in R^{\mathbf{A}}, R^{\mathbf{A}} \text{ is a relation of } \mathbf{A}\}$ where $\text{dom}(t)$

denotes the set of elements appearing in t . The treewidth of (\mathbf{A}, S) is defined to be that of its hypergraph. Dalmau, Kolaitis and Vardi [7] proved that CQ can be solved efficiently, when the treewidth of the cores of the queries is bounded.

► **Theorem 8** ([7]). *Let $k \in \mathbb{N}$ be a fixed constant. Let \mathcal{C}_k be the class of all structures with cores of treewidth at most k . Then the promise problem $\langle \text{CQ}, \mathcal{C}_k \times \Sigma^* \rangle$ is in prom-FP.*

Grohe [12] showed that this result is optimal.

► **Theorem 9** ([12]). *Let \mathcal{C} be a recursively enumerable class of structures of bounded arity. Assume $\text{FPT} \neq \text{W}[1]$. Then the following statements are equivalent:*

1. $\langle \text{CQ}, \mathcal{C} \times \Sigma^* \rangle \in \text{prom-FP}$.
2. $\langle p\text{-CQ}, \mathcal{C} \times \Sigma^* \rangle \in \text{prom-FPT}$.
3. *There is a constant c such that the cores of the structures in \mathcal{C} have treewidth at most c .*

Dalmau and Jonsson [6] considered the analogous question for #CQ for quantifier free queries and found that cores do not help in this setting.

► **Theorem 10** ([6]). *Let \mathcal{Q} be the class of all quantifier free conjunctive queries, i.e., queries of the form (\mathbf{A}, A) . Let \mathcal{C} be a recursively enumerable class of structures of bounded arity in \mathcal{Q} . Assume $\text{FPT} \neq \#\text{W}[1]$. Then the following statements are equivalent:*

1. $\langle \#\text{CQ}, \mathcal{C} \times \Sigma^* \rangle \in \text{prom-FP}$.
2. $\langle p\text{-}\#\text{CQ}, \mathcal{C} \times \Sigma^* \rangle \in \text{prom-FPT}$.
3. *There is a constant c such that the structures in \mathcal{C} have treewidth at most c .*

It is common to consider classes of queries defined by restricting their associated hypergraph. For #CQ it turns out to be helpful to also encode which vertices of a hypergraph correspond to free variables, which is formalized in the following definition. A pair (\mathcal{H}, S) where \mathcal{H} is a hypergraph and S is a subset of the vertices of \mathcal{H} is called an S -hypergraph. The S -hypergraph of a query (\mathbf{A}, S) is (\mathcal{H}, S) where \mathcal{H} is the hypergraph of \mathbf{A} .

In [8], the following version of #CQ is considered.

$\#\text{CQ}_{hyp}$

Input: An S -hypergraph (\mathcal{H}, S) and an instance $((\mathbf{A}, S), \mathbf{B})$ of #CQ where \mathcal{H} is the hypergraph of \mathbf{A} .

Problem: Compute $|\text{hom}(\mathbf{A}, \mathbf{B}, S)|$.

We define $p\text{-}\#\text{CQ}_{hyp}$ to be the parameterized problem $(\#\text{CQ}_{hyp}, \pi_1)$; here, an instance of $\#\text{CQ}_{hyp}$ is viewed as a pair $((\mathcal{H}, S), ((\mathbf{A}, S), \mathbf{B}))$, on which the operator π_1 returns (\mathcal{H}, S) .

It turns out that for $\#\text{CQ}_{hyp}$ a parameter called S -star size is of critical importance. Let $\mathcal{H} = (V, E)$ be a hypergraph and $S \subseteq V$. Let C be the vertex set of a connected component of $\mathcal{H}[V - S]$. Let E_C be the set of hyperedges $\{e \in E \mid e \cap C \neq \emptyset\}$ and $V_C := \bigcup_{e \in E_C} e$. Then $\mathcal{H}[V_C]$ is called an S -component of \mathcal{H} . The size of a biggest independent set in $\mathcal{H}[V_C \cap S]$ is called the S -star size of the S -component $\mathcal{H}[V_C]$. The S -star size of (\mathcal{H}, S) is then defined to be the maximum S -star size taken over all S -components of (\mathcal{H}, S) . By the *quantified star size* of a query (\mathbf{A}, S) we refer to the S -star size of the S -hypergraph associated to (\mathbf{A}, S) . For more explanations on these notions and examples see [15, Section 3.2].

► **Theorem 11** ([8]). *Let \mathcal{G} be a recursively enumerable class of S -hypergraphs of bounded arity. Assume that $\text{W}[1] \neq \text{FPT}$. Then the following statements are equivalent:*

1. $\langle \#\text{CQ}_{hyp}, \mathcal{G} \times \Sigma^* \rangle \in \text{prom-FP}$.
2. $\langle p\text{-}\#\text{CQ}_{hyp}, \mathcal{G} \times \Sigma^* \rangle \in \text{prom-FPT}$.

3. *There is a constant c such that for each S -hypergraph (\mathcal{H}, S) in \mathcal{G} the treewidth of \mathcal{H} and the S -star size of \mathcal{H} are at most c .*

We have seen that for the problem CQ, cores of structures are crucial, while in the classification due to Dalmau and Jonsson, they do not matter at all. Thus we introduce a notion of cores for conjunctive queries that interpolates between these two extreme cases. The idea behind the definition is that we require the homomorphisms between (\mathbf{A}, S) and its core to be the identity on the free variables, while they may map the quantified variables in any way that leads to a homomorphism. This is formalized as follows.

► **Definition 12.** For a conjunctive query (\mathbf{A}, S) where \mathbf{A} is defined on vocabulary τ , we define the *augmented structure*, denoted by $\text{aug}(\mathbf{A}, S)$, to be the structure over the vocabulary $\tau \cup \{R_a \mid a \in S\}$ where $R_a^{\text{aug}(\mathbf{A}, S)} := \{a\}$. We define the *core of a conjunctive query* of (\mathbf{A}, S) to be the core of $\text{aug}(\mathbf{A}, S)$.

► **Example 13.** Let (\mathbf{A}, S) be a query without free variables, that is, where $S = \emptyset$; then the core of (\mathbf{A}, S) is the core of \mathbf{A} . If (\mathbf{A}, S) is quantifier-free, that is, where $S = A$, then the core of (\mathbf{A}, S) equals \mathbf{A} .

The cores of conjunctive queries were essentially already studied by Chandra and Merlin in a seminal paper [2] although the notation used there is different. We give a fundamental result on conjunctive queries. We call two queries (\mathbf{A}_1, S) and (\mathbf{A}_2, S) *equivalent queries* if for each structure \mathbf{B} we have $\text{hom}(\mathbf{A}_1, \mathbf{B}, S) = \text{hom}(\mathbf{A}_2, \mathbf{B}, S)$.

► **Theorem 14** ([2]). *If two conjunctive queries (\mathbf{A}_1, S) and (\mathbf{A}_2, S) have the same core (up to isomorphism), then they are equivalent.*

4 Case Complexity

In this section we develop a version of the case complexity framework advocated in [3] which is suitable for classifying counting problems. A main motivation for this framework is the growing amount of research on parameterized problems which are restricted by the permitted values of the parameter. In particular, this kind of problem arises naturally in query answering problems where one often restricts the admissible queries for the inputs (see e.g. [12, 6, 3]). An aim of the case complexity framework as introduced in [3] is to facilitate reductions between the considered restricted parameterized problems and to show results independent of computability assumptions for the parameter.

The central notion for our framework is the following: A *case problem* consists of a problem $Q : \Sigma^* \times \Sigma^* \rightarrow \mathbb{N}$ and a subset $S \subseteq \Sigma^*$, and is denoted $Q[S]$. When $Q[S]$ is a case problem, we define the following:

- **param- $Q[S]$** is the parameterized problem (P, π_1) where $P(s, x)$ is defined as equal to $Q(s, x)$ if $s \in S$, and as 0 otherwise.
- **prom- $Q[S]$** is the promise problem $(Q, S \times \Sigma^*)$.
- **param-prom- $Q[S]$** is the parameterized promise problem $(\text{param-}Q[S], \pi_1)$.

The case problem we consider in this paper will nearly exclusively be $\#\text{CQ}[\mathcal{C}]$ where \mathcal{C} is a class a class of conjunctive queries. Nevertheless, we stress the fact that our framework is fully generic and we believe that it will in the future also be useful for presenting and proving complexity classifications for other problems.

We now introduce a reduction notion for case problems.

► **Definition 15.** A *counting slice reduction* from a case problem $Q[S]$ to a second case problem $Q'[S']$ consists of

- a computably enumerable language $U \subseteq \Sigma^* \times \wp_{\text{fin}}(\Sigma^*)$, and
- a partial function $r : \Sigma^* \times \wp_{\text{fin}}(\Sigma^*) \times \Sigma^* \rightarrow \Sigma^*$ that has domain $U \times \Sigma^*$ and is FPT-computable with respect to (π_1, π_2) via an algorithm A that, on input (s, T, y) , may make queries of the form $Q'(t, z)$ where $t \in T$,

such that the following conditions hold:

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

$$Q(s, y) = r(s, T, y).$$

As usual in counting complexity, it will often not be necessary to use the full generality of counting slice reductions. Therefore, we introduce a second, parsimonious notion of reductions for case problems which is often general enough but easier to deal with.

► **Definition 16.** A *parsimonious slice reduction* from a case problem $Q[S]$ to a second case problem $Q'[S']$ consists of

- a computably enumerable language $U \subseteq \Sigma^* \times \Sigma^*$, and
- a partial function $r : \Sigma^* \times \Sigma^* \times \Sigma^* \rightarrow \Sigma^*$ that has domain $U \times \Sigma^*$ and is FPT-computable with respect to (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

$$Q(t, y) = Q'(t', r(t, t', y)).$$

We give some basic properties of counting slice reductions. Their proofs can be found in the full version of this paper.

► **Proposition 17.** *If $Q[S]$ parsimoniously slice reduces to $Q'[S']$, then $Q[S]$ counting slice reduces to $Q'[S']$.*

► **Theorem 18.** *Counting slice reducibility is transitive.*

The next two theorems give the connection between case complexity and parameterized complexity. In particular, they show that, from a counting slice reduction, one can obtain complexity results for the corresponding parameterized problems.

► **Theorem 19.** *Let $Q[S]$ and $Q'[S']$ be case problems. Suppose that $Q[S]$ counting slice reduces to $Q'[S']$, and that both S and S' are computable. Then $\text{param-}Q[S]$ counting FPT-reduces to $\text{param-}Q'[S']$.*

► **Theorem 20.** *Let $Q[S]$ be a case problem, and let $K : \Sigma^* \times \Sigma^* \rightarrow \mathbb{N}$ be a problem. Suppose that $\text{param-prom-}Q[S]$ is in prom-FPT , S is computably enumerable, and that the case problem $K[\Sigma^*]$ counting slice reduces to $Q[S]$. Then the parameterized problem (K, π_1) is in FPT .*

In the remainder of the paper, we will show all our reductions in the case complexity framework and then use Theorem 19 and Theorem 20 to derive parameterized complexity results. This approach lets us give results on $\#\text{CQ}[\mathcal{C}]$ for different complexity assumptions on \mathcal{C} without having to deal with these assumptions in the proofs. Thus we separate the technicalities of the reductions from the assumptions on \mathcal{C} which in our opinion gives a far clearer presentation.

5 Statement of the main results

In this section we present the main results of this paper which we will then prove in the remainder of the paper. For the statement of the results we will use certain S -hypergraphs that we get as a contraction of the S -hypergraphs of conjunctive queries. When deleting a vertex v from a hypergraph, we delete v from the vertex set and all edges it appears in but keep all edges, unless they become empty after the deletion of v .

► **Definition 21.** To every S -hypergraph (\mathcal{H}, S) we define an S -hypergraph $\text{contract}(\mathcal{H}, S)$ as follows: We add an edge $\{u, v\}$ for any pair of vertices u, v that appears in a common S -component of \mathcal{H} . Then we delete the vertices in $V(\mathcal{H}) \setminus S$. To a class \mathcal{G} of S -hypergraphs we define $\text{contract}(\mathcal{G}) := \{\text{contract}(\mathcal{H}, S) \mid (\mathcal{H}, S) \in \mathcal{G}\}$.

For a conjunctive query (\mathbf{A}, S) let $\text{contract}(\mathbf{A}, S)$ be $\text{contract}(\mathcal{H}, S)$ where (\mathcal{H}, S) is the S -hypergraph of the core of (\mathbf{A}, S) . For a class \mathcal{C} of conjunctive queries, set $\text{contract}(\mathcal{C}) := \{\text{contract}(\mathbf{A}, S) \mid (\mathbf{A}, S) \in \mathcal{C}\}$.

We first present a version of our main result using the framework of case complexity.

► **Theorem 22.** *Let \mathcal{C} be a class of conjunctive queries.*

1. *If the cores of \mathcal{C} and $\text{contract}(\mathcal{C})$ are of bounded treewidth, then $\text{prom-}\#\text{CQ}[\mathcal{C}] \in \text{prom-FP}$, and hence $\text{param-prom-}\#\text{CQ}[\mathcal{C}] \in \text{prom-FPT}$.*
2. *If the cores of \mathcal{C} are of unbounded treewidth but $\text{contract}(\mathcal{C})$ is of bounded treewidth, then $\#\text{CQ}[\mathcal{C}]$ is equivalent to $\text{CLIQUE}[\mathbb{N}]$ with respect to counting slice reductions.*
3. *If the treewidth of $\text{contract}(\mathcal{C})$ is unbounded, then there is a counting slice reduction from $\#\text{CLIQUE}[\mathbb{N}]$ to $\#\text{CQ}[\mathcal{C}]$.*

We will prove Theorem 22 in Section 7.5. Using the results on case complexity, we derive from Theorem 22 two versions of the trichotomy phrased in terms of promise problems and of non-promise problems, depending on whether or not the class \mathcal{C} of conjunctive queries is assumed to be recursively enumerable or computable, respectively.

► **Theorem 23.** *Let \mathcal{C} be a class of conjunctive queries which is recursively enumerable. In the scope of this theorem, let us say that the class \mathcal{C} is tractable if $\langle \#\text{CQ}, \mathcal{C} \times \Sigma^* \rangle \in \text{prom-FP}$ and $\langle p\text{-}\#\text{CQ}, \mathcal{C} \times \Sigma^* \rangle \in \text{prom-FPT}$.*

1. *If the cores of \mathcal{C} and $\text{contract}(\mathcal{C})$ have bounded treewidth, then \mathcal{C} is tractable.*
2. *If the cores of \mathcal{C} have unbounded treewidth, then \mathcal{C} is not tractable, unless $p\text{-}\text{CLIQUE}$ is in FPT (and hence $\text{FPT} = \text{W}[1]$).*
3. *If $\text{contract}(\mathcal{C})$ has unbounded treewidth, then \mathcal{C} is not tractable, unless $p\text{-}\#\text{CLIQUE}$ is in FPT (and hence $\text{FPT} = \#\text{W}[1]$).*

Proof. (1) follows directly from item (1) of Theorem 22. (2) and (3) follow directly from the respective items of Theorem 22 and Theorem 20. ◀

From Theorem 23, one can immediately derive, as corollaries, Theorem 9 and Theorem 10.

Let us use $(p\text{-}\#\text{CQ} \upharpoonright I)$ to denote the parameterized problem which is equal to $\#\text{CQ}$ on I , and is equal to 0 elsewhere (and has the parameterization of $p\text{-}\#\text{CQ}$).

► **Theorem 24.** *Let \mathcal{C} be a class of conjunctive queries which is computable.*

1. *If the cores of \mathcal{C} and $\text{contract}(\mathcal{C})$ have bounded treewidth, then $(p\text{-}\#\text{CQ} \upharpoonright \mathcal{C} \times \Sigma^*)$ is in FPT.*
2. *If the cores of \mathcal{C} have unbounded treewidth, and $\text{contract}(\mathcal{C})$ has bounded treewidth, then $(p\text{-}\#\text{CQ} \upharpoonright \mathcal{C} \times \Sigma^*)$ is equivalent to $p\text{-}\text{CLIQUE}$ under counting FPT-reduction.*

3. If $\text{contract}(\mathcal{C})$ has unbounded treewidth, $(p\text{-}\#\text{CQ} \upharpoonright \mathcal{C} \times \Sigma^*)$ admits a counting FPT-reduction from $p\text{-}\#\text{CLIQUE}$.

Proof. For (1), the FPT algorithm is to first decide, given an instance (ϕ, \mathbf{B}) , whether or not $\phi \in \mathcal{C}$; if so, the algorithm invokes the algorithm of Theorem 22, otherwise, it returns 0. (2) and (3) follow immediately from Theorem 22 and Theorem 19. \blacktriangleleft

6 Positive complexity results

In this section, we will prove a counting version of Theorem 8. We will use a lemma that is probably well known, but as we could not find a reference, we give a proof for it in the full version of this paper.

► **Lemma 25.** *Let $k \in \mathbb{N}$ be a fixed constant. There exists a polynomial-time algorithm that, given a structure \mathbf{A} whose core has treewidth at most k , outputs a core of \mathbf{A} .*

Lemma 25 yields a counting version of Theorem 8 as an easy corollary.

► **Corollary 26.** *Let \mathcal{C} be a class of conjunctive queries such that the cores of the queries in \mathcal{C} have bounded quantified star size and bounded treewidth. Then $\text{prom-}\#\text{CQ}[\mathcal{C}] \in \text{prom-FP}$.*

Proof. Let $((\mathbf{A}, S), \mathbf{B})$ be an instance of $\text{prom-}\#\text{CQ}[\mathcal{C}]$ with domain A . By the promise, there is a constant c such that the treewidth and the quantified star size of the core of (\mathbf{A}, S) are at most c . We simply compute the core of (\mathbf{A}, S) with Lemma 25 and delete from it the relations R_a introduced when constructing $\text{aug}(\mathbf{A})$. Call the resulting query $(\hat{\mathbf{A}}, S)$. By construction, (\mathbf{A}, S) and $(\hat{\mathbf{A}}, S)$ have the same core, so by Theorem 14 are equivalent. Moreover, the treewidth and quantified star size of $(\hat{\mathbf{A}}, S)$ are bounded by c and thus Theorem 11 lets us solve the instance in polynomial time. \blacktriangleleft

Let us discuss Corollary 26.

Theorem 8 and thus also Lemma 25 crucially depends on the fact that we know by an outside promise that the treewidth of the cores we consider is bounded. If this bound is not satisfied, then the algorithm of Theorem 8 may give false positive results. Consequently, the algorithm of Lemma 25 may compute a structure that is in fact *not* the core of the input and then the algorithm of Corollary 26 gives the wrong count. Unfortunately, deciding if the core of a conjunctive query has treewidth at most k is NP-complete [7] and even the problem of deciding if a fixed structure is the core of a given structure is NP-complete. Thus there is no efficient way of realizing that the core computed by the algorithm of Lemma 25 is wrong.

Consequently, while the result of Corollary 26 is very nice from a theoretical point of view (we will see in the next section that it is in fact optimal), it is probably of limited value from a more practical perspective. We see this as evidence that in fact parameterized complexity is a framework better suited for the type of problem discussed in this paper. Note that in this more relaxed setting of parameterized complexity, computing the core of a query by brute force can easily be done in the allowed time, because the core depends only on the query which is the parameter.

We now present a counting algorithm for $\#\text{CQ}[\mathcal{C}]$ for certain classes \mathcal{C} that has oracle access to CQ, the decision version of $\#\text{CQ}$. Let ALL be the class of all conjunctive queries.

► **Lemma 27.** *Let \mathcal{C} be a class of queries such that the treewidth of the S -hypergraphs in $\text{contract}(\mathcal{C})$ is bounded by a constant c . Then there is a counting slice reduction from $\#\text{CQ}[\mathcal{C}]$ to $\text{CQ}[\text{ALL}]$.*

The idea of the proof is as follows: Since the treewidth of $\text{contract}(\mathcal{C})$ is bounded, we know that the unbounded treewidth of the cores does not originate from the structure of the free variables but only from the way the quantified variables interact in the S -components. We use the oracle for $\text{CQ}[\text{ALL}]$ to solve subqueries of the original query in order to “contract” the quantified variables into one variable per S -component. This results in an instance with the same solutions that has bounded treewidth. We then solve this instance with the algorithm of Theorem 11. The complete proof of Lemma 27 can be found in the full version of this paper.

7 Hardness results

In this section we will prove the hardness results for Theorem 22. The main idea is reducing from the hard cases of Theorem 11 in several steps.

7.1 Simulating unary relations

In this section we show that for queries whose augmented structure is a core we can simulate unary relations on the variables of the query. These additional relations will later allow us to tell the variables apart such that we can later simulate the case in which all atoms of the queries have different relation symbols.

► **Lemma 28.** *Let (\mathbf{A}, S) be a conjunctive query such that $\text{aug}(\mathbf{A}, S)$ is a core. Then every homomorphism $h : \mathbf{A} \rightarrow \mathbf{A}$ with $h|_S = \text{id}$ is a bijection.*

Proof. Clearly, h is also a homomorphism $h : \text{aug}(\mathbf{A}, S) \rightarrow \text{aug}(\mathbf{A}, S)$, because $h(a) = a \in R_a^{\text{aug}(\mathbf{A}, S)}$ for every $a \in S$. But by assumption $\text{aug}(\mathbf{A}, S)$ is a core, so there is no homomorphism from $\text{aug}(\mathbf{A}, S)$ to a proper substructure and thus h must be a bijection on $\text{aug}(\mathbf{A}, S)$ and consequently also on \mathbf{A} . ◀

We assign a structure \mathbf{A}^* to every structure \mathbf{A} :

► **Definition 29.** To a structure \mathbf{A} we assign the structure \mathbf{A}^* over the vocabulary $\tau \cup \{R_a \mid a \in A\}$ defined as $\mathbf{A}^* := \mathbf{A} \cup \bigcup_{a \in A} R_a^{\mathbf{A}^*}$ where $R_a^{\mathbf{A}^*} := \{a\}$.

Note that $\text{aug}(\mathbf{A}, S)$ and \mathbf{A}^* differ in which relations we add: For the structure $\text{aug}(\mathbf{A}, S)$ we add $R_a^{\text{aug}(\mathbf{A}, S)}$ for variables $a \in S$ while for \mathbf{A}^* we add $R_a^{\mathbf{A}^*}$ for all $a \in A$. Thus, \mathbf{A}^* in general may have more relations than $\text{aug}(\mathbf{A}, S)$.

We now formulate the main lemma of this section whose proof uses ideas from [6].

► **Lemma 30.** *Let \mathcal{C} be a class of conjunctive queries such that for each $(\mathbf{A}, S) \in \mathcal{C}$ the augmented structure $\text{aug}(\mathbf{A}, S)$ is a core. Let $\mathcal{C}^* := \{(\mathbf{A}^*, S) \mid (\mathbf{A}, S) \in \mathcal{C}\}$. Then there is a counting slice reduction from $p\text{-}\#\text{CQ}[\mathcal{C}^*]$ to $p\text{-}\#\text{CQ}[\mathcal{C}]$.*

Proof. Let $((\mathbf{A}^*, S), \mathbf{B})$ be an input for $\#\text{CQ}[\mathcal{C}^*]$. Remember that \mathbf{A}^* and \mathbf{B} are structures over the vocabulary $\tau \cup \{R_a \mid a \in A\}$. For every query (\mathbf{A}, S) , the relation U of our counting slice reduction contains $((\mathbf{A}^*, S), (\mathbf{A}, S))$. Obviously, U is computable and satisfies the coverage property.

We now will reduce the computation of the size $|\text{hom}(\mathbf{A}^*, \mathbf{B}, S)|$ to the computation of $|\text{hom}(\mathbf{A}, \mathbf{B}', S)|$ for different structures \mathbf{B}' .

Let $D := \{(a, b) \in A \times B \mid b \in R_a^{\mathbf{B}}\}$ and define a structure \mathbf{D} over the vocabulary τ with the domain D that contains for each relation symbol $R \in \tau$ the relation

$$R^{\mathbf{D}} := \{((a_1, b_1), \dots, (a_r, b_r)) \mid (a_1, \dots, a_r) \in R^{\mathbf{A}}, (b_1, \dots, b_r) \in R^{\mathbf{B}}, \\ \forall i \in [r] : (a_i, b_i) \in D\}.$$

Let again $\pi_1 : D \rightarrow A$ be the projection onto the first coordinate, i.e., $\pi_1(a, b) := a$. Observe that π_1 is by construction of \mathbf{D} a homomorphism from \mathbf{D} to \mathbf{A} .

We will several times use the following claim:

► **Claim 1.** *Let h be a homomorphism from \mathbf{A} to \mathbf{D} with $h(S) = S$. Then $\pi_1 \circ h$ is an automorphism of \mathbf{A} .*

Proof. Let $g := \pi_1 \circ h$. As the composition of two homomorphisms, g is a homomorphism from \mathbf{A} to \mathbf{A} . Furthermore, by assumption $g|_S$ is a bijection from S to S . Since S is finite, there is $i \in \mathbb{N}$ such that $g^i|_S = \text{id}$. But g^i is a homomorphism and thus, by Lemma 28, g^i is a bijection. It follows that g is a bijection.

Since A is finite, there is $j \in \mathbb{N}$ such that $g^{-1} = g^j$. It follows that g^{-1} is a homomorphism and thus g is an automorphism. ◀

Let \mathcal{N} be the set of mappings $h : S \rightarrow D$ with $\pi_1 \circ h = \text{id}$ that can be extended to a homomorphism $h' : \mathbf{A} \rightarrow \mathbf{D}$.

► **Claim 2.** *There is a bijection between $\text{hom}(\mathbf{A}^*, \mathbf{B}, S)$ and \mathcal{N} .*

Proof. For each $h^* \in \text{hom}(\mathbf{A}^*, \mathbf{B}, S)$ we define $P(h^*) := h$ by $h(a) := (a, h^*(a))$ for $a \in S$. From the extension of h^* to A we get an extension of h that is a homomorphism and thus $h \in \mathcal{N}$. Thus P is a mapping $P : \text{hom}(\mathbf{A}^*, \mathbf{B}, S) \rightarrow \mathcal{N}$.

We claim that P is a bijection. Clearly, P is injective. We will show that it is surjective as well. To this end, let $h : S \rightarrow D$ be a mapping in \mathcal{N} and let h_e be a homomorphism from \mathbf{A} to \mathbf{D} that is an extension of h . By definition of \mathcal{N} such a h_e must exist. By Claim 1 we have that $\pi_1 \circ h_e$ is an automorphism, and thus $(\pi_1 \circ h_e)^{-1}$ is a homomorphism. We set $h'_e := h_e \circ (\pi_1 \circ h_e)^{-1}$. Obviously, h'_e is a homomorphism from \mathbf{A} to \mathbf{D} , because h'_e is the composition of two homomorphisms. Furthermore, for all $a \in S$ we have $h'_e(a) = (h_e \circ (\pi_1 \circ h_e)^{-1})(a) = (h_e \circ (\pi_1 \circ h)^{-1})(a) = h_e(a) = h(a)$, so h'_e is an extension of h . Moreover $\pi_1 \circ h'_e = (\pi_1 \circ h_e) \circ (\pi_1 \circ h_e)^{-1} = \text{id}$. Hence, we have $h'_e = \text{id} \times \hat{h}$ for a homomorphism $\hat{h} : \mathbf{A} \rightarrow \hat{\mathbf{B}}$, where $\hat{\mathbf{B}}$ is the structure we get from \mathbf{B} by deleting the relations $R_a^{\mathbf{B}}$ for $a \in A$. But by definition $h'_e(a) \in D$ for all $a \in A$ and thus $\hat{h}(a) \in R_a^{\hat{\mathbf{B}}}$. It follows that \hat{h} is a homomorphism from \mathbf{A}^* to $\hat{\mathbf{B}}$. We set $h^* := \hat{h}|_S$. Clearly, $h^* \in \text{hom}(\mathbf{A}^*, \mathbf{B}, S)$ and $P(h^*) = h$. It follows that P is surjective. This proves the claim. ◀

Let I be the set of mappings $g : S \rightarrow S$ that can be extended to an automorphism of \mathbf{A} . Let \mathcal{N}' be the set of mappings $h : S \rightarrow D$ with $(\pi_1 \circ h)(S) = S$ that can be extended to homomorphisms $h' : \mathbf{A} \rightarrow \mathbf{D}$.

► **Claim 3.**

$$|\text{hom}(\mathbf{A}^*, \mathbf{B}, S)| = \frac{|\mathcal{N}'|}{|I|}.$$

Proof. Because of Claim 2 it is sufficient to show that

$$|\mathcal{N}'| = |\mathcal{N}||I|. \quad (1)$$

We first prove that

$$\mathcal{N}' = \{f \circ g \mid f \in \mathcal{N}, g \in I\}. \quad (2)$$

The \supseteq direction is obvious. For the other direction let $h \in \mathcal{N}'$. Let h' be the extension of h that is a homomorphism $h' : \mathbf{A} \rightarrow \mathbf{D}$. By Claim 1, we have that $g := \pi_1 \circ h'$ is an automorphism of \mathbf{A} . It follows that $g^{-1}|_S \in I$. Furthermore, $h \circ g^{-1}|_S$ is a mapping from S to D and $h' \circ g^{-1}$ is an extension that is a homomorphism from \mathbf{A} to \mathbf{D} . Furthermore $(\pi_1 \circ h' \circ g^{-1}|_S)(a) = (g|_S \circ g^{-1}|_S)(a) = a$ for every $a \in S$ and hence $h' \circ g^{-1}|_S \in \mathcal{N}$ and $h = h \circ g^{-1}|_S \circ g|_S$ which proves the claim (2).

To show (1), we claim that for every $f, f' \in \mathcal{N}$ and every $g, g' \in I$, if $f \neq f'$ or $g \neq g'$, then $f \circ g \neq f' \circ g'$. To see this, observe that f can always be written as $f = \text{id} \times f_2$ and thus $(f \circ g)(a) = (g(a), f_2(g(a)))$. Thus, if g and g' differ, $\pi_1 \circ f \circ g \neq \pi_1 \circ f' \circ g'$ and thus $f \circ g \neq f' \circ g'$. Also, if $g = g'$ and $f \neq f'$, then clearly $f \circ g \neq f' \circ g'$. This completes the proof of (1) and the claim. \blacktriangleleft

Clearly, the set I depends only on (\mathbf{A}, S) and thus it can be computed by an FPT-algorithm. Thus it suffices to show how to compute $|\mathcal{N}'|$ in the remainder of the proof.

For each set $T \subseteq S$ we define $\mathcal{N}_T := \{h \in \text{hom}(\mathbf{A}, \mathbf{D}, S) \mid (\pi_1 \circ h)(S) \subseteq T\}$. We have by inclusion-exclusion

$$|\mathcal{N}'| = \sum_{T \subseteq S} (-1)^{|S \setminus T|} |\mathcal{N}_T|. \quad (3)$$

Observe that there are only $2^{|S|}$ summands in (3) and thus if we can reduce all of them to #CQ with the query (\mathbf{A}, S) this will give us the desired counting slice reduction.

We will now show how to compute the $|\mathcal{N}_T|$ by interpolation. So fix a $T \subseteq S$. Let $\mathcal{N}_{T,i}$ for $i = 0, \dots, |S|$ consist of the mappings $h \in \text{hom}(\mathbf{A}, \mathbf{D}, S)$ such that there are exactly i elements $a \in S$ that are mapped to $h(a) = (a', b)$ such that $a' \in T$. Obviously, $\mathcal{N}_T = \mathcal{N}_{T,|S|}$ with this notation.

Now for each $j = 1, \dots, |S|$ we construct a new structure $\mathbf{D}_{j,T}$ over the domain $D_{j,T}$. To this end, for each $a \in T$, let $a^{(1)}, \dots, a^{(j)}$ be copies of a which are not in D . Then we set

$$D_{j,T} := \{(a^{(k)}, b) \mid (a, b) \in D, a \in T, k \in [j]\} \cup \{(a, b) \mid (a, b) \in D, a \notin T\}.$$

We define a mapping $B : D \rightarrow \wp(D_{j,T})$, where $\wp(D_{j,T})$ is the power set of $D_{j,T}$, by

$$B(a, b) := \begin{cases} \{(a^{(k)}, b) \mid k \in [j]\}, & \text{if } a \in T \\ \{(a, b)\}, & \text{otherwise.} \end{cases}$$

For every relation symbol $R \in \tau$ we define $R^{\mathbf{D}_{j,T}} := \bigcup_{(d_1, \dots, d_s) \in R^D} B(d_1) \times \dots \times B(d_s)$.

Then every $h \in \mathcal{N}_{T,i}$ corresponds to i^j mappings in $\text{hom}(\mathbf{A}, \mathbf{D}_{j,T}, S)$. Thus for each j we get $\sum_{i=1}^{|S|} i^j |\mathcal{N}_{T,i}| = |\text{hom}(\mathbf{A}, \mathbf{D}_{j,T}, S)|$. This is a linear system of equations and the corresponding matrix is a Vandermonde matrix, so $\mathcal{N}_T = \mathcal{N}_{T,|S|}$ can be computed with an oracle for #CQ on the instances $((\mathbf{A}, S), \mathbf{D}_{j,T})$. The size of the linear system depends only on $|S|$. Furthermore, $\|D^j\| \leq \|D\|j^s \leq \|D\|^{s+1}$ where s is the bound on the arity of the relations symbols in τ and thus a constant. It follows that the algorithm described above is a counting slice reduction. This completes the proof of Lemma 30. \blacktriangleleft

7.2 Reducing from hypergraphs to structures

In this section we show that we can in certain situations reduce from $\#CQ_{hyp}$ to $\#CQ$. This will later allow us to reduce from the hard cases in Theorem 11 to show the hardness results of Theorem 22.

We proceed in several steps. Let in this section \mathcal{C} be a class of conjunctive queries of bounded arity. To every query (\mathbf{A}, S) we construct a structure $\hat{\mathbf{A}}$ as follows; note that when we use this notation, S will be clear from the context. Construct the augmented structure $\text{aug}(\mathbf{A}, S)$ of \mathbf{A} and compute its core. We define $\hat{\mathbf{A}}$ to be the structure that we get by deleting the relations R_a for $a \in S$ that we added in the construction of $\text{aug}(\mathbf{A}, S)$. We set $\hat{\mathcal{C}} := \{(\hat{\mathbf{A}}, S) \mid (\mathbf{A}, S) \in \mathcal{C}\}$.

Note that in any situation where we apply both the $\hat{\cdot}$ - and \cdot^* -operators, the $\hat{\cdot}$ is applied before the \cdot^* .

► **Claim 4.** *There is a parsimonious slice-reduction from $\#CQ[\hat{\mathcal{C}}]$ to $\#CQ[\mathcal{C}]$.*

Proof. The relation U relates to every query (\mathbf{A}, S) the query $(\hat{\mathbf{A}}, S)$. Certainly, U is computable and by definition assigns to each query in $\hat{\mathcal{C}}$ a query in \mathcal{C} .

We have that $\hat{\mathbf{A}}$ is a substructure of \mathbf{A} and there is a homomorphism from \mathbf{A} to $\hat{\mathbf{A}}$, because there is a homomorphism from $\text{aug}(\mathbf{A}, S)$ to $\text{aug}(\hat{\mathbf{A}}, S)$. Hence, \mathbf{A} and $\hat{\mathbf{A}}$ are homomorphically equivalent and by Theorem 14 we have that (\mathbf{A}, S) and $(\hat{\mathbf{A}}, S)$ are equivalent. Thus setting $r((\mathbf{A}, S), (\hat{\mathbf{A}}, S), \mathbf{B}) := \mathbf{B}$ yields the desired parsimonious slice-reduction. ◀

Let $\hat{\mathcal{C}}^* := \{(\hat{\mathbf{A}}^*, S) \mid (\hat{\mathbf{A}}, S) \in \hat{\mathcal{C}}\}$. Note that, by Lemma 30, there is a counting slice reduction from $p\text{-}\#CQ[\hat{\mathcal{C}}^*]$ to $p\text{-}\#CQ[\hat{\mathcal{C}}]$.

Let now \mathcal{G} be the class of S -hypergraphs associated to the queries in $\hat{\mathcal{C}}$.

► **Claim 5.** *There is a parsimonious slice reduction from $\#CQ_{hyp}[\mathcal{G}]$ to $\#CQ[\hat{\mathcal{C}}^*]$.*

Proof. The relation U relates every S -hypergraph (\mathcal{H}, S) in \mathcal{G} to all queries $(\hat{\mathbf{A}}^*, S)$ with the hypergraph (\mathcal{H}, S) . Certainly, U is computable and by definition of \mathcal{G} it assigns to S -hypergraph in \mathcal{G} a query in $\hat{\mathcal{C}}^*$.

It remains to describe the function r . So let $((\mathbf{A}, S), \mathbf{B})$ be a $p\text{-}\#CQ$ -instance such that (\mathbf{A}, S) has the S -hypergraph $(\mathcal{H}, S) \in \mathcal{G}$. We assume w.l.o.g. that every tuple appears only in one relation of \mathbf{A} . If this is not the case, say a tuple t appears in two relations $R_1^{\mathbf{A}}$ and $R_2^{\mathbf{A}}$, then we build a new instance as follows: Delete t from $R_1^{\mathbf{A}}$ and $R_2^{\mathbf{A}}$, add a new relation $R_t^{\mathbf{A}}$ to \mathbf{A} containing only t . Finally, set $R^{\mathbf{B}} = R_1^{\mathbf{B}} \cap R_2^{\mathbf{B}}$. This operation does not change the associated S -hypergraph, so this new instance still has the S -hypergraph (\mathcal{H}, S) . Moreover it is easy to see that it has the same set of solutions.

Let the vocabulary of $(\hat{\mathbf{A}}, S)$ be τ . We construct a structure $r((\mathbf{A}, S), (\hat{\mathbf{A}}^*, S), \mathbf{B}) := \hat{\mathbf{B}}$ over the same relation symbols as $\hat{\mathbf{A}}^*$, i.e., over the vocabulary $\tau \cup \{R_a \mid a \in \hat{\mathbf{A}}\}$. The structure $\hat{\mathbf{B}}$ has the domain $\hat{B} := A \times B$ where A is the domain of \mathbf{A} and B is the domain of \mathbf{B} . For $\hat{R} \in \tau$ we set

$$\hat{R}^{\hat{\mathbf{B}}} := \{((a_1, b_1), \dots, (a_k, b_k)) \mid (a_1, \dots, a_k) \in \hat{R}^{\hat{\mathbf{A}}}, (a_1, \dots, a_k) \in R^{\mathbf{A}}, (b_1, \dots, b_k) \in R^{\mathbf{B}}\}.$$

Furthermore, for the relations symbols \hat{R}_a that are added in the construction of $\hat{\mathbf{A}}^*$ from $\hat{\mathbf{A}}$ we set $\hat{R}_a^{\hat{\mathbf{B}}} := \{(a, b) \mid b \in B\}$, where B is the domain of \mathbf{B} .

It is easy to see that from a satisfying assignment $h : \mathbf{A} \rightarrow \mathbf{B}$ we get a homomorphism $h' : \hat{\mathbf{A}}^* \rightarrow \hat{\mathbf{B}}$ by setting $h'(a) := (a, h(a))$. Furthermore, this construction is obviously

bijjective. Thus we get $|\text{hom}(\mathbf{A}, \mathbf{B}, S)| = |\text{hom}(\hat{\mathbf{A}}^*, \hat{\mathbf{B}}, S)|$. Since $\hat{\mathbf{B}}$ can be constructed in polynomial time in $\|\mathbf{A}\|$ and $\|\mathbf{B}\|$, this is a parsimonious slice reduction. \blacktriangleleft

► **Corollary 31.** *Let \mathcal{C} be a class of conjunctive queries of bounded arity and let \mathcal{G} be the class of S -hypergraphs of the cores of \mathcal{C} . Then there is a counting slice reduction from $\#\text{CQ}_{\text{hyp}}[\mathcal{G}]$ to $\#\text{CQ}[\mathcal{C}]$.*

7.3 Strict star size

In this section we introduce a notion strict S -star size to simplify some of the arguments in the next section. We define the strict S -star size of a hypergraph to be the maximum number of vertices in S that are contained in one S -component of \mathcal{H} .

► **Lemma 32.** *Let \mathcal{G} be a class of S -hypergraphs of bounded arity. If the strict S -star size of the S -hypergraphs in \mathcal{G} is unbounded, then there is a counting slice reduction from $\text{CLIQUE}[\mathbb{N}]$ to $\#\text{CQ}_{\text{hyp}}[\mathcal{G}]$.*

7.4 The main hardness results

In this section we use the results of the last sections to prove the hardness results of Theorem 22.

The proof of Theorem 11 in [8] directly yields the following result.

► **Lemma 33.** *Let \mathcal{G} be a class of S -hypergraphs of bounded arity. If the treewidth of \mathcal{G} is unbounded, then there is a counting slice reduction from $\text{CLIQUE}[\mathbb{N}]$ to $\#\text{CQ}_{\text{hyp}}[\mathcal{G}]$.*

Combining Lemma 33 and Corollary 31 yields the following corollary.

► **Corollary 34.** *Let \mathcal{C} be a class of queries such that the treewidth of the cores of the queries in \mathcal{C} is unbounded. Then there is a counting slice reduction from $\text{CLIQUE}[\mathbb{N}]$ to $\#\text{CQ}[\mathcal{C}]$.*

► **Lemma 35.** *Let \mathcal{G} be a class of hypergraphs such that $\text{contract}(\mathcal{G})$ is of unbounded treewidth. Then there is a counting slice reduction from $\#\text{CLIQUE}[\mathbb{N}]$ to $\#\text{CQ}_{\text{hyp}}[\mathcal{G}]$.*

Proof. Assume first that \mathcal{G} is of unbounded strict S -star size. Then $\#\text{CQ}_{\text{hyp}}[\mathcal{G}]$ is $\#\text{W}[1]$ -hard by Lemma 32. So we assume in the remainder of the proof that there is a constant c such that for every (\mathcal{H}, S) in \mathcal{G} every S -component of (\mathcal{H}, S) contains only c vertices from S .

From Theorem 10 it follows that there is a counting slice reduction from $\#\text{CLIQUE}[\mathbb{N}]$ to $\#\text{CQ}_{\text{hyp}}[\text{contract}(\mathcal{G})]$. Therefore, it suffices to show parsimonious slice reduction (U, r) from $\#\text{CQ}_{\text{hyp}}[\text{contract}(\mathcal{G})]$ to $\#\text{CQ}_{\text{hyp}}[\mathcal{G}]$ to show the lemma.

The relation U is defined as $U := \{(\text{contract}(\mathcal{H}, S), (\mathcal{H}, S)) \mid (\mathcal{H}, S) \in \mathcal{G}\}$. By definition, this satisfies the covering condition.

For the definition of r , consider an instance $((\mathbf{A}, S), \mathbf{B})$ of $\#\text{CQ}_{\text{hyp}}[\text{contract}(\mathcal{G})]$ and let \mathcal{H} be the hypergraph of \mathbf{A} . Moreover, let $(\mathcal{H}', S) \in \mathcal{G}$ be an S -hypergraph such that $(\mathcal{H}, S) = \text{contract}(\mathcal{H}', S)$. W.l.o.g. assume that for every edge e of \mathcal{H} , the structure \mathbf{A} contains one relation R_e containing only a single tuple \vec{e} where \vec{e} contains the elements of e in an arbitrary order. We construct an instance $r((\mathcal{H}, S), (\mathcal{H}', S), ((\mathbf{A}, S), \mathbf{B})) = ((\mathbf{A}', S), \mathbf{B}')$. Similarly to \mathbf{A} , the structure \mathbf{A}' has for every edge e in \mathcal{H}' a relation R_e that contains only a single tuple \vec{e} with the properties as before. For every S -component C of \mathcal{H} we do the following: Let D_C be the tuples encoding the homomorphisms h from $\mathbf{A}[V(C) \cap S]$ to \mathbf{B} . For every $v \in V(C) \setminus S$ we let D_C be the domain of v . Whenever two elements $u, v \in V(C) \setminus S$ appear in an edge, we set $R_e^{\mathbf{B}'}$ in such a way that for all tuples in $R_e^{\mathbf{B}'}$ the assignments to u

and v coincide. Moreover, whenever $u \in V(C) \cup S$ and $v \in V(C) \setminus S$ we allow only tuples in which the assignment to u coincides with the assignment to u that is encoded in the assignment to v . For all edges e of \mathcal{H}' with $e \setminus S \neq \emptyset$, we let $R_e^{\mathbf{B}'}$ contain all tuples that satisfy the two conditions above. Finally, for all edges e with $e \in S$ we set $R^{\mathbf{A}'} := R^{\mathbf{A}}$.

It is easy to verify that $\text{hom}(\mathbf{A}, \mathbf{B}, S) = \text{hom}(\mathbf{A}', \mathbf{B}', S)$. Thus it only remains to show that the construction can be done in polynomial time. Note first that the number of variables from S in any S -component of \mathcal{H}' is bounded by c . Thus we can compute all domains D_C in time $\|\mathbf{B}\|^{O(c)}$. The rest of the construction can then be easily done in polynomial time. ◀

► **Corollary 36.** *Let \mathcal{C} be a class of conjunctive queries such that $\text{contract}(\mathcal{C})$ is of unbounded treewidth. Then there is a counting slice reduction from $\#\text{CLIQUE}[\mathbb{N}]$ to $\#\text{CQ}[\mathcal{C}]$.*

Proof. This follows by combination of Lemma 35 and Corollary 31. ◀

7.5 Putting things together

We now finally show Theorem 22 by putting together the results of the last sections.

Proof of Theorem 22.

1. follows directly from Corollary 26 with the observation that bounded treewidth of $\text{contract}(\mathcal{C})$ implies bounded S -star size.
2. is Corollary 34 and Lemma 27 using the fact that $\text{CQ}[\text{ALL}]$ counting slice reduces to $\text{CLIQUE}[\mathbb{N}]$; this follows from [9, Section 6.1].
3. is Corollary 36. ◀

8 Conclusion

In this paper we have proved a complete classification for the counting complexity of conjunctive queries, continuing a line of work that spans several previous papers [6, 8, 11]. While this solves the bounded arity case completely, the most apparent open question is what happens for the unbounded arity case. This case is rather well understood for the decision version CQ of the problem [14], but for counting not much is known. In particular, it is not known if the results of [14] can even be adapted to the quantifier free setting.

Another interesting problem would be to go from conjunctive queries to more expressive query languages. This has been done with some success for decision problems (see e.g. [3] and the references therein), but the situation for counting is much less clear. It is known that counting and decision differ a lot at least in some settings, e.g. even very simple unions of conjunctive queries yield hard counting problems [17, 15] while CQ for these queries is very easy. Can we get a better understanding of counting complexity in this setting and how it differs from decision?

To prove our results, we have extended the case complexity framework to counting complexity. We are very optimistic that this will be helpful when studying the research areas discussed above. Moreover, due to its generic nature, we feel that this framework should also be of use outside of the query answering context and allow transparent proofs and presentations in other areas of parameterized complexity.

Acknowledgements. Chen was supported by the Spanish project TIN2013-46181-C2-2-R, by the Basque project GIU12/26, and by the Basque grant UFI11/45.

References

- 1 S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- 2 A.K. Chandra and P.M. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *STOC 1977*, pages 77–90. ACM, 1977.
- 3 Hubie Chen. The tractability frontier of graph-like first-order query sets. *CoRR*, abs/1407.3429v1, 2014. Conference version appeared in the proceedings of LICS '14.
- 4 Hubie Chen and Martin Grohe. Constraint satisfaction with succinctly specified relations. *J. Comput. Syst. Sci.*, 76(8):847–860, 2010.
- 5 Hubie Chen and Moritz Müller. One hierarchy spawns another: Graph deconstructions and the complexity classification of conjunctive queries. In *LICS*, 2014.
- 6 V. Dalmau and P. Jonsson. The complexity of counting homomorphisms seen from the other side. *Theor. Comput. Sci.*, 329(1-3):315–323, 2004.
- 7 V. Dalmau, P.G. Kolaitis, and M.Y. Vardi. Constraint Satisfaction, Bounded Treewidth, and Finite-Variable Logics. In *International Conference on Principles and Practice of Constraint Programming 2002*, pages 310–326, 2002.
- 8 A. Durand and S. Mengel. Structural tractability of counting of solutions to conjunctive queries. *Theory of Computing Systems*, pages 1–48, 2014. accepted, to appear, final version available at <http://dx.doi.org/10.1007/s00224-014-9543-y>.
- 9 J. Flum and M. Grohe. Parameterized Complexity Theory. *Texts in Theoretical Computer Science. An EATCS Series*, 2006.
- 10 Georg Gottlob, Nicola Leone, and Francesco Scarcello. Hypertree decompositions and tractable queries. *J. Comput. Syst. Sci.*, 64(3):579–627, 2002.
- 11 Gianluigi Greco and Francesco Scarcello. Counting solutions to conjunctive queries: structural and hybrid tractability. In *Proceedings of the 33rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS'14*, pages 132–143, 2014.
- 12 M. Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM*, 54(1), 2007.
- 13 P. Kolaitis and M. Vardi. Conjunctive-Query Containment and Constraint Satisfaction. *Journal of Computer and System Sciences*, 61:302–332, 2000.
- 14 Dániel Marx. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. *J. ACM*, 60(6):42, 2013.
- 15 S. Mengel. *Conjunctive Queries, Arithmetic Circuits and Counting Complexity*. PhD thesis, University of Paderborn, 2013.
- 16 C. Papadimitriou and M. Yannakakis. On the Complexity of Database Queries. *Journal of Computer and System Sciences*, 58(3):407–427, 1999.
- 17 R. Pichler and S. Skritek. Tractable counting of the answers to conjunctive queries. *Journal of Computer and System Sciences*, 2013.
- 18 Nicole Schweikardt, Thomas Schwentick, and Luc Segoufin. Database theory: Query languages. In Mikhail J. Atallah and Marina Blanton, editors, *Algorithms and Theory of Computation Handbook*, volume 2: Special Topics and Techniques, chapter 19. CRC Press, second edition, Nov 2009.