

Structural Tractability of Counting of Solutions to Conjunctive Queries

Arnaud Durand · Stefan Mengel

Published online: 23 May 2014
© Springer Science+Business Media New York 2014

Abstract We explore the complexity of counting solutions to conjunctive queries, a basic class of queries from database theory. We introduce a parameter, called the *quantified star size* of a query ϕ , which measures how the free variables are spread in ϕ . As usual in database theory, we associate a hypergraph to a query ϕ . We show that for classes of queries for which these associated hypergraphs admit good decompositions, e.g., bounded width generalized hypertree decompositions, bounded quantified star size exactly characterizes the subclasses of hypergraphs for which counting the number of solutions is tractable. In the case of bounded arity, this allows us to fully characterize the classes of hypergraphs for which counting the solutions is tractable. Finally, we also analyze the complexity of computing the quantified star size of a conjunctive query.

Keywords Conjunctive queries · Counting complexity · Hypergraph decomposition techniques

Partially supported by ANR-11-IS02-0003, project ALCOCLAN

Partially supported by DFG grants BU 1371/2-2 and BU 1371/3-1. Furthermore, the research leading to these results has received funding from the [European Community's] Seventh Framework Programme [FP7/2007-2013] under grant agreement no 238381

A. Durand (✉)

Université Paris Diderot, Institut de Mathématiques de Jussieu, UMR 7586 CNRS,
75205 Paris, France
e-mail: durand@math.univ-paris-diderot.fr

S. Mengel

Department of Mathematics, Technische Universität Berlin, 10623 Berlin, Germany

1 Introduction

Conjunctive queries are a fundamental class of queries from database theory. Equivalent to Select-Project-Join queries, they are the most basic class of database queries and at the same time play an important role in practice. Furthermore, as Kolaitis and Vardi [26] showed, conjunctive queries are intimately connected to constraint satisfaction problems, a central area from artificial intelligence. These features make conjunctive queries the best-studied type of database queries.

A CQ-instance (\mathcal{A}, ϕ) consists of a *query* ϕ , which is a logical first-order $\{\exists, \wedge\}$ -formula, also called primitive positive formula, and a finite structure (i.e., database) \mathcal{A} . The *query result* is

$$\phi(\mathcal{A}) := \{\mathbf{a} \mid (\mathcal{A}, \mathbf{a}) \models \phi(\mathbf{x})\},$$

that is, the set of assignments that make the query ϕ true.

The focus of most research on conjunctive queries has been the Boolean conjunctive query problem (short CQ) which is, given an instance, to decide if the query result of the instance is empty or not. This problem is well known to be **NP**-complete and thus the main interest of study has been to identify tractable subclasses, so-called “islands of tractability”, where CQ is tractable, i.e., can be solved in polynomial time.

One main direction in finding tractable classes of CQ-instances has been imposing structural restrictions on the queries—or often more exactly on the hypergraph associated to it—while the database is assumed to be arbitrary. In a seminal paper Yannakakis [33] proved that ACQ, the restriction of CQ to acyclic queries, is tractable. The main idea behind other structural restrictions is to extend this result by generalizing it to “nearly acyclic” queries. This has led to many different decompositions for graphs and hypergraphs and associated width measures (see e.g., [10, 18, 30]). The common approach for these decompositions is to group together vertices or edges (of the graphs or hypergraphs) into clusters of some fixed constant size and to arrange these clusters into a tree satisfying certain conditions. The resulting width measures are often sought to have two desirable properties:

- For every k the class of queries of width k should be tractable, i.e., CQ should be solvable in polynomial time.
- Given an instance, it should be possible to decide if there is a decomposition of width k and construct one if it exists.

While decomposition techniques without the first property do not make any sense in the context of conjunctive queries, the second property is sometimes relaxed. For some decomposition techniques one does not actually need the decomposition to solve CQ [8], a promise of the existence is enough. For other decompositions one only knows approximation algorithms that construct decompositions of a width that is near the optimal width, which is enough to guarantee tractability of CQ [2, 28].

More recently there has also been interest in enumerating all solutions to conjunctive queries and in the corresponding counting question #CQ which is,

given a CQ-instance, to determine the size of the query result. For enumeration of the query answers it turns out that the picture is less clear than for decision [3, 6, 21]. Also the situation for counting is more subtle: For quantifier free queries—which correspond to queries without projections in the database perspective—most commonly considered structural restrictions yield tractable counting problems (see, e.g., [31]). While this is nice, it is not fully satisfying, because quantifiers—which correspond to projections in database theory—are very natural and essential in database queries. While introducing projections does not make any difference for the complexity of CQ, the situation for #CQ is dramatically different. In [31] it is shown that even one single existentially quantified variable is enough to make counting answers to CQ-instances #P-hard, even when the associated hypergraph of the query is a tree (which implies width 1 for all commonly considered decomposition techniques). It follows that the decomposition techniques used for CQ are not enough to guarantee tractability for counting.

In a previous paper [15] the authors of this paper have proposed a way out of this dilemma for counting by introducing a parameter called *quantified star size* for acyclic conjunctive queries. This parameter measures how the free variables are spread in the query. We associated to a query $\phi(\mathbf{x})$ the list \mathbf{x} of free variables, thus extending the hypergraph $\mathcal{H} = (V, E)$ associated to $\phi(\mathbf{x})$ with a set $S \subseteq V$. Then the quantified star size is the size of a maximum independent set consisting of vertices from the set S in some specified subhypergraphs of \mathcal{H} . It turns out that, under a widely believed assumption from parameterized complexity, this measure precisely characterizes the tractable subclasses of acyclic conjunctive queries.

1.1 Overview of the Results

1.1.1 Counting Solutions to Queries

In this paper we extend the results of [15] from acyclic queries to classes of queries defined by commonly considered decomposition techniques. To do so we generalize the notion of quantified star size from acyclic queries to general conjunctive queries. We show that every class of conjunctive queries that allows efficient counting must be of bounded quantified star size—again under the same assumption from parameterized complexity. We then go on showing that for all decomposition techniques for conjunctive queries commonly considered in the literature, combining them with bounded quantified star size leads to tractable counting problems. The key feature that makes this result work is the organization of atoms and variables into a tree of clusters that is prominent in all decomposition methods for CQ known so far. Combining the results we get an exact characterization of the subclasses of CQ-instances that allow tractable counting for commonly considered classes defined by decomposition techniques. Let us illustrate these results for the example of generalized hypertree decomposition [18], which is one of the most general decomposition methods and one of the most studied too [18, 20, 30]. We have that, under the assumption that $\mathbf{FPT} \neq \#\mathbf{W}[1]$, for any (recursively enumerable)

class \mathcal{C} of hypergraphs of bounded generalized hypertreewidth the following statements are equivalent:

- #CQ for instances in \mathcal{C} can be solved in polynomial time
- \mathcal{C} is of bounded quantified star size.

In our considerations, the arity of atoms of queries is not a priori bounded. In this setting, there is no known *ultimate* measure resulting from a decomposition method that fully characterizes tractability even for the decision problem CQ. This explains why our characterizations are stated for *each* decomposition method.

For bounded arity however, the situation is different. It is known that being of bounded treewidth completely characterizes tractability for decision [22, 24] and counting [12] for quantifier free instances. Combining [22, 24] and our results from above we derive a complete characterization of tractability for #CQ in terms of tree width and quantified star size of the underlying hypergraph for the bounded arity case.

Note that our results are for counting with set semantics, i.e., we count each solution only once. Counting for bag semantics in which multiple occurrences of identical tuples are counted has already been essentially solved in [31].

1.1.2 Discovering Quantified Star Size

To exploit tractability results of the above kind it is helpful if the membership in a tractable class can be decided efficiently, i.e., in our case if computing the quantified star size of an instance is tractable. Therefore, we consider this “discovery problem” of determining the quantified star size of queries in the second part of the paper.

In [15] it is shown that quantified star size of acyclic conjunctive queries can be determined in polynomial time.

We show that computing the quantified star size of an instance is equivalent to computing maximum independent sets in hypergraphs. Consequently, we cannot expect a polynomial time algorithm for computing the quantified star size of general CQ-instances. Fortunately, it turns out that for queries ϕ of generalized hypertree width k and thus for all more restrictive decomposition techniques like hingetree width (see [10]) or treewidth, there is an algorithm that computes in time $|\phi|^{O(k)}$ the quantified star size of ϕ . We show that this is in a sense optimal, because computing the quantified star size of a given query ϕ is $\mathbf{W}[1]$ -hard parameterized by the generalized hypertree width of ϕ . Thus, under the standard assumption $\mathbf{FPT} \neq \mathbf{W}[1]$, there is no fixed-parameter algorithm for this problem.

Still, if we parameterize the computation of quantified star size by more restrictive width measures, computing the quantified star size of conjunctive queries in some cases becomes fixed-parameter tractable. We prove that this is the case queries if we parameterize by hingetree width. Because of the connection between quantified star size and maximum independent set, this result provides a new parameter of hypergraphs for which computing maximum independent sets is fixed-parameter tractable. Note that the $\mathbf{W}[1]$ -hardness result from above shows that fixed-parameter

tractability of computing maximum independent sets is unlikely to hold for other hypergraph decomposition techniques.

We then turn our attention to the approximation of quantified star size. We show that there is a polynomial time algorithm that, given a query ϕ and a decomposition of ϕ of width k , computes in time independent of k a k -approximation of the quantified star size of ϕ .

Summing these results up, quantified star size does not only imply tractable counting if combined with well known decomposition techniques, but in case the decomposition is given or can be efficiently computed (treewidth, hingetree width) or approximated (generalized hypertreewidth), then computing quantified star size is itself tractable.

We show that this is in a sense optimal, because under the assumption $\mathbf{FPT} \neq \mathbf{W}[1]$ there is no efficient (fixed parameter tractable in k) algorithm computing the quantified star size for queries parameterized by generalized hypertree width.

Finally, we investigate the problem of counting solution and computing quantified star size for queries of bounded fractional hypertree width [23, 28]. This decomposition method is of a somewhat different nature than the ones studied before so we treat it individually. We again prove that counting is tractable in this setting and that the discovery problem can be decided in $O(n^{k^{O(1)}})$, i.e., with a slightly bigger dependency in k than before.

2 Preliminaries

In this section, we introduce the basic definitions and the notation we will use throughout the paper. We start off with a formal definition of conjunctive queries, introduce some notions from parameterized complexity and then survey the graph and hypergraph decompositions we will consider.

2.1 Conjunctive Queries

We give a brief introduction to conjunctive queries. More on the subject and, in general, on database theory and finite model theory can be found in [1, 27]

A *relational vocabulary* is a set of relation symbols $\tau := \{\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_\ell\}$ where each \mathcal{R}_i has an arity r_i which we denote by $\text{arity}(\mathcal{R}_i)$. A finite (*relational*) *structure* \mathcal{A} over τ , a τ -structure for short, is a tuple $(A, \mathcal{R}_1^{\mathcal{A}}, \dots, \mathcal{R}_\ell^{\mathcal{A}})$ where A is a finite set called the *domain* of \mathcal{A} and $\mathcal{R}_i^{\mathcal{A}} \subseteq A^{r_i}$ is a relation of arity r_i called the interpretation of \mathcal{R}_i .

We denote structures by calligraphic letters, e.g., $\mathcal{A}, \mathcal{B}, \dots$. For the corresponding domains we use the corresponding roman letters, i.e., A is the domain of \mathcal{A} , B the domain of \mathcal{B} and so on.

Let $\mathcal{R} \in \tau$ with $\text{arity}(\mathcal{R}) = r$ and \bar{z} be a sequence z_1, \dots, z_r of (not necessarily distinct) variables. Then, the expression $\mathcal{R}(z_1, \dots, z_r)$ (short $\mathcal{R}(\bar{z})$) is called an *atomic formula* or *atom*. The *scope* $\text{var}(\mathcal{R}(\bar{z}))$ of this atom is defined as the set of variables appearing in (z_1, \dots, z_r) .

A *quantifier free conjunctive query* ϕ over a vocabulary τ is a logical formula of the form

$$\phi = \mathcal{R}_{i_1}(\bar{z}_1) \wedge \dots \wedge \mathcal{R}_{i_s}(\bar{z}_s),$$

where $\mathcal{R}_{i_j}(\bar{z}_j)$ are atomic formulas with $\mathcal{R}_{i_j} \in \tau$. We denote the set of variables of ϕ by $\text{var}(\phi) := \bigcup_{j \in [s]} \text{var}(\mathcal{R}_{i_j}(\bar{z}_j))$. The set of all atoms of ϕ is denoted by $\text{atom}(\phi)$.

A *conjunctive query* ϕ over τ is a formula $\phi = \exists x_1 \dots \exists x_t \phi'$ where ϕ' is a quantifier free conjunctive query over τ and $x_j \in \text{var}(\phi)$ for all $j \leq t$. The x_j are called *quantified variables*. The set of variables of ϕ is defined as $\text{var}(\phi) := \text{var}(\phi')$. The set of *free variables* of ϕ is defined as $\text{free}(\phi) := \text{var}(\phi) \setminus \{x_1, \dots, x_t\}$. We often write $\phi(\bar{y})$ where $\{\bar{y}\} = \text{free}(\phi)$, to stress the free variables.

A *conjunctive query instance* over τ , short CQ-instance, is a pair $\Phi = (\mathcal{A}, \phi)$ where \mathcal{A} is a finite τ -structure and ϕ is a conjunctive query over τ . Φ is called quantifier free if the query ϕ is quantifier free.

Let $\Phi = (\mathcal{A}, \phi)$ be a quantifier free CQ-instance. An *assignment* to Φ is a mapping $a : \text{var}(\phi) \rightarrow A$. A partial assignment to Φ is a mapping $a : X \rightarrow A$ for a subset X of $\text{var}(\phi)$. Let $a : X \rightarrow A$ and $b : Y \rightarrow A$ be two partial assignments. We call a and b *compatible*, in symbols $a \sim b$, if they agree on their common variables, i.e., for all $x \in X \cap Y$ we have $a(x) = b(x)$. Let $\mathcal{R}(z_1, \dots, z_r)$ be an atom of ϕ . We say that a *satisfies* $\mathcal{R}(z_1, \dots, z_r)$ if $(a(z_1), \dots, a(z_r)) \in \mathcal{R}^{\mathcal{A}}$. We say that a *satisfies* Φ if it satisfies all of its atoms. In this case we write $(\mathcal{A}, a) \models \phi$.

An assignment to a general, not necessarily quantifier free, CQ-instance $\Phi = (\mathcal{A}, \phi)$ is a mapping $a : \text{free}(\phi) \rightarrow A$. An assignment a can alternatively be seen as a tuple of dimension $|\text{free}(\phi)|$ indexed by the variables $\text{free}(\phi)$. Consequently, relations will also be seen either as sets of tuples or as list of assignments. An assignment $a : \text{free}(\phi) \rightarrow A$ satisfies Φ if there is an assignment $a' : \text{var}(\phi) \rightarrow A$ with $a \sim a'$ such that the quantifier free query instance (\mathcal{A}, ϕ') , where we get ϕ' by deleting all quantifiers from ϕ , is satisfied by a' . Again we write $(\mathcal{A}, a) \models \phi$. Observe that a' is in general not unique.

The *query result* $\phi(\mathcal{A})$ of a CQ-instance $\Phi = (\mathcal{A}, \phi)$ is defined as

$$\phi(\mathcal{A}) := \{a \mid (\mathcal{A}, a) \models \phi(\bar{x})\}.$$

The elements of the query result are called *solutions* of the query instance or *satisfying assignments* or *query answers*. We call two instances $\Phi = (\mathcal{A}, \phi)$, $\Phi' = (\mathcal{A}', \phi')$ *solution equivalent*, if $\text{free}(\phi) = \text{free}(\phi')$ and $\phi(\mathcal{A}) = \phi'(\mathcal{A}')$.

Let $a : X \rightarrow A$ be an assignment and $Y \subseteq X$. By $a|_Y$ we denote the restriction of a onto Y . Similarly, if \mathcal{R} is a relation indexed by X , i.e., such that each $a \in \mathcal{R}$ is interpreted as an assignment $a : X \rightarrow A$ then, $\pi_Y(\mathcal{R}) := \{a|_Y \mid a \in \mathcal{R}\}$ denotes the *projection* of \mathcal{R} onto Y . Throughout the paper we will make use of the following classical database operations on relations.

Definition 1 Let \mathcal{R}_1 and \mathcal{R}_2 be two relations indexed by the variables X and Y , respectively (X and Y being not necessarily disjoint). The *natural join* of \mathcal{R}_1 and \mathcal{R}_2 is

$$\mathcal{R}_1 \bowtie \mathcal{R}_2 := \{a : X \cup Y \rightarrow A \mid a|_X \in \mathcal{R}_1, a|_Y \in \mathcal{R}_2\}.$$

The *semi-join* of \mathcal{R}_1 and \mathcal{R}_2 is defined as $\mathcal{R}_1 \bowtie \mathcal{R}_2 := \pi_X(\mathcal{R}_1 \bowtie \mathcal{R}_2)$.

2.2 Model of Computation and Encoding of Instances

The underlying model of computation for our algorithms will be the RAM model with addition as basic operation and with unit costs measure. We assume the relations of a finite structure \mathcal{A} to be encoded by listing their tuples. Apart from this convention we will not specify an encoding but only give estimates on its size in O -notation that will be satisfied by any reasonable encoding.

Let \mathcal{A} be a τ -structure. For a relation $\mathcal{R}^{\mathcal{A}}$ let $|\mathcal{R}^{\mathcal{A}}|$ denote the cardinality of $\mathcal{R}^{\mathcal{A}}$. Then we define for $\mathcal{R}^{\mathcal{A}}$

$$\|\mathcal{R}^{\mathcal{A}}\| := \text{arity}(\mathcal{R}) \cdot |\mathcal{R}^{\mathcal{A}}|.$$

For the vocabulary τ let $|\tau|$ be the number of predicate symbols. Finally, let $|A|$ be the cardinality of a domain A . Then for a structure \mathcal{A} over the vocabulary τ with domain A we define

$$\|\mathcal{A}\| := |\tau| + |A| + \sum_{\mathcal{R} \in \tau} \|\mathcal{R}^{\mathcal{A}}\|.$$

Furthermore, we define for a conjunctive query

$$|\phi| := \sum_{\substack{\phi' \in \text{atom}(\phi), \\ \mathcal{R} \text{ relation symbol of } \phi'}} \text{arity}(\mathcal{R}).$$

Finally, for a CQ-instance $\Phi = (\mathcal{A}, \phi)$ we define

$$\|\Phi\| := |\phi| + \|\mathcal{A}\|.$$

Note that for any reasonable encoding, up to constant factors, $\|\mathcal{A}\|$ is the size of an encoding of \mathcal{A} , $|\phi|$ is the size of an encoding of ϕ and $\|\Phi\|$ is the size of an encoding of Φ . For a detailed discussion and justification of these conventions see [16, Section 2.3].

The following lemma states that the basic database operations we considered above can be performed efficiently.

Lemma 1 ([16]) *Given relations \mathcal{R}_1 and \mathcal{R}_2 and $Y \subseteq \text{var}(\mathcal{R})$, one can compute*

- $\mathcal{R}_1 \bowtie \mathcal{R}_2$ in time $O(\|\mathcal{R}_1\| + \|\mathcal{R}_2\| + \|\mathcal{R}_1 \bowtie \mathcal{R}_2\|)$,
- $\pi_Y(\mathcal{R}_1)$ in time $O(\|\mathcal{R}_1\|)$,
- $\mathcal{R}_1 \bowtie \mathcal{R}_2$ in time $O(\|\mathcal{R}_1\| + \|\mathcal{R}_2\|)$.

We will use Lemma 1 throughout the paper, most of the time without explicitly referencing it.

2.3 Query Problems

The basic computational question on CQ-instances is the *Conjunctive query answering problem*.

Conjunctive query answering problem
Input: a conjunctive query ϕ and a structure \mathcal{A} .
Problem: Compute $\phi(\mathcal{A})$.

Clearly, $\|\phi(\mathcal{A})\|$ can be exponential in $\|\Phi\|$ and thus we cannot have a polynomial time algorithm. The *Boolean conjunctive query problem* is defined as follows.

CQ
Input: a conjunctive query ϕ and a structure \mathcal{A} .
Problem: Decide if $\phi(\mathcal{A})$ is empty or not.

The main focus in this paper will be on the associated counting problem #CQ.

#CQ
Input: a conjunctive query ϕ and a structure \mathcal{A} .
Problem: Compute $|\phi(\mathcal{A})|$.

2.4 Parameterized Complexity

This section is a very short introduction to some notions from parameterized complexity used in the remainder of this paper. For more details see [17].

A *parameterized decision problem* over an alphabet Σ is a language $L \subseteq \Sigma^*$ together with a computable parameterization $\kappa : \Sigma^* \rightarrow \mathbb{N}$. The problem (L, κ) is said to be *fixed-parameter tractable*, or $(L, \kappa) \in \mathbf{FPT}$, if there is a computable function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that there is an algorithm that decides for $x \in \Sigma^*$ in time $f(\kappa(x))|x|^{O(1)}$ if x is in L .

Let (L, κ) and (L', κ') be two parameterized decision problems over the alphabets Σ , resp. Π . A *parameterized many-one reduction* from (L, κ) to (L', κ') is a function $r : \Sigma^* \rightarrow \Pi^*$ such that for all $x \in \Sigma^*$:

- $x \in L \Leftrightarrow r(x) \in L'$,
- $r(x)$ can be computed in time $f(\kappa(x))|x|^c$ for a computable function f and a constant c , and
- $\kappa'(r(x)) \leq g(\kappa(x))$ for a computable function g .

It is easy to see that **FPT** is closed under parameterized many-one reductions.

Let p -Clique be the following parameterized problem.

p -Clique
Input: a graph G and $k \in \mathbb{N}$.
Parameter: k .
Problem: Decide if G has a clique of size k .

Here the parameterization κ is simply defined by $\kappa(G, k) := k$. The class **W[1]** consists of all parameterized problems that are parameterized many-one reducible to p -Clique. A problem (L, κ) is called **W[1]-hard**, if there is a parameterized many-one reduction from p -Clique to (L, κ) .

It is widely believed that $\mathbf{FPT} \neq \mathbf{W}[1]$ and thus in particular p -Clique and all $\mathbf{W}[1]$ -hard problems are not fixed-parameter tractable.

Parameterized counting complexity theory is developed similarly to decision complexity. A *parameterized counting problem* is a function $F : \Sigma^* \times \mathbb{N} \rightarrow \mathbb{N}$, for an alphabet Σ . Let $(x, k) \in \Sigma^* \times \mathbb{N}$, then we call x the input of F and k the parameter. A parameterized counting problem F is *fixed-parameter tractable*, or $F \in \mathbf{FPT}$, if there is an algorithm computing $F(x, k)$ in time $f(k) \cdot |x|^c$ for a computable function $f : \mathbb{N} \rightarrow \mathbb{N}$ and a constant $c \in \mathbb{N}$.

Let $F : \Sigma^* \times \mathbb{N} \rightarrow \mathbb{N}$ and $G : \Pi^* \times \mathbb{N} \rightarrow \mathbb{N}$ be two parameterized counting problems. A *parameterized parsimonious reduction* from F to G is an algorithm that computes, for every instance (x, k) of F , an instance (y, l) of G in time $f(k) \cdot |x|^c$ such that $l \leq g(k)$ and $F(x, k) = G(y, l)$ for computable functions $f, g : \mathbb{N} \rightarrow \mathbb{N}$ and a constant $c \in \mathbb{N}$. A *parameterized T -reduction* from F to G is an algorithm with an oracle for G that solves any instance (x, k) of F in time $f(k) \cdot |x|^c$ in such a way that for all oracle queries the instances (y, l) satisfy $l \leq g(k)$ for computable functions f, g and a constant $c \in \mathbb{N}$.

Let p -#Clique be the following problem.

p -#Clique
Input: a graph G and $k \in \mathbb{N}$.
Parameter: k .
Problem: Compute the number of cliques of size k in G .

A parameterized problem F is in $\#\mathbf{W}[1]$ if there is a parameterized T -reduction from F to p -#Clique. F is $\#\mathbf{W}[1]$ -hard, if there is a parameterized T -reduction from p -#Clique to F . As usual, F is $\#\mathbf{W}[1]$ -complete if it is in $\#\mathbf{W}[1]$ and hard for it, too.

Again, it is widely believed that there are problems in $\#\mathbf{W}[1]$ (in particular the $\#\mathbf{W}[1]$ -complete problems) that are not fixed-parameter tractable. Thus, from showing that a problem F is $\#\mathbf{W}[1]$ -hard it follows that F can be assumed to be not fixed-parameter tractable.

We will mainly deal with two parameterized problems that are versions of CQ and #CQ parameterized by the size of the input query. This parameterization is justified by the origins from database theory. In a typical database application the query is usually far smaller than the database, so it makes sense to use the size of the query as a parameter.

p -CQ
Input: a conjunctive query ϕ and a structure \mathcal{A} .
Parameter: $|\phi|$
Problem: Decide if $\phi(\mathcal{A})$ is empty or not.

p -#CQ
Input: a conjunctive query ϕ and a structure \mathcal{A} .
Parameter: $|\phi|$
Problem: Compute $|\phi(\mathcal{A})|$.

2.5 Graphs and Hypergraphs Associated to Queries

As remarked before, CQ and #CQ are hard computational problems. One way to isolate islands of tractability, is to analyze structural aspects of the query. A key idea for this is to associated graphs and hypergraphs to queries.

A (finite) hypergraph \mathcal{H} is a pair (V, E) where V is a finite set and $E \subseteq \mathcal{P}(V)$ where $\mathcal{P}(V)$ is the power set of V . The *arity* of \mathcal{H} is $\max_{e \in E} |e|$. We associate a hypergraph $\mathcal{H} = \mathcal{H}_\phi = (V, E)$ to a query ϕ by setting $V := \text{var}(\phi)$ and $E := \{\text{var}(\phi_t) \mid \phi_t \in \text{atom}(\phi)\}$.

Example 1 Consider the query

$$\begin{aligned} \phi := & \exists u_1 \exists u_2 \exists u_3 \exists u_4 \exists u_5 \exists u_6 \exists u_7 \exists u_8 \\ & P_1(v_1, u_1) \wedge P_2(v_2, u_1, u_2) \wedge P_3(v_2, v_4, u_2, u_3) \\ & \wedge P_4(v_3, v_4, v_5, u_3, u_4, u_5) \wedge P_5(v_4, v_5, v_6, v_8) \\ & \wedge P_6(v_7, v_8, u_5, u_6) \wedge P_2(v_6, v_9, u_7) \wedge P_2(v_8, v_9, u_8) \end{aligned}$$

The associated hypergraph is illustrated in Fig. 1.

One also associates graphs to queries as follows: Let ϕ be a query with the associated hypergraph $\mathcal{H} = (V, E)$, then one associates to ϕ the so-called *primal graph* $\mathcal{H}_P = (V, E_p)$ where for every $u, v \in V$ we have $uv \in E_p$ if and only if there is an edge $e \in E$ with $u, v \in e$.


2.6 Graph and Hypergraph Decompositions

We now introduce several decomposition techniques for graphs and hypergraphs that will be used in the rest of this paper to analyze the complexity of counting solutions to queries.

2.6.1 Treewidth

We first present some basic facts on treewidth. All proofs can be found in the survey by Bodlaender [5] and the references therein.

Fig. 1 The hypergraph associated to the query ϕ of Example 1



Unless stated otherwise all graphs are nonempty, finite, undirected and simple, i.e., they have no parallel edges or loops. In contrast, trees are always assumed to be rooted and thus directed.

The treewidth of a graph G is a measure of how similar G is to a tree. There are several equivalent definitions for treewidth of which we will first present the one by Robertson and Seymour [32].

Definition 2 A *tree decomposition* of a graph $G = (V, E)$ is a pair $(\mathcal{T}, (\chi_t)_{t \in T})$ where $\mathcal{T} = (T, F)$ is a rooted tree and $\chi_t \subseteq V$ for every $t \in T$ satisfying the following properties:

1. For every $v \in V$ there is a $t \in T$ with $v \in \chi_t$.
2. For every $e \in E$ there is a $t \in T$ such that $e \subseteq \chi_t$.
3. For every $v \in V$ the set $\{t \in T \mid v \in \chi_t\}$ induces a subtree of \mathcal{T} .

The third property is called the *connectedness condition*. The sets χ_t are called *blocks* or *bags* of the decomposition.

We call $\max_{t \in T} (|\chi_t|) - 1$ the *width* of the tree composition $(\mathcal{T}, (\chi_t)_{t \in T})$. The *treewidth* $\text{tw}(G)$ of G is the minimum width over all tree decompositions of G .

To ease notation we sometimes identify a vertex $t \in T$ with the corresponding bag χ_t .

We remark that the class of graphs of treewidth 1 consists exactly of all forests, i.e., the graphs that have trees as their connected components. In particular, trees have treewidth 1.

Example 2 Given the graph G from Fig. 2, a tree decomposition of G is given in Fig. 3.

Given a graph G and an integer k , it is **NP**-complete to decide if G has treewidth G at most k , but if we take k as a parameter the problem becomes fixed-parameter tractable.




Fig. 2 The graph G of Example 2




Fig. 3 A tree decomposition of width 2 of the graph G of Fig. 2

Theorem 1 ([4]) *There is a polynomial p and an algorithm that, given a graph $G = (V, E)$, computes a tree decomposition of G of width $k := \text{tw}(G)$ in time at most $2^{p(k)}|V|$.*

We will use the following folklore results.

Lemma 2 ([13], Chap. 12) *Let $G = (V, E)$ be a graph, $C \subseteq V$ a clique in G and $(T, (\chi_t)_{t \in T})$ a tree decomposition of G . Then there is a bag χ_t such that $C \subseteq \chi_t$.*

Lemma 3 ([17], Chap. 11) *Every graph G of treewidth at most k has a vertex of degree at most k .*

We will also use an alternative definition of treewidth by so-called elimination orders.

Definition 3 Let $G = (V, E)$ be a graph with $|V| = n$. A bijection $\pi : V \rightarrow [n]$ is called an *elimination order*. We say that u is higher-numbered than v with respect to π if $\pi(u) > \pi(v)$. The *fill-in graph* G_π of G with respect to π is constructed iteratively: Starting from G , for $i = 1, \dots, |V|$ we add an edge between all pairs u, w of neighbors of $\pi^{-1}(i)$ that are higher-numbered than $\pi^{-1}(i)$.

The *width* of π is the minimum integer k such that in G_π each vertex $v \in V$ has at most k higher-numbered neighbors.

The *elimination width* $\text{elim-width}(G)$ of G is the minimum width over all elimination orders of G .

Example 3 We consider again the graph G of Fig. 2. An elimination order π is defined by the sequence $A, B, I, H, G, C, D, E, F$. The fill-in graph G_π is shown in Fig. 4. The width of π is 2.

Elimination orders give the following characterization of treewidth which appears to be folklore. A proof can be found e.g., in [5].

Lemma 4 *For every graph G we have $\text{elim-width}(G) = \text{tw}(G)$.*




Fig. 4 The fill-in graph G_π of Example 3 defined by the sequence $A, B, I, H, G, C, D, E, F$. The edges added during the construction of G_π from G are represented as *dotted lines*

2.6.2 Hypergraph Decomposition Techniques

In this section we present some well known hypergraph decomposition techniques. For more details on hypergraph decompositions see e.g., [10, 18, 30]. For all decomposition techniques defined below, the width of a CQ-instance $\Phi = (\mathcal{A}, \phi)$ is simply defined as the width of the hypergraph associated to ϕ .

The simplest idea to generalize treewidth to hypergraphs is considering primal graphs and to define the *treewidth* of a hypergraph \mathcal{H} to be treewidth of its primal graph \mathcal{H}_P . By Lemma 2, classes of hypergraphs that have unbounded edge size are of unbounded treewidth even when the hypergraphs are intuitively very simple. Consequently, treewidth is, for some considerations on hypergraphs, not the right measure of the complexity of a hypergraph. Thus research focussed on finding decompositions that work with hypergraphs directly and not with their primal graphs. The base class of hypergraphs that roughly corresponds to trees in the setting of treewidth are *acyclic hypergraphs* which are defined with the help of *join trees* which organize the edges of a hypergraph in a tree with a connectivity condition similar to that for treewidth.

Definition 4 A *join tree* of a hypergraph $\mathcal{H} = (V, E)$ is a pair $(\mathcal{T}, (\lambda_t)_{t \in T})$ where $\mathcal{T} = (T, F)$ is a rooted tree and for each $t \in T$ we have $\lambda_t \in E$ such that

- for each $e \in E$ there is a $t \in T$ such that $\lambda_t = e$,
- For each $v \in V$ the set $\{t \in T \mid v \in \lambda_t\}$ induces a subtree of T .

A hypergraph is called *acyclic* if it has a join tree.

When there is no ambiguity, we often identify vertices $t \in T$ with their edges λ_t .

Lemma 5 ([33]) *There is a polynomial time algorithm that, given a hypergraph \mathcal{H} , decides if \mathcal{H} is acyclic. Moreover, if \mathcal{H} is acyclic the algorithm computes a join tree of \mathcal{H} .*

A conjunctive query ϕ is called acyclic if its associated hypergraph is acyclic.

Acyclic conjunctive queries play an important role in database theory, because of the following result by Yannakakis [33].

Theorem 2 ([33]) *ACQ can be solved in polynomial time.*

Theorem 2 served as a starting point to finding more general classes of hypergraphs on which CQ is tractable, by trying to identify classes of “nearly” acyclic hypergraphs. There are lots of different decomposition techniques and associated width measures for hypergraphs. One of the most general width measures is *generalized hypertree width*.

The approach of generalized hypertree decomposition is similarly to that of tree decompositions: We want to organize a hypergraph into clusters that form a tree with a connectivity condition. Instead of bags that contain vertices and that must cover all edges, the basic clusters of generalized hypertree decompositions are *guarded blocks* (λ_t, χ_t) where λ_t contains edges while χ_t contains vertices. To make sure that the vertices χ_t of a guarded block form a sufficiently simple set, we demand that χ_t is covered by the edges in λ_t and that λ_t is small. To make sure that the decomposition represents the hypergraph well, we require that every edge must be contained in a set χ_t (but not necessarily in a λ_t). We now give the exact definition.

Definition 5 *A generalized hypertree decomposition of a hypergraph $\mathcal{H} = (V, E)$ is a triple $(\mathcal{T}, (\lambda_t)_{t \in T}, (\chi_t)_{t \in T})$ where $\mathcal{T} = (T, F)$ is a rooted tree and $\lambda_t \subseteq E$ and $\chi_t \subseteq V$ for every $t \in T$ satisfying the following properties:*

- For every $e \in E$ there is a $t \in T$ such that $e \subseteq \chi_t$.
- For every $t \in T$ we have $\chi_t \subseteq \bigcup_{e \in \lambda_t} e$.
- For every $v \in V$ the set $\{t \in T \mid v \in \chi_t\}$ induces a subtree of \mathcal{T} .

The third property is again called the *connectedness condition*. The sets χ_t are called *blocks* or *bags* of the decomposition, while the sets λ_t are called the *guards* of the decomposition. A pair (λ_t, χ_t) is called *guarded block*.

The *width* of a decomposition $(\mathcal{T}, (\lambda_t)_{t \in T}, (\chi_t)_{t \in T})$ is $\max_{t \in T} (|\lambda_t|)$. The *generalized hypertree width* of \mathcal{H} is the minimum width over all generalized hypertree decompositions of \mathcal{H} .

Again, we sometimes identify a guarded block (λ_t, χ_t) with the vertex t .

Example 4 Figure 5 shows a generalized hypertree decomposition of width 3 for the hypergraph from Fig. 1.

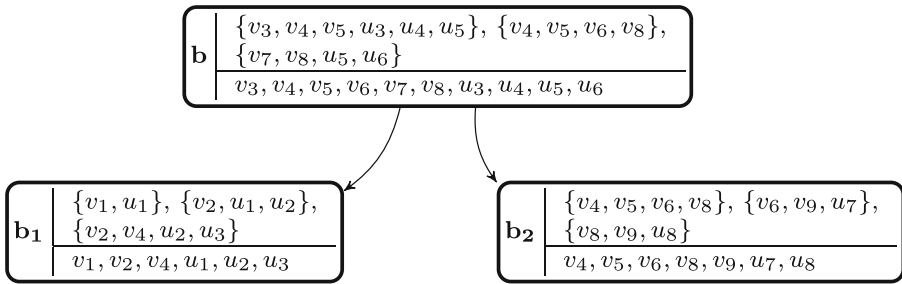


Fig. 5 A generalized hypertree decomposition of width 3 for the hypergraph from Fig. 1. The boxes are the guarded blocks. In the upper parts the guards are given while the lower parts show the blocks. On the left part, a name is given to each node

We give the following very easy upper bound for generalized hypertree width.

Observation 3 Let $\mathcal{H} = (V, E)$ be a hypergraph such that there are k edges e_1, \dots, e_k in E with $V \subseteq \bigcup_{i=1}^k e_i$. Then \mathcal{H} has generalized hypertree width at most k .

Proof We will construct a trivial width k generalized hypertree decomposition $(\mathcal{T}, (\lambda_t)_{t \in T}, (\chi_t)_{t \in T})$ of \mathcal{H} . The tree \mathcal{T} only consists of one single vertex t , the block of t is $\chi_t := V$ and the guard is $\lambda_t := \{e_1, \dots, e_k\}$. It is easily seen that this satisfies all desired properties of a hypertree decomposition. Furthermore, the decomposition has width k . □

It turns out the generalized hypertree width is strictly more general than treewidth in the following sense.

Lemma 6 ([18]) For every hypergraph \mathcal{H} the generalized hypertree width is less than or equal to $1 + \text{tw}(\mathcal{H})$. Moreover, for every ℓ there are hypergraphs of treewidth ℓ and generalized hypertree width 1.

Unfortunately, deciding if a hypergraph has generalized hypertree width at most k is NP-complete even for $k = 3$ [20]. This unpleasant result is amended by the fact that there is an approximation algorithm.

Theorem 4 ([2, 19]) There is an algorithm that, given a hypergraph \mathcal{H} of generalized hypertree width k , constructs a generalized hypertree decomposition of width $O(k)$ of \mathcal{H} in time $|\mathcal{H}|^{O(k)}$.

A hypergraph is acyclic if and only if it has generalized hypertree width 1. With this result the following lemma is easy to prove.

Lemma 7 Let $(\mathcal{T}, (\lambda_t)_{t \in T}, (\chi_t)_{t \in T})$ be a generalized hypertree decomposition of a hypergraph \mathcal{H} . Let $\mathcal{H}' = (V, E')$ where $E' := \{\chi_t \mid t \in T\}$. Then \mathcal{H}' is acyclic and $(\mathcal{T}, (\chi_t)_{t \in T})$ is a join tree of \mathcal{H}' .

Next we state a lemma that in different forms is (implicitly) used in most papers that deal with the application of hypergraph decomposition techniques to CQ, see e.g., [19]. The main observation here is that we make explicit that the construction is solution equivalent.

Lemma 8 *Given a CQ-instance Φ with associated hypergraph \mathcal{H} and a generalized hypertree decomposition of \mathcal{H} of width k , one can compute an ACQ-instance Ψ in time $\|\Phi\|^{O(k)}$ such that Φ and Ψ are solution equivalent.*

Proof Let $\Phi = (\mathcal{A}, \phi)$ and let the given generalized hypertree decomposition be $(\mathcal{T}', (\lambda_t)_{t \in T}, (\chi_t)_{t \in T})$. We construct $\Psi = (\mathcal{B}, \psi)$ with $\text{var}(\psi) = \text{var}(\phi)$ as follows: For every $t \in T$ the query ψ has an atom ψ_t with relation symbol \mathcal{R}_t and $\text{var}(\psi_t) := \chi_t$. The quantifiers are the same as for ϕ . For every $t \in T$ let ϕ_1, \dots, ϕ_s be the atoms associated to the edges in λ_t . We have $s \leq k$. Let $\phi'_1, \dots, \phi'_\ell$ be the atoms associated to the edges e with $e \subseteq \chi_t$. Then we define the relation $\mathcal{R}_t^{\mathcal{B}}$ as

$$\mathcal{R}_t^{\mathcal{B}} := \pi_{\chi_t}(\phi_1(\mathcal{A}) \bowtie \dots \bowtie \phi_s(\mathcal{A})) \bowtie \phi'_1(\mathcal{A}) \bowtie \dots \bowtie \phi'_\ell(\mathcal{A}).$$

We claim that ϕ and ψ are solution equivalent. First consider an assignment a that satisfies all atoms of ϕ . Then we have for every subset ϕ'_1, \dots, ϕ'_r of $\text{atom}(\phi)$ that the assignment a is compatible to an assignment in $\phi'_1 \bowtie \dots \bowtie \phi'_r$. It follows that for each t the new atom ψ_t is satisfied by a . Consequently, $\phi(\mathcal{A}) \subseteq \psi(\mathcal{B})$.

Let now a be an assignment that satisfies all atoms of ψ . Then a must for each t satisfy the atoms ϕ'_i from the construction of $\mathcal{R}_t^{\mathcal{B}}$. But since each edge e of \mathcal{H} is covered by a set $\chi_{t'}$, every atom in $\text{atom}(\phi)$ contributes as a ϕ'_i in the construction of a $\phi_{t'}$. Consequently, a satisfies all atoms of ϕ and thus $\psi(\mathcal{B}) \subseteq \phi(\mathcal{A})$.

We claim that this construction can be done in time $\|\Phi\|^{O(k)}$. To see this, observe that the relation $A_{\lambda_t} := \pi_{\chi_t}(\phi_1(\mathcal{A}) \bowtie \dots \bowtie \phi_s(\mathcal{A}))$ has size at most $\|\mathcal{A}\|^s \leq \|\mathcal{A}\|^k$. Since for the ϕ'_i we have $\text{var}(\phi'_i) \subseteq \chi_t$, it follows that $\mathcal{R}_t^{\mathcal{B}} \subseteq A_{\lambda_t}$ and thus consequently $\|\mathcal{R}_t^{\mathcal{B}}\| \leq \|\mathcal{A}\|^k$. With Lemma 1 it follows that we can compute $\mathcal{R}_t^{\mathcal{B}}$ in time $|\phi| \|\mathcal{A}\|^{O(k)}$. Thus computing the instance Ψ takes time $\|\Phi\|^{O(k)}$.

Finally, by Lemma 7 we have that Ψ is acyclic. □

The combination of Theorem 2 and Lemma 8 allows to solve CQ-instances in time $\|\Phi\|^{O(k)}$ provided that a generalized hypertree decomposition of width k is given. Thus the bottleneck for solving CQ-instances for many proposed decomposition techniques is the efficient computation of a good decomposition of the instance.

Let us fix some notation: For an edge set $\lambda \subseteq E$ we use the shorthand $\bigcup \lambda := \bigcup_{e \in \lambda} e$. For a decomposition $(\mathcal{T}, (\lambda_t)_{t \in T}, (\chi_t)_{t \in T})$ we write \mathcal{T}_t for the subtree of \mathcal{T} that has t as its root and denote by $V(\mathcal{T}_t) \subseteq T$ its vertex set. We also write $\chi(\mathcal{T}_t) := \bigcup_{t' \in V(\mathcal{T}_t)} \chi_{t'}$. We use these notations for tree decompositions as well.

It is sometimes helpful to consider restrictions of generalized hypertree decompositions, because those might have better structural or algorithmic properties. A

number of them have been defined and studied in the past, among others biconnected component, cycle-cutset, cycle-hypercutset, hingetree, hypertree decomposition (see [30] for a survey). Hingetree decomposition will play a role in this paper and we define it formally below.

Definition 6 A generalized hypertree decomposition is called *hingetree decomposition* if it satisfies the following conditions:

- For each pair $t_1, t_2 \in T$ with $t_1 \neq t_2$ there are edges $e_1 \in \lambda_{t_1}$ and $e_2 \in \lambda_{t_2}$ such that $\chi_{t_1} \cap \chi_{t_2} \subseteq e_1 \cap e_2$.
- For each $t \in T$ we have $\bigcup \lambda_t = \chi_t$.
- For each $e \in E$ there is a $t \in T$ such that $e \in \lambda_t$.

The *hingetree width* (also called *degree of cyclicity*) of \mathcal{H} is the minimum width over all hingetree decompositions of \mathcal{H} .

Note that this is not the original definition from [25] but an alternative, equivalent definition from [10].

Example 5 The decomposition from Fig. 5 is also a hingetree decomposition since:

- $\chi_{b_1} \cap \chi_b = \{v_4, u_3\} \subseteq \{v_2, v_4, u_2, u_3\} \cap \{v_3, v_4, v_5, u_3, u_4, u_5\} \in \lambda_{b_1} \cap \lambda_b$
- $\chi_{b_2} \cap \chi_b = \{v_4, v_5, v_6, v_8\} \in \lambda_{b_2} \cap \lambda_b$

Like treewidth, hingetree width is strictly less general than generalized hypertree width in the following sense.

Lemma 9 ([18]) *For every hypergraph the generalized hypertree width is less than or equal to the hingetree width. Moreover, there are hypergraphs for which the generalized hypertree width is strictly less than the hingetree width.*

Hingetree width makes up for this lack of generality by the fact that optimal decompositions can be computed very efficiently.

Lemma 10 ([25]) *There is an algorithm that, given a hypergraph $\mathcal{H} = (V, E)$, computes a minimum width hingetree decomposition of \mathcal{H} in time $|V| \cdot |E|^2$.*

Since the presented decomposition techniques are all easy to compute, we get the following lemma.

Lemma 11 (see e.g., [10]) *For all of the width measures defined above CQ restricted to instances Φ of width k can be solved in time $\|\Phi\|^{p(k)}$ for a polynomial p .*

3 Quantified Star Size

As proved in [31], even introducing one single existential quantifier in acyclic conjunctive queries leads to #P-complete counting problems. It follows that bounding the number of quantified variables does not yield tractable instances. In [15] we have shown a corresponding #W[1]-hardness result for p -#ACQ: We presented a class of hard instances for parameterized complexity which have a very simple form.

Lemma 12 ([15]) *Let $\phi_{star,n} := \exists z \bigwedge_{i \in [n]} \mathcal{R}_i(z, y_i)$ and let $\mathcal{C}_{star} := \{\phi_{star,n} \mid n \in \mathbb{N}\}$. Then p -#CQ is #W[1]-hard for instances restricted to queries in \mathcal{C}_{star} .*

For an alternative proof of Lemma 12 see [29].

A basic observation on the hard instances of Lemma 12 is that their associated hypergraphs are stars whose center is the single quantified variable. Abstracting this observation, we shall define the parameter called *quantified star size* which, when bounded and combined with known decomposition techniques, leads to tractable #CQ-instances. This parameter has been introduced in [15] for acyclic queries and we generalize it here to the general setting. As we will see, not the number of quantified variables in a query is crucial but how they are distributed in the associated hypergraph.

Before we introduce quantified star size, we make several other definitions.

Let $\mathcal{H} = (V, E)$ be a hypergraph and $V' \subseteq V$. A *subhypergraph* $\mathcal{H}' = (V', E')$ of \mathcal{H} is a hypergraph with $E' \subseteq \{e \cap V' \mid e \in E, e \cap V' \neq \emptyset\}$. The *induced subhypergraph* $\mathcal{H}[V']$ of \mathcal{H} is the hypergraph $\mathcal{H}[V'] = (V', E')$ where, this time, $E' = \{e \cap V' \mid e \in E, e \cap V' \neq \emptyset\}$. Let $x, y \in V$, a *path* between x and y is a sequence of vertices $x = v_1, \dots, v_k = y$ such that for each $i \in [k - 1]$ there is an edge $e_i \in E$ with $v_i, v_{i+1} \in e_i$.

A (connected) *component* of \mathcal{H} is an induced subhypergraph $\mathcal{H}[V']$ where V' is a maximal vertex set such that for each pair $x, y \in V'$ there is a path between x and y in \mathcal{H} . These definitions apply to graphs as well.

We will use the following observation on induced subgraphs of acyclic hypergraphs.

Observation 5 *Let β be any decomposition technique defined in Section 2.6. Let $\mathcal{H} = (V, E)$ be a hypergraph of β -width k . Then for every $V' \subseteq V$ the induced subhypergraph $\mathcal{H}[V']$ has β -width at most k .*

Proof Let $(\mathcal{T}, (\lambda_t)_{t \in T}, (\chi_t)_{t \in T})$ be a β -decomposition of \mathcal{H} of width k . For each guarded block (λ_t, χ_t) compute a guarded block (λ'_t, χ'_t) with $\chi_t := \chi_t \cap V'$ and $\lambda'_t := \{e \cap V' \mid e \in \lambda\}$. It is easy to check that $(\mathcal{T}, (\lambda'_t)_{t \in T}, (\chi'_t)_{t \in T})$ is a β -decomposition of width at most k . □

Since acyclic hypergraphs are the hypergraphs of generalized hypertree width 1, we get the following special case.

Observation 6 If $\mathcal{H} = (V, E)$ is an acyclic hypergraph and $V' \subseteq V$, then $\mathcal{H}[V']$ is acyclic. More specifically, let $(\mathcal{T}, (\lambda)_{t \in \mathcal{T}})$ with $\mathcal{T} = (T, F)$ be a join tree of \mathcal{H} . Then $(\mathcal{T}[T'], (\lambda_t \cap V')_{t \in T'})$ where $T' := \{t \in T \mid \lambda_t \cap V' \neq \emptyset\}$ can be made into a join tree of $\mathcal{H}[V']$ by connecting the components of $\mathcal{T}[T']$ arbitrarily.

Observe that there is no version of Observation 5 or Observation 6 for sub-hypergraphs instead of induced subhypergraphs. To see this consider an arbitrary hypergraph $\mathcal{H} = (V, E)$. Adding the edge V to E yields an acyclic hypergraph, independent of the generalized hypertree width of \mathcal{H} .

Definition 7 An S -hypergraph is a pair (\mathcal{H}, S) where $\mathcal{H} = (V, E)$ is a hypergraph and $S \subseteq V$. If \mathcal{H} is a graph, we also call (\mathcal{H}, S) an S -graph.

The S -hypergraph associated to a CQ-instance $\Phi = (\mathcal{A}, \phi)$ consists of the hypergraph associated to ϕ and $S := \text{free}(\phi)$. The *primal S -graph* of \mathcal{H} is defined as (\mathcal{H}_P, S) .

Definition 8 Let \mathcal{G} be a class of S -hypergraphs. By $\#CQ$ on \mathcal{G} we denote the restriction of $\#CQ$ to instances whose associated S -hypergraph is in \mathcal{G} . Analogously, by p - $\#CQ$ on \mathcal{G} we denote the restriction of p - $\#CQ$ to instances whose associated S -hypergraph is in \mathcal{G} .

Definition 9 We call an S -hypergraph S -connected if for every pair of vertices x, y there is a path $x = v_1, v_2, \dots, v_{k-1}, v_k = y$ such that $v_i \notin S$ for $i \notin \{1, k\}$.

Let us consider some examples of queries that have S -connected S -hypergraphs.

Example 6 Path queries (of arbitrary length), for example

$$\phi(x, y) := \exists t_1 \exists t_2 \exists t_3 \mathcal{R}(x, t_1) \wedge \mathcal{R}(t_1, t_2) \wedge \mathcal{R}(t_2, t_3) \wedge \mathcal{R}(t_3, y)$$

have as their associated S -hypergraph a path in which only the end vertices are in S . Thus their S -hypergraph is S -connected.

Example 7 The associated graph of the query $\phi_{\text{star},n}$ of Lemma 12 is the star $G_n = (V_n, E_n)$ where $V_n = \{z, y_1, \dots, y_n\}$ and $E_n = \{zy_1, \dots, zy_n\}$. Furthermore, the free variables are $S_n = \{y_1, \dots, y_n\}$. Every vertex in V_n is connected to every other vertex via $z \notin S_n$. Thus (G_n, S_n) is S_n -connected.

Definition 10 An *independent set* I in a hypergraph $\mathcal{H} = (V, E)$ is a set $I \subseteq V$ such that there are no distinct vertices $x, y \in I$ that lie in a common edge $e \in E$.

Note that independent sets defined this way are in the hypergraphs literature often called *strong independent sets*. There is also the notion of *weak independent set* where we only require that no edge lies completely in the set I . Since we will only talk about strong independent sets and there is thus no danger of confusing the two

notions, we will simply talk about *independent sets* always implying the strongness implicitly.

Definition 11 The *S*-star size of an *S*-connected *S*-hypergraph is the maximum size of an independent set consisting of vertices in *S* only. We say that such an independent set *forms* an *S*-star.

We remark that the *S*-star size of an *S*-connected *S*-hypergraph can equivalently be expressed as the the size of a maximum independent set in $\mathcal{H}[S]$.

Example 8 The *S*-hypergraphs associated to the path queries of Example 6 have *S*-star size 2, because the two end vertices of the paths are independent.

Now consider the S_n -hypergraph (G_n, S_n) from Example 7. The vertices of S_n are all independent. Consequently, the S_n -star size of (G_n, S_n) is n .

Note that while the quantified star size of instances whose associated hypergraph is a path is bounded by 2, the *S*-star size of bounded pathwidth instances is unbounded. To see this, observe that the S_n -hypergraph (G_n, S_n) from above has pathwidth 1: The decomposition $(\mathcal{P}, (\chi_t)_{t \in T})$ where \mathcal{P} is the path with vertex set $[n]$ and $\chi_i := \{z, y_i\}$, is a path decomposition of G_n of width 1.

We want to extend the notion of *S*-star size to *S*-hypergraphs that are not necessarily *S*-connected. To this end, we consider certain *S*-connected subhypergraphs that we call *S*-components. We make the following crucial definition.

Definition 12 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} .

Example 9 Let us consider the *S*-hypergraph of the query from Example 1. The associated hypergraph \mathcal{H} was already illustrated in Fig. 1. We have $S = \{v_1, \dots, v_9\}$. Then $\mathcal{H}[V \setminus S]$ has three components with the vertex sets $C_1 := \{u_7\}$, $C_2 := \{u_8\}$ and $C_3 := \{u_1, \dots, u_6\}$. Thus

- $E_{C_1} = \{\{u_7, v_6, v_9\}\}$,
- $E_{C_2} = \{\{u_8, v_8, v_9\}\}$, and
- $E_{C_3} = \{\{u_1, v_1\}, \{u_1, u_2, v_2\}, \{u_2, u_3, v_2, v_4\},$
 $\{u_3, u_4, u_5, v_3, v_4, v_5\}, \{u_5, u_6, v_7, v_8\}\}$.

Hence, the vertex sets of the components are: $V'_{C_1} = \{u_7, v_6, v_9\}$, $V'_{C_2} = \{u_8, v_8, v_9\}$ and $V'_{C_3} = \{u_1, u_2, u_3, u_4, u_5, u_6, v_1, v_2, v_3, v_4, v_5, v_7, v_8\}$. The three resulting *S*-components are depicted in Fig. 6.

The following observations are evident from the definition.

Observation 7 To an *S*-hypergraph (\mathcal{H}, S) one can in polynomial time compute all its *S*-components.




Fig. 6 The S -components of the S -hypergraph discussed in Example 9

Observation 8 *The only S -component of an S -connected S -hypergraph (\mathcal{H}, S) is \mathcal{H} . Moreover, the S -components of S -hypergraphs are S -connected.*

Observation 8 allows to extend the definition of S -star size to not necessarily S -connected hypergraphs.

Definition 13 For an S -hypergraph (\mathcal{H}, S) we define S -star size as the maximum S -star size of its S -components.

Example 10 Let us compute the S -star size of the S -hypergraph of Example 9. The S -components induced by V_{C_1}' and V_{C_2}' are completely covered by the edges $\{u_7, v_6, v_9\}$, resp., $\{u_8, v_8, v_9\}$. It follows that the S -star size of these S -components is 1. We have $V_{C_3} \cap S = \{v_1, v_2, v_3, v_4, v_5, v_7, v_8\}$. There are several maximum independent sets of vertices in $V_{C_3} \cap S$ in the S -component induced by V_{C_3} , all of size 4. An example is $\{v_1, v_2, v_3, v_7\}$. It follows that the S -star size of (\mathcal{H}, S) is 4

Now we can finally come back to CQ-instances and define the promised parameter quantified star size.

Definition 14 The *quantified star size* of a conjunctive query is defined as the S -star size of the associated S -hypergraph. The quantified star size of a CQ-instance is that of its query.

Example 11 The query from Example 1 has quantified starsize 4 as we have seen in Example 10. From Example 8 we get that the queries $\phi_{\text{star},n}$ of Lemma 12 are of quantified star size n , which is nearly the size of the query.

4 Quantified Star Size is Sufficient and Necessary for Efficient Counting

In this section we analyze the effect of quantified star size on the complexity of $\#\text{CQ}$.

4.1 An Algorithm for Instances of Bounded Quantified Star Size

In this section we show that the decomposition techniques introduced in Section 2.6 lead to efficient counting when combined with bounded quantified star size. We proceed with the following rather technical lemma.

Lemma 13 *There is an algorithm that, given a CQ-instance $\Phi = (\mathcal{A}, \phi)$ of quantified starsize ℓ and a generalized hypertree decomposition $\Xi = (\mathcal{T}, (\lambda_t)_{t \in \mathcal{T}}, (\chi_t)_{t \in \mathcal{T}})$ of Φ of width k , constructs a CQ-instance $\Psi = (\mathcal{B}, \psi)$ in time $\|\Phi\|^{p(k, \ell)}$ for a fixed polynomial p such that*

- Φ and Ψ are solution equivalent,
- Ψ is acyclic, and
- ψ is quantifier free.

In the proof we will use the following lemma from [15]. An edge cover of a hypergraph $\mathcal{H} = (V, E)$ is a set E^* of edges of \mathcal{H} such that $V \subseteq \bigcup_{e \in E^*} e$.

Lemma 14 ([15]) *For acyclic hypergraphs the size of a maximum independent set and a minimum edge cover coincide. Moreover, there is a polynomial time algorithm that given an acyclic hypergraph \mathcal{H} computes a maximum independent set I and a minimum edge cover E^* of \mathcal{H} .*

Proof of Lemma 13 Given $\Phi = (\mathcal{A}, \phi)$, we construct Ψ in several steps.

Let $\mathcal{H} = (V, E)$ be the hypergraph of ϕ . Let V_1, \dots, V_m be the vertex sets of the connected components of $\mathcal{H}[V \setminus S]$ and let V'_1, \dots, V'_m be the vertex sets of the corresponding S -components of \mathcal{H} . Clearly, we have $V_i \subseteq V'_i$ and $V'_i \setminus V_i = V'_i \cap S =: S_i$. Let Φ_i be the CQ-instance whose query ϕ_i is obtained by restricting all atoms of ϕ to the variables in V'_i and whose structure \mathcal{A}_i is obtained by projecting all relations of \mathcal{A} accordingly. The associated hypergraph of ϕ_i is $\mathcal{H}[V'_i]$. Moreover, $\mathcal{H}[V'_i]$ has a generalized hypertree decomposition Ξ_i of width at most k with a tree \mathcal{T}_i that is a subtree of \mathcal{T} (see Observation 5).

Now fix i . To Φ_i we construct a solution equivalent ACQ-instance $\Phi'_i = (\mathcal{A}'_i, \phi'_i)$ as in the proof of Lemma 8: For each $t \in \mathcal{T}_i$ we construct an atom ϕ_t in the variables χ_t . The associated relation is given by

$$\pi_{\chi_t} \left(\begin{array}{c} \text{\(\(\times\)\)} \quad \phi'(\mathcal{A}_i) \\ \phi' \in \text{atom}(\phi) : \\ \text{var}(\phi') \in \lambda_t \end{array} \right) \text{\(\(\times\)\)} \left(\begin{array}{c} \text{\(\(\times\)\)} \quad \phi'(\mathcal{A}_i) \\ \phi' \in \text{atom}(\phi_i) : \\ \text{var}(\phi') \subseteq \chi_t \end{array} \right),$$

i.e., by taking the natural join of the relations belonging to the atoms of the guard λ_t projected to χ_t and all relations of the atoms in ϕ_i whose variables lie in χ_t . The decomposition Ξ_i has width at most k so this construction can be done in time $\|\Phi\|^{O(k)}$ as seen in the proof of Lemma 8. The query ϕ'_i of Φ'_i is defined as the conjunction of the ϕ_t over all $t \in \mathcal{T}_i$ and with the same quantified variables as ϕ . By Lemma 8, Φ_i and Φ'_i are solution equivalent, we have $\|\Phi'_i\| \leq \|\Phi_i\|^{O(k)}$ and ϕ'_i is acyclic.

Let \mathcal{H}_i be the associated hypergraph of ϕ'_i , then (\mathcal{H}_i, S_i) has only one single S_i -component, because all the vertices in V_i are connected in \mathcal{H} and thus also in \mathcal{H}_i .

We claim that the S_i -star size of \mathcal{H}_i is at most the S_i -star size of $\mathcal{H}[V'_i]$. To see this, consider two independent vertices u, v in \mathcal{H}_i . The edges e of \mathcal{H}_i are equal to the blocks χ_t of \mathcal{E}_i . Because u and v are independent in \mathcal{H}_i , they do not appear in a common block χ_t in \mathcal{E}_i . But then u and v cannot lie in one common edge in $\mathcal{H}[V_i]$, because every edge in $\mathcal{H}[V'_i]$ is contained in a block χ_t by definition of generalized hypertree decompositions. So u and v are independent in $\mathcal{H}[V'_i]$ as well. Thus every independent set in \mathcal{H}_i is also independent in $\mathcal{H}[V'_i]$. So the S_i -star size of \mathcal{H}_i indeed is at most the S_i -star size of $\mathcal{H}[V'_i]$ which is at most ℓ by assumption.

Thus by Lemma 14 the vertices in S_i can be covered by at most ℓ edges e_1, \dots, e_ℓ in \mathcal{H}_i which we can compute in polynomial time. Let $\alpha_1, \dots, \alpha_\ell$ be the atoms corresponding to the edges e_1, \dots, e_ℓ .

We construct a new atomic formula ϕ''_i in the variables S_i and an associated relation \mathcal{R}''_i as follows: For each combination a_1, \dots, a_ℓ of compatible tuples in $\alpha_1(\mathcal{A}'_i), \dots, \alpha_\ell(\mathcal{A}'_i)$ let a be the single tuple in $\pi_{S_i}(\{a_1\} \bowtie \dots \bowtie \{a_\ell\})$. We fix the free variables in ϕ'_i to the constants prescribed by a . The result is a CQ-instance Φ_a with the associated hypergraph $\mathcal{H}[V_i]$. By Observation 6 Φ_a is acyclic and can thus be solved in polynomial time with Theorem 2. If Φ_a has a solution, add a to the relation \mathcal{R}''_i . This completes the construction of \mathcal{R}''_i .

Let \mathcal{A}''_i be the structure containing only the relation \mathcal{R}''_i . By construction, $\Phi''_i := (\mathcal{A}''_i, \phi''_i)$ is solution equivalent to Φ'_i and thus also to Φ_i . Observe that the instances Φ_a can be solved in polynomial time by Theorem 2. Moreover, since $\|\Phi'_i\| \leq \|\Phi\|^{O(k)}$, only $\|\Phi\|^{O(k\ell)}$ tuples a need to be considered. Thus Φ''_i can be constructed in time $\|\Phi_i\|^{p(k,\ell)}$ for a polynomial p .

We now return to the original instance Φ and eliminate the quantified variables in the query ϕ . To do so, we add the atom ϕ''_i for $i \in [m]$ and delete all atoms that contain any quantified variable. Moreover, we add the relation \mathcal{R}''_i to the structure \mathcal{A} . We call the resulting #CQ instance $\Phi'' = (\mathcal{A}'', \phi'')$. The overall runtime of the construction is at most $\|\Phi\|^{p(k,\ell)}$. Also Φ'' is solution equivalent to Φ , because $(\mathcal{A}''_i, \phi''_i)$ is solution equivalent to Φ'_i .

We claim that Φ'' has generalized hypertree width at most k . To show this we construct a generalized hypertree decomposition \mathcal{E}'' of ϕ'' by doing the following: For each $t \in T$ with $\chi_t \cap V_i \neq \emptyset$ we construct a guarded block (λ'_t, χ'_t) by deleting all edges e with $e \cap V_i \neq \emptyset$ from λ_t and adding the edge S_i for ϕ''_i . Furthermore we set $\chi'_t = (\chi_t \setminus V_i) \cup S_i$. It is easy to see that the result is indeed a generalized hypertree decomposition of ϕ'' of width at most k .

Finally, we construct an ACQ-instance $\Psi := (\mathcal{B}, \psi)$ equivalent to Φ'' with Lemma 8. □

We now directly get the desired counting result:

Corollary 1 #CQ on instances Φ of generalized hypertree width k and quantified star size ℓ can be solved in time $\|\Phi\|^{p(k,\ell)}$ for a polynomial p .

Proof Use Theorem 4 to construct a generalized hypertree decomposition of width $O(k)$, then apply Lemma 13 and count with the algorithm of [31] or [15]. \square

4.2 Bounded Quantified Star Size is Necessary

In this section we will show that bounded quantified star size is a necessary restriction for tractable $\#CQ$: Under the assumption $FPT \neq \#W[1]$, all classes \mathcal{G} of S -hypergraphs for which $p\text{-}\#CQ$ is fixed-parameter tractable must have bounded quantified star size. As polynomial time tractability trivially implies fixed-parameter tractability, it follows that bounded quantified star size must also be necessary for classes \mathcal{G} of S -hypergraphs that allow polynomial time algorithms.

Let \mathcal{G} be a class of S -hypergraphs. Remember that by $\#CQ$ on \mathcal{G} we denote the restriction of $\#CQ$ to instances whose associated S -hypergraph is in \mathcal{G} . Analogously, by $p\text{-}\#CQ$ on \mathcal{G} we denote the restriction of $p\text{-}\#CQ$ to instances whose associated S -hypergraph is in \mathcal{G} .

We will use the fact that $\#CQ$ is already hard for very restricted S -hypergraphs, namely those of the queries from the class $\mathcal{C}_{\text{star}}$ from Lemma 12.

Theorem 9 *Assume $FPT \neq \#W[1]$, and let \mathcal{G} be a recursively enumerable class of S -hypergraphs. If $p\text{-}\#CQ$ is fixed-parameter tractable for \mathcal{G} , then \mathcal{G} is of bounded S -star size.*

Before proving Theorem 9, let us take some time to discuss its assumption, because we will see this and similar assumptions throughout the rest of this paper. The reader might feel that it would be more satisfying to prove a version of this theorem not under the assumption $FPT \neq \#W[1]$ from parameterized complexity but instead to prove it based on a more standard assumption like $FP \neq \#P$. Clearly, the statement would then have to change from “fixed-parameter tractable” to “polynomial time tractable”, but this could still be preferable. Unfortunately, it is unlikely that such a version of Theorem 9 can be proved. One can show that assuming $FP \neq \#P$ there are classes of S -graphs on which $\#CQ$ is neither in FP nor $\#P$ -complete (see [29, Chapter 5.3]). Thus it seems unlikely that the theory of $\#P$ -completeness suffices to identify the classes of S -hypergraphs on which $\#CQ$ is tractable.

Furthermore, let us remark that if the reader feels uncomfortable with parameterized complexity, he can safely exchange the assumption $FPT \neq \#W[1]$ against the so-called *exponential time hypothesis* which is the following conjecture.

Conjecture 1 (Exponential time hypothesis) *3-SAT cannot be solved in time $2^{o(n)}$ where n is the number of variables of the input.*

The exponential time hypothesis implies $FPT \neq W[1]$ [14, Chapter 17] and thus also $FPT \neq \#W[1]$. Hence Theorem 9 and several other results of this paper could also be formulated with the assumption that the exponential time hypothesis is true if the reader prefers an assumption from more classical complexity theory.

We will use the following lemma to prove Theorem 9.

Lemma 15 *Let \mathcal{G} be a recursively enumerable class of S -hypergraphs of unbounded S -star size. Then $p\text{-}\#\text{CQ}$ on \mathcal{G} is $\#\mathbf{W}[1]$ -hard.*

Let $\mathcal{G}_{\text{star}}$ be the class of S -graphs (G_n, S_n) where G_n is the star with n leaves and S_n consists of all vertices but the center of G_n . Note that the S -hypergraphs of the queries $\mathcal{C}_{\text{star}}$ from Lemma 12 are the S -graphs in $\mathcal{G}_{\text{star}}$ (see Example 7). Since by Lemma 12 $p\text{-}\#\text{CQ}$ restricted to instances with queries in $\mathcal{C}_{\text{star}}$ is $\#\mathbf{W}[1]$ -hard, it follows directly that $p\text{-}\#\text{CQ}$ on $\mathcal{G}_{\text{star}}$ is $\#\mathbf{W}[1]$ -hard. The idea of the proof of Lemma 15 is to show that $\mathcal{G}_{\text{star}}$ can be embedded in an appropriate way into *any* class \mathcal{G} of S -hypergraphs of unbounded S -star size to show that $p\text{-}\#\text{CQ}$ on \mathcal{G} is $\#\mathbf{W}[1]$ -hard.

We feel that it is more transparent to show Lemma 15 for the restricted case of S -connected S -hypergraphs and to sketch afterwards how to generalize the proof to the general case. Remember that an S -hypergraph (\mathcal{H}, S) is called S -connected if for every pair of vertices x, y there is a path $x = v_1, v_2, \dots, v_{k-1}, v_k = y$ such that $v_i \notin S$ for $i \notin \{1, k\}$.

Lemma 16 *Let \mathcal{G} be a recursively enumerable class of S -connected S -hypergraphs of unbounded S -star size. Then $p\text{-}\#\text{CQ}$ on \mathcal{G} is $\#\mathbf{W}[1]$ -hard.*

Proof Let \mathcal{G} be a class of S -connected S -hypergraphs of unbounded S -star size.

Remember $\mathcal{C}_{\text{star}} := \{\phi_{\text{star},n} \mid n \in \mathbb{N}\}$ is defined with $\phi_{\text{star},n} = \exists z \bigwedge_{i \in [n]} \mathcal{R}_i(z, y_i)$ (see Lemma 12). We will show a parameterized parsimonious reduction from $p\text{-}\#\text{CQ}$, restricted to instances that have queries in $\mathcal{C}_{\text{star}}$, to $p\text{-}\#\text{CQ}$ on \mathcal{G} . As $p\text{-}\#\text{CQ}$ on the former class of instances is $\#\mathbf{W}[1]$ -hard by Lemma 12, the claim will follow.

Let $\Phi = (A, \phi)$ be an instance of $\#\text{CQ}$ restricted to queries in $\mathcal{C}_{\text{star}}$, i.e., ϕ has the form $\phi = \exists z \bigwedge_{i=1}^k \mathcal{R}_i(z, y_i)$. Because \mathcal{G} is recursively enumerable and of unbounded S -star size, there is a computable function $g : \mathbb{N} \rightarrow \mathbb{N}$ such that for given $k \in \mathbb{N}$ one can in time $g(k)$ compute an S -connected S -hypergraph $(\mathcal{H}, S) \in \mathcal{G}$ of S -star size at least k . We will embed Φ into $\mathcal{H} = (V, E)$ to construct a $\#\text{CQ}$ -instance $\Psi := (\mathcal{B}, \psi)$ of size at most $g(k)\|\Phi\|^2$. The instance Ψ will have the S -hypergraph (\mathcal{H}, S) and the same domain $B := A$ as Φ .

For each $e \in E$ let ψ_e be an atom with the relation symbol \mathcal{E}_e and the set of variables $\text{var}(\psi_e) = e$. Let

$$\psi' := \bigwedge_{e \in E} \psi_e,$$

then ψ is the query we get from ψ' by existentially quantifying all variables in $V \setminus S$. This completes the construction of the query ψ .

We now construct the structure \mathcal{B} . Let $Y = \{y_1, \dots, y_k\} \subseteq S$ be a set of independent vertices. Such a set Y must exist, because (\mathcal{H}, S) has S -star size at least k . Let d be an arbitrary but fixed element of A . We define $\mathcal{E}_e^{\mathcal{B}}$ depending on the vertices in e as follows:

Case 1 Let first $e \in E$ be an edge that contains y_i for some $i \in [k]$. Observe that y_i is uniquely determined, because no two of the vertices y_i share an edge. The atom ψ_e has the relation symbol \mathcal{E}_e and as variables the vertices of e . We assume that the order of the variables in ψ_e is as follows: y_i is the first variable, followed

by the other variables in $e \cap S$ and after those the variables in $e \setminus S$. We define $\mathcal{E}_e^{\mathcal{B}} := \{(v_2, d, \dots, d, v_1, \dots, v_1) \mid (v_1, v_2) \in \mathcal{R}_i^{\mathcal{A}}\}$, where $\mathcal{R}_i^{\mathcal{A}}$ is the relation of \mathcal{R}_i in \mathcal{A} . Observe that this forces all variables in $(e \cap S) \setminus \{y_i\}$ to be equal to the value d in satisfying assignments, while the variables in $e \setminus S$ must all have a common value v_1 . Furthermore, because y_i is uniquely determined, the relation $\mathcal{E}_e^{\mathcal{B}}$ is well defined.

Case 2 Let now $e \in E$ with $e \cap Y = \emptyset$. We assume that the variables in ψ_e are ordered such that all variables in $S \cap e$ appear before those in $e \setminus S$. Then we define $\mathcal{E}_e^{\mathcal{B}} := \{(d, \dots, d, v_1, \dots, v_1) \mid v_1 \in A\}$. Again in the satisfying assignments all variables in $e \cap S$ are forced to be equal to d , while the variables in $e \setminus S$ can take an arbitrary but equal value.

This completes the construction of \mathcal{B} and thus that of $\Psi = (\mathcal{B}, \psi)$.

Claim 1 $|\phi(\mathcal{A})| = |\psi(\mathcal{B})|$. □

Proof Let ϕ' , resp., ψ' be the quantifier free queries we get from ϕ , resp., ψ by deleting all quantifiers.

We construct a function B that to an assignment $a \in \phi'(\mathcal{A})$ constructs an assignment $B(a) := a'$ with $a' : V \rightarrow B$. We define

$$a'(x) := \begin{cases} a(x), & x \in Y, \\ d, & x \in S \setminus Y, \\ a(z), & x \in V \setminus Y \end{cases}$$

We claim that B is a bijection from $\phi(\mathcal{A})$ to $\psi(\mathcal{B})$. It is easily seen from the construction of ψ that a' satisfies all atoms of ψ and thus $a' \in \psi'(\mathcal{B})$. Furthermore, B is obviously injective. Thus it only remains to prove that B is surjective. To see this, consider $b' \in \psi(\mathcal{B})$. By construction of Ψ , we have $b'(x) = d$ for all $x \in S \setminus Y$. Because \mathcal{H} is S -connected, we have that $\mathcal{H}[V \setminus S]$ is connected. From the construction of Ψ it follows by an easy induction that there is a $v_1 \in A$ such that $b'(x) = v_1$ for all $x \in V \setminus S$. We construct an assignment $b : \text{var}(\phi) \rightarrow A$ by $b(x) := b'(x)$ for $x \in \{y_1, \dots, y_k\}$ and $b(z) := v_1$. Obviously, $B(b) = b'$. Moreover, from the construction of Ψ it follows that $b \in \phi'(\mathcal{A})$. Thus B is a bijection from $\phi'(\mathcal{A})$ to $\psi'(\mathcal{B})$.

We now construct a mapping B' from $\phi(\mathcal{A})$ to $\psi(\mathcal{B})$ as follows: For $a \in \phi'(\mathcal{A})$ we map $a|_{\text{free}(\phi)}$ to $B(a)|_{\text{free}(\psi)}$. Since B is a bijection, it follows that B' is a bijection as well. This proves the claim. □

Obviously, the S -hypergraph associated to ψ is (\mathcal{H}, S) . Moreover, by construction we have $|\psi| \leq g(k)$ and Ψ can be constructed in time at most $g(k)\|\Phi\|^2$, because \mathcal{H} has size at most $g(k)$ and the size of the relations is bounded by $|A|^2$. Thus, with Claim 1, the construction of Ψ from Φ is a parameterized parsimonious reduction. This completes the proof of Lemma 16. □

We now sketch how to extend Lemma 16 from S -connected S -hypergraphs to general S -hypergraphs in a straightforward way.

Proof of Lemma 15 (Sketch) The proof follows the same ideas as that of Lemma 16: We first compute an S -hypergraph \mathcal{H} in \mathcal{G} of S -star size at least k . Then we choose an S -component \mathcal{H}' of S -star size at least k in \mathcal{G} . We construct the relations $\mathcal{E}_e^{\mathcal{B}}$ in such a way that in every satisfying assignment every variable not in \mathcal{H} is forced to the value d . For all other variables we construct the relations as in the proof of Lemma 16. Since \mathcal{H}' is S -connected by Observation 8, the same arguments as in the proof of Lemma 16 show that the construction is a parsimonious parameterized reduction. \square

Proof of Theorem 9 Assume that p -#CQ on \mathcal{G} is fixed-parameter tractable. By Lemma 15 we directly get $\mathbf{FPT} = \#\mathbf{W}[1]$ which contradicts the assumption. \square

Combining Corollary 1 with Theorem 9 yields a characterization of classes of S -hypergraphs of bounded generalized hypertree width that allow efficient algorithms for #CQ.

Theorem 10 *Let \mathcal{G} be a recursively enumerable class of S -hypergraphs of bounded generalized hypertree width. Then (assuming $\mathbf{FPT} \neq \#\mathbf{W}[1]$) the following statements are equivalent:*

1. #CQ on \mathcal{G} is polynomial time tractable.
2. p -#CQ on \mathcal{G} is fixed-parameter tractable.
3. \mathcal{G} is of bounded S -star size.

Proof 1 \rightarrow 2 is trivial. 2 \rightarrow 3 is Theorem 9. Finally, 3 \rightarrow 1 is Corollary 1. \square

As a corollary we get that for a wide range of decomposition techniques commonly considered in the database and artificial intelligence literature, we can characterize the tractable classes of S -graphs by bounded quantified star size. For the decomposition techniques not defined here see [18].

Corollary 2 *Let β be one of the following decomposition techniques: biconnected component, cycle-cutset, cycle-hypercutset, hingetree, hypertree, or generalized hypertree decomposition. Let furthermore \mathcal{G} be a recursively enumerable class of S -hypergraphs of bounded β -width. Then (assuming $\mathbf{FPT} \neq \#\mathbf{W}[1]$), the following statements are equivalent:*

1. #CQ on \mathcal{G} is polynomial time tractable.
2. p -#CQ on \mathcal{G} is fixed-parameter tractable.
3. \mathcal{G} is of bounded S -star size.

Proof 1 \rightarrow 2 is trivial. 2 \rightarrow 3 follows from Theorem 9. For 3 \rightarrow 1 observe that for every β of the claim we have that for every hypergraph \mathcal{H} the β -width of \mathcal{H} is

bounded from below by a function in the generalized hypertree width of \mathcal{H} . Thus \mathcal{G} has bounded generalized hypertree width and the claim follows with Corollary 1. \square

5 Queries of Bounded Arity

In this section we show that for bounded arity #CQ we can exactly characterize the classes of S -hypergraphs that allow polynomial time counting. In this section all CQ-instances and all S -hypergraphs are always assumed to be of bounded arity.

We will give two different characterizations of S -hypergraphs of bounded arity that allow tractable #CQ: The first characterization is presented in Section 5.1 and uses treewidth and S -star size, following the ideas of Section 4. In Section 5.2 we introduce a notion of elimination width for conjunctive queries. It will allow us to characterize the S -hypergraphs of bounded arity that allow tractable #CQ with a single parameter.

5.1 A Characterization by Treewidth and S -star size

In this section we characterize the S -hypergraphs of bounded arity that allow tractable #CQ by treewidth and S -star size. The result of this section is based on a combination of the results of Section 4 and a result by Grohe from [22] which is a follow up of results by Grohe, Schwentick and Segoufin [24]. We state the theorem in our slightly different wording.

Theorem 11 ([22]) *Let \mathcal{G} be a recursively enumerable class of hypergraphs of bounded arity. Assume $\mathbf{FPT} \neq \mathbf{W}[1]$. Then the following three statements are equivalent:*

1. CQ on \mathcal{G} can be decided in polynomial time.
2. p -CQ on \mathcal{G} is fixed parameter tractable.
3. There is a constant c such that the hypergraphs in \mathcal{G} have treewidth at most c .

Theorem 11 is originally stated even for every fixed vocabulary. Moreover, in [22] characterizes, more generally, tractability for any class of queries (and not only in function of their underlying hypergraph) of bounded arity. In this setting a class of queries is tractable if and only if its cores, certain equivalent subqueries, are of unbounded treewidth. In particular, this yields tractable classes of queries of unbounded treewidth. With the more coarse hypergraph view, these unbounded treewidth cases cannot be captured as witnessed by Theorem 11.

Our goal is to provide a complete characterization of classes of S -hypergraphs of bounded arity that yield tractability for #CQ. Not too surprisingly, tractability depends on both treewidth and star size of the underlying S -hypergraph.

Theorem 12 *Let \mathcal{G} be a recursively enumerable class of S -hypergraphs of bounded arity. Assume that $\mathbf{W}[1] \neq \mathbf{FPT}$. Then the following statements are equivalent:*

1. #CQ on \mathcal{G} is solvable in polynomial time.

2. p -#CQ on \mathcal{G} is fixed-parameter tractable.
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 .

Let us discuss how Theorem 12 and Theorem 10 relate. First, it is not hard to see that for bounded arity hypergraphs treewidth and generalized hypertree width differ only by a constant factor. So we could have formulated Theorem 12 with generalized hypertree width instead of treewidth as well.

The key difference between Theorem 12 and Theorem 10 is that we can show here that bounded treewidth is not only sufficient for tractable counting but also necessary. As we already directly get from [31], there are by Lemma 2 families of S -hypergraphs of unbounded arity, and thus also unbounded treewidth, on which #CQ is tractable, so treewidth is not the right notion for this case. It is an intriguing question if there is a width measure that completely characterizes tractable CQ or tractable #CQ for unbounded arity, similarly to Theorem 11 and Theorem 12 in the bounded arity case.

Before giving the proof of Theorem 12 we make an observation.

Observation 13 *If there is a recursively enumerable class \mathcal{G} of S -hypergraphs of unbounded treewidth such that p -#CQ on \mathcal{G} is fixed-parameter tractable, then there is such a class \mathcal{G}' that is recursive.*

Proof Fix a Turing machine M that enumerates \mathcal{G} . Let the order in which the S -hypergraphs of \mathcal{G} are enumerated by M be $(\mathcal{H}_1, S_1), (\mathcal{H}_2, S_2), \dots$. Then define \mathcal{G}' as containing the S -hypergraphs (\mathcal{H}'_i, S'_i) where \mathcal{H}'_i is the disjoint union of the hypergraphs $\mathcal{H}_1, \dots, \mathcal{H}_i$ and $S'_i := \bigcup_{j \in [i]} S_j$.

We claim that \mathcal{G}' is recursive. Indeed the definition of \mathcal{G}' directly gives an algorithm that enumerates the elements of \mathcal{G}' ordered by size. This yields an algorithm to decide membership in \mathcal{G}' : Given an input (\mathcal{H}, S) , enumerate the elements of \mathcal{G}' until (\mathcal{H}, S) is found or an element that has more vertices than (\mathcal{H}, S) is enumerated. The treewidth of \mathcal{G}' is trivially unbounded.

Finally, we claim that #CQ on \mathcal{G}' is fixed-parameter tractable. Given an input $\Phi := (\mathcal{A}, \phi)$ first check if the associated S -hypergraph (\mathcal{H}, S) is in \mathcal{G}' . If not, stop. If yes, the query ϕ must decompose into subqueries ϕ_1, \dots, ϕ_i such that for each $j \in [i]$ the query ϕ_j has the S -hypergraph (\mathcal{H}_j, S_j) and the ϕ_j have disjoint variable sets. Using the enumerating machine M we can compute such a decomposition. Now since #CQ on \mathcal{G} is fixed-parameter tractable we can solve the instances $\Phi_j := (\mathcal{A}, \phi_j)$ in time $g(|\phi_j|) \|\Phi_j\|^c$ for a computable function g and a constant c . It follows that $|\phi(\mathcal{A})| = \prod_{j \in [i]} |\phi_j(\mathcal{A})|$ can be computed in time $\sum_{j \in [i]} g(|\phi_j|) \|\Phi_j\|^c \leq |\phi| g(|\phi|) \|\Phi\|^c$ and thus #CQ on \mathcal{G}' is fixed-parameter tractable. \square

Proof of Theorem 12 The direction $1 \rightarrow 2$ is trivial. Furthermore, $3 \rightarrow 1$ follows directly from Corollary 2. So it remains only to show $2 \rightarrow 3$.

By way of contradiction, we assume that there is a recursively enumerable class \mathcal{G} of S -hypergraphs such that counting solutions to #CQ-instances, whose S -hypergraph are in \mathcal{G} , is fixed parameter tractable, but 3 is not satisfied by \mathcal{G} . From

Theorem 9 we know that the S -starsize of \mathcal{G} must be bounded, so it follows that the treewidth of \mathcal{G} is unbounded. With Observation 13 we may assume that \mathcal{G} is recursive.

We construct a class \mathcal{G}' of hypergraphs as

$$\mathcal{G}' := \{\mathcal{H} \mid (\mathcal{H}, S) \in \mathcal{G}\}.$$

Clearly \mathcal{G}' is recursive and of unbounded treewidth. We will show that p -CQ on \mathcal{G}' is fixed-parameter tractable. This is a contradiction with Theorem 11.

Because \mathcal{G} is recursive, there is an algorithm that for each \mathcal{H} in \mathcal{G}' constructs an S -hypergraph (\mathcal{H}, S) in \mathcal{G} . For example, one can simply try all vertex sets S and check if (\mathcal{H}, S) is in \mathcal{G} . Let $f(\mathcal{H})$ be the number of steps the algorithm needs on input \mathcal{H} . The function $f(\mathcal{H})$ is well defined and computable. We then define $g : \mathbb{N} \rightarrow \mathbb{N}$ by setting $g(k) := \max_{\mathcal{H}}(f(\mathcal{H}))$, where the maximum is over all hypergraphs \mathcal{H} of size k in \mathcal{G}' . The function g is well defined and computable, because \mathcal{G}' is recursive. Thus for each \mathcal{H} in \mathcal{G}' we can compute in time $g(|\mathcal{H}|)$ an S -hypergraph (\mathcal{H}, S) in \mathcal{G} .

Now let $\Phi = (\mathcal{A}, \phi)$ be a CQ-instance with hypergraph \mathcal{H} in \mathcal{G}' . To solve it we first compute (\mathcal{H}, S) as above and construct a CQ-instance $\Psi = (\mathcal{A}, \psi)$ with (\mathcal{H}, S) as associated S -hypergraph for ψ by adding existential quantifiers for all variables not in S . Obviously Φ has solutions if and only if Ψ has one. But by assumption the solutions of Ψ can be counted in time $h(|\psi|)\|\Psi\|^{O(1)}$ for some computable function h , so Φ can be decided in time $(g(|\phi|) + h(|\phi|))\|\Phi\|^{O(1)}$. Thus p -CQ on \mathcal{G}' is fixed-parameter tractable. This is the desired contradiction to Theorem 11. \square

We remark that Dalmau and Jonsson in [12] characterize the tractable classes of quantifier free queries for #CQ by bounded treewidth. This result is proved under the weaker assumption $\#\mathbf{W}[1] \neq \mathbf{FPT}$. Showing a version of this for general #CQ seems under that assumption is likely hard since our case also contains decision problems (e.g., #CQ with no free variables).

5.2 A Characterization by Elimination Orders

While the characterization of Theorem 12 is great because it completely characterizes the tractable classes of S -hypergraphs for #CQ, it has the somewhat unpleasant property that we have to bound two different parameters of the hypergraphs instead of just one. Also, it is not clear how robust and natural the defined classes of hypergraphs are. In contrast to this, treewidth is a very robust notion that has many equivalent definitions.

In this section we improve the situation by showing that there is a notion of elimination width for S -hypergraphs that is equivalent to the combination of treewidth and S -star size.

Recall the notion of elimination orders from Section 2.6.1.

Definition 15 Let (G, S) be an S -graph. We define an *elimination order of an S -graph* π of an S -graph (G, S) as an elimination order of $G = (V, E)$ such that for each pair $v \in S, u \in V \setminus S$ such that uv is an edge in the fill-in graph G_π we have $\pi(u) < \pi(v)$.

The *elimination width* of an S -graph $\text{elim-width}(G, S)$ of (G, S) is defined as the minimum width taken over all elimination orders of (G, S) .

The elimination width $\text{elim-width}(\mathcal{H}, S)$ of an S -hypergraph (\mathcal{H}, S) is defined as the elimination width of its primal S -graph (see Definition 7). By $\mathcal{H}_{P,\pi}$ we denote the fill-in graph of the primal graph \mathcal{H}_P of \mathcal{H} with respect to π .

Remark 1 Observe that for every S -graph (G, S) we have

$$\text{elim-width}(G, S) \geq \text{elim-width}(G),$$

because every elimination order of (G, S) is an elimination order of G .

Proposition 1 *Let \mathcal{G} be a class of S -hypergraphs. Then the following statements are equivalent:*

- *The treewidth and the S -star size of the S -hypergraphs in \mathcal{G} are bounded by a constant c .*
- *The elimination width of the S -hypergraphs in \mathcal{G} is bounded by a constant c' .*

The proof of Proposition 1 is somewhat lengthy, so we prove it in two individual lemmas.

Lemma 17 *Let (\mathcal{H}, S) be an S -hypergraph of elimination width k . Then the treewidth of \mathcal{H} is at most k and the S -star size of (\mathcal{H}, S) is at most $k + 1$.*

Proof From Remark 1 and Lemma 4 it follows directly that the treewidth of \mathcal{H} is at most k . Thus we only have to show the bound on the S -star size of (\mathcal{H}, S) . To this end, we define an S -path (P, S) as a path whose end vertices are in S but all other vertices are not in S .

Claim 2 *Let u, v be the end vertices of an S -path (P, S) with $P = ux_1 \dots x_\ell v$. Then for every elimination order π of (P, S) we have $\pi(v) > \pi(x_i)$ and $\pi(u) > \pi(x_i)$ for all $i \in [\ell]$. Furthermore, uv is an edge of the fill-in graph P_π . \square*

Proof We prove this by induction on ℓ . For $\ell = 0$ there is nothing to show.

Now let $\ell \geq 1$. Let x_j be the vertex for which π is minimal, i.e., $\pi(x_j) \leq \pi(x_i)$ for all $i \in [\ell]$. By the definition of elimination orders we have $\pi(x_1) < \pi(u)$ and $\pi(x_\ell) < \pi(v)$, so $\pi(x_j) < \min(\pi(v), \pi(u))$. Let P' be the path that we get from P when deleting x_j and connecting x_{j-1} and x_{j+1} by an edge. P' is a subgraph of the fill-in graph P_π and π induces an elimination order on P' by $\pi'(w) := \pi(w) - 1$. It follows that $P'_{\pi'}$ is a subgraph of P_π . By induction $\pi'(v) > \pi'(x_i)$ and $\pi'(u) > \pi'(x_i)$ for all $i \in [\ell] \setminus \{j\}$ and thus $\pi(v) > \pi(x_i)$ and $\pi(u) > \pi(x_i)$ for all $i \in [\ell]$. Furthermore, by induction uv is an edge of $P'_{\pi'}$, and thus also of P_π . This completes the proof of the claim. \square

By definition of S -components, in every S -graph (G, S) every pair $u, v \in S$ must be connected by an S -path.

Let $\mathcal{H}_P = (V, E_P)$ be the primal graph of \mathcal{H} . Let \mathcal{H}' be an S -component of \mathcal{H} with primal S -graph $\mathcal{H}'_P = (V', E'_P)$ and let $S' := S \cap V'$.

Let π be an optimal elimination order of (\mathcal{H}, S) of width k . Then π induces for every subgraph \mathcal{H}'' an elimination order of \mathcal{H}'' of width at most k . To ease notation we will not differentiate between π and these induced elimination orders and simply call π an elimination order of all subgraphs, too.

As already remarked, all pairs $u, v \in S'$ are connected by S -paths in \mathcal{H}'_P . The fill-in graph of every subgraph of \mathcal{H}'_P is a subgraph of the fill-in graph $\mathcal{H}'_{P,\pi}$ of \mathcal{H}'_P . Thus by Claim 2 we have that the vertices in S' form a clique in $\mathcal{H}'_{P,\pi}$. Because π has width k , it follows that $|S'| \leq k + 1$. Hence the S -star size of \mathcal{H}' is at most $k + 1$. This completes the proof of Lemma 17.

For the other direction of Proposition 1 we will use the following lemma.

Lemma 18 *Let (\mathcal{H}, S) be an S -hypergraph of treewidth at most c and S -star size at most k . Then every S -component of \mathcal{H} contains at most $k(c + 1)$ vertices from S .*

Proof We prove this by induction on the S -star size k while keeping the treewidth fixed to c . If the S -star size is $k = 1$, then in every S -component all vertices from S are adjacent. But then they induce a clique in the primal graph of \mathcal{H} and thus by Lemma 2 there may be at most $c + 1$ of them.

Let now $k > 1$. Consider an S -component \mathcal{H}' of \mathcal{H} . The graph $\mathcal{H}'_P[S]$ has at most treewidth c because it is an induced subgraph of \mathcal{H}'_P which has by assumption treewidth at most c . By Lemma 3, there is a vertex v in $\mathcal{H}'_P[S]$ of degree at most c . It follows that v has at most c neighbors in S in \mathcal{H}' . Let \mathcal{H}'' be the hypergraph we get from $\mathcal{H}' = (V', E')$ by deleting v and all of its neighbors in S . We claim that $(\mathcal{H}'', S \cap V')$ has S -star size at most $k - 1$.

Assuming this is false, there are k independent vertices $v_1, \dots, v_k \in S$ in \mathcal{H}'' . But then v_1, \dots, v_k, v are $k + 1$ independent vertices from S in \mathcal{H}' , so the S -star size of \mathcal{H} is at most $k + 1$ which contradicts the assumption.

So the S -star size of \mathcal{H}'' is indeed bounded by $k - 1$. By induction \mathcal{H}'' contains at most $(k - 1)(c + 1)$ vertices from S , and since we deleted at most $c + 1$ vertices during the construction of \mathcal{H}'' we get that \mathcal{H}' contains at most $k(c + 1)$ vertices from S . \square

We now prove the second direction of Proposition 1

Lemma 19 *Let (\mathcal{H}, S) be an S -hypergraph such that the treewidth and the S -star size of (\mathcal{H}, S) are bounded by $c \in \mathbb{N}$. Then the elimination width of (\mathcal{H}, S) is at most $(c + 1)^3 + (c + 1)^2$.*

Proof Let $(\mathcal{T}, (\chi_t)_{t \in T})$ be a tree decomposition of $\mathcal{H} = (V, E)$ of minimal width ℓ . Let $S(v)$ for every $v \in V \setminus S$ be the set of vertices from S in the S -component of v . For every $t \in T$ we construct a new bag χ'_t as

$$\chi'_t := \chi_t \cup \bigcup_{v \in (V \setminus S) \cap \chi_t} S(v).$$

Because the S -star size of \mathcal{H} is at most c we get by Lemma 18 that $|S(v)| \leq c(\ell + 1)$. It follows with $\ell \leq c$ that

$$\begin{aligned} |\chi'_t| &\leq |\chi_t| + \sum_{v \in (V \setminus S) \cap \chi_t} |S(v)| \leq (\ell + 1) + (\ell + 1)c(\ell + 1) \\ &\leq (\ell + 1)(c + 1)(\ell + 1) \\ &\leq (c + 1)^3 \end{aligned}$$

It is easy to see that $(\mathcal{T}, (\chi'_t)_{t \in T})$ is a tree decomposition. Remember that for each $t \in T$ the tree \mathcal{T}_t is the subtree of \mathcal{T} with t as its root. Let V_t be the set of vertices appearing in the bags χ'_t of \mathcal{T}_t . For each $y \in V$ let $r(y)$ be the $t \in T$ with $y \in \chi'_t$ that is nearest to the root of \mathcal{T} .

Claim 3 *There exists $t \in T$ such that $\emptyset \neq V_t \setminus S \subseteq \chi'_t$ with a vertex $y \in V_t \setminus S$ with $t = r(y)$.* □

Proof We find t and y by descending in \mathcal{T} . Let r be the root of \mathcal{T} . If $V_r \setminus S \subseteq \chi'_r$ we are done. Otherwise let t_i be a child of r such that $V_{t_i} \setminus S \not\subseteq \chi'_{t_i}$. Now check if $V_{t_i} \setminus S \subseteq \chi'_{t_i}$ and if not go deeper in \mathcal{T} . Let t be the first vertex on this descent with $V_t \setminus S \subseteq \chi'_t$. Then χ'_t must contain a vertex y that is not in $\chi'_{t'}$ where t' is the parent of t in \mathcal{T} . But then $r(y) = t$ as desired. □

We construct an elimination order π of \mathcal{H} inductively as follows, starting from the empty elimination order: While any bag of the tree decomposition $(\mathcal{T}, (\chi'_t)_{t \in T})$ contains a vertex from $V \setminus S$, do the following: Choose by Claim 3 $t \in T$ such that $\emptyset \neq V_t \setminus S \subseteq \chi'_t$ with a vertex $y \in V_t \setminus S$ with $t = r(y)$, delete y from \mathcal{H} and all bags and add y as the next vertex to the elimination order π .

When the vertices from $V \setminus S$ have all been deleted, we proceed with the vertices in S in a similar fashion: While there is a non-empty bag, choose one $t \in T$ with $\emptyset \neq S \cap V_t \subseteq \chi'_t$ and $y \in S \cap V_t$ with $r(y) = t$, delete y and add y as the next vertex in π . Again, such t and y can always be found.

All vertices in $V \setminus S$ appear before all vertices in S in π , so π is an elimination order of (\mathcal{H}, S) . We will now bound the width of π .

Claim 4 *Let $x, y \in V$ with $x, y \in V \setminus S$ or $x, y \in S$ such that there exists $t \in T$ with $x, y \in \chi'_t$. If x is higher-numbered than y with respect to π , then $x \in \chi_{r(y)}$.*

Proof x and y appear in a common bag χ'_t and thus $x \in V_{r(y)}$. But x is higher-numbered, so y was deleted before x . Hence, when y was chosen to be deleted the vertex x was still in $V_{r(y)}$. But then $x \in \chi_{r(y)}$ because otherwise y would not have been chosen for deletion. □

Claim 5 *a) For every vertex $y \in V \setminus S$ the neighbors of y in $\mathcal{H}_{P,\pi}$ are vertices of the same S -component as y .*

b) When a vertex $y \in V \setminus S$ is deleted, the bag $\chi'_{r(y)}$ contains all higher-numbered neighbors of y in the fill-in graph $\mathcal{H}_{P,\pi}$ that are in $V \setminus S$.

c) When a vertex $y \in S$ is deleted, the bag $\chi'_{r(y)}$ contains all higher-numbered neighbors of y in the fill-in graph $\mathcal{H}_{P,\pi}$.

Proof We first prove a) and b) by induction along the elimination order π . So let first y be the vertex with $\pi(y) = 1$. We claim that the higher-numbered neighbors of y in $\mathcal{H}_{P,\pi}$ are simply the neighbors of y in \mathcal{H}_P . Certainly, these are all higher-numbered. Also, in the construction of $\mathcal{H}_{P,\pi}$ from \mathcal{H}_P edges incident to y may only be added by lower-numbered vertices. As there are none for y , all neighbors of y in $\mathcal{H}_{P,\pi}$ are already neighbors in \mathcal{H}_P . This proves the induction start for a). For b) consider a neighbor $x \in V \setminus S$ of y in \mathcal{H} . By the definition of tree decompositions x and y must be in one common bag χ'_t . With Claim 4 it follows that $x \in \chi_{r(y)}$.

Consider now $y \in V \setminus S$ with $\pi(y) > 1$. All neighbors $x \in V \setminus S$ of y in $\mathcal{H}_{P,\pi}$ are either already neighbors of y in \mathcal{H} and thus in the same S -component as y or they are neighbors that originate from edges added in the construction of $\mathcal{H}_{P,\pi}$ from \mathcal{H}_P . In the latter case the edge xy must have been added because of a common lower-numbered neighbor v of y and x . Because v is lower-numbered than y it follows that $v \in V \setminus S$. By induction all higher-numbered neighbors of v in $\mathcal{H}_{P,\pi}$ are in the same S -component as v in \mathcal{H} , so y, x and v are all vertices of the same S -component which completes the proof of a).

Now let $x \in V \setminus S$ be a higher-numbered neighbor of $y \in V \setminus S$ in $\mathcal{H}_{P,\pi}$. Consider first the case that xy is already an edge in \mathcal{H}_P . Then there is a bag χ'_t such that $x, y \in \chi'_t$. With Claim 4 we get $x \in \chi_{r(y)}$ as desired. If x and y are not neighbors in \mathcal{H}_P , then there is a lower-numbered vertex v of x and y in $\mathcal{H}_{P,\pi}$ that led to the introduction of the edge xy . By induction $x, y \in \chi_{r(v)}$. We conclude with Claim 4 that $x \in \chi_{r(y)}$. This completes the proof of b).

To prove c) consider a vertex $y \in S$. Let x be a higher-numbered neighbor of y in $\mathcal{H}_{P,\pi}$. By construction $x \in S$. Assume first that x and y are in a common S -component. Let $v \in V \setminus S$ be a vertex of this S -component, then, by construction of $(\mathcal{T}, (\chi'_t)_{t \in \mathcal{T}})$, the vertices x and y both appear in any bag χ'_t that contains v . We conclude with Claim 4 that $x \in \chi_{r(y)}$. If x and y are not in a common S -component, then the vertex v that leads to the introduction of the edge xy must be in S by a). Because v is a lower-numbered neighbor of x and y , we have by induction that $x, y \in \chi_{r(v)}$. By Claim 4 we get $x \in \chi_{r(y)}$ which completes the proof of Claim 5. \square

We claim that the width of π is at most $(c + 1)^3 + (c + 1)^2$. As the bags χ'_t have size at most $(c + 1)^3 + 1$, every vertex $y \in V \setminus S$ has at most $(c + 1)^3$ higher-numbered neighbors in $V \setminus S$ in $\mathcal{H}_{P,\pi}$ by Claim 5. Furthermore, $y \in V \setminus S$ has by Claim 5 and Lemma 18 at most $(c + 1)^2$ neighbors in S in $\mathcal{H}_{P,\pi}$. Finally, $y \in S$ has at most $(c + 1)^3$ higher-numbered neighbors by Claim 5 and the bound on the size of the bags. This completes the proof of Lemma 19.

From Proposition 1 and Theorem 12 we get the following alternative characterization of S -hypergraphs of bounded arity that allow tractable #CQ.

Theorem 14 *Let \mathcal{G} be a recursively enumerable class of S -hypergraphs of bounded arity. Assume that $\mathbf{W[1]} \neq \mathbf{FPT}$. Then the following statements are equivalent:*

1. *#CQ on \mathcal{G} is solvable in polynomial time.*
2. *p -#CQ on \mathcal{G} is fixed-parameter tractable.*
3. *There is a constant c such that all S -hypergraphs in \mathcal{G} have elimination width at most c .*

Let us remark that there is a similar notion of elimination width for quantified constraint satisfaction (QCSP) which is a version of CQ in which also universal quantification is allowed. Chen and Dalmou [9] introduced this measure and showed that it characterizes the tractable classes of graphs for QCSP. We consider it as likely that an equivalent characterization of the same classes of graphs could be given by treewidth and an adapted notion of S -star size. This would probably also make it possible to get a better understanding of tractable classes of hypergraphs of unbounded arity for QCSP by exchanging treewidth for e.g., generalized hypertree width.

6 Computing Star Size

In this section we consider the problem of computing the quantified star size of hypergraphs that have small width for the decomposition techniques defined in Section 2.6. Note that the computation of quantified star size is not strictly necessary for tractable counting. The algorithm of Section 4 does not need to compute the S -star size for graphs of width k but only for acyclic hypergraphs which can be done with the help of Lemma 14 (see [15] for the details). Still it is of course desirable to know the quantified star size of an instance before applying the counting algorithm, because quantified star size has an exponential influence on the runtime.

We show that for all decomposition techniques considered in this paper the quantified star size can be computed rather efficiently, in time roughly $|V|^{O(k)}$ where k is the width of the input. For small values of k , this bound is reasonable. We then proceed by showing that, on the one hand, for some decomposition measures such as treewidth or hingetree width, the computation of quantified star size is even fixed parameter tractable parameterized by the width. On the other hand, we show that for decomposition measures above hypertree width it is unlikely that fixed parameter tractability can be obtained (under standard assumptions).

Instead of tackling quantified star size directly, we consider the combinatorially less complicated notion of independent sets. This is justified by the following observation:

Observation 15 *Let β be any decomposition technique considered in this paper. Then, for every $k \in \mathbb{N}$, computing the S -star size of S -hypergraphs of β -width at most k polynomial time Turing-reduces to computing the size of a maximum independent*

set for hypergraphs of β -width at most k . Furthermore, there is a polynomial time many one reduction from computing the size of a maximum independent set in hypergraphs of β -width at most k to computing the S -star size of hypergraphs of β -width at most $k + 1$.

Proof By definition computing S -starsize reduces to the computation of independent sets of S -components. S -components are induced subhypergraphs, so we get the first direction from Observation 5.

For the other direction let $\mathcal{H} = (V, E)$ be a hypergraph for which we want to compute the size of a maximum independent set. Let $x \notin V$. We construct the hypergraph \mathcal{H}' of vertex set $V' = V \cup \{x\}$ and edge set $E' = \{e \cup \{x\} \mid e \in E\}$ and set $S := V$. The hypergraph is one single S -component, because x is in every edge. Furthermore, the S -starsize of \mathcal{H}' is obviously the size of a maximum independent set in \mathcal{H} . It is easy to see that the construction increases the treewidth of the hypergraph by at most 1 and does not increase the β -width for all other decomposition considered here at all. □

Because of Observation 15 we will not talk about S -star size in this section anymore but instead formulate everything with independent sets.

6.1 Exact Computation

Proposition 2 *There is an algorithm that, given a hypergraph $\mathcal{H} = (V, E)$ and a generalized hypertree decomposition $(\mathcal{T}, (\lambda_t)_{t \in \mathcal{T}}, (\chi_t)_{t \in \mathcal{T}})$ of width k of \mathcal{H} , computes a maximum independent set of \mathcal{H} in time $k|V|^{O(k)}$.*

Proof We apply dynamic programming along the decomposition. Let $b = (\lambda, \chi)$ be a guarded block of \mathcal{T} . Let \mathcal{T}_b be the subtree of \mathcal{T} with b as its root. We set $V_b := \chi(\mathcal{T}_b)$. Observe that $I \subseteq V_b$ is independent in \mathcal{H} if and only if it is independent in $\mathcal{H}[V_b]$ so we do not differentiate between the two notions. For each independent set $\sigma \subseteq \chi$ we will compute an independent set $I_{b,\sigma} \subseteq V_b$ that is maximum under the independent sets containing exactly the vertices σ from χ . Observe that because λ contains at most k edges that cover χ we have to compute at most kn^k independent sets $I_{b,\sigma}$ for each b .

If b is a leaf of \mathcal{T} , the construction of the $I_{b,\sigma}$ is straightforward and can certainly be done in time $k|V|^{O(k)}$.

Let now $b = (\lambda, \chi)$ be an inner vertex of \mathcal{T} with children b_1, \dots, b_r and let $b_i = (\lambda_i, \chi_i)$. For each independent set $\sigma \subseteq \chi$ we do the following: For each i , let σ_i be an independent set of χ_i such that $\sigma \cap \chi \cap \chi_i = \sigma_i \cap \chi \cap \chi_i$ and $|I_{b_i,\sigma_i}|$ is maximal. We claim that we can set $I_{b,\sigma} := \sigma \cup I_{b_1,\sigma_1} \cup \dots \cup I_{b_r,\sigma_r}$.

We first show that $I_{b,\sigma}$ defined this way is independent. Assume this is not true, then $I_{b,\sigma}$ contains x, y that are in one common edge e in $\mathcal{H}[V_b]$. But then x, y do not lie both in χ , because $I_{b,\sigma} \cap \chi = \sigma$ and σ is independent. By induction x, y do not lie in one V_{b_i} either. Assume that $x \in \chi$ and $y \in V_{b_i}$ for some i . Then certainly $x \notin V_{b_i}$ and $y \notin \chi$. But the edge e must lie in one block χ' . Because of the connectivity condition for y , the guarded block (λ', χ') must lie in the subtree with root b_i , which

contradicts $x \in e$. Finally, assume that $x \in V_{b_i}$ and $y \in V_{b_j}$ for $i \neq j$ and $x, y \notin \chi$. Then x and y cannot be adjacent because of the connectivity condition. This shows that $I_{b,\sigma}$ is indeed independent.

Now assume that $I_{b,\sigma}$ is not of maximum size and let $J \subseteq V_b$ be an independent set with $|J| > |I_{b,\sigma}|$ and $J \cap \chi = \sigma$. Because J and $I_{b,\sigma}$ are fixed to σ on χ there must be a b_i such that $|J \cap V_{b_i}| > |I_{b_i,\sigma_i}|$. This contradicts the choice of σ_i . So $I_{b,\sigma}$ is indeed of maximum size.

Because each block has at most $k|V|^k$ independent sets, all computations can be done in time $k|V|^{O(k)}$. □

6.2 Parameterized Complexity

While the algorithm in the last section is nice in that it is a polynomial time algorithm for fixed k , it is somewhat unsatisfying for some decomposition techniques: If we can compute the decomposition quickly, we would ideally want to be able to compute the S -star size efficiently, too. Naturally we cannot expect a polynomial time algorithm independent of the width k for any of the decomposition techniques we consider, because computing maximum independent sets is **NP**-hard. Instead, we can hope for independent set to be at least fixed-parameter tractable with respect to k . We will show that for general hypertree width even this is unlikely, because independent set parameterized by generalized hypertree width is **W[1]**-hard. More positively, we will show that computing maximum independent sets is fixed-parameter tractable for some other decomposition techniques, in particular tree decompositions and hingetree decompositions.

Lemma 20 *Computing maximum independent sets on hypergraphs is **W[1]**-hard parameterized by generalized hypertree width.*

Proof We will show a parameterized many-one reduction from the problem p -IndependentSet defined as follows:

p -IndependentSet
Input: a graph $G, k \in \mathbb{N}$.
Parameter: k .
Problem: Decide if G has an independent set of size k .

Because p -IndependentSet is well known to be **W[1]**-hard, this suffices to establish **W[1]**-hardness of independent sets on hypergraphs parameterized by generalized hypertree width.

So let $G = (V, E)$ be a graph and let k be a positive integer. We construct a hypergraph $\mathcal{H} = (V', E')$ in the following way: For each vertex v the hypergraph \mathcal{H} has k vertices v_1, \dots, v_k . For $i = 1, \dots, k$ we have an edge $V_i := \{v_i \mid v \in V\}$ in E' . Furthermore, for each $v \in V$ we add an edge $H_v := \{v_i \mid i \in [k]\}$. Finally we add the edge sets $E_{ij} := \{v_i v_j \mid uv \in E\}$ for $i, j \in [k]$. \mathcal{H} has no other vertices or edges. The construction is illustrated in Fig. 7.




Fig. 7 We illustrate the construction for Lemma 20 by an example. A graph G on the left with the associated hypergraph \mathcal{H} for $k = 4$ on the right. To keep the illustration more transparent the edge sets E_{ij} are not shown except for $E_{1,2}$ and $E_{2,1}$

We claim that G has an independent set of size k if and only if \mathcal{H} has an independent set of size k . Indeed, if G has an independent set v^1, \dots, v^k , then v^1, \dots, v^k is easily seen to be an independent set of size k in \mathcal{H} . Now assume that \mathcal{H} has an independent set I of size k . Then for each $v \in I$ we can choose a vertex $\pi(v) \in V$ such that $v \in H_{\pi(v)}$. Furthermore for distinct $v, u \in I$ the corresponding vertices $\pi(v), \pi(u)$ have to be distinct, too, so $\pi(I) \subseteq V$ has size k . Finally, we claim that $\pi(I)$ is independent in G . Assume this is not true, then there are vertices $\pi(v), \pi(u)$ such that $\pi(v)\pi(u) \in E$. But then $vu \in E'$ by construction which is a contradiction. So, indeed G has an independent set of size k if and only if \mathcal{H} has one.

From Observation 3 we get that \mathcal{H} has generalized hypertreewidth at most k , because V_1, \dots, V_k cover V' .

Observing that the construction of \mathcal{H} from G can be done in time polynomial in $|V|$ and k completes the proof. \square

We start our fixed-parameter tractability results with an easy observation.

Proposition 3 *Given a hypergraph \mathcal{H} computing a maximum independent set in \mathcal{H} is fixed parameter tractable parameterized by the treewidth of \mathcal{H} .*

This can be seen either by applying an optimization version of Courcelle’s Theorem [11] or by straightforward dynamic programming. Interestingly, one can show the same result also for bounded hingetree width. For this decomposition technique blocks are of unbounded size which makes the dynamic programming in the proof far more involved than for treewidth.

Proposition 4 *Given a hypergraph $\mathcal{H} = (V, E)$ of hingetree width k , a maximum independent set in \mathcal{H} can be computed in time $k2^{k^2}|E|^{O(1)}$. It follows that independent set is fixed parameter tractable parameterized by hingetree width.*

Proof First observe that minimum width hingetree decompositions can be computed in polynomial time by Lemma 10, so we simply assume that a decomposition is given in the rest of the proof.

The proof has some similarity with that of Proposition 2, so we use some notation from there. For guarded block (λ, χ) we will again compute maximum independent sets containing prescribed vertices. The difference is, that we can take these prescribed sets to be of size 1: because of the hingetree condition, only one vertex of a block may be reused in any independent set in the parent. The second idea is that we can use equivalence classes of vertices in the computation of independent sets in the considered guarded blocks, which limits the number of independent sets we have to consider. We now describe the computation in detail.

Let $\mathcal{E} = (\mathcal{T}, (\lambda_t)_{t \in \mathcal{T}}, (\chi_t)_{t \in \mathcal{T}})$ be a hingetree decomposition of \mathcal{H} of width k . Let $b = (\lambda, \chi)$ be a guarded block of \mathcal{E} and let $b' = (\lambda', \chi')$ be its parent. As before, let \mathcal{T}_b be the subtree of \mathcal{T} with b as its root and $V_b := \chi(\mathcal{T}_b)$. Set $\mathcal{H}_b := (V_b, E_b)$ with $E_b := \bigcup \lambda^*$ with the union being over all guarded blocks in \mathcal{T}_b . The main idea is to iteratively compute, for all vertices $v \in \chi' \cap \chi$, a maximum independent set $J_{v,b}$ in $\mathcal{H}_b = (V_b, E_b)$ containing v . Furthermore, we also compute an independent set $J_{\emptyset,b}$ that contains no vertices of $\chi' \cap \chi$. Note that, since $\chi \subseteq \bigcup_{e \in \lambda} e$, there are no isolated vertices in χ and the size of a maximum independent set is bounded by k in each block.

For a node $b = (\lambda, \chi)$, we organize the vertices in χ into at most 2^k equivalence classes by defining v and u to be equivalent if they lie in the same subset of edges of λ . The equivalence class of v is denoted by $\mathbf{c}(v)$. For each class, a representant is fixed. We denote by \bar{v} , the representant of the equivalence class of v and by $\bar{\chi} \subseteq \chi$, the restriction of χ on these at most 2^k representants.

Example 12 Considering node b_1 in Example 5, the equivalence classes are the following: $\bar{v}_1, \bar{u}_1, \bar{v}_2 = \bar{u}_2, \bar{v}_4 = \bar{u}_3$. Hence $\bar{\chi} = \{\bar{v}_1, \bar{u}_1, \bar{v}_2, \bar{v}_4\}$. For node b_2 , the list is: $\bar{v}_4 = \bar{v}_5, \bar{v}_6, \bar{v}_8, \bar{v}_9, \bar{u}_7, \bar{u}_8$. Note that $\{v_4, v_5, v_6, v_8\} = \chi_b \cap \chi_{b_2}$, although these vertices are spread in three distinct equivalence classes, by definition of hinge tree decomposition, they all belong to some hyperedge and no two of them can appear simultaneously in one independent set. \square

Let first b be a leaf. We first compute independent sets on $\bar{\chi}$. Observe that the independent sets are invariant under the choice of representants. For each equivalence class $\mathbf{c}(v)$, we compute $J_{\bar{v},b} \subseteq \bar{\chi}$ as a maximum independent set containing \bar{v} . Computing the classes and a choice of maximum independent sets containing each \bar{v} can be done in time $k2^{k^2}$ because independent sets cannot be bigger than k . Clearly, $J_{v,b}$, a maximum independent set containing v , can be easily computed from the set $J_{\bar{v},b}$. Thus, one can compute all the $J_{v,b}$ in time $k2^{k^2}n$. The computation of $J_{\emptyset,b}$ can be done on representants, too, by simply excluding the vertices from $\chi' \cap \chi$.

Let b now be an inner vertex and b_1, b_2, \dots, b_m be its children with $b_i = (\lambda_i, \chi_i)$, $i \in [m]$. We again consider equivalence classes on χ . Fix $v \in \chi$ and compute the list $L_{\bar{v},b}$ of all independent sets $\sigma \subseteq \bar{\chi}$ containing \bar{v} . Fix now $\sigma \in L_{\bar{v},b}$. We first compute a set $J_{v,b}^\sigma$ as a maximum independent set of \mathcal{H}_b containing v and whose vertices in χ have the representants σ . We will distinguish for a given vertex $\bar{u} \in \sigma$ if it is the

representant of a vertex belonging to the block of some (or several) children of b or if it represents vertices of $\chi \setminus (\bigcup_{i=1}^m \chi_i)$ only. Therefore we partition σ into σ', σ'' accordingly:

- $\sigma := \sigma' \cup \sigma''$
- $\sigma' := \bar{\chi} \cap \{\bar{u} \mid u \in \bigcup_{i=1}^m \chi_i\}$.
- $\sigma'' := \bar{\chi} \setminus \{\bar{u} \mid u \in \bigcup_{i=1}^m \chi_i\}$

Set $\sigma' := \{\bar{u}_1, \dots, \bar{u}_h\}$ with $h \leq m$. Let us examine the consequences of \mathcal{T} being a hingetree decomposition. We have that, for all $i \in [m]$, there exists $e_i \in \lambda$, such that $\chi \cap \chi_i \subseteq e_i$. Thus, since σ is an independent set in $\bar{\chi} \subseteq \chi$, at most one vertex in σ' is a representant of a vertex in χ_i . Thus

$$\forall u \neq u' \in \sigma : \chi_i \cap \mathbf{c}(u) = \emptyset \vee \chi_i \cap \mathbf{c}(u') = \emptyset. \tag{1}$$

We denote by $S_i = \{j \mid \mathbf{c}(u_i) \cap \chi_j \neq \emptyset\}$ and by $S = [m] \setminus \bigcup S_i$. By (1) the sets S_1, \dots, S_h, S form a partition of $[m]$. To construct $J_{v,b}^\sigma$, we now determine for each $i \leq h$, which vertex u of $\mathbf{c}(u_i)$ can contribute the most, by taking the union of all the maximum independent sets J_{u,b_j} , $j \in S_i$, that u induces at the level of the children of b .

For each fixed $u \in \mathbf{c}(u_i)$, let

$$I_{i,u} = \{u\} \cup \bigcup_{j \in S_i} J_{u,b_j},$$

where we set $J_{u,b_j} := J_{\emptyset,b_j}$ if $u \notin \chi_j$. Let then $I_i = I_{i,u}$ for some $u \in \mathbf{c}(u_i)$ for which the size of $I_{i,u}$ is maximal.

The set $J_{v,b}^\sigma$ is now obtained as follows depending on whether $\bar{v} \in \sigma''$ or $\bar{v} \in \sigma'$. If $\bar{v} \in \sigma''$, we claim that $J_{v,b}^\sigma$ can be chosen as

$$J_{v,b}^\sigma := \{v\} \cup (\sigma'' \setminus \{\bar{v}\}) \cup \bigcup_{i=1}^h I_i \cup \bigcup_{i \in S} J_{\emptyset,b_i}.$$

If $\bar{v} \in \sigma'$, say $\bar{v} = u_1$, we claim that $J_{v,b}^\sigma$ can be chosen as

$$J_{v,b}^\sigma := \sigma'' \cup \bigcup_{j \in S_1: v \in \chi_j} J_{v,b_j} \cup \bigcup_{j \in S_1: v \notin \chi_j} J_{\emptyset,b_j} \cup \bigcup_{i=2}^h I_i \cup \bigcup_{i \in S} J_{\emptyset,b_i}.$$

The set $J_{v,b}$ is taken as one of the sets $J_{v,b}^\sigma$ of maximal size for a $\sigma \in L_{v,b}$. To compute $J_{\emptyset,b}$, the arguments are similar.

We first show that all $J_{v,b}$ are indeed independent sets in \mathcal{H}_b . Clearly, it is enough to prove this for any $J_{v,b}^\sigma$. There will be no reason to distinguish whether $\bar{v} \in \sigma''$ or $\bar{v} \in \sigma'$, because our arguments will apply to all $J_{v,b}^\sigma$ independent of the choice of a distinguished element v . We will make extensive use of the two following facts.

- Let $j, j' \in [m]$ and $I \subseteq V_{b_j}, I' \subseteq V_{b_{j'}}$ independent sets of \mathcal{H}_{b_j} and $\mathcal{H}_{b_{j'}}$ respectively. By the connectivity condition for tree decomposition we have

$$I \cap I' \subseteq \chi_j \cap \chi_{j'} \cap \chi.$$

This permits to investigate the intersection of two independent sets I, I' by looking at their restriction on χ .

- Let now $I \subseteq V_{b_j}$ be an independent set of \mathcal{H}_{b_j} . Then, I remains an independent set in \mathcal{H}_b . Indeed, suppose there is a $e \in E_b \setminus E_{b_j}$ containing two vertices $y_1, y_2 \in I$. Since all edges must belong to a guard, there exists a node $b^* = (\lambda^*, \chi^*)$ such that $e \in \lambda^*$. Then, since in a hingetree decomposition we have $\chi^* = \bigcup \lambda^*$, then $\{y_1, y_2\} \subseteq e \subseteq \chi^*$. But then, by the connectivity condition it follows that $\{y_1, y_2\} \subseteq \chi$. Hence, by the intersection property of hingetree decomposition, there exists $e_j \in \chi_j$ such that

$$\{y_1, y_2\} \subseteq \chi \cap \chi_j \cap e_j$$

which implies that y_1 and y_2 are adjacent in \mathcal{H}_{b_j} . Contradiction.

We now start the proof that $J_{v,b}^\sigma$ is independent incrementally. Let $i \in [h]$, $u \in \mathbf{c}(u_i)$ and $j \in S_i$ and consider the set $I := J_{u,b_j}$. By induction, the set I is independent in \mathcal{H}_{b_j} . By the hingetree condition, there exists $e_j \in \lambda_j$ such that $\chi \cap \chi_j \subseteq e_j$. By the connectivity condition, this implies $\chi \cap I \subseteq e_j$. Then, since I is an independent set, no two vertices of χ can belong to I , i.e., $|\chi \cap I| \leq 1$. The connectivity condition also implies that, for $j' \neq j$, $V_{b_{j'}} \cap I \subseteq \chi \cap \chi_j$, hence $|V_{b_{j'}} \cap I| \leq 1$ and I is an independent set of \mathcal{H}_b . Finally, the set $I_i = \bigcup_{j \in S_i} J_{u,b_j}$ is also an independent set of \mathcal{H}_b , since for any distinct $j, j' \in S_i$:

$$J_{u,b_j} \cap J_{u,b_{j'}} \subseteq \chi_j \cap \chi_{j'} \cap \chi \subseteq e_j.$$

Hence $J_{u,b_j} \cap J_{u,b_{j'}}$ contains at most one vertex (which is in χ and could then only be u).

Let now $i, i' \in [m]$ be distinct. By the arguments above, I_i (resp. $I_{i'}$) contains at most one element u (resp. u') such that $u \in \mathbf{c}(u_i)$ (resp. $u' \in \mathbf{c}(u_{i'})$). By Eq. 1, we have that the two classes are distinct and that $u_i \neq u_{i'}$. But $u_i, u_{i'} \in \sigma$ and σ is independent in χ . Hence, $u_i, u_{i'}$ cannot be adjacent in \mathcal{H}_b . Consequently,

$$\bigcup_{i=1}^h I_i$$

is an independent set in \mathcal{H}_b .

Let $j \in S$. J_{\emptyset,b_j} is independent in \mathcal{H}_{b_j} and $J_{\emptyset,b_j} \subseteq V_{b_j} \setminus \chi$. Hence, J_{\emptyset,b_j} is independent in \mathcal{H}_b . This also implies that, given $j' \in [m]$ distinct from j , $J_{\emptyset,b_j} \cap V_{b_{j'}} = \emptyset$. Thus,

$$\bigcup_{i=1}^h I_i \cup \bigcup_{i \in S} J_{\emptyset,b_i}.$$

is independent in \mathcal{H}_b .

Finally, by construction, for all $i \in [h]$, $I_i \cap \chi = \{u\}$ with $\bar{u} = \bar{u}_i \in \sigma'$. Also $\sigma = \sigma' \cup \sigma''$ is independent in χ hence in \mathcal{H}_b . No vertices $y_1 \in I_i$ and $y_2 \in \sigma''$ can be adjacent because, again, this would imply that $\{y_1, y_2\} \subseteq \chi$ and contradict the fact that \bar{y}_1, \bar{y}_2 are independent in σ . Thus $J_{v,b}^\sigma$ is independent.

We now prove that $J_{v,b}$ is of maximum size. Observe that it suffices to show this again for each $J_{v,b}^\sigma$. Each maximum independent set J of \mathcal{H}_b that contains v and

whose vertices in χ have exactly the representants σ can be expressed as $\tau \cup J_1 \cup J_2 \cup \dots \cup J_m$. Here $\tau \subseteq \chi$ is an independent set of b containing v and whose representants are σ . Furthermore, J_i is an independent set of \mathcal{H}_b that contains only vertices of V_{b_i} . The set J_i may only contain one vertex u_i from $\chi \cap \chi_i$. But then exchanging J_i for J_{u_i, b_i} may only increase the size of the independent set, so we can assume that I has the form $\tau \cup J_{u_1, b_1} \cup J_{u_2, b_2} \cup \dots \cup J_{u_m, b_m}$ where u_i may also stand for \emptyset .

Assume now that $J_{v, b}^\sigma$ is not a maximum independent set, i.e., there is an independent set J containing v whose vertices in χ have the representants σ and J is bigger than $J_{v, b}^\sigma$. Then one of four following things must happen:

- There is an i such that $v \in \chi_i$ and $J \cap V_{b_i}$ is bigger than J_{v, b_i} . But this case cannot occur by induction.
- $v = u_1$ and there is a $j \in S_1$ such that $v \notin \chi_j$ and $|J \cap V_{b_j}| > |J_{\emptyset, b_j}|$. By induction we know that J_{\emptyset, b_j} is optimal under all independent sets of \mathcal{H}_{b_j} not containing any vertex of $\chi_j \cap \chi$, so there must be a vertex $u \in J \cap \chi \cap \chi_j$. Since J is independent, v and u share no edge in λ and then $\bar{v} \neq \bar{u}$. Since $j \in S_1$, it holds that $\mathbf{c}(v) \cap \chi_j \neq \emptyset$ and by Eq. 1, $\mathbf{c}(u) \cap \chi_j = \emptyset$. Contradiction.
- There is an $i \in S$ such that $J \cap V_{b_i}$ is bigger than J_{\emptyset, b_i} . But from $i \in S$ it follows by definition that $\chi \cap \chi_i \cap J = \emptyset$, so this case can not occur by induction, either.
- There is an $i \in [h]$ such that $|J \cap (\bigcup_{j \in S_i} V_j)| > |I_i|$. We claim that $(\bigcup_{j \in S_i} \chi_j) \cap \chi \cap J$ contains only one vertex. Assume there are two such vertices x and y . By definition, $\bar{x}, \bar{y} \in \bar{\tau}$. Since J is independent, \bar{x} and \bar{y} are not adjacent in $\bar{\chi}$ and $\bar{x} \neq \bar{y}$. At least one of these, say y , must be in $\mathbf{c}(u_i)$, because $\bar{u}_i \in \bar{\tau}$ by definition. Let $x \in V_{j'}$ with $j' \in S_i$, then there is a vertex $w \in \mathbf{c}(u_i) = \mathbf{c}(y)$ in $\chi_{j'} \cap \chi \subseteq e_j$ by definition of S_i . But then \bar{x} and \bar{y} are adjacent in $\bar{\chi}$ which is a contradiction.

So there is exactly one vertex u in $(\bigcup_{j \in S_i} \chi_j) \cap \chi \cap J$. But then $|J \cap (\bigcup_{j \in S_i} V_j)| > |I_{i, u}|$. Thus either there must be a $j \in S_i$ with $u \in V_j$ such that $|J \cap V_j| > |J_{u, b_j}|$ or there must be a $j \in S_i$ with $u \notin V_j$ such that $|J \cap V_j| > |J_{\emptyset, b_j}|$. The former clearly contradicts the optimality of J_{u, b_j} , while the latter leads to a contradiction completely analogously to the second item above.

Because only $k2^{k^2}n^2$ sets have to be considered for each guarded block, this results in an algorithm with runtime $k2^{k^2}|V|^{O(1)}$.

6.3 Approximation

We have seen that computing maximum independent sets of hypergraphs with decompositions of width k can be done in polynomial time for fixed width k and that for some decompositions it is even fixed parameter tractable with respect to k . Still, the exponential influence of k is troubling. In this section we will show that we can get rid of it if we are willing to sacrifice the optimality of the solution. We give a polynomial time k -approximation algorithm for computing maximum independent sets of graphs with generalized hypertree width k assuming that a decomposition is given. We start by formulating a lemma.

Lemma 21 *Let $\mathcal{H} = (V, E)$ be a hypergraph with a generalized hypertree decomposition $\mathcal{E} = (\mathcal{T}, (\lambda_t)_{t \in T}, (\chi_t)_{t \in T})$ of width k . Let $\mathcal{H}' = (V, E')$ where $E' := \{\chi_t \mid t \in T\}$. Let ℓ be the size of a maximum independent set in \mathcal{H} and let ℓ' be the size of a maximum independent set in \mathcal{H}' . Then*

$$\frac{\ell}{k} \leq \ell' \leq \ell.$$

Before we prove Lemma 21 we will show how to get the approximation algorithm from it.

Observation 16 *Every independent set of \mathcal{H}' is also an independent set of \mathcal{H} .*

Proof Each pair of independent vertices x, y in \mathcal{H}' is by definition in different blocks χ_t in \mathcal{H} . For each edge $e \in E$ there must (by definition of generalized hypertree decompositions) be a block χ such that $e \subseteq \chi$. Thus no edge $e \in E$ can contain both x and y , so x and y are independent in \mathcal{H} as well. □

Corollary 3 *There is a polynomial time algorithm that, given a hypergraph \mathcal{H} and a generalized hypertree decomposition of width k , computes an independent set I of \mathcal{H} such that $|I| \geq \frac{\ell}{k}$ where ℓ is the size of a maximum independent set of \mathcal{H} .*

Proof Observe that \mathcal{H}' is acyclic by Lemma 7. By Lemma 14, we compute in polynomial time a maximum independent set I of \mathcal{H}' whose size by Lemma 21 only differs by a factor $\frac{1}{k}$ from ℓ . By Observation 16, we know that I is also an independent set of \mathcal{H} . □

Proof of Lemma 21 The second inequality follows directly from Observation 16.

For the first inequality consider a maximum independent set I of \mathcal{H} . Observe that a set I' is an independent set of \mathcal{H}' if and only if it is an independent set of its primal graph \mathcal{H}'_P , so it suffices to show the same result for \mathcal{H}'_P .

Claim 6 *The graph $\mathcal{H}'_P[I]$ has treewidth at most $k - 1$.* □

Proof First observe that vertices v that appear in no edge $e \in E$ change neither the treewidth nor the generalized hypertree width of a graph or hypergraph. Thus we assume that every vertex $v \in V$ is in at least one edge $e \in E$.

We construct a tree decomposition $(\mathcal{T}', (\chi'_t)_{t \in T'})$ of $\mathcal{H}'_P[I]$ from \mathcal{E} as follows: We set $\mathcal{T}' := \mathcal{T}[T']$ where $T' := \{t \in T \mid \chi_t \cap I \neq \emptyset\}$. Furthermore, $\chi'_t := \chi_t \cap I$ for $t \in T'$. For every $v \in I$ there is an edge $e \in E$ and $t \in T$ such that $v \in e \subseteq \chi_t$ and thus $v \in \chi'_t$. It follows that the bags χ'_t cover I . Moreover, the connectivity condition for I is satisfied, because it is satisfied for \mathcal{E} . Finally, for each edge uv in $\mathcal{H}'_P[I]$ there is a guarded block (λ_t, χ_t) such that $u, v \in \chi_t$ and thus $u, v \in \chi'_t$. Hence, $(\mathcal{T}', (\chi'_t)_{t \in T'})$ is indeed a tree decomposition.

Thus we only have to show $|\chi'_t| \leq k$. To see this, observe that for each t the bag $\chi'_t \subseteq \chi_t$ is covered by λ_t . But the vertices in $\chi'_t \subseteq I$ are independent in \mathcal{H} and thus each $e \in \lambda_t$ can contain only a single vertex from χ'_t . Thus $|\chi'_t| \leq |\lambda_t| \leq k$. □

Claim 7 *The graph $\mathcal{H}'_p[I]$ has an independent set I' of size at least $\frac{|I|}{k}$.*

Proof From Claim 6 it follows with Lemma 3 that $\mathcal{H}'[I]$ and all of its induced subgraphs have a vertex of degree at most k . We construct I' iteratively by choosing a vertex of minimum degree and deleting it and its neighbors from the graph. In each round we delete at most k vertices, so we can choose a vertex in at least $\frac{|I|}{k}$ rounds. Obviously the chosen vertices are independent. \square

Every independent set of $\mathcal{H}'_p[I]$ is also an independent set of \mathcal{H}'_p which completes the proof of Lemma 21.

7 Fractional Hypertree Width

In this section we extend the main results of the paper to fractional hypertree width, which is the most general notion known so far that leads to tractable CQ [23]. In particular it is strictly more general than generalized hypertree width.

Definition 16 Let $\mathcal{H} = (V, E)$ be a hypergraph. A *fractional edge cover* of a vertex set $S \subseteq V$ is a mapping $\psi : E \rightarrow [0, 1]$ such that for every $v \in S$ we have $\sum_{e \in E: v \in e} \psi(e) \geq 1$. The weight of ψ is $\sum_{e \in E} \psi(e)$. The fractional edge cover number of S , denoted by $\rho_{\mathcal{H}}^*(S)$ is the minimum weight taken over all fractional edge covers of S .

A *fractional hypertree decomposition* of \mathcal{H} is a triple $(\mathcal{T}, (\chi_t)_{t \in T}, (\psi_t)_{t \in T})$ where $\mathcal{T} = (T, F)$ is a tree, and $\chi_t \subseteq V$ and ψ_t is a fractional edge cover of χ_t for every $t \in T$ satisfying the following properties:

1. For every $v \in V$ the set $\{t \in T \mid v \in \chi_t\}$ induces a subtree of \mathcal{T} .
2. For every $e \in E$ there is a $t \in T$ such that $e \subseteq \chi_t$.

The width of a fractional hypertree decomposition $(\mathcal{T}, (\chi_t)_{t \in T}, (\psi_t)_{t \in T})$ is defined as $\max_{t \in T} (\rho_{\mathcal{H}}^*(\chi_t))$. The *fractional hypertree width* of \mathcal{H} is the minimum width over all fractional hypertree decompositions of \mathcal{H} .

Together with the previous results of this paper, the two following Theorems will serve as key ingredients to prove the main results of this section.

Theorem 17 ([23]) *The solutions of a CQ-instance Φ with hypergraph $\mathcal{H} = (V, E)$ can be enumerated in time $\|\Phi\|^{\rho_{\mathcal{H}}^*(V)+O(1)}$.*

Theorem 18 ([28]) *Given a hypergraph \mathcal{H} and a rational number $w \geq 1$, it is possible in time $\|\mathcal{H}\|^{O(w^3)}$ to either*

- compute a fractional hypertree decomposition of \mathcal{H} with width at most $7w^3 + 31w + 7$, or
- correctly conclude that $\text{fhw}(\mathcal{H}) \geq w$.

7.1 Tractable Counting

We start of with the quantifier free case which we will use as a building block for the more general result later.

Lemma 22 *The solutions of a quantifier free CQ-instance Φ with hypergraph \mathcal{H} can be counted in time $\|\Phi\|^{\text{fhw}(\mathcal{H})^{O(1)}}$.*

Proof With Theorem 18 we can compute a fractional hypertree decomposition $(\mathcal{T}, (\chi_t)_{t \in T}, (\psi_t)_{t \in T})$ of width at most $k := O(\text{fhw}(\mathcal{H})^3)$. For each bag χ_t we can with Theorem 17 in time $\|\Phi\|^k$ compute all solutions to the CQ-instance $\Phi[\chi_t]$ that is induced by the variables in χ_t . Let these solutions form a new relation \mathcal{R}_t belonging to a new atom ϕ'_t . Then $\bigwedge_{t \in T} \phi'_t$ gives a solution equivalent, acyclic, quantifier free #CQ instance of size $\|\Phi\|^{O(k)}$. Now we can count the solutions with the algorithm [31] or [15]. □

We can now prove a version of Corollary 1 for fractional hypertree width.

Theorem 19 *There is an algorithm that given a #CQ-instance Φ of quantified star-size ℓ and fractional hypertree width k counts the solutions of Φ in time $\|\Phi\|^{p(k, \ell)}$ for a polynomial p .*

Proof This is a minor modification of the proof of Lemma 13. Let $\mathcal{H} = (V, E)$ be the hypergraph of Φ . Because of Theorem 18 we may assume that we have a fractional hypertree decomposition $\mathcal{E} := (\mathcal{T}, (\chi_t)_{t \in T}, (\psi_t)_{t \in T})$ of width $k' := k^{O(1)}$ of \mathcal{H} . For each edge $e \in E$ we let $\varphi(e)$ be the atom of Φ that induces e .

Let V_1, \dots, V_m be the vertex sets of the components of $\mathcal{H}[V \setminus S]$ and let V'_1, \dots, V'_m be the vertex sets of the S -components of \mathcal{H} . Clearly, $V_i \subseteq V'_i$ and $V'_i - V_i = V'_i \cap S =: S_i$. Let Φ_i be the restriction of Φ to the variables in V'_i and let \mathcal{E}_i be the corresponding fractional hypertree decomposition. Then \mathcal{E}_i has a tree \mathcal{T}_i that is a subtree of \mathcal{T} .

For each Φ_i we construct a new #CQ-instance Φ'_i by computing for each bag $t \in T$ an atom ϕ_t in the variables χ_t that contains the solutions of $\Phi_i[\chi_t]$ that is induced by the variables of χ_t . The decomposition \mathcal{E} has width at most k' so this can be done in time $\|\Phi\|^{O(k')}$ by Theorem 17. Obviously Φ_i and Φ'_i are solution equivalent and Φ'_i is acyclic. Furthermore, Φ'_i has only one single S_i -component, because all the vertices in V_i are connected in Φ and thus also in Φ'_i . Let \mathcal{H}_i be the hypergraph of Φ'_i , then \mathcal{H}_i has S_i -star size at most ℓ . Thus the vertices in S_i can be covered by at most ℓ edges in \mathcal{H}_i by Lemma 14.

Now we construct a CQ-instance $(\mathcal{A}''_i, \phi''_i)$ such that ϕ''_i is an atomic formula in the variables S_i exactly as in the proof of Lemma 13.

We now eliminate all quantified variables in Φ . To do so we add the atom ϕ''_i for $i \in [m]$ and delete all atoms that contain any quantified variable, i.e., we delete all

Φ'_i . Call the resulting CQ-instance Φ'' . Because $(\mathcal{A}''_i, \phi''_i)$ is solution equivalent to Φ'_i , we have that Φ and Φ'' are solution equivalent, too.

We now construct a fractional hypertree decomposition of Φ'' by doing the following: we set $\chi'_t = (\chi_t \setminus \bigcup_{i \in I_t} V_i) \cup \bigcup_{i \in I_t} S_i$ for each bag χ_t where $I_t := \{i \mid \chi_t \cap V_i \neq \emptyset\}$. For each bag χ_t we construct a fractional edge cover ψ'_t of χ'_t by setting $\psi'_t(e) := \psi_t(e)$ for all old edges and setting $\psi_t(S_i) = 1$ for $i \in I_t$ where S_i corresponds to the newly added constraint ϕ_i with $\chi_t \cap V_i \neq \emptyset$. The result is indeed a fractional edge cover, because each variable not in any S_i is still covered as before and the variables in S_i are covered by definition of ψ_t . Furthermore, we claim that the width of the cover is at most k' . Indeed, for each $i \in I$ we had for each $v \in V_i$, $\sum_{e \in E: v \in e} \psi(e) \geq 1$. None of these edges appears in the new decomposition anymore. Thus adding the edge S_i with weight 1 does not increase the total weight of the cover. It is now easy to see that doing this construction for all χ_t leads to a fractional hypertree decomposition of Φ' of width at most k' .

Applying Lemma 22 concludes the proof. □

7.2 Computing Independent Sets

Also S -star size or equivalently independent sets of bounded fractional hypertree width hypergraphs can be computed efficiently.

Lemma 23 *There is an algorithm that given a hypergraph $\mathcal{H} = (V, E)$ of fractional hypertree width at most k computes a maximum independent set of \mathcal{H} in time $|\mathcal{H}|^{k^{O(1)}}$.*

In the proof of Lemma 23 we will use the following lemma.

Lemma 24 *The independent sets of a hypergraph $\mathcal{H} = (V, E)$ can be enumerated in time $|\mathcal{H}|^{O(\rho_{\mathcal{H}}^*(V))}$.*

Proof Let $\mathcal{H} = (V, E)$. We construct a quantifier free CQ-instance $\Phi = (\mathcal{A}, \phi)$ with the hypergraph \mathcal{H} . Let V be the variables of Φ , $\{0, 1\}$ the domain and add an atom ϕ_e with relation symbol \mathcal{R}_e and scope e for each $e \in E$. The relation $\mathcal{R}_e^{\mathcal{A}}$ contains all tuples that contain at most one 1 entry. Finally, $\phi := \bigwedge_{e \in E} \phi_e$.

Clearly, Φ has indeed the hypergraph \mathcal{H} . Furthermore the solutions of Φ are exactly the characteristic vectors of independent sets of \mathcal{H} . Thus we can enumerate all independent sets of \mathcal{H} in time $|\mathcal{H}|^{O(\rho_{\mathcal{H}}^*(V))}$ with Theorem 17. □

Proof of Lemma 23 (Sketch) We proceed by dynamic programming along a fractional hypertree decomposition.

In a first step we compute a fractional hypertree decomposition $(\mathcal{T}, (\chi_t)_{t \in \mathcal{T}}, (\psi_t)_{t \in \mathcal{T}})$ of width $k' = k^{O(1)}$ of \mathcal{H} with Theorem 18. For each bag χ_t we then compute all independent sets of $\mathcal{H}[\chi_t]$ by Lemma 24; call this set I_t .

By dynamic programming similar to the proof of Lemma 2 we then compute a maximum independent set of \mathcal{H} . □

8 Conclusion

The results of this paper give a clear picture of tractability for counting solutions of conjunctive queries for structural classes that are known to have tractable decision problems. Essentially counting is tractable if and only if these classes are combined with quantified star size. So to find more general structural classes that allow tractable counting, progress for the corresponding decision question appears to be necessary.

Note that our characterizations of tractable classes of queries rely only on the underlying hypergraphs of the queries and do not use any other information. For the case of bounded arity, it is known that analyzing classes of queries directly instead of their hypergraphs allows to completely characterize the classes of queries for which the decision problem CQ is tractable [12, 22]. In particular, this includes tractable classes whose tractability is not witnessed by the hypergraph perspective alone. A refinement of our results in this style in the case of bounded arity can be found in [29]. These results will also appear in an upcoming paper by Hubie Chen and the second author.

Another way of generalizing the results of this paper would be extending the logic that the queries can be formulated in. Just recently Chen and Dalmau [9] have characterized the tractable classes of bounded arity QCSP which is essentially a version of CQ in which also universal quantifiers are allowed. They do this by introducing a new width measure for first order $\{\forall, \exists, \wedge\}$ -formulas. We conjecture that their width measure also characterizes the tractable cases for #QCSP, i.e., tractable decision and counting coincide here. It would be interesting to see how far this can be pushed for the case of unbounded arity.

Another extension of conjunctive queries appears in a recent paper by Chen [7] where he considers existential formulas that may use conjunction and disjunction. This is particularly interesting, because it corresponds to the classical select-project-join queries with union that play an important role in database theory (see e.g., the textbook [1]). One may wonder if techniques used in this paper may help to understand the complexity of this class of queries better.

Acknowledgments The authors are grateful for the very helpful feedback on this paper they got from the reviewers of the conference version. The results of this paper are a part of the PhD thesis of the second author [29]. During the writeup process of this thesis, Peter Bürgisser gave valuable feedback that helped immensely to improve the presentation of the thesis and thus also of this paper. The authors are very thankful for this.

References

1. Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison-Wesley (1995)
2. Adler, I., Gottlob, G., Grohe, G.: Hypertree width and related hypergraph invariants. *Eur. J. Comb.* **28**(8), 2167–2181 (2007)
3. Bagan, G., Durand, A., Grandjean, G.: On acyclic conjunctive queries and constant delay enumeration. In: *CSL'07, 16th Annual Conference of the EACSL, LNCS*, vol. 4646, pp. 208–222. Springer (2007)
4. Bodlaender, H.: A linear time algorithm for finding tree-decompositions of small treewidth. In: *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pp. 226–234. ACM (1993)

5. Bodlaender, H.: A partial k -arboretum of graphs with bounded treewidth. *Theor. Comput. Sci.* **209** (1-2), 1–45 (1998)
6. Bulatov, A., Dalmau, V., Grohe, M., Marx, D.: Enumerating homomorphisms. *J. Comput. Syst. Sci.* **78**(2), 638–650 (2012)
7. Chen, H.: On the Complexity of Existential Positive Queries. ArXiv e-prints (2012)
8. Chen, H., Dalmau, V.: Beyond hypertree width: decomposition methods without decompositions. In: 11th international conference principles and practice of constraint programming, CP '05, pp. 167–181 (2005)
9. Chen, H., Dalmau, V.: Decomposing quantified conjunctive (or disjunctive) formulas. *LICS* (2012)
10. Cohen, D., Jeavons, P., Gyssens, M.: A unified theory of structural tractability for constraint satisfaction problems. *J. Comput. Syst. Sci.* **74**(5), 721–743 (2008)
11. Courcelle, B.: Graph Rewriting: An Algebraic and Logic Approach. In: Handbook of theoretical computer science, volume B: Formal models and semantics (B), pp. 193–242 (1990)
12. Dalmau, V., Jonsson, P.: The complexity of counting homomorphisms seen from the other side. *Theor. Comput. Sci.* **329**(1-3), 315–323 (2004)
13. Diestel, R.: *Graph Theory*. 2005. Grad. Texts in Math (2005)
14. Downey, R.G., Fellows, M.R.: *Parameterized complexity*, Vol. 3. Springer, New York (1999)
15. Durand, A., Mengel, S.: The Complexity of Weighted Counting for Acyclic Conjunctive Queries. Arxiv preprint arXiv:1110.4201 (2011)
16. Flum, J., Frick, M., Grohe, M.: Query evaluation via tree-decompositions. *J. ACM* **49**(6), 716–752 (2002)
17. Flum, J., Grohe, M.: *Parameterized Complexity Theory*. Springer-Verlag, New York Inc (2006)
18. Gottlob, G., Leone, N., Scarcello, F.: A comparison of structural CSP decomposition methods. *Artif. Intell.* **124**(2), 243–282 (2000)
19. Gottlob, G., Leone, N., Scarcello, F.: Hypertree decompositions and tractable queries. *J. Comput. Syst. Sci.* **64**(3), 579–627 (2002)
20. Gottlob, G., Miklós, Z., Schwentick, T.: Generalized Hypertree Decompositions: NP-hardness and tractable variants. *J. ACM* **56**(6) (2009)
21. Greco, G., Scarcello, F.: Structural Tractability of Enumerating CSP Solutions. In: Proceedings of the 16th International Conference on Principles and Practice of Constraint Programming, CP '10, pp. 236–251 (2010)
22. Grohe, M.: The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM* **54**(1) (2007)
23. Grohe, M., Marx, D.: Constraint Solving via Fractional Edge Covers. In: 17th Annual ACM-SIAM Symposium on Discrete Algorithm, SODA '06, pp. 289–298. ACM, New York (2006)
24. Grohe, M., Schwentick, T., Segoufin, L.: When is the evaluation of conjunctive queries tractable? In: Proceedings of the 33rd Annual ACM Symposium on Theory of Computing, pp. 657–666. ACM (2001)
25. Gyssens, M., Jeavons, P., Cohen, D.A.: Decomposing constraint satisfaction problems using database techniques. *Artif. Intell.* **66**(1), 57–89 (1994)
26. Kolaitis, P.G., Vardi, M.Y.: Conjunctive-Query Containment and Constraint Satisfaction. *J. Comput. Syst. Sci.* **61**(2), 302–332 (2000)
27. Libkin, L.: *Elements of Finite Model Theory*, EATCS Series. Springer (2004)
28. Marx, D.: Approximating fractional hypertree width. *ACM Trans. Algorith.* **6**(2), 29:1–29:17 (2010)
29. Mengel, S.: *Conjunctive Queries, Arithmetic Circuits and Counting Complexity*. Ph.D. Thesis, University of Paderborn (2013)
30. Miklós, Z.: *Understanding Tractable Decompositions for Constraint Satisfaction*. Ph.D. Thesis, University of Oxford (2008)
31. Pichler, R., Skritek, A.: Tractable Counting of the Answers to Conjunctive Queries. AMW (2011)
32. Robertson, N., Seymour, P.: Graph minors II: algorithmic aspects of tree-width. *J. Algorith.* **7**(3) (1986)
33. Yannakakis, M.: Algorithms for Acyclic Database Schemes, pp. 82–94. VLDB (1981)