

Constant Delay Enumeration with FPT-Preprocessing for Conjunctive Queries of Bounded Submodular Width

Christoph Berkholz

Humboldt-Universität zu Berlin, Unter den Linden 6, D-10099 Berlin, Germany
berkholz@informatik.hu-berlin.de

Nicole Schweikardt

Humboldt-Universität zu Berlin, Unter den Linden 6, D-10099 Berlin, Germany
schweikn@informatik.hu-berlin.de

Abstract

Marx (STOC 2010, J. ACM 2013) introduced the notion of submodular width of a conjunctive query (CQ) and showed that for any class Φ of Boolean CQs of bounded submodular width, the model-checking problem for Φ on the class of all finite structures is fixed-parameter tractable (FPT). Note that for non-Boolean queries, the size of the query result may be far too large to be computed entirely within FPT time. We investigate the free-connex variant of submodular width and generalise Marx's result to non-Boolean queries as follows: For every class Φ of CQs of bounded free-connex submodular width, within FPT-preprocessing time we can build a data structure that allows to enumerate, without repetition and with constant delay, all tuples of the query result. Our proof builds upon Marx's splitting routine to decompose the query result into a union of results; but we have to tackle the additional technical difficulty to ensure that these can be enumerated efficiently.

2012 ACM Subject Classification Theory of computation \rightarrow Complexity theory and logic; Theory of computation \rightarrow Finite Model Theory; Theory of computation \rightarrow Fixed parameter tractability; Theory of computation \rightarrow Database theory; Theory of computation \rightarrow Database query processing and optimization (theory)

Keywords and phrases conjunctive queries, constant delay enumeration, hypertree decompositions, submodular width, fixed-parameter tractability

Digital Object Identifier 10.4230/LIPIcs.MFCS.2019.58

Funding Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SCHW 837/5-1.

1 Introduction

In the past decade, starting with Durand and Grandjean [19], the fields of logic in computer science and database theory have seen a large number of contributions that deal with the efficient enumeration of query results. In this scenario, the objective is as follows: given a finite relational structure (i.e., a database) and a logical formula (i.e., a query), after a short preprocessing phase, the query results shall be generated one by one, without repetition, with guarantees on the maximum delay time between the output of two tuples. In this vein, the best that one can hope for is *constant* delay (i.e., the delay may depend on the size of the query but not on that of the input structure) and *linear* preprocessing time (i.e., time $f(\varphi) \cdot O(N)$ where N is the size of a reasonable representation of the input structure, φ is the query, and $f(\varphi)$ is a number only depending on the query but not on the input structure). Constant delay enumeration has also been adopted as a central concept in *factorised databases* that gained recent attention [37, 36].



© Christoph Berkholz and Nicole Schweikardt;
licensed under Creative Commons License CC-BY

44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019).

Editors: Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen; Article No. 58; pp. 58:1–58:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Quite a number of query evaluation problems are known to admit constant delay algorithms preceded by linear or pseudo-linear time preprocessing. This is the case for all first-order queries, provided that they are evaluated over classes of structures of bounded degree [19, 27, 13, 30], low degree [20], bounded expansion [28], locally bounded expansion [41], and on classes that are nowhere dense [39]. Also different data models have been investigated, including tree-like data and document spanners [8, 29, 7]. Recently, also the *dynamic* setting, where a fixed query has to be evaluated repeatedly against a database that is constantly updated, has received quite some attention [31, 13, 12, 25, 14, 5, 35, 34, 6].

This paper deals with the classical, *static* setting without database updates. We focus on evaluating conjunctive queries (CQs, i.e., primitive-positive formulas) on arbitrary relational structures.¹ In the following, *FPT-preprocessing* (resp., *FPL-preprocessing*) means preprocessing that takes time $f(\varphi) \cdot N^{O(1)}$ (resp., $f(\varphi) \cdot O(N)$), and *constant delay* means delay $f(\varphi)$, where f is a computable function, φ is the query, and N is the size of the input structure.

Bagan et al. [10] showed that every free-connex acyclic CQ allows constant delay enumeration after FPL-preprocessing. More refined results in this vein are due to Bagan [9] and Brault-Baron [16]; see [40] for a survey. Bagan et al. [10] complemented their result by a conditional lower bound: assuming that Boolean matrix multiplication cannot be accomplished in time $O(n^2)$, self-join-free acyclic CQs that are *not* free-connex *cannot* be enumerated with constant delay and FPL-preprocessing. This demonstrates that even if the evaluation of *Boolean* queries is easy (as known for all acyclic CQs [42]), the enumeration of the results of non-Boolean queries might be hard (here, for acyclic CQs that are not free-connex).

Bagan et al. [10] also introduced the notion of *free-connex* (fc) treewidth (tw) of a CQ and showed that for every class Φ of CQs of bounded fc-tw, within FPT-preprocessing time, one can build a data structure that allows constant delay enumeration of the query results. This can be viewed as a generalisation, to the non-Boolean case, of the well-known result stating that the model-checking problem for classes of Boolean CQs of bounded treewidth is FPT. Note that for non-Boolean queries – even if they come from a class of bounded fc-tw – the size of the query result may be $N^{\Omega(\|\varphi\|)}$, i.e., far too large to be computed entirely within FPT-preprocessing time; and generalising the known tractability result for Boolean CQs to the non-Boolean case is far from trivial.

In a series of papers, the FPT-result for Boolean CQs has been strengthened to more and more general width-measures, namely to classes of queries of bounded generalised hypertree width (ghw) [23], bounded fractional hypertree width (fhw) [24], and bounded submodular width (subw) [33]. The result on bounded fhw has been generalised to the non-Boolean case in the context of factorised databases [37], which implies constant delay enumeration after FPT-preprocessing for CQs of bounded free-connex fractional hypertree width (fc-fhw). Related data structures that allow constant delay enumeration after FPT-preprocessing for (quantifier-free) CQs of bounded (fc-)fhw have also been provided in [18, 26].

An analogous generalisation of the result on bounded submodular width, however, is still missing. The present paper’s main result closes this gap: we show that on classes of CQs of bounded fc-subw, within FPT-preprocessing time one can build a data structure that allows constant delay enumeration of the query results. And within the same FPT-preprocessing time, one can also construct a data structure that enables to test in constant time whether an input tuple belongs to the query result. Our proof uses Marx’s splitting routine [33] to decompose the query result of φ on \mathcal{A} into the union of results of several queries φ_i on several

¹ In this paper, structures will always be finite and relational.

structures \mathcal{A}_i but we have to tackle the additional technical difficulty to ensure that the results of all the φ_i on \mathcal{A}_i can be enumerated efficiently. Once having achieved this, we can conclude by using an elegant trick provided by Durand and Strozecki [21] for enumerating, without repetition, the union of query results.

As an immediate consequence of the lower bound provided by Marx [33] in the context of Boolean CQs of unbounded submodular width, one obtains that our main result is tight for certain classes of CQs, namely, recursively enumerable classes Φ of quantifier-free and self-join-free CQs: assuming the exponential time hypothesis (ETH), such a class Φ allows constant delay enumeration after FPT-preprocessing if, and only if, Φ has bounded fc-subw.

Let us mention a related recent result which, however, is incomparable to ours. Abo Khamis et al. [3] designed an algorithm for evaluating a quantifier-free CQ φ of submodular width w within time $O(N^w) \cdot (\log N)^{f(\varphi)} + O(r \cdot \log N)$; and an analogous result is also achieved for non-quantifier-free CQs of fc-subw w [2]. Here, N is the size of the input structure, r is the number of tuples in the query result, and $f(\varphi)$ is at least exponential in number of variables of φ . In particular, the algorithm does not distinguish between a preprocessing phase and an enumeration phase and does not provide a guarantee on the delay. Also, due to the factor $(\log N)^{f(\varphi)}$, where $f(\varphi)$ is not bounded in terms of w , it seems unlikely that the approach can easily be adapted to yield a method for constant delay enumeration after FPT-preprocessing for classes of CQs of bounded fc-subw.

Outline. The rest of the paper is structured as follows. Section 2 provides basic notations concerning structures, queries, and constant delay enumeration. Section 3 recalls concepts of (free-connex) decompositions of queries, provides a precise statement of our main result, and collects the necessary tools for obtaining this result. Section 4 is devoted to the detailed proof of our main result. We conclude in Section 5. Due to space restrictions, some proof details had to be deferred to the full version of the paper.

2 Preliminaries

In this section we fix notation and summarise basic definitions.

Basic notation. We write \mathbb{N} and $\mathbb{R}_{\geq 0}$ for the set of non-negative integers and reals, respectively, and we let $\mathbb{N}_{\geq 1} := \mathbb{N} \setminus \{0\}$ and $[n] := \{1, \dots, n\}$ for all $n \in \mathbb{N}_{\geq 1}$. By 2^S we denote the power set of a set S . Whenever G denotes a graph, we write $V(G)$ and $E(G)$ for the set of nodes and the set of edges, respectively, of G . Whenever writing a to denote a k -tuple (for some arity $k \in \mathbb{N}$), we write a_i to denote the tuple's i -th component; i.e., $a = (a_1, \dots, a_k)$. For a k -tuple a and indices $i_1, \dots, i_\ell \in [k]$ we let $\pi_{i_1, \dots, i_\ell}(a) := (a_{i_1}, \dots, a_{i_\ell})$. For a set S of k -tuples we let $\pi_{i_1, \dots, i_\ell}(S) := \{\pi_{i_1, \dots, i_\ell}(a) : a \in S\}$.

If h and g are mappings with domains X and Y , respectively, we say that h and g are *joinable* if $h(z) = g(z)$ holds for all $z \in X \cap Y$. In case that h and g are joinable, we write $h \bowtie g$ to denote the mapping f with domain $X \cup Y$ where $f(x) = h(x)$ for all $x \in X$ and $f(y) = g(y)$ for all $y \in Y$. If A and B are sets of mappings with domains X and Y , respectively, then $A \bowtie B := \{h \bowtie g : h \in A, g \in B, \text{ and } h \text{ and } g \text{ are joinable}\}$.

We use the following further notation where A is a set of mappings with domain X and $h \in A$. For a set $I \subseteq X$, the projection $\pi_I(h)$ is the restriction $h|_I$ of h to I ; and $\pi_I(A) := \{\pi_I(h) : h \in A\}$. For objects z, c where $z \notin X$, we write $h \cup \{(z, c)\}$ for the extension h' of h to domain $X \cup \{z\}$ with $h'(z) = c$ and $h'(x) = h(x)$ for all $x \in X$.

Signatures and structures. A *signature* is a finite set σ of relation symbols, where each $R \in \sigma$ is equipped with a fixed *arity* $\text{ar}(R) \in \mathbb{N}_{\geq 1}$. A σ -*structure* \mathcal{A} consists of a finite set A (called the *universe* or *domain* of \mathcal{A}) and an $\text{ar}(R)$ -ary relation $R^{\mathcal{A}} \subseteq A^{\text{ar}(R)}$ for each $R \in \sigma$. The *size* $\|\sigma\|$ of a signature σ is $|\sigma| + \sum_{R \in \sigma} \text{ar}(R)$. We write $n^{\mathcal{A}}$ to denote the cardinality $|A|$ of \mathcal{A} 's universe, we write $m^{\mathcal{A}}$ to denote the number of tuples in \mathcal{A} 's largest relation, and we write $N^{\mathcal{A}}$ or $\|\mathcal{A}\|$ to denote the size of a reasonable encoding of \mathcal{A} . To be specific, let $N^{\mathcal{A}} = \|\mathcal{A}\| = \|\sigma\| + n^{\mathcal{A}} + \sum_{R \in \sigma} \|R^{\mathcal{A}}\|$, where $\|R^{\mathcal{A}}\| = \text{ar}(R) \cdot |R^{\mathcal{A}}|$. Whenever \mathcal{A} is clear from the context, we will omit the superscript $^{\mathcal{A}}$ and write n, m, N instead of $n^{\mathcal{A}}, m^{\mathcal{A}}, N^{\mathcal{A}}$. Consider signatures σ and τ with $\sigma \subseteq \tau$. The σ -*reduct* of a τ -structure \mathcal{B} is the σ -structure \mathcal{A} with $A = B$ and $R^{\mathcal{A}} = R^{\mathcal{B}}$ for all $R \in \sigma$. A τ -*expansion* of a σ -structure \mathcal{A} is a τ -structure \mathcal{B} whose σ -reduct is \mathcal{A} .

Conjunctive Queries. We fix a countably infinite set **var** of *variables*. We allow queries to use arbitrary relation symbols of arbitrary arities. An *atom* α is of the form $R(v_1, \dots, v_r)$ with $r = \text{ar}(R)$ and $v_1, \dots, v_r \in \mathbf{var}$. We write $\text{vars}(\alpha)$ to denote the set of variables occurring in α . A *conjunctive query* (CQ, for short) is of the form $\exists z_1 \cdots \exists z_\ell (\alpha_1 \wedge \cdots \wedge \alpha_d)$, where $\ell \in \mathbb{N}$, $d \in \mathbb{N}_{\geq 1}$, α_j is an atom for every $j \in [d]$, and z_1, \dots, z_ℓ are pairwise distinct elements in $\text{vars}(\alpha_1) \cup \cdots \cup \text{vars}(\alpha_d)$. For such a CQ φ we let $\text{atoms}(\varphi) = \{\alpha_1, \dots, \alpha_d\}$. We write $\text{vars}(\varphi)$ and $\sigma(\varphi)$ for the set of variables and the set of relation symbols occurring in φ , respectively. The set of *quantified* variables of φ is $\text{quant}(\varphi) := \{z_1, \dots, z_\ell\}$, and the set of *free* variables is $\text{free}(\varphi) := \text{vars}(\varphi) \setminus \text{quant}(\varphi)$. We sometimes write $\varphi(x_1, \dots, x_k)$ to indicate that x_1, \dots, x_k are the free variables of φ . The *arity* of φ is the number $k := |\text{free}(\varphi)|$. The query φ is called *quantifier-free* if $\text{quant}(\varphi) = \emptyset$, it is called *Boolean* if its arity is 0, and it is called *self-join-free* if no relation symbol occurs more than once in φ .

The semantics are defined as usual: A *valuation* for φ on a $\sigma(\varphi)$ -structure \mathcal{A} is a mapping $\beta : \text{vars}(\varphi) \rightarrow A$. A valuation β is a *homomorphism* from φ to a \mathcal{A} if for every atom $R(v_1, \dots, v_r) \in \text{atoms}(\varphi)$ we have $(\beta(v_1), \dots, \beta(v_r)) \in R^{\mathcal{A}}$. The *query result* $\llbracket \varphi \rrbracket^{\mathcal{A}}$ of a CQ φ on the $\sigma(\varphi)$ -structure \mathcal{A} is defined as the set $\{\beta \mid \beta \text{ is a homomorphism from } \varphi \text{ to } \mathcal{A}\}$. Often, we will identify the mappings $g \in \llbracket \varphi \rrbracket^{\mathcal{A}}$ with tuples $(g(x_1), \dots, g(x_k))$, where x_1, \dots, x_k is a fixed listing of the free variables of φ .

The size $\|\varphi\|$ of a query φ is the length of φ when viewed as a word over the alphabet $\sigma(\varphi) \cup \text{vars}(\varphi) \cup \{\exists, \wedge, (,)\} \cup \{, \}$.

Model of computation. For the complexity analysis we assume the RAM-model with a uniform cost measure. In particular, storing and accessing elements from a structure's universe requires $O(1)$ space and time. For an r -ary relation $R^{\mathcal{A}}$ we can construct in time $O(\|R^{\mathcal{A}}\|)$ an index that allows to enumerate $R^{\mathcal{A}}$ with $O(1)$ delay and to test for a given r -tuple a whether $a \in R^{\mathcal{A}}$ in time $O(r)$. Moreover, for every $\{i_1, \dots, i_\ell\} \subseteq [r]$ we can build a data structure where we can enumerate for every ℓ -tuple b the selection $\{a \in R^{\mathcal{A}} : \pi_{i_1, \dots, i_\ell}(a) = b\}$ with $O(1)$ delay. Such a data structure can be constructed in time $O(\|R^{\mathcal{A}}\|)$, for instance by a linear scan over $R^{\mathcal{A}}$ where we add every tuple $a \in R^{\mathcal{A}}$ to a list $\mathcal{L}_{\pi_{i_1, \dots, i_\ell}(a)}$. Using a constant access data structure of linear size, the list \mathcal{L}_b can be accessed in time $O(\ell)$ when receiving an ℓ -tuple b .

Constant delay enumeration and testing. An *enumeration algorithm* for query evaluation consists of two phases: the preprocessing phase and the enumeration phase. In the preprocessing phase the algorithm is allowed to do arbitrary preprocessing on the query φ and the input structure \mathcal{A} . We denote the time required for this phase by t_p . In the subsequent

enumeration phase the algorithm enumerates, without repetition, all tuples (or, mappings) in the query result $\llbracket \varphi \rrbracket^{\mathcal{A}}$, followed by the end-of-enumeration message **EOE**. The *delay* t_d is the maximum time that passes between the start of the enumeration phase and the output of the first tuple, between the output of two consecutive tuples, and between the last tuple and **EOE**.

A *testing algorithm* for query evaluation also starts with a preprocessing phase of time t_p in which a data structure is computed that allows to test for a given tuple (or, mapping) b whether it is contained in the query result $\llbracket \varphi \rrbracket^{\mathcal{A}}$. The *testing time* t_t of the algorithm is an upper bound on the time that passes between receiving b and providing the answer.

One speaks of *constant delay* (testing time) if the delay (testing time) depends on the query φ , but not on the input structure \mathcal{A} .

We make use of the following result from Durand and Strozecki, which allows to efficiently enumerate the union of query results, provided that each query result in the union can be enumerated and tested efficiently. Note that this is not immediate, because the union might contain many duplicates that need to be avoided during enumeration.

► **Theorem 2.1** ([21]). *Suppose that there is an enumeration algorithm \mathbb{A} that receives a query φ and a structure \mathcal{A} and enumerates $\llbracket \varphi \rrbracket^{\mathcal{A}}$ with delay $t_d(\varphi)$ after $t_p(\varphi, \mathcal{A})$ preprocessing time. Further suppose that there is a testing algorithm \mathbb{B} that receives a query φ and a structure \mathcal{A} and has $t_p(\varphi, \mathcal{A})$ preprocessing time and $t_t(\varphi)$ testing time. Then there is an algorithm \mathbb{C} that receives ℓ queries φ_i and structures \mathcal{A}_i and allows to enumerate $\bigcup_{i \in [\ell]} \llbracket \varphi_i \rrbracket^{\mathcal{A}_i}$ with $O(\sum_{i \in [\ell]} t_d(\varphi_i) + \sum_{i \in [\ell]} t_t(\varphi_i))$ delay after $O(\sum_{i \in [\ell]} t_p(\varphi_i, \mathcal{A}_i))$ preprocessing time.*

Proof (sketch). The induction start $\ell = 1$ is trivial. For the induction step $\ell \rightarrow \ell + 1$ start an enumeration of $\bigcup_{i \in [\ell]} \llbracket \varphi_i \rrbracket^{\mathcal{A}_i}$ and test for every tuple whether it is contained in $\llbracket \varphi_{\ell+1} \rrbracket^{\mathcal{A}_{\ell+1}}$. If the answer is no, then output the tuple. Otherwise discard the tuple and instead output the next tuple in an enumeration of $\llbracket \varphi_{\ell+1} \rrbracket^{\mathcal{A}_{\ell+1}}$. Subsequently enumerate the remaining tuples from $\llbracket \varphi_{\ell+1} \rrbracket^{\mathcal{A}_{\ell+1}}$. ◀

3 Main Result

At the end of this section, we provide a precise statement of our main result. Before we can do so, we have to recall the concept of free-connex decompositions of queries and the notion of submodular width. It will be convenient for us to use the following notation.

► **Definition 3.1.** *Let $\varphi = \exists z_1 \cdots \exists z_\ell (\alpha_1 \wedge \cdots \wedge \alpha_d)$ be a CQ and $S \subseteq \text{vars}(\varphi)$. We write $\varphi\langle S \rangle$ for the CQ that is equivalent to the expression*

$$(\exists y_1 \cdots \exists y_r \alpha_1) \wedge \cdots \wedge (\exists y_1 \cdots \exists y_r \alpha_d), \quad (1)$$

where $\{y_1, \dots, y_r\} = \text{vars}(\varphi) \setminus S$.

Note that $\varphi\langle S \rangle$ is obtained from φ by discarding existential quantification and projecting every atom to S , hence $\text{free}(\varphi\langle S \rangle) = S$. However, $\llbracket \varphi\langle S \rangle \rrbracket^{\mathcal{A}}$ shall not be confused with the projection of $\llbracket \varphi \rrbracket^{\mathcal{A}}$ to S . In fact, it might be that $\llbracket \varphi \rrbracket^{\mathcal{A}}$ is empty, but $\llbracket \varphi\langle S \rangle \rrbracket^{\mathcal{A}}$ is not, as the following example illustrates:

$$\varphi = E(x, y) \wedge E(y, z) \wedge E(x, z) \quad \text{and} \quad (2)$$

$$\varphi\langle \{x, z\} \rangle \equiv \exists y E(x, y) \wedge \exists y E(y, z) \wedge \exists y E(x, z) \quad (3)$$

$$\equiv E(x, z). \quad (4)$$

3.1 Constant delay enumeration using tree decompositions

We use the same notation as [22] for decompositions of queries: A *tree decomposition* (TD, for short) of a CQ φ is a tuple $TD = (T, \chi)$, for which the following two conditions are satisfied:

1. $T = (V(T), E(T))$ is a finite undirected tree.
2. χ is a mapping that associates with every node $t \in V(T)$ a set $\chi(t) \subseteq \text{vars}(\varphi)$ such that
 - a. for each atom $\alpha \in \text{atoms}(\varphi)$ there exists $t \in V(T)$ such that $\text{vars}(\alpha) \subseteq \chi(t)$, and
 - b. for each variable $v \in \text{vars}(\varphi)$ the set $\chi^{-1}(v) := \{t \in V(T) : v \in \chi(t)\}$ induces a connected subtree of T (this condition is called *path condition*).

To use a tree decomposition $TD = (T, \chi)$ of φ for query evaluation one considers, for each $t \in V(T)$ the query $\varphi\langle S \rangle$ for $S := \chi(t)$, evaluates this query on the input structure \mathcal{A} , and then combines these results for all $t \in V(T)$ along a bottom-up traversal of T . If the query is Boolean, this yields the result of φ on \mathcal{A} ; if it is non-Boolean, $\llbracket \varphi \rrbracket^{\mathcal{A}}$ can be computed by performing additional traversals of T . This approach is efficient if the result sets $\llbracket \varphi\langle \chi(t) \rangle \rrbracket^{\mathcal{A}}$ are small and can be computed efficiently (later on, we will sometimes refer to the sets $\llbracket \varphi\langle \chi(t) \rangle \rrbracket^{\mathcal{A}}$ as *projections on bags*).

The simplest queries where this is the case are acyclic queries [11, 15]. A number of equivalent characterisations of the acyclic CQs have been provided in the literature (cf. [1, 23, 25, 17]); among them a characterisation by Gottlob et al. [23] stating that a CQ is acyclic if and only if it has a tree-decomposition where every bag is covered by an atom, i.e., for every bag $\chi(t)$ there is some atom α in φ with $\chi(t) \subseteq \text{vars}(\alpha)$. The approach described above leads to a linear time algorithm for evaluating an acyclic CQ φ that is Boolean, and if φ is non-Boolean, $\llbracket \varphi \rrbracket^{\mathcal{A}}$ is computed in time linear in $\|\mathcal{A}\| + \|\llbracket \varphi \rrbracket^{\mathcal{A}}\|$. This method is known as *Yannakakis' algorithm*. But this algorithm does not distinguish between a preprocessing phase and an enumeration phase and does not guarantee constant delay enumeration. In fact, Bagan et al. identified the following additional property that is needed to ensure constant delay enumeration.

► **Definition 3.2** ([10]). *A tree decomposition $TD = (T, \chi)$ of a CQ φ is free-connex if there is a subset $U \subseteq V(T)$ that induces a connected subtree of T and that satisfies the condition $\text{free}(\varphi) = \bigcup_{t \in U} \chi(t)$.*

Bagan et al. [10] identified the *free-connex* acyclic CQs, i.e., the CQs φ that have a free-connex tree decomposition where every bag is covered by an atom, as the fragment of the acyclic CQs whose results can be enumerated with constant delay after FPL-preprocessing:

► **Theorem 3.3** (Bagan et al. [10]). *There is a computable function f and an algorithm which receives a free-connex acyclic CQ φ and a $\sigma(\varphi)$ -structure \mathcal{A} and computes within $t_p = f(\varphi)O(\|\mathcal{A}\|)$ preprocessing time and space a data structure that allows to*

- (i) *enumerate $\llbracket \varphi \rrbracket^{\mathcal{A}}$ with $f(\varphi)$ delay and*
- (ii) *test for a given tuple (or, mapping) b if $b \in \llbracket \varphi \rrbracket^{\mathcal{A}}$ within $f(\varphi)$ testing time.*

The approach of using free-connex tree decompositions for constant delay enumeration can be extended from acyclic CQs to arbitrary CQs. To do this, we have to compute for every bag $\chi(t)$ in the tree decomposition the projection $\llbracket \varphi\langle \chi(t) \rangle \rrbracket^{\mathcal{A}}$. This reduces the task to the acyclic case, where the free-connex acyclic query contains one atom α with $\text{vars}(\alpha) = \chi(t)$ for every bag $\chi(t)$ and the corresponding relation is defined by $\llbracket \varphi\langle \chi(t) \rangle \rrbracket^{\mathcal{A}}$. Because the runtime in this approach is dominated by computing $\llbracket \varphi\langle \chi(t) \rangle \rrbracket^{\mathcal{A}}$, it is only feasible if the projections are efficiently computable for every bag. If the decomposition has bounded treewidth or

bounded fractional hypertree width, then it is possible to compute $\llbracket \varphi(\chi(t)) \rrbracket^{\mathcal{A}}$ for every bag in time $f(\varphi) \cdot \|\mathcal{A}\|^{O(1)}$ [24], which in turn implies that the result can be enumerated after FPT-preprocessing time for CQs of bounded fc-tw [10] and for CQs of bounded fc-fhw [37].

3.2 Submodular width and statement of the main result

Before providing the precise definition of the submodular width of a query, let us first consider an example. The central idea behind algorithms that rely on submodular width [33, 3, 38] is to split the input structure into several parts and use for every part a different tree decomposition of φ . This will give a significant improvement over the fractional hypertree width, which uses only one tree decomposition of φ . A typical example to illustrate this idea is the following 4-cycle query (see also [3, 38]): $\varphi_4 := E_{12}(x_1, x_2) \wedge E_{23}(x_2, x_3) \wedge E_{34}(x_3, x_4) \wedge E_{41}(x_4, x_1)$.

There are essentially two non-trivial tree decompositions $TD = (T, \chi')$, $TD'' = (T, \chi'')$ of φ_4 , which are both defined over the two-vertex tree $T = (\{t_1, t_2\}, \{\{t_1, t_2\}\})$ by $\chi'(t_1) = \{x_1, x_2, x_3\}$, $\chi'(t_2) = \{x_1, x_3, x_4\}$ and $\chi''(t_1) = \{x_2, x_3, x_4\}$, $\chi''(t_2) = \{x_1, x_2, x_4\}$. Both tree decompositions lead to an optimal fractional hypertree decomposition of width $fhw(\varphi_4) = 2$. Indeed, for the worst-case instance \mathcal{A} with

$$E_{12}^{\mathcal{A}} = E_{34}^{\mathcal{A}} := [\ell] \times \{a\} \cup \{b\} \times [\ell] \quad E_{23}^{\mathcal{A}} = E_{41}^{\mathcal{A}} := [\ell] \times \{b\} \cup \{a\} \times [\ell]$$

we have $\|\mathcal{A}\| = O(\ell)$ while the projections on the bags have size $\Omega(\ell^2)$ in both decompositions:²

$$\begin{aligned} \llbracket \varphi_4(\chi'(t_1)) \rrbracket^{\mathcal{A}} &= \llbracket \varphi_4(\chi'(t_2)) \rrbracket^{\mathcal{A}} = [\ell] \times \{a\} \times [\ell] \cup \{b\} \times [\ell] \times \{b\}, \\ \llbracket \varphi_4(\chi''(t_1)) \rrbracket^{\mathcal{A}} &= \llbracket \varphi_4(\chi''(t_2)) \rrbracket^{\mathcal{A}} = [\ell] \times \{b\} \times [\ell] \cup \{a\} \times [\ell] \times \{a\}. \end{aligned}$$

However, we can split \mathcal{A} into \mathcal{A}' and \mathcal{A}'' such that $\llbracket \varphi_4 \rrbracket^{\mathcal{A}}$ is the disjoint union of $\llbracket \varphi_4 \rrbracket^{\mathcal{A}'}$ and $\llbracket \varphi_4 \rrbracket^{\mathcal{A}''}$ and the bag-sizes in the respective decompositions are small:

$$\begin{aligned} E_{12}^{\mathcal{A}'} &= E_{34}^{\mathcal{A}'} := \{b\} \times [\ell] & E_{23}^{\mathcal{A}'} &= E_{41}^{\mathcal{A}'} := [\ell] \times \{b\} \\ E_{12}^{\mathcal{A}''} &= E_{34}^{\mathcal{A}''} := [\ell] \times \{a\} & E_{23}^{\mathcal{A}''} &= E_{41}^{\mathcal{A}''} := \{a\} \times [\ell] \\ \llbracket \varphi_4(\chi'(t_1)) \rrbracket^{\mathcal{A}'} &= \llbracket \varphi_4(\chi'(t_2)) \rrbracket^{\mathcal{A}'} = \{b\} \times [\ell] \times \{b\}, \\ \llbracket \varphi_4(\chi''(t_1)) \rrbracket^{\mathcal{A}''} &= \llbracket \varphi_4(\chi''(t_2)) \rrbracket^{\mathcal{A}''} = \{a\} \times [\ell] \times \{a\}. \end{aligned}$$

Thus, we can efficiently evaluate φ_4 on \mathcal{A}' using TD' and φ_4 on \mathcal{A}'' using TD'' (in time $O(\ell)$ in this example) and then combine both results to obtain $\varphi_4(\mathcal{A})$. Using the strategy of Alon et al. [4], it is possible to split *every* database \mathcal{A} for this particular 4-cycle query φ_4 into two instances \mathcal{A}' and \mathcal{A}'' such that the bag sizes in TD' on \mathcal{A}' as well as in TD'' on \mathcal{A}'' are bounded by $\|\mathcal{A}\|^{3/2}$ and can be computed in time $O(\|\mathcal{A}\|^{3/2})$ (see [3, 38] for a detailed account on this strategy). As both decompositions are free-connex, this also leads to a constant delay enumeration algorithm for φ_4 with $O(\|\mathcal{A}\|^{3/2})$ time preprocessing, which improves the $O(\|\mathcal{A}\|^2)$ preprocessing time that follows from using one decomposition.

In general, whether such a data-dependent decomposition is possible is determined by the submodular width $subw(\varphi)$ of the query. The notion of submodular width was introduced in [33]. To present its definition, we need the following terminology. A function $g: 2^{\text{vars}(\varphi)} \rightarrow \mathbb{R}_{\geq 0}$ is

- *monotone* if $g(U) \leq g(V)$ for all $U \subseteq V \subseteq \text{vars}(\varphi)$.

² recall from Section 2 our convention to identify mappings in query results with tuples; the free variables are listed canonically here, by increasing indices

- *edge-dominated* if $g(\text{vars}(\alpha)) \leq 1$ for every atom $\alpha \in \text{atoms}(\varphi)$.
 - *submodular*, if $g(U) + g(V) \geq g(U \cap V) + g(U \cup V)$ for every $U, V \subseteq \text{vars}(\varphi)$.
- We denote by $\mathcal{S}(\varphi)$ the set of all monotone, edge-dominated, submodular functions $g: 2^{\text{vars}(\varphi)} \rightarrow \mathbb{R}_{\geq 0}$ that satisfy $g(\emptyset) = 0$, and by $\mathcal{T}(\varphi)$ the set of all tree decompositions of φ . The *submodular width* of a conjunctive query φ is

$$\text{subw}(\varphi) := \sup_{g \in \mathcal{S}(\varphi)} \min_{(T, \chi) \in \mathcal{T}(\varphi)} \max_{t \in V(T)} g(\chi(t)). \quad (5)$$

In particular, if the submodular width of φ is bounded by w , then for every submodular function g there is a tree decomposition in which every bag B satisfies $g(B) \leq w$.

It is known that $\text{subw}(\varphi) \leq \text{fhw}(\varphi)$ for all queries φ [33, Proposition 3.7]. Moreover, there is a constant c and a family of queries φ such that $\text{subw}(\varphi) \leq c$ is bounded and $\text{fhw}(\varphi) = \Omega(\sqrt{\log \|\varphi\|})$ is unbounded [32, 33]. The main result in [33] is that the submodular width characterises the tractability of Boolean CQs in the following sense.

► **Theorem 3.4** ([33]).

- (1) *There is a computable function f and an algorithm that receives a Boolean CQ φ , $\text{subw}(\varphi)$, and a $\sigma(\varphi)$ -structure \mathcal{A} and evaluates φ on \mathcal{A} in time $f(\varphi) \|\mathcal{A}\|^{O(\text{subw}(\varphi))}$.*
- (2) *Let Φ be a recursively enumerable class of Boolean, self-join-free CQs of unbounded submodular width. Assuming the exponential time hypothesis (ETH) there is no algorithm which, upon input of a query $\varphi \in \Phi$ and a structure \mathcal{A} , evaluates φ on \mathcal{A} in time $\|\mathcal{A}\|^{o(\text{subw}(\varphi)^{1/4})}$.*

The *free-connex submodular width* of a conjunctive query φ is defined in a similar way as submodular width, but this time ranges over the set $\text{fc}\mathcal{T}(\varphi)$ of all free-connex tree decompositions of φ (it is easy to see that we can assume that $\text{fc}\mathcal{T}(\varphi)$ is finite).

$$\text{fc-subw}(\varphi) := \sup_{g \in \mathcal{S}(\varphi)} \min_{(T, \chi) \in \text{fc}\mathcal{T}(\varphi)} \max_{t \in V(T)} g(\chi(t)). \quad (6)$$

Note that if φ is either quantifier-free or Boolean, we have $\text{fc-subw}(\varphi) = \text{subw}(\varphi)$. In general, this is not always the case. Consider for example the following quantified version $\varphi'_4 := \exists x_1 \exists x_3 \varphi_4$ of the quantifier-free 4-cycle query φ_4 . Here we have $\text{subw}(\varphi'_4) = \frac{3}{2}$, but $\text{fc-subw}(\varphi'_4) = 2$: one can verify $\text{fc-subw}(\varphi'_4) \geq 2$ by noting that every free-connex tree decomposition contains a bag $\{x_1, x_2, x_3, x_4\}$ and taking the submodular function $g(U) := \frac{1}{2}|U|$. Now we are ready to state the main theorem of this paper.

► **Theorem 3.5.** *For every $\delta > 0$ and $w \geq 1$ there is a computable function f and an algorithm which receives a CQ φ with $\text{fc-subw}(\varphi) \leq w$ and a $\sigma(\varphi)$ -structure \mathcal{A} and computes within $t_p = f(\varphi) \|\mathcal{A}\|^{(2+\delta)w}$ preprocessing time and space $f(\varphi) \|\mathcal{A}\|^{(1+\delta)w}$ a data structure that allows to*

- (i) *enumerate $\llbracket \varphi \rrbracket^{\mathcal{A}}$ with $f(\varphi)$ delay and*
- (ii) *test for a given tuple (or, mapping) b if $b \in \llbracket \varphi \rrbracket^{\mathcal{A}}$ within $f(\varphi)$ testing time.*

The following corollary is an immediate consequence of Theorem 3.5 and Theorem 3.4. A class Φ of CQs is said to be of *bounded free-connex submodular width* if there exists a number w such that $\text{fc-subw}(\varphi) \leq w$ for all $\varphi \in \Phi$. And by an *algorithm for Φ that enumerates with constant delay after FPT-preprocessing* we mean an algorithm that receives a query $\varphi \in \Phi$ and a $\sigma(\varphi)$ -structure \mathcal{A} and spends $f(\varphi) \|\mathcal{A}\|^{O(1)}$ preprocessing time and then enumerates $\llbracket \varphi \rrbracket^{\mathcal{A}}$ with delay $f(\varphi)$, for a computable function f .

► **Corollary 3.6.**

- (1) For every class Φ of CQs of bounded free-connex submodular width, there is an algorithm for Φ that enumerates with constant delay after FPT-preprocessing.
- (2) Let Φ be a recursively enumerable class of quantifier-free self-join-free CQs and assume that the exponential time hypothesis (ETH) holds.
Then there is an algorithm for Φ that enumerates with constant delay after FPT-preprocessing if, and only if, Φ has bounded free-connex submodular width.

4 Proof of the Main Result

To prove Theorem 3.5, we make use of Marx’s splitting routine for queries of bounded submodular width. In the following, we will adapt the main definitions and concepts from [33] to our notions. While doing this, we provide the following additional technical contributions: First, we give a detailed time and space analysis of the algorithm and improve the runtime of the consistency algorithm [33, Lemma 4.5] from quadratic to linear (see Lemma 4.2). Second, we fix an oversight in [33, Lemma 4.12] by establishing strong M -consistency (unfortunately, this fix incurs a blow-up in running time). Afterwards we prove our main theorem, where the non-Boolean setting requires us to relax Marx’s partition into *refinements* (Lemma 4.5) so that the subinstances are no longer disjoint.

Let φ be a quantifier-free CQ with $\text{vars}(\varphi) = \{x_1, \dots, x_k\}$, and let $\sigma := \sigma(\varphi)$. For every $S = \{x_{i_1}, \dots, x_{i_\ell}\} \subseteq \text{vars}(\varphi)$ where $i_1 < \dots < i_\ell$ we set $x_S := (x_{i_1}, \dots, x_{i_\ell})$ and let $R_S \notin \sigma$ be a fresh ℓ -ary relation symbol. For every collection $\mathfrak{s} \subseteq 2^{\text{vars}(\varphi)}$ we let

$$\sigma_{\mathfrak{s}} := \sigma \cup \{R_S : S \in \mathfrak{s}\} \quad \text{and} \quad (7)$$

$$\varphi_{\mathfrak{s}} := \varphi \wedge \bigwedge_{S \in \mathfrak{s}} R_S(x_S). \quad (8)$$

A *refinement* of φ and a σ -structure \mathcal{A} is a pair $(\mathfrak{s}, \mathcal{B})$, where $\mathfrak{s} \subseteq 2^{\text{vars}(\varphi)}$ is closed under taking subsets and \mathcal{B} is a $\sigma_{\mathfrak{s}}$ -expansion of \mathcal{A} . Note that if $(\mathfrak{s}, \mathcal{B})$ is a refinement of φ and \mathcal{A} , then $\llbracket \varphi_{\mathfrak{s}} \rrbracket^{\mathcal{B}} \subseteq \llbracket \varphi \rrbracket^{\mathcal{A}}$. In the following we will construct refinements that do not change the result relation, i. e., $\llbracket \varphi_{\mathfrak{s}} \rrbracket^{\mathcal{B}} = \llbracket \varphi \rrbracket^{\mathcal{A}}$. Subsequently, we will split refinements in order to partition the query result.

The following definition collects useful properties of refinements. Recall from Section 2 that for a CQ ψ and a structure \mathcal{B} , the query result $\llbracket \psi \rrbracket^{\mathcal{B}}$ actually is a set of mappings from $\text{free}(\psi)$ to B . For notational convenience we define $\mathbf{R}_S^{\mathcal{B}} := \llbracket R_S(x_S) \rrbracket^{\mathcal{B}}$ and use the set $\mathbf{R}_S^{\mathcal{B}}$ of mappings instead of the relation $R_S^{\mathcal{B}}$. In particular, by addressing/inserting/deleting a mapping $h: S \rightarrow B$ from $\mathbf{R}_S^{\mathcal{B}}$ we mean addressing/inserting/deleting the tuple $(h(x_{i_1}), \dots, h(x_{i_\ell}))$ from $R_S^{\mathcal{B}}$, where $(x_{i_1}, \dots, x_{i_\ell}) = x_S$.

► **Definition 4.1.** Let φ be a quantifier-free σ -CQ, \mathcal{A} a σ -structure, $(\mathfrak{s}, \mathcal{B})$ a refinement of φ and \mathcal{A} , and M an integer.

1. The refinement $(\mathfrak{s}, \mathcal{B})$ is consistent if

$$\mathbf{R}_S^{\mathcal{B}} = \llbracket \varphi_{\mathfrak{s}} \langle S \rangle \rrbracket^{\mathcal{B}} \text{ for all } S \in \mathfrak{s} \text{ and} \quad (9)$$

$$\mathbf{R}_S^{\mathcal{B}} = \pi_S(\mathbf{R}_T^{\mathcal{B}}) \text{ for all } S, T \in \mathfrak{s} \text{ with } S \subset T. \quad (10)$$

2. The refinement $(\mathfrak{s}, \mathcal{B})$ is M -consistent if it is consistent and

$$S \in \mathfrak{s} \iff \text{for all } T \subseteq S: |\llbracket \varphi_{\mathfrak{s}} \langle T \rangle \rrbracket^{\mathcal{B}}| \leq M. \quad (11)$$

3. The refinement $(\mathfrak{s}, \mathcal{B})$ is strongly M -consistent if it is M -consistent and

$$S \in \mathfrak{s}, T \in \mathfrak{s}, (S \cup T) \notin \mathfrak{s} \implies |\llbracket \varphi_{\mathfrak{s}} \langle S \cup T \rangle \rrbracket^{\mathcal{B}}| > M. \quad (12)$$

58:10 Constant Delay Enumeration with FPT-Preprocessing

The proof of the following lemma is deferred to the full version of the paper.

► **Lemma 4.2.** *There is an algorithm that receives a refinement $\mathcal{R} = (\mathfrak{s}, \mathcal{B})$ of φ and \mathcal{A} and computes in time $O(|\mathfrak{s}| \cdot \|\mathcal{B}\|)$ a consistent refinement $(\mathfrak{s}, \mathcal{B}')$ with $R_S^{\mathcal{B}'} \subseteq R_S^{\mathcal{B}}$ for all $S \in \mathfrak{s}$ and $\llbracket \varphi_{\mathfrak{s}} \rrbracket^{\mathcal{B}'} = \llbracket \varphi_{\mathfrak{s}} \rrbracket^{\mathcal{B}}$.*

► **Lemma 4.3.** *Let φ be a quantifier-free CQ, let \mathcal{A} be a $\sigma(\varphi)$ -structure where the largest relation contains m tuples, and let $M \geq m$. There is an algorithm that computes in time $O(2^{|\text{vars}(\varphi)|} \cdot M^2)$ and space $O(2^{|\text{vars}(\varphi)|} \cdot M)$ a strongly M -consistent refinement $(\mathfrak{s}, \mathcal{B})$ that satisfies $\llbracket \varphi \rrbracket^{\mathcal{A}} = \llbracket \varphi_{\mathfrak{s}} \rrbracket^{\mathcal{B}}$.*

■ **Algorithm 1** Computing a strongly M -consistent refinement.

```

1: INPUT: quantifier-free CQ  $\varphi(x_1, \dots, x_k)$ ,  $\sigma(\varphi)$ -structure  $\mathcal{A}$ 
2:  $\mathfrak{s} \leftarrow \emptyset$ ;  $\mathcal{B} \leftarrow \mathcal{A}$ 
3: repeat
4:   for  $\ell = 1, \dots, k$  do ▷ Step 1: Ensure condition (11).
5:     for  $S = \{x_{i_1}, \dots, x_{i_\ell}\} \subseteq \text{vars}(\varphi)$  do
6:       if  $S \notin \mathfrak{s}$  and  $S \setminus \{x\} \in \mathfrak{s}$  for all  $x \in S$  then
7:          $\mathbf{R}_S^{\mathcal{B}} \leftarrow \emptyset$ 
8:         Choose  $x \in S$  arbitrary
9:         for  $h \in \mathbf{R}_{S \setminus \{x\}}^{\mathcal{B}}$  and  $c \in \mathcal{A}$  do
10:          if  $h \cup \{(x, c)\} \in \llbracket \varphi_{\mathfrak{s}} \langle S \rangle \rrbracket^{\mathcal{B}}$  then  $\mathbf{R}_S^{\mathcal{B}} \leftarrow \mathbf{R}_S^{\mathcal{B}} \cup \{h \cup \{(x, c)\}\}$ 
11:          if  $|\mathbf{R}_S^{\mathcal{B}}| \leq M$  then  $\mathfrak{s} \leftarrow \mathfrak{s} \cup \{S\}$ 
12:
13:   for  $S, T \in \mathfrak{s}$  such that  $S \cup T \notin \mathfrak{s}$  do ▷ Step 2: Ensure condition (12).
14:     for  $g \in \mathbf{R}_S^{\mathcal{B}}$  and  $h \in \mathbf{R}_T^{\mathcal{B}}$  do
15:       if  $g \bowtie h \in \llbracket \varphi_{\mathfrak{s}} \langle S \cup T \rangle \rrbracket^{\mathcal{B}}$  then  $\mathbf{R}_{S \cup T}^{\mathcal{B}} \leftarrow \mathbf{R}_{S \cup T}^{\mathcal{B}} \cup \{g \bowtie h\}$ 
16:       if  $|\mathbf{R}_{S \cup T}^{\mathcal{B}}| > M$  then break
17:       if  $|\mathbf{R}_{S \cup T}^{\mathcal{B}}| \leq M$  then  $\mathfrak{s} \leftarrow \mathfrak{s} \cup \{S \cup T\}$ 
18:
19:    $(\mathfrak{s}, \mathcal{B}) \leftarrow \text{CONSISTENT}(\mathfrak{s}, \mathcal{B})$  ▷ Step 3: Apply Lemma 4.2 to ensure (9), (10).
20: until  $\mathfrak{s}$  remains unchanged
21: return  $(\mathfrak{s}, \mathcal{B})$ 

```

Proof. The pseudocode of the algorithm is shown in Algorithm 1. For computing the strongly M -consistent refinement we first compute all sets S where for all $T \subseteq S$ we have $|\llbracket \varphi_{\mathfrak{s}} \langle T \rangle \rrbracket^{\mathcal{B}}| \leq M$; as in [33], we say that such sets S are M -small. First note that the empty set is M -small. For nonempty sets S we know that S is only M -small if for every $x \in S$ the set $S \setminus \{x\}$ is M -small and hence already included in \mathfrak{s} . If this is the case, then $\llbracket \varphi_{\mathfrak{s}} \langle S \rangle \rrbracket^{\mathcal{B}}$ can be computed in time $O(M \cdot n)$ by testing for every $h \in \mathbf{R}_{S \setminus \{x\}}^{\mathcal{B}}$ (for an arbitrary $x \in S$) and every element c in the structure's universe, whether $h \cup \{(x, c)\} \in \llbracket \varphi_{\mathfrak{s}} \langle S \rangle \rrbracket^{\mathcal{B}}$. If $|\llbracket \varphi_{\mathfrak{s}} \langle S \rangle \rrbracket^{\mathcal{B}}| \leq M$, then we include S and $\mathbf{R}_S^{\mathcal{B}} := \llbracket \varphi_{\mathfrak{s}} \langle S \rangle \rrbracket^{\mathcal{B}}$ into our current refinement. Afterwards, we want to satisfy the condition on strong M -consistency (12) by trying all pairs of M -small sets S and T . This is the bottleneck of our algorithm and requires time $O(|\mathbf{R}_S^{\mathcal{B}}| \cdot |\mathbf{R}_T^{\mathcal{B}}|) \leq O(M^2)$. In the third step we apply Lemma 4.2 to enforce consistency of the current refinement. In particular, every set $S \cup T$ that was found in step 2 becomes M -small. Note that after deleting tuples to ensure consistency, new sets may become M -small. Therefore, we have to

repeat steps 1–3 until no more sets became M -small. Overall, we repeat the outer loop at most 2^k times, step 1 takes time $2^{O(k)} \cdot M \cdot n$, step 2 takes time $2^{O(k)} \cdot M^2$ and step 3 takes time $2^{O(k)} \cdot M$. Since $n \leq M$ this leads to the required running time. ◀

The key step in the proof of Theorem 3.5 is to compute $f(\varphi)$ strongly M -consistent refinements $(\mathfrak{s}_i, \mathcal{B}_i)$ of φ and \mathcal{A} such that $\llbracket \varphi \rrbracket^{\mathcal{A}} = \bigcup_i \llbracket \varphi_{\mathfrak{s}_i} \rrbracket^{\mathcal{B}_i}$. In addition to being strongly M -consistent, we want the structures \mathcal{B}_i to be uniform in the sense that the degree of tuples (i. e. the number of extensions) is roughly the average degree. We make this precise in a moment, but for illustration it might be helpful to consult the example from Section 3.2 again. In every relation in \mathcal{A} there is one vertex (a or b) of out-degree ℓ and there are ℓ vertices of out-degree 1. Hence the average out-degree is $2\ell/(\ell+1)$ and the vertex degrees are highly imbalanced. However, after splitting the instance in \mathcal{A}' and \mathcal{A}'' , in every relation, all vertices have either out-degree ℓ or 1 and the out-degree of every vertex matches the average out-degree of the corresponding relation. The next definition generalises this to tuples of variables. We call a refinement $(\mathfrak{s}, \mathcal{B})$ *non-trivial*, if every additional relation in the expansion \mathcal{B} is non-empty. For a non-trivial consistent refinement $(\mathfrak{s}, \mathcal{B})$ and $S, T \in \mathfrak{s}$, $S \subseteq T$, we let

$$\text{avgdeg}(S, T) := |\mathbf{R}_T^{\mathcal{B}}|/|\mathbf{R}_S^{\mathcal{B}}| \quad \text{and} \quad (13)$$

$$\text{maxdeg}(S, T) := \max_{g \in \mathbf{R}_S^{\mathcal{B}}} \{h \in \mathbf{R}_T^{\mathcal{B}} : \pi_S(h) = g\}. \quad (14)$$

Note that consistency ensures that these numbers are well-defined and non-zero. Furthermore, we can compute them from $(\mathfrak{s}, \mathcal{B})$ in time $O(|\mathfrak{s}|^2 \cdot \|\mathcal{B}\|)$. By definition we have $\text{maxdeg}(S, T) \geq \text{avgdeg}(S, T)$. The next definition states that maximum degree does not deviate too much from the average degree.

► **Definition 4.4.** *Let $(\mathfrak{s}, \mathcal{B})$ be a non-trivial consistent refinement of φ and \mathcal{A} , and let m be the number of tuples of largest relation of \mathcal{A} . The refinement $(\mathfrak{s}, \mathcal{B})$ is ε -uniform if for all $S, T \in \mathfrak{s}$ with $S \subseteq T$ we have $\text{maxdeg}(S, T) \leq m^\varepsilon \cdot \text{avgdeg}(S, T)$.*

The next lemma uses Marx’s splitting routine to obtain a partition into strongly M -consistent ε -uniform refinements, for $M := m^c$.

► **Lemma 4.5.** *Let φ be a quantifier-free CQ, let \mathcal{A} be a $\sigma(\varphi)$ -structure where the largest relation contains m tuples, and let $c \geq 1$ and $\varepsilon > 0$ be real numbers. There is a computable function f and an algorithm that computes in time $O(f(\varphi, c, \varepsilon) \cdot m^{2c})$ and space $O(f(\varphi, c, \varepsilon) \cdot m^c)$ a sequence of $\ell \leq f(\varphi, c, \varepsilon)$ strongly m^c -consistent ε -uniform refinements $(\mathfrak{s}_i, \mathcal{B}_i)$ such that $\llbracket \varphi \rrbracket^{\mathcal{A}}$ is the disjoint union of the sets $\llbracket \varphi_{\mathfrak{s}_i} \rrbracket^{\mathcal{B}_i}$.*

Proof (sketch). We follow the same splitting strategy as in [33], but use the improved algorithm from Lemma 4.3 to ensure strong m^c -consistency. Starting with the trivial refinement (\emptyset, \mathcal{A}) , in each step we first apply Lemma 4.3 to ensure strong m^c -consistency. Afterwards, we check whether the current refinement $(\mathfrak{s}, \mathcal{B})$ contains sets $S, T \in \mathfrak{s}$ that contradict ε -uniformity, i. e., $S \subseteq T$ and $\text{maxdeg}(S, T) > m^\varepsilon \cdot \text{avgdeg}(S, T)$. If this is the case, we split the refinement $(\mathfrak{s}, \mathcal{B})$ into $(\mathfrak{s}, \mathcal{B}')$ and $(\mathfrak{s}, \mathcal{B}'')$ such that $\mathbf{R}_S^{\mathcal{B}}$ is partitioned into tuples of small degree and tuples of large degree:

$$\mathbf{R}_U^{\mathcal{B}'} = \mathbf{R}_U^{\mathcal{B}''} := \mathbf{R}_U^{\mathcal{B}} \quad \text{for all } U \in \mathfrak{s} \setminus \{S\}, \quad (15)$$

$$\mathbf{R}_S^{\mathcal{B}'} := \{g \in \mathbf{R}_S^{\mathcal{B}} : |\{h \in \mathbf{R}_T^{\mathcal{B}} : \pi_S(h) = g\}| \leq m^{\varepsilon/2} \cdot \text{avgdeg}(S, T)\} \quad (16)$$

$$\mathbf{R}_S^{\mathcal{B}''} := \{g \in \mathbf{R}_S^{\mathcal{B}} : |\{h \in \mathbf{R}_T^{\mathcal{B}} : \pi_S(h) = g\}| > m^{\varepsilon/2} \cdot \text{avgdeg}(S, T)\} \quad (17)$$

It is clear that $\llbracket \varphi \rrbracket^{\mathcal{B}}$ is the disjoint union of $\llbracket \varphi \rrbracket^{\mathcal{B}'}$ and $\llbracket \varphi \rrbracket^{\mathcal{B}''}$ and that the recursion terminates at some point with a sequence of strongly m^c -consistent ϵ -uniform refinements that partition $\llbracket \varphi \rrbracket^{\mathcal{A}}$. It is also not hard to show that the height of the recursion tree is bounded by $2^{O(|\text{vars}(\varphi)|)} \cdot \frac{c}{\epsilon}$ (see [33, Lemma 4.11]). Hence, by Lemma 4.3 the procedure can be implemented in time $O(f(\varphi, c, \epsilon) \cdot m^{2c})$ and space $O(f(\varphi, c, \epsilon) \cdot m^c)$. \blacktriangleleft

The nice thing about ϵ -uniform and strongly m^c -consistent refinements is that they define, for small enough ϵ , a submodular function $g \in \mathcal{S}(\varphi)$, which in turn guarantees the existence of a tree decomposition with small projections on the bags. The following lemma from [33, Lemma 4.12] provides these functions. However, there is an oversight in Marx's proof and in order to fix this, we have to ensure *strong* m^c -consistency instead of only m^c -consistency as stated in [33, Lemma 4.12]. As suggested by Marx (personal communication), an alternative way to achieve strong m^c -consistency would be to enforce m^{2c} -consistency, which leads to the same runtime guarantees, but requires more space.

► **Lemma 4.6.** *Let $(\mathfrak{s}, \mathcal{B})$ be an ϵ -uniform strongly m^c -consistent refinement of φ and \mathcal{A} , and let $c \geq 1$ and $|\text{vars}(\varphi)|^{-3} \geq \epsilon > 0$ be real numbers. Then $g_{\mathfrak{s}, \mathcal{B}}: 2^{\text{vars}(\varphi)} \rightarrow \mathbb{R}_{\geq 0}$ is a monotone, edge-dominated, submodular function that satisfies $g_{\mathfrak{s}, \mathcal{B}}(\emptyset) = 0$:*

$$g_{\mathfrak{s}, \mathcal{B}}(U) := \begin{cases} (1 - \epsilon^{1/3}) \cdot \log_m(|\mathbf{R}_U^{\mathcal{B}}|) + h(U) & \text{if } U \in \mathfrak{s} \\ (1 - \epsilon^{1/3}) \cdot c + h(U) & \text{if } U \notin \mathfrak{s}, \end{cases} \quad (18)$$

where $h(U) := 2\epsilon^{2/3}|U| - \epsilon|U|^2 \geq 0$ for all $U \subseteq \text{vars}(\varphi)$.

The proof can be copied verbatim from Marx's proof of [33, Lemma 4.12] by using the notion of strong consistency instead of plain consistency and is provided in the full version of the paper. Now we are ready to prove our main theorem.

Proof of Theorem 3.5. We fix $c = (1 + \delta)w$ and let ϵ be the minimum of $(1 - 1/(1 + \delta))^4$ and $|\text{vars}(\varphi)|^{-4}$. Suppose that φ is of the form $\exists x_1 \cdots \exists x_k \tilde{\varphi}$ where $\tilde{\varphi}$ is quantifier-free. We apply Lemma 4.5 to $\tilde{\varphi}$, \mathcal{A} , c , ϵ to obtain in time $O(f(\varphi)m^{2c})$ a sequence of $\ell \leq f(\varphi)$ strongly m^c -consistent ϵ -uniform refinements $(\mathfrak{s}_i, \mathcal{B}_i)$ such that $\llbracket \tilde{\varphi} \rrbracket^{\mathcal{A}}$ is the disjoint union of $\llbracket \tilde{\varphi}_{\mathfrak{s}_1} \rrbracket^{\mathcal{B}_1}$, \dots , $\llbracket \tilde{\varphi}_{\mathfrak{s}_\ell} \rrbracket^{\mathcal{B}_\ell}$. By Lemma 4.6 we have $g_{\mathfrak{s}_i, \mathcal{B}_i} \in \mathcal{S}(\tilde{\varphi}) = \mathcal{S}(\varphi)$ for every $i \in [\ell]$. Hence, by the definition of free-connex submodular width (5), we know that there is a free-connex tree decomposition (T_i, χ_i) of φ such that $g_{\mathfrak{s}_i, \mathcal{B}_i}(\chi_i(t)) \leq w$ for every $t \in V(T_i)$. Note that by the choice of c , ϵ and the non-negativity of h (see Lemma 4.6) we have

$$w = c/(1 + \delta) \leq (1 - \epsilon^{1/4}) \cdot c < (1 - \epsilon^{1/3}) \cdot c + h(U). \quad (19)$$

Hence, $g_{\mathfrak{s}_i, \mathcal{B}_i}(U) \leq w$ implies $U \in \mathfrak{s}$ and therefore $|\mathbf{R}_U^{\mathcal{B}_i}| = |\llbracket \varphi_{\mathfrak{s}_i} \rrbracket^{\mathcal{B}_i}| \leq m^c$ by (9) and (11). Thus, every bag of the free-connex tree-decomposition (T_i, χ_i) is small in the i th refinement. However, (T_i, χ_i) is a tree-decomposition of φ , but not necessarily of $\varphi_{\mathfrak{s}_i}$! In fact, $\varphi_{\mathfrak{s}_i}$ can be very dense, e. g., if $\mathfrak{s}_i = 2^{\text{vars}(\varphi)}$. To take care of this, we thin out the refinement and only keep those atoms and relations that correspond to bags of the decomposition. In particular, for every $i \in [\ell]$ we define $\tilde{\psi}_i := \bigwedge_{t \in V(T_i)} R_{\chi_i(t)}(x_{\chi_i(t)})$ and let $\psi_i := \exists x_1 \cdots \exists x_k \tilde{\psi}_i$ be the quantified version. Note that ψ_i is a free-connex acyclic CQ. Additionally, we let \mathcal{C}_i be the $\sigma(\psi_i)$ -reduct of \mathcal{B}_i . We argue that $\llbracket \tilde{\varphi}_{\mathfrak{s}_i} \rrbracket^{\mathcal{B}_i} \subseteq \llbracket \tilde{\psi}_i \rrbracket^{\mathcal{C}_i} \subseteq \llbracket \tilde{\varphi} \rrbracket^{\mathcal{A}}$. The first inclusion holds because $\tilde{\varphi}_{\mathfrak{s}_i}$ and \mathcal{B}_i refine $\tilde{\psi}_i$ and \mathcal{C}_i . The second inclusion holds because every atom from $\tilde{\varphi}$ is contained in a bag of the decomposition and is hence covered by an atom in $\tilde{\psi}_i$ because of consistency. It therefore also follows that $\pi_F(\llbracket \tilde{\varphi}_{\mathfrak{s}_i} \rrbracket^{\mathcal{B}_i}) \subseteq \pi_F(\llbracket \tilde{\psi}_i \rrbracket^{\mathcal{C}_i}) \subseteq \pi_F(\llbracket \tilde{\varphi} \rrbracket^{\mathcal{A}})$ for $F := \text{free}(\varphi)$, and hence $\llbracket \varphi_{\mathfrak{s}_i} \rrbracket^{\mathcal{B}_i} \subseteq \llbracket \psi_i \rrbracket^{\mathcal{C}_i} \subseteq \llbracket \varphi \rrbracket^{\mathcal{A}}$. Overall, we have that $\llbracket \varphi \rrbracket^{\mathcal{A}} = \bigcup_{i \in [\ell]} \llbracket \psi_i \rrbracket^{\mathcal{C}_i}$, where the union is not necessarily disjoint, each ψ_i is free-connex acyclic, and $\|\mathcal{C}_i\| = O(|\text{vars}(\varphi)|^2 m^{(1+\delta)w})$. By combining Theorem 3.3 with Theorem 2.1, the theorem follows. \blacktriangleleft

5 Final Remarks

In this paper, we have investigated the enumeration complexity of conjunctive queries and have shown that every class of conjunctive queries of bounded free-connex submodular width admits constant delay enumeration with FPT-preprocessing. These are by now the largest classes of CQs that allow efficient enumeration in this sense.

For quantifier-free self-join-free CQs this upper bound is matched by Marx's lower bound [33]. I. e., recursively enumerable classes of quantifier-free self-join-free CQs of unbounded free-connex submodular width do not admit constant delay enumeration after FPT-preprocessing (assuming the exponential time hypothesis ETH).

A major future task is to obtain a complete dichotomy, or at least one for all self-join-free CQs. The gray-zone for the latter are classes of CQs that have bounded submodular width, but unbounded free-connex submodular width. An intriguing example in this gray-zone is the k -star query with a quantified center, i. e., the query ψ_k of the form $\exists z \bigwedge_{i=1}^k R_i(z, x_i)$. Here we have $subw(\psi_k) = 1$ and $fc-subw(\psi_k) = k$. It is open whether the class $\Psi = \{\psi_k : k \in \mathbb{N}_{\geq 1}\}$ admits constant delay enumeration with FPT-preprocessing.

References

- 1 Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases*. Addison-Wesley, 1995. URL: <http://webdam.inria.fr/Alice/>.
- 2 Mahmoud Abo Khamis, Hung Q. Ngo, and Dan Suciu. What do Shannon-type inequalities, submodular width, and disjunctive datalog have to do with one another? *CoRR*, abs/1612.02503, 2016. [arXiv:1612.02503](https://arxiv.org/abs/1612.02503).
- 3 Mahmoud Abo Khamis, Hung Q. Ngo, and Dan Suciu. What Do Shannon-type Inequalities, Submodular Width, and Disjunctive Datalog Have to Do with One Another? In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems PODS 2017*, pages 429–444, 2017. doi:10.1145/3034786.3056105.
- 4 Noga Alon, Raphael Yuster, and Uri Zwick. Finding and Counting Given Length Cycles. *Algorithmica*, 17(3):209–223, 1997. doi:10.1007/BF02523189.
- 5 Antoine Amarilli, Pierre Bourhis, and Stefan Mengel. Enumeration on Trees under Relabelings. In *21st International Conference on Database Theory, ICDT 2018, March 26-29, 2018, Vienna, Austria*, pages 5:1–5:18, 2018. doi:10.4230/LIPIcs.ICDT.2018.5.
- 6 Antoine Amarilli, Pierre Bourhis, Stefan Mengel, and Matthias Niewerth. Enumeration on Trees with Tractable Combined Complexity and Efficient Updates. *CoRR*, abs/1812.09519, 2018. [arXiv:1812.09519](https://arxiv.org/abs/1812.09519).
- 7 Antoine Amarilli, Pierre Bourhis, Stefan Mengel, and Matthias Niewerth. Constant-Delay Enumeration for Nondeterministic Document Spanners. In *22nd International Conference on Database Theory, ICDT 2019, March 26-28, 2019, Lisbon, Portugal*, pages 22:1–22:19, 2019. doi:10.4230/LIPIcs.ICDT.2019.22.
- 8 Guillaume Bagan. MSO queries on tree decomposable structures are computable with linear delay. In *Computer Science Logic, 20th International Workshop, CSL 2006, 15th Annual Conference of the EACSL, Szeged, Hungary, September 25-29, 2006, Proceedings*, pages 167–181, 2006. doi:10.1007/11874683_11.
- 9 Guillaume Bagan. *Algorithmes et complexité des problèmes d'énumération pour l'évaluation de requêtes logiques. (Algorithms and complexity of enumeration problems for the evaluation of logical queries)*. PhD thesis, University of Caen Normandy, France, 2009. URL: <https://tel.archives-ouvertes.fr/tel-00424232>.
- 10 Guillaume Bagan, Arnaud Durand, and Etienne Grandjean. On Acyclic Conjunctive Queries and Constant Delay Enumeration. In *Proceedings of the 16th Annual Conference of the EACSL, CSL'07, Lausanne, Switzerland, September 11–15, 2007*, pages 208–222, 2007. doi:10.1007/978-3-540-74915-8_18.
- 11 Catriel Beeri, Ronald Fagin, David Maier, and Mihalis Yannakakis. On the Desirability of Acyclic Database Schemes. *J. ACM*, 30(3):479–513, 1983. doi:10.1145/2402.322389.

- 12 Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering Conjunctive Queries under Updates. In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS'17, Chicago, IL, USA, May 14–19, 2017*, pages 303–318, 2017. Full version available at <http://arxiv.org/abs/1702.06370>. doi:10.1145/3034786.3034789.
- 13 Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering FO+MOD Queries Under Updates on Bounded Degree Databases. In Michael Benedikt and Giorgio Orsi, editors, *20th International Conference on Database Theory, ICDT 2017, March 21–24, 2017, Venice, Italy*, volume 68 of *LIPICs*, pages 8:1–8:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICs.ICDT.2017.8.
- 14 Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering UCQs under Updates and in the Presence of Integrity Constraints. In *21st International Conference on Database Theory, ICDT 2018, March 26–29, 2018, Vienna, Austria*, pages 8:1–8:19, 2018. doi:10.4230/LIPICs.ICDT.2018.8.
- 15 Philip A. Bernstein and Nathan Goodman. Power of Natural Semijoins. *SIAM J. Comput.*, 10(4):751–771, 1981. doi:10.1137/0210059.
- 16 Johann Brault-Baron. *De la pertinence de l'énumération : complexité en logiques propositionnelle et du premier ordre. (The relevance of the list: propositional logic and complexity of the first order)*. PhD thesis, University of Caen Normandy, France, 2013. URL: <https://tel.archives-ouvertes.fr/tel-01081392>.
- 17 Johann Brault-Baron. Hypergraph Acyclicity Revisited. *ACM Comput. Surv.*, 49(3):54:1–54:26, 2016. doi:10.1145/2983573.
- 18 Shaleen Deep and Paraschos Koutris. Compressed Representations of Conjunctive Query Results. In *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Houston, TX, USA, June 10–15, 2018*, pages 307–322, 2018. doi:10.1145/3196959.3196979.
- 19 Arnaud Durand and Etienne Grandjean. First-order queries on structures of bounded degree are computable with constant delay. *ACM Trans. Comput. Log.*, 8(4), 2007. doi:10.1145/1276920.1276923.
- 20 Arnaud Durand, Nicole Schweikardt, and Luc Segoufin. Enumerating answers to first-order queries over databases of low degree. In *Proceedings of the 33rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS'14, Snowbird, UT, USA, June 22–27, 2014*, pages 121–131, 2014. doi:10.1145/2594538.2594539.
- 21 Arnaud Durand and Yann Strozecki. Enumeration Complexity of Logical Query Problems with Second-order Variables. In *Computer Science Logic, 25th International Workshop / 20th Annual Conference of the EACSL, CSL 2011, September 12–15, 2011, Bergen, Norway, Proceedings*, pages 189–202, 2011. doi:10.4230/LIPICs.CSL.2011.189.
- 22 Georg Gottlob, Gianluigi Greco, Nicola Leone, and Francesco Scarcello. Hypertree Decompositions: Questions and Answers. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2016, San Francisco, CA, USA, June 26 - July 01, 2016*, pages 57–74, 2016. doi:10.1145/2902251.2902309.
- 23 Georg Gottlob, Nicola Leone, and Francesco Scarcello. Hypertree Decompositions and Tractable Queries. *J. Comput. Syst. Sci.*, 64(3):579–627, 2002. doi:10.1006/jcss.2001.1809.
- 24 Martin Grohe and Dániel Marx. Constraint Solving via Fractional Edge Covers. *ACM Trans. Algorithms*, 11(1):4:1–4:20, 2014. doi:10.1145/2636918.
- 25 Muhammad Idris, Martín Ugarte, and Stijn Vansummeren. The Dynamic Yannakakis Algorithm: Compact and Efficient Query Processing Under Updates. In *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD Conference 2017, Chicago, IL, USA, May 14–19, 2017*, pages 1259–1274, 2017. doi:10.1145/3035918.3064027.
- 26 Ahmet Kara and Dan Olteanu. Covers of Query Results. In *21st International Conference on Database Theory, ICDT 2018, March 26–29, 2018, Vienna, Austria*, pages 16:1–16:22, 2018. doi:10.4230/LIPICs.ICDT.2018.16.

- 27 Wojciech Kazana and Luc Segoufin. First-order query evaluation on structures of bounded degree. *Logical Methods in Computer Science*, 7(2), 2011. doi:10.2168/LMCS-7(2:20)2011.
- 28 Wojciech Kazana and Luc Segoufin. Enumeration of first-order queries on classes of structures with bounded expansion. In *Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2013, New York, NY, USA - June 22 - 27, 2013*, pages 297–308, 2013. doi:10.1145/2463664.2463667.
- 29 Wojciech Kazana and Luc Segoufin. Enumeration of monadic second-order queries on trees. *ACM Trans. Comput. Log.*, 14(4):25:1–25:12, 2013. doi:10.1145/2528928.
- 30 Dietrich Kuske and Nicole Schweikardt. First-order logic with counting. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*, pages 1–12, 2017. doi:10.1109/LICS.2017.8005133.
- 31 Katja Losemann and Wim Martens. MSO queries on trees: enumerating answers under updates. In *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*, pages 67:1–67:10, 2014. doi:10.1145/2603088.2603137.
- 32 Dániel Marx. Tractable Structures for Constraint Satisfaction with Truth Tables. *Theory Comput. Syst.*, 48(3):444–464, 2011. doi:10.1007/s00224-009-9248-9.
- 33 Dániel Marx. Tractable Hypergraph Properties for Constraint Satisfaction and Conjunctive Queries. *Journal of the ACM (JACM)*, Volume 60, Issue 6, Article No. 42, November 2013. doi:10.1145/2535926.
- 34 Matthias Niewerth. MSO queries on trees: Enumerating answers under updates using forest algebras. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 769–778, 2018. doi:10.1145/3209108.3209144.
- 35 Matthias Niewerth and Luc Segoufin. Enumeration of MSO Queries on Strings with Constant Delay and Logarithmic Updates. In *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Houston, TX, USA, June 10-15, 2018*, pages 179–191, 2018. doi:10.1145/3196959.3196961.
- 36 Dan Olteanu and Maximilian Schleich. Factorized Databases. *SIGMOD Record*, 45(2):5–16, 2016. doi:10.1145/3003665.3003667.
- 37 Dan Olteanu and Jakub Závodný. Size Bounds for Factorised Representations of Query Results. *ACM Trans. Database Syst.*, 40(1):2:1–2:44, 2015. doi:10.1145/2656335.
- 38 Francesco Scarcello. From Hypertree Width to Submodular Width and Data-dependent Structural Decompositions. In *Proceedings of the 26th Italian Symposium on Advanced Database Systems, Castellaneta Marina (Taranto), Italy, June 24-27, 2018.*, 2018. URL: <http://ceur-ws.org/Vol-2161/paper24.pdf>.
- 39 Nicole Schweikardt, Luc Segoufin, and Alexandre Vigny. Enumeration for FO Queries over Nowhere Dense Graphs. In *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Houston, TX, USA, June 10-15, 2018*, pages 151–163, 2018. doi:10.1145/3196959.3196971.
- 40 Luc Segoufin. Constant Delay Enumeration for Conjunctive Queries. *SIGMOD Record*, 44(1):10–17, 2015. doi:10.1145/2783888.2783894.
- 41 Luc Segoufin and Alexandre Vigny. Constant Delay Enumeration for FO Queries over Databases with Local Bounded Expansion. In *20th International Conference on Database Theory, ICDT 2017, March 21–24, 2017, Venice, Italy*, pages 20:1–20:16, 2017. doi:10.4230/LIPIcs.ICDT.2017.20.
- 42 Mihalis Yannakakis. Algorithms for Acyclic Database Schemes. In *Very Large Data Bases, 7th International Conference, September 9-11, 1981, Cannes, France, Proceedings*, pages 82–94, 1981.