Counting Answers to Unions of Conjunctive Queries: Natural Tractability Criteria and Meta-Complexity

JACOB FOCKE*, CISPA Helmholtz Center for Information Security, Germany

LESLIE ANN GOLDBERG*, Department of Computer Science, University of Oxford, UK

MARC ROTH*, School for Electronic Engineering and Computer Science, Queen Mary University of London, UK

STANISLAV ŽIVNÝ*, Department of Computer Science, University of Oxford, UK

We study the problem of counting answers to unions of conjunctive queries (UCQs) under structural restrictions on the input query. Concretely, given a class C of UCQs, the problem #UCQ(C) provides as input a UCQ $\Psi \in C$ and a database $\mathcal{D}$ and the problem is to compute the number of answers of Ψ in $\mathcal{D}$.

Chen and Mengel [PODS'16] have shown that for any recursively enumerable class C, the problem #UCQ(C) is either fixed-parameter tractable or hard for one of the parameterised complexity classes W[1] or #W[1]. However, their tractability criterion is unwieldy in the sense that, given any concrete class C of UCQs, it is not easy to determine how hard it is to count answers to queries in C. Moreover, given a single specific UCQ Ψ, it is not easy to determine how hard it is to count answers to Ψ.

In this work, we address the question of finding a *natural* tractability criterion: The combined conjunctive query of a UCQ $\Psi = \varphi_1 \vee \cdots \vee \varphi_\ell$ is the conjunctive query $\wedge(\Psi) = \varphi_1 \wedge \cdots \wedge \varphi_\ell$. We show that under natural closure properties of C, the problem #UCQ(C) is fixed-parameter tractable if and only if the combined conjunctive queries of UCQs in C, and their contracts, have bounded treewidth. A contract of a conjunctive query is an augmented structure, taking into account how the quantified variables are connected to the free variables — if all variables are free, then a conjunctive query is equal to its contract; in this special case the criterion for fixed-parameter tractability of #UCQ(C) thus simplifies to the combined queries having bounded treewidth.

Finally, we give evidence that a closure property on C is necessary for obtaining a natural tractability criterion: We show that even for a single UCQ Ψ, the meta problem of deciding whether #UCQ({Ψ}) can be solved in time $O(|\mathcal{D}|^d)$ is NP-hard for any fixed $d \geq 1$. Moreover, we prove that a known exponential-time algorithm for solving the meta problem is optimal under assumptions from fine-grained complexity theory. As a corollary of our reduction, we also establish that approximating the Weisfeiler-Leman-Dimension of a UCQ is NP-hard.

Authors' addresses: Jacob Focke, jacob.focke@cispa.de, CISPA Helmholtz Center for Information Security, Saarbrücken, Germany; Leslie Ann Goldberg, leslie.goldberg@cs.ox.ac.uk, Department of Computer Science, University of Oxford, Oxford, UK; Marc Roth, m.roth@qmul.ac.uk, School for Electronic Engineering and Computer Science, Queen Mary University of London, London, UK; Stanislav Živný, standa.zivny@cs.ox.ac.uk, Department of Computer Science, University of Oxford, Oxford, UK.

---

1 INTRODUCTION

Conjunctive queries are among the most fundamental and well-studied objects in database theory [2, 20, 47, 48, 63, 64]. A conjunctive query (CQ) φ with free variables $X = \{x_1, \ldots, x_k\}$ and quantified variables $Y = \{y_1, \ldots, y_d\}$ is of the form

$$\varphi(X) = \exists Y \, R_1(\mathbf{t}_1) \wedge \ldots \wedge R_n(\mathbf{t}_n),$$

where $R_1, \ldots, R_n$ are relational symbols and each $\mathbf{t}_i$ is a tuple of variables from $X \cup Y$. A database $\mathcal{D}$ consists of a set of elements $U(\mathcal{D})$, denoted the *universe* of $\mathcal{D}$, and a set of relations over this universe. The corresponding relation symbols are the *signature* of $\mathcal{D}$. If $R_1, \ldots, R_n$ are in the signature of $\mathcal{D}$ then an *answer* of φ in $\mathcal{D}$ is an assignment $a \colon X \to U(\mathcal{D})$ that has an extension to the existentially quantified variables Y that agrees with all the relations $R_1, \ldots, R_n$. Even more expressive is a union of conjunctive queries (UCQ). Such a union is of the form

$$\Psi(X) = \varphi_1(X) \vee \ldots \vee \varphi_\ell(X),$$

where each $\varphi_i(X)$ is a CQ with free variables X. An answer to Ψ is then any assignment that is answer to at least one of the CQs in the union.

Since evaluating a given CQ on a given database is NP-complete [20] a lot of research focused on finding tractable classes of CQs. A fundamental result by Grohe, Schwentick, and Segoufin [42] established that the tractability of evaluating all CQs of bounded arity whose Gaifman graph is in some class of graphs C depends on whether or not the treewidth in C is bounded.

More generally, finding an answer to a conjunctive query can be cast as finding a (partial) homomorphism between relational structures, and therefore is closely related to the framework of constraint satisfaction problems. In this setting, Grohe [40] showed that treewidth modulo homomorphic equivalence is the right criterion for tractability. There is also an important line of work [38, 41, 50] culminating in the fundamental work by Marx [51] that investigates the parameterised complexity for classes of queries with unbounded arity. In general, tractability of conjunctive queries is closely related to how "tree-like" or close to acyclic they are.

Counting answers to CQs has also received significant attention in the past [3, 22, 28, 30, 31, 39, 58]. Chen and Mengel [21] gave a complete classification for the counting problem on classes of CQs (with bounded arity) in terms of a natural criterion loosely based on treewidth. They present a trichotomy into fixed-parameter tractable, W[1]-complete, and #W[1]-hard cases. In subsequent work [22], this classification was extended to unions of conjunctive queries (and to even more general queries in [30]). However, for UCQs, the established criteria for tractability and intractability are implicit (see [22, Theorems 3.1 and 3.2]) in the sense that, given a specific UCQ Ψ, it is not at all clear how hard it is to count answers to Ψ based on the criteria in [22]. To make this more precise: It is not even clear whether we can, in polynomial time in the size of Ψ, determine whether answers to Ψ can be counted in linear time in the input database.

1.1 Our contributions

With the goal of establishing a more practical tractability criterion for counting answers to UCQs, we explore the following two main questions in this work:

Q1) Is there a natural criterion that captures the fixed-parameter tractability of counting answers to a class of UCQs, parameterised by the size of the query?

Q2) Is there a natural criterion that captures whether counting answers to a single fixed UCQ is linear-time solvable (in the size of a given database)?

Question Q1): Fixed-Parameter Tractability. For a class C of UCQs, we consider the problem #UCQ(C) that takes as input a UCQ Ψ from C and a database $\mathcal{D}$, and asks for the number $\mathsf{ans}(\Psi \to \mathcal{D})$ of answers of Ψ in $\mathcal{D}$. We assume that the arity of the UCQs in C is bounded, that is, there is constant c such that each relation that appears in some query in C has arity at most c. As explained earlier, due to a result of Chen and Mengel [22], there is a known but rather unwieldy tractability criterion for #UCQ(C), when the problem is parameterised by the size of the query. On a high level, the number of answers of a UCQ Ψ in a given database can be expressed as a finite linear combination of CQ answer counts, using the principle of inclusion-exclusion. This means that

$$\mathsf{ans}(\Psi \to \mathcal{D}) = \sum_i c_i \cdot \mathsf{ans}(\varphi_i \to \mathcal{D}),$$

where each φ_i is simply a conjunctive query (and not a union thereof). We refer to this linear combination as the *CQ expansion* of Ψ. Chen and Mengel showed that the parameterised complexity of computing $\mathsf{ans}(\Psi \to \mathcal{D})$ is guided by the hardest term in the respective CQ expansion. The complexity of computing these terms is simply the complexity of counting the answers of a conjunctive query, and this is well understood [21]. Hence, the main challenge for this approach is to understand the linear combination, i.e., to understand for which CQs the corresponding coefficients are non-zero. The problem is that the coefficients c_i of these linear combinations are alternating sums, which in similar settings have been observed to encode algebraic and even topological invariants [60]. This makes it highly non-trivial to determine which CQs actually contribute to the linear combination. We introduce the concepts required to state this classification informally, the corresponding definitions are given in Section 2.

We first give more details about the result of [21]. Let $\Gamma(C)$ be the class of those conjunctive queries that contribute to the CQ expansion of at least one UCQ in C, and that additionally are what we call #minimal. Intuitively, a conjunctive query φ is #minimal if there is no proper subquery φ' of φ that has the same number of answers as φ in every given database. Then the tractability criterion depends on the treewidth of the CQs in $\Gamma(C)$. It also depends on the treewidth of the corresponding class contract($\Gamma(C)$) of contracts (formally defined in Definition 18), which is an upper bound of what is called the "star size" in [31] and the "dominating star size" in [30]. Here is the formal statement of the known dichotomy for #UCQ(C).

Theorem 1 ([22])**.** *Let C be a recursively enumerable class of UCQs of bounded arity. If the treewidth of $\Gamma(C)$ and of* contract($\Gamma(C)$) *is bounded, then #UCQ(C) is fixed-parameter tractable. Otherwise, #UCQ(C) is* W[1]*-hard.*

We investigate under which conditions this dichotomy can be simplified. We show that for large classes of UCQs there is actually a much more natural tractability criterion that does not rely on $\Gamma(C)$, i.e., here the computation of the coefficients of the linear combinations as well as the concept of #minimality do not play a role. We first show a simpler classification for UCQs without existential quantifiers. To state the results we require some additional definitions: The *combined query* $\bigwedge(\Psi)$ of a UCQ $\Psi(X) = \varphi_1(X) \vee \cdots \vee \varphi_\ell(X)$ is the conjunctive query obtained from Ψ by

replacing each disjunction by a conjunction, that is

$$\bigwedge (\Psi) = \varphi_1(X) \wedge \cdots \wedge \varphi_\ell(X) \,.$$

Given a class of UCQs C, we set $\bigwedge (C) = \{\bigwedge (\Psi) \mid \Psi \in C\}$.

It will turn out that the structure of the class of combined queries $\bigwedge (C)$ determines the complexity of counting answers to UCQs in C, given that C has the following natural closure property: We say that C is *closed under deletions* if, for all $\Psi(X) = \varphi_1(X) \vee \cdots \vee \varphi_\ell(X)$ and for every $J \subseteq [\ell]$, the subquery $\bigvee_{j \in J} \varphi_j(X)$ is also contained in C. For example, any class of UCQs defined solely by the conjunctive queries admissible in the unions (such as unions of acyclic conjunctive queries) is closed under deletions. The following classification resolves the complexity of counting answers to UCQs in classes that are closed under deletions; we will see later that the closedness condition is necessary. Moreover, the tractability criterion depends solely on the structure of the combined query, and not on the terms in the CQ expansion, thus yielding, as desired, a much more concise and natural characterisation. As mentioned earlier, we first state the classification for quantifier-free UCQs.

Theorem 2. *Let C be recursively enumerable class of quantifier-free UCQs of bounded arity. If $\bigwedge (C)$ has bounded treewidth then #UCQ(C) is fixed-parameter tractable. If $\bigwedge (C)$ has unbounded treewidth and C is closed under deletions then #UCQ(C) is $\mathrm{W}[1]$-hard.*

We emphasise here that Theorem 2 is in terms of the simpler object $\bigwedge (C)$ instead of the complicated object $\Gamma(C)$.

If we allow UCQs with quantified variables in the class C then the situation becomes more intricate. Looking for a simple tractability criterion that describes the complexity of #UCQ(C) solely in terms of $\bigwedge (C)$ requires some additional effort. First, for a UCQ Ψ that has quantified variables, contract(Ψ) is not necessarily the same as Ψ, and therefore the treewidth of the contracts also plays a role. Moreover, the matching lower bound requires some conditions in addition to being closed under deletions. Nevertheless, our result is in terms of the simpler objects $\bigwedge (C)$ and contract($\bigwedge (C)$) rather than the more complicated $\Gamma(C)$ and contract($\Gamma(C)$). For Theorem 3, recall that a conjunctive query is *self-join-free* if each relation symbol occurs in at most one atom of the query.

Theorem 3. *Let C be a recursively enumerable class of UCQs of bounded arity. If $\bigwedge (C)$ and contract($\bigwedge (C)$) have bounded treewidth then #UCQ(C) is fixed-parameter tractable. Otherwise, if (I)–(III) are satisfed, then #UCQ(C) is $\mathrm{W}[1]$-hard.*

 (I) *C is closed under deletions.*
 (II) *The number of existentially quantified variables of queries in C is bounded.*
 (III) *The UCQs in C are unions of self-join-free conjunctive queries.*

In Appendix A we show that Theorem 3 is tight in the sense that, if any of these conditions is dropped, there are counterexamples to the claim that tractability is guided solely by $\bigwedge (C)$ and contract($\bigwedge (C)$).

Question Q2): Linear-Time Solvability. Now we turn to the question of linear-time solvability for a single fixed UCQ. The huge size of databases in modern applications motivates the question of which query problems are actually linear-time solvable. Along these lines, there is a lot of research for enumeration problems [7, 10, 11, 14, 19, 64].

The question whether counting answers to a conjunctive query φ can be achieved in time linear in the given database has been studied previously [52]. The corresponding dichotomy is well

known and was discovered multiple times by different authors in different contexts.[1] In these results, the tractability criterion is whether φ is *acyclic*, i.e., whether it has a join tree (see [37]). The corresponding lower bounds are conditioned on a widely used complexity assumption from fine-grained complexity, namely the Triangle Conjecture. We define all of the complexity assumptions that we use in this work in Section 2. There we also formally define the size of a database (as the sum of the size of its signature, its universe, and its relations).

It is well-known that, counting answers to quantifier-free conjunctive queries can be done in linear time if and only if the query is acyclic. The "only if" part relies on hardness assumptions from fine-grained complexity theory. Concretely, we have

Theorem 4 (See Theorem 12 in [16], and [10, 53]). *Let φ be a quantifier-free conjunctive query. If φ is acyclic, then the number of answers of φ in a given database $\mathcal{D}$ can be computed in time linear in the size of $\mathcal{D}$. Otherwise, if φ has arity* 2*, then the number of answers cannot be computed in linear time in the size of $\mathcal{D}$, unless the Triangle Conjecture fails.*

We point out that there is also a conditional lower bound for cyclic quantifier-free conjunctive queries of arity larger than 2: assuming the Hyperclique Hypothesis, counting answers to such queries can not be done in linear time — however, since the case of arity 2 is sufficient for our lower bounds and since this work already includes various lower bound assumptions, we decided not to go into the details of the higher arity case and instead refer the reader to an excellent survey by Mengel [54] and [53, Theorem 3.8].

We note that the previous theorem is false if quantified variables were allowed as this would require the consideration of semantic acyclicity[2] (see [9]).

Theorem 4 yields an efficient way to check whether counting answers to a quantifier-free conjunctive query φ can be done in linear time: Just check whether φ is acyclic (in polynomial time, see for instance [37]). We investigate the corresponding question for *unions* of conjunctive queries. In stark contrast to Theorem 4, we show that there is no efficiently computable criterion that determines the linear-time tractability of counting answers to *unions* of conjunctive queries, unless some conjectures of fine-grained complexity theory fail.

We first observe that, as in the investigation of question Q1), one can obtain a criterion for linear-time solvability by expressing UCQ answer counts as linear combinations of CQ answer counts. Concretely, by a straightforward extension of previous results, we show that, assuming the Triangle Conjecture, a linear combination of CQ answer counts can be computed in linear time if and only if the answers to each #minimal CQ in the linear combination can be computed in linear time, that is, if each such CQ is acyclic. However, this criterion is again unwieldy in the sense that, for all we know, it may take time exponential in the size of the respective UCQ to determine whether this criterion holds.

In view of our results for question Q1) about fixed-parameter tractability, one might suspect that a more natural and simpler tractability criterion exists. However, it turns out that even under strong restrictions on the UCQs that we consider, an efficiently computable criterion is unlikely. We make this formal by studying the following meta problem.[3]

---

[1]We remark that [10, Theorem 7] focuses on the special case of graphs and *near* linear time algorithms. However, in the word RAM model with $O(\log n)$ bits, a linear time algorithm is possible [19].

[2]A conjunctive query is semantically acyclic if and only if its #core (Definition 17) is acyclic.

[3]For the question Q1), considering a similar meta problem is not feasible as, in this case, the meta problem takes as input a class of graphs. If such a class were encoded as a Turing machine, the meta problem would be undecidable by Rice's theorem.

Name: META
Input: A union Ψ of quantifier-free conjunctive queries.
Output: Is it possible to count answers to Ψ in time linear in the size of any given database $\mathcal{D}$.

Restricting the input of META to quantifier-free queries is sensible as, without this restriction, the meta problem is known to be NP-hard even for conjunctive queries: If all variables are existentially quantified, then evaluating a conjunctive query can be done in linear time if and only if the query is semantically acyclic [64] (the "only if" relies on standard hardness assumptions). However, verifying whether a conjunctive query is semantically acyclic is already NP-hard [9]. In contrast, when restricted to quantifier-free conjunctive queries, the problem META is polynomially-time solvable according to Theorem 4.

We can now state our main result about the complexity of META. The hardness results hold under substantial additional input restrictions, which make these results stronger.

Theorem 5. *Assume that the Triangle Conjecture is true. When restricted to input queries of arity at most* 2, *META can be solved in time* $2^{O(\ell)} \cdot |\Psi|^{\text{poly}(\log |\Psi|)}$, *where ℓ is the number of conjunctive queries in the union. Moreover,* META *is* NP-*hard and the following conditional lower bounds apply:*

- *If ETH is true, then* META *cannot be solved in time* $2^{o(\ell)}$.
- *If the non-uniform ETH is true then* META $\notin \bigcap_{\varepsilon > 0} \text{DTime}(2^{\varepsilon \cdot \ell})$.

The lower bounds remain true even if Ψ is a union of self-join-free and acyclic conjunctive queries over a binary signature (that is, of arity 2).

We make some remarks about Theorem 5. First, it may seem counterintuitive that the algorithmic part of this result relies on some lower bound conjectures. This is explained by the fact that an algorithmic result for META is actually a classification result for the underlying counting problem. The lower bound conjectures are the reason that the algorithm for META can answer that a linear-time algorithm is *not* possible for certain UCQs. We also point out that the upper bound remains true for arbitrary arities if the Hyperclique Hypothesis is assumed (see e.g. the full version of Mengel's survey [53, Hypothesis 3]); as mentioned earlier, we omit the details in order to avoid formally introducing yet another lower bound assumption from fine-grained complexity theory.

Second, while for counting the answers to a CQ in linear time the property of being acyclic is the right criterion, note that for unions of CQs, acyclicity is not even sufficient for tractability. Even when restricted to unions of acyclic conjunctive queries, the meta problem is NP-hard.

Third, we elaborate on the idea that we use to prove Theorem 5. As mentioned before, the algorithmic part of Theorem 5 comes from the well-known technique of expressing UCQ answer counts in terms of linear combinations of CQ answer counts, and establishing a corresponding complexity monotonicity property, see Section 2.4. The more interesting result is the hardness part. Here we discover a connection between the meta question stated in META, and a topological invariant, namely, the question whether the reduced Euler characteristic of a simplicial complex is non-zero. It is known that simplicial complexes with non-vanishing reduced Euler characteristic are evasive, and as such this property is also related to Karp's Evasiveness Conjecture (see e.g. the excellent survey of Miller [55]). We use the known fact that deciding whether the reduced Euler characteristic is vanishing is NP-hard [61]. Roughly, the reduction works as follows. Given some simplicial complex Δ, we carefully define a UCQ Ψ_Δ in such a way that only one particular term in the CQ expansion of Ψ_Δ determines the linear-time tractability of counting answers to Ψ_Δ. However, the coefficient of this term is zero precisely if the reduced Euler characteristic of Δ is vanishing.

Simplicial complexes also appeared in a related context in a work by Roth and Schmitt [60]. They show a connection between the complexity of counting induced subgraphs that fulfil some

graph property and the question whether a simplicial complex associated with this graph property is non-zero. To solve their problem, it suffices to consider simplicial *graph* complexes, which are special simplicial complexes whose elements are subsets of the edges of a complete graph, and to encode these as induced subgraph counting problems. In contrast, to get our result we must encode arbitrary abstract simplicial complexes as UCQs and to show how to transfer the question about their Euler characteristic to a question about linear-time solvability of UCQs.

It turns out that, as additional consequences of our reduction in the proof of Theorem 5, we also obtain lower bounds for (approximately) computing the so-called Weisfeiler-Leman-dimension of a UCQ.

Consequences for the Weisfeiler-Leman-dimension of quantifier-free UCQs. During the last decade we have witnessed a resurge in the study of the *Weisfeiler-Leman-dimension* of graph classes and graph parameters [4, 8, 29, 36, 46, 56]. The Weisfeiler-Leman algorithm (WL-algorithm) and its higher-dimensional generalisations are important heuristics for graph isomorphism; for example, the 1-dimensional WL-algorithm is equivalent to the method of colour-refinement. We refer the reader to e.g. the EATCS Bulletin article of Arvind [4] for a concise and self-contained introduction; however, in this work we will use the WL-algorithm only in a black-box manner.

For each positive integer k, we say that two graphs G_1 and G_2 are *k-WL equivalent*, denoted by $G_1 \cong_k G_2$, if they cannot be distinguished by the k-dimensional WL-algorithm. A graph parameter π is called *k-WL invariant* if $G_1 \cong_k G_2$ implies $\pi(G_1) = \pi(G_2)$. Moreover, the *WL-dimension* of π is the minimum k for which π is k-WL invariant, if such a k exists, and ∞ otherwise (see e.g.[5]). The WL-dimension of a graph parameter π provides important information about the descriptive complexity of π [17]. Moreover, recent work of Morris et al. [56] shows that the WL-dimension of a graph parameter lower bounds the minimum dimension of a higher-order Graph Neural Network that computes the parameter.

The definitions of the WL-algorithm and the WL-dimension extend from graphs to labelled graphs, that is, directed multi-graphs with edge- and vertex-labels (see e.g. [49]). Formally, we say that a database is a *labelled graph* if its signature has arity at most 2, and if it contains no self-loops, that is, tuples of the form (v, v). Similarly, *(U)CQs on labelled graphs* have signatures of arity at most 2 and contain no atom of the form $R(v, v)$.

Definition 6 (WL-dimension). Let Ψ be a UCQ on labelled graphs. The *WL-dimension* of Ψ, denoted by $\dim_{\mathrm{WL}}(\Psi)$, is the minimum k such that, for any pair of labelled graphs $\mathcal{D}_1$ and $\mathcal{D}_2$ with $\mathcal{D}_1 \cong_k \mathcal{D}_2$, it holds that the number of answers to Ψ in $\mathcal{D}_1$ is the same as in $\mathcal{D}_2$. If no such k exists, then the WL-dimension is ∞.

Note that a CQ is a special case of a UCQ, so Definition 6 also applies when Ψ is a CQ φ.

It was shown very recently that the WL-dimension of a *quantifier-free conjunctive query* φ on labelled graphs is equal to the treewidth of the Gaifman graph of φ [49, 57]. Using known algorithms for computing the treewidth [15, 32] it follows that, for every fixed positive integer d, the problem of deciding whether the WL-dimension of φ is at most d can be solved in polynomial time (in the size of φ). Moreover, the WL-dimension of φ can be efficiently approximated in polynomial time.

In stark contrast, we show that the computation of the WL-dimension of a *UCQ* is much harder; in what follows, we say that S is an *f-approximation* of k if $k \leq S \leq f(k) \cdot k$.

Theorem 7. *There is an algorithm that computes a $O(\sqrt{\log k})$-approximation of the WL-dimension k of a quantifier-free UCQ on labelled graphs $\Psi = \varphi_1 \vee \cdots \vee \varphi_\ell$ in time $|\Psi|^{O(1)} \cdot O(2^\ell)$.*

Moreover, let $f : \mathbb{Z}_{>0} \to \mathbb{Z}_{>0}$ be any computable function. The problem of computing an f-approximation of $\dim_{\mathrm{WL}}(\Psi)$ given an input UCQ $\Psi = \varphi_1 \vee \cdots \vee \varphi_\ell$ is NP-hard, and, assuming ETH, an f-approximation of $\dim_{\mathrm{WL}}(\Psi)$ cannot be computed in time $2^{o(\ell)}$.

Finally, the computation of the WL-dimension of UCQs stays intractable even if we fix k.

Theorem 8. *Let k be any* fixed *positive integer. The problem of deciding whether the WL-dimension of a quantifier-free UCQ on labelled graphs $\Psi = \varphi_1 \vee \cdots \vee \varphi_\ell$ is at most k can be solved in time $|\Psi|^{O(1)} \cdot O(2^\ell)$.*

Moreover, the problem is NP*-hard and, assuming ETH, cannot be solved in time $2^{o(\ell)}$.*

1.2 Further Related Work

For exact counting it makes a substantial difference whether one wants to count answers to a conjunctive query or a union of conjunctive queries [22, 30]. However, for approximate counting, unions can generally be handled using a standard trick of Karp and Luby [45], and therefore fixed-parameter tractability results for approximately counting the answers to a conjunctive query also extend to unions of conjunctive queries [3, 34].

Counting and enumerating the answers to a union of conjunctive queries has also been studied in the context of dynamic databases [12, 13]. This line of research investigates the question whether linear-time dynamic algorithms are possible. Concretely, the question is whether, after a preprocessing step that builds a data structure in time linear in the size of the initial database, the number of answers to a fixed union of conjunctive queries can be returned in constant time with a constant-time update to the data structure, whenever there is a change to the database. Berkholz et al. show that for a conjunctive query such a linear-time algorithm is possible if and only if the CQ is *q-hierarchical* [12, Theorem 1.3]. There are acyclic CQs that are not *q*-hierarchical, for instance the query $\varphi(\{a, b, c, d\}) = E(a, b) \wedge E(b, c) \wedge E(c, d)$ is clearly acyclic — however, the sets of atoms that contain b and c, respectively, are neither comparable nor disjoint, and therefore φ is not *q*-hierarchical. So, there are queries for which counting in the static setting is easy, whereas it is hard in the dynamic setting. Berkholz et al. extend their result from CQs to UCQs [13, Theorem 4.5], where the criterion is whether the UCQ is *exhaustively q-hierarchical*. This property essentially means that, for every subset of the CQs in the union, if instead of taking the disjunction of these CQs we take the conjunction, then the resulting CQ should be *q*-hierarchical. Moreover, checking whether a CQ ϕ is *q*-hierarchical can be done in time polynomial in the size of ϕ. However, the straightforward approach of checking whether a UCQ is exhaustively *q*-hierarchical takes exponential time, and it is stated as an open problem in [13] whether this can be improved. In the dynamic setting this question remains open — however, in the static setting we show that, while for counting answers to CQs the criterion for linear-time tractability can be verified in polynomial time, this is not true for unions of conjunctive queries, subject to some complexity assumptions, as we have seen in Theorem 5.

2 PRELIMINARIES

Due to the fine-grained nature of the questions we ask in this work (e.g. linear time counting vs non-linear time counting), it is important to specify the machine model. We use the word RAM model with $O(\log n)$ bits as one of the lower bound conjectures we rely on in this work is stated for this model (see [1]). The exact model makes a difference. For example, it is possible to count answers to quantifier-free acyclic conjunctive queries in linear time in the word RAM model [19], while Turing machines only achieve near linear time (or expected linear time) [10].

2.1 Parameterised and Fine-grained Complexity Theory

A *parameterised counting problem* is a pair consisting of a function $P : \{0, 1\}^* \rightarrow \mathbb{N}$ and a computable[4] parameterisation $\kappa : \{0, 1\}^* \rightarrow \mathbb{N}$. For example, in the problem #CLIQUE the function maps an input (a graph G and a positive integer k, encoded as a string in $\{0, 1\}^*$ to the number of k-cliques in G. The parameter is K so $\kappa(G, k) = k$.

A parameterised counting problem (P, κ) is called *fixed-parameter tractable* (FPT) if there is a computable function f and an algorithm $\mathbb{A}$ that, given input x, computes $P(x)$ in time $f(\kappa(x)) \cdot |x|^{O(1)}$. We call $\mathbb{A}$ an *FPT-algorithm* for (P, κ).

A *parameterised Turing-reduction* from (P, κ) to (P', κ') is an algorithm $\mathbb{A}$ equipped with oracle access to P' that satisfies the following two constraints: (I) $\mathbb{A}$ is an FPT-algorithm for (P, κ), and (II) there is a computable function g such that, when the algorithm $\mathbb{A}$ is run with input x, every oracle query y to (P', κ') has the property that the parameter $\kappa'(y)$ is bounded by $g(\kappa(x))$. We write $(P, \kappa) \leq^{\mathsf{FPT}} (P'\kappa')$ if a parameterised Turing-reduction exists.

Evidence for the non-existence of FPT algorithms is usually given by hardness for the parameterised classes #W[1] and W[1], which can be considered to be the parameterised versions of #P and NP. The definition of those classes uses bounded-weft circuits, and we refer the interested reader e.g. to the standard textbook of Flum and Grohe [33] for a comprehensive introduction. For this work, it suffices to rely on the clique problem to establish hardness for those classes: A parameterised counting problem (P, κ) is *#W[1]-hard* if #CLIQUE $\leq^{\mathsf{FPT}} (P, \kappa)$, and it is *W[1]-hard* if CLIQUE $\leq^{\mathsf{FPT}} (P, \kappa)$, where CLIQUE is the decision version of #CLIQUE, that is, given G and k, the task is to decide whether there is at least one k-clique in G. As observed in previous works [21], if all variables are existentially quantified, the problem of counting answers to a conjunctive query actually encodes a decision problem. So it comes to no surprise that both complexity classes W[1] and #W[1] are relevant for its classification. It is well known (see e.g. [23, 24, 27] that W[1]-hard and #W[1]-hard problems are not fixed-parameter tractable, unless the Exponential Time Hypothesis fails. This hypothesis is stated as follows.

Conjecture 9 (ETH [43]). *3-SAT cannot be solved in time* $\exp(o(n))$, *where n denotes the number of variables of the input formula.*

We also rely on a non-uniform version of the Exponential Time Hypothesis, which is defined below.

Conjecture 10 (Non-uniform ETH [25]). *3-SAT* $\notin \bigcap_{\varepsilon > 0} \mathrm{DTime}(\exp(\varepsilon n))$, *where n denotes the number of variables of the input formula.*

Note that, clearly, non-uniform ETH implies ETH. Finally, for ruling out linear-time algorithms, we will rely on the Triangle Conjecture:

Conjecture 11 (Triangle Conjecture [1]). *There exists* $\gamma > 0$ *such that any (randomised) algorithm in the word RAM model with* $O(\log n)$ *bits which decides whether a graph with n vertices and m edges contains a triangle takes time at least* $\Omega(m^{1+\gamma})$ *in expectation.*

2.2 Structures, Homomorphisms, and Conjunctive Queries

A *signature* is a finite tuple $\tau = (R_1, \ldots, R_s)$ where each R_i is a relation symbol and comes with an arity a_i. The arity of a signature is the maximum arity of its relation symbols. A *structure* $\mathcal{A}$ over τ consists of a finite universe $U(\mathcal{A})$ and a relation $R_i^{\mathcal{A}}$ of arity a_i for each relation symbol R_i of τ. As

---

[4]Some authors require the parameterisation to be polynomial-time computable; see the discussion in the standard textbook of Flum and Grohe [33]. In this work, the parameter will always be the size of the input query, which can clearly be computed in polynomial time.

usual in relational algebra, we view *databases* as relational structures. We encode a structure by listing its signature, its universe and its relations. Therefore, given a structure $\mathcal{A}$ over τ, we set $|\mathcal{A}| = |\tau| + |U(\mathcal{A})| + \sum_{R \in \tau} |R^{\mathcal{A}}| \cdot a_R$, where a_R is the arity of R.

For example, a *graph* G is a structure over the signature (E) where E has arity 2. The *Gaifman graph* of a structure $\mathcal{A}$ has as vertices the universe $U(\mathcal{A})$ of $\mathcal{A}$, and for each pair of vertices u, v, there is an edge $\{u, v\}$ in E if and only if at least one of the relations of $\mathcal{A}$ contains a tuple containing both u and v. Note that the edge set E of a Gaifman graph is symmetric and irreflexive.

Let $\mathcal{A}$ and $\mathcal{B}$ be structures over the same signature τ. Then $\mathcal{A}$ is a *substructure* of $\mathcal{B}$ if $U(\mathcal{A}) \subseteq U(\mathcal{B})$ and, for each relation symbol R in τ, it holds that $R^{\mathcal{A}} \subseteq R^{\mathcal{B}} \cap U(\mathcal{A})^a$, where a is the arity of R. A substructure is *induced* if, for each relation symbol R in τ, we have $R^{\mathcal{A}} = R^{\mathcal{B}} \cap U(\mathcal{A})^a$. A substructure $\mathcal{A}$ of $\mathcal{B}$ with $\mathcal{A} \neq \mathcal{B}$ is a *proper substructure* of $\mathcal{B}$. We also define the *union* $\mathcal{A} \cup \mathcal{B}$ of two structures $\mathcal{A}$ and $\mathcal{B}$ as the structure over τ with universe $U(\mathcal{A}) \cup U(\mathcal{B})$ and $R^{\mathcal{A} \cup \mathcal{B}} = R^{\mathcal{A}} \cup R^{\mathcal{B}}$. Note that the union is well-defined even if the universes are not disjoint.

Homomorphisms as Answers to CQs. Let $\mathcal{A}$ and $\mathcal{B}$ be structures over signatures $\tau_{\mathcal{A}} \subseteq \tau_{\mathcal{B}}$. A *homomorphism* from $\mathcal{A}$ to $\mathcal{B}$ is a mapping $h : U(\mathcal{A}) \to U(\mathcal{B})$ such that for each relation symbol $R \in \tau_{\mathcal{A}}$ with arity a and each tuple $\vec{t} = (t_1, \ldots, t_a) \in R^{\mathcal{A}}$ we have that $h(\vec{t}) = (h(t_1), \ldots, h(t_a)) \in R^{\mathcal{B}}$. We use $\mathrm{Hom}(\mathcal{A} \to \mathcal{B})$ to denote the set of homomorphisms from $\mathcal{A}$ to $\mathcal{B}$, and we use the lower case version $\mathrm{hom}(\mathcal{A} \to \mathcal{B})$ to denote the *number* of homomorphisms from $\mathcal{A}$ to $\mathcal{B}$.

Let φ be a conjunctive query with free variables $X = \{x_1, \ldots, x_k\}$ and quantified variables $Y = \{y_1, \ldots, y_d\}$. We can associate φ with a structure $\mathcal{A}_\varphi$ defined as follows: The universe of $\mathcal{A}_\varphi$ are the variables $X \cup Y$ and for each atom $R(\vec{t})$ of φ we add the tuple $\vec{t}$ to $R^{\mathcal{A}}$. It is well-known that, for each database $\mathcal{D}$, the set of answers of φ in $\mathcal{D}$ is precisely the set of assignments $a : X \to U(\mathcal{D})$ such that there is a homomorphism $h \in \mathrm{Hom}(\mathcal{A}_\varphi \to \mathcal{D})$ with $h|_X = a$. Since working with (partial) homomorphisms will be very convenient in this work, we will use the notation from [30] and (re)define a conjunctive query as a pair consisting of a relational structure $\mathcal{A}$ together with a set $X \subseteq U(\mathcal{A})$. The size of $(\mathcal{A}, X)$ is denoted by $|(\mathcal{A}, X)|$ and defined to be $|\mathcal{A}| + |X|$. Furthermore, we define

$$\mathrm{Ans}((\mathcal{A}, X) \to \mathcal{D}) \coloneqq \{a : X \to U(\mathcal{D}) \mid \exists h \in \mathrm{Hom}(\mathcal{A} \to \mathcal{D}) : h|_X = a\}$$

as the set of answers of $(\mathcal{A}, X)$ in $\mathcal{D}$. We then use $\mathrm{ans}((\mathcal{A}, X) \to \mathcal{D})$ to denote the number of answers, i.e., $\mathrm{ans}((\mathcal{A}, X) \to \mathcal{D}) \coloneqq |\mathrm{Ans}((\mathcal{A}, X) \to \mathcal{D})|$.

We can now formally define the (parameterised) problem of counting answers to conjunctive queries. As is usual, we restrict the problem by a class C of allowed queries.

Name: #CQ(C)
Input: A conjunctive query $(\mathcal{A}, X) \in C$ together with a database $\mathcal{D}$.
Parameter: $|(\mathcal{A}, X)|$.
Output: The number of answers $\mathrm{ans}((\mathcal{A}, X) \to \mathcal{D})$.

Acyclicity and (Hyper-)Treewidth. A structure of arity 2 is *acyclic* if its Gaifman graph is acyclic. For higher arities, a structure is acyclic if it has a join-tree or, equivalently, if it has (generalised) hyper-treewidth at most 1 (see [38]). We will not need the concept of hyper-treewidth, but we refer the reader to [37] for a comprehensive treatment.

The treewidth of graphs and structures is defined as follows

Definition 12 (Tree decompositions, treewidth). Let G be a graph. A *tree decomposition* of G is a pair (T, B), where T is a (rooted) tree, and B assigns each vertex $t \in V(T)$ a *bag* B_t such that the following constraints are satisfied:

(C1) $V(G) = \bigcup_{t \in V(T)} B_t$,

(C2) For each edge $e \in E(G)$ there exists $t \in V(T)$ such that $e \subseteq B_t$, and

(C3) For each $v \in V(G)$, the subgraph of T containing all vertices t with $v \in B_t$ is connected.

The *width* of a tree decomposition is $\max_{t \in V(T)} |B_t| - 1$, and the *treewidth* of G is the minimum width of any tree decomposition of G. Finally, the treewidth of a structure is the treewidth of its Gaifman graph.

#Equivalence, #Minimality, and #Cores. In the realm of decision problems, it is well known that evaluating a conjunctive query is equivalent to evaluating the (homomorphic) core of the query, i.e., evaluating the minimal homomorphic-equivalent query (see, for instance, the seminal work of Chandra and Merlin [20]). A similar, albeit slightly different notion of equivalence and minimality is required for counting answers to conjunctive queries. In what follows, we will provide the necessary definitions and properties of equivalence, minimality and cores for counting answers to conjunctive queries, and we refer the reader to [22] and to the full version of [30] for a more comprehensive discussion. To avoid confusion between the notions in the realms of decision and counting, we will from now on use the # symbol for the counting versions (see Definition 14).

An *isomorphism* from a structure $\mathcal{A}$ to a structure $\mathcal{A}'$ is a bijection $b : U(\mathcal{A}) \to U(\mathcal{A}')$ such that for each relation symbol $R \in \tau_A$ and each tuple $\vec{t}$ of elements of $U(\mathcal{A})$, we have $\vec{t} \in R^{\mathcal{A}}$ if and only if $b(\vec{t}) \in R^{\mathcal{A}'}$.

Definition 13. Two conjunctive queries $(\mathcal{A}, X)$ and $(\mathcal{A}', X')$ are *isomorphic*, denoted by $(\mathcal{A}, X) \cong (\mathcal{A}', X')$, if there is an isomorphism b from $\mathcal{A}$ to $\mathcal{A}'$ with $b(X) = X'$.

Definition 14 (#Equivalence and #minimality (see [22, 30])). Two conjunctive queries $(\mathcal{A}, X)$ and $(\mathcal{A}', X')$ are *#equivalent*, denoted by $(\mathcal{A}, X) \sim (\mathcal{A}', X')$, if for every database $\mathcal{D}$ we have $\mathrm{ans}((\mathcal{A}, X) \to \mathcal{D}) = \mathrm{ans}((\mathcal{A}', X') \to \mathcal{D})$. A conjunctive query $(\mathcal{A}, X)$ is *#minimal* if there is no proper substructure $\mathcal{A}'$ of $\mathcal{A}$ such that $(\mathcal{A}, X) \sim (\mathcal{A}', X)$.

Observation 15. *The following are equivalent:*

(1) A conjunctive query $(\mathcal{A}, X)$ is #minimal.

(2) $(\mathcal{A}, X)$ has no #equivalent substructure that is induced by a set U with $X \subseteq U \subset U(\mathcal{A})$.

(3) Every homomorphism from $\mathcal{A}$ to itself that is the identity on X is surjective.

It turns out that #equivalence is the same as isomorphism if all variables are free, and it is the same as homomorphic equivalence if all variables are existentially quantified (see e.g. the discussion in Section 5 in the full version of [30]). Moreover, each quantifier-free conjunctive query is #minimal.

Lemma 16 (Corollary 54 in full version of [30]). *If $(\mathcal{A}, X)$ and $(\mathcal{A}', X')$ are #minimal and $(\mathcal{A}, X) \sim (\mathcal{A}', X')$ then $(\mathcal{A}, X) \cong (\mathcal{A}', X')$.*

As isomorphism trivially implies #equivalence, Lemma 16 shows that, for #minimal queries, #equivalence and isomorphism coincide.

Definition 17 (#core). A *#core* of a conjunctive query $(\mathcal{A}, X)$ is a #minimal conjunctive query $(\mathcal{A}', X')$ with $(\mathcal{A}, X) \sim (\mathcal{A}', X')$.

By Lemma 16, the #core is unique up to isomorphisms; in fact, this allows us to speak of "the" #core of a conjunctive query.

Classification of #CQ(C) via Treewidth and Contracts. It is well known that the complexity of counting answers to a conjunctive query is governed by its treewidth, and by the treewidth of its contract [21, 30], which we define as follows.

Definition 18 (Contract). Let $(\mathcal{A}, X)$ be a conjunctive query, let $Y = U(\mathcal{A}) \setminus X$, and let G be the Gaifmann graph of $\mathcal{A}$. The *contract* of $(\mathcal{A}, X)$, denoted by $\mathrm{contract}(\mathcal{A}, X)$ is obtained from $G[X]$ by adding an edge between each pair of vertices u and v for which there is a connected component S in $G[Y]$ that is adjacent to both u and v, that is, there are vertices $x, y \in S$ such that $\{x, u\} \in E(G)$ and $\{y, v\} \in E(G)$. Given a class of conjunctive queries C, we write $\mathrm{contract}(C)$ for the class of all contracts of queries in C.

We note that there are multiple equivalent ways to define the contract of a query. For our purposes, the definition in [30] is most suitable. Also, the treewidth of the contract of a conjunctive query is an upper bound of what is called the query's "star size" in [31] and its "dominating star size" in [30].

Chen and Mengel established the following classification for counting answers to conjunctive queries of bounded arity.

Theorem 19 ([21]). *Let C be a recursively enumerable class of conjunctive queries of bounded arity, and let C' be the class of #cores of queries in C. If the treewidth of C' and of $\mathrm{contract}(C')$ is bounded, then #CQ(C) is solvable in polynomial time. Otherwise, #CQ(C) is W[1]-hard.*

We point out that the W[1]-hard cases can further be partitioned into W[1]-complete, #W[1]-complete and even harder cases [21, 30].[5] However, for the purpose of this work, we are only interested in tractable and intractable cases (recall that W[1]-hard problems are not fixed-parameter tractable under standard assumptions from fine-grained and parameterised complexity theory, such as ETH).

Self-join-free Conjunctive Queries and Isolated Variables. A conjunctive query $(\mathcal{A}, X)$ is *self-join-free* if each relation of $\mathcal{A}$ contains at most one tuple. We say that a variable of a conjunctive query is *isolated* if it is not part of any relation.

Note that that adding/removing isolated variables to/from a conjunctive query does not change its treewidth or the treewidth of its #core. Further, it does not change the complexity of counting answers: Just multiply/divide by n^v, where n is the number of elements of the database and v it the number of added/removed isolated free variables. For this reason, we will allow ourselves in this work to freely add and remove isolated variables from the queries that we encounter. For the existence of homomorphisms we also observe the following.

Observation 20. *Let $(\mathcal{A}, X)$ be a conjunctive query, let X' be a superset of X and let $\mathcal{A}'$ be the structure obtained from $\mathcal{A}$ by adding an isolated variable for each $x \in X' \setminus X$. Then for all $a : X' \to U(\mathcal{D})$ we have that $a|_X \in \mathrm{Ans}((\mathcal{A}, X) \to \mathcal{D})$ iff $a \in \mathrm{Ans}((\mathcal{A}', X') \to \mathcal{D})$.*

2.3 Unions of CQs and the Homomorphism Basis

A *union of conjunctive queries* (UCQ) Ψ is a tuple of structures $(\mathcal{A}_1, \dots, \mathcal{A}_{\ell(\Psi)})$ over the same signature together with a set of designated elements X (the free variables) that are in the universe of each of the structures. For each $J \subseteq [\ell(\Psi)]$, we define $\Psi|_J = ((A_j)_{j \in J}, X)$. If we restrict to a single term of the union then we usually just write Ψ_i instead of $\Psi|_{\{i\}}$. Note that $\Psi_i = (\mathcal{A}_i, X)$ is simply a conjunctive query (rather than a union of CQs).

We will assume (without loss of generality) that, for any distinct i and i' in $[\ell(\Psi)]$, $U(\mathcal{A}_i) \cap U(\mathcal{A}_{i'}) = X$, i.e., that each CQ in the union has its own set of existentially quantified variables.

If each such conjunctive query is acyclic we say that Ψ is a union of acyclic conjunctive queries. Moreover, the arity of Ψ is the maximum arity of any of the $\mathcal{A}_i$. The size of Ψ is $|\Psi| = \sum_{i=1}^{\ell(\Psi)} |\Psi_i|$. The elements of X are the *free variables* of Ψ and $\ell(\Psi)$ is the number of CQs in the union.

---

[5]Those cases are: #W[2]-hard and #A[2]-complete.

The set of *answers* of Ψ in a database $\mathcal{D}$, denoted by $\mathrm{Ans}(\Psi \to \mathcal{D})$ is defined as follows:

$$\mathrm{Ans}(\Psi \to \mathcal{D}) = \{a : X \to U(\mathcal{D}) \mid \exists i \in [\ell] : a \in \mathrm{Ans}(\Psi_i \to \mathcal{D})\}\,.$$

Again, we use the lower case version $\mathrm{ans}(\Psi \to \mathcal{D})$ to denote the *number* of answers of Ψ in $\mathcal{D}$.

In the definition of UCQs we assume that every CQ in the union has the same set of free variables, namely X. This assumption is without loss of generality. To see this, suppose that we have a union of CQs $(\mathcal{A}_1, X_1), \ldots, (\mathcal{A}_\ell, X_\ell)$ with individual sets of free variables. Let $X = \bigcup_{i=1}^{\ell} X_i$ and, for each $i \in [\ell]$, let $\mathcal{A}_i'$ be the structure obtained from $\mathcal{A}_i$ by adding an isolated variable for each $x \in X \setminus X_i$. Then consider the UCQ $\Psi := ((\mathcal{A}_1', \ldots, \mathcal{A}_\ell'), X)$. If for some assignment $a : X \to U(\mathcal{D})$ it holds that there is an $i \in [\ell]$ such that $a|_{X_i} \in \mathrm{Ans}((\mathcal{A}_i, X_i) \to \mathcal{D})$, then, according to Observation 20, this is equivalent to $a \in \mathrm{Ans}((\mathcal{A}_i', X) \to \mathcal{D})$, which means that a is an answer of Ψ. So, without loss of generality we can work with Ψ, which uses the same set of free variables for each CQ in the union.

Now we define the parameterised problem of counting answers to UCQs. As usual, the problem is restricted by a class C of allowed queries with respect to which we classify the complexity.

Name: #UCQ(C)
Input: A UCQ $\Psi \in C$ together with a database $\mathcal{D}$.
Parameter: $|\Psi|$.
Output: The number of answers $\mathrm{ans}(\Psi \to \mathcal{D})$.

The next definition will be crucial for the analysis of the complexity of #UCQ(C).

Definition 21 (combined query $\wedge (\Psi)$). Let $\Psi = ((\mathcal{A}_1, \ldots, \mathcal{A}_\ell), X)$ be a UCQ. Then we define the *combined query* $\wedge (\Psi) = (\bigcup_{j \in [\ell]} A_j, X)$.

What follows is an easy, but crucial observation about $\wedge (\Psi|_J)$.

Observation 22. *Let $((\mathcal{A}_1, \ldots, \mathcal{A}_\ell), X)$ be a UCQ, and let $\emptyset \neq J \subseteq [\ell]$. For each database $\mathcal{D}$ and assignment $a : X \to U(\mathcal{D})$ we have*

$$a \in \mathrm{Ans}(\wedge (\Psi|_J) \to \mathcal{D}) \Leftrightarrow \forall j \in J : a \in \mathrm{Ans}(\Psi_j \to \mathcal{D}).$$

Definition 23 (Coefficient function c_Ψ). Let $\Psi = ((\mathcal{A}_1, \ldots, \mathcal{A}_\ell), X)$ be a UCQ. For each conjunctive query $(\mathcal{A}, X)$, we set $I(\mathcal{A}, X) = \{J \subseteq [\ell] \mid (\mathcal{A}, X) \sim \wedge (\Psi|_J)\}$, and we define the *coefficient function* of Ψ as follows:

$$c_\Psi(\mathcal{A}, X) = \sum_{J \in I(\mathcal{A},X)} (-1)^{|J|+1}\,.$$

Using inclusion-exclusion, we can transform the problem of counting answers to Ψ into the problem of evaluating a linear combination of CQ answer counts. We include a proof only for reasons of self-containment and note that the complexity-theoretic applications of this transformation, especially regarding lower bounds, have first been discovered by Chen and Mengel [22].

Lemma 24 ([22]). *Let $\Psi = ((\mathcal{A}_1, \ldots, \mathcal{A}_\ell), X)$ be a UCQ. For every database $\mathcal{D}$,*

$$\mathrm{ans}(\Psi \to \mathcal{D}) = \sum_{(\mathcal{A},X)} c_\Psi(\mathcal{A}, X) \cdot \mathrm{ans}((\mathcal{A}, X) \to \mathcal{D}),$$

where the sum is over all equivalence classes of $\sim$.

Proof. By inclusion-exclusion and Observation 22,

$$\mathrm{ans}(\Psi \to \mathcal{D}) = \sum_{\emptyset \neq J \subseteq [\ell]} (-1)^{|J|+1} \cdot \left| \{ a : X \to U(\mathcal{D}) \mid \forall j \in J : a \in \mathrm{Ans}((A_j, X) \to \mathcal{D}) \} \right|$$

$$= \sum_{\emptyset \neq J \subseteq [\ell]} (-1)^{|J|+1} \cdot \mathrm{ans}(\wedge (\Psi|_J) \to \mathcal{D}).$$

The claim then follows by collecting #equivalent terms. □

We conclude this subsection with the following two operations on classes of UCQs.

Definition 25 ($\Gamma(C)$ and $\wedge(C)$). Let C be a class of UCQs. $\Gamma(C)$ is the class of all $(\mathcal{A}, X)$ such that $(\mathcal{A}, X)$ is #minimal and there is $\Psi \in C$ with $c_\Psi(\mathcal{A}, X) \neq 0$. Let $\wedge(C) = \{ \wedge(\Psi) \mid \Psi \in C \}$.

2.4 Complexity Monotonicity

The principle of Complexity Monotonicity states that the computation of a linear combination of homomorphism counts is precisely as hard as computing its hardest term. It was independently discovered by Curticapean, Dell and Marx [26] and by Chen and Mengel [22]. Moreover, Chen and Mengel established the principle in the more general context of linear combinations of conjunctive queries. Formally, their result is stated below for the special case of linear combinations derived via counting answers to UCQs (see Lemma 24). We include a proof since we need to pay some special attention to running times, as we will also be interested in the question of linear time tractability.

Theorem 26 (Implicitly by [22]). *There is an algorithm $\mathbb{A}$ with the following properties:*

(1) *The input of $\mathbb{A}$ is a UCQ Ψ and a database $\mathcal{D}$.*
(2) *$\mathbb{A}$ has oracle access to the function $\mathcal{D}' \mapsto \mathrm{ans}(\Psi \to \mathcal{D}')$.*
(3) *The output of $\mathbb{A}$ is a list with entries $((\mathcal{A}, X), \mathrm{ans}((\mathcal{A}, X) \to \mathcal{D}))$ for each $(\mathcal{A}, X)$ in the support of c_Ψ.*
(4) *$\mathbb{A}$ runs in time $f(|\Psi|) \cdot O(|\mathcal{D}|)$ for some computable function f.*

Proof. A crucial operation in the construction is the Tensor product of relational structures. Let $\mathcal{A}$ and $\mathcal{B}$ be structures over the signatures τ_A and τ_B. The structure $\mathcal{A} \otimes \mathcal{B}$ is defined as follows: The signature is $\tau_A \cap \tau_B$, and the universe is $U(\mathcal{A}) \times U(\mathcal{B})$. Moreover, for every relation symbol $R \in \tau_A \cap \tau_B$ with arity r, a tuple $((u_1, v_1), \ldots, (u_r, v_r))$ is contained in $R^{\mathcal{A} \otimes \mathcal{B}}$ if and only if $(u_1, \ldots, u_r) \in R^{\mathcal{A}}$ and $(v_1, \ldots, v_r) \in R^{\mathcal{B}}$.

Observe that $\mathcal{A} \otimes \mathcal{B}$ is of size bounded by and can be computed in time $O(|\mathcal{A}||\mathcal{B}|)$.[6] The algorithm $\mathbb{A}$ proceeds as follows. Let $\Psi = ((\mathcal{A}_1, \ldots, \mathcal{A}_\ell), X)$ be the input. For a selected set of structures $\mathcal{B}_1 \ldots, \mathcal{B}_k$, specified momentarily, the algorithm queries the oracle on the Tensor products $\mathcal{D} \otimes \mathcal{B}_i$. Using Lemma 24, this yields the following equations:

$$\mathrm{ans}(\Psi \to \mathcal{D} \otimes \mathcal{B}_i) = \sum_{(\mathcal{A}, X)} c_\Psi(\mathcal{A}, X) \cdot \mathrm{ans}((\mathcal{A}, X) \to \mathcal{D} \otimes \mathcal{B}_i).$$

Next, we use the fact (see e.g. [22]) that the Tensor product is multiplicative with respect to counting answers to conjunctive queries:

$$\mathrm{ans}((\mathcal{A}, X) \to \mathcal{D} \otimes \mathcal{B}_i) = \mathrm{ans}((\mathcal{A}, X) \to \mathcal{D}) \cdot \mathrm{ans}((\mathcal{A}, X) \to \mathcal{B}_i).$$

---

[6]Compute the Cartesian product of $U(\mathcal{A})$ and $U(\mathcal{B})$ and then, for every relation $R \in \tau_A \cap \tau_B$ iterate over all pairs of tuples in $R^{\mathcal{A}}$ and $R^{\mathcal{B}}$ and add their point-wise product to $R^{\mathcal{A} \otimes \mathcal{B}}$.

In combination, the previous equations yield a system of linear equations:

$$\text{ans}(\Psi \to \mathcal{D} \otimes \mathcal{B}_i) = \sum_{(\mathcal{A},X)} c_\Psi(\mathcal{A},X) \cdot \text{ans}((\mathcal{A},X) \to \mathcal{D}) \cdot \text{ans}((\mathcal{A},X) \to \mathcal{B}_i),$$

the unknowns of which are $c_\Psi(\mathcal{A},X) \cdot \text{ans}((\mathcal{A},X) \to \mathcal{D})$. It was shown in [22] and [30] that it is always possible to find $\mathcal{B}_i$ for which the system is non-singular. Moreover, the time it takes to find the $\mathcal{B}_i$ only depends on Ψ. Finally, solving the system yields the terms $c_\Psi(\mathcal{A},X) \cdot \text{ans}((\mathcal{A},X) \to \mathcal{D})$ from which we can recover $\text{ans}((\mathcal{A},X) \to \mathcal{D})$ by dividing by $c_\Psi(\mathcal{A},X)$. The overall running time is bounded by $f(|\Psi|) \cdot O(|\mathcal{D}|)$ for some computable function f: to see this, note that the numbers in the system can be generously bounded by $f'(|\Psi|) \cdot |\mathcal{D}|^{|\Psi|}$, hence their size is bounded by $\log(f'(|\Psi|)) + |\Psi| \log(|\mathcal{D}|)$ and thus the arithmetic operations needed for solving the system do not create a super-linear overhead in $|\mathcal{D}|$. This concludes the proof.   □

Corollary 27. *Let Ψ be a UCQ. For each $d \geq 1$, computing the function $\mathcal{D} \mapsto \text{ans}(\Psi \to \mathcal{D})$ can be done in time $O(|\mathcal{D}|^d)$ if and only if for each #minimal $(\mathcal{A},X)$ with $c_\Psi(\mathcal{A},X) \neq 0$ the function $\mathcal{D} \mapsto \text{ans}((\mathcal{A},X) \to \mathcal{D})$ can be computed in time $O(|\mathcal{D}|^d)$.*

Corollary 28 (Implicitly also in [22]). *Let C be a recursively enumerable class of UCQs. The problems #UCQ(C) and #CQ($\Gamma(C)$) are interreducible with respect to parameterised Turing-reductions.*

As a consequence, in combination with Theorem 19, Chen and Mengel [22] established the following dichotomy, which we can state in a convenient way using our notation:

Theorem 1 ([22]). *Let C be a recursively enumerable class of UCQs of bounded arity. If the treewidth of $\Gamma(C)$ and of $\text{contract}(\Gamma(C))$ is bounded, then #UCQ(C) is fixed-parameter tractable. Otherwise, #UCQ(C) is W[1]-hard.*

In other words, as a summary of the above results, we know that counting answers to a UCQ is hard if and only if a conjunctive query survives in the CQ expansion, whose #core has either high treewidth or its contract has high treewidth. Unfortunately, given a concrete UCQ, or a concrete class of UCQs, it is not clear how to determine whether such high treewidth terms survive.

The central question that we ask in this work is whether this implicit tractability criterion for counting answers to UCQs can be rephrased as a more natural one. Equivalently, this means that we aim to find out whether there is a natural criterion for the existence of high treewidth terms in the CQ expansion. Such natural criteria have been found for subgraph and induced subgraph counting[7] [26, 35], and for conjunctive queries with inequalities and negations [30, 59].

3 CLASSIFICATIONS FOR DELETION-CLOSED UCQS

Let C be a class of UCQs. Recall that $\wedge(C)$ is the class of all conjunctive queries that are obtained just by substituting all $\vee$ by $\wedge$ in UCQs in C, whereas $\Gamma(C)$ in Theorem 1 is the much less natural class of #minimal queries that survive with a non-zero coefficient in the CQ expansion of a UCQ in C. The work of Chen and Mengel [22] implicitly also shows an upper bound for counting answers to UCQs from the class C in terms of the simpler objects $\wedge(C)$ and $\text{contract}(\wedge(C))$, rather than in terms of the more complicated objects $\Gamma(C)$ and $\text{contract}(\Gamma(C))$. We include a proof for completeness.

Lemma 29. *Let C be a recursively enumerable class of UCQs. Suppose that both $\wedge(C)$ and $\text{contract}(\wedge(C))$ have bounded treewidth. Then #UCQ(C) is fixed-parameter tractable.*

---

[7](Induced) Subgraph counting is a special case of counting answers to quantifier-free conjunctive queries with inequalities (and negations).

Proof. Let $\Psi \in C$. Recall from the proof of Lemma 24 that, for every databse $\mathcal{D}$,

$$\text{ans}(\Psi \to \mathcal{D}) = \sum_{\emptyset \neq J \subseteq [\ell]} (-1)^{|J|+1} \cdot \text{hom}(\wedge (\Psi|_J) \to \mathcal{D}) \,.$$

Hence $\#\text{UCQ}(C) \leq^{\text{FPT}} \#\text{CQ}(\hat{C})$ where $\hat{C}$ is $\{\wedge (\Psi|_J) \mid \Psi \in C \ \wedge \ \emptyset \neq J \subseteq [\ell(\Psi)]\}$. Finally, since $\wedge (\Psi|_J)$ is a subquery of $\wedge (\Psi)$ for each J, the treewidths of $\wedge (\Psi|_J)$ and contract$(\wedge (\Psi|_J))$ are bounded from above by the treewidths of $\wedge (\Psi)$ and contract$(\wedge (\Psi))$, respectively. Consequently, the treewidths of $\hat{C}$ and contract$(\hat{C})$ are bounded, and thus $\#\text{CQ}(\hat{C})$ is polynomial-time solvable by the classification of Chen and Mengel [21, Theorem 22] Since $\#\text{UCQ}(C) \leq^{\text{FPT}} \#\text{CQ}(\hat{C})$, the lemma follows.   □

Our goal is to relate the complexity of $\#\text{UCQ}(C)$ to the structure of $\wedge (C)$ with the hope of obtaining a more natural tractability criterion than the one given by Theorem 1. While we will see later that this seems not always possible (Appendix A), we identify conditions under which a natural criterion based on $\wedge (C)$ is possible, both in the quantifier-free case (Section 3.1), and in the general case that allows quantified variables (Section 3.2).

A class of UCQs C is *closed under deletions* if, for every $\Psi = ((\mathcal{A}_1, \ldots, \mathcal{A}_\ell), X) \in C$ and for every $\emptyset \neq J \subseteq [\ell]$, the UCQ $\Psi|_J$ is also contained in C. For example, any class of UCQs defined by prescribing the allowed conjunctive queries is closed under deletions. This includes, e.g., unions of acyclic conjunctive queries.

3.1 The Quantifier-free Case

As a warm-up, we start with the much simpler case of quantifier-free queries. Here, we only allow (unions of) conjunctive queries $(\mathcal{A}, X)$ satisfying $U(\mathcal{A}) = X$.

Lemma 30. *Let C be a recursively enumerable class of quantifier-free UCQs of bounded arity. Suppose that C is closed under deletions. If $\wedge (C)$ has unbounded treewidth then $\#UCQ(C)$ is $W[1]$-hard.*

Proof. We show that $\wedge (C) \subseteq \Gamma(C)$, which then proves the claim by Theorem 1. Recall from Definition 25 that $\Gamma(C) = \{(\mathcal{A}, X) \mid (\mathcal{A}, X) \text{ is } \#\text{minimal and there is } \Psi \in C \text{ with } c_\Psi(\mathcal{A}, X) \neq 0\}$. Let $\Psi = ((\mathcal{A}_1, \ldots, \mathcal{A}_\ell), X) \in C$. Note that, according to Observation 15, $\wedge (\Psi)$ is its own #core since it does not have existentially quantified variables. For the same reason, for each nonempty subset J of $[\ell]$, the query $\wedge (\Psi|_J)$ is its own #core. Now let $J \subseteq [\ell]$ be inclusion-minimal with the property that $\wedge (\Psi|_J)$ is isomorphic to $\wedge (\Psi)$. Since C is closed under deletions, the UCQ $\Psi|_J$ is contained in C. By the inclusion-minimality of J, Definition 23 ensures that $c_{\Psi|_J}(\wedge (\Psi)) = (-1)^{|J|+1} \neq 0$. As a consequence, $\wedge (\Psi) \in \Gamma(C)$, concluding the proof.   □

From Lemmas 30 and 29 together with the fact that the contract of a quantifier-free query is the query itself, we obtain the classification for quantifier-free UCQs, which we restate for convenience.

Theorem 2. *Let C be recursively enumerable class of quantifier-free UCQs of bounded arity. If $\wedge (C)$ has bounded treewidth then $\#UCQ(C)$ is fixed-parameter tractable. If $\wedge (C)$ has unbounded treewidth and C is closed under deletions then $\#UCQ(C)$ is $W[1]$-hard.*

3.2 The General Case

Now we consider UCQs with existentially quantified variables. Here, a corresponding hardness result (Lemma 33) can be achieved under some additional assumptions. Note that the number of existentially quantified variables in a UCQ $\Psi = ((\mathcal{A}_1, \ldots, \mathcal{A}_\ell), X)$ is equal to $\sum_{i=1}^\ell |U(\mathcal{A}_i) \setminus X|$. We first need the following two auxiliary results:

Lemma 31. *Let $(\mathcal{A}, X)$ and $(\mathcal{A}', X')$ be #equivalent conjunctive queries. Further, let G and G' be the Gaifman graphs of $\mathcal{A}$ and $\mathcal{A}'$, respectively. Then $G[X]$ and $G'[X']$ are isomorphic.*

PROOF. By [22] (see Lemma 48 in the full version of [30] for an explicit statement), there are surjective functions $s : X \to X'$ and $s' : X' \to X$ and homomorphisms $h \in \text{Hom}(\mathcal{A} \to \mathcal{A}')$ and $h' \in \text{Hom}(\mathcal{A}' \to \mathcal{A})$ such that $h|_X = s$ and $h'|_{X'} = s'$.

Clearly, s and s' are bijective. Let $e = \{u, v\}$ be an edge of $G[X]$. Then there exists a tuple $\vec{t}$ of elements of $U(\mathcal{A})$ such that

(i) $\vec{t} \in R^{\mathcal{A}}$ for some relation (symbol) R of the signature of $\mathcal{A}$, and
(ii) u and v are elements of $\vec{t}$.

Since h is a homomorphism, $h(\vec{t}) \in R^{\mathcal{B}}$. Thus $\{h(u), h(v)\} = \{s(u), s(v)\}$ is an edge of $G'[X']$. The backward direction is analogous. □

Lemma 32. *Let $(\mathcal{A}, X)$ be a self-join-free conjunctive query. Let $\mathcal{A}'$ be the structure obtained from $\mathcal{A}$ by deleting all isolated variables in $U(\mathcal{A}) \setminus X$. Then $(\mathcal{A}', X)$ is the #core of $(\mathcal{A}, X)$.*

PROOF. Clearly, $(\mathcal{A}, X)$ and $(\mathcal{A}', X)$ are #equivalent. Thus it remains to show that $(\mathcal{A}', X)$ is #minimal. Assume for contradiction by Observation 15 that $(\mathcal{A}', X)$ has an #equivalent substructure $\hat{\mathcal{A}}$ that is induced by a set U with $X \subseteq U \subset U(\mathcal{A}')$.

Since $\mathcal{A}'$ is self-join-free and it does not have isolated variables, there is a relation (symbol) R such that $R^{\mathcal{A}'}$ contains precisely one tuple, and $R^{\hat{\mathcal{A}}}$ is empty. Thus, there is no homomorphism from $\mathcal{A}'$ to $\hat{\mathcal{A}}$ and, consequently, $(\hat{\mathcal{A}}, X)$ and $(\mathcal{A}', X)$ cannot be #equivalent, yielding a contradiction and concluding the proof. □

Lemma 33. *Let C be a recursively enumerable class of unions of self-join-free conjunctive queries with bounded arity. Suppose that C is closed under deletions and that there is a finite upper bound on the number of existentially quantified variables in queries in C. If either of $\wedge(C)$ or $\text{contract}(\wedge(C))$ have unbounded treewidth then #UCQ(C) is W[1]-hard.*

PROOF. Let d be the maximum number of existentially quantified variables in a query in C. Assume first that $\wedge(C)$ has unbounded treewidth. We show that $\Gamma(C)$ has unbounded treewidth, which proves the claim by Theorem 1. To this end, let B be any positive integer. The goal is to find a conjunctive query in $\Gamma(C)$ with treewidth at least B. Since $\wedge(C)$ has unbounded treewidth, there is a UCQ $\Psi = ((\mathcal{A}_1, \ldots, \mathcal{A}_\ell), X)$ in C such that $\wedge(\Psi)$ has treewidth larger than $d + B$. Note that, although all Ψ_i are self-join-free, $\wedge(\Psi)$ is not necessarily self-join-free. Let J be inclusion-minimal among the subsets of $[\ell]$ with the property that the #core of $\wedge(\Psi|_J)$ is isomorphic to the #core of $\wedge(\Psi)$. Since C is closed under deletions, the UCQ $\Psi|_J$ is contained in C. Let $(\mathcal{A}', X')$ be the #core of $\wedge(\Psi)$.

By inclusion-minimality of J, $c_{\Psi|_J}((\mathcal{A}', X')) = (-1)^{|J|+1} \neq 0$. As a consequence, $(\mathcal{A}', X') \in \Gamma(C)$. It remains to show that the treewidth of $(\mathcal{A}', X')$ is at least B. For this, let G be the Gaifmann graph of $\wedge(\Psi)$ and let G' be the Gaifmann graph of the #core of $\wedge(\Psi)$ (the Gaifman graph of $\mathcal{A}'$). First, deletion of a vertex can decrease the treewidth by at most 1. Thus, $G[X]$ has treewidth at least $d + B - d = B$. By Lemma 31, $G[X]$ and $G'[X']$ are isomorphic.

Therefore the treewidth of $G'[X']$, i.e., the treewidth of $(\mathcal{A}', X')$, is at least B. So we have shown that if the treewidth of $\wedge(C)$ is unbounded then so is the treewidth of $\Gamma(C)$.

In the second case, we assume that the contracts of queries in $\wedge(C)$ (see Definition 18) have unbounded treewidth. We introduce the following terminology: Let $(\mathcal{A}, X)$ be a conjunctive query and let $y \in U(\mathcal{A})$. The *degree of freedom* of y is the number of vertices in X that are adjacent to y in the Gaifman graph of $\mathcal{A}$. Let $\hat{C}$ be the class of all conjunctive queries $(\mathcal{A}, X)$ such that there exists $\Psi = ((\mathcal{A}_1, \ldots, \mathcal{A}_{\ell(\Psi)}), X)$ in C with $(\mathcal{A}, X) = (\mathcal{A}_i, X)$ for some $i \in [\ell(\Psi)]$. Since C is closed

under deletions, $\hat{C} \subseteq C$. By the assumptions of the lemma, $\hat{C}$ consists only of self-join-free queries. Thus, by Lemma 32, each query in $\hat{C}$ is its own #core (up to deleting isolated variables). We will now consider the following cases:

(i) Suppose that $\hat{C}$ has unbounded degree of freedom. With Definition 18 it is straightforward to check that a quantified variable y with degree of freedom B induces a clique of size B in the contract of the corresponding query. Therefore, the contracts of the queries in $\hat{C}$ have unbounded treewidth. Consequently, #CQ($\hat{C}$) is W[1]-hard by the classification of Chen and Mengel [21, Theorem 22]. Since $\hat{C} \subseteq C$ the problem #CQ($\hat{C}$) is merely a restriction of #UCQ(C), the latter of which is thus W[1]-hard as well.

(ii) Suppose that the degree of freedom of queries in $\hat{C}$ is bounded by a constant d'. We show that $\Gamma(C)$ has unbounded treewidth, which proves the claim by Theorem 1. To this end, let B be any positive integer. The goal is to find a conjunctive query in $\Gamma(C)$ with treewidth at least B. Since contract($\wedge(C)$) has unbounded treewidth, there is a UCQ $\Psi = (\mathcal{A}_1, \ldots, \mathcal{A}_\ell), X)$ in C such that contract($\wedge(\Psi)$) has treewidth larger than $d + \binom{dd'}{2} + B$. We will show that $\wedge(\Psi)$ has treewidth larger than $d + B$, which, as we have argued previously, implies that $\Gamma(C)$ contains a query with treewidth at least B. To prove that $\wedge(\Psi)$ indeed has treewidth larger than $d + B$, let $\wedge(\Psi) = (\mathcal{A}, X)$ and let G be the Gaifman graph of $\mathcal{A}$. Let $Y = U(\mathcal{A}) \setminus X$ and recall from Definition 18 that contract($\mathcal{A}, X$) is obtained from $G[X]$ by adding an edge between any pair of vertices u and v that are adjacent to a common connected component in $G[Y]$. Let $N \subseteq X$ be the set of all vertices in X that are adjacent to a vertex in Y and observe that $|N| \leq dd'$ since the number of existentially quantified variables and the degree of freedom are bounded by d and d', respectively. Thus, contract($\mathcal{A}, X$) is obtained by adding at most $\binom{dd'}{2}$ edges to $G[X]$. The deletion of an edge can decrease the treewidth by at most 1, so

$$\mathrm{tw}(\wedge(\Psi)) = \mathrm{tw}(G) \geq \mathrm{tw}(G[X]) \geq \mathrm{tw}(\mathrm{contract}(\mathcal{A}, X)) - \binom{dd'}{2} > d + B,$$

which concludes Case (ii).

With all cases concluded, the proof is completed. □

From Lemmas 29 and 33 we directly obtain the main result of this section, which we restate for convenience.

Theorem 3. *Let C be a recursively enumerable class of UCQs of bounded arity. If $\wedge(C)$ and* contract($\wedge(C)$) *have bounded treewidth then #UCQ(C) is fixed-parameter tractable. Otherwise, if (I)–(III) are satisfied, then #UCQ(C) is* W[1]*-hard.*

 (I) *C is closed under deletions.*
 (II) *The number of existentially quantified variables of queries in C is bounded.*
 (III) *The UCQs in C are unions of self-join-free conjunctive queries.*

Remark 34. It turns out that all side conditions of Theorem 3 are necessary if we aim to classify #UCQ(C) solely via $\wedge(C)$. To this end, we provide counter examples for each missing condition in Appendix A.

4 THE META COMPLEXITY OF COUNTING ANSWERS TO UCQS

We consider the meta-complexity question of deciding whether it is possible to count the answers of a given UCQ in linear time. As pointed out in the introduction, this problem is immediately NP-hard even when restricted to conjunctive queries if we were to allow quantified variables. Therefore, we consider quantifier-free UCQs in this section. Recall the definition of META from Section 1.

Name: META
Input: A union Ψ of quantifier-free conjunctive queries.
Output: Is it possible to count answers to Ψ in time linear in the size of $\mathcal{D}$, i.e., can the function $\mathcal{D} \mapsto \text{ans}(\Psi \to \mathcal{D})$ be computed in time $O(|\mathcal{D}|)$.

Since we focus in this section solely on quantifier-free queries, it will be convenient to simplify our notation as follows. As all variables are free, we will identify a conjunctive query φ just by its associated structure, that is, we will write $\varphi = \mathcal{A}$, rather than $\varphi = (\mathcal{A}, U(\mathcal{A}))$. Similarly, we represent a union of quantifier-free conjunctive queries Ψ as a tuple of structures $\Psi = (\mathcal{A}_1, \ldots, \mathcal{A}_\ell)$.

For studying the complexity of META, it will be crucial to revisit the classification of linear-time counting of answers to quantifier-free conjunctive queries: The following theorem is well known and was discovered multiple times by different authors in different contexts.[8] This is Theorem 4 from the introduction, which we now restate in a version that expresses CQs as structures.

Theorem 35 (See Theorem 12 in [16], and [10, 53]). *Let $\mathcal{A}$ be a quantifier-free conjunctive query. If $\mathcal{A}$ is acyclic, then the function $\mathcal{D} \mapsto \text{hom}(\mathcal{A} \to \mathcal{D})$ is computable in linear time. Otherwise, if $\mathcal{A}$ has arity 2, then $\mathcal{D} \mapsto \text{hom}(\mathcal{A} \to \mathcal{D})$ cannot be solved in linear time unless the Triangle Conjecture fails.*

Theorem 35 yields an efficient way to check whether counting answers to a quantifier-free conjunctive query φ can be done in linear time: Just check whether φ is acyclic. In stark contrast, we show that no easy criterion for linear time tractability of counting answers to *unions* of conjunctive queries is possible, unless some conjectures of fine-grained complexity theory fail. In fact, our Theorem 5, which we restate here for convenience, precisely determines the complexity of META.

Theorem 5. *Assume that the Triangle Conjecture is true. When restricted to input queries of arity at most 2, META can be solved in time $2^{O(\ell)} \cdot |\Psi|^{\text{poly}(\log |\Psi|)}$, where ℓ is the number of conjunctive queries in the union. Moreover, META is NP-hard and the following conditional lower bounds apply:*

- *If ETH is true, then META cannot be solved in time $2^{o(\ell)}$.*
- *If the non-uniform ETH is true then META $\notin \bigcap_{\varepsilon > 0} \text{DTime}(2^{\varepsilon \cdot \ell})$.*

The lower bounds remain true even if Ψ is a union of self-join-free and acyclic conjunctive queries over a binary signature (that is, of arity 2).

The lower bounds in Theorem 5 imply that the exponential dependence on ℓ in our $2^{O(\ell)} \cdot |\Psi|^{\text{poly}(\log |\Psi|)}$ time algorithm for META cannot be significantly improved, unless standard assumptions fail.

The remainder of this section is devoted to the proof of Theorem 5. It is split into two parts: In the first and easier part (Section 4.1), we construct the algorithm for META. For this, all we need to do is to translate the problem into the homomorphism basis (see Section 2.3), i.e., we transform the problem of counting answers to Ψ into the problem of evaluating a linear combination of terms, each of which can be determined by counting the answers to a conjunctive query. This is done in Lemma 37. The second part (Section 4.2) concerns the lower bounds and is more challenging. The overall strategy is as follows: In the first step, we use a parsimonious reduction from 3-SAT to computing the reduced Euler Characteristic of a complex. The parsimonious reduction is due to Roune and Sáenz-de-Cabezón [61]. In combination with the Sparsification Lemma [44], this reduction becomes tight enough for our purposes. In the second step, we show how to encode a complex Δ into a union of acyclic conjunctive queries Ψ such that the following is true: The reduced Euler Characteristic of Δ is zero if and only if all terms in the homomorphism basis are acyclic. The lower bound results of Theorem 5 are established in Lemmas 51, and 52.

---

[8]We remark that [10, Theorem 7] focuses on the special case of graphs and *near* linear time algorithms. However, in the word RAM model with $O(\log n)$ bits, a linear time algorithm is possible [19].

4.1 Solving Meta via Inclusion-Exclusion

Lemma 36. *There is a deterministic algorithm that on input a quantifier-free UCQ $\Psi = (\mathcal{A}_1, \ldots, \mathcal{A}_\ell)$ decides in time $|\Psi|^{\mathrm{poly}(\log |\Psi|)} \cdot 2^{O(\ell)}$ whether all structures $\mathcal{A}$ with $c_\Psi(\mathcal{A}) \neq 0$ are acyclic.*

Proof. We rely on Babai's algorithm for isomorphism of vertex coloured graphs: This result [6, Corollary 4] implies that the isomorphism problem for vertex-coloured n-vertex graphs can be solved in time $n^{\mathrm{poly}(\log n)}$. Given a structure $\mathcal{A}$, the incidence graph of $\mathcal{A}$ is a vertex-coloured bipartite graph $B(\mathcal{A}) = (V_1, V_2, E)$ where $V_1 = U(\mathcal{A})$ and V_2 is the set of all tuples of $\mathcal{A}$. We colour each vertex in V_1 with colour U, and for each relation symbol R of the signature of $\mathcal{A}$, we colour all tuples $\vec{t} \in R^{\mathcal{A}}$ with colour R. It is then easy to see that two structures are isomorphic if and only if their incidence graphs are isomorphic (as vertex-coloured graphs). Thus, Babai's result implies that we can decide isomorphism between two structures $\mathcal{A}$ and $\mathcal{A}'$ in time $m^{\mathrm{poly}(\log m)}$ where $m = \max\{|\mathcal{A}|, |\mathcal{A}'|\}$.

We are now able to complete the proof. By Definition 23,

$$c_\Psi(\mathcal{A}) = \sum_{\substack{J \subseteq [\ell] \\ \bigwedge(\Psi|_J) \cong \mathcal{A}}} (-1)^{|J|+1}.$$

This suggests the following algorithm: for each subset $J \subseteq [\ell]$, compute $\bigwedge(\Psi|_J)$. Afterwards, using the aforementioned application of Babai's algorithm, collect the isomorphic terms and compute $c_\Psi(\bigwedge(\Psi|_J))$ for each $J \subseteq [\ell]$ in time $2^{O(\ell)} \cdot |\Psi|^{\mathrm{poly}(\log |\Psi|)}$ (observing that $|\bigwedge(\Psi|_J)| \in O(|\Psi|)$). Clearly, $c_\Psi(\mathcal{A}) = 0$ for every $\mathcal{A}$ that is not isomorphic to any $\bigwedge(\Psi|_J)$.

Finally, output 1 if each $\mathcal{A}$ with $c_\Psi(\mathcal{A}) \neq 0$ is acyclic (each of these checks can be done in linear time [62]), and output 0 otherwise. The total running time of this algorithm is bounded from above by $2^{O(\ell)} \cdot |\Psi|^{\mathrm{poly}(\log |\Psi|)}$. □

We are now able to provide our algorithm for Meta. As mentioned in the introduction, since the question asked by Meta is itself a complexity-theoretic question, we need to rely on a lower bound assumption for the correctness of the algorithm. For the case of arity 2, the Triangle Conjecture will suffice, but we remark that result below and its proof apply verbatim to the case of arbitrary arity (which is why we proved Lemma 36 for arbitrary arities); the only difference is that instead of assuming the Triangle Conjecture, the higher arity case relies on the Hyperclique Hypothesis (see [53, 54]). Since we wish to avoid introducing yet another lower bound assumption, we only present the case of arity 2.

Lemma 37. *When restricted to queries of arity 2, the problem Meta can be solved in time $|\Psi|^{\mathrm{poly}(\log |\Psi|)} \cdot 2^{O(\ell)}$, where ℓ is the number of conjunctive queries in the union, if the Triangle Conjecture is true.*

Proof. If the Triangle Conjecture is true, then Theorem 35 implies that counting answers to a quantifier-free conjunctive query of arity 2 is solvable in linear time if and only if the query is acyclic. In combination with Corollary 27 we obtain that counting answers to a UCQ $\Psi = (\mathcal{A}_1, \ldots, \mathcal{A}_\ell)$ can be done in linear time if and only if each $\mathcal{A}$ with $c_\Psi(\mathcal{A}) \neq 0$ is acyclic — recall that we consider the quantifier-free case in which #equivalence is the same as isomorphism, hence we do not need to consider #cores.

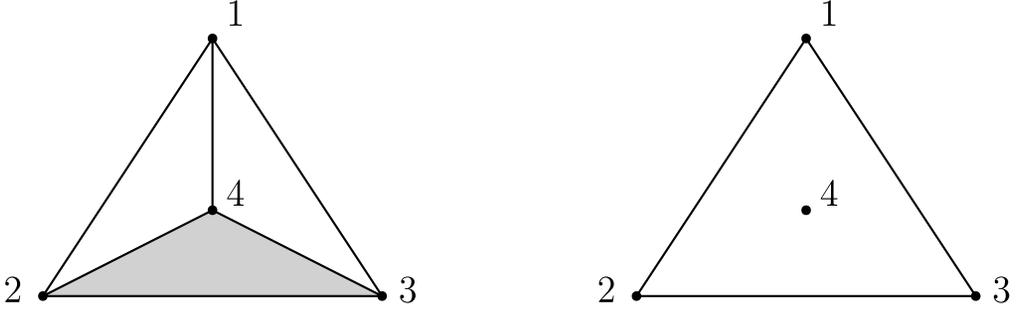Consequently, we only need to call the algorithm provided in Lemma 36. This concludes the proof. □

Fig. 1. Two complexes over the groundset $\Omega = \{1, 2, 3, 4\}$. Let Δ_1 be the complex shown on the left. It has facets $\{2, 3, 4\}, \{1, 2\}, \{1, 3\}$, and $\{1, 4\}$. Let Δ_2 be the complex shown on the right, with facets $\{1, 2\}, \{2, 3\}$, $\{1, 3\}$, and $\{4\}$. The reduced Euler characteristic of these complexes is computed as follows: Since Δ_1 has one face of size 3 ($\{2, 3, 4\}$), 6 faces of size 2, 4 faces of size 1, and the empty set as face of size 0 it holds that $\hat{\chi}(\Delta_1) = -(-1 + 6 - 4 + 1) = -2$. Similarly, we have $\hat{\chi}(\Delta_2) = -(3 - 4 + 1) = 0$.

4.2 Fine-grained Lower bounds for META

For our hardness proof for META, we will construct a reduction from the computation of the reduced Euler characteristic of an abstract simplicial complex. We begin by introducing some central notions about (abstract) simplicial complexes.

4.2.1 Simplicial Complexes.
A simplicial complex captures the geometric notion of an independence system. We will use the corresponding combinatorial description, which is also known as *abstract simplicial complex*, and defined as follows.

Definition 38. A *complex* (short for abstract simplicial complex) Δ is a pair consisting of a nonempty finite ground set Ω and a set of *faces* $\mathcal{I} \subseteq 2^\Omega$ that includes all singletons and is a downset. That is, the set of faces $\mathcal{I}$ satisfies the following criteria.

- $\forall S \subseteq \Omega : S \in \mathcal{I} \Rightarrow \forall S' \subseteq S : S' \in \mathcal{I}$, and
- $\forall x \in \Omega : \{x\} \in \mathcal{I}$.

The inclusion-maximal faces in $\mathcal{I}$ are called *facets*. Unless stated otherwise, we encode complexes by the ground set Ω and the set of facets. Then $|\Delta|$ is its encoding length.

Definition 39. The *reduced Euler characteristic* of a complex $\Delta = (\Omega, \mathcal{I})$ is defined as

$$\hat{\chi}(\Delta) := - \sum_{S \in \mathcal{I}} (-1)^{|S|} .$$

Consider for example the complexes Δ_1 and Δ_2 in Figure 1, given with a computation of their reduced Euler characteristic.

Let $\Delta = (\Omega, \mathcal{I})$ be a complex. Consider distinct elements x and y in Ω. We say that x *dominates* y if, for every $S \in \mathcal{I}$, $y \in S$ implies $S \cup \{x\} \in \mathcal{I}$. For example, x and y dominate each other if they are contained in the same facets. We start with a simple observation:

Lemma 40. *Let $\Delta = (\Omega, \mathcal{I})$ be a complex and let $x, y \in \Omega$. Then x dominates y if and only if each facet that contains y must also contain x.*

PROOF. For the forward direction suppose x dominates y. Let F be a facet that contains y. Then $F \cup \{x\} \in \mathcal{I}$. Since facets are inclusion maximal in $\mathcal{I}$ it follows that $F = F \cup \{x\}$, that is, $x \in F$.

For the backward direction suppose that each facet containing y must also contain x. Let $S \in \mathcal{I}$ with $y \in S$. Then there is a facet F with $S \subseteq F$. Since $x \in F$ we have $S \cup \{x\} \subseteq F$ and hence $S \cup \{x\} \in \mathcal{I}$. Therefore x dominates y.   □

We say that a complex $\Delta = (\Omega, \mathcal{I})$ is *irreducible* if, for every $y \in \Omega$, there is no $x \in \Omega \setminus \{y\}$ that dominates y. For example, both complexes in Figure 1 are irreducible.

Given $\Delta = (\Omega, \mathcal{I})$ and $y \in \Omega$, we define $\Delta \setminus y$ to be the complex obtained from Δ by deleting all faces that contain y and by deleting y from Ω. The following lemma seems to be folklore. We include a proof only for reasons of self-containment.

Lemma 41. *Let $\Delta = (\Omega, \mathcal{I})$ be a complex. If $y \in \Omega$ is dominated by some $x \in \Omega \setminus \{y\}$ then $\hat{\chi}(\Delta) = \hat{\chi}(\Delta \setminus y)$.*

Proof. Let $\Delta = (\Omega, \mathcal{I})$. Write $\mathcal{I}_y$ for the set of all faces containing y. Consider the mapping $b : \mathcal{I}_y \to \mathcal{I}_y$

$$b(S) := \begin{cases} S \cup \{x\} & x \notin S \\ S \setminus \{x\} & x \in S . \end{cases}$$

Note that b is well-defined since $S \cup \{x\} \in \mathcal{I}_y$ because $y \in S$ and x dominates y. Observe that b induces a partition of $\mathcal{I}_y$ in pairs $\{S, S \cup \{x\}\}$ for $x \notin S$. For those pairs, we clearly have $|S| + 1 = |S \cup \{x\}|$. Thus

$$\sum_{S \in \mathcal{I}_y} (-1)^{|S|} = 0 ,$$

and therefore

$$\hat{\chi}(\Delta) = - \sum_{S \in \mathcal{I}} (-1)^{|S|} = - \sum_{S \in \mathcal{I} \setminus \mathcal{I}_y} (-1)^{|S|} = \hat{\chi}(\Delta \setminus y) .$$

□

Definition 42. Two complexes $\Delta_1 = (\Omega_1, \mathcal{I}_1)$ and $\Delta_2 = (\Omega_2, \mathcal{I}_2)$ are *isomorphic* if there is a bijection $b : \Omega_1 \to \Omega_2$ such that $S \in \mathcal{I}_1$ if and only if $b(S) \in \mathcal{I}_2$ for each $S \subseteq \Omega$, where $b(S) = \{b(x) \mid x \in S\}$.

Finally, a complex $(\Omega, \mathcal{I})$ is called *trivial* if it is isomorphic to $(\{x\}, \{\emptyset, \{x\}\})$.

4.2.2 The Main Reduction. To begin with, we require a conjunctive query whose answers cannot be counted in linear time under standard assumptions. To this end, we define, for positive integers k and t, a binary relational structure $\mathcal{K}_t^k$ as follows. Start with a t-clique K_t and k-stretch every edge, that is, each edge of K_t is replaced by a path consisting of k edges. We denote the resulting graph by K_t^k. For each edge e of K_t^k, we introduce a relation $R_e = \{e\}$ of arity 2. The structure $\mathcal{K}_t^k$ has universe $V(K_t^k)$ and relations $(R_e)_{e \in E(K_t^k)}$.

Observation 43. *Let k and t be positive integers. The structure $\mathcal{K}_t^k$ is self-join-free and has arity 2.*

Lemma 44. *Suppose that the non-uniform ETH holds. For all positive integers d, there is a t such that for each k, the function $\mathcal{D} \mapsto \hom(\mathcal{K}_t^k \to \mathcal{D})$ cannot be computed in time $O(|\mathcal{D}|^d)$.*

Proof. We again write K_t for the t-clique and K_t^k for the k-stretch of K_t. Chen, Eickmeyer and Flum [25] proved that, under non-uniform ETH, for each positive integer d, there is a t such that determining whether a graph G contains K_t as a subgraph cannot be done in time $O(|G|^d)$.

We construct a simple linear-time reduction to computing $\mathcal{D} \mapsto \hom(\mathcal{K}_t^k \to \mathcal{D})$. Since for each edge e of the underlying graph K_t^k, the structure $\mathcal{K}_t^k$ has a separate binary single-element relation R_e, the input database $\mathcal{D}$ (of the same signature as $\mathcal{K}_t^k$) also has a relation R_e' whose elements can be interpreted as e-coloured edges. This means that the problem of computing $\mathcal{D} \mapsto \hom(\mathcal{K}_t^k \to \mathcal{D})$

can equivalently be expressed as the following problem: Given a graph G', that comes with an edge-colouring $\text{col} : E(G') \mapsto E(K_t^k)$, count the homomorphisms h from K_t^k to G' such that for each edge e of K_t^k we have $\text{col}(h(e)) = e$.

We now reduce the problem of determining whether a graph G contains K_t to this problem. Given input G, we construct an edge-coloured graph G' as follows. Each edge of G is replaced by $\binom{t}{2}$ paths of length k, and we colour the edges of the i-th of those paths with the k edges of the k-stretch of the i-th edge e_i of K_t. It is easy to see that G contains a t-clique if and only if there is at least one homomorphism from K_t^k to G' that preserves the edge-colours. Moreover, the construction of G' can clearly be done in linear time.   □

Before proving Theorem 5, we need to take a detour to examine the complexity of computing the reduced Euler Characteristic of a complex associated with a UCQ. To begin with, we introduce the notion of a "power complex".

Definition 45. Let $\mathcal{U}$ be a finite set, and let $\Omega \subseteq 2^{\mathcal{U}}$ be a set system that does not contain $\mathcal{U}$. The *power complex* $\Delta_{\Omega,\mathcal{U}}$ is a complex with ground set Ω and faces

$$\mathcal{I} = \left\{ S \subseteq \Omega \; : \; \bigcup_{A \in S} A \neq \mathcal{U} \right\} .$$

It is easy to check that $\Delta_{\Omega,\mathcal{U}}$ is a complex — for each $x \in \Omega$, the set $\{x\}$ is in $\mathcal{I}$ since Ω does not contain $\mathcal{U}$. Although we typically encode complexes by the ground set Ω and the set of facets, in the case of a power complex we list the elements of $\mathcal{U}$ and Ω.

In the first step, we show that each complex is isomorphic to a power complex.

Lemma 46. *Let $\Delta = (\Omega, \mathcal{I})$ be a non-trivial irreducible complex. It is possible to compute, in polynomial time in $|\Delta|$, a set $\mathcal{U}$ and a set $\hat{\Omega} \subseteq 2^{\mathcal{U}}$ with $\mathcal{U} \notin \hat{\Omega}$ such that Δ is isomorphic to the power complex $\Delta_{\hat{\Omega},\mathcal{U}}$.*

PROOF. Let $F_1, \ldots, F_k$ be the facets of Δ so that the encoding of Δ consists of Ω and $F_1, \ldots, F_k$, and $|\Delta|$ is the length of this encoding. For each $i \in [k]$, we introduce an element E_i corresponding to F_i. Let $\mathcal{U} = \{E_1, \ldots, E_k\}$. Next, we define a mapping $b : \Omega \to 2^{\mathcal{U}}$ as follows:

$$b(x) := \{E_i \mid x \notin F_i\} .$$

Observe that b is injective since otherwise two elements of Ω are contained in the same facets of Δ, which means that they dominate each other. Also, note that $b(x) = \mathcal{U}$ would imply that x is not contained in any facet of Δ which gives a contradiction as $\{x\} \in \mathcal{I}$ for any $x \in \Omega$.

We choose $\hat{\Omega}$ as the range of b. Then, clearly, b is a bijection from Ω to $\hat{\Omega}$. Furthermore, $\mathcal{U}$ and $\hat{\Omega}$ can be constructed in time polynomial in $|\Delta|$.

It remains to prove that b is also an isomorphism from Δ to $\Delta_{\hat{\Omega},\mathcal{U}}$. To this end, we need to show that

$$S \in \mathcal{I} \Leftrightarrow \bigcup_{A \in b(S)} A \neq \mathcal{U} ,$$

where $b(S) = \{b(x) \mid x \in S\}$.

For the first direction, let $S \in \mathcal{I}$. W.l.o.g. we have $S \subseteq F_1$. Then, for all $x \in S$, $E_1 \notin b(x)$. Hence $E_1 \notin \bigcup_{A \in b(S)} A$. For the other direction, suppose that $S \notin \mathcal{I}$. Then S is not a subset of any facet. Consequently, for every $i \in [k]$, there exists $x_i \in S$ such that $x_i \notin F_i$. By definition of b, we thus have that, for every $i \in [k]$, there exists $x_i \in S$ such that $E_i \in b(x_i)$. Thus $\bigcup_{A \in b(S)} A = \mathcal{U}$.   □

Example 47. Recall the complex Δ_1 in Figure 1. Δ_1 has facets $F_1 = \{2, 3, 4\}$, $F_2 = \{1, 2\}$, $F_3 = \{1, 3\}$, and $F_4 = \{1, 4\}$. Since Δ_1 is irreducible and non-trivial, we can apply the construction in the

previous lemma to construct an isomorphic power complex: We set $\mathcal{U} = \{E_1, E_2, E_3, E_4\}$, since Δ_1 has 4 facets. Moreover, the ground set $\hat{\Omega}$ of the power complex is the range of the function $b : \Omega \rightarrow 2^{\mathcal{U}}$ with $b(x) := \{E_i \mid x \notin F_i\}$, that is, $b(1) = \{E_1\}$ since $1 \notin F_1$, and similarly, $b(2) = \{E_3, E_4\}$, $b(3) = \{E_2, E_4\}$, and $b(4) = \{E_2, E_3\}$. Then the facets $\hat{I}$ of the power complex are $\{E_1, E_3, E_4\}$, $\{E_1, E_2, E_4\}$, $\{E_1, E_2, E_3\}$, and $\{E_2, E_3, E_4\}$. Then Δ_1 and the power complex $\Delta_{\hat{\Omega}, \mathcal{U}}$ are isomorphic, where for instance $F_2 = \{1, 2\}$ corresponds to $b(1) \cup b(2) = \{E_1, E_3, E_4\}$. When applied to Δ_2, the same construction yields a power complex with $\mathcal{U} = \{E_1, E_2, E_3, E_4\}$ and ground set $\{\{E_3, E_4\}, \{E_1, E_4\}, \{E_2, E_4\}, \{E_1, E_2, E_3\}\}$.

We now introduce the main technical result that we use to establish lower bounds for META; note that computability result in the last part of Lemma 48 will just be required for the construction of exemplary classes of UCQs in the appendix.

Lemma 48. *For each positive integer t, there is a polynomial-time algorithm $\hat{\mathbb{A}}_t$ that, when given as input a non-trivial irreducible complex $\Delta = (\Omega, \mathcal{I})$ with $\Omega \notin \mathcal{I}$, computes a union of quantifier-free conjunctive queries $\Psi = (\mathcal{B}_1, \ldots, \mathcal{B}_\ell)$ satisfying the following constraints:*

(1) $\wedge (\Psi) \cong \mathcal{K}_t^k$ for some $k \geq 1$.
(2) $c_\Psi(\wedge (\Psi)) = -\hat{\chi}(\Delta)$.
(3) For all relational structures $\mathcal{B} \ncong \wedge (\Psi)$, $c_\Psi(\mathcal{B}) \neq 0$ implies that $\mathcal{B}$ is acyclic.
(4) $\ell \leq |\Omega|$.
(5) For all $i \in [\ell]$ the conjunctive query $\mathcal{B}_i$ is acyclic and self-join-free, and has arity 2.

Moreover, the algorithm $\hat{\mathbb{A}}_t$ can be explicitly constructed from t.

PROOF. The algorithm $\hat{\mathbb{A}}_t$ applies Lemma 46 to obtain $\mathcal{U}$ and $\hat{\Omega}$ such that the power-complex $\Delta_{\hat{\Omega}, \mathcal{U}}$ is isomorphic to Δ. Let $k = |\mathcal{U}|$ and assume w.l.o.g. that $\mathcal{U} = [k]$. Let $\hat{\Omega} = \{A_1, \ldots, A_\ell\}$. Let e^j be the j-th edge of the k-stretch of e in K_t^k and recall that $\mathcal{K}_t^k$ contains the relations $R_{e^i} = \{e^i\}$ for each $e \in E(K_t)$ and $i \in [k]$. For each $i \in [k]$ we define $\mathcal{E}_i$ to be the substructure of $\mathcal{K}_t^k$ containing the universe $V(K_t^k)$ and the relations R_{e^i} (one relation for each $e \in E(K_t)$). The algorithm constructs $\Psi = (\mathcal{B}_1, \ldots, \mathcal{B}_\ell)$ as follows: For each $j \in [\ell]$, set $\mathcal{B}_j = \cup_{i \in A_j} \mathcal{E}_i$, that is, $\mathcal{B}_j$ is the substructure of $\mathcal{K}_t^k$ obtained by taking all relations in $\mathcal{E}_i$ for all $i \in A_j$. The pseudocode for $\hat{\mathbb{A}}_t$ is as follows.

We now prove the required properties of Ψ.

(1) Recall that Ω is not a facet. Since Δ and $\Delta_{\hat{\Omega}, \mathcal{U}}$ are isomorphic, $\hat{\Omega}$ is not facet. Thus $\bigcup_{i \in [\ell]} A_i = \mathcal{U}$ by the definition of power complexes. Since $\mathcal{B}_j$ contains all relations in $\mathcal{E}_i$ with $i \in A_j$, we have that $\wedge (\Psi) = \mathcal{K}_t^k$ as desired.

(2) Recall from Definition 23 that $\mathcal{I}(\mathcal{B}, X) = \{J \subseteq [\ell] \mid (\mathcal{B}, X) \sim \wedge (\Psi|_J)\}$. Since we are in the setting of quantifier-free queries ($X = U(\mathcal{B})$; we will drop the X as before), equivalence ($\sim$) is equivalent to isomorphism. Thus, we have that $c_\Psi(\wedge (\Psi))$ is equal to

$$\sum_{\substack{J \subseteq [\ell] \\ \wedge(\Psi) \cong \wedge(\Psi|_J)}} (-1)^{|J|+1} = \sum_{\substack{J \subseteq [\ell] \\ \wedge(\Psi) \ncong \wedge(\Psi|_J)}} (-1)^{|J|} = \sum_{\substack{J \subseteq [\ell] \\ \cup_{j \in J} A_j \neq \mathcal{U}}} (-1)^{|J|} = -\hat{\chi}(\Delta_{\hat{\Omega}, \mathcal{U}}) = -\hat{\chi}(\Delta),$$

where the first equality follows from the formula $\sum_{J \subseteq [\ell]} (-1)^{|J|} = 0$.

(3) If $\mathcal{B} \ncong \wedge (\Psi)$ and $c_\Psi(\mathcal{B}) \neq 0$, then $\mathcal{B}$ is a substructure of $\mathcal{K}_t^k$ that is missing at least one of the $\mathcal{E}_i$. The claim holds since each $\mathcal{E}_i$ is a feedback edge set (i.e., the deletion of all tuples in $\mathcal{E}_i$ breaks every cycle), because deleting $\mathcal{E}_i$ corresponds to deleting one edge from each stretch.

(4) Follows immediately by construction.

Algorithm 1 Pseudocode for $\hat{\mathbb{A}}_t$

1: **Input:** $\Delta = (\Omega, \mathcal{I})$
2: $(F_1, \ldots, F_k) \leftarrow$ facets of Δ
3: **for** $x \in \Omega$ **do**
4:      $A_x \leftarrow \{i \mid x \notin F_i\}$ ▷ $\mathcal{U} = [k]$ and $\hat{\Omega} = \{A_x \mid x \in \Omega\} = \{A_1, \ldots, A_\ell\}$ yield the power complex
5: **end for**
6: $K_t \leftarrow t$-clique, $m \leftarrow \binom{t}{2}$ ▷ e_i denotes the i-th edge of K_t
7: $K_t^k \leftarrow k$-stretch of K_t ▷ e_i^j denotes the j-th edge of the k-stretch of $e_i \in E(K_t)$
8: **for** $i \in [m]$, $j \in [k]$ **do**
9:      $R_{e_i^j} \leftarrow \{e_i^j\}$
10: **end for**
11: **for** $i \in [k]$ **do**
12:      $\mathcal{E}_i \leftarrow (V(K_k^t), R_{e_1^i}, \ldots, R_{e_m^i})$
13: **end for**
14: **for** $j \in [\ell]$ **do**
15:      $\mathcal{B}_j \leftarrow \bigcup_{i \in A_j} \mathcal{E}_i$
16: **end for**
17: **Output:** $\Psi = (\mathcal{B}_1, \ldots, \mathcal{B}_\ell)$

---

(5) The $\mathcal{B}_i$ are self-join-free and of arity 2 since they are substructures of $\mathcal{K}_t^k$ (see Observation 43). For acyclicity, we use the fact that for each j we have $A_j \neq \mathcal{U}$ by definition of the power complex, which implies the claim by the same argument as 3.

With all cases completed, the proof is concluded. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ □

To provide a concrete example, we apply the construction in Lemma 48 to the complexes Δ_1 and Δ_2 in Figure 1. For this example, we will choose $t = 3$. Since both Δ_1 and Δ_2 have four facets, we will choose $k = 4$. The structure $\mathcal{K}_t^k$ is depicted in Figure 2, together with six selected substructures corresponding to the $\mathcal{B}_j$ in the proof of the Lemma 48.

Let us start by applying $\hat{\mathbb{A}}_3$ to Δ_1. Recall from Example 47 that the power complex of Δ_1 has ground set $\{A_1, A_2, A_3, A_4\}$ with $A_1 = \{E_1\}$, $A_2 = \{E_3, E_4\}$, $A_3 = \{E_2, E_4\}$, and $A_4 = \{E_2, E_3\}$. Thus $\hat{\mathbb{A}}_3$ returns the UCQ $\Psi_1 = (\mathcal{S}_1, \mathcal{S}_{34}, \mathcal{S}_{24}, \mathcal{S}_{23})$. Similarly, applying $\hat{\mathbb{A}}_3$ to Δ_2 yields the UCQ $\Psi_2 = (\mathcal{S}_{24}, \mathcal{S}_{34}, \mathcal{S}_{14}, \mathcal{S}_{123})$.

For the purpose of illustration, let us write Ψ_1 and Ψ_2 as formulas. To this end, given a non-empty subset A of $\{1, 2, 3, 4\}$, define the conjunctive query φ_A as follows:

$$\varphi_A = \bigwedge_{a \in A} R_{e_1^a}(x_{a-1}, x_a) \wedge R_{e_2^a}(x_{4+a-1}, x_{4+a}) \wedge R_{e_3^a}(x_{8+a-1}, x_{8+a}),$$

where indices are taken modulo 12. Observe that the $\mathcal{S}_A$ depicted in Figure 2 are the structures associated to φ_A for $A \in \{\{1\}, \{2, 4\}, \{1, 4\}, \{3, 4\}, \{2, 3\}, \{1, 2, 3\}\}$. Then

$$\Psi_1(x_0, \ldots, x_{11}) = \varphi_1 \vee \varphi_{34} \vee \varphi_{24} \vee \varphi_{23}, \text{ and}$$
$$\Psi_2(x_0, \ldots, x_{11}) = \varphi_{24} \vee \varphi_{34} \vee \varphi_{14} \vee \varphi_{123}$$

Note that $\wedge(\Psi_1) = \wedge(\Psi_2) = \mathcal{K}_3^4$. Now recall that $\hat{\chi}(\Delta_1) \neq 0$ and $\chi(\hat{\Delta}_2) = 0$. Thus, by Item 2. of Lemma 48, it holds that $c_{\Psi_1}(\mathcal{K}_3^4) \neq 0$, which means that there is an acyclic CQ in the CQ expansion of Ψ_1. On the contrary, by Item 3. of Lemma 48, for all $(\mathcal{A}, X)$, we have that $c_{\Psi_2}(\mathcal{A}, X) \neq 0$ implies that
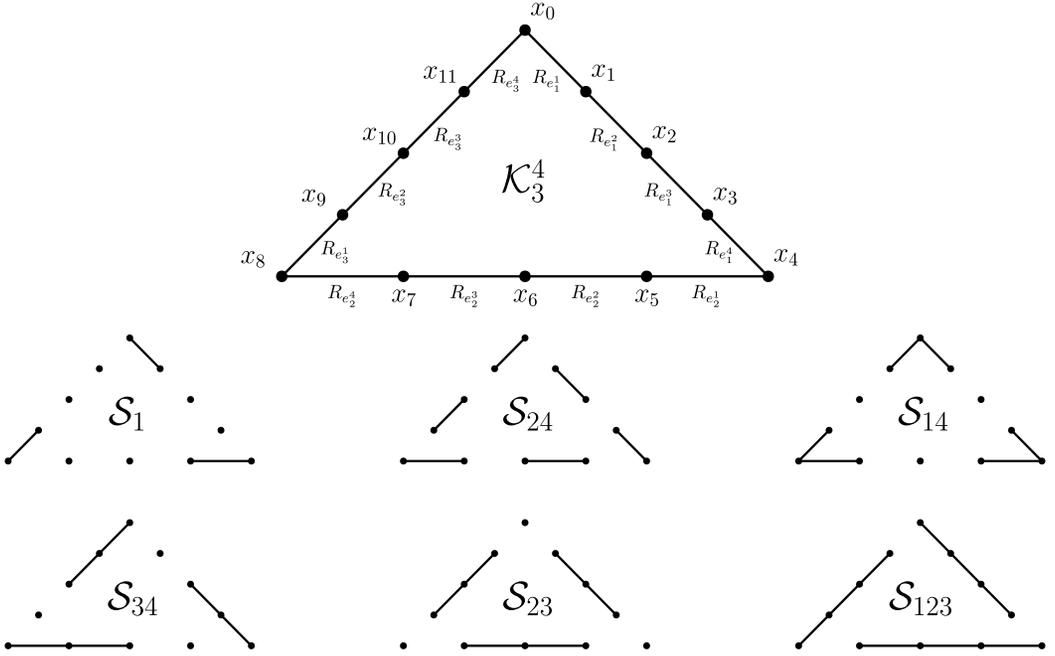
Fig. 2. *(Top:)* The structure $\mathcal{K}_3^4$. *(Bottom:)* Substructures $\mathcal{S}_A$ for some selected $A \subseteq [4]$. Observe that all of the $\mathcal{S}_A$ are *acyclic*.

$(\mathcal{A}, X)$ is acyclic. So, as a conclusion of our example, using complexity monotonicity (Corollary 27) and Theorem 35, we obtain as an immediate consequence:

Corollary 49. *While it is not possible to count answers to Ψ_1 in linear time (in the input database), unless the Triangle Conjecture fails, for Ψ_2 such a linear-time algorithm exists (even though $\wedge(\Psi_1) = \wedge(\Psi_2)$).*

PROOF. By Corollary 27, it is possible to count answers to a UCQ Ψ in linear time if and only if for each #minimal conjunctive query $(\mathcal{A}, X)$ with $c_\Psi(\mathcal{A}, X) \neq 0$, it is possible to count answers to $(\mathcal{A}, X)$ in linear time. Since Ψ_1 and Ψ_2 are quantifier-free, all conjunctive queries $\wedge(\Psi_1|_J)$ and $\wedge(\Psi_2|_J)$ are also quantifier-free and thus #minimal. For all $(\mathcal{A}, X)$, $c_{\Psi_2}(\mathcal{A}, X) \neq 0$ implies that $(\mathcal{A}, X)$ is acyclic, so it follows from Theorem 35 that it is possible to count answers to Ψ_2 in linear time. Moreover, given that $\mathcal{K}_3^4$ is not acyclic, Theorem 35 also implies that counting answers to $\mathcal{K}_3^4$ cannot be done in linear time, unless the Triangle Conjecture fails. Since $c_{\Psi_1}(\mathcal{K}_3^4) \neq 0$, counting answers to Ψ_1 cannot be done in linear time, unless the Triangle Conjecture fails.   □

We will now conclude the hardness proof for META.

Lemma 50. *For each positive integer t, there is a polynomial-time algorithm $\mathbb{A}_t$ that, when given as input a complex $\Delta = (\Omega, \mathcal{I})$, either computes $\hat{\chi}(\Delta)$, or computes a union of quantifier-free conjunctive queries $\Psi = (\mathcal{B}_1, \ldots, \mathcal{B}_\ell)$ satisfying the following constraints:*

(1) $\wedge(\Psi) \cong \mathcal{K}_t^k$ *for some $k \geq 1$.*
(2) $c_\Psi(\wedge(\Psi)) = -\hat{\chi}(\Delta)$.
(3) *For all relational structures $\mathcal{B} \not\cong \wedge(\Psi)$, $c_\Psi(\mathcal{B}) \neq 0$ implies that $\mathcal{B}$ is acyclic.*
(4) $\ell \leq |\Omega|$.

(5) *For all $i \in [\ell]$ the conjunctive query $\mathcal{B}_i$ is acyclic and self-join-free, and has arity 2.*

PROOF. Given $\Delta = (\Omega, \mathcal{I})$, we can successively apply Lemma 41 without changing the reduced Euler characteristic, until the resulting simplicial complex is irreducible. This can be done in polynomial time: By Lemma 40 it suffices to check whether there are $y \neq x \in \Omega$ such that each facet containing y also contains x. If no such pair exists, we are done. Otherwise we delete y from Ω and from all facets and continue recursively. Clearly, the number of recursive steps is bounded by $|\Omega|$ so the run time is at most a polynomial in $|\Delta|$.

If this process makes the complex trivial, we output 0 (i.e., the reduced Euler Characteristic of the trivial complex). We can furthermore assume that Ω is not a facet, i.e., that $\Omega \notin \mathcal{I}$, since in this case every subset of Ω is a face and the reduced Euler Characteristic is 0. We can thus assume that Δ is non-trivial and irreducible, and that Ω is not a facet. Therefore, we can use the algorithm $\hat{\mathbb{A}}_t$ from Lemma 48. This concludes the proof. □

We are now able to prove our lower bounds for META.

Lemma 51. *If the Triangle Conjecture is true then* META *is* NP-*hard. If the Triangle Conjecture and ETH are both true then* META *cannot be solved in time $2^{o(\ell)}$ where ℓ is the number of conjunctive queries in its input. Both results remain true even if the input to* META *is restricted to be over a binary signature.*

PROOF. For the first result, we assume the Triangle Conjecture and show that META is NP-hard. The input to META is a formula Ψ' which is a union of quantifier-free, self-join-free, and acyclic conjunctive queries. The goal is to decide whether counting answers of Ψ' (in an input database) can be done in linear time.

We reduce from 3-SAT. Let F be a 3-SAT formula. The first step of our reduction is to apply a reduction from [61]. Concretely, [61] gives a reduction that, given a 3-SAT formula F with n variables and m clauses, outputs in polynomial time a complex Δ such that F is satisfiable if and only if $\hat{\chi}(\Delta) \neq 0$. Moreover, the ground set of Δ has size $O(n + m)$.

Let $t = 3$ and let $\mathbb{A}_t$ be the algorithm from Lemma 50. Consider running $\mathbb{A}_t$ with input Δ. If $\mathbb{A}_t$ outputs $\hat{\chi}(\Delta)$ then we can check immediately whether $\hat{\chi}(\Delta) = 0$, which determines whether or not F is satisfiable. Otherwise, $\mathbb{A}_t$ outputs a formula Ψ which is a union of self-join-free, quantifier-free, and acyclic conjunctive queries of arity 2, and further has the property that $c_\Psi(\wedge(\Psi)) = -\hat{\chi}(\Delta)$. Since $\mathbb{A}_t$ is a polynomial-time algorithm, the number of conjunctive queries ℓ in Ψ is at most a polynomial in $n + m$.

We wish to show that determining whether counting answers of Ψ can be done in linear time would reveal whether or not $c_\Psi(\wedge(\Psi)) = 0$ (which would in turn reveal whether F is satisfiable).

By Corollary 27, it is possible to compute $\mathcal{D} \mapsto \text{ans}(\Psi \to \mathcal{D})$ in linear time if and only if, for each relational structure $\mathcal{A}$ with $c_\Psi(\mathcal{A}) \neq 0$, the function $\mathcal{D} \mapsto \text{hom}(\mathcal{A} \to \mathcal{D})$ can be computed in linear time.

So the intermediate problem is to check whether, for each relational structure $\mathcal{A}$ with $c_\Psi(\mathcal{A}) \neq 0$, the function $\mathcal{D} \mapsto \text{hom}(\mathcal{A} \to \mathcal{D})$ can be computed in linear time. We wish to show that solving the intermediate problem (in polynomial time) would enable us to determine whether or not $c_\Psi(\wedge(\Psi)) = 0$ (also in polynomial time).

Item 1 of Lemma 50 implies that there is a positive integer k such that $\wedge(\Psi) \cong \mathcal{K}_3^k$. Thus, $\wedge(\Psi)$ is not acyclic. However, Item 3 of Lemma 50 implies that for every relational structure $\mathcal{A} \not\cong \wedge(\Psi)$ in the intermediate problem, the structure $\mathcal{A}$ is acyclic. Theorem 35 shows (assuming the Triangle Conjecture) that, for each $\mathcal{A}$, $\mathcal{D} \mapsto \text{hom}(\mathcal{A} \to \mathcal{D})$ can be computed in linear time if and only if $\mathcal{A}$ is acyclic. We conclude that the answer to the intermediate problem is yes iff $c_\Psi(\wedge(\Psi)) = 0$, completing the proof that META is NP-hard.

To obtain the second result, we assume both the Triangle Conjecture and ETH. In this case, we apply the Sparsification Lemma [44] to the initial 3-SAT formula F before invoking the reduction. By the Sparsification Lemma, it is possible in time $2^{o(n)}$ to construct $2^{o(n)}$ formulas $F_1, F_2, \ldots$ such that F is satisfiable if and only if at least one of the F_i is satisfiable. Additionally, each F_i has $O(n)$ clauses. As before, for each such F_i we obtain in polynomial time a corresponding complex Δ_i, whose ground set has size $O(n)$. For each Δ_i, the algorithm from Lemma 50 either outputs its reduced Euler characteristic $\hat{\chi}(\Delta_i)$ or outputs a UCQ Ψ_i which has the property that that $c_\Psi(\wedge(\Psi_i)) = -\hat{\chi}(\Delta_i)$.

If, for any i, the algorithm outputs a value $\hat{\chi}(\Delta_i) \neq 0$ then F_i is satisfiable, so F is satisfiable. Let I be the set of indices i such that the algorithm outputs a UCQ Ψ_i. The argument from the first result shows that F is satisfiable if and only if there is an $i \in I$ such that counting answers to Ψ_i can be done in linear time.

We will argue that a $2^{o(\ell)}$ algorithm for META (where ℓ is the number of CQs in its input) would make it possible to determine in $2^{o(n)}$ time whether there is an $i \in I$ such that counting answers to Ψ_i can be done in linear time. This means that a $2^{o(\ell)}$ algorithm for META would make it possible to determine in $2^{o(n)}$ time whether F is satisfiable, contrary to ETH.

To do this, we just need to show that the number of CQs in Ψ_i, which we denote $\ell(\Psi_i)$, is $O(n)$. This follows since the ground set of Δ_i has size $O(n)$ and by Item 4 of Lemma 50, $\ell(\Psi_i)$ is at most the size of the ground set. Since the size of I is $2^{o(n)}$ the running time for determining whether F is satisfiable is $2^{o(n)}$ (for the sparsification) plus $|I| \operatorname{poly}(n) = 2^{o(n)}$ time (for computing the complexes Δ_i) plus $\sum_{i \in I} 2^{o(\ell(\Psi_i))} = 2^{o(n)}$ for the calls to META, contradicting ETH, as desired. □

The proof of the following lemma is analogous, with the only exception that we do not invoke Theorem 35 but apply Lemma 44 with $d = 1$ to obtain a t such that for each k the function $\mathcal{D} \mapsto \hom(\mathcal{K}_t^k \to \mathcal{D})$ cannot be evaluated in linear time.

Lemma 52. *If non-uniform ETH and the Triangle Conjecture are true then META is NP-hard and, furthermore,*

$$META \notin \bigcap_{\varepsilon > 0} \mathrm{DTime}(2^{\varepsilon \cdot \ell}) \, .$$

This remains true even if the input to META is restricted to be over a binary signature.

Theorem 5 now follows immediately from Lemmas 37, 51, and 52.

Finally, we point out that our construction shows, in fact, something much stronger than just the intractability of deciding whether we can count answers to a UCQ in linear time: For any pair (c, d) of positive integers satisfying $c \leq d$, it is hard to distinguish whether counting answers to a given UCQ can be done in time $O(n^c)$, or whether it takes time at least $\omega(n^d)$. Formally, we introduce the following gap problem:

Definition 53. Let c and d be positive integers with $c \leq d$. The problem META$[c, d]$ has as input a union of quantifier-free, self-join-free, and acyclic conjunctive queries Ψ. The goal is to decide whether the function $\mathcal{D} \mapsto \operatorname{ans}(\Psi \to \mathcal{D})$ can be computed in time $O(|\mathcal{D}|^c)$, or whether it cannot be solved in time $O(|\mathcal{D}|^d)$; the behaviour may be undefined for inputs Ψ for which the best exponent in the running time is in the interval $(c, d]$.

Theorem 54. *Assume that the Triangle Conjecture and non-uniform ETH hold. Then for each positive integer d, the problem META$[1, d]$ is NP-hard and, furthermore,*

$$META[1, d] \notin \bigcap_{\varepsilon > 0} \mathrm{DTime}(2^{\varepsilon \cdot \ell}) \, .$$

This remains true even if the input to META is restricted to be over a binary signature.

Proof. By Lemma 44 there is a positive integer t such that for all positive integers k, the function $\mathcal{D} \mapsto \hom(\mathcal{K}_t^k \to \mathcal{D})$ cannot be computed in time $O(|\mathcal{D}|^d)$. Fix this t and proceed similarly to the proof of Lemma 52.   □

Corollary 55 is an immediate consequence, since any algorithm that solves META$[c, d]$ for $1 \le c \le d$ solves, without modification, META$[1, d]$.

Corollary 55. *Assume that the Triangle Conjecture and non-uniform ETH hold. Then for every pair (c, d) of positive integers satisfying $c \le d$, the problem META$[c, d]$ is* NP*-hard and, furthermore,*

$$META[c, d] \notin \bigcap_{\varepsilon > 0} \mathrm{DTime}(2^{\varepsilon \cdot \ell}).$$

This remains true even if the input to META *is restricted to be over a binary signature.*

5   CONNECTION TO THE WEISFEILER-LEMAN-DIMENSION

Recall that we call a database a *labelled graph* if its signature has arity at most 2, and if it does not contain a self-loop, that is, a tuple of the form (v, v). Moreover, (U)CQs on labelled graphs must also have signatures of arity at most 2 and must not contain atoms of the form $R(v, v)$, where R is any relation symbol of the signature.

Neuen [57] and Lanzinger and Barceló [49] determined the WL-dimension of computing finite linear combinations of homomorphism counts to be the *hereditary treewidth*, defined momentarily. Since counting homomorphisms is equivalent to counting answers to quantifier-free conjunctive queries, and since the number of answers of a union of conjunctive queries can be expressed as a linear combination of conjunctive query answer counts (Lemma 24), we can state their results as follows.

Definition 56 (Hereditary Treewidth of UCQs). Let Ψ be a UCQ. The *hereditary treewidth* of Ψ, denoted by $\mathrm{hdtw}(\Psi)$, is defined as follows:

$$\mathrm{hdtw}(\Psi) = \max\{\mathrm{tw}(A, X) \mid c_\Psi(A, X) \ne 0\},$$

that is, $\mathrm{hdtw}(\Psi)$ is the maximum treewidth of any conjunctive query that survives with a non-zero coefficient when Ψ is expressed as a linear combination of conjunctive queries.

Then, applying the main result of Neuen, Lanzinger and Barceló to UCQs, we obtain:

Theorem 57 ([49, 57]). *Let Ψ be a quantifier-free UCQ on labelled graphs. Then $\dim_{\mathrm{WL}}(\Psi) = \mathrm{hdtw}(\Psi)$.*

This enables us to prove Theorem 7, which we restate for convenience.

Theorem 7. *There is an algorithm that computes a $O(\sqrt{\log k})$-approximation of the WL-dimension k of a quantifier-free UCQ on labelled graphs $\Psi = \varphi_1 \vee \cdots \vee \varphi_\ell$ in time $|\Psi|^{O(1)} \cdot O(2^\ell)$.*

Moreover, let $f : \mathbb{Z}_{>0} \to \mathbb{Z}_{>0}$ be any computable function. The problem of computing an f-approximation of $\dim_{\mathrm{WL}}(\Psi)$ given an input UCQ $\Psi = \varphi_1 \vee \cdots \vee \varphi_\ell$ is NP*-hard, and, assuming ETH, an f-approximation of $\dim_{\mathrm{WL}}(\Psi)$ cannot be computed in time $2^{o(\ell)}$.*

Proof. For the upper bound, we compute the coefficients

$$c_\Psi(\mathcal{A}, X) = \sum_{\substack{J \subseteq [\ell] \\ \wedge(\Psi|_J) \cong (\mathcal{A}, X)}} (-1)^{|J|+1},$$

that is, for each subset $J \subseteq [\ell]$, compute $\bigwedge (\Psi|_J)$. Afterwards, collect the isomorphic terms and compute $c_\Psi(\bigwedge (\Psi|_J))$ for each $J \subseteq [\ell]$. Clearly, $c_\Psi(\mathcal{A}, X) = 0$ for every $(\mathcal{A}, X)$ that is not isomorphic to any $\bigwedge (\Psi|_J)$. Clearly, this can be done in time $|\Psi|^{O(1)} \cdot O(2^\ell)$.

Next, for each $(\mathcal{A}, X)$ with $c_\Psi(\mathcal{A}, X) \neq 0$, we use the algorithm of Feige, Hajiaghayi, and Lee [32] to compute in polynomial time a g-approximation $S(\mathcal{A}, X)$ of the treewidth of $(\mathcal{A}, X)$, where $g(k) \in O(\sqrt{\log k})$. Finally, we output the maximum of $S(\mathcal{A}, X)$ over all $(\mathcal{A}, X)$ with $c_\Psi(\mathcal{A}, X) \neq 0$

For the lower bound, assume that there is a function $f : \mathbb{Z}_{>0} \to \mathbb{Z}_{>0}$ and an algorithm $\mathbb{A}$ that computes an f-approximation of $\dim_{\mathrm{WL}}(\Psi)$ in subexponential time in the number of conjunctive queries in the union. We will use $\mathbb{A}$ to construct a subexponential time algorithm for 3-SAT, which refutes ETH. Our construction is similar to the proof of Lemma 51. Fix any positive integer $t > f(1) + 1$.

Let F be a 3-CNF with n variables, which we can again assume to be sparse by using the Sparsification Lemma [18] (the details are identical to its application in the proof of Lemma 51). Next, using [61], we obtain a complex Δ, the reduced Euler characteristic of which is zero if and only if F is not satisfiable. Finally, we apply Lemma 50 with our choice of t. The corresponding algorithm computes in polynomial time either the reduced Euler characteristic of Δ, or otherwise outputs a UCQ Ψ such that the number ℓ of CQs in the union is bounded by $O(n)$. Moreover, the hereditary treewidth of Ψ is 1 if $\hat{\chi}(\Delta) = 0$, i.e., if F is not satisfiable; and its hereditary treewidth is at least $\mathrm{tw}(\mathcal{K}_t^k) = t - 1 > f(1)$, otherwise.

Thus, we run $\mathbb{A}$ on Ψ and report that F is satisfiable if and only if it outputs $S > f(1)$. Since $\mathbb{A}$ runs in time $2^{o(\ell)}$, the total running time is bounded by $2^{o(n)}$, which refutes ETH. NP-hardness follows likewise. □

Finally, the proof of Theorem 8 is identical with the only exception being that, since k is fixed, we can substitute the approximation algorithm for treewidth of Feige, Hajiaghayi, and Lee [32] by the exact algorithm of Bodlaender [15], which runs in polynomial time if k is fixed.

ACKNOWLEDGMENTS

REFERENCES

[1] Amir Abboud and Virginia Vassilevska Williams. 2014. Popular Conjectures Imply Strong Lower Bounds for Dynamic Problems. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*. IEEE Computer Society, 434–443. https://doi.org/10.1109/FOCS.2014.53

[2] Serge Abiteboul, Richard Hull, and Victor Vianu. 1995. *Foundations of Databases*. Addison-Wesley. http://webdam.inria.fr/Alice/

[3] Marcelo Arenas, Luis Alberto Croquevielle, Rajesh Jayaram, and Cristian Riveros. 2021. When is approximate counting for conjunctive queries tractable?. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, Samir Khuller and Virginia Vassilevska Williams (Eds.). ACM, 1015–1027. https://doi.org/10.1145/3406325.3451014

[4] Vikraman Arvind. 2016. The Weisfeiler-Lehman Procedure. *Bull. EATCS* 120 (2016). http://eatcs.org/beatcs/index.php/beatcs/article/view/442

[5] Vikraman Arvind, Frank Fuhlbrück, Johannes Köbler, and Oleg Verbitsky. 2022. On the Weisfeiler-Leman dimension of fractional packing. *Inf. Comput.* 288 (2022), 104803. https://doi.org/10.1016/j.ic.2021.104803

[6] László Babai. 2016. Graph isomorphism in quasipolynomial time [extended abstract]. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, Daniel Wichs and Yishay Mansour (Eds.). ACM, 684–697. https://doi.org/10.1145/2897518.2897542

[7] Guillaume Bagan, Arnaud Durand, and Etienne Grandjean. 2007. On Acyclic Conjunctive Queries and Constant Delay Enumeration. In *Computer Science Logic, 21st International Workshop, CSL 2007, 16th Annual Conference of the EACSL, Lausanne, Switzerland, September 11-15, 2007, Proceedings (Lecture Notes in Computer Science, Vol. 4646)*, Jacques Duparc and Thomas A. Henzinger (Eds.). Springer, 208–222. https://doi.org/10.1007/978-3-540-74915-8_18

[8] Pablo Barceló, Mikhail Galkin, Christopher Morris, and Miguel A. Romero Orth. 2022. Weisfeiler and Leman Go Relational. In *Learning on Graphs Conference, LoG 2022, 9-12 December 2022, Virtual Event (Proceedings of Machine Learning Research, Vol. 198)*, Bastian Rieck and Razvan Pascanu (Eds.). PMLR, 46. https://proceedings.mlr.press/v198/barcelo22a.html

[9] Pablo Barceló, Miguel Romero, and Moshe Y. Vardi. 2016. Semantic Acyclicity on Graph Databases. *SIAM J. Comput.* 45, 4 (2016), 1339–1376. https://doi.org/10.1137/15M1034714

[10] Suman K. Bera, Lior Gishboliner, Yevgeny Levanzov, C. Seshadhri, and Asaf Shapira. 2022. Counting Subgraphs in Degenerate Graphs. *J. ACM* 69, 3 (2022), 23:1–23:21. https://doi.org/10.1145/3520240

[11] Christoph Berkholz, Fabian Gerhardt, and Nicole Schweikardt. 2020. Constant delay enumeration for conjunctive queries: a tutorial. *ACM SIGLOG News* 7, 1 (2020), 4–33. https://doi.org/10.1145/3385634.3385636

[12] Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. 2017. Answering Conjunctive Queries under Updates. In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2017, Chicago, IL, USA, May 14-19, 2017*, Emanuel Sallinger, Jan Van den Bussche, and Floris Geerts (Eds.). ACM, 303–318. https://doi.org/10.1145/3034786.3034789

[13] Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. 2018. 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 (LIPIcs, Vol. 98)*, Benny Kimelfeld and Yael Amsterdamer (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 8:1–8:19. https://doi.org/10.4230/LIPIcs.ICDT.2018.8

[14] Christoph Berkholz and Nicole Schweikardt. 2019. Constant Delay Enumeration with FPT-Preprocessing for Conjunctive Queries of Bounded Submodular Width. In *44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, August 26-30, 2019, Aachen, Germany (LIPIcs, Vol. 138)*, Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 58:1–58:15. https://doi.org/10.4230/LIPIcs.MFCS.2019.58

[15] Hans L. Bodlaender. 1996. A Linear-Time Algorithm for Finding Tree-Decompositions of Small Treewidth. *SIAM J. Comput.* 25, 6 (1996), 1305–1317. https://doi.org/10.1137/S0097539793251219

[16] Johann Brault-Baron. 2013. *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).* Ph. D. Dissertation. University of Caen Normandy, France. https://tel.archives-ouvertes.fr/tel-01081392

[17] Jin-yi Cai, Martin Fürer, and Neil Immerman. 1992. An optimal lower bound on the number of variables for graph identification. *Comb.* 12, 4 (1992), 389–410. https://doi.org/10.1007/BF01305232

[18] Chris Calabro, Russell Impagliazzo, and Ramamohan Paturi. 2009. The Complexity of Satisfiability of Small Depth Circuits. In *Parameterized and Exact Computation, 4th International Workshop, IWPEC 2009, Copenhagen, Denmark, September 10-11, 2009, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 5917)*, Jianer Chen and Fedor V. Fomin (Eds.). Springer, 75–85. https://doi.org/10.1007/978-3-642-11269-0_6

[19] Nofar Carmeli, Shai Zeevi, Christoph Berkholz, Alessio Conte, Benny Kimelfeld, and Nicole Schweikardt. 2022. Answering (Unions of) Conjunctive Queries using Random Access and Random-Order Enumeration. *ACM Trans. Database Syst.* 47, 3 (2022), 9:1–9:49. https://doi.org/10.1145/3531055

[20] Ashok K. Chandra and Philip M. Merlin. 1977. Optimal Implementation of Conjunctive Queries in Relational Data Bases. In *Proceedings of the 9th Annual ACM Symposium on Theory of Computing, May 4-6, 1977, Boulder, Colorado, USA*, John E. Hopcroft, Emily P. Friedman, and Michael A. Harrison (Eds.). ACM, 77–90. https://doi.org/10.1145/800105.803397

[21] Hubie Chen and Stefan Mengel. 2015. A Trichotomy in the Complexity of Counting Answers to Conjunctive Queries. In *18th International Conference on Database Theory, ICDT 2015, March 23-27, 2015, Brussels, Belgium (LIPIcs, Vol. 31)*, Marcelo Arenas and Martín Ugarte (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 110–126. https://doi.org/10.4230/LIPIcs.ICDT.2015.110

[22] Hubie Chen and Stefan Mengel. 2016. Counting Answers to Existential Positive Queries: A Complexity Classification. 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*, Tova Milo and Wang-Chiew Tan (Eds.). ACM, 315–326. https://doi.org/10.1145/2902251.2902279

[23] Jianer Chen, Benny Chor, Mike Fellows, Xiuzhen Huang, David W. Juedes, Iyad A. Kanj, and Ge Xia. 2005. Tight lower bounds for certain parameterized NP-hard problems. *Inf. Comput.* 201, 2 (2005), 216–231. https://doi.org/10.1016/j.ic.2005.05.001

[24] Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. 2006. Strong computational lower bounds via parameterized complexity. *J. Comput. Syst. Sci.* 72, 8 (2006), 1346–1367. https://doi.org/10.1016/j.jcss.2006.04.007

[25] Yijia Chen, Kord Eickmeyer, and Jörg Flum. 2012. The Exponential Time Hypothesis and the Parameterized Clique Problem. In *Parameterized and Exact Computation - 7th International Symposium, IPEC 2012, Ljubljana, Slovenia, September 12-14, 2012. Proceedings (Lecture Notes in Computer Science, Vol. 7535)*, Dimitrios M. Thilikos and Gerhard J. Woeginger (Eds.). Springer, 13–24. https://doi.org/10.1007/978-3-642-33293-7_4

[26] Radu Curticapean, Holger Dell, and Dániel Marx. 2017. Homomorphisms are a good basis for counting small subgraphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, Hamed Hatami, Pierre McKenzie, and Valerie King (Eds.). ACM, 210–223. https://doi.org/10.1145/3055399.3055502

[27] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. 2015. *Parameterized Algorithms*. Springer. https://doi.org/10.1007/978-3-319-21275-3

[28] Víctor Dalmau and Peter Jonsson. 2004. The complexity of counting homomorphisms seen from the other side. *Theor. Comput. Sci.* 329, 1-3 (2004), 315–323. https://doi.org/10.1016/j.tcs.2004.08.008

[29] Holger Dell, Martin Grohe, and Gaurav Rattan. 2018. Lovász Meets Weisfeiler and Leman. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic (LIPIcs, Vol. 107)*, Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 40:1–40:14. https://doi.org/10.4230/LIPIcs.ICALP.2018.40

[30] Holger Dell, Marc Roth, and Philip Wellnitz. 2019. Counting Answers to Existential Questions. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece (LIPIcs, Vol. 132)*, Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 113:1–113:15. https://doi.org/10.4230/LIPIcs.ICALP.2019.113

[31] Arnaud Durand and Stefan Mengel. 2013. Structural tractability of counting of solutions to conjunctive queries. In *Joint 2013 EDBT/ICDT Conferences, ICDT '13 Proceedings, Genoa, Italy, March 18-22, 2013*, Wang-Chiew Tan, Giovanna Guerrini, Barbara Catania, and Anastasios Gounaris (Eds.). ACM, 81–92. https://doi.org/10.1145/2448496.2448508

[32] Uriel Feige, Mohammad Taghi Hajiaghayi, and James R. Lee. 2008. Improved Approximation Algorithms for Minimum Weight Vertex Separators. *SIAM J. Comput.* 38, 2 (2008), 629–657. https://doi.org/10.1137/05064299X

[33] Jörg Flum and Martin Grohe. 2006. *Parameterized Complexity Theory*. Springer. https://doi.org/10.1007/3-540-29953-X

[34] Jacob Focke, Leslie Ann Goldberg, Marc Roth, and Stanislav Zivný. 2022. Approximately Counting Answers to Conjunctive Queries with Disequalities and Negations. In *PODS '22: International Conference on Management of Data, Philadelphia, PA, USA, June 12 - 17, 2022*, Leonid Libkin and Pablo Barceló (Eds.). ACM, 315–324. https://doi.org/10.1145/3517804.3526231

[35] Jacob Focke and Marc Roth. 2022. Counting small induced subgraphs with hereditary properties. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, Stefano Leonardi and Anupam Gupta (Eds.). ACM, 1543–1551. https://doi.org/10.1145/3519935.3520008

[36] Martin Fürer. 2017. On the Combinatorial Power of the Weisfeiler-Lehman Algorithm. In *Algorithms and Complexity - 10th International Conference, CIAC 2017, Athens, Greece, May 24-26, 2017, Proceedings (Lecture Notes in Computer Science, Vol. 10236)*, Dimitris Fotakis, Aris Pagourtzis, and Vangelis Th. Paschos (Eds.). 260–271. https://doi.org/10.1007/978-3-319-57586-5_22

[37] Georg Gottlob, Gianluigi Greco, and Francesco Scarcello. 2014. Treewidth and Hypertree Width. In *Tractability: Practical Approaches to Hard Problems*, Lucas Bordeaux, Youssef Hamadi, and Pushmeet Kohli (Eds.). Cambridge University Press, 3–38. https://doi.org/10.1017/CBO9781139177801.002

[38] Georg Gottlob, Nicola Leone, and Francesco Scarcello. 2002. Hypertree decomposition and tractable queries. *J. Comput. System Sci.* 64, 3 (2002), 579–627. https://doi.org/10.1006/jcss.2001.1809

[39] Gianluigi Greco and Francesco Scarcello. 2014. Counting solutions to conjunctive queries: structural and hybrid tractability. In *Proceedings of the 33rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS'14, Snowbird, UT, USA, June 22-27, 2014*, Richard Hull and Martin Grohe (Eds.). ACM, 132–143. https://doi.org/10.1145/2594538.2594559

[40] Martin Grohe. 2007. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM* 54, 1 (2007), 1:1–1:24. https://doi.org/10.1145/1206035.1206036

[41] Martin Grohe and Dániel Marx. 2014. Constraint Solving via Fractional Edge Covers. *ACM Transactions on Algorithms* 11, 1 (2014), 4:1–4:20. https://doi.org/10.1145/2636918

[42] Martin Grohe, Thomas Schwentick, and Luc Segoufin. 2001. When is the evaluation of conjunctive queries tractable?. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece*, Jeffrey Scott Vitter, Paul G. Spirakis, and Mihalis Yannakakis (Eds.). ACM, 657–666. https://doi.org/10.1145/380752.380867

[43] Russell Impagliazzo and Ramamohan Paturi. 2001. On the Complexity of k-SAT. *J. Comput. Syst. Sci.* 62, 2 (2001), 367–375. https://doi.org/10.1006/jcss.2000.1727

[44] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. 2001. Which Problems Have Strongly Exponential Complexity? *J. Comput. Syst. Sci.* 63, 4 (2001), 512–530. https://doi.org/10.1006/jcss.2001.1774

[45] Richard M. Karp and Michael Luby. 1983. Monte-Carlo Algorithms for Enumeration and Reliability Problems. In *24th Annual Symposium on Foundations of Computer Science, Tucson, Arizona, USA, 7-9 November 1983*. IEEE Computer Society, 56–64. https://doi.org/10.1109/SFCS.1983.35

[46] Sandra Kiefer, Ilia Ponomarenko, and Pascal Schweitzer. 2019. The Weisfeiler-Leman Dimension of Planar Graphs Is at Most 3. *J. ACM* 66, 6 (2019), 44:1–44:31. https://doi.org/10.1145/3333003

[47] Phokion G. Kolaitis and Moshe Y. Vardi. 1995. On the Expressive Power of Datalog: Tools and a Case Study. *J. Comput. Syst. Sci.* 51, 1 (1995), 110–134. https://doi.org/10.1006/jcss.1995.1055

[48] Phokion G. Kolaitis and Moshe Y. Vardi. 2000. Conjunctive-Query Containment and Constraint Satisfaction. *J. Comput. Syst. Sci.* 61, 2 (2000), 302–332. https://doi.org/10.1006/jcss.2000.1713

[49] Matthias Lanzinger and Pablo Barceló. 2024. On the Power of the Weisfeiler-Leman Test for Graph Motif Parameters. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net. https://openreview.net/forum?id=FddFxi08J3

[50] Dániel Marx. 2010. Approximating fractional hypertree width. *ACM Trans. Algorithms* 6, 2 (2010), 29:1–29:17. https://doi.org/10.1145/1721837.1721845

[51] Dániel Marx. 2013. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. *J. ACM* 60, 6 (2013). https://doi.org/10.1145/2535926 Article No. 42.

[52] Stefan Mengel. 2021. A short note on the counting complexity of conjunctive queries. *CoRR* abs/2112.01108 (2021). arXiv:2112.01108 https://arxiv.org/abs/2112.01108

[53] Stefan Mengel. 2025. Lower Bounds for Conjunctive Query Evaluation. *CoRR* abs/2506.17702 (2025). https://doi.org/10.48550/ARXIV.2506.17702 arXiv:2506.17702

[54] Stefan Mengel. 2025. Lower Bounds for Conjunctive Query Evaluation. In *Companion of the 44th Symposium on Principles of Database Systems, PODS 2025, Berlin, Germany, June 22-27, 2025*, Floris Geerts and Benny Kimelfeld (Eds.). ACM, 5. https://doi.org/10.1145/3722234.3725824

[55] Carl A. Miller. 2013. Evasiveness of Graph Properties and Topological Fixed-Point Theorems. *Found. Trends Theor. Comput. Sci.* 7, 4 (2013), 337–415. https://doi.org/10.1561/0400000055

[56] Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. 2019. Weisfeiler and Leman Go Neural: Higher-Order Graph Neural Networks. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*. AAAI Press, 4602–4609. https://doi.org/10.1609/aaai.v33i01.33014602

[57] Daniel Neuen. 2024. Homomorphism-Distinguishing Closedness for Graphs of Bounded Tree-Width. In *41st International Symposium on Theoretical Aspects of Computer Science, STACS 2024, March 12-14, 2024, Clermont-Ferrand, France (LIPIcs, Vol. 289)*, Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 53:1–53:12. https://doi.org/10.4230/LIPICS.STACS.2024.53

[58] Reinhard Pichler and Sebastian Skritek. 2013. Tractable counting of the answers to conjunctive queries. *J. Comput. Syst. Sci.* 79, 6 (2013), 984–1001. https://doi.org/10.1016/j.jcss.2013.01.012

[59] Marc Roth. 2021. Parameterized Counting of Partially Injective Homomorphisms. *Algorithmica* 83, 6 (2021), 1829–1860. https://doi.org/10.1007/S00453-021-00805-Y

[60] Marc Roth and Johannes Schmitt. 2020. Counting Induced Subgraphs: A Topological Approach to #W[1]-hardness. *Algorithmica* 82, 8 (2020), 2267–2291. https://doi.org/10.1007/s00453-020-00676-9

[61] Bjarke Hammersholt Roune and Eduardo Sáenz-de-Cabezón. 2013. Complexity and algorithms for Euler characteristic of simplicial complexes. *J. Symb. Comput.* 50 (2013), 170–196. https://doi.org/10.1016/j.jsc.2012.07.003

[62] Robert Endre Tarjan and Mihalis Yannakakis. 1984. Simple Linear-Time Algorithms to Test Chordality of Graphs, Test Acyclicity of Hypergraphs, and Selectively Reduce Acyclic Hypergraphs. *SIAM J. Comput.* 13, 3 (1984), 566–579. https://doi.org/10.1137/0213035

[63] Moshe Y. Vardi. 2000. Constraint Satisfaction and Database Theory: a Tutorial. In *Proceedings of the Nineteenth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, May 15-17, 2000, Dallas, Texas, USA*, Victor Vianu and Georg Gottlob (Eds.). ACM, 76–85. https://doi.org/10.1145/335168.335209

[64] Mihalis Yannakakis. 1981. Algorithms for Acyclic Database Schemes. In *Very Large Data Bases, 7th International Conference, September 9-11, 1981, Cannes, France, Proceedings*. IEEE Computer Society, 82–94.

A   NECESSITY OF THE SIDE CONDITIONS IN THEOREM 3

We show that Theorem 3 is optimal in the sense that, if any of the conditions (I), (II), or (III) is dropped, the statement of the theorem becomes false (assuming that W[1]-hard problems are not fixed-parameter tractable).

Dropping condition (I).

Lemma 58. *There is a recursively enumerable class C of quantifier-free UCQs of bounded arity such that $\bigwedge (C)$ has unbounded treewidth but #UCQ(C) is fixed-parameter tractable. The class C satisfies (II) and (III).*

Proof. Let Δ be the second complex in Figure 1, that is, the ground set is $\Omega = \{1, 2, 3, 4\}$ and the facets are $\{1, 2\}$, $\{2, 3\}$, $\{3, 1\}$, and $\{4\}$. Note that Δ is irreducible (no element dominates another element), non-trivial, and Ω is not a facet. We can thus use Lemma 48 and let Ψ_t to be the output of algorithm $\hat{\mathbb{A}}_t$ given Δ. Note that Ψ_t is quantifier-free — in particular, this implies that all conjunctive queries within Ψ_t are #minimal. Since the algorithm $\hat{\mathbb{A}}_t$ can be explicitly constructed from t (see Lemma 48 and Algorithm 1) the class $C = \{\Psi_t \mid t \geq 1\}$ is recursively enumerable. Furthermore, all of the relation symbols in queries in UCQs in C have arity 2, so C has bounded arity.

By Item 1 of Lemma 48, $\bigwedge (C)$ has unbounded treewidth, since the treewidth of $\mathcal{K}_t^k$ is equal to $t - 1$. Moreover, by Item 5 of Lemma 48, all Ψ_t are unions of self-join-free conjunctive queries. We next show that #UCQ(C) is fixed-parameter tractable.

Recall that $\hat{\chi}(\Delta) = -(3 - 4 + 1) = 0$. Item 2 of Lemma 48 shows that $c_{\Psi_t}(\wedge(\Psi_t)) = 0$. Item 3 shows that for any relational structure $\mathcal{B}$ that is not isomorphic to $\wedge(\Psi_t)$ with $c_{\Psi_t}(\mathcal{B}) \neq 0$, $\mathcal{B}$ is acyclic. Recall that $\Gamma(C)$ is the class of those conjunctive queries that contribute to the CQ expansion of at least one UCQ in C. So all CQs in $\Gamma(C)$ are acyclic, which means that the treewidth of $\Gamma(C)$ is bounded by 1.

Since each query in $\Gamma(C)$ is quantifier-free, and is thus its own contract, $\Gamma(C) = \text{contract}(\Gamma(C))$. Thus, by Theorem 1, #UCQ(C) is fixed-parameter tractable.

We finish the proof by showing that C satisfies (II) and (III). Item (II) - that the number of existentially qunatified variables of queries in C is bounded - is trivial, because there are none. We have already noted (III) – that the UCQs in C are unions of self-join-free CQs.

□

Dropping condition (II).

Lemma 59. *There is a recursively enumerable and deletion-closed class C of unions of self-join-free conjunctive queries of bounded arity such that $\bigwedge (C)$ has unbounded treewidth but #UCQ(C) is fixed-parameter tractable.*

Proof. The statement of the lemma guarantees that C satisfies items (I) and (III) of Theorem 3.

Let $k \geq 3$ be a positive integer and let $\tau_k = (E_1, \ldots, E_k)$ be a signature with $\text{arity}(E_i) = 2$ for all $i \in [k]$. For any pair $i, j \in [k]$ with $i < j$, consider the conjunctive query

$$\varphi_k^{i,j}(x_1, \ldots, x_k, x_\perp) = \exists y_k^{i,j} : E_i(x_i, y_k^{i,j}) \wedge E_j(x_j, y_k^{i,j}) \wedge \bigwedge_{\ell \in [k] \setminus \{i,j\}} E_\ell(x_\ell, x_\perp).$$

Let $\Psi_k = \bigvee_{i < j \in [k]} \varphi_k^{i,j}$ and let C be obtained from the class $\{\Psi_k \mid k \geq 3\}$ by taking the closure under deletion of conjunctive queries. Clearly, C is recursively enumerable.

Note that each conjunctive query $\varphi_k^{i,j}$ is self-join-free and that Ψ_k contains $\binom{k}{2}$ existentially quantified variables. Thus the number of existentially quantified variables of queries in C is unbounded. Moreover, the treewidth of $\bigwedge (C)$ is unbounded. To see this, observe that

$$\bigwedge (\Psi_k) (x_1, \ldots, x_k) = \exists y_k^{1,2}, y_k^{1,3}, \ldots, y_k^{k-1,k} : \bigwedge_{i < j} E_i(x_i, y_k^{i,j}) \wedge E_j(x_j, y_k^{i,j}) \wedge \bigwedge_{\ell \in [k]} E_\ell(x_\ell, x_\perp).$$

Therefore, the Gaifman graph of $\bigwedge (\Psi_k)$ contains as a subgraph a subdivision of a k-clique and thus has treewidth at least $k - 1$.

It remains to show that #UCQ(C) is fixed-parameter tractable. To show this we claim that the classes $\Gamma(C)$ and $\text{contract}(\Gamma(C))$ have treewidth at most 2, and thus the problem #UCQ(C) is

fixed-parameter tractable by Theorem 1. To prove the claim, fix any $k \geq 3$ and any non-emtpy subset $J \subseteq \{(i,j) \in [k^2] \mid i < j\}$. We will show that the #minimal representatives of $\bigwedge (\Psi_k, J)$ and its contract are acyclic. To this end, assume first that $|J| = 1$. Then $\bigwedge (\Psi_k, J)$ is equal to one of the conjunctive queries $\varphi_k^{i,j}$, which is clearly acyclic. Since $\varphi_k^{i,j}$ is self-join-free and does not contain isolated variables, it is #minimal by Lemma 32. Let G be the Gaifman graph of $\varphi_k^{i,j}$, and recall that the contract of $\varphi_k^{i,j}$ is obtained from $G[X]$ by adding an edge between two free variables in X if and only if there is a connected component in the quantified variables that is adjacent to both free variables. Since the only quantified variable in $\varphi_k^{i,j}$ is $y_k^{i,j}$, which is adjacent (in G) to x_i and x_j, the contract of $\varphi_k^{i,j}$ is just the graph obtained from $G[X]$ by adding an edge between x_i and x_j, which also yields an acyclic graph.

Next assume that $|J| \geq 2$. For an index $s \in [k]$, we say that J *covers* s if each $(i,j) \in J$ satisfies $i = s$ or $j = s$. We distinguish three cases:

(A) There are distinct $s_1 < s_2$ such that J covers s_1 and s_2. Then $J = \{(s_1, s_2)\}$, contradicting the assumption that $|J| \geq 2$.

(B) There is precisely one $s \in [k]$ such that J covers s. Assume w.l.o.g. that $s = k$. Then since every $(i,j) \in J$ has $i < j$

$$\bigwedge (\Psi_k, J) = \bigwedge_{\ell \in [k-1]} E_\ell(x_\ell, x_\perp) \wedge \bigwedge_{(i,j) \in J} \exists y_k^{i,j} : E_i(x_i, y_k^{i,j}) \wedge E_j(x_j, y_k^{i,j})$$
$$= \bigwedge_{\ell \in [k-1]} E_\ell(x_\ell, x_\perp) \wedge \bigwedge_{(i,k) \in J} \exists y_k^{i,k} : E_i(x_i, y_k^{i,k}) \wedge E_k(x_k, y_k^{i,k}).$$

Observe that any answer of $\bigwedge (\Psi_k, J)$ in a database $\mathcal{D}$ is also an answer of the following query, and vice versa:

$$\psi_k := \bigwedge_{\ell \in [k-1]} E_\ell(x_\ell, x_\perp) \wedge \bigwedge_{i \in [k-1]} \exists y_k^{i,k} : E_i(x_i, y_k^{i,k}) \wedge E_k(x_k, y_k^{i,k}),$$

since all $y_k^{i,\ell}$ with $i < k$ can be mapped to the same vertex as $x_\perp$. Thus $\bigwedge (\Psi_k, J)$ and ψ_k are counting equivalent. Moreover, ψ_k is self-join-free and does not contain isolated variables. Thus, by Lemma 32, it is #minimal. Finally, deleting x_k from the Gaifman graph of ψ_k yields an acyclic graph, and the same is true for the contract of ψ_k. Therefore, the treewidth of both ψ_k and its contract are at most 2.

(C) There is no $s \in [k]$ such that J covers s. Then

$$\bigwedge (\Psi_k, J) = \bigwedge_{\ell \in [k]} E_\ell(x_\ell, x_\perp) \wedge \bigwedge_{(i,j) \in J} \exists y_k^{i,j} : E_i(x_i, y_k^{i,j}) \wedge E_j(x_j, y_k^{i,j}).$$

Observe that any answer of $\bigwedge (\Psi_k, J)$ in a database $\mathcal{D}$ is also an answer of the following query, and vice versa:

$$\psi_k := \bigwedge_{\ell \in [k]} E(x_\ell, x_\perp),$$

since all $y_k^{i,j}$ can be mapped to the same vertex as $x_\perp$.

Thus $\bigwedge (\Psi_k, J)$ and ψ_k are counting equivalent. Since ψ_k does not contain quantified variables, it must be both its own #core and its own contract. This concludes the proof of the claim since ψ_k is acyclic.

$\square$

Dropping condition (III).

Lemma 60. *There is a recursively enumerable and deletion-closed class C of quantifier-free UCQs of bounded arity such that $\bigwedge(C)$ has unbounded treewidth but #UCQ(C) is fixed-parameter tractable.*

Proof. The statement of the lemma guarantees that C satisfies items (I) and (II) of Theorem 3. We show an even stronger claim by requiring C to be a recursively enumerable class of quantifier-free CQs (instead of UCQs) of bounded arity such that $\bigwedge(C)$ has unbounded treewidth but #UCQ(C) is polynomial-time solvable. Note that each conjunctive query is a (trivial) union of conjunctive queries; moreover, this also means that C is deletion-closed.

For each $k \geq 1$, we define a conjunctive query ψ_k over the signature of graphs as follows:

$$\psi_k(x_1, \ldots, x_k, x_\perp) = \exists y : \bigwedge_{i \in [k]} E(x_i, x_\perp) \wedge E(x_i, y).$$

The query ψ_k has only one quantified variable. Moreover, the contract of ψ_k is a k-clique and thus has treewidth $k - 1$. However, ψ_k is clearly #equivalent to the query

$$\psi_k' = \bigwedge_{i \in [k]} E(x_i, x_\perp),$$

which is its own contract (since there are no quantified variables), and which is of treewidth 1. Thus, for C being the class of all ψ_k, we find that contract($\bigwedge(C)$) has unbounded treewidth, but, according to Theorem 19, the problem #UCQ(C) is solvable in polynomial time. □