# Computing the Difference of Conjunctive Queries Efficiently 

XIAO HU, University of Waterloo, Canada<br>QICHEN WANG, Hong Kong Baptist University, Hong Kong


#### Abstract

We investigate how to efficiently compute the difference result of two (or multiple) conjunctive queries, which is the last operator in relational algebra to be unraveled. The standard approach in practical database systems is to materialize the results for every input query as a separate set, and then compute the difference of two (or multiple) sets. This approach is bottlenecked by the complexity of evaluating every input query individually, which could be very expensive, particularly when there are only a few results in the difference. In this paper, we introduce a new approach by exploiting the structural property of input queries and rewriting the original query by pushing the difference operator down as much as possible. We show that for a large class of difference queries, this approach can lead to a linear-time algorithm, in terms of the input size and (final) output size, i.e., the number of query results that survive from the difference operator. We complete this result by showing the hardness of computing the remaining difference queries in linear time. Although a linear-time algorithm is hard to achieve in general, we also provide some heuristics that can provably improve the standard approach. At last, we compare our approach with standard SQL engines over graph and benchmark datasets. The experiment results demonstrate order-of-magnitude speedups achieved by our approach over the vanilla SQL engine.


CCS Concepts: • Information systems $\rightarrow$ Query optimization.
Additional Key Words and Phrases: conjunctive query, query optimization, difference operator
ACM Reference Format:
Xiao Hu and Qichen Wang. 2023. Computing the Difference of Conjunctive Queries Efficiently. Proc. ACM Manag. Data 1, 2, Article 153 (June 2023), 26 pages. https://doi.org/10.1145/3589298

## 1 INTRODUCTION

Conjunctive queries with aggregation, union, and difference (also known as negation) operators form the full relational algebra [12]. While conjunctive queries [ $15,16,24,29,33,38$ ], with aggregation [30] and unions [21, 23], have been extensively studied in the literature, the difference operator received much less attention. In modern database systems, there are several different equivalent expressions for computing the difference between two queries, such as NOT IN, NOT EXIST, EXCEPT, MINUS, DIFFERENCE, and LEFT-OUTER JOIN followed by a non-NULL filter. In contrast to its powerful expressibility, the execution plan of difference operator in existing database systems or data analytic engines (e.g., MySQL [2], Oracle [3], Postgre SQL [4], Spark SQL [6]) is quite brute-force. Given two (or multiple) conjunctive queries, their difference is simply done by materializing the answers for each participated conjunctive query separately, and then computing the difference of two (or multiple) sets. Hashing or other indexes may be built on top of the query answers to speed up the computation of the set difference at last. However, this approach is severely

[^0]
bottlenecked by evaluating every input query individually and materializing a large number of intermediate query results that do not contribute to the final results due to the difference operator.

Let's consider an example of friend recommendation in social networks (such as Twitter, Facebook, Sina Weibo). A friend recommendation is represented as a triple ( $a, b, c$ ) extracted from the network semantics, such that user $c$ is recommended to user $a$ since user $b$ is a friend of user $a$ and user $c$ is a friend of user $b$, together with other customized constraints. We also avoid the recommendation when $a$ and $c$ are already friends. The task of finding all valid recommendations can be captured by a SQL query in Example 1.1, as the difference of two sub-queries.

Example 1.1. Let Graph (src, dst) be a table storing all edges in the social network, and Triple( node 1, node2, node3) be a table storing all candidate recommendations. The following SQL query ( $Q$ ) finds all triples from Triple that do not form a triangle in the graph:

```
Q: SELECT node1, node2, node3 FROM Triple t1
    WHERE NOT EXISTS
    (SELECT * FROM Graph g1, Graph g2, Graph g3
    WHERE g1.dst = g2.src and g2.dst = g3.src and g3.dst = g1.src and g1.src = t1.node1
        and g2.src = t1.node2 and g3.src = t1.node3);
```

such that $Q$ is the difference of two sub-queries $Q_{1}$ and $Q_{2}$, where $Q_{1}$ returns all candidate recommendations from Triple and $Q_{2}$ returns all triangle friendship in the social network.

```
Q1: SELECT node1, node2, node3 FROM Triple t1
Q2: SELECT * FROM Graph g1, Graph g2, Graph g3
    WHERE g1.dst = g2.src and g2.dst = g3.src and g3.dst = g1.src and g1.src = t1.node1
        and g2.src = t1.node2 and g3.src = t1.node3;
```

We note that $Q$ can be rewritten as the following SQL query $Q^{\prime}$ :

```
Q': SELECT node1, node2, node3 FROM Triple t1
    WHERE NOT EXISTS
    (SELECT * FROM Triple t2
    WHERE EXISTS (SELECT * FROM graph g1 WHERE t2.node1 = g1.src and t2.node2 = g1.dst)
    AND EXISTS (SELECT * FROM graph g2 WHERE t2.node2 = g2.src and t2.node3 = g2.dst)
    AND EXISTS (SELECT * FROM graph g3 WHERE t2.node3 = g3.src and t2.node1 = g3.dst)
    AND t2.node1 = t1.node1 and t2.node2 = t1.node2 and t2.node3 = t1.node3)
```

such that $Q^{\prime}$ is the difference of $Q_{1}$ and another sub-query $Q_{3}$, where $Q_{3}$ finds all candidate recommendations in Triple that also form a triangle in the social network:

```
Q3: SELECT * FROM Triple t2
    WHERE EXISTS (SELECT * FROM graph g1 WHERE t2.node1 = g1.src and t2.node2 = g1.dst)
    And EXISTS (SELECT * FROM graph g2 WHERE t2.node2 = g2.src and t2.node3 = g2.dst)
    And EXISTS (SELECT * FROM graph g3 WHERE t2.node3 = g3.src and t2.node1 = g3.dst)
    AND t2.node1 = t1.node1 and t2.node2 = t1.node2 and t2.node3 = t1.node3
```

Figure 1(a) illustrates the execution plan for $Q$ generated by PostgreSQL optimizer. It first materializes all triangles in the graph as $Q_{2}$, and then computes the difference of $Q_{1}$ and $Q_{2}$ by anti-join. Moreover, hashing index is built on top of all triangles of $Q_{2}$ so that the anti-join can be executed by checking whether every candidate recommendation in $Q_{1}$ appears as a triangle in $Q_{2}$. At last, all "survived" recommendations are outputted as final answers. In plan (a), computing the set difference at last is the most time-consuming step, which is predicted to take 33.67 minutes by PostgreSQL optimizer. Although computing the subquery $Q_{2}$ is not that expensive, which only takes about 132 seconds, the number of intermediate results materialized for $Q_{2}$ is quite large as expected, which finally leads to the inefficiency of the subsequent computation on $Q_{1}-Q_{2}$. In Section 6, plan (a) actually runs in 308.175 seconds in practice.

To tackle the challenges brought by the difference operator, we take two input sub-queries as a whole into account for algorithm design. We are interested in efficient algorithms with running times linear in the final result size. This requirement rules out the standard approach of materializing the results for each input sub-query separately and then computing their set difference. Indeed, the final output size can be many magnitudes smaller than the number of intermediate results that materialized. To overcome the curse of large intermediate results, we introduce a rewritingbased approach by exploiting the joint structural properties of two input sub-queries and pushing the difference operator down as far as possible.

In Example 1.1, we can rewrite the original SQL query $Q$ into a new one $Q^{\prime}$. Instead of computing $Q_{2}$, it finds all candidate recommendations that also form a triangle in the social network as $Q_{3}$, which is exactly the intersection of $Q_{1}$ and $Q_{2}$, and then computes the difference of $Q_{1}$ and $Q_{3}$. Figure 1(b) illustrates the execution plan of this new query. We observe that computing the set difference at last is predicated to only take 21.67 seconds, which is much faster than (a). This is as expected, since the number of intermediate results generated by $Q_{3}$ is much smaller than $Q_{2}$, after
taking $Q_{1}$ into consideration, which is the key to the overall improvement. As the price, computing $Q_{3}$ is predicated to take a few more minutes than $Q_{2}$, but this is totally tolerable. In Section 6, plan (b) actually runs in 78.918 seconds, which already achieves $4 x$ speedup over (a). This significant improvement from (a) to (b) motivates us to further investigate this interesting problem for general queries.

Our contributions. In this paper, we formulate the difference of conjunctive queries ( DCQ ) problem and study the data complexity of this problem. Our contributions can be summarized as:

- Complexity Dichotomy: We give a dichotomy for computing DCQs in linear time in terms of input and output size. We characterize a class of "easy" DCQs exploiting the joint properties of two input CQs, and present a linear-time algorithm. On the other hand, we prove the hardness of obtaining a linear-time algorithm for the remaining "hard" DCQs via several well-known conjectures. (Section 3 and 4.1)
- Efficient Heuristic: We propose an efficient heuristic for computing "hard" DCQs, which does not lead to a linear-time algorithm but still improves the baseline approach greatly. The heuristic investigates the intersection of two input CQs and incorporates the state-of-the-art algorithms for CQ evaluation. (Section 4.2)
- Extension: We explore several interesting extensions. First, we design a recursive algorithm for computing the difference of multiple conjunctive queries. We also extend our algorithm to support other relational operators, such as selection, projection, join, and aggregation. At last, we investigate the DCQ problem under the bag semantics. (Section 5)
- Experimental Evaluation: We provide an experimental evaluation of our approach and standard approach on real-world datasets in both centralized and parallel database systems. The experimental results show that our approach out-performs the baseline on different classes of queries and datasets. (Section 6)
Roadmap. In Section 2, we formally define the DCQ problem and review the literature on evaluating a single CQ. In Section 3, we provide a linear-time algorithm for "easy" DCQs. In Section 4, we prove the hardness for the remaining DCQs and show efficient heuristics. In Section 5, we study several extensions of DCQs with other relational operators and bag semantics. In Section 6, we present the experimental evaluation. At last, we review related work in Section 7 and Section 8.


## 2 PRELIMINARIES

### 2.1 Problem Definition

Conjunctive Query (CQ). We consider the standard setting of multi-relational databases. Let $\mathbb{R}$ be a database schema that contains $n$ relations $R_{1}, R_{2}, \cdots, R_{n}$. Let $\mathcal{V}$ be the set of attributes in the database $\mathbb{R}$. Each relation $R_{i}$ is defined on a subset of attributes $e_{i} \subseteq \mathcal{V}$. Let $\mathcal{E}=\left\{e_{1}, e_{2}, \cdots, e_{n}\right\}$ be the set of the attributes for all relations. Let $\operatorname{dom}(x)$ be the domain of attribute $x \in \mathcal{V}$, and let $\operatorname{dom}(U)=\prod_{x \in U} \operatorname{dom}(x)$ be the domain of attributes $U \subseteq \mathcal{V}$.

Given the database schema $\mathbb{R}$, let an input instance be $D$, and the corresponding instances of $R_{1}, \cdots, R_{n}$ be $R_{1}^{D}, \cdots, R_{n}^{D}$. Where $D$ is clear from the context, we will drop the superscript and use $R_{1}, \cdots, R_{n}$ for both the schema and instances. Any tuple $t \in R_{i}$ is defined on $e_{i}$. For any attribute $x \in e_{i}, \pi_{x} t \in \operatorname{dom}(x)$ denotes the value of attribute $x$ in tuple $t$. Similarly, for a set of attributes $U \subseteq e_{i}, \pi_{U} t$ denotes the values of attributes in $U$ for $t$ with an implicit ordering on the attributes.

We consider the class of conjunctive queries without self-joins formally defined as

$$
\begin{equation*}
Q:=\pi_{\mathbf{y}}\left(\sigma_{\phi_{1}} R_{1}\left(e_{1}\right) \bowtie \cdots \sigma_{\phi_{2}} R_{2}\left(e_{2}\right) \bowtie \cdots \bowtie \sigma_{\phi_{n}} R_{n}\left(e_{n}\right)\right) \tag{1}
\end{equation*}
$$

where $\sigma$ is the selection operator, $\phi_{i}$ is the predicate defined over relation $R_{i}, \sigma_{\phi_{i}} R_{i}\left(e_{i}\right)$ selects out tuples from $R_{i}$ passing the predicate $\phi_{i}$ and $\mathbf{y} \subseteq \mathcal{V}$ denotes the output attributes. If $\mathbf{y}=\mathcal{V}$, such
a CQ query is known as full join, which represents the natural join of the underlying relations. We usually use a triple $(\mathbf{y}, \mathcal{V}, \mathcal{E})$ to represent a $\mathrm{CQ} \mathcal{Q}$, and simply use a pair $(\mathcal{V}, \mathcal{E})$ to represent a full join. Each relation $R_{i}$ in $Q$ is distinct, i.e., the CQ does not have a self-join. As a simplification, we ignore all the selection operators since it just takes $O(1)$ time to decide if a tuple passes the predicate $\phi_{i}$. Moreover, we assume every $R_{i}$ is defined on different subset of attributes; otherwise, we can simply keep the intersection of all relations defining on the same subset of attributes. Hence, we also use $R_{e}$ to denote the relation defined on $e \in \mathcal{E}$.

The result of $Q$ over instance $D$ noted as $Q(D)$, is defined as:

$$
Q(D)=\left\{t \in \operatorname{dom}(\mathbf{y}): \exists t^{\prime} \in \operatorname{dom}(\mathcal{V}), \text { s.t. } \pi_{\mathbf{y}} t^{\prime}=t, \pi_{e_{i}} t \in R_{i}, \forall i \in[n]\right\},
$$

i.e., the projection of all combinations of tuples from every relation onto $\mathbf{y}$, such that tuples in each combination have the same value(s) on the common attribute(s). Let $N=|D|$ be the input size, i.e., the total number of tuples in the input instance. Let OUT $=|Q(D)|$ be the output size, i.e., the number of query results of $Q$ over $D$.
Difference of Conjunctive Queries (DCQ). A DCQ $Q_{1}-Q_{2}$ consists of two CQs without selfjoins $Q_{1}, Q_{2}$ with the same output attributes. We also assume that the DCQ does not have a self-join, i.e., there exists no pair of relations $R_{i}$ from $Q_{1}$ and $R_{j}$ from $Q_{2}$ such that $R_{i}, R_{j}$ are the same. Note that our algorithms presented in this work also applied to the case when self-join exists in DCQ, but our lower bound assumes that no self-join exists. The input to a DCQ $Q_{1}-Q_{2}$ is a pair of database instances $D_{1}, D_{2}$ defined for $Q_{1}, Q_{2}$ respectively ${ }^{1}$. The result of $Q_{1}-Q_{2}$ over $D_{1}, D_{2}$ is $Q_{1}\left(D_{1}\right)-Q_{2}\left(D_{2}\right)$, i.e., tuples that appear in the result of $Q_{1}$ over instance $D_{1}$, but not in the result of $Q_{2}$ over instance $D_{2}$. Let $N=\left|D_{1}\right|+\left|D_{2}\right|$ be the input size, i.e., the total number of tuples in both input instances. Let OUT $=\left|Q_{1}\left(D_{1}\right)-Q_{2}\left(D_{2}\right)\right|$ be the output size.

In this paper, we adopt standard data complexity [36]; that is, we measure the complexity of algorithms with input size $N$ and output size OUT, and assume the query size as a constant.

### 2.2 Literature Review of CQ Evaluation

Before diving into the massive literature, we mention two classes of CQs that play an important role in query evaluation.

- ( $\alpha$-acyclic). A CQ $Q=(\mathbf{y}, \mathcal{V}, \mathcal{E})$ is $\alpha$-acyclic [17, 25] if there exists a tree $\mathcal{T}$ (called a join tree) such that each node in $\mathcal{T}$ corresponds to a relation in $\mathcal{E}$, and for each attribute $x \in \mathcal{V}$, the set of nodes containing $x$ form a connected subtree of $\mathcal{T}$. Moreover, we define $\operatorname{top}(x)$ as the highest node of $\mathcal{T}$ that attribute $x$ appears.
- (free-connex). A $\mathrm{CQ} Q=(\mathbf{y}, \mathcal{V}, \mathcal{E})$ is free-connex [16] if there exists a tree $\mathcal{T}$ (called a free-connex join tree) such that $\mathcal{T}$ is a join tree for $Q$, and for any pair of attributes $x_{1} \in \mathbf{y}, x_{2} \in \mathcal{V}-\mathbf{y}$, $\operatorname{top}\left(x_{2}\right)$ is not an ancestor of $\operatorname{top}\left(x_{1}\right)$. It has been proved equivalently that a $\mathrm{CQ} Q=(\mathbf{y}, \mathcal{V}, \mathcal{E})$ is free-connex if $Q$ is $\alpha$-acyclic and ( $\mathbf{y}, \mathcal{V}, \mathcal{E} \cup\{\mathbf{y}\}$ ) is also $\alpha$-acyclic.
Their relationships are illustrated in Figure 2. A free-connex CQ must be $\alpha$-acyclic. An $\alpha$-acyclic full join must be free-connex. Below, when the context is clear, we always refer "acyclic" to " $\alpha$-acyclic".

There has been a long line of research on CQ evaluation [17, 22, 31, 34, 36]. Yannakakis's seminal algorithm [17] was proposed for acyclic CQs, whose running time differs over different sub-classes of acyclic CQs. A free-connex CQ can be evaluated in $O(N+$ OUT $)$ time, which is already optimal since any algorithm needs to read input data and output all query results. On the other hand, an acyclic but non-free-connex CQ can be evaluated in $O(N \cdot$ OUT) time. Subsequent works have progressively defined different notions of "width" [26, 27], measuring how far a query is from being

[^1]

Fig. 2. Left is the classification of CQs via $\alpha$-cyclic/cyclic and full/non-full metrics. Dashed lines indicate the boundaries of different classes. Right is a join tree of an $\alpha$-acyclic full $\mathrm{CQ} Q=(\mathcal{V}, \mathcal{E})$ with $\mathcal{V}=\left\{x_{1}, x_{2}, \cdots, x_{8}\right\}$ and $\mathcal{E}=\left\{e_{1}=\left\{x_{1}, x_{2}, x_{3}\right\}, e_{2}=\left\{x_{1}, x_{4}\right\}, e_{3}=\left\{x_{2}, x_{3}, x_{5}\right\}, e_{4}=\left\{x_{5}, x_{6}\right\}, e_{5}=\left\{x_{3}, x_{7}\right\}, e_{6}=\left\{x_{5}, x_{8}\right\}\right\}$. It is also a free-connex join tree for non-full $\mathrm{CQ} Q^{\prime}=(\mathbf{y}, \mathcal{V}, \mathcal{E})$ with $\mathrm{y}=\left\{x_{1}, x_{2}, x_{3}, x_{4}\right\}$, but is not a free-connex join tree with $\mathrm{y}=\left\{x_{1}, x_{2}, x_{5}\right\}$, since top $\left(x_{3}\right)=R_{1}$ is an ancestor of $\operatorname{top}\left(x_{5}\right)=R_{3}$. The subtree in red is the reduced query of $Q^{\prime}$ after running Algorithm 1.
acyclic and tackle cyclic queries with decomposition. This line of algorithms run in $O\left(N^{w}+\right.$ OUT $)$ time, where $w$ can be the fractional hypertree width [26, 33], submodular width [14], or FAQwidth [14]. In addition, some specific classes of CQs can be speedup by fast matrix multiplication techniques [15, 18, 24], but we won't go into that direction further. CQ evaluation is still an actively researched problem; any improvement here will also improve DCQ evaluation when plugged into the baseline as well as our approach.

In the remaining, we often use $\operatorname{cost}(Q)$ to denote the time complexity of evaluating a $C Q Q$.
Implications to the Baseline Approach of DCQ Evaluation. Given a $\operatorname{DCQ} Q_{1}-Q_{2}$, the baseline approach of computing $Q_{1}, Q_{2}$ separately and then set difference incurs the following cost:

Corollary 2.1. Given two $C Q s Q_{1}$ and $Q_{2}$, the $D C Q Q_{1}-Q_{2}$ can be computed in $O\left(\operatorname{cost}\left(Q_{1}\right)+\right.$ $\left.\operatorname{cost}\left(Q_{2}\right)\right)$ time.

For example, when both $Q_{1}$ and $Q_{2}$ are free-connex, the baseline approach runs in $O\left(N+\mathrm{OUT}_{1}+\right.$ $\mathrm{OUT}_{2}$ ) time, where $\mathrm{OUT}_{1}, \mathrm{OUT}_{2}$ are the output sizes of $Q_{1}, Q_{2}$ respectively.

### 2.3 New Results of DCQ Evaluation

Our new complexity results for DCQ evaluation are summarized in Table 1. To help understand these results, we first introduce the class of linear-reducible CQs, and the reduce procedure.

Definition 2.2 (Linear-reducible). A CQ $\mathcal{Q}=(\mathbf{y}, \mathcal{V}, \mathcal{E})$ is linear-reducible if $(\mathbf{y}, \mathcal{V}, \mathcal{E} \cup\{\mathbf{y}\})$ is free-connex.

The relationship between linear-reducible CQs and existing classifications of CQs is illustrated in Figure 2. Any full or free-connex CQ must be linear-reducible. In addition, some cyclic but non-full CQs are also linear-reducible, for example, $Q=\pi_{x_{1}, x_{2}, x_{3}}\left(R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right) \bowtie R_{3}\left(x_{1}, x_{3}\right) \bowtie\right.$ $R_{4}\left(x_{3}, x_{4}\right)$ ), since adding $R_{5}\left(x_{1}, x_{2}, x_{3}\right)$ will result in a free-connex CQ (see Figure 2). It is also noted that any non-free-connex but acyclic CQ is non-linear-reducible.

Moreover, we introduce a reduce procedure in Algorithm 1, that can transform any linearreducible CQ into a full join query in $O(N)$ time, while preserving the query results. This algorithm is similar to the semi-join phase of Yannakakis algorithm [38]. Intuitively, we remove attributes or relations in a bottom-up ordering of nodes in a free-connex join tree. Recall that each node in the join tree corresponds to a relation in $\mathcal{E}$. When a node $e$ is visited, we distinguish two more cases: (line 4-5) if its output attributes are fully contained in its parent, we remove $e$ and update its parent relation via semi-joins; (line 6-7) and otherwise, we remove all non-output attributes (if

| $Q_{1}-Q_{2}$ | Baseline | Our Approach |  |
| :---: | :---: | :---: | :---: |
| $Q_{1}-Q_{2}$ is difference-linear | $\begin{gathered} \operatorname{cost}\left(Q_{1}\right) \\ + \\ \operatorname{cost}\left(Q_{2}\right) \end{gathered}$ | $N+$ OUT [Theorem 3.1] |  |
| $Q_{2}$ is linear-reducible |  | $\operatorname{cost}\left(Q_{1}\right) \quad$ [Corollary 2.5] |  |
| $Q_{2}$ is non-linear-reducible |  | $\operatorname{cost}\left(Q_{1}\right)+\min \left\{\begin{array}{l}\mathrm{OUT}_{1} \cdot \operatorname{cost}\left(Q_{2}^{\emptyset}\right) \\ \operatorname{cost}\left(Q_{2}^{\oplus}\right)\end{array}\right.$ | [Theorem 4.8] <br> [Theorem 4.10] |

Table 1. Summary of complexity results by baseline and our approach. $Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$ are two input CQs. $\left(\mathbf{y}, \mathcal{E}_{1}^{\prime}\right)$ and $\left(\mathbf{y}, \mathcal{E}_{2}^{\prime}\right)$ are reduced queries of $Q_{1}, Q_{2}$ respectively. $Q_{2}^{\emptyset}=(\emptyset, \mathcal{V}-\mathbf{y},\{e-\mathbf{y}$ : $\left.\left.e \in \mathcal{E}_{2}\right\}\right)$ and $Q_{2}^{\oplus}=\left(\mathbf{y}, \mathcal{V}_{2},\{\mathbf{y}\} \cup \mathcal{E}_{2}\right)$ are formally defined in Section 4.2. $N$ is the input size. $\mathrm{OUT}_{1}$, OUT are the output sizes of $Q_{1}, Q_{1}-Q_{2}$ respectively. $\operatorname{cost}(\cdot)$ is the time complexity of evaluating a single $C Q$.
there exists any) in $e$ via projections. The output of Algorithm 1 is a full join query $Q^{\prime}=(\mathbf{y}, \mathcal{E})$ (called the reduced query) and an instance $D^{\prime}$ (called the reduced instance) such that $Q(D)=Q^{\prime}\left(D^{\prime}\right)$. An example of reduced query is illustrated in Figure 2.

We are now ready to present the new results for DCQ evaluation.
Dichotomy for Linear-time Algorithm. Our main complexity result is a complete characterization of $Q_{1}, Q_{2}$ for which a linear algorithm can be achieved for computing $Q_{1}-Q_{2}$ :

Definition 2.3 (Difference-Linear). Given two CQs $Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$, the DCQ $Q_{1}-Q_{2}$ is difference-linear if $Q_{1}$ is free-connex, $Q_{2}$ is linear-reducible, and $\left.\left(\mathbf{y}, \mathcal{E}_{1}^{\prime} \cup\{e\}\right\}\right)$ is $\alpha$-acyclic for every $e \in \mathcal{E}_{2}^{\prime}$, where $\left(\mathbf{y}, \mathcal{E}_{1}^{\prime}\right)$ and $\left(\mathbf{y}, \mathcal{E}_{2}^{\prime}\right)$ are the reduced queries of $Q_{1}, Q_{2}$ respectively.
Theorem 2.4 (Diсhotomy). Given two $C Q s Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$, the $\operatorname{DCQ}$ $Q_{1}-Q_{2}$ can be computed in $O(N+$ OUT) time if and only if it is difference-linear.

Our proof of Theorem 2.4 consists of two steps. In Section 3, we prove the "if"-direction by designing a linear algorithm for the class of "easy" queries as characterized. In Section 4.1, we prove the "only-if" direction by showing the lower bound for the remaining class of "hard" queries, based on some well-established conjectures.

Improvement Achieved by Heuristics. For the class of "hard" DCQs on which obtaining a linear-time algorithm is hopeless, we further show some efficient heuristics that strictly improve the baseline. The complete results are presented in Section 4.2 and we mention an interesting case:

Corollary 2.5. Given two $C Q s Q_{1}$ and $Q_{2}$, if $Q_{2}$ is linear-reducible, then $D C Q Q_{1}-Q_{2}$ can be computed in $O\left(\operatorname{cost}\left(Q_{1}\right)\right)$ time.

We summarize all these results above in Table 1: (1) our approach strictly improves the baseline as long as $Q_{2}$ is linear-reducible; (2) furthermore, our approach leads to a linear-time algorithm if $Q_{1}$ also satisfies some specific conditions; (3) in the remaining case when $Q_{2}$ is non-linear-reducible, the comparison of our approach and baseline depends on specific queries or even input instances.

## 3 EASY DCQS

In this section, we show a linear-time algorithm for computing the class of "easy" DCQ characterized in Theorem 2.4. The main technique we used is simply query rewriting, but by exploiting the structures of two input queries in a non-trivial way.

Theorem 3.1. Given two $\operatorname{CQs} Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$, if $Q_{1}-Q_{2}$ is difference-linear, then $D C Q Q_{1}-Q_{2}$ can be computed in $O(N+\mathrm{OUT})$ time.

```
Algorithm 1: \(\operatorname{Reduce}(Q=(\mathbf{y}, \mathcal{V}, \mathcal{E}), D)\)
    \(\mathcal{T} \leftarrow\) the free-connex join tree of \((\mathbf{y}, \mathcal{V}, \mathcal{E} \cup\{\mathbf{y}\})\);
    foreach \(e \in \mathcal{E}\) in a bottom-up way (excluding root) of \(\mathcal{T}\) do
        Let \(e^{\prime} \leftarrow\) the (unique) parent node of \(e\);
        if \(e \cap \mathbf{y} \subseteq e^{\prime} \cap \mathbf{y}\) then
            \(R_{e^{\prime}} \leftarrow R_{e^{\prime}} \ltimes R_{e}\) and \(\mathcal{E} \leftarrow \mathcal{E}-\{e\} ;\)
        else
            if \(e \not \subset \mathrm{y}\) then \(R_{e} \leftarrow \pi_{e \cap \mathrm{y}} R_{e}\) and \(e \leftarrow e \cap \mathrm{y} ;\)
    return \(((\mathbf{y}, \mathcal{E}), D)\);
```

We start with a special class of DCQs that two input CQs share the same schema. In Section 3.1, we introduce an algorithm based on query rewriting, which always pushes the difference operator down to the input relations and avoids materializing a large number of intermediate results that do not participate in the final query result. In Section 3.2, we move to general case that $Q_{1}$ and $Q_{2}$ can have different schemas.

## $3.1 \quad Q_{1}$ and $Q_{2}$ share the same schema

We first note that if $Q_{1}, Q_{2}$ share the same schema, i.e., there is a one-to-one correspondence between the relations/attributes in $Q_{1}$ and $Q_{2}$, Theorem 3.1 degenerates to the following lemma:

Lemma 3.2. Given two $C Q s Q_{1}=Q_{2}=(\mathbf{y}, \mathcal{V}, \mathcal{E})$, if $(\mathbf{y}, \mathcal{V}, \mathcal{E})$ is free-connex, then the $D C Q Q_{1}-Q_{2}$ can be computed in $O(N+$ OUT) time.

Let's start with an example falling into this special case.
Example 3.3. Consider a $D C Q Q_{1}-Q_{2}$ with $Q_{1}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)$ and $Q_{2}=R_{1}^{\prime}\left(x_{1}, x_{2}\right) \bowtie$ $R_{2}^{\prime}\left(x_{2}, x_{3}\right)$. We can rewrite it as the union of two join queries: $Q_{1}-Q_{2}=\left(R_{1}-R_{1}^{\prime}\right) \bowtie R_{2}+R_{1} \bowtie\left(R_{2}-R_{2}^{\prime}\right)$, where the difference operator is only applied for computing $R_{1}-R_{1}^{\prime}$ and $R_{2}-R_{2}^{\prime}$. Intuitively, for every join result $(a, b, c) \in Q_{1}-Q_{2}$, it must be $(a, b) \notin R_{1}^{\prime}$ or $(b, c) \notin R_{2}^{\prime}$; otherwise, $(a, b, c) \in Q_{2}$, coming to a contradiction. The correctness of this rewriting will be formally presented in the proof of Lemma 3.4. In addition, the difference operators can be evaluated in $O(N)$ time, and the join operators can be evaluated in $O(N+$ OUT) time.

Rewrite Rule. We now generalize the rewriting rule in Example 3.9 to general DCQ $Q_{1}-Q_{2}$ for $Q_{1}=Q_{2}=(\mathbf{y}, \mathcal{V}, \mathcal{E})$, where $\mathbf{y}=\mathcal{V}$ and $(\mathcal{V}, \mathcal{E})$ is $\alpha$-acyclic. In other words, both $Q_{1}, Q_{2}$ correspond to the same acyclic full join query. For any $e \in \mathcal{E}$, let $R_{e}, R_{e}^{\prime}$ be the corresponding relations in $Q_{1}, Q_{2}$ respectively. Our rule is built on the observation that for any query result $t \in Q_{1}-Q_{2}, \pi_{e} t \in R_{e}$ must hold for every $e \in \mathcal{E}$, but $\pi_{e} t \notin R_{e}^{\prime}$ happens for some $e \in \mathcal{E}$. Applying this observation, we can rewrite such a DCQ as the (disjoint) union of a constant number of join queries as follows:

Lemma 3.4. Given two $\operatorname{CQs} Q_{1}=Q_{2}=(\mathbf{y}, \mathcal{V}, \mathcal{E})$, if $(\mathcal{V}, \mathcal{E})$ is $\alpha$-acyclic and $\mathbf{y}=\mathcal{V}, Q_{1}-Q_{2}=$ $\bigcup_{e \in \mathcal{E}}\left(\left(R_{e}-R_{e}^{\prime}\right) \bowtie\left(\bowtie_{e^{\prime} \in \mathcal{E}-\{e\}} R_{e^{\prime}}\right)\right)$.
Algorithm and Complexity. An algorithm directly follows the rewriting rule above. It first computes the difference of every pair of input relations, i.e., $R_{e}-R_{e}^{\prime}$ for each $e \in \mathcal{E}$, and then computes a full join query $\left(R_{e}-R_{e}^{\prime}\right) \bowtie Q_{1}$ derived for each $e \in \mathcal{E}$. Actually, we can handle a slightly larger class of DCQ. For $Q_{1}=Q_{2}=(\mathbf{y}, \mathcal{V}, \mathcal{E})$, if $(\mathbf{y}, \mathcal{V}, \mathcal{E})$ is free-connex, we simply remove all non-output attributes for $Q_{1}, Q_{2}$ separately in the preprocessing step, and then tackle two acyclic full joins, that share the same structure.

```
Algorithm 2: \(\operatorname{EAsyDCQ}\left(Q_{1}, Q_{2}, D_{1}, D_{2}\right)\)
    if \(\mathbf{y} \neq \mathcal{V}_{1}\) then \(\left(\left(\mathbf{y}, \mathcal{E}_{1}\right), D_{1}\right) \leftarrow \operatorname{Reduce}\left(Q_{1}, D_{1}\right)\);
    if \(\mathbf{y} \neq \mathcal{V}_{2}\) then \(\left(\left(\mathbf{y}, \mathcal{E}_{2}\right), D_{2}\right) \leftarrow \operatorname{Reduce}\left(Q_{2}, D_{2}\right)\);
    \(\mathcal{S} \leftarrow \emptyset ;\)
    foreach \(e \in \mathcal{E}_{2}\) do
        \(S_{e} \leftarrow \operatorname{Yannakakis}\left(\left(e, \mathbf{y}, \mathcal{E}_{1}\right), D_{1}\right)\);
        \(\mathcal{S} \leftarrow \mathcal{S} \cup\) Yannakakis \(\left(\left(\mathbf{y}, \mathcal{E}_{1} \cup\{e\}\right), D_{1} \cup\left\{S_{e}-R_{e}^{\prime}\right\}\right) ;\)
    return \(\mathcal{S}\);
```

As the pre-processing step and difference operators can be evaluated in $O(N)$ time, this algorithm is bottlenecked by evaluating the join query $(\mathcal{V}, \mathcal{E})$, which takes $O(N+$ OUT) time. Putting everything together, we come to Lemma 3.2.

## $3.2 \quad Q_{1}$ and $Q_{2}$ have different schemas

We next move to the general case when these two input CQs have different schemas. We focus on the case when both $Q_{1}, Q_{2}$ are full and then extend to non-full case.
DCQ with full CQs. Now, we assume that $\mathbf{y}=\mathcal{V}_{1}=\mathcal{V}_{2}=\mathcal{V}$. Theorem 3.1 simply degenerates to the Lemma 3.5.

Lemma 3.5. Given two full joins $Q_{1}=\left(\mathcal{V}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathcal{V}, \mathcal{E}_{2}\right)$, if $Q_{1}, Q_{2}$ are $\alpha$-acyclic, and $\left(\mathcal{V}, \mathcal{E}_{1} \cup\{e\}\right)$ is $\alpha$-acyclic for every $e \in \mathcal{E}_{2}$, then $Q_{1}-Q_{2}$ can be computed in $O(N+$ OUT) time.

A straightforward solution is to transform both $Q_{1}=\left(\mathcal{V}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathcal{V}, \mathcal{E}_{2}\right)$ into one auxiliary query $\left(\mathcal{V}, \mathcal{E}_{1} \cup \mathcal{E}_{2}\right)$, and then invoke the algorithm in Section 3.1 to handle the degenerated case. However, this solution does not necessarily lead to a linear-time algorithm. Let's gain some intuition from the example below.

Example 3.6. Consider a $D C Q Q_{1}-Q_{2}$ with $Q_{1}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}, x_{4}\right)$ and $Q_{2}=R_{3}\left(x_{1}, x_{2}, x_{3}\right)$ $\bowtie R_{4}\left(x_{3}, x_{4}\right)$. For an auxiliary query, we introduce the following intermediate relations $R_{5}=R_{1} \bowtie$ $\pi_{x_{2}, x_{3}} R_{2}, R_{6}=\pi_{x_{3}, x_{4}} R_{2}, R_{7}=\pi_{x_{1}, x_{2}} R_{3}$ and $R_{8}=\pi_{x_{2}, x_{3}} R_{3} \bowtie R_{4}$. Then, we can rewrite $Q_{1}, Q_{2}$ as follows:

$$
\begin{aligned}
& Q_{1}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}, x_{4}\right) \bowtie R_{5}\left(x_{1}, x_{2}, x_{3}\right) \bowtie R_{6}\left(x_{3}, x_{4}\right) \\
& Q_{2}=R_{7}\left(x_{1}, x_{2}\right) \bowtie R_{8}\left(x_{2}, x_{3}, x_{4}\right) \bowtie R_{3}\left(x_{1}, x_{2}, x_{3}\right) \bowtie R_{4}\left(x_{3}, x_{4}\right)
\end{aligned}
$$

Then, we are left with two queries that share the same schema. However, this strategy does not necessarily lead to a linear-time algorithm, since materializing the intermediate relation $R_{8}$ requires super-linear time, which could be much larger than the final output size OUT.

Careful inspection reveals that a simpler rewriting rule can avoid materializing $R_{8}$. More specifically, we keep $Q_{2}$ unchanged and rewrite $Q_{1}$ as above. Then, $Q_{1}-Q_{2}$ can be rewritten as $\left(R_{5}-R_{3}\right) \bowtie$ $R_{1} \bowtie R_{2}+R_{1} \bowtie R_{2} \bowtie\left(R_{6}-R_{4}\right)$. Intuitively, for every join result $(a, b, c, d) \in Q_{1}-Q_{2}$, it must be $(a, b, c) \notin R_{3}$ or $(c, d) \notin R_{4}$; otherwise, $(a, b, c, d) \in Q_{2}$, coming to a contradiction. The correctness of this rewriting will be formally presented in the proof of Lemma 3.7. In this case, materializing $R_{6}$ only takes $O(N)$ time, but materializing $R_{5}$ might take super-linear time. Fortunately, we can bound the size of $R_{5}$ by $O(N+$ OUT $)$. The rationale is that every tuple in $R_{5}-R_{3}$ will participate in at least on one join result of $\left(R_{5}-R_{3}\right) \bowtie R_{1} \bowtie R_{2}$, i.e., the final result of the difference query $Q_{1}-Q_{2}$, thus $\left|R_{5}-R_{3}\right| \leq$ OUT. For the difference operator, $R_{5}-R_{3}$ takes $O(N+\mathrm{OUT})$ time, and $R_{6}-R_{4}$ takes $O(N)$ time. For the join operator, both simple join queries take linear time in terms of their input size and output size. Overall, this rewriting rule can compute the example query in $O(N+$ OUT) time.

```
Algorithm 3: YANNAKAKis \((Q, D)\) [38]
    \((Q, D) \leftarrow \operatorname{Reduce}(Q, D)\);
    \(\mathcal{T} \leftarrow\) the free-connex join tree of \(Q\) rooted at \(r\);
    foreach \(v \in \mathcal{T}\) in a bottom-up way (excluding root) do
        \(R_{u} \leftarrow R_{u} \ltimes R_{v}\) for the parent node \(p(u)\) of \(u ;\)
    foreach \(u \in \mathcal{T}\) in a top-down way (excluding leaves) do
        foreach \(v\) is a child of \(u\) do \(R_{u} \leftarrow R_{u} \ltimes R_{v}\);
    foreach \(v \in \mathcal{T}\) in a bottom-up way (excluding root) do
        \(R_{u} \leftarrow R_{u} \bowtie R_{v}\) for the parent node \(p(u)\) of \(u\);
    return \(R_{r}\);
```

Rewrite Rule. Generalizing this observation, we develop the following rewriting rule for arbitrary full joins $Q_{1}, Q_{2}$. The high-level idea is to introduce an intermediate relation $S_{e}=\pi_{e} Q_{1}$ for every $e \in \mathcal{E}_{2}$, i.e., the projection of join results of $Q_{1}$ onto attributes $e$. Now we can rewrite $Q_{1}-Q_{2}$ using input relations in $Q_{1}$ and intermediate relations corresponding to $\mathcal{E}_{2}$, as well as input relations in $Q_{2}$, which results in the disjoint union of multiple full joins.

Lemma 3.7. Given two $C Q s Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$, if $\mathbf{y}=\mathcal{V}_{1}=\mathcal{V}_{2}, Q_{1}-Q_{2}=$ $\bigcup_{e \in \mathcal{E}_{2}}\left(\left(S_{e}-R_{e}^{\prime}\right) \bowtie Q_{1}\right)$, for $S_{e}=\pi_{e} Q_{1}$.

Proof. Direction $\subseteq$. Consider an arbitrary result $t \in Q_{1}-Q_{2}$. By definition, $\pi_{e_{1}} t \in R_{e_{1}}$ for every $e_{1} \in \mathcal{E}_{1}$, and $\pi_{e_{2}} t \notin R_{e_{2}}^{\prime}$ for some $e_{2} \in \mathcal{E}_{2}$. Moreover, $\pi_{e_{2}} t \in S_{e_{2}}=\pi_{e_{2}} Q_{1}$. So, $t \in\left(S_{e_{2}}-R_{e_{2}}^{\prime}\right) \bowtie Q_{1}$. Direction $\supseteq$. Consider an arbitrary $e_{2} \in \mathcal{E}_{2}$ and a query result $t \in\left(S_{e_{2}}-R_{e_{2}}^{\prime}\right) \bowtie Q_{1}$. By definition, $t \in Q_{1}$ and $t \notin R_{e_{2}}^{\prime}$, which further implies $t \notin Q_{2}$. This way, $t \in Q_{1}-Q_{2}$.

Algorithm and Complexity. An algorithm for computing the difference of two full join queries follows the rewriting rule above. For each $e \in \mathcal{E}_{2}$, it first materializes the query results of $\pi_{e} Q_{1}$, then computes the difference operator $\pi_{e} Q_{1}-R_{e}^{\prime}$, and finally the full join $\left(\pi_{e} Q_{1}-R_{e}^{\prime}\right) \bowtie Q_{1}$ by invoking the classical Yannakakis algorithm. We next analyze the complexity of the algorithm above. To establish the complexity, we first show an upper bound on the size of any intermediate relation constructed:

Lemma 3.8. $\left|S_{e}\right|=O(N+$ OUT $)$ for anye $\in \mathcal{E}_{2}$, where $S_{e}=\pi_{e} Q_{1}$.
Proof. Consider an arbitrary tuple $t \in S_{e}-R_{e}^{\prime}$. First, $t$ participates in at least one query result of $Q_{1}$. As $t \in S_{e}, t \in \pi_{e} Q_{1}$ by definition. There must exist some tuple $t^{\prime} \in Q_{1}$ such that $\pi_{e} t^{\prime}=t$. Thus, $t$ participates in some query results of $Q_{1}$. Meanwhile, $t$ does not participate in any query result of $Q_{2}$, since $t \notin R_{e}^{\prime}$. In this way, $t$ participates in at least one result in $Q_{1}-Q_{2}$, thus $\left|S_{e}-R_{e}^{\prime}\right| \leq$ OUT. Moreover, $\left|R_{e}^{\prime}\right| \leq N$. Together, we obtain $\left|S_{e}\right|=O(N+$ OUT $)$.

Let $\mathcal{V}=\mathcal{V}_{1}=\mathcal{V}_{2}$. If $\left(\mathcal{V}, \mathcal{E}_{1}\right)$ is $\alpha$-acyclic, and $\left(\mathcal{V}, \mathcal{E}_{1} \cup\{e\}\right)$ is also $\alpha$-acyclic for every $e \in \mathcal{E}_{2}$, then the constructed $\mathrm{CQ} \pi_{e} Q$ is free-connex. Implied by the existing result on CQ evaluation, $S_{e}$ can be computed in $O\left(N+\left|S_{e}\right|\right)=O(N+$ OUT $)$ time by the classic Yannakakis algorithm, where OUT is the output size of the difference query! The invocation of Yannakakis algorithm here is crucial for achieving linear complexity. For example, if $S_{e}$ is computed by first materializing the query results of $Q_{1}$ and then computing their projection onto $e$, the time complexity would be as large as $O\left(N+\mathrm{OUT}_{1}\right)$, where $\mathrm{OUT}_{1}$ is the output size of $Q_{1}$. Now, each full join $\left(S_{e}-R_{e}^{\prime}\right) \bowtie Q_{1}$ is $\alpha$-acyclic with input size $O(N+$ OUT $)$ and output size OUT, thus can be computed in $O(N+$ OUT $)$ time.

Therefore, the total time complexity is bounded by $O(N+\mathrm{OUT})$, since there are $O(1)$ sub-queries in $Q_{1}-Q_{2}$. Putting everything together, we come to Lemma 3.5.
DCQ with general CQs. Now, we are ready to present an linear-time algorithm for computing $Q_{1}-Q_{2}$, such that $Q_{1}$ is free-connex, $Q_{2}$ is linear-reducible, and $\left.\left(\mathbf{y}, \mathcal{E}_{1}^{\prime} \cup\{e\}\right\}\right)$ is $\alpha$-acyclic for every $e \in \mathcal{E}_{2}^{\prime}$, where $\left(\mathbf{y}, \mathcal{E}_{1}^{\prime}\right)$ and $\left(\mathbf{y}, \mathcal{E}_{2}^{\prime}\right)$ are the reduced queries of $Q_{1}, Q_{2}$ respectively. As described in Algorithm 2, we first apply a preprocessing step to $Q_{1}$ and $Q_{2}$ (line 1-4), which removes non-output attributes in $Q_{1}$ and $Q_{2}$ if they are non-full.

As shown in Algorithm 1, this reduce step is quite standard by first building a free-connex join tree for the derived query ( $\mathbf{y}, \mathcal{V}, \mathcal{E} \cup\{\mathbf{y}\}$ ), and then traversing the tree in a bottom-up way. In the traversal, when a relation is visited and contains some non-output attributes, we just update its parent relation by applying a semi-join and removing it. Note that if a relation does not contain any non-output attribute, then its ancestor also does not contain any, implied by the property of the free-connex join tree. Thus, the residual tree is a connected subtree that contains the root. Note that no physical relation is defined to $\mathbf{y}$, but this is not an issue since when such a relation is visited, Algorithm 1 simply skips it (line 4) as well as its ancestors. This algorithm only takes $O(N)$ time.

Then, we are left with two full joins, and invoke our rewriting rule proposed in Section 3.2 (line $6-8)$. As the reduce procedure takes $O(N)$ time, and the join phase takes $O(N+$ OUT) time implied by Lemma 3.5, we can obtain the complexity result in Theorem 3.1.
Improvement over Baseline. When $Q_{1}, Q_{2}$ fall into the class of "easy" DCQs as characterized by Theorem 3.1, our algorithm only takes $O\left(N+\right.$ OUT) time for computing $Q_{1}-Q_{2}$, while the baseline takes $O\left(N+\mathrm{OUT}_{1}+\operatorname{cost}\left(\mathrm{OUT}_{2}\right)\right)$ time, since $\operatorname{cost}\left(Q_{1}\right)=O(N+\mathrm{OUT})$ for free-connex $Q_{1}$. We next use a few examples of "easy" DCQs to illustrate the improvement achieved by our approach.

Example 3.9. Consider a DCQ with $Q_{1}=R_{1}\left(x_{1}, x_{2}, x_{3}\right)$ and $Q_{2}=R_{2}\left(x_{1}, x_{2}\right) \bowtie R_{3}\left(x_{2}, x_{3}\right) \bowtie$ $R_{4}\left(x_{1}, x_{3}\right)$. The baseline takes $O\left(N^{\frac{2 \cdot \omega}{\omega+1}}+N^{\frac{3(\omega-1)}{\omega+1}} \cdot \mathrm{OUT}_{2}^{\frac{3-\omega}{\omega+1}}\right)$ time to compute the triangle join $R_{2} \bowtie$ $R_{3} \bowtie R_{4}$ in $Q_{2}$, where $\omega$ is the exponent of fast matrix multiplication. In contrast, our approach only takes $O(N)$ time since OUT $\leq N$, improving the baseline by a factor of $O\left(N^{\frac{\omega-1}{\omega+1}}+N^{\frac{2 \omega-4}{\omega+1}} \cdot \mathrm{OUT}_{2}^{\frac{3-\omega}{\omega+1}}\right)$.

Example 3.10. Consider a DCQ with $Q_{1}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{3}, x_{4}\right)$ and $Q_{2}=R_{3}\left(x_{1}, x_{2}\right) \bowtie$ $R_{4}\left(x_{2}, x_{3}\right) \bowtie R_{5}\left(x_{1}, x_{3}\right)$. The baseline takes $O\left(N^{2}\right)$ time to materialize $Q_{1}$, which degenerates to the Cartesian product of $R_{1}$ and $R_{2}$. In contrast, our approach only requires $O(N+$ OUT) time, improving the baseline by a factor of $O\left(\frac{N^{2}}{\mathrm{OUT}}\right)$, since OUT can be much smaller than $N^{2}$.

Example 3.11. Consider a DCQ with $Q_{1}=\bowtie_{e \subseteq U:|e|=1} R_{e}\left(\left\{x_{1}\right\} \cup e\right)$ and $Q_{2}=\bowtie_{e^{\prime} \subseteq U:\left|e^{\prime}\right|=2} R_{e^{\prime}}\left(\left\{x_{1}\right\} \cup\right.$ $\left.e^{\prime}\right)$ for $U=\left\{x_{2}, \cdots, x_{k+1}\right\}$. The baseline takes $O(N)$ time to materialize $Q_{1}$, and $O\left(N^{\frac{k}{2}}\right)$ time to materialize $Q_{2}$. In contrast, our approach can compute it in $O(N+$ OUT) time, improving the baseline by a factor of $O\left(\frac{N^{k / 2}}{\text { OUT }}\right)$, since OUT can be much smaller than $N^{\frac{k}{2}}$.

## 4 HARD DCQS

In this section, we turn to the class of "hard" DCQs characterized by Theorem 2.4. We first prove the hardness of computing DCQs in linear time via some well-known conjectures, and then show an efficient heuristic for hard DCQs by further exploiting the query structures.

### 4.1 Hardness

We will prove the hardness of computing a hard DCQ $Q_{1}-Q_{2}$, in particular: (1) $Q_{1}$ is non-freeconnex; or (2) $Q_{1}$ is free-connex but $Q_{2}$ is non-linear-reducible; or (3) $Q_{1}$ is free-connex, $Q_{2}$ is linear-reducible, but there exists some $e \in \mathcal{E}_{2}^{\prime}$ such that $\left.\left(\mathbf{y}, \mathcal{E}_{1}^{\prime} \cup\{e\}\right\}\right)$ is cyclic, where $\left(\mathbf{y}, \mathcal{E}_{1}^{\prime}\right)$ and
$\left(\mathbf{y}, \mathcal{E}_{2}^{\prime}\right)$ are the reduced queries of $Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right), Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$ respectively. We will prove the hardness for each class of hard DCQs separately.
Hardness-(1). The hardness of computing DCQs in case (1) comes from computing a non-freeconnex CQ [16]. By setting the result of $Q_{2}$ as $\emptyset, Q_{1}-Q_{2}$ simply degenerates to $Q_{1}$, hence we obtain:

Lemma 4.1. For any $D C Q Q_{1}-Q_{2}$, if $Q_{1}$ is non-free-connex, any algorithm computing $Q_{1}-Q_{2}$ requires at least $\Omega(N+$ OUT) time.

The hardness of case (2) and (3) is built on the strong triangle conjecture in the literature:
Conjecture 4.2 (Strong Triangle conjecture [11]). Detecting whether an n-node m-edge graph contains a triangle requires $\Omega\left(\min \left\{n^{\omega-o(1)}, m^{2 \omega /(\omega+1)-o(1)}\right\}\right)$ time in expectation, where $\omega=2+o(1)$ is assumed as the exponent of fast matrix multiplication.

Hardness-(2). We start with two hardcore DCQs in Lemma 4.3 and Lemma 4.4. The proof of Lemma 4.5 for general DCQs in case (2) is given in the full version [10].

Lemma 4.3. Any algorithm for computing the following DCQ:

$$
Q_{1}-Q_{2}=R_{1}\left(x_{1}, x_{3}\right)-\pi_{x_{1}, x_{3}}\left(R_{2}\left(x_{1}, x_{2}\right) \bowtie R_{3}\left(x_{2}, x_{3}\right)\right)
$$

requires $\Omega(N)$ time, assuming the strong triangle conjecture.
Proof. For a graph $G=(V, E)$, we denote $m=|E|$ and $n=|V|$. Note that $n<m<n^{2}$; otherwise, we simply remove vertices that do not incident to any edges in $G$. We then construct an instance $D_{1}, D_{2}$ for $Q_{1}, Q_{2}$ by setting $R_{1}=R_{2}=R_{3}=E$. Hence, $N=m$. Note that there exists some triangle in $G$ if and only if $Q_{1} \cap Q_{2}$ is non-empty. Together with $Q_{1} \cap Q_{2}=Q_{1}-\left(Q_{1}-Q_{2}\right)$, we output "a triangle is detected in $G$ " if and only if $\left|Q_{1}-Q_{2}\right|<N$. If $Q_{1}-Q_{2}$ can be computed in $O(N)$ time, whether there exists a triangle in $G$ can be detected in $O\left(\min \left\{n^{2}, m^{4 / 3}\right\}\right)$ time, coming to a contradiction of strong triangle conjecture.

Lemma 4.4. Any algorithm for computing the following DCQ:

$$
Q_{1}-Q_{2}=R_{1}\left(x_{1}\right)-\pi_{x_{1}}\left(R_{2}\left(x_{1}, x_{3}\right) \bowtie R_{3}\left(x_{2}, x_{3}\right) \bowtie R_{4}\left(x_{1}, x_{3}\right)\right),
$$

requires $\Omega(N)$ time, assuming the strong triangle conjecture.
Proof. This is similar to the proof of Lemma 4.3. For a graph $G=(V, E)$, we construct $R_{2}=R_{3}=$ $R_{4}=E$ and $R_{1}=V$, with $m=|E|=N$ and $n=|V|$. Note that there exists some triangle in $G$ if and only if $Q_{1} \cap Q_{2}$ is non-empty. Together with $Q_{1} \cap Q_{2}=Q_{1}-\left(Q_{1}-Q_{2}\right)$, we output "a triangle is detected in $G$ " if and only if $\left|Q_{1}-Q_{2}\right|<\left|R_{1}\right|$. If $Q_{1}-Q_{2}$ can be computed in $O(N)$, whether there exists a triangle in $G$ can be detected in $O\left(\min \left\{n^{2}, m^{4 / 3}\right\}\right)$ time, coming to a contradiction of strong triangle conjecture.

Lemma 4.5. Given two $\operatorname{CQs} Q_{1}, Q_{2}$, if $Q_{1}$ is free-connex and $Q_{2}$ is non-linear-reducible, any algorithm computing $Q_{1}-Q_{2}$ requires $\Omega(N+\mathrm{OUT})$ time, assuming the strong triangle conjecture.

Hardness-(3). The hardness of evaluating a DCQ in case (3) inherits the hardness of deciding a DCQ: given a $\operatorname{DCQ} Q_{1}-Q_{1}$ and input databases $D_{1}, D_{2}$, the decidability problem asks to decide whether there exists a query result in $Q_{1}-Q_{2}$. We identify a few hardcore DCQs in Lemma 4.6. The proof of Lemma 4.7 for general DCQs in case (3) is given in the full version [10].

Lemma 4.6. Any algorithm for deciding the following DCQ

$$
\begin{aligned}
& Q_{1}-Q_{2}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)-R_{3}\left(x_{1}, x_{3}\right) \bowtie R_{4}\left(x_{2}\right) \\
& Q_{1}-Q_{2}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)-R_{3}\left(x_{1}, x_{3}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right) \\
& Q_{1}-Q_{2}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)-R_{3}\left(x_{1}, x_{3}\right) \bowtie R_{5}\left(x_{1}, x_{2}\right) \\
& Q_{1}-Q_{2}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)-R_{3}\left(x_{1}, x_{3}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right) \bowtie R_{5}\left(x_{1}, x_{2}\right)
\end{aligned}
$$

requires $\Omega(N)$ time, assuming the strong triangle conjecture.
Proof. We first focus on the first DCQ and the remaining ones can be proved similarly. Given an arbitrary graph $G=(V, E)$ with $m=|E|$ and $n=|V|$, we develop an algorithm to detect whether there exists a triangle in $G$. Note that $n<m<n^{2}$; otherwise, we simply remove vertices that do not incident to any edges in $G$. The degree $\operatorname{deg}(u)$ of a vertex $u \in V$ is defined as the size of neighbors of $u$, i.e., those incident to $u$ with an edge in $E$. We partition vertices in $V$ into two subsets: $V^{H}=\left\{v \in V: \operatorname{deg}(v)>m^{1 / 3}\right\}$ and $V^{L}=V-V^{H}$. From $G$, we construct following relations: $R=E$, $R_{0}=\left\{(u, v) \in E: u \in V^{L}\right.$ or $\left.v \in V^{L}\right\}, R_{1}=\left\{(u, v) \in E: u \in V^{H}\right\}, R_{2}=\left\{(u, v) \in E: v \in V^{H}\right\}$ and $R_{3}=V^{H} \times V^{H}-E$. Set $N=m^{4 / 3}$. It can be easily checked that each relation contains at most $m^{4 / 3}$ tuples, hence $N=m^{4 / 3}$. We further define a $\mathrm{CQ} Q$ as follows:

$$
Q=R\left(x_{1}, x_{2}\right) \bowtie R\left(x_{2}, x_{3}\right) \bowtie R_{0}\left(x_{1}, x_{3}\right)
$$

For $Q_{1}-Q_{2}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)-R_{3}\left(x_{1}, x_{3}\right) \bowtie R_{4}\left(x_{2}\right)$, we set $R_{4}=V$ and output "a triangle is detected" if and only if $Q$ or $Q_{1}-Q_{2}$ is not empty. We first prove the correctness of this algorithm, i.e., a triangle exists in $G$ if and only if $Q$ or $Q_{1}-Q_{2}$ is not empty. Direction Only-If. Consider an arbitrary triangle ( $u, v, w$ ) in $G$. We distinguish two cases: (i) at least one of $u, w$ is light; (ii) both $u$ and $w$ are heavy. In (i), assume $u$ is light. Then, $(u, v),(u, w) \in R_{0}$. We come to $(u, v, w) \in Q$. In (ii), $(u, v) \in R_{1},(v, w) \in R_{2},(u, w) \notin R_{3}$, so we come to $(u, v, w) \in Q_{1}-Q_{2}$. Direction If. If $Q \neq \emptyset$, say $(u, v, w) \in Q$, then $(u, v),(v, w),(u, w) \in E$ and therefore $(u, v, w)$ is a triangle in $G$. If $Q_{1}-Q_{2} \neq \emptyset$, say $(u, v, w) \in Q_{1}-Q_{2}$, then $(u, v) \in R_{1},(v, w) \in R_{2},(u, w) \notin R_{3}$, i.e., $(u, v),(v, w),(u, w) \in E$, and therefore $(u, v, w)$ is a triangle in $G$.

We next turn to the time complexity. All statistics and relations $R_{1}, R_{2}, R_{4}$ can be computed in $O(m)$ time. Moreover, relation $R_{3}$ can be constructed in $O\left(m^{4 / 3}\right)$ time since $\left|V^{H}\right|=O\left(m^{2 / 3}\right) . Q$ can be evaluated in $O\left(m^{4 / 3}\right)$ time, since each of $R_{0}\left(x_{1}, x_{2}\right) \bowtie R\left(x_{2}, x_{3}\right)$ generates at most $O\left(m^{4 / 3}\right)$ intermediate join results if $x_{2} \in V^{L}$, and each of $R_{0}\left(x_{1}, x_{2}\right) \bowtie R\left(x_{1}, x_{3}\right)$ generates at most $O\left(m^{4 / 3}\right)$ intermediate join results if $x_{1} \in V^{L}$. If $Q_{1}-Q_{2}$ can be decided in $O(N)$ time, whether there exists a triangle in $G$ can be decided in $O\left(m^{4 / 3}\right)$ time, coming to a contradiction to strong triangle conjecture.

For $Q_{1}-Q_{2}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)-R_{3}\left(x_{1}, x_{3}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right)$, we set $R_{4}=E$ and output "a triangle is detected" if and only if $Q$ or $Q_{1}-Q_{2}$ is not empty. For $Q_{1}-Q_{2}=R_{1}\left(x_{1}, x_{2}\right) \bowtie$ $R_{2}\left(x_{2}, x_{3}\right)-R_{3}\left(x_{1}, x_{3}\right) \bowtie R_{5}\left(x_{1}, x_{2}\right)$, we set $R_{5}=R$ and output "a triangle is detected" if and only if $Q$ or $Q_{1}-Q_{2}$ is not empty. For $Q_{1}-Q_{2}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)-R_{3}\left(x_{1}, x_{3}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right) \bowtie R_{5}\left(x_{1}, x_{2}\right)$, we set $R_{4}=R_{5}=E$ and output "a triangle is detected" if and only if $Q$ or $Q_{1}-Q_{2}$ is not empty. Similarly, $Q$ can be computed in $O\left(m^{4 / 3}\right)$ time. This way, if $Q_{1}-Q_{2}$ can be decided in $O(N)$ time, whether there exists a triangle in $G$ can be decided in $O\left(\mathrm{~m}^{4 / 3}\right)$ time, coming to a contradiction to strong triangle conjecture. Together, we have completed the proof.

Lemma 4.7. Given two $\operatorname{CQs} Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right), Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$, if $Q_{1}$ is free-connex, $Q_{2}$ is linearreducible, and there exists some $e \in \mathcal{E}_{2}^{\prime}$ such that $\left.\left(\mathbf{y}, \mathcal{E}_{1}^{\prime} \cup\{e\}\right\}\right)$ is cyclic where $\left(\mathbf{y}, \mathcal{E}_{1}^{\prime}\right)$ and $\left(\mathbf{y}, \mathcal{E}_{2}^{\prime}\right)$ are the reduced queries of $Q_{1}, Q_{2}$ respectively, any algorithm computing $Q_{1}-Q_{2}$ requires $\Omega(N+$ OUT $)$ time, assuming the strong triangle conjecture.

### 4.2 Efficient Heuristics

Although the hardness results in Section 4 have ruled out a linear-time algorithm for the "hard" DCQs, we find that it is still possible to explore efficient heuristics that can outperform the baseline approach. Our heuristic is based on a simple fact that $Q_{1}-Q_{2}=Q_{1}-Q_{1} \cap Q_{2}$. After computing the query results for $Q_{1}$, a straightforward way of deciding $Q_{1} \cap Q_{2}$ is to decide for each result $t \in Q_{1}$, whether $t \in Q_{2}$ nor not. This decidability query can be viewed as a special Boolean query by replacing every output attribute $x \in \mathbf{y}_{2}$ with a constant $\pi_{x} t$. More specifically, for $Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$, the derived a Boolean query can be represented as ( $\emptyset, \mathcal{V}_{2}-\mathbf{y},\left\{e-\mathbf{y}: e \in \mathcal{E}_{2}\right\}$ ). Putting everything together, we come to Theorem 4.8.

Theorem 4.8. Given two CQs $Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right), Q_{1}-Q_{2}$ can be computed in $O\left(\operatorname{cost}\left(Q_{1}\right)+\operatorname{OUT}_{1} \cdot \operatorname{cost}\left(Q_{2}^{\emptyset}\right)\right)$ time, where $Q_{2}^{\emptyset}=\left(\emptyset, \mathcal{V}_{2}-\mathbf{y},\left\{e-\mathbf{y}: e \in \mathcal{E}_{2}\right\}\right)$.

Remark. If $Q_{2}$ is linear-reducible, $Q_{2}$ can be reduced to a full join in $O(N)$ time by Algorithm 1. Then, $Q_{2}^{0}$ becomes empty. A faster solution is to build hashing indexes on every relation in the reduced $Q_{2}$. For each tuple $t \in Q_{1}$, it suffices to check for every $e \in \mathcal{E}_{2}$ whether $\pi_{e} t \in R_{e}^{\prime}$, which only takes $O(1)$ time. We note that the rewriting rule in Lemma 3.7 can also apply to this case and lead to the same complexity. Suppose $Q_{2}$ is reduced. Each $e \in \mathcal{E}_{2}$ induces a CQ $\left(\pi_{e} Q_{1}-R_{e}^{\prime}\right) \bowtie Q_{1}$. After materializing the results of $Q_{1}$, it suffices to check for each tuple $t \in Q_{1}$, whether $\pi_{e} t \in R_{e}^{\prime}$ or not. This is exactly how our heuristic proceeds. Hence, Corollary 2.5 follows.

Example 4.9. Consider a DCQ with $Q_{1}=\pi_{x_{1}, x_{2}, x_{3}} R_{1}\left(x_{1}, x_{4}\right) \bowtie R_{2}\left(x_{4}, x_{2}, x_{3}\right)$ and $Q_{2}=\pi_{x_{1}, x_{2}, x_{3}} R_{3}\left(x_{1}\right.$, $\left.x_{2}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right) \bowtie R_{5}\left(x_{1}, x_{3}\right) \bowtie R_{6}\left(x_{3}, x_{4}\right)$. The baseline spends $O\left(N^{\frac{2 \cdot \omega}{\omega+1}}+N^{\frac{\omega-1}{\omega+1}} \cdot\right.$ OUT $\left._{1}\right)$ time computing $Q_{1}$ and $O\left(N^{\frac{2 \omega \omega}{\omega+1}}+N^{\frac{3(\omega-1)}{\omega+1}}\right.$. $\left.\mathrm{OUT}_{2}^{\frac{3-\omega}{\omega+1}}\right)$ time computing the hidden triangle join $R_{3} \bowtie R_{4} \bowtie R_{5}$ in $Q_{2}$, where $\omega$ is the exponent of fast matrix multiplication. In contrast, our algorithm only spends $O\left(N^{\frac{2 . \omega}{\omega+1}}+N^{\frac{\omega-1}{\omega+1}} \cdot \mathrm{OUT}_{1}\right)$ time for computing $Q_{1}$, without computing the expensive $Q_{2}$, hence can improve the baseline by a factor of $O\left(N^{\frac{2(\omega-1)}{\omega+1}} \cdot \mathrm{OUT}_{2}^{\frac{3-\omega}{\omega+1}} / \mathrm{OUT}_{1}\right)$ when $N^{\frac{2(\omega-1)}{\omega+1}} \cdot \mathrm{OUT}_{2}^{\frac{3-\omega}{\omega+1}} \geq \mathrm{OUT}_{1}$.

We can show some further improvement when $Q_{2}$ is non-linear-reducible. Instead of issuing an individual Boolean query for every query result $t \in Q_{1}$, we take all the Boolean queries into account as whole. To do so, we further explore the structural property of the intersection query $Q_{2}^{\oplus}=\left(\mathbf{y}, \mathcal{V}_{2},\{\mathbf{y}\} \cup \mathcal{E}_{2}\right)$, by treating the query results of $Q_{1}$ as a single relation over attributes $\mathbf{y}$. We note that it is always cheaper (or at least not more expensive) to compute $Q_{2}^{\oplus}$ than $Q_{2}$, as one can always compute $Q_{2}^{\oplus}$ by materializing $Q_{2}$ with an additional semi-join with $Q_{1}$. The cost of this naive method evaluating $Q_{2}^{\oplus}$ is bounded by $O\left(\operatorname{cost}\left(Q_{1}\right)+\operatorname{cost}\left(Q_{2}\right)\right)$. It is possible to compute $Q_{2}^{\oplus}$ in a more efficient way.

Theorem 4.10. Given two $C Q s Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right), Q_{1}-Q_{2}$ can be computed in $O\left(\operatorname{cost}\left(Q_{1}\right)+\operatorname{cost}\left(Q_{2}^{\oplus}\right)\right)$ time, where $Q_{2}^{\oplus}=\left(\mathbf{y}, \mathcal{V}_{2},\{\mathbf{y}\} \cup \mathcal{E}_{2}\right)$.

Example 4.11. Consider a DCQ with $Q_{1}=R_{1}\left(x_{1}, x_{3}\right)$ and $Q_{2}=\pi_{x_{1}, x_{3}}\left(R_{3}\left(x_{1}, x_{2}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right)\right)$. The baseline takes $O\left(N+N \cdot \sqrt{\mathrm{OUT}_{2}}\right)$ time to materialize $Q_{2}$. The first heuristics of issuing $Q_{2}^{\emptyset}$ for each tuple $t \in R_{1}$ takes $O\left(N^{3 / 2}\right)$ time. We note that $Q_{12}=\pi_{x_{1}, x_{3}}\left(R_{1}\left(x_{1}, x_{3}\right) \bowtie R_{2}\left(x_{1}, x_{2}\right) \bowtie R_{3}\left(x_{2}, x_{3}\right)\right)$ lists edges that participate in at least one triangle. The existing best algorithm takes $O\left(N^{\frac{2 \omega}{\omega+1}}\right)$ time to compute $Q_{12}$, where $\omega$ is the exponent of fast matrix multiplication, dominating the overall complexity. Our approach will improve the baseline if $\mathrm{OUT}_{2}>N^{\frac{2(\omega-1)}{\omega+1}}$, and strictly outperforms the naive heuristic.

Example 4.12. Consider a DCQ with $Q_{1}=\pi_{x_{1}, x_{3}} R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)$ and $Q_{2}=\pi_{x_{1}, x_{3}} R_{3}\left(x_{1}, x_{2}\right) \bowtie$ $R_{4}\left(x_{2}, x_{3}\right)$. Let $R_{5}\left(x_{1}, x_{3}\right)=Q_{1}$. Here, $Q_{12}=\pi_{x_{1}, x_{3}} R_{3}\left(x_{1}, x_{2}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right) \bowtie R_{5}\left(x_{1}, x_{3}\right)$ with $\left|R_{5}\right|=$ $\mathrm{OUT}_{1}$. Similarly, the existing best algorithm takes $O\left(\mathrm{OUT}_{1}^{\frac{\omega}{\omega+1}} \cdot N^{\frac{\omega}{\omega+1}}\right)$ time to compute $Q_{12}$. It is
worth mentioning that $Q_{12} \neq \pi_{x_{1}, x_{3}} R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right) \bowtie R_{3}\left(x_{1}, x_{2}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right)$. Suppose $(a, c) \in Q_{12}$, is witnessed by $\left(a, b_{1}, c\right) \inf R_{1} \bowtie R_{2}$ and $\left(a, b_{2}, c\right) \in R_{3} \bowtie R_{4}$, but $\left(a, b_{1}, c\right) \notin R_{3} \bowtie R_{4}$ and $\left(a, b_{2}, c\right) \notin R_{1} \bowtie R_{2}$. Then, $\left(a, b_{1}, c\right),\left(a, b_{2}, c\right) \notin R_{1} \bowtie R_{2} \bowtie R_{3} \bowtie R_{4}$, hence the result $(a, c)$ will be missed in this rewriting.

## 5 EXTENSIONS

Based on the basic DCQ over two CQs discussed so far, we next consider several interesting extensions of DCQ with rich interaction between difference and other relational algebra operators.

### 5.1 Difference of Multiple CQs

The first extension is adapting our result for computing DCQ involving two CQs to multiple CQs, say $Q=Q_{1}-Q_{2}-\cdots-Q_{k}$. Given $k \operatorname{CQs} Q_{1}, Q_{2}, \cdots, Q_{k}$ with $Q_{i}=\left(\mathbf{y}, \mathcal{V}_{i}, \mathcal{E}_{i}\right)$ for $i \in\{2,3, \cdots, k\}$, we suppose $Q_{1}$ is free-connex, $Q_{i}$ for $i \geq 2$ is linear-reducible for every $i \in\{2,3, \cdots, k\}$, and the subquery induced by $\mathcal{E}_{1}^{\prime} \cup\{e\}$ is $\alpha$-acyclic for every $e \subseteq \mathcal{E}_{2}^{\prime} \cup \mathcal{E}_{3}^{\prime} \cup \cdots \cup \mathcal{E}_{k}^{\prime}$, where ( $\mathbf{y}, \mathcal{E}_{i}^{\prime}$ ) is the reduced query of $\left(\mathbf{y}, \mathcal{V}_{i}, \mathcal{E}_{i}\right)$. We can invoke our algorithm in Section 3 in a round-robin fashion, by computing $Q_{1}-Q_{i}$ for each $i \in\{2,3, \cdot, 4\}$ and stop once some $Q_{1}-Q_{i}$ finishes the computation. Suppose $Q_{1}-Q_{i}$ finishes the first. We then check for each result $t \in Q_{1}-Q_{i}$ whether $t \in Q_{j}$ for every $j \in\{2,3, \cdots, k\}-\{i\}$, and output $t \in Q_{1}-Q_{i}$ if $t \notin Q_{j}$ for every $j \in\{2,3, \cdots, k\}-\{i\}$.

Theorem 5.1. Given $k$ CQs $Q_{1}, Q_{2}, \cdots, Q_{k}$ with $Q_{i}=\left(\mathbf{y}, \mathcal{V}_{i}, \mathcal{E}_{i}\right)$ for $i \in\{2,3, \cdots, k\}$, the $D C Q$ $Q=Q_{1}-Q_{2}-\cdots-Q_{k}$ can be evaluated in $O\left(N+\min _{i \in\{2,3, \cdots, k\}} \operatorname{OUT}_{i}\right)$ time where $\mathrm{OUT}_{i}=\left|Q_{1}-Q_{i}\right|$, if $Q_{1}$ is free-connex, $Q_{i}$ for $i \geq 2$ is linear-reducible for every $i \in\{2,3, \cdots, k\}$, and the subquery induced by $\mathcal{E}_{1}^{\prime} \cup\{e\}$ is $\alpha$-acyclic for everye $\subseteq \mathcal{E}_{2}^{\prime} \cup \mathcal{E}_{3}^{\prime} \cup \cdots \cup \mathcal{E}_{k}^{\prime}$, where $\left(\mathbf{y}, \mathcal{E}_{i}^{\prime}\right)$ is the reduced query of $\left(\mathbf{y}, \mathcal{V}_{i}, \mathcal{E}_{i}\right)$.

We also mention an important property that can decompose a DCQ into the union of multiple DCQs in a special form:

Lemma 5.2. Given a $D C Q Q=Q_{1}-Q_{2}-\cdots-Q_{k}$ with full $C Q Q_{i}=\left(\mathbf{y}, \mathcal{E}_{i}\right)$ for everyi $\in\{1,2, \cdots, k\}$, then $Q=\bigcup_{\left(e_{2}, e_{3}, \cdots, e_{k}\right) \in \mathcal{E}_{2} \times \mathcal{E}_{3} \times \cdots \times \mathcal{E}_{k}}\left(Q_{1}-R_{e_{2}} \bowtie Q_{1}-R_{e_{3}} \bowtie Q_{1}-\cdots-R_{e_{k}} \bowtie Q_{1}\right)$.

Proof. Direction $\subseteq$. Consider an arbitrary tuple $t \in Q_{1}-Q_{2}-\cdots-Q_{k}$. We know $t \in Q_{1}$ and $t \notin Q_{i}$ for every $i \notin\{2,3, \cdots, k\}$. If $t \notin Q_{i}$, there must exist some $e_{i} \in \mathcal{E}_{i}$, such that $\pi_{e_{i}} t \notin R_{e_{i}}$, i.e., $t \notin R_{e_{i}} \bowtie Q_{1}$. Together, there must exist a combination $\left(e_{2}, e_{3}, \cdots, e_{k}\right) \in \mathcal{E}_{2} \times \mathcal{E}_{3} \times \cdots \times \mathcal{E}_{k}$ such that $Q_{1}-R_{e_{2}} \bowtie Q_{1}-R_{e_{3}} \bowtie Q_{1}-\cdots-R_{e_{k}} \bowtie Q_{1}$. Direction ొ. Consider an arbitrary combination $\left(e_{2}, e_{3}, \cdots, e_{k}\right) \in \mathcal{E}_{2} \times \mathcal{E}_{3} \times \cdots \times \mathcal{E}_{k}$ and an arbitrary tuple $t \in Q_{1}-R_{e_{2}} \bowtie Q_{1}-R_{e_{3}} \bowtie$ $Q_{1}-\cdots-R_{e_{k}} \bowtie Q_{1}$. By definition, $t \in Q_{1}$ and $\pi_{e_{i}} t \notin R_{e_{i}}$ for every $i \in\{2,3, \cdots, k\}$, i.e., $t \notin Q_{i}$. Hence, $t \in Q_{1}-Q_{2}-\cdots-Q_{k}$.

Then, it suffices to compute $Q_{1}-R_{e_{2}} \bowtie Q_{1}-R_{e_{3}} \bowtie Q_{1}-\cdots-R_{e_{k}} \bowtie Q_{1}$ for each combination $\left(e_{2}, e_{3}, \cdots, e_{k}\right) \in \mathcal{E}_{2} \times \mathcal{E}_{3} \times \cdots \times \mathcal{E}_{k}$. We notice that deciding whether such a query has any result can be done in $O(N)$ time, if any subquery induced by $\mathcal{E}_{1} \cup S$ is $\alpha$-acyclic for every $S \subseteq$ $\left\{e_{2}, e_{3}, \cdots, e_{k}\right\}[20]$ (see Theorem 7.5). We believe that it is promising to adapt their algorithm to further speed the evaluation of DCQ with multiple CQs, and we won't pursue this direction here.

### 5.2 Select, Project and Join

- If there is a selection operator $\sigma_{\phi}$ over $Q=Q_{1}-Q_{2}$, we can push it down such that $Q=$ $\sigma_{\phi} Q_{1}-\sigma_{\phi} Q_{2}$. If $\phi$ is a predicate on a base relation $R_{e}$ of $Q_{1}$ (resp. $Q_{2}$ ), we can we simply check if $\phi(t)$ is true for each tuple $t \in \sigma_{\phi} R_{e}$, and discard it if not. This only takes $O(N)$ time. It is challenging that $\phi$ is a predicate not on any base relation, even for a single CQ evaluation.
- If there is a projection operator $\pi_{\theta}$ over $Q=Q_{1}-Q_{2}$, we can push it down such that $Q=$ $\pi_{\theta} Q_{1}-\pi_{\theta} Q_{2}$, and handle a new $\operatorname{DCQ} Q_{1}^{\prime}-Q_{2}^{\prime}$ with $Q_{1}^{\prime}=\pi_{\theta} Q_{1}$ and $Q_{2}^{\prime}=\pi_{\theta} Q_{2}$.
- If there is a join operator over multiple DCQs, we first rewrite the join into a DCQover multiple CQs and invoke our previous algorithm in Section 5.1. More specifically, given $k$ DCQs $Q^{1}, Q^{2}, \cdots, Q^{k}$ with $Q^{i}=Q_{1}^{i}-Q_{2}^{i}$ for any $i \in[k]$, we can rewrite $Q^{1} \bowtie Q^{2} \bowtie \cdots \bowtie Q^{k}$ as a DCQ with multiple CQs: $\left(\bowtie_{i \in[k]} Q_{1}^{i}\right)-\left\{\left(\bowtie_{i \in I} Q_{1}^{i}\right) \bowtie\left(\bowtie_{j \in J} Q_{2}^{j}\right): I \subsetneq[k], J=[k]-I\right\}$.


### 5.3 Aggregation

Our algorithm for DCQ can also be extended to support aggregations over annotated relations [14, 30]. Let $(S, \oplus, \otimes)$ be a commutative ring. For a CQ $Q$ over an annotated instance $D$, every tuple $t \in R_{e}$ has an annotation $w(t) \in S$. For a full query $Q=(\mathcal{V}, \mathcal{E})$, the annotation for any join result $t \in Q(D)$ is defined as $w(t):=\underset{e \in \mathcal{E}}{\otimes} w\left(\pi_{e} t\right)$. For a non-full query $Q=(\mathbf{y}, \mathcal{V}, \mathcal{E})$, the aggregation becomes GROUP BY $\mathbf{y}$, and the annotation for each result $t \in Q$ (i.e., the aggregate of each group) is $w(t):=\underset{t^{\prime} \in \bowtie_{e \in \mathcal{E}} R_{e}: \pi_{y} t^{\prime}=t}{\oplus} w\left(t^{\prime}\right)$. Below, we introduce two commonly-used formulations. Given $Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right), Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$ and instances $D_{1}, D_{2}$, let $w_{1}, w_{2}$ be the annotations of tuples in $Q_{1}, Q_{2}$ respectively. For completeness, we set $w_{1}(t)=0$ if $t \notin Q_{1}\left(D_{1}\right)$ and $w_{2}(t)=0$ if $t \notin Q_{2}\left(D_{2}\right)$.
Relational difference. For DCQs defined on relational difference, a tuple $t$ appears in the query results of $Q_{1}-Q_{2}$ if and only if $t \in Q_{1}\left(D_{1}\right)$ and $t \notin Q_{2}\left(D_{2}\right)$. For $t \in \pi_{\mathbf{y}^{\prime}}\left(Q_{1}\left(D_{1}\right)-Q_{2}\left(D_{2}\right)\right)$, the annotation of $t$ is defined as $w(t)=\oplus_{t^{\prime} \in Q_{1}\left(D_{1}\right)-Q_{2}\left(D_{2}\right): \pi_{y^{\prime}} t^{\prime}=t^{2}\left(t^{\prime}\right) \text {. The input size is defined as }}$ $N=\left|D_{1}\right|+\left|D_{2}\right|$, and the output size is OUT $=\left|\pi_{\mathbf{y}^{\prime}}\left(Q_{1}\left(D_{1}\right)-Q_{2}\left(D_{2}\right)\right)\right|$. Again, our target is to find a linear-time algorithm in terms of $N$ and OUT. Our algorithms can be applied directly, followed by aggregation, and its complexity is bottlenecked by the output size of the difference query, i.e., $\left|Q_{1}\left(D_{1}\right)-Q_{2}\left(D_{2}\right)\right|$, which could be much larger than OUT.
Numerical difference. For DCQs defined on numerical difference, a tuple $t$ appears in the query results of $Q_{1}-Q_{2}$ if and only if $t \in Q_{1}\left(D_{1}\right)$ or $t \in Q_{2}\left(D_{2}\right)$, with annotation $w(t)=w_{1}(t)-w_{2}(t)$. Then, the aggregation operator defined over attributes $\mathbf{y}^{\prime}$ on top of $Q_{1}-Q_{2}$ can be rewritten as the numerical difference of two new annotated queries, i.e., $\pi_{y^{\prime}} Q_{1}-\pi_{y^{\prime}} Q_{2}$. The input size is defined as $N=\left|D_{1}\right|+\left|D_{2}\right|$, and the output size is OUT $=\left|\pi_{y^{\prime}} Q_{1}\left(D_{1}\right) \cup \pi_{y^{\prime}} Q_{2}\left(D_{2}\right)\right|$. Again, our target is to find an linear-time algorithm in terms of $N$ and OUT. Here, any algorithm with time complexity $O\left(N+\left|\pi_{y^{\prime}} Q_{1}\left(D_{1}\right)\right|+\left|\pi_{y^{\prime}} Q_{2}\left(D_{2}\right)\right|\right)$ is already optimal, since $\left|\pi_{\mathrm{y}^{\prime}} Q_{1}\left(D_{1}\right) \cup \pi_{\mathrm{y}^{\prime}} Q_{2}\left(D_{2}\right)\right| \geq$ $\frac{1}{2}\left(\left|\pi_{\mathrm{y}^{\prime}} Q_{1}\left(D_{1}\right)\right|+\left|\pi_{\mathrm{y}^{\prime}} Q_{2}\left(D_{2}\right)\right|\right)$. Hence, if $\pi_{\mathrm{y}^{\prime}} Q_{1}$ and $\pi_{\mathrm{y}^{\prime}} Q_{2}$ are free-connex, both our algorithm and baseline are optimal.

Theorem 5.3. Given two $C Q s Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$, and a subset of aggregation attributes $\mathbf{y}^{\prime} \subseteq \mathbf{y}$, if $\left(\mathbf{y}^{\prime}, \mathcal{V}_{1}, \mathcal{E}_{1}\right)$ and $\left(\mathbf{y}^{\prime}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$ are free-connex, $\pi_{\mathbf{y}^{\prime}}\left(Q_{1}-Q_{2}\right)$ with numerical difference can be computed in $O(N+\mathrm{OUT})$ time.

Example 5.4. Consider an example $\operatorname{DCQQ}=\pi_{x_{1}}\left(R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)-R_{3}\left(x_{1}, x_{2}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right)\right)$ over an instance in Figure 3. This DCQ can capture Q16 in the TPC-H benchmark [9] as a special case. For relational difference, the query result of $Q$ includes 2 tuples as $\left\{\left(a_{1}, 1\right),\left(a_{2}, 1\right)\right\}$. For numerical difference, the query result of $Q$ includes 3 tuples as $\left\{\left(a_{1}, 1\right),\left(a_{2}, 2\right),\left(a_{3},-2\right)\right\}$.

### 5.4 Bag Semantics

We consider the bag semantics that the set of query result is a multi-set. For simplicity, each distinct tuple $t$ is annotated with a positive integer $w(t)$ to indicate the number of copies. In a full CQ $Q=(\mathcal{V}, \mathcal{E})$, the annotation of $t \in Q$ is defined as $w(t)=\times_{e \in \mathcal{E}} w\left(\pi_{e} t\right)$. For a projection of $R_{e}$ onto attributes $e^{\prime}$, the annotation of $t \in \pi_{e^{\prime}} R_{e}$ is defined as $w(t)=\sum_{t^{\prime} \in R_{e}: \pi_{e^{\prime}} t^{\prime}=t} w\left(t^{\prime}\right)$. Given two CQs


Fig. 3. An example of DCQ $Q_{1}-Q_{2}$ with $Q_{1}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)$ and $Q_{2}=R_{3}\left(x_{1}, x_{2}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right)$. In the bag semantics, the input size is 18 , and the output size is 7. ( $a_{1}, b_{1}, c_{1}$ ) is a result under bag semantics, but not a result under set semantics. The results of $Q_{1}-Q_{2}$ computed by the old and new rewriting rules are also illustrated separately.
$Q_{1}=\left(\mathbf{y}, \mathcal{V}_{1}, \mathcal{E}_{1}\right), Q_{2}=\left(\mathbf{y}, \mathcal{V}_{2}, \mathcal{E}_{2}\right)$ and two input instances $D_{1}, D_{2}$, let $w_{1}, w_{2}$ be the annotations of tuples in $Q_{1}, Q_{2}$ respectively. For completeness, we set $w_{1}(t)=0$ if $t \notin Q_{1}\left(D_{1}\right)$ and $w_{2}(t)=0$ if $t \notin Q_{2}\left(D_{2}\right)$. A tuple $t$ is a query result of $Q_{1}-Q_{2}$ if and only if $t \in Q_{1}\left(D_{1}\right)$ and $w_{1}(t)>w_{2}(t)$. An example is given in Figure 3.

The input size is $N=\left|D_{1}\right|+\left|D_{2}\right|$, and the output size is OUT $=\sum_{t \in Q_{1}\left(D_{1}\right)} \max \left\{0, w_{1}(t)-w_{2}(t)\right\}$. Again, our target is to find an linear-time algorithm in terms of $N$ and OUT. Unfortunately, our rewriting rule in Section 3 cannot be adapted here. Figure 3 shows several incorrect behaviors: some tuple has a much higher annotation (e.g., $\left(a_{1}, b_{1}, c_{1}\right)$ ); and some tuple should not appear (e.g., $\left(a_{2}, b_{2}, c_{2}\right)$ ), which motivates us to explore new rules here.

Example 5.5. Consider a DCQ with $Q_{1}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)$ and $Q_{2}=R_{3}\left(x_{1}, x_{2}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right)$ under the bag semantics. Any result $(a, b, c) \in Q_{1}-Q_{2}$ falls into one of the three cases: $(1)(a, b) \notin R_{2}$ or $(b, c) \notin R_{4}$; (2) $w_{1}(a, b)>w_{2}(a, b)$ and $w_{1}(b, c)>w_{2}(b, c)$; (3) either $w_{1}(b, c) \leq w_{2}(b, c)$ or $w_{1}(a, b) \leq w_{2}(a, b)$, but $w_{1}(a, b) \cdot w_{1}(b, c)>w_{2}(a, b) \cdot w_{2}(b, c)$. We partition $R_{1}$ into three subsets, $R_{1 \emptyset}=\left\{t \in R_{1}: t \notin R_{3}\right\}, R_{1<}=\left\{t \in R_{1}: t \in R_{3}, w_{1}(t) \leq w_{2}(t)\right\}$ and $R_{1>}=\left\{t \in R_{1}: t \in R_{3}: w_{1}(t)>\right.$ $\left.w_{2}(t)\right\}$. Similarly, we partition $R_{2}$ into $R_{2 \emptyset}, R_{2<}, R_{2>}$ with respect to $R_{4}$. Results falling into (1) can be found by $\left(R_{1 \emptyset} \bowtie R_{2}\right) \cup\left(R_{1} \bowtie R_{2 \emptyset}\right)$. Results falling into (2) can be found by $R_{1>} \bowtie R_{2>}$. Results falling into (3) can be found by two new $\theta$-joins $\left(R_{1>} \bowtie_{\theta} R_{2<}\right) \cup\left(R_{1<} \bowtie_{\theta} R_{2>}\right)$, where a pair of tuples $\left(t_{1}, t_{2}\right)$ can be $\theta$-joined if and only if $w_{1}\left(t_{1}\right) \cdot w_{1}\left(t_{2}\right)>w_{2}\left(t_{1}\right) \cdot w_{2}\left(t_{2}\right)$.

All auxiliary relations as well as $\left(R_{1 \emptyset} \bowtie R_{2}\right) \cup\left(R_{1} \bowtie R_{2 \emptyset}\right)$ and $R_{>} \bowtie R_{2>}$ can be computed efficiently. We consider $R_{1>} \bowtie_{\theta} R_{2<}\left(\left(R_{1<} \bowtie_{\theta} R_{2>}\right)\right.$ is symmetric). The solution of checking $\theta$-condition for all combinations of tuples in $R_{1>}$ and $R_{2<}$ incurs quadratic complexity. A smarter way is to sort $R_{1>}$ and $R_{2<}$ by $B$ first, and then by the ratio of $\frac{w_{1}(\cdot)}{w_{2}(\cdot)}$ decreasingly. Then, we start with $(a, b) \in R_{1>}$ with maximum ratio, and linearly scan tuples in $R_{2<}$ with the join value $b$ until we meet some tuple $(b, c)$ such that $\frac{w_{1}(b, c)}{w_{2}(b, c)} \leq \frac{w_{2}(a, b)}{w_{1}(a, b)}$. We then stop and proceed with the next tuple in $R_{1>}$. If no join result is produced by $(a, b)$, we skip the subsequent tuples with the same join value $b$ and continue. Overall, this algorithm takes $O(N \log N+\mathrm{OUT})$ time.

Theorem 5.6. Given two $\operatorname{CQs} Q_{1}=Q_{2}=(\mathbf{y}, \mathcal{V}, \mathcal{E})$, if $(\mathbf{y}, \mathcal{V}, \mathcal{E})$ is free-connex, then $Q_{1}-Q_{2}$ under the bag semantics can be computed in $O(N \log N+$ OUT) time.

Our observation above can be extended to the case when both $Q_{1}, Q_{2}$ correspond to the same free-connex query. The proof of Theorem 5.6 is given in the full version [10]. The case when $Q_{1}, Q_{2}$ have different schemas is left as future work.

## 6 EXPERIMENTS

### 6.1 Experimental Setup

Prototype implementation. Our newly developed algorithms can be easily integrated into any SQL engine by rewriting the original SQL query. It can be further optimized if we directly integrate the rewrite procedure into the SQL parser and have customized index support. Our ultimate goal is to implement our algorithms into a system prototype with three components: a SQL parser, a query optimizer, and new indices. At the current stage, we choose to manually rewrite all SQL queries and demonstrate the power of our optimizations via the comparison with vanilla SQL queries .
Query processing engines compared. To compare the performance of all optimized techniques we proposed in the paper, we choose PostgreSQL [4], DuckDB[1], SQLite[7], MySQL[2] running in centralized settings, and Spark SQL [6] running in parallel/distributed settings, as the query processing engines. All of them are widely used in academia and industry. In the experiments, we observed that SQLite and MySQL show very poor performance, with most of the test points being timed out. Hence, we built full indices on these systems to expedite the execution. Moreover, DuckDB is a columnar-vectorized query execution engine, and indices are built when importing input data. During the experiments, we test the single-thread performance of our new optimization techniques over PostgreSQL, DuckDB, SQLite and MySQL, and parallel performance over Spark SQL. In order to separate the I/O cost from the total execution time, we load all data into the memory in advance by using pg-prewarm in PostgreSQL and cache in Spark SQL. For DuckDB and SQLite, the data need to be loaded into memory before execution, so we only count the query execution time.
Experimental environment. We perform all experiments in two machines. For experiments conducted on PostgreSQL and MySQL, we use a machine equipped with two Intel Xeon 2.1 GHz processors, each having 12 cores/ 24 threads and 416 GB memory. For all experiments on Spark SQL, DuckDB and SQLite, we use a machine equipped with two Xeon 2.0 GHz processors, each having 28 cores / 56 threads and 1TB of memory. All machines run Linux, with Scala 2.13.9 and JVM 1.8.0. We use Spark 3.3.0 and PostgreSQL 16.0. We assign 8 cores for Spark and 1 core for the rest platforms during the experiments. Each query is evaluated 10 times with each engine, and we report the average running time. Each query runs at most 10 hours to obtain meaningful results.

### 6.2 Datasets and Queries

The experiments consist of graph queries and benchmark queries.
Benchmark queries. For relational queries, we adopt two standard benchmarks (TPC-DS [8] and TPC-H [9]) in industry and select 3 queries with difference operator (TPC-H Q16, TPC-DS Q35, and TPC-DS Q69). These three benchmark queries connect DCQ with other relational operators like selection, projection, join, and aggregation. All benchmark queries can be captured by a common schema $Q=R_{1}\left(x_{1}, x_{2}\right) \bowtie\left(\pi_{x_{2}} R\left(x_{1}, x_{2}\right)-\pi_{x_{2}} R_{2}\left(x_{2}, x_{3}\right) \bowtie R_{3}\left(x_{3}, x_{4}\right)\right)$ and the joins are all primary-key foreign-key joins.

Graph queries. For graph pattern queries, we use real-world graphs (such as BitCoin, DBLP, Eponions, Google, and Wiki) from SNAP (Stanford Network Analysis Project) [5], summarized in Table 2. We store edge information as a relation Graph(src, dst) and manually create a triple

| Graph | \#edge | \#vertex | \#l2 path | \#triangle | $\#_{G 1}$ | $\# Q_{G 2}$ | $\# Q_{G 3}$ | $\# Q_{G 4}$ | $\# Q_{G 5}$ | $\# Q_{G 6}$ |
| :---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| Bitcoin | 24,186 | 3,783 | $1,256,332$ | 88,753 | 820 | $1.0 \times 10^{7}$ | 585,958 | 331,497 | $3.8 \times 10^{7}$ | $5.7 \times 10^{8}$ |
| Epinions | 508,837 | 75,879 | $3.9 \times 10^{7}$ | $3,586,405$ | 25,947 | $9.3 \times 10^{8}$ | $1.8 \times 10^{7}$ | $1.0 \times 10^{7}$ | $3.5 \times 10^{9}$ | $2.5 \times 10^{11}$ |
| DBLP | $1,049,866$ | 317,080 | $7,064,738$ | $2,224,385$ | 466,646 | $1.6 \times 10^{8}$ | $3,532,369$ | $2,203,597$ | $6.7 \times 10^{7}$ | - |
| Google | $5,105,039$ | 875,713 | $6.0 \times 10^{7}$ | $2.8 \times 10^{7}$ | 372,042 | $2.1 \times 10^{8}$ | $2.4 \times 10^{7}$ | $1.5 \times 10^{7}$ | $7.8 \times 10^{8}$ | - |
| Wiki | $2.8 \times 10^{7}$ | $2,394,385$ | $2.6 \times 10^{9}$ | $8.1 \times 10^{7}$ | 0 | $1.1 \times 10^{10}$ | $1.3 \times 10^{8}$ | $6.6 \times 10^{7}$ | - | - |

Table 2. Graph datasets and their statistics. \#edge is the input size of graph datasets. $\# Q_{G i}$ is the output size of $Q_{G i}$ over the corresponding graph datasets. '-' indicates that the output size is too huge such that all systems cannot report the output size within the time limit.

```
\(Q_{G 1}=\operatorname{Graph}\left(\operatorname{node}_{1}, \operatorname{node}_{2}\right)-\pi_{\text {node }_{1}, \text { node }_{2}}\left(\operatorname{Graph}\left(\right.\right.\) node \(\left._{1}, \operatorname{node}_{2}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{2}\), node \(\left.\left._{3}\right)\right)\)
\(Q_{G 2}=\operatorname{Graph}\left(\right.\) node \(_{1}\), node \(\left._{2}\right) \bowtie\) Triple \(\left(\right.\) node \(_{2}\), node \(_{3}\), node \(\left._{4}\right)\)
    - Triple( node \(_{1}\), node \(_{2}\), node \(\left._{3}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{3}\), node \(\left._{4}\right)\)
\(Q_{G 3}=\) Triple \(\left(\right.\) node \(_{1}\), node \(_{2}\), node \(\left._{3}\right)-\operatorname{Graph}\left(\right.\) node \(_{1}\), node \(\left._{2}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{2}\), node \(\left._{3}\right) \bowtie G r a p h\left(\right.\) node \(_{3}\), node \(\left._{1}\right)\)
\(Q_{G 4}=\operatorname{Triple}\left(\right.\) node \(_{1}\), node \(_{2}\), node \(\left._{3}\right)\)
    \(-\pi_{\text {node }_{1}, \text { node }_{2}, \text { node }_{3}} \operatorname{Graph}\left(\right.\) node \(_{1}\), node \(\left._{2}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{2}\), node \(\left._{3}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{3}\), node \(\left._{4}\right)\)
\(Q_{G 5}=\operatorname{Graph}\left(\right.\) node \(_{1}\), node \(\left._{2}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{2}\), node \(\left._{3}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{3}\), node \(\left._{4}\right)\)
    \(-\operatorname{Graph}\left(\right.\) node \(_{2}\), node \(\left._{3}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{3}\), node \(\left._{4}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{4}\), node \(\left._{1}\right)\)
\(Q_{G 6}=\operatorname{Graph}\left(\right.\) node \(_{1}\), node \(\left._{2}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{3}\), node \(\left._{4}\right)\)
    \(-\operatorname{Graph}\left(\right.\) node \(_{1}\), node \(\left._{2}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{2}\), node \(\left._{3}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{3}\), node \(\left._{1}\right) \bowtie \operatorname{Graph}\left(\right.\) node \(_{3}\), node \(\left._{4}\right)\)
```

Fig. 4. Graph queries.
relation Triple(node1, node2, node3) from the graph. Tuples in Triple are generated by following rules: (rule 1) a random length-2 path in the graph as (node1, node2, node3); or (rule 2) a random edge in the graph as (node1, node2), together with a random vertex in the graph as node3; or (rule 3) a triple (node1, node3, node5) from a random length-4 path (node1, node2, node3, node4, node5) in the graph. Triple may involve different portions of tuples generated by three rules in different queries. For a graph with $n$ length 2 paths, we set the size of Triple to be $0.05 n$ for Wiki (since it is too large to process as shown in Table 2), and $0.5 n$ for the remaining graphs. We evaluate 6 graph queries as described in Figure 5, whose original SQL queries as well as optimized SQL queries after rewriting are given in the full version [10].

More specifically, $Q_{G 1}$ finds all edges in the graph that do not participate in any length-2 path. $Q_{G 2}$ finds all length-3 paths that the third node (node ${ }_{3}$ ) is not sampled together with the edge ( node $_{1}$, node $_{2}$ ). $Q_{G 3}$ finds length-2 paths that do not form a triangle. $Q_{G 4}$ finds all generated triples that cannot extend to a length-4 path. $Q_{G 5}$ finds all length-4 paths that do not form a length-4 cycle. $Q_{G 6}$ finds all pairs of edges in the graph, which do not form a length-4 cycle.

### 6.3 Experiment Results

Running time. Figure 5 shows the running time of different engines on graph queries. The input and output size of all graphs queries are given in Table 2. All bars reaching the axis boundary indicate that the system did not finish within the 8-hour limit, or ran out of memory. As $Q_{G 6}$ contains an expensive Cartesian product as sub-query, materializing its query result exceeds the memory capacity of our machines on most datasets. PostgreSQL can only evaluate the original SQL query of $Q_{G 6}$ on Bitcoin dataset. By adding the parallelism from 8 to 80 , our optimized Spark SQL can evaluate $Q_{G 6}$ on Epinions dataset within the time limit, while the vanilla Spark SQL cannot complete the evaluation. For $Q_{G 5}$, all systems cannot finish the evaluation on Wiki dataset due


Fig. 5. Running time of graph and benchmark queries.
to the large intermediate results created. We also observe that both SQLite and MySQL cannot finish all test points for $Q_{G 2}$ and $Q_{G 6}$, and most test points over Wiki dataset. It could be the reason that both systems are not designed for analytical queries. Our optimization techniques already achieve a speedup ranging from $2 x$ to $1760 x$ on PostgreSQL, from 1.2x to 270x on Spark SQL, from 2 x to 1848x on DuckDB, from 1.25x to 1095x on SQLite, and from 1.8x to 5.1x on MySQL for graph queries, even without considering the queries that could not finish within the time limit. We also observe an unusual test point for $Q_{G 3}$ in MySQL, that our optimized SQL query takes more time than the vanilla SQL query, which may be due to some unknown deficiencies in MySQL internals. ${ }^{2}$

Figure 5 also shows the running time of all query engines on benchmark queries under different scale factors (i.e., parameters used to generate benchmark dataset, which is roughly proportional to the input data size). DuckDB and MySQL fail to finish some test points with scale factor 100 . However, the improvement in benchmark queries achieved by our optimized techniques is minor, as expected. More specifically, the vanilla benchmark query consists of two free-connex sub-queries, hence can be evaluated in $O\left(N+\mathrm{OUT}_{1}+\mathrm{OUT}_{2}\right)$ time, and its optimized query can be evaluated in $O(N+$ OUT ) time. Due to the special primary-key foreign-key joins and group-by aggregations,

[^2]

Fig. 6. Running time of $Q_{G 4}$ on Google dataset in PostgreSQL with different $\mathrm{OUT}_{1}$.
$\mathrm{OUT}_{1} \approx \mathrm{OUT}_{2} \approx$ OUT $\ll N$, such that the input contains a few hundred of million records while the query result only involves thousands of records. The improvement of our optimized techniques in SQLite, DuckDB, and MySQL is also limited. On some test points, our optimized SQL queries are even more time-consuming than vanilla SQL queries. We find that the vanilla SQL queries can greatly benefit from the indices built for primary-key foreign-key join and outperform our optimized SQL queries, which do not enjoy efficient indices for set difference or anti-join operators. How to build indices to accelerate relational operators in these systems could be interesting future work. Meanwhile, we notice that loading input data and building indices are much more timeconsuming than evaluating the query; for example, it takes DuckDB 16 minutes to load a 50G-sized TPC-DS dataset, while only 8 seconds to execute the whole query.
Impact of OUT, $\mathrm{OUT}_{1}$ and $\mathrm{OUT}_{2}$. Implied by the theoretical results, the sizes $\mathrm{OUT}_{1}, \mathrm{OUT}_{2}$ of sub-queries $Q_{1}, Q_{2}$ impact the performance of vanilla SQL queries, while only the actual output size OUT affect the performance of our optimized SQL queries. Below, we study the impact of $\mathrm{OUT}_{1}$, $\mathrm{OUT}_{2}$ and OUT on the performance of both approaches over $Q_{G 4}$.

In Figure 6, we investigate the impact of $\mathrm{OUT}_{1}$ for computing DCQ. We fix $Q_{2}$ (as well as $\mathrm{OUT}_{2}$ ) and only vary the size of Triple (as well as $N$ and OUT $_{1}$ ). Note that OUT also increases as OUT ${ }_{2}$ decreases. The running time of our optimized SQL query grows slowly with OUT, while the vanilla SQL query incurs a fixed overhead for evaluating $Q_{2}$, even when OUT $_{1}$ (as well as OUT) decreases to as small as 1 .


Fig. 7. Running time of $Q_{G 4}$ on Google dataset in PostgreSQL with different $\mathrm{OUT}_{2}$.

In Figure 7, we investigate the impact of $\mathrm{OUT}_{2}$ for computing DCQ. We fix $Q_{1}$ (as well as $N$ and $\mathrm{OUT}_{1}$ ) and vary a filter predicate applied to relation Graph in $Q_{2}$. When the predicate is more selective, $\mathrm{OUT}_{2}$ becomes smaller, and OUT becomes larger. The running time of vanilla SQL query decreases as $\mathrm{OUT}_{2}$ decreases, and the running time of our optimized SQL query does not change, which is only affected by $N$ and OUT.


Fig. 8. Running time of $Q_{G 4}$ on Google dataset in PostgreSQL with different OUT.

In Figure 8, we investigate the impact of OUT for computing DCQ. We adjust Triple by changing the proportion of tuples generated by different rules, which will only change OUT, while $\mathrm{OUT}_{1}$, $\mathrm{OUT}_{2}$, and $N$ stay the same. The running time of our optimized SQL query increases slowly as OUT increases. In contrast, the running time of vanilla SQL query remains stably high even when OUT decreases to 1 , since its running time is only impacted by $\mathrm{OUT}_{1}$ and $\mathrm{OUT}_{2}$, both of which stay unchanged.


Fig. 9. Memory usage by graph queries on Epinions.

Memory Consumption. We also test the memory consumption on both graph and benchmark queries by different engines. Due to the simplicity of memory consumption measurement, we report the results for PostgreSQL and DuckDB here. For benchmark queries, the optimized and vanilla SQL queries have similar behaviors on memory consumption, since the input size dominates the overall consumption. Below, we focus on the memory consumption of graph queries. In Figure 9, our optimized SQL queries achieve overall improvements for all graph queries on Epinions dataset in terms of space consumption. For example, our optimized SQL query only requires 6.53 GB on Spark SQL for evaluating $Q_{G 6}$, while the vanilla SQL query fails to finish evaluating $Q_{G 6}$ even using

256G memory. The improvement of our optimized SQL query is more significant on DuckDB. For $Q_{G 4}$, our optimized SQL query consumes $99.4 \%$ less memory than the vanilla SQL query. For $Q_{G 5}$ and $Q_{G 6}$, our optimized SQL queries consume roughly 2G memory. In contrast, the vanilla SQL queries fail to execute due to out-of-memory errors even after using 738G memory.

## 7 CONNECTION WITH SIGNED CONJUNCTIVE QUERY

The class of signed conjunctive queries (SCQ) [19], or noted as conjunctive queries with negation [32] in the literature, is defined as

$$
Q:=\pi_{\mathrm{y}}\left(\eta_{1} R_{1}\left(e_{1}\right) \bowtie \cdots \bowtie \eta_{2} R_{2}\left(e_{2}\right) \bowtie \cdots \bowtie \eta_{n} R_{n}\left(e_{n}\right)\right),
$$

where $\eta_{i}$ is either empty or a negation operator $\neg$. If $\eta_{i}=\neg$ for all $i \in[n]$, such an SCQ is also known as a negative conjunctive queries (NCQ). If $\eta_{i}=\emptyset$ for all $i \in[n]$, such an SCQ is also known as a CQ. Recall that $\mathcal{V}=e_{1} \cup e_{2} \cup \cdots \cup e_{n}$ and $\mathcal{E}=\left\{e_{1}, e_{2}, \cdots, e_{n}\right\}$. The query result of $Q$ over an instance $D$ denoted as $Q(D)$ is defined as

$$
Q(D)=\left\{t \in \operatorname{dom}(\mathbf{y}): \exists t^{\prime} \in \operatorname{dom}(\mathcal{V}), \pi_{e_{i}} t \in R_{i}, \forall \eta_{i}=\emptyset, \pi_{e_{j}} t \notin R_{j}, \forall \eta_{j}=\neg\right\} .
$$

We establish the connection between SCQ and DCQ via Lemma 7.1 and Lemma 7.2.
From DCQ to SCQ. Intuitively, every DCQ can be rewritten as the union of a set of SCQs. Moreover, each resulted SCQ has exactly one negated relation, and each relation of $Q_{2}$ participates in one distinct SCQ as the negated relation. For example, $Q_{1}-Q_{2}=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)-$ $R_{3}\left(x_{1}, x_{2}\right) \bowtie R_{4}\left(x_{2}, x_{3}\right)$ can be rewritten as $\left(R_{1} \bowtie R_{2} \bowtie \neg R_{3}\right) \cup\left(R_{1} \bowtie R_{2} \bowtie \neg R_{4}\right)$.

Lemma 7.1. For a $D C Q Q_{1}-Q_{2}, Q_{1}-Q_{2}=\bigcup_{e \in \mathcal{E}_{2}}\left(Q_{1} \bowtie \neg R_{e}\right)$.
Proof. Direction $\subseteq$. For every join result $t \in Q_{1}-Q_{2}$, there must exist a relation $e \in \mathcal{E}_{2}$ such that $\pi_{e} t \notin R_{e}$; otherwise, $t \in Q_{2}$, coming to a contradiction. Wlog, let $e \in \mathcal{E}_{2}$ be such a relation for $t$. Together with $t \in Q_{1}$, there must be $t \in Q_{1} \bowtie \neg R_{e}$. Direction Ə. Consider an arbitrary relation $e \in \mathcal{E}_{2}$, and an arbitrary join result $t \in Q_{1} \bowtie \neg R_{e}$. Obviously, $t \notin Q_{2}$ since $\pi_{e} t \notin R_{2}$. Together with $t \in Q_{1}$, there must be $t \in Q_{1}-Q_{2}$.

From SCQ to DCQ. On the other hand, SCQ can be rewritten as the intersection of a set of DCQs. For a SCQ $Q=(\mathcal{V}, \mathcal{E})$, let $\mathcal{E}^{+}, \mathcal{E}^{-} \subseteq \mathcal{E}$ denote the set of relations with positive, negative sign separately. Let $\mathcal{V}^{+}, \mathcal{V}^{-}$be the set of attributes that appear in positive, negative relations separately. Let $Q^{+}=\left(\bowtie_{e^{\prime} \in \mathcal{E}^{+}} R_{e^{\prime}} \times_{x \in \mathcal{V}^{-}-\mathcal{V}^{+}} \operatorname{dom}(x)\right)$ denote the positive subquery defined by positive relation as well as the whole domain of attributes which do not appear in any positive relation.

Lemma 7.2. For a $\operatorname{SCQ} Q, Q=\cap_{e \in \mathcal{E}^{-}}\left(Q^{+}-Q^{+} \bowtie R_{e}\right)$.
Proof. Direction $\subseteq$. Consider an arbitrary query result $t \in Q$. By definition, $\pi_{e} t \in R_{e}$ holds for every $e \in \mathcal{E}^{+}$and $\pi_{e} t \notin R_{e}$ holds for every $e \in \mathcal{E}^{-}$. This way, for each $e \in \mathcal{E}^{-}$, we have $t \in\left(\bowtie_{e^{\prime} \in \mathcal{E}^{+}} R_{e^{\prime}} \times_{x \in \mathcal{V}^{-}-\mathcal{V}^{+}} \operatorname{dom}(x)\right) \bowtie \neg R_{e}$. Direction $\supseteq$. Consider an arbitrary $t$ such that for every $e \in \mathcal{E}^{-}, t \in\left(\bowtie_{e^{\prime} \in \mathcal{E}^{+}} R_{e^{\prime}} \times_{x \in \mathcal{V}^{-}-\mathcal{V}^{+}} \operatorname{dom}(x)\right) \bowtie \neg R_{e}$. Then, $\pi_{e^{\prime}} t \in R_{e^{\prime}}$ for every $e^{\prime} \in \mathcal{E}^{+}$but $t \notin R_{e}$ for every $e \in \mathcal{E}^{-}$. Thus, $t \in Q$.

For example, a SCQ $Q=R_{1}\left(x_{2}, x_{3}, x_{4}\right) \bowtie R_{2}\left(x_{1}, x_{3}, x_{4}\right) \bowtie \neg R_{3}\left(x_{1}, x_{2}, x_{4}\right) \bowtie \neg R_{4}\left(x_{1}, x_{2}, x_{3}\right)$ can be rewritten as: $\left(R_{1} \bowtie R_{2}-R_{1} \bowtie R_{2} \bowtie R_{3}\right) \cap\left(R_{1} \bowtie R_{2}-R_{1} \bowtie R_{2} \bowtie R_{4}\right)$.
Decidability of SCQ. Given a SCQ $Q$, the domain of attributes, and input database $D$, the decidability problem asks to decide whether there exists a query result in $Q$. For example, a NCQ $Q=$ $\neg R_{1}\left(x_{1}, x_{2}\right) \bowtie \neg R_{2}\left(x_{2}, x_{3}\right)$ decides if there exists any tuple $(a, b, c) \in \operatorname{dom}\left(x_{1}\right) \times \operatorname{dom}\left(x_{2}\right) \times \operatorname{dom}\left(x_{3}\right)$ such that $(a, b) \notin R_{1}$ and $(b, c) \notin R_{2}$, and a $\mathrm{CQ} Q=R_{1}\left(x_{1}, x_{2}\right) \bowtie R_{2}\left(x_{2}, x_{3}\right)$ decides if there exists
any tuple $(a, b, c)$ such that $(a, b) \in R_{1}$ and $(b, c) \in R_{2}$. The decidability problem for CQ , NCQ and SCQ has been well studied separately:
Theorem 7.3 ([16]). A CQQ can be decided in linear time if and only if it is $\alpha$-acyclic.
Theorem 7.4 ([19]). A NCQ Q can be decided in linear time if and only if it is $\beta$-acyclic.
Theorem 7.5 ([20]). A SCQQ can be decided in linear time if and only if $\left(\mathbf{y}, \mathcal{E}^{+} \cup S\right)$ is $\alpha$-acyclic for every $S \subseteq \mathcal{E}^{-}$.

Note that $\beta$-acyclicity is a more restricted notion than $\alpha$-acyclicity, such that $Q$ is $\beta$-acyclic if all sub-hypergraphs of $Q$ are $\alpha$-acyclic. Obviously, $\beta$-acyclicity strictly implies $\alpha$-acyclicity. In [32], this notion of $\beta$-acyclicity has been extended to nest-set width for capturing the tractability of SCQ in terms of both query and data complexity. We won't pursue this direction further.
Decidability of DCQ. Implied by Lemma 7.1 and Theorem 7.5, we come to the following lemma:
Lemma 7.6. Given two full joins $Q_{1}=\left(\mathbf{y}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathbf{y}, \mathcal{E}_{2}\right)$, the $D C Q Q_{1}-Q_{2}$ can be decided in linear time, if $\left(\mathbf{y}, \mathcal{E}_{1}\right)$ is $\alpha$-acyclic, and $\left(\mathbf{y}, \mathcal{E}_{1} \cup\{e\}\right)$ is $\alpha$-acyclic for everye $\in \mathcal{E}_{2}$.

Lemma 7.6 can be easily proved by a linear-time algorithm. Note that $Q_{1}-Q_{2}=\bigcup_{e \in \mathcal{E}_{2}}\left(Q_{1} \bowtie \neg R_{e}\right)$. We can enumerate every tuple in $S_{e}=\pi_{e} Q_{1}$ within $O(1)$ delay, as $\left(\mathbf{y}, \mathcal{E}_{1} \cup\{e\}\right)$ is $\alpha$-acyclic. For each tuple $t \in S_{e}$ enumerated, we check whether it belongs to $R_{e}$. If $t \notin R_{e}$, a query result of $Q \bowtie \neg R_{e}$ is found; otherwise, we skip it and continue to the next one. It is easy to see that at most $O\left(\left|R_{e}\right|\right)$ tuples are checked, so this algorithm runs in $O(N)$ time.

Theorem 7.7. Given two full joins $Q_{1}=\left(\mathbf{y}, \mathcal{E}_{1}\right)$ and $Q_{2}=\left(\mathbf{y}, \mathcal{E}_{2}\right)$, the $D C Q Q_{1}-Q_{2}$ can be decided in linear time, if and only if $\left(\mathbf{y}, \mathcal{E}_{1}\right)$ is $\alpha$-acyclic, as well as $\left(\mathbf{y}, \mathcal{E}_{1} \cup\{e\}\right)$ is $\alpha$-acyclic for every $e \in \mathcal{E}_{2}$.

Proof. The if direction follows Lemma 7.6. We next distinguish two more cases for the only-if direction. (1) if $\left(\mathbf{y}, \mathcal{E}_{1}\right)$ is cyclic; and (2) if $\left(\mathbf{y}, \mathcal{E}_{1}\right)$ is $\alpha$-acyclic, and there exists some $e \in \mathcal{E}_{2}$ such that ( $\mathbf{y}, \mathcal{E}_{1} \cup\{e\}$ ) is cyclic. (1) follows Theorem 7.3 by simply setting $Q_{2}=\emptyset$. (2) follows Lemma 4.7.

## 8 RELATED WORK

Union of CQs. [21] studied the enumeration complexity of union of conjunctive queries (UCQs), i.e., the goal is to find a data structure that after linear preprocessing time, the query answers (without duplication) can be enumerated within a small delay. Their results implied a linear algorithm in terms of input and output size for the class of union-free-connex UCQs, but whether a linear algorithm can be achieved (and, if possible, how to achieve it) is unknown for the remaining class of UCQ. [23] also investigated the enumeration complexity of UCQs but in the dynamic scenario.
Selection over CQs. Recently, multiple works have studied the complexity of selections over conjunctive queries. [37] investigated the selection in the form of comparisons between two attributes or values. The work identifies an acyclic condition under which a near-linear-time algorithm can be achieved for conjunctive queries with comparisons. [13] worked on the selections over intervals, also known as intersection queries, which are special cases for comparison queries since each intersection query can be decomposed into a union of multiple comparison queries. They show a dichotomy result that an intersection join can be computed in linear time if and only if it is $l$-acyclic. [28] studied the complexity of temporal queries, where the intersection condition only exists for one global attribute. Their result suggested that a temporal query can be solved in linear time if and only if it is r-hierarchical. [35] also investigated the complexity of intersection queries in dynamic settings.

## REFERENCES

[1] DuckDB. https://duckdb.org/.
[2] MySQL. https://www.mysql.com/.
[3] Oracle. https://www.oracle.com/.
[4] PostgreSQL. https://www.postgre.org/.
[5] SNAP. https://snap.stanford.edu/snap/.
[6] SparkSQL. https://spark.apache.org/sql/.
[7] SQLite. https://www.sqlite.org/.
[8] TPC-DS. https://www.tpc.org/tpcds/.
[9] TPC-H. https://www.tpc.org/tpch/.
[10] https://arxiv.org/abs/2302.13140.
[11] Amir Abboud and Virginia Vassilevska Williams. 2014. Popular conjectures imply strong lower bounds for dynamic problems. In 2014 IEEE 55th Annual Symposium on Foundations of Computer Science. IEEE, 434-443.
[12] Serge Abiteboul, Richard Hull, and Victor Vianu. 1995. Foundations of databases. Vol. 8. Addison-Wesley Reading.
[13] Mahmoud Abo Khamis, George Chichirim, Antonia Kormpa, and Dan Olteanu. 2022. The Complexity of Boolean Conjunctive Queries with Intersection Joins. In Proceedings of the 41st ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (Philadelphia, PA, USA) (PODS '22). Association for Computing Machinery, New York, NY, USA, 53-65. https://doi.org/10.1145/3517804.3524156
[14] Mahmoud Abo Khamis, Hung Q Ngo, and Atri Rudra. 2016. FAQ: questions asked frequently. In Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems. 13-28.
[15] Rasmus Resen Amossen and Rasmus Pagh. 2009. Faster join-projects and sparse matrix multiplications. In Proceedings of the 12th International Conference on Database Theory. 121-126.
[16] Guillaume Bagan, Arnaud Durand, and Etienne Grandjean. 2007. On acyclic conjunctive queries and constant delay enumeration. In International Workshop on Computer Science Logic. Springer, 208-222.
[17] Catriel Beeri, Ronald Fagin, David Maier, and Mihalis Yannakakis. 1983. On the desirability of acyclic database schemes. Journal of the ACM (7ACM) 30, 3 (1983), 479-513.
[18] Andreas Björklund, Rasmus Pagh, Virginia Vassilevska Williams, and Uri Zwick. 2014. Listing triangles. In International Colloquium on Automata, Languages, and Programming. Springer, 223-234.
[19] Johann Brault-Baron. 2012. A Negative Conjunctive Query is Easy if and only if it is Beta-Acyclic. In Computer Science Logic (CSL'12)-26th International Workshop/21st Annual Conference of the EACSL. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
[20] Johann Brault-Baron. 2013. De la pertinence de l'énumération: complexité en logiques propositionnelle et du premier ordre. Ph. D. Dissertation. Université de Caen.
[21] Nofar Carmeli and Markus Kröll. 2019. On the Enumeration Complexity of Unions of Conjunctive Queries. In Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems. ACM, 134-148.
[22] Chandra Chekuri and Anand Rajaraman. 2000. Conjunctive query containment revisited. Theoretical Computer Science 239, 2 (2000), 211-229.
[23] Berkholz Christoph, Keppeler Jens, and Schweikardt Nicole. 2018. Answering UCQs under Updates and in the presence of integrity constraints. In Proceedings of the 21st International Conference on Database Theory (ICDT'18), Vol. 98. 1-8.
[24] Shaleen Deep, Xiao Hu, and Paraschos Koutris. 2020. Fast join project query evaluation using matrix multiplication. In Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data. 1213-1223.
[25] R. Fagin. 1983. Degrees of acyclicity for hypergraphs and relational database schemes. JACM 30, 3 (1983), 514-550.
[26] Georg Gottlob, Nicola Leone, and Francesco Scarcello. 2002. Hypertree decompositions and tractable queries. 7. Comput. System Sci. 64, 3 (2002), 579-627.
[27] Georg Gottlob, Zoltán Miklós, and Thomas Schwentick. 2009. Generalized hypertree decompositions: NP-hardness and tractable variants. fournal of the ACM (JACM) 56, 6 (2009), 1-32.
[28] Xiao Hu, Stavros Sintos, Junyang Gao, Pankaj K. Agarwal, and Jun Yang. 2022. Computing Complex Temporal Join Queries Efficiently. In Proceedings of the 2022 International Conference on Management of Data (Philadelphia, PA, USA) (SIGMOD '22). Association for Computing Machinery, New York, NY, USA, 2076-2090. https://doi.org/10.1145/3514221. 3517893
[29] Zichun Huang and Shimin Chen. 2022. Density-Optimized Intersection-Free Mapping and Matrix Multiplication for Join-Project Operations. Proc. VLDB Endow. 15, 10 (jun 2022), 2244-2256. https://doi.org/10.14778/3547305.3547326
[30] Manas R. Joglekar, Rohan Puttagunta, and Christopher Ré. 2016. AJAR: Aggregations and Joins over Annotated Relations. In Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (San Francisco, California, USA) (PODS '16). Association for Computing Machinery, New York, NY, USA, 91-106. https://doi.org/10.1145/2902251.2902293
[31] Phokion G Kolaitis and Moshe Y Vardi. 1998. Conjunctive-query containment and constraint satisfaction. In Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems. 205-213.
[32] Matthias Lanzinger. 2021. Tractability Beyond ß-Acyclicity for Conjunctive Queries with Negation. In Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems. 355-369.
[33] Hung Q Ngo, Ely Porat, Christopher Ré, and Atri Rudra. 2018. Worst-case optimal join algorithms. Journal of the ACM ( $7 A C M)$ 65, 3 (2018), 1-40.
[34] Christos H Papadimitriou and Mihalis Yannakakis. 1997. On the complexity of database queries. In Proceedings of the sixteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems. 12-19.
[35] Yufei Tao and Ke Yi. 2022. Intersection joins under updates. F. Comput. System Sci. 124 (2022), 41-64. https: //doi.org/10.1016/j.jcss.2021.09.004
[36] Moshe Y Vardi. 1982. The complexity of relational query languages. In Proceedings of the fourteenth annual ACM symposium on Theory of computing. 137-146.
[37] Qichen Wang and Ke Yi. 2022. Conjunctive Queries with Comparisons. In Proceedings of the 2022 International Conference on Management of Data (Philadelphia, PA, USA) (SIGMOD '22). Association for Computing Machinery, New York, NY, USA, 108-121. https://doi.org/10.1145/3514221.3517830
[38] Mihalis Yannakakis. 1981. Algorithms for acyclic database schemes. In VLDB, Vol. 81. 82-94.
Received October 2022; revised January 2023; accepted February 2023


[^0]:    Authors' addresses: Xiao Hu, xiaohu@uwaterloo.ca, University of Waterloo, 200 University Avenue West, Waterloo, Ontario, Canada, N2L 3G1; Qichen Wang, qcwang@hkbu.edu.hk, Hong Kong Baptist University, 224 Waterloo Road, Kowloon Tong, Hong Kong, Hong Kong.

    Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
    © 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
    2836-6573/2023/6-ART153 \$15.00
    https://doi.org/10.1145/3589298

[^1]:    ${ }^{1}$ We distinguish the input instances $D_{1}, D_{2}$ of $Q_{1}, Q_{2}$ for simplifying algorithmic description later, which is different from conventional definition.

[^2]:    ${ }^{2}$ We review the execution plan in MySQL and find that the predicated run-time of our optimized SQL query is much smaller than the vanilla SQL query, which is also consistent with our observations in other platforms. The actual running time does not match the expected cost because of some unknown deficiencies in MySQL.

