# Finding Smallest Witnesses for Conjunctive Queries

**Xiao Hu**
University of Waterloo, Canada

**Stavros Sintos**
University of Illinois Chicago, USA

## Abstract

A witness is a sub-database that preserves the query results of the original database but of much smaller size. It has wide applications in query rewriting and debugging, query explanation, IoT analytics, multi-layer network routing, etc. In this paper, we study the smallest witness problem (`SWP`) for the class of conjunctive queries (CQs) without self-joins.

We first establish the dichotomy that SWP for a CQ can be computed in polynomial time if and only if it has *head-cluster property*, unless $\texttt{P} = \texttt{NP}$. We next turn to the approximated version by relaxing the size of a witness from being minimum. We surprisingly find that the *head-domination* property - that has been identified for the deletion propagation problem [35] - can also precisely capture the hardness of the approximated smallest witness problem. In polynomial time, `SWP` for any CQ with head-domination property can be approximated within a constant factor, while `SWP` for any CQ without such a property cannot be approximated within a logarithmic factor, unless $\texttt{P} = \texttt{NP}$.

We further explore efficient approximation algorithms for CQs without head-domination property: (1) we show a trivial algorithm which achieves a polynomially large approximation ratio for general CQs; (2) for any CQ with only one non-output attribute, such as star CQs, we show a greedy algorithm with a logarithmic approximation ratio; (3) for line CQs, which contain at least two non-output attributes, we relate `SWP` problem to the directed steiner forest problem, whose algorithms can be applied to line CQs directly. Meanwhile, we establish a much higher lower bound, exponentially larger than the logarithmic lower bound obtained above. It remains open to close the gap between the lower and upper bound of the approximated `SWP` for CQs without head-domination property.

**2012 ACM Subject Classification** Theory of computation → Data provenance

**Keywords and phrases** conjunctive query, smallest witness, head-cluster, head-domination

**Digital Object Identifier** 10.4230/LIPIcs...

## 1 Introduction

To deal with large-scale data in analytical applications, people have developed a large body of data summarization techniques to reduce computational as well as storage complexity, such as sampling [12, 47, 14], sketch [18], coreset [41] and factorization [39]. The notion of witness has been studied as one form of why-provenance [11, 30, 3] that provides a *proof* for output results, with wide applications in explainable data-intensive analytics. The *smallest witness problem* was first proposed by [37] that given a query $Q$, a database $D$ and one specific query result $t \in Q(D)$, the target is to find the smallest sub-database $D' \subseteq D$ such that $t$ is witnessed by $D'$, i.e., $t \in Q(D')$. In this paper, we consider a generalized notion for all query results, i.e., our target is to find the smallest sub-database $D' \subseteq D$ such that all query results can be witnessed by $D'$, i.e., $Q(D) = Q(D')$. Our generalized smallest witness has many useful applications in practice, such as helping students learn SQL queries [37], query rewriting, query explanation, multi-layer network routing, IoT analytics on edge devices[40, 46]. We mention three application scenarios:

▶ **Example 1.** Alice located at Seattle wants to send the query results of $Q$ over a database $D$ (which is also stored at Seattle) to Bob located at New York. Unfortunately, the number of query results could be polynomially large in terms of the number of tuples in $D$. An alternative is to send the entire database $D$, since Bob can retrieve all query results by executing $Q$ over $D$ at New York. However, moving the entire database is also expensive. A natural question arises: is it necessary for Alice to send the entire $D$ or $Q(D)$? If not, what is the smallest subset of tuples to send?

▶ **Example 2.** Charlie is a novice at learning SQL in a undergraduate database course. Suppose there is a huge test database $D$, a correct query $Q$ that Charlie is expected to learn, and a wrong query $Q'$ submitted by him, where some answers in $Q(D)$ are missed from $Q'(D)$. To help Charlie debug, the instructor can simply show the whole test database $D$ to him. However, Charlie will have to dive into such a huge database to figure out where his query goes wrong. A natural question arises: is it necessary to show the entire $D$ to Charlie? If not, what is the smallest subset of tuples to show so that Charlie can quickly find all missing answers by his wrong query?

▶ **Example 3.** In a multi-layer communication or transportation network, clients and servers are connected by routers organized into layers, such that links (or edges) exist between routers residing in consecutive layers. What is the smallest subset of links needed for building a fully connected network, i.e., every client-server pair is connected via a directed path? For a given network, what is the maximum number of links that can be broken while the connectivity with respect to the client-server pairs does not change? This information could help evaluate the inherent robustness of a network to either malicious attacks or even just random failures.

Recall that our smallest witness problem finds the smallest sub-database $D' \subseteq D$ such that $Q(D) = Q(D')$. It would be sufficient for Alice to send $D'$, while Bob can retrieve all query results by executing $Q$ over $D'$, saving much transmission cost. Also, it would be sufficient for the instructor to show $D'$ to Charlie, from which all correct answers in $Q(D)$ can be recovered, saving Charlie much efforts in exploring a huge test database.[1] Moreover, the connectivity of a multi-layer network $D$ can be modeled as a line query $Q$ (formally defined in Section 5.3) with connected client-server pairs as $Q(D)$, such that $D'$ is a smallest subset of links needed for maintaining the desired connectivity, and $D - D'$ is a maximum subset of links that can be removed safely. In this paper, we aim to design algorithms that can efficiently compute or approximate the smallest witness for conjunctive queries, and understand the hardness of this problem when such algorithms do not exist.

## 1.1 Problem Definition

Let $\mathbb{R}$ be a database schema that contains $m$ relations $R_1, R_2, \cdots, R_m$. Let $\mathbb{A}$ be the set of all attributes in $\mathbb{R}$. Each relation $R_i$ is defined on a subset of attributes $\mathbb{A}_i \subseteq \mathbb{A}$. We use $A, B, C, A_1, A_2, A_3, \cdots$ etc. to denote the attributes in $\mathbb{A}$ and $a, b, c, \cdots$ etc. to denote their values. Let $\mathtt{dom}(A)$ be the domain of attribute $A \in \mathbb{A}$. The domain of a set of attributes $X \subseteq \mathbb{A}$ is defined as $\mathtt{dom}(X) = \prod_{A \in X} \mathtt{dom}(A)$. Given the database schema $\mathbb{R}$, let $D$ be a given database of $\mathbb{R}$, and let the corresponding relations of $R_1, \cdots, R_m$ be $R_1^D, \cdots, R_m^D$, where $R_i^D$ is a collection of tuples defined on $\mathtt{dom}(\mathbb{A}_i)$. The *input size* of database $D$ is denoted as $N = |D| = \sum_{i \in [m]} |R_i^D|$. Where $D$ is clear from the context, we will drop the superscript.

[1] The smallest witness for a single query result [36] has been incorporated into an educational tool (https://dukedb-hnrq.github.io/), successfully employed in Duke database courses with 1,000+ undergraduate users. Our generalized version can also be incorporated and save more efforts by showing one small witness for all answers.

We consider the class of *conjunctive queries without self-joins*:

$$Q(\mathbf{A}) : -R_1(\mathbb{A}_1), R_2(\mathbb{A}_2), \cdots, R_m(\mathbb{A}_m)$$

where $\mathbf{A} \subseteq \mathbb{A}$ is the set of *output attributes* (a.k.a. *free attributes*) and $\mathbb{A} - \mathbf{A}$ is the set of *non-output attributes* (a.k.a. *existential attributes*). A CQ is *full* if $\mathbf{A} = \mathbb{A}$, indicating the natural join among the given relations; otherwise, it is *non-full*. Each $R_i$ in $Q$ is distinct, i.e., $Q$ does not have a self-join. When a CQ $Q$ is evaluated on database $D$, its query result denoted as $Q(D)$ is the projection of natural join result of $R_1(\mathbb{A}_1) \bowtie R_2(\mathbb{A}_2) \bowtie \cdots \bowtie R_m(\mathbb{A}_m)$ onto $\mathbf{A}$ (after removing duplicates).

▶ **Definition 4.** [Smallest Witness Problem (`SWP`)] *For CQ $Q$ and database $D$, it asks to find a subset of tuples $D' \subseteq D$ such that $Q(D) = Q(D')$, while there exists no subset of tuples $D'' \subseteq D$ such that $Q(D'') = Q(D)$ and $|D''| < |D'|$.*

Given $Q$ and $D$, we denote the above problem by `SWP`$(Q, D)$. See an example in Appendix A. We note that the solution to `SWP`$(Q, D)$ may not be unique, hence our target simply finds one such solution. We study the data complexity [43] of `SWP` i.e., the sizes of database schema and query are considered as constants, and the complexity is in terms of input size $N$. For any CQ $Q$ and database $D$, the size of query results $|Q(D)|$ is polynomially large in terms of $N$, and $Q(D)$ can also be computed in polynomial time in terms of $N$. In contrast, the size of `SWP`$(Q, D)$ is always smaller than $N$, while as we see later `SWP`$(Q, D)$ may not be computed in polynomial time in terms of $N$. Again, our target is to compute the smallest witness instead of the query results. We say that `SWP` is *poly-time solvable* for $Q$ if, for an arbitrary database $D$, `SWP`$(Q, D)$ can be computed in polynomial time in terms of $|D|$. As shown later, `SWP` is not poly-time solvable for a large class of CQs, so we introduce an approximated version:

▶ **Definition 5.** [$\theta$-Approximated Smallest Witness Problem (`ASWP`)] *For CQ $Q$ and database $D$, it asks to find a subset of tuples $D' \subseteq D$ such that $Q(D') = Q(D)$ and $|D'| \leq \theta \cdot |D^*|$, where $D^*$ is a solution to `SWP`$(Q, D)$.*

Also, `SWP` is $\theta$-approximable for $Q$ if, for an arbitrary database $D$, there is a $\theta$-approximated solution to `SWP`$(Q, D)$ that can be computed in polynomial time in terms of $|D|$.

## 1.2 Our Results

Our main results obtained can be summarized as (see Figure 1):

In Section 3, we obtain a dichotomy of computing `SWP` for CQs. More specifically, `SWP` for any CQ with head-cluster property (Definition 11) can be solved by a trivial poly-time algorithm, while `SWP` for any CQ without head-cluster property is NP-hard by resorting to the *NP-hardness of set cover* problem (Section 3.2).

In Section 4, we show a dichotomy of approximating `SWP` for CQs without head-cluster property. The *head-domination* property that has been identified for *deletion propagation* problem [35], also captures the hardness of approximating `SWP`. We show a poly-time algorithm that can return a $O(1)$-approximated solution to `SWP` for CQs with head-domination property. On the other hand, we prove that `SWP` cannot be approximated within a factor of $(1 - o(1)) \cdot \log N$ for every CQ without head-domination property, unless `P` $=$ `NP`, by resorting to the *logarithmic inapproximability of set cover* problem (in Section 4.2). Interestingly, this separation on approximating `SWP` for *acyclic* CQs (in Section 2.1) coincides with the separation of *free-connex* and *non-free-connex* CQs (in Section 4.1) in the literature. In Section 5, we further explore approximation algorithms for CQs without head-domination property. Firstly,

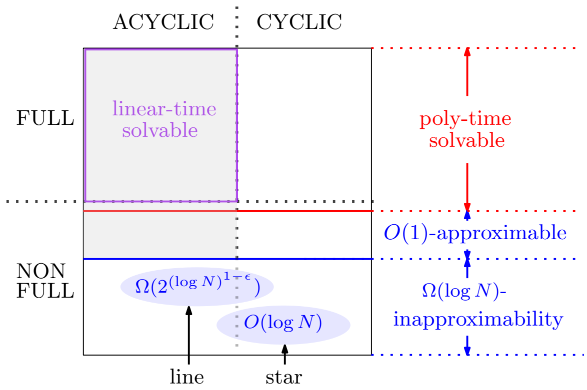


**Figure 1** Summary of our results. The shadow area is the class of free-connex CQs.

we show a baseline of returning the union of witnesses for every query result leads to a $O(N^{1-1/\rho^*})$-approximated solution, where $\rho^*$ is the fractional edge covering number[2] of the input CQ [4]. Furthermore, for any CQ with only one non-output attribute, which includes the commonly-studied *star* CQs, we show a greedy algorithm that can approximate SWP within a $O(\log N)$ factor, matching the lower bound. However, for another commonly-studied class of *line* CQs, which contains more than two non-output attributes, we prove a much higher lower bound $\Omega(2^{(\log N)^{1-\epsilon}})$ for any $\epsilon > 0$ in approximating SWP, by resorting to the *label cover* problem (Appendix E). Meanwhile, we observe that SWP problem for line queries is a special case of the *directed steiner forest* (DSF) problem ( Section 5.3), and therefore existing algorithms for DSF can be applied to SWP directly. But, how to close the gap between the upper and lower bounds on approximating SWP for line CQs remains open.

# 2 Preliminaries

## 2.1 Notations and Classifications of CQs

Extending the notation in Section 1.1, we use $\texttt{rels}(Q)$ to denote all the relations that appear in the body of $Q$, and use $\texttt{attr}(Q), \texttt{head}(Q) \subseteq \texttt{attr}(Q)$ to denote all the attributes that appear in the body, head of $Q$ separately (so, $\texttt{head}(Q) = \mathbf{A}$ in Section 1.1). Moreover, $\texttt{head}(R_i) = \texttt{head}(Q) \cap \texttt{attr}(R_i)$. For any attribute $A \in \texttt{attr}(R_i)$, $\pi_A t$ denotes the value over attribute $A$ of tuple $t$. Similarly, for a set of attributes $X \subseteq \texttt{attr}(R_i)$, $\pi_X t$ denotes values over attributes in $X$ of tuple $t$. We also mention two important classes of CQs.

▶ **Definition 6** (Acyclic CQs [7, 22]). *A CQ $Q$ is acyclic if there exists a tree $\mathbb{T}$ such that (1) there is a one-to-one correspondence between the nodes of $\mathbb{T}$ and relations in $Q$; and (2) for every attribute $A \in \texttt{attr}(Q)$, the set of nodes containing $A$ forms a connected subtree of $\mathbb{T}$. Such a tree is called the* join tree *of $Q$.*

▶ **Definition 7** (Free-connex CQs [6]). *A CQ $Q$ is free-connex if $Q$ is acyclic and the resulted CQ by adding another relation contains exactly $\texttt{head}(Q)$ to $Q$ is also acyclic.*

[2] For a CQ $Q$, a fractional edge covering is a function $W : \texttt{rels}(Q) \to [0, 1]$ with $\sum_{R_i : A \in \texttt{attr}(R_i)} W(R_i) \geq 1$ for every attribute $A \in \mathbb{A}$. The fractional edge covering number of $Q$ is the minimum value of $\sum_{R_i : R_i \in \texttt{rels}(Q)} W(R_i)$ over all fractional edge coverings.

## 2.2 SWP for One Query Result

The SWP problem for one query result is formally defined as:

▶ **Definition 8.** [SWP for One Query Result] *For CQ $Q$, database $D$ and query result $t \in Q(D)$, it asks for finding a subset of tuples $D' \subseteq D$ such that $t \in Q(D')$, while there is no subset $D'' \subseteq D$ such that $t \in Q(D'')$ and $|D''| < |D'|$.*

Given $Q$, $D$ and $t$, we denote the above problem by $\texttt{SWP}(Q, D, t)$. It has been shown by [37] that $\texttt{SWP}(Q, D, t)$ can be computed in polynomial time for arbitrary $Q$, $D$, and $t \in Q(D)$. Their algorithm [37] simply finds an arbitrary full join result $t' \in \Join_{R_i \in \texttt{rels}(Q)} R_i$ such that $\pi_{\texttt{head}(Q)} t' = t$, and returns all participating tuples in $\left\{\pi_{\texttt{attr}(R_i)} t' : R_i \in \texttt{rels}(Q)\right\}$ as the smallest witness for $t$. This primitive is used in building our SWP algorithm. The SWP problem for a Boolean CQ ($\mathbf{A} = \emptyset$, indicating whether the result of underlying natural join is empty or not) can be solved by finding SWP for an arbitrary join result in its full version.

## 2.3 Notions of Connectivity

We give three important notions of connectivity for CQs, which will play an important role in characterizing the structural properties used in SWP. See an example in Figure 2.

**Connectivity of CQ.** We capture the *connectivity* of a CQ $Q$ by modeling it as a graph $G_Q$, where each relation $R_i$ is a vertex and there is an edge between $R_i, R_j \in \texttt{rels}(Q)$ if $\texttt{attr}(R_i) \cap \texttt{attr}(R_j) \neq \emptyset$. A CQ $Q$ is *connected* if $G_Q$ is connected, and *disconnected* otherwise. For a disconnected CQ $Q$, we can decompose it into multiple connected subqueries by applying search algorithms on $G_Q$, and finding all connected components for $G_Q$. The set of relations corresponding to the set of vertices in one connected component of $G_Q$ form a connected subquery of $Q$.

Given a disconnected CQ $Q$, let $Q_1, Q_2, \cdots, Q_k$ be its connected subqueries. Given a database $D$ over $Q$, let $D_i \subseteq D$ be the corresponding sub-databases defined for $Q_i$. Observe that every witness for $Q(D)$ is the disjoint union of a witness for $Q_i(D_i)$, for $i \in [k]$. Hence, Lemma 9 follows. In the remaining of this paper, we assume that $Q$ is connected.

▶ **Lemma 9.** *For a disconnected CQ $Q$ of $k$ connected components $Q_1, Q_2, \cdots, Q_k$, SWP is poly-time solvable for $Q$ if and only if SWP is poly-time solvable for every $Q_i$, where $i \in [k]$.*

**Existential-Connectivity of CQ.** We capture the existential-connectivity of a CQ $Q$ by modeling it as a graph $G^{\exists}_Q$, where each relation $R_i$ with $\texttt{attr}(R_i) - \texttt{head}(Q) \neq \emptyset$ is a vertex, and there is an edge between $R_i, R_j \in \texttt{rels}(Q)$ if $\texttt{attr}(R_i) \cap \texttt{attr}(R_j) - \texttt{head}(Q) \neq \emptyset$. We can find the connected components of $G^{\exists}_Q$ by applying search algorithm on $G^{\exists}_Q$, and finding all connected components for $G^{\exists}_Q$. Let $E_1, E_2, \cdots, E_k \subseteq \texttt{rels}(Q)$ be the connected components of $G^{\exists}_Q$, each corresponding to a subset of relations in $Q$.

**Nonout-Connectivity of CQ.** We capture the nonout-connectivity of a CQ $Q$ by modeling it as a graph $H_Q$, where each non-output attribute $A \in \texttt{attr}(Q) - \texttt{head}(Q)$ is a vertex, and there is an edge between $A, B \in \texttt{attr}(Q) - \texttt{head}(Q)$ if there exists a relation $R_i \in \texttt{rels}(Q)$ such that $A, B \in \texttt{attr}(R_i)$. We can find the connected components of $H_Q$, and finding all connected components for $H_Q$. Let $H_1, H_2, \cdots, H_k \subseteq \texttt{attr}(Q) - \texttt{head}(Q)$ be the connected components of $H_Q$, each corresponding to a subset of non-output attributes in $Q$.

## 2.4 Head Cluster and Domination

These two important structural properties identified for characterizing the hardness of (A)SWP are directly built on the existential-connectivity of a CQ $Q$ and the notion of dominant

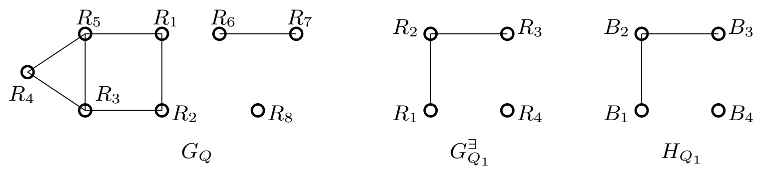


**Figure 2** The left is $G_Q$ for $Q(A_1, A_2, A_3, A_4, A_5)$ : $-$ $R_1(A_1, B_1)$, $R_2(B_1, B_2)$, $R_3(A_2, B_2, B_3)$, $R_4(A_2, A_3, B_4)$, $R_5(A_1, A_2)$, $R_6(A_4, B_5)$, $R_7(B_5, A_5)$, $R_8(B_6, B_7)$ with three sub-queries $Q_1(A_1, A_2, A_3) :- R_1(A_1, B_1)$, $R_2(B_1, B_2)$, $R_3(A_2, B_2, B_3)$, $R_4(A_2, A_3, B_4)$, $R_5(A_1, A_2)$, and $Q_2(A_4, A_5) :- R_6(A_4, B_5), R_7(B_5, A_5)$ and $Q_3 : -R_8(B_6, B_7)$. The middle is $G^{\exists}_{Q_1}$ for $Q_1$, with two connected components $\{R_1, R_2, R_3\}$, $\{R_4\}$ and dominants $R_5$, $R_4$. The right is $H_{Q_1}$ for $Q_1$, with two connected components $\{B_1, B_2, B_3\}, \{B_4\}$.

relation. For a CQ $Q$ with a subset $E \subseteq \texttt{rels}(Q)$ of relations, $R_i \in \texttt{rels}(Q)$ is a *dominant* relation for $E$ if every output attribute appearing in any relation of $E$ also appears in $R_i$, i.e., $\cup_{R_j \in E}\texttt{head}(R_j) \subseteq \texttt{head}(R_i)$.

▶ **Definition 10** (Head Domination [34])**.** *For CQ $Q$, let $E_1, E_2, \cdots, E_k$ be the connected components of $G^{\exists}_Q$. $Q$ has head-domination property if for any $i \in [k]$, there exists a dominant relation from $\texttt{rels}(Q)$ for $E_i$.*

The notion of head-domination property has been first identified for deletion propagation with side effect problem [35], which studied the smallest number of tuples to remove so that a subset of desired query results must disappear while maintain as many as remaining query results. We give a detailed comparison between SWP and deletion propagation in Appendix C, although they solve completely independent problem for CQs without self-joins.

▶ **Definition 11** (Head Cluster)**.** *For CQ $Q$, let $E_1, E_2, \cdots, E_k$ be the connected components of $G^{\exists}_Q$. $Q$ has head-cluster property if for any $i \in [k]$, every $R_j \in E_i$ is a dominant relation for $E_i$.*

There is an equivalent but simpler definition for head-cluster property: A CQ $Q$ has head-cluster property if for every pair of relations $R_i, R_j \in \texttt{rels}(Q)$ with $\texttt{head}(R_i) \neq \texttt{head}(R_j)$, there must be $\texttt{attr}(R_i) \cap \texttt{attr}(R_j) \subseteq \texttt{head}(Q)$. Here, we define head-cluster property based on dominant relation, since it is a special case of head-domination property.

## 3 Dichotomy of Exact SWP

In this section, we focus on computing SWP exactly for CQs, which can be efficiently done if *head-cluster* property is satisfied. All missing proofs are given in Appendix D.

▶ **Theorem 12.** *If a CQ $Q$ has head-cluster property, SWP is poly-time solvable; otherwise, SWP is not poly-time solvable, unless P = NP.*

### 3.1 An Exact Algorithm

We prove the first part of Theorem 12 with a poly-time algorithm. The head-cluster property implies that if two relations have different output attributes, they share no common non-output attributes. This way, we can cluster relations by output attributes. As shown, Algorithm 1 partitions all relations into $\{E_1, E_2, \cdots, E_k\}$ based on the connected components in $G^{\exists}_Q$. For the subset of relations in one connected component $E_i$, every relation is a dominant relation, i.e., shares the same output attributes. If one relation only contains output attributes $\mathbf{A}_i$

**Algorithm 1** SWP$(Q, D)$

1 $D' \leftarrow \emptyset$;
2 $(E_1, E_2, \cdots, E_k) \leftarrow$ connected components of $G^{\exists}_{Q}$;
3 $\mathbf{A}_1, \mathbf{A}_2, \cdots, \mathbf{A}_k \leftarrow$ output attributes of $E_1, E_2, \cdots, E_k$;
4 **foreach** $R_j \in \mathtt{rels}(Q)$ *with* $\mathtt{attr}(R_j) \subseteq \mathtt{head}(Q)$ **do**
5     $D' \leftarrow D' \cup \pi_{\mathtt{attr}(R_j)} Q(D)$;
6 **foreach** $i \in [k]$ **do**
7     Define $Q_i\left(\mathbf{A}_i\right) : -\{R_j(\mathbb{A}_j) : R_j \in E_i\}$;
8     **foreach** $t' \in \pi_{\mathbf{A}_i} Q(D)$ **do**
9         $D' \leftarrow D' \biguplus \mathtt{SWP}\left(Q_i, \{R_j : R_j \in E_i\}, t'\right)$;
10 **return** $D'$;

(line 4), it must appear alone as a singleton component, since we assume there is no duplicate relations in the input CQ. All tuples from such a relation that participate in any query results must be included by every witness to $Q(D)$. We next consider the remaining components containing at least two relations. In $E_i$, for each tuple $t' \in \pi_{\mathbf{A}_i} Q(D)$ in the projection of query results onto the output attributes $\mathbf{A}_i$, Algorithm 1 computes the smallest witness for $t'$ in sub-query $Q_i$ defined on relations in $E_i$. The disjoint union ($\biguplus$) of witnesses returned for all groups forms the final witness. On a CQ with head-cluster property, Algorithm 1 can be stated in a simpler way; see Algorithm 3 in Appendix D.

Algorithm 1 runs in polynomial time, as (i) $Q(D)$ can be computed in polynomial time; (ii) $|Q(D)|$ is polynomially large; and (iii) the primitive in line 9 only takes $O(1)$ time.

▶ **Lemma 13.** *For a CQ $Q$ with head-cluster property, Algorithm 1 finds a solution to* SWP$(Q, D)$ *for any database $D$ in polynomial time.*

**Remark 1.** SWP is poly-time solvable for any full CQ, since $\mathtt{attr}(R_i) \cap \mathtt{attr}(R_j) \subseteq \mathtt{head}(Q)$ holds for every pair of relations $R_i, R_j \in \mathtt{rels}(Q)$. Hence, the hardness of SWP comes from projection. On the other hand, SWP is also poly-time solvable for some non-full CQs, say $Q(A_1, A_2, A_3) : -R_1(A_1, A_2), R_2(A_2, A_3), R_3(A_1, A_3), R_4(A_1, B_1)$.

**Remark 2.** It is not always necessary to compute $Q(D)$ as Algorithm 1 does. We actually have much faster algorithms for some clases of CQs. If CQ $Q$ is full, SWP$(Q, D)$ is the set of *non-dangling* tuples in $D$, i.e., those participate in at least one query result of $Q(D)$. Furthermore, if $Q$ is an acyclic full CQ, non-dangling tuples can be identified in $O(|D|)$ time [45]. It is more expensive to identify non-dangling tuples for cyclic full CQs, for example, the PANDA algorithm [2] can identify non-dangling tuples for any full CQ in $O(N^w)$ time, where $w \geq 1$ is the sub-modular width of input query [2]. It is left as an interesting open question to compute SWP for CQs with head-cluster property more efficiently.

## 3.2 Hardness

We next prove the second part of Theorem 12 by showing the hardness for CQs without head-cluster property. Our hardness result is built on the NP-hardness of set cover [8]: Given a universe $\mathcal{U}$ of $n$ elements and a family $\mathcal{S}$ of subsets of $\mathcal{U}$, it asks to find a subfamily $\mathcal{C} \subseteq \mathcal{S}$ of sets whose union is $\mathcal{U}$ (called "cover"), while using the fewest sets. We start with $Q_{\mathsf{cover}}(A) : -R_1(A, B), R_2(B)$ and then extend to any CQ without head-cluster property.

▶ **Lemma 14.** *SWP is not poly-time solvable for $Q_{\mathsf{cover}}(A) : -R_1(A, B), R_2(B)$, unless* P = NP.

**Proof.** We show a reduction from set cover to $\texttt{SWP}$ for $Q_{\textsf{cover}}$. Given an arbitrary instance $(\mathcal{U}, \mathcal{S})$ of set cover, we construct a database $D$ for $Q_{\textsf{cover}}$ as follows. For each element $u \in \mathcal{U}$, we add a value $a_u$ to $\texttt{dom}(A)$; for each subset $S \in \mathcal{S}$, we add a value $b_S$ to $\texttt{dom}(B)$, and a tuple $(b_S)$ to $R_2$. Moreover, for each pair $(u, S) \in \mathcal{U} \times \mathcal{S}$ with $u \in S$, we add a tuple $(a_u, b_S)$ to $R_1$. Note that $Q_{\textsf{cover}}(D) = \mathcal{U}$. It is now to show that $(\mathcal{U}, \mathcal{S})$ has a cover of size $\leq k$ if and only if $\texttt{SWP}(Q_{\textsf{cover}}, D)$ has a solution $D'$ of size $\leq |\mathcal{U}| + k$. Then, if $\texttt{SWP}$ is poly-time solvable for $Q_{\textsf{cover}}$, set cover is also poly-time solvable, which is impossible unless $\texttt{P} = \texttt{NP}$. ◀

▶ **Lemma 15.** *For a CQ $Q$ without head-cluster property, $\texttt{SWP}$ is not poly-time solvable for $Q$, unless $\texttt{P} = \texttt{NP}$.*

It is always feasible to identify a pair of relations $R_i, R_j \in \texttt{rels}(Q)$ with attributes $A \in \texttt{head}(R_i) - \texttt{attr}(R_j)$ and $B \in \texttt{attr}(R_i) \cap \texttt{attr}(R_j) - \texttt{head}(Q)$. All remaining attributes contain a dummy value. Then, each relation in $Q$ degenerates to $R_1(A, B)$, or $R_2(B)$, or a dummy tuple. Our argument for $Q_{\textsf{cover}}$ can be applied here.

## 4 Dichotomy of Approximated SWP

As it is inherently difficult to compute $\texttt{SWP}$ exactly for general CQs, the next interesting question is to explore approximated solutions for $\texttt{SWP}$. In this section, we establish the following dichotomy for approximating $\texttt{SWP}$. All missing proofs are given in Appendix E.

▶ **Theorem 16.** *If a CQ $Q$ has head-domination property, $\texttt{SWP}$ is $O(1)$-approximable; otherwise, $\texttt{SWP}$ of input size $N$ is not $(1 - o(1)) \cdot \log N$-approximable, unless $\texttt{P} = \texttt{NP}$.*

### 4.1 A $O(1)$-Approximation Algorithm

Let's start by revisiting $Q_{\textsf{cover}}(A) : -R_1(A, B), R_2(B)$. Although $\texttt{SWP}$ is hard to compute exactly for $Q_{\textsf{cover}}$, it is easy to approximate $\texttt{SWP}(Q_{\textsf{cover}}, D)$ for arbitrary database $D$ within a factor of 2. Let $D^*$ be a solution to $\texttt{SWP}(Q_{\textsf{cover}}, D)$. We can simply construct an approximated solution $D'$ by picking a pair of tuples $(a, b) \in R_1, (b) \in R_2$ for every $a \in Q_{\textsf{cover}}(D)$, and show that $|D'| \leq 2 \cdot |Q_{\textsf{cover}}(D)| \leq 2 \cdot |D^*|$. This is actually not a violation to the inapproximability of set cover problem. If revisiting the proof of Lemma 14, $(\mathcal{U}, \mathcal{S})$ has a cover of size $\leq k$ if and only if $\texttt{SWP}(Q_{\textsf{cover}}, D)$ has a solution of size $\leq |\mathcal{U}| + k$. Due to the fact that $k \leq |\mathcal{U}|$, the inapproximability of set cover does not carry over to $\texttt{SWP}$ for $Q$. This observation can be generalized to all CQs with head-domination property.

As described in Algorithm 1, an approximated solution to $\texttt{SWP}(Q, D)$ for a CQ $Q$ with head-domination property consists of two parts. For every relation that only contains output attributes, Algorithm 1 includes all tuples that participate in at least one query result (line 4-5), which must be included by any witness for $Q(D)$. For the remaining relations, Algorithm 1 partitions them into groups based on the existential connectivity. Intuitively, every pair of relations across groups can only join via output attributes in their dominants. Recall that $\mathbf{A}_i$ denotes the set of output attributes appearing in relations from $E_i$. Then, for each group $E_i$, we consider each tuple $t' \in \pi_{\mathbf{A}_i} Q(D)$ and find the smallest witness for $t'$ in $Q_i$ defined by relations in $E_i$. The union of witnesses returned for all groups forms the final answer. As shown before, Algorithm 1 runs in polynomial time.

▶ **Lemma 17.** *For a CQ $Q$ with head-domination property, Algorithm 1 finds a $O(1)$-approximated solution to $\texttt{SWP}(Q, D)$ for any database $D$ in polynomial time.*

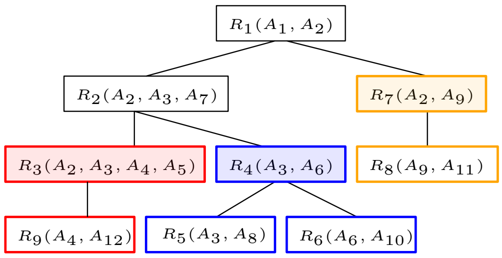


**Figure 3** A free-connex CQ $Q$ with $\texttt{head}(Q) = \{A_1, A_2, A_3, A_7\}$. A partition of relations containing non-output attributes is $\{\{R_3, R_9\}, \{R_4, R_5, R_6\}, \{R_7, R_8\}\}$, with dominant relations $R_3, R_4, R_7$ respectively.

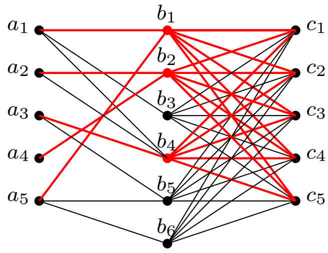


**Figure 4** A database $D$ for $Q_{\mathsf{matrix}}$ with an integral sub-database in red. Each vertex is a value in the attribute, and each edge is a tuple in $D$.

**Connection with Free-connex CQs.** We point out that every free-connex CQ has head-domination property, which is built on an important property as stated in Lemma 18.

▶ **Lemma 18** ([6]). *A free-connex CQ has a tree structure $\mathbb{T}$ such that (1) each node of $\mathbb{T}$ corresponds to $\texttt{attr}(R_i)$ or $\texttt{head}(R_i)$ for some relation $R_i \in \texttt{rels}(Q)$; and each relation $R_i \in \texttt{rels}(Q)$ corresponds to a node in $\mathbb{T}$; (2) for every attribute $A \in \texttt{attr}(Q)$, the set of nodes containing $A$ form a connected subtree of $\mathbb{T}$; (3) there is a connected subtree $\mathbb{T}_{con}$ of $\mathbb{T}$ including the root of $\mathbb{T}$, such that all attributes appearing in $\mathbb{T}$ is exactly $\texttt{head}(Q)$.*

▶ **Lemma 19.** *Every free-connex CQ has head-domination property.*

**Proof.** Given a tree structure $\mathbb{T}$ for a free-connex CQ $Q$ as descried by Lemma 18, every relation $R_i$ with $\texttt{attr}(R_i) - \texttt{head}(R_i) \neq \emptyset$ and all of its ancestors fully containing output attributes, is a dominant. Suppose $R_i$ corresponds to node $u$. As $\texttt{attr}(R_i) - \texttt{head}(R_i) \neq \emptyset$, $u \notin \mathbb{T}_{\mathsf{con}}$. Consider any descendant $v$ of $u$, that corresponds to a relation in $Q$, say $R_j$. Implied by the fact that $\mathbb{T}_{\mathsf{con}}$ is a connected subtree, $v \notin \mathbb{T}_{\mathsf{con}}$. Hence, $\texttt{head}(R_j) \subseteq \texttt{head}(R_i)$. For every relation $R_j \in \texttt{rels}(Q)$ with $\texttt{attr}(R_j) - \texttt{head}(Q) \neq \emptyset$, either itself is a dominant relation, or it has an ancestor as a dominant relation, since the root node fully contains output attributes. In our construction above, every relation $R_j \in \texttt{rels}(Q)$ with $\texttt{attr}(R_j) - \texttt{head}(Q) \neq \emptyset$ is associated with only one dominant relation. ◀

## 4.2 Logarithmic Inapproximability

Now, we turn to the class of CQs without head-domination property, and show their hardness by resorting to inapproximability of set cover [23, 20]: there is no poly-time algorithm for approximating set cover of input size $n$ within factor $(1 - o(1)) \cdot \log n$, unless $\texttt{P} = \texttt{NP}$. We identify two hardcore structures (Definition 23 and Definition 26), and prove that no poly-time algorithm can approximate $\texttt{SWP}$ for any CQ containing a hardcore within a logarithmic factor, unless $\texttt{P} = \texttt{NP}$. Lastly, we complete the proof of Theorem 16 by establishing the connection between the non-existence of a hardcore and head-domination property for CQs.

### 4.2.1 Free sequence

Let's start with the simplest acyclic but non-free-connex CQ $Q_{\mathsf{matrix}}(A, C) : -R_1(A, B), R_2(B, C)$. We show a reduction from set cover to $\texttt{SWP}$ for $Q_{\mathsf{matrix}}$ while preserving its logarithmic inapproximability. The essence of our proof is the notion of *integral* witness, such that for

$$(O1) \quad \min \sum_{b \in B} |C_b| \quad \text{s.t.} \bigcup_{b:(a,b)\in R_1} C_b = C, \forall a \in A \quad C_b \subseteq C, \forall b \in B$$

$$(O2) \quad \min \sum_{b \in B} x_b \quad \text{s.t.} \sum_{b:(a,b)\in R_1} x_b \geq 1, \forall a \in A \quad x_b \in \{0,1\}, \forall b \in B$$

**Figure 5** Optimization problems in the proof of Lemma 21.

any database $D$ where $R_2$ is a Cartesian product between $B$ and $C$, $\texttt{SWP}(Q_{\mathsf{matrix}}, D)$ always admits an integral witness.

▶ **Definition 20** (Integral Database). *For any database $D$ over $Q_{matrix}$ with $R_2 = (\pi_B R_2) \times (\pi_C R_2)$, a sub-database $D' \subseteq D$ is integral if $R_2' = (\pi_B R_2') \times (\pi_C R_2)$, where $R_2'$ is the subset of tuples of $R_2$ in $D'$.*

▶ **Lemma 21.** *For $Q_{matrix}(A, C) : -R_1(A, B), R_2(B, C)$ and any database $D$ where $R_2 = (\pi_B R_2) \times (\pi_C R_2)$, there is an integral solution to SWP$(Q_{matrix}, D)$, i.e., a smallest witness to $Q(D)$ that is also integral.*

**Proof.** Consider a database $D$ where $R_2 = (\pi_B R_2) \times (\pi_C R_2)$. We assume that there exists no dangling tuples in $R_1, R_2$, i.e., every tuple can join with some tuple from the other relation; otherwise, we simply remove these dangling tuples. With a slight abuse of notation, we denote $A = \pi_A R_1$, $B = \pi_B R_1 (= \pi_B R_2)$ and $C = \pi_C R_2$. We consider the optimization problem $O1$ in Figure 5. Intuitively, it asks to assign a subset of elements $C_b \subseteq C$ to each value $b \in B$, such that each $a$ is "connected" to all values in $C$ via tuples in $R_1, R_2$, while the total size of the assignment defined as $\sum_{b\in B} |C_b|$ is minimized. See Figure 4.

We rewrite the objective function as: $\sum_{b\in B} |C_b| = \sum_{c\in C} |\{b \in B : c \in C_b\}|$, where $\{b \in B : c \in C_b\}$ indicates the subset of values from $B$ to which $c$ is assigned. Together with the constraint that every $a \in A$ must be connected to $c$, we note that minimizing $|\{b \in B : c \in C_b\}|$ is equivalent to solving the optimization problem $(O2)$ in Figure 5. Let $x^*$ be the optimal solution of the program above, which only depends on the input relation $R_1$, and completely independent of the specific value $c$. Hence, we conclude that $\sum_{b\in B} |C_b| = |C| \cdot \sum_{b\in B} x_b^*$.

We next construct an integral sub-database $D'$ as follows. For each $b \in B$ with $x_b^* = 1$, we add tuples in $\{(b, c) \in R_2 : \forall c \in C\}$ to $D'$. For each $a \in A$, we pick an arbitrary $b \in B$ with $(a, b) \in R_1$ and $x_b^* = 1$, and add $(a, b)$ to $D'$. Any solution to $\texttt{SWP}(Q_{\mathsf{matrix}}, D)$ must contain at least $|A|$ tuples from $R_1$ and at least $|C| \cdot \sum_{b\in B} x_b^*$ tuples from $R_2$. Hence, $D'$ is an integral witness to $\texttt{SWP}(Q_{\mathsf{matrix}}, D)$. ◀

▶ **Lemma 22.** *There is no poly-time algorithm to approximate SWP for $Q_{matrix}$ within a factor of $(1 - o(1)) \cdot \log N$, unless P = NP.*

**Proof.** Consider an arbitrary instance of set cover $(\mathcal{U}, \mathcal{S})$, where $|\mathcal{U}| = n$ and $|\mathcal{S}| = n^c$ for some constant $c \geq 1$.[3] We construct a database $D$ for $Q_{\mathsf{matrix}}$ as follows. For each element $u \in \mathcal{U}$, we add a value $a_u$ to $\texttt{dom}(A)$ and $c_u$ to $\texttt{dom}(C)$; for each subset $S \in \mathcal{S}$, we add a value $b_S$ to $\texttt{dom}(B)$. Moreover, for each pair $(u, S) \in \mathcal{U} \times \mathcal{S}$ with $u \in S$, we add tuple $(a_u, b_S)$ to

[3] The inapproximability of set cover holds even when the size of the family of subsets is only polynomially large with respect to the size of the universe of elements [38, 24].

$R_1$. $R_2$ is a Cartesian product between $\texttt{dom}(B)$ and $\texttt{dom}(C)$. Every relation contains at most $n^{c+1}$ tuples. Hence $N \leq 2n^{c+1}$. Note that $Q_{\textsf{matrix}}(D)$ is the Cartesian product between $A$ and $C$. Implied by Lemma 21, it suffices to consider integral witness to $\texttt{SWP}(Q_{\textsf{matrix}}, D)$. We can show that $(\mathcal{U}, \mathcal{S})$ has a cover of size $\leq k$ if and only if $\texttt{SWP}(Q_{\textsf{matrix}}, D)$ has an integral witness of size $\leq n(k+1)$ (see Appendix E). Hence, if $\texttt{SWP}$ is $(1-o(1)) \cdot \log N$-approximable for $Q_{\textsf{matrix}}$, there is a poly-time algorithm that approximates set cover of input size $n$ within a $(1-o(1)) \cdot \log N$ factor, which is impossible unless $\texttt{P} = \texttt{NP}$. ◀

Now, we are ready to introduce the notion of *free sequence*:

▶ **Definition 23** (Free Sequence). *In a CQ $Q$, a free sequence is a sequence of attributes $P = \langle A_1, A_2, \cdots, A_k\rangle$ such that*[4]

- $A_1, A_k \in \texttt{head}(Q)$ *and* $A_2, A_3, \cdots, A_{k-1} \in \texttt{attr}(Q) - \texttt{head}(Q)$;
- *for every* $i \in [k-1]$, *there exists a relation* $R_j \in \texttt{rels}(Q)$ *such that* $A_i, A_{i+1} \in \texttt{attr}(R_j)$;
- *there exists no relation* $R_j \in \texttt{rels}(Q)$ *such that* $A_1, A_k \in R_j$.

▶ **Lemma 24.** *For a CQ $Q$ containing a free sequence, there is no poly-time algorithm that can approximate $\texttt{SWP}$ for $Q$ within a factor of $(1-o(1)) \cdot \log N$, unless $\texttt{P} = \texttt{NP}$.*

We give some high-level idea behind the proof of Lemma 24. Let $P = \langle A_1, A_2, \cdots, A_k\rangle$ be such a free sequence. In a reduction from set cover to $\texttt{SWP}$ for $Q$, $A_i$ simulates $A$ if $i = 1$, $B$ if $i \in [2 : k-1]$, and $C$ if $i = k$ as $Q_{\textsf{matrix}}$. Every remaining attribute only contains a dummy value $\{*\}$. A similar argument for $Q_{\textsf{matrix}}$ can be applied by defining integral sub-database of $D$ over $Q$, proving the existence of an integral witness to $\texttt{SWP}(Q, D)$, and establishing a correspondence between solutions to set cover and integral witness to $\texttt{SWP}(Q, D)$.

**Remark.** Any acyclic but non-free-connex CQ must have a free sequence [6]. Our characterization of $\texttt{SWP}$ for acyclic CQs coincides with the separation between free-connex and non-free-connex CQs. In short, $\texttt{SWP}$ is poly-time solvable or $O(1)$-approximable for free-connex CQs, while no poly-time algorithm can approximate $\texttt{SWP}$ for any acyclic but non-free-connex CQs within a factor of $(1-o(1)) \cdot \log N$, unless $\texttt{P} = \texttt{NP}$.

### 4.2.2 Nested Clique

Although free sequence suffices to capture the hardness of approximating $\texttt{SWP}$ for acyclic CQs, it is not enough for cyclic CQs. Let's start with the simplest cyclic CQ $Q_{\textsf{pyramid}}(A, B, C) :- R_1(A, B), R_2(A, C), R_3(B, C), R_4(A, F), R_5(B, F), R_6(C, F)$ that does not contain a free sequence, but $\texttt{SWP}$ is still difficult to approximate.

▶ **Lemma 25.** *There is no poly-time algorithm to approximate $\texttt{SWP}$ for $Q_{\textsf{pyramid}}$ within a factor of $(1-o(1)) \cdot \log N$, unless $\texttt{P} = \texttt{NP}$.*

**Proof.** Consider an instance $(\mathcal{U}, \mathcal{S})$ of set cover, with $\mathcal{U} = \{u_1, u_2, \cdots, u_n\}$ and $\mathcal{S} = \{S_1, S_2, \cdots, S_m\}$, where $m = n^c$ for some constant $c > 1$. We construct a database $D$ for $Q$ as follows. Let $\texttt{dom}(A) = \{a_1, a_2, \cdots, a_n\}$, $\texttt{dom}(B) = \{b_1, b_2, \cdots, b_n\}$, $\texttt{dom}(C) = \{c_1, c_2, \cdots, c_n\}$ and $\texttt{dom}(F) = \texttt{dom}(F^-) \times \texttt{dom}(F^+)$, where $\texttt{dom}(F^-) = \{f_1^-, f_2^-, \cdots, f_m^-\}$ and $\texttt{dom}(F^+) = \{f_1^+, f_2^+, \cdots, f_n^+\}$. Relations $R_1$, $R_2$, $R_3$ and $R_6$ are Cartesian products of their

[4] Free sequence is a slight generalized notion of free path [6] studied in the literature, which further requires that for any relation $R_j \in \texttt{rels}(Q)$, either $\texttt{attr}(R_j) \cap P = \emptyset$, or $|\texttt{attr}(R_j) \cap P| = 1$, or $\texttt{attr}(R_j) \cap P = \{A_i, A_{i+1}\}$ for some $i \in [k-1]$.

corresponding attributes. For each pair $(u_\ell, S_j) \in \mathcal{U} \times \mathcal{S}$ with $u_\ell \in S_j$, we add tuples of $\{(a_\ell, f_j^-)\} \times \texttt{dom}(F^+)$ to $R_4$. For each $i \in [n]$, we add tuples $\{b_i\} \times \texttt{dom}(F^-) \times \{f_i^+\}$ to $R_5$. It can be easily checked that the input size of $D$ is $O(n^{2c})$, hence $\log N = \Theta(\log n)$. $Q(D)$ is the Cartesian product between $A$,$B$ and $C$. Hence, every solution to $\texttt{SWP}(Q, D)$ includes all tuples in $R_1, R_2, R_3$. Below, we focus on $R_4, R_5, R_6$.

We observe that $D$ enjoys highly symmetric structure over $\texttt{dom}(B)$. More specifically, each value $b_i \in \texttt{dom}(B)$ induces a subquery $Q_i(A, C) = R_4(A, F_i) \bowtie R_5(F_i) \bowtie R_6(C, F_i)$, where $\texttt{dom}(F_i) = \texttt{dom}(F^-) \times \{f_i^+\}$, and a sub-database $D^i = \left\{R_4^i, R_5^i, R_6^i\right\}$, where $R_4^i = \{(a_\ell, f_j^-, f_i^+) : \forall \ell \in [n], j \in [m], u_\ell \in S_j\}$, $R_5^i = \texttt{dom}(F_i)$ and $R_6^i = \texttt{dom}(C) \times \texttt{dom}(F_i)$. It can be easily checked that $\texttt{SWP}(Q, D) = R_1 \uplus R_2 \uplus R_3 \uplus \left(\uplus_{i \in [n]} \texttt{SWP}(Q_i, D_i)\right)$. For any $i \in [n]$, computing $\texttt{SWP}(Q_i, D_i)$ is almost the same as $Q_{\mathsf{matrix}}$. Moreover, the solution to each $\texttt{SWP}(Q_i, D_i)$ shares the same structure, which is independent of the specific value $b_i \in \texttt{dom}(B)$. In a sub-database $D' \subseteq D$, let $R_i'$ be the corresponding sub-relation of $R_i$. A solution $D'$ to $\texttt{SWP}(Q, D)$ is *integral* if $R_4' = (\pi_{A,F^-} R_4') \times \texttt{dom}(F^+)$, $R_5' = \texttt{dom}(B) \times (\pi_{F^-} R_5') \times \texttt{dom}(F^+)$, and $R_6' = \texttt{dom}(C) \times (\pi_{F^-} R_6') \times \texttt{dom}(F^+)$. Implied by Lemma 21 and analysis above, there always exists an integral solution to $\texttt{SWP}(Q, D)$.

Moreover, $(\mathcal{U}, \mathcal{S})$ has a cover of size $\leq k$ if and only if $\texttt{SWP}(Q, D)$ has an integral witness of size $\leq 3n^2 + n(n + k + kn) = (k + 4)n^2 + kn$. If $\texttt{SWP}$ is $(1 - o(1)) \cdot \log N$-approximable for $Q$, then there is a poly-time algorithm that can approximate set cover instances of input size $n$ within a factor of $(1 - o(1)) \cdot \log n$, which is impossible unless $\texttt{P} = \texttt{NP}$. ◀

Now, we are ready to introduce the structure of *nested clique* and the *rename* procedure for capturing the hardness of cyclic CQs:

▶ **Definition 26** (Nested Clique). *In a CQ $Q$, a nested clique is a subset of attributes $P \subseteq \texttt{attr}(Q)$ such that*

- *for any pair of attributes $A, B \in P$, there is some $R_j \in \texttt{rels}(Q)$ with $A, B \in \texttt{attr}(R_j)$;*
- *$P \cap \texttt{head}(Q) \neq \emptyset$ and $P - \texttt{head}(Q) \neq \emptyset$;*
- *there is no relation $R_j \in \texttt{rels}(Q)$ with $P \cap \texttt{head}(Q) \subseteq \texttt{head}(R_j)$.*

▶ **Definition 27** (Rename). *Given the nonout-connectivity graph $H_Q$ of a CQ $Q$ with connected components $H_1, H_2, \cdots, H_k$, the rename procedure assigns one distinct attribute to all attributes in the same component. The resulted CQ $Q'$ contains the same output attributes as $Q$, and each $R_i \in \texttt{rels}(Q)$ defines a new relation $R_i' \in \texttt{rels}(Q')$ with $\texttt{attr}(R_i') = \texttt{head}(R_i) \cup \{F_j : \forall j \in [k], H_j \cap \texttt{attr}(R_i) \neq \emptyset\}$.*

In Figure 2, $Q_1(A_1, A_2, A_3) : -R_1(A_1, B_1), R_2(B_1, B_2), R_3(A_2, B_2, B_3), R_4(A_2, A_3, B_4), R_5(A_1, A_2)$ is renamed as $Q_1'(A_1, A_2, A_3) : -R_1(A_1, F_1), R_2(F_1), R_3(A_2, F_1), R_4(A_2, A_3, F_2), R_5(A_1, A_2)$.

▶ **Theorem 28.** *For a CQ $Q$, if its renamed query $Q'$ contains a nested clique, there is no poly-time algorithm that can approximate $\texttt{SWP}$ for $Q$ within a factor of $(1 - o(1)) \cdot \log N$, unless $\texttt{P} = \texttt{NP}$.*

## 4.3 Completeness

At last, we complete the proof of Theorem 16 by establishing the connection between the non-existence of hardcore structures and head-domination property:

▶ **Lemma 29.** *In a CQ $Q$, if there is neither a free sequence nor a nested clique in its renamed query, $Q$ has head-domination property.*

We have proved Lemma 29 for acyclic CQs, such that if $Q$ does not contain a free sequence, $Q$ has head-cluster property. It suffices to focus on cyclic CQs. We note that any cyclic CQ contains a *cycle* or *non-conformal clique* [9]. Our proof is based on a technical lemma that if every cycle or non-conformal clique contains only output attributes or only non-output attributes, $Q$ has head-domination property. We complete it by showing that if $Q$ does not contain a free sequence or nested clique in its renamed query, no cycle or non-conformal clique contains both output and non-output attributes.

# 5 Approximation Algorithms

We next explore possible approximation algorithms for CQs without head-domination property. All missing proofs are in Appendix F.

## 5.1 Baseline

A baseline for general CQs returns the union of smallest witness for every query result, which includes at most $\min\{N, |\mathtt{rels}(Q)| \cdot |Q(D)|\}$ tuples. Meanwhile, AGM bound [5] implies that at least $O(|Q(D)|^{1/\rho^*})$ tuples are needed to reproduce $|Q(D)|$ results, where $\rho^*$ is the fractional edge covering number of $Q$. Together, we obtain:

▶ **Theorem 30.** *SWP is $N^{1-1/\rho^*}$-approximable for any CQ $Q$, where $\rho^*$ is the fractional edge covering number of $Q$.*

This upper bound is polynomially larger than the logarithmic lower bound proved in Section 4. We next explore better approximation algorithms for some commonly-used CQs.

## 5.2 Star CQs

We look into one commonly-used class of CQs noted as *star* CQs:

$$Q_{\text{star}}(A_1, A_2, \cdots, A_m) : -R_1(A_1, B), R_2(A_2, B), \cdots, R_m(A_m, B).$$

Our approximation algorithm follows the greedy strategy developed for weighted set cover problem, where each element to be covered is a query result and each subset is a collection of tuples. Intuitively, the greedy strategy always picks a subset that minimizes the "price" for covering the remaining uncovered elements. The question boils down to specifying the universe $\mathcal{U}$ of elements, and the family $\mathcal{S}$ of subsets as well as their weights. However, naively taking every possible sub-collection of tuples from the input database as a subset, would generate an exponentially large $\mathcal{S}$, which leads to a greedy algorithm running in exponential time. Hence, it is critical to keep the size of $\mathcal{S}$ small. Let's start with $Q_{\mathsf{matrix}}$ $(m = 2)$.

**Greedy Algorithm for $Q_{\mathsf{matrix}}$.** Given $Q_{\mathsf{matrix}}$, a database $D$, and a subset of query results $\mathcal{C} \subseteq Q_{\mathsf{matrix}}(D)$, the price of a collections of tuples $(X, Y)$ for $X \subseteq \pi_A R_1$ and $Y \subseteq \pi_C R_2$ is defined as $f(\mathcal{C}, X, Y) = \frac{|X|+|Y|}{|(\mathcal{C} \cup (X \bowtie Y)) - \mathcal{C}|}$. To shrink the space of candidate subsets, a critical observation is that the subset with minimum price chosen by the greedy algorithm cannot be *divided* further, i.e., all tuples should have the same join value as captured by Lemma 31.

▶ **Lemma 31.** *Given $Q_{\mathsf{matrix}}$ and a database $D$, for two distinct values $b, b' \in \mathtt{dom}(B)$, and two pairs of subsets of tuples $(X_1, Y_1) \in 2^{\sigma_{B=b} R_1} \times 2^{\sigma_{B=b} R_2}, (X_2, Y_2) \in 2^{\sigma_{B=b'} R_1} \times 2^{\sigma_{B=b'} R_2}$, for arbitrary $\mathcal{C} \subseteq Q_{\mathsf{matrix}}(D)$, $\min\{f(\mathcal{C}, X_1, Y_1), f(\mathcal{C}, X_2, Y_2)\} \leq f(\mathcal{C}, X_1 \cup X_2, Y_1 \cup Y_2)$.*

**Algorithm 2** GREEDYSWP($Q_{\mathsf{matrix}}, D$)

1 $D' \leftarrow \emptyset$, $\mathcal{C} \leftarrow \emptyset$;
2 **foreach** $b_i \in \mathit{dom}(B)$ **do**
3 $\quad (X_i, Y_i) \leftarrow (\pi_A \sigma_{B=b_i} R_1, \pi_C \sigma_{B=b_i} R_2)$;
4 **while** $\mathcal{C} \neq Q_{\mathit{matrix}}(D)$ **do**
5 $\quad b_j \leftarrow \arg\min_{b_i \in \mathtt{dom}(B)} f(\mathcal{C}, X_i, Y_i)$;
6 $\quad D' \leftarrow D' \cup X_j \cup Y_j$, $\mathcal{C} \leftarrow \mathcal{C} \cup (X_j \times Y_j)$;
7 $\quad$ **foreach** $b_i \in \mathit{dom}(B)$ **do**
8 $\quad\quad (X_i, Y_i) \leftarrow \arg \min_{X \subseteq \pi_A \sigma_{B=b_i} R_1, Y \subseteq \pi_C \sigma_{B=b_i} R_2} f(\mathcal{C}, X, Y)$;
9 **return** $D'$;

Now, we are ready to present a greedy algorithm that runs in polynomial time. As shown in Algorithm 2, we always maintain a pair $(X_i, Y_i)$ for every value $b_i \in \mathtt{dom}(B)$, such that $(X_i, Y_i)$ minimize the function $f(\mathcal{C}, X, Y)$ for $X \subseteq \pi_A \sigma_{B=b_i} R_1$ and $Y \subseteq \pi_A \sigma_{B=b_i} R_2$. Initially when $\mathcal{C} = \emptyset$, it can be easily proved that when $\mathcal{C} = \emptyset$, the pair $(X, Y)$ that leads to minimum price must be $X = \pi_A \sigma_{B=b_i} R_1$ and $Y = \pi_C \sigma_{B=b_i} R_2$. The greedy algorithm always chooses the one pair with minimum price, pick all related tuples in this pair, and update the coverage $\mathcal{C}$. After this step, we also need to update the candidate pair $(X_i, Y_i)$ for each value $b_i \in \mathtt{dom}(B)$ and enter into the next iteration. We will stop until all query results are covered.

The remaining question is how to compute $(X_i, Y_i)$ with minimum price (line 8) efficiently. We mention a problem known as *densest subgraph in the bipartite graph* in the literature [29, 33], for which a poly-time algorithm based on max-flow has been proposed.

▶ **Definition 32** (Densest Subgraph in Bipartite Graph). *Given a bipartite graph* $(X, Y, E)$ *with* $E : X \times Y \to \{0, 1\}$*, it asks to find* $X' \subseteq X, Y' \subseteq Y$ *to maximize* $\frac{\sum_{x \in X', y \in Y'} E(x,y)}{|X'|+|Y'|}$.

Back to our problem, each value $b_i \in \mathtt{dom}(B)$ induces a bipartite graph with $X = \pi_A \sigma_{B=b_i} R_1$, $Y = \pi_C \sigma_{B=b_i} R_2$, and $E = X \times Y - \mathcal{C}$. This way, the densest subgraph in this bipartite graph corresponds to a subset of uncovered query results with minimum price.

The approximation ratio of our greedy algorithm follows the standard analysis of weighted set cover [44]. We next focus on the time complexity. The while-loop proceeds in at most $O(|Q(D)|)$ iterations, since $|\mathcal{C}|$ increases by at least 1 in every iteration. It takes $O(N)$ time to find the pair with minimum price, since there are $O(N)$ distinct values in $\mathtt{dom}(B)$. Moreover, it takes polynomial time to update $(X_i, Y_i)$ for each $b_i \in \mathtt{dom}(B)$. Overall, Algorithm 2 runs in polynomial time in terms of $N$.

**Extensions.** Our algorithm for $Q_{\mathsf{matrix}}$ can be extended to star CQs, and further to all CQs with only one non-output attribute. Similar property as Lemma 31 also holds. Our greedy algorithm needs a generalized primitive, noted as *densest subgraph in hypergraph*[5] [31] for finding the "set" with the smallest price. Following the similar analysis, we obtain:

▶ **Theorem 33.** *For any CQ Q with* $|\mathit{attr}(Q) - \mathit{head}(Q)| = 1$*, SWP is* $O(\log N)$*-approximable.*

[5] Given a hypergraph $H = (V, E)$ for $E \subseteq 2^V$, it asks to find a subset of nodes $S \subseteq V$ such that the ratio $|\{e \in E : e \subseteq S\}|/|S|$ is maximized.

### 5.3 Line CQs

We turn to another commonly-used class of CQs noted as *line* CQs:

$$Q_{\text{line}}(A_1, A_m) : -R_1(A_1, A_2), R_2(A_2, A_3), \cdots, R_m(A_m, A_{m+1})$$

with $m \geq 3$. We surprisingly find that SWP for line CQs is closely related to *directed steiner forest* (DSF) problem [21, 13, 25, 19] in the network design: Given an edge-weighed directed graph $G = (V, E)$ of $|V| = n$ and a set of $k$ demand pairs $\{(s_i, t_i) \in V \times V : i \in [k]\}$ and the goal is to find a subgraph $G'$ of $G$ with minimum weight such that there is a path in $G'$ from $s_i$ to $t_i$ for every $i \in [k]$. Observe that $\texttt{SWP}(Q_{\text{line}}, D)$ is a special case of DSF. This way, all existing algorithm proposed for DSF can be applied to SWP for $Q_{\text{line}}$. The best approximation ratio achieved is $O(\min\{k^{\frac{1}{2}+o(1)}, n^{0.5778}\})$ [25, 13, 1]. Combining the baseline in Section 5.1 (with $\rho^* = 2$ for line CQs) and existing algorithms for DSF, we obtain:

▶ **Theorem 34.** *For $Q_{line}$ and any database $D$ of input size $N$, there is a poly-time algorithm that can approximate $\texttt{SWP}(Q_{line}, D)$ within a factor of $O(\min\{|Q_{line}(D)|^{\frac{1}{2}+o(1)}, \texttt{dom}^{0.5778}, N^{\frac{1}{2}}\})$, where dom is the number of values that participates in at least one full join result.*

There is a large body of works investigating the lower bounds of DSF; and we refer interested readers to [19] for details. We mention a polynomial lower bound $\Omega(k^{1/4-o(1)})$ for DSF, but it cannot be applied to SWP, since SWP is a special case with its complexity measured by the number of edges in the graph. Instead, we built a reduction from *label cover* to SWP for $Q_{\text{line}}$ directly (which is an adaption of reduction proposed in [21, 15]) and prove the following:

▶ **Theorem 35.** *No poly-time algorithm approximates SWP for $Q_{line}$ within $\Omega(2^{(\log N)^{1-\epsilon}})$ factor for any constant $\epsilon > 0$, unless P = NP.*

## 6 Related Work

**Factorized database [39]** Factorized database studies a nested representations of query results that can be exponentially more succinct than flat query result, which has the same goal as SWP. They built a tree-based representations by exploiting the distributivity of product over union and commutativity of product and union. This notion is quite different from SWP. They measure the regularity of factorizations by readability, the minimum over all its representations of the maximum number of occurrences of any tuple in that representation, while SWP measures the size of witness.

**Data Synopses [17, 41]** In approximate query processing, people studied a lossy, compact synopsis of the data such that queries can be efficiently and approximately executed against the synopsis rather than the entire dataset, such as random samples, sketches, histograms and wavelets. These data synopses differ in terms of what class of queries can be approximately answered, space usage, accuracy etc. In computational geometry, a corset is a small set of points that approximate the shape of a larger point set, such as for shape-fitting, density estimation, high-dimensional vectors or points, clustering, graphs, Fourier transforms etc. SWP can also be viewed as data synopsis, since it also selects a representative subset of tuples. But, it is more query-dependent, as different CQs over the same database (such as $Q_1(x) : -R_1(x, y), R_2(y, z)$ and $Q_2(x, y, z) : -R_1(x, y), R_2(y, z)$) can have dramatically different SWP. Furthermore, all query results must be preserved by SWP, but this is never guaranteed in other data synopses. There is no space-accuracy tradeoff in SWP as well.

**Related to Other Problems in Database Theory.** The SWP problem also outputs the smallest provenance that can explain why all queries results are correct. This notion of why-provenance has been extensively investigated in the literature [11, 30, 3], but not from the perspective of minimizing the size of witness. The `SWP` problem is also related to the resilience problem [26, 27, 42] ( Definition 39), which intuitively finds the smallest number of tuples to remove so that the query answer turns into false. Our `SWP` problem essentially finds the maximum number of tuples to remove while the query answer does not change. We can observe clear connection here, but their solutions do not imply anything to each other.

## 7 Conclusion

In this paper, we study the data complexity of `SWP` problem for CQs without self-joins. There are several interesting problems left:

1. *Approximating `SWP` for CQs without head-domination property.* So far, the approximation of `SWP` is well-understood only on some specific class of CQs without head-domination property. For remaining CQs, both upper and lower bounds remain to be improved, which may lead to fundamental breakthrough for other related problems, such as DSF.
2. *`SWP` for CQs with self-joins.* It becomes much more challenging when self-joins exists, as one tuple appears in multiple logical copies of input relation. Similar observation has been made for the related resilience problem [27, 35]
3. *Relaxing the number of query results witnessed.* It is also possible to explore approximation on the number of query results that can be witnessed. Here, `SWP` is naturally related to the *partial set cover* problem, for which many approximation algorithms [28] have been studied.

## References


1 Amir Abboud and Greg Bodwin. Reachability preservers: New extremal bounds and approximation algorithms. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1865–1883. SIAM, 2018.

2 Mahmoud Abo Khamis, Hung Q Ngo, and Atri Rudra. Faq: questions asked frequently. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 13–28, 2016.

3 Yael Amsterdamer, Daniel Deutch, and Val Tannen. Provenance for aggregate queries. In *PODS*, pages 153–164, 2011.

4 Albert Atserias, Martin Grohe, and Dániel Marx. Size bounds and query plans for relational joins. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '08, pages 739–748, 2008.

5 Albert Atserias, Martin Grohe, and Dániel Marx. Size bounds and query plans for relational joins. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 739–748. IEEE, 2008.

6 Guillaume Bagan, Arnaud Durand, and Etienne Grandjean. On acyclic conjunctive queries and constant delay enumeration. In *International Workshop on Computer Science Logic*, pages 208–222. Springer, 2007.

7 C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *JACM*, 30(3):479–513, 1983.

8 KORTE Bernhard and JENS Vygen. Combinatorial optimization: Theory and algorithms. *Springer, Third Edition, 2005.*, 2008.

9 Johann Brault-Baron. Hypergraph acyclicity revisited. *ACM Computing Surveys (CSUR)*, 49(3):1–26, 2016.

10 Peter Buneman, Sanjeev Khanna, and Wang-Chiew Tan. On propagation of deletions and annotations through views. In *Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 150–158, 2002.

11 Peter Buneman, Sanjeev Khanna, and Tan Wang-Chiew. Why and where: A characterization of data provenance. In *International conference on database theory*, pages 316–330. Springer, 2001.

12 Surajit Chaudhuri, Rajeev Motwani, and Vivek Narasayya. On random sampling over joins. *ACM SIGMOD Record*, 28(2):263–274, 1999.

13 Chandra Chekuri, Guy Even, Anupam Gupta, and Danny Segev. Set connectivity problems in undirected graphs and the directed steiner network problem. *ACM Transactions on Algorithms (TALG)*, 7(2):1–17, 2011.

14 Yu Chen and Ke Yi. Random sampling and size estimation over cyclic joins. In *23rd International Conference on Database Theory (ICDT 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.

15 Rajesh Chitnis, Andreas Emil Feldmann, and Pasin Manurangsi. Parameterized approximation algorithms for bidirected steiner network problems. *ACM Trans. Algorithms*, 17(2), apr 2021. doi:10.1145/3447584.

16 Gao Cong, Wenfei Fan, and Floris Geerts. Annotation propagation revisited for key preserving views. In *Proceedings of the 15th ACM international conference on Information and knowledge management*, pages 632–641, 2006.

17 Graham Cormode, Minos Garofalakis, Peter J Haas, Chris Jermaine, et al. Synopses for massive data: Samples, histograms, wavelets, sketches. *Foundations and Trends® in Databases*, 4(1–3):1–294, 2011.

18 Graham Cormode and Ke Yi. *Small summaries for big data.* Cambridge University Press, 2020.

19 Irit Dinur and Pasin Manurangsi. Eth-hardness of approximating 2-csps and directed steiner network. In Anna R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference,*

*ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94 of *LIPIcs*, pages 36:1–36:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ITCS.2018.36.

20 Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 624–633, 2014.

21 Yevgeniy Dodis and Sanjeev Khanna. Design networks with bounded pairwise distance. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 750–759, 1999.

22 R. Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. *JACM*, 30(3):514–550, 1983.

23 Uriel Feige. A threshold of ln n for approximating set cover. *Journal of the ACM (JACM)*, 45(4):634–652, 1998.

24 Uriel Feige. A threshold of ln n for approximating set cover. *J. ACM*, 45(4):634–652, jul 1998. doi:10.1145/285055.285059.

25 Moran Feldman, Guy Kortsarz, and Zeev Nutov. Improved approximation algorithms for directed steiner forest. *Journal of Computer and System Sciences*, 78(1):279–292, 2012.

26 Cibele Freire, Wolfgang Gatterbauer, Neil Immerman, and Alexandra Meliou. The complexity of resilience and responsibility for self-join-free conjunctive queries. *PVLDB*, 9(3):180–191, 2015.

27 Cibele Freire, Wolfgang Gatterbauer, Neil Immerman, and Alexandra Meliou. New results for the complexity of resilience for binary conjunctive queries with self-joins. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 271–284, 2020.

28 Rajiv Gandhi, Samir Khuller, and Aravind Srinivasan. Approximation algorithms for partial covering problems. *Journal of Algorithms*, 53(1):55–84, 2004.

29 Andrew V Goldberg. Finding a maximum density subgraph. 1984.

30 Todd J Green, Grigoris Karvounarakis, and Val Tannen. Provenance semirings. In *PODS*, pages 31–40, 2007.

31 Shuguang Hu, Xiaowei Wu, and TH Hubert Chan. Maintaining densest subsets efficiently in evolving hypergraphs. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 929–938, 2017.

32 Xiao Hu, Shouzhuo Sun, Shweta Patwa, Debmalya Panigrahi, and Sudeepa Roy. Aggregated deletion propagation for counting conjunctive query answers. *Proceedings of the VLDB Endowment*, 14(2):228–240, 2020.

33 Samir Khuller and Barna Saha. On finding dense subgraphs. In *International colloquium on automata, languages, and programming*, pages 597–608. Springer, 2009.

34 Benny Kimelfeld, Jan Vondrák, and Ryan Williams. Maximizing conjunctive views in deletion propagation. In *Proceedings of the 30th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2011, June 12-16, 2011, Athens, Greece*, pages 187–198, 2011.

35 Benny Kimelfeld, Jan Vondrák, and Ryan Williams. Maximizing conjunctive views in deletion propagation. *ACM Transactions on Database Systems (TODS)*, 37(4):1–37, 2012.

36 Benny Kimelfeld, Jan Vondrák, and David P Woodruff. Multi-tuple deletion propagation: Approximations and complexity. *Proceedings of the VLDB Endowment*, 6(13):1558–1569, 2013.

37 Zhengjie Miao, Sudeepa Roy, and Jun Yang. Explaining wrong queries using small examples. In *Proceedings of the 2019 International Conference on Management of Data*, pages 503–520, 2019.

38 Dana Moshkovitz. The projection games conjecture and the np-hardness of ln n-approximating set-cover. In *APPROX-RANDOM*, pages 276–287. Springer, 2012.

39 Dan Olteanu and Maximilian Schleich. Factorized databases. *ACM SIGMOD Record*, 45(2):5–16, 2016.

**40** John Paparrizos, Chunwei Liu, Bruno Barbarioli, Johnny Hwang, Ikraduya Edian, Aaron J Elmore, Michael J Franklin, and Sanjay Krishnan. Vergedb: A database for iot analytics on edge devices. In *CIDR*, 2021.

**41** Jeff M Phillips. Coresets and sketches. In *Handbook of discrete and computational geometry*, pages 1269–1288. Chapman and Hall/CRC, 2017.

**42** Biao Qin, Deying Li, and Chunlai Zhou. The resilience of conjunctive queries with inequalities. *Information Sciences*, 613:982–1002, 2022.

**43** Moshe Y Vardi. The complexity of relational query languages. In *STOC*, pages 137–146, 1982.

**44** Vijay V Vazirani. *Approximation algorithms*, volume 1. Springer, 2001.

**45** Mihalis Yannakakis. Algorithms for acyclic database schemes. In *VLDB*, volume 81, pages 82–94, 1981.

**46** Steffen Zeuch, Eleni Tzirita Zacharatou, Shuhao Zhang, Xenofon Chatziliadis, Ankit Chaudhary, Bonaventura Del Monte, Dimitrios Giouroukis, Philipp M Grulich, Ariane Ziehn, and Volker Mark. Nebulastream: Complex analytics beyond the cloud. *Open Journal of Internet Of Things (OJIOT)*, 6(1):66–81, 2020.

**47** Zhuoyue Zhao, Robert Christensen, Feifei Li, Xiao Hu, and Ke Yi. Random sampling over joins revisited. In *Proceedings of the 2018 International Conference on Management of Data*, pages 1525–1539, 2018.

# A Missing Examples

$R_1$

| A | B |
|---|---|
| a1 | b1 |
| a2 | b2 |
| a3 | b2 |

$R_2$

| B | C |
|---|---|
| b1 | c1 |
| b2 | c3 |
| b3 | c2 |
| b3 | c3 |

$R_3$

| C | F |
|---|---|
| c1 | f1 |
| c2 | f3 |
| c3 | f3 |

$R_4$

| C | H |
|---|---|
| c1 | h1 |
| c2 | h1 |
| c3 | h1 |
| c3 | h2 |

$Q(D)$

| A | C | F |
|---|---|---|
| a1 | c1 | f1 |
| a2 | c3 | f3 |
| a3 | c3 | f3 |

**Figure 6** An example of database schema $\mathbb{R} = \{R_1, R_2, R_3, R_4\}$ (with $\mathbb{A} = \{A, B, C, F, H\}$, $\texttt{attr}(R_1) = \{A, B\}$, $\texttt{attr}(R_2) = \{B, C\}$, $\texttt{attr}(R_3) = \{C, F\}$ and $\texttt{attr}(R_4) = \{C, H\}$), a database $D$, and the result of CQ $Q(A, C, F) : -R_1(A, B), R_2(B, C), R_3(C, F), R_4(C, H)$ over $D$. $D' = \{(a_1, b_1), (b_1, c_1), (c_1, f_1), (c_1, h_1)\}$ is the solution to $\texttt{SWP}(Q, D, \langle a_1, c_1, f_1\rangle)$. $D'$ together with tuples $\{(a_2, b_2), (a_3, b_2), (b_2, c_3), (c_3, f_3), (c_3, h_2)\}$ is the solution to $\texttt{SWP}(Q, D)$.

# B Discussion on Our Improvement over Naive Solutions

At last, let's see what can be benefited from our algorithmic results such as in Examples 1, 2 and 3. Below, we compare $\texttt{SWP}(Q, D)$ with $|D|$ and $|Q(D)|$.

For a CQ $Q$ with head-cluster property, let $Q_1$, $Q_2$, $\cdots$, $Q_k$ be the subqueries defined on relations partitioned by the output attributes. For a CQ $Q$ with head-domination property, let $Q_1$, $Q_2$,$\cdots$, $Q_k$ be the subqueries defined by the connected components of $G^{\exists}_Q$. Consider an arbitrary database $D$. Let $D'_i$ be the projection of $Q(D)$ onto output attributes of $Q_i$. Let $D' = \sum_{i=1}^{k} D'_i$. We can prove for any $D$: (i) $|D'| \leq |Q(D')| = |Q(D)|$; (ii) $\texttt{SWP}(Q, D) = |D'|$, if $Q$ has head-cluster property; (iii) $\texttt{ASWP}(Q, D) \leq |\texttt{rels}(Q)| \cdot |D'|$ if $Q$ has head-domination property. In general, $|Q(D)|$ can be as large as $\Theta\left(|D'|^{\rho}\right)$, where $\rho$ is fractional edge covering number of the residual CQ by removing all non-output attributes from $Q$, and $|D|$ can also be polynomially larger than $|D'|$, so $\texttt{(A)SWP}(Q, D)$ could be polynomially smaller than both $|D|$ and $|Q(D)|$, while never be larger than them by a query-dependent constant.

For arbitrary CQ $Q$ with database $D$, let $D'$ be the projection of $Q(D)$ onto output attributes for each relation. We have proved that $\texttt{ASWP}(Q, D)$ is always smaller than $|D|$ and no larger than $|Q(D')|$ by a constant factor. The improvement would heavily depends on the specific database. As noted, in some applications (such as Example 2 and 3), a naive solution of size $|Q(D)|$ does not exist.

# C Deletion Propagation

In the literature, the classic deletion propagation ($\texttt{DP}$) of CQs without self-joins show resemblance with $\texttt{SWP}$ studied in this work. Here, we perform a comprehensive comparison between them with respect to problem formulation, complexity results and techniques.

▶ **Definition 36** (Deletion Propagation with Side Effect [35, 36]). *Given a CQ $Q$, a database $D$ and a subset of query results $J \subseteq Q(D)$, the problem asks to find a subset of tuples $D' \subseteq D$ such that $J \cap Q(D - D') = \emptyset$, and there exists no subset $D'' \subseteq D$ with $J \cap Q(D - D'') = \emptyset$ and $|Q(D - D'')| > |Q(D - D')|$.*

$\texttt{DP}$ with side effects is different from $\texttt{SWP}$ in terms of objective function, where $\texttt{DP}$ focuses on the number of remaining query results, and $\texttt{SWP}$ focuses on the number of input tuples.

▶ **Definition 37** (Deletion Propagation with Source Side Effect [10, 16])**.** *Given a CQ $Q$, a database $D$ and a subset of query results $J \subseteq Q(D)$, the problem asks to find a subset of tuples $D' \subseteq D$ such that all query results $J$ just disappear after removing all tuples in $D'$, i.e., $J \cap Q(D - D') = \emptyset$, while there exists no subset $D'' \subseteq D$ with $J \cap Q(D - D'') = \emptyset$ and $|D''| < |D'|$.*

We note that if $J = \emptyset$, i.e., no query result is desired to be removed, the trivial solution $D' = \emptyset$ is optimal, and if $J = Q(D)$, `DP` with source side effect just degenerates to the resilience problem (see Definition 39).

▶ **Definition 38** (Aggregated Deletion Propagation (`ADP`) [32])**.** *Given a CQ $Q$, a database $D$ and a parameter $1 \le k \le |Q(D)|$, the problem asks to find a subset of tuples $D' \subseteq D$ such that at least $k$ query results disappear, i.e., $|Q(D - D')| \le |Q(D)| - k$, while there exists no subset $D'' \subseteq D$ with $|Q(D - D'')| \le |Q(D)| - k$ and $|D''| < |D'|$.*

Similarly, we note that if $k = 0$, i.e., no query result is desired to be removed, the trivial solution $D' = \emptyset$ is optimal, and if $k = |Q(D)|$, `ADP` degenerates to the resilience problem. We next focus on the difference between resilience and `SWP`:

▶ **Definition 39** (Resilience [26, 27])**.** *Given a Boolean CQ $Q$ and a database $D$ with $Q(D)$ is true, the problem asks to find a subset of tuples $D' \subseteq D$ such that $Q(D - D')$ is false, while there exists no subset $D'' \subseteq D$ with $Q(D - D'')$ as false and $|D''| < |D'|$.*

These two problems have completely different constraints on $D'$, such that $Q(D') = Q(D)$ is always required by `SWP`, but $Q(D - D') = \emptyset$ is required by the resilience problem. Note that $Q(D - D') \neq Q(D) - Q(D')$! Hence, $Q(D - D') = \emptyset$ does not imply anything for $Q(D')$, so there is no connection between solutions to the resilience problem and solutions to `SWP`.

**Complexity.** As head domination property also precisely captures the hardness of deletion propagation with side effect, we next focus on the difference between our paper and [35]:

- Poly-time solvability (assume $\mathtt{P} \neq \mathtt{NP}$):
  - `SWP` is poly-time solvable iff $Q$ has head-cluster property;
  - `DP` is poly-time solvable iff $Q$ has head-domination property;
- Approximability:
  - `SWP` is $N^{1-1/\rho}$-approximable for any $Q$;
  - `DP` is $1/\min\{|\mathtt{attr}(Q)|, |\mathtt{rels}(Q)|\}$-approximable for any $Q$;
- Inapproximability for CQ without head-domination property:
  - `SWP` is NP-hard to approximate within a factor of $(1 - o(1)) \cdot \log N$;
  - `SWP` is NP-hard to approximate for line CQs within a factor of $O\left(2^{(\log N)^{1-\epsilon}}\right)$ for any constant $\epsilon > 0$;
  - `DP` is NP-hard to approximate within a factor of $\alpha$, for some constant $\alpha \in (0, 1)$, which may vary over different CQs;
- Greedy Algorithms:
  - `SWP` is $O(\log N)$-approximable for star CQs.
  - `SWP` is $O(\min\{|Q_{\text{line}}(D)|^{\frac{1}{2}+o(1)}, \mathtt{dom}^{0.5778}, N^{\frac{1}{2}}\})$-approximable for line CQs (see Theorem 34).
  - `DP` is $\frac{1}{2}$-approximable for star CQs.

**Techniques.** We are not aware of any similar techniques used for algorithms and lower bounds between this paper and [35].

**Algorithm 3** EasySWP$(Q, D)$

$D' \leftarrow \emptyset$;
$\{E_1, \cdots, E_k\} \leftarrow$ a partition of $\texttt{rels}(Q)$ by output attributes;
$\mathbf{A}_1, \mathbf{A}_2, \cdots, \mathbf{A}_k \leftarrow$ output attributes of $E_1, E_2, \cdots, E_k$;
**foreach** $i \in [k]$ **do**
  Define $Q_i(\mathbf{A}_i) : -\{R_j(\mathbb{A}_j) : j \in E_i\}$;
  **foreach** $t' \in \pi_{\mathbf{A}_i} Q(D)$ **do**
    $D' \leftarrow D' \biguplus \texttt{SWP}(Q_i, \{R_j : R_j \in E_i\}, t')$;
**return** $D'$;

## D Missing Materials in Section 3

**Proof of Lemma 13.** We prove that for any CQ $Q$ with head-cluster property and an arbitrary database $D$, Algorithm 1 returns the solution to $\texttt{SWP}(Q, D)$. Together with the fact that Algorithm 1 runs in polynomial time, we finish the proof for Lemma 13.

Let $D'$ be the solution returned by Algorithm 1. Let $D'_i \subseteq D'$ be the set of tuples from relations in $E_i$. We show that $D'$ is a witness for $Q(D)$, i.e., $Q(D') = Q(D)$. *Direction $\subseteq$.* As $Q$ is a monotone query, $Q(D') \subseteq Q(D)$ always holds for any sub-database $D' \subseteq D$. *Direction $\supseteq$.* Consider an arbitrary query result $t \in Q(D)$. Let $D'_i(t)$ denote the group of tuples returned by $\texttt{SWP}(Q_i, \{R_j : R_j \in E_i\}, \pi_{\mathbf{A}_i} t)$. We note that $t \in \pi_{\mathbf{A}} \Join_{i \in [k]} D'_i(t)$, since

- every tuple has the same value $\pi_A t$ over any output attribute $A$, if it contains attribute $A$;
- there is no non-output attribute to join for tuples across groups;
- tuples inside each group can be joined by non-output attribute; (implied by the correctness of $\texttt{SWP}$ for a single query result)

Hence, $t \in Q(D')$. Together, $Q(D') \supseteq Q(D)$.

We next show that there exists no other database $D'' \subseteq D$ such that $Q(D'') = Q(D)$ and $|D''| < |D'|$. By contradiction, assume there is some $D'' \subseteq D$ such that $Q(D'') = Q(D)$ and $|D''| < |D'|$. Let $D''_i \subseteq D''$ denote the set of tuples from relations in $E_i$. As $|D''| < |D'|$, there must exist some $i \in [k]$ such that $|D''_i| < |D'_i|$, i.e., $D'_i - D''_i \neq \emptyset$. In Algorithm 1, we can rewrite $D'_i$ as follows:

$$D'_i = \biguplus_{t' \in \pi_{\mathbf{A}_i} Q(D)} \texttt{SWP}(Q_i, \{R_i : R_i \in E_i\}, t'),$$

where $\biguplus$ denotes the disjoint union. As $D'_i - D''_i \neq \emptyset$, there must exist some $t' \in \pi_{\mathbf{A}_i} Q(D)$ such that $t' \notin Q_i(D''_i)$, i.e., $t'$ cannot be witnessed by $D''_i$. Correspondingly, there must exist some query result $t$ with $\pi_{\mathbf{A}_i} t = t'$ such that $t \notin Q(D'')$, i.e. $t$ cannot be witnessed by $D''$, contradicting the fact that $Q(D'') = Q(D)$. Hence, no such $D''$ exists and $D'$ has the smallest size. ◀

**Proof of Lemma 15.** Consider such a CQ $Q$ with a desired pair of relations $R_i, R_j \in \texttt{rels}(Q)$. We next show a reduction from $Q_{\mathsf{cover}}$ to $Q$. Given an arbitrary database $D_{\mathsf{cover}}$ over $Q_{\mathsf{cover}}$, we construct a database $D$ over $Q$ as follows. First, it is always feasible to identify attribute $A' \in \texttt{head}(R_i) - \texttt{attr}(R_j)$ and attribute $B' \in \texttt{attr}(R_i) \cap \texttt{attr}(R_j) - \texttt{head}(Q)$. We set $\texttt{dom}(A') = \texttt{dom}(A)$, $\texttt{dom}(B') = \texttt{dom}(B)$, and remaining attributes with a dummy value $\{*\}$. Each relation in $Q$ degenerates to $R_1(A, B)$, or $R_2(B)$, or a dummy tuple $\{*, *, \cdots, *\}$. It can be easily checked that there is a one-to-one correspondence between solutions to $\texttt{SWP}(Q, D)$ and solutions to $\texttt{SWP}(Q_{\mathsf{cover}}, D_{\mathsf{cover}})$. Thus, if $\texttt{SWP}$ is poly-time solvable for $Q$, then $\texttt{SWP}$ is also poly-time solvable for $Q_{\mathsf{cover}}$, leading to a contradiction of Lemma 14. ◀

**Faster Algorithm of Computing $\texttt{SWP}$ for Triangle CQ.** Let's take the triangle full CQ $Q_\triangle(A,B,C) = R_1(A,B) \bowtie R_2(B,C) \bowtie R_3(A,C)$ as example. For an arbitrary database $D$, $\texttt{SWP}(Q_\triangle, D)$ can be computed in $O\left(N^{\frac{2\omega}{\omega+1}}\right)$ time, where $\omega$ is the exponent of fast matrix multiplication algorithm. The algorithm is adapted from triangle detection. For each input relation, we proceed with the following procedure. Let's take $R_1$ as example.

This algorithm uses a parameter $\Delta$, whose value will be determined later. A value $a \in \texttt{dom}(A)$ is *heavy* if it appears in more than $\Delta$ tuples in $R_1(A,B)$ or $R_3(A,C)$, and *light* otherwise. A value $b \in \texttt{dom}(B)$ is *heavy* if it appears in more than $\Delta$ tuples in $R_1(A,B)$ or $R_2(B,C)$, and *light* otherwise. A value $c \in \texttt{dom}(C)$ is *heavy* if it appears in more than $\Delta$ tuples in $R_2(B,C)$ or $R_3(A,C)$, and *light* otherwise. We take the union of results of the following queries:

- $\pi_{A,B}\left(R_1(A,B^{\text{light}}) \bowtie R_2(B^{\text{light}},C) \ltimes R_3(A,C)\right)$: We first materialize $R_1(A,B^{\text{light}}) \bowtie R_2(B,C)$, then compute the semi-join between the intermediate join results and $R_3$, and the projection of the semi-join results onto attributes $A,B$. It takes $O(N\cdot\Delta)$ time.
- $\pi_{A,B}\left(R_1(A^{\text{light}},B) \bowtie R_3(A^{\text{light}},C) \ltimes R_2(B,C)\right)$: We apply similar procedure as above. It also takes $O(N\cdot\Delta)$ time.
- $R_1(A,B) \ltimes \left(R_2(B,C^{\text{light}}) \bowtie R_3(A,C^{\text{light}})\right)$: We materialize $R_2(B, C^{\text{light}}) \bowtie R_3(A,C^{\text{light}})$, then compute the semi-join between $R_1$ and the intermediate join results. It takes $O(N\cdot\Delta)$ time.
- $R_1(A,B) \bowtie \left(R_2(B^{\text{heavy}},C^{\text{heavy}}) \bowtie R_3(A^{\text{heavy}},C^{\text{heavy}})\right)$: We materialize $R_2(B^{\text{heavy}},C^{\text{light}}) \bowtie R_3(A^{\text{heavy}},C^{\text{light}})$ using fast matrix multiplication, then compute the semi-join between $R_1$ and the intermediate join results. It also takes $O\left((\frac{N}{\Delta})^\omega\right)$ time.

By setting $\Delta = N^{\frac{\omega-1}{\omega+1}}$, we obtain the overall run-time as $O(N^{\frac{2\omega}{\omega+1}})$.

## E Missing Proofs in Section 4

**Missing Details in the Proof of Lemma 21.** Here, we show that $(\mathcal{U},\mathcal{S})$ has a cover of size $\leq k$ if and only if the $\texttt{SWP}(Q_{\mathsf{matrix}}, D)$ has an integral solution of size $\leq |\mathcal{U}| + k\cdot|V| = n(k+1)$.

*Direction only-if.* Suppose we are given a cover $\mathcal{S}'$ of size $k$ to $(\mathcal{U},\mathcal{S})$. We construct an integral solution $D'$ to $\texttt{SWP}(Q_{\mathsf{matrix}}, D)$ as follows. Let $B' \subseteq \texttt{dom}(B)$ be the set of values that corresponding to $\mathcal{S}'$. For every $b_S \in B'$, i.e., $S \in \mathcal{S}'$, we add tuple $(b_S, c_u)$ to $D'$ for every $u \in \mathcal{U}$. For every $a_u \in \texttt{dom}(A)$, we choose an arbitrary value $b_S \in B'$ such that $u \in S$, and add tuple $(a_u, b_S)$ to $D'$. This is always feasible since $\mathcal{S}'$ is a valid set cover. It can be easily checked that $n$ tuples from $R_1$ and $k\cdot n$ tuples from $R_2$ are added to $D'$.

*Direction if.* Suppose we are given a integral solution $D'$ of size $k'$ to $\texttt{SWP}(Q_{\mathsf{matrix}}, D)$. Let $B'$ be the subset of values whose incident tuples in $R_2$ are included by $D'$. We argue that all subsets corresponding to $B'$ forms a valid cover of size $\frac{k'}{n-1}$. By definition of integral solution, $|B'| = \frac{k'}{n} - 1$. Moreover, for every value $a \in \texttt{dom}(A)$, at least one edge $(a,b)$ is included by $D'$ for some $b \in B'$, hence $B'$ must be a valid set cover. ◀

**Proof of Lemma 17.** We show that for any CQ $Q$ with head-domination property and an arbitrary database $D$, Algorithm 1 always return a $|\texttt{rels}(Q)|$-approximated solution $D'$ to $\texttt{SWP}(Q, D)$. It can be easily checked that $Q(D) = Q(D')$.

Consider the connected components $E_1, E_2, \cdots, E_k$ of $G_Q^\exists$ with dominants $\dot{R}_1, \dot{R}_2, \cdots, \dot{R}_k$. Let $Q_i$ be the subquery defined over relations in $E_i$, with output attributes $\texttt{head}(Q_i) = \texttt{attr}(\dot{R}_i)$, and $D_i = \{R_j : R_j \in E_i\}$ be the corresponding database for $Q_i$. Let $D^*$ be the solution to $\texttt{SWP}(Q, D)$. We point out some observations on $D^*$:

- For each $R_j$ with $\texttt{attr}(R_j) \subseteq \texttt{head}(Q)$, $D^*$ must include tuples in $\pi_{\texttt{attr}(R_j)} Q(D)$.
- For every dominant $\dot{R}_i$, $D^*$ must include at least $|\pi_{\texttt{head}(\dot{R}_i)} Q(D)|$ tuples from $\dot{R}_i$.

On the other hand, Algorithm 1 includes $|\texttt{rels}(Q_i)|$ tuples for each primitive at line 11, and invokes this primitive for each tuple in $\pi_{\texttt{head}(\dot{R}_i)} Q(D)$. Together, we come to:

$$
\begin{aligned}
|D'| &= \sum_{R_j : \texttt{attr}(R_j) \subseteq \texttt{head}(Q)} \left|\pi_{\texttt{attr}(R_j)} Q(D)\right| + \sum_{\dot{R}_i : i \in [k]} \left|\pi_{\texttt{head}(\dot{R}_i)} Q(D)\right| \cdot |\texttt{rels}(Q_i)| \\
&\le |D^*| + |D^*| \cdot |\texttt{rels}(Q_i)| = 2 \cdot |D^*| \cdot |\texttt{rels}(Q)|
\end{aligned}
$$

It can also be easily checked that $Q(D) = Q(D')$. Hence, Algorithm 1 always returns a $O(1)$-approximation solution for $\texttt{SWP}(Q, D)$.

Moreover, the query results $Q(D)$ can be computed in poly-time in the input size of $D$. Each primitive of finding the smallest witness for one query result takes $O(1)$ time. Hence, Algorithm 1 runs in polynomial time. ◀

**Proof of Lemma 24.** Let $P = \langle A_1, A_2, \cdots, A_k \rangle$ be such a sequence. For simplicity, let $P' = \{A_2, A_3, \cdots, A_{k-1}\}$. We next show a reduction from set cover to $\texttt{SWP}(Q, D)$. Consider an arbitrary instance of set cover with a universe $\mathcal{U}$ and a family $\mathcal{S}$ of subsets of $\mathcal{U}$ where $|\mathcal{U}| = n$ and $|\mathcal{S}| = n^c$ for some constant $c \ge 1$. We construct a database $D$ for $Q$ as follows. For each $u \in \mathcal{U}$, we add a value $a_u$ to $\texttt{dom}(A_1)$ and $\texttt{dom}(A_k)$. For each subset $S \in \mathcal{S}$, we add a value $b_S$ to $\texttt{dom}(B)$ for every $B \in P'$. We set the domain of remaining attributes in $\texttt{attr}(Q) - P$ as $\{*\}$. For each relation $R_j \in \texttt{rels}(Q)$, we distinguish the following cases:

- If $A_1 \in \texttt{attr}(R_j)$, we further distinguish two more cases:
  - if $\texttt{attr}(R_j) \cap P' = \emptyset$, we add tuple $t$ with $\pi_{A_1} t = a_u$ for every $u \in \mathcal{U}$;
  - otherwise, we add tuple $t$ with $\pi_{A_1} t = a_u$ and $\pi_{A_i} t = b_S$ for $A_i \in \texttt{attr}(R_j) \cap P'$, for every pair $(u, S) \in \mathcal{U} \times \mathcal{S}$ such that $u \in S$;
- If $A_k \in \texttt{attr}(R_j)$, we further distinguish two more cases:
  - if $\texttt{attr}(R_j) \cap P' = \emptyset$, we add tuple $t$ with $\pi_{A_k} t = a_u$ for every $u \in \mathcal{U}$;
  - otherwise, we add tuple $t$ with $\pi_{A_k} t = a_u$ and $\pi_{A_i} t = b_S$ for $A_i \in \texttt{attr}(R_j) \cap P'$, for every pair $(u, S) \in \mathcal{U} \times \mathcal{S}$;
- If $P \cap \texttt{attr}(R_j) \subseteq P - \{A_1, A_k\}$, we add a tuple $t$ with $\pi_{A_i} t = b_S$ for $A_i \in \texttt{attr}(R_j) \cap P$, for every $S \in \mathcal{S}$;
- If $P \cap \texttt{attr}(R_j) = \emptyset$, then we add a tuple $\{*\}$;

It can be easily checked that every relation contains at most $n^{c+1}$ tuples, hence $\log N = \Theta(\log n)$. The query result $Q(D)$ is exactly the Cartesian product of $\mathcal{U} \times \mathcal{U}$.

Consider a sub-database $D'$ of $D$ constructed above. Let $R'_j$ be the corresponding sub-relation of $R_j$ in $D'$. A solution $D'$ to $\texttt{SWP}(Q, D)$ is integral if $R'_j = \left(\pi_{\texttt{attr}(R_j) \cap P'} R'_j\right) \times (\pi_{A_k} R_j)$ holds for every relation $R_j$ with $A_k \in \texttt{attr}(R_j)$ and $\texttt{attr}(R_j) \cap P' \neq \emptyset$. Applying a similar argument as Lemma 21, we can show that there always exists an integral solution to $\texttt{SWP}(Q, D)$. Below, it suffices to focus on integral solutions.

It can be easily proved that $(\mathcal{U}, \mathcal{S})$ has a cover of size $\le k$ if and only if $\texttt{SWP}(Q, D)$ has an integral solution of size $\le nq_1 + knq_2 + q_3$ where $q_1, q_2, q_3 \le |\texttt{rels}(Q)|$ are query-dependent parameters. If $\texttt{SWP}$ is $(1-o(1))\cdot \log N$-approximatable for $Q$, there is a poly-time algorithm that can approximate set cover instances of input size $n$ within a $\log n$-factor, which is impossible unless $\texttt{P} = \texttt{NP}$. ◀

**Proof of Lemma 25.** Let $P \subseteq \texttt{attr}(Q)$ be the subset of attributes corresponding to the clique in the renamed query of $Q$. We identify the relation that contains the most number of output attributes of $P$, say $R_2 = \arg\max_{R_i \in \texttt{rels}(Q)} |\texttt{attr}(R_i) \cap P \cap \texttt{head}(Q)|$. As there exists no relation $R_k \in \texttt{rels}(Q)$ with $\texttt{head}(Q) \cap P \subseteq \texttt{attr}(R_k)$, it is always feasible to identify an attribute $A \in \texttt{head}(Q) \cap P - \texttt{attr}(R_2)$. For simplicity, we denote $P' = \texttt{head}(Q) \cap P - \texttt{attr}(R_2) - \{A\} = \{A_1, A_2, \cdots, A_\ell\}$, for some integer $\ell$. All non-output attributes in $P - \texttt{head}(Q)$ collapse to a single attribute $F$. All other attributes in $\texttt{attr}(Q) - P$ contain a dummy value $\{*\}$.

Consider an arbitrary instance of set cover $(\mathcal{U}, \mathcal{S})$ with $\mathcal{U} = \{u_1, u_2, \cdots, u_n\}$ and $\mathcal{S} = \{S_1, S_2, \cdots, S_m\}$, where $m = n^c$ for some constant $c > 1$. We construct a database $D$ for $Q$. Let $\texttt{dom}(A) = \{a_1, a_2, \cdots, a_n\}$, $\texttt{dom}(B) = \{b_1, b_2, \cdots, b_n\}$, $\texttt{dom}(C) = \{c_1, c_2, \cdots, c_n\}$ for every $C \in P'$, and $\texttt{dom}(F) = \texttt{dom}(F^-) \times \texttt{dom}(F^+)$ where $\texttt{dom}(F^-) = \{f_1^-, f_2^-, \cdots, f_m^-\}$ and $\texttt{dom}(F^+) = \{f^+_{h_1,h_2,\cdots,h_\ell} : \forall h_1, h_2, \cdots, h_\ell \in [n]\}$. We distinguish the following cases for a relation $R_i \in \texttt{rels}(Q)$:

- If $\texttt{attr}(R_i) \cap P \subseteq \texttt{head}(Q)$, $R_i$ is a Cartesian product over all attributes in $\texttt{head}(R_i) \cap P$;
- Otherwise, we further distinguish the following three cases:
  - $R_2$ is a Cartesian product over all attributes in $\texttt{attr}(R_2) \cap P$;
  - If $A, F \in \texttt{attr}(R_i)$, we construct sub-relation $(\pi_{A,F} R_i)$ such that for each pair $(u_\ell, S_j) \in \mathcal{U} \times \mathcal{S}$ with $u_\ell \in S_j$, we add tuples $\left\{a_\ell, f_j^-\right\} \times \texttt{dom}(F^+)$ to $(\pi_{A,F} R_i)$. If $\texttt{attr}(R_i) \cap P' \neq \emptyset$, for each tuple $\left(a_\ell, f^+_{h_1,h_2,\cdots,h_\ell}, f_j^-\right) \in (\pi_{A,F} R_i)$, we extend it by attaching value $c_{h_j}$ for attribute $A_j \in \texttt{attr}(R_i) \cap P'$. This way, we already obtain $\left(\pi_{A,F,\texttt{attr}(R_i)\cap P'} R_i\right)$. At last, we construct $R_i$ as the Cartesian product of remaining attributes in $\texttt{head}(R_i) \cap \texttt{head}(R_2) \cap P$ and $\left(\pi_{A,F,\texttt{attr}(R_i)\cap P'} R_i\right)$.
  - Otherwise, $F \in \texttt{attr}(R_i)$ but $A \notin \texttt{attr}(R_i)$. We construct sub-relation $(\pi_F R_i) = \texttt{dom}(F)$. If $\texttt{attr}(R_i) \cap P' \neq \emptyset$, for each tuple $\left(f^+_{h_1,h_2,\cdots,h_\ell}, f_j^-\right) \in \pi_F R_i$, we extend it by attaching value $c_{h_j}$ for attribute $A_j \in \texttt{attr}(R_i) \cap P'$. This way, we already obtain $\left(\pi_{F,\texttt{attr}(R_i)\cap P'} R_i\right)$. At last, we construct $R_i$ as the Cartesian product of remaining attributes in $\texttt{head}(R_i) \cap \texttt{head}(R_2) \cap P$ and $\left(\pi_{A,F,\texttt{attr}(R_i)\cap P'} R_i\right)$.

It can be checked that every relation contains $O(n^{|\texttt{head}(Q)\cap P|} \cdot m)$ tuples, hence $\log N = \Theta(\log n)$. Meanwhile, the query result $Q(D)$ is the Cartesian product over attributes in $\texttt{head}(Q) \cap P$, so every solution to $\texttt{SWP}(Q, D)$ must contain all tuples in $R_i$ if $\texttt{attr}(R_i) \cap P \subseteq \texttt{head}(Q)$, and the dummy tuple $\{*\}$ in every relation $R_i$ if $\texttt{attr}(R_i) \cap P = \emptyset$.

Here, $D$ also enjoys highly symmetric structure over every attribute $C \in P'$. More specifically, every tuple $t = (c_{i_1}, c_{i_2}, \cdots, c_{i_\ell}) \in \times_{C\in P'} \texttt{dom}(C)$ induces a subquery by removing all attributes in $P'$, and restricting $\texttt{dom}(F)$ as $\texttt{dom}(F_t) = \texttt{dom}(F^-) \times \left\{f^+_{i_1}, f^+_{i_2}, \cdots, f^+_{i_\ell}\right\}$. In a sub-database $D' \subseteq D$, let $R'_i$ be the corresponding sub-relation $R_i$. A solution $D'$ to $\texttt{SWP}(Q, D)$ is *integral* if:

- for every relation $R_i \in \texttt{rels}(Q)$ with $A, F \in \texttt{attr}(R_i)$,

$$\pi_{A,F,\texttt{head}(R_i)\cap\texttt{head}(R_2)\cap P} R'_i = \left(\pi_{A,F^-} R'_i\right) \times \texttt{dom}(F^+) \times \left(\times_{C\in\texttt{head}(R_i)\cap\texttt{head}(R_2)\cap P} \texttt{dom}(C)\right)$$

- for every relation $R_i \in \texttt{rels}(Q)$ with $A \notin \texttt{attr}(R_i)$ and $F \in \texttt{attr}(R_i)$,

$$\pi_{F,\texttt{head}(R_i)\cap\texttt{head}(R_2)\cap P} R'_i = (\pi_{F^-} R'_i) \times \texttt{dom}(F^+) \times \left(\times_{C\in\texttt{head}(R_i)\cap\texttt{head}(R_2)\cap P} \texttt{dom}(C)\right)$$

It can be easily shown that there always exists an integral solution to $\texttt{SWP}(Q, D)$. Moreover, $(\mathcal{U}, \mathcal{S})$ has a cover of size $\leq k$ if and only if $\texttt{SWP}(Q, D)$ has an integral solution of size $(q_1 k + q_2) \cdot n^{|\texttt{head}(Q) \cap P| - 1}$, where $q_1, q_2 \leq |\texttt{rels}(Q)|$ are some query-dependent parameters. Applying a similar argument as Lemma 21, we can show that if SWP is $(1 - o(1)) \cdot \log N$-approximable for $Q$, there is a poly-time algorithm that can approximate set cover instances of input size $n$ within a $\log n$-factor, which is impossible unless $\texttt{P} = \texttt{NP}$. ◀

▶ **Definition 40** (Cycle). *In a CQ $Q$, a sequence of attributes $P = \langle A_1, A_2, \cdots, A_k \rangle$ is a cycle if*

- *for every $i \in [k]$, there exists a relation $R_j \in \texttt{rels}(Q)$ such that $A_i, A_{(i+1) \bmod k} \in \texttt{attr}(R_j)$;*
- *for every $R_j \in \texttt{rels}(Q)$, either $\texttt{attr}(R_j) \cap P = \emptyset$, or $|\texttt{attr}(R_j) \cap P| = 1$ or $\texttt{attr}(R_j) \cap P = \{A_i, A_{(i+1) \bmod k}\}$ for some $i \in [k]$.*

▶ **Definition 41** (Non-Conformal Clique). *In a CQ $Q$, a non-conformal clique is a subset of attributes $P \subseteq \texttt{attr}(Q)$, such that*

- *for every pair of attributes $A_i, A_j$, there exists a relation $R_k \in \texttt{rels}(Q)$ with $A_i, A_j \in \texttt{attr}(R_k)$;*
- *there is no relation $R_\ell \in \texttt{rels}(Q)$ such that $P \subseteq \texttt{attr}(R_\ell)$.*

▶ **Lemma 42** ([9]). *Every cyclic CQ contains either a cycle or a non-conformal clique.*

**Proof of Lemma 29.** We already prove this lemma for acyclic CQs, such that if there exists no free sequence, $Q$ has head-domination property. Below, we focus on cyclic CQs. From Lemma 42, $Q$ must contain a cycle or a non-formal clique.

Consider an arbitrary cycle or clique $P$. As shown in Lemma 44, if $P \cap \texttt{head}(Q) \neq \emptyset$ and $P - \texttt{head}(Q) \neq \emptyset$, there must exist a relation $R_k \in \texttt{rels}(Q)$ such that removing $\texttt{head}(R_k)$ turns $Q$ into a disconnected CQ, where relations in $\{R_i : \texttt{attr}(R_i) \cap P - \texttt{head}(Q) \neq \emptyset\}$ is a connected subquery. Then, it is safe to remove all relations in $\{R_i : \texttt{attr}(R_i) \cap P - \texttt{head}(Q) \neq \emptyset\}$ from $Q$, and prove that the residual query (where $P$ disappears) has head-domination property.

Applying this removal procedure iteratively, we are left with cycles and non-conformal cliques that contain only output attributes, or only non-output attributes. Lemma 43 proves that $Q$ in this case has head-domination property, thus completing the proof. ◀

▶ **Lemma 43.** *In a cyclic CQ $Q$ without free sequences and nested cliques in its renamed query, if every cycle or non-conformal clique contains only output attributes, or only non-output attributes, $Q$ has head-domination property.*

**Proof of Lemma 43.** We apply the following procedure to $Q$: if there exists a cycle or non-conformal clique $P$, we add a *virtual* relation $R_p$ with $\texttt{attr}(R_p) = P$ to $\texttt{rels}(Q)$. We apply this procedure iteratively until no more cycle or non-conformal clique can be found. Let $Q'$ be the resulted query.

We first prove that for each virtual relation $R_p$ added, either $\texttt{attr}(R_p) \subseteq \texttt{head}(Q)$ or $\texttt{attr}(R_p) \subseteq \texttt{attr}(Q) - \texttt{head}(Q)$. Adding virtual relation does not increase more cycles, but some possible non-conformal cliques. We distinguish the following three cases:

- If $R_p$ is added for a cycle $P$, then $P$ must exist in $Q$ and therefore $\texttt{attr}(R_p) \subseteq \texttt{head}(Q)$ or $\texttt{attr}(R_p) \subseteq \texttt{attr}(Q) - \texttt{head}(Q)$;
- If $R_p$ is added for a non-conformal clique $P$, and $P$ exists in $Q$, then $\texttt{attr}(R_p) \subseteq \texttt{head}(Q)$ or $\texttt{attr}(R_p) \subseteq \texttt{attr}(Q) - \texttt{head}(Q)$;

**Algorithm 4** GreedySWP$(Q, D)$

```
1 D' ← ∅;
2 foreach t ∈ Q(D) do D' ← D' ∪ SWP(Q, D, t);
3 return D';
```

- Otherwise, $R_p$ is added for a non-conformal clique $P$ but $P$ does not exist in $Q$. It must be the case that after $R_{p'}$ is added for a cycle $P'$, a non-conformal clique $P$ forms with $P' \subsetneq P$. If $P' \subseteq \texttt{head}(Q)$, then $P \subseteq \texttt{head}(Q)$; otherwise, $P$ is a nested clique in the renamed query of $Q$. If $P' \subseteq \texttt{attr}(Q) - \texttt{head}(Q)$, then $P \subseteq \texttt{attr}(Q) - \texttt{head}(Q)$; otherwise, any pair of non-consecutive attributes in $P'$ will form a free sequence together with an arbitrary attribute in $P \cap \texttt{head}(Q)$, contradicting the fact that $Q$ does not contain a free sequence. In either case, $R_p$ is added as a virtual relation with $\texttt{attr}(R_p) \subseteq \texttt{head}(Q)$ or $\texttt{attr}(R_p) \subseteq \texttt{attr}(Q) - \texttt{head}(Q)$.

We next point out some critical observations on $Q'$: (1) $\texttt{head}(Q) = \texttt{head}(Q')$; (2) $Q'$ is acyclic, as it does not contain any cycle or non-conformal clique; and (3) $Q'$ also does not contain any free sequence as $Q$. Together, $Q'$ is free-connex. It is feasible to find a free-connex join tree, and identify connected components $E'_1, E'_2, \cdots, E'_k$ of $G^{\exists}_{Q'}$ with dominants $\dot{R}_1, \dot{R}_2, \cdots, \dot{R}_k \in \texttt{rels}(Q')$. Moreover, as $Q'$ is connected, $\texttt{head}(\dot{R}_i) \neq \emptyset$ holds for every $i \in [k]$. We next show that virtual relations are not dominants. By contradiction, we assume $R_p = \dot{R}_1$ for some virtual relation $R_p$. If $\texttt{attr}(R_p) - \texttt{head}(Q) = \emptyset$, $E'_1 \cup \{R_p\}$ form a connected subquery, so $Q'$ is disconnected, leading to a contradiction. If $\texttt{attr}(R_p) \subseteq \texttt{head}(Q)$, we further distinguish two more cases. If $R_p$ is added for a non-conformal clique, then $E'_1$ form a nested-clique in the rename query of $Q$; and if $R_p$ is added for a cycle, then any non-consecutive attributes of $P$ together with non-output attributes in relations of $E'_1$ form free-sequence. Hence, virtual relations cannot be dominants, and $\dot{R}_1, \dot{R}_2, \cdots, \dot{R}_k \in \texttt{rels}(Q)$.

Let $E_i \subseteq E'_i$ be the set of non-virtual relations in $E'_i$. It can be easily checked that $E_1, E_2, \cdots, E_k$ are also the connected components of $G^{\exists}_{Q}$ with dominants $\dot{R}_1, \dot{R}_2, \cdots, \dot{R}_k$. Hence, $Q$ has head-domination property. ◀

▶ **Lemma 44.** *In a cyclic CQ $Q$ without free sequences and nested cliques in its renamed query, for any cycle or non-conformal clique $P$, if $P \cap \texttt{head}(Q) \neq \emptyset$ and $P - \texttt{head}(Q) \neq \emptyset$, there must exist a relation $R_k \in \texttt{rels}(Q)$ such that removing $\texttt{head}(R_k)$ turns $Q$ into a disconnected CQ, where $\{R_i : \texttt{attr}(R_i) \cap P - \texttt{head}(Q) \neq \emptyset\}$ is a connected subquery.*

**Proof.** Consider an arbitrary non-output attribute $A \in \texttt{attr}(Q) - \texttt{head}(Q)$. Let $V_A \subseteq \texttt{attr}(Q)$ be the set of attributes renamed to attribute $A$ in this procedure. Let $N_A = \{B \in \texttt{attr}(Q) - V_A : \exists B' \in V_A, R_i \in \texttt{attr}(Q), s.t. B, B' \in \texttt{attr}(R_i)\}$ be the set of attributes that appear together with some attribute of $V_A$ in some relation. Note that $N_A \subseteq \texttt{head}(A)$; otherwise, any attribute in $N_A - \texttt{head}(A)$ will be renamed as $A$, coming to a contradiction. Moreover, for every pair of attributes $B, B' \in N_A$, there must exist a relation $R_i$ such that $B, B' \in \texttt{attr}(R_i)$; otherwise, a free sequence is identified, coming to a contradiction. Thus, $N_A$ is a clique. In addition, there must exist one relation $R_k$ such that $N_A \subseteq \texttt{attr}(R_k)$; otherwise, a nested clique $N_A \cup \{A\}$ is identified in the renamed query, coming to a contradiction. In this way, removing all attributes in $V_A \subseteq \texttt{head}(R_k)$ turns $Q$ into a disconnected CQ, where $\{R_i \in \texttt{rels}(Q) : \texttt{head}(R_i) \subseteq V_A\}$ is a connected subquery.

Let's come to cycle $P$. For simplicity, denote $P = \langle A_1, A_2, \cdots, A_k \rangle$. As there is no free sequence, $P \cap \texttt{head}(Q) = \{A_i, A_{(i+1) \bmod k}\}$ for some $i \in [k]$. Then, we can identify an

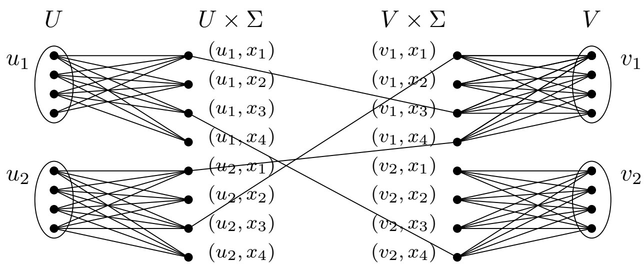


**Figure 7** An example of hard instance constructed for $Q_{\text{line3}}$.

arbitrary output attribute in $P - \texttt{head}(Q)$ as $A$, and therefore

- $P - \texttt{head}(Q) \subseteq V_A$;
- for every relation $R_j \in \texttt{rels}(Q)$ with $\texttt{attr}(R_j) \cap P - \texttt{head}(Q) \neq \emptyset$, $\texttt{head}(R_j) \subseteq V_A$.

Hence, removing $\texttt{head}(R_k)$ turns $Q$ into a disconnected CQ, where $\{R_i : \texttt{attr}(R_i) \cap P - \texttt{head}(Q) \neq \emptyset\}$ is a connected subquery.

We next turn to non-conformal clique $P$. We can identify an arbitrary output attribute in $P - \texttt{head}(Q)$ as $A$, and therefore

- $P - \texttt{head}(Q) \subseteq V_A$;
- for every relation $R_j \in \texttt{rels}(Q)$ with $\texttt{attr}(R_j) \cap P - \texttt{head}(Q) \neq \emptyset$, $\texttt{head}(R_j) \subseteq V_A$.

Hence, removing $\texttt{head}(R_k)$ turns $Q$ into a disconnected CQ, where $\{R_i : \texttt{attr}(R_i) \cap P - \texttt{head}(Q) \neq \emptyset\}$ is a connected subquery. ◀

# F Missing Proofs in Section 4

**Proof of Theorem 30.** Let $D'$ be the solution returned by Algorithm 4. For each query result $t \in Q(D)$, we add at most $|\texttt{rels}(Q)|$ tuples to $D'$, so $|D'| \leq |Q(D)| \cdot |\texttt{rels}(Q)|$. Meanwhile, $|D'| \leq N \cdot |\texttt{rels}(Q)|$. Let $D^*$ be the solution to $\texttt{SWP}(Q, D)$. Note that at most $|D^*|^\rho$ query results can be reproduced from $D^*$. Hence, we must have $|D^*|^\rho \geq |Q(D)|$, i.e., $|D^*| \geq |Q(D)|^{1/\rho}$. Below, we show that $|D'| \leq N^{1-1/\rho} \cdot |D^*|$ with the following facts:

- If $|Q(D)| \leq N$, then the approximation ratio (skipping $|\texttt{rels}(Q)|$) is $\frac{|D'|}{|D^*|} \leq \frac{|Q(D)|}{|Q(D)|^{1/\rho}} = |Q(D)|^{1-1/\rho} \leq N^{1-1/\rho}$.
- If $N < Q(D)$, then the approximation ratio (skipping $|\texttt{rels}(Q)|$) is at most $\frac{|D'|}{|D^*|} \leq \frac{N}{|Q(D)|^{1/\rho}} \leq \frac{N}{N^{1/\rho}} = N^{1-1/\rho}$.

As $Q(D)$ can be computed in polynominal time, $Q(D)$ is polynomially large and $\texttt{SWP}$ for a single tuple can be computed in poly-time, this baseline runs in polynominal time. ◀

**Proof of Theorem 35.** We focus on line-3 query $Q_{\text{line3}}(A_1, A_4) : -R_1(A_1, A_2), R_2(A_2, A_3), R_3(A_3, A_4)$, which can be generalized to line-$m$ query $Q_{\text{line}}(A_1, A_m) : -R_1(A_1, A_2), R_2(A_2, A_3), \cdots, R_m(A_m, A_{m+1})$, by putting $A_3, A_4, \cdots, A_{m-1}$ as identical copies of $A_m$.

**Label Cover.** We introduce the *Label Cover Problem* as below:

▶ **Definition 45** (Label Cover). *Given a complete bipartite graph $(U, V, E)$ with $|U| = |V| = n$, a finite alphabet $\Sigma$ with $\Sigma > n$, constraints $C_{u,v}$ for every edge $e = (u, v)$, the goal is find assignments to the vertices $M_U : U \to 2^\Sigma$ and $M_V : V \to 2^\Sigma$ that every* *edge is satisfied, i.e., there exists $x \in M_U(u), y \in M_V(v)$ such that $(x, y) \in C_{u,v}$, and $\sum_{u\in U} |M_U(u)| + \sum_{v\in V} |M_V(v)|$ is minimized. The problem size $m$ is defined as $n \times |\Sigma|$.*

▶ **Lemma 46** ([20]). *There is no poly-time algorithm that can approximate the label cover instances of size $m$ within $O\left(2^{(\log m)^{1-\epsilon}}\right)$ factor for any $\epsilon > 0$.*

**SWP for $Q_{\mathsf{line3}}$.** We set $\mathtt{dom}(A) = U \times [n^2]$, $\mathtt{dom}(B) = U \times \Sigma$, $\mathtt{dom}(C) = V \times \Sigma$ and $\mathtt{dom}(D) = V \times [n^2]$. The database $D$ includes:

- $R_1 = \{(\langle u, i\rangle, \langle u, x\rangle) : u \in V, i \in [n^2], x \in \Sigma\}$;
- $R_2 = \{(\langle u, x\rangle, \langle v, y\rangle) : u, v \in V, x, y \in \Sigma, (x, y) \in C_{uv}\}$;
- $R_3 = \{(\langle v, y\rangle, \langle v, i\rangle) : v \in V, y \in \Sigma, i \in [n^2]\}$;

In this instance, we have $N = O(n^2 \cdot |\Sigma|^2 + n^3 \cdot |\Sigma|) = O(m^2)$. It can be easily checked that $Q_{\mathsf{line3}}(D) = U \times [n^2] \times V \times [n^2]$ with $n^6$ query results. For any witness $D'$ for $\mathtt{SWP}(Q_{\mathsf{line3}}, D)$, let $D'_1, D'_2, D'_3$ be the set of tuples chosen from $R_1, R_2, R_3$ separately.

For a sub-database $D' \subseteq D$, let $R'_1, R'_3$ are the corresponding subset of tuples from $R_1, R_3$ in $D'$. A solution $D'$ to $\mathtt{SWP}(Q_{\mathsf{line3}}, D)$ is integral if for any $(u, x) \in U \times \Sigma$, either $\sigma_{B=(u,x)} R'_1 = \emptyset$ or $\sigma_{B=(u,x)} R'_1 = \{\langle u, i\rangle, \langle u, x\rangle : i \in [n^2]\}$, and for any $(v, y) \in V \times \Sigma$, either $\sigma_{C=(v,y)} R'_3 = \emptyset$ or $\sigma_{C=(v,y)} R'_3 = \{\langle v, i\rangle, \langle v, y\rangle : i \in [n^2]\}$.

▶ **Lemma 47.** *Every non-integral solution to $\mathtt{SWP}(Q_{\mathsf{line3}}, D)$ can be transformed into an integral solution to $\mathtt{SWP}(Q_{\mathsf{line3}}, D)$.*

**Proof.** Consider an arbitrary non-integral solution $D'$. W.o.l.g., assume there exists some $(u, x) \in U \times X$, such that $\sigma_{B=(u,x)} R'_1 \neq \emptyset$ and $\sigma_{B=(u,x)} R'_1 \neq \{\langle u, i\rangle, \langle u, x\rangle : i \in [n]\}$. There must exist some pair of $(i, j)$ such that $\langle u, i\rangle, \langle u, x\rangle \in D'_1$ but $\langle u, j\rangle, \langle u, x\rangle \notin D'_1$. Let $X_i = \{x \in X : \langle u, i\rangle, \langle u, x\rangle\} \in D'_1$ and $X_j = \{x \in X : \langle u, i\rangle, \langle u, x\rangle\} \in D'_1$. W.l.o.g., assume $|X_i| \le |X_j|$. As $D'$ is a solution to $\mathtt{SWP}(Q_{\mathsf{line3}}, D)$, it is feasible to remove $\langle u, j\rangle, \langle u, x\rangle \in D'_1$ for $x \in X_j$ from $D'$ and add $\langle u, j\rangle, \langle u, x\rangle \in D'_1$ for $x \in X_i$ to $D'$, which yields another solution to $\mathtt{SWP}(Q_{\mathsf{line3}}, D)$ without increasing the number of tuples. After applying this step iteratively, we can obtain a solution such that every $(u, i)$ is incident to the same values in $\mathtt{dom}(B)$ over $i \in [n]$, i.e., an integral solution. ◀

From now on, it suffices to consider the integral solution to $\mathtt{SWP}(Q_{\mathsf{line3}}, D)$. Moreover, we can assume $|D'_2| = n^2$ by picking an arbitrary edge $(\langle u, x\rangle, \langle v, y\rangle) \in R_2$ with $(x, y) \in C_{uv}$ for $(u, v) \in U \times V$. Observe that there is a solution to label cover of cost at most $c$ if and only if there is an integral solution to SWP for $Q_{\mathsf{line3}}$ of cost at most $n^2 c + n^2 = n^2(c + 1)$.

- *Direction Only-If:* Let $S = (M_U, M_V)$ be the solution to label cover. We define an integral solution $D_S$ to $\mathtt{SWP}(Q_{\mathsf{line3}}, D)$ based $S$. For each vertex $u \in U$, we add the set of tuples $(\langle u, *\rangle, x) \in R_1$ to $D_S$ if $x \in M_U(u)$. Similarly, for each vertex $v \in V$, we add the set of tuples $(\langle v, *\rangle, y) \in R_3$ to $D_S$ if $y \in M_V(v)$. Together with the $n^2$ tuples chosen from $R_2$, we have obtained a witness $D_S$ of size $c \cdot n^2 + n^2 = (c + 1)n^2$.
- *Direction If:* Let $D_S$ be an integral witness to $\mathtt{SWP}(Q_{\mathsf{line3}}, D)$ of cost $c'$. We construct a solution $(M_U, M_V)$ to label cover as follows. If $(\langle u, *\rangle, x) \in R_1$ is included by $D_S$, we add $x$ to $M_U(u)$. Similarly, if $(\langle v, *\rangle, y) \in R_3$ is included by $D_S$, then we add $y$ to $M_V(v)$. It can be easily checked that we have obtained a solution to label cover of size $\frac{c'-1}{n^2}$.

If there is a poly-time algorithm that can approximate SWP for $Q_{\mathsf{line3}}$ within $2^{(\log\sqrt{N})^{1-\epsilon}}$ factor, then there is a poly-time algorithm that can approximate label cover within $2^{(\log\sqrt{N})^{1-\epsilon}} = 2^{(\log m)^{1-\epsilon}}$ factor, coming to a contradiction of Lemma 46.

◀