




ENUMERATION ALGORITHMS FOR CONJUNCTIVE QUERIES WITH PROJECTION

SHALEEN DEEP ^a, XIAO HU ^b, AND PARASCHOS KOUTRIS ^c

^a Department of Computer Sciences, University of Wisconsin-Madison, Madison, Wisconsin, USA
e-mail address: shaleen@cs.wisc.edu

^b David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, Canada
e-mail address: xiaohu@uwaterloo.ca

^c Department of Computer Sciences, University of Wisconsin-Madison, Madison, Wisconsin, USA
e-mail address: paris@cs.wisc.edu

ABSTRACT. We investigate the enumeration of query results for an important subset of CQs with projections, namely star and path queries. The task is to design data structures and algorithms that allow for efficient enumeration with delay guarantees after a preprocessing phase. Our main contribution is a series of results based on the idea of interleaving precomputed output with further join processing to maintain delay guarantees, which maybe of independent interest. In particular, for star queries, we design combinatorial algorithms that provide instance-specific delay guarantees in linear preprocessing time. These algorithms improve upon the currently best known results. Further, we show how existing results can be improved upon by using fast matrix multiplication. We also present new results involving tradeoff between preprocessing time and delay guarantees for enumeration of path queries that contain projections. Boolean matrix multiplication is an important query that can be expressed as a CQ with projection where the join attribute is projected away. Our results can therefore also be interpreted as sparse, output-sensitive matrix multiplication with delay guarantees.

1. INTRODUCTION

The efficient evaluation of join queries over static databases is a fundamental problem in data management. There has been a long line of research on the design and analysis of algorithms that minimize the total runtime of query execution in terms of the input and output size [Yan81, NRR13, NPRR12]. However, in many data processing scenarios, it is beneficial to split query execution into two phases: the *preprocessing phase*, which computes a space-efficient intermediate data structure, and the *enumeration phase*, which uses the data structure to enumerate the query results as fast as possible, to minimize the *delay* between outputting two consecutive tuples in the result. This distinction is beneficial for several reasons. For instance, in many scenarios, the user wants to see one (or a few) results of the query as fast as possible: in this case, we want to minimize the time of the preprocessing

We are grateful to the reviewers for a careful reading of the manuscript and their feedback.

phase such that we can output the first results quickly. On the other hand, a data processing pipeline may require that the result of a query is accessed multiple times by a downstream task. In this case, spending more time during the preprocessing phase is better to guarantee a faster enumeration with a smaller delay.

Previous work in the database literature has focused on finding the class of queries that can be computed with $O(|D|)$ preprocessing time (where D is the input database instance) and constant delay during the enumeration phase. The main result in this line of work shows that full (i.e., without projections) acyclic Conjunctive Queries (CQs) admit linear preprocessing time and constant delay [BDG07]. If the CQ is not full but its free variables satisfy the *free-connex* property, the same preprocessing time and delay guarantees can still be achieved. It is also known that for any (possibly non-full) acyclic CQ, it is possible to achieve linear delay after linear preprocessing time [BDG07]. Prior work that used structural decomposition methods [GS13] generalized these results to arbitrary CQs with free variables and showed that the projected solutions could be enumerated with $O(|D|^{\text{fhw}})$ delay, where **fhw** is the *fractional hypertree width* [GGS14] of the query. Moreover, a dichotomy about the classes of conjunctive queries with fixed arities where such answers can be computed with polynomial delay (WPD) is also shown. When the CQ is full but not acyclic, factorized databases use $O(|D|^{\text{fhw}})$ preprocessing time to achieve constant delay. We should note here that we can always compute and materialize the query result during preprocessing to achieve constant delay enumeration, but this requires an exponential amount of space.

The aforementioned prior work investigates specific points in the preprocessing time-delay tradeoff space. While the story for full acyclic CQs is relatively complete, the same is not true for general CQs, even for acyclic CQs with projections. For instance, consider the simplest such query: $Q_{\text{two-path}} = \pi_{x,z}(R(x,y) \bowtie S(y,z))$, which joins two binary relations and then projects out the join attribute. For this query, [BDG07] ruled out a constant delay algorithm with linear time preprocessing unless the boolean matrix multiplication exponent is $\omega = 2$. However, we can obtain $O(|D|)$ delay with $O(|D|)$ preprocessing time. We can also obtain $O(1)$ delay with $O(|D|^2)$ preprocessing by computing and storing the full result. It is worth asking whether there are other interesting points in this tradeoff between preprocessing time and delay. Towards this end, seminal work by Kara et al. [KNOZ20] showed that for any hierarchical CQ¹ (possibly with projections), there always exists a smooth tradeoff between preprocessing time and delay. This is the first improvement over the results of Bagan et al. [BDG07] in over a decade for queries involving projections. Applied to the query $Q_{\text{two-path}}$, the main result of [KNOZ20] shows that for any $\epsilon \in [0, 1]$, we can obtain $O(|D|^{1-\epsilon})$ delay with $O(|D|^{1+\epsilon})$ preprocessing time.

In this article, we continue the investigation of the tradeoff between preprocessing time and delay for CQs with projections. We focus on two classes of CQs: *star queries*, a popular subset of hierarchical queries, and a useful subset of non-hierarchical queries known as *path queries*. We focus narrowly on these two classes for two reasons. First, star queries are of immense practical interest given their connections to set intersection, set similarity joins, and applications to entity matching (we refer the reader to [DHK20] for an overview). The most common star query seen in practice is $Q_{\text{two-path}}$ [BMT20]. The same holds for path queries, which are fundamental in graph processing. Second, as we will see in this article, the tradeoff landscape is complex and requires developing novel techniques even for the simple class of star queries. We also present a result on another subset of hierarchical CQs called

¹Hierarchical CQs are a strict subset of acyclic CQs.

left-deep. Our key insight is to design enumeration algorithms that depend not only on the input size $|D|$ but are also aware of other data-specific parameters, such as the output size. To give a flavor of our results, consider $Q_{\text{two-path}}$ as an example, and denote by OUT_{\bowtie} the output of the underlying full query (i.e., $R(x, y) \bowtie S(y, z)$). We can show the following result:

Theorem 1.1. *Given any database instance D , after $O(|D|)$ preprocessing time, the output of $Q_{\text{two-path}} = \pi_{x,z}(R(x, y) \bowtie S(y, z))$ can be enumerated with $O(|D|^2/|\text{OUT}_{\bowtie}|)$ delay.*

At this point, the reader may wonder about the improvement from the above result. [KNOZ20] implies that with preprocessing time $O(|D|)$, the delay guarantee in the worst-case is $O(|D|)$. This raises the question of whether the delay from Theorem 1.1 is truly an algorithmic improvement rather than an improved analysis of [KNOZ20]. We answer the question positively. Specifically, we show that there exists a database instance where the delay obtained from Theorem 1.1 is a polynomial improvement over the actual guarantee [KNOZ20] and not just the worst-case. Our proposed algorithm outperforms prior work by cleverly using the stored precomputed output to maintain the delay guarantee and alternate with further join processing that generates new output but is not fast enough to satisfy the delay guarantee. When the preprocessing time is linear, the delay implied by our result depends on the size of the full join. For the worst-case output size where $|\text{OUT}_{\bowtie}| = \Theta(|D|^2)$, we obtain the best possible delay, which will be constant. Compare this to the result of [KNOZ20], which would require nearly $O(|D|^2)$ preprocessing time to achieve the same guarantee. On the other hand, if $|\text{OUT}_{\bowtie}| = \Theta(|D|)$, we obtain only a linear delay guarantee of $O(|D|)$.² The reader may wonder how our result compares in general with the tradeoff in [KNOZ20] in the worst-case; we will show that we can always get at least as good of a tradeoff point as the one in [KNOZ20]. Figure 1 summarizes the prior work and the results present in this article.

Our Contribution. In this article, we improve the state-of-the-art preprocessing time-delay tradeoff for a subset of CQs with projections. We summarize our main technical contributions below (highlighted in Figure 1):

- (1) Our first contribution is a novel algorithm (Theorem 4.1) that achieves output-dependent delay guarantees for star queries (formally defined in Section 2.1) after linear preprocessing time. Specifically, we can achieve $O(|D|^{k/(k-1)}/|\text{OUT}_{\bowtie}|^{1/(k-1)})$ delay after linear preprocessing. Our key idea is to identify an appropriate degree threshold to split a relation into partitions of *heavy* and *light*, which allows us to perform efficient enumeration. For star queries, our result implies no smooth tradeoff between preprocessing time and delay guarantees as stated in [KNOZ20] for the general class of hierarchical queries.
- (2) We introduce the novel idea of *interleaving* join query computation in the context of enumeration algorithms, which forms the foundation for our algorithms and may be of independent interest. Specifically, we show that it is possible to union the output of two algorithms \mathcal{A} and \mathcal{A}' with δ delay guarantee where \mathcal{A} enumerates query results with δ delay guarantees but \mathcal{A}' does not. This technique allows us to compute a subset of a query *on the fly* when enumeration with good delay guarantees is impossible.
- (3) We use fast matrix multiplication techniques to obtain a better preprocessing time-delay tradeoff than [KNOZ20]. We also show an algorithm for left-deep hierarchical queries with linear preprocessing time and output-dependent delay guarantees.

²We do not need to consider the case where $|\text{OUT}_{\bowtie}| \leq |D|$, since then we can materialize the full result during the preprocessing time using constant delay enumeration for queries without projections [OZ15].

Queries	Preprocessing	Delay	Source
Arbitrary acyclic CQ	$O(D)$	$O(D)$	[BDG07]
Free-connex CQ (projections)	$O(D)$	$O(1)$	[BDG07]
Full CQ	$O(D ^{\mathbf{fhw}})$	$O(1)$	[OS16]
Arbitrary CQ with aggregations	$O(D ^{\mathbf{faqw}} \log D)$	$O(1)$	[AKNR16] + [OS16]
Full CQ	$O(D ^{\mathbf{subw}} \log D)$	$O(1)$	[AKNS17] + [OS16]
Hierarchical CQ (with projections)	$O(D ^{1+(w-1)\epsilon})$	$O(D ^{1-\epsilon}), \epsilon \in [0, 1]$	[KNOZ20]
Star query with k relations (with projections)	$O(D)$	$O(\frac{ D ^{k/(k-1)}}{ \mathbf{OUT}_{\bowtie} ^{1/(k-1)}})$	this article
Path query with k relations (with projections)	$O(D ^{2-\epsilon/(k-1)})$	$O(D ^\epsilon), \epsilon \in [0, 1]$	this article
Left-deep hierarchical CQ (with projections)	$O(D)$	$O(D ^k / \mathbf{OUT}_{\bowtie})$	this article
Two path query (with projections)	$O(D ^{\omega \cdot \epsilon})$	$O(D ^{1-\epsilon}), \epsilon \in [\frac{2}{\omega+1}, 1]$	this article

Figure 1: Preprocessing time and delay guarantees for different queries. $|\mathbf{OUT}_{\bowtie}|$ denotes the size of join query under consideration but without any projections. **subw** denotes the submodular width of the query [Mar13] and **faqw** denote the FAQ-width of the query [GM14]. If the CQ contains no aggregations, it is known that $\mathbf{faqw} = \mathbf{fhw}$. For each class of query, the total running time is $O(\min\{|D| \cdot |\mathbf{OUT}_{\pi}|, |D|^{\mathbf{subw}} \log |D| + |\mathbf{OUT}_{\pi}|\})$ where $|\mathbf{OUT}_{\pi}|$ denotes the size of the query result.

- (4) Finally, we present new results on preprocessing time-delay tradeoffs for non-hierarchical CQs, the class of path queries (formally defined in Section 2.1). Our results show that we can achieve delay $O(|D|^\epsilon)$ with preprocessing time $O(|D|^{2-\epsilon/(k-1)})$ for any $\epsilon \in [0, 1]$.

This article is the full version of a conference publication [DHK21]. We have added all of the proofs and intermediate results excluded from the paper. In particular, we have added the full proof of our main result – enumeration delay obtained for star queries (Theorem 4.1, Theorem 5.1, and Theorem 6.1) as well as the proofs of all the helper lemmas in Section 3. Additionally, we have added the full algorithm for the helper lemmas. We have also added Example 4.2 and Example 4.3 to improve the exposition. Finally, we have added a new result in Subsection 4.4 that transfers a result for static enumeration to the dynamic setting.

The remainder of the article is organized as follows. In Section 2, we give preliminary concepts, definitions, and notation. In Section 3, we present useful lemmas used frequently throughout the article. In Section 4, we state the main result and compare it against prior work. As a warm-up step, we present the intuition and main ideas for the two-path query, which are then generalized to star queries. Toward the end of the section, we also discuss how matrix multiplication can be useful. Section 5 and Section 6 are dedicated to studying a class of hierarchical queries that we call left-deep and path queries, respectively. Finally, we discuss past contribution measures in Section 7 and conclude in Section 8.

2. PROBLEM SETTING

In this section, we present the basic notation and terminology.

2.1. Conjunctive Queries. We will focus on the class of *conjunctive queries* (CQs):

$$Q = \pi_{\mathbf{y}}(R_1(\mathbf{x}_1) \bowtie R_2(\mathbf{x}_2) \bowtie \dots \bowtie R_n(\mathbf{x}_n))$$

Here, the symbols $\mathbf{y}, \mathbf{x}_1, \dots, \mathbf{x}_n$ are vectors that contain *variables* or *constants*. We say that Q is *full* if there is no projection. We will typically use the symbols x, y, z, \dots to denote variables, and a, b, c, \dots to denote constants. We use $Q(D)$ to denote the query's result Q over input database D . In this paper, we will focus on CQs with no constants and no repeated variables in the same atom (both cases can be handled within a linear time preprocessing step, so this assumption is without any loss of generality). Such a query can be represented equivalently as a *hypergraph* $\mathcal{H}_Q = (\mathcal{V}_Q, \mathcal{E}_Q)$, where \mathcal{V}_Q is the set of variables, and for each hyperedge $F \in \mathcal{E}_Q$ there exists a relation R_F with variables F . We will be particularly interested in two families of CQs fundamental to query processing: star and path queries. The *star query* with k relations is expressed as:

$$Q_k^* = \pi_{x_1, x_2, \dots, x_k} R_1(x_1, y) \bowtie R_2(x_2, y) \bowtie \dots \bowtie R_k(x_k, y)$$

where x_1, \dots, x_k are distinct variables. The *path query* with k relations is expressed as:

$$P_k = \pi_{x_1, x_{k+1}} R_1(x_1, x_2) \bowtie R_2(x_2, x_3) \bowtie \dots \bowtie R_k(x_k, x_{k+1})$$

Hierarchical Queries. A CQ Q is *hierarchical* if for any two of its variables, either the sets of atoms in which they occur are disjoint or one is contained in the other [SORK11]. For example, Q_k^* is hierarchical for any k , while P_k is hierarchical only when $k \leq 2$.

Join Size Bounds. Let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ be a hypergraph, and $S \subseteq \mathcal{V}$. A weight assignment $\mathbf{u} = (u_F)_{F \in \mathcal{E}}$ is called a *fractional edge cover* of S if (i) for every $F \in \mathcal{E}$, $u_F \geq 0$ and (ii) for every $x \in S$, $\sum_{F: x \in F} u_F \geq 1$. The *fractional edge cover number* of S , denoted by $\rho_{\mathcal{H}}^*(S)$ is the minimum of $\sum_{F \in \mathcal{E}} u_F$ over all fractional edge covers of S . We write $\rho^*(\mathcal{H}) = \rho_{\mathcal{H}}^*(\mathcal{V})$.

Tree Decompositions. Let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ be a hypergraph of a CQ Q . A *tree decomposition* of \mathcal{H} is a tuple $(\mathcal{T}, (\mathcal{B}_t)_{t \in V(\mathcal{T})})$ where \mathcal{T} is a tree, and every \mathcal{B}_t is a subset of \mathcal{V} , called the *bag* of t , such that

- (1) each edge in \mathcal{E} is contained in some bag; and
- (2) for each variable $x \in \mathcal{V}$, the set of nodes $\{t \mid x \in \mathcal{B}_t\}$ form a connected subtree of \mathcal{T} .

The *width* of a tree decomposition $(\mathcal{T}, (\mathcal{B}_t)_{t \in V(\mathcal{T})})$ is defined as $\max_{t \in V(\mathcal{T})} \rho^*(\mathcal{B}_t)$, where $\rho^*(\mathcal{B}_t)$ is the minimum fractional edge cover of the vertices in \mathcal{B}_t . The *fractional hypertree width* of a CQ Q , denoted as $\text{fhw}(Q)$, is the minimum width among all tree decompositions of its hypergraph. A query is *acyclic* if $\text{fhw}(Q) = 1$.

Computational Model. To measure the running time of our algorithms, we will use the uniform-cost RAM model [HUA75], where data values and pointers are of $O(1)$ size. All complexity results concern data complexity, where the query is assumed to be fixed. We will assume that a hash table or a hash map H can insert and delete a key in constant time. Further, given a key u , it can check existence, and return the value for the key (denoted by $H[u]$) in constant time. In practice, hashing can only achieve amortized constant time for some of the operations [CLRS22]. Therefore, in the paper, whenever we claim constant time operations for hash table, we mean amortized constant time.

2.2. Fast Matrix Multiplication. The matrix multiplication exponent ω is the smallest number such that for any $\epsilon > 0$, there is an algorithm that multiplies two rational $n \times n$ matrices with at most $O(n^{\omega+\epsilon})$ arithmetic operations. Let A be a $U \times V$ matrix and C be a $V \times W$ matrix over any field \mathcal{F} . $A_{i,j}$ is the shorthand notation for entry of A located in row i and column j . The matrix product is given by $(AC)_{i,j} = \sum_{k=1}^V A_{i,k}C_{k,j}$. Given their fundamental importance, algorithms for fast matrix multiplication are of great theoretical interest. We will use the following folklore lemma about rectangular matrix multiplication.

Lemma 2.1. *Let ω be the matrix multiplication exponent. Let $\beta = \min\{U, V, W\}$. Fast matrix multiplication of Boolean matrices of size $U \times V$ and $V \times W$ runs in time $O(UVW\beta^{\omega-3})$.*

In Lemma 2.1, matrix multiplication cost dominates the time required to construct the input matrices (if they have not been built already) for all $\omega \geq 2$. Fixing $\omega = 2$, rectangular matrix multiplication can be done in time $O(UVW/\beta)$. A long line of research on fast square matrix multiplication has dropped the complexity to $O(n^\omega)$, where $2 \leq \omega < 3$. The current best-known value is $\omega = 2.3715$ [WXXZ24], but it is believed that the actual value is 2, which would imply that the Boolean matrix product can be computed in time $O(n^{2+o(1)})$. We refer the reader to Section 6 in [BGS20] for a discussion on the topic.

2.3. Problem Statement. Given a CQ Q and an input database D , we want to enumerate the tuples in $Q(D)$ in any order. We will study this problem in the enumeration framework similar to that of [Seg15a], where an algorithm can be decomposed into two phases:

- **Preprocessing:** it builds a data structure that takes space S_p in *preprocessing time* T_p .
- **Enumeration:** it outputs $Q(D)$ with no repetitions. This phase has access to any data structures constructed in the preprocessing phase and can also use additional space of size S_e . The *delay* δ is the maximum time duration between outputting any pair of consecutive tuples (this includes the time to output the first tuple and the time to notify that the enumeration phase has been completed).

We aim to study the tradeoff between the preprocessing time T_p and delay δ for a given CQ Q . We want to achieve a constant delay in linear preprocessing time. As Table 1 shows, when Q is full, after $O(|D|^{\text{fhw}})$ preprocessing time, we can enumerate the results with constant delay [OS16]. When Q is acyclic, i.e., $\text{fhw} = 1$, we can achieve a constant delay with only linear preprocessing time. On the other hand, [BDG07] shows that for every acyclic CQ, we can achieve a linear delay with linear preprocessing time. Recently, [KNOZ20] showed that it is possible to get a tradeoff between the two extremes for the class of hierarchical queries, which are acyclic but not necessarily free-connex. This is the first non-trivial result that improves upon the linear delay guarantees given by [BDG07] for CQs with projections.

Theorem 2.2 [KNOZ20]. *Consider a hierarchical CQ Q with factorization width w , and an input instance D . Then, for any $\epsilon \in [0, 1]$ there exists an algorithm that can preprocess D in time $T_p = O(|D|^{1+(w-1)\epsilon})$ and space $S_p = O(|D|^{1+(w-1)\epsilon})$ such that we can enumerate the query output with $\delta = O(|D|^{1-\epsilon})$ delay and $S_e = O(1)$ space.*

The factorization width w of a CQ, originally introduced as s^\uparrow [OZ15], is a generalization of the **fhw** from Boolean to arbitrary CQs. For a star query Q_k^* , $w = k$. Observe that preprocessing time T_p must always be smaller than the time required to evaluate the full join result. This is because if $T_p = \Theta(|\text{OUT}_\infty|)$, we can evaluate the full join and deduplicate the projection output, allowing us to obtain constant delay in the enumeration phase. Therefore, the tradeoff is more meaningful when ϵ can only take values between 0 and $(\log_{|D|} |\text{OUT}_\infty| - 1)/(w - 1)$. In the worst-case, $|\text{OUT}_\infty| = |D|^w$ and ϵ can take any value between 0 and 1 (both inclusive), which is captured by the general result above.

3. HELPER LEMMAS

Before we present the proof of our main results, we discuss three useful lemmas which will be used frequently and may be of independent interest for enumeration algorithms. The first two lemmas are based on the key idea of *interleaving query results*, which we describe next. We note that the idea of interleaving computation has been explored in the past to develop dynamic algorithms with good worst-case bounds using static data structures [OVL81].

We say that an algorithm \mathcal{A} provides no delay guarantees, meaning its delay guarantee can be as large as its total execution time. In other words, if an algorithm requires time T to complete, its delay guarantee is upper bound by T . Since we use the uniform-cost RAM model, each operation takes one unit of time.

Lemma 3.1. *Consider two algorithms $\mathcal{A}, \mathcal{A}'$ and two positive integers T and T' provided as a part of the input such that*

- (1) \mathcal{A} enumerates query results in total time at most T with no delay guarantees.
- (2) \mathcal{A}' enumerates query results with delay δ and runs in total time at least T' .
- (3) The outputs of \mathcal{A} and \mathcal{A}' are disjoint.

Then, the union of the outputs of \mathcal{A} and \mathcal{A}' can be enumerated with delay $c \cdot \delta \cdot \max\{1, T/T'\}$ for some constant c .

Proof. Let η and γ denote two positive values that will be fixed later. Note that after δ time has passed, we can emit one output result from \mathcal{A}' . But since we also want to compute the output from \mathcal{A} that takes overall time T , we need to slow down the enumeration of \mathcal{A}' sufficiently so we do not run out of output from \mathcal{A}' . This can be done by interleaving the two algorithms in the following way: we run \mathcal{A}' for γ operations, pause \mathcal{A}' , then run \mathcal{A} for η operations, pause \mathcal{A} and resume \mathcal{A}' for γ operations, and so on. The pause and resume takes constant time (say c_{pause} and c_{resume}) in RAM model where the state of registers and program counter can be stored and retrieved enabling pause and resume of any algorithm. Our goal is to find a value of η and γ such that \mathcal{A}' does not terminate until \mathcal{A} has finished. This condition is satisfied when the number of iterations of \mathcal{A}' is equal to the number of iterations of \mathcal{A} . This gives us the condition that,

$$T'/\gamma \leq (\text{Time taken by } \mathcal{A}')/\gamma = (\text{Time taken by } \mathcal{A})/\eta \leq T/\eta$$

Thus, any value of $\eta \leq T \cdot \gamma/T'$ is acceptable. We fix η to be any positive integer constant and then set γ to be the smallest positive integer that satisfies the condition. The

Algorithm 1: DEDUPLICATE(J, \mathcal{A})

Input : Materialized output list J , Algorithm \mathcal{A} with known completion time T
Output : Deduplicated result of \mathcal{A}

```

1  $\delta \leftarrow O(T/J)$ ,  $\text{ptr} \leftarrow 0$ ,  $\text{dedup} \leftarrow 0$ ,  $H \leftarrow \emptyset$ ,  $H' \leftarrow \emptyset$ ;          /* empty hash-set */
2 while  $\text{ptr} < |J|$  do
3   output  $J[\text{ptr}]$ ;          /* output result from  $J$  to maintain delay guarantee */
4   insert  $J[\text{ptr}]$  in  $H'$ 
5    $\text{ptr} \leftarrow \text{ptr} + 1$ ,  $\text{counter} \leftarrow 0$ ;
6   while  $\text{counter} \leq \delta$  do
7     if  $\mathcal{A}$  has not completed then
8       Execute  $\mathcal{A}$  for  $c$  time;          /*  $c$  is a constant */
9       foreach tuple  $t$  generated (if any) do
10        if  $t \notin H'$  and  $t \notin H$  then
11          output  $t$  and insert  $t$  in  $H$ ;
12         $\text{counter} \leftarrow \text{counter} + c$ ;
```

delay is bounded by the product of worst-case number of iterations between two answers of \mathcal{A}' and the work done between each iteration which is $(\delta/\gamma) \cdot (\gamma + \eta + c_{\text{pause}} + c_{\text{resume}}) \leq \delta \cdot (1 + T/T' + (c_{\text{pause}} + c_{\text{resume}})/\gamma) = O(\delta \cdot \max\{1, T/T'\})$. \square

Lemma 3.1 tells us that as long as $T = O(T')$, the output of \mathcal{A} and \mathcal{A}' can be combined without giving up on delay guarantees by pacing the output of \mathcal{A}' . Note that we need to know the exact values of T and T' . This can be accomplished by calculating the number of operations in the algorithms \mathcal{A} and \mathcal{A}' to bound their running time. The next lemma introduces our second key idea of interleaving stored output result with *on-the-fly* query computation. Algorithm 1 describes the detailed algorithm for Lemma 3.2.

Lemma 3.2. *Suppose an algorithm \mathcal{A} can enumerate all output tuples in T time with no delay guarantees, where T is known in advance. Suppose a data structure can enumerate J output tuples without duplication with constant delay. Then, an algorithm exists that enumerates all output tuples with delay guarantee $\delta = O(T/J)$.*

Proof. Let δ be a parameter to be fixed later. Create two empty hash sets H and H' that will be used for deduplication. Using a similar interleaving strategy as above, we use the data structure to obtain one result from J in constant time, insert it into H' , and allow algorithm \mathcal{A} to run for δ time. Whenever \mathcal{A} wants to emit an output tuple, it probes the hash set H and H' , emits t only if t does not appear in H and H' , followed by inserting t in H . Inserting t in H will ensure that \mathcal{A} does not output duplicates³. Each probe and insertion takes $O(1)$ time, so the total running time of \mathcal{A} is $O(T)$. Our goal is to choose δ such that \mathcal{A} terminates before the materialized output J runs out. This condition is satisfied when $\delta \cdot J \geq O(T)$ which gives us $\delta = O(T/J)$. It can be easily checked that no duplicated result is emitted and $O(\delta)$ delay is guaranteed between every pair of consecutive results. Again, observe that we need the algorithm \mathcal{A} to be pausable, which means that we should be able to resume the execution from where we left off. This can be achieved by storing the contents of all registers in the memory and loading it when required to resume execution. \square

³If \mathcal{A} guarantees that it will generate results with no duplicates, then there is no need to use H .

Algorithm 2: MERGE($\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_m$)

```

1  $S \leftarrow \{1, 2, \dots, m\};$ 
2 foreach  $i \in S$  do  $e_i \leftarrow \mathcal{A}_i.first();$ 
3 while  $S \neq \emptyset$  do
4    $w \leftarrow \min_{i \in S} e_i;$       /* finds the smallest output (using  $\preceq$ ) over all algorithms */
5   output  $w;$ 
6   foreach  $i \in S$  do
7     if  $e_i = w$  then  $e_i \leftarrow \mathcal{A}_i.next();$ 
8     if  $e_i = \text{null}$  then  $S \leftarrow S - \{i\};$ 

```

The final helper lemma allows us to enumerate the union of (possibly overlapping) results of m different algorithms where each algorithm outputs its result according to a total order \preceq , such that the union is also enumerated according to \preceq . This lemma is based on the idea presented as Fact 3.1.4 in [Kaz13].

Lemma 3.3. *Consider m algorithms $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_m$ such that each \mathcal{A}_i enumerates its output L_i with delay $O(\delta)$ according to the total order \preceq . Then, the union of their output can be enumerated (without duplicates) with $O(m \cdot \delta)$ delay according to \preceq .*

Proof. See Algorithm 2. For simplicity of exposition, we assume that \mathcal{A}_i outputs a null value when it finishes enumeration. Note that results enumerated by one algorithm are in order. Thus, it always outputs the local minimum result (e_i) over the remaining result to be enumerated. Algorithm 2 goes over all local minimum results over all algorithms and outputs the smallest one (denoted w) as the global minimum result (line 4). Once a result is enumerated, each \mathcal{A}_i must check whether its e_i matches w . If yes, then \mathcal{A}_i needs to update its local minimum result by finding the next one. Then, Algorithm 2 repeats this loop until all algorithms finish enumeration. One distinct result is enumerated in each iteration of the while loop. It takes $O(m)$ time to find the globally minimum result and $O(m \cdot \delta)$ to update all local minimum results (line 6-line 7). Thus, Algorithm 2 has a delay of $O(m \cdot \delta)$. \square

Implied by Lemma 3.3, the *list merge* problem can be done with delay guarantees: Given m lists L_1, L_2, \dots, L_m whose elements are drawn from a common domain, if elements in L_i are distinct (i.e. no duplicates) and ordered according to \preceq , then the union of all lists $\bigcup_{i=1}^m L_i$ can be enumerated in sorted order given by \preceq with delay $O(m)$. Note that the enumeration algorithm \mathcal{A}_i degenerates to going over elements one by one in list L_i , which has $O(1)$ delay guarantee as long as indexes/pointers within L_i are well-built. Throughout the paper, we use this primitive as LISTMERGE(L_1, L_2, \dots, L_m).

4. STAR QUERIES

We now study enumeration algorithms for star queries. We start with a detailed discussion on how our result improves over prior work in Section 4.1, then a warm-up proof for two-path query in Section 4.2, and the complete proof for general star queries in Section 4.3.

Theorem 4.1. *For a star query Q_k^* and any instance D , there is an algorithm with preprocessing time $T_p = O(|D|)$ and preprocessing space $S_p = O(|D|)$, such that we can enumerate $Q_k^*(D)$ with $\delta = O\left(\frac{|D|^{k/k-1}}{|\text{OUT}_\infty|^{1/k-1}}\right)$ delay and $S_e = O(|D|)$ space.*

In the above theorem, the delay depends on the full join result size $|\text{OUT}_\bowtie| = |Q_k^*(D)|$. As the join size increases, the algorithm can obtain better delay guarantees. In the extreme case when $|\text{OUT}_\bowtie| = \Theta(|D|^k)$, it achieves constant delay with linear time preprocessing. In the other extreme, when $|\text{OUT}_\bowtie| = \Theta(|D|)$, it achieves linear delay. When $|\text{OUT}_\bowtie|$ has linear size, we can compute and materialize the query result in linear preprocessing time and achieve constant-delay enumeration. Generalizing this observation, we can always achieve constant delay when T_p is sufficient to evaluate the full join result.




Figure 2: Worst-case tradeoffs given by Theorem 2.2 without (left) and with (right) taking $|\text{OUT}_\bowtie|$ into consideration.

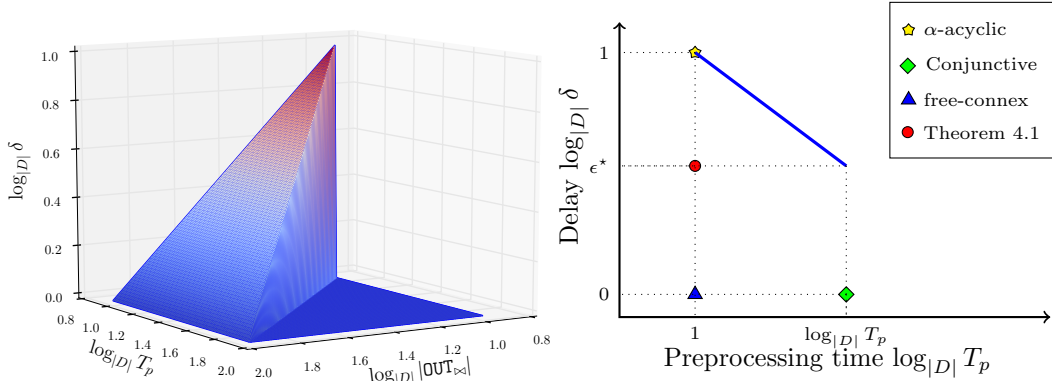


Figure 3: Theorem 4.1 for $k = 2$.

Figure 4: Trade-off in the worst-case for star query. Solid blue line shows Theorem 2.2.

4.1. Comparison with Prior Work. It is instructive now to compare the worst-case delay guarantee obtained by Theorem 2.2 for $Q_k^*(D)$ with Theorem 4.1. Suppose that we want to achieve delay $\delta = O(|D|^{1-\epsilon})$ for some $\epsilon \in [0, (\log_{|D|} |\text{OUT}_\bowtie| - 1)/(k - 1)]$. Theorem 2.2 tells us that this requires $O(|D|^{1+\epsilon(k-1)})$ preprocessing time. Then, it holds that:




Figure 5: $D_0 \cup D_1$ forms a database where Theorem 4.1 improves the delay guarantee of Theorem 2.2.

$$|D|^{1-\epsilon} \geq |D|^{1-\frac{(\log |D| |\text{OUT}_{\bowtie}| - 1)}{k-1}} = |D|^{\frac{k - \log |D| |\text{OUT}_{\bowtie}|}{k-1}} = |D|^{k/k-1} / |\text{OUT}_{\bowtie}|^{1/k-1}$$

In other words, either we have enough preprocessing time to materialize the output and achieve constant delay, or we can achieve the desirable delay with linear preprocessing time.

Figure 2, 3 and 4 show the existing and new tradeoff results. Figure 2 shows the tradeoff curve obtained from Theorem 2.2 by adding $|\text{OUT}_{\bowtie}|$ as a third dimension, and adding the optimization for constant delay when $T_p \geq O(|\text{OUT}_{\bowtie}|)$. Figure 3 shows the tradeoff obtained from our result, while Figure 4 shows other existing results for a fixed value of $|\text{OUT}_{\bowtie}|$. For a fixed value of $|\text{OUT}_{\bowtie}|$, the delay guarantee does not change in Figure 3 as we increase T_p from $|D|$ to $|\text{OUT}_{\bowtie}|$. It remains an open question to decrease the delay further if we allow more preprocessing time. Such an algorithm would correspond to a curve connecting the red point(•) and the green diamond(♦) in Figure 4. Our results thus imply that, depending on $|\text{OUT}_{\bowtie}|$, one must choose a different algorithm to achieve the optimal tradeoff between preprocessing time and delay. Since $|\text{OUT}_{\bowtie}|$ can be computed in linear time (using a simple adaptation of Yannakakis algorithm [Yan81, PP06]), this can be done without affecting the preprocessing bounds.

Next, we show how our result provides an algorithmic improvement over Theorem 2.2. Consider the instances D_0, D_1 depicted in Figure 5a and Figure 5b respectively, and assume we want to use linear preprocessing time (i.e., Theorem 2.2 is used with $\epsilon = 0$).

For D_1 , the algorithm of Theorem 2.2 materializes nothing, since no y valuation has a degree of $O(|D|^0)$, and the delay will be $\Theta(\sqrt{N})$. No materialization also occurs for D_0 , but the delay will be $O(1)$ here. It is easy to check that our algorithm matches the delay in both instances. Now, consider the instance $D = D_0 \cup D_1$. The input size for D is $\Theta(N)$, while the full join size is $N^{3/2} + N^2 = \Theta(N^2)$. The algorithm of Theorem 2.2 will again achieve only a $\Theta(\sqrt{N})$ delay since after the linear time preprocessing, no y valuations can be materialized. In contrast, our algorithm still guarantees a constant delay. This algorithmic improvement results from the careful overlapping of the constant-delay computation, for instance, D_0 , with the computation for D_1 .

The above construction can be generalized as follows. Let $\alpha \in (0, 1)$ be some constant. D_0 remains the same. For D_1 , we construct R to be the cross product of N^α x -values and $N^{1-\alpha}$ y -values, and S to be the cross product of N^α z -values and $N^{1-\alpha}$ y -values. As before, let $D = D_0 \cup D_1$. The input size for D is $\Theta(N)$, while the full join size is $N^{2-\alpha} + N^2 = \Theta(N^2)$.

Hence, our algorithm achieves constant delay with linear preprocessing time. In contrast, the algorithm of Theorem 2.2 achieves $\Theta(N^{1-\alpha})$ delay with linear preprocessing time. In fact, the $\Theta(N^{1-\alpha})$ delay occurs even if we allow $O(N^{1+\epsilon})$ preprocessing time for any $\epsilon < \alpha$. We can similarly show that there also exists an instance where achieving constant delay using Theorem 2.2 requires near quadratic preprocessing time as shown Example 4.2 below.

Example 4.2. We construct an instance where achieving constant delay with Theorem 2.2 would require close to $\Theta(|D|^2)$ computation. Let us fix N to be a power of 2. We will fix $|\mathbf{dom}(x)| = |\mathbf{dom}(z)| = N \log N$. Let D_i be the database constructed by setting $N^\alpha = 2^i$ for $i \in \{1, 2, \dots, \log N\}$ where relation R is the cross product of N^α x -values and $N^{1-\alpha}$ y -values, and S is the cross product of N^α z -values and $N^{1-\alpha}$ y -values. We also construct a database D^* consisting of a single y connected to all x and z values. Let $D = D^* \cup D_1 \cup D_2 \cup \dots \cup D_{\log N}$. It is now easy to see that $|D| = N \cdot \log N$, $|\mathbf{dom}(y)| = \sum_\alpha N^{1-\alpha} \leq 2N = \Theta(|D|/\log |D|)$ and $|\mathbf{OUT}_\bowtie| = \sum_\alpha N^{1+\alpha} + N^2 \log^2 N = \Theta(|D|^2)$. On this instance, Theorem 2.2 achieves $\Theta(|D|/\log |D|)$ after linear time preprocessing. Suppose we target constant-delay enumeration. Let us fix this constant as c^* (a power of 2) for simplicity. Then, we need enough preprocessing time to materialize the join result of all database instances D_i where $i \in \{1, 2, \dots, \log(N/c^*)\}$ to ensure that the number of heavy y values that remain is at most c^* . Hence, $T_p > \sum_{i \in \{1, 2, \dots, \log(N/c^*)\}} N \cdot 2^i > N^2/c^* = \Theta(|D|^2/\log |D|)$. Theorem 2.2 requires near quadratic computation to achieve constant-delay enumeration.

4.2. Warm-up: Two-Path Query. As warm-up, we show an algorithm for $Q_{\text{two-path}} = \pi_{x,z}(R(x, y) \bowtie S(y, z))$ that achieves $O(|D|^2/|\mathbf{OUT}_\bowtie|)$ delay with linear preprocessing time. At a high level, we will decompose the join into two subqueries with disjoint outputs. The subqueries will be generated based on whether a valuation for x is *light* or not based on its degree in relation R . For all light valuations of x (degree at most δ), we will show that their enumeration is achievable with delay δ . For the heavy x valuations, we will show that they also can be computed *on-the-fly* while maintaining the delay guarantees.

Preprocessing Phase. We first process the input relations by removing dangling⁴ tuples. During the preprocessing phase, we will store the input relations as a hash map and sort the valuations for x in increasing order of their degree. Using any comparison based sorting technique requires $\Omega(|D| \log |D|)$ time in general. Thus, if we wish to remove the $\log |D|$ factor, we must use non-comparison-based sorting algorithms. In this paper, we will use count sort [CLRS09], which has complexity $O(|D| + r)$ where r is the range of the non-negative key values. However, we need to ensure that all relations in the database D satisfy the bounded range requirement. This can be easily accomplished by introducing a bijective function $f : \mathbf{dom}(D) \rightarrow \{1, 2, \dots, |D|\}$ that maps all values in the active domain of the database to some integer between 1 and $|D|$ (both inclusive). Both f and its inverse f^{-1} can be stored as hash tables: suppose there is a counter $c \leftarrow 1$. We perform a linear pass over the database and check if some value $v \in \mathbf{dom}(D)$ has been mapped or not (by checking if there exists an entry $f(v)$). If not, we set $f(v) = c$, $f^{-1}(c) = v$ and increment c . Once the hash tables f and f^{-1} have been created, we modify the input relation R (and S similarly) by replacing every tuple $t \in R$ with tuple $t' = f(t)$. Since the mapping is a relabeling scheme, such a transformation preserves the degree of all the values. The codomain of f is also equipped

⁴A tuple $t \in R_\ell(\mathbf{x}_\ell)$ is said to be dangling if $R_1(\mathbf{x}_1) \bowtie \dots \bowtie R_{\ell-1}(\mathbf{x}_{\ell-1}) \bowtie t \bowtie R_{\ell+1}(\mathbf{x}_{\ell+1}) \bowtie \dots \bowtie R_n(\mathbf{x}_n)$ is empty.

Algorithm 3: ENUMTWOPATH

```

1 for  $i = 1, \dots, i^*$  do
2   Let  $\pi_y(\sigma_{x=v_i}(R)) = \{u_1, u_2, \dots, u_\ell\}$ ;
3   output  $(v_i, f^{-1}(\text{LISTMERGE}(\pi_{z\sigma_{y=u_1}}S, \pi_{z\sigma_{y=u_2}}S, \dots, \pi_{z\sigma_{y=u_\ell}}S)))^5$ 
4 for  $i = i^* + 1, \dots, n$  do
5   Let  $\pi_y(\sigma_{x=v_i}(R)) = \{u_1, u_2, \dots, u_\ell\}$ ;
6   output  $(v_i, f^{-1}(\text{LISTMERGE}(\pi_{z\sigma_{y=u_1}}S, \pi_{z\sigma_{y=u_2}}S, \dots, \pi_{z\sigma_{y=u_\ell}}S)))$ 

```


run for $O(1)$ time
then switchrun for $O(1)$ time
then switch

Figure 6: Example for two path query enumeration

with a total order \preceq (we will use \leq). Note that f is not an order-preserving transformation in general, but our algorithms do not require this property.

Next, for every tuple $t \in R(x, y)$, we create a hash map with key $\pi_x(t)$, and the value is a list to which $\pi_y(t)$ is appended; and for every tuple $t \in S(y, z)$, we create a hash map with key $\pi_y(t)$ and the value is a list to which $\pi_z(t)$ is appended. For the second hash map, we sort the value list using sort order \preceq for each key once each tuple $t \in S(y, z)$ has been processed. Finally, we sort all values in $\pi_x(R)$ in increasing order of their degree in R (i.e., $|\sigma_{x=v_i}R(x, y)|$ is the sort key). Let $\mathcal{L} = \{v_1, \dots, v_n\}$ denote the ordered set of these values sorted by their degree and let d_1, \dots, d_n be their respective degrees. Creating the sorted list \mathcal{L} takes $O(|D|)$ time since the degrees d_i satisfy the bounded range requirement (i.e. $1 \leq d_i \leq |D|$). Next, we identify the smallest index i^* such that

$$\sum_{v: \{v_1, v_2, \dots, v_{i^*}\}} |R(v, y) \bowtie S(y, z)| \geq \sum_{v: \{v_{i^*+1}, \dots, v_n\}} |R(v, y) \bowtie S(y, z)| \quad (4.1)$$

This can be computed by doing a linear pass on \mathcal{L} using a simple adaptation of Yannakakis algorithm [Yan81, PP06]. This entire phase takes time $O(|D|)$.

Enumeration Phase. The enumeration algorithm interleaves the following two loops using the construction in Lemma 3.1. Specifically, it will spend an equal amount of time (a constant) before switching to the computation of the other loop. The algorithm alternates between low-degree and high-degree values in \mathcal{L} . The main idea is that, for a given $v_i \in \mathcal{L}$, we can enumerate the result of the subquery $\sigma_{x=v_i}(Q_{\text{two-path}})$ with delay $O(d_i)$. This can be accomplished by Algorithm 2, since the subquery is equivalent to list merging.

⁵Abusing the notation, $f^{-1}(\mathcal{B})$ for some ordered list (or tuple) \mathcal{B} returns an ordered list (tuple) \mathcal{B}' where $\mathcal{B}'(i) = f^{-1}(\mathcal{B}(i))$.

Example 4.3. Consider relations R and S as shown in Figure 6a and Figure 6b. Figure 6c shows the sorted valuations a_2 and a_1 by their degree and the valuations for Z as sorted lists $S[b_1]$, $S[b_2]$ and $S[b_3]$. For both a_1 and a_2 , the pointers point to the head of the lists. We will now show how $\text{LISTMERGE}(S[b_1], S[b_2], S[b_3])$ is executed for a_1 . Since three sorted lists need to be merged, the algorithm finds the smallest valuation across the three lists. c_1 is the smallest valuation and the algorithm outputs (a_1, c_1) . Then, we need to increment pointers of all lists pointing to c_1 ($S[b_1]$ is the only list containing c_1). Figure 6d shows the state of pointers after this step. The pointer for $S[b_1]$ points to c_2 , and all other pointers still point to the list head. Next, we merge the list by finding each list's smallest valuation. Both $S[b_1]$ and $S[b_2]$ pointers are pointing to c_2 and the algorithm outputs (a_1, c_2) . The pointers for both $S[b_1]$ and $S[b_2]$ are incremented, and the enumeration for both lists is complete as shown in Figure 6e. In the last step, only $S[b_3]$ list remains and we output (a_1, c_3) and increment the pointer for $S[b_3]$. All pointers are now past the end of the lists and the enumeration is now complete.

Theorem 4.4. For $Q_{\text{two-path}}$ and any instance D , there is an algorithm with preprocessing time $T_p = O(|D|)$ and preprocessing space $S_p = O(|D|)$, such that the query result $Q_{\text{two-path}}(D)$ can be enumerated with delay $\delta = O(|D|^2/|\text{OUT}_{\bowtie}|)$ and space $S_e = O(|D|)$.

Proof. To prove this result, we will apply Lemma 3.1, where \mathcal{A}' is the first loop (the one with light-degree values), and \mathcal{A} is the second loop (the one with high-degree values).

Let δ denote the degree of the valuation v_{i^*} . First, we claim that the delay of \mathcal{A}' will be $O(\delta)$. Indeed, LISTMERGE will output a result every $O(\delta)$ time since the degree of each valuation in the first loop is at most δ . Let $J_h = \sum_{i > i^*} |R(v_i, y) \bowtie S(y, z)|$ and $J_\ell = \sum_{i \leq i^*} |R(v_i, y) \bowtie S(y, z)|$. Then, \mathcal{A}' runs in time at least J_ℓ , and \mathcal{A} in time at most $c^* \cdot J_h$. Here, c^* is an upper bound on the number operations in each loop iteration in Algorithm 3. Since by construction $J_\ell \geq J_h$, Lemma 3.1 obtains a total delay of $O(\delta)$.

It now remains to bound δ . First, as i^* is the smallest index that satisfies Equation 4.1, it must be that $J_\ell - J_h \leq |D|$ (if not, shifting the smallest index by one decreases the LHS by at most $|D|$ and increases the RHS by at most $|D|$ while still satisfying the condition that $J_\ell \geq J_h$). Combined with the fact that $J_\ell + J_h = |\text{OUT}_{\bowtie}|$, we get that $J_h \geq |\text{OUT}_{\bowtie}|/2 - |D|/2 \geq 1/4 \cdot |\text{OUT}_{\bowtie}|$ assuming $|\text{OUT}_{\bowtie}| \geq 2 \cdot |D|$. Finally, $J_h \leq |D|^2/\delta$ since there are at most $|D|/\delta$ heavy values, and each heavy value can join with at most $|D|$ tuples for the full join. Combining the two inequalities gives us the claimed guarantee. \square

The reader should note that the delay of $\delta = O(|D|^2/|\text{OUT}_{\bowtie}|)$ is only an upper bound. Depending on the skew in the database instance, Algorithm 3 may achieve much better delay guarantees in practice, as shown in Example 4.5 below.

Example 4.5. Consider relation $R(x, y)$ of size $O(N)$, with values v_1, \dots, v_N in x . Suppose that each of v_1, \dots, v_{N-1} has a degree exactly 1, and each is connected to a distinct value in y . Also, v_N has degree $N - 1$ and is connected to all $N - 1$ values in y . Let relation $S(z, y) = \{t \in R(x, y)\}$ (i.e., relation S is a copy of relation R). Suppose we want to compute $Q_{\text{two-path}}$. It is easy to see that $\text{OUT}_{\bowtie} = \Theta(N)$. Thus, applying the bound of $\delta = O(N^2/|\text{OUT}_{\bowtie}|)$ gives us $O(N)$ delay. However, Algorithm 3 achieves a constant delay guarantee since all of v_1, \dots, v_{N-1} are processed by the left pointer in $O(1)$ delay as they produce exactly one output result, while the right pointer processes v_N on the fly in $O(N)$ time.

4.3. Proof of Theorem 4.1. We are ready to prove Theorem 4.1 formally. We now generalize Algorithm 3 for any star query. At a high level, we will decompose the join query Q_k^* into a union of $k + 1$ subqueries whose output is a partition of the result of the original query. These subqueries will be generated based on whether a value for some x_i is *light* or not. We will show that if any of the values for x_i are light, the enumeration delay is small. The $(k + 1)$ -th subquery will contain heavy values for all attributes. Our key idea again is to interleave the join computation of the *heavy* subquery with the remaining light subqueries.

Preprocessing Phase. Suppose the dangling tuples have been removed from all input relation, which can be achieved in linear time [Yan81]. The full join size $|\text{OUT}_\bowtie|$ can also be computed in linear time. Similar to the preprocessing phase in the previous section, we construct the hash tables f, f^{-1} to perform the domain compression and modify all the input relations by replacing tuple t with $f(t)$. Set $\Delta = (2 \cdot |D|^k / |\text{OUT}_\bowtie|)^{\frac{1}{k-1}}$. For each relation R_i , a value v for attribute x_i is *heavy* if its degree (i.e. $|\pi_{y \sigma_{x_i=v}} R(x_i, y)|$) is greater than Δ , and *light* otherwise. Moreover, a tuple $t \in R_i$ is identified as heavy or light depending on whether $\pi_{x_i}(t)$ is heavy or light. In this way, each relation R is divided into two relations R^h and R^ℓ , containing heavy and light tuples, respectively, in time $O(|D|)$. The original query can be decomposed into subqueries of the following form: $\pi_{x_1, x_2, \dots, x_k}(R_1^? \bowtie R_2^? \bowtie \dots \bowtie R_k^?)$, where $?$ can be either h, ℓ or \star . Here, R_i^\star denotes the original relation R_i . However, care must be taken to generate the subqueries so that there is no overlap between the output of any subquery. In order to do so, we create k subqueries of the form

$$Q_i = \pi_{x_1, \dots, x_k}(R_1^h \bowtie \dots \bowtie R_{i-1}^h \bowtie R_i^\ell \bowtie R_{i+1}^\star \bowtie \dots \bowtie R_k^\star)$$

In Q_i , relation R_i has superscript ℓ , all relations R_1, \dots, R_{i-1} have superscript h and relations R_{i+1}, \dots, R_k have superscript \star . The $(k + 1)$ -th query with all $?$ as h is denoted by Q_H . Note that each output tuple t is generated by exactly one of the Q_i ; thus, the output of all subqueries is disjoint. This implies that each $f^{-1}(t)$ is also generated by exactly one subquery. Similar to the preprocessing phase of the two path query, we store all R_i^ℓ and R_i^h in hashmaps where the values in the maps are lists sorted in lexicographic order.

Enumeration Phase. We next describe how enumeration is performed. We will show that for $Q_L = Q_1 \cup \dots \cup Q_k$, the result can be enumerated with delay $O(\Delta)$. As Q_H contains all heavy valuations from all relations, we compute its join *on-the-fly* by alternating between some subquery in Q_L and Q_H . This will ensure that we can give some output to the user with delay guarantees and also make progress on computing the full join of Q_H . Our goal is to reason about the running time of enumerating Q_L (denoted by T_L) and the running time of Q_H (denoted by T_H) and make sure that while we compute Q_H , we do not run out of the output from Q_L . Next, we introduce the algorithm that enumerates output for any specific valuation v of attribute x_i , described in Lemma 4.6. This algorithm can be viewed as another instantiation of Algorithm 2.

Lemma 4.6. *For a star query Q_k^* , any instance D , and an arbitrary value $v_i \in \text{dom}(x_i)$ with degree d in $R_i(x_i, y)$, its query result $\sigma_{x_i=v_i} Q_k^*(D)$ can be enumerated with $O(d)$ delay.*

Proof. Consider some tuple $(v_i, u) \in R_i$. Each u is associated with a list of valuations over attributes $(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k)$, which is a Cartesian product of $k - 1$ sub-lists $\sigma_{y=u} R_j(x_j, y)$. Note that such a list is not materialized as that for a two-path query but is present in a factorized form. We next define the enumeration algorithm \mathcal{A}_u for each $u \in \pi_{y \sigma_{x_i=v_i}} R_i(x_i, y)$, with lexicographical ordering of attributes $(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k)$.

Note that elements in each list $\pi_{x_j} \sigma_{y=u} R_j^?(x_j, y)$ can be enumerated with $O(1)$ delay. Then, \mathcal{A}_u enumerates all results in $\times_{j \neq i: j \in \{1, 2, \dots, k\}} \sigma_{y=u} R_j^?(x_j, y)$ by $k - 1$ level of nested loops in lexicographic order $(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k)$, which has $O(k - 1) = O(1)$ delay. After applying Algorithm 2, we can obtain an enumeration algorithm that enumerates the union of query results over all neighbors with $O(d)$ delay guarantee. For each output tuple t generated by Algorithm 2, we return $f^{-1}(t)$ to the user. \square

Let c^* be an upper bound on the number of operations in each iteration of LISTMERGE. This can be calculated by counting the number of operations in the exact implementation of the algorithm. Directly implied by Lemma 4.6, the result of any subquery in Q_L can be enumerated with delay $O(\Delta)$. Let Q_H^* denote the corresponding full query of Q_H , i.e., the head of Q_H^* also includes the variable y (Q_L^* is defined similarly). Then, Q_H^* can be evaluated in time $T_H \leq c^* \cdot |Q_H^*| \leq c^* \cdot |\text{OUT}_\infty|/2$ by using LISTMERGE on subquery Q_H . This follows from the bound $|Q_H^*| \leq |D| \cdot (|D|/\Delta)^{k-1}$ and our choice of $\Delta = (2 \cdot |D|^k / |\text{OUT}_\infty|)^{\frac{1}{k-1}}$. Since $|Q_H^*| + |Q_L^*| = |\text{OUT}_\infty|$, it holds that $|Q_L^*| \geq |\text{OUT}_\infty|/2$ given our choice of Δ . Also, the running time T_L is lower bounded by $|Q_L^*|$ (since we need at least one operation for every result). Thus, $T_L \geq |\text{OUT}_\infty|/2$. We are now ready to apply Lemma 3.1 with the following parameters:

- (1) \mathcal{A} is the full join computation of Q_H and $T = c^* \cdot |\text{OUT}_\infty|/2$.
- (2) \mathcal{A}' is the enumeration algorithm applied to Q_L with delay guarantee $\delta = O(\Delta)$ and $T' = |\text{OUT}_\infty|/2$.
- (3) T and T' are fixed once $|\text{OUT}_\infty|$, Δ , and the constant c^* are known.

By construction, the outputs of Q_H and Q_L are also disjoint. Thus, the conditions of Lemma 3.1 apply, and we obtain a delay of $O(\Delta)$. \square

4.4. Interleaving with Join Computation. Theorem 4.1 obtains poor delay guarantees when the full join size $|\text{OUT}_\infty|$ is close to input size $|D|$. In this section, we present an alternate algorithm that provides good delay guarantees. The algorithm is an instantiation of Lemma 3.2 on the star query, which degenerates to computing as many distinct output results as possible in limited preprocessing time. An observation is that for each valuation u of attribute y , the Cartesian product $\times_{i \in \{1, 2, \dots, k\}} \pi_{x_i} \sigma_{y=u} R_i(x_i, y)$ is a subset of output results without duplication. Thus, this subset of the output result is readily available since no deduplication needs to be performed. Similarly, after all relations are reduced, it is also guaranteed that each valuation of attribute x_i of relation R_i generates at least one output result. Thus, $\max_{i=1}^k |\text{dom}(x_i)|$ results are also readily available that do not require deduplication. We define J as the larger of the two quantities, i.e., $J = \max \left\{ \max_{i=1}^k |\text{dom}(x_i)|, \max_{u \in \text{dom}(y)} \prod_{i=1}^k |\sigma_{y=u} R_i(x_i, y)| \right\}$. Together with these observations, we can achieve the following theorem.

Theorem 4.7. *For a star query Q_k^* and any instance D , there exists an algorithm with preprocessing time $O(|D|)$ and space $O(|D|)$, such that the query result $Q_k^*(D)$ can be enumerated with delay $\delta = O\left(|\text{OUT}_\infty|/|\text{OUT}_\pi|^{1/k}\right)$ and space $S_e = O(|D|)$.*

In the above theorem, we obtain delay guarantees that depend on both the size of the full join result OUT_∞ and the projection output size OUT_π . However, one does not need to know $|\text{OUT}_\infty|$ or $|\text{OUT}_\pi|$ to apply the result. We first compare the result with Theorem 4.1.

First, observe that both Theorem 4.7 and Theorem 4.1 require $O(|D|)$ preprocessing time. Second, the delay guarantee provided by Theorem 4.7 can be better than Theorem 4.1. This happens when $|\text{OUT}_\infty| \leq |D| \cdot J^{1-1/k}$, a condition that can be easily checked in linear time.

We now proceed to describe the algorithm. First, we compute all the statistics for computing J in linear time. If $J = |\text{dom}(x_j)|$ for some integer $j \in \{1, 2, \dots, k\}$, we just materialize one result for each valuation of x_j . Otherwise, $J = \prod_{i=1}^k |\sigma_{y=u} R_i(x_i, y)|$ for some valuation u in attribute y . Note that we do not need to materialize the Cartesian product explicitly but only need to store the tuples $\sigma_{y=u} R_i(x_i, y)$ for all $i \in \{1, \dots, k\}$. As mentioned before, each output in $\times_{i=1}^k (\pi_{x_i} \sigma_{y=u} R_i(x_i, y))$ can be enumerated with $O(1)$ delay and clearly does not have any duplicate output. This preprocessing phase takes $O(|D|)$ time and $O(|D|)$ space. We can now invoke Lemma 3.2 to achieve the claimed delay. The final observation is to express J in terms of $|\text{OUT}_\pi|$. Note that $|\text{OUT}_\pi| \leq \prod_{i \in [k]} |\text{dom}(x_i)|$ which implies that $\max_{i \in [k]} |\text{dom}(x_i)| \geq |\text{OUT}_\pi|^{1/k}$. Thus, it holds that $J \geq |\text{OUT}_\pi|^{1/k}$ which gives us the desired bound on the delay guarantee.

4.5. Dynamic setting. Remarkably, it is also possible to do a simple adaptation of the algorithm for the dynamic setting but only in the case of self-joins, where $R_1 = R_2 = \dots R_k = R$. In this setting, single tuple updates are allowed to the underlying relation, i.e., a tuple t over variables x, y can be either inserted or deleted from the relation $R(x, y)$ (note that a relation does not allow duplicates).

We will store the relation $R(x, y)$ in a bidirectional hash table (i.e., given a u , we can retrieve $\pi_y \sigma_{x=u} R$ and $|\pi_y \sigma_{x=u} R|$, and given a v , we can retrieve $\pi_x \sigma_{y=v} R$ and $|\pi_x \sigma_{y=v} R|$). We will also use a max heap structure that admits constant amortized time insertion and updates (i.e., changing the weight of a key in the heap), and amortized $O(\log N)$ time to remove the element at the top of the heap with N elements. Fibonacci heap [FT87] is an example of such a data structure.

Algorithm 4 shows the procedures for preprocessing, update, and enumeration. In the preprocessing phase, we populate the set J , which stores precomputed results that will be used for interleaving, and compute the full join size and store it in the `joinsize` variable. The hash maps storing the relation can be maintained in constant time under single-tuple updates for an update made to the relation. The set J is also maintained by adding a new tuple whenever a new domain value for x is seen and removing a tuple (a, a) from J whenever $|\sigma_{x=a} R|$ becomes zero.

For the enumeration phase, we apply Lemma 3.2 with stored hash set J , since the hash set is a data structure that allows constant-delay enumeration of its contents and does not contain any duplicates. We can now state the result formally.

Theorem 4.8. *For a self-join star query Q_k^* and any single relation R , there exists an algorithm with preprocessing time $O(|R|)$ and preprocessing space $O(|R|)$, such that $Q_k^*(R)$ can be enumerated with delay $\delta = O\left(|\text{OUT}_\infty|/|\text{OUT}_\pi|^{1/k}\right)$ and enumeration space $S_e = O(|R|)$ with $O(\log |R|)$ amortized update time for single-tuple updates.*

Proof. It is easy to see that the procedure `PREPROCESS` requires $O(|R|)$ time and correctly computes the full join size `fulljoinsize` as $\sum_{v \in \text{dom}(y)} |\sigma_{y=v} R|^k$. Each y valuation in the active domain is inserted into the heap, taking constant time per insertion. Each operation in the procedure `UPDATE` also requires constant time by our assumption on the insertion and

Algorithm 4: DYNAMICSTARQUERYSELFJOIN(Q_k^*, R)

```

1 fulljoinsize  $\leftarrow 0$ ,  $J \leftarrow \emptyset$ , heap  $\leftarrow \emptyset$ ;                                /* empty hash set and heap */
2 procedure PREPROCESS()
3   foreach  $t \in R$  do
4     add  $(\pi_x t, \pi_x t)$  to  $J$ 
5     fulljoinsize  $\leftarrow$  fulljoinsize +  $|\sigma_{y=\pi_y t} R|^k$ ;
6     add key  $\pi_y t$  with weight  $|\sigma_{y=\pi_y t} R|^k$  to heap
7 procedure UPDATE( $t, \text{op}$ )
8   oldcontribution  $\leftarrow |\sigma_{y=\pi_y t} R|^k$ 
9   newcontribution  $\leftarrow 0$ 
10  if  $\text{op} = \text{insert}$  then
11    insert  $t$  in  $R$ ;                                /* insert in bidirectional hash tables */
12    add  $(\pi_x t, \pi_x t)$  to  $J$ 
13    newcontribution  $\leftarrow |\sigma_{y=\pi_y t} R|^k$ 
14  if  $\text{op} = \text{delete}$  then
15    delete  $t$  from  $R$                                 /* delete from bidirectional hash tables */
16    if  $|\sigma_{x=\pi_x t} R| = 0$  then delete  $(\pi_x t, \pi_x t)$  from  $J$ 
17    fulljoinsize  $\leftarrow$  fulljoinsize + newcontribution – oldcontribution;
18    update key  $\pi_y t$  to weight  $|\sigma_{y=\pi_y t} R|^k$  in heap
19    Rebuild heap by removing all keys with weight 0 after every  $O(|R|)$  updates
20 procedure ENUMERATE()
21    $s \leftarrow$  key at the top of the heap
22   Let  $J'$  be the set with larger cardinality out of  $J$  or  $\times_{i \in [k]} \pi_{x_i} \sigma_{y=s} R_i$ 
23   Call Lemma 3.2 with stored hash set  $J'$ ,  $\mathcal{A}$  as any worst-case optimal join
    algorithm on the input query, and  $T = c^* \cdot \max\{\text{fulljoinsize}, |\times_{i \in [k]} \pi_{x_i} \sigma_{y=s} R_i|\}$ 
    /*  $c^*$  is the constant in  $O(\cdot)$  of the worst-case optimal join algorithm */

```

deletion time complexity of the hash table and the hash set. To maintain the full join size correctly, we need to consider the new degree of the y valuation after inserting/deleting t in R . Line 17 does exactly that by subtracting the contribution of the y valuation before the update and adding the new contribution of the y valuation to the full join size after the update. Since we also need to update the weight of $\pi_y t$ in the heap, line 18 updates the weight of the key appropriately in constant time. Finally, procedure ENUMERATE invokes Lemma 3.2. The main idea here is to pick the larger of J or $|\times_{i \in [k]} \pi_{x_i} \sigma_{y=s} R_i|$ so that the enumeration delay is smaller by larger output that can be interleaved with the query execution. Finding the value s (which is at the top of the heap) is a constant time operation. Applying the same reasoning as Theorem 4.7, we can get desired delay $\delta = O(|\text{OUT}_\infty|/|\text{OUT}_\pi|^{1/k})$.

To reason about the amortized update time guarantee, note that on line 19, we rebuild the heap after every $O(|R|)$ updates operations. The step is important to make sure that the space usage of the heap remains linear in the size of the input. The rebuilding step repeatedly pops all entries from the heap and then creates the heap with only the entries that have non-zero weight. Thus, the step requires $O(|R| \log |R|)$ time but since it is done only once after every $O(|R|)$ update operations, the cost amortized cost is $O(\log |R|)$. Since

all other update operations are constant time, the heap rebuilding dominates the overall update guarantee. \square

To achieve constant update time, the main result that applies to all hierarchical queries (including self-joins) from [KNOZ20] requires linear preprocessing time and achieves worst-case linear delay. Since $|\text{OUT}_\bowtie|/|\text{OUT}_\pi|^{1/k} \leq |D|$ [AP09], our strategy provides an alternate algorithm that achieves the same worst-case linear delay (with $O(\log |D|)$ update time instead of constant update time) but may perform better on certain database instances.

Optimality for dynamic self-joins. Next, we show that the delay guarantee obtained by the algorithm above is optimal conditioned on the popular Boolean matrix multiplication (BMM) conjecture [WW18]: In the Word RAM model with words of $O(\log n)$ bits, given two $n \times n$ Boolean matrices, there exists no combinatorial algorithm for multiplying the matrices in $O(n^{3-\epsilon})$ time for any constant $\epsilon > 0$. BMM conjecture is also hard when both input matrices are identical⁶. Although the term combinatorial is not formally defined, it refers to an algorithm that does not use any algebraic techniques (such as the ones used for fast matrix multiplication).

Lemma 4.9. *For a self-join star query Q_2^* , any single relation R and parameter $\delta = O(|\text{OUT}_\bowtie|/\sqrt{|\text{OUT}_\pi|})$, there exists no combinatorial algorithm with preprocessing time $O(|R|)$ and space $O(|R|)$, such that the query result $Q_2^*(R)$ can be enumerated with delay $O(\delta^{1-\epsilon})$, for any constant $\epsilon > 0$, with $O(\log |R|)$ amortized update time for single-tuple updates, assuming the Boolean matrix multiplication conjecture.*

Proof. Consider the input matrix M for the BMM problem where we want to compute the product of M with itself. As shown below, we construct a matrix M' of size $2n \times 2n$.

$$M' = \begin{bmatrix} \mathbf{0} & M \\ \mathbf{1} & \mathbf{0} \end{bmatrix}$$

Here, $\mathbf{0}$ and $\mathbf{1}$ represent a $n \times n$ matrix with all entries as 0 and 1 respectively. We will now create a relation $R(x, y)$ from M' by calling the UPDATE procedure for every non-zero entry (i, j) of M' . Since there are $O(n^2)$ entries and every entry is processed in $O(\log n)$ amortized time, the total time required is $O(n^2 \log n)$. Next, observe that due to the $\mathbf{1}$ matrix in the lower left half of M' , there is a y value (in fact, there are n such y values) that generates $\Theta(n^2)$ amount of output. Also, note that $|\text{OUT}_\bowtie| = \Theta(n^3)$. Therefore, Algorithm 4 would provide a delay guarantee of $O(n^3/n^2) = O(n)$. Suppose an algorithm could enumerate the result with $O(n^{1-\epsilon})$ delay. The total output of the join can be as large as $O(2n \cdot 2n) = O(n^2)$. The original output of multiplying M with itself can be recovered by only looking at output tuples of the join result (z, w) where both $z \leq n$ and $w \leq n$. Thus, the entire output can be enumerated in total time $O(n^2 \cdot n^{1-\epsilon}) = O(n^{3-\epsilon})$ time, which contradicts the BMM conjecture. \square

⁶Given two $n \times n$ matrices M_1 and M_2 , we can create a matrix

$$M = \begin{bmatrix} \mathbf{0} & M_1 \\ M_2 & \mathbf{0} \end{bmatrix}$$

to show that if multiplication of M with itself can be achieved in $O(n^{3-\epsilon})$ time, then BMM can also be solved in $O(n^{3-\epsilon})$ time. Here $\mathbf{0}$ is an $n \times n$ matrix with all 0 entries.

Why non-self-join is hard. At this point, the reader may wonder what breaks down for the non-self-join case. First, note that our lower bound argument directly extends to the non-self-join case as well. However, in Algorithm 4, we can no longer maintain J under updates in the non-self-join case. Indeed, in the presence of self-joins, each $(a, b) \in R$ guarantees that (a, a, \dots, a) is a part of the output. However, no such claim can be made when the relations are independent. Algorithm 4 keeps track of the y valuation that generates the largest amount of output by using the heap and use it for interleaving (see line 22 where pick the larger of J and the y valuation that generates the most output), but in that case, we do not get any closed form output dependent expression for the delay guarantee. The observation that self-joins can make the evaluation of queries easier (but in the static setting) has also been made by prior work [CS23]. Thus, constructing an algorithm for the non-self-join that leads to an output-sensitive guarantee is a non-trivial problem that likely requires the development of new technical machinery.

4.6. Fast Matrix Multiplication. So far, we have focused on combinatorial algorithms. We next show how fast matrix multiplication can obtain a tradeoff between preprocessing time and delay better than Theorem 2.2 for some delay values.

Theorem 4.10. *For a star query Q_k^* and any instance D , there exists an algorithm with pre-processing time $T_p = O((|D|/\delta)^{\omega+k-2})$ and space $S_p = O((|D|/\delta)^{\omega+k-2})$ such that the query result $Q_k^*(D)$ can be enumerated with delay $O(\delta)$ and enumeration space $S_e = O(|D|)$, for $1 \leq \delta \leq |D|^{(\omega+k-3)/(\omega+2\cdot k-3)}$.*

Proof. Let δ be the degree threshold for deciding whether a valuation is heavy or light. We can partition the original query into the subqueries of the form:

$$\pi_{x_1, \dots, x_k}(R_1(x_i^?, y^?) \bowtie R_2(x_2^?, y^?) \bowtie \dots \bowtie R_k(x_k^?, y^?))$$

where $?$ can be either h, ℓ or \star . We will reuse the definition of heavy and light values for variables x_i and y as defined in Section 4.3. The first observation is that the result of the k subqueries where superscript of $x_1 \dots x_{i-1}$ is h , superscript for x_i is ℓ , and it is \star for the remaining variables (including y), can be enumerated with $O(\delta)$ delay using the LISTMERGE subroutine for each valuation v of x_i^ℓ that is light. Partitioning of each relation R_i into $R_i(x_i^h, y^\star)$ and $R_i(x_i^\ell, y^\star)$, as well as their indexing into a bidirectional hash table, can be done in linear time. These k queries can be directly enumerated in the enumeration phase with the desired delay.

For the remaining queries, we need to perform preprocessing. The first subquery to process is Q' with superscript for all $x_i, i \in [k]$ is \star and it is ℓ for y . We define a valuation v for variable y to be light if $|\pi_{x_i \sigma_{y=v}} R(x_i, y)| \leq \delta$ for each $i \in \{1, \dots, k\}$. Using any worst-case optimal join algorithm, this subquery can be evaluated in time $O(\sum_{\text{all light valuations } v} |\pi_{x_i \sigma_{y=v}} R(x_i, y)|) \leq O(\delta^{k-1} \cdot \sum_{\text{all light valuations } v} |\pi_{x_i \sigma_{y=v}} R(x_k, y)|) = O(|D| \cdot \delta^{k-1})$.

The second query we need to preprocess is the subquery Q'' , where all variables are heavy, i.e., superscript for $x_i, i \in [k]$ and y is h . For this subquery, we will use matrix multiplication using the construction outlined in [DHK20]. Specifically, we first create a rectangular matrix M of dimensions $(|D|/\delta)^{\lceil k/2 \rceil} \times |D|/\delta$ and a matrix N of dimensions $|D|/\delta \times (|D|/\delta)^{\lfloor k/2 \rfloor}$ such that,

$$M_{(a_1, a_2, \dots, a_{\lceil k/2 \rceil}), b} = \begin{cases} 1, & (a_1, b) \in R_1, \dots, (a_{\lceil k/2 \rceil}, b) \in R_{\lceil k/2 \rceil} \\ 0, & \text{otherwise} \end{cases}$$

$$N_{b, (a_{\lceil k/2 \rceil+1}, \dots, a_k)} = \begin{cases} 1, & (a_{\lceil k/2 \rceil+1}, b) \in R_{\lceil k/2 \rceil+1}, \dots, (a_k, b) \in R_k \\ 0, & \text{otherwise} \end{cases}$$

Observe that the non-zero entries in the product of matrices M and N correspond to the output of the subquery Q'' . In other words, a non-zero entry for row $(a_1, a_2, \dots, a_{\lceil k/2 \rceil})$ and column $(a_{\lceil k/2 \rceil+1}, \dots, a_k)$ implies that there exists a value b , such that $(a_i, b) \in R_i$ for each input relation and thus $(a_1, a_2, \dots, a_k) \in Q''$. By construction, it is also guaranteed that every output tuple in the result of Q'' will have a non-zero entry in the product of M and N . Using Lemma 2.1, we can construct and multiply the matrices in time $O((|D|/\delta)^{\lceil k/2 \rceil+1+\lceil k/2 \rceil} \cdot (|D|/\delta)^{\omega-3}) = O((|D|/\delta)^{\omega+k-2})$ since the smallest dimension $\beta = (|D|/\delta)$ as $k \geq 2$. The output of Q' and Q'' can be deduplicated using a hash set.

We are ready to enumerate the query result. The key observation is that the materialized and deduplicated output of Q' and Q'' can be enumerated with constant delay, each of k subqueries where $x_i = \ell$ can be enumerated with delay $O(\delta)$, and there is no overlap in the output between any of the queries during the enumeration phase. It remains to argue the total preprocessing time required for processing Q' and Q'' . Processing the two queries requires $T_p = O(|D| \cdot \delta^{k-1} + (|D|/\delta)^{\omega+k-2})$. By ensuring that $|D| \cdot \delta^{k-1} \leq (|D|/\delta)^{\omega+k-2}$, we get the desired range of $1 \leq \delta \leq |D|^{(\omega+k-3)/(\omega+2k-3)}$. \square

For the two-path query and the current best value of $\omega = 2.3715$ [WXXZ24], we get the tradeoff as $T_p = O((|D|/\delta)^{2.3715})$ and a delay guarantee of $O(\delta)$ for $|D|^{0.15} < \delta \leq |D|^{0.406}$. If we choose $\delta = |D|^{0.40}$, the worst-case preprocessing time is $T_p = O(|D|^{1.4229})$. In contrast, Theorem 2.2 requires a worst-case preprocessing time of $T_p = O(|D|^{1.6})$, which is suboptimal compared to the above theorem. On the other hand, since $T_p = O(|D|^{1.4229})$, we can safely assume that $|\text{OUT}_\infty| > |D|^{1.4229}$, otherwise one can simply compute the full join in time $c^* \cdot |D|^{1.4229}$ using LISTMERGE, deduplicate and get constant delay enumeration. Applying Theorem 4.1 with $|\text{OUT}_\infty| > |D|^{1.4229}$ tells us that we can obtain delay as $O(|D|^2/|\text{OUT}_\infty|) = O(|D|^{0.5771})$. Thus, we can offer the user both choices, and the user can decide which algorithm to use.

5. LEFT-DEEP HIERARCHICAL QUERIES

In this section, we will apply our techniques to another subset of hierarchical queries, which we call *left-deep*. A left-deep hierarchical query is of the following form:

$$Q_{\text{leftdeep}}^k = \pi_{w_1, w_2, \dots, w_k} R_1(w_1, x_1) \bowtie R_2(w_2, x_1, x_2) \bowtie \dots \bowtie R_k(w_k, x_1, \dots, x_k)$$


Q_{leftdeep}^k is a hierarchical query for any $k \geq 1$. For $k = 2$, we get the two-path query. For $k = 3$, we get $\pi_{x_1, x_2, x_3} R_1(w_1, x_1) \bowtie R_2(w_2, x_1, x_2) \bowtie R_3(w_3, x_1, x_2, x_3)$ (hypergraph shown in Figure 7).

Algorithm 5 shows the preprocessing and enumeration phases for left-deep hierarchical queries. Let us first look at the preprocessing step. Similar to star queries, the algorithm begins by computing the size of the full join and performing the domain transformation to make sure all values in the relation are between 1 and $|D|$. Then, we create a bidirectional index for relation R_k . For each valuation v of w_k , we store the list $\pi_{x_1, \dots, x_k} \sigma_{w_k=v} R_k$. The

Algorithm 5: LEFTDEEPHIERARCHICAL(Q_{leftdeep}^k, D)

```

1 procedure PREPROCESS
2   Compute  $|\text{OUT}_{\infty}|$ 
3   Fix  $\delta = 2 \cdot |D|^k / |\text{OUT}_{\infty}|$ 
4   Apply the domain compression transformation using  $f$  and  $f^{-1}$  on each relation
      as described in Subsection 4.2
5   For relation  $R_k$ , create two additional indexes: (i) using a hash table  $H_k^{\ell}$  with
      key as all light valuations  $v$  for variable  $w_i$  and value as the list
       $\pi_{x_1, \dots, x_i} \sigma_{w_i=v} R_i$ ; (ii) using a hash table  $H_k^h$  with key as all heavy valuations  $v$ 
      for variable  $w_i$  and value as the list  $\pi_{x_1, \dots, x_i} \sigma_{w_i=v} R_i$ 
6   Create an index (called  $H_i(u)$ ) for each input relation  $R_i, i \in \{1, \dots, k-1\}$ 
      using a hash table with key as valuation  $u$  for variables  $x_1, \dots, x_i$  and value as
      the sorted list  $\pi_{w_i} \sigma_{x_1=u[x_1], \dots, x_i=u[x_i]} R_i$ 
7 procedure ENUMERATEFIXED $w_k(v)$ 
8   Consider each  $u \in \pi_{x_1, \dots, x_k} \sigma_{w_k=v} R_k$ . Let  $\mathcal{A}_u^v$  be the algorithm that enumerates
       $\times_{j=1}^{k-1} (H_j(\pi_{x_1, \dots, x_j} u))$  in lexicographic order with constant delay
9   Invoke Algorithm 2 with input as  $\mathcal{A}_u^v, u \in \pi_{x_1, \dots, x_k} \sigma_{w_k=v} R_k$ . Capture the output
       $w$  emitted on line 5 of Algorithm 2.
10  output  $f^{-1}(v, w)$ 
11 procedure ENUMERATE
12  Let  $\mathcal{A}'$  be the algorithm that calls ENUMERATEFIXED $w_k(v)$  for all light  $w_k$ 
      valuations, requiring total time  $T' = |\text{OUT}_{\infty}|/2$ 
13  Let  $\mathcal{A}$  be the algorithm that calls ENUMERATEFIXED $w_k(v)$  for all heavy  $w_k$ 
      valuations, requiring total time  $T = c^* \cdot |\text{OUT}_{\infty}|/2$  /*  $c^*$  is the upper bound on
      the number operations required by one invocation of ENUMERATEFIXED $w_k(v)$  */
14  Apply Lemma 3.1 with  $\mathcal{A}', \mathcal{A}, T, T'$  as input
  
```

Figure 7: Hypergraph for query Q_{leftdeep}^3 .

light valuations have the list stored in the hash table H_k^{ℓ} and the heavy valuations have the lists stored in H_k^h . For the remaining relations, we store a sorted list of w_i for each tuple $\pi_{x_1, \dots, x_i} R_i$ in the hash table H_i .

During the enumeration phase, like star queries, we will interleave the algorithm's execution that processes the light w_k valuations with the heavy w_k valuations. In particular, procedure $\text{ENUMERATEFIXED}_{w_k}(v)$ shows how to generate the output for a fixed w_k valuation v . In particular, for any fixing u of $\pi_{x_1, \dots, x_k} \sigma_{w_k=v} R_k$, the cartesian product of the $H_j(\pi_{x_1, \dots, x_k} u)$ can be enumerated with constant delay in lexicographic order. Therefore, now we can use the MERGE subroutine directly (Algorithm 2) to combine the output for all $u \in \pi_{x_1, \dots, x_k} \sigma_{w_k=v} R_k$. Since only the light w_k valuations will give us a small delay guarantee, the interleaving with the heavy valuations will help us ensure we can get a possibly sublinear overall guarantee. We show that the following result holds:

Theorem 5.1. *For a left-deep query Q_{leftdeep}^k and any instance D , there exists an algorithm with preprocessing time $T_p = O(|D|)$ and preprocessing space $S_p = O(|D|)$ such that the query results $Q_{\text{leftdeep}}^k(D)$ can be enumerated with delay $O(|D|^k/|\text{OUT}_{\bowtie}|)$ and enumeration space $S_e = O(|D|)$.*

Proof. Let the degree threshold be $\delta = 2 \cdot |D|^k/|\text{OUT}_{\bowtie}|$. First, observe that each procedure step PREPROCESS requires only $O(|D|)$ time. Indeed, computing $|\text{OUT}_{\bowtie}|$, domain compression, and index creation all require a linear pass on each relation R_i . Next, we argue that procedure $\text{ENUMERATEDFIXED}_{w_k}(v)$ can enumerate output with delay $O(|\pi_{x_1, \dots, x_k} \sigma_{w_k=v} R_k|)$. First, note that for a fixed $u \in \pi_{x_1, \dots, x_k} \sigma_{w_k=v} R_k$, \mathcal{A}_u^v enumerates the cartesian product with constant delay. Thus, since we invoke $|\pi_{x_1, \dots, x_k} \sigma_{w_k=v} R_k|$ instances of Algorithm 2, one for each value of u , each of which enumerates sorted results in constant delay, the overall delay is $O(|\pi_{x_1, \dots, x_k} \sigma_{w_k=v} R_k|)$ according to Lemma 3.3.

First, we analyze the total time T required by \mathcal{A} that processes all heavy valuations. As the degree threshold is δ , there are at most $|D|/\delta$ heavy w_k valuations, and each of them may generate $|D|^{k-1}$ output for the $k-1$ remaining projection variables. Therefore, $T \leq c^* \cdot |D|^{k-1} \cdot |D|/\delta = c^* \cdot |D|^k/\delta = c^* \cdot |\text{OUT}_{\bowtie}|/2$, where c^* is the upper bound on the number of operations performed by $\text{ENUMERATEDFIXED}_{w_k}(v)$ for generating one output tuple. Further, if Q_H and Q_L denote the subqueries without projections evaluated by \mathcal{A} and \mathcal{A}' , it holds that $|Q_H| + |Q_L| = |\text{OUT}_{\bowtie}|$ since both algorithms are doing list merging and go over every tuple in the full join once. Therefore, it holds that $|Q_L| = |\text{OUT}_{\bowtie}| - |Q_H| = |\text{OUT}_{\bowtie}| - |D|^k/\delta \geq |\text{OUT}_{\bowtie}|/2$. Therefore, a lower bound on T' (i.e., time required to process all light valuations) is $|\text{OUT}_{\bowtie}|/2$.

We can now apply Lemma 3.1 with the following parameters

- \mathcal{A}' with $T' = |\text{OUT}_{\bowtie}|/2$ and delay guarantee of \mathcal{A}' as $O(\delta)$
 - \mathcal{A} with $T = c^* \cdot |\text{OUT}_{\bowtie}|/2$
 - T, T' are fixed once $|\text{OUT}_{\bowtie}|$ and constants c^* has been determined by analyzing the program.
- Lemma 3.1 implies the delay $O(\delta \cdot \max\{1, T/T'\}) = O(\delta)$, achieving the desired guarantee. \square

In the above theorem, OUT_{\bowtie} is the full join size of Q_{leftdeep}^k . The AGM exponent for Q_{leftdeep}^k is k . Observe that Theorem 5.1 is of interest when $|\text{OUT}_{\bowtie}| > |D|^{k-1}$ to ensure that the delay is smaller than $O(|D|)$. When the condition $|\text{OUT}_{\bowtie}| > |D|^{k-1}$ holds, the delay obtained by Theorem 5.1 is also better than the one given by the tradeoff in Theorem 2.2. In the worst-case when $|\text{OUT}_{\bowtie}| = \Theta(|D|^k)$, we can achieve constant delay enumeration after linear preprocessing time, compared to Theorem 2.2 that would require $\Theta(|D|^k)$ preprocessing time to achieve the same delay. The decision of when to apply Theorem 5.1 or Theorem 2.2 can be made in linear time by checking whether $|D|^k/|\text{OUT}_{\bowtie}|$ is smaller or larger than the actual delay guarantee obtained by the algorithm of Theorem 2.2 after linear time preprocessing.

6. PATH QUERIES

In this section, we will study path queries P_k . Recall that for $k \geq 3$, P_k is not a hierarchical query; hence, the tradeoff from [KNOZ20] does not apply. A subset of path queries, namely 3-path and 4-path counting queries, were considered in [KNN⁺19]. The algorithm used for counting the answers of 3-path and 4-path queries under updates constructed a set of views that can be used to enumerate the query results under the static setting. Our result extends the same idea to apply to arbitrary length path queries, which we state next.

Theorem 6.1. *For a path query P_k , any instance D and parameter $\epsilon \in [0, 1)$, there exists an algorithm with preprocessing time $T_p = O(|D|^{2-\epsilon/(k-1)})$ and preprocessing space $S_p = O(|D|^{2-\epsilon/(k-1)})$, such that the query result $P_k(D)$ can be enumerated with delay $O(|D|^\epsilon)$ and enumeration space $S_e = O(|D|)$.*

Proof. Let Δ be a parameter that we will fix later. In the preprocessing phase, we first perform a full reducer pass to remove dangling tuples, apply the domain transformation technique by creating f and f^{-1} and then create a hash map for each relation $R_i(x_i, x_{i+1})$ with key x_i , and all its corresponding x_{i+1} values sorted for each key entry. (We also store the degree of each value.) Next, for every $i = 1, \dots, k$, and every heavy value a of x_i in R_i (with degree $> \Delta$), we compute the query $\pi_{x_{k+1}}(R_i(a, x_{i+1}) \bowtie \dots \bowtie R_k(x_k, x_{k+1}))$, and store its result sorted in a hash map with key a . Note that each such query can be computed in time $O(|D|)$ through a sequence of semijoins and projections, and sorting in linear time using count sort. Since there are at most $|D|/\Delta$ heavy values for each x_i , the total running time (and space necessary) of this step is $O(|D|^2/\Delta)$.

We will present the enumeration algorithm using induction. In particular, for each $i = k, \dots, 1$ and for every value a of x_i , the subquery $\pi_{x_{k+1}}(R_i(a, x_{i+1}) \bowtie \dots \bowtie R_k(x_k, x_{k+1}))$ can be enumerated (using the same order) with delay $O(\Delta^{k-i})$. This implies that our target path query can be enumerated with delay $O(\Delta^{k-1})$ by simply iterating through all values of x_1 in R_1 . Finally, we can obtain the desired result by choosing $\Delta = |D|^{\epsilon/(k-1)}$. Indeed, for the base case ($i = k$) it is trivial to see that we can enumerate $\pi_{x_{k+1}}(R_k(a, x_{k+1}))$ with constant delay $O(1)$ using the stored hash map. Consider some i and a value a for x_i in R_i for the inductive step. If the value a is heavy, we can enumerate all the x_{k+1} 's with constant delay by probing the hash map we computed during the preprocessing phase. If the value is light, then there are at most Δ values of x_{i+1} . For each such value b , the inductive step provides an algorithm that enumerates all x_{k+1} with delay $O(\Delta^{k-i-1})$. Observe that the order across all b 's will be the same. Thus, we can apply Lemma 3.3 to obtain that we can enumerate the union of the results with delay $O(\Delta \cdot \Delta^{k-i-1}) = O(\Delta^{k-i})$. Finally, For each output tuple t generated, we return $f^{-1}(t)$ to the user. \square

When $\epsilon = 1$, we can obtain a delay $O(|D|)$ using only linear preprocessing time $O(|D|)$ using the result of [BDG07] since the query is acyclic, while when ϵ approaches to 1, the above theorem would give preprocessing time $O(|D|^{2-1/(k-1)})$. Hence, for $k \geq 3$, we observe a discontinuity in the time-delay tradeoff. A second observation following from Theorem 6.1 is that as $k \rightarrow \infty$, the tradeoff collapses to only two extremal points: one where we get constant delay with $T_p = O(|D|^2)$, and the other where we get linear delay with $T_p = O(|D|)$.

7. RELATED WORK

We overview prior work on static query evaluation for acyclic join-project queries. The result of any acyclic CQ can be enumerated with constant delay after linear time preprocessing if and only if it is free-connex [BDG07]. This is based on the conjecture that Boolean multiplication of $n \times n$ matrices cannot be done in $O(n^2)$ time. Acyclicity itself is necessary for having constant delay enumeration: A CQ admits constant delay enumeration after linear-time preprocessing if and only if it is free-connex acyclic [BB13]. This is based on a stronger hypothesis that the existence of a triangle in a hypergraph of n vertices cannot be tested in time $O(n^2)$ and that for any k , testing the presence of a k -dimensional tetrahedron cannot be tested in linear time. We refer readers to an overview of problems and progress related to constant delay enumeration in [Seg15b]. Prior work also exhibits a dependency between the space and enumeration delay for conjunctive queries with access patterns [DK18]. It constructs a succinct representation of the query result that allows for the enumeration of tuples over some variables under value bindings for all other variables. As noted by [KNOZ20], it does not support enumeration for queries with free variables, which is also its main contribution. Our work demonstrates that for a subset of hierarchical queries, the tradeoff shown in [KNOZ20] is not optimal. Our work introduces fundamentally new ideas that may be useful in improving the tradeoff for arbitrary hierarchical queries and enumeration of UCQs. There has also been some experimental work by the database community on problems related to enumerating join-project query results efficiently but without any formal delay guarantees. Seminal work [XKD15, XD17, XSD17, DK21] has studied how compressed representations can be created apriori that allow for faster enumeration of query results. For the two path query, the fastest evaluation algorithm (with no delay guarantees) evaluates the projection join output in time $O(|D| \cdot |\text{OUT}_\pi|^{\frac{(\omega-1)}{(\omega+1)}} + |D|^{\frac{2(\omega-1)}{(\omega+1)}} \cdot |\text{OUT}_\pi|^{\frac{2}{(\omega+1)}})$ [DHK20, AP09]. For star queries, there is no closed form expression but fast matrix multiplication can be used to obtain instance dependent bounds on running time. Also related is the problem of dynamic evaluation of hierarchical queries. Recent work [KNN⁺19, KNOZ20, BKS17, BKS18] has studied the tradeoff between amortized update time and delay guarantees. Some of our techniques may also lead to new insights and improvements in existing algorithms. Prior work in differential privacy [RCWH⁺20] and directed graphical models [CRJ20] may also benefit from some of our techniques.

8. CONCLUSION AND OPEN PROBLEMS

In this paper, we studied the problem of enumerating query results for an important subset of CQs with projections, namely star and path queries. We presented data-dependent algorithms that improve upon existing results by achieving non-trivial delay guarantees in linear preprocessing time. Our results are based on interleaving join query computation to achieve meaningful delay guarantees. Further, we showed how non-combinatorial algorithms (fast matrix multiplication) can be used for faster preprocessing to improve the tradeoff between preprocessing time and delay. We also presented new results on time-delay tradeoffs for a subset of non-hierarchical queries for the class of path queries. Our results also open several new tantalizing questions that open up possible directions for future work.

- *More preprocessing time for star queries.* One major open question is whether Theorem 4.1 can benefit from more preprocessing time to achieve lower delay guarantees. For instance,

if we can afford the algorithm preprocessing time $T_p = O(|\text{OUT}_\infty|/|D|^\epsilon + |D|)$ time, can we expect to get delay $\delta = O(|D|^\epsilon)$ for all $\epsilon \in (0, 1)$?

- *Sublinear delay guarantees for the two-path query.* It is not known whether we can achieve a sublinear delay guarantee after linear preprocessing time for $Q_{\text{two-path}}$. This question is equivalent to the following problem: for what values of $|\text{OUT}_\pi|$ can Q_{path} be evaluated in linear time? If $|\text{OUT}_\pi| = |D|^\epsilon$, the best known algorithms can evaluate $Q_{\text{two-path}}$ in time $O(|D|^{1+\epsilon/3})$ (using fast matrix multiplication) [DHK20] but this is still superlinear.
- *Space-delay bounds.* The last question is to study the tradeoff between space vs delay for arbitrary hierarchical and path queries. Using some of our techniques, it may be possible to smartly materialize a certain subset of joins that could be used to achieve delay guarantees by interleaving with join computation. We also believe the space-delay tradeoff implied by prior work can be improved for specific delay ranges using the ideas introduced in this paper.

REFERENCES

- [AKNR16] 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.
- [AKNS17] Mahmoud Abo Khamis, Hung Q Ngo, and Dan Suciu. What do shannon-type inequalities, submodular width, and disjunctive datalog have to do with one another? In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 429–444, 2017.
- [AP09] Rasmus Resen Amossen and Rasmus Pagh. Faster join-projects and sparse matrix multiplications. In *Proceedings of the 12th International Conference on Database Theory*, pages 121–126. ACM, 2009.
- [BB13] Johann Brault-Baron. *De la pertinence de l'énumération: complexité en logiques propositionnelle et du premier ordre*. PhD thesis, Université de Caen, 2013.
- [BDG07] 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.
- [BGS20] Christoph Berkholz, Fabian Gerhardt, and Nicole Schweikardt. Constant delay enumeration for conjunctive queries: a tutorial. *ACM SIGLOG News*, 7(1):4–33, 2020.
- [BKS17] Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering conjunctive queries under updates. In *proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI symposium on Principles of database systems*, pages 303–318. ACM, 2017.
- [BKS18] Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering fo+ mod queries under updates on bounded degree databases. *ACM Transactions on Database Systems (TODS)*, 43(2):7, 2018.
- [BMT20] Angela Bonifati, Wim Martens, and Thomas Timm. An analytical study of large sparql query logs. *The VLDB Journal*, 29(2):655–679, 2020.
- [CLRS09] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Chapter 8.2, Introduction to algorithms*. MIT press, 2009.
- [CLRS22] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2022.
- [CRJ20] Amrita Roy Chowdhury, Theodoros Rekatsinas, and Somesh Jha. Data-dependent differentially private parameter learning for directed graphical models. In *International Conference on Machine Learning*, pages 1939–1951. PMLR, 2020.
- [CS23] Nofar Carmeli and Luc Segoufin. Conjunctive queries with self-joins, towards a fine-grained enumeration complexity analysis. In *Proceedings of the 42nd ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 277–289, 2023.
- [DHK20] Shaleen Deep, Xiao Hu, and Paraschos Koutris. Fast join project query evaluation using matrix multiplication. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 1213–1223, 2020.
- [DHK21] Shaleen Deep, Xiao Hu, and Paraschos Koutris. Enumeration algorithms for conjunctive queries with projection. In *24th International Conference on Database Theory (ICDT 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- [DK18] Shaleen Deep and Paraschos Koutris. Compressed representations of conjunctive query results. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 307–322. ACM, 2018.
- [DK21] Shaleen Deep and Paraschos Koutris. Ranked enumeration of conjunctive query results. *To appear in Joint 2021 EDBT/ICDT Conferences, ICDT '21 Proceedings*, 2021.
- [FT87] Michael L Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the ACM (JACM)*, 34(3):596–615, 1987.
- [GGS14] Georg Gottlob, Gianluigi Greco, and Francesco Scarcello. Treewidth and hypertree width. *Tractability: Practical Approaches to Hard Problems*, 1, 2014.
- [GM14] Martin Grohe and Dániel Marx. Constraint solving via fractional edge covers. *ACM Transactions on Algorithms (TALG)*, 11(1):1–20, 2014.
- [GS13] Gianluigi Greco and Francesco Scarcello. Structural tractability of enumerating csp solutions. *Constraints*, 18(1):38–74, 2013.

- [HUA75] John E Hopcroft, Jeffrey D Ullman, and AV Aho. The design and analysis of computer algorithms, 1975.
- [Kaz13] Wojciech Kazana. *Query evaluation with constant delay*. PhD thesis, 2013.
- [KNN⁺19] Ahmet Kara, Hung Q Ngo, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. Counting triangles under updates in worst-case optimal time. In *22nd International Conference on Database Theory*, 2019.
- [KNOZ20] Ahmet Kara, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. Trade-offs in static and dynamic evaluation of hierarchical queries. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 375–392, 2020.
- [Mar13] Dániel Marx. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. *Journal of the ACM (JACM)*, 60(6):1–51, 2013.
- [NPRR12] Hung Q Ngo, Ely Porat, Christopher Ré, and Atri Rudra. Worst-case optimal join algorithms. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems*, pages 37–48. ACM, 2012.
- [NRR13] Hung Q. Ngo, Christopher Ré, and Atri Rudra. Skew strikes back: new developments in the theory of join algorithms. *SIGMOD Record*, 42(4):5–16, 2013. doi:10.1145/2590989.2590991.
- [OS16] Dan Olteanu and Maximilian Schleich. Factorized databases. *ACM SIGMOD Record*, 45(2):5–16, 2016.
- [OVL81] Mark H Overmars and Jan Van Leeuwen. Dynamization of decomposable searching problems yielding good worst-case bounds. In *Theoretical Computer Science*, pages 224–233. Springer, 1981.
- [OZ15] Dan Olteanu and Jakub Závodný. Size bounds for factorised representations of query results. *ACM Transactions on Database Systems (TODS)*, 40(1):1–44, 2015.
- [PP06] Anna Pagh and Rasmus Pagh. Scalable computation of acyclic joins. In *Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 225–232, 2006.
- [RCWH⁺20] Amrita Roy Chowdhury, Chenghong Wang, Xi He, Ashwin Machanavajjhala, and Somesh Jha. Crypt?: Crypto-assisted differential privacy on untrusted servers. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 603–619, 2020.
- [Seg15a] Luc Segoufin. Constant delay enumeration for conjunctive queries. *SIGMOD Record*, 44(1):10–17, 2015. doi:10.1145/2783888.2783894.
- [Seg15b] Luc Segoufin. Constant delay enumeration for conjunctive queries. *ACM SIGMOD Record*, 44(1):10–17, 2015.
- [SORK11] Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch. Probabilistic databases, synthesis lectures on data management. *Morgan & Claypool*, 2011.
- [WW18] Virginia Vassilevska Williams and R Ryan Williams. Subcubic equivalences between path, matrix, and triangle problems. *Journal of the ACM (JACM)*, 65(5):1–38, 2018.
- [WXXZ24] Virginia Vassilevska Williams, Yinzhan Xu, Zixuan Xu, and Renfei Zhou. New bounds for matrix multiplication: from alpha to omega. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3792–3835. SIAM, 2024.
- [XD17] Konstantinos Xirogiannopoulos and Amol Deshpande. Extracting and analyzing hidden graphs from relational databases. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 897–912. ACM, 2017.
- [XKD15] Konstantinos Xirogiannopoulos, Udayan Khurana, and Amol Deshpande. Graphgen: Exploring interesting graphs in relational data. *Proceedings of the VLDB Endowment*, 8(12):2032–2035, 2015.
- [XSD17] Konstantinos Xirogiannopoulos, Virinchi Srinivas, and Amol Deshpande. Graphgen: Adaptive graph processing using relational databases. In *Proceedings of the Fifth International Workshop on Graph Data-management Experiences & Systems, GRADES’17*, pages 9:1–9:7, New York, NY, USA, 2017. ACM. doi:10.1145/3078447.3078456.
- [Yan81] Mihalis Yannakakis. Algorithms for acyclic database schemes. In *VLDB*, volume 81, pages 82–94, 1981.