

Fine-Grained Complexity of Regular Path Queries*

Katrin Casel¹ and Markus L. Schmid²

¹Hasso Plattner Institute, University of Potsdam, Potsdam, Germany,
Katrin.Casel@hpi.de

²Humboldt-Universität zu Berlin, Unter den Linden 6, D-10099, Berlin, Germany,
MLSchmid@MLSchmid.de

January 22, 2021

Abstract

A regular path query (RPQ) is a regular expression q that returns all node pairs (u, v) from a graph database that are connected by an arbitrary path labelled with a word from $L(q)$. The obvious algorithmic approach to RPQ-evaluation (called PG-approach), i. e., constructing the product graph between an NFA for q and the graph database, is appealing due to its simplicity and also leads to efficient algorithms. However, it is unclear whether the PG-approach is optimal. We address this question by thoroughly investigating which upper complexity bounds can be achieved by the PG-approach, and we complement these with conditional lower bounds (in the sense of the fine-grained complexity framework). A special focus is put on enumeration and delay bounds, as well as the data complexity perspective. A main insight is that we can achieve optimal (or near optimal) algorithms with the PG-approach, but the delay for enumeration is rather high (linear in the database). We explore three successful approaches towards enumeration with sub-linear delay: super-linear preprocessing, approximations of the solution sets, and restricted classes of RPQs.

1 Introduction

An essential component of graph query languages (to be found both in academical prototypes as well as in industrial solutions) are *regular path queries* (RPQs). Abstractly speaking, a regular expression q over some alphabet Σ is interpreted as query that returns from a Σ -edge-labelled, directed graph \mathcal{D} (i. e., a *graph database*) the set $q(\mathcal{D})$ of all node pairs (u, v) that are connected by a q -path, i. e., a path labelled with a word from q 's language (and possibly also a witness path per node pair, or even all such paths). This simple, yet relevant concept has heavily been studied in database theory (the following list is somewhat biased towards recent work): results on RPQs [25, 7, 34, 36, 39, 16], conjunctive RPQs [15, 43, 9] and extensions thereof [33, 12, 33, 28], questions of static analysis [27, 11, 26, 30, 44], experimental analyses [17, 19, 38], and surveys of this research area [10, 54, 5, 23].

In the simplest setting, where we are only interested in the node pairs (but no paths) connected by *arbitrary* q -paths (instead of, e. g., simple paths), evaluation can be done efficiently. Deviating from this simple setting, however, leads to intractability: if we ask for nodes connected by *simple paths* (no repeated nodes), or connected by *trails* (no repeated arcs), then RPQ evaluation is NP-hard even in data-complexity (see [40, 7, 37] and [36], respectively). Note that the simple path and trail semantics are mostly motivated by the fact that under these semantics there is only a finite number of q -paths per node pair. If we move to conjunctions of RPQs (CRPQs) or even

*This is the preliminary full version of an article accepted for publication in the proceedings of the 24th International Conference on Database Theory (ICDT 2021). The first author has been funded by the Federal Ministry of Education and Research of Germany (BMBF) in the KI-LAB-ITSE framework – project number 01IS19066. The second author has been funded by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) – project number 416776735 (gefördert durch die Deutsche Forschungsgemeinschaft (DFG) – Projektnummer 416776735).

more powerful extensions motivated by practical requirements, then also with the arbitrary path semantics evaluation becomes intractable in combined complexity (i. e., they inherit hardness from relational conjunctive queries (CQs)).

In order to guide practical developments in the area of graph databases, the computational hard cases of RPQ (and CRPQ) evaluation have been thoroughly investigated in database theory. However, with respect to arbitrary q -paths, research seems to have stopped at the conclusion that efficient evaluation is possible by the following simple *PG-approach*: given graph database $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$ and NFA M_q for q with state set V_q , construct the *product graph* $G_{\mathcal{D},q}$ with nodes $V_{\mathcal{D}} \times V_q$ and an arc $((u, p), (v, p'))$ iff, for some $\mathbf{a} \in \Sigma$, \mathcal{D} has an arc (u, \mathbf{a}, v) and M_q has an arc (p, \mathbf{a}, p') , and then use simple graph-searching techniques on $G_{\mathcal{D},q}$.

The PG-approach is explicitly defined in several papers, e. g., [40, 10, 37], and mainly used to prove a worst-case upper bound (actually $O(|q||\mathcal{D}|)$ for Boolean evaluation); in [37] it is used for enumerating q -paths between two given nodes. But it is also very appealing from a practical point of view due to its simplicity: we are just coupling well-understood algorithmic concepts like finite automata and graph reachability algorithms. Arguably, implementing the PG-approach is an exercise suitable for a first year programming course (making it feasible and cost-efficient for industrial systems). As it seems, putting RPQ evaluation with arbitrary path semantics and the respective PG-approach into the focus of a thorough theoretical study has not yet been done. This paper is devoted to this task. In particular, we wish to investigate the following two (somewhat overlapping) aspects:

1. Applicability of the PG-approach: the PG-approach is suited for solving simple evaluation problems like checking $q(\mathcal{D}) = \emptyset$ or $(u, v) \in q(\mathcal{D})$ (for given $u, v \in V_{\mathcal{D}}$), but is it also appropriate for more relevant tasks like computing, counting or enumerating $q(\mathcal{D})$?
2. Optimality of the PG-approach: Does the PG-approach lead to optimal algorithms, or can it be beaten by conceptionally different techniques?

Answering these questions provides a better theoretical understanding of RPQ-evaluation (which, as mentioned above, are at the heart of many graph query languages). But also for the more powerful CRPQs and more practically motivated graph query languages, we can derive valuable insights from our investigation. Let us mention two such examples (a complete summary of our results follows further down). As noted in [10], we can reduce CRPQ evaluation to the evaluation of relational CQs by first constructing all tables represented by the single RPQs and then evaluating a CQ over this database. To do this, we first have to compute the results of all RPQs, so it seems helpful to know the best algorithms for this intermediate task. Moreover, if we want to benefit from the existing CQ evaluation techniques (e. g., exploiting acyclicity etc.) we are more or less forced to this two-step approach. With respect to *enumerations* of CQs, it is known that linear preprocessing and constant delay enumeration is possible provided that the CQs satisfy certain acyclicity properties (see [8, 14], or the surveys [13, 45]). Unfortunately, these techniques do not carry over to CRPQs since, as we show, linear preprocessing and constant delay enumeration is not possible even for single RPQs (conditional to some complexity assumptions).

Since the problem we investigate can be solved in polynomial time (also in combined complexity), we cannot show lower bounds in terms of hardness results for complexity classes like NP or PSPACE. Instead, we make use of the framework of *fine-grained complexity*, which allows to prove lower bounds that are conditional on some algorithmic assumptions (see the surveys [51, 22, 52]). In particular, fine-grained complexity is a rather successful toolbox for giving evidence that the *obvious* algorithmic approach to some basic problem, is also the *optimal* one. This is exactly our setting here, with respect to RPQ-evaluation and the PG-approach. To the knowledge of the authors, such conditional lower bounds are not yet a well-established technique in database theory (however, see [13, Section 6] for a survey of conditional lower bounds in the context of CQ enumeration).

A main challenge is that fine-grained complexity is not exactly tailored to either the data-complexity perspective or to enumeration problems. We will next outline our results.

Non-enum. Results		BOOLE, TEST, WITNESS	EVAL	COUNT
upper bounds		$O(\mathcal{D} q)$	$O(V_{\mathcal{D}} \mathcal{D} q)$ $O((V_{\mathcal{D}} q)^{\omega})$	$O(V_{\mathcal{D}} \mathcal{D} q)$ $O((V_{\mathcal{D}} q)^{\omega})$
lower bounds	OV & com-BMM	$O((\mathcal{D} q)^{1-\epsilon})$	—	—
	OV	—	—	$O^{\text{dc}}((V_{\mathcal{D}} \mathcal{D})^{1-\epsilon})$
	SBMM	—	$O^{\text{dc}}(q(\mathcal{D}) + \mathcal{D})$	—
	com-BMM	—	$O^{\text{dc}}((V_{\mathcal{D}} \mathcal{D})^{1-\epsilon})$	—

Table 1: All upper bounds can be achieved as running times of some algorithm, while the lower bounds cannot be achieved as running time by any algorithm, unless the displayed hypothesis fails. The exponent ω denotes the best known matrix multiplication exponent.

1.1 Our Contribution

All investigated RPQ-evaluation problems are summarised on page 7. In the following, $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$ is the graph database, q is the RPQ, and $\epsilon > 0$. With the notation $O^{\text{dc}}(\cdot)$, we hide factors $f(|q|)$ for some function f (i. e., it is used for stating data-complexities). All lower bounds mentioned in the following are conditional to some of the algorithmic assumptions summarised in Section 4 (we encourage the reader less familiar with fine-grained complexity hypotheses to have a look at this section first, which can be read independently). For presentational reasons, we do not always explicitly mention this in the rest of the introduction and when we say that a certain running time is “not possible”, this statements is always conditional in this sense (see Tables 1 and 2 for the actual hypotheses). As common in fine-grained complexity, we rule out *true* sub-linear ($O(n^{1-\epsilon})$), sub-quadratic ($O(n^{2-\epsilon})$), or sub-cubic ($O(n^{3-\epsilon})$) running times, but not possible improvements by logarithmic factors.

Non-Enumeration Variants

The following results are summarised in Table 1. For the simple problems BOOLE (checking $q(\mathcal{D}) = \emptyset$), TEST (checking $(u, v) \in q(\mathcal{D})$) and WITNESS (computing some element from $q(\mathcal{D})$), the PG-approach yields an upper bound of $O(|\mathcal{D}||q|)$, which is optimal (since linear) in data complexity, and we can show lower bounds demonstrating its optimality also in combined complexity. For EVAL (computing the set $q(\mathcal{D})$) the PG-approach yields a data complexity upper bound of $O^{\text{dc}}(|V_{\mathcal{D}}||\mathcal{D}|)$, which cannot be improved by *combinatorial* algorithms, although $O^{\text{dc}}(|V_{\mathcal{D}}|^{2.37})$ is possible by fast matrix multiplication (see Section 4 for a discussion of the meaning of the term “combinatorial”). In addition, we can show that linear time data complexity, i. e., $O^{\text{dc}}(|q(\mathcal{D})| + |\mathcal{D}|)$, is not possible even for non-combinatorial algorithms. For COUNT (computing $|q(\mathcal{D})|$), we get $O^{\text{dc}}(|V_{\mathcal{D}}||\mathcal{D}|)$ as upper and lower bound, not restricted to combinatorial algorithms.

Enumeration

Our results for RPQ-enumeration are summarised in Table 2. An entry “O(delay)” in column “preprocessing” means that the preprocessing is bounded by the delay (which means that no preprocessing is required). The column “sorted” indicates whether the enumeration is produced lexicographically sorted.

In comparison to the non-enumeration problem variants, the picture is less clear and deserves more explanation. The PG-approach yields a simple enumeration algorithm with delay $O(|\mathcal{D}||q|)$, that also trivially supports updates in constant time, since the preprocessing fits into the delay bound. Our lower bounds for BOOLE also mean that this delay cannot be improved in terms of *combined complexity*. While this lower bound was interesting for problems like BOOLE etc., it now gives a correct answer to the wrong question. The main goal now should be to find out whether we can remedy the linear dependency of the delay on $|\mathcal{D}|$, at the expense of spending more time in terms of $|q|$, or of losing the ability of handling updates, or even of allowing a slightly super-linear preprocessing.

In this regard, the strongest result would be linear preprocessing $O(|\mathcal{D}|f(|q|))$ and constant delay $O(f(|q|))$. However, we can rule this out even for algorithms not capable of handling updates.

Enum. Results		ENUM			
		preprocessing	delay	sorted	updates
upper bound		$O(\text{delay})$	$O(\mathcal{D} q)$	\checkmark	$O(1)$
lower bounds	OV & com-BMM	$O(\text{delay})$	$O((\mathcal{D} q)^{1-\epsilon})$	\times	\times
	SBMM	$O^{\text{dc}}(\mathcal{D})$	$O^{\text{dc}}(1)$	\times	\times
	com-BMM	$O^{\text{dc}}(\mathcal{D})$	$O^{\text{dc}}(V_{\mathcal{D}} ^{1-\epsilon})$	\times	\times
	OMv	arbitrary	$O^{\text{dc}}(V_{\mathcal{D}} ^{1-\epsilon})$	\times	$O^{\text{dc}}(V_{\mathcal{D}} ^{1-\epsilon})$
	com-BMM	$O^{\text{dc}}(\mathcal{D})$	$O^{\text{dc}}(V_{\mathcal{D}} ^{2-\epsilon})$	\times	$O^{\text{dc}}(V_{\mathcal{D}} ^{2-\epsilon})$

Table 2: All upper bounds can be achieved as running times of some algorithm, while the lower bounds cannot be achieved as running time by any algorithm, unless the displayed hypothesis fails. The exponent ω denotes the best known matrix multiplication exponent.

Then, the next question is which non-constant delays can be achieved that are strictly better than linear. For example, none of our lower bounds for the non-enumeration variants suggest that linear preprocessing and a delay bounded by, e. g., $|V_{\mathcal{D}}|$ or the degree of \mathcal{D} , should not be possible. We are *not* able to answer this question in its general form (and believe it to be very challenging), but we are able to provide several noteworthy insights.

For linear preprocessing, a delay of $O(|V_{\mathcal{D}}|)$ (if possible at all) cannot be beaten by combinatorial algorithms (even without updates). This can be strengthened considerably, if we also require updates in some reasonable time: for general algorithms (i. e., not necessarily combinatorial) delay *and* update time strictly better than $O^{\text{dc}}(|V_{\mathcal{D}}|)$ is not possible even with arbitrary preprocessing, and for combinatorial algorithms with linear preprocessing even delay *and* update time of $O(|\mathcal{D}|)$ cannot be beaten. This last result nicely complements the upper bound at least for combinatorial algorithms and in the dynamic case.

In summary, for linear preprocessing, $O(|V_{\mathcal{D}}|)$ is a lower bound for the delay and if we can beat $O(|\mathcal{D}|)$, we should not be able to also support updates.

Enumeration of Restricted Variants

Finally, we obtain restricted problem variants that can be solved with delay strictly better than $O(|\mathcal{D}|)$ (in data complexity). We explore three different approaches:

1. by allowing super-linear preprocessing of $O^{\text{dc}}(\overline{\Delta}(\mathcal{D}) \log(\overline{\Delta}(\mathcal{D}))|\mathcal{D}|)$ (where $\overline{\Delta}(\mathcal{D})$ is the average degree of \mathcal{D}), we can achieve a delay of $O(|V_{\mathcal{D}}|)$;
2. in linear preprocessing and constant delay, we can enumerate a representative subset of $q(\mathcal{D})$ instead of the whole set $q(\mathcal{D})$;
3. for a subclass of RPQs, we can solve RPQ-ENUM with linear preprocessing and delay $O(\Delta(\mathcal{D}))$ (where $\Delta(\mathcal{D})$ is the maximum degree of \mathcal{D}).

2 Main Definitions

Let $\mathbb{N} = \{1, 2, 3, \dots\}$ and $[n] = \{1, 2, \dots, n\}$ for $n \in \mathbb{N}$. For a finite alphabet A , A^+ denotes the set of non-empty words over A and $A^* = A^+ \cup \{\varepsilon\}$ (where ε is the empty word). For a word $w \in A^*$, $|w|$ denotes its length; $w^1 = w$ and $w^k = ww^{k-1}$ for every $k \geq 2$. For $L, K \subseteq A^*$, let $L \cdot K = \{w_1 \cdot w_2 \mid w_1 \in L, w_2 \in K\}$, let $L^1 = L$ and $L^k = L \cdot L^{k-1}$ for every $k \geq 2$, let $L^+ = \bigcup_{k \geq 1} L^k$ and $L^* = L^+ \cup \{\varepsilon\}$.

2.1 Σ -Graphs

We now define the central graph model that is used to represent graph databases as well as finite automata. Let Σ be a finite alphabet of constant size. A Σ -graph is a directed, edge labelled multigraph $G = (V, E)$, where V is the set of *vertices* (or *nodes*) and $E \subseteq V \times (\Sigma \cup \{\varepsilon\}) \times V$ is the set of *edges* (or *arcs*). For $u \in V$ and $x \in \Sigma \cup \{\varepsilon\}$, $E_x(u) = \{v \mid (u, x, v) \in E\}$ is

the set of x -successors of u . A path from $w_0 \in V$ to $w_k \in V$ of length $k \geq 0$ is a sequence $p = (w_0, a_1, w_1, a_2, w_2, \dots, w_{k-1}, a_k, w_k)$ with $(w_{i-1}, a_i, w_i) \in E$ for every $i \in [k]$. We say that p is *labelled* with the *word* $a_1 a_2 \dots a_k \in \Sigma^*$. According to this definition, for every $v \in V$, (v) is a path from v to v of length 0 that is labelled by ε . Hence, every node v of a Σ -graph has an ε -labelled path to itself, even though there might not be an ε -arc from v to v . Moreover, due to ε as a possible edge-label, paths of length k may be labelled with words w with $|w| < k$. The size of $G = (V, E)$ is $|G| = \max\{|V|, |E|\}$.

For any Σ -graph $G = (V, E)$, we call $(V, \{(u, v) \mid u \neq v \wedge \exists x \in \Sigma \cup \{\varepsilon\} : (u, x, v) \in E\})$ the *underlying graph* of G (note that the underlying graph is simple, non-labelled and has no loops). In particular, by a slight abuse of notation, we denote by E^* the reflexive-transitive closure of the underlying graph of G . Since we always assume $|\Sigma|$ to be a constant, we have that $|G| = \Theta(\max\{|V|, |\{(u, v) \mid u \neq v \wedge \exists x \in \Sigma \cup \{\varepsilon\} : (u, x, v) \in E\}|\})$ (i. e., $|G|$ is asymptotically equal to the size of its underlying graph). For every $u \in V$, the *degree of u* is $\Delta(u) = |\bigcup_{x \in \Sigma \cup \{\varepsilon\}} E_x(u)|$ (so $\Delta(u)$ is actually the out-degree), and the *maximum degree* of G is $\Delta(G) = \max\{|\Delta(u)| \mid u \in V\}$. The *average degree* of G is $\bar{\Delta}(G) = \frac{1}{|V|} \sum_{u \in V} |\Delta(u)|$. Obviously, $\bar{\Delta}(G) \leq \Delta(G) \leq |V|$.

Since Σ -graphs are the central data structures for our algorithms, we have to discuss implementational aspects of Σ -graphs in more detail. The set V of a Σ -graph $G = (V, E)$ is represented as a list, and, for every $u \in V$ and for every $x \in \Sigma \cup \{\varepsilon\}$, we store a list of all x -successors of u , which is called the *x -adjacency list* for u . We assume that we can check in constant time whether a list is empty and we can insert elements in constant time. However, finding and deleting an element from a list requires linear time. Furthermore, we assume that we always store together with a node a pointer to its adjacency list (thus, we can always retrieve the x -adjacency list for a given node in constant time).

Remark 2.1. All lower bounds presented in this paper hold for any graph representation that can be constructed in time linear in $|G| = \max\{|V|, |E|\}$. For the upper bounds, we chose the simple representation with adjacency lists as it emerged as the natural structure for our enumeration approach; let us point out here that since we always store pointers to the adjacency lists along with the nodes, we can perform a breadth-first search (BFS) from any given start node u in time $O(|G|)$. It is a plausible assumption that most specific graph representations can be transformed into our list-based representation without much effort. This ensures a certain generality of our upper complexity bounds in the sense that the corresponding algorithms are, to a large extent, independent from implementational details. Note also that the list-based structure only requires space linear in $|G|$.

In the adjacency list representation, we do not have random access to specific nodes in the graph database, or to specific neighbours of a given node. Thus, we have to measure a non-constant running-time for performing such operations. However, the algorithms for our upper bounds are independent from this aspect, i. e., the total running times would not change if we assume random access to nodes in constant time.

An exception to this is Theorem 7.1, for which we can obtain some small improvement by applying the technique of lazy array initialization (see Remark 7.2).

For a Σ -graph $G = (V, E)$, we denote by $G^R = (V, E^R)$ the Σ -graph obtained from G by reversing all arcs, i. e., $E^R = \{(v, x, u) \mid (u, x, v) \in E\}$.

Lemma 2.2. *Let $G = (V, E)$ be a Σ -graph. Then G^R can be computed in time $O(|G|)$.*

Proof. Since $G^R = (V, E^R)$, it is sufficient to show how the adjacency lists can be computed that represent E^R . To avoid confusion with respect to whether we talk about the Σ -graph G or the Σ -graph G^R to be constructed, we denote the x -adjacency lists by x - G -adjacency lists or x - G^R -adjacency lists, respectively.

We first move through the list for V and, for every $u \in V$, we store this node in an array along with an empty x - G^R -adjacency for every $x \in \Sigma \cup \{\varepsilon\}$. This requires time $O(|V|)$. Then, for every $u \in V$ and $x \in \Sigma \cup \{\varepsilon\}$, we move through the x - G -adjacency list for u , and for every element v that we encounter, we add u to the x - G^R -adjacency list for v . Since we can access all x - G^R -adjacency lists in constant time, we only have to add an element to some x - G^R -adjacency list for $|E|$ times. Consequently, we can construct all x - G^R -adjacency lists in total time $O(|V| + |E|) = O(|G|)$. \square

2.2 Graph Databases and Regular Path Queries

A *nondeterministic finite automaton* (NFA for short) is a tuple $M = (G, S, T)$, where $G = (V, E)$ is a Σ -graph (the nodes $q \in V$ are also called *states*), $S \subseteq V$ with $S \neq \emptyset$ is the set of *start states* and $T \subseteq V$ with $T \neq \emptyset$ is the set of *final states*. The language $\mathcal{L}(M)$ of an NFA M is the set of all labels of paths from some start state to some final state. For a Σ -graph $G = (V, E)$, any subsets $S, T \subseteq V$ with $S \neq \emptyset \neq T$ induce the NFA (G, S, T) . If $S = \{s\}$ and $T = \{t\}$ are singletons, then we also write (G, s, t) instead of $(G, \{s\}, \{t\})$.

The set RE_Σ of *regular expressions* (over Σ) is recursively defined as follows: $a \in \text{RE}_\Sigma$ for every $a \in \Sigma \cup \{\varepsilon\}$; $(\alpha \cdot \beta) \in \text{RE}_\Sigma$, $(\alpha \vee \beta) \in \text{RE}_\Sigma$, and $(\alpha)^+ \in \text{RE}_\Sigma$, for every $\alpha, \beta \in \text{RE}_\Sigma$. For any $\alpha \in \text{RE}_\Sigma$, let $\mathcal{L}(\alpha)$ be the regular language described by the regular expression α defined as usual: for every $a \in \Sigma \cup \{\varepsilon\}$, $\mathcal{L}(a) = \{a\}$, and for every $\alpha, \beta \in \text{RE}_\Sigma$, $\mathcal{L}(\alpha \cdot \beta) = \mathcal{L}(\alpha) \cdot \mathcal{L}(\beta)$, $\mathcal{L}(\alpha \vee \beta) = \mathcal{L}(\alpha) \cup \mathcal{L}(\beta)$ and $\mathcal{L}(\alpha^+) = \mathcal{L}(\alpha)^+$. We also use α^* as short hand form for $\alpha^+ \vee \varepsilon$. By $|\alpha|$, we denote the length of α represented as a string.

Proposition 2.3. *Every regular expression α can be transformed in time $O(|\alpha|)$ into an equivalent NFA $M = (G, p_0, p_f)$ with $|G| = O(|\alpha|)$.*

Proof. Let $\alpha \in \text{RE}_\Sigma$. We first construct the syntax tree T_α of α with node set V_{T_α} . Obviously, T_α has size $O(|\alpha|)$ and we can obtain V_{T_α} from α in time $O(|\alpha|)$ (for example, we can transform the expression α to prefix notation and then construct T_α while moving through the prefix notation of α from left to right). We now construct an NFA $M_\alpha = (G = (V, E), \{p_0\}, \{p_f\})$ from T_α as follows (recall that G is a Σ -graph and therefore it should adhere to our representations of Σ -graphs). We first construct an array of size $2|V_{T_\alpha}|$ that contains the nodes of $V = \{t_1, t_2 \mid t \in V_{T_\alpha}\}$ and we initialise empty x -adjacency lists for all these nodes and for every $x \in \Sigma \cup \{\varepsilon\}$. Then we move through T_α top-down and if the current node t is an inner node with two children r and s , then we do the following:

- If t corresponds to a concatenation \cdot , then we add r_1 to the ε -adjacency list of t_1 , we add s_1 to the ε -adjacency list of r_2 , and we add t_2 to the ε -adjacency list of s_2 .
- If t corresponds to an alternation \vee , then we add r_1 to the ε -adjacency list of t_1 , we add s_1 to the ε -adjacency list of t_1 , we add t_2 to the ε -adjacency list of r_2 , and we add t_2 to the ε -adjacency list of s_2 .

If t is an inner node with one child r (which means it necessarily corresponds to a $+$), then we add r_1 to the ε -adjacency list of t_1 , we add t_2 to the ε -adjacency list of r_2 , and we add t_1 to the ε -adjacency list of t_2 . If t is a leaf labelled with $x \in \Sigma \cup \{\varepsilon\}$, then we add t_2 to the x -adjacency list of t_1 . Finally, if r is the root of T_α , we relabel r_1 by p_0 and we relabel r_2 by p_f .

It can be easily verified that $\mathcal{L}(M_\alpha) = \mathcal{L}(\alpha)$. By construction $|V| = O(|T_\alpha|) = O(|\alpha|)$ and, since G has constant degree, we also have that $|G| = O(|V|) = O(|\alpha|)$. Moreover, in the construction we spend constant time per each arc that is added, so the whole construction of M_α can be done in time $O(|\alpha|)$. \square

In the following, when we speak about an *automaton* (or an NFA) for a regular expression α , we always mean an NFA equivalent to α with the properties asserted by Proposition 2.3.

A Σ -graph without ε -arcs is also called a *graph database* (over Σ); in the following, we denote graph databases by $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$. Since $V_{\mathcal{D}}$ is represented as a list, any graph database implicitly represents a linear order on $V_{\mathcal{D}}$ (i. e., the order induced by the list that represents $V_{\mathcal{D}}$), which we denote by $\preceq_{\mathcal{D}}$, or simply \preceq if \mathcal{D} is clear from the context. A graph database \mathcal{D} is *sparse* if $|E_{\mathcal{D}}| = O(|V_{\mathcal{D}}|)$.

Regular expressions q (over alphabet Σ) are interpreted as *regular path queries* (RPQ) for graph databases (over Σ). The *result* $q(\mathcal{D})$ of an RPQ q on a graph database $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$ over Σ is the set $q(\mathcal{D}) = \{(u, v) \mid u, v \in V_{\mathcal{D}}, \mathcal{L}((\mathcal{D}, u, v)) \cap \mathcal{L}(q) \neq \emptyset\}$.

If we interpret q as a *Boolean* RPQ, then the result is $q_{\text{B}}(\mathcal{D}) = \text{true}$ if $q(\mathcal{D}) \neq \emptyset$ and $q_{\text{B}}(\mathcal{D}) = \text{false}$ otherwise. We consider the RPQ-evaluation problems summarised in Table 3. By *sorted* RPQ-ENUM (or *semi-sorted* RPQ-ENUM), we denote the variant of RPQ-ENUM, where the pairs of $q(\mathcal{D})$ are to be enumerated in lexicographical order with respect to $\preceq_{\mathcal{D}}$ (or ordered only with respect to their left elements, while successive pairs with the same right element can be ordered arbitrarily, respectively).

Name	Input	Task
RPQ-BOOLE	\mathcal{D}, q	Decide whether $q_{\mathcal{B}}(\mathcal{D}) = \text{true}$.
RPQ-TEST	\mathcal{D}, q, u, v	Decide whether $(u, v) \in q(\mathcal{D})$.
RPQ-WITNESS	\mathcal{D}, q	Compute a witness $(u, v) \in q(\mathcal{D})$ or report that none exists.
RPQ-EVAL	\mathcal{D}, q	Compute the whole set $q(\mathcal{D})$
RPQ-COUNT	\mathcal{D}, q	Compute $ q(\mathcal{D}) $.
RPQ-ENUM	\mathcal{D}, q	Enumerate the whole set $q(\mathcal{D})$.

Table 3: The investigated RPQ-evaluation problems (\mathcal{D} is a graph database, q an RPQ and u, v two nodes from \mathcal{D}).

Remark 2.4. If an order \preceq' on $V_{\mathcal{D}}$ is explicitly given as a bijection $\pi : V_{\mathcal{D}} \rightarrow \{1, \dots, n\}$, then we can modify \mathcal{D} (in $O(|V_{\mathcal{D}}|)$) such that $\preceq_{\mathcal{D}} = \preceq'$. In this regard, sorted RPQ-ENUM just models the case where we wish the enumeration to be sorted according to some order. In particular, by assuming the order $\preceq_{\mathcal{D}}$ to be implicitly represented by \mathcal{D} , we *do not* hide the complexity of sorting n elements.

A graph database \mathcal{D} is *well-formed* if $V_{\mathcal{D}} = [n]$ for some $n \in \mathbb{N}$ and $\preceq_{\mathcal{D}}$ corresponds to \leq on $[n]$.

Lemma 2.5. *Let \mathcal{D} be a graph database with $|V_{\mathcal{D}}| = n$. Then we can construct in time $O(|\mathcal{D}|)$ a well-formed graph database \mathcal{D}' and an isomorphism $\pi : [n] \rightarrow V_{\mathcal{D}}$ between \mathcal{D}' and \mathcal{D} .*

Proof. Let \mathcal{D} be a graph database with $|V_{\mathcal{D}}| = n$. We initialise an array A of size n that can be addressed with the elements of $V_{\mathcal{D}}$ and the entries of which can store numbers of $[n]$, and an array B of size n that can be addressed with the elements of $[n]$ and the entries of which can store elements of $V_{\mathcal{D}}$. Furthermore, we initialise a counter $c = 1$. Then we move through the list for $V_{\mathcal{D}}$ from left to right and for each element u that we encounter, we set $A[u] = c$, $B[c] = u$ and increment c . Obviously, A describes an isomorphism $\pi : V_{\mathcal{D}} \rightarrow [n]$ and B describes π^{-1} . Moreover, this can be done in time $O(|V_{\mathcal{D}}|)$.

Next, we make a copy \mathcal{D}' of \mathcal{D} but replace each $u \in V_{\mathcal{D}}$ by the number $A[u]$ (note that \mathcal{D} is just a collection of lists that store elements of $V_{\mathcal{D}}$ along with pointers to lists). This can be done in time $O(|\mathcal{D}|)$. Moreover, \mathcal{D}' is obviously isomorphic to \mathcal{D} and B describes an isomorphism $[n] \rightarrow V_{\mathcal{D}}$. \square

Lemma 2.5 means that with an overhead of time $O(|\mathcal{D}|)$, we can always assume that our input graph databases are well-formed; in particular, note that with the isomorphism π ensured by the lemma, we can always translate elements from $q(\mathcal{D}')$, where \mathcal{D}' is the well-formed graph database isomorphic to \mathcal{D} , back to the corresponding elements of $q(\mathcal{D})$. Thus, Lemma 2.5 justifies that whenever we spend at least $O(|\mathcal{D}|)$ in some preprocessing, we can always assume that the input graph database is well-formed.

2.3 General Algorithmic Framework for RPQ-Evaluation

We assume the RAM-model with logarithmic word-size as our computational model. Let us next discuss our algorithmic framework for RPQ-evaluation. The input to our algorithms is a graph database $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$ and an RPQ q (and, for solving the problem RPQ-TEST, also a pair $(u, v) \in V_{\mathcal{D}}$).

In the case of RPQ-ENUM, the algorithms have routines `preprocess` and `enum`. Initially, `preprocess` performs some preliminary computations on the input or constructs some auxiliary data-structures; the performance of `preprocess` is measured in its running-time depending on the input size as usual (i.e., we treat `preprocess` as an individual algorithm). Then `enum` will produce an enumeration $(u_1, v_1), (u_2, v_2), \dots, (u_{\ell}, v_{\ell})$ such that $q(\mathcal{D}) = \{(u_i, v_i) \mid 1 \leq i \leq \ell\}$, no element occurs twice, and the algorithm reports when the enumeration is done. We measure the performance of `enum` in terms of its *delay*, which describes the time that (in the worst-case) elapses between enumerating two consecutive elements, between the start of the enumeration and the first element, and between the last element and the end of the enumeration (or between start and end in case

that $q(\mathcal{D}) = \emptyset$). We say that (variants of) RPQ-ENUM can be solved *with preprocessing p and delay d* , where p and d are functions bounding the preprocessing running-time and the delay. In the case that $p = O(d)$, the preprocessing complexity is absorbed by the delay; in this case, we say that (variants of) RPQ-ENUM can be solved *with delay d* and do not mention any bound on the preprocessing.

We also consider RPQ-ENUM in the *dynamic setting*, i. e., there is the possibility to perform *update* operations on the input graph database \mathcal{D} , which triggers a routine *update*. After an update and termination of the *update* routine, invoking *enum* is supposed to enumerate $q(\mathcal{D}')$, where \mathcal{D}' is the updated graph database. The performance of an algorithm for RPQ-ENUM is then measured in the running-times of routines *preprocess* (to be initially carried out only once) and *update*, as well as the delay. We only consider the following types of individual updates: inserting a new arc between existing nodes, deleting an arc, adding a new (isolated) node, deleting an (isolated) node. In particular, deleting or adding a single non-isolated node u may require a non-constant number of updates.

3 The Product Graph Approach

The PG-approach has already been informally described in the introduction; for our fine-grained perspective, we need to define it in detail. Let $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$ be a graph database over some alphabet Σ and let q be an RPQ over Σ . Furthermore, let (G_q, p_0, p_f) with $G_q = (V_q, E_q)$ be an automaton for q . Recall that, according to Proposition 2.3, G_q can be obtained in time $O(|q|)$ and it has $O(|q|)$ states and $O(|q|)$ arcs. The *product graph* of \mathcal{D} and G_q is the Σ -graph $G_{\boxtimes}(\mathcal{D}, q) = (V_{\boxtimes}(\mathcal{D}, q), E_{\boxtimes}(\mathcal{D}, q))$, where $V_{\boxtimes}(\mathcal{D}, q) = \{(u, p) \mid u \in V_{\mathcal{D}}, p \in V_q\}$ and

$$E_{\boxtimes}(\mathcal{D}, q) = \{((u, p), x, (v, p')) \mid (u, x, v) \in E_{\mathcal{D}}, (p, x, p') \in E_q\} \cup \{((u, p), \varepsilon, (u, p')) \mid u \in V_{\mathcal{D}}, (p, \varepsilon, p') \in E_q\}.$$

Remark 3.1. The arc labels in $G_{\boxtimes}(\mathcal{D}, q)$ are superfluous in the sense that we only need the underlying graph of $G_{\boxtimes}(\mathcal{D}, q)$ (see Lemma 3.3). We define it nevertheless as Σ -graph, since then all our definitions and terminology for Σ -graphs introduced above apply as well.

Lemma 3.2. $|V_{\boxtimes}(\mathcal{D}, q)| = O(|V_{\mathcal{D}}||q|)$, $|E_{\boxtimes}(\mathcal{D}, q)| = O(|\mathcal{D}||q|)$ and $G_{\boxtimes}(\mathcal{D}, q)$ can be computed in time $O(|G_{\boxtimes}(\mathcal{D}, q)|) = O(|\mathcal{D}||q|)$.

Proof. We first note that $|V_{\boxtimes}(\mathcal{D}, q)| \leq |V_{\mathcal{D}}||q|$ directly follows from the definition. Moreover, the following is also immediate by definition:

$$\begin{aligned} |\{((u, p), \varepsilon, (u, p')) \mid u \in V_{\mathcal{D}}, (p, \varepsilon, p') \in E_q\}| &= O(|V_{\mathcal{D}}||E_q|) = O(|\mathcal{D}||q|) \text{ and} \\ |\{((u, p), x, (v, p')) \mid (u, x, v) \in E_{\mathcal{D}}, (p, x, p') \in E_q\}| &= O(|E_{\mathcal{D}}||E_q|) = O(|\mathcal{D}||q|). \end{aligned}$$

Consequently, $|E_{\boxtimes}(\mathcal{D}, q)| = O(|\mathcal{D}||q|)$.

For the question how $G_{\boxtimes}(\mathcal{D}, q)$ can be computed, we have to keep in mind our list-based implementation of Σ -graphs (see Section 2). For every $u \in V_{\mathcal{D}}$ and every $p \in V_q$, we add (u, p) to the list that stores $V_{\boxtimes}(\mathcal{D}, q)$. This can be done by moving $|V_{\mathcal{D}}|$ times through the list for V_q . Thus, time $O(|V_{\mathcal{D}}||V_q|) = O(|\mathcal{D}||q|)$ is sufficient.

In order to construct the adjacency lists, we proceed as follows. Let $u \in V_{\mathcal{D}}$, $p \in V_q$ and $x \in \Sigma$. Then we add all (v, p') to the x -adjacency list of (u, p) , where v is an element of the x -adjacency list of u , and p' is an element of the x -adjacency list of p . Moreover, we add all (u, p') to the ε -adjacency list of (u, p) , where p' is an element of the ε -adjacency list of p . This can be done by moving once through the lists of adjacency lists for $E_{\mathcal{D}}$ and, for each encountered element, to move through the lists of adjacency lists for E_q . Since each insertion to a list can be done in constant time, the whole procedure can be carried out in time $O(|\Sigma||\mathcal{D}||q|) = O(|\mathcal{D}||q|)$. \square

The following lemma, which is an immediate consequence of the construction, shows how $G_{\boxtimes}(\mathcal{D}, q)$ can be used for solving RPQ-evaluation tasks (recall that E^* is the reflexive-transitive closure of the underlying unlabelled graph).

Lemma 3.3. *For every $u, v \in V_{\mathcal{D}}$, $(u, v) \in q(\mathcal{D})$ if and only if $((u, p_0), (v, p_f)) \in (E_{\boxtimes}(\mathcal{D}, q))^*$.*

4 Fine-grained Complexity and Conditional Lower Bounds

We now state several computational problems along with hypotheses regarding their complexity, which are commonly used in the framework of fine-grained complexity to obtain conditional lower bounds. We discuss some details and give background information later on.

- Orthogonal Vectors (OV): Given sets A, B each containing n Boolean-vectors of dimension d , check whether there are vectors $\vec{a} \in A$ and $\vec{b} \in B$ that are orthogonal.
OV-Hypothesis: For every $\epsilon > 0$ there is no algorithm solving OV in time $O(n^{2-\epsilon} \text{poly}(d))$.
- Boolean Matrix Multiplication (BMM): Given Boolean $n \times n$ matrices A, B , compute $A \times B$.
com-BMM-Hypothesis: For every $\epsilon > 0$ there is no combinatorial algorithm that solves BMM in time $O(n^{3-\epsilon})$.
- Sparse Boolean Matrix Multiplication (SBMM): Like BMM, but all matrices are represented as sets $\{(i, j) \mid A[i, j] = 1\}$ of 1-entries.
SBMM-Hypothesis: There is no algorithm that solves SBMM in time $O(m)$ (where m is the total number of 1-entries, i. e., $m = |\{(i, j) \mid A[i, j] = 1\}| + |\{(i, j) \mid B[i, j] = 1\}| + |\{(i, j) \mid (A \times B)[i, j] = 1\}|$).
- Online Matrix-Vector Multiplication (OMv): Given Boolean $n \times n$ -matrix M and a sequence $\vec{v}^1, \vec{v}^2, \dots, \vec{v}^n$ of n -dimensional Boolean vectors, compute sequence $M\vec{v}^1, M\vec{v}^2, \dots, M\vec{v}^n$, where $M\vec{v}^i$ is produced as output before \vec{v}^{i+1} is received as input.
OMv-Hypothesis: For every $\epsilon > 0$ there is no algorithm that solves OMv in time $O(n^{3-\epsilon})$.

We will reduce these problems to variants of RPQ evaluation problems in such a way that algorithms with certain running-times for RPQ evaluation would break the corresponding hypotheses mentioned above. Thus, we obtain lower bounds for RPQ enumeration that are conditional to these hypotheses. In the following, we give a very brief overview of the relevance of these problems and corresponding hypotheses in fine-grained complexity.

The problem OV can be solved by brute-force in time $O(n^2d)$ and the hypothesis that there is no subquadratic algorithm is well-established. It exists in slightly different variants and has been formulated in several different places in the literature (e. g., [21, 22, 51]). The variant used here is sometimes referred to as *moderate dimension* OV-hypothesis in contrast to *low dimension* variants, where d can be assumed to be rather small in comparison to n . The relevance of the OV-hypothesis is due to the fact that it is implied by the Strong Exponential Time Hypothesis (SETH) [47, 48], and therefore it is a convenient tool to prove SETH lower bounds that has been applied in various contexts.

One of the most famous computational problems is BMM, which, unfortunately, is a much less suitable basis for conditional lower bounds. The straightforward algorithm solves it in time $O(n^3)$, but there are *fast matrix multiplication* algorithms that run in time $O(n^{2.373})$ [50, 29]. It is unclear how much further this exponent can be decreased and there is even belief that BMM can be solved in time $n^{2+o(1)}$ (see [49] and [13, Section 6]). However, these theoretically fast algorithms cannot be considered efficient in a practical sense, which motivates the mathematically informal notion of “*combinatorial*” algorithms (see, e. g., [53]).¹ So far, no truly subcubic *combinatorial* BMM-algorithm exists and it has been shown in [53] that BMM is contained in a class of problems (including other prominent examples like Triangle Finding (also mentioned below) and Context-Free Grammar Parsing) which are all equivalent in the sense that if one such problem is solvable in truly subcubic time by a combinatorial algorithm, then all of them are. Consequently, it is often possible to argue that the existence of a certain combinatorial algorithm for some problem would imply a major (and unlikely) algorithmic breakthrough with respect to BMM, Parsing, Triangle Finding, etc. Despite the defect of relying on the vague notion of *combinatorial* algorithms, this lower bound technique is a common approach in fine-grained complexity (see, e. g., [53, 31, 2, 3, 1, 32]). Whenever we use the com-BMM-hypothesis, our reductions will always be combinatorial, which is necessary; moreover, whenever we say that a certain running

¹The term “combinatorial algorithm” is not well-defined, but intuitively such algorithms have running-times with low constants in the O-notation, and are feasibly implementable.

time cannot be achieved unless the **com-BMM**-hypothesis fails, we mean, of course, that it cannot be achieved by a combinatorial algorithm.

In order to make **BMM** suitable as base problem for conditional lower bounds (that does not rely on *combinatorial* algorithms) one can formulate the weaker (i. e., more plausible) hypothesis that **BMM** cannot be solved in time linear in the number of 1-entries of the matrices (therefore called *sparse BMM* since matrices are represented in a sparse way); see [4, 55]. Another approach is to require the output matrix $A \times B$ to be computed column by column, i. e., formulating it as the online-version **OMv**. For **OMv**, subcubic algorithms are not known and would yield several major algorithmic breakthroughs (see [31]).

A convenient tool to deal with **BMM** is the problem **Triangle**: check whether a given undirected graph G has a triangle. This is due to the fact that these two problems are subcubic equivalent with respect to combinatorial algorithms (see [53]), i. e., the **com-BMM**-hypothesis fails if and only if **Triangle** can be solved by a combinatorial algorithm in time $O(n^{3-\epsilon})$ for some $\epsilon > 0$. Thus, for lower bounds conditional to the **com-BMM**-hypothesis, we can make use of both these problems. There is also a (non-combinatorial) **Triangle**-hypothesis that states that **Triangle** cannot be solved in linear time in the number of edges, but we were not able to apply it in the context of RPQ-evaluation (see [2] for different variants of **Triangle**).

5 Bounds for the Non-Enumeration Problem Variants

We now investigate how well the PG-approach performs with respect to the non-enumeration variants of RPQ-evaluation, and we give some evidence that, in most cases, it can be considered optimal or almost optimal (subject to the algorithmic hypotheses of Section 4).

5.1 Boolean Evaluation, Testing and Computing a Witness

It is relatively straightforward to see that the problems **RPQ-TEST** and **RPQ-BOOLE** are equivalent and can both be reduced to **RPQ-WITNESS**. Hence, upper bounds for **RPQ-WITNESS** and lower bounds for **RPQ-TEST** or **RPQ-BOOLE** automatically apply to all three problem variants, which simplifies the proofs for such bounds. We shall now formally prove this.

Lemma 5.1. *Let (\mathcal{D}, q) be an **RPQ-BOOLE**-instance. Then we can construct an equivalent **RPQ-TEST**-instance (\mathcal{D}', q', u, v) with $|\mathcal{D}'| = O(|\mathcal{D}|)$ and $|q'| = O(|q|)$ in time $O(|\mathcal{D}| + |q|)$.*

Proof. Let \mathcal{D} and q be defined over Σ . We transform \mathcal{D} into a graph database \mathcal{D}' over $\Sigma \cup \{\#\}$ (where $\# \notin \Sigma$) by adding new nodes u and v , and new arcs $(u, \#, x)$ and $(x, \#, v)$ for every $x \in V_{\mathcal{D}}$. Moreover, we set $q' = \#q\#$. This construction can be carried out in time $O(|\mathcal{D}| + |q|)$ and we can also note that $|\mathcal{D}'| = O(|\mathcal{D}|)$ and $|q'| = O(|q|)$.

It remains to show that $q_{\mathbb{B}}(\mathcal{D}) = \text{true}$ if and only if $(u, v) \in q'(\mathcal{D}')$. If $q_{\mathbb{B}}(\mathcal{D}) = \text{true}$, then there is some $(u', v') \in q(\mathcal{D})$, which means that in \mathcal{D} there is a path u', \dots, v' that is labelled with a word $w \in \mathcal{L}(q)$. Since there are arcs $(u, \#, u')$ and $(v', \#, v)$ in \mathcal{D}' , there is a path u, u', \dots, v', v in \mathcal{D}' that is labelled with $\#w\# \in \mathcal{L}(q')$. Therefore, $(u, v) \in q'(\mathcal{D}')$. On the other hand, if $(u, v) \in q'(\mathcal{D}')$, then there is some path u, u', \dots, v', v in \mathcal{D}' that is labelled with $\#w\# \in \mathcal{L}(q')$, so therefore also $(u', v') \in q(\mathcal{D})$ and $q_{\mathbb{B}}(\mathcal{D}) = \text{true}$ (note that $u' = v'$ and therefore $w = \varepsilon$ is also possible). \square

Lemma 5.2. *Let (\mathcal{D}, q, u, v) be an **RPQ-TEST**-instance. Then we can construct an equivalent **RPQ-BOOLE**-instance (\mathcal{D}', q') with $|\mathcal{D}'| = O(|\mathcal{D}|)$ and $|q'| = O(|q|)$ in time $O(|\mathcal{D}| + |q|)$.*

Proof. Let \mathcal{D} and q be defined over Σ . We transform \mathcal{D} into a graph database \mathcal{D}' over $\Sigma \cup \{\#\}$ (where $\# \notin \Sigma$) by adding new nodes s and t with arcs $(s, \#, u)$ and $(v, \#, t)$, and we define $q' = \#q\#$. This construction can be carried out in time $O(|\mathcal{D}| + |q|)$ and we can also note that $|\mathcal{D}'| = O(|\mathcal{D}|)$ and $|q'| = O(|q|)$.

It remains to show that $(u, v) \in q(\mathcal{D})$ if and only if $q'_{\mathbb{B}}(\mathcal{D}') = \text{true}$. If $(u, v) \in q(\mathcal{D})$, then there is a path from u to v in \mathcal{D} that is labelled with a word from $\mathcal{L}(q)$. Since there are arcs $(s, \#, u)$ and $(v, \#, t)$ in \mathcal{D}' , there is a path in \mathcal{D}' from s to t labelled with a word from $\mathcal{L}(\#q\#)$; thus $(s, t) \in q'(\mathcal{D}')$ and therefore $q'_{\mathbb{B}}(\mathcal{D}') = \text{true}$. On the other hand, if $q'_{\mathbb{B}}(\mathcal{D}') = \text{true}$, then we can conclude that $(s, t) \in q'(\mathcal{D}')$, which is due to the fact that $q' = \#q\#$, q does not contain any

occurrence of $\#$ and the only arcs labelled with $\#$ have source s and target t . This implies that there is a path s, u, \dots, v, t labelled with a word $\#w\#$ with $w \in \mathcal{L}(q)$, which implies that there is a path from u to v labelled with a word from $\mathcal{L}(q)$, which means that $(u, v) \in q(\mathcal{D})$. \square

Theorem 5.3. *Let $f : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ be some polynomial function with $f(n_1, n_2) = \Omega(n_1 + n_2)$.*

1. RPQ-TEST can be solved in time $f(|\mathcal{D}|, |q|)$ if and only if RPQ-BOOLE can be solved in time $f(|\mathcal{D}|, |q|)$.
2. If RPQ-WITNESS can be solved in time $f(|\mathcal{D}|, |q|)$, then RPQ-BOOLE and RPQ-TEST can be solved in time $f(|\mathcal{D}|, |q|)$.

Proof. Lemmas 5.1 and 5.2 directly imply Point 1. Moreover, an algorithm that solves RPQ-WITNESS implicitly solves RPQ-BOOLE as well. Together with Point 1, this proves Point 2. \square

We are now ready to prove that the PG-approach yields the following upper bound.

Theorem 5.4. *RPQ-TEST, RPQ-BOOLE and RPQ-WITNESS can be solved in time $O(|\mathcal{D}||q|)$.*

Proof. We only show the upper bound for RPQ-WITNESS (due to Theorem 5.3 it applies to RPQ-TEST and RPQ-BOOLE as well). To this end, let $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$ be a graph database over Σ and let q be an RPQ over Σ . We construct $G_{\boxtimes}(\mathcal{D}, q)$, which, according to Lemma 3.2, can be done in time $O(|\mathcal{D}||q|)$. In the following considerations, we interpret $G_{\boxtimes}(\mathcal{D}, q)$ as its underlying non-labelled graph.

We add to $G_{\boxtimes}(\mathcal{D}, q)$ a node v_{source} with an arc to each (u, p_0) with $u \in V_{\mathcal{D}}$. This can be done in time $O(|V_{\mathcal{D}}||q|)$ as follows. First, we add v_{source} to the list that represents $V_{\boxtimes}(\mathcal{D}, q)$, which requires constant time. Then, we move through the list that represents $V_{\boxtimes}(\mathcal{D}, q)$ and every node (u, p_0) that we encounter is added to the adjacency list for v_{source} .

Next, we perform a special kind of BFS in v_{source} . First, we construct an array S of size $|V_{\boxtimes}(\mathcal{D}, q)|$ the entries of which can store values from $V_{\boxtimes}(\mathcal{D}, q) \cup \{0\}$ and which can be addressed by the nodes from $V_{\boxtimes}(\mathcal{D}, q)$. Moreover, S is initialised with every entry storing 0. Then, for every $u \in V_{\mathcal{D}}$, we set $S[(u, p_0)] = u$. This can be done in time $O(|V_{\boxtimes}(\mathcal{D}, q)|) = O(|\mathcal{D}||q|)$. We perform a BFS from v_{source} and whenever we traverse an arc $((u, p), (u', p'))$, we set $S[(u', p')] = S[(u, p)]$. This can be done in time $O(|G_{\boxtimes}(\mathcal{D}, q)|) = O(|\mathcal{D}||q|)$. Since $S[(u, p_0)] = u$ for every $u \in V_{\mathcal{D}}$, we can conclude by induction that, for every $(v, p) \in V_{\boxtimes}(\mathcal{D}, q)$, if (v, p) is reachable from some node (u, p_0) , then $S[(v, p)] = u'$ for some $u' \in V_{\mathcal{D}}$ such that (v, p) is reachable from (u', p_0) , and if (v, p) is not reachable from any node (u, p_0) , then $S[(v, p)] = 0$. Consequently, if there is some $v \in V_{\mathcal{D}}$ with $S[(v, p_f)] = u \neq 0$, then there is a path from (u, p_0) to (v, p_f) in $G_{\boxtimes}(\mathcal{D}, q)$, which, according to Lemma 3.3, means that $(u, v) \in q(\mathcal{D})$ and therefore, we can produce the output (u, v) . On the other hand, if there is no $v \in V_{\mathcal{D}}$ with $S[(v, p_f)] \neq 0$, then there are no $u, v \in V_{\mathcal{D}}$ with a path from (u, p_0) to (v, p_f) ; thus, $q(\mathcal{D}) = \emptyset$. Checking whether there is some $v \in V_{\mathcal{D}}$ with $S[(v, p_f)] \neq 0$ can be done in time $O(|V_{\mathcal{D}}|)$. \square

More interestingly, we can complement this upper bound with lower bounds as follows.

Theorem 5.5. *If any of the problems RPQ-TEST, RPQ-BOOLE and RPQ-WITNESS can be solved in time $O(|\mathcal{D}|^{2-\epsilon} + |q|^2)$ or $O(|\mathcal{D}|^2 + |q|^{2-\epsilon})$ for some $\epsilon > 0$, then the OV-hypothesis fails. This lower bound also holds for the restriction to sparse graph databases.*

Proof. We prove the lower bound for RPQ-TEST only, since by Theorem 5.3 it also applies to RPQ-WITNESS and RPQ-BOOLE. We first devise a general reduction from the OV-problem to RPQ-TEST (which is similar to the reduction from [6] used for proving conditional lower bounds of regular expression matching):

Let $A = \{\vec{a}^1, \dots, \vec{a}^n\}$ and $B = \{\vec{b}^1, \dots, \vec{b}^n\}$ be an instance for the OV-problem. We define an RPQ q and a graph database \mathcal{D} over the alphabet $\Sigma = \{0, 1, \#\}$ as follows. For every $i \in [n]$, let $w_i = \vec{b}^i[1]\vec{b}^i[2] \dots \vec{b}^i[d]$, and let $q = \#(w_1 \vee w_2 \vee \dots \vee w_n)\#$. Moreover, let $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$, where $V_{\mathcal{D}}$ contains nodes s and t and, for every $i \in [n]$, nodes $v_{i,0}, v_{i,1}, \dots, v_{i,d}$. For every $i \in [n]$ and $j \in \{0\} \cup [d-1]$, there is an arc from $v_{i,j}$ to $v_{i,j+1}$ labelled with 0 and, if $\vec{a}^i[j+1] = 0$, also an arc from $v_{i,j}$ to $v_{i,j+1}$ labelled with 1. Finally, for every $i \in [n]$, there are arcs labelled with $\#$ from

s to $v_{i,0}$ and from $v_{i,d}$ to t . It can be easily verified that there are orthogonal $\vec{a} \in A$ and $\vec{b} \in B$ if and only if $(s, t) \in q(\mathcal{D})$. Moreover, $|\mathcal{D}| = O(|A|d) = O(nd)$ and $|q| = O(|B|d) = O(nd)$, and, furthermore, \mathcal{D} and q can also be constructed in time $O(nd)$, and, since $|E_{\mathcal{D}}| = O(|V_{\mathcal{D}}|)$, \mathcal{D} is a sparse graph database.

We now assume that RPQ-TEST can be solved in time $O(|\mathcal{D}|^2 + |q|^{2-\epsilon})$ for some $\epsilon > 0$. Let again $A = \{\vec{a}^1, \dots, \vec{a}^n\}$ and $B = \{\vec{b}^1, \dots, \vec{b}^n\}$ be an instance for the OV-problem, and let ϵ' be arbitrarily chosen with $0 < \epsilon' < \epsilon$. We divide A into $A_1, A_2, \dots, A_{\lceil n^{\epsilon'} \rceil}$ with $|A_i| = \lceil n^{1-\epsilon'} \rceil$ for every $i \in [\lceil n^{\epsilon'} \rceil]$ and such that $\bigcup_{i=1}^{\lceil n^{\epsilon'} \rceil} A_i = A$. Obviously, (A, B) is a positive OV-instance if and only if at least one of $(A_1, B), (A_2, B), \dots, (A_{\lceil n^{\epsilon'} \rceil}, B)$ is a positive OV-instance. We can now separately reduce each (A_i, B) to an RPQ-TEST-instance (\mathcal{D}_i, q) as described above and we note that $|\mathcal{D}_i| = O(|A_i|d) = O(\lceil n^{1-\epsilon'} \rceil d)$ and $|q| = O(|B|d) = O(nd)$. Then we solve each instance in the assumed time bound $O(|\mathcal{D}|^2 + |q|^{2-\epsilon})$, which requires total time of $O(\lceil n^{\epsilon'} \rceil ((\lceil n^{1-\epsilon'} \rceil d)^2 + (nd)^{2-\epsilon}) = O((n^{2-\epsilon'} + n^{2-\epsilon+\epsilon'}) \text{poly}(d)) = O(n^{2-(\epsilon-\epsilon')} \text{poly}(d))$, where $(\epsilon - \epsilon') > 0$. This contradicts the OV-hypothesis.

The assumption that RPQ-TEST can be solved in time $O(|\mathcal{D}|^{2-\epsilon} + |q|^2)$ for some $\epsilon > 0$ can be handled analogously. We divide again A into $A_1, A_2, \dots, A_{\lceil n^{\epsilon'} \rceil}$ as described above, but then we reduce the OV-instances $(B, A_1), (B, A_2), \dots, (B, A_{\lceil n^{\epsilon'} \rceil})$ to the RPQ-TEST-instances $(\mathcal{D}, q_1), (\mathcal{D}, q_2), \dots, (\mathcal{D}, q_{\lceil n^{\epsilon'} \rceil})$. Note that $|\mathcal{D}| = O(nd)$ and $|q_i| = O(|A_i|d) = O(\lceil n^{1-\epsilon'} \rceil d)$. By assumption, each of these instances can be solved in time $O(|\mathcal{D}|^{2-\epsilon} + |q_i|^2) = O((nd)^{2-\epsilon} + (\lceil n^{1-\epsilon'} \rceil d)^2)$, which again leads to a total running-time of $O(\lceil n^{\epsilon'} \rceil ((nd)^{2-\epsilon} + (\lceil n^{1-\epsilon'} \rceil d)^2)) = O(n^{2-(\epsilon-\epsilon')} \text{poly}(d))$, where $(\epsilon - \epsilon') > 0$. This contradicts the OV-hypothesis. \square

Since $(|\mathcal{D}||q|)^{1-\epsilon} \leq ((\max\{|\mathcal{D}|, |q|\})^2)^{1-\epsilon} = \max\{|\mathcal{D}|^{2-2\epsilon}, |q|^{2-2\epsilon}\} \leq |\mathcal{D}|^{2-\epsilon} + |q|^2$, Theorem 5.5 also rules out running times of the form $O((|\mathcal{D}||q|)^{1-\epsilon})$ and $O(\max\{|\mathcal{D}|, |q|\}^{2-\epsilon})$, but does not exclude a running time of $O(|\mathcal{D}|^{2-\epsilon} + f(|q|))$ with $f(|q|) = \Omega(|q|^2)$. However, for $\epsilon < 1$ this is super-linear in $|\mathcal{D}|$ (and therefore inferior to $O(|\mathcal{D}||q|)$ under the assumption $|q| \ll |\mathcal{D}|$), and for $\epsilon = 1$, we would obtain $O(|\mathcal{D}| + f(|q|))$, which, under the assumption $|q| \ll |\mathcal{D}|$, is a small and arguably negligible asymptotic improvement over $O(|\mathcal{D}||q|)$.

If the size of \mathcal{D} is expressed in terms of $|V_{\mathcal{D}}|$, then Theorem 5.4 also gives an upper bound of $O(|V_{\mathcal{D}}|^2|q|)$. In this regard, we can show the following lower bound.

Theorem 5.6. *If any of the problems RPQ-TEST, RPQ-BOOLE and RPQ-WITNESS can be solved in time $O(|V_{\mathcal{D}}|^{3-\epsilon} + |q|^{3-\epsilon})$ for some $\epsilon > 0$, then the com-BMM-hypothesis fails.*

Proof. We prove the lower bound for RPQ-BOOLE only, since by Theorem 5.3 it also applies to RPQ-WITNESS and RPQ-BOOLE.

We assume that RPQ-BOOLE can be solved by a combinatorial algorithm in time $O(|V_{\mathcal{D}}|^{3-\epsilon} + |q|^{3-\epsilon})$ for some $\epsilon > 0$. We can then solve Triangle on some instance $G = (V, E)$ with $V = \{v_1, v_2, \dots, v_n\}$ as follows (the result will then follow from the combinatorial subcubic equivalence of BMM and Triangle (see Section 4)). We construct the graph database \mathcal{D} over $\{\mathbf{a}, \#\}$ with nodes

$$V_{\mathcal{D}} = \{s', t'\} \cup \{s_j, t_j, \mid 1 \leq j \leq n\} \cup \{(u, i) \mid 0 \leq i \leq 3, u \in V\}$$

and arcs

$$\begin{aligned} E_{\mathcal{D}} = & \{((u, i), \mathbf{a}, (u', i+1)) \mid 0 \leq i \leq 2, (u, u') \in E\} \cup \\ & \{(s', \#, s_1), (s_i, \mathbf{a}, s_{i+1}), (t_i, \mathbf{a}, t_{i+1}), (t_n, \#, t') \mid 1 \leq i \leq n-1\} \cup \\ & \{(s_i, \mathbf{a}, (v_i, 0)), ((v_i, 3), \mathbf{a}, t_i) \mid 1 \leq i \leq n\}. \end{aligned}$$

Furthermore, we define the RPQ $q = \#\mathbf{a}^{n+4}\#$.

By definition, $q(\mathcal{D}) = \{(s', t')\}$ or $q(\mathcal{D}) = \emptyset$. We call any path p from s' to t' an i - j -path for some $i, j \in [n]$ if there are $\ell_1, \ell_2 \in [n]$ such that $p = (s', s_1, \dots, s_i, v_{i,0}, v_{\ell_1,1}, v_{\ell_2,2}, v_{j,3}, t_j, \dots, t_n, t')$. It can be easily seen that any path from s' to t' is an i - j -path for some $i, j \in [n]$, and that any i - j -path is labelled with $\#\mathbf{a}^i \mathbf{a}^3 \mathbf{a}^{n-j+1} \# = \#\mathbf{a}^{i+n-j+4} \#$. Hence, $q(\mathcal{D}) \neq \emptyset$ if and only if there is an i - i -path for some $i \in [n]$. Since, for every $i \in [n]$, there is an i - i -path if and only if there is a

path $(v_{i,0}, v_{\ell_1,1}, v_{\ell_2,2}, v_{i,3})$, we see that, for every $i \in [n]$, there is an i - i -path if and only if G has a triangle that contains v_i . Consequently, $q(\mathcal{D}) \neq \emptyset$ if and only if G has a triangle.

This means that by first constructing \mathcal{D} and q in time $O(|\mathcal{D}| + |q|) = O(|G| + |V|)$ and then solving RPQ-BOOLE for instance (\mathcal{D}, q) in the assumed time $O(|V_{\mathcal{D}}|^{3-\epsilon} + |q|^{3-\epsilon})$ for some $\epsilon > 0$, means that we can solve Triangle by a combinatorial algorithm in time $O(|V_{\mathcal{D}}|^{3-\epsilon} + |q|^{3-\epsilon}) = O(|V|^{3-\epsilon})$.

With Lemma 5.3, we can conclude that the assumptions that RPQ-TEST or RPQ-WITNESS can be solved in time $O(|V_{\mathcal{D}}|^{3-\epsilon} + |q|^{3-\epsilon})$ for some $\epsilon > 0$, leads to a combinatorial $O(|V|^{3-\epsilon})$ algorithm for Triangle as well. Thus, due to the combinatorial subcubic equivalence of BMM and Triangle (see Section 4), there is a combinatorial $O(n^{3-\epsilon})$ algorithm for BMM, contradicting the com-BMM-hypothesis. \square

Since $O((|\mathcal{D}||q|)^{1-\epsilon}) \subseteq O((|V_{\mathcal{D}}|^2|q|)^{1-\epsilon}) \subseteq O(|V_{\mathcal{D}}|^{3-\epsilon} + |q|^{3-\epsilon})$, such running-times are also ruled out under the com-BMM-hypothesis. Especially, a combinatorial algorithm with running-time $O((|\mathcal{D}||q|)^{1-\epsilon})$ refutes *both* the OV- and the com-BMM-hypothesis; thus, such an algorithm does not exist provided that at least one of these hypotheses is true (basing lower bounds on several hypotheses is common in fine-grained complexity, see, e. g., [3]).

The lower bounds discussed above are only meaningful for combined complexity. However, the upper bound of Theorem 5.4 already yields the optimum of *linear* data complexity.

5.2 Full Evaluation and Counting

The following upper bound is again a straightforward application of the PG-approach.

Theorem 5.7. *RPQ-EVAL can be solved in time $O(|V_{\mathcal{D}}||\mathcal{D}||q|)$.*

Proof. Let $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$ be a graph database over Σ and let q be an RPQ over Σ . We construct $G_{\boxtimes}(\mathcal{D}, q)$, which, according to Lemma 3.2, can be done in time $O(|\mathcal{D}||q|)$. We interpret $G_{\boxtimes}(\mathcal{D}, q)$ as its underlying non-labelled graph. According to Lemma 3.3, we can now compute $q(\mathcal{D})$ by performing a BFS from each node (u, p_0) with $u \in V_{\mathcal{D}}$. Each such BFS requires time $O(|G_{\boxtimes}(\mathcal{D}, q)|) = O(|\mathcal{D}||q|)$, which means that the total running-time is $O(|V_{\mathcal{D}}||G_{\boxtimes}(\mathcal{D}, q)|) = O(|V_{\mathcal{D}}||\mathcal{D}||q|)$. \square

Instead of using graph-searching techniques on $G_{\boxtimes}(\mathcal{D}, q)$, we could also compute the complete transitive closure of $G_{\boxtimes}(\mathcal{D}, q)$ with fast matrix multiplication.

Theorem 5.8. *If BMM can be solved in time $O(n^\omega)$ with $\omega \geq 2$, then RPQ-EVAL can be solved in time $O(|V_{\mathcal{D}}|^\omega |q|^\omega)$.*

Proof. We assume that BMM can be solved in time $O(n^\omega)$. Let $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$ be a graph database over Σ and let q be an RPQ over Σ . We construct $G_{\boxtimes}(\mathcal{D}, q)$, which, according to Lemma 3.2, can be done in time $O(|\mathcal{D}||q|)$, and we interpret it as its underlying non-labelled graph. Obviously, we can turn the adjacency list-based representation of $G_{\boxtimes}(\mathcal{D}, q)$ into an adjacency matrix-based representation in time $O(|V_{\boxtimes}(\mathcal{D}, q)|^2) = O((|V_{\mathcal{D}}||q|)^2) = O((|V_{\mathcal{D}}||q|)^\omega)$. Then, we compute the transitive closure $(E_{\boxtimes}(\mathcal{D}, q))^*$ of $G_{\boxtimes}(\mathcal{D}, q)$, which can be done in time $O(|V_{\boxtimes}(\mathcal{D}, q)|^\omega) = O((|V_{\mathcal{D}}||q|)^\omega)$ (see [42]). In order to obtain $q(\mathcal{D})$, it is sufficient to go through all elements $((u, p), (v, p')) \in (E_{\boxtimes}(\mathcal{D}, q))^*$ and add (u, v) to a new set if and only if $p = p_0$ and $p' = p_f$ (see Lemma 3.3). Since $|(E_{\boxtimes}(\mathcal{D}, q))^*| \leq |V_{\boxtimes}(\mathcal{D}, q)|^2 = O(|V_{\boxtimes}(\mathcal{D}, q)|^\omega) = O((|V_{\mathcal{D}}||q|)^\omega)$, this can be done in time $O((|V_{\mathcal{D}}||q|)^\omega)$. \square

We mention this theoretical upper bound for completeness, but stress the fact that our main interest lies in *combinatorial* algorithms. In addition to the limitation that algorithms for fast matrix multiplication are not practical, we also observe that the approach of Theorem 5.8 is only better if the graph database is not too sparse, i. e., only if $|V_{\mathcal{D}}||\mathcal{D}| = \Omega(|V_{\mathcal{D}}|^\omega)$.

Next, we investigate the question whether $O(|V_{\mathcal{D}}||\mathcal{D}||q|)$ is optimal for RPQ-EVAL at least with respect to combinatorial algorithms. Since for RPQ-EVAL the PG-approach does not yield an algorithm that is linear in data complexity (like it was the case with respect to the problems of Section 5.1), the question arises whether the $|V_{\mathcal{D}}||\mathcal{D}|$ part can be improved at the cost of spending more time in $|q|$. It seems necessary that respective *data complexity lower bounds* need reductions

that do not use q to represent a non-constant part of the instance, as it was the case for both the OV and the Triangle reduction from Section 5.1.

It is not difficult to see that multiplying two $n \times n$ Boolean matrices reduces to RPQ-EVAL as formally stated by the next lemma.

Lemma 5.9. *If RPQ-EVAL can be solved in time $O(|V_{\mathcal{D}}|^{\omega} f(|q|))$ for some function f and $\omega \geq 2$, then BMM can be solved in time $O(n^{\omega})$.*

Proof. We assume that RPQ-EVAL can be solved in time $O(|V_{\mathcal{D}}|^{\omega} f(|q|))$ for some function f . Let A and B be $n \times n$ Boolean matrices. Then we construct the graph database $\mathcal{D}_{A,B}$ over $\{\mathbf{a}\}$ with $V_{\mathcal{D}_{A,B}} = \{(i, 0), (i, 1), (i, 2) \mid i \in [n]\}$ and $E_{\mathcal{D}_{A,B}} = \{((i, 0), \mathbf{a}, (j, 1)) \mid A[i, j] = 1\} \cup \{((i, 1), \mathbf{a}, (j, 2)) \mid B[i, j] = 1\}$, and the RPQ $q = \mathbf{aa}$. Obviously, $q(\mathcal{D}_{A,B}) \subseteq \{(i, 0), (j, 2) \mid i, j \in [n]\}$ and $((i, 0), (j, 2)) \in q(\mathcal{D}_{A,B})$ if and only if $(A \times B)[i, j] = 1$. Hence, we can construct $A \times B$ from $q(\mathcal{D}_{A,B})$ in time $O(|q(\mathcal{D}_{A,B})|) = O(n^2)$. By assumption, the set $q(\mathcal{D})$ can be computed in time $O(|V_{\mathcal{D}_{A,B}}|^{\omega} f(|q|)) = O(n^{\omega} f(2)) = O(n^{\omega})$. \square

From this, we can conclude the following lower bound.

Theorem 5.10. *If RPQ-EVAL can be solved in time $O((|V_{\mathcal{D}}| |\mathcal{D}|)^{1-\epsilon} f(|q|))$ for some function f and some $\epsilon > 0$, then the com-BMM-hypothesis fails.*

Proof. If RPQ-EVAL can be solved by a combinatorial algorithm in time $O((|V_{\mathcal{D}}| |\mathcal{D}|)^{1-\epsilon} f(|q|))$ for some function f and some $\epsilon > 0$, then, since $O((|V_{\mathcal{D}}| |\mathcal{D}|)^{1-\epsilon} f(|q|)) = O(|V_{\mathcal{D}}|^{3-\epsilon} f(|q|))$, Lemma 5.9 implies that BMM can be solved in time $O(n^{3-\epsilon})$. Thus, the com-BMM-hypothesis fails. \square

If we drop the restriction to combinatorial algorithms, we can nevertheless show (with more or less the same construction) that linear time in data complexity is impossible, unless the SBMM-hypothesis fails. However, since the size of the output $q(\mathcal{D})$ might be super-linear in $|\mathcal{D}|$, we should interpret *linear* as linear in $|\mathcal{D}| + |q(\mathcal{D})|$.

Lemma 5.11. *If RPQ-EVAL can be solved in time $O((|q(\mathcal{D})| + |\mathcal{D}|)^{\omega} f(|q|))$ for some function f and $\omega \geq 1$, then SBMM can be solved in time $O(m^{\omega})$.*

Proof. We assume that RPQ-EVAL can be solved in time $O((|q(\mathcal{D})| + |\mathcal{D}|)^{\omega} f(|q|))$ for some function f . Let A and B be $n \times n$ Boolean matrices given as sets of their 1-entries and let m be the total number of 1-entries in A , B and $A \times B$. Then we can construct a graph database $\mathcal{D}_{A,B}$ over $\{\mathbf{a}\}$ as follows. For every $i \in [n]$, the set $V_{\mathcal{D}_{A,B}}$ contains a node $(i, 0)$ if the i^{th} row of A contains at least one 1-entry, a node $(i, 1)$ if the i^{th} column of A contains at least one 1-entry or the i^{th} row of B contains at least one 1-entry, and a node $(i, 2)$ if the i^{th} column of B contains at least one 1-entry. The set of arcs is defined by $E_{\mathcal{D}_{A,B}} = \{((i, 0), \mathbf{a}, (j, 1)) \mid A[i, j] = 1\} \cup \{((i, 1), \mathbf{a}, (j, 2)) \mid B[i, j] = 1\}$. We observe that $|E_{\mathcal{D}_{A,B}}| = O(m)$ and, since every node in $V_{\mathcal{D}_{A,B}}$ has degree at least 1, $|V_{\mathcal{D}_{A,B}}| = O(|E_{\mathcal{D}_{A,B}}|)$. Consequently, $|\mathcal{D}_{A,B}| = O(m)$ and $\mathcal{D}_{A,B}$ can be constructed in time $O(m)$. Further, we define $q = \mathbf{aa}$.

We can observe that $q(\mathcal{D}_{A,B}) \subseteq \{(i, 0), (j, 2) \mid i, j \in [n]\}$ and, for every $i, j \in [n]$, $((i, 0), (j, 2)) \in q(\mathcal{D}_{A,B})$ if and only if $(A \times B)[i, j] = 1$. Thus, $|q(\mathcal{D}_{A,B})| = O(m)$ and we can obtain a set of exactly the 1-entries of $A \times B$ from $q(\mathcal{D}_{A,B})$ in time $O(|q(\mathcal{D}_{A,B})|) = O(m)$.

By assumption, we can compute $q(\mathcal{D}_{A,B})$ in time $O((|q(\mathcal{D}_{A,B})| + |\mathcal{D}_{A,B}|)^{\omega} f(|q|))$ for some function f and $\omega \geq 1$. This implies that we can obtain a set of exactly the 1-entries of $A \times B$ in time $O((|q(\mathcal{D}_{A,B})| + |\mathcal{D}_{A,B}|)^{\omega} f(|q|)) = O(m^{\omega} f(2)) = O(m^{\omega})$. \square

Lemma 5.11 directly implies the following lower bound.

Theorem 5.12. *If RPQ-EVAL can be solved in time $O((|q(\mathcal{D})| + |\mathcal{D}|) f(|q|))$ for some function f , then the SBMM-hypothesis fails.*

Surprisingly, we can obtain a more complete picture for the problem RPQ-COUNT. First, we observe that obviously all upper bounds carry over from RPQ-EVAL to RPQ-COUNT. On the other hand, a combinatorial $O((|V_{\mathcal{D}}| |\mathcal{D}|)^{1-\epsilon} f(q))$ algorithm or a general $O((|q(\mathcal{D})| + |\mathcal{D}|) f(|q|))$ algorithm for RPQ-COUNT does not seem to help for solving Boolean matrix multiplication (and therefore, the lower bounds do not carry over). Fortunately, it turns out that OV is a suitable problem to reduce to RPQ-COUNT, although by a rather different reduction compared to the one used for Theorem 5.5.

Theorem 5.13. *If RPQ-COUNT can be solved in time $O(|\mathcal{D}|^{2-\epsilon}f(|q|))$ for some function f and $\epsilon > 0$, then the OV-hypothesis fails.*

Proof. We assume that RPQ-COUNT can be solved in time $O(|\mathcal{D}|^{2-\epsilon}f(|q|))$ for some function f and $\epsilon > 0$. Let $A = \{\vec{a}_1, \vec{a}_2, \dots, \vec{a}_n\}$ and $B = \{\vec{b}_1, \vec{b}_2, \dots, \vec{b}_n\}$ be an OV-instance, i.e., for every $i \in [n]$, \vec{a}_i and \vec{b}_i are d -dimensional Boolean vectors. Now let A' be the Boolean matrix having rows $\vec{a}_1, \vec{a}_2, \dots, \vec{a}_n$ and let B' be the Boolean matrix having columns $\vec{b}_1, \vec{b}_2, \dots, \vec{b}_n$. It can be easily seen that, for every $i, j \in [n]$, $(A' \times B')[i, j] = 0$ if and only if \vec{a}_i and \vec{b}_j are orthogonal. Moreover, A' and B' can be constructed in time $O(nd)$.

Next, we construct $\mathcal{D}_{A', B'}$ as in the proof of Lemma 5.9, and we note that $V_{\mathcal{D}_{A', B'}} = O(n + d)$ and $E_{\mathcal{D}_{A', B'}} = O(nd)$; in particular, we construct $\mathcal{D}_{A', B'}$ in time $O(nd)$. Since, for every $i, j \in [n]$, $((i, 0), (j, 2)) \in q(\mathcal{D}_{A, B})$ if and only if $(A \times B)[i, j] = 1$, we know that, for every $i, j \in [n]$, $((i, 0), (j, 2)) \in q(\mathcal{D}_{A, B})$ if and only if \vec{a}_i and \vec{b}_j are not orthogonal. Furthermore, since $q(\mathcal{D}_{A, B}) \subseteq \{((i, 0), (j, 2)) \mid i, j \in [n]\}$, $|q(\mathcal{D}_{A, B})| = n^2$ if and only if there are no $\vec{a} \in A$ and $\vec{b} \in B$ that are orthogonal. Consequently, we can check whether there are no $\vec{a} \in A$ and $\vec{b} \in B$ that are orthogonal by computing $|q(\mathcal{D}_{A, B})|$, which, by assumption, can be done in time $O(|\mathcal{D}_{A, B}|^{2-\epsilon}f(|q|)) = O((nd)^{2-\epsilon}f(2)) = O(n^{2-\epsilon}\text{poly}(d))$. \square

Since Theorem 5.13 also excludes running time $O((|V_{\mathcal{D}}||\mathcal{D}|)^{1-\epsilon}f(|q|))$ for any function f and $\epsilon > 0$ (without restriction to combinatorial algorithms), it also shows that, subject to the OV-hypothesis, $O(|V_{\mathcal{D}}||\mathcal{D}|)$ is a tight bound for the data complexity of RPQ-COUNT.

6 Bounds for the Enumeration of RPQs

By using the PG-approach for enumeration, we can obtain the following upper bound.

Theorem 6.1. *Sorted RPQ-ENUM can be solved with $O(|\mathcal{D}||q|)$ delay and $O(1)$ updates.*

Proof. Let \mathcal{D} be a graph database over Σ , let q be an RPQ over Σ and let \preceq be the order on $V_{\mathcal{D}}$. We assume that $V_{\mathcal{D}} = [n]$ with $1 \preceq 2 \preceq \dots \preceq n$ (see Lemma 2.5).

Preprocessing:

1. We compute $G_{\boxtimes}(\mathcal{D}, q)$, which, according to Lemma 3.2, can be done in time $O(|\mathcal{D}||q|)$, and we interpret $G_{\boxtimes}(\mathcal{D}, q)$ as its underlying non-labelled graph.
2. We construct two arrays S and T of size $|V_{\mathcal{D}}|$ such that, for every $i \in [n]$, $S[i] = (i, p_0)$ and $T[i] = (i, p_f)$. Note that this also means that pointers to the corresponding adjacency lists are stored along with the nodes in S and T . Computing S and T can be done in time $O(|V_{\mathcal{D}}||q|)$ as follows. We move through the list that represents $V_{\boxtimes}(\mathcal{D}, q)$ and for every node (i, p_0) , we set $S[i] = (i, p_0)$, and for every node (i, p_f) , we set $T[i] = (i, p_f)$.
3. We modify $G_{\boxtimes}(\mathcal{D}, q)$ by adding a new node v_{sink} with an arc from each (i, p_f) with $i \in [n]$. This can be done in time $O(|V_{\mathcal{D}}|)$ as follows. We first add v_{sink} to the list that represents $V_{\boxtimes}(\mathcal{D}, q)$, which requires constant time. Then, for every node (i, p_f) of T , we add v_{sink} to the adjacency list for (i, p_f) . Again, this can be done in time $O(|V_{\mathcal{D}}|)$.
4. We obtain $(G_{\boxtimes}(\mathcal{D}, q))^R$, which, according to Lemma 2.2, can be done in time $O(|G_{\boxtimes}(\mathcal{D}, q)|) = O(|\mathcal{D}||q|)$. Then we initialise a Boolean array S' of size $|V_{\mathcal{D}}|$ with entries 0 everywhere. We then perform a BFS in $(G_{\boxtimes}(\mathcal{D}, q))^R$ starting in v_{sink} and, for every visited node (i, p_0) , we set $S'[i] = 1$. For every $i \in V_{\mathcal{D}}$, $S'[i] = 1$ if and only if v_{sink} and therefore some node from T can be reached from (i, p_0) . This step can be done in time $O(|\mathcal{D}||q|)$.
5. We initialise a Boolean array T' of size $|V_{\mathcal{D}}|$ with entries 0 everywhere, which can be done in time $O(|V_{\mathcal{D}}|)$.

Enumeration: In the enumeration phase, we carry out the following procedure.

- For every $i = 1, 2, \dots, n$:

- If $S'[i] = 1$, then
 - * perform a BFS in (i, p_0) and for every (j, p_f) that we visit, we set $T'[j] = 1$,
 - * for every $j = 1, 2, \dots, n$: produce (i, j) as output if $T'[j] = 1$,
 - * Set all entries of T' to 0.

Correctness: In the enumeration phase we perform a BFS from each (i, p_0) that can reach at least one node (j, p_f) (i. e., from each (i, p_0) with $S'[i] = 1$) and after termination of this BFS, we output exactly the pairs (i, j) for which (j, p_f) is visited in this BFS. This directly shows that we correctly enumerate $q(\mathcal{D})$ without repetitions. Furthermore, since we consider the start vertices for the BFSs in increasing order with respect to \preceq , and since after termination of a BFS started in (i, p_0) we output the pairs (i, j) in increasing order by the second element, the enumeration is sorted by lexicographic order.

Next, we estimate the delay of the enumeration. Each iteration of the main loop with $S'[i] = 1$ requires time $O(|G_{\boxtimes}(\mathcal{D}, q)|)$ for performing the BFS, $O(|V_{\mathcal{D}}|)$ for producing the output and $O(|V_{\mathcal{D}}|)$ for resetting T' . Thus, the total running-time is $O(|G_{\boxtimes}(\mathcal{D}, q)|) = O(|\mathcal{D}||q|)$ in the case that $S'[i] = 1$. If, on the other hand, $S'[i] = 0$, then the iteration terminates after constant time. Moreover, since $S'[i] = 1$ implies that i can reach v_{sink} and since v_{sink} can only be reached via some (j, p_f) , we produce at least one output in such an iteration. Consequently, the delay between two outputs is bounded by the total running time for one iteration of the main loop, which is $O(|\mathcal{D}||q|)$.

Updates: If \mathcal{D} is changed to \mathcal{D}' by an update, then we can again perform the whole preprocessing (with respect to \mathcal{D}') followed by the enumeration procedure. Technically, the preprocessing is done as a first step of the enumeration procedure (since our algorithmic framework does not allow to re-run the preprocessing after an update), which is possible since its running time of $O(|\mathcal{D}'||q|)$ is completely subsumed by the time available for the first delay. \square

This enumeration algorithm is easy to implement and has some nice features like linear preprocessing (in data complexity), sorted enumeration and constant updates. Unfortunately, these features come more or less for free with the disappointing delay bound. Is the PG-approach therefore the wrong tool for RPQ-enumeration? Or can we give evidence that linear delay is a barrier we cannot break? The rest of this work is devoted to this question.

Since running an algorithm for RPQ-ENUM until we get the first element yields an algorithm for RPQ-BOOLE (with preprocessing plus delay as running time), and since running such an algorithm completely solves RPQ-EVAL in time preprocessing plus $|q(\mathcal{D})|$ times delay, we can inherit several lower bounds directly from Section 5.

Theorem 6.2. *If, for some function f and $\epsilon \geq 0$, RPQ-ENUM can be solved with*

1. *delay $O(|\mathcal{D}|^{2-\epsilon} + |q|^2)$ or $O(|\mathcal{D}|^2 + |q|^{2-\epsilon})$, then the *OV-conjecture* fails.*
2. *delay $O(|V_{\mathcal{D}}|^{3-\epsilon} + |q|^{3-\epsilon})$, then the *com-BMM-hypothesis* fails.*
3. *preprocessing $O(|\mathcal{D}|f(|q|))$ and delay $O(f(|q|))$, then the *SBMM-hypothesis* fails.*
4. *preprocessing $O(|V_{\mathcal{D}}|^{3-\epsilon}f(|q|))$ and delay $O(|V_{\mathcal{D}}|^{1-\epsilon}f(|q|))$, then the *com-BMM-hypothesis* fails.*

Proof. We prove the four lower bounds of Theorem 6.2 separately.

1. Let us assume that RPQ-ENUM can be solved with delay $O(|\mathcal{D}|^{2-\epsilon} + |q|^2)$ (or $O(|\mathcal{D}|^2 + |q|^{2-\epsilon})$) for some $\epsilon > 0$. Then we can solve RPQ-BOOLE on an instance (\mathcal{D}, q) as follows. We run the assumed enumeration algorithm for RPQ-ENUM on (\mathcal{D}, q) , but stop as soon as the first pair (if any) is produced. By assumption, this requires time $O(|\mathcal{D}|^{2-\epsilon} + |q|^2)$ (or $O(|\mathcal{D}|^2 + |q|^{2-\epsilon})$). Moreover, (\mathcal{D}, q) is a positive RPQ-BOOLE instance if and only if the interrupted enumeration algorithm produces a pair. Thus, we can solve RPQ-BOOLE in time $O(|\mathcal{D}|^{2-\epsilon} + |q|^2)$ (or $O(|\mathcal{D}|^2 + |q|^{2-\epsilon})$) for some $\epsilon > 0$, which, according to Theorem 5.5 means that the *OV-hypothesis* fails.

2. Let us assume that RPQ-ENUM can be solved with delay $O(|V_{\mathcal{D}}|^{3-\epsilon} + |q|^{3-\epsilon})$ for some $\epsilon > 0$. Then we can solve RPQ-BOOLE on an instance (\mathcal{D}, q) as follows. We run the assumed enumeration algorithm for RPQ-ENUM on (\mathcal{D}, q) , but stop as soon as the first pair (if any) is produced. By assumption, this requires time $O(|V_{\mathcal{D}}|^{3-\epsilon} + |q|^{3-\epsilon})$. Moreover, (\mathcal{D}, q) is a positive RPQ-BOOLE instance if and only if the interrupted enumeration algorithm produces a pair. Thus, we can solve RPQ-BOOLE in time $O(|V_{\mathcal{D}}|^{3-\epsilon} + |q|^{3-\epsilon})$ for some $\epsilon > 0$, which, according to Theorem 5.6 means that the com-BMM-hypothesis fails.
3. We assume that RPQ-ENUM can be solved with preprocessing $O(|\mathcal{D}|f(|q|))$ and delay $O(f(|q|))$ for some function f . By completely running the corresponding enumeration algorithm, we can solve RPQ-EVAL in time $O(|\mathcal{D}|f(|q|) + O(|q(\mathcal{D})|f(|q|)) = O((|q(\mathcal{D})| + |\mathcal{D}|)f(|q|))$. According to Theorem 5.12 this means that SBMM can be solved in time $O(m)$.
4. Assume that RPQ-ENUM can be solved with preprocessing $O(|V_{\mathcal{D}}|^{3-\epsilon}f(|q|))$ and delay $O(|V_{\mathcal{D}}|^{1-\epsilon}f(|q|))$ for some function f and $\epsilon \geq 0$. By completely running the corresponding enumeration algorithm, we can solve RPQ-EVAL in time $O(|V_{\mathcal{D}}|^{3-\epsilon}f(|q|) + |q(\mathcal{D})||V_{\mathcal{D}}|^{1-\epsilon}f(|q|)) = O(|V_{\mathcal{D}}|^{3-\epsilon}f(|q|))$. According to Lemma 5.9, this means that BMM can be solved in time $O(n^{3-\epsilon})$.

□

The first two bounds only tell us that we might not be able to lower the delay $O(|\mathcal{D}||q|)$ in terms of combined complexity. While this point of view was justified for the problems discussed in Section 5, it does not say anything regarding delays of the form $O(|\mathcal{D}|^{1-\epsilon}f(|q|))$. The third bound, conditional to the SBMM-hypothesis, is much more relevant, since it suggests that the optimum of linear preprocessing and constant delay is not reachable. For combinatorial algorithms, the fourth bound at least answers our main question with $|\mathcal{D}|$ replaced by $|V_{\mathcal{D}}|$: with linear preprocessing, we cannot get below a delay of $O(|V_{\mathcal{D}}|)$.

These lower bounds can be improved significantly, if we also want to handle updates (within some reasonable time bounds).

Theorem 6.3. *If RPQ-ENUM can be solved with*

1. *arbitrary preprocessing, $O(|V_{\mathcal{D}}|^{1-\epsilon}f(|q|))$ updates and $O(|V_{\mathcal{D}}|^{1-\epsilon}f(|q|))$ delay for some function f and $\epsilon \geq 0$, then the OMv-hypothesis fails.*
2. *$O(|V_{\mathcal{D}}|^{3-\epsilon}f(|q|))$ preprocessing, $O(|V_{\mathcal{D}}|^{2-\epsilon}f(|q|))$ updates and $O(|V_{\mathcal{D}}|^{2-\epsilon}f(|q|))$ delay for some function f and $\epsilon \geq 0$, then the com-BMM-hypothesis fails.*

Proof. We prove the two parts of Theorem 6.3 separately.

1. Assume that there is an algorithm that solves RPQ-ENUM with preprocessing t_p , delay t_d and update time t_u . We can then solve OMv on some instance M and $\vec{v}^1, \vec{v}^2, \dots, \vec{v}^n$ as follows. We initialise the empty graph database (over alphabet $\{\mathbf{a}\}$) and the RPQ $q = \mathbf{aa}$, and then we perform the preprocessing in constant time. Then, by performing $O(n^2)$ updates, we transform the current graph database into the graph database $\mathcal{D}_{M, \vec{v}^1}$ over $\{\mathbf{a}\}$ with nodes $\{u_i, v_i, w \mid 1 \leq i \leq n\}$ and arcs $\{(u_i, \mathbf{a}, v_j) \mid M[i, j] = 1\} \cup \{(v_j, \mathbf{a}, w) \mid \vec{v}_1[j] = 1\}$, which requires time $O(n^2 t_u)$. Since $q(\mathcal{D}_{M, \vec{v}^1}) = \{(u_i, w) \mid (M\vec{v}^1)[i] = 1\}$, we can now compute $M\vec{v}^1$ by enumerating $q(\mathcal{D}_{M, \vec{v}^1})$ in time $O(nt_d)$. In the same way, for every i with $2 \leq i \leq n$, we can obtain $M\vec{v}^i$ by enumerating $q(\mathcal{D}_{M, \vec{v}^i})$, where every $\mathcal{D}_{M, \vec{v}^{i+1}}$ can be obtained from $\mathcal{D}_{M, \vec{v}^i}$ by n updates in time $O(nt_u)$. Consequently, the whole procedure requires time $O(n^2 t_u + nt_d + (n-1)(nt_u + nt_d))$, and the number of nodes always satisfies $|V_{\mathcal{D}}| \in O(n)$. With the assumption that $t_d = O(|V_{\mathcal{D}}|^{1-\epsilon}f(|q|))$ and $t_u = O(|V_{\mathcal{D}}|^{1-\epsilon}f(|q|))$ for some $\epsilon > 0$ and some function f , this leads to a total running-time of $O(n^2 |V_{\mathcal{D}}|^{1-\epsilon}f(|q|)) = O(n^{3-\epsilon})$.
2. Assume that there is an algorithm that solves RPQ-ENUM with preprocessing t_p , delay t_d and update time t_u . We can then solve Triangle on some instance $G = (V, E)$ with $V = \{v_1, v_2, \dots, v_n\}$ as follows (the result will then follow from the combinatorial subcubic equivalence of BMM and Triangle (see Section 4)). We construct the graph database \mathcal{D}_{v_1} over $\{\mathbf{a}\}$ with nodes $V_{\mathcal{D}_{v_1}} = \{s, t, (u, i) \mid 0 \leq i \leq 3, u \in V\}$ and arcs $E_{\mathcal{D}_{v_1}} = \{((u, i), \mathbf{a}, (u', i+1)) \mid$

$0 \leq i \leq 2, (u, u') \in E\} \cup \{(s, \mathbf{a}, (v_1, 0)), ((v_1, 3), \mathbf{a}, t)\}$. Furthermore, we define the RPQ $q = \mathbf{aaaaa}$. It can be easily seen that $q(\mathcal{D}_{v_1}) = \{(s, t)\}$ if there is a triangle that contains v_1 , and $q(\mathcal{D}_{v_1}) = \emptyset$ otherwise. In particular, this means that enumerating $q(\mathcal{D}_{v_1})$ can be done in time $O(t_d)$.

Analogously, by constructing $\mathcal{D}_{v_2}, \dots, \mathcal{D}_{v_n}$ and enumerating $q(\mathcal{D}_{v_2}), \dots, q(\mathcal{D}_{v_n})$, we can therefore check whether G contains a triangle. For every i with $2 \leq i \leq n$, we can obtain $\mathcal{D}_{v_{i+1}}$ from \mathcal{D}_{v_i} by performing a constant number of updates (we just have to remove arcs $(s, \mathbf{a}, (v_i, 0))$ and $((v_i, 3), \mathbf{a}, t)$ and instead add arcs $(s, \mathbf{a}, (v_{i+1}, 0))$ and $((v_{i+1}, 3), \mathbf{a}, t)$). Consequently, this whole procedure requires the preprocessing for a graph database with $O(|V|)$ nodes and $O(|E|)$ arcs, $|V|$ many delays and $|V|$ many updates, i. e., the total running time of this procedure is $O(t_p + |V|(t_d + t_u))$. With the assumption that $t_p = O(|V_{\mathcal{D}}|^{3-\epsilon} f(|q|))$ and $t_d = t_u = O(|V_{\mathcal{D}}|^{2-\epsilon} f(|q|))$ for some $\epsilon > 0$ and some function f , this leads to a total running-time of $O(|V|^{3-\epsilon} f(|q|) + |V||V_{\mathcal{D}}|^{2-\epsilon} f(|q|)) = O(|V|^{3-\epsilon})$.

□

The first bound rules out that we can get below $O(|V_{\mathcal{D}}|)$ for delay *and* update time, regardless of the preprocessing; the second one analogously rules out anything below $O(|\mathcal{D}|)$ (for combinatorial algorithms and linear preprocessing). While all these lower bounds suggest that improving on the linear delay may be rather difficult, they leave the following case open.

Question 6.4. *Can RPQ-ENUM be solved with $O(|\mathcal{D}|)$ preprocessing and $O(|V_{\mathcal{D}}|)$ delay in data complexity?*

What is a reasonable conjecture with respect to this question? A combinatorial algorithm that answers it in the positive does not seem to have any unlikely consequences. Indeed, it would just entail an $O(|\mathcal{D}||V_{\mathcal{D}}|)$ algorithm for RPQ-EVAL (which exactly fits to Theorem 5.7), an $O(|V||G|)$ algorithm for computing transitive closures and an $O(n^3)$ algorithm for multiplying Boolean $n \times n$ matrices, which is the state of the art for combinatorial algorithms (this is due to the obvious reductions from these problems to RPQ-ENUM). However, since Question 6.4 is about enumeration, a positive answer also means that after $O(|G|)$ preprocessing, we can enumerate the transitive closure of a graph with delay $O(|V|)$, and that after $O(n^2)$ preprocessing, all 1-entries of the Boolean matrix multiplication can be enumerated with delay $O(n)$. Are such enumeration algorithms unlikely, so that we should rather expect a negative answer to Question 6.4? In fact not, since for the simple RPQs $q = \mathbf{a}^*$ or $q = \mathbf{aa}$, which are sufficient to encode transitive closures and Boolean matrix multiplications, linear preprocessing and delay $O(|V_{\mathcal{D}}|)$ is indeed possible, as we shall obtain as byproducts of the results in the next section.

We close this section by the following remark that points out some similarity (and differences) of reductions used in Sections 5 and 6.

Remark 6.5. As already mentioned in Section 5.1, the reduction used for the combined complexity lower bounds of Theorems 5.5 (and therefore Point 1 of Theorem 6.2) is similar to the reduction from [6] used for proving conditional lower bounds of regular expression matching. Moreover, the OV-lower bound for RPQ-COUNT of Theorem 5.13 is similar to a lower bound on counting the results of certain conjunctive queries from [14], and the OMv-lower bound from Point 1 of Theorem 6.3 is similar to a lower bound on enumerating certain conjunctive queries from [14].

The quite simple observation that Boolean matrix multiplication can be expressed as querying a bipartite graph (with conjunctive queries) has also been used in [8] (see also [13, Section 6]) and is also the base for the OMv-lower bound of [14]. In the context of this paper, this connection has been used in the bounds of Theorems 5.10, 5.12 and Points 3 and 4 of 6.2.

The obvious connection between evaluating (non-acyclic) conjunctive queries and finding triangles (or larger cliques) has already been observed in [20] (see also [13, Section 6]). However, the Triangle-lower bounds of this paper are quite different, since RPQs cannot explicitly express the structure of a triangle (or a larger clique, for that matter) by using conjunction. Therefore, our respective lower bounds (Theorems 5.6, Point 2 of Theorem 6.2, and Point 2 of Theorem 6.3) need to encode this aspect in a different way. With respect to Theorems 5.6 and Point 2 of Theorem 6.2 this is done by using non-constant queries (which explains why the lower bounds are not for data complexity, in contrast to the case of conjunctive queries), and with respect to Point 2

of Theorem 6.3, which *does* work for data complexity, it is done by using updates. Moreover, our Triangle-lower bounds do not seem to extend to larger cliques like it is the case for conjunctive queries (see [13, Section 6]).

Finally, we wish to point out that although some of the reductions used in this paper are similar to reductions used in the context of conjunctive queries, due to the difference of RPQs and CQs, none of the lower bounds directly carry over. Furthermore, note that the lower bound reductions in [8, 14] have been used for obtaining dichotomies and therefore have been stated in a much more general way.

7 Enumeration with Sub-Linear Delay

We now explore three different approaches towards enumeration of $q(\mathcal{D})$ with delay strictly better than $O(|\mathcal{D}|)$ (in data complexity):

1. allowing super-linear preprocessing,
2. enumerating a representative subset of $q(\mathcal{D})$, and
3. restricting the RPQs.

7.1 Super-Linear Preprocessing

Regarding the first approach, we can improve the delay from $O(|\mathcal{D}|)$ to $O(|V_{\mathcal{D}}|)$ by increasing the linear preprocessing by a factor of $\Delta(\mathcal{D}) \log(\overline{\Delta}(\mathcal{D}))$ (in data complexity). Recall that $\overline{\Delta}(\mathcal{D})$ is the average degree of \mathcal{D} .

Theorem 7.1. *Sorted RPQ-ENUM can be solved with $O(|q|^2 \log(\overline{\Delta}(\mathcal{D})|q|) \overline{\Delta}(\mathcal{D})|\mathcal{D}|)$ preprocessing and $O(|V_{\mathcal{D}}|)$ delay.*

Proof. Let \mathcal{D} be a graph database over Σ , let q be an RPQ over Σ and let \preceq be the order on $V_{\mathcal{D}}$. We assume that $V_{\mathcal{D}} = [n]$ with $1 \preceq 2 \preceq \dots \preceq n$ (see Lemma 2.5). Recall that $\overline{\Delta}(\mathcal{D})$ denotes the average degree of \mathcal{D} , which in particular means that $O(|E_{\mathcal{D}}|) = O(|V_{\mathcal{D}}| \overline{\Delta}(\mathcal{D}))$.

Preprocessing:

1. Compute $G_{\boxtimes}(\mathcal{D}, q)$ and interpret it as its underlying non-labelled graph. We also compute the number $\overline{\Delta}(\mathcal{D})$, which can be done by moving through the adjacency lists of \mathcal{D} and counting all edges.
2. Let $H = (V_H, E_H)$ be the DAG of $G_{\boxtimes}(\mathcal{D}, q)$'s strongly connected components. Without loss of generality, we assume that connected components are numbered with $1, 2, \dots, \ell$ for some $\ell \in \mathbb{N}$, i. e., $V_h = [h]$, and we further assume that $1, 2, \dots, \ell$ is a reverse topological sorting of H . We can compute H as follows.
 - (a) By the algorithm of Tarjan [46], we first compute the strongly connected components $K_1, K_2, \dots, K_{\ell} \subseteq V_{\boxtimes}(\mathcal{D}, q)$ of $G_{\boxtimes}(\mathcal{D}, q)$. Since the finishing times of the depth-first searches of this algorithm determine a reverse topological sorting of H , we can assume that $K_1, K_2, \dots, K_{\ell}$ describes a reverse topological sorting. We represent the strongly connected components $K_1, K_2, \dots, K_{\ell}$ by an array L of size $|V_{\mathcal{D}}||q|$ addressable by the nodes of $V_{\boxtimes}(\mathcal{D}, q)$, i. e., for every $(i, p) \in V_{\boxtimes}(\mathcal{D}, q)$, we set $L[(i, p)] = j$ if and only if $(i, p) \in K_j$.
 - (b) In order to construct the actual graph H , we first initialise an empty adjacency list for every node $j \in V_H = [\ell]$ and also an additional empty list P_j . Then we iterate through the list that represents $V_{\boxtimes}(\mathcal{D}, q)$ and add every node (i, p) that we encounter to the list $P_{L[(i, p)]}$. For every $j \in [\ell]$, P_j now represents K_j as a list. Next, we initialise an empty integer array W of size ℓ . Then, for each $j \in [\ell]$, we iterate through the list P_j and for each node (i, p) that we encounter, we do the following. For every (i', p') in the adjacency list of (i, p) , if $W[L[i', p']] \neq j$, then we add $L[(i', p')]$ to the adjacency list of j and set $W[L[i', p']] = j$ (the array W makes sure that we do not add the same arc

between strongly connected components more than once, i. e., for each individual $j \in [\ell]$ considered in the above procedure, W stores for each strongly connected component $j' \in [\ell]$ whether we have already added the arc (j, j') .

3. Construct H^R .
4. Initialise for each node j in H^R an empty AVL tree A_j (or any search tree that allows the operations insert, delete, lookup, delete-min and delete-max in worst-case time that is logarithmic in its size), and store pointers to the roots of these trees and a counter for their size in an array A of size ℓ , such that finding the tree associated to a specific node can be done in constant time. For every $i \in [n]$, insert i into the tree $A_{L[(i,p_f)]}$, and at each point where the size of $A_{L[(i,p_f)]}$ exceeds $\bar{\Delta}(\mathcal{D})|q|$ during these insertions, delete the node of largest index from $A_{L[(i,p_f)]}$.
5. For $j = 1, \dots, \ell$ go through the adjacency list of node j in H^R . For each arc (j, j') in this list, insert one by one all entries from A_j into $A_{j'}$ (first lookup to avoid duplicates). Update also the size counter of the tree $A_{j'}$ accordingly, and at each point where this size exceeds $\bar{\Delta}(\mathcal{D})|q|$ during these insertions, delete the node of largest index from $A_{j'}$. This ensures that $A_{j'}$ is kept at a size at most $\bar{\Delta}(\mathcal{D})|q|$.
6. Initialise a Boolean array S of size ℓ with entries 0 everywhere. Create for each $j \in [\ell]$ a sorted array Z_j of all entries in A_j by successively extracting and deleting the smallest entry. If Z_j is not empty, set $S[j] = 1$.
7. Initialise a Boolean array T of size $|V_{\mathcal{D}}|$ with entries 0 everywhere.

Enumeration: In the enumeration phase, we carry out the following procedure.

- For every $i = 1, 2, \dots, n$:
 - If $S[L[(i, q_0)]] = 1$, then
 - * If $Z_{L[(i, q_0)]}$ has size less than $\bar{\Delta}(\mathcal{D})|q|$, output all (i, j) for all j stored in Z_i (in the sorted order).
 - * If $Z_{L[(i, q_0)]}$ has size $\bar{\Delta}(\mathcal{D})|q|$ perform a BFS in $G_{\boxtimes}(\mathcal{D}, q)$ starting at node (i, p_0) and for every visited (i', p_f) , set $T[i'] = 1$. Output during this BFS, all pairs (i, i') for all i' stored in $Z_{L[(i, q_0)]}$ (in the sorted order) to keep the desired delay, i.e., start with $k = 1$ and whenever the delay of $|V_{\mathcal{D}}|k$ has expired output $(i, Z_i[k])$ and set $k = k + 1$.
 - * for every $i' = Z_{L[(i, q_0)]}[k], \dots, n$: produce (i, i') as output if $T[i'] = 1$,
 - * Set all entries of T to 0.

Running time: We first note that Steps 1 to 3 can be done in time $O(|G_{\boxtimes}(\mathcal{D}, q)|) = O(|\mathcal{D}||q|)$: Lemma 3.2 ensures that $G_{\boxtimes}(\mathcal{D}, q)$ can be computed in $O(|\mathcal{D}||q|)$, Lemma 2.2 ensures that H^R can be computed in $O(|H|) = O(|\mathcal{D}||q|)$, and it can be easily seen that in order to construct H , we only have to move through $G_{\boxtimes}(\mathcal{D}, q)$ a constant number of times; thus, time $O(|G_{\boxtimes}(\mathcal{D}, q)|) = O(|\mathcal{D}||q|)$ is sufficient.

Since the search trees never exceed a size of $\bar{\Delta}(\mathcal{D})|q|$, each operation supported by the search trees can be carried out in time $O(\log(\bar{\Delta}(\mathcal{D})|q|))$. In Step 4, we first initialise the search trees, which can be done in $O(|V_H|) = O(|V_{\boxtimes}(\mathcal{D}, q)|)$. Then, adding i to the search tree $A_{L[(i,p_f)]}$ for every $i \in [n]$ can be done in time $O(|V_{\mathcal{D}}| \log(\bar{\Delta}(\mathcal{D})|q|))$. In Step 5, we have to perform $O(\bar{\Delta}(\mathcal{D})|q|)$ search tree operations per arc of H , i. e., we need time $O(|E_H| \bar{\Delta}(\mathcal{D})|q| \log(\bar{\Delta}(\mathcal{D})|q|)) = O(|G_{\boxtimes}(\mathcal{D}, q)| \bar{\Delta}(\mathcal{D})|q| \log(\bar{\Delta}(\mathcal{D})|q|))$.

Step 6, we have to copy, for every $j \in [\ell]$, the elements from A_j to Z_j , which requires a total of $\ell \bar{\Delta}(\mathcal{D})|q|$ search tree operations, so time $O(|G_{\boxtimes}(\mathcal{D}, q)| \bar{\Delta}(\mathcal{D})|q| \log(\bar{\Delta}(\mathcal{D})|q|))$ is sufficient (note that the time needed for initialising and filling S is clearly subsumed by this). Finally, Step 7 can obviously be carried out in time $O(|V_{\mathcal{D}}|)$. Hence, the total running time for the preprocessing is $O(|G_{\boxtimes}(\mathcal{D}, q)| \bar{\Delta}(\mathcal{D})|q| \log(\bar{\Delta}(\mathcal{D})|q|)) = O(|\mathcal{D}||q|^2 \bar{\Delta}(\mathcal{D}) \log(\bar{\Delta}(\mathcal{D})|q|))$.

Correctness: We first show that after the preprocessing terminates, for every $i \in [n]$, $Z_{L[(i,p_0)]}$ contains the first $\overline{\Delta}(\mathcal{D})|q|$ nodes i' (“first” in the sense of smallest index) such that the node (i', p_f) is reachable from (i, p_0) . In other words, $Z_{L[(i,p_0)]}$ contains the first $\overline{\Delta}(\mathcal{D})|q|$ elements of a lexicographically sorted enumeration of $\{(i, i') \mid i' \in [n], (i, i') \in q(\mathcal{D})\}$. To this end, we will show if i' is among the $\overline{\Delta}(\mathcal{D})|q|$ smallest reachable nodes from (i, p_0) , then i' is at some point inserted into $A_{L[i,p_0]}$, which means that i' is among the nodes that are copied to $Z_{L[(i,p_0)]}$ in Step 6 (recall that we only remove largest elements from some A_j if its size exceeds $\overline{\Delta}(\mathcal{D})|q|$).

If (i, p_0) and (i', p_f) are in the same strongly connected component in $G_{\boxtimes}(\mathcal{D}, q)$, then $L[(i, p_0)] = L[(i', p_f)]$ which means that in Step 4 the index i' is inserted into $A_{L[(i', p_f)]} = A_{L[(i, p_0)]}$. Otherwise there is a path in H from $L[(i, p_0)]$ to $L[(i', p_f)]$, and reversing this path yields a path from $L[(i', p_f)]$ to $L[(i, p_0)]$ in H^R . Denote by $L[(i', p_f)] = i_1, i_2, \dots, i_x = L[(i, p_0)]$ the vertices on the path from $L[(i', p_f)]$ to $L[(i, p_0)]$ in H^R . By the reverse topological sorting on H , it follows that $i_1 < i_2 < \dots < i_x$. In Step 4, i' is put in $A_{i_1} = A_{L[(i', p_f)]}$, and iteratively in Step 5, i' is copied from A_{i_h} to $A_{i_{h+1}}$ for $1 \leq h < x$, and therefore also into $A_{i_x} = A_{L[(i, p_0)]}$.

In the enumeration procedure, we go through *phases* $i = 1, 2, \dots, n$ and we will next show that in phase i , we enumerate without duplicates exactly the pairs $(i, i') \in q(\mathcal{D})$ sorted by their right element (this means that the enumeration procedure enumerates $q(\mathcal{D})$ without duplicates and lexicographically sorted).

If the size of $Z_{L[(i,p_0)]}$ is less than $\overline{\Delta}(\mathcal{D})|q|$, then, by the considerations above, we know that $Z_{L[(i,p_0)]}$ contains *all* i' with $(i, i') \in q(\mathcal{D})$. Thus, we enumerate all pairs (i, i') with $i' \in Z_{L[(i,p_0)]}$ sorted by their right elements. On the other hand, if the size of $Z_{L[(i,p_0)]}$ is at least $\overline{\Delta}(\mathcal{D})|q|$, we perform a BFS starting in (i, p_0) . After termination of this BFS, we output exactly the pairs (i, i') for which (i', p_f) is visited in this BFS and that are not already contained in $Z_{L[(i,p_0)]}$. This shows that we correctly enumerate $q(\mathcal{D})$ without duplicates. Furthermore, since we list again the pairs (i, i') by increasing index i' (the first $\overline{\Delta}(\mathcal{D})|q|$ by the sorting provided by $Z_{L[(i,p_0)]}$, and the possibly further by the sorted array T), the enumeration is again sorted by the right elements of the pairs.

It only remains to estimate the delay of the enumeration procedure, and it is sufficient to do this for each phase i separately. If $S[L[(i, q_0)]] = 0$ then phase i terminates in constant time, which also means that in case that $S[L[(i, q_0)]] = 0$ for all phases i , we need at most time in $O(|V_{\mathcal{D}}|)$ to determine that the enumeration has terminated. If the size of $Z_{L[(i,p_0)]}$ is less than avg , then we enumerate all pairs with constant delay. If, on the other hand, the size of $Z_{L[(i,p_0)]}$ is $\overline{\Delta}(\mathcal{D})|q|$, then this array contains enough solutions to pay for the BFS we run for (i, p_0) , in the sense that producing the output from $Z_{L[(i,p_0)]}$ with delay $|V_{\mathcal{D}}|$ yields time $|V_{\mathcal{D}}| \overline{\Delta}(\mathcal{D})|q| = |E_{\mathcal{D}}||q|$ which is enough to completely run the BFS. The possibly further solutions stored in T are then produced with delay at most $|V_{\mathcal{D}}|$ (for running through this list to find the non-zero entries). \square

Obviously, in the worst case we can have $\overline{\Delta}(\mathcal{D}) = \Omega(|V_{\mathcal{D}}|)$ and then the preprocessing of the algorithm of Theorem 7.1 is $\Omega(|V_{\mathcal{D}}| |\mathcal{D}|)$ (in data complexity) and therefore no improvement over just computing the complete set $q(\mathcal{D})$ in time $O(|V_{\mathcal{D}}| |\mathcal{D}|)$. However, for graph databases with low average degree, we can decrease the delay significantly from $O(|\mathcal{D}|)$ to $O(|V_{\mathcal{D}}|)$ at the cost of a slightly super-linear preprocessing time. As another remark about Theorem 7.1, we observe that the pre-computed information becomes worthless if \mathcal{D} is updated.

Finally, we discuss a minor modification of the algorithm of Theorem 7.1, which yields a slightly different result that is interesting in its own right (we state this result as a corollary after the remark).

Remark 7.2. At the cost of additional space (in $O(|V_{\mathcal{D}}|^2)$) and giving up on sorted enumeration, we can use the lazy array initialisation technique (see, e. g., the textbook [41]) in order to reduce the preprocessing time in Theorem 7.1 by a factor $\log(\overline{\Delta}(\mathcal{D})|q|)$ to $O(|q|^2 \overline{\Delta}(\mathcal{D}) |\mathcal{D}|)$. This can be done by replacing each AVL-tree A_j in step 4 by an array A_j (with lazy initialization) of size $|V_{\mathcal{D}}|$, and completely dropping step 6. Instead of the arrays Z_j , we use in the enumeration phase the unsorted lists \widehat{A}_j of the elements stored in the arrays A_j . In step 4 and 5, we simply stop filling the arrays A_j as soon as their counter exceeds $\overline{\Delta}(\mathcal{D})|q|$. This reduces the running time for these steps to $O(|V_{\mathcal{D}}|)$, and $O(|G_{\boxtimes}(\mathcal{D}, q)| \overline{\Delta}(\mathcal{D})|q|)$, respectively. Note that the lazy initialization allows us to avoid duplicates in constant time (which before was achieved by looking up the elements in the AVL-trees which required time $O(\log(\overline{\Delta}(\mathcal{D})|q|))$). In every phase i of the enumeration procedure

instead of a sorted list $Z_{L[(i,p_0)]}$ with the (at most) $\overline{\Delta}(\mathcal{D})|q|$ smallest nodes i' with $(i, i') \in q(\mathcal{D})$ we only have an unsorted list $\widehat{A}_{L[(i,p_0)]}$ of at most $\overline{\Delta}(\mathcal{D})|q|$ nodes i' (not necessarily the smallest ones) with $(i, i') \in q(\mathcal{D})$. In the case where less than $\overline{\Delta}(\mathcal{D})|q|$ elements are stored in $\widehat{A}_{L[(i,p_0)]}$, we can just produce all (i, i') for all i' stored in $\widehat{A}_{L[(i,p_0)]}$. On the other hand, if $\widehat{A}_{L[(i,p_0)]}$ stores more elements, then again we use the elements stored in $\widehat{A}_{L[(i,p_0)]}$ in order to pay for a BFS and store in the array T the new elements found by the BFS. However, since we do not produce the elements in a sorted way, we use T to also track which elements have been produced already. Whenever we output (i, i') with i' from $\widehat{A}_{L[(i,p_0)]}$ during the run of the BFS for i , we set $T[i'] = 2$ (and never overwrite a 2-entry in T). In order to list the pairs produced by the BFS after its termination, we pass once through T and only produce (i, i') if $T[i'] = 1$. As last technical detail, we change step 7 to initialise T as integer array (so we can actually have 2-entries in T).

Corollary 7.3. RPQ-ENUM can be solved with $O(|q|^2 \overline{\Delta}(\mathcal{D})|\mathcal{D}|)$ preprocessing and $O(|V_{\mathcal{D}}|)$ delay.

7.2 Enumeration of Representative Subsets

Evaluating an RPQ on a graph database \mathcal{D} aims to find for each node u all its q -successors, i.e., nodes reachable by a q -path. It is therefore a natural restriction to ask for only *at least one* (if any) such successor. Likewise, we could also ask for at least one (if any) q -predecessor for every node. More precisely, instead of the whole set $q(\mathcal{D})$, the task is to enumerate a $q(\mathcal{D})$ -approximation, which is a set $A \subseteq q(\mathcal{D})$ such that, for every $u, v \in V_{\mathcal{D}}$, if $(u, v) \in q(\mathcal{D})$, then also $(u, v'), (u', v) \in A$ for some $u', v' \in V_{\mathcal{D}}$. Such a set is representative for $q(\mathcal{D})$, since it contains for *every* node the information, whether it is involved as a source and whether it is involved as a target in some reachable pair from $q(\mathcal{D})$. The problem of enumerating any $q(\mathcal{D})$ -approximation for given \mathcal{D} and q will be denoted by APP-RPQ-ENUM.

Lemma 7.4. Given a graph database \mathcal{D} and an RPQ q , we can compute a $q(\mathcal{D})$ -approximation in time $O(|\mathcal{D}||q|)$.

Proof. Let \mathcal{D} be a graph database over Σ , let q be an RPQ over Σ and let \preceq be the order on $V_{\mathcal{D}}$. We assume that $V_{\mathcal{D}} = [n]$ with $1 \preceq 2 \preceq \dots \preceq n$ (see Lemma 2.5).

We compute $G_{\boxtimes}(\mathcal{D}, q)$, which, according to Lemma 3.2, can be done in time $O(|\mathcal{D}||q|)$, and we interpret $G_{\boxtimes}(\mathcal{D}, q)$ as its underlying non-labelled graph. We construct two arrays S and T of size $|V_{\mathcal{D}}|$ such that, for every $i \in [n]$, $S[i] = (i, p_0)$ and $T[i] = (i, p_f)$. Note that this also means that pointers to the corresponding adjacency lists are stored along with the nodes in S and T . Computing S and T can be done in time $O(|V_{\mathcal{D}}||q|)$ as follows. We move through the list that represents $V_{\boxtimes}(\mathcal{D}, q)$ and for every node (i, p_0) , we set $S[i] = (i, p_0)$, and for every node (i, p_f) , we set $T[i] = (i, p_f)$.

We modify $G_{\boxtimes}(\mathcal{D}, q)$ by adding a new node v_{source} with an arc to each (i, p_0) with $i \in [n]$, and a new node v_{sink} with an arc from each (i, p_f) with $i \in [n]$. This can be done in time $O(|V_{\mathcal{D}}|)$ as follows. We first add v_{source} and v_{sink} to the list that represents $V_{\boxtimes}(\mathcal{D}, q)$, which requires constant time. Then we add all vertices from S to the adjacency list for v_{source} , which requires time $O(|V_{\mathcal{D}}|)$. Finally, for every node (i, p_f) in T , we add v_{sink} to the adjacency list for (i, p_f) . Again, this can be done in time $O(|V_{\mathcal{D}}|)$.

Let S' and T' be arrays of size $|V_{\boxtimes}(\mathcal{D}, q)|$ the entries of which can store values from $[n] \cup \{0\}$ and can be addressed by the nodes from $V_{\boxtimes}(\mathcal{D}, q)$. We wish to fill these arrays such that they satisfy the following properties. For every $(u, p) \in V_{\boxtimes}(\mathcal{D}, q)$, if (u, p) is reachable from some node from S , then $S'[(u, p)] = i$ for some $i \in [n]$ such that (u, p) is reachable from (i, p_0) , and if (u, p) is not reachable from any node from S , then $S'[(u, p)] = 0$. Analogously, for every $(u, p) \in V_{\boxtimes}(\mathcal{D}, q)$, if (u, p) can reach some node from T , then $T'[(u, p)] = i$ for some $i \in [n]$ such that (u, p) can reach (i, p_f) , and if (u, p) cannot reach any node from T , then $T'[(u, p)] = 0$.

This means that S' and T' contain all the information we need to construct a $q(\mathcal{D})$ -approximation. More precisely, let $A = \{(i, T'[(i, p_0)]) \mid i \in [n], T'[(i, p_0)] \neq 0\} \cup \{(S'[(i, p_f)], i) \mid i \in [n], S'[(i, p_f)] \neq 0\}$. It can be easily seen that $A \subseteq q(\mathcal{D})$ and, for every $u, v \in V_{\mathcal{D}}$,

$$(u, v) \in q(\mathcal{D}) \implies (\exists v' \in V_{\mathcal{D}} : (u, v') \in A) \wedge (\exists u' \in V_{\mathcal{D}} : (u', v) \in A).$$

Thus, A is a $q(\mathcal{D})$ -approximation. Moreover, provided that we have the arrays S' and T' at our disposal, we can compute A in time $O(|V_{\mathcal{D}}|)$. Consequently, in order to conclude the proof, it remains to explain how the arrays S' and T' can be computed in time $O(|\mathcal{D}||q|)$.

We assume that S' and T' are initialised with every entry storing 0. For every $i \in [n]$, we set $S'[(i, p_0)] = i$ and $T'[(i, p_f)] = i$. This can be done in time $O(|V_{\boxtimes}(\mathcal{D}, q)|) = O(|V_{\mathcal{D}}||q|)$. Next, we perform a BFS from v_{source} and whenever we traverse an arc $((u, p), (u', p'))$, we set $S'[(u', p')] = S'[(u, p)]$. This can be done in time $O(|G_{\boxtimes}(\mathcal{D}, q)|) = O(|\mathcal{D}||q|)$. Interpreting the initialisation that ensures $S'[(i, p_0)] = i$, for every $i \in [n]$, as the base of an induction and the assignments $S'[(u', p')] = S'[(u, p)]$ for traversed arcs $((u, p), (u', p'))$ as the step of the induction, we can directly conclude by induction that after termination of this BFS, S' has the desired property. In order to compute T' , we first have to construct $(G_{\boxtimes}(\mathcal{D}, q))^R$, which, according to Lemma 2.2, can be done in time $O(|G_{\boxtimes}(\mathcal{D}, q)|) = O(|\mathcal{D}||q|)$. Then we perform a BFS from v_{sink} and whenever we traverse an arc $((u, p), (u', p'))$, we set $T'[(u', p')] = T'[(u, p)]$. Hence, computing S' and T' can be done in time $O(|G_{\boxtimes}(\mathcal{D}, q)|) = O(|\mathcal{D}||q|)$. \square

Lemma 7.4 directly implies the following result.

Theorem 7.5. *APP-RPQ-ENUM can be solved with linear preprocessing and constant delay.*

Proof. Let \mathcal{D} be a graph database over Σ and let q be an RPQ over Σ . In the preprocessing, we can compute a $q(\mathcal{D})$ -approximation A , which, according to Lemma 7.4, can be done in time $O(|\mathcal{D}||q|)$. In the enumeration, we simply enumerate the set A with constant delay. \square

Interestingly, it seems rather difficult to also support updates while keeping a low delay (the following bounds are due to the same reductions used for Theorem 6.3, simply because these reductions produce instances for which all $q(\mathcal{D})$ -approximations are equal to $q(\mathcal{D})$).

Theorem 7.6. *If APP-RPQ-ENUM can be solved with*

1. *arbitrary preprocessing, $O(|V_{\mathcal{D}}|^{1-\epsilon}f(|q|))$ updates and $O(|V_{\mathcal{D}}|^{1-\epsilon}f(|q|))$ delay, then the OMv-hypothesis fails.*
2. *$O(|V_{\mathcal{D}}|^{3-\epsilon}f(|q|))$ preprocessing, $O(|V_{\mathcal{D}}|^{2-\epsilon}f(|q|))$ updates and $O(|V_{\mathcal{D}}|^{2-\epsilon}f(|q|))$ delay for some function f and $\epsilon \geq 0$, then the com-BMM-hypothesis fails.*

Proof. We prove the two lower bounds of Theorem 7.6 separately.

1. We observe that exactly the same reduction used in the proof of Point 1 of Theorem 6.3 works as well. Since we assume the same bound on the time required for updates, we can construct the graph database $\mathcal{D}_{M, \bar{v}^1}$ in the same way as in the proof of Point 1 of Theorem 6.3. Recall that $q(\mathcal{D}_{M, \bar{v}^1}) = \{(u_i, w) \mid (M\bar{v}^1)[i] = 1\}$. However, since for every $i \in [n]$, either no node is reachable by a path labelled with 11 or only the node w is reachable by a path labelled with 11, we know that $A \subseteq q(\mathcal{D})$ is a $q(\mathcal{D}_{M, \bar{v}^1})$ -approximation if and only if $A = q(\mathcal{D}_{M, \bar{v}^1})$. Thus, we can compute $M\bar{v}^1$ by enumerating a $q(\mathcal{D}_{M, \bar{v}^1})$ -approximation in time $O(nt_d)$. Repeating this step in the same way as done in the proof of Point 1 of Theorem 6.3 leads to an algorithm solving OMv in time $O(n^{3-\epsilon})$.
2. We observe that exactly the same reduction used in the proof of Point 2 of Theorem 6.3 works as well. We construct the graph database \mathcal{D}_{v_1} over $\{\mathbf{a}\}$ and the RPQ $q = \mathbf{aaaa}$ in the same way. Recall that $q(\mathcal{D}_{v_1}) = \{(s, t)\}$ if there is a triangle that contains v_1 , and $q(\mathcal{D}_{v_1}) = \emptyset$ otherwise. This also means that any $q(\mathcal{D}_{v_1})$ -approximation A is equal to $\{(s, t)\}$ if there is a triangle that contains v_1 , and $A = \emptyset$ otherwise. Thus, the assumption that APP-RPQ-ENUM can be solved with preprocessing $O(|V_{\mathcal{D}}|^{3-\epsilon}f(|q|))$, $O(|V_{\mathcal{D}}|^{2-\epsilon}f(|q|)$ updates and $O(|V_{\mathcal{D}}|^{2-\epsilon}f(|q|))$ delay for some function f and $\epsilon \geq 0$, leads to combinatorial subcubic algorithm for Triangle in the same way as in the proof of Point 2 of Theorem 6.3. \square

7.3 Restricted RPQs

As our third approach, we show that for restricted RPQ, we can solve RPQ-ENUM with delay much smaller than $O(|\mathcal{D}|)$. We first need some definitions. For any class $Q \subseteq \text{RPQ}$, we denote by $\text{ENUM}(Q)$ the problem RPQ-ENUM where the input RPQ is from Q . Moreover, $\bigvee(Q)$ is the set of all RPQs of the form $(q_1 \vee \dots \vee q_m)$ with $q_i \in Q$ for every $i \in [m]$. An RPQ q over Σ is a *basic transitive* RPQ (over Σ) if $q = (x_1 \vee \dots \vee x_k)^*$ or $q = (x_1 \vee \dots \vee x_k)^+$, where $x_1, \dots, x_k \in \Sigma$; and q is a *short* RPQ (over Σ) if $q = (x_1 \vee \dots \vee x_k)$ or $q = (x_1 \vee \dots \vee x_k)(y_1 \vee \dots \vee y_{k'})$, where $x_1, \dots, x_k, y_1, \dots, y_{k'} \in \Sigma$. By BT-RPQ and S-RPQ, we denote the class of basic transitive RPQ and the class of short RPQ, respectively.

We show that for the class $\bigvee(\text{S-RPQ} \cup \text{BT-RPQ})$ (which, e.g., contains RPQs of the form $q = (\mathbf{ab} \vee \mathbf{c}^* \vee \mathbf{b}(\mathbf{c} \vee \mathbf{d}) \vee (\mathbf{a} \vee \mathbf{b} \vee \mathbf{d})^+)$), semi-sorted RPQ-ENUM can be solved with linear preprocessing and delay $O(\Delta(\mathcal{D}))$ in data complexity (recall *semi-sorted* from Section 2). To this end, we first prove a general upper bound for semi-sorted $\text{ENUM}(\bigvee(Q))$ in terms of the $\text{ENUM}(Q)$.

Lemma 7.7. *Let $Q \subseteq \text{RPQ}$ be such that semi-sorted $\text{ENUM}(Q)$ can be solved with preprocessing $p(|\mathcal{D}|, |q|)$ and delay $d(|\mathcal{D}|, |q|)$, then semi-sorted $\text{ENUM}(\bigvee(Q))$ can be solved with preprocessing $O((|q|p(|\mathcal{D}|, |q|) + |V_{\mathcal{D}}|))$ and delay $O(|q|d(|\mathcal{D}|, |q|))$.*

Proof. We assume that semi-sorted $\text{ENUM}(Q)$ can be solved with preprocessing $p(|\mathcal{D}|, |q|)$ and delay $d(|\mathcal{D}|, |q|)$. Let $q \in \bigvee(Q)$, i.e., $q = (q_1 \vee q_2 \vee \dots \vee q_m)$ with $q_i \in Q$ for every $i \in [m]$. By assumption, there are enumeration algorithms A_1, A_2, \dots, A_m that, for any input graph database \mathcal{D} , perform some preprocessing in time $p(|\mathcal{D}|, |q_i|)$ and then enumerate the pairs $q_i(\mathcal{D})$ ordered by their first elements with delay $d(|\mathcal{D}|, |q_i|)$. We describe now an enumeration algorithm A that, for any input graph database \mathcal{D} , performs some preprocessing in time $O(|q|p(|\mathcal{D}|, |q|))$ and then enumerates the pairs in $q(\mathcal{D})$ ordered by their first elements with delay $O(|q|d(|\mathcal{D}|, |q|))$. Note that $q(\mathcal{D}) = \bigcup_{i=1}^m q_i(\mathcal{D})$.

Let \mathcal{D} be the input graph database over Σ . We assume that $V_{\mathcal{D}} = [n]$ with $1 \preceq 2 \preceq \dots \preceq n$ (see Lemma 2.5).

Intuitive Explanation: The algorithm A uses size- m arrays L and R that store nodes from $V_{\mathcal{D}} = [n]$, and size- n Boolean arrays X and Y that represent subsets of $[n]$. We construct these in the preprocessing, in addition to performing the preprocessing procedures of the algorithms A_1, A_2, \dots, A_m . Since the enumeration algorithms A_1, A_2, \dots, A_m provide a semi-ordered enumeration, each enumeration of $q_j(\mathcal{D})$ can be seen as a table with two columns whose rows correspond to the pairs from $q_j(\mathcal{D})$, ordered by the first column. The enumeration then produces the rows of this table from top to bottom. Therefore, the enumeration proceeds in n phases, where the i^{th} phase consists in producing all rows with i in the left column (note that for some i such rows might not exist). Consequently, we can also interpret phase i of the algorithms A_j as an enumeration of only single elements from $V_{\mathcal{D}}$, since the left elements of the pairs are always i . The enumeration procedure of A also proceeds in n phases, where the i^{th} phase consists in having the enumeration algorithms A_1, A_2, \dots, A_m perform their i^{th} phases in an interleaved manner, i.e., we let each A_i produce the next element one after the other, and then we repeat this step until all A_j have finished their i^{th} phase. We maintain in array L the current phases (i.e., the current left element of the produced pairs) of the algorithms A_j (this is necessary, since only a subset may participate in phase i for a fixed i and some A_j may finish it earlier than others) and we store in array R the elements produced most recently by the algorithms A_j . Once all A_j (that are still in phase i) have produced one new element, we disregard all those elements v among them such that A has already produced (i, v) in phase i (these already produced elements will be stored in Y), and we store all other elements in X . Then A picks some element $v \in X$, produces (i, v) (and therefore “buys” another delay), marks v as already produced by adding it to Y and removes it from X . Then, again each A_j will produce the next element and this goes on until all A_j have finished phase i , and therefore we can move on to the next phase.

Due to the sets X and Y , we do not produce duplicates. However, in order to bound the delay, we have to produce at least one pair in each iteration of the main loop in each phase, and therefore we have to show that it cannot happen that X is empty when A needs to produce the next pair. We shall now define this algorithm more formally.

Preprocessing: First, A performs all preprocessing procedures of the algorithms A_1, A_2, \dots, A_m ,

which can be done in time $\sum_{i=1}^m p(|\mathcal{D}|, |q_i|) = O(|q|p(|\mathcal{D}|, |q|))$. Then we construct arrays L and R each of size m the entries of which can store values from $[n] \cup \{0, n+1\}$, and we initialise all entries with 0. Furthermore, we construct Boolean arrays X and Y of size n that initially store 0 in every entry (note that these arrays shall be used for storing subsets of $[n]$). This can clearly be done in time $O(|V_{\mathcal{D}}| + |q|)$. Consequently, the total preprocessing time is $O(|q|p(|\mathcal{D}|, |q|) + |V_{\mathcal{D}}|)$.

Enumeration: Firstly, we set $c = 0$, where c will be a counter that indicates the current phase. Then we iterate the following main loop until it terminates:

1. For every $i \in [m]$ with $L[i] = c$, we request the next element from the enumeration procedure of A_i . If such an element (u, v) is returned, we set $L[i] = u$ and $R[i] = v$, and if no element is returned (i. e., the enumeration of A_i is done), we set $L[i] = R[i] = n + 1$.
2. If $c < \min\{L[i] \mid i \in [m]\}$, then
 - (a) for every $v \in X$, remove v from X and produce (c, v) ,
 - (b) set $Y = \emptyset$,
 - (c) set $c = \min\{L[i] \mid i \in [m]\}$.
3. If $c = n + 1$, interrupt.
4. For every $i \in [m]$, if $L[i] = c$ and $R[i] \notin Y$, add $R[i]$ to X .
5. Choose some element $v \in X$, produce (c, v) as output, remove v from X and add v to Y .
6. Move to Step 1.

Correctness: We first note that in the first iteration of the main loop, we have $L[i] = c = 0$ for every $i \in [m]$, so we request the first pairs from *all* A_j and store the respective left and right elements in L and R , respectively (or store $n+1$ in both L and R , if the enumeration procedure has already terminated without producing any pair). In any subsequent iteration, this only happens with respect to those A_j that are in phase c . If condition $c < \min\{L[i] \mid i \in [m]\}$ is satisfied in Step 2, then this means that by performing Step 1, we have reached a new phase (note that also in the very first iteration, we reach a new phase, i. e., phase i for the smallest left element $i \in [n]$ in any pair produced by some A_j) and therefore we produce all elements of X as output (paired with c) and then empty the sets X and Y (which were only storing elements that are relevant for phase c that just terminated). Moreover, we have to determine the next phase, which is done by setting $c = \min\{L[i] \mid i \in [m]\}$. In the very first iteration, X is empty, so nothing happens in these steps except setting c to the current phase (which must be strictly larger than 0, but not necessarily 1). Next, Step 3 would now interrupt the whole procedure, if $c = n + 1$, which means that all enumeration procedures of the algorithms A_j have already terminated. Note that if this happens for the first time, i. e., the last still active enumeration procedures of some A_j terminate by requesting their last elements in Step 1, then we necessarily also have $c < \min\{L[i] \mid i \in [m]\}$ in Step 2, which means that we have produced all elements from X as output before the interruption is invoked in Step 3. If, on the other hand, we reach Step 4, then we collect in the set X the right elements from all new pairs produced by A_j that are in phase c , but only if these are not already stored in Y , since then they would have already been produced as right element in a pair with c as left element. Step 5 then chooses and actually produces (paired with c) one of the elements from X (and removes it from X and adds it to Y to store that it has already been produced). Finally, in Step 6, we move back to Step 1, which triggers a new iteration.

These considerations show that the procedure from above will enumerate $q(\mathcal{D})$. In particular, note that if we enter Step 5 with X being empty, then we do not produce an output at this point, but, due to Step 2, we will nevertheless completely enumerate $q(\mathcal{D})$. Since we synchronise the enumerations of the A_j with respect to their phases, the enumeration of $q(\mathcal{D})$ produces *all* pairs from $q(\mathcal{D})$ with left element i , then all pairs with left element $i' > i$ and so on; thus, the enumeration of $q(\mathcal{D})$ is semi-ordered. Furthermore, the book keeping done in sets X and Y guarantees that we do not produce duplicates. It only remains to analyse the delay of this enumeration procedure.

We first observe that Step 2b, i. e., setting $Y = \emptyset$, is problematic since it requires time $O(n)$ (we have to set $Y[i] = 0$ for every $i \in [n]$). Therefore, we implement the array Y as follows. Instead of

letting it be Boolean, we assume that it can store elements from $[n] \cup \{0\}$. The idea is that $Y[i] = 0$ means that $i \notin Y$ (just as for the Boolean case), while $Y[i] = j$ with $j \in [n]$ means $i \in Y$ in the case that we are currently in phase j , i. e., $c = j$, and $i \notin Y$ otherwise. With this interpretation, Step 2b is not necessary anymore, since setting $c = \min\{L[i] \mid i \in [m]\}$ in Step 2c has the same effect as erasing all elements from Y . Consequently, we can ignore Step 2b altogether (or rather interpret as a mere comment in the pseudo code above to indicate what is happening at Step 2b). In particular, we note that with this implementation of Y , we can still check for both X and Y whether they contain a specific element, and we can add or erase specific elements in constant time (adding i to Y in phase c just means to set $Y[i] = c$ instead of $Y[i] = 1$).

However, in Steps 2a and 5, we also have to retrieve *some* element from X . In order to do this efficiently (and not by moving through the array from left to right to find some elements, which requires time $O(n)$), we also store the elements of X as an unsorted list (which is initialised in the preprocessing). This means that we can always obtain some element of X in constant time (by just retrieving the first list element). Keeping the array and the list for X synchronised is no problem: Whenever we add some i to X (Step 4), we set $X[i] = 1$ and add x at the end of the list for X ; whenever we want to retrieve some element from X (Steps 2a and 5), we retrieve and remove the first list element, say i , and then we set $X[i] = 0$. Consequently, all operations with respect to the lists X and Y can be performed in constant time.

We estimate the running-time for each of the separate steps of an iteration. Step 1 requires time $O(\sum_{j=1}^m d(|\mathcal{D}|, |q_j|)) = O(md(|\mathcal{D}|, |q|))$. The total running time of Steps 2 to 2c is $O(m+k)$, where k is the number of pairs (c, v) produced in Step 2a and $O(m)$ is needed to compute $\min\{L[i] \mid i \in [m]\}$. Note that the pairs that are produced in Step 2 are output with constant delay in Step 2a and pay for the running time dependence on k , hence in the worst case, $k = 0$ where Step 2 requires time in $O(m)$. Step 4 requires time $O(m)$. All other steps can be carried out in constant time. This means that if in each iteration at least one pair is produced by Step 6, then the delay of the whole enumeration procedure of algorithm A is $O(md(|\mathcal{D}|, |q|))$. Obviously, if we can never reach the situation that $X = \emptyset$ in Step 5, then in each iteration at least one pair is produced. Therefore, it is sufficient to prove this property.

Let us assume that the enumeration procedure has just finished Step 3 and we are in some iteration of phase i , i. e., $c = i$. Moreover, we assume that so far, we have not encountered the situation that $X = \emptyset$ in Step 5. For every $j \in [m]$, let $a_j = R[j]$ if $L[j] = c$ where $a_j = \perp$ indicates the situation that element a_j does not exist. This means that all existing elements a_j with $j \in [m]$ are exactly those elements that have most recently been produced in phases i of the enumeration procedures from the algorithms A_j (this can have happened in Step 1 of the same iteration or, if this is the first iteration of phase i , also in applications of Step 1 in previous iterations). In particular, the existing elements a_j with $j \in [m]$ have not yet been handled in the sense of Step 4, i. e., we have not yet checked whether they have already been produced as output and, if not, have added them to X . We can also note that there must be at least one $j \in [m]$ with $a_j \neq \perp$, since otherwise $c < \min\{L[i] \mid i \in [m]\}$.

In addition to these elements a_j , for every $j \in [m]$, let $b_{j,1}, b_{j,2}, \dots, b_{j,\ell_j}$ (note that $\ell_j = 0$ is possible) be exactly the elements already produced in phase i of the enumeration procedure of A_j in some previous iterations. In other words, for every $j \in [m]$, we have requested in applications of Step 1 exactly the elements $b_{j,1}, b_{j,2}, \dots, b_{j,\ell_j}, a_j$ from phase i of the enumeration procedure of A_j (note that $a_j = \perp$ is possible, which means that a_j does not exist and therefore has not been requested). Moreover, this has happened in $\ell = \max\{\ell_j \mid j \in [m]\}$ previous (i. e., not counting the current one) iterations of the main loop of the enumeration procedure of algorithm A . In particular, this means that we have in phase i so far only produced ℓ pairs with i as left element.

Let $K = \{b_{j,p} \mid j \in [m], p \in [\ell_j]\}$ and let $M = \{a_j \mid j \in [m]\}$. Since $\ell = \max\{\ell_j \mid j \in [m]\}$, there is at least one $j' \in [m]$ with $\ell_{j'} = \ell$ and therefore $|\{b_{j',1}, b_{j',2}, \dots, b_{j',\ell_{j'}}\}| = \ell$. This is true since the elements $b_{j',1}, b_{j',2}, \dots, b_{j',\ell_{j'}}$ are part of the i^{th} phase of the enumeration of $A_{j'}$ and therefore must be distinct. Thus, $|K| \geq \ell$. Furthermore, we can choose j' such that $a_{j'} \neq \perp$. This is the case since we have $a_j = \perp$ if and only if $\ell_j < \ell$ (i. e., phase i of the enumeration of A_j has already terminated) and, as observed above, there must be at least one $j \in [m]$ with $a_j \neq \perp$.

For every $b \in K$, either (i, b) has been produced as output, or $b \in X$. Since so far we have only produced ℓ pairs as output, this directly implies that if $|K| > \ell$, then $X \neq \emptyset$, which means that we reach Step 5 with $X \neq \emptyset$. If, on the other hand, $|K| = \ell$, then $K = \{b_{j',1}, b_{j',2}, \dots, b_{j',\ell_{j'}}\}$,

which also means that $a_{j'} \notin K$ since $b_{j',1}, b_{j',2}, \dots, b_{j',\ell_{j'}}, a_{j'}$ is an enumeration of distinct elements provided by algorithm $A_{j'}$ (recall that $a_{j'} \neq \perp$, as observed above). Consequently, we also have $a_{j'} \notin Y$ and therefore $a_{j'}$ is added to X in Step 4. Hence, we reach Step 5 with $X \neq \emptyset$. \square

Now, we give upper bounds for $\text{ENUM}(\text{BT-RPQ})$ and $\text{ENUM}(\text{S-RPQ})$ separately.

Theorem 7.8. *Semi-sorted $\text{ENUM}(\text{BT-RPQ})$ can be solved with delay $O(\Delta(\mathcal{D}))$.*

Proof. Let \mathcal{D} be a graph database and let $q = (x_1 \vee x_2 \vee \dots \vee x_k)^*$ and $q' = (x_1 \vee x_2 \vee \dots \vee x_k)^+$, where $x_1, x_2, \dots, x_k \in \Sigma$. It can be easily seen that enumerating $q(\mathcal{D})$ or $q'(\mathcal{D})$ is the same as enumerating the reflexive-transitive closure $(E_{\mathcal{D}'})^*$ or the transitive closure $(E_{\mathcal{D}'})^+$, where \mathcal{D}' is obtained from \mathcal{D} by deleting all x -adjacency lists with $x \notin \{x_1, x_2, \dots, x_k\}$. In [24], it is shown for general directed graphs $G = (V, E)$ how to enumerate E^* and E^+ sorted by first coordinate (denoted by “row-wise”) with delay $O(\Delta(G))$. This approach translates to a semi-sorted enumeration and can be used on \mathcal{D} in such a way that all x -adjacency lists with $x \notin \{x_1, x_2, \dots, x_k\}$ are ignored (so without preprocessing). Thus, by using this procedure on \mathcal{D} , we can enumerate $q(\mathcal{D})$ and $q'(\mathcal{D})$ semi-sorted with delay $O(\Delta(\mathcal{D}))$. \square

Theorem 7.9. *Semi-sorted $\text{ENUM}(\text{S-RPQ})$ can be solved with preprocessing $O(|\mathcal{D}|)$ and delay $O(\Delta(\mathcal{D}))$.*

Proof. Let \mathcal{D} be a graph database over Σ . We assume that $V_{\mathcal{D}} = [n]$ with $1 \preceq 2 \preceq \dots \preceq n$ (see Lemma 2.5).

Let $q = (x_1 \vee x_2 \vee \dots \vee x_k)$ with $x_1, x_2, \dots, x_k \in \Sigma$. Then $q(\mathcal{D}) = \{(u, v) \mid (u, x_i, v) \in E_{\mathcal{D}}, 1 \leq i \leq k\}$. Obviously, we can in time $O(|\mathcal{D}|)$ compute this whole set by moving through the list for $V_{\mathcal{D}}$ and for every encountered node u and every $i \in [k]$, we access the x_i -adjacency list for u and for all v it contains, we add (u, v) to a list. This list contains exactly the elements from $q(\mathcal{D})$ and is also semi-sorted. Thus, we can compute $q(\mathcal{D})$ completely in the preprocessing and therefore enumerate it semi-sorted with constant delay.

Next, we assume that $q = (x_1 \vee x_2 \vee \dots \vee x_k)(y_1 \vee y_2 \vee \dots \vee y_{k'})$ with $x_1, \dots, x_k, y_1, \dots, y_{k'} \in \Sigma$, for which the algorithm is more difficult.

Preprocessing: First, we construct $G_{\boxtimes}(\mathcal{D}, q)$ in time $O(|\mathcal{D}||q|) = O(|\mathcal{D}|)$. Since G_q , the underlying Σ -graph of the NFA for q , can be assumed to have nodes $\{1, 2, 3\}$, where 1 is the initial and 3 the accepting state, and arcs $\{(1, x_i, 2) \mid 1 \leq i \leq k\} \cup \{(2, y_i, 3) \mid 1 \leq i \leq k'\}$, we can assume that the underlying non-labelled graph of $G_{\boxtimes}(\mathcal{D}, q)$ has the following simple structure: $V_{\boxtimes}(\mathcal{D}, q) = V_1 \cup V_2 \cup V_3$ with $V_i = \{(u, i) \mid u \in V_{\mathcal{D}}\}$ for every $i \in [3]$, and $E_{\boxtimes}(\mathcal{D}, q) = E_1 \cup E_2$ with $E_1 = \{((u, 1), (v, 2)) \mid \exists i \in [k] : (u, x_i, v) \in E_{\mathcal{D}}\}$ and $E_2 = \{((u, 2), (v, 3)) \mid \exists i \in [k'] : (u, y_i, v) \in E_{\mathcal{D}}\}$. The task of enumerating $q(\mathcal{D})$ reduces now to the task of enumerating all reachable pairs $((u, 1), (v, 3))$ or, equivalently, to enumerate $(E_{\boxtimes}(\mathcal{D}, q))^* \cap (V_1 \times V_3)$. We further assume that the sets V_1, V_2 and V_3 are stored in individual lists. This can be easily achieved during the construction of $G_{\boxtimes}(\mathcal{D}, q)$: for every $u \in V_{\mathcal{D}}$ we add $(u, 1)$ to a list that stores V_1 , then we construct lists for V_2 and V_3 analogously.

Finally, we modify $G_{\boxtimes}(\mathcal{D}, q)$ as follows. We construct a Boolean array S_{in} of size $|V_{\mathcal{D}}|$ (that can be addressed by the elements from $V_{\mathcal{D}}$) initialised with 0 in all entries. Then, we move through the list for V_1 and for every $(u, 1)$ we encounter, we move through its adjacency list and for every $(v, 2)$ we encounter in this adjacency list, we set $S_{\text{in}}[v] = 1$. This procedure can be carried out in $O(|\mathcal{D}|)$ and after its termination, we have that $S_{\text{in}}[v] = 0$ if and only if the in-degree of $(v, 2)$ is 0. Next, we construct a Boolean array S_{out} of size $|V_{\mathcal{D}}|$ (that can be addressed by the elements from $V_{\mathcal{D}}$) initialised with 0 in all entries. Then, we move through the list for V_2 and for every $(u, 2)$ we encounter that has a non-empty adjacency list, we set $S_{\text{out}}[u] = 1$. This procedure can be carried out in $O(|V_{\mathcal{D}}|)$ and after its termination, we have that $S_{\text{out}}[u] = 0$ if and only if the out-degree of $(u, 2)$ is 0. We move again through the list for V_2 and we remove every encountered node $(v, 2)$ if $S_{\text{in}}[v] = 0$ or $S_{\text{out}}[v] = 0$; moreover, we move through the list for V_1 and for every $(u, 1)$ we encounter, we move through its adjacency list and we remove every $(v, 2)$ that we encounter if $S_{\text{in}}[v] = 0$ or $S_{\text{out}}[v] = 0$. These two steps can again be carried out in time $O(|\mathcal{D}|)$. We have now removed all nodes from V_2 with in-degree or out-degree 0, and all arcs adjacent to such nodes. Next, we remove all isolated nodes from V_1 and V_3 . This can be done by first computing all nodes

from V_1 with out-degree 0 and all nodes from V_3 with in-degree 0, which can be done in the same way as we did for nodes from V_2 in time $O(|\mathcal{D}|)$.

Enumeration: Let Q be an empty queue. For every $i = 1, 2, \dots, n$, we start a BFS from $(i, 1)$. Such a BFS is carried out until all nodes from the neighbourhood of $(i, 1)$ have been visited; thus, the BFS proceeds by visiting all the neighbourhoods $N((j, 2))$ for every $(j, 2) \in N((i, 1))$. From now on, we add (i, k) to Q for every node $(k, 3)$ that we visit for the first time. Whenever $\Delta(\mathcal{D})$ steps are made by the BFS after the last output while visiting a neighbourhood $N((j, 2))$, we produce after visiting this whole neighbourhood the first element from Q as output and remove this element from Q . In order to do this efficiently, we store the elements of Q not just as a list but also maintain two lists of size $[n]$ over $\{1, \dots, n\}$ storing the elements currently in and the elements already produced and deleted from Q , respectively. When the BFS terminates, i. e., all neighbourhoods of the nodes from $N((i, 1))$ have been visited, then we produce the remaining elements in Q (and thus emptying it) and proceed with the BFS for the next node from V_1 .

Correctness: The enumeration procedure goes through n phases, where in each such phase we perform a BFS in some vertex $u \in V_1$ and produce only pairs of the form (i, j) , where $u = (i, 1)$. This means it is sufficient to show for each phase separately that we will produce all pairs of the form (i, j) (where $u = (i, 1) \in V_1$ is the BFS-start-node of this phase) without duplicates and with the desired delay. Moreover, note that the enumeration is obviously semi-sorted. Further, if Q is never empty when we request the next element, the delay is obviously in $O(\Delta(\mathcal{D}))$ since visiting one neighbourhood $N((j, 2))$ requires at most $\Delta(\mathcal{D})$ steps.

Let $u \in V_1$ and consider the phase of the enumeration procedure that performs the BFS in u . The BFS visits all nodes of the neighbourhood $N(u) = \{v_1, v_2, \dots, v_h\}$ and marks them as visited, which can be done in time $O(h) = O(\Delta(\mathcal{D}))$. Then it visits the complete neighbourhood $N(v_1)$, then the complete neighbourhood $N(v_2)$ and so on until all nodes $\bigcup_{i=1}^h N(v_i)$ are visited. For each $i \in [h]$, the BFS performs $k_i = \sum_{j=1}^i |N(v_j)|$ steps to visit all neighbourhoods $N(v_1), \dots, N(v_i)$, while producing at least $q_i = \max\{|N(v_j)| \mid 1 \leq j \leq i\}$ different elements for Q . To show that Q is never empty, it suffices to show that $\sum_{j=1}^i k_j \leq q_i \Delta(\mathcal{D})$ for all i ; note that we request an element from Q exactly after $\sum_{j=1}^{r_i} k_j$ steps for some $r_1, \dots, r_x \in [h]$ where $1 = r_1 < r_2 < \dots, r_x$ and $\sum_{j=r_i}^{r_{i+1}+1} k_j \geq \Delta(\mathcal{D})$ for each $i \in [x]$. Since the degree of u is at most $\Delta(\mathcal{D})$, we know that $h \leq \Delta(\mathcal{D})$ which means $\sum_{j=1}^i k_j \leq i \max\{|N(v_j)| \mid 1 \leq j \leq i\} \leq \Delta(\mathcal{D})q_i$. \square

Finally, by using Lemma 7.7, we can plug together Theorems 7.8 and 7.9 in order to obtain the following upper bound for $\text{ENUM}(\vee(\text{S-RPQ} \cup \text{BT-RPQ}))$.

Theorem 7.10. *Semi-sorted $\text{ENUM}(\vee(\text{S-RPQ} \cup \text{BT-RPQ}))$ can be solved with preprocessing $O(|q|^2 |\mathcal{D}|)$ and delay $O(|q|^2 \Delta(\mathcal{D}))$.*

Proof. Due to Theorems 7.8 and 7.9, both semi-sorted $\text{ENUM}(\text{S-RPQ})$ and semi-sorted $\text{ENUM}(\text{BT-RPQ})$ can be solved with preprocessing $O(|q| |\mathcal{D}|)$ and delay $O(\Delta(\mathcal{D}))$. This obviously also means that semi-sorted $\text{ENUM}(\text{S-RPQ} \cup \text{BT-RPQ})$ can be solved with preprocessing $O(|q| |\mathcal{D}|)$ and delay $O(\Delta(\mathcal{D}))$. Lemma 7.7 now implies that semi-sorted $\text{ENUM}(\vee(\text{S-RPQ} \cup \text{BT-RPQ}))$ can be solved with preprocessing $O(|q|^2 |\mathcal{D}|)$ and delay $O(|q|^2 \Delta(\mathcal{D}))$. \square

Since $q = \text{aa} \in \text{S-RPQ}$ is sufficient to express BMM as RPQ-evaluation, Theorem 7.10 implies that enumerating (the 1-entries of) Boolean matrix products can be solved with linear preprocessing and $O(n)$ delay, but, on the other hand, this also immediately implies a matching data complexity lower bound for the upper bound of Theorem 7.10.

Theorem 7.11. *If $\text{RPQ-ENUM}(\text{S-RPQ})$ can be solved with prep. $O(|V_{\mathcal{D}}|^{3-\epsilon} f(|q|))$ and delay $O(|\Delta(\mathcal{D})|^{1-\epsilon} f(|q|))$ for some function f and $\epsilon \geq 0$, then the com-BMM-hypothesis fails.*

Proof. If $\text{RPQ-ENUM}(\text{S-RPQ})$ can be solved with preprocessing $O(|V_{\mathcal{D}}|^{3-\epsilon} f(|q|))$ and with delay $O(|\Delta(\mathcal{D})|^{1-\epsilon} f(|q|))$ for some function f and $\epsilon \geq 0$, then we can also compute $q(\mathcal{D})$ in total time $O(|V_{\mathcal{D}}|^{3-\epsilon} f(|q|) + (|q(\mathcal{D})| |\Delta(\mathcal{D})|^{1-\epsilon} f(|q|)))$. For the reduction of Lemma 5.9 this implies that we can solve com-BMM in time $O(n^{3-\epsilon} + (n^2 n^{1-\epsilon})) = O(n^{3-\epsilon})$. \square

Compared to the full class of RPQs, the class $\bigvee(\text{S-RPQ} \cup \text{BT-RPQ})$ is quite restricted. However, comprehensive experimental analyses of query logs suggest that quite restricted RPQs are still practically relevant: in the corpus of more than 50 million RPQs analysed in [18, Table 4], roughly 50% of the RPQs are from BT-RPQ and another 25% are of the form $q = x_1x_2 \dots x_k$, many of which with $k \leq 2$ making them S-RPQs [35].

References

- [1] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. If the current clique algorithms are optimal, so is valiant’s parser. *SIAM J. Comput.*, 47(6):2527–2555, 2018.
- [2] Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 434–443, 2014.
- [3] Amir Abboud, Virginia Vassilevska Williams, and Huacheng Yu. Matching triangles and basing hardness on an extremely popular conjecture. *SIAM J. Comput.*, 47(3):1098–1122, 2018.
- [4] Rasmus Resen Amossen and Rasmus Pagh. Faster join-projects and sparse matrix multiplications. In *Database Theory - ICDT 2009, 12th International Conference, St. Petersburg, Russia, March 23-25, 2009, Proceedings*, pages 121–126, 2009.
- [5] Renzo Angles, Marcelo Arenas, Pablo Barceló, Aidan Hogan, Juan L. Reutter, and Domagoj Vrgoc. Foundations of modern query languages for graph databases. *ACM Comput. Surv.*, 50(5):68:1–68:40, 2017.
- [6] Arturs Backurs and Piotr Indyk. Which regular expression patterns are hard to match? In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 457–466, 2016.
- [7] Guillaume Bagan, Angela Bonifati, and Benoît Groz. A trichotomy for regular simple path queries on graphs. In *Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2013, New York, NY, USA - June 22 - 27, 2013*, pages 261–272, 2013.
- [8] Guillaume Bagan, Arnaud Durand, and Etienne Grandjean. On acyclic conjunctive queries and constant delay enumeration. In *Computer Science Logic, 21st International Workshop, CSL 2007, 16th Annual Conference of the EACSL, Lausanne, Switzerland, September 11-15, 2007, Proceedings*, pages 208–222, 2007.
- [9] Jean-François Baget, Meghyn Bienvenu, Marie-Laure Mugnier, and Michaël Thomazo. Answering conjunctive regular path queries over guarded existential rules. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 793–799, 2017.
- [10] Pablo Barceló. Querying graph databases. In *Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2013, New York, NY, USA - June 22 - 27, 2013*, pages 175–188, 2013.
- [11] Pablo Barceló, Diego Figueira, and Miguel Romero. Boundedness of conjunctive regular path queries. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, pages 104:1–104:15, 2019.
- [12] Pablo Barceló, Leonid Libkin, Anthony Widjaja Lin, and Peter T. Wood. Expressive languages for path queries over graph-structured data. *ACM Transactions on Database Systems (TODS)*, 37(4):31:1–31:46, 2012.
- [13] Christoph Berkholz, Fabian Gerhardt, and Nicole Schweikardt. Constant delay enumeration for conjunctive queries: a tutorial. *ACM SIGLOG News*, 7(1):4–33, 2020.

- [14] Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering conjunctive queries under updates. In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2017, Chicago, IL, USA, May 14-19, 2017*, pages 303–318, 2017.
- [15] Meghyn Bienvenu, Magdalena Ortiz, and Mantas Simkus. Regular path queries in lightweight description logics: Complexity and algorithms. *J. Artif. Intell. Res.*, 53:315–374, 2015.
- [16] Meghyn Bienvenu and Michaël Thomazo. On the complexity of evaluating regular path queries over linear existential rules. In *Web Reasoning and Rule Systems - 10th International Conference, RR 2016, Aberdeen, UK, September 9-11, 2016, Proceedings*, pages 1–17, 2016.
- [17] Angela Bonifati, Wim Martens, and Thomas Timm. An analytical study of large SPARQL query logs. *PVLDB*, 11(2):149–161, 2017.
- [18] Angela Bonifati, Wim Martens, and Thomas Timm. Navigating the maze of wikidata query logs. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pages 127–138, 2019.
- [19] Angela Bonifati, Wim Martens, and Thomas Timm. An analytical study of large SPARQL query logs. *VLDB J.*, 29(2-3):655–679, 2020.
- [20] Johann Brault-Baron. *De la pertinence de l'énumération : complexité en logiques propositionnelle et du premier ordre. (The relevance of the list: propositional logic and complexity of the first order)*. PhD thesis, University of Caen Normandy, France, 2013.
- [21] Karl Bringmann. Why walking the dog takes time: Frechet distance has no strongly sub-quadratic algorithms unless SETH fails. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 661–670, 2014.
- [22] Karl Bringmann. Fine-grained complexity theory (tutorial). In *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany*, pages 4:1–4:7, 2019.
- [23] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Moshe Y. Vardi. Reasoning on regular path queries. *SIGMOD Record*, 32(4):83–92, 2003.
- [24] Katrin Casel, Tobias Friedrich, Stefan Neubert, and Markus L. Schmid. Shortest distances as enumeration problem. *CoRR*, abs/2005.06827, 2020.
- [25] Isabel F. Cruz, Alberto O. Mendelzon, and Peter T. Wood. A graphical query language supporting recursion. In *Proceedings of the Association for Computing Machinery Special Interest Group on Management of Data 1987 Annual Conference, San Francisco, California, May 27-29, 1987*, pages 323–330, 1987.
- [26] Diego Figueira. Containment of UC2RPQ: the hard and easy cases. In *23rd International Conference on Database Theory, ICDT 2020, March 30-April 2, 2020, Copenhagen, Denmark*, pages 9:1–9:18, 2020.
- [27] Diego Figueira, Adwait Godbole, Shankara Narayanan Krishna, Wim Martens, Matthias Niewerth, and Tina Trautner. Containment of simple regular path queries. *CoRR*, abs/2003.04411, 2020.
- [28] Dominik D. Freydenberger and Nicole Schweikardt. Expressiveness and static analysis of extended conjunctive regular path queries. *J. Comput. Syst. Sci.*, 79(6):892–909, 2013.
- [29] François Le Gall. Powers of tensors and fast matrix multiplication. In Katsusuke Nabeshima, Kosaku Nagasaka, Franz Winkler, and Ágnes Szántó, editors, *International Symposium on Symbolic and Algebraic Computation, ISSAC '14, Kobe, Japan, July 23-25, 2014*, pages 296–303. ACM, 2014.

- [30] Grzegorz Gluch, Jerzy Marcinkowski, and Piotr Ostropolski-Nalewaja. The first order truth behind undecidability of regular path queries determinacy. In *22nd International Conference on Database Theory, ICDT 2019, March 26-28, 2019, Lisbon, Portugal*, pages 15:1–15:18, 2019.
- [31] Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 21–30, 2015.
- [32] Monika Henzinger, Andrea Lincoln, Stefan Neumann, and Virginia Vassilevska Williams. Conditional hardness for sensitivity problems. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, pages 26:1–26:31, 2017.
- [33] Leonid Libkin, Wim Martens, and Domagoj Vrgoc. Querying graphs with data. *Journal of the ACM*, 63(2):14:1–14:53, 2016.
- [34] Katja Losemann and Wim Martens. The complexity of regular expressions and property paths in SPARQL. *ACM Transactions on Database Systems (TODS)*, 38(4):24:1–24:39, 2013.
- [35] Wim Martens. Personal communication by email, October 28, 2019.
- [36] Wim Martens, Matthias Niewerth, and Tina Trautner. A trichotomy for regular trail queries. In *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, March 10-13, 2020, Montpellier, France*, pages 7:1–7:16, 2020.
- [37] Wim Martens and Tina Trautner. Evaluation and enumeration problems for regular path queries. In *21st International Conference on Database Theory, ICDT 2018, March 26-29, 2018, Vienna, Austria*, pages 19:1–19:21, 2018.
- [38] Wim Martens and Tina Trautner. Bridging theory and practice with query log analysis. *SIGMOD Rec.*, 48(1):6–13, 2019.
- [39] Wim Martens and Tina Trautner. Dichotomies for evaluating simple regular path queries. *ACM Trans. Database Syst.*, 44(4):16:1–16:46, 2019.
- [40] Alberto O. Mendelzon and Peter T. Wood. Finding regular simple paths in graph databases. *SIAM Journal on Computing (SICOMP)*, 24(6):1235–1258, 1995.
- [41] Bernard M. E. Moret and Henry D. Shapiro. *Algorithms from P to NP (Vol. 1): Design and Efficiency*. Benjamin-Cummings Publishing Co., Inc., USA, 1991.
- [42] J. Ian Munro. Efficient determination of the transitive closure of a directed graph. *Inf. Process. Lett.*, 1(2):56–58, 1971.
- [43] Juan L. Reutter, Miguel Romero, and Moshe Y. Vardi. Regular queries on graph databases. *Theory of Computing Systems (ToCS)*, 61(1):31–83, 2017.
- [44] Miguel Romero, Pablo Barceló, and Moshe Y. Vardi. The homomorphism problem for regular graph patterns. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*, pages 1–12, 2017.
- [45] Luc Segoufin. Constant delay enumeration for conjunctive queries. *SIGMOD Record*, 44(1):10–17, 2015.
- [46] Robert Endre Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1(2):146–160, 1972.
- [47] Ryan Williams. A new algorithm for optimal constraint satisfaction and its implications. In *Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, Turku, Finland, July 12-16, 2004. Proceedings*, pages 1227–1237, 2004.

- [48] Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, 2005.
- [49] Virginia Vassilevska Williams. Algorithms column: An overview of the recent progress on matrix multiplication. *SIGACT News*, 43(4):57–59, 2012.
- [50] Virginia Vassilevska Williams. Multiplying matrices faster than coppersmith-winograd. In Howard J. Karloff and Toniann Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 887–898. ACM, 2012.
- [51] Virginia Vassilevska Williams. Hardness of easy problems: Basing hardness on popular conjectures such as the strong exponential time hypothesis (invited talk). In *10th International Symposium on Parameterized and Exact Computation, IPEC 2015, September 16-18, 2015, Patras, Greece*, pages 17–29, 2015.
- [52] Virginia Vassilevska Williams. Some open problems in fine-grained complexity. *SIGACT News*, 49(4):29–35, 2018.
- [53] Virginia Vassilevska Williams and R. Ryan Williams. Subcubic equivalences between path, matrix, and triangle problems. *J. ACM*, 65(5):27:1–27:38, 2018.
- [54] Peter T. Wood. Query languages for graph databases. *SIGMOD Rec.*, 41(1):50–60, 2012.
- [55] Raphael Yuster and Uri Zwick. Fast sparse matrix multiplication. *ACM Trans. Algorithms*, 1(1):2–13, 2005.