

## DIRECT ACCESS FOR ANSWERS TO CONJUNCTIVE QUERIES WITH AGGREGATION

IDAN ELDAR <sup>a</sup>, NOFAR CARMELI <sup>b</sup>, AND BENNY KIMELFELD <sup>a</sup>

<sup>a</sup>Technion – Israel Institute of Technology, Haifa, Israel  
*e-mail address:* idel@campus.technion.ac.il, bennyk@cs.technion.ac.il

<sup>b</sup>Inria, LIRMM, University of Montpellier, CNRS, France  
*e-mail address:* nofar.carmeli@inria.fr

**ABSTRACT.** We study the fine-grained complexity of conjunctive queries with grouping and aggregation. For common aggregate functions (e.g., min, max, count, sum), such a query can be phrased as an ordinary conjunctive query over a database annotated with a suitable commutative semiring. We investigate the ability to evaluate such queries by constructing in loglinear time a data structure that provides logarithmic-time direct access to the answers ordered by a given lexicographic order. This task is nontrivial since the number of answers might be larger than loglinear in the size of the input, so the data structure needs to provide a compact representation of the space of answers. In the absence of aggregation and annotation, past research established a sufficient tractability condition on queries and orders. For queries without self-joins, this condition is not just sufficient, but also necessary (under conventional lower-bound assumptions in fine-grained complexity).

We show that all past results continue to hold for annotated databases, assuming that the annotation itself does not participate in the lexicographic order. Yet, past algorithms do not apply to the count-distinct aggregation, which has no efficient representation as a commutative semiring; for this aggregation, we establish the corresponding tractability condition. We then show how the complexity of the problem changes when we include the aggregate and annotation value in the order. We also study the impact of having all relations but one annotated by the multiplicative identity (one), as happens when we translate aggregate queries into semiring annotations, and having a semiring with an idempotent addition, such as the case of min, max, and count-distinct over a logarithmic-size domain.

### 1. INTRODUCTION

Consider a query  $Q$  that may have a large number of answers, say cubic in the number of tuples of the input database  $D$ . By answering  $Q$  via *direct access*, we avoid the materialization of the list of answers, and instead, construct a compact data structure  $S$  that supports random access: given an index  $i$ , retrieve the  $i$ th answer. Hence, direct access evaluation for a query  $Q$  consists of two algorithms, one for the structure construction (with the input  $D$ ), called *preprocessing*, and one for fast access to the answers (with the input  $S$  and  $i$ ). This task is nontrivial when  $S$  is considerably cheaper to construct than  $Q(D)$ . Similarly to past work on direct access [CTG<sup>+</sup>23], we adopt the tractability requirement of linear or quasi-linear time to construct  $S$ , and logarithmic time per access. Hence, up to a poly-logarithmic factor,

*Key words and phrases:* aggregate queries, conjunctive queries, provenance semirings, commutative semirings, annotated databases, direct access, ranking function, answer orderings, query classification.

the required construction time is what it takes to read the database (i.e., linear time), and the access time is constant. The structure  $S$  can be viewed as a compact representation of  $Q(D)$ , in the general sense of Factorized Databases [OS16], since its size is necessarily quasi-linear and it provides fast access.

Direct access solutions have been devised for Conjunctive Queries (CQs), first as a way to establish algorithms for enumerating the answers with linear preprocessing time and constant delay [BB13] (and FO queries with restrictions on the database [BDGO08]); the preprocessing phase constructs  $S$ , and the enumeration phase retrieves the answers by accessing  $S$  with increasing indices  $i$ . Later, direct access had a more crucial role within the task of enumerating the answers in a uniformly random order [CZB<sup>+</sup>22]. As a notion of query evaluation, direct access is interesting in its own right, since we can view  $S$  itself as the “result” of the query in the case where array-like access is sufficient for downstream processing (e.g., to produce a sample of answers, to return answers by pages, to answer  $q$ -quantile queries, etc.). But then  $S$  has the benefit that it is considerably smaller and faster to produce than the materialized set of answers. Indeed, recent work has studied the complexity of direct access independently (regardless of any enumeration context) [BCM25], and specifically studied which *orders* over the answers allow for such evaluation [CTG<sup>+</sup>23]. In this paper, we continue with the line of work by Carmeli et al. [CTG<sup>+</sup>23] and investigate the ability to support query evaluation via direct access for *aggregate queries*, while focusing on lexicographic orderings of answers.

For illustration, consider the following example, inspired by the FIFA World Cup. Suppose that we have a database of players of teams (countries), sponsors of teams, and goals scored in different games. Specifically, we have three relations: TEAMS(*player, country*), SPONSORS(*org, country*), and GOALS(*game, player, time*). The following CQ finds times when sponsors got exposure due to goals of supported teams:

$$Q_1(c, o, p, t) :- \text{TEAMS}(p, c), \text{SPONSORS}(o, c), \text{GOALS}(g, p, t)$$

Suppose also that we would like the answers to be ordered lexicographically by their order in the head: first by  $c$  (country), then by  $o$  (organization), then by  $p$  (player), and lastly by  $t$  (time). Note that  $o$ ,  $c$ ,  $p$  and  $t$  are *free* variables and  $g$  is an *existential* variable. Carmeli et al. [CTG<sup>+</sup>23] studied the ability to evaluate such ordered queries with direct access. In the case of  $Q_1$ , the results of Carmeli et al. show that there is an efficient direct access evaluation (since the query is free-connex and there is no “disruptive trio”).

Now, suppose that we would like to count the goals per sponsorship and player. In standard CQ notation (e.g., Cohen et al. [CNS07]), we can phrase this query as follows.

$$Q_2(c, o, p, \text{Count}()) :- \text{TEAMS}(p, c), \text{SPONSORS}(o, c), \text{GOALS}(g, p, t)$$

Here, the free variables  $c$ ,  $o$ , and  $p$  are treated as the *grouping variables*, where each combination of values defines a group of tuples over  $(c, o, p, g, t)$  and  $\text{Count}()$  simply counts the tuples in the group. Again, we would like to answer this query via direct access. This introduces two challenges. The first challenge is *aggregate construction*: when we access a tuple using  $S$ , the aggregate value should be quickly produced as well. The second challenge is *ordering by aggregation*: how can we incorporate the aggregation in the lexicographic order of the answers if so desired by the query? As an example, we may wish to order the answers first by  $c$ , then by  $\text{Count}()$ , and then by  $o$  and  $p$ ; in this case, we would phrase the head accordingly as  $Q_2(c, \text{Count}(), o, p)$ .

As previously done in the context of algorithms for aggregate queries [RS09, KCM<sup>+</sup>20], we also study a semiring alternative to the above formalism of aggregate queries. Specifically, we can adopt the well-known framework of *provenance semiring* of Green, Karvounarakis, and Tannen [GKT07] and phrase the query as an ordinary CQ with the annotation carrying the aggregate value (e.g., the number of goals in our example). To reason about random-access evaluation, we found it more elegant, general, and insightful to support CQs over annotated databases rather than SQL-like aggregate functions. For illustration, we can phrase the above aggregate query  $Q_2$  as the following CQ  $Q_3$ , but for a database that is annotated using a specific commutative semiring.

$$Q_3(c, o, p) :- \text{TEAMS}(p, c), \text{SPONSORS}(o, c), \text{GOALS}(g, p, t)$$

In a nutshell (the formal definition is in Section 2), the idea is that each tuple is annotated with an element of the semiring, the annotation of each tuple in the group is the product of the participating tuple annotations, and the annotation of the whole group is the sum of all tuple annotations in the group’s tuples. In the case of our example with  $Q_3$ , we use the counting semiring  $(\mathbb{N}, +, \cdot, 0, 1)$ , and each tuple is annotated simply with the number 1. We can use different semirings and annotations to compute different aggregate functions like sum, min, and max. Here again, we have challenges analogous to the aggregate case: *annotation construction* and *ordering by annotation*. The previous example becomes ordering by  $c$ , then by the annotation, and then by  $o$  and  $p$ . Notationally, we specify the annotation position by the symbol  $\star$  and phrase the query as  $Q_3(c, \star, o, p) :- \text{TEAMS}(p, c), \text{SPONSORS}(o, c), \text{GOALS}(g, p, t)$ . We refer to such a query as a CQ $^\star$ .

In this paper, we study queries in both formalisms—CQs enhanced with aggregate functions and ordinary CQ $^\star$ s over annotated databases. We usually devise algorithms and upper bounds on general commutative semirings (possibly with additional conditions), as positive results carry over to the aggregate formalism, and we prove cases of specific intractable queries with specific aggregate functions over ordinary (non-annotated) databases.

Our analysis is done in two parts. In Section 4, we study the case where the annotation or aggregation is *not* a part of the lexicographic order; we show that under reasonable assumptions about the complexity of the semiring operations, the known dichotomy for ordinary databases [CTG<sup>+</sup>23] continues to hold in the presence of annotation (hence, we can efficiently solve the aforementioned first challenge, namely annotation construction). We conclude the analogous tractability frontier for the common aggregate functions (count, sum, min, max, average). A notable exception is the count-distinct aggregation, which cannot be expressed efficiently as a semiring annotation; we show that the class of tractable queries for count-distinct is indeed more restricted, and we establish the precise tractability condition (dichotomy) for this aggregation.

In Section 5, we study the ability to include the annotation or aggregation in nontrivial positions within the lexicographic order (i.e., the second challenge). We establish a sufficient tractability condition for CQ $^\star$ s. Moreover, we prove that this tractability condition is also necessary (under standard assumptions in fine-grained complexity theory) for commonly studied semirings, when the CQ $^\star$  has no self-joins. We also investigate special cases of aggregate CQs where the complexity is not resolved by the dichotomy on CQ $^\star$ s. Interestingly, the translation to CQ $^\star$ s of CQs with the aforementioned aggregate functions has a special property: in all relations but one, the annotation is constantly the multiplicative identity (i.e., the “one” element). We refer to annotated databases with this property as *locally annotated*. In Section 6, we examine the implication of local annotation on the complexity

of direct access, that is, whether this property can be utilized to establish tractable queries that are hard in the absence of this property, and to what extent. We answer the first question affirmatively. Moreover, we establish a dichotomy similarly to the previous ones (i.e., the hardness side assumes that there are no self-joins) for the class of full CQ\*s. We also study the implication of local annotation in conjunction of another property, namely that the addition operation is *idempotent*. This covers the min and max aggregations, as well as count-distinct in the case of a small (logarithmic-size) domain of counted entities. We establish the corresponding dichotomy for the class of CQ\*s (with projection).

As closely related work, we note that Keppeler [Kep20] has proposed algorithms for direct access with aggregation, where he considers the class of q-hierarchical CQs (that support efficient updates of the direct-access structure [BKS17]). Yet, he does not discuss the dependence of tractability on the lexicographic order. Moreover, our positive results (e.g., Theorems 4.1 and 5.1) apply to the class of free-connex CQs, which is more general than that of the q-hierarchical CQs.

An abridged version of this manuscript appeared in a conference proceedings [ECK24]. Compared to the conference version, this manuscript has significant additions and generalizations. First, we include here the full proofs of all results. Second, we generalized two results into full dichotomy theorems: the lower bound for a specific count-distinct query [ECK24, Theorem 8] is generalized to the dichotomy of Theorem 4.5 here, and the tractability condition for annotated databases (with the computed annotation taking part of the order) [ECK24, Theorem 12] is now Corollary 5.2 of the new Theorem 5.1 here (that extends the tractability criterion and shows its necessity). Third, we reorganized the results and proofs by providing new concepts (e.g., the restriction  $Q|_V$  of a query  $Q$  to a subset  $V$  of its variables) and machinery (e.g., Lemma 3.2 and Theorem 3.4 that give general reducibility conditions that are used throughout the article and are of interest independently of this work). Fourth, we rephrased the results on locally annotated databases (Section 6) to state the precise tractability condition rather than a procedure for its detection [ECK24, Theorem 20]. In particular, this rephrasing allowed us to add Corollary 6.11 for the min and max aggregations.

The remainder of the manuscript is organized as follows. After preliminary concepts and notation in Section 2, Section 3 defines the challenge of direct access with an underlying order and recalls the state of affairs for ordinary CQs over ordinary databases. In Section 4, we analyze queries where the aggregation or annotation does not participate in the lexicographic order. We study the incorporation of the aggregation and annotation in the order in Section 5, and in Section 6 we study the implication of local annotation (with a general addition and an idempotent addition) on the complexity of direct access. We conclude in Section 7.

## 2. PRELIMINARIES

We begin with preliminary notation and terminology that we use throughout the paper.

**Databases and conjunctive queries.** A *schema*  $\mathbf{S}$  is a finite set  $\{R_1, \dots, R_k\}$  of relation symbols. Each relation symbol  $R$  is associated with an arity  $\text{ar}(R)$ , which is a natural number. We assume a countably infinite set  $\text{Vals}$  of *values* that appear in databases. A *database*  $D$  over a schema  $\mathbf{S}$  maps every relation symbol  $R$  of  $\mathbf{S}$  to a finite relation  $R^D \subseteq \text{Vals}^{\text{ar}(R)}$ . If  $(c_1, \dots, c_k)$  is a tuple of  $R^D$  (where  $k = \text{ar}(R)$ ), then we call the expression  $R(c_1, \dots, c_k)$  a *fact* of  $D$ .

A *Conjunctive Query (CQ)* over the schema  $\mathbf{S}$  has the form  $Q(\vec{x}) :- \varphi_1(\vec{x}, \vec{y}), \dots, \varphi_\ell(\vec{x}, \vec{y})$  where  $\vec{x}$  and  $\vec{y}$  are disjoint sequences of variables, and each  $\varphi_i(\vec{x}, \vec{y})$  is an *atomic query*  $R(z_1, \dots, z_k)$  such that  $R \in \mathbf{S}$  with  $\text{ar}(R) = k$  and each  $z_i$  is a variable in  $\vec{x}$  or  $\vec{y}$ . Each  $\varphi_i(\vec{x}, \vec{y})$  is an *atom* of  $Q$ , and we denote by  $\text{atoms}(Q)$  the set of atoms of  $Q$ . We call  $Q(\vec{x})$  the *head* of the query and  $\varphi_1(\vec{x}, \vec{y}), \dots, \varphi_\ell(\vec{x}, \vec{y})$  the *body* of the query. The variables of  $\vec{x}$  are the *free* variables of  $Q$ , and those of  $\vec{y}$  are the *existential* variables of  $Q$ , and every variable occurs at least once in the body. We use  $\text{vars}(Q)$  and  $\text{free}(Q)$  to denote the set of all variables and all free variables of  $Q$ , respectively. If  $\varphi \in \text{atoms}(Q)$ , then  $\text{vars}(\varphi)$  is the set of variables in  $\varphi$ . We say that  $Q$  is *full* if it has no existential variables, that is  $\text{vars}(Q) = \text{free}(Q)$ .

We refer to a database  $D$  over the schema  $\mathbf{S}$  of the CQ  $Q$  as a *database over  $Q$* . A *homomorphism* from a CQ  $Q$  to a database  $D$  over  $Q$  is a mapping  $h$  from the variables of  $Q$  into values of  $D$  such that for each atom  $R(z_1, \dots, z_k)$  of  $Q$  it holds that  $R(h(z_1), \dots, h(z_k))$  is a fact of  $D$ . We denote by  $\text{Hom}(Q, D)$  the set of all homomorphisms from  $Q$  to  $D$ . If  $h \in \text{Hom}(Q, D)$  then we denote by  $h(\vec{x})$  the tuple obtained by replacing every variable  $x$  with the value  $h(x)$ , and we denote by  $h(\varphi_i(\vec{x}, \vec{y}))$  the fact that is obtained from the atom  $\varphi_i(\vec{x}, \vec{y})$  by replacing every variable  $z$  with the value  $h(z)$ . An *answer* to  $Q$  over  $D$  is a tuple of the form  $h(\vec{x})$  where  $h \in \text{Hom}(Q, D)$ . The *result* of  $Q$  over  $D$ , denoted  $Q(D)$ , is  $Q(D) := \{h(\vec{x}) \mid h \in \text{Hom}(Q, D)\}$ .

In our proofs, we will use the following definition of when two facts  $f$  and  $f'$  agree in the context of two queries. Intuitively, facts agree if they assign the same variables with the same values. Let  $Q$  and  $Q'$  be CQs over the schemas  $\mathbf{S}$  and  $\mathbf{S}'$ , respectively. Let  $\varphi$  and  $\varphi'$  be atoms of  $Q$  and  $Q'$  over the relation symbols  $R$  and  $R'$ , respectively. Let  $f$  and  $f'$  be facts over  $R$  and  $R'$ , respectively. We say that  $f$  and  $f'$  *agree* (w.r.t.  $\varphi$  and  $\varphi'$ ) if there exists a homomorphism  $h : \text{vars}(\varphi) \cup \text{vars}(\varphi') \rightarrow \text{Vals}$  such that  $f$  and  $f'$  are obtained from  $\varphi$  and  $\varphi'$ , respectively, by replacing each variable  $x$  with the value  $h(x)$ .

Consider a CQ  $Q(\vec{x}) :- \varphi_1(\vec{x}, \vec{y}), \dots, \varphi_\ell(\vec{x}, \vec{y})$ . We may refer to  $Q$  as  $Q(\vec{x})$  to specify the sequence of free variables in the head. In this work, the *order* of the sequence  $\vec{x}$  has a crucial role, since it determines the desired order of answers. Specifically, we will assume that the desired order of answers is *lexicographic* in the left-to-right order of  $\vec{x}$ . For example, the CQ  $Q(x_1, x_2) :- R(x_1, x_2), S(x_2, y)$  differs from the CQ  $Q'(x_2, x_1) :- R(x_1, x_2), S(x_2, y)$  not only in the order of values within each answer tuple but also in the order over the answers. For  $Q(x_1, x_2)$  we order the answers first by  $x_1$  and then by  $x_2$ , and for  $Q'(x_2, x_1)$  we order first by  $x_2$  and then by  $x_1$ .

As usual, we associate a CQ  $Q$  with the hypergraph  $H(Q) = (V_Q, E_Q)$  where  $V_Q = \text{vars}(Q)$  and  $E_Q = \{\text{vars}(\varphi) \mid \varphi \in \text{atoms}(Q)\}$ . We say that  $Q$  is *acyclic* if  $H(Q)$  is an acyclic ( $\alpha$ -acyclic) hypergraph. Recall that a hypergraph  $H = (V, E)$  is acyclic if there is a tree  $T = (E, E_T)$ , called a *join tree* of  $H$ , with the running intersection property: for each vertex  $v \in V$ , the set of hyperedges that contain  $v$  induces a connected subtree of  $T$ . If  $H$  is acyclic and  $S \subseteq V$ , then we say that  $H$  is *S-connex* if  $H$  remains acyclic even if we add  $S$  to the set of hyperedges [BB13]. An acyclic CQ  $Q$  is *free-connex* if  $H(Q)$  is acyclic and  $\text{free}(Q)$ -connex.

A hypergraph  $H' = (V, E')$  is an *inclusive extension* of a hypergraph  $H = (V, E)$  if  $E \subseteq E'$  and for every edge  $e' \in E'$  there is an edge  $e \in E$  such that  $e' \subseteq e$ . It is known that  $H$  is acyclic  $S$ -connex if and only if  $H$  has an inclusive extension with a join tree  $T$  such that  $S$  is precisely the set of all variables contained in the vertices of some subtree of  $T$  [BDG07]. We call such a tree *ext-S-connex tree*. When  $S$  is the set of free variables of the CQ, and the CQ is clear from the context, we call such a tree *ext-free-connex*.

Let  $Q$  be a CQ. Two variables of  $Q$  are *neighbors* (in  $Q$ ) if they occur jointly in at least one of the atoms. If  $\vec{x}$  is a set of variables in  $\text{vars}(Q)$ , then we denote by  $N_Q(\vec{x})$  the *neighborhood* of  $\vec{x}$  that consists of every variable in  $\vec{x}$  and every neighbor of every variable in  $\vec{x}$ . If  $Q$  is clear from the context, then we may omit it and write just  $N(\vec{x})$ .

The notion of a *disruptive trio* has been introduced previously in the context of direct access to the answers of CQs [CTG<sup>+</sup>23]. A disruptive trio of a CQ  $Q(\vec{x})$  is a set of three distinct free variables  $x_1$ ,  $x_2$ , and  $x_3$  such that  $x_1$  and  $x_2$  neighbor  $x_3$  but not each other, and  $x_3$  succeeds both  $x_1$  and  $x_2$  in  $\vec{x}$ .

Aggregate queries. By an *aggregate function* we refer to a function that takes as input a bag of tuples over  $\mathbf{Vals}$  and returns a single value in  $\mathbf{Vals}$ . We adopt the notation of Cohen et al. [CNS07], as follows. An *aggregate query* here is an expression of the form

$$Q(\vec{x}, \alpha(\vec{w}), \vec{z}) := \varphi_1(\vec{x}, \vec{y}, \vec{z}), \dots, \varphi_\ell(\vec{x}, \vec{y}, \vec{z})$$

such that  $Q'(\vec{x}, \vec{z}) := \varphi_1(\vec{x}, \vec{y}, \vec{z}), \dots, \varphi_\ell(\vec{x}, \vec{y}, \vec{z})$  is a CQ,  $\alpha$  an aggregate function, and  $\vec{w}$  a sequence of variables from  $\vec{y}$ . An example is  $Q(x_1, x_2, \text{Sum}(y_2), z) := R(x_1, x_2, y_1), S(y_1, y_2, z)$ . We refer to such a query as an *Aggregate CQ* or *AggCQ* for short. Given a database  $D$  over  $Q'$ , the result  $Q(D)$  is defined by  $Q(D) := \{(\vec{a}, \alpha(B(\vec{a}, \vec{b})), \vec{b}) \mid (\vec{a}, \vec{b}) \in Q'(D)\}$  where  $B(\vec{a}, \vec{b})$  is the bag that is obtained by collecting the tuples  $h(\vec{w})$  from every  $h \in \text{Hom}(Q', D)$  with  $h(\vec{x}) = \vec{a}$  and  $h(\vec{z}) = \vec{b}$ . Note that our database and query model use set semantics, and we use bag semantics only to define the aggregate functions (in order to capture important functions such as count and sum).

We say that  $Q$  is *acyclic* if  $Q'$  is acyclic. Similarly,  $Q$  is *free-connex* if  $Q'$  is free-connex. A *disruptive trio* of  $Q$  is a disruptive trio of  $Q'$ ; in other words, the definition of a disruptive trio remains unchanged when introducing aggregates, while we consider only the free variables and not the aggregate function.

**Remark 2.1.** We remark on two aspects in our definition of AggCQs. First, the reason for using both  $\vec{x}$  and  $\vec{z}$  as sequences of free variables is to determine a position for the aggregate value  $\alpha(\vec{w})$  and, consequently, define its position in the lexicographic order over the answers. Second, the reader should note that, in our notation, an AggCQ has a single aggregate function. While this is important for some of our results, other results can be easily extended to multiple aggregate functions  $\alpha(\vec{w}_1), \dots, \alpha(\vec{w}_k)$ . We will mention this extension when relevant.  $\square$

In this work, we restrict the discussion to the common aggregate functions **Count**, **CountD** (count distinct), **Sum**, **Avg** (average), **Min**, and **Max**. All aggregate functions take a single column as input (i.e.,  $\vec{w}$  is of length one) except for **Count** that counts the tuples in the group and takes no argument. For instance, the query  $Q_2$  in the Introduction uses **Count()** and it could also use **CountD**( $g$ ) for counting the distinct games with scored goals.

Commutative semirings. A *commutative monoid* is an algebraic structure  $(\mathbb{K}, \cdot)$  over a domain  $\mathbb{K}$ , with a binary operation  $\cdot$  that satisfies *associativity*:  $(a \cdot b) \cdot c = a \cdot (b \cdot c)$  for any  $a, b, c \in \mathbb{K}$ , *commutativity*:  $a \cdot b = b \cdot a$  for any  $a, b \in \mathbb{K}$ , and *identity element*: there exists an element  $\emptyset \in \mathbb{K}$  such that  $a \cdot \emptyset = a$  for any  $a \in \mathbb{K}$ . A *commutative semiring* is an algebraic structure  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  over a domain  $\mathbb{K}$ , with two binary operations  $\oplus$  and  $\otimes$  and two distinguished elements  $\bar{0}$  and  $\bar{1}$  in  $\mathbb{K}$  that satisfy the following conditions: (a)  $(\mathbb{K}, \oplus)$  is a commutative monoid with the identity element  $\bar{0}$ ; (b)  $(\mathbb{K}, \otimes)$  is a commutative

monoid with the identity element  $\bar{1}$ ; (c)  $a \otimes (b \oplus c) = (a \otimes b) \oplus (a \otimes c)$  for all  $a, b, c \in \mathbb{K}$ ; and (d)  $a \otimes \bar{0} = \bar{0}$  for all  $a \in \mathbb{K}$ .

We refer to  $\oplus$  as the *addition* operation,  $\otimes$  as the *multiplication* operation,  $\bar{0}$  as the *additive identity* and  $\bar{1}$  as the *multiplicative identity*. We give examples of commutative semirings at the end of this section.

Annotated databases and query answers. Let  $\mathbf{S}$  be a schema and  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  a commutative semiring. A  $\mathbb{K}$ -*database* (over  $\mathbf{S}$ ) is a pair  $(D, \tau)$  where  $D$  is a database over  $\mathbf{S}$  and  $\tau : D \rightarrow \mathbb{K}$  is function that maps every fact  $f$  of  $D$  to an element  $\tau(f)$  of  $\mathbb{K}$ , called the *annotation* of  $f$ .

The annotation of a database propagates to the query answers by associating a semiring operation with each algebraic operation [GKT07]. In the case of CQs, the relevant operations are *joins* and *projection*. For join, the annotation of the result is the product of the annotation of the operands. For projection, the annotation is the sum of the annotations of the tuples that give rise to the answer. In our terminology, we have the following.

Let  $Q(\vec{x}) :- \varphi_1(\vec{x}, \vec{y}), \dots, \varphi_\ell(\vec{x}, \vec{y})$  be a CQ and  $(D, \tau)$  an annotated database. For a homomorphism  $h$  from  $Q$  to  $D$ , we denote by  $\otimes h$  the product of the annotations of the facts in the range of  $h$ , that is  $\otimes h := \tau(h(\varphi_1(\vec{x}, \vec{y}))) \otimes \dots \otimes \tau(h(\varphi_\ell(\vec{x}, \vec{y})))$ . An *answer* to  $Q$  over  $(D, \tau)$  is a pair  $(\vec{c}, a)$  such that  $\vec{c} \in Q(D)$  and

$$a = \oplus \{ \otimes h \mid h \in \text{Hom}(Q, D) \wedge h(\vec{x}) = \vec{c} \}$$

where, for  $A = \{a_1, \dots, a_n\} \subseteq \mathbb{K}$ , we define  $\oplus A = a_1 \oplus \dots \oplus a_n$ . As before, the *result* of  $Q$  over  $(D, \tau)$ , denoted  $Q(D, \tau)$ , is the set of answers  $(\vec{c}, a)$  to  $Q$  over  $(D, \tau)$ . We will make use of the fact that, over commutative semirings, projections and joins are commutative [GKT07].

In this work, we study the ability to incorporate the annotation in the order over the answers. More precisely, we will investigate the complexity of involving the annotation in the lexicographic order over the answers, as if it were another value in the tuple. So, when we consider a CQ  $Q(\vec{x})$ , we need to specify where the annotation goes inside  $\vec{x}$ . Similarly to the AggCQ notation, we do so by replacing  $\vec{x}$  with a sequence  $(\vec{x}, \star, \vec{z})$  where  $\star$  represents the annotation value. We refer to a CQ of this form as a CQ $^\star$ . An example of a CQ $^\star$  is  $Q(x_1, x_2, \star, z) :- R(x_1, x_2, y_1), S(y_1, y_2, z)$  where the lexicographic order is by  $x_1$ , then by  $x_2$ , then by the annotation, and then by  $z$ .

**Remark 2.2.** The literature on the complexity of query evaluation over annotated databases sometimes considers answers  $(\vec{c}, a)$  as valid only if  $a \neq \bar{0}$ . For instance, this convention is adopted in the recent work of Muñoz, Riveros, and Vansummeren [MRV26] on conjunctive queries for linear algebra over matrices over a semiring. In contrast, our semantics allows answers with an annotation of zero. This variation is significant in our context, since the number of answers with a zero annotation can be large. The techniques developed in this manuscript do not readily extend to supporting this restriction, and we leave this direction for future work.  $\square$

Let  $Q$  be a CQ $^\star$ , and let  $Q'$  be the CQ obtained from  $Q$  by removing  $\star$  from the head. As in the case of AggCQs,  $Q$  is *acyclic* if  $Q'$  is acyclic,  $Q$  is *free-connex* if  $Q'$  is free-connex, and a *disruptive trio* of  $Q$  is a disruptive trio of  $Q'$ .

Aggregate functions can often be captured by annotations of answers in annotated databases, where each aggregate function might require a different commutative semiring:

- **Sum:** the numeric semiring  $(\mathbb{Q}, +, \cdot, 0, 1)$ .
- **Count:** the counting semiring  $(\mathbb{N}, +, \cdot, 0, 1)$ .

TEAMS		GOALS			REPLAYS			TEAM			GOALS				REPLAYS		
$p$	$c$	$g$	$p$	$t$	$g$	$t$		$p$	$c$	$\tau_+$	$g$	$p$	$t$	$\tau_+$	$g$	$t$	$\tau_+$
1	5	1	1	31	1	1	$\Rightarrow$	1	5	1	1	1	31	1	1	1	1
2	5	1	3	50	1	31		2	5	1	1	3	50	1	1	31	31
3	6	1	3	75	1	50		3	6	1	1	3	75	1	1	50	50
4	7	2	4	90	2	5		4	7	1	2	4	90	1	2	5	5
5	8	2	4	9	1	90		5	8	1	2	4	9	1	1	90	90

Figure 1: An example of a  $\mathbb{Q}$ -database over the numerical semiring constructed to evaluate the AggCQ  $Q(c, \text{Sum}(t)) :- \text{TEAMS}(p, c), \text{GOALS}(g, p, t), \text{REPLAYS}(g, t)$ .

- **Max**: the max tropical semiring  $(\mathbb{Q} \cup \{-\infty\}, \max, +, -\infty, 0)$ .
- **Min**: the min tropical semiring  $(\mathbb{Q} \cup \{\infty\}, \min, +, \infty, 0)$ .

The translation is straightforward (and known, e.g., [RS09, KCM<sup>+</sup>20]), as we illustrate in Figure 1: the aggregated value becomes the annotation on one of the relations, the annotation outside of this relation is the multiplicative identity (as we later term “locally annotated”), and the addition operation captures the aggregate function. Note that in the case of the numeric and min/max tropical semirings, we are using the domain  $\mathbb{Q}$  of rational numbers rather than all real numbers to avoid issues of numerical presentation in the computational model.

In general, the translation is directly applicable whenever the aggregate function is defined as a commutative monoid over the elements of a column. It may also be applicable indirectly. For example, **Avg** can be computed using **Sum** and **Count**. **CountD** (count distinct) cannot be captured by a semiring, as the result of  $\oplus$  cannot be computed from two intermediary annotations in the domain. We can, however, capture a semantically similar concept with the set semiring  $(\mathcal{P}(\mathcal{X}), \cup, \cap, \emptyset, \mathcal{X})$  by annotating each fact with the actual set of distinct elements. However, in such cases, we will need our complexity analysis to be aware of the cost of the operations.

### 3. THE DIRECT-ACCESS PROBLEM

In this paper, we study CQs with lexicographic orders over the answers. As said earlier, the lexicographic order for the CQ  $Q(\vec{x})$  is left to right according to  $\vec{x}$ . We will also investigate lexicographic orders that involve the annotation or aggregation when the query is a CQ<sup>\*</sup>  $Q(\vec{x}, \star, \vec{z})$  or an AggCQ  $Q(\vec{x}, \alpha(\vec{w}), \vec{z})$ , respectively. We refer uniformly to the annotation of an answer (over an annotated database) and to the aggregate value of the answer’s group (over an ordinary database) as the *computed value*.

Let  $Q$  be a CQ, CQ<sup>\*</sup>, or an AggCQ. A *direct access* solution for  $Q$  consists of two algorithms: one for *preprocessing* and one for *access*.

- The preprocessing algorithm takes as input a database  $D$  over  $Q$  and constructs a data structure  $S_D$ .
- The access algorithm takes as input  $S_D$  and an index  $i$ , and returns the  $i$ th answer of  $Q(D)$  in the lexicographic order. Note that this answer includes the computed value, when it exists. If  $i > |Q(D)|$  then the algorithm should return *null*.

To define the complexity requirements of *efficient* direct access, we first describe the complexity model that we adopt. We use *data complexity* as a yardstick of tractability.

Hence, complexity is measured in terms of the size of the database, while the size of the query is fixed (and every query is a separate computational problem). Assuming the input is of size  $n$ , we use the RAM model of computation with  $O(\log n)$ -bit words and uniform-cost operations. Notably, this model allows us to assume perfect hash tables can be constructed in linear time, and they provide access in constant time [GJ22].

Regarding the complexity of the semiring operations, the RAM model allows us to assume that the numeric, counting, min tropical, and max tropical semirings use constant space for representing values and constant time for the operations  $\oplus$  and  $\otimes$ . In fact, it suffices for our results to assume that the operations take logarithmic time, and later we will make use of this relaxed assumption (within a special case of CountD). We refer to a semiring with this property as a *logarithmic-time (commutative) semiring*.

Let  $T_p$  and  $T_a$  be numeric functions. A direct-access algorithm is said to be in  $\langle T_p, T_a \rangle$  if, for every input database  $D$ , the preprocessing phase takes  $O(T_p(|D|))$  time and each access takes  $O(T_a(|D|))$  time. For example,  $\langle \text{loglinear}, \log \rangle$  states that preprocessing constructs in  $O(|D| \log |D|)$  time a data structure that provides  $O(\log |D|)$ -time access. In this work, a query  $Q$  has *efficient* direct access (and  $Q$  is deemed tractable) if it has a direct access algorithm in  $\langle \text{loglinear}, \log \rangle$ .

Carmeli et al. [CTG<sup>+</sup>23] established a dichotomy in the tractability of the CQs and lexicographic orders. This dichotomy relies on the following hypotheses.

- **SparseBMM**: two binary matrices  $A^{n \times n}$  and  $B^{n \times n}$  represented by lists of their non-zero entries cannot be multiplied in  $O(m \text{ polylog}(m))$  time where  $m$  is the total number of non-zero entries in  $A$ ,  $B$  and  $A \times B$ .
- **HYPERCLIQUE**: for all  $k \geq 2$ , there is no  $O(m \text{ polylog}(m))$ -time algorithm that, given a hypergraph with  $m$  hyperedges, determines whether there exists a set of  $k + 1$  vertices such that every subset of  $k$  vertices among them forms a hyperedge.

**Theorem 3.1** [CTG<sup>+</sup>23]. *Let  $Q$  be a CQ.*

- (1) *If  $Q$  is free-connex with no disruptive trio, then direct access for  $Q$  is in  $\langle \text{loglinear}, \log \rangle$ .*
- (2) *Otherwise, if  $Q$  is also self-join-free, then direct access for  $Q$  is not in  $\langle \text{loglinear}, \log \rangle$ , assuming the HYPERCLIQUE hypothesis (in case  $Q$  is cyclic) and the SparseBMM hypothesis (in case  $Q$  is acyclic).*

The following lemma claims that introducing new variables into a hard query cannot make it easier (as long as the connections between the original variables remain unchanged).

To formalize this claim, we use the following definitions. Given a self-join-free query  $Q$  and a subset  $V$  of its variables, we denote by  $Q_{|V}$  the query obtained by removing all variables that are not in  $V$  from  $Q$  (thus, the arity of the atoms or head may decrease). We say that two queries  $Q_1$  and  $Q_2$  are *structurally equivalent* if there exists a query  $Q'_1$  that is isomorphic to  $Q_1$  (i.e.,  $Q'_1$  is obtained by renaming the variables of  $Q_1$ ) such that: (1) the variables of every atom of  $Q'_1$  are contained in an atom of  $Q_2$  and vice versa; and (2)  $Q'_1$  and  $Q_2$  have the same head.

**Lemma 3.2.** *Let  $Q$  and  $Q'$  be queries such that they are both CQs, both CQ\*s over a logarithmic-time semiring, or both AggCQs, and  $Q$  has no self-joins. If there exists  $V \subseteq \text{vars}(Q)$  such that  $Q_{|V}$  and  $Q'$  are structurally equivalent, and direct access for  $Q$  is in  $\langle T_p, T_a \rangle$ , then direct access for  $Q'$  is in  $\langle \text{loglinear} + T_p, T_a \rangle$ .*

As an example for the use of Lemma 3.2, it can reprove the fact that, given that  $Q_{\text{trio}}(x_1, x_2, x_3) :- R_1(x_1, x_3), R_2(x_2, x_3)$  does not admit efficient direct access (assuming the

SparseBMM hypothesis), neither does any other CQ with a disruptive trio. We will later use Lemma 3.2 similarly with different hard queries instead of  $Q_{\text{trio}}$ .

*Proof of Lemma 3.2.* For the simplicity of presentation, we assume that the variables of  $Q$  are renamed such that the variables of every atom of  $Q_{|\text{vars}(Q')}$  are contained in an atom of  $Q'$  and vice versa, and  $Q_{|\text{vars}(Q')}$  and  $Q'$  have the same head.

We show an *order-preserving exact reduction* from  $Q'$  to  $Q$ . That is, given a database  $D'$  for  $Q'$  we show how to construct a database  $D$  for  $Q$  in loglinear time such that there is an order-preserving bijection from  $Q(D)$  to  $Q'(D')$  and, moreover, the translation of a tuple from  $Q(D)$  to its corresponding tuple from  $Q'(D')$  can be done in constant time. This proves that direct access for  $Q$  implies direct access for  $Q'$  with the same time guarantees after loglinear preprocessing.

For every atom  $R(\vec{u})$  of  $Q$ , take an atom  $R'(\vec{v})$  of  $Q'$  that contains  $\vec{u}$  restricted to  $\text{vars}(Q')$ . Construct a relation for  $R$  from  $R'$  according to these atoms while padding facts with the value  $c$  where a value is missing. That is, treat every fact  $R'(\vec{a})$  as a function  $f$  mapping  $v_i$  to  $a_i$ , extend this function to map all other variables to  $c$ , and add the fact  $R(f(\vec{u}))$  to the construction.

Note that, for an atom  $R(\vec{u})$  of  $Q$ , multiple atoms  $R'(\vec{v})$  of  $Q'$  may contain  $\vec{u}$  restricted to  $\text{vars}(Q')$ . Since we select only one  $R'(\vec{v})$ , there may be some atoms of  $Q'$  that do not participate in the above process. For every atom  $R'(\vec{v})$  of  $Q'$  that was not used in the process, take an atom  $R(\vec{u})$  of  $Q$  that contains its variables, and filter its relation in  $D$  according to  $R'$ . That is, for every fact  $R(\vec{a})$  in our constructed instance, treat it as a function  $f$  mapping  $v_i$  to  $a_i$ , and check whether there is a corresponding fact  $R'(f(\vec{v}))$  in the input; if none exists, then remove the fact  $R(\vec{a})$  from the construction.

In case the queries are CQ\*s, the database  $D'$  is accompanied by the annotations  $\tau'$ , and we construct the annotations  $\tau$  for  $D$  as follows:

- (1) We first annotate all facts in the constructed instance with  $\bar{1}$ .
- (2) Then, for every atom  $R'(\vec{v})$  in  $Q'$ , we take an atom  $R(\vec{u})$  in  $Q$  containing its variables, and multiply the annotation of every fact in  $R$  with the annotation of the corresponding fact in  $R'$  (if a corresponding fact exists; if none exist, then this fact of  $R'$  would not be used in an answer, and the annotation we choose for this fact does not matter).

Note that the construction uses a linear number of operations. As the semiring multiplication takes logarithmic time, and the other operations take constant time, the construction takes  $O(|D'| \log |D'|)$  time in total.

We claim that there is a bijection between homomorphisms from  $Q'$  to  $D'$  and homomorphisms from  $Q$  to  $D$  given by extending the homomorphism by mapping all other variables to  $c$ . In addition, in case the queries are CQ\*s, multiplying all annotations of facts matching the source homomorphism and doing the same for the target homomorphism gives the same result. In case the queries are AggCQs, this implies that the aggregation is applied to the same set of assignments, and therefore the aggregates computed in both cases are the same. Together with the fact that  $Q_{|\text{vars}(Q')}$  and  $Q'$  have the same heads, this means that projecting  $Q(D)$  to the variables of  $Q'$  results in  $Q'(D')$  with the same answer order (and the same annotations as in case of CQ\*s), and this projection can be done in constant time.  $\square$

Next, we use Lemma 3.2 to eliminate the existential variables from our queries, by setting  $V = \text{free}(Q)$ . It is folklore that free-connex CQs (over non-annotated databases) can

be transformed into *full* acyclic CQs in linear time [IUV17, OZ15]. The following lemma states that the same holds for free-connex CQ\*s over annotated databases.

**Lemma 3.3.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a logarithmic-time commutative semiring, and let  $Q(\vec{x}, \star, \vec{z})$  be a self-join-free free-connex CQ\*. Then,  $Q_{|\text{free}(Q)}$  is a full acyclic CQ\*, and there is an  $O(|D| \log |D|)$ -time construction that maps  $\mathbb{K}$ -databases  $(D, \tau)$  of  $Q$  to  $\mathbb{K}$ -databases  $(D', \tau')$  of  $Q_{|\text{free}(Q)}$  such that  $Q_{|\text{free}(Q)}(D', \tau') = Q(D, \tau)$ .*

*Proof.* We show an  $O(|D| \log |D|)$ -time construction that maps  $\mathbb{K}$ -databases  $(D, \tau)$  of  $Q$  to  $\mathbb{K}$ -databases  $(D', \tau')$  of  $Q_{|\text{free}(Q)}$  such that  $Q_{|\text{free}(Q)}(D', \tau') = Q(D, \tau)$ .

Since  $Q$  is free-connex, it has an ext-free-connex join tree  $T$  with a subtree  $T'$  which contains precisely the free variables of  $Q$ . We call the vertices of  $T'$  the *free vertices*. The proof idea is as follows: we first adapt the database to have a relation for each vertex of  $T$ , we then build a database with the same answers using only the vertices of  $T'$ , and finally we translate this into a database for  $Q_{|\text{free}(Q)}$ .

The first step is to make it so that each vertex of the tree has a unique associated atom in the query and a unique associated relation in the database. For every vertex of the tree that does not correspond to a relation, we take the relation of a vertex that contains it (without the annotations), project it accordingly, and annotate the facts with  $\bar{1}$ . The query answers are unchanged by the addition of these annotated relations as they contain projections of existing relations annotated with the multiplicative identity.

Next, our goal is to only use relations matching the free vertices and maintain the same query answers. We claim that as long as  $T$  has vertices that are not in the set of free vertices, we can remove such a vertex from  $T$  and remove its associated atom from the query while adapting the database so that the query results remain unchanged.

As long as  $T$  has a vertex not in the set of free vertices, eliminate such a leaf vertex as follows:

- (1) In the relation associated with the vertex, project out every variable not shared between the vertex and its neighbor while summing ( $\oplus$ ) annotations.
- (2) Join the obtained relation into the relation of the neighbor vertex while multiplying ( $\otimes$ ) annotations.
- (3) Remove the leaf from the tree, its associated atom from the query, and its associated relation from the database.

More specifically, denote by  $S$  the variables shared between the treated leaf vertex  $v$  and its neighbor. During step (1), we initialize a lookup table where for each unique assignment of  $S$  matching a fact in the relation of  $v$ , we set the semiring value  $\bar{0}$ . Then, we iterate over each fact of the relation of  $v$  and add its annotation (using the semiring  $\oplus$ ) to the entry matching its assignment to  $S$ . Next, during step (2), we iterate over every fact in the relation matching the neighbor of  $v$  and multiply its existing annotation (using the semiring  $\otimes$ ) with the value stored for its assignment to  $S$ . Once we are done, each fact in the relation matching the neighbor of  $v$  had its annotation multiplied by the sum of the annotations of the facts that agree with it in the relation of  $v$ . In terms of complexity, we perform the elimination of a single vertex in loglinear time. Indeed, the computation model allows constructing in linear time lookup tables that can be accessed in constant time, and we perform a linear number of semiring operations, each taking logarithmic time. Since the number of elimination steps is bounded by the size of the query, the entire elimination takes loglinear time.

As for the correctness, consider the variables that appear in  $v$  but not in its neighbor. Since  $v$  is chosen to be a leaf, the running intersection property guarantees that these variables only appear in  $v$ . This means that they do not appear in the free vertices, so they are existential and must be projected out. As joins and projections are commutative over annotated databases [GKT07], we can do the projection first. As these variables appear only in  $v$ , the projection can be done locally. Thus, the results of the modified query over the modified database remain unchanged.

Finally, we adapt the relations matching the free vertices into a database for  $Q|_{\text{free}(Q)}$ . We first claim that for each atom  $\varphi$  of  $Q$ , we get that  $\text{free}(\varphi)$  is contained in some free vertex. As  $T$  is an ext-free-connex tree,  $\text{vars}(\varphi)$  is a vertex of  $T$ . If  $\text{vars}(\varphi)$  is a free vertex, then  $\text{vars}(\varphi) = \text{free}(\varphi)$  is a free vertex. Otherwise, consider the nearest free vertex to  $\text{vars}(\varphi)$ . Since  $T'$  contains all free variables of  $Q$ , in particular  $\text{free}(\varphi)$ , then from the running intersection property, we know that the nearest vertex must also contain  $\text{free}(\varphi)$ .

Consider a vertex of  $T'$  which does not match an atom of  $Q|_{\text{free}(Q)}$ . Since it is a vertex of  $T'$ , we know that it only contains free variables, and that its variables are contained in some atom of  $Q$ . This means that its variables are also contained in some atom of  $Q|_{\text{free}(Q)}$  and thus some other vertex of  $T'$ . Therefore, we can eliminate  $v$  (and its relation counterpart) by joining its relation into the relation of the containing vertex in the same fashion as in step (2). Next, consider an atom of  $Q|_{\text{free}(Q)}$  which does not have a matching vertex in  $T'$ . We know that  $T'$  has a vertex  $u$  containing the variables of this atom. For this atom, we take the relation of  $u$  (without the annotations), project it accordingly, and annotate the facts with  $\bar{1}$ . As we only added facts projected from existing relations and annotated by  $\bar{1}$ , the query answers and annotations remain unchanged.  $\square$

We note that in the case where the semiring operations require only constant time, or in the case that Lemma 3.2 is applied to CQs or AggCQs, the construction time of Lemma 3.2 and Lemma 3.3 is linear time instead of loglinear time. In addition, when using Lemma 3.3, two free variables are neighbors in  $Q|_{\text{free}(Q)}$  if and only if they are neighbors in  $Q$ , and since the order remains the same,  $Q|_{\text{free}(Q)}$  will have a disruptive trio if and only if  $Q$  has a disruptive trio.

Lemma 3.3 shows that, if direct access for  $Q|_{\text{free}(Q)}$  over  $\mathbb{K}$ -databases is in  $\langle T_p, T_a \rangle$ , then direct access for  $Q$  over  $\mathbb{K}$ -databases is in  $\langle \text{loglinear} + T_p, T_a \rangle$ . By combining this with Lemma 3.2 (for  $V = \text{free}(Q)$ ), we get that  $Q$  and  $Q|_{\text{free}(Q)}$  are equally hard for direct access.

**Theorem 3.4.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a logarithmic-time commutative semiring, and let  $Q(\vec{x}, \star, \vec{z})$  be a self-join-free free-connex CQ $^*$ . Then, direct access for  $Q$  is in  $\langle T_p, T_a \rangle$  if and only if direct access for  $Q|_{\text{free}(Q)}$  is in  $\langle T_p, T_a \rangle$ , assuming  $T_p = \Omega(\text{loglinear})$ .*

The positive side of Theorem 3.4 can easily be used also in case  $Q$  contains self-joins by removing the self-joins while duplicating the relevant relations. We say that a query  $Q^{\text{sf}}$  is a self-join-free version of a query  $Q$  if it can be obtained from  $Q$  by replacing the relation symbols such that each relation symbol appears at most once in  $Q^{\text{sf}}$ .

**Corollary 3.5.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a logarithmic-time commutative semiring, and let  $Q(\vec{x}, \star, \vec{z})$  be a free-connex CQ $^*$ . Let  $Q^{\text{sf}}$  be a self-join-free version of  $Q$ . If direct access for  $Q|_{\text{free}(Q)}$  is in  $\langle \text{loglinear}, \log \rangle$ , then direct access for  $Q$  is in  $\langle \text{loglinear}, \log \rangle$ .*

## 4. INCORPORATING ANNOTATION AND AGGREGATION IN THE ANSWERS

In this section, we discuss the existence of efficient direct access in the case where the order *does not* involve the computed value, that is, the annotation (for CQ<sup>\*</sup>s) or the aggregate values (for AggCQs). Equivalently, these are queries where the computed value is last in order, that is, the vector  $\vec{z}$  in the head is empty. Hence, we focus on CQ<sup>\*</sup>s of the form  $Q(\vec{x}, \star)$  and AggCQs of the form  $Q(\vec{x}, \alpha(\vec{w}))$ . In other words, the problem is similar to the CQ case, except that the access algorithm should also retrieve the aggregated value from the data structure. In Section 4.1, we will use annotated databases to identify the cases where this can be done efficiently for min, max, count, sum, and average. By contrast, in Section 4.2 we will show that for count-distinct, the class of tractable cases is more restricted unless the domain of the elements we count is small (logarithmic size).

**4.1. Generalized Dichotomies.** We now show that Theorem 3.1 extends to databases with annotations, and so, also to queries with aggregate functions that can be efficiently simulated by annotations.

**Theorem 4.1.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a logarithmic-time commutative semiring, and let  $Q(\vec{x}, \star)$  be a CQ<sup>\*</sup>.*

- (1) *If  $Q$  is free-connex with no disruptive trio, then direct access for  $Q$  is in  $\langle \text{loglinear}, \log \rangle$  on  $\mathbb{K}$ -databases.*
- (2) *Otherwise, if  $Q$  is also self-join-free, then direct access for  $Q$  is not in  $\langle \text{loglinear}, \log \rangle$ , assuming the HYPERCLIQUE and SparseBMM hypotheses.*

*Proof.* For the negative side of the dichotomy, we simply use the negative side of Theorem 3.1. This can be done since each answer to a CQ<sup>\*</sup> contains the ordinary (non-annotated) answer to the CQ obtained by removing  $\star$ , and the answers have the same order.

For the positive side of the dichotomy, Corollary 3.5 implies that it is enough to show an algorithm for  $Q'(\vec{x}, \star) = Q_{\text{free}(Q)}^{\text{sf}}$ , which is a full self-join-free acyclic CQ<sup>\*</sup> with no disruptive trios. Using Theorem 3.1, we have direct access for  $Q'$  in  $\langle \text{loglinear}, \log \rangle$  if we ignore the annotations. Since  $Q'$  is full, from an answer to  $Q'$  we compute the annotation by finding the tuples that agree with the answer in every relation of  $Q'$  and multiply the annotations of these tuples; this can be done in logarithmic time.  $\square$

From the positive side of Theorem 4.1, we conclude efficient direct access for AggCQs  $Q(\vec{x}, \alpha(\vec{w}))$ , as long as we can efficiently formulate the aggregate function as an annotation over some logarithmic-time commutative semiring. This is stated in the following corollary of Theorem 4.1.

**Corollary 4.2.** *Consider an AggCQ  $Q(\vec{x}, \alpha(\vec{w})) := \varphi_1(\vec{x}, \vec{y}), \dots, \varphi_\ell(\vec{x}, \vec{y})$  where  $\alpha$  is one of Min, Max, Count, Sum, and Avg.*

- (1) *If the CQ  $Q'(\vec{x}) := \varphi_1(\vec{x}, \vec{y}), \dots, \varphi_\ell(\vec{x}, \vec{y})$  is free-connex with no disruptive trio, then direct access for  $Q$  is in  $\langle \text{loglinear}, \log \rangle$ .*
- (2) *Otherwise, if  $Q$  is also self-join-free, then direct access for  $Q$  is not in  $\langle \text{loglinear}, \log \rangle$ , assuming the HYPERCLIQUE and SparseBMM hypotheses.*

*Proof.* For the positive side, we simply apply Theorem 4.1 with the corresponding semiring. In the case where  $\alpha$  is Avg, we compute Sum and Count separately and divide the results. The negative side carries over from Theorem 3.1 since a direct access solution for  $Q$  in

$\langle \text{loglinear}, \text{log} \rangle$  is also a direct access solution for  $Q'$  in  $\langle \text{loglinear}, \text{log} \rangle$  if we ignore the aggregated values.  $\square$

**Remark 4.3.** Corollary 4.2 can be easily extended to support multiple aggregate functions  $\alpha_1(\vec{w}_1), \dots, \alpha_k(\vec{w}_k)$ . For that, we can simply solve the problem for each  $\alpha_i(\vec{w}_i)$  separately, and extract the aggregate values of an answer from the  $k$  data structures that we construct in the preprocessing phase. (Moreover, a practical implementation can handle all aggregate values in the same structure.)  $\square$

**4.2. Count Distinct.** Can we generalize Corollary 4.2 beyond the stated aggregate functions? The most notable missing aggregate function is CountD (count-distinct). Next, we show that we *cannot* have similar tractability for count-distinct, and we illustrate it with a specific query (Lemma 4.4). After that, we show how precisely Corollary 4.2 changes for count-distinct, and particularly which AggCQs move from the tractable to the intractable side (Theorem 4.5). We later explain that the tractable cases of Corollary 5.10 are fully restored if we assume that the domain size is logarithmic in that of the input (Remark 4.7).

The negative results in this section require the *small-universe Hitting Set Conjecture* (HSC) [VW15]. In HSC, we are given two sets  $\mathcal{U}$  and  $\mathcal{V}$  of size  $N$ , each containing sets over the universe  $\{1, 2, \dots, d\}$ , and the goal is to determine whether  $\mathcal{U}$  contains a set that hits (i.e., shares an element with) every set in  $\mathcal{V}$ . HSC states that the problem takes  $N^{2-o(1)}$  time for every function  $d = \omega(\log(N))$ . (In fact, it is conjectured that even a randomized algorithm for this problem needs  $N^{2-o(1)}$  time in expectation [VW15].)

**Lemma 4.4.** *Assuming HSC, for all  $k \geq 1$ , direct access is not in  $\langle \text{loglinear}, \text{log} \rangle$  for  $Q_H^k(x, \text{CountD}(w)) :- R_0(x, y_1), R_1(y_1, y_2), \dots, R_{k-1}(y_{k-1}, y_k), R_k(y_k, w)$ .*

*Proof.* Let  $\mathcal{U} = \{U_1, U_2, \dots, U_N\}$  and  $\mathcal{V} = \{V_1, V_2, \dots, V_N\}$  be sets of sets of elements of the universe  $\{1, 2, \dots, d\}$  where  $d = N^c$  for some  $0 < c < 1$ . Indeed,  $d = \omega(\log N)$ , as required by the conjecture. We construct a database with the fact  $R_0(i, j)$  for all  $j \in U_i$ , the fact  $R_k(j, i)$  for all  $j \in V_i$ , and the facts  $R_i(j, j)$  for all  $j$  in the universe and  $0 < i < k$ .

Next, we assume efficient direct access for  $Q_H^k$ , and use that to solve the hitting-set problem. Each query answer is a pair  $(i, n)$  where  $i$  is the index of a set  $U_i$  from  $\mathcal{U}$  and  $n$  is the number of sets  $V_j$  that  $U_i$  hits. By accessing all query answers, we can check all sets in  $\mathcal{U}$ , one by one, and test whether any is hitting all  $N$  sets. The number of facts in  $D$  is  $O(dN + d^2) \leq O(dN)$ , and the number of query answers (and so the number of access calls) is  $N$ . Hence, direct access for  $Q_H^k$  in  $\langle \text{loglinear}, \text{log} \rangle$  would imply a solution to the hitting-set problem in better than  $N^{2-o(1)}$  time. Indeed, for any  $d = O(N^c)$ , we get a solution in  $O(N^{1+c} \cdot \log N)$  time, which contradicts HSC for  $c < 1$ .  $\square$

Importantly, the reduction used in the proof of Lemma 4.4 does not involve the order, and hence, the theorem holds true even for direct access, or enumeration, without any order requirement.

Next, we generalize the hardness proof from Lemma 4.4 to a general class of AggCQs with count-distinct, and we devise an algorithm for other count-distinct queries. We show that this results in a dichotomy for count-distinct. The dichotomy states the following. For the aggregate functions of Corollary 4.2, we have established that the same dichotomy for standard CQs applies when we “ignore” the aggregation in the head and treat it as an *existential* variable. The case of count-distinct turns out to be the opposite: the dichotomy for standard CQs applies if we treat the aggregated variable as a *free* variable.

**Theorem 4.5.** Consider an AggCQ  $Q(\vec{x}, \text{CountD}(w)) :- \varphi_1(\vec{x}, \vec{y}), \dots, \varphi_\ell(\vec{x}, \vec{y})$ .

- (1) If the CQ  $Q'(\vec{x}, w) :- \varphi_1(\vec{x}, \vec{y}), \dots, \varphi_\ell(\vec{x}, \vec{y})$  is free-connex with no disruptive trio, then direct access for  $Q$  is in  $\langle \text{loglinear}, \text{log} \rangle$ .
- (2) Otherwise, if  $Q$  is also self-join-free, then direct access for  $Q$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming the HYPERCLIQUE, SparseBMM, and HSC hypotheses.

*Proof.* We start with the positive side of the dichotomy, and assume we are given a database  $D$  as input. First, we remove self-joins and eliminate the existential variables that are not part of the aggregation. We can remove self-joins by replacing the relation symbols with unique relation names and copying the relations into their new names. Since  $Q'(\vec{x}, w)$  is free-connex with no disruptive trios, we can build in  $O(|D| \log |D|)$  time a full acyclic CQ  $Q^{\text{full}}(\vec{x}, w)$ , without self-joins or disruptive trios, and a database  $D'$  for  $Q^{\text{full}}$  such that  $Q^{\text{full}}(D') = Q'(D)$ . This transformation is folklore and can be seen as a special case of our construction from Lemma 3.3 when we ignore the annotation. We illustrate the construction in Example 4.6 following the proof.

As  $Q^{\text{full}}$  has no disruptive trio, we know that it has an atom that contains its last free variable and all of its neighbors [CTG<sup>+</sup>23, proof of Lemma 3.9]. In other words, when considering all the atoms that contain  $w$ , there is a maximal atom  $\varphi$  that contains all of their variables. We filter the relation of  $\varphi$  according to all of these atoms. That is, we treat every fact as a function mapping variables to values, and we remove a fact  $f_1$  from the relation of  $\varphi$  if there exists another atom containing  $w$  that does not contain a fact that agrees with  $f_1$  on their common variables. Notice that this filtering does not affect the query answers, and that the other atoms that contain  $w$  are now no longer needed: the query will have the same answers if they are removed.

Consider the CQ  $Q^{\text{count}}(\vec{x}, w)$  with the same body as  $Q^{\text{full}}$  but without all atoms that contain  $w$ , and with the new atom  $\varphi'$ , where the relation for  $\varphi'$  is obtained from the filtered  $\varphi$  by replacing the  $w$  values with their counts. That is, we sort  $\varphi$  according to all variables other than  $w$ , then scan the relation and replace every set of tuples that agree on the other variables with one tuple where the  $w$  value is the number of tuples we removed in this replacement.

Now we can use the known algorithm from Theorem 3.1 for free-connex CQs with no disruptive trios to solve  $Q^{\text{count}}$ . As we argued that all the transformations above maintain the query answers, this algorithm results in exactly  $Q(D)$ . We also notice that these transformations can be done in  $O(|D| \log |D|)$  time.

Next, we prove the negative side of the dichotomy. That is, we assume that  $Q'$  is either not free-connex or contains a disruptive trio and prove the hardness of  $Q$ . For this, we use the following characterization of free-connex CQs. Given a CQ  $Q$ , a *free-path* is a chordless path  $(x_1, y_1, \dots, y_k, x_2)$  in  $H(Q)$  with  $k \geq 1$ , such that  $x_1, x_2 \in \text{free}(Q)$ , and  $y_1, \dots, y_k \notin \text{free}(Q)$ . Here, a *chordless path* is a simple path (i.e., a sequence of distinct vertices where every two consecutive vertices co-occur in at least one hyperedge) where no two non-consecutive vertices are neighbors. An acyclic CQ  $Q$  is free-connex if and only if it does not contain a free-path [BDG07].

Consider the CQ  $Q''(\vec{x})$  with the same body as  $Q$  and  $Q'$ , but where we removed the aggregated value from the query head. The first case is that  $Q''$  is not free-connex or has disruptive trios. In this case, according to Theorem 3.1, direct access for  $Q''$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ . Since each answer to  $Q$  contains an answer to  $Q''$  and the answers have the same order, we can use  $Q$  to solve  $Q''$ . So, we conclude that  $Q$  has no direct access in

$\langle \text{loglinear}, \text{log} \rangle$  either. The second case is that  $Q''$  is free-connex with no disruptive trios. If  $Q'$  has a disruptive trio, its last variable has to be  $w$  since  $Q''$  has no disruptive trios. In other words, the disruptive trio comprises of  $\{x_i, x_j, w\}$  where  $x_i, x_j \in \vec{x}$  and  $(x_i, w, x_j)$  is a chordless path. This chordless path constitutes a free-path in  $Q''$ , contradicting the fact that  $Q''$  is free-connex. So, we conclude that  $Q'$  has no disruptive trios. This means that  $Q'$  is acyclic (since  $Q''$  is acyclic and they have the same body) but not free-connex. Thus,  $Q'$  has a free-path that is not a free-path in  $Q''$ , and so it must use  $w$ . We conclude that this path is of the form  $(x_i, v_1, \dots, v_k, w)$  with  $x_i \in \vec{x}$  and  $v_1, \dots, v_k \in \vec{y}$ . Since all queries we discussed have the same body, this chordless path also appears in  $Q$ .

Since  $(x_i, v_1, \dots, v_k, w)$  is a chordless path, the AggCQ  $Q_{\{x_i, v_1, \dots, v_k, w\}}$  is structurally equivalent to

$$Q_H^k(x, \text{CountD}(w)) :- R_0(x, y_1), R_1(y_1, y_2), \dots, R_{k-1}(y_{k-1}, y_k), R_k(y_k, w)$$

where  $x_i, v_j$ , and  $w$  take the roles of  $x, y_j$ , and  $w$  respectively, for all  $1 \leq j \leq k$ . Lemma 4.4 determines that  $Q_H^k$  has no direct access in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming HSC. Using Lemma 3.2, we conclude that direct access for  $Q$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ .  $\square$

The following example illustrates the construction in the positive side of the proof.

**Example 4.6.** Consider the following AggCQ:

$$Q(x_1, x_2, x_3, \text{CountD}(w)) :- R(x_1, w, y), R(x_1, x_2, w), S(x_2, x_3)$$

For this query we construct

$$Q^{\text{full}}(x_1, x_2, x_3, w) :- R_1(x_1, w), R_2(x_1, x_2, w), S(x_2, x_3),$$

which is self-join-free query and without existential variables, by replacing  $R$  with unique names and removing  $y$ . We then get

$$Q^{\text{count}}(x_1, x_2, x_3, w) :- R'_2(x_1, x_2, w), S(x_2, x_3)$$

by keeping only the maximal atom containing  $w$ . Given  $R$  and  $S$  as follows, the following relations will be built.  $R_1$  is obtained from  $R$  by projecting away the third column.  $R'_2$  is obtained by filtering  $R_2$  according to  $R_1$  (which removes the fact  $R_2(a_1, b_2, b_3)$ ), then grouping by the first two columns and counting.

$S$	$R = R_2$	$R_1$	$R'_2$
$b_1 \ c_1$	$a_1 \ b_1 \ b_1$	$a_1 \ b_1$	$a_1 \ b_1 \ 2$
$b_2 \ c_1$	$a_1 \ b_1 \ b_2$	$a_1 \ b_2$	$a_1 \ b_2 \ 1$
$b_2 \ c_2$	$a_1 \ b_2 \ b_1$	$a_2 \ b_2$	$a_2 \ b_2 \ 1$
	$a_1 \ b_2 \ b_3$		
	$a_2 \ b_2 \ b_2$		

**Remark 4.7.** Theorem 4.5 shows which AggCQs belong to the positive side of Corollary 4.2, but become intractable when the aggregation is replaced with count-distinct. Nevertheless, it is important to observe that there are cases where Theorem 4.1 can be used to compute AggCQs with count-distinct even for the intractable cases of Theorem 4.5: when the size of the domain  $\mathcal{X}$  of the distinct elements we count is bounded by a logarithm in the input size  $|D|$ . Such an assumption can be realistic when we count, say, distinct categories from a small ontology (e.g., item categories in a sales context), distinct countries from a small collection of countries, and so on. In such cases, we can compute the exact set of distinct

elements, and not just their count, by using the set semiring  $(\mathcal{P}(\mathcal{X}), \cup, \cap, \emptyset, \mathcal{X})$  since the operations  $\cup$  and  $\cap$  can be performed in logarithmic time for logarithmic domains.  $\square$

## 5. INCORPORATING THE ANNOTATION AND AGGREGATION IN THE ORDER

The results of the previous section apply when the lexicographic order does not include the computed value, or equivalently when the computed value is last in the head of the query. In this section, we explore the ability to include the computed value earlier in the lexicographic order. To this end, we assume that the underlying commutative semiring has an ordered domain. When the domain is numerical, we will implicitly assume the natural order (i.e., smaller numbers precede larger numbers) without mentioning it. In terms of the computational model, we assume that we can compare two given elements of the domain in time logarithmic in the size of the input.

Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a commutative semiring with an underlying order  $\succeq$  over  $\mathbb{K}$ . The semiring is said to be  $\otimes$ -monotone if for every  $c \in \mathbb{K}$  the function  $f_c : \mathbb{K} \rightarrow \mathbb{K}$ , defined by  $f_c(y) = c \otimes y$ , is monotone. This means that either  $c \otimes a \succeq c \otimes b$  whenever  $a \succeq b$ , or  $c \otimes a \preceq c \otimes b$  whenever  $b \succeq a$ . Computationally, we assume that we can determine efficiently (in logarithmic time in the input) whether a given  $c$  is such that the function  $f_c(x) = c \otimes x$  is non-decreasing or non-increasing. All specific semirings that we mention in the paper are  $\otimes$ -monotone. Our main result for this section is the following dichotomy. Recall that the 3SUM conjecture [GO95, Pat10] states it takes  $N^{2-o(1)}$  time to determine whether a given set of  $N$  elements from  $\{-N^3, \dots, N^3\}$  contains distinct elements  $a, b, c$  such that  $a + b = c$ .

**Theorem 5.1.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a  $\otimes$ -monotone logarithmic-time commutative semiring, and  $Q(\vec{x}, \star, \vec{z})$  a  $CQ^\star$ .*

- (1) *If  $Q$  is free-connex with no disruptive trio and there exists an atom of  $Q$  containing  $N(\vec{z}) \cap \text{free}(Q)$ , then direct access for  $Q$  is in  $\langle \log\text{linear}, \log \rangle$ .*
- (2) *Otherwise, if  $Q$  is also self-join-free and  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  is one of the counting, numerical, max tropical, or min tropical semirings, then direct access for  $Q$  is not in  $\langle \log\text{linear}, \log \rangle$ , assuming the 3SUM, SparseBMM, and HYPERCLIQUE hypotheses.*

Recall that the notation  $N(\vec{z})$  contains also all variables of  $\vec{z}$ . Before we prove Theorem 5.1, we illustrate the tractability condition with a simple special case: If  $Q(\vec{x}, \star, \vec{z})$  is free-connex and all variables of  $\vec{z}$  act “as a single variable,” then direct access for  $Q$  is in  $\langle \log\text{linear}, \log \rangle$ . More formally:

**Corollary 5.2.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a  $\otimes$ -monotone logarithmic-time commutative semiring, and  $Q(\vec{x}, \star, \vec{z})$  a free-connex  $CQ^\star$  with no disruptive trio. If every atom of  $Q$  contains either all variables of  $\vec{z}$  or none of them, then direct access for  $Q$  is in  $\langle \log\text{linear}, \log \rangle$ .*

*Proof.* We first use Corollary 3.5 to eliminate the existential variables of  $Q$  and work with a full acyclic  $CQ^\star$   $Q_{|\text{free}(Q)}(\vec{x}, \star, \vec{z})$ . We notice that  $Q_{|\text{free}(Q)}$ , like  $Q$ , is such that every atom contains either all variables of  $\vec{z}$  or none of them and there is no disruptive trio. If  $\vec{z}$  is empty, then every atom contains  $N(\vec{z})$ . Otherwise, denote the last variable in  $\vec{z}$  as  $z_{\text{last}}$ . It has been established that a full acyclic  $CQ$  with no disruptive trios has an atom that contains the last free variable and all of its neighbors [CTG<sup>+</sup>23, proof of Lemma 3.9]; hence,  $Q_{|\text{free}(Q)}$  has an atom  $\varphi$  that contains  $z_{\text{last}}$  and all of its neighbors. Since each atom of  $Q$  contains either all or none of the variables of  $\vec{z}$ , we know that  $N(\vec{z}) = N(z_{\text{last}})$ , hence  $\varphi$  contains every

variable in  $N(\vec{z})$ . From Theorem 5.1 we can now conclude that direct access for  $Q|_{\text{free}(Q)}$  is in  $\langle \text{loglinear}, \text{log} \rangle$ .  $\square$

For example, Corollary 5.2 (and Theorem 5.1) imply that direct access for the following CQ $^*$  is in  $\langle \text{loglinear}, \text{log} \rangle$  over databases annotated with the numerical semiring.

$$Q(x_1, x_2, \star, z_1, z_2) :- R(x_1, x_2), S(x_2, z_1, z_2), T(z_1, z_2)$$

The variables that follow the computed value  $\star$  are  $z_1$  and  $z_2$ , and indeed, every atom either contains both  $z_1$  and  $z_2$  (as the second and third atoms) or contains none of them (as the first atom). Hence, direct access is tractable according to Corollary 5.2. Similarly, the atom  $S$  contains  $N(\vec{z})$  (i.e., the variables  $x_2, z_1, z_2$ ), so Theorem 5.1 also identifies this query as tractable.

However, Corollary 5.2 fails to capture the tractability of the following example.

$$Q(x_1, x_2, \star, z_1, z_2) :- R(x_1, x_2), S(x_2, z_1, z_2), T(z_2)$$

The variables  $z_1$  and  $z_2$  do not occur jointly in all atoms, as only  $z_2$  appears in  $T(z_2)$ . Yet, there is an atom that contains both  $z_1$  and  $z_2$ , so direct access for  $Q$  is in  $\langle \text{loglinear}, \text{log} \rangle$  according to Theorem 5.1.

As a negative example, consider the simplest possible Cartesian product CQ $^*$   $R \times S$  for unary  $R$  and  $S$ , where we wish to have the annotation *first* in the order.

$$Q_{\star \times}(\star, x, y) :- R(x), S(y). \quad (5.1)$$

This CQ $^*$  falls in the hard side of Theorem 5.1, since there is no atom that contains both  $x$  and  $y$ . In contrast, Theorem 5.1 implies that whenever  $\star$  is not first, as is the case with  $(x, \star, y)$ , the query is tractable.

In the remainder of this section, we prove Theorem 5.1. We begin with the tractability side (Section 5.1), then complete the proof with the hardness side (Section 5.2), and discuss the implications for AggCQs (Section 5.3).

**5.1. Tractability Side of Theorem 5.1.** The following lemma states the tractability side of Theorem 5.1. The proof includes the algorithm.

**Lemma 5.3.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a  $\otimes$ -monotone logarithmic-time commutative semiring, and  $Q(\vec{x}, \star, \vec{z})$  a free-connex CQ $^*$  with no disruptive trio. If  $Q$  has an atom that contains  $N(\vec{z}) \cap \text{free}(Q)$ , then direct access for  $Q$  is in  $\langle \text{loglinear}, \text{log} \rangle$ .*

*Proof.* The proof idea is as follows. After eliminating the projection (Lemma 3.3), an answer is obtained by joining a selected fact from each relation, and its annotation is obtained by the multiplication of the annotations of the selected facts. Denote by  $A_{\text{before}}$  the set of atoms that contain only variables of  $\vec{x}$ , and by  $A_{\text{after}}$  the set of remaining atoms. The annotation of each answer can be expressed as a product  $c_{\text{before}} \otimes c_{\text{after}}$ , where  $c_{\text{before}}$  is determined once  $\vec{x}$  is assigned values, and  $c_{\text{after}}$  is determined once the remaining variables are assigned values. Once  $c_{\text{before}}$  is determined, sorting by  $\star = c_{\text{before}} \otimes c_{\text{after}}$  is the same as sorting by  $c_{\text{after}}$ . Let  $\varphi$  be an atom that contains  $N(\vec{z}) \cap \text{free}(Q)$ .

We show that we can extend  $\varphi$  with a fresh variable, assign this variable with the annotation  $c_{\text{after}}$ , and then use a known direct-access solution on the extended database and query. We explain this idea in the remainder of the proof.

We first eliminate the existential variables of  $Q$  to work with a full acyclic CQ $^*$   $Q|_{\text{free}(Q)}(\vec{x}, \star, \vec{z})$ . According to Lemma 3.3, a  $\mathbb{K}$ -database  $(D, \tau)$  can be transformed in

$O(|D| \log |D|)$  time to a  $\mathbb{K}$ -database  $(D_{\text{full}}, \tau_{\text{full}})$  such that  $Q(D, \tau) = Q_{|\text{free}(Q)}(D_{\text{full}}, \tau_{\text{full}})$ . We notice that the obtained  $Q_{|\text{free}(Q)}$ , like  $Q$ , has an atom  $\varphi_{\text{full}}$  that contains  $N_Q(\vec{z}) \cap \text{free}(Q) = N_{Q_{|\text{free}(Q)}}(\vec{z})$ . Every atom in  $A_{\text{after}}$  contains at least one variable in  $\vec{z}$ , thus  $\varphi_{\text{full}}$  contains every variable that occurs in any atom in  $A_{\text{after}}$ .

Next, we add a new variable  $y$  to  $Q_{|\text{free}(Q)}$ :

- We add  $y$  to the atom  $\varphi_{\text{full}}$ ;
- Replace the head  $(\vec{x}, \star, \vec{z})$  with  $(\vec{x}, y, \vec{z}, \star)$ .

We denote the resulting CQ $^*$  by  $Q_{|\text{free}(Q)}^{+y}$ .

We first prove that direct access for  $Q_{|\text{free}(Q)}^{+y}$  is in  $\langle \text{loglinear}, \text{log} \rangle$ . Due to Theorem 4.1, it suffices to show that  $Q_{|\text{free}(Q)}^{+y}$  is acyclic and that it does not contain a disruptive trio. To see that  $Q_{|\text{free}(Q)}^{+y}$  is acyclic, we can use the same join tree as  $Q_{|\text{free}(Q)}$ ; the running-intersection property is preserved since we add  $y$  to only one atom. Assume, by way of contradiction, that there is a disruptive trio. As  $Q_{|\text{free}(Q)}$  has no disruptive trio, and the variable sequence of  $Q_{|\text{free}(Q)}^{+y}$  differs from that of  $Q_{|\text{free}(Q)}$  by  $y$  alone, any disruptive trio in  $Q_{|\text{free}(Q)}^{+y}$  must include  $y$ .

Suppose that  $y$  forms a disruptive trio together with  $x_a$  and  $x_b$ . Recall that  $y$  was added to only one atom, namely  $\varphi_{\text{full}}$ , and so all of its neighbors are neighbors of each other. Therefore, the only possibility for  $\{y, x_a, x_b\}$  to form a disruptive trio is the following:

- $x_a$  and  $x_b$  are neighbors in an atom  $\varphi'$ ;
- $x_a$  and  $y$  are non-neighbors;
- $x_b$  and  $y$  are neighbors in  $\varphi_{\text{full}}$ ;
- $x_b$  appears after  $y$  and  $x_a$  in the head of  $Q_{|\text{free}(Q)}^{+y}$ .

Since  $x_b$  appears after  $y$ , we know that  $x_b$  is in  $\vec{z}$ . Hence,  $\varphi'$  is in  $A_{\text{after}}$ . As  $\varphi_{\text{full}}$  contains the variables of all atoms in  $A_{\text{after}}$ , it also contains  $x_a$ , contradicting the fact that  $x_a$  and  $y$  are non-neighbors. In conclusion, direct access for  $Q_{|\text{free}(Q)}^{+y}$  is in  $\langle \text{loglinear}, \text{log} \rangle$ , as claimed.

Next, we show how we can use direct access to  $Q_{|\text{free}(Q)}^{+y}$  to obtain direct access for  $Q_{|\text{free}(Q)}$ . Given a  $\mathbb{K}$ -database  $(D_{\text{full}}, \tau_{\text{full}})$  for  $Q_{|\text{free}(Q)}$ , we start by modifying the  $\mathbb{K}$ -database to obtain a new  $\mathbb{K}$ -database  $(D_{\text{full}}^{+y}, \tau_{\text{full}}^{+y})$  for  $Q_{|\text{free}(Q)}^{+y}$ . Let  $R_\varphi$  be the relation symbol of  $\varphi_{\text{full}}$ . For any  $R_\varphi$ -fact  $f$  in  $D_{\text{full}}$ , and for every atom  $\varphi'$  in  $A_{\text{after}}$ , if  $D_{\text{full}}$  does not contain a fact  $f'$  such that  $f$  and  $f'$  agree w.r.t.  $\varphi_{\text{full}}$  and  $\varphi'$ , then we can ignore  $f$ . Also, recall that  $\varphi_{\text{full}}$  contains all variables that occur in  $A_{\text{after}}$ . Hence, for every  $R_\varphi$ -fact  $f$  of  $D_{\text{full}}$  and atom  $\varphi' \in A_{\text{after}}$ , there is at most one fact  $f'$  of  $\varphi'$  that agrees with  $f$ . We set the value of  $y$  for every  $R_\varphi$ -fact  $f$  of  $D_{\text{full}}^{+y}$  to be the product of the annotations of the facts that agree with  $f$  in the relations of atoms in  $A_{\text{after}}$ . Since  $|A_{\text{after}}|$  is a constant in terms of data complexity, we can compute the product in constant time per fact, and the construction takes linear time. The annotations of facts of  $D_{\text{full}}^{+y}$  will be identical to the annotations of the corresponding facts of  $D_{\text{full}}$  as the annotations are unchanged.

We apply the direct-access algorithm on  $D_{\text{full}}^{+y}$  twice: once normally, and once when reversing the order over the domain of  $y$ . This gives us two answers,  $\vec{a}$  and  $\vec{a}_r$ , that have the same assignment to  $\vec{x}$ . We compute the result of  $c_{\text{before}}$  of these answers and check whether multiplying by  $c_{\text{before}}$  is monotone non-decreasing or monotone non-increasing (which we assume is doable in logarithmic time). If it is monotone non-decreasing, then sorting by  $c_{\text{after}}$  is identical to sorting by  $c_{\text{before}} \otimes c_{\text{after}}$  and we choose  $\vec{a}$  as the answer. If it is monotone

non-increasing, then sorting by  $c_{\text{after}}$  in reverse order is identical to sorting by  $c_{\text{before}} \otimes c_{\text{after}}$ , and we choose  $\vec{a}_r$  as the answer. In either case, we return the answer after projecting out variable the value corresponding to  $y$ .<sup>1</sup> Hence, we get direct access for  $Q_{|\text{free}(Q)}$  in  $\langle \text{loglinear}, \text{log} \rangle$ , and so, direct access for  $Q$  in  $\langle \text{loglinear}, \text{log} \rangle$ , as required.  $\square$

**5.2. Hardness Side of Theorem 5.1.** To prove the hardness side of the theorem, we first show hardness of the CQ $^*$   $Q_{\star \times}$  of (5.1). The next lemma states that under the 3SUM conjecture, direct access for  $Q_{\star \times}$  is impossible over  $\mathbb{K}$ -databases with semirings that gave positive results in the previous section.

**Lemma 5.4.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be one of the counting, numerical, max tropical, or min tropical semirings. Direct access for the CQ $^*$   $Q_{\star \times}$  (of (5.1)) is not in  $\langle \text{loglinear}, \text{log} \rangle$  over  $\mathbb{K}$ -databases, assuming the 3SUM conjecture. Moreover, this holds even if the annotation is in the domain  $[1, |D|^4 - 1]$ .*

*Proof.* We will show a reduction from the 3SUM problem. Given a set  $V$  of  $N$  values over the universe  $\{-N^3, \dots, N^3\}$ , we construct a logarithmic-time, homomorphism  $h : \mathbb{Z} \rightarrow \mathbb{K}$ , which we define later on, such that for every  $a, b \in V$  we get  $h(a + b) = h(a) \otimes h(b)$ . The  $\mathbb{K}$ -relation  $(D, \tau)$  that we construct consists of the facts  $R(a)$  and  $S(a)$  for every  $a \neq 0$ . As for the annotation, we define  $\tau(R(a)) = \tau(S(a)) = h(a)$ . Hence, the annotation of an answer  $(a, b)$  is  $h(a) \otimes h(b)$ .

Now, assume that we have efficient direct access to  $Q_{\star \times}(D)$ . For every value  $c \in V$ , we compute  $h(c)$  and use binary search on  $Q(D)$  to find an answer  $(a, b)$  with the annotation  $h(c)$ . If we find such an answer, this means that there exists some  $a, b \in V$  such that  $h(a) \otimes h(b) = h(c)$ . We will show that  $h$  is injective within a subset of  $\mathbb{Z}$  that contains  $a, b$ , and  $a + b$  for all  $a, b \in V$ . As a result, we get that  $h(a) \otimes h(b) = h(c)$  implies  $a + b = c$ . If  $a \neq b$ , then we have found the required triple in  $V$ , and we are done. If  $a = b$ , then we check whether the preceding or the following answer also has the annotation  $h(c)$ . If so, then we have found our triple since at most one of these answers satisfies  $a = b$  (i.e.,  $a = b = c/2$ ). If we do not find the triple we seek after iterating over  $V$ , then we conclude that the answer to the 3SUM problem is negative.

As there are  $N^2$  answers to  $Q_{\star \times}$ , binary search for an answer by the annotation requires  $O(\log N)$  access calls. Given a value  $c \in V$ , using logarithmic-time direct access we can find an answer, if exists, in  $O(\log^2 N)$  time. In total, we handle all values in  $V$  and solve the 3SUM problem in  $O(N \cdot \log^2 N)$  time, contradicting the 3SUM conjecture.

The rest of this proof is dedicated to finding a suitable logarithmic-time homomorphism  $h$ , given  $V$ . When discussing the min tropical or max tropical semiring, we can simply use the identity as the homomorphism since  $\otimes$  is the numerical addition (+). As for the numerical and counting semiring,  $\otimes$  is defined as  $\cdot$ , so our homomorphism is more involved. We handle these in the remainder of the proof.

In the following paragraphs, we make use of standard notations of number theory (see, e.g., [Coh93]). We denote by  $(\mathbb{Z}/m\mathbb{Z})$ , or  $\mathbb{Z}_m$  for short, the group of integers modulo  $m$  with

<sup>1</sup>It is possible to only use one direct-access structure and support accessing a variable in both orders using a suffix-sum in addition to a prefix-sum at preprocessing. During access, we first determine the value for the first  $i$  variables, and then decide in which direction to access the next variable. This would give a practical improvement but the same complexity. Since we do not give the details of the direct access algorithm for CQs in this paper, for the sake of simplicity of arguments, we use the algorithm twice in the proof instead.

the addition (+) operator. We denote with  $(\mathbb{Z}/m\mathbb{Z})^*$ , or  $\mathbb{Z}_m^*$  for short, the group of integers modulo  $m$  with the multiplication ( $\cdot$ ) operator. Recall that  $|\mathbb{Z}_m^*| = \phi(m)$  where  $\phi$  stands for the Euler's totient function.

To define the homomorphism  $h$ , we choose a value  $m_0$  with the following properties:

- (1)  $\phi(m_0) \geq 4N^3 + 1$ .
- (2) The multiplicative modular group  $\mathbb{Z}_{m_0}^*$  has 2 for a generator.
- (3)  $\log m_0 = O(\log N)$ .

We later show how we use these properties of  $m_0$  to define  $h$ . We choose  $m_0 = 3^{\lceil \log_3(4N^3+1) \rceil + 1}$ , and before we proceed, we will prove that indeed all three properties hold for this  $m_0$ . By the definition of the totient function,  $\phi(m_0) = 2 \cdot 3^{\lceil \log_3(4N^3+1) \rceil} \geq 4N^3 + 1$ , so Item 1 holds. Since  $m_0$  is a power of the prime 3, it holds that  $\mathbb{Z}_{m_0}^*$  is cyclic [Coh93, Theorem 1.4.1]. In addition, the number 2 is a generator for each power of 3 because 2 is a generator for  $\mathbb{Z}_{3^2}^*$  [Coh93, Lemma 1.4.5]. Hence, Item 2 holds. Lastly,  $m_0 = 3^{\lceil \log_3(4N^3+1) \rceil + 1} \leq 9 \cdot 3^{\log_3(4N^3+1)} \leq 9 \cdot (4N^3 + 1)$ , and so  $\log m_0 = O(\log N)$  and Item 3 holds.

It remains to construct  $h : \mathbb{Z} \rightarrow \mathbb{K}$  using  $m_0$ , and we do so as follows. Define  $h := h_3 \circ h_2 \circ h_1$  (that is,  $h(a) = h_3(h_2(h_1(a)))$ ) for the following homomorphisms  $h_1$ ,  $h_2$  and  $h_3$ :

- $h_1 : \mathbb{Z} \rightarrow \mathbb{Z}_{\phi(m_0)}$  is defined by  $h_1(a) = a \pmod{\phi(m_0)}$ . For every  $a, b \in V$ , we know that  $a$ ,  $b$ , and  $a + b$  are all in the range of  $[-2N^3, \dots, 2N^3]$ . Due to Item 1, we have  $\phi(m_0) \geq 4N^3 + 1$ , and therefore  $h_1$  is injective for any possible  $a, b, a + b$  originating from  $a, b \in V$ .
- $h_2 : \mathbb{Z}_{\phi(m_0)} \rightarrow \mathbb{Z}_{m_0}^*$  is defined as follows. Due to Item 2, the number 2 is a generator for  $\mathbb{Z}_{m_0}^*$ , and so we can use as  $h_2$  the well-known isomorphism between modular additive groups and cyclic groups:  $h_2(a) = 2^a \pmod{m_0}$ .
- $h_3 : \mathbb{Z}_{m_0}^* \rightarrow \mathbb{K}$  is simply the identity function. Note when discussing the numerical semiring,  $\mathbb{K} = \mathbb{Q}$ , and when discussing the counting semiring,  $\mathbb{K} = \mathbb{N}$ . Therefore,  $h_3$  is well defined in the case of each of our semirings.

Note that  $h_1$  is injective over the range of possible values for  $a, b, a + b$  for  $a, b \in V$ , and  $h_2$  and  $h_3$  are injective in their entire domain. As a composition, for every possible value  $a, b \in V$  the function  $h$  is injective and is a homomorphism:  $h(a + b) = h(a) \otimes h(b)$ . In addition, we can compute  $h(a)$  in logarithmic time: modular exponentiation in  $\mathbb{Z}_{m_0}^*$  can be done using  $O(\log m_0)$  multiplications, so the computation of  $h$  can be done in  $O(\log m_0)$  time. Due to Item 3,  $\log m_0 = O(\log N)$ , and therefore,  $h_1$ ,  $h_2$ , and  $h_3$  all take logarithmic time. In addition, due to Item 3, the values computed by  $h$  can be represented in a logarithmic number of bits and fit in the registers of our RAM model.

Lastly, recall that the values  $h$  outputs are the values used as annotations in the database. Since the output of  $h_2$  is in  $\mathbb{Z}_{m_0}^*$ , and  $m_0 = 3^{\lceil \log_3(4N^3+1) \rceil + 1}$ , then the output of  $h_2$  is in the range of  $[1, \dots, 12N^3 + 2]$ . As  $h_3$  is defined as the identity, we get that the output of  $h$  is in that same range,  $[1, \dots, 12N^3 + 2]$ . For every  $N \geq 13$ , the output of  $h$  is in the range of  $[1, \dots, N^4 - 1]$ .  $\square$

Next, we build on the hardness in Lemma 5.4 to reason about a more general case.

**Lemma 5.5.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be one of the counting, numerical, max tropical, or min tropical semirings and let  $Q(\vec{x}, \star, \vec{z})$  be a full, acyclic, self-join-free  $CQ^*$  with no disruptive trio. If direct access for  $Q$  over  $\mathbb{K}$ -databases is in  $\langle \log\text{linear}, \log \rangle$ , then all variables in  $\vec{z}$  neighbor each other, assuming the 3SUM conjecture.*

*Proof.* We will prove the contrapositive of the lemma: we assume that there are two variables  $z_1, z_2$  in  $\vec{z}$  that are not neighbors and show that direct access for  $Q$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ . More specifically, we show how to use direct access for  $Q$  to obtain direct access for  $Q_{\star \times}(\star, x, y) :- R(x), S(y)$ . Given a  $\mathbb{K}$ -database  $(D_\star, \tau_\star)$  for  $Q_{\star \times}$ , we show how to construct a  $\mathbb{K}$ -database  $(D, \tau)$  for  $Q$  such that  $Q(D, \tau) = Q_{\star \times}(D_\star, \tau_\star)$ . Since by Lemma 5.4 direct access for  $Q_{\star \times}$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ , we get that neither is direct access for  $Q$ .

Since  $z_1$  and  $z_2$  are not neighbors, the atoms of  $Q$  can be partitioned into those containing  $z_1$ , those containing  $z_2$ , and those containing none of them. Denote the sets of these atoms by  $\varphi_{z_1}$ ,  $\varphi_{z_2}$ , and  $\varphi_\emptyset$ , respectively. For every fact  $R(c)$  in  $D_\star$  we add for every atom  $\varphi$  in  $\varphi_{z_1}$  a fact  $f_{c,\varphi}$  to  $D$  derived from  $\varphi$  by replacing  $z_1$  with  $c$  and every other variable with an arbitrary value  $a_0$ . We then choose an arbitrary atom  $\varphi'$  in  $\varphi_{z_1}$  and define  $\tau(f_{c,\varphi'}) = \tau_\star(R(c))$ , and for every other atom  $\varphi$  in  $\varphi_{z_1}$  we define  $\tau(f_{c,\varphi}) = \bar{1}$ . Similarly, for every fact  $S(c)$  in  $D_\star$  we add for every atom  $\varphi$  in  $\varphi_{z_2}$  a fact  $f_{c,\varphi}$  to  $D$  derived from  $\varphi$  by replacing  $z_2$  with  $c$  and every other variable with  $a_0$ . We then choose an arbitrary atom  $\varphi'$  in  $\varphi_{z_2}$  and define  $\tau(f_{c,\varphi'}) = \tau_\star(S(c))$ , and for every other atom  $\varphi$  in  $\varphi_{z_2}$  we define  $\tau(f_{c,\varphi}) = \bar{1}$ . For every atom  $\varphi$  in  $\varphi_\emptyset$  we add a single fact derived from  $\varphi$  by replacing every variable with  $a_0$ , annotated by  $\bar{1}$ .

Now note that there is a bijection from  $Q(D, \tau)$  to  $Q_{\star \times}(D_\star, \tau_\star)$  mapping  $(c_1, c_2)$  to a tuple of the form  $(a_0, \dots, c_1, \dots, c_2, \dots, a_0)$  where  $c_1$  and  $c_2$  are in place of  $z_1$  and  $z_2$  in the head of  $Q_{\star \times}$ . Moreover,  $\tau_\star((c_1, c_2)) = \tau((a_0, \dots, c_1, \dots, c_2, \dots, a_0))$ . This means that we can use  $Q$ ,  $D$  and  $\tau$  to provide direct access for  $Q_{\star \times}$ ,  $D_\star$  and  $\tau_\star$ . As  $|Q|$  is constant and we constructed  $D$  in  $|D_\star| |Q|$  time, we get direct access for  $Q_{\star \times}$  in  $\langle \text{loglinear}, \text{log} \rangle$  time.  $\square$

We also build on the hardness of queries with disruptive trios, as stated in Theorem 3.1, to deduce the hardness in case there is a variable before the annotation that neighbors some but not all variables after the annotation.

**Lemma 5.6.** *Let  $Q(\vec{x}, \star, \vec{z})$  be a full, acyclic, self-join-free  $\text{CQ}^\star$  with no disruptive trio such that all variables in  $\vec{z}$  neighbor each other. If direct access for  $Q$  over  $\mathbb{K}$  databases is in  $\langle \text{loglinear}, \text{log} \rangle$ , then every variable  $x \in \vec{x}$  that neighbors a single variable in  $\vec{z}$  neighbors all variables in  $\vec{z}$ , assuming the SparseBMM hypothesis.*

*Proof.* We will prove the contrapositive of the lemma. We assume that there are variables  $x \in \vec{x}$  and  $z_i, z_j \in \vec{z}$  such that  $x$  neighbors  $z_j$  but not  $z_i$ , and show that direct access for  $Q$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ . Let  $Q'$  be the  $\text{CQ}$   $Q'(\vec{x}, z_i, \vec{z})$  with the same body as  $Q$ . Note that  $z_i$  appears twice in the head of  $Q'$ , and therefore the second appearance does not affect the order. From Theorem 3.1, since  $Q$  is acyclic and  $x, z_i$  and  $z_j$  form a disruptive trio, we know that direct access for  $Q'(\vec{x}, z_i, \vec{z})$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming the SparseBMM hypothesis. Therefore, if we show that we can use direct access for  $Q$  to obtain direct access for  $Q'$ , then we get that direct access for  $Q$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ .

Given a database  $D'$  for  $Q'$ , we show how to construct a  $\mathbb{K}$ -database  $(D, \tau)$  for  $Q$ . We begin by performing the Yannakakis full reduction to eliminate dangling tuples [Yan81]. We can then assume that every fact  $f$  of  $D$  is consistent with  $Q$  in the sense that we can obtain  $f$  from the corresponding atom of  $Q$  by assigning some values to the variables; we then denote by  $f[x]$  the value of  $f$  in the positions of the variable  $x$ . We then construct  $D$  by first adding every fact in  $D'$  to  $D$ . Then we choose a single arbitrary relation in  $D$  that corresponds to an atom that contains  $z_i$ , and for every fact  $f$  in the relation, we define  $\tau(f) = f[z_i]$ . For every other fact  $f$  in  $D$  we define  $\tau(f) = \bar{1}$ .

Since the bodies of  $Q'$  and  $Q$  are identical, there is a bijection from  $Q'(D')$  to  $Q(D, \tau)$  mapping  $f'$  to  $f$  such that  $f[w] = f'[w]$  for every variable, where  $\tau(f) = f[z_i]$ . Therefore, ordering  $Q(D, \tau)$  by  $\vec{x}, \star, \vec{z}$  is the same as ordering  $Q'(D')$  by  $\vec{x}, z_i, \vec{z}$ . For any database  $D'$ , we can obtain direct access for  $Q'(D')$  by producing the equivalent database  $(D, \tau)$  and then providing direct access for  $Q(D, \tau)$ .  $\square$

We will use the following lemma, often seen as folklore [CTG<sup>+</sup>23, Lemma 5.4] [CKKK21, Proposition 5.3]. For completeness, we offer a proof suitable for our formalization.

**Lemma 5.7.** *Let  $Q$  be an acyclic CQ and let  $X$  be a set of variables of  $Q$ . If all pairs of variables of  $X$  are neighbors in  $Q$ , then there exists an atom  $\varphi$  in  $Q$  such that  $X \subseteq \text{vars}(\varphi)$ .*

*Proof.* Let  $T$  be a join tree of  $Q$ . Due to the running-intersection property, for each variable  $x_i$  in  $X$ , the induced subgraph of vertices containing  $x_i$  is a (connected) subtree, and let us denote it by  $T_{x_i}$ . It is known that every set of subtrees of  $T$  fulfills the *Helly property* [Gol04]: if the intersection of every two subtrees in the set is non-empty, then the intersection of all subtrees in the set is non-empty. Since every two variables  $x_i$  and  $x_j$  in  $X$  are neighbors, the subtrees  $T_{x_i}$  and  $T_{x_j}$  share some node of  $T$ . Therefore, the intersection of all subtrees contains some node of the join tree. That means that there is an atom which induced that node, whose variables contain  $X$ .  $\square$

We can now prove the negative side of the dichotomy for the case of free-connex CQ<sup>\*</sup>s with no disruptive trios.

**Lemma 5.8.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be the counting, numerical, max tropical, or min tropical semirings, and  $Q(\vec{x}, \star, \vec{z})$  a free-connex, self-join-free CQ<sup>\*</sup> with no disruptive trio. If direct access for  $Q$  is in  $\langle \text{loglinear}, \text{log} \rangle$ , then there exists an atom of  $Q$  that contains  $\mathbf{N}(\vec{z}) \cap \text{free}(Q)$ , assuming the 3SUM and SparseBMM hypotheses.*

*Proof.* Since direct access for  $Q$  is in  $\langle \text{loglinear}, \text{log} \rangle$ , by Theorem 3.4 we know that direct access for  $Q|_{\text{free}(Q)}$  is in  $\langle \text{loglinear}, \text{log} \rangle$ . Since  $Q|_{\text{free}(Q)}$  contains all free variables of  $Q$  and every atom of  $Q|_{\text{free}(Q)}$  is contained in an atom of  $Q$ , it is enough to show that there exists an atom in  $Q|_{\text{free}(Q)}$  that contains  $\mathbf{N}(\vec{z})$ . Let us consider  $Q|_{\text{free}(Q)}$  from now on. As direct access for  $Q|_{\text{free}(Q)}$  is in  $\langle \text{loglinear}, \text{log} \rangle$ , by Lemma 5.5 we know that all variables in  $\vec{z}$  neighbor each other. By Lemma 5.6, we also know that every variable in  $\vec{x}$  that neighbors a single variable in  $\vec{z}$  neighbors all variables in  $\vec{z}$ .

Let  $x_1, x_2 \in \vec{x}$  be two neighbors of  $\vec{z}$ . We claim that  $x_1$  neighbors  $x_2$ . Let  $z_1$  be an arbitrary variable in  $\vec{z}$ . As  $x_1$  and  $x_2$  are in  $\mathbf{N}(\vec{z})$  we know that both  $x_1$  and  $x_2$  neighbor all variables in  $\vec{z}$  and in particular  $z_1$ . Both  $x_1$  and  $x_2$  neighbor  $z_1$ , and  $z_1$  appears after  $x_1$  and  $x_2$  in the order of  $Q|_{\text{free}(Q)}$ . Since we know there are no disruptive trios, it must mean that  $x_1$  neighbors  $x_2$ .

In conclusion, let  $y_1, y_2$  be two variables in  $\mathbf{N}(\vec{z})$ . If  $y_1, y_2$  are in  $\vec{z}$ , they neighbor each other as we concluded from Lemma 5.5. If  $y_1 \in \vec{z}$  and  $y_2 \in \vec{x}$ , then they neighbor each other as we concluded from Lemma 5.6. If  $y_1, y_2$  are in  $\vec{x}$ , they neighbor each other as we showed in the previous paragraph. Overall, every two variables in  $\mathbf{N}(\vec{z})$  are neighbors, and from Lemma 5.7 we know that there exists an atom in  $Q|_{\text{free}(Q)}$  that contains  $\mathbf{N}(\vec{z})$ .  $\square$

It is left to show the hardness of CQ<sup>\*</sup>s that are not free-connex or that contain disruptive trios, which easily follows from the hardness of CQs without annotations.

**Lemma 5.9.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a commutative semiring, and  $Q(\vec{x}, \star, \vec{z})$  a self-join-free CQ $^*$ . If  $Q$  is not free-connex or has a disruptive trio, then direct access for  $Q$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming the HYPERCLIQUE and SparseBMM hypotheses.*

*Proof.* Consider the CQ  $Q'(\vec{x}, \vec{z})$  with the same body as  $Q$ . According to Theorem 3.1, direct access for  $Q'$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming the HYPERCLIQUE and SparseBMM hypotheses. We can construct an input for  $Q$  given an input for  $Q'$  by using the same database and annotating all facts with  $\bar{0}$ . Due to the semiring properties, we have that  $\bar{0} \oplus \bar{0} = \bar{0}$  and  $\bar{0} \otimes \bar{0} = \bar{0}$  (since  $a \oplus \bar{0} = a$  and  $a \otimes \bar{0} = \bar{0}$  for all  $a \in \mathbb{K}$ ), so  $Q$  will have the same answers as  $Q'$ , all annotated by  $\bar{0}$ , and therefore also in the same order. This way, we can use any direct access algorithm for  $Q$  to get direct access for  $Q'$  with the same time guarantees. As direct access for  $Q'$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ , neither is direct access for  $Q$ .  $\square$

Combining Lemma 5.3, Lemma 5.8, and Lemma 5.9 completes the proof of Theorem 5.1.

**5.3. Discussion on AggCQs.** Similarly to Corollary 4.2, we conclude from the tractability side of Theorem 5.1 the following corollary for AggCQs.

**Corollary 5.10.** *Let  $Q(\vec{x}, \alpha(\vec{w}), \vec{z}) :- \varphi_1(\vec{x}, \vec{y}), \dots, \varphi_\ell(\vec{x}, \vec{y})$  be a free-connex AggCQ with no disruptive trio. Suppose that  $\alpha$  is one of Min, Max, Count, and Sum. If  $Q$  has an atom that contains  $N(\vec{z}) \cap \text{free}(Q)$ , then direct access for  $Q$  is in  $\langle \text{loglinear}, \text{log} \rangle$ .*

Note that, contrasting Corollary 4.2, the aggregation Avg is not included in Corollary 5.10, since the argument that we can support Avg by solving Sum and Count separately no longer holds when the aggregation is part of the order.

In addition, we cannot similarly deduce the hardness of AggCQs based on the negative side of Theorem 5.1. The proof of Theorem 5.1 is based on Lemma 5.4 that states the hardness of direct access for  $R \times S$  by the order of the annotation. For that, we needed to use the power of the annotation, namely, the annotation of an answer  $(a, b)$  is the product of the annotations of  $R(a)$  and  $S(b)$ . This does not necessarily imply that we have a similar hardness when the computed value is within an AggCQ, say using Count. For example, direct access for  $Q(\text{Count}(), x, y) :- R(x), S(y)$  is clearly in  $\langle \text{loglinear}, \text{log} \rangle$  since the computed value has no impact (as it is always 1).

Nevertheless, we can show that incorporating the computed value in the order introduces hardness for another fixed AggCQ  $Q_c(\text{Count}(), x, y)$ . Moreover, this AggCQ is tractable if the order was  $x, y, \text{Count}()$ , due to Theorem 4.1. We prove the hardness of  $Q_c(\text{Count}(), x, y)$  using a reduction from  $Q_{\star \times}$  (of (5.1)) with the counting semiring  $(\mathbb{N}, +, \cdot, 0, 1)$ .

**Theorem 5.11.** *There exists a free-connex AggCQ  $Q_c(\text{Count}(), x, y)$  such that direct access for  $Q$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming the 3SUM conjecture.*

*Proof.* We show a reduction from the CQ $^*$   $Q_{\star \times}(\star, x, y)$  of Equation (5.1) over annotated databases  $(D, \tau)$  with the semiring  $(\mathbb{N}, +, \cdot, 0, 1)$ . Lemma 5.4 states that, assuming 3SUM,  $Q_{\star \times}$  is not in  $\langle \text{loglinear}, \text{log} \rangle$  even under the assumption that, for  $N = |D|$ , all annotations are within the domain  $[1, N^{d+1} - 1]$  for some fixed degree  $d$ . Let  $(D, \tau)$  be such an input database for  $Q_{\star \times}$ .

We construct an acyclic AggCQ  $Q_R(x, \text{Count}()) :- \varphi(x, \vec{w}_1)$  over a new schema, and show how to construct in linear time a new database  $D_R$  for  $Q_R$  such that for every fact  $R(a) \in D$  it holds that  $(a, n) \in Q_R(D_R)$  if and only if  $\tau(R(a)) = n$ . We similarly define

$R'$					$T'$		$T_1$		$T_2$		$T_3$	
$x$	$u$	$z_1$	$z_2$	$z_3$	$u$	$v$	$z_1$	$w_1$	$z_2$	$w_2$	$z_3$	$w_3$
2	2	1	1	1	1	1	1	1	1	1	1	1
13	1	1	1	1	2	1	4	1	4	1	4	1
13	1	4	1	1	2	2	4	2	4	2	4	2
13	2	4	1	1			4	3	4	3	4	3
64	1	4	4	4			4	4	4	4	4	4
192	1	4	4	4								
192	2	4	4	4								

$$Q_R(x, \text{Count}()) :- R'(x, u, z_1, \dots, z_d), T'(u, v), T_1(z_1, w_1), T_2(z_2, w_2), T_3(z_3, w_3)$$

Figure 2: Example of  $Q_R$  and  $D_R$  when  $N = 4$  and  $d = 3$ . The input  $(D, \tau)$  is defined by  $D = \{R(2), R(13), R(64), R(192)\}$  where  $\tau(R(a)) = a$  for every fact  $R(a)$ .

$Q_S(y, \text{Count}()) :- \psi(y, \vec{w}_2)$  and  $D_S$  for the relation  $S$ , so that  $Q_R$  and  $Q_S$  use disjoint sets of relation symbols and variables. Then, the AggCQ  $Q_c$  of the theorem is given by

$$Q_c(\text{Count}(), x, y) :- \varphi(x, \vec{w}_1), \psi(y, \vec{w}_2).$$

The database that we construct in the reduction is  $D_R \cup D_S$ . In the remainder of the proof, we show how  $Q_R$  and  $D_R$  are constructed. The construction of  $Q_S$  and  $D_S$  is analogous (using a fresh set of variables).

The AggCQ  $Q_R$  is the following:

$$Q_R(x, \text{Count}()) :- R'(x, u, z_1, \dots, z_d), T'(u, v), T_1(z_1, w_1), \dots, T_d(z_d, w_d)$$

For ease of explanation, we first handle the case of  $d = 0$ . In this case, the annotations are in  $\{1, \dots, N - 1\}$  and  $Q_R$  is  $Q_R(x, \text{Count}()) :- R'(x, u), T'(u, v)$ . Consider a fact  $R(a)$  of  $D$  and its annotation  $\tau(R(a))$ . As the annotation is in the range  $\{1, \dots, N - 1\}$ , it can be represented in binary form using  $\log N$  bits. Let  $I \subseteq \{0, \dots, \log(N) - 1\}$  be the set of positions of 1-bits in the binary representation of the annotation  $\tau(R(a))$ . We have that  $\tau(R(a)) = \sum_{i \in I} 2^i$ . For every  $i \in I$ , we put the fact  $R'(a, i)$  in  $D_R$ . We also make sure that there are  $2^i$  facts where variable  $u$  is assigned value  $i$  in  $T'$ . More precisely, for every  $i \in \{0, \dots, \log(N) - 1\}$  and every  $j \in \{1, \dots, 2^i\}$ , we add to  $D_R$  the fact  $T'(i, j)$ . This construction ensures that the answers to  $Q_R$  are exactly the answers to  $Q$  where the count is equal to the annotation.

We now define the construction of  $Q_R$  and  $D_R$  for an arbitrary  $d \geq 0$ . Consider a fact  $R(a)$  and its annotation  $\tau(R(a))$ . As the annotation is in the range  $\{1, \dots, N^{d+1} - 1\}$ , it can be represented in base  $N$  using  $d + 1$  digits. Each digit in base  $N$  can be represented in binary using  $\log N$  bits. Let  $p_j$  be the  $j$ th digit in the base- $N$  representation of  $\tau(R(a))$  and let  $I_j \subseteq \{0, \dots, \log(N) - 1\}$  be the set of positions of 1-bits in the binary representation of  $p_j$ . For every  $i \in I_j$  we add to  $D_R$  the following fact:

$$R'(a, i, \underbrace{N, \dots, N}_j, \underbrace{1, \dots, 1}_{d-j})$$

For every  $i \in 1, \dots, d$  we make sure there are  $N$  tuples starting with  $N$  and one tuple starting with 1 in  $T_i^{D_R}$ . For that, we add to  $D_R$  the fact  $T_i(1, 1)$ , and for every  $q = 1, \dots, N$

we add the fact  $T_i(N, q)$ . We create  $T'$  as in the case when  $d = 0$ . Note that the schema of  $Q_R$  depends only on  $d$ . Also note that the facts of  $T'$  and  $T_i$  depend only on  $N$ , regardless of the facts of  $D$  or their annotations. Finally, the reader can verify that this construction ensures that  $(a, n) \in Q_R(D_R)$  if and only if  $\tau(R(a)) = n$ . (See also the following Example 5.12 for an illustration.)

Finally, let us inspect the size of  $D_R$ . Every fact of  $D$  is used at most  $d \log N$  times in  $R'^{D_R}$ . We also have that  $|T'^{D_R}| = \sum_{i=0}^{\log N - 1} 2^i = N - 1$ , and  $|T_i^{D_R}| = N + 1$ . Hence, we add to  $D_R$  a total of  $O(d \log(N) + N)$  facts.  $\square$

The following example illustrates the construction used in the proof of Theorem 5.11.

**Example 5.12.** Figure 2 depicts an example of the construction for the database

$$D = \{R(2), R(13), R(64), R(192)\}$$

where  $\tau(R(a)) = a$  for every fact  $R(a)$ . In the example, we have  $N = |R^D| = 4$  and  $d = 3$ , and thus the construction supports annotations in  $[1, \dots, 4^4 - 1 = 255]$ . We illustrate our construction by showing that for every fact  $R(a)$  with the annotation  $a$ , we will get the answer  $(a, a)$  in  $Q_R(D_R)$ .

- The fact  $R(2)$  has the annotation 2, which can be represented in binary using a single 1-bit, so we add to  $D_R$  the single fact  $R'(2, 2, 1, 1, 1)$  and make sure there are two facts in  $T'^{D_R}$  starting with 2. So  $(2, 2)$  will be in  $Q_R(D_R)$ .
- For  $R(13)$ , the annotation 13 can be represented by the polynomial  $3 \cdot 4 + 1$ . Therefore, we transform  $R(13)$  to three facts in  $R'^{D_R}$ , one with a count of  $1 \cdot 1 \cdot 1 \cdot 1$ , one with a count of  $1 \cdot 4 \cdot 1 \cdot 1$ , and one with a count of  $2 \cdot 4 \cdot 1 \cdot 1$ . As a result, the tuple  $(13, 13)$  will be in  $Q_R(D_R)$  as  $1 + 4 + 8 = 13$ .
- As for  $R(64)$ , its annotation of 64 can be represented as  $1 \cdot 4^3$ . As 1 contains a single active bit, we transform it to the fact  $R'(1, 4, 4, 4)$  which has a count of  $1 \cdot 4 \cdot 4 \cdot 4 = 64$ .
- Now consider the fact  $R(192)$  with the annotation 192. We can represent 192 as a polynomial of  $N = 4$  as  $3 \cdot 4^3$  and 3 can be represented in binary using two 1-bits. So, the fact  $R(192)$  is transformed into the facts  $R'(192, 1, 4, 4, 4)$  and  $R'(192, 2, 4, 4, 4)$  in  $D_R$ . The former of these facts has a count of  $1 \cdot 4 \cdot 4 \cdot 4 = 64$  and the latter has a count of  $2 \cdot 4 \cdot 4 \cdot 4 = 128$ , so the  $Q_R(D_R)$  will contain the tuple  $(192, 192)$  since  $64 + 128 = 192$ .

Hence, we get the answer  $(a, a)$  for each  $R(a)$ , as we claimed.  $\square$

Lemma 5.4 states hardness of  $Q_{\star \times}(\star, x, y) :- R(x), S(y)$ , while in Theorem 5.11 we only claimed the existence of the hard AggCQ  $Q_c(\text{Count}(), x, y)$ , which is more involved than  $Q_{\star \times}(\star, x, y)$  (and we construct it in the proof). The reader might wonder whether we could also phrase Lemma 5.4 over the Cartesian product  $Q(\text{Count}(), x, y) :- R(x), S(y)$  or alike. Clearly, incorporating  $\text{Count}$  in  $R(x) \times S(y)$  would be meaningless since every answer appears exactly once (and has a count of 1). The next theorem shows that the reason goes deeper: even if we add to  $Q(\text{Count}(), x, y) :- R(x), S(y)$  existential variables, the query remains in  $\langle \text{loglinear}, \text{log} \rangle$ .

**Proposition 5.13.** *For  $Q(\text{Count}(), x, y) :- R(x, w), S(y, z)$ , direct access is in  $\langle \text{loglinear}, \text{log} \rangle$ .*

*Proof.* For ease of following the proof, we rename  $Q$  as follows so that the two atoms have matching structures:

$$Q(\text{Count}(), x, x') :- R(x, w), R'(x', w')$$

$R$	$R'$		$L$
$x$	$w$		$c$
$a$	$1$		$1$
$b$	$1$		$1$
$b$	$2$	$\longrightarrow$	$2$
$c$	$1$		$3$
$c$	$2$		$2$
$c$	$3$		$3$
$a'$	$1$		$[a]$
$b'$	$1$		$[a]$
$c'$	$1$		$[b]$
$c'$	$2$		$[c]$
$d'$	$1$		$[b]$
$d'$	$2$		$[c]$
$[a', b']$	$[a', b']$		$[a', b']$
$[c', d']$	$[c', d']$		$[c', d']$
$[a', b']$	$[a', b']$		$[a', b']$
$[c', d']$	$[c', d']$		$[c', d']$
$[c', d']$	$[c', d']$		$[c', d']$

$$Q(\text{Count}(), x, x') :- R(x, w), R'(x', w')$$

Figure 3: Example of the construction in the proof of Proposition 5.13: direct access for the AggCQ  $Q(\text{Count}(), x, x') :- R(x, w), R'(x', w')$ .

In the remainder of this proof, we fix an input database  $D$  for  $Q$ . Note that the answers are of the form  $(c \cdot c', a, a')$  where  $c$  is the count of the facts  $R(a, \cdot)$  that have  $a$  as the first element, and  $c'$  is the count of the facts  $R'(a', \cdot)$  that have  $a'$  as the first element.

Let us describe the preprocessing step. First, we compute the number of facts  $R(a, \cdot)$  for every possible value of  $a$ . We use the result to create the set of all counts per possible value of  $x$ , and denote this set by  $C$ . For all  $c \in C$ , we also keep a list  $X_c$  with the set of all values  $a$  that appear  $c$  times in the left column of  $R$ . The list  $X_c$  is sorted so that we can easily locate a given  $a$ . We do the same for  $R'$  to obtain  $C'$  and a sorted list  $X'_{c'}$  for every  $c' \in C'$ .

Next, we create a list  $L$ , where for every pair of counts  $(c, c') \in C \times C'$  it holds the tuple  $(c, c', X_c, X'_{c'})$  where  $X_c$  and  $X'_{c'}$  are represented as *pointers* to the corresponding lists. See Figure 3 for an example of the construction of  $L$  from an input database. We perform direct access on  $L$  sorted by  $c \cdot c'$  and weighted by  $|X_c| \cdot |X'_{c'}|$  using prefix sum, similarly to the way established by Carmeli et al. [CTG<sup>+</sup>23] for a single relation, as we briefly describe next.

We first sort  $L$  by the product of the first two elements of each tuple,  $c \cdot c'$ . Notice that  $L$  indeed represents the answers in the order we desire. For example, if the first fact in  $L$  is  $(c, c', X_c, X'_{c'})$ , then the first  $|X_c| \cdot |X'_{c'}|$  answers have the count  $c \cdot c'$  and assign to  $x$  and  $x'$  the values that occur in  $X_c$  and  $X'_{c'}$ , respectively. We then iterate over  $L$ , and for a tuple index  $i$  we compute the sum of  $|X_c| \cdot |X'_{c'}|$  over all tuples in indices  $1, \dots, i-1$  in  $L$ . We denote this sum by  $l_i$ . Note that  $l_{i+1} > l_i$  for all  $i = 1, \dots, |L|$ .

We now describe the access procedure. Suppose that we are requested to fetch result number  $d$ . We perform a binary search on  $L$  to find the tuple in index  $i$  such that  $l_i < d \leq l_{i+1}$ . Assume that this tuple is  $(c, c', X_c, X'_{c'})$ . The count we return is  $c \cdot c'$ . Next, we need to access the  $(d - l_i)$ th element  $(x, x')$  in  $X_c \times X'_{c'}$ , sorted lexicographically by  $x$  and then by  $x'$ . As in multidimensional arrays, we assign to  $x'$  the element with the index  $(d - l_i) \bmod |X'_{c'}|$  of  $X'_{c'}$ , and we assign to  $x$  the element with the index  $\lfloor \frac{d - l_i}{|X'_{c'}|} \rfloor$  of  $X_c$ .

We now analyze the execution time. Processing each of the  $a$  counts and the  $a'$  counts separately requires only  $O(|D|)$  time. The concern is the time it takes to build and sort the list  $L$ , which might be of size  $|C| \cdot |C'|$ . Assume, without loss of generality, that  $|C| \geq |C'|$ . We claim that  $|R^D| = \Omega(|C|^2)$ . If  $R^D$  has the smallest possible number of facts to result in  $|C|$  distinct counts, then the counts are  $1, \dots, |C|$ . The number of facts in this case is  $\sum_{i=1}^{|C|} i = \frac{|C|(|C|+1)}{2}$ . So,  $|L| = |C| \cdot |C'| \leq |C|^2 = O(|R^D|)$ . We conclude that the algorithm runs in  $O(|D| \log |D|)$  preprocessing time and  $O(\log |D|)$  access time. In conclusion, direct access for  $Q$  is in  $\langle \text{loglinear}, \log \rangle$ , as claimed.  $\square$

In the next section, we study how additional assumptions on the annotated database, useful for AggCQs, can lead to additional opportunities to efficiently incorporate the computed value in the ordering.

## 6. LOCALLY ANNOTATED DATABASES

We have seen in Lemma 5.4 that even the simple CQ $^*$   $Q_{\star \times}(\star, x, y) :- R(x), S(y)$  is intractable, and we have seen in Section 5.3 that this hardness does not necessarily apply to AggCQs. As another example, consider the AggCQ

$$Q(\text{Sum}(w), x, y) :- R(x, w), S(y). \quad (6.1)$$

When translating into an annotated database, we obtain the CQ $^*$   $Q_{\star \times}$  over  $\mathbb{Q}$ -databases annotated by the numerical semiring  $(\mathbb{Q}, +, \cdot, 0, 1)$ . Hence, we translate the problem into an intractable one. Nevertheless, direct access for  $Q$  by  $(\star, x, y)$  is, in fact, in  $\langle \text{loglinear}, \log \rangle$ , as we will show in Theorem 6.4. This discrepancy stems from the fact that the hardness of  $Q_{\star \times}$  (established in the proof of Lemma 5.4) relies on the annotation of tuples from both  $R$  and  $S$ . Yet, in our translation, all  $S$ -facts are annotated by 1, and only  $R$ -facts have a nontrivial annotation. The resulting  $\mathbb{K}$ -database is such that every fact is annotated by  $\bar{1}$  (the multiplicative identity), with the exception of one relation. We call such a  $\mathbb{K}$ -database *locally annotated* or *R-annotated* (when we need to specify  $R$ ). We now focus on such databases and show how the assumption of local annotations can be used for efficient access.

In the remainder of this section, we restrict the discussion to queries without self-joins<sup>2</sup> and fix a logarithmic-time commutative semiring  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$ .

**6.1. Full CQ $^*$ s.** We first discuss full CQ $^*$ s (i.e., without existential variables). In the next sections, we also discuss the implications on (non-full) AggCQs. We begin with some examples that demonstrate the results that follow later in this section.

**Example 6.1.** In some cases, incorporating the annotation in an otherwise tractable order may introduce hardness. Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a logarithmic-time commutative semiring, and consider the full CQ $^*$

$$Q(\vec{x}, \star, \vec{z}) :- R(v_1, v_3), S(v_2, v_3)$$

over  $S$ -annotated  $\mathbb{K}$ -databases. Note that  $Q$  has a disruptive trio if  $v_3$  appears after both  $v_1$  and  $v_2$  in the order. Let us consider orders where this is not the case. The lexicographic

<sup>2</sup>In fact, for the algorithm, we only need that at most one atom refers to a relation with unrestricted annotations, so it suffices that the relation with unrestricted annotations is not part of more than one atom. The lower bounds do rely on the assumption that every refers to a distinct relation, just like the lower bound of Theorem 3.1 and all the other lower bounds in this paper.

order  $(\star, v_2, v_3, v_1)$  is not covered by Theorem 5.1, as  $v_1$  appears only in  $R$ , but  $v_2$  does not. However, we will show in Theorem 6.4 that, when considering  $S$ -annotated  $\mathbb{K}$ -databases, direct access for  $Q$  by this order is in  $\langle \text{loglinear}, \text{log} \rangle$ , while direct access for  $Q$  by  $(\star, v_1, v_3, v_2)$  is *not* in  $\langle \text{loglinear}, \text{log} \rangle$ .  $\square$

**Example 6.2.** It may also happen that we incorporate the annotation non-trivially and the query remains tractable. Consider the full CQ $^*$

$$Q'(\vec{x}, \star, \vec{z}) :- R(v_1, v_3), S(v_2, v_3), T(v_3)$$

over  $T$ -annotated  $\mathbb{K}$ -databases. Note that this case is a slight variation on Example 6.1. For any lexicographic order  $(\vec{x}, \star, \vec{z})$ , if  $v_3$  appears after  $v_1$  and  $v_2$ , then  $Q'$  has a disruptive trio and, therefore, direct access for  $Q'$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ . As we show in Theorem 6.4, for  $Q'$  over  $T$ -annotated  $\mathbb{K}$ -databases, the lack of a disruptive trio is a sufficient condition for tractability. That is, if  $v_3$  does not appear after  $v_1$  and  $v_2$ , then it holds that direct access for  $Q'$  (and for  $Q$  of Example 6.1) is in  $\langle \text{loglinear}, \text{log} \rangle$ .  $\square$

When dealing with  $R$ -annotated databases, we can replace the annotation of the facts of  $R$  with a new extra attribute, added to  $R$ , and then reason about orders that involve the annotation by considering orders that involve the new attribute instead. To do so, we introduce the following variations of a query and order. Let  $Q$  be a full acyclic CQ $^*$  without self-joins. For a relation symbol  $R$  of  $Q$ , we define the  $R$ -deannotation of  $Q(\vec{x}, \star, \vec{z})$  to be the CQ  $Q_R$  obtained as follows, where we denote by  $\varphi_S$  the atom of a relation symbol  $S$ .

- In the head, replace  $\star$  with a new variable  $y$ .
- If, in addition,  $\text{vars}(\varphi_R)$  contains only variables from  $\vec{x}$ , then in the head of  $Q_R$ , advance  $y$  to be immediately after the last variable of  $\varphi_R$ .
- For each relation  $S$  of  $Q$ , if  $\text{vars}(\varphi_R) \subseteq \text{vars}(\varphi_S)$  then concatenate  $y$  to the variable sequence of  $\varphi_S$ .

**Example 6.3.** Consider the CQ $^*$   $Q'$  of Example 6.2. The  $R$ -deannotation of  $Q'$  is

$$Q_R(\vec{x}_R, y, \vec{z}_R) :- U(v_1, v_3, y), V(v_2, v_3, y), R(v_3, y)$$

where  $\vec{x}_R$  and  $\vec{z}_R$  are adjustments of  $\vec{x}$  and  $\vec{z}$ : if  $v_3$  is in  $\vec{x}$ , the suffix of  $\vec{x}$  that follows  $v_3$  is moved to the beginning of  $\vec{z}$ .  $\square$

When  $Q$  is full, we can reduce direct access for  $Q$  to direct access for  $Q_R$ . This is stated in the next theorem. The theorem also says that, whenever the annotation domain contains the natural numbers, this reduction is optimal in the sense that, if we got an intractable  $Q_R$ , then direct access for  $Q$  was hard to begin with.

**Theorem 6.4.** Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a logarithmic-time commutative semiring. Let  $Q$  be a full CQ $^*$  without self-joins and  $Q_R$  the  $R$ -deannotation of  $Q$  for a relation symbol  $R$  of  $Q$ .

- (1) If  $Q_R$  is acyclic and has no disruptive trio, then direct access for  $Q$  is in  $\langle \text{loglinear}, \text{log} \rangle$  on  $R$ -annotated  $\mathbb{K}$ -databases.
- (2) Otherwise, if  $\mathbb{N} \subseteq \mathbb{K}$ , then direct access for  $Q$  is not in  $\langle \text{loglinear}, \text{log} \rangle$  on  $R$ -annotated  $\mathbb{K}$ -databases, assuming the HYPERCLIQUE and SparseBMM hypotheses.

We note that the negative side of Theorem 6.4 applies to any domain other than  $\mathbb{N}$ , as long as we can generate infinitely many elements according to the underlying order of the semiring. Before we prove the theorem, let us give a usage example.

**Example 6.5.** Theorem 6.4 gives us a useful tool to analyze the previous examples. Recall that in Example 6.3 we showed the  $R$ -deannotation of  $Q'$  from Example 6.2. Since  $y$  appears in every atom, it cannot be part of any disruptive trio. In particular, the disruptive trios of  $Q_R$  are exactly the disruptive trios of  $Q$ . Therefore, given an order  $(\vec{x}, \star, \vec{z})$ , checking whether direct access for  $Q$  is in  $\langle \text{loglinear}, \text{log} \rangle$  boils down to verifying that  $v_1, v_2$  and  $v_3$  do not form a disruptive trio.  $\square$

To prove Theorem 6.4, we first prove the following lemma.

**Lemma 6.6.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a logarithmic-time commutative semiring. Let  $Q$  be a full CQ $^*$  without self-joins and  $Q_R$  the  $R$ -deannotation of  $Q$  for a relation symbol  $R$  of  $Q$ . If  $Q$  is acyclic, then  $Q_R$  is acyclic.*

*Proof.* Since  $Q$  is acyclic it has a join tree  $T$ . We add the variable  $y$  to every vertex of  $T$  that contains all variables of  $R$ . As per the definition of  $R$ -deannotation each vertex in the tree now corresponds to an atom of  $Q_R$ . We know that every variable in  $T$ , besides  $y$ , satisfies the running-intersection property. This is true, in particular, for every variable in  $R$ . Since the intersection of subtrees is in itself a subtree, we get that  $y$  is added to a connected set of vertices, and it too satisfies the running-intersection property. Hence, the new tree is a join tree of  $Q_R$ , and so,  $Q_R$  is acyclic.  $\square$

With Lemma 6.6 at hand, we can now prove Theorem 6.4

*Proof.* We prove each part separately.

Part 1. First, we show that if  $Q_R$  is acyclic and has no disruptive trio, then direct access for  $Q$  on  $R$ -annotated  $\mathbb{K}$ -databases is in  $\langle \text{loglinear}, \text{log} \rangle$ . Let  $(D, \tau)$  be an  $R$ -annotated  $\mathbb{K}$ -database. Since  $Q$  is full, the answer  $Q(D, \tau)$  is simply the join of the relations of all atoms in  $Q$ . Since  $D$  is also  $R$ -annotated, the annotation of an answer is the annotation of its supporting fact in  $R^D$ . Yet, the annotation of an  $R$ -fact is functionally dependent on the other attributes. We will use a recent result on direct access in the presence of Functional Dependencies (FDs) [CTG<sup>+</sup>23].

We first construct from the CQ $^*$   $Q$  and the annotated database  $(D, \tau)$  a CQ  $Q_{\text{fd}}$  and a database  $D_{\text{fd}}$ , as follows:

- Introduce a relation symbol  $R_e$  of arity  $\text{ar}(R) + 1$ .
- Modify the atom containing  $R$  in  $Q$  by replacing  $R$  with  $R_e$  and concatenating a new variable  $y$ .
- Construct  $R_e^{D_{\text{fd}}}$  from  $R^D$  by setting the value of  $y$  in each fact  $f$  to be  $\tau(f)$ .
- In the head, replace  $\star$  with  $y$ .
- Add the FD  $R_e : 1, \dots, \text{ar}(R) \rightarrow \text{ar}(R) + 1$  to the schema of  $Q_{\text{fd}}$ . (Hence, here we allow the schema to include an integrity constraint.)

From the discussion above, answers in  $Q_{\text{fd}}(D_{\text{fd}})$  are identical to answers in  $Q(D, \tau)$ , except for the annotations that are now represented as values for  $y$ . Hence, if we prove that direct access for  $Q_{\text{fd}}$  is in  $\langle \text{loglinear}, \text{log} \rangle$ , then so is direct access for  $Q$ . We do so next.

We follow the algorithm suggested by Carmeli et al.[CTG<sup>+</sup>23] to obtain what they refer to as the *FD-reordered extension* of  $Q_{\text{fd}}$ . By doing so, we get *precisely* our  $R$ -deannotation  $Q_R$  of  $Q$ . Carmeli et al. [CTG<sup>+</sup>23] proved that given a CQ and its FD-reordered extension, if direct access for the FD-reordered extension is in  $\langle \text{loglinear}, \text{log} \rangle$ , then so is direct access for the original CQ, which is  $Q_{\text{fd}}$  in our case. Therefore if direct access for  $Q_R$  is in  $\langle \text{loglinear}, \text{log} \rangle$ , then so is direct access for  $Q_{\text{fd}}$ , and so is direct access for  $Q$  on  $R$ -annotated

$\mathbb{K}$ -databases. Since  $Q_R$  is acyclic and has no disruptive trio, by Theorem 3.1 direct access for  $Q_R$  is, indeed, in  $\langle \text{loglinear}, \text{log} \rangle$ .

Part 2. Next, we assume that  $\mathbb{N} \subseteq \mathbb{K}$  and either  $Q_R$  is cyclic or  $Q_R$  has a disruptive trio. We analyze all possible cases and show that in every case direct access for  $Q$  on  $R$ -annotated  $\mathbb{K}$ -databases is not in  $\langle \text{loglinear}, \text{log} \rangle$  (under the corresponding complexity assumptions). We denote by  $X_R$  the set of variables that occur in the atom of  $R$  in  $Q$ .

**Case 1:**  $Q_R$  is cyclic. In Lemma 6.6 we proved that if  $Q$  is acyclic, so is  $Q_R$ . Therefore, if  $Q_R$  is cyclic, so is  $Q$ , and by Theorem 4.1 direct access for  $Q$  on  $R$ -annotated  $\mathbb{K}$ -databases is not in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming the HYPERCLIQUE hypothesis. In the remaining cases, we assume that  $Q_R$  is acyclic.

**Case 2:** *There exists a disruptive trio that does not involve  $y$ .* Next, we assume that  $x_1$ ,  $x_2$ , and  $x_3$  form a disruptive trio in  $Q_R$ , and  $y$  is not one of  $x_1, x_2, x_3$ . We create a new CQ  $Q_0$  by removing  $y$  from  $Q_R$ . The variables  $x_1, x_2, x_3$  still form a disruptive trio in  $Q_0$ , and we know by Theorem 3.1 that direct access for  $Q_0$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming SparseBMM. Assume, by way of contradiction, that direct access for  $Q$  on  $R$ -annotated  $\mathbb{K}$ -databases is in  $\langle \text{loglinear}, \text{log} \rangle$ . The lexicographic order of the CQ\*  $Q$  differs from that of the CQ  $Q_0$  by the existence of  $\star$ ; other than that,  $Q_0$  is identical to  $Q$ . Therefore, we can provide direct access for  $Q_0$  on a given database  $D$  by annotating each fact in  $D$  with  $\bar{1}$ , and providing direct access for  $Q$  on the annotated database. We conclude that direct access for  $Q_0$  is in  $\langle \text{loglinear}, \text{log} \rangle$ , which is a contradiction.

**Case 3:** *There is a disruptive trio that involves  $y$ , and  $y$  is a neighbor of both other trio members.* Suppose that  $x_1$ ,  $x_2$ , and  $y$  form such a disruptive trio in  $Q_R$ . Then  $y$  is a neighbor of both  $x_1$  and  $x_2$ , the variables  $x_1$  and  $x_2$  are non-neighbors, and  $y$  appears after  $x_1$  and  $x_2$  in the order. Recall that  $X_R$  denotes the set of variables in the atom of  $R$ . Denote by  $x_f$  the variable in  $X_R$  that appears last in  $Q_R$ . We claim that  $x_1$ ,  $x_2$ , and  $x_f$  form a disruptive trio in  $Q_R$ . Since  $y$  was added only to atoms containing all variables in  $X_R$ , we have that  $x_f$  appears in all atoms containing  $y$ , and so, both  $x_1$  and  $x_2$  are neighbors of  $x_f$ . Since  $x_1$  and  $x_2$  are not neighbors of each other, we have that  $x_1 \neq x_f$  and  $x_2 \neq x_f$ . We know that  $x_f$  appears either after  $y$  or immediately before  $y$  in the head since this is how the head of  $Q_R$  was reordered. Since  $x_f$  appears in  $Q_R$  either after  $y$  or immediately before  $y$ , it necessarily appears after  $x_1$  and  $x_2$ . This proves that  $x_1$ ,  $x_2$ , and  $x_f$  form a disruptive trio in  $Q_R$ . As we saw in the previous case, a disruptive trio that does not include  $y$  means direct access for  $Q$  on  $R$ -annotated  $\mathbb{K}$ -databases is not in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming SparseBMM.

**Case 4:** *There is a disruptive trio that involves  $y$ , and  $y$  is a neighbor of only one of the other trio members.* Assume that  $x_2$  is the neighbor of both  $x_1$  and  $y$ , that  $x_1$  and  $y$  are not neighbors, and that  $x_2$  appears after  $x_1$  and  $y$  in the order.

We claim that since  $x_1$  and  $y$  are not neighbors, there exists a variable  $x' \in X_R$  that is not a neighbor of  $x_1$  in  $Q_R$ . Assume, by way of contradiction, that no such variable exists and so every variable of  $X_R$  is a neighbor of  $x_1$ . By Lemma 5.7, there exists an atom in  $Q_R$ , and by extension  $Q$  that contains  $\{x_1\} \cup X_R$ . Recall that we constructed  $Q_R$  as the  $R$ -deannotation of  $Q$  by adding  $y$  to every atom that contains  $X_R$ . Therefore,  $x_1$  and  $y$  occur jointly in the atom that contains  $\{x_1\} \cup X_R$  and are neighbors. We conclude that

since  $x_1$  and  $y$  are not neighbors, there exists a variable  $x' \in X_R$  such that  $x_1$  and  $x'$  are not neighbors.

As  $y$  appears only in atoms that contain  $X_R$ , and  $x'$  is a variable in  $X_R$ , we have that  $x'$  is in every atom that contains  $y$ . Therefore, since  $x_2$  is a neighbor of  $y$ , it is also a neighbor of  $x'$ . Overall, we have that  $x_2$  is a neighbor of both  $x'$  and  $x_1$ , but  $x_1$  is not a neighbor of  $x'$ . Assume first that  $x'$  appears before  $y$  in the head of  $Q_R$ . As  $y$  appears before  $x_2$ , we have that  $x'$  appears before  $x_2$ . Therefore, the variables  $x_1$ ,  $x'$ , and  $x_2$  form a disruptive trio with respect to  $Q_R$ . We know that if there exists a disruptive trio that does not involve  $y$ , then direct access for  $Q$  by on  $R$ -annotated  $\mathbb{K}$ -databases is not in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming **SparseBMM**. So now, we are left with the case where  $x'$  appears after  $y$ . In the remainder of the proof, we handle this case.

We create a new CQ  $Q_0$  from  $Q_R$  by removing  $y$  and moving  $x'$  to the position that  $y$  occupied in the head. As  $x_2$  appears in  $Q_R$  after  $x_1$  and  $y$ , it appears in  $Q_0$  after  $x_1$  and  $x'$ . Since  $x_2$  is a neighbor to both  $x_1$  and  $x'$ , and  $x'$  and  $x_1$  are not neighbors, it holds that  $x_1$ ,  $x'$ , and  $x_2$  form a disruptive trio in  $Q_0$ . From Theorem 3.1, we know that direct access for  $Q_0$  is not in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming **SparseBMM**. To conclude the proof, we show a reduction from direct access for  $Q_0$  to direct access for  $Q$  on  $R$ -annotated  $\mathbb{K}$ -databases. This would immediately imply that direct access for  $Q$  on  $R$ -annotated  $\mathbb{K}$ -databases is not in  $\langle \text{loglinear}, \text{log} \rangle$ , assuming **SparseBMM**.

Recall our assumption that the semiring domain includes  $\mathbb{N}$ . We take all values of  $x'$  in  $R^D$ , place them in an array, denoted by  $A$ , and sort them. We denote by  $h(f)$  the position of  $f$  in  $A$ . Then we add annotations to  $D$ . For every fact  $f$  in  $R^D$ , we define that  $\tau(f) = h(f)$ . For every fact  $f$  in  $D$  but not in  $R^D$ , we define  $\tau(f) = \bar{1}$ . Note that, with this construction, ordering by  $\star$  is the same as ordering by  $x'$ . As  $x'$  is in  $X_R$  and appears after  $\star$  in the order, we have that  $y$  is not advanced during the construction of  $Q_R$ , and so,  $\star$  appears in  $Q$  in the same position where  $y$  appears in  $Q_R$ . Since  $x'$  appears after  $y$  in  $Q_R$ , it also appears after  $\star$  in  $Q$ . Due to the way we annotate the facts, sorting the results first by  $\star$  and then by  $x'$  is the same as sorting only by  $\star$ . This means that, for all  $i$ , the  $i$ -th answer for  $Q_0(D)$  is the same as the  $i$ -th answer for  $Q(D, \tau)$ , where the value of the annotation in  $Q(D, \tau)$  is the position in  $A$  of the value assigned to  $x'$  in  $Q_0(D)$ , and the value for  $x'$  appears in the answer of  $Q(D, \tau)$  in a different position. When the  $i$ -th answer for  $Q_0(D)$  is requested, we fetch the  $i$ -th answer for  $Q(D, \tau)$  and replace the annotation with the value in the position of  $x'$ . Hence, we get the claimed reduction from direct access for  $Q_0$  to direct access for  $Q$  on  $R$ -annotated  $\mathbb{K}$ -databases, where the database construction takes  $O(|D| \log |D|)$  time, and the translation of an answer during access takes constant time.  $\square$

**6.2. Queries with projection in the case of idempotence.** Next, we extend Theorem 6.4 beyond full CQ $^*$ s. In this section, we will be working with queries without self-joins. For simplicity, we will often mention a relation symbol to refer to the atom using this relation symbol. Consider the following AggCQ:

$$Q(\text{Max}(w_2), x_1, x_2, x_3) :- R(x_1, x_3, w_3), S(x_2, x_3), T(x_3, w_1), U(w_1, w_2) \quad (6.2)$$

We can solve direct access for  $Q$  using direct access for the CQ $^*$

$$Q'(\star, x_1, x_2, x_3) :- R(x_1, x_3, w_3), S(x_2, x_3), T(x_3, w_1), U'(w_1)$$

over  $U'$ -annotated  $\mathbb{Q}$ -databases and the max tropical semiring. We can then eliminate existential variables as we did in Lemma 3.3 and reduce direct access for  $Q'$  to direct access

for the full CQ<sup>\*</sup>  $Q'_{|\text{free}(Q')}$ .

$$Q'_{|\text{free}(Q')}(\star, x_1, x_2, x_3) :- R_{|\text{free}(Q')}(x_1, x_3), S_{|\text{free}(Q')}(x_2, x_3), T_{|\text{free}(Q')}(x_3)$$

As we will show in Lemma 6.8, when the input database for  $Q'$  is  $U'$ -annotated over the max tropical semiring, the suitable database for  $Q'_{|\text{free}(Q')}$  is  $T_{|\text{free}(Q')}$ -annotated. From  $Q'_{|\text{free}(Q')}$  we define  $Q_0$  as the  $T_{|\text{free}(Q')}$ -deannotation of  $Q'_{|\text{free}(Q')}$ .

$$Q_0(y, x_1, x_2, x_3) :- R_0(x_1, x_3, y), S_0(x_2, x_3, y), T_0(x_3, y)$$

Since  $Q'_{|\text{free}(Q')}$  is full and  $Q_0$  is acyclic with no disruptive trios, Theorem 6.4 implies that  $Q'_{|\text{free}(Q')}$  is in  $(\text{loglinear}, \text{log})$  on  $T_{|\text{free}(Q')}$ -annotated  $\mathbb{Q}$ -databases. We therefore get that direct access for  $Q'$  on  $U$ -annotated  $\mathbb{Q}$ -databases is in  $(\text{loglinear}, \text{log})$ , and as a consequence so is direct access for  $Q$ .

As we explain next, the argument above is specific to **Max** and not all aggregate functions since we rely on **Max** being *idempotent*. An operation  $\oplus$  is said to be idempotent if for every  $a$  in the domain  $\mathbb{K}$  we have that  $a \oplus a = a$ . We say that a commutative semiring is an  $\oplus$ -*idempotent semiring* if its addition operation,  $\oplus$ , is idempotent.

Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be some logarithmic-time commutative semiring. The process of existential variable elimination as given in Lemma 3.3 takes a free-connex CQ<sup>\*</sup>  $Q$  without self-joins and a  $\mathbb{K}$ -database  $(D, \tau)$  and translates it to a full acyclic CQ<sup>\*</sup>  $Q_{|\text{free}(Q)}$  and a  $\mathbb{K}$ -database  $(D_{\text{free}}, \tau_{\text{free}})$ . When working over an  $\oplus$ -idempotent semiring, if the input database is locally annotated, then we can show that the output database is also guaranteed to be locally annotated. The semirings used for **CountD** (over logarithmic domain), **Min**, and **Max** are  $\oplus$ -idempotent and therefore provide such a guarantee, while the semirings used for **Sum** and **Count** do not. In the case of the query of (6.2), if the aggregate function was **Sum** instead of **Max**, once we eliminate existential variables, the database for  $Q_{|\text{free}(Q)}$  is no longer guaranteed to be locally-annotated since projecting out the existential variable  $w_3$  may cause  $R$ , in addition to  $T'$ , to be annotated with values other than  $\bar{1}$ .

Let  $Q$  be a free-connex CQ<sup>\*</sup> without self-joins using a relation symbol  $R$ . We define the *carrying variables* of  $R$  to be the free variables  $x$  such that there is a path in  $H(Q)$  from  $x$  to a variable of  $R$  where all variables on the path other than  $x$  are existential. Notice that all free variables of the atom using  $R$  are carrying, as witnessed by a path of length 0. Notice also that there might be no carrying variables in case the query is Boolean or in case  $R$  contains only existential variables and the query hypergraph is disconnected (e.g.,  $Q(y) :- R(x), S(y)$ ).

**Lemma 6.7.** *Let  $Q$  be a free-connex CQ<sup>\*</sup> without self-joins and using a relation symbol  $R$ . There exists an atom in  $Q$  whose free variables are precisely the carrying variables of  $R$ .*

*Proof.* Let us first consider the case that the  $R$  atom contains only free variables. In this case, the carrying variables are exactly the variables of  $R$ , so the  $R$  atom satisfies the lemma requirements, and we are done. Assume next that  $R$  contains existential variables.

Since  $Q$  is free-connex, there is a join tree  $T$  containing a vertex for each of its atoms and another vertex  $v$  containing exactly  $\text{free}(Q)$ . Consider the path  $P$  in  $T$  from  $v$  to the vertex of  $R$ . We claim that the first vertex on  $P$  containing an existential variable that is connected by a path of only existential variables to a variable of  $R$  satisfies the lemma.<sup>3</sup> We

<sup>3</sup>This definition may seem overly complicated, but it is there to handle cases like the query  $Q(x_1, x_2) :- S(x_1, x_2, y_1), R(x_1, y_2)$  and the extended join tree  $(x_1, x_2) - (x_1, x_2, y_1) - (x_1, y_2)$ . Here, only  $x_1$  is carrying.

denote this vertex by  $u$ . Such a vertex exists since we assume that  $R$  contains existential variables, and so the vertex of  $R$  has such a path of length 0.

We can easily see that all free variables in  $u$  are carrying. Indeed, let  $x$  be a free variable of  $u$ . Since  $u$  contains an existential variable that has a path of only existential variables to a variable of  $R$ , concatenating  $x$  to the beginning of this path is a witness to the fact that  $x$  is carrying.

It is left to prove that all carrying variables appear in  $u$ . We treat  $T$  as rooted in  $v$ , and denote by  $T_u$  its subtree rooted in  $u$ . Notice that the  $R$  atom is in  $T_u$ . Let  $w_1$  be a carrying variable, and consider a path  $w_1, \dots, w_k$  witnessing this (i.e.,  $w_i$  is existential for all  $i > 1$ , and  $w_k$  is in the  $R$  atom). We claim by induction on decreasing  $i$  that every variable  $w_i$  on this path appears in  $T_u$ . First,  $w_k$  appears in  $T_u$  as it appears in the  $R$  atom. Next, consider  $w_i$  with  $i > 1$  that appears in  $T_u$ . The variable  $w_i$  cannot appear in the parent of  $u$  because  $u$  is the first vertex on  $P$  that is allowed to contain a variable starting a path of existential variables such as  $w_i, \dots, w_k$ . Since  $w_i$  appears in  $T_u$  but not in the parent of  $u$ , by the running intersection property, we have that all atoms containing  $w_i$  appear in  $T_u$ . This includes the atom witnessing the fact that  $w_i$  and  $w_{i-1}$  are neighbors. Thus, we have that  $w_{i-1}$  also appears in  $T_u$ . From this induction, we deduce that  $w_1$  appears in  $T_u$ . Since it appears in  $T_u$  and in  $v$ , by the running intersection property, it must appear in  $u$ .  $\square$

Lemma 6.7 shows that there is an atom of  $Q_{|\text{free}(Q)}$  that contains exactly the carrying variables. We denote the relation symbol of this atom by  $R_{\text{carry}}$ . Next, we prove the positive side of the dichotomy using a construction similar to Lemma 3.3 but specific to locally annotated databases with  $\oplus$ -idempotent semirings.

**Lemma 6.8.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a logarithmic-time  $\oplus$ -idempotent commutative semiring, and let  $R$  be a relation symbol of a free-connex  $CQ^*$   $Q$  without self-joins. If direct access for  $Q_{|\text{free}(Q)}$  over  $R_{\text{carry}}$ -annotated  $\mathbb{K}$ -databases is in  $\langle T_p, T_a \rangle$ , then direct access for  $Q$  over  $R$ -annotated  $\mathbb{K}$ -databases is in  $\langle \text{loglinear} + T_p, T_a \rangle$ .*

To prove Lemma 6.8, we notice that when our construction from Lemma 3.3 is applied to a locally annotated database with  $\oplus$ -idempotent semirings, the result is a locally annotated database, where the annotation is on an atom that contains all carrying variables and depends only on them. We can then move the annotations to guarantee that the construction results in an  $R_{\text{carry}}$ -annotated database.

*Proof.* We show an  $O(|D| \log |D|)$ -time construction that maps  $R$ -annotated  $\mathbb{K}$ -databases  $(D, \tau)$  of  $Q$  to  $R_{\text{carry}}$ -annotated  $\mathbb{K}$ -databases  $(D', \tau')$  of  $Q_{|\text{free}(Q)}$  such that  $Q_{|\text{free}(Q)}(D', \tau') = Q(D, \tau)$ . Given an  $R$ -annotated  $\mathbb{K}$ -databases of  $Q$ , we use our construction from Lemma 3.3. Recall that as a first step, we assign a relation to every vertex of an ext-free-connex tree. These relations are annotated either by the original annotations or by  $\bar{1}$ , so at this point, only the one vertex matching  $R$  in the tree is annotated by non-trivial annotations. Next, until we are left with only free variables, we repeatedly eliminate some leaf vertices as follows:

- (1) In the relation associated with the vertex, project out every variable not shared between the vertex and its neighbor while summing ( $\oplus$ ) annotations.
- (2) Join the obtained relation into the relation of the neighbor vertex while multiplying ( $\otimes$ ) annotations.
- (3) Remove the leaf from the tree, its associated atom from the query, and its associated relation from the database.

We first claim that our construction results in a locally-annotated database. If we eliminate a relation that is only annotated by  $\bar{1}$ , then since the semiring is  $\oplus$ -idempotent (and in particular  $\bar{1} \oplus \bar{1} = \bar{1}$ ), we get that its projection is still only annotated by  $\bar{1}$ , and when we join the projection into another relation, if it was only annotated by  $\bar{1}$ , it remains this way. When we eliminate a relation with unrestricted annotations, we know that it is the only relation annotated by values other than  $\bar{1}$ ; thus, at the end of the step, the database is still locally annotated, except now the relation with unrestricted annotation is its parent in the tree.

We next claim that the obtained annotation does not depend on non-carrying free variables. In fact, we prove a slightly stronger claim. We say that a variable is relevant if it is either carrying or it is an existential variable that has a path to an existential variable of  $R$  that only goes through existential variables. Notice that any variables that appear in the same atom as an existential relevant variable are also relevant. We claim that, during the elimination process, whenever a relation contains two facts that differ only on non-relevant variables, these two facts have the same annotation. Since the elimination process results in only free variables, we get that at its end, the annotation depends only on carrying variables.

This claim is trivially true in the beginning since all variables in  $R$  are relevant and all other relations are always annotated by  $\bar{1}$ . Consider the step of eliminating a vertex  $v_e$  into its parent  $v_p$ , and assume as an induction hypothesis that the claim holds before this step. We want to show that the claim also holds in  $v_p$  after this step (this is enough because  $v_e$  is eliminated in this step and all other relations remain unchanged). We consider two cases.

- (1) The first case is that all variables in  $v_e$  are relevant. In particular, all shared variables between  $v_p$  and  $v_e$  are relevant. Thus, when joining the projected  $v_e$  into  $v_p$ , facts in  $v_p$  that disagree only on non-relevant variables, agree on the shared variables with  $v_e$ , will be multiplied by the same annotation, and will keep having identical annotations to one another.
- (2) The second case is that  $v_e$  contains a non-relevant variable. From the induction hypothesis, before the projection of  $v_e$ , the annotations in  $v_e$  do not depend on the shared non-relevant variables. Since  $v_e$  contains a non-relevant variable, we have that  $v_e$  does not contain a relevant existential variable, or in other words, all existential variables in  $v_e$  are non-relevant. From the induction hypothesis, before the projection of  $v_e$ , the annotations in  $v_e$  do not depend on any existential variables. Due to the idempotence, the projection does not change the annotation of the resulting facts. Thus, after the projection, the annotation in  $v_e$  still does not depend on the shared non-relevant variables. When joining the projected  $v_e$  into  $v_p$ , facts in  $v_p$  that disagree only on non-relevant variables would be multiplied by the same annotations, and so will keep having identical annotations to one another.

Since the nontrivial annotations appear in only one atom and may only depend on carrying variables, we can move them to any atom containing all carrying variables, and in particular  $R_{\text{carry}}$ .  $\square$

Lemma 6.8 proves the positive direction. That is, assuming a  $\oplus$ -idempotent commutative semiring, if we want to know that  $Q$  over  $R$ -annotated databases admits efficient direct access, then it is enough to require that  $Q_{\text{free}(Q)}$  over  $R_{\text{carry}}$ -annotated databases admits efficient direct access. We next prove the other direction.

**Lemma 6.9.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a commutative semiring, and let  $R$  be a relation symbol of a free-connex  $CQ^*$   $Q$  without self-joins. If direct access for  $Q$  over  $R$ -annotated  $\mathbb{K}$ -databases is in  $\langle T_p, T_a \rangle$ , then direct access for  $Q_{|\text{free}(Q)}$  over  $R_{\text{carry}}$ -annotated  $\mathbb{K}$ -databases is in  $\langle \text{linear} + T_p, T_a \rangle$ .*

We first notice that we cannot simply use the same construction as in the proof of Lemma 3.2 to prove Lemma 6.9. There, we assigned a value to all existential variables, so we do not necessarily have that answers with distinct annotations use distinct facts of  $R$ , and we cannot simply move the annotations from  $R_{\text{carry}}$  to  $R$ . Instead, we use a different construction that propagates the values assigned to the variables of  $R_{\text{carry}}$  through the existential variables into the variables of  $R$ , thus allowing us to move the annotations from  $R_{\text{carry}}$  to  $R$ .

*Proof.* We show an  $O(|D'|)$ -time construction that maps  $R_{\text{carry}}$ -annotated  $\mathbb{K}$ -databases  $(D', \tau')$  of  $Q_{|\text{free}(Q)}$  to  $R$ -annotated  $\mathbb{K}$ -databases  $(D, \tau)$  of  $Q$  such that  $Q_{|\text{free}(Q)}(D', \tau') = Q(D, \tau)$ . As a first step, we perform the Yannakakis full reduction [Yan81], and then we can assume that every fact  $f'$  of  $D'$  is consistent with  $Q_{|\text{free}(Q)}$  in the sense that we can obtain  $f'$  from the corresponding atom of  $Q_{|\text{free}(Q)}$  by assigning some values to the variables. Then, we can denote by  $f'[x]$  the value of  $f'$  in the positions of the variable  $x$ . We also denote by  $\langle f' \rangle$  a value that encodes the tuple of the fact  $f'$ .

Recall the definition of carrying variables. The atom of  $R_{\text{carry}}$  contains all free variables of  $Q$  that have a path in  $H(Q)$  to a variable of the atom of  $R$  where every variable on the path, aside from the carrying variable, is existential. Denote by  $V$  the set of all existential variables that are part of such a path. Note that if an atom  $\varphi$  of  $Q$  contains a variable in  $V$ , then every free variable in  $\varphi$  must be a carrying variable by definition, so the free variables of  $\varphi$  all appear in the atom of  $R_{\text{carry}}$ . Recall that  $Q_{|\text{free}(Q)}$  is created directly from  $Q$  by restricting  $Q$  to the free variables, so every atom  $\varphi$  in  $Q$  has a matching atom  $\varphi_{|\text{free}(Q)}$  in  $Q_{|\text{free}(Q)}$ . Let us now construct  $(D, \tau)$  as follows.

- 
- 1: Remove dangling facts using the Yannakakis algorithm [Yan81]
  - 2: **for all** atoms  $\varphi$  of  $Q$  **do**
  - 3:   **if**  $\varphi$  has a variable in  $V$  **then**
  - 4:     source :=  $R_{\text{carry}}$
  - 5:   **else**
  - 6:     source := the relation symbol of  $\varphi_{|\text{free}(Q)}$  in  $Q_{|\text{free}(Q)}$
  - 7:   **for all** facts  $f'$  in the relation of source in  $D'$  **do**
  - 8:     **for all**  $x \in \text{vars}(\varphi)$  **do**
  - 9:       **if**  $x \in \text{free}(\varphi)$  **then**
  - 10:          $c_x := f'[x]$
  - 11:       **else if**  $x \in V$  **then**
  - 12:          $c_x := \langle f' \rangle$
  - 13:       **else**
  - 14:          $c_x := a_0$
  - 15:     Add to  $D$  a new fact  $f$  obtained from  $\varphi$  by replacing each variable  $x$  with  $c_x$
  - 16:     **if** the relation symbol of  $\varphi$  is  $R$  **then**
  - 17:        $\tau(f) :=$  the annotation of the fact agreeing with  $f'$  in  $R_{\text{carry}}$
  - 18:     **else**
  - 19:        $\tau(f) := \bar{1}$
-




Figure 4: An example for the construction from Lemma 6.9 for  $R$ -annotated databases on the query  $Q(x_1, x_2) :- R(w_1, w_2), S(w_2, x_1), T(x_1, x_2, w_3), U(x_2, w_4, w_5)$ . Here,  $x_1$  is the carrying variable,  $R_{\text{carry}} = S|_{\text{free}(Q)}$ , and  $V = \{w_1, w_2\}$ .

To create  $D$ , we iterate over every atom of  $Q$  (Line 2) and construct its relation. We begin by choosing the source of the data for the relation (Line 3). If  $\varphi$  contains a variable in  $V$ , we set its data source to be  $R_{\text{carry}}$ . Otherwise, we set its data source to be  $\varphi|_{\text{free}(Q)}$ . Then, we iterate over the facts of the data source in  $D'$  (Line 7), and for each fact  $f'$ , we create in the relation of  $\varphi$  in  $D$  a new fact  $f$ . The new fact  $f$  is created so that it agrees on the free variables with  $f'$  (Line 10), has an encoding of  $\langle f' \rangle$  in place of the variables of  $V$  (Line 12), and has an arbitrary value  $a_0$  in place of all other existential variables (Line 14). As for the annotations, we annotate all facts with  $\bar{1}$  (Line 19) except for those of  $R$ , which are annotated the same as their source facts. Therefore,  $(D, \tau)$  is  $R$ -annotated. See Figure 4 for an illustration of this construction.

We claim next that the construction time is  $O(|D'|)$ . Indeed, the Yannakakis algorithm requires linear time (using perfect hash tables in the RAM model of computation). Then, Line 2 iterates over a constant number of atoms, Line 7 iterates over a linear number of facts, and Line 8 iterates over a constant number of variables. Overall, we have a linear number of operations, each requiring constant time. It is left to prove the correctness.

First, we show that every answer tuple  $\vec{a}$  in  $Q(D)$  is also in  $Q|_{\text{free}(Q)}(D')$ . We claim that the mapping from the free variables to the domain given by  $\vec{a}$  is a homomorphism from  $Q|_{\text{free}(Q)}$  to  $D'$ . Consider an atom  $\varphi$  of  $Q$ , and denote its relation symbol by  $S$ . First, consider the case that  $\varphi$  does not contain a variable in  $V$ . Since  $D$  contains some  $S$ -fact that supports  $\vec{a}$  and each  $S$ -fact was created from an  $S|_{\text{free}(Q)}$ -fact in  $D'$  that agrees with it on the free variables, that means that  $D'$  contains an  $S|_{\text{free}(Q)}$ -fact that supports  $\vec{a}$ . The second case

is that  $\varphi$  contains a variable in  $V$ . That means that every free variable in  $\varphi$  is a carrying variable that appears also in  $R_{\text{carry}}$ . In this case, each  $S$ -fact was created from an  $R_{\text{carry}}$ -fact that agrees with it on the free variables of  $S$ . As we removed dangling tuples, there must also be an  $S_{|\text{free}(Q)}$ -fact in  $D'$  that agrees with that  $R_{\text{carry}}$ -fact on the free variables of  $S$ . This  $S_{|\text{free}(Q)}$ -fact supports  $\vec{a}$ .

Next, we show that every answer tuple  $\vec{a}$  in  $Q_{|\text{free}(Q)}(D')$  is also in  $Q(D)$ . Since  $\vec{a}$  is an answer to  $Q_{|\text{free}(Q)}(D')$ , there exists in  $D'$  an  $R_{\text{carry}}$ -fact  $f_{\text{carry}}$  that supports  $\vec{a}$ . We claim that there is a homomorphism from  $Q$  to  $D$  that assigns the free variables according to  $\vec{a}$ , assigns  $\langle f_{\text{carry}} \rangle$  to the variables of  $V$ , and assigns  $a_0$  to all other existential variables. Indeed, consider some atom  $\varphi$  of  $Q$ . If  $\varphi$  contains a variable of  $V$ , then by construction its relation contains a fact that: agrees with  $f_{\text{carry}}$  on the carrying variables, has  $\langle f_{\text{carry}} \rangle$  in place of any existential variable in  $V$ , and has  $a_0$  in place of any other existential variable. If the atom  $\varphi$  does not contain a variable of  $V$ , we rely on the fact that the relation of  $\varphi_{|\text{free}(Q)}$  in  $D'$  contains some fact  $f'$  that supports  $\vec{a}$  as an answer. We added a fact to the relation of  $\varphi$  in  $D$  that agrees with  $f'$  on the free variables and has  $a_0$  in place of any existential variables. We conclude that the claimed homomorphism exists, and  $\vec{a}$  is also an answer for  $Q(D)$ .

It is left to show that, given an answer tuple  $\vec{a}$  in  $Q_{|\text{free}(Q)}(D')$  and  $Q(D)$ , we have that  $\tau(\vec{a}) = \tau'(\vec{a})$ . As  $Q_{|\text{free}(Q)}$  is full, there is a single  $R_{\text{carry}}$ -fact  $f_{\text{carry}}$  in  $D'$  that supports  $\vec{a}$ , and since  $Q_{|\text{free}(Q)}$  is  $R_{\text{carry}}$  annotated, we know that  $\tau'(\vec{a}) = \tau'(f_{\text{carry}})$ . In our construction, existential variables not in  $V$  are always assigned the value  $a_0$ , and existential variables in  $V$  are always assigned the value  $\langle f_{\text{carry}} \rangle$ , where  $f_{\text{carry}}$  is a fact of  $R_{\text{carry}}$ . Given an answer  $\vec{a}$  to  $Q$ , there is exactly one homomorphism from  $Q$  to  $D$  that supports it, the one where free variables are assigned according to  $\vec{a}$ , the  $V$  variables are assigned  $\langle f_{\text{carry}} \rangle$ , where  $f_{\text{carry}}$  is the fact of  $R_{\text{carry}}$  that supports  $\vec{a}$ , and the other existential variables are assigned  $a_0$ . Since  $D$  is locally annotated,  $\tau(\vec{a}) = \tau(f_R)$ , where  $f_R$  is the fact supporting  $\vec{a}$  in  $R$ . If  $R$  contains a variable of  $V$ , then its source is  $R_{\text{carry}}$ , and  $\tau(f_R) = \tau'(f_{\text{carry}})$ . Otherwise, its source is  $R_{|\text{free}(Q)}$ . In this case, since  $R$  does not contain any variable of  $V$ , it means that  $V$  is empty by its definition,  $R$  contains only free variables, and the carrying variables are exactly the variables of  $R$ , so  $R_{\text{carry}}$  and  $R_{|\text{free}(Q)}$  have the same variables. In this case, if  $f'$  is the fact from  $R_{|\text{free}(Q)}$  that constructs  $f_R$ , we have that  $f_{\text{carry}}$  is the fact agreeing with  $f'$  in  $R_{\text{carry}}$ , and our construction sets  $\tau(f_R) = \tau'(f_{\text{carry}})$ . Overall, we have that  $\tau(\vec{a}) = \tau(f_R) = \tau'(f_{\text{carry}}) = \tau'(\vec{a})$ .  $\square$

We now have a classification of the CQ\*s without self-joins over locally-annotated databases in the case of an idempotent addition. By Lemma 6.8 and Lemma 6.9, direct access for  $Q$  over  $R$ -annotated  $\mathbb{K}$ -databases is in  $\langle \text{loglinear}, \text{log} \rangle$  if and only if direct access for  $Q_{|\text{free}(Q)}$  over  $R_{\text{carry}}$ -annotated  $\mathbb{K}$ -databases is in  $\langle \text{loglinear}, \text{log} \rangle$ . Then, we only need to apply Theorem 6.4 on the full  $Q_{|\text{free}(Q)}$  to determine the complexity of direct access for  $Q_{|\text{free}(Q)}$  based on the structure of its  $R_{\text{carry}}$ -deannotation.

**Theorem 6.10.** *Let  $(\mathbb{K}, \oplus, \otimes, \bar{0}, \bar{1})$  be a logarithmic-time  $\oplus$ -idempotent commutative semiring. Let  $R$  be a relation symbol of a CQ\*  $Q$  without self-joins, and let  $Q_R$  be the  $R_{\text{carry}}$ -deannotation of  $Q_{|\text{free}(Q)}$ .*

- (1) *If  $Q_R$  is acyclic and has no disruptive trio, then direct access for  $Q$  is in  $\langle \text{loglinear}, \text{log} \rangle$  on  $R$ -annotated  $\mathbb{K}$ -databases.*
- (2) *Otherwise, if  $\mathbb{N} \subseteq \mathbb{K}$ , then direct access for  $Q$  is not in  $\langle \text{loglinear}, \text{log} \rangle$  on  $R$ -annotated  $\mathbb{K}$ -databases, assuming the HYPERCLIQUE and SparseBMM hypotheses.*

Notice that we can use the positive side of this dichotomy even if  $Q$  has self-joins by first assigning different atoms with different relation symbols and copying the relevant relations. We can use Theorem 6.10 to find tractable cases for AggCQs that have a translation to a  $\oplus$ -idempotent semiring.

**Corollary 6.11.** *Consider an AggCQ  $Q(\vec{x}, \alpha(w), \vec{z}) :- \varphi_1(\vec{x}, \vec{y}, \vec{z}), \dots, \varphi_\ell(\vec{x}, \vec{y}, \vec{z})$  where either  $\alpha$  is Min or Max or  $\alpha$  is CountD and  $w$  has a logarithmic domain. Let  $Q'(\vec{x}, \star, \vec{z})$  be a CQ $^\star$  with the same body as  $Q$ , and let  $Q_R$  be the  $R_{\text{carry}}$ -deannotation of  $Q'_{|\text{free}(Q')}$ , where  $R$  is an atom containing  $w$ . If  $Q_R$  is acyclic and has no disruptive trio, then direct access for  $Q$  is in  $\langle \text{loglinear}, \text{log} \rangle$ .*

## 7. CONCLUDING REMARKS

Direct access provides an opportunity to efficiently evaluate queries with ordering, even when the number of answers is enormous. Past research studied the feasibility of direct access for CQs, and here we embarked on the exploration of this problem for queries that involve grouping and aggregation, either as standard (SQL) aggregate functions (AggCQs) or annotation with commutative semirings (CQ $^\star$ s). For CQ $^\star$ s (and AggCQs that can be efficiently simulated by CQ $^\star$ s), we showed that the past classification holds as long as the computed value is not involved in the order. This result also holds for count-distinct, but under the more restricted condition that the counted attribute is treated as an ordinary free variable (coming last in the lexicographic order). Involving the computed value in the order introduces hardness pretty quickly, and we established the narrower tractability condition in this case. We also showed how the condition generalizes when restricting our attention to databases locally annotated by a semiring with an idempotent addition.

An important direction for future work is the exploration of queries beyond free-connex ones; for that, we need to allow for broader yardsticks of efficiency, as done for direct access for CQs without aggregation [BCM25] and as done for Functional Aggregate Queries (FAQ) for CQs with aggregation and traditional query evaluation [KNR16]. Moreover, we plan to explore the extension of our results to queries with self-joins, and we believe that recent results [BCM25] can be used towards such an extension. It is also left for future work to better understand the limits of computation and establish lower bounds (and dichotomies) for general classes of queries, commutative semirings, and aggregate functions. Another important direction is to explore the ability to efficiently maintain the direct-access structure for CQ $^\star$ s and AggCQs through updates of the database, as previously studied in the context of aggregate queries [Kep20] as well as non-aggregate queries [BKS18, SVZ18]. Finally, we plan to investigate the practical behavior of our algorithms and understand how well the theoretical acceleration is realized in comparison to existing query engines.

## ACKNOWLEDGMENT

The work of Idan Eldar and Benny Kimelfeld was supported by the German Research Foundation (DFG) grant KI 2348/1-1 (DIP Program).

## REFERENCES

- [BB13] Johann Brault-Baron. *De la pertinence de l'énumération : complexité en logiques propositionnelle et du premier ordre*. Theses, Université de Caen, April 2013.
- [BCM25] Karl Bringmann, Nofar Carmeli, and Stefan Mengel. Tight fine-grained bounds for direct access on join queries. *ACM Transactions on Database Systems*, 50(1):1–44, 2025.
- [BDG07] Guillaume Bagan, Arnaud Durand, and Etienne Grandjean. On acyclic conjunctive queries and constant delay enumeration. In *International Workshop on Computer Science Logic*, pages 208–222. Springer, 2007.
- [BDGO08] Guillaume Bagan, Arnaud Durand, Etienne Grandjean, and Frédéric Olive. Computing the  $j$ th solution of a first-order query. *RAIRO-Theoretical Informatics and Applications*, 42(1):147–164, 2008.
- [BKS17] Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering conjunctive queries under updates. In *PODS*, pages 303–318. ACM, 2017.
- [BKS18] Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering FO+MOD queries under updates on bounded degree databases. *ACM Transactions on Database Systems*, 43(2):7:1–7:32, 2018.
- [CKKK21] Nofar Carmeli, Batya Kenig, Benny Kimelfeld, and Markus Kröll. Efficiently enumerating minimal triangulations. *Discrete Applied Mathematics*, 303:216–236, 2021.
- [CNS07] Sara Cohen, Werner Nutt, and Yehoshua Sagiv. Deciding equivalences among conjunctive aggregate queries. *Journal of the ACM (JACM)*, 54(2):5–es, 2007.
- [Coh93] Henri Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate texts in mathematics*. Springer, 1993.
- [CTG<sup>+</sup>23] Nofar Carmeli, Nikolaos Tziavelis, Wolfgang Gatterbauer, Benny Kimelfeld, and Mirek Riedewald. Tractable orders for direct access to ranked answers of conjunctive queries. *ACM Transactions on Database Systems*, 48(1):1–45, 2023.
- [CZB<sup>+</sup>22] Nofar Carmeli, Shai Zeevi, Christoph Berkholz, Alessio Conte, Benny Kimelfeld, and Nicole Schweikardt. Answering (unions of) conjunctive queries using random access and random-order enumeration. *ACM Transactions on Database Systems*, 47(3):9:1–9:49, 2022.
- [ECK24] Idan Eldar, Nofar Carmeli, and Benny Kimelfeld. Direct access for answers to conjunctive queries with aggregation. In *ICDT*, pages 4–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.
- [GJ22] Étienne Grandjean and Louis Jachiet. Which arithmetic operations can be performed in constant time in the ram model with addition? *arXiv preprint arXiv:2206.13851*, 2022.
- [GKT07] Todd J Green, Grigoris Karvounarakis, and Val Tannen. Provenance semirings. In *PODS*, pages 31–40. ACM, 2007.
- [GO95] Anka Gajentaan and Mark H Overmars. On a class of  $O(n^2)$  problems in computational geometry. *Computational geometry*, 5(3):165–185, 1995.
- [Gol04] Martin Charles Golumbic. *Algorithmic graph theory and perfect graphs*, volume 57. Elsevier, 2004.
- [IUV17] Muhammad Idris, Martín Ugarte, and Stijn Vansummeren. The dynamic yannakakis algorithm: Compact and efficient query processing under updates. In *SIGMOD*, pages 1259–1274, 2017.
- [KCM<sup>+</sup>20] Mahmoud Abo Khamis, Ryan R. Curtin, Benjamin Moseley, Hung Q. Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. Functional aggregate queries with additive inequalities. *ACM Transactions on Database Systems*, 45(4):17:1–17:41, 2020.
- [Kep20] Jens Keppeler. *Answering Conjunctive Queries and FO+MOD Queries under Updates*. PhD thesis, Humboldt University of Berlin, Germany, 2020.
- [KNR16] Mahmoud Abo Khamis, Hung Q. Ngo, and Atri Rudra. FAQ: questions asked frequently. In *PODS*, pages 13–28. ACM, 2016.
- [MRV26] Thomas Muñoz, Cristian Riveros, and Stijn Vansummeren. Enumeration and updates for conjunctive linear algebra queries through expressibility. *Log. Methods Comput. Sci.*, 22(1), 2026.
- [OS16] Dan Olteanu and Maximilian Schleich. Factorized databases. *SIGMOD Rec.*, 45(2):5–16, 2016.
- [OZ15] Dan Olteanu and Jakub Závodný. Size bounds for factorised representations of query results. *ACM Transactions on Database Systems (TODS)*, 40(1):1–44, 2015.
- [Pat10] Mihai Patrascu. Towards polynomial lower bounds for dynamic problems. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 603–610, 2010.
- [RS09] Christopher Ré and Dan Suciu. The trichotomy of having queries on a probabilistic database. *The VLDB Journal*, 18(5):1091–1116, 2009.

- [SVZ18] Thomas Schwentick, Nils Vortmeier, and Thomas Zeume. Dynamic complexity under definable changes. *ACM Transactions on Database Systems*, 43(3):12:1–12:38, 2018.
- [VW15] Virginia Vassilevska Williams. Hardness of easy problems: Basing hardness on popular conjectures such as the strong exponential time hypothesis (invited talk). In *IPEC*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
- [Yan81] Mihalis Yannakakis. Algorithms for acyclic database schemes. In *VLDB*, pages 82–94. IEEE Computer Society, 1981.