


# 1 Enumeration and Updates for Conjunctive Linear 2 Algebra Queries Through Expressibility


3 **Thomas Muñoz Serrano** ✉ 

4 UHasselt, Data Science Institute, Belgium

5 **Cristian Riveros** ✉ 

6 Pontificia Universidad Católica de Chile, Chile

7 Millennium Institute for Foundational Research on Data, Chile

8 **Stijn Vansummeren** ✉ 

9 UHasselt, Data Science Institute, Belgium

## 10 — Abstract —

11 Due to the importance of linear algebra and matrix operations in data analytics, there is significant  
12 interest in using relational query optimization and processing techniques for evaluating (sparse) linear  
13 algebra programs. In particular, in recent years close connections have been established between  
14 linear algebra programs and relational algebra that allow transferring optimization techniques of the  
15 latter to the former. In this paper, we ask ourselves which linear algebra programs in MATLANG  
16 correspond to the free-connex and q-hierarchical fragments of conjunctive first-order logic. Both  
17 fragments have desirable query processing properties: free-connex conjunctive queries support  
18 constant-delay enumeration after a linear-time preprocessing phase, and q-hierarchical conjunctive  
19 queries further allow constant-time updates. By characterizing the corresponding fragments of  
20 MATLANG, we hence identify the fragments of linear algebra programs that one can evaluate with  
21 constant-delay enumeration after linear-time preprocessing and with constant-time updates. To  
22 derive our results, we improve and generalize previous correspondences between MATLANG and  
23 relational algebra evaluated over semiring-annotated relations. In addition, we identify properties  
24 on semirings that allow to generalize the complexity bounds for free-connex and q-hierarchical  
25 conjunctive queries from Boolean annotations to general semirings.

26 **2012 ACM Subject Classification** Theory of computation → Database theory

27 **Keywords and phrases** Query evaluation, conjunctive queries, linear algebra, enumeration algorithms.

28 **Digital Object Identifier** 10.4230/LIPIcs.ICDT.2024.12

29 **Related Version** *Full paper version with proofs:* <https://arxiv.org/abs/2310.04118> [38]

30 **Acknowledgements** Cristian Riveros was funded by ANID – Millennium Science Initiative Program –  
31 Code ICM17\_0 and ANID Fondecyt Regular project 1230935. Thomas Muñoz and Stijn Vansummeren  
32 were supported by the Bijzonder Onderzoeksfonds (BOF) of Hasselt University (Belgium) under  
33 Grants No. BOF21OWB13 and BOF20ZAP02 as well as by the Research Foundation Flanders  
34 (FWO) under research project Grant No. G019222N.

## 35 **1** Introduction

36 Linear algebra forms the backbone of modern data analytics, as most machine learning  
37 algorithms are coded as sequences of matrix operations [1, 4, 19, 20, 42]. In practice, linear  
38 algebra programs operate over matrices with millions of entries. Therefore, efficient evaluation  
39 of linear algebra programs is a relevant challenge for data management systems which has  
40 attracted research attention with several proposals in the area [30, 33, 34, 37, 41].

41 To optimize and evaluate linear algebra programs, we must first agree on the language in  
42 which such programs are expressed. There has been a renewed interest in recent years for  
43 designing query languages for specifying linear algebra programs and for understanding their



© Thomas Muñoz, Cristian Riveros and Stijn Vansummeren;  
licensed under Creative Commons License CC-BY 4.0

27th International Conference on Database Theory (ICDT 2024).

Editors: Graham Cormode and Michael Shekelyan; Article No. 12; pp. 12:1–12:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

44 expressive power [6,12,13,23,40]. One such proposal is MATLANG [12], a formal matrix query  
 45 language that consists of only the basic linear algebra operations and whose extensions (e.g.,  
 46 for-MATLANG) achieve the expressive power of most linear algebra operations [23]. Although  
 47 MATLANG is a theoretical query language, it includes the core of any linear algebra program  
 48 and, thus, the optimization and efficient evaluation of MATLANG could have a crucial impact  
 49 on today’s machine learning systems.

50 In this work, we study the efficient evaluation of MATLANG programs over sparse matrices  
 51 whose entries are taken from a general semiring. We consider MATLANG evaluation in both  
 52 the static and dynamic setting. For static evaluation, we want to identify the fragment that  
 53 one can evaluate by preprocessing the input in linear time to build a data structure for  
 54 enumerating the output entries with constant-delay. For dynamic evaluation, we assume that  
 55 matrix entries are updated regularly and we want to maintain the output of a MATLANG  
 56 query without recomputing it. For this dynamic setting, we aim to identify the MATLANG  
 57 fragment that one can evaluate by taking linear time in the size of the update to refresh the  
 58 aforementioned data structure so that it supports constant-delay enumeration of the modified  
 59 output entries. These guarantees for both scenarios have become the holy grail for algorithmic  
 60 query processing since, arguably, it is the best that one can achieve complexity-wise in terms  
 61 of the input, the output, and the cost of an update [5,8,11,16,17,31,35,36,39].

62 To identify the MATLANG fragments with these guarantees, our approach is straightfor-  
 63 ward but effective. Instead of developing evaluation algorithms from scratch, we establish  
 64 a direct correspondence between linear algebra and relational algebra to take advantage  
 65 of the query evaluation results for conjunctive queries. Indeed, prior work has precisely  
 66 characterized which subfragments of conjunctive queries can be evaluated and updated  
 67 efficiently [5,8,9,31]. Our main strategy, then, is to link these conjunctive query fragments  
 68 to corresponding linear algebra fragments. More specifically, our contributions are as follows.

- 69 1. We start by understanding the deep connection between positive first-order logic ( $\text{FO}^+$ )  
 70 over binary relations and *sum-MATLANG* [23], an extension of MATLANG. We formalize  
 71 this connection by introducing schema encodings, which specify how relations simulate  
 72 matrices and vice-versa, forcing a lossless relationship between both. Using this machinery,  
 73 we show that *sum-MATLANG* and positive first-order logic are equally expressive over any  
 74 relation, matrix, and matrix dimension (including non-rectangular matrices). Moreover,  
 75 we show that conjunctive queries (CQ) coincide with *sum-MATLANG* without matrix ad-  
 76 dition, which we call *conj-MATLANG*. This result forms the basis for linking both settings  
 77 and translating the algorithmic results from CQ to subfragments of *conj-MATLANG*.
- 78 2. We propose free-connex MATLANG (*fc-MATLANG*) for static evaluation, a natural  
 79 MATLANG subfragment that we show to be equally expressive as *free-connex CQ* [5], a  
 80 subfragment of CQ that allows linear time preprocessing and constant-delay enumera-  
 81 tion. To obtain our expressiveness result, we show that free-connex CQs over binary  
 82 relations are equally expressive as the two-variable fragment of conjunctive  $\text{FO}^+$ , a logical  
 83 characterization of this class that could be of independent interest.
- 84 3. For the dynamic setting we introduce the language *qh-MATLANG*, a MATLANG fragment  
 85 that we show equally expressive to *q-hierarchical CQ* [9,31], a fragment of CQ that allows  
 86 constant update time and constant-delay enumeration.
- 87 4. Both free-connex and q-hierarchical CQ are known to characterize the class of CQs  
 88 that one can evaluate efficiently on Boolean databases. We are interested, however, in  
 89 evaluating MATLANG queries on matrices featuring entries in a *general semiring*. To  
 90 obtain the complexity bounds for *fc-MATLANG* and *qh-MATLANG* on general semirings,  
 91 therefore, we show that the upper and lower bounds for free-connex and q-hierarchical

92 CQs generalize from Boolean annotations to classes of semirings which includes most  
 93 semirings used in practice, like the reals. The tight expressiveness connections established  
 94 in this paper then prove that for such semirings  $\text{fc-MATLANG}$  and  $\text{qh-MATLANG}$  can  
 95 be evaluated with the same guarantees as their CQ counterparts and that they are  
 96 optimal: one cannot evaluate any other  $\text{conj-MATLANG}$  query outside this fragment  
 97 under complexity-theoretic assumptions [9].

98 An extended version of this paper that includes full proofs of formal statements is available  
 99 online [38].

100 **Related work.** In addition to the work that has already been cited above, the following  
 101 work is relevant. Brijder et al. [13] have shown equivalence between  $\text{MATLANG}$  and  $\text{FO}_3^+$ ,  
 102 the 3-variable fragment of positive first order logic. By contrast, we show equivalence  
 103 between  $\text{sum-MATLANG}$  and  $\text{FO}^+$ , and study the relationship between the free-connex and  
 104 q-hierarchical fragments of  $\text{MATLANG}$  and  $\text{FO}^\wedge$ , the conjunctive fragment of positive first  
 105 order logic.

106 Geerts et al. [23] previously established a correspondence between  $\text{sum-MATLANG}$  and  
 107  $\text{FO}^+$ . However, as we illustrate in [38] their correspondence is (1) restricted to square  
 108 matrices, (2) asymmetric between the two settings, and (3) encodes matrix instances as  
 109 databases of more than linear size, making it unsuitable to derive the complexity bounds.

110 Eldar et al. [18] have recently also generalized complexity bounds for free-connex CQs  
 111 from Boolean annotations to general semirings. Nevertheless, this generalization is with  
 112 respect to *direct access*, not enumeration. In their work the focus is to compute aggregate  
 113 queries, which is achieved by providing direct access to the answers of a query even if the  
 114 annotated value (aggregation result) is zero. By contrast, in our setting, zero-annotated  
 115 values must not be reported during the enumeration of query answers. This difference in  
 116 the treatment of zero leads to a substantial difference in the properties that a semiring must  
 117 have in order to generalize the existing complexity bounds.

118 There are deep connections known between the treewidth and the number of variables of a  
 119 conjunctive  $\text{FO}^+$  formula ( $\text{FO}^\wedge$ ). For example, Kolaitis and Vardi established the equivalence  
 120 of boolean queries in  $\text{FO}_k^\wedge$ , the  $k$ -variable fragment of  $\text{FO}^\wedge$ , and boolean queries in  $\text{FO}^\wedge$  of  
 121 treewidth less than  $k$ . Because they focus on boolean queries (i.e., without free variables),  
 122 this result does not imply our result that for binary queries free-connex  $\text{FO}^\wedge$  equals  $\text{FO}_2^\wedge$ .  
 123 Similarly, Geerts and Reutter [24] introduce a tensor logic  $\text{TL}$  over binary relations and  
 124 show that conjunctive expressions in this language that have treewidth  $k$  can be expressed  
 125 in  $\text{TL}_{k+1}$ , the  $k$ -variable fragment of  $\text{TL}$ . While they do take free variables into account,  
 126 we show in [38] that there are free-connex conjunctive queries with 2 free variables with  
 127 treewidth 2 in their formalism—for which their result hence only implies expressibility in  
 128  $\text{FO}_3^\wedge$ , not  $\text{FO}_2^\wedge$  as we show here.

129 Several proposals [30, 33, 34, 37, 41] have been made regarding the efficient evaluation of  
 130 linear algebra programs in the last few years. All these works focused on query optimization  
 131 without formal guarantees regarding the preprocessing, updates, or enumeration in query  
 132 evaluation. To the best of our knowledge, this is the first work on finding subfragments of a  
 133 linear algebra query language (i.e.,  $\text{MATLANG}$ ) with such efficient guarantees.

## 134 2 Preliminaries

135 In this section we recall the main definitions of  $\text{MATLANG}$ , a query language on matrices,  
 136 and first order logic ( $\text{FO}$ ), a query language on relations.

## 12:4 Enumeration and Updates for Conjunctive Linear Algebra Queries

137 **Semirings.** We evaluate both languages over arbitrary commutative and non-trivial semirings.  
 138 A (commutative and non-trivial) *semiring*  $(K, \oplus, \odot, \mathbf{0}, \mathbf{1})$  is an algebraic structure  
 139 where  $K$  is a non-empty set,  $\oplus$  and  $\odot$  are binary operations over  $K$ , and  $\mathbf{0}, \mathbf{1} \in K$  with  $\mathbf{0} \neq \mathbf{1}$ .  
 140 Furthermore,  $\oplus$  and  $\odot$  are associative operations,  $\mathbf{0}$  and  $\mathbf{1}$  are the identities of  $\oplus$  and  $\odot$ ,  
 141 respectively,  $\oplus$  and  $\odot$  are commutative operations,  $\odot$  distributes over  $\oplus$ , and  $\mathbf{0}$  annihilates  
 142  $K$  (i.e.  $\mathbf{0} \odot k = k \odot \mathbf{0} = \mathbf{0}$ ). We use  $\bigoplus_L$  and  $\bigodot_L$  to denote the  $\oplus$  and  $\odot$  operation over all  
 143 elements in  $L \subseteq K$ , respectively. Typical examples of semirings are the reals  $(\mathbb{R}, +, \times, 0, 1)$ ,  
 144 the natural numbers  $(\mathbb{N}, +, \times, 0, 1)$ , and the boolean semiring  $\mathbb{B} = (\{\mathbf{t}, \mathbf{f}\}, \vee, \wedge, \mathbf{f}, \mathbf{t})$ .

145 Henceforth, when we say “semiring” we mean “commutative and non-trivial” semiring.  
 146 We fix such an arbitrary semiring  $K$  throughout the document. We denote by  $\mathbb{N}_{>0}$  the set of  
 147 non-zero natural numbers.

148 **Matrices and size symbols.** A  $K$ -*matrix* (or just matrix) of dimension  $m \times n$  is a  $m \times n$   
 149 matrix with elements in  $K$  as its entries. We write  $\mathbf{A}_{ij}$  to denote the  $(i, j)$ -entry of  $\mathbf{A}$ .  
 150 Matrices of dimension  $m \times 1$  are *column vectors* and those of dimension  $1 \times n$  are *row vectors*.  
 151 We also refer to matrices of dimension  $1 \times 1$  as *scalars*.

152 We assume a collection of *size symbols* denoted with greek letters  $\alpha, \beta, \dots$  and assume  
 153 that the natural number 1 is a valid size symbol. A *type* is a pair  $(\alpha, \beta)$  of size symbols.  
 154 Intuitively, types represent sets of matrix dimensions. In particular, we obtain dimensions  
 155 from types by replacing size symbols by elements from  $\mathbb{N}_{>0}$ , where the size symbol 1 is always  
 156 replaced by the natural number 1. So,  $(\alpha, \beta)$  with  $\alpha \neq 1 \neq \beta$  represents the set of dimensions  
 157  $\{(m, n) \mid m, n \in \mathbb{N}_{>0}\}$ , while  $(\alpha, \alpha)$  represents the dimensions  $\{(m, m) \mid m \in \mathbb{N}_{>0}\}$  of square  
 158 matrices; and  $(\alpha, 1)$  represents the dimensions  $\{(m, 1) \mid m \in \mathbb{N}_{>0}\}$  of column vectors and  
 159  $(1, 1)$  represents the dimension  $(1, 1)$  of scalars.

160 **Schemas and instances.** We assume a set  $\mathcal{M} = \{\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{V}, \dots\}$  of *matrix symbols*, disjoint  
 161 with the size symbols and denoted by bold uppercase letters. Each matrix symbol  $\mathbf{A}$  has a  
 162 fixed associated type. We write  $\mathbf{A} : (\alpha, \beta)$  to denote that  $\mathbf{A}$  has type  $(\alpha, \beta)$ .

163 A matrix *schema*  $\mathcal{S}$  is a finite set of matrix and size symbols. We require that the special  
 164 size symbol 1 is always in  $\mathcal{S}$ , and that all size symbols occurring in the type of any matrix  
 165 symbol  $\mathbf{A} \in \mathcal{S}$  are also in  $\mathcal{S}$ . A matrix *instance*  $\mathcal{I}$  over a matrix schema  $\mathcal{S}$  is a function  
 166 that maps each size symbol  $\alpha$  in  $\mathcal{S}$  to a non-zero natural number  $\alpha^{\mathcal{I}} \in \mathbb{N}_{>0}$ , and maps each  
 167 matrix symbol  $\mathbf{A} : (\alpha, \beta)$  in  $\mathcal{S}$  to a  $K$ -matrix  $\mathbf{A}^{\mathcal{I}}$  of dimension  $\alpha^{\mathcal{I}} \times \beta^{\mathcal{I}}$ . We assume that for  
 168 the size symbol 1, we have  $1^{\mathcal{I}} = 1$ , for every instance  $\mathcal{I}$ .

169 **Sum-Matlang.** Let  $\mathcal{S}$  be a matrix schema. Before defining the syntax of sum-MATLANG,  
 170 we assume a set  $\mathcal{V} = \{\mathbf{u}, \mathbf{v}, \mathbf{w}, \mathbf{x}, \dots\}$  of *vector variables* over  $\mathcal{S}$ , which is disjoint with matrix  
 171 and size symbols in  $\mathcal{S}$ . Each such variable  $\mathbf{v}$  has a fixed associated type, which must be a  
 172 vector type  $(\gamma, 1)$  for some size symbol  $\gamma \in \mathcal{S}$ . We also write  $\mathbf{v} : (\gamma, 1)$  in that case.

173 The syntax of sum-MATLANG expressions [23] over  $\mathcal{S}$  is defined by the following grammar:

174	$e ::=$	$\mathbf{A} \in \mathcal{S}$	(matrix symbol)		$\mathbf{v} \in \mathcal{V}$	(vector variable)
		$e^T$	(transpose)		$e_1 \cdot e_2$	(matrix multiplication)
		$e_1 + e_2$	(matrix addition)		$e_1 \times e_2$	(scalar multiplication)
		$e_1 \odot e_2$	(pointwise multiplication)		$\Sigma \mathbf{v}.e$	(sum-iteration).

175 In addition, we require that expressions  $e$  are *well-typed*, in the sense that its *type*  $\text{type}(e)$  is

176 correctly defined as follows:

$$\begin{aligned}
& \text{type}(\mathbf{A}) := (\alpha, \beta) \text{ for a matrix symbol } \mathbf{A}: (\alpha, \beta) \\
& \text{type}(\mathbf{v}) := (\gamma, 1) \text{ for vector variable } \mathbf{v}: (\gamma, 1) \\
& \text{type}(e^T) := (\beta, \alpha) \text{ if } \text{type}(e) = (\alpha, \beta) \\
& \text{type}(e_1 \cdot e_2) := (\alpha, \gamma) \text{ if } \text{type}(e_1) = (\alpha, \beta) \text{ and } \text{type}(e_2) = (\beta, \gamma) \\
& \text{type}(e_1 + e_2) := (\alpha, \beta) \text{ if } \text{type}(e_1) = \text{type}(e_2) = (\alpha, \beta) \\
& \text{type}(e_1 \times e_2) := (\alpha, \beta) \text{ if } \text{type}(e_1) = (1, 1) \text{ and } \text{type}(e_2) = (\alpha, \beta) \\
& \text{type}(e_1 \odot e_2) := (\alpha, \beta) \text{ if } \text{type}(e_1) = \text{type}(e_2) = (\alpha, \beta) \\
& \text{type}(\Sigma \mathbf{v}.e) := (\alpha, \beta) \text{ if } e: (\alpha, \beta) \text{ and } \text{type}(\mathbf{v}) = (\gamma, 1).
\end{aligned}$$

178 In what follows, we always consider well-typed expressions and write  $e: (\alpha, \beta)$  to denote that  
179  $e$  is well typed, and its type is  $(\alpha, \beta)$ .

180 For an expression  $e$ , we say that a vector variable  $\mathbf{v}$  is *bound* if it is under a sum-iteration  
181  $\Sigma \mathbf{v}$ , and *free* otherwise. To evaluate expressions with free vector variables, we require the  
182 following notion of a valuation. Fix a matrix instance  $\mathcal{I}$  over  $\mathcal{S}$ . A *vector valuation* over  $\mathcal{I}$  is  
183 a function  $\mu$  that maps each vector symbol  $\mathbf{v}: (\gamma, 1)$  to a column vector of dimension  $\gamma^{\mathcal{I}} \times 1$ .  
184 Further, if  $\mathbf{b}$  is a vector of dimension  $\gamma^{\mathcal{I}} \times 1$ , then let  $\mu[\mathbf{v} := \mathbf{b}]$  denote the *extended* vector  
185 valuation over  $\mathcal{I}$  that coincides with  $\mu$ , except that  $\mathbf{v}: (\gamma, 1)$  is mapped to  $\mathbf{b}$ .

186 Let  $e: (\alpha, \beta)$  be a sum-MATLANG expression over  $\mathcal{S}$ . When one evaluates  $e$  over a matrix  
187 instance  $\mathcal{I}$  and a matrix valuation  $\mu$  over  $\mathcal{I}$ , it produces a matrix  $\llbracket e \rrbracket(\mathcal{I}, \mu)$  of dimension  
188  $\alpha^{\mathcal{I}} \times \beta^{\mathcal{I}}$  such that each entry  $i, j$  satisfies:

$$\begin{aligned}
& \llbracket \mathbf{A} \rrbracket(\mathcal{I}, \mu)_{ij} := \mathbf{A}_{ij}^{\mathcal{I}} \text{ for } \mathbf{A} \in \mathcal{S} \\
& \llbracket \mathbf{v} \rrbracket(\mathcal{I}, \mu)_{ij} := \mu(\mathbf{v})_{ij} \text{ for } \mathbf{v} \in \mathcal{V} \\
& \llbracket e^T \rrbracket(\mathcal{I}, \mu)_{ij} := \llbracket e \rrbracket(\mathcal{I}, \mu)_{ji} \\
& \llbracket e_1 + e_2 \rrbracket(\mathcal{I}, \mu)_{ij} := \llbracket e_1 \rrbracket(\mathcal{I}, \mu)_{ij} \oplus \llbracket e_2 \rrbracket(\mathcal{I}, \mu)_{ij} \\
& \llbracket e_1 \odot e_2 \rrbracket(\mathcal{I}, \mu)_{ij} := \llbracket e_1 \rrbracket(\mathcal{I}, \mu)_{ij} \odot \llbracket e_2 \rrbracket(\mathcal{I}, \mu)_{ij} \\
& \llbracket \Sigma \mathbf{v}.e \rrbracket(\mathcal{I}, \mu)_{ij} := \bigoplus_{k=1}^{\gamma^{\mathcal{I}}} \llbracket e \rrbracket(\mathcal{I}, \mu[\mathbf{v} := \mathbf{b}_k^{\gamma^{\mathcal{I}}}] )_{ij} \\
& \llbracket e_1 \cdot e_2 \rrbracket(\mathcal{I}, \mu)_{ij} := \bigoplus_k \llbracket e_1 \rrbracket(\mathcal{I}, \mu)_{ik} \odot \llbracket e_2 \rrbracket(\mathcal{I}, \mu)_{kj} \\
& \llbracket e_1 \times e_2 \rrbracket(\mathcal{I}, \mu)_{ij} := a \odot \llbracket e_2 \rrbracket(\mathcal{I}, \mu)_{ij} \text{ with } \llbracket e_1 \rrbracket(\mathcal{I}, \mu) = [a]
\end{aligned}$$

190 where  $\mathbf{v}: (\gamma, 1)$  and  $\mathbf{b}_1^n, \mathbf{b}_2^n, \dots, \mathbf{b}_n^n$  are the  $n$ -dimension canonical vectors, namely, the vectors  
191  $[\mathbf{1} \ \mathbf{0} \ \dots \ \mathbf{0}]^T, [\mathbf{0} \ \mathbf{1} \ \dots \ \mathbf{0}]^T, \dots, [\mathbf{0} \ \mathbf{0} \ \dots \ \mathbf{1}]^T$ , respectively.

192 ► **Example 1.** Let  $\mathcal{S} = \{\mathbf{A}\}$  where  $\mathbf{A}: (\alpha, \alpha)$ . Let  $\mathcal{I}$  be an instance over  $\mathcal{S}$  such that

$$193 \alpha^{\mathcal{I}} = 3 \text{ and } \mathbf{A}^{\mathcal{I}} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}. \text{ Let } \mathbf{v} \in \mathcal{V} \text{ where } \mathbf{v}: (\alpha, 1). \text{ The expression } \Sigma \mathbf{v}.\mathbf{A} \cdot \mathbf{v} \text{ is}$$

194 well-typed and

$$195 \llbracket \Sigma \mathbf{v}.\mathbf{A} \cdot \mathbf{v} \rrbracket(\mathcal{I}, \emptyset) = \mathbf{A} \cdot \mathbf{b}_1^3 + \mathbf{A} \cdot \mathbf{b}_2^3 + \mathbf{A} \cdot \mathbf{b}_3^3 = \begin{bmatrix} a_{11} \\ a_{21} \\ a_{31} \end{bmatrix} + \begin{bmatrix} a_{12} \\ a_{22} \\ a_{32} \end{bmatrix} + \begin{bmatrix} a_{13} \\ a_{23} \\ a_{33} \end{bmatrix}.$$

196 ► **Example 2.** Let  $\mathcal{S}$  and  $\mathcal{I}$  be as in Example 1. Let  $\mathbf{v} \in \mathcal{V}$  where  $\mathbf{v}: (\gamma, 1)$ . The expression  
197  $\Sigma \mathbf{v}.\mathbf{A}$  is well-typed and

$$198 \llbracket \Sigma \mathbf{v}.\mathbf{A} \rrbracket(\mathcal{I}, \emptyset) = \underbrace{\mathbf{A} + \dots + \mathbf{A}}_{\gamma^{\mathcal{I}} \text{ times}}.$$

## 12:6 Enumeration and Updates for Conjunctive Linear Algebra Queries

199 **Matlang.** MATLANG is a linear algebra query language that is a fragment of sum-MATLANG.  
 200 Specifically, define the syntax of MATLANG expressions over  $\mathcal{S}$  by the following grammar:

$$201 \quad e ::= \mathbf{A} \in \mathcal{S} \mid e^T \mid e_1 \cdot e_2 \mid e_1 + e_2 \mid e_1 \times e_2 \mid e_1 \odot e_2 \mid \mathbf{1}^\alpha \mid \mathbf{I}^\alpha$$

202 for every size symbol  $\alpha \in \mathcal{S}$ . Here,  $\mathbf{1}^\alpha$  and  $\mathbf{I}^\alpha$  denote the *ones-vector* and *identity-matrix*,  
 203 respectively, of type  $\text{type}(\mathbf{1}^\alpha) = (\alpha, 1)$  and  $\text{type}(\mathbf{I}^\alpha) = (\alpha, \alpha)$ . Their semantics can be  
 204 defined by using sum-MATLANG as:

$$205 \quad \llbracket \mathbf{1}^\alpha \rrbracket(\mathcal{I}, \mu) := \llbracket \Sigma \mathbf{v} \cdot \mathbf{v} \rrbracket(\mathcal{I}, \mu) \quad \llbracket \mathbf{I}^\alpha \rrbracket(\mathcal{I}, \mu) := \llbracket \Sigma \mathbf{v} \cdot \mathbf{v} \cdot \mathbf{v}^T \rrbracket(\mathcal{I}, \mu)$$

206 Note that the original MATLANG version introduced in [12] included an operator for the  
 207 *diagonalization of a vector*. This operator can be simulated by using the  $\mathbf{I}^\alpha$ -operator and  
 208 vice versa. Furthermore, we have included the pointwise multiplication  $\odot$  in MATLANG, also  
 209 known as the *Hadamard product*. This operation will be essential for our characterization  
 210 results. In [12, 23], the syntax of MATLANG was more generally parameterized by a family  
 211 of  $n$ -ary functions that could be pointwise applied. Similarly to [13, 22] we do not include  
 212 such functions here, but leave their detailed study to future work.

213 **Sum-Matlang queries.** A sum-MATLANG *query*  $\mathbf{Q}$  over a matrix schema  $\mathcal{S}$  is an expression  
 214 of the form  $\mathbf{H} := e$  where  $e$  is a well-typed sum-MATLANG expression without free vector  
 215 variables,  $\mathbf{H}$  is a “fresh” matrix symbol that does not occur in  $\mathcal{S}$ , and  $\text{type}(e) = \text{type}(\mathbf{H})$ .  
 216 When evaluated on a matrix instance  $\mathcal{I}$  over schema  $\mathcal{S}$ ,  $\mathbf{Q}$  returns a matrix instance  $\mathcal{E}$  over  
 217 the extended schema  $\mathcal{S} \cup \{\mathbf{H}\}$ :  $\mathcal{E}$  coincides with  $\mathcal{I}$  for every matrix and size symbol in  $\mathcal{S}$   
 218 and additionally maps  $\mathbf{H}^\mathcal{E} = \llbracket e \rrbracket(\mathcal{I}, \emptyset)$  with  $\emptyset$  denoting the empty vector valuation. We  
 219 denote the instance resulting from evaluating  $\mathbf{Q}$  by  $\llbracket \mathbf{Q} \rrbracket(\mathcal{I})$ . If  $\mathcal{S}$  is a matrix schema and  $\mathbf{Q}$  a  
 220 sum-MATLANG query over  $\mathcal{S}$  then we use  $\mathcal{S}(\mathbf{Q})$  to denote the extended schema  $\mathcal{S} \cup \{\mathbf{H}\}$ .

221  **$K$ -relations.** A  $K$ -relation over a domain of data values  $\mathbb{D}$  is a function  $f: \mathbb{D}^a \rightarrow K$  such  
 222 that  $f(\vec{d}) \neq \emptyset$  for finitely many  $\vec{d} \in \mathbb{D}^a$ . Here, ‘ $a$ ’ is the *arity* of  $R$ . Since we want to compare  
 223 relational queries with sum-MATLANG queries, we will restrict our attention in what follows  
 224 to  $K$ -relations where the domain  $\mathbb{D}$  of data values is the set  $\mathbb{N}_{>0}$ . In this context, we may  
 225 naturally view a  $K$ -matrix of dimensions  $n \times m$  as a  $K$ -relation such that the entry  $(i, j)$  of  
 226 the matrix is encoded by the  $K$ -value of the tuple  $(i, j)$  in the relation (see also Section 3).

227 **Vocabularies and databases.** We assume an infinite set of *relation symbols* together with  
 228 an infinite and disjoint set of *constant symbols*. Every relation symbol  $R$  is associated with a  
 229 number, its *arity*, which we denote by  $\text{ar}(R) \in \mathbb{N}$ . A *vocabulary*  $\sigma$  is a finite set of relation and  
 230 constant symbols. A *database* over  $\sigma$  is a function  $db$  that maps every constant symbol  $c \in \sigma$   
 231 to a value  $c^{db}$  in  $\mathbb{N}_{>0}$ ; and every relation symbol  $R \in \sigma$  to a  $K$ -relation  $R^{db}$  of arity  $\text{ar}(R)$ .

232 **Positive first order logic.** As our relational query language, we will work with the positive  
 233 fragment of first order logic ( $\text{FO}^+$ ). In contrast to the standard setting in database theory,  
 234 where the only atomic formulas are relational atoms of the form  $R(\vec{x})$ , we also allow the  
 235 ability to compare variables with constant symbols. To this end, the following definitions  
 236 are in order. We assume an infinite set of variables, which we usually denote by  $x, y, z$ . We  
 237 denote tuples of variables by  $\vec{x}, \vec{y}$ , and so on. A *relational atom* is expression of the form  
 238  $R(x_1, \dots, x_k)$  with  $R$  a relation symbol of arity  $k$ . A *comparison atom* is of the form  $x \leq c$   
 239 with  $x$  a variable and  $c$  a constant symbol. A *positive first order logic formula* ( $\text{FO}^+$  formula)  
 240 over a vocabulary  $\sigma$  is an expression generated by the following grammar:

$$241 \quad \varphi ::= R(\bar{x}) \mid x \leq c \mid \exists \bar{y}. \varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi$$

242 where  $R$  and  $c$  range over relations and constants in  $\sigma$ , respectively. We restrict ourselves to  
 243 *safe* formulas, in the sense that in a disjunction  $\varphi_1 \vee \varphi_2$ , we require that  $\varphi_1$  and  $\varphi_2$  have the  
 244 same set of free variables [2]. The notions of *free and bound variables* are defined as usual  
 245 in FO. We denote the set of free variables of  $\varphi$  by  $\text{free}(\varphi)$  and its multiset of atoms by  $\text{at}(\varphi)$ .

246 We evaluate formulas over  $K$ -relations as follows using the well-known semiring seman-  
 247 tics [26]. A *valuation* over a set of variables  $X$  is a function  $\nu: X \rightarrow \mathbb{N}_{>0}$  that assigns a  
 248 value in  $\mathbb{N}_{>0}$  to each variable in  $X$  (recall that  $\mathbb{D} = \mathbb{N}_{>0}$ ). We denote by  $\nu: X$  that  $\nu$  is a  
 249 valuation on  $X$  and by  $\nu|_Y$  the restriction of  $\nu$  to  $X \cap Y$ . As usual, valuations are extended  
 250 point-wise to tuples of variables, i.e.,  $\nu(x_1, \dots, x_n) = (\nu(x_1), \dots, \nu(x_n))$ . Let  $\varphi$  be an  $\text{FO}^+$   
 251 formula over vocabulary  $\sigma$ . When evaluated on a database  $db$  over vocabulary  $\sigma$ , it defines a  
 252 mapping  $\llbracket \varphi \rrbracket_{db}$  from valuations over  $\text{free}(\varphi)$  to  $K$  inductively defined as:

$$\begin{aligned} \llbracket R(x_1, \dots, x_k) \rrbracket_{db}(\nu) &:= R^{db}(\nu(x_1), \dots, \nu(x_k)) \\ \llbracket x \leq c \rrbracket_{db}(\nu) &:= \begin{cases} 1 & \text{if } \nu(x) \leq c^{db} \\ 0 & \text{otherwise} \end{cases} \\ \llbracket \varphi_1 \wedge \varphi_2 \rrbracket_{db}(\nu) &:= \llbracket \varphi_1 \rrbracket_{db}(\nu|_{\text{free}(\varphi_1)}) \odot \llbracket \varphi_2 \rrbracket_{db}(\nu|_{\text{free}(\varphi_2)}) \\ \llbracket \varphi_1 \vee \varphi_2 \rrbracket_{db}(\nu) &:= \llbracket \varphi_1 \rrbracket_{db}(\nu) \oplus \llbracket \varphi_2 \rrbracket_{db}(\nu) \\ \llbracket \exists y. \varphi \rrbracket_{db}(\nu) &:= \bigoplus_{\mu: \text{free}(\varphi) \text{ s.t. } \mu|_{\text{free}(\varphi) \setminus \{y\}} = \nu} \llbracket \varphi \rrbracket_{db}(\mu) \end{aligned}$$

254 **FO queries.** An  $\text{FO}^+$  *query*  $Q$  over vocabulary  $\sigma$  is an expression of the form  $H(\bar{x}) \leftarrow \varphi$   
 255 where  $\varphi$  is an  $\text{FO}^+$  formula over  $\sigma$ ,  $\bar{x} = (x_1, \dots, x_k)$  is a sequence of (not necessarily distinct)  
 256 free variables of  $\varphi$ , such that every free variable of  $\varphi$  occurs in  $\bar{x}$ , and  $H$  is a “fresh” relation  
 257 symbol not in  $\sigma$  with  $\text{ar}(H) = k$ . The formula  $\varphi$  is called the *body* of  $Q$ , and  $H(\bar{x})$  its *head*.

258 When evaluated over a database  $db$  over  $\sigma$ ,  $Q$  returns a database  $\llbracket Q \rrbracket_{db}$  over the extended  
 259 vocabulary  $\sigma \cup \{H\}$ . This database  $\llbracket Q \rrbracket_{db}$  coincides with  $db$  for every relation and constant  
 260 symbol in  $\sigma$ , and maps the relation symbol  $H$  to the  $K$ -relation of arity  $k$  defined as follows.  
 261 For a sequence of domain values  $\bar{d} = (d_1, \dots, d_k)$ , we write  $\bar{d} \models \bar{x}$  if, for all  $i \neq j$  with  
 262  $x_i = x_j$  we also have  $d_i = d_j$ . Clearly, if  $\bar{d} \models \bar{x}$  then the mapping  $\{x_1 \rightarrow d_1, \dots, x_k \rightarrow d_k\}$  is  
 263 well-defined. Denote this mapping by  $\bar{x} \mapsto \bar{d}$  in this case. Then

$$264 \quad \llbracket Q \rrbracket_{db}(H) := \bar{d} \mapsto \begin{cases} \llbracket \varphi \rrbracket_{db}(\bar{x} \mapsto \bar{d}) & \text{if } \bar{d} \models \bar{x} \\ 0 & \text{otherwise} \end{cases}$$

265 In what follows, if  $Q$  is a query, then we will often use the notation  $Q(\bar{x})$  to denote that the  
 266 sequence of the variables in the head of  $Q$  is  $\bar{x}$ . If  $Q$  is a query over  $\sigma$  and  $H(\bar{x})$  its head,  
 267 then we write  $\sigma(Q)$  for the extended vocabulary  $\sigma \cup \{H\}$ .

268 We denote by  $\text{FO}^\wedge$  the fragment of  $\text{FO}^+$  formulas in which disjunction is disallowed. A  
 269 query  $Q = H(\bar{x}) \leftarrow \varphi$  is an  $\text{FO}^\wedge$  query if  $\varphi$  is in  $\text{FO}^\wedge$ . If additionally  $\varphi$  is in prenex normal  
 270 form, i.e.,  $Q: H(\bar{x}) \leftarrow \exists \bar{y}. a_1 \wedge \dots \wedge a_n$  with  $a_1, \dots, a_n$  (relational or comparison) atoms,  
 271 then  $Q$  is a *conjunctive query* (CQ). Note that, while the classes of conjunctive queries and  
 272  $\text{FO}^\wedge$  queries are equally expressive, for our purposes conjunctive queries are hence formally a  
 273 syntactic fragment of  $\text{FO}^\wedge$  queries.

274 An  $\text{FO}^+$  query is *binary* if every relational atom occurring in it (body and head) has  
 275 arity at most two. Because in sum-MATLANG both the input and output are matrices, our  
 276 correspondences between sum-MATLANG and  $\text{FO}^+$  will focus on binary queries.

277 **Discussion.** We have added comparison atoms to  $\text{FO}^+$  in order to establish its correspon-  
 278 dence with  $\text{sum-MATLANG}$ . To illustrate why we will need comparison atoms, consider the  
 279  $\text{sum-MATLANG}$  expression  $\mathbf{I}^\alpha$  of type  $(\alpha, \alpha)$  that computes the identity matrix. This can  
 280 be expressed by means of the following CQ  $Q : I(x, x) \leftarrow x \leq \alpha$ . We hence use comparison  
 281 atoms to align the dimension of the matrices with the domain size of relations.

282 To make the correspondence hold, we note that in  $\text{MATLANG}$  there is a special size  
 283 symbol,  $1$ , which is always interpreted as the constant  $1 \in \mathbb{N}$ . This size symbol is used in  
 284 particular to represent column and row vectors, which have type  $(\alpha, 1)$  and  $(1, \alpha)$  respectively.  
 285 We endow CQs with the same property in the sense that we will assume in what follows that  
 286  $1$  is a valid constant symbol and that  $1^{db} = 1$  for every database  $db$ .

### 287 3 From matrices to relations and back

288 Geerts et al. [23] previously established a correspondence between  $\text{sum-MATLANG}$  and  $\text{FO}^+$ .  
 289 However, as we illustrate in the full version [38] their correspondence is (1) for square matrices,  
 290 (2) asymmetric, and (3) encodes matrix instances as databases of more than linear size,  
 291 making it unsuitable to derive the complexity bounds that we are interested in here. In  
 292 this section, we revisit and generalize the connection between  $\text{sum-MATLANG}$  and  $\text{FO}^+$  by  
 293 providing translations between the two query languages that works for any matrix schema,  
 294 are symmetric, and ensure that matrices are encoded as databases of linear size. Towards  
 295 this goal, we introduce next all the formal machinery to link both settings. We start by  
 296 determining precisely in what sense relations can encode matrices, or matrices can represent  
 297 relations, and how this correspondence transfers to queries. Then we show how to generalize  
 298 the expressibility results in [23] for any matrix sizes and every encoding between schemas.

299 **How we relate objects.** Let  $\mathbf{A}$  be a matrix of dimension  $m \times n$ . There exist multiple  
 300 natural ways to encode  $\mathbf{A}$  as a relation, depending on the dimension  $m \times n$ .

- 301 ■ We can always encode  $\mathbf{A}$ , whatever the values of  $m$  and  $n$ , as the binary  $K$ -relation  $R$   
 302 such that (1)  $\mathbf{A}_{i,j} = R(i, j)$  for every  $i \leq m, j \leq n$  and (2)  $R(i, j) = \mathbb{0}$  if  $i > m$  or  $j > n$ .
  - 303 ■ If  $\mathbf{A}$  is a column vector ( $n = 1$ ) then we can also encode it as the unary  $K$ -relation  $R$   
 304 such that  $\mathbf{A}_{i,1} = R(i)$  for every  $i \leq m$  and  $R(i) = \mathbb{0}$  if  $i > m$ .
  - 305 ■ Similarly, if  $\mathbf{A}$  is a row vector ( $m = 1$ ) then we can encode it as the unary  $K$ -relation  $R$   
 306 with  $\mathbf{A}_{1,j} = R(j)$  for every  $j \leq n$  and  $R(i) = \mathbb{0}$  if  $j > n$ .
  - 307 ■ If  $\mathbf{A}$  is a scalar ( $m = n = 1$ ), we can encode it as a nullary  $K$ -relation  $R$  with  $\mathbf{A}_{1,1} = R()$ .
- 308 Note that if  $\mathbf{A}$  is scalar then we can hence encode it by means of a binary relation, a unary  
 309 relation, or a nullary relation; and if it is a vector we can encode it by a binary or unary  
 310 relation. In what follows, we write  $\mathbf{A} \simeq R$  to denote that  $R$  encodes  $\mathbf{A}$ .

311 Conversely, given a (nullary, unary, or binary)  $K$ -relation  $R$  we may interpret this as  
 312 a matrix of appropriate dimension. Specifically, we say that relation  $R$  is *consistent* with  
 313 dimension  $m \times n$  if there exists a matrix  $\mathbf{A}$  of dimension  $m \times n$  such that  $\mathbf{A} \simeq R$ . This is  
 314 equivalent to requiring that relation is  $\mathbb{0}$  on entries outside of  $m \times n$ . Note that, given  $R$   
 315 that is consistent with  $m \times n$  there is exactly one matrix  $\mathbf{A} : m \times n$  such that  $\mathbf{A} \simeq R$ .

316 **How we relate schemas.** A *matrix-to-relational schema encoding* from a matrix schema  $\mathcal{S}$   
 317 to a relational vocabulary  $\sigma$  is a function  $\text{Rel} : \mathcal{S} \rightarrow \sigma$  that maps every matrix symbol  $\mathbf{A}$  in  $\mathcal{S}$   
 318 to a unary or binary relation symbol  $\text{Rel}(\mathbf{A})$  in  $\sigma$ , and every size symbol  $\alpha$  in  $\mathcal{S}$  to a constant  
 319 symbol  $\text{Rel}(\alpha)$  in  $\sigma$ . Here,  $\text{Rel}(\mathbf{A})$  can be unary only if  $\mathbf{A}$  is of vector type, and nullary only  
 320 if  $\mathbf{A}$  is of scalar type. Intuitively,  $\text{Rel}$  specifies which relation symbols will be used to store

the encodings of which matrix symbols. In addition, we require that  $Rel(1) = 1$  and that  $Rel$  is a bijection between  $\mathcal{S}$  and  $\sigma$ . This makes sure that we can always invert  $Rel$ . In what follows, we will only specify that  $Rel$  is a matrix-to-relational schema encoding on  $\mathcal{S}$ , leaving the vocabulary  $\sigma$  unspecified. In that case, we write  $Rel(\mathcal{S})$  for the relational vocabulary  $\sigma$ .

Conversely, we define a *relational-to-matrix schema encoding* from  $\sigma$  into  $\mathcal{S}$  as a function  $Mat: \sigma \rightarrow \mathcal{S}$  that maps every relation symbol  $R$  to a matrix symbol  $Mat(R)$  and every constant symbol  $c$  to a size symbol  $Mat(c)$ . We require that all unary relations are mapped to matrix symbols of vector type, either row or column, and that all nullary relations are mapped to matrix symbols of scalar type. Furthermore,  $Mat$  must map  $1 \mapsto 1$  and be bijective. Similarly, we denote by  $Mat(\sigma)$  the matrix schema  $\mathcal{S}$  mapped by  $Mat$ .

Note that the bijection assumption over  $Mat$  imposes some requirements over  $\sigma$  to encode it as matrices. For example,  $Mat$  requires the existence of at least one constant symbol in  $\sigma$  for encoding matrices dimensions, since every matrix symbol has at least one size symbol in its type, and that size symbol is by definition in  $\mathcal{S}$ . The bijection between constant and size symbols is necessary in order to have lossless encoding between both settings.

Given that  $Rel$  and  $Mat$  are bijections between  $\mathcal{S}$  and  $\sigma$ , their inverses  $Rel^{-1}$  and  $Mat^{-1}$  are well defined. Furthermore, by definition we have that  $Rel^{-1}$  and  $Mat^{-1}$  are relational-to-matrix and matrix-to-relational schema encodings, respectively.

**How we relate instances.** We start by specifying how to encode matrix instances as database instances. Fix a matrix-to-relational schema encoding  $Rel$  between  $\mathcal{S}$  and  $Rel(\mathcal{S})$ . Let  $\mathcal{I}$  be a matrix instance over  $\mathcal{S}$  and  $db$  a database over  $Rel(\mathcal{S})$ . We say that  $db$  is a *relational encoding of  $\mathcal{I}$  w.r.t.  $Rel$* , denoted by  $\mathcal{I} \simeq_{Rel} db$ , if

- $\mathbf{A}^{\mathcal{I}} \simeq Rel(\mathbf{A})^{db}$  for every matrix symbol  $\mathbf{A}$  in  $\mathcal{S}$ , and
- $Rel(\alpha)^{db} = \alpha^{\mathcal{I}}$  for every size symbol  $\alpha$  in  $\mathcal{S}$ .

Note that, given  $\mathcal{I}$  and  $Rel$ , the relational encoding  $db$  is uniquely defined. As such, we also denote this database by  $Rel(\mathcal{I})$ .

We now focus on interpreting database instances as matrix instances, which is more subtle. Fix a relational-to-matrix schema encoding  $Mat$  from  $\sigma$  to  $Mat(\sigma)$ . We need to first leverage the consistency requirement from relations to databases. Formally, we say that a database  $db$  over  $\sigma$  is *consistent with  $Mat$*  if for every relation symbol  $R$  in  $\sigma$ ,  $R^{db}$  is consistent with dimension  $c^{db} \times d^{db}$  where  $Mat(R): (Mat(c), Mat(d))$ . In other words, a consistent database specifies the value of each dimension, and the relations are themselves consistent with them.

Let  $db$  be a database over  $\sigma$ , consistent with  $Mat$  and let  $\mathcal{I}$  be a matrix instance of  $Mat(\sigma)$ . We say that  $\mathcal{I}$  is a *matrix encoding of  $db$  w.r.t.  $Mat$* , denoted  $db \simeq_{Mat} \mathcal{I}$ , if

- $Mat(R)^{\mathcal{I}} \simeq R^{db}$  for every relation symbol  $R \in \sigma$ ; and
- $c^{db} = Mat(c)^{\mathcal{I}}$  for every constant symbol  $c \in \sigma$ .

Given  $Mat$  and a consistent database  $db$ , the matrix encoding  $\mathcal{I}$  is uniquely defined. As such, we also denote this instance by  $Mat(db)$ .

From the previous definitions, one notes an asymmetry between both directions. Although an encoding always holds from matrices to relations, we require that the relations are consistent with the sizes (i.e., constants) from relations to matrices. Nevertheless, this asymmetry does not impose a problem when we want to go back and forth, as the next result shows.

► **Proposition 3.** *Let  $Rel$  and  $Mat$  be matrix-to-relational and relational-to-matrix schema encodings from  $\mathcal{S}$  to  $\sigma$  and from  $\sigma$  to  $\mathcal{S}$ , respectively, such that  $Mat = Rel^{-1}$ . Then*

- $Rel^{-1}(Rel(\mathcal{S})) = \mathcal{S}$  and  $Mat^{-1}(Mat(\sigma)) = \sigma$ ;
- $Rel(\mathcal{I})$  is consistent with  $Rel^{-1}$ , for every instance  $\mathcal{I}$  over  $\mathcal{S}$ ;
- $Rel^{-1}(Rel(\mathcal{I})) = \mathcal{I}$ , for every instance  $\mathcal{I}$  over  $\mathcal{S}$ ; and

## 12:10 Enumeration and Updates for Conjunctive Linear Algebra Queries

368 ■  $Mat^{-1}(Mat(db)) = db$ , for every  $db$  consistent with  $Mat$ .

369 The previous proposition is a direct consequence of the definitions; however, it shows that  
 370 the consistency requirement and schema encodings provide a lossless encoding between the  
 371 relational and matrix settings. This fact is crucial to formalize the expressiveness equivalence  
 372 between **sum-MATLANG** and  $FO^+$ , and their subfragments in the following sections.

373 **From sum-Matlang to positive-FO.** We first aim to simulate every **sum-MATLANG** query  
 374 with an  $FO^+$  query w.r.t. some matrix-to-relational schema encoding. This was already  
 375 proven in [23] for a different but related setting, and only for *Rel* encodings that map matrix  
 376 symbols with vector types into unary relations. Here we generalize it to arbitrary encodings.

377 In what follows, fix a matrix schema  $\mathcal{S}$ . Let  $\mathbf{Q}$  be a **sum-MATLANG** query over  $\mathcal{S}$ , and  
 378 *Rel* be a matrix-to-relational schema encoding on  $\mathcal{S}(\mathbf{Q})$ . We say that  $FO^+$  query  $Q$  *simulates*  
 379  $\mathbf{Q}$  w.r.t. *Rel* if  $Rel(\llbracket \mathbf{Q} \rrbracket(\mathcal{I})) = \llbracket Q \rrbracket_{Rel(\mathcal{I})}$  for every matrix instance  $\mathcal{I}$  over  $\mathcal{S}$ . Note that the  
 380 definition implies that the output matrix symbol of  $\mathbf{Q}$  must be mapped to the output relation  
 381 symbol of  $Q$  by *Rel*, since *Rel* is a bijection and the condition must hold for every matrix  
 382 instance. Indeed, it is equivalent to  $\llbracket \mathbf{Q} \rrbracket(\mathcal{I}) = Rel^{-1}(\llbracket Q \rrbracket_{Rel(\mathcal{I})})$ , namely, that one can evaluate  
 383  $\mathbf{Q}$  by first evaluating  $\llbracket Q \rrbracket_{Rel(\mathcal{I})}$  and then mapping the results back.

384 Next, we show that we can simulate every **sum-MATLANG** query in the relational setting.

385 ► **Proposition 4.** *For every sum-MATLANG query  $\mathbf{Q}$  over  $\mathcal{S}$  and every matrix-to-relational  
 386 schema encoding *Rel* on  $\mathcal{S}(\mathbf{Q})$ , there exists an  $FO^+$  query  $Q$  that simulates  $\mathbf{Q}$  w.r.t. *Rel*.*

387 **From positive-FO to sum-Matlang.** We now aim to simulate every  $FO^+$  query with a  
 388 **sum-MATLANG** query. Contrary to the previous direction, the expressiveness result here is  
 389 more subtle and requires more discussion and additional notions.

390 Fix a vocabulary  $\sigma$ . Let  $Q$  be a  $FO^+$  query over  $\sigma$  and let *Mat* be relational-to-matrix  
 391 schema encoding on  $\sigma(Q)$ . We say that a matrix query  $\mathbf{Q}$  *simulates*  $Q$  w.r.t. *Mat* if  
 392  $Mat(\llbracket \mathbf{Q} \rrbracket_{db}) = \llbracket Q \rrbracket(Mat(db))$  for every database  $db$  consistent with *Mat*. We note again that  
 393 this definition implies that the input vocabulary and output relation symbol of  $Q$  coincides  
 394 with the input schema and output matrix symbol of  $\mathbf{Q}$ , respectively. Further, it is equivalent  
 395 that  $\llbracket \mathbf{Q} \rrbracket_{db} = Mat^{-1}(\llbracket \mathbf{Q} \rrbracket(Mat(db)))$ .

396 Before stating how to connect  $FO^+$  with **sum-MATLANG**, we need to overcome the  
 397 following problem: a  $FO^+$  query can use the same variable within different relational atoms,  
 398 which can be mapped to matrix symbols of different types. For an illustrative example of  
 399 this problem, consider the query

$$400 \quad Q: H(x, y) \leftarrow R(x, y), S(y, z)$$

401 and a relational-to-matrix schema encoding such that *Mat* maps  $R$  and  $S$  to symbols of type  
 402  $(\alpha, \beta)$ ,  $H$  to a symbol of type  $(\beta, \beta)$ , and  $c$  and  $d$  to  $\alpha$  and  $\beta$ , respectively. For a consistent  
 403 database  $db$  w.r.t. *Mat*, we could have that  $R$  and  $S$  are consistent with  $c^{db} \times d^{db}$ , but  $H$  is  
 404 not consistent with  $d^{db} \times d^{db}$  if  $d^{db} < c^{db}$ . Moreover, *Mat* bounds variable  $y$  with different  
 405 sizes  $c^{db}$  and  $d^{db}$ . It is then problematic to simulate  $Q$  under *Mat* in **sum-MATLANG** because  
 406 **sum-MATLANG** expressions need to be well-typed.

407 Given the previous discussion, the *well-typedness* definition of a  $FO^+$  formula is necessary.  
 408 Let *Mat* be a relational-to-matrix schema encoding on  $\sigma$ . Given a  $FO^+$  formula  $\varphi$  over  $\sigma$  and  
 409 a function  $\tau$  from  $free(\varphi)$  to size symbols in  $Mat(\sigma)$ , define the rule  $Mat \vdash \varphi: \tau$  inductively as  
 410 shown in Figure 1, where  $\tau_1 \sim \tau_2$  if and only if  $\tau_1(x) = \tau_2(x)$  for every  $x \in dom(\tau_1) \cap dom(\tau_2)$ .

$$\begin{array}{c}
\frac{\text{type}(\text{Mat}(R)) = (\alpha, \beta)}{\text{Mat} \vdash R(x_1, x_2) : \{x_1 \mapsto \alpha, x_2 \mapsto \beta\}} \qquad \frac{\text{type}(\text{Mat}(R)) = (\alpha, 1) \text{ or } (1, \alpha)}{\text{Mat} \vdash R(x) : \{x \mapsto \alpha\}} \\
\frac{\text{type}(\text{Mat}(R)) = (1, 1)}{\text{Mat} \vdash R() : \{\}} \qquad \frac{}{\text{Mat} \vdash x \leq c : \{x \mapsto \text{Mat}(c)\}} \qquad \frac{\text{Mat} \vdash \varphi : \tau}{\text{Mat} \vdash \exists \bar{y}. \varphi : \tau|_{\text{free}(\varphi) \setminus \bar{y}}} \\
\frac{\text{Mat} \vdash \varphi_1 : \tau_1, \text{Mat} \vdash \varphi_2 : \tau_2 \text{ and } \tau_1 \sim \tau_2}{\text{Mat} \vdash \varphi_1 \wedge \varphi_2 : \tau_1 \cup \tau_2} \qquad \frac{\text{Mat} \vdash \varphi_1 : \tau_1, \text{Mat} \vdash \varphi_2 : \tau_2 \text{ and } \tau_1 \sim \tau_2}{\text{Mat} \vdash \varphi_1 \vee \varphi_2 : \tau_1 \cup \tau_2}
\end{array}$$

■ **Figure 1** Well-typedness of  $\text{FO}^+$  formulas under relational-to-matrix mapping  $\text{Mat}$ .

411 We say that  $\varphi$  over  $\sigma$  is *well-typed* w.r.t.  $\text{Mat}$  if there exists such a function  $\tau$  such that  
412  $\text{Mat} \vdash \varphi : \tau$ . Note that if  $\varphi$  is well-typed, then there is a unique  $\tau$  such that  $\text{Mat} \vdash \varphi : \tau$ .

413 Now, let  $Q: H(\bar{x}) \leftarrow \varphi$  be a binary  $\text{FO}^+$  query over  $\sigma$  and  $\text{Mat}$  be a matrix encoding  
414 specification on  $\sigma(Q)$ . We say that  $Q$  is *well-typed* w.r.t.  $\text{Mat}$  if  $\varphi$  is well-typed w.r.t.  $\text{Mat}$   
415 and for  $\tau$  such that  $\text{Mat} \vdash \varphi : \tau$ , we have:

- 416 ■ if  $\bar{x} = (x_1, x_2)$ , then  $\text{type}(\text{Mat}(H)) = (\tau(x_1), \tau(x_2))$ ; or
- 417 ■ if  $\bar{x} = (x)$ , then  $\text{type}(\text{Mat}(H))$  is either  $(\tau(x), 1)$  or  $(1, \tau(x))$ .

418 We write  $\text{Mat} \vdash Q : \tau$  to indicate that  $Q$  is well-typed w.r.t.  $\text{Mat}$ , and  $\tau$  is the unique function  
419 testifying to well-typedness of  $\text{FO}^+$  formula of  $Q$ . We note that that we can show that the  
420 query obtained by Proposition 4 is always well-typed.

421 The next proposition connects well-typedness with consistency.

422 ► **Proposition 5.** *For binary  $\text{FO}^+$  query  $Q: H(\bar{x}) \leftarrow \varphi$  over a vocabulary  $\sigma$ , if  $\text{Mat} \vdash Q : \tau$   
423 then for any  $db$  consistent with  $\text{Mat}$  we have:*

- 424 ■ If  $\bar{x} = (x_1, x_2)$  then  $\llbracket Q \rrbracket_{db}(H)$  is consistent with dimension  $\tau(x_1)^{db} \times \tau(x_2)^{db}$ .
- 425 ■ If  $\bar{x} = (x)$  then  $\llbracket Q \rrbracket_{db}(H)$  is consistent with both dimension  $\tau(x)^{db} \times 1$  and  $1 \times \tau(x)^{db}$ .

426 We have now all the formal machinery to state how to simulate every  $\text{FO}^+$  query over  
427 relations with a sum-MATLANG query over matrices.

428 ► **Proposition 6.** *For every binary  $\text{FO}^+$  query  $Q$  over a vocabulary  $\sigma$  and every relational-  
429 to-matrix schema encoding  $\text{Mat}$  on  $\sigma(Q)$  such that  $Q$  is well typed w.r.t.  $\text{Mat}$  there exists a  
430 sum-MATLANG query  $\mathbf{Q}$  that simulates  $Q$  w.r.t.  $\text{Mat}$ .*

431 **Conjunctive Matlang.** Taking into account the correspondence between sum-MATLANG  
432 and  $\text{FO}^+$  established by Propositions 5 and 6, in what follows we say that matrix query  
433 language  $\mathcal{L}_M \subseteq \text{sum-MATLANG}$  and relational language  $\mathcal{L}_R \subseteq \text{FO}^+$  are *equivalent* or *equally*  
434 *expressive* if (1) for every matrix query  $\mathbf{Q} \in \mathcal{L}_M$  over a matrix schema  $\mathcal{S}$  and every matrix-to-  
435 relational schema encoding  $\text{Rel}$  on  $\mathcal{S}(\mathbf{Q})$  there exists a query  $Q \in \mathcal{L}_R$  that simulates  $\mathbf{Q}$  w.r.t.  
436  $\text{Rel}$  and is well-typed w.r.t.  $\text{Rel}^{-1}$ ; and (2) for every binary query  $Q \in \mathcal{L}_R$  over a vocabulary  
437  $\sigma$  and every relational-to-matrix schema encoding  $\text{Mat}$  such that  $Q$  is well-typed w.r.t.  $\text{Mat}$   
438 there exists  $\mathbf{Q} \in \mathcal{L}_M$  that simulates  $Q$  w.r.t.  $\text{Mat}$ .

439 Let conj-MATLANG be the sum-MATLANG fragment that includes all operations except  
440 matrix addition (+). Then we can derive the following characterization of CQs.

441 ► **Corollary 7.** *conj-MATLANG and conjunctive queries are equally expressive.*

442 While this result is a consequence of the connection between sum-MATLANG and  $\text{FO}^+$ , it  
443 provides the basis to explore the fragments of conj-MATLANG that correspond to fragments of  
444 CQ, like free-connex or q-hierarchical CQ. We determine these fragments in the next sections.

445 **4 The fragment of free-connex queries**

446 In this section, we specialize the correspondence of Corollary 7 between conj-MATLANG and  
 447 CQs to *free-connex* CQs [5]. Free-connex CQs are a subset of acyclic CQs that allow efficient  
 448 enumeration-based query evaluation: in the Boolean semiring and under data complexity they  
 449 allow to enumerate the query result  $\llbracket Q \rrbracket_{db}(H)$  of free-connex CQ  $Q: H(\bar{x}) \leftarrow \varphi$  with *constant*  
 450 *delay* after a preprocessing phase that is linear in  $db$ . In fact, under complexity-theoretic  
 451 assumptions, the class of CQs that admits constant delay enumeration after linear time  
 452 preprocessing is precisely the class of free-connex CQs [5].

453 **Acyclic and free-connex CQs.** A CQ  $Q: H(\bar{x}) \leftarrow \exists \bar{y}. a_1 \wedge \dots \wedge a_n$  is called *acyclic* [7,10,21]  
 454 if it has a *join-tree*, i.e. an undirected tree  $T = (V, E)$  with  $V$  the set  $\{a_1, \dots, a_n\}$  of atoms  
 455 in the body and where for each variable  $z$  occurring in  $Q$  the set  $\{a_i \in V \mid z \in \text{vars}(a_i)\}$   
 456 induces a connected subtree of  $T$ . Note that we consider inequality predicates as unary.  
 457 Furthermore,  $Q$  is *free-connex* [5,11] if it is acyclic and the query  $Q'$  obtained by adding  
 458 the head atom  $H(\bar{x})$  to the body of  $Q$  is also acyclic. The second condition forbids queries  
 459 like  $H(x, y) \leftarrow \exists z. A(x, z) \wedge B(z, y)$  since when we adjoin the head to the body we get  
 460  $\exists z. A(x, z) \wedge B(z, y) \wedge H(x, y)$ , which is cyclic. We will usually refer to free-connex CQs  
 461 simply as fc-CQs in what follows.

462 To identify the sum-MATLANG fragment that corresponds to fc-CQs, we find it convenient  
 463 to first observe the following correspondence between fc-CQs and  $\text{FO}_2^\wedge$ , the two-variable  
 464 fragment of  $\text{FO}^\wedge$ . Here, a formula  $\varphi$  in  $\text{FO}^\wedge$  is said to be in  $\text{FO}_2^\wedge$  if the set of all variables  
 465 used in  $\varphi$  (free or bound) is of cardinality at most two. So,  $\exists y \exists z. A(x, y) \wedge B(y, z)$  is not  
 466 in  $\text{FO}_2^\wedge$ , but the equivalent formula  $\exists y. (A(x, y) \wedge \exists x. B(y, x))$  is. An  $\text{FO}_2^\wedge$  query is a binary  
 467 query whose body is an  $\text{FO}_2^\wedge$  formula. Recall that a query is binary if every relational atom  
 468 occurring in its body and head have arity at most two.

469 **► Theorem 8.** *Binary free-connex CQs and  $\text{FO}_2^\wedge$  queries are equally expressive.*

470 We find this a remarkable characterization of the fc-CQs on binary relations that, to the  
 471 best of our knowledge, it is new. Moreover, this result motivates the fragment of MATLANG  
 472 that characterizes fc-CQs.

473 **Free-connex MATLANG.** Define fc-MATLANG to be the class of all MATLANG expressions  
 474 generated by the grammar:

$$475 \quad e ::= \mathbf{A} \mid \mathbf{1}^\alpha \mid \mathbf{I}^\alpha \mid e^T \mid e_1 \times e_2 \mid e_1 \odot e_2 \mid e_1 \cdot v_2 \mid v_1 \cdot e_2$$

476 where  $v_1$  and  $v_2$  are fc-MATLANG expressions with type  $(\alpha, 1)$  or  $(1, \alpha)$ . In other words,  
 477 matrix multiplication  $e_1 \cdot e_2$  is only allowed when at least one of  $e_1$  or  $e_2$  has a row or column  
 478 vector type.

479 Interestingly, we are able to show that  $\text{FO}_2^\wedge$  and fc-MATLANG are equally expressive.

480 **► Theorem 9.** *fc-MATLANG and  $\text{FO}_2^\wedge$  are equally expressive.*

481 From Theorem 8 and Theorem 9 we obtain:

482 **► Corollary 10.** *fc-MATLANG and binary free-connex CQs are equally expressive.*

## 5 The fragment of q-hierarchical queries

483

484 We next specialize the correspondence between conj-MATLANG and CQs to *q-hierarchical*  
 485 CQs [5]. The q-hierarchical CQs form a fragment of the free-connex CQs that, in addition to  
 486 supporting constant delay enumeration after linear time preprocessing, have the property  
 487 that every single-tuple update (insertion or deletion) to the input database can be processed  
 488 in constant time, after which the enumeration of the updated query result can again proceed  
 489 with constant delay [9].

490 **Q-hierarchical CQs.** Let  $Q: H(\bar{x}) \leftarrow \exists \bar{y}. a_1 \wedge \dots \wedge a_n$  be a CQ. For every variable  $x$ , define  
 491  $at(x)$  to be the set  $\{a_i \mid x \in \text{var}(a_i)\}$  of relational atoms that mention  $x$ . Note that, contrary  
 492 to acyclic queries, here we make the distinction between relational atoms and inequalities.  
 493 Then  $Q$  is q-hierarchical if for any two variables  $x, y$  the following is satisfied:

- 494 1.  $at(x) \subseteq at(y)$  or  $at(x) \supseteq at(y)$  or  $at(x) \cap at(y) = \emptyset$ , and
- 495 2. if  $x \in \bar{x}$  and  $at(x) \subsetneq at(y)$  then  $y \in \bar{x}$ .

496 For example,  $H(x) \leftarrow \exists y. A(x, y) \wedge U(x)$  is q-hierarchical. By contrast, the variant  $H(x) \leftarrow$   
 497  $\exists y. A(x, y) \wedge U(y)$  is not q-hierarchical, as it violates the second condition. Furthermore,  
 498  $H(x, y) \leftarrow A(x, y) \wedge U(x) \wedge V(y)$  violates the first condition, and is also not q-hierarchical.  
 499 Note that all these examples are free-connex. We refer to q-hierarchical CQs simply as qh-CQs.

500 **Q-hierarchical Matlang.** The fragment of sum-MATLANG that is equivalent to qh-CQs is a  
 501 two-layered language where expressions in a higher layer can only be built from the lower  
 502 layer. This lower layer, called **simple-MATLANG**, is a fragment of **fc-MATLANG** defined as:

$$503 \quad e ::= \mathbf{A} \mid \mathbf{1}^\alpha \mid \mathbf{I}^\alpha \mid e^T \mid e_1 \times e_2 \mid e_1 \odot e_2 \mid e \cdot \mathbf{1}^\alpha.$$

504 Note that in **simple-MATLANG** matrix multiplication is further restricted to matrix-vector  
 505 multiplication with the ones vector. Intuitively, all **simple-MATLANG** expressions can already  
 506 define q-hierarchical CQs like  $H(x) \leftarrow \exists y. A(x, y) \wedge U(x)$ , but it cannot define cross-products  
 507 like  $H(x, y) \leftarrow A(x) \wedge B(y)$ , which are q-hierarchical. For this reason, we need to enhance  
 508 **simple-MATLANG** with the higher layer. Specifically, we define **qh-MATLANG** as follows:

$$509 \quad e ::= e_1 \mid e_1 \odot (e_2 \cdot (\mathbf{1}^\alpha)^T) \mid (\mathbf{1}^\alpha \cdot e_1) \odot e_2 \mid (\mathbf{1}^\alpha \cdot e_1) \odot (e_2 \cdot (\mathbf{1}^\alpha)^T)$$

510 where  $e_1$  and  $e_2$  are **simple-MATLANG** expressions. Note that the subexpressions  $\mathbf{1}^\alpha \cdot e_1$  and  
 511  $e_2 \cdot (\mathbf{1}^\alpha)^T$  are valid if  $e_1$  and  $e_2$  have a row and column vector type, respectively. Then, both  
 512 subexpressions are useful for expanding vector-type expressions into a matrix-type expression.

513 The **qh-MATLANG** syntax does not allow expressions like, for example,  $\mathbf{1}^\alpha \cdot e_1$  where  $e_1$  is  
 514 a **simple-MATLANG** expression. Nevertheless, one can define this expression alternatively as  
 515  $(\mathbf{1}^\alpha \cdot e_1) \odot (\mathbf{1}^\alpha \cdot (\mathbf{1}^\alpha)^T)$ . For presentational purposes, we decided to define **qh-MATLANG** as  
 516 simple as possible, leaving out some expressions in **fc-MATLANG** that are not in **qh-MATLANG**,  
 517 although an equivalent **qh-MATLANG** expression defines it.

518 ► **Theorem 11.** *qh-MATLANG and binary q-hierarchical CQs are equally expressive.*

## 6 Efficient evaluation of free-connex and q-hierarchical queries

519

520 Now that we have precise connections between subfragments of **MATLANG** and subfragments  
 521 of CQ, we can use these connections to derive efficient evaluation algorithms for **MATLANG**.  
 522 Unfortunately, to apply the algorithms for CQ, we must first face two problems: (1) the

## 12:14 Enumeration and Updates for Conjunctive Linear Algebra Queries

523 evaluation algorithms for fc-CQs and qh-CQs are usually restricted to the Boolean semiring,  
 524 and (2) these algorithms are for CQ without inequalities (comparison atoms). To overcome  
 525 these problems, we need to revisit the fc-CQs and qh-CQs evaluation problem, generalize the  
 526 algorithmic results to other semirings (e.g.,  $\mathbb{R}$ ), and extend the results when queries have also  
 527 inequalities. Only then can we derive efficient algorithms for the fragments of fc-MATLANG  
 528 and qh-MATLANG.

529 **The evaluation setting.** In our setting, we consider query evaluation as an enumeration  
 530 problem. Specifically, let  $Q$  be a CQ over vocabulary  $\sigma$ ,  $H$  its relation symbol in the head,  
 531 and  $\mathcal{K}$  a semiring. We define the evaluation problem  $\text{Eval}(Q, \sigma, \mathcal{K})$  as follows:

532 **Problem:**  $\text{Eval}(Q, \sigma, \mathcal{K})$   
**Input:** A  $\mathcal{K}$ -database  $db$  over  $\sigma$   
**Output:** Enumerate  $\{(\bar{d}, \llbracket Q \rrbracket_{db}(H)(\bar{d})) \mid \llbracket Q \rrbracket_{db}(H)(\bar{d}) \neq \mathbb{0}\}$ .

533 In other words, the evaluation problems ask to retrieve the set of all tuples output by  $Q$   
 534 on  $db$ , together with their non-zero annotations. Similarly, we can define the evaluation  
 535 problem  $\text{Eval}(\mathbf{Q}, \mathcal{S}, \mathcal{K})$  for a conj-MATLANG query  $\mathbf{Q}$  over a matrix schema  $\mathcal{S}$  where we  
 536 aim to enumerate the non-zero entries of  $\llbracket \mathbf{Q} \rrbracket(\mathcal{I})$  given as input a matrix instance  $\mathcal{I}$  over  $\mathcal{S}$ ,  
 537 represented sparsely as the set of its non-zero entries.

538 As it is standard in the area, we consider enumeration algorithms on the Random Access  
 539 Machine (RAM) with uniform cost measure [3]. We assume that semiring values can be  
 540 represented in  $\mathcal{O}(1)$  space and that the semiring operations  $\oplus$  and  $\odot$  can be evaluated  
 541 in  $\mathcal{O}(1)$  time. We say that  $\text{Eval}(\mathbf{Q}, \mathcal{S}, \mathcal{K})$  can be evaluated with *linear-time preprocessing*  
 542 *and constant-delay*, if there exists an enumeration algorithm that takes  $\mathcal{O}(\|db\|)$  time to  
 543 preprocess the input database  $db$ , and then retrieves each output  $(\bar{d}, \llbracket Q \rrbracket_{db}(H)(\bar{d}))$  one by  
 544 one, without repetitions, and with constant-delay per output. Here, the size of database  
 545  $db$  is defined to be the number of non-zero entries. The same extends to  $\text{Eval}(\mathbf{Q}, \mathcal{S}, \mathcal{K})$  as  
 546 expected. Note that we measure the time and delay in *data complexity* as is standard in the  
 547 literature [5, 8, 9, 32].

548 **Evaluation of free-connex queries.** We are ready to state the algorithmic results for fc-CQ  
 549 and fc-MATLANG. A semiring  $(K, \oplus, \odot, \mathbb{0}, \mathbb{1})$  is *zero-divisor free* if, for all  $a, b \in K$ ,  $a \odot b = \mathbb{0}$   
 550 implies  $a = \mathbb{0}$  or  $b = \mathbb{0}$ . A zero-divisor free semiring is called a *semi-integral domain* [25].  
 551 Note that semirings used in practice, like  $\mathbb{B}$ ,  $\mathbb{N}$ , and  $\mathbb{R}$ , are semi-integral domains.

552 **► Theorem 12.** *Let  $\mathcal{K}$  be a semi-integral domain. For every free-connex query  $Q$  over  $\sigma$ ,*  
 553  *$\text{Eval}(Q, \sigma, \mathcal{K})$  can be evaluated with linear-time preprocessing and constant-delay. In particu-*  
 554 *lar,  $\text{Eval}(\mathbf{Q}, \mathcal{S}, \mathcal{K})$  can also be evaluated with linear-time preprocessing and constant-delay*  
 555 *for every fc-MATLANG  $\mathbf{Q}$  over  $\mathcal{S}$ .*

556 The semi-integral condition is necessary to ensure that zero outputs could only be  
 557 produced by some zero entries. For instance, consider a semiring  $(K, \oplus, \odot, \mathbb{0}, \mathbb{1})$  such that  
 558 there exist  $a, b \in K$  where  $a \neq \mathbb{0}$ ,  $b \neq \mathbb{0}$  and  $a \odot b = \mathbb{0}$ . Further, consider the query  
 559  $Q : H(x, y) \leftarrow R(x) \wedge S(y)$  over the previous semiring. Let  $R$  and  $S$  be relation symbols  
 560 with arity one and  $db$  a database over  $\sigma = \{R, S\}$  such that  $R(1) = a; R(2) = a; S(1) = b$   
 561 and  $S(2) = b$ . Then, the output of  $\text{Eval}(Q, \sigma, \mathcal{K})$  with input  $db$  is empty, although the body  
 562 can be instantiated in four different ways, all of them producing  $\mathbb{0}$  values.

563 We derive the enumeration algorithm for  $\text{Eval}(Q, \sigma, \mathcal{K})$  by extending the algorithms in [5]  
 564 for any semi-integral domain  $\mathcal{K}$  and taking care of the inequalities in  $Q$ . We derive the  
 565 enumeration algorithm for  $\text{Eval}(\mathbf{Q}, \mathcal{S}, \mathcal{K})$  by reducing to  $\text{Eval}(Q, \sigma, \mathcal{K})$ , using Corollary 10.

566 To illustrate the utility of Theorem 12, consider the fc-MATLANG query  $\mathbf{A} \odot (\mathbf{U} \cdot \mathbf{V}^T)$   
 567 where  $\mathbf{U}, \mathbf{V}$  are column vectors. When this query is evaluated in a bottom-up fashion, the  
 568 subexpression  $(\mathbf{U} \cdot \mathbf{V}^T)$  will generate partial results of size  $\|\mathbf{U}\| \|\mathbf{V}\|$ , causing the entire  
 569 evaluation to be of complexity  $\Omega(\|\mathbf{A}\| + \|\mathbf{U}\| \|\mathbf{V}\|)$ . By contrast, Theorem 12 tells us that  
 570 we may evaluate the query in time  $\mathcal{O}(\|\mathbf{A}\| + \|\mathbf{U}\| + \|\mathbf{V}\|)$ .

571 For lower bounds, we can also extend the results in [5] and [8] for the relational setting  
 572 under standard complexity assumptions. As in previous work, our lower bounds are for  
 573 CQ *without self-joins* and we need some additional restrictions for inequalities. Specifically,  
 574 we say that an inequality  $x \leq c$  in  $Q$  is *covered*, if there exists a relational atom in  $Q$  that  
 575 mention  $x$ . We say that  $Q$  is *constant-disjoint* if (i) for all covered inequalities  $x \leq c$  we have  
 576  $c \neq 1$ , and (ii) for all pairs  $(x \leq c, y \leq d)$  in  $Q$  of covered inequality  $x \leq c$  and non-covered  
 577 inequality  $y \leq d$  if  $c = d$  then  $y \notin \text{free}(Q)$ . In other words, if a constant symbol other than 1  
 578 occurs in both a covered and non-covered inequality in  $Q$ , then it occurs with a bound  
 579 variable in the non-covered inequality.

580 A semiring  $\mathcal{K} = (K, \oplus, \odot, \mathbf{0}, \mathbf{1})$  is *zero-sum free* [27, 28] if for all  $a, b \in K$  it holds that  
 581  $a \oplus b = \mathbf{0}$  implies  $a = b = \mathbf{0}$ . We are ready to extend the lower bound in [5, 8] as follows.

582 **► Theorem 13.** *Let  $Q$  be a CQ over  $\sigma$  without self-joins and constant-disjoint. Let  $\mathcal{K}$*   
 583 *be a semiring such that the subsemiring generated by  $\mathbf{0}_{\mathcal{K}}$  and  $\mathbf{1}_{\mathcal{K}}$  is zero-sum free. If*  
 584  *$\text{Eval}(Q, \sigma, \mathcal{K})$  can be evaluated with linear-time preprocessing and constant-delay, then  $Q$  is*  
 585 *free-connex, unless either the Sparse Boolean Matrix Multiplication, the Triangle Detection,*  
 586 *or the  $(k, k + 1)$ -Hyperclique conjecture is false.*

587 It is important to note that most semirings used in practice, like  $\mathbb{B}$ ,  $\mathbb{N}$ , and  $\mathbb{R}$ , are such that  
 588 the subsemiring generated by  $\mathbf{0}_{\mathcal{K}}$  and  $\mathbf{1}_{\mathcal{K}}$  is zero-sum free. Furthermore, the Sparse Boolean  
 589 Matrix Multiplication conjecture, the Triangle Detection conjecture, and the  $(k, k + 1)$ -  
 590 Hyperclique conjecture are standard complexity assumptions used by previous works [5, 8]  
 591 (see the formal statements in the full version [38]).

592 Unfortunately, given the asymmetry between the relational and matrix settings, the lower  
 593 bounds do not immediately transfer from the relational to the matrix setting. Specifically,  
 594 we need a syntactical restriction for conj-MATLANG that implies the constant-disjointedness  
 595 restriction in the translation of Theorem 8. Intuitively, this happens when both dimensions  
 596 of a conj-MATLANG query  $\mathbf{H} := e$  are fixed by matrix symbols in  $\mathcal{S}$ . For example, the  
 597 expressions  $\mathbf{A} \odot (\mathbf{U} \cdot \mathbf{V}^T)$  and  $\Sigma \mathbf{v}. \mathbf{A} \cdot \mathbf{v}$  have both dimensions (i.e., row and column) fixed  
 598 by  $\mathbf{A}$  where  $\mathbf{U}, \mathbf{V}$  are column vectors. Instead, the expression  $\mathbf{U} \cdot (\mathbf{1}^\alpha)^T$  does not, since its  
 599 column dimension depends on the value assigned for  $\alpha$  and is not necessarily fixed by  $\mathbf{U}$ .  
 600 Formally,  $\text{FixDim}(e)$  is the set of *fixed dimensions* of a conj-MATLANG expression  $e$  and it is  
 601 inductively defined as follows:

$$\begin{aligned}
 & \text{FixDim}(\mathbf{A}) := \{0, 1\} \\
 & \text{FixDim}(\mathbf{v}) := \{1\} \\
 & \text{FixDim}(e^T) := \{(i + 1) \bmod 2 \mid i \in \text{FixDim}(e)\} \\
 602 & \text{FixDim}(e_1 \cdot e_2) := (\text{FixDim}(e_1) \setminus \{1\}) \cup (\text{FixDim}(e_2) \setminus \{0\}) \\
 & \text{FixDim}(e_1 \times e_2) := \text{FixDim}(e_2) \\
 & \text{FixDim}(e_1 \odot e_2) := \text{FixDim}(e_1) \cup \text{FixDim}(e_2) \\
 & \text{FixDim}(\Sigma \mathbf{v}. e) := \text{FixDim}(e).
 \end{aligned}$$

603 An expression  $e$  has *guarded dimensions* if  $\text{FixDim}(e) = \{0, 1\}$ .

604 The former is straightforwardly extended to a conj-MATLANG query  $\mathbf{Q}$ , where  $\text{FixDim}(\mathbf{Q})$   
 605 stands for  $\text{FixDim}(e)$ . The following lower bound is now attainable.

606 ► **Corollary 14.** *Let  $\mathbf{Q}$  be a conj-MATLANG query over  $\mathcal{S}$  such that  $\mathbf{Q}$  does not repeat matrix*  
 607 *symbols and  $\mathbf{Q}$  has guarded dimensions. Let  $\mathcal{K}$  be a semiring such that the subsemiring*  
 608 *generated by  $0_{\mathcal{K}}$  and  $1_{\mathcal{K}}$  is zero-sum free. If  $\text{Eval}(\mathbf{Q}, \mathcal{S}, \mathcal{K})$  can be evaluated with linear-time*  
 609 *preprocessing and constant-delay, then  $\mathbf{Q}$  is equivalent to a fc-MATLANG query, unless either*  
 610 *the Sparse Boolean Matrix Multiplication, the Triangle Detection, or the  $(k, k+1)$ -Hyperclique*  
 611 *conjecture is false.*

612 **The dynamic evaluation setting.** We move now to the dynamic query evaluation both in  
 613 the relational and matrix scenarios. Specifically, we consider the following set of updates.  
 614 Recall that  $\mathcal{K} = (K, \oplus, \odot, \mathbf{0}, \mathbf{1})$  is a semiring and  $\sigma$  a vocabulary.

- 615 ■ A single-tuple *insertion* (over  $\mathcal{K}$  and  $\sigma$ ) is an operation  $u = \text{insert}(R, \bar{d}, k)$  with  $R \in \sigma$ ,  $\bar{d}$   
 616 a tuple of arity  $\text{ar}(R)$ , and  $k \in K$ . When applied to a database  $db$  it induces the database  
 617  $db + u$  that is identical to  $db$ , but  $R^{db+u}(\bar{d}) = R^{db}(\bar{d}) \oplus k$ .
- 618 ■ A single-tuple *deletion* (over  $\mathcal{K}$  and  $\sigma$ ) is an expression  $u = \text{delete}(R, \bar{d})$  with  $R \in \sigma$  and  
 619  $\bar{d}$  a tuple of arity  $\text{ar}(R)$ . When applied to a database  $db$  it induces the database  $db + u$   
 620 that is identical to  $db$ , but  $R^{db+u}(\bar{d}) = \mathbf{0}$ .

621 Notice that if every element in  $\mathcal{K}$  has an additive inverse (i.e.,  $\mathcal{K}$  is a ring), one can simulate  
 622 a deletion with an insertion. However, if this is not the case (e.g.,  $\mathbb{B}$  or  $\mathbb{N}$ ), then a single-tuple  
 623 deletion is a necessary operation.

624 We define the dynamic query evaluation problem  $\text{DynEval}(Q, \sigma, \mathcal{K})$  as the extension of  
 625  $\text{Eval}(Q, \sigma, \mathcal{K})$  under the set of updates  $U$  of all single-tuple insertions and deletions over  
 626  $\mathcal{K}$  and  $\sigma$ . We say that  $\text{DynEval}(Q, \sigma, \mathcal{K})$  can be evaluated dynamically with *constant-time*  
 627 *update and constant-delay*, if  $\text{Eval}(Q, \sigma, \mathcal{K})$  can be evaluated with linear-time preprocessing  
 628 and constant delay, and, moreover, for every update  $u \in U$ , it takes constant-time to update  
 629 the state of the algorithm from  $db$  to  $db + u$  so that, immediately after, we can retrieve  
 630 each output  $(\bar{d}, \llbracket Q \rrbracket_{db+u}(H)(\bar{d}))$  one by one, without repetitions, and with constant-delay  
 631 per output. Similarly, we define the dynamic query evaluation problem  $\text{DynEval}(\mathbf{Q}, \mathcal{S}, \mathcal{K})$  of  
 632  $\text{Eval}(\mathbf{Q}, \mathcal{S}, \mathcal{K})$  for the matrix setting, requiring the same dynamic guarantees.

633 Note that updates allow to modify the contents of relations, but not of constant symbols.  
 634 Similarly, in the linear algebra setting updates only affect matrix entry values, not matrix  
 635 dimensions. An interesting line of future work is to consider dimension updates, which  
 636 correspond to allow updates of constant symbols in CQs.

637 **Evaluation of q-hierarchical queries.** Similar than for free-connex queries, we can provide  
 638 dynamic evaluation algorithms for qh-CQ and qh-MATLANG queries. However, for this  
 639 dynamic setting, we require some additional algorithmic assumptions over the semiring. Let  
 640  $\mathcal{K} = (K, \oplus, \odot, \mathbf{0}, \mathbf{1})$  be a semiring and  $M$  be the set of all multisets of  $K$ . For any  $k \in K$   
 641 and  $m \in M$ , define  $\text{ins}(k, m)$  and  $\text{del}(k, m)$  to be the multisets resulting from inserting or  
 642 deleting  $k$  from  $m$ , respectively. Then we say that  $\mathcal{K}$  is *sum-maintainable* if there exists a data  
 643 structure  $\mathcal{D}$  to represent multisets of  $K$  such that the empty set  $\emptyset$  can be built in constant  
 644 time, and if  $\mathcal{D}$  represents  $m \in M$  then: **(1)** the value  $\bigoplus_{k \in m} k$  can always be computed from  
 645  $\mathcal{D}$  in constant time; **(2)** a data structure that represents  $\text{ins}(k, m)$  can be obtained from  $\mathcal{D}$   
 646 in constant time; and **(3)** a data structure that represents  $\text{del}(k, m)$  can be obtained from  
 647  $\mathcal{D}$  in constant time. One can easily notice that if each element of  $\mathcal{K}$  has an additive inverse  
 648 (i.e.,  $\mathcal{K}$  is a ring), then  $\mathcal{K}$  is sum-maintainable, like  $\mathbb{R}$ . Other examples of sum-maintainable  
 649 semirings (without additive inverses) are  $\mathbb{B}$  and  $\mathbb{N}$ .

650 ► **Theorem 15.** *Let  $\mathcal{K}$  be a sum-maintainable semi-integral domain. For every q-hierarchical*  
 651 *CQ  $Q$ ,  $\text{DynEval}(Q, \sigma, \mathcal{K})$  can be evaluated dynamically with constant-time update and constant-*

652 *delay. In particular,  $\text{DynEval}(\mathbf{Q}, \mathcal{S}, \mathcal{K})$  can also be evaluated dynamically with constant-time*  
 653 *update and constant-delay for every qh-MATLANG  $\mathbf{Q}$  over  $\mathcal{S}$ .*

654 Similar than for free-connex CQ, we can extend the lower bound in [9] when the subsemiring  
 655 generated by  $\mathbf{0}_{\mathcal{K}}$  and  $\mathbf{1}_{\mathcal{K}}$  is zero-sum free, by assuming the Online Boolean Matrix-Vector  
 656 Multiplication (OMv) conjecture [29].

657 ► **Theorem 16.** *Let  $Q$  be a CQ over  $\sigma$  without self-joins and constant-disjoint. Let  $\mathcal{K}$*   
 658 *be a semiring such that the subsemiring generated by  $\mathbf{0}_{\mathcal{K}}$  and  $\mathbf{1}_{\mathcal{K}}$  is zero-sum free. If*  
 659  *$\text{DynEval}(Q, \sigma, \mathcal{K})$  can be evaluated dynamically with constant-time update and constant-delay,*  
 660 *then  $Q$  is  $q$ -hierarchical, unless the OMv conjecture is false.*

661 This transfers to conj-MATLANG, similarly to the lower bound in the free-connex case.

662 ► **Corollary 17.** *Let  $\mathbf{Q}$  be a conj-MATLANG query over  $\mathcal{S}$  such that  $\mathbf{Q}$  does not repeat matrix*  
 663 *symbols and  $\mathbf{Q}$  has guarded dimensions. Let  $\mathcal{K}$  be a semiring such that the subsemiring*  
 664 *generated by  $\mathbf{0}_{\mathcal{K}}$  and  $\mathbf{1}_{\mathcal{K}}$  is zero-sum free. If  $\text{DynEval}(\mathbf{Q}, \mathcal{S}, \mathcal{K})$  can be evaluated dynamically*  
 665 *with constant-time update and constant-delay, then  $\mathbf{Q}$  is equivalent to a qh-MATLANG query,*  
 666 *unless the OMv conjecture is false.*

## 667 7 Conclusions and future work

668 In this work, we isolated the subfragments of conj-MATLANG that admit efficient evaluation in  
 669 both static and dynamic scenarios. We found these algorithms by making the correspondence  
 670 between CQ and MATLANG, extending the evaluation algorithms for free-connex and  $q$ -  
 671 hierarchical CQ, and then translating these algorithms to the corresponding subfragments,  
 672 namely, fc-MATLANG and qh-MATLANG. To the best of our knowledge, this is the first  
 673 work that characterizes subfragments of linear algebra query languages that admit efficient  
 674 evaluation. Moreover, this correspondence improves our understanding of its expressibility.

675 Regarding future work, a relevant direction is to extend fc-MATLANG and qh-MATLANG  
 676 with disjunction, namely, matrix summation. This direction is still an open problem even for  
 677 CQ with union [14]. Another natural extension is to add point-wise functions and understand  
 678 how they affect expressibility and efficient evaluation. Finally, improving the lower bounds to  
 679 queries without self-join would be interesting, which is also an open problem for CQ [8, 15].

## 680 — References —

- 681 1 Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu  
 682 Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg,  
 683 Rajat Monga, Sherry Moore, Derek Gordon Murray, Benoit Steiner, Paul A. Tucker, Vijay  
 684 Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: A  
 685 system for large-scale machine learning. In Kimberly Keeton and Timothy Roscoe, editors,  
 686 *12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016,*  
 687 *Savannah, GA, USA, November 2-4, 2016*, pages 265–283. USENIX Association, 2016. URL:  
 688 <https://www.usenix.org/conference/osdi16/technical-sessions/presentation/abadi>.
- 689 2 Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases*. Addison-Wesley,  
 690 1995. URL: <http://webdam.inria.fr/Alice/>.
- 691 3 Alfred V. Aho, John E. Hopcroft, and Jeffrey D. Ullman. *The Design and Analysis of Computer*  
 692 *Algorithms*. Addison-Wesley, 1974.
- 693 4 Michael J. Anderson, Shaden Smith, Narayanan Sundaram, Mihai Capota, Zheguang Zhao,  
 694 Subramanya Dullloor, Nadathur Satish, and Theodore L. Willke. Bridging the gap between  
 695 HPC and big data frameworks. *Proc. VLDB Endow.*, 10(8):901–912, 2017. URL: <http://www.vldb.org/pvldb/vol10/p901-anderson.pdf>, doi:10.14778/3090163.3090168.

- 697 5 Guillaume Bagan, Arnaud Durand, and Etienne Grandjean. On acyclic conjunctive queries and  
698 constant delay enumeration. In Jacques Duparc and Thomas A. Henzinger, editors, *Computer*  
699 *Science Logic, 21st International Workshop, CSL 2007, 16th Annual Conference of the EACSL,*  
700 *Lausanne, Switzerland, September 11-15, 2007, Proceedings*, volume 4646 of *Lecture Notes in*  
701 *Computer Science*, pages 208–222. Springer, 2007. doi:10.1007/978-3-540-74915-8\_18.
- 702 6 Pablo Barceló, Nelson Higuera, Jorge Pérez, and Bernardo Subercaseaux. On the ex-  
703 pressiveness of LARA: A unified language for linear and relational algebra. In Carsten  
704 Lutz and Jean Christoph Jung, editors, *23rd International Conference on Database The-*  
705 *ory, ICDT 2020, March 30-April 2, 2020, Copenhagen, Denmark*, volume 155 of *LIPICs*,  
706 pages 6:1–6:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. URL: <https://doi.org/10.4230/LIPICs.ICDT.2020.6>, doi:10.4230/LIPICs.ICDT.2020.6.
- 707 7 Catriel Beeri, Ronald Fagin, David Maier, and Mihalis Yannakakis. On the desirability of  
708 acyclic database schemes. *J. ACM*, 30(3):479–513, 1983. doi:10.1145/2402.322389.
- 709 8 Christoph Berkholz, Fabian Gerhardt, and Nicole Schweikardt. Constant delay enumeration  
710 for conjunctive queries: a tutorial. *ACM SIGLOG News*, 7(1):4–33, 2020. doi:10.1145/  
711 3385634.3385636.
- 712 9 Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering conjunctive queries  
713 under updates. In Emanuel Sallinger, Jan Van den Bussche, and Floris Geerts, editors,  
714 *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database*  
715 *Systems, PODS 2017, Chicago, IL, USA, May 14-19, 2017*, pages 303–318. ACM, 2017.  
716 doi:10.1145/3034786.3034789.
- 717 10 Philip A. Bernstein and Nathan Goodman. Power of natural semijoins. *SIAM J. Comput.*,  
718 10(4):751–771, 1981. doi:10.1137/0210059.
- 719 11 Johann Brault-Baron. *De la pertinence de l'énumération : complexité en logiques proposi-*  
720 *tionnelle et du premier ordre. (The relevance of the list: propositional logic and complex-*  
721 *ity of the first order)*. PhD thesis, University of Caen Normandy, France, 2013. URL:  
722 <https://tel.archives-ouvertes.fr/tel-01081392>.
- 723 12 Robert Brijder, Floris Geerts, Jan Van den Bussche, and Timmy Weerwag. On the expressive  
724 power of query languages for matrices. *ACM Trans. Database Syst.*, 44(4):15:1–15:31, 2019.  
725 doi:10.1145/3331445.
- 726 13 Robert Brijder, Marc Gyssens, and Jan Van den Bussche. On matrices and k-relations. In  
727 Andreas Herzig and Juha Kontinen, editors, *Foundations of Information and Knowledge*  
728 *Systems - 11th International Symposium, FoIKS 2020, Dortmund, Germany, February 17-21,*  
729 *2020, Proceedings*, volume 12012 of *Lecture Notes in Computer Science*, pages 42–57. Springer,  
730 2020. doi:10.1007/978-3-030-39951-1\_3.
- 731 14 Nofar Carmeli and Markus Kröll. On the enumeration complexity of unions of conjunctive  
732 queries. *ACM Trans. Database Syst.*, 46(2):5:1–5:41, 2021. doi:10.1145/3450263.
- 733 15 Nofar Carmeli and Luc Segoufin. Conjunctive queries with self-joins, towards a fine-grained  
734 enumeration complexity analysis. In Floris Geerts, Hung Q. Ngo, and Stavros Sintos, editors,  
735 *Proceedings of the 42nd ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database*  
736 *Systems, PODS 2023, Seattle, WA, USA, June 18-23, 2023*, pages 277–289. ACM, 2023.  
737 doi:10.1145/3584372.3588667.
- 738 16 Arnaud Durand and Etienne Grandjean. First-order queries on structures of bounded degree  
739 are computable with constant delay. *ACM Trans. Comput. Log.*, 8(4):21, 2007. doi:10.1145/  
740 1276920.1276923.
- 741 17 Arnaud Durand, Nicole Schweikardt, and Luc Segoufin. Enumerating answers to first-order  
742 queries over databases of low degree. In Richard Hull and Martin Grohe, editors, *Proceedings*  
743 *of the 33rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems,*  
744 *PODS 2014, Snowbird, UT, USA, June 22-27, 2014*, pages 121–131. ACM, 2014. doi:  
745 10.1145/2594538.2594539.
- 746

- 747 18 Idan Eldar, Nofar Carmeli, and Benny Kimelfeld. Direct access for answers to conjunctive  
748 queries with aggregation. *CoRR*, abs/2303.05327, 2023. URL: [https://doi.org/10.48550/  
749 arXiv.2303.05327](https://doi.org/10.48550/arXiv.2303.05327), [arXiv:2303.05327](https://arxiv.org/abs/2303.05327), doi:10.48550/ARXIV.2303.05327.
- 750 19 Tarek Elgamal, Shangyu Luo, Matthias Boehm, Alexandre V. Evfimievski, Shirish Tatikonda,  
751 Berthold Reinwald, and Prithviraj Sen. SPOOF: sum-product optimization and operator  
752 fusion for large-scale machine learning. In *8th Biennial Conference on Innovative Data Sys-  
753 tems Research, CIDR 2017, Chaminade, CA, USA, January 8-11, 2017, Online Proceedings*.  
754 [www.cidrdb.org](http://cidrdb.org), 2017. URL: [http://cidrdb.org/cidr2017/papers/p3-elgamal-cidr17.  
755 pdf](http://cidrdb.org/cidr2017/papers/p3-elgamal-cidr17.pdf).
- 756 20 Ahmed Elgohary, Matthias Boehm, Peter J. Haas, Frederick R. Reiss, and Berthold Reinwald.  
757 Scaling machine learning via compressed linear algebra. *SIGMOD Rec.*, 46(1):42–49, 2017.  
758 doi:10.1145/3093754.3093765.
- 759 21 Ronald Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. *J. ACM*,  
760 30(3):514–550, 1983. doi:10.1145/2402.322390.
- 761 22 Floris Geerts. On the expressive power of linear algebra on graphs. In Pablo Barceló and Marco  
762 Calautti, editors, *22nd International Conference on Database Theory, ICDT 2019, March  
763 26-28, 2019, Lisbon, Portugal*, volume 127 of *LIPICs*, pages 7:1–7:19. Schloss Dagstuhl - Leibniz-  
764 Zentrum für Informatik, 2019. URL: <https://doi.org/10.4230/LIPICs.ICDT.2019.7>, doi:  
765 10.4230/LIPICs.ICDT.2019.7.
- 766 23 Floris Geerts, Thomas Muñoz, Cristian Riveros, and Domagoj Vrgoc. Expressive power of  
767 linear algebra query languages. In Leonid Libkin, Reinhard Pichler, and Paolo Guagliardo,  
768 editors, *Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles  
769 of Database Systems, PODS 2021, Virtual Event, China, June 20-25, 2021*, pages 342–354.  
770 ACM, 2021. doi:10.1145/3452021.3458314.
- 771 24 Floris Geerts and Juan L. Reutter. Expressiveness and approximation properties of graph  
772 neural networks. In *The Tenth International Conference on Learning Representations, ICLR  
773 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL: [https://openreview.  
774 net/forum?id=wIzUeM3TAU](https://openreview.net/forum?id=wIzUeM3TAU).
- 775 25 Jonathan S Golan. *Semirings and their Applications*. Springer Science & Business Media,  
776 2013.
- 777 26 Todd J. Green, Gregory Karvounarakis, and Val Tannen. Provenance semirings. In Leonid  
778 Libkin, editor, *Proceedings of the 26th ACM SIGACT-SIGMOD-SIGART Symposium on  
779 Principles of Database Systems, PODS 2007, Beijing, China, June 11-13, 2007*, pages 31–40.  
780 ACM, 2007. doi:10.1145/1265530.1265535.
- 781 27 Udo Hebisch and Hanns Joachim Weinert. *Semirings: algebraic theory and applications in  
782 computer science*, volume 5. World Scientific, 1998.
- 783 28 Udo Hebisch and Hans Joachim Weinert. Semirings and semifields. *Handbook of Algebra*,  
784 1:425–462, 1996.
- 785 29 Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak.  
786 Unifying and strengthening hardness for dynamic problems via the online matrix-vector  
787 multiplication conjecture. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *Proceedings of  
788 the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland,  
789 OR, USA, June 14-17, 2015*, pages 21–30. ACM, 2015. doi:10.1145/2746539.2746609.
- 790 30 Botong Huang, Shivnath Babu, and Jun Yang. Cumulon: optimizing statistical data analysis in  
791 the cloud. In Kenneth A. Ross, Divesh Srivastava, and Dimitris Papadias, editors, *Proceedings  
792 of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2013, New  
793 York, NY, USA, June 22-27, 2013*, pages 1–12. ACM, 2013. doi:10.1145/2463676.2465273.
- 794 31 Muhammad Idris, Martín Ugarte, and Stijn Vansummeren. The dynamic yannakakis algorithm:  
795 Compact and efficient query processing under updates. In Semih Salihoglu, Wenchao Zhou,  
796 Rada Chirkova, Jun Yang, and Dan Suciu, editors, *Proceedings of the 2017 ACM International  
797 Conference on Management of Data, SIGMOD Conference 2017, Chicago, IL, USA, May  
798 14-19, 2017*, pages 1259–1274. ACM, 2017. doi:10.1145/3035918.3064027.

- 799 32 Muhammad Idris, Martín Ugarte, Stijn Vansummeren, Hannes Voigt, and Wolfgang Lehner.  
800 General dynamic yannakakis: conjunctive queries with theta joins under updates. *VLDB*  
801 *J.*, 29(2-3):619–653, 2020. URL: <https://doi.org/10.1007/s00778-019-00590-9>, doi:10.  
802 1007/S00778-019-00590-9.
- 803 33 Dimitrije Jankov, Shangyu Luo, Binhang Yuan, Zhuhua Cai, Jia Zou, Chris Jermaine, and  
804 Zekai J. Gao. Declarative recursive computation on an RDBMS: or, why you should use a  
805 database for distributed machine learning. *SIGMOD Rec.*, 49(1):43–50, 2020. doi:10.1145/  
806 3422648.3422659.
- 807 34 Konstantinos Kanellopoulos, Nandita Vijaykumar, Christina Giannoula, Roknoddin Azizi,  
808 Skanda Koppula, Nika Mansouri-Ghiasi, Taha Shahroodi, Juan Gómez-Luna, and Onur  
809 Mutlu. SMASH: co-designing software compression and hardware-accelerated indexing for  
810 efficient sparse matrix operations. In *Proceedings of the 52nd Annual IEEE/ACM International*  
811 *Symposium on Microarchitecture, MICRO 2019, Columbus, OH, USA, October 12-16, 2019*,  
812 pages 600–614. ACM, 2019. doi:10.1145/3352460.3358286.
- 813 35 Ahmet Kara, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. Trade-offs in static and  
814 dynamic evaluation of hierarchical queries. In Dan Suciu, Yufei Tao, and Zhewei Wei, editors,  
815 *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database*  
816 *Systems, PODS 2020, Portland, OR, USA, June 14-19, 2020*, pages 375–392. ACM, 2020.  
817 doi:10.1145/3375395.3387646.
- 818 36 Wojciech Kazana and Luc Segoufin. Enumeration of first-order queries on classes of structures  
819 with bounded expansion. In Richard Hull and Wenfei Fan, editors, *Proceedings of the 32nd*  
820 *ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS*  
821 *2013, New York, NY, USA, June 22 - 27, 2013*, pages 297–308. ACM, 2013. doi:10.1145/  
822 2463664.2463667.
- 823 37 Shangyu Luo, Zekai J. Gao, Michael N. Gubanov, Luis Leopoldo Perez, and Christopher M.  
824 Jermaine. Scalable linear algebra on a relational database system. *SIGMOD Rec.*, 47(1):24–31,  
825 2018. doi:10.1145/3277006.3277013.
- 826 38 Thomas Muñoz, Cristian Riveros, and Stijn Vansummeren. Enumeration and updates for  
827 conjunctive linear algebra queries through expressibility. *CoRR*, abs/2310.04118, 2023.  
828 URL: <https://doi.org/10.48550/arXiv.2310.04118>, arXiv:2310.04118, doi:10.48550/  
829 ARXIV.2310.04118.
- 830 39 Nicole Schweikardt, Luc Segoufin, and Alexandre Vigny. Enumeration for FO queries over  
831 nowhere dense graphs. In Jan Van den Bussche and Marcelo Arenas, editors, *Proceedings*  
832 *of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems,*  
833 *PODS 2018, Houston, TX, USA, June 10-15, 2018*, pages 151–163. ACM, 2018. doi:10.1145/  
834 3196959.3196971.
- 835 40 Amir Shaikhha, Mohammed Elseidy, Stephan Mihaila, Daniel Espino, and Christoph Koch.  
836 Synthesis of incremental linear algebra programs. *ACM Trans. Database Syst.*, 45(3):12:1–12:44,  
837 2020. doi:10.1145/3385398.
- 838 41 Yisu Remy Wang, Shana Hutchison, Dan Suciu, Bill Howe, and Jonathan Leang. SPORES:  
839 sum-product optimization via relational equality saturation for large scale linear algebra.  
840 *Proc. VLDB Endow.*, 13(11):1919–1932, 2020. URL: [http://www.vldb.org/pvldb/vol13/  
841 p1919-wang.pdf](http://www.vldb.org/pvldb/vol13/p1919-wang.pdf).
- 842 42 Fan Yang, Yuzhen Huang, Yunjian Zhao, Jinfeng Li, Guanxian Jiang, and James Cheng.  
843 The best of both worlds: Big data programming with both productivity and performance.  
844 In Semih Salihoglu, Wenchao Zhou, Rada Chirkova, Jun Yang, and Dan Suciu, editors,  
845 *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD*  
846 *Conference 2017, Chicago, IL, USA, May 14-19, 2017*, pages 1619–1622. ACM, 2017. doi:  
847 10.1145/3035918.3058735.