

# Information Theory Strikes Back: New Development in the Theory of Cardinality Estimation

Mahmoud Abo Khamis  
RelationalAI

Dan Olteanu  
University of Zurich

Vasileios Nakos  
National University of Athens

Dan Suciu  
University of Washington

## ABSTRACT

Estimating the cardinality of the output of a query is a fundamental problem in database query processing.

In this article, we overview a recently published contribution that casts the cardinality estimation problem as linear optimization and computes guaranteed upper bounds on the cardinality of the output for any full conjunctive query. The objective of the linear program is to maximize the joint entropy of the query variables and its constraints are the Shannon information inequalities and new information inequalities involving  $\ell_p$ -norms of the degree sequences of the join attributes.

The bounds based on arbitrary norms can be asymptotically lower than those based on the  $\ell_1$  and  $\ell_\infty$  norms, which capture the cardinalities and respectively the max-degrees of the input relations. They come with a matching query evaluation algorithm, are computable in exponential time in the query size, and are provably tight when each degree sequence is on one join attribute.

## 1. INTRODUCTION

After five decades of research, the cardinality estimation problem remains one of the unsolved central problems in database systems. It is a crucial component of a query optimizer, as it allows to select a query plan that minimizes the size of the intermediate results and therefore the necessary time and memory to compute the query. Yet traditional estimators present in virtually all database management systems routinely underestimate or overestimate the true cardinality

This work was partially supported by NSF-BSF 2109922, NSF IIS 2314527, NSF SHF 2312195, Swiss NSF 200021-231956, and RelationalAI. The authors would like to acknowledge Christoph Mayer, Luis Torrejón Machado, and Haozhe Zhang for their help with the preliminary experiments reported in Section 6 of this paper.

© 2025 Copyright held by the authors. Publication rights licensed to ACM. This is a minor revision of the paper entitled “Join Size Bounds using  $\ell_p$ -Norms on Degree Sequences”, published in Proc. ACM Manag. Data 2, 2 (PODS), Article 96 (May 2024). DOI: <http://dx.doi.org/10.1145/3651597>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

Copyright 2025 ACM 0001-0782/24/0X00 ...\$5.00.

by orders of magnitude, which can lead to inefficient query plans [1, 2, 3, 4].

The past two decades introduced information-theoretic *worst-case upper bounds* on the output size of a full conjunctive query. The first such bound is the *AGM bound*, which is a function of the sizes of the input tables [5]. It was further refined in the presence of functional dependencies [6, 7]. A more general bound is the *PANDA bound*, which is a function of both the sizes of the input tables and the max degrees of join attributes in these tables [8]. Yet these information-theoretic upper bounds have not had practical impact. One reason is that most queries in practice are acyclic, where these upper bounds become trivial: the AGM bound is simply the multiplication of the sizes of some of the tables, while the PANDA bound is the multiplication of the size of one table with the *maximum* degrees of the joining tables. The latter is not new for a practitioner: standard estimators do the same, but use the *average* degrees instead of the max degrees. A second, related reason, is that these bounds use essentially the same type of statistics as existing cardinality estimators: cardinalities and max or average degrees. They have been implemented under the name *pessimistic cardinality estimators* [9, 10], but their empirical evaluation showed that they remain less accurate than other estimators [11, 3].

In this paper, we discuss upper bounds on the query output size that use more refined data statistics: *the  $\ell_p$ -norms of degree sequences of the join attributes in the relations*. They were originally introduced in our prior work [12].

The *degree* of an attribute value is the number of tuples in the table with that attribute value. The *degree sequence* of a join attribute is the sorted list of the degrees of the distinct attribute values,  $\mathbf{d} = (d_1 \geq \dots \geq d_m)$ , where  $d_1$  the largest degree and  $d_m$  the smallest degree. The  $\ell_p$ -norm of a degree sequence  $\mathbf{d}$  is defined as  $\|\mathbf{d}\|_p = (d_1^p + \dots + d_m^p)^{1/p}$ . In particular, the  $\ell_1$ -norm is  $\|\mathbf{d}\|_1 = \sum_{i=1}^m d_i$ , i.e., the cardinality of the table, while the  $\ell_\infty$ -norm is  $\|\mathbf{d}\|_\infty = d_1$ , i.e., the max degree. The upper bounds based on  $\ell_p$ -norms, or  $\ell_p$ -bounds for short, strictly generalize previous bounds based on cardinalities and max-degrees [8], since they can use any set of  $\ell_p$ -norms, not only the  $\ell_1$  and  $\ell_\infty$  norms.

Like the AGM [5] and the PANDA [8] bounds, these new  $\ell_p$ -bounds rely on information inequalities. The computed bound is the optimal solution of a linear program, whose objective is to maximize the joint entropy of the query variables, and whose constraints are the Shannon information inequalities and new information inequalities involving  $\ell_p$ -norms of the degree sequences of the join attributes. The linear program in the original paper [12] has a number of

unknowns that is exponential in the number of query variables (the unknowns are the joint entropies of any subset of the set of query variables). Follow-up work [13] introduces two improvements: smaller equivalent linear programs and support for a richer query language. We can construct equivalent linear programs whose number<sup>1</sup> of unknowns is quadratic in the number of query variables and linear in the number of statistics. We can also generalize the  $\ell_p$ -bounds from full conjunctive queries to queries with group-by clauses and equality and range predicates.

For more background on cardinality estimation, we point the reader to a recent overview on traditional and pessimistic estimators [15], a recent survey on learned cardinality estimators [16], and a recent monograph on query optimizers, with an discussion on the cardinality estimation in Microsoft SQL Server [17].

## 1.1 The $\ell_p$ -Bounds in Action

We illustrate the  $\ell_p$ -bounds for two simple queries: the join of two relations, as the simplest acyclic query, and the triangle join, as the simplest cyclic query. Further examples are discussed in the original paper [12].

EXAMPLE 1.1. Let us first consider the following simple join of two relations:

$$Q_1(X, Y, Z) = R(X, Y) \wedge S(Y, Z) \quad (1)$$

Traditional cardinality estimators (as found in textbooks [18], see also [1]) use the formula

$$|Q_1| \approx \frac{|R| \cdot |S|}{\max(|\Pi_Y(R)|, |\Pi_Y(S)|)} \quad (2)$$

Since  $\frac{|R|}{|\Pi_Y(R)|}$  is the average degree of  $R(X|Y)$ , (2) is equivalent to

$$|Q_1| \approx \min(|S| \cdot \text{avg}(\text{deg}_R(X|Y)), |R| \cdot \text{avg}(\text{deg}_S(Z|Y))) \quad (3)$$

Here, we use  $\text{deg}_R(X|Y) = (d_1, \dots, d_m)$  to denote the degree sequence of  $Y$  in  $R$ :  $d_i$  is the number of occurrences of the  $i$ 'th most frequent value  $Y = y$ . Similarly,  $\text{deg}_S(Z|Y)$  is the degree sequence of  $Y$  in  $S$ .

We now turn our attention to upper bounds. The AGM bound for  $Q_1$  is  $|R| \cdot |S|$ . A better bound is the PANDA bound, which replaces avg with max in (3):

$$|Q_1| \leq \min(|S| \cdot \|\text{deg}_R(X|Y)\|_\infty, |R| \cdot \|\text{deg}_S(Z|Y)\|_\infty) \quad (4)$$

Our framework derives several new upper bounds, by using  $\ell_p$ -norms other than  $\ell_1$  and  $\ell_\infty$ . We start with the simplest:

$$|Q_1| \leq \|\text{deg}_R(X|Y)\|_2 \cdot \|\text{deg}_S(Z|Y)\|_2 \quad (5)$$

This is the Cauchy-Schwartz inequality. Depending on the data, (5) can be asymptotically better than (4). A simple example where this happens is when  $Q_1$  is a self-join, i.e.,  $Q_1(X, Y, Z) = R(X, Y) \wedge R(Z, Y)$ . Then, the two degree sequences are equal,  $\text{deg}_R(X|Y) = \text{deg}_R(Z|Y)$ , and (5) becomes an equality, because  $|Q_1| = \|\text{deg}_R(X|Y)\|_2^2$ . Thus, (5) becomes equality, while (4) continues to be an over approximation of  $|Q_1|$ , and can be asymptotically worse.

<sup>1</sup> This improvement draws on a recent work [14] that reduces linear programs using  $\ell_1$  and  $\ell_\infty$  norms to an equivalent collection of network flow problems. We extended this idea to the case of arbitrary  $\ell_p$ -norms.

A more sophisticated inequality for  $Q_1$  is the following, which holds for all  $p, q \geq 0$  s.t.  $\frac{1}{p} + \frac{1}{q} \leq 1$ :

$$|Q_1| \leq \|\text{deg}_R(X|Y)\|_p \cdot \|\text{deg}_S(Z|Y)\|_q^{\frac{q}{p(q-1)}} |S|^{1 - \frac{q}{p(q-1)}} \quad (6)$$

Depending on the concrete statistics on the data, this new bound can be much better than both (4) and (5). The new bounds (5)-(6) are just three examples and other inequalities can be derived using  $\ell_p$ -bounds.

EXAMPLE 1.2. Let us now consider the triangle query:

$$Q_2(X, Y, Z) = R(X, Y) \wedge S(Y, Z) \wedge T(Z, X). \quad (7)$$

The AGM bound [5] uses the  $\ell_1$  norm:

$$|Q_2| \leq (\|\text{deg}_R(Y|X)\|_1 \cdot \|\text{deg}_S(Z|Y)\|_1 \cdot \|\text{deg}_T(X|Z)\|_1)^{1/2} \\ = (|R| \cdot |S| \cdot |T|)^{1/2}. \quad (8)$$

The PANDA bound [8] uses the  $\ell_1$  and  $\ell_\infty$  norms, two examples of this bound are as follows:

$$|Q_2| \leq \|\text{deg}_R(Y|X)\|_1 \cdot \|\text{deg}_S(Z|Y)\|_\infty \quad (9) \\ |Q_2| \leq \|\text{deg}_T(X|Z)\|_\infty \cdot \|\text{deg}_S(Z|Y)\|_1$$

Here,  $\text{deg}_R(Y|X)$ ,  $\text{deg}_S(Z|Y)$ , and  $\text{deg}_T(X|Z)$  are the degree sequences of  $X$  in  $R$ ,  $Y$  in  $S$ , and  $Z$  in  $T$  respectively.

If the  $\ell_2$  and  $\ell_3$  norms of the degree sequences are also available, then we can derive new upper bounds, for example:

$$|Q_2| \leq (\|\text{deg}_R(Y|X)\|_2^2 \cdot \|\text{deg}_S(Z|Y)\|_2^2 \cdot \|\text{deg}_T(X|Z)\|_2^2)^{1/3} \quad (10)$$

$$|Q_2| \leq (\|\text{deg}_R(Y|X)\|_3^3 \cdot \|\text{deg}_S(Z|Y)\|_3^3 \cdot |T|^5)^{1/6} \quad (11)$$

There can be many possible upper bounds that can be derived using the available  $\ell_p$ -norms. Our approach returns the smallest such bound, which depends on the data. The bounds based on arbitrary  $\ell_p$ -norms can be asymptotically tighter than the AGM and PANDA bounds, even for a single join. Preliminary experiments reported in Sec. 6 show that the upper bounds based on  $\ell_p$ -norms can be closer to the true cardinalities than the traditional cardinality estimators (e.g., used by POSTGRES, DUCKDB, and a commercial database system DBX), the theoretical AGM and PANDA bounds, and a learned cardinality estimator based on probabilistic graphical models. To achieve the best upper bound with our method, we used the  $\ell_p$ -norms for  $p \in \{1, \dots, 10, \infty\}$ .

In Section 3, we show how some of the above inequalities can be derived manually using our framework and that the best upper bound is the optimal solution of a linear program that uses the available  $\ell_p$ -norms of the degree sequences of the join attributes.

## 2. PRELIMINARIES

We first introduce the class of queries and data statistics under consideration and then give a brief background in information theory.

### 2.1 Queries and Data Statistics

For a number  $n$ , let  $[n] \stackrel{\text{def}}{=} \{1, 2, \dots, n\}$ . We use upper case  $X$  for variable names, and lower case  $x$  for values of these variables. We use boldface for sets of variables, e.g.,  $\mathbf{X}$ , and of constants, e.g.,  $\mathbf{x}$ .

$R$		
$X$	$Y$	
1	a	$\text{deg}_R(Y X) = \text{deg}_R(X Y) = (3, 2, 2, 1)$
1	b	$\ \text{deg}_R(Y X)\ _1 = 3 + 2 + 2 + 1 = 8$
1	c	
2	a	$\ \text{deg}_R(Y X)\ _\infty = \max(3, 2, 2, 1) = 3$
2	b	
3	b	$\ \text{deg}_R(Y X)\ _2 = (3^2 + 2^2 + 2^2 + 1^2)^{1/2} \leq 4.25$
3	c	
4	d	

**Figure 1: Relation  $R$  used in Example 2.1. The degree sequence  $\text{deg}_R(Y|X)$  states how many  $Y$ -values are paired with each  $X$ -value in  $R$ . The degree sequence  $\text{deg}_R(X|Y)$  states how many  $X$ -values are paired with each  $Y$ -value in  $R$ .**

A full conjunctive query (or query for short) is defined by:

$$Q(\mathbf{X}) = \bigwedge_{j \in [m]} R_j(\mathbf{Y}_j) \quad (12)$$

where  $\mathbf{Y}_j$  is the tuple of variables in  $R_j$  and  $\mathbf{X} = \bigcup_{j \in [m]} \mathbf{Y}_j$  is the set of  $n \stackrel{\text{def}}{=} |\mathbf{X}|$  variables in the query  $Q$ .

For a relation  $S$  and subsets  $\mathbf{U}, \mathbf{V}$  of its attributes, let  $\text{deg}_S(\mathbf{V}|\mathbf{U})$  be the degree sequence of  $\mathbf{U}$  in the projection  $\Pi_{\mathbf{U}\mathbf{V}}S$ . Formally, let  $G \stackrel{\text{def}}{=} (\Pi_{\mathbf{U}}(S), \Pi_{\mathbf{V}}(S), E)$  be the bipartite graph whose edges  $E$  are all pairs  $(\mathbf{u}, \mathbf{v}) \in \Pi_{\mathbf{U}\mathbf{V}}(S)$ . Then  $\text{deg}_S(\mathbf{V}|\mathbf{U}) \stackrel{\text{def}}{=} (d_1, d_2, \dots, d_m)$  is the degree sequence of the  $\mathbf{U}$ -nodes of the graph.

Fix  $\mathbf{X}$  a set of variables. An *abstract conditional*, or simply *conditional*, is an expression of the form  $\sigma = (\mathbf{V}|\mathbf{U})$ . We say that  $\sigma$  is *guarded* by a relation  $R(\mathbf{Y})$  if  $\mathbf{U}, \mathbf{V} \subseteq \mathbf{Y}$ ; then we write  $\text{deg}_R(\sigma) \stackrel{\text{def}}{=} \text{deg}_R(\mathbf{V}|\mathbf{U})$ . An *abstract statistics* is a pair  $\tau = (\sigma, p)$ , where  $p \in (0, \infty]$ . If  $B \geq 1$  is a real number, then we call the pair  $(\tau, B)$  a *concrete statistics*, and call  $(\tau, b)$ , where  $b \stackrel{\text{def}}{=} \log B$ , a *concrete log-statistics*. If  $R$  is a relation guarding  $\sigma$ , then we say that  $R$  *satisfies*  $(\tau, B)$  if  $\|\text{deg}_R(\sigma)\|_p \leq B$ . When  $p = 1$  then the statistics is a cardinality assertion on  $|\Pi_{\mathbf{U}\mathbf{V}}(R)|$ , and when  $p = \infty$  then it is an assertion on the maximum degree. We write  $\Sigma = \{\tau_1, \dots, \tau_s\}$  for a set of abstract statistics, and  $\mathbf{B} = \{B_1, \dots, B_s\}$  for an associated set of real numbers; thus, every pair  $(\tau_i, B_i)$  is a concrete statistics. We will call the pair  $(\Sigma, \mathbf{B})$  a *set of (concrete) statistics*, and call  $(\Sigma, \mathbf{b})$ , where  $b_i \stackrel{\text{def}}{=} \log B_i$ , a set of concrete *log-statistics*. We say that  $\Sigma$  is guarded by a relational schema  $\mathbf{R} = (R_1, \dots, R_m)$  if every  $\tau_i \in \Sigma$  has a guard  $R_{j_i}$ , and we say that a database instance  $\mathbf{D} = (R_1^D, \dots, R_m^D)$  *satisfies* the statistics  $(\Sigma, \mathbf{B})$ , denoted by  $\mathbf{D} \models (\Sigma, \mathbf{B})$ , if  $\|\text{deg}_{R_{j_i}^D}(\sigma_i)\|_{p_i} \leq B_i$  for all  $i \in [s]$ , where  $R_{j_i}$  is the guard of  $\sigma_i$ .

**EXAMPLE 2.1.** Fig. 1 depicts a relation  $R$  and two degree sequences: the degree sequence  $\text{deg}_R(Y|X)$  on column  $X$  and the degree sequence  $\text{deg}_R(X|Y)$  on column  $Y$ . They are both the same for this specific relation:  $(3, 2, 2, 1)$ . There are 4 distinct  $X$ -values: 1 occurs 3 times (so its degree is 3), 2 occurs 2 times, 3 occurs 2 times, and 4 occurs 1 time. Similarly, there are 4 distinct  $Y$ -values:  $b$  occurs 3 times,  $a$  occurs 2 times,  $c$  occurs 2 times, and  $d$  occurs 1 time.

The figure also exemplifies three norms on  $\text{deg}_R(Y|X)$ : The  $\ell_1$  norm amounts to summing up the degrees in the

sequence, and therefore gives the size of  $R$ , the  $\ell_\infty$  norm gives the maximum degree in the sequence, and the  $\ell_2$  norm is the square root of the sum of the squares of the degrees.

Using the terminology introduced in this section,  $(Y|X)$  and  $(X|Y)$  are abstract conditionals and are guarded by  $R$ . Fig. 1 shows the abstract statistics  $((Y|X), 1)$ ,  $((Y|X), \infty)$ , and  $((Y|X), 2)$  and concrete statistics based on them:  $((Y|X), 1, 8)$ ,  $((Y|X), \infty, 3)$ , and  $((Y|X), 2, 4.25)$  that are satisfied by  $R$ .

A degree sequence can be recovered exactly from its  $\ell_p$ -norms. In practice, a few norms are often enough to approximate well the degree sequence. In our experiments with the real-world dataset IMDB, the degree sequences have sizes up to a few thousands and a long tail of degree 1. For the purpose of cardinality estimation [13], they can be approximated well by the  $\ell_p$ -norms for  $p \in \{1, \dots, 10, \infty\}$ .

## 2.2 Background in Information Theory

Consider a finite probability space  $(D, P)$ , where  $P : D \rightarrow [0, 1]$ ,  $\sum_{x \in D} P(x) = 1$ , and denote by  $X$  the random variable with outcomes in  $D$ . The *entropy* of  $X$  is:

$$H(X) \stackrel{\text{def}}{=} - \sum_{x \in D} P(x) \log P(x) \quad (13)$$

If  $N \stackrel{\text{def}}{=} |D|$ , then  $0 \leq H(X) \leq \log N$ . The equality  $H(X) = 0$  holds iff  $X$  is deterministic, and  $H(X) = \log N$  holds iff  $X$  is uniformly distributed. Given  $n$  jointly distributed random variables  $\mathbf{X} = \{X_1, \dots, X_n\}$ , we denote by  $\mathbf{h} \in \mathbb{R}_+^{2^{[n]}}$  the following vector:  $h_\alpha \stackrel{\text{def}}{=} H(\mathbf{X}_\alpha)$  for  $\alpha \subseteq [n]$ , where  $\mathbf{X}_\alpha$  is the joint random variable  $(X_i)_{i \in \alpha}$ , and  $H(\mathbf{X}_\alpha)$  is its entropy; such a vector  $\mathbf{h} \in \mathbb{R}_+^{2^{[n]}}$  is called *entropic*. We will blur the distinction between a vector in  $\mathbb{R}_+^{2^{[n]}}$ , a vector in  $\mathbb{R}_+^{2^{\mathbf{X}}}$ , and a function  $2^{\mathbf{X}} \rightarrow \mathbb{R}_+$ , and write interchangeably  $\mathbf{h}_\alpha$ ,  $\mathbf{h}_{\mathbf{X}_\alpha}$ , or  $h(\mathbf{X}_\alpha)$ . A *polymatroid* is a vector  $\mathbf{h} \in \mathbb{R}_+^{2^{[n]}}$  that satisfies the following *basic Shannon inequalities*:

$$h(\emptyset) = 0 \quad (14)$$

$$h(\mathbf{U} \cup \mathbf{V}) \geq h(\mathbf{U}) \quad (15)$$

$$h(\mathbf{U}) + h(\mathbf{V}) \geq h(\mathbf{U} \cup \mathbf{V}) + h(\mathbf{U} \cap \mathbf{V}) \quad (16)$$

The last two inequalities are called *monotonicity* and *submodularity* respectively.

The *conditional* of a vector  $\mathbf{h}$  is defined as:

$$h(\mathbf{V}|\mathbf{U}) \stackrel{\text{def}}{=} h(\mathbf{UV}) - h(\mathbf{U})$$

where  $\mathbf{U}, \mathbf{V} \subseteq \mathbf{X}$ . If  $\mathbf{h}$  is a polymatroid, then  $h(\mathbf{V}|\mathbf{U}) \geq 0$ . If  $\mathbf{h}$  is entropic and realized by some probability distribution, then:

$$h(\mathbf{V}|\mathbf{U}) = \mathbb{E}_{\mathbf{u}}[h(\mathbf{V}|\mathbf{U} = \mathbf{u})] \quad (17)$$

where  $h(\mathbf{V}|\mathbf{U} = \mathbf{u})$  is the standard entropy of the random variable  $\mathbf{V}$  conditioned on  $\mathbf{U} = \mathbf{u}$ .

An *information inequality* is a linear inequality of the form:

$$\mathbf{c} \cdot \mathbf{h} \geq 0, \text{ where } \mathbf{c} \in \mathbb{R}^{2^{[n]}}. \quad (18)$$

## 3. BOUNDS ON QUERY OUTPUT SIZE

In recent work [12], we solved Problem 1 below for any query  $Q$ , database  $\mathbf{D}$ , and statistics  $(\Sigma, \mathbf{B})$  consisting of  $\ell_p$ -norms of degree sequences:

PROBLEM 1. Given a query  $Q$  and a set of statistics  $(\Sigma, \mathbf{B})$  guarded by (the schema of)  $Q$ , find a bound  $U \in \mathbb{R}$  such that for all database instances  $\mathbf{D}$ , if  $\mathbf{D} \models (\Sigma, \mathbf{B})$ , then  $|Q(\mathbf{D})| \leq U$ .

The upper bound  $U$  is *tight*, if there exists a database instance  $\mathbf{D}$  such that  $\mathbf{D} \models (\Sigma, \mathbf{B})$  and  $U = O(|Q(\mathbf{D})|)$ .

The key observation is that the concrete statistics

$$\|\mathbf{deg}(\mathbf{V}|\mathbf{U})\|_p \leq B$$

implies the following inequality in information theory:

$$\frac{1}{p}h(\mathbf{U}) + h(\mathbf{V}|\mathbf{U}) \leq \log B \quad (19)$$

where  $h$  is the entropy of some probability distribution on a relation  $R$  that guards the conditional  $(\mathbf{V}|\mathbf{U})$ .

Using (19) we prove the following general upper bound on the query output size. In the following, the query variables become random variables when we take a probability distribution over the tuples in the query output:

THEOREM 3.1 ([12]). Let  $Q$  be a query (12),  $\mathbf{U}_i, \mathbf{V}_i \subseteq \mathbf{X}$  be sets of variables, for  $i \in [s]$ , and suppose that the following information inequality is valid for all entropic vectors  $\mathbf{h}$  with variables  $\mathbf{X}$ :

$$\sum_{i \in [s]} w_i \left( \frac{1}{p_i} h(\mathbf{U}_i) + h(\mathbf{V}_i|\mathbf{U}_i) \right) \geq h(\mathbf{X}) \quad (20)$$

where  $w_i \geq 0$ , and  $p_i \in (0, \infty]$ , for all  $i \in [s]$ . Assume that each conditional  $(\mathbf{V}_i|\mathbf{U}_i)$  in (20) is guarded by some relation  $R_{j_i}$  in  $Q$ . Then, for any database instance  $\mathbf{D} = (R_1^{\mathbf{D}}, R_2^{\mathbf{D}}, \dots)$ , the following upper bound holds on the query output size:

$$|Q| \leq \prod_{i \in [s]} \|\mathbf{deg}_{R_{j_i}^{\mathbf{D}}}(\mathbf{V}_i|\mathbf{U}_i)\|_{p_i}^{w_i} \quad (21)$$

Thus, one approach to find an upper bound on the query output is to find an inequality of the form (20), prove it using Shannon inequalities, then conclude that (21) holds.

EXAMPLE 3.2. Let us prove the bound (5):

$$\|\mathbf{deg}_R(X|Y)\|_2 \cdot \|\mathbf{deg}_S(Z|Y)\|_2 \geq |Q_1|$$

using the following data statistics inequalities:

$$\begin{aligned} \log \|\mathbf{deg}_R(X|Y)\|_2 &\geq \frac{1}{2}h(Y) + h(X|Y) \\ \log \|\mathbf{deg}_S(Z|Y)\|_2 &\geq \frac{1}{2}h(Y) + h(Z|Y) \end{aligned}$$

We next sum up these two inequalities:

$$\begin{aligned} \log (\|\mathbf{deg}_R(X|Y)\|_2 \cdot \|\mathbf{deg}_S(Z|Y)\|_2) &\geq \frac{1}{2}h(Y) + h(X|Y) + \frac{1}{2}h(Y) + h(Z|Y) \\ &= h(Y) + h(X|Y) + h(Z|Y) \\ &= h(XY) + h(Z|Y) \\ &\geq h(XY) + h(Z|XY) \\ &= h(XYZ) = \log |Q_1| \end{aligned}$$

The last inequality follows from the observation that conditioning  $h(Z|Y)$  further does not increase its entropy (but it can either keep it the same or decrease it).

EXAMPLE 3.3. We now prove the bound (11):

$$(\|\mathbf{deg}_R(Y|X)\|_3^3 \cdot \|\mathbf{deg}_S(Y|Z)\|_3^3 \cdot \|\mathbf{deg}_T(X|Z)\|_1^5)^{1/6} \geq |Q_2|$$

using the following data statistics inequalities:

$$\log \|\mathbf{deg}_R(Y|X)\|_3 \geq \frac{1}{3}h(X) + h(Y|X) \quad (22)$$

$$\log \|\mathbf{deg}_S(Y|Z)\|_3 \geq \frac{1}{3}h(Z) + h(Y|Z) \quad (23)$$

$$\log \|\mathbf{deg}_T(X|Z)\|_1 \geq h(Z) + h(X|Z) = h(XZ) \quad (24)$$

We next sum up three times the first inequality, three times the second inequality, and five times the last inequality:

$$\log (\|\mathbf{deg}_R(Y|X)\|_3^3 \cdot \|\mathbf{deg}_S(Y|Z)\|_3^3 \cdot \|\mathbf{deg}_T(X|Z)\|_1^5) \geq \quad (25)$$

$$h(X) + 3h(Y|X) + h(Z) + 3h(Y|Z) + 5h(XZ) = \quad (26)$$

$$h(XY) + h(YZ) + h(XZ) + \quad (27)$$

$$2h(Y|X) + 2h(XZ) + \quad (28)$$

$$2h(Y|Z) + 2h(XZ) \quad (29)$$

In the above, we decompose the sum (26) into three sums (27)-(29) and show that each of these three sums upper bounds  $2h(XYZ)$ . This concludes that (25) upper bounds  $6h(XYZ)$ , which is the same as  $6 \log |Q_2|$  or  $\log (|Q_2|^6)$ . Since the log function is monotone, this implies the desired bound (11).

We prove the bound for the sum (27):

$$\begin{aligned} h(XY) + h(YZ) + h(XZ) &= h(X) + h(Y|X) + h(YZ) + h(XZ) \\ &\geq h(X) + h(Y|XZ) + h(YZ) + h(XZ) \\ &= h(X) + h(YZ) + h(XYZ) \geq h(XYZ) + h(XYZ) \end{aligned}$$

We explain the derivation steps. We decompose  $h(XY)$  using the definition of conditional entropy. We then condition  $h(Y|X)$  further to obtain  $h(Y|XZ)$ , which cannot be larger than  $h(Y|X)$ . We use the definition of entropy to obtain  $h(XYZ) = h(Y|XZ) + h(XZ)$ . Finally, we apply the submodularity inequality (16):  $h(X) + h(YZ) \geq h(XYZ) + h(\emptyset)$ .

We next prove the bound for the sums (28) and (29):

$$2h(Y|X) + 2h(XZ) \geq 2h(Y|XZ) + 2h(XZ) = 2h(XYZ)$$

$$2h(Y|Z) + 2h(XZ) \geq 2h(Y|XZ) + 2h(XZ) = 2h(XYZ)$$

One way to describe the solution to Problem 1 is as follows. Consider a set of statistics  $(\Sigma, \mathbf{B})$ . Any valid information inequality (20) implies some bound on the query output size, namely  $|Q| \leq \prod_{i \in [s]} B_{j_i}^{w_i}$ . The best bound is their *minimum*, over all valid information inequalities (20). Because characterizing all such inequalities is a major open problem in information theory, we only consider a subset of them, namely Shannon inequalities. We denote the log of this minimum by *Log-U-Bound*. This describes the solution to Problem 1 as a minimization problem. In order to compute *Log-U-Bound*, we describe below an alternative, dual characterization, as a *maximization* problem, by considering the following quantity:

$$\text{Log-L-Bound} = \max_{\mathbf{h} \models (\Sigma, \mathbf{b})} h(\mathbf{X}) \quad (30)$$

where  $\mathbf{X}$  is the set of all variables in the query  $Q$ , and  $\mathbf{h}$  ranges over all polymatroids that “satisfy” the concrete log-statistics  $(\Sigma, \mathbf{b})$ , i.e., Inequality (19) is satisfied for every statistics in  $\Sigma$ . The following theorem follows from the duality of linear programming:

THEOREM 3.4. *Log-L-Bound = Log-U-Bound.*

Theorems 3.1 and 3.4 imply:

$$\log |Q| \leq \text{Log-L-Bound} \quad (31)$$

We refer to *Log-L-Bound* as the *polymatroid bound*. A bound is *tight* if there exists a database  $\mathbf{D}$  such that the size of the query output is  $|Q| \geq c \cdot 2^{\text{Log-L-Bound}}$ , where  $c$  is a constant that depends only on the query  $Q$ . The polymatroid bound is not tight in general. Yet it becomes tight for *simple* degree sequences. A degree sequence  $\text{deg}_R(\mathbf{V}|\mathbf{U})$  is *simple* if  $|\mathbf{U}| \leq 1$ . In case of simple degree sequences, the worst-case database  $\mathbf{D}$  has a special form, called a *normal database* [12].

Equation (30) gives us an effective method for solving Problem 1, when (20) are restricted to Shannon inequalities, because in that case, *Log-L-Bound* is the optimal value of a linear program. We exemplify two such linear programs for the simple join and triangle queries.

EXAMPLE 3.5. The linear program for the simple join  $Q_1(X, Y, Z) = R(X, Y) \wedge S(Y, Z)$  in (1) is:

$$\begin{aligned} & \max h(X, Y, Z) \text{ subject to:} \\ & \quad // \text{ basic Shannon inequalities: Eq. (14) - (16)} \\ & \quad // \forall \mathbf{U}, \mathbf{V} \subseteq \{X, Y, Z\} \\ & \quad h(\mathbf{U}) \leq h(\mathbf{U} \cup \mathbf{V}) \\ & \quad h(\mathbf{U} \cup \mathbf{V}) + h(\mathbf{U} \cap \mathbf{V}) \leq h(\mathbf{U}) + h(\mathbf{V}) \\ & \quad h(\emptyset) = 0 \\ & \quad // \text{ data statistics inequalities } p_1, p_2 \in \{1, 2, \dots, \infty\} \\ & \quad \frac{1}{p_1} h(Y) + h(X|Y) \leq \log \|\text{deg}_R(X|Y)\|_{p_1} \\ & \quad \frac{1}{p_2} h(Y) + h(Z|Y) \leq \log \|\text{deg}_S(Z|Y)\|_{p_2} \end{aligned}$$

The linear program uses  $\ell_p$ -norms on the degree sequences  $\text{deg}_R(X|Y)$  and  $\text{deg}_S(Z|Y)$  of the join columns of the two relations. For each available  $\ell_p$ -norm, the program has one constraint that is an instantiation of the key inequality (19).

EXAMPLE 3.6. The linear program for the triangle query  $Q_2(X, Y, Z) = R(X, Y) \wedge S(Y, Z) \wedge T(Z, X)$  in (7) is:

$$\begin{aligned} & \max h(X, Y, Z) \text{ subject to:} \\ & \quad // \text{ basic Shannon inequalities in Eq. (14) - (16)} \\ & \quad \dots \quad \dots \\ & \quad // \text{ data statistics inequalities } p_1, p_2 \in \{1, 2, \dots, \infty\} \\ & \quad \frac{1}{p_1} h(X) + h(Y|X) \leq \log \|\text{deg}_R(Y|X)\|_{p_1} \\ & \quad \frac{1}{p_2} h(Y) + h(X|Y) \leq \log \|\text{deg}_R(X|Y)\|_{p_2} \\ & \quad \text{Similarly for } \text{deg}_S(Z|Y) \text{ and } \text{deg}_S(Y|Z) \\ & \quad \text{Similarly for } \text{deg}_T(X|Z) \text{ and } \text{deg}_T(Z|X) \end{aligned}$$

Each relation joins on each of its two columns, so the above linear program uses the degree sequences of each of these join columns:  $\text{deg}_R(Y|X)$  and  $\text{deg}_R(X|Y)$  for relation  $R$  and similarly for relations  $S$  and  $T$ . For each available  $\ell_p$ -norm on each of these degree sequences, there is one constraint that is an instantiation of the key inequality (19).

In the above linear programs, which data statistics inequalities are used depends on the availability of concrete statistics, e.g., which  $\ell_p$ -norms are available for which degree sequences.




Figure 2: Relations  $(R_p)_{p \geq 1}$  defined as the union of  $k$  Cartesian products of blocks of values for the attributes  $A, B_1, \dots, B_p$ .

#### 4. DEMYSTIFYING THE $\ell_p$ -NORM INFORMATION INEQUALITY

In this section, we look closer at the key inequality (19) behind the  $\ell_p$ -bound approach, instantiated for the case of an arbitrary binary relation  $R_1$  with attributes  $A$  and  $B_1$ :

$$\frac{1}{p} h(A) + h(B_1|A) \leq \log \|\text{deg}_{R_1}(B_1|A)\|_p. \quad (32)$$

Using  $R_1$ , we construct the relations  $R_p$ , parameterized by  $p > 1$ , with attributes  $A, B_1, \dots, B_p$ , as depicted in Fig. 2. In particular,  $R_1(A, B_1)$  contains  $k$  different values for  $A$ , namely  $a_1, \dots, a_k$ , where the  $i$ -th value  $a_i$  has  $n_i$  different values for  $B_1$ , namely  $b_1^{(i)}, \dots, b_{n_i}^{(i)}$ . In contrast,  $R_p(A, B_1, \dots, B_p)$  contains  $k$  different values for  $A$ , and for each value  $a_i$  of  $A$ , it contains the Cartesian product of  $p$  copies of the set  $\{b_1^{(i)}, \dots, b_{n_i}^{(i)}\}$ ; one copy for each attribute  $B_1, \dots, B_p$ . By construction, the size of  $R_p$  is  $|R_p| = \sum_{i \in [k]} n_i^p$ . Also, for any  $A$ -value  $a$ ,  $\pi_{B_i}(\sigma_{A=a} R_p) = \pi_{B_j}(\sigma_{A=a} R_p)$  for all  $i, j \in [p]$ . Furthermore,  $R_p$  satisfies the join dependency  $(AB_1, \dots, AB_p)$ , i.e.,  $\bowtie_{i \in [p]} \pi_{AB_i} R_p = R_p$ , and therefore  $B_1$  to  $B_p$  are pairwise independent given  $A$ .

The construction of  $R_p$  is used to relate the  $\ell_1$ -norm of the degree sequence  $\text{deg}_{R_p}(B_1 \dots B_p|A)$  and the  $\ell_p$ -norm of the degree sequence  $\text{deg}_{R_1}(B_1|A)$ :

$$|R_p| = \|\text{deg}_{R_p}(B_1 \dots B_p|A)\|_1 = \sum_{i \in [k]} n_i^p = \|\text{deg}_{R_1}(B_1|A)\|_p^p.$$

Let  $P : R_p \rightarrow [0, 1]$  be any probability distribution whose outcomes are the tuples in  $R_p$ . This also defines a probability distribution over the tuples in  $R_i$  for  $i \in [p]$ . Let  $\mathbf{h} : 2^{\{A, B_1, \dots, B_p\}} \rightarrow \mathbb{R}_+$  be the entropic vector of  $P$ , so its elements are the joint entropies of the subsets of the set  $\{A, B_1, \dots, B_p\}$  of the attributes of  $R_p$ .

We first explain why (32) holds for  $p = 1$  and  $p = \infty$ . We then discuss the remaining cases.

For  $p = 1$ , the inequality becomes

$$h(A) + h(B_1|A) = h(AB_1) \leq \log \|\text{deg}_{R_1}(B_1|A)\|_1 = \log |R_1|.$$

To see why this holds, recall that the maximum joint entropy  $h(AB_1)$  is obtained for the uniform probability distribution

over the tuples of  $R_1$ , so each tuple  $t \in R_1$  has the same probability  $P(t) = \frac{1}{|R_1|}$ . Then, by the definition of the entropy,  $h(AB_1) = -\sum_{t \in R_1} P(t) \cdot \log P(t) = -\sum_{t \in R_1} \frac{1}{|R_1|} \cdot \log \frac{1}{|R_1|} = \frac{1}{|R_1|} \cdot \log |R_1| \cdot \sum_{t \in R_1} 1 = \log |R_1|$ .

For  $p = \infty$ , the inequality becomes

$$h(B_1|A) \leq \log \|\mathbf{deg}_{R_1}(B_1|A)\|_\infty = \max_{i \in [k]} n_i.$$

To see why this holds, recall from (17) that  $h(B_1|A) = \mathbb{E}_a[h(B_1|A = a)]$ . Furthermore,  $h(B_1|A = a) \leq \log |\sigma_{A=a} R_1|$ , as discussed above for  $p = 1$ . The number of tuples in  $R_1$  with  $A = a$  is the degree of the  $A$ -value  $a$  in  $R_1$ , thus:  $h(B_1|A = a) \leq \log \mathbf{deg}_{R_1}(B_1|A = a) \leq \log \max_{a'} \mathbf{deg}_{R_1}(B_1|A = a') = \log \|\mathbf{deg}_{R_1}(B_1|A)\|_\infty$ .

For any other value of  $p$ , the challenge is to connect the entropy, which is expressed using the probability distribution over the tuples in  $R_1$ , and the  $\ell_p$ -norm of the degree sequence  $\mathbf{deg}_{R_1}(B_1|A)$ . This can be achieved by connecting the  $\ell_p$ -norm of this degree sequence and the  $\ell_1$ -norm of the degree sequence  $\mathbf{deg}_{R_p}(B_1 \dots B_p|A)$ .

Let us first exemplify for  $p = 2$ . The relation  $R_2$  satisfies the join dependency  $(AB_1, AB_2)$  and  $B_1$  and  $B_2$  are independent and identically distributed given  $A$ . In the language of entropy, this means:

$$\begin{aligned} h(AB_1B_2) &= h(A) + h(B_1B_2|A) \\ &= h(A) + h(B_1|A) + h(B_2|A) \\ &= h(A) + 2h(B_1|A). \end{aligned}$$

By construction,  $|R_2| = \sum_{i \in [k]} n_i^2 = \|\mathbf{deg}_{R_1}(B_1|A)\|_2^2$ . We also know that  $h(B_1B_2|A) \leq \log |R_2|$  using the above argument for  $R_1$  and  $p = 1$ . We now put the aforementioned observations together to obtain:

$$\begin{aligned} h(A) + 2h(B_1|A) &= h(AB_1B_2) \leq \\ \log |R_2| &= \log \|\mathbf{deg}_{R_1}(B_1|A)\|_2^2 = 2 \log \|\mathbf{deg}_{R_1}(B_1|A)\|_2 \end{aligned}$$

We divide both sides of the inequality by 2 and obtain the desired inequality:

$$\frac{1}{2}h(A) + h(B_1|A) \leq \log \|\mathbf{deg}_{R_1}(B_1|A)\|_2.$$

For arbitrary  $p \in (1, \infty)$ , the argument is similar to  $p = 2$ :

$$\begin{aligned} h(AB_1 \dots B_p) &= h(A) + h(B_1 \dots B_p|A) \\ &= h(A) + h(B_1|A) + \dots + h(B_p|A) \\ &= h(A) + ph(B_1|A). \end{aligned}$$

Furthermore,  $|R_p| = \sum_{i \in [k]} n_i^p = \|\mathbf{deg}_{R_1}(B_1|A)\|_p^p$ . Then,

$$\begin{aligned} h(A) + ph(B_1|A) &= h(AB_1 \dots B_p) \leq \\ \log |R_p| &= \log \|\mathbf{deg}_{R_1}(B_1|A)\|_p^p = p \log \|\mathbf{deg}_{R_1}(B_1|A)\|_p \end{aligned}$$

By dividing both sides of the inequality by  $p$ , we obtain (32). For an alternative proof, which works for any real value  $p > 0$ , we refer the reader to the original paper [12].

## 5. ALGORITHM MEETING THE BOUND

A second application of the  $\ell_p$ -bounds is for query evaluation: If Inequality (20) holds for all polymatroids, then we can evaluate the query in time bounded by (21) multiplied by a poly-logarithmic factor in the data and an exponential factor in the sum of the  $p$  values of the statistics.

This new evaluation algorithm for conjunctive queries generalizes the PANDA algorithm [8, 19] from  $\ell_1$  and  $\ell_\infty$  norms to arbitrary norms. Recall that PANDA starts from an inequality of the form (20), where every  $p_i$  is either 1 or  $\infty$ , and computes the query in time  $O(\prod_i B_i^{w_i})$  if the database satisfies  $|\Pi_{U_i V_i}(R_{j_i})| \leq B_i$  when  $p_i = 1$  and  $\|\mathbf{deg}_{R_{j_i}}(V_i|U_i)\|_\infty \leq B_i$  when  $p_i = \infty$ . Our algorithm uses PANDA as a black box, as follows. It first partitions the relations on the join columns so that, within each partition, all degrees are within a factor of two, and each statistics defined by some  $\ell_p$ -norm on the degree sequence of the join column can be expressed alternatively using only  $\ell_1$  and  $\ell_\infty$ . The original query becomes a union of queries, one per combination of parts of different relations. The algorithm then evaluates each of these queries using the PANDA algorithm.

We describe next the details of data partitioning and the reduction to PANDA. Consider a relation  $R$  with attributes  $\mathbf{X}$  and a concrete statistics  $(\tau, B)$ , where  $\tau = ((\mathbf{V}|U), p)$ . We say that  $R$  *strongly satisfies*  $(\tau, B)$ , in notation  $R \models_s (\tau, B)$ , if there exists a number  $d > 0$  such that  $\|\mathbf{deg}_R(\mathbf{V}|U)\|_\infty \leq d$  and  $|\Pi_U(R)| \leq B^p/d^p$ . If  $R \models_s (\tau, B)$  then  $R \models (\tau, B)$  because:

$$\|\mathbf{deg}_R(\mathbf{V}|U)\|_p^p \leq |\Pi_U(R)| \cdot \|\mathbf{deg}_R(\mathbf{V}|U)\|_\infty^p \leq \frac{B^p}{d^p} d^p = B^p$$

In other words,  $R$  strongly satisfies the  $\ell_p$  statistics  $(\tau, B)$  if it satisfies an  $\ell_1$  and an  $\ell_\infty$  statistics that imply  $(\tau, B)$ .

**LEMMA 5.1.** *Fix a join query  $Q$ , and suppose that inequality (20) holds for all polymatroids. Let  $\Sigma = \{(V_i|U_i, p_i) \mid i \in [s]\}$  be the abstract statistics and  $w_i \geq 0$  be the coefficients in (20). If a database  $\mathbf{D}$  strongly satisfies the concrete statistics  $(\Sigma, \mathbf{B})$ , then the query output can be computed in time  $O(\prod_{i \in [s]} B_i^{w_i} \text{polylog } N)$ , where  $N$  is the size of the active domain of  $\mathbf{D}$ .*

In order to use Lemma 5.1, we need the following:

**LEMMA 5.2.** *Let  $R$  be a relation that satisfies an  $\ell_p$ -statistics,  $R \models ((\mathbf{V}|U), p, B)$ . Then we can partition  $R$  into  $\lceil 2^p \rceil \log N$  disjoint relations,  $R = R_1 \cup R_2 \cup \dots$ , such that each  $R_i$  strongly satisfies the  $\ell_p$ -statistics,  $R_i \models_s (((\mathbf{V}|U), p), B)$ .*

To compute  $Q$  in a runtime upper-bounded by the  $\ell_p$ -bound in (20), we can proceed as follows. Using Lemma 5.2, for each  $\ell_{p_i}$ -norm, we partition  $\mathbf{D}$  into a union of  $2^{p_i}$  databases  $\mathbf{D}_1 \cup \mathbf{D}_2 \cup \dots$ , where each  $\mathbf{D}_j$  strongly satisfies  $(\Sigma, \mathbf{B})$ . Resolving  $s$  such norms like this partitions  $\mathbf{D}$  into  $c$  parts. We then apply Lemma 5.1 to each part. This implies:

**THEOREM 5.3.** *There exists an algorithm that, given a join query  $Q$ , an inequality (20) holding for all polymatroids, and a database  $\mathbf{D}$  satisfying the concrete statistics  $(\Sigma, \mathbf{B})$ , computes the query output in time  $O(c \cdot \prod_{i \in [s]} B_i^{w_i} \text{polylog } N)$ , where  $c = \prod_{i \in [s]} \lceil 2^{p_i} \rceil$  and  $p_1, \dots, p_s$  are the norms in  $\Sigma$ .*

## 6. PRELIMINARY EXPERIMENTS

In this section, we report on preliminary experiments with three types of cardinality estimators: pessimistic, traditional, and learned cardinality estimators.

We consider three pessimistic estimators: LPBOUND, which implements the  $\ell_p$ -bounds approach discussed in this paper

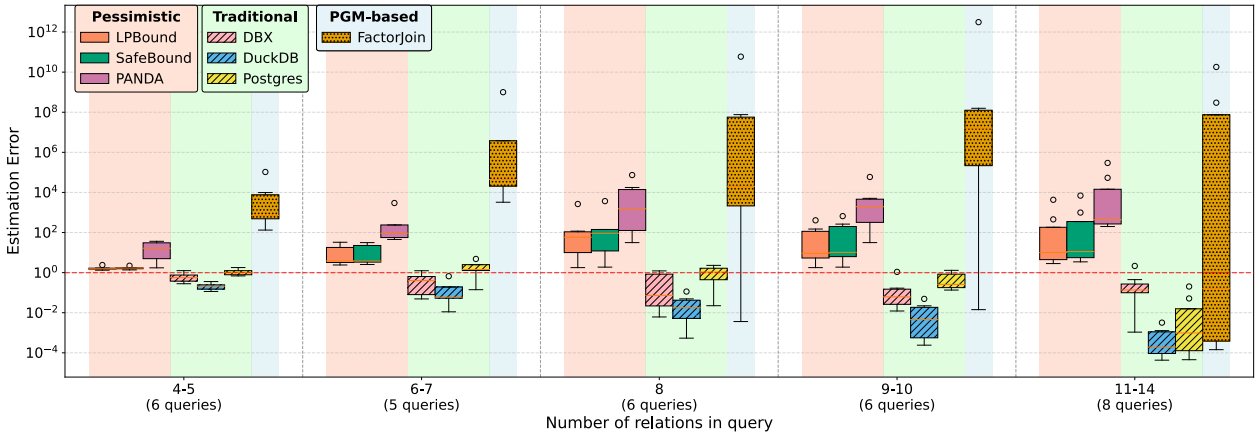


Figure 3: Estimation errors for 31 Berge-acyclic full conjunctive queries in the JOBjoin benchmark.

and uses the  $\ell_p$  norms on the degree sequences of the join columns, for  $p \in \{1, \dots, 10, \infty\}$ ; SAFEBOUND [20], which uses a lossy compression of the degree sequences on the join columns and only works for full Berge-acyclic conjunctive queries; and PANDA, which is LPBOUND restricted to the  $\ell_1$  and  $\ell_\infty$  norms. The AGM bounds, which only use the  $\ell_1$  norm, are too high (up to 30 orders of magnitude higher) relative to LPBOUND and not shown. We also consider the three traditional estimators used in the two open-source systems POSTGRES 13.14 and DUCKDB 0.10.1 and in a commercial system DBX. We finally consider FACTORJOIN [21], a learned estimator based on probabilistic graphical models. At the time of writing, this was the only data-driven learned cardinality estimator that worked on our workload.

For the aforementioned cardinality estimators, we measure their accuracy by means of *estimation error*, which is the estimation produced by the estimator divided by the true query output size (in case the true output size is 0, the error is just the estimation).

## 6.1 Experiments with Berge-Acyclic Queries

Fig. 3 compares the estimation error of these estimators on 32 full conjunctive queries (Berge-acyclic, no filter predicates, no group-by) constructed as the natural join of some of the relations in the IMDB dataset (3.7GB). The estimation error is shown for the queries divided into five groups depending on the number of their relations.

The pessimistic estimators (LPBOUND, SAFEBOUND, and PANDA) provide guaranteed upper bounds, as expected. The effect of using more norms is evident when comparing PANDA and LPBOUND: PANDA’s errors are at least one order of magnitude higher than LPBOUND’s. SAFEBOUND performs slightly worse than LPBOUND: Both estimators employ lossy compressions of the degree sequences, LPBOUND uses the  $\ell_p$ -norms whereas SAFEBOUND uses a piecewise linear approximation. The figure shows that the approximation by  $\ell_p$ -norms is more accurate than the piecewise linear approximation; a different experiment also shows that the former also takes less space than the latter.

The traditional estimators both underestimate and overestimate. As we increase the number of relations in the query, their underestimation increases. The learned estimator FAC-

TORJOIN both underestimates and overestimates and has a larger error than the other estimators.

## 6.2 Experiments with Triangle Join Queries

The  $\ell_p$ -bounds can be also effective for cyclic queries. We report on a limited exploration of the usefulness of  $\ell_p$ -bounds for two triangle join queries, which differ in the third copy of the edge relation:

$$Q_3(X, Y, Z) = E(X, Y) \wedge E(Y, Z) \wedge E(Z, X)$$

$$Q_4(X, Y, Z) = E(X, Y) \wedge E(Y, Z) \wedge E(X, Z)$$

where  $E$  is the edge relation of one of the 7 graphs from the SNAP repository [22]<sup>2</sup>. Four of these graphs are directed (soc-Epinions, soc-LiveJournal, facebook, and twitter). The two queries have different outputs on the directed graphs and the same output on the remaining undirected graphs.

Table 1 reports the estimation errors for the  $\ell_p$ -bound approach, where we only use the  $\ell_1$ -norm to recover the AGM bound [5], the  $\ell_{\{1, \infty\}}$  norms to recover the PANDA’s polymatroid bound [8], and the  $\ell_2$ -norm alone. We also report the estimation errors for the traditional estimators DUCKDB, POSTGRES<sup>3</sup>, and DBX. SAFEBOUND and FACTORJOIN do not support cyclic queries.

Although the  $\ell_2$ -bound is the same for both  $Q_3$  and  $Q_4$  on each of the graphs, it uses different, albeit equivalent, upper bound formulas for the two queries, as explained next. Out of the degree sequences on the two possible columns of the edge relation, the  $\ell_2$ -bound formulas always use the degree sequence with the smallest  $\ell_2$ -norm. For the directed graphs except facebook, this is the first column; for the undirected graphs, the two degree sequences are the same.

For  $Q_3$ , the  $\ell_2$ -bound formula is:

$$(\|\text{deg}_E(Y|X)\|_2 \cdot \|\text{deg}_E(Z|Y)\|_2 \cdot \|\text{deg}_E(X|Z)\|_2)^{1/3}$$

<sup>2</sup> We removed the duplicates in the twitter SNAP dataset before processing, the other datasets do not have duplicates.

<sup>3</sup> POSTGRES’s accuracy is sensitive to its sampling size used to determine the domain sizes of the data columns in the input database. We chose a sampling size of 3M. For sampling sizes 3K, 30K, 300K, 1.5M, POSTGRES gives different accuracies for different datasets, with no clear winner.

$Q_3(X, Y, Z) = E(X, Y) \wedge E(Y, Z) \wedge E(Z, X)$						
Dataset	{1}	{1, $\infty$ }	{2}	DuckDB	POSTGRES	DBX
ca-GrQc	1.6E+06	7.5E+05	1.6E+05	1.4E+05	2.1E+02	1.3E+04
ca-HepTh	6.9E+01	2.0E+01	3.8E+00	5.2E+00	8.8E+01	3.0E-01
facebook	2.6E+07	2.6E+07	5.4E+06	2.8E+07	2.8E+04	9.4E+04
soc-Epinions	1.6E+02	1.6E+02	2.5E+01	7.6E+01	4.9E+00	2.2E-01
soc-LiveJournal	7.8E+02	7.8E+02	1.0E+01	3.3E+01	3.5E+03	1.0E-01
soc-pokec	2.1E+03	2.1E+03	2.8E+01	1.6E+02	7.9E+02	4.0E-01
twitter	1.1E+02	1.1E+02	7.2E+00	5.6E+01	1.4E+00	8.5E-02

$Q_4(X, Y, Z) = E(X, Y) \wedge E(Y, Z) \wedge E(X, Z)$						
Dataset	{1}	{1, $\infty$ }	{2}	DuckDB	POSTGRES	DBX
ca-GrQc	3.3E+01	1.6E+01	3.4E+00	3.0E+00	5.3E+01	2.8E-01
ca-HepTh	6.9E+01	2.0E+01	3.8E+00	5.2E+00	8.8E+01	3.0E-01
facebook	1.6E+01	1.4E+01	3.3E+00	1.7E+01	9.5E+00	5.8E-02
soc-Epinions	1.0E+02	1.0E+02	1.5E+01	4.7E+01	2.5E+00	1.4E-01
soc-LiveJournal	6.1E+02	6.1E+02	7.7E+00	2.6E+01	1.5E+03	8.0E-02
soc-pokec	1.8E+03	1.8E+03	2.4E+01	1.3E+02	5.8E+02	6.2E-01
twitter	7.3E+01	6.6E+01	4.7E+00	3.6E+01	1.1E+00	5.5E-02

**Table 1: Estimation errors for AGM ( $\ell_1$ ), PANDA ( $\ell_1, \ell_\infty$ ),  $\ell_p$ -bounds (only  $\ell_2$ ), DuckDB, Postgres, and DbX for the triangle join queries  $Q_3(X, Y, Z) = E(X, Y) \wedge E(Y, Z) \wedge E(Z, X)$  and  $Q_4(X, Y, Z) = E(X, Y) \wedge E(Y, Z) \wedge E(X, Z)$  on SNAP graph datasets. The best estimation error for each dataset is in bold. The estimation errors in the range of above three (two, one) orders of magnitude are highlighted in red (orange, blue), while the errors less than one order of magnitude are not highlighted.**

Since  $\deg_E(Y|X) = \deg_E(Z|Y) = \deg_E(X|Z)$ , the  $\ell_2$ -bound formula is simply  $\|\deg_E(Y|X)\|_2$ .

For  $Q_4$ , the  $\ell_2$ -bound formula is:

$$(\|\deg_E(Y|X)\|_2 \cdot \|\deg_E(Z|X)\|_2)^{1/2} = \|\deg_E(Y|X)\|_2$$

following the same argument as for  $Q_3$ .

All other estimators, except for POSTGRES on all datasets and DBX on soc-pokec, yield the same estimate for both queries. The errors reported in Table 1 for the two queries differ, however, since the query outputs have different sizes for all datasets except ca-HepTh. In particular, on ca-GrQc and facebook the output of  $Q_3$  is empty whereas the output of  $Q_4$  is not empty. For all other datasets, the output of  $Q_3$  is smaller than the output of  $Q_4$ , and all estimators except DBX have larger errors for  $Q_3$ .

The error of the  $\ell_2$ -bound is under one order of magnitude for 5/7 datasets for  $Q_4$  and 2/7 datasets for  $Q_3$ , and it is under two orders of magnitude for 7/7 datasets for  $Q_4$  and 5/7 datasets for  $Q_3$ . Notably, the error is very large for  $Q_3$  on ca-GrQc and facebook, since the query output is empty. The other estimators also perform poorly in these two cases. Whenever the  $\ell_2$ -bound is not the best, it is within a small factor from the best, unless the query output is empty.

The  $\ell_2$ -bound errors can be up to 2 orders of magnitude lower than the errors of the  $\ell_1$ -bound and the  $\ell_{\{1, \infty\}}$ -bound. In the majority of the cases, the addition of the  $\ell_\infty$ -norm does not improve the estimation obtained using the  $\ell_1$ -norm.

For these triangle queries, the best  $\ell_p$ -bound is obtained using the  $\ell_2$ -norm, regardless how many further norms are available. If we were to use the  $\ell_3$ -norm instead, then the bound would be 1.3 to 4.7 times worse, yet still much better than the  $\ell_1$ -bound and the  $\ell_{\{1, \infty\}}$ -bound.

The traditional estimators DUCKDB and POSTGRES consistently overestimate the cardinality of the triangle queries, whereas DBX underestimates in all but the two cases of empty query output.

More extensive experiments with a variety of query workloads, including acyclic and cyclic queries with equality and range predicates and group-by clauses (free variables), are reported in subsequent work [13].

## 7. CONCLUSIONS

In this paper, we overviewed a new framework for upper bounds on the output size of conjunctive queries [12]. These bounds use  $\ell_p$ -norms of degree sequences, are based on information inequalities, and can be computed by optimizing a linear program whose size is exponential in the number of variables in the query. They are asymptotically tight in the case when each degree sequence is on one column. The bounds represent non-trivial generalizations of a series of prior results on cardinality upper bounds [5, 6, 7, 8], in particular for acyclic queries for which the previous AGM and PANDA bounds degenerate to trivial observations. We also reported on preliminary experiments with a workload of full conjunctive queries on a real dataset. The experiments highlight the promise of the  $\ell_p$ -bounds when compared with previously introduced upper bounds, but also traditional estimators and a recent data-driven learned estimator. More extensive experiments with an extension of this framework to arbitrary conjunctive queries with equality and range predicates are reported in follow-up work [13].

The cardinality estimator is complemented by a query evaluation algorithm whose runtime matches the size bound.

The key limitation of the  $\ell_p$ -bounds, which is shared with existing pessimistic cardinality estimators including the AGM and PANDA bounds and SAFEBOUND, is that they can overestimate significantly in case the input data is very large yet the query output is very small or even empty. Such situations happen when the input relations are miscalibrated, which is common when there are selective predicates on some of the input relations. Addressing this challenge is an important direction for future work.

## 8. REFERENCES

- [1] Viktor Leis, Andrey Gubichev, Atanas Mirchev, Peter A. Boncz, Alfons Kemper, and Thomas Neumann. How good are query optimizers, really? *Proc. VLDB Endow.*, 9(3):204–215, 2015.
- [2] Viktor Leis, Bernhard Radke, Andrey Gubichev, Atanas Mirchev, Peter A. Boncz, Alfons Kemper, and Thomas Neumann. Query optimization through the looking glass, and what we found running the join order benchmark. *VLDB J.*, 27(5):643–668, 2018.
- [3] Yuxing Han, Ziniu Wu, Peizhi Wu, Rong Zhu, Jingyi Yang, Liang Wei Tan, Kai Zeng, Gao Cong, Yanzhao Qin, Andreas Pfadler, Zhengping Qian, Jingren Zhou, Jiangneng Li, and Bin Cui. Cardinality estimation in DBMS: A comprehensive benchmark evaluation. *Proc. VLDB Endow.*, 15(4):752–765, 2021.
- [4] Kyoungmin Kim, Jisung Jung, In Seo, Wook-Shin Han, Kangwoo Choi, and Jaehyok Chong. Learned cardinality estimation: An in-depth study. In *SIGMOD*, pages 1214–1227, 2022.
- [5] Albert Atserias, Martin Grohe, and Dániel Marx. Size bounds and query plans for relational joins. *SIAM J. Comput.*, 42(4):1737–1767, 2013.
- [6] Georg Gottlob, Stephanie Tien Lee, Gregory Valiant, and Paul Valiant. Size and treewidth bounds for conjunctive queries. *J. ACM*, 59(3):16:1–16:35, 2012.
- [7] Mahmoud Abo Khamis, Hung Q. Ngo, and Dan Suciu. Computing join queries with functional dependencies. In *PODS*, pages 327–342, 2016.
- [8] Mahmoud Abo Khamis, Hung Q. Ngo, and Dan Suciu. What do shannon-type inequalities, submodular width, and disjunctive datalog have to do with one another? In *PODS*, pages 429–444, 2017.
- [9] Walter Cai, Magdalena Balazinska, and Dan Suciu. Pessimistic cardinality estimation: Tighter upper bounds for intermediate join cardinalities. In *SIGMOD*, pages 18–35, 2019.
- [10] Axel Hertzschuch, Claudio Hartmann, Dirk Habich, and Wolfgang Lehner. Simplicity done right for join ordering. In *CIDR*, 2021.
- [11] Jeremy Chen, Yuqing Huang, Mushi Wang, Semih Salihoglu, and Kenneth Salem. Accurate summary-based cardinality estimation through the lens of cardinality estimation graphs. *Proc. VLDB Endow.*, 15(8):1533–1545, 2022.
- [12] Mahmoud Abo Khamis, Vasileios Nakos, Dan Olteanu, and Dan Suciu. Join size bounds using  $l_p$ -norms on degree sequences. *Proc. ACM Manag. Data*, 2(2):96, 2024.
- [13] Haozhe Zhang, Christoph Mayer, Mahmoud Abo Khamis, Dan Olteanu, and Dan Suciu. LpBound: Pessimistic cardinality estimation using  $l_p$ -norms of degree sequences. To appear in *Proc. ACM Manag. Data (SIGMOD)*, 2025.
- [14] S. Im, B. Moseley, H. Ngo, and K. Pruhs. Efficient algorithms for cardinality estimation and conjunctive query evaluation with simple degree constraints, 2025. To appear in *Proc. ACM Manag. Data (PODS)*.
- [15] Mahmoud Abo Khamis, Kyle Deeds, Dan Olteanu, and Dan Suciu. Pessimistic cardinality estimation. *SIGMOD Rec.*, 53(4):1–17, 2024.
- [16] Bolin Ding, Rong Zhu, and Jingren Zhou. Learned query optimizers. *Found. Trends Databases*, 13(4):250–310, 2024.
- [17] Bailu Ding, Vivek R. Narasayya, and Surajit Chaudhuri. Extensible query optimizers in practice. *Found. Trends Databases*, 14(3-4):186–402, 2024.
- [18] Raghu Ramakrishnan and Johannes Gehrke. *Database management systems (3. ed.)*. McGraw-Hill, 2003.
- [19] Mahmoud Abo Khamis, Hung Q. Ngo, and Dan Suciu. PANDA: Query Evaluation in Submodular Width. *arXiv e-prints*, page arXiv:2402.02001, February 2024.
- [20] Kyle B. Deeds, Dan Suciu, and Magdalena Balazinska. Safebound: A practical system for generating cardinality bounds. *Proc. ACM Manag. Data*, 1(1):53:1–53:26, 2023.
- [21] Ziniu Wu, Parimarjan Negi, Mohammad Alizadeh, Tim Kraska, and Samuel Madden. Factorjoin: A new cardinality estimation framework for join queries. *Proc. ACM Manag. Data*, 1(1):41:1–41:27, 2023.
- [22] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.