

Rapid Approximate Aggregation with Distribution-Sensitive Interval Guarantees

Stephen Macke^{*†}, Maryam Aliakbarpour[‡], Ilias Diakonikolas[§], Aditya Parameswaran[†], Ronitt Rubinfeld[‡]

^{*}University of Illinois (UIUC) [†]UC Berkeley [‡]MIT [§]University of Wisconsin, Madison
{smacke, adityagp}@berkeley.edu {maryama, ronitt}@csail.mit.edu ilias.diakonikolas@gmail.com

Abstract—Aggregating data is fundamental to data analytics, data exploration, and OLAP. Approximate query processing (AQP) techniques are often used to accelerate computation of aggregates using samples, for which confidence intervals (CIs) are widely used to quantify the associated error. CIs used in practice fall into two categories: techniques that are *tight but not correct*, i.e., they yield tight intervals but only offer asymptotic guarantees, making them unreliable, or techniques that are *correct but not tight*, i.e., they offer rigorous guarantees, but are overly conservative, leading to confidence intervals that are too loose to be useful. In this paper, we develop a CI technique that is both correct and tighter than traditional approaches. Starting from conservative CIs, we identify two issues they often face: *pessimistic mass allocation (PMA)* and *phantom outlier sensitivity (PHOS)*. By developing a novel *range-trimming* technique for eliminating PHOS and pairing it with known CI techniques without PMA, we develop a technique for computing CIs with strong guarantees that requires fewer samples for the same width. We implement our techniques underneath a sampling-optimized in-memory column store and show how they accelerate queries involving aggregates on real datasets with typical speedups on the order of 10× over both traditional AQP-with-guarantees and exact methods, all while obeying accuracy constraints.

I. INTRODUCTION

Primitives for aggregation like AVG, SUM, and COUNT are key to making sense of and drawing insights from large volumes of data, powering applications in OLAP, exploratory data analysis, and visual analytics. Accelerating their computation is therefore of great importance. Approximate Query Processing (AQP) is commonly used to accelerate computation of these aggregates by estimating them on a subset or sample of the full data. Reasoning about the error of the estimates as introduced by approximation is crucial: consumers of approximate answers—ranging from human decision makers to automated processes—rely on confidence intervals (CIs) or error bounds as the foundation for understanding the quality of the approximate answer. Therefore, many AQP techniques come with CIs to allow for more confident or informed decisions made using approximate estimates.

Error bounding, or CI computation techniques take a confidence parameter $\delta \in [0, 1]$, with the semantics that the returned intervals $[g_\ell, g_r]$ fail to enclose the true aggregate g^* at most δ of the time. One can tune δ to be as small as needed at the cost of requiring more samples to achieve the same interval width ($g_r - g_\ell$). Likewise, for a given δ , taking more samples typically causes the error bounding procedure to return a narrower confidence interval. Since δ is typically small, we use the phrase “with high probability” (w.h.p.) as shorthand for “with probability greater than $(1 - \delta)$ ”. CI computation techniques need to satisfy two goals: **(i) compactness:** by minimizing the interval width $g_r - g_\ell$, and **(ii) correctness:** by ensuring that $g^* \in [g_\ell, g_r]$ with high probability. However, achieving both compactness and correctness is difficult.

Existing techniques either prefer compactness over correctness (*asymptotic* techniques) or vice versa (*conservative* techniques):

```
SELECT Origin, AVG(DepDelay) FROM flights
GROUP BY Origin HAVING AVG(DepDelay) < 0
```

Fig. 1: Origin airports with negative average delay. In this query, the AVG aggregates are consumed both by the user and by the system.

Compactness without Correctness. *Asymptotic* error bounding techniques such as bootstrap CIs [1, 2] or central limit theorem (CLT)-based CIs [3, 4] make assumptions about the distribution taken by the data given a “large enough” sample size. These procedures typically give CIs that are much tighter (and therefore more useful for drawing inferences about the query results), and have enjoyed numerous applications in database and visual analytics systems [5, 6, 7, 8, 9, 10], including Aqua [11], BlinkDB [12, 13], DBO [14], and online aggregation [15], and have furthermore seen a number of DBMS-specific extensions [2, 16].

However, these asymptotic techniques result in intervals that only enclose the true aggregate w.h.p. in the limit as the size of the sample grows to infinity; i.e., they provide no real guarantees for any given finite instance, potentially leading to failures downstream. For example, consider the query in Figure 1, which determines origin airports whose departing flights are ahead-of-schedule, on average. An AQP system could use CIs to facilitate early stopping by using them to infer on which side of the HAVING threshold the various groups appear. If such a system relies on asymptotic CIs, it is prone to serious types of error [6] whereby certain tuples may be missing, and other tuples may appear spuriously.

Correctness without Compactness. Recognizing the downsides of asymptotic approaches, recent work [17, 18, 19, 20, 21] has begun to adopt *conservative* error bounders, which leverage concentration inequalities to compute CIs. These procedures return bounds that follow *probably approximately correct* (PAC) [22] semantics: given $\delta \in [0, 1]$, the probability that the procedure returns lower and upper bounds $[g_\ell, g_r]$ around the approximate aggregate \hat{g} that fail to enclose the true aggregate g^* should be *at most* δ for *any* sample size (in contrast with asymptotic techniques, for which the probability converges to δ given a large enough sample). These techniques have been used in online aggregation [15, 23] and more recently in work on visual analytics [18, 19, 20, 21].

In general, conservative methods such as those based on Hoeffding’s inequality [24] or on the Hoeffding-Serfling inequality [25] rely on a-priori knowledge of *range bounds* a and b between which the data fall (typically inferred during data loading). Although they achieve the correctness goal of error bounders, when used for AVG, the CI width for Hoeffding-based error bounders scales with the range size ($b - a$), creating at least two major issues in the context of a relational database, illustrated in Figure 2. **(i) First, the presence of a very few outliers can significantly widen the range $[a, b]$** (and therefore the CI width), even though most of the data may lie in a much smaller range. **(ii)**

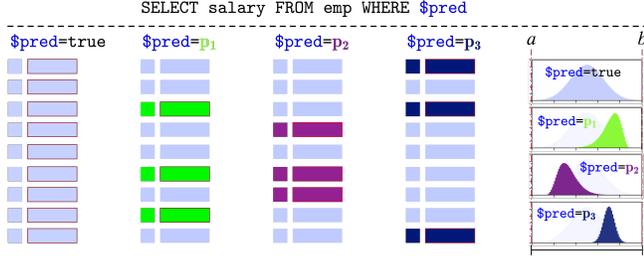


Fig. 2: Few points may lie near the range bounds a and b , and with filters applied, the true range could be significantly smaller than $(b-a)$.

Second, predicates and groupings may be applied during data exploration, so that the filtered data lies in a smaller range than $[a, b]$.

Key Research Challenges and Contributions. We aim to *preserve correctness* (or safety) of conservative error bounds for AVG, SUM, and COUNT aggregates *while also providing compactness* (for speed). We encounter a number of challenges toward this end:

1. *Identifying conservative error bounder pathologies.* To improve the viability of approaches with strict correctness guarantees, we must first determine the circumstances under which conservative error bounders are *too* conservative.

Our contribution: We identify two issues in range-based concentration inequalities that cause unnecessary looseness when used to compute conservative error bounds for AVG. The first, *pessimistic mass allocation* (PMA), refers to the unnecessary placement of unseen probability mass at endpoints a and b of the range enclosing the data. The second, *phantom outlier sensitivity* (PHOS), occurs when computation of the lower confidence bound g_ℓ depends on the upper range bound b *even without observed samples near b* , and vice versa for a dependency from a to g_r . PHOS captures the intuition that unobserved large (small) values should not loosen g_ℓ (g_r).

2. *Correcting error bounder pathologies.* After identifying correctable issues with existing conservative error bounders, we need to develop novel techniques that address these issues, while ensuring that these techniques are efficient in terms of computation and memory.

Our contribution: We develop a simple and general error bounding technique, *range trimming*, that corrects PHOS without sacrificing desirable PAC semantics. At a high level, range trimming operates by making error bounders *asymmetric*, so that g_ℓ depends only on the MAX value seen (and not on b), and g_r depends only on the MIN value seen, yielding tighter intervals when $(\text{MAX}-\text{MIN})$ is smaller than $(b-a)$. Range trimming can be used with any existing conservative range-based error bounder (i.e., an error bounder whose only assumption is that data falls in $[a, b]$). We show how range trimming can be used to develop an error bounder for AVG (and by extension SUM) with neither PHOS nor PMA by using it alongside a bounder based on Bernstein’s inequality.

3. *Minimizing sampling overhead.* In order to enjoy the benefits of early termination for queries with multiple aggregates, we need to ensure that termination is not bottlenecked on any single aggregate, allowing query processing to adaptively sample from the most informative locations on physical storage while simultaneously minimizing overhead.

Our contribution: We show how to couple our approach with a sampling-optimized column store that takes without-replacement samples in a locality-aware manner, and that leverages bitmap indexes to prioritize samples that enable earlier termination for GROUP BYs.

Impact. We develop error bounding techniques that more effectively leverage distributional information of the underlying data, and

Symbols / Terms	Descriptions
\mathcal{D}, N, S, m	Dataset, num. points in dataset (i.e. $ \mathcal{D} $), sample, num. points in sample (i.e. $ S $)
$g^*, \hat{g}, g_\ell, g_r$	True aggregate, estimate, error bounds
$a, b, \sigma^2, \hat{\sigma}^2, \delta, \epsilon$	Range bounds, variance, empirical variance, error probability upper bound, error
Lbound, Rbound	Confidence lower (resp. upper) bounding routines parameterized on a, b, N , and other sample state.
SSI, PMA, PHOS	Sample-size-independent, pessimistic mass allocation, phantom outlier sensitivity

TABLE I: Glossary of terms / notation.

that therefore often lead to tighter error bounds as compared with those yielded by typical conservative error bounders. When used in conjunction with a sampling-optimized column store for in-memory analytics, we demonstrate typical speedups on the order of $10\times$ over both exact techniques and traditional conservative approximate techniques, all without sacrificing strong correctness guarantees.

Extensibility. While our presentation focuses on confidence intervals for queries over a single table with simple AVG aggregates, we note that our techniques are more general and can be used to facilitate SUM and COUNT aggregates, queries over views formed from joins in a snowflake schema, and queries with general UDFs — we discuss these extensions in Section IV-A and in our extended technical report [26].

Outline. The rest of this paper is organized as follows. Section II discusses existing conservative error bounders and their prior usage in the DBMS literature, and develops a conceptual framework for identifying issues with these error bounders. In Section III we develop the theory behind our RangeTrim technique, and show how to fix issues with previous error bounders in Section II. Section IV addresses systems issues that appear when sampling without replacement and develops FastFrame, our sampling-optimized column store, and Section V empirically evaluates our techniques in the context of this system. We survey additional related work in Section VI.

II. DBMS ERROR BOUND INTEGRATION

In this section, we first describe applications of confidence intervals for facilitating query processing in a database system (§II-A). Next, we survey methods for computing error bounds with guarantees applicable to DBMS aggregates (§II-B) identify their shortcomings (§II-C) and conclude with a formal problem statement (§II-D).

A. DBMS CI Applications

Consider the query in Figure 1. In this query, AVG aggregates are both *displayed* as output in the query results, and are also used to *filter* the set of tuples in the output. This reflects two major applications of confidence intervals in a DBMS setting: CIs that are *explicitly used* downstream, i.e., by an analyst, or CIs that are *implicitly used* by automated processes.

Explicit Use of Downstream CIs. When approximating aggregates in a DBMS, CIs can be included in the output displayed to users. For example, the AVG aggregates belonging to the groups output by the query in Figure 1 are augmented with CIs and included in the output. Such CIs help users *reason about uncertainty in approximate answers* during analysis [9, 15].

Implicit Use of Downstream CIs. Confidence intervals have been applied towards various downstream applications, for example, to enable early stopping. Example applications include high-level accuracy contracts [7, 16] (i.e., guaranteeing query results are within ϵ of the correct), ranking query results [19, 21], and bounding relative error [18]. In all cases, the user need not ever observe the interval:

the goal is to provide *early stopping* while ensuring *correct results*. We consider these applications later in our experiments in [Section V](#).

Goal. In this paper, we are primarily concerned with enabling CI *compactness* (to reduce query latency) without sacrificing CI *correctness* (thereby ensuring safety), for both explicit and implicit applications of CIs. *The major goal is therefore to develop CI techniques that are as tight as possible, while always enclosing the quantity in question.* Throughout this section and [Section III](#), we will focus our discussion on CIs for AVG aggregates; we will cover SUM and COUNT aggregates in [Section IV](#).

B. Computing CIs in a DBMS

We now describe methods for computing error bounds with accuracy guarantees in a database system, along with any assumptions required. Relevant notation is summarized in [Table I](#). We begin by defining error bounders, bounds, and confidence intervals.

Definition 1 [(1- δ) error bounders and bounds]. A procedure P that returns error bounds $[g_\ell, g_r]$ for some aggregate g^* given a sample is a (1- δ) error bouncer if, across all possible samples, $\mathbb{P}(g^* \notin [g_\ell, g_r]) < \delta$. $[g_\ell, g_r]$ is called the (1- δ) confidence interval for g^* , and g_ℓ and g_r are collectively referred to as (1- δ) error or confidence bounds.

In contrast with asymptotic error bounders that only satisfy $\mathbb{P}(g^* \notin [g_\ell, g_r]) \approx \delta$ for large-enough sample sizes, the (1- δ) error bounders from [Definition 1](#) always satisfy $\mathbb{P}(g^* \notin [g_\ell, g_r]) < \delta$ for any sample size, so we call them *sample-size-independent (SSI)*.

1) *Assumptions Applicable to Data in a DBMS:* In the case of AVG aggregates, all error bounding procedures require some prior knowledge about the data over which they operate – otherwise, outliers can have arbitrarily strong effects on the aggregate in question. Weaker assumptions are more general, but typically yield more conservative bounds.

In this paper, we make two assumptions about the data \mathcal{D} over which queries operate: first, that every datapoint $x \in \mathcal{D}$ lies in some interval $[a, b]$; second, that datapoints can be effectively sampled *without replacement* from \mathcal{D} . We now discuss these assumptions in the context of prior work and show that they can be implemented effectively within real systems.

Known Range Bounds. As in prior work [[15](#)], we assume that the database catalog maintains *range bounds* a and b for the MIN and MAX of each continuous column, inferred, for example, during data loading. (Note that we do not require $[a, b] = [\text{MIN}, \text{MAX}]$, but only that $[a, b] \supseteq [\text{MIN}, \text{MAX}]$.) We refer to bounders that assume knowledge of a and b as *range-based error bounders* throughout this paper.

Sampling Without Replacement. Estimates for AVG aggregates generally converge faster for samples taken without replacement than samples taken with replacement [[25](#), [27](#)]. In the context of a DBMS, sampling with replacement has traditionally been considered easier than sampling without replacement, since the system does not need to “remember” the samples already taken [[19](#), [28](#)]. Sampling as traditionally implemented, however, also has poor locality properties, as nearly every read operation results in a cache miss. Another approach taken in prior work [[21](#), [29](#), [30](#), [31](#)] is to materialize samples ahead-of-time by performing a single up-front shuffle of the entire relation, so that sampling without replacement can be implemented via a scan of the data *regardless* of any applied filters or other transformations. Since this approach is valid for multiple queries executed during ad-hoc, exploratory workloads (in contrast with approaches that use workload assumptions to pre-materialize stratified samples [[12](#), [32](#)]), we design our system architecture around this approach, described in more detail in [Section IV](#).

Error Bouncer	PMA	PHOS	Sampling	Memory
Hoeffding(-Serfling)	✓	✓	R* (NR)	$\mathcal{O}(1)$
Berstein(-Serfling)	✓	✓	R* (NR)	$\mathcal{O}(1)$
Anderson/DKW [26]	✓		R, NR	$\mathcal{O}(m)$

TABLE II: Bouncer properties. R = sampling with replacement, NR = without. A * indicates that the non-Serfling variant also holds for NR sampling.

2) *Error Bounds for Finite and Bounded Data:* In this section, we review some techniques for computing confidence intervals that leverage only the assumptions discussed previously: that samples are taken without-replacement from data bounded in some a priori-known range $[a, b]$. Further details about these bounders, such as implementation pseudocode and full restatements of relevant theorems, are available in our extended technical report [[26](#)].

Hoeffding-Serfling-based Bouncer. An error bouncer based on the Hoeffding-Serfling inequality [[25](#)] computes CIs whose widths depend only on the range $(b-a)$ and the number of samples m , and that have size $\mathcal{O}((b-a)/\sqrt{m})$. While asymptotically optimal for worst-case data distributed with half of the points at a and the other half at b , it is needlessly wide in practice, when few points occur near a or b .

Empirical Bernstein-Serfling-based Bouncer. The *empirical Bernstein-Serfling inequality* is another concentration inequality for sampling without replacement [[27](#)]. A corresponding error bouncer yields (1- δ) CIs whose widths have size $\mathcal{O}(\hat{\sigma}/\sqrt{m} + (b-a)/m)$, where $\hat{\sigma}^2 = \frac{1}{m} \sum_{t=1}^m (X_t - \bar{X})^2$ is the maximum likelihood estimator for $\text{VAR}(\mathcal{D})$. Although $\hat{\sigma}$ is a random quantity, it concentrates near σ ; thus, we see that error bounds derived from the Bernstein-Serfling inequality can be significantly tighter than those derived from Hoeffding-Serfling (which has widths of size $\mathcal{O}((b-a)/\sqrt{m})$) when σ is small compared to $(b-a)$. In fact, we will see shortly ([§II-C](#)) that Bernstein-based bounders do not suffer from one of the problems that causes Hoeffding-based bounders to be *overly-conservative*.

Applications in Prior DB Literature. To our knowledge, Hoeffding and Hoeffding-Serfling-based bounders are the only SSI bounders that have seen extensive use in the DB literature for computing error bounds for AVG [[15](#), [18](#), [19](#), [23](#)]. We are aware of one incorrect application of the empirical Bernstein-Serfling inequality [[33](#)] (incorrect because the procedure given in [[33](#)] continuously recomputes confidence (1- δ) intervals as more samples are taken, so that the overall procedure is no longer guaranteed to fail with probability at most δ). Overall it is somewhat surprising that error bounders derived from the empirical Bernstein-Serfling inequality [[27](#)] have not seen more widespread usage, as they are nearly as simple to compute as those derived from the Hoeffding-Serfling inequality and typically yield error bounds that are much tighter.

C. Error Bouncer Pathologies

We identify two problems that cause SSI error bounders to be *too* conservative. These pathologies, which we refer to as *pessimistic mass allocation (PMA)* and *phantom outlier sensitivity (PHOS)*, are based on simple intuitions about how error bounders should behave: namely, they should return tighter bounds when observing samples with fewer extreme values, and error lower bounds (respectively error upper bounds) should only be looser due to potential large values (resp. small values) if such values are actually observed.

1) *Pessimistic Mass Allocation:* PMA captures the intuition that error bounders should be sensitive to the observed sample values:

Definition 2 [PMA]. An error bounding procedure P exhibits pessimistic mass allocation (PMA) if there exists a dataset \mathcal{D} bounded in $[a, b]$, a value a' with $a < a' < b$, and a set $S \subseteq \mathcal{D}$ with values in

$[a, a']$ such that, for $S' = \{\max(x, a') : x \in S\}$, P returns a confidence interval of the same width for both S and S' . P likewise exhibits PMA if there exists some b' with $a < b' < b$ and an S with values in $(b', b]$ such that, for $S' = \{\min(x, b') : x \in S\}$, P returns a confidence interval of the same width for both S and S' .

That is, for an error bouncer P with PMA, we can replace the smallest (largest) elements in a sample with something larger (resp. smaller) without shrinking the width of P 's returned confidence interval. For example, for data known to lie in $[0, 1]$, P might yield an interval of the same width for both a sample split evenly between 0 and 1 as well as a sample split evenly between 0.25 and 0.75, even though the latter sample should clearly give rise to a tighter interval. Intuitively, P is overly-pessimistic about how mass in the distribution from which it is sampling is allocated, despite contrary evidence observed in the sample.

2) *Phantom Outlier Sensitivity*: PHOS captures the intuition that unobserved extreme values should not affect both the lower and the upper error bounds computed by some error bouncer P :

Definition 3 [PHOS]. An error bounding procedure P exhibits phantom outlier sensitivity (PHOS) if, for data falling in $[a, b]$, P 's returned confidence lower bound g_ℓ depends on the value of b , and similarly if the g_r returned by P depends on a .

To understand PHOS intuitively, consider the case of computing a confidence lower bound. Given a sample S , the worse P “believes” S could be shifted (on average) toward larger values as compared to \mathcal{D} , the smaller of a confidence lower bound it should return. In what ways could S be shifted toward higher values? One possibility is if small elements are underrepresented in S . The other possibility, and the one we are interested in, is if large elements are overrepresented in S . For this reason, a confidence lower bound should only be affected by datapoints near the upper range bound b if it actually observes them, and the appearance of b in the computation of a confidence lower bound is a potential source of unnecessary conservativeness.

3) *Examples of PMA and PHOS in Error Bounders*: In this section, we give examples of PMA and PHOS in the context of previously-discussed error bounders. Table II summarizes pathologies exhibited by various SSI error bounders.

Hoeffding-based. Hoeffding-based error bounders suffer from both PMA and PHOS. They have PMA since their returned CIs have widths depending only on the range of the data, $(b - a)$, and the number of samples. As such, replacing values in the sample with larger or smaller values does not affect the width of the returned error bounds. Such bounders also have PHOS since they have symmetric error, with both ends of the confidence interval depending on both a and b .

Bernstein-based. Bernstein-based error bounders do *not* suffer from PMA. To see this, notice that increasing the smallest values in some sample will also reduce the sample variance, affecting the width of the returned confidence interval, and likewise for decreasing the largest values in the sample. These bounders do, however, suffer from PHOS. Like Hoeffding-based bounders, they return confidence intervals with symmetric error, so that each end of the confidence interval is affected by both ends of the data range a and b .

In our extended technical report, we also consider one other error bouncer based on the DKW inequality, which interestingly suffers from PMA but not PHOS; please see [26] for details. Because it requires storing the sample in memory, we do not consider it subsequently.

D. Problem Statement

We are now ready to give a formal problem statement.

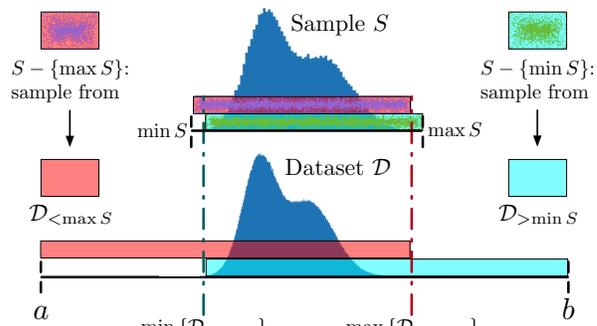


Fig. 3: Range trimming eliminates PHOS for range-based error bounders.

Problem 1. Design an SSI error bouncer that, given a without-replacement sample from any \mathcal{D} with elements from $[a, b] \subseteq \mathbb{R}$, suffers from neither PMA nor PHOS when computing $(1 - \delta)$ error bounds for $\text{AVG}(\mathcal{D})$, for any $0 < \delta < 1$.

Our solution to Problem 1 is given in Section III and relies on a technique we call *range trimming* in order to systematically eliminate PHOS from any range-based error bouncer.

III. FIXING BOUNDER PATHOLOGIES

From our discussion in Section II-C, we see that there do exist error bounders without one of either PMA or PHOS, but not without both. We first argue that error bounders without PHOS must be *asymmetric*; that is, they cannot compute bounds of the form $\hat{g} \pm \varepsilon$, where the same ε is both added and subtracted to the sample average \hat{g} in order to compute bounds. Next, we describe how to use a process we call *range trimming* to convert any symmetric, ranged-based error bouncer to an asymmetric one without PHOS.

A. Decoupling Lower and Upper Bounds

Excepting an error bouncer based on DKW, all of the error bounders surveyed suffer from PHOS. This is because all the other error bounders are based on concentration inequalities with *symmetric error* — that is, they return confidence intervals $[g_\ell, g_r]$ of the form $[\hat{g} - \varepsilon, \hat{g} + \varepsilon]$. At a high level, it is precisely this symmetry that causes PHOS. Although a confidence lower bound should not have any dependency on b , it is intuitively unavoidable that it has some dependency on a . Reiterating, an estimate \hat{g} could be an overestimate because of (i) not enough observed values near a , or (ii) too many observed values near b . A similar statement holds regarding confidence upper bounds, with the roles of a and b reversed.

We hypothesize that it is impossible for any confidence lower bound (resp. upper bound) to completely eliminate the dependency on a (resp. b), since it is always possible that the confidence bounding procedure got “unlucky” and operated on a sample in which values near a (resp. b) were underrepresented. Taking this hypothesis as given, this means that any symmetric confidence bounding procedure that returns bounds of the form $[\hat{g} - \varepsilon, \hat{g} + \varepsilon]$ will have ε dependent on both a and b — that is, any symmetric confidence bounding procedure will have PHOS. As such, the first step to eliminating PHOS from range-based confidence bounders is to accept asymmetric error as a hard requirement: that is, we must consider confidence bounding procedures that return bounds of the form $[\hat{g} - \varepsilon_\ell, \hat{g} + \varepsilon_r]$ for which ε_ℓ and ε_r are not necessarily equal.

B. Range Trimming

Our approach to deriving an error bouncer with neither PMA nor PHOS is to start with a symmetric bouncer without PMA (such as

Algorithm 1: The RangeTrim meta-algorithm

Input : Dataset \mathcal{D} of N values in $[a, b]$, error prob. δ , sample size m
Output : Error bounds that fail to enclose $\text{AVG}(\mathcal{D})$ with probability $< \delta$

```

1  $S_\ell \leftarrow \text{init\_state}()$ ;
2  $S_r \leftarrow \text{init\_state}()$ ;
3  $a' \leftarrow \text{sample\_without\_replacement}(\mathcal{D})$ ;
4  $b' \leftarrow a'$ ;
5 for  $i = 1$  to  $m-1$  do
6    $v \leftarrow \text{sample\_without\_replacement}(\mathcal{D})$ ;
7    $S_\ell \leftarrow \text{update\_state}(S_\ell, \min(v, b'))$ ;
8    $S_r \leftarrow \text{update\_state}(S_r, \max(v, a'))$ ;
9    $a' \leftarrow \min(a', v)$ ;
10   $b' \leftarrow \max(b', v)$ ;
11 end
12 return  $[\text{Lbound}(S_\ell, a, b', N-1, \frac{\delta}{2}), \text{Rbound}(S_r, a', b, N-1, \frac{\delta}{2})]$ ;

```

a Bernstein-based bounder) and “asymmetrize” it so that Lbound becomes independent of b , and Rbound becomes independent of a , thereby eliminating PHOS. The result, given in [Algorithm 1](#), composes any range-based error bounder that exposes the following interface:

- 1 $\text{init_state}()$: Initializes state needed for error bounds.
- 2 $\text{update_state}(S, v)$: Given the current state S and a newly-seen value v , compute new state S' .
- 3 $\text{Lbound}(S, a, b, N, \delta)$: Return a confidence lower bound for a sample whose relevant statistics are captured in state S , assuming the sample came from a finite dataset \mathcal{D} of N values in $[a, b]$. The probability that the sample leads to this function returning a value greater than $\text{AVG}(\mathcal{D})$ is $< \delta$.
- 4 $\text{Rbound}(S, a, b, N, \delta)$: Symmetric to Lbound for the confidence upper bound. Can typically be implemented in terms of Lbound after a suitable transformation of S .

The state S captures information such as the count of tuples examined and the current running average, as well as anything else required by Lbound and Rbound . The state initialization and update logic is analogous to state maintenance logic for aggregate functions as implemented in existing commercial database systems [34].

Besides the memory required to maintain state for the left and right error bounders, S_ℓ and S_r , [Algorithm 1](#) requires $\mathcal{O}(1)$ extra memory to maintain the MIN and MAX element seen so far (which replace a and b when computing Rbound and Lbound , respectively).

When \mathcal{D} contains unique elements, [Algorithm 1](#) conceptually performs the following steps:

- 1) Sample S without replacement from \mathcal{D} .
- 2) Use Lbound to compute a $1 - \frac{\delta}{2}$ lower confidence bound for $\text{AVG}(\mathcal{D}_{<\max S})$, with $S - \{\max S\}$ as the sample, and with a and $\max S$ in place of the normal range bounds a and b , respectively.
- 3) Use Rbound to compute a $1 - \frac{\delta}{2}$ upper confidence bound for $\text{AVG}(\mathcal{D}_{>\min S})$, with $S - \{\min S\}$ as the sample, and with $\min S$ substituted for the range bound lower bound a .

Note that we use $\mathcal{D}_{<x}$ and $\mathcal{D}_{>x}$ as shorthand for $\mathcal{D} \cap (-\infty, x)$ and $\mathcal{D} \cap (x, \infty)$, respectively. The primary difference between these high-level steps and the pseudocode presented in [Algorithm 1](#) is that [Algorithm 1](#) maintains $\min S$ and $\max S$ in an online, streaming fashion (so that the sample S does not need to be stored in memory), and that the confidence interval returned by [Algorithm 1](#) is valid even when \mathcal{D} contains duplicates (although the returned confidence bounds will bound the AVG of sets that differ slightly from $\mathcal{D}_{<\max S}$ and $\mathcal{D}_{>\min S}$). That said, we restrict our discussion and analysis to the case where \mathcal{D} contains unique elements, for simplicity.

Correctness of [Algorithm 1](#) crucially depends on the fact that, conditioned on the value of $\max S$ (and for any such value), the remaining

elements in S (namely $S - \{\max S\}$) constitute a uniform without-replacement sample from $\mathcal{D}_{<\max S}$, with a symmetric statement for $\min S$ and $S - \{\min S\}$. At a high level, this means that a confidence lower bound computed over $S - \{\max S\}$ is a valid confidence lower bound for $\text{AVG}(\mathcal{D}_{<\max S})$, and since $\text{AVG}(\mathcal{D}_{<\max S}) \leq \text{AVG}(\mathcal{D})$, it is also a valid confidence lower bound for $\text{AVG}(\mathcal{D})$, with symmetric statements holding for the confidence upper bound, $S - \{\min S\}$, and $\mathcal{D}_{>\min S}$. These core ideas are illustrated in [Figure 3](#).

Tradeoffs. When either of the range bounds a or b are actually observed in a sample, [Algorithm 1](#) will compute CIs that are looser than necessary. For example, suppose that, for $[a, b] = [0, 1]$, a sample with a single 0 is observed. When computing the upper confidence bound, [Algorithm 1](#) sets $a' = 0$, which is the same as the original a ; however, it also throws away this 0 point which would have pulled the upper confidence bound lower. The upper confidence bound actually computed will end up being larger than necessary, corresponding to an unnecessarily loose interval. Fortunately, we have found that, in practice, losing this single sample does not significantly increase the number of samples needed to achieve some desired CI width.

C. Proof of Correctness

In this section, we prove correctness of [Algorithm 1](#) (that is, that it returns intervals that fail to enclose $\text{AVG}(\mathcal{D})$ with probability less than δ). For the sake of simplicity, our analysis assumes that \mathcal{D} contains no duplicate values, although it is possible to remove this assumption. To begin, we first prove a crucial lemma about the sampling distribution of $S - \{\max S\}$, given that S was sampled uniformly without-replacement from \mathcal{D} .

Lemma 1. *Given a dataset \mathcal{D} of N unique real values in $[a, b]$ and a uniform without-replacement sample S of m values from \mathcal{D} , if we denote $b' = \max S$, the set $S - \{b'\}$ takes the distribution of a uniform without-replacement sample from $\mathcal{D}_{<b'} = \mathcal{D} \cap [a, b')$, for any applicable value of $b' \in \mathcal{D}$.*

Proof. Because S is drawn uniformly without-replacement from \mathcal{D} , any particular instance satisfies

$$\mathbb{P}_{\mathcal{D}}[S = s] = \frac{|\mathcal{D}|}{|s|} \mathbb{I}\{s \subseteq \mathcal{D}\} = \binom{N}{m}^{-1} \mathbb{I}\{s \subseteq \mathcal{D}\}$$

where we use the notation $\mathbb{P}_{\mathcal{D}}[S = s]$ to denote the probability that s was drawn uniformly without-replacement from \mathcal{D} , and $\mathbb{I}\{\cdot\}$ denotes the indicator function. We need to show that, for any $b' \in \mathcal{D}$,

$$\mathbb{P}_{\mathcal{D}}[S = s | \max S = b'] = \mathbb{P}_{\mathcal{D}_{<b'}}[S = s - \{b'\}] \mathbb{I}\{\max(s) = b'\}$$

First, letting s' be any set such that $|s'| = m-1$, we have that

$$\mathbb{P}_{\mathcal{D}_{<b'}}[S = s'] = \binom{|\mathcal{D}_{<b'}|}{m-1}^{-1} \mathbb{I}\{s' \subseteq \mathcal{D}_{<b'}\}$$

Next, consider $\mathbb{P}_{\mathcal{D}}[S = s | \max S = b']$. Bayes' rule gives that

$$\mathbb{P}_{\mathcal{D}}[S = s | \max S = b'] = \frac{\mathbb{P}_{\mathcal{D}}[S = s \wedge \max S = b']}{\mathbb{P}_{\mathcal{D}}[\max S = b']}$$

We have $\mathbb{P}_{\mathcal{D}}[S = s \wedge \max S = b'] = \mathbb{P}_{\mathcal{D}}[S = s] \mathbb{I}\{\max(s) = b'\}$ which is a known quantity, so the key is to compute the denominator $\mathbb{P}_{\mathcal{D}}[\max S = b']$. Using the assumption that \mathcal{D} contains unique elements, we may proceed by analogy with binary strings. The rank of b' within \mathcal{D} (starting from the smallest element) is $1 + |\mathcal{D}_{<b'}|$, so we need to compute the number of binary strings of length N containing m 1's and $(N-m)$ 0's such that position $1 + |\mathcal{D}_{<b'}|$ has a 1, and the remaining $(m-1)$ 1's are all at positions less than $1 + |\mathcal{D}_{<b'}|$. This

is precisely the same as the number of binary strings of length $|\mathcal{D}_{<b'}|$ with $(m-1)$ 1's and $(|\mathcal{D}_{<b'}|-m+1)$ 0's. Putting everything together,

$$\begin{aligned} \mathbb{P}_{\mathcal{D}}[S = s | \max S = b'] &= \frac{\mathbb{P}_{\mathcal{D}}[S = s] \mathbb{I}\{\max(s) = b'\}}{\mathbb{P}_{\mathcal{D}}[\max S = b']} \\ &= \frac{\binom{N}{m}^{-1} \mathbb{I}\{s \subseteq \mathcal{D} \wedge \max(s) = b'\}}{\binom{|\mathcal{D}_{<b'}|}{m-1} / \binom{N}{m}} \\ &= \binom{|\mathcal{D}_{<b'}|}{m-1}^{-1} \mathbb{I}\{s \subseteq \mathcal{D} \wedge \max(s) = b'\} \\ &= \binom{|\mathcal{D}_{<b'}|}{m-1}^{-1} \mathbb{I}\{s - \{b'\} \subseteq \mathcal{D}_{<b'} \wedge \max(s) = b'\} \\ &= \mathbb{P}_{\mathcal{D}_{<b'}}[S = s - \{b'\}] \mathbb{I}\{\max(s) = b'\} \end{aligned}$$

which is precisely what we wanted to show. \square

Wrinkle in Lemma 1 and Fix. The proof of Lemma 1 assumes unique values; please see our extended technical report [26] for how to remove this assumption without loss of generality.

We next give a symmetric statement for $S - \{\min S\}$ and $\mathcal{D}_{>\min S}$ as the below corollary:

Corollary 1. *Given a dataset \mathcal{D} of N unique real values in $[a, b]$ and a uniform without-replacement sample S of m values from \mathcal{D} such that $\min S = a'$, the set $S - \{a'\}$ is a uniform without-replacement sample from $\mathcal{D}_{>a'} = \mathcal{D} \cap (a', b]$, for any applicable value of $a' \in \mathcal{D}$.*

We now give a theorem on the correctness of Algorithm 1.

Theorem 1. *Given SSI range-based bounds L_{bound} and R_{bound} for computing lower (resp. upper) confidence bounds and a dataset \mathcal{D} of N unique values known to all fall in the interval $[a, b] \subseteq \mathbb{R}$, Algorithm 1 returns a $(1-\delta)$ confidence interval for $\text{AVG}(\mathcal{D})$.*

Proof. Please see our extended technical report [26]. \square

IV. SYSTEM CONSIDERATIONS

In this section, we address a number of implementation issues that become pertinent when applying techniques of previous sections in a real system. Although the techniques presented in this section are auxiliary to our primary contribution and can be used with any CI approach, they are developed with SSI error bounders and strong probabilistic guarantees in mind. First, we describe how to augment the techniques of Section III, which apply for a fixed sample size taken without replacement from a finite dataset of known size, with locality-aware *scan-based* without-replacement sampling that need not know N , and we further describe how to use this layout to facilitate SUM and COUNT aggregations (§IV-A). Next, we describe an *optional stopping* routine that does not require a sample size to be specified up-front (§IV-B). Finally, we describe an *active scanning* architectural optimization that prioritizes samples that facilitate early termination (§IV-C), all without losing guarantees proved in Section III.

These system details are implemented within the context of FastFrame, which is our general relational column store for approximate report generation with guarantees. FastFrame uses the error bounders from Section III and pairs them with a practical architecture for without-replacement sampling. FastFrame uses block-based bitmaps over categorical attributes (similar to [21]) for efficient processing of queries with predicates or groups. Furthermore, for continuous attributes, FastFrame stores the minimum and maximum values in a catalog, to be used as the range bounds a and b for the desired range-based error bounder.

A. Scan-Based Sampling for DB Aggregates

We now describe how FastFrame implements without-replacement sampling in a locality-aware manner by *scanning* over pre-shuffled data, and furthermore how this approach can be used to compute CIs for COUNT and SUM. The up-front shuffling cost need only be paid once in order to facilitate many queries, although care must be taken to set the error probability δ small enough when running multiple queries to avoid losing error bounder guarantees. This approach is not new and has been used in prior work [21, 29, 30, 31, 35]. We begin by introducing *scrambles* and *aggregate views*:

Definition 4 [Scramble]. *A scramble is an ordered copy of a relational table that has been permuted randomly, allowing for scan-based without-replacement sampling.*

Note that there exist external shuffling techniques that can handle data too large to fit in memory [36].

Scanning a continuous column in a scramble is equivalent to sampling without replacement. In fact, scanning any subset of data in a continuous column in a scramble (assuming the subset is chosen without knowledge of the order of data) is also equivalent to sampling without replacement, so that scanning a scramble can be used to sample without replacement for any aggregate appearing in a query containing arbitrary filters or GROUP BY clauses. We call such subsets *aggregate views*:

Definition 5 [Aggregate View]. *An aggregate view for some aggregate A appearing in a query (possibly belonging to a group induced by a GROUP BY clause) is the set of values in a scramble that contribute toward the computation of A .*

Choosing δ to Maintain Guarantees. Note that δ must be divided by the number of aggregate views in a query (or an upper bound) to preserve error guarantees (via a union bound). Furthermore, to guarantee (via another union bound) that the probability of one or more queries in a workload giving incorrect results is at most δ , we must further divide by (an upper bound on) the number of queries in the workload before supplying it as a parameter to an error bounder. For this reason, we choose $\delta = 10^{-15}$ when we discuss our empirical study in Section V, since even for large workloads comprised of, say, 10^7 queries, the probability of one or more mistakes will still be at most 10^{-8} . Indeed, our techniques are specifically targeted toward the case wherein it is known ahead of time that a large (but possibly unknown) number of exploratory queries will be issued, so that the cost to materialize the scramble is justified. Note that, provided δ is properly decayed, guarantees hold when the same scramble is used across an entire query workload; i.e., it is not necessary to reshuffle the scramble between queries.

Computing CIs for COUNT. Ensuring that data in a scramble are permuted randomly makes it easy to compute bounds on the selectivities of aggregate views, and by extension on the COUNT of tuples in each aggregate view, using existing techniques [37, 38]. To outline the basic idea, consider that one can conceptually assign each row of a scramble a 1 if it belongs to the aggregate view of interest, and a 0 otherwise. The AVG of this “derived” view (over the whole scramble) is exactly the selectivity of the aggregate view, and we can use a Hoeffding-Serfling-based bounder to compute a CI for the selectivity (using range bounds of $a = 0$ and $b = 1$). Multiplying these bounds by the total number of rows in the scramble then yields a CI for the COUNT of rows that participate in the aggregate view.

Computing CIs for AVG with unknown COUNT. Recall that the range-based error bounders we consider in Section V all take as input the number of tuples in an aggregate view N_i . In fact, an upper bound

on N_i suffices to preserve correctness. Thus, the same method for computing CIs for COUNT aggregates can be used in conjunction with error bounders for AVG by using an upper bound on COUNT in place of the exact dataset size N_i .

Computing CIs for SUM. Now that we have established how to compute CIs for AVG and COUNT, we briefly describe how to combine these two techniques to compute CIs for SUM. Given a $(1 - \frac{\delta}{2})$ confidence interval for COUNT as $[c_\ell, c_r]$ and a $(1 - \frac{\delta}{2})$ confidence interval for SUM as $[g_\ell, g_r]$, union bounding gives $[c_\ell \cdot g_\ell, c_r \cdot g_r]$ as a $(1 - \delta)$ confidence interval for SUM.

Scramble Use and Maintenance under Updates. There exist a few ways to leverage the scramble for analytical queries even for hybrid workloads that additionally involve deletes, updates, and insertions. First, deletes can be handled by simply writing a tombstone to any deleted tuples, and (non-insertion) updates can be handled by simply modifying the updated tuple in-place. In this case, processing the tuples sequentially while ignoring tombstones will remain equivalent to sampling without replacement.

For insertions, there are two options: the first is to leave some holes in the scramble to allow for insertions in random positions, at the cost of increasing space and query time. The second option is to write all updates to a separate smaller “insertion table” that is scanned in full for every query, and then combined with the results of the scramble in a manner which preserves guarantees. This insertion table can then be periodically merged with the scramble and reshuffled. Overall, with some minor extensions inspired from existing big data systems (such as Bigtable [39]) that handle updates by keeping them separate and periodically merging them, scrambles can be readily retrofitted to handle non-read-only workloads.

B. Optional Stopping

The techniques discussed in Section III describe how to compute high-probability bounds on error given statistics computed from a particular sample of m datapoints. Fixing a sample size ahead of time is oftentimes impractical, since it is usually unknown how many samples are needed to ensure CIs that are “just tight enough” to facilitate downstream applications on the part of the user or the system. For example, one approach (e.g. taken by VerdictDB [16]) is to first compute error bounds around an approximate aggregate, and then run an exact query if these bounds are too loose.

Another approach, which we take in this paper, is to continue taking samples until a bound on the error is provably small enough. For this approach, care must be taken to avoid losing guarantees offered by range-based error bounders, since the tighter of two $(1 - \delta)$ confidence intervals for a particular aggregate is itself not necessarily a $(1 - \delta)$ confidence interval. As the technique employed is orthogonal to the primary contribution of this work, which is to eliminate PMA and PHOS from range-based error bounders, we omit discussion here and instead give a treatment in the technical report [26].

Stopping Conditions. We consider several stopping conditions used in our system implementation:

- ❶ **Sufficient Relative Accuracy** ($\max\{\frac{g_r - \hat{g}}{g_r}, \frac{\hat{g} - g_\ell}{g_\ell}\} < \epsilon$): The interval width is sufficiently small (relative to the possible correct values implied by the interval).
- ❷ **Threshold Side Determined** ($v \notin [g_\ell, g_r]$): The interval does not contain some threshold value v , indicating that the true AVG is w.h.p. either less than or greater than the threshold v .

Dataset	Size	Tuples	Columns	Replications
FLIGHTS	32 GB	606 mil.	5	5×
TAXI	36 GB	679 mil.	4	4×
POLICE	12 GB	292 mil.	3	72×

TABLE III: Dataset descriptions.

- ❸ **Top- or Bottom- K Separated:** In a query with multiple groups, the error bounds of the groups with either K smallest or largest aggregates do not intersect those of any of the remaining groups.
- ❹ **Groups Ordered Correctly:** In a query with multiple groups, the error bounds for each group intersect none of the other groups’ error bounds, indicating that the correct ordering of group aggregates has been determined [19].

Different stopping conditions apply to different queries. For example, stopping conditions ❶ and ❷ might be used for the query in Figure 1.

C. Active Scanning

For queries with GROUP BYs, different groups may require different numbers of samples to achieve stopping conditions of the types considered in Section IV-B. For simple scans that simply read blocks of the scramble in the order in which they appear, it is impossible to control the relative number of tuples for each group, leading to potential inefficiencies. For example, consider one of the queries in our experiments, F-q2, which selects airlines with average delay above some threshold. This query uses stopping condition ❷ in order to determine when to terminate, since, when this stopping condition has been achieved, it has been determined w.h.p. whether each airline has average delay above or below the threshold. Those groups (airlines) for which the average delay is near $\$thresh$ require more samples than those for which the average delay is far from $\$thresh$ in order to achieve condition ❷. If these groups are sparse within the scramble, a scan will look at much more data than necessary.

For this reason, we process queries that perform GROUP BYs with an adaptive sampling approach using *active scanning*. Active scanning uses block-based bitmap indexes to efficiently check whether a block contains tuples for any active group — such groups are marked for processing, and blocks without tuples for any active group are skipped, since they are unlikely to help achieve early termination. The notion of an active group depends on the stopping condition, but in brief, active groups are groups that should be prioritized. Please see our extended technical report [26] for a description of active groups for each applicable stopping condition.

We furthermore accelerate active scanning with a lookahead technique from prior work [21], which we briefly describe here. Instead of checking each block one by one for whether it contains tuples for any active group, active scanning *with lookahead* checks, for each active group, whether each block in a batch of 1024 blocks contains any tuples for that group. By iterating over an entire batch of 1024 blocks for a given active group, bitmaps for the group tend to be in cache more often, making the index lookup operation more efficient. We refer the reader to [21] for more details. In our experiments in Section V, we set the block size to 64 rows, so a batch of 1024 blocks contains a total of 65536 rows.

V. EMPIRICAL STUDY

In this section, we perform an extensive empirical evaluation of various error bounders and sampling strategies on real data.

A. Datasets and Queries

We evaluate various error bounding techniques on publicly available FLIGHTS, TAXI, and POLICE datasets [40, 41, 42]. For

Query	Stop When	Parameters Varied	Defaults
F-q1	(⊙) $\max\{\frac{g_r - \hat{g}}{g_r}, \frac{\hat{g} - g_\ell}{g_\ell}\} < \epsilon$	\$airport (Figure 5), ϵ (Figure 6(a))	\$airport=ORD $\epsilon=0.5$
F-q2	(⊙) $\text{Stresh} \notin [g_\ell, g_r]$	\$stresh (Figure 6(b))	\$stresh=0
F-q3	(⊙) bottom-2 separated	\$min_dep_time (Figure 7)	\$min_dep_time =10:50pm

TABLE IV: Summary of stopping conditions and template parameters used for queries provided in Figure 4. Template variable arguments shown in blue.

F-q1: avg delay for \$airport <code>SELECT AVG(DepDelay) FROM flights WHERE Origin = \$airport</code>
F-q2: airlines with avg delay above \$stresh <code>SELECT Airline FROM flights GROUP BY Airline HAVING AVG(DepDelay) > \$stresh</code>
F-q3: 2 airlines with min avg delay after \$min_dep_time <code>SELECT Airline FROM flights WHERE DepTime > \$min_dep_time GROUP BY Airline ORDER BY AVG(DepDelay) ASC LIMIT 2</code>

Fig. 4: SQL for the first three FLIGHTS queries. Template parameters are shown in blue.

FLIGHTS, we extract five attributes corresponding to the origin airport, airline, departure delay, departure time, and day of week. To ensure sufficient scale of the data, we perform 5 replications, giving a 32 GiB dataset of 606 million tuples in total. For TAXI, we extract four attributes (for hour of day, passenger count, trip distance, and fare amount) and perform 4 replications. For POLICE, we extract 3 attributes (for number of violations, officer race, and driver age) and perform 72 replications. These datasets are summarized in Table III.

Queries and Query Templates. We evaluate our techniques on a diverse set of queries that include various filters and GROUP BY clauses and exercise all the stopping conditions described in Section IV-B. The first 3 of the 13 queries in our set are presented in full in Figure 4 and are additionally parameterized (shown in blue in Figure 4) in order to reveal interesting data-dependent behavior. The stopping conditions for these queries are summarized in Table IV. Descriptions of the remaining 10 queries are presented in the technical report [26] to save space.

B. Experimental Setup

The core of our experiments consists of two ablation studies, intended to evaluate the impact of both our error bounder innovations and that of our architectural innovations. In particular, we evaluate various error bounders with and without our RangeTrim technique developed in Section III, and for the best error bounder (Bernstein+RT), we furthermore evaluate the impact of leaving out features of our active scanning sampling strategy described in Section IV.

We set $\delta = 10^{-15}$ as the default for all queries unless otherwise noted, as we expect users of with-guarantees AQP to desire results that are correct in an *effectively deterministic* manner.

Approaches. We used the following strategies to bound error when running queries in Figure 4:

Error Bounders.

- Bernstein+RT. This uses the empirical Bernstein-Serfling error bounder described in Section II-C, coupled with our RangeTrim technique described in Section III, which eliminates PHOS.
- Bernstein. Same as the previous, but without RangeTrim. Bernstein and Bernstein+RT are included to evaluate the impact of an error bounder without PMA.
- Hoeffding+RT. This uses the Hoeffding-Serfling error bounder described in Section II-C, coupled with our RangeTrim technique

described in Section III, which eliminates PHOS from Hoeffding (but does not fix PMA).

- Hoeffding. Same as the previous, but without RangeTrim.
- Exact. This strawman approach eschews approximation and runs queries exactly, to serve as a simple baseline.

We furthermore used the following strategies for sampling when running queries in Figure 4:

Sampling Strategies.

- ActivePeek. This uses the active scanning technique to prioritize groups that are preventing satisfaction of various stopping conditions, along with cache-efficient queries to bitmaps with lookahead (see Section IV-C for details).
- ActiveSync. This uses active scanning, but processes each block synchronously when deciding whether to read it, incurring high overhead since queries to bitmaps typically result in cache misses.
- Scan. This strategy does not leverage bitmaps in order to decide whether to read a block for active scanning (but may leverage bitmaps for evaluation of whether a block contains tuples that satisfy a fixed predicate, such as the one appearing in F-q1). Without any predicate, this approach simply processes all blocks in the scramble sequentially. Note that the Exact baseline described previously always uses Scan, as only approximate approaches can prune groups.

Environment. Experiments were run on single Intel Xeon E5-2630 node with 125 GiB of RAM and with 8 physical cores (16 logical) each running at 2.40 GHz, although we restrict our experiments to a single thread, noting that our techniques can be easily parallelized. The Level 1, Level 2, and Level 3 CPU cache sizes are, respectively: 512 KiB, 2048 KiB, and 20480 KiB. We ran Linux with kernel version 2.6.32. We report results for data stored in-memory, since the cost of main memory has decreased to the point that many interactive workloads can be performed entirely in-core. Each approximate query was started from a random position in the shuffled data. We found wall clock time to be stable for all approaches, and report times as the average of 3 runs for all methods.

C. Metrics

We gather several metrics in order to test two hypothesis: one, that our error bounding strategies in conjunction with our sampling strategies lead to speedups over simpler baselines; and two, that they do so without sacrificing correctness of query results.

Correctness of Query Results. The most important metric is the fraction of queries run that returned correct results, which we discuss briefly here. *Across all methods, all queries, and all parameter settings, results either matched the ground truth determined from an Exact evaluation, or were within error tolerance in the case of F-q1 and F-q7.* This is expected, given that we are consider SSI error bounders with strong probabilistic guarantees, coupled with the fact that our RangeTrim technique and system architecture do not compromise these guarantees. As such, we expect fewer than $\delta = 10^{-15}$ fraction of queries to yield incorrect results, which rounds down to a cool 0.

For the remaining experiments, we focus on the following metrics:

Estimate Error. For a given requested error bound ϵ supplied to applicable queries, we measure the actual error. The observed error should always fall within the requested error bound.

Wall-Clock Time. Our primary metric evaluates the end-to-end time required for various error bounders and various sampling strategies

	Time (s)	Avg Speedup over Exact (raw time in (s))			
		Exact	Hoeffding	Hoeffding+RT	Bernstein
F-q1	20.7	58.0× (0.4)	59.6× (0.3)	2129× (0.01)	2160× (0.01)
F-q2	42.2	233× (0.2)	316× (0.1)	2405× (0.02)	4683× (0.01)
F-q3	25.3	1.1× (23.4)	1.7× (15.1)	6.5× (3.9)	13.8× (1.8)
F-q4	20.7	12.8× (1.6)	12.7× (1.6)	922× (0.02)	925× (0.02)
F-q5	47.7	0.7× (68.8)	1.1× (44.4)	2.2× (21.4)	4.2× (11.5)
F-q6	66.4	1.1× (62.4)	1.1× (58.7)	11.2× (6.0)	16.7× (4.0)
F-q7	29.8	1.0× (30.0)	1.0× (29.1)	2.6× (11.6)	2.9× (10.4)
F-q8	47.7	2.0× (23.3)	1.0× (47.6)	9.1× (5.3)	9.5× (5.0)
F-q9	38.2	0.9× (41.0)	1.0× (38.6)	123× (0.3)	130× (0.3)
T-q1	32.4	2.2× (14.6)	3.4× (9.5)	17.8× (1.8)	29.4× (1.1)
T-q2	39.6	1.6× (24.5)	2.1× (19.0)	19.3× (2.0)	27.0× (1.5)
P-q1	20.9	17.8× (1.2)	17.9× (1.2)	311× (0.07)	314× (0.07)
P-q2	22.5	11.7× (1.9)	11.3× (2.0)	18.9× (1.2)	20.8× (1.1)

TABLE V: Summary of average query speedups and latencies for various error bounders.

(where the Exact baseline is included as a “sampling strategy”), across all the queries considered.

Number of Blocks Fetched. We also measure the number of blocks fetched from main memory into CPU cache when using various approaches. This is mainly due to the fact that error bounders incur additional CPU overhead and therefore wall-clock time, with Bernstein and Bernstein+RT incurring the highest overhead, so measuring blocks fetched for these approaches removes this confounding variable by decoupling performance from CPU attributes.

D. Results

In this section, we present results of our empirical study.

1) Impact of Error Bounder Used:

Summary. Using the Bernstein+RT error bounder resulted in typical speedups of at least 10× over Exact (for 10 out of 13 queries) and Hoeffding (for 9 out of 13 queries), and additionally was almost always on par or better than Bernstein (up to 2× faster).

We evaluate Hoeffding+RT and Bernstein+RT error bounders, along with Hoeffding and Bernstein (to ablate our RangeTrim technique) and an Exact query processor (to ablate any benefits due to approximation) against all the queries in Figure 4, with the resulting time measurements summarized in Table V.

First we note that all error bounders incur additional overhead — in the case of F-q5 where techniques like Hoeffding and Hoeffding+RT needed to process all the data in order to terminate (due to PMA), they actually ran *more slowly* than Exact. Using Bernstein, which does not suffer from PMA, yielded significant benefits over Exact, Hoeffding, and Hoeffding+RT across all queries. In cases where Hoeffding and Hoeffding+RT showed improvements over Exact, Bernstein amplified these improvements (F-q1, F-q2).

Using RangeTrim in conjunction with both Hoeffding and Bernstein typically led to similar performance, with a few queries exhibiting clearly superior performance (F-q5, F-q6, and F-q3). These queries have the following in common: they all have sparse groups with low selectivity (either because of the large number of groups in the case of F-q5 and F-q3, or because of the restrictive filter in the case of F-q5), and they are all “easy” to approximate, in that none of the groups require too many samples in order to achieve the relevant stopping condition. (F-q8 also has many groups, but some of them require many samples due to a large number of airports with average delay near the max.) This is an ideal condition for Bernstein+RT to show benefit: sparse groups will bottleneck the query, but RangeTrim will achieve termination faster since these sparse groups tend to have fewer outliers than do non-sparse groups. For such bottlenecking sparse groups, the range bounds for the DepDelay column are overly-conservative and dominate the sampling complexity. In this case, Bernstein, which

	Avg Speedup over Scan (time in seconds)		
	Scan	ActiveSync	ActivePeek
F-q5	40.1	2.0× (19.6)	3.5× (11.4)
F-q6	4.2	1.1× (3.8)	1.1× (3.9)
F-q7	10.3	1.0× (10.1)	1.0× (10.4)
F-q8	40.4	3.2× (12.7)	8.2× (5.0)
T-q1	1.2	1.2× (1.0)	1.1× (1.1)
T-q2	1.6	1.2× (1.4)	1.1× (1.5)
P-q2	11.3	7.1× (1.6)	10.4× (1.1)

TABLE VI: Summary of query speedups and latencies for various sampling strategies, restricted to GROUP BY queries that take more than 500ms for Scan with Bernstein+RT.

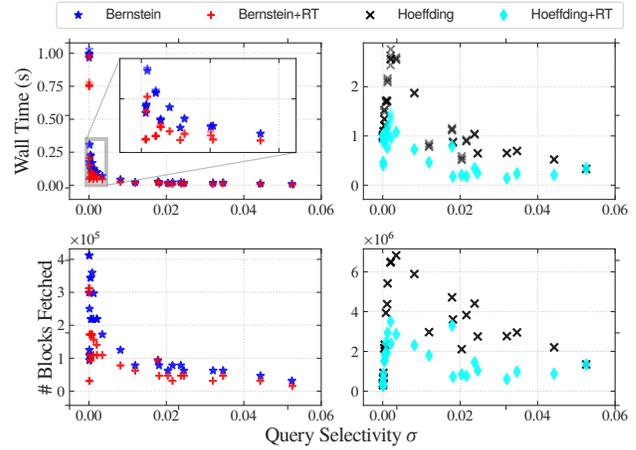


Fig. 5: Effect of query selectivity on wall time and blocks fetched, for selectivity determined by varying the origin airport used to filter F-q1 [ε = .5].

has PHOS, will require twice as many samples for such groups — and since these groups are the bottleneck, it will require roughly twice as much time, an intuition reflected in Table V.

2) Impact of Sampling Strategy Used:

Summary. ActiveSync sampling was typically on par or better than Scan, and ActivePeek sampling was typically on par or better than ActiveSync. ActiveSync significantly outperformed Scan on F-q5, F-q8, and P-q1 (by more than 3×), and ActivePeek significantly outperformed ActiveSync on the same set of queries.

We evaluate the impact of various sampling strategies when used in conjunction with the Bernstein+RT error bounder, the results of which are summarized in Table VI. In some cases (F-q5 and F-q8), the performance of the Scan baseline when used in conjunction with Bernstein+RT was on par with that of the Exact baseline, indicating that some form of block skipping can be crucial for queries with GROUP BYs. When implementing active scanning block by block as in ActiveSync, the improvement was most significant for F-q5, F-q8, and P-q1, with ActivePeek showing even further improvements for these queries. It is no coincidence that these are the same queries for which Scan performance using approximation is similar to Exact performance. This indicates that there were a few sparse groups preventing termination when Scan is used, which is the very case for which the greatest benefit can be derived from (an efficient implementation of) block skipping.

3) Impact of Data and Query Characteristics: To better understand various data- and query-dependent aspects of our techniques, we now study the effect of varying the parameters supplied to F-q1, F-q2, and F-q3.

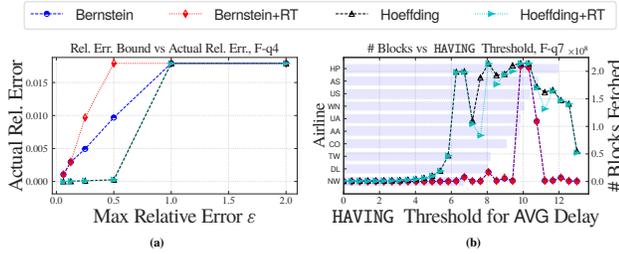


Fig. 6: (a): Effect of requested maximum relative error ϵ on actual relative error achieved for F-q1. (b): Data required for different HAVING thresholds used in F-q2. The group aggregates for also displayed for comparison.

Selectivity σ of filter.

Summary. As the fraction of tuples passing F-q1’s filter increases, wall clock time and blocks fetched both increase rapidly, then decrease, with RangeTrim giving the most benefit in the case of predicates with intermediate selectivity.

Different Origin attribute values used for filtering F-q1 have different selectivities. By varying the filter attribute value, we reveal interesting behavior impacted by the selectivity of the filter. (We consider selectivity as a number and not a quality, so that larger proportions of tuples satisfy predicates with higher selectivity.) For all four error bounding techniques considered, wall time and blocks fetched are plotted versus query selectivity in Figure 5. Bernstein and Bernstein+RT are plotted separately from Hoeffding and Hoeffding+RT for presentation.

For Hoeffding and Hoeffding+RT bounders, as selectivity increases, both wall-clock time and number of blocks fetched first increase rapidly, then decrease, in a strongly correlated fashion. This is likely because the sparsest filters require examining all the data before terminating, obviating early stopping benefits. After a certain point, however, early termination kicks in, happening more quickly as fewer tuples are filtered. The selectivity threshold for early termination appears to be much lower for Bernstein and Bernstein+RT bounders, which explains why wall-clock time and number of blocks fetched appear to be strictly decreasing with selectivity. The performance gap between techniques with and without RangeTrim generally decreases with increasing selectivity — perhaps because filters with higher selectivity tend to have range bounds that are not as conservative when compared with the a priori range bounds known to hold for the entire column.

ϵ for stopping condition 1.

Summary. For different upper bounds on relative error, the actual relative error in the query result is always within the requested error, for all error bounders applied to F-q1. The achieved relative error drops to 0 more quickly for the more conservative bounders Hoeffding and Hoeffding+RT as the requested error is decreased.

By varying the requested maximum relative error ϵ for query F-q1, we reveal its impact on the relative error achieved for various error bounders, shown in Figure 6(a). The main takeaways are that, for all error bounders, the achieved relative error generally decreases as the requested error bound decreases, with Hoeffding-based bounders dropping more quickly, as they are more conservative due to PMA.

HAVING threshold for stopping condition 2.

Summary. HAVING thresholds that are closer to group aggregates require more samples in order to achieve stopping condition 2, and Hoeffding-based error bounders in particular are more sensitive than Bernstein-based error bounders for the same threshold.

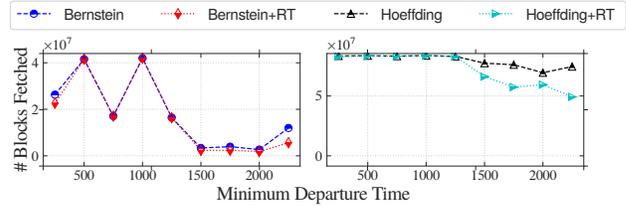


Fig. 7: Effect of minimum departure time on blocks fetched for F-q3.

By varying the HAVING threshold used to filter groups / airlines post-aggregate in F-q2 and measuring its effect on the number of blocks fetched for a particular query, we reveal interesting data-dependent behavior impacted by the true aggregates for each airline, depicted in Figure 6(b). This figure also plots the group aggregates using a horizontal bar chart sharing the same x-axis as the HAVING threshold, revealing that it is “harder” to determine which side of the HAVING threshold a given group is if its aggregate is close to the threshold. Indeed, from Figure 6(b), we see that the initial thresholds near 0 are very easy for all groups, allowing for very fast termination. The first spike in number of blocks fetched occurs between 6 and 7, corresponding to the aggregate for airline NW. At and after this point, we see spikes in blocks fetched for both Hoeffding-based and Bernstein-based error bounders whenever the threshold approaches one or more airline aggregate values, although we note that Bernstein-based error bounders appear to be more robust, requiring the threshold to be much closer before they are adversely affected as compared to Hoeffding-based bounders.

Minimum departure time for F-q3.

Summary. As the minimum departure time is increased, the spread of average delay between airlines increases, making it easier to separate the two airlines with the minimum average delays and achieve stopping condition 3 earlier. At the same time, termination becomes bottlenecked on sparse airlines, increasing the gap between similar bounders with and without RangeTrim.

By varying the minimum departure time $\$min_dep_time$ in F-q3, we reveal its impact on the number of blocks fetched for various error bounders, shown in Figure 7. This plot exhibits two interesting data-dependent behaviors worth unpacking. First, as the $\$min_dep_time$ increases, the variance in average delay between different airlines increases, perhaps because some airlines tend to have flights that are delayed more for later flights as compared with other airlines. This makes it easier to achieve stopping condition 3, since the average delays become more spread out with increasing minimum departure time, so we observe a decreasing trend in the number of blocks fetched. At the same time, as $\$min_dep_time$ increases, the selectivity of the various groups decreases. Since all the groups are sparser, the groups for which stopping condition 3 is bottlenecked are also sparser. Since we have an “easy” query (due to the higher variance between groups) for which sparse groups are bottlenecking termination, we tend to see a bigger performance gap between bounders with and without our RangeTrim technique.

Impact of ϵ on latency.

Summary. Figure 8 demonstrates that latency increases super-exponentially as ϵ decreases to 0, across all error bounders for F-q1[$\$origin=ORD$].

By varying the maximum relative error bound ϵ , we reveal its impact on latency for various bounders in F-q1[$\$origin=ORD$]. As depicted in Figure 8, latency increases super-exponentially with decreasing relative

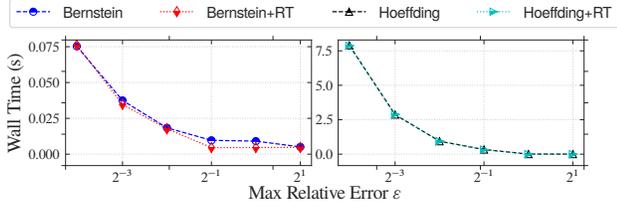


Fig. 8: Effect of ε on wall time for $F\text{-}q1[\text{origin}=\text{ORD}]$

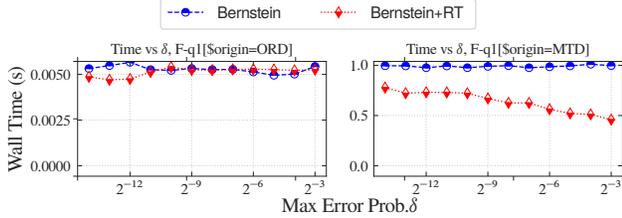


Fig. 9: Effect of δ on wall time for $F\text{-}q1[\text{origin}=\text{ORD}]$ and $F\text{-}q1[\text{origin}=\text{MTJ}]$, respectively.

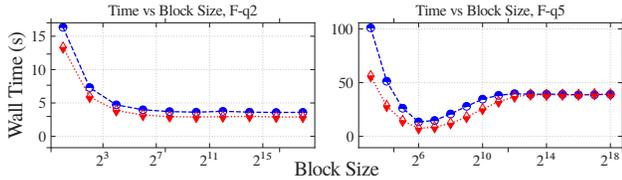


Fig. 10: Effect of block size on wall time for $F\text{-}q2$ and $F\text{-}q5$, respectively.

error threshold ε (note the log scale on the x-axis). From inspection of formulas for Hoeffding and Bernstein-style bounds, we would expect exponential behavior; the super-exponential behavior can be attributed to the conservative optional stopping procedure, which sacrifices some statistical efficiency. This figure illustrates the importance of being able to leverage other stopping conditions besides that of condition 1 in the case of queries that do not actually need to make use of the aggregate but merely use it for downstream decision making, since such queries can sometimes tolerate very large relative errors.

Impact of δ on latency.

Summary. Figure 9 illustrates that latency is robust to small δ 's, with sublinear slowdowns for exponentially decreasing δ .

By varying δ , we reveal its impact on latency for bounds Bernstein and Bernstein+RT on $F\text{-}q1[\text{origin}=\text{ORD}]$ and $F\text{-}q1[\text{origin}=\text{MTJ}]$. In the case of a high-selectivity predicate (i.e., that for $F\text{-}q1[\text{origin}=\text{ORD}]$), we see virtually no impact of δ on latency; for this high-selectivity predicate, the query is able to terminate after the first bounds computation for both Bernstein and Bernstein+RT. For the sparser predicate in $F\text{-}q1[\text{origin}=\text{MTJ}]$, δ does not impact the number of rounds of bounds computation needed for Bernstein, although a larger δ does gradually decrease this quantity for Bernstein+RT, albeit insignificantly considering the exponential increase in δ . Overall, this illustrates that the dependence of each error bound on $\sqrt{\log(1/\delta)}$ translates to query latency that is highly robust in δ , motivating our choice of 10^{-15} as the default δ .

Impact of block size on latency.

Summary. Figure 10 illustrates a quasiconvex relationship between query latency and block size.

By varying the block size, we reveal its impact on latency for $F\text{-}q2$ and $F\text{-}q5$. For both of these queries, smaller block sizes are associated with higher latency due to poorer cache locality and more frequent bounds recomputation. Higher block sizes, however, lose out on benefits from block skipping. This is not an issue for $F\text{-}q2$, which has relatively few groups (and none of which are sparse); for $F\text{-}q5$, however, it impacts latency adversely before leveling off. The level-off in both cases can be attributed to saturating the bitmap indexes, since larger block sizes will be more likely to contain tuples passing each group's filter.

VI. RELATED WORK

In this section, we survey related literature and highlight similarities and differences with this work.

Approximate Query Processing (AQP). We survey the AQP literature along two dimensions: first, online versus offline; second, approaches with strong versus asymptotic guarantees.

Online versus Offline AQP. Online sampling-based AQP schemes select samples as queries are issued, contrasted with offline schemes which compute strata ahead of time. Although our approach does perform a shuffle offline, it is nevertheless closer to online schemes, as it uses the scramble to compute samples on the fly as in [21, 29, 30, 31, 35]. Online schemes can use index structures like bitmaps to materialize relevant samples on-the-fly [15, 19, 20, 21], or obey an accuracy constraint for computing predefined aggregates without indices [43, 44]. Offline schemes, on the other hand, materialize samples ahead-of-time [12, 13, 17, 32] based off workload assumptions, sometimes tuning the computed strata as new workload information is available [12, 32].

While we implement our error bounders without PMA or PHOS in the context of a system for online AQP, our core algorithmic techniques are orthogonal to the exact approach, and could be paired with either online or offline schemes.

Sample-size-independent versus Asymptotic Guarantees. Most of the AQP systems from prior work have traditionally leveraged asymptotic error bounders [11, 12, 13, 16], though some have mentioned allowing either approach as an option [15]. Other approaches have leveraged deterministic [37, 45] or concentration-based error bounding techniques [17, 18, 19, 20, 21] under range-based or other very mild assumptions. In some cases, novel asymptotic error bounding techniques have been developed [2, 16, 37] to be used in conjunction with existing systems. Our approach is analogous to these, but instead of basing our techniques on asymptotic methods, we develop error bounding techniques with guarantees independent of sample size, starting from existing concentration-based methods and systematically ameliorating various pathologies.

Of particular note is the work of Agarwal et al. [46], which, as in this paper, recognizes both the pessimism of SSI techniques and the error-proneness of asymptotic techniques. They propose to run a *diagnostic procedure* in conjunction with asymptotic techniques in order to determine when such techniques are unable to yield accurate answers; however, the diagnostic procedure itself has no guarantees when used for query processing and can give both false positives and false negatives, which we consider unacceptable for the purposes of this paper.

Access Patterns for Informative Samples. A number of techniques have been developed to optimize access to relevant data for analytical queries. Please see the technical report [26] for a more extensive survey; here we focus on two techniques in particular: *outlier indexing* and *priority sampling*.

Outlier indexing [47] works by computing approximate aggregates derived by combining an estimate from the main table and an exact aggregate from the so-called “outlier index”, which stores all the rows with outlier values. The benefit of the outlier index is that it shrinks the range of the data from which samples are taken, allowing for faster convergence of approximate answers. One could think of the outlier index as an offline analogy of our own RangeTrim technique. Outlier indexing has some additional limitations that RangeTrim does not have; namely, it cannot be used to facilitate queries with aggregates involving arbitrary expressions, since such expressions can drastically change the set of outlying values. That said, for simple aggregates the two approaches are orthogonal, and could be leveraged together.

Priority sampling [48, 49] is also particularly useful for coping with outliers. While priority sampling applies in the presence of arbitrary filters and can furthermore be modified to allow for computation of AVG aggregates in addition to SUM, it has the drawback that the attribute or expression being aggregated must be known ahead of time (to say nothing of arbitrary expressions), so that the tuples can be sorted in descending order of priority, a limitation our techniques do not have.

Statistical Estimators and Confidence Intervals. The well-known error bounders in statistics and probability leverage asymptotic techniques [1, 3, 4, 50]. We already surveyed relatively more obscure SSI bounders in Section II when we discussed the empirical Bernstein-Serfling error bounder developed by Bardenet et al. [27], which we adapt for use in a database setting with our RangeTrim technique.

VII. CONCLUSION AND FUTURE WORK

We categorized existing conservative error bounders in terms of two pathologies, PMA and PHOS, and developed a technique, RangeTrim, for eliminating PHOS from any range-based error bounder. We showed how an SSI Bernstein-based bounder without PMA can significantly accelerate approximate queries, and how our RangeTrim technique, which eliminates PHOS, leads to an additional $2\times$ speedup in the best case, without ever hurting performance in the worst case. By implementing our distribution-aware techniques in the context of FastFrame, which prioritizes groups that require more samples in order to facilitate early termination, we demonstrate significant speedups (on the order of $10\times$ over both exact processing and traditional techniques based on Hoeffding) without losing guarantees. This suggests a viable path toward practical with-guarantees AQP for workload-agnostic analytics.

REFERENCES

- [1] B. Efron et al., “Bootstrap methods: Another look at the jackknife,” *The Annals of Statistics*, vol. 7, no. 1, pp. 1–26, 1979.
- [2] K. Zeng, S. Gao et al., “The analytical bootstrap: a new method for fast error estimation in approximate query processing,” in *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, 2014, pp. 277–288.
- [3] Student, “The probable error of a mean,” *Biometrika*, pp. 1–25, 1908.
- [4] J. Hájek, “Limiting distributions in simple random sampling from a finite population,” *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, vol. 5, pp. 361–374, 1960.
- [5] A. Pol and C. Jermaine, “Relational confidence bounds are easy with the bootstrap,” in *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, 2005, pp. 587–598.
- [6] B. Mozafari, “Approximate query engines: Commercial challenges and research opportunities,” in *SIGMOD*. ACM, 2017, pp. 521–524.
- [7] B. Mozafari and N. Niu, “A handbook for building an approximate query engine,” *IEEE Data Eng. Bull.*, vol. 38, no. 3, pp. 3–29, 2015.
- [8] K. Li and G. Li, “Approximate query processing: What is new and where to go?” *Data Science and Engineering*, vol. 3, no. 4, pp. 379–397, 2018.
- [9] D. Fisher, “Incremental, approximate database queries and uncertainty for exploratory visualization,” in *2011 IEEE Symposium on Large Data Analysis and Visualization*. IEEE, 2011, pp. 73–80.
- [10] B. C. Kwon, J. Verma, P. J. Haas et al., “Sampling for scalable visual analytics,” *IEEE computer graphics and applications*, vol. 37, no. 1, pp. 100–108, 2017.
- [11] S. Acharya, P. B. Gibbons et al., “The aqua approximate query answering system,” in *ACM Sigmod Record*, vol. 28, no. 2. ACM, 1999, pp. 574–576.
- [12] S. Agarwal et al., “Blinkdb: Queries with bounded errors and bounded response times on very large data,” in *EuroSys*. New York, NY, USA: ACM, 2013, pp. 29–42. [Online]. Available: <http://doi.acm.org/10.1145/2465351.2465355>
- [13] —, “Blink and it’s done: interactive queries on very large data,” 2012.
- [14] C. Jermaine et al., “Scalable approximate query processing with the dbo engine,” *ACM Transactions on Database Systems (TODS)*, vol. 33, no. 4, p. 23, 2008.
- [15] J. M. Hellerstein, P. J. Haas, and H. J. Wang, “Online aggregation,” *ACM SIGMOD Record*, vol. 26, no. 2, pp. 171–182, jun 1997. [Online]. Available: <http://dl.acm.org/citation.cfm?id=253262.253291>
- [16] Y. Park, B. Mozafari, J. Sorenson, and J. Wang, “Verdictdb: Universalizing approximate query processing,” in *Proceedings of the 2018 International Conference on Management of Data*, 2018, pp. 1461–1476.
- [17] B. Ding, S. Huang et al., “Sample + seek: Approximating aggregates with distribution precision guarantee,” in *SIGMOD*, 2016.
- [18] D. Alabi and E. Wu, “Pfunk-h: approximate query processing using perceptual models,” in *Proc. 1st Workshop on Human-In-the-Loop Data Analytics*, 2016, p. 10.
- [19] A. Kim, E. Blais, A. Parameswaran, P. Indyk, S. Madden, and R. Rubinfeld, “Rapid sampling for visualizations with ordering guarantees,” *PVLDB*, vol. 8, no. 5, pp. 521–532, Jan. 2015. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2735479.2735485>
- [20] S. Rahman, M. Aliakbarpour, K. Kong et al., “I’ve seen ‘enough’: Incrementally improving visualizations to support rapid decision making,” in *VLDB*, 2017.
- [21] S. Macke et al., “Adaptive sampling for rapidly matching histograms,” *Proceedings of the VLDB Endowment*, vol. 11, no. 10, pp. 1262–1275, 2018.
- [22] L. Valiant, *Probably Approximately Correct: Nature’s Algorithms for Learning and Prospering in a Complex World*. Basic Books (AZ), 2013.
- [23] P. J. Haas, *Hoeffding inequalities for join-selectivity estimation and online aggregation*. IBM, 1996.
- [24] W. Hoeffding, “Probability inequalities for sums of bounded random variables,” *Journal of the American statistical association*, vol. 58, no. 301, pp. 13–30, 1963.
- [25] R. J. Serfling, “Probability inequalities for the sum in sampling without replacement,” *The Annals of Statistics*, pp. 39–48, 1974.
- [26] S. Macke, M. Aliakbarpour, I. Diakonikolas, A. Parameswaran, and R. Rubinfeld, “Rapid approximate aggregation with distribution-sensitive interval guarantees,” Available at: <https://smacke.net/papers/ddavg.pdf>, Tech. Rep., 2020.
- [27] R. Bardenet, O.-A. Maillard et al., “Concentration inequalities for sampling without replacement,” *Bernoulli*, vol. 21, no. 3, pp. 1361–1385, 2015.
- [28] F. Olken, “Random sampling from databases,” Ph.D. dissertation, 1993.
- [29] C. Qin and F. Rusu, “Pf-ola: a high-performance framework for parallel online aggregation,” *Distributed and Parallel Databases*, vol. 32, no. 3, pp. 337–375, 2014.
- [30] S. Wu, B. C. Ooi, and K.-L. Tan, “Continuous sampling for online aggregation over multiple queries,” in *SIGMOD*. ACM, 2010, pp. 651–662.
- [31] K. Zeng, S. Agarwal, A. Dave et al., “G-ola: Generalized on-line aggregation for interactive analysis on big data,” in *SIGMOD*. ACM, 2015, pp. 913–918.
- [32] V. Ganti, M.-L. Lee, and R. Ramakrishnan, “Icicles: Self-tuning samples for approximate query answering,” in *VLDB*, vol. 176, no. 187, 2000.
- [33] C. Chen, W. Wang, X. Wang, and S. Yang, “Effective order preserving estimation method,” in *Australasian Database Conference*. Springer, 2016, pp. 369–380.
- [34] S. Gupta, S. Purandare, and K. Ramachandra, “Aggify: Lifting the curse of cursor loops using custom aggregates,” in *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, 2020, pp. 559–573.
- [35] X. Feng, A. Kumar, B. Recht, and C. Ré, “Towards a unified architecture for in-rdbms analytics,” in *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, 2012, pp. 325–336.
- [36] D. Lemire, “External-memory shuffling in linear time?” 2010 (accessed August 12, 2020), <https://lemire.me/blog/2010/03/15/external-memory-shuffling-in-linear-time/>.
- [37] P. J. Haas, “Large-sample and deterministic confidence intervals for online aggregation,” in *Proceedings. Ninth International Conference on Scientific and Statistical Database Management (Cat. No. 97TB100150)*. IEEE, 1997, pp. 51–62.
- [38] P. J. Haas and J. M. Hellerstein, “Ripple joins for online aggregation,” *ACM SIGMOD Record*, vol. 28, no. 2, pp. 287–298, 1999.
- [39] F. Chang, J. Dean et al., “Bigtable: A distributed storage system for structured data,” *ACM Transactions on Computer Systems (TOCS)*, vol. 26, no. 2, pp. 1–26, 2008.
- [40] “Flight Records,” <http://stat-computing.org/dataexpo/2009/the-data.html>, 2009.
- [41] “NYC Taxi Trip Records,” <https://github.com/toddschneider/nyc-taxi-data>, 2015.
- [42] “WA Police Stop Records,” <https://stacks.stanford.edu/file/druid:py883nd2578/WA-clean.csv.gz>, 2017.
- [43] W.-C. Hou, G. Ozsoyoglu, and B. K. Taneja, “Statistical estimators for relational algebra expressions,” in *Proceedings of the seventh ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, 1988, pp. 276–287.
- [44] —, “Processing aggregate relational queries with hard time constraints,” in *ACM SIGMOD Record*, vol. 18, no. 2. ACM, 1989, pp. 68–77.
- [45] N. Potti and J. M. Patel, “Daq: a new paradigm for approximate query processing,” *Proceedings of the VLDB Endowment*, vol. 8, no. 9, pp. 898–909, 2015.
- [46] S. Agarwal et al., “Knowing when you’re wrong,” in *SIGMOD*. New York, New York, USA: ACM Press, Jun. 2014, pp. 481–492. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2588555.2593667>
- [47] S. Chaudhuri, G. Das, M. Datar, R. Motwani, and V. Narasayya, “Overcoming limitations of sampling for aggregation queries,” in *ICDE*. IEEE, 2001, pp. 534–542.
- [48] N. Duffield, C. Lund, and M. Thorup, “Priority sampling for estimation of arbitrary subset sums,” *Journal of the ACM (JACM)*, vol. 54, no. 6, pp. 32–es, 2007.
- [49] N. Alon, N. Duffield, C. Lund, and M. Thorup, “Estimating arbitrary subset sums with few probes,” in *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, 2005, pp. 317–325.
- [50] B. Efron, “Bootstrap methods: another look at the jackknife,” in *Breakthroughs in statistics*. Springer, 1992, pp. 569–593.