

Predicting Query Execution Time: Are Optimizer Cost Models Really Unusable?

Wentao Wu[†] Yun Chi[‡] Shenghuo Zhu[‡] Junichi Tatemura[‡]
Hakan Hacigümüş[‡] Jeffrey F. Naughton[†]

[†]*Computer Sciences Department, University of Wisconsin, Madison, WI, USA*
{wentao, naughton}@cs.wisc.edu

[‡]*NEC Laboratories America, Cupertino, CA, USA*
{ychi, zsh, tatemura, hakan}@sv.nec-labs.com

Abstract—Predicting query execution time is useful in many database management issues including admission control, query scheduling, progress monitoring, and system sizing. Recently the research community has been exploring the use of statistical machine learning approaches to build predictive models for this task. An implicit assumption behind this work is that the cost models used by query optimizers are insufficient for query execution time prediction. In this paper we challenge this assumption and show while the simple approach of scaling the optimizer’s estimated cost indeed fails, a properly calibrated optimizer cost model is surprisingly effective. However, even a well-tuned optimizer cost model will fail in the presence of errors in cardinality estimates. Accordingly we investigate the novel idea of spending extra resources to refine estimates for the query plan after it has been chosen by the optimizer but before execution. In our experiments we find that a well calibrated query optimizer model along with cardinality estimation refinement provides a low overhead way to provide estimates that are always competitive and often much better than the best reported numbers from the machine learning approaches.

I. INTRODUCTION

Predicting query execution time has always been desirable if somewhat elusive capability for database management systems. This capability has received a flurry of attention recently, perhaps because it has become increasingly important in the context of offering databases as a service (DaaS). A DaaS provider has to manage infrastructure costs as well as honor service level agreements (SLAs), and many system management decisions can benefit from prediction of query execution time, including:

- *Admission Control*: Knowing the execution time of an incoming query can enable cost-based decisions on admission control [26], [28].
- *Query Scheduling*: Knowing the query execution time is crucial in deadline and latency aware scheduling [8], [13].
- *Progress Monitoring*: Knowing the execution time of an incoming query can help avoid “rogue queries” that are submitted in error and take an unreasonably long time to execute [22].
- *System Sizing*: Knowing query execution time as a function of hardware resources can help in system sizing [27].

Recent work on predicting query execution time [4], [11], [26], [28] has focused on various machine learning techniques,

which treat the database system as a black box and try to learn a query running time prediction model. This move toward black box machine learning techniques is implicitly and sometimes explicitly motivated by a belief that query optimizers’ cost estimations are not good enough for run time prediction. For example, in [11], the authors found that using linear regression to map the cost from Neoview’s commercial query optimizer to the actual running was not effective (see Figure 17 of [11]). In [4], the same approach was used to map PostgreSQL’s estimate to the actual execution time, and similar disappointing results were obtained (see Figure 5 of [4]).

It is clear from this previous work that post-processing the optimizer cost estimate is not effective. However, we argue in this paper that this does not imply that optimizer estimates are not useful — to the contrary, our experiments show that if the optimizer’s internal cost model is tuned before making the estimate, the optimizer’s estimates are competitive with and often superior to those obtained by more complex approaches. In more detail, for specificity consider the cost model used by the PostgreSQL query optimizer:

Example 1 (PostgreSQL’s Cost Models): PostgreSQL’s optimizer uses a vector of five parameters (referred to as *cost units*) in its cost model: $\mathbf{c} = (c_s, c_r, c_t, c_i, c_o)^T$, defined as follows:

- 1) c_s : *seq_page_cost*, the I/O cost to sequentially access a page.
- 2) c_r : *random_page_cost*, the I/O cost to randomly access a page.
- 3) c_t : *cpu_tuple_cost*, the CPU cost to process a tuple.
- 4) c_i : *cpu_index_tuple_cost*, the CPU cost to process a tuple via index access.
- 5) c_o : *cpu_operator_cost*, the CPU cost to perform an operation such as hash or aggregation.

The cost C_O of an operator O in a query plan is then computed by a linear combination of c_s , c_r , c_t , c_i , and c_o :

$$C_O = \mathbf{n}^T \mathbf{c} = n_s \cdot c_s + n_r \cdot c_r + n_t \cdot c_t + n_i \cdot c_i + n_o \cdot c_o. \quad (1)$$

The values $\mathbf{n} = (n_s, n_r, n_t, n_i, n_o)^T$ here represent *the number of pages sequentially scanned, the number of pages randomly accessed, and so forth, during the execution of the operator*

O . The total estimated cost of a query plan is then simply the sum of the costs of the individual operators in the query plan. \square

The accuracy of C_O hence depends on both the accuracy of the c 's and the n 's. In PostgreSQL, by default, $c_s = 1.0$, $c_r = 4.0$, $c_t = 0.01$, $c_i = 0.005$, and $c_o = 0.0025$. The cost C_O in Equation (1) is thus reported in units of sequential page I/O cost (since $c_s = 1.0$). Note that these cost units were somewhat arbitrarily set by the optimizer designers with no knowledge of the system on which the query is actually being run. Using linear regression to map an estimate so obtained will only work if the ratios among these units is correct, and not surprisingly, these default ratios were far from correct on our systems.

Of course, the accuracy of C_O also depends on the quantities n_s , n_r , n_t , n_i , and n_o . Determining accurate values for these quantities is not a matter of calibration — rather, it is a matter of good cardinality estimation in the optimizer. Hence one could say that we have reduced the problem of query time prediction to the previously unsolved problem of cardinality estimation. In a sense this is true, but further reflection reveals that we are solving a subtly but significantly different problem.

In their traditional role, cardinality estimates are required for every cardinality encountered as the optimizer searches thousands or tens of thousands of alternative plans. This of course means that the estimation process itself must be extremely efficient, or long optimization times will result. But our problem is different: we must determine cardinalities for the single plan that the optimizer has already chosen. The fact that we are working on a single plan means we can afford to spend some extra time to improve the original optimizer estimates. Specifically, in this paper we consider using sampling-based approaches to refine these estimates. We believe that although sampling-based approaches may be too expensive to be used while *searching* for good query plans, they can be practically used for *correcting* the erroneous cardinality estimates in a ready-to-be-executed query plan.

Our experiments show that if we correctly calibrate the constants in the optimizer's cost model, it yields good query execution time estimates when the cardinality estimates are good (as is the case in, for example, the uniformly distributed data set variants of the TPC-H benchmark.) Furthermore, "expensive" techniques such as sampling can be effectively used to improve the cardinality estimates for the chosen plan. Putting the two together yields cost estimates that are as good or better than those obtained by previously studied "black box" machine learning approaches.

The rest of the paper is organized as follows. We first give an overview of our cost-model based approach in Section II. We then discuss the two error-correction steps, i.e., calibrating cost units and cardinality estimates, in Sections III and IV, respectively. We further conduct extensive experimental evaluations and present our results in Section V. We summarize related work in Section VI and conclude the paper in Section VII.

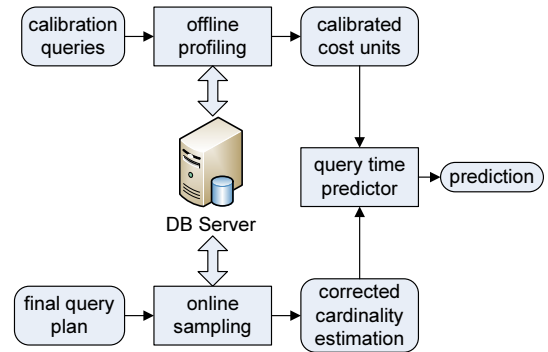


Fig. 1. The architecture of our framework.

II. OUR FRAMEWORK

Example 1 demonstrates that errors in c and n might prevent us from leveraging $C_O = \mathbf{n}^T \mathbf{c}$ to predict query execution time. Our basic idea is simply to correct these errors in a principled fashion.

As illustrated in Figure 1, our framework consists of two error-correction stages, namely, an *offline* profiling stage to calibrate the c , and an *online* sampling stage to refine the n :

- *Offline profiling to calibrate c:*
The errors in c reflect an inaccurate view of the underlying hardware and database system. To correct this, instead of using the default values assigned by the query optimizer, we calibrate them by running a family of profiling queries on the system on which the queries will be run. Note that this profiling stage is offline. Moreover, it only needs to be run once as long as the underlying hardware and software configuration does not change. We describe the details of the profiling stage in Section III.
- *Online sampling to refine n:*
The errors in n reflect errors in cardinality estimation. Once the query plan has been chosen by the query optimizer, we re-estimate cardinalities, if necessary, using a sampling-based approach. Although using sampling for cardinality estimation is well-known, current DBMS optimizers exclude sampling from their implementations, perhaps due to the additional overhead sampling incurs. However, since we only have to estimate cardinalities for one plan, the overhead of sampling is affordable in practice. We describe the details of the sampling stage in Section IV. We note that the important idea is that there is an opportunity to spend extra time refining cardinality estimates once a query plan has been chosen; sampling is one example of how that could be done, and finding other techniques is an interesting area for future work.

The advantages of our framework for predicting query execution time include the following:

- *Lightweight:* The profiling step is fast and so can be conducted in a new hardware environment quickly. The sampling step, as will be shown, introduces small (usually $< 10\%$) and tunable overhead.

- *No training data needed*: Unlike machine-learning-based approaches, which heavily rely on training data representative of the actual workload, our framework does not rely on such a training data set and so can handle *ad hoc* queries well.
- *White-box approach*: Instead of sophisticated statistical models (e.g., SVM and KCCA), which are often difficult to understand for non-experts, our framework adopts an intuitive white-box approach that fits naturally into the existing paradigm of query optimization and evaluation in relational database systems.

III. CALIBRATING COST UNITS

In this section we consider the task of calibrating the cost units in the optimizer cost model to match the true performance of the hardware and software on which the query will be run. It turns out that closely related problems have been studied in the context of heterogeneous DBMS [9], DB resource virtualization [24], and storage type selection [29]. Previous work, however, has either focused on cost models for particular operators (e.g., selections and 2-way joins in [9]), or on a subset of cost units dedicated to a particular subsystem of the DBMS (e.g., CPU in [24] and I/O in [29]). We build on this previous work following their technique of using a set of *calibration queries*. The basic idea is that for each cost unit to be calibrated, one designs some particular query that isolates this parameter from the others. In practice, this is not completely straightforward in that not every cost unit can be isolated in a single query.

A. Guiding Principles

Ideally, we wish to have one calibration query for each parameter. However, this is not always possible. For instance, there is no SQL query which involves the *cpu_operator_cost* but not the *cpu_tuple_cost*. A natural generalization is then to use k calibration queries for k parameters, as was done in [24]. The following example illustrates this idea.

Example 2 (Calibration Query): Suppose R is some relation that is buffer pool resident. We can use the following two calibration queries to obtain the parameters *cpu_tuple_cost* and *cpu_operator_cost*:

- Q1: `SELECT * FROM R`
- Q2: `SELECT COUNT(*) FROM R`

Since R is memory resident, there is no I/O cost for Q1 or Q2. Q1 only involves the parameter *cpu_tuple_cost*, while Q2 involves both *cpu_tuple_cost* and *cpu_operator_cost* (due to the `COUNT` aggregation). Suppose the execution time of Q1 and Q2 are t_1 and t_2 , respectively. Since the overhead due to *cpu_tuple_cost* for Q1 and Q2 are the same, we can then infer *cpu_operator_cost* with respect to the execution time $t_2 - t_1$.

Specifically, let the number of tuples processed be n_t , and the number of CPU operations be n_o , as in Equation (1). In PostgreSQL’s cost model, a CPU operation means things like *adding two integers*, *hashing an attribute*, and so on. n_o is thus the number of such operations performed. On the other hand, n_t is the number of input *tuples* (sometimes the output

tuples or the sum of both, depending on the specific operator). Here, for Q1, the cost model will only charge one c_t per tuple, since the CPU merely reads in the tuple without any further processing. For Q2, in addition to charging one c_t per tuple for reading it, the cost model will also charge one c_o per tuple for doing the aggregation (i.e., `COUNT`), and hence the total CPU cost is estimated to be $n_t c_t + n_o c_o$. Note that for this particular query Q2, we coincidentally have $n_t = n_o = |R|$. In general, n_t and n_o could be different. For example, for the *sort* operator, n_t is the number of input tuples, and n_o is the number of comparisons made. In this case, $n_o = n_t \log n_t$.

Now suppose that T_t is the time for the CPU to process one tuple, and T_o is the time for one CPU operation. We then have

$$\begin{aligned} t_1 &= T_t \cdot n_t, \\ t_2 &= T_t \cdot n_t + T_o \cdot n_o. \end{aligned}$$

Solving these two equations gives us the values of T_t and T_o , in turn determines c_t and c_o in Equation (1). \square

In general, with a set of calibration queries Q , we first estimate \mathbf{n}_i for each $q_i \in Q$ and then measure its execution time t_i . With $\mathbf{N} = (\mathbf{n}_1, \dots, \mathbf{n}_k)^\top$ and $\mathbf{t} = (t_1, \dots, t_k)^\top$, we can solve the following equation for \mathbf{c} :

$$\mathbf{N}\mathbf{c} = \mathbf{t},$$

which is just a system of k equations.

The next question is then, given a set of optimizer cost units \mathbf{c} , how to design a set of calibration queries Q . Our goal is to design a set with the following properties:

- *Completeness*: Each cost unit in \mathbf{c} should be covered by at least one calibration query $q \in Q$.
- *Conciseness*: Each query $q \in Q$ should be necessary to guarantee completeness. In other words, any subset $Q' \subset Q$ is not complete.
- *Simplicity*: Each query $q \in Q$ should be as simple as possible when Q is both complete and concise.

Clearly “completeness” is mandatory, while the others are just “desirable”. Since the possible number of SQL queries on a given database is infinite, we restrict our attention to concise complete subsets. However, there is still infinite number of sets Q that are both complete and concise. Simpler queries are preferred over more complex, because it is easier to obtain correct values for the cardinalities such as n_t and n_o in Example 2 (getting exact values for such cardinalities may be difficult for operators embedded in deep query trees).

B. Implementation

We designed the 5 calibration queries for the PostgreSQL optimizer as follows. We chose queries q_i for Q by introducing individual cost units one by one:

- *cpu_tuple_cost*: We use query q_1 :
`SELECT * FROM R`
as the calibration query. The relation R is first paged into the buffer pool, and hence there is no I/O cost involved:
 $\mathbf{n}_1 = (0, 0, n_{t1}, 0, 0)^\top$.

- *cpu_operator_cost*: We use query q_2 :
`SELECT COUNT (*) FROM R`
as another calibration query. We then use the method illustrated in Example 2. Again, R is memory resident: $\mathbf{n}_2 = (0, 0, n_{t2}, 0, n_{o2})^\top$ and $n_{t2} = n_{t1}$.
- *cpu_index_tuple_cost*: We use query q_3 :
`SELECT * FROM R WHERE R.A < a`
where $R.A$ has a clustered index and we pick a so that the optimizer will choose an index scan. This query involves *cpu_tuple_cost*, *cpu_index_tuple_cost*, and *cpu_operator_cost*. Once again, R is memory resident: $\mathbf{n}_3 = (0, 0, n_{t3}, n_{i3}, n_{o3})^\top$.
- *seq_page_cost*: We use query q_4 :
`SELECT * FROM R`
as the calibration query. This query will be executed in a sequential scan, and the cost model only involves overhead in terms of *seq_page_cost* and *cpu_tuple_cost*: $\mathbf{n}_4 = (n_{s4}, 0, n_{t4}, 0, 0)^\top$.
- *rand_page_cost*: We use query q_5 :
`SELECT * FROM R where R.B < b`
as the calibration query. Here $R.B$ is some attribute of the relation R on which an *unclustered* index is built. The values of B are uniformly generated, and we pick b so that the optimizer chooses an index scan. Ideally, we would like that the qualified tuples were *completely* randomly distributed so that we could isolate the parameter *rand_page_cost*. However, in practice, *pure* random access is difficult to achieve, since the execution subsystem can first determine the pages that need to be accessed based on the qualified tuples before it actually accesses the pages. In this sense, local sequential accesses are unavoidable, and the query plan involves more or less overhead in terms of *seq_page_cost*. In fact, a typical query plan of this query will contain all the five parameters: $\mathbf{n}_5 = (n_{s5}, n_{r5}, n_{t5}, n_{i5}, n_{o5})^\top$.

Notice that \mathbf{n}_i can be estimated relatively accurately due to simplicity of q_i . Furthermore, the 5 equations generated by the 5 queries are *independent*, which guarantees the existence of a unique solution for \mathbf{c} . This can be easily seen by observing the matrix $\mathbf{N} = (\mathbf{n}_1, \dots, \mathbf{n}_5)^\top$, namely,

$$\mathbf{N} = \begin{pmatrix} 0 & 0 & n_{t1} & 0 & 0 \\ 0 & 0 & n_{t2} & 0 & n_{o2} \\ 0 & 0 & n_{t3} & n_{i3} & n_{o3} \\ n_{s4} & 0 & n_{t4} & 0 & 0 \\ n_{s5} & n_{r5} & n_{t5} & n_{i5} & n_{o5} \end{pmatrix}.$$

Note that the determinant $|\mathbf{N}|$ satisfies $|\mathbf{N}| > 0$, since by rearranging the columns of \mathbf{N} , we can make it a triangular matrix.

To make this approach more robust, our implementation uses multiple queries for each q_i and finds the best-fitting of \mathbf{c} . This is done by picking different relations R and different values for the a in the predicates of the form $R.A < a$.

IV. REFINING CARDINALITY ESTIMATION

We discuss how to refine \mathbf{n} in this section. To make this paper self-contained, we first discuss how the optimizer obtains \mathbf{n} for a given query plan. We then propose a sampling-based method of refining the cardinality estimates (and hence the \mathbf{n}) of the final plan chosen by the optimizer. We describe the details of the algorithm and our current implementation.

A. Optimizer's Estimation of \mathbf{n}

The optimizer estimates query execution cost by aggregating the cost estimates of the operators in the query plan. To distinguish blocking and non-blocking operators, this cost model comprises of the *start_cost* and *total_cost* of each operator:

- *start_cost* (sc) is the cost before the operator can produce its first output tuple;
- *total_cost* (tc) is the cost after the operator generates all of its output tuples.

Note that the cost of an operator includes the cost of its child operators.

As an example, we show how \mathbf{n} is derived for the *in-memory sort* and *nested-loop join* operators in PostgreSQL. These operators are representative of blocking and non-blocking operators, respectively. In the following illustration, *run_cost* (rc for short) is defined as $rc = tc - sc$, and N_t is the (estimated) number of input tuples for the operator. Observe that the costs are given as linear combinations of \mathbf{c} .

Example 3 (In-Memory Sort): *Quick sort* is used for tables that optimizer estimates can be completely held in memory. The values sc and rc are estimated as follows:

$$\begin{aligned} sc &= 2 \cdot c_o \cdot N_t \cdot \log N_t + tc \text{ of child,} \\ rc &= c_t \cdot N_t. \end{aligned}$$

□

Example 4 (Nested-Loop Join): The *nested-loop join* operator joins two input relations. The sc and rc are estimated as follows:

$$\begin{aligned} sc &= sc \text{ of outer child} + sc \text{ of inner child,} \\ rc &= c_t \cdot N_t^o \cdot N_t^i + N_t^o \cdot rc \text{ of inner child.} \end{aligned}$$

Here N_t^o and N_t^i are the number of input tuples from the outer and inner child operator, respectively. □

Notice that the main uncertainty in \mathbf{n} comes from the estimated input cardinality N_t in both of these examples. In general, the logic flow in the cost models of PostgreSQL optimizer can be summarized with five steps:

- 1) estimate the input/output cardinality;
- 2) compute the CPU cost based on cardinality estimates;
- 3) estimate the number of accessed pages according to the cardinality estimates;
- 4) compute the I/O cost based on estimates of accessed pages;
- 5) compute the total cost as the sum of CPU and I/O cost.

Hence, our main task in calibrating \mathbf{n} is to refine the input/output cardinalities for each operator.

B. Cardinality Refinement

As mentioned in the introduction, traditionally cardinality estimation has had to satisfy strict performance constraints because it is done for every plan considered by the optimizer. This has led to compromises that may produce inaccuracies in estimates that are too large for run time estimation.

Our goal is to *refine* the cardinality estimates for the plan chosen by the optimizer. Clearly, this will increase the overhead of the optimization phase. However, the key insight here is that because the refinement procedure only needs to be performed *once* per query, rather than once per plan, we can afford to spend more time than is possible for traditional cardinality estimation.

C. A Sampling-Based Approach

In principle, any approach that can improve cardinality estimation can be applied. We use a generalized version of the *sequential-sampling* estimator proposed in [14] for the following two reasons:

- It incorporates a tunable trade-off between efficiency (i.e., the number of samples taken) and effectiveness (i.e., the precision of the estimates);
- It can simultaneously estimate cardinalities for multiple operators in the query plan.

In this paper, we extend the framework in [14] in the following two aspects:

- The estimator described in [14] is for join queries. We generalize the framework to queries with arbitrary number of selections and joins. We prove that this extension preserves the two key properties, namely, *unbiasedness* and *strong consistency*, of the original estimator.
- The framework described in [14] uses random disk accesses to take samples. In comparison, we propose to take samples offline, store them as materialized views, and directly use them at runtime. This greatly reduces the runtime overhead and requires very minimal changes to the database engine (e.g., a few hundred lines of C code in the case of PostgreSQL). We further show that this offline sampling preserves the semantics of the original online sampling.

We next first describe the estimator in its generalized form, and then describe our cardinality refinement algorithm and its implementation details.

D. The Estimator

Let \mathcal{D} be a database consisting of K relations R_1, \dots, R_K . Suppose that R_k is partitioned into m_k blocks each with size N_k , namely, $|R_k| = m_k N_k$. Consider the two basic relational operators: *selection* σ_F (F is a boolean formula representing the selection condition), and *cross-product* \times . For σ_F , we define the output of an input block B to be $\sigma_F(B)$. For \times , we define the output of the two input blocks B and B' to be $B \times B'$. Instead of estimating the cardinality of the output relation directly, the estimator will estimate the *selectivity* of the operator, which is defined as the output cardinality divided

by the input cardinality. Specifically, the selectivity of the selection operator σ_F is $\rho_R = |\sigma_F(R)|/|R|$ where R is the input relation. Moreover, the selectivity of σ_F on a particular block B of R is $\rho_B = |\sigma_F(B)|/|B|$. On the other hand, the selectivity of the cross-product operator \times is always 1. It is then straightforward to obtain the output cardinality once we know the selectivity of the operator¹.

In the following, without loss of generality, we will assume that the query considered is over relations R_1, \dots, R_K , and we use the notation $\mathbf{R} = R_1 \times \dots \times R_K$. Let $B(k, j)$ be the j -th block of relation k ($1 \leq j \leq m_k$, and $1 \leq k \leq K$). We use $\mathbf{B}(L_{1,i_1}, \dots, L_{K,i_K})$ to represent $B(1, L_{1,i_1}) \times \dots \times B(K, L_{K,i_K})$, where $B(k, L_{k,i_k})$ is the block (with index L_{k,i_k}) randomly picked from the relation R_k in the i_k -th sampling step. Moreover, we use the notation \mathbf{B}_i if $i_1 = i_2 = \dots = i_K = i$.

Lemma 1: Consider $\sigma_F(\mathbf{R})$. Let $\mathbf{B}_1, \dots, \mathbf{B}_n$ be a sequence of n random samples (with replacement) from \mathbf{R} . Define $\rho_{\mathbf{B}_i} = |\sigma_F(\mathbf{B}_i)|/|\mathbf{B}_i|$ ($1 \leq i \leq n$). Then $E[\rho_{\mathbf{B}_i}] = \rho_{\mathbf{R}}$.

Proof: We have

$$\rho_{\mathbf{R}} = \frac{|\sigma_F(\mathbf{R})|}{|\mathbf{R}|} = \frac{\sum_{k=1}^K \sum_{j=1}^{m_k} |\sigma_F(B(k, j))|}{\prod_{k=1}^K m_k N_k}.$$

Since

$$E[|\sigma_F(\mathbf{B}_i)|] = \frac{\sum_{k=1}^K \sum_{j=1}^{m_k} |\sigma_F(B(k, j))|}{\prod_{k=1}^K m_k},$$

we thus have

$$E[\rho_{\mathbf{B}_i}] = \frac{E[|\sigma_F(\mathbf{B}_i)|]}{\prod_{k=1}^K N_k} = \frac{\sum_{k=1}^K \sum_{j=1}^{m_k} |\sigma_F(B(k, j))|}{\prod_{k=1}^K m_k N_k} = \rho_{\mathbf{R}}. \quad \blacksquare$$

Define

$$\tilde{\rho}_{\mathbf{R}} = \frac{1}{n} \sum_{i=1}^n \rho_{\mathbf{B}_i}.$$

Then it is easy to see from Lemma 1 that $E[\tilde{\rho}_{\mathbf{R}}] = \rho_{\mathbf{R}}$. Moreover, since the random variables $\rho_{\mathbf{B}_i}$ are i.i.d., by the strong law of large numbers, we have $Pr[\lim_{n \rightarrow \infty} \tilde{\rho}_{\mathbf{R}} = \rho_{\mathbf{R}}] = 1$. We summarize this result in the following lemma:

Lemma 2: $E[\tilde{\rho}_{\mathbf{R}}] = \rho_{\mathbf{R}}$, and $Pr[\lim_{n \rightarrow \infty} \tilde{\rho}_{\mathbf{R}} = \rho_{\mathbf{R}}] = 1$.

Lemma 2 can be generalized to queries with arbitrary number of selections and joins:

¹Note that the input cardinality is already known before the estimation procedure runs. It is simply the product of the cardinalities of the underlying relations that are input to the operator, which can be directly obtained from the statistics stored in system catalogs.

Theorem 1: Let q be any query involving only selections and joins over \mathbf{R} , and let ρ_q be the selectivity of q . Then $E[\tilde{\rho}_q] = \rho_q$, and $Pr[\lim_{n \rightarrow \infty} \tilde{\rho}_q = \rho_q] = 1$.

Proof: The proof is easy by noticing that q can be written as its normal form [2]: $\sigma_F(\mathbf{R})$. We then apply Lemma 1 to complete the proof. ■

In statistical terminology, the estimator $\tilde{\rho}_q$ is *unbiased*, and *strongly consistent* for ρ_q : the more samples we take, the closer $\tilde{\rho}_q$ is to ρ_q . This gives us a way to control the trade-off between the estimation accuracy and the number of samples we take.

The estimator we just described takes samples from each relation uniformly and independently (called *independent sampling* [14]). Therefore, after n steps, we have n observations in total. In [14], the authors further discussed another alternative called *cross-product sampling*. The idea is that, at the i -th step, assuming the K blocks taken from the K relations are $B(1, L_{1,i}), \dots, B(K, L_{K,i})$, we can actually join each $B(k, L_{k,i})$ with each $B(k', L_{k',i'})$ such that $1 \leq i' \leq i$ and $k' \neq k$ (note that in the case of independent sampling, we only join among the blocks with $i' = i$). In this way, we can obtain n^K observations after n steps.

Define

$$\tilde{\rho}_{\mathbf{R}}^{cp} = \frac{1}{n^K} \sum_{i_1=1}^n \cdots \sum_{i_K=1}^n \rho_{\mathbf{B}(L_{1,i_1}, \dots, L_{K,i_K})}. \quad (2)$$

From Lemma 1, it is clear that $\tilde{\rho}_{\mathbf{R}}^{cp}$ is still unbiased, i.e., $E[\tilde{\rho}_{\mathbf{R}}^{cp}] = \rho_{\mathbf{R}}$. However, since now the $\rho_{\mathbf{B}(L_{1,i_1}, \dots, L_{K,i_K})}$'s are no longer independent, we cannot directly apply the strong law of large numbers to show the strong consistency of $\tilde{\rho}_{\mathbf{R}}^{cp}$, although it still holds here:

Lemma 3: $E[\tilde{\rho}_{\mathbf{R}}^{cp}] = \rho_{\mathbf{R}}$, and $Pr[\lim_{n \rightarrow \infty} \tilde{\rho}_{\mathbf{R}}^{cp} = \rho_{\mathbf{R}}] = 1$.

Proof: The proof idea is briefly sketched in Appendix D of [14]. Define $X(i, j) = 1$ if the block we take from the relation R_1 in the i -th sampling step is $B(1, j)$, and $X(i, j) = 0$ otherwise, where $1 \leq j \leq m_1$. Define $\mathbf{X}_i = (X(i, 1), \dots, X(i, m_1))$. Since we sample with replacement, $\mathbf{X}_1, \dots, \mathbf{X}_n$ are i.i.d. random vectors with (common) expectation

$$\mu_{\mathbf{X}} = E[\mathbf{X}_i] = \left(\frac{1}{m_1}, \dots, \frac{1}{m_1}\right).$$

Let

$$\bar{\mathbf{X}}_n = \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i = (\bar{X}(n, 1), \dots, \bar{X}(n, m_1)),$$

where $\bar{X}(n, j) = \frac{1}{n} \sum_{i=1}^n X(i, j)$. Then by the strong law of large numbers, $Pr[\lim_{n \rightarrow \infty} \bar{\mathbf{X}}_n = \mu_{\mathbf{X}}] = 1$.

On the other hand, define

$$Y(n, j) = \frac{1}{n^{K-1}} \sum_{i_2=1}^n \cdots \sum_{i_K=1}^n \rho_{\mathbf{B}(j, L_{2,i_2}, \dots, L_{K,i_K})},$$

and $\mathbf{Y}_n = (Y(n, 1), \dots, Y(n, m_1))$. We have

$$\mu_{\mathbf{Y}} = E[\mathbf{Y}_n] = (\rho_1, \rho_2, \dots, \rho_{m_1}),$$

where

$$\rho_j = \frac{1}{\prod_{k=2}^K m_k} \sum_{i_2=1}^{m_2} \cdots \sum_{i_K=1}^{m_K} \rho_{\mathbf{B}(j, i_2, \dots, i_K)}.$$

Note that we can now write Equation (2) as

$$\tilde{\rho}_{\mathbf{R}}^{cp} = \bar{\mathbf{X}}_n \cdot \mathbf{Y}_n = \sum_{j=1}^{m_1} \bar{X}(n, j) Y(n, j).$$

Our next goal is to show that $Pr[\lim_{n \rightarrow \infty} \mathbf{Y}_n = \mu_{\mathbf{Y}}] = 1$. If this is true, then we have

$$Pr[\lim_{n \rightarrow \infty} \tilde{\rho}_{\mathbf{R}}^{cp} = \mu_{\mathbf{X}} \cdot \mu_{\mathbf{Y}}] = 1.$$

Since

$$\begin{aligned} \mu_{\mathbf{X}} \cdot \mu_{\mathbf{Y}} &= \frac{1}{m_1} \sum_{j=1}^{m_1} \frac{1}{\prod_{k=2}^K m_k} \sum_{i_2=1}^{m_2} \cdots \sum_{i_K=1}^{m_K} \rho_{\mathbf{B}(j, i_2, \dots, i_K)} \\ &= \frac{1}{\prod_{k=1}^K m_k} \sum_{i_1=1}^{m_1} \cdots \sum_{i_K=1}^{m_K} \rho_{\mathbf{B}(i_1, i_2, \dots, i_K)} = \rho_{\mathbf{R}}, \end{aligned}$$

the strong consistency of $\tilde{\rho}_{\mathbf{R}}^{cp}$ is therefore proved.

We now show that $Pr[\lim_{n \rightarrow \infty} \mathbf{Y}_n = \mu_{\mathbf{Y}}] = 1$ by induction on K . When $K = 2$,

$$Y(n, j) = \frac{1}{n} \sum_{i_2=1}^n \rho_{\mathbf{B}(j, L_{2,i_2})}.$$

Note that, with a fixed j , the $\rho_{\mathbf{B}(j, L_{2,i_2})}$'s are i.i.d. random variables. So by the strong law of large numbers, $Pr[\lim_{n \rightarrow \infty} Y(n, j) = \rho_j] = 1$, and hence $Pr[\lim_{n \rightarrow \infty} \mathbf{Y}_n = \mu_{\mathbf{Y}}] = 1$. Now assume that the strong consistency holds for K relations. Consider the case of $K + 1$ relations:

$$Y(n, j) = \frac{1}{n^K} \sum_{i_2=1}^n \cdots \sum_{i_{K+1}=1}^n \rho_{\mathbf{B}(j, L_{2,i_2}, \dots, L_{K+1,i_{K+1}})}.$$

Define

$$Y(n, j, i_2) = \frac{1}{n^{K-1}} \sum_{i_3=1}^n \cdots \sum_{i_{K+1}=1}^n \rho_{\mathbf{B}(j, L_{2,i_2}, \dots, L_{K+1,i_{K+1}})}.$$

By induction hypothesis,

$$Pr[\lim_{n \rightarrow \infty} Y(n, j, i_2) = \rho_{j, i_2}] = 1,$$

where

$$\rho_{j, i_2} = \frac{1}{\prod_{k=3}^{K+1} m_k} \sum_{i_3=1}^{m_3} \cdots \sum_{i_{K+1}=1}^{m_{K+1}} \rho_{\mathbf{B}(j, i_2, \dots, i_{K+1})}.$$

Since

$$Y(n, j) = \frac{1}{n} \sum_{i_2=1}^n Y(n, j, i_2),$$

and for fixed j , $Y(n, j, i_2)$ are i.i.d. random variables with (common) mean

$$\begin{aligned} E[Y(n, j, i_2)] &= \frac{1}{m_2} \sum_{i_2=1}^{m_2} \rho_{j, i_2} \\ &= \frac{1}{\prod_{k=2}^{K+1} m_k} \sum_{i_2=1}^{m_2} \cdots \sum_{i_{K+1}=1}^{m_{K+1}} \rho_{\mathbf{B}(j, i_2, \dots, i_{K+1})} \\ &= \rho_j, \end{aligned}$$

we hence have

$$Pr\left[\lim_{n \rightarrow \infty} Y(n, j) = \rho_j\right] = 1,$$

by applying the strong law of large numbers again. This completes the proof of the lemma. \blacksquare

Therefore, Theorem 1 still holds in the case of cross-product sampling. It is also shown that cross-product sampling is always superior to independent sampling because of its lower sample variance (see Theorem 2 of [14]). Therefore, our cardinality refinement algorithm discussed next is based on cross-product sampling instead of independent sampling.

E. The Cardinality Refinement Algorithm

There are several considerations when designing our refinement algorithm based on the sampling-based estimator.

First, the estimator needs to access disk to take samples. However, since samples should be randomly taken, this means significant random reads may be required during the sampling phase, which may be too costly in practice. To overcome this issue, as has been suggested in previous applications of sampling in DBMS (e.g., [23]), we take samples offline and store them as materialized views in the database. We found in our experiments that the number of samples required is quite small and therefore can be cached in memory during runtime.

Second, the estimator we discussed so far focuses on estimating the selectivity (or equivalently, cardinality) for a single operator. However, in practice, a query plan may contain more than one operator, and for our purpose of refining the cost estimates of this plan, we need to estimate the cardinality for each operator. Another good property of the estimator is that, for a query plan with a fixed join order, which is always the case when refinement is performed, we can estimate all selection and join operators in the plan simultaneously. Consider, for example, a three-way join query $q = R_1 \bowtie R_2 \bowtie R_3$. We need to estimate the cardinality for both $q' = R_1 \bowtie R_2$ and q . However, after we are done with q' , we can estimate for q by directly evaluating $q' \bowtie R_3$. This means, we can estimate the cardinality for each operator by simply invoking the query plan q over the sample relations and then apply the estimator to each operator (see Theorem 2 below).

Theorem 2: Let $q = \sigma_F(R_1 \times \cdots \times R_K)$ be an arbitrary query with only selections and joins. For every subquery $q_i = \sigma_{F_i}(R_1 \times \cdots \times R_i)$ ($1 \leq i \leq K$, and F_i is the selection condition only involving R_1, \dots, R_i), $E[\tilde{\rho}_{q_i}] = \rho_{q_i}$ and $Pr\left[\lim_{n \rightarrow \infty} \tilde{\rho}_{q_i} = \rho_{q_i}\right] = 1$.

Proof: The proof is simply applying Theorem 1 to each q_i . \blacksquare

Third, while the estimator discussed above is both unbiased and strongly consistent, it only works for queries involving selections and joins. In practice, SQL queries can contain additional operators. A particularly common class of such operators we encountered in TPC-H queries is *aggregates*, for which we need to estimate the number of distinct values in the input relation. Aggregates basically *collapse* the underlying data distribution, so the estimator cannot work for queries containing aggregates. As a result, we can only apply the sampling-based estimator to the part of the query plan that does not involve aggregates. For aggregate operators, we simply rely on PostgreSQL's models for estimating output cardinalities. However, note that, since the refinement phase may change the input estimates for the aggregate, the output estimates for the aggregate may change as well. We observed in our experimental evaluation (see Section V) that the current approach already leads to promising prediction of execution time in practice. We leave the problem of further integrating state-of-the-art estimators (e.g., the GEE estimator in [6]) for estimating the number of distinct values as future work.

Our cardinality refinement algorithm is illustrated in Algorithm 1. For the input query q , we first call the optimizer to obtain its query plan P_q (line 30). We then modify P_q by replacing the relations it touches with the corresponding sample relations (i.e., materialized views), and run P_q over the sample relations (line 31 to 34). After that, we call the procedure *RecomputeCardinality* to refine the cardinality estimation for each operator in P_q (line 36).

The procedure *RecomputeCardinality* (line 11 to 27) works as follows. It first invokes *EstimateCardinality* on the child operators (if any) of the current operator O (line 12 to 17). It then checks the flag *HasAgg*, which indicates whether O has any *descendant* operator that is an aggregate. If the flag is set, then it simply calls the optimizer's own model to do cardinality estimation for O (line 18 to 19), since our estimator refinement cannot be applied in this case, as discussed above. If, on the other hand, the flag is not set, then it further checks whether O is itself an aggregate. If so, it again calls the optimizer's cardinality estimation model for O , and sets the flag *HasAgg* (line 21 to 23). If not, it invokes *EstimateCardinality* to estimate the cardinality of O (line 25). Note that due to the order that *EstimateCardinality* is invoked, the estimator is applied to each operator in a bottom-up manner. This guarantees that the input cardinality of any operator will be estimated after its child operators (if any). While this is not necessary if we only need to refine the cardinalities, it is necessary since we need to further estimate based on the cardinality (for example, the number of pages accessed), which enforces the same bottom-up ordering here since the cost of an operator covers the cost of its child operators as well (recall Example 3 and 4).

The procedure *EstimateCardinality* (line 3 to 9) implements the estimator. *GetSubPlan* returns the subtree P_O of

Algorithm 1: Cardinality Refinement

Input: q , a SQL query
Output: P_q , query plan of q with refined cardinalities

- 1 $HasAgg \leftarrow False$;
- 2
- 3 **EstimateCardinality**(O):
- 4 $P_O \leftarrow GetSubPlan(O)$;
- 5 $N_s \leftarrow \prod_{R^s \in SampleRelations(P_O)} |R^s|$;
- 6 $E_s \leftarrow CardinalityBySampling(O)$;
- 7 $N_O \leftarrow \prod_{R_O \in Relations(P_O)} |R_O|$;
- 8 $E_O \leftarrow N_O \cdot \frac{E_s}{N_s}$;
- 9 Treat E_O as the cardinality estimate for O ;
- 10
- 11 **RecomputeCardinality**(O):
- 12 **if** O has left child O_{lc} **then**
- 13 | $EstimateCardinality(O_{lc})$;
- 14 **end**
- 15 **if** O has right child O_{rc} **then**
- 16 | $EstimateCardinality(O_{rc})$;
- 17 **end**
- 18 **if** $HasAgg$ **then**
- 19 | Use optimizer’s model to estimate for O ;
- 20 **else**
- 21 | **if** O is aggregate **then**
- 22 | | Use optimizer’s model to estimate for O ;
- 23 | | $HasAgg \leftarrow True$;
- 24 | **else**
- 25 | | $EstimateCardinality(O)$;
- 26 | **end**
- 27 **end**
- 28
- 29 **Main:**
- 30 $P_q \leftarrow GetPlanFromOptimizer(q)$;
- 31 **foreach** $R \in Relations(P_q)$ **do**
- 32 | Replace R with its sample relation R^s ;
- 33 **end**
- 34 Run the plan P_q over the sample relations;
- 35 $O \leftarrow GetRootOperator(P_q)$;
- 36 $RecomputeCardinality(O)$;
- 37 **return** P_q ;

the query plan with the current operator O as the root. N_s is the product of the cardinalities of the sample relations involved in P_O , and E_s is the exact output cardinality of O when the plan is evaluated over sample relations. Therefore, the estimated selectivity is $\frac{E_s}{N_s}$, and the output cardinality of O over the original relations is then $N_O \cdot \frac{E_s}{N_s}$, where N_O is the product of the cardinalities of the original input relations. In Theorem 3, we further show that *EstimateCardinality* is a particular implementation of the estimator conforming to the semantics of cross-product sampling, with a special *tuple-level* partitioning scheme, where each block contains only a single tuple of the relation.

Theorem 3: The procedure *EstimateCardinality* estimates the cardinality of the operator O according to the semantics of cross-product sampling.

Proof: Suppose that the query $q = \sigma_F(R_1 \times \dots \times R_K)$. Consider the following *tuple-level* partitioning scheme: for each relation R_k ($1 \leq k \leq K$), treat each tuple of R_k a single partition. In this case, $m_k = |R_k|$ and $N_k = 1$. Since

$$\begin{aligned} E_s &= |\sigma_F(R_1^s \times \dots \times R_K^s)| \\ &= \sum_{i_1=1}^{|R_1^s|} \dots \sum_{i_K=1}^{|R_K^s|} |\sigma_F(\{t_{1,i_1}\} \times \dots \times \{t_{K,i_K}\})|, \end{aligned}$$

where t_{k,i_k} is the tuple (i.e., block) from R_k taken in the i_k -th sampling step, and $N_s = \prod_{k=1}^K |R_k^s|$, we estimate

$$\tilde{\rho}_q = \frac{E_s}{N_s} = \frac{\sum_{i_1=1}^{|R_1^s|} \dots \sum_{i_K=1}^{|R_K^s|} |\sigma_F(\{t_{1,i_1}\} \times \dots \times \{t_{K,i_K}\})|}{\prod_{k=1}^K |R_k^s|}.$$

This has the same semantics as the cross-product estimator defined in Equation (2), except that here the $|R_k^s|$ may be different for different R_k . It is straightforward to extend Equation (2) to the case where different relations have different sampling steps (i.e., different n ’s), and Lemma 3 (and hence Theorem 1) still holds in this case. ■

V. EVALUATION

In this section, we describe our experimental settings and report our results.

A. Experimental Settings

We implemented Algorithm 1 inside PostgreSQL 9.0.4 by modifying the query optimizer. In addition, we added instrumentation code to the optimizer to collect the input cardinalities for each operator. Our software setup was PostgreSQL on Linux Kernel 2.6.18, and we tested our method on both TPC-H 1GB and 10GB databases.

Our experiments were conducted on two different hardware configurations:

- *PC1*: configured with a 1-core 2.27GHz Intel CPU and 2GB memory;
- *PC2*: configured with an 8-core 2.40GHz Intel CPU and 16GB memory.

We randomly drew 10 queries from each of the 21 query templates², and we ran each query 5 times. Our error metric is computed based on the mean execution time of the queries. We cleared both the filesystem and DB buffers between each run of each query.

Since the original TPC-H database generator uses uniform distributions, to test the robustness of different approaches on different data distributions, we also used a skewed TPC-H database generator [1]. This database generator populates a TPC-H database using a Zipf distribution. This distribution has a parameter z that controls the degree of skewness. $z = 0$

²We excluded the template Q15 because it creates a view before the query runs, which is not supported by our current implementation of Algorithm 1.

generates a uniform distribution, and as z increases, the data becomes more and more skewed. We created skewed databases generated using $z = 1$.

B. Calibrating Cost Units

We use the approach described in Section III to generate calibration queries. The calibrated values for the 5 PostgreSQL optimizer parameters on PC1 and PC2 are shown in Table I and II, respectively. Except for *rand.page.cost*, the cost units show very small variance when profiled under different relations.

Calibrating the *rand.page.cost* is more difficult. As discussed in Section III-B, achieving purely random reads in a query appears difficult in practice. In addition, the number of random pages accessed as estimated by optimizer is based on statistics about correlations between the order of keys stored in the unclustered index and their actual order in the corresponding data file. Therefore, there is some inherent uncertainty in this estimation. It is interesting future work to see whether the *rand.page.cost* could be calibrated more accurately with different methods than the one described in this paper.

Optimizer Parameter	Calibrated $\mu \pm \sigma$ (ms)	Default
<i>seq.page.cost</i>	$5.53\text{e-}2 \pm 3.09\text{e-}3$	1.0
<i>rand.page.cost</i>	$6.50\text{e-}2 \pm 2.32\text{e-}2$	4.0
<i>cpu.tuple.cost</i>	$1.67\text{e-}4 \pm 5.83\text{e-}6$	0.01
<i>cpu.index.tuple.cost</i>	$3.41\text{e-}5 \pm 2.30\text{e-}5$	0.005
<i>cpu.operator.cost</i>	$1.12\text{e-}4 \pm 1.30\text{e-}6$	0.0025

TABLE I

ACTUAL VALUES OF POSTGRESQL OPTIMIZER PARAMETERS ON PC1

Optimizer Parameter	Calibrated $\mu \pm \sigma$ (ms)	Default
<i>seq.page.cost</i>	$5.03\text{e-}2 \pm 3.82\text{e-}3$	1.0
<i>rand.page.cost</i>	$4.89\text{e-}1 \pm 7.44\text{e-}2$	4.0
<i>cpu.tuple.cost</i>	$1.41\text{e-}4 \pm 1.35\text{e-}5$	0.01
<i>cpu.index.tuple.cost</i>	$3.34\text{e-}5 \pm 3.85\text{e-}5$	0.005
<i>cpu.operator.cost</i>	$7.10\text{e-}5 \pm 1.52\text{e-}5$	0.0025

TABLE II

ACTUAL VALUES OF POSTGRESQL OPTIMIZER PARAMETERS ON PC2

Note that the default settings of the parameters fail to accurately reflect the actual relative magnitudes. For example, on PC1, the ratio of calibrated *cpu.tuple.cost* to *seq.page.cost* is about 0.003 instead of 0.01.

Clearly, the overhead of this profiling stage depends on how many calibration queries we use. In our experiments on the TPC-H database, we used the 5 largest relations as the R in $\text{SELECT } * \text{ FROM } R$ and $\text{SELECT COUNT}(*) \text{ FROM } R$, respectively. For $\text{SELECT } * \text{ FROM } R \text{ WHERE } R.A < a$, we used the largest relation (*lineitem*), and generated 10 queries where the predicate $R.A < a$ had different selectivities in each. Under this setting, the profiling stage usually finishes in less than an hour, which is substantially less than the long training stage of machine-learning-based approaches.

Moreover, our profiling stage was conducted on top of the uniform TPC-H database. Note that we do not need to run it again for the skewed TPC-H database. This is because the

values of the cost units only depend on the specific hardware configuration. After the cost units are calibrated, they can be used as long as the hardware configuration does not change. On the other hand, machine-learning-based approaches usually need to collect new training data and rebuild the predictive model if the underlying data distribution significantly changes.

C. Prediction Results

We evaluated the accuracy of prediction in terms of the *mean relative error* (MRE), a metric used in [4]. MRE is defined as

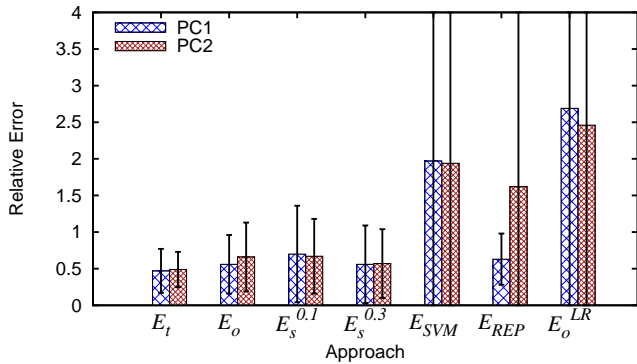
$$\frac{1}{M} \sum_{i=1}^M \frac{|T_i^{\text{pred}} - T_i^{\text{act}}|}{T_i^{\text{act}}},$$

where M is the number of testing queries, T_i^{pred} and T_i^{act} are the predicted and actual execution time of the testing query i , respectively.


We compare the prediction accuracy of our approach with several state-of-the-art machine-learning-based solutions: *Plan-level modeling with SVM* [4], *Plan-level modeling with REP trees* [28], and *Operator-level modeling with Multivariate Linear Regression (MLR)* [4]. We use the same set of features as described in [4]. We focus on the settings of the so-called *dynamic workload* in [4]. The idea of plan-level modeling was also tried in [11], and the authors chose to use Kernel Canonical Correlation Analysis (KCCA) [5] instead of SVM as the machine-learning approach. We do not compare our techniques with theirs, for it has been shown in [4] that both the plan-level and operator-level modeling approach of [4] are superior to the KCCA-based approach for dynamic workloads.

To generate a dynamic workload, we conducted the following “leave-one-template-out” experiment as in [4]. Among the N TPC-H query templates, we chose one template to generate the queries whose execution time is to be predicted, and the other $N-1$ templates were used to generate the training queries used by the machine learning methods to train their predictive models.

Figure 2 shows the results on the uniform (i.e., $z = 0$) 1GB TPC-H database. As presented in [4], operator-level modeling requires that the testing queries do not contain operators not used by the training queries. Since some TPC-H templates include specific operators not found in the other templates (e.g., *hash-semi-join*), we excluded these templates from our experiments. The same argument applies to TPC-H templates containing PostgreSQL-specific structures (i.e., *INITPLAN* and *SUBQUERY*). The authors of [4] also excluded these queries for the same reason. However, we note here that this is a problem due to the particular choice of the workload and database system, not due to the operator-level modeling itself. If TPC-H were a more varied workload, we would not have this restriction. For example, if it had multiple queries that used the hash-semi-join operator, we could have incorporated queries with that operator in our experiments. This leaves 11 TPC-H templates participating in the dynamic workload experiment when operator-level modeling is leveraged (see Figure 2(b)).



(a) Without E_{Oper} , 21 templates



(b) With E_{Oper} , 11 templates

Fig. 2. Uniform TPC-H 1GB database

In Figure 2, the x -axis represents the approaches we tested in the experiments, and the y -axis shows the average error and the standard deviation (shown with the error bars) over the TPC-H templates. The detailed results for each query template can be found from Table III to Table VI in Appendix A. Here, E_t is the prediction error of our approach when the *true* cardinalities are used (the true cardinalities are measured in an artificial “pre-running” of the query — we present this number to provide insight into what could be achievable if we were able to get perfect cost estimates). E_o is the prediction error of our approach when the cardinalities from the optimizer are used without refinement. E_s^f is the prediction error of our approach when the cardinalities are estimated via sampling, where f is the sampling ratio (e.g., $f = 0.1$ means we take a 10% sample from each underlying table). In our experiments, we tested sampling ratios $f = 0.05, 0.1, 0.2, 0.3, 0.4$. Due to space limitations, we only present the results of $f = 0.1$ and $f = 0.3$. E_{SVM} , E_{REP} , and E_{Oper} are the prediction errors of the three machine-learning based approaches, i.e., plan-level modeling with SVM, plan-level modeling with REP, and operator-level modeling, respectively. Finally, as a baseline, E_o^{LR} is the prediction error of mapping the original cost estimates from the optimizer to the execution time via linear regression, as was done in previous work.

We have several observations. First, in the case of uniform data, the cost models with properly tuned c ’s already work well (the E_o in Figure 2 is close to E_t). Sampling does not help




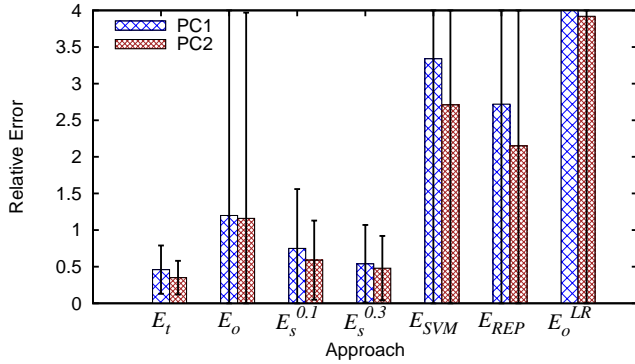
Fig. 3. Queries projected on the 3 dominating principal components

much in improving the prediction accuracy. This is reasonable, because the assumptions like uniformity and independence leveraged by the optimizer for cardinality estimation usually hold in this case.


Second, the performance of machine-learning based approaches is not consistent. For some queries, their predictions are good. However, for the other queries, their predictions are far away from the true values. This can be observed by noticing the big error bars in the figures. As an example, on PC1, the E_{SVM} in Figure 2(a) varies between 0.03 for Q7 and 12.16 for Q17 (see Table III in Appendix A). One possible reason for this is: most machine learning methods assume that the testing queries should be *similar* to the queries used in training the model. More specifically, the feature vectors of the testing queries should be close to those feature vectors of the training queries, in terms of the distance in the feature space. Unfortunately, this assumption is not valid for dynamic workloads.

To see this, we further apply Principal Component Analysis (PCA) [15] on the query features (47 features in total) and project the queries onto the subspace spanned by the three most dominating principal components. It is well known that these dominating components reveal the *major* directions in the feature space. While it is true that more principal components are better, we restrict ourselves to 3-dimensional space for the purpose of visualization.

Figure 3 shows the queries in the projected space, where each combination of color and shape represents one query template. From the figure we can see that the templates can be grouped into several clusters. About half of the templates fall into the rightmost cluster, and each of the remaining templates usually forms a singleton cluster. The distances between clusters are quite big. Note that PCA will not increase the distances between feature vectors after the projection, which means the distances between feature vectors in the original 47-dimensional space can only be the same or even bigger. This suggests that there is little similarity among the TPC-H templates within different clusters. Therefore, if we test the queries from a template within a singleton cluster, by



(a) Without E_{Oper} , 21 templates



(b) With E_{Oper} , 11 templates


Fig. 4. Skewed TPC-H 1GB database

using the model trained with the other templates, then there is little hope for us to observe good predictions.

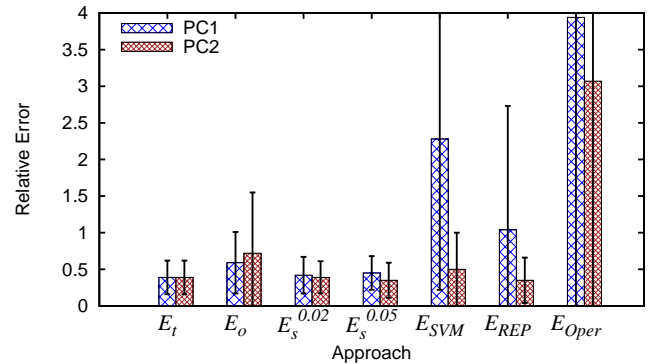
Third, machine-learning approaches are sensitive to the set of queries used in training. By comparing Table III and IV in Appendix A, we can see that the prediction errors for some queries fluctuate dramatically when different sets of training queries are used³. For instance, E_{REP} for Q8 is 0.44 when 20 templates are used in training (as in Figure 2(a)), but it will increase to 2.87 when only 10 templates are used (as in Figure 2(b)). Picking a set of proper training queries hence is critical in practice when using machine-learning approaches. However, it seems quite difficult in the environment when workload is not known in advance.

Figure 4 further presents the results on the skewed TPC-H 1GB database, and the details for each query template are presented from Table VII to Table X in Appendix B. As expected, when data becomes skewed, the cardinality estimates from the optimizer become inaccurate, and hence the predictive power is weakened. However, by leveraging the sampling-based cardinality correction, the prediction accuracy is improved. Moreover, more samples usually means better prediction accuracy, as long as the overhead on sampling is acceptable (see Section V-D). We note that the sampling overhead can be up to 20% on this data set. As we will see, this can be viewed as a problem that arises on small

³Recall that, to be fair, we only use 10 templates in training when comparing with operator-level modeling (as in Figure 2(b)).



(a) Without E_{Oper} , 21 templates



(b) With E_{Oper} , 11 templates


Fig. 5. Uniform TPC-H 10GB database

data sets, as the overhead due to sampling for the 10GB data set is much lower. On the other hand, the performance of machine-learning based approaches becomes even worse. This is perhaps partially because of the worse distortion of the assumption that training and testing queries should be similar.


Similar results on the TPC-H 10GB database are observed in our experiments, as shown in Figure 5 and Figure 6. We also include details of the errors on each template from Table XI to Table XIV in Appendix C for uniform data, and from Table XV to Table XVIII in Appendix D for skewed data. Here, to make the overall experiment time controllable, as done in [4], we kill the query if it runs longer than an hour. This leaves us with 18 templates participating in the evaluation. We tested sampling ratios $f = 0.01, 0.02, 0.05, 0.1$, and present the results of $f = 0.02$ and $f = 0.05$, for space constraints. Note that, while the database size scales up by a factor of 10, the required absolute number of samples to achieve predictions close to the ideal case (compare $E_s^{0.05}$ and E_t in the figures) remains almost the same. We need $0.05 \times 10GB = 0.5GB$ samples here, while we need $0.3 \times 1GB = 0.3GB$ samples in the case of 1GB database. Therefore, the additional overhead of taking samples becomes ignorable when the database is larger (see Section V-D).

D. Overhead of Sampling

Figure 7 shows the additional runtime overhead due to sampling for the 1GB TPC-H database. Here r_s^f is defined



(a) Without E_{Oper} , 21 templates




(b) With E_{Oper} , 11 templates

Fig. 6. Skewed TPC-H 10GB database


as $r_s^f = T_s/T$, where T_s and T are the time to run the queries over the sample tables and original tables, respectively, and f is the sampling ratio as before. For each sampling ratio, we report the average r_s^f as well as the standard deviation (shown with the error bars) over the participating TPC-H templates. We refer the readers to Appendix E (from Table XIX to Table XXII) for the details of the overhead for different query templates.

We can see that for the sampling ratio $f = 0.3$, which allows us to achieve close prediction accuracy to what if the true cardinalities were used on both the uniform and skewed data, the average additional runtime overhead is around 20% of the actual execution time of the query. Note that for query optimization, 20% is prohibitively high. For example, it means that we can only consider $1/20\% = 5$ plans during optimization before the estimation cost dominates the query execution time, since sampling should be invoked for every query plan considered. But for our purposes, where we are trying to estimate the running time of a single query plan, this amount of overhead may be acceptable. Perhaps more importantly, this overhead drops dramatically when we move to the 10GB data set.

Figure 8 further presents the results over 10GB TPC-H database. Again, the details for each query template are presented in Appendix F, from Table XXIII to Table XXVI. It confirms that the additional overhead introduced by sampling is even smaller, compared with the overhead of running the




(a) On uniform data




(b) On skewed data

Fig. 7. Additional overhead of sampling on TPC-H 1GB database



(a) On uniform data



(b) On skewed data

Fig. 8. Additional overhead of sampling on TPC-H 10GB database

original query. For the case where good prediction can be achieved (i.e., $f = 0.05$, see Figure 5 and Figure 6), the additional overhead is below 4% on average. This demonstrates the practicality of incorporating sampling for the purpose of query time prediction.

VI. RELATED WORK

Query optimizers have built-in cost models that provide cardinality/cost estimates for a given query. There is a lot of previous work on this topic, including methods based on sampling (e.g., [16], [20]), methods based on histograms (e.g., [17]), and methods based on machine learning (e.g., [12], [25]). However, the purpose of these estimates is to help the optimizer pick a relatively good plan, not to predict the *actual* execution time of the query. Therefore, these estimates need not to be very accurate as long as the optimizer can leverage them to distinguish good plans from bad ones. As shown in [4], [11], without proper calibration, directly leveraging these cost estimates cannot provide good predictions of execution time.

Previous work has explored the issue of calibrating optimizer parameters, for different purposes such as query optimization in heterogeneous DBMS [9], DB resource virtualization [24], and storage type selection [29]. To the best of our knowledge, we are not aware of any work that tries to predict the execution time of SQL queries based on calibrating the cost models of query optimizers.

Another related research direction is *query progress indicators* [7], [18], [19], [21]. The task of a progress indicator is to dynamically monitor the percentage of work that has been done so far for the query. The key difference from query execution time prediction is that progress indicators usually rely on *runtime* statistics obtained *during* the actual execution of the query, which are not available if the prediction is restricted to be made before query execution. A query running time predictor could be useful in providing the very first estimate for a progress indicator (one used before the query starts executing).

Quite surprisingly, the problem of predicting *actual* execution time of a query has been specifically addressed only recently [11]. Existing work [3], [4], [10], [11], [26] usually employs predictive frameworks based on statistical machine learning techniques.

In [11], each query is represented as a set of features containing an instance count and cardinality sum for each possible operator. Kernel Canonical Correlation Analysis (KCCA) [5] modeling techniques are then used to map the queries from their feature space onto their performance space. One main limitation of this approach is that its prediction is based on taking the average of the k (usually 3) nearest neighbors in the training set, which means that the prediction can never exceed the longest execution time observed during training stage. Hence, when the query to be predicted takes significantly longer time than all the training queries observed, the model is incapable of giving reasonable predictions.

In [4], a similar idea of using features extracted from the entire plan to represent a query is leveraged, and the authors

propose to use SVM instead of KCCA. However, the SVM approach still suffers from the same generalization problem. To alleviate this, the authors further apply this idea at the operator-level. But from the reported experimental results (both in [4] and Section V of this paper), it seems that operator-level modeling is still quite vulnerable to workload changes. Our approach in this paper avoids this generalization problem, for it does not rely on any particular training queries.

In [10], the authors study the problem of predicting the execution time when the query is concurrently running with the other queries, with a linear multivariate regression model to capture the interaction between queries. Similar problems in admission control and query scheduling are also studied in [26] and [3], respectively. One key limitation of this line of work is that they all assume a *closed-world* workload scenario, where all possible queries are needed to be known in ahead, which is hardly to be the case in practice. Although this paper does not address the prediction problem in the presence of concurrent query execution, it is interesting to see how to extend the techniques here to provide alternative solutions to this challenge.

VII. CONCLUSION

In this paper, we studied the problem of leveraging optimizer's cost models to predict query execution time. We show that, after proper calibration, the current cost models used by query optimizers can be more effective for predicting query execution time than reported by previous work.

Of course, it is possible that a new machine learning technique, perhaps with improved feature selection, will outperform the techniques presented here. On the other hand, further improvements are also possible in optimizer-based running time prediction. Perhaps the most interesting aspect of this work is the basic question it raises: should query running time prediction treat the DBMS as a black box (the machine learning approach), or should we exploit the fact that we actually know exactly what is going on inside the box (the optimizer based approach)? We regard this paper as an argument that the latter approach shows promise, and expect that exploring the capabilities of the two very different approaches will be fertile ground for future research.

REFERENCES

- [1] Skewed tpc-h data generator. <ftp://ftp.research.microsoft.com/users/viveknar/TPCDSkew/>.
- [2] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [3] M. Ahmad, A. Aboulnaga, S. Babu, and K. Munagala. Interaction-aware scheduling of report-generation workloads. *The VLDB Journal*, 20:589–615, 2011.
- [4] M. Akdere, U. Çetintemel, M. Riondato, E. Upfal, and S. Zdonik. Learning-based query performance modeling and prediction. In *ICDE*, 2012.
- [5] F. R. Bach and M. I. Jordan. Kernel independent component analysis. *Journal of Machine Learning Research*, 3:1–48, 2002.
- [6] M. Charikar, S. Chaudhuri, R. Motwani, and V. R. Narasayya. Towards estimation error guarantees for distinct values. In *PODS*, pages 268–279, 2000.
- [7] S. Chaudhuri, V. R. Narasayya, and R. Ramamurthy. Estimating progress of execution for SQL queries. In *SIGMOD*, 2004.

- [8] Y. Chi, H. J. Moon, and H. Hacigümüş. iCBS: Incremental costbased scheduling under piecewise linear slas. *PVLDB*, 4(9):563–574, 2011.
- [9] W. Du, R. Krishnamurthy, and M.-C. Shan. Query optimization in a heterogeneous dbms. In *VLDB*, pages 277–291, 1992.
- [10] J. Duggan, U. Çetintemel, O. Papaemmanouil, and E. Upfal. Performance prediction for concurrent database workloads. In *SIGMOD*, 2011.
- [11] A. Ganapathi, H. A. Kuno, U. Dayal, J. L. Wiener, A. Fox, M. I. Jordan, and D. A. Patterson. Predicting multiple metrics for queries: Better decisions enabled by machine learning. In *ICDE*, 2009.
- [12] L. Getoor, B. Taskar, and D. Koller. Selectivity estimation using probabilistic models. In *SIGMOD*, 2001.
- [13] S. Guirguis, M. A. Sharaf, P. K. Chrysanthos, A. Labrinidis, and K. Pruks. Adaptive scheduling of web transactions. In *ICDE*, 2009.
- [14] P. J. Haas, J. F. Naughton, S. Seshadri, and A. N. Swami. Selectivity and cost estimation for joins based on random sampling. *J. Comput. Syst. Sci.*, 52(3):550–569, 1996.
- [15] T. Hastie, R. Tibshirani, and J. H. Friedman. *The Elements of Statistical Learning*. Springer, 2003.
- [16] W.-C. Hou and G. Ozsoyoglu. Statistical estimators for aggregate relational algebra queries. *ACM Trans. Database Syst.*, 16:600–654, 1991.
- [17] Y. E. Ioannidis. The history of histograms (abridged). In *VLDB*, pages 19–30, 2003.
- [18] A. C. König, B. Ding, S. Chaudhuri, and V. R. Narasayya. A statistical approach towards robust progress estimation. *PVLDB*, 5(4):382–393, 2011.
- [19] J. Li, R. V. Nehme, and J. F. Naughton. GSLPI: A cost-based query progress indicator. In *ICDE*, pages 678–689, 2012.
- [20] R. J. Lipton, J. F. Naughton, and D. A. Schneider. Practical selectivity estimation through adaptive sampling. In *SIGMOD*, 1990.
- [21] G. Luo, J. F. Naughton, C. J. Ellmann, and M. Watzke. Toward a progress indicator for database queries. In *SIGMOD*, 2004.
- [22] C. Mishra and N. Koudas. The design of a query monitoring system. *ACM Trans. Database Syst.*, 34(1), 2009.
- [23] R. Ramamurthy and D. J. DeWitt. Buffer-pool aware query optimization. In *CIDR*, pages 250–261, 2005.
- [24] A. A. Soror, U. F. Minhas, A. Aboulmaga, K. Salem, P. Kokosielis, and S. Kamath. Automatic virtual machine configuration for database workloads. In *SIGMOD*, 2008.
- [25] M. Stillger, G. M. Lohman, V. Markl, and M. Kandil. LEO - DB2’s learning optimizer. In *VLDB*, 2001.
- [26] S. Tozer, T. Brecht, and A. Aboulmaga. Q-Cop: Avoiding bad query mixes to minimize client timeouts under heavy loads. In *ICDE*, 2010.
- [27] T. J. Wasserman, P. Martin, D. B. Skillicorn, and H. Rizvi. Developing a characterization of business intelligence workloads for sizing new database systems. In *DOLAP*, 2004.
- [28] P. Xiong, Y. Chi, S. Zhu, J. Tatemura, C. Pu, and H. Hacigümüş. ActiveSLA: a profit-oriented admission control framework for database-as-a-service providers. In *SOCC*, 2011.
- [29] N. Zhang, J. Tatemura, J. M. Patel, and H. Hacigümüş. Towards cost-effective storage provisioning for DBMSs. *PVLDB*, 5(4):274–285, 2011.

APPENDIX

In this section, we present the complete experimental results over TPC-H templates.

A. Prediction Errors on Uniform TPC-H 1GB Database

Table III and IV present the prediction errors for the participating approaches over the TPC-H templates on the uniform 1GB database deployed on PC1. Errors above 1.0 are highlighted. In Table III, the errors E_{SVM} and E_{REP} are obtained by using 20 templates in training. In Table IV, the errors E_{SVM} , E_{REP} , and E_{Oper} are based on using 10 templates in training, for fairness of comparison, since operator-level modeling requires testing queries do not contain operators that are not used by training queries, as described in Section V-C. We observe that the prediction errors of machine-learning approaches are sensitive to the training queries in use (e.g. E_{REP} increases from 0.44 in Table III to 2.87 in

Q	E_t	E_o	$E_s^{0.1}$	$E_s^{0.3}$	E_{SVM}	E_{REP}	E_o^{LR}
1	0.17	0.17	0.17	0.17	0.74	0.05	0.57
2	0.32	0.35	1.15	0.19	1.01	1.26	1.33
3	0.37	0.94	0.12	0.09	0.53	0.78	0.29
4	0.57	0.56	0.55	0.57	2.19	0.03	2.97
5	0.24	0.23	0.23	0.24	0.79	1.18	0.64
6	0.71	0.71	0.71	0.71	0.07	0.49	0.17
7	0.07	0.04	0.09	0.09	0.03	0.52	0.54
8	0.37	0.35	0.40	0.37	2.10	0.44	2.25
9	0.14	0.96	0.28	0.12	0.82	0.97	0.73
10	0.43	0.22	0.39	0.34	0.18	0.56	0.29
11	0.54	0.21	0.20	0.21	6.56	0.35	5.08
12	0.65	0.67	0.65	0.65	0.20	0.14	0.10
13	0.83	0.83	0.77	0.82	0.72	0.63	0.55
14	0.88	0.88	0.88	0.87	0.71	0.90	0.54
16	0.71	0.83	0.73	0.72	0.46	0.66	0.41
17	0.02	0.07	0.21	0.09	12.16	1.13	25.96
18	0.14	1.59	1.47	1.78	0.73	0.45	0.08
19	1.03	1.03	1.03	1.04	0.20	0.51	0.42
20	0.64	0.65	0.64	0.64	4.66	0.45	5.71
21	0.88	0.13	3.06	1.96	0.33	0.67	0.40
22	0.13	0.28	0.90	0.04	6.20	0.99	7.41
Avg	0.47	0.56	0.70	0.56	1.97	0.63	2.69

TABLE III

ON PC1 (w/o E_{Oper} , 21 TEMPLATES)

Q	E_t	E_o	$E_s^{0.1}$	$E_s^{0.3}$	E_{SVM}	E_{REP}	E_{Oper}
1	0.17	0.17	0.17	0.17	0.73	0.64	0.49
3	0.37	0.94	0.12	0.09	0.52	0.79	1.45
5	0.24	0.23	0.23	0.24	1.63	1.24	1.87
6	0.71	0.71	0.71	0.71	0.16	0.46	0.42
7	0.07	0.04	0.09	0.09	0.20	4.05	2.18
8	0.37	0.35	0.40	0.37	3.55	2.87	3.05
9	0.14	0.96	0.28	0.12	0.82	0.82	0.83
12	0.65	0.67	0.65	0.65	0.02	0.55	0.37
14	0.88	0.88	0.88	0.87	1.51	0.90	0.86
19	1.03	1.03	1.03	1.04	0.23	0.54	1.34
Avg	0.46	0.60	0.46	0.43	0.94	1.29	1.28

TABLE IV

ON PC1 (w/ E_{Oper} , 10 TEMPLATES)

Q	E_t	E_o	$E_s^{0.1}$	$E_s^{0.3}$	E_{SVM}	E_{REP}	E_o^{LR}
1	0.25	0.25	0.24	0.25	0.70	0.57	0.59
2	0.17	0.12	1.87	0.49	2.39	3.75	2.88
3	0.26	1.21	0.01	0.02	0.29	1.34	0.09
4	0.60	0.59	0.58	0.59	2.56	0.42	3.30
5	0.34	0.33	0.33	0.34	0.29	1.17	0.08
6	0.75	0.74	0.75	0.75	0.03	0.73	0.13
7	0.57	0.56	0.56	0.56	0.43	0.18	0.22
8	0.56	0.55	0.56	0.56	1.59	0.66	2.27
9	0.32	0.72	0.26	0.32	0.81	0.78	0.73
10	0.50	0.39	0.48	0.46	0.49	0.25	0.26
11	0.55	0.25	0.27	0.25	7.08	2.10	7.65
12	0.67	0.67	0.67	0.67	0.25	1.44	0.06
13	0.71	0.71	0.84	0.74	0.34	0.08	0.13
14	0.85	0.85	0.85	0.85	1.53	1.57	1.94
16	0.75	0.85	0.76	0.76	0.40	0.76	0.33
17	0.29	0.24	0.16	0.30	8.47	12.20	16.64
18	0.12	2.32	2.15	2.30	0.58	1.35	0.17
19	0.95	0.95	0.93	0.93	0.02	0.04	0.32
20	0.65	0.65	0.66	0.67	5.11	2.99	6.08
21	0.22	0.54	0.57	0.17	0.68	0.77	0.50
22	0.23	0.36	0.62	0.10	6.74	0.99	7.30
Avg	0.49	0.66	0.67	0.57	1.94	1.62	2.46

TABLE V

ON PC2 (w/o E_{Oper} , 21 TEMPLATES)

Table IV). Also, not surprisingly, the predictions made by the

Q	E_t	E_o	$E_s^{0.1}$	$E_s^{0.3}$	E_{SVM}	E_{REP}	E_{Oper}
1	0.25	0.25	0.24	0.25	0.59	0.57	0.34
3	0.26	1.21	0.01	0.02	0.04	1.34	0.07
5	0.34	0.33	0.33	0.34	0.03	1.32	2.32
6	0.75	0.74	0.75	0.75	0.40	2.15	0.35
7	0.57	0.56	0.56	0.56	0.14	0.18	1.42
8	0.56	0.55	0.56	0.56	2.79	2.58	2.87
9	0.32	0.72	0.26	0.32	0.73	0.79	0.68
12	0.67	0.67	0.67	0.67	0.08	1.43	0.58
14	0.85	0.85	0.85	0.85	2.69	1.65	0.86
19	0.95	0.95	0.93	0.93	0.46	0.03	1.16
Avg	0.55	0.68	0.52	0.52	0.80	1.20	1.07

TABLE VI
ON PC2 (w/ E_{Oper} , 10 TEMPLATES)

optimizer without proper calibration (i.e., E_o^{LR}) are almost always the worst among all the approaches in our experiments. Table V and VI present similar results on PC2.

B. Prediction Errors on Skewed TPC-H 1GB Database

Q	E_t	E_o	$E_s^{0.1}$	$E_s^{0.3}$	E_{SVM}	E_{REP}	E_o^{LR}
1	0.17	0.17	0.17	0.17	0.77	0.68	0.63
2	0.97	0.93	3.65	1.91	13.81	0.73	17.98
3	0.61	0.34	0.43	0.11	0.57	1.26	0.31
4	0.13	0.13	0.11	0.12	6.34	0.53	8.23
5	0.33	0.21	0.33	0.13	1.73	0.70	2.59
6	0.28	0.28	0.28	0.28	1.27	0.27	1.96
7	0.03	0.31	0.06	0.03	0.30	0.37	0.58
8	0.39	0.36	0.39	0.11	0.04	3.08	0.17
9	0.26	0.92	0.49	0.26	0.79	0.84	0.65
10	0.43	0.26	0.36	0.22	0.37	0.24	0.78
11	0.38	0.41	0.31	0.33	14.24	34.37	18.23
12	0.26	0.24	0.27	0.26	0.81	1.03	1.34
13	0.82	0.82	0.76	0.81	0.75	0.88	0.59
14	0.76	0.77	0.76	0.70	3.02	3.60	4.49
16	0.76	0.86	0.77	0.76	0.58	0.88	0.31
17	0.36	15.04	1.03	0.74	5.15	1.32	6.93
18	0.23	1.25	1.15	1.35	0.81	0.55	0.68
19	1.26	1.28	1.28	1.14	0.63	2.44	0.99
20	0.24	0.26	0.25	0.24	10.77	0.49	13.54
21	0.91	0.07	1.88	1.46	0.23	2.23	0.07
22	0.08	0.24	1.03	0.09	7.19	0.65	9.35
Avg	0.46	1.20	0.75	0.54	3.34	2.72	4.31

TABLE VII
ON PC1 (w/o E_{Oper} , 21 TEMPLATES)

Q	E_t	E_o	$E_s^{0.1}$	$E_s^{0.3}$	E_{SVM}	E_{REP}	E_{Oper}
1	0.17	0.17	0.17	0.17	0.78	0.72	0.56
3	0.61	0.34	0.43	0.11	0.58	0.88	0.23
5	0.33	0.21	0.33	0.13	2.24	0.78	2.12
6	0.28	0.28	0.28	0.28	1.70	0.33	0.45
7	0.03	0.31	0.06	0.03	0.42	0.41	2.06
8	0.39	0.36	0.39	0.11	0.10	2.40	0.95
9	0.26	0.92	0.49	0.26	0.79	0.80	0.83
10	0.43	0.26	0.36	0.22	0.58	0.46	1.59
12	0.26	0.24	0.27	0.26	1.08	0.17	1.39
14	0.76	0.77	0.76	0.70	3.87	3.59	0.94
19	1.26	1.28	1.28	1.14	0.64	4.37	9.95
Avg	0.44	0.47	0.44	0.31	1.16	1.35	1.92

TABLE VIII
ON PC1 (w E_{Oper} , 11 TEMPLATES)

Table VII and VIII show the prediction errors over skewed TPC-H 1GB database on PC1. We can observe that the

Q	E_t	E_o	$E_s^{0.1}$	$E_s^{0.3}$	E_{SVM}	E_{REP}	E_o^{LR}
1	0.25	0.25	0.25	0.25	0.75	0.60	0.63
2	0.46	0.43	2.20	1.07	7.14	4.81	10.52
3	0.25	1.38	0.04	0.03	0.39	0.12	0.13
4	0.19	0.11	0.14	0.17	5.13	0.57	8.15
5	0.20	0.22	0.23	0.19	0.33	0.26	0.08
6	0.36	0.36	0.35	0.36	1.01	0.70	1.91
7	0.33	0.33	0.33	0.34	0.34	0.69	0.05
8	0.08	0.04	0.09	0.09	0.05	0.48	0.37
9	0.40	0.89	0.56	0.53	0.76	0.71	0.64
10	0.38	0.28	0.35	0.34	0.44	0.80	0.16
11	0.30	0.49	0.29	0.24	12.46	0.96	18.51
12	0.37	0.36	0.37	0.37	0.24	0.73	0.78
13	0.74	0.73	0.85	0.76	0.46	0.04	0.21
14	0.49	0.49	0.50	0.49	5.88	7.13	9.28
16	0.79	0.87	0.80	0.80	0.57	0.97	0.37
17	0.27	13.26	0.95	0.70	4.20	24.13	6.58
18	0.01	1.89	1.73	1.88	0.69	0.48	0.55
19	0.95	0.95	0.95	0.96	0.16	0.57	0.21
20	0.28	0.30	0.31	0.30	9.28	0.28	13.80
21	0.13	0.46	0.30	0.14	0.70	0.05	0.54
22	0.19	0.33	0.78	0.03	5.98	0.06	8.94
Avg	0.35	1.16	0.59	0.48	2.71	2.15	3.92

TABLE IX
ON PC2 (w/o E_{Oper} , 21 TEMPLATES)

Q	E_t	E_o	$E_s^{0.1}$	$E_s^{0.3}$	E_{SVM}	E_{REP}	E_{Oper}
1	0.25	0.25	0.25	0.25	0.69	0.69	0.47
3	0.25	1.38	0.04	0.03	0.22	0.07	0.38
5	0.20	0.22	0.23	0.19	0.15	0.13	2.53
6	0.36	0.36	0.35	0.36	1.84	4.04	0.21
7	0.33	0.33	0.33	0.34	0.16	1.57	2.12
8	0.08	0.04	0.09	0.09	0.21	0.40	1.28
9	0.40	0.89	0.56	0.53	0.70	0.61	0.74
10	0.38	0.28	0.35	0.34	0.31	0.07	18.42
12	0.37	0.36	0.37	0.37	0.76	0.37	1.73
14	0.49	0.49	0.50	0.49	8.64	2.10	0.59
19	0.95	0.95	0.95	0.96	0.08	2.49	0.49
Avg	0.37	0.50	0.37	0.36	1.25	1.14	2.63

TABLE X
ON PC2 (w E_{Oper} , 11 TEMPLATES)

prediction errors E_o of the optimizer with calibrated cost units increase when data becomes skewed. However, by using sampling to refine cardinality estimates, we are able to achieve comparable prediction accuracy as observed on the uniform data, even when the data becomes skewed. On the other hand, the predictions made by machine-learning based approaches are even worse. Similar results on PC2 are reported in Table IX and X.

C. Prediction Errors on Uniform TPC-H 10GB Database

Table XI to XIV show the results on uniform TPC-H 10GB database. Table XI and XII present the prediction errors of different approaches on PC1, while Table XIII and XIV present their errors on PC2, respectively.

We observe very similar results here, compared with the results on uniform TPC-H 1GB database, as shown in Appendix A. Note that we achieve this comparable performance without increasing the number of samples too much. We observe that $E_s^{0.05}$ is close to E_t for the 10GB database, while it is $E_s^{0.3}$ for 1GB database. This means, we need $0.05 \times 10GB = 0.5GB$ sample for 10GB database, while

Q	E_t	E_o	$E_s^{0.02}$	$E_s^{0.05}$	E_{SVM}	E_{REP}	E_o^{LR}
1	0.06	0.06	0.06	0.06	3.22	0.50	3.53
2	0.41	0.23	0.10	0.24	11.08	3.29	13.63
3	0.34	0.25	0.27	0.34	1.45	0.24	2.05
4	0.29	0.86	0.77	0.59	2.14	0.10	1.65
5	0.55	0.51	0.55	0.55	0.66	0.47	0.84
6	0.54	0.54	0.54	0.54	0.41	2.90	0.29
7	0.53	0.54	0.55	0.54	1.31	0.95	1.37
8	0.61	0.61	0.73	0.72	2.37	0.17	2.39
9	0.06	0.91	0.38	0.44	0.75	0.63	0.68
10	0.18	1.64	0.06	0.20	1.70	0.75	3.07
11	0.48	0.44	0.67	0.70	5.96	1.16	6.94
12	0.46	0.46	0.48	0.48	0.56	0.79	0.49
13	0.93	0.94	1.57	0.85	3.67	0.94	5.03
14	0.76	0.76	0.80	0.80	0.82	0.57	0.80
16	0.56	0.50	0.56	0.56	3.46	0.60	3.62
18	0.38	3.71	3.14	0.42	0.77	0.21	0.42
19	0.22	0.22	0.22	0.22	0.36	1.62	0.25
22	0.47	0.65	2.13	0.07	5.99	1.41	8.65
Avg	0.43	0.77	0.76	0.46	2.59	0.96	3.09

TABLE XI
ON PC1 (w/o E_{Oper} , 18 TEMPLATES)

Q	E_t	E_o	$E_s^{0.02}$	$E_s^{0.05}$	E_{SVM}	E_{REP}	E_{Oper}
1	0.06	0.06	0.06	0.06	6.33	5.83	6.33
3	0.34	0.25	0.27	0.34	3.33	0.45	0.27
5	0.55	0.51	0.55	0.55	1.72	0.42	2.05
6	0.54	0.54	0.54	0.54	0.20	0.08	1.92
7	0.53	0.54	0.55	0.54	2.77	0.06	6.46
8	0.61	0.61	0.73	0.72	4.75	0.48	4.07
9	0.06	0.91	0.38	0.44	0.75	0.74	0.88
10	0.18	1.64	0.06	0.20	3.68	2.09	13.07
12	0.46	0.46	0.48	0.48	0.54	0.50	0.57
14	0.76	0.76	0.80	0.80	0.82	0.71	0.64
19	0.22	0.22	0.22	0.22	0.18	0.05	7.13
Avg	0.39	0.59	0.42	0.45	2.28	1.04	3.94

TABLE XII
ON PC1 (w E_{Oper} , 11 TEMPLATES)

we need $0.3 \times 1GB = 0.3GB$ samples for 1GB database.

Q	E_t	E_o	$E_s^{0.02}$	$E_s^{0.05}$	E_{SVM}	E_{REP}	E_o^{LR}
1	0.42	0.42	0.42	0.37	0.27	1.57	1.62
2	0.70	0.21	0.17	0.17	4.32	9.14	9.20
3	0.33	0.55	0.26	0.16	0.57	2.19	2.25
4	0.15	0.78	0.45	0.45	0.29	0.41	0.46
5	0.38	0.33	0.38	0.28	0.21	1.44	1.49
6	0.52	0.52	0.52	0.46	0.78	2.62	2.68
7	0.53	0.53	0.53	0.47	0.05	0.90	0.96
8	0.91	0.91	0.91	0.91	0.92	0.92	0.82
9	0.06	0.57	0.32	0.15	0.93	0.95	0.85
10	0.21	3.16	0.07	0.02	0.33	0.15	1.75
11	0.27	0.20	0.20	0.11	7.10	66.07	14.53
12	0.46	0.46	0.46	0.41	0.29	0.23	1.66
13	0.07	0.06	0.35	0.17	0.53	1.35	2.17
14	0.19	0.19	0.19	0.20	1.01	0.54	3.14
16	0.61	0.49	0.60	0.58	2.05	0.84	4.84
18	0.40	0.79	0.59	0.59	0.51	0.62	0.15
19	0.26	0.26	0.26	0.39	2.28	0.41	5.30
22	0.52	0.95	0.33	0.50	0.47	0.83	0.25
Avg	0.39	0.63	0.39	0.36	1.27	5.07	3.01

TABLE XIII
ON PC2 (w/o E_{Oper} , 18 TEMPLATES)

D. Prediction Errors on Skewed TPC-H 10GB Database

Table XV to XVIII further show the results on skewed TPC-H 10GB database. Table XV and XVI present the prediction

Q	E_t	E_o	$E_s^{0.02}$	$E_s^{0.05}$	E_{SVM}	E_{REP}	E_{Oper}
1	0.42	0.42	0.42	0.37	0.05	0.19	5.17
3	0.33	0.55	0.26	0.16	0.27	0.31	0.18
5	0.38	0.33	0.38	0.28	0.02	0.11	5.41
6	0.52	0.52	0.52	0.46	0.50	0.42	0.10
7	0.53	0.53	0.53	0.47	0.34	0.28	11.08
8	0.91	0.91	0.91	0.91	0.94	0.83	0.77
9	0.06	0.57	0.32	0.15	0.95	0.87	0.92
10	0.21	3.16	0.07	0.02	0.08	0.02	1.53
12	0.46	0.46	0.46	0.41	0.04	0.03	4.96
14	0.19	0.19	0.19	0.20	0.69	0.13	0.14
19	0.26	0.26	0.26	0.39	1.57	0.71	3.50
Avg	0.39	0.72	0.39	0.35	0.50	0.35	3.07

TABLE XIV
ON PC2 (w E_{Oper} , 11 TEMPLATES)

errors of different approaches on PC1, while Table XVII and XVIII present their errors on PC2, respectively. Once again, the results are very similar to those reported on skewed TPC-H 1GB database, as presented in Appendix B.

Q	E_t	E_o	$E_s^{0.02}$	$E_s^{0.05}$	E_{SVM}	E_{REP}	E_o^{LR}
1	0.08	0.08	0.08	0.08	0.23	0.50	0.19
2	0.33	0.37	0.23	0.40	2.35	4.03	2.74
3	0.35	0.82	0.21	0.22	0.28	0.35	0.31
4	0.75	0.79	0.74	0.75	0.98	1.95	0.84
5	0.08	0.01	0.08	0.09	0.12	0.32	0.04
6	0.02	0.02	0.02	0.02	1.09	2.11	1.00
7	0.41	0.40	0.41	0.41	0.24	0.13	0.25
8	0.48	0.43	0.48	0.48	0.57	0.37	0.47
9	0.55	0.50	0.95	0.37	0.77	0.69	0.47
10	0.17	1.29	0.11	0.10	0.20	0.25	0.44
11	0.81	0.76	0.78	0.78	0.24	0.56	0.41
12	0.11	0.11	0.11	0.11	0.36	0.75	0.39
13	0.77	0.77	1.46	0.99	0.70	1.59	0.95
14	0.32	0.34	0.32	0.32	0.59	1.41	0.71
16	0.61	0.53	0.60	0.61	0.57	1.38	0.44
18	0.50	2.03	1.91	1.87	0.80	0.65	0.62
19	0.26	0.26	0.26	0.26	0.49	0.03	0.56
22	0.22	0.74	1.87	0.42	2.62	1.70	2.93
Avg	0.38	0.57	0.59	0.46	0.73	1.04	0.76

TABLE XV
ON PC1 (w/o E_{Oper} , 18 TEMPLATES)

Q	E_t	E_o	$E_s^{0.02}$	$E_s^{0.05}$	E_{SVM}	E_{REP}	E_{Oper}
1	0.08	0.08	0.08	0.08	0.18	0.36	0.17
3	0.35	0.82	0.21	0.22	0.21	0.17	1.21
5	0.08	0.01	0.08	0.09	0.06	0.17	0.22
6	0.02	0.02	0.02	0.02	1.46	0.28	0.32
7	0.41	0.40	0.41	0.41	0.24	1.86	3.20
8	0.48	0.43	0.48	0.48	0.49	0.34	0.40
9	0.55	0.50	0.95	0.37	0.77	0.76	0.70
10	0.17	1.29	0.11	0.10	0.13	0.33	0.25
12	0.11	0.11	0.11	0.11	0.47	0.96	0.22
14	0.32	0.34	0.32	0.32	1.16	2.21	0.33
19	0.26	0.26	0.26	0.26	0.68	1.47	0.26
Avg	0.26	0.39	0.28	0.22	0.53	0.81	0.66

TABLE XVI
ON PC1 (w E_{Oper} , 11 TEMPLATES)

E. Sampling Overhead on TPC-H 1GB Database

In this section, we present the detailed results of the additional runtime overhead due to sampling on the TPC-H 1GB database. Table XIX and XX summarize the results over

Q	E_t	E_o	$E_s^{0.02}$	$E_s^{0.05}$	E_{SVM}	E_{REP}	E_o^{LR}
1	0.10	0.10	0.10	0.10	0.10	0.36	0.42
2	1.69	0.94	1.11	1.11	2.76	2.85	4.16
3	0.04	2.63	0.21	0.22	0.05	0.05	0.87
4	0.02	0.17	0.02	0.02	0.48	0.76	0.21
5	0.15	0.25	0.15	0.15	0.09	0.13	0.47
6	0.15	0.15	0.15	0.15	0.95	0.82	1.73
7	0.09	0.08	0.10	0.09	0.24	0.29	0.20
8	0.11	0.03	0.12	0.11	0.23	0.28	0.24
9	0.32	0.06	0.32	0.20	0.93	0.93	0.82
10	0.21	3.84	0.30	0.28	0.14	0.19	0.76
11	0.72	0.93	0.85	0.87	6.99	6.64	9.79
12	0.15	0.15	0.15	0.15	0.18	0.14	0.79
13	1.09	1.08	1.66	1.27	0.40	0.36	1.25
14	0.65	0.65	0.64	0.64	1.61	1.45	2.59
16	0.54	0.39	0.53	0.54	0.46	0.80	1.11
18	0.06	1.51	1.23	1.23	0.67	3.73	0.85
19	0.44	0.44	0.44	0.44	0.45	0.50	1.27
22	0.47	0.94	0.24	0.40	0.69	0.80	0.54
Avg	0.39	0.80	0.46	0.44	0.97	1.17	1.56

TABLE XVII
ON PC2 (w/o E_{Oper} , 18 TEMPLATES)

Q	E_t	E_o	$E_s^{0.02}$	$E_s^{0.05}$	E_{SVM}	E_{REP}	E_{Oper}
1	0.10	0.10	0.10	0.10	0.13	0.99	0.27
3	0.04	2.63	0.21	0.22	0.07	1.12	3.15
5	0.15	0.25	0.15	0.15	0.11	1.04	0.57
6	0.15	0.15	0.15	0.15	0.96	3.24	0.27
7	0.09	0.08	0.10	0.09	0.29	0.68	12.38
8	0.11	0.03	0.12	0.11	0.28	0.71	0.22
9	0.32	0.06	0.32	0.20	0.93	0.91	0.91
10	0.72	0.93	0.85	0.87	0.16	0.08	10.14
12	0.15	0.15	0.15	0.15	0.15	0.52	0.52
14	0.65	0.65	0.64	0.64	1.62	0.34	0.35
19	0.44	0.44	0.44	0.44	0.53	0.23	0.80
Avg	0.27	0.50	0.29	0.28	0.47	0.90	2.69

TABLE XVIII
ON PC2 (w E_{Oper} , 11 TEMPLATES)

uniform and skewed data on PC1, respectively. Table XXI and XXII report corresponding results on PC2.

Q	$r_s^{0.05}$	$r_s^{0.1}$	$r_s^{0.2}$	$r_s^{0.3}$	$r_s^{0.4}$
1	0.0495	0.0996	0.2000	0.3007	0.4008
2	0.0055	0.0138	0.0401	0.0756	0.1244
3	0.0794	0.1233	0.2607	0.3299	0.4314
4	0.0247	0.0552	0.1241	0.1971	0.2706
5	0.0293	0.0416	0.0826	0.1252	0.1780
6	0.0171	0.0344	0.1151	0.2173	0.3224
7	0.0028	0.0081	0.0284	0.0602	0.1039
8	0.0117	0.0273	0.0576	0.0873	0.1310
9	0.0005	0.0013	0.0042	0.0095	0.0188
10	0.0049	0.0105	0.0261	0.0531	0.1034
11	0.0028	0.0080	0.0270	0.0535	0.0892
12	0.0179	0.0366	0.1177	0.2157	0.3190
13	0.0106	0.0219	0.0794	0.1756	0.2668
14	0.0778	0.1120	0.1704	0.1948	0.2853
16	0.0097	0.0197	0.0523	0.0979	0.1618
17	0.0140	0.0302	0.0662	0.1052	0.1651
18	0.0299	0.0558	0.1044	0.3616	0.4670
19	0.0694	0.1073	0.2132	0.2942	0.3859
20	0.0070	0.0157	0.0379	0.0644	0.0983
21	0.0306	0.0633	0.1333	0.2125	0.2941
22	0.0478	0.0985	0.1990	0.2966	0.3945
Avg	0.0259	0.0469	0.1019	0.1680	0.2387

TABLE XIX
ON UNIFORM DATA (PC1)

Q	$r_s^{0.05}$	$r_s^{0.1}$	$r_s^{0.2}$	$r_s^{0.3}$	$r_s^{0.4}$
1	0.0495	0.1002	0.1997	0.3002	0.4002
2	0.0176	0.0368	0.0831	0.1432	0.2120
3	0.1541	0.2901	0.2360	0.3339	0.4208
4	0.0383	0.0818	0.1727	0.2685	0.3619
5	0.1989	0.3085	0.1022	0.1558	0.2213
6	0.0300	0.0579	0.1117	0.1897	0.3016
7	0.0027	0.0082	0.0274	0.0589	0.1018
8	0.1312	0.1881	0.1948	0.2854	0.3809
9	0.0007	0.0018	0.0056	0.0106	0.0254
10	0.0050	0.0119	0.0283	0.0532	0.1107
11	0.0035	0.0108	0.0394	0.0708	0.1119
12	0.0312	0.0619	0.1217	0.2015	0.3098
13	0.0104	0.0218	0.0745	0.1714	0.2631
14	0.2239	0.2399	0.2953	0.3668	0.4443
16	0.0458	0.0358	0.0468	0.0953	0.1515
17	0.0053	0.0102	0.0227	0.0505	0.1116
18	0.0160	0.0484	0.1030	0.3164	0.4170
19	0.1128	0.1160	0.2090	0.3123	0.4142
20	0.0144	0.0310	0.0726	0.1151	0.1696
21	0.0336	0.0685	0.1447	0.2203	0.3039
22	0.0497	0.1031	0.2069	0.3072	0.4063
Avg	0.0559	0.0873	0.1190	0.1918	0.2686

TABLE XX
ON SKEWED DATA (PC1)

Q	$r_s^{0.05}$	$r_s^{0.1}$	$r_s^{0.2}$	$r_s^{0.3}$	$r_s^{0.4}$
1	0.0497	0.0993	0.1978	0.298	0.3962
2	0.0211	0.0442	0.1056	0.1762	0.2575
3	0.0462	0.0926	0.1838	0.2807	0.378
4	0.0249	0.0555	0.1228	0.1974	0.2697
5	0.0284	0.0585	0.1221	0.1958	0.2782
6	0.0163	0.0324	0.1112	0.2151	0.3224
7	0.0257	0.0495	0.1252	0.2177	0.313
8	0.0267	0.0543	0.1125	0.1753	0.2441
9	0.0054	0.011	0.0244	0.041	0.0622
10	0.0234	0.0463	0.0946	0.1481	0.2159
11	0.0034	0.0094	0.03	0.0594	0.1157
12	0.0216	0.0444	0.1255	0.2289	0.3334
13	0.0364	0.0779	0.1681	0.2622	0.3573
14	0.0229	0.0462	0.0864	0.1385	0.233
16	0.0062	0.0154	0.0458	0.0908	0.1526
17	0.0106	0.0239	0.0506	0.0908	0.148
18	0.028	0.0555	0.1188	0.3204	0.4339
19	0.0491	0.0971	0.193	0.2876	0.3899
20	0.0075	0.0165	0.0381	0.0675	0.1042
21	0.0364	0.0714	0.1402	0.2111	0.2888
22	0.0479	0.095	0.2013	0.2883	0.3914
Avg	0.0256	0.0522	0.1142	0.1900	0.2707

TABLE XXI
ON UNIFORM DATA (PC2)

We observe that, although the underlying data distribution and hardware configuration are quite different, the additional overheads due to sampling are consistent. Sampling incurs about 20% additional time (i.e., the sampling ratio $f = 0.3$) to achieve predictions close to what if the actual cardinalities were available. This overhead will dramatically decrease when we move to bigger datasets.

F. Sampling Overhead on TPC-H 10GB Database

We further show the additional overhead of sampling on the TPC-H 10GB database from Table XXIII to Table XXVI. As before, the results on PC1 are shown in Table XXIII and XXIV, and the results on PC2 are shown in Table XXV and XXVI.

Q	$r_s^{0.05}$	$r_s^{0.1}$	$r_s^{0.2}$	$r_s^{0.3}$	$r_s^{0.4}$
1	0.0497	0.1027	0.2021	0.3015	0.4019
2	0.0234	0.047	0.1048	0.1762	0.2591
3	0.0465	0.0947	0.1876	0.2915	0.3729
4	0.0381	0.0819	0.1691	0.2629	0.3586
5	0.0292	0.0617	0.1319	0.2107	0.2999
6	0.0268	0.0526	0.1017	0.1748	0.2896
7	0.0312	0.0623	0.13	0.2319	0.3254
8	0.0442	0.0893	0.1797	0.2745	0.3759
9	0.0019	0.0039	0.0097	0.0167	0.0306
10	0.0277	0.0559	0.1143	0.1764	0.247
11	0.0053	0.0144	0.0438	0.0791	0.1238
12	0.0352	0.0687	0.1373	0.2205	0.3282
13	0.0356	0.0802	0.1675	0.2692	0.3597
14	0.0523	0.0977	0.1876	0.2793	0.3567
16	0.0058	0.014	0.0383	0.0931	0.1473
17	0.0064	0.0112	0.0266	0.0545	0.1164
18	0.0235	0.0471	0.1077	0.278	0.3828
19	0.0483	0.0973	0.1994	0.2982	0.3912
20	0.0147	0.0302	0.0731	0.1135	0.1739
21	0.0402	0.0851	0.1641	0.2421	0.3248
22	0.0496	0.1033	0.208	0.3163	0.4125
Avg	0.0303	0.0620	0.1278	0.2077	0.2894

TABLE XXII
ON SKEWED DATA (PC2)

Q	$r_s^{0.01}$	$r_s^{0.02}$	$r_s^{0.05}$	$r_s^{0.1}$
1	0.0090	0.0160	0.0406	0.0822
2	0.0060	0.0165	0.0252	0.0491
3	0.0080	0.0181	0.0452	0.0478
4	0.0021	0.0059	0.0197	0.0680
5	0.0261	0.0096	0.0177	0.1025
6	0.0089	0.0116	0.0303	0.0735
7	0.0179	0.0273	0.0656	0.0892
8	0.0039	0.0068	0.0123	0.0241
9	0.0042	0.0063	0.0193	0.0427
10	0.0198	0.0291	0.0457	0.0808
11	0.0014	0.0024	0.0132	0.0344
12	0.0112	0.0199	0.0569	0.0995
13	0.0076	0.0123	0.0333	0.0651
14	0.0233	0.0216	0.0498	0.1106
16	0.0195	0.0087	0.0224	0.0361
18	0.0045	0.0245	0.0584	0.1075
19	0.0103	0.0333	0.0780	0.1377
22	0.0053	0.0187	0.0769	0.1779
Avg	0.0105	0.0160	0.0395	0.0794

TABLE XXIII
ON UNIFORM DATA (PC1)

We have shown that, while the database is 10 times bigger, the number of samples required to achieve predictions that are close to the ideal case where perfect cardinalities are known remains almost the same. This explains the dramatic drop on the relative overhead sampling incurs. For sampling ratio $f = 0.05$, the average additional runtime overhead caused by sampling is about 4%. Note that this is consistent on both uniform and skewed data, and on both PC1 and PC2. This demonstrates the practicality of using sampling for the purpose of predicting query execution time. We also note here that even 4% may be still too costly if we want to use sampling for cardinality estimation in query optimization, since it means that the optimizer can consider only $1/4\% = 25$ plans even if it is given the time within which the query can finish execution.

Q	$r_s^{0.01}$	$r_s^{0.02}$	$r_s^{0.05}$	$r_s^{0.1}$
1	0.0074	0.0186	0.0448	0.0902
2	0.0065	0.0127	0.0299	0.0642
3	0.0136	0.0218	0.0537	0.0928
4	0.0057	0.0122	0.0351	0.0808
5	0.0177	0.0218	0.0372	0.0792
6	0.0103	0.0201	0.0526	0.1019
7	0.0214	0.0305	0.0780	0.1259
8	0.0089	0.0197	0.0455	0.1108
9	0.0020	0.0034	0.0047	0.0124
10	0.0135	0.0190	0.0475	0.0905
11	0.0008	0.0022	0.0056	0.0372
12	0.0101	0.0209	0.0517	0.1172
13	0.0059	0.0123	0.0294	0.0588
14	0.0261	0.0418	0.0974	0.1560
16	0.0047	0.0078	0.0203	0.0362
18	0.0038	0.0086	0.0412	0.0945
19	0.0155	0.0253	0.0748	0.1474
22	0.0083	0.0181	0.0802	0.2034
Avg	0.0101	0.0176	0.0461	0.0944

TABLE XXIV
ON SKEWED DATA (PC1)

Q	$r_s^{0.01}$	$r_s^{0.02}$	$r_s^{0.05}$	$r_s^{0.1}$
1	0.0116	0.0213	0.0524	0.0999
2	0.009	0.0167	0.0332	0.0667
3	0.009	0.018	0.0459	0.0891
4	0.0082	0.0243	0.0706	0.1546
5	0.0092	0.0182	0.0445	0.0919
6	0.0059	0.0112	0.0303	0.0608
7	0.0058	0.0118	0.029	0.0558
8	0.0006	0.0014	0.0049	0.0148
9	0.0006	0.0012	0.003	0.007
10	0.0069	0.0137	0.0305	0.0579
11	0.0129	0.0208	0.0507	0.0984
12	0.0066	0.0134	0.0351	0.0697
13	0.0091	0.017	0.0399	0.0805
14	0.0101	0.0204	0.0477	0.0975
16	0.0069	0.0087	0.0195	0.0417
18	0.0066	0.0135	0.0297	0.0518
19	0.0151	0.0217	0.0513	0.1044
22	0.0022	0.0083	0.0423	0.1208
Avg	0.0076	0.0145	0.0367	0.0757

TABLE XXV
ON UNIFORM DATA (PC2)

Q	$r_s^{0.01}$	$r_s^{0.02}$	$r_s^{0.05}$	$r_s^{0.1}$
1	0.0121	0.0227	0.0529	0.0975
2	0.0077	0.0186	0.0444	0.0881
3	0.0084	0.0172	0.0474	0.1008
4	0.0128	0.0287	0.0795	0.1554
5	0.0085	0.0176	0.0469	0.0943
6	0.0109	0.0211	0.0505	0.099
7	0.0076	0.014	0.0372	0.0745
8	0.0085	0.0161	0.0431	0.0898
9	0.0015	0.0031	0.0079	0.019
10	0.0096	0.0172	0.0389	0.0755
11	0.0107	0.0218	0.0522	0.1221
12	0.0119	0.0214	0.0503	0.1021
13	0.0091	0.0168	0.0363	0.0746
14	0.0166	0.0309	0.0676	0.1274
16	0.0038	0.0065	0.017	0.0361
18	0.0072	0.0131	0.0349	0.0536
19	0.0108	0.0201	0.0531	0.1035
22	0.0025	0.0078	0.0484	0.1502
Avg	0.0089	0.0175	0.0449	0.0924

TABLE XXVI
ON SKEWED DATA (PC2)