



Серия «Математика»
2025. Т. 52. С. 137–152

Онлайн-доступ к журналу:
<http://mathizv.isu.ru>

ИЗВЕСТИЯ

Иркутского
государственного
университета

Research article

УДК 518.517

MSC 68T05, 68P15

DOI <https://doi.org/10.26516/1997-7670.2025.52.137>

Adaptive Cost Model for Query Optimization

Nikita K. Vasilenko^{1✉}, Alexander V. Demin^{1✉}, Denis
K. Ponomaryov^{1✉}

¹ Ershov Institute of Informatics Systems SB RAS, Novosibirsk, Russian Federation

✉ vasilenko.nikita.research@gmail.com

✉ alexandredemin@yandex.ru

✉ ponom@iis.nsk.su

Abstract. The principal component of conventional database query optimizers is a cost model that is used to estimate expected performance of query plans. The accuracy of the cost model has direct impact on the optimality of execution plans selected by the optimizer and thus, on the resulting query latency. Several common parameters of cost models in modern DBMS are related to the performance of CPU and I/O and are typically set by a database administrator upon system tuning. However these performance characteristics are not stable and therefore, a single point estimation may not suffice for all DB load regimes. In this paper, we propose an Adaptive Cost Model (ACM) which dynamically optimizes CPU- and I/O-related plan cost parameters at DB runtime. By continuously monitoring query execution statistics and the state of DB buffer cache ACM adjusts cost parameters without the need for manual intervention from a database administrator. This allows for responding to changes in the workload and system performance ensuring more optimal query execution plans. We describe the main ideas in the implementation of ACM and report on a preliminary experimental evaluation showing 20% end-to-end latency improvement on TPC-H benchmark.

Keywords: query optimization, cost model, online machine learning

For citation: Vasilenko N. K., Demin A. V., Ponomaryov D. K. Adaptive Cost Model for Query Optimization. *The Bulletin of Irkutsk State University. Series Mathematics*, 2025, vol. 52, pp. 137–152.

<https://doi.org/10.26516/1997-7670.2025.52.137>

Научная статья

**Адаптивная стоимостная модель для оптимизации
запросов**

Н. К. Василенко^{1✉}, А. В. Демин^{1✉}, Д. К. Пономарев^{1✉}

¹ Институт систем информатики им. А. П. Ершова СО РАН, Новосибирск, Российская Федерация

✉ vasilenko.nikita.research@gmail.com

✉ alexandremin@yandex.ru

✉ ponom@iis.nsk.su

Аннотация. Рассматривается важная компонента традиционных оптимизаторов запросов в базах данных — стоимостная модель, которая используется для оценки ожидаемой эффективности планов исполнения запросов. Точность модели непосредственно влияет на оптимальность планов, выбираемых оптимизатором, и, как следствие, на итоговую скорость исполнения запросов. Несколько параметров стоимостной модели в современных СУБД, как правило, связаны с производительностью ЦПУ и скоростью ввода/вывода хранения данных и обычно задаются администратором СУБД при настройке системы. Однако эти параметры производительности нестабильны во времени, и потому оценка в отдельной точке может не быть оптимальной для всех режимов работы БД. Предлагается адаптивная стоимостная модель (АСМ), которая позволяет динамически оптимизировать параметры стоимости планов, связанные с ЦПУ и хранилищем, во время работы БД. Модель непрерывно анализирует статистику исполнения запросов и состояние буферного кэша СУБД и динамически выставляет значение соответствующих параметров стоимости планов без необходимости ручного вмешательства со стороны администратора. Это позволяет адаптироваться к изменениям в нагрузке на систему и выбирать более оптимальные планы для запросов. Описываются основные идеи в реализации АСМ и приводятся результаты предварительных экспериментов, показывающих 20

Ключевые слова: оптимизация запросов, стоимостная модель, онлайн машинное обучение

Ссылка для цитирования: Vasilenko N. K., Demin A. V., Ponomaryov D. K. Adaptive Cost Model for Query Optimization // Известия Иркутского государственного университета. Серия Математика. 2025. Т. 52. С. 137–152.
<https://doi.org/10.26516/1997-7670.2025.52.137>

1. Introduction

The convenience of (declarative) query languages goes along with the complexity of finding an optimal execution plan for a query. It is well-known that for a SQL query there may exist an exponential number of execution plans (in the size of the query) and the problem to find one with a minimal latency is NP-complete. In the search for near-optimal execution plans, a typical query optimizer relies on a cost model which provides an estimation of the expected plan latency in terms of an abstract cost. In conventional implementations, a cost model is essentially a collection of formulas providing integer costs to the physical implementation of query operators and formulas to aggregate these costs into the resulting cost of a

plan. The accuracy of the cost model has direct impact on the optimality of plans selected by the optimizer and therefore, on the end-to-end query latency and overall DB performance. A higher cost would typically mean a higher execution time for a plan, but in practice these values are often decorrelated, non-optimal plans are selected by the optimizer and some queries appear to be slow.

A typical cost-based query optimizer relies on three main components: cardinality estimation, CPU and I/O cost estimation, and plan enumeration. Cardinality estimation means predicting the number of records in the output of a particular plan node. Numerous techniques to improve the accuracy of cardinality estimation have been proposed in the literature including [2–4;6]. The component related to CPU and I/O cost estimation is also important for accurate estimation of query plan performance, because performance of a plan (when seen as a program) is obviously affected by CPU and I/O load at the execution time [7;14]. Finally, plan enumeration is to find an execution plan with the lowest cost. By employing all these components, the task of an optimizer is to find the most efficient execution plan for a query, thereby improving the overall DB performance. However this approach has a significant drawback: in modern DBs cost estimation may be wrong by **several orders of magnitude** (e.g., when estimating cost of complex queries with multiple joins [14]).

Recently, an idea of learning-based optimizer leveraging prior query execution experience has been proposed [1; 10; 16]. This approach tries to mitigate the issues of conventional cost-based optimizers by employing learning from complex prior examples. The downside of this method is that it is well-suited for static workloads, but in the case of data drift, changes in query patterns, or data schema updates the optimizer has to be retrained.

In [8], the authors showed that to reduce cost estimation errors, one has to improve both, cardinality estimation accuracy and cost model parameter tuning. In this work, we argue that dynamically tuning CPU- and I/O-related cost model parameters alone can improve optimizer decisions in the search for an optimal plan.

Typically, a database administrator configures these cost model parameters manually. Finding the most suitable values for these parameters requires a deep understanding of the database structure, workload characteristics, and system resources. Inaccurate parameter setting can lead to suboptimal execution plans and degraded performance. Furthermore, cost parameters may need to be periodically adjusted to adapt to changes in workload, data access patterns, and system performance. This makes selecting optimal cost parameters complicated.

Another approach is to execute queries from a specifically designed calibration workload and collect execution statistics on it [14]. Then cost model parameters can be adjusted based on these statistics. However in this approach one has to periodically execute a synthetic calibration

workload in addition to the real DB workload and even in this case the obtained parameter estimations may be inadequate for the real workload due differences with calibration queries.

In this work, we propose an **adaptive cost model (ACM)**, a novel approach for optimizing CPU- and I/O-related plan cost parameters which implements lightweight ML models to analyze statistics on executed queries and makes prediction on the parameter values for next queries. The solution consists of two components: a model for estimating cost parameters related to CPU performance and a model for estimating disk-related parameters. The approach implements online fine-tuning which allows for dynamically adapting to workload changes. It avoids the need of using calibration queries, thereby reducing the overall load on database and also on DBAs by improving query plans and reducing the need to manually tune a number of configuration parameters.

We begin our exposition in Section 2 with a brief summary on related work and the main ingredients of cost models in conventional database optimizers. Then we describe in Section 3 the main ideas and design solutions in Adaptive Cost Model: the CPU- and disk-related components. Finally, in 4 we provide results of a preliminary experimental evaluation on TPC-H benchmark and in Section 5 we conclude.

2. Background

2.1. RELATED WORK

In the literature, there is a number of approaches to configuring grand unified configuration parameters (GUCs) of databases, including those related to cost model. The most general methods rely on search techniques like random search, simulated annealing, and Reinforcement Learning [5; 9; 12; 13; 15], to name a few. In these approaches, cost model parameters are treated as general database configuration parameters and it is suggested to optimize them with search techniques just like any other DB parameters.

In [12], the authors proposed to use a language model combined with Reinforcement Learning techniques to identify DB configuration parameters from manuals and posts in open DBA discussion groups. In [15], a deep deterministic policy gradient method was proposed to find optimal configurations in high-dimensional continuous spaces with Reinforcement Learning.

In [9], the authors propose to consider three levels of tuning with Reinforcement Learning. At the query level, GUCs are tuned separately for each query prior to execution by representing a query as a vector and recommending a configuration. At the workload level, the configuration is tuned for the entire workload by combining all query vectors into one and

generating a recommended configuration. At the cluster level, queries are separated into clusters and configurations are tuned for each cluster.

In [5], the authors identify the most important GUCs for tuning by first using Latin Hypercube Sampling (LHS) to sample configurations for all GUCs. They then collect statistics for these samples by running benchmarks for each configuration. Then by using a random forest, they rank GUCs to determine the most significant ones.

In the approach presented in [13], the authors propose to begin with selecting the closest matching configuration for a database. During an initial observation period, they characterize the workload and hardware by using internal runtime metrics and by running queries. Factor analysis and K-means clustering are used to reduce the value of the metric and the result is used for system profiling. The most important knobs are selected by using Lasso and ranking. Tuning proceeds incrementally, increasing the number of knobs during the process. Initially, they find the closest similar workload from the repository, and then they determine the best configuration by using a Gaussian Process.

Another class of approaches is based on the idea of using calibration queries and formulas to adjust cost parameters. For example, in [14], based on the observed latency of calibration queries, the values of the parameters are derived analytically from a set of formulas.

We can summarize all these approaches as follows:

- Search-based approaches:
 - provide more accurate parameter values
 - are time-costly
- Calibration-based approaches:
 - allow for quickly estimating parameter values
 - face the problem of constructing a calibration query for each parameter
 - may provide relatively high errors preventing accurate calculation of optimal values

Both approaches share the common disadvantage: they require a large number of trials to assess the impact of a set of parameter values on workload execution time.

2.2. COST MODEL IN DB OPTIMIZERS

Conventional open-source cost-based databases with bottom-up optimization (e.g., relational DBs including PostgreSQL, MySQL, openGauss, etc.) employ a cost model which includes the following five key parameters:

1. *cpu_tuple_cost*, the CPU cost to process a tuple,
2. *cpu_index_tuple_cost*, the CPU cost to process a tuple via index access,

3. *cpu_operator_cost*, the CPU cost to perform an operation such as hashing or aggregation,
4. *seq_page_cost*, the I/O cost to sequentially access a page,
5. *random_page_cost*, the I/O cost to randomly access a page.

The cost of the query execution plan is calculated as the sum of the costs of the operators in the plan. The cost of operators is calculated as a linear combination of basic parameters (cost-relevant GUC parameters being numeric constants) and variable values such as the number of data blocks read from disk, etc.

In general, the cost of a SQL operator is computed as the sum of the products of cost model parameters and estimated cardinalities:

$$Cost_{operator} = c_t \cdot n_t + c_o \cdot n_o + c_s \cdot n_s + c_i \cdot n_i + c_r \cdot n_r, \quad (2.1)$$

where:

- c_t - cost of processing a tuple,
- n_t - number of tuples processed,
- c_o - cost of performing an operation,
- n_o - number of operations performed,
- c_s - cost of a sequential disk page fetching,
- n_s - number of disk pages fetched sequentially,
- c_i - cost of processing an index entry,
- n_i - number of index entries processed,
- c_r - cost of fetching a disk page randomly,
- n_r - number of disk pages fetched randomly.

For example, the cost of sequential access to a table (SeqScan operator) without filter can be defined as:

$$Cost_{SeqScan} = c_t \cdot n_t + c_s \cdot n_s$$

and the cost of aggregation operator is:

$$Cost_{Agg} = c_t \cdot n_t$$

The optimal choice of values for the parameters above depends on the hardware configuration, workload type, as well as on the ‘warm-up degree’ of the database, i.e., the amount of data which resides in the buffer cache.

In the above mentioned list of parameters, *seq_page_cost* is typically set to the time required to process one page in sequential scan. The rest of the parameters are set as fractions of this base value. For instance, *cpu_tuple_cost* set to 0.01 means that processing 100 rows requires the same time as reading one page sequentially from disk. Similarly, *random_page_cost* equal to 4 means that the time required for random access to a page exceeds the sequential access time by a factor of 4.

Thus, to increase the cost of sequential page scan, one can decrease the values of the other four parameters, and vice versa, to increase the cost of all operations except *seq_page_cost*, it suffices to decrease only *seq_page_cost*. Therefore, to adjust disk and CPU-related configuration of a cost model it suffices to change only four parameters. W.l.o.g. we assume in this paper that *seq_page_cost* is 1.

The primary objective of cost model parameter tuning is to increase the correlation between the cost and the execution time of a node, which allows for estimating the execution time of plans more accurately. Thus, the task looks like:

$$\text{corr}(\text{cost}, \text{time}) \rightarrow 1 \quad (2.2)$$

where

- *cost* is a cost for a plan operator,
- *time* is a real execution time for this operator.

3. Design of ACM

ACM implements computation of the cost parameters associated with disk I/O operations (I/O cost) and cost parameters related to the processor utilization (CPU cost) in two different ways. Having two different approaches is natural, because CPU parameters do not depend on the current state of the database, while disk-related parameters can change depending on the state of the database buffer cache, which varies nearly after every query.

ACM computes optimal I/O cost parameters before every query by assessing the volume of requested data already present in the database buffer cache and by estimating the amount of data that needs to be retrieved from the disk. The model employs historical data on the execution of past queries to generate accurate predictions.

To calculate optimal CPU-related cost parameters, the model is continuously monitoring query operator execution statistics and computes parameters to align the estimated execution time (determined at the planning stage) as close as possible to the actual execution time (obtained post query execution). This way ACM continuously adjusts the cost model parameters to provide more accurate estimates of plan execution costs, which guides the database optimizer towards more optimal query plans.

In the following, we provide a detailed description of the two components of ACM: the disk-related model and CPU-related model.

3.1. DISK-RELEVANT PARAMETERS

Heuristic-based formulas of the conventional cost model take into account the difference between the cost of sequential and random page access by differences in cost parameters. However, this approach does not allow for

estimating the number of requested pages that are already in the buffer cache and thus, do not have to be fetched from disk. Since the optimizer is not aware of how much of the requested data already resides in the memory, it makes estimation errors. ACM improves the accuracy of cost estimation by dynamically adjusting the cost of random disk access based on the information about the availability of data in the buffer.

When a query is executed, ACM collects statistics on the usage of data from the buffer (buffer hit) and disk (read). This relationship is important because it gives an idea of how much of the requested data was in the database memory. First of all, the proposed model separately calculates the *random_page_cost* parameter for each table, as one table might reside entirely in the buffer cache, while another table might be absent. Having the entire table in the buffer cache allows for setting *random_page_cost* equal to *seq_page_cost*, suggesting to the optimizer that the read time will be the same. Thus, ACM implements a dynamic computation of *random_page_cost* relative to tables. As already mentioned, an exact estimate for *random_page_cost* depends on the information on how much of the requested data will be taken from the buffer and how much will be read from disk. This information can be expressed by the 'hit ratio' defined as:

$$R(q) = \frac{hit}{hit + read} \quad (3.1)$$

where

- q is a query,
- hit is the number of pages that was requested by the query and read from the memory,
- $read$ is the number of pages that was requested by the query and read from the disk.

Typically, this information becomes available after query execution, but ACM tries to estimate this ratio by using statistics on the execution of previous queries. For each table, the model collects hit ratio statistics:

$$R(q_t), R(q_{t-1}), R(q_{t-2}), \dots \quad (3.2)$$

where

- $R(q_t)$ is the hit ratio of query q_t executed at time t .

To estimate the hit ratio $R(q_{t+1})$ for a new query q_{t+1} one can use various models to predict the hit ratio based on previous values $R(q_t)$, $R(q_{t-1})$, For example, one can use various time series prediction models like ARIMA, exponential smoothing, etc.. But in ACM we employ a simpler solution which analyses only the previous hit ratio and a degradation factor. In fact, in our experiments, this simpler model showed results as good as more complex ones:

$$R(q_{t+1}) = R(q_t) \cdot D, \quad (3.3)$$

where D is a degradation factor defined as:

$$D = \frac{1 + (qc - tc)}{1 + (qc - tc)^2} \quad (3.4)$$

and

- qc is a query counter (current total number of accesses to all tables; it gets incremented upon any table read),
- tc is a counter that indicates the last access to the target table. If target table is accessed, the current value of qc is written to tc .

The meaning of the degradation factor is explained as follows. In general, after a query referring to a table T_1 is executed, some other queries referring to tables T_2, \dots, T_n may be processed, which may cause the data of table T_1 to be evicted from the buffer. The more requests to different tables have been executed since the last access to table T_1 , the less likely its data will remain in the buffer. The degradation factor takes this effect into account.

Next, it is necessary to relate $R(q_{t+1})$ to *random_page_cost* so that at maximum the predicted hit ratio *random_page_cost* equals *seq_page_cost*. We set:

$$\begin{aligned} random_page_cost_{new} = & random_page_cost_{def} \cdot \\ & (1 - R(q_{t+1})) + seq_page_cost \cdot R(q_{t+1}) \end{aligned} \quad (3.5)$$

where

- $random_page_cost_{def}$ is the default value of *random_page_cost*,
- $random_page_cost_{new}$ is the new value for *random_page_cost* that is used to calculate costs of scans for a particular table.

Figure 1 shows the schema of computing buffer usage statistics at query execution. The more pages are read for the table from the memory, the lower the cost of subsequent random accesses to the memory will be, since most pages can be read from the buffer. Thus, the cost of random access to the page is adjusted for each data table.

The number of pages read from disk and memory is collected for each table until the amount of data reaches the level when it is possible to predict the hit ratio for a next table access. After that, at query planning, the cost of random access to the page is computed by using the degradation factor as defined above and the obtained costs are injected. ACM considers the size of the time interval between the last access to the table and the scheduled one, changing the number of hits to the degradation factor. With cost injection, query planning is standard and after a query is executed new parameter values will be recorded for all the tables with the data.

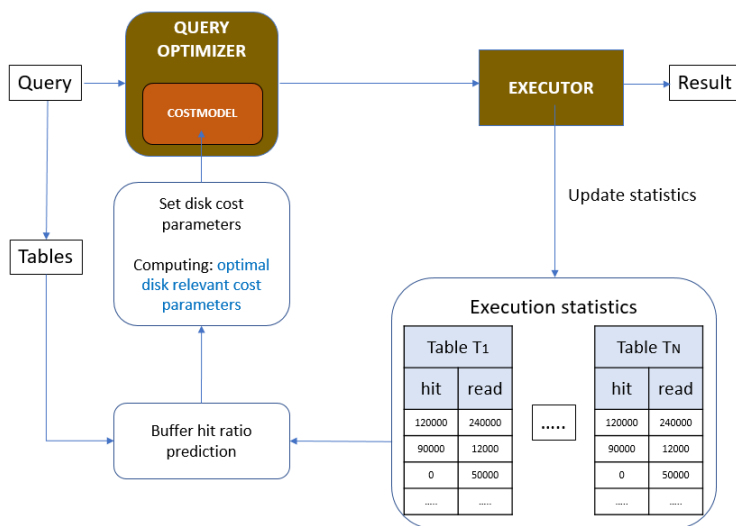


Figure 1. Computation of disk-relevant cost parameters

3.2. CPU-RELEVANT PARAMETERS

The other three parameters in the standard cost model refer to CPU operations, such as the cost of processing one tuple (tuple cost), performing one operation (operator cost), and processing one index tuple (index tuple cost). Standard parameter values do not reflect specifics of various hardware setups and database configurations which again may lead to inaccurate estimates of query costs. Also, standard cost models do not provide individual cost computation formulas for different operators which is another source of inaccuracy. To address these issues, ACM provides a method to collecting statistics and dynamically set parameters to account for a DB configuration and hardware setup individually for each plan operator. Figure 2 shows the logic of statistics collection on CPU usage at query execution.

After a query is executed, the system collects key statistics including the execution time for operators and the number of tuples processed by each operator. Statistics for each type of operator are collected separately. The purpose of using these statistics is to reduce the mismatch between the output values of operator cost formulas and the execution times for these operators. The minimization function uses linear lightweight models, which have a minimal computation overhead. An operator cost is defined as:

$$F(c_t, c_o, c_i) = c_t \cdot n_t + c_o \cdot n_o + c_i \cdot n_i + s, \quad (3.6)$$

$$s = c_s \cdot n_s + c_r \cdot n_r$$

where

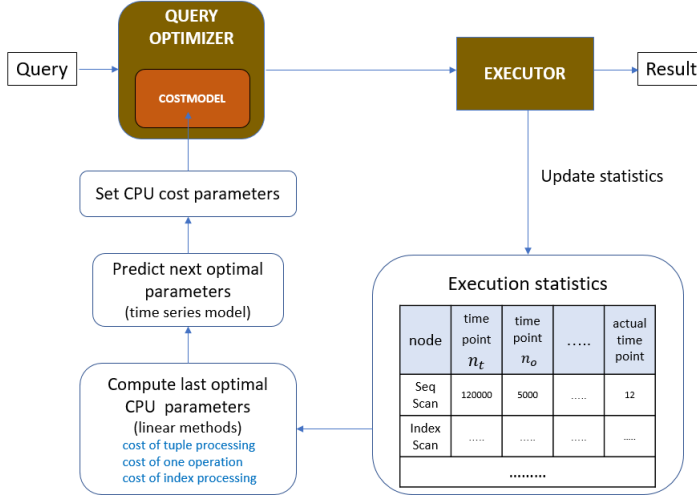


Figure 2. Computation of CPU-relevant cost parameters

- c_t is a cost of processing a tuple,
- n_t a number of tuples processed,
- c_o a cost of performing an operation,
- n_o a number of operations performed,
- c_i a cost of processing an index entry,
- n_i a number of index entries,
- s a cost of sequential and random reads for a node with GUC parameters derived from disk-relevant model.

Based on this cost formula, a simple linear model is defined:

$$L(c, x) = x^T c \quad (3.7)$$

where the vector $c = (c_t, c_o, c_i, 1)$ of model parameters and $x = (n_t, n_o, n_i, s)$ are independent variables. Note that s is treated as an independent variable because its value is set by the disk-relevant model.

Suppose we have a sample of observations X of variable values collected by the system at query execution:

$$X = \begin{pmatrix} n_t^1 & n_o^1 & n_i^1 & s^1 \\ n_t^2 & n_o^2 & n_i^2 & s^2 \\ \dots & \dots & \dots & \dots \\ n_t^n & n_o^n & n_i^n & s^n \end{pmatrix} \quad (3.8)$$

We can rewrite the model 3.7 in the matrix form:

$$L(c, X) = Xc \quad (3.9)$$

and then the task of finding optimal cost parameters is reduced to the linear regression problem:

$$L(c, X) \rightarrow time \cdot scale_factor \quad (3.10)$$

where

- $time = \begin{pmatrix} t^1 \\ t^2 \\ \dots \\ t^n \end{pmatrix}$ is a collected execution time for an operator,
- $scale_factor$ is a constant representing a scaling factor between execution time and cost.

Various techniques can be used here to estimate model parameters. In our implementation we used the classical least squares method.

After collecting statistics and calculating cost parameters, ACM saves the results for predicting cost parameters for future queries. When adjusting parameters dynamically, it is important to respond in a timely manner to changes in database configuration, as well as to other factors that may affect cost. ACM uses time series analysis to estimate future cost parameters when they are individually set for each operator type. In our implementation we used the exponential moving average, but we note that other models (e.g. ARIMA, etc.) can be used. The exponential moving average is defined as:

$$c_{pred,n} = (1 - \alpha) \cdot c_{n-1} + \alpha \cdot c_{pred,n-1} \quad (3.11)$$

where

- $c_{pred,i}$ is a predicted value for cost parameter on step i ,
- c_i a cost parameter value derived from the linear model (3.9) on step i ,
- α a smoothing factor.

In this way, ACM employs query execution statistics to dynamically adjust 5 essential cost model parameters, which enables the DB optimizer to estimate cost more accurately for each type of operator. Information analyzed by ACM includes (but is not limited to) buffer state information, execution time of queries and nodes, the number of tuples produced and other key feature values.

4. Preliminary Experiments

We now present first results of an experimental evaluation of ACM on the standard TPC-H benchmark [11] in which we answer the following questions:

- Is correlation between time and cost for nodes increased with ACM?
- Does ACM improve benchmark execution latency?

For experiments, we used a 10Gb TPC-H dataset consisting of 22 template-based queries. We implemented ACM in the open-source database openGauss, which has a cost model similar to PostgreSQL. The benchmark

was run several times to ‘warm up’ buffer followed by the execution of test queries with measurements.

The correlation between cost and execution time of plans for TPC-H queries without/with ACM is shown in Figures 3. One can notice two groups of nodes on the left part of the figure: one group has time within 200 ms and cost value up to 5000, while the second group lies significantly higher, with costs ranging from 20,000 to 25,000 and times from 200 to 400 ms. The correlation between cost and time is 0.29.

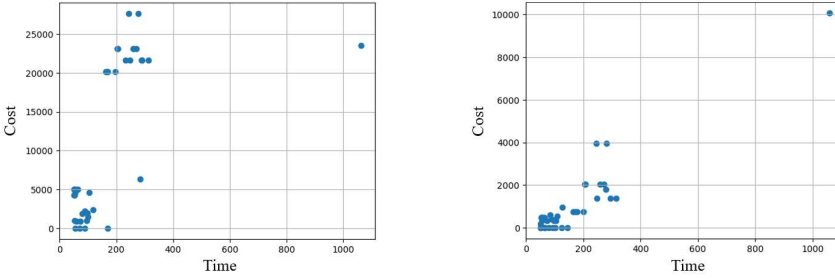


Figure 3. Correlation between node cost and execution time: with standard openGauss cost model (left) and with ACM(right)

When using ACM, the correlation is increased to 0.92, which means a big improvement of the cost model accuracy.

Figure 4 shows end-to-end latency of TPC-H queries.




Figure 4. E2E latency of TPC-H queries: summary

The query plan is changed for queries 3, 5, 7, 9, 10, 12, and 17. For the remaining queries, the execution plan remains unchanged. The overall latency improvement for queries with a modified plan is 46% and the latency improvement for the entire benchmark is 20%. We provide detailed latency results (in milliseconds) in Table 1.

Table 1

E2E latency values for TPC-H queries

Query No.	Baseline	ACM	Query No.	Baseline	ACM
1	128235	127906	12	51593	39639
2	1325	1423	13	33872	33795
3	42453	22632	14	4685	4442
4	5969	5956	15	5893	5949
5	43756	9860	16	8992	8816
6	4290	4254	17	1971	1225
7	42803	21306	18	52151	51435
8	6819	6808	19	36333	36371
9	93829	61407	20	71459	71743
10	35248	11726	21	47928	47944
11	2308	2225	22	3909	3878

We note that some queries could not be sped-up with ACM which was primarily caused by errors in cardinality estimation. Even though ACM improves key parameters of the cost model it may not be able to shift the choice of optimizer towards better plans in all cases due to inaccuracy of cardinality estimation.

5. Conclusion

In this paper, we have proposed a method to dynamically adjust CPU- and I/O-related parameters of a conventional cost model of a database optimizer. Our approach differs from other solutions known in the literature in that it is online, it does not require pre-training or periodic running of calibration queries on the database, and it is based on light-weight learning models with a minimal computational overhead, which makes it practical for implementation in DB kernels. The proposed solution takes into account several nuances for a more accurate cost estimation including buffer hit ratio prediction for accessed tables and per operator cost estimation. Even though improving the accuracy of CPU- and disk-related cost model parameters alone can not improve the optimality of plans in all cases, we have demonstrated in preliminary experiments on TPC-H benchmark that it may have a significant positive impact on query latency due to improved correlation between cost and execution time. The next natural step would be to evaluate this approach in conjunction with methods to improve cardinality estimation and planning in general, and we leave this for future work.

References

1. Akdere M., Cetintemel U., Riondato M., Upfal E., Zdonik S. B. Learning-based query performance modeling and prediction. *Proc. IEEE 28th International Conference on Data Engineering*, 2012, pp. 390–401.
2. Dutt A., Wang C., Nazi A., Kandula S., Narasayya V., Chaudhuri S. Selectivity estimation for range predicates using lightweight models. *Proc. VLDB Endowment*, 2019, vol. 12, pp. 1044–1057.
3. Halford M., Saint-Pierre P., Morvan F. Selectivity correction with online machine learning. *arXiv e-prints*, Sept. 2020, arXiv:2009.09884.
4. Hasan S., Thirumuruganathan S., Augustine J., Koudas N., Das G. Deep learning models for selectivity estimation of multi-attribute queries. *Proc. ACM SIGMOD International Conference on Management of Data*, New York, USA, 2020, pp. 1035–1050.
5. Kanellis K., Alagappan R., Venkataraman S. Too many knobs to tune? Towards faster database tuning by pre-selecting important knobs. *Proc. 12th USENIX Conference on Hot Topics in Storage and File Systems (HotStorage)*, USA, 2020.
6. Kipf A., Kipf T., Radke B., Leis V., Boncz P., Kemper A. Learned Cardinalities: Estimating Correlated Joins with Deep Learning. *arXiv e-prints*, Sept. 2018, arXiv:1809.00677.
7. Lan H., Bao Z., Peng Y. A Survey on Advancing the DBMS Query Optimizer: Cardinality Estimation, Cost Model, and Plan Enumeration. *arXiv e-prints*, Jan. 2021, arXiv:2101.01507.
8. Leis V., Radke B., Gubichev A., Mirchev A., Boncz P., Kemper A., Neumann T. Query optimization through the looking glass, and what we found running the join order benchmark. *VLDB Journal*, 2018, vol. 27, pp. 643–668.
9. Li G., Zhou X., Li S., Gao B. Qtune: a query-aware database tuning system with deep reinforcement learning. *Proc. VLDB Endowment*, 2019, vol. 12, pp. 2118–2130.
10. Marcus R., Negi P., Mao H., Zhang C., Alizadeh M., Kraska T., Papaemmanouil O., Tatbul N. Neo: a learned query optimizer. *Proc. VLDB Endowment*, 2019, vol. 12, pp. 1705–1718.
11. Poess M., Nambiar R. TPC H benchmark standard specification. URL: https://www.tpc.org/tpc_documents_current_versions/pdf/tpc-h_v2.17.1.pdf
12. Trummer I. DB-BERT: a Database Tuning Tool that “Reads the Manual”. *arXiv e-prints*, Dec. 2021, arXiv:2112.10925.
13. Van Aken D., Pavlo A., Gordon G.J., Zhang B. Automatic database management system tuning through large-scale machine learning. *Proc. ACM International Conference on Management of Data*, New York, USA, 2017, pp. 1009–1024.
14. Wu W., Chi Y., Zhu S., Tatemura J., Hacigumus H., Naughton J. F. Predicting query execution time: Are optimizer cost models really unusable? *Proc. IEEE 29th International Conference on Data Engineering*, 2013, pp. 1081–1092.
15. Zhang J., Liu Y., Zhou K., Li G., Xiao Z., Cheng B., Xing J., Wang Y., Cheng T., Liu L., Ran M., Li Z. An end-to-end automatic cloud database tuning system using deep reinforcement learning. *Proc. International Conference on Management of Data*, New York, USA, 2019, pp. 415–432.
16. Zhu R., Chen W., Ding B., Chen X., Pfadler A., Wu Z., Zhou J. Lero: A learning-to-rank query optimizer. *Proc. VLDB Endowment*, 2023, vol. 16, pp. 1466–1479.

Об авторах**Василенко Никита**

Константинович, аспирант,
Институт систем информатики им.
А. П. Ершова СО РАН,
Новосибирск, 630090, Российская
Федерация,
<https://orcid.org/0009-0003-8727-3000>

Демин Александр Викторович,

канд. физ.-мат. наук, Институт
систем информатики им. А. П.
Ершова СО РАН, Новосибирск,
630090, Российская Федерация,
alexandredemin@yandex.ru

Пономарев Денис

Константинович, канд. физ.-мат.
наук, Институт систем информатики
им. А. П. Ершова СО РАН,
Новосибирск, 630090, Российская
Федерация, ponom@iis.nsk.su

About the authors

Nikita K. Vasilenko, Postgraduate,
Ershov Institute of Informatics
Systems SB RAS, Novosibirsk, 630090,
Russian Federation,
<https://orcid.org/0009-0003-8727-3000>

Alexander V. Demin, Cand. Sci.
(Phys.-Math.), Ershov Institute of
Informatics Systems SB RAS,
Novosibirsk, 630090, Russian
Federation,
alexandredemin@yandex.ru

Denis K. Ponomaryov, Cand. Sci.
(Phys.-Math.), Ershov Institute of
Informatics Systems SB RAS,
Novosibirsk, 630090, Russian
Federation, ponom@iis.nsk.su

Поступила в редакцию / Received 06.09.2024

Поступила после рецензирования / Revised 20.11.2024

Принята к публикации / Accepted 27.11.2024