

Instance-Optimal Cloud DB Through Adaptive Indexing

Keywords: Instance-optimality, adaptive indexing, SPA

Problem: Instance-optimality in OLAP processing is an “optimal” goal that every database wants to achieve. It means for every dataset given, we aim for the minimum overhead by processing only the necessary data, generate only the necessary intermediate results. For cloud databases with multiple workers and data movements between workers, we also want to minimize the data transfer overheads and make every worker do some work until all there’s no work left.

This instance-optimality is now being achieved in a very limited way, just for acyclic join queries [2, 3, 4], where theoretical foundations were built many decades ago [1]. The focus of recent papers [2, 3, 4] is to implement the theoretical works in an efficient way. However, we see a lot of reinventions of wheels here, for example, the idea of postponing the generation unnecessary intermediate results was from 2016 in a slightly different domain [5].

In the cloud setting, data is stored in blocks, and simple min/max filtering (just storing min/max values for each column/attribute per block) is used in general without an index, as the index size can be significant for such a large cloud data. If the min/max filters do not prune a data block, this block is read into memory and scanned to fetch the actual tuples/rows. However, because of the min/max values are too broad, uninformative, no tuple can be fetched after reading into the memory, wasting the time. Therefore, an experimental work [6] has revealed that adaptive indexing can reduce the scan cost further by maintaining a lightweight index that helps min/max filtering. This offers robust scan performance. Still, the implementation of [6] is based on a simulation, not deployed into an actual open-source block-based engine/cloud database.

We aim to extend this adaptive indexing further with two novel ideas, to extend this paradigm to joins and reinforce the simple min/max filtering. Starting from a baseline cloud database with block-based execution, we integrate [6] into the baseline, and implement our ideas on top of it.

As this project gains high interests from industry and academy, we have potential collaborators and connections including Snowflake, Google, Microsoft, and other universities.

[1] Yannakakis algorithm

[2] Robust Join Processing with Diamond Hardened Joins

[3] Instance-Optimal Acyclic Join Processing Without Regret: Engineering the Yannakakis Algorithm in Column Stores

[4] Predicate Transfer: Efficient Pre-Filtering on Multi-Join Queries

[5] Efficient Subgraph Matching by Postponing Cartesian Products

[6] SPA: Economical and Workload-Driven Indexing for Data Analytics in the Cloud

DIAS: Data-Intensive Applications and Systems Laboratory

School of Computer and Communication Sciences

Ecole Polytechnique Fédérale de Lausanne

Building BC, Station 14

CH-1015 Lausanne

URL: <http://dias.epfl.ch/>



Project: In this project, the student will 1) survey and run baseline block-based execution databases, 2) implement the experimental idea of adaptive indexing of [6] into the baseline, and 3) implement our novel ideas for further improvement.

Plan:

1. Setup and run a baseline database, measure the initial performance.
2. Implement the adaptive indexing in [6] into the baseline.
3. Implement our ideas of adaptive indexing for joins and adaptive min/max filtering.

Supervisor: Prof. Anastasia Ailamaki, anastasia.ailamaki@epfl.ch

Responsible collaborator(s): Kyoungmin Kim, kyoung-min.kim@epfl.ch

Duration: 3-6 months