

CS591A1 Spring 2019 - Research Project

Title: *Design of workload-aware range-partitioning scheme for columnar layouts*

Background: One of the key design decisions of modern data analytics systems is the organization of data. For read-intensive analytic workloads, data is stored in a columnar layout on a leading column (or a group of columns) [1]. As data is typically stored in column chunks in a column layout strategy, intuitively, read queries require more partitions and updates will require fewer. For example, in a dense sorted column, read cost is computed as logarithmic cost to find the value in question, followed by sequential scan, whereas an insert may trigger moving large proportion of a dense column. On the contrary, if the stored data has no particular order, read cost subjects to a complete sequential scan, while entries are inserted at the end of the column. In such cases, range-partitioning scheme can strike a balance between the read and write performance, given a workload.

Objective: The objective of this project is to generate a workload-aware range-partitioning scheme. The workflow of the project is as the following.

- (a) Characterize the workload by studying a sample from it
- (b) Based on the distribution of the values in the dataset and the distribution of the access patterns of the workload, output a partitioning scheme that minimizes the overall cost
- (c) Try to make the solution robust, in case the actual workload behavior is different from that of the sample workload

[1] Daniel Abadi, Peter A. Boncz, Stavros Harizopoulos, Stratos Idreos, Samuel Madden. **The Design and Implementation of Modern Column-Oriented Database Systems**. Foundations and Trends in Databases 5(3): 197-280 (2013).