

Trading off Accuracy for Speed in Powerdrill

Authors : Filip Buruiana, Alexander Hall, Reimar Hofmann, Thomas Hofmann, Silviu Ganceanu, Alexandru Tudorica

Abstract : In-memory column-stores make interactive analysis feasible for many big data scenarios. PowerDrill is a system used internally at Google for exploration in logs data. Even though it is a highly parallelized column-store and uses in memory caching, interactive response times cannot be achieved for all datasets (note that it is common to analyze data with 50 billion records in PowerDrill). In this paper, we investigate two orthogonal approaches to optimize performance at the expense of an acceptable loss of accuracy. Both approaches can be implemented as outer wrappers around existing database engines and so they should be easily applicable to other systems. For the first optimization we show that memory is the limiting factor in executing queries at speed and therefore explore possibilities to improve memory efficiency. We adapt some of the theory behind data sketches to reduce the size of particularly expensive fields in our largest tables by a factor of 4.5 when compared to a standard compression algorithm. This saves 37% of the overall memory in PowerDrill and introduces a 0.4% relative error in the 90th percentile for results of queries with the expensive fields. We additionally evaluate the effects of using sampling on accuracy and propose a simple heuristic for annotating individual result-values as accurate (or not). Based on measurements of user behavior in our real production system, we show that these estimates are essential for interpreting intermediate results before final results are available. For a large set of queries this effectively brings down the 95th latency percentile from 30 to 4 seconds.

Keywords : big data, in-memory column-store, high-performance SQL queries, approximate SQL queries

Conference Title : ICDE 2016 : International Conference on Data Engineering

Conference Location : Chicago, United States

Conference Dates : September 19-20, 2016