Tastes Great, Less Filling: Low-Impact OLAP MapReduce Queries on High-Performance OLTP Systems

Xin Jia
Brown University
xin@cs.brown.edu

Andrew Pavlo
Brown University
pavlo@cs.brown.edu

Stan Zdonik
Brown University
sbz@cs.brown.edu

ABSTRACT

The previous decade saw the rise of separate, dedicated database management systems (DBMS) for online transaction processing (OLTP) and online analytical processing (OLAP) workloads [3]. The former are focused on executing short-lived, small-footprint transactions with high throughput and strong consistency guarantees. OLAP DBMSs typically target longer running and more complex queries that examine the database after it is offloaded from the front-end OLTP DBMS. For many, the latency overhead of transferring data between these two systems, as well as their administrative costs, is too onerous.

A burgeoning alternative is to use a hybrid approach where an OLTP system is able execute OLAP-style queries alongside the transactional workload [1]. This provides users the ability to execute business intelligence and other analytical queries in "real-time" (i.e., without waiting for data to be copied to the OLAP system). Such an approach has its own drawbacks, however, especially in a clustered environment. If the data is spread across multiple machines, then the OLAP queries must be executed as heavy-weight distributed transactions, which are well-known to significantly reduce the overall throughput of an OLTP system [2].

To overcome this problem, we propose a novel method of executing OLAP workloads on a parallel OLTP DBMS using the MapReduce programming model. OLAP queries are decomposed into map and reduce operations that are executed as separate transactions with either strong or weak consistency guarantees across the cluster. We implemented this model in the H-Store OLTP system [2] and evaluated it with a mixed OLTP/OLAP workload [1].

BODY

OLAP queries executed as MapReduce jobs in an OLTP DBMS achieve same latency as distributed transactions but improve throughput by 20%–50%

REFERENCES

[1] A. Kemper and T. Neumann. HyPer: A Hybrid OLTP&OLAP Main Memory Database System Based on Virtual Memory Snapshots. In *ICDE '11*, pages 195–206, 2011.

[2] A. Pavlo, C. Curino, and S. Zdonik. Skew-aware Automatic Database Partitioning in Shared-nothing, Parallel OLTP Systems. In *SIGMOD '12*, pages 61–72, 2012.

[3] M. Stonebraker, S. Madden, D. J. Abadi, S. Harizopoulos, N. Hachem, and P. Helland. The End of an Architectural Era: (It's Time for a Complete Rewrite). In *VLDB '07*, pages 1150–1160, 2007.

Volume 1 of Tiny Transactions on Computer Science
This content is released under the Creative Commons Attribution-NonCommercial ShareAlike License. Permission to make digital or hard copies of all or part of this work is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.
CC BY-NC-SA 3.0: http://creativecommons.org/licenses/by-nc-sa/3.0/.