Big Data Benchmarking: Data Intensive Research and Explorative Data Analysis

Stefan Manegold
Centrum Wiskunde & Informatica (CWI)
Amsterdam, The Netherlands

Background

Established data management benchmarks are generally based on the same assumption that are also fundamental in database managements systems (DBMSs).

- The data is already present in the database, loaded entirely from CSV files into the database before the benchmark starts, or produced inside the DBMS during the benchmark. In other words, query processing happens inside the DBMS, only.
- The entire physical design, including all relevant indexes, is optimized completely before query processing starts.
- A benchmark's workload consists of database queries that are to be evaluated literally and entirely to always produce complete and correct results.

Undoubtedly, these assumptions do hold and are required in traditional business-inspired database application scenarios.

However, in the light of Big Data, there is a clear need for interactive exploration of extremely large databases, especially in the area of scientific data management where ingestion of multiple Terabytes on a daily basis is foreseen.

To tackle these challenges, we need to rethink some of the strict requirements database systems adopted in the past. One direction proposed in [6] is that next generation database systems should *interpret queries by their intent*, rather than as a contract carved in stone for complete and correct answers. The result set should aid the user in understanding the database's content and provide guidance to continue the data exploration journey. A scientist can stepwise explore deeper and deeper into the database, and stop when the result content and quality reaches his satisfaction point. At the same time, response times should be close to instant such that they allow a scientist to *interact* with the system and explore the data in a contextualized way.

Another proposed step towards interactive explorative data analysis suggests novel approximation techniques that give precise control over runtime and quality of query answering [7]. These techniques differ from previous sampling approaches in their *bias* towards the focal point of the scientific data exploration, their *multi-layer* design, and their *adaptiveness* to shifting query workloads.

Further, workload-adaptive partial and incremental data loading from external files, index creation and optimization, and query processing techniques [4, 5, 3, 1] have been proposed and are being developed to lower the threshold and response time for a-priori unknown or rapidly changing workloads.

Requirements and Challenges

Big Data benchmark scenarios need to reflect the revise requirements in order to provide suitable means to assess and distinguish future big data management techniques that will diverge significantly from database technologies as we are used to today.

Requirements include but are not limited to:

- In addition to generating the data inside a DBMS or providing at as CSV files, data should also be available in domain-specific standardized file format, as, e.g., used remote sensing, large-scale scientific experiments or multi-media.

- Benchmark workflows need to relief the strict three-phase paradigm of data loading, index creation, query processing. Instead, benchmarks need to allows for holistic approaches where the systems decides automatically and according to the given workload when and how to load and index which and how much data in order to answer queries.
- Benchmark workloads need to be designed in such a way that they simulate a-priori unknown workloads and unpredictable workload changes [2]. However, in between changes, the workloads should be focused on specific areas of the data.
- In addition to classical performance metrics ranging from execution time and throughput to resource and energy consumption, big data benchmarks should provide additional metrics. With explorative data analysis more and more relying on iterative approaches that start with only informative and approximate answers before converging to complete results for limited areas of the data, the quality of the initial approximate answers as well as the speed of convergence towards a final complete answer need to be reflected in the benchmark results. The same holds for the convergence of adaptive query processing techniques towards state that is fully optimized for a given but a priori unknown workload [2].

References

[1] I. Alagiannis, R. Borovica, M. Branco, S. Idreos, and A. Ailamaki. NoDB: Efficient Query Execution on Raw Data Files. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*, Scottsdale, AZ, USA, May 2012.

[2] G. Graefe, S. Ideros, H. Kuno, and S. Manegold. Benchmarking Adaptive Indexing. In *Proceedings of TPC Technology Conference on Performance Evaluation & Benchmarking (TPC TC)*, Singapore, Sept. 2010.

[3] S. Idreos, I. Alagiannis, R. Johnson, and A. Ailamaki. Here are my Data Files. Here are my Queries. Where are my Results? In *Proceedings of the Biennial Conference on Innovative Data Systems Research (CIDR)*, pages 57–68, Asilomar, CA, USA, Jan. 2011.

[4] S. Idreos, S. Manegold, H. Kuno, and G. Graefe. Merging What's Cracked, Cracking What's Merged: Adaptive Indexing in Main-Memory Column-Stores. *Proceedings of the VLDB Endowment (PVLDB)*, 4(9):585–597, June 2011.

[5] M. Ivanova, M. Kersten, and S. Manegold. Data Vaults: a Symbiosis between Database Technology and Scientific File Repositories. In *Proceedings of the International Conference on Scientific and Statistical Database Management (SSDBM)*, Chania, Crete, Greece, June 2012.

[6] M. L. Kersten, S. Idreos, S. Manegold, and E. Liarou. The Researcher's Guide to the Data Deluge: Querying a Scientific Database in Just a Few Seconds. *Proceedings of the VLDB Endowment (PVLDB)*, 4(12):1474–1477, Aug. 2011.

[7] L. Sidirourgos, M. L. Kersten, and P. A. Boncz. SciBORQ: Scientific data management with Bounds On Runtime and Quality. In *Proceedings of the Biennial Conference on Innovative Data Systems Research (CIDR)*, pages 296–301, Asilomar, CA, USA, Jan. 2011.