

DIAS: Data-Intensive Applications and Systems
Laboratory
School of Computer and Communication Sciences
Ecole Polytechnique Federale de Lausanne
Building BC, Station 14
CH-1015 Lausanne
URL: <http://dias.epfl.ch>



Master/Intern/Semester Project

TPCH Implementation and Analysis on top of Shore-MT

Keywords: Shore-MT, OLAP, TPCH, Profiling

Problem: TPC-H is one of the most popular query processing benchmarks used in both industry and academia today. Shore-MT [1] is an open-source state-of-the-art storage manager, which is proved to have performance and scalability better/comparable with the existing storage managers for OLTP workloads. However, its performance under OLAP workloads, like TPC-H, has not been analyzed thoroughly up till now.

Project: This project is about implementing the TPC-H queries on top of Shore-MT storage manager and later analyzing their performance to see how Shore-MT behaves under these workloads and find potential bottlenecks.

Plan:

1. Familiarize with Shore-MT storage manager and its application layer Shore-Kits (1 week)
2. Implement the TPC-H queries at Shore-Kits (5 weeks)
3. Analyze the performance for the implemented and tested TPC-H queries (1 week)
4. Write a report and present the work (1 week)

Supervisor: Prof. Anastasia Ailamaki

Responsible person: Pinar Tozun (email to pinar.tozun@epfl.ch for contact)

Duration: 2 months

References

1. R. Johnson, I. Pandis, N. Hardavellas, A. Ailamaki, and B. Falsafi. "Shore-MT: A scalable storage manager for the multi-core era" In Proc. EDBT, 2009