

SQL parsing and execution in high-performance in-memory HTAP DBMS

Keywords: DBMS, OLTP, OLAP, HTAP

Problem: Generic SQL parser and execution presents itself a trade-off between generalization and performance of the DBMS engine. SQL parsing and execution is the first entry point in all database management systems. Parsing statements (select/inserts/updates/deletes), executing in conjunction with concurrency control (CC) protocols and handling commits and aborts have presented a challenge between flexibility of execution engine and performance of the engine itself. Furthermore, in hybrid transactional and analytical (HTAP) setting, routing analytical and transactional statements to appropriate interface while guaranteeing consistency and snapshot isolation is also part of main SQL parser.

Project: In this project, student will implement a generic transactional SQL parser and executor. Apache Calcite will be used as the frontend for SQL parsing and conversion into a query plan which will be parsed and executed by an OLTP engine. For transactional guarantees, the execution will follow the engine's concurrency control protocol and evaluate statements accordingly.

Plan:

1. Implement generic query-plan parser in OLTP engine
2. Implement DML-statements execution in OLTP engine with appropriate concurrency control (MV2PL).
3. Implement DDL-statements execution in DBMS.
4. Implement DQL-statement routing in HTAP-setting.
5. Performance analysis of standard benchmarks through SQL interface (OLTPbench).

Supervisor: Prof. Anastasia Ailamaki, anastasia.ailamaki@epfl.ch

Responsible collaborator(s): Hamish Nicholson, hamish.nicholson@epfl.ch

Duration: 6 months