



Enterprise Platform and Integration Concepts: Research Group of Prof. Dr. Hasso Plattner

Masterproject Sommerterm 2014

Code better, run faster

Tools for Performance-Driven Enterprise Application Development

Motivation

Unfolding the full potential of a columnar in-memory database and its underlying hardware requires application developers to write program flows and queries that can be efficiently executed by the database engine. This is a non-trivial task, especially for non-trivial database operations. The goal of this project is to create and evaluate intuitive tools for expressing efficient database operators at design time and supporting application developers in achieving maximum runtime performance.

Especially when it comes to outlining potential bottlenecks and understanding unexpected behavior, information about expected result set sizes, number of cache misses, core utilization and execution times is key for building highly responsive applications. Future tools shall not only reactively deliver such information on demand, but should proactively supply developers with helpful information and visualizations while coding. Making this information accessible in an easy and consumable way without having to worry much about setting up and installing additional tools for statistical evaluation and presentation of the results could significantly improve the developer experience and lead to better software.

This project aims to extend an existing browser-based interface for the prototypical open-source in-memory database Hyrise. Helping the developer to implement, test and evaluate high performance applications currently needs a lot of low-level work to increase performance. In this project, we want to integrate the implementation and tweaking, for example of stored procedures or custom operators, into the existing development environment. We will verify our concepts with the goal to write efficient algorithms for analyzing sensor data that was obtained from soccer matches of the German national team.

Project Goals

1. Familiarization with in-memory data management and hands-on experience with Hyrise
2. Design of tools to improve the development experience
3. Prototypical realization and validation using the browser-based IDE for Hyrise
4. Verification of concepts in the context of analyzing real-world sensor data of a soccer match (60 million data points)

Technology & Skills

Hyrise is the foundation and backend for the tools to be developed during the project. Technologies and languages for the frontend are intentionally not restricted and will be determined during the requirements engineering. This openness with regards to technologies requires a broad set of skills. Hence, participants should have interest in any of the following:

- Web Development (JavaScript, HTML5, Python, Ruby, ...)
- C++
- Database Technologies (In-Memory Data Management, SQL, Stored-Procedures, ...)
- User Interface Design

Contact

You are welcome to visit us in the “Villa” or reach out to one of the contacts listed below. For further information, we also invite you to an upfront meeting at room V2.16 on March 10, 2014 at 11am.

Dr. Matthias Uflacker (matthias.uflacker@hpi.uni-potsdam.de)

Martin Faust (martin.faust@hpi.uni-potsdam.de)

Keven Richly (keven.richly@hpi.uni-potsdam.de)

David Schwalb (david.schwalb@hpi.uni-potsdam.de)

Ralf Teusner (ralf.teusner@hpi.uni-potsdam.de)