



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

Master Project Summer Term 2018

C++ Low-Level Performance Engineering for Database Systems

“Premature optimization is the root of all evil”

- Donald Knuth

Motivation

Databases are only about performance. Users expect high throughput and low response times for their queries. To achieve this, both the general algorithms and the particular implementation have to be developed with performance in mind. In order to increase the performance, developers must have a good understanding on where bottlenecks are. Trying to optimize blindly wastes time and can even be detrimental.

In this project, we want to take an existing code base of our research database Hyrise, identify the bottlenecks, and improve the overall performance. This can include issues like locking and waiting, false cache line sharing, algorithmic inefficiencies, as well as inefficiencies in the binary generated by the C++ compiler.

Project Goals

1. Investigate and apply different performance profiling approaches
2. Use multiple tool chains (e.g., Intel vTune, valgrind) and evaluate their strengths and weaknesses
3. Learn to pinpoint performance bottlenecks
4. Understand how C++ code, compiler, and the underlying hardware determine execution speed
5. Learn how to act on the identified bottlenecks and how to improve the code

Technology & Skills

We will use our research database Hyrise (<https://github.com/hyrise/hyrise>) as the platform on which we will explore different profiling and optimization methods. A benchmark framework is already in place, so we do not have to spend a lot of time setting up the workspace. Continuous support by the Hyrise core development team is provided.

Prior understanding of the fundamentals of databases (e.g., from the Datenbanksysteme lecture, the Trends and Concepts online class, or the Build your own Database seminar) is expected as well as knowledge of C++.

Contact

You are welcome to visit us in the “Villa” or reach out to one of the following contacts:

Dr. Matthias Uflacker, Markus Dreseler, Jan Koßmann