Seminar: Implementation Techniques for Main Memory Database Systems

Kickoff Meeting

Prof. Dr. Jana Giceva
Jan Böttcher, M.Sc.
Dominik Durner, M.Sc.
Philipp Fent, M.Sc.
Michael Freitag, M.Sc.
Maximilian Schüle, M.Sc.

July 13, 2020

Overview

Weekly Meeting

- Monday, 16:00 - 17:30, starting October 12, 2020
- Room MI 02.09.014
- 2 presentations per meeting
- There will be an attendance log

Required Work

- Seminar paper ($\leq$ 5 pages)
- Sample implementation (C++)
- Presentation (20 minutes + 10 minutes discussion)
- Moderate one discussion (act as the "devil's advocate", you should pair up for this)

Organization

Registration through matching system

- **Register for the seminar on https://matching.in.tum.de!**

After matching: Check in with the supervisor for your preferred topic

1. Check in **soon after matching** for paper recommendations (preferences considered FCFS)
2. Check in when rough structure is planned
3. Check in when first draft is ready

Due Dates

- Structure: ca. 4 weeks prior to presentation date
- Presentation slides: 1 week prior to presentation date
- Seminar paper and finished implementation: 2 weeks after presentation date

Preliminary Topic List

Topic	Supervisor
• Scalable and Robust Latches for Database Systems (1) • Scalable and Robust Latches for Database Systems (2)	Jan Böttcher
• Towards Scalable Dataframe Systems / When sweet and cute isn't enough anymore: Solving scalability issues in Python Pandas with Grizzly • PolarFS: An Ultra-low Latency and Failure Resilient Distributed File System for Shared Storage Cloud Database	Dominik Durner
• BB-Tree: A practical and efficient main-memory index structure for multidimensional workloads • Interpolation-friendly B-trees: Bridging the Gap Between Algorithmic and Learned Indexes • External Merge Sort for Top-K Queries	Philipp Fent
• Updateable HyperLogLog Sketches • Leapfrog Triejoin: A Simple, Worst-Case Optimal Join Algorithm • White-box Compression:Learning and Exploiting Compact Table Representations	Michael Freitag
• HetExchange: Encapsulating heterogeneous CPU-GPU parallelism in JIT compiled engines (**GPU required**) • DB4ML - An In-Memory Database Machine Learning Support	Maximilian Schüle

https://db.in.tum.de/teaching/ws2021/seminarHauptspeicherdbs/