

Seminar: Implementation Techniques for Main Memory Database Systems

Kickoff Meeting

Prof. Dr. Thomas Neumann

Jan Böttcher, M.Sc.

Dominik Durner, M.Sc.

Philipp Fent, M.Sc.

Michael Freitag, M.Sc.

Alice Rey, M.Sc.

Maximilian Schüle, M.Sc.

July 7, 2021



Overview

Weekly Meeting

- Monday, 16:00 - 17:30, starting October 18, 2021
- Room MI 02.09.014
- 2 presentations per meeting
- [There will be an attendance log](#)

Required Work

- Seminar paper (≤ 5 pages)
- Sample implementation (C++)
- Presentation (20 minutes + 10 minutes discussion)

Organization

Registration through matching system

- Write an email to freitagm@in.tum.de if you are interested
 - Subject should be: [DBSeminar] Kickoff Meeting
 - Emails will be filtered based on this subject
- **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
<ul style="list-style-type: none"> ● Optimistic Lock Coupling: A Scalable and Efficient General-Purpose Synchronization Method ● On-demand-fork: A Microsecond Fork for Memory-Intensive and Latency-Sensitive Applications ● Scalable garbage collection for in-memory MVCC systems 	Jan Böttcher
<ul style="list-style-type: none"> ● Releasing Locks As Early As You Can: Reducing Contention of Hotspots by Violating Two-Phase Locking ● Rethink the Scan in MVCC Databases ● Scalable Structural Index Construction for JSON Analytics 	Dominik Durner
<ul style="list-style-type: none"> ● Small Selectivities Matter: Lifting the Burden of Empty Samples ● Contention and Space Management in B-Trees ● FLAT: Fast, Lightweight and Accurate Method for Cardinality Estimation 	Philipp Fent
<ul style="list-style-type: none"> ● FSST: Fast Random Access String Compression ● Interleaving with Coroutines: A Practical Approach for Robust Index Joins 	Michael Freitag
<ul style="list-style-type: none"> ● Filter Representation in Vectorized Query Execution ● KallaxDB: A Table-less Hash-based Key-Value Store on Storage Hardware with Built-in Transparent Compression 	Alice Rey
<ul style="list-style-type: none"> ● Automatic Optimization of Matrix Implementations for Distributed Machine Learning and Linear Algebra ● LIMA: Fine-grained Lineage Tracing and Reuse in Machine Learning Systems ● Fast CSV Loading Using GPUs and RDMA for In-Memory Data Processing (GPU required) 	Maximilian Schüle

<https://db.in.tum.de/teaching/ws2122/seminarHauptspeicherdbs/>