Seminar: Implementation Techniques for Main Memory Database Systems

Organizational Meeting

Michael Freitag, M.Sc., Prof. Alfons Kemper, Ph.D.

Garching, June 28, 2018

Overview

Weekly Meeting

- Monday, 16:00 - 17:30, starting October 22, 2018
- 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 & Due Dates

Check in via email (freitagm@in.tum.de) or personally (room MI 02.11.034)

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 sample implementation: 2 weeks after presentation date

Topics

Block 1: Cardinality & Frequent Item Estimation

- Flajolet-Martin Sketches
- Index-Based Join Sampling
- Count and Count-Min Sketches

Block 2: SIMD-Acceleration & Parallelization

- SIMD-Accelerated Regular Expression Matching
- SIMD-Accelerated Hash Tables: Linear Probing and Double Hashing
- SIMD-Accelerated Hash Tables: Cuckoo Hashing
- Main Memory Hash Join Algorithms for Multi-Core CPUs
- Parallel Array-Based Single- and Multi-Source Breath First Search

Topics

Block 3: Storage

- SLACID: Sparse Linear Algebra in a Column-Oriented In-Memory Database System
- ArrayStore: A Storage Manager for Complex Parallel Array Processing
- Data Blocks: Hybrid OLTP and OLAP on Compressed Storage

Block 4: Compression

- Database Compression on Graphics Processors (**requires NVIDIA GPU**)
- MILC: Inverted List Compression in Memory
- BitWeaving: Fast Scans for Main Memory Data Processing

Block 5: Other Topics

- Column Imprints: A Secondary Index Structure
- Result Set Serialization

https://db.in.tum.de/teaching/ws1819/seminarHauptspeicherdbs/

Michael Freitag

freitagm@in.tum.de

MI 02.11.034

Have fun!