Seminar: Modern Database Systems

Organizational Meeting

Chair of Database Systems
Chair of Data Science & Engineering
Garching, February 4, 2019

Overview

Weekly Meeting

- Thursday, 16:00 - 17:30, starting April 25, 2019
- 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 ({boettcher,radke,renen,winterch}@in.tum.de) or personally

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:

- Adaptive-Size Reservoir Sampling over Data Streams
- A Dip in the Reservoir: Maintaining Sample Synopses of Evolving Datasets
- StreamSamp
- Hash-Merge Join:A Non-blocking Join Algorithm for Producing Fast and Early Join Results
- Approximate Join Processing Over Data Streams

Block 2:

- Write Behind Logging
- BzTree: A High-Performance Latch-free Range Index for Non-Volatile Memory
- Exploiting Coroutines to Attack the "Killer Nanoseconds"

Topics

Block 3:

- Hekaton, compilation in Hekaton
- HyPer + Fork
- Silo
- Cicadia
- Efficient Window Calculation using Segment Trees
- Fast Serializable Multi-Version Concurrency Control for Main-Memory Database Systems

Topics

Block 4:

- IKKBZ Removing the `log` Factor
- Tighter Upper Bounds for Join Size Estimation
- Polynomial Heuristics for Query Optimization (bsizepp)
- Optimization of Conjunctive Predicates for Main Memory Column Stores
- Faster Plan Generation through Consideration of Functional Dependencies and Keys
- Solving the Join Ordering Problemvia Mixed Integer Linear Programming

http://db.in.tum.de/teaching/ss19/seminarModernDatabaseSystems/

Jan Böttcher

Bernhard Radke

Alexander van Renen

Christian Winter

{boettcher,radke,renen,winterch}@in.tum.de

Have fun!