Seminar: Modern Database Systems

Kickoff Meeting

Prof. Dr. Jana Giceva
Per Fuchs, M.Sc.

February 04, 2021

Overview

Weekly Meeting

- Monday, 14:00 - 16:00, starting April 26th, 2021
- in this BBB room
- 2 presentations per meeting
- There will be an attendance log and participation is part of your grade

Required Work

- Seminar paper ($\leq$ 5 pages) - 60% of the grade
- including a short description of your own idea for future research and the expected challenges
- Presentation (20 minutes + 10 minutes discussion) - 30% of the grade
- Moderate one discussion (act as the "devil's advocate", you should pair up for this) - 10% of the grade

Organization

Registration through matching system

- Write an email to per.fuchs@cs.tum.edu 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: choose a paper to present

1. Email 3 preferences soon after the matching has been finalized (preferences considered FCFS)
2. Check in when rough structure is planned
3. Check in when final draft is ready

Due Dates

- Structure: ca. 4 weeks prior to presentation date
- Presentation slides: 1 week prior to presentation date
- Seminar paper final handin: 2 weeks after presentation date

Topics List

Block 1: Distributed Systems Challenges

- A Critique of the CAP Theorem (arXiv 2015)
- Paxos made simple (SIGACT 2001)
- In search of an understandable consensus algorithm (Raft, Usenix 2014)
- Hermes: a Fast, Fault-Tolerant and Linearizable Replication Protocol (ASPLOS 2020)

Topics List

Block 2: Consistency

- Consistency, Availability, and Convergence – Mahajan et al. (TR from UT Austin, 2011)
- Don’t settle for Eventual: Scalable Causal Consistency for Wide-Area Storage with COPS – Lloyd et al. (SOSP’11)
- A comprehensive study of Convergent and Commutative Replicated Data Types – Shapiro et al. (2011)

Topics List

Block 3: Database Systems

- Anna: a KVS for Any Scale – Wu et al. (ICDE'18)
- Cloud-Native Database Systems at Alibaba: Opportunities and Challenges (VLDB 2019)
- Building an Elastic Query Engine on Disaggregated Storage. - Vuppalapati et al. (NSDI 2020)
- POLARIS: Distributed SQL Engine in Azure Synapse. - Aguilar-Saborit (VLDB 2020)
- CockroachDB: The Resilient Geo-Distributed SQL Database (SIGMOD 2020)

Topics List

Block 4: Transactions in the Cloud

- Obladi: Oblivious Serializable Transactions in the Cloud - Crooks et al. (OSDI'18)
- FaSST: Fast, Scalable and Simple Distributed Transactions with Two-Sided (RDMA) Datagram RPCs - Kalia et al. (OSDI'16)
- Highly Available Transactions: Virtues and Limitations – Bailis et al. (VLDB'14)
- No compromises: distributed transactions with consistency, availability, and performance - Dragojevic et al (SIGMOD'15)

Topics List

Block 5: Serverless Data Processing

- Shuffling , Fast and Slow : Scalable Analytics on Serverless Infrastructure - Pu et al. (NSDI'19)
- Lambada: Interactive Data Analytics on Cold Data Using Serverless Cloud Infrastructure - Mueller et al. (SIGMOD'20)
- Starling: A Scalable Query Engine on Cloud Functions Matthew - Perron (SIGMOD'20)
- Magnet: push-based shuffle service for large-scale data processing (VLDB 2020)

https://db.in.tum.de/teaching/ss21/seminarModernDatabaseSystems/