

# Kick-Off

# Course Goals

## Primary goals:

- Learn how to write a scientific paper
- Learn how to give a scientific talk
- Structure ideas from multiple research papers

## Secondary goals:

- Practice to read scientific papers
- Understand modern cloud database system architectures
- Have interesting discussions with your peers

# What you get out of this course

In this course you will get:

- At the beginning of the seminar we will give three introduction lectures:
  - Introduction to distributed cloud databases
  - How to write a scientific text
  - How to give a scientific talk
- Two peer reviews for your paper
- Presentation feedback from your peers

# Prerequisites

Hard prerequisite:

- Introduction to Fundamentals of Databases (IN0008) or equivalent

Beneficial previous knowledge:

- Database Systems on Modern CPU Architectures (IN2118)
- Query Optimization (IN2219)
- Most courses from Prof. Neumann, Prof. Giceva, Prof. Kemper, and Prof. Leis

# Deliverables

During the course each of you will create:

- A report proposal that summarizes your reports topic and focus (<1 page)
- A research report **draft** (5-8 pages + references)
- 2 peer reviews for your fellow students (<1 page each)
- A research report (5-8 pages + references)
- A pre-session protocol that shows your preparation for the presentation. (<1 page)
- A presentation (13-15 minutes)
- Meaningful contribution to the panel discussion in your presentations session

# Grading

Rough estimate of grading contributions:

- $\approx 45\%$  Report
- $\approx 25\%$  Presentation
- $\approx 10\%$  Report Proposal
- $\approx 10\%$  Peer reviews
- $\approx 10\%$  Pre-session protocol and panel discussion

**This is subject to change!**

# Topics (1)

Preliminary list of topics:

- **Neojoin:** How to do distributed joins fast: NeoJoins and evolution
- **Flowjoin:** How to handle skew in distributed joins
- **Snowflake:** A distributed OLAP cloud Unicorn
- **Firebolt:** Building a good system fast
- **Redshift:** How Amazon scales for larger datasets
- **Polaris:** How Microsoft scales virtually infinitely
- **Trino:** How Meta analyzes large data from many sources
- **Velox:** Meta's faster Trino
- **Lambada:** Completely serverless
- **Client-Server:** Hand off computation to clients

## Topics (2)

- **Anyblob:** Single node storage decoupling
- **Microsoft's Optimizers:** Early ideas SQL server PDW and heuristics
- **MemSQL:** Query Optimization is hard
- **Vertica:** uses vertical fragmentation
- **Eigen:** Alibaba's cluster scheduling
- **Key Recommendation:** for better schemata
- **EdgeFrame:** Worst-Case Optimal Joins and Graph Pattern Matching
- **Scyper:** single node + log replication is fast
- **Socrates:** Scaling OLTP
- **Noria:** fast analytics in OLTP systems
- **RAV Scheduling:** NUMA-aware query plan deployment
- **Self-Tuning Query Scheduling:** Scheduling for analytical workloads
- **CHARM:** Chiplet-aware scheduling for OLAP

Your ideas for related topics are very welcome!

# Timeline

Preliminary timeline:

- Mo 20.04.2026 introduction lecture 01
- Mo 27.04.2026 introduction lecture 02 | **submit** topic preferences
- Mo 04.05.2026 introduction lecture 03
- Mo 11.05.2026 ∅ | **submit** report proposal
- Mo 18.05.2026 ∅
- ~~Mo 25.05.2026~~ Pfingstferien
- Mo 01.06.2026 ∅
- Mo 08.06.2026 ∅
- Mo 15.06.2026 ∅ | **submit** report draft
- Mo 22.06.2026 presentation session 01
- Mo 29.06.2026 presentation session 02 | **submit** peer reviews
- Mo 06.07.2026 presentation session 03
- Mo 13.07.2026 presentation session 04 | **submit** final report
- **Attendance to all sessions in presence is mandatory**

# Organization

- Attendance to all sessions in presence is mandatory
- You have to write your paper in LaTeX using **our template**
- Register on our gitlab server with your TUM email to get access
- Communication will take place using our **Mattermost** instance

# Matching

Register for the course through the matching platform (<https://matching-in.cit.tum.de/>)

# Contact

- Website: <http://db.in.tum.de/teaching/ss26/seminarModernDatabaseSystems>
- Team (preferred): [mds.db@cit.tum.de](mailto:mds.db@cit.tum.de)
- Maximilian Rieger: [riegerm@in.tum.de](mailto:riegerm@in.tum.de)
- Christoph Rotte: [christoph.rotte@in.tum.de](mailto:christoph.rotte@in.tum.de)