

Architecture and Implementation of Database Systems (Summer 2018)

Jens Teubner, DBIS Group
`jens.teubner@cs.tu-dortmund.de`

Summer 2018


A Few Words About Me

Jens Teubner

DBIS Group (LS 6)

Otto-Hahn-Strasse 14, Room 333

jens.teubner@cs.tu-dortmund.de



1996–2001 Diploma in Physics, U Konstanz

2001–2005 Research assistant, DBIS Group, U Konstanz

2005–2007 Research assistant, Database Group, TU München

Oct 2006 PhD in Computer Science (XML query processing)

2007–2008 Postdoc, IBM T. J. Watson Research Center, NY, USA

2008–2013 Senior Researcher, Systems Group, ETH Zurich

since 4/2013 Full Professor, DBIS Group, TU Dortmund University


Topic: Database systems on modern computing hardware

What this Course is About

TPC-E Industry Benchmark (modeling a brokerage firm)

- 2.3M customers; 11.5M accounts; 2B holdings; 40B trades
- 18.5 TB initial database size


Current record: **11,059 tpsE (110,590 transactions/sec total)**!



Numerical Quantities Summary				
Reported Throughput: 11,058.99 tpsE	Configured Customers: 5,750,000			
Response Time (in seconds)	Minimum	Average	90 th Percentile	Maximum
Broker-Volume	0.01	0.01	0.03	1.02
Customer-Position	0.01	0.01	0.02	2.96
Market-Feed	0.01	0.01	0.02	2.93
Market-Watch	0.01	0.01	0.02	1.00
Security-Detail	0.01	0.01	0.01	0.82
Trade-Lookup	0.01	0.05	0.08	1.03

**How can databases provide such throughput,
yet be transaction-safe?**

Architecture of a DBMS / Course Outline



Lecture

- Mondays, 8–10h, Room OH 12, E.003
- Wednesdays, 10–12h, Room OH 12, E.003
- Course website: <http://dbis.cs.tu-dortmund.de/cms/en/teaching/ss18/arch-dbms/>
Please visit this website **regularly**. I will frequently post new information during the semester.

Exercises

- Mondays, 14–16h, Room OH 12, E.003
 - Tuesdays, 8–10h, Room OH 12, 3.031
- } starting **next week!**
- TA: Michael Mller (michael3.mueller@tu-dortmund.de)
 - Register via **AsSESS** by **Fri 13, 2018**. (Registration will open shortly.)
 - Participation is not mandatory, but **highly** recommended.

There will be a **written exam**.

- Thursday, July 26, 2018; 11:00h-12:30h; Room HG 2, HS5
- Thursday, September 20, 2018; 13:30-15:00h; Room HG 2, HS3
- duration: 90 minutes
- More information during the semester

- I'd like to make this course highly **interactive**.
 - Please speak up, discuss, ask questions!
- The material we discuss is relevant in **practice**.
 - We'll provide practical examples and exercises.
 - You get most out of this course by actively participating in (and solving) the **exercises**.

- Raghu Ramakrishnan and Johannes Gehrke. *Database Management Systems*. McGraw-Hill.
- Philip M. Lewis, Arthur Bernstein, and Michael Kifer. *Databases and Transaction Processing*. Addison-Wesley.
- Alfons Kemper and André Eickler. *Datenbanksysteme: Eine Einführung*. Oldenbourg Verlag.
- Dennis Shasha and Philippe Bonet. *Database Tuning*. Morgan-Kaufmann.
- ... or pick your own favourite in the library — most text books about advanced database topics will do.

I'll also sometimes point to specific research papers and you're welcome to search for additional background reading. Search engines like Google Scholar are your friend.

You might also be interested in...

... the course **Data Processing on Modern Hardware**

- Winter Semester
- Many similar concepts, geared to exploit the characteristics of modern computing hardware.

... doing your **Master Thesis** with me

- Chance to participate in the forefront of research.
- Many ways to join, approach me if you're interested.