Advanced Topics in Databases

Gunter Saake
David Broneske, Gabriel Campero Durand, Bala Gurumurthy,
Andreas Meister, Marcus Pinnecke, Roman Zoun

OTTO VON GUERICKE UNIVERSITÄT MAGDEBURG

INF FAKULTÄT FÜR INFORMATIK

Aim of the Course

- Familiarize students with current developments in database research
- Topics chosen:
  - First solutions currently making their way into database management systems and applications → practical relevance
  - Solutions not yet fully developed and where open problems still exist → research relevance
- Possible starting points for scientific work, e.g. master thesis, position in academia, Ph.D. thesis, industry R&D, etc.

Audience & Prerequisites

What you should need to know already

- Database introductory course (e.g., Database Concepts)
- Recommended: Database implementation techniques

What you'll learn in this lecture

- Impact of modern hardware on main-memory database systems
  - Database operators
  - Query optimization
  - Index structures
- HTAP database management systems
- AI techniques for data management
- Analytics in document-stores

Motivation for this Course

PART I

Yesterday's DBMS Landscape

ORACLE

IBM DB2

Microsoft SQL Server

TERADATA

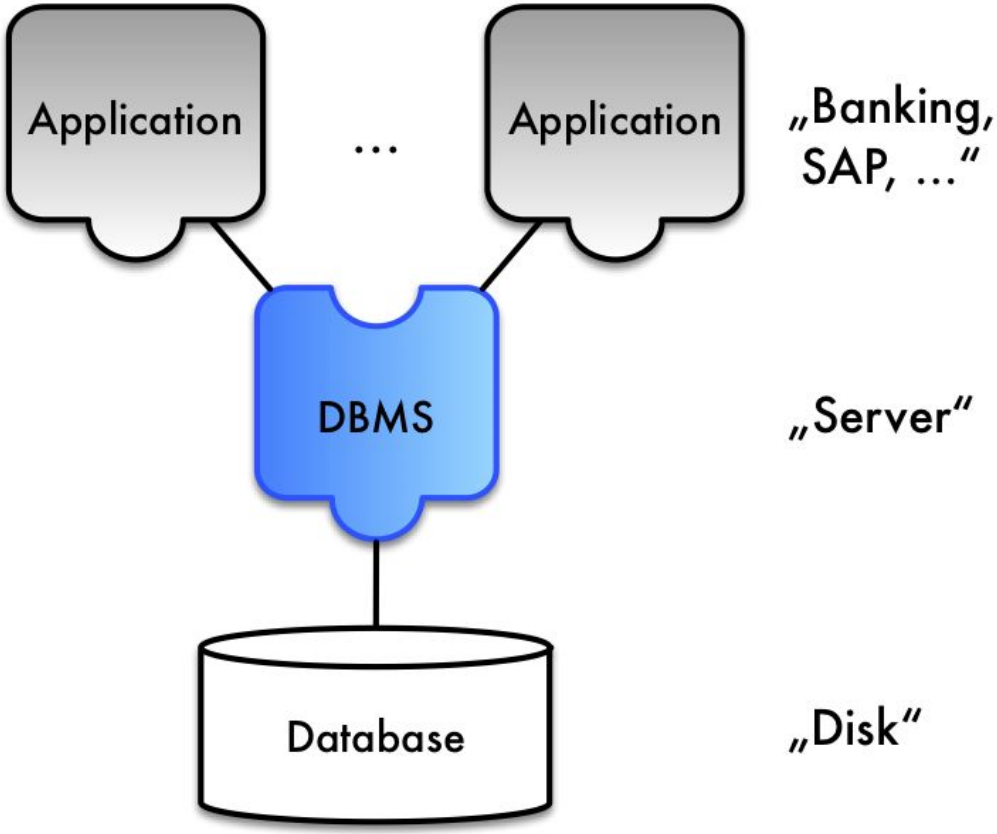

Yesterday's DBMS Hardware

Picture taken from [1]

Small main memory

Picture taken from [2]

Disk-based systems

Assumptions of Yesterday's DBMS's

- Capacity of main memory <1% of the stored data
- Fixed block size based on the transfer unit between disks and main memory
- Central scheduler to schedule transactions
- No redundant data storage in main memory
- Pipelining is always beneficial (no storage of intermediate results)
- Compiling of SQL for one processor architecture → Reuse of compiled plan

Data Platforms Landscape Map – February 2014

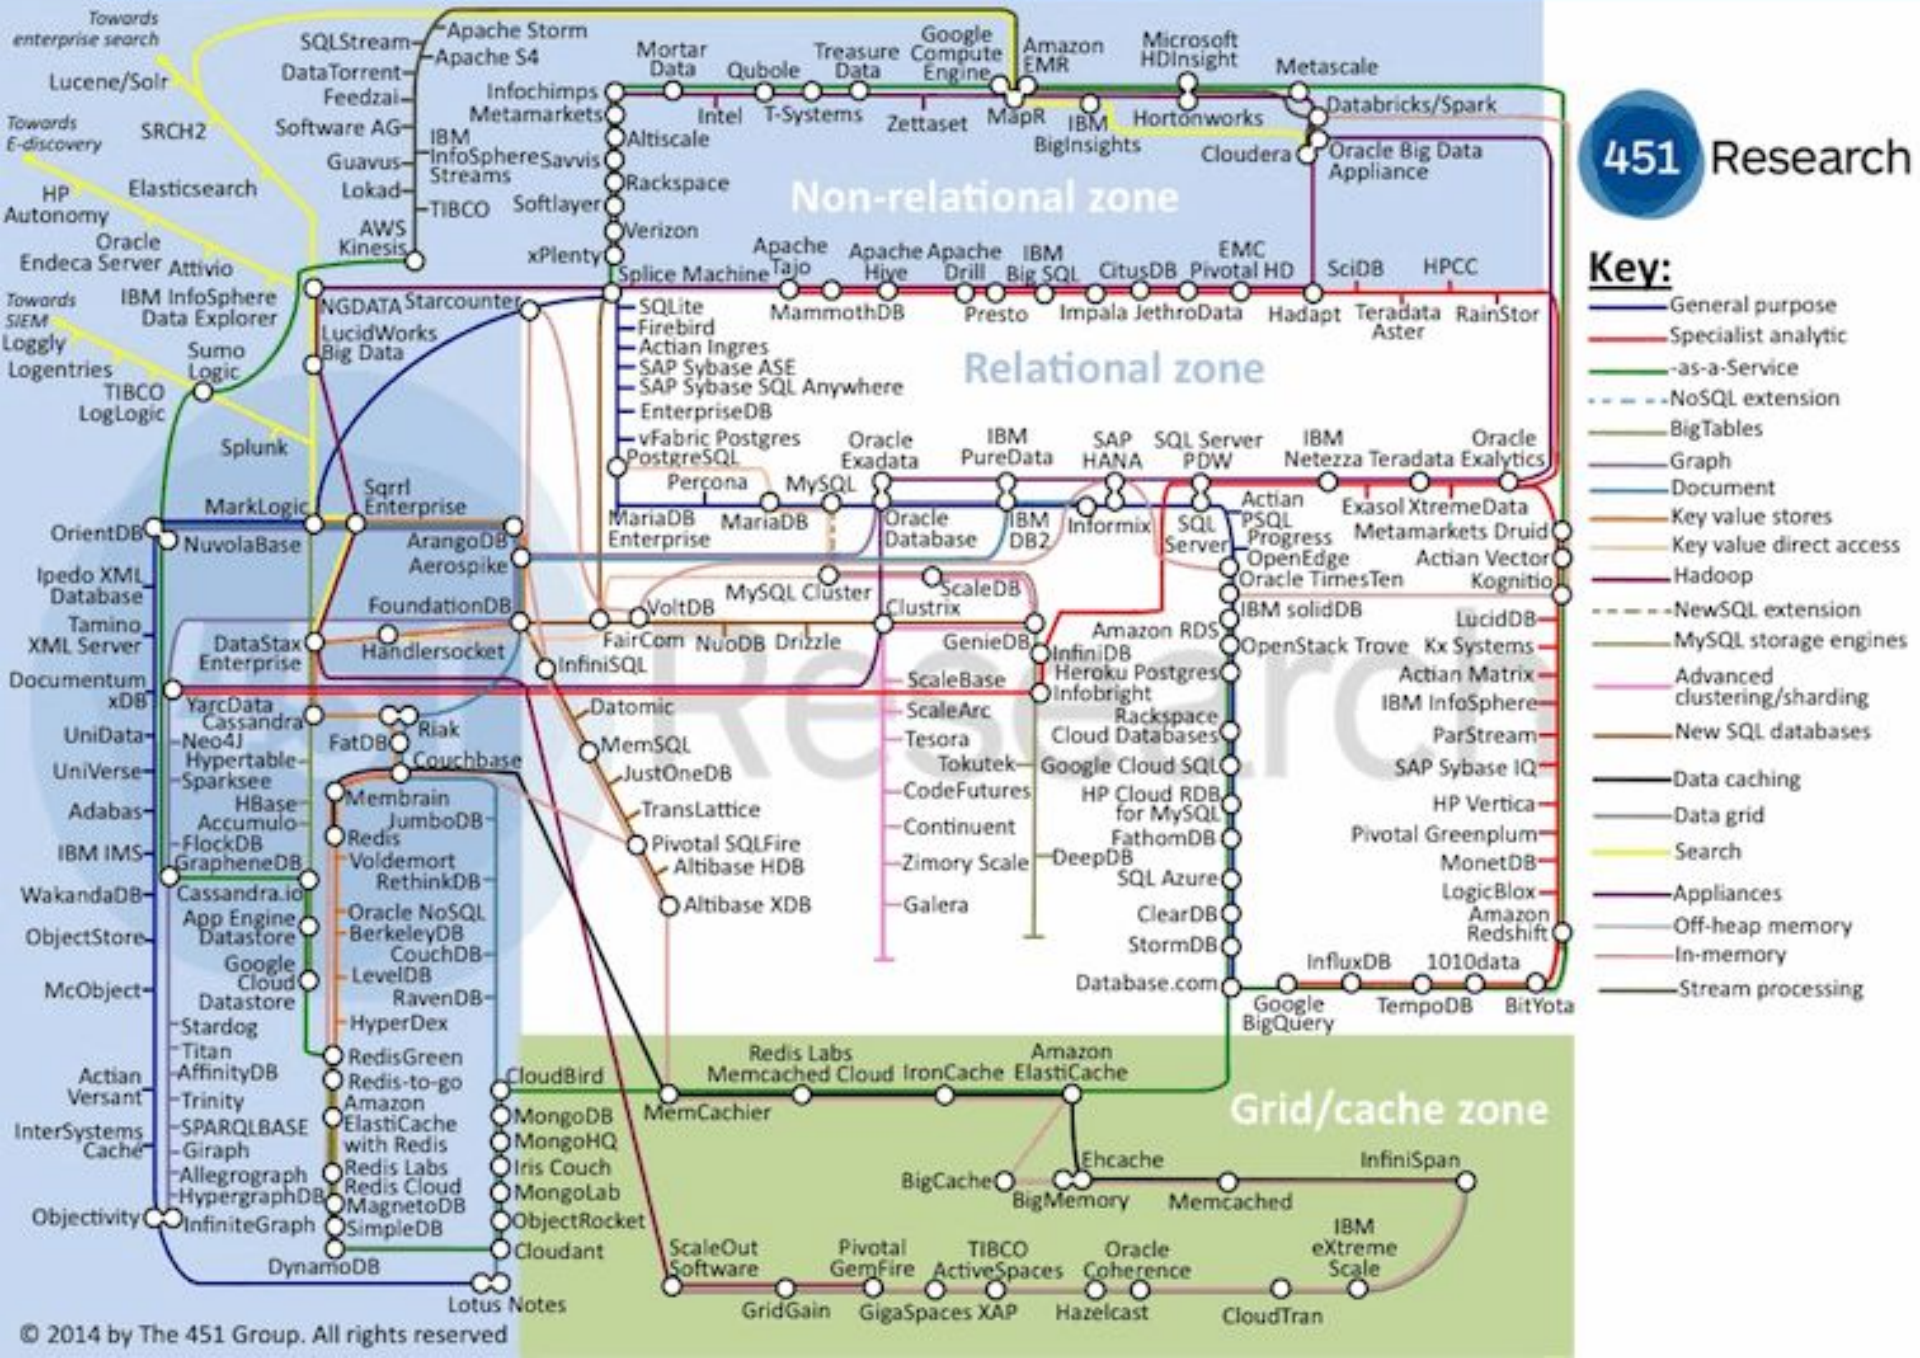

Today's Hot Topics

Today's DBMS Infrastructure

- Large-scale query/data flow engines
- Stream-based query engines
- **In-Memory Storage**
- MPP DBs, cloud EDWs, **GPU DBs**
- NewSQL: Large-scale OLTP and **HTAP** DBs
- NoSQL: Column-families, graph data, key-value stores, **documents**, time series, etc.
- Specialized data transformation & integration tools

mattturck.com/bigdata2018

Final 2018 version, updated 07/15/2018

Today's DBMS Analytics

- Statistical analysis and Data science workloads backed by DBs
- Interactive visual data exploration & BI tools
- Specialized ML systems with their own data solutions
- Search engines
- Web, Commerce, Social and Log analytics
- Speech and NLP

mattturck.com/bigdata2018

Final 2018 version, updated 07/15/2018

Today's DBMS Hardware

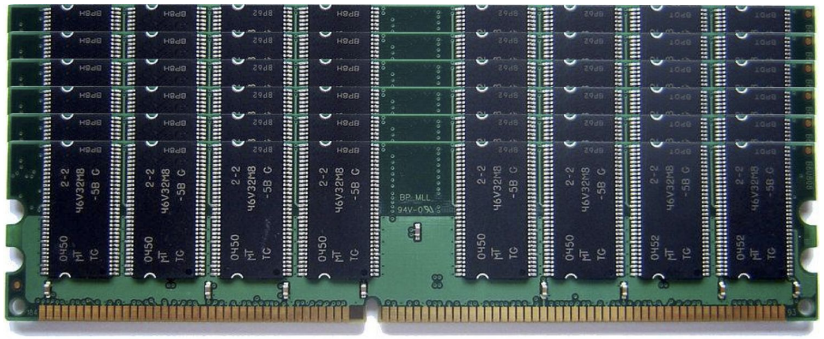

Picture taken from [1]

Large main memory

Picture taken from [3]

Multi-core CPUs

Picture taken from [4]

Solid state disks

Picture taken from [5]

Co-processors

Future DBMS's

- Capacity of main memory <1% of the stored data
  - DB in main memory
- Fixed block size based on the transfer unit
  - Direct access of data on all devices
- Central scheduler to schedule transactions
  - Which processor should do the job?
- No redundant data storage in main memory
  - Redundant data at co-processors
- Pipelining is always beneficial
  - Co-processors like GPUs support massive parallelism
- Reuse of compiled plan
  - Load-balancing between co-processors requires different plans

The Goals of a ”Databaser”

The Goals of a ”Databaser”

Picture taken from [6]

- Performance

The Goals of a ”Databaser”

Picture taken from [6]

- Performance
- Performance

The Goals of a ”Databaser”

Picture taken from [6]

- Performance
- Performance
- Performance

The Goals of a ”Databaser”

Picture taken from [6]

- Performance
- Performance
- Performance

How can we achieve more performance?

Are DBMSs written for yesterday's hardware efficient on today's hardware as well?

"30 years of Moore's law has antiquated the disk-oriented relational architecture for OLTP applications"

[Stonebraker et al.]

Data Access – Yesterday's Bottleneck

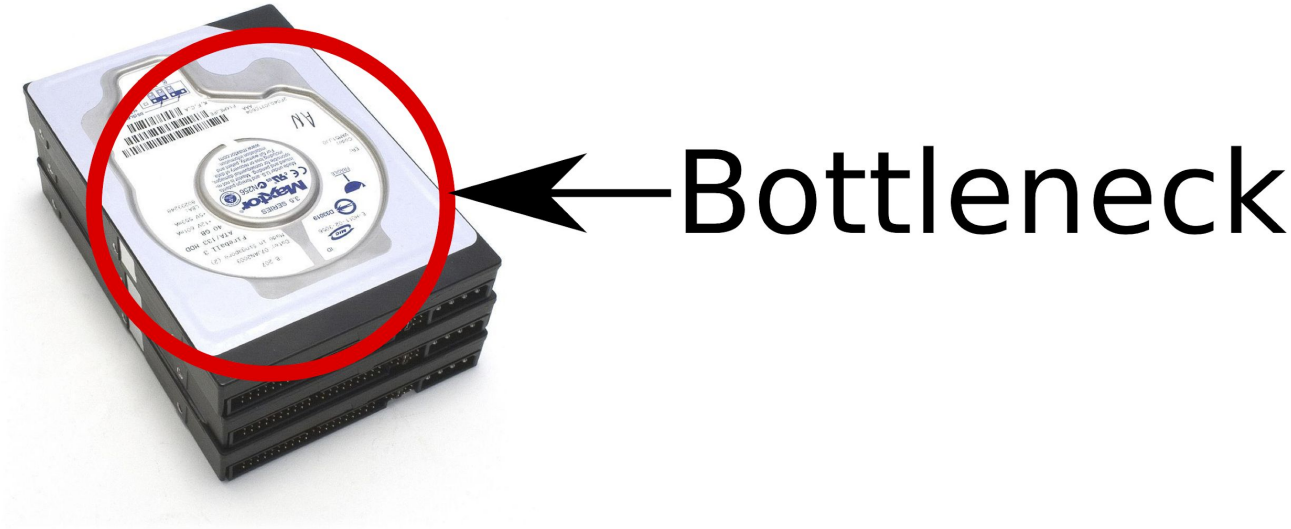

Data Access – Today's Bottleneck

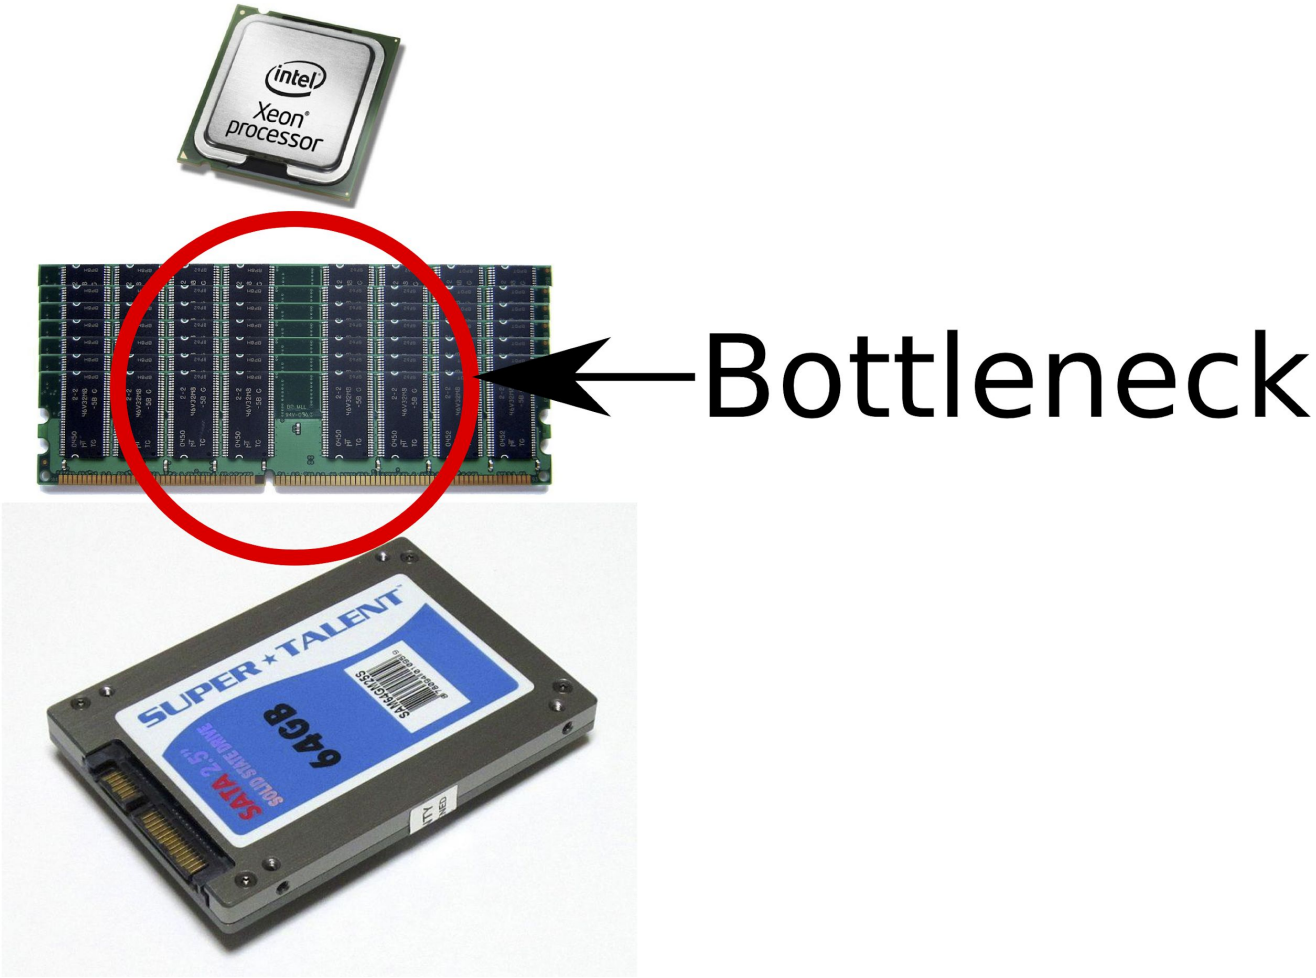

The World of Co-Processors

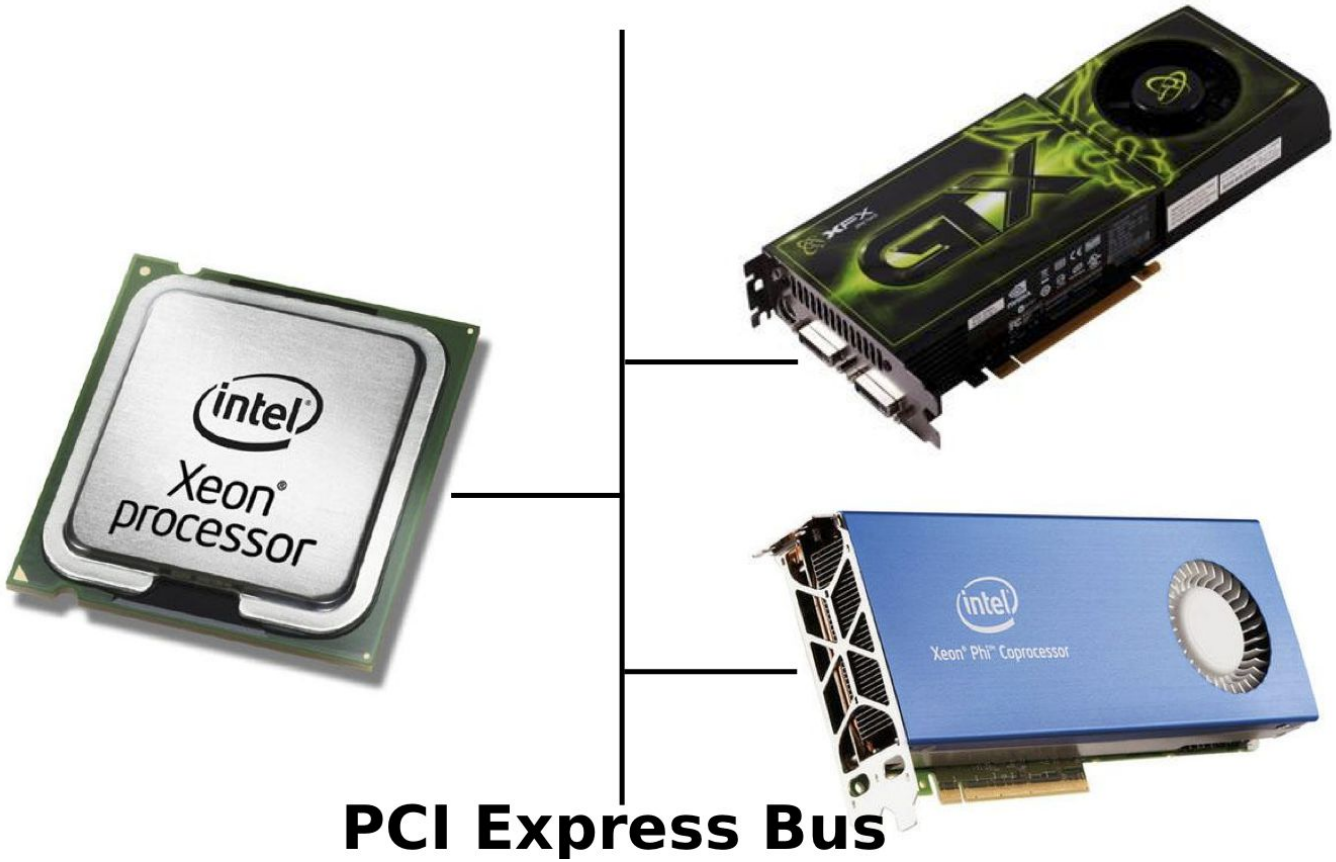

Picture taken from [7]

What do we have to change in DBMSs' architecture to **exploit new hardware capabilities** and to **meet tomorrow's challenges** and **applications**?

Topic Outline

PART I

Topic Categorization

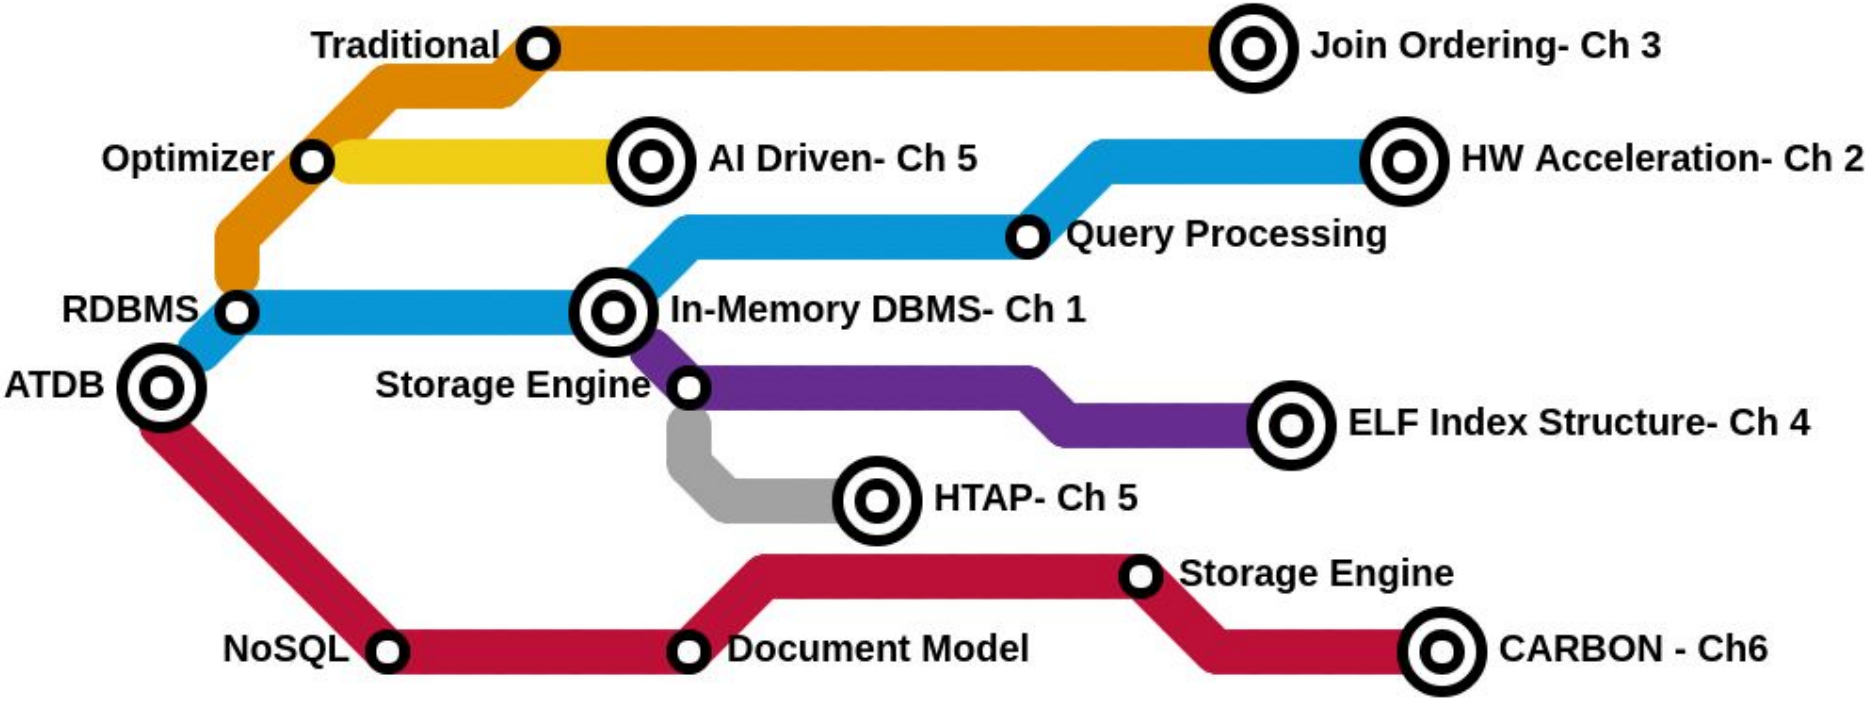

Main-Memory Database Systems

2019/April/05

- **Computer and Database Systems Architecture**
  Changes in hardware and their implications for database systems
- **Cache Awareness**
  How do caches work and how to optimize for them?
- **Processing Models**
  How do database systems execute an operation on a number of tuples?
- **Storage Models**
  How to store a two-dimensional table in one-dimensional memory?

Parallel Join Ordering

2019/April/26

- **Query Processing**
  Overview of the process of query processing
- **Join ordering**
  Overview of join ordering
- **Dynamic programming for join ordering**
  Discussion about sequential and dynamic programming variants

Hardware-Sensitive DBMS Operations

2019/May/10, 2019/May/17

- **Hardware in DBMS**
  Overview on different eras of H/W evolution and their capabilities
- **CPU - Code Optimization**
  Introduction to implementing hardware sensitive DBMS operations
- **GPU Accelerated Processing**
  Introduction to GPU architecture and kernel-based execution

Index Structures for Main-Memory Database Systems

2019/May/24

- **Query Processing Basics**
  Recap about query optimizer and selections
- **Accelerated Full-Table Scans**
  Tuning scans to the underlying hardware
- **Tree-Based Index Structures for Main Memory**
  Hardware-sensitive tree-based index structures optimized for SIMD and cache consciousness

HTAP Data Management

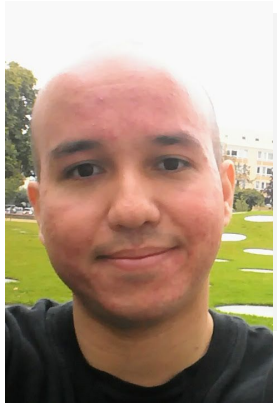

TBD

- **DBMS Design for Main-Memory OLTP**
  Overview about organization choices, OLTP indexes, versioning
- **Design Choices for HTAP**
  How do HTAP systems balance OLAP and OLTP designs?
  Illustrations from production DBMSs
- **Beyond Static HTAP Designs**
  How can databases automatically adapt to shifting workloads?

Physical Design for Document Store Analytics

2019/June/07
2019/June/14
2019/June/21

- **Document Data Model and Document Stores**
  Get in touch with JSON, MongoDB, CouchDB, and what it means
- **Document Store Storage Engine Internals**
  MongoDB/WiredTiger & CouchDB storage internals incl. records
- **Columnar Binary-Encoded JSON (Carbon) Archives**
  Get conceptual (and low-level technical) insights into our research
- **Overview on Current State and Your Points to Join**
  Get an overview on open projects (thesis, individual projects,...)

AI Techniques for Data Management

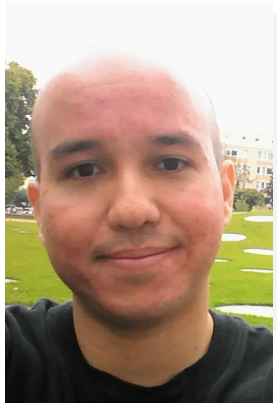

TBD

Fathers of the Deep Learning Revolution Receive ACM A.M. Turing Award

Bengio, Hinton and LeCun Ushered in Major Breakthroughs in Artificial Intelligence

ACM named Yoshua Bengio, Geoffrey Hinton, and Yann LeCun recipients of the 2018 ACM A.M. Turing Award for conceptual and engineering breakthroughs that have made deep neural networks a critical component of computing. Bengio is Professor at the University of Montreal and Scientific Director at Mila, Quebec's Artificial Intelligence Institute; Hinton is VP and Engineering Fellow of Google, Chief Scientific Adviser of The Vector Institute, and University Professor Emeritus at the University of Toronto; and LeCun is Professor at New York University and VP and Chief AI Scientist at Facebook.

- **How can developments from ML (machine learning) be used for next-gen database optimization problems?**
  Introduction to the nascent field of ML for data management
  Overview of core problems being tackled
  Examples of applications
- **Background on ML techniques gaining interest**
  Introduction to deep reinforcement learning

Organization

PART I

Tutor

Andreas Meister
PhD student
andreas.meister@ovgu.de

Organization

13 Lectures (each with an exercise sheet)

New exercise sheets: on Friday

Begin of exercises: from 2019/April/10 to 2019/July/03

12 Exercise Sheets

Registration to tutorials: Groups of 4 students until 2019/April/12

We expect you to be prepared before a tutorial starts.

Questions

Ask your fellow students first > then your tutor > then the main organizer > then the professor

Points & Assignments

- Exercises are optional, but recommended for being successful in the exam
  - Presenting task by task
  - Discussing student solutions and alternative solutions
  - Short introductory exercise at 2019/April/10
- Each student team has to submit and successfully solve **2** out of 4 programming tasks
- Programming tasks will be presented in end of April (including registration for it)
- Limited amount of teams per task!
- Final submission: **2019/July/05**

Programming Tasks

1. **Extending Main-Memory Index Structures with Special Selection Capabilities**
   C/C++ Framework
2. **Improving a Deep Reinforcement Learning Index Advisor**
   Horizon Framework for Deep Reinforcement Learning, PostgreSQL
3. **Single Column Selection in a Interpretation-Based System**
   C/C++ framework
4. **Accelerating Analytics in CARBON**
   ANSI C, CARBON Framework

Additional Material

Elf code repository

- Our main-memory index structure for multi-column selection predicates
- https://git.iti.cs.ovgu.de/dbronesk/ICDE-elf

Libcarbon code repository

- A C library for creating, modifying and querying Columnar Binary-Encoded JSON (Carbon) files
- http://github.com/protolabs/libcarbon

Web Resources

- [1] http://commons.wikimedia.org/wiki/File:RAM_module_SDRAM_1GiB.jpg
- [2] http://commons.wikimedia.org/wiki/File:Hard_disks.jpg
- [3] http://www.flickr.com/photos/25757823@N07/2719552544
- [4] http://commons.wikimedia.org/wiki/File:Super_Talent_2.5in_SATA_SSD_SAM64GM25S.jpg
- [5] http://commons.wikimedia.org/wiki/File:Gtx260.jpg
- [6] http://commons.wikimedia.org/wiki/File:Travis_Race_car.jpg
- [7] http://www.flickr.com/photos/denieseclariz/7412854696

Summary

Andreas Meister

http://www.dbse.ovgu.de/Lehre/ATDB.html

andreas.meister@ovgu.de

Have Fun and Good Luck!