SCIENCE
PASSION
TECHNOLOGY

Data Management
01 Introduction and Overview

Matthias Boehm

Graz University of Technology, Austria
Computer Science and Biomedical Engineering
Institute of Interactive Systems and Data Science
BMVIT endowed chair for Data Management

Last update: Mar 02, 2020

ISDS

Announcements/Org

- **#1 Video Recording**
  - Link in TeachCenter & TUbe (lectures will be public)

- **#2 Course Registration** (as of Mar 02)
  - Data Management VO: **457**
  - Data Management KU: 451 ☺
  - Databases VU: **89**

Total: **546**

- **#3 Bac::Mas Thesis Fair** (Mar 5, 10am-1pm, INF 25d Foyer)
  - STV-organized fair for open bachelor/master topics at institutes

- **#4 CS Talks x7** (Mar 10, 5pm, Aula Alte Technik)
  - Claudia Müller-Birn (Freie Universität of Berlin)
  - Title: Collaboration is Key – Human-Centered Design of Computational Systems

Agenda

- **Data Management Group**
- **Course Motivation, Goals, and Outline**
- **Course Organization and Logistics**
- **History of Data Management**

Data Management Group

About Me

- **09/2018 TU Graz**, Austria
  - BMVIT endowed chair for data management
  - **Data management for data science**
    (ML systems internals, end-to-end data science lifecycle)
- **2012-2018 IBM Research – Almaden**, USA
  - Declarative large-scale machine learning
  - Optimizer and runtime of **Apache SystemML**
- **2011 PhD TU Dresden**, Germany
  - Cost-based optimization of integration flows
  - Systems support for time series forecasting
  - In-memory indexing and query processing

https://github.com/tugraz-isds/systemds

DB group

Data Management Courses

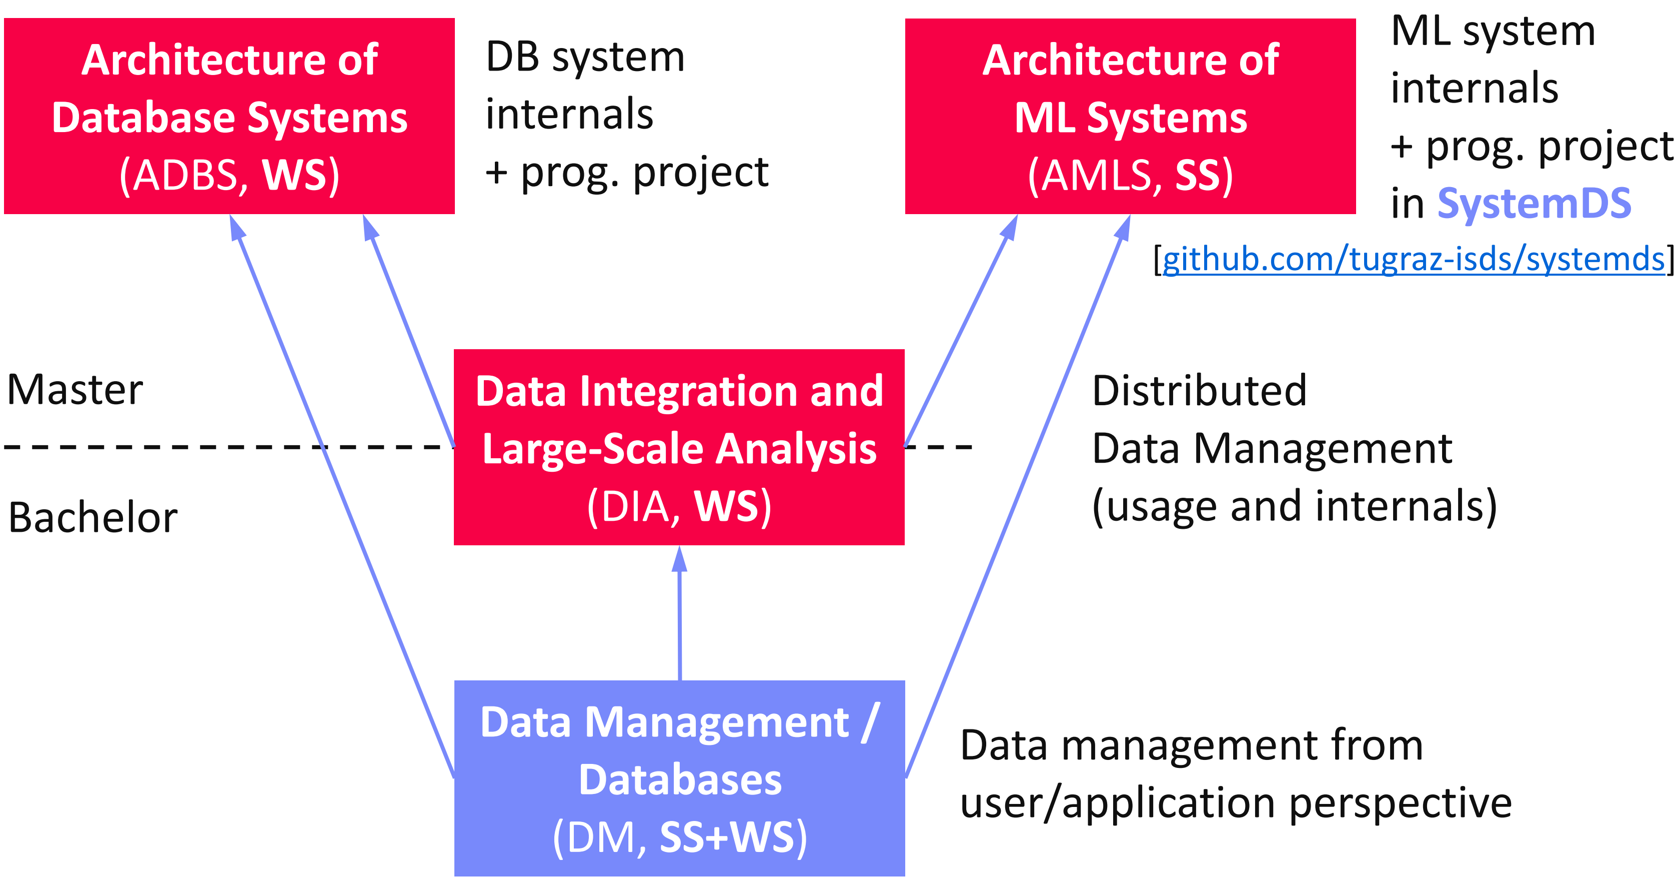

Course Motivation, Goals, and Outline

Database Systems and Modern Data Management

Definition and Impact

- **Def: Database System**
  - Overall system of DBMS + DBs
  - DBMS: Database Management System (SW to handle DBs)
  - DBs: Database (data/metadata collection of conceptual mini-world)
  - Note: DB also a short for DBS/DBMS

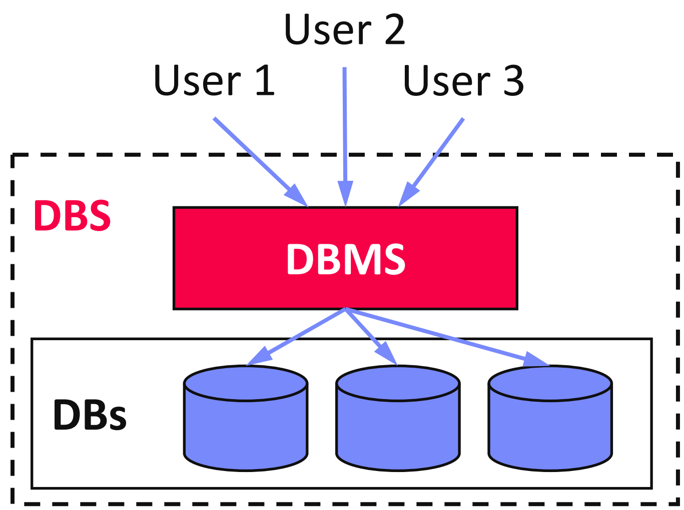

- **Importance in Practice**
  - Market Volume: **10-100B $US**
  - Foundation of many applications in various domains

[Marianne Winslett: Bruce Lindsay speaks out: […]. SIGMOD Record 34(2), **2005**]

“Relational databases are the foundation of western civilization”

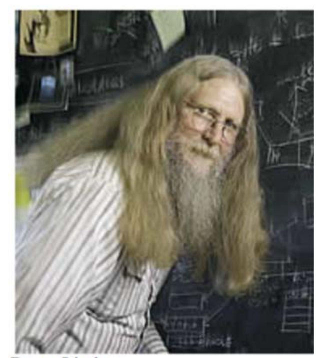
Bruce Lindsay

Motivation Database Systems

- **Application development and maintenance costs**
  - **Declarative queries** (what not how) and **data independence**
  - Efficient, correct, and independent data organization, size, access
- **Multi-user operations and access control**
  - Synchronization of **concurrent user queries and updates**
  - Enforce access control (e.g., permissions on tables, views)
- **Consistency and data integrity**
  - Eliminates **redundancy** and thus, enforces consistency
  - Enforces **integrity constraints** (e.g., semantic rules)
- **Logging and Recovery**
  - Recovery of consistent state after HW or SW failure
- **Performance and Scalability**
  - **High performance** for large datasets or high transaction throughput
  - Scale to large datasets with low memory requirements

Goals

- **Course Goals**
  - **A:** Understanding of **database systems** (from user perspective)
  - **B:** Understanding of **modern data management** (from user perspective)

	INF.01017UF (VO) Data Mgmt.	INF.02018UF (KU) Data Mgmt.
706.010 (VU) Databases	Part A 9 Lectures	Part A 3 Exercises
3(2) ECTS	Part B 3 Lectures	Part B 1 Exercise
	3 ECTS	1 ECTS

- **Meta Goals**
  - **Understand**, **use**, **debug**, and **evaluate** data management systems
  - Awareness of **system alternatives** and their **tradeoffs**
  - Fundamental concepts as basis for advanced courses and other areas

Part A: Database System Fundamentals

- **01 Introduction and Overview** [Mar 02]
- **02 Conceptual Architecture and Design** [Mar 09]
- **03 Data Models and Normalization** [Mar 16]

→ Exercise 1: Data Modeling **[Mar 31]**

- **04 Relational Algebra and Tuple Calculus** [Mar 23]
- **05 Query Languages (SQL, XML, JSON)** [Mar 30]
- **06 APIs (ODBC, JDBC, OR frameworks)** [Apr 20]

→ Exercise 2: Queries **[Apr 28]**

- **07 Physical Design and Tuning** [Apr 27]
- **08 Query Processing** [May 04]
- **09 Transaction Processing and Concurrency** [May 11]

→ Exercise 3: Tuning **[May 19]**

Part B: Modern Data Management

- **10 NoSQL (key-value, document, graph, time series)** [May 18]
- **11 Distributed file systems and object storage** [May 25]
- **12 Data-parallel computation (MapReduce, Spark)** [May 25]
- **13 Data stream processing systems** [Jun 08]

→ Exercise 4: Spark **[Jun 16]**

- **14 Q&A and exam preparation** [Jun 15]
- **Final written exam** [TBD; e.g., Jun 22, Jun 29]

Course Organization

Basic Course Organization

- **Staff**
  - **Lecturer:** Univ.-Prof. Dr.-Ing. Matthias Boehm, ISDS
    Assistant Lecturer: M.Tech. Arnab Phani, ISDS

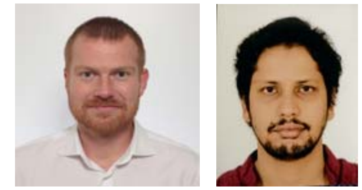

  - **Teaching Assistants:**
    Alina Herderich, Dardan Dermaku, Olga Ovcharenko, Oliver Nikolic, Melanie Willfurth, Paul Mirtl

- **Language**
  - Lectures and slides: **English**
  - Communication and examination: **English**/**German**

- **Course Format**
  - DM VO + KU 2/1 (**3+1 ECTS**), DB VU 1/1 (**3(2) ECTS**)
  - **Weekly lectures** (**start 4.10pm**, including **Q&A**), **attendance optional**
  - **4/3 exercises** (introduced in lecture) as individual assignments

Course Logistics

- **Communication**
  - **Informal language** (first name is fine)
  - Please, **immediate feedback** (unclear content, missing background)
  - **Newsgroup:** news://news.tugraz.at/tu-graz.lv.dbase (email for private issues)
  - **Office hours: Mo 1pm**, or after lecture
- **Website**
  - https://mboehm7.github.io/teaching/ss20_dbs/index.htm
  - All course material (lecture slides, exercises) and dates
- **Exam**
  - **Completed mandatory exercises** (Mar 31, Apr 28, May 19, [Jun 16])
  - **Final written exam** (TBD, doodle voting)
  - **DB Grading** (30% exercises, 70% final)

Course Logistics, cont.

- **Exercises**
  - Written and programming assignments, submitted through **TeachCenter**
  - Assignments have **25 points + ? bonus points** (capped for DB at **80/75**)
  - Assignment **completed if >50% points in total** (**but all submitted**)
  - Deadlines are important (**at most 7 late days in total**)
  - Individual assignments (**academic honesty / no plagiarism**)

- **SW Tools and Languages**
  - **Open Source PostgreSQL** DBMS (setup on your own)
  - Distributed FS/object storage and **Apache Spark** for distributed computation
  - Languages for local/distributed programs (of your choice): e.g., **Python**, **Java**, **Scala**, C, C++, C#, Rust, Go, etc.

Exercises: DBLP Publications

New

- **Dataset**
  - CC0-licensed, derived (extracted, cleaned) from **DBLP** (https://dblp.org Feb 1, 2020) for publication year ≥ 2011
  - Clone or download your copy from https://github.com/tugraz-isds/datasets.git

- **Exercises**
  - **01** Data modeling (relational schema)
  - **02** Data ingestion and SQL query processing
  - **03** Tuning, query processing, and transaction processing
  - **04** Large-scale data analysis (distributed data ingestions and query processing)

persons.csv: The persons file contains author information inclu websites. It's detailed structure and examples look as follows.

```
#PID | name | aliases | affiliation | url
A261789|Matthias Boehm 0001|Matthias Böhm 0001|Graz Unive
A1537639|Stefanie N. Lindstaedt|Stefanie N. Lindstädt||ht
A977823|Denis Helic||Graz University of Technology, Austr
```

theses.csv: The theses file contains the information of public Pł look as follows.

```
#TKey | author | title | year | type | school | pages | i
T25621|A261789|Cost-based optimization of integration flo
T30052|A1399369|An Architecture for Fast and General Data
```

pubs.csv: The pubs file (or better, its individual parts) contains detailed structure and examples look as follows.

```
#PKey | authors| title | year | type | journal | volumne
P519327|A382693:A261789:A261428:A2051042:A69590|MNC: Stru
P1640801|A261789:A2051042:A2047447:A472485:A261428:A38856
P12485|A1399369:A1703306:A1416241:A557115:A650354:A863102
```

confs.csv: The confs file contains the information on conferenc likely be further improved soon.

```
#CKey | title | editors | year | isbn
C8036|Proceedings of the 2019 International Conference on
C76|Proceedings of the 9th USENIX Symposium on Networked
```

Exercises: DBLP Publications, cont.

New

- **DBLP Statistics**
  - **4,782,347** pubs
  - **2,438,282** authors
  - **75,435** PhD theses
  - **43,218** conferences

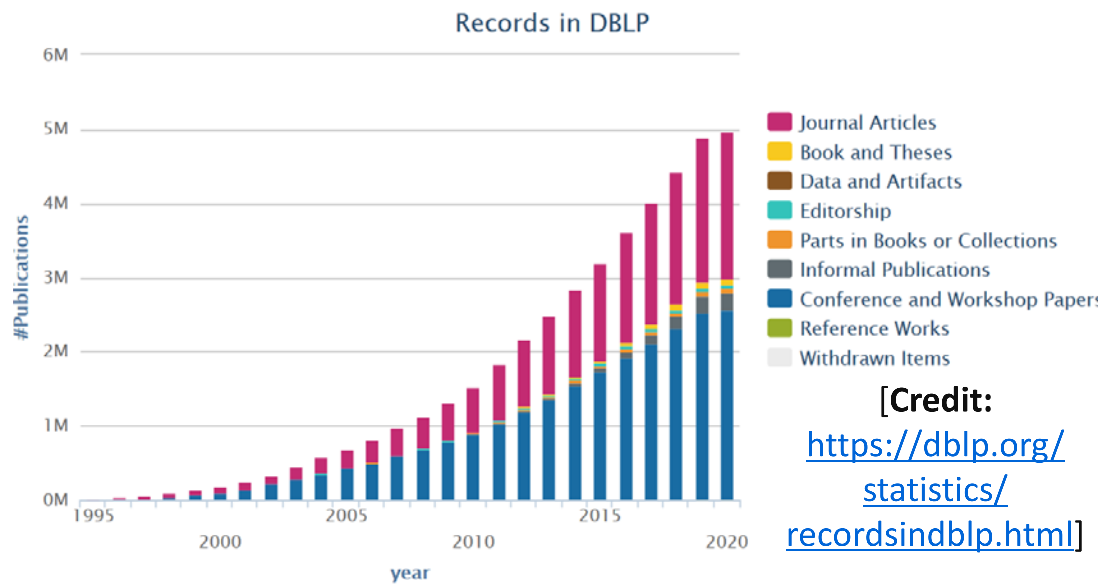

[**Credit:** https://dblp.org/statistics/recordsindblp.html]

- **Our Exercise Dataset**
  - Subset w/ **year ≥ 2011** and **selected features**
  - **2,607,587** pubs
  - **1,716,612** authors
  - **32,534** PhD theses
  - **22,730** conferences

→ **438 MB in uncompressed text files (CSV)**

Literature

- **Not needed for lectures / exercises** (course is self-contained), **but second perspective on covered topics of first part**

- Raghu Ramakrishnan, Johannes Gehrke: Database Management Systems (3. ed.). McGraw-Hill 2003, ISBN 978-0-07-115110-8, pp. I-XXXII, 1-1065
- Jeffrey D. Ullman, Jennifer Widom: A first course in database systems (2. ed.). Prentice Hall 2002, ISBN 978-0-13-035300-9, pp. I-XVI, 1-511
- Ramez Elmasri, Shamkant B. Navathe: Fundamentals of Database Systems, 3rd Edition. Addison-Wesley-Longman 2000, ISBN 978-0-8053-1755-8, pp. I-XXVII, 1-955
- Alfons Kemper, André Eickler: Datenbanksysteme - Eine Einführung, 10. Auflage. De Gruyter Studium, de Gruyter Oldenbourg 2015, ISBN 978-3-11-044375-2, pp. 1-879

History of Data Management

History 1960/70s (pre-relational)

CODASYL … Conference on Data Systems Languages

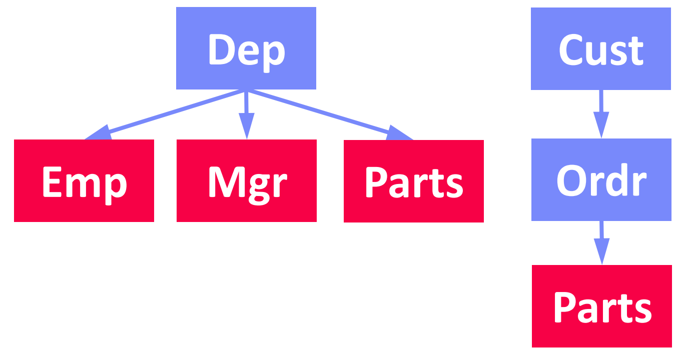

- **Hierarchical Model**
  - Tree of records
  - E.g., IBM Information Management System (IMS) – IMS 15 (Oct 2017)

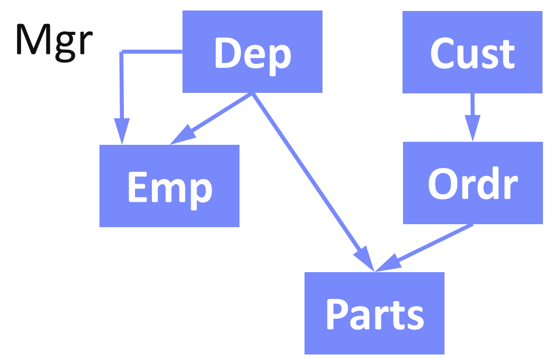

- **Network Model**
  - CODASYL (COBOL, DB interfaces)
  - Graph of records
  - Charles Bachman (**Turing Award '73**)
  - E.g., Integrated Data Store (IDS)
- **Pros and Cons** (see **NoSQL Doc-Stores**)
  - Performance by **directly traversing static links**
  - **Duplicates** → inconsistencies on updates, **data dependence**

History 1970/80s (relational)

Ingres @ UC Berkeley (Stonebraker et al., Turing Award '14)

↑

QUEL - - → SEQUEL

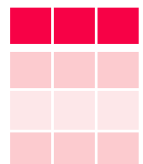

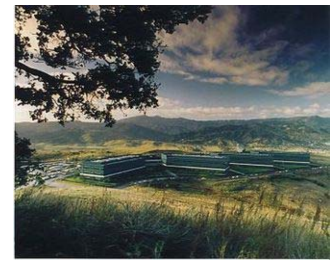

System R @ IBM Research – Almaden (Jim Gray et al., Turing Award '98)

↑

SEQUEL

↑

SQL Standard (SQL-**86**)

Oracle, IBM DB2, Informix, Sybase → MS SQL

Tuple Calculus

Relational Algebra

Relational Model

Goal: **Data Independence** (physical data independence)

- Ordering Dependence
- Indexing Dependence
- Access Path Depend.

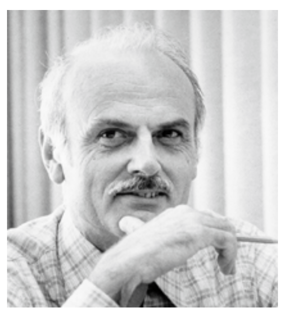

Edgar F. "Ted" Codd @ IBM Research (Turing Award '81)

[E. F. Codd: A Relational Model of Data for Large Shared Data Banks. Comm. ACM 13(6), **1970**]

Success of SQL / Relational Model

Query:

```
SELECT O_OID, sum(O_Price)
FROM Orders, Lineitem, Customer
WHERE O_OID = L_OID AND O_CID = C_CID
   AND O_Odate >= '2018-11-14'
   AND C_Msegment = 'AUTOMOBILE'
GROUP BY O_OID
```

#1 Declarative: what not how

#2 Flexibility: closure property → composability

Logical Query Plans

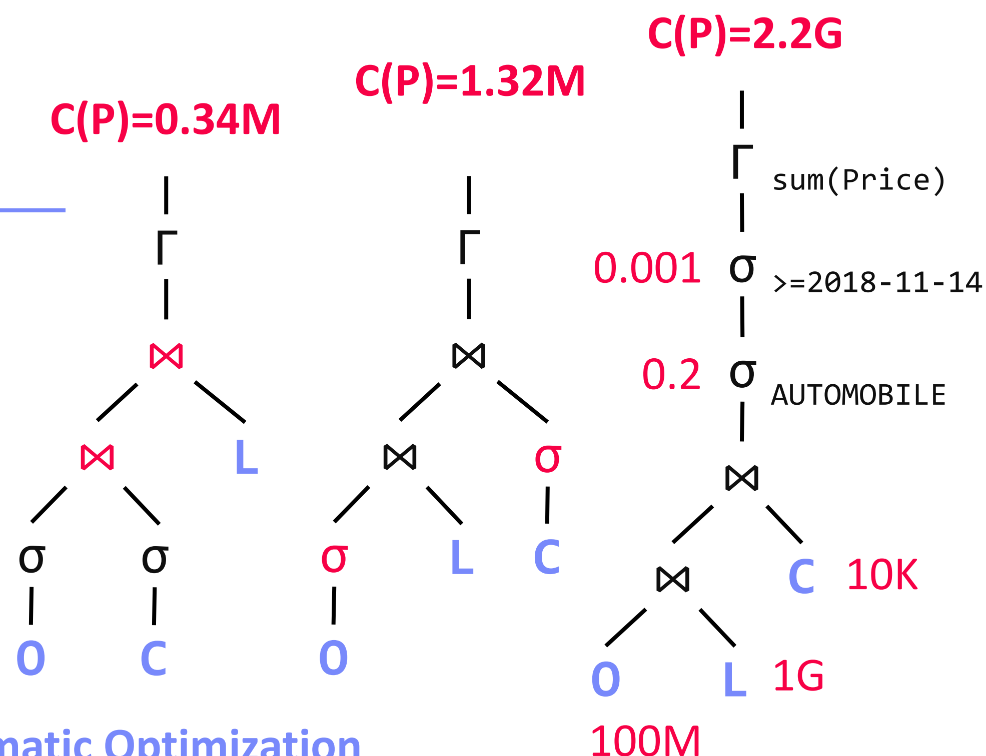

#3 Automatic Optimization

Physical Query Plan

#4 Physical Data Independence

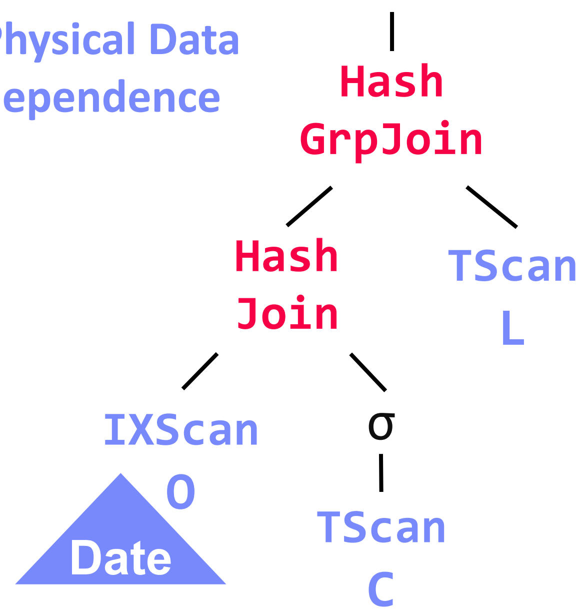

Excursus: PostgreSQL

- **History of PostgreSQL** (used in the exercises)
  - Postgres is the successor project of commercialized Ingres
  - Focus on abstract data types, commercialized as Illustra
  - Prototype w/ SQL open sourced as **Postgres95** → **PostgreSQL**
  - Heavily used as **basis for research projects / startups**

- **Recommended Reading**
  - Michael Stonebraker: The land sharks are on the squawk box. Commun. ACM 59(2): 74-83 (2016), Turing Award Lecture, https://dl.acm.org/citation.cfm?doid=2886013.2869958
  - **Video:** http://www.youtube.com/watch?v=sEPTZVGk3WY
  - **Slides:** http://vldb.org/2015/wp-content/uploads/2015/09/stonebraker.pdf

History 1980/90/2000s

OLTP … Online Transaction Processing
OLAP … Online Analytical Processing
ETL … Extract, Transform, Load

- **Enterprise DBMS**
  - **Heavy investment in research and development** ➔ adoption
  - Oracle, IBM DB2, Informix, Sybase, MS SQL, PostgreSQL, MySQL
  - **Other technologies:** OODBMS, Multimedia, Spatiotemporal, Web, XML
- **Information/Data Warehousing (DWH)**
  - **Workload separation into OLTP and OLAP**
  - **Classical DWH architecture:** operational, staging, DWH, data marts + mining
  - **ETL Process** (Extract, Transform, Load)

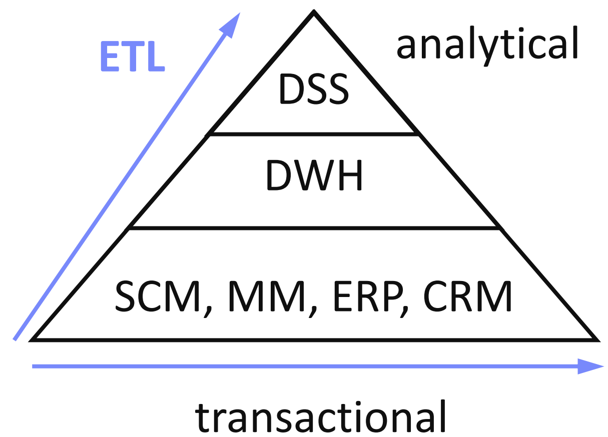

- **Different Personas**
  - Domain Experts (e.g., BI Tools, SAP R/3)
  - DB Application Developers (e.g., ABAP)
  - DB Developers and DB Admins

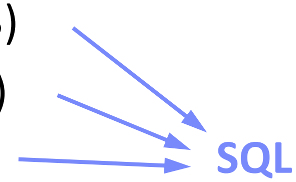

History 2000s / Early 2010s

[M. Stonebraker, S. Madden, D. J. Abadi, S. Harizopoulos, N. Hachem, P. Helland: **The End of an Architectural Era** (It's Time for a Complete Rewrite). **VLDB 2007**]

- **Specialized Systems**
  - **Column stores + compression** for OLAP
  - **Main memory systems** for OLTP and OLAP
  - Data streaming, scientific and graph databases
  - Information extraction / retrieval, and XML
- **Other Research Trends**
  - Approximate QP / **Adaptive QP** / tuning tools
  - **Large-scale data management** (DFS, MR) / **cloud computing**
- **Toward Flexible, Large-Scale Data Management** (DWH … a bygone era)
  - **MAD Skills (magnetic, agile, deep)**, MADlib
  - Integration of R, Python in data analysis
  - Open data and its integration
  - Query processing over raw data files

[J. Cohen, B. Dolan, M. Dunlap, J. M. Hellerstein, C. Welton: **MAD Skills**: New Analysis Practices for Big Data. **PVLDB 2(2) 2009**]

History 2010s – Present

- **Two Key Drivers of DB Research**
  - **New analysis workloads** (NLP, key/value, RDF/graphs, documents, time series, ML) and applications
  - **New HW/infrastructure** (multi-/many-core, cloud, scale-up/ scale-out, NUMA/HBM, RDMA, SSD/NVM, FPGA/GPU/ASIC)

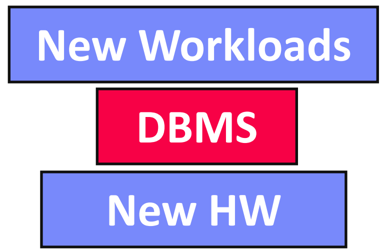

- **Excursus: A retrospective view of specialized systems**
  - Goal **#1**: Avoid boundary crossing ➔ **General-purpose**
  - Goal **#2**: New workload + Performance ➔ **Specialized systems**
  - Some Examples

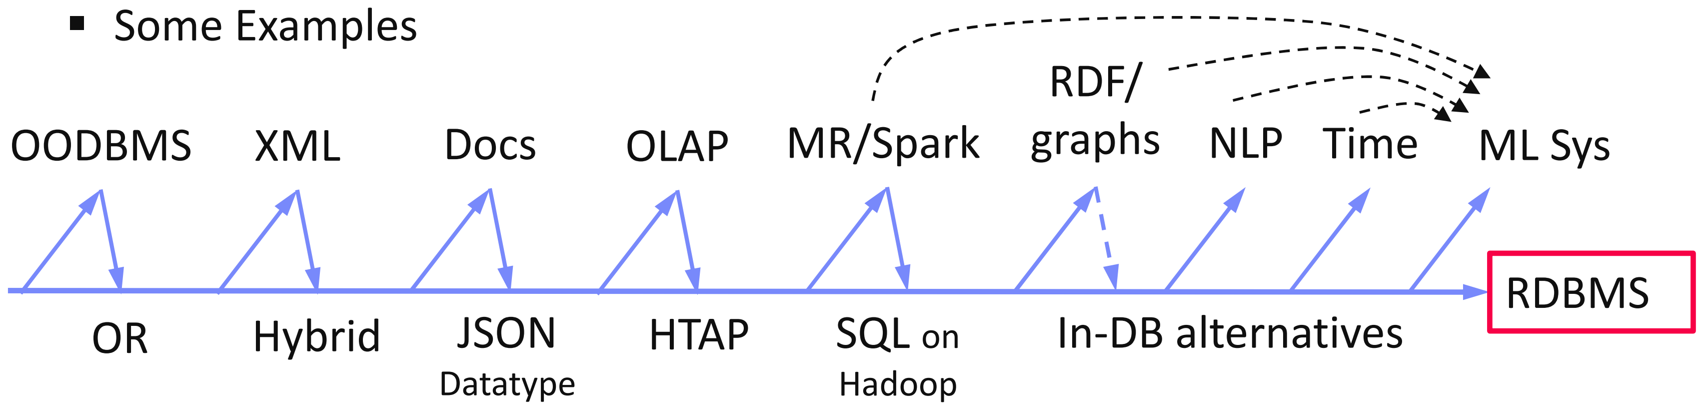

History 2010s – Present (2)

- **Motivation NoSQL Systems**
  - **Flexible schema** (no upfront costs), **scalability**, or specific data types
  - Relaxed **ACID** (atomicity, consistency, isolation, durability) requirements
    ➔ **BASE** (basically available, soft state, eventual consistency)

- **Example NoSQL Systems** (local and distributed)**:**
  - **Key/Value-Stores:** simple put/get/delete, massive scalability
  - **Document-Stores:** store nested documents (tree)
  - **RDF Stores:** store subject-predicate-object triples
  - **Graph DBs:** store nodes/edges/attributes, vertex-centric
  - **Time Series DBs:** store sequences of observations

History 2010s – Present (3)

- **Motivation Large-Scale Data Management**
  - Massive scalability (data/compute) on demand, fault tolerance, flexibility
  - Example Facebook 2014: **300PB** DWH, **600TB** daily ingest
  - Cost-effective commodity hardware
  - Error rate increases with increasing scale

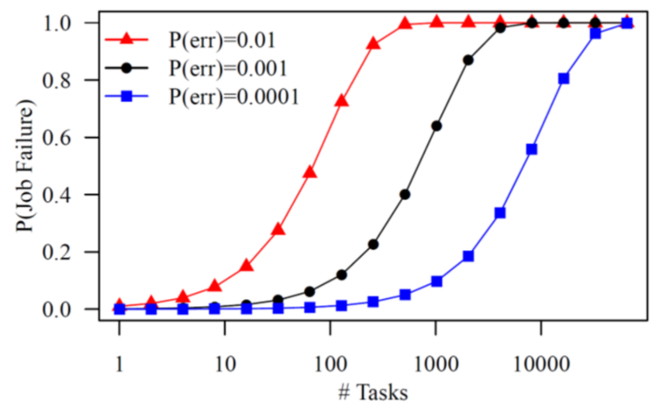

- **Examples Large-Scale Data Management**
  - Distributed file systems w/ replication (e.g., GPFS, HDFS)
  - Cloud object storage (e.g., Amazon s3, OpenStack Swift)
  - Data-parallel data analysis with Spark/Flink, incl streaming
  - Automatic cloud resource elasticity (pay as you go)

} Data Lake

Summary and Q&A

- **#1 Database Systems**
  - Mature and established technology → broadly applicable & eco system
  - **General concepts:** abstraction, data modeling, query optimization & processing, transaction processing and recovery, physical design and tuning
- **#2 Modern Data Management**
  - Multiple specialized systems for specific scale / data types
  - General trend toward less upfront cost, flexibility, and higher scalability

→ Variety of data management tools → Course meta goals

  - **Understand, use, debug**, and evaluate data management systems
  - **Fundamental concepts** as basis for advanced courses and other areas

- **Upcoming**
  - **02** **Conceptual Architecture and Design** [Mar 09] (ER Diagrams)
  - **03** **Data Models and Normalization** [Mar 16] (ERD -> Relational Model)