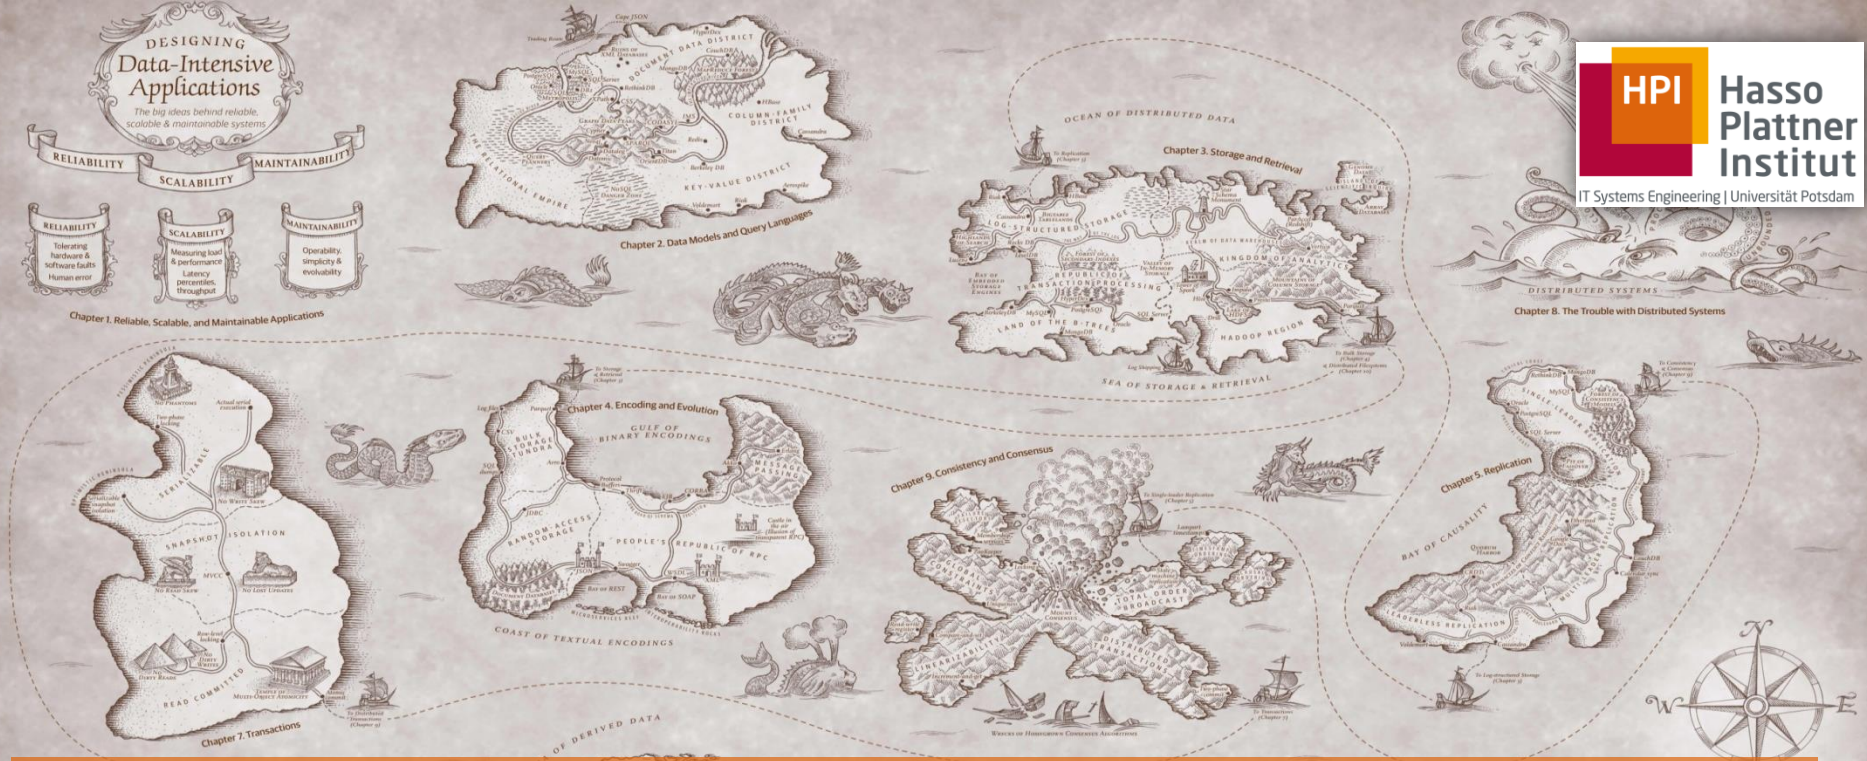

Distributed Data Management
Introduction

Thorsten Papenbrock
Felix Naumann

F-2.03/F-2.04, Campus II

Hasso Plattner Institut

O'REILLY
Martin Kleppmann: *Designing Data-Intensive Applications*
O'Reilly Media, 2017
ISBN 978-1-4493-7332-0


Outdated Data
Lost & Invalid Messages
Concurrency
Consensus
Termination

Dr. Thorsten **Papenbrock**

Diana **Stephan**

Prof. Felix **Naumann**

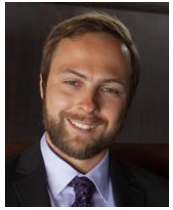

Dr. Ralf **Krestel**

Tim **Repke**

Julian **Risch**

Hazar **Harmouch**

Michael **Loster**

John **Koumarelas**

Tobias **Bleifuß**

Leon **Bornemann**

Lan **Jiang**

Konstantina **Lazariduo**

Service-Oriented Systems

project **DuDe**

Data Fusion

Data Scrubbing

project **Stratosphere**

Entity Search

Data Profiling

Information Integration

Data as a Service

Web Science

Duplicate Detection

project **Metanome**

Data Cleansing

Text Mining

Information Quality

Web Data

Dependency Detection

Linked Open Data

RDF Data Mining

Entity Recognition

Data Change

Opinion Mining

ETL Management

Data Preparation

project **DataChEx**

Change Exploration

Introduction: Audience

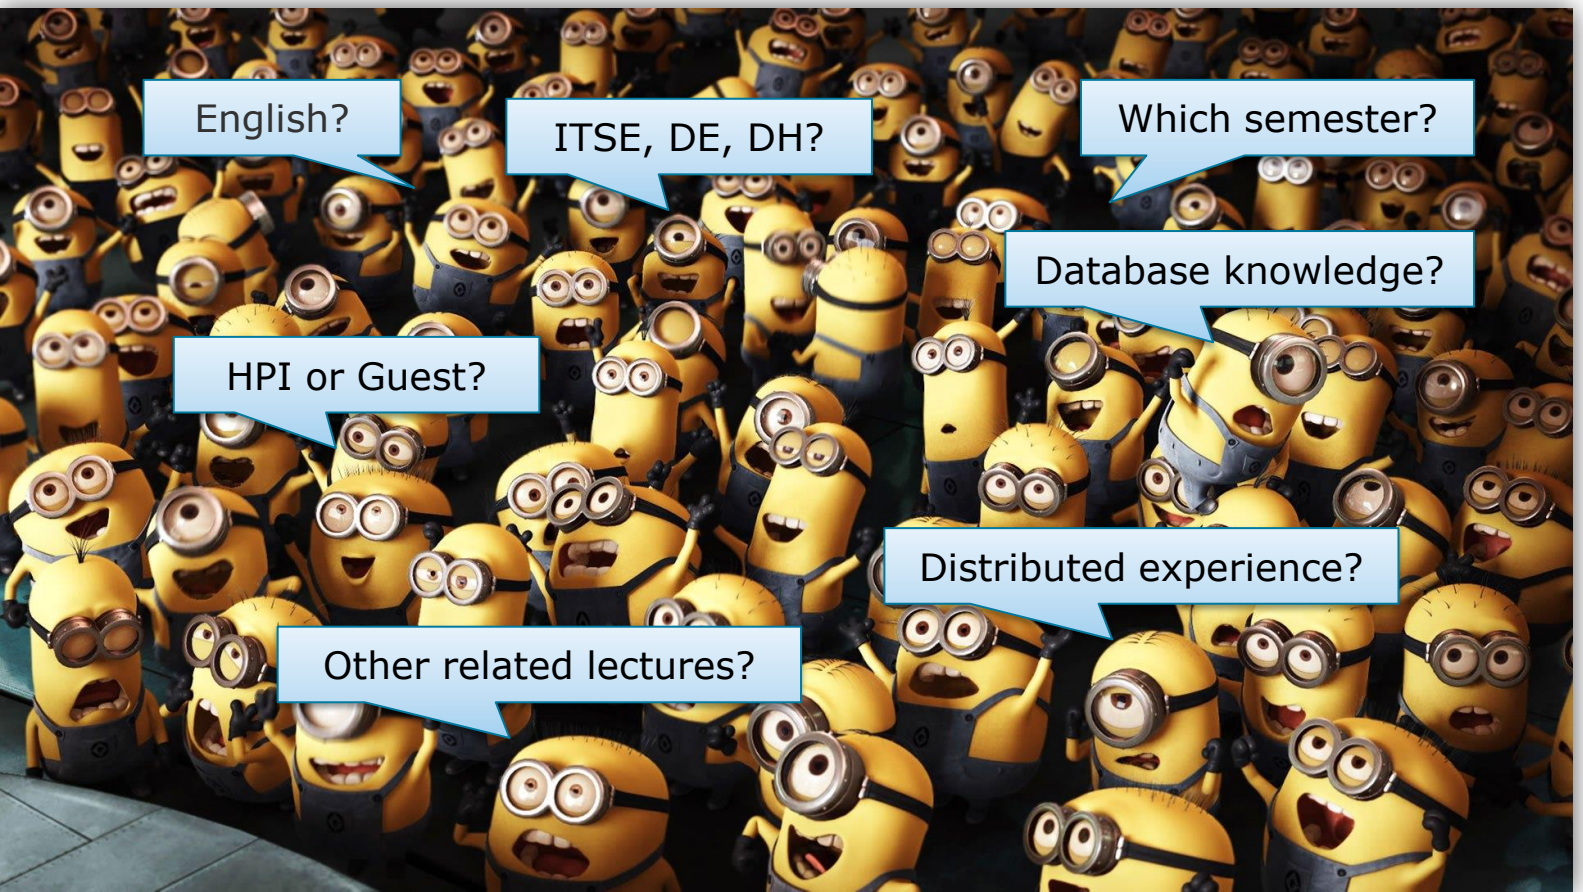

Courses 2018/2019

Lectures

- Datenbanksysteme II (Bachelor)
- Deep Learning for Text Mining (Master)
- Distributed Data Management (Master)

Seminars

- Methoden der Forschung (Master)
- Data Preparation for Science (Master)
- Recommender Systems (Master)
- Social Network Analysis in Practice (Master)

Domain Expertise (e.g., Industry 4.0, Medicine, Physics, Engineering, Energy, Logistics)

Mathematical Programming

Linear Algebra

Statistics

Text Mining

Graph Mining

Signal Processing

Stochastic Gradient Descent

Machine Learning

Error Estimation

Active Sampling

Monte Carlo

Regression

Predictive Analytics

Sketches

Data Obfuscation

Convergence

Decoupling

Iterative Algorithms

Curse of Dimensionality

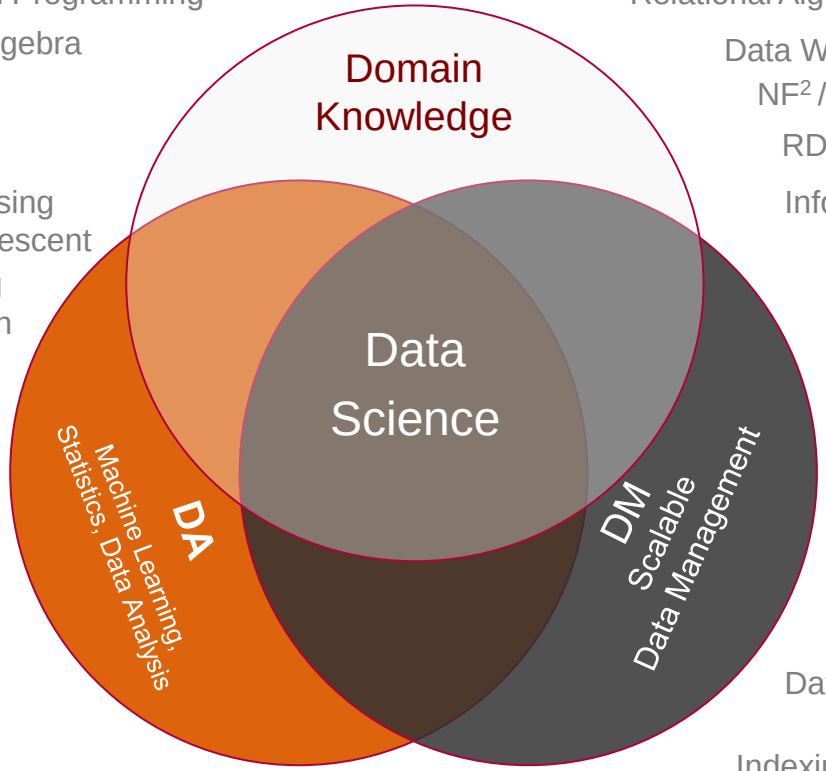

Business Models

Legal Aspects

Relational Algebra / SQL

Data Warehouse/OLAP

NF^2 /XQuery

RDF / SparQL

Information Integration

Information Extraction

Visual Analytics

Privacy

Memory Management

Parallelization

Scalability

Memory Hierarchy

Fault Tolerance

Security

Data Analysis Languages

Query Optimization

Indexing

Data Flow

Control Flow

Real-Time

This Lecture

Lecture

- For master students
  (IT-Systems Engineering,
  Digital Health, Data Engineering)
- 6 credit points, 4 SWS
- Mondays 13:30 – 15:00
  Tuesdays 9:15 – 10:45

Exercises

- Interleaved with lectures

Slides

- On website

Website

- https://hpi.de/naumann/teaching/teaching/ws-1819/distributed-data-management-vl-master.html

Prerequisites

- To participate:
  A little background and interest in databases (e.g. DBS I lecture)
- For exam:
  Attending lectures, participation in exercises, and completion of exercise homework tasks

Antirequisite

- Distributed Data Analytics

Exam

- Written exam
- Probably first week after lectures

Feedback

Question any time please!

- During lectures
- Visit us: Campus II, Rooms F-2.03 and F-2.04
- Email:
  - thorsten.papenbrock@hpi.de
  - felix.naumann@hpi.de

Also: Give feedback about …

- improving lectures
- informational material
- organization

Official evaluation

- At the end of this semester
- … too late for important feedback!

Distributed Data Management

Distributed Data Management

Motivation: "Distributed"

Paradigm Shift in Software-Writing

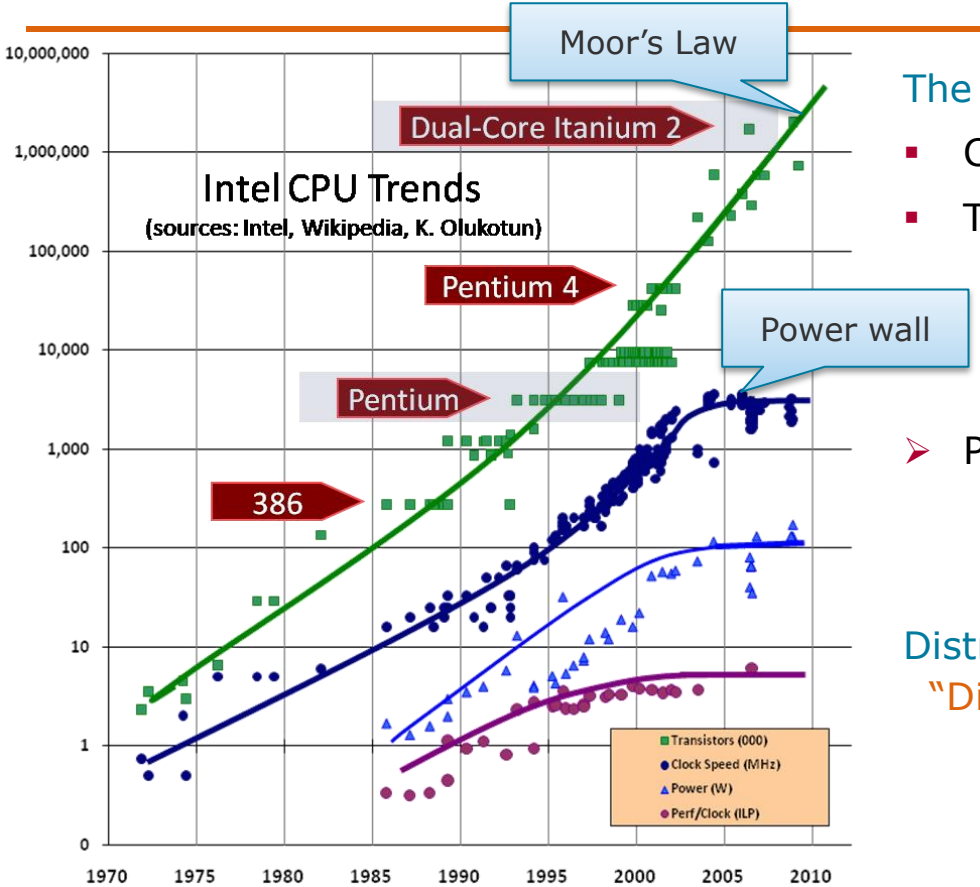

The free lunch is over!

- Clock speeds stall
- Transistor numbers still increase
  - Cores in CPUs/GPUs
    CPUs/GPUs in compute nodes,
    compute nodes in clusters
  - Paradigm Shift:
    - Earlier: optimize code for a single thread
    - Now: solve tasks in parallel

Distributed computing

"Distribution of work on (potentially) physically isolated compute nodes"

http://www.gotw.ca/publications/concurrency-ddj.htm

Motivation: “Distributed”

Surpassing Moor’s Law

With **distributed machines**, we can already build systems with any number of transistors!

(don’t even need to wait for a new processors)

Moore’s Law (Observation)

“The number of transistors in integrated circuits doubles approximately every two years”

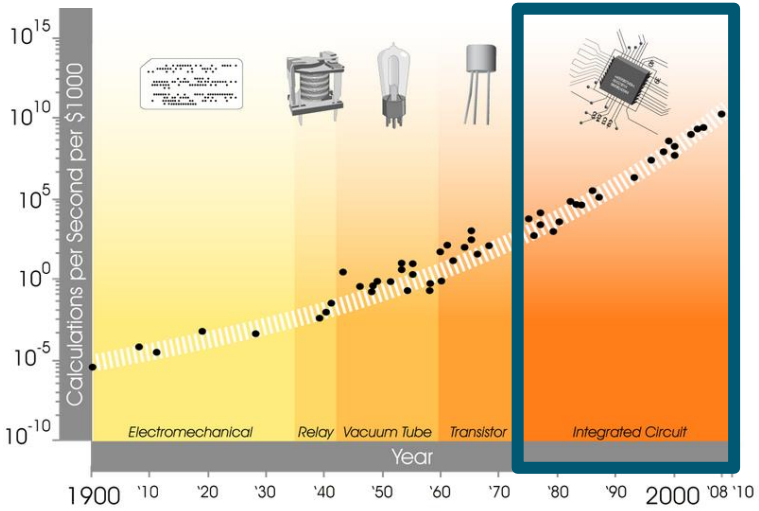

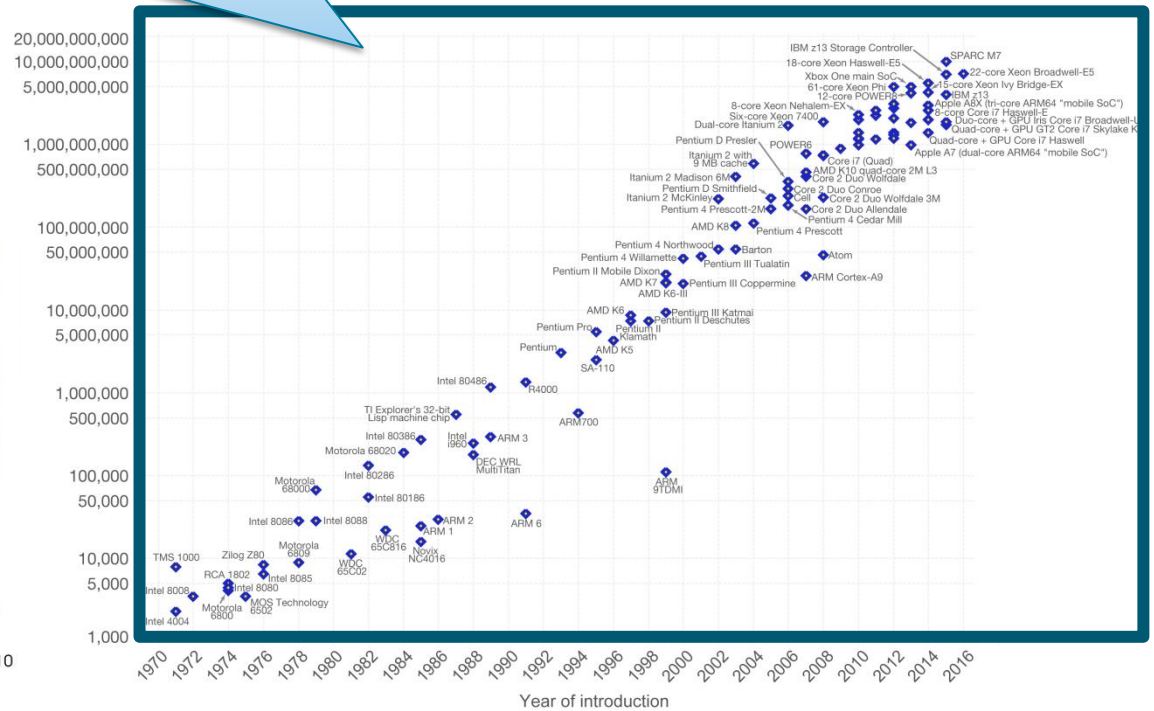

Motivation: "Distributed"

A Rule to Acknowledge

Even **distributed parallelization** cannot work around this law!

Amdahl's Law

"The speedup of a program using multiple processors for parallel computing is limited by the sequential fraction of the program"

$$Speedup(s) = \frac{1}{(1-p) + \frac{p}{s}}$$

s: degree of parallelization (e.g. #cores)

p: percentage of the algorithm that profits from parallelization

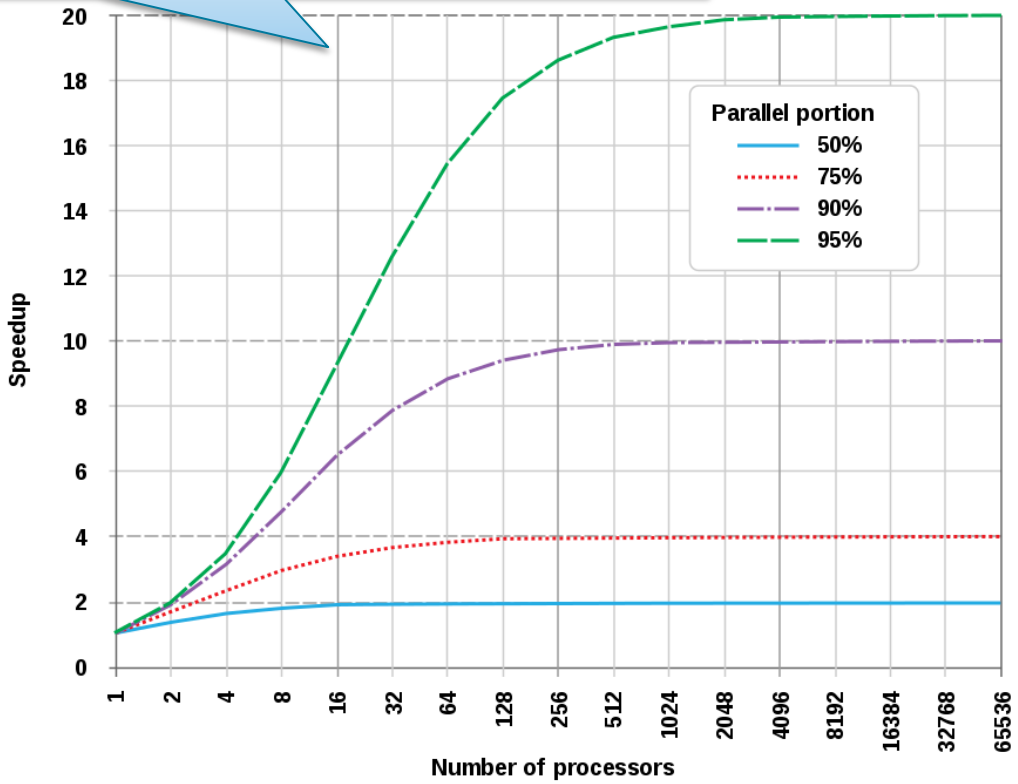

Motivation: “Distributed”

New Technologies

Distributed Computing

...

Distributed Storage

...

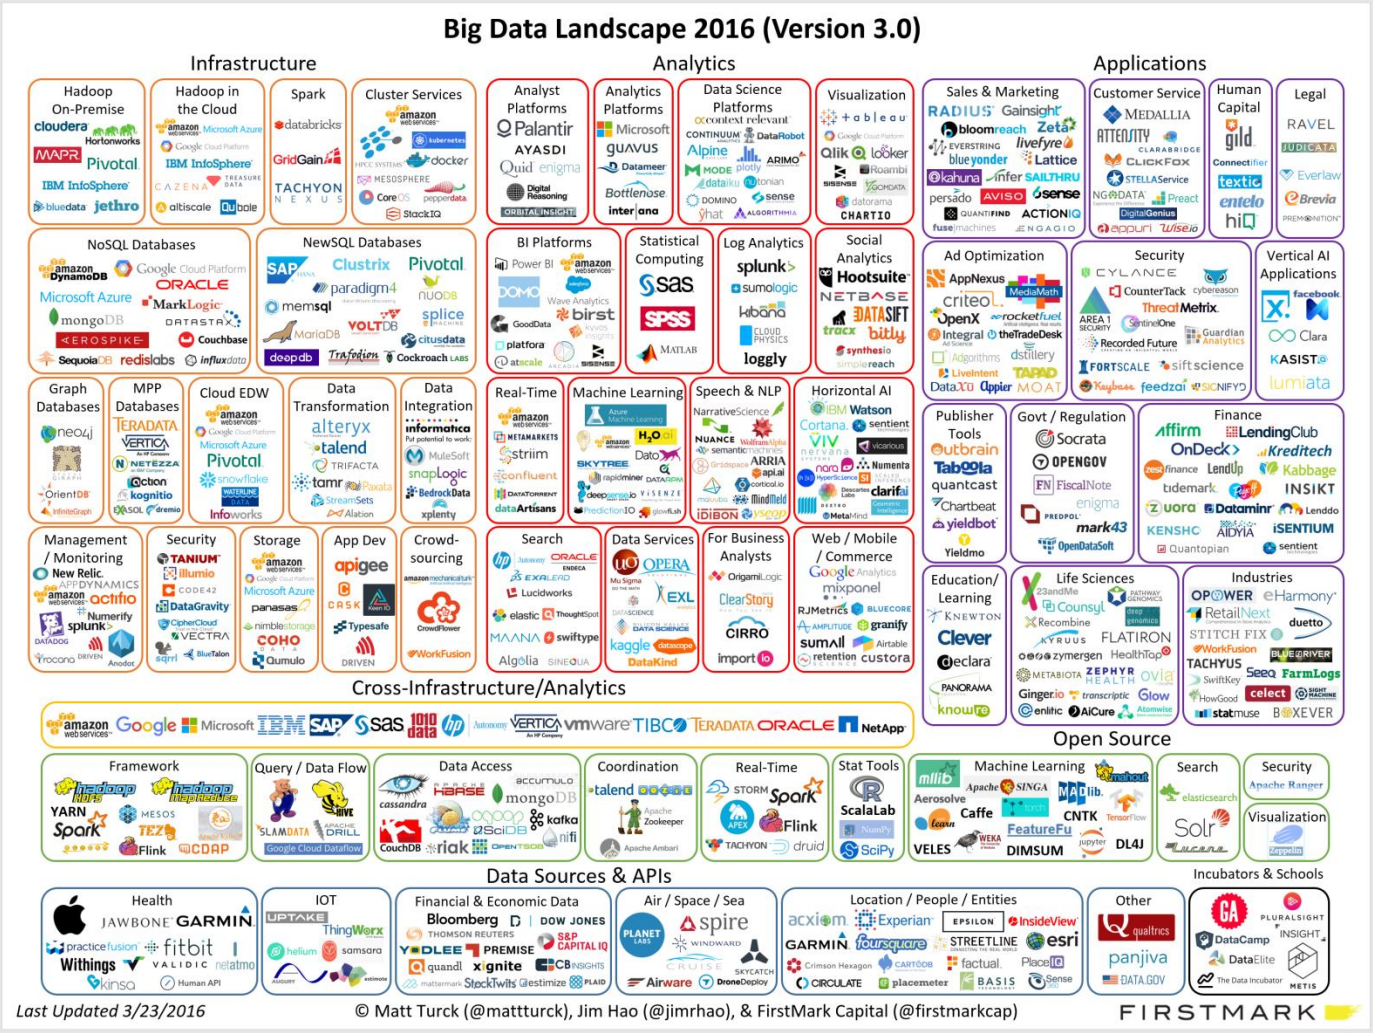

HPI Hasso Plattner Institut

Distributed Data Management

Introduction

ThorstenPapenbrock

Slide **15**

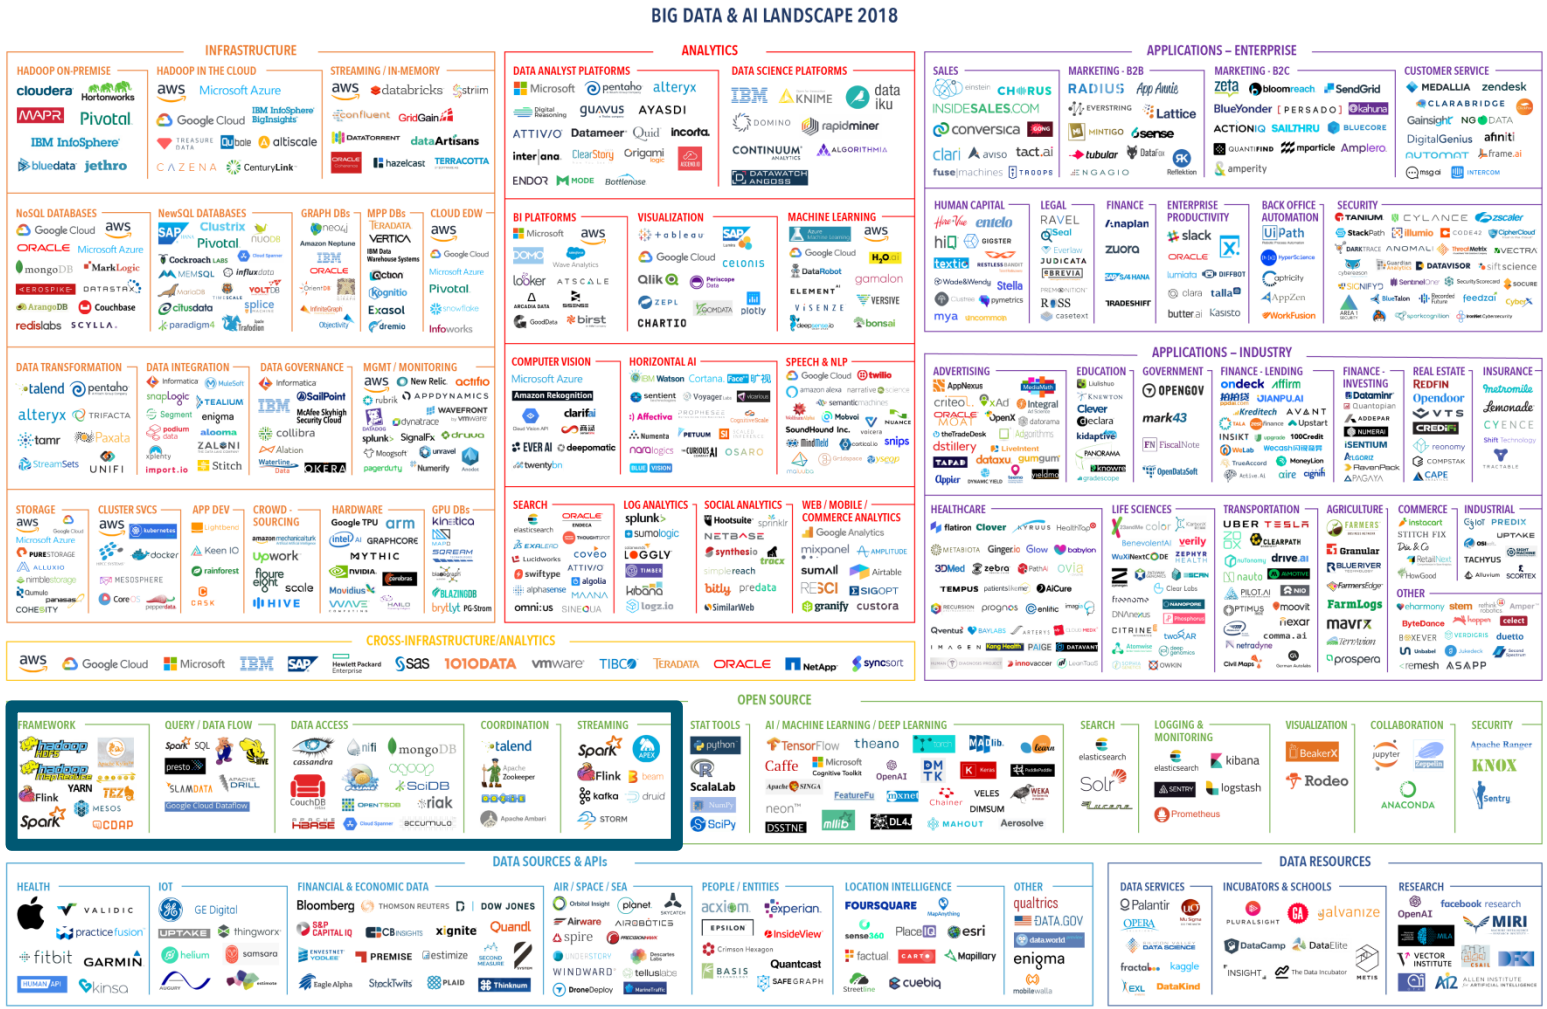

HPI Hasso Plattner Institut

Distributed Data Management

Introduction

ThorstenPapenbrock
Slide **16**

FIRSTMARK EARLY STAGE VENTURE CAPITAL

Motivation: “Distributed”

Small and Medium Scale

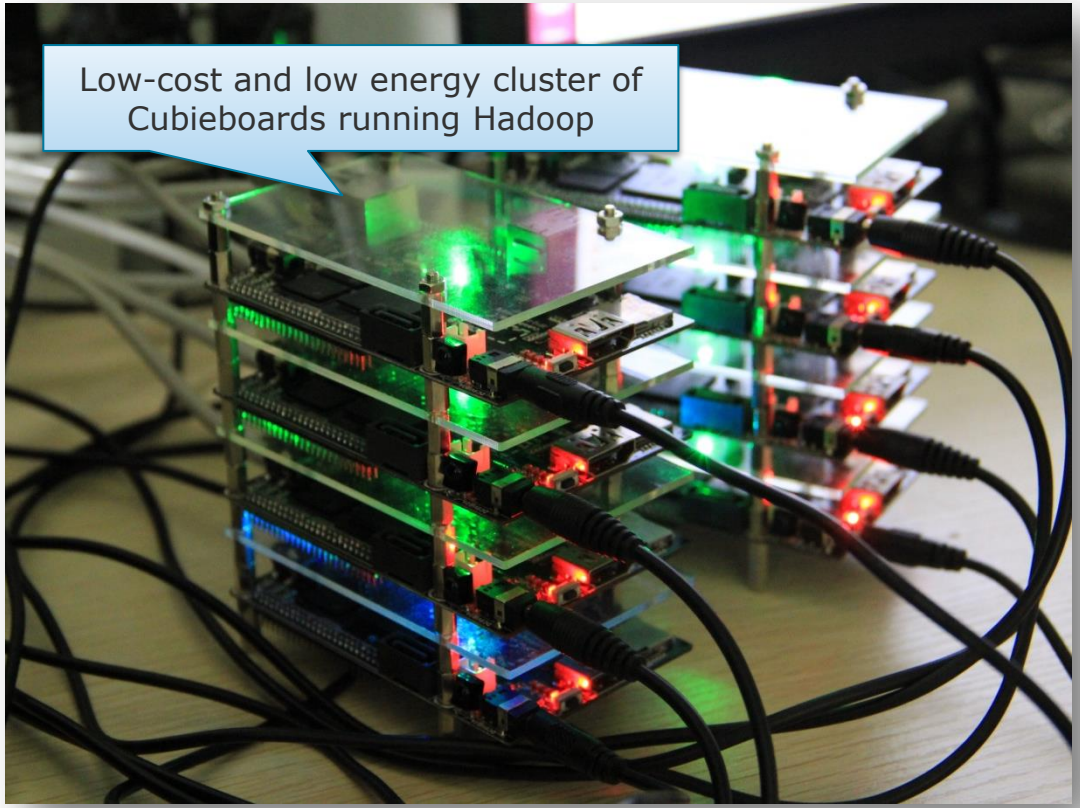

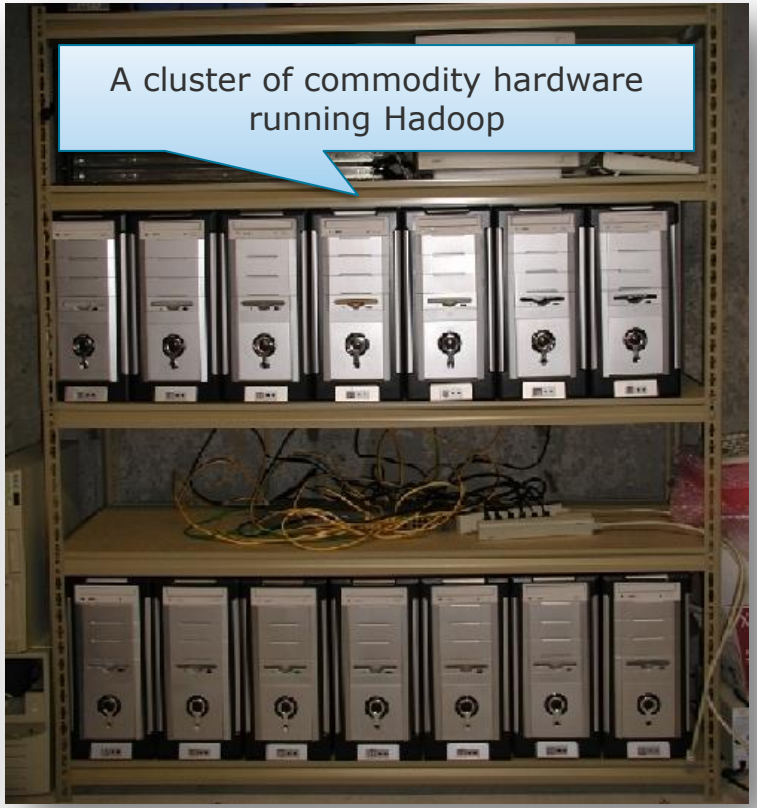

Motivation: “Distributed”

Large Scale

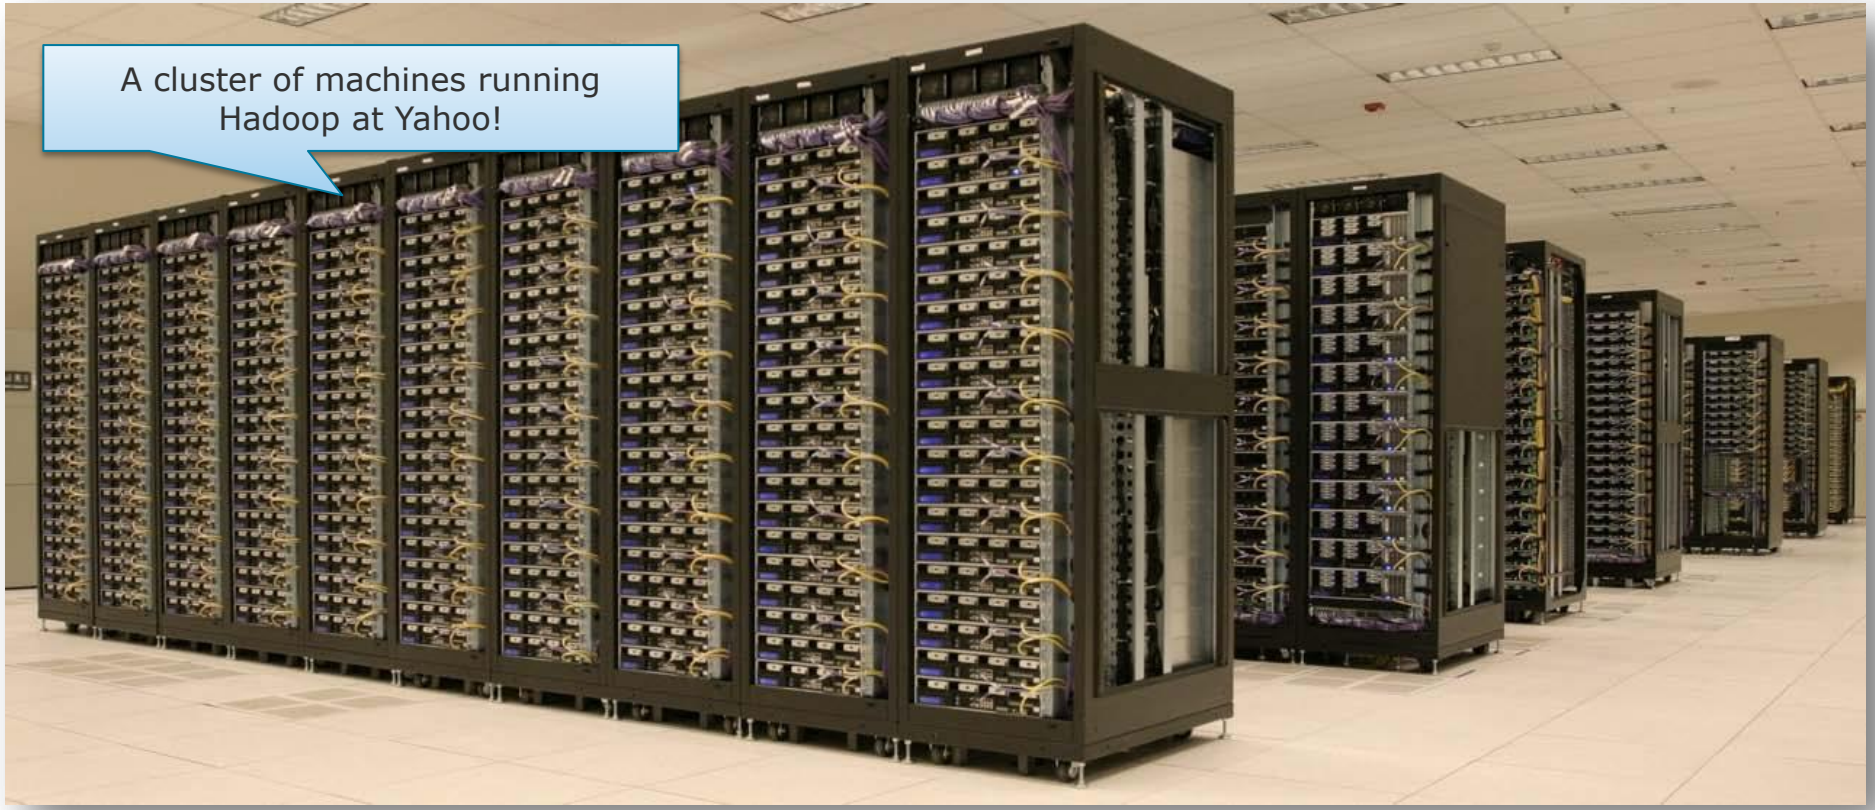

Motivation: "Distributed"

Super Large Scale

Top 10 Super Computers 2017

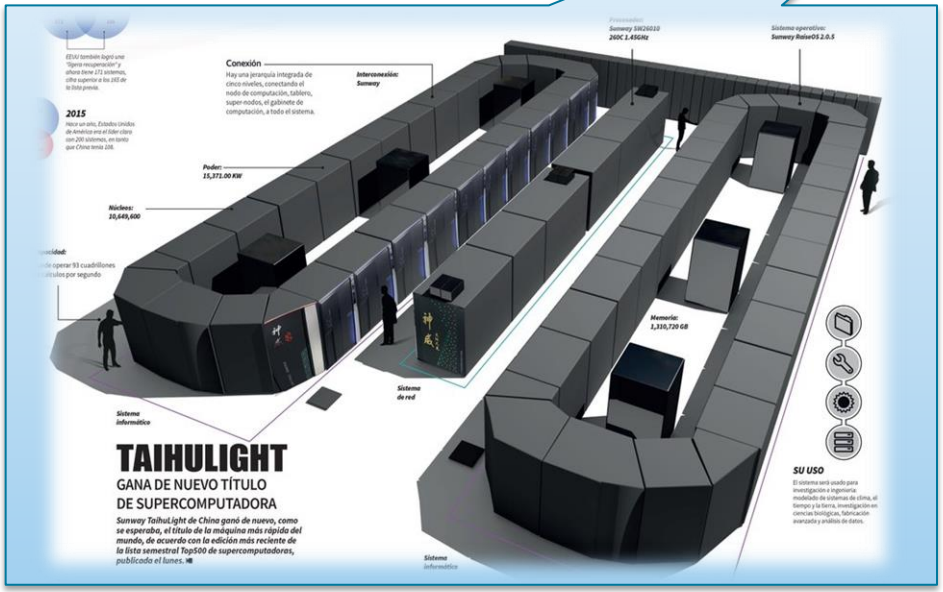

All **distributed systems**!

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
2	**Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P , NUDT National Super Computer Center in Guangzhou China	3,120,000	33,862.7	54,902.4	17,808
3	**Piz Daint** - Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100 , Cray Inc. Swiss National Supercomputing Centre (CSCS) Switzerland	361,760	19,590.0	25,326.3	2,272
4	**Titan** - Cray XK7, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x , Cray Inc. DOE/SC/Oak Ridge National Laboratory United States	560,640	17,590.0	27,112.5	8,209
5	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom , IBM DOE/NNSA/LLNL United States	1,572,864	17,173.2	20,132.7	7,890
6	**Cori** - Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect , Cray Inc. DOE/SC/LBNL/NERSC United States	622,336	14,014.7	27,880.7	3,939
7	**Oakforest-PACS** - PRIMERGY CX1640 M1, Intel Xeon Phi 7250 68C 1.4GHz, Intel Omni-Path , Fujitsu Joint Center for Advanced High Performance Computing Japan	556,104	13,554.6	24,913.5	2,719
8	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect , Fujitsu RIKEN Advanced Institute for Computational Science (AICS) Japan	705,024	10,510.0	11,280.4	12,660
9	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom , IBM DOE/SC/Argonne National Laboratory United States	786,432	8,586.6	10,066.3	3,945
10	**Trinity** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect , Cray Inc. DOE/NNSA/LANL/SNL United States	301,056	8,100.9	11,078.9	4,233

https://www.top500.org/lists/2017/06/

Motivation: "Distributed"

Super Large Scale

Top 10 Super Computers 2017

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
2	**Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P , NUDT National Super Computer Center in Guangzhou China	3,120,000	33,862.7	54,902.4	17,808
3	**Piz Daint** - Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100 , Cray Inc. Swiss National Supercomputing Centre (CSCS) Switzerland	361,760	19,590.0	25,326.3	2,272
4	**Titan** - Cray XK7, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x , Cray Inc. DOE/SC/Oak Ridge National Laboratory United States	560,640	17,590.0	27,112.5	8,209
5	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom , IBM DOE/NNSA/LLNL United States	1,572,864	17,173.2	20,132.7	7,890
6	**Cori** - Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect , Cray Inc. DOE/SC/LBNL/NERSC United States	622,336	14,014.7	27,880.7	3,939
7	**Oakforest-PACS** - PRIMERGY CX1640 M1, Intel Xeon Phi 7250 68C 1.4GHz, Intel Omni-Path , Fujitsu Joint Center for Advanced High Performance Computing Japan	556,104	13,554.6	24,913.5	2,719
8	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect , Fujitsu RIKEN Advanced Institute for Computational Science (AICS) Japan	705,024	10,510.0	11,280.4	12,660
9	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom , IBM DOE/SC/Argonne National Laboratory United States	786,432	8,586.6	10,066.3	3,945
10	**Trinity** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect , Cray Inc. DOE/NNSA/LANL/SNL United States	301,056	8,100.9	11,078.9	4,233

All **distributed systems**!

https://www.top500.org/lists/2017/06/

Motivation: "Distributed"

Super Large Scale

Top 10 Super Computers 2017

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
2	**Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P , NUDT National Super Computer Center in Guangzhou China	3,120,000	33,862.7	54,902.4	17,808
3	**Piz Daint** - Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100 , Cray Inc. Swiss National Supercomputing Centre (CSCS) Switzerland	361,760	19,590.0	25,326.3	2,272
4	**Titan** - Cray XK7, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x , Cray Inc. DOE/SC/Oak Ridge National Laboratory United States	560,640	17,590.0	27,112.5	8,209
5	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom , IBM DOE/NNSA/LLNL United States	1,572,864	17,173.2	20,132.7	7,890
6	**Cori** - Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect , Cray Inc. DOE/SC/LBNL/NERSC United States	622,336	14,014.7	27,880.7	3,939
7	**Oakforest-PACS** - PRIMERGY CX1640 M1, Intel Xeon Phi 7250 68C 1.4GHz, Intel Omni-Path , Fujitsu Joint Center for Advanced High Performance Computing Japan	556,104	13,554.6	24,913.5	2,719
8	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect , Fujitsu RIKEN Advanced Institute for Computational Science (AICS) Japan	705,024	10,510.0	11,280.4	12,660
9	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom , IBM DOE/SC/Argonne National Laboratory United States	786,432	8,586.6	10,066.3	3,945
10	**Trinity** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect , Cray Inc. DOE/NNSA/LANL/SNL United States	301,056	8,100.9	11,078.9	4,233

All **distributed systems**!

https://www.top500.org/lists/2017/06/

Distributed Data Management

"The ability of extracting different kinds of information from data!"

- Structural information
- Explicit information
- Implicit/derived information

Analytics

BIG Data & Analytics Software Vendors

Top 10 LinkedIn followers per employee:
1. Dundas
2. Predixion Software
3. iDashboards
4. Chartio
5. CXO-Cockpit
6. Kognitio
7. TARGIT
8. Prophix
9. Alpine Analytics
10. Jedox

A market worth $122 billion in 2016 with a growth of 11.3% per year!

Excellent job opportunities in many companies!

For a world that created an entire zettabyte (which is exactly 10^{12} GB) of data in the 2010 alone!

Business users (Log-scale)

Customers (Log-scale)

Shape: Deployment types

On Prem

Cloud — Mobile

Cloud (Mobile)

On Prem

MicroStrategy, Alteryx, IBM, Oracle, SAP, Microsoft, Information Builders, Informatica, OpenText, SAS, Cloudera, Hortonworks, FICO, Tibco, Tableau, Qlik, Dun & Bradstreet, Birst, GoodData, Zoho, Infor, Domo, Logi Analytics, TARGIT, Board, Tagetik, Bureau van Dijk, iDashboard, Tidemark, Dundas, 1010data, Host Analytics, Prophix, Adaptive Insight, Workday/Platfora, Anaplan, Phocas Software, Pentaho, Xactly, Sisense, Pyramid Analytics, Jedox, Stibo Systems, Jet Reports, LongView, Iptor, Prognoz, TimeXtender, InetSoft, EXASOL, Toolpack, ClearStory Data, Chartio, Infobright, MapR, Looker, MarkLogic, DataWatch, MongoDB, KNIME, RapidMiner, Panorama Software, Dell, Angoss, Collibra, CallidusCloud, Predixion Software, Corporater, Salesforce, Metapraxis, Lavastorm, MemSQL, SEAMS, CXO-Cockpit, HaloBI, Transwarp, ZAP, Axiom, Solver - BI360, InfoSuite, Actian, Dimensional Insight, Rapid Insight, eQ Technologic, Alpine Data, Yellowfin, Ducen, Neotix, InetSoft, Megaputer, Aaro, prevero, AnyChart, MITS, BizWiz, cubus, Bilbeo, InsightSoftware, Exago, Silvon, Strategy Companion, Jinfonet Software, BI4Dynamics, Cyberscience, evidenza, Bissantz, Konsolidator, Cubeware Cockpit

Top 5 Self Service BI Vendors:
1. Tableau
2. Microsoft
3. Qlik
4. MicroStrategy
5. GoodData

Top 10 Business Analytics Vendors:
1. Microsoft
2. MicroStrategy
3. IBM
4. SAP
5. Tableau
6. Qlik
7. Alteryx
8. SAS Institute
9. TIBCO
10. GoodData

Top 5 Predictive Analytics Vendors:
1. IBM
2. SAS
3. RapidMiner
4. KNIME
5. Microsoft

Top 5 Data Warehouse Vendors:
1. Teradata
2. SAP
3. IBM
4. Oracle
5. Microsoft

Color = Client focus
(Blue = Enterprise ; Orange = SMB ; Green: Mix)

Size = Completeness of vision
(BIG Data, Analytics, CPM, Data Warehousing)

https://www.idc.com/getdoc.jsp?containerId=prUS41826116
http://sigmacareer.com/big-data-what-is-it-and-what-are-the-trends

VLDB 2017 Program

International conference "Very Large Data Bases"

29 Aug (Tue)								
8:30-10:00		Welcome Messages + Wolfgang Lehner keynote						
10:00-10:30	break							
10:30-12:00		Panel		DB Engines 1	Data Cleaning	Spatial Data Management 1	Graphs and Networks 1	Tutorial: Caching at the Web Scale
12:00-13:30	lunch							
13:30-15:00		Query Processing and Optimization		Stream Processing 1	Data Formats	Privacy and Security	Event Processing	Tutorial: Summarizing Static and Dynamic Big Graphs
15:00-15:30	break							
:00		Transactions		Spatial Data management 2	Graphs and Networks 2	Information Integration	Applications	Tutorial: Blockchains and Databases
:00			poster reception					
(Wed)								
8:30-10:00		Endowment Update + VLDB Awards + Michael Franklin keynote						
10:00-10:30	break							
10:30-12:00		High Performance Query Processing		Recommendations and Skylines	Approximation Structures	Stream Processing 2		Tutorial: New Trends on Exploratory Methods for Data Analytics
12:00-13:30	lunch							
13:30-15:00		Scalable Storage		Crowdsourcing	Stream Processing 3	Graphs and Networks 3	Social Network Analysis	Tutorial: Geometric Approaches for Top-k Queries
15:00-15:30	break							
15:30-17:00		Concurrency Control		Query Optimization	Specialized Data Management	Graphs and Networks 4	Data Mining and Analytics	Tutorial: Spatial Crowdsourcing: Challenges, Techniques, and Applications
17:00-18:00			poster reception					
19:00 - ...	**"VLDB Octoberfest" banquet at the Hofbräuhaus**							
31 Aug (Thu)								
8:30-10:30		Endowment Awards + Demo Award + Jens Dittrich plenary						
10:30-10:55	break							
10:55-12:00		Transactions and Persistence		Data Access	Data Statistics	Potpourri		
12:00-13:30	lunch							
13:30-15:00		Data Partitioning		Graphs and Networks 5	Visualization	Distributed Systems and Cloud 1	Estimation and Approximation	Tutorial: The Era of Big Spatial Data
15:00-15:30	break							
15:30-17:00		DB Engines 2		Query Processing	Text and Semi-Structured	Spatial Data Management 3	Distributed Systems and Cloud 2	Tutorial: Complex Event Recognition in the Big Data Era
17:00-18:00			poster reception					

VLDB 2017 Program

29 Aug (Tue)								
8:30-10:00		Welcome Messages + Wolfgang Lehner keynote						
10:00-10:30	break							
10:30-12:00		Panel		DB Engines 1	Data Cleaning	Spatial Data Management 1	Graphs and Networks 1	Tutorial: Caching at the Web Scale
12:00-13:30	lunch							
13:30-15:00		Query Processing and Optimization		Stream Processing 1	Data Formats	Privacy and Security	Event Processing	Tutorial: Summarizing Static and Dynamic Big Graphs
15:00-15:30	break							
15:30-17:00		Transactions		Spatial Data management 2	Graphs and Networks 2	Information Integration	Applications	Tutorial: Blockchains and Databases
17:00-18:00			poster reception					
30 Aug (Wed)								
8:30-10:00		Endowment Update + VLDB Awards + Michael Franklin keynote						
10:00-10:30	break							
10:30-12:00		High Performance Query Processing		Recommendations and Skylines	Approximation Structures	Stream Processing 2		Tutorial: New Trends on Exploratory Methods for Data Analytics
12:00-13:30	lunch							
13:30-15:00		Scalable Storage		Crowdsourcing	Stream Processing 3	Graphs and Networks 3	Social Network Analysis	Tutorial: Geometric Approaches for Top-k Queries
[illegible]	break							
[illegible]:00		Concurrency Control		Query Optimization	Specialized Data Management	Graphs and Networks 4	Data Mining and Analytics	Tutorial: Spatial Crowdsourcing: Challenges, Techniques, and Applications
[illegible]:00			poster reception					
[illegible]	**"VLDB Octoberfest" banquet at the Hofbräuhaus**							
[illegible] (Thu)								
8:30-10:30		Endowment Awards + Demo Award + Jens Dittrich plenary						
10:30-10:55	break							
10:55-12:00		Transactions and Persistence		Data Access	Data Statistics	Potpourri		
12:00-13:30	lunch							
13:30-15:00		Data Partitioning		Graphs and Networks 5	Visualization	Distributed Systems and Cloud 1	Estimation and Approximation	Tutorial: The Era of Big Spatial Data
15:00-15:30	break							
15:30-17:00		DB Engines 2		Query Processing	Text and Semi-Structured	Spatial Data Management 3	Distributed Systems and Cloud 2	Tutorial: Complex Event Recognition in the Big Data Era
17:00-18:00			poster reception					

With strong relation to **data analytics** and/or **distributed processing**

Motivation: "Data Analytics"

Back to Super Computers

For what tasks are these computer used?

- Weather forecasting
- Market analysis
- Crash simulation
- Disaster simulation
- Brute force decryption
- Molecular dynamics modeling
- …

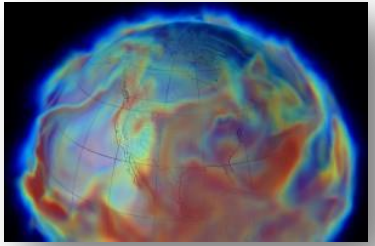

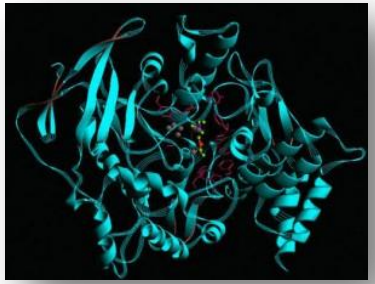

Data Analytics tasks!

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
2	**Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P , NUDT National Super Computer Center in Guangzhou China	3,120,000	33,862.7	54,902.4	17,808
3	**Piz Daint** - Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100 , Cray Inc. Swiss National Supercomputing Centre (CSCS) Switzerland	361,760	19,590.0	25,326.3	2,272
4	**Titan** - Cray XK7, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x , Cray Inc. DOE/SC/Oak Ridge National Laboratory United States	560,640	17,590.0	27,112.5	8,209
5	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom , IBM DOE/NNSA/LLNL United States	1,572,864	17,173.2	20,132.7	7,890
6	**Cori** - Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect , Cray Inc. DOE/SC/LBNL/NERSC United States	622,336	14,014.7	27,880.7	3,939
7	**Oakforest-PACS** - PRIMERGY CX1640 M1, Intel Xeon Phi 7250 68C 1.4GHz, Intel Omni-Path , Fujitsu Joint Center for Advanced High Performance Computing Japan	556,104	13,554.6	24,913.5	2,719
8	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect , Fujitsu RIKEN Advanced Institute for Computational Science (AICS) Japan	705,024	10,510.0	11,280.4	12,660
9	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom , IBM DOE/SC/Argonne National Laboratory United States	786,432	8,586.6	10,066.3	3,945
10	**Trinity** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect , Cray Inc. DOE/NNSA/LANL/SNL United States	301,056	8,100.9	11,078.9	4,233

Distributed Data Management

"The ability to efficiently read, transform, and store large amounts of data!"

- Static (block) data
- Volatile (streaming) data

Processing

Motivation: “Data Processing”

Data Processing in Practice

Startup

- Everyone starts small but be prepared for the hype!

Motivation: “Data Processing”

Data Processing in Practice

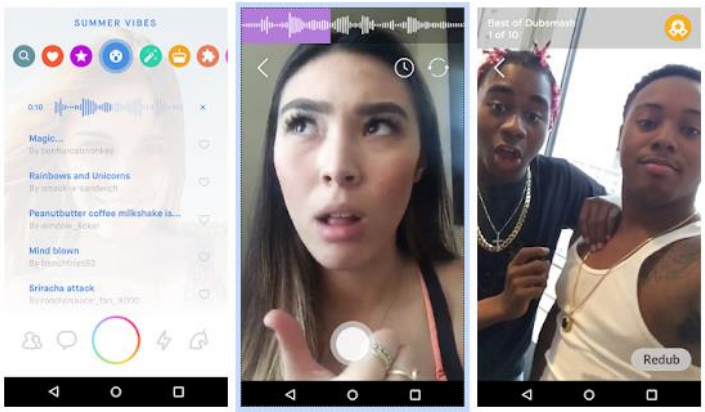

Example: Mobile Motion GmbH

Dubsmash

- An HPI-Startup of 2013
- Founders:
  - Jonas Drüppel, Roland Grenke, Daniel Taschik

November 19, 2014:
- **Launch** of the Dubsmash app

November 26, 2014:
- Dubsmash reached the **number one** downloaded app in Germany

June 1, 2015:
- Dubsmash had been downloaded over **50 million times in 192 countries**

Motivation: “Data Processing”

Data Processing in Practice

Distributed Data Management

Introduction

Data Processing in Practice

Successful IT-Startups in recent years are masters of data:

1. **AirBnB**
2. **Instagram**
3. **Pinterest**
4. **Angry Birds**
5. **Linkedin**
6. **Uber**
7. **Snapchat**
8. **WhatsApp**
9. **Twitter**
10. **Facebook**
11. **…**

Peta- to Exabytes of …

- profile data (names, addresses, friends, …)
- content data (images, videos, messages, …)
- event data (logins, interactions, games, …)
- …

Challenged with …

- streaming
- persistence
- analytics
- load-balancing
- …

Motivation: "Data Processing"

Driving Forces

- Data volumes increase:
  business data, sensor data, social media data, …
- Data analytics gains importance:
  downtime-less, real-time, predictive
- Parallelization paradigm shifts:
  multi-core and network speeds increase while CPU clock speeds stall
- Computation resources become more available:
  IaaS, PaaS, SaaS
- Free and open source software gains popularity:
  setting standards, utilizing external development resources, improving software quality, avoiding vendor locks …

Related Topics

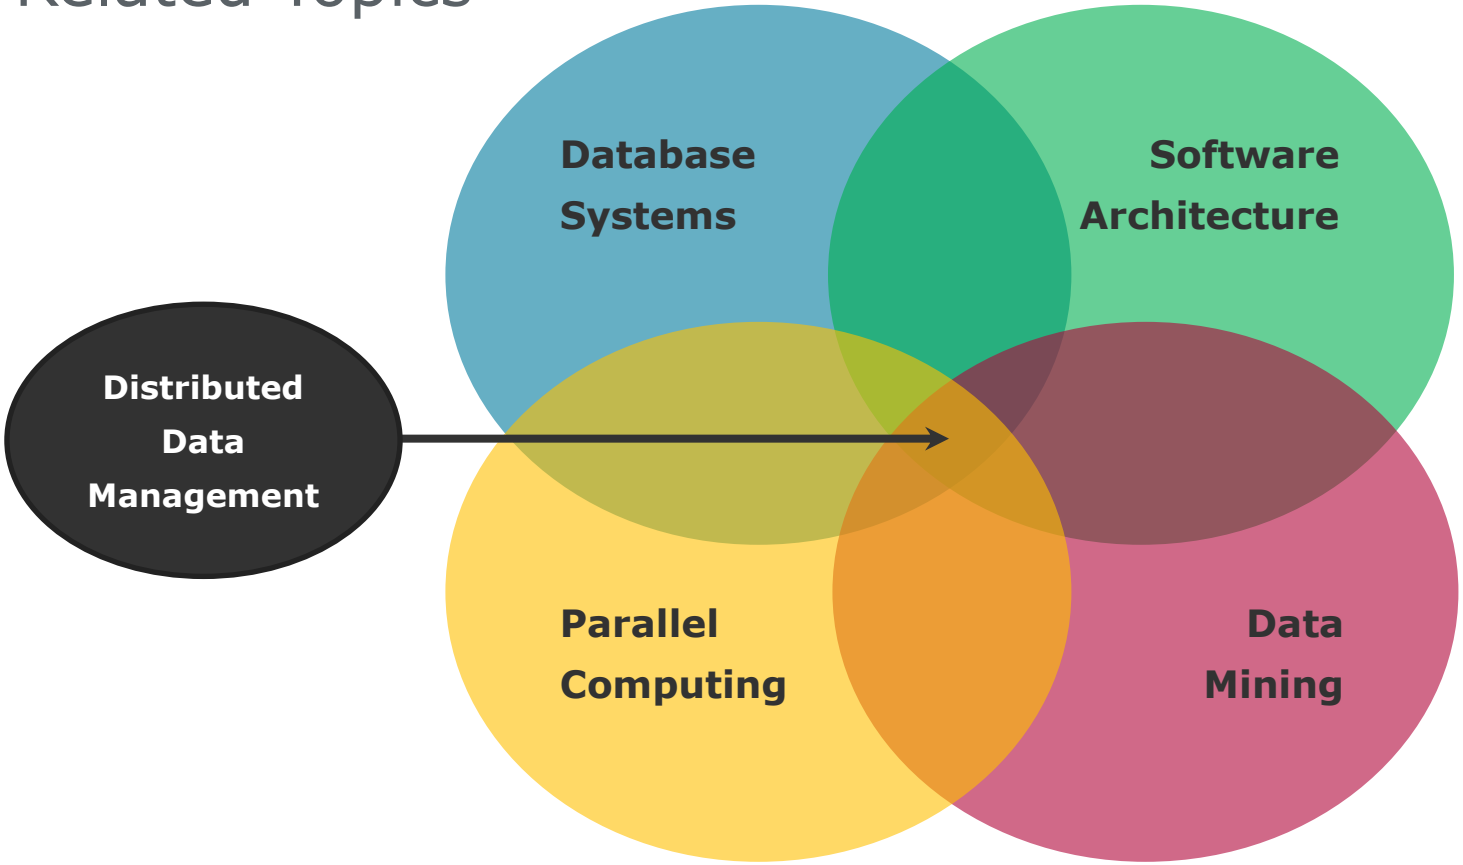

Related Topics

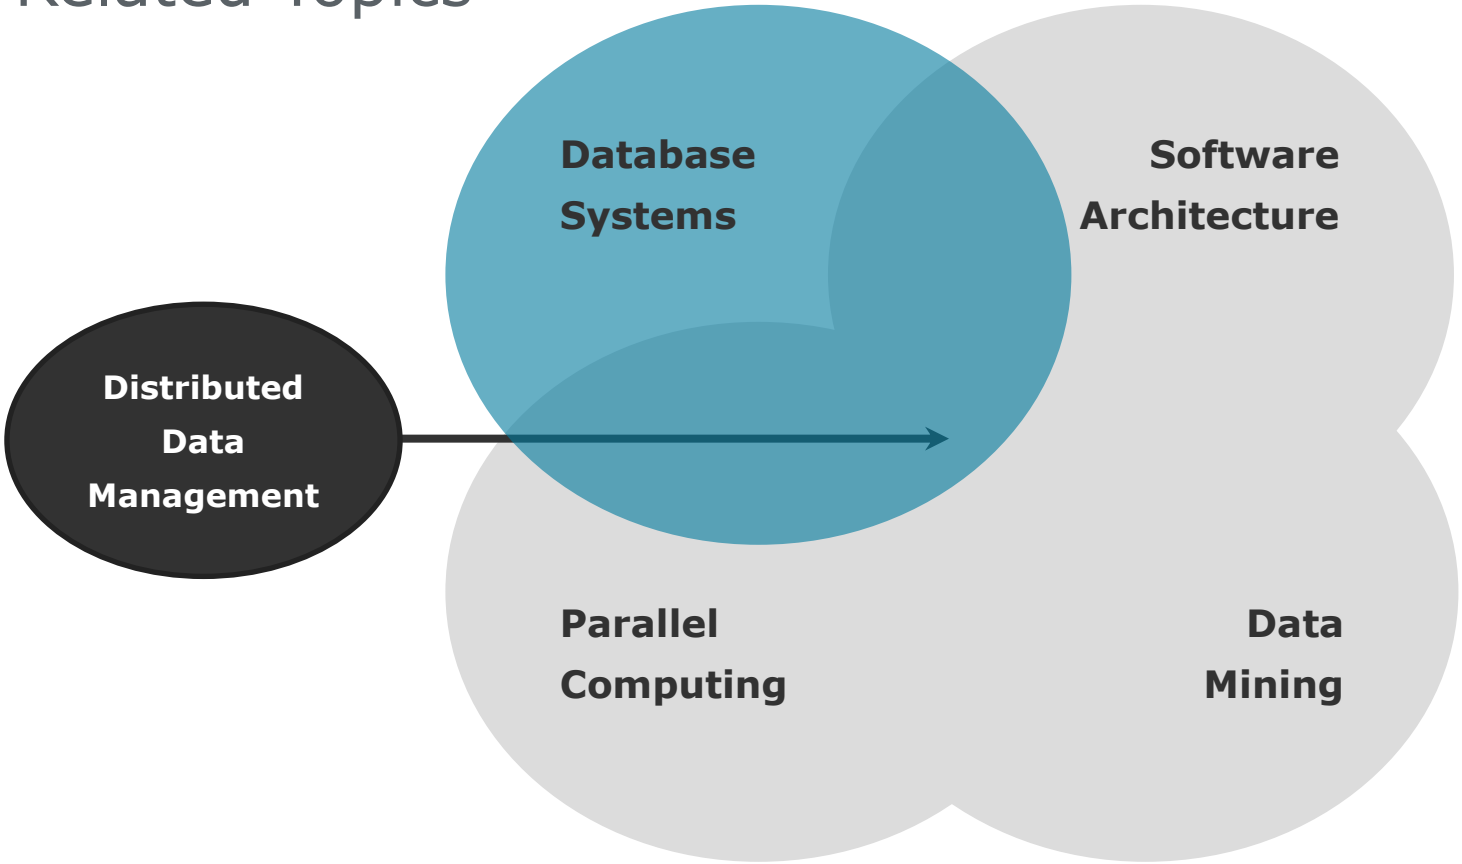

Database Systems

Touch points

- Data models, query languages, and consistency guarantees
- Distributed storage and retrieval of data
- Index structures

Not in this lecture

- Physical data storage
- In-depth presentations language constructs and transactions
- Core database technology, e.g., query optimizer

More focused lectures

- Database Systems I + II (Prof. Naumann)
- Trends and Concepts in Software Industry (Prof. Plattner)

Related Topics

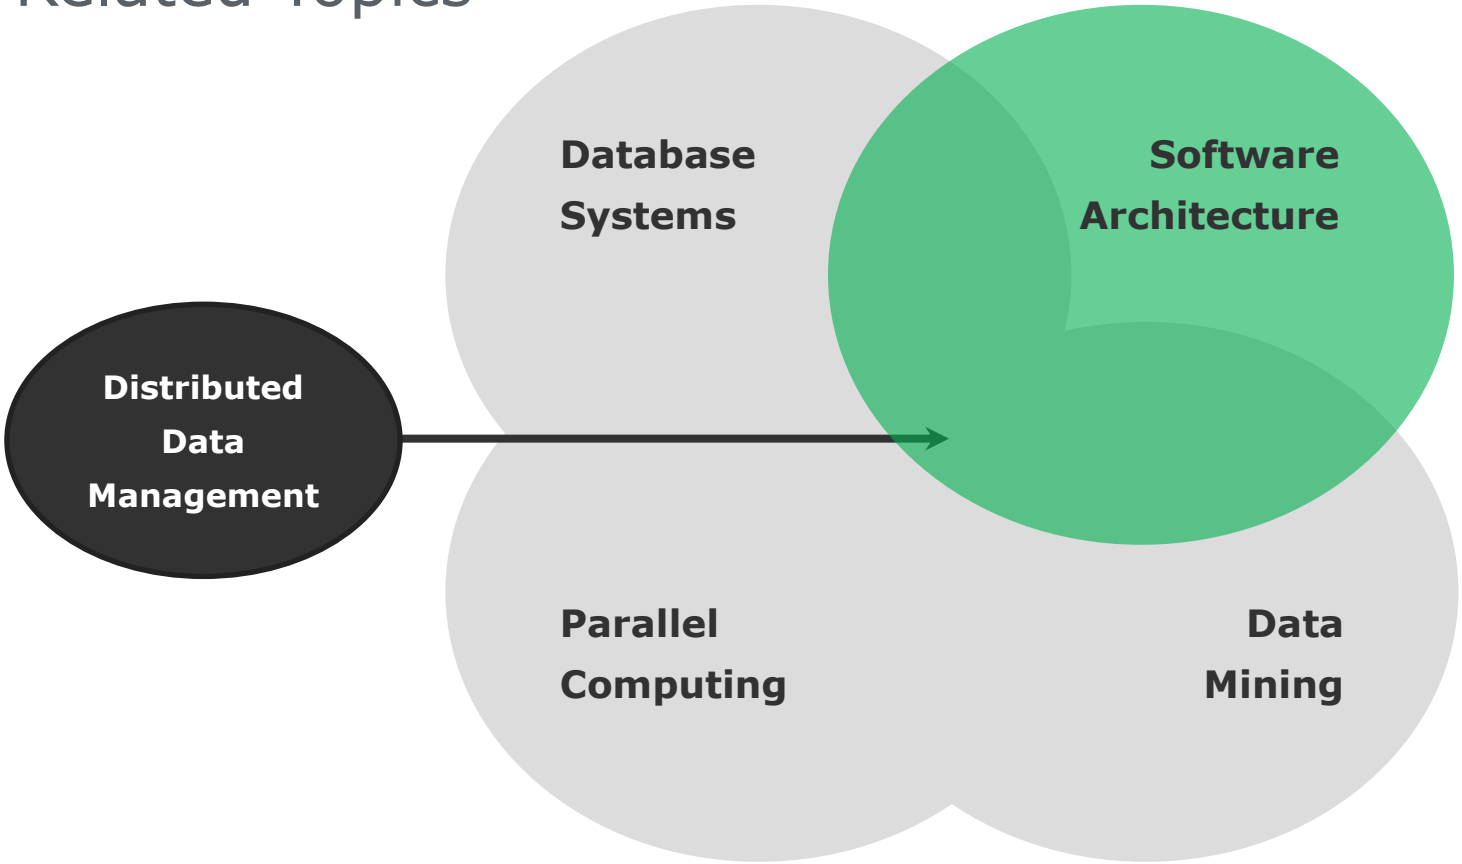

Software Architectures

Touch points

- Requirements, design, and architecture of distributed systems
- Pros and cons of different technologies for distributed systems

Not in this lecture

- Non-distributed systems
- Agile software development techniques
- Software patterns

More focused lectures

- Software Architecture (Dr. Uflacker)
- Software Technique (Dr. Uflacker)

Related Topics

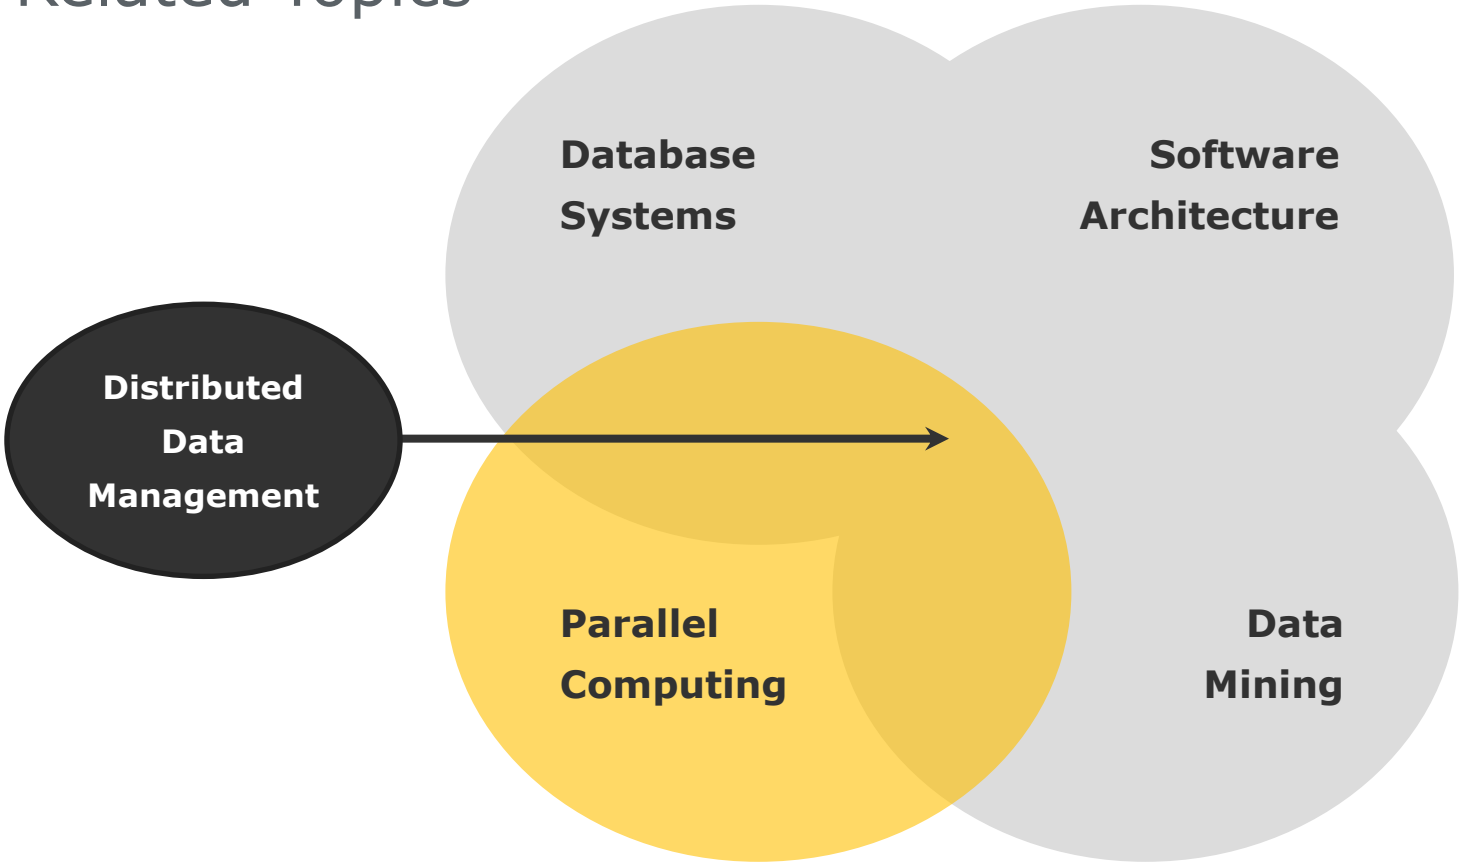

Parallel Computing

Touch points

- Distributed data storage concepts
- Distributed programming models, e.g., actor programming and MapReduce

Not in this lecture

- Parallel, non-distributed programming languages, e.g., CUDA or OpenMP
- Core parallel computing concepts, e.g., scheduling or shared memory
- Processor architectures, cache hierarchies, GPU programming, …

More focused lectures

- Parallel Programming (Dr. Tröger)
- Programmierung paralleler und verteilter Systeme (Dr. Feinbube)

Related Topics

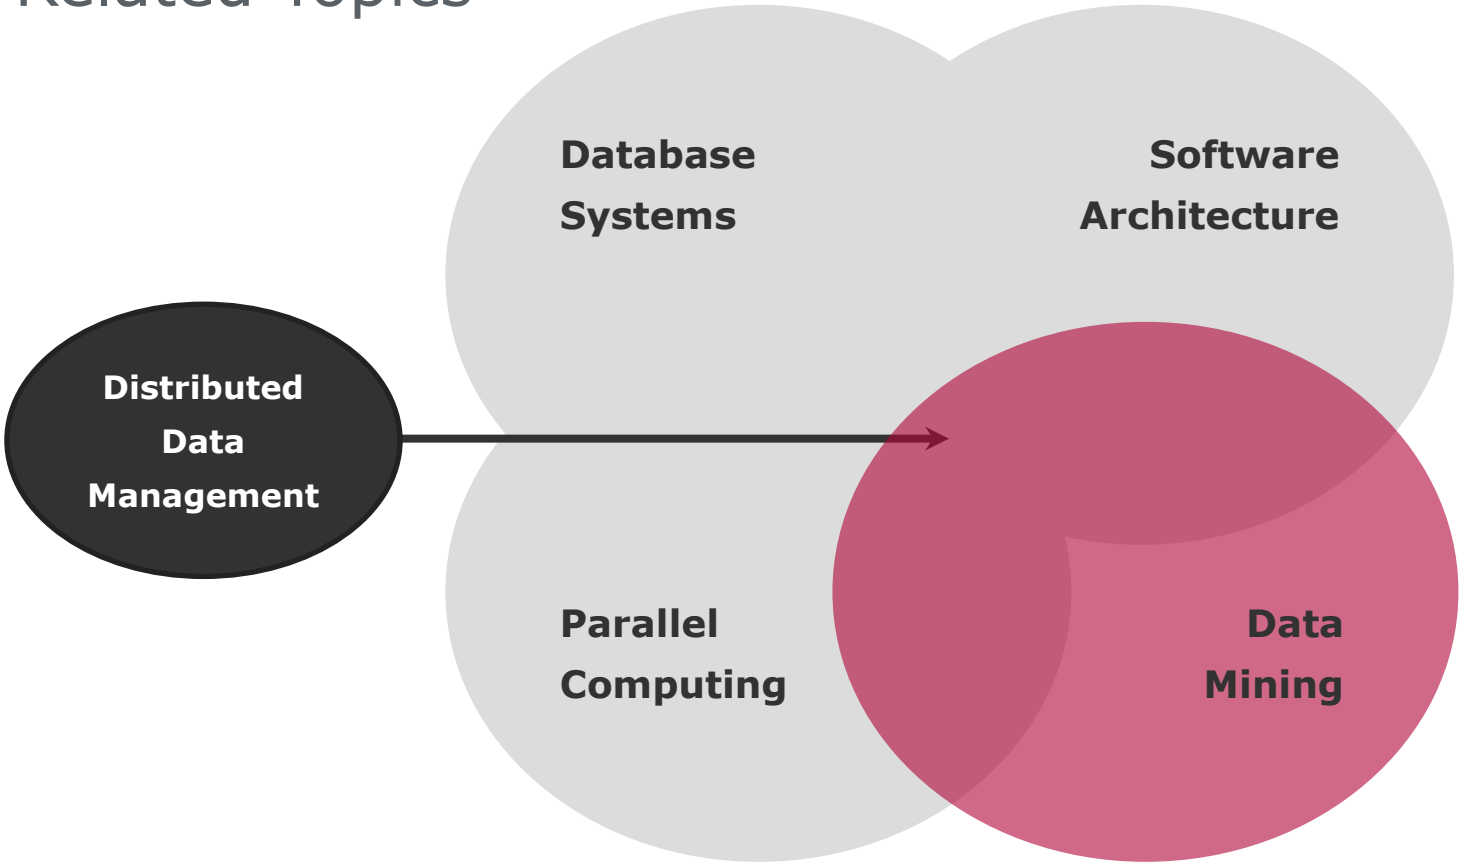

Data Mining

Touch points

- Data analytics: aggregation queries and basic data mining algorithms

Not in this lecture

- Detailed introduction to machine learning, e.g., neuronal networks, (un)supervised learning, or Bayesian classification
- Statistics, linear algebra, and most sophisticated mining algorithms

More focused lectures

- Big Data Analytics (Prof. Müller)
- Graph Mining (Prof. Müller)
- Data Profiling (Prof. Naumann)
- Data Mining and Probabilistic Reasoning (Dr. Krestel)

Related Topics

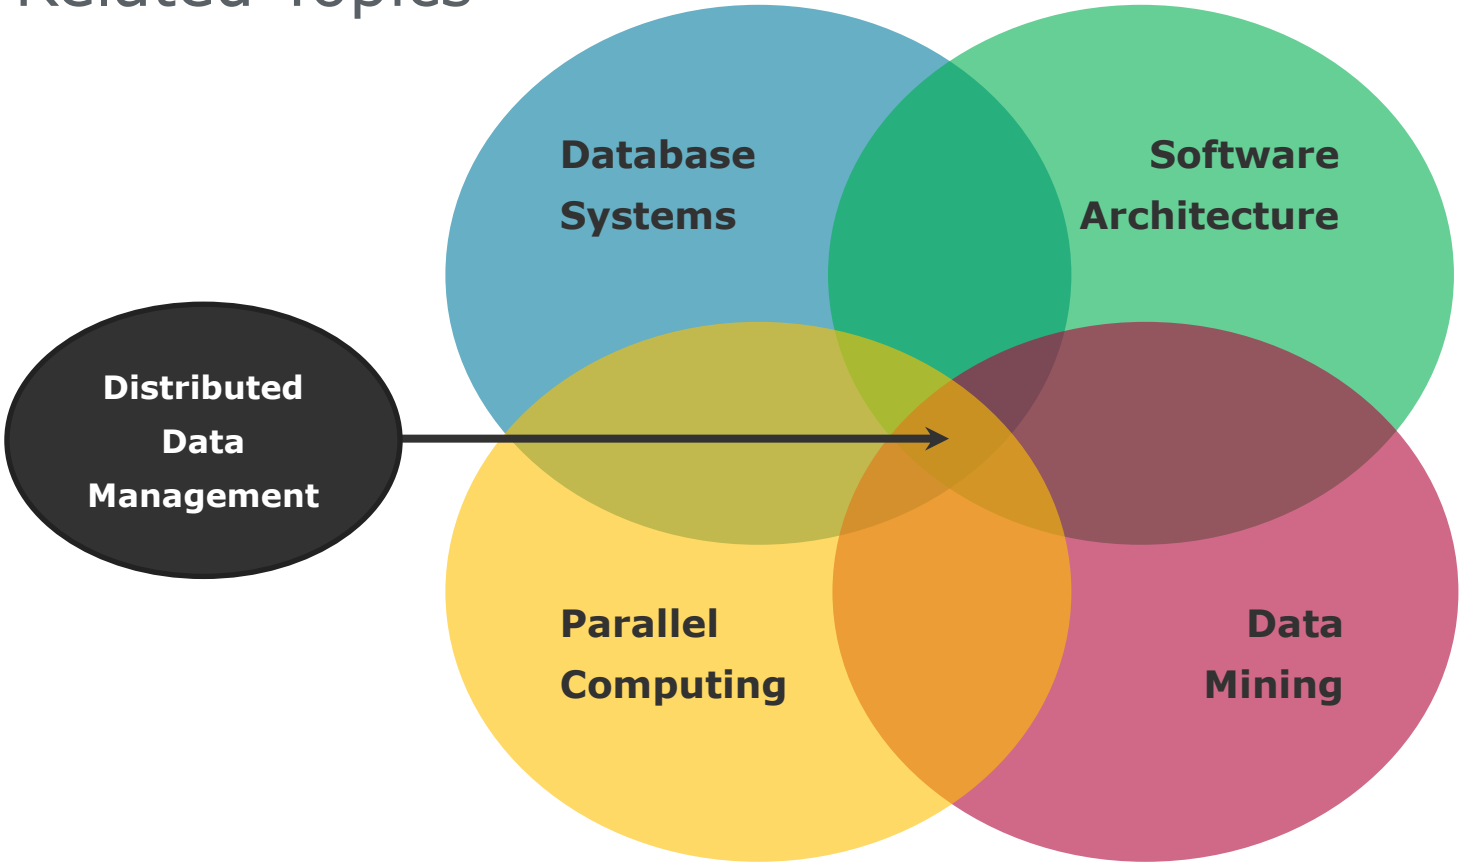

Lecture Goals

Sorting the buzzwords

- NoSQL, Big Data, OLAP, Web-scale, ACID, Sharding, MapReduce, …

Understanding the systems

- How and why do distributed systems work?
- What are their individual core technologies?
- Which advantages and disadvantages do the technologies have?
- How can different technologies be combined?

Exercising in data analytics

- Distributed processing of data-parallel tasks
- Data mining on streams

Lecture Outline

1. Introduction
2. Foundations
3. OLAP and OLTP
4. Encoding and Evolution
5. Hands-On: Akka
6. Data Models and Query Languages
7. Storage and Retrieval
8. Replication
9. Partitioning
10. Batch Processing
11. Hands-On: Spark
12. Distributed Systems
13. Consistency and Consensus
14. Transactions
15. Stream Processing
16. Hands on: Flink
17. Mining Data Streams
18. Distributed Algorithms
19. Services and Containerization
20. Cloud-based Data Systems
21. Lecture Summary and Exam Preparation

On-site and by **dataArtisans**!

Lecture Outline – Homework

1. Introduction
2. Foundations
3. OLAP and OLTP
4. Encoding and Evolution
5. Hands-On: Akka
6. Data Models and Query Languages
7. Storage and Retrieval
8. Replication
9. Partitioning
10. Batch Processing
11. Hands-On: Spark

12. Distributed Systems
13. Consistency and Consensus
14. Transactions
15. Stream Processing
16. Hands on: Flink
17. Mining Data Streams
18. Distributed Algorithms
19. Services and Containerization
20. Cloud-based Data Systems
21. Lecture Summary and Exam Preparation

Literature: Course Book

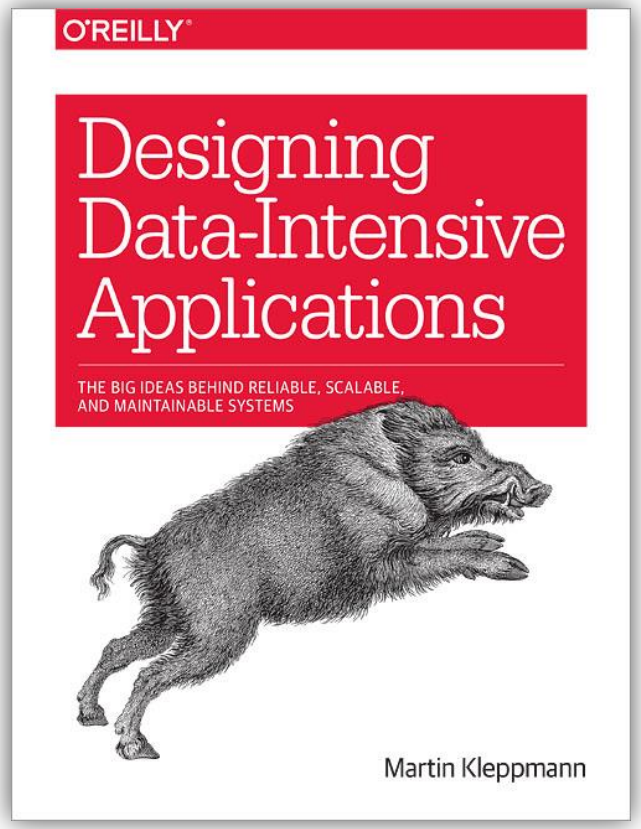

Designing Data-Intensive Applications

- Author: Martin Klappmann
- Date: March 2017
- Publisher: O'Reilly Media, Inc
- ISBN: 978-1-449-37332-0
- References: https://github.com/ept/ddia-references

Scope for this lecture

- Distributed and parallel systems for data analytics and data mining
- Big data storage
- Batch and stream processing

Literature: Further Reading

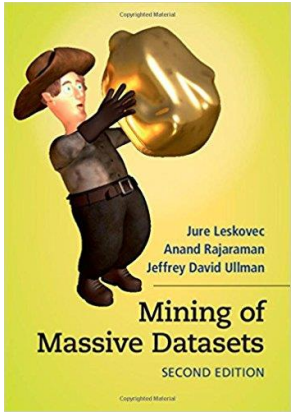

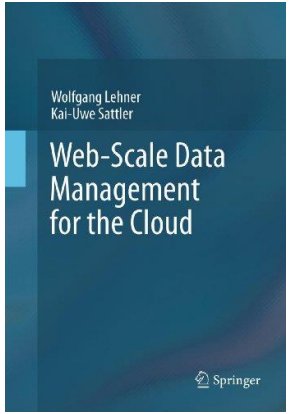

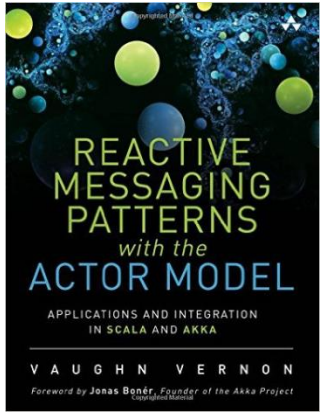

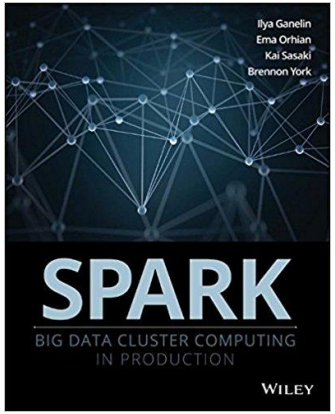

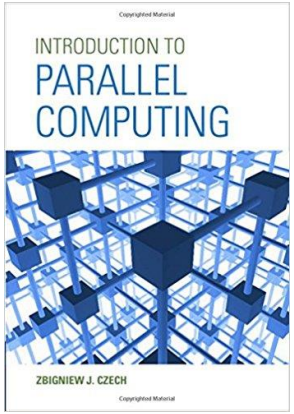

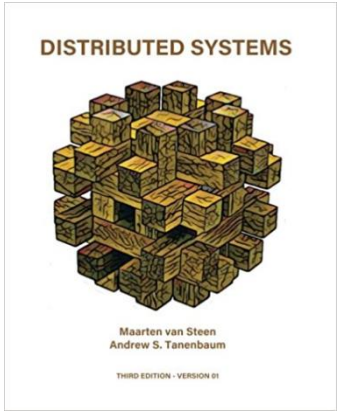

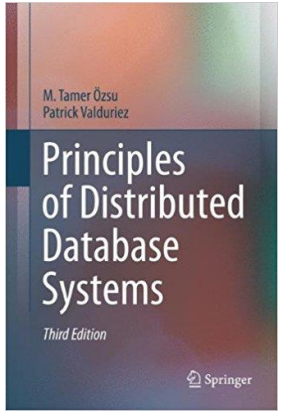

Web-links are given on the slides during the lecture.

Feedback: Distributed Data Analytics

EvaP

Winter term 2017/2018 / Distributed Data Analytics

Distributed Data Analytics (Winter term 2017/2018)

View as | Myself | Public

Overview | Final grades

Degree	Type	Responsible	Voters	Distribution and average grade
MA ITSE	Lecture	Prof. Dr. Felix Naumann	18/28 64%	1.5

I learned a lot in the course.	17	1.7
… gave the course in an interesting way.	17	1.9
I was satisfied how the exercise submission worked.	17	1.8

Feedback: Distributed Data Analytics

EvaP

Winter term 2017/2018 / Distributed Data Analytics

What did you like?	• Ich fand gut, dass du im Laufe der Veranstaltung nicht mehr so nervös warst wie am Anfang. • Thorsten hat eine sehr ruhige Art, was es möglich macht ihm gut zuzuhören. Man hat auch das Gefühl, dass er wirklich Herzblut in die Veranstaltung gesteckt hat. Dankeschön! • Du hast eine sehr angenehme und ruhige Art zu präsentieren. Man merkt, dass dir die Studenten und die Lehre wichtig sind. • Did not let nervousness affect the lecture, did grow more and more comfortable. I feel like Thorsten hast a lot of know how in the area. • Du bist sehr freundlich und sehr respektvoll gegenüber den Studenten. Bitte behalte dies bei! :)
What could be improved?	• Maybe try a more project based approach as well instead of the two hands ons? Or a project seminar that fits the lecture. • Klarer skizzieren wieso bestimmte Vorlesungspunkte wichtig sind. Wie würde man Spark o.ä. in einer kompletten Pipeline verwenden? Also angenommen man erhält kontinuierlich Daten und soll diese analysieren und das Ergebnis ausgeben - wie kann man dies erreichen?

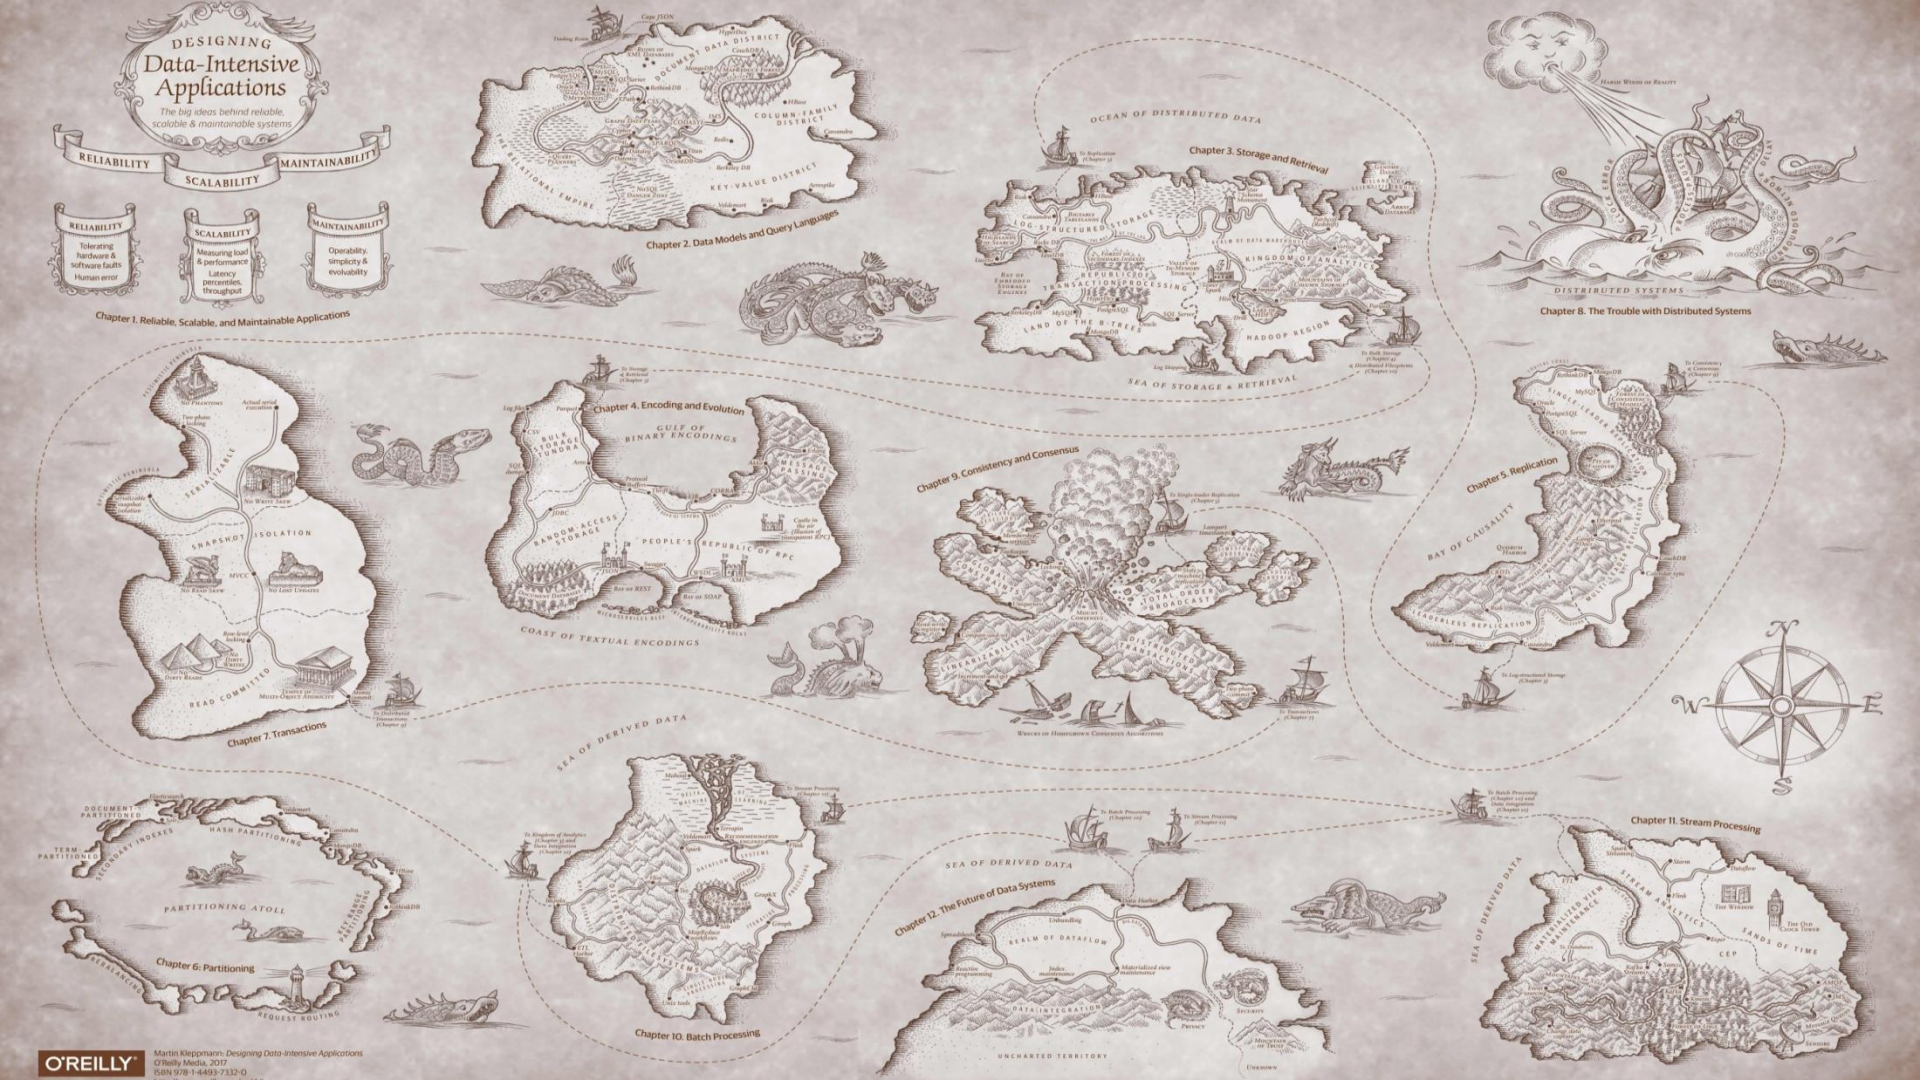

O'REILLY
Martin Kleppmann: *Designing Data-Intensive Applications*
O'Reilly Media, 2017
ISBN 978-1-4493-7332-0