

# Systems Infrastructure for Data Science

Web Science Group

Uni Freiburg

WS 2013/14

# Introduction

# About myself

- Since Fall 2011:  
Assistant/Junior Professor for Web Science
- Before Senior Researcher (“Oberassistent”) at  
Systems Group, ETH Zurich
- Research Interests:
  - Social Media Analytics
  - Realtime Analytics
  - Temporal Data Management
  - Stream/Event Processing
- Systems-oriented approach

Contact: [peter.fischer@cs.uni-freiburg.de](mailto:peter.fischer@cs.uni-freiburg.de)

# Basic Course Information

- Credits: 3V + 1U (= 6 ECTS)
- Language: English (feel free to ask in German)
- Time and Location:
  - Tuesday 16:15-18:00 (HS 00 006, Building 082)
  - Thursday 14:15-15:00 (HS 00 006, Building 082)
  - Exercise: Thursday 15:00-15:45 (also HS 00 006)
- Webpage:
  - <https://websci.informatik.uni-freiburg.de/teaching/ws201415/infosys/infosys>

# Workload & Grading

- Exercises
  - Weekly exercise sheets with questions related to the lecture coverage
  - Not graded
  - Attendance to exercise sessions is not mandatory, but it is highly recommended to do well in the exam.
- Exam
  - No prerequisites to participate
  - Written or oral dependent on number of participants

# Exercise Sessions

- Content
  - Explain the new exercise sheet
  - Provide solutions for the previous exercise sheet
  - Answer your questions
- First sheets will be made available on October 23rd
- The first exercise sessions will take place on October 30<sup>th</sup>

# Course Objectives

- Overall Objective:
  - to understand how different platforms for data management and analysis work
- Partial goals:
  - Understand the internals of the architecture, implementation, and optimization of a relational database system
  - Understand the basics of distributed data processing
  - Understand concepts and implementations of novel platforms

# Course Motivation: New Analytics

- No longer just structured, “clean” business data:
  - Text data, photos, videos
  - Social media: social networks, social streams
  - Science
  - ...
- Much broader range of analytics
  - Information Retrieval
  - Machine Learning: Classification, Mining
  - Statistics
  - Human Interaction: Crowdsourcing, Interactive exploration
- Much larger volumes (think Google, Facebook!)
- Unpredictable workloads
- Results required in real time



# Course Motivation: New Platforms

- Increasing CPU core count: Massive Parallelism
- Increasing RAM, “slower” disks, new storage
- Faster Networks and massive Distribution
  - Racks and Datacenters as new basic building blocks
  - Global Replication, Consistency and Access
- New Processing paradigms:
  - Map/Reduce
  - Key/Value Stores
  - Event, Data Stream Processing

# Topics

- Classical Databases: “Complete” package for moderate workloads
  - Storage and Indexing
  - Query Processing and Optimization
  - Performance Tuning and Benchmarking
- Distributed and Parallel Databases: Scaling with DB means
  - General Architecture
  - Distribution
  - Query Processing
- Map-Reduce: Highly scalable, unstructured data, simple programming model
- Key-Value Stores: Storing and retrieving data at speed and scale
- Stream Processing: Processing instantly without storing
- (Graph Processing, Machine Learning)
- ...

# Relation to other lectures

- Information Systems:
  - DB Intro: Foundations, Transactions
  - Distributed Systems: focused on data consistency and distributed transactions
  - Data Models and Query Languages: covers models, languages and theory
- Other areas:
  - Operating Systems, Networks: same foundations, sometimes same problems

# Starting point: Classical DB

- Still useful for moderate-sized workloads (few TBs, standard queries, transactions)
- Guidepost for technologies
  - Nearly all aspects of data management covered
  - Decades of experience and refinement
  - (Many aspects being re-discovered by “cool new platforms”)

# Architecture of a Database System

