

# Cloud Computing (Seminar)

Dipl.-Inf. Thomas Hornung

20. April 2010

# Administrative Preamble

## Outline

- **ECTS:** 4
- **Room:** SR 01-016, Building 101
- **Language:** German/English
- **Time:** Wednesday, 2 p.m. - 4 p.m. c.t. weekly (starting with 1<sup>st</sup> talk)

# Administrative Preamble cont'ed

## Grade Criteria

- Attendance in all presentations
- Written report handed in until July 24th, 2010 (latest)
- Presentation of approx. 45 min (either in German or English), where:
  - approx. 35 minutes presentation, and
  - approx. 10 minutes questions and discussion
- Preparation of two questions for two different talks (i.e. 4 questions) are randomly assigned (have to be send by e-mail to the adviser two days before the respective talk)
- Appointment with adviser to discuss slides at latest 2 weeks prior the presentation date
- Lastly, grading is based on 70% presentation and 30% written report

# Paper 1 – 5

- M. Isard et. al: **Dryad: Distributed Data-Parallel Programs from Sequential Building Blocks**. EuroSys 2007 ([open](#))  
→ Presentation date: 16/06/2010
- J. Dean and S. Ghemawat: **MapReduce: Simplified Data Processing on Large Clusters**. OSDI 2004 ([Asha Nagendra](#))  
→ Presentation date: 16/06/2010
- H. Yang et. al: **Map-Reduce-Merge: Simplified Relational Data Processing on Large Clusters**. SIGMOD 2007 ([open](#))  
→ Presentation date: 23/06/2010
- F. N. Afrati and J. D. Ullman: **Optimizing Joins in a Map-Reduce Environment**. EDBT 2010 ([open](#))  
→ Presentation date: 23/06/2010
- R. Pike et. al: **Interpreting the Data: Parallel Analysis with Sawzall**. Scientific Programming 2005 ([open](#))  
→ Presentation date: 30/06/2010

## Paper 6 – 10

- C. Olston et. al: **Pig Latin: a Not-so-foreign Language for Data Processing**. SIGMOD 2008 (**Alibek Kulzhabayev**)  
→ Presentation date: 30/06/2010
- J. Urbani et. al: **Scalable Distributed Reasoning using MapReduce**. ISWC 2009 (**open**)  
→ Presentation date: 07/07/2010
- S. Matthews and T. Williams: **MrsRF: an Efficient MapReduce Algorithm for Analyzing Large Collections of Evolutionary Trees**. APBC 2010 (**open**)  
→ Presentation date: 07/07/2010
- A. Abouzeid et. al: **HadoopDB: An Architectural Hybrid of MapReduce and DBMS Technologies for Analytical Workloads**. VLDB 2009 (**open**)  
→ Presentation date: 14/07/2010
- E. Friedman et. al: **SQL/MapReduce: a Practical Approach to Self-describing, Polymorphic, and Parallelizable User-defined Functions**. VLDB 2009 (**Raphael Winkler**)  
→ Presentation date: 14/07/2010