



Wintersemester 2012/2013

Seminare

Bachelor Informatik CS 3702
„Datenbanken und Anfrageverarbeitung“

Master Informatik
„Advanced Topics of Database Systems“

CS 5840 - Fachübergreifende Kompetenzen = englischsprachiges Seminar

**CS 5480 - für den Bereich Software Systems Engineering = Seminar
Software Systems Engineering**

Privatdozent Dr. rer. nat. habil. Sven Groppe

groppe@ifis.uni-luebeck.de

Dipl.-Inf. Stefan Werner

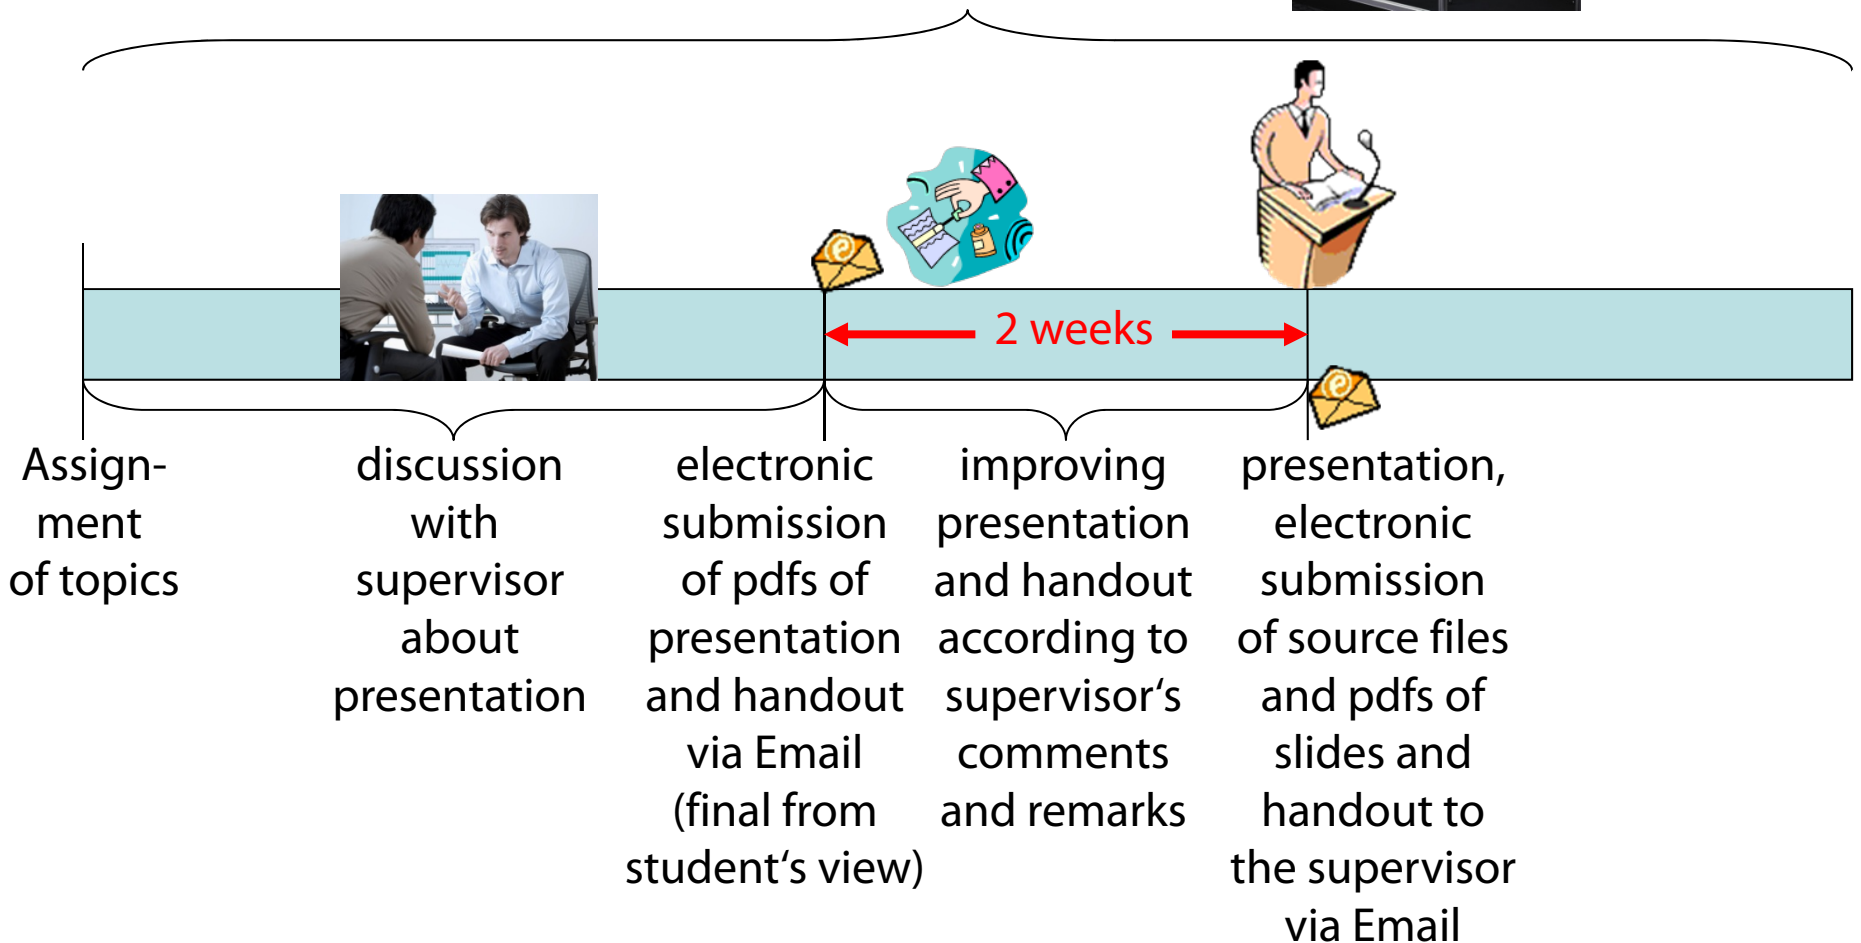
werner@ifis.uni-luebeck.de

Students' Duties

- Preparation of Slides
- Preparation of Handout
 - 2 to 3 pages, to be delivered to all participants and to the supervisors directly before the presentation
- Presentation
 - Approx. 1 hour (inclusive discussion)
- Attending presentations of all other students
 - Contributions to a lively discussion after each presentation

Timeline

participating in all other presentations
and contributing to lively discussions



Topics Bachelor/Master Seminars

- Green topics -> seminar for bachelor students
- Blue topics -> seminar for master students
- Topics may be exchanged to be discussed with the supervisors
 - Bachelor students choosing blue topics
 - Master students choosing green topics




Overview Topics

- FPGA
- Cloud Computing
- Semantic Web
- B⁺-tree
- Compression
- Regular Expressions
- Relational Databases
- Recursive Queries

FPGA (Field Programmable Gate Arrays)

- integrated circuit (IC)
- configuration after manufacturing
- hardware description language (HDL)
- complex functional blocks, arranged in periodic structure
- interconnection network



FPGA (Field Programmable Gate Arrays)

- Advantages towards General CPUs
 - inherent parallelism
 - heavy throughput at low clock rate
 - high energy efficiency
- Reconfigurable computing = offload „expensive“ tasks from software to FPGA

FPGA (Field Programmable Gate Arrays)

- Topics -

- **A Memory-Balanced Linear Pipeline Architecture for Trie-based IP Lookup (HOTI 2007)**
- **High Throughput and Large Capacity Pipelined Dynamic Search Tree on FPGA (FPGA 2010)**
- **Massively Parallel XML Twig Filtering Using Dynamic Programming on FPGAs (ICDE 2010)**
- **How Soccer Players Would Do Stream Joins (SIGMOD 2011)**
- **Evaluating FPGA-acceleration for Real-time Unstructured Search (ISPASS 2012)**
- **Skeleton Automata for FPGAs: Reconfiguring without Reconstruction (SIGMOD 2012)**

Cloud Computing

- typically, cloud provider offers online-services
- multiple server-based computational resources via a digital network like internet
- applications are provided and managed by the cloud server
- data is stored remotely in the cloud configuration



Cloud Computing

Having secure access to all your applications and data from any network device

Cloud computing - Topics

- **Walnut: A Unified Cloud Object Store (SIGMOD 2012)**
 - Developed at Yahoo! with the goal of serving as a common low-level storage layer for a variety of cloud data management systems (e.g. Hadoop, MObStor, PNUTS)
=> Enabling sharing hardware resources between different types of clouds.
- **CloudRAMSort: Fast and Efficient Large-Scale Distributed RAM Sort on Shared-Nothing Cluster (SIGMOD 2012)**
 - Sort algorithm, which sorts 1 Terabyte of data in 4.6 seconds
- **PIQL: Success-Tolerant Query Processing in the Cloud (VLDB 2012)**
 - Newly-released web applications often have a „Success Disaster“, where overloaded database systems destroy a previously good user experience, i.e. too much success causes performance problems.
 - The authors propose PIQL for providing *scale independance* for queries, such that applications become *success-tolerant*.

Semantic Web

- Idea
 - "web of data" that enables machines to understand the semantics, or meaning, of information on the World Wide Web.
 - extends the network of hyperlinked human-readable web pages by inserting machine-readable metadata about pages and how they are related to each other
- Semantic Web databases can be seen as graph databases for labelled and directed graphs
- Additional implicit data is considered based on ontology inference
 - E.g. *Peter* has the *driverLicenseNo 456* (which is a fact) and only a person can have a *driverLicenseNo* (expressed in an ontology), then *Peter* is a *person* (implicitly present)

Semantic Web - Topics


- **DBpedia SPARQL Benchmark –Performance Assessment with Real Queries on Real Data (ISWC 2011)**
 - DBpedia contains data extracted from Wikipedia
 - Benchmark queries are determined from frequently used queries for a SPARQL querying web service (*endpoint*)
- **ANAPSID: An Adaptive Query Processing Engine for SPARQL Endpoints (ISWC 2011)**
 - ANAPSID provides physical operators that detect when a source becomes blocked or data traffic is bursty, and produce results as quickly as data arrives from the sources
- **Castor: A Constraint-Based SPARQL Engine with Active Filter Processing (ESWC 2011)**
 - Constraint Programming was originally designed for solving NP-hard problems, but the authors adapt it to SPARQL processing by exploiting filters early-on instead of a post-processing step

Semantic Web - Topics

- **View Selection in Semantic Web Databases (VLDB 2012)**
 - Use existing views (=materialized results of previous queries) to answer a new query
 - Simple ontology inference is considered
- **The Complexity of Evaluating Path Expressions in SPARQL (PODS 2012)**
 - SPARQL 1.1 allows to query nodes in arbitrary depth. This paper studies its complexity for W3C-conform evaluation and presents an alternative.
- **Static Analysis and Optimization of Semantic Web Queries (PODS 2012)**
 - Study of Containment, Equivalence and Evaluation of SPARQL queries

B⁺-tree

- Variant of the B-tree
- Self-balancing, block-oriented search tree for large datasets



B⁺-tree - Topics

- **B+-tree Index Optimization by Exploiting Internal Parallelism of Flash-based Solid State Drives (VLDB 2011)**
 - B⁺-tree variant for SSDs using parallel I/O

Compression - Topics

- **Relative Lempel-Ziv Factorization for Efficient Storage and Retrieval of Web Collections (VLDB 2011)**
 - Instead of window-based dictionaries, a dictionary is generated by sampling the whole document

Regular expressions

- Regular expressions contain operators for
 - concatenation,
 - union,
 - choice and
 - Kleene star
- Used to define searches/matchings in
 - file systems
 - databases
 - programming languages
(e.g. tokens while compiling and api support)
 - ...

Regular Expressions - Topics

- **Deterministic Regular Expressions in Linear Time (PODS 2012)**
 - Algorithms to test determinism and matching

Relational Database

- Most widely used type of database
- Data model
 - Relation (table)

<i>Students</i>			<i>Lectures</i>	
Name	ID	Adress	ID	Lecture-ID
student1	1	HL	1	DB1
student2	2	HH	2	MVDB2

- Query language
 - SQL

Relational Database - Topics

- **NoDB: Efficient Query Execution on Raw Data Files (SIGMOD 2012)**
 - NoDB avoids data loading while still maintaining the whole feature set of modern databases, i.e., NoDB works on raw data files.
- **SharedDB: Killing One Thousand Queries With One Stone (VLDB 2012)**
 - deals with multi-query optimization, i.e., query execution plans for many queries to be executed at the same time are generated.
 - High throughput for high workloads

Recursive Queries

- **Datalog as subset of prolog is a rule language for deductive databases**

- **Example:**

Facts: parent(bill,mary).
 parent(mary,john).

Rules: ancestor(X,Y) :- parent(X,Y).
 ancestor(X,Y) :- parent(X,Z), ancestor(Z,Y).

Query: ?- ancestor(bill,X).

- **Topic:**

More Efficient Datalog Queries: Subsumptive Tabling Beats Magic Sets (SIGMOD 2011)

- top-down evaluation method with more reuse of answers than the dominant tabling strategy

Recursive Queries

- **Recursive query constructs part of ANSI SQL 99**

- **Topic:**

Adaptive Optimization of Recursive Queries in Teradata (SIGMOD '12)

- Query execution plan is optimized within iterations of recursive queries

FPGA

- **A Memory-Balanced Linear Pipeline Architecture for Trie-based IP Lookup (HOTI 2007)**
- **High Throughput and Large Capacity Pipelined Dynamic Search Tree on FPGA (FPGA 2010)**
- **Massively Parallel XML Twig Filtering Using Dynamic Programming on FPGAs (ICDE 2010)**
- **How Soccer Players Would Do Stream Joins (SIGMOD 2011)**
- **Evaluating FPGA-acceleration for Real-time Unstructured Search (ISPASS 2012)**
- **Skeleton Automata for FPGAs: Reconfiguring without Reconstruction (SIGMOD 2012)**

Cloud Computing

- **Walnut: A Unified Cloud Object Store (SIGMOD 2012)**
- **CloudRAMSort: Fast and Efficient Large-Scale Distributed RAM Sort on Shared-Nothing Cluster (SIGMOD 2012)**
- **PIQL: Success-Tolerant Query Processing in the Cloud (VLDB 2012)**

Semantic Web

- **DBpedia SPARQL Benchmark –Performance Assessment with Real Queries on Real Data (ISWC 2011)**
- **ANAPSID: An Adaptive Query Processing Engine for SPARQL Endpoints (ISWC 2011)**
- **Castor: A Constraint-Based SPARQL Engine with Active Filter Processing (ESWC 2011)**

Semantic Web

- **View Selection in Semantic Web Databases (VLDB 2012)**
- **The Complexity of Evaluating Path Expressions in SPARQL (PODS 2012)**
- **Static Analysis and Optimization of Semantic Web Queries (PODS 2012)**

B⁺-tree

- **B+-tree Index Optimization by Exploiting Internal Parallelism of Flash-based Solid State Drives (VLDB 2011)**

Compression

- **Relative Lempel-Ziv Factorization for Efficient Storage and Retrieval of Web Collections (VLDB 2011)**

Regular Expressions

- **Deterministic Regular Expressions in Linear Time (PODS 2012)**

Relational Database

- **NoDB: Efficient Query Execution on Raw Data Files (SIGMOD 2012)**
- **SharedDB: Killing One Thousand Queries With One Stone (VLDB 2012)**

Recursive Queries

- **More Efficient Datalog Queries: Subsumptive Tabling Beats Magic Sets (SIGMOD 2011)**
- **Adaptive Optimization of Recursive Queries in Teradata (SIGMOD '12)**