

Introduction to Database Systems

Course Outline and Organisation

Werner Nutt

Aims

To be able to use Database Management Systems (DBMSs) successfully, one has to understand the concepts on which they are based.

The aims of this course are to

- familiarise you with the basic concepts underlying a DBMS;
- show how they are realized in specific systems such as the PostgreSQL DBMS;
- give you some hands-on experience in using a DBMS.

Course Content (1)

- Fundamental Database Concepts
- The Entity Relationship (= ER) Model
(the most common approach to conceptual database design)
- The Relational Data Model
 - Relations
 - Integrity Constraints (keys, foreign keys, etc.)
- Logical Database Design
(ER to relational schemas)
- Relational Algebra
(an algebraic query language for the relational model)

Course Content (2)

- SQL: Querying and Manipulating Data
 - SQL Data Definition Language
 - Single Block Queries
 - Aggregation
 - Joins and Outer Joins
 - Nesting
 - Negation
- Transaction Management and Concurrency Control
- Database Access from a Programming Language: JDBC

Course Content (3)

- Data Storage and Indexing
 - File Organisation and Indexes
 - Tree-structured Indexing: B+-trees
 - Hash-based Indexing
 - Indexes in PostgreSQL
- Query Evaluation
 - Sorting
 - Evaluation of Relational Operators
 - Query Optimisation
 - Physical Database Design

Course Content (4)

- Query Plans in PostgreSQL
- Functional Dependencies and Normalisation

Course Format

- Lectures
 - introduce new concepts, give examples
- Labs
 - exercises (→ preparation for exam questions)
 - support for group projects
- Group Projects
 - develop a toy database application
 - 3 students

Lectures

- Main textbook
A First Course in Database Systems
by Jeff Ullman and Jennifer Widom
- Lectures on data storage and indexing will follow
Database Management Systems
by Raghuram Ramakrishnan and Johannes Gehrke
- Slides will be made available at course web site
www.inf.unibz.it/~nutt/IDBs0910

Labs and Teaching Assistants

- Start in week 2
- Teaching assistants
 - Michail Kazimianec
 - Damiano Somenzi
- All members of a project group attend the same lab
- Support for projects during lab session
- Also, meetings by appointment

Group Projects

- Groups of three students
- You choose your group as you like
- Each group develops an “Individual Database Application” on a topic of their choice
- Each group will have a joint account on the faculty PostgreSQL server

Project Steps

Mimic the development of a “real” database

- Writing up **data requirements**
- Designing a **conceptual model** in the form of an Entity Relationship diagram
- Translating the conceptual model into a **relational schema**
- Implementing the relational schema in **PostgreSQL** and **populating** the database
- **Querying** and **modifying** the database by SQL statements
- Writing a **Java client** that accesses the DB via JDBC
- Optimising the access to data by adding **indexes** to the relational schema

Project Home Page

- Each group will set up a project home page with documents on the project and progress reports
 - **Special Web space** will be allocated on the faculty file server for each group
- There will be a **pointer** from the course home page to the project home pages so that you can learn from the work of other groups

Milestones

Week 2: Group registered and topic of project defined

Week 4: **Data requirements** and **conceptual model**

Week 6: Translation into **relational schema**, implementation of the schema, **population** of the schema with data

Week 8: SQL **queries** over the database

Week 10: **Physical design**, analysis of query execution plans, performance analysis

Week 12: **JDBC client** runs transactions on the database

Submission

- The deadline for each milestone is
Monday 10.30 hrs
following the respective week
- You submit your work by publishing it at your project home page
- You will receive a mark for the work that can be found at that time at your home page
- The tutors will also publish comments on your submissions (but marks are confidential)

Registration

- To register your group, send a mail to

kazimianec@inf.unibz.it

with an XML document (see course home page)
containing

- the names and email addresses of the group members
 - the topic of the project
 - a short description
- All groups have to work on different topics
 - If two groups choose the same topic, the group that registered second will have to choose a different one

Presentation in Lab

The lab in week 3 will be devoted to the conceptual model and the relational schema

- Each group gives a short presentation
- All participants of the tutorial are encouraged to discuss the projects presented
- The tutor asks questions about the planned project so that problems with the design can be identified at an early stage

Assessment

- Each project gets a mark, which is also the mark for each member of the group
- There will also be a written exam
- Final mark = $\max \{ 30\% \times \text{project} + 70\% \times \text{exam}, \text{exam} \}$