

INFO-H-417 : Database System Architecture Course Information

The objectives in brief

In contrast to a typical introductory course in database systems where one learns to design and query relational databases, the goal of this course is to get a fundamental insight into the implementation aspects of database systems. In particular, we take a look under the hood of relational database management systems, with a focus on query and transaction processing.

With respect to query processing, we study the whole workflow of how a typical relational database management system optimizes and executes SQL queries. By having an in-depth understanding of the query-optimization-and-execution pipeline, one becomes more proficient in administering DBMSs, and hand-optimizing SQL queries for fast execution.

With respect to transaction processing we study how a typical relational database management systems ensures recovery from errors and controls concurrent access to the data.

Contacts

- Lecturer : Stijn Vansummeren (UB4.125, stijn.vansummeren@ulb.ac.be)
- Assistant : François Picalausa (UB4.133, fpicalau@ulb.ac.be)
- Course Web Page : <http://cs.ulb.ac.be/public/teaching/infoh417>

Schedule

Date	Time		Content
Mon 17 Sep	14h-17	Lect. 1	Course introduction Reminder of the operations of the Relational Algebra Translation of SQL into the Relational Algebra
Mon 24 Sep	14h-17	Lect. 2 Ex. 1	Logical Query Optimization Exercises on SQL-to-Relational-Algebra translation
Mon 1 Oct	14h-17	R.A. 1 Ex. 2	Reading assignment : Physical Data Organization Exercises on Logical Query Optimization
Mon 8 Oct	14h-17	Lect. 3	Index Structures (1)
Mon 15 Oct	14h-16	Lect. 4	Index Structures (2) Introduction to DBMS architecture Project assignment
	16h-17	Lab 1	Lab on Physical data organization
Fri 26 Oct	10h-12 14h-16	Lect. 5 Lab 2	Multidimensional index structures Lab on Physical data organization and BTrees
Fri 9 Nov	10h-12	Lect. 6	Physical Operators
Fri 16 Nov	10h-12 14h-15	Lect 7 Ex. 4	Cost-based plan selection Exercises on physical operators
Fri 23 Nov	10h-12 14h-15	Demo 1 Lect. 8 Ex. 5	Demo : cost-based plan selection in practice Coping with system failures Exercises on cost-based plan selection
Fri 30 Nov	10h-12 14h-15	Lect. 9 Ex. 6.	Transaction processing : concurrency control Integrated exercises
Fri 7 Dec	10h-12 14h-15	Lect. 10 Ex. 7	More about transaction processing Exercises on system failures & concurrency control
Fri 14 Dec	10h-12	Demo 2 Ex. 8	Demo : Transaction processing in practice Question & answer session Overview exercises
Sat 22 Dec			Project deadline

Detailed Objectives

Upon successful completion of this course, the student should master the following competences :

1. Translating a given SQL expression into the Relational Algebra
2. Improving a relational algebra expression by, where possible, removing redundant joins in select-project-join subexpressions
3. Improving a relational algebra expression by, where possible, (a) replacing cartesian products by joins ; and (b) pushing selections and projections
4. Describing and being able to implement traditional secondary-memory index structures (BTrees, Hashing)
5. Being able to describe and demonstrate the shortcomings of traditional index structures with respect to multi-dimensional search keys. In addition, explaining the studied multi-dimensional indexes by means of an example
6. Describing the most important implementation algorithms (one-pass, sorting, hashing, index) for each of the relational algebra operators, as well as judging the cost of each operator, and knowing their limitations of applicability
7. Given a logical query plan and given base statistics about the size and distributions of the database relations, constructing a heuristically optimal physical query plan, by estimating the sizes of the intermediate results and correspondingly comparing the possible implementations. When joins can be reordered, choosing the order with the least cost.
8. Solving exercises on logging like the ones in sections 17.2.6, 17.3.5 and 17.4.4
9. Solving exercises on concurrency control like exercises 18.2.4, 18.2.5, 18.8.1, 18.8.2, and 18.9.1
10. Solving exercises on recoverability like exercises 19.1.1, 19.1.2 and 19.1.3
11. Being able to reconstruct the studied proofs

Method of evaluation

The project work contributes 6/20 points to the overall score, and the written exam contributes the remaining 14/20 points. Participation in both the project work and the written exam are mandatory requirements for passing the course.