

Final Exam 2006-12-20

DATABASE TECHNOLOGY - 1MB025, 1MD026, 1DL116, 1DL124

Date Wednesday, Dec 20, 2006
Time 09:00-14:00
Teacher on duty Kjell Orsborn, phone 471 11 54 or 070 425 06 91
Exam aids calculator

Instructions:

- Read through the complete exam and note any unclear directives before you start solving the questions. The following guidelines hold:
 - Write clear and neat answers! Answers that cannot be read can obviously not result in any points and unclear formulations can be misunderstood.
 - Assumptions outside of what is stated in the question must be explained. Any assumptions made should not alter the given question.
 - Write your answer on only one side of the paper and use a new paper for each new question to simplify the correction process and to avoid possible misunderstandings.
- A passing grade requires about 50% of the maximum number of points.

1. **Database terminology:** 2 pts
Explain the following database concepts:
 - (a) entity integrity constraint
 - (b) participation constraint

2. **Database system architecture:** 4 pts
Describe the three-schema architecture for database management systems and explain how it supports different forms of data independence.

3. **Enhanced entity-relationship modeling:** 6 pts
Explain the concepts specialization, generalization and aggregation (sv. specialisering, generalisering och aggregering) within enhanced entity-relationship modeling.

4. **Normalization:** 4 pts
Explain:
 - (a) *partial* functional dependency (sv. *partiellt* funktionellt beroende) and (2pts)
 - (b) *transitive* functional dependency (sv. *transitivt* funktionellt beroende). (2pts)

5. **Physical database design:** 4 pts
Describe the basic principles of *external* hashing and how it can be used to store and retrieve data records in files.

6. **Query optimization:** 4 pts
 - (a) What is selectivity and why is it needed in cost-based query optimization? (1p)
 - (b) What is the worst case complexity of cost-based query optimization and how does one avoid this cost as user? (1p)
 - (c) Why does cost-based query optimization pay off despite its complexity? (1p)
 - (d) Give examples of two operators in an 'execution plan' that are not in the relational algebra. (1p)

7. **Object-oriented and object-relational databases:** 4 pts
 - (a) What three kinds of user-definable database extensibility mechanisms are available in an object-relational database system? (2pts)

- (b) Which of the above extensibility mechanisms are lacking or weak in a simple object-oriented database system (i.e. in an object store)? (2pts)

8. Data warehouses:

4 pts

A university needs to analyze the quality of its education and wants to use data warehousing technology. They want to summarize grades, courses, and students using a data cube containing grades of students per course, year, and department.

- (a) Define a star schema to represent the data cube in a relational database. (2pts).
- (b) Define a data cube query to summarize the average grades per department over years 2002-2006. (2pts)

Good Luck and Merry, Merry Christmas!

/ Kjell och Tore