T–76.1143 Database Management, Fall 2012

Exam December 19th, 2013 **No calculators or extra material allowed**

Please write your name, your student ID, the name and the code of this course, the date, and the total number of papers you submit on top of each paper you submit.

1. a) (8 p) Construct an E/R diagram for a chain of shops and its credit customers (i.e. customers who can pay their purchases using a special credit card and pay their bills according to certain rules).

   Each customer has a unique ID, a name and an address. A customer can have at most one charge account (i.e. the account in which the purchases and payments are registered). However, the same charge account can have several owners. For example, a couple can have a common charge account, but each of them is registered as a separate customer. The information about charge account includes, in addition to its owners, its unique ID, balance, interest rate, credit limit and minimum monthly payment.

   The chain of shops consists of several shops. Information on each shop includes a unique ID, a name and an address. If the customer visits a certain shop often enough, he / she is registered as a regular customer of this shop. The same customer can be a regular customer of several shops, but he / she has only one charge account, which he / she can use in all shops of this chain.

   Each purchase made using the charge account is registered by a unique ID. However, these IDs are only unique in a certain shop, not in the whole chain. The information on the purchase also includes the customer, the shop, the date and the value of the purchase. The system also registers all payments the customer has made to her / his charge account. The information on a payment includes the date, the sum, the unique ID (index number) and the charge account.

   b) (2 p) Convert the E/R diagram from part (a) into relations. Write the relation schemas and underline the names of the key attributes.

2. Consider the following database schema which presents students, courses, and teachers in a university.

```
Student(studentID, name, department)
Course(code, coursename, credits)
Teacher(teacherID, teachername, department, salary)
Teaches(teacherID, coursecode, term)
Passed(studentID, coursecode, grade, term)
```

   Relation `Teaches` tells which teacher teaches which course in a certain term. Relation `Passed` includes information about the courses the students have passed. For the case of simplicity, we assume that each student can take a certain course only once and relation `Passed` does not contain any failed grades. The credits and salaries are decimal numbers and the grades are integers. The values of all other attributes are strings. The terms are expressed like `'Fall2012'` or `'Spring2011'`. You may assume that the tuples of the relations do not contain NULL values.

   Write the following SQL queries:

   a) (2 p) The student IDs and names of those students who have passed the course CSE-A1200 (course code) in Spring 2013.

   b) (2 p) The names and course codes of those courses which have had at least two different teachers in Spring 2013.

   c) (2 p) The student IDs and names of those students who have passed the course Mathematics L1 (course name), but not the course Mathematics L2.

   d) (2 p) The students whose total number of credits (from the courses their have passed) in Spring 2013 is less than 15. For each such student, the query should produce the name, the student ID and the sum of the credits of this student in Spring 2013.

Turn the paper!

Write the following queries as expressions of the relational algebra:

e) (2 p) All courses passed by the student with the name Teemu Teekkari (or students if there are several with the same name) in Fall 2013. For those courses, list the course code, the name and the number of the credits of the course.

f) (2 p) The teacher IDs and names of those teachers, who have taught exactly one course (not two or more) in Spring 2013.

3. Consider a relation R with schema $R(A, B, C, D, E)$ and functional dependencies $A \rightarrow C$, $AD \rightarrow B$ and $BD \rightarrow E$.

   a) (1 p) Explain why this relation is not in Boyce-Codd normal form (BCNF).

   b) (6 p) Decompose the relation using the BCNF decomposition algorithm taught in this course and in the text book. Give a short justification for each new relation. Continue the decomposition until the final relations are in BCNF. Explain why the final relations are in BCNF.

4. Answer the following questions (a few sentences for each item).

   a) (2 p) What does it mean that the transactions are serializable? Give an example of a problem which can arise, if the transactions are not serializable.

   b) (2 p) What does the atomicity of transactions mean? Give a short example.

   c) (3 p) What does multivalued dependency (denoted by $\twoheadrightarrow$) mean? You do not have to give an exact definition. It is enough to explain the idea. Give a short example. Can a relation have multivalued dependencies if it is in Boyce-Codd normal form? Give a short justification.

5. (4 p) Assume that you have a database consisting of several relations (tables). You are selecting which indexes you should create. How do you judge which indexes are needed? Give a short example to illustrate your answer. (Your example does not have to be complete as long as it shows essential points.)