University of Helsinki, Department of computer science
Introduction to databases, 4.5.2006, H. Laine

Write the name of the course, the date of the exam, your name clearly, your date of birth, and your signature on each separate answer paper.

GIVE THE ANSWERS OF TASKS 1-2 ON ONE ANSWER PAPER AND THE ANSWERS OF TASKS 3-5 ON ANOTHER ANSWER PAPER.

Tasks 1-3 deal with the following recipe database. Column *noOfServings* in table *Course* indicates the number of servings in the recipe. *CookingTime* is expressed in minutes. Categories of courses include *soup*, *salad*, *appetiser*, *dessert* and *main course*. Column *type* in table *Material* contains values like '*fish*', '*pork*', and '*vegetable*'. *Amount* in *Ingredients* is expressed in units specified in table *Material* (for example *kg*, *table spoon*, *apiece*). It contains the amount of material needed for the whole recipe. Table *biggest_courseid* contains the biggest *courseID* value in use. Notation *x -> y* indicates that *x* is a foreign key that refers to table *y*.

```
course (ID, name, easeOfPreparation, noOfServings, cookingTime) [1000 rows]
categories (course->course, category) [3000 rows]
material (materialID, name, type, unit, unitPrice) [200 rows]
ingredients(course->course, materialID->material, amount) [20000 rows]
instruction(course->course, phaseNo, description) [10000 rows]
biggest_courseid(highvalue) [1 row]
```

1.

a) How many rows there are in the result of the natural join of tables *course* and *instruction*?
b) How many rows there are in the result of the projection $\pi_{\textbf{materialID}}$ **(*ingredients*)**?
c) How are the cardinalities (number of rows) of projections $\pi_{\textbf{type}}$ **(*material*)** and $\pi_{\textbf{type,unit}}$ ***(material)*** related to each other ?
d) How many rows there are in the result of the join ***course*** $\bowtie_{\textbf{courseID=course}}$ **ingredients** ?
e) A row is deleted from table *categories*. What should be done to the other tables to preserve the referential integrity of the database?

If you are not able to deduce the exact number of rows, give an estimate and justify it. (10p)

2. Consider the database defined above. Express the following queries in SQL. Define a proper order for the result tables.

a) Make a list of courses that are categorized both as desserts and as soups.
b) List the courses that do not contain onions.
c) Find out in how many courses each material is used as an ingredient. Order the result in degreasing order on the number of courses. (12p)

Turn the paper for tasks 3-5 and give their answer on a separate answer paper.

3. Consider the recipe database. A 'Veggie pea soup' is introduced as a variation of 'Pea soup' (courseID= 230). In this variation the ingredient *pork* (materialID=1330) is substituted with an equal amount of *carrots* (materialID=111). Otherwise the ingredients are the same. Explain what database modification operations are needed to register this variation in the database by making full use of the already existing recipe for the pea soup. In addition to explaining the operations give them also in SQL. (8p)

4. The following schema defines a part of a conference database in the conceptual level. Stereotype <<id>> indicates an identifying attribute or relationship. How should this information be stored as tables in a relational database? Define the schema using the notation used in the beginning of this paper. (9p)

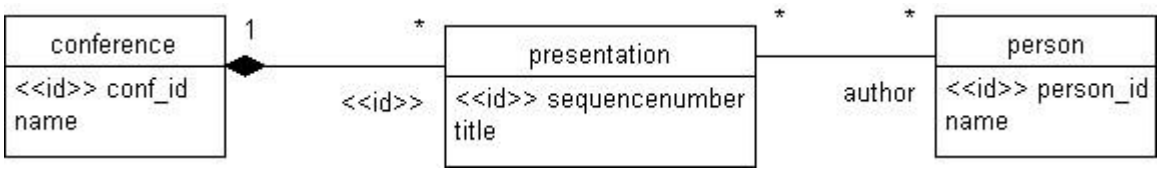

5. Let's consider a relation that contains information on purchase orders

```
purchase_order (personID, person_name, person_address, order_number,
        delivery_address, order_item_number, product_id, amount_ordered,
        dateOrdered).
```

a) Explain the meaning of the functional dependency *personID* à *order_number*

b) Express the rule 'A person can make only one order on the same day' as a functional dependency. (6p)

Turn the paper for tasks 1-2 and give their answer on a separate paper.