Information page for written examinations at Linköping University

Examination date	2011-08-15
Room (1) If the exam is given in different rooms you have to attach an information paper for each room and mark intended place	TER1
Time	14-18
Course code	TDDD37
Exam code	TEN1
Course name **Exam name**	Databasteknik Skriftlig tentamen
Department	IDA
Number of questions in the examination	8
Teacher responsible/contact person during the exam time	Patrick Lambrix, Lena Strömbäck, Jose Pena
Contact number during the exam time	2605
Visit to the examination room approx.	15.15, 16.30
Name and contact details to the course administrator (name + phone nr + mail)	
Equipment permitted	None
Other important information	
Which type of paper should be used, cross-ruled or lined	
Number of exams in the bag	

LiU, Linköping University
IDA, Department of Computer and Information Systems
2011-08-15

EXAM

TDDD37 Database technology

15th August 2011, 14.00-18.00

Room
TER1

Help
None

Pass Grade
You can get max 33 points. To pass the exam, grade 3, you need 7,5 points on the practical part and 9 points on the theoretical part of the exam.

Questions
During the exam, the following teachers are available: Lena Strömbäck 013 282324, 0709 396776, Jose M. Peña, tel. 013 281651, Patrick Lambrix, tel. 013 282605.

Instructions
Write clearly. Give relevant and motivated answers only to the questions asked. State the assumptions you make besides those in the questions. None of these additional assumptions should change the exercises. You can answer in Swedish or English.

Good luck!

Part 1: Practice

Exercise 1. EER modeling (4+1=5 p):

In a school not mentioned here, there are teachers that make exams for several courses. To make the work easier and vary the questions in a good way they want to create a database with exam questions and which exams they have been used in. In the database they want to store:

- For each question: The teacher that created it, at which date it was created, how many credits it give, which exams it has been used in and the text of the question.
- Every question is classified according to its topic. For database courses, for example, there are the following topics: EER-modeling. EER→ relations, SQL, normalization, index, transactions, optimization, recovery and database integration.
- For each exam: the date it was given, responsible teacher, total credits on the exam, and credits needed for pass.
- All database exams contain exactly one question from each of the following topic areas: . EER→ relations, SQL, normalization, index, transactions, and recovery
- DBADV exams also contain one question about optimization
- DBBIO exams also contain one question about database integration

a. Create an EER model for the exam database.
b. Explain how you can guarantee that each exam contains exact the correct exercises, i.e. one from each required topic area. (In your model or in an implementation of it.)

Exercise 2. Translation to relational model (4 p):

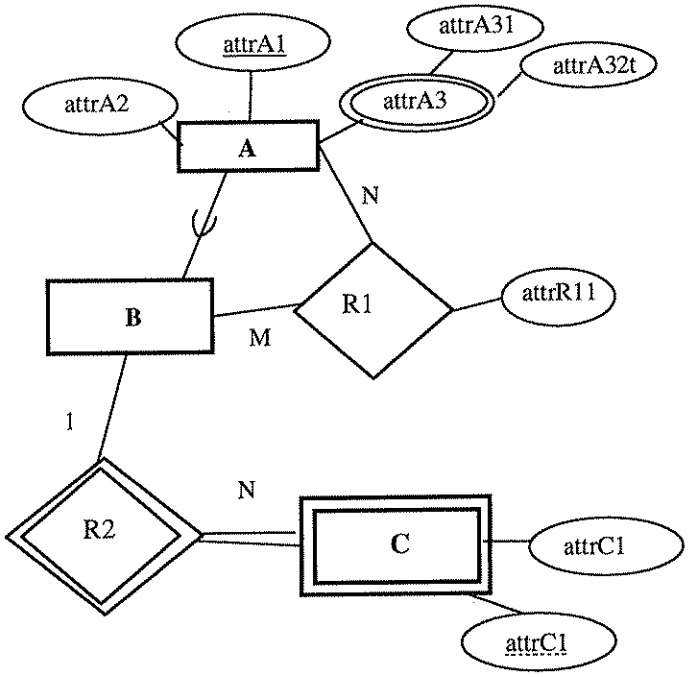

Translate the EER diagram above into a relational model (you have to follow the algorithm seen in the course). Mark the primary keys with solid underlining and the foreign keys with dotted underlining and an arrow from the foreign key to the attribute(s) pointed by the foreign key.

Exercise 3. SQL (1 + 1.5 + 1.5 + 2 = 6 p):

Study the following relations describing how shops order goods:

Supplier:

Id	**Name**
1	Estrella
2	OLW
3	Eldorado
...	...

Customer:

Id	**Name**
1	ICA
2	Coop
...	...

Order:

Supplier	**Customer**	**Date**	**Amount (kkr)**
1	1	16/06/2011	100
1	1	08/06/2011	95
3	1	08/06/2011	105
1	2	16/06/2011	100
2	2	08/06/2011	80
...	...	...	...

Supplier is a foreign key referring to Supplier (id)
Customer is a foreign key referring to Custom (id)

Write SQL queries for the following:

a) List the names of suppliers that have at least one order with ICA (1p)
b) List the names of suppliers that do not have any order with ICA. (1.5p)
c) List the names of suppliers that have more than 10 orders with ICA. (1.5p)
d) List pairs of suppliers that do not have orders with the same customer. For the data above we expect the following answer: (2p)

Supplier1	**Supplier2**
3	2
2	3

Part 2: Theory

Exercise 4. Normalization (1+1+2=4 p):

Assume the relation R(A,B,C,D,E,F) with the following functional dependencies:
BC → A, B→ D, D→BE and E → F

a) Give a primary key for the relation. Show how you arrive at your solution. If there is more than one candidate key for the relation, list them all.
b) Which normal form is the relation in? Motivate.
c) Normalize into BCNF. Show every step in the normalization procedure.

Exercise 5. Data structures (1+1+1+1 = 4 p):

We have a file with 30000 records. Each record is 5 bytes long. The records have a key attribute Y. The file is not ordered. The database uses a block size of B=150 bytes and unspanning allocation.

How many blocks do you have to access to find a record with a given value for Y

a) when you do not use any index ?
b) when you use a secondary index ? Each index record is 3 bytes long.
c) when you use a multilevel index ? Each index record is 3 bytes long.

d) Explain why the multilevel index is the fastest.

Recall that $\log_2 2^x = x$.

x	1	2	3	4	5	6	7	8	9	10	11	12
2^x	2	4	8	16	32	64	128	256	512	1024	2048	4096

Exercise 6. Transactions and concurrency control (2 + 1 + 1 = 4 p):

a) Give a transaction schedule that is NOT serializable. Show that the schedule is not serializable.
b) Use the two phase locking protocol on the transactions in your schedule. Why do we use the two phase locking protocol ?
c) Show that there can be deadlocks with the two phase locking protocol. You do not need to use the same transactions as in your schedule above, i.e. you can make up a new example.

Exercise 7. Database recovery (3 p):

Use the three recovery methods we have seen in the course in the system log below. Show all the operations (in the right order) that must be done to recover from the crash.

Part of system log:

Start-transaction T1
Write-item T1, A, 10, 20
Start-transaction T2
Write-item T1, B, 10, 20
Write-item T2, C, 10, 20
Write-item T2, C, 20, 40
Commit T1
Commit T2
Checkpoint
Start-transaction T3
Write-item T3, D, 20, 30
Commit T3
→system crash

Exercise 8. Query optimization (2+1=3 p):

a) Optimize the following MySQL query:

```
SELECT T1.a, T2.c
FROM T1, T2, T3
WHERE T1.pk=T2.pk AND T2.pk=T3.pk AND T2.b=13;
```

b) The heuristic query optimization algorithm seen in the course pushes selections and projections as far down as possible in the query tree. Why ?