Försättsblad till skriftlig tentamen vid Linköpings universitet

(fylls i av ansvarig)

Datum för tentamen	15/10/2009
Sal	TER1
Tid	14-18
Kurskod	TDDD37 + TDDC94
Provkod	TEN1
Kursnamn/benämning	Databasteknik / Database Technology
Institution	*IDA*
Antal uppgifter som ingår i tentamen	9
Antal sidor på tentamen (inkl. försättsbladet)	6
Jour/Kursansvarig	Patrick Lambrix
Telefon under skrivtid	2605
Besöker salen ca kl.	15.15, 16.30
Kursadministratör (namn + tfnnr + mailadress)	Madeleine Dahlqvist tel. 2360 madha@ida.liu.se
Tillåtna hjälpmedel	lexikon, miniräknare / dictionary, calculator
Övrigt (exempel när resultat kan ses på webben, betygsgränser, visning, övriga salar tentan går i m.m.)	
Vilken typ av papper ska användas, rutigt eller linjerat	
Antal exemplar i påsen	

Institutionen för datavetenskap
Linköpings universitet

TENTAMEN
TDDD37/TDDC94 Database Technology
15 october 2009, 14.00-18.00

Jour: Patrick Lambrix (2605)

Grades: The exam consists of 2 parts. For a pass grade you need to obtain 50% of the total points on **each** part. When a pass grade is obtained, the final grade is based on the total result and not on the different parts.

Instructions: In addition to the instructions on the cover page:

- Write clearly.
- Start the answers to a question on a new page.
- If you make assumptions that are not given in a question, then clearly describe these assumptions. (Of course, these assumptions cannot change the exercise.)

Tools: dictionary, calculator.

LYCKA TILL!

Practical part (14 points)

Question 1. Data modeling with EER diagram (5 p):

Read the whole exercise before you start.

The hotel Nicesthotel needs help with their bookings. In particular they need to handle:
Guests: Specified by a name, address and a phone number for each guest.
Rooms: Specified by the room number and the available number of beds.
Bookings: Specified by arrival and departure dates, and a unique booking number. For each booking we need to keep track of which guests that will come and which room they will get.

The hotel also has a number of luxury rooms and guests can book a 'luxury event'. This always includes a luxury room and meals.
When all ordinary rooms have been booked, the hotel may decide to assign an ordinary guest (who did not book the luxury event) to a luxury room. In this case meals are not included.

Your task is to build an EER model that they can use for creating a database. Clearly write down your choices and assumptions in case you find that something in the information above is not clear. (For full credits you need to use inheritance.)

Question 2. SQL (1 + 1 + 2 = 5 p):

Team

id	name	Arena	founded

Player

id	name	position	age

Playing

id	team	player	season	points

team is a foreign key reference to *id* in table Team.
player is a foreign key reference to *id* in table Player.

A player can only play for one team in a season.
points is the total number of points a player scored in the season.

1. List the names of the teams founded before 1980.

2. List the name and age of all players that play for "LiU Club" during the season 09/10.

3. For each player who is younger than the player *Magnus Johansson*, list the number of seasons they played. (Some junior players may have not played in any season yet!)

Question 3. Translation EER to relational schema (3 + 1 = 4 p):

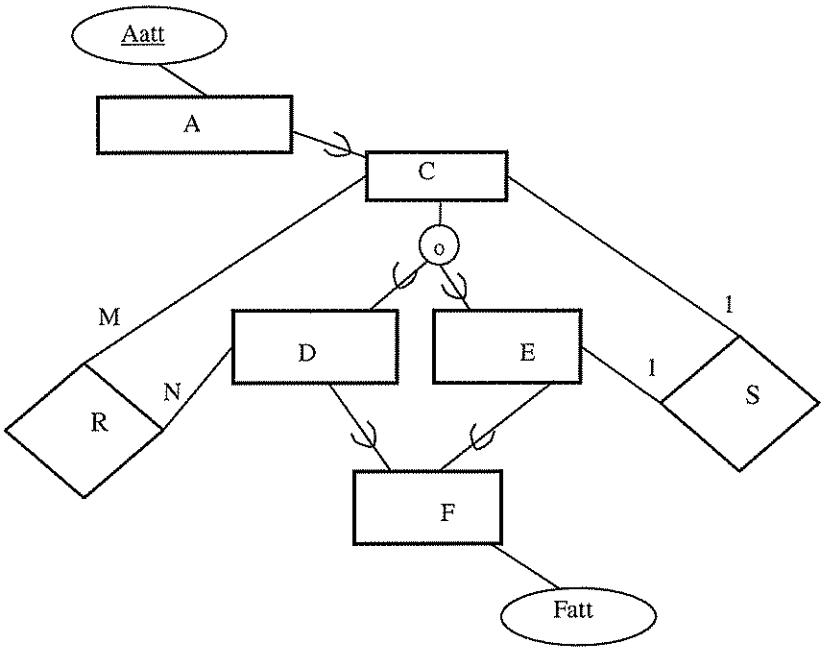

a. Translate the EER diagram to a relational schema. Use as few relations (tables) as possible (but use the algorithm seen in the course).
b. Discuss the advantages and disadvantages of your translation when compared with an alternative translation where more relations are used.

Theoretical part (18 points)

Question 4. Normalization (2 p):

Normalize (1NF→2NF→3NF→BCNF) the relation R(A, B, C, D, E, F, G, H) with functional dependencies F={ABC→DEFGH, D→CEF, EF→GH, G→F, H→G}. *Explain your solution step by step.*

Question 5. Data structures (4 p):

Assume a data file has 1,000,000 records. The file is ordered on the key field ***A***. The size of each record is 100 bytes. The disk block is of size 4096 bytes, and records are stored unspanning. An index entry is of size 10 bytes. A database application always needs to query records randomly based on a key field ***B***.

Illustrate how a single level index can speed up your queries in this application. You need to make a sketch of your solution (index) and calculate the number of block accesses needed for a query *with and without* the index.

Question 6. Concurrency control (2 + 1 = 3 p):

Given the following transactions.

TRANSACTION1
Read-item(Other-account1);
Read-item(My-account);
Other-account1 = Other-account1 – 2000;
My-account = My-account + 2000;
Write-item(My-account);
Write-item(Other-account1);

TRANSACTION2
Read-item(Other-account2);
Read-item(My-account);
Other-account2 = Other-account2 + 3000;
Write-item(Other-account2);
My-account = My-account – 3000;
Write-item(My-account);

a) Give all possible conflicts and draw a serializable (but not serial) schedule. Prove that the schedule is serializable.
b) Can a deadlock occur in a schedule with only the two *given* transactions? Explain your answer.

Question 7. Recovery (3 + 1 = 4 p):

a) Describe the method for database recovery with immediate update. Use the variant that requires both values in the system log for each write-item. Use the system log below to exemplify the method. Show all operations that are performed when the database is recovered. In the right order!

b) Would the use of checkpoints be advantageous for this method? Explain your answer.

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
Commit T1
Start-transaction T3
Start-transaction T4
Write-item T3, D, 20, 30
Write-item T4, E, 50, 40
Write-item T2, C, 20, 40
Commit T2
→system crash

Question 8. Optimization (2 + 1 = 3 p):

a. Let A, B and C be three tables with 10 attributes each. Each of the attributes has the UNIQUE constraint. Optimize the following MySQL query:

```
SELECT A.a, C.c
FROM A, B, C
WHERE A.pk=B.pk AND B.pk=C.pk AND B.b=13;
```

b. Assume that the tables do not contain any NULL value. Assume also that each table contains 10 tuples and that each attribute is of size 1 byte. Show that the optimized query tree is more efficient than the canonial query tree.

Question 9. Stored procedures and triggers (2 p):

Design a trigger such that, when a new department is inserted into jbdept, the trigger calls an already implemented stored procedure AUX(employee_id) with the id of each employee in jbemployee as input parameter.