



# Försättsblad till skriftlig tentamen vid Linköpings universitet

(fylls i av ansvarig)

<b>Datum för tentamen</b>	7 januari 2010
<b>Sal</b>	
<b>Tid</b>	14-18
<b>Kurskod</b>	TDDD37
<b>Provkod</b>	
<b>Kursnamn/benämning</b>	Databasteknik / Database Technology
<b>Institution</b>	<i>IDA</i>
<b>Antal uppgifter som ingår i tentamen</b>	9
<b>Antal sidor på tentamen (inkl. försättsbladet)</b>	7
<b>Jour/Kursansvarig</b>	Patrick Lambrix
<b>Telefon under skrivtid</b>	2605
<b>Besöker salen ca kl.</b>	15, 16.30
<b>Kursadministratör (namn + tfnr + mailadress)</b>	
<b>Tillåtna hjälpmedel</b>	dictionary, calculator
<b>Övrigt (exempel när resultat kan ses på webben, betygsgränser, visning, övriga salar tentan går i m.m.)</b>	
<b>Vilken typ av papper ska användas, rutigt eller linjerat</b>	
<b>Antal exemplar i påsen</b>	

Institutionen för datavetenskap  
Linköpings universitet

# TENTAMEN

## TDDD37 Databasteknik / Database Technology

7 januari 2010, kl 14-18

*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.)
- You can answer in English or Swedish.

*Tools:* dictionary, calculator.

LYCKA TILL!

**Practical part (15 points)**

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

Read the whole exercise before starting.

After Christmas the hard part of Santa's work starts. Now he has to handle all complaints on Christmas presents. He does not like this part of his work and wants to minimize it for coming years. He wants to create a database over complaints and actions. In this way he hopes to avoid complaints in the future by simply avoiding handing out presents that can lead to problems. Santa needs to save the following information:

- For givers and receivers of presents their idnr, name and address. Observe that most persons both give and receive presents but there are persons that only give presents or only receive presents.
- Id and a description of things that have been or can be the content of a gift.
- Information about complaints, who complained, and a text describing the problem.

Santa processes the complaints and solves the problems, but to have more information for next year he wants to save the following information:

- For each solution, a short description about what was the problem and a categorization of the problem to one of the following cases:
  - o The main problem was the gift and information about what was the gift.
  - o The main problem was the giver and information about who was the giver.
  - o The main problem was the receiver and information about who was the receiver.

Help Santa to create an ER/EER diagram for the solution. Note that it is important for Santa to save whether it was the person as giver or receiver that caused the complaint or whether it was the gift. For full points you need a good solution for this in your diagram.

**Question 2. SQL (1 + 2 + 1 + 2 = 6 p):**

Study the following relations describing activities after an alarm:

Rescue Unit:

UnitId	Name	Type	OnMission
1	Ambulance1	Ambulance	1
2	Ambulance2	Ambulance	NULL
3	Police1	Police	1
4	Police 2	Police	NULL

OnMission is a foreign key referring to MissionId.

Mission:

MissionId	Description	MissionStatus
1	Traffic accident	InProgress
2	Robbery	NEW
3	Fire	Solved

Person:

PersonId	Name	InUnit
1	Kalle Karlsson	1
2	Lars Larsson	1
3	Per Persson	3
4	Jonas Jonsson	NULL

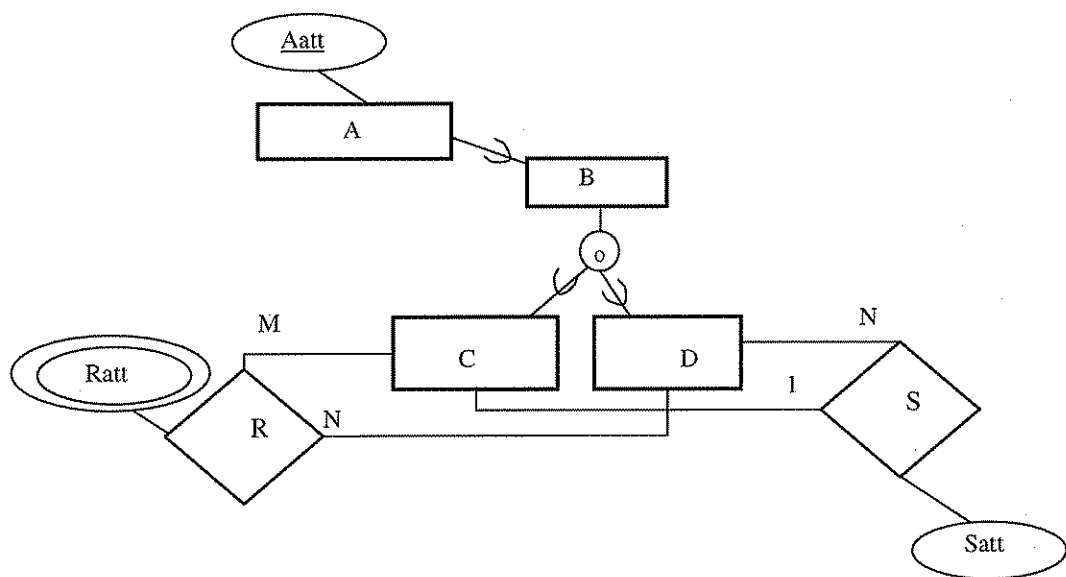
InUnit is a foreign key referring to UnitId

Rescue units can be sent on missions. A unit with *null* in *OnMission* is free for new missions. Missions may have different status, *new* for newly incoming, *InProgress* if a unit has been assigned and *solved* when the mission is finished. For every person there is information about which unit he works in for the moment. Persons with no unit are free.

Write SQL queries for the following:

- List all ongoing missions. (1p)
- List the names of persons participating in each mission grouped by mission (2p)
- The total number of solved missions (1p)
- For every person, list the unit he works in and the mission they are on. Note that free persons and units not on mission must also be in your list. (2p)

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



- Translate the diagram above to relations (you have to follow the algorithm seen in the course).
- Discuss the advantages and disadvantages of your translation when compared with alternative translations of the diagram above.

**Theoretical part (20 points)**

**Question 4. Normalization (1 + 2 = 3 p):**

- a. Prove the additive/union inference rule by using only the reflexive, transitive and augmentation inference rules.
- b. Normalize (1NF→2NF→3NF→BCNF) the relation R(A, B, C, D, E, F, G, H) with functional dependencies {A→BCDEFGH, BC→A, C→H, DE→B, DE→C, E→F, F→G}.  
*Explain your solution step by step.*

**Question 5. Data structures (1 + 1 + 2 = 4 p):**

Assume a table with 500 000 records. Each record is 400 bytes long. The database uses block size  $B = 4\ 096$  bytes and the records are stored unspanning.

- i) How many blocks are needed to store the table?
- ii) Assume we create a primary index where each index record uses 8 bytes (4 bytes for the key and 4 bytes for the disk pointer). How many blocks are needed to store the index?
- ii) How many block accesses are needed to find a record with a given key
  - a. When no index is used.
  - b. When a primary index is used.

**Question 6. Transactions (2 + 2 = 4 p):**

- a. Describe the ACID properties for transactions. Give for each property who/what part of the database management system is responsible for maintaining the property.
- b. Explain AND give example for *lost update problem*, *dirty read problem*, *incorrect summary problem*, *unrepeatable read problem*.

**Question 7. Database recovery (3 + 1 = 4 p):**

- a) Describe the method for recovery with immediate update. Use the variant that works for the system log below. Use the system log below to exemplify the method. Show all operations that are performed during the recovery. In the correct order!
- b) Is the use of checkpoints advantageous in this method? Explain your answer.

Part of system log:

```
Start-transaction T1
Write-item T1, A, 10
Start-transaction T2
Write-item T1, B, 10
Write-item T2, C, 10
Commit T1
Start-transaction T3
Start-transaction T4
Write-item T3, D, 20
Write-item T4, E, 50
Write-item T2, C, 20
Commit T2
→system crash
```

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

- 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 canonical query tree.

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

Design a trigger such that, when a new employee is inserted into jbemployee, the trigger counts in a global variable how many of the already existing employees have a salary greater than the new employee. You are not allowed to use the aggregate function COUNT() to solve this exercise. You should use a cursor, instead.