



Försättsblad till skriftlig tentamen vid Linköpings universitet

(fylls i av ansvarig)

Datum för tentamen	31a. maj 2012
Sal	TER1
Tid	14.00-18.00
Kurskod	TDDD46
Provkod	TEN1
Kursnamn/benämning	Databasteknik
Institution	<i>IDA</i>
Antal uppgifter som ingår i tentamen	7
Antal sidor på tentamen (inkl. försättsbladet)	5
Jour/Kursansvarig	Jose M. Peña
Telefon under skrivtid	013 281651
Besöker salen ca kl.	15.00 och 17.00
Kursadministratör (namn + tfnr + mailadress)	Madeleine Häger Dahlqvist, 013 282360, madeleine.hager.dahlqvist@liu.se
Tillåtna hjälpmedel	Ordbok
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	

EXAM
TDDD12 Databasteknik
TDDD46 Databasteknik

May 31, 2012, 14.00-18.00

Help

Dictionary.

Grades

You can get max 30 points. To pass the exam, grade 3, you need 7.5 points in both the practical and theoretical part of the exam. For grade 4 and 5, you need 21 and 27 points, respectively.

Questions

During the exam, there exists the possibility of asking questions and clarifications from Jose M. Peña, tel. 013 281651 and Fang Wei-Kleiner, tel. 013 284604. They will also visit the room at 15.00 and 17.00.

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 spirit of the exercises. Please, answer in English.

Good luck!

Practical part (15 points)

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

We want to create a database to store information about a new mentorship program. The program states that every child and teenager will have a mentor according to the following rules:

- Every teenager has a single mentor.
- The mentor of a teenager is always an adult.
- Every child has a single mentor.
- The mentor of a child is either a teenager or a child whose mentor is a teenager.

For every child and teenager, we want to store who his/her mentor is. We also want to store who his/her parents are. We assume that his/her parents are always adults, although the parents do not need to be his/her mentor. Your task is to build an EER model that we can use for creating the database. **You should not use more than one regular entity type, otherwise we will withdraw one point (however, you can use as many subclasses as you want).** Clearly write down your choices and assumptions in case you find that something in the information above is not clear.

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

Consider the following database schema

Country(Name, Code, Capital, Area, Population)

Organization(Name, Abbreviation, Established)

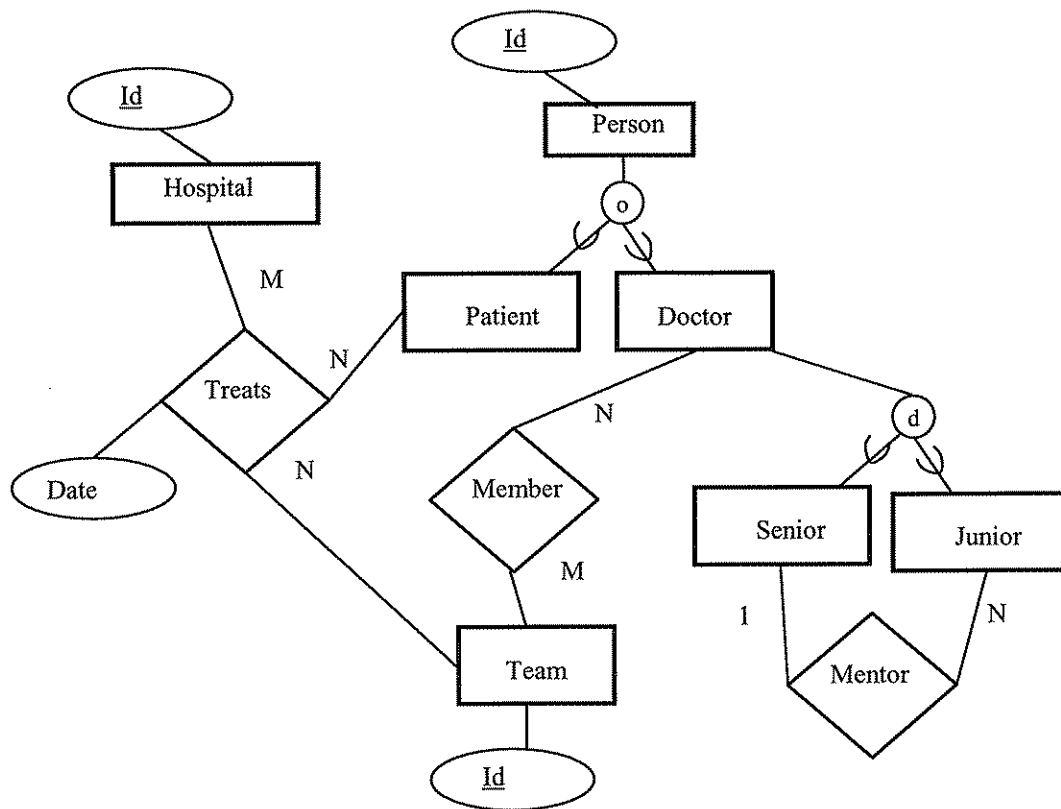
IsMember(Organization, Country, Type)

The attribute Organization in the table IsMember is a foreign key reference to the attribute Abbreviation in the table Organization.

The attribute Country in the table IsMember is a foreign key reference to the attribute Code in the table Country.

1. List all the organization names which Sweden ('SWE') is a member of.
2. For each organization, compute the sum of the population of its member countries. List the organization in descending order of this sum.
3. List the country names which are members of at least one organization, which Sweden is also a member of.

Question 3. Translation of EER diagram into relational schema (5 p):



Translate the EER diagram to a relational schema (use the algorithm seen in the course).

Theoretical part (15 points)

Question 4. Normalization (3 p):

Normalize (1NF→2NF→3NF→BCNF) the relation R(A, B, C, D, E, F, G) with functional dependencies F={ABC→DEFG, A→FG, B→D, DE→BC, F→G}. *Explain your solution step by step.*

Question 5. Data structures (2 + 3 = 5 p):

We have a file with 1000000 records. Each record is 10 bytes long. The records have two key attributes X and Y. The file is ordered on X. The database uses a block size of B=1000 bytes and unspanning allocation. Each index record is 4 bytes long.

1. Calculate the average number of block access needed to find a record with a given value for X when using (1) the primary access method and (2) a single level index.
2. Calculate the average number of block access needed to find a record with a given value for Y when using (1) the primary access method, (2) a single level index and (3) static multi-level index.

Recall that $\log_2 2^x = x$. That is, $\log_2 1 = 0$, $\log_2 2 = 1$, $\log_2 4 = 2$, $\log_2 8 = 3$, $\log_2 16 = 4$, $\log_2 32 = 5$, $\log_2 64 = 6$, $\log_2 128 = 7$, $\log_2 256 = 8$, $\log_2 512 = 9$, $\log_2 1024 = 10$, $\log_2 2048 = 11$, etc.

Question 6. Transactions and concurrency control (1 + 1 + 1 = 3 p):

Consider the following schedule:

T1	T2	T3
write(x)		write(x)
	write(x)	
	write(y)	
write(z)		write(y)

1. The schedule is not serializable. Justify this claim.
2. Is it possible to obtain a serializable schedule by deleting only one operation from the above schedule? How many possibilities are there?
3. If the answer from 2 is yes, consider all the possible new serializable schedules. Which one(s) from them permit(s) the two-phase locking protocol, i.e. can you apply the protocol so that the transactions interleave as in the schedule above? Justify your answer.

Question 7. Database recovery (2 + 1 + 1 = 4 p):

1. Apply the two immediate update recovery methods seen in the course to the system log below. Show all operations that are performed during the recovery. In the correct order!

Part of system log:
Start-transaction T2
Write-item T2, B, 3, 4
Start-transaction T3
Write-item T3, A, 7, 8
Checkpoint
Write-item T3, A, 8, 1
Commit T2
Checkpoint
Write-item T3, A, 1, 5
Start-transaction T4
Write-item T4, B, 4, 5
Write-item T4, B, 5, 10
Commit T3
Checkpoint
Start-transaction T1
Write-item T1, C, 8, 9
Commit T4
→system crash

2. The cache can buffer up to three disk blocks. A transaction modifies on average four disk blocks. Which database recovery method do you recommend to use ?
3. Which database recovery strategy does not need that the checkpoints are stored in the system log ?