Försättsblad till skriftlig tentamen vid Linköpings universitet

	(fylls i av ansvarig)
Datum för tentamen	2012-01-09
Sal	
Tid	kl 14-18
Kurskod	TDDD12
Provkod	TEN1
Kursnamn/benämning	Databasteknik
Institution	*IDA*
Antal uppgifter som ingår i tentamen	7
Antal sidor på tentamen (inkl. försättsbladet)	5 + försättsbladet
Jour/Kursansvarig	Patrick Lambrix, Jose Pena
Telefon under skrivtid	2605, 1651
Besöker salen ca kl.	15.15, 16.45
Kursadministratör (namn + tfnnr + mailadress)	
Tillåtna hjälpmedel	lexikon
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
TDDD12/TDDD46 Database Technology
January 9, 2012, 14.00-18.00

Jour: Patrick Lambrix (2605), José Pena (1651)

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

LYCKA TILL!

Practical part (14 points)

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

Read the whole exercise before you start.

We want to create a database to store information about the Swedish National Health System. Specifically, we want to store information about doctors, hospitals and patients. We also want to store information about which doctor or team of doctors treated which patient in which hospital. Notice that not only single doctors but also teams of doctors can treat patients. Furthermore, we assume doctors can work in several hospitals and that patients can be treated in several hospitals. Notice that doctors can also be patients. We also want to distinguish between senior and junior doctors. Every junior doctor has a senior doctor as mentor. We want to store who is the mentor of whom.

Your task is to build an EER model that they can use for creating the database. 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):

Team

id	name	Arena	founded

Player

id	name	position	age

Playing

id	team	player	year	points

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

points is the total number of points a player scored for a team in a year.
Note that a player can play for more than one team in the same year.

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

2. List the name of each player that has played for more than one team during the year 2011.

3. For each player, show the first year she played and the total number of points she scored in that year.

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

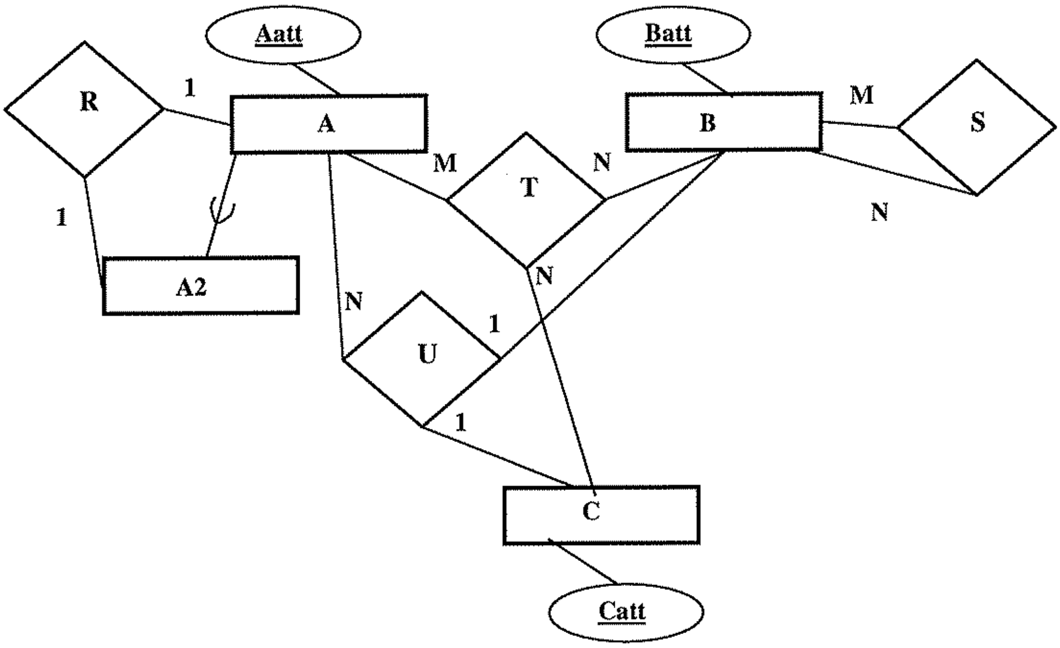

Translate the EER diagram to a relational schema (use the algorithm seen in the course).

Theoretical part (13 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}. *Explain your solution step by step.*

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

We have a file with 30000 records. Each record is 5 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=100 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 the primary access method and when using 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 the primary access method and when using a single level index. T
3. Explain why you obtain different results in the question 2 depending on whether you use an index or not.

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 (2 + 1 = 3 p):

1. Is the following transaction schedule serializable? Motivate your answer.

T1	T2	T3
read(x)		
x:=x+1		
write(x)		
		read(x)
		x:=x+1
		write(x)
	read(x)	
	x:=x+1	
	write(x)	
read(y)		
y:=y+1		
write(y)		
	read(y)	
	y:=y+1	
	write(y)	

2. Does this schedule permit 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 (3 p):

Apply the three 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
Write-item T3, A, 8, 1
Commit T2
Start-transaction T4
Write-item T4, B, 4, 5
Write-item T4, B, 5, 10
Write-item T3, A, 1, 5
Checkpoint
Start-transaction T1
Commit T3
Write-item T1, C, 8, 9
→system crash