CHALMERS UNIVERSITY OF TECHNOLOGY
Department of Computer Science and Engineering

Examination in Databases, TDA357/DIT620

Thursday 16 April 2009, 08:30-12:30

Examiner:	Graham Kemp (telephone 772 5411, room 6475 EDIT) The examiner will visit the exam room at 09:30 and 11:30.
Results:	Will be published by the end of April at the latest.
Exam review:	see course web page for time and place http://www.cs.chalmers.se/Cs/Grundutb/Kurser/dbas/DbasHT2008/index.cgi
Grades:	Grades for Chalmers students (TDA357) are normally determined as follows: ≥ 48 for grade 5; ≥ 36 for grade 4; ≥ 24 for grade 3. Grades for GU students (DIT620) are normally determined as follows: ≥ 42 for grade VG; ≥ 24 for grade G.
Help material:	One A4 sheet with hand-written notes. You may write on both sides of that sheet. That sheet must be handed in with your answers to the exam questions. English language dictionaries are allowed.

Specific instructions:

- Please answer in English where possible. You may clarify your answers in Swedish if you are not confident you have expressed yourself correctly in English.
- Begin the answer to each question on a new page.
- Write clearly; unreadable = wrong!
- Fewer points are given for unnecessarily complicated solutions.
- Indicate clearly if you make any assumptions that are not given in the question.
- Write the page number and question number on every page.

Question 1. Consider the following domain description.

10 p

> The items sold by a chain of furniture stores are identified by the item's model name and colour. Each store in the chain is in a different city. Within a particular store, the place where an item can be found is given by a section name and a shelf number (e.g. section "A", shelf "4"). While most items are sold in a single package, some items are so large that they are sold in two or more packages. Each package has a part number (e.g. if an item is sold in two packages, these packages will be identified by part numbers "1" and "2"). Each package also has a length, width, height and weight. Each purchase made from a store consists of one or more items and may include several instances of the same item (e.g. 4 chairs with the same model name and colour). The owners of the chain of furniture stores want to record which purchases were made from each store. To help with this, each purchase is given a unique purchase reference number. When recording a purchase, the quantity of each item bought must be recorded, and also the purchase date. In the case of purchases made using a credit card, the credit card number and credit card type must also be recorded.

a) Draw an E-R diagram that correctly models this domain. (5p)

b) Translate this E-R diagram into a set of relations, clearly marking all references and keys. (5p)

Question 2. 11 p

a) Suppose we have relation $R(A, B, C, D)$ and functional dependencies $BD \rightarrow A$, $C \rightarrow D$.

i) By considering the closures of all subsets of attributes, find **all** non-trivial FDs, superkeys and keys. State which FDs violate BCNF. (4p)

ii) Decompose relation R to BCNF in two different ways:

— by first decomposing on $BD \rightarrow A$

— by first decomposing on $C \rightarrow D$.

Show each step in the normalisation process, and at each step indicate which functional dependency is being used. (3p)

b) In the following relation, column **teacher** contains the name(s) of the teachers who teach the course whose code is in column **course**. Column **author** contains the names of the authors of the course textbook.

teacher	**author**	**course**
Kemp	Garcia-Molina	TDA357
Kemp	Ullman	TDA357
Kemp	Widom	TDA357
Broberg	Garcia-Molina	TDA357
Broberg	Ullman	TDA357
Broberg	Widom	TDA357
Holmström	Tardos	TIN092
Holmström	Kleinberg	TIN092

State which multi-valued dependencies hold for this relation.

Suppose tuple (Kemp, Tardos, TIN092) is added to this relation.
Which other tuple(s) must be added to the relation?

Decompose this relation to 4th normal form, and show the rows in the resulting relations.

(4p)

Question 3. Suppose relation Employees is as follows:
5 p

name	branch	salary
Andersson	3	20000
Johnsson	3	25000
Larsson	3	32000
Persson	2	28000
Svensson	2	35000

and relation Branches is as follows:

number	city
1	Stockholm
2	Paris
3	London
4	Berlin
5	Rome

a) Write two different relational algebra expressions that evaluate to:

name	city
Larsson	London
Svensson	Paris

i) one of these relational algebra expressions must use the Cartesian product operator.

ii) one of these relational algebra expressions must **not** use the Cartesian product operator.

(3p)

b) Write a relational algebra expression that computes the average salary at each branch, and sorts the results in increasing order of the average salary.

(2p)

Question 4. 8 p

Consider the following relation that contains summary information about the sales made by a company in different countries each month:

$Sales(item, quantity, country, year, month)$

Assume that the month is represented by an integer (1-12).

a) Write an SQL query that finds the total quantity of item "p001" that was sold in the first half (January to June) of 2008. (2p)

b) Write an SQL query that finds those items that had higher sales in Norway than in Denmark in March 2009. (2p)

c) Create a view $V(item, quantity, country, year)$ which contains the total amount of sales on each item in each country in each year. (2p)

d) Write an SQL query that finds those countries where the sales of item "p001" has been higher December than in July in every year for which there are sales records. (2p)

Question 5. 10 p

A database system for managing information about papers submitted to a conference has the following relations:

Papers(refNo, title)
PaperAuthors(refNo, email, name, institute)
Reviewers(refNo, email)
ProgrammeCommittee(email, name, institute)

Each paper submitted to the conference is assigned a unique reference number. Each paper can have several authors, and each author may be an author of more than one paper. Members of the programme committee are each allocated several papers to review, and each paper will be reviewed by several different reviewers. Author names and reviewer names might not be unique, however email addresses are guaranteed to be unique. The value stored in the *institute* columns is the name of the institute where the author or programme committee member works.

a) i) Suggest keys and references for these relations.
Write SQL statements that create these relations with constraints in a DBMS.
(4p)

ii) Motivate the update and delete policies that you choose for foreign key in your answer to part (i).
(2p)

b) A member of the programme committee might be an author of one or more papers that are submitted to the conference. Write an assertion that checks that no reviewer is allocated a paper to review where they are one of the authors.
(2p)

c) Define a view, *Conflicts(refNo, email)*, that lists conflicts of interest that occur when a programme committee member works at the same institute as one of the paper's authors. In this view, *refNo* is a paper reference number and *email* is the email address of a programme committee member.
(2p)

Question 6. 4 p

Suppose we have relation $Accounts(accNo, custNo, balance)$, and that this relation contains the following tuples:

accNo	custNo	balance
A001	C001	500
A002	C001	500

Suppose that customer C001 transfers 100 from account A001 to A002 (transaction T1) and withdraws 200 account A002 (transaction T2). Suppose also that a bank official runs a query to find the total amount in customer C001's accounts (transaction T3). These transactions consist of the following steps:

T1:

$Step\ T1_A:\ get\ balance\ of\ A001\ into\ X$
$Step\ T1_B:\ set\ balance\ of\ A001\ to\ X - 100$
$Step\ T1_C:\ get\ balance\ of\ A002\ into\ Y$
$Step\ T1_D:\ set\ balance\ of\ A002\ to\ Y + 100$

T2:

$Step\ T2_A:\ get\ balance\ of\ A002\ into\ Y$
$Step\ T2_B:\ set\ balance\ of\ A002\ to\ Y - 200$

T3:

$Step\ T3_A:\ get\ balance\ of\ A001\ into\ V$
$Step\ T3_B:\ get\ balance\ of\ A002\ into\ W$
$Step\ T3_C:\ print\ V + W$

a) What are the possible outcomes if all three transactions are run as serializable transactions? Explain your answer. (2p)

b) What are the possible outcomes if T1 and T3 are run as serializable transactions, but T2 is not? Explain your answer. (2p)

Question 7. 6 p

a) Suppose relation R contains the following rows:

a1	d	80
a1	k	20
a1	r	30
a2	m	60
a4	g	90
a5	d	60
a6	m	40
a7	c	80
a8	k	60
a8	s	30

i) Draw a picture that shows a *dense primary index* on the first column of relation R. (2p)

ii) Draw a picture that shows a *secondary index* on one of the other columns of relation R. (2p)

b) Give two reasons why creating an index on a particular column in a particular relation could result in worse overall performance. (2p)

Question 8. Consider the following piece of XML.
6 p

```
<?xml version="1.0" standalone="no">

<Universities>
 <City name="Gothenburg">
  <University name="CTH" type="Technical" />
  <University name="GU"  type="Comprehensive" />
 </City>
 <City name="Stockholm">
  <University name="KTH" type="Technical" />
  <University name="SU"  type="Comprehensive" />
  <University name="KI"  type="Medical" />
 </City>
 <City name="Uppsala">
  <University name="UU"  type="Comprehensive" />
 </City>
</Universities>
```

a) Write a Document Type Definition (DTD) for this XML example. (2p)

b) Write XPath expressions that find:
   - the names of all universities;
   - all University elements for the technical universities. (2p)

c) Write an XQuery expression that gives the city name and university name of the technical universities, i.e.

```
<Result>Gothenburg: CTH</Result>
<Result>Stockholm: KTH</Result>
```

(2p)