CHALMERS UNIVERSITY OF TECHNOLOGY
UNIVERSITY OF GOTHENBURG
Department of Computer Science and Engineering

Examination in Databases, TDA357/DIT620

Thursday 24 August 2017, 14:00-18:00

Solutions

Updated 2017-09-12

Question 1. a) E-R diagram:
12 p

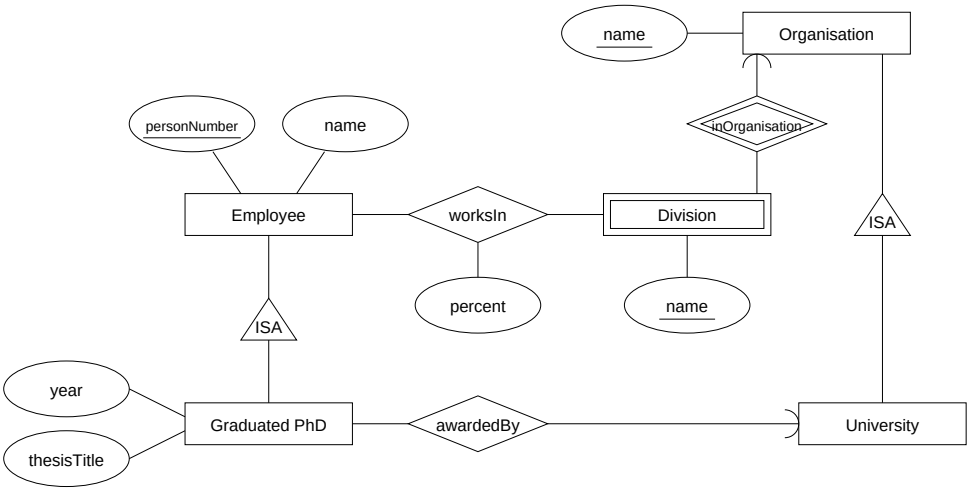

b) $Organisation(\underline{name})$

$University(\underline{name})$

$Division(\underline{name}, \underline{organisation})$
$organisation \rightarrow Organisation.name$

$Employee(\underline{personNumber}, name)$

$graduatedPhD(\underline{personNumber}, year, thesisTitle)$
$personNumber \rightarrow Employee.personNumber$

$worksIn(\underline{employee}, \underline{division}, \underline{organisation}, percent)$
$employee \rightarrow Employee.personNumber$
$(division, organisation) \rightarrow Division.(name, organisation)$

Question 2. Suppose we have relation $R(A, B, C, D, E)$ with functional dependencies $AB \to C$, $C \to B$, $C \to D$, $C \to E$, $D \to E$.
10 p

a) i) All of the given FDs except $AB \to C$ violate BCNF, since their left sides aer not superkeys of R.

ii)
```
Decompose R on C -> D
{C}+ = {B,C,D,E}

        R1(B,_C,D,E)
        R2(_C,_A)
            C -> R1.C

Decompose R1 on D -> E

        R11(_D,E)
        R12(B,_C,D)
            D -> R11.D
```

iii) Problem with dependency preservation — we can no longer enforce $AB \to C$ since no relation contains all three of these attributes.

b) i)
```
R1(A,B,C)
R2(B,C,D)
R3(D,E)
```

ii) This decomposition does not have the same problem as the decomposition in (a)(ii) — 3NF guarantees dependency preservation. But there is some redundancy (e.g. $C \to B$ and we have attributes B and C are together in two relations).

iii) Relation R1 is not in BCNF, since we have FD $C \to B$ and C is not a superkey of R1.

Question 3. 8 p

a)
```
CREATE TABLE Clients (
    cid         INT PRIMARY KEY,
    name        VARCHAR(30),
    telephone   VARCHAR(15)
);

CREATE TABLE Properties (
    ref         VARCHAR(10) PRIMARY KEY,
    address     VARCHAR(30),
    area        INT,
    seller      INT,
    sold        CHAR(1) DEFAULT 'N' CHECK (sold IN ('Y','N')),
    saleDate    DATE,
    FOREIGN KEY (seller) REFERENCES Clients(cid)
);

CREATE TABLE Bids (
    property    VARCHAR(10),
    client      INT,
    amount      INT,
    PRIMARY KEY (property, client, amount),
    FOREIGN KEY (property) REFERENCES Properties(ref),
    FOREIGN KEY (client) REFERENCES Clients(cid)
);
```

b) The triggerering event is an insert into the Bids relation. The trigger should be awakened before this triggering event. The condition should be checked and an exception raised if the constraint would become violated by the insert.

Question 4. 6 p

a) $\tau_{area}(\pi_{name,address}(Clients \bowtie_{cid=seller} (\sigma_{area>300}(Properties))))$

b)
$$\pi_{address}(Properties \bowtie_{ref=property} (\sigma_{numHighBids>3}(\gamma_{property,COUNT(amount)\rightarrow numHighBids}(\sigma_{amount>4000000}(Bids)))))$$

Question 5. 11 p

a)
```
SELECT   name, telephone
FROM     Properties, Clients
WHERE    cid=seller AND sold='N'
```

b)
```
SELECT   property, address, MAX(amount)
FROM     Bids JOIN Properties ON property=ref
GROUP BY property, address
```

c)
```
-- Find those who bid in 2012
-- Find those who bid in 2011
-- Subtract these
```

Question 6. 4 p

Suppose users A and B both run transaction T at the same time, and that A completes steps T_1 and T_2 before B starts step T_1. User B will see a list of available seats that does not include the seat that had been chosen by user A (say, seat "9F"). If A then performs a rollback, seat "9F" will be available. Thus, user B has read a data item (the selection of seat "9F") that was never committed and so never really existed.

Although user B missed the chance to book seat "9F", the consequences of the dirty read are not serious for either the airline or the customers. While dirty reads could be avoided by using other isolation levels, the disadvantage of lower throughput of tranactions could outweigh the benefits of avoiding dirty reads, e.g. if SERIALIZABLE is chosen the resulting performance of the system would be unacceptable.

Question 7. 9 p

a)

```
<!DOCTYPE University [

<!ELEMENT University (Departments, People) >

<!ELEMENT Departments (Department*) >

  <!ELEMENT Department (Programme*) >
    <!ATTLIST Department
      name   ID       #REQUIRED >

    <!ELEMENT Programme EMPTY >
      <!ATTLIST Programme
        code  ID     #REQUIRED
        type  CDATA     #REQUIRED >

  <!ELEMENT People (Person*) >

    <!ELEMENT Person (Teacher?,Student?) >
      <!ATTLIST Person
        name   ID        #REQUIRED
        address CDATA    #IMPLIED >

      <!ELEMENT Teacher EMPTY >
        <!ATTLIST Teacher
          position  CDATA   #REQUIRED
          dept      IDREF   #REQUIRED >

      <!ELEMENT Student (Registered*) >
        <!ATTLIST Student
          prog      IDREF   #REQUIRED >

        <!ELEMENT Registered (#PCDATA) >

]>
```

b) //Person[Student/Registered="DAT620"]

c) (Here is one suggestion; variants might be accepted.)
Solution should have classes Department, Programme, Course and Person. Classes Teacher and Student are subclasses of Person. Many-to-many relationship between students and courses. Many-to-one relationship between teachers and departments. Many-to-one relationship between students and programmes.