Databases Exam

TDA357 (Chalmers), DIT620 (University of Gothenburg)

8 de enero de 2018

Department of Computer Science and Engineering

Grades Chalmers: 24 for 3, 36 for 4, 48 for 5. GU: 24 for G, 42 for VG.

Help material One cheat sheet, which is an A4 sheet with hand-written notes. You may write on both sides of that sheet. If you bring a sheet, it must be handed in with your answers to the exam questions. One English language dictionary is also allowed.

Specific instructions Answer questions in English. Begin the answer to each question (numbers 2 to 5) on a new page. The a,b,c,... parts with the same number can be on the same page. The multiple choice question must be answered in the corresponding boxes.

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. In SQL questions, use standard SQL or PostgreSQL. If you use any other variant (such as Oracle or MySQL), say this; but full points are not guaranteed since this may change the nature of the question.

Multiple Choice Marks right answer = +1, bad answer = -0,5; blank = 0.

1. Multiple Choice Questions (10p, 1p each subquestion)

If a relationship has a cascade updates constraint, then if _____ in the parent table is changed, then the same change will automatically be made to any corresponding foreign key value.:

a) the primary key
b) any alternate key
c) a surrogate key
d) a foreign key

A

Which is the SQL statement that corresponds to $\pi_{title,year}(\sigma_{year=2005}(\text{Articles}))$ (Duplicate elimination should be performed):

a) SELECT UNIQUE title, authorID FROM Articles WHERE year=2005;
b) SELECT DISTINCT title, authorID FROM Articles WHERE year=2005;
c) SELECT title, authorID FROM Articles WHERE year=2005;
d) SELECT title, authorID FROM Articles WHERE (year=2005 AND (title,authorID) is UNIQUE;)

B

In which of the following is a single-entity instance of one type of related to a single-entity instance of another type?

a) 1:1 Relationship
b) 1:N Relationship
c) N:N Relationship
d) Composite Relationship

A

Which of the following columns is(are) are required in a table?:

a) A foreign key
b) A surrogate key.
c) A primary key
d) An alternate key

C

Given this SQL query:

SELECT a , b **FROM** p , q **WHERE** d>e **AND** f=g ;

Which is the relation algebra statement that corresponds to that query?

a) $\pi_{a,b}(\sigma_{(d>e)\wedge(f=g)}(\text{p}\bowtie\text{q}))$
b) $\sigma_{a,b}(\pi_{(d>e)\wedge(f=g)}(\text{p}\times\text{q}))$
c) $\pi_{a,b}(\sigma_{(d>e)\wedge(f=g)}(\text{p}\times\text{q}))$
d) $\gamma_{a,b}(\sigma_{(d>e)\wedge(f=g)}(\text{p}\bowtie\text{q}))$

C

Given this SQL:

```
SELECT customer_T.customer_ID, order_T.customer_ID, name, order_ID
FROM customer_T, order_T;
```

Which is the type of that join?

a) Equi-join

b) Natural join

c) Outer join

d) Cartesian join

D

Considering the following actions:

Transactions...
Commit;
Rollback;
What does Rollback do?

a) Nothing

b) Undoes the transaction before commit

c) Redoes the transactions before commit

d) Clears all transactions

A

Given these definitions:

```
CREATE TABLE User (
 userID INTEGER,
 name CHAR(20),
 password CHAR(30),
PRIMARY KEY (userID)
 );
```

```
CREATE TABLE Follows (
 follower INTEGER,
 followed INTEGER,
PRIMARY KEY (follower,
              followed),
FOREIGN KEY (follower)
        REFERENCES User
          ON DELETE CASCADE,
FOREIGN KEY (followed)
        REFERENCES User
          ON DELETE CASCADE
 );
 );
```

What is "*Follows*" in the E-R model?

a) an entity

b) a relationship

c) a week entity

d) a week relationship

B

Given the following database:

Farm (fcode, description, nid, gcode)
Farmer (nid, name, address, telephoneNr)
Cattle (gcode, description, amount, nid)

The restriction between *Farmer* and *Farm* and between *Farmer* and *Cattle* is **DELETE CASCADE**. Which of the follwing statemets is correct?

a) Restriction between Cattle and Farm must be **DELETE RESTRICT**
b) Restriction between Cattle and Farm must be **DELETE SET NULL**
c) Restriction between Cattle and Farm must be **DELETE CASCADE**
d) None of the above

A

Given the relation R=$\{A, B, C, D, E, F\}$ and the following Functional Dependencies:
FD = $\{BC \to D; D \to C; CE \to A; F \to E; A \to C; AD \to B\}$. Which are the possible keys?

a) $\{FD, FAB\}$
b) $\{FA, FAB, FBC\}$
c) $\{FD, FB, FBC\}$
d) $\{FD, FAB, FBC\}$

D

---

2. Entity Relationship (2 parts, 8p)

3a. Given the following database schema:

Author (a_code, name, nationality)
Book (b_code, title, price, t_code)
Writes (a_code, b_code, date)
Topic (t_code, description)

Draw an Entity-Relationship diagram for the former schema (8p)

SOLUTION:

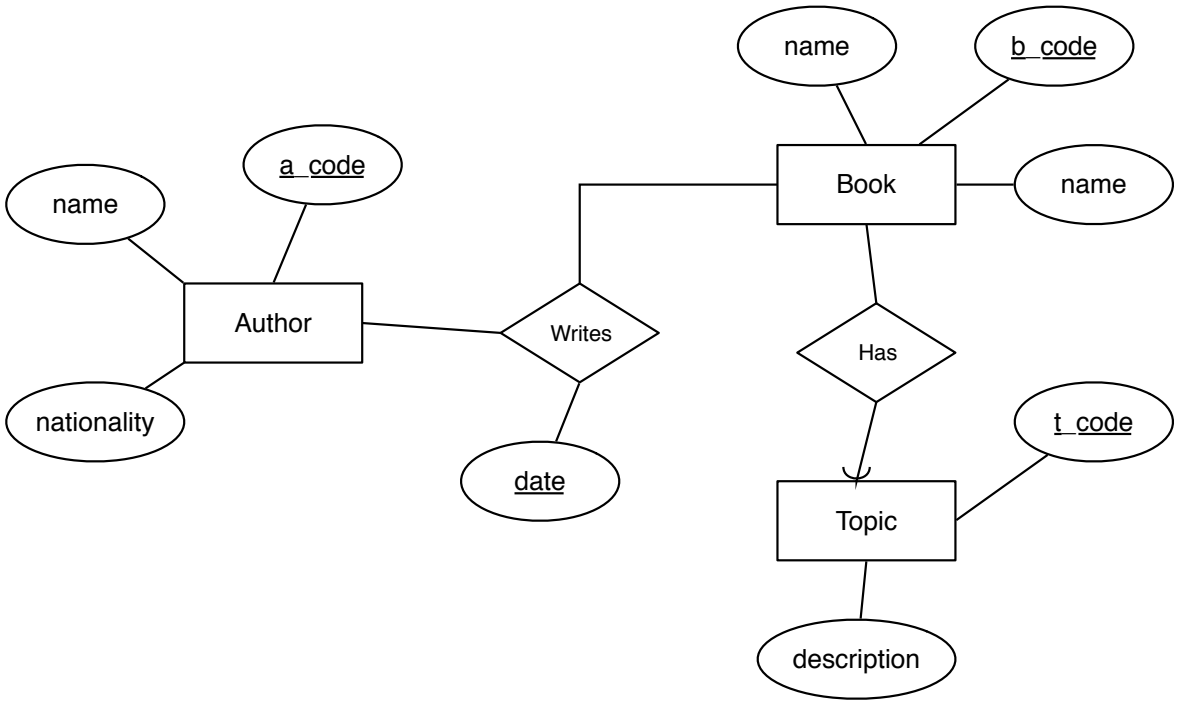

3. Functional Dependencies and Normal Forms (3 parts, 10p)

Suppose we have relation `R(A, B, C, D, E, F)` and functional dependencies FD=$\{A \to F, CD \to A, F \to B, AB \to E, CB \to D\}$.

2a. Extract the keys of the relation `R` and give the Normal Form of `R`. (3p)

SOLUTION:
Keys = {CA,CB,CD,CF}

Since AB is not a super key and E is nor prime in this FD: $AB \to E$:
Normal Form = 2NF

2b. State, with reasons, which FD(s) of relation `R` violate BCNF. (2p)

SOLUTION:
There exist three FD that violates BCNF: $A \to F$, $F \to B$ and $AB \to E$ because their LHS are neither keys nor superkeys.

2c. Decompose relation `R` to BCNF. Show each step in the normalisation process, and at each step indicate which functional dependency is being used. (5p)

SOLUTION:
The first element we choose to normalise is $AB \to E$ because is the only FD which we do not loose other FDs:

- $R_1 =< T_1, L_1 >: T_1 = \{A, B, E\}; L_1 = \{AB \to E\}$ where AB is the key and it is in BCNF
- $R_2 =< T_2, L_2 >: T_2 = \{A, B, C, D, F\}; L_2 = \{A \to F, F \to B, CD \to A, CB \to D\}$. Now the keys are CA, CB, CD, CF and there are 2 more FD in L_2 which violate BCNF: $A \to F$, $F \to B$

Now we choose $A \to F$. We could have chosen the other FDs because, either way we will loose one FD. In this case we will loose $F \to B$ (if we had chosen $F \to B$ then we would loose $A \to B$):

- $R_3 =< T_3, L_3 >: T_3 = \{A, F\}; L_3 = \{A \to F\}$ where A is the key and it is in BCNF
- $R_4 =< T_4, L_4 >: T_4 = \{A, B, C, D\}; L_4 = \{A \to B, CD \to A, CB \to D\}$. Now the keys are CA, CB, CD and there is only more FD in L_4 which violates BCNF: $A \to B$

Finally, we choose $A \to B$ because this is the only FD which violates the BCNF in L_4:

- $R_5 =< T_5, L_5 >: T_5 = \{A, B\}; L_5 = \{A \to F\}$ where A is the key and it is in BCNF
- $R_6 =< T_6, L_6 >: T_6 = \{A, C, D\}; L_6 = \{CD \to A, CA \to D\}$. Now the keys are CA, CD and R_6 and they are in BCNF.

Note now that R_3 and R_5 are identical so we have to join them: $R_{35} =< T_{35}, L_{35} >: T_{35} = \{A, B, F\}; L_{35} = \{A \to F, A \to B\}$. Unorfunatelly, due to this union, the original FD ($F \to B$) arises again because F and B are together again. However, the consequence of having $\{A \to F, A \to B, F \to B\}$ is that $A \to B$ is redundant so the final R_{35} is:

- $R_{35} =< T_{35}, L_{35} >: T_{35} = \{A, B, F\}; L_{35} = \{A \to F, F \to B\}$ where A is the key and it is **not** in BCNF because $F \to B$ is a transitive dependency.

Now we choose the FD that violates the BCNF: $F \to B$:

- $R_{351} =< T_{351}, L_{351} >: T_{351} = \{F, B\}; L_{351} = \{F \to B\}$ where F is the key and it is in BCNF.
- $R_{352} =< T_{352}, L_{352} >: T_{352} = \{A, F\}; L_{352} = \{A \to F\}$ where A is the key and it is in BCNF.

To conclude, the schemas in BCNF are: $R_1 =< T_1, L_1 >$, $R_{351} =< T_{351}, L_{351} >$, $R_{352} =< T_{352}, L_{352} >$ and $R_6 =< T_6, L_6 >$

4. Relational algebra (2 parts, 10p)

Assuming we have the following Relational DataBase:

```
Car(Nrplate, brand, Colour, engine, fuel)
Option(oID, description, price)
Have(Nrplate, oID)
```

4a. Write a relational algebra expression that finds the brand of those cars with all options. (5p)

SOLUTION: $\pi_{brand}(\text{Cars} \times(\text{Have} \div \pi_{oID}\ (\text{Options}))$

4b. Write a relational algebra expression that finds the description of those options that only have the cars whose brand is "*Ovlov*". (5p)

SOLUTION:
$\pi_{description}\ ((\pi_{Nrplate}(\text{Cars}) - \pi_{Nrplate}(\sigma_{brand<>'Ovlov'}(\text{Cars})) \bowtie \text{Have}) \bowtie \text{Options})$

5. SQL, Views, Triggers (2 parts, 10p)

Given the following database schema:

Author (a_code, name, nationality)

Book (b_code, title, price, t_code)

Writes (a_code, b_code, date)

Topic (t_code, description)

5a. Write an SQL query to get the average price of those books writen by Swedish authors. This query must be writen by using nesting queries (HIT: ***IN***) (3p)

SOLUTION:

```
SELECT AVG(price)
FROM Book
WHERE b_code IN (SELECT b_code
        FROM Writes
        WHERE a_code IN (SELECT a_code
                FROM Author
                WHERE nationality='Swedish'));
```

5b. Write an SQL query that lists the authors who have only writen about one and only one topic (5p)

SOLUTION:

```
SELECT name
FROM Author, Writes, Book, Topic
WHERE Author.a_code = Writes.a_code
        AND Writes.b_code = Book.b_code
        AND Book.t_code = Topic.t_code
GROUP BY name
HAVING COUNT (DISTINCT t_code)=1;
```