Birkbeck

(University of London)

BSc/FD EXAMINATION

Department of Computer Science and Information Systems

Database Management (COIY028H6)

CREDIT VALUE: 15 credits

Date of examination: 26 May 2021, 10:00
Duration of paper: 2 hours (in a 3 hour window)

There are **four** questions on this paper.

Answer any **three** of the four questions.

If you answer more than three questions, only the best three answers will count.

Each question carries **20** marks.

The maximum number of marks overall is therefore **60**.

The paper is not prior-disclosed.

This examination is open-book, meaning that you can consult module notes as well as other online material when answering questions.

You are not allowed to discuss the paper with anybody else for the duration of the examination.

For questions that ask you to give textual answers, the examiners will assess the correctness, originality and depth of understanding demonstrated in the answer.

1. (a) Consider the following entity relationship diagram (ERD):

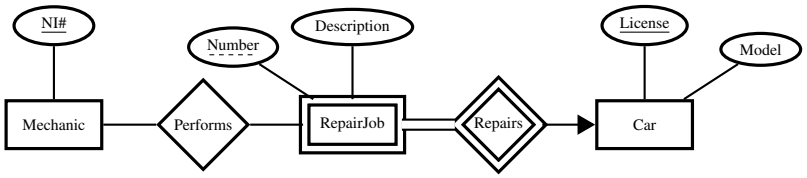

i. Explain what assumptions have been made regarding values for the Number attribute which then make RepairJob a *weak entity type*, as shown.

(2 marks)

ii. When translating from the above ERD to relational tables, what attributes would be placed in the table corresponding to the each of the following, and what would be the primary key in each case?

A. the RepairJob weak entity type

(4 marks)

B. the Performs relationship type

(4 marks)

(b) Assume we want to design a database for a simplified view of airline flights. We have decided that we should include the following seven attributes: FlightNumber, DepAirport, ArrAirport, FlightDate, PassportNumber, SeatNumber and Price.

We have determined that the following functional dependencies (FDs) hold:

FlightNumber $\rightarrow$ DepAirport, ArrAirport
FlightNumber, FlightDate, PassportNumber $\rightarrow$ SeatNo
FlightNumber, SeatNo $\rightarrow$ Price

Our initial relation schema U includes all seven attributes. Schema U is not in Third Normal Form (3NF).

i. Find all the keys for U, demonstrating why they are keys.

(2 marks)

ii. For each FD, explain why it does or does not satisfy the definition of 3NF in U.

(3 marks)

iii. Find a 3NF decomposition of U.

(3 marks)

iv. Is each of your decomposed relation schemas in Boyce-Codd Normal Form (BCNF) or not? Explain your answer.

(2 marks)

2. (a) When defining a column C as a foreign key in a table T, we can specify the condition/action "on delete set null."

i. Explain precisely what this condition/action means (i.e., when it is acted on and what happens as a result), with reference to C and T as well any other columns or tables involved.

(3 marks)

ii. In what situation, with respect to column C, would the action not be able to be carried out and why?

(2 marks)

(b) Consider the following table representing the results so far in a small-scale tennis tournament. (A tennis tournament is a knock-out event, so if somebody loses they don't get to play again.)

Round	Player_1	Player_2	Score
1	A	B	6-4, 6-2
1	C	D	7-6, 3-6, 6-4
1	E	F	6-4, 6-2
1	G	H	6-1, 6-0
2	A	C	7-5, 6-1

Assume that the winner of each match is always placed in the Player_1 column, while the loser is always in the Player_2 column..

i. Identify two functional dependencies with *different lefthand sides* which are *violated* in the above table.

(2 marks)

ii. Identify two functional dependencies with *different lefthand sides* which are *satisfied* in the above table.

(2 marks)

iii. Which of the functional dependencies that you identified in (ii) would you expect to *hold*, in any completed tournament? Justify your answer.

(4 marks)

(c) The following questions relate to transaction management.

i. Explain the difference between the execution of two transactions being *serial* as compared to being *serialisable*.

(3 marks)

ii. Give an example of a situation (which can be described in English) in which two transactions do *not* preserve the desirable property of *isolation*, explaining why isolation has not been preserved.

(4 marks)

3. (a) Consider the following SQL query:

```
select * FROM Part
where city='Paris' or (city='London' and weight > 12)
```

What would the value of the (compound) condition in the `where` clause be for a particular row in `Part` if

i. its `city` value is 'London' and its `weight` value is null?

(1 mark)

ii. its `city` value is null and its `weight` value is 15?

(1 mark)

(b) Assume that we have a table `Sells` with attributes `Store`, `Item` and `Price`, and a view defined as follows:

```
CREATE VIEW euroSells AS
(SELECT Store, Item, Price*1.15 AS Euros FROM Sells);
```

i. Executing the following SQL command on a MySQL database

```
INSERT INTO euroSells VALUES ('OnIT', 'Widget', 115);
```

results in the error "Column 'Euros' is not updatable." Explain why this error message is returned.

(2 marks)

ii. Despite the error message, explain how the system *could* in fact insert the data into the database.

(2 marks)

(c) Explain three ways in which *object-relational* systems can violate the first normal form requirement of "pure" relational systems.

(3 marks)

(d) In PHP (and other languages), one can either pass an SQL query to the database system for immediate execution using the `query` method, or one can use the `prepare` method followed by the `execute` method.

i. In the situation in which user input is being included in the query at runtime, the prepare-and-execute approach is said to be safer because it reduces the possibility of SQL-injection attacks. Explain why this is the case.

(5 marks)

ii. The `prepare` method is said to cause the database system to *compile* the query. Explain the steps involved in this compilation process.

(6 marks)

4. (a) Consider a table `person` with attributes `id` and `name`, along with a table `taken` with attributes `person` (the `id` of a person) and `course` (which represents which people have taken which courses). Now consider the following SQL query:

```
(select id from person)
except
(select person as id from taken)
```

i. Explain why `as id` is required in the above query.

(2 marks)

ii. Describe what the query is asking for, i.e., which people are returned as answers?

(2 marks)

iii. MySQL does not support the `except` operator, so describe *two* alternative ways of producing the same result in SQL, one of which should involve consideration of null values. (You do not need to write down the SQL, but can if you prefer.)

(4 marks)

iv. Assume that a B-tree index has been created on each of the `id` and `person` attributes. Describe what you think would be an efficient method by which the database system could evaluate the query involving `except` (note that we did not explicitly cover such queries when we discussed query processing, but you should still be able to apply methods that we did discuss).

(2 marks)

v. Assume that the `person` table has primary key `id` and has n rows, while the `taken` table has primary key `(person, course)` and has m rows. If the `person` attribute in the `taken` table is a foreign key referencing the `person` table, how many rows would there be in

A. `person join taken on id=person`

B. `person left join taken on id=person`

Explain your answers.

(4 marks)

(b) Assume we have the set of attributes $\{A, B, C, D, E\}$ and the set of functional dependencies (FDs): $\{A \rightarrow B, B \rightarrow A, AC \rightarrow D, D \rightarrow C, E \rightarrow C\}$. Find the closure of each of the sets of attributes AC and BE, showing how you derived your answers. Is either of them a key? Explain your answer.

(6 marks)

