Birkbeck

(University of London)

BSc/FD EXAMINATION

Department of Computer Science and Information Systems

Database Management (COIY028H6)

CREDIT VALUE: 15 credits

Date of examination: 12 June 2020
Duration of paper: (14:00–16:00)

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.

 ©Birkbeck College 2020

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

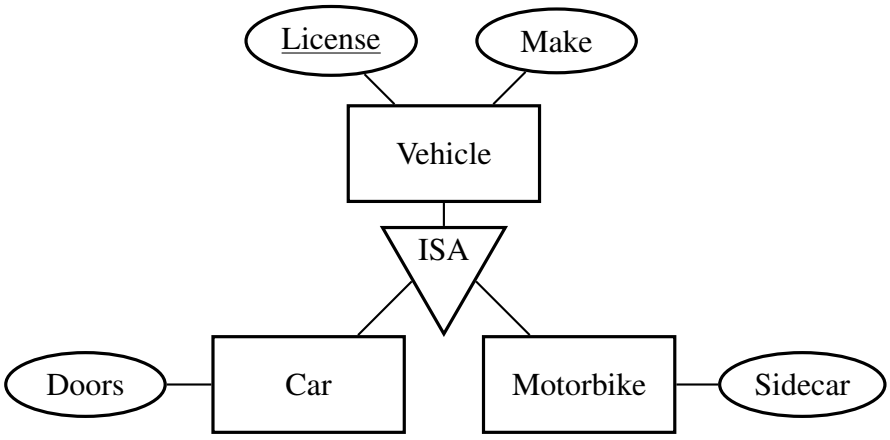

Assume that we want to translate the three entity types in the diagram to tables in a relational system. Explain in what ways the table definitions would be similar and in what ways they would differ when translating to a conventional *relational* system compared to an *object-relational* system.

(8 marks)

(b) Assume that we have a table `Supply` with attributes `Supplier` (a string), `Part` (a string) and `Quantity` (an integer), where the combination of `Supplier` and `Part` is the primary key, and a view defined as follows:

```
CREATE VIEW nonSuppliers AS
SELECT Supplier, Part FROM Supply WHERE Quantity = 0;
```

i. Consider the following SQL command:

```
INSERT INTO nonSuppliers VALUES ('ABC Ltd', 'Axle');
```

What conditions must be true in order for this command to execute without generating an error message? What is the result of the command, assuming that it executes successfully?

(5 marks)

ii. Explain how the `nonSuppliers` view could be used to find suppliers who supply only non-zero quantities of parts.

(4 marks)

(c) In transaction management, a *dirty read* is defined as one transaction, say T_1, reading data written by another, uncommitted transaction, say T_2. Explain why this may or may not cause a problem for transaction T_1.

(3 marks)

 ©Birkbeck College 2020

2. (a) Consider an application which stores information about music files. The attributes are *Track*, *TrackTitle*, *Album*, *AlbumTitle*, *Artist* and *Role*. The functional dependencies (FDs) present are as follows:

$$Track \rightarrow TrackTitle$$
$$Album \rightarrow AlbumTitle$$
$$Track \rightarrow Album$$

i. Determine a key for the full set of attributes, demonstrating why it is a key.

(4 marks)

ii. Produce a database design in Third Normal Form (3NF), explaining the steps you have taken in order to produce the design.

(5 marks)

(b) Assume that we have a web page which uses PHP on a web server to run database queries. The page uses an HTML form with a single text box, with a `name` attribute value of `tablename`, into which the user is expected to enter the name of a table. When the form is submitted, a page containing PHP code is processed by the web server.

i. Now consider the following line of code from the PHP program:

```
$query = "select * from $table";
```

The variable `$table` is expected to contain what the user typed in the HTML text box. Explain how this value is obtained by the PHP code (you can write the necessary PHP if you wish).

(3 marks)

ii. Assume that the PHP code does not check the value entered by the user. Explain why this is a security risk, using an example of what the user might enter in the text box.

(2 marks)

iii. Say that we decided to use an *anonymous placeholder* instead of the `$table` variable in the SQL query above. Write down what the new query would look like. Would doing this solve the problem identified in 2(b)ii above? Explain your answer.

(3 marks)

(c) Consider the SQL condition

```
not(supplier='ABC Ltd') and quantity > 50.
```

appearing in an SQL WHERE clause. Assume that the condition `quantity > 50` is true, and that the attribute `supplier` can take NULL values. Explain the various values to which the above compound condition could evaluate, by considering all the possible values for the condition `supplier='ABC Ltd'`.

(3 marks)

 ©Birkbeck College 2020

3. Consider a table `Player`, with attributes `name`, `team` and `nationality`, and a table `Team`, with attributes `name` and `country`. The `team` attribute in `Player` gives the `name` of the team for which the player plays. We assume that in each table the `name` attribute is the primary key. Consider also the following SQL query:

```
SELECT Player.name, Team.name
FROM Team, Player
WHERE Team.name=Player.team AND
      country='England' AND nationality='British';
```

(a) How would you define the `team` attribute as a foreign key in SQL (you do not need to provide the correct SQL syntax, but should describe what would need to be specified)? Assuming the default constraints associated with foreign keys, how would this affect insertions and deletions of rows in the two tables?

(6 marks)

(b) Describe in English what the above query returns.

(3 marks)

(c) Rewrite the FROM clause in the SQL query so that it uses a `JOIN` expression instead (the WHERE clause would also change, but you do not need to include that in your answer).

(2 marks)

(d) Describe three different ways in which a database system might evaluate the above query, comparing the methods in terms of their relative, expected efficiency.

(9 marks)

 ©Birkbeck College 2020

4. (a) Assume we have the set of attributes $\{A, B, C, D, E\}$ and the set of functional dependencies (FDs): $\{AB \rightarrow C, C \rightarrow A, C \rightarrow D, E \rightarrow D\}$. Find the closure of each of the sets of attributes AB and BC, showing how you derived your answers.

(4 marks)

(b) Consider a relation schema *Bread* which provides information about bread recipes for a home bread-making machine. The attributes of the schema are:

- *type*: the type of flour used (e.g., brown or white),
- *size*: e.g., small, medium or large,
- *programme*: e.g., normal or rapid-bake,
- *flour*: the amount of flour in grams,
- *yeast*: the amount of yeast in teaspoons,
- *water*: the amount of water in millilitres.

Assume that the following four FDs (functional dependencies) hold in *Bread*:

type size $\rightarrow$ *water*
type size programme $\rightarrow$ *yeast*
size $\rightarrow$ *flour*
flour $\rightarrow$ *size*

i. Determine all the keys for *Bread*, demonstrating why they are keys.

(4 marks)

ii. Does the FD *size* $\rightarrow$ *flour* in *Bread* violate (A) Boyce-Codd Normal Form (BCNF), (B) Third Normal Form (3NF)? Explain your answer in each case.

(4 marks)

iii. Identify a redundancy problem in *Bread*.

(2 marks)

iv. Show, step by step, how to produce a BCNF decomposition of *Bread*. Explain why each step is necessary.

(6 marks)

 ©Birkbeck College 2020