BCS THE CHARTERED INSTITUTE FOR IT

BCS HIGHER EDUCATION QUALIFICATIONS
BCS Level 5 Diploma in IT

DATABASE SYSTEMS

Monday 7th October 2024 – Morning

Answer **any** FOUR questions out of SIX. All questions carry equal marks.

Time: TWO hours

Answer any Section A questions you attempt in Answer Book A
Answer any Section B questions you attempt in Answer Book B

The marks given in brackets are **indicative** of the weight given to each part of the question.

Calculators are **NOT** allowed in this examination.

THIS PAGE IS INTENTIONALLY BLANK

Section A
Answer Section A questions in Answer Book A

A1.

a) Define the following concepts used in the context of relational databases and provide an example for **each**:

i. Tuple.

(2 marks)

ii. Attribute.

(2 marks)

iii. n:m relationship.

(3 marks)

b) Database design involves four steps: requirements capture, conceptual modelling, logical modelling and physical modelling. Explain the focus of the following two steps and identify for **each**, **two** key features they involve.

i. Conceptual modelling.

(3 marks)

ii. Physical modelling.

(3 marks)

c) Consider the following requirements in a takeout food delivery company context and create an ERD (Entity Relationship Diagram). Ensure that you show entities, attributes as well their relations and the cardinalities of the relations.

Requirements:

1. A customer has a name and address.
2. Customers can order items from a number of menus (such as lunch menu from provider X or dinner menu from provider Y). Each menu has a name and a name (of the provider).
3. Each menu contains a number of items (a menu with no items cannot exist).
4. Each item has a name, description and price and occurs on at least one menu.
5. Identical menu items can occur on more than one menu (e.g. Chicken Korma might be available from several providers).
6. An order contains one or more items, and a customer could order several of the same item in an order.
7. An order belongs to one customer, but a customer can place multiple orders.

(12 marks)

[Turn Over]

A2.

A railway operator has asked for a new database to be designed for them to keep track of trains. They have been provided with the following sample table capturing their requirements:

Route	Train No and Type	Engine	Scheduled departures	Scheduled Arrivals
London-Paris	103 Express	Class 3	London 10:02 Ashford 13:05	Ashford 13:00 Paris 16:00
London-Glasgow	104 Express	Class 3	London 7:00 Birmingham 9:00 Manchester 10:00	Birmingham 8:45 Manchester 9:45 Glasgow 14:00
	105 Express	Class 3	London 9:00 Birmingham 11:00 Manchester 13:00	Birmingham 10:45 Manchester 12:45 Glasgow 17:00
London-Birmingham	207 Local	Class 5	London 7:00	Birmingham 13:00
London-Southhampton	308 Local	Class 5	London 14:00	Southhampton 19:00

Consider the table and answer the following questions:

a) The table is in unnormalised form.

i. Explain the characteristics of an unnormalised table.

(1 mark)

ii. What problem can occur when searching for data about departures from a particular station? Give an example from the table.

(2 marks)

iii. What problem can occur when a row is deleted from this unnormalised table? Give an example from the table.

(2 marks)

b) Normalise the above table to 1NF.

(3 marks)

c) Identify the partial dependencies in the 1NF table from part b).

(5 marks)

d) Convert the table from part b) to 3NF.

(7 marks)

c) In the context of the designed access control policy from part b), provide examples of how GRANT statements can be used to implement some of the permissions you have identified. There should be **three** GRANT statements, each relating to a different activity (such as select or delete). Decide on suitable role and table names for your statements (you do not need to define any tables or users).

(6 marks)

d) Outline **three** steps that should be taken in the event of a security incident or breach in the database.

(6 marks)

B6.

a) Briefly explain the concept of database recovery.

(2 marks)

b) Identify **four** actions that should be taken in a database recovery procedure addressing media or machine failure, name **each** and give a brief explanation of what they involve.

(8 marks)

c) Describe the concept of 2 Phase Locking and how it provides a solution to the 'lost update' problem.

(4 marks)

d) Define the concept of transaction as used in RDBMS.

(5 marks)

e) Consider the following undo/redo log from a database system that crashed (maybe due to a disk failure) and answer questions below.

Undo/redo log:

```
Timestamp  1 <START T1>
Timestamp  2     <T1, UPDATE(XXX)>
Timestamp  3 <START T2>
Timestamp  4 <COMMIT T1>
Timestamp  5     <T2, UPDATE(YYY)>
Timestamp  6     <T2, UPDATE(ZZZ)>
Timestamp  7 <START  T3>
Timestamp  8     <T3, UPDATE(AAA)>
Timestamp  9 <COMMIT  T2>
Timestamp 10!!!!System crash!!!!
```

Decide for **each** of the transactions (T_1, T_2 and T_3) whether they need to be undone or redone and briefly justify your decisions.

(6 marks)

END OF EXAMINATION

Section B
Answer Section B questions in Answer Book B

B4.

a) Consider the concepts of physical and logical data independence, and answer the following:

i. Briefly explain the concept of logical data independence explaining the key advantage it provides.

(3 marks)

ii. What support for physical data independence is available in a typical RDBMS.

(8 marks)

b) Consider the following scenario and answer the questions below.

An able programmer is keeping data on their social connections on paper record cards but has decided to digitise their record keeping. Their current cards are one per contact, showing some details such as name and contact details but also a list of their other connections (a kind of 'friend of a friend').

i. What would a file-based approach to digitising the data look like and what are its advantages and disadvantages?

(5 marks)

ii. Describe the concept of a graph database.

(4 marks)

iii. What advantage would a graph-based database have for the given example?

(5 marks)

B5.

a) Define the following concepts in the context of database security and provide an example of them:

i. Authentication.

(2 marks)

ii. Privileges.

(2 marks)

iii. Principle of least privilege.

(3 marks)

b) Consider a scenario in an educational institution where the database holds tables with student personal data, their admissions documents, their grades on courses and their fee payment records. The institution has academic staff, admissions officers and finance officers. Define corresponding privileges for people in the above types of jobs with regard to the four data sets identified above.

(6 marks)

e) The company has additional tables as follows for their service engineers:

Engines

Engine ID	Engine Type	Service Interval (miles)	Last Service Mileage	Current Milage	Last service Engineer
1004	Class 3	10000	100000	105000	2
2009	Class 5	10000	1130000	1135600	5
4567	Class 3	10000	260000	263000	2

Engineers

ID	Name
2	Engineer A
5	Engineer B

Define the SQL DDL statement that creates these table (remember to include the required primary and foreign key constraints).

(5 marks)

[Turn Over]

A3.

a) Consider the following tables and answer the questions below:

Contractors Table:

ContractorID	Name	TaxRegistrationNo
1	A	ABC123
2	B	DEF345
3	C	XYZ123
4	D	RST456

Jobs Table:

JobID	Start Date	Contractor ID	Completion date
101	January 2020	3	August 2025
102	September 2021	4	January 2023
103	April 2025		December 2025
104	April 2025	4	December 2027

i. Write an SQL query that finds jobs with no contractor assigned.

(2 marks)

ii. Explain the purpose of this query (written in relational algebra) and convert the relational algebra to SQL.

$\pi_{Name}(\sigma_{JobID\ =\ 1}(Contractors \bowtie Jobs))$

(2 marks)

iii. Explain the purpose of the following query and write the query using relational algebra (assume that we currently have May 2024).

```
SELECT taxregistrationno
FROM contractors
INNER JOIN jobs
ON jobs.contractorid = contractors.contractorid
WHERE completiondate > May 2024 AND contractor.name = 'C';
```

(2 marks)

iv. Consider the following query, given as relational algebra and show the resulting table if it were to be executed.

$\pi_{Name,\ NumberOfJobs}(Contractors \bowtie_{Contractors.ContractorID\ =\ Jobs.ContractorID}$
$(\pi_{ContractorID,\ COUNT(JobID)\ as\ NumberOfJobs}(\sigma_{ContractorID\ =4}(Jobs))))$

(3 marks)

b) Keeping with the contractor example from part a) please answer the following:

i. The SQL DDL definition of contractors (the `CREATE TABLE` statement) contains `TaxRegistrationNo VARCHAR(6) NOT NULL`. Explain what NOT NULL achieves here and why it is desirable.

(2 marks)

ii. The SQL DDL definition of the `jobs` table contains ON UPDATE RESTRICT as part of the foreign key definition for `ContractorsID`. What does this do and why is it desirable?

(3 marks)

c) Consider the following queries written by a novice exploring different joins and show the resulting tables.

i.
```
SELECT * FROM contractors CROSS JOIN jobs WHERE
jobs.contractorID=3;
```

(3 marks)

ii.
```
SELECT * FROM contractors
LEFT OUTER JOIN jobs
ON contractors.contractorsID = jobs.contractordid;
```

(3 marks)

d) Write a query that returns a table showing names and TaxRegistrationNo of all contractors with a total of the number of jobs to which they were or are assigned.

(5 marks)

[Turn Over]