

BCS THE CHARTERED INSTITUTE FOR IT

BCS HIGHER EDUCATION QUALIFICATIONS
BCS Level 6 Professional Graduate Diploma in IT

ADVANCED DATABASE MANAGEMENT SYSTEMS

Monday 2nd October 2023 - Afternoon

Answer **any THREE** questions out of FIVE. All questions carry equal marks.

Time: THREE 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.

Section A
Answer Section A questions in Answer Book A

A1.

- a) Superficially the Object Oriented data model and the Relational data model are similar but in reality they are functionally very different. Discuss whether:
- i. A Class in a UML class diagram is equivalent to an Entity Type in a relational model.
 - ii. A Class in a UML class diagram is equivalent to a Relation/Table in a relational model.
 - iii. An object in a UML class diagram is equivalent to a tuple in a relational model.

(9 marks)

- b) Object Oriented database systems based entirely on Object Oriented principles have not been commercially successful, instead Object Oriented concepts have been integrated into the Relational data model. This requires the mismatch between the Object Oriented approach and the Relational approach to be overcome.

- i. Discuss the principles behind the "Object-Relational Mapping" (ORM) approach used to overcome the mismatch between Object Oriented concepts and programming languages with Relational databases.

(9 marks)

- ii. The SQL-99 ANSI standard provides support for Object Oriented concepts. Explain the following SQL statements. In each case, state which Object Oriented concept is being used.

```
CREATE TYPE Person_Type AS OBJECT
    (person_title VARCHAR(10),
     person_first_name VARCHAR(20),
     person_last_name VARCHAR(20))
NOT FINAL;

CREATE TABLE Person_Table OF Person_Type;
CREATE TYPE Student_Type UNDER Person_Type (
    Enrolment_number VARCHAR(20));

CREATE TABLE Student_Table OF Student_Type(
    Enrolment_number PRIMARY KEY);
```

(4 marks)

- c) Consider the 2-phase commit protocol in distributed databases. Consider a scenario with a Coordinator C and 2 participating RDBMSs, R1 and R2. What is the next action to be sent in the following scenario:

1. C sends prepare Transaction T1 to R1 and R2;
2. R1 sends prepared to C;
3. R2 sends abort to C.

(3 marks)

END OF EXAMINATION

B5.

- a) A Relational Database Management System (RDBMS) is used by many users simultaneously. This can lead to various challenges especially when data is updated. A number of techniques and ideas have been developed to manage these.

With this in mind, answer the questions below.

- i. Expand the acronym ACID (state what each letter stands for). Explain the concept behind “I”.
- ii. Long running processes/transactions, such as those behind business transactions involving the transport of physical items, use compensation actions instead of ACID compliance. Explain why ACID is not suitable for transactions that might run for several days.

(6 marks)

b)

- i. Consider transactions TX1 and TX2. Describe what will happen when TX2 attempts to acquire either a read lock or write lock on a data item when TX1 already has a read or a write lock in place on that data item.

(3 marks)

- ii. Briefly explain the concept of row level locking and give an example in which it is preferable to table level locking.

(4 marks)

- iii. Consider the following schedule for transactions (processed simultaneously by two people, Alice and Bob, who both withdraw money from their shared account). Without any locking, this schedule (sequence of actions) leads to a lost update problem.

The first and last columns show the sequence of Alice’s and Bob’s actions respectively while the middle column shows the current value in the database

Expand the transaction schedule by adding details of when and how locks are taken that prevent the ‘lost update’ problem using;

1. an optimistic locking strategy;
2. a pessimistic locking strategy.

Alice (TX1)	Database	Bob (TX2)
Read balance x	Balance x = 50	
		Read balance x
	Balance x = 20	Update balance x to balance x - 30
Update balance x to balance x - 40	Balance x = 10	

(7 marks)

- iv. Conservative 2 phase locking is a special case of 2 phase locking (2PL) in which all locks are obtained before a transaction begins. What is the advantage of using this technique over 2PL?

(2 marks)

- iii. Study the following two SQL INSERT statements.

```
INSERT INTO Person_Table VALUES('Mr', 'fred', 'Smith');

INSERT INTO Student_Table VALUES('E23453','Mr',
                                   'fred','Smith');
```

Show and explain the output produced by running each of the following SQL SELECT statements after the execution of the above INSERT statements.

```
SELECT * FROM Student_Table;

SELECT * FROM Person_Table;
```

(3 marks)

A2.

- a) Using appropriate examples, describe the defining characteristics of a Data Warehouse and how it differs in content and purpose from an OLTP database. **(8 marks)**
- b) Using suitable examples where appropriate, explain how a Data Warehouse is populated. You should include details of the Extract Transform Load (ETL) process, the key parts of each stage, and the main challenges and common problems associated with the initial population of a Data Warehouse. **(10 marks)**
- c) It has been said that “Metadata turns information into an asset. The more valuable the information asset, the more critical it is to manage the metadata about it”. With reference to the above statement and given the context of a Data Warehouse, explain what is meant by “metadata”.

Give examples of a range of metadata that could be collected during the ETL process.

(7 marks)

[Turn Over]

A3.

Refer to the Employee table below (Fig A2.1) and the Database Trigger (Fig A2.2) below.

Fig A2.1 Employee table (primary key is employeeID)

EMPLOYEE_ID	SURNAME	JOB_ID	SALARY	MANAGER_ID
100	King	PRESIDENT	24000	NULL
124	Mourgos	MANAGER	5800	100
148	Cambrault	MANAGER	11000	100
149	Zlotkey	MANAGER	10500	100
168	Ozer	REP	10500	148
174	Abel	REP	11000	149

Fig A2.2 Database Trigger

```
CREATE TRIGGER salary_difference_trg
AFTER UPDATE OF salary ON employee
FOR EACH ROW
DECLARE salary_diff NUMBER;
BEGIN
    salary_diff := :NEW.salary - :OLD.salary
END;
```

- a) Explain the overall function of the trigger code (Fig A2.2). Show the value of salary_diff when the trigger is invoked by an update that changes the salary of the employee having EmployeeID = 168 to \$10800. **(6 marks)**
- b) A database trigger can be replaced by a stored procedure that has the same functionality. How does a stored procedure differ from a database trigger? **(3 marks)**
- c) Write a SQL SELECT statement that could be used to show details of any employee that has a salary higher than their manager's salary. Your query should output the following result set (Fig A2.3).

Fig A2.3 Result set

EMPLOYEE_ID	SURNAME	JOB_ID	SALARY	MANAGER_ID	MGR_SALARY
174	Abel	REP	11000	149	10500

(6 marks)

- d) Tuning of SQL queries is one of the most important tasks of database management. State the overall objective and discuss the main challenges of Tuning SQL queries. Particular attention should be paid to the use of the following techniques and features available on database systems:
 - Indexes
 - Statistics
 - Execution plans
 - Memory optimisation
 - Caching techniques

(10 marks)

Section B
Answer Section B questions in Answer Book B

B4.

- a) Consider the following tasks related to database and cyber security issues.
 - i. Explain the concept of SQL injection and discuss how it can be prevented.
 - ii. Briefly explain the responsibility of the database administrator with regards to database security.
 - iii. Nullification, shuffling and substitution are methods for data masking. Describe these methods and explain why you would make use of data masking. **(12 marks)**
- b) The concept of data fragmentation is used in distributed databases. It divides large tables into smaller segments that can be physically stored in different locations. Fragmentation is transparent to the user.
 - i. Name **TWO** advantages of table fragmentation. **(2 marks)**
 - ii. Describe vertical fragmentation. **(3 marks)**
 - iii. Consider the following table of registered students (the primary key is underlined). The university has a finance department and academic departments which all access only some of the columns in the table. Show the **TWO** tables that would be produced if vertical fragmentation was to be used.

StudNo	Name	Course	Address	Fees	Marks
<u>2020001</u>	Joe Smith	Computer Science	IT	14000	12,45,69
<u>2020002</u>	Jane Doe	Law	European Law	12000	45,45,60

(5 marks)

- c) Briefly explain the options available for ensuring that data inserted into the database conforms in an expected format. **(3 marks)**

[Turn Over]