



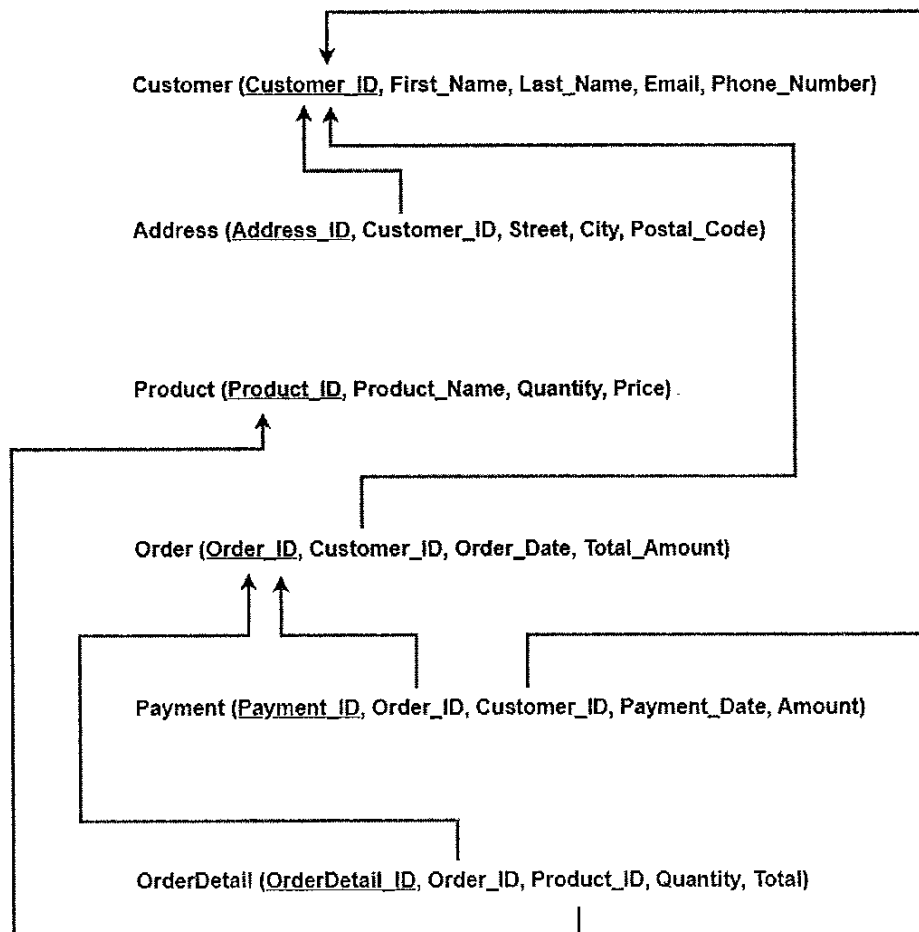
Study Programme	: Bachelor of Software Engineering Honours
Name of the Examination	: Final Examination
Course Code and Title	: EEI5466 Advanced Database Systems
Academic Year	: 2023/2024
Date	: 24 th August 2024
Time	: 13:30 -16:30 hrs
Duration	: 3 Hours

General Instructions

1. This examination paper has five (05) questions in four (04) pages.
2. Answer all five (05) questions.
3. This is a closed book exam, and no reference books and materials are allowed.

Consider the following scenario to answer **Question 01 & 02**

An online shopping platform needs to manage products, customers, orders, and order details. The **relational schema** of this shopping platform is given below.



Question 1 – Database Programming - (20 Marks)

Write the following queries in SQL.

- Write appropriate SQL query to retrieve all customer details who have ordered Product_ID = "P001". (05 Marks)
- Write a stored procedure named "InsertProduct" to insert data into the **Product** table. The stored procedure must use parameters to accept the product details. (06 Marks)
- Write a trigger named "Update_Order_Total" which automatically updates the Total_Amount in the **Order** table whenever a new entry is added to the **OrderDetail** table. The Total_Amount should be the sum of the Total from all **OrderDetail** entries associated with that order. (09 Marks)

Question 2 – Relational Algebra – (20 Marks)

Write the following queries in Relational Algebra.

- List the "Customer_ID" of customers who have made a payment on "24-08-2024". (04 Marks)
- List the "Product_ID" for all products that have a quantity greater than 10 of all products in the **Product** table. (04 Marks)
- List the "Product_ID", "Product_Name", and "Order_ID" of all products that have a price greater than Rs. 1000 and have been ordered by customers. (04 Marks)
- List the "First_Name" and "Last_Name" of customers who have ordered a product with Product_ID = "P005". (04 Marks)
- List the "Order_ID" and "Total_Amount" of all orders placed by the customer with Customer_ID = "C123". (04 Marks)

Question 3 – Object Relational Databases – (20 Marks)

An e-commerce object relational database system has the following requirements:

Products should have a unique identifier, name, price, and categories (which can be a **VARRAY** of strings).

Attribute Name	Product_ID	Product_Name	Price	Categories
Data Type	Number	Varchar (100)	Number	VARRAY (10) OF Varchar2(50)

Customers should have a unique identifier, name, email, and a **nested table** of addresses.

Attribute Name	Customer_ID	Customer_Name	Email	Addresses
Data Type	Number	Varchar2 (100)	Varchar2 (100)	Address_Table

Address should have a street, city, and a postal_code.

Attribute Name	Street	City	Postal_Code
Data Type	Varchar2 (100)	Varchar2 (50)	Varchar2 (10)

Orders should have a unique identifier, a reference to the customer, a list of ordered products (which can be a **nested table**), and the total amount.

Attribute Name	Order_ID	Customer	OrderItems	Total_Amount
Data Type	Number	Customer_Type	OrderItem_Table	Number

OrderItem should have a product REF product object type and quantity.

Attribute Name	Product	Quantity
Data Type	REF Product_Type	Number

- Create object types for **Product**, **Customer**, **Address**, **Order** and **OrderItem** and nested tables for **Address** and **OrderItem**. (11 Marks)
- Create tables for **Customer**, **Products**, and **Order**. (05 Marks)
- Write a query to insert a new product with Product_ID = 1, Product_Name = 'Laptop', Price = 1500, and Categories ('Electronics', 'Computers') into the **Product** table. (02 Marks)
- Write a query to retrieve all orders placed by the customer with **Customer_ID** = 1. (02 Marks)

Question 4 – Transactions & Concurrency Control – (20 Marks)

- Name the four properties of transaction and briefly explain them. (02 Marks)
- Briefly explain what an unrepeatable read is. Give an example of it. (03 Marks)
- State whether the following schedule is serializable or not by comparing the net values of given schedule with the serialized schedules of it.

T1	T2
Read(X)	
Write(X)	
	Read(X)
	Write(X)
	Read(X)
	Write(X)
	Commit T2
Read(Y)	
Write(Y)	
Commit T1	

T1 – Transfer 500 from X to Y
T2 – Calculate 15% interest.

At the initial state of the schedule the balance in account X & Y are 2000 & 1000 respectively. Show all the necessary steps of your answer.

- Briefly explain deadlock prevention & deadlock detection. Explain approaches for each methodology. (03 Marks)

- e) Consider the following sequence of actions, listed in the order of which they are submitted to the DBMS. The Strict 2PL has been used for concurrency control.

T1: R(P), T2: W(Q), T2: W(P), T3: R(R), T3: R(Q), T1: W(R)

T1	T2	T3
S(P)		
R(P)		
	X(Q)	
	W(Q)	
	X(P)	
	W(P)	
		S(R)
		R(R)
		S(Q)
		R(Q)
X(R)		
W(R)		

Assume that older transactions always have higher priority always. The DBMS processes actions in the order shown.

- I. Follow wait-die policy to deal with deadlock in above schedule. Draw schedule again. (03 Marks)
- II. Follow deadlock detection approach to deal with deadlocks in the above schedules. Draw a wait-for graph for the schedule given above to detect the deadlock and briefly explain how to identify the deadlock and remove it. (03 Marks)

Question 5 – Disk, Files, Indexes & Query Processing – (20 Marks)

- a) Explain the terms related to disk access; **Seek time, Rotational delay, Data transfer time.** (03 Marks)
- b) Consider a disk with sector size of 512 bytes, 4000 tracks per surface, 20 sectors per track, and five single-sided platters. Each disk block spans two sectors. The disk platters rotate at 6000 rpm (revolutions per minute).
 - I. What is the capacity of the disk in bytes? (03 Marks)
 - II. What is the average access time for a disk block if the average seek time is 100 msec (0.1 seconds)? Assume $\frac{1}{2}$ revolution for average rotational delay. You can transfer 20 sectors in a single revolution. (04 Marks)

Note: Disk Rotational Time = Seek Time + Rotational Delay + Transfer Time
- c) What are the steps in Query Processing? Explain each step. (04 Marks)
- d) What is the justification for using I/O costs as the main measure in relational query optimizer. (02 Marks)
- e) Briefly explain clustered indexes and un-clustered indexes. (04 Marks)