

ALCCS

FEBRUARY 2009

Code: CS33

Subject: DATABASE MANAGEMENT SYSTEMS

Time: 3 Hours

Max. Marks: 100

NOTE:

- Question 1 is compulsory and carries 28 marks. Answer any FOUR questions from the rest. Marks are indicated against each question.
 - Parts of a question should be answered at the same place.
-

Q.1

(7 x 4)

- Explain the difference between logical and physical data independence.
- Consider the following tables T_1 and T_2 :

T_1			T_2		
P	Q	R	P	S	T
10	a	5	10	b	6
15	b	8	25	c	3
25	a	6	10	b	5

Show the result of the following operations:

(i) $T_1 \bowtie T_2$

(ii) $T_1 \Join T_2$

- What is a data model? List the types of data models.
- What are dangling tuples? What problems may arise because of them?
- In multiple-granularity locking, what is the difference between implicit and explicit locking?
- Define OLAP? What is it used for?
- Among SQL statements DROP TABLE and DELETE, which is DML and which is DDL? Also mention what is the difference in the working of these statements?

Q.2

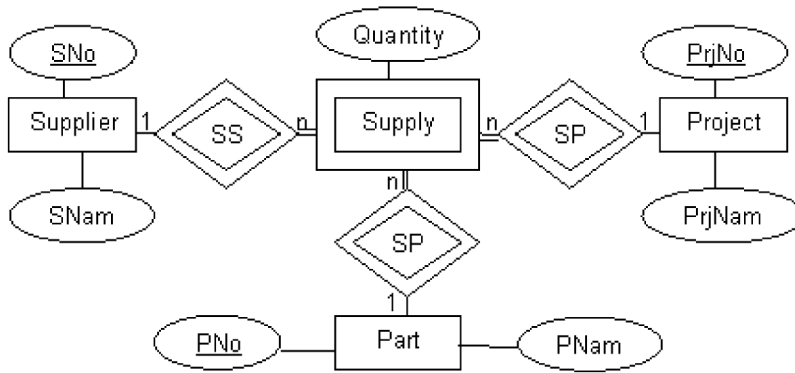
- Create an ER diagram for the following problem definition:

Each company operates four departments, and each department belongs to one company. Each department employs one or more employees, and each employee works for one department. Each of the employees may or may not have one or more dependants, and each dependant belongs to one employee.

(6)

- Design a relational database corresponding to the ER diagram given below. Also specify primary key and foreign key of each relation. For foreign keys, mention the relation in which they are the primary key.

(6)



c. What do you understand by an attribute of an entity? Provide examples of simple, composite, single-valued, multi-valued, and derived attributes. Also draw the symbols used for them in ER diagrams. **(6)**

Q.3

Consider the following relational database

- Hotel (hotelNo, hotelName, city)
- Room (roomNo, hotelNo, type, price)
- Booking (hotelNo, guestNo, dateFrom, dateTo, roomNo)
- Guest (guestNo, guestName, guestAddress)

Where

- Hotel contains hotel details and hotelNo is the primary key;
- Room contains room details for each hotel and (roomNo, hotelNo) forms the primary key;
- Booking contains details of the bookings and (hotelNo, guestNo, dateFrom) forms the primary key;
- Guest contains guest details and guestNo is the primary key.

(a)

produced by the following relational algebra operations:

- (i) $\sigma_{\text{Hotel.hotelNo} = \text{Room.hotelNo}}(\text{Hotel} \times \text{Room})$
- (ii) $\Pi_{\text{hotelName}}(\text{Hotel} \bowtie_{\text{Hotel.hotelNo} = \text{Room.hotelNo}} (\sigma_{\text{price} > 50}(\text{Room})))$
- (iii) $\sigma_{\text{dateTo} \geq '1-Jan-2009'}(\text{Booking})$

What is the output

$\sigma_{\text{Hotel.hotelNo} =$

$\Pi_{\text{hotelName}}(\text{Hotel}$

$\text{Guest} (\sigma_{\text{dateTo} \geq '1-Jan-}$

(3 x 3)

(b)

following:

- (i) List the number of rooms in each hotel in Delhi where there are atleast 10 rooms.
- (ii) List the total income from bookings for the Ambassador Hotel today.
- (iii) List all guests currently staying at the Ambassador Hotel.

Write SQL queries for the

List the number of rooms

List the total income from

List all guests currently

(3 x 3)

Q.4 a. What is the use of GROUP BY clause of SQL? What is the difference between WHERE and HAVING clause in SELECT statement? **(6)**

b. Consider the relation REFRIG(MODEL#, YEAR, PRICE, MANUF_PLANT, COLOR), which is abbreviated as REFRIG(M, Y, P, MP, C), and the following set of F of functional dependencies: $F = \{M \rightarrow MP, \{M, Y\} \rightarrow P, MP \rightarrow C\}$

(i) Find whether the following is a candidate key of relation REFRIG

- a) {M}
- b) {M, Y}

(3)

(ii) Consider the decomposition of REFRIG into $D = \{R1(M,Y,P), R2(M,MP,C)\}$. Is this decomposition lossless? Justify? (3)

(iii) Is the relation REFRIG in 3NF and in BCNF. Justify? (3)

c. Use Armstrong's axioms to prove $\{X \rightarrow Y, XY \rightarrow Z\} \models \{X \rightarrow Z\}$ (3)

Q.5 a. Consider the three transactions T1, T2, and T3, and the schedules S1 and S2 given below.:

T1: r1(x); r1(z); w1(x)

T2: r2(z); r2(y); w2(z); w2(y)

T3: r3(x); r3(y); w3(y)

S1: r1(x); r2(z); r1(x); r3(x); r3(y); w1(x); w3(y); r2(y); w2(z); w2(y)

S2: r1(x); r2(z); r3(x); r1(z); r2(y); r3(y); w1(x); w2(z); w3(y); w2(y)

Draw the serializability (precedence) graphs for S1 and S2 and state whether each schedule is serializable or not. If a schedule is serializable, write down the equivalent serial schedule(s). (7)

b. Briefly explain the deferred update technique of recovery. What are the advantages and disadvantages of this technique? Why is it called the NO-UNDO/REDO method? (7)

c. What do you understand by multiversion concurrency schemes? (4)

Q.6 a. Compute the canonical cover of the following set of functional dependencies for the relational schema $R = (A, B, C, D, E, F)$

$AB \rightarrow C$

$C \rightarrow A$

$BC \rightarrow D$

$ACD \rightarrow B$

$BE \rightarrow C$

$CE \rightarrow A$

$CE \rightarrow F$

$CF \rightarrow B$

$CF \rightarrow D$

$D \rightarrow E$

$D \rightarrow F$

(8)

b. Describe the concept of transitive dependency and explain how this concept is used to define 3NF. (4)

c. Discuss the two phase locking protocol? What are its variations? (6)

Q.7 a. Discuss the key control measures that are used to provide security to data in databases. (6)

b. Discuss the relationship between data warehousing and data mining. (6)

c. Discuss fragmentation transparency, replication transparency and location transparency in distributed databases. (6)