DBMS Gate Questions

Q. 1 Given the relations
employee (name, salary, deptno), and
department (deptno, deptname, address)
Which of the following queries cannot be expressed using the basic relational algebra operations ($\sigma, \pi, \bowtie \cup, \cap -$) ?

(A) Department address of every employee

(B) Employee whose name is the same as their department name

(C) The sum of all employee salaries

(D) All employees of a given department

Q. 2 Given the following relation instance.

X	Y	Z
1	4	2
1	5	3
1	6	3
3	2	2

Which of the following functional dependencies are satisfied by the instance ?

(A) $XY \rightarrow Z$ and $Z \rightarrow Y$ (B) $YZ \rightarrow X$ and $Y \rightarrow Z$

(C) $YZ \rightarrow X$ and $X \rightarrow Z$ (D) $XZ \rightarrow Y$ and $Y \rightarrow X$

Q. 3 Given relations $r(w,x)$ and $s(y,z)$, the result of select distinct w,x from r,s:
is guaranteed to be same as r, provided :

(A) r has no duplicates and s is non empty

(B) r and s have no duplicates

(C) s has no duplicates and r is non empty

(D) r and s have the same number of tuples

Q. 4 Consider a schema R(A,B,C,D) and functional dependencies $A \rightarrow B$ and $C \rightarrow D$. Then the decomposition of R into $R_1(AB)$ and $R_2(CD)$ is :

(A) Dependency preserving and lossless join

(B) Lossless join but not dependency preserving

(C) Dependency preserving but not lossless join

(D) Not dependency preserving and not lossless join

Q. 5 Suppose the adjacency relation of vertices in a graph is represented in a table Adj (X, Y). Which of the following queries cannot be expressed by a relational algebra expression of constant length ?

(A) List all vertices adjacent to a given vertex.

(B) List all vertices which have self loops

(C) List all vertices which belong to cycles of less than three vertices

(D) List all vertices reachable from a given vertex

Q. 6 Let r and s be two relations over the relation schemes R and S respectively, and let A be an attribute in R. Then the relational algebra expression $\sigma_{A=a}(r \bowtie s)$ is always equal to :

(A) $\sigma_{A=a}(r)$ (B) r

(C) $\sigma_{A=a}(r \bowtie s)$ (D) None of the above

Q. 7 R,(A,B,C,D) is a relation. Which of the following does not have a lossless join, dependency preserving BCNF decomposition ?

(A) $A \rightarrow B, B \rightarrow CD$ (B) $A \rightarrow B, B \rightarrow C, C \rightarrow D$

(C) $AB \rightarrow C, C \rightarrow AD$ (D) $A \rightarrow BCD$

Q. 8 Which of the following relational calculus expressions is not safe ?

(A) $\{r \mid \exists u \in R_1(t[A]) = u[A] \wedge \neg\exists s \in R_2(t[A] = s[A])\}$

(B) $\{r \mid \forall u \in R_1(u[A]) = \text{"}x\text{"} \Rightarrow \exists s \in R_2(t[A] = s[A] \wedge s[A] = u[A]))\}$

(C) $\{t \mid \neg(t \in R_1)\}$

(D) $\{t \mid \exists u \in R_1(t[A] = u[A]) \wedge \exists s \in R_2 \vartheta t[A] = s[A])\}$

Q. 9 Consider a relation geq which represents “greater than or equal to”, that is, $(x,y) \in$ geq only if $y \leq x$:

Create table gaq

(*Ib* integer not null

ub integer not null

primary key *Ib*

foreign key (*ub*) references geq on delete cascade):

Which of the following is possible if a tuple (x,y) is deleted ?

(A) A tuple (z,w) with $z > y$ is deleted (B) A tuple (z,w) with $z > x$ is deleted

(C) A tuple (z,w) with $w < x$ is deleted (D) The deletion of (x,y) is prohibited

Q. 10 Relation R with an associated set of functional dependencies, F, is decomposed into $BCNF$. The redundancy (arising out of functional dependencies) in the resulting set of relations is.

(A) Zero

(B) More than zero but less than that of an equivalent $3NF$ decomposition

(C) Proportional to the size of F^+

(D) Indetermine.

Q. 11 With regard to the expressive power of the formal relational query languages, which of the following statements is true ?

(A) Relational algebra is more powerful than relational calculus.

(B) Relational algebra has the same power as relational calculus.

(C) Relational algebra has the same power as safe relational calculus.

(D) None of the above.

Q. 12 AB^+-tree index is to be built on the Name attribute of the relation $STUDENT$. Assume that all student names are of length 8 bytes, disk blocks are of size 512 bytes, and index pointers are of size 4 bytes. Given this scenario, what would be the best choice of the degree (i.e. the number of pointers per node) of the B^+-tree

(A) 16 (B) 42

(C) 43 (D) 44

Q. 13 Relation R is decomposed using a set of functional dependencies, F ,and relation S is decomposed using another set of functional dependencies, G . One decomposition is definitely $BCNF$, the other is definitely. $3NF$, but it is not known which is which. To make a guaranteed identification, which one of the following tests should be used on the decompositions ? (Assume that the closures of F and G are available).

(A) Dependency-preservation (B) Lossless-join

(C) $BCNF$ definition (D) $3NF$ definition

Q. 14 From the following instance of relation schema $R(A, B, C)$, we can conclude that :

A	B	C
1	1	1
1	1	0
2	3	2
2	3	2

(A) A functionally determines B and B functionally determines C

(B) A functionally determines B and B does not functionally determines C .

(C) B does not functionally determines C

(D) A does not functionally B and B does not functionally determines.

Q. 15 Which of the following scenarios may lead to an irrecoverable error in a database system?

(A) A transaction writes a data item after it is read by an uncommitted transaction

(B) A transaction read a data item after it is read by an uncommitted transaction

(C) A transaction read a data item after it is written by an committed transaction

(D) A transaction read a data item after it is written by an uncommitted transaction

Q. 16 Consider the following SQL query

select distinct $a_1, a_2, \ldots\ldots a_n$

from $r_1, r_2, \ldots\ldots r_m$

where P

For an arbitrary predicate P, this query is equivalent to which of the following relational algebra expressions?

(A) $\prod_{a_1, a_2, \ldots a_n} \sigma_\rho (r_1 \times r_2 \times \ldots\ldots r_m)$

(B) $\prod_{a_1, a_2, \ldots a_n} \sigma_\rho (r_1 \triangleright\triangleleft r_2 \triangleright\triangleleft \ldots\ldots \triangleright\triangleleft r_m)$

(C) $\prod_{a_1, a_2, \ldots a_n} \sigma_\rho (r_1 \cup r_2 \cup \ldots\ldots \cup r_m)$

(D) $\prod_{a_1, a_2, \ldots a_n} \sigma_\rho (r_1 \cap r_2 \cup \ldots\ldots \cap r_m)$

Q. 17 Consider the following functional dependencise in a database:

Data_of_Birth→Age

Age→Eligibility

Name→Roll_number

Roll_number→Name

Course_number→Course_name

Course_number→Instructor

(Roll_number,Course_number)→Grade

The relation(Roll)number,Name,Date_of_brith,Age)is

(A) in second normal normal form but not in third normal form

(B) in third normal form but not in BCNF

(C) in BCNF

(D) in none of the above

Q. 18 Consider the set of relations shown below and the SQL query that follow:

Students:(Roll_number,Name,Date_of_birth)

Courses:(Course_number,Course_name,Instructor)

Grades:(Roll_number,Course_number,Grade)

select distrinct Name

from Students, Courses, Grades

Where Students,Roll_number=Grades. Toll_number

and Courses. Instructor=Korth

and Courses. Course_number=Grades. Course_number

and Grades.grade=A

Which of the following sets is computed by the above query?

(A) Names of students who have got an A grade in all courses taught by Korth

(B) Names of students who have got an A grade in all courses

(C) Name of students who have got an A grade in at least one of the courses taught by Korth

(D) None of the above

Q. 19 Consider three data items $D1$,$D2$ and $D3$ and the following execution schedule of transactions $T1$, $T2$ and $T3$. In the diagram, $R(D)$ and $W(D)$ denote the actions reading and writing the data item D respectively.

T1	T2	T3
	R(D3);	
	R(D2);	
	R(D2);	
		R(D2);
		R(D3);
R(D1);		
R(D1);		
		W(D2);
		W(D3);
	R(D1);	
R(D2);		
W(D2);		
	W(D1);	

(A) The schedule is serializable as T2; T3;T1;

(B) The schedule is serializable as T2; T1;T3;

(C) The schedule is serializable as T3; T2; T1;

(D) The schedule is not serializable

Q. 20 Let $R_1(A, B, C)$ and $R_2(D, E)$ be two relation schema, where the primary keys are shown underlined, and let C be a foreign key in R_1 referring to R_2. Suppose there is no violation of the above referential integrity constraint in the corresponding relation instances r_1 and $rl2$. Which one of the following relational algebra expressions would necessarily produce an empty relation?

(A) $\Pi_D(r_1) - \Pi_C(r_1)$ (B) $\Pi_C(r_1) - \Pi_D(r_1)$

(C) $\Pi_D(r_1 \bowtie_{C \neq D} R_2) - \Pi_C(r_1)$ (D) $\Pi_C(r_1 \bowtie_{C=D} R_2)$

Q. 21 Consider the following relation schema pertaining to a students database:
Student(<u>rollno</u>,name,address)
Enroll(<u>rollno</u>,courseno, coursename)
where the primary keys are shown underlined. The number of tuples in the student and Enroll tables are 120 and 8 respectively. What are the maximum and minimum number of tuples that can be present in (Student*Enroll), where '*' denotes natural join?

(A) 8,8 (B) 120,8

(C) 960,8 (D) 960,120

Q. 22 It is desired to design an object-oriented employee record system for a company. Each employee has a name, unique id and salary. Employees belong to different categories and their salary is determined by their category. The functions get Name., getld and compute Salary are required. Given the class hierarchy below, possible locations for these functions are:

(i) getld is implemented in the superclass
(ii) getld is implemented in the suclass
(iii) getName is an abstract function in the superclass
(iv) getName is implemented in the superclass
(v) getName is implemented in the subclass
(vi) getSalary is an abstract function in the superclass
(vii) getSalary is implemented in the superclass
(viii) getSalary is implemented in the subclass

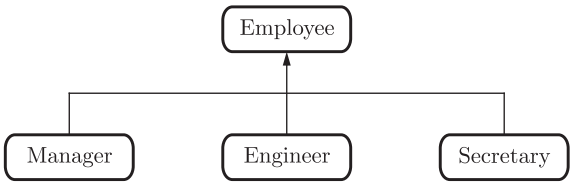

Choose the best design

(A) (i),(iv),(vi).(viii)

(B) (i),(iv),(vii)

(C) (i),(iii),(v),(vi),(viii)

(D) (ii),(v),(viii)

Q. 23 The relation scheme student Performance (name, courselNo, rollNo, grade) has the following functional dependencies:

name, courseNo →grade
RollNo, courseNo →grade
name →rollNo
rollNo →name

The highest normal form of this relation scheme is

(A) 2 NF (B) 3NF

(C) BCNF (D) 4 NF

Q. 24 Consider the relation Student (name, sex, marks), where the primary key is shown underlined, pertaining to students in a class that has at least one boy and one girl. What does the following relational algebra expression produce?

$\Pi_{\text{name}}(r_{\text{sex}=\text{females}}(\text{Student}))\, P_{\text{name}}(\text{Student} \underset{\substack{(\text{sex}=\text{female} \\ \wedge\, x=\text{male} \\ \wedge \text{marks} \le m)}}{\bowtie} r_{n,x,m}(\text{student}))$

(A) names of girl students with the highest marks

(B) names of girl students with more marks than some boy student

(C) names of girl students with marks not less than some boy student

(D) names of girl students with more marks than all the boy students

Q. 25 The order of an internal node in a B* tree index is the maximum number of children it can have. Suppose that a child pointer takes 6 bytes, the search field value takes 14 bytes., and the block size is 512 bytes. What is the order of the internal node?

(A) 24 (B) 25

(C) 26 (D) 27

Q. 26 The employee information in a company is stored in the relation

Employee (name, sex, salary, deptName)

Consider the following SQL query

```
select deptname
    from Employee
    where sex='M'
group by deptName
having avg (salary)>
    (select avg(salary)from Employee)
```

It returns the names of the department in which

(A) the average salary is more than the average salary in the company

(B) the average salary of male employees is more than the average salary of all male employees in the company

(C) the average salary of male employees is more than the average salary of employees in the same department

(D) the average salary of made employees is more than the average salary in the company

Q. 27 Which one of the following is a key factor for preferring B^+-trees to binary search trees for indexing database relation?

(A) Database relations have a large number of record

(B) Database relations are sorted on the primary key

(C) B^+-trees require less memory than binary search trees

(D) Data transfer from disks is in blocks

Q. 28 Which-one of the following statements about normal forms is FALSE?

(A) BCNF is stricter than 3 NF

(B) Loss less, dependency-preserving decomposition into 3 NF is always possible

(C) Loss less, dependency-preserving decomposition into BCNF is always possible

(D) Any relation with two attributes is BCNF

Q. 29 Let r be a relation instance with schema $R = (A, B, C, D)$. WE DEFINE $R_1 = \Pi_{A,B,C}(r)$ and $r_2 = \Pi_{AD}(r)$. let $S = r_1 * r_2$ where * denotes natural join. Given that the decomposition of r into r_1 and r_2 is lossy, which one of the following is TRUE?

(A) $s \subset r$ (B) $r \subset s = r$

(C) $r \subset s$ (D) $r * s = s$

Q. 30 Let E_1 and E_2 be two entities in an E/R diagram with simple single-valued attributes. R_1 and R_2 are two relationships between E_1 and E_2 where R_1 is one-to-many and R_2 is many-to-many. R_1 and R_2 do not have any attributes of their own. What is the minimum number of tables required to represent this situation in the relational model?

(A) 2 (B) 3

(C) 4 (D) 5

Q. 31 The following table has two attributes A and C where A is the primary key and C is the foreign key referencing A with on-delete cascade.

A	C
2	4
3	4
4	3
5	2
7	2
9	5
6	4

The set of all tuples that must be additionally deleted to preserve referential integrity when the tuple (2,4) is deleted is:

(A) (3,4) and (6,4)

(B) (5,2) and (7,2)

(C) (5,2)(7,2) and (9,5)

(D) 1

Q. 32 The relation book (title, price) contains the titles and prices of different books. Assuming that no two books have the same price, what does the following SQL

```
select title
from book as B
where (select count(*)
       from book as T
       where T. price>B.Price)<5
```

(A) Titles of the four most expensive books

(B) Title of the fifth most inexpensive book

(C) Title of the fifth most expensive book

(D) Titles of the five most expensive books

Q. 33 Consider a relation scheme $R = (A, B, C, D, E, H)$ on which the following functional dependencies hold:

$\{A \rightarrow B,.\ BC \rightarrow D,\ E \rightarrow C, D \rightarrow A\}$

What are the candidate keys of R?

(A) AE, BE

(B) AE, BE, DE

(C) AEH, BEH, BCH

(D) AEH, BEH, DEH

Q. 34 Consider the following log sequence of two transactions on a bank account, with initial balance 12000,that transfer 2000 to a mortgage payment and, then apply a 5% interest.

1. T1 start
2. T1 B old = 12000 new = 10000
3. T1 M old = 0 ne = 2000
4. T1 commit
5. T2 start
6. T2 B old = 10000 new = 10500
7. T2 commit

Suppose the database system crashed just before log record 7 is written. When the system is restarted, which one statement is true of the recovery procedure?

(A) We must redo log record 6 to set B to 10500

(B) We must undo log record 6 to set B to 10000 and then redo log records 2 and 3

(C) We need not redo log records 2 and 3 because transaction T1 has committed

(D) We can apply redo and undo operations in arbitrary order because they are idempotent

Q. 35 Consider the relation account (customer, balance) where customer is a primary key and there are no mall values. We would like to rank customers according to decreasing balance. The customer with the largest balance gets rank 1. Ties are not broken but ranks are skipped: if exactly two customers have the largest balance they each get rank 1 and rank 2 is not assigned.

Query 1 : Select A. customer, count (B. customer) from account A, account B where A. customer

Query 2 : Select A. customer, 1+count(B. customer)from account A, accountB where A, balance<B, balance 7 group by A. customer

Consider these statements about Query 1 and Query 2.

1. Query 1will produce the same row set as Query 2 for some but not all databases
2. Both Query 1 Query 2 are correct implementations of the specification
3. Query 1 is a correct implementation of the specification but Query 2 is not
4. Neither query 1 nor Query 2 is a correct implementation of the specification
5. Assigning rank with a pure relational Query takes less time than scanning in decreasing balance order the assigning ranks using ODBC

Which two of the above statements are correct?

(A) 2 and 5 (B) 1 and 3

(C) 1 and 4 (D) 3 and 5

Q. 36 Consider the relation enrolled (student, course) in which student, course) is the primary key, and the relation paid (student, amount) where student is the primary key . Assume no null values and no foreign keys or integrity constraints. Given the following four queries:

Query 1: Select from enrolled where student in (select student form paid)
Query 2: Select student from paid where student in (select student from enrolled)
Query 3: Select E. student from enrolled E, paid P where E. student = P student
Query 4: Se3lect student from paid where exists (select*from enrolled where enrolled student = paid.student
Which one of the following statements is correct?

(A) All queries return identical row sets for any database

(B) Query 2 and Query 4 return identical row sets for all databases but there exist database for which Query 1 and Query 2 retrun different row sets

(C) There exist databases for which Query 3 returns strictly fewer rows than Query 2

(D) There exist databases for which Query 4 will encounter an intergrity violation at runtime

Q. 37 Consider the relation enrolled (student, course) in which (student, course) is the primary key, and the relation paid (student, amount) where student is the primary key. Assume no null values and no foreign keys or integrity constraints. Assume that amounts 6000, 7000,8000,9000 and 10000 were each paid by 20% of the students. Consider these query plans (Plan 1 on left, Plan 2 on right) to "list all courses taken by students who have paid more than x."

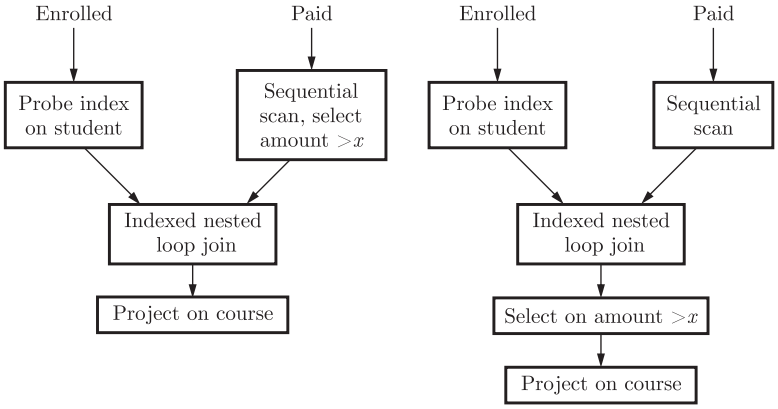

A disk seek takes 4 ms. disk data transfer bank width is 300 MB/s and checking a tuple to see if amount is greater x takes $10\mu s$. Which of the following statements

is correct?

(A) Plan 1 and Plan 2 will not output identical row sets for all databases

(B) A course may be listed more than once int the output of Plan 1 for some databases

(C) For x = 5000, Plan 1 executes faster than Plan 2 for all databases

(D) For x = 9000, Plan 1 executes slower than Plan 2 for all databases

Q. 38 The following functional dependencies are given:
$AB \rightarrow CF, AF \rightarrow D, DE \rightarrow F, C \rightarrow G, F \rightarrow E, G \rightarrow A$.
Which one of the following options is false?

(A) $\{CF\}^+ = \{ACFEFG\}$ (B) $\{BG\}^+ = \{ABCDG\}$

(C) $\{AF\}^+ = \{ACDEFG\}$ (D) $\{AB\}^+ = \{ACDFG\}$

Q. 39 Information about a collection of students is given by the relation studInfo (studId, name, sex). The relation enroll (studID, CourseId) gives which student has enrolled for (or taken) what course(s). Assume that every course is taken by at least one male and at least one female student. What does the following relational algebra expression represent?

$\Pi_{courseId}((\Pi_{studId}(\sigma_{sex="female"}(studInfo)) \times \Pi_{courseId}(enroII)) - enroII)$

(A) Courses in which all the female students are enrolled

(B) Courses in which a proper subset of female students are enrolled

(C) Courses in which only male students are enrolled

(D) None of the above.

Q. 40 Consider the relation employee (name, sex, supervisorName (with name as the key. supervisor Name-gives the name of the supervisor of the employee under consideration. What does the following Tuple Relational Calculus query produce?

{e.name|employee(e)∨
(∀ ×)[¬employee(x) ∨ x.supervisorName
≠ e.name ∨ x.sex = "male"]}

(A) Names of employees with a male supervisor

(B) Names of employees with no immediate male subordinates

(C) Names of employees with no immediate female subordinates

(D) Names of employees with a female supervisor

Q. 41 Consider the table employee (empId, name, department, salary) and the two queries Q_1, Q_2 below. Assuming that department 5 has more than one employee, and we want to find the employees who get higher salary than anyone in the department 5, which one of the statements is TRUE for any arbitrary employee table?

Q_1 : Select e. empId

From employee e

Where not exists

(Select*From employee s Where s. department=”5” and s.salay>=e.salary)

$Q._2$: Select e. empId

From employee e

Where e.salary>Any

(Select distinct salary From employee s Where s. department=”5”)

(A) Q_1 is the correct query.

(B) Q_2 is the correct query

(C) Both Q_1 and Q_2 produce the same answer

(D) Neither Q_1 nor Q_2 is the correct query

Q. 42 Which one of the following statements is FALSE?

(A) Any relation with two attributes is in BCNF

(B) A relation in which every key has only one attribute is in 2NF

(C) A prime attribute can be transitively dependent on a key in 3NF relation

(D) A prime attribute can be transitively dependent on a key in a BNCF relation.

Q. 43 The order of a leaf node in a B^+-tree is the maximum number of (value, data record pointer) pairs it can hold. Given that the block size is 1K bytes, data record pointer is 7 bytes long, the value field is 9 bytes long and a block pointer is 6 bytes long, what is the order of the leaf node?

(A) 63 (B) 64

(C) 67 (D) 68

Q. 44 Consider the following schedules involving two transactions. Which one of the following statements is TRUE?

$S_1 r_1(X); r_1(Y); r_2(X); r_2(Y); w_2(Y); w_1(X)$

$S_2 r_1(X); r_2(X); r_2(Y); w_2(Y); r_1(Y); w_1(X)$

(A) Both S_1 and S_2 are conflict serializable

(B) S_1 is conflict serializable and S_2 is not conflict serializable

(C) S_1 is not conflict serializable and S_2 is conflict serializable

(D) Both S_1 and S_2 are not conflict serializable

Q. 45 A clustering index is defined on the fields which are of type

(A) Non-key and ordering (B) Non-key and non-ordering

(C) key and ordering (D) Key and non-ordering

Q. 46 Let R and S be two relations with the following schema

$R(\underline{P}, \underline{Q}, R1, R2, R3)$

$S(\underline{P}, \underline{Q}, S1, S2)$

Where $\{P, Q\}$ is the key for both schemes. Which of the following queries are equivalent?

I $\Pi_P(\bowtie S)$

II $\Pi_P(R) \bowtie \Pi_P(S)$

III $\Pi_P(\Pi_{P,Q}(R) \cap \Pi_{P,Q}(S))$

IV $\Pi_P(\Pi_{P,Q}(R) - (\Pi_{P,Q}(R) - \Pi_{P,Q}(S)))$

(A) Only I and II (B) Only I and II

(C) Only I, II and III (D) Only I, II and IV

Q. 47 Consider the following relational schemes for a library database:

Book (Title, Author, Catalog_no, Publisher, Year, price)

Collection (Title, Author, Catalog_no)

Which the following functional dependencies:

I. Title Author →Catalog_no

II. Catalog_no →Title Author Publisher Year

III. Publisher Title Year→price

Assume {Author, Title} is the key for both schemes: which of the following statements is true?

(A) Both Book and Collection are in BCNF

(B) Both Book and Collection are in 3NF only

(C) Book is in 2NF and Collection is in 3NF

(D) Both Book and Collection are in 2NF only

Q. 48 Consider a file of 1684 records. Each record is 32 bytes long and its key field is of size 6 bytes. The file is ordered on a non-key field, and the file organization is unspanned. The file is stored in a file system with block size 1024 bytes, and the size of a block pointer is 10 bytes. If the secondary index is built on the key field of the file, and a multi-level index scheme is used to store the secondary index, the number of first-level and second-level blocks in the multi-level index are respectively

(A) 8 and 0 (B) 128 and 6

(C) 256 and 4 (D) 512 and 5

Common Data For Q. 49 & 50 :

Solve the problems and choose the correct answers.
Consider the following ER diagram

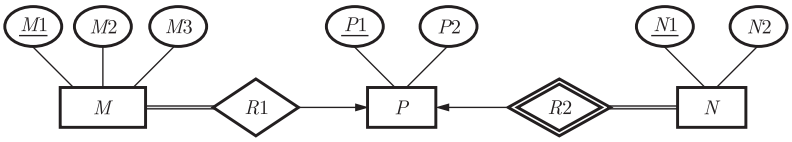

Q. 49 The minimum number of tables needed to represent $M, N, P, R1, R2$ is

(A) 2 (B) 3

(C) 4 (D) 5

Q. 50 Which of the following is a correct attribute set for one of the tables for the correct answer to the above question?

(A) {M1,M2,M3,P1} (B) {M1,P1,N1,N2}

(C) {M1,P1,N1} (D) {M1,P1}

Q. 51 Consider two transactions T_1 and T_2 and four schedules S_1, S_2, S_3, S_4 of T_1 and T_2 as given below :

$T1$: $R1[x]\,W1[x]\,W1[y]$

$T2$: $R2[x]\,R2[y]\,W2[y]$

$S1$: $R1[x]\,R2[x]\,R2[y]\,W1[x]\,W1[y]\,W2[y]$

$S2$: $R1[x]\,R2[x]\,R2[y]\,W1[x]\,W2[y]\,W1[y]$

$S3$: $R1[x]\,W1[x]\,R2[x]\,W1[y]\,R2[y]\,W2[y]$

$S4$: $R2[x]\,R2[y]\,R1[x]\,W1[x]\,W1[y]\,W2[y]$

Which of the above schedules are conflict-serializable?

(A) S1 and S2 (B) S2 and S3

(C) S3 only (D) S4 only

Q. 52 The following key values are inserted into a $B+-$tree in which order of the internal nodes is 3, and that of the leaf nodes is 2, in the sequence given below. The order of internal nodes is the maximum number of tree pointers in each node, and the order of leaf nodes is the maximum number of data items that can be stored in it. The $B+-$tree is initially empty.

10,3,6,8,4,2,1

The maximum number of times leaf nodes would get split up as a result of these insertions is

(A) 2 (B) 3

(C) 4 (D) 5

Q. 53 Let R and S be relation schemes such that $R = \{a, b, c\}$ and $S = \{c\}$. Now consider the following queries on the database :

I. $\pi_{R-S}(r) - \pi_{R-S}(\pi_{R-S}(r) \times S - \pi_{R-S,S}(r))$

II. $\{t \mid t \in \pi_{R-S}(r) \wedge \forall u \in s(\exists v \in r(u = v[s] \wedge t = v[R - S]))\}$

III. $\{t \mid t \in \pi_{R-S}(r) \wedge \forall v \in r(\exists u \in s(u = v[s] \wedge t = v[R - S]))\}$

IV. Select $R.a, R.b$
From R, S
Where $R.c = S.c$

Which of the above queries are equivalent?

(A) I and II (B) I and III

(C) II and IV (C) III and IV

Common Data For Q. 54 & 55 :

Consider the following relational schema:
Suppliers(sid: integer, sname:string, city:string, street:string)
Parts(pid:integer, pname:string, color:string)
Catalog(sid:integer, pid:integer, cost:real)

Q. 54 Consider the following relational query on the above database :

```
SELECT S.sname
FROM Suppliers S
WHERE S.sid NOT IN ( SELECT C.sid
                     FROM  Catalog C
                      WHERE  C.pid NOT IN ( SELECT P.pid
                      FROMParts P
                     WHERE  P.color◇'blue'))
```

Assume that relations corresponding to the above schema are not empty. Which one of the following is the correct interpretation of the above query?

(A) Find the names of all suppliers who have supplied a non-blue part.

(B) Find the names of all suppliers who have not supplied a non-blue part.

(C) Find the names of all suppliers who have supplied only blue parts.

(D) Find the names of all suppliers who have not supplied only blue parts.

Q. 55 Assume that, in the suppliers relation above, each supplier and each street within a city has a unique name, and (same, city) forms a candidate key. No other functional dependencies are implied other than those implied by primary and candidate keys. Which one of the following is TRUE about the above schema ?

(A) The schema is in BCNF. (B) The schema is in 3NF but not in BCNF.

(C) The schema is in 2NF but not in 3NF.

(D) The schema is not in 2NF.

Q. 56 Consider a B^+-tree in which the maximum number of keys in a node is 5. What is the minimum number of keys in any non-root node?

(A) 1 (B) 2

(C) 3 (D) 4

Q. 57 A relational schema for a train reservation database is given below.

Passenger (pid, pname, age)

Reservation (pid, class, tid)

Table : passenger

pid	pname	Age
0	'Sachin'	65
1	'Rahul'	66
2	Sourav'	67
3	'Anil'	69

Table : Reservation

pid	class	tid
0	'AC'	8200
1	'AC'	8201
2	'SC'	8201
3	'AC'	8203
4	'SC'	8204
5	'AC'	8202

What pids are returned by the following SQL query for the above instance of the tables ?

```
SELECT pid
FROM Reservation
WHERE class= 'AC' AND
EXISTS (SELECT*
FROM Passenger
WHERE age>65 AND
Passenger.pid = Reservation.pid)
```

(A) 1,0 (B) 1, 2

(C) 1, 3 (D) 1, 5

Q. 58 Which of the following concurrency control protocols ensure both conflict serializability and freedom from deadlock?

I. 2-phase locking

II Time-stamp ordering

(A) I only (B) II only

(C) Both I and II (D) Neither I nor II

Q. 59 Consider the following schedule for transactions T1, T2, and T3 :

```
T1          T2          T3
Read (X)
            Read (Y)
                        Read (Y)
Write (X)
                        Write (X)
            Read (X)
            Write (X)
```

Which one of the schedules below is the correct serialization of the above ?

(A) $T1 \rightarrow T3 \rightarrow T2$

(B) $T2 \rightarrow T1 \rightarrow T3$

(C) $T2 \rightarrow T3 \rightarrow T1$

(D) $T3 \rightarrow T1 \rightarrow T2$

Q. 60 The following functional dependencies hold for relations R(A,B,C) and S(B,D,E):

$B \rightarrow A$

$A \rightarrow C$

The relation R contains 200 tuples and the relation S contains 100 tuples. What is the maximum number of tuples possible in the natural join $R \bowtie S$?

(A) 100

(B) 200

(C) 300

(D) 2000

Q. 61 Consider a relational table with a single record from each registered student with the following attributes.

1. Registration_Num: Unique registration number of each registered student
2. UID: Unique identity number, unique at the national level for each citizen
3. BankAccount_Num: Unique account number at the bank. A student can have multiple accounts or joint accounts. This attribute stores the primary account number.
4. Name: Name of the student
5. Hostel_Room: Room number of the hostel

Which of the following options is INCORRECT?

(A) BandAccount_Num is a candidate key

(B) Registration_Num can be a primary key

(C) UID is a candidate key if all students are from the same country

(D) If S is a superkey such that S $\cap$ UID is NULL the S $\cup$ UID is also a superkey

Q. 62 Consider a database table T containing two columns X and Y each of type integer. After the creation of the table, one record (X=1, Y=1) is inserted in the table.
Let MX and MY denote the respective maximum values of X and Y among all records in the table at any point in time. Using MX and MY, new records are inserted in the table 128 times with X and Y values being MX + 1, 2*MY + 1 respectively. It may be noted that each time after the insertion, values of MX and MY change.
What will be the output of the following SQL query after the steps mentioned above are carried out?
SELECT Y FROM T WHERE X = 7;
(A) 127
(B) 255
(C) 129
(D) 257

Q. 63 Database table by name `Loan_records` is given below:

Borrower	Bank_Manager	Loan_Amount
Ramesh	Sunderajan	10000.00
Suresh	Ramgopal	5000.00
Mahesh	Sunderajan	7000.00

What is the output of the following SQL query?

```
SELECT count(*)
FROM (
     (SELECT Borrower, Bank_Manager FROM Loan_Records)
     AS S NATURAL JOIN
     SELECT Bank-Manager, Loan_Amount FROM Loan_Records)
     AST
);
```

(A) 3
(B) 9
(C) 5
(D) 6

Q. 64 Which of the following is TRUE?
(A) Every relation in 3NF is also in BCNF
(B) A relation R is in 3NF if every non-prime attribute of R is fully functionally dependent on every key of R
(C) Every relation in BCNF is also in 3NF
(D) No relation can be in both BCNF and 3NF

Q. 65 Given the basic ER and relational models, which of the following is INCORRECT?
(A) An attribute of an entity can have more than one value
(B) An attribute of an entity can be composite
(C) In a row of a relational table, an attribute can have more than one value
(D) In a row of a relational table, an attribute can have exactly one value or a NULL value

Q. 66 Which of the following statements are TRUE about an SQL query?

P: An SQL query can contain a HAVING clause even if it does not have a GROUP BY clause

Q: An SQL query can contain a HAVING clause only if it has a GROUP BY clause

R: All attributes used in the GROUP BY clause must appear in the SELECT clause

S: Not all attributes used in the GROUP BY clause need to appear in the SELECT clause

(A) P and R (B) P and S

(C) Q and R (D) Q and S

Q. 67 Consider the following transactions with data items P and Q initialized to zero:

```
T1: Read (P);
    red (0);
    if P = 0 then Q: = Q + 1;
    write (Q);
T2: read (0);
    read (P);
    if Q = 0 then P: = P + 1;
    write (P);
```

Any non-serial interleaving of T_1 and T_2 for concurrent execution leads to

(A) a serializable schedule

(B) a schedule that is no conflict serializable

(C) a conflict serializable schedule

(D) a schedule for which a precedence graph cannot be drawn

Q. 68 Suppose $R_1(A,B)$ and $R_2(C,D)$ are two relation schemes. Let $r1$ and $r2$ be the corresponding relation instances. B is a foreign key that refers to C in $R2$. If data in $r1$ and $r2$ satisfy referential integrity constraints, which of the following is ALWAYS TRUE?

(A) $\Pi_B(r_1) - \Pi_C(r_2) = \varnothing$ (B) $\Pi_C(r_2) - \Pi_B(r_1) = \varnothing$

(C) $\Pi_B(r_1) = \Pi_C(r_2)$ (D) $\Pi_B(r_1) - \Pi_C(r_2) \neq \varnothing$

Common Data For Q. 69 and 70 :

Consider the following relations A, B and C:

A.

Id	Name	Age
12	Arun	60
15	Shreya	24
99	Rohit	11

B.

Id	Name	Age
15	Shreya	24
25	Hari	40
98	Rohit	20
99	Rohit	11

C.

Id	Phone	Age
10	2200	02
99	2100	01

Q. 69 How many tuples does the result of the following relational algebra expression contain? Assume that the schema of $A \cup B$ is the same as that of A.

$(A \cup B) \bowtie_{A.Id > 40 \vee C.Id < 15} C$

(A) 7 (B) 4

(C) 5 (D) 9

Q. 70 How many tuples does the result of the following SQL query contain?

```
SELECT A.Id
FROM A
WHERE A.Age>
All (SELECT B.Age
FROM B
WHERE B.Name='Arun')
```

(A) 4 (B) 3

(C) 0 (D) 1

****

ANSWER KEY

Databases									
1	2	3	4	5	6	7	8	9	10
(C)	(B)	(A)	(C)	(C)	(C)	(D)	(C)	(C)	(A)
11	12	13	14	15	16	17	18	19	20
(C)	(D)	(C)	(B)	(D)	(A)	(D)	(C)	(D)	(A)
21	22	23	24	25	26	27	28	29	30
(A)	(A)	(B)	(D)	(C)	(D)	(D)	(C)	(C)	(B)
31	32	33	34	35	36	37	38	39	40
(C)	(D)	(D)	(C)	(C)	(A)	(C)	(C)	(D)	(C)
41	42	43	44	45	46	47	48	49	50
(B)	(D)	(B)	(C)	(A)	(C)	(C)	(C)	(A)	(A)
51	52	53	54	55	56	57	58	59	60
(B)	(C)	(C)	(A)	(A)	(B)	(C)	(B)	(A)	(A)
61	62	63	64	65	66	67	68	69	70
(A)	(A)	(C)	(C)	(C)	(A)	(B)	(A)	(A)	(B)