# Institute of Engineering Studies (IES, Bangalore)

**Leading Institute in Bangalore for GATE/IES Classroom Coaching &**
**Leading all over India for Correspondence & Online Courses**

**Old GATE Questions along with Key from 1991-2014**
**In CS & IT Department**
**(Questions of other Departments are also added for Combined Syllabus)**

## Index- Databases






***Published by:***

Institute of Engineering Studies (IES, Bangalore),

Branches: Jayanagar & Malleshwaram of Bangalore

www.GateIesPsu.com Helpdesk: (+91) 99003-99699

## Database Design

1) Which one of the following is a key factor for preferring $B^+$ – trees to binary search trees for indexing database relations ?

**2 Marks GATE-CSE/IT-2005( )**

[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

2) A $B^+$ – 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?

**2 Marks GATE-CSE/IT-2002( )**

[A] 16
[B] 42
[C] 43
[D] 44

Database Design

**Key Paper**

| | | | |
|---|---|---|---|
| **1.** | **D** | **2.** | **B** |

## Normal forms

1) Consider a schema R(A,B,C,D) and functional dependencies A→ B and C →D. Then the decomposition of R into R1 $R_1$ (AB) and $R_2$ (CD) is

1 Marks GATE-CSE/IT-2001( )

[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

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

1 Marks GATE-CSE/IT-2001( )

[A] A →B, B →CD
[B] A →B, B →C, C →D
[C]AB →C, C →AD
[D]A →BCD

3) 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?

2 Marks GATE-CSE/IT-2000( )

[A] XY –→ Z and Z →Y
[B] YZ →X and Y →Z
[C]YZ →X and X →Z
[D]XZ →Y and Y →X

4) Let R = (a, b, c, d, e, f) be a relation scheme with the following dependencies c →f, e →a, ec →d. a →b. Which of the following is a key for R ?

1 Marks GATE-CSE/IT-1999( )

[A] CD
[B] EC
[C]AE
[D]AC

5) Consider the schema R = (S T U V) and the dependencies S → T, T →U. U →V and V →S. Let R = (R1 and R2) be a decomposition such that R1 ∩ R2 = ∅ . The decomposition is

2 Marks GATE-CSE/IT-1999( )

[A] not in 2NF
[B] in 2NF but not 3NF
[C] in 3NF but not in 2NF
[D]in both 2NF and 3NF

6) Which normal form is considered adequate for normal relational database design?

1 Marks GATE-CSE/IT-1998( )

[A] 2 NF
[B] 5 NF
[C]4 NF
[D]3 NF

7) There are five records in a database

| Name | Age | Occupation | Category |
|---|---|---|---|
| Rama | 27 | CON | A |
| Abdul | 22 | ENG | A |
| Jeniffer | 28 | DOC | B |
| Maya | 32 | SER | D |
| Dev | 24 | MUS | C |

There is an index file associated with this and it contains the values 1,3,2,5 and 4.Which one of the fields is the index built from?

1 Marks GATE-CSE/IT-1998( )

[A] Age
[B] Name
[C]Occupation
[D]Category

8) For a database relation R(a,b,c,d), where the domains of a, b, c, d include only atomic values, only the following functional dependencies and those that can be inferred from them hold:

a →c

b →d This relation is

## Normal forms

2 Marks GATE-CSE/IT-1997( )

[A] in first normal form but not in second normal form
[B] in second normal form but not in third normal form
[C] in third normal form
[D] None of the above

9) Let R (a, b, c) and S(d, e, f) be two relations in which d is the foreign key of S that refers to the primary key of R. Consider the following four operations R and S
(a) Insert into R
(b) Insert into S
(c) Delete from R
(d) Delete from S
Which of the following is true about the referential integrity constraint above?

2 Marks GATE-CSE/IT-1997( )

[A] None of (a), (b), (c) or (d) can cause its violation
[B] All of (a), (b), (c) and (d) can cause its violation
[C] Both (a) and (d) can cause its violation
[D] Both (b) and (c) can cause its violation

10) Consider the following functional dependencies 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_ birth, Age)is

2 Marks GATE-CSE/IT-2003( )

[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

11) the relation scheme studentPerformance (name, course No, rollNo, grade) has the following functionaldependencies :
Name, course No →grade
RollNo, course No →grade
name →rollNo
rollNo →name
the highest normal form ofthis relation scheme is

2 Marks GATE-CSE/IT-2004( )

[A] 2 NF
[B] 3 NF
[C] BCNF
[D] 4 NF

12) The following functionaldependencies are given :
AB CD, AF D, DE F,C G,F E,G A
which one of the followingoptions is false ?

1 Marks GATE-CSE/IT-2006( )

[A] $\{CF\}^+ = \{ACDEFG\}$
[B] $\{BG\}^+ = \{ABCDG\}$
[C] $\{AF\}^+ = \{ACDEFG\}$
[D] $\{AB\}^+ = \{ABCDG\}$

13) Consider the followingrelational schemes for a librarydatabase:
Book (Title, Author, Catalog_no, Publisher, Year , price)
Collection (Title, Author,Catalog_no)
Which the followingfunctional dependencies :
I. Title Author → Catalog _no
II. Catalog_no → Title Author Publisher Year
III. Publisher TitleYear → Price
Assume {Author, Title} is the key for bothschemes : Which of the following statements is true ?

2 Marks GATE-CSE/IT-2008( )

[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

14) Assume that, in thesuppliers relation above, each supplier and each street within a city has aunique name, and (sname, city) forms a candidate key. No other functional dependencies are impliedother than those implied by primary and candidate keys. Which one of the following is TRUE about theabove schema?

2 Marks GATE-CSE/IT-2009( )

[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 3 NF
[D] The schema is not in 2NF

## Normal forms

15) The relation scheme student Performance (name, course No, rollNo, grade) has the following functional dependencies :

Name, course No →grade
RollNo, course No →grade
name →rollNo
rollNo →name

the highest normal form of this relation scheme is

2 Marks GATE-CSE/IT-2004( )

[A] 2 NF
[B] 3 NF
[C]BCNF
[D] 4 NF

16) Which-one of the following statements about normal forms is FALSE ?

1 Marks GATE-CSE/IT-2005( )

[A](a) BCNF is stricter than 3 NF
[B] (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

17) The following functional dependencies are given :
AB → CD, AF→ D, DE →F,C →G,F→ E, G → A
which one of the following options is false ?

2 Marks GATE-CSE/IT-2006( )

[A]$\{CF\}^+ = \{ACDEFG\}$
[B] $\{BG\}^+ = \{ABCDG\}$
[C] $\{AF\}^+ = \{ACDEFG\}$
[D]$\{AB\}^+ = \{ACDFG\}$

18) Which one of the following statements is FALSE ?

2 Marks GATE-CSE/IT-2007( )

[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 BCNF relation

19) 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 ?

2 Marks GATE-CSE/IT-2008( )

[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

20) Which of the following is TRUE?

1 Marks GATE-CSE/IT-2012( )

[A] Every relation is 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

21) 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

1 Marks GATE-CSE/IT-2002( )

[A] Zero
[B] More than zero but less than that of an equivalent 3NF decomposition
[C]Proportional to the size of F+
[D]Indetermine

22) With regard to the expressive power of the formal relational query languages, which of the following statements is true?

1 Marks GATE-CSE/IT-2002( )

[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

23) The binary relation S = ϕ (empty set) on set A = {1,2,3} is

2 Marks GATE-CSE/IT-2002( )

[A] Neither reflexive nor symmetric

[B] Symmetric and reflexive

[C] Transitive and reflexive

[D] Transitive and symmetric

24) 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).

2 Marks GATE-CSE/IT-2002( )

[A] Dependency-preservation

[B] Lossless-join

[C] BCNF definition

[D] 3NF definition

25) From the following instance of a 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 |

2 Marks GATE-CSE/IT-2002( )

[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 determines B and B does not functionally determines C

## Normal forms

**Key Paper**

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| **1.** | **C** | **2.** | **C** | **3.** | **A** | **4.** | **B** | **5.** | **D** |
| **6.** | **D** | **7.** | **C** | **8.** | **A** | **9.** | **D** | **10.** | **D** |
| **11.** | **B** | **12.** | **C** | **13.** | **C** | **14.** | **B** | **15.** | **B** |
| **16.** | **C** | **17.** | **C** | **18.** | **D** | **19.** | **C** | **20.** | **C** |
| **21.** | **B** | **22.** | **B** | **23.** | **D** | **24.** | **C** | **25.** | **B** |

## ER diagrams

**Common Data for Q1 and Q2 is given below**

Consider the following ER diagram

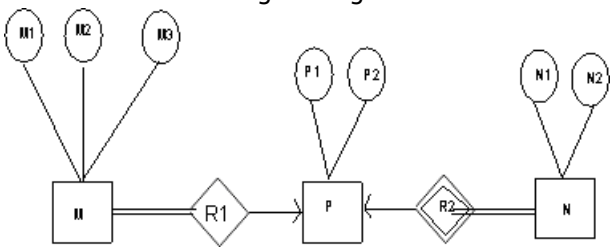


1) The minimum number of tables needed to represent M,N,P,R1, R2 is

**2 Marks GATE-CSE/IT-2008,GATE-CSE/IT-2008( )**

[A] 2 [B] 3
[C] 4 [D] 5

2) Which of the following is a correct attribute set for one of the tables for the correct answer to the above question ?

**2 Marks GATE-CSE/IT-2008( )**

[A] {M1, M2, M3,P1} [B] {M1, P1, N1, N2}
[C] {M1, P1, N1} [D] {M1, P1}

**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)

3) Consider the following relational query on the above database:

```
SELECT S.name
FROM Suppliers S
WHERE S. Sid NOT IN (SELECT C.sid
                     FROM Catalog C
                     WHERE C. Pid NOT IN (SELECT P. pid
                     FROM Parts 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?

**2 Marks GATE-CSE/IT-2009( )**

[A] Find the names of all supplies 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

**Common Data for Q5 and Q4 is given below**

Consider the following ER diagram

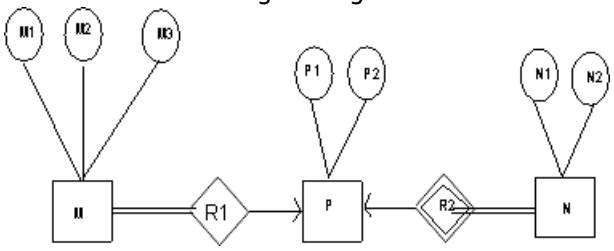


4) The minimum number of tables needed to represent M,N,P,R1, R2 is

ER diagrams

**2 Marks GATE-CSE/IT-2008( )**

[A] 2 [B] 3
[C] 4 [D] 5

5) Which of the following is a correct attribute set for one of the tables for the correct answer to the above question ?

**2 Marks GATE-CSE/IT-2008( )**

[A] {M1, M2, M3, P1} [B] {M1, P1, N1, N2}
[C] {M1, P1, N1} [D] {M1, P1}

6) It is desired to design anobject-oriented employee record system for a company. Each employee has a name,unique id and salary, Employees belongto different categories and their salary is determined by their category. The functions get Name, geld and compute Salaryare required. Given the class hierarchybelow, possible locations for these functions are :
(i) getld is implemented inthe superclass
(ii) Getld is implemented inthe subclass
(iii) get Name is an abstractfunction in the superclass
(iv) get Name is implementedin the superclass
(v) getName is implemented inthe subclass
(vi) getsalary is an abstractfunction in the superclass
(vii) getSalary isimplemented in the superclass
(viii) get Salary isimplemented in the subclass

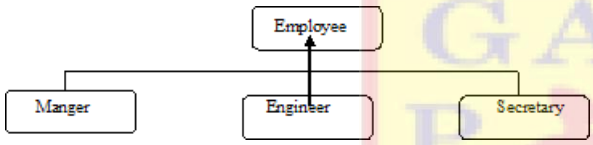


**2 Marks GATE-CSE/IT-2004( )**

[A] (i), (iv), (vi), (viii) [B] (i), (iv), (vii)
[C] (i), (iii), (v), (vi), (viii) [D] (ii), (v), (vii)

7) Le E1 and E2 be two entities in an E/R diagram with simple single valued attributes. R1 and R2 are two relationships between E1 and E2, where R1 is one-to-many and R2 is many-to-many. R1 and R2 do not have any attributes of their own. What is the minimum number of tables required to represent this situation in the relational model ?

**2 Marks GATE-CSE/IT-2005( )**

[A] 2 [B] 3
[C] 4 [D] 5

8) Le $E_1$ 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 ?

**2 Marks GATE-CSE/IT-2005( )**

[A] 2 [B] 3
[C] 4 [D] 5

9) Consider a relation scheme R= (A,B,C,D,E,H) on which the following functional dependencies hold: {A→ B, BC→ D,E→ C, D→ A}. What are the candidate keys of R ?

**2 Marks GATE-CSE/IT-2005( )**

[A] AE, BE [B] AE, BE, DE
[C] AEH, BEH, BCH [D] AEH, BEH, DEH

ER diagrams

10) 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 ?

**1 Marks GATE-CSE/IT-2006( )**

[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 record 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

11) 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 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

**2 Marks GATE-CSE/IT-2009( )**

[A] 2

[B] 3

[C] 4

[D] 5

12) Given the basic ER and relational models, which of the following is INCORRECT?

**1 Marks GATE-CSE/IT-2012( )**

[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

13) Given max heap with level order elements as 10, 8, 5, 3, 2 in order. Insert 1 and 7 into the heap tree and the BFS of resultant tree.

**2 Marks GATE-CSE/IT-2014( )**

[A] #

[B] #

[C] 10, 8, 7, 3, 2, 1, 5

[D] #

ER diagrams

**Key Paper**

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| **1.** | **B** | **2.** | **A** | **3.** | **A** | **4.** | **B** | **5.** | **A** |
| **6.** | **A** | **7.** | **B** | **8.** | **B** | **9.** | **D** | **10.** | **C** |
| **11.** | **B** | **12.** | **C** | **13.** | **C** | | | | |

## Transaction management

1) For the schedule given below, which of the following is correct:

| | | |
|---|---|---|
| 1 | Read | |
| 2 | | Read B |
| 3 | Write A | |
| 4 | | Read A |
| 5 | | Write A |
| 6 | | Write B |
| 7 | Read B | |
| 8 | Write B | |

2 Marks GATE-CSE/IT-1999( )

[A] This schedule is serialized and can occur in a scheme using 2PL protocol

[B] This schedule is serializable but cannot occur in a scheme using 2PL protocol

[C] This schedule is not serialiable but can occur in a scheme using 2PL protocol

[D] This schedule is not seralisable and cannot occur in a scheme using 2PL protocol.

2) Consider the following logsequence 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 systemcrashed just before log record 7 is written. When the system is restarted,which one statement is true of the recovery procedure ?

2 Marks GATE-CSE/IT-2006( )

[A] We must undo log record 6 to set B to 10000 and then redo log record 2 and 3

[B] We need not redo log records 2 and 3 because transaction T1 has committed

[C] We can apply redo and undo operations in arbitrary order because they are idempotent

[D] We must redo log record 6 to set B to 10500

3) Consider the followingschedules involving two transactions.Which one of the following statements is TRUE ?

S1 : r1(X);r1(Y); r2(X); r2 (Y); w2(Y); w1(X)

S2 : r1(X); r2(X); r2(Y); W2(Y);r1(Y); w1(X)

2 Marks GATE-CSE/IT-2007( )

[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

4) 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); | |
| | W(D2); | |
| | | R(D2); |
| | | R(D3) |
| R(D1); | | |
| W(D1); | | W(D2); |
| | | W(D3); |
| | R(D1); | |
| R(D2); | | |
| W(D2); | | |
| | W(D1); | |

2 Marks GATE-CSE/IT-2003( )

[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

5) Consider the following schedules involving two transactions.
Which one of the following statements is TRUE ?
$S_1$ : $r_1(X)$; $r_1(Y)$; $r_2(X)$; $r_2(Y)$; $w_2(Y)$; $w_1(X)$
$S_2$ : $r_1(X)$; $r_2(X)$; $r_2(Y)$; $W_2(Y)$; $r_1(Y)$; $w_1(X)$

**2 Marks GATE-CSE/IT-2007( )**

[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

6) 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 :
$T_1$: $R_1[x]$ $W_1[x]$ $W_1[y]$
$T_2$: $R_2[x]$ $R_2[y]$ $W_2[y]$
$S_1$: $R_1[x]$ $R_2[x]$ $R_2[y]$ $W_1[x]$ $W_1[y]$ $W_2[y]$
$S_2$: $R_1[x]$ $R_2[x]$ $R_2[y]$ $W_1[x]$ $W_2[y]$ $W_1[y]$
$S_3$: $R_1[x]$ $W_1[x]$ $R_2[x]$ $W_1[y]$ $W_1[y]$
$S_4$: $R_2[x]$ $R_2[y]$ $R_1[x]$ $W_1[x]$ $W_1[y]$ $W_2[y]$

Which of the above schedules are conflict – serializable ?

**2 Marks GATE-CSE/IT-2009( )**

[A] $S_1$ and $S_2$

[B] $S_2$ and $S_3$

[C] $S_3$ only

[D] $S_4$ only

Transaction management

**Key Paper**

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| **1.** | **D** | **2.** | **B** | **3.** | **C** | **4.** | **D** | **5.** | **C** |
| 6. | **B** | | | | | | | | |

## File structure in DBMS

1) The order of a leaf node in aB+- tree is the maximum number of (value, data record pointer) pairsit can hold. Given that the block size is 1K bytes. Data record pointer is 7bytes long, the value field is 9 bytes long and a block pointer is 6 byteslong, what is the order of the leaf node ?

**2 Marks GATE-CSE/IT-2007( )**

[A]63 [B]64
[C]67 [D]68

2) A clustering index is definedon the fields which are of type

**1 Marks GATE-CSE/IT-2008( )**

[A] Non-key and ordering [B] Non-key and non-ordering
[C]Key and ordering [D]Key and non-ordering

3) Consider a file of 1684 records. Each record is32 bytes long and its key field is of size 6 bytes. The file is ordered on a non-key field, andthe file organization is unspanned. Thefile is stored in a file system with block size 1024 bytes, and the size of ablock pointer is 10 bytes. If thesecondary index is built on the key field of the file, and a multi-level index scheme is used to store the secondaryindex, the number of first-level and second-level blocks in the multi-level indexare respectively

**2 Marks GATE-CSE/IT-2008( )**

[A] 8 and 0 [B]128 and 6
[C]256 and 4 [D]512 and 5

4) Which one of the following is a key factor for preferring B+- trees to binary search trees for indexing database relations ?

**2 Marks GATE-CSE/IT-2005,GATE-CSE/IT-2005( )**

[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

5) 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 ?

**2 Marks GATE-CSE/IT-2007( )**

[A]63 [B]64
[C]67 [D]68

6) A clustering index is defined on the fields which are of type

**1 Marks GATE-CSE/IT-2008( )**

[A] Non-key and ordering [B] Non-key and non-ordering
[C] Key and ordering [D] Key and non-ordering

7) 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

**2 Marks GATE-CSE/IT-2008( )**

[A] 8 and 0 [B]128 and 6
[C]256 and 4 [D]512 and 5

File structure in DBMS

**Key Paper**

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| **1.** | **A** | **2.** | **A** | **3.** | **C** | **4.** | **D** | **5.** | **B** |
| **6.** | **A** | **7.** | **C** | | | | | | |

## Relationa Algera & Calculus

1) The relational algebra expression equivalent to the following tuple calculus expression
{t | t ∈ r ∧ (t [ A] = 10 ∧ t [B] = 20)} is

1 Marks GATE-CSE/IT-1999( )

[A] $\sigma_{(A=10 \vee B=20)}(r)$
[B] $\sigma_{(A=10)}(r) \cup \sigma_{(B=20)}(r)$
[C] $\sigma_{(A=10)}(r) \cap \sigma_{(B=20)}(r)$
[D] $\sigma_{(A=10)}(r) - \sigma_{(B=20)}(r)$

2) Given two union compatible relations $R_1$(A,B) and $R_2$(C,D), what is the result of the operation $R_1$A = CAB = D$R_2$(a)

1 Marks GATE-CSE/IT-1998( )

[A] $R_1 \cup R_2$
[B] $R_1 \times R_2$
[C] $R_1 - R_2$
[D] $R_1 \cap R_2$

3) Let R1 (A,B,C) and R2(D,E)be two relation schema, where the primary keys are shown underlined, and let Cbe a foreign key in R1 referring to R2. Si[[pse tjee os mp violation of theabove referential integrity constraint in the corresponding relation
instances r1 and r2. Which one of the following relational algebraexpressions would necessarily produce an empty relation ?

2 Marks GATE-CSE/IT-2004( )

[A] $(r_2) - (r_1)$
[B] $(r_1) - (r_2)$
[C] $(r_1 R_2) - (r_1)$
[D] $(r_1 R_2)$

4) Let R and S be two relationswith the following schema
R(P,Q,R1,R2,R3)
S(P,Q,S1 S2)
Where{P,Q} is the key for both schemas. Which of thefollowing queries are equivalent ?

I. ( S)
II. (R) (S)
III. ( (R) (S))
IV. ( (R)– ( (R) – (S)))

2 Marks GATE-CSE/IT-2008( )

[A] Only I and II
[B] Only I and III
[C] Only I, II and III
[D] Only I, III, and IV

5) Consider the following SQL query
select distinct a1, a2,….,an
from r1, r2…, rm
where P

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

1 Marks GATE-CSE/IT-2003( )

[A] $a_1, a_2, \ldots \pi, a_n \sigma_p (r_1 \times r_2 \times \ldots \times r_m)$
[B] $a_1, a_2, \ldots \pi, a_n \sigma_p (r_1 \bowtie r_2 \bowtie \ldots \bowtie r_m)$
[C] $a_1 a_2, \ldots \pi, \sigma_p (r_1 \cup r_2 \cup \ldots \cup r_m)$
[D] $a_1, a_2, \ldots \pi, \sigma_p (r_1 \cap r_2 \cap \ldots \cap r_m)$

6) An index is clustered, if

1 Marks GATE-CSE/IT-2013( )

[A] it is on a set of fields that form a candidate key.
[B] it is on a set of fields that include the primary key.
[C] the data records of the file are organized in the same order as the data entries of the index.
[D] the data records of the file are organized not in the same order as the data entries of the index

7) 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 $r_2$. Which one of the following relational algebra expressions would necessarily produce an empty relation ?

1 Marks GATE-CSE/IT-2004( )

[A] $\pi_D (r_2) - \pi_C (r_1)$
[B] $\pi_C (r_1) - \pi_D (r_2)$
[C] $\pi_D (r_1 \bowtie_{C=d} R_2) - \pi_C (r1)$
[D] $\pi_C (r_1 \bowtie_{C=d} R_2)$

8) 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?

$$\pi_{name}(r_{sex=females}(Student))\, P_{name}\,(Student \triangleright\!\triangleleft r_{n,x,m}(Student))$$
(sex=females
^x =male
^marks ≤m)

2 Marks GATE-CSE/IT-2004( )

[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

9) 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?

2 Marks GATE-CSE/IT-2004( )

[A] 24 [B] 25

[C] 26 [D] 27

10) Let r be a relation instance with schema R= (A, B,C,D). We define $r_1 = {}^{\pi A,B,C}(r)$ and $r_2 = {}^{\pi A,D}(r)$. 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 he following is TRUE?

1 Marks GATE-CSE/IT-2004( )

[A] $s \subset r$ [B] $r \cup s = r$

[C] $r \subset s$ [D] r*s =s

11) Information about a collection of students is given by the relation studInfo(studId, name, sex0 The relation enroll (studID, CoursId) 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?

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

2 Marks GATE-CSE/IT-2007( )

[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

12) Consider the relation employee (name, sex, supervisor Name(with name as the key. Supervisor Name- gives the name of the supervisor of the employee under consideration. What does the following Tuple {e.name | employee (e)

$$(\forall x)[\neg employee(x) \vee x.supervisorName$$

"male"]}

2 Marks GATE-CSE/IT-2007( )

[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

13) Let R and S be two relations with the following schema
R(P,Q,R1,R2,R3)
S(P,Q,S1 S2)
Where{P,Q} is the key for both schemas. Which of the following queries are equivalent ?
I. $\prod_{P}(\triangleright\triangleleft S)$
II. $\prod_{P}(R) \triangleright\triangleleft \prod_{P}(S)$
III. $\prod_{P}(\prod_{P,Q,}(R) \bigcap \prod_{P,Q}(S))$
IV. $\prod_{P}(\prod_{P,Q,}(R)(\prod_{P,Q,}(R), \prod_{P,Q}(S)))$

**2 Marks GATE-CSE/IT-2008( )**

[A] Only I and II
[B] Only I and III
[C] Only I, II and III
[D] Only I, II, and IV

14) Suppose $R_1$ (A, B) and $R_2$ (C, D) are two relation schemas. Let $r_1$ and $r_2$ be the corresponding relation instances. B is a foreign key that refers to C in $R_2$. If data in $r_1$ and $r_2$ satisfy referential integrity constrains, which of the following is ALWAYS TRUE?

**2 Marks GATE-CSE/IT-2012( )**

[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] $\Pi_B(r_1) - \Pi_C(r_2) \neq \varnothing$

15) Consider a relational table r with sufficient number of records, having attributes $A_1$, $A_2$,..., $A_n$ and let $1 \leq p \leq n$. Two queries Q1 and Q2 are given below.
Q1 : $\pi_{A_1 \ldots A_n}(\sigma_{A_p = c}(r))$ where c is a const
Q2 : $\pi_{A_1 \ldots A_n}(\sigma_{c_1 \leq A_p \leq c_2}(r))$ where c1 and c2 are constants
The database can be configured to do ordered indexing on $A_p$ or hashing on $A_p$. Which of the following statements is TRUE?

**2 Marks GATE-CSE/IT-2011( )**

[A] Ordered indexing will always outperform hashing for both queries
[B] Hashing will always outperform ordered indexing for both queries
[C] Hashing will outperform ordered indexing on Q1, but not on Q2
[D] Hashing will outperform ordered indexing on Q2, but not on Q1

16) 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 $\sigma_{A=a}(r \bowtie s)$ always equal to

**1 Marks GATE-CSE/IT-2001( )**

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

Relationa Algera & Calculus

**Key Paper**

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| **1.** | **C** | **2.** | **D** | **3.** | **B** | **4.** | **D** | **5.** | **A** |
| **6.** | **C** | **7.** | **A** | **8.** | **D** | **9.** | **C** | **10.** | **C** |
| **11.** | **C** | **12.** | **C** | **13.** | **C** | **14.** | **A** | **15.** | **C** |
| **16.** | **C** | | | | | | | | |

## SQL queries

**Common Data for Q1 and Q2 is given below**

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 | Area |
|---|---|---|
| 10 | 220 | 02 |
| 99 | 2100 | 01 |

1) 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')
```

**2 Marks GATE-CSE/IT-2012,GATE-CSE/IT-2012( )**

[A] 4 [B] 3

[C] 0 [D] 1

2) How many tuples does the result of the following relational algebra expression contain? Assume that the schema of A ∪ B is the same as that of A.

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

**2 Marks GATE-CSE/IT-2012( )**

[A] 7 [B] 4

[C] 5 [D] 9

3) Consider a relational table with a single record for each registered student with the following attributes.

1. Registration_Number: Unique registration number for each registered student
2. UID: Unique Identity number, unique at the national level for each citizen
3. BankAccount_Number: Unique account number at the bank. A student can have multiple accounts or joint accounts. This attributes 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?

**1 Marks GATE-CSE/IT-2011( )**

[A] BankAccount_Number is a candidate key

[B] Registration_Number 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 ∩ UID is NULL then S ∪ UID is also a superkey

## SQL queries

4) 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) AS T
);

**2 Marks GATE-CSE/IT-2011( )**

[A] 3 [B] 9
[C] 5 [D] 6

5) 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=l) 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;

**2 Marks GATE-CSE/IT-2011( )**

[A] 127 [B] 255
[C] 129 [D] 257

6) A relational schema for a train reservation database is given below
Passenger (pid, pname, age)
Reservation (pid, cass, 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 |
| 5 | 'AC' | 8203 |
| 1 | 'SC' | 8204 |
| 3 | 'AC' | 8202 |

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

```
SELECT    pid
 FROM     Re servation
WHERE        class = 'AC' AND

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

**1 Marks GATE-CSE/IT-2010( )**

[A] 1,0 [B] 1,2
[C] 1,3 [D] 1,5

## SQL queries

7) 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 200tuples and the relation S contains 100tuples. What is the maximum number of tuples possible in the natural join $R \bowtie S$?

**2 Marks GATE-CSE/IT-2010( )**

[A] 100 [B] 200
[C] 300 [D] 2000

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

**2 Marks GATE-CSE/IT-2001( )**

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

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

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

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

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

```
create table geq
(
lb          integer     not null
ub          integer     not null
primary key lb
foreign key (ub) references geq on delete cascade)
```

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

**2 Marks GATE-CSE/IT-2001( )**

[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

10) 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

**2 Marks GATE-CSE/IT-2000( )**

[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

11) In SQL, relations can contain null values, and comparisons with null values are treated as unknown. Suppose all comparisons with a null value are treated as false. Which of the following pairs is not equivalent?

**2 Marks GATE-CSE/IT-2000( )**

[A] $x = 5$ not (not $(x = 5)$)
[B] $x = 5$ $x > 4$ and $x < 6$, where x is an integer
[C] $x \neq 5$ not $(x = 5)$
[D] None of the above

12) Consider the join of a relation R with a relation S. If R has m tuples and S has n tuples then the maximum and minimum sizes of the join respectively are

**1 Marks GATE-CSE/IT-1999( )**

[A] $m + n$ and 0 [B] $mn$ and 0
[C] $m + n$ and $|m - n|$ [D] $mn$ and $m + n$

13) Which of the following is/are correct?

**2 Marks GATE-CSE/IT-1999( )**

[A] An SQL query automatically eliminates duplicates
[B] An SQL query will not work if there are no indexes on the relations
[C] SQL permits attribute names to be repeated in the same relation
[D] None of the above

## SQL queries

14) Consider the set of relationsshown 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 distinct Name
From Students, courses,Grades
Where Students. Roll_number=Grades. Roll_ 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 ?

**2 Marks GATE-CSE/IT-2003( )**

[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 graded 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.

15) 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 ?

**2 Marks GATE-CSE/IT-2004( )**

[A]24
[B]25
[C]26
[D]27

16) 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 (selectcount(*)
      from book as T
       where T. price >B. Price) <5
```

**2 Marks GATE-CSE/IT-2005( )**

[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]Title of the five most expensive books

17) Consider the relation account(customer, balance) where customer is a primary key and there are no nullvalues. We would like to rank customersaccording to decreasing balance. Thecustomer with the largest balance gets rank. 1. Ties are not broken but ranks are skipped: if exactly two customers have the largest balance theyeach get rank 1 and rank 2 is not assigned.
Query 1: Select A. customer, count (B. customer)fromaccount A, account B where A. customer
Query 2 : Select A. customer,1 + count (B. customer) from account
A, account Bwhere A. balance <B. balance 7 group by A. customer
Consider these statementsabout Query 1 and Query 2.
1. Query 1 will produce thesame row set as Query 2 for some but not all databases
2. Both Query 1 Query 2 are correct implementations of thespecification
3. Query 1 is a correctimplementation of the specification but Query 2 is not
4. Neither Query 1 nor Query2 is a correct implementation of the specification
5. Assigning rank with a purerelational Query takes less time than scanning in decreasing balance order andassigning ranks using ODBC
Which two of the abovestatements are correct ?

**2 Marks GATE-CSE/IT-2006( )**

[A]1 and 3
[B]1 and 4
[C]3 and 5
[D]2and 5

## SQL queries

18) Consider the relationenrolled (student, course) in which (student, course) is the primary key, andthe relation paid(Student,amount) wherestudent is the primary key. Assume nonull values and no foreign key. Assumeno null values and no foreign keys or integrity constraints, Given the following four queries :
Query 1: select student fromenrolled where student in select student form paid)
Query 2: Select student frompaid where student in (select student from enrolled)
Query 3: Select E. studentfrom enrolled E, paid P where E. student= P student
Query 4: Select student frompaid where exists (select * from enrolled where enrolled. Student = paid. student
Which one of the followingstatements is correct ?

**2 Marks GATE-CSE/IT-2006( )**

[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 databases for which Query 1 and Query 2 return 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 integrity violation at runtime.

19) Information about a collection of students is given by the relation studInfo(studId, name, sex0 The relation enroll (studID, CoursId) 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 ?

**2 Marks GATE-CSE/IT-2007( )**

[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

20) Consider the relation employee (name, sex, supervisor Name(with name as the key. Supervisor Name- gives the name of the supervisor of the employee under consideration. What does the following Tuple
{e.name | employee (e)

__(x) [ employee (x) x. supervisor Name

e.name x.sex = "male"]}

**2 Marks GATE-CSE/IT-2007( )**

[A] Names of employees with no immediate male subordinates
[B] Names of employees with no immediate female subordinates
[C] Names of employees with a female supervisor
[D]names of employees with a male supervisor

21) Consider the table employee(empId, name, department, salary) and the two queries Q1, Q2 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 ?

Q1 : Select e. empId
From employee e
Where not exists
(Select * Fro employee s Where s. department = "5" and
s. salary >= e. salary)

Q2: Select e. empId
From employee e
Where e. salary > Any

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

**2 Marks GATE-CSE/IT-2007( )**

[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

22) Which one of the following statements is FALSE ?

**2 Marks GATE-CSE/IT-2007( )**

[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 BCNF relation

SQL queries

23) 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?

$\prod_{nsme}(r_{sex=femals}))P_{name}(Student) \bowtie_{(sex=female)} r_{nxm}(Student))$

^ x=male

^ marks <= m)

2 Marks GATE-CSE/IT-2004( )

[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

24) Which-one of the following statements about normal forms is FALSE?

1 Marks GATE-CSE/IT-2005( )

[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

25) Let r be a relation instance with schema R= (A, B,C,D). We define $r_1 = \prod_{A,B,C}(r)$ and $r_1 = \prod_{A,D}(r)$ let S= r1 * r2 where * denotes natural join. Given that the decomposition of r into r1 and r2 is lossy, which one of he following is TRUE?

1 Marks GATE-CSE/IT-2005( )

[A] s⊂ r

[B] r U s = r

[C] r ⊂s

[D] r*s =s

26) 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 | B |
|---|---|
| 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 :

2 Marks GATE-CSE/IT-2005( )

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

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

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

[D] 1

27) Consider the following relation schema pertaining to a students database;
Student (rollno, name, address)
Enroll(rollno, 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 ?

1 Marks GATE-CSE/IT-2004( )

[A] 8,8

[B] 120,8

[C] 960,8

[D] 960,120

## SQL queries

28) The employee information in a company is stored in the relation Employee(name, sex, salary, dept Name)
Consider the following SQL query

```
Select dept Name
        From Employee
        where sex = 'M'
  group by dept Name
  having avg (salary)
        (select avg(salary) from Employee)
```

It returns the names of the department in which

**2 Marks GATE-CSE/IT-2004( )**

[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 male employees is more than the average salary in the company.

29) 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 | B |
|---|---|
| 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 :

**2 Marks GATE-CSE/IT-2005( )**

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

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

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

[D] 1

30) 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
```

**2 Marks GATE-CSE/IT-2005( )**

[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] Title of the five most expensive books

## SQL queries

31) 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 key. Assume no null values and no foreign keys or integrity constraints, Given the following four queries :
Query 1: select student 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: Select student from paid where exists (select * from enrolled where enrolled. Student = paid. student
Which one of the following statements is correct ?

**2 Marks GATE-CSE/IT-2006( )**

[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 databases for which Query 1 and Query 2 return 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 integrity violation at runtime.

32) 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".

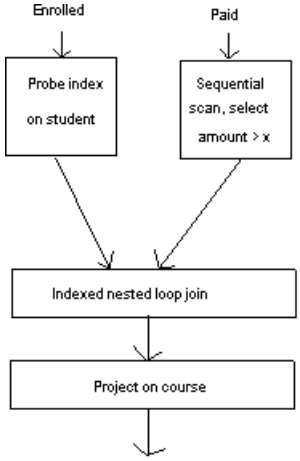


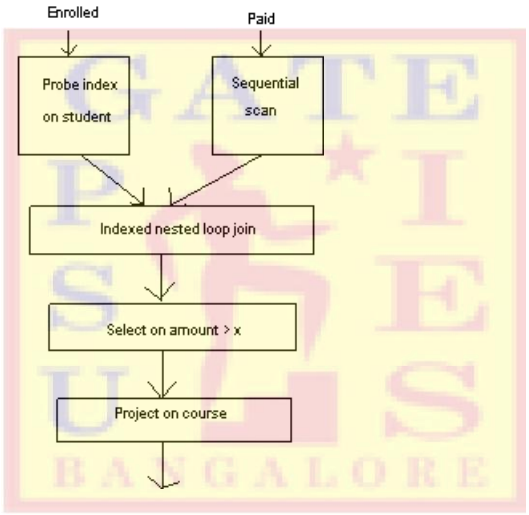


A disk seek takes 4 ms. Disk data transfer bandwidth is 300 MB/s and checking a tuple to see if amount us greater than x takes 10 s Which of the following statements is correct ?

**2 Marks GATE-CSE/IT-2006( )**

[A] plan 1 and Plan 2 will not output identical row sets for all databases

[B] A course may be listed more than once in 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

33) 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
Where not exists
(Select * Fro employee s Where s. department = "5" and s. salary >= e. salary)

$Q_2$: Select e. empId
From employee
Where e. salary >Any
(Select distinct salary From employee s Where s. department = "5")

## SQL queries

**2 Marks GATE-CSE/IT-2007( )**

[A] Q1 is the correct query.
[B] Q2 is the correct query
[C] Both Q1 and Q2 produce the same answer
[D] Neither $Q_1$ nor $Q_2$ is the correct query

34) Consider the following relational schema.
Students(rollno:integer, sname: string)
Courses(courseno:integer, cname: string)
Registration(rollno:integer,courseno;integer, percent: real)
Which of the following queries are equivalent to this query in English?
"Find the distinct names of all students who score more than 90% in the course numbered 107"
(I) SELECT DISTINCT S.sname
FROM Students as S, Registration as R
WHERE R.rollno=S.rollno AND R.Courseno=107 AND R.percent>90
(II) $\Pi_{sname}(\sigma_{courseno=107 \wedge percent > 90}(Registration \quad Student))$
(III) {T | ∃S ∈ Students, ∃R ∈ Registration (S.rollno = R.rollno ∧R.courseno = 107 ∧ R.percent >90 ∧ T.sname = S.name)}
(IV) $\{< S_N > | \exists S_R \exists R_P (< S_R , S_N > \in \text{Students} \wedge < S_R, 107, R_P > \in \text{Registration} \wedge R_P > 90)\}$

**2 Marks GATE-CSE/IT-2013( )**

[A] I, II, III and IV
[B] I, II and III only
[C] I, II and IV only
[D] II, III and IV only

35) 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 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

**1 Marks GATE-CSE/IT-2012( )**

[A] P and R
[B] P and S
[C] Q and R
[D] Q and S

36) Suppose a circular queue of capacity (n – 1) elements is implemented with an array of n elements. Assume that the insertion and deletion operations are carried out using REAR and FRONT as array index variables, respectively. Initially, REAR = FRONT = 0. The conditions to detect queue full and queue empty are

**2 Marks GATE-CSE/IT-2012( )**

[A] full: (REAR+1) mod n==FRONT
empty: REAR ==FRONT
[B] full:(REAR+1)mod n==FRONT
empty: (FRONT+1)mod n==REAR
[C] full:
REAR==FRONT
empty: (REAR+1) mod n ==FRONT
[D] full:(FRONT+1)mod n==REAR
empty: REAR ==FRONT

**Statement for Linked answer Q37 and Q38 is given below**

37) Relation R has eight attributes ABCDEFGH. Fields of R contain only atomic values.
F = {CH → G, A → BC, B → CFH, E → A, F → EG} is a set of functional dependencies (FDs) so that F+ is exactly the set of FDs that hold for R

Q.How many candidate keys does the relation R have?

**2 Marks GATE-CSE/IT-2013,GATE-CSE/IT-2013( )**

[A] 3
[B] 4
[C] 5
[D] 6

38) The relation R is

**2 Marks GATE-CSE/IT-2013( )**

[A] in INF, but not in 2NF
[B] in 2NF, but not in 3NF
[C] in 3NF, but not in BCNF
[D] in BCNF

## SQL queries

**Key Paper**

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| **1.** | **B** | **2.** | **A** | **3.** | **A** | **4.** | **C** | **5.** | **A** |
| **6.** | **C** | **7.** | **A** | **8.** | **C** | **9.** | **B** | **10.** | **A** |
| **11.** | **C** | **12.** | **B** | **13.** | **D** | **14.** | **C** | **15.** | **D** |
| **16.** | **D** | **17.** | **B** | **18.** | **B** | **19.** | **C** | **20.** | **B** |
| **21.** | **B** | **22.** | **D** | **23.** | **D** | **24.** | **C** | **25.** | **C** |
| **26.** | **C** | **27.** | **A** | **28.** | **D** | **29.** | **C** | **30.** | **D** |
| **31.** | **A** | **32.** | **C** | **33.** | **B** | **34.** | **A** | **35.** | **A** |
| **36.** | **A** | **37.** | **B** | **38.** | **A** | | | | |