ACE Engineering Publications
SINCE 2011

GATE | PSUs

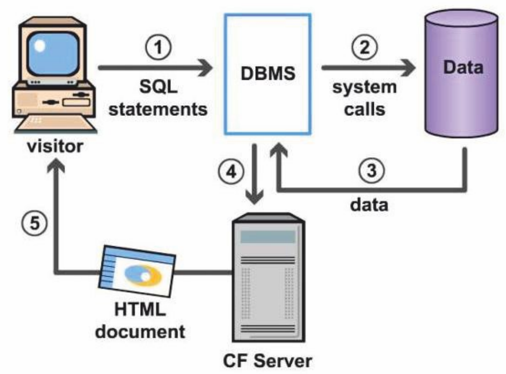

Computer Science & Information Technology

DATA BASE MANAGEMENT SYSTEMS

Text Book:
Theory with worked out Examples and Practice Questions

Database Management Systems

(Solutions for Text Book Practice Questions)

Chapter 2

ER and Relational Model

01. Ans: (b)

Sol: Derived attribute is an attribute that derives its value from one or more attributes.

02. Ans: (c)

Sol: A State might have several Districts and each district belongs to one state.

03. Ans: (a)

Sol: As every specialized entity is subset of generalized entity, then the deletion of generalized entity requires the deletion of specialized entity.

04. Ans: (c)

Sol:

- Composite attribute is an attribute which is composed of other attributes.
- Multi valued attribute represented with double ellipse.
- Derived attributes represented with dotted Ellipse

05. Ans: (c)

Sol:

- Participation of both department and employee is partial, so both options (a) and (d) are incorrect.
- The number on the cardinality specifies the maximum but not the minimum therefore option (b) is incorrect and option (c) is correct.

06. Ans: (b) & (d)

Sol: Based on the participation constraints on Relationship type, double line shows the total participation and single line represents the partial participation.

07. Ans: (a)

Sol:

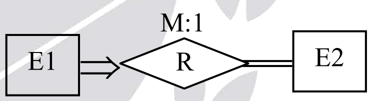

08. Ans: (a) & (b)

Sol: {12, Jessica} is a relation instance of the schema student (id#, emp name)

09. Ans: (b)

Sol: A superkey is one which contains a candidate key.

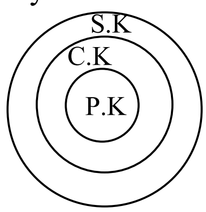

A key which contains a candidate key (primary key) VY then it will be considered as super key.

In options (a), (c) and (d) it contains candidate key (primary key) VY.
But option (b) doesn't contain candidate key (primary key) VY. So it is not super key.

10. Ans: (b)

Sol: All the values present in Foreign key must present in primary key of the referenced relation.

11. Ans: (c)

Sol: S1: A relation schema can have more than one foreign key references to more than one table.

S2: A foreign key in a relation schema R can be used to refer to tuples of R.

∴ Both statements S1 and S2 are false.

12. Ans: (c)

Sol: When parent is update, it requires child table to be updated simultaneously

13. Ans: (c)

Sol: 1. On removal of row (2,4), row (5,2) and (7,2) must also be deleted as they depend on value

2. On removal of row (5,2), row (9,5) must also be deleted as it depends on value 5.

14. Ans: 0

Sol: When <3, 8> is deleted, its related tuples in T_2 is (8, 3) and 3 is set to null. Hence the number of additional tuples to delete is 0

15. Ans: (a)

Sol:

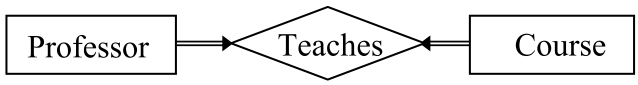

16. Ans: (b)

Sol:

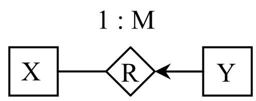

The ER model can be represented with two relations X and YR, therefore Y side include primary key of X side as Foreign key.

17. Ans: (b)

Sol: There are some tuples of course may not participate with any tuple of professor, then cid is a key for the relation.

18. Ans: (a)

Sol: (AR_1B) will be one table as there is total participation and key constraint.
(CR_2) will be the second table as there is a key constraint.

19. Ans: (b)

Sol: Strong entities E_1 and E_2 are represented as separate tables, in addition to that many to many relationship (R_2) must be converted as separate table by having primary key of E_1 and E_2 as foreign key. One to many relationship must be transferred to 'many' side table by having primary key of one side as foreign key. Hence we will have minimum of 3 tables.

20. Ans: (b)

Sol: Strong entities E1 and E2 are converted as separate tables. Since A23 is a multi valued attribute it should also be converted as separate table. Relationship R is transferred to 'm' side (E2).

21. Ans: 3

Sol: E-R model is

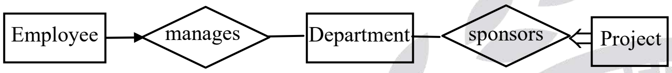

The minimum number of relations in relational model is 3.

1. (Employee, manages)
2. Department
3. (Project, sponsors)

22. Ans: (c)

Sol: From both Employee and Department we have total participation and key constraint possible, Which can be converted into one relation in relational model.

23. Ans: (b)

Sol: M, P are strong entities hence they must be represented by separate tables. M table is modified to include primary key of P side (i.e P1). N is weak entity, and it is modified to include primary key of P (i.e P1).

24. Ans: (a)

Sol: M and P are strong entities hence they must be represented as separate tables. To include R1, M table is modified to accommodate primary key of P side (i.e P1) as foreign key.

N is weak entity, so modify N to accommodate primary key of P (i.e P1) as foreign key.

Therefore tables are (M1, M2, M3, P1), (P1, P2), (N1, N2, P1). So correct answer is (M1, M2, M3, P1).

25. Ans: (c)

Sol:

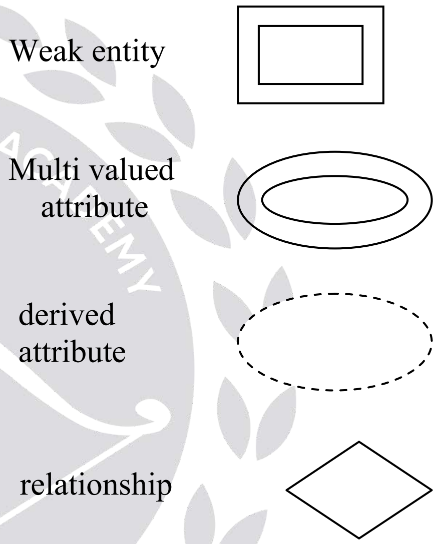

Chapter 3 Functional Dependencies

01. Ans: (a), (c) & (d)

Sol: A particular employee id determines more than one department, hence employee id → department is invalid on the above table.

02. Ans: (b)

Sol: p cannot uniquely determine q, because for the same value of p multiple values of r are obtained. So p → q is false. (a) and (d) are false hence. r → p is also false because unique value of r determines multiple values of p. So (c) is false. (b) is true because for unique values of q, unique value of p is determined and unique value of p determines unique value of r.

03. Ans: (a), (b), (c) & (d)

Sol: For any value of 'a' the dependency A→B is true

04. Ans: (d)

Sol: AC^+ =A, C, B, E, F, G

05. Ans: (c)

Sol: AF^+ = AFDE not ACDEFG as given.

06. Ans: (b)

Sol: (A) closure of {A, C} = {A, C, E}; hence, {A, C} does not determine F.

(B) is true as {A, B, C} is a candidate key.

(C) closure of {A, B, D} = {A, B, D, E, F}; hence, {A, B, D} does not determine C.

07. Ans: (b)

Sol: CD^+ from functional dependencies (FDs) = **C**DE**A**B, it includes RHS attributes AC, so it can be derived from FDs BD^+ from functional dependencies (FDs) = BD only, RHS attributes CD are not included in the closure. Hence it cannot be derived BC^+ from functional dependencies (FDs) = BCDEA, it includes RHS attributes CD, so it can be derived from FDs AC^+ from functional dependencies (FDs) = ACBDE, it includes RHS attributes BC so it can be derived from FDs

08. Ans: (a), (b) & (c)

Sol: (a) P+ = PQR ∴ P → R possible

(b) PS+ = PSQRT ∴ PS → T possible

(c) PS+ = PSQRT ∴ PS→ Q possible

(d) R+ = R ∴ R → T is not possible

09. Ans: (c)

Sol: Every dependency of F can be determined using G and every dependency of F can be determined using F.

10. Ans: (c)

Sol: D→C in set-2 and C→D in set-1 not covered by each other.

11. Ans: (d)

Sol: AB→C, A→BC both can be determined from remaining set of FD's.

12. Ans: 5

Sol: AC → D can be eliminated, it can be derived from A → B and CB → D using augmentation and transitive rule.

$A \rightarrow B \Rightarrow AC \rightarrow BC$

$\Rightarrow AC \rightarrow D$

And remaining FD's are not possible to eliminate

∴ 5 FD's are there in minimal cover.

13. Ans: (d)

Sol: BC→A is inessential as it can be determined from the remaining set of dependencies.

14. Ans: (a)

Sol: As V → W, delete W from VW → X results in V → X

As V → X, delete X from Y → VX results in Y → V

The irreducible set is

V → W

V → X

Y → V

Y → Z

15. Ans: 24

Sol: $2^4 + 2^4 - 2^3 = 24$

16. Ans: (d)

Sol: {A, B}+ = {A, B, C, D, E, F, G} and {B, C}+ = {B, C, D, A, E, F, G}. Other than AB and BC, no other combination of attributes determines all other attributes. Only attributes A, B and C are present in the candidate keys (i.e. AB, BC) of the relation R.

17. Ans: (a)

Sol: Only EC+ contains all attributes of the relation, then EC is key for R.

18. Ans: (b)

Sol: As 'K' is independent attribute, key is ABDK.

19. Ans: (d)

Sol: ABD^+ = A, B, C, D, E.

20. Ans: (b)

Sol: $ACEH^+$ contains all the attributes of R.

21. Ans: (d)

Sol: Closure of $AEH^+ = BEH^+ = DEH^+$ = A, B, C, D, E, H. If any closure includes all attributes of a table then it can become candidate key of the table. Closure of AEH, BEH, DEH includes all attributes of table. Hence they are candidate keys.

22. Ans: (b)

Sol: A^+ = ABCEFGH

B^+ = ABCEFGH

E^+ = ABCEFGH

F^+ = ABCEFGH

All of the above attribute closures contain all attributes of R, except D. Hence the candidate keys are AD, BD, ED and FD. i.e, the number of candidate keys are 4.

23. Ans: 3

Sol: The candidate keys of a relation is: AB, AD and C.

24. Ans: (a), (b) & (c)

Sol: From the given set of functional dependencies, we can see that (Reg No)+ = {Reg No, Name, Address, Phone, Class ID}. Thus, Reg No is a candidate key. According to FD2 and FD4, {Address, Phone} and Class ID, are also the candidate keys.

25. Ans: (c)

Sol: A candidate key always determines any attribute of a relation. A superset of a candidate key is called super key and it can determine all the attributes of a relation

Chapter 4 Normalization

01. Ans: (a)

Sol: Join between the tables of II is returning:

p	q	r
a	1	x
a	1	y
b	1	x
b	1	y

which is not the original table, hence the decomposition is lossy.

02. Ans: (d)

Sol: (1) is lossy as Join operation between the tables does not return the table which is same as STUDENT.

(2) is lossy as there is no common attribute between the decomposed relations.

03. Ans: (b)

Sol: R1= A, B

R2= B, C

R3= B, D

R2 $\cap$ R3 = B and it is key in R2 (B$\to$C).

(R2 $\cup$ R3)$\cap$R1 = (B, C, D) $\cap$ (A, B) = B.

B is a key in (B, C, D) as B$\to$C, C$\to$D. Hence it is lossless join but C$\to$D is not preserved.

04. Ans: (c)

Sol: Decomposition D1

R1(PQST) R2 (PTX)R3(QY) R4(YZW)

R1∩R2 = PT+ = PTXYZW; determining attributes of R2

[(R1joinR2)∩R3] = Q + = QYZW; determining attributes of R3

[(R1joinR2joinR3)∩R4] = Y+ = YZW; determining attributes of R4

∴The decomposition D1 is lossless.

Decomposition D2

R1(PQS) R2(TX) R3(QY) R4(YZW)

(R1joinR3joinR4) ∩R2 = ø

∴ The decomposition D2 is lossy.

05.

Sol: R is in 1NF ∴ decompose to 2NF

A^+ = {A, D, E, I, J} R_1 = 2NF

B^+ = {B, F, G, H} R_2 = 2NF

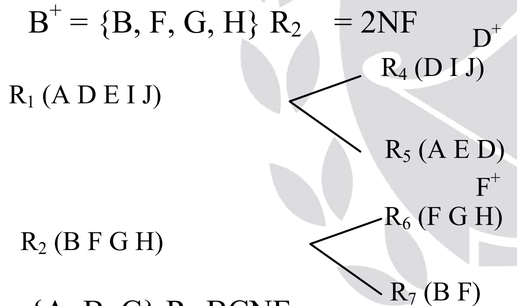

{A, B, C} R_3 BCNF

Then decompose into 2NF

R_1 (ADEIJJ)

R_2 (BFGH)

R_3 (ABC)

3NF also in BCNF

R_3 (ABC)

R_4 (DIJ)

R_5 (AED)

R_6 (FGH)

R_7 (BF)

AB^+ is key.

06.

Sol: Candidate key: AC

A^+ = (ABE) R_1, C^+ = (CD) R_2

(ACF) R_3

07. Ans: (c)

Sol: R is in 1NF as A→FC and B→E are partial dependencies

08.

Sol: (1) C → D

C → A

B → C

C.K: B, 2NF but not 3NF

(2) 2NF but not 3NF as no partial dependency CK: BD.

(3) R is in 3NF but not in BCNF

(4) C.K = A

(5) Candidate Keys = AB, CD, BC, AD

R is in 3NF but not in BCNF.

09. Ans: (a)

Sol: F = {QR → S, R→P, S→Q}

The decomposed relations Y(PR) and Z(QRS) satisfying the dependencies {R→P} and {QR→S, S→Q} respectively.

Relation Y is in BCNF but relation Z is not in BCNF **because** in S→Q ; S is not a super key. All the dependencies of relation X is satisfying on relations Y and Z.

10. Ans: (b)

Sol: S1 is true because it contains one attribute then no partial dependencies possible; hence a relation is always in 2NF.

S2 is false because partial dependencies are not allowed in 2NF and other higher normal forms beyond 3NF also satisfy the conditions required for 2NF.

11. Ans: (c)

Sol: {A, B} = {A, B, C, D}. Hence {AB} is the primary key. A and B are prime attributes. In the FD AB→CD the determinant is a super key. In the FD D→B, the attribute B is a prime attribute. Hence R2 is in 3NF but not in BCNF.

12. Ans: (a)

Sol: Candidate keys of the relation are A, BC and E. As all determinants are keys, the relation is in BCNF.

13. Ans: (b)

Sol: (Volume, Number) → Year is a partial functional dependency. So, the given relation is in 1 NF but not in 2 NF.

14. Ans: (b)

Sol: To simply the process assume

A=name, B=courseNo, C=RollNo, D=grade. Candidate keys are AB, CB. If we select AB as the primary key, then C→ A is allowed in 3 NF (either left side is a key or right hand side is part of the key). A→C is allowed since it is prime attribute (an attribute that is part of any candidate key)

15. Ans: (b)

Sol: From the customer requirement, we can infer FD: Seniority → Salary and is considered as Transitive dependency. Hence the relation is in 2NF but not in 3NF and BCNF.

16. Ans: (a), (b) & (c)

Sol: Every dependency of R satisfies 3NF, that is if X→A is a functional dependency then either X is super key or A is a prime attribute. A relation in 3NF also satisfies 2NF and 1NF.

17. Ans: (d)

Sol: If relation consists only two attributes always it satisfies BCNF (no partial and transitive dependencies). Hence option (a) is correct.

If every key consists only one attribute, then there is no partial dependency hence option (b) is correct.

If there is dependency X → Y, it is allowed in 3 NF if either X is a key or Y is part of some key (prime attribute). Hence option (c) is correct.

If there is dependency X → Y, it is allowed in BCNF if either X is a key or all determinants are keys. Hence option (d) is incorrect.

18. Ans: (c)

Sol: A table is said to be in BCNF if it is already in 3 NF and all determinants are keys.

19. Ans: (a)

Sol: A relation with two attributes is always in BCNF

If all attributes of a relation are prime attributes then the relation is in 3NF

A relation need not to have non-prime attributes.

BCNF Decomposition may not preserves functional dependencies.

20. Ans: (c)

Sol: 2NF deals with partial dependencies.

If $X \rightarrow Y$ allowed in 3NF then either X is superkey **or** Y is part of the key.

Chapter 5 Relational Algebra & Calculus

01. Ans: (b)

Sol: This relational algebra queries will be solved by using the select and project operator from the above mentioned relational schemas. At first it will select all the tuples from the relation whose name is Aditya using select operator then it will select phone number and address of 'Aditya' through the project operator

02. Ans: (a) & (c)

Sol: The common tuples between the two tables comes as output. Hence intersection between the tables has been performed. So option (a) is correct. Again $R1 \cap R2 = R1-(R1-R2)$. Hence option (c) is also correct.

03. Ans: (b) & (d)

Sol: The given operation is equivalent to the intersection of Zone1 and Zone2. Hence, options (b) and (d) are correct.

04. Ans: (a)

Sol: $\Pi_B(r_1) - \Pi_C(r_2) = \phi$ is always true. Because 'B' is foreign key referencing 'C' , so 'C' must be a primary key, 'B' cannot have a value that is not available in 'C'. Hence operation $\Pi_B(r_1)-\Pi_C(r_2)$ is always ϕ.

05. Ans: (a)

Sol: Common column between tables 'R' and 'S' is attribute B. In table 'R' B is primary key (B→A, A→C). In table 'S' B is foreign key so join is performed on attribute B. Therefore maximum tuples possible in the output is equal to rows in Table S (as it has less number of rows, provided B values are not repeated in table 'S').

06. Ans: (a)

Sol: R in r1 (P,Q,R) is foreign key with 2000 tuples references R (primary key) in r2 (R,S,T) with 2500 tuples. So natural matching rows are 2000

07. Ans: (a)

Sol: As relation 'r' need to satisfy both conditions F_1 and F_2, we replace the expression $\sigma_{F_1}\left(\sigma_{F_2}(r)\right)$ with $\sigma_{F_1 \wedge F_2}(r)$

Using A_1 only sufficient in the selection because $A_1 \subset A_2$.

08. Ans: (c)

Sol: 'A' is a column in relation R, then instead of joining all the tuples of R, with S we join only few tuples of R that satisfying the condition <u>A = a with S</u> which is the optimized query.

09. Ans: (b)

Sol: bal < 0 filter rows from account ⋈ depositor from which we can operate on few rows to filter b city = "Agra".

10. Ans: (d)

Sol:

empAge

empno	age
1	20
2	19
3	18

empAge

empno1	age1
1	20
2	19
3	18

The query return

empno
1
2

∴ The query returns Employee number of all employees whose age is not the minimum.

11. Ans: (d)

Sol: Minus operator indicates rows available in LHS table but not in RHS table. In this expression, LHS table produces all female students, RHS table consists students with less marks hence it produces names of all girl students with more marks than all the boy students.

12. Ans: (a)

Sol: Relational algebra expression $\Pi_{P,Q(\sigma(R=3)V(R=5)(S1))}$ will return tuples {(1,2), (2,4)} based on the condition R = 3 and tuples {(3,6),(1,2)} based on the condition R = 5, resulting in tuples {(1,2),(2,4),(3,6)} as the final outcome of the mentioned expression.

Relational algebra expression $\Pi_{T,U(\sigma(V\neq 2)V(V\neq 3)(S2))}$ will return tuple {(1,2)} based on the condition V! = 2 and tuples

{(2,4),(1,2),(3,6)} based on the condition V! = 3, resulting in tuples {(1,2),(2,4),(3,6)} as the final outcome of the mentioned expression.

Hence the complete relation algebra expression will return nothing.

13. Ans: (a) & (c)

Sol: As per the rules of division operation.

14. Ans: (b)

Sol: The result will have Q attribute values from S1 which are associated with all the pairs of P, R values present in S2 and are present in S1 as triples, hence outcome will have a relation with scheme (Q) and tuples {(2), (4)}. Note that 6 does not have the above property as only (1, 6, 2) is present in S1 but (2, 6, 4) is not present in S1.

15. Ans: 4

Sol: The output of T_1 is: courseName

CA
CB
CC

the output of T_2 is: StudentName

SA
SC
SD
SF

16. Ans: (a) & (b)

Sol: Π_{eId} ($\Pi_{eId,bId}$ (Own)/ Π_{bId} (Brand))

Π_{eId} (Own) − Π_{eId} ((Π_{eId} (Own)× Π_{bId} (Brand)) −$\Pi_{eId,bId}$ (Own))

17. Ans: (c)

Sol: The 'Λ' operator in tuple calculus will have same effect as the '$\cap$' intersection operator in relational algebra.

18. Ans: (d)

Sol: Retrieve rollNo of students who have enrolled for exactly one course

19. Ans: (c)

Sol: Here, an existentially bound tuple variable e1 is used to make sure that if roll number and name of a student are returned by the query then the student definitely had enrolled for some course.

20. Ans: (c)

Sol: In negative queries TRC produces infinite results hence it is not considered as safe.

21. Ans: 2

Sol: Relational calculus eliminate the duplicates.

{T/∃B∈ Book (T.Title = B.Title)}

22. Ans: (a)

Sol: SQL, Relational algebra, tuple relational calculus and Domain relational calculus all is representing the same. i.e., all these expressions representing to find the distinct names of all students who score more than 90% in the course numbered 107.

Chapter 6

Structured Query Language (SQL)

01. Ans: (d)

Sol: To search for a specified pattern in a column we use LIKE operator in the WHERE clause. The % sign represent zero or multiple characters.

Option (a) displays the name and loan number of customers whose names are exactly 'A'.

Option (c) displays the name and loan number of customers whose names ends with 'A'.

Option (b) gives error message, invalid relational operator for the word LIKES.

Since option (d) satisfies the requirements, it is correct.

02. Ans: (a), (b)

Sol: Option (c) is incorrect, since option (c) finds the students whose registration number is either 1 or 10 which also includes the tuple (1,ALEX, CIVIL).

03. Ans: (b)

Sol: The result of the query is

A	B	C
1	4	4
2	5	1
2	5	3
3	5	2

04. Ans: (d)

Sol: The query finding the tuples in r_1 but not in r_2

05. Ans: (a)

Sol: Table after the update query is

Rollno	Marks
456	70
457	53
458	63
459	null

Output of the average query is (70+53+63)/3 = 62

06. Ans: (c)

Sol: sum (rating)/count(0) is smaller value than avg(rating).

07. Ans: (c)

Sol: Union operator eliminates the duplicates.

08. Ans: (a)

Sol: All the three queries return the same results.

09. Ans: (b)

Sol: Select clause contains either aggregate function or the attributes that appear in group by clause.

10. Ans: (b)

Sol: The having clause selects only those names, which are duplicated. Therefore the guests is finding duplicate names of employees.

11. Ans: 2

Sol: It returns two rows.

Student – Name	**Sum(P.Marks)**
Raj	4
Rohit	2

12. Ans: 3

Sol: The query finds the details of the staff of age below 40 and with at least one more staff of opposite gender having same supervisor.

The query find the details of the staff with name Alice, Mathew and Jane

13. Ans: (d)

Sol: R(A,B,C) will be

A	B	C
4	8	null
3	6	3
2	4	null
null	3	6
null	2	4

14. Ans: (c)

Sol:

R_1

A	B
1	5
3	7

R_2

A	C
1	7
4	9

$R_1 \bowtie R_2$

A	B	C
1	5	7
3	7	null
4	null	9

15. Ans: (d)

Sol:

- Inner join returns the rows that have matching rows of both the relation.
- Left outer join returns all the rows from left side relation even if there is no matching row in the right side relation.
- Right outer join returns all the rows from right side relation even if there is no matching row in the left side relation.
- Full outer join returns all the rows from both the relation even if there is no matching row in the other relation.
- Query 4 returns a result, which is superset of Query1, Query2 and Query3.

16. Ans: (b)

Sol: All requires the given condition to be true for all values in the set.

ANY requires the given condition to be true for atleast one value. Therefore Q1 is subset of Q2.

17. Ans: (b)

Sol: The condition is B>any (1,2,1,3,2,4) and the output will be 4.

18. Ans: (a)

Sol: '=any' operator is same as 'in' operator

19. Ans: (a)

Sol: The inner query returns all values of capacity and P1.capacity is true only for the maximum capacity.

20. Ans: (b)

Sol: The ALL keyword specifies that the search condition is TRUE if the comparison is TRUE for every value that the sub query returns. If the sub query returns no value, the condition is TRUE.

21. Ans: 7

Sol: The output of the query is

ta.player
Klose
Ronaldo
G muller
Fontaine
Pele
Klismann
Kocsis

22. Ans: (c)

Sol: The given query is a correlated query.
For each row in outer query, all rows in inner query are processed. [like nested for loop]

Inner Query:
For each employee [each employee is listed in outer query], "It counts the number of employees whose age is less than that employee".
For the youngest employee, this count is "0"
For the second youngest employee, this count is "1"
Like wise, for the seventh youngest employee this count is "6"
For 8th youngest [& all others who are elder than him], this is not less than 7.
So, they don't satisfy, inner query condition.
So, the answer is Emp_ids of the seven youngest employees.

23. Ans: 2

Sol: The query finds name of those passengers whose age is above 65 and has some reservation for 'AC' class.
The output of the query is:

pname
Rohan
Anil

24. Ans: (c)

Sol: In the given query AND operator is mentioned so both the condition should be matched. Hence the above Query returns the number of employees who have distinct salaries.

25. Ans: 2

Sol: The given query is correlated subquery and for each row of the outer query inner query is going to be executed once.
first inner query is independent of outer query and the output is {c11, c12}.
Second inner query is dependent on the row of outer query, for the student <S01, James, D01> it returns {C11, C12}. The except operation between two inner queries return empty set. NOT EXISTS (empty set) returns true and student <S01, James, D01> in the final output.
The same reason applicable for student <S05, milli, D02>
final output of the query is

Sno	Sname	Dno
S01	James	D01
S05	milli	D02

The number of rows returned as the result of the above query is '2'.

26. Ans: 2

Sol:

Total		Total – avg
name	**capacity**	**capacity**
Ajmeer	20	25
Bikaner	40	
Churu	30	
Dungargarh	10	

The result of the query is:

name
Bikaner
Churu

Chapter 7

Transactions & Concurrency Control

01. Ans: (d)

Sol: A: Atomicity C: Consistency
I: Isolation D: Durability

02. Ans: (b)

Sol: The data base system must be consistent before and after the transaction.

03. Ans: (b)

Sol: The number of serial schedules are 2

The number of concurrent schedules are $= \frac{(5+3)!}{5!*3!} = 56$

Then, the total number of non serial schedules are = (number of concurrent schedules – number of serial schedules)
$= 56 - 2 = 54$

04. Ans: (d)

Sol: Transaction T_3 perform read on A, which is updated by T_1 and committed before T_1 does.

05. Ans: (a)

Sol: As $R_3(x)$ is dirty operations which read $W_1(x)$ and is committed before T_1. Hence schedule is non-recoverable.

06. Ans: (c)

Sol: A recoverable schedule is one where for each pair of transactions T_i and T_j such that T_j reads a data item previously written by T_i, the commit operation of T_i appear before the read operation of T_j.

07. Ans: (a), (b)

Sol: T2 performs dirty read on T3, so T2 should commit after T3 is committed (or) T3 commits before T2

08. Ans: (c)

Sol: A schedule is said to be strict if a value written by a transaction T is to be read or written by another transaction until either T commits or aborts.

09. Ans: (a), (c)

Sol:

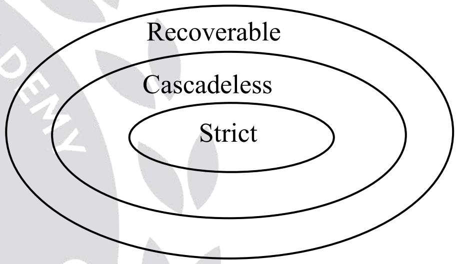

10. Ans: (a) & (b)

Sol: If we rewrite the read and write operations serially in a table format and draw the precedence graph, the precedence graph does not have any cycle. So, the schedule is conflict serializable.

In a recoverable schedule, for each pair of Transactions Ti and Tj such that Tj reads data item previously written by Ti, the commit operation of Ti appears before the commit operation Tj. This is true in the given schedule.

In the given schedule, T4 reads data items X and Y which are previously written by T1 and T2. But the commit operation of T1 appears after the read operation of T4. So, the schedule is not cascadeless and not Strict.

11. Ans: (a)

Sol: If a schedule is serializable, the topological order of a graph (precedence graph) yields a serial schedule.

12. Ans: (a)

Sol: The precedence graph for the given schedule is

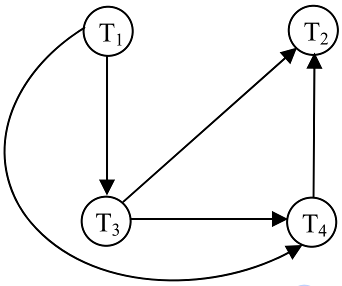

The topological sort of the graph is $T_1–T_3–T_4–T_2$

13. Ans: (a)

Sol: S1:

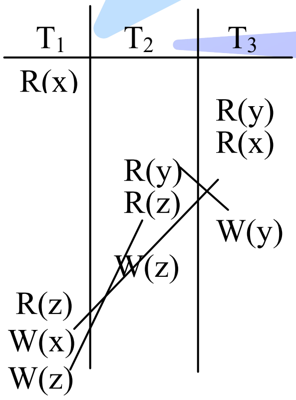

Precedence graph

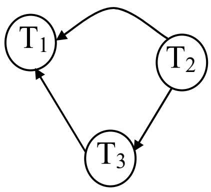

S1 is conflict serializable to $T_2 \rightarrow T_3 \rightarrow T_1$

S2:

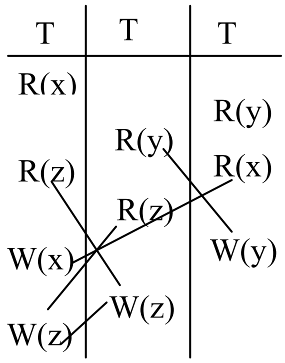

Precedence graph

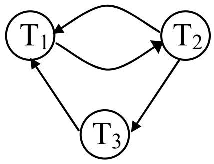

S2 is not conflict serializable

14. Ans: (d)

Sol:

T_1	T_2	T_3
R(A)		
	W(A)	
		R(A)
W(A)		
		W(A)

S1 and S2 are conflict equivalent to serial schedule T_2, T_3, T_1.

S3 is not conflict equivalent as 2RA, 3WA ($T_2<T_3$) and 3WA, 2WA ($T_3<T_2$) are the conflict operations. There is no serial schedule that satisfies both $T_2<T_3$ and $T_3<T_2$.

15. Ans: (a)

Sol: The two schedules are said to be conflict equivalent if all the conflicts in both the schedules are same.

All the conflicts in the given schedule are same in the schedule of option (a).

16. Ans: (c)

Sol: Precedence graph is

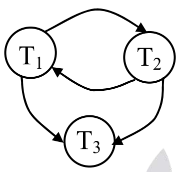

17. Ans: (a)

Sol: Draw the precedence graph of the transactions and observe that the graph has a cycle. So, the above schedule is not a conflict serializable schedule.

Next, we have to check whether the schedule is view serializable and the conflict operations should follow the time stamp order T2-T3-T1 and the order is violating. Therefore the schedule is neither view nor conflict serializable schedule.

18. Ans : (c)

Sol: To perform W(y), T2 acquires Exclusive lock on y, which will be released only after commit in strict 2pl. Then request for T1 is processed only after commit of T2.

19. Ans: (c)

Sol: In strict 2 PL, all excusive locks (write locks) must be released after commit statements. Hence Option (c) is correct.

20. Ans: (b) & (c)

Sol: Wound-Wait Deadlock Prevention Scheme:

When TA1 requests data item held by TA2(older means smaller timestamp), two cases may arise.

- If TA1 older than TA2 then TA1 wounds TA2(TA2 will be aborted).
- If TA1 younger than TA2 then TA1 will wait to release the data item held by TA2.

In this case Transaction T_1 is older and T_3 is younger in respect of Transaction T_2. Hence options (b) and (c) are correct.

21. Ans: (c)

Sol: In wait-die deadlock prevention strategy and older transaction will wait for younger transaction to release the lock where as an younger transaction aborts if requesting a lock held by an older one.

In wound-wait deadlock prevention strategy a younger transactions need to wait for an older transaction to release the lock but an older transition requesting a lock held by an younger one preempts younger transaction to abort.

22. Ans: (d)

Sol: Figure1 is a wait-for graph without a cycle, so it doesn't shows possibility of deadlock, while Figure2 wait-for graph has cycle, so it shows possibility of deadlock.

23. Ans: (b)

Sol: 2 PL is based on locks and hence not free from deadlock but ensures conflict serializability. Timestamp ordering protocol is based on timestamps ensures conflict serializability and also free from deadlock.

24. Ans: (b)

Sol: The order of TimeStamp is $T_1 \rightarrow T_3 \rightarrow T_2$.

The timestamp ordering requires that all the conflicts to be processed in the order of their timestamps.

$W_2(B) - W_3(B)$ is violating time stamp order and the schedule is not possible under timestamp protocol, But allowed under Thomas Write Rule which ignores $W_3(B)$ called Obsolete Write.

25. Ans: (d)

Sol: When T_2 performs $R_2(A)$ and $TS(T_2)<W\text{-}TS(A)$ then $R_2(A)$ is rejected and rolled back. Therefore the above schedule is not allowed in both Basic timestamp protocol and Thomas write rule.

Chapter 8 Recovery Management System

01. Ans: (b)

Sol: Before T_i executes write (X) operation, a log record $< T_i, X, V_1, V_2 >$ is written where V_1 is the value of X before the write (the old value) and V_2 is the value to be written to X (the new value).

02. Ans: (a) & (b)

Sol: The immediate modification scheme allows updates of an uncommitted transaction to be made to the buffer, or the disk itself, before transaction commits.

03. Ans: (b)

Sol: Here, T0 and T1 both are committed but T2 is not committed at the time of crash.

So, undo (T2) will be done first and the value of A is set to 300.

Transactions T0 and T1 need to be redone because log contains both the start and commit for both of them.

04. Ans: (c)

Sol: In the immediate database modification scheme, during recovery after a crash, a transaction needs to be redone if and only if both $(T_i$, start), $(T_i$, commit) are So, option (c) is correct.

05. Ans: (d)

Sol: In the deferred database modification scheme all modifications are recorded to the log, but all the writes are done after partial commit.

During recovery after a crash, a transaction needs to be redone if and only if both (T_i start) and (T_i commit) are present in the log.

In this scheme there is no need of any undo operation because all write operations are deferred until partial commit of that transaction.

If the system crash after step 6 then only redo of T1 is required but no action is required for T0 because partial commit of T0 is not done.

06. Ans: (a)

Sol: In the deferred database modification scheme, all modifications are recorded in the log, but all the writes are done after partial commit.

During recovery after a crash, a transaction needs to be redone if and only if both (T_i, start) and (T_i commit) are present in the log.

In this scheme, there is no need of any undo operation because all write operations are deferred until partial commit of that transaction.

So, option (a) is correct.

07. Ans: (d)

Sol: As per the process of transaction recovery.

08. Ans: (c)

Sol: As per the process of transaction recovery

09. Ans: (c)

Sol: In order to recover the database even on multiple crashes the same undo and redo list will be used.

10. Ans: (b)

Sol: In log based recovery we must perform Redo operation for those transactions that contains both <start> and <commit> log record.

We perform Undo operation for those transactions that contains only <start> but not <commit> log record.

Therefore we perform Redo of T_1 and Undo of T_2.

Chapter 9

Indexing

01. Ans: (a)

Sol: Cluster index is created on a column with cluster of values (like year, age, branch where group of students fall under one group), so it is non-key but requires ordering.

02. Ans: (c)

Sol: **Primary index:** in a sequentially ordered file, the index whose search key specifies the sequential order of the file.

Secondary index: an index whose search key specifies an order different from the sequential order of the file.

Dense index : Index record appears for every search-key value in the file.

Sparse Index: contains index records for only some search-key values.

The indexing shown holds index records for some search key values and not all, hence it is sparse indexing. Moreover indexing on search key mentions sequential order of file. So it is primary indexing. Hence option (c) is correct.

03. Ans: (c)

Sol: A clustering index as the name suggests is created when the data can be grouped in the form of clusters of non-key and order. So, option (c) is correct.

04. Ans: (d)

Sol: Number of blocks needed to store data file with 1000 records = 1000/5 = 200

Number of blocks needed to store dense file index = 1000/15= 67

Total blocks required = 200 + 67 = 267 blocks.

So, option (d) is correct.

05. Ans: 1286

Sol: Number of data file blocks =

$$\left\lceil \frac{\#\text{Records}}{\text{Data File Blocking Factor}} \right\rceil = \left\lceil \frac{1310720}{10} \right\rceil$$

= 131072

Number of index file blocks =

$$\left\lceil \frac{\#\text{Index Entries}}{\text{Index File Blocking Factor}} \right\rceil = \left\lceil \frac{131072}{102} \right\rceil$$

= 1286

06. Ans: 698

Sol: Size of index record = 12 + 7 = 19

No. of index records = 1, 50, 000

$$\text{Blocking factor } = \frac{4096}{15} = 215 \text{ record/block}$$

$$\text{No. of index blocks} = \frac{1,50,000}{215}$$

= 698.

07. Ans: (c)

Sol: Since block size is 2^{10} (=1024 bytes), $2^m \times 2^n = 2^{10}$. Only option (c) satisfies with m = 8 and n = 2.

08. Ans: 3

Sol: Blocking Factor= 1024/100 =10

No. of data blocks = 30000/10 =3000

Blocking factor for the index file =1024/ 15 = 68

No. of first level blocks = 3000/68 = 45

No. of second level blocks = 45/68 = 1

Total number of blocks required to access a record by searching multilevel index = 2 + 1 = 3

09. Ans: (c)

Sol: $n*5+(n-1)*(10+8) \leq 512$

$5n+16n-18 \leq 512$

$23n \leq 530$

$n \leq 23.$

10. Ans: (c)

Sol: The order of the B-Tree is

$n* 6 + (n-1) * (9+7) \leq 512$

$6n + 16n - 16 \leq 512$

$22n \leq 528$

$n \leq 24$

- Root has 1 node, 24 pointers and 23 index records
- Level 1 has 24 nodes, 576 pointer and 552 records
- Level 2 has 576 nodes, 13,824 pointer and 13248 records

11. Ans: (a)

Sol: Insertion of new key is leading to insertion of new node at all 4 levels, in turn it is leading to insertion of new root node. Hence the maximum number of nodes that could be created are 5

12. Ans: (c) & (d)

Sol: The tree after key 40 inserted

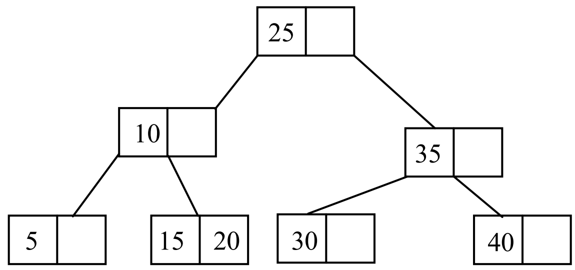

13. Ans: (B)

Sol: If the key is deleted from the leaf node, then borrow the keys from its sibling node.

14. Ans: 5

Sol: The nodes to access all records with a "search key greater than or equal to 7 and less than 15" is (9), (5), (5, 7) (9, 11) and (13, 15).

15. Ans: (b)

Sol:
- B^+ Tree is a **height** balanced search tree
- non leaf nodes have pointers to the next level nodes but not to the data records
- All the leaf nodes are connected with a pointer $\mathbf{P_{next}}$.
- All the key values in each node are kept in sorted order.

16. Ans: 52

Sol: Key = 8, Block size = 512,

Block pointer = 2 bytes, the order of B^+ tree is maximum number of block pointers in it. (Let 'n')

$n * 2 + (n - 1) 8 \leq 512$

$2n + 8n - 8 \leq 512$

$10n \leq 520$

$n \leq 52$

17. Ans: (a)

Sol: $n * (k+P_r) + P \leq B$,

(Where n is order of the tree, P is block pointer, k is key value and B is block size).

$n*(9 + 7) + 6 \leq 1024$

$16n \leq 1018$

$N \leq 63$

18. Ans: 50

Sol: Order of non-leaf node is

$(n\times8) + (n-1)\ 12 \leq 1024$

$8n + 12n - 12 \leq 1024$

$20n \leq 1036$

$n \leq 51$

maximum number of keys possible is :50

19. Ans: (b)

Sol: We begin the traversal from the root and identify the leaf node (20, 22) as the possible node to insert. But this node is full because the order of leaf is 2. So we split this node to insert 15. The resulting nodes are (15, 20) and (22). So the number of leaf nodes is incremented by 1. The element 20 is moved to the internal node (22). The pair (20, 22) is the modified internal node. So the number of internal nodes doesn't change.

20. Ans: (d)

Sol: Deleting '10' from internal node requires 10 to be replaced with copy of 13.

21. Ans: (a)

Sol: Insert 15

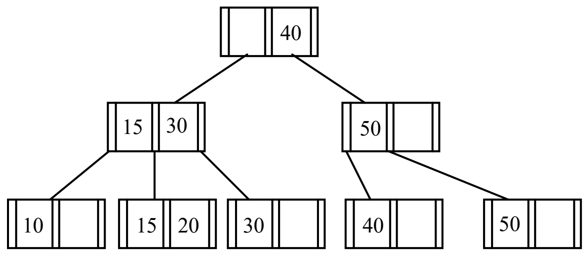

Insert 25

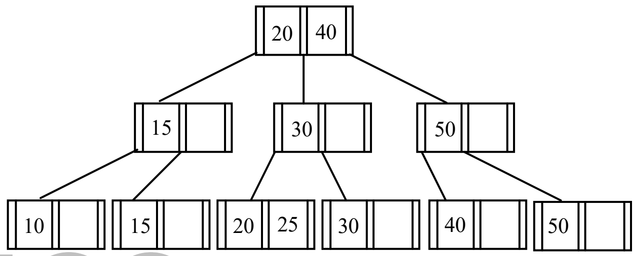

22. Ans: (a)

Sol: If we remove K50 in index node, there is no change in height as still root, index and leaf nodes exist.

Leaf nodes are now, after Delete 50, the B^+ is:

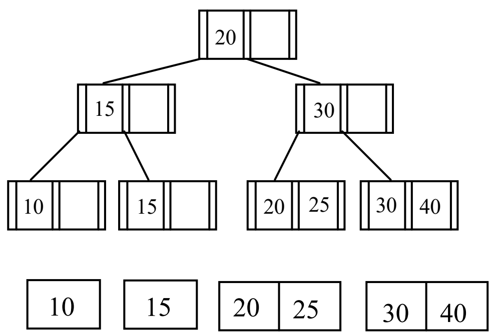

Hence Root now consists: 20

(i) is true, (ii) is true but (iii) is not true