CBCS SCHEME

USN | 2 | V | D | 2 | 3 | C | I | 0 | 0 | 7 | BCS403

Fourth Semester B.E./B.Tech. Degree Examination, June/July 2025
Database Management Systems

Time: 3 hrs. Max. Marks: 100

Note: 1. Answer any FIVE full questions, choosing ONE full question from each module.
2. M : Marks , L: Bloom's level , C: Course outcomes.

		Module – 1	M	L	C
Q.1	a.	Explain the types of attributes with example.	4	L2	CO1
	b.	Define database. Explain the main characteristics of the database approach	8	L2	CO1
	c.	Show the ER diagram for an EMPLOYEE database by assuming your own entities (minimum 4) attributes and relationships, mention cardinality ratios wherever appropriate.	8	L3	CO2
		OR			
Q.2	a.	Describe the three schema architecture.	4	L2	CO1
	b.	Explain the component models of DBMS and their interaction with the help of diagram.	8	L2	CO1
	c.	Design ER diagram for a university database by assuming your own entities (4). Mention primary key , constraints and relationships.	8	L3	CO2
		Module – 2			
Q.3	a.	Explain relational model constraints.	6	L2	CO1
	b.	Explain the characteristics of relations with suitable example for each.	6	L2	CO1
	c.	Considering the following schema : Sailors (sid , sname , rating , age) Boats (bid , bname , color) Reserves (sid , bid , day) Write a relational algebra queries for the following : i) Find the names of sailors, who have reserved red and a green boat. ii) Find the names of sailors who have reserved a red boat. iii) Find the names of sailors who have reserved a red or green boat. iv) Find the names of sailors who have reserved all boats.	8	L3	CO1
		OR			
Q.4	a.	Explain the steps to convert the basic ER model to relational Database schema.	6	L2	CO1
	b.	Explain Unary relational operations with example.	6	L2	CO1

	c.	Consider the relation schema Employee database. EMPLOYEE (Fname ,Minit , Lname , SSn , Bdates , Address , Sex , Salary Super_SSn , Dno) DEPARTMENT (Dname , Dnumber , Mgr_SSn , Mgr_start_date) PROJECT (Pname , PNumber , Plocation , Dnum) WORKS_ON (Essn , Pno , Hours) DEPENDENT (Essn , Dependent_name , sex, Bdate , Relationship) Write relational algebra queries for the following . i) Retrieve the name and address of all employees who work for the 'Research' department. ii) List the names of all employees with 2 or more dependents. iii) Find the names of employees who work on all the projects controlled by department number 5. iv) List the names of employees who have no dependents.	8	L3	CO3
		Module – 3			
Q.5	a.	What is the need for normalization? Explain second and third normal form with examples.	6	L2	CO4
	b.	Outline constraints in SQL.	6	L2	CO1
	c.	Identify the given Relation R(ABCDE) and its instance, check whether FDS given hold or not. Give reasons i) $A \rightarrow B$ ii) $B \rightarrow C$ iii) $D \rightarrow E$ iv) $CD \rightarrow E$. A B C D E a_1 b_1 c_1 d_1 e_1 a_1 b_2 c_1 d_1 e_1 a_2 b_2 c_1 d_2 e_3 a_2 b_3 c_3 d_2 e_2	8	L3	CO4
		OR			
Q.6	a.	What is Multivalued dependency? Explain 4NF and 5NF with suitable example.	6	L2	CO4
	b.	Outline the informal design guidelines for relational schema.	6	L2	CO4
	c.	Consider relation R with following function dependency : EMPPROJ (SSn , Pnumber , Hours , Ename , Pname , Plocation) SSN , Pnumber → Hours, SSN → Ename Pnumber → Pname , Plocation. Is it 2NF? Verify? If no give reason.	8	L3	CO4

		Module – 4			
Q.7	**a.**	Consider the following schema for a company database : Employee (FName , LName , SSn , Adderss , Sex , Salary , Dno , Super_SSn) Department (Dname , Dnumber , mgr_SSn, mgr_st_date) Project (Pname , Pnumber , Plocation , Dnum) WORKS_on (Essn , Pno , Hours) DEPENDENT (Essn , Dependent name , Sex , Bdate, relationship) Write the SQL queries for the following : i) List the names of managers who have atleast one dependent (use correlated nested). ii) Retrieve the name of each employee who has a dependent with the same first name and is the same sex as the employee. iii) For each project retrieve the project number , project name and the number of employees who work on that project. iv) Retrieve the SSN of all employees who work on project number 1, 2 or 3. (Use IN). v) Find the sum of the salaries of all employees of the 'Research' department as well as maximum salary , minimum salary , average salary in this department.	**10**	**L3**	**CO3**
	b.	Why concurrency control is needed? Demonstrate with an example.	**10**	**L2**	**CO5**
		OR			
Q.8	**a.**	Consider the following schedule. The actions are listed in the order they are scheduled and prefixed with the transaction name. S1 : T1 : R(X) , T2 : R(X) T1 : W(Y) , T2 : W(Y) , T1 : R(Y) , T2 : R(Y) S2 : T3 : W(X) , T1 : R(X) , T1 : W(Y) , T2 : R(Z) , T2 : W(Z) , T3 : R(Z) For each schedule answer the following : i) What is the precedence graph for the schedule? ii) Is the schedule conflict_serializable? If so what are all the conflicts equivalent serial schedules? iii) Is the schedule view serializable? If so what are all the view equivalent serial schedules?	**10**	**L3**	**CO5**
	b.	Explain triggers with example write a trigger in SQL to call a procedure "Inform_Supervisor" whenever an employees salary is greater than the salary of his or her direct supervisor in the COMPANY database.	**10**	**L3**	**CO5**
		Module – 5			
Q.9	**a.**	Describe the two – phase locking protocol for concurrency control provide example to illustrate how it ensures serializability in transaction schedule.	**10**	**L2**	**CO5**
	b.	Explain the characteristics of NOSQL system.	**10**	**L2**	**CO6**
		OR			
Q.10	**a.**	Explain binary locks and shared lock with algorithm.	**10**	**L2**	**CO5**
	b.	Explain MongoDB data model, CRUD operations and distributed system characteristics.	**10**	**L2**	**CO6**

* * * * *

Q.1 b) Define database. Explain the main characteristics of database approach.

Soln. A database is a collection of related data.
The main characteristics of the database application

1) Self describing nature of the data base system.
→ data base system stores both data and its complete definition.
→ DBMS catalog contain structure of each file, name of the data item, data types, storage formats, size position of each data item, constraints on the data item.

2) Insulation between programs and data and data-abstraction.
→ DBMS based system store the structure of data files seperately in the DBMS catalog.
→ This seperation between program and data definition is called program data independent.

3) Support multiple views of the data.
→ A data base system typically stores multiple types of users with different roles and requirments
→ Each user may need a different perspective of the same database.
→ A view is a customization representation of the database to a specific user or application.

4) Sharing data and multi user transaction
→ Multiple users can access the same database simultaniously
→ A DBMS include concurrency control mechanism to manage simultaneous access to a shared data.

Q.1 c] show a ER diagram for an EMPLOYEE database by assuming own entities, attributes, relationships, mention cardinality ratios, wherever appropriate.

Q1

a) Explain the types of attribute with example.

① Composite attributes are those which can be divided into smaller subparts, which represent more basic attribut with independent meaning.

Ex: Address
- street_address → Number, street, Apartment number.
- City
- State
-

② Simple attributes: are attributes that are drawn from an atomic value or attribute that are not divisible.

Ex: Apartment number, street.

③ Single valued attributes: which can take more than one value for a given entity from an entity set.

Ex: A person who can speak in more than one language.

④ Multivalued attributes: which can take only one value for a given entity from an entity set

Ex: Age.

⑤ Derived Attribute: which can be derived from other attribute

Ex: Age of student is derived from date of birth.

⑥ Key attributes: which can identify an entity uniquely in a set.

Ex: Roll no. of a student.

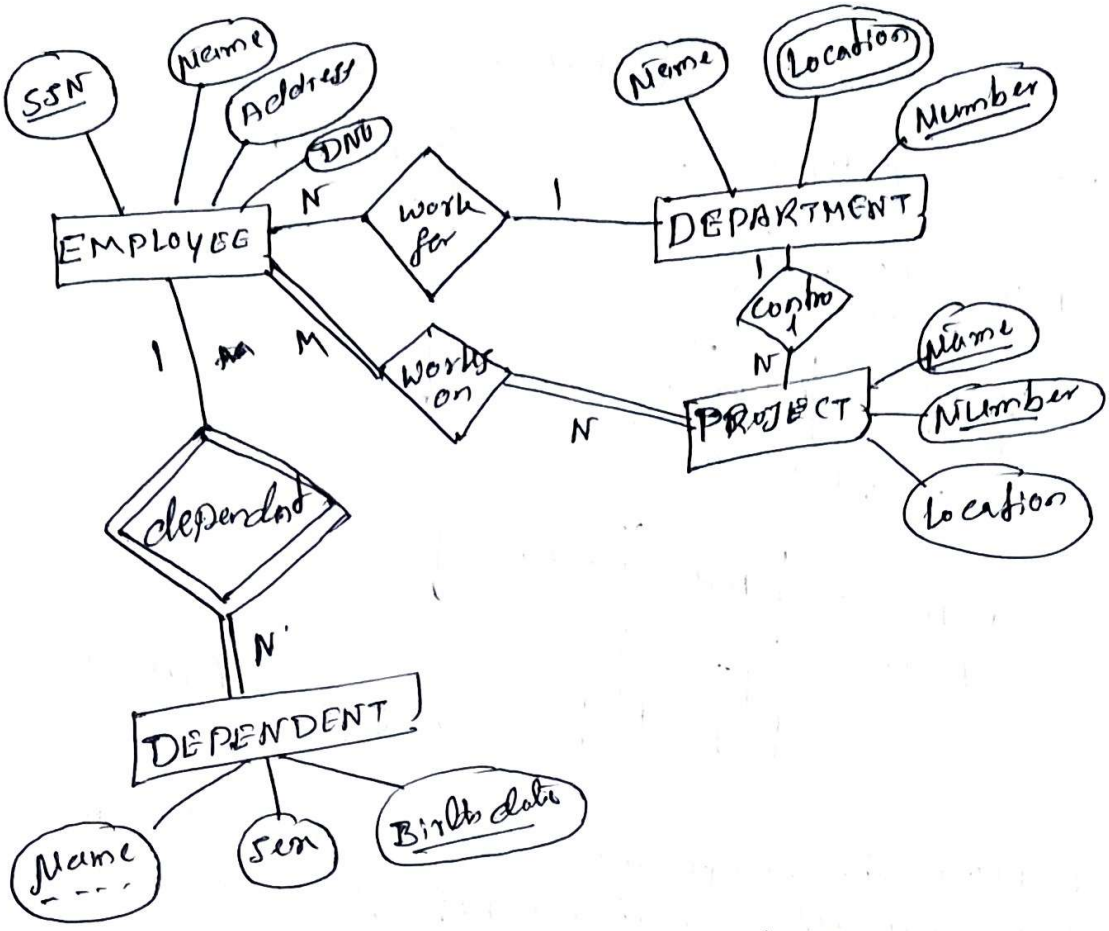

ER diagram of employee database

Q. 2 a) Describe the three schema architecture.

Soln The three schema architecture can be defined at the following three levels.

(1) The internal level has an internal schema, which describe the physical storage structure of the database.

(2) The conceptual level has a conceptual schema which describe the structure of the whole database for a community of users.
It hides the details of physical storage structure and describing entities, data types, relationships, user constraints.

(3) External view or view level include a number of external schemas or user views. Each external schema describes a part of the database, that a particular user group is interested in & hide the rest of the database from the users.

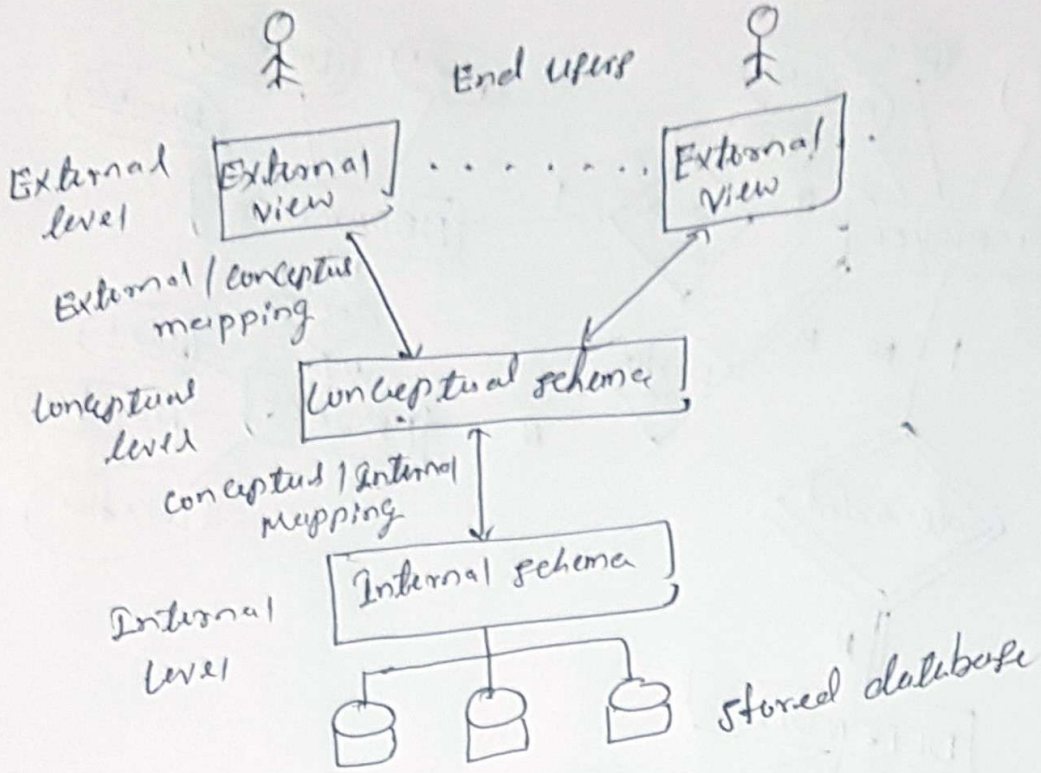

Q 2. (b) Explain the component modules of DBMS by assuming and their interaction with the help of diagram.

Soln → A DBMS a collection of software modules that manage database creation, access and maintainence.

→ DBMS components are divided into two major parts.
  → User level components
  → Internal DBMS components

→ DBA controls and manage the entire database system that define database schema, storage structures, constraints access rights.

→ Casual users access the database using interactive query interfaces.

→ Application programmers write database programs using host languages ~~such~~ such as Java, C, Python.

→ Parametric user execute predefined or command transaction

→ DDL compiler processes database schema definition.

→ The query compiler checks SQL query syntax and validates attribute and table names.

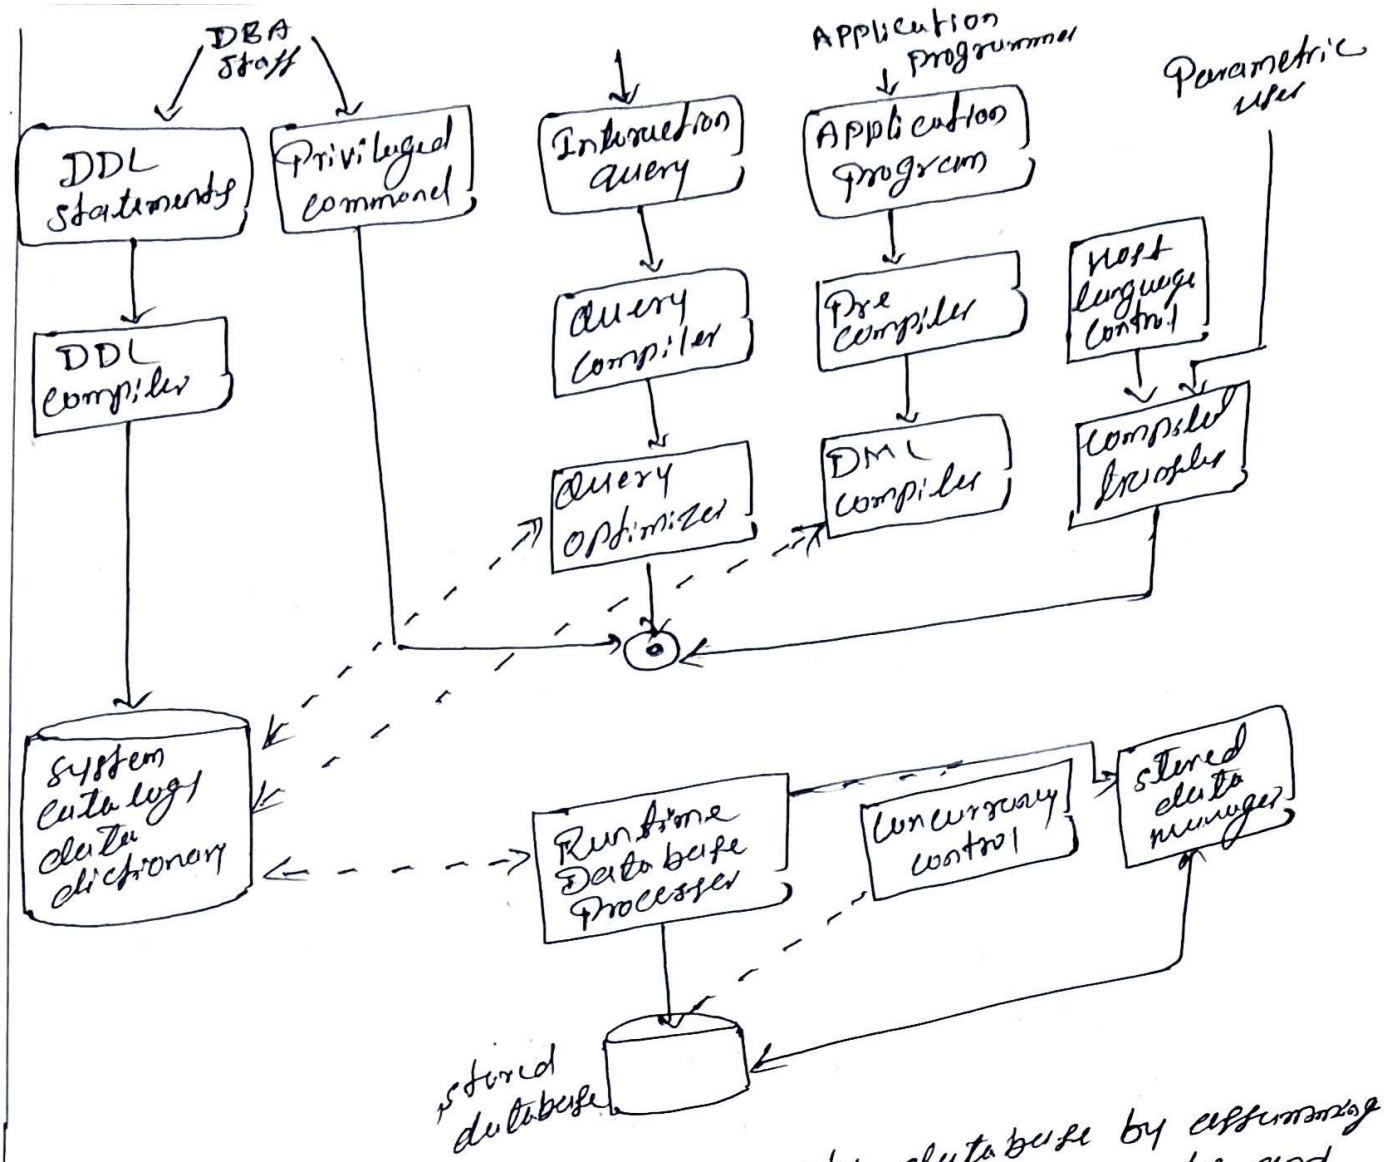

Q 2.

c) Design ER diagram for a university database by assuming your own entities, mention primary key, constraints and Relationships.

Sol.

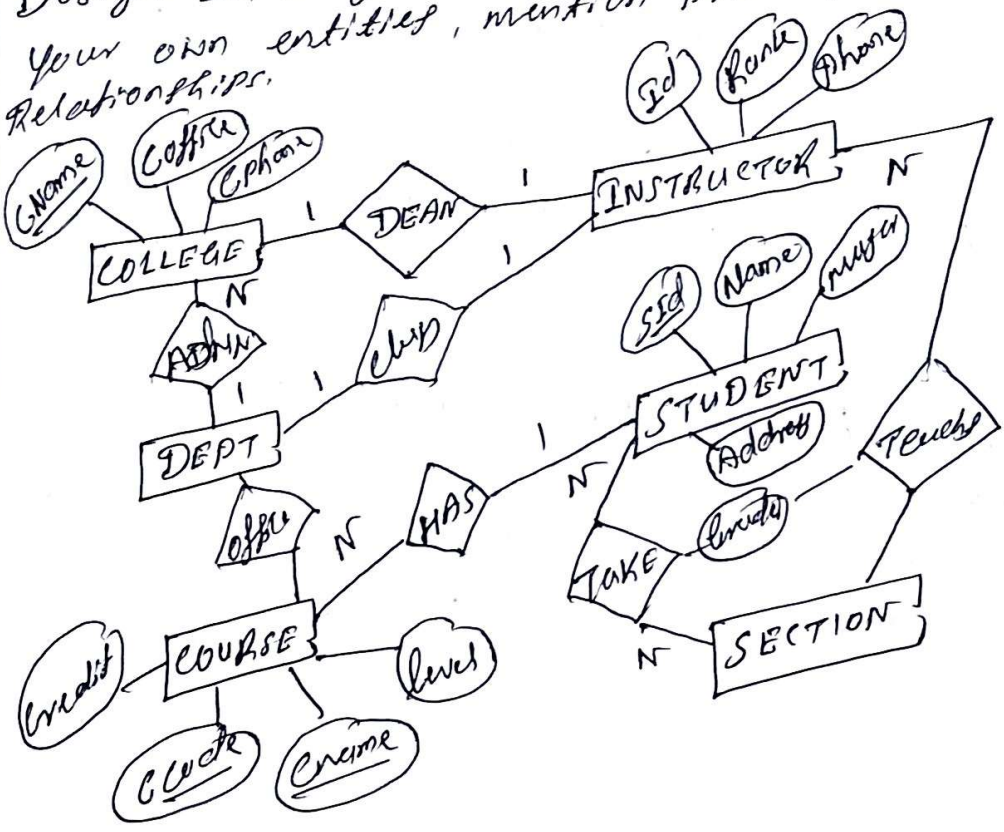

Q3 a) soln

Explain Relational model constraints.

constraints on database can generally be divided into three main categories.

(1) constraints that are inherent in the data model called implicit constraints.

(2) constraints that can be directly expressed in the schemas of the data model called data definition language called explicit constraints.

(3) constraints that can't be directly expressed in the schema of the data model, called Application based or constraint rules.

* Domain constraints specify that each attribute value must be an atomic value from its predefined domain.
It defines data types, value range, enumerated values, special types.

* Key constraints and constraints on null values.

→ A relation is called as a set of tuples, All elements are distinct. No two tuples have the same combination of values for all there attributes.

→ super key defines uniqueness constraints that no two distinct tuples in any state r of R can have the same value for SK.
i.e $t_1[SK] \neq t_2[SK]$

→ A key is a minimal super key with additional property that no two tuple share the same key value and disjoining any attribute remains uniqueness

→ A relation schema may have more than one key. each key is called a candidate key.

→ one of the candidate key is designated as primary key. of the relation.

→ other candidate keys are designated as unique keys

b) Explain the characteristics of relationships with suitable example for each.

Soln

(a) ordering of tuples in a relation: A relation is defined as a set of tuples.
Elements in a relation do not have any particular order.

(b) ordering of values within a tuples and an attribute definition of a relation.
A relation schema $R = \{A_1, A_2 \ldots A_n\}$ is a set of attributes and a relation state $r(R)$ a finite set of mappings $r = \{t_1, t_2 \ldots t_n\}$ where each tuple t_i is a mapping from R to D and D is the union of the attribute domain.

(c) Values and NULL in the tuples:- Each value in the tuple is an atomic value, that it is not divisible into components.

(d) Interpretation of Relation:- The relation schema can be interpreted as a declaration or a type of assertion.
Each tuple can be interpreted as a fact in the relation

c) Consider the following schema
Sailors (Sid, Sname, Rating, age)
Boats (bid, bname, color)
Reserves (Sid, bid, day)

(1) Find the name of the sailors, who have secured red and green boat.

$Temp_{red} \leftarrow \Pi_{sid}(\sigma_{color='red'}(Boats \bowtie Reserves))$

$Temp_{green} \leftarrow \Pi_{sid}(\sigma_{color='green'}(Boats \bowtie Reserves))$

$\Pi_{sname}((Temp_{red} \cap Temp_{green}) \bowtie Sailors)$.

(2) Find the name of sailors, who have reserved a red boat

$Sailors \leftarrow Sailors \bowtie Reserves$

$\Pi_{sname}(\sigma_{color='red'}(Boats) \bowtie Sailors)$.

(3) Find the name of sailors who have reserved a red or green boat.

Tempboats ← ($\sigma_{color='red'}$ Boats) ∪ ($\sigma_{color='green'}$ Boats)

Π_{sname} (Tempboats ⋈ (Reserves ⋈ Sailors))

(4) Find the name of the sailors who have reserved all boats

Tempsids ← ($\Pi_{sid, bid}$ Reserves) ÷ (Π_{bid} Boats)

Result ← Π_{sname} (Tempsids ⋈ Sailors).

Q. 4 a) Explain the steps to convert the basic ER model to relational database schema.

Soln

Step 1: mapping of regular entity types:- For each regular entity type E in ER schema, create a relation R. That include all the simple attributes of E. Choose one of the key attribute of E as the primary key for R

Step 2:- mapping of weak entity ~~type~~ type.

For each weak entity ~~type~~ type W in the ER schema with owner entity type E.

Create a relation R and include all simple attributes of W as attribute of R

Step 3:- Mapping of binary 1:1 mapping relation types.

For each binary 1:1 relationship type R in the ER schema, identify the relations S and T that corresponds to the entity types participating in R

Step 4:- Mapping of the Binary 1:N Relationship type

(1) foreign key approach:- For each ~~binary~~ regular binary 1:N relationship type R identify the relation S. That representing participating entity type at the N side of the relationship type.

② Relationship relation approach: is to use the cross reference option as in the third option for binary 1:1 relationship

Step 5: Mapping of the binary M:N Relationship types:-
In the traditional relational model with no multivalued attributes. the only option for M:N relationship is the relationship relation option.

Step 6: Mapping of the multivalued Attributes
For each multivalued attribute A create a new relation R this relation R will include an attribute corresponding to A plus the primary key attribute K. as a foreign key in R.

Step 7: Mapping of N-ary relationship types
For each n-ary relationship type R, where n>2 Create a foreign key attributes in S, the primary keys of the relations that represents the participating entity.

b) Explain Unary relational operations with example

Soln

Unary Select operation:-
* the select operation is used to choose a subset of the tuples from a relation that satisfies a selection condition.
* we can consider the select operation to select the tuple in a relation to only those tuples that satisfy the condition.

Ex:- To select the employee tuples where department is 4 or those whose salary is greater than $30,000

$\sigma_{Dno=4}(Employee)$
$\sigma_{Dsalary>30000}(Employee)$

* the Project operation: It selects certain colomns from the table and discards the other table columns. If you are intrested in only certain attribute of a relation, we use the Project operation to project the relation over these attributes

Ex:- To list each employee first and last name and salary
Use the project operations as follows.

$\Pi_{Fname, Lname, Salary}(Employee)$.

6) Consider relation schema Employee database

Employee (Fname, Minit, Lastname, SSN, Bdate, Address, Sex, Salary, Super-ssn, Dno)

Department (Dname, Dnumber, mgr-ssn, mgr-start-date)

Project (Pname, Pnumber, Plocation, Dnum)

Works-on (Essn, Pno, hours)

Dependent (Essn, Dependent-name, Sex, Bdate, Relationship)

Write a relation queries for the following

(a) Retrive the name and address of all employee who works for the Research Department.

$\Pi_{Fname, Lname, Address}(\sigma_{Dname = 'Research'}(Department) \bowtie_{Dnumber = Dno} Employee)$

(b) List all the names of all employee with 2 or more dependent.

$\Pi_{Fname, Lname}(Employee \bowtie_{SSN = Essn} \rho_{(Essn, count)}\sigma_{count > 2}(_{Department\_name}\gamma_{count(Dependent)}))$

(c) Find the name of employee who works on all the project controlled by department no = 5

$\Pi_{Fname, Lname}(Employee \ltimes_{Ssn = Essn} (\Pi_{Essn, Pno}(Works\_on) \div \Pi_{Pno}(\sigma_{Dnum = 5}(project))))$

(d) List the name of employees who have no dependents

$\Pi_{Fname, Lname}(Employee) - \Pi_{Fname, Lname}(Employee \bowtie_{Ssn = Essn} Dependent)$

Q 5 a) What is the need for normalization? Explain 2nd and 3NF with examples

Soln Normalization is a systematic process of organizing in a relational database to reduce redundancy and improve data integrity. It provides database designers with the following.

* A formal framework for analyzing relational schemas based on there keys and on the functional dependencies among their attributes.
* A series of normal forms, test that can be carried out individual relation schemas so that the relational database can be managed to any desired degree.

2NF: A relation schema R is in 2NF if every non prime attribute A in R is fully functionally dependent on the primary key of R

It is based on the concept of full functional dependency

A functional dependency $X \rightarrow Y$ is full functional dependency if removal of any attribute A from X means that the dependency does not hold any more.

Ex. Enrollment (studentid, courseid, Sname, coursename)

Functional Dependencies

studentid → student_name

courseid → course_name

conversion into 2NF

Student (studentid, student_name)

course (courseid, course_name)

enrollment (studentid, courseid)

third normal form: a relation schema R is in 3NF, if it satisfies 2NF and no non prime attribute of R is transitively dependent on the primary key.

Ex: Employee (Empid, Ename, dipid, dept_name)

conversion to 3NF

Employee (Empid, Emp_name, dept_id)

Dept (Dept_id, Dept_name)

b) Outline the constraints in SQL.

(1) Specifying attribute constraints and attribute default.

→ A constraint NOT NULL may be specified if NULL is not permitted for a particular attribute.

→ Default value for an attribute by appending the clause DEFAULT <value> to an attribute definition.

(2) Specifying Key and Referencial integrity constraints:-

→ The primary key clause specifies one or more attributes that make up the primary key of a relation.

→ unique key clause specifies alternate, known as candidate key if it is a single attribute.

→ Referential integrity is specified via the FOREIGN key clause

(3) Giving names to constraints:-

A constraint name is used to identify a particular constraints in case the constraint must be dropped later or replaced with another constraints

c) Consider relation R with following function dependency

EMPPROJ (Ssn, Pno, hours, Ename, Pname, Plocation)

Ssn, Pnumber → Hours

Ssn → Ename

Pnumber → Pname, Plocation. If it in 2NF? verify.

Soln

(a) The value of an employees ssn uniquely determines the employee name (Ename)

(b) The value of a projects numbers (Pnumber) uniquely determine the project name (Pname) & Plocation.

(c) A combination of SSN & Pnumber uniquely determine the number of ~~for~~ hours the employee currently works on the project per work.

~~Therefore~~ Therefore the relation is NOT in 2nd normal form.

Q 7

a) i) List the name of managers who have atleast dependent

```
select E.Fname, E.Lname
from Employee e
where E.Ssn in ( select D.mgr_ssn
                 from Department D
                 where exists ( select * from
                                Dependent DP
                                where DP.Essn = D.mgr_ssn ));
```

ii) Retrieve the name of each employee who has a dependent with the same first name & same sex

```
select E.name, E.bname
from employee as E
where Exists ( select *
               from Dependent D
               where D.Essn = E.Ssn
               AND D.Dependent_name = E.Fname
               AND D.Sex = E.Sex);
```

iii) For each project, retrieve project no, project name, and number of employee working on it.

```
Select P.number, P.Pname, count (W.Essn) AS No_of_Employees
from project P, works_on W
where P.Pno = W.Pno
Group by P.Pno, Pname;
```

iv) Retrieve the ssn of all employee who work on project no 1, 2 & 3

```
select Distinct Essn
from works_on
where Pno in (1,2,3);
```

v) Find sum, min, max, avg salary of employees in Research dept.

```
Select Sum(E.Salary), MAX(E.Salary), Min(E.Salary),
AVG(E.Salary)
from Employee E, department D
where E.Dno = D.Dnumber AND D.Dname = 'Research';
```

Q 7 b) Why concurrency if needed? Demonstrate with example

Soln Several problems can occur when concurrent transaction execute in an uncontrolled manner

(1) The lost update problem:- occurs where two transaction that access the same database items have their operation interleaved in a way that makes the value of some database item incorrect.

(2) The temporary update problem:- occurs when one transaction update a database item and then the transaction fails for some reason. meanwhile the updated item is accessed by another transaction before it is changed back to its original value.

(3) The incorrect summary problem:- If one transaction is calculating an aggregate summary function on a number of database items while other transactions are updating some value before they are updated and others after they are updated.

(4) The unrepeatable Read problem:- May occur, where a transaction T reads the same item twice and the item is changed by another transaction T' between the two reads, Hence T receives different values for its two reads of the same item.

Q.8) a) Given schedules

S_1
T_1 : R(X)
T_2 : R(X)
T_1 : W(X)
T_2 : W(Y)
T_1 : R(Y)
T_1 : R(Y)

S_2
T_3 : W(X)
T_1 : R(X)
T_1 : W(Y)
T_2 : R(Z)
T_2 : W(Z)
T_3 : R(Z)

i) What is the Precedence graph for the schedule?

$T_1 : W(x) \rightarrow T_2 : W(y)$
$T_1 : W(y) \rightarrow T_2 : R(y)$
$T_2 : W(y) \rightarrow T_1 : R(y)$

Graph : $T_1 \longleftrightarrow T_2$

ii) Is the schedule conflict serializable? If so what are the conflict equivalent serial schedules?

Conflict operation

① $T_3 : W(x) \rightarrow T_1 : R(x)$
$\Rightarrow$ edge $T_3 \rightarrow T_1$

② $T_2 : W(z) \rightarrow T_3 : R(z)$
$\Rightarrow$ edge $T_2 \rightarrow T_3$

The Precedence graph is acyclic and view serializable
The view serial schedule : $T_2 \rightarrow T_3 \rightarrow T_1$

Q. 8) b) Explain triggers with example. Write a trigger in SQL to call a procedure "Inform supervisor" whenever an employee salary is greater than the salary of his or her direct supervisor in the company database.

Soln A trigger is a special kind of stored procedure that is automatically created by the database management system whenever a specified event occurs in the database.

A trigger consist of three main components : an event, a condition, an action

Event specifies the database operation that activates the trigger, such as INSERT, DELETE or UPDATE on a particular relation

A condition is a boolean expression that is evaluated once the event occurs.

The Action part may consist of one of SQL statement including invocation of stored Procedure.

Consider the company database, where each employee has a direct supervisor and a salary. Suppose an employee's salary become greater than the salary of supervisor, the database

system must automatically call a stored procedure named inform-supervisor.
the trigger can be written as follows

```
Create trigger check-salary
After update the salary on employee
for each row
where (new.salary >
        (select salary
         from employee
         where ssn = new.super-ssn))
begin
    call inform-supervisor;
end;
```

Q 9
a) Describe the two phase locking protocol for concurrency control provide example to illustrate how it ensures serializability in transaction schedule.

Soln It is a concurrency control protocol that ensures conflict serializability of transaction scheduler by controlling how locks are acquired and released.
→ A transaction is said to follow the two phase locking protocol with first unlock operation in the transaction.
→ Such transaction can be divided into two phases
① an expanding, during which new locks on items can be acquired with, but none can be released.
② Shrinking, during which existing locks can be released but no how locks can be acquired.

T1	T2
read_lock (A);	read_lock (A);
read_item (A);	read_item (A);
unlock (A);	unlock (A);

initial value x=20, y=30
Result serial schedule T_1
followed T_2: x=50, y=80
Result of serial schedule T_2
followed by T_1: x=70, y=50.

Transaction T_1 and T_2 do not follow the two phase locking protocol because write_lock(x) operation follows the unlock(y) operation in T_1 and similarly the write operation follows the unlock(x) in T_2

```
T1
read_lock(y);
read_item(y);
write_lock(x);
unlock(y);
read_item(x);
x := x+y;
write_item(x);
unlock(x);
```

```
T2
read_lock(y);
read_item(x);
write_lock(y);
unlock(x);
read_item(y);
y := x+y;
write_item(y);
unlock(y);
```

If every one transaction in a schedule follows the two phase locking protocol, the schedule is guaranteed to be serializable

b) Explain the characteristics of Novell System.

(1) Scalability: In Novell system horizontal-scalability is generally used, where the distributed system is expanded by adding more nodes for data storage & processing as the volume of data grows.

(2) Availability, Replication and Eventual consistency.
Many applications that use Novell system require continuous system availability. To accomplish this data is replicated two or more nodes ~~for~~ in transparent manner, so that if one node fails, the data is still available on other nodes.

Replication improve data availability and can also improve read performance because read request can often be ~~served~~ processed from any of the replicated data nodes

(3) Replication models: Two major replication models are used in NoSQL systems, master-slave and master-master replication.
Master-slave replication requires one copy of the master copy. All the write operations must be applied to the and then propagated to the slave copies usually cause eventual consistency.

(4) Sharing of files: sharing serves to distributing the load of accessing the file records to multiple nodes.

(5) High performance data access: It is necessary to find individual records or objects from among the millions of data records or objects in a file. Two techniques used.

(1) Hashing: Hashing hash is applied to the key K. and location of the object with key K is determined by value of h(K).

(2) Range partitioning:- the location is determined via a range of key values.

10 a) Explain binary locks and shared lock with algorithms

Soln Binary locks can have two states or values locked or unlocked A distinct lock is associated with each data base item x. If lock on x if 1; item cannot be accessed by database operation. if '0' item can be accessed when requested.

```
Lock_item (x):
B: if lock (x) = 0
     then Lock (x) ← 1
   else
   begin
   wait (until Lock (x) = 0
   go to B
   end
Unlock_item (x)
   lock (x) ← 0
```

If every transaction are waiting, then wakeup one of the waiting transaction.

Shared / Exclusive locks:- If a transaction is to write item x it must have exclusive access to x, and shared / exclusive or read / write locks. A lock is associated with an item x. lock(x), has three possible states, read-locked, write locked, or unlocked.

read lock item, because other transactions are allowed to read the item.

write lock item, because a single transactions exclusive holds the locks on the item.

```
read_lock(y):
B: if Lock(x) = "unlocked"
 then begin lock(x) ← "read locked"
    no-of-reads(x) ← 1
 end
 else if lock(x) = "read-locked"
 then no-of-reads(x) ← no.of-reads(x)+1
 else begin
 wait (until lock(x) = "unlocked"
 and the lock manager wakes up the transaction go to B
 end;
```

b) Explain mongoDB data model CRUD operation and distributed system characteristics.

Soln → The mongoDB is a documented oriented NoSQL database that stores data in the form of documents rather than tables and rows

→ documents are grouped into collections which are analogous to relation in database.

→ mongoDB supports embedded documents and arrays allowing complex data structures to be stored within a single document.

CRUD Operations:

(1) Create operation: data is inserted into a collection. Using insert operation, which adds new document to a collection.

(2) Read operation: Data retrieval is performed using query operation that select documents based on specified condition.

(3) Update operation: Existing documents can be modified using update operation.

(4) Delete operation: Documents can be deleted using delete operation based on given criteria.

(5) MongoDB supports rich query expression like indexing, aggregation for efficient data access

(Bheerappa S.)