

Bachelor Of Engineering In Information Technology

2nd Year 1st Semester, Semester Examination-2025

Subject Name – (IT/PC/B/T/213) Database Management Systems

Time: 3 Hrs.

Full Marks=100


CO1 [10]	Q1. (a) Explain the three-schema architecture in a database system with examples. How does it enhance abstraction? (b) Illustrate the differences between logical and physical data independence. Provide real-world examples. (c) What challenges might arise if referential integrity is not maintained in a database system? Provide examples. OR Explain the role of DBA. [4+4+2=10]																																			
CO2 [20]	Q2. (a) Explain how redundant data in a database schema can lead to anomalies. Provide examples to illustrate your points. (b) Consider the following instance of the relational schema: Pet_Store (Pet_ID, Pet_Name, Categories, Price, Breed) <table border="1"><thead><tr><th colspan="5">Pet_Store</th></tr><tr><th>Pet_ID</th><th>Pet_Name</th><th>Category</th><th>Price</th><th>Breed</th></tr></thead><tbody><tr><td>1281</td><td>Tom</td><td>Dog</td><td>10,000</td><td>Beagle</td></tr><tr><td>1282</td><td>Jenny</td><td>Cat</td><td>5,000</td><td>Persian</td></tr><tr><td>1281</td><td>Tom</td><td>Dog</td><td>20,000</td><td>X</td></tr><tr><td>1282</td><td>Tom</td><td>Dog</td><td>15,000</td><td>Labrador</td></tr><tr><td>1281</td><td>Lucy</td><td>Bird</td><td>8,000</td><td>Macaw</td></tr></tbody></table> Identify that the following two functional dependencies do hold or not on Pet_Store. i) Pet_ID → Pet_Name ii) {Category, Price} → Pet_Name (c) How does understanding functional dependencies help in determining whether a schema satisfies Boyce-Codd Normal Form (BCNF)? (d) Consider the relational schema: DeviceLogs(Device, OperatingSystem, Logins, DateOfRecord, DeviceLocation) with the following Functional Dependency set: {Device, OperatingSystem} → DeviceLocation {Logins, DateOfRecord} → OperatingSystem DeviceLocation → DateOfRecord Identify all candidate key(s) for the above relation. What is the highest normal form of DeviceLogs? [5+6+3+6=20]	Pet_Store					Pet_ID	Pet_Name	Category	Price	Breed	1281	Tom	Dog	10,000	Beagle	1282	Jenny	Cat	5,000	Persian	1281	Tom	Dog	20,000	X	1282	Tom	Dog	15,000	Labrador	1281	Lucy	Bird	8,000	Macaw
Pet_Store																																				
Pet_ID	Pet_Name	Category	Price	Breed																																
1281	Tom	Dog	10,000	Beagle																																
1282	Jenny	Cat	5,000	Persian																																
1281	Tom	Dog	20,000	X																																
1282	Tom	Dog	15,000	Labrador																																
1281	Lucy	Bird	8,000	Macaw																																

CO3

Q3.

[20]

(a) Identify the primary key, partial key, composite key, multi-valued attribute, and derived attribute (if any). Additionally, determine the weak entity, strong entity, type of participation, and type of relationship from the given **Entity Relationship Diagram** (if any).



(b) Explain the **division operation** in relational algebra using a suitable example.

(c) Consider the following relational schema:

Employee(EID, ENAME, DEPARTMENT, SALARY)

Project(PID, PNAME, BUDGET)

Assigned(EID, PID, HOURS_WORKED)

i) Write a **SQL query** to display the names of employees who are working on the “ML-NLP” project.

ii) Write a **SQL query** to display the total number of hours worked by each employee across all projects.

(d) Consider the following relational schema:

Suppliers(sID, sName, address)

Parts(pID, pName, colour)

Catalog(sID, pID, price)

Write down the **relational algebra expression** equivalent to the following statement given below.

i) Retrieve the names of suppliers who do not supply any parts with the colour “green”.

ii) Retrieve the names of suppliers who supply the most expensive part. [6+2+6+6=20]

CO4

Q4.

[20]

(a) Write down the difference between primary and secondary indexes in DBMS.

OR

Justify this statement given below:

“Size of non-leaf nodes of B+ tree is larger than the leaf nodes”.

(b) Why must secondary indexes be dense in a database management system (DBMS)?

(c) Consider a disk pack with the following specifications: 16 double-sided platters, 256 tracks per surface, 512 sectors per track and 1024 bytes per sector. Find out the **capacity** of the disk.

(d) Identify the cost estimation of a query evaluation plan, if 500 blocks are required to be transferred from the disk and the required number of disk seeks is 20. Time to transfer one block (t_r) is 2 milliseconds. Time for one seek (t_s) is 0.2 seconds. [4+4+7+5=20]

CO5

[20]

Q5.

(a) Consider the following schedule S involving five transactions T_1, T_2, T_3, T_4, T_5 .

T_1	T_2	T_3	T_4	T_5
R(X)				
		R(X)		
		W(X)		
	W(X)			
				W(Z)
			R(Z)	
			W(Z)	
R(Z)				

R(X) denotes read operation on data item X by transaction T_i .W(X) denotes write operation on data item X by transaction T_i .

Identify whether the provided schedule S is a **conflict serializable schedule**, a **view serializable schedule**, **both**, or **none of these**. In the case of **conflict serializability**, determine the possible number of conflict serializable schedules, and for each identified schedule, provide the **correct order of execution**.

OR

Consider two schedules **S1** and **S2** as follows. Please note that if any schedule suffers from deadlock, some operations of the transactions in that schedule may not be executed.

S_1	
T_1	T_2
lock-X(A)	
read(A)	
write(A)	
unlock(A)	
	lock-S(B)
	read(B)
lock-S(B)	
read(B)	
unlock(B)	
	lock-X(A)
	read(A)
	write(A)
	unlock(A)
	unlock(B)

S_2	
T_1	T_2
lock-X(A)	
read(A)	
write(A)	
	lock-X(B)
	read(B)
lock-S(B)	lock-X(A)
read(B)	read(A)
unlock(B)	write(A)
unlock(A)	
	unlock(A)
	unlock(B)

Justify these two statements:

i) "S1 will suffer from deadlock, S2 will not suffer from deadlock."

ii) "S1 will not suffer from deadlock, S2 will suffer from deadlock."

(b) Assume an **immediate database modification scheme**. Consider the following log records for transactions T1, T2, T3 and T4:

steps	Details of log
1	⟨T1,start⟩
2	⟨T2,start⟩
3	⟨T1,A,500,600⟩
4	⟨T1,commit⟩
5	⟨T2,B,200,400⟩
6	⟨checkpoint{T2}⟩
7	⟨T3,start⟩
8	⟨T2,commit⟩
9	⟨T3,C,200,500⟩
10	⟨T4,start⟩
11	⟨T4,D,600,1000⟩
12	⟨T4,commit⟩

If there is a **crash** just after step 12 and the recovery of the system is successfully completed, identify the **final values of all data items, undo list, redo list and no action list** of transactions if any for the above scenario.

(c) Explain the role of a transaction in a database system. Discuss how the ACID properties contribute to transaction management.

[8+7+5=20]

CO6 Q6. (Answer any two)

[10] (a) Why do we need a distributed database?

(b) Why do Asynchronous Distributed Databases exhibit comparatively faster response times?

(c) Write down the 5V's (characteristics) of big data.

[5+5=10]

(IT/PC/B/T/213) Database Management Systems:

After completing this course, the students should be able to:

CO1: Explain the basic Database concepts and different data models. (K2)

CO2: Find the available functional dependencies to apply normalization concepts in typical scenarios. (K3)

CO3: Design queries using relational algebra operations and SQL. (K3)

CO4: Explain principles of Physical Data Storage and Query Optimization. (K3)

CO5: Comprehend transaction processing and concurrency control techniques and apply them in various problems (K3)

CO6: Discuss different types of advanced databases. (K2)