

BACHELOR OF INFORMATION TECHNOLOGY ENGG. EXAMINATION, 2019
2nd year, 1st semester
Database Management Systems

Time: 3 Hours

Marks: 100

Note: Read the questions carefully and answer all six COs.

CO1: Answer any one question	Full marks: 10
<p>1A. (i) Describe the three-schema architecture. Why do we need the mappings between different schema levels? (ii) Explain the disadvantage of Network database model. (iii) What are the disadvantages of database system? Explain them briefly. (4+2+2+2)</p> <p>1B. Draw the ER diagram for the system given below and also indicate the constraints on the relation. An organization is executing number of projects. Periodically, experts review the projects and submit reports. A project is reviewed by one or more experts. An expert may review multiple projects. Each reviewer provides separate report for each project he/she reviews. System must keep the information about experts, projects and reports. It must be able to find out which report corresponds to which expert and which project. Also notes which project is reviewed by which expert and on which date. (10)</p>	
CO2: Answer any one question	Full marks: 15
<p>2A. (i) 'Armstrong's Axioms are sound and complete' — explain. (ii) With an example discuss why referential integrity constraint and functional dependency constraint need to be satisfied for any relational schema. (iii) BCNF may not be dependency preserving. Why? (iv) Give a set of functional dependencies for the relation schema R(A,B,C,D,E) with primary key AB under which R is in 2NF but not in 3NF. (3+4+3+5)</p> <p>2B. (i) Explain the anomalies that may occur when two integrity rules are violated. (ii) What is spurious tuple? (iii) Consider a schema to store the following information for each student: Roll, Name, Address, Coursecode, Coursename, Courseduration, Coursefee, and for each subject taught in the course SubjectCode, SubjectName, FullMarks, PassMarks, MarksScored. Each student has unique Roll and further assume the following FDs : Coursecode → Coursename, Courseduration, Coursefee. SubjectCode → SubjectName, FullMarks, PassMarks. Roll, SubjectCode → MarksScored. Normalize upto 3NF. Show steps and indicate foreign keys. (3+2+10)</p>	
CO3: Answer any two question	Full marks: 10
<p>3A. (i) What is a view? (ii) What is the difference between WHERE and Having Clause? (2+3)</p> <p>3B. Consider the relations and write down the relational algebra expressions: DEPT (DCODE, DNAME) and EMP (ECODE, ENAME, DCODE, BASIC). (i) Find out the name of the employee(s) with highest basic. (ii) Display the name of the employees and name of the department (if it is available) for those with BASIC > 10000. (2+3)</p>	

3C. Consider the following schemas and write down the SQL statement for the following :

STUDENT (ROLL, NAME)
 SUBJECT (SCODE, SNAME, PASS_MARKS)
 RESULT (ROLL, SCODE, MARKS)

- (i) Show for all subjects: the subject name and number of students passed in that subject.
- (ii) Find out the name of the students who have scored 80 or more in all subjects.

(2+3)

CO4: Answer any three questions

Full marks: 15

- 4A. (i) What are the steps involved in query processing?
- (ii) How do you measure the cost of query evaluation?

(3+2)

B. What is the heuristics optimization technique? Compare it with the cost based optimization techniques.

(2+3)

C. Find out the optimized query tree for the given SQL query:

*SELECT S.name FROM Student S, Book B, Checkout C
 WHERE S.sid = C.sid AND B.bid = C.bid AND B.author = 'Olden Fames'
 AND S.age > 12 AND S.age < 20;*

Student (sid, name, age, address)

Book (bid, title, author)

Checkout (sid, bid, date)

Also assume there are 10, 000 Student records and 50, 000 Book records with 700 authors and 30,000 Checkout records.

(5)

D. Why is accessing a disk block expensive? Discuss the time components involved in accessing a disk block.

(2+3)

CO5: Answer any two questions

Full marks: 40

5A. (i) What are the ACID properties of a transaction?

(ii) Discuss the actions taken by the read_item and write_item operations on a database.

(iii) Discuss a read only transaction with one example.

(iv) What is a serializable schedule?

(v) Is this schedule conflict-serializable? Give reason for your decision.

	T1	T2	T3	T4
	Write(B)	Read(C)		
	Read(A)		Read(A)	
	Read(C)			Read(A)
	Read(A)		Write(B)	Write(A)
		Write(C)		

(4+4+4+2+6)

5B. (i) What is 2-phase locking protocol? Explain, whether it is free from deadlock or not with an appropriate example.

(ii) In time stamp based protocol, if a transaction T_i wants to modify data item X, under which condition it can modify X.

(iii) In multi-version technique of concurrency control, when does a new version of Q will be created?

(iv) A partial schedule for several data items for transactions with timestamps 1, 2, 3, and 4. Check whether this schedule will be allowed to execute or not? Support your answer.

T1	T2	T3	T4
Read(Q) Write(Q)			Read(P)
Read(P)	Read(P)	Read(Q) Read(R)	Read(Q)
Write(P)	Write(P)	Write(R)	Write(Q)

(v) Write an algorithm to avoid deadlock.

(5+4+3+5+3)

5C. (i) What is the advantage of using Stable-Storage in recovery technique.

(ii) "Redo and undo operations must be idempotent." Explain it.

(iii) What is write-ahead logging?

(iv) Describe the shadow paging recovery technique. How it is different from Log based recovery technique?

(v) A Database management system uses deferred database modification technique. In an instances of time, the following log entries were there:

```

<T1 start>
<T1, X, 10, 50>
<T2 start>
<T0 commit>
<T2, Y, 200, 100>
<T2, Z, 100, 120>
<checkpoint {T1, T2}>
<T3 start>
<T2 commit>
<T4 start>
<T4, X, 50, 20>
<T3, Y, 100, 10>
<T3 commit>
<T4, Z, 120, 150>
Crashed...

```

Identify the transactions that will be rollback or redone?

(2+3+3+4+4+4)

C06: Write a short note on any two question

Full Marks: 10

6A. Data Warehouse Architecture

B. Distributed database environments

C. Data Mining Technologies

D. Market Basket Analysis

(5+5)