KONERU LAKSHMAIAH EDUCATION FOUNDATION
(Deemed to be University, Estd. u/s. 3 of UGC Act 1956)

B.Tech - Odd Sem : End Semester Exam
Academic Year:2020-2021

19CS2108S - DATABASE MANAGEMENT SYSTEMS

Set No: 1

Time: **Max.Marks: 100**

S.NO	Answer All Questions	Choice	Options	Marks	CO
1.	Consider the following database schema (University Management System): student (ID, name, dept_name, tot_cred) ;advisor (s_id, i_id) ;takes (ID, course_id, sec_id, semester, year, grade) ;classroom (building, roomnumber, capacity) ;timeslot (timeslot_id, day, start_time, end_time); section (course_id, sec_id, semester, year, building, roomnumber, timeslot_id);teaches (ID, course_id, sec_id, semester, year); department (dept_name, building, budget) ;instructor(ID, name, dept_name, salary);course(course_id, title, dept_name, credits) ;prereq(course_id, prereq_id).; Design an ER diagram for the above scenario by identifying various constraints and relationships between the entities.	choice Q-2		10Marks	CO1
2.	Differentiate specialization hierarchy & specialization lattice			10Marks	CO1
3.	The loan office in a bank receives from various parties requests to investigate the credit status of a customer. Each credit request is identified by a Request ID and is described by a Request Date and Requesting Party Name. The loan office also received results of credit checks. A credit check is identified by a Credit Check ID and is described by the Credit Check Date and the Credit Rating. The loan office matches credit requests with credit check results. A credit request may be recorded before its result arrives; a particular credit result may be used in support of several credit requests. Design an ER Diagram for this situation. Now, assume that credit results may not be reused for multiple credit requests. Redesign the ER Diagram for this new situation using two entity types, and then redraw it again using one entity type. Justify which among these two versions do you prefer, Justify?	choice Q-4		15Marks	CO1
4.	Design Relational schema to the below mentioned ER Diagram 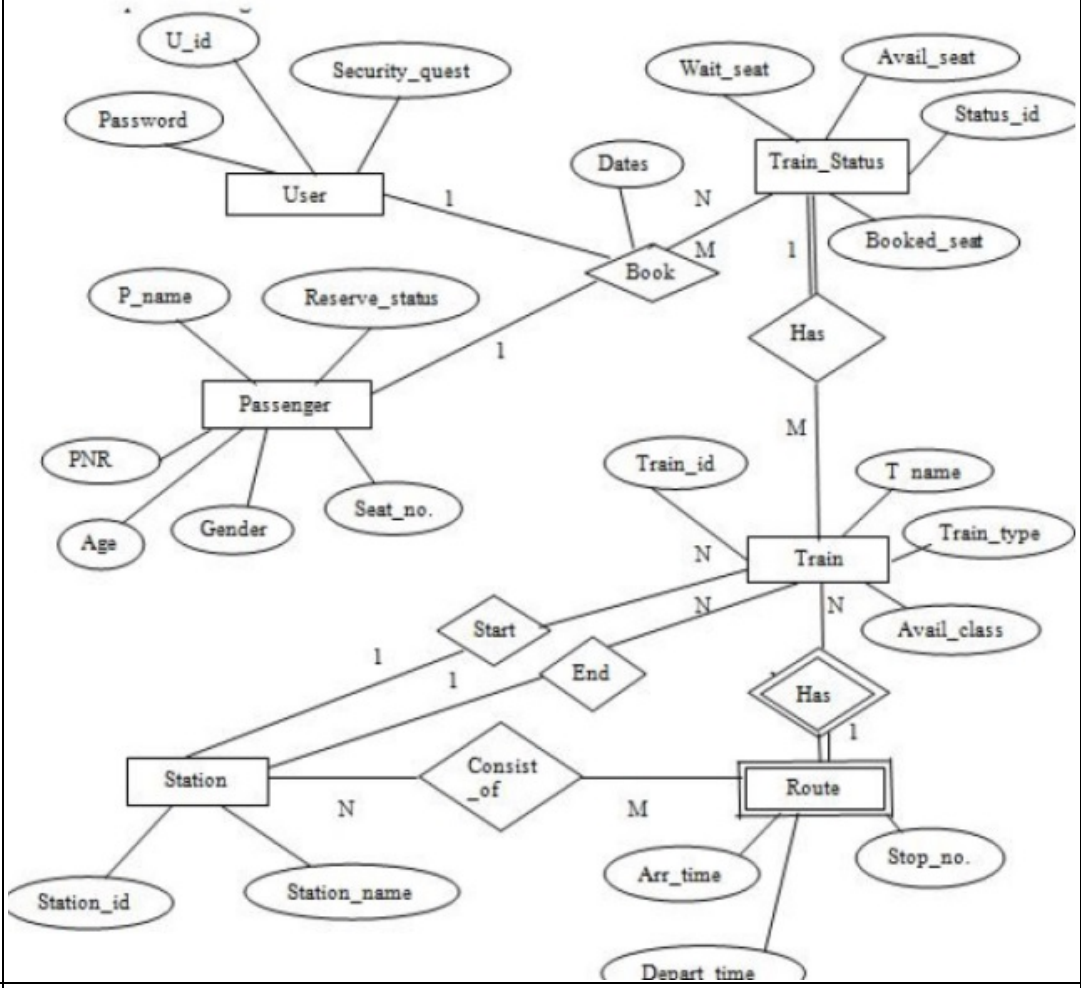			15Marks	CO1
5.	Consider a schema with two relations, R(A, B) and S(B, C), where all values are integers. Make no assumptions about keys. Consider the following three relational algebra expressions: Two of the three expressions are equivalent (i.e., produce the same answer on all databases), while one of them can produce a different answer. Which query can produce a different answer? Give the simplest database instance you can think of where a different answer is produced	choice Q-6		10Marks	CO2

	a. $\pi_{A,C}(R \bowtie \sigma_{B=1} S)$ b. $\pi_A(\sigma_{B=1} R) \times \pi_C(\sigma_{B=1} S)$ c. $\pi_{A,C}(\pi_A R \times \sigma_{B=1} S)$				
6.	Consider the following relational schema Student(id, name), EnrolledIn (id, code) Subject (code, lecturer). Given the above relation, and our university schema, write each of the following queries in SQL. i. Find the names of students enrolled in cs3020? ii. Which subjects is Hector taking? iii. List the lecturers who teach cs1500? iv. What are the names of all the students? v. List all the lecturers who teach at least two different subjects? vi. Find the names of students enrolled for a subject taught by Roger? vii. Find the names of students in both cs1500 and cs1200? viii. What are the names of students in at least two different subjects? ix. What are the codes of all the subjects taught?			10Marks	CO2
7.	Consider the following schema: Suppliers(sid: integer, sname: string, address: string) ,Parts(pid: integer, pname: string, color: string) ,Catalog(sid: integer, pid: integer, cost: real) ,The Catalog relation lists the prices charged for parts by Suppliers. Write Relational Algebra expressions for the following 1. Find the pnames of parts for which there is some supplier. 2. Find the snames of suppliers who supply every part. 3. Find the snames of suppliers who supply every red part. 4. Find the pnames of parts supplied by Acme Widget Suppliers and no one else. 5. Find the sids of suppliers who charge more for some part than the average cost of that part (averaged over all the suppliers who supply that part). 6. For each part, find the sname of the supplier who charges the most for that part. 7. Find the sids of suppliers who supply only red parts. 8. Find the sids of suppliers who supply a red part and a green part. 9. Find the sids of suppliers who supply a red part or a green part. 10. For every supplier that only supplies green parts, print the name of the supplier and the total number of parts that she supplies.	choice Q-8		15Marks	CO2
8.	Model the following queries in the Relational Algebra using below relational schema? Product (maker, model, type), PC(model, speed, ram, hd, price), Laptop(model, speed, ram, hd, screen, price), Printer(model, color, type, price). i. Find the model number, speed, and hard-disk size for all PC's whose price is under $1000. ii. Find the manufacturers of printers. ii. Find the model number, memory size, and screen size for laptops costing more than $1500. iv. Find all the PCs whose cost is in the range 20000 to 40000. v. Find the model number and hard-disk size for those PC's that have a speed of 3.2 and a price less than $2000. vi. Give the manufacturer and speed of laptops with a hard disk of at least thirty gigabytes. vii. Find the model number and price of all products (of any type) made by manufacturer B. viii. Find the makers of PC's with a speed of at least 3.0. ix. Find the makers who produces both PC and Laptop with ram greater than 2GB			15Marks	CO2
9.	Consider a relation R with five attributes A,B,C,D,E. You are given the following dependencies: A --t B, Be ~ B, and BD ~ A.1. List all keys for R.2. Is R in :3NF?3. Is R in BCNF?	choice Q-10		10Marks	CO3
10.	Illustrate Extendable hashing with a suitable example.			10Marks	CO3
11.	Consider the disk specifications, a disk with a sector size of 512 bytes, 2000 tracks per surface, 50 sectors per track, five double-sided platters, and average seek time of 10 msec.and suppose that a block size of 1024 bytes is chosen. Suppose that a file containing 100,000 records of 100 bytes each is to be stored on such a disk and that no record is allowed to span two blocks. L How many records fit onto a block? 2. How many blocks are required to store the entire file? If the file is arranged sequentially on disk, how many surfaces are needed? 3. How many records of 100 bytes each can be stored using this disk? 4. If pages are stored sequentially on disk, with page 1 on block 1 of track 1, what page is stored on block 1 of track 1 on the next disk surface? How would your answer change if the disk were capable of reading and writing from all heads in parallel? 5. What time is required to read a file containing 100,000 records of 100 bytes each sequentially? Again, how would your answer change if the disk were capable of reading/writing from all heads in parallel (and the data was arranged optimally)? 6. That is the time required to read a file containing 100,000 records of 100 bytes each in a random order? To read a record, the block containing the record has to be fetched from disk. Assume that each block request incurs the average seek time and rotational delay.	choice Q-12		15Marks	CO3
12.	Suppose you are given a relation StaffPropertylnspection (propertyNo, iDate, iTime, pAddress, comments, staffNo, sName, carReg) with the following functional dependencies: : FD1: propertyNo, iDate -> iTime, comments, staffNo, sName, carReg, FD2 : propertyNo -> pAddress ,FD3 : staffNo -> sName , FD4 : staffNo, iDate -> carReg , FD4 : carReg, iDate, iTime -> propertyNo, pAddress, comments, staffNo, sName , FD6 : staffNo, iDate, iTime -> propertyNo, pAddress, comments i. Find all candidate keys? ii. Identify the best normal form that StaffPropertylnspection satisfies (1NF, 2NF, 3NF, or BCNF)? iii. If the relation is not in BCNF, decompose it until it becomes BCNF. At each step, identify a new relation, decompose and recompute the keys and the normal forms they satisfy?			15Marks	CO3
13.	Analyze the following schedules and determine which of them is (conflict) serializable? For each serializable schedule, determine the equivalent serial schedules. a. r1(X); r3(X); w1(X); r2(X); w3(X); b. r1(X); r3(X); w3(X); w1(X); r2(X); c. r3(X); r2(X); w3(X); r1(X); w1(X); d. r3(X); r2(X); r1(X); w3(X); w1(X);	choice Q-14		10Marks	CO4
14.	Consider the following two transactions: T1 : Read_item(A),Write_item(A), Read_item(B),Write_item(B), Commit T2 :Read_item(B), Read_item(C),Write_item(C),Write_item(B), Commit Consider the following interleaved schedule of the two transactions: Read_itemT1(A), Read_itemT2(B), Read_itemT2(C),Write_itemT1(A), Read_itemT1(B),Write_itemT1(B),Write_itemT2(C),Write_itemT2(B),CommitT1,CommitT2. Is the schedule serializable? If you claim yes, write an equivalent serial (non-interleaved) execution of the two transactions. If you claim no, explain why it is not serializable			10Marks	CO4
15.	Consider the execution of two transactions T1 and T2 assume that if the initial values of X, Y, M and N are 100, 800, 10, 45 respectively. i. Write the final values of X and Y as per schedule A. Is this a serializable schedule? ii. Write the final values of X and Y for all possible serial schedules as per schedule B.	choice Q-16		15Marks	CO4

ScheduleA

T1	T2
Read_item(X)	
X := X-N	
	Read_item(X)
	X :=X+ M
Write_item(X)	
Read_item(Y)	
	Write_item(X)
Y := Y+N	
Write-item(Y)	

Schedule B

T1	T2
Read_item(X)	
X := X-N	
Write_item(X)	
Read_item(Y)	
Y := Y+N	
Write-item(Y)	
	Read_item(X)
	X :=X+ M
	Write_item(X)

16. A DBMS attempts to run the following schedule. Show: (I) How conservative 2PL would prevent deadlock. (II) .How ordering all data items would prevent deadlock. (III) How the wait-for scheme would prevent deadlock. (IV) How the wound-wait scheme would prevent deadlock

operation	T1	T2	operation
1	Read_lock(Y);		
2	Read_item(Y);		
3	Y:=Y+3;		
		Read_lock(Y);	1
		Read_item(Y);	2
		Y:=1.2*Y;	3
4	Write_lock(Y);		
5	Write_item(Y);		
6	Read_lock(X);		
7	Read_item(X);		
8	X:=X+Y;		
9	Write_lock(X);		
10	Write_item(X);		
11	Unlock(X);		
12	Unlock(Y);		
		Write_lock(Y);	4
		Write_item(Y);	5
		Read_lock(X);	6
		Read_item(X);	7
		X:=X+Y;	8
		Write_lock(X);	9
		Write_item(X);	10
		Unlock(X);	11
		Unlock(Y);	12

15Marks CO4

[object HTMLDivElement]