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

19CS2108A - Database Management Systems (Advanced)

Set No: 3

Time: **Max.Marks: 100**

S.NO	Answer All Questions	Choice	Options	Marks	CO
1.	A university database contains information about professors (identified by social security number, or SSN) and courses (identified by courseid). Professors teach courses; each of the following situations concerns the Teaches relationship set. For each situation, draw an ER diagram that describes it (assuming that no further constraints hold). 1. Professors can teach the same course in several semesters, and each offering must be recorded. 2. Professors can teach the same course in several semesters, and only the most recent such offering needs to be recorded. (Assume this condition applies in all subsequent questions.) 3. Every professor must teach some course. 4. Every professor teaches exactly one course (no more, no less). 5. Every professor teaches exactly one course (no more, no less), and every course must be taught by some professor. 6. Now suppose that certain courses can be taught by a team of professors jointly, but it is possible that no one professor in a team can teach the course. Model this situation, introducing additional entity sets and relationship sets if necessary	choice Q-2		10Marks	CO1
2.	A company database needs to store information about employees (identified by ssn, with salary and phone as attributes); departments (identified by dno, with dname and budget as attributes); and children of employees (with name and age as attributes). Employees work in departments; each department is managed by an employee; a child must be identified uniquely by name when the parent (who is an employee; assume that only one parent works for the company) is known. We are not interested in information about a child once the parent leaves the company. Draw an ER diagram that captures this information.			10Marks	CO1
3.	Answer each of the following questions briefly. The questions are based on the following relational schema: Emp(eid: integer, ename: string, age: integer, salary: real) Works(eid: integer, did: integer, pct time: integer) Dept(did: integer, dname: string, budget: real, managerid: integer) 1. Give an example of a foreign key constraint that involves the Dept relation. What are the options for enforcing this constraint when a user attempts to delete a Dept tuple? 2. Write the SQL statements required to create the above relations, including appropriate versions of all primary and foreign key integrity constraints. 3. Define the Dept relation in SQL so that every department is guaranteed to have a manager. 4. Write an SQL statement to add 'John Doe' as an employee with eid = 101, age = 32 and salary = 15, 000. 5. Write an SQL statement to give every employee a 10% raise. 6. Write an SQL statement to delete the 'Toy' department. Given the referential integrity constraints you chose for this schema, explain what happens when this statement is executed.	choice Q-4		15Marks	CO1
4.	Students(sid: string, name: string, login: string, age: integer, gpa: real) Faculty(fid: string, fname: string, sal: real) Courses(cid: string, cname: string, credits: integer) Rooms(rno: integer, address: string, capacity: integer) Enrolled(sid: string, cid: string, grade: string) Teaches(fid: string, cid: string) Meets_In(cid: string, rno: integer, time: string) Consider the relations Students, Faculty, Courses, Rooms, Enrolled, Teaches, and Meets_In given above. 1. List all the foreign key constraints among these relations. 2. Give an example of a (plausible) constraint involving one or more of these relations that is not a primary key or foreign key constraint.			15Marks	CO1
5.	Consider the following relation schema: Sailors(sid: integer, sname: string, rating: integer, age: real) Boat(bid: integer, bname: string, color: string) Reserves(sid: integer, bid: integer, day: date) Write the following queries in SQL. (i) Find the average age of the sailor who are eligible for voting for each rating level that has at least two sailors. (ii) Find the name of sailors who have reserved both red and a green boat. (iii) Find the sailor_id of sailors who have reserved a red boat	choice Q-6		10Marks	CO2
6.	Consider the following relation R(A,B,C,D,E) and FD's ABC, CA,DE, FA, ED. Is the decomposition of R into R1(A, C, D), R2(B, C, D) AND R3(E,F,D) lossless?			10Marks	CO2
7.	Consider the following collection of relations and dependencies. Assume that each relation is obtained through decomposition from a relation with attributes ABCDEFGHI and that all the known dependencies over relation ABCDEFGHI are listed for each question. (The questions are independent of each other, obviously, since the given dependencies over ABCDEFGHI are different.) For each (sub) relation: (a) State the strongest normal form that the relation is in. (b) If it is not in BCNF, decompose it into a collection of BCNF relations. 1. R1(A,C,B,D,E), A → B, C → D 2. R2(A,B,F), AC → E, B → F 3. R3(A,D), D → G, G → H 4. R4(D,C,H,G), A → I, I → A 5. R5(A,I,C,E)	choice Q-8		15Marks	CO2
8.	Suppose that we have the following three tuples in a legal instance of a relation schema S with three attributes ABC (listed in order): (1,2,3), (4,2,3), and (5,3,3). 1. Which of the following dependencies can you infer does not hold over schema S? (a) A → B (b) BC → A (c) B → C 2. Can you identify any dependencies that hold over S?.			15Marks	CO2
9.	Consider the Linear Hashing index shown in Figure. Assume that we split whenever an overflow page is created. Answer the following questions about this index: 1. What can you say about the last entry that was inserted into the index? 2. What can you say about the last entry that was inserted into the index if you know that there have been no deletions from this index so far? 3. Suppose you know that there have been no deletions from this index so far. What can you say about the last entry whose insertion into the index caused a split? 4. Show the index after inserting an entry with hash value 4. 5. Show the original index after inserting an entry with hash value 15	choice Q-10		10Marks	CO3

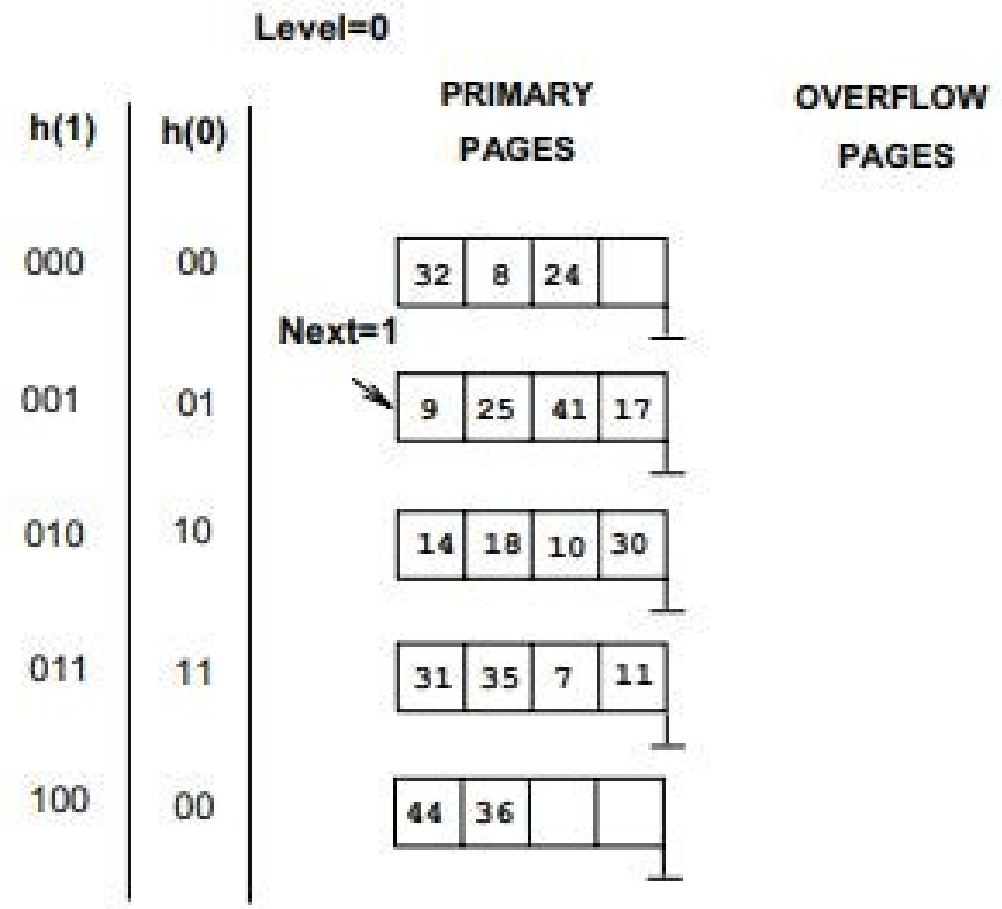

10.	Consider the snapshot of the Linear Hashing index shown in Figure. Assume that a bucket split occurs whenever an overflow page is created. 1. What is the maximum number of data entries that can be inserted (given the best possible distribution of keys) before you have to split a bucket? Explain very briefly. 2. Show the file after inserting a single record whose insertion causes a bucket split.			10Marks	CO3

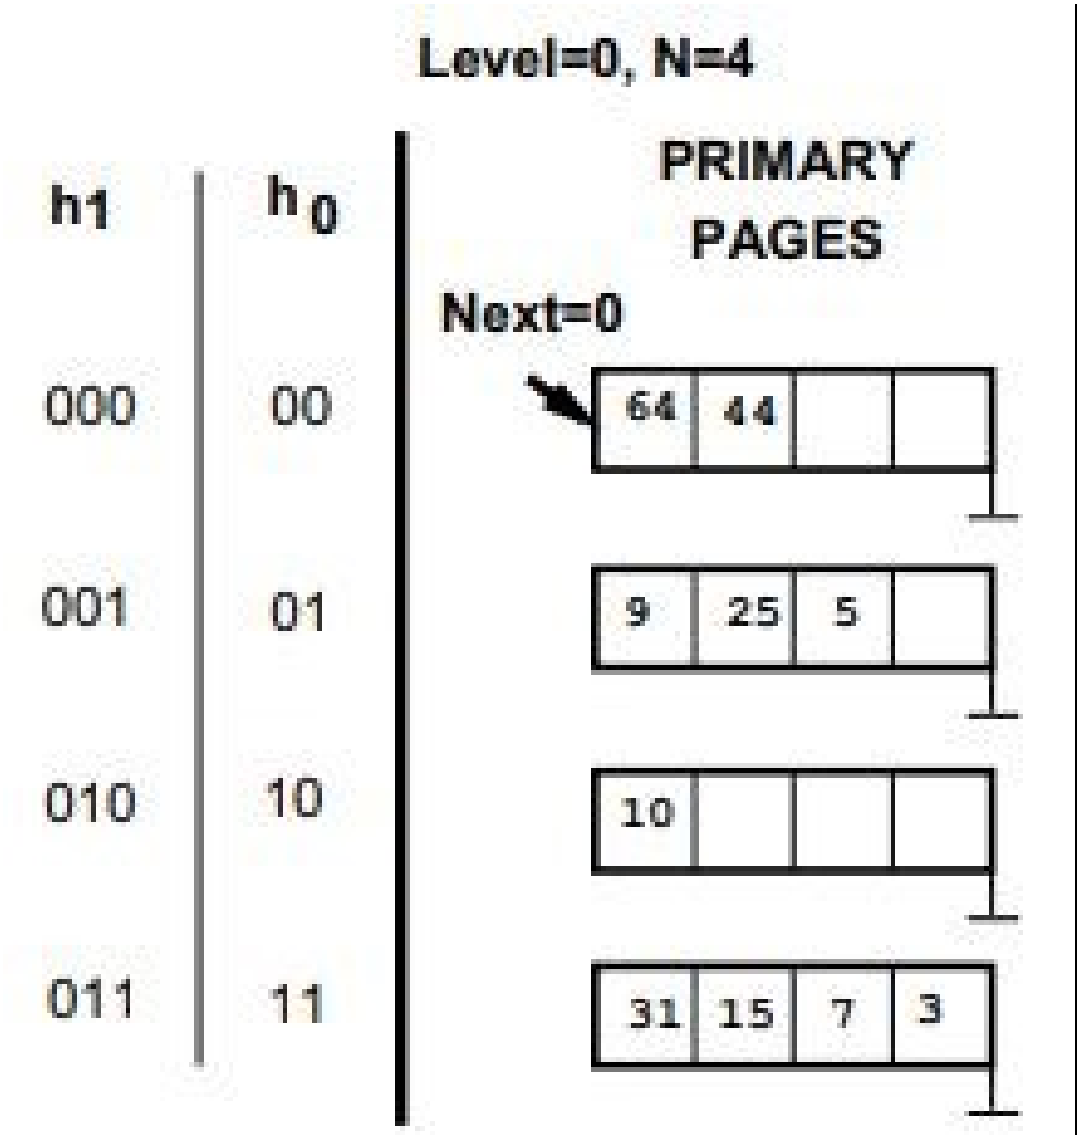

11.	Consider the following actions taken by transaction T1 on database objects X and Y: R(X), W(X), R(Y), W(Y) 1. Give an example of another transaction T2 that, if run concurrently to transaction T without some form of concurrency control, could interfere with T1. 2. Explain how the use of Strict 2PL would prevent interference between the two transactions. 3. Strict 2PL is used in many database systems. Give two reasons for its popularity.	choice Q-12		15Marks	CO3
12.	Consider the following (incomplete) schedule S: T1:R(X), T1:R(Y), T1:W(X), T2:R(Y), T3:W(Y), T1:W(X), T2:R(Y) 1. Can you determine the serializability graph for this schedule? Assuming that all three transactions eventually commit, show the serializability graph. 2. For each of the following, modify S to create a complete schedule that satisfies the stated condition. If a modification is not possible, explain briefly. If it is possible, use the smallest possible number of actions (read, write, commit, or abort). You are free to add new actions anywhere in the schedule S, including in the middle. (a) Resulting schedule avoids cascading aborts but is not recoverable. (b) Resulting schedule is recoverable. (c) Resulting schedule is conflict-serializable.			15Marks	CO3
13.	Show that there are schedules that are possible under the two-phase locking protocol, but are not possible under the timestamp protocol, and vice versa.	choice Q-14		10Marks	CO4
14.	Describe the consistent hashing schema for data distribution, replication, and sharding. How are consistency and versioning handled in Voldemort?			10Marks	CO4
15.	Consider a database with objects X and Y and assume that there are two transactions T1 and T2. T1 first reads X and Y and then writes X and Y. T2 reads and writes X then reads and writes Y. (a) Give an example schedule that is not serializable. Explain why your schedule is not serializable. (b) Show that strict 2PL disallows this schedule.	choice Q-16		15Marks	CO4
16.	Consider a database with objects X and Y and assume that there are two transactions Tl and T2. Transaction T1 reads objects X and Y and then writes object X. Transaction T2 reads objects X and Y and then writes objects X and Y. a. Give an example schedule with actions of transactions T1 and T2 on objects X and Y that results in a write-read conflict. b. Give an example schedule with actions of transactions T1 and T2 on objects X and Y that results in a read-write conflict. c. Give an example schedule with actions of transactions T1 and T2 on objects X and Y that results in a write-write conflict. d. For each of the three schedules, show that Strict 2PL disallows the schedule.			15Marks	CO4

[object HTMLDivElement]