Reg No.:_____ Name:_____

APJ ABDUL KALAM TECHNOLOGICAL UNIVERSITY

FOURTH SEMESTER B.TECH DEGREE EXAMINATION, APRIL 2018

Course Code: CS208

Course Name: PRINCIPLES OF DATABASE DESIGN (CS, IT)

Max. Marks: 100 Duration: 3 Hours

Limit answers to the required points.

PART A

Answer all questions, each carries 3 marks. Marks

1 List any *three* categories of database users, highlighting any *one* important characteristic of *each* category. (3)

2 In a relationship of *degree* 2, how can we decide if an attribute of the relationship can be moved to one of the entity sets? (3)

3 Give suitable examples for *multi-valued*, *composite* and *multi-valued composite* attributes. (3)

4 Given the relations PROFESSOR(PID,PNAME, DEPT, SALARY) and STUDENT(ROLLNO, NAME, CLASS, ADVISER) where ADVISER is a foreign key to PROFESSOR, write *one* equivalent relational algebra expressionfor *each* of the following: (3)

i. $\pi_{PNAME}(\sigma_{SALARY>30000}(PROFESSOR))$

ii. $\pi_{SNAME}(\sigma_{DEPT='CSE'}(STUDENT \underset{ADVISER=PID}{\bowtie} PROFESSOR))$

PART B

Answer any two full questions, each carries 9 marks.

5 a) Distinguish between *total* and *partial* participation constraints with the help of real examples. (4)

b) Describe the real-world situation described by the following ER diagram. (3)

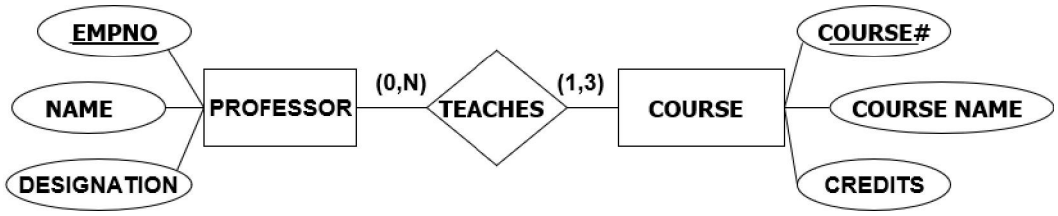

c) Re-draw the ER diagram in 5(b) replacing the (*min,max*) notation with the conventional notation showing cardinality and participation. (2)

6 The relational schema for a library describing members, books and issue information is given below. Foreign keys have the same name as primary keys.

BOOKS(ACC-NO, ISBN, TITLE, EDITION, YEAR)

MEMBERS(MEMBERID, MEMBERNAME, MEMBERTYPE)
ISSUEDTO(ACC-NO, MEMBERID, DATE OF ISSUE)

Write relational algebra expressions for the following queries:

a) Accession Number(s) and Name(s) of third edition books published in 2018. (2)

b) Names and dates of issue of books taken by a member with name 'PRIYA'. (3)

c) Names of books *not* taken by any member. (4)

7 a) Use the standard synthesis procedure to generate the set of relations corresponding to the ER diagram below. Identify primary and foreign keys of the generated relations. (4)

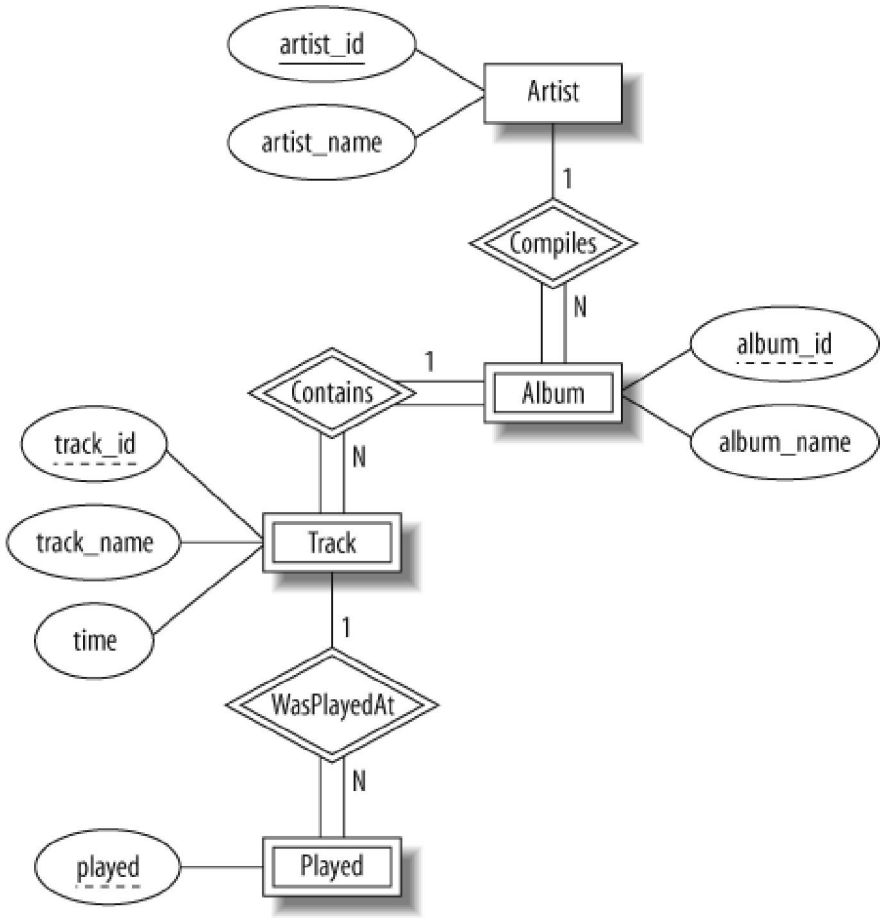

b) What is meant by referential integrity? How is it implemented using foreign key? Illustrate using a real example. (5)

PART C

Answer all questions, each carries 3 marks.

8 Consider two tables STUDENT(ROLLNO,NAME,CLASS) and ENROLLMENT(ROLLNO,COURSENAME) where ROLLNO in ENROLLMENT is a foreign key referring to STUDENT. It is required that every time a STUDENT tuple is deleted, all the ENROLLMENT tuples referring (3)

to the deleted STUDENT tuple are also deleted. Write SQL statements to specify this foreign key requirement.

9 Illustrate use of *assertions* with an example. (3)

10 When do you say that two sets of functional dependencies are *equivalent*? Give a brief explanation. (3)

11 What is meant by *transitive dependency*? Given an example. (3)

PART D

Answer any two full questions, each carries 9 marks.

12 Consider the following relations:

FACULTY(FNO, NAME, GENDER, AGE, SALARY, DNUM)
DEPARTMENT(DNO, DNAME, DPHONE)
COURSE(CNO, CNAME, CREDITS, ODNO)
TEACHING(FNO, CNO, SEMESTER)

DNUM is a foreign key that identifies the department to which a faculty belongs. ODNO is a foreign key identifying the department that offers a course.

Write SQL expressions for the following queries:

a) Course numbers and names of 3-credit courses offered by 'CS' department. (2)

b) Names of faculty members teaching *maximum*3 courses. (3)

c) Names of departments along with number of courses offered by each of them, in the increasing order of number of courses; exclude departments which do not offer any course. (4)

13 Given a relation R(A1,A2,A3,A4,A5) with functional dependencies A1→A2A4 and A4→A5, check if the decomposition R1(A1,A2,A3), R2(A1,A4), R3(A2,A4,A5)is lossless. (9)

14 a) For the relations listed below, write SQL statements for the updates that follow. (Assume suitable domains for attributes.) (5)

ALBUMS(ALBUM-ID, ALBUM-NAME, PRODUCED-BY, YEAR)
SONGS(SONG-ID, SONG-START, DURATION, ALBUM-ID)

Update the year of the album with name 'SUHANA RATH' to 2018.
Delete the album 'YADON KI BAARISH' along with all the songs in it.

b) Briefly discuss 3NF and BCNF with suitable real examples. (4)

PART E

Answer any four full questions, each carries 10 marks.

15 a) Distinguish between *denseindex* and *sparseindex* and give examples for each. (4)

b) Consider an EMPLOYEE file with 10000 records where each record is of size 80 bytes. The file is sorted on employee number (15 bytes long), which is the primary key. Assuming un-spanned organization, block size of 512 bytes and block pointer size of 5 bytes, compute the number of block accesses needed for retrieving an employee record based on employee number if (i) No index is used (ii)Multi-level primary index is used (6)

16 a) How is a B-Tree structurally different from a B+-Tree? (3)

b) Write an explanatory note on *clustering index* by quoting an example. (4)

c) Show *two non-canonical query trees* for the following relational algebra expression: (3)

$$\pi_{ROLLNO,\ CID}\ (COURSE \underset{COURSE.CID=ENROLL.CNO}{\bowtie} ENROLL \underset{EROLL.ROLL\#=STUDENT.ROLLNO}{\bowtie} STUDENT))$$

17 Consider three tables COURSE(CNO,CNAME,CREDITS), STUDENT(ROLLNO,NAME,ADDRESS,SEM) and ENROLLMENT (CNO,ROLLNO,GRADE). Foreign keys have the *same* name as primary keys. Identify one initial canonical query tree for the following SQL expression and show the steps to optimize it using heuristics. Assume that CNAME is a candidate key. (10)

SELECT S.NAME, S.ADDRESS, E.GRADE FROM COURSE C, STUDENT S, ENROLLMENT E WHERE S.ROLLNO =E.ROLLNO AND C.CNO = E.CNO AND CNAME='PDBD'

18 a) Determine if the following schedule is recoverable. Is the schedule cascade-less? Justify your answer. (4)

r1(X), r2(Z), r1(Z), r3(X), r3(Y), w1(X), c1, w3(Y), c3, r2(Y), w2(Z), w2(Y), c2

(Note: ri(X)/wi(X) means transaction Ti issues read/write on item X; ci means transaction Ti commits.)

b) Discuss the four ACID properties and their importance. (6)

19 a) Argue that *two-phase locking* protocol ensures serializability. (2)

b) Illustrate how *conflict serializability* is checked using *precedence graph* with the help of an example. (4)

c) What is wrong with the following concurrent schedule? What is the solution? (4)

T1 issues read lock on Y and reads Y, T2 issues read lock X and reads X, T1 issues write lock on X, T2 issues write lock on Y.

20 a) Write explanatory notes on the following: (7)

i) GIS ii) Big Data

b) How does RDF support semantic web technology? (3)

****