

BITS, Pilani – Dubai
International Academic City, Dubai
Second Semester 2007 – 08


Course Number & Title	CS UC 352 – Database Systems
Component Name	Comprehensive Exam (Closed Book)
Weightage	40 %
Duration	3 hours
Date	2-06-2008
Max. Marks	40 Marks

Note:-1. Answer all questions.


1. Consider a transaction T initiated at a home computer that asks bank B to transfer \$10,000 from an account at B to another bank C.
- What are the components of distributed transaction T? What should the components at B and C do?
 - What can go wrong if there is not \$10,000 in the account at B?
 - What can happen if one or both computers crash or the network is disconnected? 3M

2. Consider the tree shown in the Figure Describe the steps involved in executing each of the following operations according to the tree-index concurrency control algorithm, in terms of the order in which nodes are locked, unlocked, read, and written. Be specific about the kind of lock obtained and answer each part independently of the others, always starting with the tree shown in the Figure.

- Search for data entry 40*.
- Search for all data entries k^* with $k \leq 40$.
- Insert data entry 62*.
- Insert data entry 40*.
- Insert data entries 62* and 75*. 3M



3. Answer the questions based on the transaction given



- What is the value of the LSN stored in the master log record?
- What is done during Analysis?
- Show the log when recovery is complete, including all non-null prevLSN and undonextLSN values in log records. 4M

4. With the help of clear figures show the working of the Extensible hash tables 3M

5. What does swizzling mean? Mention the different types of swizzling. 3M

6. Consider the relation for published books BOOK (Book Title, Author Name, Book Type, List Price, Author Affiliation, Publisher). Suppose the following dependencies exist.

Book Title → { Publisher, Book Type }
 Book Type → Book Price
 Author Name → Author Affiliation

- What normal form is the relation in?
- Apply normalization until you cannot decompose the relation further, state the reason behind each decomposition 3M

7 Notown Records has decided to store information about musicians who perform on its albums (as well as other company data) in a database. The

company has wisely chosen to hire you as a database designer (at your usual consulting fee of \$2500/day).

Each musician that records at Notown has an SSN, a name, an address, and a phone number. Poorly paid musicians often share the same address, and no address has more than one phone. Each instrument used in songs recorded at Notown has a unique identification number, a name (e.g., guitar, synthesizer, flute) and a musical key (e.g., C, B-flat, E-flat).

Each album recorded on the Notown label has a unique identification number, a title, a copyright date, a format (e.g., CD or MC), and an album identifier. Each song recorded at Notown has a title and an author. Each musician may play several instruments, and a given instrument may be played by several musicians. Each album has a number of songs on it, but no song may appear on more than one album. Each song is performed by one or more musicians, and a musician may perform a number of songs. Each album has exactly one musician who acts as its producer. A musician may produce several albums, of course.

Design a conceptual schema for Notown and draw an ER diagram for your schema.

The preceding information describes the situation that the Notown database must model. Be sure to indicate all key and cardinality constraints and any assumptions you make. Identify any constraints you are unable to capture in the ER diagram and briefly explain why you could not express them. 3M

8. Consider the following relations:

Student(snum: integer, sname: string, major: string, level: string, age: integer)

Class(name: string, meets at: string, room: string, fid: integer)

Enrolled(snum: integer, cname: string)

Faculty(fid: integer, fname: string, deptid: integer)

The meaning of these relations is straightforward; for example, Enrolled has one record per student-class pair such that the student is enrolled in the class. Write the following queries in SQL. No duplicates should be printed in any of the answers.

- a. Find the names of all Juniors (level = JR) who are enrolled in a class taught by I. Teach.
- b. Find the age of the oldest student who is either a History major or enrolled in a course taught by I. Teach.
- c. Find the names of all classes that either meet in room R128 or have five or more students enrolled.

3 X 2 = 6M

9. Write short notes on

- a. What is a transaction? In what ways is it different from an ordinary program (in a language such as C)?
- b. What is the phantom problem? Can it occur in a database where the set of database objects is fixed and only the values of objects can be changed?

(PTO)

- c. Use of Quad trees for multidimensional data storage.
- d. Document Type definition used in XML.
- e. Explain the term schema in databases.
- f. Use of bitmap indexes

2X 6 = 12M

***** ALL THE BEST

BITS, Pilani – Dubai
Second Semester 2007 – 08

Course Number & Title	CS UC 352 – Database Systems
Component Name	Test 2 (Open Book)
Weightage	20 %
Duration	50mins
Date	13-04-2008
Max. Marks	20 Marks

Note:-1. Answer all questions

Consider the following relational schema. Underlined keys indicate primary keys.

EMPLOYEE

FNAME	MINIT	LNAME	<u>SSN</u>	BDATE	ADDRESS	SEX	SALARY	SUPERSSN	DNUM
-------	-------	-------	------------	-------	---------	-----	--------	----------	------

DEPARTMENT

DNAME	<u>DNUM</u>	MGRSSN	MGRSTDATE
-------	-------------	--------	-----------

PROJECT

PNAME	<u>PNUM</u>	PLOCATION	DNUM
-------	-------------	-----------	------

WORKS_ON

<u>SSN</u>	<u>PNUM</u>	HOURS
------------	-------------	-------

DEPENDENT


<u>SSN</u>	<u>DEPENDENT_NAME</u>	SEX	BDATE	RELATIONSHIP
------------	-----------------------	-----	-------	--------------

Answer the following queries in relational algebra

- a. Retrieve the names of employees who work on all the projects that 'Pradeep Shukla' works on. 2M
 - b. Make a list of project numbers that involve an employee whose last name is 'Reddy' either as a worker or as a manager of the department that controls the project. 2M
 - c. List the names of employees who have atleast one dependent. 2M
2. Answer the following queries in SQL.

- Retrieve the birthdate and address of the employee whose name is 'Neha B Sarkar'. 1M
- List all employees whose address is in Delhi, India. 1M
- List the names of managers who have atleast one dependent. (answer as a correlated query) 3M
- Find the sum of salaries of all employees of the HR department, also find the maximum salary, minimum salary and average salary. 2M

3. Consider the B+ tree index shown in the Figure below. Each intermediate node can hold up to five pointers and four key values. Each leaf can hold up to four records, and leaf nodes point to the sibling node although these links are not shown in the figure. Answer the following questions.



- Name all the tree nodes that must be fetched to answer the following query: "Get all records with search key greater than 38." 1M
- Show the B+ tree that would result from inserting a record with search key 109 into the tree. 2M
- Show the B+ tree that would result from deleting the record with search key 81 from the original tree. 2M

4. Give details about how websites like Google organize search of websites.

2M

BITS, Pilani – Dubai
Second Semester 2007 – 08

Course Number & Title	CS UC 352 – Database Systems
Component Name	Test 1 (Closed Book)
Weightage	25 %
Duration	50mins
Date	2-03-2008
Max. Marks	25 Marks

Note:-1. Answer all questions.

1. What are the roles of a Database Administrator 2M
2. What are database languages and how are they used ? 3M
3. Explain the following terms giving clear examples
a. Arity
b. Aggregation
c. Overlap Constraint
d. one – many relationship 2 X 4 = 8M

4. 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, *pcttime*: integer)

Dept(did: integer, *dname*: string, *budget*: real, *managerid*: integer)

- a. Give an example of a foreign key constraint that involves the Dept relation, give SQL statements. What are the options for enforcing this constraint when a user attempts to delete a Dept tuple? 2M
- b. Write the SQL statements required to create the preceding relations, including appropriate versions of all primary and foreign key integrity constraints. 3M
- c. Define the Dept relation in SQL so that every department is guaranteed to have a manager. 1M

5. 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. 3M

6. Consider the following relation CAR SALE (car number, date sold, salesman number, commission percentage, discount amount). Assume that a car may be sold by multiple salesmen hence car number and salesman number together are the primary key. Additional dependencies are a. date sold → discount amount and b. salesman number → commission % . Based on the given primary key is this relation in 1NF, 2NF or 3NF. Why or why not? How would you successively normalize it completely? 3M

*****BEST OF LUCK*****