

Comprehensive Examination Question Paper
BITS, Pilani – Dubai Campus, DIAC
III Year SECOND Semester 2010-2011

Degree: B.E.(Hons.) Branch: C.S.

Course No : CSC 352 Course Title: Database Systems

Date: 08/06/11 Wednesday Time: 3 hours Total marks: 80 Weight-age: 40%

Data provided are complete. *Closed Book.*

This question paper has 4 pages.

PART A: (write in PART A Answer booklet)

Answer all Questions

A.1. List the different ways in which the type declaration system of a language such as Java or C++ differs from the data definition language used in a database. 3M

A.2. Suppose you want to build a video site similar to YouTube. Consider each of the points listed below, as disadvantages of keeping data in a file-processing system. Discuss the relevance of each of these points to the storage of actual video data, and to metadata about the video, such as title, the user who uploaded it, tags, and which users viewed it.

- a. Data redundancy and inconsistency
- b. Difficulty in accessing data
- c. Data isolation
- d. Integrity problems
- e. Atomicity problems
- f. Concurrent users
- g. Security

4M

PTO

A.3. The purpose of this form shown below is to record the medication given to a particular patient at the *Wellmeadows Hospital*. Examine the data shown on the form. *Discuss* how the data shown in this form could be represented in *tables* to represent *functional dependencies*. *Label* all *transitive* and *partial* dependencies.

Wellmeadows Hospital Patient Medication Form							
		Patient Number : P10034					
Full Name :		Robert MacDonald		Ward Number :		Ward 11	
Bed Number :		84		Ward Name :		Orthopaedic	
Drug Number	Name	Description	Dosage	Method of Admin	Units per Day	Start Date	Finish Date
10223	Morphine	Pain Killer	10mg/ml	Oral	50	24/03/96	24/04/96
10334	Tetracycline	Antibiotic	0.5mg/ml	IV	10	24/03/96	17/04/96
10223	Morphine	Pain Killer	10mg/ml	Oral	10	25/04/96	02/05/96

- (a) Is this table in 3NF? Explain why or why not ?
- (b) Describe and illustrate the process of normalizing the data shown in this table to third normal form (3NF).
- (c) Identify the primary, candidate and foreign keys in your 3NF relations. (1+3+2M)

A.4. i) Given the following descriptions, create an appropriate ER diagram for each of the specified relationships *separately*.

(a) Each company operates four departments, and each department belongs to one company.

(b) Each department in part (a) employs one or more employees, and every employee is employed by one department.

(c) Each of the employees in part (b) may or may not have one or more dependants, and each dependant belongs to one employee.

(d) Each employee in part (c) may or may not have an employment history.

- ii) Represent all of the relationships described in Question 4.i) above as a single ER diagram and convert the same to the relational model. (4+2+2M)

PTO

A.5 Given below are three tables, answer the following queries using relational algebra.

PLAYER(PlayerCode, PlayerName, Instrument, Age, OrchestraCode)

ORCHESTRA(OrchestraCode, OrchestraName, Town, Conductor)

EVENT(Event#, EventDate, Venue, OrchestraCode)

- a. Get a list of player names with the town their orchestra is based in.
- b. Get a list of venues and event dates where the conductor is 'P Mancini'.
- c. Get a list of player names and orchestra names with venues and event dates that the 'Omagh Brass Ensemble' which is the name of the orchestra is playing.

(3X2=6M)

A.6 i) Consider the following relational tables , underlined attributes indicate the primary key.

Customer(customer_name, customer_street, customer_city)

Branch(branch_name, branch_city, assets)

Account(account_number, branch_name, balance)

Depositor(customer_name, account_number)

Answer the following questions in SQL

- a. Find the names of all branches that have assets greater than at least one branch located at Brooklyn
- b. Find the average balance for each customer who lives in Harrison and has atleast three accounts.
- c. Find a branch whose average balance is greater than or equal to all average balances.

2X3=6M

ii) Create a trigger to make a copy of all the records stored into a table called Staff, the table has the fields (staffno, name, position, sex, age, salary, branch-no), the new table created is called Staff audit.

(3M)

A.7. Write a brief technical note on each of the following terms w.r.t. mechanics of disks, with a suitable diagram:

- a) Disk Assembly b) Head Assembly c) Cylinder d) Sector.

[4 M]

P.T.O

PART B: (write in PART B Answer booklet)

- B.1) In a distributed database, when is it useful to have
i) Replication of data. ii) Fragmentation of data ? [4 M]
- B.2) Explain in brief the following w.r.t. temporal query languages:
i) SNAPSHOT ii) TEMPORAL SELECTION
iii) TEMPORAL PROJECTION iv) TEMPORAL JOIN [5 M]
- B.3) Write a brief Technical Note on Geographical Data (Raster Data and Vector data) with a suitable example in each category. [5 M]
- B.4.) Illustrate with an example, Hash-Based Indexing in Databases. [5 M]
- B.5) Explain with an example each of the following Data Elements:
a) Record with Variable Length Fields
b) Record with Repeating Fields [5 M]
- B.6) Draw a **Tree** showing the *Query Evaluation Plan*, for the following query:
(the order in which the basic operations are performed)
SELECT S.sname
FROM Reserves R, Sailors S
WHERE R.sid = S.sid
AND R.bid = 200 AND S.rating > 4 [6 M]
- B.7) Explain the following in brief w.r.t. TRANSACTION MANAGEMENT in database systems:
a) ATOMICITY
b) CONSISTENCY
c) ISOLATION
d) DURABILITY
e) *Serializable schedule* [10 M]

TEST II Question Paper

BITS, Pilani – Dubai Campus, Academic City, Dubai.

III Year SECOND SEMESTER 2010-2011

Degree: B.E. (Hons.) Branch: C.S.

Course No : CS C352 Course Title: Database Systems

Date: 01, May, 2011 Sunday Time: 50 min. Total marks: 40

Data provided are complete. **OPEN Book.**

Text / Reference Book and hand written class notes permitted.

This question paper has 2 pages.

Answer all Questions.

1. A patient record consists of the following fixed length fields:
The patient's *date of birth*, *social security number* and the *patient ID*, each of **10 bytes** long.
It has the following variable length fields: *name*, *address* and *patient history*.
If each of the *pointers* within a *record* requires **4 bytes** and the *record length* is a **4 byte** integer,
- How many bytes**, exclusive of the space needed for the variable length fields, are needed for the record? You may assume that no alignment of fields is required.
 - Assume that the variable length fields *name*, *address* and *patient history*, each have a length that is uniformly distributed. For the *name*, the range is **10-50** bytes; for *address* it is **20-80** bytes and for history it is **0-1000** bytes. What is the average length of the patient record?

[10 M]

2) As a Database Designer, You are required to create an INDEX-ORGANIZED File HASHED on IDNO for the given below student data. The hashing function is defined as follows:

$$h(\text{IDNO}) = \text{IDNO} \text{ MOD } 11.$$

Consider a suitable alternative for storing the data entry in an index.

(here you have to **draw the diagram only**).

[10 M]

STUDENT DATA

<u>IDNO</u>	<u>CGPA</u>	<u>IDNO</u>	<u>CGPA</u>
15	7.6	18	9.2
10	9.2	20	9.4
19	8.6	25	9.5
7	7.1		
12	7.7		
17	7.3		
23	8.5		
4	7.9		
9	7.2		
11	7.3		
13	7.5		
16	7.6		

3. Consider the following relational tables

Job_current(contact_id, title, salary, start_date)

My_contacts(contact_id, last_name, first_name, phone, email, city)

Job_desired(contact_id, title, salary_low, salary_high, years_exp)

Note: title here refers to job title. Ex: web developer

Write SQL queries for each of the following

- Get the average salary of a Web Developer from the Job_current table.
- Get the first name, last name, and salary of all Web Developers .
- Display how much more or less than the average, each web developer makes .
- Display titles in alphabetical order from Job_current and Job_desired tables.
- Add an extra field gender to the table my_contacts and ensure that the value of this field is either 'M' or 'F'. (1+2+3+2+2)

4. Given two relational tables Child and Toys

Child

Toys

Child id	Child	Toy id
1	Jane	3
2	Sally	4
3	Cindy	1

Toy id	toy
1	Toy soldier
2	Huala hoop
3	Barbie
4	Tinker toys
5	Pokemon cards
6	Skipping rope

Give the output table obtained after executing the following query.

```
SELECT C.CHILD, T.TOY
FROM TOYS T
LEFT OUTER JOIN CHILD C
ON C.TOY_ID = T.TOY_ID;
```

(4M)

5. Two tables are created using the given SQL statements.

```
CREATE TABLE TodaysRecords(TodaysMaxBarometricPressure FLOAT,
TodaysMinBarometricPressure FLOAT)
```

```
CREATE TABLE OurCitysRecords(RecordMaxBarometricPressure FLOAT,
RecordMinBarometricPressure FLOAT)
```

Write the SQL command to **create a trigger** which is activated when the reading TodaysMaxBarometricPressure in the TodaysRecords table exceeds the RecordMaxBarometricPressure in the OurCitysRecords and it updates the value of the field RecordMaxBarometricPressure in the OurCitysRecords. (6M)

BITS Pilani, Dubai Campus, Academic City, Dubai.

III Year SECOND SEMESTER 2010-2011

Degree: B.E. (Hons.) Branch: C.S.

TEST I Question Paper

Course No : CS C352 Course Title: Database Systems

Date: 10, March, 2011 Thursday Time: 50 min. Total marks: 50

Data provided are complete. **Closed Book.**

This question paper has 2 pages.

Answer all Questions.

1. Suppose we have a relation $R(A,B,C)$ with a multi valued dependency $A \twoheadrightarrow B$. The tuples are of the form (a,b,c) , where b is any of b_1, b_2 , or b_3 , and c is any of c_1, c_2 , or c_3 . If we know that the tuples (a, b_1, c_1) , (a, b_2, c_2) and (a, b_3, c_3) are in the current instance of R , **what other tuples** do you know must also be in R ? [6 M]

2. For the given relation schema and MVDs, **Decompose** the relation into a collection of relation schemas in **4NF**:

$R(A, B, C, D)$ with multi valued dependencies $A \twoheadrightarrow B$ and $A \twoheadrightarrow C$. [6 M]

3. Normalize **each** one of the following Relation Schemas (FD's for each are also given) into **3NF relations**:

i) $R(A, B, C, D)$
FDs: $C \rightarrow D$, $BC \rightarrow A$

ii) $R(W, X, Y, Z)$
FDs: $WZ \rightarrow XYZ$ [5 M]

4. Given the following **Functional Dependencies**:

$X \rightarrow Y$, $X \rightarrow W$, $WY \rightarrow Z$,

SHOW that the functional dependency $X \rightarrow Z$ can be **IMPLIED**
From the given FDs using **ARMSTRONG's axiom(s)**. [3 M]

P.T.O.

5. A General Hospital consists of a number of specialized wards (such as Maternity, Pediatric, Oncology, etc). Each ward hosts a number of patients, who were admitted on the recommendation of their own GP and confirmed by a consultant employed by the Hospital. On admission, the personal details of every patient are recorded. A separate register is to be held to store the information of the tests undertaken and the results of a prescribed treatment. A number of tests may be conducted for each patient. Each patient is assigned to one leading consultant but may be examined by another doctor, if required. Doctors are specialists in some branch of medicine and may be leading consultants for a number of patients, not necessarily from the same ward.

Represent all the above requirements in a neatly drawn **ER diagram**. **Convert the above ER diagram to the relational model** and *show all the obtained tables with their attributes*, indicate the **primary key** of each table. (8 + 8 = 16M)

6. With an example explain a **weak entity set** in an ER diagram. (4M)

7. With reference to the data stored in the database, **explain in brief** the following terms

i) **Work of the buffer manager** ii) **Metadata** iii) **Indexes** (3 X 2 = 6M)

8. Write in brief a technical note on the following terms:

i) **Data Definition Language** ii) **Data Manipulation Language** (2 X 2 = 4M)

QUIZ I

Course No : CS C352 Course Title: Data Base Systems

Date: 6, Apr., 2011 Wednesday Time: 20 min. Total marks: = 16

Weightage: 8% Venue : seating arrangement **Closed Book**.

This question paper has 2 pages [use **back page** for **rough work** only]

IDNO:

Name:

SET A

Write answers in the space provided in question paper. Answer all questions.

I. Consider a supplier-parts database. The database has three relations, with the schema listed below.

Supplier (*SupplierName, City*)

Part (*PartName, Colour, Weight*)

Shipment (*SupplierName, PartName, Date*)

Formulate the following queries in **relational algebra**. [1+1+2+2]

1. What are the names of all suppliers?
2. What are the names of suppliers located in Paris?
3. What are the names of parts that are red and weigh more than 20 kilos?
4. What are the names of parts that are neither red nor weigh more than 20 kilos?

II. Briefly define the following term w.r.t. Operations on disks:
Seek Time

[2 M]

QUIZ I

Course No : CS C352 Course Title: Data Base Systems

Date: 6, Apr., 2011 Wednesday Time: 20 min. Total marks: = 16

Weightage: 8% Venue : seating arrangement *Closed Book*.

This question paper has 2 pages [use *back page* for *rough work* only]

IDNO:

Name:

SET A

III. Briefly mention the steps in MODIFYING a BLOCK on DISK. [3M]

IV Consider the following relations

Employee(*person_name*, *street*, *city*)

Company(*company_name*, *city*)

Works(*person_name*, *company_name*, *salary*)

Manages(*person_name*, *manager_name*)

Give the expressions in relational algebra for each of the queries

1. Name the employees who work for First Bank Corporation. 1M
2. Find the names and cities of all employees who work for First Bank Corporation. 2M
3. Find the names of employees who do not work for First Bank Corporation. 2M