Reg. No.



I SEMESTER M.TECH. (CSE / CSIS)

END SEMESTER EXAMINATIONS, NOV/DEC 2016

SUBJECT: ADVANCED DATABASE SYSTEMS [CSE5102] REVISED CREDIT SYSTEM (26/11/2016)

Time: 3 Hours

MAX. MARKS: 50

3M

3M

3M

Instructions to Candidates:

✤ Answer ALL the questions.

✤ Missing data may be suitable assumed.

- **1A.** Explain the role of the following constructs in PL/SQL with an example. i. EXCEPTION ii. CURSOR iii. TRIGGER
- **1B.** Consider the following Sales relational schema:
 - Customer(<u>cid:integer</u>, cname: string, rating: integer, age:integer) Shop(Sid:integer, Location:string, District:string) Item(<u>lid:integer</u>, Iname: string, <u>Cost: integer</u>, Brand: string) Buys(cid: integer, Sid:integer, Iid: integer)

Write the following queries in SQL:

- i. Find the name of the customers who have bought an item from a shop in Udupi district.
- ii. Find the Shop which has sold every item in 'Item table' to one or another customer.
- iii. Find all the Districts where the total no. of shops is greater than the average of the total no. of shops at all districts **4M**
- **1C.** Give an algorithm for Block Nested-loop join. What is the worst case and best case cost estimate? How the performance of the algorithm can be further improved? **3M**
- **2A.** How to estimate the size of following operations?
 - i. Complex selection with conjunction / disjunction ii. Join
- **2B.** Illustrate Nesting and Unnesting constructs available in ORDB.

2C. Consider a object relational database schema with a relation 'Student' whose attributes are as shown below, with types specified for multivalued attributes.
 Student = (<u>SId</u>, Name, SubjectSet multiset(Subject), Location(City, District, State), Batch(StartYear, EndYear), Dept(DName, Institute, University), CGPA))
 Subject = (Suld, SuName, Type, Credits, MarksObtained)
 Define the above schema in SQL 2003 with appropriate types for each attribute

- i. Define the above schema in SQL:2003, with appropriate types for each attribute.
- ii. Using the above schema, write the following queries in SQL:2003
 - a. List the name of all 2010(StartYear) batch UP state students.
 - b. List the name of 'MIT' students who have scored 100 in 'Advanced Database Systems'
 - c. Get the average marks obtained in 4 credit, Core 'type' subjects by the student with SID, 100 **4M**
- **3A.** What is speed-up and scale-up? What are the ideal values for these in parallel database **3M** systems? Practically what is achievable and why? Explain.

3B. Explain fragment and replicate join algorithm. Using this algorithm, how can you optimize the evaluation if the join condition for r and s is of the form | r. A − s . B | ≤ k, where k is a small constant?

3M

3C. Explain the majority protocol used in distributed lock manager approach. It has got potential for deadlock even with single item. Justify. Under what condition, quorum consensus protocol works like majority protocol? How the majority protocol can be can be modified to work even if some sites are unavailable?

4M

4A. Build a star schema for the following Bank relational schema, with Balance(fixed), avg_Balance as the measures and branch and customer as dimensions. Define the resultant star schema in DMQL.



Starting with base cuboid what OLAP Operations should be performed

- i. To find the total balance of the customers staying in 'Bangalore' city.
- ii. To find the total balance in 'South India' region, if we extend the branch dimension with an hierarchy branch_city \rightarrow District \rightarrow State \rightarrow Region **3M**
- **4B.** Illustrate the extended aggregate constructs i. pivot and ii. group by cube. How the 'null' can be replaced by 'all' in group by cube using grouping() and decode()?
- **4C.** Considering the following University DTD.

<!DOCTYPE university [

<!ELEMENT university ((department|course|instructor|teaches)+)>

- <!ELEMENT department (dept name, building, budget)>
- <!ELEMENT course (course id, title, dept name, credits)>
- <!ELEMENT instructor (IID, name, dept name, salary)>
- <!ELEMENT teaches (IID, course id)>
- <!ELEMENT dept name(#PCDATA)>
- <!ELEMENT building(#PCDATA)>
- <!ELEMENT budget(#PCDATA)>
- <!ELEMENT course id (#PCDATA)>
- <!ELEMENT title (#PCDATA)>
- <!ELEMENT credits(#PCDATA)>
- <!ELEMENT IID(#PCDATA)>
- <!ELEMENT name(#PCDATA)>
- <!ELEMENT salary(#PCDATA)>
-]>
- i. Write an XPATH expression to get instructors teaching more than 2 courses
- ii. Write an Xquery to join instructor, teaches and course elements
- iii. Write an Xquery to find the total salary, across all instructors, at each department
- iv. Give the equivalent DTD for Unversity using ID and IDREF to specify the relationship between department course and instructor course.
 5A. Give an overview of execution of MapReduce program. How it handles node failure?
- How to implement grouping and aggregation operation using map reduce?
- **5B.** Illustrate the constructs for the following in the context of MongoDB: **3M**
 - i. Aggregation ii. Mapreduce iii. Cursor
- **5C.** Create a column family to store basic information of Students such as Roll No, Student Name, Student Date of Birth and Student Address with two to three student records.
 - i. Alter the table to include the subject preferences and hobbies of each student. There should be a minimum of two subject preferences and a maximum of four. The order of preferences as given by the student should be preserved. The hobbies as given by the student should be arranged in alphabetical order.
 - ii. Give a query to replace the first subject preference by 'Big Data and Analytics' for Roll No=1.
 - iii. List the name of students whose subject preferences include 'Machine Learning'

3M