

THE OPEN UNIVERSITY OF SRI LANKA
DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING

FINAL EXAMINATION 2015

BACHELOR OF SOFTWARE ENGINEERING– LEVEL 5

ECI5266 – ADVANCED DATABASE SYSTEMS



DATE: SEPTEMBER 16, 2015


TIME: 0930-1230 HRS

This question paper consists of eight questions. Answer **FIVE** questions **ONLY**.

Question 1 – Database Programming

(20 marks)

Consider the following schema:



The *Employee* relation contains information about employees in the ABC. The *Department* relation consists of information of all departments. The *project* relation consists of information related to projects conducted by each department. The *works_on* relation contains information about assignment of employees for the projects. Dependants of employees are recorded in the *Dependent* table.

- (i.) Write a SQL query to list the name of the department and its manager's first name for those departments that manage a total budget of Rs. 100,000.
(7 marks)
- (ii.) Write a SQL statement to create a view (called *v_projectInfo*) that displays the project name, budget of the project and name of the department that needs the budget.
(4 marks)
- (iii.) Write a trigger (called *tr_checkRecentMothers*) to ensure that tech lead of a project does not assign mothers having babies less than one year old.
(9 marks)

Question 2 – Object Relational Databases**(20 marks)**

Consider the following object relational schema for a University database:

Object types:

course_t (cid: integer, title: varchar(15), credits_req: integer)

unit_t (uid: integer, credits: integer)

offering_t (unit: ref unit_t, semester: number(1), year: number(4))

enrolled_t (unitoffer: ref offering_t, mark: integer)

enrolled_list table of enrolled_t

student_t (sid: integer, name: varchar(15), phone: varchar(10), course: ref course_t, enrolments: enrolled_list)

Tables:

Courses of course_t (cid primary key)

Units of unit_t (uid primary key)

Offered of offering_t (unit not null references units)

Students of student_t (sid primary key, course references courses)

Nested table enrolments store as enrolledlist_ntab

- The Courses table of course_t has attributes of course id (cid), title, and credit points required. It contains tuples for all courses.
 - Units table of unit_t contains tuples for all units, and has the attributes unit number (uid), and number of credits.
 - The Offered table of offering_t records the information of units offered in a particular semester and consists of attributes for unit reference, semester and year.
 - The Students table of student_t contains student id, name, phone, course enrolled and the list of enrolled units (as an attribute of nested table type) for each student. Each tuple of the nested table contains a reference to unit offering, and the mark obtained by the student.
 - The attribute types are specified in the type descriptions above, as also are the primary keys and referential constraints in the table schema.
- (a) Write Oracle OBJECT SQL statements for the following queries (use columns of REF type instead of joins to link tables):
- i. Find the name and course title of students who have completed the total credits required for their course (i.e., the total credits of units in which the student obtained at least 50 marks is greater than or equal to the credits required).
(4 marks)
 - ii. For students who are enrolled in units offered during 2015, find the name of the student, course title and the average mark for the units taken in 2015.
(3 marks)

- (b) Add a new enrolment for student with *sid* 1256432 in the unit with *uid* 302636 in semester 2 of year 2015. Assume that the unit offering already exists in the *Offered* table and the given student also exists in the *students* table with other enrolments. The mark should be entered as null.

(4 marks)

- (c) It is required to add a member method called *credcomp* to calculate the number of credits completed by a student. Only units where student received a mark of at least 50 are counted as completed.

Write Oracle SQL statements to modify the object type *student_t* by adding this method specification.

(6 marks)

- (d) Using the method defined above, write an Oracle SQL statement to display for each course, the course title and the number of students who have completed the required credits for their courses.

(3 marks)

Question 3 – Relational Algebra

(20 marks)

Consider the following relational database schema of a university marks database:

Student (studentNo, studentName, gpa, level)



Marks (studentNo, subjectNo, grade)



Subject (subjectNo, subjectName, credits)

Student relation contains information of students. *Subject* relation contains information of subjects offered in the university. *Marks* relation contains the grades obtained by students for different subjects. A student who has completed a subject will have a tuple in the *Marks* relation stating the *studentNo* and *subjectNo* of completed subjects.

Write the **relational algebra** expressions for the following queries:

- (a.) List the names of students who have earned a “B” grade for “Software Engineering.

(5 marks)

- (b.) Print name and number of students who have not completed at least one subject.

(5 marks)

- (c.) Print the student numbers of those who have already completed 6 subjects.

(10 marks)

Question 4 – Disk and Files**(20 marks)**

- a) Explain the terms related to disk access; *Seek time, Rotational delay, Data transfer time* (3 marks)
- b) Consider a disk with the following characteristics: block size 512 bytes, number of blocks per track 20, number of tracks per surface 400. A disk consists of 15 double-sided platters.
- i. What is the total capacity of a track in Kbytes? (2 marks)
 - ii. How many cylinders are there? (1 mark)
 - iii. What is the total capacity of a cylinder in Kbytes? (2 marks)
 - iv. What is the total capacity of disk in Kbytes? (2 marks)
 - v. Suppose the disk drive rotates the platters at a speed of 2400 rpm (revolutions per minute); what is the average rotational delay and data transfer time for a disk block in msec? (2 marks)
 - vi. Suppose the average seek time is 30 msec. How much time does it take (on the average) in msec to locate and transfer a single block given its block address? (2 marks)
 - vii. Calculate the average time it would take to transfer 20 random blocks and compare it with the time it would take to transfer 20 consecutive blocks in same track. (2 marks)
- c) Briefly describe the following record formats giving an example for each; *Fixed length record, Variable length record* (4 marks)

Question 5 – XML Databases

(20 marks)

- (a) List and briefly describe each of the components of a FLWOR expression.

(4 marks)

Consider the following XML document for the questions given below:

```
<dept bldg = "101">
  <employee id="135">
    <name>Ramesh</name>
    <phone>0714123123</phone>
    <office>344</office>
  </employee>
  <employee id="137">
    <name>Prabash</name>
    <phone>0777123123</phone>
    <office>345</office>
  </employee>
</dept>
```

Note that the above XML document is stored in *Department*(*id int, emp xml*) table created in MS SQL Server and it contains only a single record.

- (b) What is the output of the following XPath expression:

```
SELECT emp.query('//employee[office="344" or @id="135"]/name')
FROM department
```

(3 marks)

- (c) What would be the output of the following XPath expression?

```
SELECT emp.query('/dept/employee[2]/name')
FROM department
```

(3 marks)

- (d) What would be the output of the following XQuery expression?

```
SELECT emp.query('
  for $d in /dept
  let $emp:=$d//employee/name
  where $d/@bldg > 95
  order by $d/@bldg
  return
  <EmpList>{$d/@bldg, $emp}</EmpList>')
FROM department
```

(5 marks)

- (e) Write an expression to insert the following employee after employee id 135 in the above XML document.

```
<employee id="137">
  <name>Saman</name>
  <phone>0755123123</phone>
  <office>532</office>
</employee>
```

(5 marks)

Question 6 – Indexing and Query Processing

(20 marks)

- (a) Why does it cost more to use an unclustered index than a clustered index to retrieve tuples that satisfy a range selection? (4 marks)
- (b) Briefly describe what happens in external merge sort in the first two passes. (4 marks)
- (c) Consider a relation named *parts* with primary key *pid* and 100,000 records stored 10 records per block or page. If 1% of records are randomly accessed by *pid* value every day, compare the costs in number of disk accesses between a hash index and a B+ tree index on *pid*. (6 marks)
- (d) Assume that we want to join two relations *Units* and *Enrolments* with the following schema and associated statistics (relation *Students* is not used for this question):
- Students(sno, sname, address, gpa) 3000 pages with 10 tuple per page.
- Units(unno, cname, credits, dept) 400 pages with 10 tuples per page.
- Enrolments(unno, sno, grade) 4,000 pages with 100 tuples per page.

If there were no indexes on the two tables, what would be the cost of a *sort merge join*? Assume that there are 5 buffer pages to sort each table.

(6 marks)

Question 7 – Transactions and Concurrency Control

(20 marks)

- (a.) Briefly explain the properties of a transaction. (4 marks)
- (b.) Briefly explain the rules in Strict 2 Phase Locking Protocol. (3 marks)
- (c.) Explain why DBMSs interleave actions of multiple transactions. (2 marks)

(d.) Briefly explain the terms *deadlock prevention* and *deadlock detection*. Explain approaches for each methodology.

(3 marks)

(e.) Consider the following part of the schedule.

T_1	T_2	T_3	T_4
S(A)			
R(A)			
	X(B)		
	W(B)		
S(B)			
			S(D)
			R(D)
	X(D)		
		S(C)	
		R(C)	
			X(A)

Assume that Transaction T_i has a higher priority than transaction T_{i+1} (i.e. transaction T_1 has higher priority than T_2 ; T_2 has higher priority than T_3 ; and T_3 has higher priority than T_4).

i. Draw a **Wait-For-Graph** for the given schedule.

(2 marks)

ii. List the steps for deadlocks detection and the approach to break the deadlock after identifying?

(2 mark)


iii. Draw the above schedule again considering deadlock prevention algorithm: *Wait-Die approach*

(4 marks)

Question 8 – Crash Recovery

(20 marks)

(a.) Consider the diagram below:



What is the desired state that the database should be in after Crash Recovery? Briefly explain.

(5 marks)

(b.) What is meant by Write-Ahead Logging. Why is it necessary to implement this protocol in a DBMS?

(5 marks)

(c.) What are the roles of Analysis, Redo and Undo phases in ARIES?

(5 marks)

(d.) What are the differences between update log records and CLRs?

(5 marks)