

**UCSC****University of Colombo, Sri Lanka***University of Colombo School of Computing***DEGREE OF BACHELOR OF INFORMATION TECHNOLOGY
(EXTERNAL)**Academic Year 2025 — 3rd Year Examination — Semester 6**IT6405 (R) — Database Systems II
(Repeat Paper)***Structured Question Paper*
(2 Hours)**To be completed by the candidate****Index Number**

--	--	--	--	--	--	--

Important Instructions

- The duration of the paper is **2 hours**.
- The medium of instructions and questions is English. Students should answer in the medium of English language only.
- This paper has **4 questions** on **12 pages**. Answer **all** questions.
- All questions carry **equal** marks.
- Write your answers **only on the space provided** on this question paper.
- Do not tear off any part of this question paper. Under no circumstances may this paper (or any part of this paper), used or unused, be removed from the Examination Hall by a candidate.
- Note that questions appear on both sides of the paper. If a page or part of a page is not printed, please inform the supervisor/invigilator immediately.
- Any electronic device capable of storing and retrieving text, including electronic dictionaries, smartwatches, mobile phones and Programmable calculators are **not allowed**.
- Non- Programmable Calculators are **allowed**.
- *All Rights Reserved*. This question paper can NOT be used without proper permission from the University of Colombo School of Computing.

**To be completed by
the examiners**

1	
2	
3	
4	
Total	

- 1) (a) Write the SQL query to create the Student table with the following constraints.

Student (StudentId, Name, Address, Telephone, Age)

- StudentId is a number used as the primary key
- Name cannot be null
- Default Address to be 'Colombo'
- Telephone specifies a unique number
- Age should be a number between 5-18

[6 marks]

ANSWER IN THIS BOX

- (b) Consider the following relational schema to answer the question below.

Employee (eno, ename, hiredate, salary, saldifference)

Write a trigger to calculate and update the *saldifference* field (difference between the old and the new salary) when the salary is updated.

[6 marks]

ANSWER IN THIS BOX

- (c) Write down **two (02)** differences between triggers and stored procedures.

[4 marks]

ANSWER IN THIS BOX

- (d) Suppose relation R1 (A,B) has tuples {(a,b), (a,b), (c,d)}, and relation R2 (B,C) has tuples {(b,e), (b,e), (d,f), (g,h)}. Consider the following SQL query.

```
SELECT * FROM R1 RIGHT OUTER JOIN R2 ON R1.B = R2.B;
```

How many tuples are in the result of the above SQL query?

[4 marks]

ANSWER IN THIS BOX

- (e) Consider the relations Employee and Evaluation as given below used to keep track of employees, and their job evaluation records.

Employee (Empid, First_name, Job_id, Salary, Department_id)

Evaluation (Empid, Evaluation_date , Job_id, Department_id, Score)

Write down a trigger named ***Salary_Increment*** to increase Salary of an employee as given below (based on the Score value) upon insertion of his/her evaluation record into the Evaluation table.

- 15% Salary increment if Score is ≥ 9
- 10% Salary increment if Score ≥ 7 and Score < 9

[5 marks]

ANSWER IN THIS BOX

- 2) (a) Suppose that we have an ordered file with $r = 30,000$ records stored in a disk with block size $B = 1024$ bytes. File records are of fixed size and unspanned, with record size $R = 100$ bytes.

- (i) Calculate the blocking factor (BFR).

[2 marks]

ANSWER IN THIS BOX

- (ii) How many block accesses are required to search a record in the data file using the binary search?

[2 marks]

ANSWER IN THIS BOX

- (iii) Suppose that the ordering key field of the file is $V = 9$ bytes long, a block pointer is $P = 6$ bytes long, and we have constructed a primary index for the file. How many block accesses are required to search a record using the primary index?

[4 marks]

ANSWER IN THIS BOX

- (b) “*Selection of certain indexes may cause excessive overhead on the Database Management System and lead to poor performance.*” Is this statement correct or incorrect? Justify your answer.

[4 marks]

ANSWER IN THIS BOX

- (c) Briefly describe the *Primary Index* and the *Secondary Index* types.

[4 marks]

ANSWER IN THIS BOX

- (d) Illustrate the flow of *transaction states* using a diagram.

[6 marks]

ANSWER IN THIS BOX

- (e) Briefly explain using a suitable example, the *Dirty Read problem* that can occur due to the concurrent execution of transactions.

[3 marks]

ANSWER IN THIS BOX

- 3) (a) Consider the following schedule S, consisting of three concurrent transactions T1, T2, and T3. Note that each R and W denotes read, and write operations respectively.

T1	T2	T3
		R(Y)
		R(Z)
R(X)		
W(X)		
		W(Y)
		W(Z)
	R(Z)	
R(Y)		
W(Y)		
	R(Y)	
	W(Y)	
	R(X)	
	W(X)	

- (i) Draw the precedence graph for S.

[5 marks]

ANSWER IN THIS BOX

- (ii) Is S conflict serializable? Justify your answer.

[3 marks]

ANSWER IN THIS BOX

- (iii) Is S view serializable? Justify your answer.

[3 marks]

ANSWER IN THIS BOX

- (b) Consider the following T1 and T2 transactions. Note that each Lock-s and Lock-x denotes shared lock and exclusive lock respectively.

T1	T2
Lock-s(A)	Lock-s(A)
Read(A)	Read (A)
Lock-x(B)	Lock-x(B)
Unlock(A)	Read(B)
Read(B)	Write(B)
Write(B)	Commit
Commit	Unlock(B)
Unlock(B)	Unlock(A)

- (i) What is the concurrency control protocol used in transaction T1?

[3 marks]

ANSWER IN THIS BOX

- (ii) What is the concurrency control protocol used in transaction T2?

[3 marks]

ANSWER IN THIS BOX

- (c) Write two (02) limitations of *Two Phase Locking protocol*.

[4 marks]

ANSWER IN THIS BOX

- (d) Briefly explain the *Wait-Die* deadlock prevention protocol.

[4 marks]

ANSWER IN THIS BOX

- 4 (a) Construct a B-tree of order three by inserting the following eight key values sequentially. Draw the final B-tree after all insertions are completed.

Key values: 8, 5, 1, 7, 3, 12, 9, 6

[5 marks]

ANSWER IN THIS BOX

- (b) State two (02) disadvantages of *distributed databases*.

[4 marks]

ANSWER IN THIS BOX

- (c) Draw the architectural diagram for a distributed database management system and identify its key components.

[7 marks]

ANSWER IN THIS BOX

- (d) Consider the following relational schema.

Employee (EmpId, Fname, Lname, Address, Gender, Salary, Bdate)

Project (Pnumber, Pname, location)

Works_on (Pno, Empid)

Consider the following SQL query.

```
Select E.Lname
From Employee E, Works_on W, Project P
Where P.Pname = 'Synergy' And
P.Pnumber= W.Pno And E.EmpId=W.Empid
And E.Bdate > '2004-08-24';
```

- (i) Draw the initial query tree for the SQL query given above.

[4 marks]

ANSWER IN THIS BOX

(ii)

Draw the optimized query tree for the SQL query given above.

[5 Marks]


