



University of Ruhuna- Faculty of Technology
Bachelor of Information & Communication Technology Honours Degree
Level 03 (Semester II) Examination, November/December 2025
Academic year 2023/2024

Course Unit: ICT3212 - Advanced Database Management Systems (Written)
(New Curriculum) Duration: 02 hours

.....
This Theory examination paper contains **04 pages** including this instruction page.

IMPORTANT INSTRUCTIONS:

1. The medium of this examination is **English**.
2. This is a Closed Book examination.
3. This Examination consists of **four (04) questions** that are given equal marks.
4. You must answer **all four (04) questions** in this examination.
5. Nonprogrammable calculators are allowed for this Examination.

Question 01

- a) Briefly define the concept of **file organization** in the context of databases. [15 Marks]
- b)
- Write down *two (02)* advantages of **heap file** organization. [10 Marks]
 - Write down *two (02)* disadvantages of **sequential file** organization. [10 Marks]
- c) Explain briefly **main drawback** of the **Static Hashing**. [10 Marks]
- d) List down *three (03)* properties of a good hash function. [15 Marks]
- e) A hash table of *size 10* uses linear probing to resolve collisions. Insert the following eight keys into the hash table in the given order using the following hash function. [24 Marks]
Keys: 23, 45, 12, 39, 27, 32, 52, 29
The hash function: $h(\text{key}) = \text{key} \bmod 10$
- f) List down *two (02)* main disadvantages of the Extendable Hashing. [16 Marks]

Question 02

- a) Why is the concept of **indexing** important in the context of the database systems? [10 Marks]
- b) State whether the following statements are *true or false*. [30 Marks]
- Any field of the file can be used to create an index.
 - The index typically stores each value of the indexed field along with a list of pointers to all disk blocks that contain records with that field value.
 - Secondary indexes can be specified on any non-ordering field.
 - There is only one secondary index for a file.
 - Clustering indexes specified on non-ordering field of a file.
- c) Differentiate between **B-tree** and **B+ tree** with respect to their *node structure* and the *height of the tree*. [20 Marks]
- d) Suppose that we have an ordered file with $r = 45,000$ records stored on a disk with block size $B = 1024$ bytes. File records are of fixed size and are unspanned, with record length $R = 16$ bytes, and there are no indexing on the file. [40 Marks]

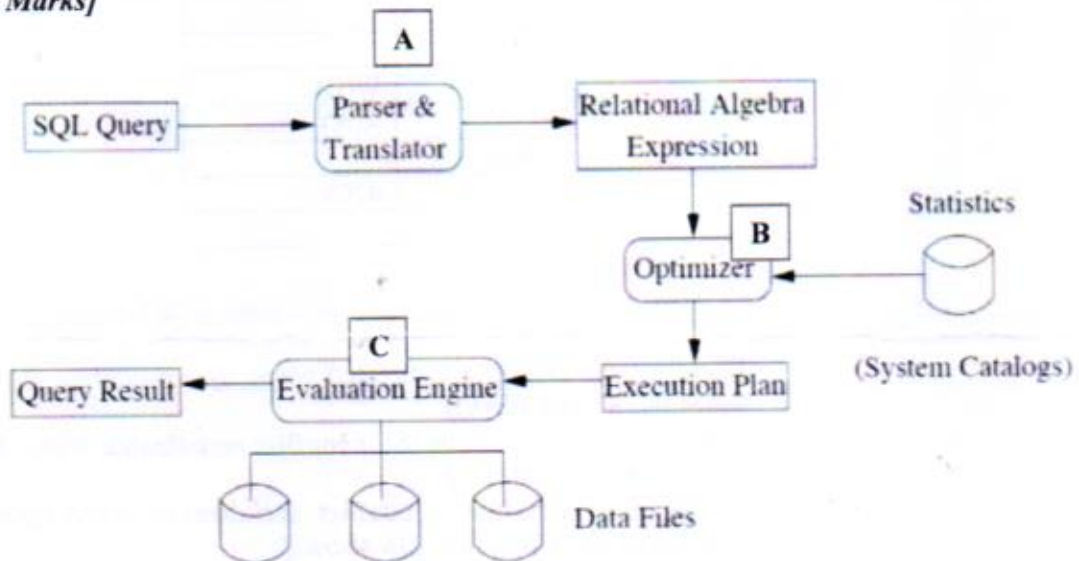
Note:

- For Blocking factor: Round your answers to lower limit integer value, example: $5.6 = 5$.
 - For Number of Blocks accesses / required: Round your answers to upper limit integer value, example: $5.6 = 6$.
- Calculate the Blocking Factor (bfr)? [10 Marks]
 - How many blocks are required to store the above data file? [10 Marks]

- iii) Suppose if we perform a binary search on the data file how many blocks access is required to search a record on the data file. [10 Marks]
- iv) Now suppose that the ordering key field of the file is $V = 5$ bytes long, a block pointer is $P = 3$ bytes long, and we have constructed a primary index for the file.
 - a) Calculate the blocking factor for the index file (bfr_i)? [05 Marks]
 - b) How many blocks are required to store the above primary index file? [05 Marks]

Question 03

- a) Briefly define what is meant by **query optimization**? [10 Marks]
- b) Consider the Diagram Given below, and explain the role of each unit named **A**, **B**, and **C**. [20 Marks]



- c) Briefly explain how a **hash join** reduces the number of comparisons during the join process. [10 Marks]
- d) Mention **one (01)** difference between **Materialized evaluations** and **Pipeline evaluations** in Query optimization. [10 Marks].
- e) Consider two database tables: Table T1 has 4000 records and occupies 40 disk blocks, and Table T2 has 3000 records and occupies 30 disk blocks. These two tables have to be joined as per a specified join condition that need to be evaluated for every pair of records from these two tables. Assume that no indexes are available on either table. If nested loop join algorithm is employed to perform this join:
 - i) How many blocks access are required in the worst case? [10 Marks]
 - ii) How many seeks operations are required in the worst case? [10 Marks]
 - iii) How many blocks access are required in the best case? [05 Marks]
 - iv) How many seeks operations are required in the best case? [05 Marks]

Question 04

- a) Briefly define what is meant by a **transaction** in the context of the database systems. *[10 Marks]*
- b) Describe the properties of a transaction. *[20 Marks]*
- c) Consider the **three (03)** transactions T1, T2, T3 and the schedule 'S' given below. *[40 Marks]*

T1	T2	T3
R(A)		
W(A)		
	R(A)	
	W(B)	
R(C)		
	R(C)	
		W(B)
		R(A)
	W(C)	
		R(C)

Schedule S

- i. Construct the precedence directed graph by considering the given schedule 'S'. *[10 Marks]*
 - ii. State whether the schedule 'S' is **conflict serializable** or not. Justify your answer. *[10 Marks]*
 - iii. If the given schedule is **conflict serializable**, write down the equivalent serial schedule. *[10 Marks]*
 - iv. State whether the schedule 'S' is **view serializable** or not. Justify your answer. *[10 Marks]*
- d) Briefly explain the following types of locking mechanisms. *[14 Marks]*
- i. Binary Locks
 - ii. Shared/exclusive Locks.
- e) Undo and Redo are two operations that performs under database recovery. Briefly explain how each of these operations handles the recovery process of a transaction. *[16 Marks]*

*****End of the Paper*****