

# King Fahd University of Petroleum and Minerals

Information and Computer Science Department

ICS 334: Database Systems, Section 01

## FINAL EXAM

**Instructor:** Dr. Salahadin Mohammed

---

---

<b>ID:</b>						
------------	--	--	--	--	--	--

<b>Name:</b>	
--------------	--

<b>Grades</b>		
<b>Section</b>	<b>Max</b>	<b>Scored</b>
<b>A</b>	<b>5</b>	
<b>B</b>	<b>30</b>	
<b>C</b>	<b>41</b>	
<b>D</b>	<b>24</b>	
<b>TOTAL</b>	<b>100</b>	

**Section A – True False (5 marks - 1 mark per question)**

**Circle the correct answer.**

1. Primary key index can NOT be a clustered index
  - a. TRUE
  - b. FASLE
  
2. Hash files contain both index blocks and data blocks
  - a. TRUE
  - b. FASLE
  
3. The rules of tuning queries are NOT applicable to views
  - a. TRUE
  - b. FALSE
  
4. Binary search is appropriate for sequential ordered files
  - a. TRUE
  - b. FALSE
  
5. RAID 4 is of better performance than RAID 5
  - a. TRUE
  - b. FALSE

**Section B - Multiple Choice (30 marks – 2.5 marks per question)**

In the following multiple choice questions, more than one choice may be correct. Circle all correct answers.

1. Informal measures of quality for relational schema design are:
  - a. Semantics of the relation attributes
  - b. Having a good DBA and writing efficient SQL
  - c. Reducing the null values in tuples
  - d. Having primary keys in each relation
  - e. Disallowing generating spurious tuples
  
2. Improper grouping of attributes into a relational schema can result in
  - a. Storage wastage
  - b. Update anomaly
  - c. Fast update
  - d. Fewer nulls
  - e. None of the above
  
3. What are the problems with null values
  - a. Save space in a disk
  - b. Problems in understanding the meaning of attributes
  - c. They have two interpretations
  - d. Problems in specifying join operations
  - e. None of the above
  
4. Objective of normalization is to help
  - a. Solve problems associated with redundancy
  - b. Minimize the number of relations
  - c. Identify various types of update anomalies
  - d. The always improve performance
  - e. None of the above

5. Which of the followings is/are NOT true
- A table can be in the third normal form but not in the second normal form.
  - For a relation with one candidate key, a third normal form is also in BCNF
  - A relation is in BCNF if and only if every determinant is a candidate key
  - Violation on BCNF can only occur in a relation with 2 or more composite keys.
  - None of the above
6. Which of the following is/are TRUE about foreign keys
- A foreign key is always a primary key in the owner relation.
  - A foreign key can be null
  - A foreign key must be unique and not necessarily a primary key in the owner relation.
  - A foreign key must be part of the primary key
  - A foreign key must be unique
7. Which of the following is/are NOT goals of a database design
- Provide natural and easy to understanding of the information
  - Satisfy the data requirements of user or application
  - Support processing requirements and performance objectives
  - Save disk space
  - None of the above
8. A database system handles a user query through a number of phases. Which of the following is/are NOT part of the phases
- Logical processing of the user requests
  - Transformation of user request to an access plan
  - Data transfer
  - Correction of user queries
  - None of the above
9. Which of the following is NOT true about heap organization
- Records are placed in the file in the order in which they are inserted.
  - Binary search is the most appropriate
  - A new record is placed at the end of the file
  - Delayed and immediate deletion of records is possible
  - None of the above

10. Which of the following is TRUE about indices

- a. Secondary index can be used for primary key
- b. Secondary index can be clustering index
- c. Primary index is always dense index
- d. Secondary index can be sparse index
- e. Primary index can NOT be multi-level index

11. The goal of database tuning can include

- a. To make an application run faster
- b. To lower the response of database queries
- c. To prevent unauthorized users
- d. To check query syntax
- e. None of the above

12. Which of the following is NOT one of the basic properties of a transaction?

- a. Consistency preservation
- b. Atomicity
- c. Durability
- d. Isolation
- e. None of the above

**Section C – Answer all the following short questions**

1. Explain why we need less number of disks in RAID 3 than in RAID 2. **[5 points]**
2. Why do we have categories? Give an example which shows their benefits. **[6 points]**
3. In the context of concurrency control, explain the incorrect analysis problem. **[6 points]**

4. Give two reasons of why we need to tune indexes **[6 points]**
5. What is the use of vertical and horizontal partitioning **[6 points]**
6. In RAID 5 how many disk accesses do you need to update one record?  
**[6 points]**
7. Why creating a primary index takes more time than creating a secondary index.  
**[6 points]**

## Section D: Answer the following long questions

1. Consider the following student lines relational schema. (Assume a referential integrity constraint between attributes of the same name)

CNO	Cname	<b>Course</b>
-----	-------	---------------

LNO	Lname	<b>Lecturer</b>
-----	-------	-----------------

SNO	CNO	LNO	<b>Offering</b>
-----	-----	-----	-----------------

SNO	Name	DNO	<b>STU</b>
-----	------	-----	------------

DNO	Dname	<b>Dept</b>
-----	-------	-------------

- a. Indicate the keys in each relation. **[4 points]**
  
- b. Write an SQL query that retrieves the SNO of the ICS students who are taking ICS 334 with Ejaz. **[4 points]**
  
- c. Step-by-step optimize the SQL that you wrote in answering question 1.b starting with the canonical tree. In each step mention the heuristic rule (strategy) used to draw the respected query tree. **[12 points]**
  
- d. Write the optimized query in relational Algebra. **[4 points]**