

**KING FAHD UNIVERSITY OF PETROLEUM AND MINERALS**  
**Department of Information and Computer Science**

**063**  
**ICS 424-01: Advanced Database Systems**  
**MAJOR EXAM 2**

**Time Allowed: 90 Minutes**

Student ID: \_\_\_\_\_

Name: \_\_\_\_\_

Instructor: Muhammad Waheed Aslam

<b>Problem</b>	<b>Points</b>	<b>Score</b>
<b>Q1: Query Processing</b>	<b>50</b>	
<b>Q2: Transaction Processing</b>	<b>34</b>	
<b>Q3: Concurrency Control</b>	<b>30</b>	
<b>Q4: Recovery</b>	<b>50</b>	
<b>Q5: SQL</b>	<b>36</b>	
<b>Total</b>	<b>200</b>	

**Notes:**

- Please skim through all the questions, make sure that you understand them, and then attempt to answer them with a time-allocation in mind. If any question is not clear, get it clarified during the first fifteen minutes.
- If you need to make any assumptions, please document them as part of your answers.
- There are five questions in this exam each focusing on a distinct aspect. You are expected to answer all of them.

*August 4, 2007*

**1. [Query Processing – (50 points)]**

- a) List typical steps to process a high level query? (5 Points)

Consider the query from Company database “*Find First name of the employees born after 1980 work on a project named ‘Alpha’*”.

- b) Write query in SQL. (5 Points)

- c) Draw initial query tree for the query then using Heuristic Algebraic Optimization Algorithm, optimize the query (Show all steps of your work) (40 Points)

Q3 Cont...

**2. [Transaction processing (40 points)]**

a) Explain the lost update problem with the help of practical example.

(10 points)

b) What is a schedule? Explain serial, nonserial, and serializable schedules.

(16 Points)

c) Why is a serial schedule considered correct? Why is a serializable schedule considered correct? (8 Points)

**3. [Concurrency Control (30 points)]**

a) What is the two-phase locking protocol? Briefly describe its phases.

(10 Points)

b) Below find a schedule S of two transaction  $T_1$  and  $T_2$  that exhibit the lost-update problem. Use locking and the 2PL to rewrite the schedule S as another schedule S' that does not have this problem. Verify your answer by filling the column  $bal_x$  in the two schedules S and S' (20 points)

<b>S</b>			
Time	$T_1$	$T_2$	$bal_x$
$t_1$		begin_transactio	<b>550</b>
$t_2$	begin_transactio	read ( $bal_x$ )	
$t_3$	read ( $bal_x$ )	$bal_x = bal_x + 100$	
$t_4$	$bal_x = bal_x - 50$	write ( $bal_x$ )	
$t_5$	write ( $bal_x$ )	committ	
$t_6$	committ		

<b>S'</b>			
Time	$T_1$	$T_2$	$Bal_x$
		begin_transactio	<b>550</b>
	begin_transactio		

**4. [Recovery – (50 points)]**

a) Discuss the UNDO and REDO operations and the recovery technique that each use. (10 Points)

b) How can recovery handle transaction operations that do not affect the database such as printing of reports by the transaction? (10 Points)

c) Consider the following four transactions:

(30 Points)

**Transaction T<sub>1</sub>:** read-item(A); read-item(D); write-item(D);

**Transaction T<sub>2</sub>:** read-item(B); write-item(B); read-item(D); write-item(D);

**Transaction T<sub>3</sub>:** read-item(A); write-item(A); read-item(C); write-item(C);

**Transaction T<sub>4</sub>:** read-item(B); write-item(B); read-item(A); write-item(A);

Following is a log corresponding to a particular schedule at the point of a system crash.

[start-transaction, T<sub>1</sub>] [read-item, T<sub>1</sub>, A] [read-item, T<sub>1</sub>, D] [write-item, T<sub>1</sub>, D, 20] [commit, T<sub>1</sub>] [**checkpoint**] [start-transaction, T<sub>2</sub>] [read-item, T<sub>2</sub>, B] [write-item, T<sub>2</sub>, B, 12] [start-transaction, T<sub>4</sub>] [read-item, T<sub>4</sub>, B] [write-item, T<sub>4</sub>, B, 15] [start-transaction, T<sub>3</sub>] [write-item, T<sub>3</sub>, A, 30] [read-item, T<sub>4</sub>, A] [write-item, T<sub>4</sub>, A, 20] [commit, T<sub>4</sub>] [read-item, T<sub>2</sub>, D] [write-item, T<sub>2</sub>, D, 25] **System crash**

Suppose that we use the *immediate update protocol* with checkpointing. Describe the recovery process from the system crash. Specify which transactions are rolled back, which operations in the log are redone and which (if any) are undone, and whether any cascading rollback takes place.

**5. [SQL – (36 points)]**

Given the following relations:

**Hotel**

<u>Hotel-No</u>	Name	Address
H1	Meridian	Khobar
H2	Sheraton	Dammam

**Room**

<u>Room-No</u>	<u>Hotel-No</u>	Type	Price
R1	H1	Single	300
R2	H1	Double	400
R3	H1	Single	350
R1	H2	Double	500
R2	H2	Single	350
R3	H2	Double	400

**Booking**

<u>Hotel-No</u>	Guest-No	Date-From	Date-To	Room-No
H2	G1	Nov. 20	Nov. 22	R1
H2	G3	Jan. 1	Jan. 5	R2
H1	G4	Jan 3	Jan 5	R1
H1	G2	Dec. 21	Dec. 30	R3

**Guest**

Guest-No	Name	Address
G1	Ahmad Ali	Jeddah
G2	Mohammad Kamel	Taif
G3	Ali Barakat	Mekka
G4	Said Khaled	Madina

a) write the appropriate SQL statements to answer the queries below: **(17 points)**

i) Find all single rooms that costs more than SR300.

ii) Find the names and addresses of all guests who reserved rooms starting from or after 'Jan. 5'(Jan is in next year,ie. Nov, Dec and then Jan)



iii) Find the type and price of the room that is reserved by 'Mohammad Kamel'.

b) Using SQL, create a view, **GUEST\_VIEW**, which will display the names of the guests occupying the rooms, with the room number, in the Meridian Hotel. **(8 points)**

c) Insert following records in to HOTEL table

**(6 points)**

H3	Hilton	Dhahran
H4	Intercontinental	Jubail

d) Delete the record of the guest, who belongs to Taif

**(4 points)**