

Final Exam for Distributed Database Systems (2012)

1. (20 marks) Suppose we have three transactions T1, T2 and T3 scheduled by the locking based method as follows:

	T1	T2	T3
(1)	R1(x)		
(2)			R3(y)
(3)			W3(y)
(4)	R1(y)		
(5)			W3(x)
(6)		R2(y)	
(7)			C3
(8)	W1(x)		
(9)		R2(x)	
(10)		W2(x)	
(11)	C1		
(12)		C2	

- (i) Tell if this schedule is serializable or non-serializable and give your explanation.
- (ii) If in the schedule W3(x) of T3 is changed into R3(x), is this changed schedule serializable or non-serializable? Why?

2. (20 marks) Transactions T1 and T2 have following schedule in a database

	T1	T2
(1)	R1(x)	
(2)	R1(y)	
(3)		R2(y)
(4)	W1(x)	
(5)		R2(z)
(6)	C1	
(7)		W2(y)
(8)		C2

Suppose the interaction between the Local Recovery Manager and the Buffer Manager is No-Fix/Flush, tell what actions should be taken for T1 and T2 in following cases when the database system restarts

- (i) The database system fails at step (4) after executing W1(x).
- (ii) The database system fails at step (7) after executing W2(y).

3. (20 marks) In a distributed transaction commit following 3PC protocol, three sites -- one coordinator C and two participant sites P1 and P2, are involved. To commit a transaction T, C first sends the PREPARE message to P1 and P2, and both P1 and P2 reply to C by a VOTE-COMMIT message. So C decides to globally commit T and sends a PREPARE-TO-COMMIT message to P1 and P2. However, C fails after sending the message to P1 but before sending the message to P2 while P1 and P2 work normally.
- (i) Tell what status that P1 and P2 would remain when C fails as indicated.
- (ii) For transaction T, what actions should be taken for P1 and P2 to terminate and what actions should be taken for C to recover?
4. (20 marks) In the following diagram, the left column shows a list of events which are either from user or from system failures, and the right column shows a list of actions taken by LRM (Local Recovery Manager) to deal with such events; connect the box of every event in the left column with the box of corresponding command for dealing with it in the right column.

EVENTS		LRM COMMANDS
Abort	<input type="checkbox"/>	<input type="checkbox"/> Fetch
Commit	<input type="checkbox"/>	<input type="checkbox"/> Flush
SQL Query	<input type="checkbox"/>	<input type="checkbox"/> Partial undo
Operating System Failure	<input type="checkbox"/>	<input type="checkbox"/> Global redo
Hardware Failure	<input type="checkbox"/>	<input type="checkbox"/> Global undo

5. (20 marks) We are now in a Big Data Era, where data possesses high volume, velocity (streaming), variety, and uncertainty characteristics and multiple users concurrently access databases.
- (i) Which kinds of data management technologies do you think important in dealing with Big Data?
- (ii) Can you identify a specific technique developed by the traditional database field which is not applicable any more or needs some further improvement?