

Homework 7: Transactions, Logging and Recovery
(due April 20th, 2016, 4:00pm, in class—hard-copy please)

Reminders:

- a. Out of 100 points. Contains 4 pages.
- b. Rough time-estimates: 4-6 hours.
- c. Please type your answers. Illegible handwriting may get no points, at the discretion of the grader. Only drawings may be hand-drawn, as long as they are neat and legible.
- d. There could be more than one correct answer. We shall accept them all.
- e. Whenever you are making an assumption, please state it clearly.
- f. Lead TA for this HW: Sorour Amiri.

Q1. Serializability [30 points]

Consider the following schedules (S1 and S2). The actions are listed in the order they are scheduled, and prefixed with the transaction name.

S1: T1:R(A), T1:W(A), T3:R(A), T3:W(A), T2:R(A), T1:R(B), T3:R(B), T1:W(B),
T3:W(B), T2:R(B)

S2: T1:R(A), T1:W(B), T2:R(B), T2:W(C), T3:R(C), T3:W(A)

Q1.1. (15 points) Consider the schedule S1.

- A. (5 points) Draw the precedence graph for S1.
- B. (5 points) Is S1 a conflict serializable schedule? If yes, what is the equivalent serial schedule? If no, explain in 1-2 lines.
- C. (5 points) From the following options, choose which of the anomalies is present in S1 (if any). Explain in 1-2 lines which actions cause that anomaly.
 - a. Dirty read (WR Conflict)
 - b. Unrepeatable read (RW Conflict)
 - c. Lost update (WW Conflict)

Q1.2. (15 points) Consider the schedule S2.

- A. (5 points) Draw the precedence graph for S2.
- B. (5 points) Is S2 a conflict serializable schedule? If yes, what is the equivalent serial schedule? If no, explain in 1-2 lines.

- C. (5 points) From the following options, choose which of the anomalies is present in S2 (if any). Explain in 1-2 lines which actions cause that anomaly.
- Dirty read (WR Conflict)
 - Unrepeatable read (RW Conflict)
 - Lost update (WW Conflict)

Note: The anomalies due to interleaved execution of transactions are explained on slide 38-44 of the lecture 16 or page 526-529 of the textbook. The concept of conflict serializable schedules is explained on slide 46-47 of lecture 16 or page 550-551. The precedence graphs are explained on slide 51-52 of the lecture 16 or page 550-551 of the textbook.

Q2. Locking Protocols [20 points]

Consider the following schedules. The actions are listed in the order they are scheduled, and prefixed with the transaction name.

S1: T1:R(X), T2:W(Y), T2:R(X), T1:W(Y), T1:Commit, T2:Commit

S2: T1:R(X), T1:R(Y), T1:R(Z), T1:Commit, T2:W(X), T2:Commit, T3:W(Y), T3:W(Z), T3:Commit

S3: T1:R(X), T1:W(X), T3:R(X), T3:W(X), T2:R(X), T1:R(Y), T3:R(Y), T1:W(Y), T3:W(Y), T2:R(Y)

Q2.1. (6 points) Consider schedule S1.

- (2 points) Write the schedule S1 in a table format and draw the precedence graph.
- (4 points) Will the actions of S1 be allowed by the following concurrency control protocols?
 - 2PL
 - Strict 2PL
 If YES, show in the table form where the lock requests can happen; If NO, explain briefly in 1-2 lines.

Q2.2. (7 points) Consider schedule S2.

- (3 points) Write the schedule S2 in a table format and draw the precedence graph.
- (4 points) Will the actions of S2 be allowed by the following concurrency control protocols?
 - 2PL
 - Strict 2PL
 If YES, show in the table form where the lock requests can happen; If NO, explain briefly in 1-2 lines.

- Q2.3. (7 points) Consider schedule S3.
- A. (3 points) Write the schedule S3 in a table format and draw the precedence graph.
 - B. (4 points) Will the actions of S3 be allowed by the following concurrency control protocols?
 - a. 2PL
 - b. Strict 2PL
 If YES, show in the table form where the lock requests can happen; If NO, explain briefly in 1-2 lines.

Note: By table format we mean the style given in slide 28-30 in lecture 17. "2PL" and "Strict 2PL" protocols are explained on slide 19-34 of the lecture 17 or textbook page 550-552. "Lock Management" which explains how the lock requests are granted is explained on slide 36-41 of the lecture 17 or on page 553-554 of the textbook.

Q3. Deadlock Management [20 points]

Consider the following sequence of actions, listed in the order it is submitted to the DBMS (S is a shared lock, X is an exclusive lock):

S1: T1:S(A), T2:X(A), T3:X(B), T1:X(B), T3:S(A)

S2: T1:S(A), T2:X(A), T3:X(B), T3:X(A), T1:S(B)

For S1 and S2 as given above, answer the following questions:

- Q3.1. (4 points) For S1, write whether lock requests of its actions will be granted or blocked by the lock manager.
- Q3.2. (4 points) Draw the waits-for graphs for S1 and write if the schedule will result in a deadlock condition. If there is no deadlock condition, write the order of completion of the schedule. Explain in 1-2 lines.
- Q3.3. (6 points) For S2, write whether the lock requests of its actions will be granted or blocked by the lock manager.
- Q3.4. (6 points) Draw the waits-for graphs for S2 and write if the schedule will result in a deadlock condition. If there is no deadlock condition, write the order of completion of the schedule. Explain in 1-2 lines.

Note: The implementation of lock and unlock requests is given in detail on slide 35-41 of lecture 17 or on page 554 of the textbook. The waits-for graphs are explained on slide 43 of lecture 17 or on page 556 of the textbook.

Q4. Logging [18 points]

Consider following Log table.

LSN	Log entry	prevLSN	undonextLSN
00	Begin checkpoint		
10	Update: T1 writes P1		
20	Update: T2 writes P2		
30	Update: T3 writes P3		
40	T2 commit		
50	Update: T3 write P2		
60	End checkpoint		
70	T2 end		
80	Update: T1 writes P5		
90	Update: T4 write P4		
100	update: T3 writes P5		
110	Update: T4 write P4		
120	Update: T4 write P2		
130	T3 abort		

Q4.1. (5 points) Write the appropriate values in the prevLSN and undonextLSN columns.

Note: The Log record and CLR are explained in detail on page 582-585 of the textbook and also in Lecture 19.

Q4.2. (5 points) Describe the actions taken to rollback transaction T3.

Q4.3. (8 points) Show the log after T3 is rolled back, including all the prevLSN and undonextLSN values in the log records.

Q5. Crashing and Recovery [10 points]

Consider the same log as given in Q4 above, but now with a CRASH after LSN 130. The recovery manager now sees the log (of course with the correct prevLSN and undonextLSN values filled in) after the CRASH.

Q5.1 (1 point) What is the redo set?

Q5.2 (1 point) What is the undo set?

Q5.3 (8 points) Show the content of the transaction table and the dirty page table at the end of the analysis phase.

Note: Assume that both transaction table and dirty page table are empty at the beginning of the log.

Q6. Anagrams [2 points]

What is 'these churn air' an anagram for?