

**CAS CS 460: Introduction to Database Systems – Fall 2020 – Written Assignment #7**

Due: 12/06/20 11:59 PM in Gradescope

**Problem 1 [60 pts]**

- a. For each of the following schedules, draw the precedence graph. Explain whether or not the schedule is conflict serializable. If the schedule is conflict serializable, give one possible equivalent serial schedule.

Note that,  $R_i$  denotes transaction  $i$  reads an item, while  $W_i$  denotes a write.

- $R_1(A) - W_1(A) - R_2(B) - R_2(C) - W_3(A) - W_2(B) - R_4(B) - R_4(A) - R_4(E) - R_2(A) - R_3(B) - W_1(B)$
- $R_1(A) - R_4(B) - W_1(B) - W_3(A) - R_2(B) - R_2(A) - R_4(C) - R_4(B) - R_2(A) - R_3(C)$

- b. Assume we have a strict 2PL protocol. For each of the following schedules, insert the appropriate locks (both shared and exclusive). Additionally, explain what happens to the scheduler as each schedule is being executed.

Note that, if a transaction blocks because of an operation, the transaction with the next operation in the schedule will continue. When a transaction unblocks, it resumes its operations.

If you have a deadlock, you need to choose a transaction to abort and follow all the proper protocols to allow the rest of the schedule to continue. Additionally, aborted transactions will be required to restart again at some point along the schedule.

Lastly write the *actual* executed schedule at the end.

- $R_1(A) - R_2(A) - W_3(C) - R_1(B) - W_2(A) - W_3(B)$
- $R_1(A) - R_2(B) - R_3(C) - R_1(B) - R_2(C) - R_3(A) - W_1(A) - W_2(B) - W_3(D)$

**Problem 2 [40 pts]**

a. Consider the log below

LSN	LOG
00	update: T1 writes P2
10	update: T1 writes P3
20	update: T1 writes P1
30	update: T1 writes P2
40	begin_checkpoint
45	end_checkpoint
50	update: T3 writes P3
60	T1 commit
70	update: T3 writes P4
80	T1 end
90	update: T2 writes P2
100	T2 commit
	CRASH, RESTART

In this log, we store information about 3 transactions. After the log record with LSN 100, the system crashes and then we restart. We use the ARIES recovery algorithm discussed in Chapter 18 in the book. Based on that, answer the following questions:

- i) What is done during the Analysis phase?
- ii) What is done during the Redo phase?
- iii) What is done during the Undo phase?"
- iv) Show the log when recovery is complete, including all non-null prevLSN and undonextLSN values in the log records.