Assignment 8

Reading –Read Sections 8.5-end of Chapter 8 of the textbook, plus Chapter 4.

Problem 1

Consider a system that uses LSN-based operation logging with fuzzy checkpoints and logging of undo's, as described in the lecture slides and in Chapter 8 of the textbook. Assume that there is no analysis pass before recovery (so there is no dirty page table in the checkpoint record), and that CLR's (i.e. undo records) are treated as normal updates and do *not* splice out a portion of the log (as shown on lecture slide 35).

The following sequence of records is found in a database log after a system failure. Data values, such as before and after images, are omitted. The notation P1/r1 means record r1 on page P1.

LSN	Trans	Operation Type	Page/Record	Trans backpointer
1000	T0	Update	P0/r0	null
1100	T1	Update	P1/r1	null
1200	T1	Update	P2/r2	1100
1300	checkpoint log record			
	Active transactions: [T0, LSN 1000], [T1, LSN 1200]			
1400	T1	Update	P3/r3	1200
1500	T2	Update	P1/r4	null
1600	T2	Update	P4/r5	1500
1700	T3	Update	P4/r6	null
1700	checkpoint log record			
	Active transactions: [T1, LSN 1400], [T2, LSN 1600], [T3, LSN 1700]			
1800	T1	CLR	P3/r3	1400
1900	T1	CLR	P2/r2	1800
2000	T1	CLR	P1/r1	1900
2100	T1	End Abort		
2200	T3	Commit		
2300	T0	Commit		

Answer each of the following. In each case, explain briefly why it's the right answer.

a. Show the log records that must be written by the recovery process, in the proper order, and briefly explain why they must be written. The new log records should have LSNs numbered sequentially starting with 2300.
b. What LSNs are on pages P1, P2 and P3 after recovery?
c. Looking only at the log and not the database state, which updates and CLRs *might* have to be redone?
d. What pages are fetched from disk by the recovery process?
e. Does the log give you enough information to tell whether record-level or page-level lock granularity is being used? If so, which is it and how can you tell? If not, explain why not.

Now suppose we modify the example so that it uses an analysis pass.
Each checkpoint record now includes a dirty page table as follows:

- LSN 1300: Dirty Page Table: [P1, 1100], [P2, 1200]
- LSN 1700: Dirty Page Table: [P1, 1500], [P3, 1400], [P4, 1700]

The concepts of Dirty Page Table and analysis pass was described on slides 37-38 of DB Recovery, but don't appear in the textbook.

f. At the time of the second checkpoint, what LSNs are on pages P0 – P4 on disk?
g. Which updates and CLRs may have to be redone?