Assignment 4

Read Chapter 7 of the textbook.

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 7 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 33).

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. We have five transactions T0-T4 and four database pages P0-P3.

LSN	Trans	Operation Type	Page/Record	Trans backpointer
21	T0	Update	P0/r0	null
22	T1	Update	P1/r1	null
23	T2	Update	P2/r2	null
24	T2	Update	P0/r3	23
25	T2	CLR	P0/r3	24
26	T1	Update	P0/r4	22
27	T0	Update	P3/r5	21
28	checkpoint log record, Active transactions: [T0, 27], [T1, 26], [T2, 25]			
29	T1	Commit		
30	T3	Update	P2/r6	null
31	T0	Update	P0/r0	27
32	T2	CLR	P2/r2	25
33	checkpoint log record, Active transactions: [T0, 31], [T2, 32], [T3, 30]			
34	T2	End Abort		
35	T3	Update	P0/r3	30
36	T4	Update	P2/r2	null

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 37.
b. What LSN is on each page after recovery completes?
c. Based on what you see in the log, what is the smallest LSN of any log record that 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.

f. Would it have been legal for T2 to have written an End Abort record in between LSN 32 and 33? Why?
g. Suppose the system failed immediately after LSN 33, so that records 34 – 36 were not written to the log. What log records would be written during recovery in order to finish the abort of T2?

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

- In LSN 28, Dirty page table = [P0:26, P1:22, P3:27]
- In LSN 33, Dirty page table = [P0:31, P2: 32]

h. Immediately after the second checkpoint is flushed, what LSNs could be on each page on disk?
i. What would it mean if the dirty page table in LSN 33 did not have an entry for P2?