Homework 4

LSN	Txn	Operation Type	Page/Record	Trans backpointer
11	T0	Update	P0/r0	null
12	T1	Update	P1/r1	null
13	T2	Update	P2/r2	null
14	T2	Update	P0/r3	13
15	T2	CLR	P0/r3	14
16	T1	Update	P0/r4	12
17	T0	Update	P3/r5	11
18	checkpoint log record, Active transactions: [T0,17], [T1,16], [T2,15]			
19	T1	Commit		
20	T3	Update	P2/r6	null
21	T0	Update	P0/r0	17
22	T2	CLR	P2/r2	15
23	checkpoint log record, Active transactions: [T0,21], [T2,22], [T3,20]			
24	T2	End Abort		
25	T3	Update	P0/r3	20
26	T4	Update	P2/r2	null

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 records that need to be written are those required to abort T0, T3, and T4:

27	T4	CLR	P2/r2	26
28	T4	End Abort		
29	T3	CLR	P0/r3	25
30	T3	CLR	P2/r6	29
31	T3	End Abort		
32	T0	CLR	P0/r0	21
33	T0	CLR	P3/r5	32
34	T0	CLR	P0/r0	33
35	T0	End Abort		

The above assumes that txns are aborted in the inverse order in which their last log record appears in the log.

b. What LSN is on each page after recovery?

P0 34 P1 12 P2 30 P3 33

For each page, you want the LSN of the last log record that updated the page. For P0, P2, and P3, these are log records that were written by the recovery procedure.

c. Based on what you see in the log, what is the smallest LSN of any log record that might have to be redone?

LSN 20, which is the first update record that follows the penultimate checkpoint.

d. What pages are fetched from disk by the recovery process?

P0 and P2 during the redo pass. P3 during the undo pass.

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.

It uses record-level locking, because (for example) in LSN 11 T0 wrote to P0 and in LSN 14, T2 wrote to P0 even though T0 is still active. LSNs 16, 20, 21, 25, and 26 also show this.

f. Would it have been legal for T2 to have written an End Abort record in between LSN 22 and 23? Why?

Yes, because all of T2's CLRs have been written. But in this case, T2 would have to be deleted from the checkpoint record that follows.

g. Suppose the system failed immediately after LSN 23, so that records 24 – 26 were not written to the log. What log records would be written during recovery in order to finish the abort of T2?

Since the system is not doing the optimization to splice out completed undo records, the undo records also have to be undone, leading to the following:

24	T2	CLR	P2/r2	22
25	T2	CLR	P0/r3	24
26	T2	CLR	P0/r3	25
27	T2	CLR	P2/r2	26
28	T2	End Abort		

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 18, Dirty page table = [P0:15, P1:12, P3:17]
- In LSN 23, Dirty page table = [P0:21, P2:22]

h. At the time of the second checkpoint, what LSNs could be on each page on disk?

- The page LSN is the LSN of the last update record preceding the one pointed to by a dirty page table entry. Since there's no entry for P1 and P3, they must be clean in cache, which means their LSNs are those of the last update record to each page before the checkpoint, namely 12 and 17 resp.
- Entry P0:21 says that LSN 21 needs to be redone. So on disk, P0 could have LSN 16, if it wasn't flushed after the last checkpoint, or LSN 21 or 25 if it was flushed later.
- Given entry P2:22, P2's LSN on disk could be 20, 22, or 26. Notice that it cannot be 13, because entry P2:22 ensures that P2 was flushed after LSN 20.

i. What would it mean if the dirty page table in LSN 23 did not have an entry for P2?

It means that P2 was flushed to disk after LSN 22 was written and before the checkpoint at LSN 23 was performed.