

Assignment 4 - Solution

Problem 1

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

37	T4	CLR	P2/r2	36
38	T4	End Abort		
39	T3	CLR	P0/r3	35
40	T3	CLR	P2/r6	39
41	T3	End Abort		
42	T0	CLR	P0/r0	31
43	T0	CLR	P3/r5	42
44	T0	CLR	P0/r0	43
45	T0	End Abort		

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

b. P0 44

P1 22

P2 40

P3 43

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. LSN 30, which is the first update record that follows the penultimate checkpoint.

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

e. It uses record-level locking because different records on the same page were modified by concurrent transactions. For example in LSN 21, T0 wrote to P0 and in LSN 24, T2 wrote to P0 even though T0 was still active. As another example, T3 in LSN 30 wrote to P2 even though T2 wrote to P0 in LSN 23 and was still active, and later T4 in LSN 36 wrote to P0 even though T3 was still active.

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

g. 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:

34	T2	CLR	P2/r2	32
35	T2	CLR	P0/r3	34
36	T2	CLR	P0/r3	35
37	T2	CLR	P2/r2	36
38	T2	End Abort		

h. The LSN on each page is the LSN of the last update record in the log preceding the log record pointed to by a dirty page table entry. If a page is not in the dirty page table, it is clean.

We have four pages P0-P3. Since pages P1 and P3 have no entries in the second checkpoint, they must be clean in cache, which means their LSNs are those of the last update record to each page before the checkpoint, namely 22 and 27 respectively.

Entry P0:31 says that LSN 31 needs to be redone as the page was dirty at checkpoint time. So on disk, P0 could have LSN 26, if it wasn't flushed after the last checkpoint, or LSN 31 if it was flushed later.

Given entry P2:32, P2's LSN on disk could be 30 or 32. It could be 32 because P2 might be flushed after the checkpoint activity freezes the content of the checkpoint record and before the moment "immediately after the second checkpoint record is flushed." Notice that it cannot be 23, because entry P2:32 ensures that P2 was flushed after LSN 30.

i. It means that P2 was flushed to disk after LSN 32 was written and before the checkpoint at LSN 33 was performed.