

CS W186 Databases - Fall 2019

Guerilla Section 5: Recovery

Saturday, November 23, 2019

1 Conceptual Recovery

1. Consider a scenario where we update the recLSN in the dirty page table to reflect each update to a page, regardless of when the page was brought into the buffer pool. What bugs might you see after recovery? Select all that apply. Explain your reasoning.
 - (a) Some writes of committed transactions would be lost.
 - (b) Some writes of aborted transactions would be visible in the database.
 - (c) The system tries to commit or abort a transaction that is not in the transaction table.

2. Suppose that you are forced to flush pages in the DPT to disk upon making a checkpoint. Which of the following cases are now guaranteed? There is one correct answer. Explain your reasoning.
 - (a) We can skip one of the three phases (analysis/redo/undo) completely
 - (b) We must start analysis from the beginning of the log
 - (c) Redo will start at the checkpoint.
 - (d) Redo must start from the beginning of the log
 - (e) Undo can start at the checkpoint
 - (f) Undo must run until the beginning of the log

3. If the buffer pool is large enough that uncommitted data are never forced to disk, is UNDO still necessary? How about REDO? Explain your reasoning.

4. If updates are always forced to disk when a transaction commits, is UNDO still necessary? Will ARIES perform any REDOs? Explain your reasoning.

2 Recovery Practice

For this question, you will want to have the details of the ARIES protocol handy, so we suggest you have the slides or some notes to look at while doing this question.

The year is 2029. Power outages in Berkeley are so common now that PG&E does not even send out warnings anymore - instead, they just pull the plug whenever they want.

Our database has just restarted from one such power outage. You look at the logs on disk, and this is what you see:

LSN	Record		
10	T1	update	P1
20	T2	update	P2
30	T1	update	P2
40	T1	update	P3
50	begin-checkpoint		
60	T1	update	P4
70	end-checkpoint		
80	T1	commit	
90	T2	update	P1

You load up the checkpoint and see:

Transaction Table			Dirty Page Table	
Txn ID	Last LSN	Txn status	Page	recLSN
T1	40	running	P1	10
T2	20	running	P2	30

1. What is the latest LSN that this checkpoint is guaranteed to be up-to-date to?
2. What do the transaction table and dirty page table look like at the end of analysis, and what log records do we write during analysis?

3. The next phase of ARIES is redo. What LSN do we start the redo from?
4. From that record, we will redo the effects of all the following records, except we will not redo certain records. What are the LSNs of the records we do NOT redo?
5. The last phase of ARIES is undo. What do we do for this phase? Answer this question by writing out the log records that will be recorded for each step.
Stop after you write your first CLR record (make sure your CLR record specifies the nextLSN!).

Click! The lights go out, and you realize PG&E has pulled the power yet again... during ARIES recovery no less!

Five minutes later, the power comes back online. You inspect the log, and are glad to see that **all the records you wrote have made it to disk.**

6. You load up the checkpoint. What does the transaction table and dirty page table look like?
7. You run the analysis phase. What do the transaction table and dirty page look like at the end of analysis?
8. You run the redo phase. In order, what are the LSNs that we redo?
9. Now we run the undo phase. What do we do? (Answer again with the log records that you have to add.)