

# 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.

Answer: a, b. a is correct because during the REDO phase of recovery, some UPDATE log records that reflect writes that never made it to disk will be skipped. Similarly, b is correct, because some CLR's that reflect UNDO's that never made it to disk will be skipped. c is incorrect because even if REDO begins at a later LSN, the system does not add any new transactions to the transaction table during REDO.

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

Answer: c. In general, we redo everything from the earliest recLSN in the DPT to get back unflushed changes from before crash. Since we can guarantee that all changes up to a checkpoint have been flushed, all unflushed changes from before the crash happened after the checkpoint. Therefore, we can redo can start from the checkpoint.

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.

Answer: UNDO isn't necessary in terms of undoing operations on disk. Having a buffer pool large enough to hold all uncommitted data means we don't have to STEAL (allow an uncommitted transaction to overwrite the most recent committed value of an object on disk). Since all the updates will be

sitting in the buffer pool at the time of crash, no changes will be made to disk, so no operations need to be undone. REDO is still necessary. REDO is needed to get back unflushed changes from before the crash. If everything is held in the buffer, this must be redone.

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

Answer: Only UNDO is necessary. UNDO is necessary because we still have to abort transactions that are in progress and aren't committing. The updates that these transactions have gotten through so far must be rolled back. As for REDOs, ARIES might perform some redoes because there may be transactions still in progress at the time of a crash, but these will be undone in the UNDO phase.

## 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	
<b>Txn ID</b>	<b>Last LSN</b>	<b>Txn status</b>	<b>Page</b>	<b>recLSN</b>
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?

Answer: LSN 50 - the **begin-checkpoint** record.

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?

<table border="1" style="margin: auto;"> <thead> <tr> <th>LSN</th> <th>Record</th> </tr> </thead> <tbody> <tr> <td>100</td> <td>T1 end</td> </tr> </tbody> </table>	LSN	Record	100	T1 end	<table border="1" style="margin: auto;"> <thead> <tr> <th colspan="3">Transaction Table</th> </tr> <tr> <th>Txn ID</th> <th>Last LSN</th> <th>Txn status</th> </tr> </thead> <tbody> <tr> <td>T2</td> <td>90</td> <td>running</td> </tr> </tbody> </table>	Transaction Table			Txn ID	Last LSN	Txn status	T2	90	running	<table border="1" style="margin: auto;"> <thead> <tr> <th colspan="2">Dirty Page Table</th> </tr> <tr> <th>Page</th> <th>recLSN</th> </tr> </thead> <tbody> <tr> <td>P1</td> <td>10</td> </tr> <tr> <td>P2</td> <td>30</td> </tr> <tr> <td>P4</td> <td>60</td> </tr> </tbody> </table>	Dirty Page Table		Page	recLSN	P1	10	P2	30	P4	60
LSN	Record																								
100	T1 end																								
Transaction Table																									
Txn ID	Last LSN	Txn status																							
T2	90	running																							
Dirty Page Table																									
Page	recLSN																								
P1	10																								
P2	30																								
P4	60																								

3. The next phase of ARIES is redo. What LSN do we start the redo from?

Answer: LSN 10 - the oldest recLSN in the dirty page table.

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?

Answer:

LSN 20 is not redone; the recLSN of P2 is already higher than 20, so the effects of LSN 20 are already on disk.

LSN 40 is not redone; page P3 is not in the dirty page table, and thus also already on disk.

LSNs 50, 70, and 80 are not update operations, so we don't do anything for them.

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!).

LSN	Record
110	T2 abort
120	T2 CLR nextLSN: 20

**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 log 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?

Answer: Same as the checkpoint from before! We never made another checkpoint.

7. You run the analysis phase. What do the transaction table and dirty page look like at the end of analysis?

Transaction Table			Dirty Page Table	
Txn ID	Last LSN	Txn status	Page	recLSN
T2	120	aborting	P1	10
			P2	30
			P4	60

8. You run the redo phase. In order, what are the LSNs that we redo?

Answer: 10, 30, 60, 90, 120. Importantly, note that we redid T1 even though it committed, and that we ARE redoing CLR's.

9. Now we run the undo phase. What do we do? (Answer again with the log records that you have to add.)

LSN	Record
130	T2 CLR nextLSN: null
140	T2 end