

CARNEGIE MELLON UNIVERSITY  
DEPARTMENT OF COMPUTER SCIENCE  
15-445/645 – DATABASE SYSTEMS (FALL 2017)  
PROF. ANDY PAVLO

Homework 6 (by Sivaprasad Sudhir) – Solutions  
Due: **Monday Nov 27, 2017 @ 11:59pm**

**IMPORTANT:**

- **Upload this PDF** with your answers to **Gradescope by 11:59pm on Monday Nov 27, 2017.**
- **Plagiarism:** Homework may be discussed with other students, but all homework is to be completed **individually.**
- **You have to use this PDF for all of your answers.**

For your information:

- Graded out of **100** points; **3** questions total
- Rough time estimate:  $\approx$  1 - 2 hours

*Revision : 2017/12/10 13:52*

Question	Points	Score
Write-Ahead Logging	35	
ARIES	37	
Miscellaneous	28	
Total:	100	

**Question 1: Write-Ahead Logging.....[35 points]**

Consider a DBMS using write-ahead logging with physical log records with the STEAL and NO-FORCE buffer pool management policy. Assume the DBMS executes a non-fuzzy checkpoint where all dirty pages are written to disk.

Its transaction recovery log contains log records of the following form:

<txnId, objectId, beforeValue, afterValue>

The log also contains checkpoint, transaction begin, and transaction commit records.

The database contains three objects (i.e., X, Y, and Z).

The DBMS sees records as in Figure 1 in the WAL on disk after a crash.

LSN	WAL Record
1	<T1 BEGIN>
2	<T1, X, 1, 2>
3	<T2 BEGIN>
4	<T3 BEGIN>
5	<T2, Y, 1, 2>
6	<T2 COMMIT>
7	<T1, Y, 2, 3>
8	<T3, Z, 1, 2>
9	<CHECKPOINT>
10	<T1, X, 2, 3>
11	<T1, Y, 3, 4>
12	<T3, Z, 2, 3>
13	<T3 COMMIT>
14	<T1, Z, 3, 4>

Figure 1: WAL

(a) [10 points] What are the values of X, Y, and Z in the database stored on disk before the DBMS recovers the state of the database?

- X=1, Y=1, Z=1
- X=1, Y=2, Z=3
- X=2, Y=3, Z=2
- X=2, Y=2, Z=3
- X=3, Y=1, Z=2
- X=3, Y=4, Z=3
- X=3, Y=4, Z=4

■ Not possible to determine

**Solution:** The checkpoint flushed everything to disk, but then all the data objects were modified by transactions after the checkpoint.

Since we are using NO-FORCE, any dirty page could be written to disk, so therefore we don't know the contents of the database on disk at the crash.

(b) [5 points] What should be the correct action on T1 when recovering the database from WAL?

■ **undo all of T1's changes**

- redo all of T1's changes
- do nothing to T1

**Solution:** T1 never committed. All of its changes should only be undone.

(c) [5 points] What should be the correct action on T2 when recovering the database from WAL?

- undo all of T2's changes
- redo all of T2's changes

■ **do nothing to T2**

**Solution:** T2 committed before the checkpoint. All of its changes were written to disk. There is nothing to redo or undo.

(d) [5 points] What should be the correct action on T3 when recovering the database from WAL?

- undo all of T3's changes
- **redo all of T3's changes**
- do nothing to T3

**Solution:** T3 committed after the checkpoint, so that means the DBMS has to redo all of its changes.

(e) [10 points] Assume that the DBMS flushes all dirty pages when the recovery process finishes. What are the values of X, Y, and Z after the DBMS recovers the state of the database from the WAL in Figure 1?

- X=1, Y=1, Z=1
- **X=1, Y=2, Z=3**
- X=2, Y=3, Z=2
- X=2, Y=2, Z=3
- X=3, Y=1, Z=2
- X=3, Y=4, Z=3
- X=3, Y=4, Z=4
- Not possible to determine

**Solution:** X = 1 (rollback the beforeValue from T1)  
Y = 2 (rollback to the afterValue made by T2)

Z = 3 (rollback to the afterValue made by T3)

**Question 2: ARIES..... [37 points]**

Consider a DBMS using ARIES protocol for logging records.

Its transaction recovery log contains log records of the following form:

<txnId, objectId, beforeValue, afterValue>

The log also contains checkpoint begin, checkpoint end, transaction begin, transaction commit, transaction end and undo action records.

The database contains four objects (i.e., X, Y, Z and W).

The DBMS sees records as in Figure 2 in the log on disk after a crash. Notice that there are no dirty pages nor active transactions, during the first and only checkpoint.

LSN	ARIES Record
1	<BEGIN CHECKPOINT>
2	<END CHECKPOINT>
3	<T1, X, 1, 2>
4	<T2, Y, 3, 4>
5	<T1 COMMIT>
6	<T1 END>
7	<T3, Z, 5, 6>
8	<T4, W, 5, 6>
9	<T2 COMMIT>

Figure 2: ARIES

(a) [10 points] Which transaction(s) will be undone, if any?

- T1
- T2
- T3
- T4
- None of the above

**Solution:** T3, T4

(b) After the recovery has ended successfully and assuming that all the dirty pages have been flushed to disk,

i. [3 points] What will be the value of X on disk?

- 1
- 2
- Unknown

**Solution:** New value of X will be on disk as T1 will be redone

ii. [3 points] What will be the value of Y on disk?

- 3
- 4
- Unknown

**Solution:** New value of Y will be on disk as T2 will be redone

iii. [3 points] What will be the value of Z on disk?

- 5
- 6
- Unknown

**Solution:** Old value of Z will be on disk as T3 will be undone

iv. [3 points] What will be the value of W on disk?

- 5
- 6
- Unknown

**Solution:** Old value of W will be on disk as T4 will be undone

(c) [15 points] After the recovery has ended, what will be the contents of the log? Specify the log records after the crash.

**Solution:** 10: <T2 end>  
11: <CLR undo T4 LSN 8>  
12: <T4 end ;>  
13: <CLR undo T3 LSN 7>  
14: <T3 end>

**Question 3: Miscellaneous ..... [28 points]**

All the questions below refer to recovery algorithms that use a write-ahead log (WAL). State True or False for each of the following questions.

- (a) **[4 points]** Under NO-STEAL + FORCE policy, a DBMS will never need to redo the changes of a committed transaction during recovery.
- True  
 False
- (b) **[4 points]** Under NO-STEAL + FORCE policy, a DBMS will have to undo the changes of an aborted transaction during recovery.
- True  
 False
- (c) **[4 points]** Under the NO-STEAL policy, a DBMS will need to store the whole table in RAM if a transaction updates all the records of that table.
- True  
 False
- (d) **[4 points]** While doing deferred updates with WAL, if we prevent the DBMS from writing dirty records to disk until the transaction commits, then we do not need to store their original values.
- True  
 False
- (e) **[4 points]** Most production systems use STEAL + NO-FORCE policy as it has the fastest recovery performance.
- True  
 False
- (f) **[4 points]** In ARIES, log records are immediately flushed on the log, as soon as they are produced.
- True  
 False
- (g) **[4 points]** Without checkpoints, the redo phase of ARIES recovery should process the whole log.
- True  
 False