Exercises

(Course: Database Management Systems)

Chapter 6

Database Recovery Techniques

1. Exercise 23.21 in the text book ("Fundamentals of Database Systems- 6th Edition", Elmasri et al.)

Suppose that the system crashes before the [read_item, T_3, A] entry is written to the log in Figure 23.1(b). Will that make any difference in the recovery process?

(a)

T_1	T_2	T_3
read_item(*A*)	read_item(*B*)	read_item(*C*)
read_item(*D*)	write_item(*B*)	write_item(*B*)
write_item(*D*)	read_item(*D*)	read_item(*A*)
	write_item(*D*)	write_item(*A*)

(b)

		A	*B*	*C*	*D*
		30	15	40	20
	[start_transaction,T_3]				
	[read_item,T_3,*C*]				
*	[write_item,T_3,*B*,15,12]		12		
	[start_transaction,T_2]				
	[read_item,T_2,*B*]				
**	[write_item,T_2,*B*,12,18]		18		
	[start_transaction,T_1]				
	[read_item,T_1,*A*]				
	[read_item,T_1,*D*]				
	[write_item,T_1,*D*,20,25]				25
	[read_item,T_2,*D*]				
**	[write_item,T_2,*D*,25,26]				26
	[read_item,T_3,*A*]				

← System crash

Figure 23.1
Illustrating cascading rollback (a process that never occurs in strict or cascadeless schedules). (a) The read and write operations of three transactions. (b) System log at point of crash. (c) Operations before the crash.

* T_3 is rolled back because it did not reach its commit point.

** T_2 is rolled back because it reads the value of item *B* written by T_3.

2. Exercise 23.22 in the text book ("Fundamentals of Database Systems- 6th Edition", Elmasri et al.)

Suppose that the system crashes before the [write_item, T_2, D, 25, 26] entry is written to the log in Figure 23.1(b). Will that make any difference in the recovery process?

3. Exercise 23.23 in the text book ("Fundamentals of Database Systems- 6th Edition", Elmasri et al.)

Figure 23.6 shows the log corresponding to a particular schedule at the point of a system crash for four transactions T_1, T_2, T_3, and T_4. Suppose that we use the *immediate update protocol* with checkpointing. Describe the recovery process from the system crash. Specify which transactions are rolled-back, which operations in the log are redone and which (if any) are undone, and whether any cascading rollback takes place.

[start_transaction, T_1]
[read_item, T_1, A]
[read_item, T_1, D]
[write_item, T_1, D, 20, 25]
[commit, T_1]
[checkpoint]
[start_transaction, T_2]
[read_item, T_2, B]
[write_item, T_2, B, 12, 18]
[start_transaction, T_4]
[read_item, T_4, D]
[write_item, T_4, D, 25, 15]
[start_transaction, T_3]
[write_item, T_3, C, 30, 40]
[read_item, T_4, A]
[write_item, T_4, A, 30, 20]
[commit, T_4]
[read_item, T_2, D]
[write_item, T_2, D, 15, 25] ← System crash

Figure 23.6
A sample schedule and its corresponding log.

4. Exercise 23.24 in the text book ("Fundamentals of Database Systems- 6th Edition", Elmasri et al.)

Suppose that we use the deferred update protocol for the example in Figure 23.6. Show how the log would be different in the case of deferred update by removing the unnecessary log entries; then describe the recovery process, using your modified log. Assume that only REDO operations are applied, and specify which operations in the log are redone and which are ignored.

5. After a crash, we find the following log and transaction & dirty page tables at time of checkpoint:

Table 1 – The log at point of crash

LSN	LAST_LSN	TRAN_ID	TYPE	PAGE_ID	…
1	0	T1	update	C	…
2	1	T1	update	B	…

3	0	T2	update	C	…
4	begin_checkpoint				
5	end_checkpoint				
6	2	T1	commit		…
7	0	T3	update	A	…

Table 1 – Transaction & Dirty Page Tables at time of checkpoint

TRANSACTION TABLE			DIRTY PAGE TABLE	
TRANSACTION ID	LAST LSN	STATUS	PAGE ID	LSN
T1	2	in progress	C	1
T2	3	in progress	B	2

Specify the transaction and dirty page tables after the analysis phase.