


Chapter 23

Database Recovery Techniques



Addison-Wesley
is an imprint of

PEARSON

(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]			

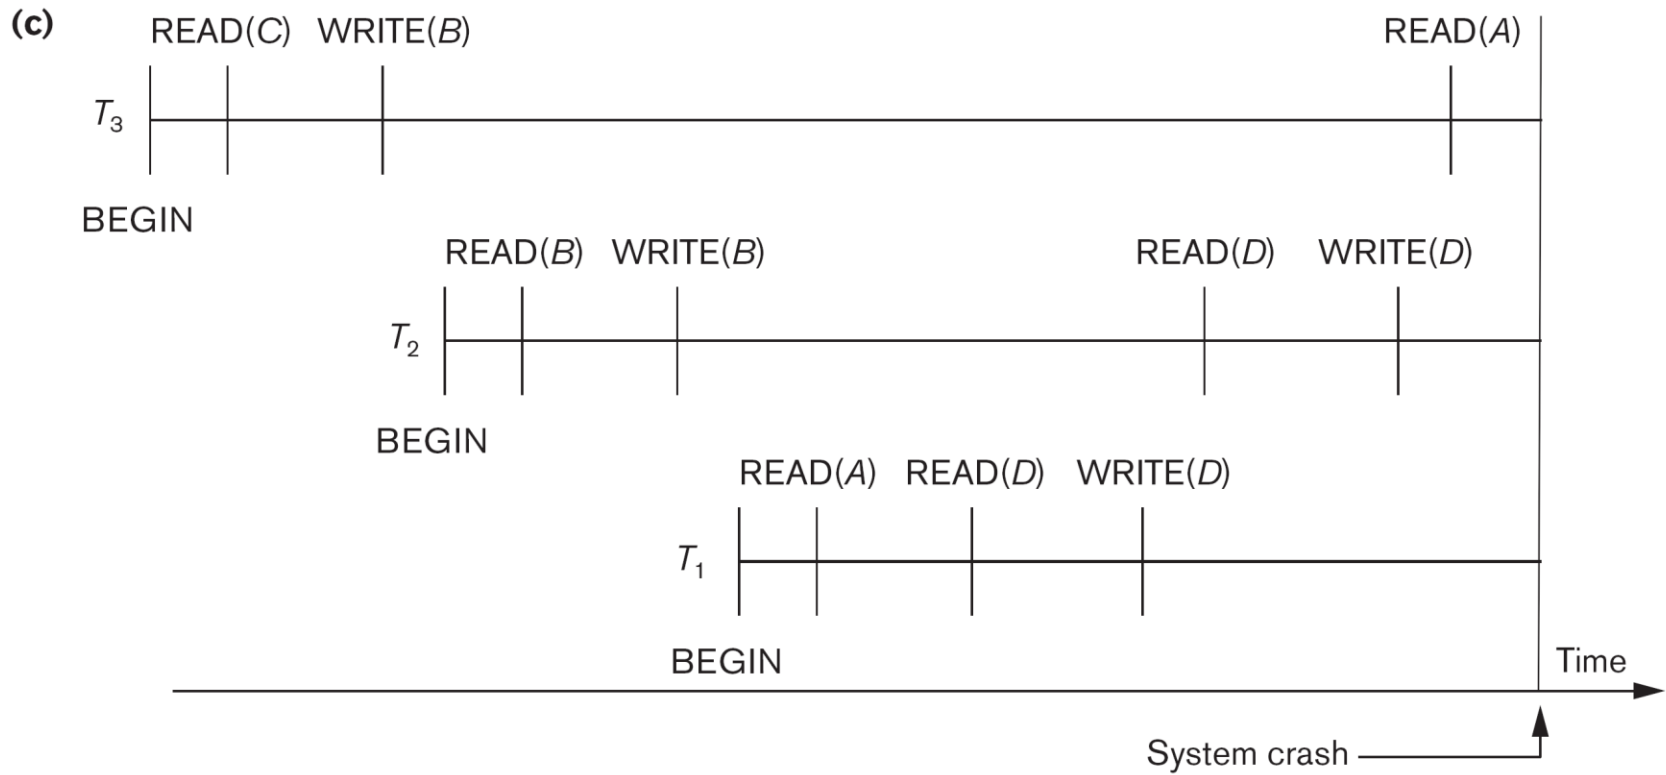
← System crash

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 .






Figure 23.2
 An example of a recovery timeline to illustrate the effect of checkpointing.

(a)

T_1	T_2	T_3	T_4
read_item(A)	read_item(B)	read_item(A)	read_item(B)
read_item(D)	write_item(B)	write_item(A)	write_item(B)
write_item(D)	read_item(D)	read_item(C)	read_item(A)
	write_item(D)	write_item(C)	write_item(A)

(b)

[start_transaction, T_1]
[write_item, T_1 , D, 20]
[commit, T_1]
[checkpoint]
[start_transaction, T_4]
[write_item, T_4 , B, 15]
[write_item, T_4 , A, 20]
[commit, T_4]
[start_transaction, T_2]
[write_item, T_2 , B, 12]
[start_transaction, T_3]
[write_item, T_3 , A, 30]
[write_item, T_2 , D, 25]

← System crash

T_2 and T_3 are ignored because they did not reach their commit points.

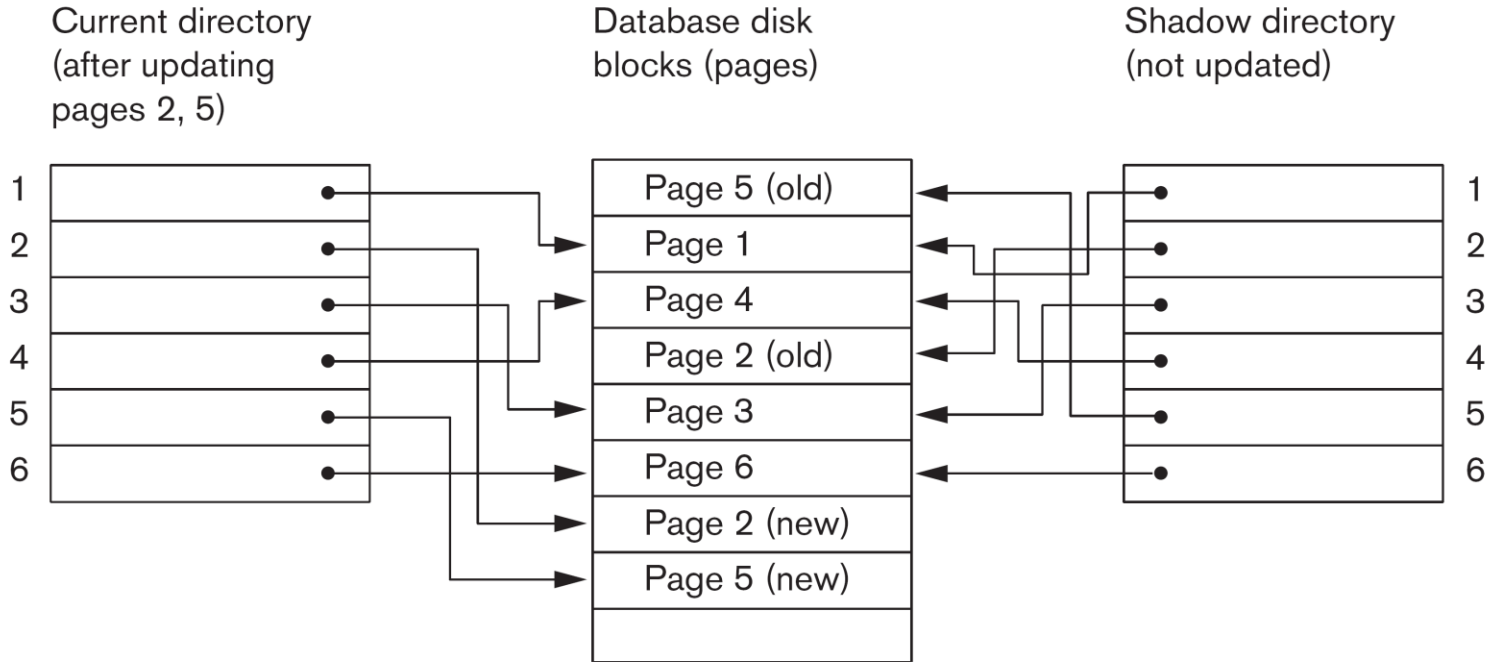
T_4 is redone because its commit point is after the last system checkpoint.

Figure 23.3

An example of recovery using deferred update with concurrent transactions. (a) The READ and WRITE operations of four transactions. (b) System log at the point of crash.

Figure 23.4

An example of shadow paging.



⁵The directory is similar to the page table maintained by the operating system for each process.

(a)

Lsn	Last_lsn	Tran_id	Type	Page_id	Other_information
1	0	T_1	update	<i>C</i>	...
2	0	T_2	update	<i>B</i>	...
3	1	T_1	commit		...
4	begin checkpoint				
5	end checkpoint				
6	0	T_3	update	<i>A</i>	...
7	2	T_2	update	<i>C</i>	...
8	7	T_2	commit		...

TRANSACTION TABLE

(b)

Transaction_id	Last_lsn	Status
T_1	3	commit
T_2	2	in progress

DIRTY PAGE TABLE

Page_id	Lsn
<i>C</i>	1
<i>B</i>	2

TRANSACTION TABLE

(c)

Transaction_id	Last_lsn	Status
T_1	3	commit
T_2	8	commit
T_3	6	in progress

DIRTY PAGE TABLE

Page_id	Lsn
<i>C</i>	1
<i>B</i>	2
<i>A</i>	6

Figure 23.5

An example of recovery in ARIES. (a) The log at point of crash. (b) The Transaction and Dirty Page Tables at time of checkpoint. (c) The Transaction and Dirty Page Tables after the analysis phase.

[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.