

# Database System Internals Transactions: Recovery (part 1)

Paul G. Allen School of Computer Science and Engineering University of Washington, Seattle

February 19, 2021

CSE 444 - Winter 2020

Main textbook (Garcia-Molina)

- Ch. 17.2-4, 18.1-3, 18.8-9
- Second textbook (Ramakrishnan)
- Ch. 16-18

Also: M. J. Franklin. Concurrency Control and Recovery. The Handbook of Computer Science and Engineering, A. Tucker, ed., CRC Press, Boca Raton, 1997.

## **Transaction Management**

Two parts:

- Concurrency control: ACID
- Recovery from crashes: <u>ACID</u>

We already discussed concurrency control You are implementing locking in lab3

Today, we start recovery

Type of Crash	Prevention
Wrong data entry	Constraints and Data cleaning
Disk crashes	Redundancy: e.g. RAID, archive
Data center failures	Remote backups or replicas
System failures: e.g. power	DATABASE RECOVERY



## Each transaction has internal state

- When system crashes, internal state is lost
  - Don't know which parts executed and which didn't
  - Need ability to undo and redo

## **Buffer Manager Review**



## **Buffer Manager Review**

- Enables higher layers of the DBMS to assume that needed data is in main memory
- Caches data in memory. Problems when crash occurs:
  - 1. If committed data was not yet written to disk
  - 2. If uncommitted data was flushed to disk

## Transactions

- Assumption: the database is composed of <u>elements</u>.
- I element can be either:
  - 1 page = physical logging
  - 1 record = logical logging
- In Lab 4 we use page-level elements

## Primitive Operations of Transactions

- READ(X,t)
  - copy element X to transaction local variable t
- WRITE(X,t)
  - copy transaction local variable t to element X
- INPUT(X)
  - read element X to memory buffer
- OUTPUT(X)
  - write element X to disk

```
BEGIN TRANSACTION
READ(A,t);
t := t*2;
WRITE(A,t);
READ(B,t);
t := t*2;
WRITE(B,t)
COMMIT;
```

Initially, A=B=8.

<u>Atomicity</u> requires that either (1) T commits and A=B=16, or (2) T does not commit and A=B=8.

```
BEGIN TRANSACTION
READ(A,t);
t := t*2;
                                Initially, A=B=8.
WRITE(A,t);
                                Atomicity requires that either
READ(B,t);
                                (1) T commits and A=B=16, or
                                (2) T does not commit and A=B=8.
    Will look at various crash scenarios
    What behavior do we want in each case?
```

Fe

Transactio		n Buffe	r pool	Disk	
Action	t	Mem A	Mem B	Disk A	Disk B
INPUT(A)		8		8	8
READ(A,t)					
t:=t*2					
WRITE(A,t)					
INPUT(B)					
READ(B,t)					
t:=t*2					
WRITE(B,t)					
OUTPUT(A)					
OUTPUT(B)					
COMMIT					

	Transaction	n Buffei	r pool	D	isk
Action	t	Mem A	Mem B	Disk A	Disk B
INPUT(A)		8		8	8
READ(A,t)	8	8		8	8
t:=t*2					
WRITE(A,t)					
INPUT(B)					
READ(B,t)					
t:=t*2					
WRITE(B,t)					
OUTPUT(A)					
OUTPUT(B)					
 COMMIT					

Fel

	Transactio	n Buffei	r pool	D	isk
Action	t	Mem A	Mem B	Disk A	Disk B
INPUT(A)		8		8	8
READ(A,t)	8	8		8	8
t:=t*2	16	8		8	8
WRITE(A,t)					
INPUT(B)					
READ(B,t)					
t:=t*2					
WRITE(B,t)					
OUTPUT(A)					
OUTPUT(B)					

	Transaction	n Buffei	r pool	D	isk
Action	t	Mem A	Mem B	Disk A	Disk B
INPUT(A)		8		8	8
READ(A,t)	8	8		8	8
t:=t*2	16	8		8	8
WRITE(A,t)	16	16		8	8
INPUT(B)					
READ(B,t)					
t:=t*2					
WRITE(B,t)					
OUTPUT(A)					
OUTPUT(B)					
 COMMIT					

	Transaction	n Buffei	r pool	D	isk
Action	t	Mem A	Mem B	Disk A	Disk B
INPUT(A)		8		8	8
READ(A,t)	8	8		8	8
t:=t*2	16	8		8	8
WRITE(A,t)	16	16		8	8
INPUT(B)	16	16	8	8	8
READ(B,t)					
t:=t*2					
WRITE(B,t)					
OUTPUT(A)					
OUTPUT(B)					
 COMMIT					

	Transaction	n Buffei	r pool	D	isk
Action	t	Mem A	Mem B	Disk A	Disk B
INPUT(A)		8		8	8
READ(A,t)	8	8		8	8
t:=t*2	16	8		8	8
WRITE(A,t)	16	16		8	8
INPUT(B)	16	16	8	8	8
READ(B,t)	8	16	8	8	8
t:=t*2					
WRITE(B,t)					
OUTPUT(A)					
OUTPUT(B)					
 COMMIT					

	Transaction	n Buffei	r pool	D	isk
Action	t	Mem A	Mem B	Disk A	Disk B
INPUT(A)		8		8	8
READ(A,t)	8	8		8	8
t:=t*2	16	8		8	8
WRITE(A,t)	16	16		8	8
INPUT(B)	16	16	8	8	8
READ(B,t)	8	16	8	8	8
t:=t*2	16	16	8	8	8
WRITE(B,t)					
OUTPUT(A)					
OUTPUT(B)					
 COMMIT					

		Transaction	n Buffei	Buffer pool		isk
	Action	t	Mem A	Mem B	Disk A	Disk B
	INPUT(A)		8		8	8
	READ(A,t)	8	8		8	8
	t:=t*2	16	8		8	8
	WRITE(A,t)	16	16		8	8
	INPUT(B)	16	16	8	8	8
	READ(B,t)	8	16	8	8	8
	t:=t*2	16	16	8	8	8
	WRITE(B,t)	16	16	16	8	8
	OUTPUT(A)					
	OUTPUT(B)					
	COMMIT					
Feb	ruary 19 2021			Winter 2020		

	Transaction		n Buffei	Buffer pool		isk
	Action	t	Mem A	Mem B	Disk A	Disk B
	INPUT(A)		8		8	8
	READ(A,t)	8	8		8	8
	t:=t*2	16	8		8	8
	WRITE(A,t)	16	16		8	8
	INPUT(B)	16	16	8	8	8
	READ(B,t)	8	16	8	8	8
	t:=t*2	16	16	8	8	8
	WRITE(B,t)	16	16	16	8	8
	OUTPUT(A)	16	16	16	16	8
	OUTPUT(B)					
	COMMIT					
Feb	$r_{\rm H}$ ary 19 2021			Winter 2020		

		Transaction	n Buffei	Buffer pool		isk
	Action	t	Mem A	Mem B	Disk A	Disk B
	INPUT(A)		8		8	8
	READ(A,t)	8	8		8	8
	t:=t*2	16	8		8	8
	WRITE(A,t)	16	16		8	8
	INPUT(B)	16	16	8	8	8
	READ(B,t)	8	16	8	8	8
	t:=t*2	16	16	8	8	8
	WRITE(B,t)	16	16	16	8	8
	OUTPUT(A)	16	16	16	16	8
	OUTPUT(B)	16	16	16	16	16
	COMMIT					
Feb	ruary 19 2021		CSE AAA =	Winter 2020		

	Action	t	Mem A	Mem B	Disk A	Disk B	
	INPUT(A)		8		8	8	
	READ(A,t)	8	8		8	8	
	t:=t*2	16	8		8	8	
	WRITE(A,t)	16	16		8	8	
	INPUT(B)	16	16	8	8	8	
	READ(B,t)	8	16	8	8	8	
	t:=t*2	16	16	8	8	8	
	WRITE(B,t)	16	16	16	8	8	
	OUTPUT(A)	16	16	16	16	8	Crash !
	OUTPUT(B)	16	16	16	16	16	
	COMMIT						
Febi	ruary 19, 2021		CSE 444 - \	Winter 2020			

#### Yes it's bad: A=16, B=8....

	Action	t	Mem A	Mem B	Disk A	Disk B	
	INPUT(A)		8		8	8	
	READ(A,t)	8	8		8	8	
	t:=t*2	16	8		8	8	
	WRITE(A,t)	16	16		8	8	
	INPUT(B)	16	16	8	8	8	
	READ(B,t)	8	16	8	8	8	
	t:=t*2	16	16	8	8	8	
	WRITE(B,t)	16	16	16	8	8	
	OUTPUT(A)	16	16	16	16	8	Crash !
	OUTPUT(B)	16	16	16	16	16	
	COMMIT						
Feb	ruary 19, 2021		CSE 444 - \	Winter 2020	-		25

	Action	t	Mem A	Mem B	Disk A	Disk B	
	INPUT(A)		8		8	8	
	READ(A,t)	8	8		8	8	
	t:=t*2	16	8		8	8	
	WRITE(A,t)	16	16		8	8	
	INPUT(B)	16	16	8	8	8	
	READ(B,t)	8	16	8	8	8	
	t:=t*2	16	16	8	8	8	
	WRITE(B,t)	16	16	16	8	8	
	OUTPUT(A)	16	16	16	16	8	
	OUTPUT(B)	16	16	16	16	16	Cras
	COMMIT						
Feb	ruary 19, 2021		CSE 444 - `	Winter 2020			

#### Yes it's bad: A=B=16, but not committed

Z

27

Action	t	Mem A	Mem B	Disk A	Disk B
INPUT(A)		8		8	8
READ(A,t)	8	8		8	8
t:=t*2	16	8		8	8
WRITE(A,t)	16	16		8	8
INPUT(B)	16	16	8	8	8
READ(B,t)	8	16	8	8	8
t:=t*2	16	16	8	8	8
WRITE(B,t)	16	16	16	8	8
OUTPUT(A)	16	16	16	16	8
OUTPUT(B)	16	16	16	16	16
COMMIT					

	Action	t	Mem A	Mem B	Disk A	Disk B	
	INPUT(A)		8		8	8	
	READ(A,t)	8	8		8	8	
	t:=t*2	16	8		8	8	
	WRITE(A,t)	16	16		8	8	
	INPUT(B)	16	16	8	8	8	
	READ(B,t)	8	16	8	8	8	
	t:=t*2	16	16	8	8	8	
	WRITE(B,t)	16	16	16	8	8	Crash !
	OUTPUT(A)	16	16	16	16	8	
	OUTPUT(B)	16	16	16	16	16	
	COMMIT						
Febi	ruary 19, 2021		CSE 444 - \	Winter 2020			28

#### No: that's OK

	Action	t	Mem A	Mem B	Disk A	Disk B	
	INPUT(A)		8		8	8	
	READ(A,t)	8	8		8	8	
	t:=t*2	16	8		8	8	
	WRITE(A,t)	16	16		8	8	
	INPUT(B)	16	16	8	8	8	
	READ(B,t)	8	16	8	8	8	
	t:=t*2	16	16	8	8	8	
	WRITE(B,t)	16	16	16	8	8	Crash !
	OUTPUT(A)	16	16	16	16	8	
	OUTPUT(B)	16	16	16	16	16	
	COMMIT						
Feb	ruary 19, 2021		CSE 444 - Y	Winter 2020			29

#### OUTPUT can also happen after COMMIT (details coming)

Action	t	Mem A	Mem B	Disk A	Disk B
INPUT(A)		8		8	8
READ(A,t)	8	8		8	8
t:=t*2	16	8		8	8
WRITE(A,t)	16	16		8	8
INPUT(B)	16	16	8	8	8
READ(B,t)	8	16	8	8	8
t:=t*2	16	16	8	8	8
WRITE(B,t)	16	16	16	8	8
COMMIT					
OUTPUT(A)	16	16	16	16	8
OUTPUT(B)	16	16	16	16	16

#### OUTPUT can also happen after COMMIT (details coming)

Action	t	Mem A	Mem B	Disk A	Disk B	
INPUT(A)		8		8	8	
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	
INPUT(B)	16	16	8	8	8	
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	
COMMIT						
OUTPUT(A)	16	16	16	16	8	Crash !
OUTPUT(B)	16	16	16	16	16	

### FORCE or NO-FORCE

 Should all updates of a transaction be forced to disk before the transaction commits?

## STEAL or NO-STEAL

 Can an update made by an uncommitted transaction overwrite the most recent committed value of a data item on disk?

# Force/No-steal (most strict)

FORCE: Pages of committed transactions must be forced to disk before commit

### NO-STEAL: Pages of uncommitted transactions cannot be written to disk

## Easy to implement (how?) and ensures atomicity

# No-Force/Steal (most strict)

 NO-FORCE: Pages of committed transactions need not be written to disk

STEAL: Pages of uncommitted transactions may be written to disk

# In both cases, need a Write Ahead Log (WAL) to provide atomicity in face of failures

The Log: append-only file containing log records

- Records every single action of every TXN
- Forces log entries to disk as needed
- After a system crash, use log to recover
- Three types: UNDO, REDO, UNDO-REDO
- Aries: is an UNDO-REDO log

	NO-STEAL	STEAL
FORCE	Lab 3	Undo Log
NO-FORCE	Redo Log	Undo-Redo Log

# "UNDO" Log

FORCE and STEAL

February 19, 2021

CSE 444 - Winter 2020
Log records

- START T>
  - transaction T has begun
- <COMMIT T>
  - T has committed
- ABORT T>
  - T has aborted
- <t
  - T has updated element X, and its <u>old</u> value was v
  - Idempotent, physical log records

Action	t	Mem A	Mem B	Disk A	Disk B	UNDO Log
						<start t=""></start>
INPUT(A)		8		8	8	
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	<t,a,<mark>8&gt;</t,a,<mark>
INPUT(B)	16	16	8	8	8	
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	<t,b,8></t,b,8>
OUTPUT(A)	16	16	16	16	8	
OUTPUT(B)	16	16	16	16	16	
COMMIT						<commit t=""></commit>

	Action	t	Mem A	Mem B	Disk A	Disk B	UNDO Log
							<start t=""></start>
	INPUT(A)		8		8	8	
	READ(A,t)	8	8		8	8	
	t:=t*2	16	8		8	8	
	WRITE(A,t)	16	16		8	8	<t,a,8></t,a,8>
	INPUT(B)	16	16	8	8	8	
	READ(B,t)	8	16	8	8	8	
	t:=t*2	16	16	8	8	8	
	WRITE(B,t)	16	16	16	8	8	<t,b,<mark>8&gt;</t,b,<mark>
	OUTPUT(A)	16	16	16	16	8	Croch
	OUTPUT(B)	16	16	16	16	16	
	COMMIT						<commit t=""></commit>
Feb	WHAT DO WE DO ?       bruary 19, 2021     CSE 444 - Winter 2020     40						

	Action	t	Mem A	Mem B	Disk A	Disk B	UNDO Log	
							<start t=""></start>	
	INPUT(A)		8		8	8		
	READ(A,t)	8	8		8	8		
	t:=t*2	16	8		8	8		
	WRITE(A,t)	16	16		8	8	<t,a,8></t,a,8>	
	INPUT(B)	16	16	8	8	8		
	READ(B,t)	8	16	8	8	8		
	t:=t*2	16	16	8	8	8		
	WRITE(B,t)	16	16	16	8	8	<t,b,8></t,b,8>	
	OUTPUT(A)	16	16	16	16	8		
	OUTPUT(B)	16	16	16	16	16	Crash !	
	COMMIT						<commit t=""></commit>	
Febr	WHAT DO WE DO ? We UNDO by setting B=8 and A=8							

February 19, 2021

41

Action	t	Mem A	Mem B	Disk A	Disk B	UNDO Log		
						<start t=""></start>		
INPUT(A)		8		8	8			
READ(A,t)	8	8		8	8			
t:=t*2	16	8		8	8			
WRITE(A,t)	16	16		8	8	<t,a,8></t,a,8>		
INPUT(B)	16	16	8	8	8			
READ(B,t)	8	16	8	8	8			
t:=t*2	16	16	8	8	8			
WRITE(B,t)	16	16	16	8	8	<t,b,8></t,b,8>		
OUTPUT(A)	16	16	16	16	8			
OUTPUT(B)	16	16	16	16	16			
COMMIT						<commit t=""></commit>		
What do v	What do we do now ?							

February 19, 2021

Action	t	Mem A	Mem B	Disk A	Disk B	UNDO Log
						<start t=""></start>
INPUT(A)		8		8	8	
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	<t,a,8></t,a,8>
INPUT(B)	16	16	8	8	8	
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	<t,b,8></t,b,8>
OUTPUT(A)	16	16	16	16	8	
OUTPUT(B)	16	16	16	16	16	
COMMIT						<commit t=""></commit>
What do v	ve do now	? 	Noth E 444 - Winter 2	ing: log col	ntains CON	ИМІТ

February 19, 2021

• This is all we see (for example):

Disk A	Disk B	<start t=""></start>
8	16	<t,a,8></t,a,8>
		<t,b,8></t,b,8>

• This is all we see (for example):

Disk A	Disk B	<start t=""></start>
8	16	<t,a,8></t,a,8>
		<t,b,8></t,b,8>

- This is all we see (for example):
- Need to step through the log

Disk A	Disk B	<start t=""></start>
8	16	<t,a,8></t,a,8>
		<t,b,8></t,b,8>

- This is all we see (for example):
- Need to step through the log



• What direction?

- This is all we see (for example):
- Need to step through the log



- What direction?
- In UNDO log, we start at the most recent and go backwards in time

- This is all we see (for example):
- Need to step through the log



- What direction?
- In UNDO log, we start at the most recent and go backwards in time

- This is all we see (for example):
- Need to step through the log



- What direction?
- In UNDO log, we start at the most recent and go backwards in time

- This is all we see (for example):
- Need to step through the log



- What direction?
- In UNDO log, we start at the most recent and go backwards in time

- This is all we see (for example):
- Need to step through the log



- What direction?
- In UNDO log, we start at the most recent and go backwards in time

### If we see NO Commit statement:

- We UNDO both changes: A=8, B=8
- The transaction is atomic, since none of its actions have been executed

### In we see that T has a Commit statement

- We don't undo anything
- The transaction is atomic, since both it's actions have been executed

After system's crash, run recovery manager

- Decide for each transaction T whether it is completed or not
  - <START T>....<COMMIT T>.... = yes
  - <START T>....<ABORT T>.... = yes
  - <START T>..... = no
- Undo all modifications by incomplete transactions

Recovery manager:

Read log from the end; cases:
<COMMIT T>: mark T as completed
<ABORT T>: mark T as completed
<T,X,v>: if T is not completed
then write X=v to disk
else ignore
<START T>: ignore

• • •

<T6,X6,v6>

• • •

... <START T5> <START T4> <T1,X1,v1> <T5,X5,v5> <T4,X4,v4> <COMMIT T5> <T3,X3,v3> <T2,X2,v2> Question1: Which updates are undone ?

Question 2:

How far back do we need to read in the log ?

**Question 3:** 

What happens if second crash during recovery?

February 19, 2021

Crash !

• • •

<T6,X6,v6>

• • •

... <START T5> <START T4> <T1,X1,v1> <T5,X5,v5> <T4,X4,v4> <COMMIT T5> <T3,X3,v3> <T2,X2,v2> Question1: Which updates are undone ?

Question 2:

How far back do we need to read in the log ? To the beginning.

**Question 3**:

What happens if second crash during recovery?

February 19, 2021

Crash !

•••

<T6,X6,v6>

• • •

... <START T5> <START T4> <T1,X1,v1> <T5,X5,v5> <T4,X4,v4> <COMMIT T5> <T3,X3,v3> <T2,X2,v2> Question1: Which updates are undone ?

Question 2:

How far back do we need to read in the log ? To the beginning.

#### **Question 3**:

What happens if second crash during recovery? No problem! Log records are idempotent. Can reapply.

February 19, 2021

Crash !

Action	t	Mem A	Mem B	Disk A	Disk B	UNDO Log
						<start t=""></start>
INPUT(A)		V	When mu	ust	8	
READ(A,t)	8	V	ve force	pages	8	
t:=t*2	16	8	O OISK ?		8	
WRITE(A,t)	16	16		8	8	< <b>T,A</b> ,8>
INPUT(B)	16	16	8	8	8	
READ(B,t)	8	16	8	8	8	2
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	< <mark>T,B</mark> ,8>
OUTPUT(A)	<b>1</b> 6	16	16	16	8	
OUT <mark>P</mark> UT(B)	<b>1</b> 6	16	16	16	16	
COMMIT						<commit t=""></commit>

Action	t	Mem A	Mem B	Disk A	Disk B	UNDO Log
						<start t=""></start>
INPUT(A)		8		8	8	
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	<
INPUT(B)	16	16	8	8	8	
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	<t,b,8></t,b,8>
OUTPUT(A)	16	16	16	16	8	
OUTPUT(B)	16	16	16	16	16	
COMMIT				FOR	CE	

RULES: log entry *before* OUTPUT *before* COMMIT

February 19, 2021

# **Undo-Logging Rules**

- U1: If T modifies X, then <T,X,v> must be written to disk before OUTPUT(X)
- U2: If T commits, then OUTPUT(X) must be written to disk before <COMMIT T>
- Hence: OUTPUTs are done <u>early</u>, before the transaction commits



Checkpoint the database periodically

- Stop accepting new transactions
- Wait until all current transactions complete
- Flush log to disk
- Write a <CKPT> log record, flush
- Resume transactions

# Undo Recovery with Checkpointing



### **Nonquiescent Checkpointing**

- Problem with checkpointing: database freezes during checkpoint
- Would like to checkpoint while database is operational
- Idea: nonquiescent checkpointing

Quiescent = being quiet, still, or at rest; inactive Non-quiescent = allowing transactions to be active

### Nonquiescent Checkpointing

- Write a <START CKPT(T1,...,Tk)> where T1,...,Tk are all active transactions. Flush log to disk
- Continue normal operation
- When all of T1,...,Tk have completed, write <END CKPT>, flush log to disk

# Undo with Nonquiescent Checkpointing



# Implementing ROLLBACK

- Recall: a transaction can end in COMMIT or ROLLBACK
- Idea: use the undo-log to implement ROLLBACK
- How ?
  - LSN = Log Sequence Number
  - Log entries for the same transaction are linked, using the LSN's
  - Read log in reverse, using LSN pointers

### Implementing ROLLRACK



CK

### **REDO**

#### NO-FORCE and NO-STEAL

Action	t	Mem A	Mem B	Disk A	Disk B
READ(A,t)	8	8		8	8
t:=t*2	16	8		8	8
WRITE(A,t)	16	16		8	8
READ(B,t)	8	16	8	8	8
t:=t*2	16	16	8	8	8
WRITE(B,t)	16	16	16	8	8
COMMIT					
OUTPUT(A)	16	16	16	16	8
OUTPUT(B)	16	16	16	16	16

#### Is this bad?

Action	t	Mem A	Mem B	Disk A	Disk B	
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	
COMMIT						
OUTPUT(A)	16	16	16	16	8	Crash !
OUTPUT(B)	16	16	16	16	16	

#### Is this bad?

#### Yes, it's bad: A=16, B=8

Action	t	Mem A	Mem B	Disk A	Disk B	
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	
COMMIT						
OUTPUT(A)	16	16	16	16	8	Crash !
OUTPUT(B)	16	16	16	16	16	

#### Is this bad?

Action	t	Mem A	Mem B	Disk A	Disk B	
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	
COMMIT						Crash
OUTPUT(A)	16	16	16	16	8	
OUTPUT(B)	16	16	16	16	16	
## Is this bad?

## Yes, it's bad: lost update

Action	t	Mem A	Mem B	Disk A	Disk B	
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	
COMMIT						Crash
OUTPUT(A)	16	16	16	16	8	
OUTPUT(B)	16	16	16	16	16	

## Is this bad?

Action	t	Mem A	Mem B	Disk A	Disk B	
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	Crash
COMMIT						
OUTPUT(A)	16	16	16	16	8	
OUTPUT(B)	16	16	16	16	16	

## Is this bad?

Action	t	Mem A	Mem B	Disk A	Disk B	
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	Crash
COMMIT						
OUTPUT(A)	16	16	16	16	8	
OUTPUT(B)	16	16	16	16	16	