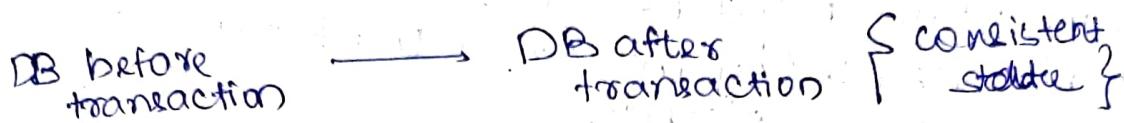


## Module - 5

### Transaction Management

- It is a program unit whose execution may change the content of database.
- like read, write, update or delete



A 300	B 400	100 from A → B
Read(A)	Read(B)	Before: 700 After: 700
A = A - 100 [200]	B = B + 100	
Write(A)	Write(B) [500]	

Transactions satisfy the ACID properties.

#### ACID Properties

##### ① Atomicity:

- It states that the transaction must be atomic i.e. all will be executed or none.
- If any part of transaction fails, everything will be rollback

Abort: Changes to DB are not visible

Commit: Changes are visible.

Ex:

T <sub>1</sub> & T <sub>2</sub> tran.	100% From A to B
A : 500	B : 200
T <sub>1</sub>	
Read(A)	T <sub>2</sub>
A = A - 100	Read(B)
write(A)	B = B + 100
Commit	write(B)
	Commit

## ③ Consistency:

Transaction should maintain data integrity, regardless of whether it succeeds or fails.

## ④ Isolation:

Transaction must be isolated i.e. a transaction cannot be used by the second transaction until the first one is completed.

→ Must be independent i.e. concurrent transaction do not interfere with each other.

## ⑤ Durability:

→ Once a transaction is committed, its changes are permanent, even in the event of system failure.

## ⑥ Operations of Transaction:

### (i) Read (X)

→ Used to read the value of X

→ Store it in a buffer in the main memory  
for further processing

### (ii) Write (X)

→ Used to modify the database

→ Change the state of the database as part of the transaction

### (iii) Commit

→ Used to maintain the data integrity in the database.

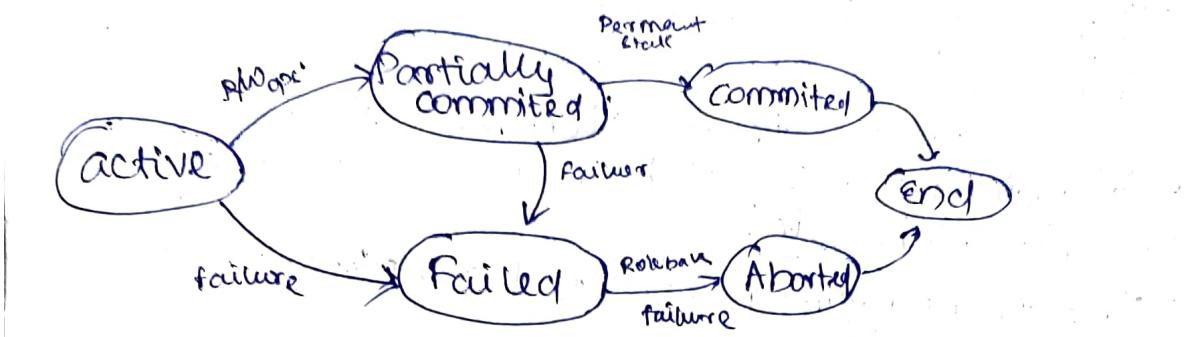
→ Makes all changes permanent in the database.

### (iv) Rollback:

→ Aborts the transaction and prevent changes made to the database.

→ Ensures atomicity -

# States Of Transaction



## (i) Active State:

- The active state is the first state of every transaction.
- In this state, transaction is being executed.
- Stores in buffer (local memory)

## (ii) Partially Committed:

- If all the query execution satisfies the ACID properties then it is committed
- If any system failures occurs it is terminated.

## (iii) Commit State:

- A transaction is said to be in a committed state if it executes all the operations successfully.

## (iv) Failed State:

- If any of the checks made by the database recovery system fails, then the transaction is said to be in failed state.

## (v) Aborted State

- In the aborted state, it clears all the buffer (local memory) and make the database in its previous consistent state.
- Kills the transaction  
Restart the transaction

And finally it comes to end

## Deferred Update

- Updates are applied only at commit time
- Changes not visible
- Only redo is req. for recovery
- Slower
- Concurrency control are easier to implement

## Immediate Update

- Updates are applied during transaction execution
- Change may visible
- Both redo & undo req.
- Faster
- Requires stronger concurrency mechanism.

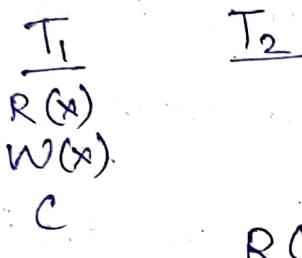
## Transaction Schedule

→ By default, commit is executed at last step  
→ Abort instruction on the last statement

- Chronological order of execution of multiple transaction.
- Series of one transaction to another.
- To maintain the database consistency & integrity

### ① Serial Schedule:

One transaction is executed completely before starting another.

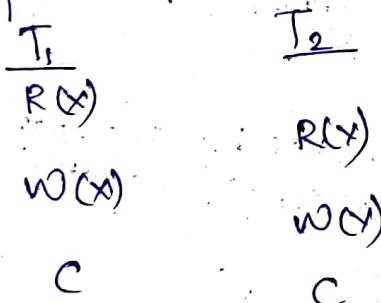


Adv:  
\* consistency

→ concurrent schedule

### ② Non-Serial Schedule:

- Operations from different transaction are interleaved.
- Improve performance but ~~risks~~ risks inconsistency



### ③ Serializable Schedule:

# A Serializable Schedule is a non-serial schedule that ensures the same result as a serial schedule.

- The serializability property ensures the correctness of a non-serial schedule by verifying that its outcome is equivalent to some serial schedule.
- It improves both resource utilization & CPU.

"Serial execution of a set of transactions preserves database consistency".

So a schedule is said to be serializable if it's equivalent to serial schedule.

## Testing of Serializability :-

→ Serialization graph is used for testing  
 ↓  
 Precedence graph

$$G = (V, E)$$

graph      set of vertices      set of edges

$E$  contains all edges  $T_i \rightarrow T_j$ :

① Create a node  $T_i \rightarrow T_j$  if  $T_i$  executes  $W(Q)$  before  $T_j$  executes  $R(Q)$

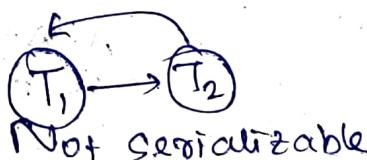
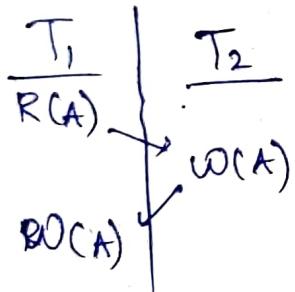
( $W \rightarrow R$ )

② Create node  $T_i \rightarrow T_j$  if  $T_i$  executes  $R(Q)$  before  $T_j$  executes  $W(Q)$

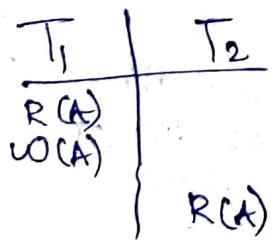
( $R \rightarrow W$ )

③ Create node  $T_i \rightarrow T_j$  if  $T_i$  executes  $W(Q)$  before  $T_j$  executes  $W(Q)$

( $W \rightarrow W$ )



Not Serializable



Serializable

## Concurrency:

→ It is a procedure in DBMS which help us in the management of two simultaneous process to execute without conflicts between each other.

\* Type of problems in concurrency.

① Dirty read.



II Dirty read

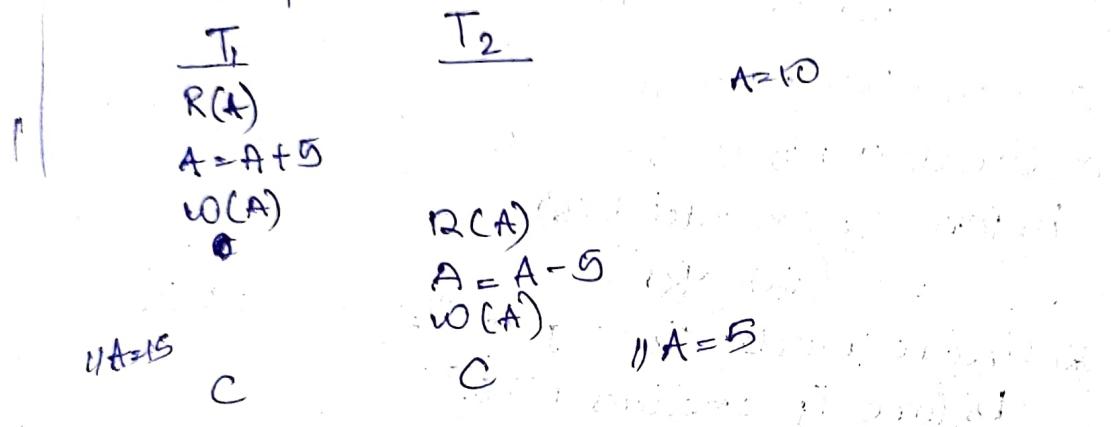
failure

## ④ Incorrect sum:

- Started when b/w two transaction when one start performs aggregate func<sup>n</sup>.

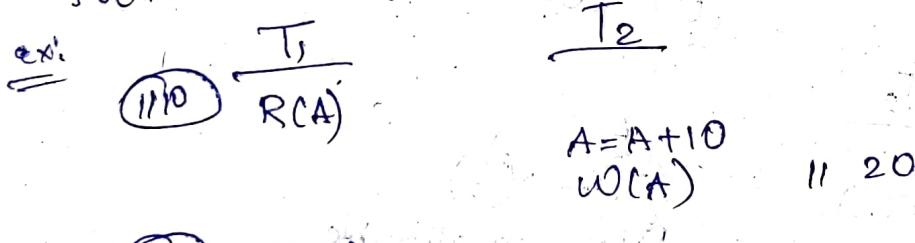
## ⑤ Lost Update:

Occurs when two transactions overwrote each other's update without proper synchronization.



## ⑥ Unrepeatable read:

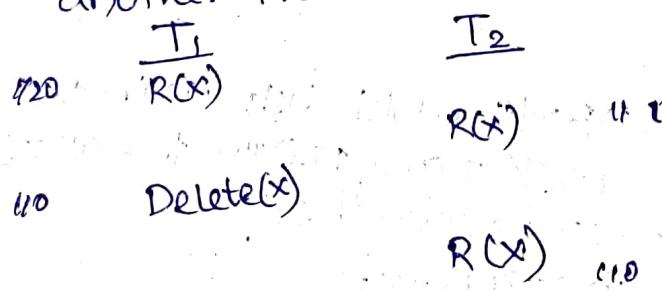
→ Occurs when a transaction reads the same data item twice, but gets diff. results due to another transaction's update.



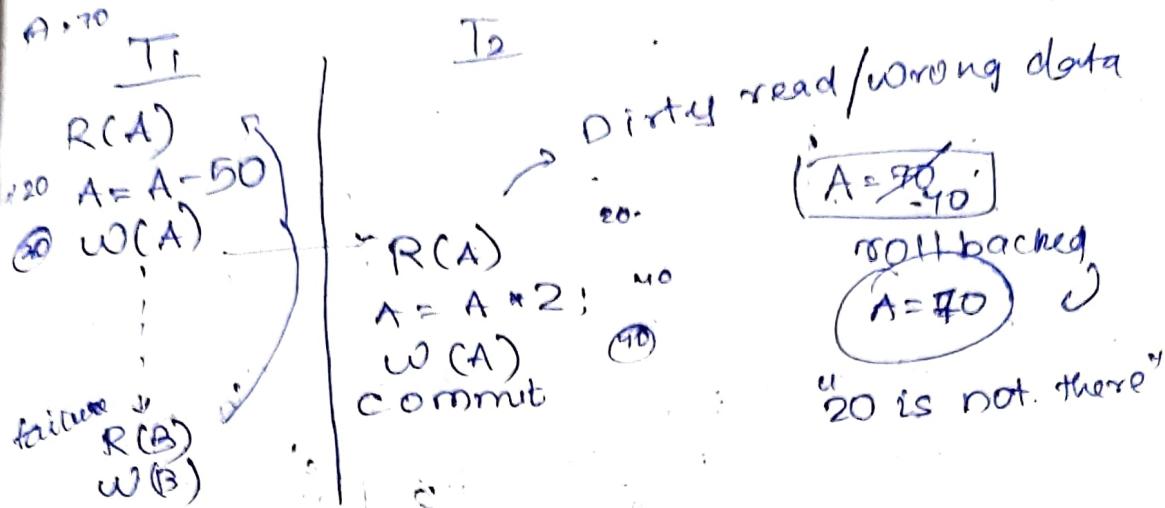
causing inconsistency.

## ⑦ Phantom read:

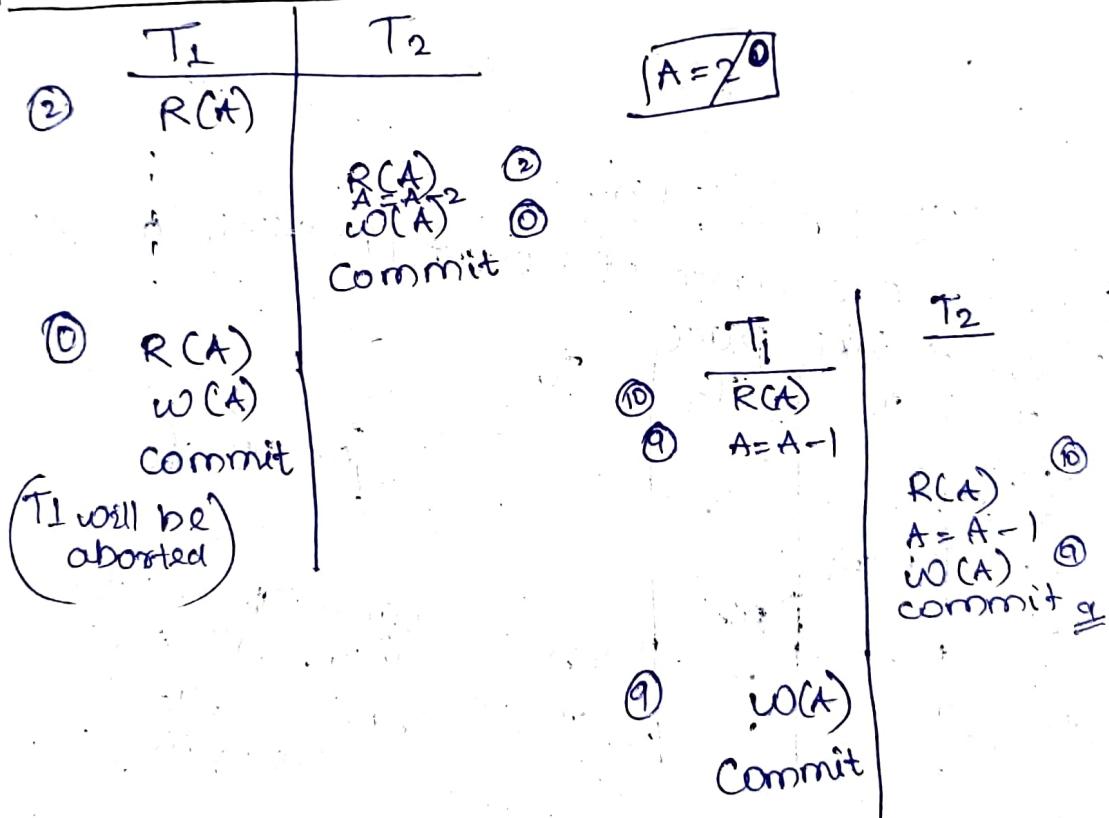
→ Occurs when a transaction retrieves different sets of rows in repeated queries due to another transaction's insertion or deletion.



## \* Write Read Conflict :-



## \* Read Write Conflict :-



## Conflict Pairs :-

$R(A) \quad R(A) \leftarrow \text{Non conflict}$

TWO are reduce, but only 1 is reduced.

$R(A) \quad W(A)$  }  
 $W(A) \quad R(A)$  }  
 $W(A) \quad W(A)$  }

Conflict

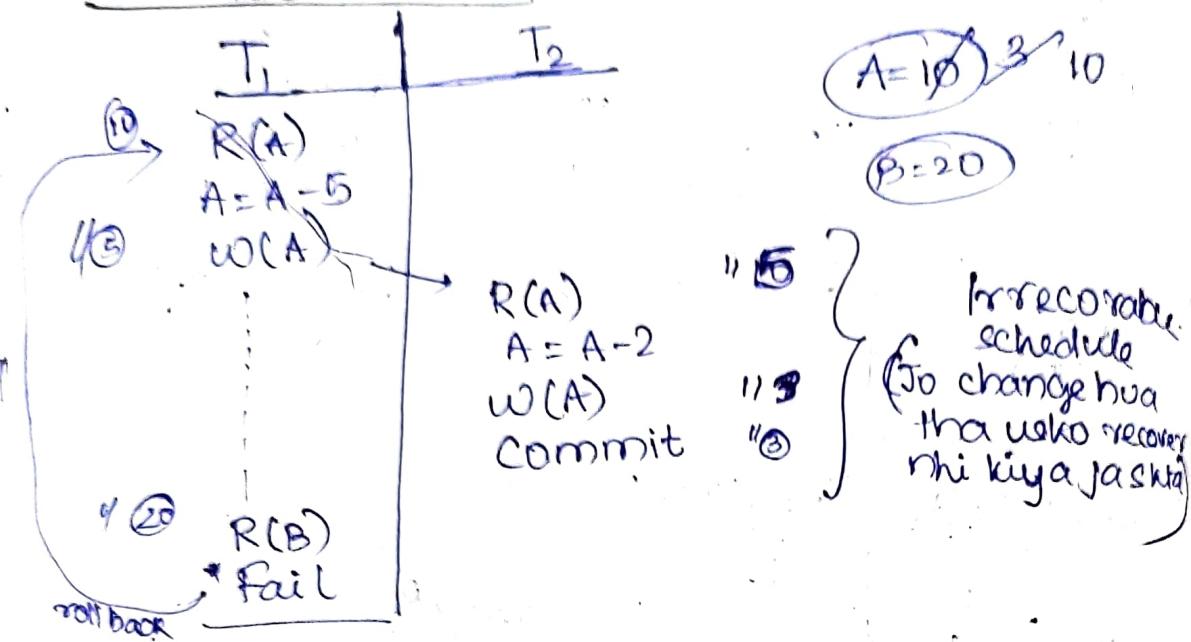
$R(B) \quad R(A)$  }  
 $W(B) \quad R(A)$  }  
 $R(B) \quad W(B)$  }  
 $W(A) \quad W(B)$  }

Non-conflict

$$A = 10 \text{ or } 9$$

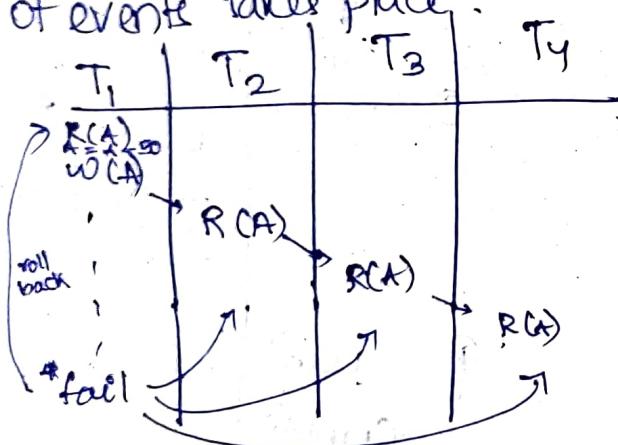
## ④ Recoverability

### \* Non-Recoverable:



### Cascading Schedule

Due to occurrence of 1 event multiple occurring of events take place

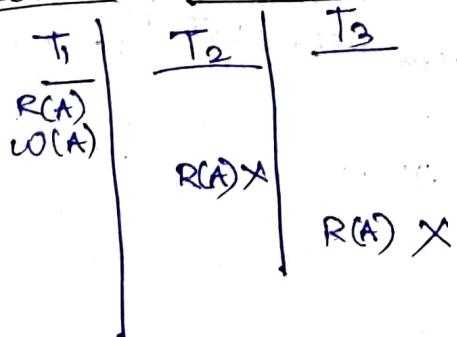


$$A = 10 \cancel{6} / 100$$

We will tell to abort all these  
This is called cascading

DisAdv: Failure

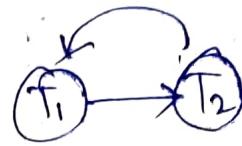
### Cascadless Schedule



Not allows for read until  $T_1$  is committed.  
But can write so it is a problem

## # Serializability:

$T_1$	$T_2$
$R(A)$	
$W(A)$	
	$R(A)$
	$W(A)$
	$W(A)$



Parallel (loops)  
↓ clone convert  
serializable

## Serializability

Conflict

View

No. of ways to make it

$$3! = 6$$

For 3 transaction serializable are:

$$T_1 \rightarrow T_2 \rightarrow T_3$$

$$T_1 \rightarrow T_3 \rightarrow T_2$$

$$T_2 \rightarrow T_3 \rightarrow T_1$$

$$T_2 \rightarrow T_1 \rightarrow T_3$$

$$T_3 \rightarrow T_2 \rightarrow T_1$$

$$T_3 \rightarrow T_1 \rightarrow T_2$$

## Conflict Equivalent

Conflict:  
 $R(A) \quad W(*)$   
 $R(*) \quad R(A)$   
 $W(*) \quad W(*)$

→ If a schedule  $S$  can be transformed into  $S'$  by a series of swapping of non-conflict instruction

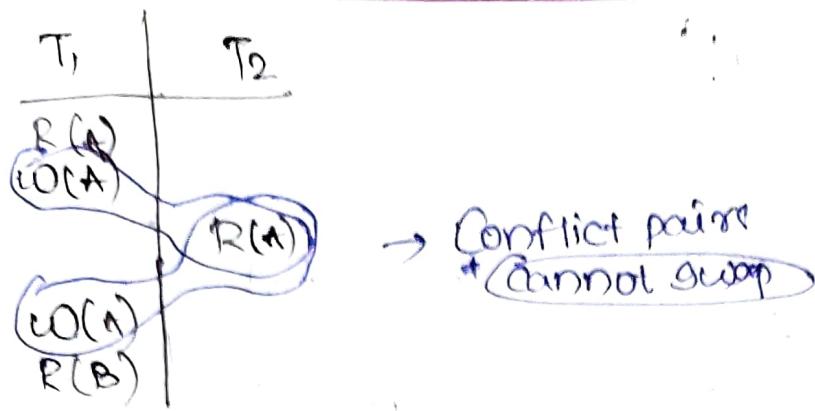
<u><math>S</math></u>	
$T_1$	$T_2$
$R(A)$	$W(A)$
$W(A)$	

$T_1$	$T_2$
$R(A)$	$W(A)$
$R(B)$	$R(A)$
$C(A)$	

<u><math>S'</math></u>	
$T_1$	$T_2$
$R(A)$	$W(A)$
$W(A)$	
$R(B)$	
	$R(A)$
	$W(A)$

Adjacent non-conflict pair  
Swap their pos?

$S \rightarrow S'$  serializable



T <sub>1</sub>	T <sub>2</sub>	T <sub>3</sub>
-R(A)		
	R(A) R(X)	R(C) R(X)
	R(C) W(Z)	W(C)
R(Z)		
W(X)		
W(Z)		

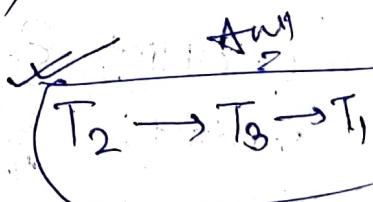
check conflict pairs  
on other two transaction

W(C)  $\rightarrow$  {R(C)} ?

No loop exist  
Conflict Serializable

Find In degree = 0.

$\hookrightarrow T_2$  (Remove it)

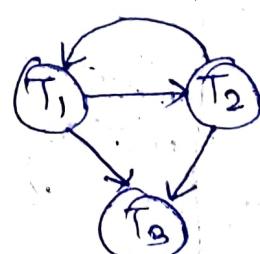


Indegree = 0

$\hookrightarrow T_3$  (Remove)

\* View Serializability

T <sub>1</sub>	T <sub>2</sub>	B	S
-R(A)			
	+C(A)		
W(A)			
		W(A)	



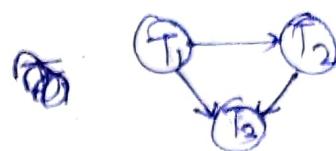
Not a conflict  
serializability  
as cycle occurs.

To overcome it we use view set.

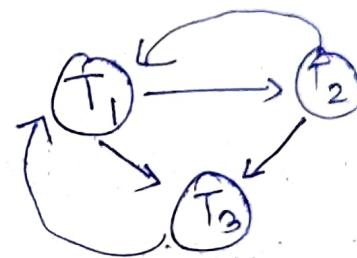
$S'$	$T_1$	$T_2$	$T_3$
	$R(A)$		
	$W(A)$		$W(A)$

$$S = S'$$

$$T_1 \rightarrow T_2 \rightarrow T_3$$



	$T_1$	$T_2$	$T_3$
	$R(A)$		
		$W(A)$	
	$W(A)$		
			$R(A)$
			$W(A)$



Non conflict  
serializable

### Concurrent Execution

If two transaction are running concurrently, the OS may execute  $T_1$  for a little while then perform a context switch, execute  $T_2$  for sometime and then switch back to  $T_1$  and so on.

→ Thus the CPU time is shared among all the transactions

### Advantage :-

- (i) Multiple transaction are allowed to run concurrently in the system.
- (ii) Increase processor & disk utilization
- (iii) Reduce avg response time
- (iv) Concurrency control schemes

## \* Concurrency Control

- Procedure in DBMS for managing simultaneous operations of transactions which execute without conflicting with each other.
- Used to address conflict pairs.

### Problems



### ④ Concurrency control Schemes :-

#### ① Lock-Based Protocols

- Mechanism to control concurrent access to a data item by various transaction is called lock.
- Data items can be locked in two modes:
  - ↳ Shared (S) mode : It can only be read
  - ↳ Exclusive (X) mode : It can be both read & write as well
- A transaction can proceed with its execution only after the request is granted.

#### # Compatibility Function

Request →

Grant ↓	G	S	X
		Yes	No
X	No	No	

S - S  
R - R ✓  
R - W ✗  
W - R ✗  
RW - W/R ✗

T <sub>1</sub>	T <sub>2</sub>
G(A)	X(A)
R(A)	R(A)
U(A)	U(A)

→ We can observe that if any transaction holds exclusive lock on the item then no other transaction can hold any lock on the same item until the first transaction does not release its X-lock on the item.

T<sub>1</sub>

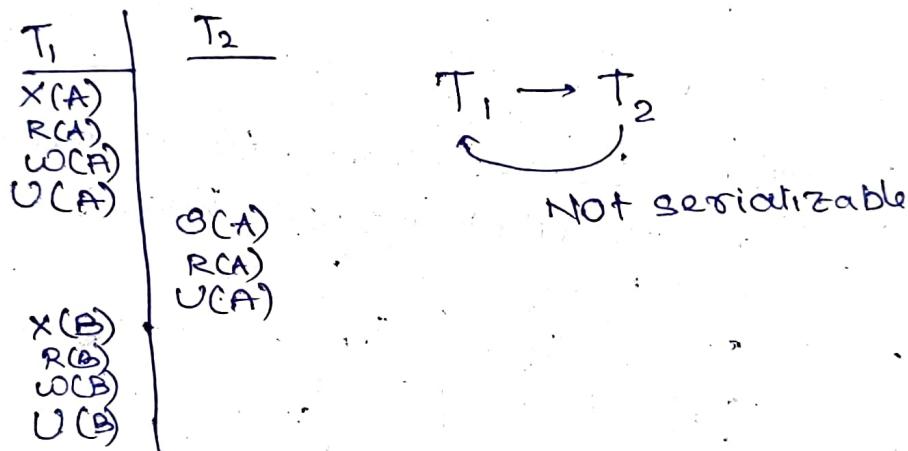
lock-X(B)  
read(A)  
 $B = B - 50$   
write(B)  
unlock(B)

T<sub>2</sub>

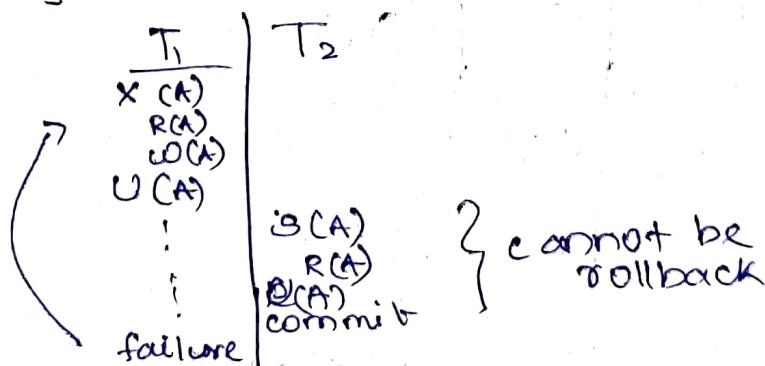
lock-S(A)  
read(A)  
unlock(A)  
lock-S(B)  
read(B)  
unlock(B)

### ④ Problems in Lock-Based protocols :

- 1) May not sufficient to produce only serializable Schedule

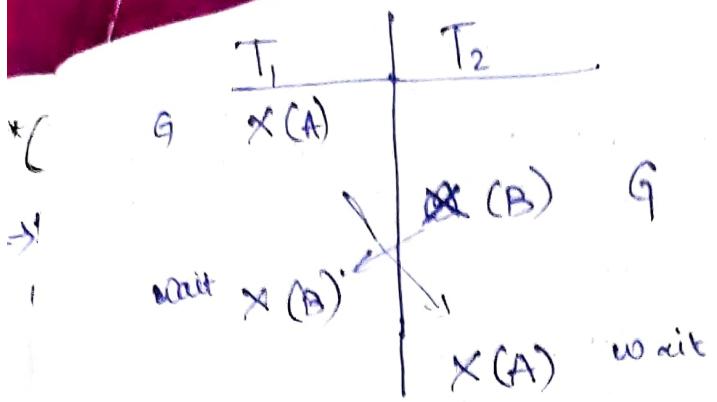


- 2) May not free from Irrecoverability



- 3) May not free from deadlock

"Deadlock" occurs when two or more transactions are waiting for each other to release locks on resources, none of them proceed as a result of causing Indefinite halt.



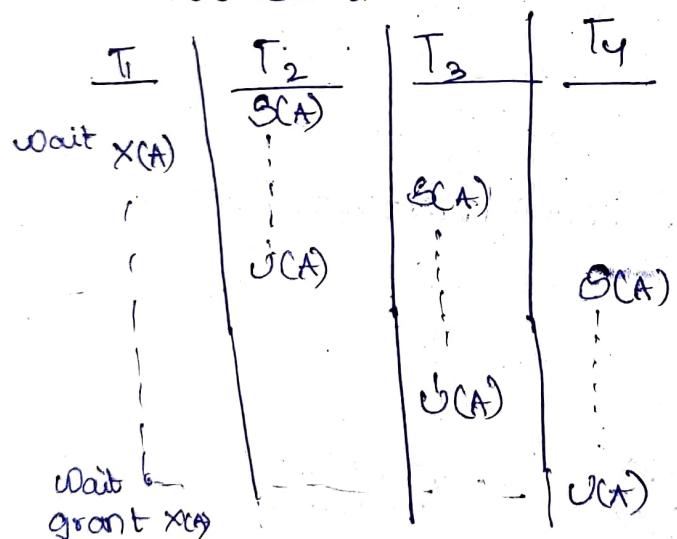
T<sub>1</sub> has resource A and wait for request resp.

T<sub>2</sub> has resource B & wait for req resp.

Neither T<sub>1</sub> & T<sub>2</sub> can proceed; as both are waiting for other to release a resource. This creates a circular wait, leading to deadlock

(4) May not free from Starvation

- Starvation: Occurs when a transaction waits till higher-priority transaction keep acquiring the resources it needs



\* Two phase locking Protocols:-

Works in two phases:-

(i) Growing Phase: Tran. can obtain locks but cannot release locks.

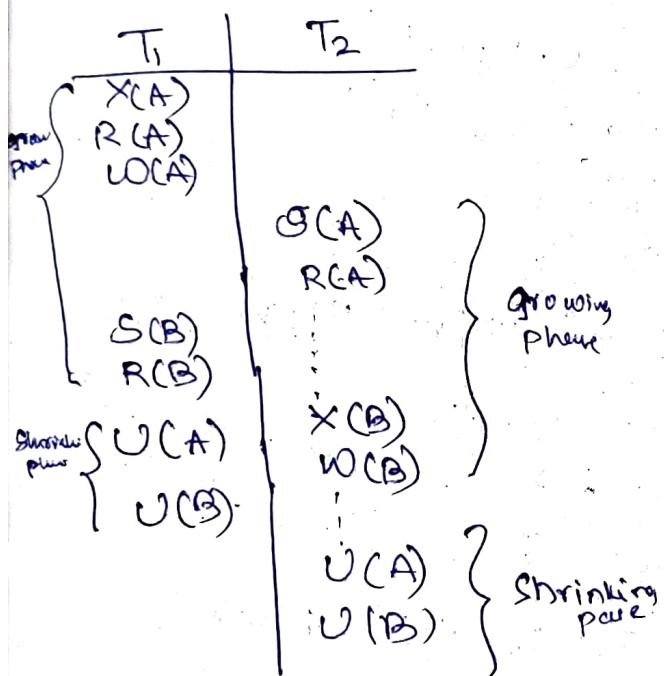
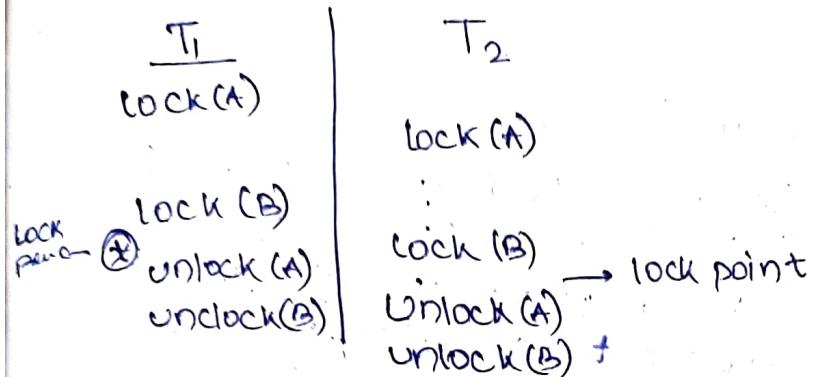
(ii) Shrinking Phase: Transaction can release locks but cannot obtain locks.

Ensures Serializability.

### Lock Points:

It is the exact moment in a transaction where acquiring of lock stops and releasing of acquired lock started.

Based on the order of execution



In  $T_1$ ,  $X(A)$  is granted then  $S(B)$  is granted

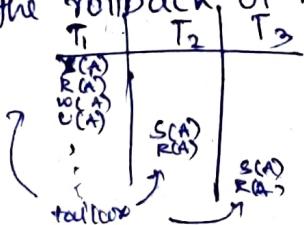
In  $T_2$ ,  $S(A)$  is wait till the  $X(A)$  completed then it will start.  
So it ensures serializability

$T_1 \rightarrow T_2$

### # Drawbacks of 2PL

- 1) May not free from unrecoverability
  - ↳ Not durability of commit.
- 2) Does not ensure free from deadlock
- 3) Not free from Starvation
- 4) Not free from Cascading Rollback

Cascading Rollback:- When a single transaction's failure leads to the rollback of multiple dependent transactions.



\* (vi)  $w(B)$  by  $T_2$

$$\rightarrow \left\{ \begin{array}{l} R\_TS(B) > T_2 \\ BO > 20 \end{array} \right. \begin{array}{l} (\text{True}) \\ \text{roll back } T_2 \end{array}$$

\* (vii)  $w(A)$  by  $T_3$

$$\left\{ \begin{array}{l} R\_TS(A) > T_3 \\ IO > 30 \end{array} \right. \begin{array}{l} (\text{False}) \\ \left. \begin{array}{l} w\_TS(A) > T_3 \\ IO > 30 \end{array} \right. (\text{false}) \end{array}$$

$$\text{Set } w\_TS(A) = TS(F_2) = 30$$

{ Database Recovery :-

- Failure Classification :-

- (i) Transaction failure:
  - ↳ logical error (Bad input, overflow, resource limit)
  - ↳ System error (deadlock, starvation)

(ii) System Crash

- (iii) Disk failure

# Recovery Algorithm :-

→ Techniques to ensure database consistency and transaction atomicity & durability despite failure.

Log Based Recovery :-

→ Most widely used structure for recording database modification is the log.

→ Log is a sequence of log records, recording all the update activities in the database.

It contains info about:

(i) Transaction Start  $\langle T_1, \text{start} \rangle$

(ii) Write Operations  $\langle T_1, x, \text{OldValue}, \text{NewValue} \rangle$

(iii) Commit  $\langle T_1, \text{commit} \rangle$

(iv) Abort  $\langle T_1, \text{Abort} \rangle$

### Deferred

- Records all modifications to log, but defers all the write op. until the transaction commits
- Only redo logs are req.
- Simpler recovery process
- Perf. may be slower
- Safer against inconsistency

### Immediate

- Updates are applied to the database as soon as a write operation occurs
- Both redo & undo are req.
- Complex one
- Faster
- More prone to inconsistency

### Undo Database Recovery :-

Restores the value of all data items that are updated, to their old value.

$\langle T_1, A, \text{oldValue}, \text{newValue} \rangle$

$T_1$  New Value = 200 fails ~~Before Commit~~  
Undo restores Old Value = 100 to A

→ Works backward from the point of failure

### Redo Database Recovery :-

Get the value of all data item that are updated to new values

$\langle T_1, A, \text{OldVal}, \text{NewVal} \rangle$

$T_1$  NewVal = 400 , commits, but change wasn't flushed to disk due to crash

Redo writes New Value = 400 to B

→ Redo works forward from the point of failure

### Strict 2PL

basic 2PL all exclusive, holds <sup>locke are</sup> until commit/Abort

### Rigorous 2PL

Basic 2PL with all exclusive & shared locks are held until commit/Abort

$T_1$	$T_2$
$x(A)$	
$r(A)$	
$w(A)$	$s(A)$
$c$	$r(A)$
$v(A)$	

overlapping roll back demand

core cables  
strict recoverable

### Conservative 2PL

Before starting the transaction takes all the locks

resolve deadlock problem