


TRANSACTION & Concurrency Control

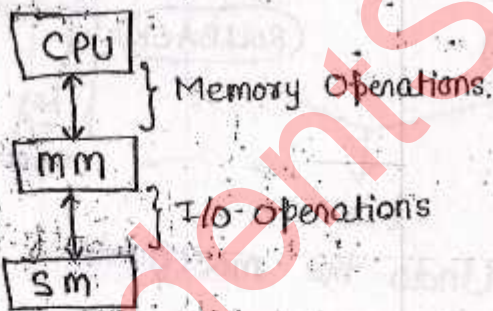
Transaction: set of logically related operations which is used to perform some "unit of works"
 ↳ "Set of logically related operations."

» READ(A): Accessing data item A from DB to programmed variable.



» Write operation: updating the details into DB.

Read/Write operations and I/O operations




» Assume some transaction T, t/f Rs 500 from A/c A to A/c B.

Trans T: Transfer ₹ 500 from A to B

```

Begin Trans
Read A
A = A - 500
Write A
Read B
B = B + 500
Write B
Commit
    
```



» Commit () is a memory operation that specifies successful completion of transaction.

» To maintain INTEGRITY, transaction should satisfy following conditions.

(ACID Properties)

Atomicity :- Transaction should execute all the operations successfully (commit) or none of them.
 ↳ sw modules (automated)

Begin Trans
 Read A
 $A = A - 500$
 Write A
 Read B
 $B = B + 500$
 Write B
 Commit


~~A = 1000~~
 A = 500

~~B = 2000~~
 B = 2000

Total: 3000
 Total: 2500

FAILURE

To maintain ATOMICITY



Recovery mgmt Component
 Log File maintenance

i.e
 Undo the modifications done till now

can be due to


- Power failure
- sw crash
- H/w crash etc
- DISK crash

Log file is not permanent

Recovery :-

Shadow Recovery :- Whenever a transaction begins, it maintains a log file corresponding to transaction, if transaction fails, log file searched to check the previous fresh point and if transaction commit, log file deleted.

Consistency :- DB should be consistent before and after the transaction execution.



→ taken care by the user.

→ if failure occur due to some major problem (disk crash)

the automated recovery may not be possible, in such cases user takes care of data consistency.

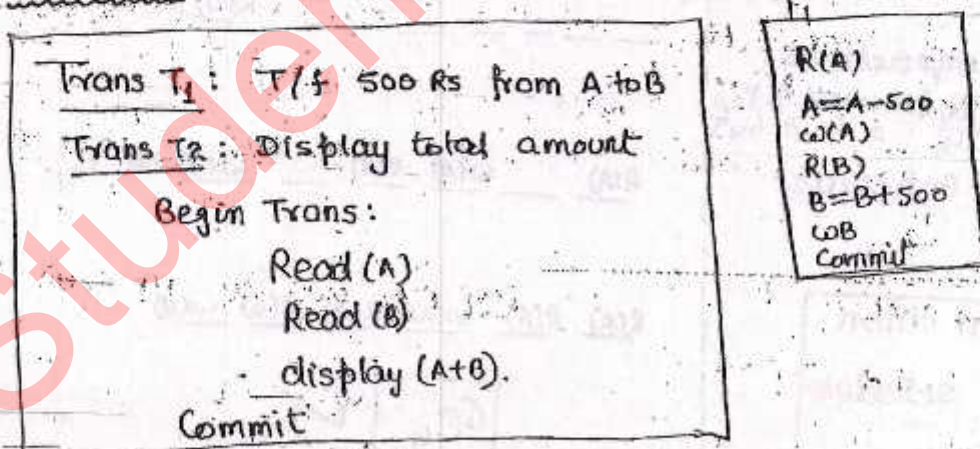
(Eg. using fresh disk etc)

Factors affecting consistency:

- ① Atomicity
- ② Isolation
- ③ Durability

ISOLATION :-

Schedule : Time order sequence of two or more trans.



Schedule means these ^{two} transactions can execute simultaneously.

Schedule

Serial schedule

Concurrent schedule

Execution of one trans. possible only after completion of other

Interleaved execution of two or more transactions

Execution of 2 trans., one after other

T ₁	T ₂
R(A)	
W(A)	
R(B)	
W(B)	
	R(A)
	W(A)



OR

T ₁	T ₂
	R(A)
	W(A)
R(A)	
W(A)	
R(B)	
W(B)	



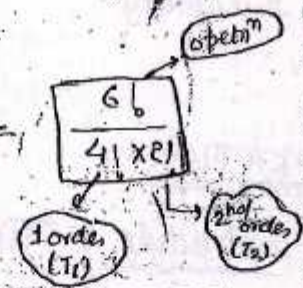
T ₁	T ₂
R(A)	
W(A)	
	R(A)
R(B)	
W(B)	
	W(A)

T ₁	T ₂
R(A)	
R(B)	
	R(A)
	W(A)
	R(B)
	W(B)

T ₁	T ₂
R(A)	
W(A)	
	R(A)
R(B)	
W(B)	

T ₁	T ₂
R(A)	
W(A)	
R(B)	
W(B)	
	R(A)
	W(A)

Total no. of schedules possible with 2 trans.



T₁, T₂ Two transactions then T₁ followed by T₂ or T₂ followed by T₁

Two serial schedule

$$R(A) \quad W(A) \quad R(B) \quad W(B) \Rightarrow {}^6C_4$$

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

$${}^6C_4 \quad \checkmark$$

X It is not ${}^6C_4 \times 4! \times 2 \times 2!$

T₁ ∴ R(A) W(A) R(B) W(B) — 1 order only
 T₂ ∴ R(B) R(B) — 1 order only

In General: n transaction then n! serial schedule possible

If m and n operations of Trans T_1 and T_2

Then

How many Concurrent schedules are possible

⇒ $\boxed{\binom{m+n}{m}}$ or $\boxed{\binom{m+n}{n}}$

If $m, n,$ and p operations of Trans T_1, T_2, T_3

Concurrent Schedules = $\binom{m+n+p}{m} * \binom{n+p}{n} * \binom{p}{p}$

out of $(m+n+p)$, m people made to sit

= $\binom{m+n+p}{m} * \binom{n+p}{n} * 1$ #

out of remaining $n+p$, n made to sit

- There is no inconsistency if schedule executed in serial order.
- Serial order leads to less throughput.
- To avoid this problem, we go for concurrent schedule. CPU ideal time now is reduced, hence leads to better throughput!

A=1000
B=2000
2500

T_1	T_2
R(A)	
W(A)	
R(B)	
W(B)	
	R(A) 500
	R(B) 2500
	<u>3000</u>

Consistent

T_1	T_2
R(A)	
W(A)	
	R(A) 500
	R(B) 2000
	<u>2500</u> X
R(B)	
W(B)	

Inconsistent Concurrent Schedule

Concurrency Control Component should not allow any inconsistent schedule.

How??

How it come to know about a particular schedule consistent or inconsistent??

Serial Schedule ^{always} results in consistency.

If any concurrent schedule reducible or is similar to any serial schedule (equivalent) then it is consistent.

If any concurrent transaction T_i equal to any serial schedule $(T_1 \rightarrow T_2)$ or $(T_2 \rightarrow T_1)$ it is said to follow ISOLATION.

ELSE, it is not a consistent concurrent schedule and may violate isolation (produce inconsistent result).

⇒ Isolation means concurrent execution of two or more transactions should be equal to any serial schedule. Then only we say transaction is isolated.

④ DURABILITY :- Database should be able to recover under any case of failure.

⇒ Redundancy introduced to enforce durability

⇒ Both h/w and sw responsibility " "

⇒ RAID: Redundant Array of Independent Disk.

RAID-0 — No redundancy, only one copy

RAID-1 — Mirroring — one fail other used

RAID-2 —

↑ % of risk | update at 2 places | Cost doubles

PROBLEMS Because of Concurrent Execution:-

① WR Problem:

(Read after write operation)

T ₁	T ₂
W(A)	
⋮	R(A)

(Uncommitted Read/Transaction) OR Dirty Read

Uncommitted-Read: Transaction

T₂ reads data item A that is updated by Un-committed transaction.

T ₁	T ₂
R(A)	
W(A)	
	R(A)
	R(B)
R(B)	
W(B)	

Uncommitted read

↓
leads to inconsistency

T ₁	T ₂
R(A)	
W(A)	
	R(A)
R(B)	
W(B)	
	R(B)

Uncommitted read still Transactions Consistent.

⇒ Uncommitted read may create problem sometimes, but are effective in awareness to people.

② RW Problem (Write-After Read)

Ex:- Library DB

A: # of copies of DB book

```

R(A)
if(A > 0)
  A = A - 1
  W(A)
commit
else
  Abort
  
```


T ₁	T ₂
⋮	⋮
W(A)	R(A)
R(A)	W(A)

T ₁	T ₂


(III) Write-Write Problem (write after write problem)

Assume Two Transactions are there

Before commit T_1 updates A and B which is already updated by T_2 .



T_1



$A, B \text{ are } 0$
 $A, B \text{ are } 20$

or $T_2 \rightarrow T_1$

How assume execution

So T_1 follows T_2 ($w(A) \rightarrow T_2$), final updation done by T_2

occuring concurrent

problems caused by write
execution

Write Read Problem : (Read after write problem)

T ₁	T ₂
W(A)	R(A) (Uncommitted Read) or (Dirty Read)

Uncommitted Read

Transaction T₂ reads data A that is updated by uncommitted transaction T₁. Uncommitted read results in inconsistency.

T ₁	T ₂
W(A) commit	R(A) (committed Read)

Uncommitted Read may not create inconsistency

T ₁	T ₂
R(A) W(A)	
R(B) W(B)	R(A) R(B)

inconsistent schedule

T ₁	T ₂
R(A) W(A)	(R(A))
R(B) W(B)	R(B)

Uncommitted Read

Consistent schedule

Read Write Problem write after read

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

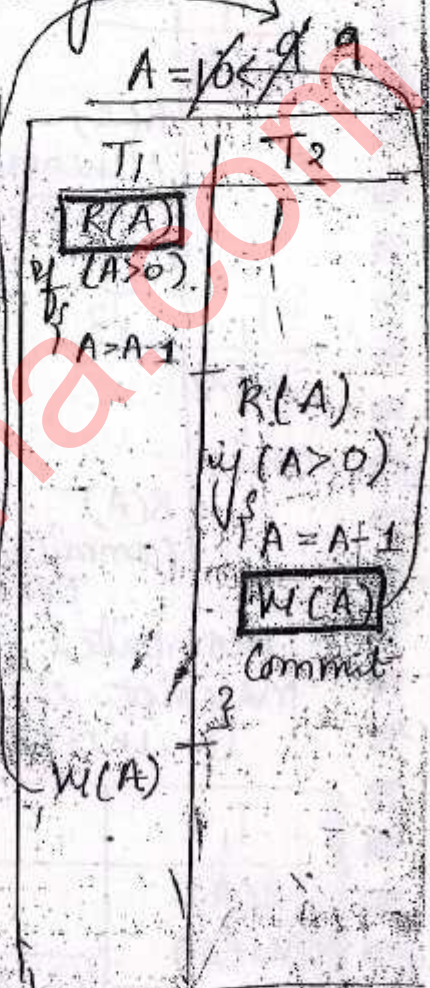
Ex. library Database

A : # of copies of DBMS books

Trans T_1

```

R(A)
if (A > 0)
    A = A - 1
W(A)
    commit
else
    No Books
    
```



If serially executed
no inconsistent
result.

ie. Before commit
of T_1, T_2
perform write
operations.


Write Write Problem

T_1	Transaction T_2
$W(A)$	updating data item (A) that is already updated by uncommitted Trans T_1 .
$W(A)$	

T_1 : Set A, B as 10 (T_1 : $W(A), W(B)$)
 T_2 : Set A, B as 20 (T_2 : $W(A), W(B)$)


$T_1 \rightarrow T_2 \Rightarrow A = B = 20$. T_1 followed by T_2
 $T_2 \rightarrow T_1 \Rightarrow A = B = 10$. T_2 followed by T_1

$[W-W]$ may or may not suffer from consistency



A = 20
B = 10

Inconsistent Schedule



A = 20
B = 20

Consistent Schedule

Schedule

Inconsistent

Consistent

- RW (OK)
- WR (OK)
- WW problem.

→ No problem

WW problem (Not equal to any serial schedule)

T1	T2
R(A)	
W(A)	
	W(A)
	Commit

- (1) RW Problem } concurrent & not equal to any serial execution
- WR Problem }
- WW Problem }

(2) Other than above problems we have lost-update problem
A = 0

no concurrent execution problem because it is

B) T1 → T2 serial schedule

B) → if failure occurred

T1 log
A_{old} = 0
A_{new} = 10

T2 log
A_{old} = 10
A_{new} = 20

A = 20

Rollback occurs

Because of T1 failure, updation of T2 is lost. Problem because of this.

Very lost update problem ^{or} write-write problem
 Simultaneous write-write operation is lost update problem and schedules non-serializable possible.

Simultaneous operation W-W is called as W-W problem. \rightarrow NO

Two conditions required to be W-W problem

- 1) WW operation exists & if
- 2) 'S' Schedule not equal to any serial schedule.

T ₁	T ₂
W(A)	W(A)
W(B)	W(B)

A = 20
B = 20

Consistent

← lost update possible but not W-W problem

- \rightarrow WW operation ✓
- \rightarrow S not equal to serial schedule result ✓

Will be W-W problem if schedule gives inconsistent result.

Classification of Schedule

* Schedule is said to be correct only if it satisfies

- (1) Recoverability (2P)
- (2) Serializability

which is a consistent serial schedule

T ₁	T ₂
W(A)	
W(B)	
	W(A)
	W(B)

Consistent or Inconsistent

T ₁	T ₂
w(A)	w(A)
w(B)	w(B)

Consistent

#

When we say

RW Problem } exists — it means
 WR " }
 ww " }

Concurrent execution of transactions, not equal to serial execution.

④ Lost Update Problem :-

A=0 — Initially

T ₁	T ₂
R(A)	
w(A)	
	w(A)
	Commit

T₁.log
A_{old}=0
A_{new}=10

~~T₂.log
A_{old}=0
A_{new}=20~~


→ T₂ commit then T₂.log lost.
If T₁ commits

Failure

⇒ Transaction T_2 lost updation becoz if T_1 failure.

⇒ Problem occurs becoz simultaneous write-write operations.

Classification of Schedules:




∴ Protocols must satisfy both recoverability & L, serializability.

Recoverability:-

IRRECOVERABLE: Rollback of committed transaction is irrecoverable.

∴ It is not possible to recover a committed transaction. Because once a T commit, its log no more exists.

→ Why rollback of transaction required???



DEPENDENCY :-

- ①
- | T ₁ | T ₂ |
|----------------|----------------|
| W(A) | |
| ⋮ | |
| Commit | R(A) |
- Updated value of A, read before Commit. So it is Uncommitted Read.
Hence T_2 depend on T_1
- ②
- | T ₁ | T ₂ |
|----------------|----------------|
| W(A) | |
| Commit | |
| | R(A) |
- A value read after commit, no dependency here b/w two transacks. No dependency
- ③
- | T ₁ | T ₂ |
|----------------|----------------|
| W(A) | |
| Rollback | |
| | R(A) |
- A value modified [A=0] but is Rollbacked soon to the back. near older value i.e (A=0) only No dep
- ④
- | T ₁ | T ₂ |
|----------------|----------------|
| W(A) | |
| | W(A) |
- No modification depend on other
→ A modified by T₁ then updated by T₂
No problem. No dep
- ⑤
- | T ₁ | T ₂ |
|----------------|----------------|
| W(A) | |
| | W(A) |
| | R(A) |
- ⇒ Here it is not uncommitted read. uncommitted
∵ Uncommitted read, specify modification by one transaction and reading by another. No dep
- ⑥
- | T ₁ | T ₂ |
|----------------|----------------|
| W(A) | W(B) |
| R(B) | R(A) |
- ⇒ T_1 depends on T_2 } Dependency.
 T_2 depends on T_1 }

Possibility to get Irrecoverability :-

T ₁	T ₂
W(A)	
	R(A)
C/R	Commit

Commit Rollback

T ₁	T ₂
W(A)	↑
Rollback	RIA) Commit


IRRECOVERABLE

T ₁	T ₂
W(A)	
Commit	RIA) Commit

IRRECOVERABLE

T ₁	T ₂
W(A)	RIA)
c/r	Rollback

RECOVERABLE



HOW TO GET RECOVERABLE SCHEDULE:-

RECOVERABLE SCHEDULE:-

(Rollbacking of Uncommitted Transaction)

Restriction to put on schedule, to make it RECOVERABLE ONLY.

If T₂ depend on T₁, commit of T₂ should be delayed until commit/rollback of T₁.

↳ If T₁ rollbacked/commit → T₂ also can be rolled back/commit.

T ₁	T ₂
W(A)	
C/R	RIA)

Commit

⇒ " If Tran T₂ reads data item A that is updated by T₁ and T₂ commit operation should be delayed until C/R of T₁. " Then T₁ is **RECOVERABLE**

NOTE:

Recoverable schedule may / may not be free from

- RW Problem
- WR "
- WW "
- Lost Update Problem

Illustration:

T ₁	T ₂
R(A)	
W(A)	R(A)
	R(B)
R(B)	
W(B)	
Commit	Commit

⇒ In this scenario, T₂ depends on T₁, i.e. Uncommitted read.
 Commit of T₂ delayed, till commit/rollback of T₁. Thus schedule is **RECOVERABLE**

BUT

It suffer from **WR PROBLEM**

2) Cascading rollbacks are possible:-

T ₁	T ₂	T ₃	T ₄
W(A)			
↑	R(A)		
↑	W(A)		
↑		R(A)	
↑			R(A)
Roll back			
* Failed			

only T₁ failed, but due to dependency among transactions all T₁, T₂, T₃, T₄ ... rollbacked.

Whatever the work done by other transaction is lost. Hence a wastage of CPU execution time and I/O cost.

⇒ Recoverable schedule may cause cascading rollbacks:-

Cascadeless Rollback Recoverable Schedule :-

(Recoverable schedule && No cascading Rollbacks)

T ₁	T ₂
W(A)	
C/R	R(A)

Trans. T₂ should not be allowed to read until modified value committed/rolledback.

⇒ Not allowing any uncommitted Reads

⇒ If one Trans (T₁) updates some value, that value should not be read (allowed to read) by other transaction until former transaction commit/rollback.

⇒ "If Trans. T₁ updates a data item A, then other transaction T₂ is not allowed to read the data item until commit/rollback of T₁"

T ₁	T ₂
R(A)	R(A)
	W(A)
R(B)	W(B)
R(B)	C/R

Cascadeless Rec. Sch.

T ₁	T ₂
W(A)	R(B)
	W(A)
Commit	Commit

→ Cascadeless Rollback Recoverable Schedule.

NOTE :-

① Cascadeless Rollback Recoverable schedule is free from


Write Read (WR) Problem.

But RW, WW, lost update problem may be possible.


due to non-existence of equivalent serial schedule for concurrent schedule.

Concern of Recoverability

→ These problems controlled by



Whenever WW possible \exists possibility of lost update problem.



If some Trans. T perform write on data item A, then other Trans. T' not allowed to perform read of A or write of A (W(A)/R(A)) until Commit or Rollback of Trans. T.

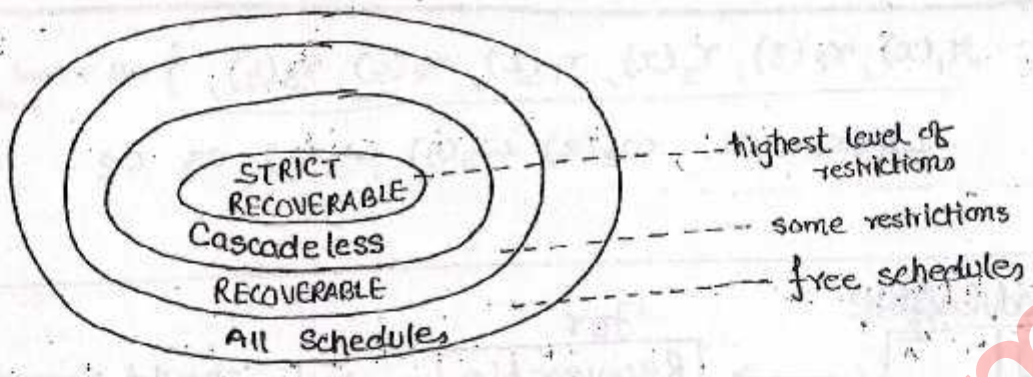
STRICT RECOVERABLE SCHEDULE

- » Removal of lost update problem leads to removal of WW problem also.
- » In STRICT RECOVERABLE SCHEDULE, there are
 - Recoverable Schedule
 - \exists No cascading Rollbacks
 - No WR Problem
 - \exists No lost Update Problem
 - \exists no WW problem

BUT RW Problem is still possible.

#

Classification of Schedules



Q:

Which is TRUE??

- a) Irrecoverable
- b) Recoverable but not cascadeless
- c) Cascadeless but not strict
- d) strict schedule


Q:

SI: $r_1(x) r_2(z) r_3(z) r_3(x) r_3(y) w_1(x) w_3(y) r_2(y)$
 $w_2(z) w_2(y) c_1, c_2, c_3$
 Schedule Commit 1 Commit 2 Commit 3

⇒ 3 Trans. involved.

T ₁	T ₂	T ₃
R(x)		
R(z)	R(z)	
w(x)		R(x) R(y)
	R(y) w(z) w(y)	w(y)
Commit	Commit	Commit

⇒ Uncommitted Read

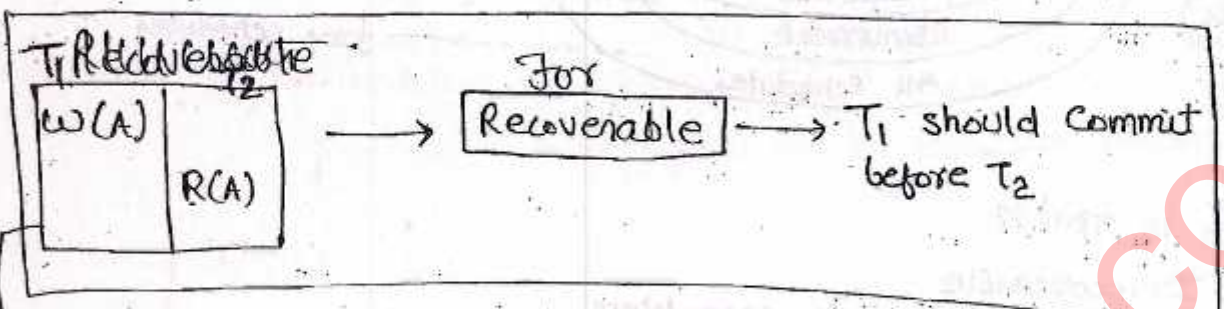


Hence it is

Irrecoverable

ii) S2: $r_1(x), r_2(z), r_3(x), r_1(z), r_2(y), r_3(y)$, } all read no problem


$w_1(x) \quad c_1 \quad w_2(z) \quad w_3(y) \quad w_2(y) \quad c_3 \quad c_2$



∃ no Write Read sequence, so it is

- ✓ — Recoverable and
- ✓ — Cascadeless

If WW sequence (write after write) also not there then it will become Strict Recoverable (Here it is not)



T_1	T_2
$w(A)$	
c/R	$R(A)$

↓
if this is not followed then it is not Cascadeless.

S2: → Recoverable ✓
→ Cascadeless ✓
→ Strict ✗

Imp

(iii) S3: $r_3(x), r_2(x), \omega_3(y), \omega_2(y), r_2(y), e_1, e_2$
 $\omega_3(x), r_2(x), \omega_1(y), r_2(y), \omega_2(x), c_3, e_1, c_2$
└ Recoverable ✓
└ Cancellable X

(iv) S4: $r_1(x), r_2(x), \omega_1(y), \omega_2(y), r_2(y), e_1, c_2$
└ Rec ✓
└ Cas ✓
└ Strict X

StudentSuvidha.com

Schedule

Classification Based on Serializability

Serial Schedule	Serializable Schedule
<ul style="list-style-type: none"> * No Concurrency * Every serial schedule are correct/consistent 	<ul style="list-style-type: none"> * concurrency allowed * concurrent schedule should be equivalent to any serial schedule execution * All serializable schedules are consistent * No $\begin{bmatrix} RW \\ WR \\ WW \end{bmatrix}$ problems, not possible, but cascading rollback (irrecoverability) and lost update may be possible.

OTE

⇒ If a schedule is BOTH RECOVERABLE and SERIALIZABLE, then

⇒ No WR, RW, WW problem

⇒ No Irreversibility

BUT

⇒ Lost Update

⇒ Cascading Rollback } are possible

⇒ STRICT ~~RECOVERABLE~~ RECOVERABLE + SERIALIZABLE

GOAL OF CONCURRENCY CONTROL

» GOAL OF CONCURRENCY CONTROL TECHNIQUE

STRICT RECOVERABLE AND SERIALIZABILITY

only suffer from

RW problem

which is ultimately overcome by Serializability.

6th Nov 2011


Schedules based on SERIALIZABILITY

① Conflict serializable schedule :-

⇒ Conflict pair :- On swapping operations of schedule equivalent. Then it is called non-conflict pair.



S1 : T₁: R[A] T₂: R[A]

S1 ≡ S2



Both schedules equal so we call them NON-CONFLICT PAIR.

» Initial Read ⇒ Read operation from Database.




S1 ≠ S2

CONFLICT PAIR

Conflict Pairs:-

- (1) T_1 :
- (2) T_1 :
- (3) T_1 :
- (4) $T_1: W(A) \quad T_2: W(A)$ ————— Conflict Pair
- (5) $T_1: R(A)/W(A) \quad T_2: R(B)/W(B)$



Pair of operations is said to be Conflict only if

- ① Atleast one write operation and
- ② Both operation (Read/write) should be on same data item, and
- ③ Both operation should not be from same transaction.

Conflict Equivalent Schedule:-

S' results after swapping consecutive non-conflict pairs of S . Then

S' and S are conflict equivalent #

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

Conflict pair (between $W(A)$ and $R(A)$)
Non-conflict pair (between $R(B)$ and $R(A)$)

Swapping non conflict pairs


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

S'


Q: 11
Pg 43

Which one conflict eq. to which???


S1:



S2:



S3:




swapping non conflict pair


swapping non conflict pair

By swapping non conflict pairs of S1, we get S2
so $S1 \equiv S2$ or $S1 = S2$


S2:



S3:



non conflict pair swapped



final updation (A) by T3




Final updation (A) by T2

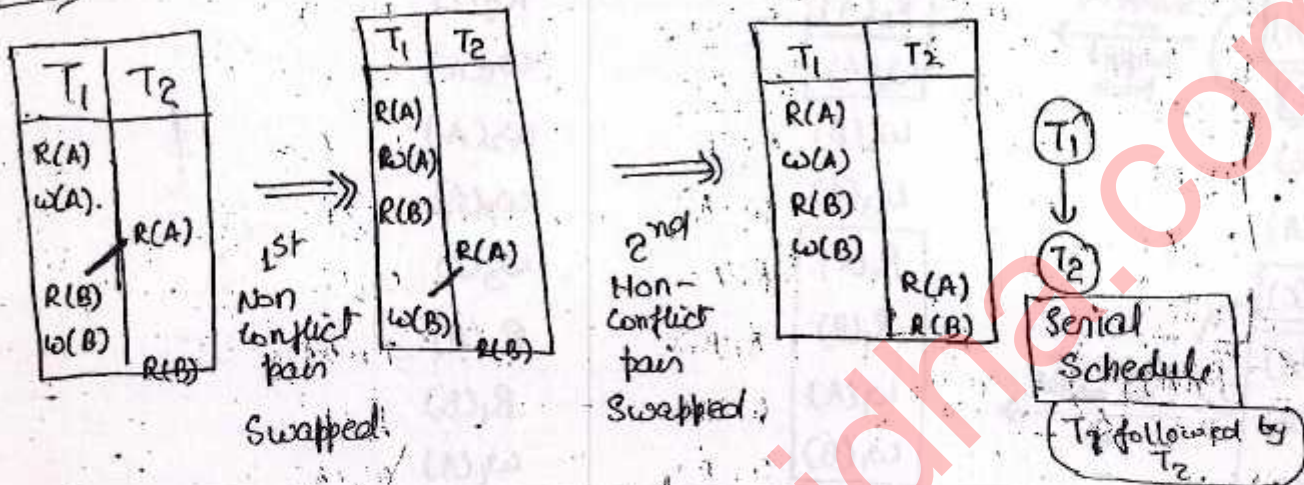
Conflict pairs swapped, hence

$S2 \neq S3$

Conflict Schedule Pair



Given schedule is Conflict serializable schedule
 if \exists Conflict equivalent serial schedule of given schedule.



Aise (non-conflict) pairs ko swap karte-2 agar hme koi serial schedule mil jae then we say that given schedule is Conflict-serializable.

To check Conflict Serializability

Precedence Graph

Graph: (V, E)

Venices: Transactions in the schedule.

Edges: $T_i \rightarrow T_j$ $i \neq j$


only if \exists conflict pair with T_i preceded then T_j.

RW	T _i : R(A)	T _j : W(A)
WR	T _i : W(A)	T _j : R(A)
WW	T _i : W(A)	T _j : W(A)

#

①

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



If the Precedence Graph contain **No Cycle** then it is Conflict Serializable

#


②

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



③

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



Not Conflict Serializable



NOTE :- TESTING Condition

⇒ If precedence graph ACYCLIC, then schedule is Conflict Serializable.

⇒ Equivalent serial schedule is Topological sequence of acyclic precedence graph.

⇒ If precedence graph cyclic, then schedule not **NOT Conflict Serializable**



③
 S_1
 $R_1(A)$
 $W_1(A)$



Acyclic graph

It is Conflict serializable

③



Not conflict serializable
 Schedule

Topological Sequence


(Traversing Acyclic Graph)

If no vertex with indegree = 0 there exists CYCLE in GRAPH

[1] Visit vertex (v) indegree "0" and delete vertex 'v' and outdegree edges of "v"





[2] Repeat (1), till graph becomes empty.

Ex




Start with indeg(0) = T1

Delete T1 & outdegree of T1



Schedules



Two schedule possible

① $T_1 \rightarrow T_2 \rightarrow T_3 \rightarrow T_4 \rightarrow T_5$

② $T_1 \rightarrow T_2 \rightarrow T_4 \rightarrow T_3 \rightarrow T_5$

Q:- Check whether conflict serializable or not ??

S

$R_4(A)$

$R_2(A)$



$R_3(A)$

$W_1(B)$

$W_2(B)$

$W_3(B)$


$W_2(B)$



So not conflict serializable.


Q:-

S





Q: How many conflict equivalent serial schedule possible ??

(# graph cyclic = 0 conflict schedule)



This schedule is equivalent to four conflict equivalent serial schedule.

Sequence = ?



So possible

Sequences are:

- ① T1 T2 T3 T4
- ② T2 T1 T3 T4
- ③ T2 T3 T1 T4

#	T ₁	T ₂
①	X(A) R(A)	
	W(A)	
X(B)	U(A)	S(A) R(A)
	U(B)	S(B) ← denied already locked. R(B)
	R(B) W(B)	

Not allowed by 2PL

Explanation


T ₁	T ₂
X(A) R(A) W(A)	
X(B) U(A)	S(A) R(A)
R(B) W(B) U(B)	S(B) R(B)
	U(B) U(A)

Allowed to execute by 2PL

Q Allowed by 2PL or not??

T ₁	T ₂	T ₃
X(A) R(A) W(A)		S(C)
		R(C) S(A) R(A)
W(A) W(B)	W(C)	

This schedule not allowed by 2PL.



T₃ (Indeg=0) → T₁ → T₂

T₃ → T₂ → T₁

Lock Upgrading Techniques:

Shared Lock can be upgraded to Exclusive Lock by requesting 'X' lock without unlock of shared lock.

T
S(A)
R(A)
⋮
U(A)
X(A)
W(A)


(E.g.)

T ₁	T ₂	T ₃
S(A) R(A)		S(C) R(C)
	X(C) W(C) U(C)	R(A) U(A)
X(A) W(A) X(B) W(B) U(B) U(A)		R(A)

2PL + updating Lock Technique

allowed

higher degree of conc. than 2PL



NOTE

S/X && 2PL

↳ always ensures CSS

↳ may not be free from irrecoverability.

T ₁	T ₂
R(A) W(A)	
	R(A)
Commit	Commit

T ₁	T ₂
X(A) R(A) W(A) U(A)	
	S(A) R(A) U(A)
Commit	Commit

Allowed by 2PL

↳ so definitely serializable

is IRRECOVERABLE

T₂ commit before T₁


T₁ updating data item hence to be recoverable, it should commit before T₂. But here T₂ is committing before T₁ hence Irrecoverable

Imp

Conflict serializable = ??

	T ₁	T ₂	T ₃
R(A)			
W(A)		W(A)	
			W(A)

Solⁿ



Precedence graph contain cycle. So
 ↓
 Schedule

Not Conflict SS

↓ equivalent Serial schedule

	T ₁	T ₂	T ₃
R(A)			
W(A)		W(A)	
			W(A)

⇒ For S₁, equal serial schedule

T₁ → T₂ → T₃

NOTE

For S ∃ S₁ which is serial schedule
 So S is SERIALIZABLE

⇒ In Both S and S₁

- R(A) done from database
 - Final updation by T₃ in both S and S₁
- No plbm
or, no plbm

NOTE:

If the Schedule is Conflict serializable schedule, we can say serializable schedule. But if schedule is not conflict serializable, may or may not be serializable.

if (Acyclic Precedence Graph) // Only sufficient condition
 {
 Conflict SS
 ↳ serializable
 }
 else // Not necessary condition
 {
 not conflict serializable schedule
 ↳ may/may not be serializable
 }

VIEW SERIALIZABLE SCHEDULE

(Both necessary & sufficient condition)

```
if (S satisfied VIEW SS condition)
{
    Schedule is Serializable
}
else
{
    Not serializable schedule
}
```

↳ View Serializable Schedule :- View Equivalent Serial Schedule.

View Equivalent :-


S1 and S2 said to be View Equivalent

Only if

[1] Initial reads of S1 and S2 should be same.



Example



$S_1 \neq S_2$

[2]

$W_i(A) \quad R_j(A)$

Write-Read sequence should also be equal.
(Updated Reads should be same)

T_i	T_j
$W(A)$	$R(A)$

T_i	T_j
$W(A)$	$R(A)$

Example

S_1

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

Updated Read
 $T_2 \rightarrow T_3$

S_2


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

Updated read
 $T_1 \rightarrow T_3$

$S_1 \neq S_2$

[3]

Final Updates for every data item should be same in S_1 and S_2 .



Example

S_1


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

Final updation
A \rightarrow T_3
B \rightarrow T_3

S_2

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

Final Updation



$S_1 \neq S_2$

⇒ When all above three conditions satisfied, S schedule is said to be VIEW SERIALIZABLE.

#

	T ₁	T ₂	T ₃
w(A)			
.		w(A)	
.			w(A)

(S₁)

T ₁	T ₂	T ₃
	w(A)	
w(A)		
		w(A)

(S₂)

- ⇒ S₁ and S₂ are NON-CONFLICT. Equivalent solutions.
- ⇒ S₁ and S₂ are VIEW SERIALIZABLE

T ₁	T ₂	T ₃
w(A)		
	w(A)	
		w(A)

↳ says that if final updation done by T₃ order of T₁ & T₂ doesnot matter.

↳ But if ∃ any read R(A) b/w three consecutive updations, then schedule order (T₁ → T₂) (T₂ → T₁) matters.

••• If ∃ no read b/w three consecutive w(A) from three transactions, first two updations can be in any order.

Q: Schedule serializable on matrix

- a) $T_1 \rightarrow T_2 \rightarrow T_3$
- b) $T_2 \rightarrow T_1 \rightarrow T_3$
- c) $T_2 \rightarrow T_3 \rightarrow T_1$
- d) Non-serializable (all the above)

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


	CSS	
a	X	X
b	X	✓
c	X	X
d	✓	✓

(All of above)
 ↳ view serializable ↳ possible if view serializable



\Rightarrow S Conflict. serializable as :- check CSS
 \Rightarrow S serializable as :- check for view SS

VIEW

① Final updation :-



↳ updation on A by all but finally by T_3
 \Rightarrow In its equivalent serial schedule T_3 must execute at last.



tilt here.


(2) Initial Read

Write Operations

A: (T_2) $T_1 T_2 T_3$
 B: (T_2) T_1

↑
 A and B both initially read by T_2 .

— means T_2 should execute before other two write operations.




✓ — ok till now

— To preserve Initial Read, T_2 should execute before other trans.

$T_2 \rightarrow T_1$

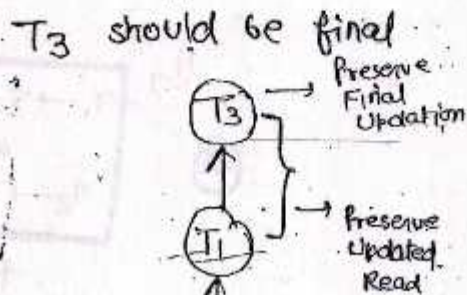
→ To preserve ~~write~~ ^{Read} followed by ~~Read~~ ^{write}:
 (WR)

↳ T_3 Reading Value Updated by T_1



⇒ From above three conditions, we have

- # {
 a) $(T_1, T_2) \rightarrow (T_3)$
 $(T_2 - T_1) \rightarrow T_3$
 b) $T_2 \rightarrow (T_1, T_3)$
 c) $T_1 \rightarrow T_3$




become

If some give and other give $T_1 \rightarrow T_3$ } Schedule is
 $T_3 \rightarrow T_1$ } Non-serializable Schedule

Q Pg 90
 Q41

(i) Final Updation

$D_1 : T_2$
 $D_2 : T_1$
 $D_3 : T_3$ } write



Both can't be true at the right time

Non-serializable

(ii) Initial Reads

$D_1 : T_1$
 $D_2 : T_2$
 $D_3 : T_2$

(iii) WR (updated Read)

T_1	T_2	T_3
	$w(D_2)$	
$R(D_3)$		




T_1	T_2	T_3

Pg 92
 Q.9

T_1	T_2
$R(A)$	$R(A)$ $w(A)$ $R(B)$
$w(A)$ $R(B)$ $w(B)$	$w(B)$

(i) Final Updation

Write oper^m




Both not possible simultaneously. Thus

Non-serializable

Strict R.S. may sometimes suffer from Read Write Problem

Based on R/W operation + Commit/Rollback



Serializability
 (No RW, WR, W/W) problem
 ∃ a Serial schedule for concurrent execution of Trans.
 only based on R/W operations only

#	T ₁	T ₂
	R(A)	
	W(A)	R(A)
		R(B)
	W(B)	
	C	C

Recoverable

#	T ₁	T ₂
	R(A)	
	W(A)	R(A)
		R(B)
	W(B)	
	C	C


Irrecoverable

Q:- Write a schedule which is Recoverable but not serializable

- ① Recoverable not serializable
- ② Cascadeless

- ④ CSS & Strict
- ⑦ CSS & CL but strict
- ⑧
- ⑨

7/11/11

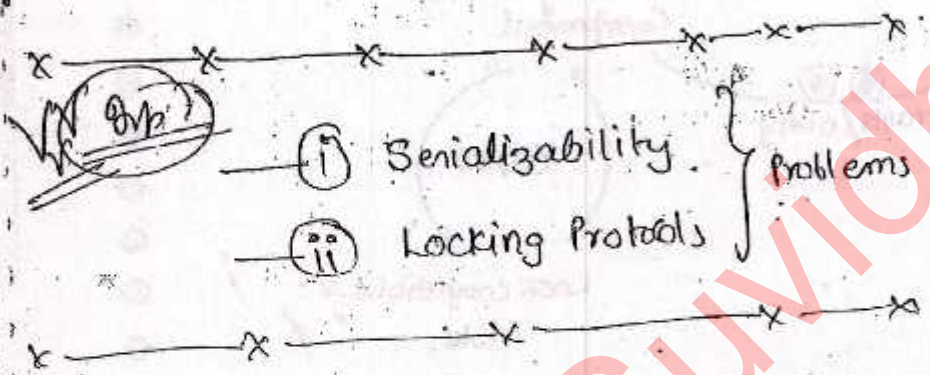


⇒ A schedule is free from all problems (RW, RR, WW, IR, recoverability)

if it is **STRICT RECOVERABLE and Serializable**

GOAL OF TRANSACTION CONTROL TECHNIQUES

- lock based protocols
- Timestamp " "



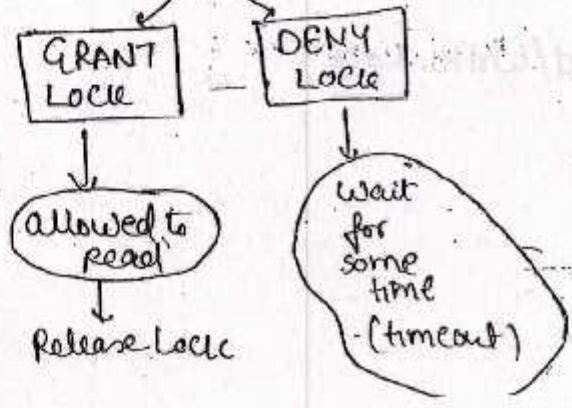
LOCK BASED PROTOCOLS

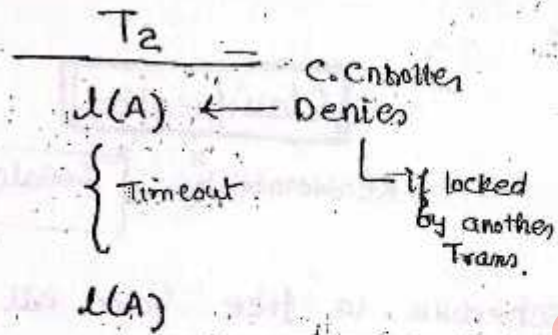
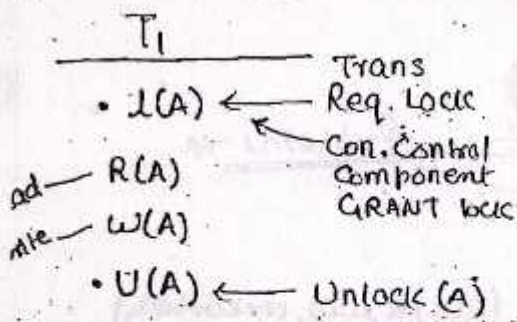
LOCK :- Variable used to identify status of the data item.

* Trans. should request lock before usage of any data item.

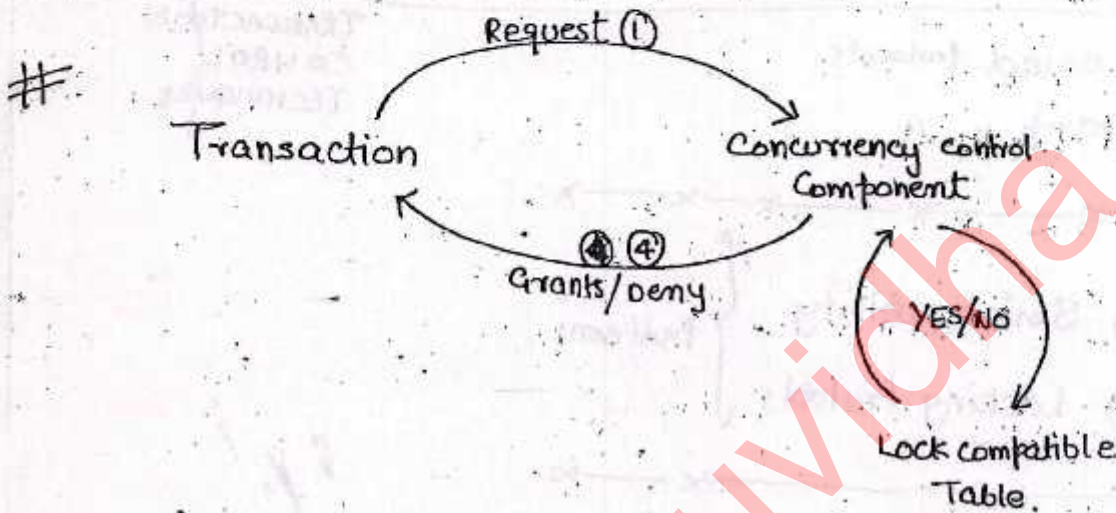
⇒ Trans. Request Lock

It is responsibility of CONCURRENCY CONTROL TECHNIQUE.





(Tran. Legal)



Shared / Exclusive Locking

T
S(A)
R(A)
W(A)

shared lock

T having shared lock on data item can read it but can't write the data item.

Shared Lock [S]: → can Read but not write

Exclusive Lock [X]: (Read/write Lock)

T
X(A)
R(A)
W(A)
⋮

⇒ Given schedule S

T ₁	T ₂
R(A)	
W(A)	
	R(A)
	R(B)
R(B)	
W(B)	

Implementing using Locks

T₁ → Read + write

Require Exclusive Lock

T₂ → Only Read shared Lock

Implementation

T ₁	T ₂
X(A)	
R(A)	
W(A)	
U(A)	
	S(A)
	R(A)
	U(A)
	S(B)
	R(B)
	U(B)
X(B)	
R(B)	
W(B)	
U(B)	

Legal Schedule

Non-serializable Schedule
(WR problem exist)

Strict Recoverable + Serializable not achieved

Lock/Compatible Table

To overcome this incapability

TWO PHASE LOCKING PROTOCOL

Trans. T not allowed

to REQUEST lock on data item if T has already performed some unlock operation.

T
l(A)
l(B)
l(C)
l(D)
⋮
u(C)
l(E)
u(A)
u(B)
u(D)

Locking PHASE (GROWING PHASE)

Lock Point

once unlock Not allowed to request anymore locks.

UNLOCKING PHASE (SHRINKING PHASE)

Lock compatible Table

	S	X	→ Hold ⁱ
S	Yes	No	
X	No	No	

Request_j

Lock point :- Position of last lock operation or position of first unlock operation.

- Transition phase from lock to unlock phase


» 2PL always allowed ~~to visit~~ to execute Serializable schedule only.
(conflict serializable schedule)

» If schedule (S) allowed to execute by 2PL then S surely conflict serializable.



» Equivalent serial schedule of 2PL allowed schedule is based on LOCK POINTS.


» If assume T_1, T_2, T_3



If it is 2PL schedule

Equivalent serial schedule

$T_2 : T_1 : T_3$




? How to test S satisfied by 2PL.

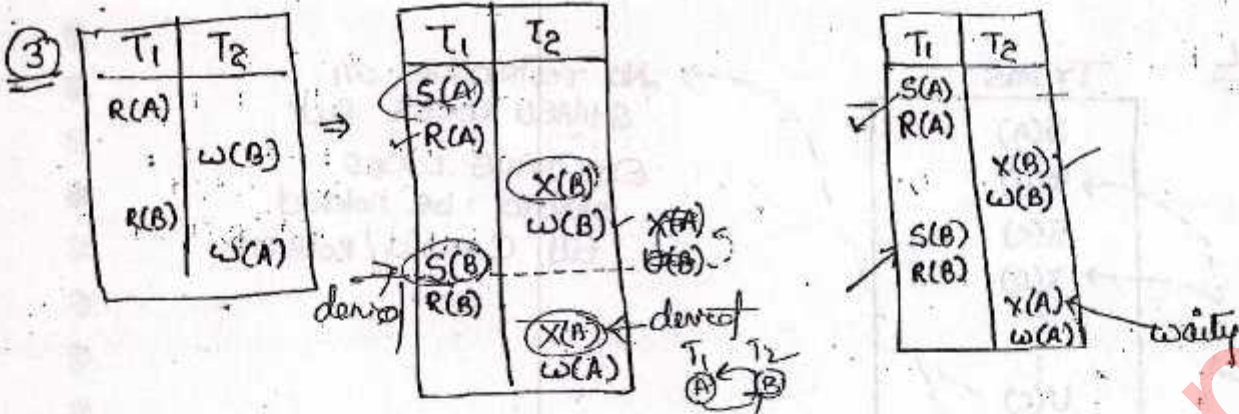
└ check whether conflict serializable

YES
Conflict serializable

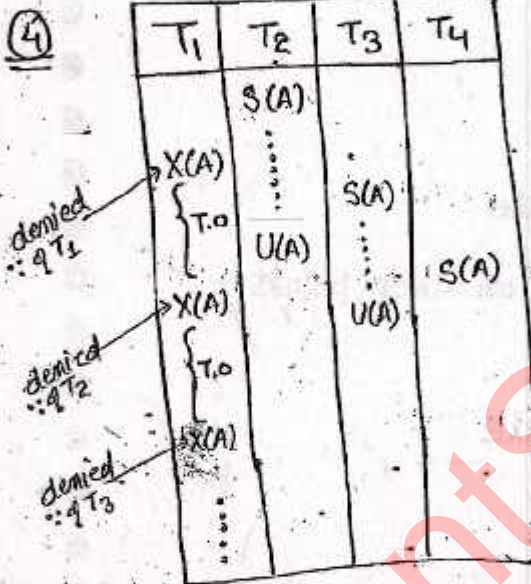
NOT
Conflict serializable

may/may not be allowed by 2PL





Note: 2PL may not be free from **DEADLOCK**

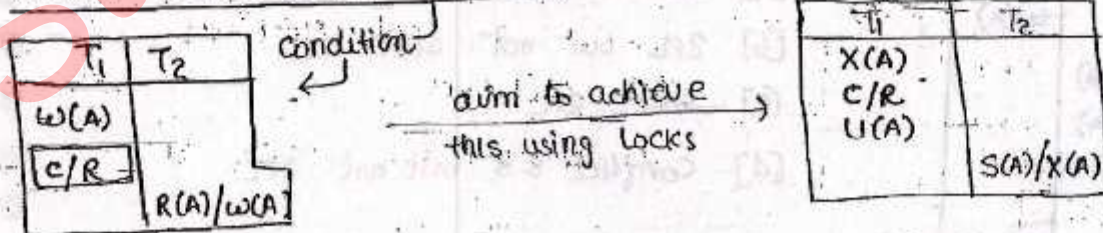


⇒ 2PL may suffer from **STARVATION**

STRICT 2PL PROTOCOL

↳ we are trying for strict recoverable

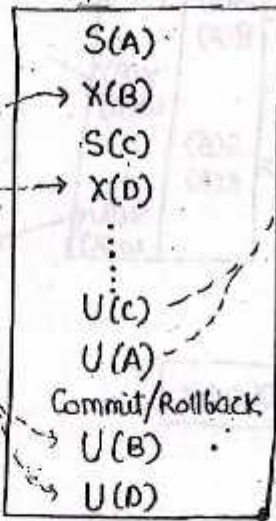
STRICT Recoverable Schedule



⇒ All Exclusive Locks should be hold until commit / rollback.

Strict 2PL

Trans



Held
Inlocked
after
Commit
or
Rollback

No restrictions on
SHARED LOCKS But
EXCLUSIVE LOCKS
should be held
till Commit/Rollback.

STRICT 2PL PROTOCOL

- Ensures conflict serializable schedule
- Equivalent serial schedule, based on lock points and
- Ensures Strict Recoverable Schedule
- May not be free from
 - ⊙ Deadlock &L
 - ⊙ STARVATION

Q:-

T ₁	T ₂	T ₃
R(A)		W(A)
	R(B)	
	R(A)	
	C	
C		
		Commit

→ Given trans. is

- [a] not in 2PL
- [b] 2PL but not Strict 2PL
- [c] Strict 2PL
- [d] Conflict ss but not 2PL

Sol^m (2PL) ?? =>

T ₁	T ₂	T ₃
S(A)		
R(A)		
U(A)		
	S(B)	X(A)
	R(B)	W(A)
	S(A)	U(A)
	R(A)	
	C	
C		C

2PL ✓

STRICT 2PL

T ₁	T ₂	T ₃
S(A)		
R(A)		
U(A)		
	S(B)	
	R(B)	
	S(A)	
	R(A)	
	C	
C		
		U(A)
		C

(STRICT 2PL)

exclusive lock delayed till Commit / Rollback

conflict

So not strict 2PL

③ REGORIOUS 2PL

2PL & every lock (shared/exclusive) should be hold untill commit


» Ensures serializability (conflict)

» Equivalent serial schedule based on the order of commit

Trans
S(A)
X(B)
X(C)
S(B)
⋮
C/R
U(B)
U(A)
U(D)
U(C)

T ₁	T ₂	T ₃
	C	
C		C

Serial schedule = T₂ → T₁ → T₃




Time STAMP BASED PROTOCOLS

Time Stamp Value :- Unique value assigned by DBMS to every transaction in ascending order.

If Time stamp : 10 30 20 40 60 50
 Trans : T_1 T_2 T_3 T_4 T_5 T_6

Then Schedule



Time stamp value can be used to

- ① Transaction ID
- ② Set the priority

Time stamp value of data item :-

Read_TS(A) :- Highest Trans. time stamp value that has performed R(A) operation successfully.

10	30	20	40
T_1	T_2	T_3	T_4
R(A)	R(A)		

$RTS(A) = 10$ ——— Initially

$RTS(A) = 30$ ——— [max(10,30)]
 ——— finally

Q. Write $TS(A)$:- Highest Transaction Time stamp value that has performed $w(A)$ operation successfully.

10	30	20
T_1	T_2	T_3
R(A) W(A)	R(A)	W(A)

$RTS(A) = \cancel{10}$ (30) \rightarrow Youngest T.S for Read

$WTS(A) = \cancel{40}$ (40) \rightarrow Youngest T.S for Write

BASIC GOALS OF TIMESTAMP ORDERING PROTOCOL

(i)

10	20	30
T_1	T_2	T_3

 \Rightarrow Concurrency allowed only if equivalent serial schedule should be based on T.S ordering

even though

$(T_3 \rightarrow T_1 \rightarrow T_2)$
 $(T_2 \rightarrow T_1 \rightarrow T_3)$

\rightarrow T.S.O. Protocol not allowed

\rightarrow Concurrent ex = $T_1 \rightarrow T_2 \rightarrow T_3$ (serial sch)
Then allowed.
If not equal then Rollback.
It should strictly be equal to $T_1 \rightarrow T_2 \rightarrow T_3$

Q. IS S allowed to execute using T.S. ordering protocol?

S: _____
 $T_s: (T_1, T_2, T_3) : (20, 10, 30)$

Sol^m If \exists serial order $(T_2 \rightarrow T_1 \rightarrow T_3)$ for S, then S is allowed to execute using T.S ordering protocol.

\rightarrow (i) check for CSS

\rightarrow (ii) If correspond, serial order = $T_2 \rightarrow T_1 \rightarrow T_3$ - ok.


Q. S: $R_1(A), R_2(B), W_1(A), W_1(B)$

$T_s: \{10, 20\}$

Allowed using T.S. o ??

\Rightarrow


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



\rightarrow Topological Seq = $T_1 \rightarrow T_2$

But we expect

$(T_1 \rightarrow T_2)$
based on T.S value
So this schedule not allowed to execute using T.S.O. Protocol



This conc. exe should be = serial schedule


$$T_1 \rightarrow T_3 \rightarrow T_2$$

Allowed conflict pairs


- T₁ → T₂
- T₃ → T₂
- T₁ → T₃

If any conflict pair other than this then T.S.O protocol can't be applied.


not a allowed conflict pair Hence w₁(A) causing problem so T₁ - Rollback



Possibility of Rollbacks:



T₁ → T₂ ok (First open in not a problem)



T₂ → T₁ So T₁ Rollbacked ∵ it is violation WTS=20

NOTE :-

If Younger Trans. updates data item. Then older trans. not allowed to read.

T₁: Issues R(A):-
 $WTS(A) > TS(T_1) \Rightarrow$ Rollback T₁

②

10	20
T ₁	T ₂
	R(A)

WTS = 0
 TS(T₁) = 20
 ↳ Issue Read
 0 > 20 false so not rollbacked.

②

10	20
T ₁	T ₂
↑ W(A)	R(A)

⇒ we are getting T₂ → T₁ instead of expected T₁ → T₂

⇒ If YOUNGER TRANS. performed the R(A) then older Trans. not allowed to update A.
 T₁: Issue W(A):
 RTS(A) > TS(T₁)
 Rollback T₁

③

10	20
T ₁	T ₂
↑ W(A)	W(A)

T₁ → T₂ — expected
 T₂ → T₁ — actually got.
 Rollback T₁

⇒ If Younger trans. performed write operation (update on data item) then older Trans. Not allowed to update A.
 T₁: Issue W(A):
 WTS(A) > TS(T₁)
 Rollback T₁

Basic Timestamp Protocol :-

[1] Transaction T issues R(A) : operation

- a) if $WTS(A) > TS(T)$ Then Rollback T
- b) Otherwise allowed to execute R(A) by T
 Set $Read_TS(A) = \max \{ RTS(A), TS(T) \}$

② Trans. T issues W(A) operation :-

a) if $RTS(A) > TS(T)$ Then Rollback T

b) if $WTS(A) > TS(T)$ Then Rollback T

c) Otherwise execute W(A) operation by T.
Set $WTS(A) = TS(T)$

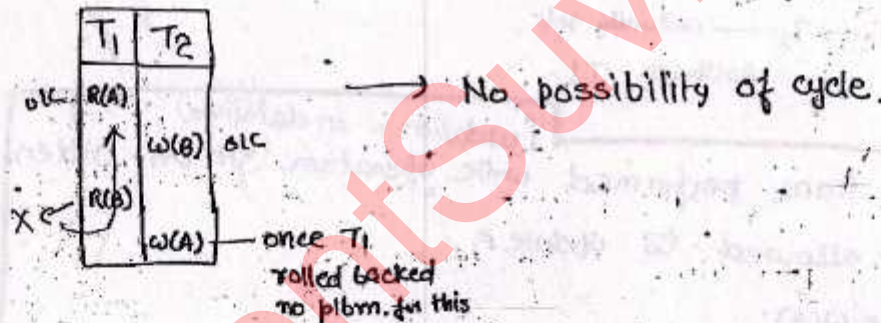
B.T.O. Protocol

① Ensures Serializability

Equivalent serial schedule should be always based on TS ordering.

② Timestamp Ordering protocol free from deadlock.

Ex



③ Not free from Irrecoverability (Irrecoverable schedule)

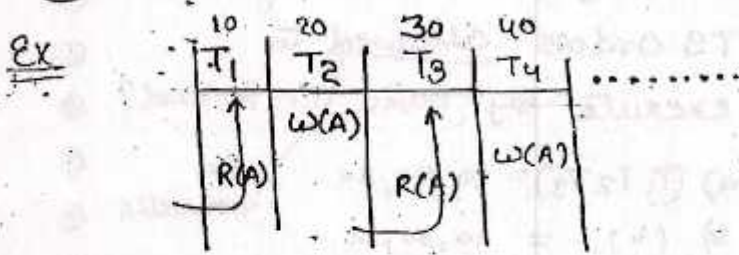
Ex

	T ₁	T ₂
1o	W(A)	
2o		R(A)
	C	C

→ S allowed to execute using B.T.O protocol
Thus it is a conflict serializable schedule.

But it is Irrecoverable ∵ T₁ performing updation getting committed after T₂.

④ may not be free from **STARVATION**



8th Nov 2011

BASIC Time Stamp Protocol :-

- ① ensures serializability
- ② Deadlock free
- ③ Starvation may be possible
- ⇒ ④ Irrecoverable may be possible

STRICT TO Protocol

- ① ensure serializability
- ② Deadlock free
- ③ starvation may be possible
- ④ Always result strict recoverable schedule.

STRICT Time Stamp Ordering Protocol :

t_0	2_0
T_1	T_2
$w(A)$	
c/R	$R(A)/w(A)$

if $Trans(T_2) \text{ } R(A)/w(A)$
 $WTS(A) < TS(T_2)$

$R(A)/w(A)$ of T_2
 should be delayed until $c/R(T_1)$

Strict Recoverability

⇒ S TSO Protocol :

• Follow Basic TO Protocol

& R

• Trans T_2 issues $R(A)/w(A)$ operations :-

if $WTS(A) < TS(T_2)$ then $R(A)/w(A)$ of T_2 should be delayed until $Trans(T_1)$ that has performed $w(A)$ commit/rollbacks.

S:

T ₁	T ₂	T ₃
R(A)		
	R(B)	
w(C)		R(B)
		R(C)
	w(B)	
		w(A)

Which of the following TS orders allowed to execute by Basic T.O Protocol?


- a) (T₁ T₂ T₃) = 10, 20, 30
 - b) (") = 10, 30, 20
 - c) (") = 30, 10, 20
 - d) (") = 30, 20, 10
- no rollbacks

Equivalent Serial schedule acc. to options.

- a) T₁ → T₂ → T₃
- b) T₁ → T₃ → T₂
- c) T₃ → T₁ → T₂
- d) T₃ → T₂ → T₁

- allowed conflict pairs
- ① T₁ → T₂ ② T₂ → T₃ ③ T₁ → T₃
 - ② T₁ → T₃ , T₃ → T₂ T₁ → T₂
 - ③ T₃ → T₁ T₁ → T₂ T₃ → T₂
 - ④ T₃ → T₂ T₂ → T₁ T₃ → T₁

Now



Conflict equivalent to = T₁ → T₃ → T₂
equivalent serial schedule.

Q: Schedule Given
Given Time stamp
Which trans. Rollbacks???

Q) Sequence

10	20	30
T ₁	T ₂	T ₃
R(A)		
	R(B)	
w(C)		R(B)
		R(C)
	w(B)	
		w(A)

allowable Conflict Pair

- T₁ → T₂
- T₂ → T₃
- T₁ → T₃

any trans. violating it is Rollbacked

no conflict

1st conflict pair T₁ → T₃ allowed


2nd conflict pair T₃ → T₂ X

T₁ → T₃ → allowed


no mean

kill


Once a trans. is Rollbacked (Killed) } here T_2 Rollbacked
 then there is no need to consider its operations anymore.
 (In T_2 afterwards $w(A)$ comes, we can neglect that
 its mean nothing $\because T_2$ has already being rollbacked)



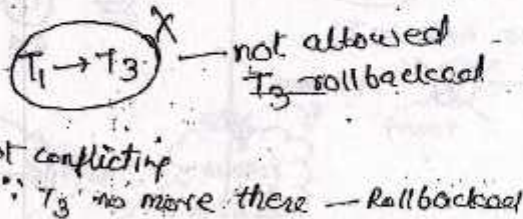
T_1	T_2	T_3
R(A)		
	R(B)	
w(C)		
		R(B)
		R(C)
	w(B)	
		w(A)



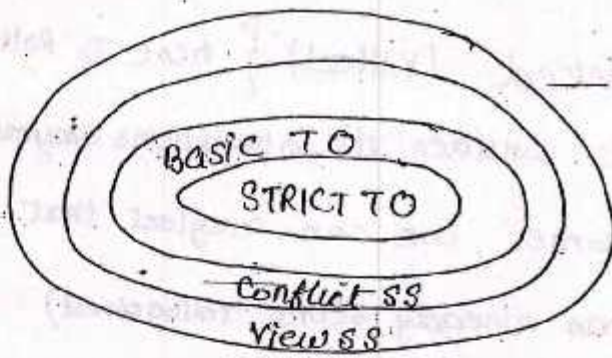
* no need to check this as T_3 already killed.
 not conflict
 \because possibility was conflicting with T_3 but T_3 has already being Rollbacked.




T_1	T_2	T_3
R(A)		
	R(B)	
w(C)		
		R(B)
		R(C)
	w(B)	
		w(A)



Rollbacked
 T_3 (circled)



90% GATE
90% of question
asked in Trans
asked from
these two topics



② Locking Protocols

S/X, 2PL, Strict 2PL

③ RW/WR/WO Problems & Recoverability

④ Time stamp ordering

check
if Recov
Cascadec
Strict




main
hears
Tues:

Deadlock


Deadlock Prevention algo: use time stamp value to avoid deadlocks in locking protocols.

Wait-Die

- Trans in Ascending order of TS value $TS(T_i) < TS(T_j)$
- ① if T_i required resource held by T_j





- ② if T_j require res. held by T_i




Wound-Wait

- (Opp. to wait-die)
- ① $TS(T_i) < TS(T_j)$
- ② ③ if T_i depends on T_j

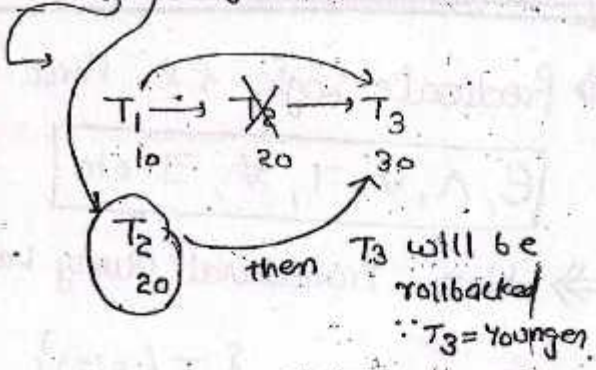


Rollback
& if
Res. is freed



Both are **deadlock** and **starvation** free protocols.

taken care by
Restarting with same timestamp



StudentSuvidha.com