

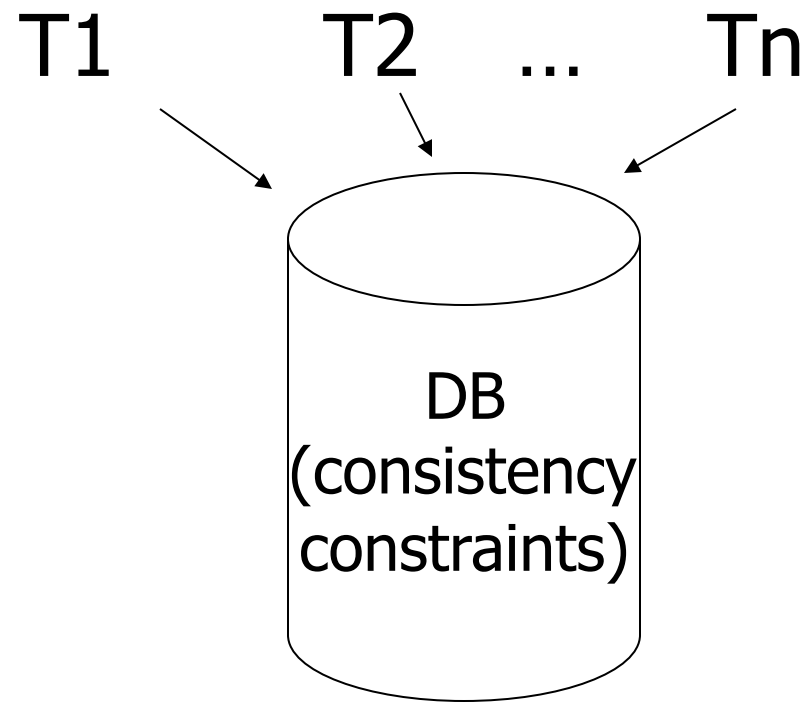
Concurrency Control

Parke Godfrey

Thanks to

- These slides are authored by Hector Garcia Molina (Stanford), 2002.
- They follow the class textbook (“Stanford”).

Chapter 18 [18] Concurrency Control



Example:

T1: Read(A)
A \leftarrow A+100
Write(A)
Read(B)
B \leftarrow B+100
Write(B)

T2: Read(A)
A \leftarrow A \times 2
Write(A)
Read(B)
B \leftarrow B \times 2
Write(B)

Constraint: A=B

Schedule A

T1

Read(A); $A \leftarrow A+100$

Write(A);

Read(B); $B \leftarrow B+100$;

Write(B);

T2

Read(A); $A \leftarrow A \times 2$;

Write(A);

Read(B); $B \leftarrow B \times 2$;

Write(B);

Schedule A

T1	T2	A	B
		25	25
Read(A); A ← A+100			
Write(A);		125	
Read(B); B ← B+100;			
Write(B);			125
	Read(A); A ← A×2;		
	Write(A);	250	
	Read(B); B ← B×2;		
	Write(B);		250
		250	250

Schedule B

T1

Read(A); $A \leftarrow A+100$
Write(A);
Read(B); $B \leftarrow B+100$;
Write(B);

T2

Read(A); $A \leftarrow A \times 2$;
Write(A);
Read(B); $B \leftarrow B \times 2$;
Write(B);

Schedule B

T1	T2	A	B
		25	25
	Read(A); $A \leftarrow A \times 2$;		
	Write(A);	50	
	Read(B); $B \leftarrow B \times 2$;		
	Write(B);		50
Read(A); $A \leftarrow A + 100$			
Write(A);		150	
Read(B); $B \leftarrow B + 100$;			
Write(B);			150
		150	150

Schedule C

T1

Read(A); $A \leftarrow A+100$
Write(A);

Read(B); $B \leftarrow B+100$;
Write(B);

T2

Read(A); $A \leftarrow A \times 2$;
Write(A);

Read(B); $B \leftarrow B \times 2$;
Write(B);

Schedule C

T1	T2	A	B
		25	25
Read(A); $A \leftarrow A+100$			
Write(A);		125	
	Read(A); $A \leftarrow A \times 2$;		
	Write(A);	250	
Read(B); $B \leftarrow B+100$;			
Write(B);			125
	Read(B); $B \leftarrow B \times 2$;		
	Write(B);		250
		250	250

Schedule D

T1

Read(A); $A \leftarrow A+100$
Write(A);

Read(B); $B \leftarrow B+100$;
Write(B);

T2

Read(A); $A \leftarrow A \times 2$;
Write(A);
Read(B); $B \leftarrow B \times 2$;
Write(B);

Schedule D

T1	T2	A	B
		25	25
Read(A); $A \leftarrow A+100$			
Write(A);		125	
	Read(A); $A \leftarrow A \times 2$;		
	Write(A);	250	
	Read(B); $B \leftarrow B \times 2$;		
	Write(B);		50
Read(B); $B \leftarrow B+100$;			
Write(B);			150
		250	150

Schedule E

Same as Schedule D
but with new T2'

T1

Read(A); $A \leftarrow A+100$
Write(A);

Read(B); $B \leftarrow B+100$;
Write(B);

T2'

Read(A); $A \leftarrow A \times 1$;
Write(A);
Read(B); $B \leftarrow B \times 1$;
Write(B);

Schedule E

Same as Schedule D
but with new T2'

T1	T2'
Read(A); A ← A+100 Write(A);	Read(A); A ← A×1; Write(A);
Read(B); B ← B+100; Write(B);	Read(B); B ← B×1; Write(B);

A	B
25	25
125	
125	
	25
	125
125	125


- Want schedules that are “good”, regardless of
 - initial state and
 - transaction semantics
- Only look at order of read and writes

Example:

$Sc = r_1(A)w_1(A)r_2(A)w_2(A)r_1(B)w_1(B)r_2(B)w_2(B)$

Example:

$$Sc = r_1(A)w_1(A)r_2(A)w_2(A)r_1(B)w_1(B)r_2(B)w_2(B)$$



$$Sc' = r_1(A)w_1(A) r_1(B)w_1(B)r_2(A)w_2(A)r_2(B)w_2(B)$$



T_1

T_2

The Transaction Game

A								
B								
T1								
T2								

The Transaction Game

A	r	w	r	w				
B					r	w	r	w
T1	r	w			r	w		
T2			r	w			r	w

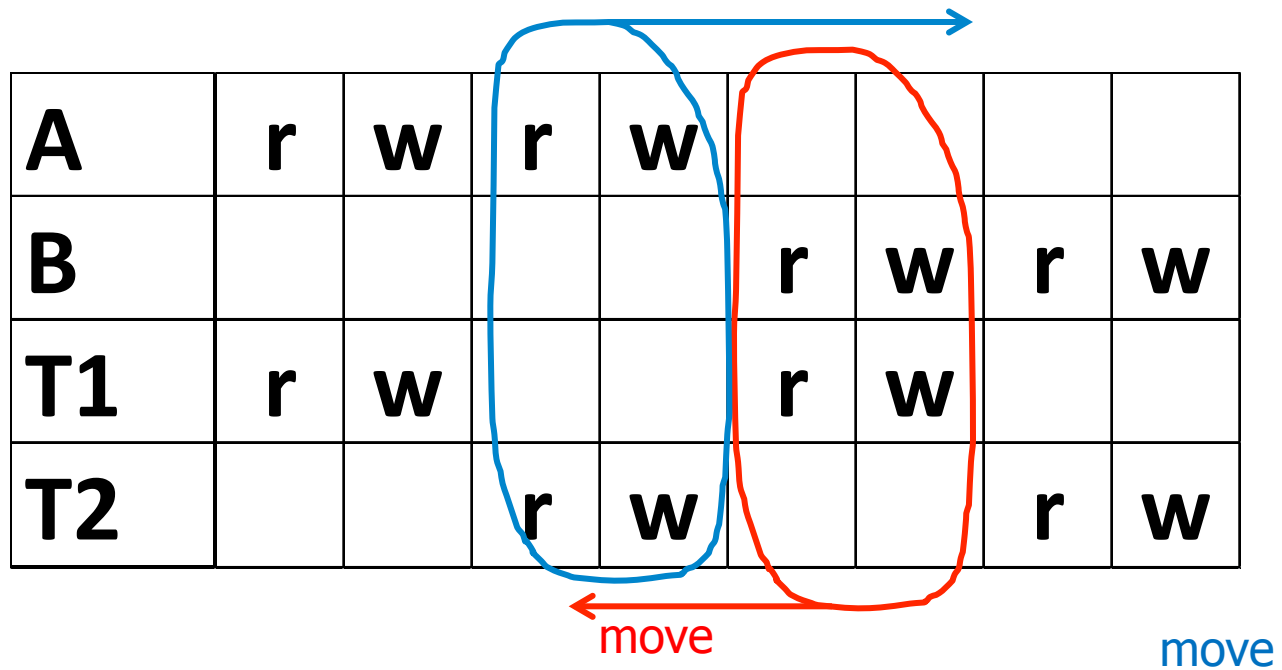
The Transaction Game

A	r	w	r	w				
B					r	w	r	w
T1	r	w			r	w		
T2			r	w			r	w

until column

hits something

can move column



A	r	w			r	w		
B			r	w			r	w
T1	r	w	r	w				
T2					r	w	r	w

Schedule D


A	r	w	r	w				
B					r	w	r	w
T1	r	w					r	w
T2			r	w	r	w		

However, for S_d :

$$S_d = \underbrace{r_1(A)w_1(A)r_2(A)w_2(A)}_{\text{Process 1}} \underbrace{r_2(B)w_2(B)r_1(B)w_1(B)}_{\text{Process 2}}$$

- as a matter of fact,
 T_2 must precede T_1
in any equivalent schedule,
i.e., $T_2 \rightarrow T_1$

- $T_2 \rightarrow T_1$
- Also, $T_1 \rightarrow T_2$



S_d cannot be rearranged
into a serial schedule




S_d is not “equivalent” to
any serial schedule




S_d is “bad”

Returning to Sc



Returning to Sc



• no cycles \Rightarrow Sc is “equivalent” to a serial schedule
(in this case T_1, T_2)

Concepts

Transaction: sequence of $r_i(x)$, $w_i(x)$ actions

Conflicting actions:

$r_1(A)$ $w_2(A)$ $w_1(A)$
 $w_2(A)$ $r_1(A)$ $w_2(A)$


Schedule: represents chronological order
in which actions are executed

Serial schedule: no interleaving of actions
or transactions


Is it OK to model reads & writes as occurring at a single point in time in a schedule?

- $S = \dots r_1(x) \dots w_2(b) \dots$

What about conflicting, concurrent actions on same object?



What about conflicting, concurrent actions on same object?



- Assume equivalent to either $r_1(A) w_2(A)$
or $w_2(A) r_1(A)$
- \Rightarrow low level synchronization mechanism
- Assumption called "atomic actions"

Definition

S_1, S_2 are conflict equivalent schedules if S_1 can be transformed into S_2 by a series of swaps on non-conflicting actions.

Definition

A schedule is conflict serializable if it is conflict equivalent to some serial schedule.

Precedence graph $P(S)$ (S is schedule)

Nodes: transactions in S

Arcs: $T_i \rightarrow T_j$ whenever

- $p_i(A), q_j(A)$ are actions in S
- $p_i(A) <_S q_j(A)$
- at least one of p_i, q_j is a write

Another Exercise:

- What is $P(S)$ for
 $S = w_1(A) r_2(A) r_3(A) w_4(A) ?$

Lemma

S_1, S_2 conflict equivalent $\Rightarrow P(S_1)=P(S_2)$

Lemma

S_1, S_2 conflict equivalent $\Rightarrow P(S_1)=P(S_2)$

Proof:

Assume $P(S_1) \neq P(S_2)$

$\Rightarrow \exists T_i: T_i \rightarrow T_j$ in S_1 and not in S_2

$\Rightarrow S_1 = \dots p_i(A) \dots q_j(A) \dots$
 $S_2 = \dots q_j(A) \dots p_i(A) \dots$ $\left\{ \begin{array}{l} p_i, q_j \\ \text{conflict} \end{array} \right.$

$\Rightarrow S_1, S_2$ not conflict equivalent

Note: $P(S_1)=P(S_2) \not\Rightarrow S_1, S_2$ conflict equivalent

Note: $P(S_1)=P(S_2) \not\Rightarrow S_1, S_2$ conflict equivalent

Counter example:

$S_1 = w_1(A) \ r_2(A) \quad w_2(B) \ r_1(B)$

$S_2 = r_2(A) \ w_1(A) \quad r_1(B) \ w_2(B)$

Theorem

$P(S_1)$ acyclic $\iff S_1$ conflict serializable

Theorem

$P(S_1)$ acyclic $\iff S_1$ conflict serializable

(\Leftarrow) Assume S_1 is conflict serializable


$\Rightarrow \exists S_s: S_s, S_1$ conflict equivalent

$\Rightarrow P(S_s) = P(S_1)$

$\Rightarrow P(S_1)$ acyclic since $P(S_s)$ is acyclic

Theorem

$P(S_1)$ acyclic $\iff S_1$ conflict serializable



Theorem

$P(S_1)$ acyclic $\iff S_1$ conflict serializable


(\implies) Assume $P(S_1)$ is acyclic

Transform S_1 as follows:

(1) Take T_1 to be transaction with no incident arcs

(2) Move all T_1 actions to the front

$S_1 = \dots q_j(A) \dots p_1(A) \dots$



(3) we now have $S_1 = \langle T_1 \text{ actions} \rangle \langle \dots \text{rest} \dots \rangle$


(4) repeat above steps to serialize rest!

How to enforce serializable schedules?

Option 1: run system, recording $P(S)$;
at end of day, check for $P(S)$
cycles and declare if execution
was good

How to enforce serializable schedules?

Option 2: prevent P(S) cycles from occurring




A locking protocol

Two new actions:

lock (exclusive): $li(A)$

unlock: $ui(A)$



Rule #1: Well-formed transactions

$T_i: \dots l_i(A) \dots p_i(A) \dots u_i(A) \dots$

Rule #2 Legal scheduler

$S = \dots li(A) \dots ui(A) \dots$

\longleftrightarrow
no $lj(A)$

Exercise:

- What schedules are legal?

What transactions are well-formed?

$S1 = l_1(A)l_1(B)r_1(A)w_1(B)l_2(B)u_1(A)u_1(B)$
 $r_2(B)w_2(B)u_2(B)l_3(B)r_3(B)u_3(B)$

$S2 = l_1(A)r_1(A)w_1(B)u_1(A)u_1(B)$
 $l_2(B)r_2(B)w_2(B)l_3(B)r_3(B)u_3(B)$

$S3 = l_1(A)r_1(A)u_1(A)l_1(B)w_1(B)u_1(B)$
 $l_2(B)r_2(B)w_2(B)u_2(B)l_3(B)r_3(B)u_3(B)$

Exercise:

- What schedules are legal?

What transactions are well-formed?

S1 = $l_1(A)l_1(B)r_1(A)w_1(B)l_2(B)u_1(A)u_1(B)$
 $r_2(B)w_2(B)u_2(B)l_3(B)r_3(B)u_3(B)$

S2 = $l_1(A)r_1(A)w_1(B)u_1(A)u_1(B)$
 $l_2(B)r_2(B)w_2(B)l_3(B)r_3(B)u_3(B)u_2(B)?$

S3 = $l_1(A)r_1(A)u_1(A)l_1(B)w_1(B)u_1(B)$
 $l_2(B)r_2(B)w_2(B)u_2(B)l_3(B)r_3(B)u_3(B)$

Schedule F

T1

$l_1(A); \text{Read}(A)$

$A \leftarrow A + 100; \text{Write}(A); u_1(A)$

$l_1(B); \text{Read}(B)$

$B \leftarrow B + 100; \text{Write}(B); u_1(B)$

T2

$l_2(A); \text{Read}(A)$

$A \leftarrow Ax2; \text{Write}(A); u_2(A)$

$l_2(B); \text{Read}(B)$


$B \leftarrow Bx2; \text{Write}(B); u_2(B)$

Schedule F


		A	B
T1	T2	25	25
$l_1(A); \text{Read}(A)$			
$A \leftarrow A + 100; \text{Write}(A); u_1(A)$		125	
	$l_2(A); \text{Read}(A)$		
	$A \leftarrow A \times 2; \text{Write}(A); u_2(A)$	250	
	$l_2(B); \text{Read}(B)$		
	$B \leftarrow B \times 2; \text{Write}(B); u_2(B)$		50
$l_1(B); \text{Read}(B)$			
$B \leftarrow B + 100; \text{Write}(B); u_1(B)$			150
		250	150

Rule #3 Two phase locking (2PL) for transactions


$T_i = \dots li(A) \dots ui(A) \dots$



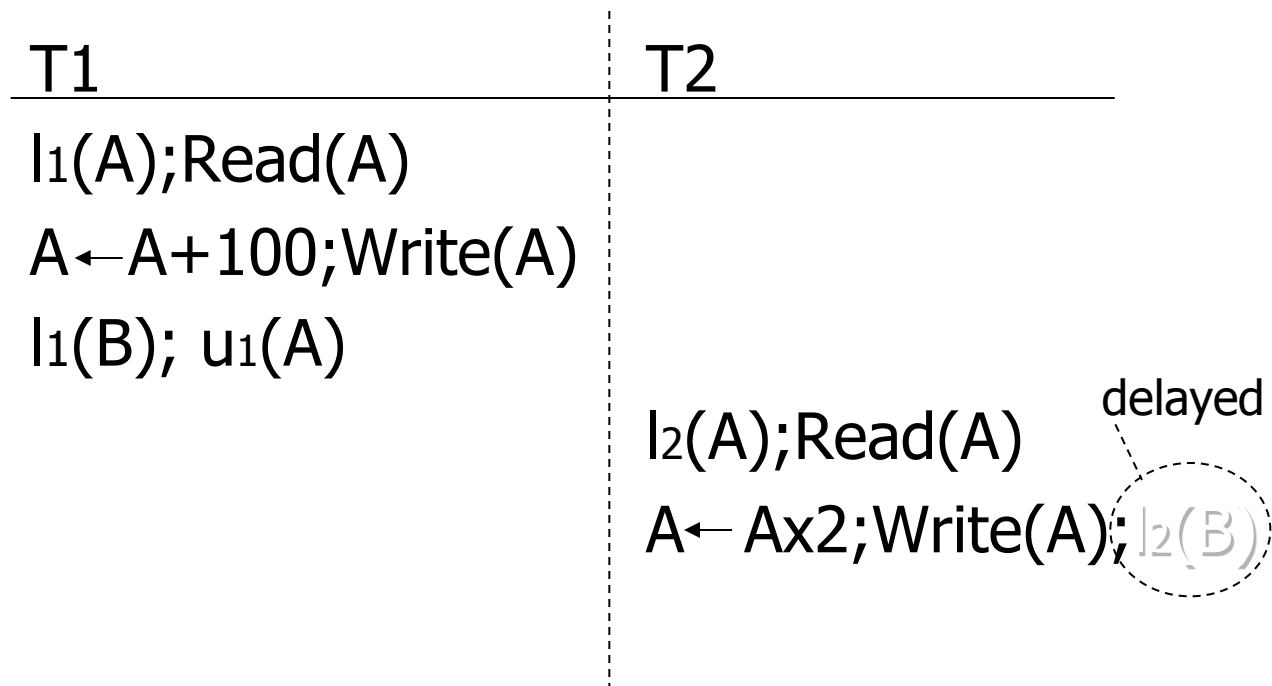
no unlocks



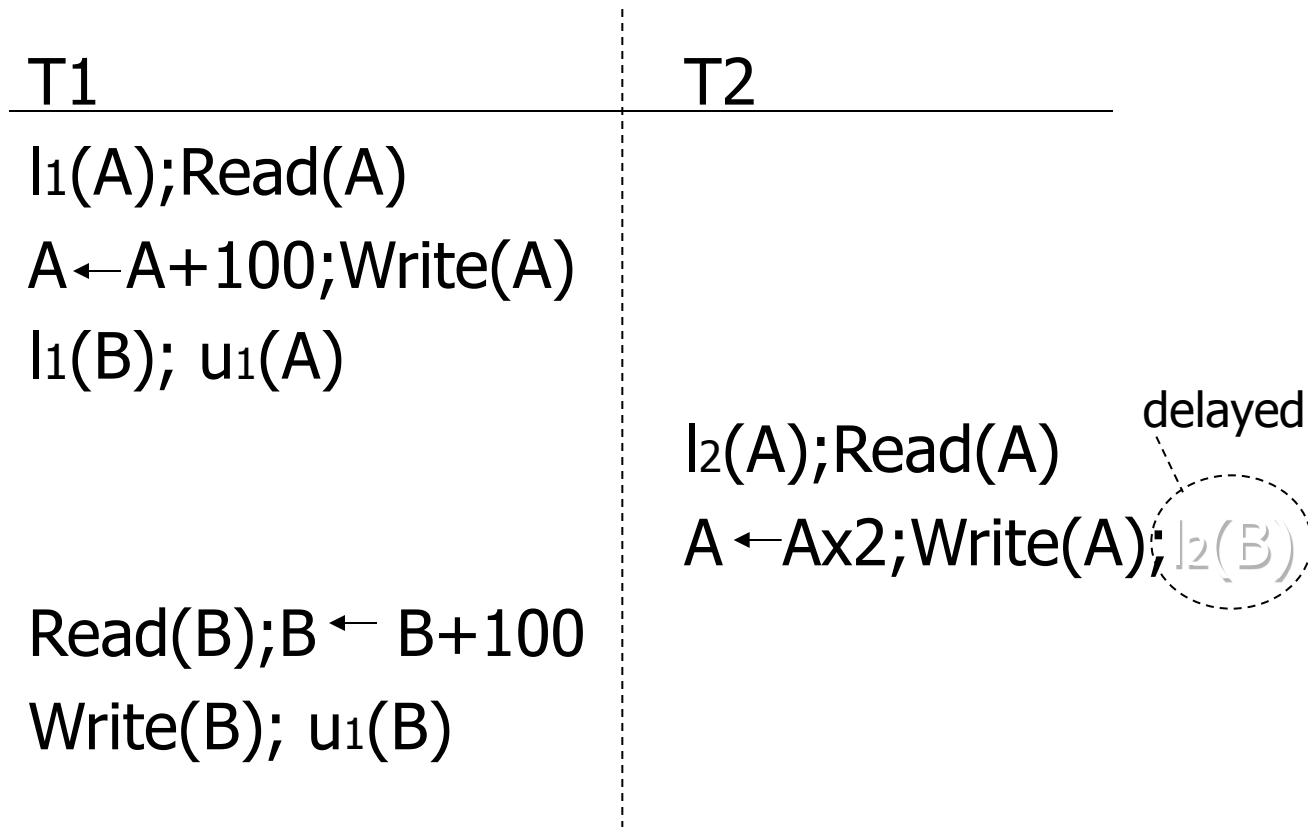
no locks



Schedule G



Schedule G



Schedule G

T1

$l_1(A); \text{Read}(A)$

$A \leftarrow A + 100; \text{Write}(A)$

$l_1(B); u_1(A)$

$\text{Read}(B); B \leftarrow B + 100$

$\text{Write}(B); u_1(B)$

T2

$l_2(A); \text{Read}(A)$


$A \leftarrow Ax2; \text{Write}(A); l_2(B)$

delayed


$l_2(B); u_2(A); \text{Read}(B)$

$B \leftarrow Bx2; \text{Write}(B); u_2(B);$

Schedule H (T2 reversed)



- Assume deadlocked transactions are rolled back
 - They have no effect
 - They do not appear in schedule

E.g., Schedule H = 
This space intentionally left blank!

Next step:

Show that rules #1,2,3 \Rightarrow conflict-
serializable
schedules

Conflict rules for $l_i(A), u_i(A)$:

- $l_i(A), l_j(A)$ conflict
- $l_i(A), u_j(A)$ conflict

Note: no conflict $\langle u_i(A), u_j(A) \rangle, \langle l_i(A), r_j(A) \rangle, \dots$

Theorem Rules #1,2,3 \Rightarrow conflict
(2PL) serializable
schedule

Theorem Rules #1,2,3 \Rightarrow conflict
(2PL) serializable
schedule

To help in proof:

Definition $\text{Shrink}(T_i) = \text{SH}(T_i) =$
first unlock

action of T_i

Lemma

$$T_i \rightarrow T_j \text{ in } S \Rightarrow SH(T_i) <_S SH(T_j)$$

Lemma

$T_i \rightarrow T_j \text{ in } S \Rightarrow SH(T_i) <_S SH(T_j)$

Proof of lemma:

$T_i \rightarrow T_j$ means that

$S = \dots p_i(A) \dots q_j(A) \dots; \quad p, q \text{ conflict}$

By rules 1,2:

$S = \dots p_i(A) \dots u_i(A) \dots l_j(A) \dots q_j(A) \dots$

Lemma

$T_i \rightarrow T_j \text{ in } S \Rightarrow SH(T_i) <_S SH(T_j)$


Proof of lemma:

$T_i \rightarrow T_j$ means that

$S = \dots p_i(A) \dots q_j(A) \dots$; p, q conflict

By rules 1,2:

$S = \dots p_i(A) \dots u_i(A) \dots l_j(A) \dots q_j(A) \dots$



By rule 3: $SH(T_i)$ $SH(T_j)$

So, $SH(T_i) <_S SH(T_j)$

Theorem Rules #1,2,3 \Rightarrow conflict
(2PL) serializable
schedule

Proof:

(1) Assume $P(S)$ has cycle

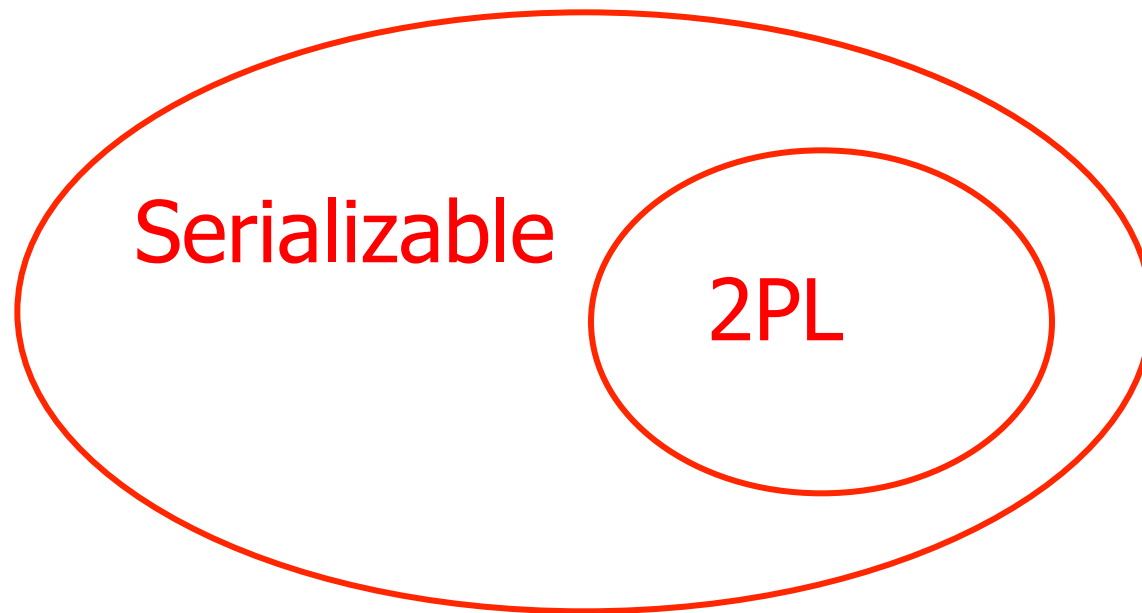
$$T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T_n \rightarrow T_1$$


(2) By lemma: $SH(T_1) < SH(T_2) < \dots < SH(T_1)$

(3) Impossible, so $P(S)$ acyclic

(4) $\Rightarrow S$ is conflict serializable

2PL subset of Serializable





S1: w1(x) w3(x) w2(y) w1(y)

S1: w1(x) w3(x) w2(y) w1(y)

- S1 cannot be achieved via 2PL:
The lock by T1 for y must occur after w2(y), so the unlock by T1 for x must occur after this point (and before w1(x)). Thus, w3(x) cannot occur under 2PL where shown in S1 because T1 holds the x lock at that point.
- However, S1 is serializable (equivalent to T2, T1, T3).

If you need a bit more practice:

Are our schedules S_C and S_D 2PL schedules?

S_C : w1(A) w2(A) w1(B) w2(B)

S_D : w1(A) w2(A) w2(B) w1(B)

- Beyond this simple 2PL protocol, it is all a matter of improving performance and allowing more concurrency....
 - Shared locks
 - Multiple granularity
 - Inserts, deletes and phantoms
 - Other types of C.C. mechanisms

Shared locks

So far:

$S = \dots l_1(A) r_1(A) u_1(A) \dots l_2(A) r_2(A) u_2(A) \dots$




Do not conflict

Shared locks

So far:

$S = \dots l_1(A) r_1(A) u_1(A) \dots l_2(A) r_2(A) u_2(A) \dots$



Instead:

$S = \dots l_{s1}(A) r_1(A) l_{s2}(A) r_2(A) \dots u_{s1}(A) u_{s2}(A)$

Lock actions

$l-t_i(A)$: lock A in t mode (t is S or X)

$u-t_i(A)$: unlock t mode (t is S or X)

Shorthand:

$u_i(A)$: unlock whatever modes

T_i has locked A

Rule #1 Well formed transactions

$T_i = \dots I-S_1(A) \dots r_1(A) \dots u_1(A) \dots$

$T_i = \dots I-X_1(A) \dots w_1(A) \dots u_1(A) \dots$

- What about transactions that read and write same object?

Option 1: Request exclusive lock

$T_i = \dots l-X_1(A) \dots r_1(A) \dots w_1(A) \dots u(A) \dots$

- What about transactions that read and write same object?

Option 2: Upgrade

(E.g., need to read, but don't know if will write...)


$T_i = \dots I-S_1(A) \dots r_1(A) \dots I-X_1(A) \dots w_1(A) \dots u(A) \dots$

Think of

- Get 2nd lock on A, or
- Drop S, get X lock


Rule #2 Legal scheduler

$S = \dots l-S_i(A) \dots \dots u_i(A) \dots$



no $l-X_j(A)$

$S = \dots l-X_i(A) \dots \dots u_i(A) \dots$



no $l-X_j(A)$

no $l-S_j(A)$

A way to summarize Rule #2

Compatibility matrix

Comp

	S	X
S	true	false
X	false	false

Rule # 3 2PL transactions

No change except for upgrades:

(I) If upgrade gets more locks

(e.g., $S \rightarrow \{S, X\}$) then no change!

(II) If upgrade releases read (shared) lock (e.g., $S \rightarrow X$)

- can be allowed in growing phase

Theorem Rules 1,2,3 \Rightarrow Conf.serializable
for S/X locks schedules

Proof: similar to X locks case

Detail:

$l-t_i(A), l-r_j(A)$ do not conflict if $\text{comp}(t,r)$

$l-t_i(A), u-r_j(A)$ do not conflict if $\text{comp}(t,r)$


Lock types beyond S/X

Examples:

- (1) increment lock
- (2) update lock

Example (1): increment lock

- Atomic increment action: $IN_i(A)$
 $\{\text{Read}(A); A \leftarrow A+k; \text{Write}(A)\}$
- $IN_i(A), IN_j(A)$ do not conflict!



Comp


	S	X	I
S			
X			
I			

Comp

	S	X	I
S	T	F	F
X	F	F	F
I	F	F	T

Update locks

A common deadlock problem with upgrades:



Solution

If T_i wants to read A and knows it may later want to write A , it requests update lock (not shared)

New request

Comp

Lock
already
held in

	S	X	U
S			
X			
U			

New request

Comp

Lock
already
held in

	S	X	U
S	T	F	T
X	F	F	F
U	TorF	F	F

-> symmetric table?

Note: object A may be locked in different modes at the same time...

$$S_1 = \dots | - S_1(A) \dots | - S_2(A) \dots | - U_3(A) \dots \left\{ \begin{array}{l} | - S_4(A) \dots ? \\ | - U_4(A) \dots ? \end{array} \right.$$

Note: object A may be locked in different modes at the same time...

$$S_1 = \dots I-S_1(A) \dots I-S_2(A) \dots I-U_3(A) \dots \left\{ \begin{array}{l} I-S_4(A) \dots ? \\ I-U_4(A) \dots ? \end{array} \right.$$


- To grant a lock in mode t, mode t must be compatible with all currently held locks on object

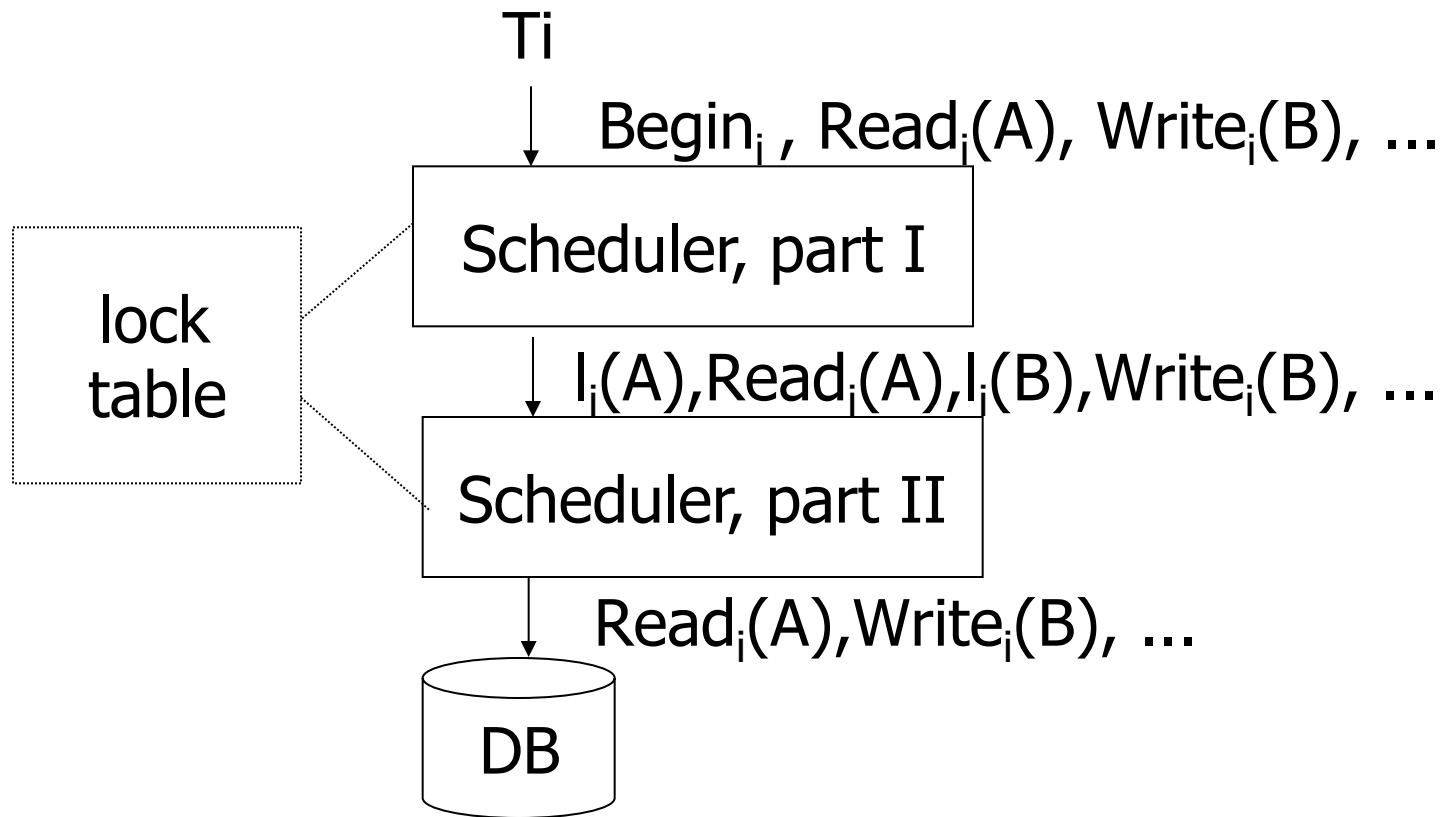
How does locking work in practice?

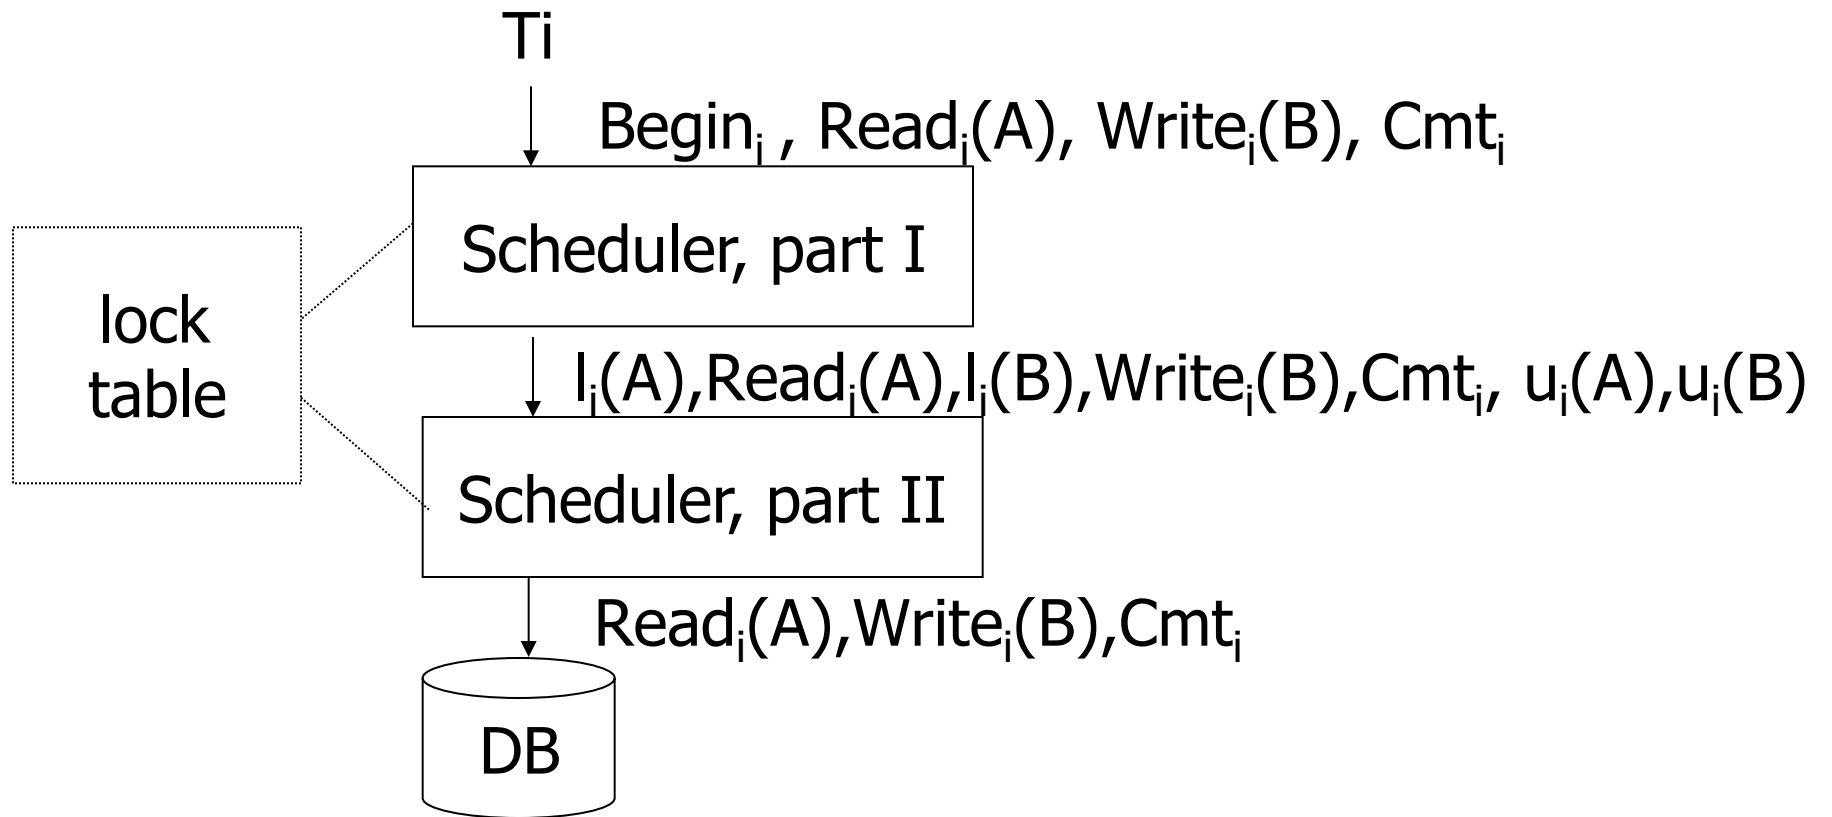
- Every system is different
(E.g., may not even provide
CONFLICT-SERIALIZABLE schedules)
- But here is one (simplified) way ...

Sample Locking System:

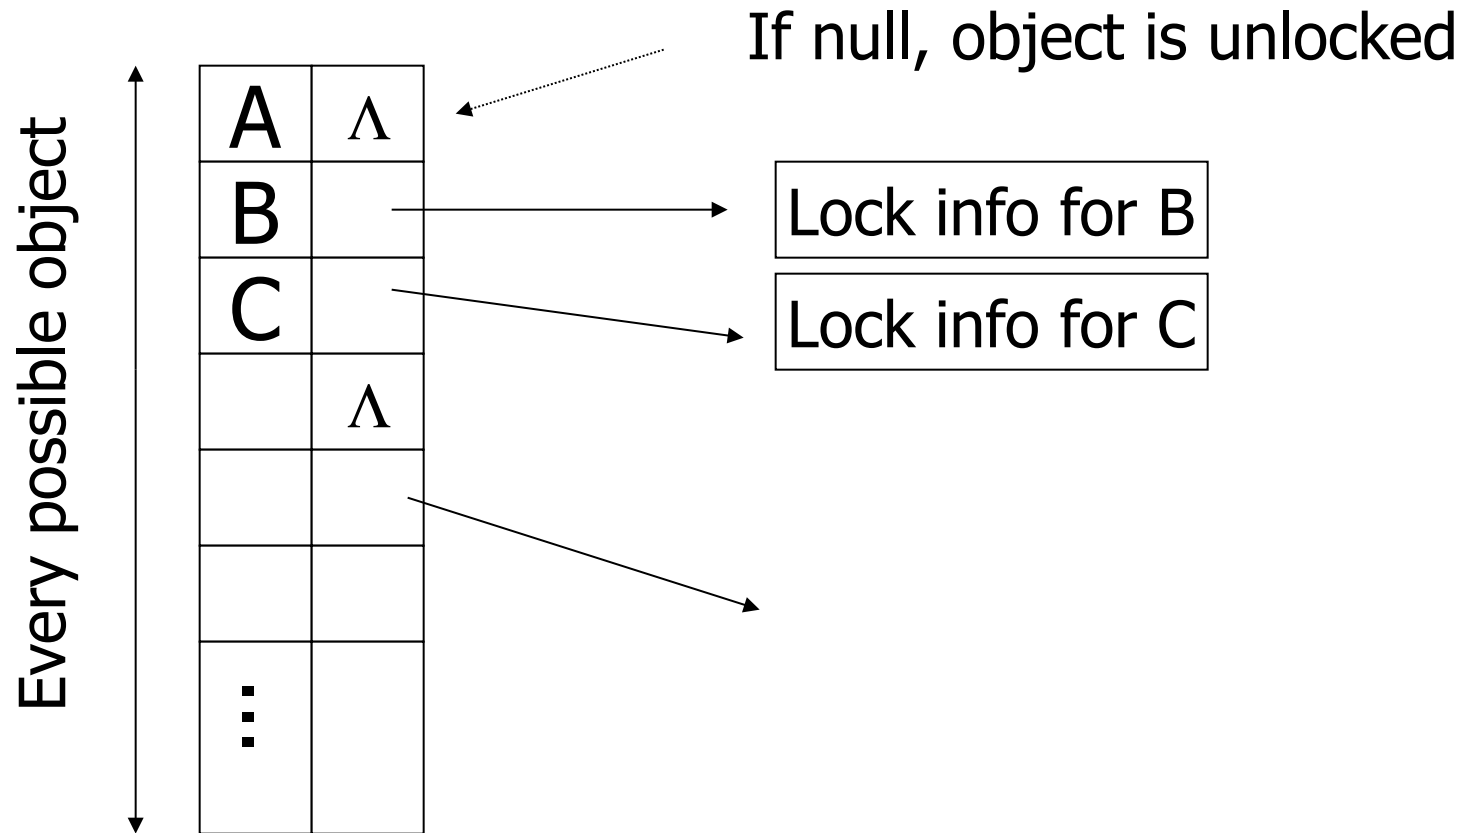
- (1) Don't trust transactions to request/release locks
- (2) Hold all locks until transaction commits








Lock table Conceptually

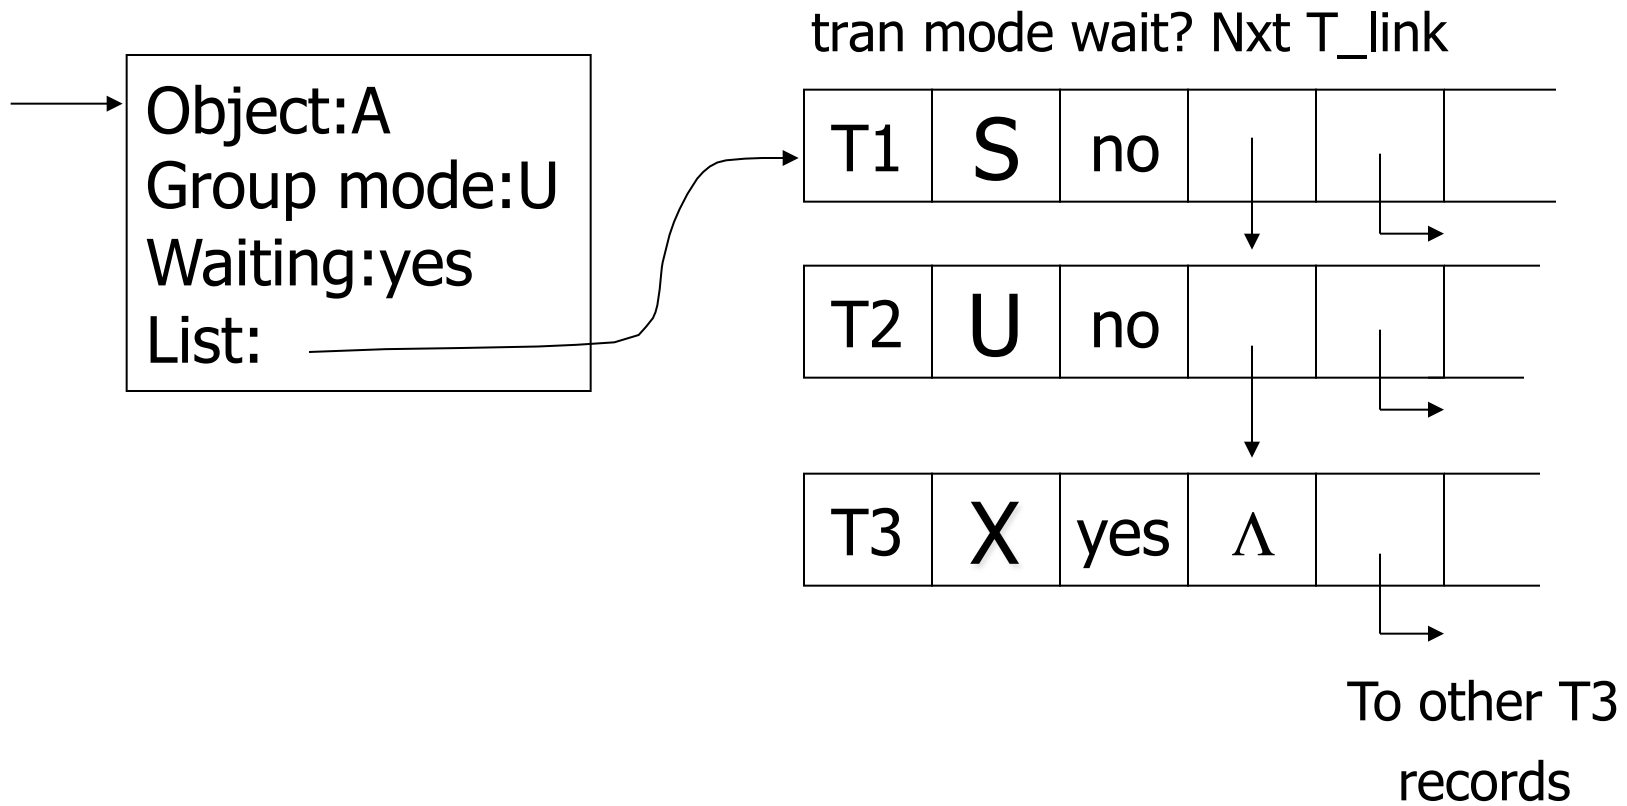


But use hash table:




If object not found in hash table, it is unlocked


Lock info for A - example




What are the objects we lock?



DB



DB



DB


?

- Locking works in any case, but should we choose small or large objects?


- Locking works in any case, but should we choose small or large objects?
- If we lock large objects (e.g., Relations)
 - Need few locks
 - Low concurrency
- If we lock small objects (e.g., tuples, fields)
 - Need more locks
 - More concurrency

We can have it both ways!!


Ask any janitor to give you the solution...




Example




Example




Example



Example (b)



Example



Multiple granularity

Comp

Requestor

IS IX S SIX X

Holder

IS					
IX					
S					
SIX					
X					


Multiple granularity

Comp


Requestor

		IS	IX	S	SIX	X
Holder	IS	T	T	T	T	F
	IX	T	T	F	F	F
	S	T	F	T	F	F
	SIX	T	F	F	F	F
	X	F	F	F	F	F

Parent locked in	Child can be locked in
IS	
IX	
S	
SIX	
X	



Parent locked in	Child can be locked by same transaction in
IS	IS, S
IX	IS, S, IX, X, SIX
S	none
SIX	X, IX, [SIX]
X	none



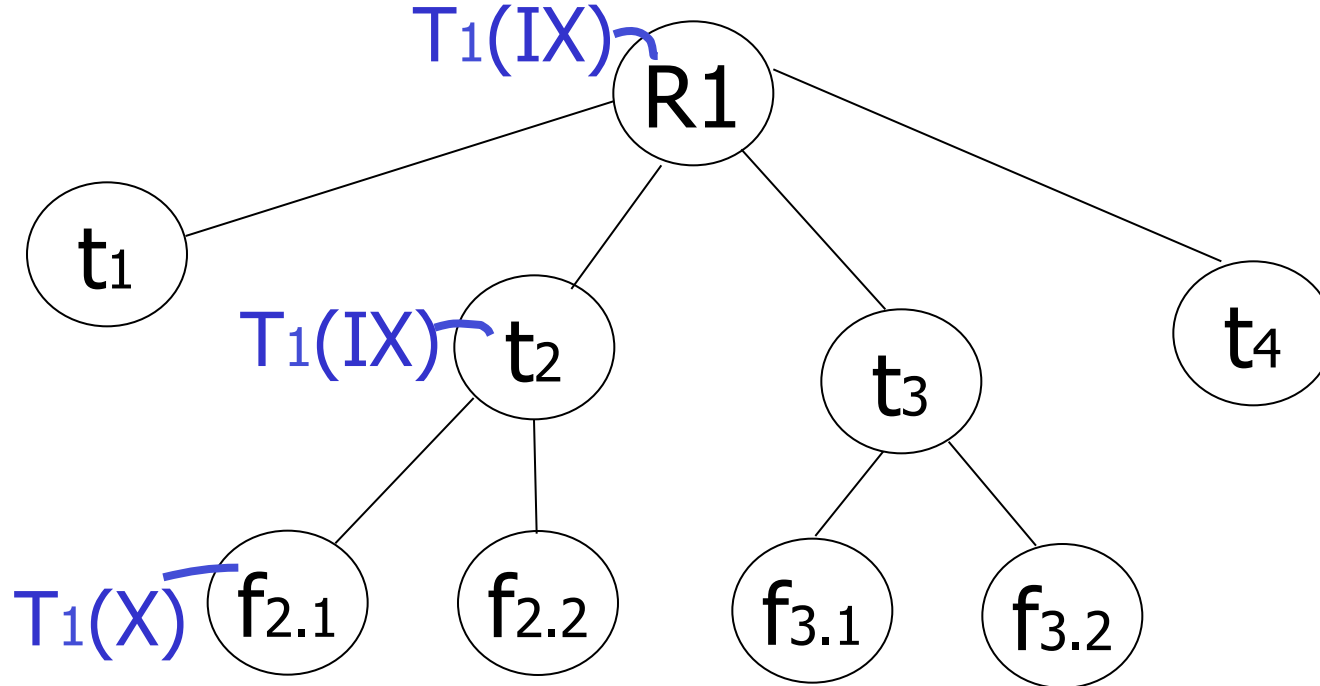
not necessary

Rules

- (1) Follow multiple granularity comp function
- (2) Lock root of tree first, any mode
- (3) Node Q can be locked by T_i in S or IS only if $\text{parent}(Q)$ locked by T_i in IX or IS
- (4) Node Q can be locked by T_i in X,SIX,IX only if $\text{parent}(Q)$ locked by T_i in IX,SIX
- (5) T_i is two-phase
- (6) T_i can unlock node Q only if none of Q's children are locked by T_i


Exercise:

- Can T2 access object f2.2 in X mode?
What locks will T2 get?



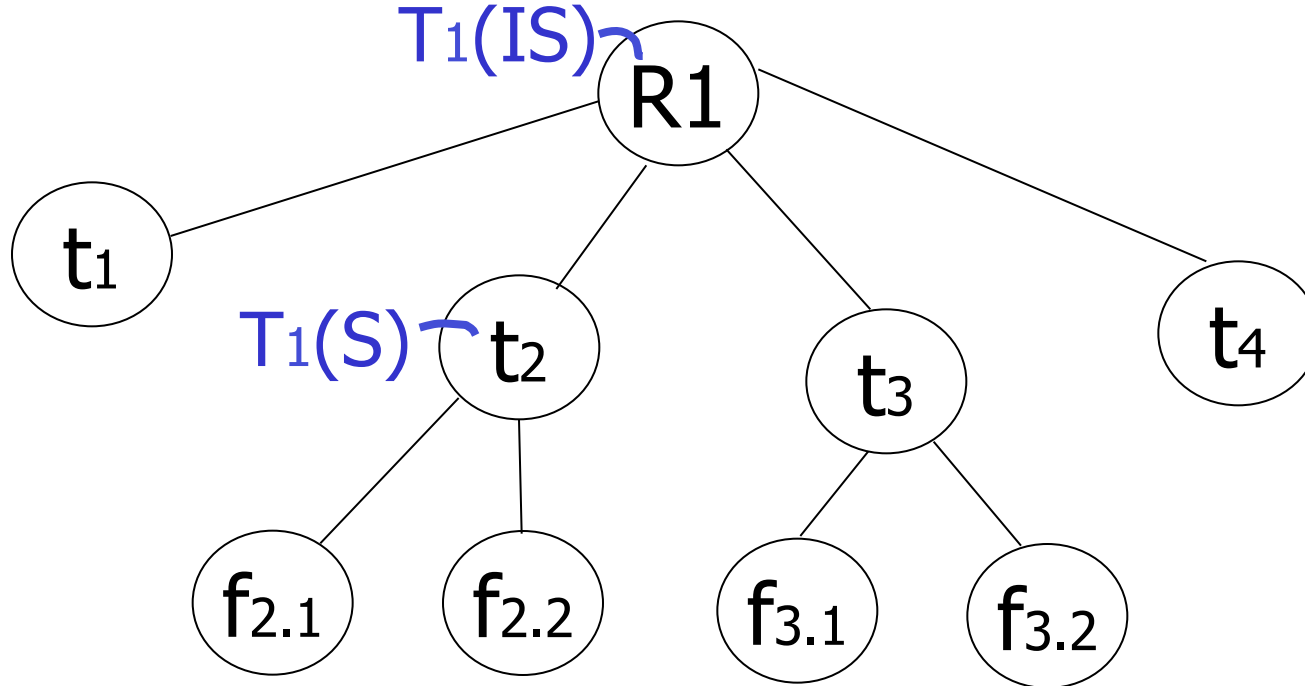
Exercise:

- Can T2 access object f2.2 in X mode?
What locks will T2 get?



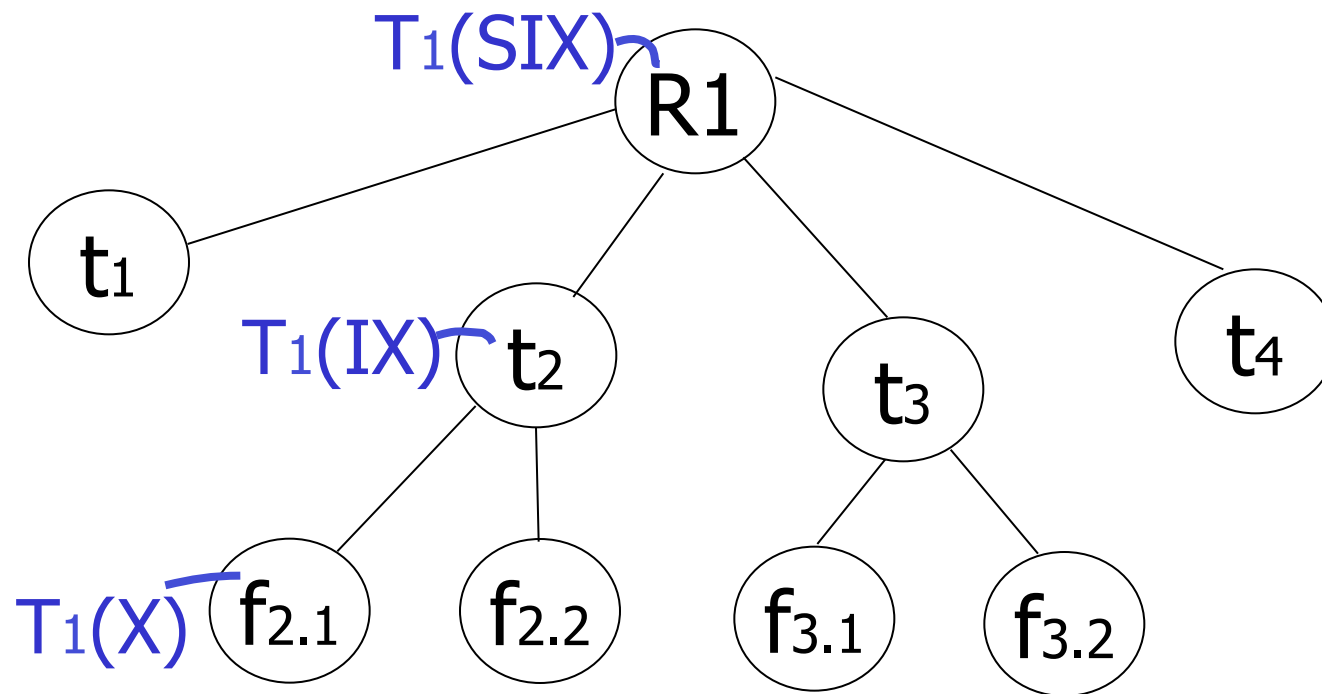
Exercise:

- Can T2 access object f3.1 in X mode?
What locks will T2 get?




Exercise:

- Can T2 access object f2.2 in S mode?
What locks will T2 get?




Exercise:

- Can T2 access object f2.2 in X mode?
What locks will T2 get?



Insert + delete operations



Modifications to locking rules:

- (1) Get exclusive lock on A before deleting A
- (2) At insert A operation by T_i ,
 T_i is given exclusive lock on A

Still have a problem: **Phantoms**

Example: relation R (E#,name,...)

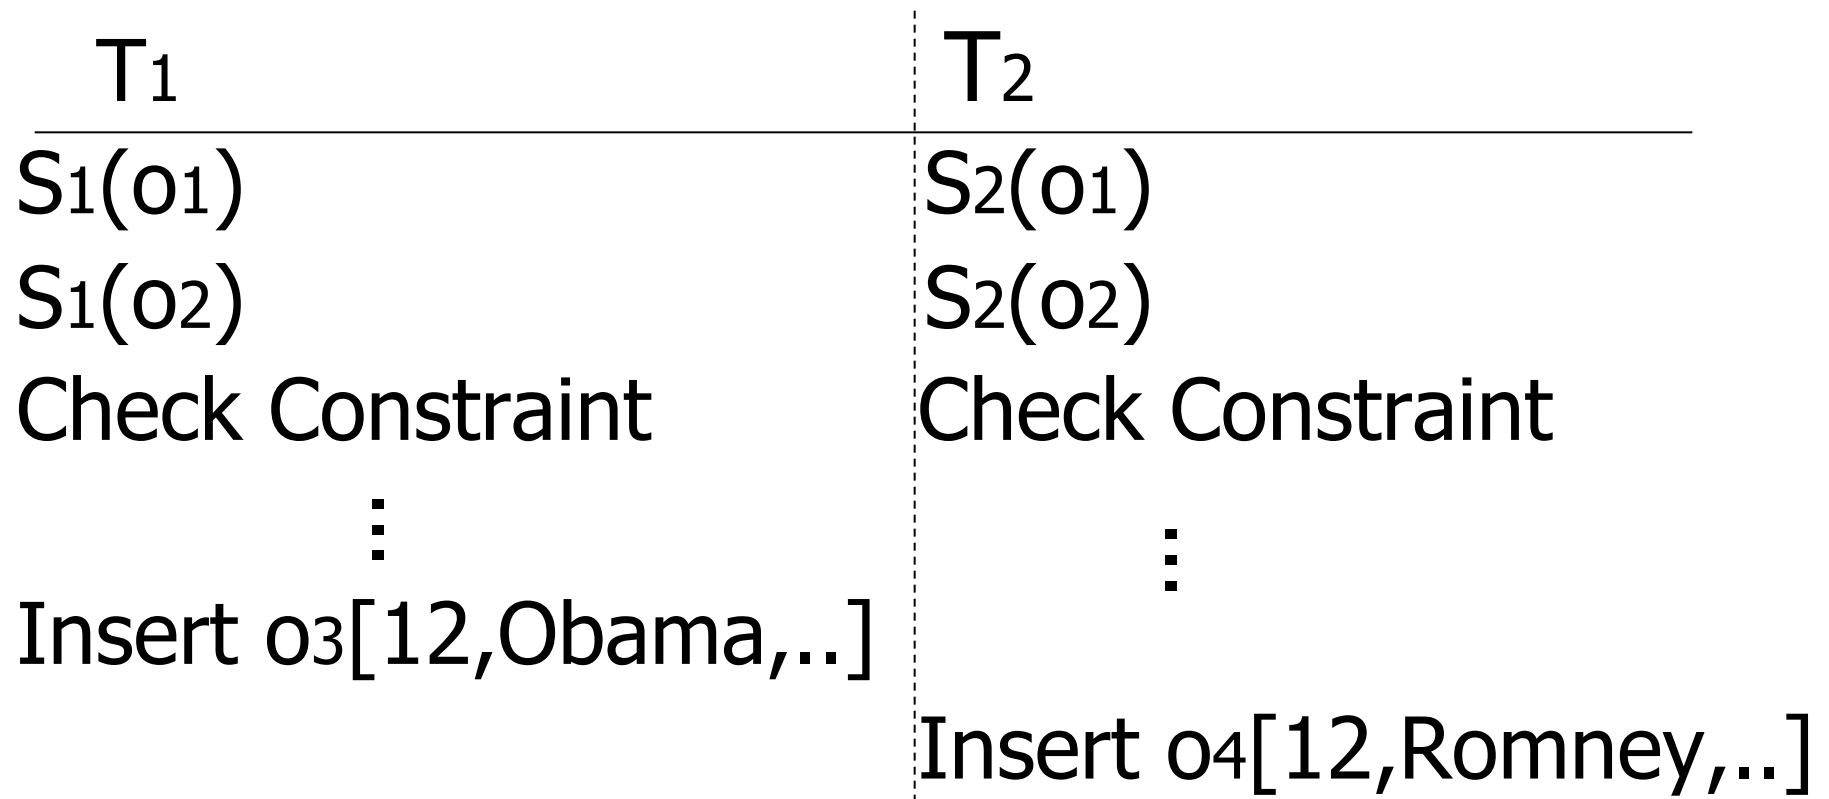
constraint: E# is key

use tuple locking

R	E#	Name
o1	55	Smith	
o2	75	Jones	


T₁: Insert <12,Obama,...> into R

T₂: Insert <12,Romney,...> into R



Solution

- Use multiple granularity tree
- Before insert of node Q,
lock parent(Q) in
X mode



Back to example

T₁: Insert<12,Obama>

T₁

X₁(R)

Check constraint

Insert<12,Obama>

U₁(R)

T₂: Insert<12,Romney>

T₂

X₂(R) ← *delayed*

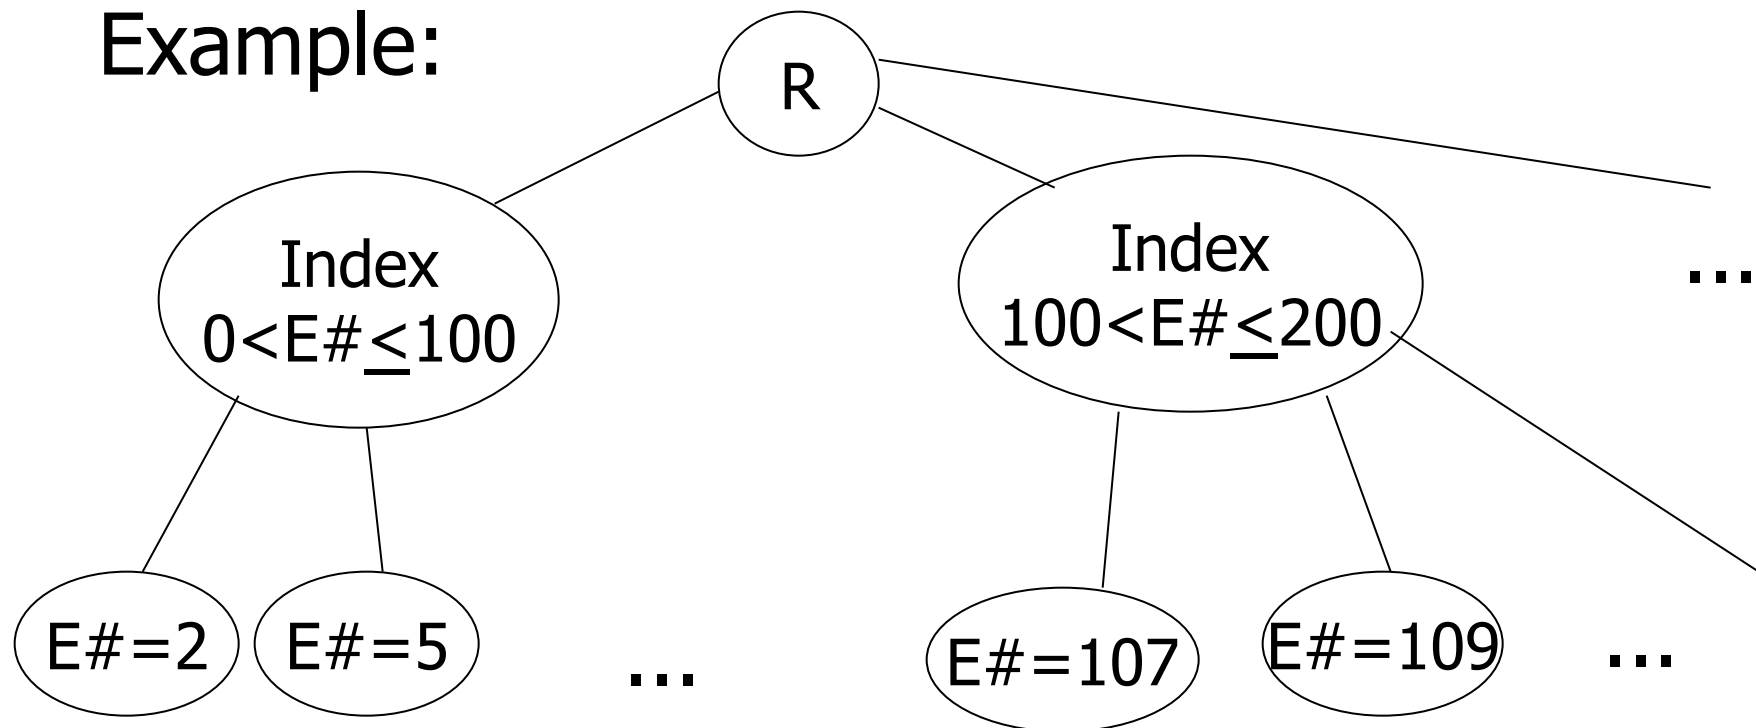
X₂(R)

Check constraint

Oops! e# = 12 already in R!

Instead of using R, can use index on R:

Example:




- This approach can be generalized to multiple indexes...

Next:

- Tree-based concurrency control
- Validation concurrency control


Example

- all objects accessed through root, following pointers




Example

- all objects accessed through root, following pointers





Example

- all objects accessed through root, following pointers





- can we release A lock if we no longer need A??



Idea: traverse like "Monkey Bars"



Idea: traverse like "Monkey Bars"




Idea: traverse like "Monkey Bars"



Why does this work?

- Assume all T_i start at root; exclusive lock
- $T_i \rightarrow T_j \Rightarrow T_i$ locks root before T_j




- Actually works if we don't always start at root

Rules: tree protocol (exclusive locks)

- (1) First lock by T_i may be on any item
- (2) After that, item Q can be locked by T_i
only if $\text{parent}(Q)$ locked by T_i
- (3) Items may be unlocked at any time
- (4) After T_i unlocks Q , it cannot relock Q

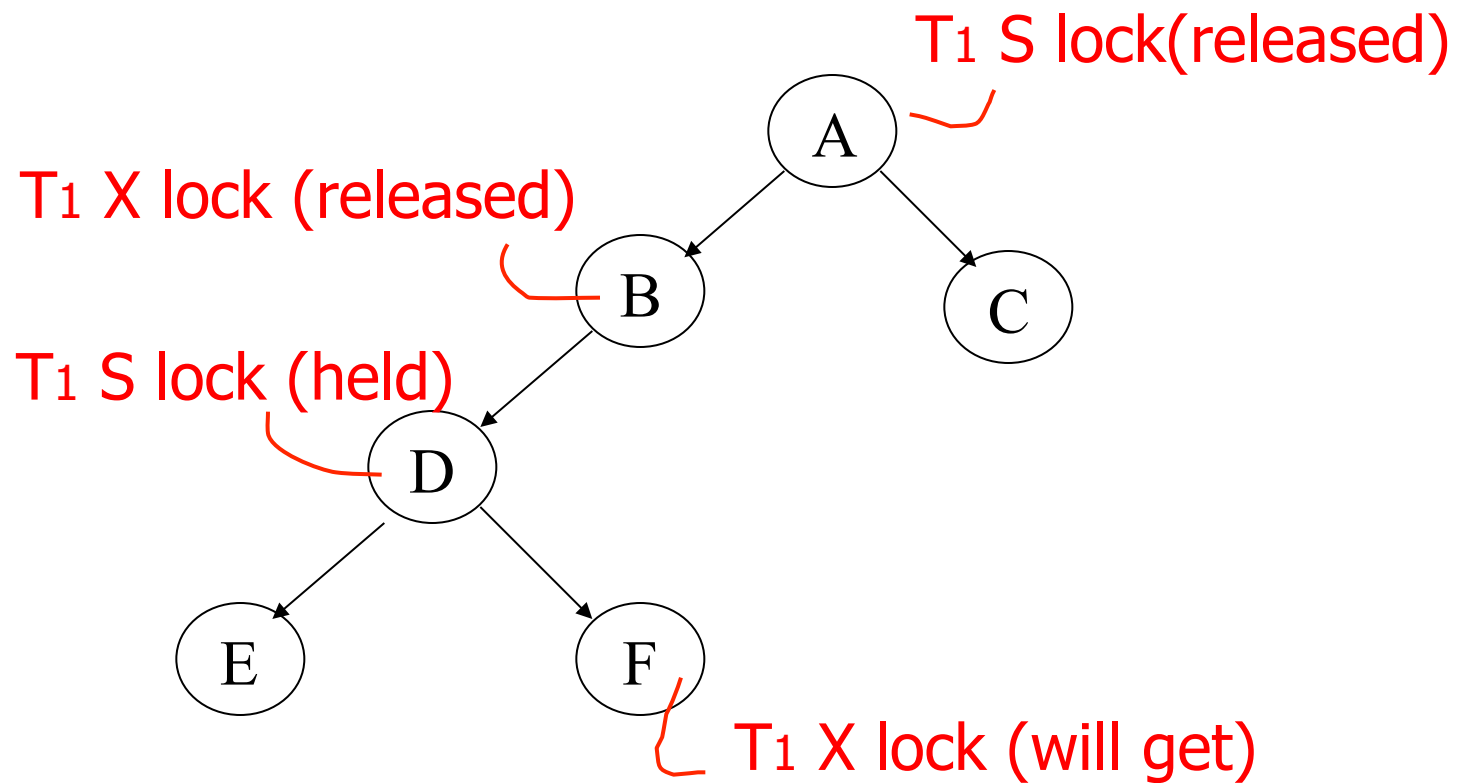
- Tree-like protocols are used typically for B-tree concurrency control



E.g., during insert, do not release parent lock, until you are certain child does not have to split

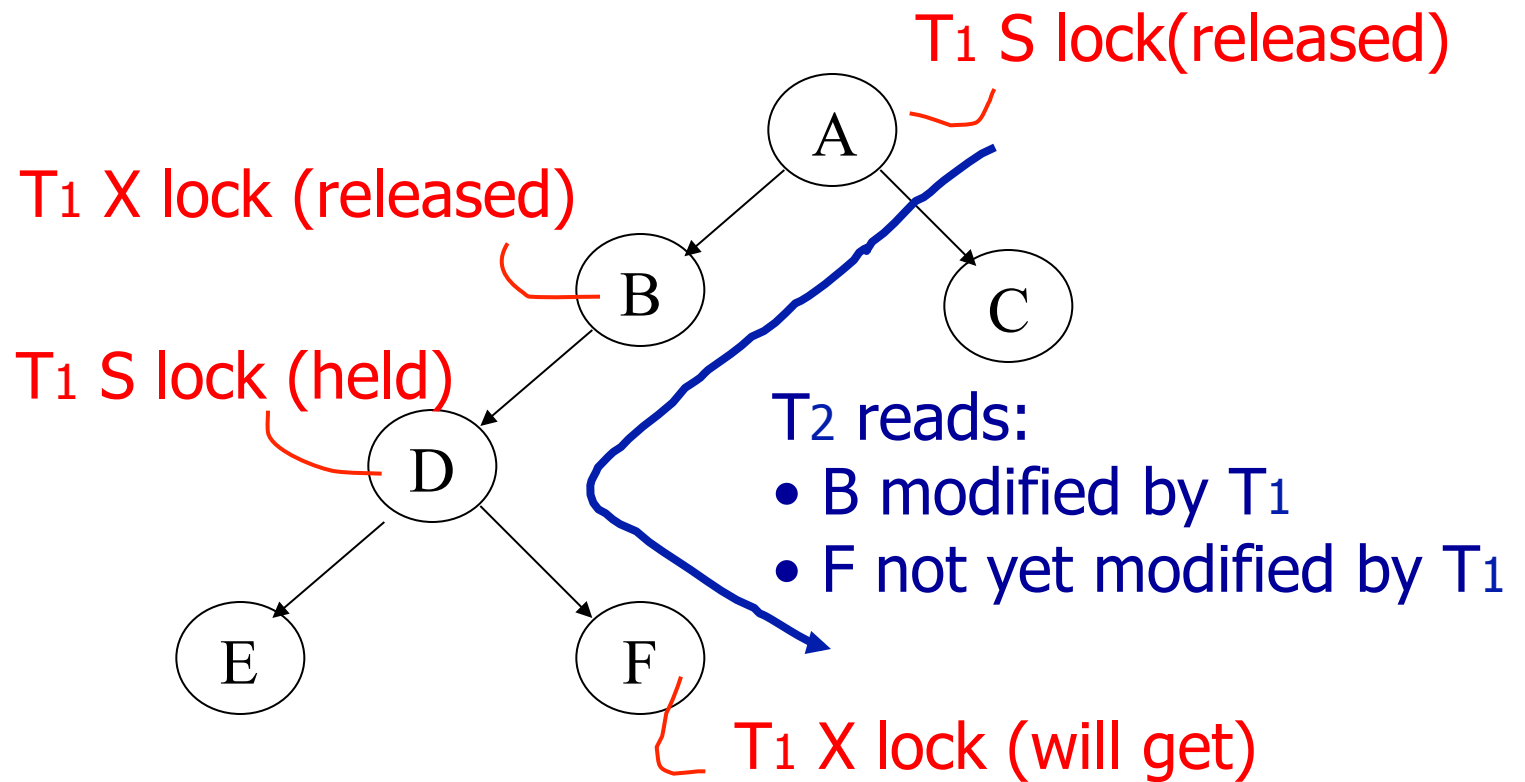
Tree Protocol with Shared Locks

- Rules for shared & exclusive locks?



Tree Protocol with Shared Locks

- Rules for shared & exclusive locks?



Tree Protocol with Shared Locks

- Need more restrictive protocol
- Will this work??
 - Once T_1 locks one object in X mode, all further locks down the tree must be in X mode

Validation

Transactions have 3 phases:

(1) Read

- all DB values read
- writes to temporary storage
- no locking

(2) Validate

- check if schedule so far is serializable

(3) Write

- if validate ok, write to DB

Key idea


- Make validation atomic
- If T_1, T_2, T_3, \dots is validation order, then resulting schedule will be conflict equivalent to $S_s = T_1 T_2 T_3 \dots$

To implement validation, system keeps two sets:

- FIN = transactions that have finished phase 3 (and are all done)
- VAL = transactions that have successfully finished phase 2 (validation)

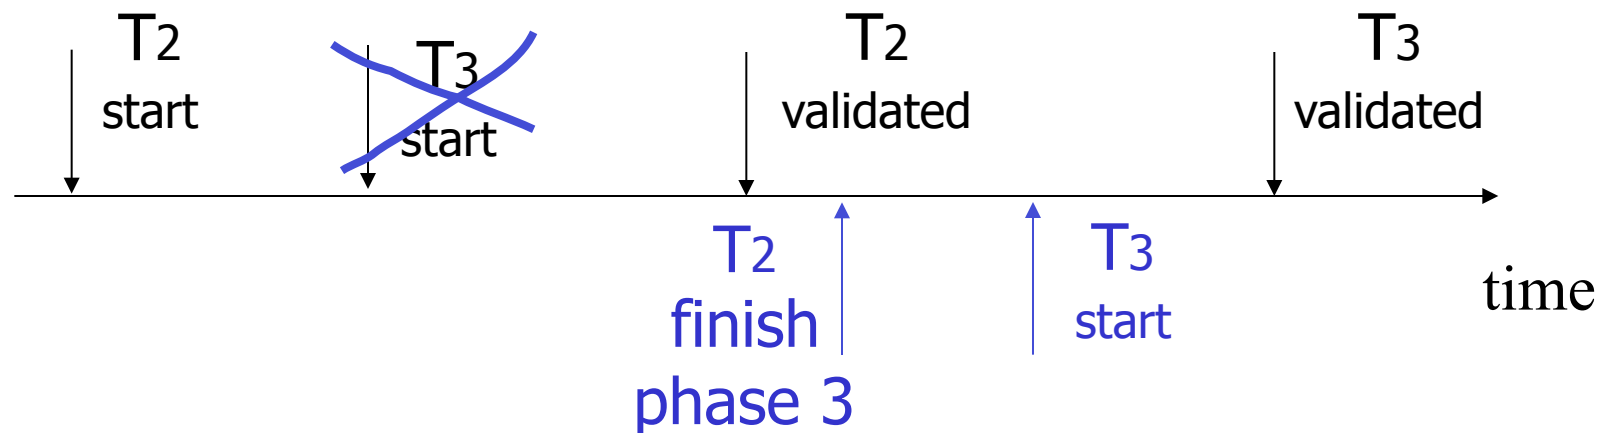
Example of what validation must prevent:

$$\begin{array}{l} RS(T_2) = \{B\} \\ WS(T_2) = \{B, D\} \end{array} \cap \begin{array}{l} RS(T_3) = \{A, B\} \neq \emptyset \\ WS(T_3) = \{C\} \end{array}$$



Example of what validation must ~~prevent~~ allow:

$$\begin{array}{l} RS(T_2) = \{B\} \\ WS(T_2) = \{B, D\} \end{array} \cap \begin{array}{l} RS(T_3) = \{A, B\} \neq \emptyset \\ WS(T_3) = \{C\} \end{array}$$




Another thing validation must prevent:

$$RS(T_2) = \{A\}$$

$$RS(T_3) = \{A, B\}$$

$$WS(T_2) = \{D, E\}$$

$$WS(T_3) = \{C, D\}$$



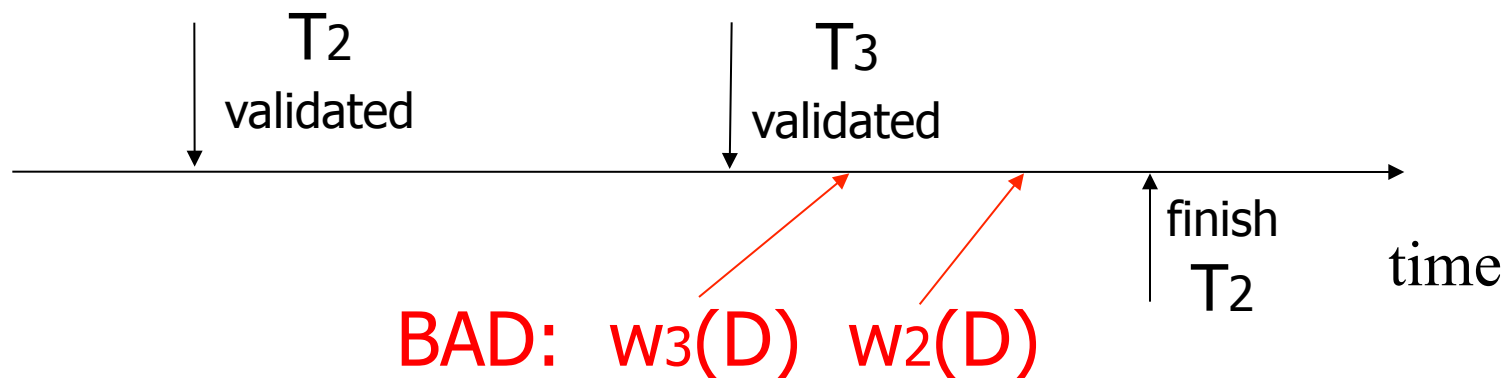
Another thing validation must prevent:

$RS(T_2) = \{A\}$

$RS(T_3) = \{A, B\}$

$WS(T_2) = \{D, E\}$

$WS(T_3) = \{C, D\}$



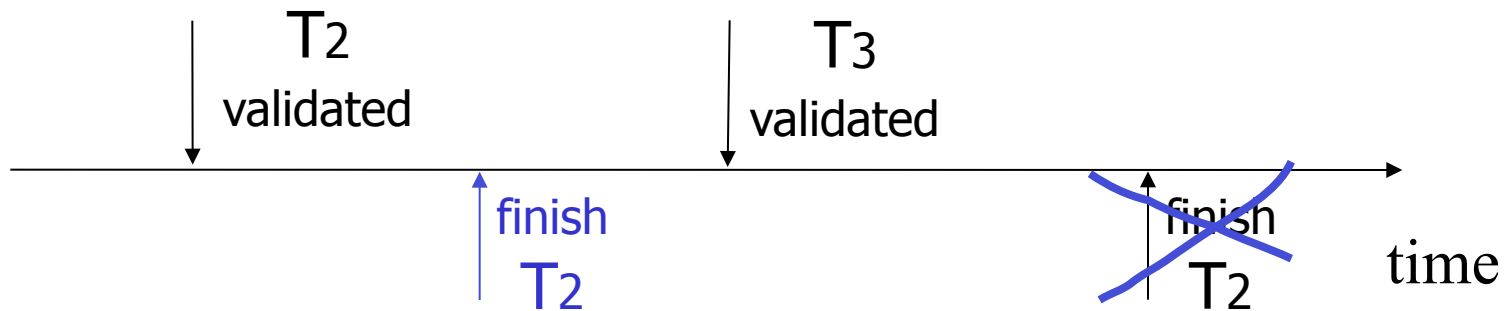
Another thing validation must ~~prevent~~ allow:

$$RS(T_2) = \{A\}$$

$$RS(T_3) = \{A, B\}$$

$$WS(T_2) = \{D, E\}$$

$$WS(T_3) = \{C, D\}$$



Validation rules for T_j :

(1) When T_j starts phase 1:

$\text{ignore}(T_j) \leftarrow \text{FIN}$

(2) at T_j Validation:

if check (T_j) then

[$\text{VAL} \leftarrow \text{VAL} \cup \{T_j\};$

do write phase;

$\text{FIN} \leftarrow \text{FIN} \cup \{T_j\}$]

Check (T_j):

For $T_i \in \text{VAL} - \text{IGNORE}(T_j)$ DO

IF [$\text{WS}(T_i) \cap \text{RS}(T_j) \neq \emptyset$ OR

$T_i \notin \text{FIN}$] THEN RETURN false;

RETURN true;

Check (T_j):

For $T_i \in \text{VAL} - \text{IGNORE}(T_j)$ DO

IF [$\text{WS}(T_i) \cap \text{RS}(T_j) \neq \emptyset$ OR

$T_i \notin \text{FIN}$] THEN RETURN false;

RETURN true;

Is this check too restrictive ?

Improving Check(T_j)

For $T_i \in \text{VAL} - \text{IGNORE}(T_j)$ DO


IF [$\text{WS}(T_i) \cap \text{RS}(T_j) \neq \emptyset$ OR

$(T_i \notin \text{FIN} \text{ AND } \text{WS}(T_i) \cap \text{WS}(T_j) \neq \emptyset)$]

THEN RETURN false;


RETURN true;

Exercise:



U: $RS(U) = \{B\}$
 $WS(U) = \{D\}$


W: $RS(W) = \{A, D\}$
 $WS(W) = \{A, C\}$



T: $RS(T) = \{A, B\}$
 $WS(T) = \{A, C\}$

V: $RS(V) = \{B\}$
 $WS(V) = \{D, E\}$

Is Validation = 2PL?



S2: $w_2(y) \quad w_1(x) \quad w_2(x)$

- Achievable with 2PL?
- Achievable with validation?

S2: w2(y) w1(x) w2(x)

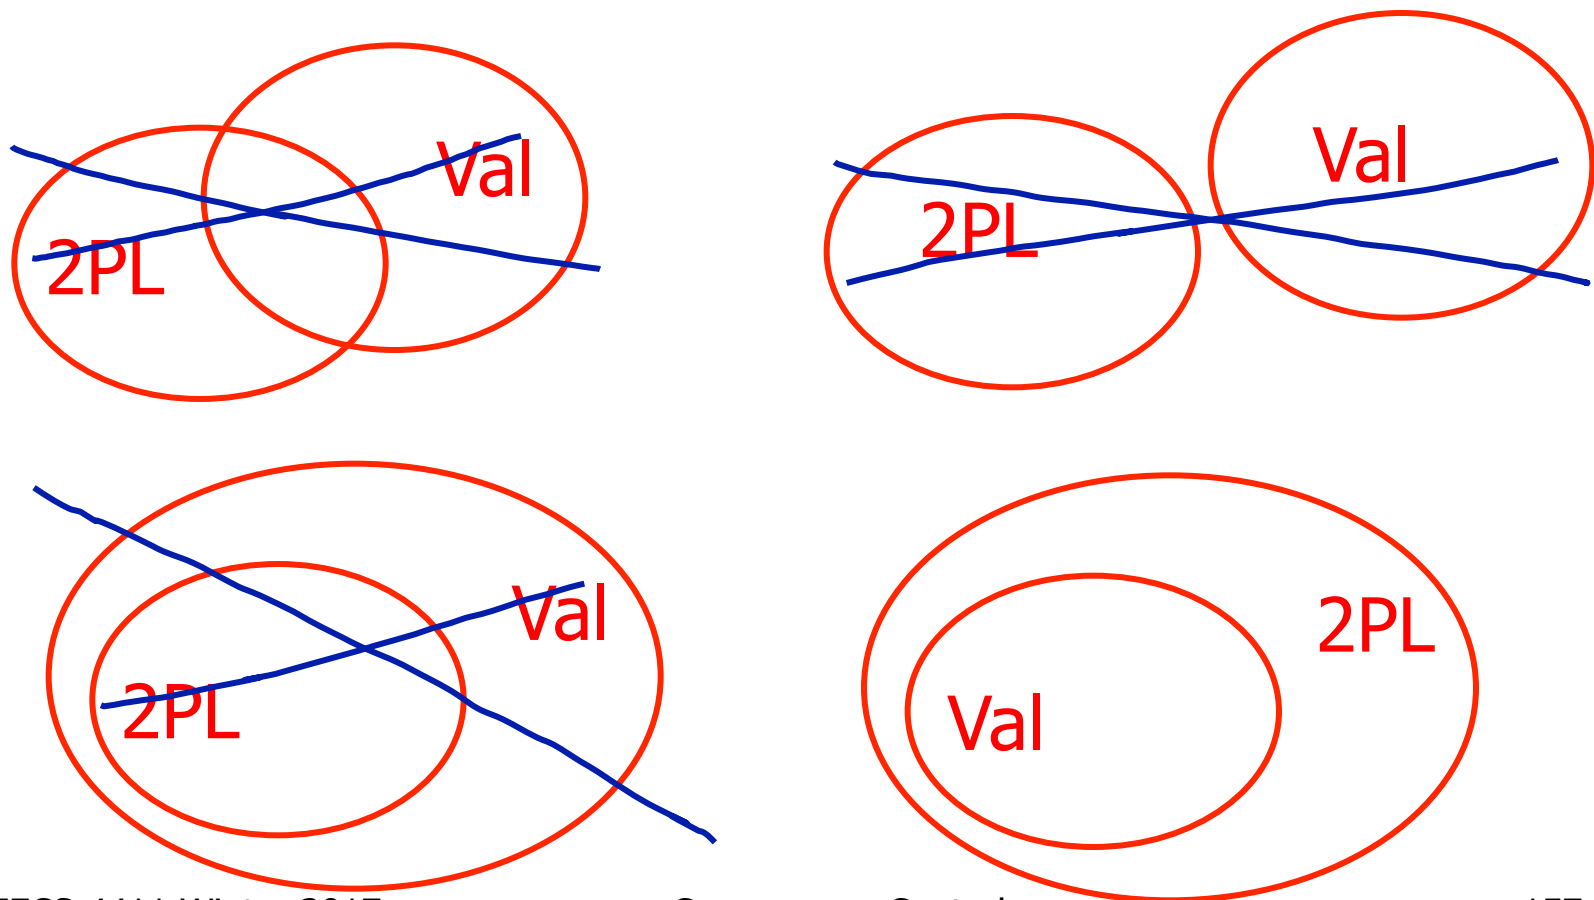
- S2 can be achieved with 2PL:
l2(y) w2(y) l1(x) w1(x) u1(x) l2(x) w2(x) u2(y) u2(x)
- S2 cannot be achieved by validation:
The validation point of T2, val2 must occur before w2(y) since transactions do not write to the database until after validation. Because of the conflict on x, val1 < val2, so we must have something like
S2: val1 val2 w2(y) w1(x) w2(x)
With the validation protocol, the writes of T2 should not start until T1 is all done with its writes, which is not the case.

Validation subset of 2PL?

- Possible proof (Check!):
 - Let S be validation schedule
 - For each T in S insert lock/unlocks, get S' :
 - At T start: request read locks for all of $RS(T)$
 - At T validation: request write locks for $WS(T)$; release read locks for read-only objects
 - At T end: release all write locks
 - Clearly transactions well-formed and 2PL
 - Must show S' is legal (next page)

- Say S' not legal (due to w-r conflict):
 $S': \dots l1(x) \quad w2(x) \quad r1(x) \quad val1 \quad u1(x) \dots$
 - At $val1$: T2 not in Ignore(T1); T2 in VAL
 - T1 does not validate: $WS(T2) \cap RS(T1) \neq \emptyset$
 - contradiction!
- Say S' not legal (due to w-w conflict):
 $S': \dots val1 \quad l1(x) \quad w2(x) \quad w1(x) \quad u1(x) \dots$
 - Say T2 validates first (proof similar if T1 validates first)
 - At $val1$: T2 not in Ignore(T1); T2 in VAL
 - T1 does not validate:
 $T2 \notin FIN \quad \text{AND} \quad WS(T1) \cap WS(T2) \neq \emptyset$
 - contradiction!

Conclusion: Validation subset 2PL



Validation (also called optimistic concurrency control) is useful in some cases:

- Conflicts rare
- System resources plentiful
- Have real time constraints

Summary

Have studied C.C. mechanisms used in practice

- 2 PL
- Multiple granularity
- Tree (index) protocols
- Validation