



CSCD43: Database Systems Technology

Lecture 13

Wael Aboulsaadat

Acknowledgment: these slides are based on Prof. Garcia-Molina & Prof. Ullman slides accompanying the course's textbook.



Enforcing Serializability

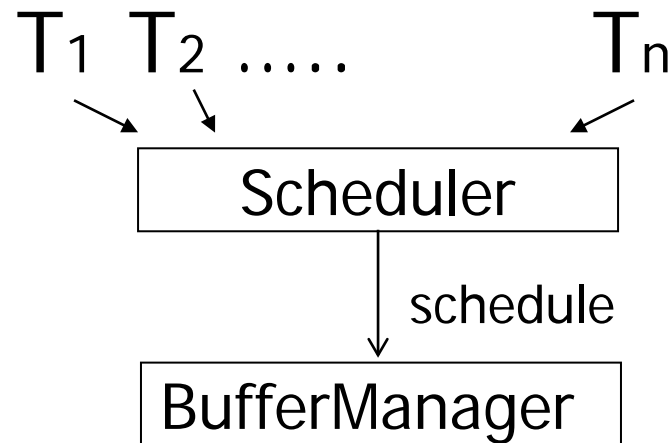


How to enforce serializable schedules?

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

How to enforce serializable schedules?

Option 2: prevent P(S) cycles from occurring



But how ???

2.A) Buffer transactions during n seconds, stop DBMS, make schedule, execute schedule, repeat...


unrealistic...!

2.B) Use a locking protocol!

Two new actions:

lock (exclusive): $l_i(A)$

unlock: $u_i(A)$





Rule #1: Well-formed transactions

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

Locking Example

T1: Read(A); A =A+100; Write(A);
 Read(B); B=B+100; Write(B);

T2: Read(A); A =A*2; Write(A);
 Read(B); B=B*2; Write(B);



Serial Schedule

T1	T2
Read(A)	
$A \leftarrow A + 100$; Write(A)	
Read(B);	
$B \leftarrow B + 100$; Write(B);	
	Read(A)
	$A \leftarrow Ax2$; Write(A);
	Read(B)
	$B \leftarrow Bx2$; Write(B);

A	B
25	25
125	
	125
250	
	250
250	250

Schedule A

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 A

		A	B
T1	T2	25	25
$I_1(A); \text{Read}(A)$			
$A \leftarrow A + 100; \text{Write}(A); u_1(A)$		125	
	$I_2(A); \text{Read}(A)$		
	$A \leftarrow A \times 2; \text{Write}(A); u_2(A)$	250	
	$I_2(B); \text{Read}(B)$		
	$B \leftarrow B \times 2; \text{Write}(B); u_2(B)$		50
$I_1(B); \text{Read}(B)$			
$B \leftarrow B + 100; \text{Write}(B); u_1(B)$			150
		250	150




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


$T_i = \dots \dots \text{li}(A) \dots \dots \text{ui}(A) \dots \dots$



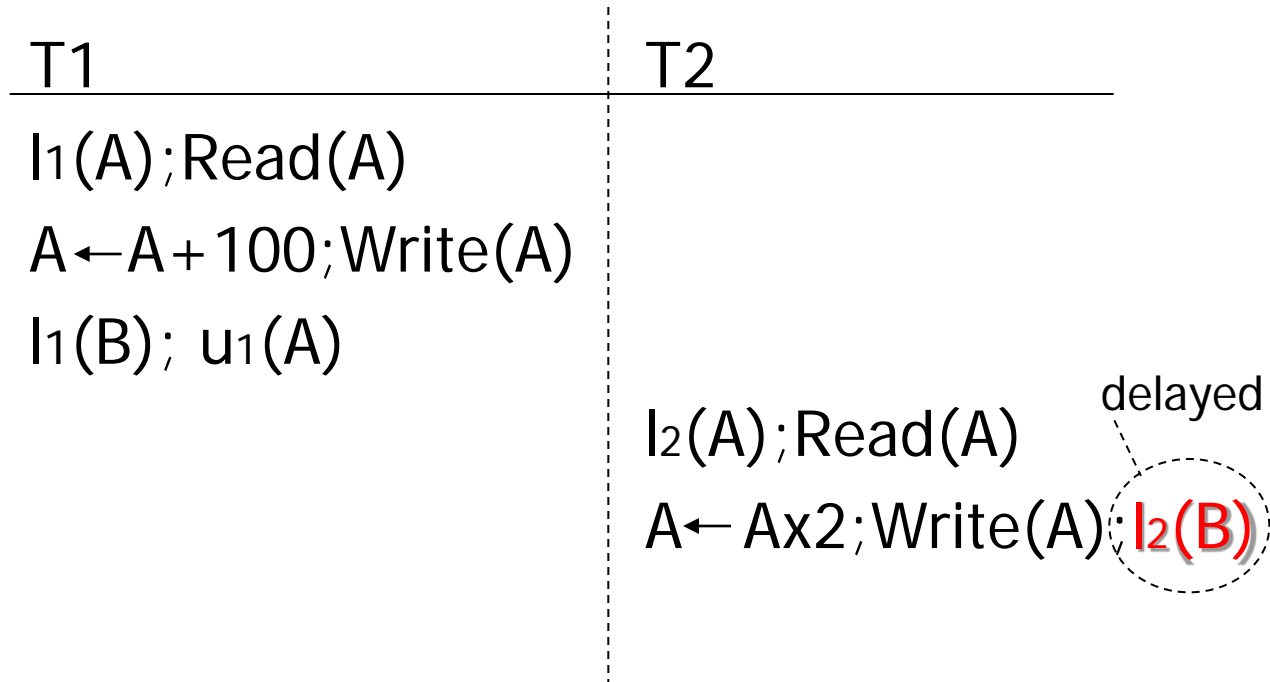
no unlocks



no locks




Schedule B



Schedule B

T1	T2
$I_1(A); \text{Read}(A)$	
$A \leftarrow A + 100; \text{Write}(A)$	
$I_1(B); u_1(A)$	
	$I_2(A); \text{Read}(A)$
	$A \leftarrow Ax2; \text{Write}(A); I_2(B)$
$\text{Read}(B); B \leftarrow B + 100$	
$\text{Write}(B); u_1(B)$	

delayed



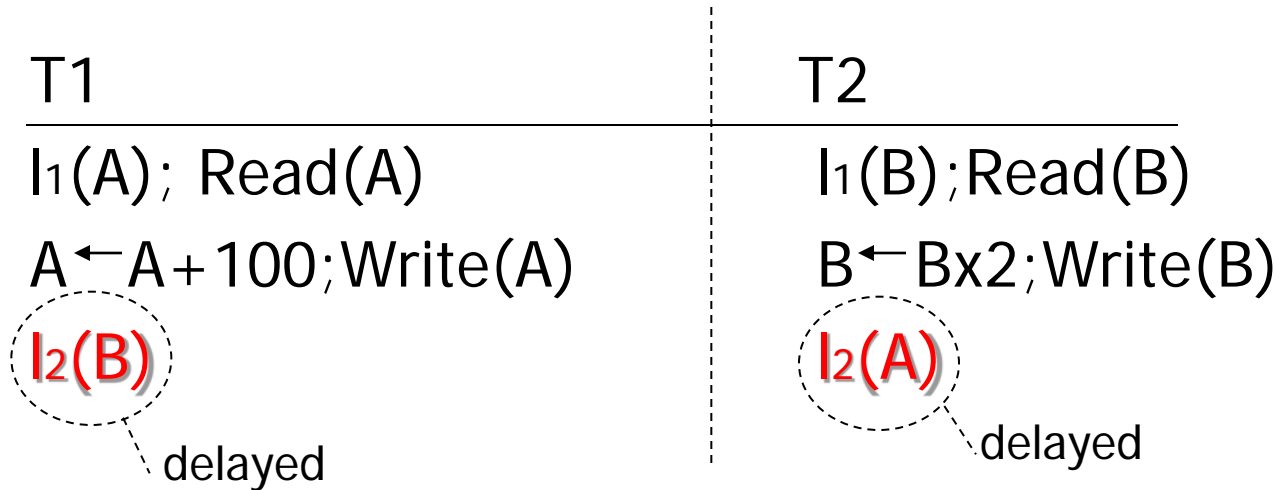


Schedule B

T1	T2
$I_1(A); \text{Read}(A)$	
$A \leftarrow A + 100; \text{Write}(A)$	
$I_1(B); u_1(A)$	
	$I_2(A); \text{Read}(A)$ delayed
	$A \leftarrow A \times 2; \text{Write}(A); I_2(B)$
$\text{Read}(B); B \leftarrow B + 100$	
$\text{Write}(B); u_1(B)$	
	$I_2(B); u_2(A); \text{Read}(B)$
	$B \leftarrow B \times 2; \text{Write}(B); u_2(B);$

A	B
25	25
125	
250	
	125
	250
250	250

Schedule C (T₂ reversed)





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

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

To help in proof:

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

Lemma

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

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:

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



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



- 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

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



- What about transactions that read and write same object?

Option 1: Request exclusive lock

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



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 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

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


no $I-X_j(A)$

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


no $I-X_j(A)$

no $I-S_j(A)$

A way to summarize Rule #2

Compatibility matrix

		New request	
		S	X
Lock already held in	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: $IN_i(A)$
 $\{Read(A); A \leftarrow A+k; Write(A)\}$
- $IN_i(A), IN_j(A)$ do not conflict!







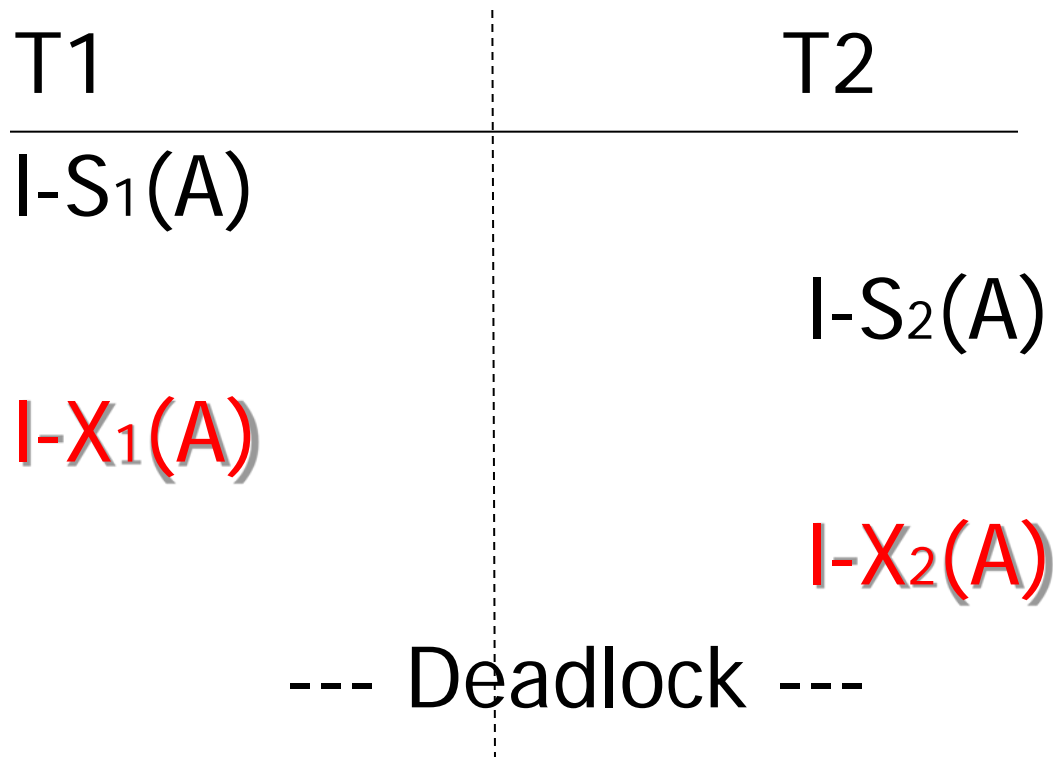
Comp

Lock already held in

		New request		
		S	X	I
S	S	T	F	F
	X	F	F	F
	I	F	F	T

Update locks

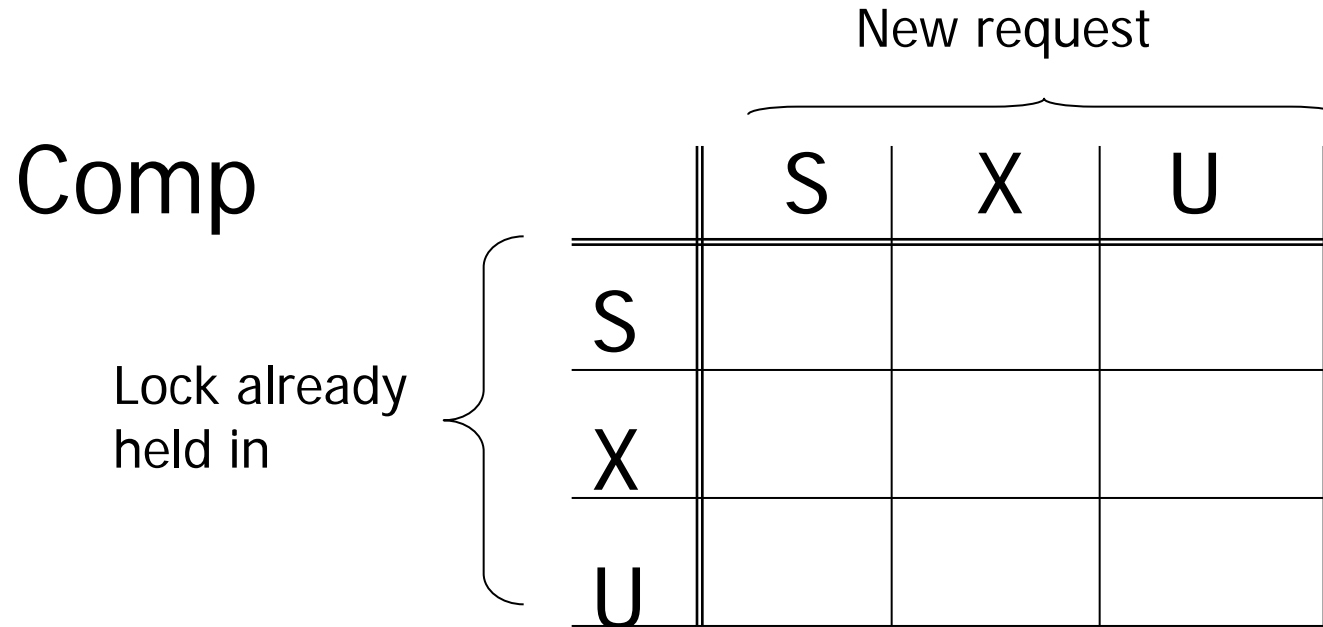
A common deadlock problem with upgrades:





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	T	F	T
X	F	F	F
U	TorF	F	T



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




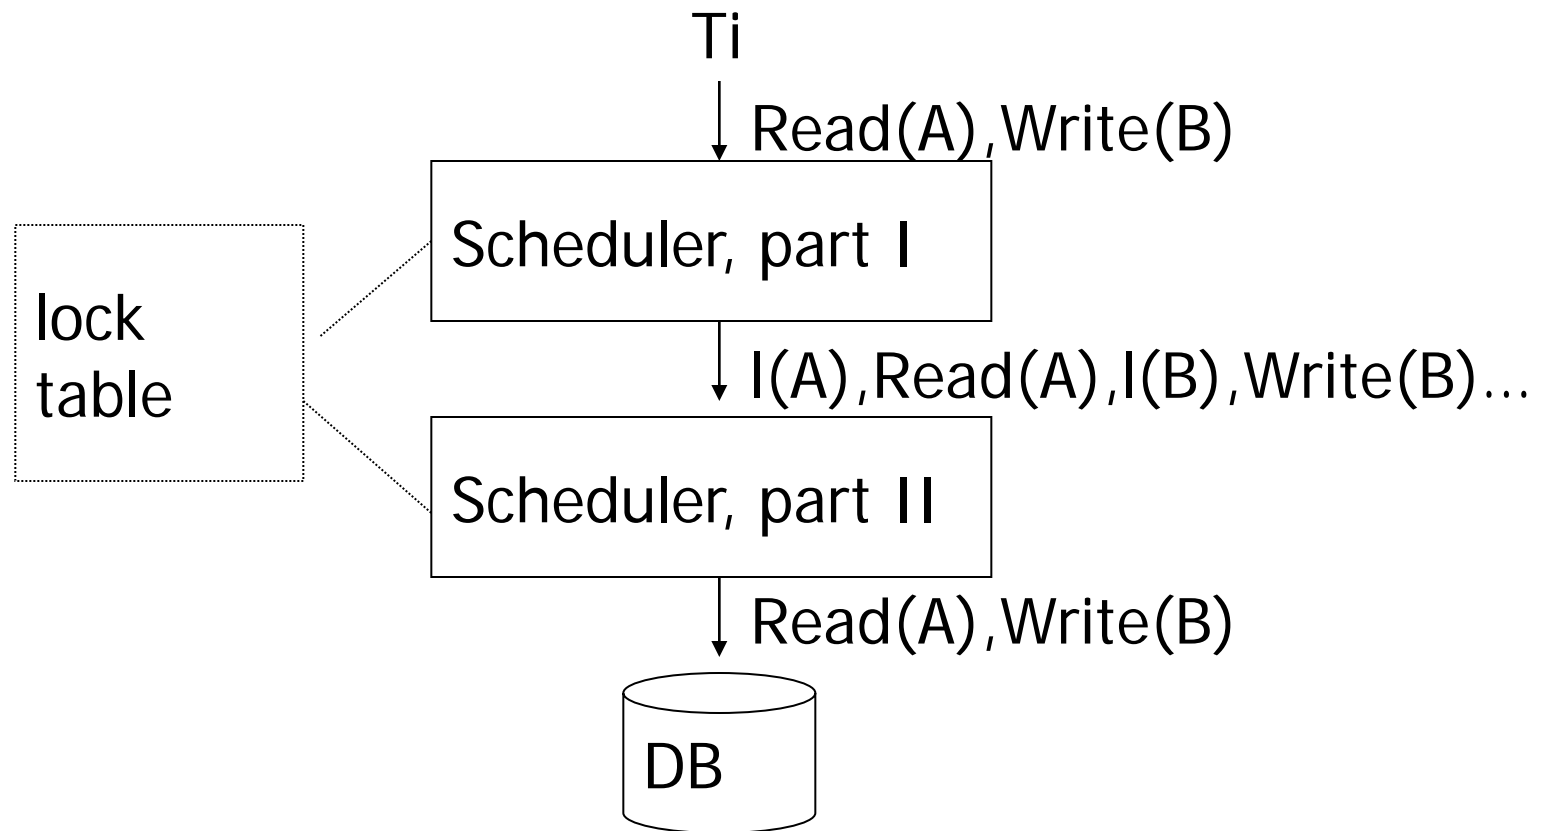
How does locking work in practice?

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

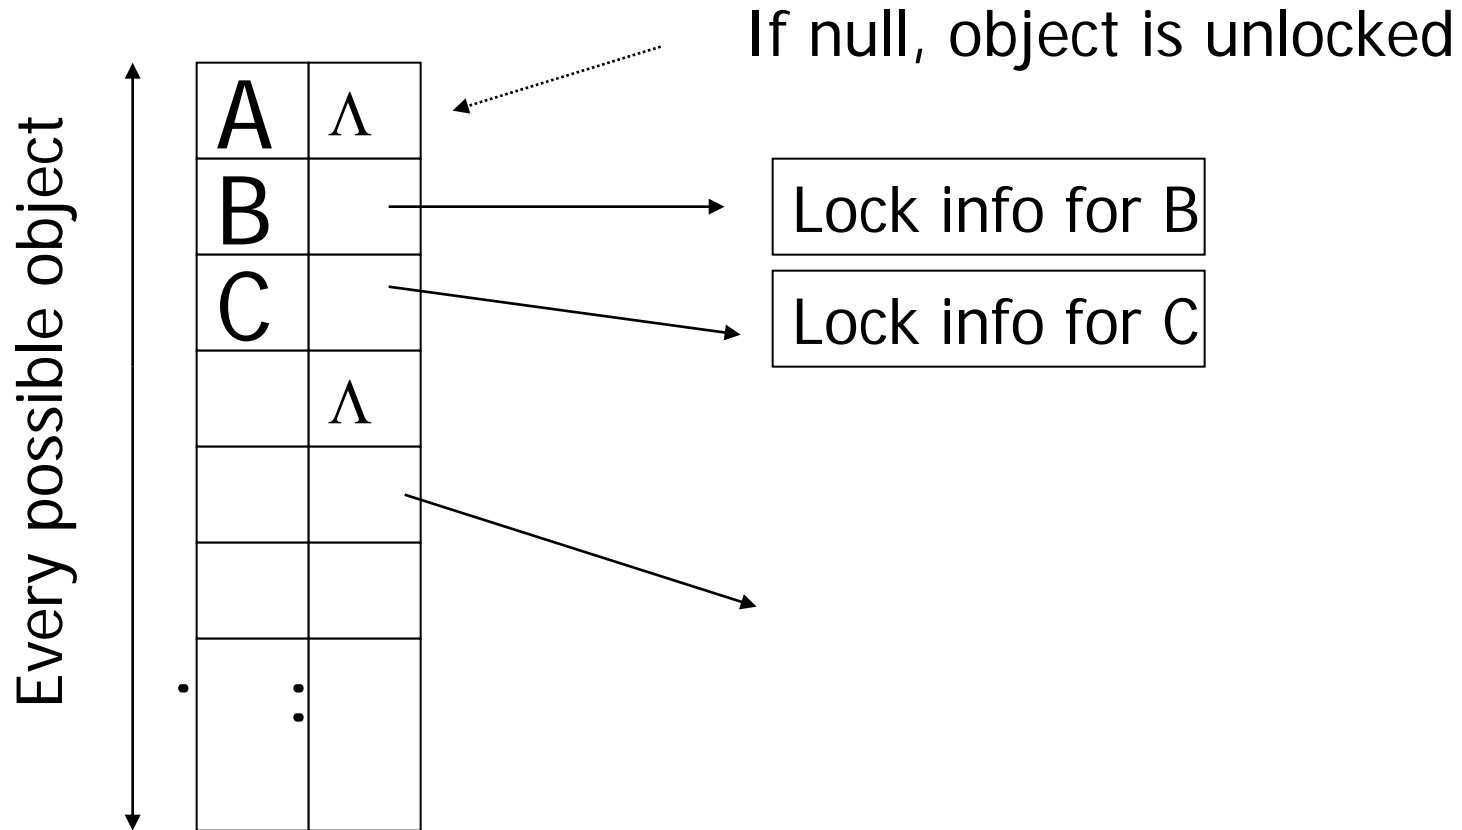
Sample Locking System:

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






Lock table Conceptually

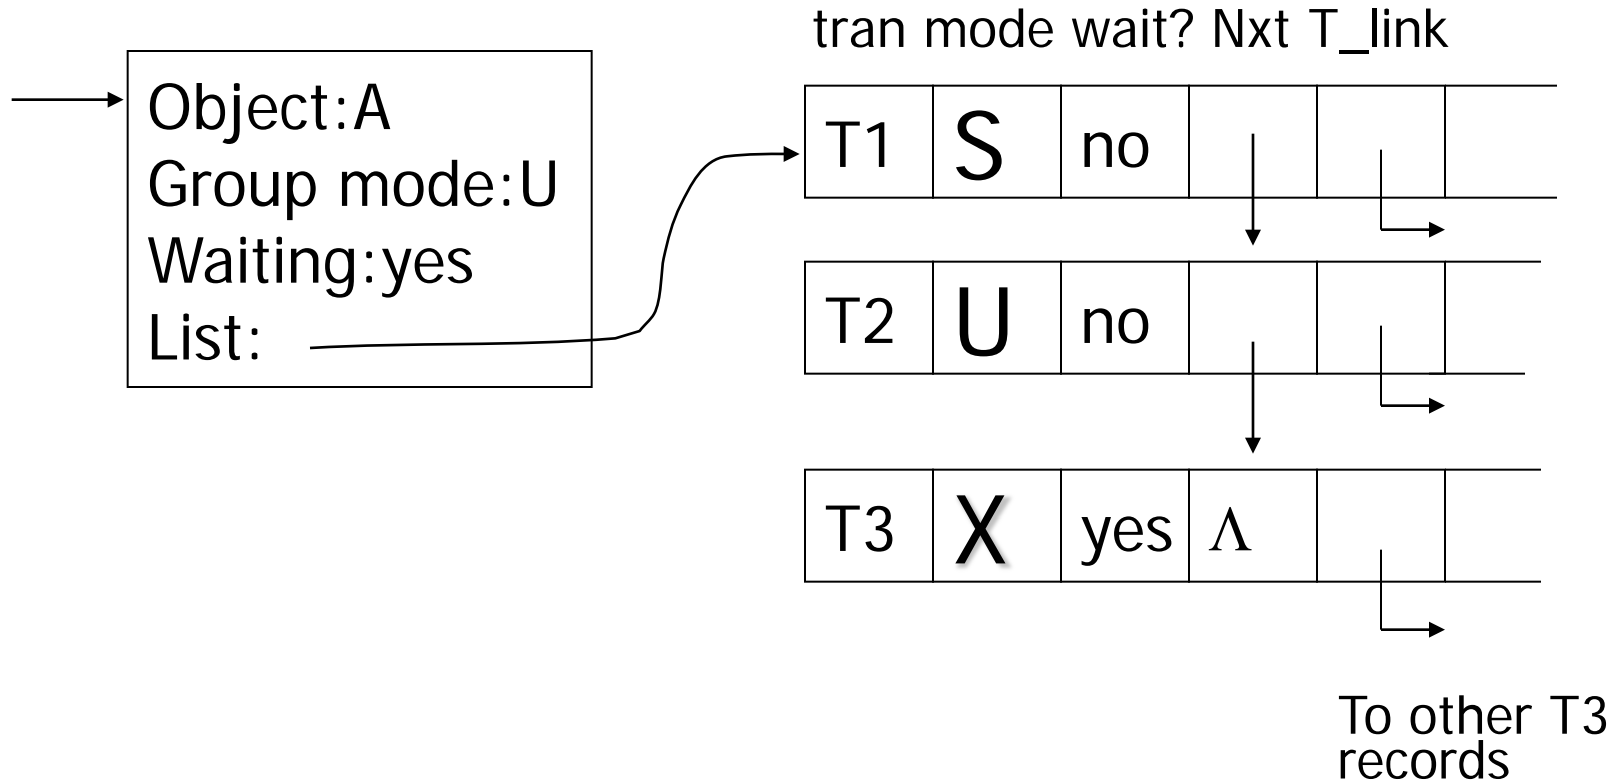


But use hash table:




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?
- 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