

# CS216: Data-Intensive Computing Systems

## **Concurrency Control (II)**


Shivnath Babu

# 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):  $l_i(A)$

unlock:  $u_i(A)$



# Rule #1: Well-formed transactions

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

## 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 =  $l_1(A)r_1(A)w_1(B)u_1(A)u_1(B)$

$l_2(B)r_2(B)w_2(B)l_3(B)r_3(B)u_3(B)$

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

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

$A \leftarrow A \times 2; \text{Write}(A); u_2(A)$

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


$B \leftarrow B \times 2; \text{Write}(B); u_2(B)$

# Schedule F


		A	B
T1	T2	25	25
<hr/>			
$l_1(A); \text{Read}(A)$			
$A \leftarrow A + 100; \text{Write}(A); u_1(A)$		125	
	$l_2(A); \text{Read}(A)$		
	$A \leftarrow Ax2; \text{Write}(A); u_2(A)$	250	
	$l_2(B); \text{Read}(B)$		
	$B \leftarrow Bx2; \text{Write}(B); u_2(B)$		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

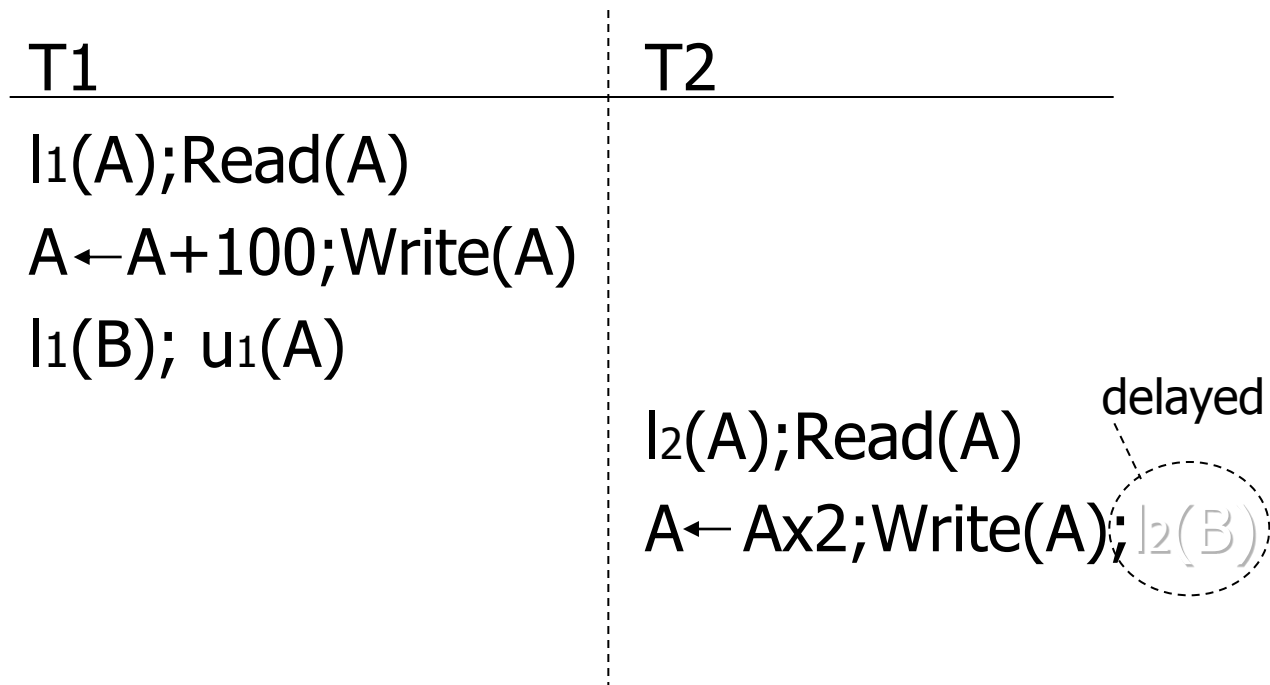


no locks

# locks held by Ti



# 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

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

$A \leftarrow A \times 2; \text{Write}(A); l_2(B)$

delayed



# 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

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


$A \leftarrow A \times 2; \text{Write}(A); l_2(B)$

delayed

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


$B \leftarrow B \times 2; \text{Write}(B); u_2(B);$

# Schedule H (T<sub>2</sub> 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

To help in proof:

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

## Lemma

$T_i \rightarrow T_j$  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

- 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