

# CS 44800: Introduction To Relational Database Systems

## *Locking*

Prof. Chris Clifton  
4 November 2021




## 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




11/9/2021

3

## A locking protocol

- Two new actions:
  - lock (exclusive):  $l_i(A)$
  - unlock:  $u_i(A)$



11/9/2021

4

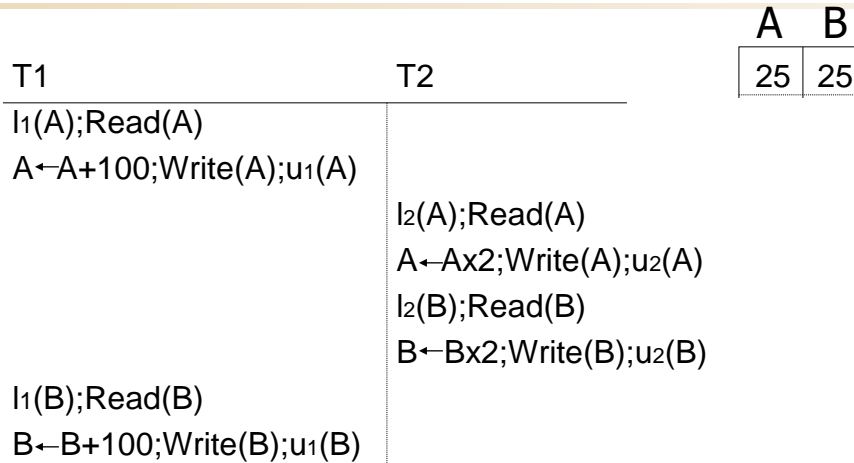
## Rules for Locking

- Rule #1: Well-formed transactions
  - $T_i: \dots l_i(A) \dots p_i(A) \dots u_i(A) \dots$   
( $p_i$  is a read or write)
- Rule #2: Legal scheduler
  - $S = \dots l_i(A) \xleftrightarrow{\hspace{2cm}} u_i(A) \dots$   
no  $l_j(A)$

## Exercise:

- What schedules are legal?  
What transactions are well-formed?
- $S_1 = l_1(A)l_1(B)r_1(A)w_1(B)l_2(B)u_1(A)u_1(B)r_2(B)w_2(B)u_2(B)l_3(B)r_3(B)u_3(B)$
- $S_2 = l_1(A)r_1(A)w_1(B)u_1(A)u_1(B)l_2(B)r_2(B)w_2(B)l_3(B)r_3(B)u_3(B)$
- $S_3 = l_1(A)r_1(A)u_1(A)l_1(B)w_1(B)u_1(B)l_2(B)r_2(B)w_2(B)u_2(B)l_3(B)r_3(B)u_3(B)$


# Schedule F



11/9/2021

7


# Rule #3: Two phase locking (2PL)



11/9/2021

9

## Two phase locking (2PL)




11/9/2021

10




### The Two-Phase Locking Protocol


- A protocol which ensures conflict-serializable schedules.
- Phase 1: **Growing Phase**
  - Transaction may obtain locks
  - Transaction may not release locks
- Phase 2: **Shrinking Phase**
  - Transaction may release locks
  - Transaction may not obtain locks
- The protocol assures serializability. It can be proved that the transactions can be serialized in the order of their **lock points** (i.e., the point where a transaction acquired its final lock).



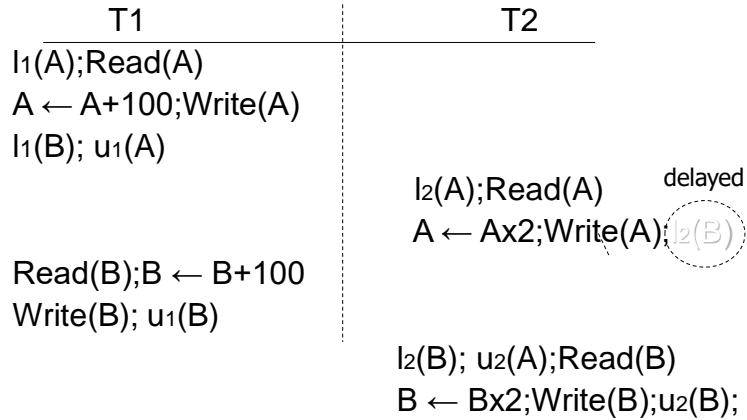
# Schedule G



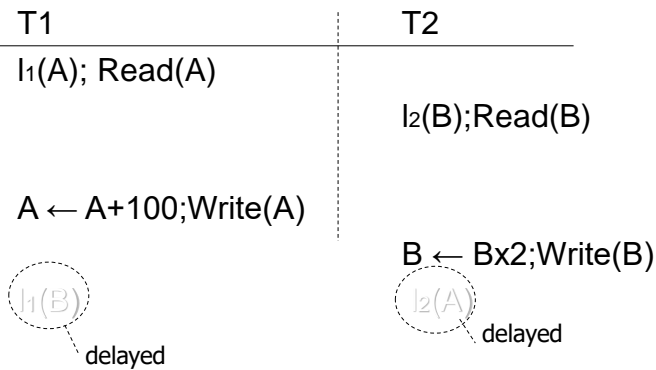
# Schedule G



## Schedule G



## Schedule H (T2 reversed)





## The Two-Phase Locking Protocol (Cont.)

- Two-phase locking is not a necessary condition for serializability
  - There are conflict serializable schedules that cannot be obtained if the two-phase locking protocol is used.
- In the absence of extra information (e.g., ordering of access to data), two-phase locking is necessary for conflict serializability *in the following sense*:
  - Given a transaction  $T_i$  that does not follow two-phase locking, we can find a transaction  $T_j$  that uses two-phase locking, and a schedule for  $T_i$  and  $T_j$  that is not conflict serializable.

$T_1$	$T_2$
lock-X(B)	
read(B)	
$B := B - 50$	
write(B)	
unlock(B)	lock-S(A)
	read(A)
	unlock(A)
	lock-S(B)
	read(B)
	unlock(B)
	display(A + B)
lock-X(A)	
read(A)	
$A := A + 50$	
write(A)	
unlock(A)	



PURDUE  
UNIVERSITY

Department of Computer Science

## Schedule H: Deadlock

- 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!





## Deadlock: Another Example

- Consider the partial schedule

$T_3$	$T_4$
lock-X( $B$ )	
read( $B$ )	
$B := B - 50$	
write( $B$ )	
	lock-S( $A$ )
	read( $A$ )
	lock-S( $B$ )
lock-X( $A$ )	

- Neither  $T_3$  nor  $T_4$  can make progress — executing **lock-S( $B$ )** causes  $T_4$  to wait for  $T_3$  to release its lock on  $B$ , while executing **lock-X( $A$ )** causes  $T_3$  to wait for  $T_4$  to release its lock on  $A$ .
- Such a situation is called a **deadlock**.
  - To handle a deadlock one of  $T_3$  or  $T_4$  must be rolled back and its locks released.



## Deadlock (Cont.)

- The potential for deadlock exists in most locking protocols. Deadlocks are a necessary evil.
- Starvation** is also possible if concurrency control manager is badly designed. For example:
  - A transaction may be waiting for an X-lock on an item, while a sequence of other transactions request and are granted an S-lock on the same item.
  - The same transaction is repeatedly rolled back due to deadlocks.
- Concurrency control manager can be designed to prevent starvation.



## The Two-Phase Locking Protocol (Cont.)

- Two-phase locking *does not* ensure freedom from deadlocks
- Extensions to basic two-phase locking needed to ensure recoverability of freedom from cascading roll-back
  - **Strict two-phase locking:** a transaction must hold all its exclusive locks till it commits/aborts.
    - Ensures recoverability and avoids cascading roll-backs
  - **Rigorous two-phase locking:** a transaction must hold *all* locks till commit/abort.
    - Transactions can be serialized in the order in which they commit.
- Most databases implement rigorous two-phase locking, *but refer to it as simply two-phase locking*



PURDUE  
UNIVERSITY

Department of Computer Science

## 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 Shrink( $T_i$ ) = SH( $T_i$ ) =  
 first unlock action of  $T_i$

11/9/2021

22

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:  $\leftarrow$  SH( $T_i$ )      SH( $T_j$ )  $\rightarrow$

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

11/9/2021

23

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

Proof:

(1) Assume P(S) has cycle

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

(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

11/9/2021

24



Department of Computer Science

- 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

11/9/2021

25