

# Transaction Management

## Concurrency Control (2)

# Conflict Actions

- A pair of consecutive actions in a schedule constitutes a **conflict** if swapping these actions may change the effect of at least one of the transactions involved.
- Most pairs of actions do not cause a conflict.
- $r_i(X)$  and  $r_j(Y)$  never cause a conflict, even if  $X = Y$ , since they do not modify the DB state.
- $r_i(X)$  and  $w_j(Y)$  do not cause a conflict if  $X \neq Y$ .
- $w_i(X)$  and  $r_j(Y)$  do not cause a conflict if  $X \neq Y$ .
- $w_i(X)$  and  $w_j(Y)$  do not cause a conflict if  $X \neq Y$ .

# Conflict Actions (cont.)

- The following three situations do cause a conflict:
- **Actions of the same transaction**, i.e.  $i = j$ .
- Two writes of the same database element by different transactions, i.e.  **$w_i(X)$  and  $w_j(X)$** ,  $i \neq j$ .  
Depending on the schedule, the results of either  $w_i(X)$  or  $w_j(X)$  survive, which may be different.
- A read and a write of the same database element by different transactions, i.e.  **$r_i(X)$  and  $w_j(X)$** ,  $i \neq j$ .  $r_i(X)$  may read a different version of  $X$ .

# Conflict Equivalent/Serializable


- 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.

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

# Review: Schedule C

$S(C) = r_1(A)w_1(A)r_2(A)w_2(A)r_1(B)w_1(B)r_2(B)w_2(B)$



$S(C)' = r_1(A)w_1(A) r_1(B)w_1(B)r_2(A)w_2(A)r_2(B)w_2(B)$

**T<sub>1</sub>**

**T<sub>2</sub>**

# Conflict-Serializability


- If transactions  $T_i$  and  $T_j$  contain at least two pairs of conflicting actions, then for each of these pairs the action of  $T_i$  has to be performed before that of  $T_j$  (or always  $T_j$  before  $T_i$ ).
- Given a schedule  $S$ ,  $T_i$  *takes precedence over*  $T_j$ , denoted by  $T_i <_S T_j$ , if there are actions  $p_i(A)$  of  $T_i$  and  $q_j(A)$  of  $T_j$  such that
  - $p_i(A)$  is ahead of  $q_j(A)$  in  $S$ ,
  - both  $p_i(A)$  and  $q_j(A)$  involve the same database element, and at least one of them is a **write**.

# Conflict-Serializability

- If  $T_i$  takes precedence over  $T_j$ , then a schedule  $S'$  that is conflict equivalent to  $S$  must have  $p_i(A)$  before  $q_j(A)$ .
- *Precedence graph*: directed graph with *nodes* representing the transactions of  $S$ ,  
*edges* representing precedence relationships,  
i.e. edge from node  $T_i$  to  $T_j$  if  $T_i <_S T_j$ .
- Notation:  $P(S)$

# Examples (1)

- What is  $P(S)$  for  
 $S=w3(A)w2(C)r1(A)w1(B)r1(C)w2(A)r4(A)w4(D)$




**Is  $S$  serializable?**



# Examples (2)

- What is  $P(S)$  for  
 $S=r_1(A)w_1(B)r_1(C)w_2(C)w_2(A)w_3(A)r_4(A)w_4(D)$



**S serial  $\rightarrow$  P(S) acyclic!**

# Conflict-Serializability

- Lemma 1

$S1, S2$  conflict equivalent  $\Rightarrow P(S1) = P(S2)$

- Proof

Assume  $P(S1) \neq P(S2)$

$\Rightarrow \exists Ti, Tj: Ti \rightarrow Tj$  in  $P(S1)$  and not in  $P(S2)$

$\Rightarrow S1 = \dots pi(A) \dots qj(A) \dots$   $\left\{ \begin{array}{l} pi, qj \\ \text{in conflict} \end{array} \right.$   
 $S2 = \dots qj(A) \dots pi(A) \dots$

$\Rightarrow S1, S2$  not conflict equivalent

# Conflict-Serializability

- Note

$P(S1)=P(S2) \not\Rightarrow S1, S2$  conflict equivalent

- Counter example

$S1 = w1(A) r2(A) w2(B) r1(B)$

$S2 = r2(A) w1(A) r1(B) w2(B)$

$P(S1)=P(S2) = T1 \rightleftarrows T2$

$S1$  not conflict equivalent to  $S2$ , since  $w1(A)$  and  $r2(A)$  cannot be swapped

# Conflict-Serializability

- Theorem 2

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

- Proof

(i)  $\Leftarrow$

Assume  $S$  is conflict serializable.

$\Rightarrow \exists S'$ :  $S'$  is serial,  $S$  conflict equivalent to  $S'$ .

$\Rightarrow P(S') = P(S)$  according to Lemma 1.

$P(S')$  is acyclic because  $S'$  is serial.

$\Rightarrow P(S)$  is acyclic.


# Conflict-Serializability

## ■ Proof

(ii)  $\Rightarrow$

Assume  $P(S)$  is acyclic.

Transform  $S$  as follows:



(1) Take  $T_1$  to be transaction with no incoming edges.

$T_1$  exists, since  $P(S)$  is acyclic.

(2) Move all  $T_1$  actions to the front:

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

This does not create any conflicts, since there is no  $T_j$  with  $T_j \rightarrow T_1$ .

(3) We now have  $S' = \langle T_1 \text{ actions} \rangle \langle \dots \text{rest} \dots \rangle$ .

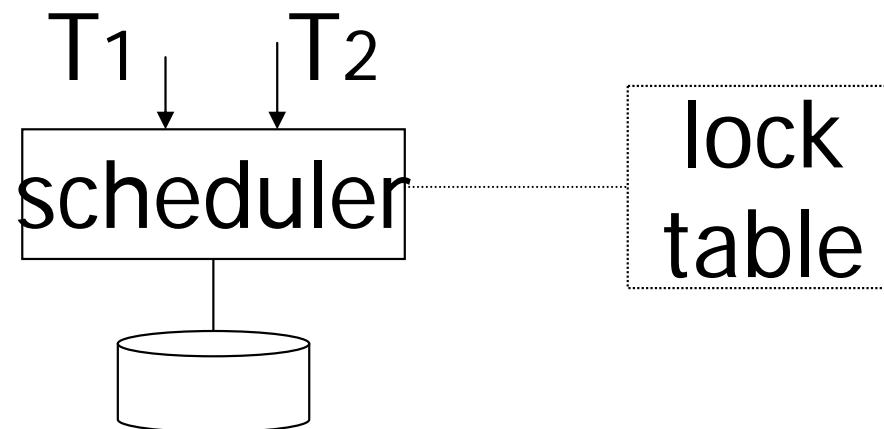
(4) Repeat above steps to serialize rest.

# Conflict-Serializability

- How to enforce that only conflict-serializable schedules are executed?
- There are two alternative approaches:
  - *Pessimistic concurrency control*  
Lock data elements to prevent P(S) cycles from occurring.
  - *Optimistic concurrency control*  
Detect P(S) cycles and undo participating transactions, if necessary.

# Enforcing Serializability by Locks

- Before accessing a database element, a transaction requests a *lock* on that element in order to prevent other transactions from accessing the same database element at the “same” time.
- Typically, different types of locks are used for different types of access operations, but we first introduce a simplified lock protocol with only one type of lock.



# Enforcing Serializability by Locks

- We introduce two new actions:
  - $l_i(X)$ : *lock* database element  $X$
  - $u_i(X)$ : *unlock* database element  $X$ , i.e. release lock.
- A locking protocol must guarantee the *consistency of transactions*:
  - A transaction can only read or write database  $X$  element if it currently holds a lock on  $X$ .
  - A transaction must unlock all database elements that it has locked at some later time.
- A consistent transaction is also called *well-formed*.

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



# Enforcing Serializability by Locks

- A locking protocol must also guarantee the *legality of schedules*:

At most one transaction can hold a lock on database element  $X$  at a given point of time.

- If there are actions  $l_i(X)$  followed by  $l_j(X)$  in some schedule, then there must be an action  $u_i(X)$  somewhere between these two actions.

$S = \dots \dots \dots l_i(A) \xrightarrow{\dots \dots \dots} u_i(A) \dots \dots \dots$   
no  $l_j(A)$

# Enforcing Serializability by Locks

## ■ Example

S1 = l1(A)l1(B)r1(A)w1(B)l2(B)u1(A)u1(B)

r2(B)w2(B)u2(B)l3(B)r3(B)u3(B)

→ S1 illegal, because T2 locks B before T1 has unlocked it

S2 = l1(A)r1(A)w1(B)u1(A)u1(B)

l2(B)r2(B)w2(B)l3(B)r3(B)u3(B)

→ T1 inconsistent, because T1 writes B before locking it

S3 = l1(A)r1(A)u1(A)l1(B)w1(B)u1(B)

l2(B)r2(B)w2(B)u2(B)l3(B)r3(B)u3(B)

→ schedule legal and all transactions consistent

# To-Do-List

- Do a research on how the currency control and logging recovery are related.