

# Concurrency Control

Instructor: Matei Zaharia

[cs245.stanford.edu](https://cs245.stanford.edu)

# Outline

What makes a schedule serializable?

Conflict serializability

Precedence graphs

**Enforcing serializability via 2-phase locking**

- » Shared and exclusive locks
- » Lock tables and multi-level locking

Optimistic concurrency with validation

Concurrency control + recovery

# Lock Modes Beyond S/X

Examples:

(1) increment lock

(2) update lock

# Example 1: Increment Lock

Atomic addition action:  $IN_i(A)$

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

$IN_i(A)$ ,  $IN_j(A)$  do not conflict, because addition is commutative!

# Compatibility Matrix

compat

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

# Update Locks

A common deadlock problem with upgrades:

T1	T2
I-S <sub>1</sub> (A)	
	I-S <sub>2</sub> (A)
I-X <sub>1</sub> (A)	
	I-X <sub>2</sub> (A)

--- Deadlock ---

# Solution

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

# Compatibility Matrix

compat

New request

Lock already held in

	S	X	U
S	T	F	
X	F	F	
U			

The diagram illustrates a compatibility matrix for locks. The vertical axis represents the 'Lock already held in' and the horizontal axis represents the 'New request'. Both axes have three states: S (Shared), X (Exclusive), and U (Update). The matrix cells contain compatibility results: 'T' (True) for (S, S), and 'F' (False) for (S, X), (X, S), and (X, X). All other combinations (S, U), (U, S), (U, X), and (U, U) are empty.



# Compatibility Matrix

compat

New request

Lock already held in

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

Note: asymmetric table!


# How Is Locking Implemented 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 ask transactions to request/release locks: just get the weakest lock for each action they perform
2. Hold all locks until transaction commits




# Sample Locking System

Under the hood: lock manager that keeps track of which objects are locked


» E.g. hash table

Also need a good way to block transactions until locks are available, and find deadlocks


# Which Objects Do We Lock?



DB



DB



DB

# Which Objects Do We Lock?

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

# Which Objects Do We Lock?

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 2



# Example 2



# Example 3



# Multiple Granularity Locks

compat

Requestor

IS IX S SIX X

Holder

	IS	IX	S	SIX	X
IS					
IX					
S					
SIX					
X					

# Multiple Granularity Locks

compat


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



# Rules Within A Transaction

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




# Multi-Granularity 2PL Rules

1. Follow multi-granularity compat 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  $T_2$  access object  $f_{2.2}$  in X mode? What locks will  $T_2$  get?




# Exercise:

Can  $T_2$  access object  $f_{2.2}$  in X mode? What locks will  $T_2$  get?




# Exercise:

Can  $T_2$  access object  $f3.1$  in X mode? What locks will  $T_2$  get?




# Exercise:

Can  $T_2$  access object  $f_{2.2}$  in S mode? What locks will  $T_2$  get?

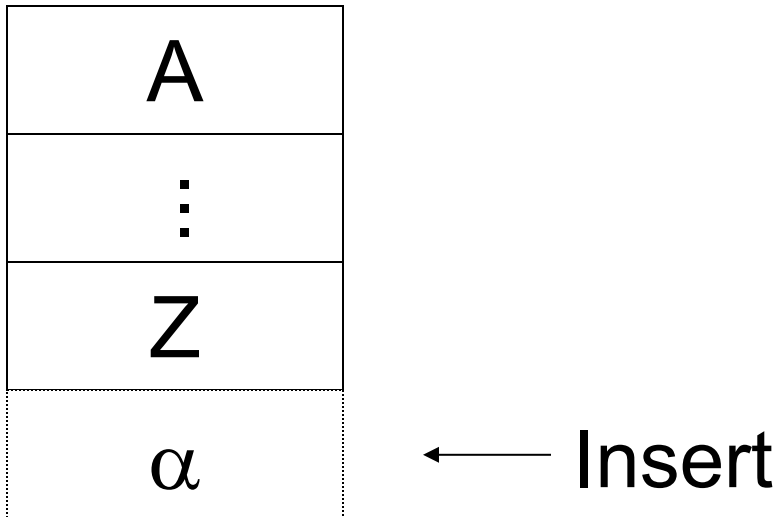


# Exercise:

Can  $T_2$  access object  $f_{2.2}$  in X mode? What locks will  $T_2$  get?



# Insert + Delete Operations





# Changes 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 Problem: Phantoms

Example: relation R (id, name,...)  
constraint: id is unique key  
use tuple locking

R            id   Name            ....

o1	55	Smith	
o2	75	Jones	


**T1: Insert <12,Mary,...> into R**  
**T2: Insert <12,Sam,...> into R**

T1	T2
S1(o1)	S2(o1)
S1(o2)	S2(o2)
Check Constraint	Check Constraint
⋮	⋮
Insert o3[12,Mary,..]	Insert o4[12,Sam,..]

# Solution

Use multiple granularity tree

Before insert of node N,  
lock parent(N) in X mode



# Back to Example

T1: Insert<12,Mary>

T1

T2: Insert<12,Sam>

T2

X1(R)

X2(R) ← delayed

Check constraint

Insert<12,Mary>

U1(R)


X2(R)

Check constraint

Oops! e# = 12 already in R!

# Instead of Using R, Can Use Index Nodes for Ranges

Example:



# Outline

What makes a schedule serializable?

Conflict serializability

Precedence graphs

Enforcing serializability via 2-phase locking

- » Shared and exclusive locks
- » Lock tables and multi-level locking

Optimistic concurrency with validation

Concurrency control + recovery

# Validation Approach

Transactions have 3 phases:

## 1. Read

- » Read all DB values needed
- » Write to temporary storage
- » No locking

## 2. Validate

- » Check whether 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 the validation order, then resulting schedule will be conflict equivalent to  $S_s = T_1, T_2, T_3, \dots$

# Implementing Validation


System keeps track of two sets:

**FIN** = transactions that have finished phase 3  
(write phase) and are all done

**VAL** = transactions that have successfully  
finished phase 2 (validation)


# Example That Validation Must Prevent:

$$\begin{array}{l} \text{RS}(T2)=\{B\} \quad \text{RS}(T3)=\{A,B\} \neq \emptyset \\ \text{WS}(T2)=\{B,D\} \quad \text{WS}(T3)=\{C\} \end{array}$$



# Example That Validation Must ~~Allow~~ Prevent:


$$\begin{aligned} RS(T2) &= \{B\} & RS(T3) &= \{A, B\} \neq \emptyset \\ WS(T2) &= \{B, D\} & WS(T3) &= \{C\} \end{aligned}$$



# Another Thing Validation Must Prevent:

$RS(T_2) = \{A\}$        $RS(T_3) = \{A, B\}$

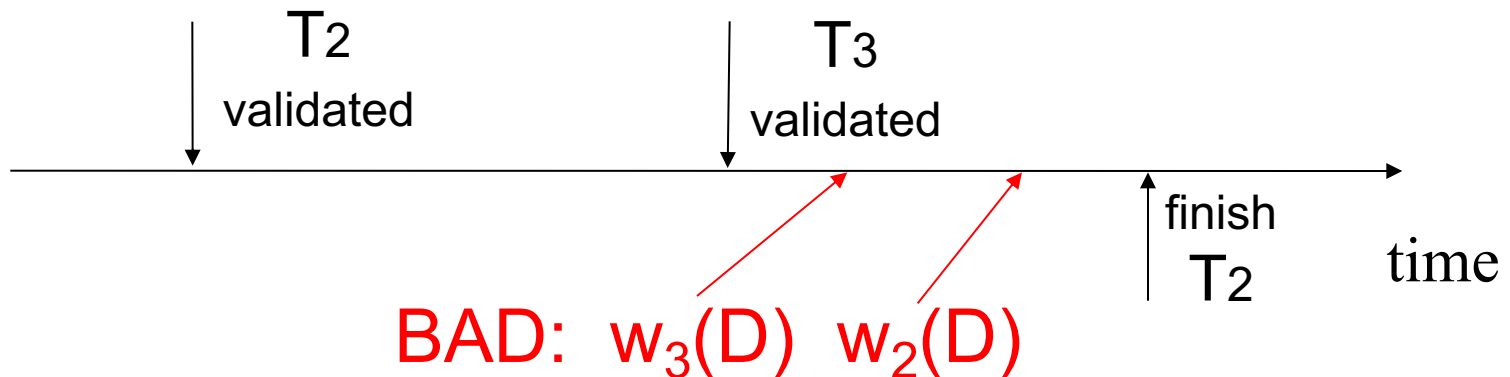
$WS(T_2) = \{D, E\}$        $WS(T_3) = \{C, D\}$



# Another Thing Validation Must Prevent:

$RS(T_2) = \{A\}$        $RS(T_3) = \{A, B\}$

$WS(T_2) = \{D, E\}$      $WS(T_3) = \{C, D\}$

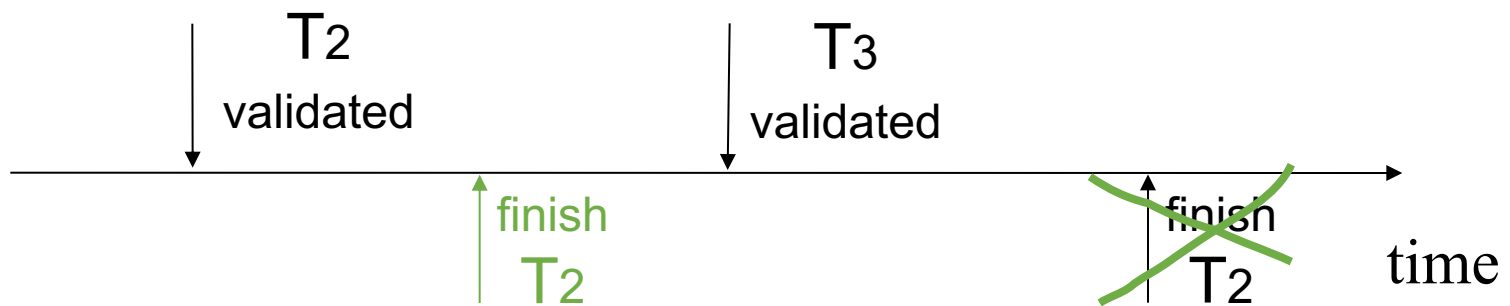


Allow

# Another Thing Validation Must Prevent:

$RS(T_2) = \{A\}$        $RS(T_3) = \{A, B\}$

$WS(T_2) = \{D, E\}$        $WS(T_3) = \{C, D\}$



# Validation Rules for Tj:

when Tj starts phase 1:

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

at Tj Validation:

if  $\text{Check}(T_j)$  then

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

do write phase

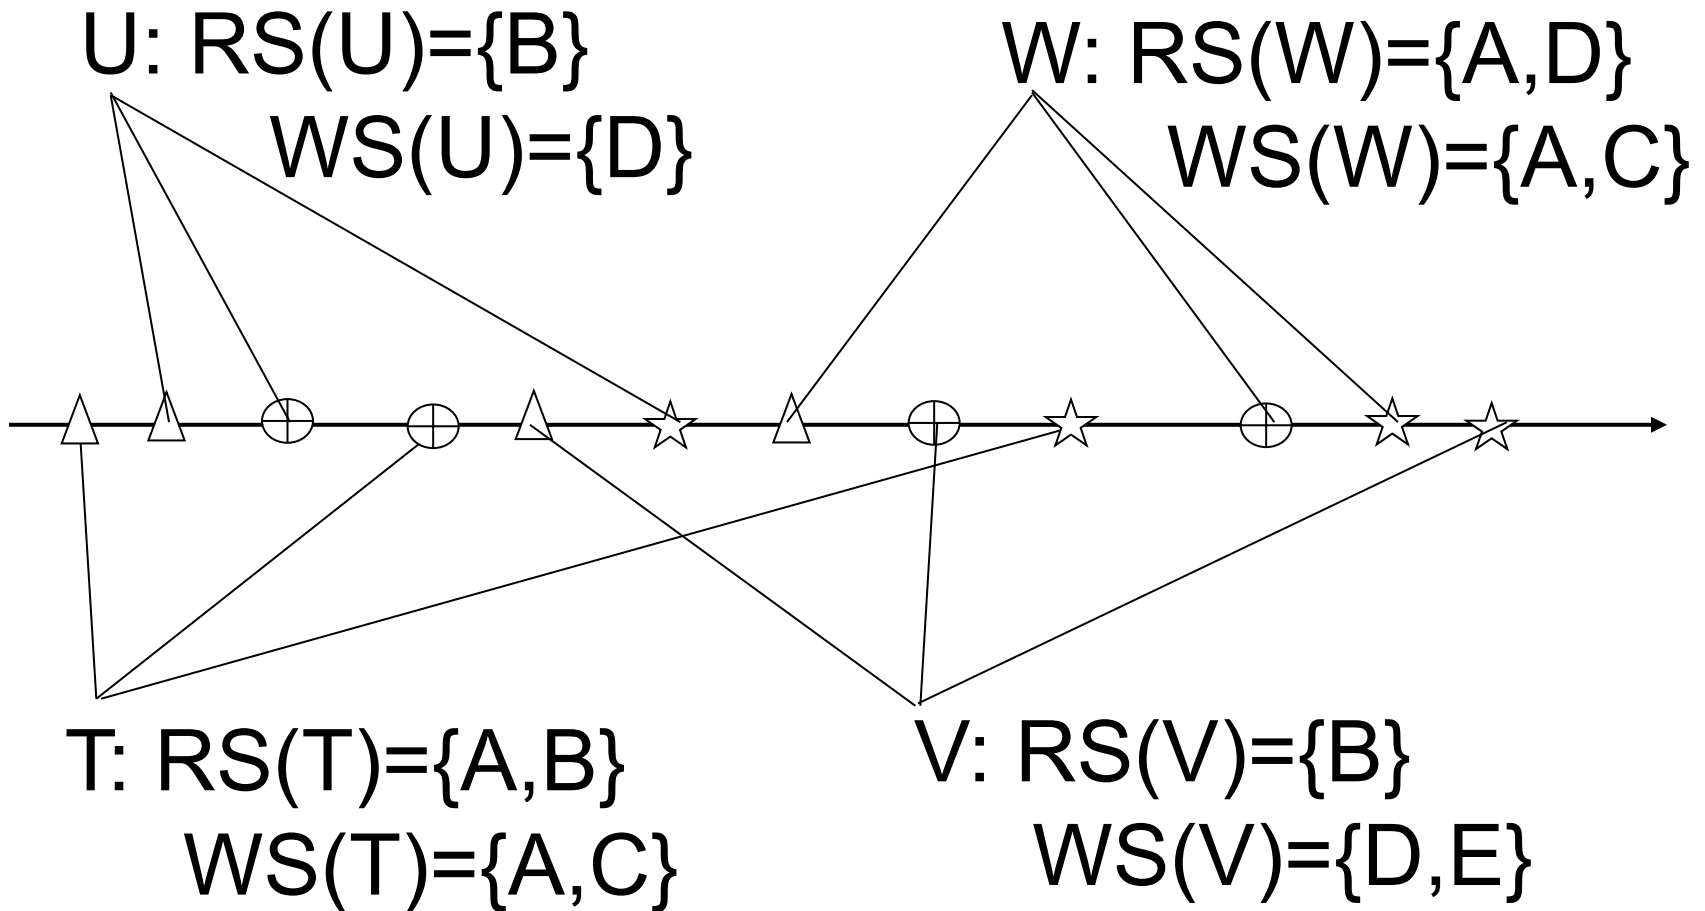

$\text{FIN} \leftarrow \text{FIN} \cup \{T_j\}$







# Check(Tj)

```
for Ti ∈ VAL – ignore(Tj) do
    if (WS(Ti) ∩ RS(Tj) ≠ ∅ or
        (Ti ∉ FIN and WS(Ti) ∩ WS(Tj) ≠ ∅))
        then return false
return true
```

# Exercise



# Is Validation = 2PL?



**S:  $w_2(y) w_1(x) w_2(x)$**

Achievable with 2PL?

Achievable with validation?

**S:  $w_2(y) w_1(x) w_2(x)$**

**S can be achieved with 2PL:**

$l_2(y) w_2(y) l_1(x) w_1(x) u_1(x) l_2(x) w_2(x) u_2(x) u_2(y)$

**S cannot be achieved by validation:**

The validation point of  $T_2$ ,  $val_2$ , must occur before  $w_2(y)$  since transactions do not write to the database until after validation. Because of the conflict on  $x$ ,  $val_1 < val_2$ , so we must have something like:

S:  $val_1 val_2 w_2(y) w_1(x) w_2(x)$

With the validation protocol, the writes of  $T_2$  should not start until  $T_1$  is all done with 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 slide)

# Validation Subset of 2PL?

Say S' not legal (due to w-r conflict):

S': ... l1(x) w2(x) r1(x) val1 u1(x) ...





- » 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': ... val1 l1(x) w2(x) w1(x) u1(x) ...

- » 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 AND  $WS(T1) \cap WS(T2) \neq \emptyset$
- » contradiction!

# Is Validation = 2PL?





# When to Use Validation?

Validation performs better than locking when:

- » Conflicts are rare
- » System resources are plentiful
- » Have tight latency constraints

# Summary

Have studied several concurrency control mechanisms used in practice

- » 2 PL
- » Multiple granularity
- » Validation

**Next:** how does concurrency control interact with failure recovery?

# Outline

What makes a schedule serializable?

Conflict serializability

Precedence graphs

Enforcing serializability via 2-phase locking

- » Shared and exclusive locks
- » Lock tables and multi-level locking

Optimistic concurrency with validation

Concurrency control + recovery