

# Database System Internals Concurrency Control - Locking

Paul G. Allen School of Computer Science and Engineering University of Washington, Seattle

February 10, 2020

#### Lab 2 due tonight

- Before final submission, clone fresh repo on attu and run "ant test-report"
- Lab 1+2 quiz on Wednesday in-class
  - Closed book. Calculator allowed but you won't need one.
- 544M Paper 2 due next week

## Conflicts

# Write-Read – WR

- Read-Write RW
- Write-Write WW

# **Conflict Serializability**

### Conflicts:

Two actions by same transaction  $T_i$ :  $r_i(X)$ 

$$r_i(X); w_i(Y)$$

Two writes by  $T_i$ ,  $T_j$  to same element



Read/write by T<sub>i</sub>, T<sub>i</sub> to same element





#### **Definition** A schedule is <u>conflict serializable</u> if it can be transformed into a serial schedule by a series of swappings of adjacent non-conflicting actions

- Every conflict-serializable schedule is serializable
- The converse is not true in general

### Testing for Conflict-Serializability

#### **Precedence graph:**

- A node for each transaction T<sub>i</sub>,
- An edge from T<sub>i</sub> to T<sub>j</sub> whenever an action in T<sub>i</sub> conflicts with, and comes before an action in T<sub>i</sub>
- No edge for actions in the same transaction
- The schedule is serializable iff the precedence graph is acyclic

### Testing for Conflict-Serializability

Important:

Always draw the full graph, unless ONLY asked if (yes or no) the schedule is conflict serializable



CSE 444 - Winter 2020

 $r_2(A)$ ;  $r_1(B)$ ;  $w_2(A)$ ;  $r_3(A)$ ;  $w_1(B)$ ;  $w_3(A)$ ;  $r_2(B)$ ;  $w_2(B)$ 

Example 1







1) (2) (3)

February 10, 2020

#### February 10, 2020

CSE 444 - Winter 2020

 $r_2(A) || r_1(B)$ 

 $r_2(A)$ ;  $r_1(B)$ ;  $w_2(A)$ ;  $r_3(A)$ ;  $w_1(B)$ ;  $w_3(A)$ ;  $r_2(B)$ ;  $w_2(B)$ 

2

3

10

 $r_2(A)$  $r_1(B)$ 

No edge because no conflict (A != B)

 $r_2(A)$ ;  $r_1(B)$ ;  $w_2(A)$ ;  $r_3(A)$ ;  $w_1(B)$ ;  $w_3(A)$ ;  $r_2(B)$ ;  $w_2(B)$ 

 $) \qquad (2) \qquad (3)$ 

CSE 444 - Winter 2020

 $r_2(A) || w_2(A)$ 

 $r_2(A)$ ;  $r_1(B)$ ;  $w_2(A)$ ;  $r_3(A)$ ;  $w_1(B)$ ;  $w_3(A)$ ;  $r_2(B)$ ;  $w_2(B)$ 

2

3

12

$$r_2(A)$$
  $w_2(A)$ 

No edge because same txn (2)

 $r_2(A)$ ;  $r_1(B)$ ;  $w_2(A)$ ;  $r_3(A)$ ;  $w_1(B)$ ;  $w_3(A)$ ;  $r_2(B)$ ;  $w_2(B)$ 



$$[r_{2}(A)] r_{3}(A) ?$$

$$r_{2}(A); r_{1}(B); w_{2}(A); r_{3}(A); w_{1}(B); w_{3}(A); r_{2}(B); w_{2}(B)$$



$$r_2(A)$$
  $w_1(B)$  ?  
 $r_2(A); r_1(B); w_2(A); r_3(A); w_1(B); w_3(A); r_2(B); w_2(B)$ 

1) (2) (3)

$$r_2(A)$$
  $w_3(A)$  ?  
 $r_2(A)$ ;  $r_1(B)$ ;  $w_2(A)$ ;  $r_3(A)$ ;  $w_1(B)$ ;  $w_3(A)$ ;  $r_2(B)$ ;  $w_2(B)$ 

1) (2) (3)

$$\begin{array}{c|c} r_2(A) & w_3(A) & {}_{T2 \text{ to } T3} \\ \hline r_2(A); & r_1(B); & w_2(A); & r_3(A); & w_1(B); \\ \hline w_3(A); & r_2(B); & w_2(B) \\ \hline \end{array}$$

1 F

$$\begin{array}{c|c} r_2(A) & w_3(A) & {}^{Edge! \ Conflict \ from} \\ r_2(A); \ r_1(B); \ w_2(A); \ r_3(A); \ w_1(B); \ w_3(A); \ r_2(B); \ w_2(B) \end{array}$$



$$r_2(A)$$
  $r_2(B)$  ?

 $r_2(A); r_1(B); w_2(A); r_3(A); w_1(B); w_3(A); r_2(B) w_2(B)$ 

#### And so on until compared every pair of actions... $1 \qquad (2) \qquad (3)$







More edges, but repeats of the same directed edge not necessary







#### This schedule is **conflict-serializable**



### r<sub>2</sub>(A); r<sub>1</sub>(B); w<sub>2</sub>(A); r<sub>2</sub>(B); r<sub>3</sub>(A); w<sub>1</sub>(B); w<sub>3</sub>(A); w<sub>2</sub>(B)











#### This schedule is NOT conflict-serializable

 A serializable schedule need not be conflict serializable, even under the "worst case update" assumption

$$W_1(X); W_2(X); W_2(Y); W_1(Y); W_3(Y);$$

Is this schedule conflict-serializable ?

 A serializable schedule need not be conflict serializable, even under the "worst case update" assumption

$$W_1(X); W_2(X); W_2(Y); W_1(Y); W_3(Y);$$

Is this schedule conflict-serializable ?



 A serializable schedule need not be conflict serializable, even under the "worst case update" assumption





#### Serializable, but not conflict serializable

February 10, 2020

Two schedules S, S' are *view equivalent* if:

- If T reads an initial value of A in S, then T reads the initial value of A in S'
- If T reads a value of A written by T' in S, then T reads a value of A written by T' in S'
- If T writes the final value of A in S, then T writes the final value of A in S'

# A schedule is view serializable if it is view equivalent to a serial schedule

Remark:

- If a schedule is conflict serializable, then it is also view serializable
- But not vice versa

### Schedules with Aborted Transactions

- When a transaction aborts, the recovery manager undoes its updates
- But some of its updates may have affected other transactions !

#### Schedules with Aborted Transactions



### Schedules with Aborted Transactions



#### Cannot abort T1 because cannot undo T2

February 10, 2020

A schedule is *recoverable* if:

- It is conflict-serializable, and
- Whenever a transaction T commits, all transactions that have written elements read by T have already committed

A schedule is *recoverable* if:

- It is conflict-serializable, and
- Whenever a transaction T commits, all transactions that have written elements read by T have already committed


#### **Recoverable Schedules**



February 10, 2020

### **Cascading Aborts**

- If a transaction T aborts, then we need to abort any other transaction T' that has read an element written by T
- A schedule avoids cascading aborts if whenever a transaction reads an element, the transaction that has last written it has already committed.

#### We base our locking scheme on this rule!

### **Avoiding Cascading Aborts**



With cascading aborts

Without cascading aborts

### Serializability

#### Recoverability

- Serial
- Serializable
- Conflict serializable
- View serializable

- Recoverable
- Avoids cascading deletes

- The scheduler:
- Module that schedules the transaction's actions, ensuring serializability
- Two main approaches
- Pessimistic: locks
- Optimistic: timestamps, multi-version, validation

### **Pessimistic Scheduler**

Simple idea:

- Each element has a unique lock
- Each transaction must first acquire the lock before reading/writing that element
- If the lock is taken by another transaction, then wait
- The transaction must release the lock(s)



 $L_i(A)$  = transaction  $T_i$  acquires lock for element A  $U_i(A)$  = transaction  $T_i$  releases lock for element A

#### A Non-Serializable Schedule





T1 T2 L<sub>1</sub>(A); READ(A, t) t := t+100 WRITE(A, t); U<sub>1</sub>(A);  $L_2(A)$ ; READ(A,s) s := s\*2 WRITE(A,s);  $U_2(A)$ ;  $L_2(B)$ ; READ(B,s) s := s\*2 WRITE(B,s); U<sub>2</sub>(B); L<sub>1</sub>(B); READ(B, t)

t := t+100 WRITE(B,t); U<sub>1</sub>(B);

Locks did not enforce conflict-serializability !!! What's wrong ?

February 10, 2020

## Two Phase Locking (2PL)

The 2PL rule:

- In every transaction, all lock requests must precede all unlock requests
- This ensures conflict serializability ! (will prove this shortly)

#### Example: 2PL transactions

T1  $L_1(A); L_1(B); READ(A, t)$  t := t+100WRITE(A, t); U<sub>1</sub>(A)

> L<sub>2</sub>(A); READ(A,s) s := s\*2 WRITE(A,s); L<sub>2</sub>(B); DENIED...

READ(B, t) t := t+100 WRITE(B,t); U<sub>1</sub>(B);

> ...GRANTED; READ(B,s) s := s\*2 WRITE(B,s); U<sub>2</sub>(A); U<sub>2</sub>(B);

Now it is conflict-serializable

February 10, 2020

CSE 444 - Winter 2020

T2

# Example with Multiple Transactions



# Equivalent to each transaction executing entirely the moment it enters shrinking phase

### Two Phase Locking (2PL)

#### **Theorem**: 2PL ensures conflict serializability

**Proof**. Suppose not: then there exists a cycle in the precedence graph.



**Proof**. Suppose not: then there exists a cycle in the precedence graph.

Then there is the following <u>temporal</u> cycle in the schedule:



**Proof**. Suppose not: then there exists a cycle in the precedence graph.



Then there is the following <u>temporal</u> cycle in the schedule:  $U_1(A) \rightarrow L_2(A)$  why?

**Proof**. Suppose not: then there exists a cycle in the precedence graph.



Then there is the following <u>temporal</u> cycle in the schedule:  $U_1(A) \rightarrow L_2(A)$  $L_2(A) \rightarrow U_2(B)$  why?

**Proof**. Suppose not: then there exists a cycle in the precedence graph.



Then there is the following temporal cycle in the schedule:  $U_1(A) \rightarrow L_2(A)$  $L_2(A) \rightarrow U_2(B)$  $U_2(B) \rightarrow L_3(B)$ \_<sub>3</sub>(B)→U<sub>3</sub>(C)  $U_3(C) \rightarrow L_1(C)$  $C) \rightarrow U_1(A)$  Contradiction

#### A New Problem:

T1  $L_1(A); L_1(B); READ(A, t)$  t := t+100WRITE(A, t); U<sub>1</sub>(A)

L<sub>2</sub>(A); READ(A,s) s := s\*2 WRITE(A,s); L<sub>2</sub>(B); DENIED...

T2

READ(B, t) t := t+100 WRITE(B,t); U<sub>1</sub>(B);

> ...GRANTED; READ(B,s) s := s\*2 WRITE(B,s); U<sub>2</sub>(A); U<sub>2</sub>(B); Commit

#### **Abort**

# Strict 2PL

- Strict 2PL: All locks held by a transaction are released when the transaction is completed; release happens at the time of COMMIT or ROLLBACK
- Schedule is recoverable
- Schedule avoids cascading aborts

T1		T2	
L <sub>1</sub> (A); READ(A)			
A :=A+100			
WRITE(A);			
		L <sub>2</sub> (A); DENIED	
L <sub>1</sub> (B); READ(B)			
B :=B+100			
WRITE(B);			
U <sub>1</sub> (A),U <sub>1</sub> (B); Rollba	ack		
		GRANTED; READ(A)	
		A := A*2	
		WRITE(A):	
		$L_2(B)$ : READ(B)	
		B := B*2	
		WRITE(B);	
C	SE 444 - Winter 2020	$U_2(A); U_2(B); Commit$	
ary 10, 2020			

# Summary of Strict 2PL

- Ensures serializability, recoverability, and avoids cascading aborts
- Issues?

# Summary of Strict 2PL

- Ensures serializability, recoverability, and avoids cascading aborts
- Issues: implementation, lock modes, granularity, deadlocks, performance

Task 1: -- act on behalf of the transaction

Add lock/unlock requests to transactions

- Examine all READ(A) or WRITE(A) actions
- Add appropriate lock requests
- On COMMIT/ROLLBACK release all locks
- Ensures Strict 2PL !

### The Locking Scheduler

- Task 2: -- act on behalf of the system Execute the locks accordingly
- Lock table: a big, critical data structure in a DBMS !
- When a lock is requested, check the lock table
  - Grant, or add the transaction to the element's wait list
- When a lock is released, re-activate a transaction from its wait list
- When a transaction aborts, release all its locks
- Check for deadlocks occasionally

#### S = shared lock (for READ)

X = exclusive lock (for WRITE)

Lock compatibility ma	trix:		
	None	S	X
None	OK	OK	OK
S	OK	ОК	Conflict
X	OK	Conflict	Conflict

### Lock Granularity

Fine granularity locking (e.g., tuples)

Coarse grain locking (e.g., tables, predicate locks)

February 10, 2020

•

•

### Lock Granularity

#### Fine granularity locking (e.g., tuples)

- High concurrency
- High overhead in managing locks
- Coarse grain locking (e.g., tables, predicate locks)

•

•

### Lock Granularity

#### Fine granularity locking (e.g., tuples)

- High concurrency
- High overhead in managing locks
- Coarse grain locking (e.g., tables, predicate locks)
  - Many false conflicts
  - Less overhead in managing locks

# Deadlocks

#### • Cycle in the wait-for graph:

- T1 waits for T2
- T2 waits for T3
- T3 waits for T1
- Deadlock detection
  - Timeouts
  - Wait-for graph
- Deadlock avoidance
  - Acquire locks in pre-defined order
  - Acquire all locks at once before starting



#### # Active Transactions

- So far we have assumed the database to be a static collection of elements (=tuples)
- If tuples are inserted/deleted then the *phantom* problem appears

### **Phantom Problem**

#### T1

T2

#### SELECT \* FROM Product WHERE color='blue'

INSERT INTO Product(name, color) VALUES ('gizmo','blue')

SELECT \* FROM Product WHERE color='blue'

Is this schedule serializable?

#### T1

T2

#### SELECT \* FROM Product WHERE color='blue'

INSERT INTO Product(name, color) VALUES ('gizmo','blue')

#### SELECT \* FROM Product WHERE color='blue'

Suppose there are two blue products, X1, X2:

## R1(X1),R1(X2),W2(X3),R1(X1),R1(X2),R1(X3)

#### T1

T2

#### SELECT \* FROM Product WHERE color='blue'

INSERT INTO Product(name, color) VALUES ('gizmo','blue')

#### SELECT \* FROM Product WHERE color='blue'

Suppose there are two blue products, X1, X2:

## R1(X1),R1(X2),W2(X3),R1(X1),R1(X2),R1(X3)

This is conflict serializable ! What's wrong ??
#### T1

T2

#### SELECT \* FROM Product WHERE color='blue'

INSERT INTO Product(name, color) VALUES ('gizmo','blue')

#### SELECT \* FROM Product WHERE color='blue'

Suppose there are two blue products, X1, X2:

## R1(X1),R1(X2),W2(X3),R1(X1),R1(X2),R1(X3)

Not serializable due to *phantoms* 

CSE 444 - Winter 2020

# Phantom Problem

- A "phantom" is a tuple that is invisible during part of a transaction execution but not invisible during the entire execution
- In our example:
  - T1: reads list of products
  - T2: inserts a new product
  - T1: re-reads: a new product appears !

- In a <u>static</u> database:
  - Conflict serializability implies serializability
- In a <u>dynamic</u> database, this may fail due to phantoms
- Strict 2PL guarantees conflict serializability, but not serializability

# **Dealing With Phantoms**

- Lock the entire table, or
- Lock the index entry for 'blue'
  - If index is available
- Or use predicate locks
  - A lock on an arbitrary predicate

## Dealing with phantoms is expensive !

## **Isolation Levels in SQL**

- 1. "Dirty reads" SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED
- 2. "Committed reads" SET TRANSACTION ISOLATION LEVEL READ COMMITTED
- 3. "Repeatable reads" SET TRANSACTION ISOLATION LEVEL REPEATABLE READ
- 4. Serializable transactions SET TRANSACTION ISOLATION LEVEL SERIALIZABLE (



## 1. Isolation Level: Dirty Reads

#### "Long duration" WRITE locks

- Strict 2PL
- No READ locks
  - Read-only transactions are never delayed

### Possible pbs: dirty and inconsistent reads

#### "Long duration" WRITE locks

- Strict 2PL
- "Short duration" READ locks
  - Only acquire lock while reading (not 2PL)

Unrepeatable reads When reading same element twice, may get two different values

#### "Long duration" WRITE locks

• Strict 2PL

### "Long duration" READ locks

• Strict 2PL

### This is not serializable yet !!!



# 4. Isolation Level Serializable

- "Long duration" WRITE locks
  - Strict 2PL
- "Long duration" READ locks
  - Strict 2PL
- Predicate locking
  - To deal with phantoms

Client 1: START TRANSACTION **INSERT INTO SmallProduct(name, price) SELECT** pname, price **FROM** Product WHERE price <= 0.99 **DELETE FROM Product** WHERE price <= 0.99 COMMIT Client 2: SET TRANSACTION READ ONLY START TRANSACTION **SELECT** count(\*) **FROM** Product **SELECT** count(\*) **FROM** SmallProduct

COMMIT

February 10, 2020

May improve

performance

# **Commercial Systems**

### Always check documentation!

- DB2: Strict 2PL
- SQL Server:
  - Strict 2PL for standard 4 levels of isolation
  - Multiversion concurrency control for snapshot isolation
- PostgreSQL: Snapshot isolation; recently: seralizable Snapshot isolation (!)
- Oracle: Snapshot isolation