

# Introduction to Database Systems CSE 414

## Lecture 24: Implementation of Transactions

CSE 414 - Spring 2018

1

## Conflict Serializability

- A schedule is *conflict serializable* 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 (why?)

CSE 414 - Spring 2018

2

## Testing for Conflict-Serializability

### Precedence graph:


- A node for each transaction  $T_i$ ,
- An edge from  $T_i$  to  $T_j$  whenever an action in  $T_i$  conflicts with, and comes before an action in  $T_j$
- **The schedule is conflict-serializable iff the precedence graph is acyclic**

CSE 414 - Spring 2018

3

## Example 2

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



This schedule is **NOT** conflict-serializable

CSE 414 - Spring 2018

4

## More Notations

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

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

CSE 414 - Spring 2018

5

## A Non-Serializable Schedule

<p>T1</p> <p>READ(A)</p> <p>A := A+100</p> <p>WRITE(A)</p> <p>READ(B)</p> <p>B := B+100</p> <p>WRITE(B)</p>	<p>T2</p> <p>READ(A)</p> <p>A := A*2</p> <p>WRITE(A)</p> <p>READ(B)</p> <p>B := B*2</p> <p>WRITE(B)</p>
--	---

CSE 414 - Spring 2018

6

### Example

<p>T1</p> <pre>L<sub>1</sub>(A); READ(A) A := A+100 WRITE(A); U<sub>1</sub>(A); L<sub>1</sub>(B)</pre> <pre>READ(B) B := B+100 WRITE(B); U<sub>1</sub>(B);</pre>	<p>T2</p> <pre>L<sub>2</sub>(A); READ(A) A := A*2 WRITE(A); U<sub>2</sub>(A); L<sub>2</sub>(B); <b>BLOCKED...</b></pre> <pre>...GRANTED; READ(B) B := B*2 WRITE(B); U<sub>2</sub>(B);</pre>
---	--

*time* ↓

Scheduler has ensured a conflict-serializable schedule

7

### But...

<p>T1</p> <pre>L<sub>1</sub>(A); READ(A) A := A+100 WRITE(A); U<sub>1</sub>(A);</pre> <pre>L<sub>1</sub>(B); READ(B) B := B+100 WRITE(B); U<sub>1</sub>(B);</pre>	<p>T2</p> <pre>L<sub>2</sub>(A); READ(A) A := A*2 WRITE(A); U<sub>2</sub>(A); L<sub>2</sub>(B); READ(B) B := B*2 WRITE(B); U<sub>2</sub>(B);</pre>
--	--

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

8

## Two Phase Locking (2PL)

The 2PL rule:

In every transaction, all lock requests must precede all unlock requests

CSE 414 - Spring 2018 9

### Example: 2PL transactions

<p>T1</p> <pre>L<sub>1</sub>(A); L<sub>1</sub>(B); READ(A) A := A+100 WRITE(A); U<sub>1</sub>(A)</pre> <pre>READ(B) B := B+100 WRITE(B); U<sub>1</sub>(B);</pre>	<p>T2</p> <pre>L<sub>2</sub>(A); READ(A) A := A*2 WRITE(A); L<sub>2</sub>(B); <b>BLOCKED...</b></pre> <pre>...GRANTED; READ(B) B := B*2 WRITE(B); U<sub>2</sub>(A); U<sub>2</sub>(B);</pre>
---	--

Now it is conflict-serializable

10

## Two Phase Locking (2PL)

**Theorem: 2PL ensures conflict serializability**

11

## Two Phase Locking (2PL)

**Theorem: 2PL ensures conflict serializability**

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

```

graph TD
    T1((T1)) -- A --> T2((T2))
    T2 -- B --> T3((T3))
    T3 -- C --> T1
  
```

10

### Two Phase Locking (2PL)

**Theorem:** 2PL ensures conflict serializability

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

Then there is the following **temporal** cycle in the schedule:

13

### Two Phase Locking (2PL)

**Theorem:** 2PL ensures conflict serializability

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

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

$U_1(A)$  happened strictly *before*  $L_2(A)$

14

### Two Phase Locking (2PL)

**Theorem:** 2PL ensures conflict serializability

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

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

$L_2(A)$  happened strictly *before*  $U_1(A)$

15

### Two Phase Locking (2PL)

**Theorem:** 2PL ensures conflict serializability

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

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

16

### Two Phase Locking (2PL)

**Theorem:** 2PL ensures conflict serializability

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

Then there is the following **temporal** cycle in the schedule:  
 $U_1(A) \rightarrow L_2(A)$   
 $L_2(A) \rightarrow U_2(B)$   
 $U_2(B) \rightarrow L_3(B)$  why?

17

### Two Phase Locking (2PL)

**Theorem:** 2PL ensures conflict serializability

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

Then there is the following **temporal** cycle in the schedule:  
 $U_1(A) \rightarrow L_2(A)$   
 $L_2(A) \rightarrow U_2(B)$   
 $U_2(B) \rightarrow L_3(B)$   
 .....etc.....

18

## Two Phase Locking (2PL)

**Theorem:** 2PL ensures conflict serializability

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



Then there is the following **temporal** cycle in the schedule:

$U_1(A) \rightarrow L_2(A)$

$L_2(A) \rightarrow U_2(B)$

$U_2(B) \rightarrow L_3(B)$

$L_3(B) \rightarrow U_3(C)$

$U_3(C) \rightarrow L_1(C)$

$L_1(C) \rightarrow U_1(A)$

Cycle in time:  
Contradiction

## A New Problem: Non-recoverable Schedule

<p>T1 L<sub>1</sub>(A); L<sub>1</sub>(B); READ(A) A := A+100 WRITE(A); U<sub>1</sub>(A)</p> <p>READ(B) B := B+100 WRITE(B); U<sub>1</sub>(B);</p> <p>Rollback</p>	<p>T2</p> <hr/> <p>L<sub>2</sub>(A); READ(A) A := A*2 WRITE(A); L<sub>2</sub>(B); <b>BLOCKED...</b></p> <p>...GRANTED; READ(B) B := B*2 WRITE(B); U<sub>2</sub>(A); U<sub>2</sub>(B); Commit</p>
---	--

CSE 414 - Spring 2018

20

## A New Problem: Non-recoverable Schedule

<p>T1 L<sub>1</sub>(A); L<sub>1</sub>(B); READ(A) A := A+100 WRITE(A); U<sub>1</sub>(A)</p> <p>READ(B) B := B+100 WRITE(B); U<sub>1</sub>(B);</p> <p>Rollback</p>	<p>T2</p> <hr/> <p>L<sub>2</sub>(A); READ(A) A := A*2 WRITE(A); L<sub>2</sub>(B); <b>BLOCKED...</b></p> <p>...GRANTED; READ(B) B := B*2 WRITE(B); U<sub>2</sub>(A); U<sub>2</sub>(B); Commit</p>
---	--

Elements A, B written by T1 are restored to their original value.

Spring 2018

21

## A New Problem: Non-recoverable Schedule

<p>T1 L<sub>1</sub>(A); L<sub>1</sub>(B); READ(A) A := A+100 WRITE(A); U<sub>1</sub>(A)</p> <p>READ(B) B := B+100 WRITE(B); U<sub>1</sub>(B);</p> <p>Rollback</p>	<p>T2</p> <hr/> <p>L<sub>2</sub>(A); READ(A) A := A*2 WRITE(A); L<sub>2</sub>(B); <b>BLOCKED...</b></p> <p>...GRANTED; READ(B) B := B*2 WRITE(B); U<sub>2</sub>(A); U<sub>2</sub>(B); Commit</p>
---	--

Dirty reads of A, B lead to incorrect writes.

Elements A, B written by T1 are restored to their original value.

Spring 2018

22

## A New Problem: Non-recoverable Schedule

<p>T1 L<sub>1</sub>(A); L<sub>1</sub>(B); READ(A) A := A+100 WRITE(A); U<sub>1</sub>(A)</p> <p>READ(B) B := B+100 WRITE(B); U<sub>1</sub>(B);</p> <p>Rollback</p>	<p>T2</p> <hr/> <p>L<sub>2</sub>(A); READ(A) A := A*2 WRITE(A); L<sub>2</sub>(B); <b>BLOCKED...</b></p> <p>...GRANTED; READ(B) B := B*2 WRITE(B); U<sub>2</sub>(A); U<sub>2</sub>(B); Commit</p>
---	--

Dirty reads of A, B lead to incorrect writes.

Elements A, B written by T1 are restored to their original value.

Spring 2018

Can no longer undo!

## Strict 2PL

<p>T1 L<sub>1</sub>(A); READ(A) A := A+100 WRITE(A);</p> <p>L<sub>1</sub>(B); READ(B) B := B+100 WRITE(B); Rollback &amp; U<sub>1</sub>(A); U<sub>1</sub>(B);</p>	<p>T2</p> <hr/> <p>L<sub>2</sub>(A); <b>BLOCKED...</b></p> <p>...GRANTED; READ(A) A := A*2 WRITE(A); L<sub>2</sub>(B); READ(B) B := B*2 WRITE(B); Commit &amp; U<sub>2</sub>(A); U<sub>2</sub>(B);</p>
---	--

24

## Strict 2PL

The Strict 2PL rule:

All locks are held until commit/abort:  
All unlocks are done together with commit/abort.

With strict 2PL, we will get schedules that are both conflict-serializable and recoverable

CSE 414 - Spring 2018

25

## Another problem: Deadlocks

- $T_1$ : R(A), W(B)
- $T_2$ : R(B), W(A)
- $T_1$  holds the lock on A, waits for B
- $T_2$  holds the lock on B, waits for A

This is a deadlock!

CSE 414 - Spring 2018

26

## Another problem: Deadlocks

To detect a deadlocks, search for a cycle in the *waits-for graph*:

- $T_1$  waits for a lock held by  $T_2$ ;
- $T_2$  waits for a lock held by  $T_3$ ;
- ...
- $T_n$  waits for a lock held by  $T_1$

Relatively expensive: check periodically, if deadlock is found, then abort one transaction.  
need to continuously re-check for deadlocks

27

## A "Solution": Lock Modes

- S = shared lock (for READ)
- X = exclusive lock (for WRITE)

Lock compatibility matrix:

	None	S	X
None			
S			
X			

CSE 414 - Spring 2018

28

## A "Solution": Lock Modes

- S = shared lock (for READ)
- X = exclusive lock (for WRITE)

Lock compatibility matrix:

	None	S	X
None	✓	✓	✓
S	✓	✓	✗
X	✓	✗	✗

CSE 414 - Spring 2018

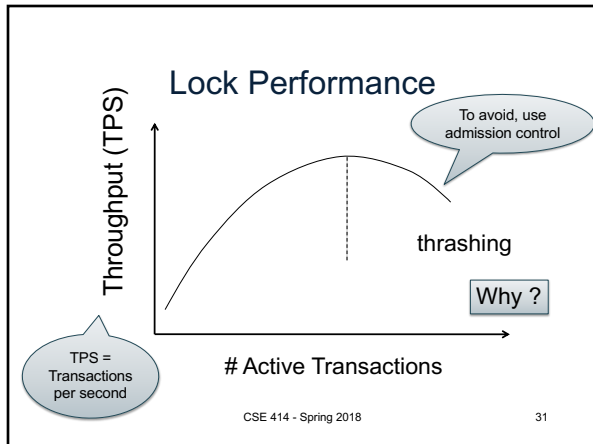
29

## Lock Granularity

- **Fine granularity locking** (e.g., tuples)
  - High concurrency
  - High overhead in managing locks
  - E.g., SQL Server
- **Coarse grain locking** (e.g., tables, entire database)
  - Many false conflicts
  - Less overhead in managing locks
  - E.g., SQL Lite
- **Solution: lock escalation changes granularity as needed**

CSE 414 - Spring 2018

30



### Phantom Problem

- 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

CSE 414 - Spring 2018 32

Suppose there are two blue products, A1, A2:

### Phantom Problem

T1	T2
SELECT * FROM Product WHERE color='blue'	INSERT INTO Product(name, color) VALUES ('A3','blue')
SELECT * FROM Product WHERE color='blue'	

Is this schedule serializable ?

CSE 414 - Spring 2018 33

Suppose there are two blue products, A1, A2:

### Phantom Problem

T1	T2
SELECT * FROM Product WHERE color='blue'	INSERT INTO Product(name, color) VALUES ('A3','blue')
SELECT * FROM Product WHERE color='blue'	

R<sub>1</sub>(A1);R<sub>1</sub>(A2);W<sub>2</sub>(A3);R<sub>1</sub>(A1);R<sub>1</sub>(A2);R<sub>1</sub>(A3)

CSE 414 - Spring 2018 34

Suppose there are two blue products, A1, A2:

### Phantom Problem

T1	T2
SELECT * FROM Product WHERE color='blue'	INSERT INTO Product(name, color) VALUES ('A3','blue')
SELECT * FROM Product WHERE color='blue'	

R<sub>1</sub>(A1);R<sub>1</sub>(A2);W<sub>2</sub>(A3);R<sub>1</sub>(A1);R<sub>1</sub>(A2);R<sub>1</sub>(A3)

W<sub>2</sub>(A3);R<sub>1</sub>(A1);R<sub>1</sub>(A2);R<sub>1</sub>(A1);R<sub>1</sub>(A2);R<sub>1</sub>(A3)

CSE 414 - Spring 2018 35

### 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 !

CSE 414 - Spring 2018 36

## Dealing With Phantoms

- Lock the entire table
- 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 !

CSE 414 - Spring 2018

37

## Summary of Serializability

- Serializable schedule = equivalent to a serial schedule
- (strict) 2PL guarantees *conflict serializability*
  - What is the difference?
- **Static database:**
  - *Conflict serializability* implies serializability
- **Dynamic database:**
  - This no longer holds

CSE 414 - Spring 2018

38

## 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`

ACID

CSE 414 - Spring 2018

39

## 1. Isolation Level: Dirty Reads

- "Long duration" WRITE locks
  - Strict 2PL
- No READ locks
  - Read-only transactions are never delayed

Possible problems: dirty and inconsistent reads

CSE 414 - Spring 2018

40

## 2. Isolation Level: Read Committed

- "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

CSE 414 - Spring 2018

41

## 3. Isolation Level: Repeatable Read

- "Long duration" WRITE locks
  - Strict 2PL
- "Long duration" READ locks
  - Strict 2PL

Why ?

This is not serializable yet !!!

CSE 414 - Spring 2018

42

## 4. Isolation Level Serializable

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

CSE 414 - Spring 2018

43

## Beware!

In commercial DBMSs:

- Default level is often NOT serializable
- Default level differs between DBMSs
- Some engines support subset of levels!
- Serializable may not be exactly ACID
  - Locking ensures isolation, not atomicity
- Also, some DBMSs do NOT use locking and different isolation levels can lead to different pbs
- Bottom line: RTFM for your DBMS!

CSE 414 - Spring 2018

44