

# Introduction to Database Systems CSE 414

## Lecture 28: Transactions Wrap-up

CSE 414 - Autumn 2018

1

# Announcements

- 2 late days for HW 8 are now free
  - No more than 2 late days allowed. Monday Dec. 10 is the hard cut off
- Office hours changes
  - Ryan tomorrow at 11am instead of 10:30
  - Andrew additional office hours Friday

CSE 414 - Autumn 2018

2

## A New Problem: Non-recoverable Schedule

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

CSE 414 - Autumn 2018

3

## A New Problem: Non-recoverable Schedule

T1	T2
L <sub>1</sub> (A); L <sub>1</sub> (B); READ(A) A := A+100 WRITE(A); U <sub>1</sub> (A)	L <sub>2</sub> (A); READ(A) A := A*2 WRITE(A); L <sub>2</sub> (B); <b>BLOCKED...</b>
READ(B) B := B+100 WRITE(B); U <sub>1</sub> (B);	... <b>GRANTED</b> ; READ(B) B := B*2 WRITE(B); U <sub>2</sub> (A); U <sub>2</sub> (B); Commit
Rollback	Elements A, B written by T1 are restored to their original value.

Autumn 2018

4

## A New Problem: Non-recoverable Schedule

T1	T2
L <sub>1</sub> (A); L <sub>1</sub> (B); READ(A) A := A+100 WRITE(A); U <sub>1</sub> (A)	L <sub>2</sub> (A); READ(A) A := A*2 WRITE(A); L <sub>2</sub> (B); <b>BLOCKED...</b>
READ(B) B := B+100 WRITE(B); U <sub>1</sub> (B);	... <b>GRANTED</b> ; READ(B) B := B*2 WRITE(B); U <sub>2</sub> (A); U <sub>2</sub> (B); Commit
Rollback	Dirty reads of A, B lead to incorrect writes.
Elements A, B written by T1 are restored to their original value.	

Autumn 2018

5

## A New Problem: Non-recoverable Schedule

T1	T2
L <sub>1</sub> (A); L <sub>1</sub> (B); READ(A) A := A+100 WRITE(A); U <sub>1</sub> (A)	L <sub>2</sub> (A); READ(A) A := A*2 WRITE(A); L <sub>2</sub> (B); <b>BLOCKED...</b>
READ(B) B := B+100 WRITE(B); U <sub>1</sub> (B);	... <b>GRANTED</b> ; READ(B) B := B*2 WRITE(B); U <sub>2</sub> (A); U <sub>2</sub> (B); Commit
Rollback	Dirty reads of A, B lead to incorrect writes.
Elements A, B written by T1 are restored to their original value.	Can no longer undo!

Autumn 2018

## 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 - Autumn 2018

7

## Strict 2PL

<p>T1</p> <p>L1(A); READ(A) A := A+100 WRITE(A);</p> <p>L1(B); READ(B) B := B+100 WRITE(B);</p> <p>Rollback &amp; U1(A); U1(B);</p>	<p>T2</p> <p>L2(A); BLOCKED...</p> <p>...GRANTED; READ(A) A := A*2 WRITE(A); L2(B); READ(B) B := B*2 WRITE(B);</p> <p>Commit &amp; U2(A); U2(B);</p>
---	--

8

## Strict 2PL

- Lock-based systems always use strict 2PL
- Easy to implement:
  - Before a transaction reads or writes an element A, insert an L(A)
  - When the transaction commits/aborts, then release all locks
- Ensures both conflict serializability and recoverability

CSE 414 - Autumn 2018

9

## Another problem: Deadlocks

- T<sub>1</sub>: R(A), W(B)
- T<sub>2</sub>: R(B), W(A)
- T<sub>1</sub> holds the lock on A, waits for B
- T<sub>2</sub> holds the lock on B, waits for A

This is a deadlock!


CSE 414 - Autumn 2018

10

## Another problem: Deadlocks

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

- T<sub>1</sub> waits for a lock held by T<sub>2</sub>;
- T<sub>2</sub> waits for a lock held by T<sub>3</sub>;
- . . .
- T<sub>n</sub> waits for a lock held by T<sub>1</sub>



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

11

## 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 - Autumn 2018

12

## A "Solution"?: Lock Modes

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

Lock compatibility matrix:


	None	S	X
X			

Can only fix deadlocks if transactions declare exclusive locks in advance.

CSE 414 - Autumn 2018

13

## Lock Performance



CSE 414 - Autumn 2018

14

## 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 - Autumn 2018

15

## 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 - Autumn 2018

16

Suppose there are two blue products, A1, A2:

## Phantom Problem

<p>T1</p> <hr/> <p>SELECT * FROM Product WHERE color='blue'</p> <p>SELECT * FROM Product WHERE color='blue'</p>	<p>T2</p> <hr/> <p>INSERT INTO Product(name, color) VALUES ('A3', 'blue')</p>
---	---

Is this schedule serializable ?

CSE 414 - Autumn 2018

17

Suppose there are two blue products, A1, A2:

## Phantom Problem

<p>T1</p> <hr/> <p>SELECT * FROM Product WHERE color='blue'</p> <p>SELECT * FROM Product WHERE color='blue'</p>	<p>T2</p> <hr/> <p>INSERT INTO Product(name, color) VALUES ('A3', 'blue')</p>
---	---

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 - Autumn 2018

18

Suppose there are two blue products, A1, A2:

### Phantom Problem

<p>T1</p> <pre>SELECT * FROM Product WHERE color='blue'</pre> <p>SELECT * FROM Product WHERE color='blue'</p>	<p>T2</p> <pre>INSERT INTO Product(name, color) VALUES ('A3','blue')</pre>
---	--

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)<sup>19</sup>

### 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 !

### 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 !

### 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

### Isolation Levels in SQL

For better performance

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

### 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

## 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 - Autumn 2018

25

## 3. Isolation Level: Repeatable Read

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

This is not serializable yet !!!

Why ?

CSE 414 - Autumn 2018

26

## 4. Isolation Level Serializable

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

CSE 414 - Autumn 2018

27

## 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 - Autumn 2018

28