

## Database Systems CSE 414

### Lecture 22: Transaction Implementations

CSE 414 - Spring 2017

1

## Announcements

- WQ7 (last!) due on Sunday
- HW7:
  - due on Wed, May 24
  - using JDBC to execute SQL from Java
  - using SQL Server via Azure

CSE 414 - Spring 2017

2

## Recap

- What are transactions?
  - Why do we need them?
- Maintain ACID properties via schedules
  - We focus on the **isolation** property
  - We briefly discussed **consistency & durability**
  - We do not discuss **atomicity**
- Ensure conflict-serializable schedules with locks

CSE 414 - Spring 2017

3

## Implementing a Scheduler

Major differences between database vendors

- **Locking Scheduler**
  - Aka “pessimistic concurrency control”
  - SQLite, SQL Server, DB2
- **Multiversion Concurrency Control (MVCC)**
  - Aka “optimistic concurrency control”
  - Postgres, Oracle

We discuss only locking in 414

CSE 414 - Spring 2017

4

## Locking Scheduler

Simple idea:

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

By using locks, scheduler **can** ensure conflict-serializability

CSE 414 - Spring 2017

5

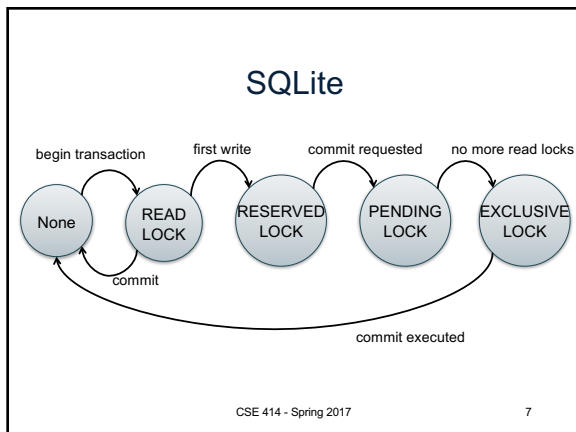
## What Data Elements are Locked?

Major differences between vendors:

- Lock on the entire database
  - SQLite
- Lock on individual records
  - SQL Server, DB2, etc.
  - can be even more fine-grained by having **different types** of locks (allows more txns to run simultaneously)

CSE 414 - Spring 2017

6



## Locks in the Abstract

CSE 414 - Spring 2017 8

## Notation

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

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

CSE 414 - Spring 2017 9

## A Non-Serializable Schedule

<p>T1</p> <p>READ(A)</p> <p>A := A+100</p> <p>WRITE(A)</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 2017 10

## Example

<p>T1</p> <p><math>L_1(A)</math>; READ(A)</p> <p>A := A+100</p> <p>WRITE(A); <math>U_1(A)</math>; <math>L_1(B)</math></p>	<p>T2</p> <p><math>L_2(A)</math>; READ(A)</p> <p>A := A*2</p> <p>WRITE(A); <math>U_2(A)</math>;</p> <p><math>L_2(B)</math>; <b>BLOCKED...</b></p>
---	---

<p>READ(B)</p> <p>B := B+100</p> <p>WRITE(B); <math>U_1(B)</math>;</p>	<p><b>...GRANTED</b>; READ(B)</p> <p>B := B*2</p> <p>WRITE(B); <math>U_2(B)</math>;</p>
--	---

Scheduler has ensured a conflict-serializable schedule

CSE 414 - Spring 2017 11

## But...

<p>T1</p> <p><math>L_1(A)</math>; READ(A)</p> <p>A := A+100</p> <p>WRITE(A); <math>U_1(A)</math>;</p>	<p>T2</p> <p><math>L_2(A)</math>; READ(A)</p> <p>A := A*2</p> <p>WRITE(A); <math>U_2(A)</math>;</p> <p><math>L_2(B)</math>; READ(B)</p> <p>B := B*2</p> <p>WRITE(B); <math>U_2(B)</math>;</p>
---	---

<p><math>L_1(B)</math>; READ(B)</p> <p>B := B+100</p> <p>WRITE(B); <math>U_1(B)</math>;</p>	<p><b>...</b></p>
---	-------------------

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

CSE 414 - Spring 2017 12

## Two Phase Locking (2PL)

The 2PL rule:

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

2PL approach developed by Jim Gray

CSE 414 - Spring 2017 13

## Example: 2PL transactions

<p>T1</p> <p><math>L_1(A); L_1(B); \text{READ}(A)</math></p> <p><math>A := A+100</math></p> <p><math>\text{WRITE}(A); U_1(A)</math></p> <p><math>\text{READ}(B)</math></p> <p><math>B := B+100</math></p> <p><math>\text{WRITE}(B); U_1(B);</math></p>	<p>T2</p> <p><math>L_2(A); \text{READ}(A)</math></p> <p><math>A := A*2</math></p> <p><math>\text{WRITE}(A);</math></p> <p><math>L_2(B); \text{BLOCKED}...</math></p> <p><math>...GRANTED; \text{READ}(B)</math></p> <p><math>B := B*2</math></p> <p><math>\text{WRITE}(B); U_2(A); U_2(B);</math></p>
---	---

Now it is conflict-serializable

CSE 414 - Spring 2017 14

## Two Phase Locking (2PL)

**Theorem:** 2PL ensures conflict serializability

15

## 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
    
```

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:

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:

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

18

### 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
    
```

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?

19

### 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
    
```

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

$U_1(A) \rightarrow L_2(A)$   
 $L_2(A) \rightarrow U_2(B)$   
 $U_2(B) \rightarrow L_3(B)$   
 $L_3(B) \rightarrow U_3(C)$   
 $U_3(C) \rightarrow L_1(C)$   
 $L_1(C) \rightarrow U_1(A)$

**Contradiction**

20

### A New Problem: Non-recoverable Schedule

T1	T2
$L_1(A); L_1(B); READ(A)$ $A := A+100$ $WRITE(A); U_1(A)$  $READ(B)$ $B := B+100$ $WRITE(B); U_1(B);$  <span style="color: blue;">Rollback</span>	$L_2(A); READ(A)$ $A := A*2$ $WRITE(A);$ $L_2(B); \text{BLOCKED}...$  <span style="color: red;">...GRANTED;</span> $READ(B)$ $B := B*2$ $WRITE(B); U_2(A); U_2(B);$ <span style="color: blue;">Commit</span>

CSE 414 - Spring 2017 21

### Strict 2PL

The Strict 2PL rule:

All locks are held until the transaction commits or aborts.

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

CSE 414 - Spring 2017 22

### Strict 2PL

T1	T2
$L_1(A); READ(A)$ $A := A+100$ $WRITE(A);$  $L_1(B); READ(B)$ $B := B+100$ $WRITE(B);$ $ROLLBACK; U_1(A), U_1(B)$	$L_2(A); \text{BLOCKED}...$  <span style="color: red;">...GRANTED;</span> $READ(A)$ $A := A*2$ $WRITE(A);$ $L_2(B); READ(B)$ $B := B*2$ $WRITE(B);$ <span style="color: blue;">COMMIT;</span> $U_2(A); U_2(B)$

CSE 414 - Spring 2017 23

### Another problem: Deadlocks

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

SQL Lite: there is only one exclusive lock; thus, never deadlocks

SQL Server: checks periodically for deadlocks and aborts one TXN

CSE 414 - Spring 2017 24

### Lock Modes

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

Lock compatibility matrix:

	None	S	X
None			
S			
X			

CSE 414 - Spring 2017 25

### Lock Modes

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

Lock compatibility matrix:

	None	S	X
None	✓	✓	✓
S	✓	✓	✗
X	✓	✗	✗

CSE 414 - Spring 2017 26

### 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 2017 27

### Lock Performance

CSE 414 - Spring 2017 28

### 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 2017 29

### Phantom Problem

Suppose there are two blue products, A1, A2:

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

SELECT \*  
FROM Product  
WHERE color='blue'

Is this schedule serializable ?

CSE 414 - Spring 2017 30

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'	

R1(A1),R1(A2),W2(A3),R1(A1),R1(A2),R1(A3)

CSE 414 - Spring 2017 31

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'	

R1(A1),R1(A2),W2(A3),R1(A1),R1(A2),R1(A3)

W2(A3),R1(A1),R1(A2),R1(A1),R1(A2),R1(A3)

CSE 414 - Spring 2017 31

### 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 2017 33

### 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 2017 34

## Locking & SQL

CSE 414 - Spring 2017 35

### Isolation Levels in SQL

- "Dirty reads"  
SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED
- "Committed reads"  
SET TRANSACTION ISOLATION LEVEL READ COMMITTED
- "Repeatable reads"  
SET TRANSACTION ISOLATION LEVEL REPEATABLE READ
- Serializable transactions  
SET TRANSACTION ISOLATION LEVEL SERIALIZABLE

ACID

CSE 414 - Spring 2017 36

## 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 2017

37

## 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 2017

38

## 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 2017

39

## 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 2017

40

## Beware!

In commercial DBMSs:

- Default level is often **NOT** serializable (SQL Server!)
- 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 probs
- **Bottom line: Read the doc for your DBMS!**

CSE 414 - Spring 2017

41

Next two slides: try them on Azure

CSE 414 - Spring 2017

42

## Demonstration with SQL Server

### Application 1:

```
create table R(a int);
insert into R values(1);
set transaction isolation level serializable;
begin transaction;
select * from R; -- get a shared lock
waitfor delay '00:01'; -- wait for one minute
```

### Application 2:

```
set transaction isolation level serializable;
begin transaction;
select * from R; -- get a shared lock
insert into R values(2); -- blocked waiting on exclusive lock
-- App 2 unblocks and executes insert after app 1 commits/aborts
```

CSE 414 - Spring 2017

43

## Demonstration with SQL Server

### Application 1:

```
create table R(a int);
insert into R values(1);
set transaction isolation level repeatable read;
begin transaction;
select * from R; -- get a shared lock
waitfor delay '00:01'; -- wait for one minute
```

### Application 2:

```
set transaction isolation level repeatable read;
begin transaction;
select * from R; -- get a shared lock
insert into R values(3); -- gets an exclusive lock on new tuple
-- If app 1 reads now, it blocks because read dirty
-- If app 1 reads after app 2 commits, app 1 sees new value
```

CSE 414 - Spring 2017

44