

Introduction to Data Management CSE 344

Lecture 22: Transaction Implementations

CSE 344 - Winter 2016

1

Announcements

- WQ6 (last!) due on Wednesday evening
- HW7 due on Monday
 - Make sure you can run starter code
 - Start early, there is little time!

CSE 344 - Winter 2016

2

Recap

- What are transactions
 - And why do we need them
- How to maintain ACID properties via schedules
 - We focus on the **isolation** property
 - We learn about atomicity & durability in 444
- How to ensure conflict-serializable schedules with locks

CSE 344 - Winter 2016

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 344

CSE 344 - Winter 2016

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 ensures conflict-serializability

CSE 344 - Winter 2016

5

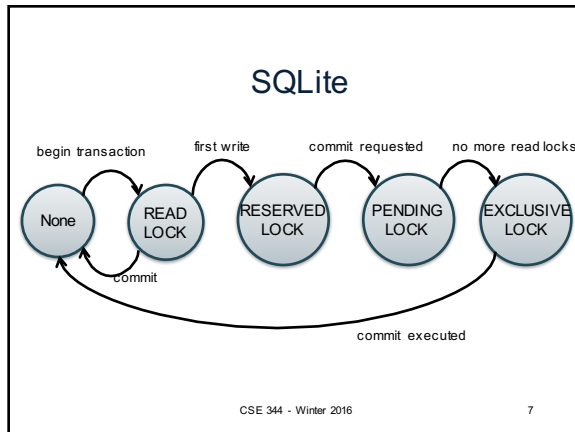
What Data Elements are Locked?

Major differences between vendors:

- Lock on the entire database
 - SQLite
- Lock on individual records
 - SQL Server, DB2, etc

CSE 344 - Winter 2016

6



Now for something more serious...

CSE 344 - Winter 2016

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 344 - Winter 2016

9

A Non-Serializable Schedule

T1	T2
READ(A)	
A := A+100	
WRITE(A)	
	READ(A)
	A := A*2
	WRITE(A)
	READ(B)
	B := B*2
	WRITE(B)
READ(B)	
B := B+100	
WRITE(B)	

CSE 344 - Winter 2016

10

Example

T1	T2
$L_1(A)$; READ(A)	
A := A+100	
WRITE(A); $U_1(A)$; $L_1(B)$	
	$L_2(A)$; READ(A)
	A := A*2
	WRITE(A); $U_2(A)$;
	$L_2(B)$; BLOCKED...
READ(B)	
B := B+100	
WRITE(B); $U_1(B)$;	
	...GRANTED ; READ(B)
	B := B*2
	WRITE(B); $U_2(B)$;

Scheduler has ensured a conflict-serializable schedule

11

But...

T1	T2
$L_1(A)$; READ(A)	
A := A+100	
WRITE(A); $U_1(A)$;	
	$L_2(A)$; READ(A)
	A := A*2
	WRITE(A); $U_2(A)$;
	$L_2(B)$; READ(B)
	B := B*2
	WRITE(B); $U_2(B)$;
$L_1(B)$; READ(B)	
B := B+100	
WRITE(B); $U_1(B)$;	

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

12

Two Phase Locking (2PL)

The 2PL rule:

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

CSE 344 - Winter 2016

13

Example: 2PL transactions

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

Now it is conflict-serializable

CSE 344 - Winter 2016

14

A New Problem: Non-recoverable Schedule

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

CSE 344 - Winter 2016

15

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 344 - Winter 2016

16

Strict 2PL

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

CSE 344 - Winter 2016

17

Another problem: Deadlocks

- T₁ waits for a lock held by T₂;
- T₂ waits for a lock held by T₃;
- T₃ waits for
- . . .
- T_n waits for a lock held by T₁

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

SQL Server: checks periodically for deadlocks and aborts one TXN

CSE 344 - Winter 2016

18

Lock Modes

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

Lock compatibility matrix:

	None	S	X
None			
S			
X			

CSE 344 - Winter 2016

19

Lock Modes

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

Lock compatibility matrix:

	None	S	X
None	✓	✓	✓
S	✓	✓	✗
X	✓	✗	✗

CSE 344 - Winter 2016

20

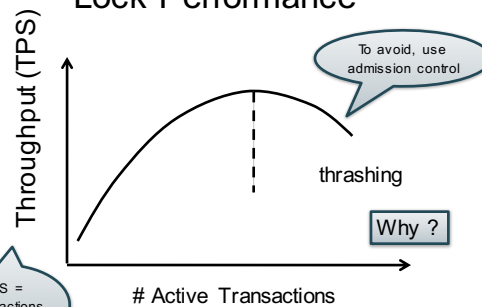
Lock Granularity

- **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 344 - Winter 2016

21

Lock Performance



CSE 344 - Winter 2016

22

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 344 - Winter 2016

23

Suppose there are two blue products, A1, A2:

Phantom Problem

<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 344 - Winter 2016

24

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 344 - Winter 2016

25

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)

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 344 - Winter 2016

27

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 344 - Winter 2016

28

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 344 - Winter 2016

29

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 344 - Winter 2016

30

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 344 - Winter 2016

31

3. Isolation Level: Repeatable Read

- “Long duration” WRITE locks
 - Strict 2PL
- “Long duration” READ locks
 - Strict 2PL

This is not serializable yet !!!

Why ?

CSE 344 - Winter 2016

32

4. Isolation Level Serializable

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

CSE 344 - Winter 2016

33

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: Read the doc for your DBMS!

CSE 344 - Winter 2016

34

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
```

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 344 - Winter 2016

35

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
```

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 344 - Winter 2016

36