

Introduction to Database Systems CSE 414

Lecture 27: Implementation of Transactions

CSE 414 - Autumn 2018

1

Announcements

- Fix quotes in Flights data
 - See email/Piazza post
 - <https://piazza.com/class/jmftm54e88t2kk?cid=729>
- Final exam Thu, Dec 13 – 2:30 here
 - Will test concepts from entire class but emphasis on post-midterm
 - Previous finals are for reference only, better to study lecture and section materials

CSE 414 - Autumn 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 - Autumn 2018

3

Example 1

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




CSE 414 - Autumn 2018

4

Example 1

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



This schedule is **conflict-serializable**

CSE 414 - Autumn 2018

5

Example 2

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




CSE 414 - Autumn 2018

6

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 - Autumn 2018

7

Implementing Transactions

CSE 414 - Autumn 2018

8

Scheduler

- **Scheduler** = the module that schedules the transaction's actions, ensuring serializability
- Also called **Concurrency Control Manager**
- We discuss next how a scheduler may be implemented

CSE 414 - Autumn 2018

9

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"

We discuss only locking schedulers in this class

CSE 414 - Autumn 2018

10

Locking 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)

By using locks scheduler ensures conflict-serializability

What Data Elements are Locked?

Major differences between vendors:

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

CSE 414 - Autumn 2018

12

More Notations

$L_i(A)$ = 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
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)	

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

But what if...

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 ?

Two Phase Locking (2PL)

The 2PL rule:

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

Example: 2PL transactions

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

Now it is conflict-serializable

Two Phase Locking (2PL)


Theorem: 2PL ensures conflict serializability

19

Two Phase Locking (2PL)

Theorem: 2PL ensures conflict serializability


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



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:

21

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** $L_2(A)$

22

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** $U_1(A)$

23

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?

24

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?

25

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.....

26

Two Phase Locking (2PL)

Theorem: 2PL ensures conflict serializability

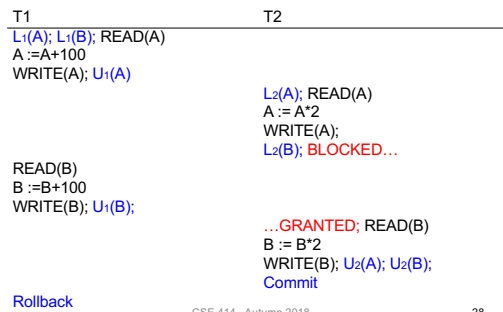
Proof. Suppose not: then there exists a cycle in the precedence graph.



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

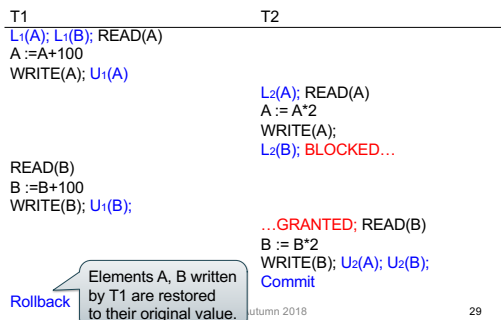
A New Problem: Non-recoverable Schedule



CSE 414 - Autumn 2018

28

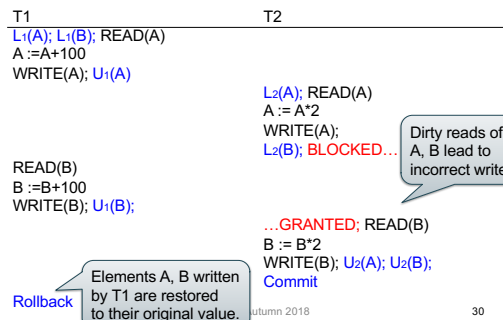
A New Problem: Non-recoverable Schedule



Autumn 2018

29

A New Problem: Non-recoverable Schedule



Autumn 2018

30

A New Problem: Non-recoverable Schedule

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

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

Dirty reads of A, B lead to incorrect writes.

Can no longer undo!

CSE 414 - Autumn 2018 32

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 32

Strict 2PL

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

33

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 34

Another problem: Deadlocks

- T₁: R(A), W(B)
- T₂: R(B), W(A)

- T₁ holds the lock on A, waits for B
- T₂ holds the lock on B, waits for A

This is a deadlock!

CSE 414 - Autumn 2018 35

Another problem: Deadlocks

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

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

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

36

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

37

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

38

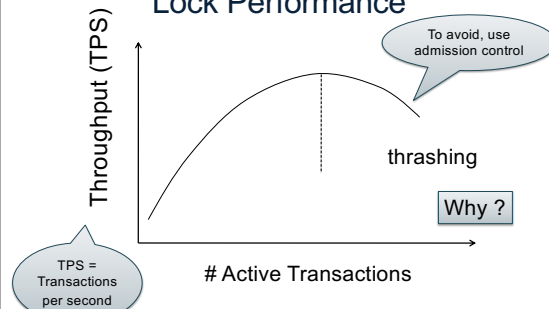
Lock Granularity

- **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

39

Lock Performance



CSE 414 - Autumn 2018

40

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

41

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 - Autumn 2018

42

Suppose there are two blue products, A1, A2:

Phantom Problem

T1	T2
<hr/>	
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 - Autumn 2018

43

Suppose there are two blue products, A1, A2:

Phantom Problem

T1	T2
<hr/>	
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)⁴⁴

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 - Autumn 2018

45

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 - Autumn 2018

46

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 - Autumn 2018

47

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 - Autumn 2018

48

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

3. Isolation Level: Repeatable Read

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

Why ?

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

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!