

Introduction to Database Systems CSE 444

Lecture 13
Transactions: concurrency control
(part 1)

Magda Balazinska - CSE 444, Fall 2010 1

Outline

- Serial and Serializable Schedules (18.1)
- Conflict Serializability (18.2)
- Locks (18.3)

Magda Balazinska - CSE 444, Fall 2010 2

The Problem

- Multiple transactions are running concurrently
 T_1, T_2, \dots
- They read/write some common elements
 A_1, A_2, \dots
- How can we prevent unwanted interference ?
- The SCHEDULER is responsible for that

Magda Balazinska - CSE 444, Fall 2010 3

Some Famous Anomalies

- What could go wrong if we didn't have concurrency control:
 - Dirty reads (including inconsistent reads)
 - Unrepeatable reads
 - Lost updates

Many other things can go wrong too

Magda Balazinska - CSE 444, Fall 2010 4

Dirty Reads

Write-Read Conflict

T_1 : WRITE(A) T_1 : ABORT	T_2 : READ(A)
---	-----------------

Magda Balazinska - CSE 444, Fall 2010 5

Inconsistent Read

Write-Read Conflict

T_1 : A := 20; B := 20; T_1 : WRITE(A) T_1 : WRITE(B)	T_2 : READ(A); T_2 : READ(B);
---	--------------------------------------

Magda Balazinska - CSE 444, Fall 2010 6

Unrepeatable Read

Read-Write Conflict

T₁: WRITE(A)

T₂: READ(A);
T₂: READ(A);

Magda Balazinska - CSE 444, Fall 2010 7

Lost Update

Write-Write Conflict

T₁: READ(A)
T₁: A := A+5
T₁: WRITE(A)

T₂: READ(A);
T₂: A := A*1.3
T₂: WRITE(A);

Magda Balazinska - CSE 444, Fall 2010 8

Schedules

- Given multiple transactions
- A *schedule* is a sequence of interleaved actions from all transactions

Magda Balazinska - CSE 444, Fall 2010 9

Example

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

Magda Balazinska - CSE 444, Fall 2010 10

A Serial Schedule

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

Magda Balazinska - CSE 444, Fall 2010 11

Serializable Schedule

- A schedule is *serializable* if it is equivalent to a serial schedule

Magda Balazinska - CSE 444, Fall 2010 12

A Serializable Schedule

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

Notice:
This is NOT a serial schedule

Magda Balazinska - CSE 444, Fall 2010 13

A Non-Serializable Schedule

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

Magda Balazinska - CSE 444, Fall 2010 14

Ignoring Details

- Sometimes transactions' actions can commute accidentally because of specific updates
 - Serializability is undecidable !
- Scheduler should not look at transaction details
- Assume worst case updates
 - Only care about reads $r(A)$ and writes $w(A)$
 - Not the actual values involved

Magda Balazinska - CSE 444, Fall 2010 15

Notation

$$T_1: r_1(A); w_1(A); r_1(B); w_1(B)$$

$$T_2: r_2(A); w_2(A); r_2(B); w_2(B)$$

Magda Balazinska - CSE 444, Fall 2010 16

Conflict Serializability

Conflicts:

Two actions by same transaction T_i : $r_i(X); w_i(Y)$

Two writes by T_i, T_j to same element $w_i(X); w_j(X)$

Read/write by T_i, T_j to same element $w_i(X); r_j(X)$

$r_i(X); w_j(X)$

Magda Balazinska - CSE 444, Fall 2010 17

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

Example:

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

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

Magda Balazinska - CSE 444, Fall 2010 18

The Precedence Graph Test

Is a schedule conflict-serializable ?
Simple test:

- Build a graph of all transactions T_i
- Edge from T_i to T_j if T_i makes an action that conflicts with one of T_j and comes first
- The test: if the graph has no cycles, then it is conflict serializable !

Magda Balazinska - CSE 444, Fall 2010 19

Example 1

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

```

    graph LR
      1((1)) -- B --> 2((2))
      2 -- A --> 3((3))
    
```

This schedule is conflict-serializable

Magda Balazinska - CSE 444, Fall 2010 20

Example 2

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

```

    graph LR
      1((1)) -- B --> 2((2))
      2 -- A --> 3((3))
      1 -- B --> 1
    
```

This schedule is NOT conflict-serializable

Magda Balazinska - CSE 444, Fall 2010 21

Conflict Serializability

- A serializable schedule need not be conflict serializable, even under the “worst case update” assumption

Lost write

$w_1(Y); w_2(Y); w_2(X); w_1(X); w_3(X);$

$w_1(Y); w_1(X); w_2(Y); w_2(X); w_3(X);$

Equivalent, but can't swap

Magda Balazinska - CSE 444, Fall 2010 22

Scheduler

- The scheduler is the module that schedules the transaction's actions, ensuring serializability
- How ? We discuss three techniques in class:
 - Locks
 - Time stamps (next lecture)
 - Validation (next lecture)

Magda Balazinska - CSE 444, Fall 2010 23

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)

Magda Balazinska - CSE 444, Fall 2010 24

Notation

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

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

Magda Balazinska - CSE 444, Fall 2010 25

Example

<p>T1</p> <p>$L_1(A)$; READ(A, t)</p> <p>t := t+100</p> <p>WRITE(A, t); $U_1(A)$; $L_1(B)$</p>	<p>T2</p> <p>$L_2(A)$; READ(A,s)</p> <p>s := s*2</p> <p>WRITE(A,s); $U_2(A)$;</p> <p>$L_2(B)$; DENIED...</p>
<p>READ(B, t)</p> <p>t := t+100</p> <p>WRITE(B,t); $U_1(B)$;</p>	<p>...GRANTED; READ(B,s)</p> <p>s := s*2</p> <p>WRITE(B,s); $U_2(B)$;</p>

Scheduler has ensured a conflict-serializable schedule 26

Example

<p>T1</p> <p>$L_1(A)$; READ(A, t)</p> <p>t := t+100</p> <p>WRITE(A, t); $U_1(A)$;</p>	<p>T2</p> <p>$L_2(A)$; READ(A,s)</p> <p>s := s*2</p> <p>WRITE(A,s); $U_2(A)$;</p> <p>$L_2(B)$; READ(B,s)</p> <p>s := s*2</p> <p>WRITE(B,s); $U_2(B)$;</p>
<p>$L_1(B)$; READ(B, t)</p> <p>t := t+100</p> <p>WRITE(B,t); $U_1(B)$;</p>	

Locks did not enforce conflict-serializability !!! 27

Two Phase Locking (2PL)

The 2PL rule:

- In every transaction, all lock requests must precede all unlock requests
- This ensures conflict serializability ! (why?)

Magda Balazinska - CSE 444, Fall 2010 28

Example: 2PL transactions

<p>T1</p> <p>$L_1(A)$; $L_1(B)$; READ(A, t)</p> <p>t := t+100</p> <p>WRITE(A, t); $U_1(A)$</p>	<p>T2</p> <p>$L_2(A)$; READ(A,s)</p> <p>s := s*2</p> <p>WRITE(A,s);</p> <p>$L_2(B)$; DENIED...</p>
<p>READ(B, t)</p> <p>t := t+100</p> <p>WRITE(B,t); $U_1(B)$;</p>	<p>...GRANTED; READ(B,s)</p> <p>s := s*2</p> <p>WRITE(B,s); $U_2(A)$; $U_2(B)$;</p>

Now it is conflict-serializable 29

What about Aborts?

- 2PL enforces conflict-serializable schedules
- But what if a transaction releases its locks and then aborts?
- Serializable schedule definition only considers transactions that commit
 - Relies on assumptions that aborted transactions can be undone completely

Magda Balazinska - CSE 444, Fall 2010 30

Example with Abort

T1	T2
$L_1(A)$; $L_1(B)$; $READ(A, t)$ $t := t+100$ $WRITE(A, t)$; $U_1(A)$ $READ(B, t)$ $t := t+100$ $WRITE(B, t)$; $U_1(B)$	$L_2(A)$; $READ(A, s)$ $s := s*2$ $WRITE(A, s)$; $L_2(B)$; DENIED... ... GRANTED ; $READ(B, s)$ $s := s*2$ $WRITE(B, s)$; $U_2(A)$; $U_2(B)$; Commit

Abort **Commit**

Strict 2PL

- Strict 2PL: All locks held by a transaction are released when the transaction is completed
- Ensures that schedules are recoverable
 - Transactions commit only after all transactions whose changes they read also commit
- Avoids cascading rollbacks

Magda Balazinska - CSE 444, Fall 2010 32

Deadlock

- Transaction T_1 waits for a lock held by T_2 ;
- But T_2 waits for a lock held by T_3 ;
- While T_3 waits for
-
- . . . and T_3 waits for a lock held by T_1 !!
- Could be avoided, by ordering all elements (see book); or deadlock detection + rollback

Magda Balazinska - CSE 444, Fall 2010 33

Lock Modes

- S = shared lock (for READ)
- X = exclusive lock (for WRITE)
- U = update lock
 - Initially like S
 - Later may be upgraded to X
- I = increment lock (for $A := A + \text{something}$)
 - Increment operations commute

Recommended reading: chapter 18.4

Magda Balazinska - CSE 444, Fall 2010 34

The Locking Scheduler

Task 1:

- Add lock/unlock requests to transactions
- Examine all $READ(A)$ or $WRITE(A)$ actions
- Add appropriate lock requests
- Ensure 2PL !

Recommended reading: chapter 18.5

Magda Balazinska - CSE 444, Fall 2010 35

The Locking Scheduler

Task 2:

- Execute the locks accordingly
- Lock table: a big, critical data structure in a DBMS !
- When a lock is requested, check the lock table
 - Grant, or add the transaction to the element's wait list
- When a lock is released, re-activate a transaction from its wait list
- When a transaction aborts, release all its locks
- Check for deadlocks occasionally

Recommended reading: chapter 18.5

Magda Balazinska - CSE 444, Fall 2010 36