

Lecture 17: Concurrency Control

Friday, February 17, 2006

1

Outline

- Serial and Serializable Schedules (18.1)
- Conflict Serializability (18.2)
- Locks (18.3)
- Multiple lock modes (18.4)
- The tree protocol (18.7)

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

3

Three Famous Anomalies

What can go wrong if we didn't have
concurrency control:

- Dirty reads
- Lost updates
- Inconsistent reads

Many other things may go wrong, but have no names

4

Dirty Reads

T₁: WRITE(A)

T₁: ABORT

T₂: READ(A)

5

Lost Update

T₁: READ(A)

T₁: A := A+5

T₁: WRITE(A)

T₂: READ(A);

T₂: A := A*1.3

T₂: WRITE(A);

6

Inconsistent Read

T₁: A := 20; B := 20;

T₁: WRITE(A)

T₁: WRITE(B)

T₂: READ(A);

T₂: READ(B);

7

Schedules

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

8

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)

9

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)

10

Serializable Schedule

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

11

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

12

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)	

13

Ignoring Details

- Sometimes transactions' actions may commute accidentally because of specific updates
 - Serializability is undecidable !
- The scheduler shouldn't look at the transactions' details
- Assume worst case updates, only care about reads $r(A)$ and writes $w(A)$

14

Notation

$T_1: r_1(A); w_1(A); r_1(B); w_1(B)$
 $T_2: r_2(A); w_2(A); r_2(B); w_2(B)$

15

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 $w_i(X); r_j(X)$

$r_i(X); w_j(X)$

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:

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




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

Conflict Serializability

- Any conflict serializable schedule is also a serializable schedule (why ?)
- The converse is not true, 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

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 !

19

Example 1

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




This schedule is conflict-serializable

20

Example 2

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



This schedule is NOT conflict-serializable

21

Scheduler

- The scheduler is the module that schedules the transaction's actions, ensuring serializability
- How ? Three techniques:
 - Locks
 - Time stamps
 - Validation

22

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)

23

Notation

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

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

24

Example

T1	T2
<p>L₁(A); READ(A, t) t := t+100 WRITE(A, t); U₁(A); L₁(B)</p> <p>READ(B, t) t := t+100 WRITE(B,t); U₁(B);</p>	<p>L₂(A); READ(A,s) s := s*2 WRITE(A,s); U₂(A); L₂(B); DENIED...</p> <p>...GRANTED; READ(B,s) s := s*2 WRITE(B,s); U₂(B);</p>
The scheduler has ensured a conflict-serializable schedule 25	

Example

T1	T2
<p>L₁(A); READ(A, t) t := t+100 WRITE(A, t); U₁(A);</p> <p>L₁(B); READ(B, t) t := t+100 WRITE(B,t); U₁(B);</p>	<p>L₂(A); READ(A,s) s := s*2 WRITE(A,s); U₂(A); L₂(B); READ(B,s) s := s*2 WRITE(B,s); U₂(B);</p>

Locks did not enforce conflict-serializability !!!

26

Two Phase Locking (2PL)

The 2PL rule:

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

27

Example: 2PL transactions

T1

L₁(A); L₁(B); READ(A, t)

t := t+100

WRITE(A, t); U₁(A)

READ(B, t)

t := t+100

WRITE(B,t); U₁(B);

T2

L₂(A); READ(A,s)

s := s*2

WRITE(A,s);

L₂(B); **DENIED...**

...GRANTED; READ(B,s)

s := s*2

WRITE(B,s); U₂(A); U₂(B);

Now it is conflict-serializable

28

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_7 waits for a lock held by T_1 !!

Could be avoided, by ordering all elements (see book); or deadlock detection plus rollback

29

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
- READ CHAPTER 18.4 !

30

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 !

31

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

32

The Tree Protocol

- An alternative to 2PL, for tree structures
- E.g. B-trees (the indexes of choice in databases)

33

The Tree Protocol

Rules:

- The first lock may be any node of the tree
- Subsequently, a lock on a node A may only be acquired if the transaction holds a lock on its parent B
- Nodes can be unlocked in any order (no 2PL necessary)

The tree protocol is NOT 2PL, yet ensures conflict-serializability !

34