

Introduction to Database Systems

CSE 444

Lecture 11

Transactions: concurrency control (part 1)

Outline

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

Next time:

- Concurrency control by timestamps (18.8)
- Concurrency control by validation (18.9)

Some additional material not in the book

The Problem

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

Some Famous Anomalies

- What could go wrong if we didn't have concurrency control?
 - Dirty reads
 - Inconsistent reads
 - Unrepeatable reads
 - Lost updates

Many other things can go wrong too

Conflicts

- Write-Read – WR
- Read-Write – RW
- Write-Write – WW

Dirty Reads

Write-Read Conflict

T_1 : WRITE(A)

T_1 : ABORT

T_2 : READ(A)

Inconsistent Read

Write-Read Conflict

```
T1: A := 20; B := 20;  
T1: WRITE(A)
```

```
T1: WRITE(B)
```

```
T2: READ(A);  
T2: READ(B);
```

Unrepeatable Read

Read-Write Conflict

T_1 : WRITE(A)

T_2 : READ(A);

T_2 : READ(A);

Lost Update

Write-Write Conflict

T_1 : READ(A)

T_1 : $A := A + 5$

T_1 : WRITE(A)

T_2 : READ(A);

T_2 : $A := A * 1.3$

T_2 : WRITE(A);

Schedules

- Given multiple transactions
 - A *schedule* is a sequence of interleaved actions from all transactions
 - A *serial schedule* is one in which transactions appear one after the other in some order with no overlap

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)

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)

Serializable Schedule

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

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

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)	

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

Notation

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

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:

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



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

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!

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

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

View Equivalence

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

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

Lost write



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

Equivalent, but can't swap

View Equivalent

T1	T2	T3	T1	T2	T3
W1(X)			W1(X)		
	W2(X)		W1(Y)		
	W2(Y)		CO1		
	CO2			W2(X)	
W1(Y)				W2(Y)	
CO1				CO2	
		W3(Y)			W3(Y)
		CO3			CO3

Lost

Serializable, but not conflict serializable

View Equivalence

Two schedules S, S' are *view equivalent* if:

- If T reads an initial value of A in S , then T also reads the initial value of A in S'
- If T reads a value of A written by T' in S , then T also reads a value of A written by T' in S'
- If T writes the final value of A in S , then it writes the final value of A in S'

Schedules with Aborted Transactions

- When a transaction aborts, the recovery manager undoes its updates
- But some of its updates may have affected other transactions !

Schedules with Aborted Transactions

T1	T2
R(A)	
W(A)	
	R(A)
	W(A)
	R(B)
	W(B)
	Commit
Abort	

Cannot abort T1 because cannot undo T2

Recoverable Schedules

- A schedule is *recoverable* if whenever a transaction T commits, all transactions who have written elements read by T have already committed

Recoverable Schedules

T1	T2
R(A)	
W(A)	
	R(A)
	W(A)
	R(B)
	W(B)
	Commit
Abort	

Nonrecoverable

T1	T2
R(A)	
W(A)	
	R(A)
	W(A)
	R(B)
	W(B)
Abort	
	Commit

Recoverable

Cascading Aborts

- If a transaction T aborts, then we need to abort any other transaction T' that has read an element written by T
- A schedule is said to *avoid cascading aborts* if whenever a transaction read an element, the transaction that has last written it has already committed.

Avoiding Cascading Aborts

T1	T2
R(A)	
W(A)	
	R(A)
	W(A)
	R(B)
	W(B)
...	...

With cascading aborts

T1	T2
R(A)	
W(A)	
Commit	
	R(A)
	W(A)
	R(B)
	W(B)
	...

Without cascading aborts

Review of Schedules

Serializability

- Serial
- Serializable
- Conflict serializable
- View equivalent to serial

Recoverability

- Recoverable
- Avoiding cascading deletes

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)

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)

Notation

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

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

Example

T1

$L_1(A)$; READ(A, t)
t := t+100
WRITE(A, t); $U_1(A)$; $L_1(B)$

READ(B, t)
t := t+100
WRITE(B,t); $U_1(B)$;

T2

$L_2(A)$; READ(A,s)
s := s*2
WRITE(A,s); $U_2(A)$;
 $L_2(B)$; **DENIED...**

...GRANTED; READ(B,s)
s := s*2
WRITE(B,s); $U_2(B)$;

Scheduler has ensured a conflict-serializable schedule

Example

T1

$L_1(A)$; READ(A, t)

t := t+100

WRITE(A, t); $U_1(A)$;

$L_1(B)$; READ(B, t)

t := t+100

WRITE(B,t); $U_1(B)$;

T2

$L_2(A)$; READ(A,s)

s := s*2

WRITE(A,s); $U_2(A)$;

$L_2(B)$; READ(B,s)

s := s*2

WRITE(B,s); $U_2(B)$;

Locks did not enforce conflict-serializability !!!

Two Phase Locking (2PL)

The 2PL rule:

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

Example: 2PL transactions

T1

$L_1(A)$; $L_1(B)$; READ(A, t)
t := t+100
WRITE(A, t); $U_1(A)$

READ(B, t)

t := t+100

WRITE(B,t); $U_1(B)$;

T2

$L_2(A)$; READ(A,s)

s := s*2

WRITE(A,s);

$L_2(B)$; **DENIED...**

...**GRANTED**; READ(B,s)

s := s*2

WRITE(B,s); $U_2(A)$; $U_2(B)$;

Now it is conflict-serializable

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

A Non-Recoverable Schedule

T1

$L_1(A)$; $L_1(B)$; READ(A, t)
 $t := t+100$
WRITE(A, t); $U_1(A)$

READ(B, t)
 $t := t+100$
WRITE(B,t); $U_1(B)$;

Abort

T2

$L_2(A)$; 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

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**

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_{73} waits for a lock held by T_1 !!

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

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

Lock Granularity

- **Fine granularity locking** (e.g., tuples)
 - High concurrency
 - High overhead in managing locks
- **Coarse grain locking** (e.g., tables, predicate locks)
 - Many false conflicts
 - Less overhead in managing locks
- **Alternative techniques**
 - Hierarchical locking (and intentional locks) [commercial DBMSs]
 - Lock escalation

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 !

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