CSE 444: Database Internals

Lectures 13 Transaction Schedules

Announcements

- Lab 2 is due TODAY
 - Lab 3 will be released tomorrow, part 1 due next Monday
- HW4 is due on Wednesday
 - HW3 has been released, due next week
- 544M: Paper 3 reading is due TODAY
 - Papers 4 and 5 are due on same day in a few weeks
 - Write-up should be 2 to 3 pages long since 2 papers

Motivating Example



Transaction

<u>**Definition</u>**: a transaction is a sequence of updates to the database with the property that either all complete, or none completes (all-or-nothing).</u>



In ad-hoc SQL: each statement = one transaction This is referred to as autocommit

Motivating Example

START TRANSACTION

UPDATE Budget SET money=money-100 WHERE pid = 1

UPDATE Budget SET money=money+60 WHERE pid = 2

UPDATE Budget SET money=money+40 WHERE pid = 3 COMMIT (or ROLLBACK)



Transactions

- Major component of database systems
- Critical for most applications; arguably more so than SQL
- Turing awards to database researchers:
 - Charles Bachman 1973
 - Edgar Codd 1981 for inventing relational dbs
 - Jim Gray 1998 for inventing transactions
 - Mike Stonebraker 2015 for INGRES and Postgres
 - And many other ideas after that

ROLLBACK

- If the app gets to a place where it can't complete the transaction successfully, it can execute ROLLBACK
- This causes the system to "abort" the transaction
 - Database returns to a state without any of the changes made by the transaction
- Several reasons: user, application, system

ACID Properties

- Atomicity: Either all changes performed by transaction occur or none occurs
- Consistency: A transaction as a whole does not violate integrity constraints
- Isolation: Transactions appear to execute one after the other in sequence
- Durability: If a transaction commits, its changes will survive failures

What Could Go Wrong?

Why is it hard to provide ACID properties?

- Concurrent operations
 - Isolation problems
 - We saw one example earlier
- Failures can occur at any time
 - Atomicity and durability problems
 - Later lectures
- Transaction may need to abort

Different Types of Problems



Different Types of Problems



What could go wrong ?

Lost update

Different Types of Problems



What could go wrong ?

Dirty reads

Types of Problems: Summary

- Concurrent execution problems
 - Write-read conflict: dirty read (includes inconsistent read)
 - A transaction reads a value written by another transaction that has not yet committed
 - Read-write conflict: unrepeatable read
 - A transaction reads the value of the same object twice. Another transaction modifies that value in between the two reads
 - Write-write conflict: lost update
 - Two transactions update the value of the same object. The second one to write the value overwrite the first change
- Failure problems
 - DBMS can crash in the middle of a series of updates
 - Can leave the database in an inconsistent state

Terminology Needed For Lab 3 Buffer Manager Policies

- STEAL or NO-STEAL
 - Can an update made by an uncommitted transaction overwrite the most recent committed value of a data item on disk?

FORCE or NO-FORCE

- Should all updates of a transaction be forced to disk before the transaction commits?
- Easiest for recovery: NO-STEAL/FORCE (lab 3)
- Highest performance: STEAL/NO-FORCE (lab 5)
- We will get back to this next week

Outline

- Transactions motivation, definition, properties
- Concurrency Control (the C in ACID)
 - This week
- Recovery from failures (the A in ACID)
 Next week

Schedules

A <u>schedule</u> is a sequence of interleaved actions from all transactions

A and B are elements in the database Example t and s are variables in tx source code **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 <u>serializable</u> if it is equivalent to a serial schedule

A Serializable Schedule



A Non-Serializable Schedule



Serializable Schedules

• The role of the scheduler is to ensure that the schedule is serializable

Q: Why not run only serial schedules ? I.e. run one transaction after the other ?

Serializable Schedules

• The role of the scheduler is to ensure that the schedule is serializable

Q: Why not run only serial schedules ? I.e. run one transaction after the other ?

A: Because of very poor throughput due to disk latency.

Lesson: main memory databases may schedule TXNs serially

Still Serializable, but...



...we don't expect the scheduler to schedule this

Ignoring Details

- Assume worst case updates:
 - We never commute actions done by transactions
- As a consequence, we only care about reads and writes
 - Transaction = sequence of R(A)'s and W(A)'s

Conflicts

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

Conflicts:

Two actions by same transaction T_i:

$$r_i(X); w_i(Y)$$

Two writes by T_i , T_i to same element



Read/write by T_i , T_i to same element



Definition A schedule is <u>conflict serializable</u> if it can be transformed into a serial schedule by a series of swappings of adjacent non-conflicting actions

- Every conflict-serializable schedule is serializable
- The converse is not true in general

Example:

r₁(A); w₁(A); r₂(A); w₂(A); r₁(B); w₁(B); r₂(B); w₂(B)

Example:

r₁(A); w₁(A); r₂(A); w₂(A); r₁(B); w₁(B); r₂(B); w₂(B)

r₁(A); w₁(A); r₁(B); w₁(B); r₂(A); w₂(A); r₂(B); w₂(B)



r₁(A); w₁(A); r₁(B); w₁(B); r₂(A); w₂(A); r₂(B); w₂(B)



r₁(A); w₁(A); r₁(B); w₁(B); r₂(A); w₂(A); r₂(B); w₂(B)





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_i
- The schedule is serializable iff the precedence graph is acyclic

Example 1

r₂(A); r₁(B); w₂(A); r₃(A); w₁(B); w₃(A); r₂(B); w₂(B)



Example 1



 $1 \xrightarrow{B} 2 \xrightarrow{A} 3$

This schedule is conflict-serializable

Example 2

r₂(A); r₁(B); w₂(A); r₂(B); r₃(A); w₁(B); w₃(A); w₂(B)



Example 2 $r_2(A); r_1(B); w_2(A); r_2(B); r_3(A); w_1(B); w_3(A); w_2(B)$ В Α 2 3 B

This schedule is NOT conflict-serializable

 A serializable schedule need not be conflict serializable, even under the "worst case update" assumption

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

Is this schedule conflict-serializable ?

 A serializable schedule need not be conflict serializable, even under the "worst case update" assumption

$$W_1(X); W_2(X); W_2(Y); W_1(Y); W_3(Y);$$

Is this schedule conflict-serializable ?



 A serializable schedule need not be conflict serializable, even under the "worst case update" assumption

Equivalent, but not conflict-equivalent



Serializable, but not conflict serializable 42

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

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

View-Serializability

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

Remark:

- If a schedule is *conflict serializable*, then it is also *view serializable*
- But not vice versa

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



Schedules with Aborted Transactions



Cannot abort T1 because cannot undo T2

Recoverable Schedules

A schedule is *recoverable* if:

- It is conflict-serializable, and
- Whenever a transaction T commits, all transactions who have written elements read by T have already committed

Recoverable Schedules



Re	covera	ble S	chedules
T1	T2	Т3	T4
R(A) W(A)			
	R(A)		
	W(A)		
	к(b) W(B)		
		R(B)	
		W(B)	
		R(C) W(C)	
		VV(O)	R(C)
			W(C)
			R(D)
Abort			VV(D)
	How do we recover ?		

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 avoids cascading aborts if whenever a transaction reads an element, the transaction that has last written it has already committed.

Avoiding Cascading Aborts



Review of Schedules

Serializability

Recoverability

- Serial
- Serializable
- Conflict serializable
- View serializable

- Recoverable
- Avoids cascading deletes

Scheduler

- The scheduler:
- Module that schedules the transaction's actions, ensuring serializability
- Two main approaches
- Pessimistic: locks
- Optimistic: time stamps, MV, validation