Introduction to Data Management CSE 344

Lecture 21: More Transactions

Announcements

- HW6 due tonight
- HW7:
 - Some Java programming required
 - Plus connection to SQL Azure
 - Due Monday, March 7

Outline

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

Review: Transactions

- Problem: An application must perform *several* writes and reads to the database, as a unit
- Solution: multiple actions of the application are bundled into one unit called a *Transaction*
- Turing awards to database researchers
 - Charles Bachman 1973 for CODASYL
 - Edgar Codd 1981 for relational databases
 - Jim Gray 1998 for transactions

CSE 344 - Winter 2016

Review: TXNs in SQL

BEGIN TRANSACTION [SQL statements] COMMIT or ROLLBACK (=ABORT)



Review: ACID

- Atomic
 - State shows either all the effects of txn, or none of them
- Consistent
 - Txn moves from a state where integrity holds, to another where integrity holds
- Isolated
 - Effect of txns is the same as txns running one after another (i.e., looks like batch mode)
- Durable
 - Once a txn has committed, its effects remain in the database

Isolation: The Problem

- Multiple transactions are running concurrently $T_1, T_2, ...$
- They read/write some common elements
 A₁, A₂, …
- How can we prevent unwanted interference ?
- The SCHEDULER is responsible for that

Schedules

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

Serial Schedule

- A <u>serial schedule</u> is one in which transactions are executed one after the other, in some sequential order
- Fact: nothing can go wrong if the system executes transactions serially
 - But database systems don't do that because we need better performance

CSE 344 - Winter 2016



A and B are elements in the database t and s are variables in txn source code

Time





Time

Serializable Schedule

A schedule is <u>serializable</u> if it is equivalent to a serial schedule



CSE 344 - Winter 2016

A Non-Serializable Schedule



How do We Know if a Schedule is Serializable?

Notation

T₁: $r_1(A)$; $w_1(A)$; $r_1(B)$; $w_1(B)$ T₂: $r_2(A)$; $w_2(A)$; $r_2(B)$; $w_2(B)$

Key Idea: Focus on *conflicting* operations

Conflicts

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

Conflicts: (it means: cannot be swapped) $r_i(X); w_i(Y)$ Two actions by same transaction T_i :

Two writes by T_i , T_i to same element

$$w_i(X); w_j(X)$$

Read/write by T_i , T_i to same element





CSF 344 - Winter 2016

- 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
- A serializable schedule may not necessarily be conflict-serializable

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)



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

Example:

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

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

CSE 344 - Winter 2016



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



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 2

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



Example 2 r₂(A); r₁(B); w₂(A); r₂(B); r₃(A); w₁(B); w₃(A); w₂(B) В Α 2 3 R

This schedule is NOT conflict-serializable

CSE 344 - Winter 2016

Scheduler

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

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"
 - Postgres, Oracle

We discuss only locking in 344

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

Let's Study SQLite First

- SQLite is very simple
- More info: <u>http://www.sqlite.org/atomiccommit.html</u>
- Lock types
 - READ LOCK (to read)
 - RESERVED LOCK (to write)
 - PENDING LOCK (wants to commit)
 - EXCLUSIVE LOCK (to commit)

Step 1: when a transaction begins

- Acquire a **READ LOCK** (aka "SHARED" lock)
- All these transactions may read happily
- They all read data from the database file
- If the transaction commits without writing anything, then it simply releases the lock

Step 2: when one transaction wants to write

- Acquire a **RESERVED LOCK**
- May coexists with many READ LOCKs
- Writer TXN may write; these updates are only in main memory; others don't see the updates
- Reader TXN continue to read from the file
- New readers accepted
- No other TXN is allowed a RESERVED LOCK

Step 3: when writer transaction wants to commit, it needs *exclusive lock*, which can't coexists with *read locks*

Acquire a PENDING LOCK

Why not write to disk right now?

- May coexists with old READ LOCKs
- No new READ LOCKS are accepted
- Wait for all read locks to be released

Step 4: when all read locks have been released

- Acquire the EXCLUSIVE LOCK
- Nobody can touch the database now
- All updates are written permanently to the database file
- Release the lock and COMMIT



SQLite Demo

create table r(a int, b int); insert into r values (1,10); insert into r values (2,20); insert into r values (3,30);

T1: begin transaction; select * from r; -- T1 has a READ LOCK T2: begin transaction; select * from r; -- T2 has a READ LOCK

```
T1:
update r set b=11 where a=1;
-- T1 has a RESERVED LOCK
```

T2:

update r set b=21 where a=2; -- T2 asked for a RESERVED LOCK: DENIED

T3:

begin transaction;

select * from r;

commit;

-- everything works fine, could obtain READ LOCK

T1:

commit;

- -- SQL error: database is locked
- -- T1 asked for PENDING LOCK -- GRANTED
- -- T1 asked for EXCLUSIVE LOCK -- DENIED

T3': begin transaction; select * from r; -- T3 asked for READ LOCK-- DENIED (due to T1)

T2:

commit;

-- releases the last READ LOCK; T1 can commit

Review: 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

Dirty Reads

Write-Read Conflict





Inconsistent Read

Write-Read Conflict



Unrepeatable Read

Read-Write Conflict

T₁: WRITE(A)

CSE 344 - Winter 2016

Write-Write Conflict
$$T_1: READ(A)$$
 $T_2: READ(A);$ $T_1: A := A+5$ $T_2: READ(A);$ $T_1: WRITE(A)$ $T_2: WRITE(A);$

Lost Update

50