L16: Transactions & Concurrency Control Part 1: Transactions & Logging Part 2: Concurrency Control

CS3200 Database design (fa18 s2)

https://northeastern-datalab.github.io/cs3200/

Version 11/1/2018

Announcements!

- Exam 2 is on Monday
 - content is everything seen until today
 - setup like for Exam1: laptop SQL + paper database design + paper transactions
- HW10 = Past semester's final exam (ignore section 4 on internals / sort)
- Notice our changed OHs this week (no THU, but TUE and FRI)
- Today
 - Transactions
- CURN CCIS Undergraduate Research Night: Monday 6:30-8:30
- Jupyter exercises: let's try it once more together

CURN (CCIS Undergraduate Research Night): Mon, Nov 5 UNDERGRADUATE RESEARCH NIGHT

NOVEMBER 5 6:30 PM - 8:30 PM EST



Title: Undergraduate Research Night

Location: Northeastern University, 102 West Village H, 440 Huntington Avenue, Boston, Massachusetts 02115

Abstract

Learn about how you can participate in cutting-edge computer science research! Featuring a panel with fellow undergraduate students who have participated in research to discuss their experience, and academic advisors to discuss how you can integrate research into your curriculum (as a co-op and course credit). Afterward, attend the poster session to interact with researchers discussing ongoing research projects in the college and how you can get involved. This will be the fourth annual UG Research Night that CCIS has hosted.

Free food provided!

ACID: Isolation

- A transaction executes concurrently with other transactions
- <u>Isolation</u>: the effect is as if each transaction executes in isolation of the others.
 - E.g. Should not be able to observe changes from other transactions during the run

Isolation failure

Write-Write Conflict



ACI<u>D</u>: Durability

- The effect of a TXN must continue to exist ("persist") after the TXN
 - And after the whole program has terminated
 - And even if there are power failures, crashes, etc.
 - And etc...
- Means: Write data to disk



Change on the horizon? Non-Volatile Ram (NVRam). Byte addressable.

Challenges for ACID properties

- In spite of failures: Power failures, but not media failures
- Users may abort the program: need to "<u>rollback</u> the changes"
 - Need to log what happened
- Many users executing concurrently
 - Can be solved via locking

And all this with... Performance!!

A Note: ACID is contentious!

- Many debates over ACID, both historically and currently
- Many newer "NoSQL" DBMSs relax ACID
- In turn, now "NewSQL" reintroduces ACID compliance to NoSQL-style DBMSs...



ACID is an extremely important & successful paradigm, but still debated!

Our first Goal: Ensuring Atomicity & Durability

- <u>A</u>tomicity:
 - TXNs should either happen completely or not at all
 - If abort / crash during TXN, no effects should be seen
- <u>D</u>urability:
 - If DBMS stops running, changes due to completed TXNs should all persist
 - Just store on stable disk

| TXN 1 | Crash / abort |
|-------------------------------------|---------------|
| | |
| <u>No</u> changes persist | |
| TXN 2 | |
| | |
| <u>All</u> changes persist | |

We'll focus on how to accomplish atomicity (via logging)

The Log

- Is a list of modifications
- Log is <u>duplexed and archived</u> on stable storage.

Assume we don't lose it!

- Can <u>force write</u> entries to disk
 - A page goes to disk.
- All log activities handled transparently by the DBMS.
 - ("transparently" = without the user being aware)

Basic Idea: (Physical) Logging

- Record UNDO information for every update!
 - Sequential writes to log
 - Minimal info (diff) written to log
- The log consists of an <u>ordered list of actions</u>
 - Log record contains:
 - <XID, location, old data, new data>

This is sufficient to UNDO any transaction!

3. Atomicity & Durability via Logging

An animation of commit protocols

A picture of logging





A picture of logging



Operation recorded in log in main memory!



What is the correct way to write this all to disk?

- We'll look at the <u>Write-Ahead Logging (WAL)</u> protocol
- We'll see why it works by looking at other protocols which are incorrect!

Remember: Key idea is to ensure durability *while* maintaining our ability to "undo"!

Write-Ahead Logging (WAL) TXN Commit Protocol

Transaction Commit Process

- FORCE Write commit record to log
- All log records up to last update from this TX are FORCED
- Commit() returns

Transaction is committed once commit log record is on stable storage











Let's try committing *after* we've written data but *before* we've written log to disk...

OK, Commit!



Improved Commit Protocol (WAL)

Write-ahead Logging (WAL) Commit Protocol



Write-ahead Logging (WAL) Commit Protocol



• DB uses Write-Ahead Logging (WAL) Protocol:

- 1. Must force log record for an update before the corresponding data page goes to storage
- 2. Must write all log records for a TX before commit

Each update is logged! Why not reads?

 \rightarrow <u>Atomicity</u>

 \rightarrow <u>Durability</u>

Logging Summary

- If DB says TX commits, TX effect remains after database crash
- DB can <u>undo actions</u> and help us with <u>atomicity</u>

• This is only half the story...

Concurrency & Locking

 Concurrency, scheduling & anomalies
Locking: 2PL, conflict serializability, deadlock detection

1. Concurrency, Scheduling & Anomalies

What we will learn next

- Interleaving & scheduling
- Conflict & anomaly types
- ACTIVITY: TXN viewer

Concurrency: Isolation & Consistency

- The DBMS must handle concurrency s.t. ...
 - <u>Isolation</u> is maintained: Users must be able to execute each TXN <u>as if they were the only user</u>
 - DBMS handles the details of interleaving various TXNs

- <u>Consistency</u> is maintained: TXNs must leave the DB in a consistent state
 - DBMS handles the details of enforcing integrity constraints

The hard part is the effect of *interleaving* transactions and *crashes*.





```
T1: START TRANSACTION
UPDATE Accounts
SET Amt = Amt + 100
WHERE Name = 'A'
UPDATE Accounts
SET Amt = Amt - 100
WHERE Name = 'B'
```

COMMIT

T1 transfers \$100 from B's account to A's account

T2: START TRANSACTION UPDATE Accounts SET Amt = Amt * 1.06 COMMIT

T2 credits both accounts with a 6% interest payment

We can look at the TXNs in a timeline view- serial execution:

Time

T1 transfers \$100 from B's account to A's account

T2 credits both accounts with a 6% interest payment

The TXNs could occur in either order... DBMS allows!

T₁
$$A += 100$$
 $B -= 100$
T₂ $A *= 1.06$ $B *= 1.06$

Time

T2 credits both accounts with a 6% interest payment

T1 transfers \$100 from B's account to A's account

The DBMS can also interleave the TXNs



Time

T1 transfers \$100 to A's accout, then T2 credits A's account with 6% interest payment T2 credits B's account with a 6% interest payment, then T1 transfers \$100 from B's account...

The DBMS can also interleave the TXNs



Time

What goes wrong here??

Recall: Three Types of Regions of Memory

- Local: In our model each process in a DBMS has its own local memory, where it stores values that only it "sees"
- 2. <u>Global</u>: Each process can read from / write to shared data in main memory
- 3. <u>Disk</u>: Global memory can read from / flush to disk
- 4. <u>Log</u>: Assume on stable disk storage- spans both main memory and disk...



Log is a *sequence* from main memory -> disk

<u>"Flushing</u> to disk" = writing to disk from main memory

Why Interleave TXNs?

• Interleaving TXNs might lead to anomalous outcomes... why do it?

- Several important reasons:
 - Individual TXNs might be slow- don't want to block other users during!
 - Disk access may be slow- let some TXNs use CPUs while others accessing disk!

All concern large differences in *performance*

Interleaving & Isolation

• The DBMS has freedom to interleave TXNs

 However, it must pick an interleaving or schedule such that isolation and consistency are maintained

"With great power comes great responsibility"

- Must be as if the TXNs had executed serially!



DBMS must pick a schedule which maintains isolation & consistency

| | Starting | Α | В | |
|--|----------|-------|-------|---------|
| Carial cabadula T. T. | Balance | \$50 | \$200 | |
| <u>senai schedule 1₁,1₂:</u> | | | | |
| T ₁ A += 100 B -= 100 | | | | |
| T ₂ A *= 1.06 B *= 1.0 | 06 | Α | В | |
| · Z | | \$159 | \$106 | |
| | | | | |
| Interleaved schedule 1: | | | | Same |
| T $A = 100$ B = 100 | | | | result! |
| 1 | | | | |
| A *= 1.06 B *= 1.0 | 6 | Α | В | |
| ' 2 | | \$159 | \$106 | |

| | Starting | Α | В | |
|------------------------------------|----------|------------|-------|---------------------------------------|
| Serial schedule T. T.: | Balance | \$50 | \$200 | |
| | | | | |
| T ₁ A += 100 B -= 100 | | | | |
| T ₂ A *= 1.06 B *= 1. | 06 | А | В | |
| ζ | | \$159 | \$106 | |
| Interleaved schedule 2: | | | | Different result than |
| T ₁ A += 100 B -= 10 | 00 | | | serial 1 ₁ ,1 ₂ |
| T ₂ A *= 1.06 B *= 1.06 | | A \$159 | \$112 | |



⁸¹

This schedule is different than *any serial order!* We say that it is <u>not</u> <u>serializable</u>

Interleaved schedule 2:

Scheduling Definitions

- A serial schedule is one that does not interleave the actions of different transactions
- Schedules X and Y are equivalent schedules if, for any database state, the effect on DB of executing X is <u>identical</u> to the effect of executing Y
- A serializable schedule is a schedule that is equivalent to <u>some</u> serial execution of the transactions.

The word "**some"** makes this definition powerful & tricky!

Serializable?







| А | В | |
|--------------|--------------|--|
| 1.06*(A+100) | 1.06*(B-100) | |

Same as a serial schedule for all possible values of A, B = <u>serializable</u>

Serializable?



Not *equivalent* to any serializable schedule = *not* <u>serializable</u>

Serial schedules:

What else can go wrong with interleaving?

- Various <u>anomalies</u> which break isolation / serializability
 - Often referred to by name (recall last class: "lost updated", "unrepeateable reads", etc.)
 - We see them again in a bit
- Occur because of / with certain "conflicts" between interleaved TXNs



Conflict Types

Two actions <u>conflict</u> if they are part of different TXNs, involve the same variable, and at least one of them is a write

- Thus, there are three types of conflicts:
 - Read-Write conflicts (RW)
 - Write-Read conflicts (WR)
 - Write-Write conflicts (WW)

Why no "RR Conflict"?

Interleaving anomalies occur with / because of these conflicts between TXNs (but these conflicts can occur without causing anomalies!)

Classic Anomalies with Interleaved Execution

"Unrepeatable read"





1. $T_1 reads$ some data from A

2. T_2 <u>writes</u> to A

3. Then, T₁ reads from A again and now gets a different / inconsistent value

Occurring with / because of a RW conflict

"Dirty read" (Reading uncommitted data)



1. $T_1 \underline{\text{writes}}$ some data to A

- 2. T₂ <u>reads</u> from A, then writes back to A & commits
- 3. T_1 then aborts- now T_2 's result is based on an obsolete / inconsistent value

Occurring with / because of a WR conflict

"Inconsistent read" (Reading partial commits)



1. $T_1 \underline{\text{writes}}$ some data to A

- T₂ <u>reads</u> from A *and B*, and then writes some value which depends on A & B
- 3. T_1 then writes to B- now T_2 's result is based on an incomplete commit

Again, occurring because of a WR conflict

Partially-lost update





1. $T_1 \underline{blind}$ writes some data to A

- 2. T₂ <u>blind writes</u> to A and B
- 3. T₁ then <u>blind writes</u> to B; now we have T₂'s value for B and T₁'s value for A- not equivalent to any serial schedule!

Occurring because of a WW conflict

Activity-31.ipynb

2. Conflict Serializability, Locking & Deadlock

What we will learn next

- RECAP: Concurrency
- Conflict Serializability
- DAGs & Topological Orderings
- Strict 2PLDeadlocks

Recall: Concurrency as Interleaving TXNs



 For our purposes, having TXNs occur concurrently means <u>interleaving their</u> <u>component actions (R/W)</u>

We call the particular order of interleaving a <u>schedule</u>

Recall: "Good" vs. "bad" schedules



We want to develop ways of discerning "good" vs. "bad" schedules

Ways of Defining "Good" vs. "Bad" Schedules

- Recall: we call a schedule <u>serializable</u> if it is equivalent to some serial schedule
 - We used this as a notion of a "good" interleaved schedule, since <u>a</u> serializable schedule will maintain isolation & consistency
- Now, we'll define a stricter, but very useful variant:
 - Conflict serializability

We'll need to define *conflicts* first..



Two actions <u>conflict</u> if they are part of different TXNs, involve the same variable, and at least one of them is a write





Two actions <u>conflict</u> if they are part of different TXNs, involve the same variable, and at least one of them is a write



All "conflicts"!

Conflict Serializability

- Two schedules are <u>conflict equivalent</u> if:
 - They involve the same actions of the same TXNs
 - Every pair of conflicting actions of two TXNs are ordered in the same way

 Schedule S is <u>conflict serializable</u> if S is conflict equivalent to some serial schedule

Conflict serializable \Rightarrow serializable

So if we have conflict serializable, we have consistency & isolation!

Recall: "Good" vs. "bad" schedules



Conflict serializability also provides us with an operative notion of "good" vs. "bad" schedules!

Note: Conflicts vs. Anomalies

- <u>Conflicts</u> are things we talk about to help us characterize different schedules
 - Present in both "good" and "bad" schedules

- <u>Anomalies</u> are instances where isolation and/or consistency is broken because of a "bad" schedule
 - We often characterize different anomaly types by what types of conflicts predicated them

The Conflict Graph

• Let's now consider looking at conflicts at the TXN level

• Consider a graph where the nodes are TXNs, and there is an edge from $T_i \rightarrow T_j$ if any actions in T_i precede and conflict with any actions in T_j



What can we say about "good" vs. "bad" conflict graphs?



What can we say about "good" vs. "bad" conflict graphs?



<u>Theorem</u>: Schedule is **conflict serializable** if and only if its conflict graph is <u>acyclic</u>

DAGs & Topological Orderings

• A <u>topological ordering</u> of a directed graph is a linear ordering of its vertices that respects all the directed edges

- A <u>directed acyclic graph (DAG)</u> always has one or more topological orderings
 - (And there exists a topological ordering if and only if there are no directed cycles)

DAGs & Topological Orderings

• Ex: What is one possible topological ordering here?



Ex: 0, 1, 2, 3 (or: 0, 1, 3, 2)

DAGs & Topological Orderings

• Ex: What is one possible topological ordering here?



There is none!

Connection to conflict serializability

 In the conflict graph, a topological ordering of nodes corresponds to a serial ordering of TXNs

• Thus an <u>acyclic conflict graph</u> \rightarrow conflict serializable!

<u>Theorem</u>: Schedule is **conflict serializable** if and only if its conflict graph is <u>acyclic</u>