

0. Introduction

- The problem: We have multiple transactions -- T_1, T_2, \dots, T_N -- all of which must be atomic, and all of which can have multiple steps. We want to schedule the steps of these transactions so that it appears as if they ran sequentially.
- Naive solution: run transaction sequentially with a single global lock.
 - Very poor performance
- Better solution: fine-grained locking. But we all agreed that this was error prone back in the OS section. What to do?

1. Serializability

- What does it mean for transactions to "appear" as if they were run in sequence?
 - That the final written state is the same?
 - That the final written state + intermediate reads are the same?
 - Something else?
- It depends! There are different types of "serializability". The right one depends on what your application is doing.

2. Conflict Serializability

- Two operations conflict if
 1. They both operate on the same object
 2. At least one of them is a write
- Definition highlights that concurrent reads are generally fine, but problems arise as soon as writes get involved.
- A schedule is conflict serializable if the order of all of its conflicts is the same as the order of the conflicts in some sequential schedule.
 - By "order of conflicts" we mean the ordering of the steps in each individual conflict.
- See slides for examples. A schedule can be final-state serializable but not conflict serializable.

3. Conflict graphs

- Nodes are transactions
- Edges are directed
- There is an edge from T_i to T_j iff
 - T_i and T_j have a conflict between them
 - The first step in the conflict occurs in T_i
- Example (this one not covered in lecture; it's good practice!)


T1: read(x)
T2: write(x)
T3: read(y)
T4: read(y)

T1: write(y)
 T2: write(y)
 T3: write(z)

The conflicts here are:

T1: read(x); T2: write(x)
 T3: read(y); T1: write(y)
 T3: read(y); T2: write(y)
 T4: read(y); T1: write(y)
 T4: read(y); T2: write(y)
 T1: write(y); T2: write(y)

The conflict graph is:



- Acyclic conflict graph \Leftrightarrow conflict-serializable

4. Interlude

- We're going to explore conflict serializability in more depth because it's useful in practice.
- As of right now, we have no methodical way to create conflict serializable schedules
- We can check if a schedule is conflict serializable, but if we want a conflict serializable schedule the best we can do right now is keep generating random schedules and testing until we find one with an acyclic conflict graph

5. Two-phase locking (2PL)

- So how do we generate conflict-serializable schedules?
- Via two-phase locking:
 - Each shared variable has a lock (fine-grained locking).
 - Before any operation on the variable, the transaction must acquire the corresponding lock.
 - After a transaction releases a lock, it may not acquire any other locks.
- Note: fine-grained locking but in a systematic way
- Why "two-phase"? We see two phases:
 - Acquire phase, where transactions acquire locks
 - Release phase, where transactions release locks
- Immediate problem: 2PL can result in deadlock:
- One solution to deadlock: global ordering on locks. Not very modular.
- Better solution: take advantage of atomicity and abort one of the

transactions.

- Seems like we're punting, but is actually very elegant; given atomicity, aborting is a-okay.
- Detecting deadlock is possible
 - Use "wait-dependency" graphs, which capture the locks each transaction has and the ones it wants. Cycle in wait-dependency graph => deadlock
 - Or just abort a transaction after X seconds (X "reasonably" large). Simpler.

6. Performance Improvement: Reader-writer locks

- Reader-writer locks.
- Rules:
 - Can acquire a reader lock at the same time as other readers.
 - Can acquire a writer lock only if there are no other readers or writers.
- What about fairness? If readers keep acquiring the lock, and a writer is waiting?
 - Typically: if writer is waiting, new readers wait too.
- Reader-writer locks improve performance
 - As described, they allow for concurrent reads.
 - We can also release read locks prior to commit
 - Why? Once a transaction T has acquired all its locks (reached its "lock point") any conflict transaction will run later
 - If T reaches its lock point and will no longer access X, releasing read locks on X will be fine
 - Hold write locks until commit, though, in case the transaction aborts.
- Can also improve performance by relaxing our requirements for serializability/isolation.
 - Read-committed or snapshot isolation (see future hands-on)
- Important: there are a ton of tradeoffs between performance and isolation semantics.


7. Another Possible Performance improvement: Giving up on Conflict Serializability

- Sometimes conflict serializability can seem like too strict a requirement
- Example:

T1	T2	T3
read(x)		
	write(x)	
write(x)		write(x)

Conflict graph:

T1 <--> T2



- Not acyclic => not conflict serializable
 - But compare it to running T1, then T2, then T3 (sequentially)
 - Final-state is fine
 - Intermediate reads are fine
 - So what's wrong? Why shouldn't we allow this schedule?
 - This schedule is view serializable, but not conflict serializable
 - Informally: a schedule is view serializable if the final written state as well as intermediate reads are the same as in some sequential schedule.
 - Formally, for those interested (this will NOT be on any exam)
 - Two schedules S and S' are view equivalent if:
 - If T_i in S reads an initial value for X, so does T_i in S'
 - If T_i in S reads the value written by T_j in S for some X, so does T_i in S'
 - If T_i in S does the final write to X, so does T_i in S'
 - A schedule is view serializable if it is view equivalent to some sequential schedule.
 - Why focus on conflict serializability when it seems too strict? Why not focus on view serializability?
 - View serializability is hard to test for.
 - Conflict serializability is not, since checking whether a graph is acyclic is fast.
 - We have an easy way to generate conflict serializable schedules.
 - Conflict serializable schedules are also view serializable, so technically this means we have an easy way to generate view serializable schedules. But we don't have an easy way to generate view schedules that allows for ones like the example above.
 - Schedules that are view serializable but not conflict serializable involve blind writes: writes that are ultimately not read. These are not common in practice.
 - Basically: conflict serializability has practical benefits
8. Summary
- Now: Have atomicity and isolation working on a single machine (concurrent transactions, good performance, etc.)
 - Next week: distributed transactions