

CSE 486/586 Distributed Systems Concurrency Control --- 2

Steve Ko
Computer Sciences and Engineering
University at Buffalo

CSE 486/586

Recap

- Question: How to support transactions (with locks)?
 - Multiple transactions share data.
- Complete serialization is correct, but performance and abort are two issues.
- Interleaving transactions for performance
 - Problem: Not all interleavings produce a correct outcome

CSE 486/586

2

Recap: Conflicting Operations

- Two operations are said to be in conflict, if their *combined effect* depends on the order they are executed, e.g., read-write, write-read, write-write (all on same variables). NOT read-read, not on different variables.

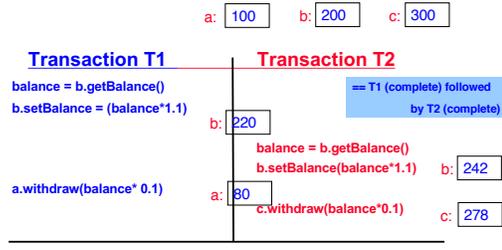
Operations of different transactions	Conflict	Reason
read read	No	Because the effect of a pair of <i>read</i> operations does not depend on the order in which they are executed
read write	Yes	Because the effect of a <i>read</i> and a <i>write</i> operation depends on the order of their execution
write write	Yes	Because the effect of a pair of <i>write</i> operations depends on the order of their execution

CSE 486/586

3

Recap: Serial Equivalence

- An interleaving of the operations of 2 or more transactions is said to be *serially equivalent* if the combined effect is the same as if these transactions had been performed sequentially (in some order).



CSE 486/586

4

Recap: Serial Equivalence

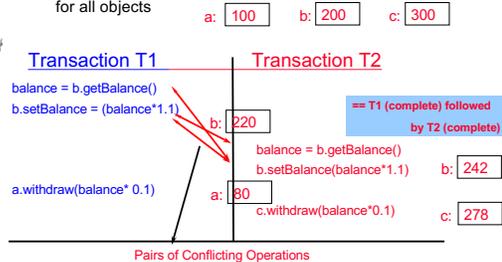
- How to provide serial equivalence with conflicting operations?
 - Execute all pairs of conflicting operations in the same order for all objects

CSE 486/586

5

Recap: Serial Equivalence

- How to provide serial equivalence with conflicting operations?
 - Execute all pairs of conflicting operations in the same order for all objects



CSE 486/586

6

Handling Abort()

- What can go wrong?

Transaction V:		Transaction W:	
<code>a.withdraw(100);</code> <code>b.deposit(100)</code>		<code>aBranch.branchTotal()</code>	
<code>a.withdraw(100);</code>	\$100	<code>total = a.getBalance()</code>	\$100
<code>b.deposit(100)</code>	\$300	<code>total = total + b.getBalance()</code> <code>total = total + c.getBalance()</code> ...	\$400

CSE 486/586

7

Strict Executions of Transactions

- Problem of interleaving for abort()
 - Intermediate state visible to other transactions, i.e., other transactions could have used some results already.
- For abort(), transactions should *delay both their read and write operations* on an object (until commit time)
 - Until all transactions that previously wrote that object have either committed or aborted
 - This way, we avoid making intermediate states visible before commit, just in case we need to abort.
 - This is called *strict executions*.
- This further restricts which interleavings of transactions are allowed.
- Thus, correctness criteria for transactions:
 - Serial equivalence
 - Strict execution

CSE 486/586

8

Story Thus Far

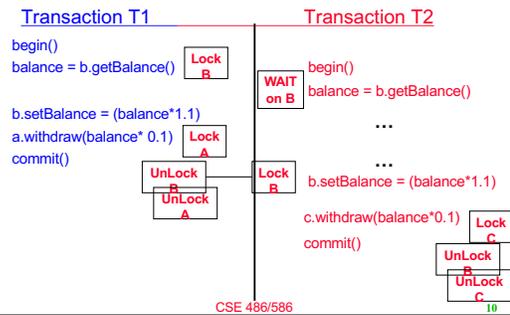
- Question: How to support transactions?
 - With multiple transactions sharing data
- First strategy: Complete serialization
 - One transaction at a time with one big lock
 - Correct, but at the cost of performance
- How to improve performance?
 - Let's see if we can interleave different transactions.
- Problem: Not all interleavings produce a correct outcome
 - Serial equivalence & strict execution must be met.
- Now, how do we meet the requirements?
 - Overall strategy: using more and more fine-grained locking
 - No silver bullet. Fine-grained locks have their own implications.

CSE 486/586

9

Using Exclusive Locks

- Exclusive Locks (Avoiding One Big Lock)

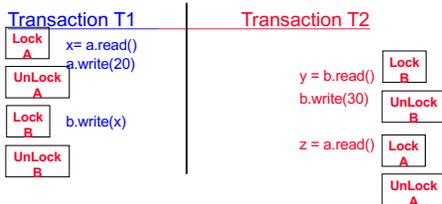


CSE 486/586

10

How to Acquire/Release Locks

- Can't do it naively



- Serially equivalent?
- Strict execution?

CSE 486/586

11

Using Exclusive Locks

- Two phase locking
 - To satisfy serial equivalence
 - First phase (growing phase): new locks are acquired
 - Second phase (shrinking phase): locks are only released
 - A transaction is not allowed to acquire any new lock, once it has released any one lock
- Strict two phase locking
 - To satisfy strict execution, i.e., to handle abort() & failures
 - Locks are only released at the end of the transaction, either at `commit()` or `abort()`, i.e., the second phase is only executed at `commit()` or `abort()`.
- The example shown before does both.

CSE 486/586

12

CSE 486/586 Administrivia

- Midterm grades will be posted today.

CSE 486/586

13

Can We Do Better?

- What we saw was “exclusive” locks.
- Non-exclusive locks: break a lock into a read lock and a write lock
- Allows more concurrency
 - Read locks can be shared (no harm to share)
 - Write locks should be exclusive

CSE 486/586

14

Non-Exclusive Locks

non-exclusive lock compatibility

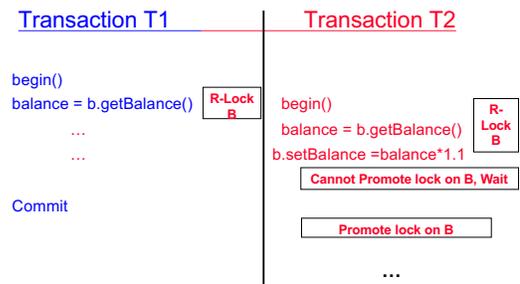
Lock already set	Lock requested	
	read	write
none	OK	OK
read	OK	WAIT
write	WAIT	WAIT

- A read lock is **promoted** to a write lock when the transaction needs write access to the same object.
- A read lock **shared** with other transactions' read lock(s) cannot be promoted. Transaction waits for other read locks to be released.
- Cannot **demote** a write lock to read lock during transaction – violates the 2P principle

CSE 486/586

15

Example: Non-Exclusive Locks

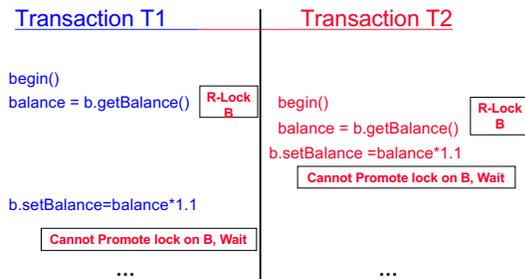


CSE 486/586

16

2PL: a Problem

- What happens in the example below?

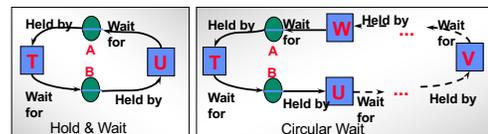


CSE 486/586

17

Deadlock Conditions

- Necessary conditions
 - Non-sharable resources (locked objects)
 - No lock preemption
 - Hold & wait or circular wait



CSE 486/586

18

Preventing Deadlocks

- Acquiring all locks at once
- Acquiring locks in a predefined order
- Not always practical:
 - Transactions might not know which locks they will need in the future
- One strategy: timeout
 - If we design each transaction to be short and fast, then we can abort() after some period of time.

CSE 486/586

19

Two-Version Locking

- Three types of locks: read lock, write lock, commit lock
 - Transaction cannot get a read or write lock if there is a commit lock
 - Read and write (from different transactions) can go together.
 - Acquiring a commit lock only happens at commit().

lock compatibility

Lock already set	Lock requested		
	read	write	commit
none	OK	OK	OK
read	OK	OK	WAIT
write	OK	WAIT	WAIT
commit	WAIT	WAIT	WAIT

CSE 486/586

20

Two-Version Locking

- Three types of locks: read lock, write lock, commit lock
 - Transaction cannot get a read or write lock if there is a commit lock
- At commit(),
 - Promote all the write locks of the transaction into commit locks
 - If any objects have outstanding read locks, transaction must wait until the transactions that set these locks have completed and locks are released
- Compare with read/write locks:
 - Read operations are delayed only while transactions are committed
 - Read operations of one transaction can cause a delay in the committing of other transactions

CSE 486/586

21

Extracting Even More Concurrency

- Allow writing *tentative versions* of objects
 - Letting other transactions read from the previously committed version
- At commit(),
 - Promote all the write locks of the transaction into commit locks
 - If any objects have outstanding read locks, transaction must wait until the transactions that set these locks have completed and locks are released
- Allow read and write locks to be set together by different transactions
 - Unlike non-exclusive locks
- Disallow commit if other uncompleted transactions have read the objects
 - These transactions must wait until the reading transactions have committed

CSE 486/586

22

Extracting Even More Concurrency

- This allows for more concurrency than read-write locks.
- Writing transactions risk waiting or rejection when commit
- Read operations wait only if another transaction is committing the same object
- Read operations of one transaction can cause a delay in the committing of other transactions

CSE 486/586

23

Summary

- Strict Execution
 - Delaying both their read and write operations on an object until all transactions that previously wrote that object have either committed or aborted
- Strict execution with exclusive locks
 - Strict 2PL
- Increasing concurrency
 - Non-exclusive locks
 - Two-version locks
 - Etc.

CSE 486/586

24

Acknowledgements

- These slides contain material developed and copyrighted by Indranil Gupta (UIUC).