

CSE 486/586 Distributed Systems Concurrency Control --- 1

Steve Ko
Computer Sciences and Engineering
University at Buffalo

CSE 486/586, Spring 2014

Banking Example (Once Again)

- Banking transaction for a customer (e.g., at ATM or browser)
 - Transfer \$100 from saving to checking account
 - Transfer \$200 from money-market to checking account
 - Withdraw \$400 from checking account
- Transaction
 1. `savings.deduct(100)`
 2. `checking.add(100)`
 3. `mnymkt.deduct(200)`
 4. `checking.add(200)`
 5. `checking.deduct(400)`
 6. `dispense(400)`

CSE 486/586, Spring 2014

2

Wait...We've Seen This Before...

- What are some things that can go wrong?
 - Multiple clients
 - Multiple servers
- How do you solve this?
 - Group everything as if it's a single step
- Where have we seen this?
 - Mutual exclusion lecture
- So, we're done?
 - No, we're not satisfied.

CSE 486/586, Spring 2014

3

Concurrent Transactions

• Process 1

```
lock(mutex);
savings.deduct(100);
checking.add(100);
mnymkt.deduct(200);
checking.add(200);
checking.deduct(400);
dispense(400);
unlock(mutex);
```


• Process 2

```
lock(mutex);
savings.deduct(200);
checking.add(200);
unlock(mutex);
```

CSE 486/586, Spring 2014

4

Why Not Satisfied?



CSE 486/586, Spring 2014

5

Recap: Locks & Transactions

- What we discussed in mutual exclusion is one big lock.
 - Everyone else has to wait.
 - It does not necessarily deal with failures.
- Performance
 - Observation: we can interleave some operations from different processes.
- Failure
 - If a process crashes while holding a lock
- Let's go beyond simple locking!

CSE 486/586, Spring 2014

6

Transaction

- Abstraction for **grouping multiple operations into one**
- A transaction is **indivisible (atomic)** from the point of view of other transactions
 - No access to intermediate results/states
 - Free from interference by other operations
- Primitives
 - begin(): begins a transaction
 - commit(): tries completing the transaction
 - abort(): aborts the transaction
- Implementing transactions
 - **Performance**: finding out what operations we can interleave
 - **Failure**: dealing with failures, rolling back changes if necessary

CSE 486/586, Spring 2014

7


Properties of Transactions: ACID

- **Atomicity**: All or nothing
- **Consistency**: if the server starts in a consistent state, the transaction ends with the server in a consistent state.
- **Isolation**: Each transaction must be performed without interference from other transactions, i.e., the non-final effects of a transaction must not be visible to other transactions.
- **Durability**: After a transaction has completed successfully, all its effects are saved in permanent storage.

CSE 486/586, Spring 2014

8

What Can Go Wrong?




- T1/T2's update on the shared object, "b", is lost

CSE 486/586, Spring 2014

9

Lost Update Problem

- One transaction causes loss of info. for another: consider three account objects




- T1/T2's update on the shared object, "b", is lost

CSE 486/586, Spring 2014

10

What Can Go Wrong?




- T1's partial result is used by T2, giving the wrong result

CSE 486/586, Spring 2014

11

Inconsistent Retrieval Problem

- Partial, incomplete results of one transaction are retrieved by another transaction.




- T1's partial result is used by T2, giving the wrong result

CSE 486/586, Spring 2014

12

What is "Correct"?

- How would you define correctness?




CSE 486/586, Spring 2014

13

Concurrency Control: Providing "Correct" Interleaving

- 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, Spring 2014

14

CSE 486/586 Administrivia


- Midterm: 3/6 (Wednesday) in class
 - Everything up to leader election
 - 1-page cheat sheet is allowed.

CSE 486/586, Spring 2014

15

Providing Serial Equivalence

- What operations are we considering?
 - Read/write
- What operations matter for correctness?
 - When write is involved



CSE 486/586, Spring 2014

16

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 read operations does not depend on the order in which they are executed
read write	Yes	Because the effect of a read and a write operation depends on the order of their execution
write write	Yes	Because the effect of a pair of write operations depends on the order of their execution

CSE 486/586, Spring 2014

17

Conditions for Correct Interleaving

- What should we need to do to guarantee serial equivalence with conflicting operations?
- Case 1
 - T1.1 -> T1.2 -> T2.1 -> T2.2 -> T1.3 -> T2.3
- Case 2
 - T1.1 -> T2.1 -> T2.2 -> T1.2 -> T1.3 -> T2.3
- Which one's correct and why?



CSE 486/586, Spring 2014

18

Conflicting Operations

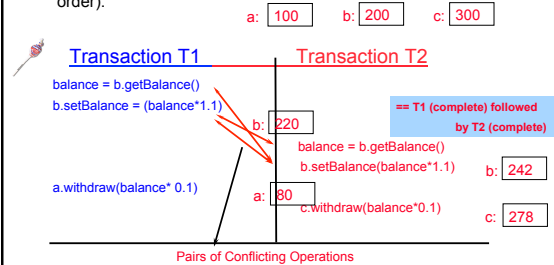
- Insight for serial equivalence
 - Outcomes of write operations in one transaction to all shared objects should be *either consistently visible to the other transaction or the other way round*.
- The effect of an operation refers to
 - The value of an object set by a write operation
 - The result returned by a read operation.
- Two transactions are *serially equivalent* if and only if *all pairs of conflicting operations* (pair containing one operation from each transaction) *are executed in the same order* (transaction order) for *all objects (data) they both access*.

CSE 486/586, Spring 2014

19

Example of Conflicting Operations

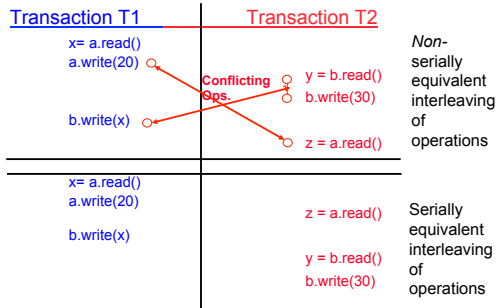
- An interleaving of the operations of 2 or more transactions is said to be *serially equivalent* if the combined effect is the same as if these transactions had been performed sequentially (in some order).



CSE 486/586, Spring 2014

20

Another Example



CSE 486/586, Spring 2014

21

Inconsistent Retrievals Problem

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

Both withdraw and deposit contain a write operation

CSE 486/586, Spring 2014

22

Serially-Equivalent Ordering

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

CSE 486/586, Spring 2014

23

Summary

- Transactions need to provide ACID
- Serial equivalence defines correctness of executing concurrent transactions
- It is handled by ordering conflicting operations

CSE 486/586, Spring 2014

24

Acknowledgements

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