

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

Steve Ko
Computer Sciences and Engineering
University at Buffalo

CSE 486/586, Spring 2013

Recap: 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);
```

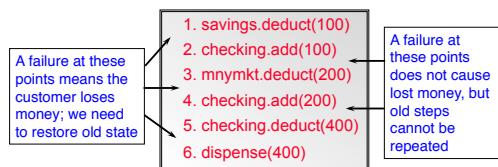
- Process 2

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

CSE 486/586, Spring 2013

2

Why Not Satisfied?



CSE 486/586, Spring 2013

3

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 2013

4

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 2013

5

Properties of Transactions: ACID

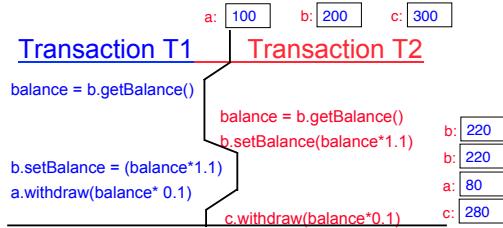
- **Atomicity:** All or nothing
- **Consistency:** if the server starts in a consistent state, the transaction ends with the server in a consistent state.
- **Isolation:** Each transaction must be performed without interference from other transactions, i.e., the non-final effects of a transaction must 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 2013

6



What Can Go Wrong?



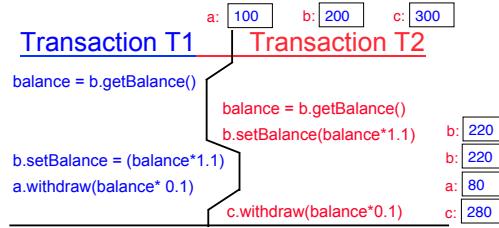
- T1/T2's update on the shared object, "b", is lost

CSE 486/586, Spring 2013

7

Lost Update Problem

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



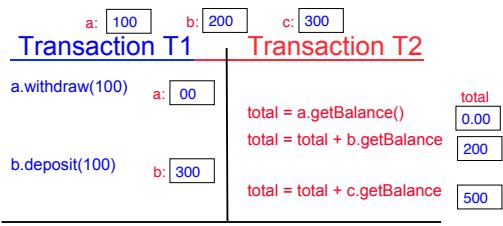
- T1/T2's update on the shared object, "b", is lost

CSE 486/586, Spring 2013

8



What Can Go Wrong?



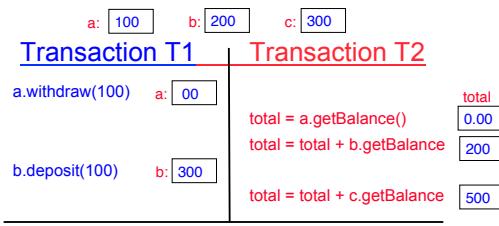
- T1's partial result is used by T2, giving the wrong result

CSE 486/586, Spring 2013

9

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

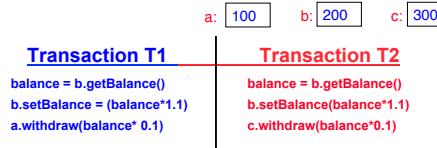
CSE 486/586, Spring 2013

10



What is "Correct"?

- How would you define correctness?

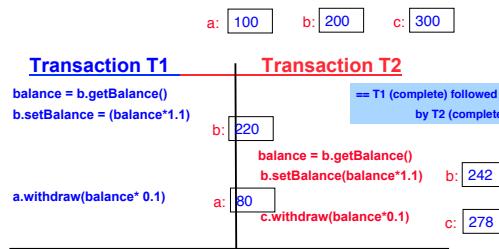


CSE 486/586, Spring 2013

11

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 2013

12

CSE 486/586 Administrivia

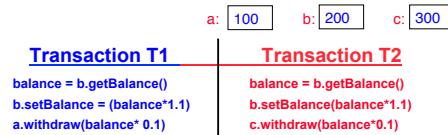
- Midterm: 3/6 (Wednesday) in class
 - 45 minutes
 - Everything up to leader election
 - 1-page cheat sheet is allowed.
- Tech Talk: Dave Parfitt (Basho) Tonight March 4 at 6PM in Davis 338A
- PA3 is out.
- No recitations this week
- Anonymous feedback form still available.
- Please come to me!

CSE 486/586, Spring 2013

13

Providing Serial Equivalence

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



CSE 486/586, Spring 2013

14

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 Conflict transactions			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, Spring 2013

15

Conditions for Correct Interleaving

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



CSE 486/586, Spring 2013

16

Conflicting Operations

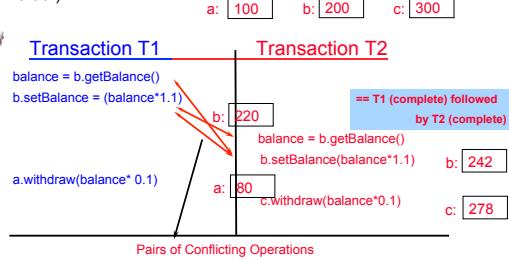
- Insight for serial equivalence
 - Outcomes of write operations in one transaction to all shared objects should be *either consistently visible to the other transaction or the other way round*.
- The effect of an operation refers to
 - The value of an object set by a write operation
 - The result returned 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 2013

17

Example of Conflicting Operations

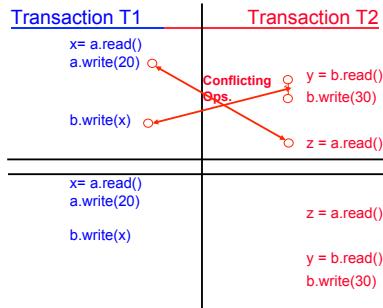
- An interleaving of the operations of 2 or more transactions is said to be *serially equivalent* if the combined effect is the same as if these transactions had been performed sequentially (in some order).



CSE 486/586, Spring 2013

18

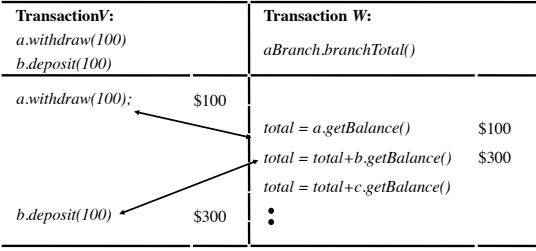
Another Example



CSE 486/586, Spring 2013

19

Inconsistent Retrievals Problem

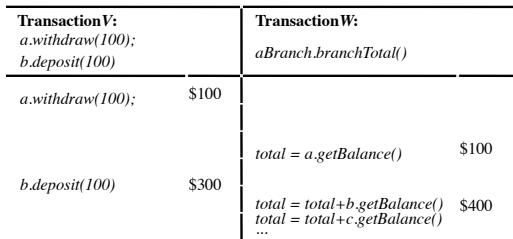


Both withdraw and deposit contain a write operation

CSE 486/586, Spring 2013

20

Serially-Equivalent Ordering



CSE 486/586, Spring 2013

21

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 2013

22

Acknowledgements

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

CSE 486/586, Spring 2013

23