

# CSE 486/586 Distributed Systems

## Concurrency Control (part 1)

Slides by Steve Ko  
Computer Sciences and Engineering  
University at Buffalo

# 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
  - `savings.deduct(100)`
  - `checking.add(100)`
  - `monymarket.deduct(200)`
  - `checking.add(200)`
  - `checking.deduct(400)`
  - `dispense(400)`

# 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 as if nothing happened
- Why `abort()`?
  - A failure happens in the middle of execution.
  - A transaction is part of a bigger transaction (*i.e.*, it's a sub-transaction), and the bigger transaction needs abort.
  - *etc.*

# Properties of Transactions: ACID

- **Atomicity**: Every transaction is all or nothing
- **Consistency**: If the system starts in a consistent state, the transaction ends with the system 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 of its effects are saved in permanent storage. (*E.g.*, powering off the machine doesn't mean the result is gone.)

# This Week

- Question: How can we support multiple transactions?
  - When multiple transactions **share data**.
  - Assume a **single processor** (one instruction at a time).
- What would be your first strategy (hint: locks)?
  - One transaction at a time with one big lock, *i.e.*, complete serialization
- Two issues
  - Performance
  - Abort

# Performance?

- Process 1

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

- Process 2

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

# Abort?

An abort at these points means the customer loses money; we need to restore old state

1. savings.deduct(100)
2. checking.add(100)
3. mnymkt.deduct(200)
4. checking.add(200)
5. checking.deduct(400)
6. dispense(400)

An abort at these points does not cause lost money, but old steps cannot be repeated

# This Week

- Question: How can we support multiple transactions?
  - Multiple transactions share data.
- What would be your first strategy (hint: locks)?
  - Complete serialization
  - One transaction at a time with one big lock
  - Two issues: **performance** and **abort**
- First, let's see how we can **improve performance**.



# Possibility: Interleaving Transactions for Performance

- Process 1

savings.deduct(100);

checking.add(100);

moneymarket.deduct(200);

checking.add(200);

checking.deduct(400);

dispense(400);

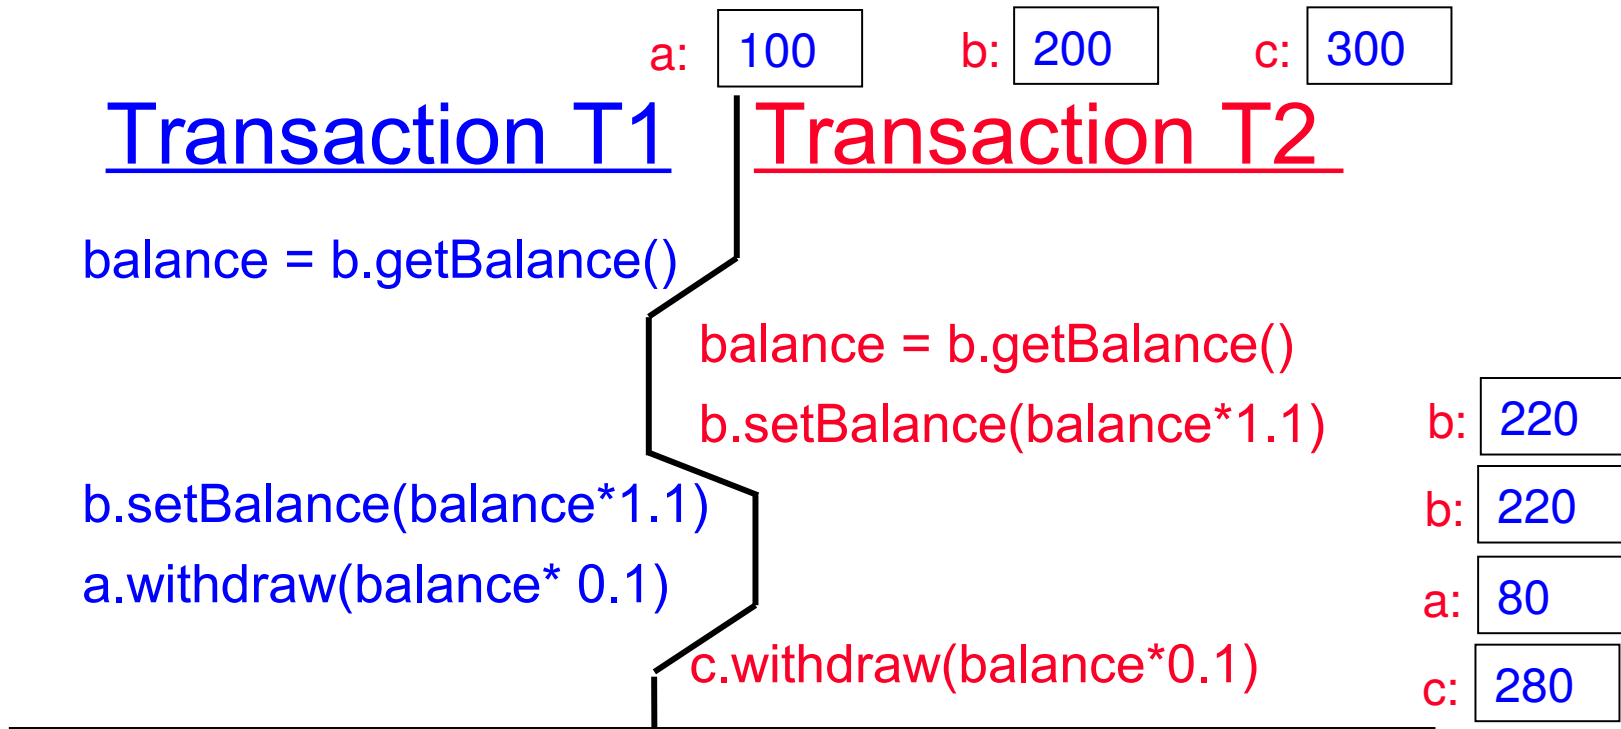
- Process 2

savings.deduct(200);

checking.add(200);

- P2 will not have to wait until P1 finishes.

# What Can Go Wrong?



- T1/T2's update on the shared object **b** is lost
- This is the **lost update problem**.

# What Can Go Wrong?

## Transaction T1

## Transaction T2

a:

100

b:

200

c:

300

a.withdraw(100)

a:

00

b.deposit(100)

b:

300

total = a.getBalance()

total = total + b.getBalance

total = total + c.getBalance

total

0.00

200

500

- T1's **partial result** is used by T2, giving the wrong result
- This is the **inconsistent retrieval problem**

# What This Means

- How do we support transactions using locks?
  - Interesting only when multiple transactions have shared data.
- Complete serialization is correct, but **performance** and **abort handling** are two issues.
- Concurrent execution of transactions would improve performance
  - But **not all interleavings** produce a correct outcome!

# What is “Correct”?

- How would you define correctness?
- For example, two independent transactions made by me and my wife on our three accounts.
- What do we care about **at the end of the day**?
  - Correct **final balance** for each account

## Transaction T1

```
balance = b.getBalance()  
b.setBalance = (balance*1.1)  
a.withdraw(balance* 0.1)
```

## Transaction T2

```
balance = b.getBalance()  
b.setBalance(balance*1.1)  
c.withdraw(balance*0.1)
```

a:

100

b:

200

c:

300

# 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**.

a: 100      b: 200      c: 300

## Transaction T1

balance = b.getBalance()  
b.setBalance = (balance\*1.1)

a.withdraw(balance\* 0.1)

## Transaction T2

b: 220

balance = b.getBalance()  
b.setBalance(balance\*1.1)

a: 80

c.withdraw(balance\*0.1)

== T1 (complete) followed  
by T2 (complete)

b: 242

c: 278

# Providing Serial Equivalence

- What operations are we considering?
  - Read/write
- What operations matter for correctness?
  - Operations on any state that is written at any time

a: 100      b: 200      c: 300

## Transaction T1

balance = b.getBalance()  
b.setBalance = (balance\*1.1)  
a.withdraw(balance\* 0.1)

## Transaction T2

balance = b.getBalance()  
b.setBalance(balance\*1.1)  
c.withdraw(balance\*0.1)

# Conflicting Operations

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

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



# Conditions for Correct Interleaving

- What do 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 is correct and why?

## Transaction T1

1. `balance = b.getBalance()`
2. `b.setBalance = (balance*1.1)`
3. `a.withdraw(balance* 0.1)`

## Transaction T2

1. `balance = b.getBalance()`
2. `b.setBalance(balance*1.1)`
3. `c.withdraw(balance*0.1)`

# Observation

- Case 1
  - T1.1 -> T1.2 -> T2.1 -> T2.2 -> T1.3 -> T2.3
- Correct: for a shared object (b), T1 produces its final outcome, and then T2 accesses it.

## Transaction T1

1. `balance = b.getBalance()`
2. `b.setBalance = (balance*1.1)`
3. `a.withdraw(balance* 0.1)`

## Transaction T2

1. `balance = b.getBalance()`
2. `b.setBalance(balance*1.1)`
3. `c.withdraw(balance*0.1)`

# Observation

- Case 2
  - T1.1 -> T2.1 -> T2.2 -> T1.2 -> T1.3 -> T2.3
- Incorrect: for a shared object (b), T2 uses T1's intermediate outcome, which T1 later modifies.

## Transaction T1

1. `balance = b.getBalance()`
2. `b.setBalance = (balance*1.1)`
3. `a.withdraw(balance* 0.1)`

## Transaction T2

1. `balance = b.getBalance()`
2. `b.setBalance(balance*1.1)`
3. `c.withdraw(balance*0.1)`

# Generalizing the Observation

- Insight for serial equivalence
  - Only the **final outcome of a shared object** from one transaction should be visible to another transaction.
  - The above should be the case for **each and every shared object in the same order**.
  - E.g., if T1's final outcome **on one shared object** becomes visible to T2, then **for each and every other shared object**, T1 should produce the final outcome **before T2 uses it**.
  - The other way round is also correct, *i.e.*, T2 first then T1.
- What is it called when one transaction modifies a shared object and another transaction uses it?
  - **Conflicting operations**

# Serial Equivalence and Conflicting Operations

- 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**.

# Serial Equivalence Example

- 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).

a: 100      b: 200      c: 300

## Transaction T1

balance = b.getBalance()  
b.setBalance = (balance\*1.1)

a.withdraw(balance\* 0.1)

## Transaction T2

b: 220

balance = b.getBalance()

b.setBalance(balance\*1.1)

a: 80

c.withdraw(balance\*0.1)

== T1 (complete) followed  
by T2 (complete)

b: 242

c: 278

Pairs of Conflicting Operations

# Another Example

## Transaction T1

x = a.read()

a.write(20)

b.write(x)

**Conflicting  
Ops.**

## Transaction T2

y = b.read()

b.write(30)

z = a.read()

*Non-serially  
equivalent  
interleaving  
of operations*

x = a.read()

a.write(20)

b.write(x)

z = a.read()

y = b.read()

b.write(30)

*Serially  
equivalent  
interleaving  
of operations*

# Inconsistent Retrievals Problem

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

Both withdraw and deposit contain a write operation



# Serially-Equivalent Ordering

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

# Summary

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

# References

- Textbook sections 16.1-16.4. **Required Reading.**

# Acknowledgements

- These slides originally by Steve Ko, lightly modified and used with permission by Ethan Blanton.
- These slides contain material developed and copyrighted by Indranil Gupta (UIUC).