

Computer Science 425

Distributed Systems

CS 425 / CSE 424 / ECE 428

Fall 2012

Indranil Gupta (Indy)

October 18, 2012

Lecture 16

Concurrency Control

Reading: Chapter 16.1,2,4 and 17.1,2,3,5 (relevant parts)

Transactions



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 (invoked at client):

1. `savings.withdraw(100)`

2. `checking.deposit(100)`

3. `mnymkt.withdraw(200)`

4. `checking.deposit(200)`

5. `checking.withdraw(400)`

6. `dispense(400)`

7. `commit`

/ Every step is an RPC */*

/ includes verification */*

/ depends on success of 1 */*

/ includes verification */*

/ depends on success of 3 */*

/ includes verification */*

Bank Server: Coordinator Interface

- ❖ All the following are RPCs from a client to the server
- ❖ Transaction calls that can be made at a client, and return values from the server:

openTransaction() -> *trans*;

starts a new transaction and delivers a unique transaction identifier (TID) *trans*. This TID will be used in the other operations in the transaction.

closeTransaction(trans) -> (*commit*, *abort*);

ends a transaction: a *commit* return value indicates that the transaction has committed; an *abort* return value indicates that it has aborted.

abortTransaction(trans);

aborts the transaction.

- ❖ TID can be passed implicitly (for other operations between open and close) with CORBA

Bank Server: Account, Branch interfaces

Operations of the Account interface

deposit(amount)

deposit amount in the account

withdraw(amount)

withdraw amount from the account

getBalance() -> *amount*

return the balance of the account

setBalance(amount)

set the balance of the account to amount

Operations of the Branch interface

create(name) -> *account*

create a new account with a given name

lookup(name) -> *account*

return a reference to the account with the given name

branchTotal() -> *amount*

return the total of all the balances at the branch

Transaction

- ❖ Sequence of operations that forms a single step, transforming the server data from one consistent state to another.
 - ❑ All or nothing principle: a transaction either completes successfully, and the effects are recorded in the objects, or it has no effect at all. (even with multiple clients, or crashes)
- ❖ A transaction is indivisible (atomic) from the point of view of other transactions
 - ❖ No access to intermediate results/states of other transactions
 - ❖ Free from interference by operations of other transactions

But...

- ❖ Transactions could run concurrently, i.e., with multiple clients
- ❖ Transactions may be distributed, i.e., across multiple servers

Transaction Failure Modes

Transaction:

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

A failure at these points means the customer loses money; we need to restore old state

A failure at these points does not cause lost money, but old steps cannot be repeated

This is the point of no return

A failure after the commit point (ATM crashes) needs corrective action; no undoing possible.

Transactions in Traditional Databases (ACID)

- ❖ **A**tomicity: All or nothing
 - ❖ **C**onsistency: if the server starts in a consistent state, the transaction ends the server in a consistent state.
 - ❖ **I**solation: 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.
 - ❖ **D**urability: After a transaction has completed successfully, all its effects are saved in permanent storage.
-
- ❖ **Atomicity**: store tentative object updates (for later undo/redo) – many different ways of doing this
 - ❖ **Durability**: store entire results of transactions (all updated objects) to recover from permanent server crashes.

Concurrent Transactions: Lost Update Problem

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

a: 100 b: 200 c: 300

Transaction T1

Transaction T2

balance = b.getBalance()

balance = b.getBalance()

b.setBalance(balance*1.1)

b: 220

b.setBalance(balance*1.1)

b: 220

a.withdraw(balance*0.1)

a: 80

c.withdraw(balance*0.1)

c: 280

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

Conc. Trans.: Inconsistent Retrieval Prob.

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

a: 100 b: 200 c: 300

Transaction T1

a.withdraw(100)

a: 00

b.deposit(100)

b: 300

Transaction T2

total = a.getBalance()

total

0.00

total = total + b.getBalance

200

total = total + c.getBalance

500

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

Concurrency Control: “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).

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

Checking Serial Equivalence – Conflicting Operations

- ❑ 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 operations are said to be conflicting operations, 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.
- ❑ 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.*
 - ❑ *Why? Can start from original operation sequence and swap the order of non-conflicting operations to obtain a series of operations where one transaction finishes completely before the second transaction starts*
- ❑ Why is the above result important? Because: **Serial equivalence is the basis for concurrency control protocols for transactions.**

Read *and* Write Operation Conflict Rules

Operations of different transactions			Conflict	Reason
<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

Concurrency Control: “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).

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

Conflicting Operators Example

Transaction T1

`x = a.read()`

`a.write(20)`

`b.write(x)`

Transaction T2

`y = b.read()`

`b.write(30)`

`z = a.read()`

Conflicting
Ops.

`x = a.read()`

`a.write(20)`

`b.write(x)`

`z = a.read()`

`y = b.read()`

`b.write(30)`

Non-serially equivalent interleaving of operations

Serially equivalent interleaving of operations (why?)

Inconsistent Retrieval Prob – Caught!

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

a: 100 b: 200 c: 300

Transaction T1

Transaction T2

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 for T2

A Serially Equivalent Interleaving of T1 and T2

Transaction <i>T1</i>		Transaction <i>T2</i>	
<i>a.withdraw(100);</i> <i>b.deposit(100)</i>		<i>aBranch.branchTotal()</i>	
<i>a.withdraw(100);</i>	\$100		
<i>b.deposit(100)</i>	\$300	<i>total = a.getBalance()</i>	\$100
		<i>total = total+b.getBalance()</i>	\$400
		<i>total = total+c.getBalance()</i>	

Implementing Concurrent Transactions

- ♣ How can we prevent isolation from being violated?
- ♣ Concurrent operations must be consistent:
 - ♣ If trans. T has executed a *read* operation on object A, a concurrent trans. U must not *write* to A until T commits or aborts.
 - ♣ If trans. T has executed a *write* operation on object A, a concurrent U must not *read or write* to A until T commits or aborts.
- ♣ How to implement this?
 - ♣ First cut: locks

Example: Concurrent Transactions

❖ Exclusive Locks

Transaction T1

OpenTransaction()

balance = b.getBalance()

Lock
B

b.setBalance(balance*1.1)

a.withdraw(balance* 0.1)

Lock
A

CloseTransaction()

UnLock
B

UnLock
A

Transaction T2

OpenTransaction()

balance = b.getBalance()

WAIT
on B

...

...

b.setBalance(balance*1.1)

c.withdraw(balance*0.1)

Lock
C

CloseTransaction()

UnLock
B

UnLock
C

Basic Locking

- ♣ Transaction managers (on server side) set locks on objects they need. A concurrent trans. cannot access locked objects.
- ♣ **Two phase locking:**
 - ♣ In the first (growing) phase of the transaction, new locks are only acquired, and in the second (shrinking) phase, locks are only released.
 - ♣ A transaction is not allowed acquire *any* new locks, once it has released any one lock.
- ♣ **Strict two phase locking:**
 - ♣ Locking on an object is performed only before the first request to read/write that object is about to be applied.
 - ♣ Unlocking is performed by the commit/abort operations of the transaction coordinator.
 - ♣ To prevent dirty reads and premature writes, a transaction waits for another to commit/abort
- ♣ However, use of separate **read** and **write** locks leads to more concurrency than a single **exclusive** lock – Next slide

2P Locking: Non-exclusive lock (per object)

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

Locking Procedure in Strict-2P Locking

♣ When an operation accesses an object:

- ◆ if you can promote a lock (nothing -> read -> write)
- ◆ Don't promote the lock if it would result in a conflict with another transaction's already-existing lock
 - ◆ wait until all shared locks are released, then lock & proceed

♣ When a transaction commits or aborts:

- release all locks that were set by the transaction

Example: Concurrent Transactions

❖ Non-exclusive Locks

Transaction T1

OpenTransaction()

balance = b.getBalance()

R-Lock
B

Commit

Transaction T2

OpenTransaction()

balance = b.getBalance()

b.setBalance(balance*1.1)

R-Lock
B

Cannot Promote lock on B, Wait

Promote lock on B

...

Example: Concurrent Transactions

❖ What happens in the example below?

Transaction T1

OpenTransaction()

balance = b.getBalance()

R-Lock
B

b.setBalance=balance*1.1

Cannot Promote lock on B, Wait

...

Transaction T2

OpenTransaction()

balance = b.getBalance()

b.setBalance(balance*1.1)

R-
Lock
B

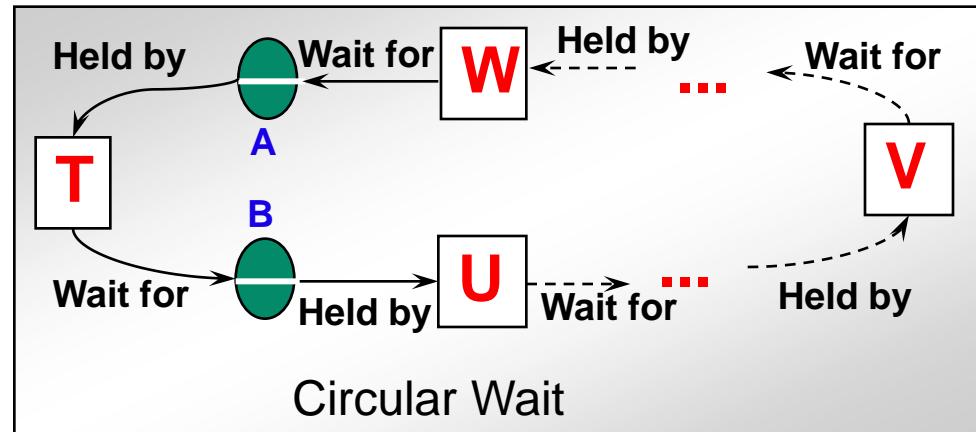
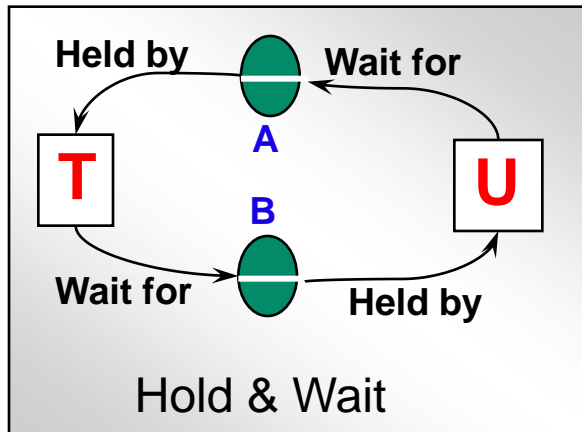
Cannot Promote lock on B, Wait

...

Deadlocks

❖ Necessary conditions for deadlocks

- ❑ Non-shareable resources (exclusive lock modes)
- ❑ No preemption on locks
- ❑ Hold & Wait or Circular Wait



Naïve Deadlock Resolution Using Timeout

Transaction T		Transaction U	
Operations	Locks	Operations	Locks
<i>a.deposit(100);</i>	write lock A		
<i>b.withdraw(100)</i>		<i>b.deposit(200)</i>	write lock B
• • •	waits for U 's lock on B (timeout elapses)	<i>a.withdraw(200);</i>	waits for T's lock on A
	T 's lock on A becomes vulnerable, unlock A , abort T	• • •	
		• • •	
		<i>a.withdraw(200);</i>	write locks A unlock A B

Disadvantages?

Lecture 16-25

Strategies to Fight Deadlock

- ❑ Lock timeout (costly and open to false positives)
- ❑ Deadlock **Prevention**: violate one of the necessary conditions for deadlock (from 2 slides ago), e.g., lock all objects before transaction starts, aborting if any fails
- ❑ Deadlock **Avoidance**: Have transactions declare max resources they will request, but allow them to lock at any time (Banker's algorithm)
- ❑ Deadlock **Detection**: detect cycles in the wait-for graph, and then abort one or more of the transactions in cycle

Summary

- **Increasing concurrency important because it improves throughput at server (means more revenue \$\$\$)**
- **Applications are willing to tolerate temporary inconsistency and deadlocks in turn**
 - Need to detect and prevent these
- **Driven and validated by actual application characteristics – mostly-read transactions abound**

Midterm Statistics

- Graduate:
 - Count: 34
 - **Avg: 87.06**, StdDev: 11.3603275
 - Median: 91
 - **Max: 99 Min: 45**
- Undergraduate:
 - Count: 42
 - **Avg: 83.64**, StdDev: 10.79085301
 - Median: 87.75
 - **Max: 100 Min: 55**