

Transactions

Distributed Software
Systems

Transactions

- Motivation

- Provide atomic operations at servers that maintain shared data for clients
- Provide recoverability from server crashes

- Properties

- Atomicity, Consistency, Isolation, Durability (ACID)

- Concepts: commit, abort

Concurrency control

- Motivation: without concurrency control, we have lost updates, inconsistent retrievals, dirty reads, etc. (see following slides)
- Concurrency control schemes are designed to allow two or more transactions to be executed correctly while maintaining serial equivalence
 - Serial Equivalence is correctness criterion
 - Schedule produced by concurrency control scheme should be equivalent to a serial schedule in which transactions are executed one after the other
- Schemes: locking, optimistic concurrency control, time-stamp based concurrency control

Figure 12.1 A client's banking transaction

Transaction:T:
Bank\$Withdraw(A, 100);
Bank\$Deposit(B, 100);
Bank\$Withdraw(C, 200);
Bank\$Deposit(B, 200);

Figure 12.2 Transactional service operations.

OpenTransaction \rightarrow *Trans*

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

CloseTransaction(Trans) \rightarrow (*Commit*, *Abort*)

ends a transaction: a *Commit* returned value indicates that the transaction has committed; an *Abort* returned value indicates that it has aborted.

AbortTransaction(Trans)

aborts the transaction.

Figure 12.3 Transaction life histories.

<i>Successful</i>	<i>Aborted by client</i>	<i>Aborted by server</i>
<i>OpenTransaction</i> <i>operation</i> <i>operation</i> <ul style="list-style-type: none"> • • <i>operation</i>	<i>OpenTransaction</i> <i>operation</i> <i>operation</i> <ul style="list-style-type: none"> • • <i>operation</i>	<i>OpenTransaction</i> <i>operation</i> <i>operation</i> <ul style="list-style-type: none"> • • <i>operation ERROR</i> <i>reported to client</i>
<i>CloseTransaction</i>	<i>AbortTransaction</i>	<p style="text-align: center;">SERVER ABORTS →</p>

Figure 12.4 The lost update problem.

Transaction T:		Transaction U:	
<i>Bank\$Withdraw(A, 4);</i>		<i>Bank\$Withdraw(C, 3);</i>	
<i>Bank\$Deposit(B, 4)</i>		<i>Bank\$Deposit(B, 3)</i>	
<i>balance := A.Read()</i>	\$100	<i>balance := C.Read()</i>	\$300
<i>A.Write (balance - 4)</i>	\$96	<i>C.Write (balance - 3)</i>	\$297
<i>balance := B.Read()</i>	\$200	<i>balance := B.Read()</i>	\$200
<i>B.Write (balance + 4)</i>	\$204	<i>B.Write (balance + 3)</i>	\$203

Figure 12.5 The inconsistent retrievals problem.

Transaction T:		Transaction U:	
<i>Bank\$Withdraw(A, 100);</i>		<i>Bank\$BranchTotal()</i>	
<i>Bank\$Deposit(B, 100)</i>			
<i>balance := A.Read()</i>	\$200	<i>balance := A.Read()</i>	\$100
<i>A.Write (balance - 100)</i>	\$100	<i>balance := balance + B.Read()</i>	\$300
		<i>balance := balance + C.Read()</i>	\$300+.
<i>balance := B.Read()</i>	\$200	•	
<i>B.Write (balance + 100)</i>	\$300	•	

Figure 12.6 A serially equivalent interleaving of T and U.

Transaction T:		Transaction U:	
<i>Bank\$Withdraw(A, 4);</i>		<i>Bank\$Withdraw(C, 3);</i>	
<i>Bank\$Deposit(B, 4)</i>		<i>Bank\$Deposit(B, 3)</i>	
<i>balance := A.Read()</i>	\$100	<i>balance := C.Read()</i>	\$300
<i>A.Write(balance - 4)</i>	\$96	<i>C.Write(balance - 3)</i>	\$297
<i>balance := B.Read()</i>	\$200	<i>balance := B.Read()</i>	\$204
<i>B.Write (balance + 4)</i>	\$204	<i>B.Write(balance + 3)</i>	\$207

Figure 12.7 An inconsistent retrievals solution

Transaction T:		Transaction U:	
<i>Bank\$Withdraw(A, 100);</i>		<i>Bank\$BranchTotal()</i>	
<i>Bank\$Deposit(B, 100)</i>			
<i>balance := A.Read()</i>	\$200		
<i>A.Write(balance - 100)</i>	\$100		
<i>balance := B.Read()</i>	\$200		
<i>B.Write(balance + 100)</i>	\$300		
		<i>balance := A.Read()</i>	\$100
		<i>balance := balance + B.Read()</i>	\$400
		<i>balance := balance + C.Read()</i>	\$400+
		...	

Figure 12.8 A dirty read when transaction T aborts.

Transaction T: <i>Bank.\$Deposit(A, 3)</i>	Transaction U: <i>Bank.\$Deposit(A, 5)</i>
<i>balance := A.Read()</i> \$100	
<i>A.Write(balance + 3)</i> \$103	
	<i>balance := A.Read()</i> \$103
	<i>A.Write (balance + 5)</i> \$108
	<i>Commit transaction</i>
<i>Abort transaction</i>	

Figure 12.9 Over-writing uncommitted values.

Transaction T: <i>Bank\$Deposit(A, 3)</i>		Transaction U: <i>Bank\$Deposit(A, 5)</i>	
<i>balance := A.Read()</i>	\$100		
<i>A.Write(balance + 3)</i>	\$103		
		<i>balance := A.Read()</i>	\$103
		<i>A.Write (balance + 5)</i>	\$108
		<i>Abort transaction</i>	

Figure 12.11 Nested *Transfer* transaction.



