

Many of the slides through slide 21are figures from Coulouris, Dollimore, Kindberg, and Blair.





Figure 16.1. A pair of interfaces to be used in upcoming examples.



Figure 16.2: A sample banking transaction using the account interface.

Operations in	s in Co nterfac	ordinator e
openTransaction() -> ti starts a new transa trans. This identifier operations in the tra	rans; ction and o r will be us ansaction.	delivers a unique TID sed in the other
<i>closeTransaction(trans</i> ends a transaction: that the transaction value indicates that	s) -> (comr a commit has com t it has abo	<i>mit, abort);</i> return value indicates mitted; an <i>abort</i> return orted.
abortTransaction(trans aborts the transacti	;) <i>;</i> ion.	
CS 138	XVIII–5	Copyright © 2016 Thomas W. Doeppner. All rights reserved.

Figure 16.3

Successiui	Aborted by client	Aborte	ed by server
openTransaction	openTransaction		openTransaction
operation	operation		operation
operation	operation		operation
•	:	server aborts transaction	→ :
operation	operation		operation ERROR
closeTransaction	abortTransaction		

Figure 16.4

nitial balances: a: \$100, b:	\$200, c: \$300
Transaction T:	Transaction <i>U</i> :
balance = b.getBalance(); b.setBalance(balance*1.1); a.withdraw(balance/10)	<pre>balance = b.getBalance(); b.setBalance(balance*1.1); c.withdraw(balance/10)</pre>
balance = b.getBalance(); \$200	
	balance = b.getBalance(); \$200
	b.setBalance(balance*1.1); \$220
b.setBalance(balance*1.1); \$220	
a.withdraw(balance/10) \$80	
	c.withdraw(balance/10) \$280

Figure 16.5

The in	consi pr	stent retrieval oblem	S
Transaction <i>V:</i>		Transaction W:	
a.withdraw(100) b.deposit(100)		aBranch.branchTotal()	
a.withdraw(100);	\$100	total = a.getBalance() total = total+b.getBalance total = total+c.getBalance	\$1 e() \$3 e()
o.deposit(100)	\$300	:	
		XVIII–8 Copyright © 2016 Thomas V	V. Doeppne

Figure 16.6



Transaction <i>T</i> :	Transaction <i>U</i> :
balance = b.getBalance()	balance = b.getBalance()
b.setBalance(balance*1.1) a.withdraw(balance/10)	b.setBalance(balance*1.1) c.withdraw(balance/10)
balance = b.getBalance() \$200 b.setBalance(balance*1.1)\$220	
	balance = b.getBalance() \$220
a.withdraw(balance/10) \$80	b.setBalance(balance*1.1)\$242
	c withdraw(balance/10) \$278

Figure 16.7

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

Figure 16.8

<i>Read</i> and <i>write</i> operation conflict rules			
Operations transa	of different ctions	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 <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
 S 138			XVIII—12 Copyright © 2016 Thomas W. Doeppner. All rights res

Figure 16.9

	Serially E	Equivalent?	
	Transaction R:	Transaction S:	
	x = read(i) write(i, 10)	y = read(j) write(j, 30)	
	write(j, 20)	z = read (i)	
Rule: equiv exect	Two transaction ex alent iff all conflictin ited in the same orc	ecutions are seria g operations of the ler.	lly e two are
CS 138	:	XVIII–13 Copyright © 2016 Th	nomas W. Doeppner. All rights reserved.

Figure 16.10: Note that all of R's accesses to i come before S's, and all of S's accesses to j come before R's. Nevertheless this is not a serially equivalent concurrent execution of R and S.



Transaction U: a.getBalance() a.setBalance(balance + 20)
a.getBalance() a setBalance(balance + 20)
balance = a.getBalance() \$110 a.setBalance(balance + 20) \$130 commit transaction

Figure 16.11: The isolation property requires that transactions do not see the uncommitted state of other transactions.

Overwi	riting u val	uncommitted ues	
Transaction <i>T</i> :		Transaction <i>U</i> :	
a.setBalance(105)		a.setBalance(110)	
a.setBalance(105)	\$100 \$105		
		a.setBalance(110)	\$110
		abort transaction	

Figure 16.12: Suppose U aborts after overwriting A.

Transactions <i>T</i> and <i>U</i> with exclusive locks				
Transaction T: balance = b.getBal b.setBalance(bal*1 a.withdraw(bal/10)	ance() .1)	Transaction U: balance = b.getBala b.setBalance(bal*1. c.withdraw(bal/10)	nce() 1)	
Operations	Locks	Operations	Locks	
openTransaction bal = b.getBalance() b.setBalance(bal*1.1) a.withdraw(bal/10)	lock B lock A	openTransaction bal = b.getBalance(() waits for T	's lock on I
closeTransaction	unlock A, B	•••		
			lock B	
		b.setBalance(bal*1. c.withdraw(bal/10) closeTransaction	1) lock C unlock B,0	C
8		XVIII–17 Copyright @	© 2016 Thomas W. Doep	pner. All rights re

Figure 16.14





Transaction A completes before transaction B. Two-phase locking was used. Is it necessarily the case that their concurrent execution is equivalent to first executing all of A, then all of B?

	Transaction Steps
•	Accumulate changes – store as "tentative versions" Make sure everything is ok Commit or abort – move tentative versions to actual
	or – delete tentative versions
CS 138	XVIII—20 Copyright © 2016 Thomas W. Doeppner. All rights reserve



Figure 16.13



Much of the remainder of this lecture is adapted from the textbook by Tanenbaum and Van Steen and from Chapter 7 of *Concurrency Control and Recovery in Database Systems*, by P. Bernstein, V. Hadzilacos, and N. Goodman, Addison-Wesley (1987). The latter text is available at http://research.microsoft.com/en-us/people/philbe/ccontrol.aspx.





This is adapted from Bernstein et al.





The labeling of the arcs (A/B) means that if "A" occurs, then perform action "B" and follow the arc to the next state.









This implies that the participants know one another's identities. They could be supplied by the coordinator in the initial vote request.







Note that NB does not hold for two-phase commit!



For our upcoming discussion of three-phase commit, we assume that the only sort of failure is that of a machine crashing (then recovering). In particular, communication failures do not happen.







If a participant times out in its init state while waiting for a vote req from the coordinator, it may safely unilaterally abort.



If the coordinator times out in its wait state while waiting to receive votes from participants, it should send aborts to all operational participants.



If the coordinator times out while in its precommit state, waiting to receive acks from the participants, it may safely commit, since it had received commit votes from all. The failed participants will learn about the commit when they reboot.



If a participant times out in its uncertain state waiting to hear from the coordinator, it must communicate with the other operational participants to determine if it should commit or abort. In particular, if any other participant has aborted, it should abort. But what if this is not the case? (Go on to the next slides ...)



If a participant times out in its precommit state, waiting to hear from the coordinator, shouldn't it simply assume it may commit? The answer is no, because that might violate NB: there may be some other participant that's still in the uncertain state.

The situation we're concerned about is that, after committing, the participant might fail, while some other participant (perhaps the only other participant) remains operational, but in the uncertain state. That participant, now not knowing anything about the states of the others, should be allowed to abort by virtue of NB.





Note that the newly elected coordinator could fail. If so, a new one is elected. (Participants will time-out waiting for a message.)





For details, see chapter 7 from Bernstein et al.





See Bernstein et al. for details.



See Bernstein et al. for details.