

Figure 12.1 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

Figure 12.2 A client's banking transaction

Transaction T:
a.withdraw(100);
b.deposit(100);
c.withdraw(200);
b.deposit(200);

Figure 12.3 Operations in *Coordinator* interface

openTransaction() → *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) → (*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.

Figure 12.4 Transaction life histories

<i>Successful</i>	<i>Aborted by client</i>	<i>Aborted by server</i>
<i>openTransaction</i>	<i>openTransaction</i>	<i>openTransaction</i>
<i>operation</i>	<i>operation</i>	<i>operation</i>
<i>operation</i>	<i>operation</i>	<i>operation</i>
•	•	server aborts
•	•	transaction →
<i>operation</i>	<i>operation</i>	•
		•
		<i>operation ERROR</i>
		<i>reported to client</i>
<i>closeTransaction</i>	<i>abortTransaction</i>	

Figure 12.5 The lost update problem

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

Figure 12.6 The inconsistent retrievals problem

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

Figure 12.7 A serially equivalent interleaving of *T* and *U*

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

Figure 12.8 A serially equivalent interleaving of V and W

Transaction V:		Transaction W:	
<i>a.withdraw(100);</i>		<i>aBranch.branchTotal()</i>	
<i>b.deposit(100)</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>	
		...	

Figure 12.9 *Read and write operation conflict rules*

<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

Figure 12.10 A non-serially equivalent interleaving of operations of transactions T and U

Transaction T:	Transaction U:
$x = read(i)$ $write(i, 10)$	$y = read(j)$ $write(j, 30)$
$write(j, 20)$	$z = read(i)$

Figure 12.11 A dirty read when transaction *T* aborts

Transaction <i>T</i>:	Transaction <i>U</i>:
<i>a.getBalance()</i>	<i>a.getBalance()</i>
<i>a.setBalance(balance + 10)</i>	<i>a.setBalance(balance + 20)</i>
<i>balance = a.getBalance()</i> \$100	
<i>a.setBalance(balance + 10)</i> \$110	
	<i>balance = a.getBalance()</i> \$110
	<i>a.setBalance(balance + 20)</i> \$130
	<i>commit transaction</i>
<i>abort transaction</i>	

Figure 12.12 Overwriting uncommitted values

Transaction T:		Transaction U:	
<i>a.setBalance(105)</i>		<i>a.setBalance(110)</i>	
	\$100		
<i>a.setBalance(105)</i>	\$105	<i>a.setBalance(110)</i>	\$110

Figure 12.13 Nested transactions

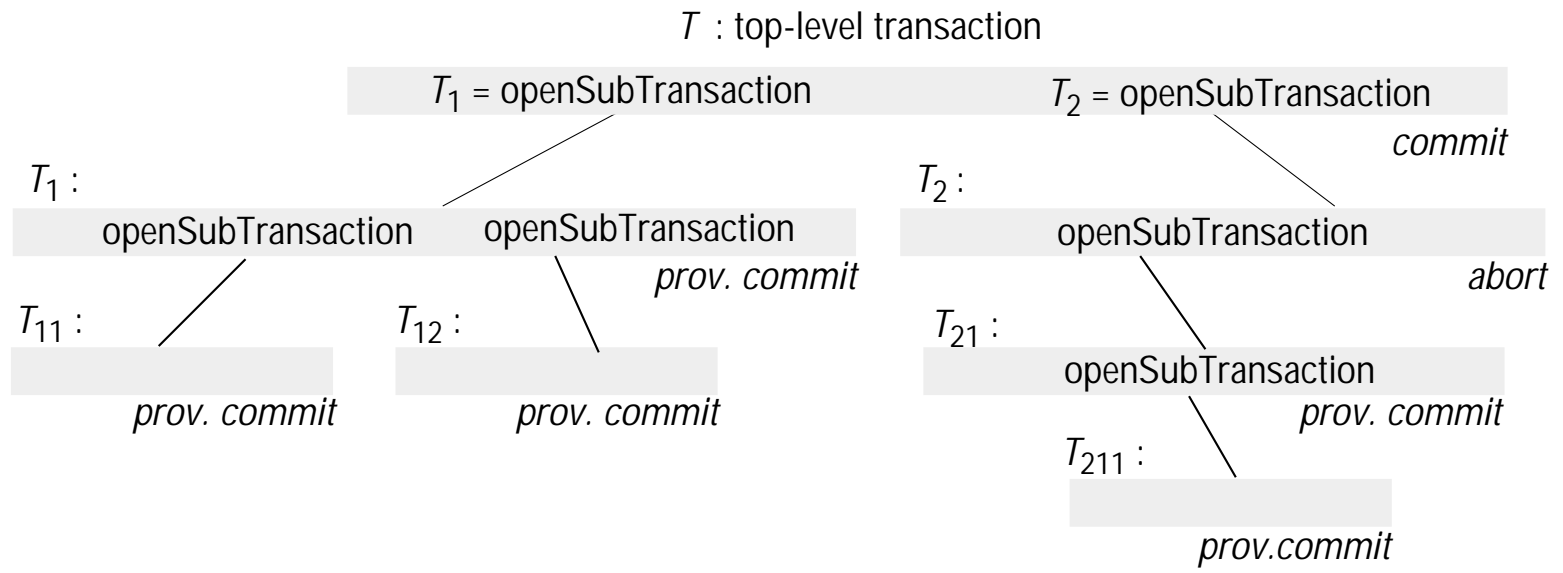


Figure 12.14 Transactions *T* and *U* with exclusive locks

Transaction <i>T</i>:		Transaction <i>U</i>:	
<i>balance = b.getBalance()</i>		<i>balance = b.getBalance()</i>	
<i>b.setBalance(bal*1.1)</i>		<i>b.setBalance(bal*1.1)</i>	
<i>a.withdraw(bal/10)</i>		<i>c.withdraw(bal/10)</i>	
Operations	Locks	Operations	Locks
<i>openTransaction</i>		<i>openTransaction</i>	
<i>bal = b.getBalance()</i>	lock <i>B</i>	<i>bal = b.getBalance()</i>	waits for <i>T</i> 's lock on <i>B</i>
<i>b.setBalance(bal*1.1)</i>		...	
<i>a.withdraw(bal/10)</i>	lock <i>A</i>		lock <i>B</i>
<i>closeTransaction</i>	unlock <i>A, B</i>	<i>b.setBalance(bal*1.1)</i>	
		<i>c.withdraw(bal/10)</i>	lock <i>C</i>
		<i>closeTransaction</i>	unlock <i>B, C</i>

Figure 12.15

Lock compatibility

<i>For one object</i>		<i>Lock requested</i>	
		<i>read</i>	<i>write</i>
<i>Lock already set</i>	<i>none</i>	OK	OK
	<i>read</i>	OK	wait
	<i>write</i>	wait	wait

Figure 12.16 Use of locks in strict two-phase locking

1. When an operation accesses an object within a transaction:
 - (a) If the object is not already locked, it is locked and the operation proceeds.
 - (b) If the object has a conflicting lock set by another transaction, the transaction must wait until it is unlocked.
 - (c) If the object has a non-conflicting lock set by another transaction, the lock is shared and the operation proceeds.
 - (d) If the object has already been locked in the same transaction, the lock will be promoted if necessary and the operation proceeds. (Where promotion is prevented by a conflicting lock, rule (b) is used.)
2. When a transaction is committed or aborted, the server unlocks all objects it locked for the transaction.

Figure 12.17

Lock class

```
public class Lock {  
    private Object object; // the object being protected by the lock  
    private Vector holders; // the TIDs of current holders  
    private LockType lockType; // the current type  
  
    public synchronized void acquire(TransID trans, LockType aLockType ){  
        while(/*another transaction holds the lock in conflicting mode*/) {  
            try {  
                wait();  
            }  
            catch ( InterruptedException e){/*...*/ }  
        }  
        if(holders.isEmpty()) { // no TIDs hold lock  
            holders.addElement(trans);  
            lockType = aLockType;  
        } else if(/*another transaction holds the lock, share it*/ ) {  
            if(/* this transaction not a holder*/) holders.addElement(trans);  
        } else if (/* this transaction is a holder but needs a more exclusive lock*/) {  
            lockType.promote();  
        }  
    }  
}
```

// this figure continues on the next slide

Figure 12.17 continued

```
public synchronized void release(TransID trans ){  
    holders.removeElement(trans); // remove this holder  
    // set locktype to none  
    notifyAll();  
    }  
}
```

Figure 12.18 *LockManager* class

```
public class LockManager {  
    private Hashtable theLocks;  
  
    public void setLock(Object object, TransID trans, LockType lockType){  
        Lock foundLock;  
        synchronized(this){  
            // find the lock associated with object  
            // if there isn't one, create it and add to the hashtable  
        }  
        foundLock.acquire(trans, lockType);  
    }  
  
    // synchronize this one because we want to remove all entries  
    public synchronized void unlock(TransID trans) {  
        Enumeration e = theLocks.elements();  
        while(e.hasMoreElements()){  
            Lock aLock = (Lock)(e.nextElement());  
            if( /* trans is a holder of this lock */ ) aLock.release(trans);  
        }  
    }  
}
```

Figure 12.19 Deadlock with write locks

Transaction <i>T</i>		Transaction <i>U</i>	
Operations	Locks	Operations	Locks
<i>a.deposit(100);</i>	write lock <i>A</i>	<i>b.deposit(200)</i>	write lock <i>B</i>
<i>b.withdraw(100)</i>		<i>a.withdraw(200);</i>	
...	waits for <i>U</i> 's lock on <i>B</i>	...	waits for <i>T</i> 's lock on <i>A</i>
...		...	
...		...	

Figure 12.20 The wait-for graph for Figure 12.19




Figure 12.21 A cycle in a wait-for graph




Figure 12.22 Another wait-for graph





Figure 12.23 Resolution of the deadlock in Figure 12.19

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

Figure 12.24 Lock compatibility (*read, write and commit locks*)

<i>For one object</i>		<i>Lock to be set</i>		
		<i>read</i>	<i>write</i>	<i>commit</i>
<i>Lock already set</i>	<i>none</i>	OK	OK	OK
	<i>read</i>	OK	OK	wait
	<i>write</i>	OK	wait	–
	<i>commit</i>	wait	wait	–

Figure 12.25 Lock hierarchy for the banking example




Figure 12.26 Lock hierarchy for a diary




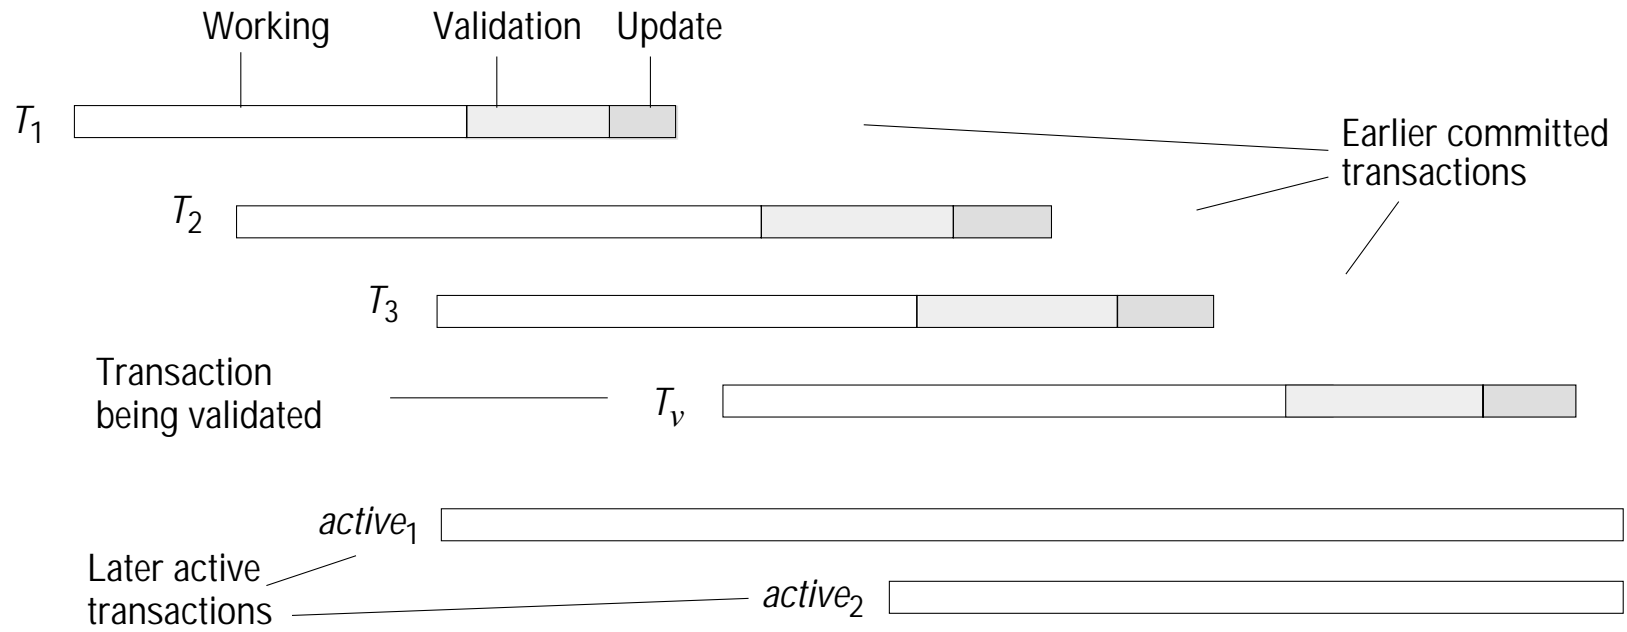
Figure 12.27 Lock compatibility table for hierarchic locks

<i>For one object</i>		<i>Lock to be set</i>			
		<i>read</i>	<i>write</i>	<i>I-read</i>	<i>I-write</i>
<i>Lock already set</i>	<i>none</i>	OK	OK	OK	OK
	<i>read</i>	OK	wait	OK	wait
	<i>write</i>	wait	wait	wait	wait
	<i>I-read</i>	OK	wait	OK	OK
	<i>I-write</i>	wait	wait	OK	OK

Serializability of transaction T with respect to transaction T_i

T_v	T_i	Rule
<i>write</i>	<i>read</i>	1. T_i must not read objects written by T_v .
<i>read</i>	<i>write</i>	2. T_v must not read objects written by T_i .
<i>write</i>	<i>write</i>	3. T_i must not write objects written by T_v and T_v must not write objects written by T_i .

Figure 12.28 Validation of transactions



Backward validation of transaction T_v

```
boolean valid = true;
for (int  $T_i = startT_{n+1}; T_i \leq finishT_n; T_i++$ ) {
    if (read set of  $T_v$  intersects write set of  $T_i$ ) valid = false;
}
```

Forward validation of transaction T_v


```
boolean valid = true;
for (int  $T_{id} = active_1; T_{id} \leq active_N; T_{id}++$ ) {
    if (write set of  $T_v$  intersects read set of  $T_{id}$ ) valid = false;
}
```

Figure 12.29 Operation conflicts for timestamp ordering


<i>Rule</i>	T_c	T_i	
1.	<i>write</i>	<i>read</i>	T_c must not <i>write</i> an object that has been <i>read</i> by any T_i where $T_i > T_c$ this requires that $T_c \geq$ the maximum read timestamp of the object.
2.	<i>write</i>	<i>write</i>	T_c must not <i>write</i> an object that has been <i>written</i> by any T_i where $T_i > T_c$ this requires that $T_c >$ write timestamp of the committed object.
3.	<i>read</i>	<i>write</i>	T_c must not <i>read</i> an object that has been <i>written</i> by any T_i where $T_i > T_c$ this requires that $T_c >$ write timestamp of the committed object.

Figure 12.30 Write operations and timestamps


(a) T_3 write



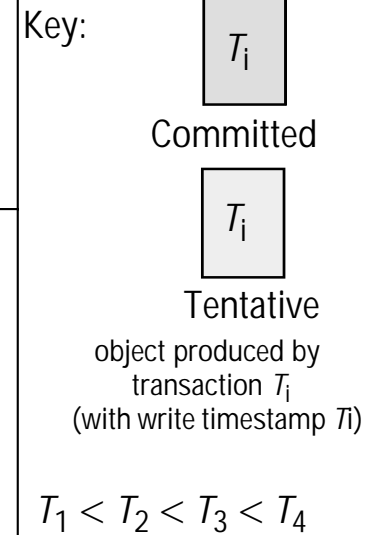

(b) T_3 write



(c) T_3 write



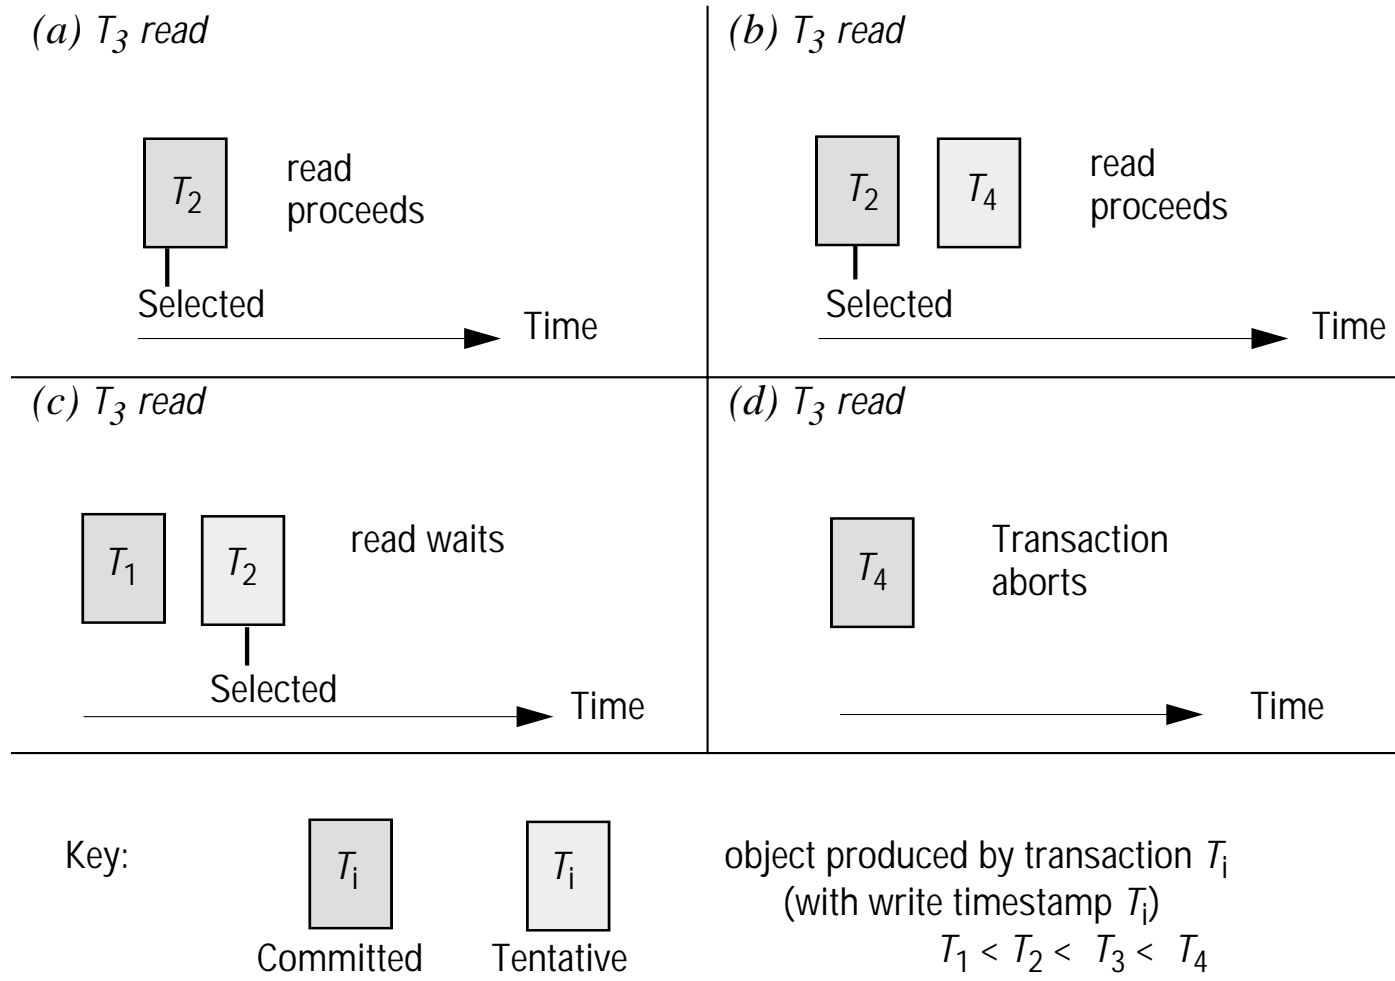
(d) T_3 write



Timestamp ordering write rule

```
if ( $T_c \geq$  maximum read timestamp on  $D$  &&  
     $T_c >$  write timestamp on committed version of  $D$ )  
    perform write operation on tentative version of  $D$  with write timestamp  $T_c$   
else /* write is too late */  
Abort transaction  $T_c$ 
```


Figure 12.31 Read operations and timestamps



Timestamp ordering read rule

```
if (  $T_c >$  write timestamp on committed version of  $D$  ) {  
    let  $D_{\text{selected}}$  be the version of  $D$  with the maximum write timestamp  $\leq T_c$   
    if ( $D_{\text{selected}}$  is committed)  
        perform read operation on the version  $D_{\text{selected}}$   
    else  
        Wait until the transaction that made version  $D_{\text{selected}}$  commits or aborts  
        then reapply the read rule  
} else  
    Abort transaction  $T_c$ 
```

Figure 12.32 Timestamps in transactions *T* and *U*

		<i>Timestamps and versions of objects</i>					
<i>T</i>	<i>U</i>	<i>A</i>		<i>B</i>		<i>C</i>	
		<i>RTS</i>	<i>WTS</i>	<i>RTS</i>	<i>WTS</i>	<i>RTS</i>	<i>WTS</i>
		{}	S	{}	S	{}	S
<i>openTransaction</i>							
<i>bal = b.getBalance()</i>				{ <i>T</i> }			
	<i>openTransaction</i>						
<i>b.setBalance(bal*1.1)</i>					S, T		
	<i>bal = b.getBalance()</i>						
	<i>wait for T</i>						
<i>a.withdraw(bal/10)</i>	...		S, T				
<i>commit</i>	...		T		T		
	<i>bal = b.getBalance()</i>						
	<i>b.setBalance(bal*1.1)</i>			{ <i>U</i> }			
	<i>b.setBalance(bal*1.1)</i>				T, U		
	<i>c.withdraw(bal/10)</i>						S, U

Figure 12.33 Late *write* operation would invalidate a *read*

