

# Chapter 22

---

## Transaction Management

# Transaction Support

---

## Transaction

**Action, or series of actions, carried out by user or application, which reads or updates contents of database.**

- ◆ **Logical unit of work on the database.**
- ◆ **Application program is series of transactions with non-database processing in between.**
- ◆ **Transforms database from one consistent state to another, although consistency may be violated during transaction.**

# Example Transaction

---

```
read(staffNo = x, salary)
salary = salary * 1.1
write(staffNo = x, new_salary)
```

(a)

```
delete(staffNo = x)
for all PropertyForRent records, pno
begin
  read(propertyNo = pno, staffNo)
  if (staffNo = x) then
    begin
      staffNo = newStaffNo
      write(propertyNo = pno, staffNo)
    end
  end
end
```

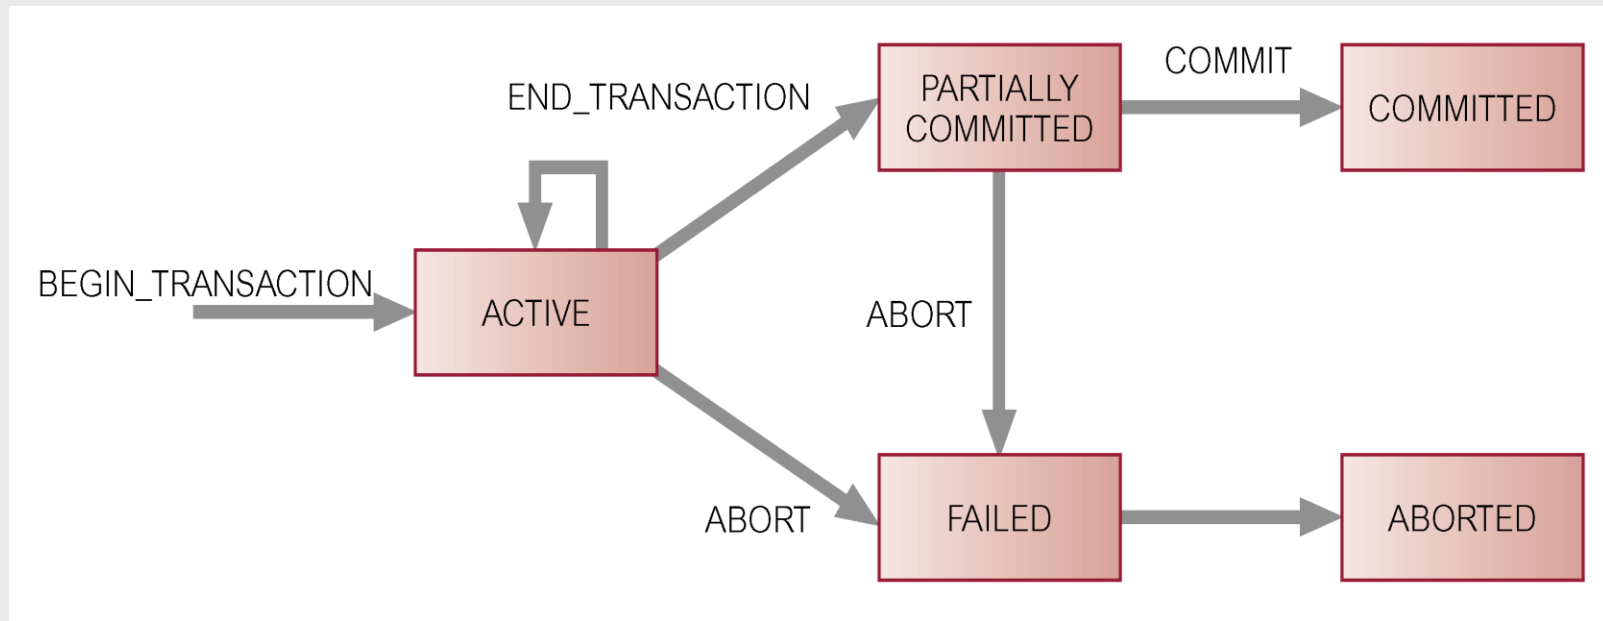
(b)

# Transaction Support

---

- ◆ **Can have one of two outcomes:**
  - **Success** - transaction *commits* and database reaches a new consistent state.
  - **Failure** - transaction *aborts*, and database must be restored to consistent state before it started.
  - Such a transaction is *rolled back* or *undone*.
- ◆ **Committed transaction cannot be aborted.**
- ◆ **Aborted transaction that is rolled back can be restarted later.**

# State Transition Diagram for Transaction



# Properties of Transactions

---

◆ Four basic (*ACID*) properties of a transaction are:

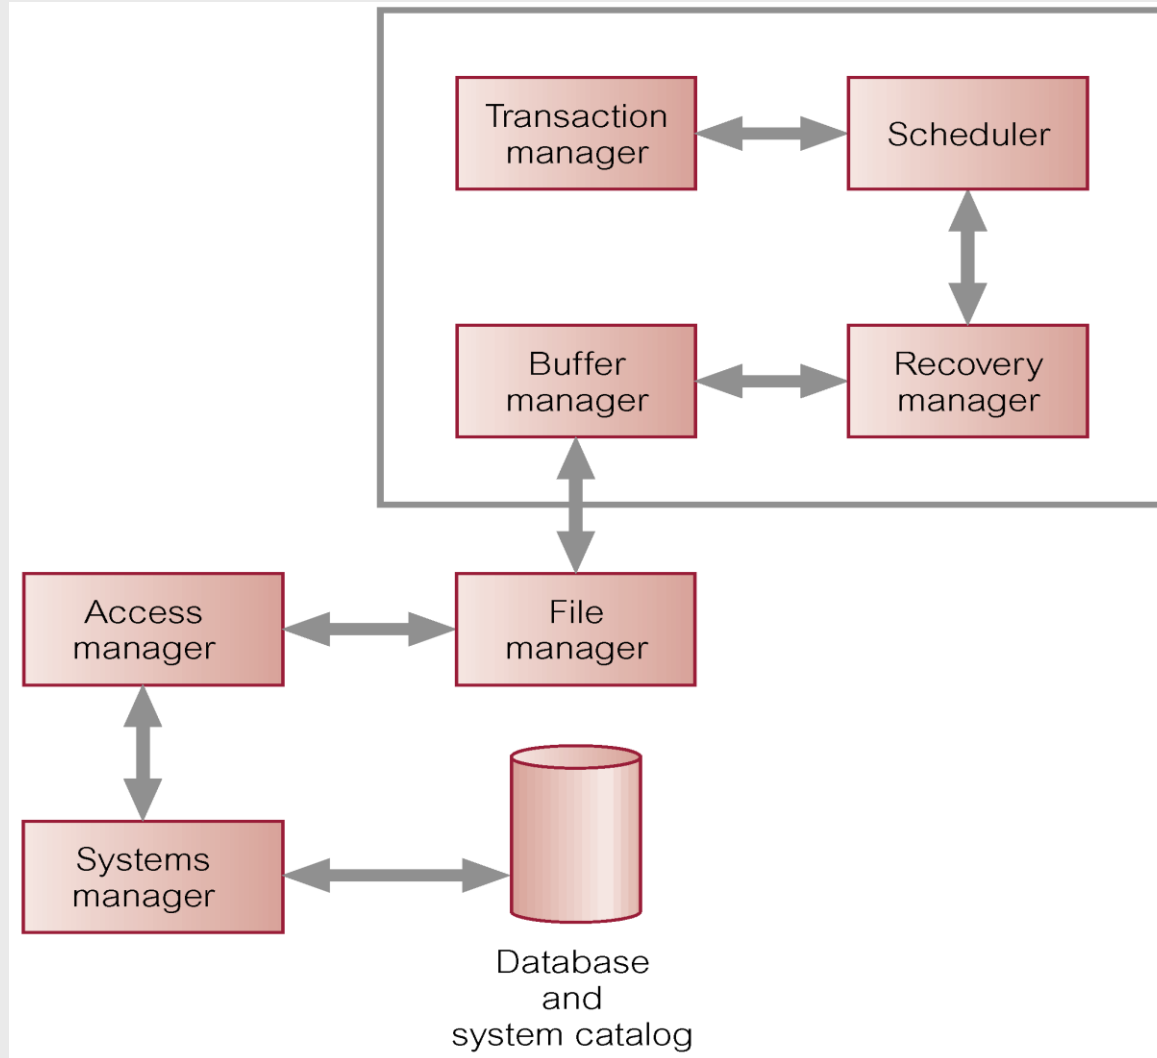
Atomicity      ‘All or nothing’ property.

Consistency    Must transform database from one consistent state to another.

Isolation      Partial effects of incomplete transactions should not be visible to other transactions.

Durability     Effects of a committed transaction are permanent and must not be lost because of later failure.

# DBMS Transaction Subsystem



# Concurrency Control

---

**Process of managing simultaneous operations on the database without having them interfere with one another.**

- ◆ **Prevents interference when two or more users are accessing database simultaneously and at least one is updating data.**
- ◆ **Although two transactions may be correct in themselves, interleaving of operations may produce an incorrect result.**



# Need for Concurrency Control

---

- ◆ **Three examples of potential problems caused by concurrency:**
  - **Lost update problem.**
  - **Uncommitted dependency problem.**
  - **Inconsistent analysis problem.**

# Lost Update Problem

---

- ◆ **Successfully completed update is overridden by another user.**
- ◆  **$T_1$  withdrawing £10 from an account with  $bal_x$ , initially £100.**
- ◆  **$T_2$  depositing £100 into same account.**
- ◆ **Serially, final balance would be £190.**

# Lost Update Problem

Time	T <sub>1</sub>	T <sub>2</sub>	bal <sub>x</sub>
t <sub>1</sub>		begin_transaction	100
t <sub>2</sub>	begin_transaction	read(bal <sub>x</sub> )	100
t <sub>3</sub>	read(bal <sub>x</sub> )	bal <sub>x</sub> = bal <sub>x</sub> + 100	100
t <sub>4</sub>	bal <sub>x</sub> = bal <sub>x</sub> - 10	write(bal <sub>x</sub> )	200
t <sub>5</sub>	write(bal <sub>x</sub> )	commit	90
t <sub>6</sub>	commit		90

- ◆ **Loss of T<sub>2</sub>'s update avoided by preventing T<sub>1</sub> from reading bal<sub>x</sub> until after update.**

# Uncommitted Dependency Problem

---

- ◆ Occurs when one transaction can see intermediate results of another transaction before it has committed.
- ◆  $T_4$  updates  $bal_x$  to £200 but it aborts, so  $bal_x$  should be back at original value of £100.
- ◆  $T_3$  has read new value of  $bal_x$  (£200) and uses value as basis of £10 reduction, giving a new balance of £190, instead of £90.

# Uncommitted Dependency Problem

Time	T <sub>3</sub>	T <sub>4</sub>	bal <sub>x</sub>
t <sub>1</sub>		begin_transaction	100
t <sub>2</sub>		read(bal <sub>x</sub> )	100
t <sub>3</sub>		bal <sub>x</sub> = bal <sub>x</sub> + 100	100
t <sub>4</sub>	begin_transaction	write(bal <sub>x</sub> )	200
t <sub>5</sub>	read(bal <sub>x</sub> )	:	200
t <sub>6</sub>	bal <sub>x</sub> = bal <sub>x</sub> - 10	rollback	100
t <sub>7</sub>	write(bal <sub>x</sub> )		190
t <sub>8</sub>	commit		190

- ◆ Problem avoided by preventing T<sub>3</sub> from reading bal<sub>x</sub> until after T<sub>4</sub> commits or aborts.

## Inconsistent Analysis Problem

---

- ◆ Occurs when transaction reads several values but second transaction updates some of them during execution of first.
- ◆ Sometimes referred to as *dirty read* or *unrepeatable read*.
- ◆  $T_6$  is totaling balances of accounts:  
 $x$  (£100),  $y$  (£50), and  $z$  (£25).
- ◆ Meantime,  $T_5$  has transferred £10 from  $bal_x$  to  $bal_z$ , so  $T_6$  now has wrong result (£10 too high).

# Inconsistent Analysis Problem

Time	T <sub>5</sub>	T <sub>6</sub>	bal <sub>x</sub>	bal <sub>y</sub>	bal <sub>z</sub>	sum
t <sub>1</sub>		begin_transaction	100	50	25	
t <sub>2</sub>	begin_transaction	sum = 0	100	50	25	0
t <sub>3</sub>	read(bal <sub>x</sub> )	read(bal <sub>x</sub> )	100	50	25	0
t <sub>4</sub>	bal <sub>x</sub> = bal <sub>x</sub> - 10	sum = sum + bal <sub>x</sub>	100	50	25	100
t <sub>5</sub>	write(bal <sub>x</sub> )	read(bal <sub>y</sub> )	90	50	25	100
t <sub>6</sub>	read(bal <sub>z</sub> )	sum = sum + bal <sub>y</sub>	90	50	25	150
t <sub>7</sub>	bal <sub>z</sub> = bal <sub>z</sub> + 10		90	50	25	150
t <sub>8</sub>	write(bal <sub>z</sub> )		90	50	35	150
t <sub>9</sub>	commit	read(bal <sub>z</sub> )	90	50	35	150
t <sub>10</sub>		sum = sum + bal <sub>z</sub>	90	50	35	185
t <sub>11</sub>		commit	90	50	35	185

- ◆ Problem avoided by preventing T<sub>6</sub> from reading bal<sub>x</sub> and bal<sub>z</sub> until after T<sub>5</sub> completed updates.

# Serializability

---

- ◆ **Objective of a concurrency control protocol is to schedule transactions in such a way as to avoid any interference.**
- ◆ **Could run transactions serially, but this limits degree of concurrency or parallelism in system.**
- ◆ **Serializability identifies those executions of transactions guaranteed to ensure consistency.**



# Serializability

---

## Schedule

**Sequence of reads/writes by set of concurrent transactions.**

## Serial Schedule

**Schedule where operations of each transaction are executed consecutively without any interleaved operations from other transactions.**

- ◆ **No guarantee that results of all serial executions of a given set of transactions will be identical.**

# Nonserial Schedule

---

- ◆ **Schedule where operations from set of concurrent transactions are interleaved.**
- ◆ **Objective of serializability is to find nonserial schedules that allow transactions to execute concurrently without interfering with one another.**
- ◆ **In other words, want to find nonserial schedules that are equivalent to *some* serial schedule. Such a schedule is called *serializable*.**

# Serializability

---

- ◆ **In serializability, ordering of read/writes is important:**
  - (a) If two transactions only read a data item, they do not conflict and order is not important.**
  - (b) If two transactions either read or write completely separate data items, they do not conflict and order is not important.**
  - (c) If one transaction writes a data item and another reads or writes same data item, order of execution is important.**

# Example of Conflict Serializability

---

Time	T <sub>7</sub>	T <sub>8</sub>
t <sub>1</sub>	begin_transaction	
t <sub>2</sub>	read( <b>bal<sub>x</sub></b> )	
t <sub>3</sub>	write( <b>bal<sub>x</sub></b> )	
t <sub>4</sub>		begin_transaction
t <sub>5</sub>		read( <b>bal<sub>x</sub></b> )
t <sub>6</sub>		write( <b>bal<sub>x</sub></b> )
t <sub>7</sub>	read( <b>bal<sub>y</sub></b> )	
t <sub>8</sub>	write( <b>bal<sub>y</sub></b> )	
t <sub>9</sub>	commit	
t <sub>10</sub>		read( <b>bal<sub>y</sub></b> )
t <sub>11</sub>		write( <b>bal<sub>y</sub></b> )
t <sub>12</sub>		commit

# Example of Conflict Serializability

---

T <sub>7</sub>	T <sub>8</sub>	T <sub>7</sub>	T <sub>8</sub>
begin_transaction		begin_transaction	
read( <b>bal<sub>x</sub></b> )		read( <b>bal<sub>x</sub></b> )	
write( <b>bal<sub>x</sub></b> )		write( <b>bal<sub>x</sub></b> )	
	begin_transaction	read( <b>bal<sub>y</sub></b> )	
	read( <b>bal<sub>x</sub></b> )	write( <b>bal<sub>y</sub></b> )	
read( <b>bal<sub>y</sub></b> )		commit	
	write( <b>bal<sub>x</sub></b> )		begin_transaction
write( <b>bal<sub>y</sub></b> )			read( <b>bal<sub>x</sub></b> )
commit			write( <b>bal<sub>x</sub></b> )
	read( <b>bal<sub>y</sub></b> )		read( <b>bal<sub>y</sub></b> )
	write( <b>bal<sub>y</sub></b> )		write( <b>bal<sub>y</sub></b> )
	commit		commit

# Serializability

---

- ◆ **Conflict serializable schedule orders any conflicting operations in same way as some serial execution.**
- ◆ **Under *constrained write rule* where:**
  - **A transaction updates data item based on its old value, which is first read.**
- ◆ **Use *precedence graph* to test for serializability.**

# Precedence Graph

---

- ◆ **Create:**
  - **node for each transaction;**
  - **a directed edge  $T_i \rightarrow T_j$ , if  $T_j$  reads the value of an item written by  $T_i$ ;**
  - **a directed edge  $T_i \rightarrow T_j$ , if  $T_j$  writes a value into an item after it has been read by  $T_i$ .**
  - **a directed edge  $T_i \rightarrow T_j$ , if  $T_j$  writes a value into an item after it has been written by  $T_i$ .**
- ◆ **If precedence graph contains a cycle, schedule is not conflict serializable.**

## Example - Non-conflict serializable schedule

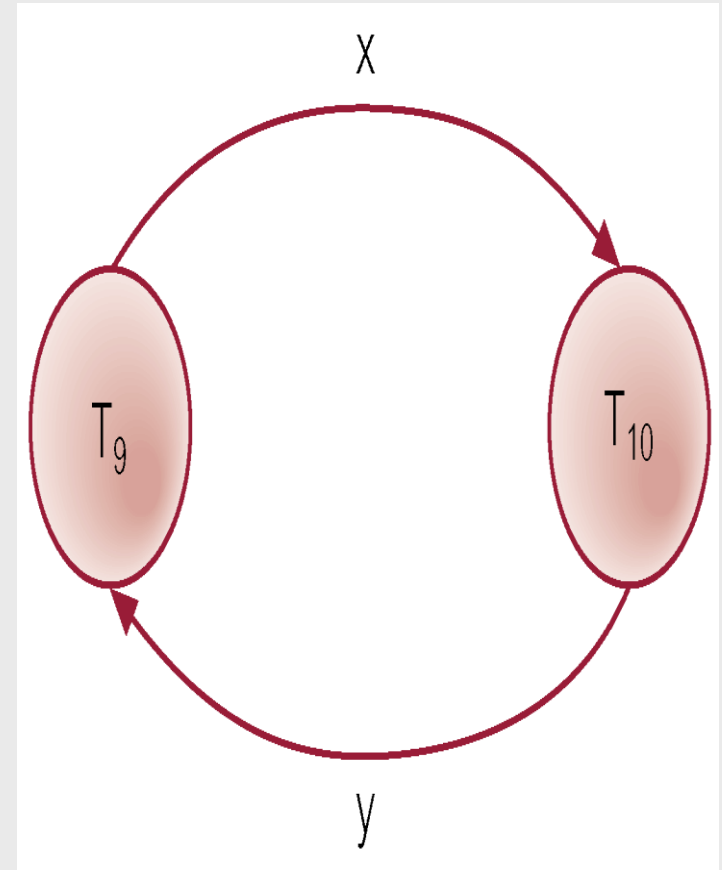
---

- ◆  $T_9$  is transferring £100 from one account with balance  $bal_x$  to another account with balance  $bal_y$ .
- ◆  $T_{10}$  is increasing balance of these two accounts by 10%.
- ◆ Precedence graph has a cycle and so is not serializable.



# Example - Non-conflict serializable schedule

Time	T <sub>9</sub>	T <sub>10</sub>
t <sub>1</sub>	begin_transaction	
t <sub>2</sub>	read( <b>bal<sub>x</sub></b> )	
t <sub>3</sub>	<b>bal<sub>x</sub> = bal<sub>x</sub> + 100</b>	
t <sub>4</sub>	write( <b>bal<sub>x</sub></b> )	begin_transaction
t <sub>5</sub>		read( <b>bal<sub>x</sub></b> )
t <sub>6</sub>		<b>bal<sub>x</sub> = bal<sub>x</sub> * 1.1</b>
t <sub>7</sub>		write( <b>bal<sub>x</sub></b> )
t <sub>8</sub>		read( <b>bal<sub>y</sub></b> )
t <sub>9</sub>		<b>bal<sub>y</sub> = bal<sub>y</sub> * 1.1</b>
t <sub>10</sub>		write( <b>bal<sub>y</sub></b> )
t <sub>11</sub>	read( <b>bal<sub>y</sub></b> )	commit
t <sub>12</sub>	<b>bal<sub>y</sub> = bal<sub>y</sub> - 100</b>	
t <sub>13</sub>	write( <b>bal<sub>y</sub></b> )	
t <sub>14</sub>	commit	



# Recoverability

---

- ◆ **Serializability identifies schedules that maintain database consistency, assuming no transaction fails.**
- ◆ **Could also examine recoverability of transactions within schedule.**
- ◆ **If transaction fails, atomicity requires effects of transaction to be undone.**
- ◆ **Durability states that once transaction commits, its changes cannot be undone (without running another, compensating, transaction).**

## Recoverable Schedule

---

**A schedule where, for each pair of transactions  $T_i$  and  $T_j$ , if  $T_j$  reads a data item previously written by  $T_i$ , then the commit operation of  $T_i$  precedes the commit operation of  $T_j$ .**

# Concurrency Control Techniques

---

- ◆ **Two basic concurrency control techniques:**
  - **Locking,**
  - **Timestamping.**
- ◆ **Both are conservative approaches: delay transactions in case they conflict with other transactions.**
- ◆ **Optimistic methods assume conflict is rare and only check for conflicts at commit.**

# Locking

---

Transaction uses locks to deny access to other transactions and so prevent incorrect updates.

- ◆ Most widely used approach to ensure serializability.
- ◆ Generally, a transaction must claim a *shared (read)* or *exclusive (write)* lock on a data item before read or write.
- ◆ Lock prevents another transaction from modifying item or even reading it, in the case of a write lock.

## Locking - Basic Rules

---

- ◆ **If transaction has shared lock on item, can read but not update item.**
- ◆ **If transaction has exclusive lock on item, can both read and update item.**
- ◆ **Reads cannot conflict, so more than one transaction can hold shared locks simultaneously on same item.**
- ◆ **Exclusive lock gives transaction exclusive access to that item.**

## Locking - Basic Rules

---

- ◆ **Some systems allow transaction to upgrade read lock to an exclusive lock, or downgrade exclusive lock to a shared lock.**

## Example - Incorrect Locking Schedule

Time	T <sub>9</sub>	T <sub>10</sub>
t <sub>1</sub>	begin_transaction	
t <sub>2</sub>	read(bal <sub>x</sub> )	
t <sub>3</sub>	bal <sub>x</sub> = bal <sub>x</sub> + 100	
t <sub>4</sub>	write(bal <sub>x</sub> )	begin_transaction
t <sub>5</sub>		read(bal <sub>x</sub> )
t <sub>6</sub>		bal <sub>x</sub> = bal <sub>x</sub> * 1.1
t <sub>7</sub>		write(bal <sub>x</sub> )
t <sub>8</sub>		read(bal <sub>y</sub> )
t <sub>9</sub>		bal <sub>y</sub> = bal <sub>y</sub> * 1.1
t <sub>10</sub>		write(bal <sub>y</sub> )
t <sub>11</sub>	read(bal <sub>y</sub> )	commit
t <sub>12</sub>	bal <sub>y</sub> = bal <sub>y</sub> - 100	
t <sub>13</sub>	write(bal <sub>y</sub> )	
t <sub>14</sub>	commit	

◆ For these two transactions, a valid schedule using these rules is:  $S = \{$

**T09-** write\_lock(bal<sub>x</sub>), read(bal<sub>x</sub>), write(bal<sub>x</sub>), unlock(bal<sub>x</sub>).

**T10-** write\_lock(bal<sub>x</sub>), read(bal<sub>x</sub>), write(bal<sub>x</sub>), unlock(bal<sub>x</sub>).

**T10-** write\_lock(bal<sub>y</sub>), read(bal<sub>y</sub>), write( bal<sub>y</sub>), unlock(bal<sub>y</sub>), commit(T<sub>10</sub>)

**T09-** write\_lock(bal<sub>y</sub>), read(bal<sub>y</sub>), write(bal<sub>y</sub>), unlock(bal<sub>y</sub>), commit(T<sub>9</sub>)}



## Example - Incorrect Locking Schedule

---

- ◆ If at start,  $\text{bal}_x = 100$ ,  $\text{bal}_y = 400$ , result should be:
  - $\text{bal}_x = 220$ ,  $\text{bal}_y = 330$ , if  $T_9$  executes before  $T_{10}$ ,  
or
  - $\text{bal}_x = 210$ ,  $\text{bal}_y = 340$ , if  $T_{10}$  executes before  $T_9$ .
- ◆ However, result gives  $\text{bal}_x = 220$  and  $\text{bal}_y = 340$ .
- ◆ S is not a serializable schedule.

## Example - Incorrect Locking Schedule

---

- ◆ **Problem is that transactions release locks too soon, resulting in loss of total isolation and atomicity.**
- ◆ **To guarantee serializability, need an additional protocol concerning the positioning of lock and unlock operations in every transaction.**

## Two-Phase Locking (2PL)

---

**Transaction follows 2PL protocol if all locking operations precede first unlock operation in the transaction.**

- ◆ **Two phases for transaction:**
  - **Growing phase - acquires all locks but cannot release any locks.**
  - **Shrinking phase - releases locks but cannot acquire any new locks.**

# Preventing Lost Update Problem using 2PL

Time	T <sub>1</sub>	T <sub>2</sub>	bal <sub>x</sub>
t <sub>1</sub>		begin_transaction	100
t <sub>2</sub>	begin_transaction	write_lock( <b>bal<sub>x</sub></b> )	100
t <sub>3</sub>	write_lock( <b>bal<sub>x</sub></b> )	read( <b>bal<sub>x</sub></b> )	100
t <sub>4</sub>	WAIT	<b>bal<sub>x</sub></b> = <b>bal<sub>x</sub></b> + 100	100
t <sub>5</sub>	WAIT	write( <b>bal<sub>x</sub></b> )	200
t <sub>6</sub>	WAIT	commit/unlock( <b>bal<sub>x</sub></b> )	200
t <sub>7</sub>	read( <b>bal<sub>x</sub></b> )		200
t <sub>8</sub>	<b>bal<sub>x</sub></b> = <b>bal<sub>x</sub></b> - 10		200
t <sub>9</sub>	write( <b>bal<sub>x</sub></b> )		190
t <sub>10</sub>	commit/unlock( <b>bal<sub>x</sub></b> )		190

# Preventing Uncommitted Dependency Problem using 2PL

Time	T <sub>3</sub>	T <sub>4</sub>	bal <sub>x</sub>
t <sub>1</sub>		begin_transaction	100
t <sub>2</sub>		write_lock( <b>bal<sub>x</sub></b> )	100
t <sub>3</sub>		read( <b>bal<sub>x</sub></b> )	100
t <sub>4</sub>	begin_transaction	<b>bal<sub>x</sub> = bal<sub>x</sub> + 100</b>	100
t <sub>5</sub>	write_lock( <b>bal<sub>x</sub></b> )	write( <b>bal<sub>x</sub></b> )	200
t <sub>6</sub>	WAIT	rollback/unlock( <b>bal<sub>x</sub></b> )	100
t <sub>7</sub>	read( <b>bal<sub>x</sub></b> )		100
t <sub>8</sub>	<b>bal<sub>x</sub> = bal<sub>x</sub> - 10</b>		100
t <sub>9</sub>	write( <b>bal<sub>x</sub></b> )		90
t <sub>10</sub>	commit/unlock( <b>bal<sub>x</sub></b> )		90

# Preventing Inconsistent Analysis Problem using 2PL

Time	T <sub>5</sub>	T <sub>6</sub>	bal <sub>x</sub>	bal <sub>y</sub>	bal <sub>z</sub>	sum
t <sub>1</sub>		begin_transaction	100	50	25	
t <sub>2</sub>	begin_transaction	sum = 0	100	50	25	0
t <sub>3</sub>	write_lock( <b>bal<sub>x</sub></b> )		100	50	25	0
t <sub>4</sub>	read( <b>bal<sub>x</sub></b> )	read_lock( <b>bal<sub>x</sub></b> )	100	50	25	0
t <sub>5</sub>	<b>bal<sub>x</sub> = bal<sub>x</sub> - 10</b>	WAIT	100	50	25	0
t <sub>6</sub>	write( <b>bal<sub>x</sub></b> )	WAIT	90	50	25	0
t <sub>7</sub>	write_lock( <b>bal<sub>z</sub></b> )	WAIT	90	50	25	0
t <sub>8</sub>	read( <b>bal<sub>z</sub></b> )	WAIT	90	50	25	0
t <sub>9</sub>	<b>bal<sub>z</sub> = bal<sub>z</sub> + 10</b>	WAIT	90	50	25	0
t <sub>10</sub>	write( <b>bal<sub>z</sub></b> )	WAIT	90	50	35	0
t <sub>11</sub>	commit/unlock( <b>bal<sub>x</sub>, bal<sub>z</sub></b> )	WAIT	90	50	35	0
t <sub>12</sub>		read( <b>bal<sub>x</sub></b> )	90	50	35	0
t <sub>13</sub>		sum = sum + <b>bal<sub>x</sub></b>	90	50	35	90
t <sub>14</sub>		read_lock( <b>bal<sub>y</sub></b> )	90	50	35	90
t <sub>15</sub>		read( <b>bal<sub>y</sub></b> )	90	50	35	90
t <sub>16</sub>		sum = sum + <b>bal<sub>y</sub></b>	90	50	35	140
t <sub>17</sub>		read_lock( <b>bal<sub>z</sub></b> )	90	50	35	140
t <sub>18</sub>		read( <b>bal<sub>z</sub></b> )	90	50	35	140
t <sub>19</sub>		sum = sum + <b>bal<sub>z</sub></b>	90	50	35	175
t <sub>20</sub>		commit/unlock( <b>bal<sub>x</sub>, bal<sub>y</sub>, bal<sub>z</sub></b> )	90	50	35	175

# Cascading Rollback

---

- ◆ If *every* transaction in a schedule follows 2PL, schedule is serializable.
- ◆ However, problems can occur with interpretation of when locks can be released.

# Cascading Rollback

Time	T <sub>14</sub>	T <sub>15</sub>	T <sub>16</sub>
t <sub>1</sub>	begin_transaction		
t <sub>2</sub>	write_lock( <b>bal<sub>x</sub></b> )		
t <sub>3</sub>	read( <b>bal<sub>x</sub></b> )		
t <sub>4</sub>	read_lock( <b>bal<sub>y</sub></b> )		
t <sub>5</sub>	read( <b>bal<sub>y</sub></b> )		
t <sub>6</sub>	<b>bal<sub>x</sub> = bal<sub>y</sub> + bal<sub>x</sub></b>		
t <sub>7</sub>	write( <b>bal<sub>x</sub></b> )		
t <sub>8</sub>	unlock( <b>bal<sub>x</sub></b> )	begin_transaction	
t <sub>9</sub>	:	write_lock( <b>bal<sub>x</sub></b> )	
t <sub>10</sub>	:	read( <b>bal<sub>x</sub></b> )	
t <sub>11</sub>	:	<b>bal<sub>x</sub> = bal<sub>x</sub> + 100</b>	
t <sub>12</sub>	:	write( <b>bal<sub>x</sub></b> )	
t <sub>13</sub>	:	unlock( <b>bal<sub>x</sub></b> )	
t <sub>14</sub>	:	:	
t <sub>15</sub>	rollback	:	
t <sub>16</sub>		:	begin_transaction
t <sub>17</sub>		:	read_lock( <b>bal<sub>x</sub></b> )
t <sub>18</sub>		rollback	:
t <sub>19</sub>			rollback



# Cascading Rollback

---

- ◆ Transactions conform to 2PL.
- ◆  $T_{14}$  aborts.
- ◆ Since  $T_{15}$  is dependent on  $T_{14}$ ,  $T_{15}$  must also be rolled back. Since  $T_{16}$  is dependent on  $T_{15}$ , it too must be rolled back.
- ◆ This is called *cascading rollback*.
- ◆ To prevent this with 2PL, leave release of *all* locks until end of transaction.

# Deadlock

An impasse that may result when two (or more) transactions are each waiting for locks held by the other to be released.

Time	T <sub>17</sub>	T <sub>18</sub>
t <sub>1</sub>	begin_transaction	
t <sub>2</sub>	write_lock( <b>bal<sub>x</sub></b> )	begin_transaction
t <sub>3</sub>	read( <b>bal<sub>x</sub></b> )	write_lock( <b>bal<sub>y</sub></b> )
t <sub>4</sub>	<b>bal<sub>x</sub> = bal<sub>x</sub> - 10</b>	read( <b>bal<sub>y</sub></b> )
t <sub>5</sub>	write( <b>bal<sub>x</sub></b> )	<b>bal<sub>y</sub> = bal<sub>y</sub> + 100</b>
t <sub>6</sub>	write_lock( <b>bal<sub>y</sub></b> )	write( <b>bal<sub>y</sub></b> )
t <sub>7</sub>	WAIT	write_lock( <b>bal<sub>x</sub></b> )
t <sub>8</sub>	WAIT	WAIT
t <sub>9</sub>	WAIT	WAIT
t <sub>10</sub>	⋮	WAIT
t <sub>11</sub>	⋮	⋮

# Deadlock

---

- ◆ **Only one way to break deadlock: abort one or more of the transactions.**
- ◆ **Deadlock should be transparent to user, so DBMS should restart transaction(s).**
- ◆ **Three general techniques for handling deadlock:**
  - **Timeouts.**
  - **Deadlock prevention.**
  - **Deadlock detection and recovery.**