


Lock

Binary
(boolean value for the lock)


0 - unlocked

1 - locked.

There is no distinction
between read lock
& write lock



Read lock (counter)
& write lock (boolean)



RL & WL are
mutually exclusive.

Binary Lock

Lock (item)

Unlock (item)

Lock(item):

```
start: if ( Lock(item) == 0 )
    { // item is unlocked
      Lock(item) = 1
    }
else // item is locked.
    {
      wait (until Lock(item) == 0) // process goes to sleep
      Lock manager wakes up the transaction
      when the lock is released.
    }
goto start
```

unlock(item):

Lock (item) = 0

Signal (activate) Lock Manager

so that Lock Manager will wake up

transactions that are waiting for this item.

Read Lock & Write Lock

Mode:

Read Locked (Shared) with no. of readers.

Write Locked (Exclusive): 0 or 1

Unlocked

ReadLock(item), WriteLock(item), Unlock(item)

ReadLock(item):

```
start: if (Lock(item) == 'Unlocked') {
        Lock(item) = 'Readlocked'
        no_of_readers = 1
    }
    else if (Lock(item) == 'Readlocked') {
        no_of_readers++
    }
    else // item is write locked
    {
        wait (until Lock(item) == 'Unlocked')
        Lock Manager wakes up the
        transaction
    }
    goto start
```

WriteLock(item):

```
start: if (Lock(item) == 'Unlocked')
    {
        Lock(item) = 'write locked'
    }
    else {
        wait (until Lock(item) == 'Unlocked')
        Lock Manager wakes up the transaction
    }
    goto start
```

Unlock(item):

Unlock(item): (Read lock & write lock)

```
if (Lock(item) == 'writelocked')
{
    Lock(item) = 'Unlocked'
    Signal (activate) Lock Manager
    so that it will wake up transactions
    that are waiting for this item
}
else // read locked
{
    No_of_readers--
    if (No_of_readers == 0)
    {
        Lock(item) = 'Unlocked'
        Signal (activate) Lock Manager
        so that it will wake up transactions
        that are waiting for this item
    }
}
```

Granularity of Locks

- Server - level

- Database Level

- Object (table) - Level

- row - level

- column - level

Intention Locks (meaningful in a hierarchy of objects)

IS (Intention shared)

placed at a higher level unit

and

Shared lock can be

placed at a lower level

e.g.

Company DB — IS

Employee Table — Shared lock

IE (Intention exclusive)

e.g.

T₄: Employee Table — ~~IX~~ IX

on a specific row — Exclusive lock.

T₆: [read on ~~some~~ ^{all rows} row of Employee Table
readlock(row)

→ Check the status on Employee Table: ~~IX~~

T₆ understands all may not be available for read.

Server

higher

Database

Object

Row

Column

Lower

Lost update Example

not following ^{T1} 2-PL

+ Lock(X)
read(X)
 $X = X - N$
~~Write(X)~~
~~Unlock(X)~~

+ Lock(Y)
Read(Y)
 $Y = Y + N$
Write(Y)
- Unlock(Y)

T2

+ Lock(X)
Read(X)
 $X = X + M$
Write(X)
- Unlock(X)

T1
Basic
Following 2-PL

↑ + Lock(X)
read(X)
 $X = X - N$
Write(X)
↓ + Lock(Y)
* Read(Y)
- Unlock(X)
Shrink
phase $Y = Y + N$
Write(Y)
↓ - Unlock(Y)

Two phase Locking Protocol

Each transaction has two distinct phases:

- ① Growth phase
Acquires locks for DB items
w/o releasing any locks
- ② Shrink phase
Releases locks of DB items
w/o acquiring new locks

Benefit: Any schedule made of transactions that follows two-phase locking protocol, is guaranteed to be ~~not~~ conflict serializable.

However a deadlock ~~may~~ can occur with two phase locking protocol transactions.

T1
Conservative 2PL

+ Lock(X)
+ Lock(Y)

read(X)
X = X - N
write(X)
- Unlock(X)

Read(Y)
Y = Y + N
write(Y)
- Unlock(Y)

This will not lead to Deadlock

T3
Strict 2PL

+ RLock(A)
Read(A)

+ WLock(B)
Read(B)
B = B + 100

- Unlock(A)
write(B)
Commit.

- Unlock(B)

↑ Growth Phase
↓ Shrink Phase

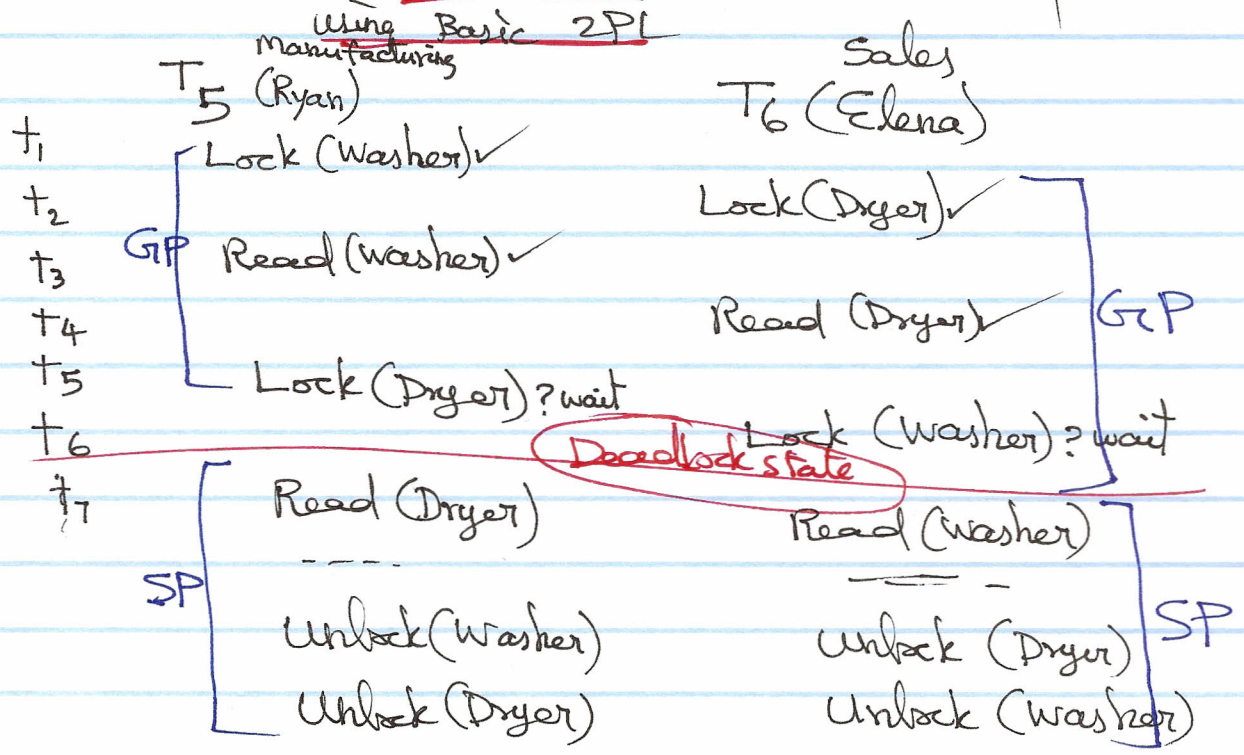
T3
Rigorous 2PL

+ RLock(A)
Read(A)

+ WLock(B)
Read(B)
B = B + 100
write(B)
Commit

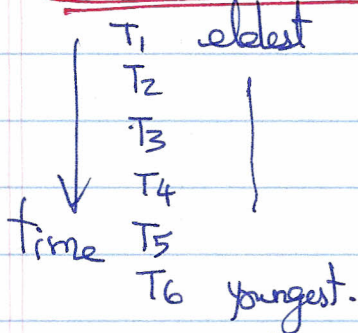
- Unlock(B)
- Unlock(A)

Deadlock



Deadlock Prevention

With Time stamp (Chronological order)



a) Wait-Die

	Elder T ₅ (Ryan)	Younger T ₆ (Elena)
t ₁	Lock(Washer) ✓	
t ₂		Lock(Dryer) ✓
t ₃	Read(Washer) ✓	Read(Dryer) ✓
t ₄	Lock(Dryer)? wait	Lock(Washer)? self abort
t ₅		
t ₆	Read(Dryer) ✓ unlock(Washer) ✓ unlock(Dryer) ✓	

↑ unlock(Dryer)
 ↑ rollback
 T₆ gets resubmitted
 restarted
 with previous
 timestamp T₆

b) Wound-wait

	T ₅ (Ryan)	T ₆ (Elena)
t ₁	Lock(Washer) ✓	
t ₂		Lock(Dryer) ✓
t ₃	Read(Washer) ✓	Read(Dryer) ✓
t ₄		
t ₅	Lock(Dryer)? abort → T ₆ gets aborted	
t ₆	Read(Dryer) ✓ unlock(Washer) ✓ unlock(Dryer) ✓	

↑ unlock(Dryer)
 ↑ rollback
 T₆ gets restarted
 with the same
 timestamp

Timestamp Ordering

(no locks are used)

For each ^{active} DB item (i.e. DB item in memory)

a read timestamp (RTS)

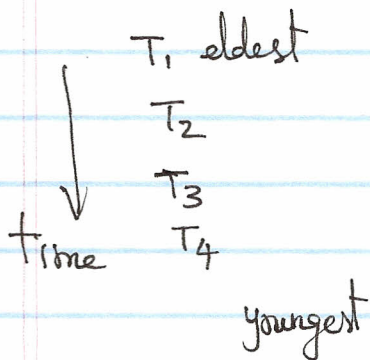
and a write timestamp (WTS) are associated with the item.

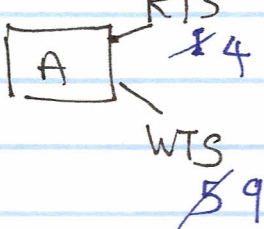
RTS: for an item:

The TS of the youngest transaction that read the item.

WTS: for an item:

The TS of the youngest transaction that performed write operation on the item.



t ₁	T ₁ . Read(A) ✓	RTS 1	
t ₂	T ₄ . Read(A) ✓	4	
t ₃	T ₂ . Read(A) -	4	
t ₄	T ₅ . Write(A)	WTS 5	
t ₅	T ₉ . Write(A)	9	
t ₆	T ₇ . Write(A)	9	