

## 12 Concurrency control

### 12.1 Serializability and Concurrency Control

### 12.2 Locking

#### Lock protocols

- Two phase locking
- Strict transactional protocols
- Lock conflicts and Deadlocks
- Lock modes
- Deadlock detection, resolution, avoidance

### 12.3 Nonlocking concurrency control

#### 12.3.1 Time stamp ordering

#### 12.3.2 Optimistic cc methods


#### 12.3.3 Multiversion cc

Lit.: Eickler/ Kemper chap 11.6-11.12, Elmasri /Navathe chap. 20, Garcia-Molina, Ullman, Widom: chap. 18

## Concurrency control...and serializability

### Wanted:

effective **real-time scheduling** of operations with guaranteed serializability of the resulting execution sequence.



12-CC-2

## Concurrency control

### Def.: Concurrency control (\*) in DBS

Methods for **scheduling the steps (operations) of database transactions** in a way which isolates concurrent transactions in order to guarantee serializability. ("between system start and shutdown").

(\*) "Synchronisierung"

12-CC-3

## Methods

### Approaches:

- (1) **Pessimistic**: Scheduler has to check if next incoming operation can be executed without compromising isolation.
- (2) **Optimistic**: Check for potential conflicts at the end of a TA. If yes: abort, else write effects into DB
- (3) **Multiversion cc (MVCC)**: orthogonal to (1), (2). More than one version of each data object allowed. May employ (1) or (2). Main advantage: readers preferred.  $\Rightarrow$  very useful for "read-mostly" databases

12-CC-4

## CC methods

### Primary concurrency control methods

1. **Locking** (most important)
2. **Non-locking protocols**:
  - Optimistic concurrency control
  - Time stamps
  - Multiversion CC (locking and non-locking variants)
3. **MVCC** more and more important

12-CC-5

## Concurrency control

### No explicit locking in application programs

- error prone,
  - responsibility of scheduler (and lock manager)
- In most DBS also explicit locking allowed in addition to implicit locking by scheduler. Use with care!

There are **read lock (shared locks)** and **write locks (exclusive)** - makes sense: **no read-read conflicts**

Not considered here: transaction independent locking, e.g. writing a page p to disk requires a short term lock (a "latch") on p


12-CC-6

## Optimistic vs. pessimistic

### Locking is pessimistic

Scenario: during operation  $op(x)$  of TA1 a (potentially) conflicting operation  $op'(x)$  of TA2 may access the same object  $x$ .

Avoided by locking  $x$  before accessing this object.



How long should a data objects be locked?

12-CC-7

## T12.2. Lock protocols

### Def.: Standard locking

1. Each object referenced by  $TA_i$  has to be locked before usage
2. Existing locks of other TA's will be respected i.e. wait until lock released.
3. Locks are released eventually.
4. A requests of a lock by a TA which it already holds, has no effect.
5. No preemption of a lock by another TA.

12-CC-8

## Lock protocols

### Standard object locking does not help....

Lock each object before reading / writing, unlock when operation finished

⇒ schedule may not be serializable (why?)

Example

$l1(x) r1(x) ul1(x)$        $l1(x) w1(x) ul(x)$   
 $l2(x) w2(x) ul2(x)$

⇒ lost update ⇒ useless

$l(x), ul(x)$  = lock / unlock  $x$

12-CC-9


## Lock protocols

### Preclaiming

- Acquire all locks needed before performing an operation
- **release, if you do not get all of them.** Try again. *race condition, transaction could starve!*
- Execute transaction
- Release locks

Preclaiming serializable?

Why (not) ?



**Bad:** objects to be processed may not be known in advance.

**Not used in DBS.**

12-CC-10

## Two phase locking (2PL)

### Def.: Two phase locking

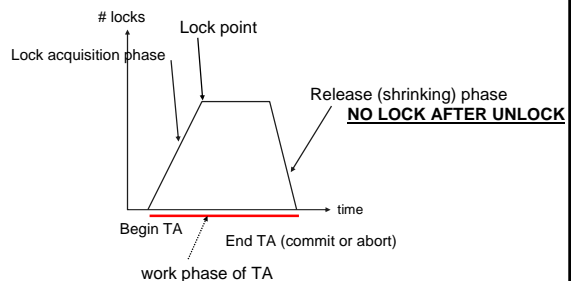
1. All rules of standard locking hold.
2. If a transaction TA **releases a lock** it holds on an object, TA **must not acquire a new lock** on any object.

**NO LOCK AFTER UNLOCK !**

Standard method for concurrency control in many database systems

12-CC-11

## Concurrency control: 2PL



Locked objects may be read / written already in lock acquisition phase !

12-CC-12

Concurrency control 2PL Freie Universität Berlin

### Why no lock after unlock?

TA1

```

l1(x)
r1(x)
x=x*10,
u11(x)

l1(y)
w1(y=x+y)
u11(y)
            
```

inconsistent (wrong!) y

TA2

```

l2(x),
r2(x),
x:=x+1,
w2(x)
u12(x)
            
```

TA1

Lock profile of TA1 is NOT 2PL

12-CC-13

Concurrency control 2PL Freie Universität Berlin

### 2-Phase locking theorem

If all transactions follow the **2-phase locking** protocol, the resulting history is **serializable**.

Proof sketch:

Suppose a resulting schedule is not serializable.  
 $\Rightarrow$  conflict graph contains a cycle  $\Rightarrow$  there are transactions TA1 and TA2 with conflict pairs (p,q) and (q', p').  
 $\Rightarrow$  One of them must have violated the "no lock after unlock" rule

(assuming a cycle of length 1, induction for the general case)

$\rightarrow$  12-CC-14

Concurrency control 2PL Freie Universität Berlin

Let e.g. (p,q) = (r1(x), w2(x)),  
 (q',p') = (w2(y), w1(y))

Analyze all of the possible execution sequences:

p, q, q', p

p, q', q, p'

q', p, q, p'

q', p, p', q

q', p', p, q

T2: l(y), T1: l(x), T2: l(x), T1: l(y)

T1 must have released lock on x and acquired one on y (or T2 must have acquired after release)  
 Violates 2-phase rule!  
 Contradiction to assumption that all TAs use 2PL protocol

Same holds for the other possible sequences  $\Rightarrow$  Theorem

12-CC-15

2PL = Serializable? Freie Universität Berlin

The **converse** of the 2PL theorem **does not hold**:  
 There are histories which are serializable, but the TAs did not lock according to 2 PL.

l1(x)r1[x] u11(x) l2(x) w2[x] ul2(x) c2 l(y)w1[y] ul(y) c1

Conflict graph

12-CC-16

Strict concurrency protocols: motivation Freie Universität Berlin

A different transaction TA2 could have used an object x which was unlocked by TA1 in the release phase.  
 ... no problem, if TA1 commits,  
 but abort... ??

Recursive situation  $\Rightarrow$  **cascading abort**

12-CC-17


Nonstrict example: even worse Freie Universität Berlin

Worse situation, why?

12-CC-18

## Strict 2PL

Locking protocol is **strict** if locks are released at commit / abort.



12-CC-19

## Lock modes

### Primary goal

**no harmful effects** (lost update, ...)

### Secondary goal

**Degree of parallelism** should be as high as possible, even when locking is used

**Low deadlock probability**, if any

Ways to increase parallelism

**Compatible locks** (read versus write semantics)

Different **lock granularity**

Application semantics

No locks: e.g. **time stamps**, optimistic cc

12-CC-20

## Lock modes

### Lock modes and lock compatibility

RW-(SX) – model: read (R) and eXclusive(W) locks (or: write locks)

	holder	R	W
requester	R	+	-
	W	-	-

R-lock = Shared (S) lock  
W-lock = X-lock

### Lock compatibility matrix

**Lock compatibility** in the RW model:

Objects locked in R-mode may be locked in R-mode by other transactions(+)

Objects locked in W-mode may not be locked by any other transaction in any mode.  
Lock conflict  $\Rightarrow$  requesting TA has to wait

12-CC-21

## Lock modes

### Hierarchical locking

– One single **lock granularity** (e.g. records) insufficient, large overhead when many rows have to be locked

$\Rightarrow$  Most DBS have at least two lock granularities: **row locks** and **table locks**

Issue: TA<sub>i</sub> wants to lock table R

- some rows of R locked by different transactions


$\Rightarrow$  different lock conflict as before: TA<sub>i</sub> is **waiting for** release of **all record locks**

- No other TA should be able to lock a record, otherwise TA<sub>i</sub> could starve

12-CC-22

## Concurrency control

### Locks of different granularity



Red square: Lock held by other transactions<sub>n</sub>


Yellow square: Lock request: must not be granted until lock on D is released

**Efficient implementation of this type of situation??**

12-CC-23

## Lock modes: Hierarchical locking

### Intention locks



Feature of **intention locks** for hierarchical locking: **for each lock mode, there is an intention lock**, e.g. for RX-lock modes: IR and IX

### Semantics:

A TA holds an Intention<sub>p</sub>-lock on an object D on level i, if and only if it holds an p-lock on an object D' on level j > i subordinate to D

12-CC-24


## Concurrency control

### Hierarchical locking

An object O on level i contains all objects x on level i+1

Locks of O lock all subordinate objects x

If a subordinate object x (level i+1) is locked, this is indicated by an intention lock on level i



### Lock escalation

If too many objects x on level i+1 are locked by a transaction, it may be converted into one lock on level i

12-CC-25

## Lock modes

### Hierarchical locking (cont)

Advantage: **one lookup** is sufficient to check if a lock on higher level (say on a table) can be granted

**Protocol:** if a TA wants to lock an object on level i in mode <M> (X or R), lock all objects on higher level (on the path to root) in I<M>-mode

Easy to check, if the locks on all subordinate objects are released: simply implement I<M>-lock as a counter

requester	holder			
	IR	IX	R	X
IR	+	+	+	-
IX	+	+	-	-
R	+	-	+	-
X	-	-	-	-

Compatibility matrix

12-CC-26

## Lock conflicts and deadlocks

### Lock conflict

Two or more processes request an exclusive lock for the same object

### Deadlock (\*)

Locking: always threat of deadlock if

- No preemption
- No lock release in case of lock conflicts

⇒ **Two-Phase locking may cause deadlocks**

$l_i(x)$  = Transaction i **requests** lock on x

$u_i(x)$  = Transaction i **releases** lock on x

Lock sequence:

$l_1(x), l_2(y), \dots, l_1(y), l_2(x)$  causes deadlock

(\*) Verklemmung

12-CC-27

## Deadlock detection and resolution

### Deadlocks

Release of a lock could break rule 4

$w_1(x), w_2(y), w_1(y) \rightarrow TA1: WAIT \text{ for } w_2(y), w_2(x) \rightarrow TA2: WAIT \text{ for } w_1(x)$

Note: **deadlocks very different from lock conflicts:**

$\dots w_1(x), w_2(y), w_1(y) \rightarrow TA1: WAIT \text{ for } w_2(y), w_2(z), w_2(y), w_2(z), \dots$

Lock conflict, y is locked by TA2, TA1 waits for unlock

Lock conflict resolved by  $w_2(x)$ , TA1 proceeds

Not schedules, but call sequences and lock / unlock operations provided by the scheduler

12-CC-28

## Deadlock

**Def.:** Wait-for graph  $WF = (T, E)$  where

- node set T is the set of running transactions
- edges  $(T_i, T_j) \in E \Leftrightarrow T_i$  required a lock on some object x which has been locked by  $T_j$  in an incompatible mode.

There is a deadlock between transactions

⇔ WF contains a cycle

⇒ System has to check WF for cycles periodically.

12-CC-29

## Deadlock resolution

### Resolving deadlocks

In case of cycle:

- One of the waiting transaction ("victim") has to be rolled back
- Which one? Heuristic decision: youngest, TA with least write activity, ...

### Timeout: an alternative?

- If TA has been waiting longer than the time limit, it is aborted.
- No: efficient but **may roll back innocent victims** (deadlock does not exist)

Oracle: WF-graph in central DB, timeout in distributed

12-CC-30