

Transactions Isolation

## Concurrency means trouble

 Several processes manipulating the same data at the same time can lead to inconsistencies



## Transactions

START TRANSACTION

#### COMMIT/ROLLBACK

- Atomic: Either the whole transaction is run, or nothing.
   Consistent: Database constraints are preserved.
- Isolated: Different transactions may not interact with each other.
- Durable: Effects of a transaction are not lost in case of a system crash.

## Serializability

 If the outcome of a number of transactions is equal to the outcome of the transactions executed without time overlap

$T_1$	$T_2$	A	B
		25	25
READ(A,t)			
t := t + 100		10000	
WRITE(A,t)		125	
	READ(A,s)		
	s := s*2		
	WRITE(A,s)	250	
	READ(B,s)		
	s := s*2		
	WRITE(B,s)		50
READ(B,t)			
t := t + 100			
WRITE(B,t)			150

To run transactions serially
 ...would solve the problem...

...but is often practically impossible

 Therefore we use different isolation levels



#### Lock compatibility

- To achieve different levels of isolation, you can use locks
- Shared (read) lock
- Exclusive (write) lock
- Locks are put on relevant parts of the data
- Lock requests are queued

LOCK	Write	Read
Write	No	No
Read	No	Yes

Isolation 0 – read uncommitted
Does not ask for read lock

 Therefore it can read data that another transaction has a write lock on

 "Dirty reads" – data modified by another transaction, but not yet committed

## Isolation 1 – read committed

 Asks for read lock, but releases it after reading

 "Non-repeatable reads" – the same query can give different results



## Isolation 2 – repeatable read

 If a query has a WHERE clause spanning a range, a read lock is acquired only for the result, not the entire range (no range lock)

 The result can't be changed, but new data can be added – so called "Phantoms"

<b>T1</b>	T2
SELECT * FROM users WHERE age BETWEEN 10 AND 30;	
	INSERT INTO users VALUES ( 3, 'Bob', 27 ); COMMIT;
SELECT * FROM users WHERE age BETWEEN 10 AND 30;	

## Isolation 3 - serializable

 No other transactions are allowed to interact with the data

Range locks are used

 All locks collected are kept until after COMMIT

## None of this is of course true...

### …at least not in ORACLE

## Different approaches

#### "Pessimistic"

Locks

#### "Optimistic"

 "We'd better make sure nothing funny is going on"  "If something funny is going on, we can take care of that later"

 Check before COMMIT that everything is in order, otherwise abort

## Which is better?

 Optimistic approach never blocks concurrent transactions

 But if conflicts happen often, the cost for aborting will be high

## How to do it? (1<sup>st</sup> example) Timestamp ordering

 Every transaction is given a timestamp T<sub>t</sub>

 Every object has two timestamps

 Last read T<sub>r</sub>
 Last write T<sub>w</sub>

  Read is only allowed if T<sub>t</sub>>T<sub>w</sub>, otherwise abort
 Set T<sub>r</sub>=T<sub>t</sub>

- Write is only allowed if T<sub>t</sub>>T<sub>r</sub> , otherwise abort
  - Set  $T_w = T_t$
  - SKIP write if T<sub>t</sub><T<sub>w</sub>! (Thomas write rule)

How to do it? (2<sup>nd</sup> example)
 Multiversion concurrency control
 Several timestamped versions of data exist

 If T<sub>t</sub>>T<sub>w</sub> the transaction can pick an older version to read, and does not have to abort (p 940)

 ORACLE uses something called Snapshot Isolation, based on MVCC

 A simple way to think of Oracle read consistency is to imagine each user operating a private copy of the database

# Not even the locks are that simple...

	NL	SCH-S	SCH-M	s	U	×	IS	IU	IX	SIU	SIX	UIX	BU	RS-S	RS-U	RI-N	RI-S	RI-U	RI-X	RX-S	RX-U	RX-X
NL	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	Ν	N	N	N	N
SCH-S	N	N	С	N	N	N	N	N	N	N	N	N	N	I	I	I	I	I	I	I	I	I
SCH-M	N	C	С	С	С	С	C	С	С	С	С	С	С	I	I	I	Ι	Ι	I	I	I	I
S	N	N	С	N	N	С	N	N	С	N	С	С	С	N	N	N	N	N	С	N	N	С
U	N	N	С	N	С	С	N	С	С	С	С	С	С	N	С	N	N	С	С	N	С	С
х	N	N	С	С	C	С	С	С	С	С	С	С	C	C	С	N	С	С	С	С	С	С
IS	N	N	С	N	N	С	N	N	N	N	N	N	С	I	I	I	I	Ι	I	I	I	I
IU	N	N	С	N	С	С	N	N	N	N	N	С	С	I	Ι	I	I	I	I	I	I	I
IX	N	N	С	С	С	С	N	N	N	С	С	С	С	I	I	I	Ι	I	I	I	I	I
SIU	N	N	С	N	C	С	N	N	С	N	С	С	C	I	Ι	I	I	I	I	I	I	I
SIX	N	N	С	С	С	С	N	N	С	С	С	С	С	I	I	I	Ι	Ι	I	I	I	Ι
UIX	N	N	С	С	C	С	N	С	С	C	С	С	C	I	Ι	I	I	I	Ι	I	I	I
BU	N	N	С	С	С	С	С	С	С	С	С	С	N	I	I	Ι	Ι	Ι	I	I	I	I
RS-S	N	I	I	N	N	С	I	Ι	I	I	Ι	I	I	N	N	С	С	С	С	С	С	С
RS-U	N	I	I	N	С	C	I	I	I	I	I	I	I	N	С	С	C	C	C	С	С	C
RI-N	N	I	I	N	N	N	I	Ι	I	I	I	I	Ι	C	С	N	N	N	N	С	С	С
RI-S	N	I	I	N	N	С	I	I	I	I	Ι	Ι	I	С	С	N	N	N	С	С	С	С
RI-U	N	I	I	N	C	С	I	Ι	I	I	I	I	I	C	С	N	N	С	С	С	С	С
RI-X	N	I	I	С	С	С	I	I	I	I	Ι	I	I	С	С	N	С	C	С	С	С	С
RX-S	N	I	I	N	N	С	I	Ι	I	I	I	I	Ι	C	С	С	С	С	С	С	С	С
RX-U	N	I	I	N	С	С	I	I	I	I	Ι	Ι	I	С	С	С	С	С	С	С	С	С
RX-X	N	I	I	С	C	C	I	I	I	I	I	I	I	C	C	C	C	C	С	С	C	C

#### Key

N	No Conflict	SIU	Share with Intent Update
I	Illegal	SIX	Shared with Intent Exclusive
С	Conflict	UIX	Update with Intent Exclusive
		BU	Bulk Update
NL	No Lock	RS-S	Shared Range-Shared
SCH-S	Schema Stability Locks	RS-U	Shared Range-Update
SCH-M	Schema Modification Locks	RI-N	Insert Range-Null
S	Shared	RI-S	Insert Range-Shared
U	Update	RI-U	Insert Range-Update
х	Exclusive	RI-X	Insert Range-Exclusive
IS	Intent Shared	RX-S	Exclusive Range-Shared
IU	Intent Update	RX-U	Exclusive Range-Update
IX	Intent Exclusive	RX-X	Exclusive Range-Exclusive