



DBTech VET / DBTechNet SQL Transactions*

Database Transactions Summit 2013
*HAAGA-HELIA University of Applied Sciences,
4 September 2013, Helsinki, Finland.*

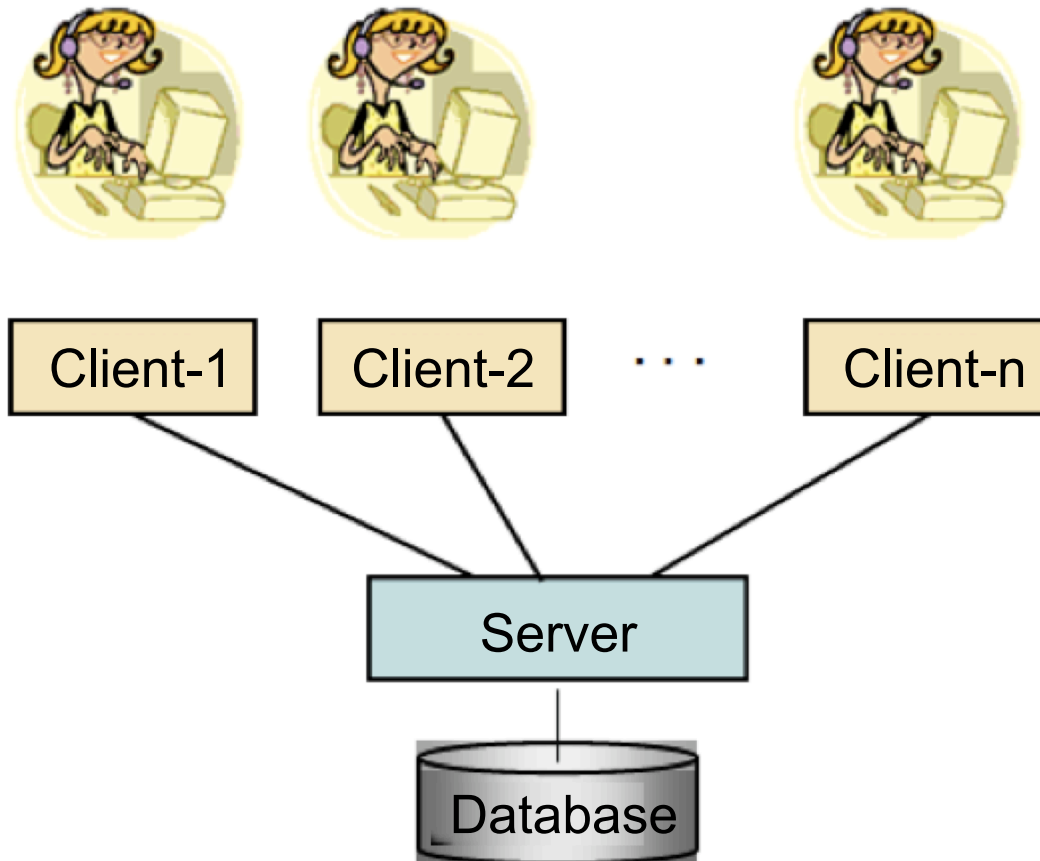
M. Laiho, D.A. Dervos, K. Silpiö
www.dbtechnet.org



The educational and training content of the present DBTech VET tutorial and hands-on laboratory session are licensed under a *Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License* (<http://creativecommons.org/licenses/by-nc-sa/3.0/deed.en>). Attributions must refer to the DBTech VET “SQL Transactions” course as a whole, in accordance with the directions provided at <http://www.dbtechnet.org/DBTechVET-CC-attributions-auidelines.PDF>.



Concurrently executing transactions





IF ... THEN

IF

- The DBMS is lacking the support of basic concurrency control (CC) services, or
- The programmer is lacking the knowledge of how to make proper use of the DBMS supported CC services

THEN

Data update operations may end up corrupting the DB data content



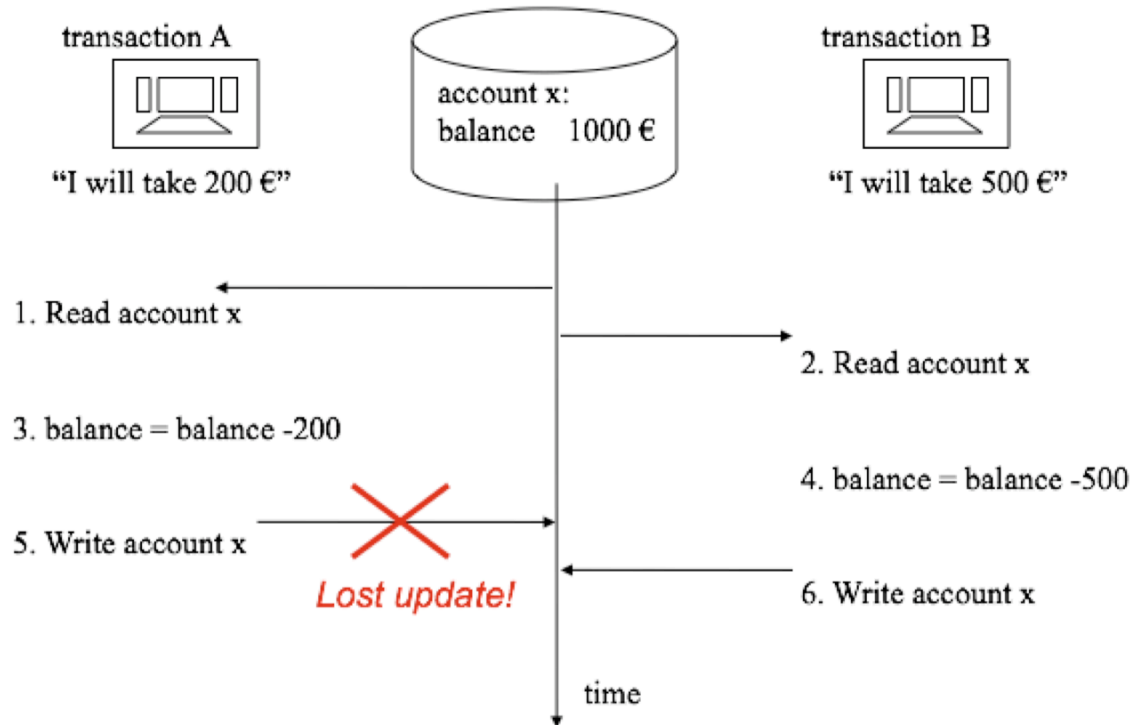
Concurrency problems (anomalies)

- **Lost update**
- Dirty read
- Non-repeatable read
- Phantom read



The lost update problem

"Tellers"



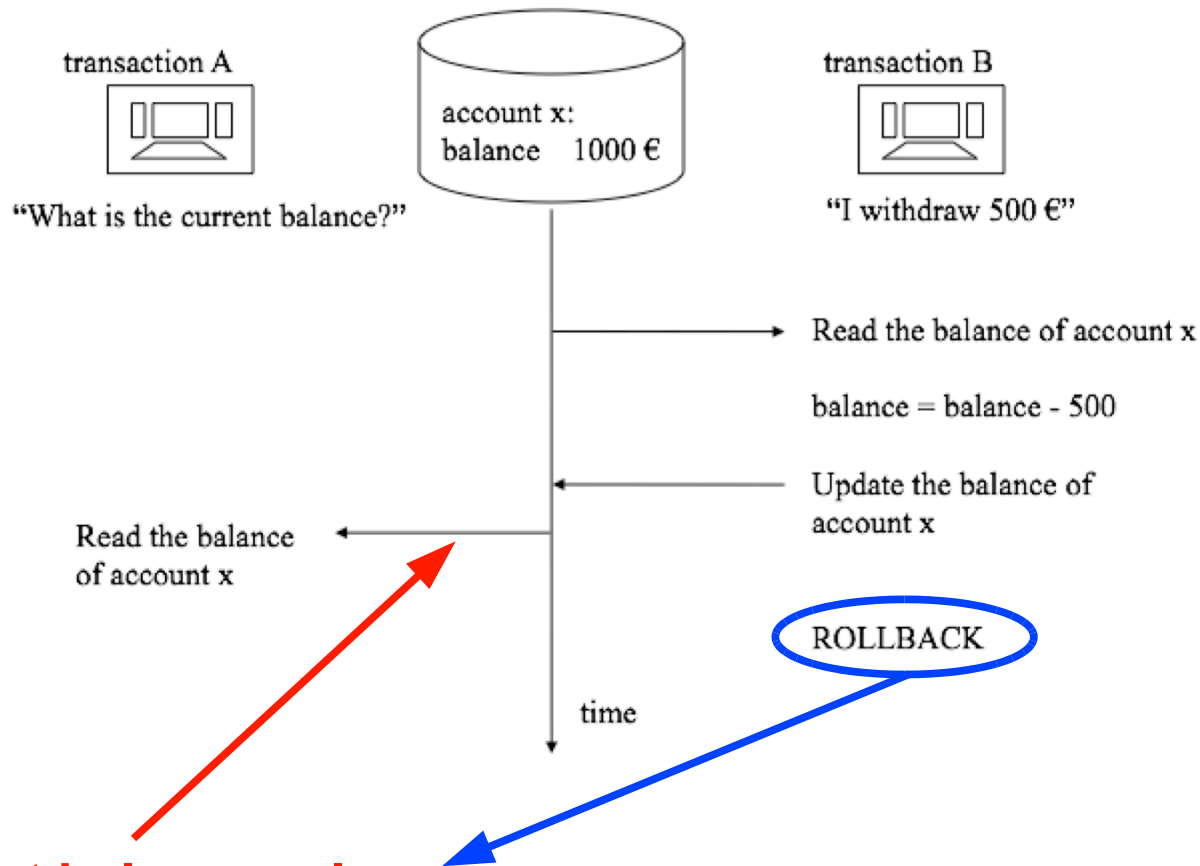


Concurrency problems (anomalies)

- Lost update
- Dirty read
- Non-repeatable read
- Phantom read



The dirty read problem



account balance value that never existed!




Concurrency problems (anomalies)

- Lost updates
- Dirty reads
- **Non-repeatable reads**
- Phantom reads





Non-repeatable reads

transaction A




transaction B






Non-repeatable reads



transaction A



transaction B




1. SELECT ... FROM table
WHERE ... ;






Non-repeatable reads


transaction A




1. SELECT ... FROM table
WHERE ... ;



transaction B




2. UPDATE table
SET c = ...
WHERE ... ;
DELETE FROM table
WHERE ... ;
...
COMMIT;






Non-repeatable reads


transaction A




1. SELECT ... FROM table
WHERE ... ;




3. SELECT ... FROM table
WHERE ... ;



4. COMMIT;



transaction B



2. UPDATE table
SET c = ...
WHERE ... ;
DELETE FROM table
WHERE ... ;
...
COMMIT;



Non-repeatable reads (NRR) vs. dirty reads (DR)

- The transaction “feels” changes made by other transactions (both NRR and DR)
- Repeating the same read operation may yield different results (both NRR and DR)
- Dirty reads (DR): the transaction “feels” changes made by other (concurrently running) transactions while the latter are still active (i.e. it is not yet known whether they will commit or rollback next)
- Non-repeatable reads (NRR): the transaction “feels” changes made by other (concurrently running) transactions only after they commit




Concurrency problems (anomalies)



- Lost updates
- Dirty reads
- Non-repeatable reads
- Phantom reads



transaction A




transaction B





Phantom reads



transaction A



transaction B




1. SELECT ... FROM table
WHERE ... ;

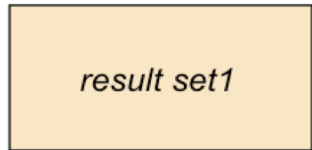




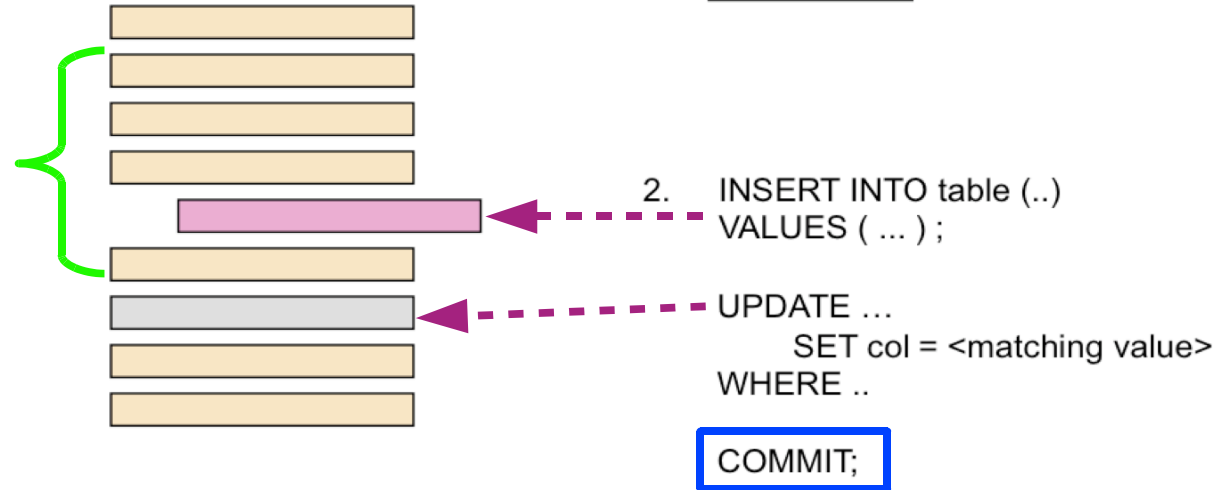

transaction A



1. SELECT ... FROM table
WHERE ... ;




transaction B

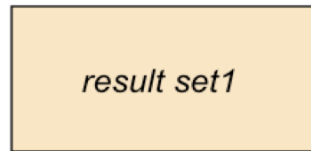





transaction A



1. SELECT ... FROM table
WHERE ... ;




3. SELECT ... FROM table
WHERE ... ;



4. COMMIT


transaction B



2. INSERT INTO table (..) VALUES (...) ;

UPDATE ... SET col = <matching value> WHERE ..

COMMIT;





Non-repeatable reads (NR) vs. phantom reads (PR)

- Rows out of nowhere (phantoms) do appear in both NRR and PR resultsets
- In NRR the affected transaction is assumed to be using the same search criterion (WHERE ...), repeatedly
- PR is more general: the affected transaction launches a new search criterion (WHERE ...) each one time.
- Phantom 'reads' because the targeted data/table regions may also involve 'ghost' rows (to be defined next)



A.C.I.D. properties

A transaction should execute in ...

Atomicity

... an ALL or NOTHING fashion

... Consistency

with regard to all the DBMS imposed data integrity rules

... Isolation

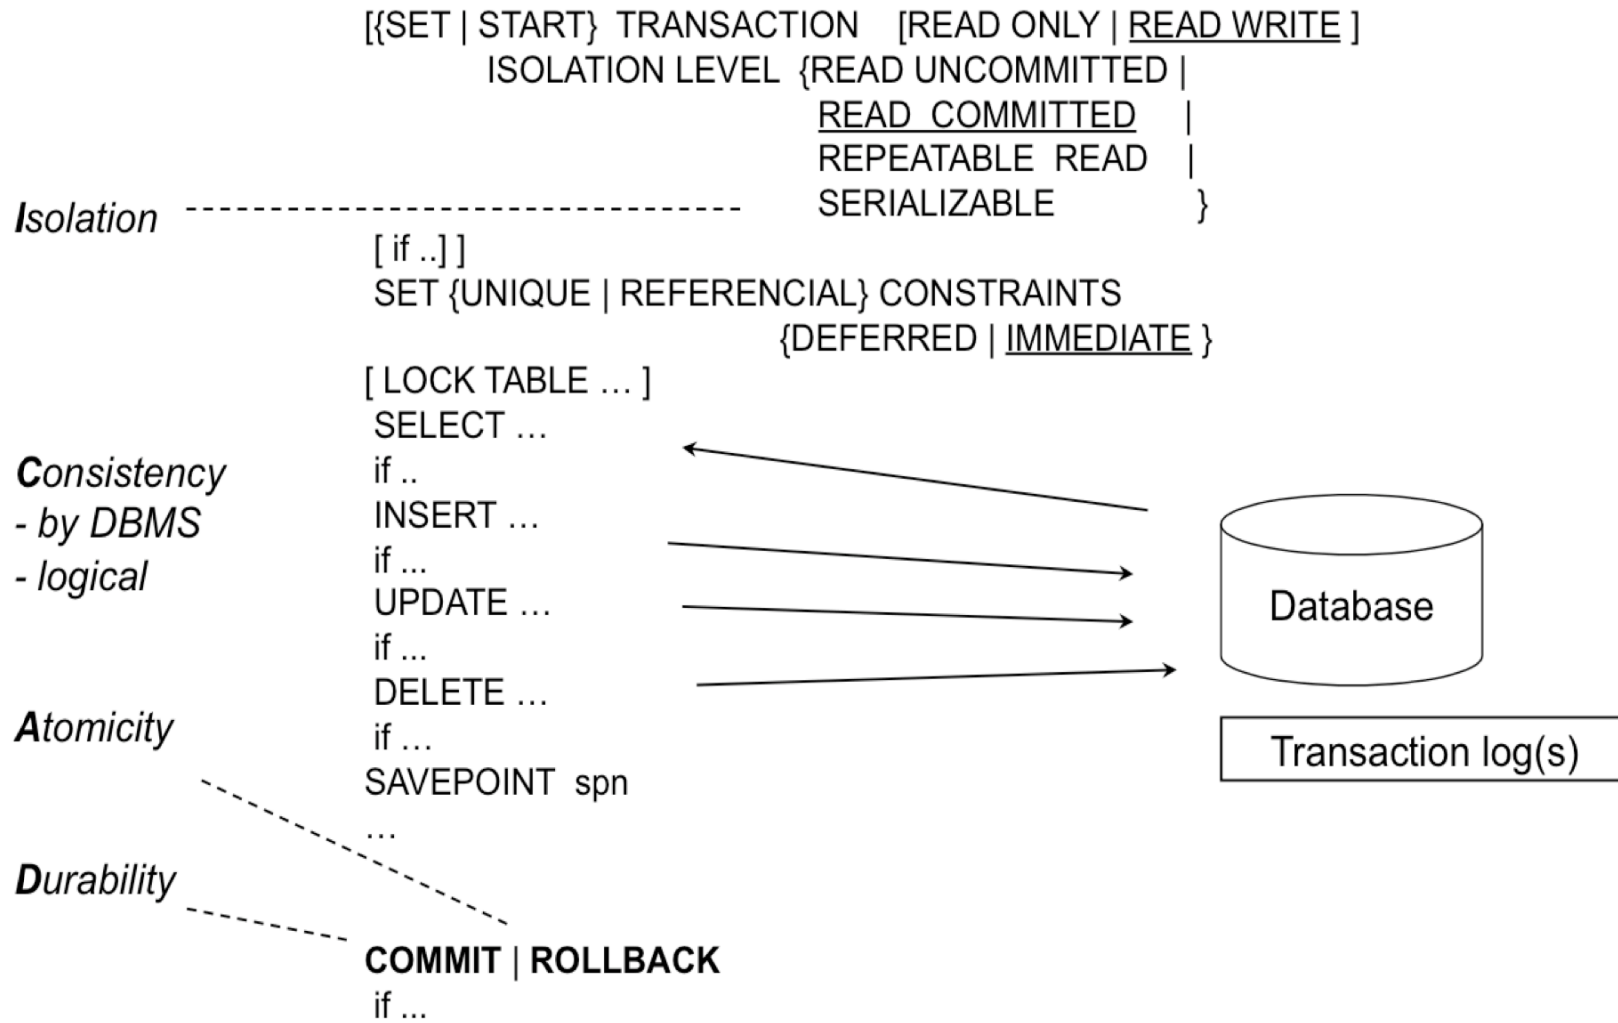
from what other concurrently running transactions do to the database content

Durability

... a way in which its COMMIT is guaranteed to make persistent all the changes made to the database

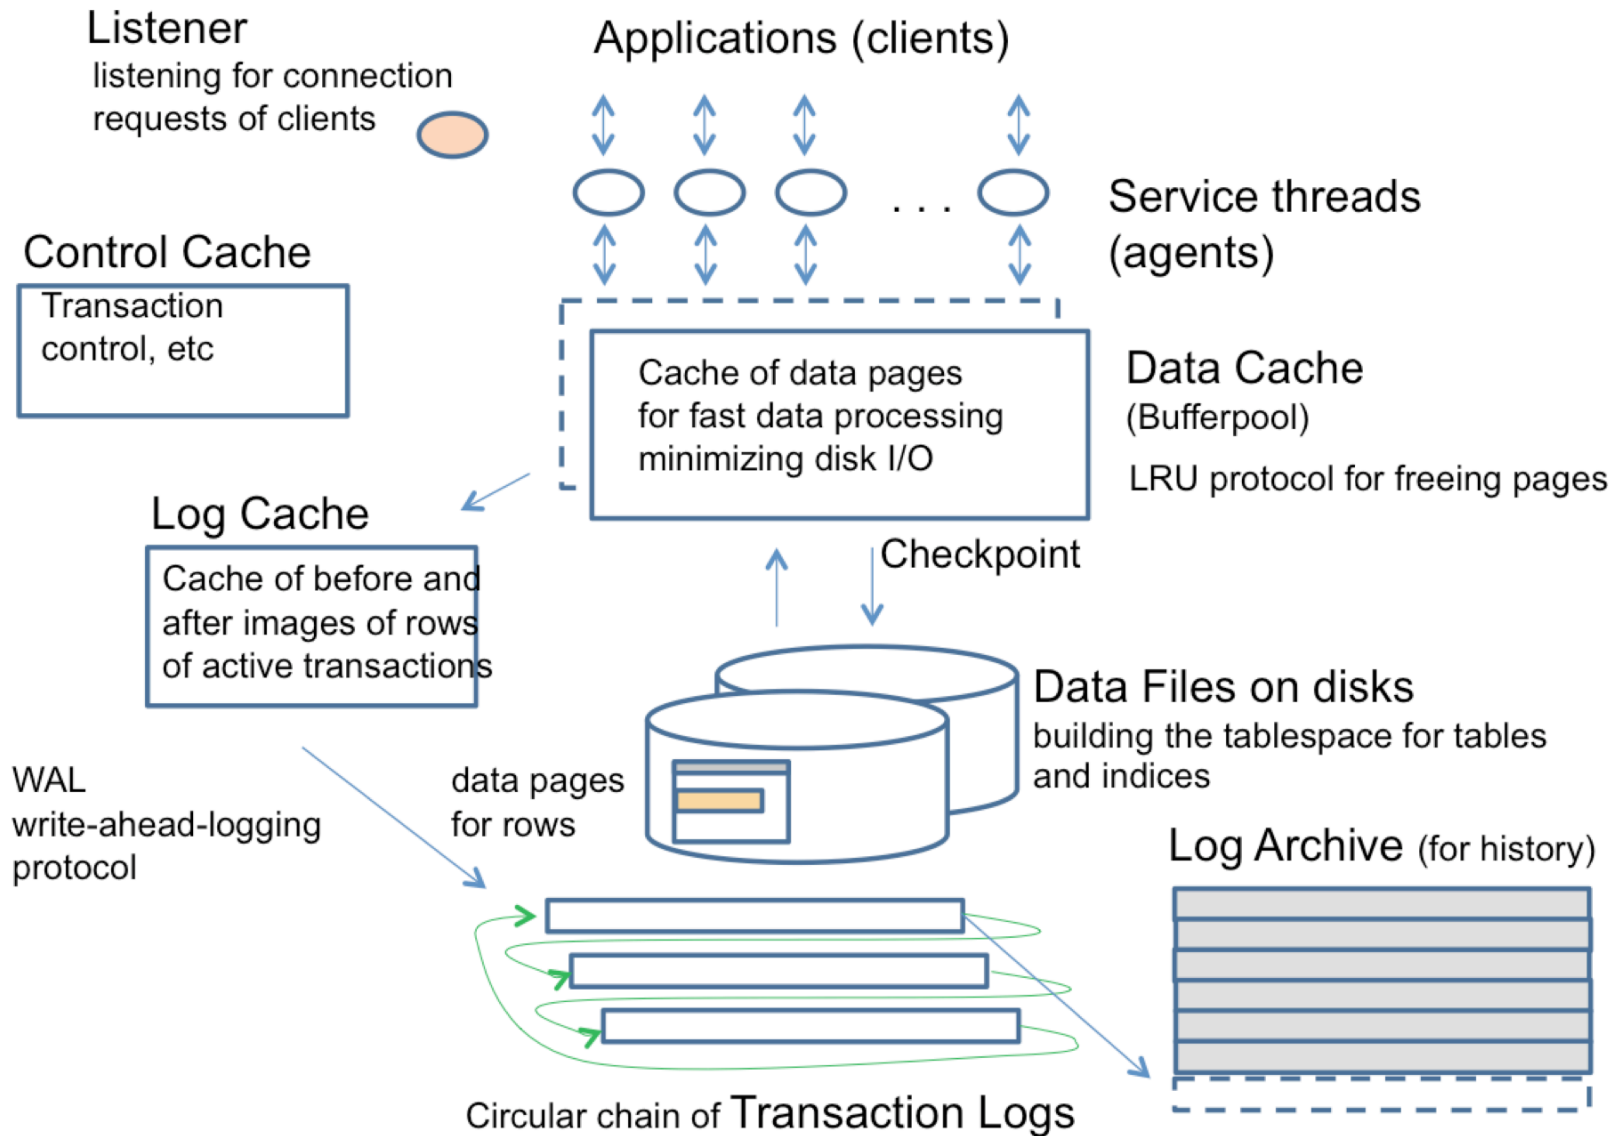


ISO SQL Transaction



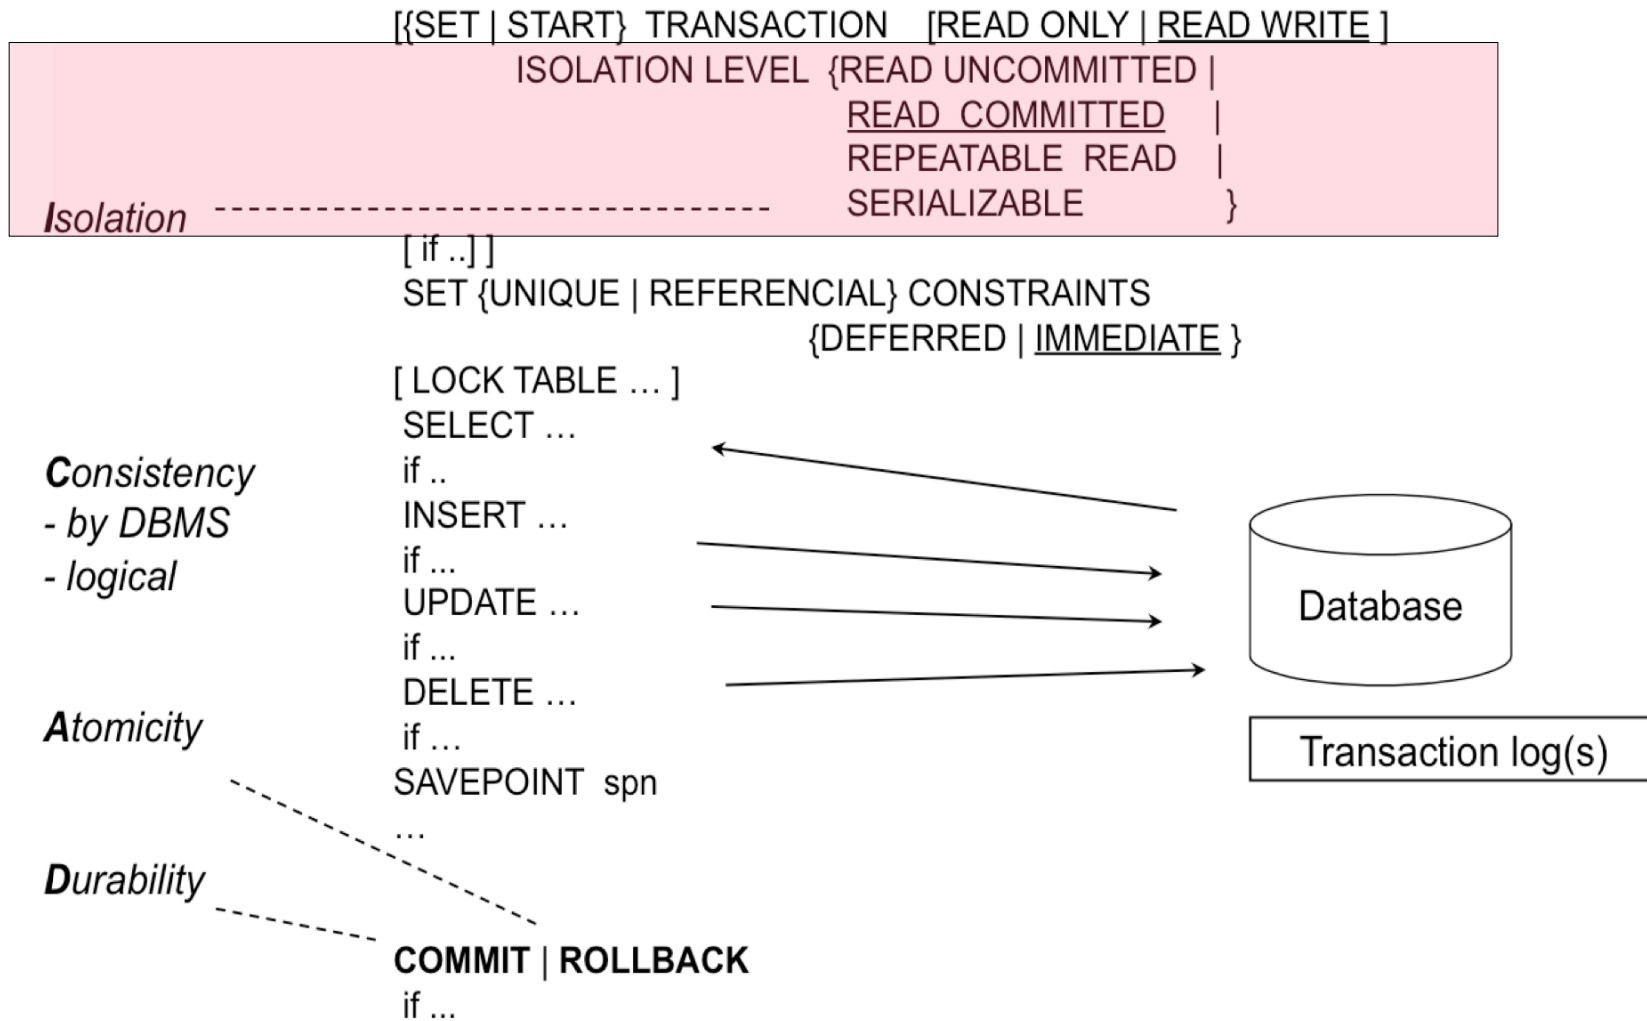


A & D: the underlying technology





ISO SQL Transaction





ISO SQL isolation levels defined

<i>Isolation Level:</i>	<i>Anomalies:</i>	<i>Lost Update</i>	<i>Dirty Read</i>	<i>Nonrepeatable Reads</i>	<i>Phantom Reads</i>
<i>READ UNCOMMITTED</i>		<i>NOT possible</i>	<i>Possible !</i>	<i>Possible !</i>	<i>Possible !</i>
<i>READ COMMITTED</i>		<i>NOT possible</i>	<i>NOT possible</i>	<i>Possible !</i>	<i>Possible !</i>
<i>REPEATABLE READ</i>		<i>NOT possible</i>	<i>NOT possible</i>	<i>NOT possible</i>	<i>Possible !</i>
<i>SERIALIZABLE</i>		<i>NOT possible</i>	<i>NOT possible</i>	<i>NOT possible</i>	<i>NOT possible</i>



- Multi-Granular Locking (MGL)
- Multi-Versioning (MVCC)
- Optimistic (OCC)



Compatibility of S and X locks:


Lock of transaction A to object o

Lock request of transaction B to object o

	<u>S</u> hared	e <u>X</u> clusive
<u>S</u> hared	Grant	Wait !
e <u>X</u> clusive	Wait !	Wait !

- S-lock grants read access to object
- X-lock grants write access to object
- X-lock request after getting S-lock is called as lock promotion

Locking granularity:





READ UNCOMMITTED:

- No S-lock protection for reading, long duration X-locks

READ COMMITTED:

- Short duration S-locks, long duration X-locks


REPEATABLE READ and SERIALIZABLE:

- Long duration S- and X-locks



A picture's worth a thousand words

PHENOMENA



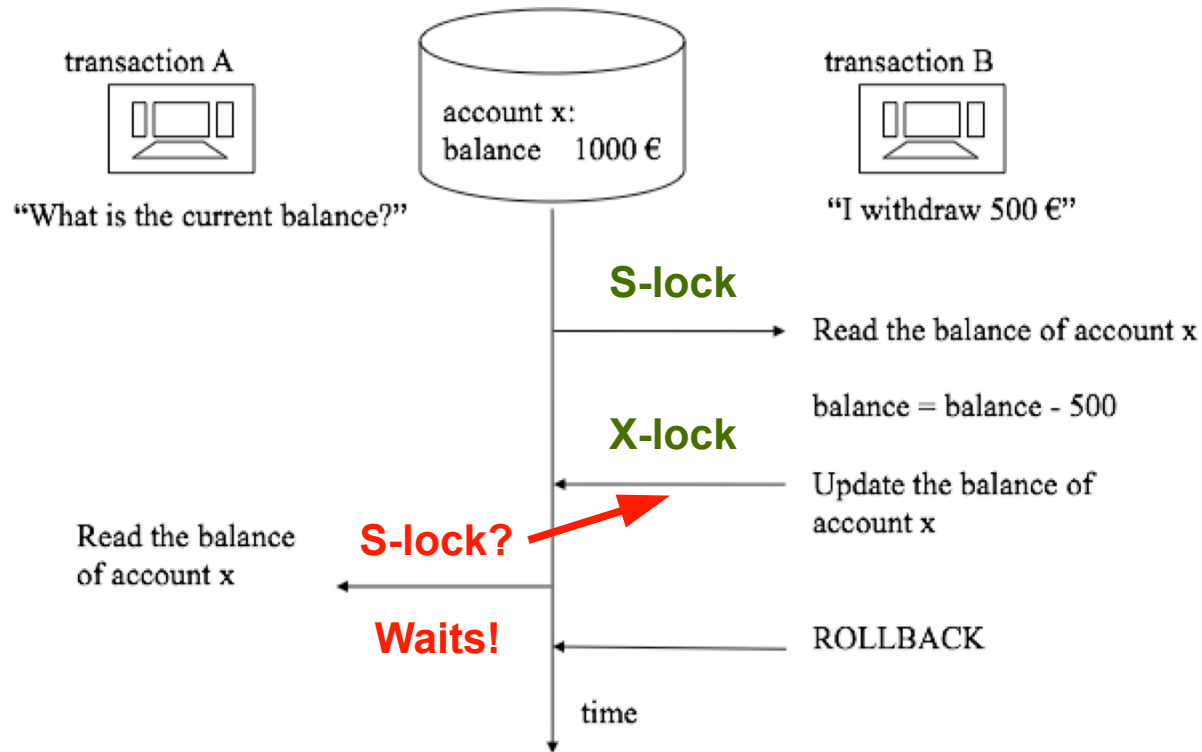
SQL-92 TRANSACTION ISOLATION LEVELS

Lock Durations

Exclusive Locks	long	long	long	long
Shared Locks (read locks)	<u>not used</u>	<u>short</u>	<u>long</u>	<u>long</u> (based on <u>the predicate</u>)



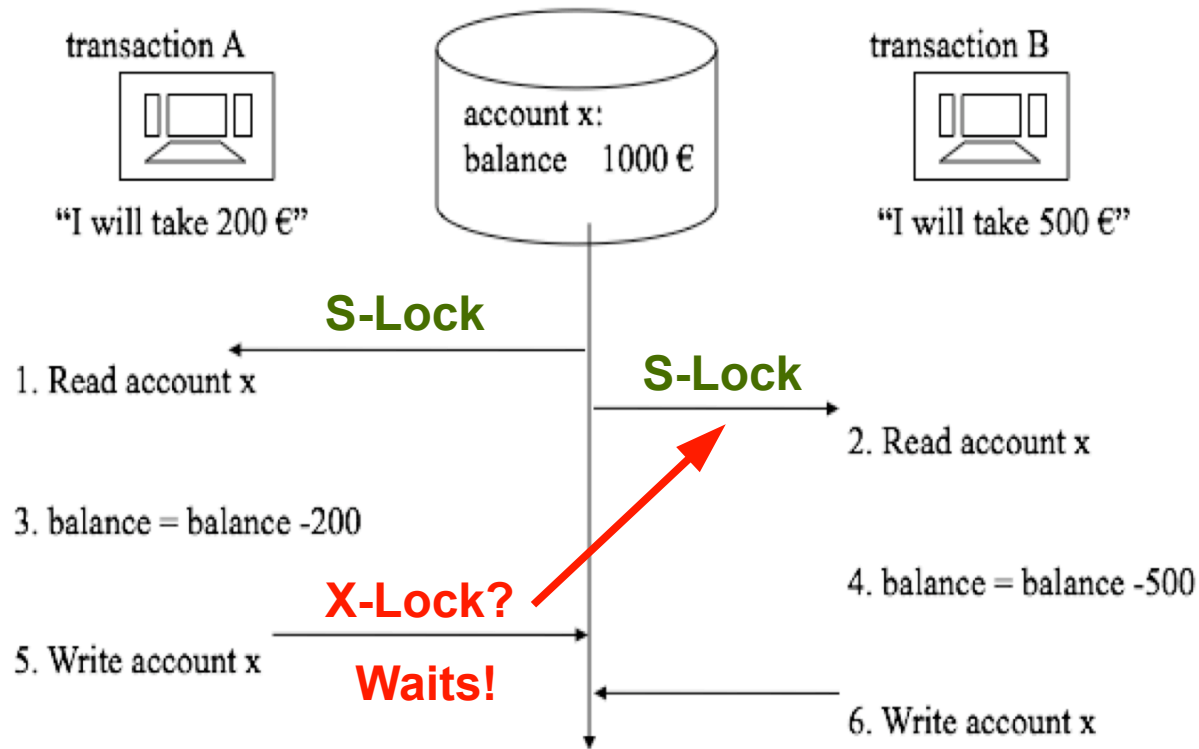
Dirty read with locks...



... problem resolved!

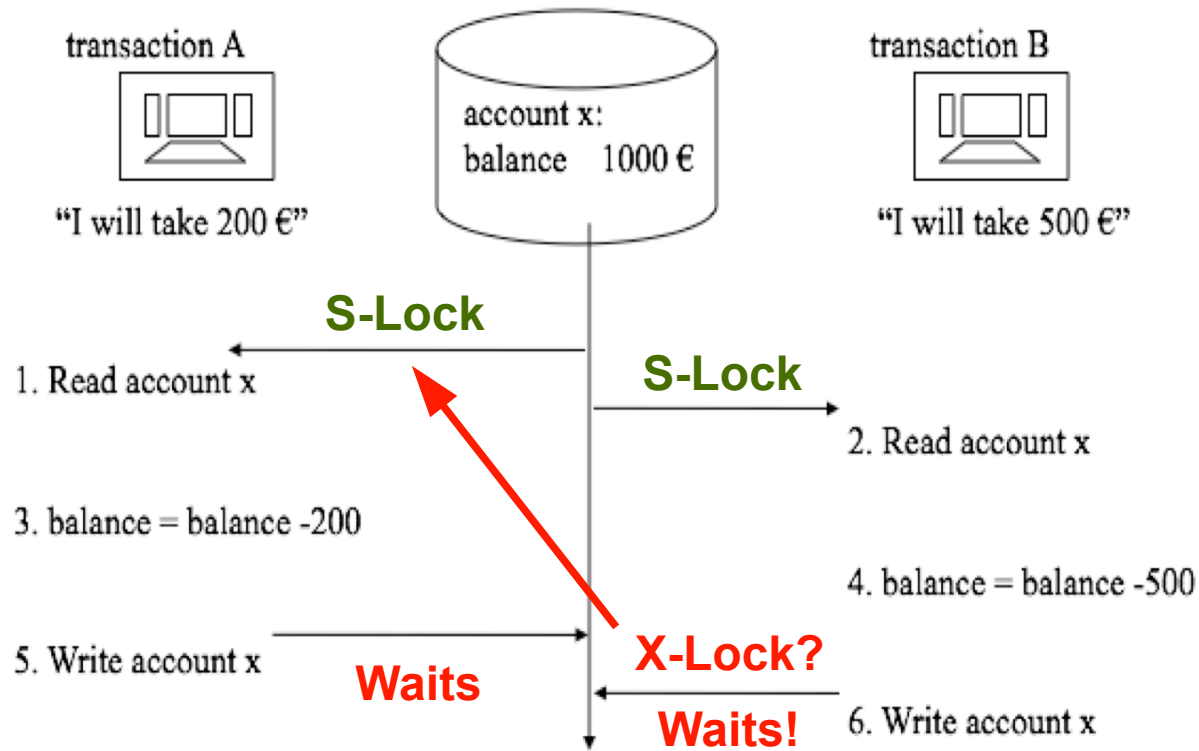


Lost update with locks



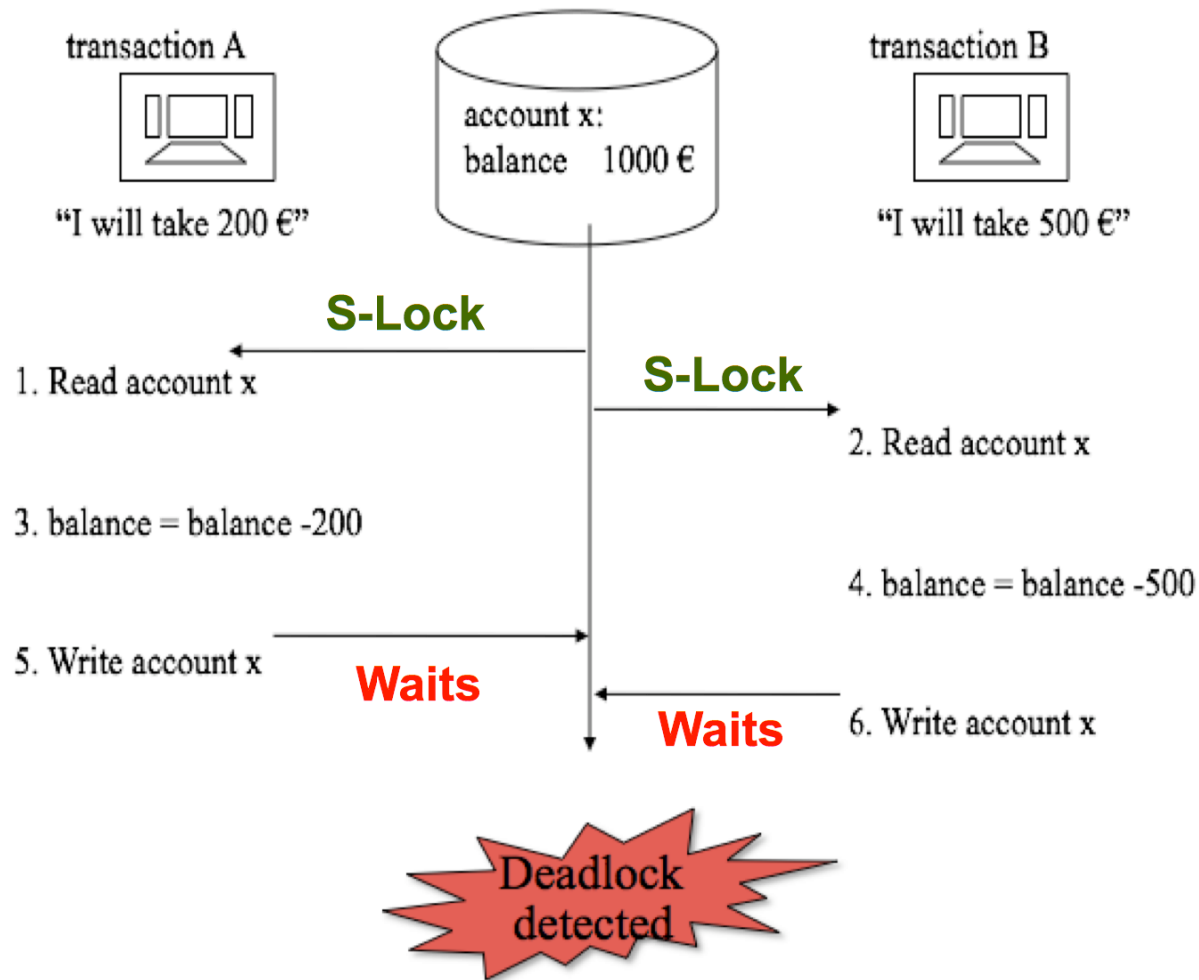


Lost update with locks





Lost update with locks



... one of (A,B) need be **rolled back**



Using ANSI/ISO SQL, the “lost update” scenario is implemented with transactions A and B making use of (single) SQL UPDATE statements, instead of conducting (separate) READ and WRITE operations, e.g.:

```
UPDATE Accounts SET balance = balance – 200  
WHERE acctID = 100;
```

In most of today's DBMS's: lock-based CC protection is in effect, and the above gets resolved in a deadlock-free manner...

...resolved?



- Concurrent transactions involving more than one sensitive update SQL statements, plus
- Clumsy programming at the application side...

...may very well lead to having the lost update anomaly appear, even when lock-based concurrency control is implemented at the server side.

The solution to the problem: the system's state need be carefully inspected, in systematic way, following the execution of each one (any) SQL statement: GET DIAGNOSTICS



- part of the ISO SQL Standard
- implemented in MySQL v5.6, and MariaDB

e.g. in MySQL v5.6 (to be considered during the HoL session):

```
GET DIAGNOSTICS @rowcount = ROW_COUNT;
```

```
GET DIAGNOSTICS CONDITION 1 @sqlstate = RETURNED_SQLSTATE,  
@sqlcode = MYSQL_ERRNO ;
```

```
SELECT @sqlstate, @sqlcode, @rowcount;
```



- Sample variants of lock compatibility matrices

Lock granules:

database


(tablespace)

table

(extent)

page

row



Lock requested:	Lock already granted to some other process				
	IS	IX	S	SIX	X
IS	grant	grant	grant	grant	wait
IX	grant	grant	wait	wait	wait
S	grant	wait	grant	wait	wait
SIX	grant	wait	wait	wait	wait
X	wait	wait	wait	wait	wait

$$SIX = S + IX$$

1. Intent locks
IS for S on row
IX for X on row



2. Lock on row

Lock requested:	Lock already granted to some other process			
	none	S	U	X
S	grant	grant	grant ³	wait
U	grant	grant	wait	wait
X	grant	wait	wait	wait

Shared locks (S) allow reading.
eXclusive locks (X) allow writing and are kept up to end of transaction eliminating lost updates.

Other locks on index ranges, schemas



Concurrency control: implementation of

- Multi-Granular Locking (MGL)
- Multi-Versioning (MVCC)
- Optimistic (OCC)



A database where several tables updated frequently:

- Want to read uncommitted, but do not wish to compromise on data consistency
- Equivalently: wish to conduct data reading on a (logical) snapshot of the DB content, taken at begin time of the read-only transaction



- The DB server makes use of timestamps to **maintain history chains (versions), one for each row that is being updated**
- Considering the above, any one transaction may either read the latest committed version of each row **at read time (READ COMMITTED)**, or the latest committed version of each one row **at its (transaction) begin time (SNAPSHOT)**

Consequently, readers and writers do not block each other:
improved performance

** extra overhead, or...*



Snapshot isolation (SI)

- Improved performance (readers and writers do not block each other, fewer deadlocks)
- Write-write conflicts are handled by policies that depend on the DBMS used:
 - (a) ORACLE uses some type of locking; the second writer waits,
 - (b) under SolidDB, the first writer is the only one who is allowed to commit



- A transaction never accesses phantom rows, i.e. rows that have been inserted by other transactions after its begin time(stamp): compare/contrast with MGL
- A transaction may access ghost rows, i.e. rows that have been deleted or updated by other transactions after its begin time(stamp): compare/contrast with MGL



- SI is not appropriate for implementing the ISO SERIALIZABLE isolation level
- MySQL uses it only for implementing the REPEATABLE READ isolation level

Example*

T1: copies x into y

T2: copies y into x

Initially: $(x,y) = (8,10)$

Possible outcomes

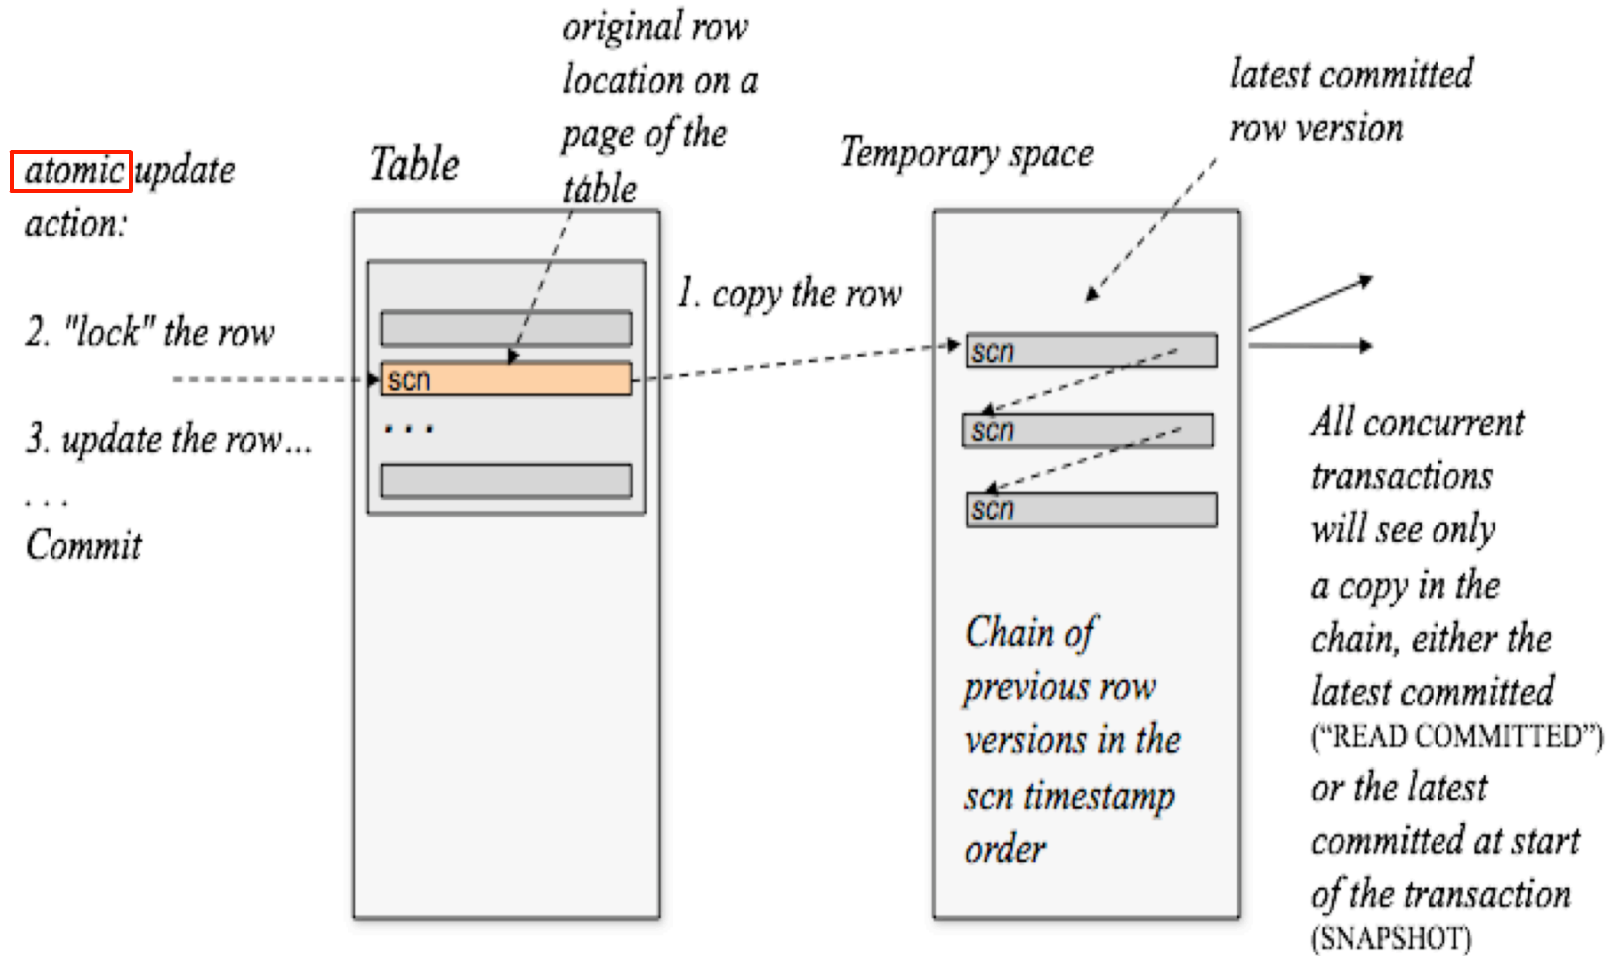
Under ISO SERIALIZABLE: $(8,8)$, $(10,10)$

With SI: $(8,8)$, $(10,10)$, $(10,8)$

* Philip A. Bernstein and Eric Newcomer, *Principles of Transaction Processing, 2nd Edition*, Morgan Kaufmann, 2009



MVCC implementation: Oracle



scn: system change number



- **READ UNCOMMITTED:** MVCC (*read latest written*)
- **READ COMMITTED:** MVCC (*read latest committed*)
- **REPEATABLE READ:** MVCC (*snapshot isolation*)
- **SERIALIZABLE:** MGL (*MV with long locks on the latest version of the rows read/written*)



- Appropriate for situations where DB data are cached outside the DB server's cache memory
- The application reads data from the (remote) cache memory in an optimistic way (i.e. hoping that the corresponding DB server residing data have not been altered/updated in the meantime)
- Variations exist (e.g. MS-SQL Server's *optimistic with values*, and *optimistic with versions*)
- Original OCC (Phyrrho DBMS): all changes made to data by a transaction are kept separate from the database and they get registered to/synchronized with the latter at commit time



?



The DBTechNET DebianDB VM