Introduction to Database Systems
CSE 444

Lectures 13-14

Transactions: Isolation & ARIES

Today's Outline

1. User interface:
   1. Read-only transactions
   2. Weak isolation levels
   3. Transaction implementation in commercial DBMSs
2. The ARIES recovery method

- Reading: M. J. Franklin. "Concurrency Control and Recovery". Posted on class website

READ-ONLY Transactions

```
Client 1: START TRANSACTION
          INSERT INTO SmallProduct(name, price)
                SELECT pname, price
                FROM Product
                WHERE price <= 0.99

          DELETE  FROM Product
                  WHERE price <=0.99
          COMMIT

Client 2: SET TRANSACTION READ ONLY
          START TRANSACTION
          SELECT count(*)
          FROM Product

          SELECT count(*)
          FROM SmallProduct
          COMMIT
```

Can help DBMS improve performance

Isolation Levels in SQL

1. "Dirty reads"

   SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED

2. "Committed reads"

   SET TRANSACTION ISOLATION LEVEL READ COMMITTED

3. "Repeatable reads"

   SET TRANSACTION ISOLATION LEVEL REPEATABLE READ

4. Serializable transactions

   SET TRANSACTION ISOLATION LEVEL SERIALIZABLE

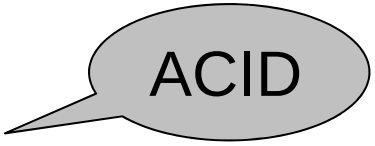

Choosing Isolation Level

- Trade-off: efficiency vs correctness
- DBMSs give user choice of level

Always read DBMS docs!

Beware!!
- Default level is often NOT serializable
- Default level differs between DBMSs
- Some engines support subset of levels!
- Serializable may not be exactly ACID

1. Isolation Level: Dirty Reads

Implementation using locks:

- "Long duration" WRITE locks
  - A.k.a Strict Two Phase Locking (you knew that !)
- No READ locks
  - Read-only transactions are never delayed

Possible problems: dirty and inconsistent reads

2. Isolation Level: Read Committed

Implementation using locks:

- "Long duration" WRITE locks
- "Short duration" READ locks
  - Only acquire lock while reading (not 2PL)

- Possible problems: unrepeatable reads
  - When reading same element twice, may get two different values

2. Read Committed in Java

In the handout: Lecture13.java - Transaction 1:

```
db.setTransactionIsolation(Connection.TRANSACTION_READ_COMMITTED);
db.setAutoCommit(false);
readAccount();
Thread.sleep(5000);
readAccount();
db.commit();
```

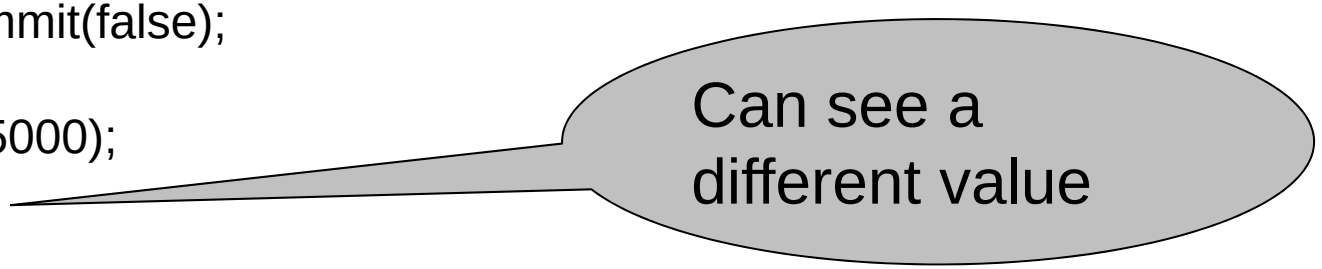

In the handout: Lecture13.java – Transaction 2:

```
db.setTransactionIsolation(Connection.TRANSACTION_READ_COMMITTED);
db.setAutoCommit(false);
writeAccount();
db.commit();
```

3. Isolation Level: Repeatable Read

Implementation using locks:

- "Long duration" READ and WRITE locks
  - Full Strict Two Phase Locking
- This is not serializable yet !!! (Why?)

3. Repeatable Read in Java

In the handout: Lecture13.java - Transaction 1:

```
db.setTransactionIsolation(Connection.TRANSACTION_REPEATABLE_READ);
db.setAutoCommit(false);
readAccount();
Thread.sleep(5000);
readAccount();
db.commit();
```

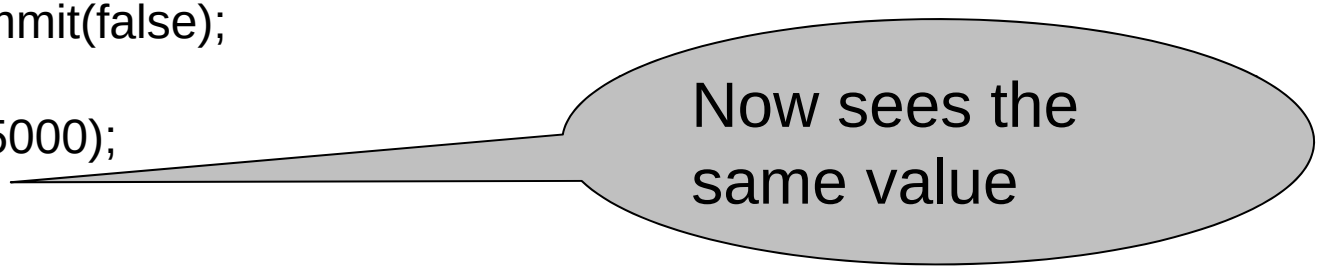

In the handout: Lecture13.java – Transaction 2:

```
db.setTransactionIsolation(Connection. TRANSACTION_REPEATABLE_READ);
db.setAutoCommit(false);
writeAccount();
db.commit();
```

3. Repeatable Read in Java

In the handout: Lecture13.java – Transaction 3:

```
db.setTransactionIsolation(Connection.TRANSACTION_REPEATABLE_READ);
db.setAutoCommit(false);
countAccounts();
Thread.sleep(5000);
countAccounts();
db.commit();
```

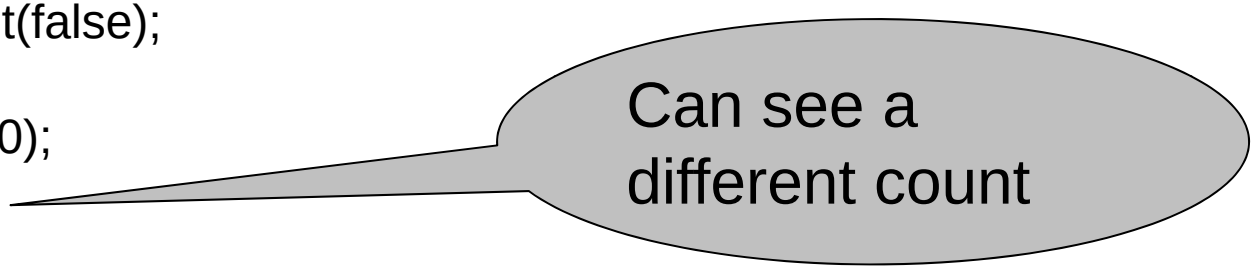

In the handout: Lecture13.java – Transaction 4:

```
db.setTransactionIsolation(Connection.TRANSACTION_REPEATABLE_READ);
db.setAutoCommit(false);
insertAccount();
db.commit();
```

Note: In PostgreSQL will still see the same count.
But not serializable in general (i.e., other DBs).

4. Serializable in Java

In the handout: Lecture13.java – Transaction 3:

```
db.setTransactionIsolation(Connection. TRANSACTION_SERIALIZABLE);
db.setAutoCommit(false);
countAccounts();
Thread.sleep(5000);
countAccounts();
db.commit();
```

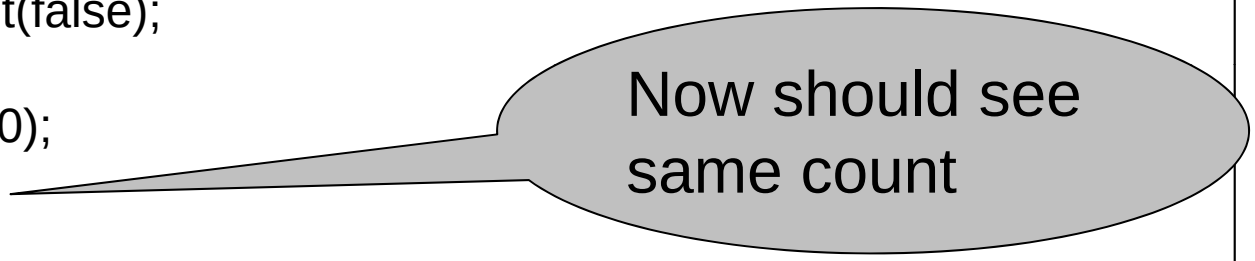

In the handout: Lecture13.java – Transaction 4:

```
db.setTransactionIsolation(Connection. TRANSACTION_SERIALIZABLE);
db.setAutoCommit(false);
insertAccount();
db.commit();
```

Commercial Systems

- DB2: Strict 2PL
- SQL Server:
  - Strict 2PL for standard 4 levels of isolation
  - Multiversion concurrency control for snapshot isolation
- PostgreSQL:
  - Multiversion concurrency control
- Oracle
  - Snapshot isolation even for SERIALIZABLE(!)

Today's Outline

1. User interface
2. The ARIES recovery method

- Reading: M. J. Franklin. "Concurrency Control and Recovery". Posted on class website

ARIES Overview

- Undo/redo log with lots of clever details
- Physiological logging
- Each log entry has unique Log Sequence Number, LSN

Granularity in ARIES

- Physical logging for REDO (element=one page)
- Logical logging for UNDO (element=one record)
- Result: logs logical operations within a page
- This is called *physiological logging*
- Why this choice?
  - Must do physical REDO since cannot guarantee that db is in an action-consistent state after crash
  - Must do logical undo because ARIES will only undo loser transactions (this also facilitates ROLLBACKs)

The LSN

- Each log entry receives a unique *Log Sequence Number*, LSN
  - The LSN is written in the log entry
  - Entries belonging to the same transaction are chained in the log via **prevLSN**
  - LSN's help us find the end of a circular log file:

After crash, log file = (22, 23, 24, 25, 26, 18, 19, 20, 21)
Where is the end of the log ? 18

Aries Data Structures

- Each page on disk has **pageLSN**:
  = LSN of the last log entry for that page
- Transaction table: each entry has **lastLSN**
  = LSN of the last log entry for that transaction
  Transaction table tracks all active transactions
- Dirty page table: each entry has **recoveryLSN**
  = LSN of earliest log entry that made it dirty
  Dirty page table tracks all dirty pages
- Txn and dirty page tables in main memory

Checkpoints

- Write into the log
  - Contents of transactions table
  - Contents of dirty page table
- Very fast ! No waiting, no END CKPT
- But, effectiveness is limited by dirty pages
  - There is a background process that periodically sends dirty pages to disk

ARIES Recovery in Three Steps

- **Analysis pass**
  - Figure out what was going on at time of crash
  - List of dirty pages and running transactions
- **Redo pass (repeating history principle)**
  - Redo all operations, even for transactions that will not commit
  - Get back state at the moment of the crash
- **Undo pass**
  - Remove effects of all uncommitted transactions
  - Log changes during undo in case of another crash during undo

ARIES Method Illustration

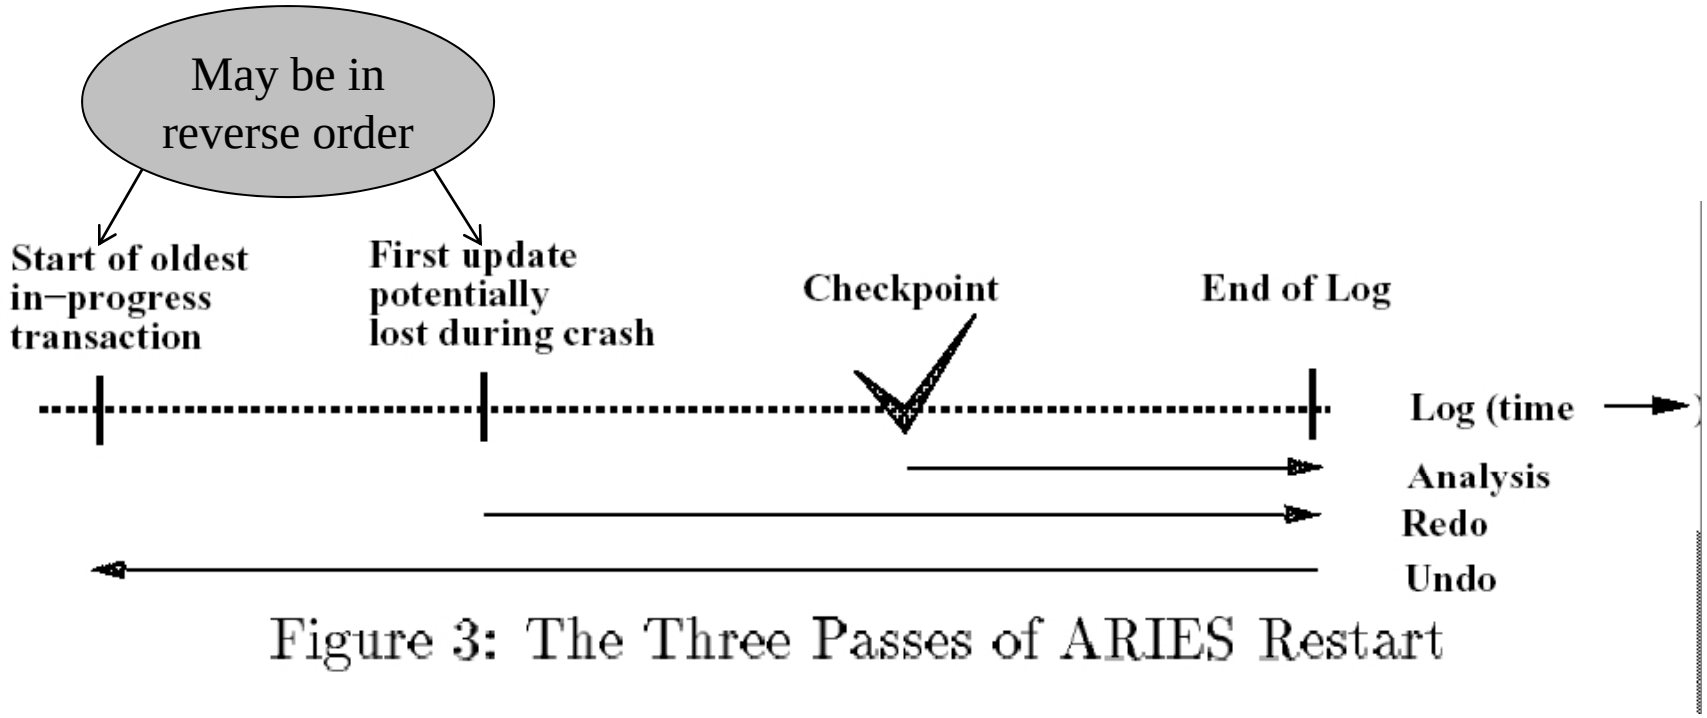

Figure 3: The Three Passes of ARIES Restart

[Franklin97]

Analysis Phase

- Goal
  - Determine point in log where to start REDO
  - Determine set of dirty pages when crashed
    - Conservative estimate of dirty pages
  - Identify active transactions when crashed

- Approach
  - Rebuild transactions table and dirty pages table
  - Start from the latest checkpoint
  - Scan the log, and update the two tables accordingly
  - Find oldest recoveryLSN (firstLSN) in dirty pages tables

Redo Phase

- Goal: redo all updates since firstLSN
- For each log record
  - If affected page is not in the Dirty Page Table then **do not update**
  - If affected page is in the Dirty Page Table but recoveryLSN > LSN of record, then **no update**
  - Else need to read the page from disk; if pageLSN > LSN, then **no update**
  - Otherwise perform update

Undo Phase

- Goal: undo effects of aborted transactions
- Identifies all loser transactions in trans. table
- Scan log backwards
  - Undo all operations of loser transactions
  - Undo each operation unconditionally
  - All ops. logged with compensation log records (CLR)
  - Never undo a CLR
    - Look-up the UndoNextLSN and continue from there

Handling Crashes during Undo

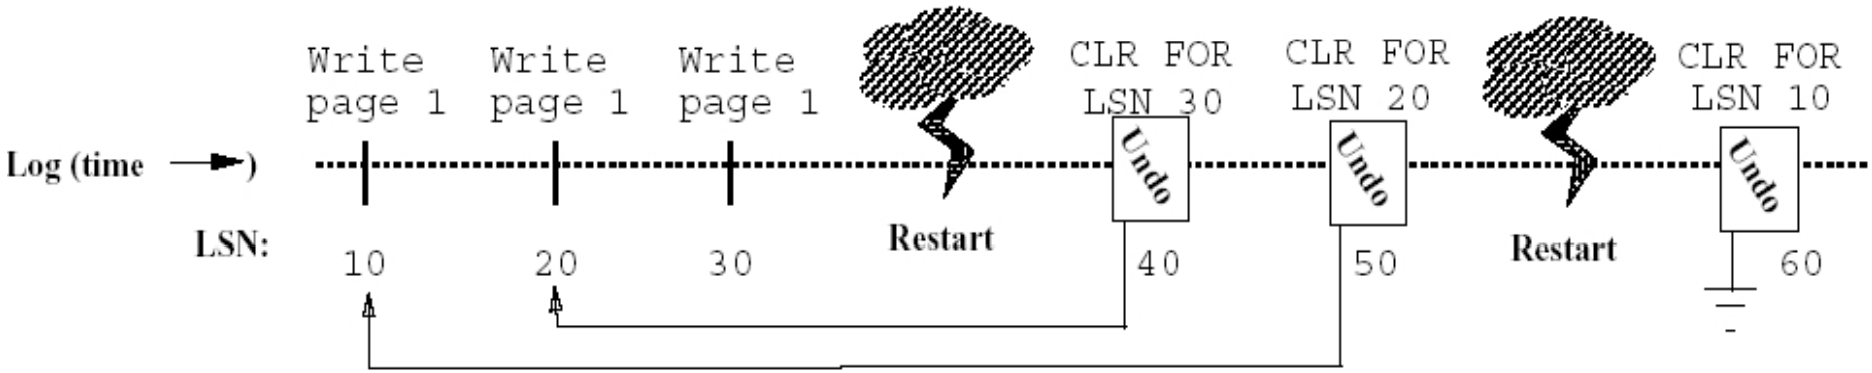

Figure 4: The Use of CLRs for UNDO

[Franklin97]