

Transaction Processing

Enterprise Scale Data Management

Divy Agrawal
Department of Computer Science
University of California at Santa Barbara

The Transaction Concept

- Multiple online users:
 - Gives rise to the **concurrency problem**.
- Component unreliability:
 - Gives rise to the **failure problem**.
- Database designers confronted these problems in the context of managing persistent data:
 - Online transaction processing system
 - Design, implementation, and operation of large application system with hundreds of terminals, tens of computers, providing service with no downtime, guaranteeing application correctness and data consistency.

Winter 2013: CMPSC 274 1/9/13

2

The Transaction Concept

- Transactions provide an integrative framework in the presence of many "moving parts".
- Distributed transaction-oriented systems are the enabling technology:
 - Distributed and Networked applications
 - E-commerce and Workflow systems
 - Large-scale Information Infrastructures
- Without transactions, distributed systems/networked applications cannot be made to work.

Winter 2013: CMPSC 274 1/9/13

3

+ The Transaction Concept

- Transactions were originally developed in the context of DBMS as a paradigm to deal with:
 - Concurrent access to shared data
 - Failures of different kinds/types.
- Typical and canonical application scenarios in the context of banking application: Debit/Credit operations, and fund Transfers.
- The key problem solved in an elegant manner:
 - Subtle and difficult issue of keeping data consistent in the presence of concurrency and failures

while ensuring performance, reliability, and availability.

Winter 2013: CMPSC 274 1/9/13

+ OLTP Example: Debit/Credit

```

void main () {
    EXEC SQL BEGIN DECLARE SECTION
    int BAL, AID, amount;
    EXEC SQL END DECLARE SECTION;

    scanf ("%d %d", &AID, &amount); // USER INPUT

    EXEC SQL Select Balance into :BAL From Account
    Where Account Id = :AID; // READ FROM DB

    BAL = BAL + amount; // update BALANCE

    EXEC SQL Update Account
    Set Balance = :b Where Account Id = :AID; // WRITE TO DB
    EXEC SQL Commit Work;
}
    
```

Winter 2013: CMPSC 274 1/9/13

Concurrent Executions: Lost Update Anomaly

DEBIT(\$50)	Time	CREDIT(\$100)
Select Balance Into :b ₁ From Account Where Account Id = :a	1	
		Select Balance Into :b ₂ From Account Where Account Id = :a
	2	
b ₁ = b ₁ - 50	3	
	4	
Update Account Set Balance = :b ₁ Where Account Id = :a	5	Update Account Set Balance = :b ₂ Where Account Id = :a
	6	

Observation: concurrency or parallelism may cause inconsistencies, requires concurrency control for "isolation"

Winter 2013: CMPSC 274 1/9/13

+ Funds Transfer: Inconsistent DATA

```

void main () {
  /* read user input */
  scanf ("%d %d %d", &srcid, &tgtid, &amount);
  /* subtract amount from source account */
  EXEC SQL Update Account
    Set Balance = Balance - :amount Where AccId = :srceid;
  /* add amount to target account */
  EXEC SQL Update Account
    Set Balance = Balance + :amount Where AccId = :tgtid;
  EXEC SQL Commit Work; }

```

CRASH ↓

Observation: failures may cause inconsistencies, require recovery for "atomicity" and "durability"

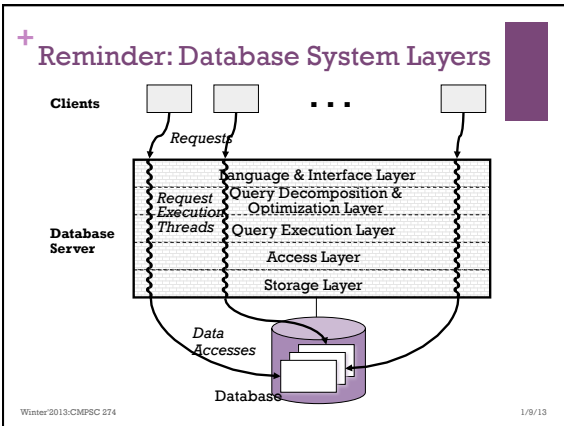
Winter 2013:CMPS 274 1/9/13

+ Database Correctness

Database Fundamentals

Divy Agrawal
Department of Computer Science
UC Santa Barbara

1/9/13 Winter 2013:CMPS 274



+ Basic Ingredients

- Elementary Operations (read and write)
- Transactions (i.e., transaction program executions)
- Execution histories
- Characterization of correct executions
- Protocols (i.e., online algorithms to ensure correctness)

10

Winter 2013: CMPSC 274 1/9/13

+ Transaction Page Model: Syntax

Page Model of Transaction:

A **transaction** T is a partial order of steps (actions) of the form $r[x]$ or $w[x]$, where $x \in D$ and reads and writes as well as multiple writes applied to the same object are ordered.

We write $T = (op, <)$ for transaction T with step set op and partial order $<$.

Example: $r[x] w[x] r[y] w[y]$

Winter 2013: CMPSC 274 1/9/13

+ Transaction Page Model: Semantics

Interpretation of j^{th} step, p_j , of T :

If $p_j = r[x]$, then interpretation is assignment $v_j := x$ to local variable v_j

If $p_j = w[x]$ then interpretation is assignment $x := f_j(v_{j_1}, \dots, v_{j_k})$ with unknown function f_j and j_1, \dots, j_k denoting T 's prior read steps.

Winter 2013: CMPSC 274 1/9/13

+ Lost Update Problem

P1	Time	P2
$r(x)$	1	
$x := x + 100$	2	$r(x)$
$w(x)$	3	$x := x + 200$
	4	
	5	$w(x)$
	6	
	7	

/* x = 100 */
 /* x = 200 */
 /* x = 300 */

↑
update "lost"

Observation: problem is the interleaving $r_1(x) r_2(x) w_1(x) w_2(x)$

Winter 2013: CMPSC 274 1/9/13

+ Dirty Read Problem

P1	Time	P2
$r(x)$	1	
$x := x + 100$	2	
$w(x)$	3	
	4	$r(x)$
	5	$x := x - 100$
failure & rollback	6	
	7	$w(x)$

↑
cannot rely on validity of previously read data

Observation: transaction rollbacks could affect concurrent transactions

Winter 2013: CMPSC 274 1/9/13

+ Correctness Requirements: ACID

15

- **ATOMICITY:**
 - All-or-none property of user programs
- **CONSISTENCY**
 - User program is a consistent unit of execution
- **ISOLATION**
 - User programs are isolated with the side-effects of other user programs
- **DURABILITY:**
 - Effects of user programs are persistent forever

Winter 2013: CMPSC 274 1/9/13

+ Transactions Executions: History

16

- History:
 - Contains all operations from all transactions
 - Distinct termination for every transaction
 - Preserves the order of operations of all transactions
 - Termination is the final step
 - Conflicting operations are ordered

Winter 2013: CMPSC 274 1/9/13

+ Notion of Transaction Histories

17

- Goal:
 - A technique/algorithm/scheduler that prevents incorrect or bad execution.
- Develop the notion of correctness – or characterize what does correct execution means.
- This characterization will be based on the histories of transaction execution:


```

graph LR
  H --> Good
  H --> Bad
        
```

Winter 2013: CMPSC 274 1/9/13

+ Transaction Executions: Histories

Let $T = \{T_1, \dots, T_n\}$ be a set of transactions, where each $T_i \in T$ has the form $T_i = (op_i, <_i)$.

A history for T is $H = (op(H), <_H)$ such that:

- $op(s) \subseteq \cup_{i=1..n} op_i \cup \cup_{i=1..n} \{a_i, c_i\}$
- for all $i, 1 \leq i \leq n$: $c_i \in op(s) \Leftrightarrow a_i \notin op(s)$
- $\cup_{i=1..n} <_i \subseteq <_s$
- for all $i, 1 \leq i \leq n$, and all $p \in op_i$: $p <_H c_i$ or $p <_H a_i$
- for all $p, q \in op(s)$ s.t. at least one of them is a write and both access the same data item: $p <_s q$ or $q <_s p$

Winter 2013: CMPSC 274 1/9/13

+ History Example

Winter 2013:CMPS 274 19

+ Correctness

- Syntactical semantics for schedules based on an intuitive notion:
 - Each transaction is a correct mapping, i.e.,

Hence, serial execution of transactions will be correct.

Winter 2013:CMPS 274 20

+ Serial History

- A history H is serial if for any two transactions T_i and T_j in H, all operations of T_i are ordered in H before all operations of T_j or vice-versa.

Winter 2013:CMPS 274 21

+ General Idea

- Notion of equivalence of two histories H_1 and H_2 .
- Use this notion of equivalence to accept all histories which are “equivalent” to some serial history as being correct.
- How to establish this equivalence notion?

Winter 2013: CMPSC 274 1/9/13

22

+ Semantics

- Equivalence via a notion of semantics:
 - We do not know the semantics of transaction programs
 - We need a general notion that can capture all potential transaction semantics
- ➔ Need a general enough and powerful notion that can capture all possible semantics of transactions.

Winter 2013: CMPSC 274 1/9/13

23
