

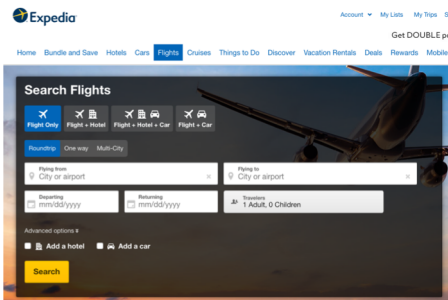
Introduction to Database Systems CSE 414

Lecture 26: More Transactions

Announcements

- Web quiz due tonight
- HW7 due tonight
- HW8 out, make sure to do setup early

HW8



What can go wrong?

- Manager: balance budgets among projects
 - Remove \$10k from project A
 - Add \$7k to project B
 - Add \$3k to project C
- CEO: check company's total balance
 - SELECT SUM(money) FROM budget;
- This is called a dirty / inconsistent read aka a **WRITE-READ** conflict

What can go wrong?

- App 1:
SELECT inventory FROM products WHERE pid = 1
- App 2:
UPDATE products SET inventory = 0 WHERE pid = 1
- App 1:
SELECT inventory * price FROM products WHERE pid = 1
- This is known as an unrepeatable read aka **READ-WRITE** conflict

What can go wrong?

- Account 1 = \$100
Account 2 = \$100
Total = \$200
- App 1:
 - Set Account 1 = \$200
 - Set Account 2 = \$0
 - App 2:
 - Set Account 2 = \$200
 - Set Account 1 = \$0
 - At the end:
 - Total = \$200
 - At the end:
 - Total = \$0

This is called the lost update aka **WRITE-WRITE** conflict

What can go wrong?

- Buying tickets to the next Bieber concert:
 - Fill up form with your mailing address
 - Put in debit card number
 - Click submit
 - Screen shows money deducted from your account
 - [Your browser crashes]

Lesson:
Changes to the database
should be **ALL** or **NOTHING**

CSE 414 - Autumn 2018

7

Transactions

- Collection of statements that are executed atomically (logically speaking)

```
BEGIN TRANSACTION
[SQL statements]
COMMIT      or
ROLLBACK (=ABORT)
```

```
[single SQL statement]
```

If BEGIN... missing,
then TXN consists
of a single instruction

CSE 414 - Autumn 2018

Know your chemistry transactions: ACID

- **Atomic**
 - State shows either all the effects of txn, or none of them
- **Consistent**
 - Txn moves from a DBMS state where integrity holds, to another where integrity holds
 - remember integrity constraints?
- **Isolated**
 - Effect of txns is the same as txns running one after another (i.e., looks like batch mode)
- **Durable**
 - Once a txn has committed, its effects remain in the database

CSE 414 - Autumn 2018

9

Transaction Schedules

CSE 414 - Autumn 2018

10

Schedules

A **schedule** is a sequence
of interleaved actions
from all transactions

CSE 414 - Autumn 2018

11

Serial Schedule

- A *serial schedule* is one in which transactions are executed one after the other, in some sequential order
- **Fact:** nothing can go wrong if the system executes transactions serially
 - (up to what we have learned so far)
 - But DBMS don't do that because we want better overall system performance

CSE 414 - Autumn 2018

12

Example

A and B are elements in the database
t and s are variables in txn source code

T1	T2
READ(A, t)	READ(A, s)
t := t+100	s := s*2
WRITE(A, t)	WRITE(A, s)
READ(B, t)	READ(B, s)
t := t+100	s := s*2
WRITE(B, t)	WRITE(B, s)

Example of a (Serial) Schedule

T1	T2
READ(A, t)	
t := t+100	
WRITE(A, t)	
READ(B, t)	
t := t+100	
WRITE(B, t)	
	READ(A, s)
	s := s*2
	WRITE(A, s)
	READ(B, s)
	s := s*2
	WRITE(B, s)

Time ↓

Another Serial Schedule

T1	T2
	READ(A, s)
	s := s*2
	WRITE(A, s)
	READ(B, s)
	s := s*2
	WRITE(B, s)
READ(A, t)	
t := t+100	
WRITE(A, t)	
READ(B, t)	
t := t+100	
WRITE(B, t)	

Time ↓

Review: Serializable Schedule

A schedule is **serializable** if it is equivalent to a serial schedule

A Serializable Schedule

T1	T2
READ(A, t)	
t := t+100	
WRITE(A, t)	
	READ(A, s)
	s := s*2
	WRITE(A, s)
READ(B, t)	
t := t+100	
WRITE(B, t)	
	READ(B, s)
	s := s*2
	WRITE(B, s)

Time ↓

A Serializable Schedule

T1	T2
READ(A, t)	
t := t+100	
WRITE(A, t)	
	READ(A, s)
	s := s*2
	WRITE(A, s)
READ(B, t)	
t := t+100	
WRITE(B, t)	
	READ(B, s)
	s := s*2
	WRITE(B, s)

Time ↓

A Serializable Schedule

<p>T1</p> <p>READ(A, t)</p> <p>t := t+100</p> <p>WRITE(A, t)</p> <p>READ(B, t)</p> <p>t := t+100</p> <p>WRITE(B, t)</p>	<p>T2</p> <p>READ(A, s)</p> <p>s := s*2</p> <p>WRITE(A, s)</p> <p>READ(B, s)</p> <p>s := s*2</p> <p>WRITE(B, s)</p>	<p>Time</p>
---	---	-------------

CSE 414 - Autumn 2018 19

A Serializable Schedule

<p>T1</p> <p>READ(A, t)</p> <p>t := t+100</p> <p>WRITE(A, t)</p> <p>READ(B, t)</p> <p>t := t+100</p> <p>WRITE(B, t)</p>	<p>T2</p> <p>READ(A, s)</p> <p>s := s*2</p> <p>WRITE(A, s)</p> <p>READ(B, s)</p> <p>s := s*2</p> <p>WRITE(B, s)</p>	<p>Time</p>
---	---	-------------

This is a **serializable** schedule.
This is NOT a serial schedule

CSE 414 - Autumn 2018 20

A Non-Serializable Schedule

<p>T1</p> <p>READ(A, t)</p> <p>t := t+100</p> <p>WRITE(A, t)</p> <p>READ(B, t)</p> <p>t := t+100</p> <p>WRITE(B, t)</p>	<p>T2</p> <p>READ(A, s)</p> <p>s := s*2</p> <p>WRITE(A, s)</p> <p>READ(B, s)</p> <p>s := s*2</p> <p>WRITE(B, s)</p>
---	---

CSE 414 - Autumn 2018 21

A Non-Serializable Schedule

<p>T1</p> <p>READ(A, t)</p> <p>t := t+100</p> <p>WRITE(A, t)</p> <p>READ(B, t)</p> <p>t := t+100</p> <p>WRITE(B, t)</p>	<p>T2</p> <p>READ(A, s)</p> <p>s := s*2</p> <p>WRITE(A, s)</p> <p>READ(B, s)</p> <p>s := s*2</p> <p>WRITE(B, s)</p>
---	---

CSE 414 - Autumn 2018 22

How do We Know if a Schedule is Serializable?

Notation:

T₁: r₁(A); w₁(A); r₁(B); w₁(B)
T₂: r₂(A); w₂(A); r₂(B); w₂(B)

Key Idea: Focus on *conflicting* operations

CSE 414 - Autumn 2018 23

Conflicts

- Write-Read – WR
- Read-Write – RW
- Write-Write – WW
- **Read-Read?**

CSE 414 - Autumn 2018 24

Conflict Serializability

Conflicts: (i.e., swapping will change program behavior)

Two actions by same transaction T_i : $r_i(X); w_i(Y)$

Two writes by T_i, T_j to same element $w_i(X); w_j(X)$

Read/write by T_i, T_j to same element $w_i(X); r_j(X)$
 $r_i(X); w_j(X)$

CSE 414 - Autumn 2018

25

Conflict Serializability

- A schedule is **conflict serializable** if it can be transformed into a serial schedule by a series of swappings of adjacent non-conflicting actions
- Every conflict-serializable schedule is serializable
- The converse is not true (why?)
 - Conflict serializable only looks at conflicts, not values
 - Schedules might have conflicts but would have the same output no matter the order depending on the values

CSE 414 - Autumn 2018

26

Conflict Serializability

Example:

$r_1(A); w_1(A); r_2(A); w_2(A); r_1(B); w_1(B); r_2(B); w_2(B)$

CSE 414 - Autumn 2018

27

Conflict Serializability

Example:

$r_1(A); w_1(A); r_2(A); w_2(A); r_1(B); w_1(B); r_2(B); w_2(B)$



$r_1(A); w_1(A); r_1(B); w_1(B); r_2(A); w_2(A); r_2(B); w_2(B)$

CSE 414 - Autumn 2018

28

Conflict Serializability

Example:

$r_1(A); w_1(A); r_2(A); w_2(A); r_1(B); w_1(B); r_2(B); w_2(B)$



$r_1(A); w_1(A); r_1(B); w_1(B); r_2(A); w_2(A); r_2(B); w_2(B)$

CSE 414 - Autumn 2018

29

Conflict Serializability

Example:

$r_1(A); w_1(A); r_2(A); w_2(A); r_1(B); w_1(B); r_2(B); w_2(B)$



$r_1(A); w_1(A); r_2(A); r_1(B); w_2(A); w_1(B); r_2(B); w_2(B)$




$r_1(A); w_1(A); r_1(B); w_1(B); r_2(A); w_2(A); r_2(B); w_2(B)$

CSE 414 - Autumn 2018

30

Conflict Serializability

Example:



CSE 414 - Autumn 2018

31

Testing for Conflict-Serializability

Precedence graph:

- A node for each transaction T_i ,
 - An edge from T_i to T_j whenever an action in T_i conflicts with, and comes before an action in T_j
- The schedule is conflict-serializable iff the precedence graph is acyclic

CSE 414 - Autumn 2018

32

Example 1

$r_2(A); r_1(B); w_2(A); r_3(A); w_1(B); w_3(A); r_2(B); w_2(B)$




CSE 414 - Autumn 2018

33

Example 1

$r_2(A); r_1(B); w_2(A); r_3(A); w_1(B); w_3(A); r_2(B); w_2(B)$



This schedule is **conflict-serializable**

CSE 414 - Autumn 2018

34

Example 2

$r_2(A); r_1(B); w_2(A); r_2(B); r_3(A); w_1(B); w_3(A); w_2(B)$




CSE 414 - Autumn 2018

35

Example 2

$r_2(A); r_1(B); w_2(A); r_2(B); r_3(A); w_1(B); w_3(A); w_2(B)$



This schedule is **NOT conflict-serializable**

CSE 414 - Autumn 2018

36

Implementing Transactions

CSE 414 - Autumn 2018

37

Scheduler

- **Scheduler** = the module that schedules the transaction's actions, ensuring serializability
- Also called **Concurrency Control Manager**
- We discuss next how a scheduler may be implemented

CSE 414 - Autumn 2018

38

Implementing a Scheduler

Major differences between database vendors

- **Locking Scheduler**
 - Aka "pessimistic concurrency control"
 - SQLite, SQL Server, DB2
- **Multiversion Concurrency Control (MVCC)**
 - Aka "optimistic concurrency control"
 - Postgres, Oracle: Snapshot Isolation (SI)

We discuss only locking schedulers in this class

CSE 414 - Autumn 2018

39

Locking Scheduler

Simple idea:

- Each element has a unique **lock**
- Each transaction must first **acquire** the lock before reading/writing that element
- If the lock is taken by another transaction, then wait
- The transaction must **release** the lock(s)

By using locks scheduler ensures conflict-serializability

What Data Elements are Locked?

Major differences between vendors:

- Lock on the entire database
 - SQLite
- Lock on individual records
 - SQL Server, DB2, etc

CSE 414 - Autumn 2018

41

More Notations

$L_i(A)$ = transaction T_i acquires lock for element A

$U_i(A)$ = transaction T_i releases lock for element A

CSE 414 - Autumn 2018

42

A Non-Serializable Schedule

T1	T2
READ(A)	
A := A+100	
WRITE(A)	
	READ(A)
	A := A*2
	WRITE(A)
	READ(B)
	B := B*2
	WRITE(B)
READ(B)	
B := B+100	
WRITE(B)	

CSE 414 - Autumn 2018

43

Example

T1	T2
$L_1(A)$; READ(A)	
A := A+100	
WRITE(A); $U_1(A)$; $L_1(B)$	
	$L_2(A)$; READ(A)
	A := A*2
	WRITE(A); $U_2(A)$;
	$L_2(B)$; BLOCKED...
READ(B)	
B := B+100	
WRITE(B); $U_1(B)$;	
	...GRANTED ; READ(B)
	B := B*2
	WRITE(B); $U_2(B)$;

Scheduler has ensured a conflict-serializable schedule

44

But what if...

T1	T2
$L_1(A)$; READ(A)	
A := A+100	
WRITE(A); $U_1(A)$;	
	$L_2(A)$; READ(A)
	A := A*2
	WRITE(A); $U_2(A)$;
	$L_2(B)$; READ(B)
	B := B*2
	WRITE(B); $U_2(B)$;
$L_1(B)$; READ(B)	
B := B+100	
WRITE(B); $U_1(B)$;	

Locks did not enforce conflict-serializability !!! What's wrong ?

45

Two Phase Locking (2PL)

The 2PL rule:

In every transaction, all lock requests must precede all unlock requests

CSE 414 - Autumn 2018

46

Example: 2PL transactions

T1	T2
$L_1(A)$; $L_1(B)$; READ(A)	
A := A+100	
WRITE(A); $U_1(A)$	
	$L_2(A)$; READ(A)
	A := A*2
	WRITE(A);
	$L_2(B)$; BLOCKED...
READ(B)	
B := B+100	
WRITE(B); $U_1(B)$;	
	...GRANTED ; READ(B)
	B := B*2
	WRITE(B); $U_2(A)$; $U_2(B)$;

Now it is conflict-serializable

CSE 414 - Autumn 2018

47

Two Phase Locking (2PL)


Theorem: 2PL ensures conflict serializability

48

Two Phase Locking (2PL)

Theorem: 2PL ensures conflict serializability

Proof. Suppose not: then there exists a cycle in the precedence graph.




Two Phase Locking (2PL)

Theorem: 2PL ensures conflict serializability

Proof. Suppose not: then there exists a cycle in the precedence graph.

Then there is the following **temporal** cycle in the schedule:




50

Two Phase Locking (2PL)

Theorem: 2PL ensures conflict serializability

Proof. Suppose not: then there exists a cycle in the precedence graph.

Then there is the following **temporal** cycle in the schedule:
 $U_1(A) \rightarrow L_2(A)$ why?



$U_1(A)$ happened strictly **before** $L_2(A)$


51

Two Phase Locking (2PL)

Theorem: 2PL ensures conflict serializability

Proof. Suppose not: then there exists a cycle in the precedence graph.

Then there is the following **temporal** cycle in the schedule:
 $U_1(A) \rightarrow L_2(A)$
 $L_2(A) \rightarrow U_2(B)$ why?



$L_2(A)$ happened strictly **before** $U_1(A)$


52

Two Phase Locking (2PL)

Theorem: 2PL ensures conflict serializability

Proof. Suppose not: then there exists a cycle in the precedence graph.

Then there is the following **temporal** cycle in the schedule:
 $U_1(A) \rightarrow L_2(A)$
 $L_2(A) \rightarrow U_2(B)$ why?




53

Two Phase Locking (2PL)

Theorem: 2PL ensures conflict serializability

Proof. Suppose not: then there exists a cycle in the precedence graph.

Then there is the following **temporal** cycle in the schedule:
 $U_1(A) \rightarrow L_2(A)$
 $L_2(A) \rightarrow U_2(B)$
 $U_2(B) \rightarrow L_3(B)$ why?




54

Two Phase Locking (2PL)

Theorem: 2PL ensures conflict serializability

Proof. Suppose not: then there exists a cycle in the precedence graph.



Then there is the following **temporal** cycle in the schedule:


$U_1(A) \rightarrow L_2(A)$
 $L_2(A) \rightarrow U_2(B)$
 $U_2(B) \rightarrow L_3(B)$
etc.....

55

Two Phase Locking (2PL)

Theorem: 2PL ensures conflict serializability

Proof. Suppose not: then there exists a cycle in the precedence graph.



Then there is the following **temporal** cycle in the schedule:

$U_1(A) \rightarrow L_2(A)$
 $L_2(A) \rightarrow U_2(B)$
 $U_2(B) \rightarrow L_3(B)$
 $L_3(B) \rightarrow U_3(C)$
 $U_3(C) \rightarrow L_1(C)$
 $L_1(C) \rightarrow U_1(A)$

Cycle in time:
Contradiction

A New Problem: Non-recoverable Schedule

<p>T1</p> <p>$L_1(A); L_1(B); \text{READ}(A)$ $A := A + 100$ $\text{WRITE}(A); U_1(A)$</p> <p>$\text{READ}(B)$ $B := B + 100$ $\text{WRITE}(B); U_1(B)$</p> <p>Rollback</p>	<p>T2</p> <p>$L_2(A); \text{READ}(A)$ $A := A * 2$ $\text{WRITE}(A);$ $L_2(B); \text{BLOCKED}...$</p> <p>$... \text{GRANTED}; \text{READ}(B)$ $B := B * 2$ $\text{WRITE}(B); U_2(A); U_2(B);$ Commit</p>
---	---

CSE 414 - Autumn 2018

57

A New Problem: Non-recoverable Schedule

<p>T1</p> <p>$L_1(A); L_1(B); \text{READ}(A)$ $A := A + 100$ $\text{WRITE}(A); U_1(A)$</p> <p>$\text{READ}(B)$ $B := B + 100$ $\text{WRITE}(B); U_1(B)$</p> <p>Rollback</p>	<p>T2</p> <p>$L_2(A); \text{READ}(A)$ $A := A * 2$ $\text{WRITE}(A);$ $L_2(B); \text{BLOCKED}...$</p> <p>$... \text{GRANTED}; \text{READ}(B)$ $B := B * 2$ $\text{WRITE}(B); U_2(A); U_2(B);$ Commit</p>
---	---

Elements A, B written by T1 are restored to their original value.

Autumn 2018

58

A New Problem: Non-recoverable Schedule

<p>T1</p> <p>$L_1(A); L_1(B); \text{READ}(A)$ $A := A + 100$ $\text{WRITE}(A); U_1(A)$</p> <p>$\text{READ}(B)$ $B := B + 100$ $\text{WRITE}(B); U_1(B)$</p> <p>Rollback</p>	<p>T2</p> <p>$L_2(A); \text{READ}(A)$ $A := A * 2$ $\text{WRITE}(A);$ $L_2(B); \text{BLOCKED}...$</p> <p>Dirty reads of A, B lead to incorrect writes.</p> <p>$... \text{GRANTED}; \text{READ}(B)$ $B := B * 2$ $\text{WRITE}(B); U_2(A); U_2(B);$ Commit</p>
---	--

Elements A, B written by T1 are restored to their original value.

Autumn 2018

59

A New Problem: Non-recoverable Schedule

<p>T1</p> <p>$L_1(A); L_1(B); \text{READ}(A)$ $A := A + 100$ $\text{WRITE}(A); U_1(A)$</p> <p>$\text{READ}(B)$ $B := B + 100$ $\text{WRITE}(B); U_1(B)$</p> <p>Rollback</p>	<p>T2</p> <p>$L_2(A); \text{READ}(A)$ $A := A * 2$ $\text{WRITE}(A);$ $L_2(B); \text{BLOCKED}...$</p> <p>Dirty reads of A, B lead to incorrect writes.</p> <p>$... \text{GRANTED}; \text{READ}(B)$ $B := B * 2$ $\text{WRITE}(B); U_2(A); U_2(B);$ Commit</p>
---	--

Elements A, B written by T1 are restored to their original value.

Autumn 2018

Can no longer undo!

Strict 2PL

The Strict 2PL rule:

All locks are held until commit/abort:
All unlocks are done together with commit/abort.

With strict 2PL, we will get schedules that are both conflict-serializable and recoverable

CSE 414 - Autumn 2018

61

Strict 2PL

T1	T2
<code>L1(A); READ(A)</code>	
<code>A := A+100</code>	
<code>WRITE(A);</code>	
	<code>L2(A); BLOCKED...</code>
<code>L1(B); READ(B)</code>	
<code>B := B+100</code>	
<code>WRITE(B);</code>	
<code>Rollback & U1(A);U1(B);</code>	
	<code>...GRANTED; READ(A)</code>
	<code>A := A*2</code>
	<code>WRITE(A);</code>
	<code>L2(B); READ(B)</code>
	<code>B := B*2</code>
	<code>WRITE(B);</code>
	<code>Commit & U2(A); U2(B);</code>

62

Strict 2PL

- Lock-based systems always use strict 2PL
- Easy to implement:
 - Before a transaction reads or writes an element A, insert an L(A)
 - When the transaction commits/aborts, then release all locks
- Ensures both conflict serializability and recoverability

CSE 414 - Autumn 2018

63