

# CSE 444: Database Internals

## Lectures 13 Transaction Schedules

CSE 444 - Winter 2018

1

## About Lab 3

- In lab 3, we implement transactions
- Focus on concurrency control
  - Want to run many transactions at the same time
  - Transactions want to read and write same pages
  - Will use locks to ensure conflict serializable execution
  - Use strict 2PL
- Build your own lock manager
  - Understand how locking works in depth
  - Ensure transactions rather than threads hold locks
    - Many threads can execute different pieces of the same transaction
    - Need to detect deadlocks and resolve them by aborting a transaction
  - But use Java synchronization to protect your data structures

CSE 444 - Winter 2018

2

## Motivating Example

Client 1:  
`UPDATE Budget`  
`SET money=money-100`  
`WHERE pid = 1`  
  
`UPDATE Budget`  
`SET money=money+60`  
`WHERE pid = 2`  
  
`UPDATE Budget`  
`SET money=money+40`  
`WHERE pid = 3`

Client 2:  
`SELECT sum(money)`  
`FROM Budget`

Would like to treat  
each group of  
instructions as a unit

CSE 444 - Winter 2018

3

## Transaction

**Definition:** a transaction is a sequence of updates to the database with the property that either all complete, or none completes (all-or-nothing).

`START TRANSACTION`

[SQL statements]

`COMMIT` or `ROLLBACK (=ABORT)`

May be omitted if  
autocommit is off:  
first SQL query  
starts txn

In ad-hoc SQL: each statement = one transaction  
This is referred to as autocommit

CSE 444 - Winter 2018

4

## Motivating Example

`START TRANSACTION`  
`UPDATE Budget`  
`SET money=money-100`  
`WHERE pid = 1`  
  
`UPDATE Budget`  
`SET money=money+60`  
`WHERE pid = 2`  
  
`UPDATE Budget`  
`SET money=money+40`  
`WHERE pid = 3`  
`COMMIT (or ROLLBACK)`

`SELECT sum(money)`  
`FROM Budget`

With autocommit and  
without `START TRANSACTION`,  
each SQL command  
is a transaction

CSE 444 - Winter 2018

5

## ROLLBACK

- If the app gets to a place where it can't complete the transaction successfully, it can execute **ROLLBACK**
- This causes the system to "abort" the transaction
  - Database returns to a state without any of the changes made by the transaction
- Several reasons: user, application, system

CSE 444 - Winter 2018

6

## Transactions

- Major component of database systems
- Critical for most applications; arguably more so than SQL
- Turing awards to database researchers:
  - Charles Bachman 1973
  - Edgar Codd 1981 for inventing relational dbs
  - Jim Gray 1998 for inventing transactions
  - Mike Stonebraker 2015 for INGRES and Postgres
    - And many other ideas after that

CSE 444 - Winter 2018

7

## ACID Properties

- **Atomicity**: Either all changes performed by transaction occur or none occurs
- **Consistency**: A transaction as a whole does not violate integrity constraints
- **Isolation**: Transactions appear to execute one after the other in sequence
- **Durability**: If a transaction commits, its changes will survive failures

CSE 444 - Winter 2018

8

## What Could Go Wrong?

Why is it hard to provide ACID properties?

- **Concurrent** operations
  - Isolation problems
  - We saw one example earlier
- **Failures** can occur at any time
  - Atomicity and durability problems
  - Later lectures
- Transaction may need to **abort**

CSE 444 - Winter 2018

9

## Terminology Needed For Lab 3 Buffer Manager Policies

- **STEAL or NO-STEAL**
  - Can an update made by an uncommitted transaction overwrite the most recent committed value of a data item on disk?
- **FORCE or NO-FORCE**
  - Should all updates of a transaction be forced to disk before the transaction commits?
- Easiest for recovery: NO-STEAL/FORCE (lab 3)
- Highest performance: STEAL/NO-FORCE (lab 4)
- We will get back to this next week

CSE 444 - Winter 2018

10

## Transaction Isolation

CSE 444 - Winter 2018

11

## Concurrent Execution Problems

- **Write-read conflict: dirty read, inconsistent read**
  - A transaction reads a value written by another transaction that has not yet committed
- **Read-write conflict: unrepeatable read**
  - A transaction reads the value of the same object twice. Another transaction modifies that value in between the two reads
- **Write-write conflict: lost update**
  - Two transactions update the value of the same object. The second one to write the value overwrites the first change

CSE 444 - Winter 2018

12

## Schedules

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

CSE 444 - Winter 2018

13

## Example

A and B are elements in the database  
t and s are variables in tx 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)

CSE 444 - Winter 2018

14

## 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)

CSE 444 - Winter 2018

15

## Serializable Schedule

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

CSE 444 - Winter 2018

16

## 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)

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

CSE 444 - Winter 2018

17

## A Non-Serializable Schedule

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

CSE 444 - Winter 2018

18

## Serializable Schedules

- The role of the scheduler is to ensure that the schedule is serializable

**Q:** Why not run only serial schedules ?  
I.e. run one transaction after the other ?

CSE 444 - Winter 2018

19

## Serializable Schedules

- The role of the scheduler is to ensure that the schedule is serializable

**Q:** Why not run only serial schedules ?  
I.e. run one transaction after the other ?

**A:** Because of very poor throughput due to disk latency.

**Lesson:** main memory databases *may* schedule TXNs serially

CSE 444 - Winter 2018

20

## Still Serializable, but...

<p>T1</p> <p>READ(A, t)</p> <p>t := t+100</p> <p>WRITE(A, t)</p>	<p>T2</p> <p>READ(A, s)</p> <p>s := s + 200</p> <p>WRITE(A, s)</p> <p>READ(B, s)</p> <p>s := s + 200</p> <p>WRITE(B, s)</p>
--	---

READ(B, t)

t := t+100

WRITE(B, t)

Schedule is serializable  
because t=t+100 and  
s=s+200 commute

...we don't expect the scheduler to schedule this

## To Be Practical

- Assume worst case updates:
  - Assume cannot commute actions done by transactions
- Therefore, we only care about reads and writes
  - Transaction = sequence of R(A)'s and W(A)'s

T<sub>1</sub>: r<sub>1</sub>(A); w<sub>1</sub>(A); r<sub>1</sub>(B); w<sub>1</sub>(B)  
T<sub>2</sub>: r<sub>2</sub>(A); w<sub>2</sub>(A); r<sub>2</sub>(B); w<sub>2</sub>(B)

CSE 444 - Winter 2018

22

## Conflicts

- Write-Read – WR
- Read-Write – RW
- Write-Write – WW

CSE 444 - Winter 2018

23

## Conflict Serializability

### Conflicts:

Two actions by same transaction T<sub>i</sub>:

r<sub>i</sub>(X); w<sub>i</sub>(Y)

Two writes by T<sub>i</sub>, T<sub>j</sub> to same element

w<sub>i</sub>(X); w<sub>j</sub>(X)

Read/write by T<sub>i</sub>, T<sub>j</sub> to same element

w<sub>i</sub>(X); r<sub>j</sub>(X)

r<sub>i</sub>(X); w<sub>j</sub>(X)

CSE 444 - Winter 2018

24

## Conflict Serializability

**Definition** 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 in general

CSE 444 - Winter 2018

25

## 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 444 - Winter 2018

26

## Conflict Serializability

Example:

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



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

CSE 444 - Winter 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 444 - Winter 2018

28

## Conflict Serializability

Example:

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



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



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

CSE 444 - Winter 2018

29

## Conflict Serializability

Example:

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



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



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



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

CSE 444 - Winter 2018

30

## 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 serializable iff the precedence graph is acyclic

CSE 444 - Winter 2018

31

## 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 444 - Winter 2018

32

## Example 1

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



This schedule is **conflict-serializable**

CSE 444 - Winter 2018

33

## 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 444 - Winter 2018

34

## Example 2

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



This schedule is **NOT conflict-serializable**

CSE 444 - Winter 2018

35

## View Equivalence

- A serializable schedule need not be conflict serializable, even under the “worst case update” assumption

$w_1(X); w_2(X); w_2(Y); w_1(Y); w_3(Y);$

Is this schedule conflict-serializable ?

CSE 444 - Winter 2018

36

## View Equivalence

- A serializable schedule need not be conflict serializable, even under the “worst case update” assumption

$w_1(X); w_2(X); w_2(Y); w_1(Y); w_3(Y);$

Is this schedule conflict-serializable ?

No...

CSE 444 - Winter 2018

37

## View Equivalence

- A serializable schedule need not be conflict serializable, even under the “worst case update” assumption

$w_1(X); w_2(X); w_2(Y); w_1(Y); w_3(Y);$


Lost write

$w_1(X); w_1(Y); w_2(X); w_2(Y); w_3(Y);$

Equivalent, but not conflict-equivalent

38

## View Equivalence



Serializable, but not conflict serializable

39

## View Equivalence

Two schedules  $S, S'$  are **view equivalent** if:

- If  $T$  reads an **initial value** of  $A$  in  $S$ , then  $T$  reads the **initial value** of  $A$  in  $S'$
- If  $T$  reads a value of  $A$  **written by  $T'$**  in  $S$ , then  $T$  reads a value of  $A$  **written by  $T'$**  in  $S'$
- If  $T$  writes the **final value** of  $A$  in  $S$ , then  $T$  writes the **final value** of  $A$  in  $S'$

CSE 444 - Winter 2018

40

## View-Serializability

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

Remark:

- If a schedule is **conflict serializable**, then it is also **view serializable**
- But not vice versa

CSE 444 - Winter 2018

41


## Schedules with Aborted Transactions

- When a transaction aborts, the recovery manager undoes its updates
- But some of its updates may have affected other transactions !

CSE 444 - Winter 2018

42


## Schedules with Aborted Transactions



CSE 444 - Winter 2018

43

## Schedules with Aborted Transactions



Cannot abort T1 because cannot undo T2

## Recoverable Schedules


A schedule is *recoverable* if:

- It is conflict-serializable, and
- Whenever a transaction T commits, all transactions who have written elements read by T have already committed

CSE 444 - Winter 2018

45


## Recoverable Schedules



CSE 444 - Winter 2018

46

## Recoverable Schedules



47

## Cascading Aborts

- If a transaction T aborts, then we need to abort any other transaction T' that has read an element written by T
- A schedule *avoids cascading aborts* if whenever a transaction reads an element, the transaction that has last written it has already committed.

CSE 444 - Winter 2018

48



## Avoiding Cascading Aborts

T1	T2	T1	T2
R(A)		R(A)	
W(A)		W(A)	
	R(A)	Commit	
	W(A)		R(A)
	R(B)		W(A)
	W(B)		R(B)
...			W(B)
	...		...

With cascading aborts

Without cascading aborts

CSE 444 - Winter 2018

49

## Review of Schedules

### Serializability

- Serial
- Serializable
- Conflict serializable
- View serializable

### Recoverability

- Recoverable
- Avoids cascading deletes

CSE 444 - Winter 2018

50

## Scheduler

- The scheduler:
- Module that schedules the transaction's actions, ensuring serializability
- Two main approaches
  - **Pessimistic:** locks
  - **Optimistic:** timestamps, multi-version, validation

CSE 444 - Winter 2018

51