

# CSE 444: Database Internals

## Lectures 13 Transaction Schedules

## 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

## 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

## 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

## 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

## 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

## 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

## 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**

## 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

## Transaction Isolation

## 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

## Schedules

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

## 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 2017

13

## 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 2017

14

## Serializable Schedule

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

CSE 444 - Winter 2017

15

## 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 2017

16

## 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 2017

17

## 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 2017

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 ?

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

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

CSE 444 - Winter 2017

19

## Still Serializable, but...

T1  
 READ(A, t)  
 t := t+100  
 WRITE(A, t)

T2  
 READ(A,s)  
 s := s + 200  
 WRITE(A,s)  
 READ(B,s)  
 s := s + 200  
 WRITE(B,s)

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

READ(B, t)  
 t := t+100  
 WRITE(B,t)

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

## Ignoring Details

- Assume worst case updates:
  - We never 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 2017

21

## Conflicts

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

CSE 444 - Winter 2017

22

## 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)

Read/write by T<sub>i</sub>, T<sub>j</sub> to same element r<sub>i</sub>(X); w<sub>j</sub>(X)

CSE 444 - Winter 2017

23

## 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 2017

24

## Conflict Serializability

Example:

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

## Conflict Serializability

Example:

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



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

## Conflict Serializability

Example:

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



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

## Conflict Serializability

Example:

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



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



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

## Conflict Serializability

Example:

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



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



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



....

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

## 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

### Example 1

---

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

①
②
③

CSE 444 - Winter 2017 31

### Example 1

---

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

This schedule is **conflict-serializable**

CSE 444 - Winter 2017 32

### Example 2

---

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

①
②
③

CSE 444 - Winter 2017 33

### Example 2

---

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

This schedule is **NOT conflict-serializable**

CSE 444 - Winter 2017 34

### 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 2017 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 ?

No...

CSE 444 - Winter 2017 36

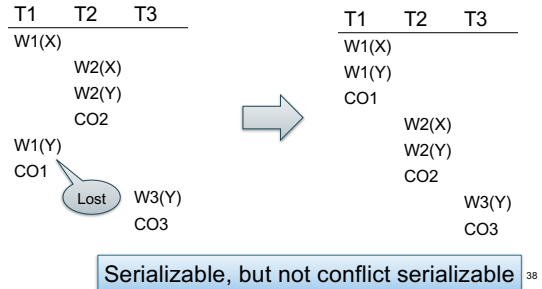
## View Equivalence

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



37

## View Equivalence



38

## 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 2017

39

## 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 2017

40


## 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 2017

41


## Schedules with Aborted Transactions



CSE 444 - Winter 2017

42

## 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

## Recoverable Schedules




Nonrecoverable



Recoverable

## Recoverable Schedules




How do we recover ?

## 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.

## Avoiding Cascading Aborts



With cascading aborts

Without cascading aborts



## Review of Schedules

### Serializability

- Serial
- Serializable
- Conflict serializable
- View serializable

### Recoverability

- Recoverable
- Avoids cascading deletes

CSE 444 - Winter 2017

49

## 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 2017

50