

CSE 344: Section 11

Transactions

December 7th, 2017



Administrivia

Deadline of **HW 8** pushed back to **Next Wednesday**

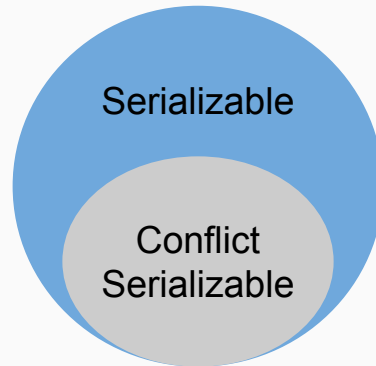
Hard Deadline (No late days allowed)

Long and difficult

Serializability



Conflict serializable is stricter than serializable

I.e. Any schedule that is conflict serializable must be serializable.



Serializability

Checking for conflict serializability -> precedence graph and cycle checking



Serializability

S1: w1(Y); w2(Y); w1(X); w2(X); w3(X)


Are these serializable?
Conflict serializable?

S2: w1(Y); w2(Y); w2(X); w1(X); w3(X)

Serializability


S1: $w1(Y)$; $w2(Y)$; $w1(X)$; $w2(X)$; $w3(X)$

Conflict Serializable



S2: $w1(Y)$; $w2(Y)$; $w2(X)$; $w1(X)$; $w3(X)$

Serializable (but not conflict serializable)



2PL v.s. Strict 2PL

2PL:

- In every transaction, all lock requests must precede all unlock requests
- Ensure Conflict Serializability
- Might not be able to recover (Dirty Read: Read on some write that gets rolled back)

Strict 2PL:

- Every lock each transaction is held until commit or abort
- Ensure Conflict Serializability
- Recoverable as each transaction does not affect others until commit/abort

2PL v.s. Strict 2PL

A New Problem: Non-recoverable Schedule

T1	T2
<u>L₁(A); L₁(B);</u> 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); Commit

Rollback

Isolation Level: Read Uncommitted

Write Locks? Strict 2PL

Read Locks? No (Immediate Read)

Problem: Dirty-Read

Reading uncommitted data that can be rolled back

Isolation Level: Read Uncommitted

Example Transaction:

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

T2 is reading value of A updated by T1's write on A, but T1 has not committed yet.

The value of A read by T2 might not even be in the result.

Then T2's action can be influenced by such uncommitted data.

Isolation Level: Read Committed

Write Locks? Strict 2PL

Read Locks? Obtain before read, release after (No more dirty read)

Problem: Unrepeatable Read

The values of 2 reads on the same tuple can be different in the same transaction

Isolation Level: Read Committed

Example Transaction:

$$S = \begin{array}{cc} & \begin{array}{c} T1 \\ T2 \end{array} \\ \begin{array}{c} R(A) \\ \\ R(A) \\ W(A) \\ Com. \end{array} & \begin{array}{c} \\ R(A) \\ W(A) \\ Com. \end{array} \end{array}$$

T1's first R(A) and T1's second R(A) might have different results.

Updated by T2's W(A).

Isolation Level: Repeatable Read

Write Locks? Strict 2PL

Read Locks? Strict 2PL (No more unrepeatable read)

Same as Serializable if no insert or delete

Problem: Phantom Read

In the same transaction, some tuples appear sometimes and disappear other times

Isolation Level: Repeatable Read

Suppose there are two blue products, A1, A2:

Phantom Problem

T1

T2

```
SELECT *  
FROM Product  
WHERE color='blue'
```

```
INSERT INTO Product(name, color)  
VALUES ('A3','blue')
```

```
SELECT *  
FROM Product  
WHERE color='blue'
```

Isolation Level: Serializable

Not the same thing as Serializable schedule!!!

Write Locks: Strict 2PL

Read Locks: Strict 2PL

Predicate Lock/Table Lock (No Phantom)

Isolation Level: Serializable

Predicate Lock Example:

In Transaction T, we have a statement:

```
SELECT * FROM People WHERE age > 18;
```

In this case, the transaction will grab a predicate lock that prevent inserting and deleting tuples that can affect the predicate/statement.

In this case, the lock prevents inserting and deleting tuples with age > 18.

Isolation Level: Summary

Isolation Level	Dirty Reads	Nonrepeatable Reads	Phantoms
Read Uncommitted	Allowed	Allowed	Allowed
Read Committed	Not Allowed	Allowed	Allowed
Repeatable Read	Not Allowed	Not Allowed	Allowed
Serializable	Not Allowed	Not Allowed	Not Allowed