

Database System Internals
Concurrency Control - Locking

Paul G. Allen School of Computer Science and Engineering
University of Washington, Seattle

February 10, 2020 CSE 444 - Winter 2020 1

1

Announcements

- Lab 2 due tonight
 - Before final submission, clone fresh repo on attu and run "ant test-report"
- Lab 1+2 quiz on Wednesday in-class
 - Closed book. Calculator allowed but you won't need one.
- 544M Paper 2 due next week

February 10, 2020 CSE 444 - Winter 2020 2

2

Conflicts

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

February 10, 2020 CSE 444 - Winter 2020 3

3

Conflict Serializability

Conflicts:

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 $w_i(X); r_j(X)$
 $r_i(X); w_j(X)$

February 10, 2020 CSE 444 - Winter 2020 4

4

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

February 10, 2020 CSE 444 - Winter 2020 5

5

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
- No edge for actions in the same transaction

▪ The schedule is **serializable** iff the precedence graph is acyclic

February 10, 2020 CSE 444 - Winter 2020 6

6

Testing for Conflict-Serializability

Important:

Always draw the full graph, unless ONLY asked if (yes or no) the schedule is conflict serializable

February 10, 2020 CSE 444 - Winter 2020 7

7

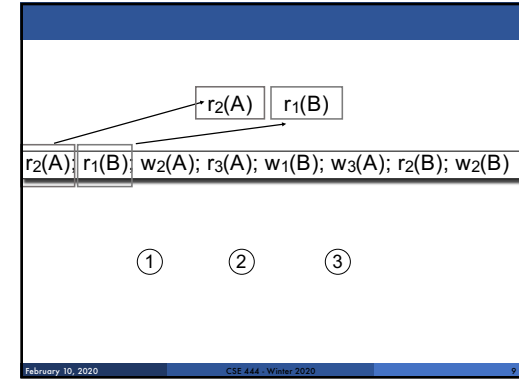
Example 1

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

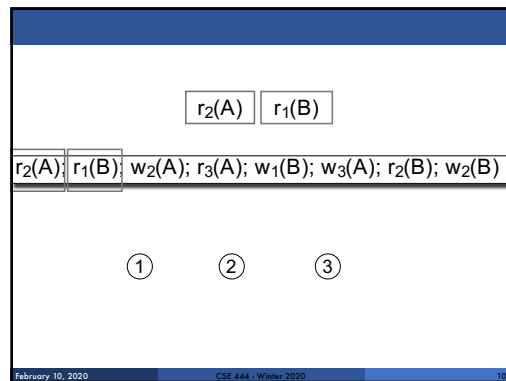
① ② ③

February 10, 2020 CSE 444 - Winter 2020 8

8

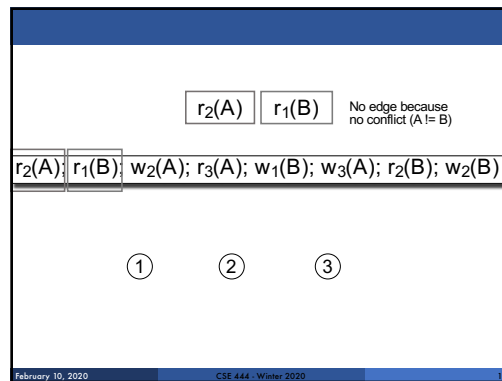


9



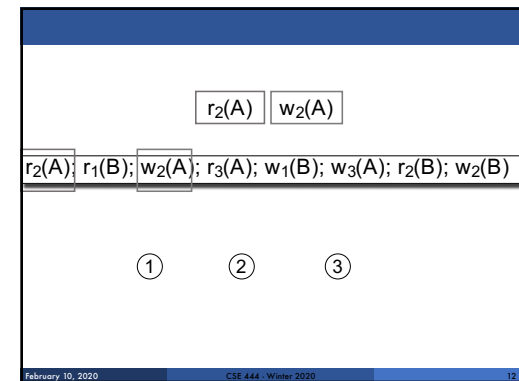
February 10, 2020 CSE 444 - Winter 2020 10

10



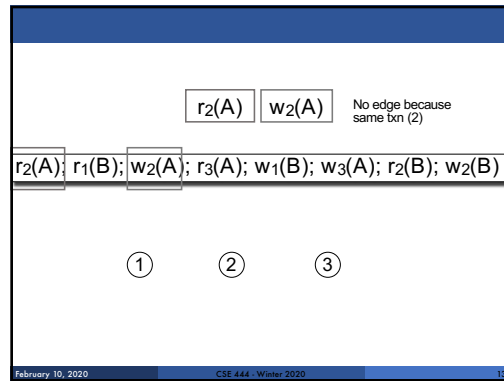
February 10, 2020 CSE 444 - Winter 2020 11

11

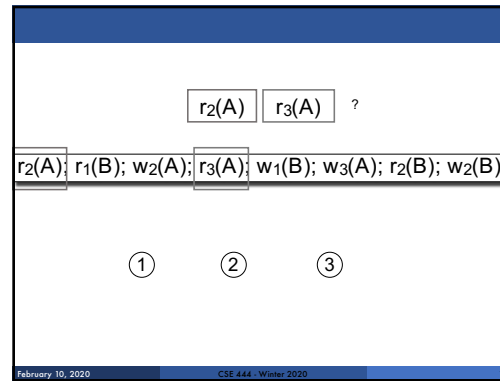


February 10, 2020 CSE 444 - Winter 2020 12

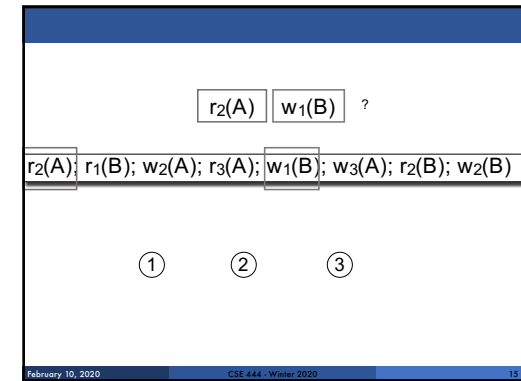
12



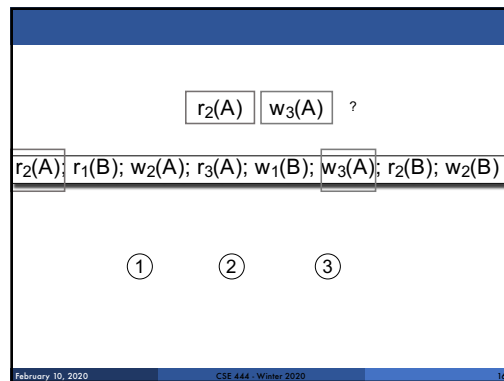
13



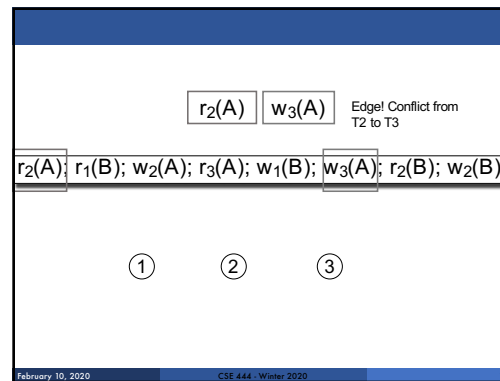
14



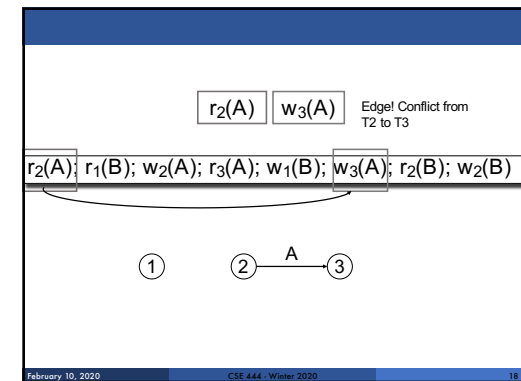
15



16



17



18

Diagram illustrating a sequence of actions: $r_2(A)$, $r_2(B)$, and a question mark. Below this, a sequence of actions is shown: $r_2(A)$, $r_1(B)$, $w_2(A)$, $r_3(A)$, $w_1(B)$, $w_3(A)$, $r_2(B)$, and $w_2(B)$. A text box states: "And so on until compared every pair of actions..." Below the text box, three nodes labeled 1, 2, and 3 are shown, with a directed edge from 2 to 3.

19

Example 1

Diagram illustrating a sequence of actions: $r_2(A)$, $r_1(B)$, $w_2(A)$, $r_3(A)$, $w_1(B)$, $w_3(A)$, $r_2(B)$, and $w_2(B)$. Below this, three nodes labeled 1, 2, and 3 are shown, with a directed edge from 1 to 2 labeled B, and a directed edge from 2 to 3 labeled A. A text box states: "More edges, but repeats of the same directed edge not necessary".

20

Example 1

Diagram illustrating a sequence of actions: $r_2(A)$, $r_1(B)$, $w_2(A)$, $r_3(A)$, $w_1(B)$, $w_3(A)$, $r_2(B)$, and $w_2(B)$. Below this, three nodes labeled 1, 2, and 3 are shown, with a directed edge from 1 to 2 labeled B, and a directed edge from 2 to 3 labeled A. A text box states: "This schedule is **conflict-serializable**".

21

Example 2

Diagram illustrating a sequence of actions: $r_2(A)$, $r_1(B)$, $w_2(A)$, $r_2(B)$, $r_3(A)$, $w_1(B)$, $w_3(A)$, and $w_2(B)$. Below this, three nodes labeled 1, 2, and 3 are shown, with no directed edges between them.

22

Example 2

Diagram illustrating a sequence of actions: $r_2(A)$, $r_1(B)$, $w_2(A)$, $r_2(B)$, $r_3(A)$, $w_1(B)$, $w_3(A)$, and $w_2(B)$. Below this, three nodes labeled 1, 2, and 3 are shown, with a directed edge from 1 to 2 labeled B, and a directed edge from 2 to 3 labeled A. A text box states: "This schedule is **NOT conflict-serializable**".

23

Example 2

Diagram illustrating a sequence of actions: $r_2(A)$, $r_1(B)$, $w_2(A)$, $r_2(B)$, $r_3(A)$, $w_1(B)$, $w_3(A)$, and $w_2(B)$. Below this, three nodes labeled 1, 2, and 3 are shown, with a directed edge from 1 to 2 labeled B, and a directed edge from 2 to 3 labeled A. A text box states: "This schedule is **NOT conflict-serializable**".

24

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 ?

February 10, 2020 CSE 444, Winter 2020 25

25

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...

February 10, 2020 CSE 444, Winter 2020 26

26

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

February 10, 2020 CSE 444, Winter 2020 27

27

View Equivalence



Serializable, but not conflict serializable

February 10, 2020 CSE 444, Winter 2020 28

28

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'

February 10, 2020 CSE 444, Winter 2020 29

29

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

February 10, 2020 CSE 444, Winter 2020 30

30

Schedules with Aborted Transactions

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

February 10, 2020 CSE 444 - Winter 2020 31

31

Schedules with Aborted Transactions

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

What's wrong?

February 10, 2020 CSE 444 - Winter 2020 32

32

Schedules with Aborted Transactions

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

What's wrong?

February 10, 2020 CSE 444 - Winter 2020 33

33

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 that have written elements read by T have already committed

February 10, 2020 CSE 444 - Winter 2020 34

34

Recoverable Schedules

A schedule is *recoverable* if:

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

February 10, 2020 CSE 444 - Winter 2020 35

35

Recoverable Schedules

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

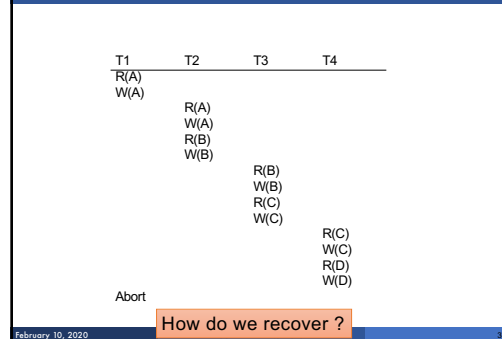
Nonrecoverable

Recoverable

February 10, 2020 CSE 444 - Winter 2020 36

36

Recoverable Schedules



37

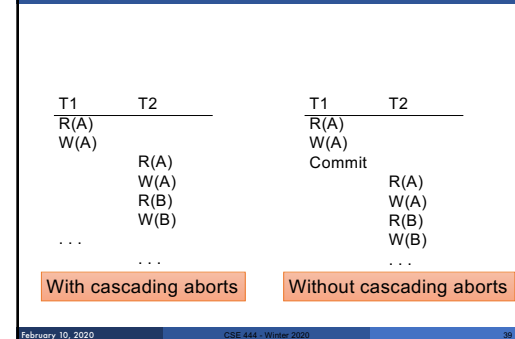
Cascading Aborts

- 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*.

We base our locking scheme on this rule!

38

Avoiding Cascading Aborts



39

Serializability

- Serial
- Serializable
- Conflict serializable
- View serializable

Recoverability

- Recoverable
- Avoids cascading deletes

40

Scheduler

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

41

Pessimistic 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)

42

Notation

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

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

February 10, 2020 CSE 444 - Winter 2020 43

43

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)	

February 10, 2020 CSE 444 - Winter 2020 44

44

Example

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

Scheduler has ensured a conflict-serializable schedule 15

February 10, 2020 CSE 444 - Winter 2020 45

45

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

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

February 10, 2020 CSE 444 - Winter 2020 46

46

Two Phase Locking (2PL)

The 2PL rule:

- In every transaction, all lock requests must precede all unlock requests
- This ensures conflict serializability ! (will prove this shortly)

February 10, 2020 CSE 444 - Winter 2020 47

47

Example: 2PL transactions

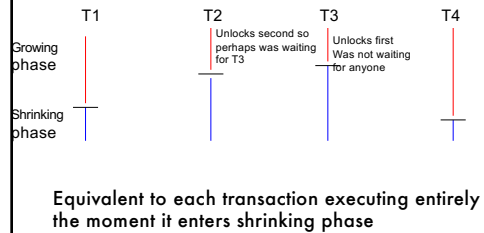
T1	T2
$L_1(A)$; $L_1(B)$; READ(A, t)	
t := t+100	
WRITE(A, t); $U_1(A)$	
	$L_2(A)$; READ(A, s)
	s := s*2
	WRITE(A, s);
	$L_2(B)$; DENIED...
READ(B, t)	
t := t+100	
WRITE(B, t); $U_1(B)$;	
	...GRANTED ; READ(B, s)
	s := s*2
	WRITE(B, s); $U_2(A)$; $U_2(B)$;

Now it is conflict-serializable

February 10, 2020 CSE 444 - Winter 2020 48

48

Example with Multiple Transactions



49

Two Phase Locking (2PL)


Theorem: 2PL ensures conflict serializability

50

Two Phase Locking (2PL)

Theorem: 2PL ensures conflict serializability

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




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:




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:
 $U_1(A) \rightarrow L_2(A)$ 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:
 $U_1(A) \rightarrow L_2(A)$
 $L_2(A) \rightarrow U_2(B)$ 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:

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

Contradiction

55

A New Problem:

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

56

Strict 2PL

- Strict 2PL: All locks held by a transaction are released when the transaction is completed; release happens at the time of COMMIT or ROLLBACK
- Schedule is **recoverable**
- Schedule **avoids cascading aborts**

57

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

58

Summary of Strict 2PL

- Ensures serializability, recoverability, and avoids cascading aborts
- Issues?

59

Summary of Strict 2PL

- Ensures serializability, recoverability, and avoids cascading aborts
- Issues: implementation, lock modes, granularity, deadlocks, performance

60

The Locking Scheduler

Task 1: -- act on behalf of the transaction

- Add lock/unlock requests to transactions
- Examine all READ(A) or WRITE(A) actions
- Add appropriate lock requests
- On COMMIT/ROLLBACK release all locks
- Ensures Strict 2PL !

February 10, 2020 CSE 444 - Winter 2020 61

61

The Locking Scheduler

Task 2: -- act on behalf of the system
Execute the locks accordingly

- Lock table: a big, critical data structure in a DBMS !
- When a lock is requested, check the lock table
 - Grant, or add the transaction to the element's wait list
- When a lock is released, re-activate a transaction from its wait list
- When a transaction aborts, release all its locks
- Check for deadlocks occasionally

February 10, 2020 CSE 444 - Winter 2020 62

62

Lock Modes

- S = shared lock (for READ)
- X = exclusive lock (for WRITE)

Lock compatibility matrix:

	None	S	X
None	OK	OK	OK
S	OK	OK	Conflict
X	OK	Conflict	Conflict

February 10, 2020 CSE 444 - Winter 2020 63

63

Lock Granularity

- Fine granularity locking (e.g., tuples)
 -
 -
- Coarse grain locking (e.g., tables, predicate locks)
 -
 -

February 10, 2020 CSE 444 - Winter 2020 64

64

Lock Granularity

- Fine granularity locking (e.g., tuples)
 - High concurrency
 - High overhead in managing locks
- Coarse grain locking (e.g., tables, predicate locks)
 -
 -

February 10, 2020 CSE 444 - Winter 2020 65

65

Lock Granularity

- Fine granularity locking (e.g., tuples)
 - High concurrency
 - High overhead in managing locks
- Coarse grain locking (e.g., tables, predicate locks)
 - Many false conflicts
 - Less overhead in managing locks

February 10, 2020 CSE 444 - Winter 2020 66

66


Deadlocks

- **Cycle in the wait-for graph:**
 - T1 waits for T2
 - T2 waits for T3
 - T3 waits for T1
- **Deadlock detection**
 - Timeouts
 - Wait-for graph
- **Deadlock avoidance**
 - Acquire locks in pre-defined order
 - Acquire all locks at once before starting

February 10, 2020 CSE 444 - Winter 2020 71

71

Lock Performance



February 10, 2020 CSE 444 - Winter 2020 72

72

Phantom Problem

- So far we have assumed the database to be a *static* collection of elements (=tuples)
- If tuples are inserted/deleted then the *phantom problem* appears

February 10, 2020 CSE 444 - Winter 2020 75

75

Phantom Problem

T1	T2
SELECT *	
FROM Product	
WHERE color='blue'	
	INSERT INTO Product(name, color)
	VALUES ('gizmo','blue')
SELECT *	
FROM Product	
WHERE color='blue'	

Is this schedule serializable ?

February 10, 2020 CSE 444 - Winter 2020 76

76

Phantom Problem

T1	T2
SELECT *	
FROM Product	
WHERE color='blue'	
	INSERT INTO Product(name, color)
	VALUES ('gizmo','blue')
SELECT *	
FROM Product	
WHERE color='blue'	

Suppose there are two blue products, X1, X2:

R1(X1),R1(X2),W2(X3),R1(X1),R1(X2),R1(X3)

February 10, 2020 CSE 444 - Winter 2020 77

77

Phantom Problem

T1	T2
SELECT *	
FROM Product	
WHERE color='blue'	
	INSERT INTO Product(name, color)
	VALUES ('gizmo','blue')
SELECT *	
FROM Product	
WHERE color='blue'	

Suppose there are two blue products, X1, X2:

R1(X1),R1(X2),W2(X3),R1(X1),R1(X2),R1(X3)

This is conflict serializable ! What's wrong ??

February 10, 2020 CSE 444 - Winter 2020 78

78

Phantom Problem

T1	T2
SELECT *	
FROM Product	
WHERE color='blue'	
	INSERT INTO Product(name, color)
	VALUES ('gizmo','blue')
SELECT *	
FROM Product	
WHERE color='blue'	

Suppose there are two blue products, X1, X2:

R1(X1),R1(X2),W2(X3),R1(X1),R1(X2),R1(X3)

Not serializable due to **phantoms**

February 10, 2020 CSE 444 - Winter 2020 79

79

Phantom Problem

- A "phantom" is a tuple that is invisible during **part** of a transaction execution but not invisible during the **entire** execution
- In our example:
 - T1: reads list of products
 - T2: inserts a new product
 - T1: re-reads: a new product appears !

February 10, 2020 CSE 444 - Winter 2020 80

80

Phantom Problem

- In a **static** database:
 - Conflict serializability implies serializability
- In a **dynamic** database, this may fail due to phantoms
- Strict 2PL guarantees conflict serializability, but not serializability

February 10, 2020 CSE 444 - Winter 2020 81

81

Dealing With Phantoms

- Lock the entire table, or
- Lock the index entry for 'blue'
 - If index is available
- Or use predicate locks
 - A lock on an arbitrary predicate

Dealing with phantoms is expensive !

February 10, 2020 CSE 444 - Winter 2020 82

82

Isolation Levels in SQL

- "Dirty reads"
 - SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED
- "Committed reads"
 - SET TRANSACTION ISOLATION LEVEL READ COMMITTED
- "Repeatable reads"
 - SET TRANSACTION ISOLATION LEVEL REPEATABLE READ
- Serializable transactions
 - SET TRANSACTION ISOLATION LEVEL SERIALIZABLE

ACID

February 10, 2020 CSE 444 - Winter 2020 83

83

1. Isolation Level: Dirty Reads

- "Long duration" WRITE locks
 - Strict 2PL
- No READ locks
 - Read-only transactions are never delayed

Possible pbs: dirty and inconsistent reads

February 10, 2020 CSE 444 - Winter 2020 84

84

- "Long duration" WRITE locks
 - Strict 2PL
- "Short duration" READ locks
 - Only acquire lock while reading (not 2PL)

Unrepeatable reads
When reading same element twice,
may get two different values

85

- "Long duration" WRITE locks
 - Strict 2PL
- "Long duration" READ locks
 - Strict 2PL

This is not serializable yet !!!

Why ?

86

4. Isolation Level Serializable

- "Long duration" WRITE locks
 - Strict 2PL
- "Long duration" READ locks
 - Strict 2PL
- Predicate locking
 - To deal with phantoms

87

READ-ONLY Transactions

```

Client 1: START TRANSACTION
INSERT INTO SmallProduct(name, price)
SELECT pname, price
FROM Product
WHERE price <= 0.99

DELETE FROM Product
WHERE price <= 0.99
COMMIT

Client 2: SET TRANSACTION READ ONLY
START TRANSACTION
SELECT count(*)
FROM Product

SELECT count(*)
FROM SmallProduct
COMMIT
  
```

May improve performance

88

Commercial Systems

Always check documentation!

- DB2: Strict 2PL
- SQL Server:
 - Strict 2PL for standard 4 levels of isolation
 - Multiversion concurrency control for snapshot isolation
- PostgreSQL: Snapshot isolation; recently: serializable Snapshot isolation (!)
- Oracle: Snapshot isolation

89