

# Introduction to Data Management

## CSE 344

### Lecture 20: Transactions

# Where We Are

- HW5 due tonight!
- No more quizzes!
- HW6 to be posted – no late days!

**Transaction** = a collection of instructions to be executed all-or-nothing

**ACID**: atomic, consistent, isolated, durable

- **Atomic** = the recovery manager
- **Isolated** = the scheduler

# Schedules

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

# Example

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

# A Serial Schedule

T1

T2

---

READ(A)

A := A+100

WRITE(A)

READ(B)

B := B+100

WRITE(B)

READ(A)

A := A\*2

WRITE(A)

READ(B)

B := B\*2

WRITE(B)

# Serializable Schedule

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

# A Serializable Schedule

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

This is NOT a serial schedule,  
but is serializable

# 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)	

This is *non-serializable*



# Lock-Based Scheduler

Simple idea:

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

# Notation

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

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

# 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)	

# Example

T1

$L_1(A)$ ; READ(A)

A := A+100

WRITE(A);  $U_1(A)$ ;  $L_1(B)$

READ(B)

B := B+100

WRITE(B);  $U_1(B)$ ;

T2

$L_2(A)$ ; READ(A)

A := A\*2

WRITE(A);  $U_2(A)$ ;

$L_2(B)$ ; DENIED...

...GRANTED; READ(B)

B := B\*2

WRITE(B);  $U_2(B)$ ;

Scheduler has ensured a conflict-serializable schedule

# But...

T1

$L_1(A)$ ; READ(A)  
A := A+100  
WRITE(A);  $U_1(A)$ ;

$L_1(B)$ ; READ(B)  
B := B+100  
WRITE(B);  $U_1(B)$ ;

T2

$L_2(A)$ ; READ(A)  
A := A\*2  
WRITE(A);  $U_2(A)$ ;  
 $L_2(B)$ ; READ(B)  
B := B\*2  
WRITE(B);  $U_2(B)$ ;

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

# Two Phase Locking (2PL)

The 2PL rule:

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

# Example: 2PL transactions

T1

$L_1(A)$ ;  $L_1(B)$ ; READ(A)

A := A+100

WRITE(A);  $U_1(A)$

READ(B)

B := B+100

WRITE(B);  $U_1(B)$ ;

T2

$L_2(A)$ ; READ(A)

A := A\*2

WRITE(A);

$L_2(B)$ ; DENIED...

...GRANTED; READ(B)

B := B\*2

WRITE(B);  $U_2(A)$ ;  $U_2(B)$ ;

Now it is serializable

# A New Problem: Non-recoverable Schedule

T1

---

$L_1(A)$ ;  $L_1(B)$ ; READ(A)  
A := A+100  
WRITE(A);  $U_1(A)$

READ(B)  
B := B+100  
WRITE(B);  $U_1(B)$

Rollback

T2

$L_2(A)$ ; READ(A)  
A := A\*2  
WRITE(A);  
 $L_2(B)$ ; DENIED...

...GRANTED; READ(B)  
B := B\*2  
WRITE(B);  $U_2(A)$ ;  $U_2(B)$ ;  
Commit



# Strict 2PL

The Strict 2PL rule:

All locks are held until the transaction commits or aborts.

# Strict 2PL

T1

$L_1(A)$ ; READ(A)

A := A+100

WRITE(A);

$L_1(B)$ ; READ(B)

B := B+100

WRITE(B);

$U_1(A), U_1(B)$ ;

Rollback

T2

$L_2(A)$ ; DENIED...

...GRANTED; READ(A)

A := A\*2

WRITE(A);

$L_2(B)$ ; READ(B)

B := B\*2

WRITE(B);  $U_2(A)$ ;  $U_2(B)$ ;

Commit

# Deadlocks

- $T_1$  waits for a lock held by  $T_2$ ;
- $T_2$  waits for a lock held by  $T_3$ ;
- $T_3$  waits for . . . .
- . . . .
- $T_n$  waits for a lock held by  $T_1$

SQL Lite: there is only one exclusive lock; thus, never deadlocks

SQL Server: checks periodically for deadlocks and aborts one TXN

# Lock Modes

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

Lock compatibility matrix:

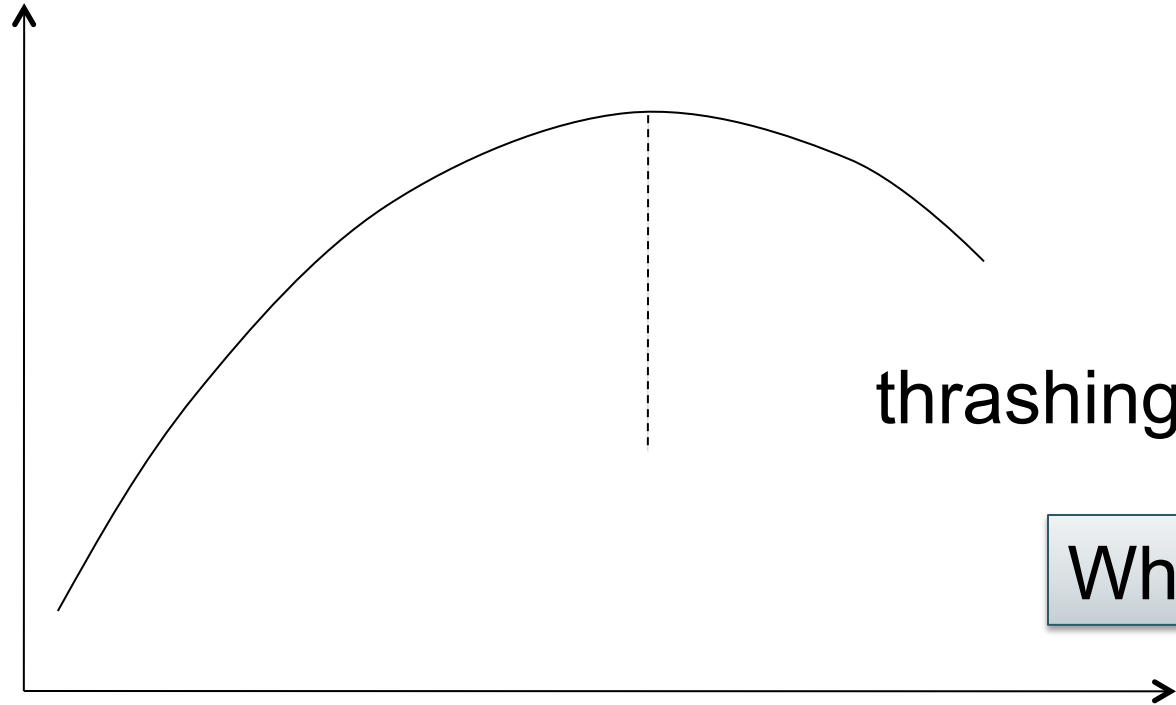
	None	S	X
None	OK	OK	OK
S	OK	OK	Conflict
X	OK	Conflict	Conflict

# Lock Granularity

- **Fine granularity locking** (e.g., tuples)
  - High concurrency
  - High overhead in managing locks
  - E.g. SQL Server
- **Coarse grain locking** (e.g., tables, entire database)
  - Many false conflicts
  - Less overhead in managing locks
  - E.g. SQL Lite

# Lock Performance

Throughput (TPS)



thrashing

Why ?

TPS =  
Transactions  
per second

# Active Transactions

# 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

# 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, A1, A2:

Is this schedule serializable ?



# Phantom Problem

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

Suppose there are two blue products, A1, A2:

Is this schedule serializable ?

NO: T1: sees 2 products the first time, then sees 3 products the second time

# 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, A1, A2:

R1(A1),R1(A2),W2(A3),R1(A1),R1(A2),R1(A3)

# Phantom Problem

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

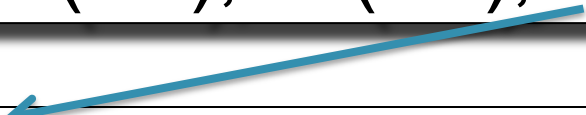
SELECT \*  
FROM Product  
WHERE color='blue'

When seen as a sequence of R/W,  
the schedule appears serializable.  
Locks ***cannot*** prevent this schedule.

Suppose there are two blue products, A1, A2:

R1(A1),R1(A2),W2(A3),R1(A1),R1(A2),R1(A3)

W2(A3),R1(A1),R1(A2),R1(A1),R1(A2),R1(A3)



# 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 !

# 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 !

# Isolation Levels in SQL

1. “Dirty reads”

SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED

2. “Committed reads”

SET TRANSACTION ISOLATION LEVEL READ COMMITTED

3. “Repeatable reads”

SET TRANSACTION ISOLATION LEVEL REPEATABLE READ

4. Serializable transactions

SET TRANSACTION ISOLATION LEVEL SERIALIZABLE



ACID

# 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

## 2. Isolation Level: Read Committed

- “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



### 3. Isolation Level: Repeatable Read

- “Long duration” WRITE locks
  - Strict 2PL
- “Long duration” READ locks
  - Strict 2PL

This is not serializable yet !!!

Why ?

# 4. Isolation Level Serializable

- “Long duration” WRITE locks
  - Strict 2PL
- “Long duration” READ locks
  - Strict 2PL
- Also deals with phantoms