

CSC 742 Database Management Systems

Topic #14: Concurrency Control – 2 Phase Locking

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

1

Locks

Locks are objects that describe the usage status of a data item

- A data item may be
 - ◆ a record
 - ◆ a field
 - ◆ a page
 - ◆ an index
 - ◆ a table
 - ◆ the whole DB
- Granularity determines concurrency and overhead (hence a trade-off).

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

2

Kinds of Locks

- Binary locks:
 - ◆ Conceptually, each data item x needs a lock
 - ◆ Two operations:
 - ◆ lock(x)
 - ◆ unlock(x)
 - ◆ Must be atomic
 - ◆ Gives mutex but restrictive
 - ◆ Implementation
 - ◆ Lock table: Stores the active locks
 - ◆ Lock manager: maintain lock table

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

3

Use Binary Lock for Transactions

- A transaction T
 - ◆ Lock(x) before read(x) or write(x)
 - ◆ Unlock(x) after all read(x) and write(x) are completed
 - ◆ Will not issue lock(x) if it already has the lock on x
 - ◆ Will not unlock(x) unless it already has the lock on x .
- Question:
 - ◆ What if no transaction write(x)?

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

4

Kinds of Locks (Cont'd)

- Multimode
 - ◆ Intuition: distinguish locks for read(x) and write(x)
 - ◆ shared-lock(x) read(x): multiple transactions can read x concurrently.
 - ◆ exclusive-lock(x) write(x): only one transaction can write x at each time.

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

5

Use Multimode Locks

- A transaction T
 - ◆ Read_lock(x) or write_lock(x) before read(x)
 - ◆ Write_lock(x) before write(x)
 - ◆ Unlock(x) after all read(x) and write(x) are completed
 - ◆ Will not issue read_lock(x) if it already has a read lock on x
 - ◆ Will not issue write_lock(x) if it already has a write lock on x
 - ◆ Will not unlock(x) unless it already has a read or write lock on x .

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

6

Lock Conversion

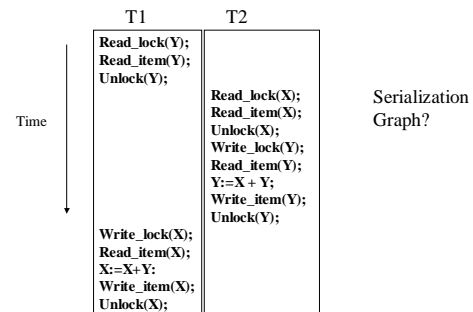
- Lock conversion:
 - ◆ can be upgraded (read to write)
 - ◆ or downgraded (write to read)

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

7

Does locking guarantee serializability?



Spring 2002

CSC 742: DBMS by Dr. Peng Ning

8

Two-Phase Locking

Moral: can't release locks too soon


- 2PL: All locking operations precede the first unlock operation.
 - ◆ growing phase
 - ◆ shrinking phase
- Guarantees serializability, but can lead to deadlock

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

9

Are these transactions using 2PL?



Spring 2002

CSC 742: DBMS by Dr. Peng Ning

10

Basic 2PL

- Rules for basic 2PL scheduler
 - ◆ For any operation $p_i(x)$ (p is read or write), test if $p_lock_i(x)$ conflicts with some $q_lock_j(x)$ that is already set. If so, it delays $p_i(x)$ until it can set $p_lock_i(x)$. If not, set $p_lock_i(x)$.
 - ◆ No concurrent access to the same item.
 - ◆ Once the scheduler has set $p_lock_i(x)$, it may not release it at least until $p_i(x)$ has been performed.
 - ◆ Further guarantee no concurrent access.
 - ◆ Once the scheduler has released a lock for T_i , it may not obtain any more locks for T_i .
 - ◆ Two phase rule


Spring 2002

CSC 742: DBMS by Dr. Peng Ning

11

2PL

- 2PL guarantees serializability.
- Deadlock



Spring 2002

CSC 742: DBMS by Dr. Peng Ning

12

Conservative 2PL

- Conservative or static 2PL
 - ◆ Obtain all locks before any operation
 - ◆ Make transaction wait (without any lock) if not all the locks can be obtained.
 - ◆ No deadlock: If T is waiting for a lock held by T', then T has no lock.
 - ◆ Disadvantage: you have to know what locks a transaction needs
 - ◆ How to get Read set and write set?

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

13

Strict 2PL

- Strict 2PL
 - ◆ Release all locks at once when the transaction commits or aborts
 - ◆ ensures strict schedules
 - ◆ but can deadlock

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

14

Deadlock Prevention

- Pessimistic: prevent deadlock from even becoming possible by restricting access when T_i tries to get an element locked by T_j
- Deadlock prevention using timestamps (TS)
 - ◆ An older transaction has smaller TS.
 - ◆ Two variations:
 - ◆ Wait-die
 - ◆ Wound-wait

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

15

Deadlock Prevention (Cont'd)

- Suppose T_i tries to lock x but is not able to because x is locked by T_j with a conflicting lock.
 - ◆ *wait-die*:
 - ◆ If $TS(T_i) < TS(T_j)$ then wait T_i
 - ◆ else abort T_i and restart with same time
 - ◆ Old transactions are allowed to wait.
 - ◆ How can wait-die prevent deadlock?

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

16

Deadlock Prevention (Cont'd)

- Suppose T_i tries to lock x but is not able to because x is locked by T_j with a conflicting lock.
 - ◆ *wound-wait*:
 - ◆ If $TS(T_i) < TS(T_j)$ abort T_j and restart with some timestamp,
 - ◆ else T_i wait
 - ◆ Young transactions are allowed to wait.
 - ◆ How can wound-wait prevent deadlock?

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

17

Deadlock Prevention (Cont'd)

- Prevent deadlock by Limiting Waiting
 - ◆ *No waiting*: abort transaction immediately if lock not obtained
 - ◆ *Cautious waiting*: abort transaction only if current lock holder is itself blocked

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

18

Deadlock Detection

- Optimistic strategy
- Detect a cycle in *waits-for graph*
- Choose a *victim* transaction
- Abort it thereby removing the deadlock
- Potentially unfair: the same victim is repeatedly chosen

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

19


Deadlock Detection

- Wait-for Graph
 - ◆ One node for each transaction
 - ◆ An edge from T_i to T_j if T_i is waiting to lock x that is currently locked by T_j .
 - ◆ Cycle means deadlock.

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

20



Spring 2002

CSC 742: DBMS by Dr. Peng Ning

21

Multiversion 2PL

- Basic idea:
 - ◆ Maintain up to two versions of each data item x .
 - ◆ Each x must have one committed version, supplied to transactions that read x .
 - ◆ Create a new version when T needs to write x
 - ◆ Once T that writes x is ready to commit, it must obtain a *certify lock* on all items that it currently holds write locks on before it can commit.
 - ◆ To install new versions.

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

22

Multi-version 2PL (Cont'd)

- Lock compatibility tables

- ◆ 2PL

	Read	Write
Read	Yes	No
Write	No	No

What do we gain via multi-version 2PL?

- ◆ Multi-version 2PL


	Read	Write	Certify
Read	Yes	Yes	No
Write	Yes	No	No
Certify	No	No	No

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

23

Is deadlock possible in multi-version 2PL?



Spring 2002

CSC 742: DBMS by Dr. Peng Ning

24

Multi-granularity locking

- Granularity: the size of a data item
 - ◆ Database
 - ◆ Database file
 - ◆ Disk block
 - ◆ Relation
 - ◆ Tuple

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

25

Multi-granularity locking (Cont'd)

- Transaction 1: update 75% of the tuples in relation Employee.
- Transaction 2: update 1 tuple in relation Employee.
- How should we set the granularity of data items?
 - ◆ Coarse: less concurrency
 - ◆ Fine: more locks


Spring 2002

CSC 742: DBMS by Dr. Peng Ning

26

Multi-granularity locking (Cont'd)


- Basic idea:
 - ◆ Support multiple granularities.



Spring 2002

CSC 742: DBMS by Dr. Peng Ning

27



T1: I want to read_lock(r111). Is there any conflicting lock?

T2: I want to read_lock(f1). Is there any conflicting lock?

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

28

Multi-granularity locking (Cont'd)

- Solution to reducing search for conflicting locks
 - ◆ Intention lock:
 - ◆ For the nodes along the path from the root to the item of choice (excluding the final node)
 - ◆ Indicate what types of lock T wants to obtain for the current node's descendants

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

29

Multi-granularity locking (Cont'd)

- Intention locks:
 - ◆ Intention-shared (IS): a shared lock will be requested on some descendants
 - ◆ Intention-exclusive (IX): an exclusive lock will be requested on some descendants
 - ◆ Shared-intention-exclusive (SIX): the current node is locked in shared mode, but an exclusive lock will be requested on some descendants.

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

30

T1: I want to read (r111).

T2: I want to write (r111).

T3: I want to go through the Employee relation stored in f1 and update the tuples with Salary > 30000.

What locks will be requested?

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 31

Compatibility matrix for multi-granularity locking

	IS	IX	S	SIX	X
IS					
IX					
S					
SIX					
X					

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 32

Multi-granularity locking protocol

1. The lock compatibility matrix must be adhered to.
2. The root of the tree must be locked first, in any mode.
3. A node N can be locked by T in S or IS only if the parent node is already locked by T in IS or IX.
4. A node N can be locked by T in X, IX, or SIX mode only if the parent is already locked by T in IX or SIX mode.
5. T can lock a node only if it has not unlocked any node (2-phase rule).
6. T can unlock a node N only if none of the children of N are locked by T (2-phase rule).

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 33

Phantom problem

- Phantom problem occurs when there are insertions.
 - ◆ When a new record being inserted by T satisfies a condition that a set of records accessed by T' must satisfy.

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 34

Accounts			Assets	
Account#	Location	Balance	Location	Balance
111	Raleigh	100	Raleigh	100
222	Apex	200	Apex	500
333	Apex	300		

What's the Result?

T1

```

Read(Accounts[111]);
Read(Accounts[222]);
Read(Accounts[333]);

Compute assets[Raleigh];
Compute assets[Apex];
Write(Assets[Raleigh]);
write(Assets[Apex]);

```

T2

```

Insert(Accounts[444,
Raleigh, 100])

```

What's wrong?

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 35

Phantom problem (Cont'd)

- Solutions
 - ◆ Index locking
 - ◆ Predicate locking

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 36

Optimistic Concurrency Control

- Three phases of a transaction T
 - ◆ *Read phase*: T reads data, updates local copies
 - ◆ *Validation phase*: check to ensure that serializability will not be violated if the updates are applied to the DB
 - ◆ *Write phase*: if valid, write to DB
- Basic idea: do all checks at once.
- write-set(T): items written by T
- read-set(T): items read by T

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

37

Optimistic Protocol

- Validate T_i w.r.t. any T_j that committed or is being validated
 - ◆ T_j completed its write phase before T_i began its read phase
 - ◆ Serial transactions
 - ◆ T_i starts its write phase after T_j completes its write phase, and $\text{read_set}(T_i) \cap \text{write_set}(T_j) = \emptyset$.
 - ◆ All possible conflicting pairs of operations are from T_j to T_i .
 - ◆ T_j completed its read phase before T_i completes its read phase, $\text{read_set}(T_i) \cap \text{write_set}(T_j) = \emptyset$, and $\text{write_set}(T_i) \cap \text{write_set}(T_j) = \emptyset$.
 - ◆ All possible conflicting pairs of operations are from T_j to T_i .

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

38