



Chapter 10 : Concurrency Control

modified from:

Database System Concepts, 6th Ed.

©Silberschatz, Korth and Sudarshan

See www.db-book.com for conditions on re-use



Intuition of Lock-based Protocols

- n Transactions have to acquire locks on data items before accessing them
- n If a lock is hold by one transaction on a data item this restricts the ability of other transactions to acquire locks for that data item
- n By locking a data item we want to ensure that no access to that data item is possible that would lead to non-serializable schedules
- n The trick is to design a lock model and protocol that guarantees that
- n Lock-based concurrency protocols are a form of **pessimistic concurrency control mechanism**
 - | We avoid ever getting into a state that can lead to a non-serializable schedule
- n Alternative concurrency control mechanism do not avoid conflicts, but determine later on (at commit time) whether committing a transaction would cause a non-serializable schedule to be generated
 - | **Optimistic concurrency control mechanism**



Lock-Based Protocols

- n A lock is a mechanism to control concurrent access to a data item
- n Data items can be locked in two modes :
 1. *exclusive (X) mode*. Data item can be both read as well as written. X-lock is requested using **lock-X** instruction.
 2. *shared (S) mode*. Data item can only be read. S-lock is requested using **lock-S** instruction.
- n Lock requests are made to concurrency-control manager.
 - | Transaction do not access data items before having acquired a lock on that data item
 - | Transactions release their locks on a data item only after they have accessed a data item



Lock-Based Protocols (Cont.)

n Lock-compatibility matrix

	S	X
S	true	false
X	false	false

- n A transaction may be granted a lock on an item if the requested lock is compatible with locks already held on the item by other transactions
- n Any number of transactions can hold shared locks on an item,
 - | but if any transaction holds an exclusive lock on the item no other transaction may hold any lock on the item.
- n If a lock cannot be granted, the requesting transaction is made to wait till all incompatible locks held by other transactions have been released. The lock is then granted.



Lock-Based Protocols (Cont.)

- n Example of a transaction performing locking:

```
 $T_2$ : lock-S(A);  
      read (A);  
      unlock(A);  
      lock-S(B);  
      read (B);  
      unlock(B);  
      display(A+B)
```

- n Locking as above is not sufficient to guarantee serializability — if A and B get updated in-between the read of A and B , the displayed sum would be wrong.
- n A **locking protocol** is a set of rules followed by all transactions while requesting and releasing locks. Locking protocols restrict the set of possible schedules.



Pitfalls of Lock-Based Protocols

- n Consider the partial schedule

T_3	T_4
lock-x (B)	
read (B)	
$B := B - 50$	
write (B)	
	lock-s (A)
	read (A)
	lock-s (B)
lock-x (A)	

- n Neither T_3 nor T_4 can make progress — executing **lock-S(B)** causes T_4 to wait for T_3 to release its lock on B , while executing **lock-X(A)** causes T_3 to wait for T_4 to release its lock on A .
- n Such a situation is called a **deadlock**.
 - l To handle a deadlock one of T_3 or T_4 must be rolled back and its locks released.



Pitfalls of Lock-Based Protocols (Cont.)

- n The potential for deadlock exists in most locking protocols. Deadlocks are a necessary evil.
- n **Starvation** is also possible if the concurrency control manager is badly designed. For example:
 - | A transaction may be waiting for an X-lock on an item, while a sequence of other transactions request and are granted an S-lock on the same item.
 - | The same transaction is repeatedly rolled back due to deadlocks.
- n Concurrency control managers can be designed to prevent starvation.



The Two-Phase Locking Protocol

- n This is a protocol which ensures conflict-serializable schedules.
- n Phase 1: Growing Phase
 - | transaction may obtain locks
 - | transaction may not release locks
- n Phase 2: Shrinking Phase
 - | transaction may release locks
 - | transaction may not obtain locks
- n The protocol assures serializability. It can be proved that the transactions can be serialized in the order of their **lock points** (i.e. the point where a transaction acquired its final lock).



The Two-Phase Locking Protocol (Cont.)

- n Two-phase locking *does not* ensure freedom from deadlocks
- n Cascading roll-back is possible under two-phase locking. To avoid this, follow a modified protocol called **strict two-phase locking (S2PL)**. Here a transaction must hold all its exclusive locks till it commits/aborts.
- n **Rigorous two-phase locking (SS2PL)** is even stricter: here *all* locks are held till commit/abort. In this protocol transactions can be serialized in the order in which they commit.



The Two-Phase Locking Protocol (Cont.)

- n There can be conflict serializable schedules that cannot be obtained if two-phase locking is used.
- n However, in the absence of extra information (e.g., ordering of access to data), two-phase locking is needed for conflict serializability in the following sense:

Given a transaction T_i that does not follow two-phase locking, we can find a transaction T_j that uses two-phase locking, and a schedule for T_i and T_j that is not conflict serializable.



Lock Conversions

- n Two-phase locking with lock conversions:
 - First Phase:
 - | can acquire a lock-S on item
 - | can acquire a lock-X on item
 - | can convert a lock-S to a lock-X (upgrade)
 - Second Phase:
 - | can release a lock-S
 - | can release a lock-X
 - | can convert a lock-X to a lock-S (downgrade)
- n This protocol assures serializability. But still relies on the programmer to insert the various locking instructions.



Automatic Acquisition of Locks

- n A transaction T_i issues the standard read/write instruction, without explicit locking calls.
- n The operation **read**(D) is processed as:
 - if** T_i has a lock on D
 - then**
 - read(D)
 - else begin**
 - if necessary wait until no other transaction has a **lock-X** on D
 - grant T_i a **lock-S** on D ;
 - read(D)
 - end**



Automatic Acquisition of Locks (Cont.)

- n **write(D)** is processed as:
 - if T_i has a **lock-X** on D
 - then
 - write(D)
 - else begin
 - if necessary wait until no other trans. has any lock on D ,
 - if T_i has a **lock-S** on D
 - then
 - upgrade lock on D to **lock-X**
 - else
 - grant T_i a **lock-X** on D
 - write(D)
 - end;
 - n All locks are released after commit or abort




Implementation of Locking

- n A **lock manager** can be implemented as a separate process to which transactions send lock and unlock requests
- n The lock manager replies to a lock request by sending a lock grant messages (or a message asking the transaction to roll back, in case of a deadlock)
- n The requesting transaction waits until its request is answered
- n The lock manager maintains a data-structure called a **lock table** to record granted locks and pending requests
- n The lock table is usually implemented as an in-memory hash table indexed on the name of the data item being locked



Lock Table



- n Lock table also records the type of lock granted or requested
- n New request is added to the end of the queue of requests for the data item, and granted if it is compatible with all earlier locks
- n Unlock requests result in the request being deleted, and later requests are checked to see if they can now be granted
- n If transaction aborts, all waiting or granted requests of the transaction are deleted
 - | lock manager may keep a list of locks held by each transaction, to implement this efficiently



Deadlock Handling

- n Consider the following two transactions:

T_1 : write (X) T_2 : write(Y)
 write(Y) write(X)

- n Schedule with deadlock

T_1	T_2
lock-X on A write (A)	
	lock-X on B write (B) wait for lock-X on A
wait for lock-X on B	



Deadlock Handling

- n System is deadlocked if there is a set of transactions such that every transaction in the set is waiting for another transaction in the set.
- n **Deadlock prevention** protocols ensure that the system will *never* enter into a deadlock state. Some prevention strategies :
 - | Require that each transaction locks all its data items before it begins execution (predeclaration).
 - ▶ Not practical
 - | Impose partial ordering of all data items and require that a transaction can lock data items only in the order specified by the partial order (graph-based protocol).



More Deadlock Prevention Strategies

- n Following schemes use transaction timestamps for the sake of deadlock prevention alone.
 - | **Preemptive**: Transaction holding a lock is aborted to make lock available
- n **wait-die** scheme — non-preemptive
 - | older transaction may wait for younger one to release data item. Younger transactions never wait for older ones; they are rolled back instead.
 - | a transaction may die several times before acquiring needed data item
- n **wound-wait** scheme — preemptive
 - | older transaction *wounds* (forces rollback) of younger transaction instead of waiting for it. Younger transactions may wait for older ones.
 - | may be fewer rollbacks than *wait-die* scheme.



Deadlock prevention (Cont.)

- n Both in *wait-die* and in *wound-wait* schemes, a rolled back transactions is restarted with its original timestamp. Older transactions thus have precedence over newer ones, and starvation is hence avoided.
- n **Timeout-Based Schemes:**
 - | a transaction waits for a lock only for a specified amount of time. After that, the wait times out and the transaction is rolled back.
 - | thus deadlocks are not possible
 - | simple to implement; but starvation is possible. Also difficult to determine good value of the timeout interval.




Deadlock Detection


- n Deadlocks can be described as a *wait-for graph*, which consists of a pair $G = (V, E)$,
 - | V is a set of vertices (all the transactions in the system)
 - | E is a set of edges; each element is an ordered pair $T_i \rightarrow T_j$.
- n If $T_i \rightarrow T_j$ is in E , then there is a directed edge from T_i to T_j , implying that T_i is waiting for T_j to release a data item.
- n When T_i requests a data item currently being held by T_j , then the edge $T_i \rightarrow T_j$ is inserted in the wait-for graph. This edge is removed only when T_j is no longer holding a data item needed by T_i .
- n The system is in a deadlock state if and only if the wait-for graph has a cycle. Must invoke a deadlock-detection algorithm periodically to look for cycles.



Deadlock Detection (Cont.)



Wait-for graph without a cycle



Wait-for graph with a cycle



Deadlock Recovery

- n When deadlock is detected :
 - | Some transaction will have to rolled back (made a victim) to break deadlock. Select that transaction as victim that will incur minimum cost.
 - | Rollback -- determine how far to roll back transaction
 - ▶ **Total rollback**: Abort the transaction and then restart it.
 - ▶ More effective to roll back transaction only as far as necessary to break deadlock.
 - | Starvation happens if same transaction is always chosen as victim. Include the number of rollbacks in the cost factor to avoid starvation



Weak Levels of Consistency

- n **Degree-two consistency:** differs from two-phase locking in that S-locks may be released at any time, and locks may be acquired at any time
 - | X-locks must be held till end of transaction
 - | Serializability is not guaranteed, programmer must ensure that no erroneous database state will occur]
- n **Cursor stability:**
 - | For reads, each tuple is locked, read, and lock is immediately released
 - | X-locks are held till end of transaction
 - | Special case of degree-two consistency



Weak Levels of Consistency in SQL

- n SQL allows non-serializable executions
 - | **Serializable**: is the default
 - | **Repeatable read**: allows only committed records to be read, and repeating a read should return the same value (so read locks should be retained)
 - ▶ However, the phantom phenomenon need not be prevented
 - T1 may see some records inserted by T2, but may not see others inserted by T2
 - | **Read committed**: same as degree two consistency, but most systems implement it as cursor-stability
 - | **Read uncommitted**: allows even uncommitted data to be read
- n In many database systems, read committed is the default consistency level
 - | has to be explicitly changed to serializable when required
 - ▶ **set isolation level serializable**



Recap

n Concurrency Control

- | **Pessimistic:** Prevent bad things from happening
 - ▶ Locking Protocols
- | **Optimistic:** Detect that bad things have happened and resolve the problem

n Two-Phase Locking (2PL)

- | Two types of locks:
 - ▶ Shared (S) locks for read-only access
 - ▶ Exclusive (X) locks for write + read access
- | Lock compatibility
- | Transactions cannot acquire locks after they have released a lock
 - ▶ Divides transaction into growing and shrinking phase
- | **Ensures conflict-serializability**
- | **Cascading rollbacks are possible**
- | **Deadlocks are possible**



Recap

- n **Strict Two-Phase Locking (S2PL)**
 - | Exclusive locks are held until transaction commit
 - | **Prevents cascading rollbacks**
 - | **Deadlocks are still possible**
- n **Strict Strong Two-Phase Locking (SS2PL)**
 - | All locks are held until transaction commit
 - | **Enables serializability in commit order**
- n **Deadlocks**
 - | **Deadlock Prevention**
 - ▶ **Wait-die:** Younger transaction that waits for older is rolled back
 - ▶ **Wound-wait:** If older waits for younger, then younger is rolled back
 - | **Deadlock Detection**
 - ▶ Cycle Detection in Waits-for graph
 - Expensive
 - ▶ Timeout



End of Chapter

Thanks to Alan Fekete and Sudhir Jorwekar for Snapshot Isolation examples

modified from:

Database System Concepts, 6th Ed.

©Silberschatz, Korth and Sudarshan

See www.db-book.com for conditions on re-use