

Lecture 17: Two-Phase Locking

15-445/645 Database Systems (Fall 2017)

Carnegie Mellon University

Prof. Andy Pavlo

Review

1. Conflict Serializable

- (a) Verify using either the swapping method or dependency graphs.
- (b) Any DBMS that says they support serializable isolation does this.

2. View Serializable

- (a) No efficient way to verify.
- (b) Requires knowledge of the program.
- (c) Probably no DBMS does this.

Locks vs. Latches

1. Locks

- (a) Protect the logical contents from other transactions.
- (b) Held for duration of transaction.
- (c) Need to be able to rollback changes.

2. Latches

- (a) Protects the critical section of the internal data structure from other threads
- (b) Held for operation duration
- (c) Do not need to be able to rollback changes

Transaction Locks

The DBMS contains a centralized lock manager that decides decisions whether a transaction can have a lock or not. It has a global view of whats going on inside the system.

1. **Shared Lock (S-LOCK):** A lock that allows multiple transactions to read the same object at the same time. If one transaction holds a shared lock, then another transaction can acquire that same shared lock.

2. **Exclusive Lock (X-LOCK):** Allows a transaction to modify an object. This lock is not compatible for any other lock. Only one transaction can hold an exclusive lock at a time.
3. Shared locks can be shared with shared locks, but exclusive locks can't be shared with anything.
4. Executing with locks:
 - (a) Transactions request locks (or upgrades) from the lock manager.
 - (b) The lock manager grants or blocks requests based on what locks are currently held by other transactions.
 - (c) Transactions release locks when they no longer need them.
 - (d) The lock manager updates its internal lock-table and then gives locks to waiting transactions.

Two-Phase Locking

Two-Phase locking (2PL) is a concurrency control protocol that determines whether a transaction is allowed to access an object in the database on the fly. The protocol does not need to know all of the queries that a transaction will execute ahead of time.

1. Phase #1: Growing

- (a) Each transaction requests the locks that it needs from the DBMS's lock manager.
- (b) The lock manager grants/denies lock requests.

2. Phase #2: Shrinking

- (a) The transaction enters this phase immediately after it releases its first lock.
 - (b) The transaction is allowed to only release locks that it previously acquired. It cannot acquire new locks in this phase.
3. On its own, 2PL is sufficient to guarantee conflict serializability. It generates schedules whose precedence graph is acyclic.
 4. It is susceptible to **cascading aborts**, which is when a transaction aborts and now another transaction must be rolled back, which results in wasted work.
 5. There are potential schedules that are serializable but would not be allowed by 2PL (locking can limit concurrency).

Strict Two-Phase Locking

The transaction only releases locks when it finishes. There really isn't a shrinking phase like in regular 2PL.

1. A schedule is **strict** if a value written by a transaction is not read or overwritten by other transactions until that transaction finishes.
2. The advantage of this approach is that the DBMS does not incur cascading aborts. The DBMS can also reverse the changes of an aborted transaction by just restoring original values of modified tuples.

2PL Deadlock Handling

A **deadlock** is a cycle of transactions waiting for locks to be released by each other. There are two approaches to handling deadlocks in 2PL: detection and prevention.

Approach #1: Deadlock Detection

The DBMS creates a **waits-for** graph: Nodes are transactions, and edge from T_i to T_j if transaction T_i is waiting for transaction T_j to release a lock. The system will periodically check for cycles in waits-for graph and then make a decision on how to break it.

1. When the DBMS detects a deadlock, it will select a “victim” transaction to rollback to break the cycle.
2. The victim transaction will either restart or abort depending on how the application invoked it
3. There are multiple transaction properties to consider when selecting a victim. There is no one choice that is better than others. 2PL DBMSs all do different things:
 - (a) By age (newest or oldest timestamp).
 - (b) By progress (least/most queries executed).
 - (c) By the # of items already locked.
 - (d) By the # of transactions that we have to rollback with it.
 - (e) # of times a transaction has been restarted in the past
4. **Rollback Length:** After selecting a victim transaction to abort, the DBMS can also decide on how far to rollback the transaction’s changes. Can be either the entire transaction or just enough queries to break the deadlock.

Approach #2: Deadlock Prevention

When a transaction tries to acquire a lock, if that lock is currently held by another transaction, then perform some action to prevent a deadlock. Assign priorities based on timestamps (e.g., older means higher priority). These schemes guarantee no deadlocks because only one type of direction is allowed when waiting for a lock. When a transaction restarts, its (new) priority is its old timestamp.

1. **Wait-Die (“Old waits for Young”):** If T_1 has higher priority, T_1 waits for T_2 . Otherwise T_1 aborts
2. **Wound-Wait (“Young waits for Old”):** If T_1 has higher priority, T_2 aborts. Otherwise T_1 waits.

Lock Granularities

If a transaction wants to update a billion tuples, it has to ask the DBMS’s lock manager for a billion locks. This will be slow because we have to take latches in the lock manager’s internal lock table data structure. Instead, we want to use a lock hierarchy that allow a transaction to take more coarse-grained locks in the system. Acquiring a lock for something in this hierarchy implicitly acquires a lock for all its children.

Intention locks allow a higher level node to be locked in shared or exclusive mode without having to check all descendant nodes. If a node is in an intention mode, then explicit locking is being done at a lower level in the tree.

1. **Intention-Shared (IS):** Indicates explicit locking at a lower level with shared locks.

2. **Intention-Exclusive (IX):** Indicates explicit locking at a lower level with exclusive or shared locks.
3. **Shared+Intention-Exclusive (SIX):** The subtree rooted at that node is locked explicitly in shared mode and explicit locking is being done at a lower level with exclusive-mode locks.

Conclusion

1. The application does not typically set locks manually. But sometimes it can provide the DBMS with hints to help it improve concurrency:
SELECT . . . FOR UPDATE: Perform a select and then sets an exclusive lock on fetched tuples
2. 2PL is used in almost all DBMS. It automatically provides correct interleavings of transaction operations. But it must handle deadlocks.