

# Lecture #16: Two-Phase Locking

15-445/645 Database Systems (Fall 2022)

<https://15445.courses.cs.cmu.edu/fall2022/>

Carnegie Mellon University

Andy Pavlo

## 1 Transaction Locks

---

A DBMS uses *locks* to dynamically generate an execution schedule for transactions that is serializable without knowing each transaction's read/write set ahead of time. These locks protect database objects during concurrent access when there are multiple readers and writes. The DBMS contains a centralized *lock manager* that decides whether a transaction can acquire a lock or not. It also provides a global view of what's going on inside the system.

There are two basic types of locks:

- **Shared Lock (S-LOCK):** A shared 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 also acquire that same shared lock.
- **Exclusive Lock (X-LOCK):** An exclusive lock allows a transaction to modify an object. This lock prevents other transactions from taking any other lock (S-LOCK or X-LOCK) on the object. Only one transaction can hold an exclusive lock at a time.

Transactions must request locks (or upgrades) from the lock manager. The lock manager grants or blocks requests based on what locks are currently held by other transactions. Transactions must release locks when they no longer need them to free up the object. The lock manager updates its internal lock-table with information about which transactions hold which locks and which transactions are waiting to acquire locks.

The DBMS's lock-table does not need to be durable since any transaction that is active (i.e., still running) when the DBMS crashes is automatically aborted.

Just the usage of locks does not automatically resolve all the issues associated with concurrent transactions. Locks need to be complemented by a concurrency control protocol.

## 2 Two-Phase Locking

---

Two-Phase locking (2PL) is a pessimistic concurrency control protocol that uses locks to determine 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.

**Phase #1– Growing:** In the growing phase, each transaction requests the locks that it needs from the DBMS's lock manager. The lock manager grants/denies these lock requests.

**Phase #2– Shrinking:** Transactions enter the shrinking phase immediately after it releases its first lock. In the shrinking phase, transactions are only allowed to release locks. They are not allowed to acquire new ones.

On its own, 2PL is sufficient to guarantee conflict serializability. It generates schedules whose precedence graph is acyclic. But 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.

2PL can still have dirty reads and it can also lead to deadlocks. There are also potential schedules that are serializable but would not be allowed by 2PL (locking can limit concurrency).

### Strong Strict Two-Phase Locking

A schedule is *strict* if any value written by a transaction is never read or overwritten by another transaction until the first transaction commits. *Strong Strict 2PL* (also known as *Rigorous 2PL*) is a variant of 2PL where the transactions only release locks when they commit.

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 restoring the original values of modified tuples. However, Strict 2PL generates more cautious/pessimistic schedules that limit concurrency.

### Universe of Schedules

$SerialSchedules \subset StrongStrict2PL \subset ConflictSerializableSchedules \subset ViewSerializableSchedules \subset AllSchedules$

## 3 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

To detect deadlocks, the DBMS creates a *waits-for* graph where transactions are nodes, and there exists a directed 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 the waits-for graph (usually with a background thread) and then make a decision on how to break it. Latches are not needed when constructing the graph since if the DBMS misses a deadlock in one pass, it will find it in the subsequent passes. Note that there is a tradeoff between the frequency of deadlock checks (uses cpu cycles) and the wait time till a deadlock is broken.

When the DBMS detects a deadlock, it will select a “victim” transaction to abort to break the cycle. The victim transaction will either restart or abort depending on how the application invoked it.

The DBMS can consider multiple transaction properties when selecting a victim to break the deadlock:

1. By age (newest or oldest timestamp).
2. By progress (least/most queries executed).
3. By the # of items already locked.
4. By the # of transactions needed to rollback with it.
5. # of times a transaction has been restarted in the past (to avoid starvation).

There is no one choice that is better than others. Many systems use a combination of these factors.

After selecting a victim transaction to abort, the DBMS can also decide on how far to rollback the transaction’s changes. It can either rollback the entire transaction or just enough queries to break the deadlock.

### Approach #2: Deadlock Prevention

Instead of letting transactions try to acquire any lock they need and then deal with deadlocks afterwards, deadlock prevention 2PL stops transactions from causing deadlocks before they occur. When a transaction tries to acquire a lock held by another transaction (which could cause a deadlock), the DBMS kills one of them. To implement this, transactions are assigned priorities based on timestamps (older transactions have

higher priority). These schemes guarantee no deadlocks because only one type of direction is allowed when waiting for a lock. When a transaction restarts, the DBMS reuses the same timestamp.

There are two ways to kill transactions under deadlock prevention:

- **Wait-Die (“Old Waits for Young”)**: If the requesting transaction has a higher priority than the holding transaction, it waits. Otherwise, it aborts.
- **Wound-Wait (“Young Waits for Old”)**: If the requesting transaction has a higher priority than the holding transaction, the holding transaction aborts and releases the lock. Otherwise, the requesting transaction waits.

## 4 Lock Granularities

---

If a transaction wants to update one billion tuples, it has to ask the DBMS’s lock manager for a billion locks. This will be slow because the transaction has to take latches in the lock manager’s internal lock table data structure as it acquires/releases locks.

To avoid this overhead, the DBMS can use to use a lock hierarchy that allows a transaction to take more coarse-grained locks in the system. For example, it could acquire a single lock on the table with one billion tuples instead of one billion separate locks. When a transaction acquires a lock for an object in this hierarchy, it implicitly acquires the locks for all its children objects.

### Database Lock Hierarchy:

1. Database level (Slightly Rare)
2. Table level (Very Common)
3. Page level (Common)
4. Tuple level (Very Common)
5. Attribute level (Rare)

**Intention locks** allow a higher level node to be locked in shared mode 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.

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