


Concurrency Control


Chapter 17



Lock Management

- ❖ Lock and unlock requests are handled by the lock manager
- ❖ Lock table entry:
 - Transactions currently holding a lock
 - Type of lock held (shared or exclusive)
 - Pointer to queue of lock requests
- ❖ Locking and unlocking have to be atomic operations
- ❖ Lock upgrade: transaction that holds a shared lock can be upgraded to hold an exclusive lock

Example



T1: S(A) R(A)

S(B) R(B)

T2: X(B) W(B)

X(C) W(C)

T3: S(C) R(C)


X(A) W(A)


T4: X(B) W(B)

Lock table:

Object	X	S	Queue
A		T1	T3
B	T2		T1, T4
C		T3	T2


Waits-for Graph:






Deadlocks

- ❖ Deadlock: Cycle of transactions waiting for locks to be released by each other.
- ❖ Two ways of dealing with deadlocks:
 - Deadlock prevention
 - Deadlock detection




Deadlock Prevention

- ❖ Assign priorities based on timestamps.
Assume T_i wants a lock that T_j holds. Two policies are possible:
 - **Wait-Die:** If T_i has higher priority, T_i waits for T_j ; otherwise T_i aborts
 - **Wound-wait:** If T_i has higher priority, T_j aborts; otherwise T_i waits
- ❖ If a transaction re-starts, make sure it has its original timestamp



Deadlock Detection



- ❖ Create a **waits-for graph**:
 - Nodes are transactions
 - There is an edge from T_i to T_j if T_i is waiting for T_j to release a lock
- ❖ Periodically check for cycles in the waits-for graph




Deadlock Detection (Continued)

Example:


T1: S(A), R(A), S(B)
T2: X(B), W(B) X(C)
T3: S(C), R(C) X(A)
T4: X(B)






Optimistic CC

- ❖ Locking is a conservative approach in which conflicts are prevented. Disadvantages:
 - Lock management overhead.
 - Deadlock detection/resolution.
 - Lock contention for heavily used objects.
- ❖ If conflicts are rare, we might be able to gain concurrency by not locking, and instead checking for conflicts before Xacts commit.




Optimistic CC Model

- ❖ Xacts have three phases:
 - **READ:** Xacts read from the database, but make changes to private copies of objects.
 - **VALIDATE:** Check for conflicts.
 - **WRITE:** Make local copies of changes public.




Validation


- ❖ Test conditions that are **sufficient** to ensure that no conflict occurred.
- ❖ Each Xact is assigned a numeric id.
 - Just use a **timestamp**.
- ❖ Xact ids assigned at end of READ phase, just before validation begins.
- ❖ **ReadSet(Ti)**: Set of objects read by Xact Ti.
- ❖ **WriteSet(Ti)**: Set of objects modified by Ti.



Test 1


- ❖ For all i and j such that $T_i < T_j$, check that T_i completes before T_j begins.






Test 2

- ❖ For all i and j such that $T_i < T_j$, check that:
 - T_i completes before T_j begins its Write phase +
 - $\text{WriteSet}(T_i) \cap \text{ReadSet}(T_j)$ is empty.




Does T_j read dirty data? Does T_i overwrite T_j 's writes?




Test 3

- ❖ For all i and j such that $T_i < T_j$, check that:
 - T_i completes Read phase before T_j does +
 - $WriteSet(T_i) \cap ReadSet(T_j)$ is empty +
 - $WriteSet(T_i) \cap WriteSet(T_j)$ is empty.




Does T_j read dirty data? Does T_i overwrite T_j 's writes?



Summary

- ❖ There are several lock-based concurrency control schemes (Strict 2PL, 2PL). Conflicts between transactions can be detected in the dependency graph
- ❖ The lock manager keeps track of the locks issued. Deadlocks can either be prevented or detected.



Summary (Contd.)

- ❖ Multiple granularity locking reduces the overhead involved in setting locks for nested collections of objects (e.g., a file of pages); should not be confused with tree index locking!
- ❖ Optimistic CC aims to minimize CC overheads in an “optimistic” environment where reads are common and writes are rare.
- ❖ Optimistic CC has its own overheads however; most real systems use locking.
- ❖ SQL-92 provides different isolation levels that control the degree of concurrency