


Carnegie Mellon Univ.
Dept. of Computer Science
15-415/615 - DB Applications


C. Faloutsos – A. Pavlo
 Lecture#22: Concurrency Control – Part 2
 (R&G ch. 17)



Last Class

- A **concurrency control** scheme uses locks and aborts to ensure correctness.
- Conflict vs. View Serializability
- (Strict) 2PL is popular.
- We need to handle deadlocks in 2PL:
 - **Detection:** *Waits-for* graph
 - **Prevention:** Abort some txns, defensively

Faloutsos/Pavlo CMU SCS 15-415/615 2



Last Class Assumption

- We assumed that the database was **fixed** collection of **independent** objects.
 - No objects are added or deleted.
 - No relationship between objects.
 - No indexes.

Faloutsos/Pavlo CMU SCS 15-415/615 3

CMU SCS

Today's Class

- Lock Granularities
- Locking in B+Trees
- The Phantom Problem
- Transaction Isolation Levels

Faloutsos/Pavlo CMU SCS 15-415/615 4

CMU SCS

Lock Granularities

- When we say that a txn acquires a “lock”, what does that actually mean?
 - On a field? Record? Page? Table?
- Ideally, each txn should obtain fewest number of locks that is needed...

Faloutsos/Pavlo CMU SCS 15-415/615 5


CMU SCS

Database Lock Hierarchy


```

graph TD
  Database[Database] --> Table1[Table 1]
  Database --> Table2[Table 2]
  Table1 --> Tuple1[Tuple 1]
  Table1 --> Tuple2[Tuple 2]
  Table1 --> TupleN[Tuple n]
  Tuple1 --> Attr1[Attr 1]
  Tuple1 --> Attr2[Attr 2]
  Tuple1 --> AttrN[Attr n]
  
```

Faloutsos/Pavlo CMU SCS 15-415/615 6


CMU SCS 

Example




- **T1:** Get the balance of Christos' shady off-shore bank account.
- **T2:** Increase all account balances by 1%.
- **Q:** What locks should they obtain?

Faloutsos/Pavlo CMU SCS 15-415/615 7


CMU SCS 

Example




- **Q:** What locks should they obtain?
- **A: Multiple**
 - **Exclusive + Shared** for leaves of lock tree.
 - Special **Intention** locks for higher levels

Faloutsos/Pavlo CMU SCS 15-415/615 8

CMU SCS 

Intention Locks




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

Faloutsos/Pavlo CMU SCS 15-415/615 9

CMU SCS

Intention Locks



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

Faloutsos/Pavlo CMU SCS 15-415/615 10

CMU SCS

Compatibility Matrix

		T2 Wants				
		IS	IX	S	SIX	X
T1 Holds	IS	✓	✓	✓	✓	✗
	IX		✓	✗	✗	✗
	S			✓	✗	✗
	SIX				✗	✗
	X					✗

Faloutsos/Pavlo CMU SCS 15-415/615 11

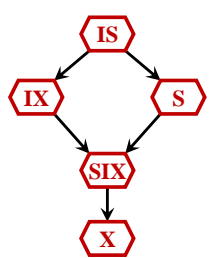
CMU SCS

Multiple Granularity Protocol

↑ Weaker

↑ Privileges

↓ Stronger



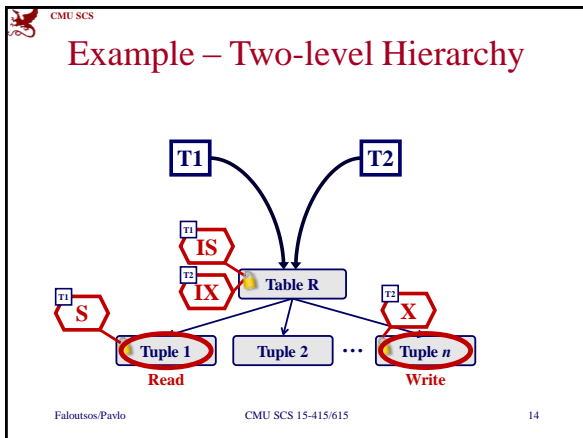
Faloutsos/Pavlo CMU SCS 15-415/615 12

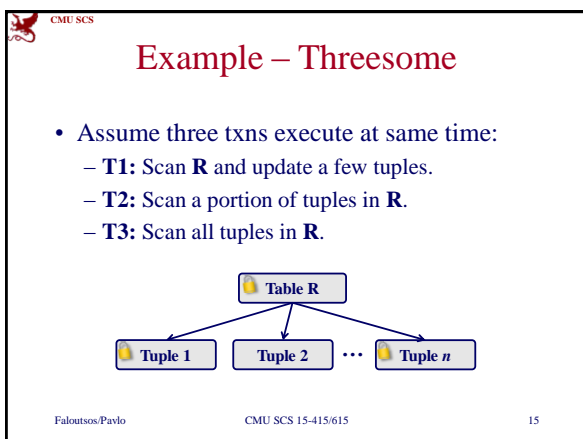
CMU SCS

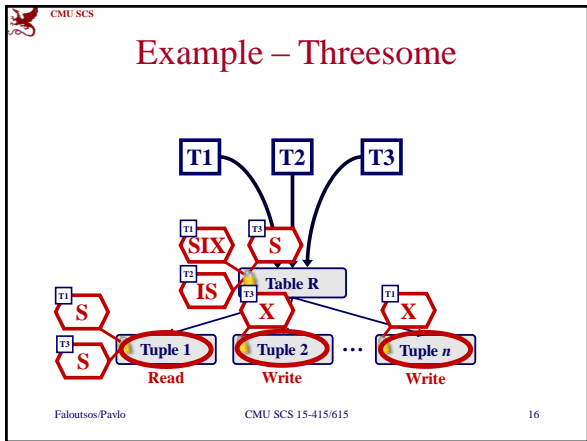
Locking Protocol

- Each txn obtains appropriate lock at highest level of the database hierarchy.
- To get **S** or **IS** lock on a node, the txn must hold at least **IS** on parent node.
 - What if txn holds **SIX** on parent? S on parent?
- To get **X**, **IX**, or **SIX** on a node, must hold at least **IX** on parent node.

Faloutsos/Pavlo CMU SCS 15-415/615 13







- ### Example - Threesome
- **T1:** Get an **SIX** lock on **R**, then get **X** lock on tuples that are updated.
 - **T2:** Get an **IS** lock on **R**, and repeatedly get an **S** lock on tuples of **R**.
 - **T3:** Two choices:
 - T3 gets an **S** lock on **R**.
 - OR, T3 could behave like T2; can use *lock escalation* to decide which.

Lock Escalation

- Lock escalation dynamically asks for coarser-grained locks when too many low level locks acquired.

CMU SCS

Multiple Lock Granularities

- Useful in practice as each txn only needs a few locks.
- Intention locks help improve concurrency:
 - **Intention-Shared (IS)**: Intent to get **S** lock(s) at finer granularity.
 - **Intention-Exclusive (IX)**: Intent to get **X** lock(s) at finer granularity.
 - **Shared+Intention-Exclusive (SIX)**: Like **S** and **IX** at the same time.

Faloutsos/Pavlo CMU SCS 15-415/615 19

CMU SCS

Today's Class

- Lock Granularities
- ➔ • Locking in B+Trees
- The Phantom Problem
- Transaction Isolation Levels

Faloutsos/Pavlo CMU SCS 15-415/615 20

CMU SCS

Locking in B+Trees

- **Q**: What about locking indexes?
- **A**: They are not quite like other database elements so we can treat them differently:
 - It's okay to have non-serializable concurrent access to an index as long as the accuracy of the index is maintained.

Faloutsos/Pavlo CMU SCS 15-415/615 21

CMU SCS

Example

- T1 wants to insert in H
- T2 wants to insert in I
- **Q:** Why not plain 2PL?
- **A:** Because txns have to hold on to their locks for too long!

Faloutsos/Pavlo CMU SCS 15-415/615 22

CMU SCS

Lock Crabbing

- Improves concurrency for B+Trees.
- Get lock for parent; get lock for child; release lock for parent if “safe”.
- **Safe Nodes:** Any node that won't split or merge when updated.
 - Not full (on insertion)
 - More than half-full (on deletion)

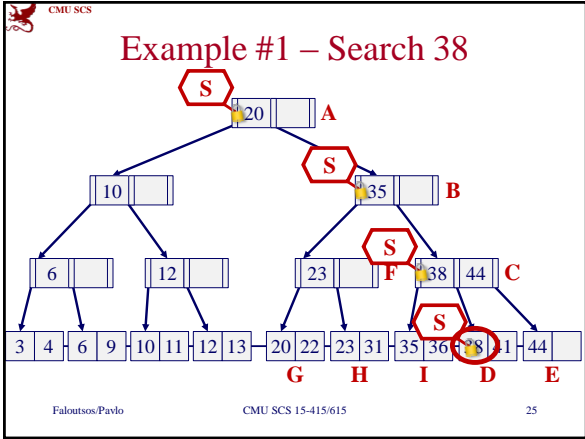
Faloutsos/Pavlo CMU SCS 15-415/615 23

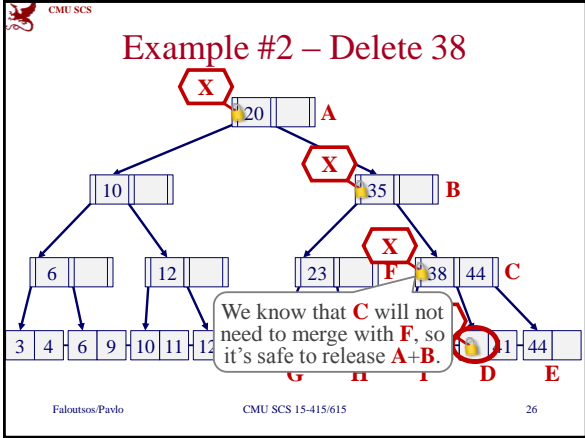
CMU SCS

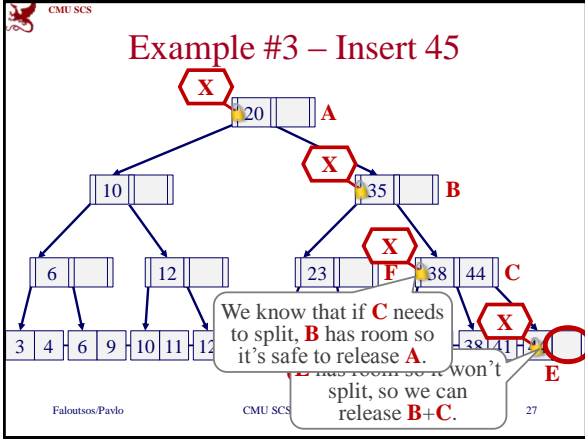
Lock Crabbing

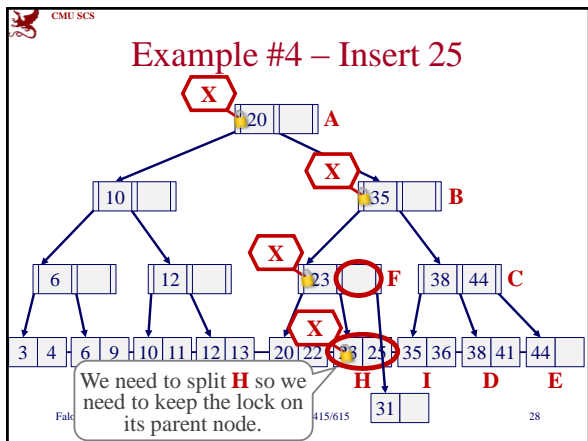
- **Search:** Start at root and go down; repeatedly,
 - **S** lock child
 - then unlock parent
- **Insert/Delete:** Start at root and go down, obtaining **X** locks as needed. Once child is locked, check if it is safe:
 - If child is safe, release all locks on ancestors.

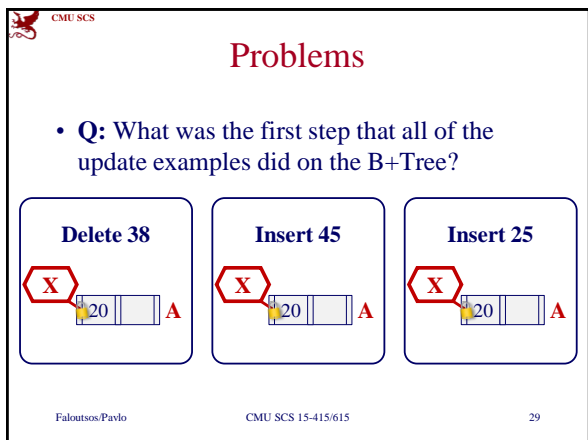
Faloutsos/Pavlo CMU SCS 15-415/615 24

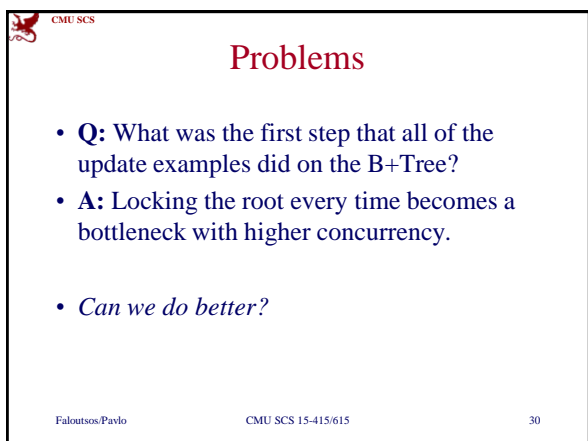













CMU SCS

Better Tree Locking Algorithm

- **Main Idea:**
 - Assume that the leaf is ‘safe’, and use S-locks & crabbing to reach it, and verify.
 - If leaf is not safe, then do previous algorithm.
- **Rudolf Bayer, Mario Schkolnick:**
Concurrency of Operations on B-Trees.
 Acta Inf. 9: 1-21 (1977)



Faloutsos/Pavlo CMU SCS 15-415/615

CMU SCS

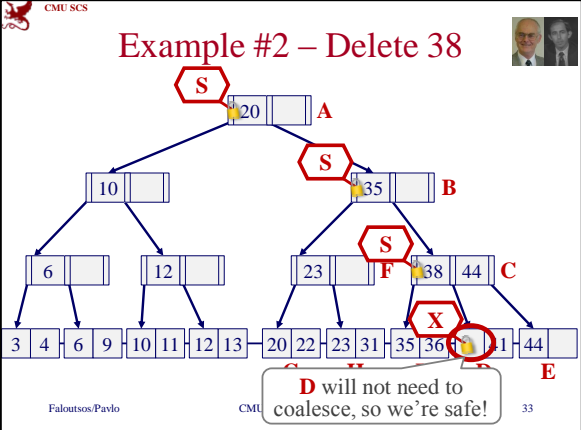
Better Tree Locking Algorithm

- **Search:** Same as before.
- **Insert/Delete:**
 - Set locks as if for search, get to leaf, and set **X** lock on leaf.
 - If leaf is not safe, release all locks, and restart txn using previous Insert/Delete protocol.
- Gambles that only leaf node will be modified; if not, **S** locks set on the first pass to leaf are wasteful.

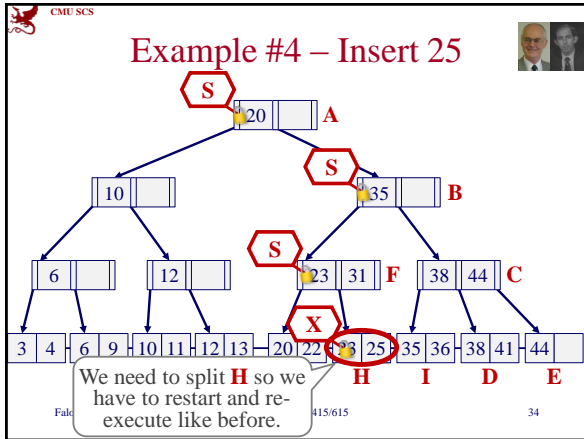
Faloutsos/Pavlo CMU SCS 15-415/615 32

CMU SCS

Example #2 – Delete 38



Faloutsos/Pavlo CMU 33



Another Alternative

- Textbook has a third variation, that uses lock-upgrades instead of restarting.
- This approach may lead to deadlocks.

Additional Points

- **Q:** Which order to release locks in multiple-granularity locking?
- **A:** From the bottom up
- **Q:** Which order to release locks in tree-locking?
- **A:** As early as possible to maximize concurrency.

CMU SCS

Today's Class

- Lock Granularities
- Locking in B+Trees
- ➔ • The Phantom Problem
- Transaction Isolation Levels

Faloutsos/Pavlo CMU SCS 15-415/615 37

CMU SCS

Dynamic Databases

- Recall that so far we have only dealing with transactions that read and update data.
- But now if we have insertions deletions, we have new problems...

Faloutsos/Pavlo CMU SCS 15-415/615 38

CMU SCS

The Phantom Problem

Schedule

	T1		T2
<div style="writing-mode: vertical-rl; transform: rotate(180deg);">TIME</div>	BEGIN	BEGIN	BEGIN
	SELECT MAX(age) FROM sailors WHERE rating=1	72	INSERT INTO sailors (age=96, rating=1)
	SELECT MAX(age) FROM sailors WHERE rating=1	96	COMMIT
	COMMIT	COMMIT	

Faloutsos/Pavlo CMU SCS 15-415/615 39

CMU SCS

How did this happen?

- Because T1 locked only existing records and not ones under way!
- Conflict serializability on reads and writes of individual items guarantees serializability only if the set of objects is fixed.
- Solution?

Faloutsos/Pavlo CMU SCS 15-415/615 40

CMU SCS

Predicate Locking

- Lock records that satisfy a logical predicate:
 - Example: **rating=1**.
- In general, predicate locking has a lot of locking overhead.
- **Index locking** is a special case of predicate locking that is potentially more efficient.

Faloutsos/Pavlo CMU SCS 15-415/615 41

CMU SCS

Index Locking

- If there is a dense index on the **rating** field then the txn can lock index page containing the data with **rating=1**.
- If there are no records with **rating=1**, the txn must lock the index page where such a data entry would be, if it existed.

Faloutsos/Pavlo CMU SCS 15-415/615 42

CMU SCS

Locking without an Index

- If there is no suitable index, then the txn must obtain:
 - A lock on every page in the table to prevent a record's **rating** from being changed to 1.
 - The lock for the table itself to prevent records with **rating=1** from being added or deleted.

Faloutsos/Pavlo CMU SCS 15-415/615 43

CMU SCS

Phantom Problem

Faloutsos/Pavlo CMU SCS 15-415/615 44

CMU SCS

Today's Class

- Lock Granularities
- Locking in B+Trees
- The Phantom Problem
- ➔ • Weaker Levels of Consistency

Faloutsos/Pavlo CMU SCS 15-415/615 45

CMU SCS

Weaker Levels of Consistency

- Serializability is useful because it allows programmers to ignore concurrency issues.
- But enforcing it may allow too little concurrency and limit performance.
- We may want to use a weaker level of consistency to improve scalability.

Faloutsos/Pavlo CMU SCS 15-415/615 46

CMU SCS

Isolation Levels

- Controls the extent that a txn is exposed to the actions of other concurrent txns.
- Provides for greater concurrency at the cost of exposing txns to uncommitted changes:
 - Dirty Reads
 - Unrepeatable Reads
 - Phantom Reads

Faloutsos/Pavlo CMU SCS 15-415/615 47

CMU SCS

Isolation Levels

Isolation (High—Low)

- **SERIALIZABLE:** No phantoms, all reads repeatable, no dirty reads.
- **REPEATABLE READS:** Phantoms may happen.
- **READ COMMITTED:** Phantoms and unrepeatable reads may happen.
- **READ UNCOMMITTED:** All of them may happen.

Faloutsos/Pavlo CMU SCS 15-415/615 48

CMU SCS

Isolation Levels

	Dirty Read	Unrepeatable Read	Phantom
READ UNCOMMITTED	Maybe	Maybe	Maybe
READ COMMITTED	No	Maybe	Maybe
REPEATABLE READ	No	No	Maybe
SERIALIZABLE	No	No	No

Faloutsos/Pavlo CMU SCS 15-415/615 49

- CMU SCS
- ### Isolation Levels
- **SERIALIZABLE:** Obtain all locks first; plus index locks, plus strict 2PL.
 - **REPEATABLE READS:** Same as above, but no index locks.
 - **READ COMMITTED:** Same as above, but **S** locks are released immediately.
 - **READ UNCOMMITTED:** Same as above, but allows dirty reads (no **S** locks).
- Faloutsos/Pavlo CMU SCS 15-415/615 50

CMU SCS

SQL-92 Isolation Levels

SET TRANSACTION ISOLATION LEVEL
<isolation-level>;

- Default: **SERIALIZABLE**
- Not all DBMS support all isolation levels in all execution scenarios (e.g., replication).

Faloutsos/Pavlo CMU SCS 15-415/615 51

CMU SCS

Access Modes

- You can also provide hints to the DBMS about whether a txn will modify the database.
- Only two possible modes:
 - **READ WRITE**
 - **READ ONLY**

Faloutsos/Pavlo CMU SCS 15-415/615 52

CMU SCS

SQL-92 Access Modes

SQL-92
`SET TRANSACTION <access-mode>;`

Postgres + MySQL 5.6
`START TRANSACTION <access-mode>;`

- Default: **READ WRITE**
- Not all DBMSs will optimize execution if you set a txn to in **READ ONLY** mode.

Faloutsos/Pavlo CMU SCS 15-415/615 53

CMU SCS

Transaction Demo

Faloutsos/Pavlo CMU SCS 15-415/615 54

CMU SCS

Summary

- Multiple granularity locking: leads to few locks, at appropriate levels
- Tree-structured indexes:
 - Lock crabbing and safe nodes
- Important distinction:
 - Multiple granularity locking releases locks bottom-up.
 - Tree-locking releases top-down to maximize concurrency.

Faloutsos/Pavlo CMU SCS 15-415/615 55

CMU SCS

Summary

- The Phantom Problem occurs if insertions/deletions
- Use Predicate locking to prevent this:
 - Index Locking
 - Table Locking

Faloutsos/Pavlo CMU SCS 15-415/615 56
