

CMUSCS

Carnegie Mellon Univ.  
Dept. of Computer Science  
15-415 - Database Applications

Concurrency Control  
Part 2 (R&G ch. 17)

Faloutsos SCS 15-415 #1

CMUSCS

## Outline

- conflict/view serializability
- Two-phase locking (2PL); strict 2PL (== 2PL-C, for 'Commit')
- deadlocks prevention & detection
- Locking granularity
- Tree locking protocols
- Phantoms & predicate locking

Faloutsos SCS 15-415 #2

CMUSCS

## Review questions

- conflict serializability?
- 2PL theorem?
- what is strict 2PL? why do we need it?
  - 'dirty read'?
  - cascading aborts?
- who generates the lock requests?

Faloutsos SCS 15-415 #3

CMUSCS

## Not in book: 'Lost update' problem

	T1	T2
time	Read(N)	Read(N)
↓	N=N-1	N=N-1
	Write(N)	Write(N)

Faloutsos SCS 15-415 #4

CMUSCS

## Major conclusions so far:

- **(strict) 2PL: extremely popular**
- Deadlock may still happen
  - detection: wait-for graph
  - prevention: abort some xacts, defensively
- philosophically: concurrency control uses:
  - locks
  - and aborts

Faloutsos SCS 15-415 #5

CMUSCS

## Outline

- conflict/view serializability
- Two-phase locking (2PL); strict 2PL (== 2PL-C, for 'Commit')
- deadlocks prevention & detection
- ➔ Locking granularity
- Tree locking protocols
- Phantoms & predicate locking

Faloutsos SCS 15-415 #6

CMU SCS

## Lock granularity?

- lock granularity
  - field? record? page? table?
- Pros and cons?
- (Ideally, each transaction should obtain a few locks)

Faloutsos


SCS 15-415

#7

CMU SCS

## Multiple granularity

- Eg:



Faloutsos

SCS 15-415

#8

CMU SCS

## What would you do?

- T1: read Smith's salary,
- while T2: give 10% raise to everybody
- what locks should they obtain?



Faloutsos

SCS 15-415

#9

CMU SCS

## What types of locks?

- X/S locks for leaf level +
- 'intent' locks, for higher levels

Faloutsos

SCS 15-415

#10

CMU SCS

## What types of locks?

- X/S locks for leaf level +
- 'intent' locks, for higher levels
- IS: intent to obtain S-lock underneath
- IX: intent .... X-lock ...
- S: shared lock for this level
- X: ex- lock for this level
- SIX: shared lock here; + IX

Faloutsos

SCS 15-415

#11

CMU SCS

## Protocol

- each xact obtains appropriate lock at highest level
- proceeds to desirable lower levels

Faloutsos

SCS 15-415

#12

CMU SCS

### Compatibility matrix

T2 wants T1 has	IS	IX	S	SIX	X
IS					
IX					
S					
SIX					
X					

Faloutsos SCS 15-415 #13

CMU SCS

### Compatibility matrix

T2 wants T1 has	IS	IX	S	SIX	X
IS	Y	Y	Y	Y	N
IX					
S					
SIX					
X					

Faloutsos SCS 15-415 #14

CMU SCS

### Compatibility matrix

T2 wants T1 has	IS	IX	S	SIX	X
IS	Y	Y	Y	Y	N
IX		Y	N	N	N
S					
SIX					
X					

Faloutsos SCS 15-415 #15

CMU SCS

### Compatibility matrix

T2 wants T1 has	IS	IX	S	SIX	X
IS	Y	Y	Y	Y	N
IX		Y	N	N	N
S			Y	N	N
SIX					
X					

Faloutsos SCS 15-415 #16

CMU SCS

### Compatibility matrix

T2 wants T1 has	IS	IX	S	SIX	X
IS	Y	Y	Y	Y	N
IX		Y	N	N	N
S			Y	N	N
SIX				N	N
X					N

Faloutsos SCS 15-415 #17

CMU SCS

### Multiple Granularity Lock Protocol

- Each Xact: lock root.
- To get S or IS lock on a node, must hold **at least** IS on parent node.
  - What if Xact holds SIX on parent? S on parent?
- To get X or IX or SIX on a node, must hold **at least** IX on parent node.
- Must release locks in bottom-up order.

Faloutsos SCS 15-415 #18

CMU SCS

## Multiple granularity protocol

stronger  
(more privileges)

weaker

Faloutsos SCS 15-415 #19

CMU SCS

## Examples – 2 level hierarchy

- T1 scans R, and updates a few tuples:

Tables  
|  
Tuples

Faloutsos SCS 15-415 #20

CMU SCS

## Examples – 2 level hierarchy

- T1 scans R, and updates a few tuples:
- T1 gets an SIX lock on R, then get X lock on tuples that are updated.

Faloutsos SCS 15-415 #21

CMU SCS

## Examples – 2 level hierarchy

- T2: find avg salary of ‘Sales’ employees

Faloutsos SCS 15-415 #22

CMU SCS

## Examples – 2 level hierarchy

- T2: find avg salary of ‘Sales’ employees
- T2 gets an IS lock on R, and repeatedly gets an S lock on tuples of R.

Faloutsos SCS 15-415 #23

CMU SCS

## Examples – 2 level hierarchy

- T3: sum of salaries of everybody in ‘R’:

Faloutsos SCS 15-415 #24

CMU/SCS

## Examples – 2 level hierarchy

- T3: sum of salaries of everybody in ‘R’:
- T3 gets an S lock on R.
- OR, T3 could behave like T2; can use **lock escalation** to decide which.
  - Lock escalation dynamically asks for coarser-grained locks when too many low level locks acquired

Faloutsos SCS 15-415 #25

CMU/SCS

## Multiple granularity

- Very useful in practice
- each xact needs only a few locks

Faloutsos SCS 15-415 #26

CMU/SCS

## Outline

- ...
- Locking granularity
- ➔ Tree locking protocols
- Phantoms & predicate locking

Faloutsos SCS 15-415 #27

CMU/SCS

## Locking in B+ Trees

- What about locking indexes?

Faloutsos SCS 15-415 #28

CMU/SCS

## Example B+tree

- T1 wants to insert in H
- T2 wants to insert in I
- why not plain 2PL?

Faloutsos SCS 15-415 #29

CMU/SCS

## Example B+tree

- T1 wants to insert in H
- T2 wants to insert in I
- why not plain 2PL?
- Because: X/S locks for too long!

Faloutsos SCS 15-415 #30

CMU SCS

## Two main ideas:

- ‘crabbing’: get lock for parent; get lock for child; release lock for parent (if ‘safe’)
- ‘safe’ nodes == nodes that won’t split or merge, ie:
  - not full (on insertion)
  - more than half-full (on deletion)

Faloutsos SCS 15-415 #31

CMU SCS

## Example B+tree

- T1 wants to insert in H
- crabbing:

Faloutsos SCS 15-415 #32

CMU SCS

## Example B+tree

- T1 wants to insert in H

Faloutsos SCS 15-415 #33

CMU SCS

## Example B+tree

- T1 wants to insert in H
- (if ‘B’ is ‘safe’)

Faloutsos SCS 15-415 #34

CMU SCS

## Example B+tree

- T1 wants to insert in H
- continue ‘crabbing’

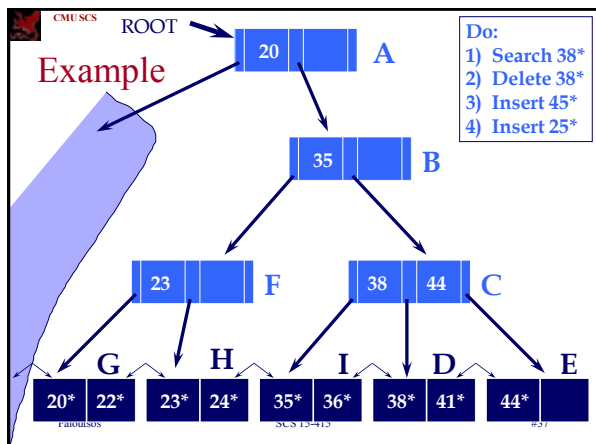
Faloutsos SCS 15-415 #35

CMU SCS

## A Simple Tree Locking Algorithm: “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 SCS 15-415 #36



- Answers:**
- Search 38\*  
- 'crabbing': S A, S B, U A, S C, U B, S D, U C
  - Delete 38\*  
- X A, X B, X C; U A, U B, X D, U C
  - Insert 45\*  
- X A, X B; U A, X C, X E, U C
  - Insert 25\*  
- X A, X B, U A, X F, U B, X H

- Answers:**
- Search 38\*  
- 'crabbing': S A, S B, U A, S C, U B, S D, U C
  - Delete 38\*  
- ~~X A~~ X B, X C; U A, U B, X D, U C
  - Insert 45\*  
- ~~X A~~ X B; U A, X C, X E, U C
  - Insert 25\*  
- ~~X A~~ X B, U A, X F, U B, X H
- CAN WE DO BETTER?**

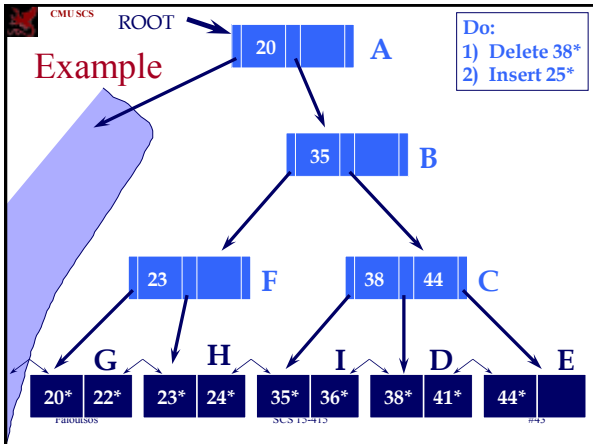
- Can we do better?**
- Yes [Bayer and Schkolnik]:
  - Idea: hope that the leaf is 'safe', and use S-locks & crabbing to reach it, and verify
  - (if false, do previous algo)

**Can we do better?**

- Yes [Bayer and Schkolnik]:

**Rudolf Bayer, Mario Schkolnik: Concurrency of Operations on B-Trees. Acta Inf. 9: 1-21 (1977)**

- advanced**
- A Better Tree Locking Algorithm (From Bayer-Schkolnick paper)**
- Search:** 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 Xact 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. In practice, better than previous alg.



**Answers:**

- Delete 38\*
  - SA, SB, UA, SC, UB, XD, UC
- Insert 25\*
  - SA, SB, UA, SF, UB, XH; UH;
  - XA, XB, UA, XF, UB, XH

**Notice:**

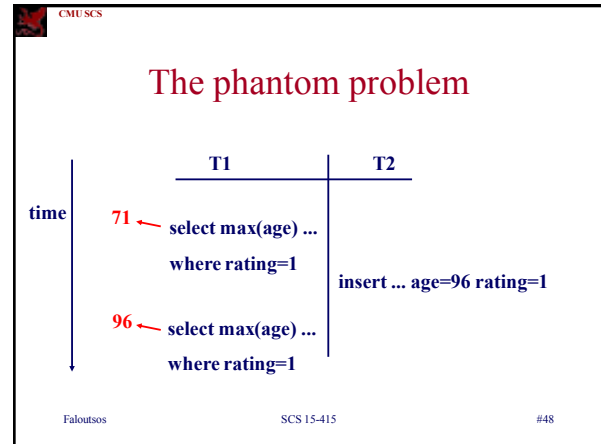
- Textbook has a third variation, that uses lock-upgrades (and may lead to deadlocks)

**Outline**

- Locking granularity
- Tree locking protocols
- ➔ Phantoms & predicate locking

**Dynamic Databases – The “Phantom” Problem**

- so far: only reads and updates – no insertions/deletions
- with insertions/deletions, new problems:





CMU SCS

## Why?

- because T1 locked only \*existing\* records  
– not ones under way!
- Solution?

Faloutsos SCS 15-415 #49

CMU SCS

## Solution

theoretical solution:

- ‘predicate locking’: e.g., lock all records (current or incoming) with rating=1  
– VERY EXPENSIVE

Faloutsos SCS 15-415 #50

CMU SCS

## Solution

practical solution:

- index locking: if an index (on ‘rating’) exists, lock the appropriate entries (rating=1 in our case)
- otherwise, lock whole table (and thus block insertions/deletions)

Faloutsos SCS 15-415 #51

CMU SCS

## Transaction Support in SQL-92

*recommended* **SERIALIZABLE** – No phantoms, all reads repeatable, no “dirty” (uncommitted) reads.

- REPEATABLE READS – phantoms may happen.
- READ COMMITTED – phantoms and unrepeatable reads may happen
- READ UNCOMMITTED – all of them may happen.

Faloutsos SCS 15-415 #52

CMU SCS

## Transaction Support in SQL-92

- **SERIALIZABLE** : obtains all locks first; plus index locks, plus strict 2PL
- **REPEATABLE READS** – as above, but no index locks
- **READ COMMITTED** – as above, but S-locks are released immediately
- **READ UNCOMMITTED** – as above, but allowing ‘dirty reads’ (no S-locks)

Faloutsos SCS 15-415 #53

CMU SCS

## Transaction Support in SQL-92

SET TRANSACTION ISOLATION LEVEL  
SERIALIZABLE READ ONLY

Defaults:

SERIALIZABLE	←	isolation level
READ WRITE	←	access mode

Faloutsos SCS 15-415 #54

CMU SCS

## Summary

- Multiple granularity locking: leads to few locks, at appropriate levels
- Tree-structured indexes:
  - ‘crabbing’ and ‘safe nodes’

Faloutsos SCS 15-415 #55

CMU SCS

## Summary

- “phantom problem”, if insertions/deletions
  - (Predicate locking prevents phantoms)
  - Index locking, or table locking

Faloutsos SCS 15-415 #56