

CS 245: Database System Principles

Notes 10: More TP

Hector Garcia-Molina
(Some modifications by Chris Olston)

CS 245

Notes 10

1

Sections to Skim:

- Chapter 8:
 - skim 8.5
- Chapter 9:
 - skim 9.6, 9.8
- Chapter 10:
 - skim 10.4, 10.5, 10.6
- Chapter 11: none (read all sections)

CS 245

Notes 10

2

Chapter 10 More on transaction processing

Topics:

- Cascading rollback, recoverable schedule
- Deadlocks
 - Prevention
 - Detection
- View serializability
- Distributed transactions
- Long transactions (nested, compensation)

CS 245

Notes 10

3

Concurrency control & recovery

Example:

T_j	T_i
\vdots	\vdots
$W_j(A)$	\vdots
\vdots	$r_i(A)$
\vdots	Commit T_i
\vdots	\vdots
Abort T_j	\vdots

➡ Need to "uncommit" T_i (Bad!)

CS 245

Notes 10

4

- Schedule is conflict serializable
- $T_j \rightarrow T_i$
- But not recoverable
 - I.e., not always possible to abort a transaction without affecting the status of some transaction that already committed

CS 245

Notes 10

5

- Need to make "final" decision for each transaction:
 - **commit decision** - system guarantees transaction will or has completed, no matter what
 - **abort decision** - system guarantees transaction will or has been rolled back (has no effect)

CS 245

Notes 10

6

To model this, two new actions:


- C_i - transaction T_i commits
- A_i - transaction T_i aborts

Note: in transactions, reads and writes precede commit or abort

- \Rightarrow If $C_i \in T_i$, then $r_i(A) < C_i$
 $w_i(A) < C_i$
- \Rightarrow If $A_i \in T_i$, then $r_i(A) < A_i$
 $w_i(A) < A_i$

- Also, one of C_i, A_i per transaction

Back to example:



Definition

T_i reads from T_j in S ($T_j \Rightarrow_S T_i$) if


- (1) $w_j(A) <_S r_i(A)$
- (2) $a_j \not<_S r_i(A)$ ($\not<$: does not precede)
- (3) If $w_j(A) <_S w_k(A) <_S r_i(A)$ then $a_k <_S r_i(A)$

Definition


Schedule S is recoverable if whenever $T_j \Rightarrow_S T_i$ and $j \neq i$ and $C_i \in S$ then $C_j <_S C_i$

- S is recoverable if each transaction *commits* only after all transactions from which it read have committed.
- S avoids cascading rollback if each transaction may *read* only those values written by committed transactions.

- S is strict if each transaction may *read and write* only items previously written by committed transactions.



Where are serializable schedules?




Examples

- Recoverable:
 - $w_1(A) w_1(B) w_2(A) r_2(B) c_1 c_2$
- Avoids Cascading Rollback:
 - $w_1(A) w_1(B) w_2(A) c_1 r_2(B) c_2$ Assumes $w_2(A)$ is done without reading
- Strict:
 - $w_1(A) w_1(B) c_1 w_2(A) r_2(B) c_2$

How to achieve schedules that avoid cascading rollback (and are thus also recoverable) ?

⇒ With 2PL, hold write locks until after commit (strict 2PL)



⇒ With validation, no change!

Deadlocks

- Detection
 - Timeout
 - Wait-for graph
- Prevention
 - Resource ordering
 - Wait-die
 - Wound-wait

CS 245

Notes 10

19

Timeout (Guess...)

- If transaction waits more than L sec., roll it back!
- Simple scheme
- Hard to select L
 - May wait too long before rolling back deadlocked transactions
 - May rollback transactions that were not deadlocked!


CS 245

Notes 10

20

Real Deadlock Detection

- Build Wait-For graph
- Use lock table structures
- Build incrementally or periodically
- When cycle found, rollback victim



CS 245

Notes 10

21

Deadlock Prevention ...

CS 245

Notes 10

22

Prevention(1): Resource Ordering

- Order all elements A_1, A_2, \dots, A_n
- A transaction T can lock A_i after A_j only if $i > j$

Problem : Ordered lock requests not realistic in most cases

CS 245

Notes 10

23

Prevention(2): Wait-die


- Transactions given a timestamp when they arrive $ts(T_i)$
- T_i can only wait for T_j if $ts(T_i) < ts(T_j)$...else die

CS 245

Notes 10

24

Example:



Prevention(3): Wound-wait

- Transactions given a timestamp when they arrive ... $ts(T_i)$
- T_i wounds T_j if $ts(T_i) < ts(T_j)$
else T_i waits

“Wound”: T_j rolls back and gives lock to T_i


Locking

Locking

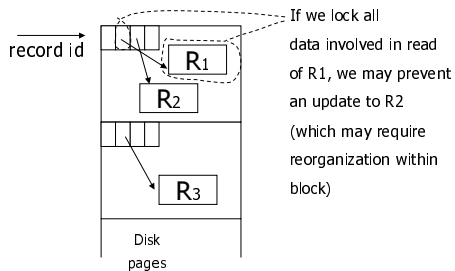
What are we really locking?



Example:



But underneath:



Solution: view DB at two levels

Top level: record actions
record locks
undo/redo actions — logical

e.g., Insert record(X,Y,Z)
Redo: insert(X,Y,Z)
Undo: delete


Low level: deal with physical details
latch page during action
(release at end of action)

CS 245

Notes 10

31

Note: undo does not return physical DB
to original state; only same logical state



CS 245

Notes 10

32

We've looked at the low level guts of
locking and concurrency control

Let's pop our heads out and look at things
from the point of view of users

(Always a good exercise ...)



CS 245

Notes 10

33

User/Program commands

Lots of variations, but in general

- Begin_work
- Commit_work
- Rollback_work (Abort_work)

CS 245

Notes 10

34

Transactions

User program (or SQL prompt):

⋮

Begin_work;

⋮

⋮

If results_ok, then commit work
else abort_work

CS 245

Notes 10

35

Problem with Regular Transactions

Granularity of commit/abort is too coarse

Desirable behavior: see video games ...

<save> <try to slay dragon> <save>
<try to jump abyss> <save> ...

CS 245

Notes 10

36

Nested transactions

User program:

```
⋮
Begin_work;
  Begin_work;
  ⋮
  If results_ok, then commit work
  else {abort_work; try something else...}
⋮
If results_ok, then commit work
else abort_work
```


CS 245

Notes 10

37

Parallel Nested Transactions

```
T1: begin_work
⋮
parallel:
T11: begin_work
⋮
commit_work
T12: begin_work
⋮
commit_work
⋮
commit_work
```



CS 245

Notes 10

38

Another Problem with Transactions

- Some transactions take a long time!
 - data analysis / mining
 - CAD systems
 - workflow
- Want to avoid locking out other transactions altogether
- At same time, avoid complete anarchy

CS 245

Notes 10

39

Idea #1: Different Levels of Isolation

- For some long-running analysis or mining queries, can tolerate some inconsistencies
- For such transactions, do not enforce 2PL (i.e., permit unlock then lock)
- Permits other transactions to modify the data while the analysis query is running
- Lose guarantee of "repeatable read"

CS 245

Notes 10

40

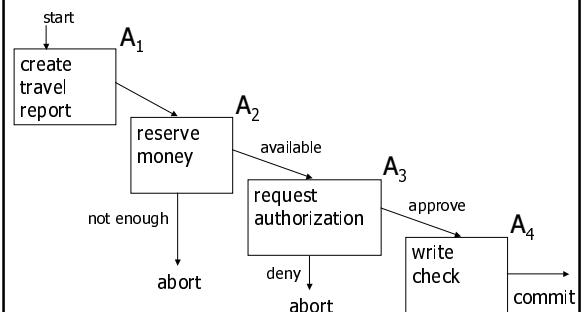
What about long-running transactions that modify the database?

CS 245

Notes 10

41

Workflow example: expense report



CS 245

Notes 10

42

Compensating Transactions

- For every action A , an action A^{-1} that "undoes" the effects of A , regardless of intervening actions
- Formally: $A\alpha A^{-1} \equiv \alpha$
- A_1^{-1} = delete travel report
- A_2^{-1} = put money back in account
- A_3^{-1} = delete authorization form
- A_4^{-1} = cancel check (never necessary)

CS 245

Notes 10

43

Semantic atomicity:

execute one of

- A_1, A_2, A_3, A_4
- $A_1, A_2, A_3, A_3^{-1}, A_2^{-1}, A_1^{-1}$
- $A_1, A_2, A_2^{-1}, A_1^{-1}$
- A_1, A_1^{-1}
- nothing

called a saga

CS 245

Notes 10

44

Careful!

- Compensating actions are tricky
- Remember requirement: $A\alpha A^{-1} \equiv \alpha$
- EX: start with an account worth \$2000
A = add \$1000; A^{-1} = deduct \$1000
C = deduct \$2500

Constraint: balance can't be negative

$ACA^{-1} \rightarrow A^{-1}$ fails
 $C \rightarrow C$ fails

$A\alpha A^{-1} \not\equiv \alpha$

CS 245

Notes 10

45

Summary

- Cascading rollback
Recoverable schedule
- Deadlock
 - Prevention
 - Detection
- Logical vs. physical locking
- Nested transactions
- Long running transactions

CS 245

Notes 10

46