Transactions; Concurrency; Recovery

Amol Deshpande CMSC424

Spring 2020 – Online Instruction Plan

- Week 1: File Organization and Indexes
- Week 2: Query Processing
- Week 3: Query Optimization; Parallel Databases 1
- Week 4: Parallel Databases; Mapreduce; Transactions 1
- Week 5: Transactions 2 (Homework Due May 1)
 - Transactions: Serializability, Recoverability
 - Transactions: Concurrency 1
 - Transactions: Concurrency 2: Other Concurrency Schemes
 - Transactions: Recovery
- Week 6: Distributed Transactions; Miscellaneous Topics (Homework Due May 8)

Transactions: Overview

- Book Chapters
 - *****14.6*,* 14.7
- Key topics:
 - ★ Conflict equivalence of schedules
 - Conflict serializability and checking by drawing precedence graphs
 - ★ View serializability
 - ★ Recoverability

Another schedule

_T1	T2	
read(A) A = A -50 write(A)		Is this schedule okay ?
	read(A) tmp = A*0.1 A = A tmp	Lets look at the final effect
	write(A)	Effect: <u>Before</u> <u>After</u> A 100 45
read(B) B=B+50 write(B)		B 50 105
	read(B) B = B+ tmp	Further, the effect same as the serial schedule 1.
	write(B)	Called <u>serializable</u>

Conflict Serializability

- Two read/write instructions "conflict" if
 - ★ They are by different transactions
 - ★ They operate on the same data item
 - ★ At least one is a "write" instruction
- Why do we care ?
 - If two read/write instructions don't conflict, they can be "swapped" without any change in the final effect
 - However, if they conflict they CAN'T be swapped without changing the final effect

Equivalence by Swapping

T1	T2	T1	T2
read(A)		read(A)	
A = A -50		A = A -50	
write(A)		write(A)	
	read(A)		read(A)
	$tmp = A^*0.1$		$tmp = A^*0.1$
	A = A - tmp		A = A - tmp
	write(A)		
		read(B)	
read(B)			write(A)
B=B+50		B=B+50	
write(B)		write(B)	
	read(B)		read(B)
	B = B + tmp		B = B + tmp
	write(B)		write(B)
Effect: Bef	ore After	Effect	Before After
A 1	00 45	A	100 45
R P	50 105	== /\ B	50 105

Equivalence by Swapping

T1	T2		T2
read(A)		read(A)	
A = A -50		A = A -50	
write(A)		write(A)	
	read(A)		read(A)
	tmp = $A^*0.1$		$tmp = A^*0.1$
	A = A - tmp		A = A - tmp
	write(A)		write(A)
read(B)		read(B)	
B=B+50		B=B+50	
write(B)			read(B)
	read(B)	write(B)	
	B = B + tmp		B = B + tmp
	write(B)		write(B)
Effect: Bef	ore After	Effect: B	efore After
A 10	$\frac{0.0}{10}$ $\frac{7.00}{45}$	<u>спост</u> <u>с</u>	100 45
B F	50 105	!== // B	50 55

Conflict Serializability

- Conflict-equivalent schedules:
 - If S can be transformed into S' through a series of swaps, S and S' are called conflict-equivalent
 - * conflict-equivalent guarantees same final effect on the database
- A schedule S is conflict-serializable if it is conflict-equivalent to a serial schedule

Equivalence by Swapping

T1	T2	T1	T2
read(A)		read(A)	
A = A -50		A = A -50	
write(A)		write(A)	
	read(A)		read(A)
	$tmp = A^*0.1$		tmp = $\dot{A}^{*}0.1$
	A = A - tmp		A = A - tmp
	write(A)		
		read(B)	
read(B)		B=B+50	
B=B+50			write(A)
write(B)		write(B)	WIIIC(A)
	read(B)		read(B)
	$B = B \perp tmn$		$B = B \downarrow tmp$
	D = D + (III)		D = D + (III)
	WIILE(D)		WITE(D)
Effect: <u>Bef</u>	<u>ore</u> <u>After</u>	Effect	<u>Before</u> <u>After</u>
A 1	00 45	==	A 100 45
B	50 105		B 50 105

Equivalence by Swapping

T1	T2	T1	T2
read(A)		read(A)	
A = A -50		A = A -50	
write(A)		write(A)	
	read(A)		
	$tmp = A^*0.1$	read(B)	
	A = A - tmp	B=B+50	
	write(A)	write(B)	
			read(A)
read(B)			$tmp = A^*0.1$
B=B+50			A = A - tmp
write(B)			write(A)
	read(B)		read(B)
	B - B + tmn		$B = B_{\perp} tmp$
	write(R)		$B = B + timpwrite(B)$
Effect: Bet	<u>ore</u> <u>Atter</u>	Effect:	Betore <u>Atter</u>
A 1	00 45	A	100 45
B	50 105	В	50 105

Example Schedules (Cont.)

A "bad" schedule



Serializability

In essence, following set of instructions is not conflict-serializable:

T_3	T_4
read(Q)	
	write (Q)
write (Q)	

View-Serializability

Similarly, following not conflict-serializable

T_3	T_4	T_6
read(Q)		
write(Q)	write(Q)	
		write (Q)

- BUT, it is serializable
 - ★ Intuitively, this is because the *conflicting write instructions* don't matter
 - ★ The final write is the only one that matters
- View-serializability allows these
 - ★ Read up

Other notions of serializability

T_1	T_5
read(A)	
A := A - 50	
write(A)	
	read(B)
	B := B - 10
	write(B)
read(B)	
B := B + 50	
write(B)	
	read(A)
	A := A + 10
	write(A)

- Not conflict-serializable or view-serializable, but serializable
- Mainly because of the +/- only operations
 - Requires analysis of the actual operations, not just read/write operations
- Most high-performance transaction systems will allow these

Testing for conflict-serializability

Given a schedule, determine if it is conflict-serializable

- Draw a precedence-graph over the transactions
 - A directed edge from T1 and T2, if they have conflicting instructions, and T1's conflicting instruction comes first
- If there is a cycle in the graph, not conflict-serializable
 - Can be checked in at most O(n+e) time, where n is the number of vertices, and e is the number of edges
- If there is none, conflict-serializable
- Testing for view-serializability is NP-hard.

Example Schedule (Schedule A) + Precedence Graph



Recap so far...

We discussed:

- ★ Serial schedules, serializability
- ★ Conflict-serializability, view-serializability
- ★ How to check for conflict-serializability
- We haven't discussed:
 - ★ How to guarantee serializability ?
 - Allowing transactions to run, and then aborting them if the schedules wasn't serializable is clearly not the way to go
 - We instead use schemes to guarantee that the schedule will be conflict-serializable

★ Also, *recoverability* ?

Recoverability

Serializability is good for

consistency

	T1	T2
But what if transactions fail ?	read(A)	
T2 has already committed	A = A - 50	
A user might have been notified	write(A)	read(A)
Now T1 abort creates a problem		$tmp = A^*0.1$
T2 has seen its effect, so just aborting T1 is not enough. T2 must be aborted as well (and possibly restarted)		A = A – tmp write(A) COMMIT
➢ But T2 is committed	read(B) B=B+50 write(B) ABORT	

Recoverability

Recoverable schedule: If T1 has read something T2 has written, T2 must commit before T1

★ Otherwise, if T1 commits, and T2 aborts, we have a problem

Cascading rollbacks: If T10 aborts, T11 must abort, and hence T12 must abort and so on.

T_{10}	T_{11}	<i>T</i> ₁₂
read(A)		
read(B)		
write(A)		
	read(A)	
	write(A)	
		read(A)

Recoverability

- Dirty read: Reading a value written by a transaction that hasn't committed yet
- Cascadeless schedules:
 - ★ A transaction only reads *committed* values.
 - ★ So if T1 has written A, but not committed it, T2 can't read it.
 - > No dirty reads
- Cascadeless \rightarrow No cascading rollbacks
 - ★ That's good
 - ★ We will try to guarantee that as well

Recap so far...

We discussed:

- ★ Serial schedules, serializability
- ★ Conflict-serializability, view-serializability
- ★ How to check for conflict-serializability
- ★ Recoverability, cascade-less schedules
- We haven't discussed:
 - ★ How to guarantee serializability ?
 - Allowing transactions to run, and then aborting them if the schedules wasn't serializable is clearly not the way to go
 - We instead use schemes to guarantee that the schedule will be conflict-serializable