

CSE 444: Database Internals

Lectures 15 and 16
Transactions: Optimistic
Concurrency Control

CSE 444 – Spring 2019

1

Announcements

- Quiz 1+2 Friday in class
 - Last quarter's quiz linked in calendar
 - 1 page (2 sides) of notes allowed
 - Special attention to your implementation:
 - Understand SimpleDB operator parameters
- Lab 3 part 1 due *Saturday* 11pm
 - Less times for labs in general

CSE 444 – Spring 2019

2

Pessimistic vs. Optimistic

- **Pessimistic CC** (locking)
 - Prevents unserializable schedules
 - Never abort for serializability (but may abort for deadlocks)
 - Best for workloads with high levels of contention
- **Optimistic CC** (timestamp, multi-version, validation)
 - Assume schedule will be serializable
 - Abort when conflicts detected
 - Best for workloads with low levels of contention

CSE 444 – Spring 2019

3

Outline

- **Concurrency control by timestamps (18.8)**
- Concurrency control by validation (18.9)
- Snapshot Isolation

CSE 444 – Spring 2019

4

Timestamps

- Each transaction receives unique timestamp **TS(T)**

Could be:

- The system's clock
- A unique counter, incremented by the scheduler

CSE 444 – Spring 2019

5

Timestamps

Main invariant:

The timestamp order defines
the serialization order of the transaction

Will generate a schedule that is view-equivalent
to a serial schedule, and recoverable

CSE 444 – Spring 2019

6

Timestamps

With each element X , associate

- $RT(X)$ = the highest timestamp of any transaction U that read X
- $WT(X)$ = the highest timestamp of any transaction U that wrote X
- $C(X)$ = the commit bit: true when transaction with highest timestamp that wrote X committed

CSE 444 – Spring 2019

7

Main Idea

For any $r_T(X)$ or $w_T(X)$ request, check for conflicts:

- $w_U(X) \dots r_T(X)$
- $r_U(X) \dots w_T(X)$
- $w_U(X) \dots w_T(X)$

How do we check if Read too late?

Write too late?

CSE 444 – Spring 2019

8

Main Idea

For any $r_T(X)$ or $w_T(X)$ request, check for conflicts:

- $w_U(X) \dots r_T(X)$
- $r_U(X) \dots w_T(X)$
- $w_U(X) \dots w_T(X)$

How do we check if Read too late?

Write too late?

When T requests $r_T(X)$, need to check $TS(U) \leq TS(T)$

CSE 444 – Spring 2019

9

Read Too Late

- T wants to read X

START(T) ... START(U) ... $w_U(X) \dots r_T(X)$

CSE 444 – Spring 2019

10

Read Too Late

- T wants to read X

START(T) ... START(U) ... $w_U(X) \dots r_T(X)$

If $WT(X) > TS(T)$ then need to rollback T !

CSE 444 – Spring 2019

11

Write Too Late

- T wants to write X


START(T) ... START(U) ... $r_U(X) \dots w_T(X)$

CSE 444 – Spring 2019

12

Write Too Late

- T wants to write X



If $RT(X) > TS(T)$ then need to rollback T !


CSE 444 – Spring 2019

13

Thomas' Rule

But we can still handle it:

- T wants to write X



If $RT(X) \leq TS(T)$ and $WT(X) > TS(T)$
then don't write X at all !


Why does this work?

14

Thomas' Rule

But we can still handle it:

- T wants to write X



If $RT(X) \leq TS(T)$ and $WT(X) > TS(T)$
then don't write X at all !

Why does this work?

View-serializable
schedule

View-Serializability

- By using Thomas' rule we do obtain a view-serializable schedule

CSE 444 – Spring 2019

16

Summary So Far

Only for transactions that do not abort
Otherwise, may result in non-recoverable schedule

Transaction wants to READ element X
If $WT(X) > TS(T)$ then ROLLBACK
Else READ and update $RT(X)$ to larger of $TS(T)$ or $RT(X)$

Transaction wants to WRITE element X
If $RT(X) > TS(T)$ then ROLLBACK
Else if $WT(X) > TS(T)$ ignore write & continue (Thomas Write Rule)
Otherwise, WRITE and update $WT(X) = TS(T)$

CSE 444 – Spring 2019

17

Ensuring Recoverable Schedules

Recall:

- Schedule avoids cascading aborts if whenever a transaction reads an element, then the transaction that wrote it must have already committed
- Use the commit bit $C(X)$ to keep track if the transaction that last wrote X has committed


CSE 444 – Spring 2019

18

Ensuring Recoverable Schedules

Read dirty data:

- T wants to read X, and $WT(X) < TS(T)$
- Seems OK, but...



If $C(X)=\text{false}$, T needs to wait for it to become true


CSE 444 – Spring 2019

19

Ensuring Recoverable Schedules

Thomas' rule needs to be revised:

- T wants to write X, and $WT(X) > TS(T)$
- Seems OK not to write at all, but ...



If $C(X)=\text{false}$, T needs to wait for it to become true

CSE 444 – Spring 2019

20

Timestamp-based Scheduling

- When a transaction T requests $r_T(X)$ or $w_T(X)$, the scheduler examines $RT(X)$, $WT(X)$, $C(X)$, and decides one of:
 - To grant the request, or
 - To rollback T (and restart with later timestamp)
 - To delay T until $C(X) = \text{true}$

CSE 444 – Spring 2019

21

Timestamp-based Scheduling

RULES including commit bit

- There are 4 long rules in Sec. 18.8.4
- You should be able to derive them yourself, based on the previous slides
- Make sure you understand them !

READING ASSIGNMENT:
Garcia-Molina et al. 18.8.4

CSE 444 – Spring 2019

22

Timestamp-based Scheduling (Read 18.8.4 instead!)

Transaction wants to READ element X

If $WT(X) > TS(T)$ then ROLLBACK
Else if $C(X) = \text{false}$, then WAIT
Else READ and update $RT(X)$ to larger of $TS(T)$ or $RT(X)$

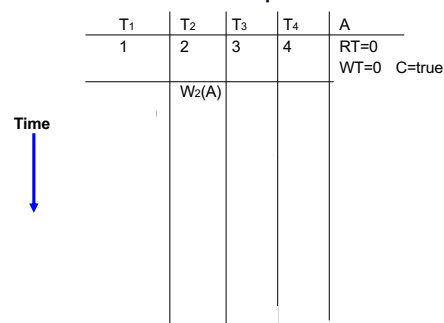
Transaction wants to WRITE element X

If $RT(X) > TS(T)$ then ROLLBACK
Else if $WT(X) > TS(T)$
Then If $C(X) = \text{false}$ then WAIT
else IGNORE write (Thomas Write Rule)
Otherwise, WRITE, and update $WT(X)=TS(T)$, $C(X)=\text{false}$

CSE 444 – Spring 2019

23

Basic Timestamps with Commit Bit



CSE 444 – Spring 2019

24

Basic Timestamps with Commit Bit

T ₁	T ₂	T ₃	T ₄	A
1	2	3	4	RT=0 WT=0 C=true
R ₁ (A) Abort	W ₂ (A) C	R ₃ (A) Delay		WT=2 C=false RT=0
		R ₃ (A)		C=true
		W ₄ (A)		RT=3 WT=4 C=false
		W ₃ (A) delay		
		W ₃ (A) abort		WT=2 C=true WT=3 C=false

CSE 444 – Spring 2019

25

Summary of Timestamp-based Scheduling

- View-serializable
- Avoids cascading aborts (hence: recoverable)
- Does NOT handle phantoms
 - These need to be handled separately, e.g. predicate locks

CSE 444 – Spring 2019

26

Multiversion Timestamp

- When transaction T requests r(X) but $WT(X) > TS(T)$, then T must rollback
- Idea: keep multiple versions of X: $X_t, X_{t-1}, X_{t-2}, \dots$

$$TS(X_t) > TS(X_{t-1}) > TS(X_{t-2}) > \dots$$

CSE 444 – Spring 2019

27

Details

- When $w_T(X)$ occurs, if the write is legal then create a **new version**, denoted X_t where $t = TS(T)$
 - When $r_T(X)$ occurs, find **most recent version** X_t such that $t \leq TS(T)$
- Notes:
- $WT(X_t) = t$ and it never changes
 - $RT(X_t)$ must still be maintained to check legality of writes
- Can delete X_t if we have a later version X_{t+1} and all active transactions T have $TS(T) > t$

CSE 444 – Spring 2019

28

Example w/ Basic Timestamps

T ₁	T ₂	T ₃	T ₄	A
Timestamps: 150	200	175	225	RT=0 WT=0
R ₁ (A) W ₁ (A)				RT=150 WT=150
	R ₂ (A) W ₂ (A)			RT=200 WT=200
		R ₃ (A) Abort		
			R ₄ (A)	RT=225

CSE 444 – Spring 2019

30

Example w/ Multiversion

T ₁	T ₂	T ₃	T ₄	A ₀	A ₁₅₀	A ₂₀₀
150	200	175	225			
R ₁ (A) W ₁ (A)				RT=150		
	R ₂ (A) W ₂ (A)				Create RT=200	
		R ₃ (A) W ₃ (A) abort			RT=200	Create
			R ₄ (A)			RT=225

CSE 444 – Spring 2019

31

Outline

- Concurrency control by timestamps (18.8)
- Concurrency control by validation (18.9)
- **Snapshot Isolation**
 - Not in the book, but good overview in Wikipedia

CSE 444 – Spring 2019

38

Snapshot Isolation

- A type of multiversion concurrency control algorithm
- Provides yet another level of isolation
- Very efficient, and very popular
 - Oracle, PostgreSQL, SQL Server 2005
- Prevents many classical anomalies BUT...
- Not serializable (!), yet ORACLE and PostgreSQL use it even for SERIALIZABLE transactions!
 - But "serializable snapshot isolation" now in PostgreSQL

CSE 444 – Spring 2019

39

Snapshot Isolation Overview

- Each transactions receives a timestamp $TS(T)$
- Transaction T sees snapshot at time $TS(T)$ of the database
- Write/write conflicts resolved by "first committer wins" rule
 - Loser gets aborted
- Read/write conflicts are ignored

CSE 444 – Spring 2019

40

Snapshot Isolation Details

- Multiversion concurrency control:
 - Versions of X: $X_{t1}, X_{t2}, X_{t3}, \dots$
- When T reads X, return $X_{TS(T)}$.
- When T writes X (to avoid lost update):
 - If latest version of X is $TS(T)$ then **proceed**
 - If $C(X) = \text{true}$ then **abort**
 - If $C(X) = \text{false}$ then **wait**
- When T commits, write its updates to disk

CSE 444 – Spring 2019

41

What Works and What Not

- No dirty reads (Why ?)
- No inconsistent reads (Why ?)
- No lost updates ("first committer wins")
- Moreover: no reads are ever delayed
- However: read-write conflicts not caught !

CSE 444 – Spring 2019

42

Write Skew

```
T1:
READ(X);
if X >= 50
  then Y = -50; WRITE(Y)
COMMIT
```

```
T2:
READ(Y);
if Y >= 50
  then X = -50; WRITE(X)
COMMIT
```

In our notation:

$R_1(X), R_2(Y), W_1(Y), W_2(X), C_1, C_2$

Starting with $X=50, Y=50$, we end with $X=-50, Y=-50$.
Non-serializable !!!

CSE 444 – Spring 2019

43

Write Skews Can Be Serious

- Acidicland had two viceroys, Delta and Rho
- Budget had two registers: taXes, and spendYng
- They had high taxes and low spending...

Delta: READ(taXes); if taXes = 'High' then { spendYng = 'Raise'; WRITE(spendYng) } COMMIT	Rho: READ(spendYng); if spendYng = 'Low' then { taXes = 'Cut'; WRITE(taXes) } COMMIT
--	---

... and they ran a deficit ever since.

CSE 444 – Spring 2019

44

Discussion: Tradeoffs

- **Pessimistic CC: Locks**
 - Great when there are many conflicts
 - Poor when there are few conflicts
- **Optimistic CC: Timestamps, Validation, SI**
 - Poor when there are many conflicts (rollbacks)
 - Great when there are few conflicts
- **Compromise**
 - READ ONLY transactions → timestamps
 - READ/WRITE transactions → locks

CSE 444 – Spring 2019

45

Commercial Systems

Always check documentation!

- **DB2**: Strict 2PL
- **SQL Server**:
 - Strict 2PL for standard 4 levels of isolation
 - Multiversion concurrency control for snapshot isolation
- **PostgreSQL**: SI; recently: serializable SI (!)
- **Oracle**: SI

CSE 444 – Spring 2019

46