

# CS 764: Topics in Database Management Systems Lecture 7: Optimistic Concurrency Control

Xiangyao Yu 9/28/2020

### Announcement

Guest lecture on Wednesday (Sep. 30) by Shasank Chavan from Oracle on "Hardware Acceleration with Oracle Database In-Memory"

Student round-table discussion after the talk (2:30-3:30)

### Today's Paper: Optimistic Concurrency Control

### On Optimistic Methods for Concurrency Control

H.T. KUNG and JOHN T. ROBINSON Carnegie-Mellon University

Most current approaches to concurrency control in database systems rely on locking of data objects as a control mechanism. In this paper, two families of nonlocking concurrency controls are presented. The methods used are "optimistic" in the sense that they rely mainly on transaction backup as a control mechanism, "hoping" that conflicts between transactions will not occur. Applications for which these methods should be more efficient than locking are discussed.

Key Words and Phrases: databases, concurrency controls, transaction processing CR Categories: 4.32, 4.33

#### 1. INTRODUCTION

Consider the problem of providing shared access to a database organized as a collection of objects. We assume that certain distinguished objects, called the roots, are always present and access to any object other than a root is gained only by first accessing a root and then following pointers to that object. Any sequence of accesses to the database that preserves the integrity constraints of the data is called a *transaction* (see, e.g., [4]).

If our goal is to maximize the throughput of accesses to the database, then there are at least two cases where highly concurrent access is desirable.

#### ACM Trans. Database Syst. 1981

# Agenda

Pessimistic concurrency control

Optimistic concurrency control

### **Concurrency Control**

**Concurrency control** ensures the <u>correctness</u> for concurrent operations

Assume **serializable** isolation level for this lecture

**Concurrency control** ensures the <u>correctness</u> for concurrent operations

Assume **serializable** isolation level for this lecture

**Pessimistic**: Resolve conflicts eagerly

**Optimistic**: Ignore conflicts during a transaction's execution and resolve conflicts lazily only when at a transaction's completion time

**Concurrency control** ensures the <u>correctness</u> for concurrent operations

Assume **serializable** isolation level for this lecture

**Pessimistic**: Resolve conflicts eagerly

**Optimistic**: Ignore conflicts during a transaction's execution and resolve conflicts lazily only when at a transaction's completion time

Other common concurrency control protocols

- Timestamp ordering (T/O)
- Multi-version concurrency control (MVCC)

Strict two-phase locking (2PL)

- Acquire the right type of locks before accessing data
- Release locks when the transaction commits

Strict two-phase locking (2PL)

Time

- Acquire the right type of locks before accessing data
- Release locks when the transaction commits



Strict two-phase locking (2PL)

- Acquire the right type of locks before accessing data
- Release locks when the transaction commits



Strict two-phase locking (2PL)

- Acquire the right type of locks before accessing data
- Release locks when the transaction commits



# Conflicts in 2PL



Solution 1: T2 waits for T1 to release lock (e.g., **wait-die**, **deadlock-detection**) Solution 2: T2 self aborts (e.g., **wait-die**, **no-wait**) Solution 3: T2 forces T1 to abort (e.g., **wound-wait**)

**T1** Begin Read(X) T2 Time ↓







Deadlock detection (DL\_DETECT)

 Maintain a wait-for graph among transactions; abort a transaction if a cycle is formed

Deadlock detection (DL\_DETECT)

 Maintain a wait-for graph among transactions; abort a transaction if a cycle is formed

NO\_WAIT

• The requesting transaction self aborts when a conflict occurs

#### Deadlock detection (DL\_DETECT)

 Maintain a wait-for graph among transactions; abort a transaction if a cycle is formed

NO\_WAIT

• The requesting transaction self aborts when a conflict occurs

#### WAIT\_DIE

 The requesting transaction waits if its priority is higher than the lock owner (wait), otherwise the requesting transaction self aborts (die)

#### Deadlock detection (DL\_DETECT)

 Maintain a wait-for graph among transactions; abort a transaction if a cycle is formed

NO\_WAIT

• The requesting transaction self aborts when a conflict occurs

#### WAIT\_DIE

• The requesting transaction waits if its priority is higher than the lock owner (**wait**), otherwise the requesting transaction self aborts (**die**)

#### WOUND\_WAIT

• The requesting transaction forces the lock owner to abort (**wound**) if its priority is higher than the lock owner, otherwise the requesting transaction waits (**wait**)

# Issues with Pessimistic CC

Overhead

- Overhead of acquiring/releasing locks and maintaining lock metadata
- Even read-only transactions acquire locks

Deadlocks

Limited concurrency

Locks are held till the end of a transaction

Real workloads have low contention

• Locking is unnecessary if no contention exists

# **Optimistic Concurrency Control (OCC)**

Goal: eliminating pessimistic locking

Three executing phases:

Read



Fig. 1. The three phases of a transaction.

*n* = *tcreate* 

tcreate = ( n := create;  $create set := create set \cup \{n\};$ **return** n)

n = tcreate twrite(n, i, v)

```
twrite(n, i, v) = (
if n \in create \ set
then write(n, i, v)
else if n \in write \ set
then write(copies[n], i, v)
else (
m := copy(n);
copies[n] := m;
write \ set := write \ set \cup \{n\};
write(copies[n], i, v)))
```

n = tcreate

```
twrite(n, i, v)
```

value = tread(n, i)

```
tread(n, i) = (
read set := read set \cup \{n\};
if n \in write set
then return read(copies[n], i)
else
return read(n, i))
```

n = tcreate twrite(n, i, v) value = tread(n, i) tdelete(n)

> tdelete(n) = ( $delete set := delete set \cup \{n\}).$

n = tcreate twrite(n, i, v) value = tread(n, i) tdelete(n)

All changes (i.e., inserts, updates, deletes) are kept local to the transaction without updating the database

### Write Phase

All written values become "global"

for  $n \in write set do exchange(n, copies[n])$ .

All created nodes become accessible All deleted nodes become inaccessible

A transaction *i* is assigned a transaction number *t(i)* when it enters the validation phase

• t(i) < t(j) = exists a serial schedule where  $T_i$  is before  $T_j$ 

- t(i) < t(j) => exists a serial schedule where  $T_i$  is before  $T_j$
- For t(i) < t(j), one of the following must be true
  - 1.  $T_i$  completes its write phase before  $T_i$  starts its read phase.
  - 2. The write set of  $T_i$  does not intersect the read set of  $T_j$ , and  $T_i$  completes its write phase before  $T_i$  starts its write phase.
  - 3. The write set of  $T_i$  does not intersect the read set or the write set of  $T_j$ , and  $T_i$  completes its read phase before  $T_i$  completes its read phase.

- t(i) < t(j) => exists a serial schedule where  $T_i$  is before  $T_j$
- For t(i) < t(j), one of the following must be true
  - 1.  $T_i$  completes its write phase before  $T_i$  starts its read phase.
  - 2. The write set of T<sub>i</sub> does not intersect the read set of T<sub>j</sub>, and T<sub>i</sub> completes its write phase before T<sub>i</sub> starts its write phase.
  - 3. The write set of  $T_i$  does not intersect the read set or the write set of  $T_j$ , and  $T_i$  completes its read phase before  $T_i$  completes its read phase.



- t(i) < t(j) => exists a serial schedule where  $T_i$  is before  $T_j$
- For t(i) < t(j), one of the following must be true
  - 1.  $T_i$  completes its write phase before  $T_i$  starts its read phase.
  - 2. The write set of  $T_i$  does not intersect the read set of  $T_j$ , and  $T_i$  completes its write phase before  $T_i$  starts its write phase.
  - 3. The write set of  $T_i$  does not intersect the read set or the write set of  $T_j$ , and  $T_i$  completes its read phase before  $T_i$  completes its read phase.



- t(i) < t(j) = exists a serial schedule where  $T_i$  is before  $T_j$
- For t(i) < t(j), one of the following must be true
  - 1.  $T_i$  completes its write phase before  $T_i$  starts its read phase.
  - 2. The write set of T<sub>i</sub> does not intersect the read set of T<sub>j</sub>, and T<sub>i</sub> completes its write phase before T<sub>i</sub> starts its write phase.
  - 3. The write set of  $T_i$  does not intersect the read set or the write set of  $T_j$ , and  $T_i$  completes its read phase before  $T_i$  completes its read phase.



# **Serial Validation**

tbegin = (	
$start \ tn := tnc$ ) tend = (	Critical Section
<pre>{finish tn := tnc; valid := true; for t from start tn + 1 to finish tn do</pre>	er t intersects read set) Which transactions will T2, T3,
	Problem: The entire validated
$I_2 \qquad F \qquad $	process happens in the critical section

# Improved Serial Validation

tend := (mid tn := tnc; valid := true: for t from start tn + 1 to mid th do **if** (*write set of transaction with transaction number t intersects read set*) then valid := false:  $\langle finish tn := tnc; \rangle$ for t from mid tn + 1 to finish the do **if** (*write set of transaction with transaction number t intersects read set*) then valid := false; if valid then ((write phase): tnc := tnc + 1: tn := tnc)): if valid **Critical Section** then (cleanup) else (backup)).

Part of the validation process happens outside the critical section

The optimization can be applied repeatedly

Readonly transactions do not enter the critical section



### **Parallel Validation**



Validation against other transactions and writes both happen outside the critical section

Length of the critical section is independent of the number of validating transactions

Leading to unnecessary aborts



# Q/A - OCC

Why write and validation phases likely take place in RAM?

Hybrid CC that combines OCC and 2PL?

Yes. Checkout <u>MOCC</u> and <u>CormCC</u>

Concurrent way to deal with unnecessary aborts in parallel validation? tbegin vs. tcreate?

Why any serial order of transactions acceptable? Shouldn't it be the submission order?

 Strict serializability: If T1 finishes before T2 starts, T1 is before T2 in the global serial order

Practical systems using 2PL vs. OCC?

OCC vs. 2PL in performance?

# **Group Discussion**

What are the downsides of OCC compared to 2PL?

# **Before Next Lecture**

Submit discussion summary to <a href="https://wisc-cs764-f20.hotcrp.com">https://wisc-cs764-f20.hotcrp.com</a>

- Title: Lecture 7 discussion. group ##
- Authors: Names of students who joined the discussion
- Summary submission Deadline: Tuesday 11:59pm

Submit review for

- Philip L. Lehman, S. Bing Yao: <u>Efficient Locking for Concurrent Operations</u> on <u>B-Trees</u>. ACM Trans. Database Syst. 1981.
- Before next Monday (Oct. 5)