



# CS 764: Topics in Database Management Systems

## Lecture 10: Optimistic Concurrency Control

Xiangyao Yu

10/11/2021

# Announcement

---

Guest lecture next Monday (Oct. 18) from Oracle

- The lecture is offered in **online mode**

**Round-table discussion** right after the talk (2:00—3:00 PM)

- Good opportunity if you are looking for internship or full-time positions

Office hour is pushed to 3:00–3:30 PM

# Announcement

---

Project proposal deadline: **Oct. 25**

Make sure to cover the following aspects (in ~1 page)

- Project name
- Author list
- Background and motivation (why is the problem important? what are the challenges)
- Task plan (what will you do in the project? what are your key contributions?)
- Timeline

Submission website: <https://wisc-cs764-f21.hotcrp.com>

ACM format: <https://www.acm.org/publications/proceedings-template>

# 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.

# Agenda

---

Downsides of pessimistic concurrency control

Optimistic concurrency control

- Read phase
- Write phase
- Validation phase

# Concurrency Control

---

**Concurrency control** ensures the correctness for concurrent operations

Assume **serializable** isolation level for this lecture

# Concurrency Control

---

**Concurrency control** ensures the correctness 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

---

**Concurrency control** ensures the correctness 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)



# Pessimistic Concurrency Control

---

## Strict two-phase locking (2PL)

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

# Pessimistic Concurrency Control

---

## Strict two-phase locking (2PL)

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

## Downsides of pessimistic concurrency control

- Locking overhead, even for read-only transactions
- Deadlocks
- Limited concurrency due to (1) congestion and (2) holding locks till the end of a transaction

# Pessimistic Concurrency Control

---

## Strict two-phase locking (2PL)

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

## Downsides of pessimistic concurrency control

- Locking overhead, even for read-only transactions
- Deadlocks
- Limited concurrency due to (1) congestion and (2) holding locks till the end of a transaction

**Observation:** Locking is needed only if contention exists; real workloads have low contention

# Optimistic Concurrency Control (OCC)

Goal: eliminating pessimistic locking

Three executing phases:

- Read
- Validation
- Write




Fig. 1. The three phases of a transaction.

# Read Phase

---

*n = tcreate*

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

# Read Phase

---

*n = tcreate*

*twrite(n, i, v)*

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

# Read Phase

---

*n = tcreate*

*twrite(n, i, v)*

*value = tread(n, i)*

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

# Read Phase

---

*n = tcreate*

*twrite(n, i, v)*

*value = tread(n, i)*

*tdelete(n)*

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



# Read Phase

---

*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*  $n \in \text{write set}$  **do** *exchange*( $n$ , *copies*[ $n$ ]).

All created nodes become accessible

All deleted nodes become inaccessible

# Validation Phase

---

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

–  $t(i) < t(j) \Rightarrow$  exists a serial schedule where  $T_i$  is before  $T_j$


# Serial Validation

```
tbegin = (  
  start tn := tnc)
```

```
tend = (  
  <finish tn := tnc;
```

## Critical Section

```
  valid := true;  
  for t from start tn + 1 to finish tn 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  
    then (cleanup)  
    else (backup)).
```




Which transactions will T2, T3, and T4 be validated against?

**Problem:** Both *validate* and *write* phases happen in the critical section

# Improved Serial Validation

```
tend := (  
  mid tn := tnc;  
  valid := true;  
  for t from start tn + 1 to mid tn do  
    if (write set of transaction with transaction number t intersects read set)  
      then valid := false;  
  <finish tn := tnc;  
  for t from mid tn + 1 to finish tn 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  
    then (cleanup)  
    else (backup)).
```

## Critical Section



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


```
tend = (  
  (finish tn := tnc;  
  finish active := (make a copy of active);  
  active := active ∪ {id of this transaction});  
  valid := true;  
  for t from start tn + 1 to finish tn do  
    if (write set of transaction with transaction number t intersects read set)  
      then valid := false;  
  for i ∈ finish active do  
    if (write set of transaction Ti intersects read set or write set)  
      then valid := false;  
  if valid  
    then (  
      (write phase);  
      (tnc := tnc + 1;  
      tn := tnc;  
      active := active − {id of this transaction});  
      (cleanup))  
    else (  
      (active := active − {id of transaction});  
      (backup)).
```

## Critical Sections

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  
– Abort due to conflict with an aborted transaction



# Parallel Validation

```
tend = (  
  ⟨finish tn := tnc;  
  finish active := (make a copy of active);  
  active := active ∪ {id of this transaction}⟩;  
  valid := true;  
  for t from start tn + 1 to finish tn do  
    if (write set of transaction with transaction number t intersects read set)  
      then valid := false;  
  for i ∈ finish active do  
    if (write set of transaction Ti intersects read set or write set)  
      then valid := false;  
  if valid  
    then (  
      (write phase);  
      ⟨tnc := tnc + 1;  
      tn := tnc;  
      active := active — {id of this transaction}⟩;  
      (cleanup))  
    else (  
      ⟨active := active — {id of transaction}⟩;  
      (backup)).
```

**Question:** Why need to consider both read set and write set when validating against transactions in *finish active*? Can you think of a solution to avoid considering write set?

# 2PL vs. OCC

---

Revisit the motivation of OCC:

- **Locking overhead**, even for read-only transactions
- **Deadlocks**
- **Limited concurrency** due to (1) congestion and (2) holding locks till the end of a transaction

Comments:

- Optimized locks have low overhead, relative to disk and network cost
- When 2PL has limited concurrency, OCC may have high abort rate



# Q/A – OCC

---

Is OCC used in practice?

Locking required to achieve “exchange”?

How to guard critical sections?

How to keep tnc in a distributed system?

Stricter isolation than serializability?

Timestamp-based vs. MVCC vs. optimistic vs. pessimistic?

Do real workloads have low contention?

Combine pessimistic and optimistic?

# Discussion

---

What are the downsides of OCC compared to 2PL?

# Before Next Lecture

---

Submit review for

- Stephen Tu, et al., [Speedy transactions in multicore in-memory databases.](#)  
SOSP, 2013