



CS 764: Topics in Database Management Systems

Lecture 9: Isolation

Xiangyao Yu

10/6/2020

Today's Paper: Isolation

A Critique of ANSI SQL Isolation Levels

Hal Berenson	Microsoft Corp.	haroldb@microsoft.com
Phil Bernstein	Microsoft Corp.	philbe@microsoft.com
Jim Gray	U.C. Berkeley	gray@crl.com
Jim Melton	Sybase Corp.	jim.melton@sybase.com
Elizabeth O'Neil	UMass/Boston	eoneil@cs.umb.edu
Patrick O'Neil	UMass/Boston	poneil@cs.umb.edu

Abstract: ANSI SQL-92 [MS, ANSI] defines Isolation Levels in terms of *phenomena*: Dirty Reads, Non-Repeatable Reads, and Phantoms. This paper shows that these phenomena and the ANSI SQL definitions fail to properly characterize several popular isolation levels, including the standard locking implementations of the levels covered. Ambiguity in the statement of the phenomena is investigated and a more formal statement is arrived at; in addition new phenomena that better characterize isolation types are introduced. Finally, an important multiversion isolation type, called Snapshot Isolation, is defined.

1. Introduction

Running concurrent transactions at different isolation levels allows application designers to trade off concurrency and throughput for correctness. Lower isolation levels increase transaction concurrency at the risk of allowing transactions to observe a fuzzy or incorrect database state. Surprisingly, some transactions can execute at the highest isolation level (perfect serializability) while concurrently executing transactions running at a lower isolation level can access states that are not yet committed or that postdate states the transaction read earlier [GLPT]. Of course, transactions running at lower isolation levels can produce invalid data. Application designers must guard against a later transaction running at a higher isolation level accessing this invalid data and propagating such errors.

The ANSI/ISO SQL-92 specifications [MS, ANSI] define four isolation levels: (1) READ UNCOMMITTED, (2) READ COMMITTED, (3) REPEATABLE READ, (4) SERIALIZABLE. These levels are defined with the classical serializability definition, plus three prohibited operation subsequences, called *phenomena*: Dirty Read, Non-repeatable Read, and Phantom. The concept of a *phenomenon* is not explicitly defined in the ANSI specifications, but the specifications suggest that phenomena are operation subsequences that may lead to anomalous (perhaps non-serializable) behavior. We refer to *anomalies* in what follows when making suggested additions to the set of ANSI phenomena. As shown later, there is a technical distinction between anomalies and phenomena, but this distinction is not crucial for a general understanding.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association of Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.
SIGMOD '95, San Jose, CA USA
© 1995 ACM 0-89791-731-6/95/0005..\$3.50

The ANSI isolation levels are related to the behavior of lock schedulers. Some lock schedulers allow transactions to vary the scope and duration of their lock requests, thus departing from pure two-phase locking. This idea was introduced by [GLPT], which defined *Degrees of Consistency* in three ways: locking, data-flow graphs, and anomalies. Defining isolation levels by phenomena (anomalies) was intended to allow non-lock-based implementations of the SQL standard.

This paper shows a number of weaknesses in the anomaly approach to defining isolation levels. The three ANSI phenomena are ambiguous, and even in their loosest interpretations do not exclude some anomalous behavior that may arise in execution histories. This leads to some counter-intuitive results. In particular, lock-based isolation levels have different characteristics than their ANSI equivalents. This is disconcerting because commercial database systems typically use locking implementations. Additionally, the ANSI phenomena do not distinguish between a number of types of isolation level behavior that are popular in commercial systems. Additional phenomena to characterize these isolation levels are suggested here.

Section 2 introduces the basic terminology of isolation levels. It defines the ANSI SQL and locking isolation levels. Section 3 examines some drawbacks of the ANSI isolation levels and proposes a new phenomenon. Other popular isolation levels are also defined. The various definitions map between ANSI SQL isolation levels and the *degrees of consistency* defined in 1977 in [GLPT]. They also encompass Chris Date's definitions of Cursor Stability and Repeatable Read [DAT]. Discussing the isolation levels in a uniform framework reduces misunderstandings arising from independent terminology.

Section 4 introduces a multiversion concurrency control mechanism, called *Snapshot Isolation*, that avoids the ANSI SQL phenomena, but is not serializable. Snapshot Isolation is interesting in its own right, since it provides a reduced-isolation level approach that lies between READ COMMITTED and REPEATABLE READ. A new formalism (available in the longer version of this conference paper [OOBBGM]) connects reduced isolation levels for multiversioned data to the classical single-version locking serializability theory.

Section 5 explores some new anomalies to differentiate the isolation levels introduced in Sections 3 and 4. The extended ANSI SQL phenomena proposed here lack the power to characterize Snapshot isolation and Cursor Stability. Section 6 presents a Summary and Conclusions.

Agenda

ANSI isolation levels

Cursor stability and snapshot isolation

Discussions

Recap: Degree of Consistency

Degree 3: Serializability (assuming no phantom effect)

- Two-phase with respect to both reads and writes

Degree 2: Read Committed

- Two-phase with respect to writes
- Short read locks

Degree 1: Read Uncommitted

- Two-phase with respect to writes
- No read locks (may observe dirty data)

Degree 0:

- Short write locks
- No read locks

ANSI Isolation Levels

Isolation Level	P1 (or A1) Dirty Read	P2 (or A2) Fuzzy Read	P3 (or A3) Phantom
Degree 1 → ANSI READ UNCOMMITTED	Possible	Possible	Possible
Degree 2 → ANSI READ COMMITTED	Not Possible	Possible	Possible
ANSI REPEATABLE READ	Not Possible	Not Possible	Possible
Degree 3 → ANOMALY SERIALIZABLE	Not Possible	Not Possible	Not Possible

ANSI SQL-92 defines four isolation levels by phenomena

The original definitions were ambiguous

This lecture focuses on the “correct” definitions

Notation

w1[x]: transaction 1 writes record x

r2[y]: transaction 2 reads record y

w1[P] (r1[P]): transaction 1 writes (reads) records that satisfy predicate P

c1: commit of transaction 1

a1: abort of transaction 1

Locking-Based Definition

Well-formed: lock (on tuple or predicate) before reading/writing records

Long locks: hold the lock until transaction commits or aborts

Consistency Level = Locking Isolation Level	Read Locks on Data Items and Predicates (the same unless noted)	Write Locks on Data Items and Predicates (always the same)
Degree 3 = Locking SERIALIZABLE	Well-formed Reads Long duration Read locks (both)	Well-formed Writes, Long duration Write locks

Locking-Based Definition

Well-formed: lock (on tuple or predicate) before reading/writing records

Long locks: hold the lock until transaction commits or aborts

Consistency Level = Locking Isolation Level	Read Locks on Data Items and Predicates (the same unless noted)	Write Locks on Data Items and Predicates (always the same)
Locking REPEATABLE READ	Well-formed Reads Long duration data-item Read locks Short duration Read Predicate locks	Well-formed Writes, Long duration Write locks
Degree 3 = Locking SERIALIZABLE	Well-formed Reads Long duration Read locks (both)	Well-formed Writes, Long duration Write locks

Phenomenon P3: Phantom

$r1[P] \dots w2[y \text{ in } P] \dots$ (c1 or a1) and (c2 or a2) any order)

– Anomalous behavior: multiple $r[P]$'s return different results

P3 is allowed in *repeatable read* but forbidden in *serializable*

Locking-Based Definition

Well-formed: lock (on tuple or predicate) before reading/writing records

Long locks: hold the lock until transaction commits or aborts

Consistency Level = Locking Isolation Level	Read Locks on Data Items and Predicates (the same unless noted)	Write Locks on Data Items and Predicates (always the same)
Degree 2 = Locking READ COMMITTED	Well-formed Reads Short duration Read locks (both)	Well-formed Writes, Long duration Write locks
Locking REPEATABLE READ	Well-formed Reads Long duration data-item Read locks Short duration Read Predicate locks	Well-formed Writes, Long duration Write locks

Phenomenon P2: Fuzzy Read

$r1[x] \dots w2[x] \dots$ (c1 or a1) and (c2 or a2) any order)

– Anomalous behavior: multiple $r[x]$'s return different results

P2 is allowed in *read committed* but forbidden in *repeatable read*

Locking-Based Definition

Well-formed: lock (on tuple or predicate) before reading/writing records

Long locks: hold the lock until transaction commits or aborts

Consistency Level = Locking Isolation Level	Read Locks on Data Items and Predicates (the same unless noted)	Write Locks on Data Items and Predicates (always the same)
Degree 1 = Locking READ UNCOMMITTED	none required	Well-formed Writes Long duration Write locks
Degree 2 = Locking READ COMMITTED	Well-formed Reads Short duration Read locks (both)	Well-formed Writes, Long duration Write locks

Phenomenon P1: Dirty Read

$w_1[x] \dots r_2[x] \dots$ (c1 or a1) and (c2 or a2) any order)

– Anomalous behavior: transaction reads data that was never committed

P1 is allowed in *read uncommitted* but forbidden in *read committed*

Locking-Based Definition

Well-formed: lock (on tuple or predicate) before reading/writing records

Long locks: hold the lock until transaction commits or aborts

Consistency Level = Locking Isolation Level	Read Locks on Data Items and Predicates (the same unless noted)	Write Locks on Data Items and Predicates (always the same)
Degree 0	none required	Well-formed Writes Short duration Write locks
Degree 1 = Locking READ UNCOMMITTED	none required	Well-formed Writes Long duration Write locks

Phenomenon P0: Dirty Write

$w_1[x] \dots w_2[x] \dots$ (c1 or a1) and (c2 or a2) any order)

– Anomalous behavior: when transaction 1 rolls back x, unclear what value to roll back to

P0 is forbidden in all ANSI isolation levels

Equivalent Definitions

Table 3. ANSI SQL Isolation Levels Defined in terms of the four phenomena				
Isolation Level	P 0 Dirty Write	P 1 Dirty Read	P 2 Fuzzy Read	P 3 Phantom
READ UNCOMMITTED	Not Possible	Possible	Possible	Possible
READ COMMITTED	Not Possible	Not Possible	Possible	Possible
REPEATABLE READ	Not Possible	Not Possible	Not Possible	Possible
SERIALIZABLE	Not Possible	Not Possible	Not Possible	Not Possible

Consistency Level = Locking Isolation Level	Read Locks on Data Items and Predicates (the same unless noted)	Write Locks on Data Items and Predicates (always the same)
Degree 1 = Locking READ UNCOMMITTED	none required	Well-formed Writes Long duration Write locks
Degree 2 = Locking READ COMMITTED	Well-formed Reads Short duration Read locks (both)	Well-formed Writes, Long duration Write locks
Locking REPEATABLE READ	Well-formed Reads Long duration data-item Read locks Short duration Read Predicate locks	Well-formed Writes, Long duration Write locks
Degree 3 = Locking SERIALIZABLE	Well-formed Reads Long duration Read locks (both)	Well-formed Writes, Long duration Write locks

Discussion

Why define isolation levels? Why not serializability?

Why define isolation levels in this particular way?

Hierarchy of Isolation Levels

Isolation level L1 is **weaker** than isolation level L2, denoted **L1 << L2**, if all non-serializable histories that obey the criteria of L2 also satisfy L1 and there is at least one non-serializable history that can occur at level L1 but not at level L2.

Read Uncommitted

<< Read Committed

<< Repeatable Read

<< Serializability

Cursor Stability

Consistency Level = Locking Isolation Level	Read Locks on Data Items and Predicates (the same unless noted)	Write Locks on Data Items and Predicates (always the same)
Degree 2 = Locking READ COMMITTED	Well-formed Reads Short duration Read locks (both)	Well-formed Writes, Long duration Write locks
Cursor Stability (see Section 4.1)	Well-formed Reads <u>Read locks held on current of cursor</u> Short duration Read Predicate locks	Well-formed Writes, Long duration Write locks
Locking REPEATABLE READ	Well-formed Reads Long duration data-item Read locks Short duration Read Predicate locks	Well-formed Writes, Long duration Write locks

Cursor: can be viewed as a pointer to one row in a set of rows. The cursor can only reference one row at a time, but can move to other rows of the result set as needed

Phenomenon P4: Lost Update

$r1[x] \dots w2[x] \dots w1[x] \dots c1$

- Anomalous behavior: transaction 2's update is overwritten by transaction 1

Snapshot Isolation



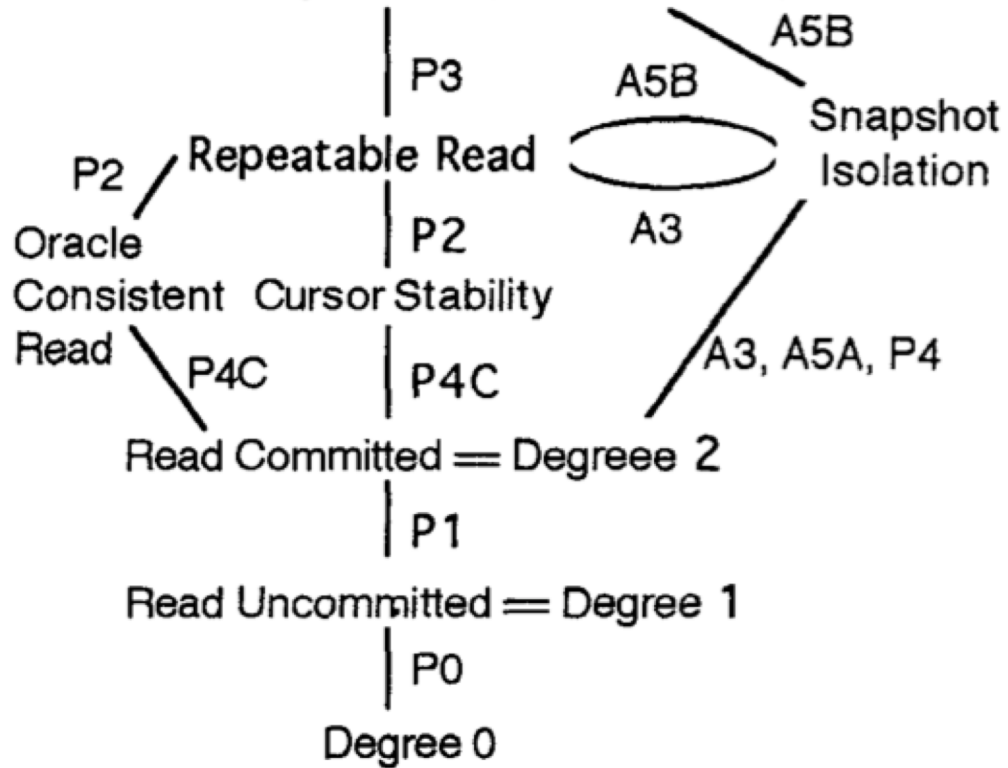
All reads see a **snapshot** of data as of the time the transaction started (t1)

A transaction can commit if records in **write set** are not modified by other transactions between t1 and t2

At commit time, apply all writes with timestamp t2

Hierarchy of Isolation Levels

Serializable = Degree 3 = {Date, DB2} Repeatable Read



Again, why define isolation levels in this particular way?

Isolation is Complex



Initially
checking.balance = 1000



```
bal = read(balance)
If bal > 100
    bal = bal - 100
    write(balance, bal)
    dispense cash
else
    reject
```

Isolation is Complex



Initially
checking.balance = 1000



```
bal = read(balance)
```

1

```
If bal > 100
```

```
    bal = bal - 100
```

```
    write(balance, bal)
```

```
    dispense cash
```

```
else
```

```
    reject
```

```
bal = read(balance)
```

2

```
If bal > 100
```

```
    bal = bal - 100
```

```
    write(balance, bal)
```

```
    dispense cash
```

```
else
```

```
    reject
```

Isolation is Complex



Initially
checking.balance = 1000



```
bal = read(balance)
If bal > 100
    bal = bal - 100
    write(balance, bal)
    dispense cash
else
    reject
```

1

3

```
bal = read(balance)
If bal > 100
    bal = bal - 100
    write(balance, bal)
    dispense cash
else
    reject
```

2

4

Isolation is Complex



Initially
checking.balance = 1000



```
bal = read(balance)
If bal > 100
    bal = bal - 100
    write(balance, bal)
dispense cash
else
    reject
```

1

3

```
bal = read(balance)
If bal > 100
    bal = bal - 100
    write(balance, bal)
dispense cash
else
    reject
```

2

4

balance = 900

ACID: Isolation – Why Strong Isolation?

MongoDB & Bitcoin: How NoSQL design flaws brought down two exchanges

 DZone April 2014

Attackers stole **896 Bitcoins ≈ 46 million US dollars**

ACID: Isolation – Why Strong Isolation?

MongoDB & Bitcoin: How NoSQL design flaws brought down two exchanges

 DZone April 2014

Attackers stole **896 Bitcoins** \approx **46 million US dollars**

Why you should pick strong consistency, whenever possible

 Google Cloud January 2018



*Systems that don't provide strong consistency ...
create a burden for application developers*



ACID: Isolation – Why Strong Isolation?

MongoDB & Bitcoin: How NoSQL design flaws brought down two exchanges

 DZone April 2014

Attackers stole **896 Bitcoins ≈ 46 million US dollars**

Q: “What is the biggest mistake in your life as an engineer?”

A: (from **Jeff Dean**)

March 2016



Not putting distributed transactions in BigTable.

In retrospect lots of teams wanted that capability and built their own with different degrees of success.



ACID: Isolation – Why Strong Isolation?

MongoDB & Bitcoin: How NoSQL design flaws brought down two exchanges

 DZone April 2014

Attackers stole **896 Bitcoins ≈ 46 million US dollars**

Q: “What is the biggest mistake in your life as an engineer?”

A: (from **Jeff Dean**)

March 2016



Not putting distributed transactions in BigTable.

In retrospect lots of teams wanted that capability and built their own with different degrees of success.



SQL (before 2000) -> NoSQL (since 2000) -> NewSQL (since 2010s)

Revisit Definition

Serializability: A transaction schedule is serializable if its outcome is equal to the outcome of its transactions executed serially

Consistency Level = Locking Isolation Level	Read Locks on Data Items and Predicates (the same unless noted)	Write Locks on Data Items and Predicates (always the same)
Degree 3 = Locking SERIALIZABLE	Well-formed Reads Long duration Read locks (both)	Well-formed Writes, Long duration Write locks

Are both definitions above equivalent?

Revisit Definition

Serializability: A transaction schedule is serializable if its outcome is equal to the outcome of its transactions executed serially

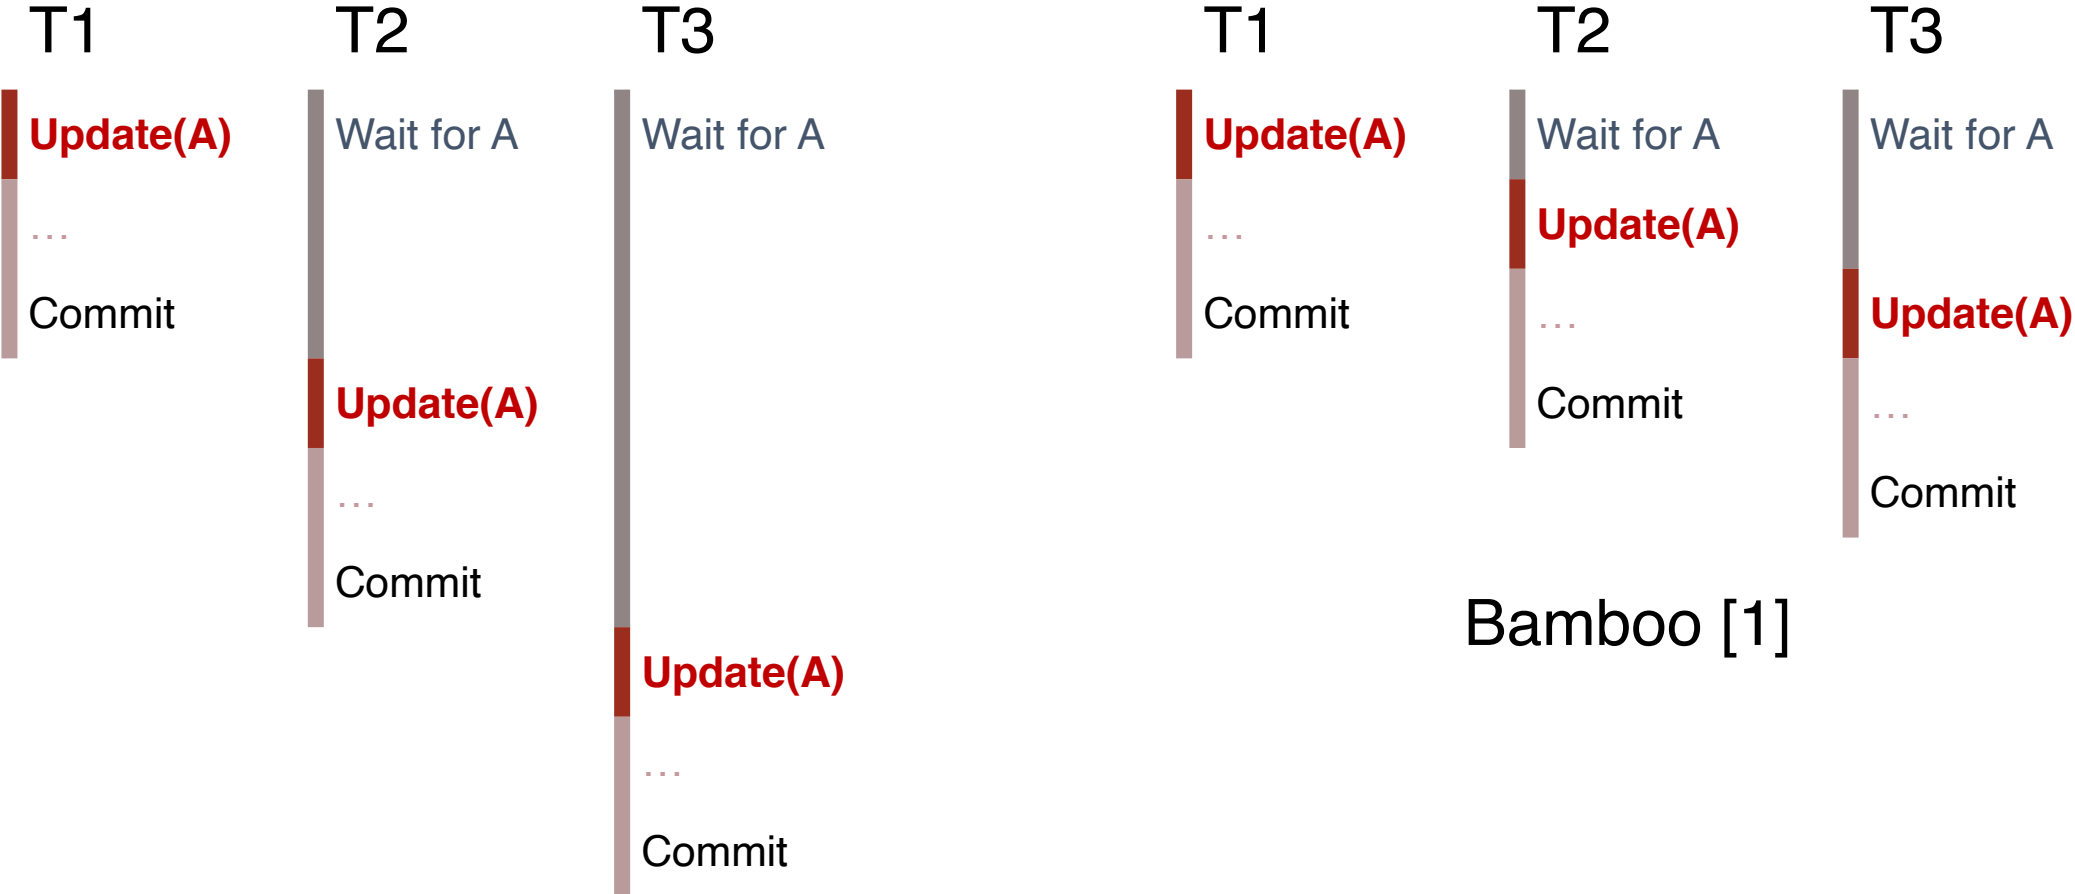
Consistency Level = Locking Isolation Level	Read Locks on Data Items and Predicates (the same unless noted)	Write Locks on Data Items and Predicates (always the same)
Degree 3 = Locking SERIALIZABLE	Well-formed Reads Long duration Read locks (both)	Well-formed Writes, Long duration Write locks

Are both definitions above equivalent?

Consider:

- T1: w1[x], r1[y], r1[z]
- T2: w2[x], r2[y], r2[z]
- Legal schedule? w1[x], w2[x], r1[y], r1[z], r2[y], r2[z], c1, c2

Optimize for Hotspots

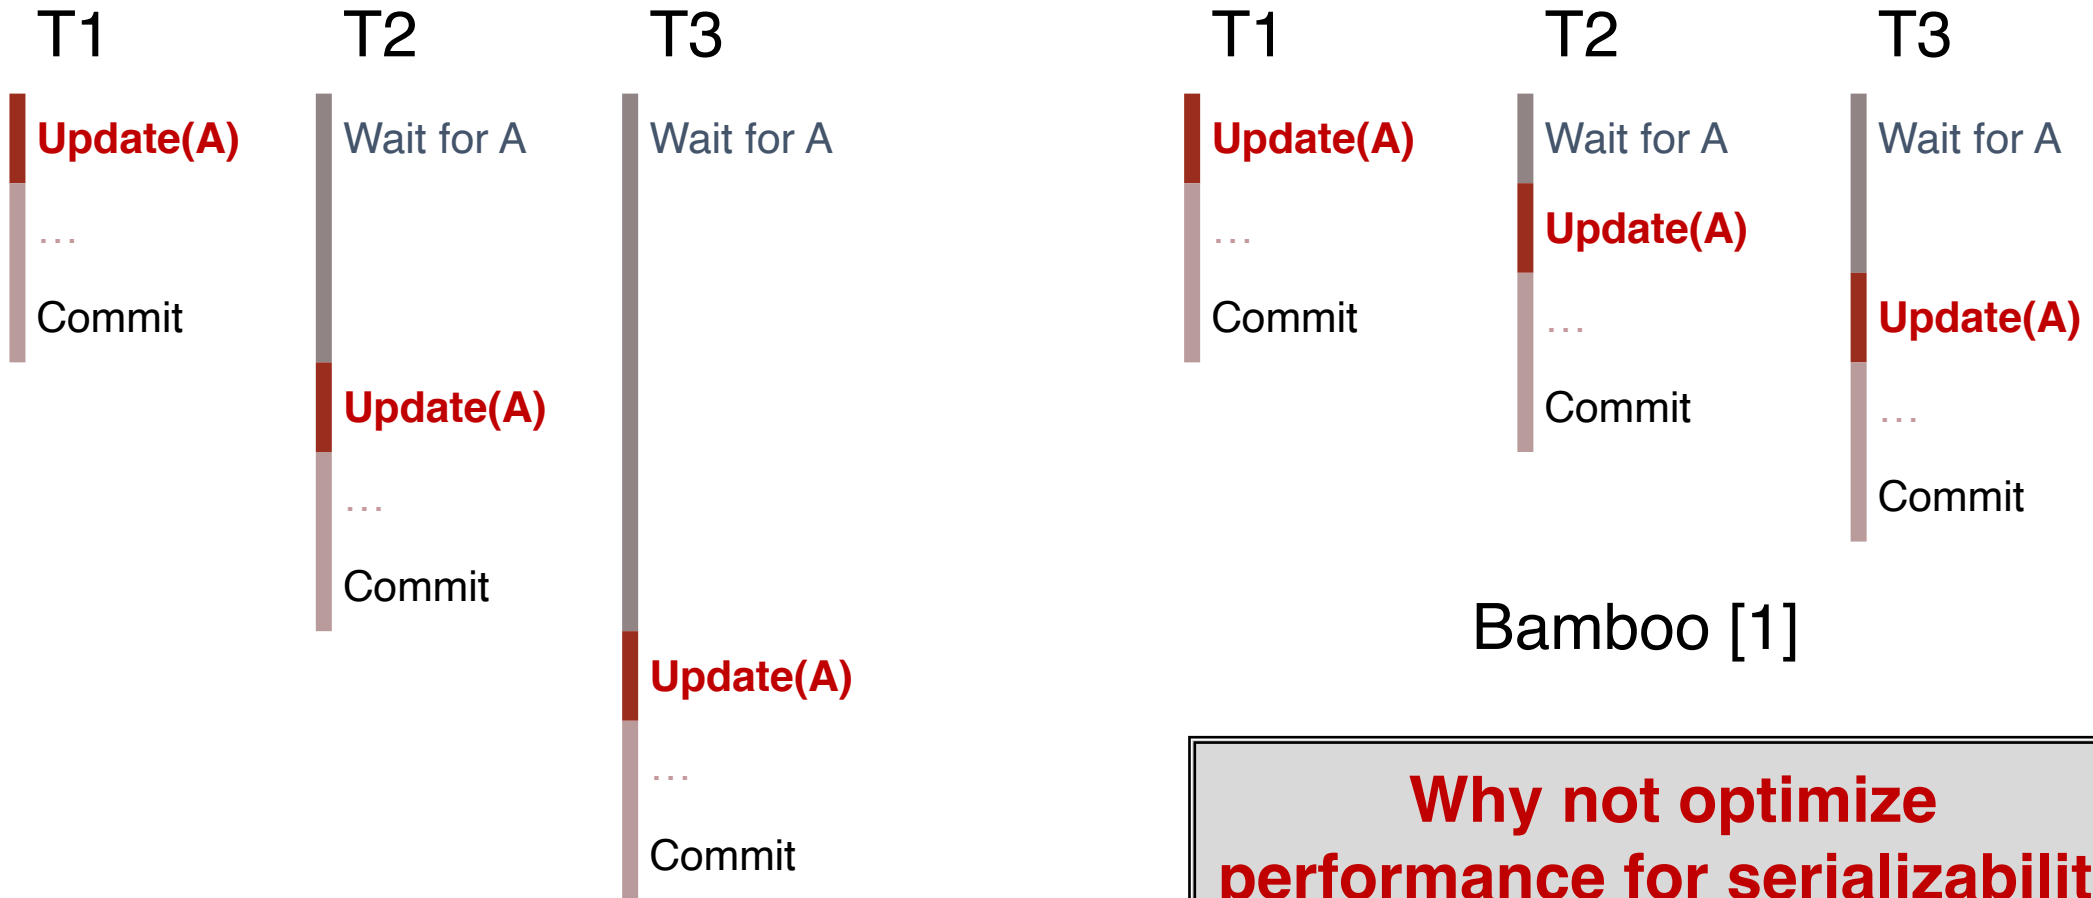


Conventional 2PL

Bamboo [1]

[1] Zhihan Guo, et al. *Releasing Locks As Early As You Can: Reducing Contention of Hotspots by Violating Two-Phase Locking*, SIGMOD 2021

Optimize for Hotspots



Conventional 2PL

Bamboo [1]

Why not optimize performance for serializability instead of relaxing it?

[1] Zhihan Guo, et al. *Releasing Locks As Early As You Can: Reducing Contention of Hotspots by Violating Two-Phase Locking*, SIGMOD 2021

Q/A – Isolation

Isolation levels in commercial DBMS today?

– Which are used? Are they all used?

The current ANSI standard?

Long running transactions in Snapshot Isolation?

Why do we need all these isolation levels?

Ongoing research to discover new phenomenon or isolation levels?

Before Next Lecture

Submit review for

- H. T. Kung, John T. Robinson, [On Optimistic Methods for Concurrency Control](#). ACM Transactions on Database Systems, 1981