

# TRANSACTIONS: PESSIMISTIC

CSE 4/562: Database Systems | Lecture 19

---

**DB. Sys.: T.C.B.:** Ch. 18.1-18.2, 19.1

- 1. Checkpoint 4 released**
- 2. Schedule checkpoint 3 code reviews**

# Quiz

## Conflict Equivalence

Two schedules are conflict equivalent if there is a sequence of pairwise “flips” of reads or operations on different objects that transforms one schedule into the other.

## Conflict Serializability

A schedule is conflict serializable if it is conflict equivalent to a serial schedule.

**How do we decide  
whether a schedule is  
conflict serializable?**

## Conflicts

write/write, read/write, and write/read operation pairs can't be flipped.

Let's give these types of operation pairs a name: **Conflicts**<sup>1</sup>.

---

<sup>1</sup>A note on language: A conflict does **not** indicate a problem.

<b>Time</b>	<b>T1</b>	<b>T2</b>
↓	W(A)	
↓		W(A)

<b>Time</b>	<b>T1</b>	<b>T2</b>
↓	W(A)	
↓		W(A)

This conflict requires that T1 **happens before** T2

Conflicts create a partial order over the conflict-equivalent serial schedules.

<b>Time</b>	<b>T1</b>	<b>T2</b>
↓		W(B)
↓	R(B)	
↓		W(A)
↓	W(A)	

In this schedule...

<b>Time</b>	<b>T1</b>	<b>T2</b>
↓		W(B)
↓	R(B)	
↓		W(A)
↓	W(A)	

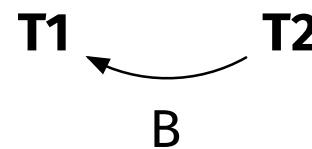
In this schedule...

- **T2**'s write to B “happens before” **T1**'s read.

<b>Time</b>	<b>T1</b>	<b>T2</b>
↓		W(B)
↓	R(B)	
↓		W(A)
↓	W(A)	

In this schedule...

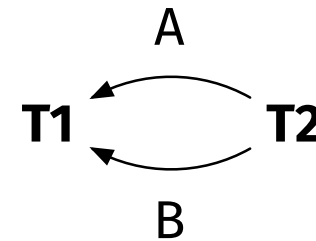
- **T2**'s write to B "happens before" **T1**'s read.



<b>Time</b>	<b>T1</b>	<b>T2</b>
↓		W(B)
↓	R(B)	
↓		W(A)
↓	W(A)	

In this schedule...

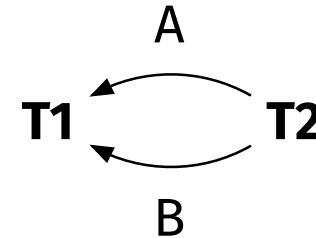
- **T2**'s write to B “happens before” **T1**'s read.
- **T2**'s write to A “happens before” **T1**'s write.



Time	T1	T2
↓		W(B)
↓	R(B)	
↓		W(A)
↓	W(A)	

In this schedule...

- **T2**'s write to B “happens before” **T1**'s read.
- **T2**'s write to A “happens before” **T1**'s write.



No cycles in the “conflict graph”  
Schedule is Conflict Serializable

<b>Time</b>	<b>T1</b>	<b>T2</b>
↓		W(B)
↓	R(B)	
↓	W(A)	
↓		W(A)

In this schedule...

<b>Time</b>	<b>T1</b>	<b>T2</b>
↓		W(B)
↓	R(B)	
↓	W(A)	
↓		W(A)

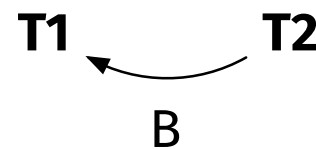
In this schedule...

- **T2**'s write to B “happens before” **T1**'s read.

<b>Time</b>	<b>T1</b>	<b>T2</b>
↓		W(B)
↓	R(B)	
↓	W(A)	
↓		W(A)

In this schedule...

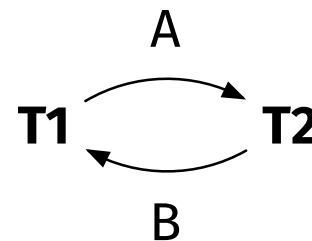
- **T2**'s write to B "happens before" **T1**'s read.



<b>Time</b>	<b>T1</b>	<b>T2</b>
↓		W(B)
↓	R(B)	
↓	W(A)	
↓		W(A)

In this schedule...

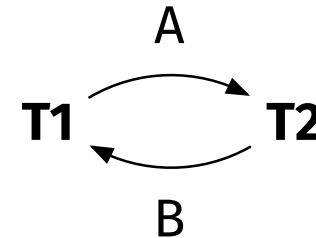
- **T2**'s write to B “happens before” **T1**'s read.
- **T1**'s write to A “happens before” **T2**'s write.



Time	T1	T2
↓		W(B)
↓	R(B)	
↓	W(A)	
↓		W(A)

In this schedule...

- **T2**'s write to B “happens before” **T1**'s read.
- **T1**'s write to A “happens before” **T2**'s write.



Cycle = Schedule is **not** Conflict Serializable

<b>Time</b>	<b>T1</b>	<b>T2</b>	<b>T3</b>
↓	R(A)		
↓		R(B)	
↓		W(A)	
↓			W(A)
↓	W(B)		

In this schedule...

**T1**

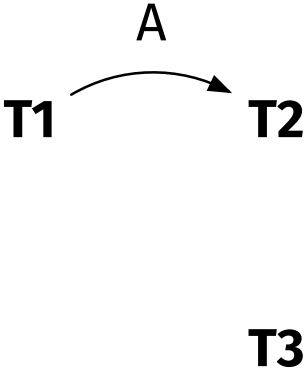
**T2**

**T3**

Time	T1	T2	T3
↓	R(A)		
↓		R(B)	
↓		W(A)	
↓			W(A)
↓	W(B)		

In this schedule...

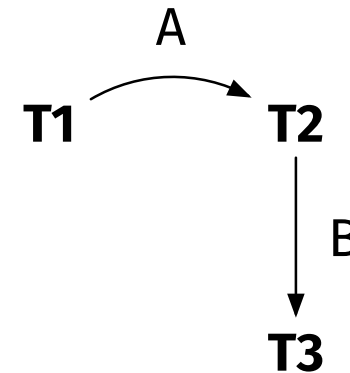
- **T1**'s read of A "happens before" **T2**'s write



Time	T1	T2	T3
↓	R(A)		
↓		R(B)	
↓		W(A)	
↓			W(A)
↓	W(B)		

In this schedule...

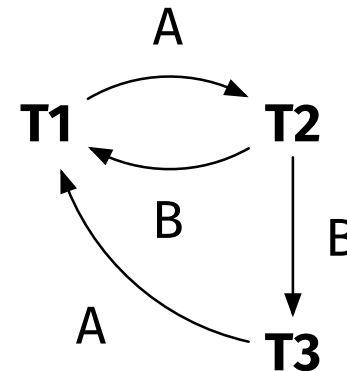
- **T1**'s read of A “happens before” **T2**'s write
- **T2**'s read of B “happens before” **T3**'s write



Time	T1	T2	T3
↓	R(A)		
↓		R(B)	
↓		W(A)	
↓			W(A)
↓	W(B)		

In this schedule...

- **T1**'s read of A “happens before” **T2**'s write
- **T2**'s read of B “happens before” **T3**'s write
- **T2**'s write to B and **T3**'s write to A “happen before” **T1**'s write

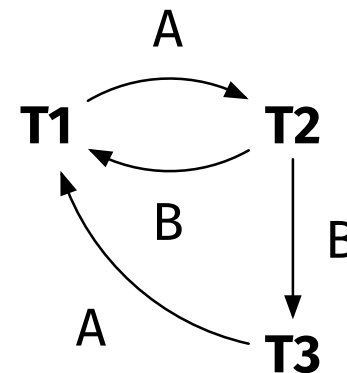


Time	T1	T2	T3
↓	R(A)		
↓		R(B)	
↓		W(A)	
↓			W(A)
↓	W(B)		

Cycle!

In this schedule...

- **T1**'s read of A “happens before” **T2**'s write
- **T2**'s read of B “happens before” **T3**'s write
- **T2**'s write to B and **T3**'s write to A “happen before” **T1**'s write



1. Add one node for every transaction
2. For every pair of operations, add an edge from the later operation's transaction to the earlier operation's transaction if:
  - ... the operations are from different transactions
  - ... the operations act on the same object
  - ... at least one operation is a write

If the resulting schedule is acyclic, the schedule is **conflict serializable**

If the schedule is conflict serializable, it is serializable.

If the schedule is not conflict serializable, it may still be serializable...

... but we don't have the tools to prove it

You will be expected to know whether a schedule is (conflict) serializable...

You will be expected to know whether a schedule is (conflict) serializable...  
... but this is usually not the point.

You will be expected to know whether a schedule is (conflict) serializable...  
... but this is usually not the point.

By learning to work with serializability, we can prove to ourselves that a new concurrency control scheme can **guarantee** that it will only ever permit transactions to execute according to a (conflict) serializable schedule.

Allow concurrency, but...

## **Locking (Pessimistic)**

... pause a transaction before it ever risks non-serializable behavior.

## **Snapshot Isolation (Optimistic)**

... check whether a transaction had a serializable schedule prior to committing and abort it if not.

## **Timestamp Concurrency Control (Optimistic)**

... detect serializability violations while the transaction is running and abort.

Allow concurrency, but...

## **Locking (Pessimistic)**

... pause a transaction before it ever risks non-serializable behavior.

## **Snapshot Isolation (Optimistic)**

... check whether a transaction had a serializable schedule prior to committing and abort it if not.

## **Timestamp Concurrency Control (Optimistic)**

... detect serializability violations while the transaction is running and abort.

**Observation:** All conflicts require access to the same object.

**Idea:** When a second transaction tries to access an object, block it until the first `COMMITs` or `ABORTs` first.

**Idea:** When a second transaction tries to access an object, block it until the first `COMMITs` or `ABORTs` first.

## Locking

1. A transaction must **Lock** every object it accesses before the access
2. All locks are released when the transaction `COMMITs` or `ABORTs`

<b>Time</b>	<b>T1</b>	<b>T2</b>	<b>T3</b>
↓	R(A)		
↓		R(B)	
↓		W(A)	
↓		COMMIT	
↓			W(A)
↓			COMMIT
↓	W(B)		
↓	COMMIT		

<b>Time</b>	<b>T1</b>	<b>T2</b>	<b>T3</b>
↓	Lock(A)		
↓	R(A)		
↓		R(B)	
↓		W(A)	
↓		COMMIT	
↓			W(A)
↓			COMMIT
↓	W(B)		
↓	COMMIT		

<b>Time</b>	<b>T1</b>	<b>T2</b>	<b>T3</b>
↓	Lock(A)		
↓	R(A)		
↓		Lock(B)	
↓		R(B)	
↓		W(A)	
↓		COMMIT	
↓			W(A)
↓			COMMIT
↓	W(B)		
↓	COMMIT		

Time	T1	T2	T3
↓	Lock(A)		
↓	R(A)		
↓		Lock(B)	
↓		R(B)	
↓		Lock(A)	
↓		W(A)	
↓		COMMIT	
↓			W(A)
↓			COMMIT
↓	W(B)		
↓	COMMIT		

Time	T1	T2	T3
↓	Lock(A)		
↓	R(A)		
↓		Lock(B)	
↓		R(B)	
↓		😞	
↓		W(A)	
↓		COMMIT	
↓			W(A)
↓			COMMIT
↓	W(B)		
↓	COMMIT		

Time	T1	T2	T3
↓	Lock(A)		
↓	R(A)		
↓		Lock(B)	
↓		R(B)	
↓		😞	
↓		W(A)	
↓		COMMIT	
↓			W(A)
↓			COMMIT
↓	W(B)		
↓	COMMIT		

This schedule is not permitted!

Time	T1	T2	T3	Locks
↓	R(A)			
↓	W(B)			
↓	COMMIT			
↓		R(B)		
↓			W(A)	
↓			COMMIT	
↓		W(A)		
↓		COMMIT		

<b>Time</b>	<b>T1</b>	<b>T2</b>	<b>T3</b>	<b>Locks</b>
↓	Lock(A)			<b>T1</b> holds { A }
↓	R(A)			
↓	W(B)			
↓	COMMIT			
↓		R(B)		
↓			W(A)	
↓			COMMIT	
↓		W(A)		
↓		COMMIT		

Time	T1	T2	T3	Locks
↓	Lock(A)			<b>T1</b> holds { A }
↓	R(A)			
↓	Lock(B)			<b>T1</b> holds { A, B }
↓	W(B)			
↓	COMMIT			
↓		R(B)		
↓			W(A)	
↓			COMMIT	
↓		W(A)		
↓		COMMIT		

Time	T1	T2	T3	Locks
↓	Lock(A)			<b>T1</b> holds { A }
↓	R(A)			
↓	Lock(B)			<b>T1</b> holds { A, B }
↓	W(B)			
↓	COMMIT			<b>T1</b> holds $\emptyset$
↓		R(B)		
↓			W(A)	
↓			COMMIT	
↓		W(A)		
↓		COMMIT		

Time	T1	T2	T3	Locks
↓	Lock(A)			<b>T1</b> holds { A }
↓	R(A)			
↓	Lock(B)			<b>T1</b> holds { A, B }
↓	W(B)			
↓	COMMIT			<b>T1</b> holds $\emptyset$
↓		Lock(B)		<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }
↓		R(B)		
↓			W(A)	
↓			COMMIT	
↓		W(A)		
↓		COMMIT		

Time	T1	T2	T3	Locks
↓	Lock(A)			<b>T1</b> holds { A }
↓	R(A)			
↓	Lock(B)			<b>T1</b> holds { A, B }
↓	W(B)			
↓	COMMIT			<b>T1</b> holds $\emptyset$
↓		Lock(B)		<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }
↓		R(B)		
↓			Lock(A)	<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }, <b>T3</b> holds { A }
↓			W(A)	
↓			COMMIT	
↓		W(A)		
↓		COMMIT		

Time	T1	T2	T3	Locks
↓	Lock(A)			<b>T1</b> holds { A }
↓	R(A)			
↓	Lock(B)			<b>T1</b> holds { A, B }
↓	W(B)			
↓	COMMIT			<b>T1</b> holds $\emptyset$
↓		Lock(B)		<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }
↓		R(B)		
↓			Lock(A)	<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }, <b>T3</b> holds { A }
↓			W(A)	
↓			COMMIT	<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }, <b>T3</b> holds $\emptyset$
↓		W(A)		
↓		COMMIT		

Time	T1	T2	T3	Locks
↓	Lock(A)			<b>T1</b> holds { A }
↓	R(A)			
↓	Lock(B)			<b>T1</b> holds { A, B }
↓	W(B)			
↓	COMMIT			<b>T1</b> holds $\emptyset$
↓		Lock(B)		<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }
↓		R(B)		
↓			Lock(A)	<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }, <b>T3</b> holds { A }
↓			W(A)	
↓			COMMIT	<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }, <b>T3</b> holds $\emptyset$
↓		Lock(A)		<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B, A }, <b>T3</b> holds $\emptyset$
↓		W(A)		
↓		COMMIT		

Time	T1	T2	T3	Locks
↓	Lock(A)			<b>T1</b> holds { A }
↓	R(A)			
↓	Lock(B)			<b>T1</b> holds { A, B }
↓	W(B)			
↓	COMMIT			<b>T1</b> holds $\emptyset$
↓		Lock(B)		<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }
↓		R(B)		
↓			Lock(A)	<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }, <b>T3</b> holds { A }
↓			W(A)	
↓			COMMIT	<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }, <b>T3</b> holds $\emptyset$
↓		Lock(A)		<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B, A }, <b>T3</b> holds $\emptyset$
↓		W(A)		
↓		COMMIT		<b>T1</b> holds $\emptyset$ , <b>T2</b> holds $\emptyset$ , <b>T3</b> holds $\emptyset$

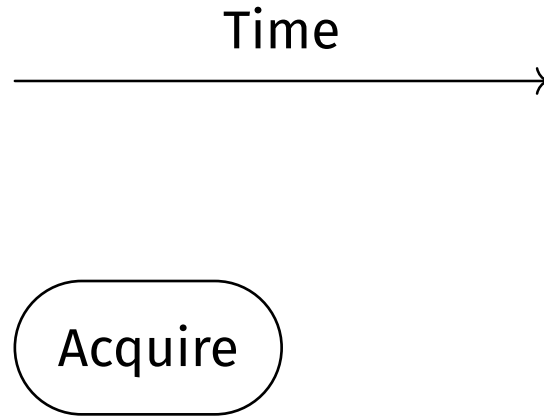
Time	T1	T2	T3	Locks
↓	Lock(A)			<b>T1</b> holds { A }
↓	R(A)			
↓	Lock(B)			<b>T1</b> holds { A, B }
↓	W(B)			
↓	COMMIT			<b>T1</b> holds $\emptyset$
↓		Lock(B)		<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }
↓		R(B)		
↓			Lock(A)	<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }, <b>T3</b> holds { A }
↓			W(A)	
↓			COMMIT	<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B }, <b>T3</b> holds $\emptyset$
↓		Lock(A)		<b>T1</b> holds $\emptyset$ , <b>T2</b> holds { B, A }, <b>T3</b> holds $\emptyset$
↓		W(A)		
↓		COMMIT		<b>T1</b> holds $\emptyset$ , <b>T2</b> holds $\emptyset$ , <b>T3</b> holds $\emptyset$

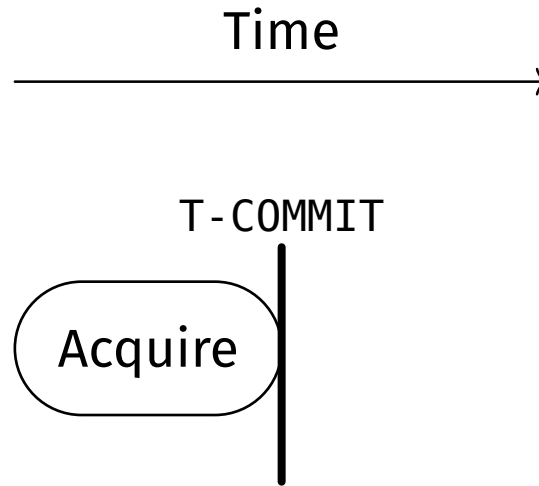
This schedule **is** permitted!

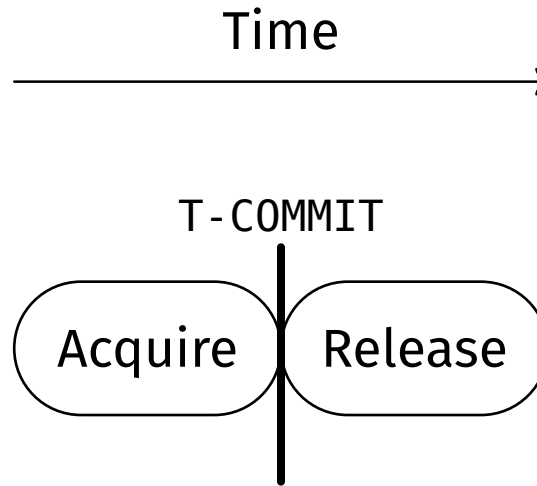
Locking is usually used to protect critical sections.

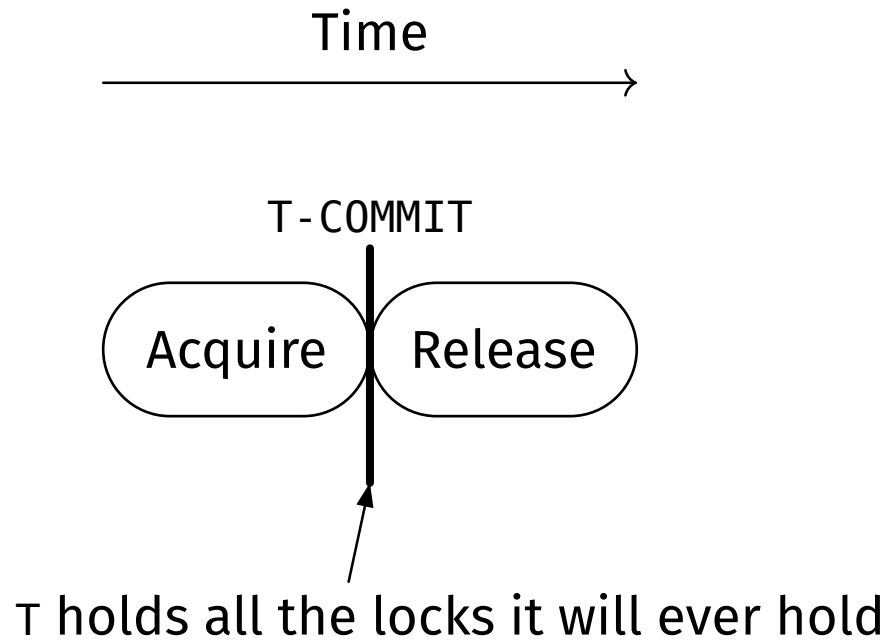
**That is NOT what is happening here**

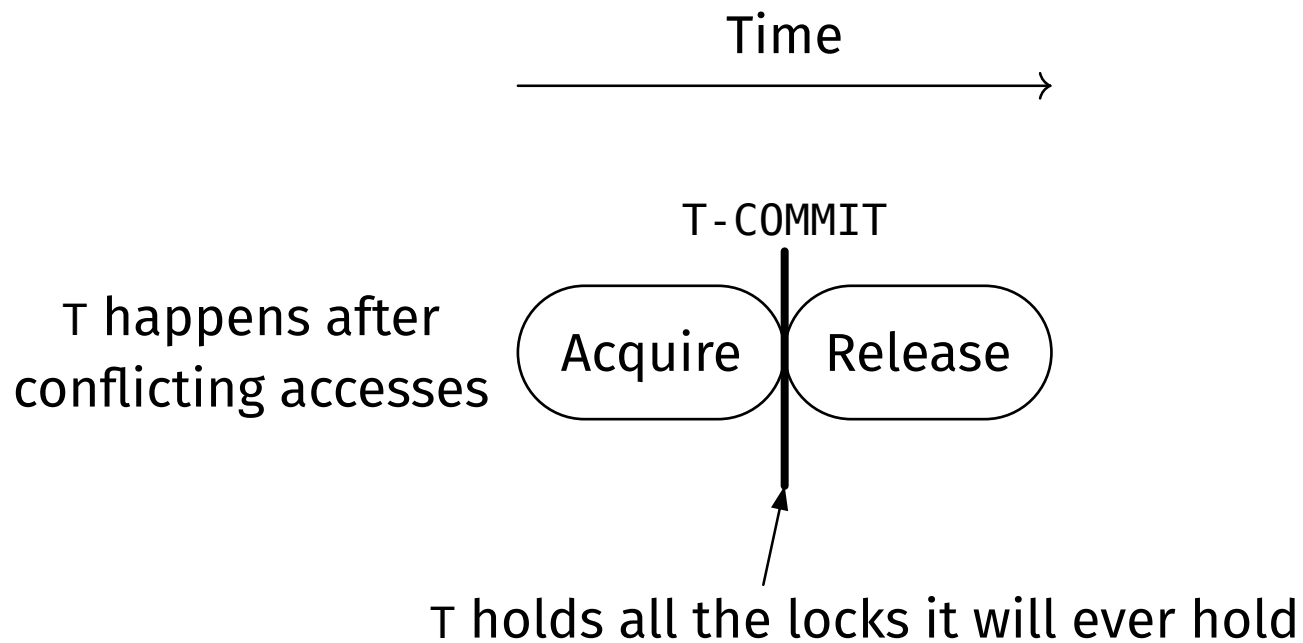
We are using locking to enforce an order over the transactions.

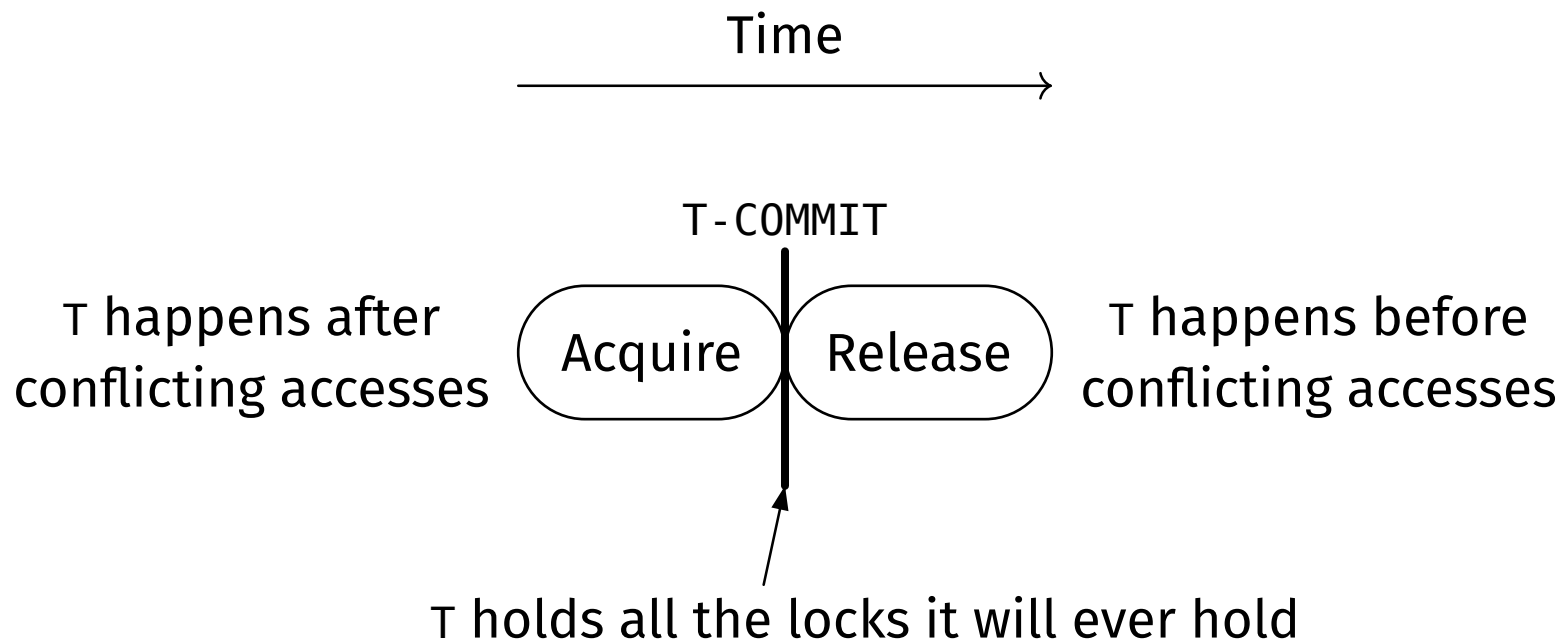












We can be a bit more general...

A transaction must hold (have acquired, but not released) a lock to access a variable.

- 1. Acquire Phase:** Transaction can only acquire locks (but not release them)
  - T “Happens after” any permitted concurrent writes
- 2. Release Phase:** Transaction can only release locks (but not acquire them)
  - T “Happens before” any permitted concurrent writes

This is called 2-Phase Locking

**2-Phase Locking only permits conflict serializable schedules**

General 2-Phase locking is stronger than “Release at COMMIT” locking.

<b>Time</b>	<b>T1</b>	<b>T2</b>	<b>T3</b>
↓	W(A)		
↓		W(A)	
↓	W(B)		
↓			W(A)
↓		W(B)	
↓			W(B)

General 2-Phase locking is stronger than “Release at COMMIT” locking.

Time	T1	T2	T3
↓	Lock(A)		
↓	Lock(B)		
↓	W(A)		
↓	Unlock(A)		
↓		Lock(A)	
↓		W(A)	
↓	W(B)		
↓	Unlock(B)		
↓		Lock(B)	
↓		Unlock(A)	
↓			Lock(A)
↓			W(A)
↓		W(B)	
↓		Unlock(B)	
↓			Lock(B)
↓			W(B)
↓			Unlock(A)
↓			Unlock(B)

... although some conflict serializable schedules may not be permitted

<b>Time</b>	<b>T1</b>	<b>T2</b>	<b>T3</b>
↓	W(A)		
↓		W(A)	
↓			W(A)
↓	W(B)		
↓		W(B)	
↓			W(B)

**A schedule admitted by 2-Phase locking**

... is always ...

**A conflict serializable schedule**

... is always ...

**A serializable schedule**

(But the reverse is not true)

# Optimizations

**Observation:** Read-Read conflicts shouldn't be a problem

**Observation:** Read-Read conflicts shouldn't be a problem

**Solution:** Reader/Writer Locks

(Any number of readers **or** one writer can hold the lock simultaneously)

**Observation:** Read-Read conflicts shouldn't be a problem

**Solution:** Reader/Writer Locks

(Any number of readers **or** one writer can hold the lock simultaneously)

... also called Shared (S) / Exclusive (X) locks

		Requested	
		S	X
Held	None	Allow	Allow
	S	Allow	Block
	X	Block	Block

## A little abstraction

I'm going to use the word “object” to mean some database “thing”.

- ... a table
- ... a record
- ... a collection of records (e.g., a page)
- ... a column of data
- ... a specific field in a specific record

## A little abstraction

I'm going to use the word “object” to mean some database “thing”.

- ... a table
- ... a record
- ... a collection of records (e.g., a page)
- ... a column of data
- ... a specific field in a specific record

Which should we actually implement?

**Observation 1:** If the locked objects are too coarse, we lose concurrency opportunities.

**Observation 2:** If the locked objects are too fine, we have to do more work locking them.

**Idea:** Separate locks for tables, pages, rows, cells, etc...

~~**Idea:** Separate locks for tables, pages, rows, cells, etc...~~

**Doesn't work:** Locking row on a page should block an attempt to lock the page.

~~**Idea:** Separate locks for tables, pages, rows, cells, etc...~~

**Doesn't work:** Locking row on a page should block an attempt to lock the page.

**Better Idea:** Lock the page (and table) before locking the row.

~~**Idea:** Separate locks for tables, pages, rows, cells, etc...~~

**Doesn't work:** Locking row on a page should block an attempt to lock the page.

~~**Better Idea:** Lock the page (and table) before locking the row.~~

**Doesn't work:** Kills concurrent access to different rows.

## **Exclusive (X)**

The holding thread can safely modify the object

## **Shared (S)**

The holding thread can safely read the object

## **Exclusive (X)**

The holding thread can safely modify the object

## **Shared (S)**

The holding thread can safely read the object

## **Intent-Exclusive (IX)**

The holding thread can safely acquire an exclusive lock on a smaller component of the object

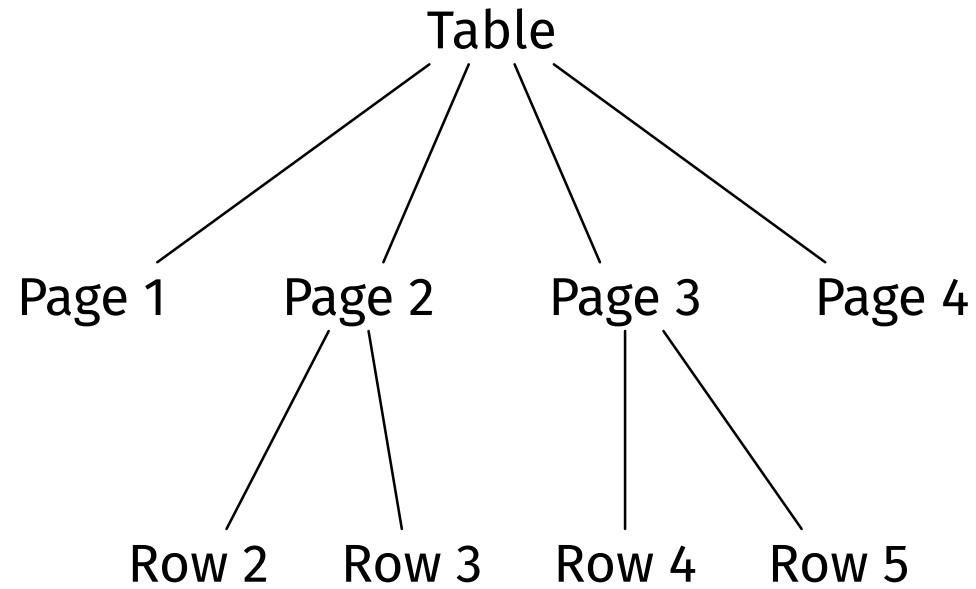
## **Intent-Shared (IS)**

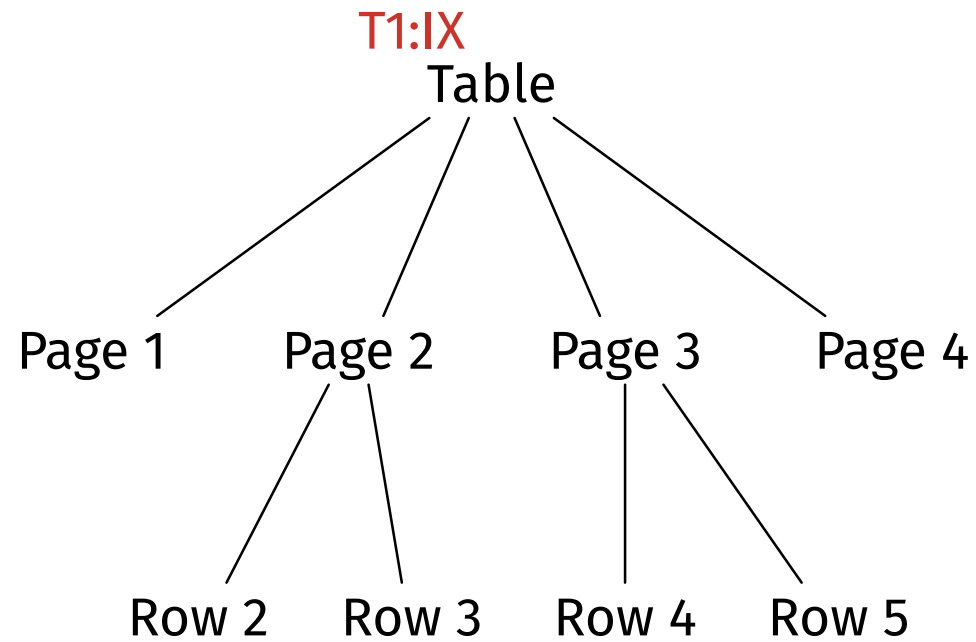
The holding thread can safely acquire a shared lock on a smaller component of the object

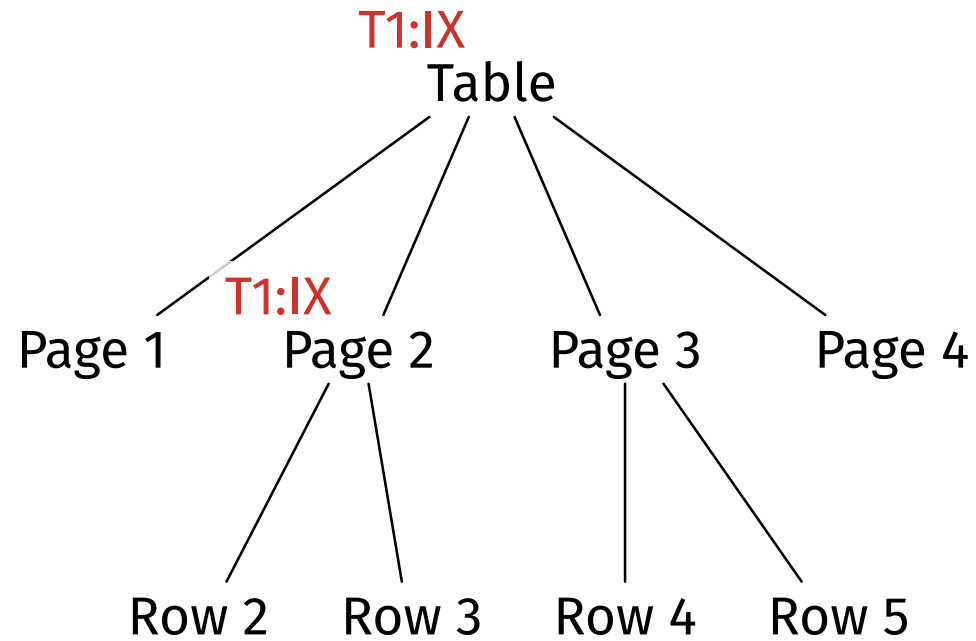
To acquire a shared (exclusive) lock, first acquire an intent-shared (intent-exclusive) lock on all containing objects.

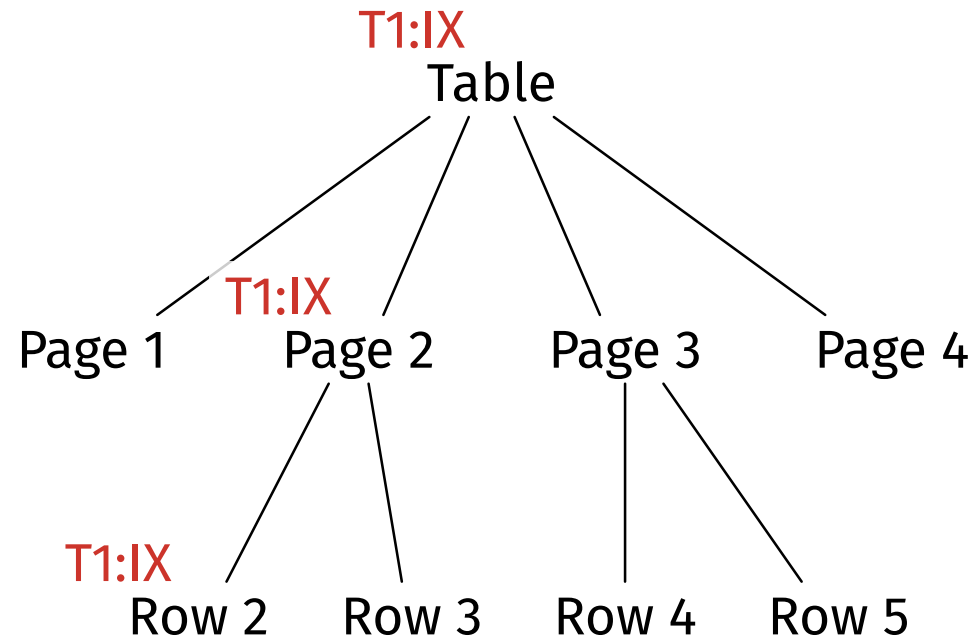
1. Acquire-IX(Table)
2. Acquire-IX(Page)
3. Acquire-IX(Row)
4. Acquire-X(Cell)
5. Modify Cell
6. Release all locks in reverse order

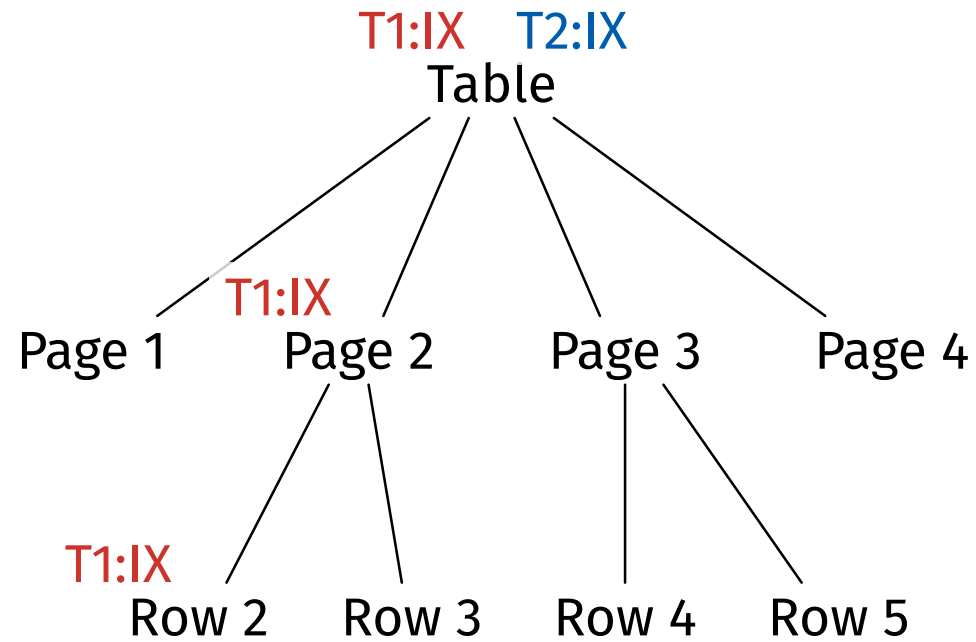
		Requested			
		IS	IX	S	X
Held	None	Allow	Allow	Allow	Allow
	IS	Allow	Allow	Allow	Block
	IX	Allow	Allow	Block	Block
	S	Allow	Block	Allow	Block
	X	Block	Block	Block	Block

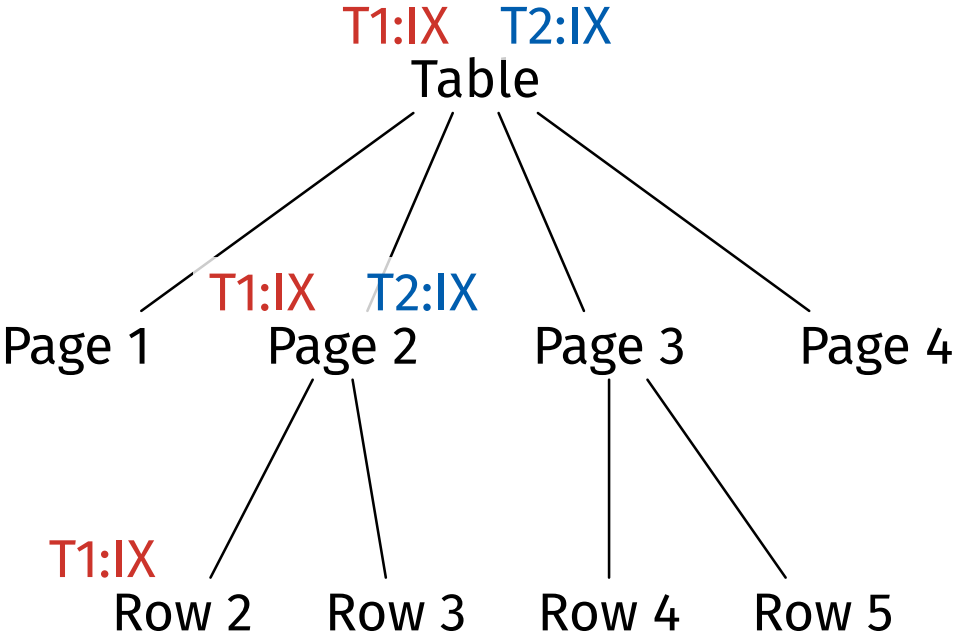


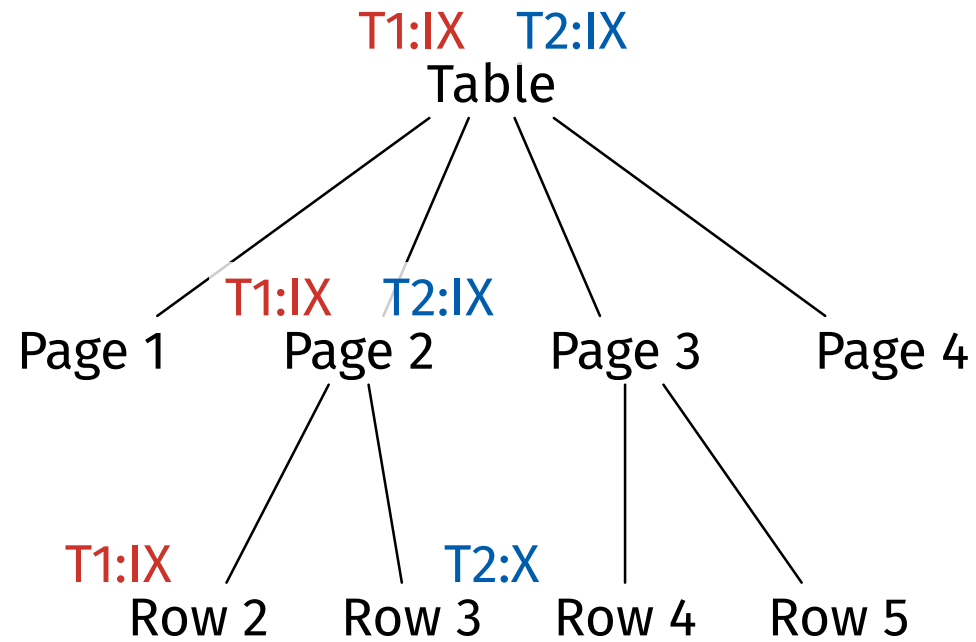


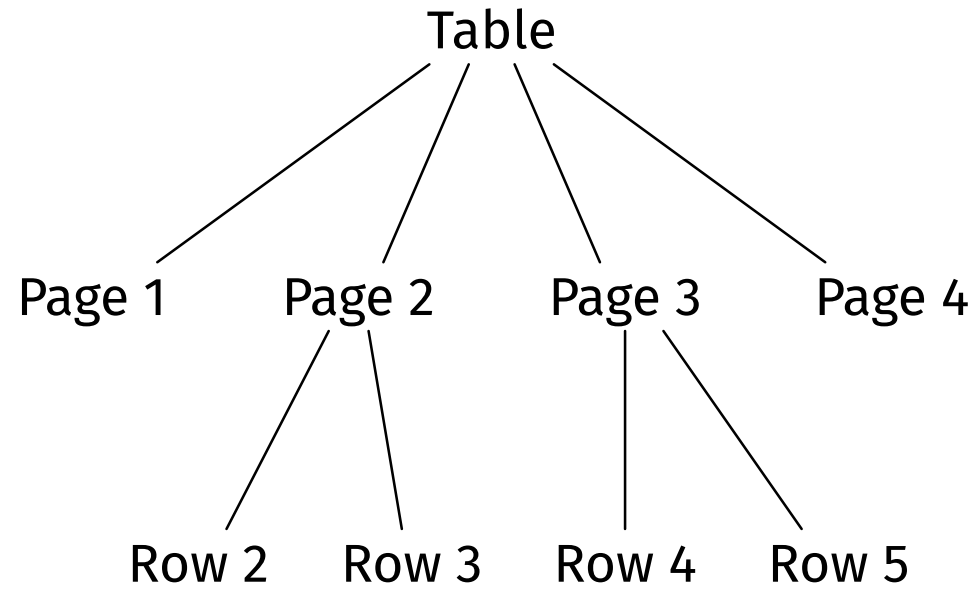


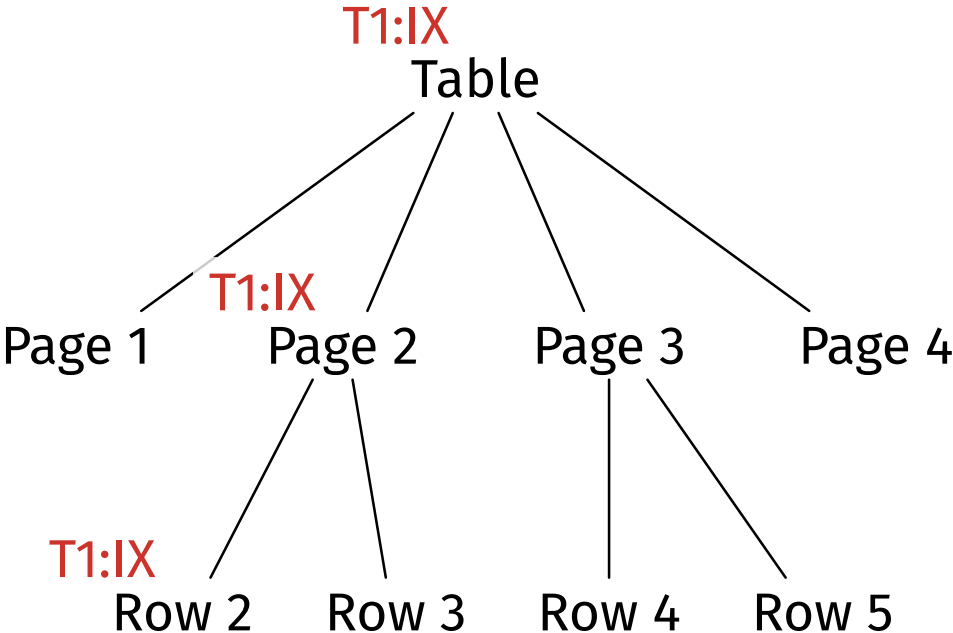


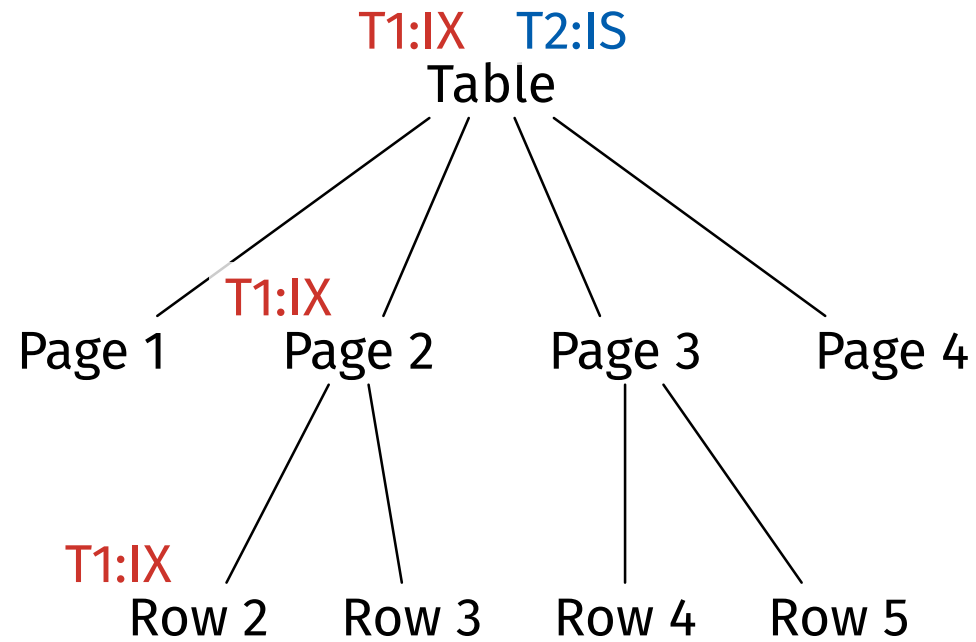


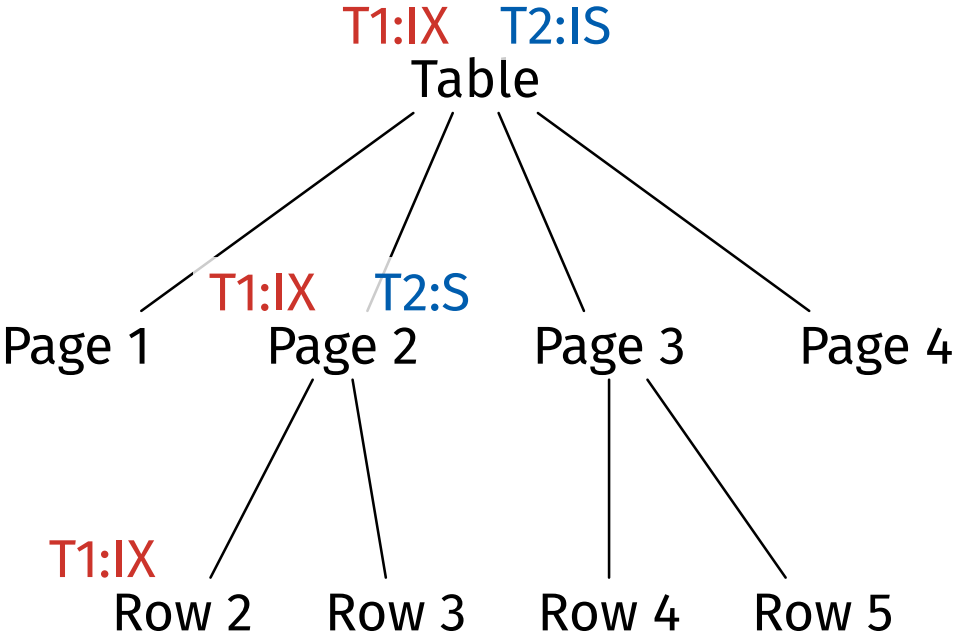


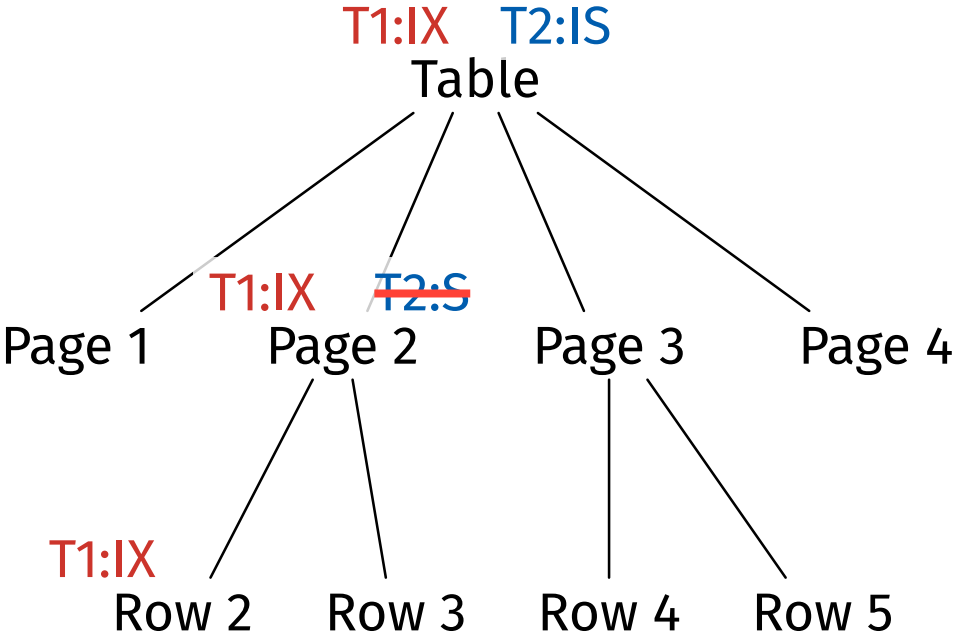












# Deadlock

<b>Time</b>	<b>T1</b>	<b>T2</b>
↓	Lock(A)	
↓	R(A)	
↓		Lock(B)
↓		R(B)
↓	???	
↓	W(A)	
↓		???
↓		W(B)

## **Optimistic**

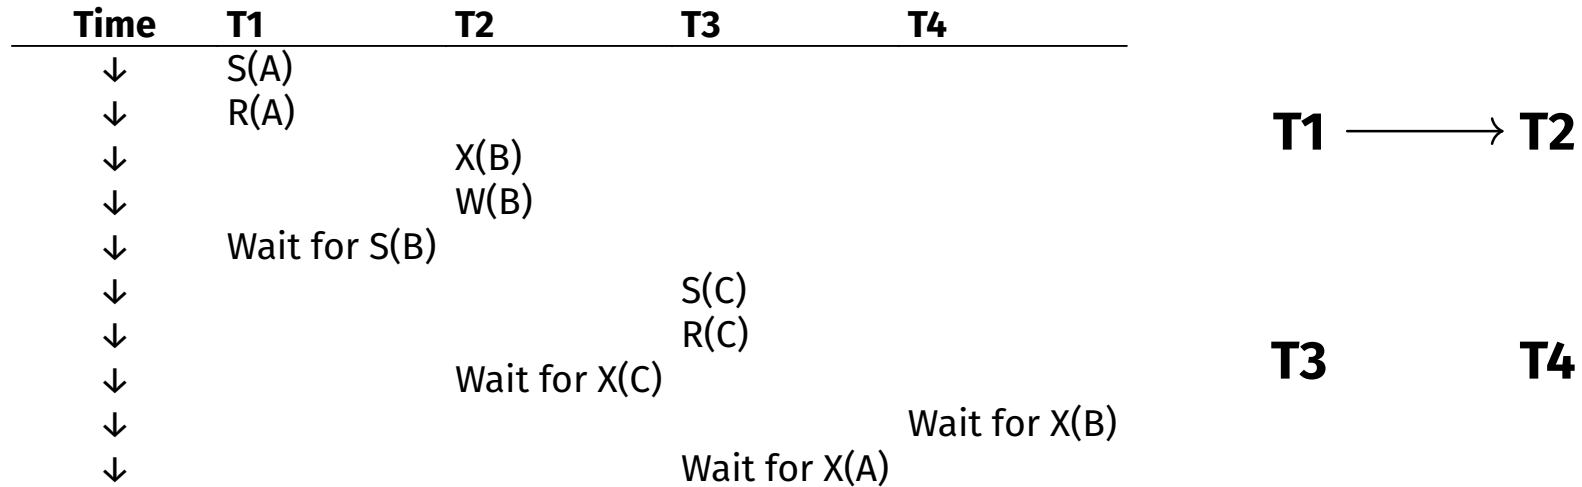
Detect deadlock situations when they happen and abort the deadlocked transaction.

## **Pessimistic**

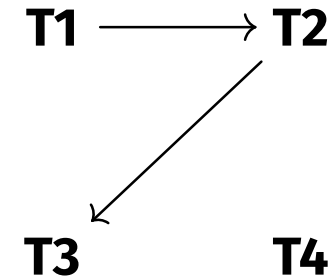
Anticipate deadlock situations and abort preemptively

- Add one node for every **running** transaction
- Add one edge from a blocked transaction to the transaction(s) holding the conflicting lock.

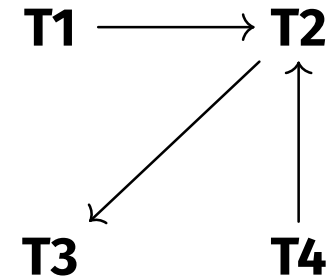
<b>Time</b>	<b>T1</b>	<b>T2</b>	<b>T3</b>	<b>T4</b>		
↓	S(A)					
↓	R(A)					
↓		X(B)			<b>T1</b>	<b>T2</b>
↓		W(B)				
↓	Wait for S(B)					
↓			S(C)			
↓			R(C)			
↓		Wait for X(C)			<b>T3</b>	<b>T4</b>
↓				Wait for X(B)		
↓			Wait for X(A)			



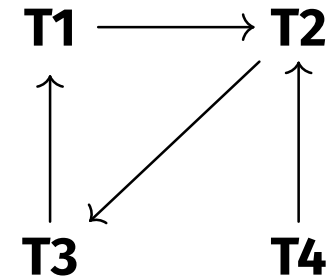
Time	T1	T2	T3	T4
↓	S(A)			
↓	R(A)			
↓		X(B)		
↓		W(B)		
↓	Wait for S(B)			
↓			S(C)	
↓			R(C)	
↓		Wait for X(C)		
↓				Wait for X(B)
↓			Wait for X(A)	



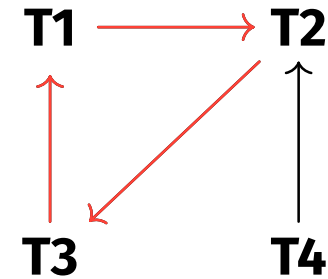
Time	T1	T2	T3	T4
↓	S(A)			
↓	R(A)			
↓		X(B)		
↓		W(B)		
↓	Wait for S(B)			
↓			S(C)	
↓			R(C)	
↓		Wait for X(C)		
↓				Wait for X(B)
↓			Wait for X(A)	



Time	T1	T2	T3	T4
↓	S(A)			
↓	R(A)			
↓		X(B)		
↓		W(B)		
↓	Wait for S(B)			
↓			S(C)	
↓			R(C)	
↓		Wait for X(C)		
↓				Wait for X(B)
↓			Wait for X(A)	



Time	T1	T2	T3	T4
↓	S(A)			
↓	R(A)			
↓		X(B)		
↓		W(B)		
↓	Wait for S(B)			
↓			S(C)	
↓			R(C)	
↓		Wait for X(C)		
↓				Wait for X(B)
↓			Wait for X(A)	



A cycle is a set of transactions that will never finish.

1. Periodically check for cycles in the Waits For Graph

1. Periodically check for cycles in the Waits For Graph
2. Abort one transaction from each cycle until there are no more

**Cycle detection is  
expensive!  $O(N + E)$**

**Idea:** Define an order over objects and only allow locks to be acquired in order

Only allow locks to be acquired in alphabetical order

<b>Time</b>	<b>T1</b>	<b>T2</b>	<b>T3</b>	<b>T4</b>
↓	S(A)			
↓	R(A)			
↓		X(B)		
↓		W(B)		
↓	Wait for S(B)			
↓			S(C)	
↓			R(C)	
↓		Wait for X(C)		
↓				Wait for X(B)
↓			Wait for X(A)	

Only allow locks to be acquired in alphabetical order

Time	T1	T2	T3	T4
↓	S(A)			
↓	R(A)			
↓		X(B)		
↓		W(B)		
↓	Wait for S(B)			
↓			S(C)	
↓			R(C)	
↓		Wait for X(C)		
↓				Wait for X(B)
↓			Wait for X(A)	

**T3** already has **C**, so it can't wait for **A**

Only allow locks to be acquired in alphabetical order

Time	T1	T2	T3	T4
↓	S(A)			
↓	R(A)			
↓		X(B)		
↓		W(B)		
↓	Wait for S(B)			
↓			S(C)	
↓			R(C)	
↓		Wait for X(C)		
↓				Wait for X(B)
↓			Wait for X(A)	

**T3** already has **C**, so it can't wait for **A**

Only allow locks to be acquired in alphabetical order

Time	T1	T2	T3	T4
↓	S(A)			
↓	R(A)			
↓		X(B)		
↓		W(B)		
↓	Wait for S(B)			
↓			S(C)	
↓			R(C)	
↓		Wait for X(C)		
↓				Wait for X(B)
↓			Wait for X(A)	

**T3** already has **C**, so it can't wait for **A**

Only allow locks to be acquired in alphabetical order

Time	T1	T2	T3	T4
↓	S(A)			
↓	R(A)			
↓		X(B)		
↓		W(B)		
↓	Wait for S(B)			
↓			S(C)	
↓			R(C)	
↓		Wait for X(C)		
↓				Wait for X(B)
↓			Wait for X(A)	

**T3** already has **C**, so it can't wait for **A**

Only allow locks to be acquired in alphabetical order

Time	T1	T2	T3	T4
↓	S(A)			
↓	R(A)			
↓		X(B)		
↓		W(B)		
↓	Wait for S(B)			
↓			S(C)	
↓			R(C)	
↓		Wait for X(C)		
↓				Wait for X(B)
↓			Wait for X(A)	

**T3** already has **C**, so it can't wait for **A**

**Variant Idea:** Acquire all locks at the start of the transaction

## **Pros**

- No deadlocks... ever
- No expensive cycle detection

## **Cons**

- Not all transactions can bound the set of objects they'll need upfront

**Idea:** (Rare) False positives are ok

1. Define an order over transactions as they arrive
2. Only allow newer transactions to block on older transactions

## Variant 1

## Variant 1

**T1 holds a lock on A**

T2 tries to acquire a lock on A and blocks

## Variant 1

### **T1 holds a lock on A**

T2 tries to acquire a lock on A and blocks

- Ok because T1 is “older”

## Variant 1

### **T1 holds a lock on A**

T2 tries to acquire a lock on A and blocks

- Ok because T1 is “older”

### **T2 holds a lock on A**

T1 tries to acquire a lock on A and blocks

## Variant 1

### **T1 holds a lock on A**

T2 tries to acquire a lock on A and blocks

- Ok because T1 is “older”

### **T2 holds a lock on A**

T1 tries to acquire a lock on A and blocks

- ABORT T1 and restart it as a “younger” transaction

## Variant 1

### **T1 holds a lock on A**

T2 tries to acquire a lock on A and blocks

- Ok because T1 is “older”

### **T2 holds a lock on A**

T1 tries to acquire a lock on A and blocks

- ABORT T1 and restart it as a “younger” transaction

“Wait-Die”

## Variant 2

## Variant 2

### **T1 holds a lock on A**

T2 tries to acquire a lock on A and blocks

- Ok because T1 is “older”

## Variant 2

### **T1 holds a lock on A**

T2 tries to acquire a lock on A and blocks

- Ok because T1 is “older”

### **T2 holds a lock on A**

T1 tries to acquire a lock on A and blocks

## Variant 2

### **T1 holds a lock on A**

T2 tries to acquire a lock on A and blocks

- Ok because T1 is “older”

### **T2 holds a lock on A**

T1 tries to acquire a lock on A and blocks

- ABORT T2 and restart it at the same age

## Variant 2

### **T1 holds a lock on A**

T2 tries to acquire a lock on A and blocks

- Ok because T1 is “older”

### **T2 holds a lock on A**

T1 tries to acquire a lock on A and blocks

- ABORT T2 and restart it at the same age

## Variant 2

### **T1 holds a lock on A**

T2 tries to acquire a lock on A and blocks

- Ok because T1 is “older”

### **T2 holds a lock on A**

T1 tries to acquire a lock on A and blocks

- ABORT T2 and restart it at the same age

“Wait-Wound”

## **Wait Die**

Abort older transactions that block on younger ones

## **Wait Wound**

Abort younger transactions if an older transaction would block on it

**Remember serializability  
containment**

**A schedule admitted by 2-Phase locking**

... is always ...

**A conflict serializable schedule**

... is always ...

**A serializable schedule**

(But the reverse is not true)