

# Introduction to Data Management

## CSE 344

### Lecture 21: More Transactions

# Announcements

- HW6, WQ6 due tonight
- HW7 will be released today
  - Some Java programming required
  - Connecting to SQL Azure
  - Due Wednesday, November 30
- WQ7 (final one!) released
  - Due Tuesday, November 29

# Outline

- Serial and Serializable Schedules (18.1)
- Conflict Serializability (18.2)
- Transaction implementation using locks (18.3)


# Review: Transactions

- **Problem:** An application must perform *several* writes and reads to the database, as a unit
- **Solution:** multiple actions of the application are bundled into one unit called a *Transaction*

# Turing Awards in Data Management



Charles Bachman, 1973  
*IDS and CODASYL*



Ted Codd, 1981  
*Relational model*



Jim Gray, 1998  
*Transaction processing*



Michael Stonebraker, 2014  
*INGRES and Postgres*

# Review: Transactions in SQL

**BEGIN TRANSACTION**  
[SQL statements]  
**COMMIT** or  
**ROLLBACK (=ABORT)**

[single SQL statement]

If BEGIN... missing,  
then TXN consists  
of a single instruction

# Review: ACID

- **Atomic**
  - State shows either all the effects of txn, or none of them
- **Consistent**
  - Txn moves from a state where integrity holds, to another where integrity holds
- **Isolated**
  - Effect of txns is the same as txns running one after another (i.e., looks like batch mode)
- **Durable**
  - Once a txn has committed, its effects remain in the database

# Isolation: The Problem

- Multiple transactions are running concurrently  
 $T_1, T_2, \dots$
- They read/write some common elements  
 $A_1, A_2, \dots$
- How do we prevent unwanted interference?
- The **SCHEDULER** is responsible for that



# Schedules

A **schedule** is a sequence of interleaved actions from all transactions

# Serial Schedule

- A serial schedule is one in which transactions are executed one after the other, in some sequential order
- Review: nothing can go wrong if the system executes transactions serially
  - But DBMS don't do that because we want better overall system performance

A and B are elements  
in the database  
t and s are variables  
in txn source code

## Example

DB element (rows, cell, table)

var

T1

T2

READ(A, t)

READ(A, s)

t := t+100

s := s\*2

WRITE(A, t)

WRITE(A,s)

READ(B, t)

READ(B,s)

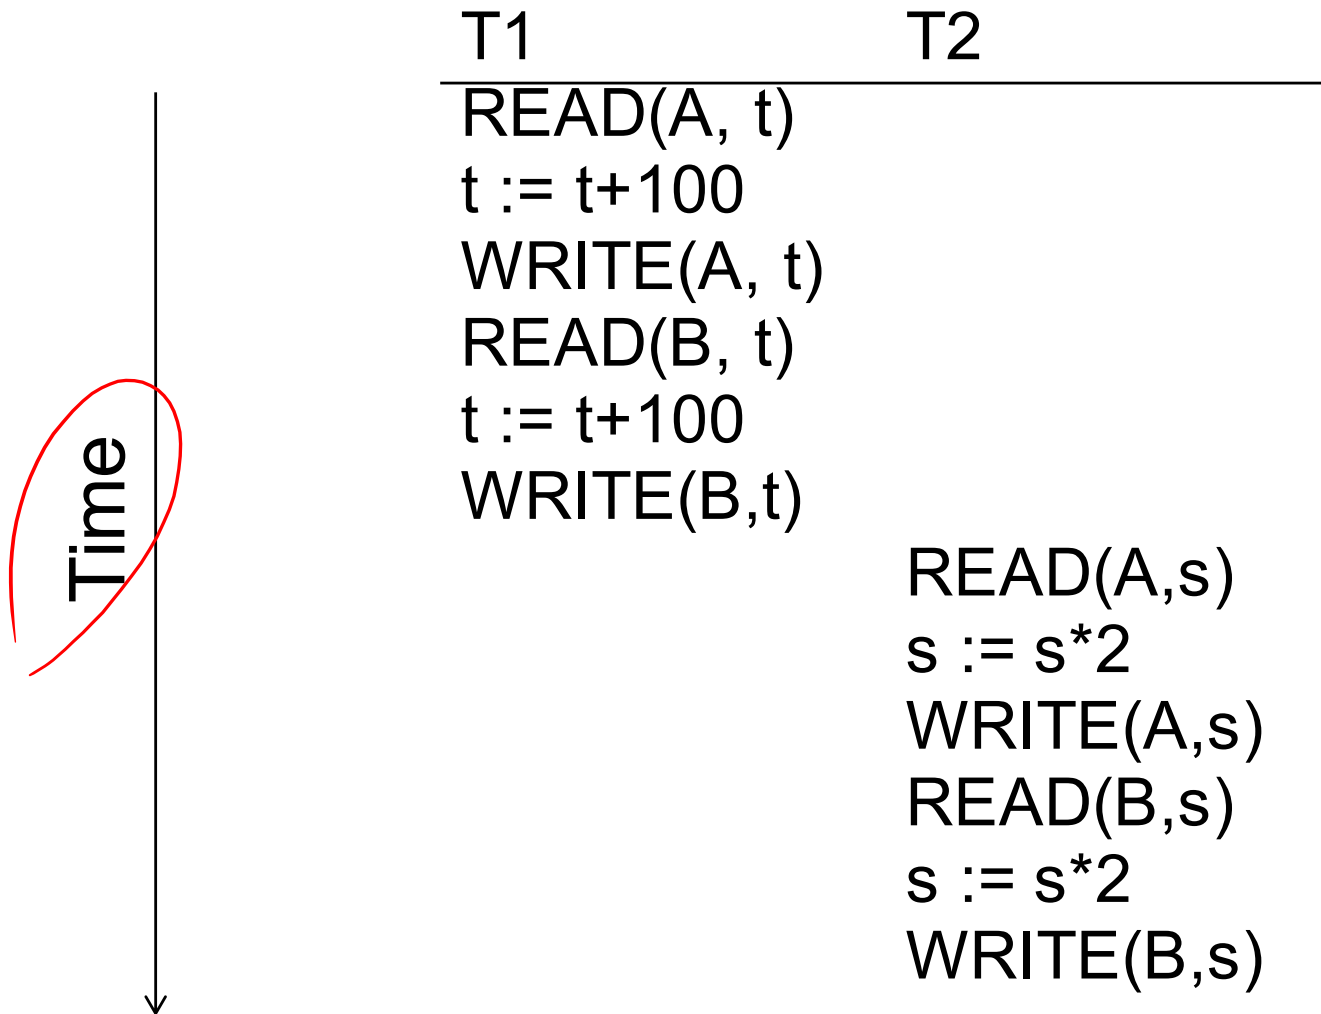
t := t+100

s := s\*2


WRITE(B,t)

WRITE(B,s)

# Example of a (Serial) Schedule



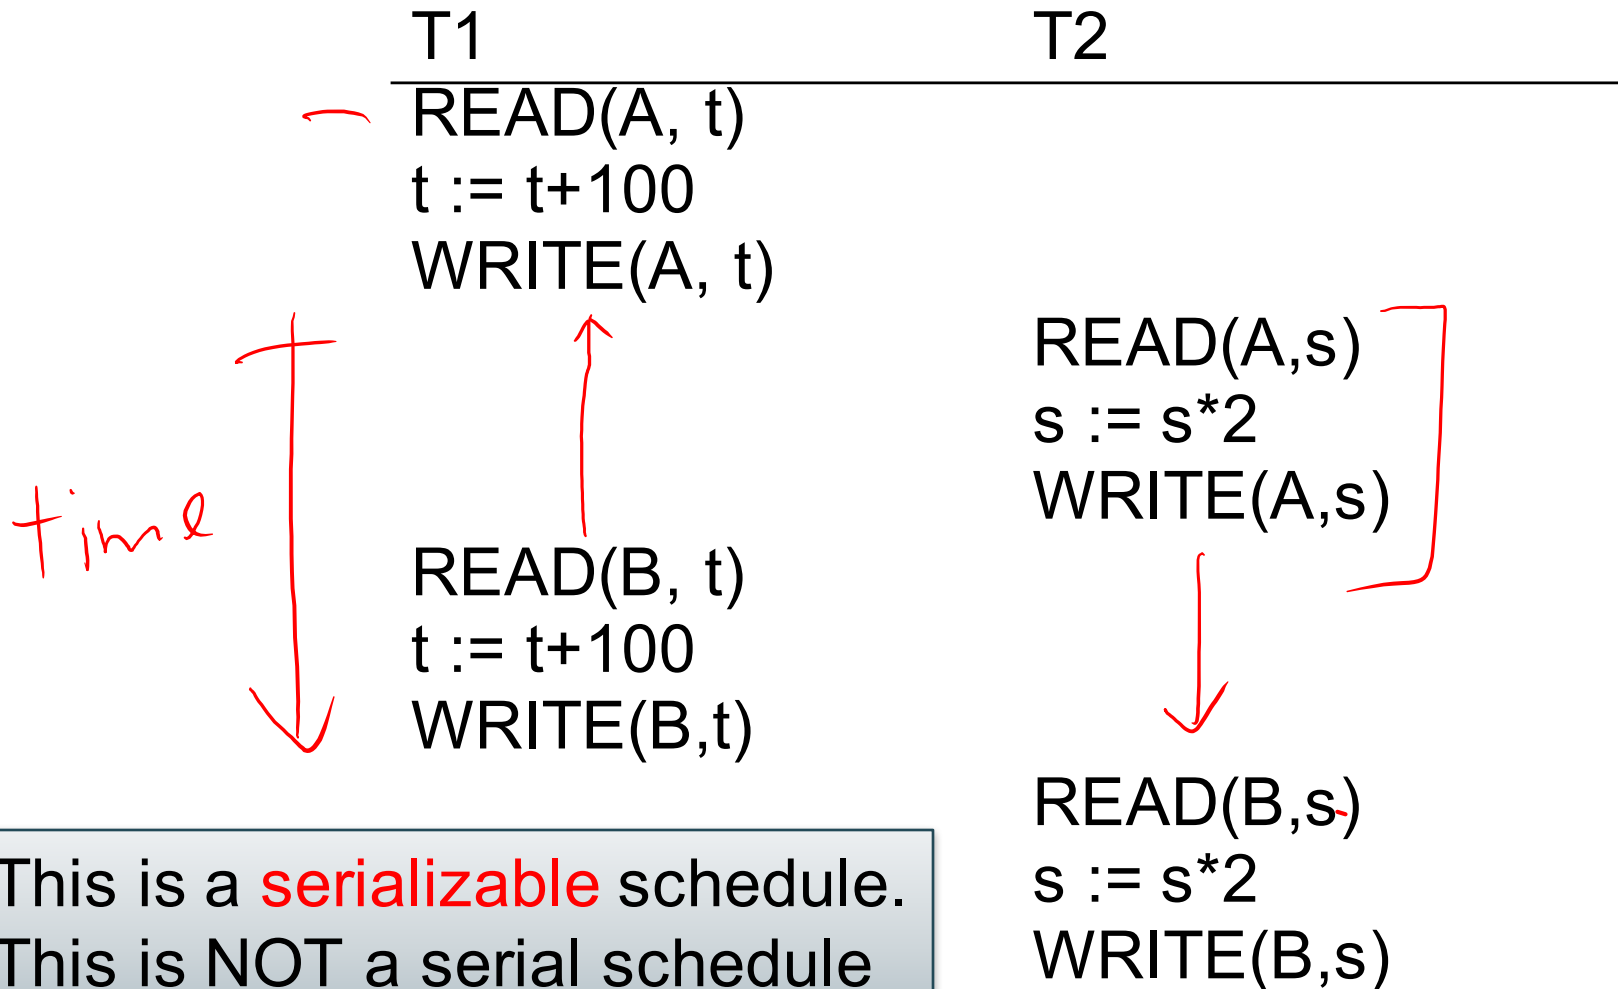
# Another Serial Schedule



# Serializable Schedule

A schedule is **serializable** if it is equivalent to a serial schedule

# A Serializable Schedule



# A Non-Serializable Schedule

T1	T2
READ(A, t)	
t := t+100	
WRITE(A, t)	
	READ(A,s)
	s := s*2
	WRITE(A,s)
	READ(B,s)
	s := s*2
	WRITE(B,s)
READ(B, t)	
t := t+100	
WRITE(B,t)	



# How do We Know if a Schedule is Serializable?

Notation:

$T_1: r_1(A); w_1(A); r_1(B); w_1(B)$   
 $T_2: r_2(A); w_2(A); r_2(B); w_2(B)$

Key Idea: Focus on *conflicting* operations

# Conflicts

- Write-Read – WR
- Read-Write – RW
- Write-Write – WW
- Read-Read?

# Conflict Serializability

**Conflicts:** (i.e., swapping will change program behavior)

Two actions by same transaction  $T_i$ :

$r_i(X); w_i(Y)$

Two writes by  $T_i, T_j$  to same element

$w_i(X); w_j(X)$

Read/write by  $T_i, T_j$  to same element

$w_i(X); r_j(X)$

$r_i(X); w_j(X)$

# Conflict Serializability

- A schedule is conflict serializable if it can be transformed into a serial schedule by a series of swappings of adjacent non-conflicting actions
- Every conflict-serializable schedule is serializable
- A serializable schedule may not necessarily be conflict-serializable

$w_1(x, 0); w_2(x, 0); R_1(x); R_2(x) \leftarrow$   
 $w_1(x, 0); R_1(x); w_2(x, 0); R_2(x) \leftarrow$

# Conflict Serializability

Example:

*time* →

$r_1(A); w_1(A); r_2(A); w_2(A); r_1(B); w_1(B); r_2(B); w_2(B)$

# Conflict Serializability

Example:

$r_1(A); w_1(A); r_2(A); w_2(A); r_1(B); w_1(B); r_2(B); w_2(B)$



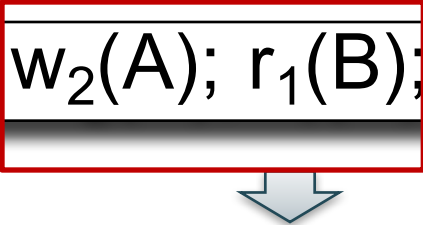
$r_1(A); w_1(A); r_1(B); w_1(B); r_2(A); w_2(A); r_2(B); w_2(B)$

*Serial  
schedule*

# Conflict Serializability

Example:

$r_1(A); w_1(A); r_2(A); w_2(A); r_1(B); w_1(B); r_2(B); w_2(B)$



$r_1(A); w_1(A); r_1(B); w_1(B); r_2(A); w_2(A); r_2(B); w_2(B)$

# Conflict Serializability

Example:

$r_1(A); w_1(A); r_2(A); w_2(A); r_1(B); w_1(B); r_2(B); w_2(B)$

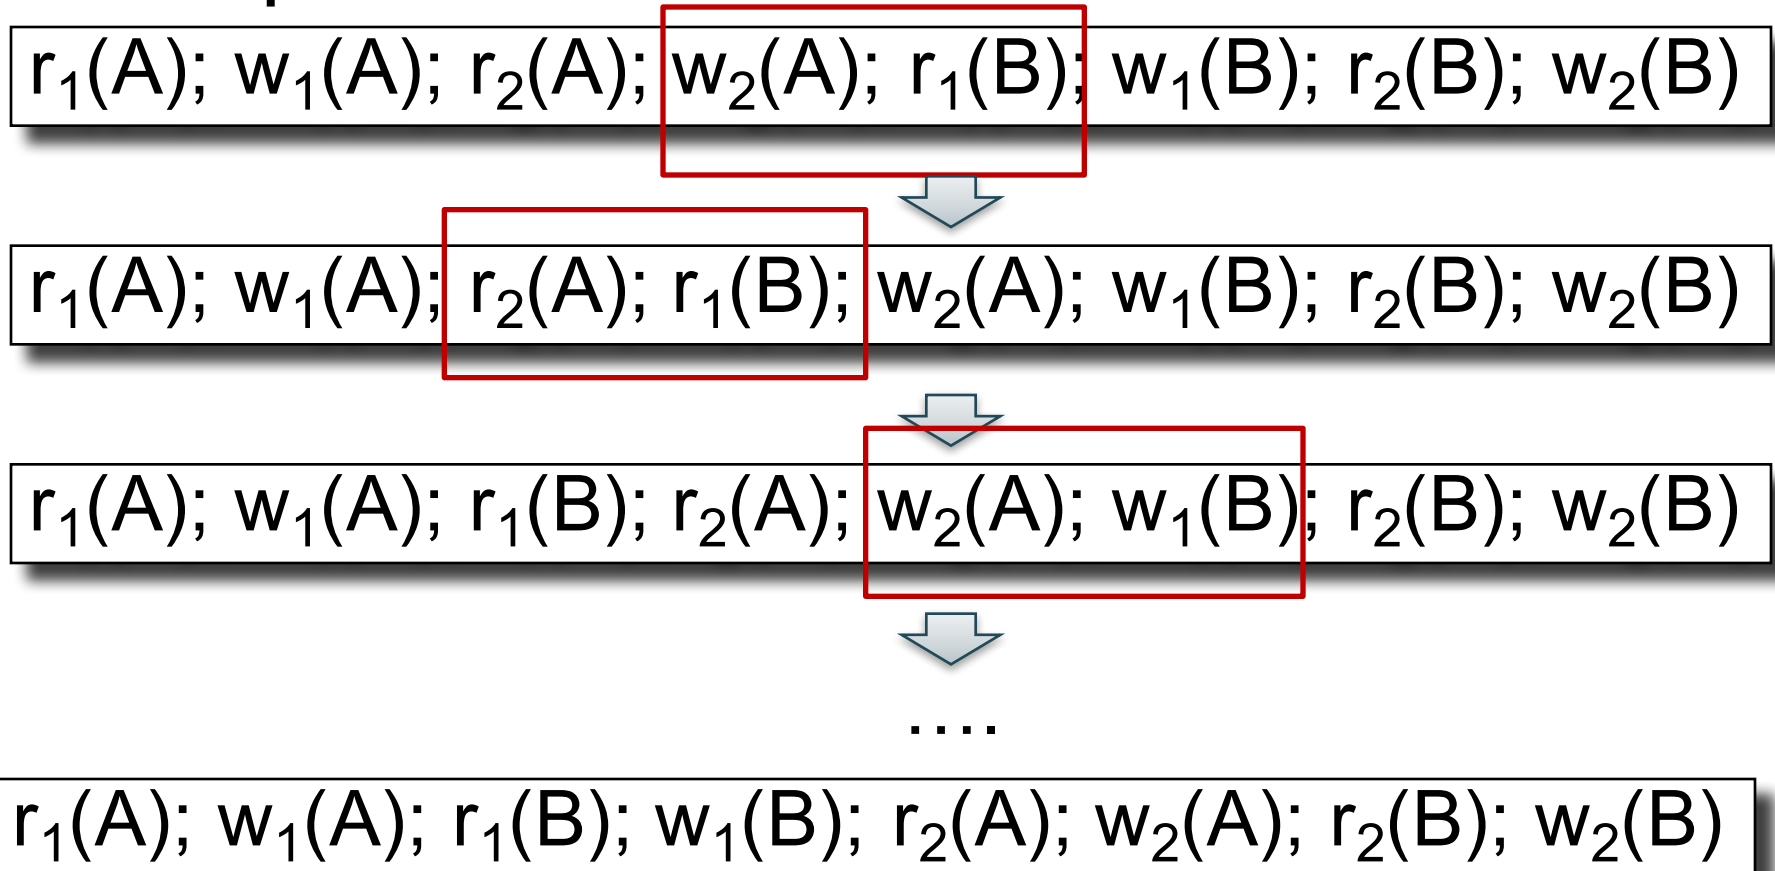
$r_1(A); w_1(A); r_2(A); r_1(B); w_2(A); w_1(B); r_2(B); w_2(B)$

$r_1(A); w_1(A); r_1(B); w_1(B); r_2(A); w_2(A); r_2(B); w_2(B)$



# Conflict Serializability

Example:



# Testing for Conflict-Serializability

## Precedence graph:

- A node for each transaction  $T_i$ ,
- An edge from  $T_i$  to  $T_j$  whenever an action in  $T_i$  conflicts with, and comes before an action in  $T_j$
- The schedule is conflict-serializable iff the precedence graph is acyclic

# Example 1

$r_2(A); r_1(B); w_2(A); r_3(A); w_1(B); w_3(A); r_2(B); w_2(B)$


①

②

③

# Example 1

$r_2(A); r_1(B); w_2(A); r_3(A); w_1(B); w_3(A); r_2(B); w_2(B)$



This schedule is **conflict-serializable**

# Example 2

$r_2(A); r_1(B); w_2(A); r_2(B); r_3(A); w_1(B); w_3(A); w_2(B)$


①

②

③

## Example 2

$r_2(A); r_1(B); w_2(A); r_2(B); r_3(A); w_1(B); w_3(A); w_2(B)$



This schedule **is NOT conflict-serializable**

# Scheduler

- **Scheduler** = the module that schedules the transaction's actions, ensuring serializability
- Also called **Concurrency Control Manager**
- We discuss next how a scheduler may be implemented

# Implementing a Scheduler

Major differences between database vendors

- **Locking Scheduler**
  - Aka “pessimistic concurrency control”
  - SQLite, SQL Server, DB2
- **Multiversion Concurrency Control (MVCC)**
  - Aka “optimistic concurrency control”
  - Postgres, Oracle

We discuss only locking schedulers in 344



# Locking Scheduler

Simple idea:

- Each element has a unique **lock**
- Each transaction must first **acquire** the lock before reading/writing that element
- If the lock is taken by another transaction, then wait
- The transaction must **release** the lock(s)

By using locks scheduler ensures conflict-serializability