

departamento de informática  
FACULDADE DE CIÊNCIAS E TECNOLOGIA  
UNIVERSIDADE NOVA DE LISBOA

# Concurrency Control (1)

Concurrency and Parallelism — 2017-18  
Masters in Computer Science  
(Mestrado Integrado em Eng. Informática)


Joao Lourenço <joao.lourenco@fct.unl.pt>

Base on slides from: [https://users.cs.duke.edu/~shivnath/courses/fall06/Lectures/11\\_serial.ppt](https://users.cs.duke.edu/~shivnath/courses/fall06/Lectures/11_serial.ppt)

# Concurrency Control

- Contents:

- Transactional model
- Serializability
- Conflicting operations



- Reading list:

- Chap 17 of Database management systems (3rd Ed.)  
McGraw-Hill Education  
Raghu Ramakrishnan, Johannes Gehrke  
ISBN: 0-07-123151-X


# Transaction

---

- Programming abstraction
- Implement real-world transactions
  - Banking transaction
  - Airline reservation

# Transaction: Programmer's Role

- Consistency




# Transaction: System's Role

---

- Atomicity
  - All changes of the transaction take effect or none at all
- Durability
  - All future transactions see the changes made by this transaction if it completes
- Isolation
  - Same effect as if the transaction executed in isolation

# Transaction: States



# Transactions

---


- Historical note:
  - Turing Award for Transaction concept
  - Jim Gray (1998)
- Interesting reading:

Transaction Concept: Virtues and Limitations  
by Jim Gray

<http://www.hpl.hp.com/techreports/tandem/TR-81.3.pdf>

# Issues with Concurrency: Example

Bank database: 3 Accounts



Property:  $A + B + C = 1500$

Money does not enter the system

Money does not leave the system

# Issues with Concurrency: Example

- Transaction T1: Transfer 100 from A to B

$A = 500, B = 500, C = 500$   $\longrightarrow$

Read (A, t)  
 $t = t - 100$   
Write (A, t)  
Read (B, t)  
 $t = t + 100$   
Write (B, t)

$A = 400, B = 600, C = 500$   $\longrightarrow$

# Issues with Concurrency: Example

---

- Transaction T2: Transfer 100 from A to C

Read (A, s)

$s = s - 100$

Write (A, s)

Read (C, s)

$s = s + 100$

Write (C, s)

Transaction T1	Transaction T2	A	B	C
Read (A, t)		500	500	500
$t = t - 100$				
Write (A, t)		400	500	500
	Read (A, s)			
	$s = s - 100$			
	Write (A, s)	300	500	500
Read (B, t)				
$t = t + 100$				
Write (B, t)		300	600	500
	Read (C, s)			
	$s = s + 100$			
	Write (C, s)	300	600	600

$$300 + 600 + 600 = 1500$$

Transaction T1	Transaction T2	A	B	C
Read (A, t)		500	500	500
$t = t - 100$				
	Read (A, s)			
	$s = s - 100$			
	Write (A, s)	400	500	500
Write (A, t)		400	500	500
Read (B, t)				
$t = t + 100$				
Write (B, t)		400	600	500
	Read (C, s)			
	$s = s + 100$			
	Write (C, s)	400	600	600

$$400 + 600 + 600 = 1600$$

# Terminology

---

- Schedule:
  - The exact sequence of (relevant) actions of one or more transactions

# Problems

---

- Which schedules are “correct”?
  - Mathematical characterization
- How to build a system that allows only “correct” schedules?
  - Efficient procedure to enforce correctness

# Correct Schedules: Serializability

---

- Initial database state is consistent
- Transaction:
  - consistent state → consistent state
- Serial execution of transactions:
  - Initial state → consistent state
- **Serializable schedule:**
  - A schedule equivalent to a serial schedule
  - Always “correct”

# Serial Schedule

		A	B	C
	Read (A, t)	500	500	500
	$t = t - 100$			
T1	Write (A, t)			
	Read (B, t)			
	$t = t + 100$			
	Write (B, t)	400	600	500
	Read (A, s)			
	$s = s - 100$			
	Write (A, s)			
T2	Read (C, s)			
	$s = s + 100$			
	Write (C, s)	300	600	600

$$300 + 600 + 600 = 1500$$

# Serial Schedule

T2

Read (A, s)

$s = s - 100$

Write (A, s)

Read (C, s)

$s = s + 100$

Write (C, s)

Read (A, t)

$t = t - 100$

Write (A, t)

Read (B, t)

$t = t + 100$


Write (B, t)

T1

	A	B	C
	500	500	500
	400	500	600
	300	600	600

$$300 + 600 + 600 = 1500$$

# Serial Schedule



# Is this Serializable?

Read (A, t)

$t = t - 100$

Write (A, t)



Read (B, t)

$t = t + 100$

Write (B, t)

Read (A, s)

$s = s - 100$

Write (A, s)



Read (C, s)

$s = s + 100$

Write (C, s)

Transaction T1

Transaction T2

# Equivalent Serial Schedule

Read (A, t)

$t = t - 100$

Write (A, t)

Read (B, t)

$t = t + 100$

Write (B, t)

Read (A, s)

$s = s - 100$

Write (A, s)

Read (C, s)

$s = s + 100$

Write (C, s)

Transaction T1

Transaction T2

# Is this Serializable?

Read (A, t)

$t = t - 100$

Write (A, t)

Read (B, t)

$t = t + 100$

Write (B, t)

Transaction T1

Read (A, s)

$s = s - 100$

Write (A, s)

Read (C, s)

$s = s + 100$

Write (C, s)

Transaction T2

No. In fact, it leads to inconsistent state

# Is this Serializable?

Read (A, t)

$t = t - 100$

Write (A, t)

Read (B, t)

$t = t + 100$

Write (B, t)

Read (A, s)

~~$s = s - 100$~~  0

Write (A, s)

Read (C, s)

~~$s = s + 100$~~  0

Write (C, s)

Transaction T1

Transaction T2

# Is this Serializable?

Read (A, t)

$t = t - 100$

Write (A, t)

Read (B, t)

$t = t + 100$

Write (B, t)

Transaction T1

Read (A, s)

$s = s - 0$

Write (A, s)

Read (C, s)

$s = s + 0$

Write (C, s)

Transaction T2

Yes, T2 is no-op

# Serializable Schedule

Read (A, t)

$t = t - 100$

Write (A, t)

Read (B, t)

$t = t + 100$

Write (B, t)

Read (A, s)

$s = s - 0$

Write (A, s)

Read (C, s)

$s = s + 0$


Write (C, s)

Transaction T1

Transaction T2

Serializability depends  
on code details

# Serializable Schedule



Still Serializable!

# Serializability

---

- General Serializability:
  - Hard to determine
- Goal: weaker serializability
  - Determined from database operations alone
- Database Operations:
  - Reads, Writes, Inserts, ...

# Simpler Notation

---

$r_T(X)$       Transaction T reads X

$w_T(X)$       Transaction T writes X

# What is $X$ in $r(X)$ ?

- $X$  could be any component of a database:
  - Attribute of a tuple
  - Tuple
  - Block in which a tuple resides
  - A relation
  - ...

# New Notation: Example Schedule

---

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




# Conflict Serializability

---

- Weaker notion of serializability
- Depends only on reads and writes


# Conflict Serializability

---




# Conflict Serializable Schedule

Transformations: swap **non-conflicting** actions



# Transformation: Example

r1(A) w1(A) r2(A) w2(A) r1(B) w1(B) r2(B) w2(B)



r1(A) w1(A) r2(A) r1(B) w2(A) w1(B) r2(B) w2(B)

# Non-Conflicting Actions

---

Two actions are **non-conflicting** if whenever they occur consecutively in a schedule, swapping them does not affect the final state produced by the schedule. Otherwise, they are **conflicting**.

# Conflicting Actions: General Rules

---

- Two actions of the same transaction conflict:
  - $r1(A) w1(B)$
  - $r1(A) r1(B)$
- Two actions over the same database element conflict, if one of them is a write
  - $r1(A) w2(A)$
  - $w1(A) w2(A)$

# Testing Conflict Serializability

---

- Construct **precedence graph**  $G$  for given schedule  $S$
- $S$  is conflict-serializable iff  $G$  is **acyclic**

# View Serializability

---

- A schedule  $S$  is view serializable if there exists a serial schedule  $S'$ , such that the source of all reads in  $S$  and  $S'$  are the same.

# View Serializability Example


---

View Serializable Schedule

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


# View Serializability Example

View Serializable Schedule



# View Serializability Example

View Serializable Schedule



---

Serial Schedule

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


# View Serializability Example

View Serializable Schedule

  
r2(B) w2(A) r1(A) r3(A) w1(B) w2(B) w3(B)

---

Serial Schedule

r2(B) w2(A) w2(B) r1(A) w1(B) r3(A) w3(B)  


# View Serializability Example

View Serializable Schedule


r2(B) w2(A) r1(A) r3(A) w1(B) w2(B) w3(B)



---

Serial Schedule

r2(B) w2(A) w2(B) r1(A) w1(B) r3(A) w3(B)



# View Serializability Example

View Serializable Schedule

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


Serial Schedule

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

**YES**

# View Serializability

---



# Problems


---

- Which schedules are “correct”?
  - Serializability theory

- How to build a system that allows only “correct” schedules?
  - Efficient procedure to enforce ~~correctness~~ serializable schedules

# Enforcing Serializability

Strategy:  
Prevent precedence  
graph cycles?



# Next class

---

- Enforcing serializability
  - Locking-based techniques
  - Timestamp-based techniques

# The END

---