

Database Management Systems

Madhavan Mukund

<https://www.cmi.ac.in/~madhavan>

Sai University

Lecture 20, 8 November 2023

Transactions — ACID properties

- Atomicity
- Consistency
- Isolation
- Durability

All or nothing

```
read(A);  
A := A - 50;  
write(A);  
read(B);  
B := B + 50;  
write(B).
```

2

1

AUDIT

↓

read(A)
read(B)

new A
old B

Atomicity vs isolation

Transactions — ACID properties

- Atomicity
- Consistency
- Isolation
- Durability

Keys

Referential integrity

Value checks ...

Invariant →

```
read(A);  
A := A - 50;  
write(A);  
read(B);  
B := B + 50;  
write(B).
```

Invariant →

Application constraints

AUDIT - Balance total must not change

Transactions — ACID properties

- Atomicity
- Consistency
- Isolation
- Durability

```
read(A);  
A := A - 50;  
write(A);  
read(B);  
B := B + 50;  
write(B).
```

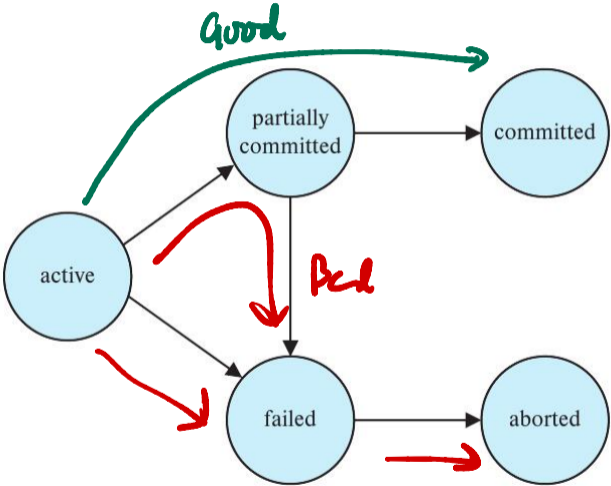
All updates on disk
persist

"Committed"

States of a transaction

- After abort, restart or kill
- Be careful about observable external writes

Dispense cash
Send SMS
⋮



Transaction logs

- Log each update **before** it happens
- Rollback updates in case of failure

Concurrent execution and schedules

T_1 : read(A);
A := A - 50; 1
write(A); 2
read(B); 4
B := B + 50; 6
write(B); 10

T_2 : read(A);
temp := A * 0.1; 3
A := A - temp; 5
write(A); 7
read(B); 8
B := B + temp; 9
write(B); 11

Schedule

Sequence of "low level" operations
How do ops interleave?

Concurrent execution and schedules

1000 A
2000 B

T_1 : read(A);
A := A - 50;
write(A);
read(B);
B := B + 50;
write(B).

only read & write - also inserts / delete

T_1	T_2
read(A)	
A := A - 50	
write(A)	
read(B)	
B := B + 50	
write(B)	
commit	
	read(A)
	temp := A * 0.1
	A := A - temp
	write(A)
	read(B)
	B := B + temp
	write(B)
	commit

950
2050

Explicitly note end of transaction

95

T_2 : read(A);
temp := A * 0.1;
A := A - temp;
write(A);
read(B);
B := B + temp;
write(B).

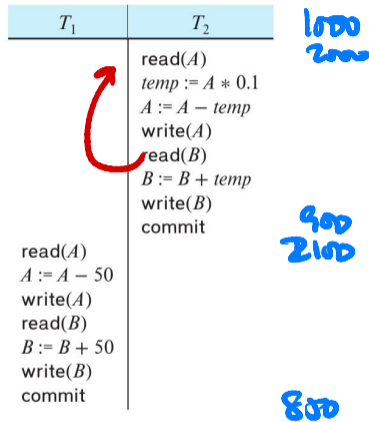
855
2145

Serial schedule 1

Concurrent execution and schedules

T_1 : read(A);
 $A := A - 50$;
write(A);
read(B);
 $B := B + 50$;
write(B).

T_2 : read(A);
 $temp := A * 0.1$;
 $A := A - temp$;
write(A);
read(B);
 $B := B + temp$;
write(B).



Serial schedule 2

Concurrent execution and schedules

T_1 : read(A);
 $A := A - 50$;
 write(A);
 read(B);
 $B := B + 50$;
 write(B).

T_2 : read(A);
 $temp := A * 0.1$;
 $A := A - temp$;
 write(A);
 read(B);
 $B := B + temp$;
 write(B).

T_1	T_2
read(A)	
$A := A - 50$	
write(A)	
read(B)	
$B := B + 50$	
write(B)	
commit	
	read(A)
	$temp := A * 0.1$
	$A := A - temp$
	write(A)
	read(B)
	$B := B + temp$
	write(B)
	commit

Serial schedule 1

equivalent

T_1	T_2
read(A)	
$A := A - 50$	
write(A)	
	read(A)
	$temp := A * 0.1$
	$A := A - temp$
	write(A)
read(B)	
$B := B + 50$	
write(B)	
commit	
	read(B)
	$B := B + temp$
	write(B)
	commit

Consistent concurrent schedule

Concurrent execution and schedules

T_1 : read(A);
 $A := A - 50$;
 write(A);
 read(B);
 $B := B + 50$;
 write(B).

T_2 : read(A);
 $temp := A * 0.1$;
 $A := A - temp$;
 write(A);
 read(B);
 $B := B + temp$;
 write(B).

T_1	T_2
read(A)	
$A := A - 50$	
write(A)	
read(B)	
$B := B + 50$	
write(B)	
commit	
	read(A)
	$temp := A * 0.1$
	$A := A - temp$
	write(A)
	read(B)
	$B := B + temp$
	write(B)
	commit

Serial schedule 1

T_1	T_2
read(A)	
$A := A - 50$	
write(A)	
	read(A)
	$temp := A * 0.1$
	$A := A - temp$
	write(A)
	read(B)
	$B := B + 50$
	write(B)
	commit
	read(B)
	$B := B + temp$
	write(B)
	commit

Consistent concurrent schedule

Interleave operations
Internally preserve order



Concurrent execution and schedules

T_1 : read(A);
 $A := A - 50$;
 write(A);
 read(B);
 $B := B + 50$;
 write(B).

T_2 : read(A);
 $temp := A * 0.1$;
 $A := A - temp$;
 write(A);
 read(B);
 $B := B + temp$;
 write(B).

T_1	T_2
read(A)	
$A := A - 50$	
write(A)	
read(B)	
$B := B + 50$	
write(B)	
commit	
	read(A)
	$temp := A * 0.1$
	$A := A - temp$
	write(A)
	read(B)
	$B := B + temp$
	write(B)
	commit

Serial schedule 1

T_1	T_2
read(A)	
$A := A - 50$	
	read(A) ? 1000
	$temp := A * 0.1$
	$A := A - temp$
	write(A) ← 950
	read(B)
950 → write(A)	
	2000
2050 → read(B)	
	write(A)
	$B := B + 50$
	write(B)
	commit
	$B := B + temp$
	write(B) ← 2150
	commit

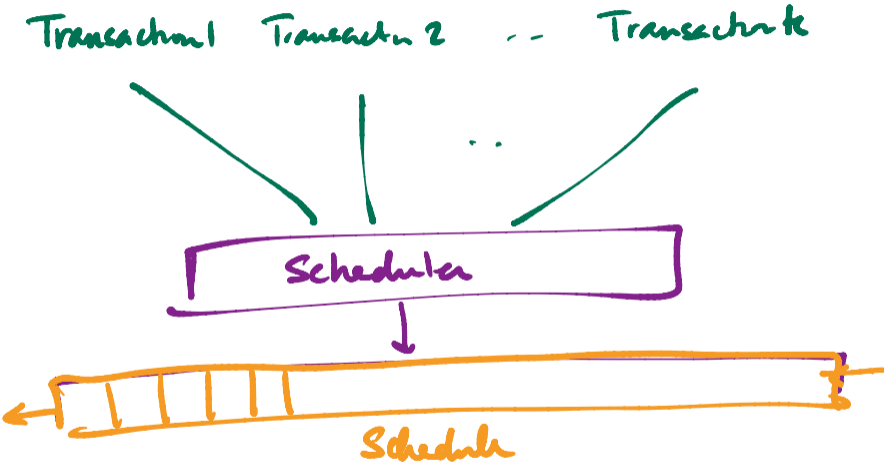
Inconsistent concurrent schedule

A concurrent schedule is consistent if it is equivalent to some serial schedule

Said to be Serializable

Given a concurrent schedule - is it serializable?

Serializability



Detecting if a schedule is serializable .

Problem? Interference of reads & writes of same value

Two concurrent reads do not interfere

$T_1: \text{write}(A)$
 $T_2: \text{read}(A)$ \neq $T_2: \text{read}(A)$
 $T_1: \text{write}(A)$ \parallel conflicts

Also

T1: write(A)

vs

T1: read(A)

T2: write(A)

overwrite

T2: read(A)

interchangeable

Conflicting pairs of actions

T1: read(A)

T1: write(A)

T2: write(A)

T2: read(A)

T1: write(A)

T2: write(A)

Conflict serializability



Swap adjacent operations that do not conflict

└ read(A) read(A)

└ read(A) write(B) -different data

Repeatedly swap non conflicting ops to get a serial schedule

Conflict serializability

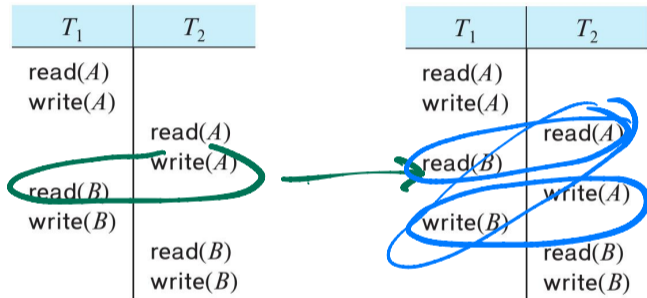
Transfer 50% Transfer 10%

T_1	T_2
read(A)	
write(A)	read(A)
read(B)	write(A)
write(B)	
	read(B)
	write(B)

conflict

? Interchangeable

Conflict serializability



Conflict serializability

T_1	T_2
read(A)	
write(A)	
	read(A)
	write(A)
read(B)	
write(B)	
	read(B)
	write(B)

T_1	T_2
read(A)	
write(A)	
	read(A)
read(B)	
	write(A)
write(B)	
	read(B)
	write(B)

3 swaps
→

T_1	T_2
read(A)	
write(A)	
read(B)	
write(B)	
	read(A)
	write(A)
	read(B)
	write(B)

Sufficient, not necessary

Serial

Testing for conflict serializability

Schedule

T1 —
T2 —
P3 —
T1 —
T1 —
T2 —
⋮
⋮
⋮

Conflict
Serializable?

Do not want
to try all
swaps!

Serial Schedule

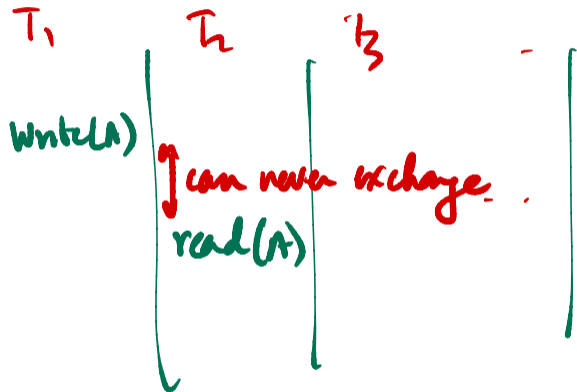
T_i happens before T_j

Write(A) influences read(A)

← If T1 writes (A) before
T2 read (A)

T1 must come before T2

Testing for conflict serializability

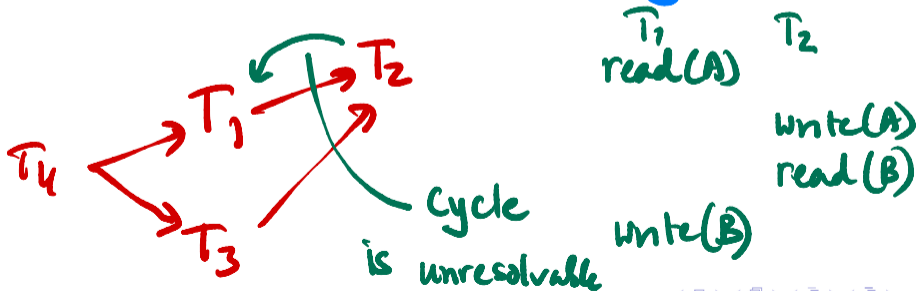


Record

$T_1 \rightarrow T_2$

Edge $T_i \rightarrow T_j$ if some earlier op in T_i conflicts some later op in T_j

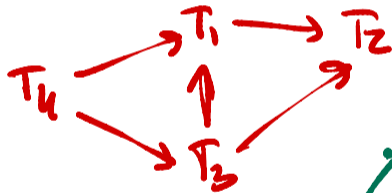
Combine all such constraints in a graph



Conflict graph G

- If G has a cycle - no consistent serial schedule exists
- If G is acyclic - it is conflict serializable

Directed acyclic graph



There must be a node with no incoming edge

Proof by contradiction

$T_0' \rightarrow T_2' \rightarrow T_1'$

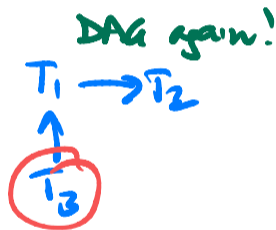
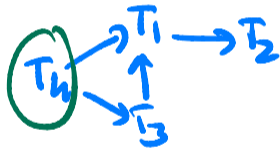
$T_i' \rightarrow T_{i+1}'$

But there a fixed number of nodes, n

Testing for conflict serializability

DAAG - directed acyclic graph

Enumerate a "starting" node & delete it



Any DAG can be "sequentialized" respecting dependencies

"Topological Sort"

Algorithms

1. Check if a directed graph has cycles
2. If DAG, generate a topological sort

Testing for conflict serializability

