

Database Management Systems

Madhavan Mukund

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

Sai University

Lecture 19, 3 November 2023

Application form - Assign a new App ID to each applicant

DB has the current max ID ?

20230001

20230002

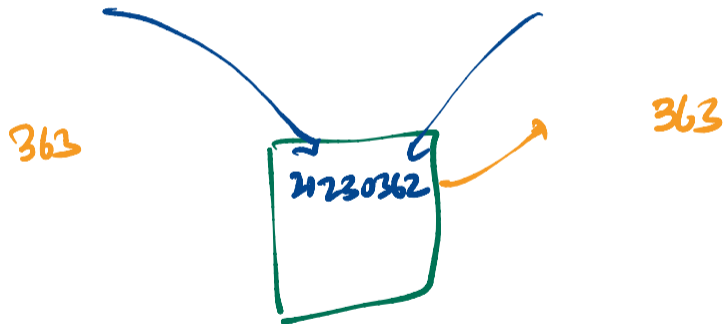
⋮

~~20230362~~

20230363

20230363

Two applicants access in parallel



read current- Max ID

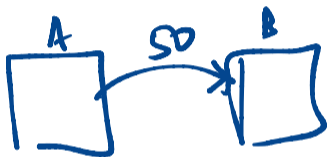
add 1

write back new Max ID

A unit

2 ops

Bank transfer



Reduce A's balance by 50
Increase B's balance by 50

UNIT

read ()

write ()

transfer from disk to memory

transfer from memory to disk

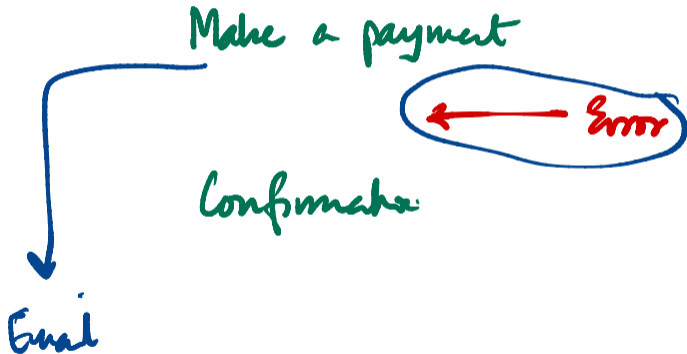
Transaction is a unit

All or nothing ...

Desirable properties

- Atomicity

All or nothing



Desirable properties

- Atomicity
- Consistency

$A \xrightarrow{SD} B$

read(A)
A = A - SD
write(A)
~~~~~  
read(B)  
B = B + SD  
write(B)

A+B does  
not  
change



# Desirable properties

- Atomicity
- Consistency
- Isolation

Transaction 1  
read (A)

$A = A - 50$

write (A)

---

read (B)

$B = B + 50$

write (B)

Audit Transaction

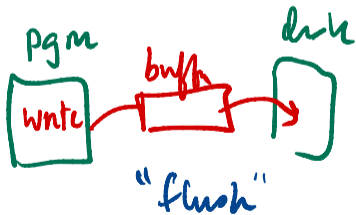
read (A) new

read (B) old

Behaviour affected  
by T1 being  
in parallel

# Desirable properties

- Atomicity
- Consistency
- Isolation
- Durability



Effect is permanent

Volatile & non volatile storage



Memory



Disk

"Committing" a transaction

RAID disks

# Desirable properties

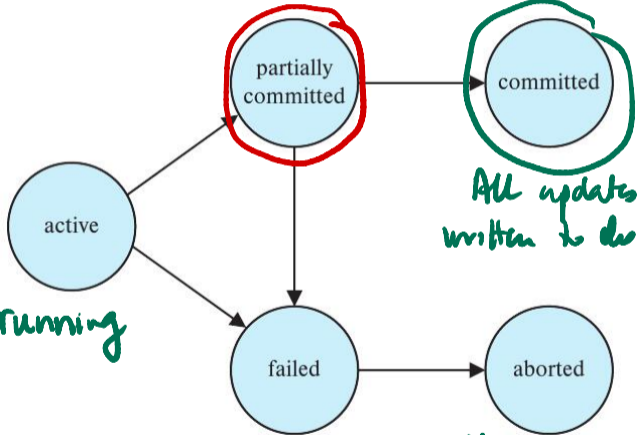
- Atomicity
- Consistency
- Isolation
- Durability
- ACID properties

# States of a transaction

Sequence of actions  
reads & writes

Abort  
↓  
Retry    Abandon  
Decided by "user"

Computation is over — All ops are done  
But data is not written at



How do you undo the effect of an aborted transaction?

Record each update

Where      Old Value      New Value      → On disk

When to log?

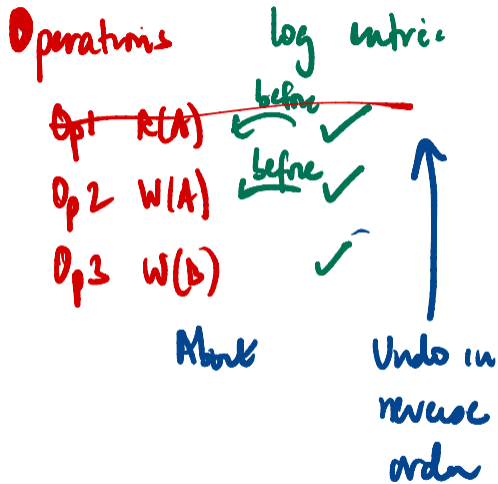
Crash!

① Update database

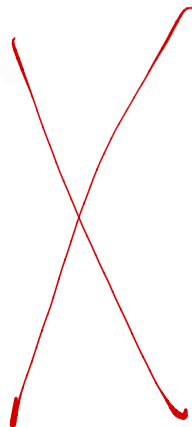
② Write log entry

Log before update

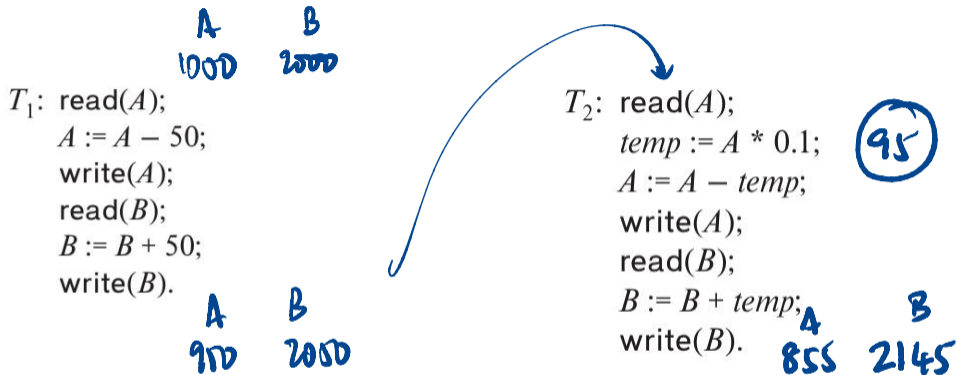
log ——— crash  
Update



Worry about  
external  
communications



# Concurrent execution and schedules



## Schedule

Order in which  
the operations  
are executed

All of  $T_1$ ,  
before all of  $T_2$

|            | $T_1$                                                                                  | $T_2$                                                                                                           |
|------------|----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| 1000, 2000 | read(A)<br>$A := A - 50$<br>write(A)<br>read(B)<br>$B := B + 50$<br>write(B)<br>commit |                                                                                                                 |
| 900, 1950  |                                                                                        | read(A)<br>$temp := A * 0.1$<br>$A := A - temp$<br>write(A)<br>read(B)<br>$B := B + temp$<br>write(B)<br>commit |
| 855, 1945  |                                                                                        |                                                                                                                 |



# Concurrent execution and schedules

All of  $T_2$  before  
all of  $T_1$

Two schedules have  
different outcomes

| $T_1$         | $T_2$                                              |
|---------------|----------------------------------------------------|
|               | read(A) <span style="color: red;">1000 2500</span> |
|               | $temp := A * 0.1$                                  |
|               | $A := A - temp$                                    |
|               | write(A)                                           |
|               | read(B)                                            |
|               | $B := B + temp$                                    |
|               | write(B) <span style="color: red;">900 2100</span> |
|               | commit                                             |
| read(A)       |                                                    |
| $A := A - 50$ |                                                    |
| write(A)      |                                                    |
| read(B)       |                                                    |
| $B := B + 50$ |                                                    |
| write(B)      |                                                    |
| commit        | <span style="color: red;">800 2150</span>          |

# Concurrent execution and schedules

This concurrent schedule  
"behaves" like  
 $T_1$  then  $T_2$

When is a concurrent  
schedule good?

|      | $T_1$         | $T_2$                     |
|------|---------------|---------------------------|
| 1000 | read(A)       |                           |
|      | $A := A - 50$ |                           |
| 950  | write(A)      |                           |
|      |               | read(A)                   |
|      |               | <u>temp := A * 0.1</u> 95 |
|      |               | $A := A - temp$           |
|      |               | write(A) 885              |
|      | read(B)       |                           |
|      | $B := B + 50$ |                           |
|      | write(B)      |                           |
| 2050 | commit        |                           |
|      |               | read(B)                   |
|      |               | $B := B + temp$           |
|      |               | write(B)                  |
|      |               | commit 2145               |

# Concurrent execution and schedules

Like my applicant ID  
problem

Inconsistent (bad)  
concurrent  
schedule

