

DBMS - IV

Transactions & Indexes

① Transactions

→ ACID

→ InnoDB - ACID

② Indexes

→ Type of indexes

→ Implementation of indexes

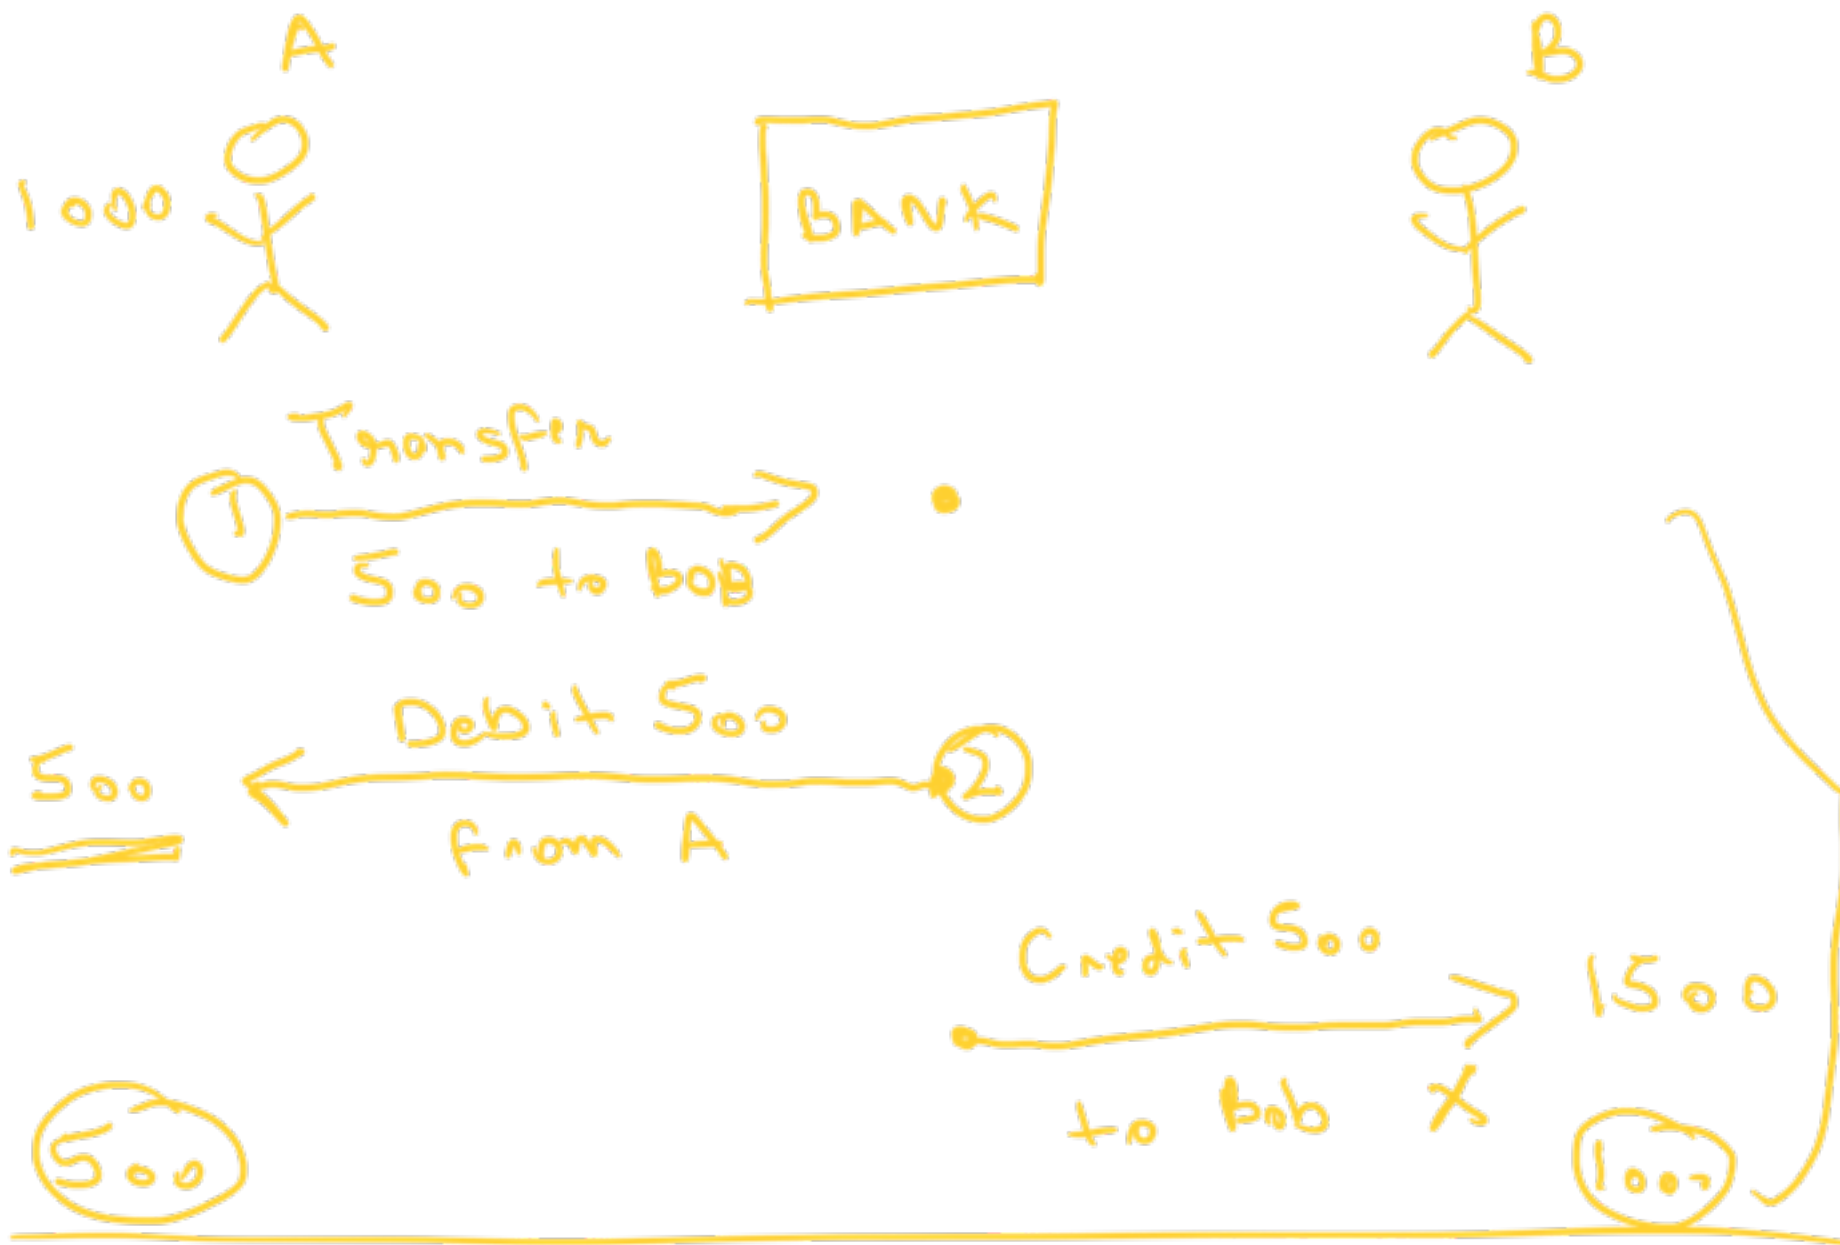
Transactions

Banking

- Alice - 1000

- Bob - 1000

Alice $\xrightarrow{500}$ Bob



① low balance

- ② One operation fails
- ③ Server crash

all or none - atomic

group of related operations
that should be executed as
one unit of work

→ Debit } Transaction
→ Credit }

BMS

26A

→ Book a seat

→ M.h. D.

} all or none

Make Payment

Transaction S - ACID

① Atomic

→ all or none

DONE vs NOT-DONE

- Simple - undivisible

→ BEGIN Transaction

→ Transfer from A to B

→ Debit from Alice

→ Credit to Bob

begin

Balance := read (ALICE)

if balance < 500:

ABORT

A = 10

T1

A = 20

rollback

ROLLBACK

A ⇒ 10

- Begin
- rollback

begin

balance = ALICE

DEBIT

CREDIT

COMMIT

A = 10

begin

A = 20

COMMIT

A = 20

→ begin
→ rollback } About

→ Commit)

ABORT - cancel the transaction

Rollback - revert changes
- keep the transaction alive

Bank

① HDFC
→ HDFC down
→ rollback

② ICICI → re run

ABORT + rollback → flag

END

HDFC

ICICI

A = 10

BEGIN

A = 20

error

roll back

A = 10

Students

ID	NAME
1	John
2	Mary

T1 = John Watson

T2 - John

3	Tony
---	------

T1 → A = 10

A = 20 ← committed values

T2

un
T1

Memory

1	John
	W.

1	John
---	------

Disk


ID	Name
1	John

Commit


commit

Shadow paging - atomically

↳ 80's - Couch DB



→ memory - copy of a copy



InnoDB - Write-ahead logging
- redo log



own time

— how to do an undo
update

↓
redo log

A = 20 — ① write in log

② update in error

if roll back ① read from
log how to
revert

WAL

ACID

↳ Atomicity

↳ all or none

↳ BEGIN

↳ ROLLBACK

}

↳ COMMIT ✓

→ Shadow paging

→ DB keeps a copy of the records in a transaction

→ Transaction uses the copy (shadow) while others use the main

→ WAL (redo log)

→ stores the operation, how to undo/redo

→ In case of failure, DB just uses the log to re-run the operations



Consistency



begin

A = "string" domain

1000 Alice → 1000 bob • 2000

begin

DEBIT → Alice (500)

Failure → 1500


before $A + B \Rightarrow 2000$
After $A + B \Rightarrow 1500$

2000 ① constraints are met
② some business validation
→ $A + B = A' + B'$

① referential integrity


② Application

Isolation



interfere with

each other
sum of all the
accounts



1000

1000

T1

Bank

T2

Debits 500 from A

A = 300
B = 1000

A + B = 1500 ?

Isolations

- helps us to manage how multiple transactions execute

A = 1000

B = 1000

T1

Bank

T2

A = A - 500

A = 500

B = 1000

sum
A + B = 1000
1000 + 1000
= 2000

Isolation levels



Issues

Phenomena

1 P1

Dirty read


↓
updated



PO - Dirty write

dirty bit

Dirty read (PI)



Dirty read

→ T2 needs uncommitted values

→ T1 - Reads balance as 1000

→ T1 → Updates balance = 500

→ T2 → Reads balance as 500

Dirty read → P1

PO - Dirty write

T1 → Reads bal. = 1000

T1 → Updates = 500

T2 → balance = 300

T1 → Commits { 3000 }

T2 → roll back

Dirty write

PO (w) PI (r)

P2 — Non-repeatable reads

P3 — Phantom reads

①

T1 — Read Bal = 1000

T2 — Updates Bal = 500

T1 — Read = ⑤00

T1 — 1000 } non-repeatable
— 500 } need

P3 - Phantom (Ghost)

- Insertion or deletion

T1 - select * from s = 100 rows

T2 - inserts 1 row
commit

T1 - select * from \Rightarrow 101

phantom read lol
aa

P0	- Dirty write	} uncommitted
P1	- Dirty read	
P2	- non-repeatable	- A = 1000 A' = 500
	- phantom	- low(N) + low(N')

Isolation level

- ① Read uncommitted - P0, P1, P2, P3
- ② Read committed - NO P0, P1
- ③ Repeatable reads - NO P0, P1, P2
- ④ Serializable
→ serially

T1

T2

T3

→ No issues (P0, P1, P2, P3)



↑
Concurrency VS performance
↓

Isolation levels

	PO
	P1
	P2
	P3
	AS ^{PS}

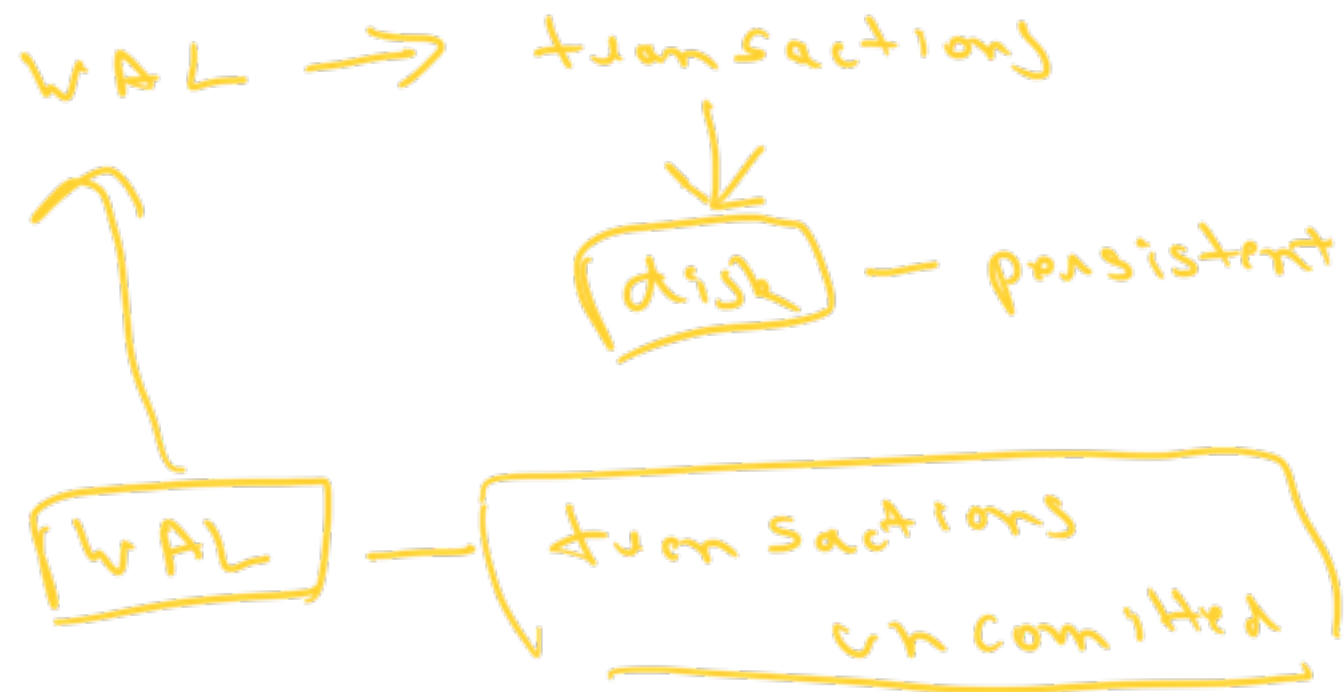
6:13 - 6:20
10:43 - 10:50

DB - isolation level
- transaction

in SQL - non table ...

My SQL - repeatable reads

SQL serve - Read Committed



-
- ① Cache
 - ② Isolation

Availability

DURABILITY

→ your results must be permanent

→ Begin
UPDATE

COMMIT

;
; ++ 10

DB crashes

permanent

- Disk

WAL
binary log

disk

log

binlog

→ log file

Update

↳ Store UPDATE A=10

↳ begin transaction

COMMIT ✗

→ bin log (status)

→ uncommitted changes

→ write to disk

→ rollback those operations

Distributed

→ distributed log file





binlog

transaction

WAL - redo log

log - periodically cleaned

ACID

Atomic - all or none

Consistency - constraints
- db state correct

Isolation - concurrent
transactions

↳ Durability

— permanent after
commit

DB — ACID compliant

RDBMS

NoSQL — BASE

eventually consistent

In desc

Book — Consensus


→ backtracking

↳ Page 1, 10, 20

→ Index

back tracking - Page 1, 10, 20

Database

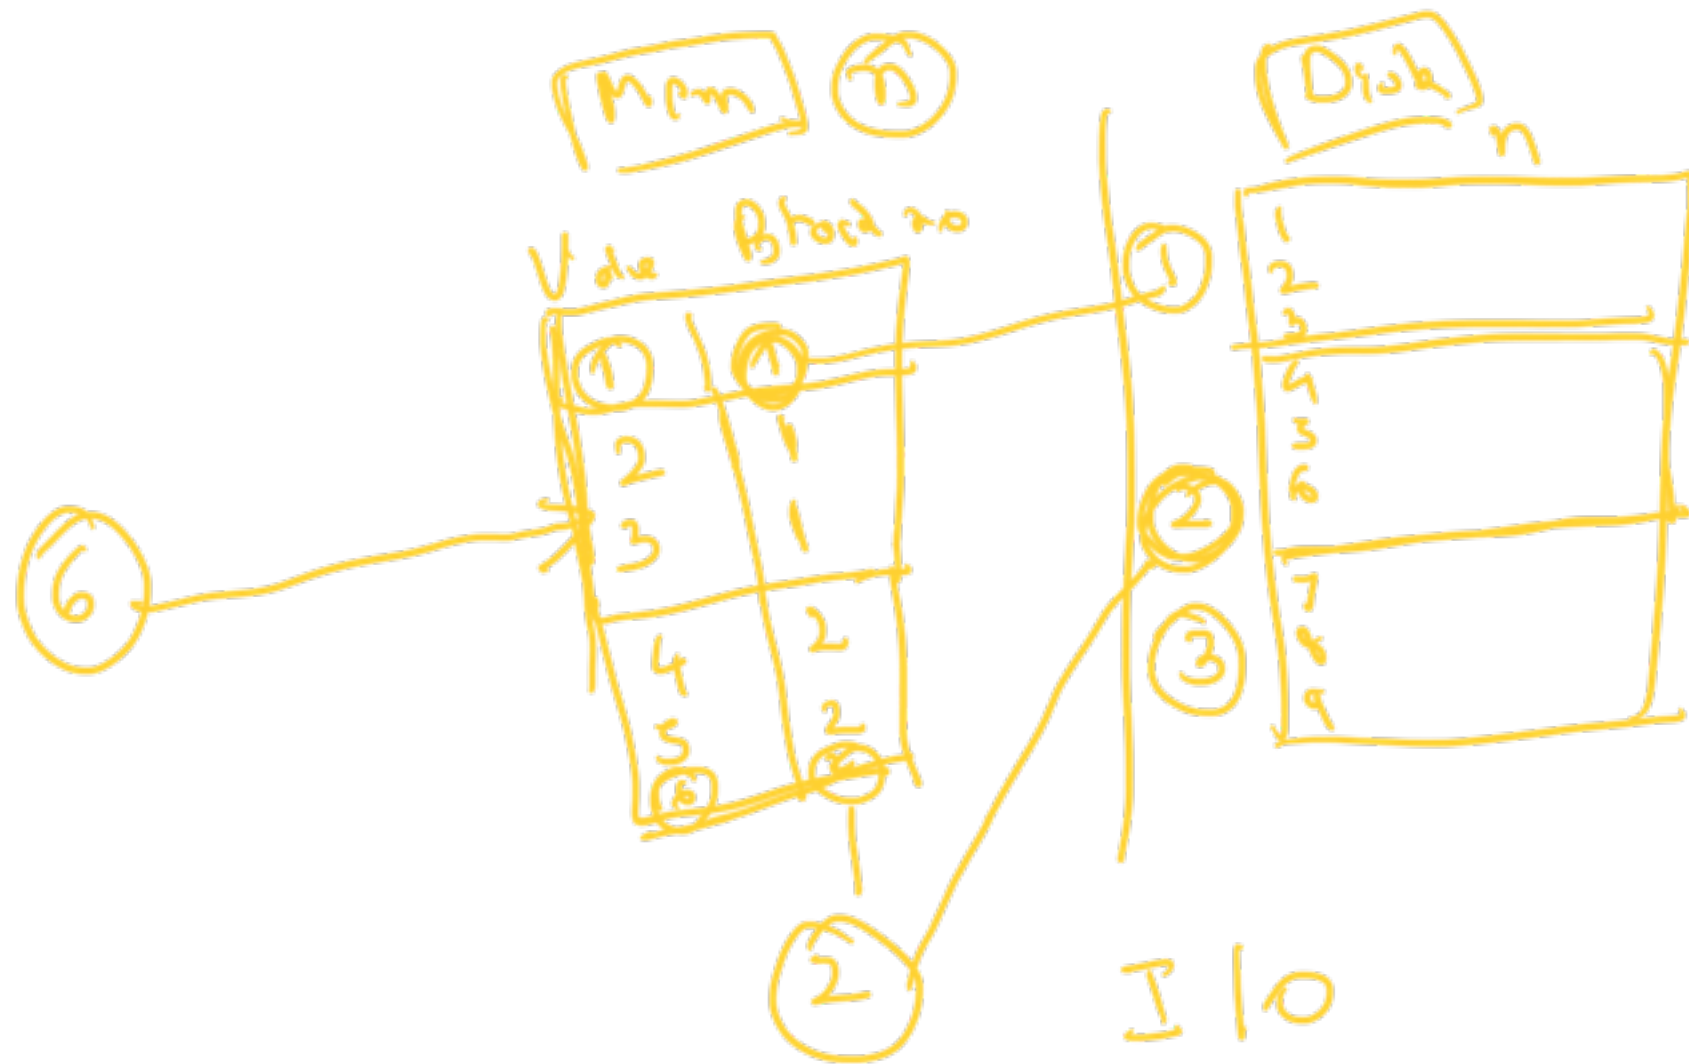


Sequential

Disk → RAM

I/O operation

extremely slow



Index - Data structure that helps us reduce I/O operations

Dense index

1:1 b/w value : address


→ large storage

Type of values

- unique ✓
- ordered ✓

⑥





	<u>unique</u>	non-unique
<u>Ordered</u>	Primary Index	clustered in dex
Unordered	Secondary	Secondary

PK - Primary In dex

In dex

→ reduce I/O

→ increase database

book keeping


→ write slower

→ Clustered

→ Secondary

Clustered → Ordered

→ non-unique





flag
- trace

is flag true?

if yes, fetch next block




Sequential - 4 I/O operations

A - 1

B - 2



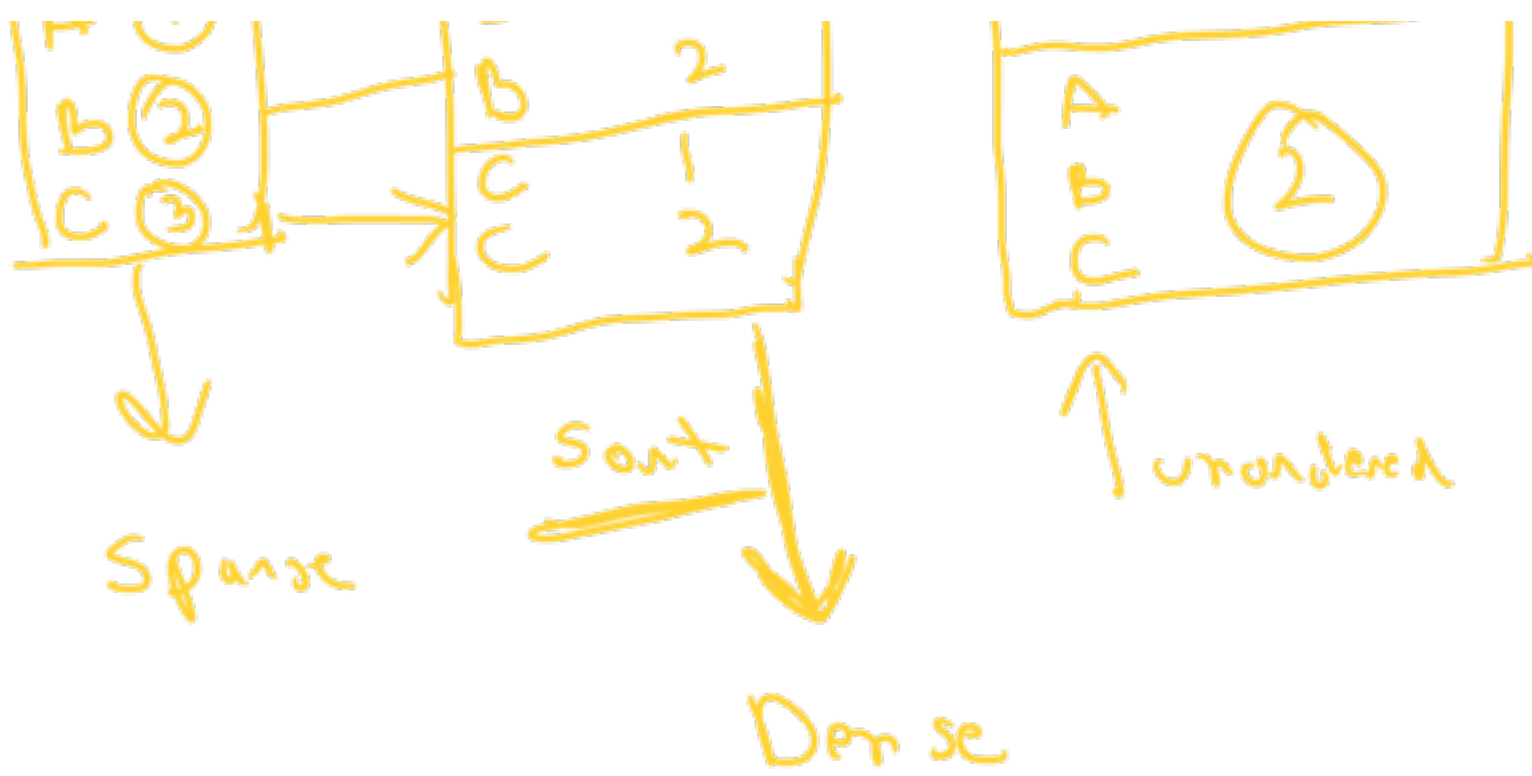


① In docs
reduce IO operations

Sec on dang

- unsorted
- non-unique





Sec on day

Types - Primary - unique sorted
 - Sparse
 starting value
 to address

- Clustered - 0 but not unique
 - Sparse + trace

- Secondary - unordered

- two levels

- 1st dense - sort

- 2nd sparse

- lookup

When to use an index

- PK

- select * from - where name
phone

When not to use an index

- null values

- small table

Transactions - ACID

Indexes - I/O operations

- Sparse / Dense

- Primary

- Secondary

- Clustered

Query execution

- MySQL

- Workbench

- Table plus






B+ trees



B+ trees

- ordered keys
- multiple children





balancing

Implementation details

A C I D

consistency

Alice \rightarrow Bob

A = 1000

Bob = 1000

Debit 500 from Alice] $A + B = 2000$

$A = 500$, $B = 1000$ - $A + B = 1500$

Failure

1500

2000

Correct state



business logic

constraints

should be met

$$1000 + 1000 = 2000$$

$$A + B = 1500$$

$$= A' + B'$$

int - (A) - int

① - Lock currency control

→ locks

- Read
- Read write
- range

Too many indexes

- ↳ large memory
- ↳ writes

Using indexes

FULLTEXT



data gets corrupted



hash()

1 ≠ 2 → corrupt

checksum