Subqueries, views and query optimisation

1. Sub-queries
   - → correlated
2. Views
   - → implementation
3. Query optimisation
   - → life cycle of a query
   - → EXPLAIN
   - → execution plan
   - → Scans

1. Window functions
2. Stored procedures
3. Triggers
4. Advanced — shadow paging
   — WAL
5. privileges

HLD — no-sql
— sharding

OS → Locking

Sub queries

→ I want all the students whose iq is > Average Iq

SID	Name	Iq
1	A	120
2	B	150
3	C	150

140

Select * From students

Select [illegible] X

where IQ > Avg(IQ)

① Get average IQ

② Use average IQ in where clause }

Sub queries

Two queries → Outer query

→ Inner query

② Select * from students
where ID →

(Select AVG(ID)
From students)

);

① Allows you to run multiple

(2) Intuitive + easy tone.d

5:57 - 6:02

- 10:32

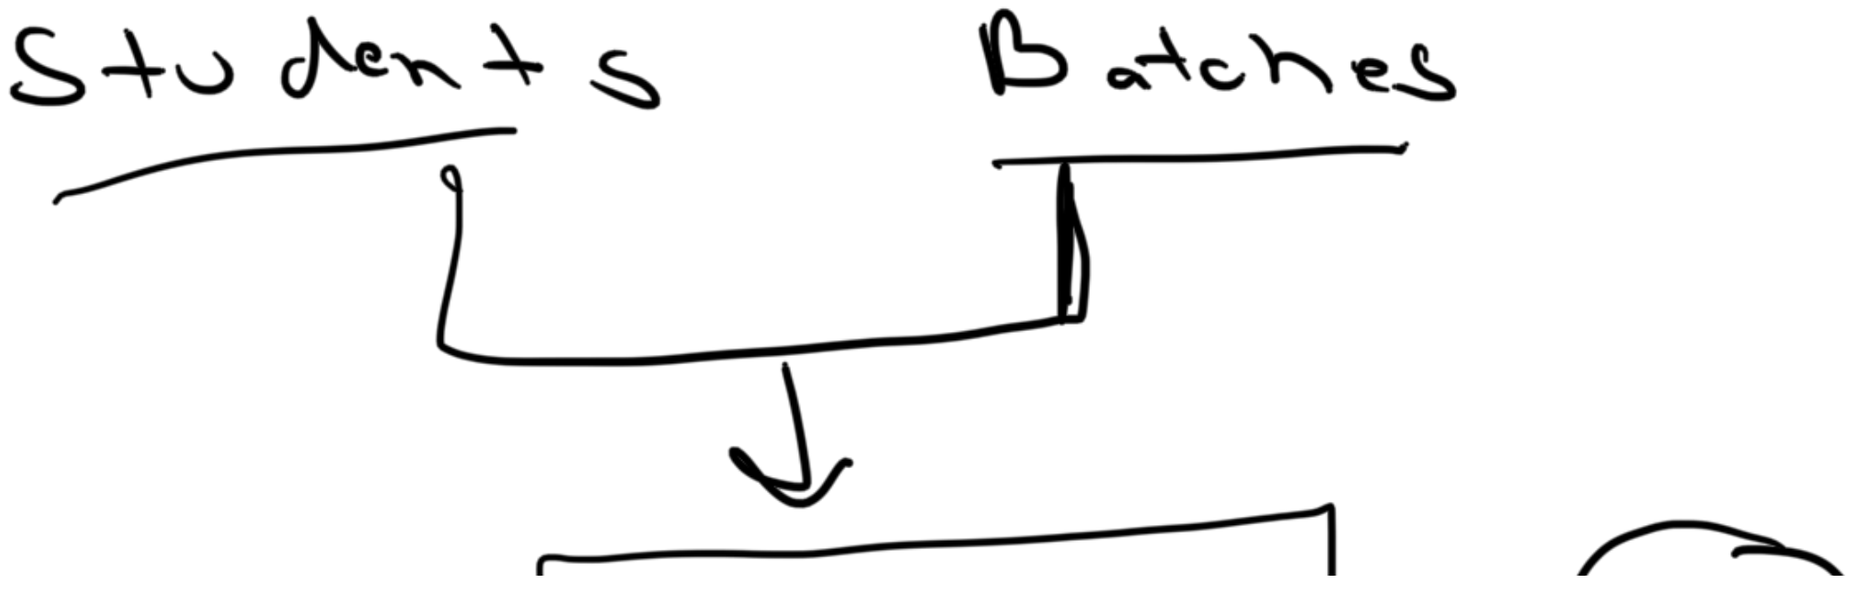

batch_id → FK

Indexes → PK

→ FK

Query → Application

→ FK

→ Index

→ Adhoc → no need if not frequent

Student S → id
→ email
→ phone
→ indexes

App → Newrelic
→ Sentry
→ Slow queries

Join → 10 times day

when → 1000

Student → name, phone

Where

→ name = ? AND phone = ?

phone = ? AND name = ?

(student.phone)

Student phone X

Student A phone = null

email, phone, name

email, phone

email

Views

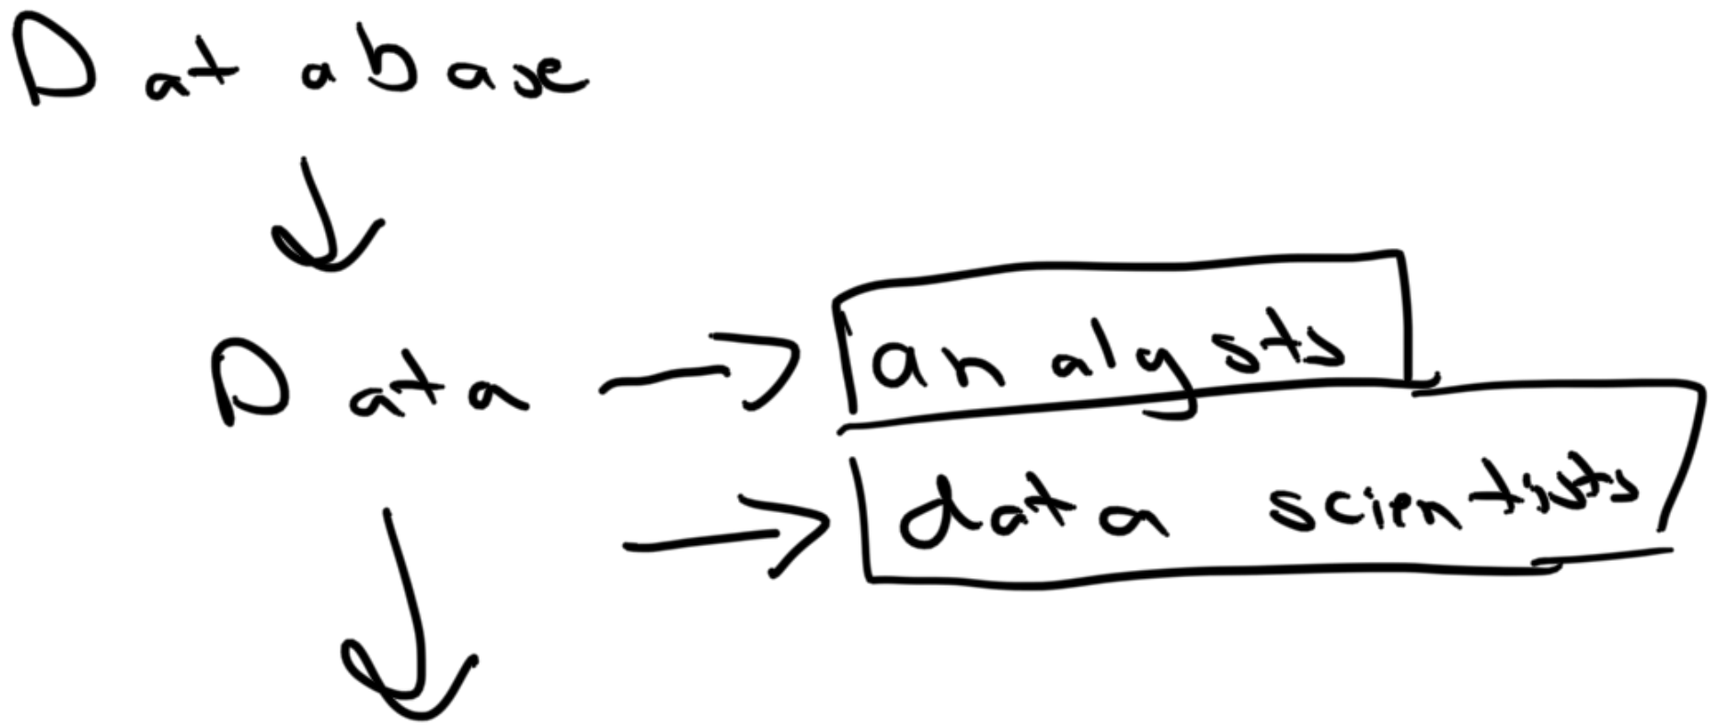

Excel

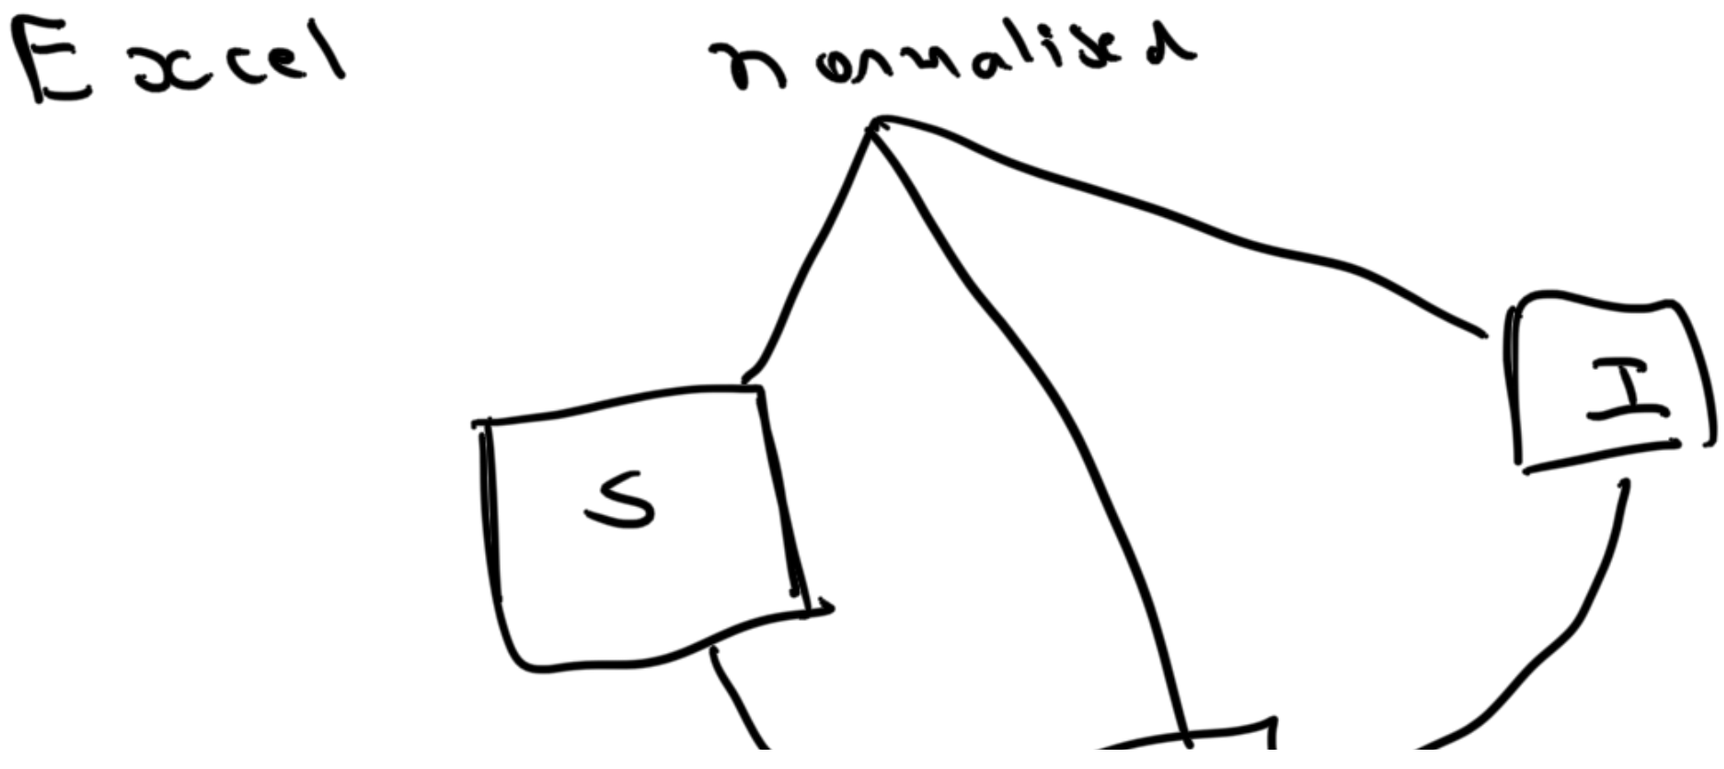

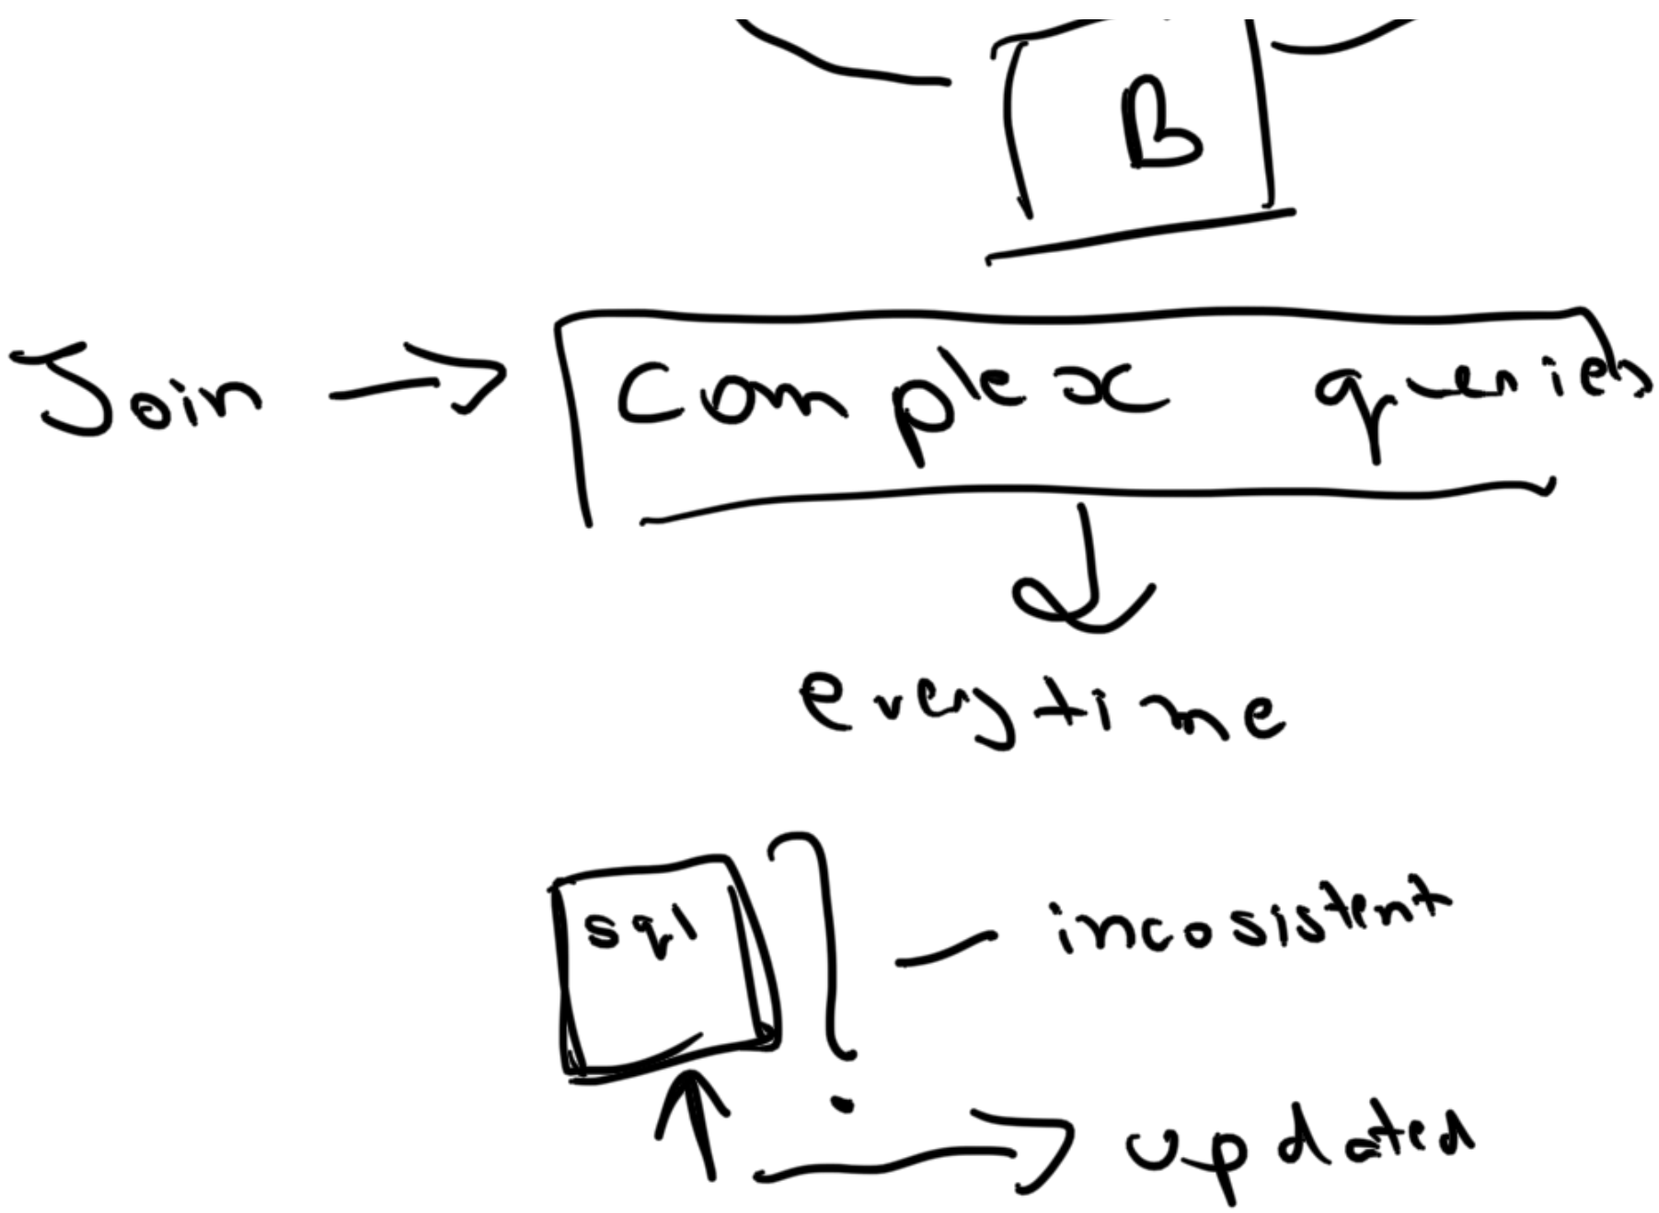

(2) Security | abstraction

→ If I want to get datas
I need to know about the

structure | schema.

(3) Denormalised

→ easy to work with

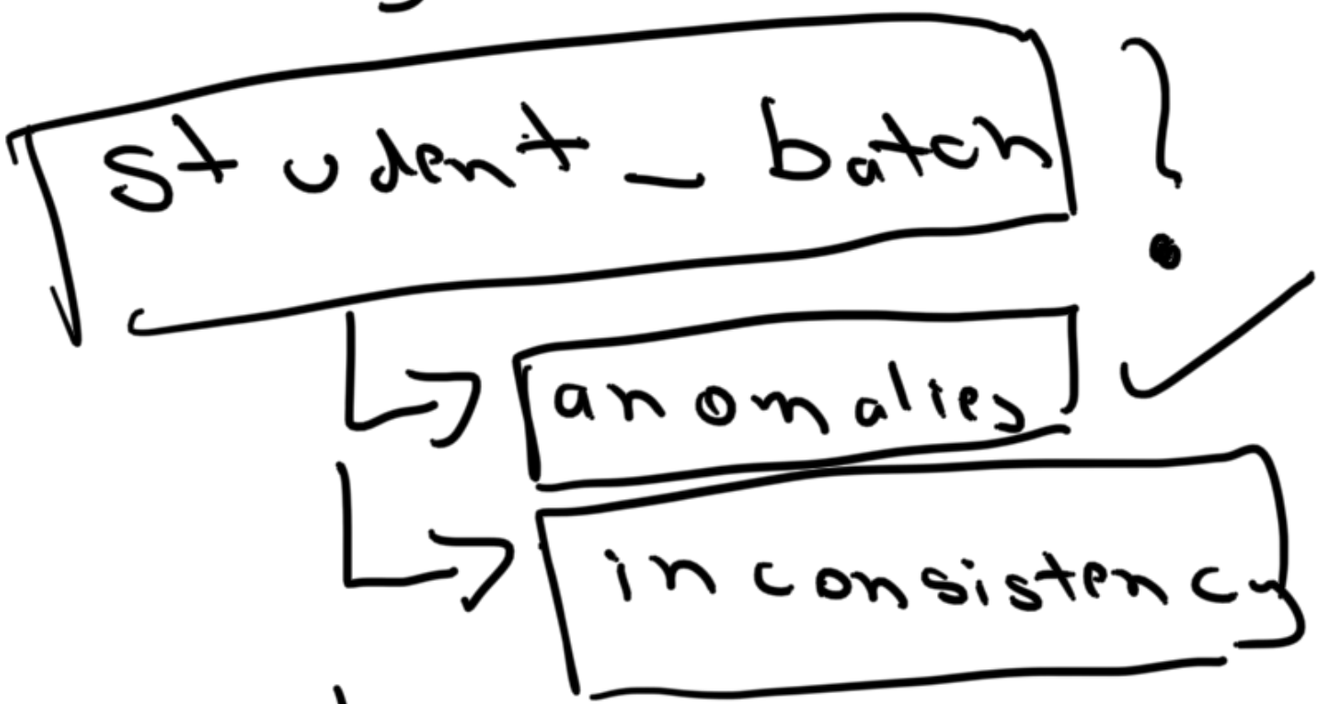

View

↳ a temp. place to store the results of a query

Select *

FROM students

JOIN

Batches

view

View → abstraction over your query

student_batch

ALIAS

FOR A QUERY

No need to

(1) Rewrite the query

(2) Hide the complexity

(3) Security → View → Third party

(4) Maintainability is improved

→ S_name, batch_name → batch

→ UPDATE VIEW

↓

RECREATE

Developer

↓

Creates the table

Data Analyst

↓

Creates the query

↓

Store query in a file

Data analyst

↳ Create the query

Developer

Maintain the database

→ abstraction for a query

→ result set from the query

↳ limitations

Views → CRUD

Reading

DDL , DML

CREATE VIEW NAME AS

Select ———

FROM ———

WHERE ———

i

Select * from view_name;

View

↳ Users

↳ Roles

→ DEV

→ 3P

→ Viewers.

↳ Tables

→ privileges

↳ DEV *

Role based access control

L7 3P VIEWS

---

* Normalised
* NOT easy to work with
* Join

* Run queries every time X
* How do we maintain! X

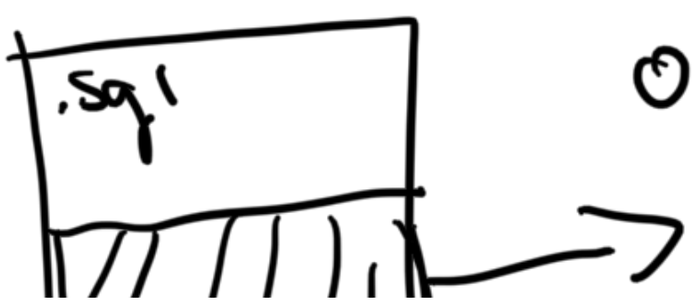

Out of sync

tied to

a version of schema

error

* Denorm table

→ waste of space

→ maintain the table

→ inconsistency

View

Application

DB

100

① Filtered rows

② Pagination

① WHERE

CROSS JOIN

WHERE

② Subqueries

FROM

(

WHERE

Join

10,000

LIMIT 20

OFFSET 20

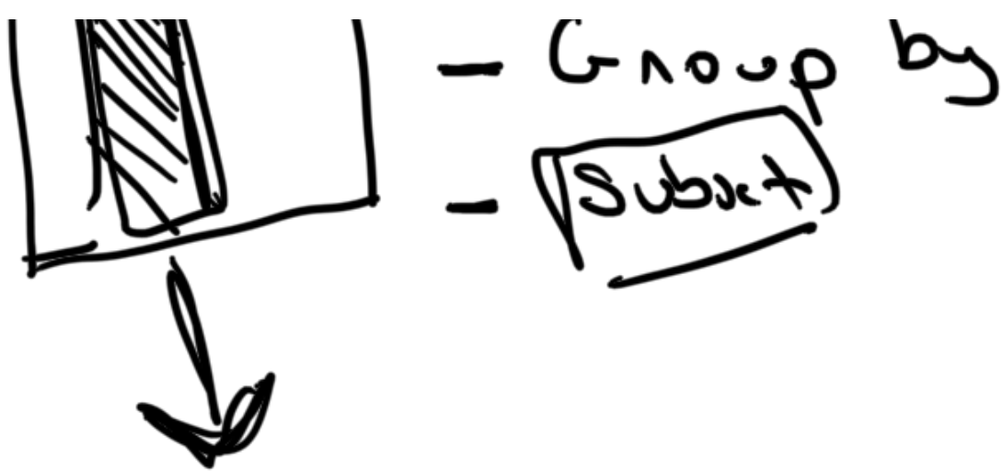
- Group by
- Subset