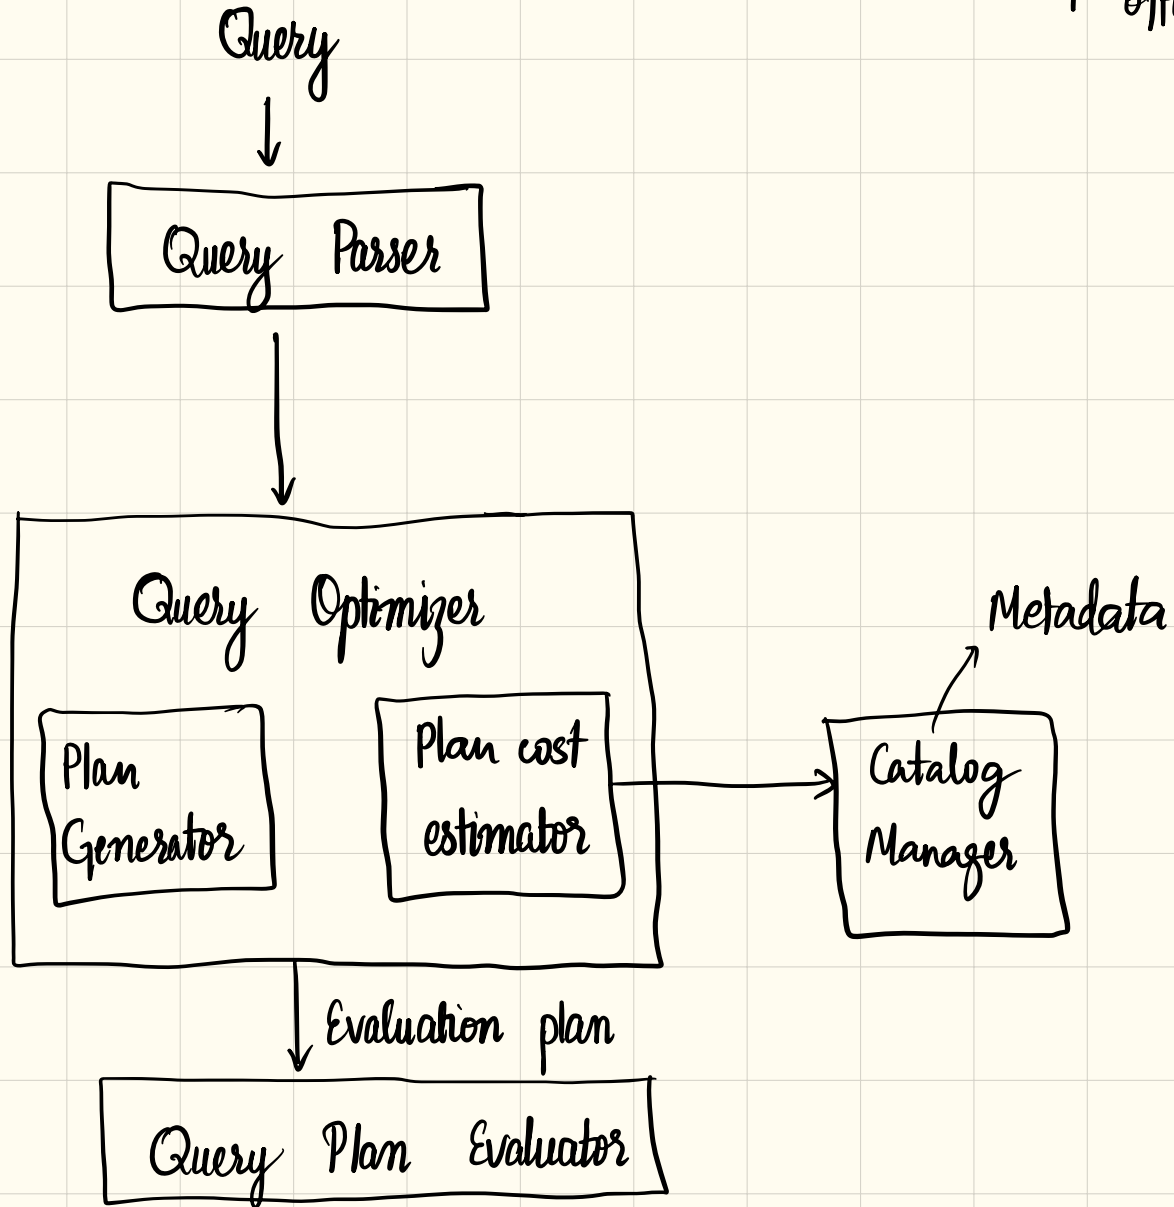


03 Mar 2025 - DBMS-II - Week 09

After break
1 offline class

Query Optimization



Query is essentially treated as a σ - π - γ algebra expression.

Query \rightarrow relational algebra



can you rewrite
so that output same? } Some
alternatives



Analyse cost and
choose the best


Optimization itself is a subject of study.

Query Evaluation Plans

Select S.name
 From Reserves R, Sailors S
 where R.sid = S.sid
 and R.bid = 100 and S.Rating > 5.

$\pi_{sname} (\sigma_{bid=100 \wedge rating > 5} (Reserves \bowtie_{sid=sid} Sailors))$

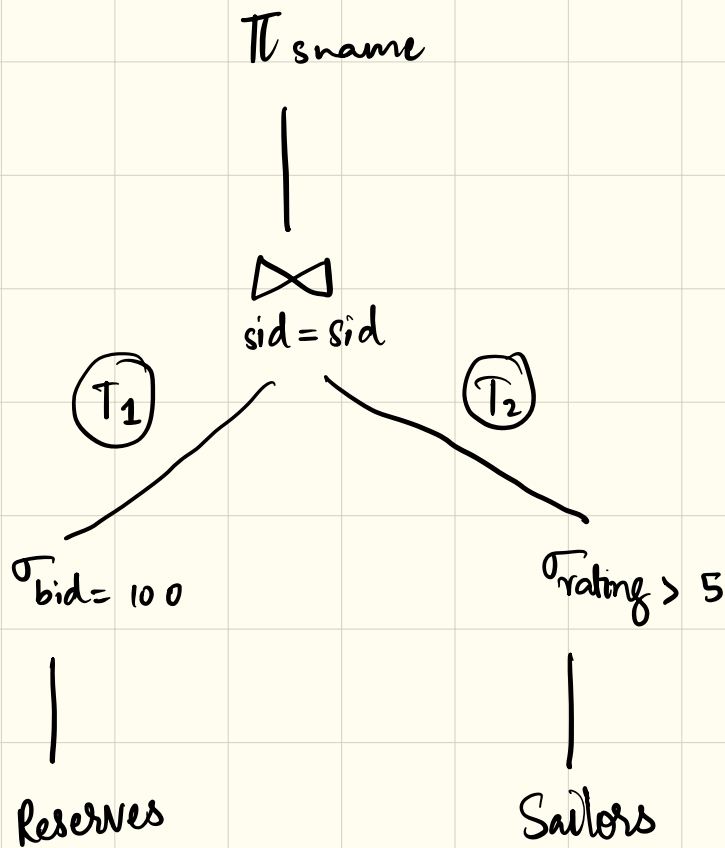
→ Parse tree



(B)

Scan:
write to
temp T₁

file scan




(A), (B) → same output

(B) may be better if no. of tuples in output is less (10%)
T₂ → may fit in the main memory

If $\sigma_{bid=100}$, $\sigma_{rating > 5}$ } → very less, then (B) good.

We push the selection operation below.



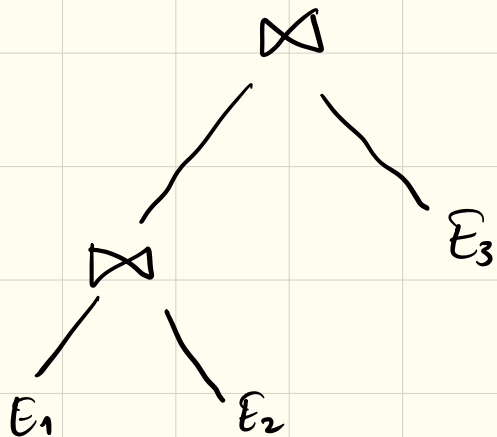

→ cannot be done always → there must be some rules when op is equivalent



$(M) \sim E_1 \bowtie E_2$
 if any one
 fits in
 Main Memory

$(N) \sim M < N$

$M + N$



$(E_1 \bowtie E_2) \bowtie E_3$

$E_1 \bowtie (E_2 \bowtie E_3)$

(T_1)

Suppose E_1 and E_2 are large, but $E_1 \bowtie E_2$ is small, then $(E_1 \bowtie E_2) \bowtie E_3$ might be easier.

Equivalence Rules

① $\sigma_{\theta_1 \wedge \theta_2}(E) = \sigma_{\theta_1}(\sigma_{\theta_2}(E))$ may be better if we have index on θ_2

Conjunctive selection

② $\sigma_{\theta_1}(\sigma_{\theta_2}(E)) = \sigma_{\theta_2}(\sigma_{\theta_1}(E))$

Selection is commutative

$$\textcircled{3} \quad \pi_{L_1} (\pi_{L_2} (\dots (\pi_{L_n} (E)) \dots)) = \pi_{L_1} (E)$$

Projection is also expensive (you need to duplicate)

$$\textcircled{6} \quad \text{a.} \quad \sigma_0 (E_1 \times E_2) = E_1 \bowtie_0 E_2$$

b.

$$\textcircled{7} \quad \text{a)} \quad \underbrace{\sigma_0 (E_1 \bowtie_0 E_2)}_{\substack{\text{may not} \\ \text{fit in} \\ \text{memory}}} = \underbrace{(\sigma_0 (E_1)) \bowtie_0 E_2}_B$$

$T_1 \rightarrow$ may be very small compared to E_1

\swarrow has attr only of A

$$\text{b)} \quad \sigma_{\theta_1 \wedge \theta_2} (E_1 \bowtie_0 E_2) = (\sigma_{\theta_1} (E_1)) \bowtie_0 (\sigma_{\theta_2} (E_2))$$

⑧

⑨

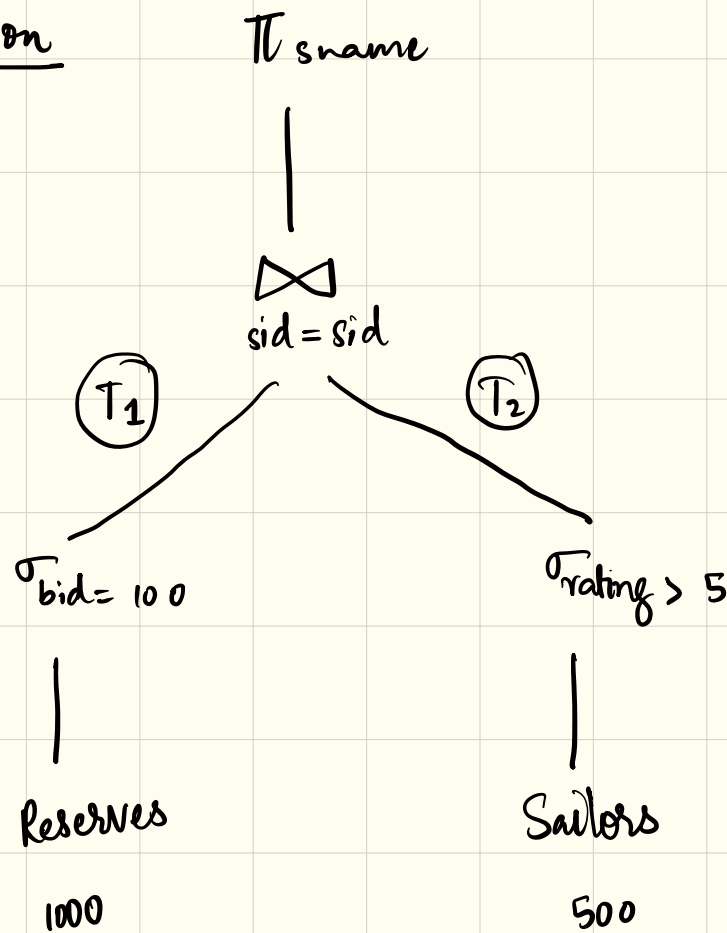
⑩

⑪



Main TB
easy rules

Query Optimization



Assume

No. of pages in

$$T_1 = 10$$

$$T_2 = 250$$

$$\text{buffer} = 5$$

Cost 1: applying sort merge join for T_1 and T_2 :
the cost of selection $(1000 + 10 + 500 + 250 = 1760)$
and the cost of the join $(40 + 2,000 + 260 = 2,300)$
 $= 4,060$ page I/Os.

2. Applying block nested join

$$1760 + 1010 = 2770 \text{ page I/Os}$$

sort merge join — better

block nested join X

If query is sorted go for
block nested join

skipped projection join.


Using Indexes

Π_{name}

|

$\sigma_{\text{rate} > 5}$

|



Indep nested loops
with pipelining.

hash index on
bid.
 $\sigma_{bid=100}$
Reserves.

Sailors hash index on sid

Compare.

07 Mar 2025

Transactions

- ① Advanced SQL
- ② Database Design
- ③ Storage, Indexing and Query evaluation
- ④ Transactions → now, last (3 more lectures)

A **transaction** is a unit of program that accesses and possibly updates various data items.

e.g. website → new customer
→ form → submit

Information may be inserted in multiple tables } → Concurrency issues

transaction complete \Rightarrow all records inserted

Suppose DB crashes while inserting Phone no.

\$50 from account A to acc. B:

1. read(A)

2. $A := A - 50$

3. write(A)

server crashes

4. read(B)

5. $B := B + 50$

6. write(B)

'assumption: von neumann architecture.

\rightarrow sequential execution in the processes

More problems: simultaneous transaction

1. Atomicity requirement
either 1-6 complete or nothing

Normally, each stored procedure in MySQL is a transaction

2. Durability requirement
updates to the database must persist even if there are software or hardware failures

Written in cache but not to main-memory

If not committed → roll back

3. Consistency Requirement :

- Sum of A and B is unchanged by the execution of the transaction
- Integrity constraints

4. Isolation

- Multi-core and Single-core
Parallel Concurrent

Time

T_1 \rightarrow 1, 2, 3 satisfied

read(A)

A: A - 50

write(A)

⋮

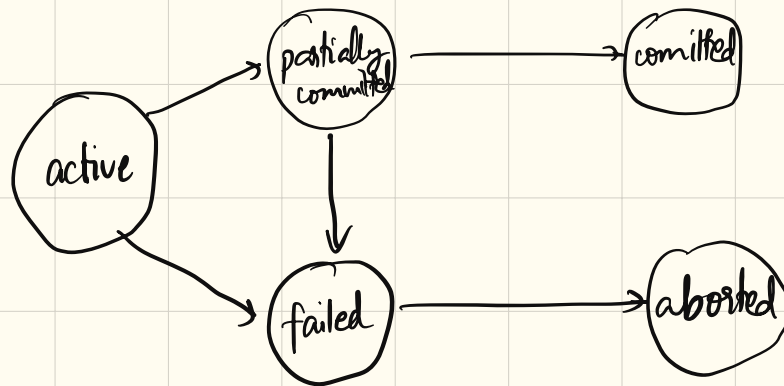
read(A), read(B), print(A+B)
50 less

ACID Properties

→ Common interview question

- Atomicity
- Consistency
- Isolation
- Durability

Transaction State
(similar to process state)



System logs
in secondary
memory

→ Recovery is
not covered here

Concurrent Execution

Schedules

- Serial schedule
- Concurrent schedule

Quiz 4

$$P_R + \frac{P_R}{(B-A) \cdot P_S}$$

$$2000 + \frac{2}{2000} \\ 20 * 1000$$

200000

80000

$$80000 + \frac{100000 \times 80000}{20}$$