

# Fundamentals of Database Systems

## [Query Optimization – II]

Malay Bhattacharyya

Assistant Professor

Machine Intelligence Unit  
Indian Statistical Institute, Kolkata

September, 2019

## 1 Query Equivalence

## 2 Query Tuning

- Motivation
- Logical and Physical Plan
- Optimizing Search Strategies

# What is query equivalence?

Two relational algebra expressions are said to be *equivalent* if on every legal database instance (i.e., a relation) the two expressions generate the same relation (i.e., the same set of tuples).

Query equivalence relations are used for tuning a query into an optimized form.

# Query equivalence – On projection and selection

- Cascade property of projection:  $\pi_{X_1}(\pi_{X_2}(\dots(\pi_{X_n}(R))\dots)) \equiv \pi_{X_1}(R)$ .
- Cascade property of selection:  $\sigma_{\theta_1 \wedge \theta_2}(R) \equiv \sigma_{\theta_1}(\sigma_{\theta_2}(R))$ .
- Commutative property of selection:  $\sigma_{\theta_1}(\sigma_{\theta_2}(R)) \equiv \sigma_{\theta_2}(\sigma_{\theta_1}(R))$ .
- Selection can be combined with Cartesian product and theta join in the following way:
  - 1  $\sigma_{\theta}(R_1 \times R_2) \equiv R_1 \bowtie_{\theta} R_2$  (this is simply the definition of theta join).
  - 2  $\sigma_{\theta_1}(R_1 \bowtie_{\theta_2} R_2) \equiv R_1 \bowtie_{\theta_1 \wedge \theta_2} R_2$ .

## Query equivalence – On theta-join

- Commutative property of theta-join:  $R_1 \bowtie_{\theta} R_2 \equiv R_2 \bowtie_{\theta} R_1$ .
- Theta joins are associative in the following way:  
 $(R_1 \bowtie_{\theta_1} R_2) \bowtie_{\theta_2 \wedge \theta_3} R_3 \equiv R_1 \bowtie_{\theta_1 \wedge \theta_3} (R_2 \bowtie_{\theta_2} R_3)$ , where  $\theta_2$  involves attributes only from  $R_2$  and  $R_3$ . Any of these conditions may be empty, and hence it follows that the Cartesian product operation is also associative.
- The selection operation distributes over the theta-join operation under the following two conditions:
  - 1 It distributes when all the attributes in selection condition  $\theta_0$  involve only the attributes of one of the relations (say  $R_1$ ) being joined. i.e.  $\sigma_{\theta_0}(R_1 \bowtie_{\theta} R_2) \equiv (\sigma_{\theta_0}(R_1) \bowtie_{\theta} R_2)$ .
  - 2 It distributes when selection condition  $\theta_1$  involves only the attributes of  $R_1$  and  $\theta_2$  involves only the attributes of  $R_2$ . i.e.  $\sigma_{\theta_1 \wedge \theta_2}(R_1 \bowtie_{\theta} R_2) \equiv (\sigma_{\theta_1}(R_1) \bowtie_{\theta} \sigma_{\theta_2}(R_2))$ .

## Query equivalence – On theta-join

- The projection operation distributes over the theta-join operation under the following two conditions:
  - 1 Let  $L_1$  and  $L_2$  be the attributes of  $R_1$  and  $R_2$ , respectively, and the join condition  $\theta$  involves the attributes only in  $L_1 \cup L_2$ .  
Then we have

$$\pi_{L_1 \cup L_2}(R_1 \bowtie_{\theta} R_2) \equiv (\pi_{L_1}(R_1)) \bowtie_{\theta} (\pi_{L_2}(R_2)).$$

- 2 Consider a join operation  $R_1 \bowtie_{\theta} R_2$  and suppose  $L_1$  and  $L_2$  be the sets of attributes from  $R_1$  and  $R_2$ , respectively. Further assume that  $L_3$  and  $L_4$  denote the attributes of  $R_1$  and  $R_2$ , respectively, that are involved in join condition  $\theta$ , but are not in  $L_1 \cup L_2$ . Then we have

$$\pi_{L_1 \cup L_2}(R_1 \bowtie_{\theta} R_2) \equiv \pi_{L_1 \cup L_2}((\pi_{L_1 \cup L_3}(R_1)) \bowtie_{\theta} (\pi_{L_2 \cup L_4}(R_2))).$$

## Query equivalence – On natural join

- Commutative property of natural join:  $R_1 \bowtie R_2 \equiv R_2 \bowtie R_1$ .
- Associative property of natural join:  $(R_1 \bowtie R_2) \bowtie R_3 \equiv R_1 \bowtie (R_2 \bowtie R_3)$ .

**Note:** The commutativity and associativity of join operations are important for join reordering in query optimization.

## Query equivalence – On set operations

- Commutative property of set union:  $R_1 \cup R_2 \equiv R_2 \cup R_1$ .
- Associative property of set union:  $(R_1 \cup R_2) \cup R_3 \equiv R_1 \cup (R_2 \cup R_3)$ .
- Commutative property of set intersection:  $R_1 \cap R_2 \equiv R_2 \cap R_1$ .
- Associative property of set intersection:  $(R_1 \cap R_2) \cap R_3 \equiv R_1 \cap (R_2 \cap R_3)$ .
- The selection operation distributes over the union, intersection and set difference operations:  $\sigma_\theta(R_1 - R_2) \equiv \sigma_\theta(R_1) - \sigma_\theta(R_2)$  (replacing ‘-’ with either  $\cup$  or  $\cap$  also holds).
- Again we have the equivalence relation:  $\sigma_\theta(R_1 - R_2) \equiv \sigma_\theta(R_1) - R_2$  (replacing ‘-’ with  $\cap$  also holds, but not for  $\cup$ ).
- Distributive property of projection over union:  $\pi_X(R_1 \cup R_2) \equiv (\pi_X(R_1)) \cup (\pi_X(R_2))$ .











## Physical plans – Hash join

Hash joins are used while joining large tables. The optimizer uses smaller of the two tables to build a hash table in memory and then scans the larger table and compares its hash values (of tuples from larger table) with the hash table to find the joined rows.

The algorithm of hash join is divided in two parts as follows:

- 1** Build an in-memory hash table on smaller of the two tables.
- 2** Probe this hash table with hash value for each row in the other table



## Physical plans – Sort Merge join

Sort merge join is used to join two independent data sources. They perform better than the nested loop when the volume of data is big in tables.

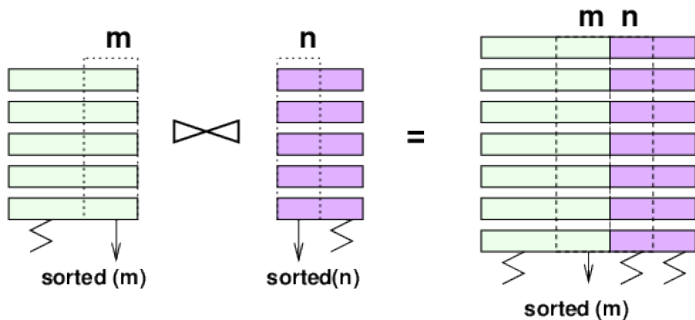
They perform better than hash join when the join condition is either an inequality condition or if sorting is anyways required due to some other attribute (other than join) like *order by*.

The full operation is done in two parts as follows:

- 1 Sort join operation
- 2 Merge join operation

**Note:** If the data is already sorted, first step is avoided.

# Physical plans – Sort Merge join





# Optimizing search strategies

There are several ways to optimize the searching mechanism as listed below:

- Heuristics
- Heuristics + Cost-based join order search
- Randomized algorithms
- Stratified search
- Unified search

# Optimization based on heuristics

Define static rules that transform logical operators to a physical plan. Some of these as follows:

- Perform most restrictive selection early
- Perform all selections before joins
- Predicate/Limit/Projection pushdowns
- Join ordering based on cardinality

**Note:** Original versions of INGRES and Oracle (until mid 1990s) used this.

# Optimization based on heuristics – An example

Consider the following three relations ACTOR, MOVIE and ACTS with primary keys {AID}, {MID} and {AID, MID}, respectively.

AID	Name
1	Amitabh Bachchan
2	Taapsee Pannu
3	Aksay Kumar
4	Prabhas
5	Shraddha Kapoor
6	Kangana Ranaut

MID	Name
1	Badla
2	Kesari
3	Saaho
4	Mental Hai Kya

AID	MID
1	1
2	1
3	2
5	3
4	3
6	4

## Optimization based on heuristics – An example

Suppose we want to retrieve the names of actors who acted in the movie *Saaho*.

The following query serves the purpose:

```
select ACTOR.Name from ACTOR, ACTS, MOVIE where
ACTOR.AID = ACTS.AID and ACTS.MID = MOVIE.MID and
MOVIE.Name = "Saaho";
```

Let us see how a heuristics based optimizer works!!!

# Optimization based on heuristics – An example

**Step 1:** Decompose a complex query into single-variable queries

Q

```
select ACTOR.Name from ACTOR, ACTS, MOVIE where
ACTOR.AID = ACTS.AID and ACTS.MID = MOVIE.MID and
MOVIE.Name = "Saaho";
```

Q1

```
select MOVIE.MID into
TEMP1 from MOVIE where
MOVIE.Name = "Saaho";
```

Q2

```
select ACTOR.Name from
ACTOR, ACTS, TEMP1
where ACTOR.AID =
ACTS.AID and ACTS.MID =
TEMP1.MID;
```

# Optimization based on heuristics – An example

**Step 1:** Decompose a complex query into single-variable queries  
(Continued ...)

Q2

```
select ACTOR.Name from ACTOR, ACTS, TEMP1 where
ACTOR.AID = ACTS.AID and ACTS.MID = TEMP1.MID;
```

Q3

```
select ACTS.AID into TEMP2
from ACTS, TEMP1 where
ACTS.MID = TEMP1.MID;
```

Q4

```
select ACTOR.Name from
ACTOR, TEMP2 where
ACTOR.AID = TEMP2.AID;
```

# Optimization based on heuristics – An example

**Step 2:** Substitute the values in the order Q1 → Q3 → Q4

Q1

```
select MOVIE.MID into TEMP1 from MOVIE where MOVIE.Name = "Saaho";
```

<b>MID</b>
3

# Optimization based on heuristics – An example

**Step 2:** Substitute the values in the order Q1 → Q3 → Q4  
(Continued ...)

Q3

```
select ACTS.AID into TEMP2 from ACTS, TEMP1 where  
ACTS.MID = 3;
```

AID
5
4



# Optimization based on heuristics – An example

**Step 2:** Substitute the values in the order Q1 → Q3 → Q4  
(Continued ...)

Q4

```
select ACTOR.Name from ACTOR, TEMP2 where ACTOR.AID =  
TEMP2.AID;
```

Name
Shraddha Kapoor
Prabhas

# Optimization based on heuristics – Pros and cons

## Advantages:

- 1 Easy to implement and debug.
- 2 Works reasonably well and is fast for simple queries.

## Disadvantages:

- 1 Relies on magic constants that predict the efficacy of a planning decision.
- 2 Nearly impossible to generate good plans when operators have complex inter-dependencies.

# Optimization based on heuristics + cost-based join order search

Use static rules to perform initial optimization. Then use dynamic programming to determine the best join order for tables.

- First cost-based query optimizer
- Bottom-up planning (forward chaining) using a divide-and-conquer search method

**Note:** System R, early IBM DB2, most open-source DBMSs used this.

# System R style optimization

- 1 Break query up into blocks and generate the logical operators for each block.
- 2 For each logical operator, generate a set of physical operators that implement it.
  - All combinations of join algorithms and access paths
- 3 Then iteratively construct a *left-deep* tree that minimizes the estimated amount of work to execute the plan.

## System R style optimization – An example

Suppose we want to retrieve the names of actors who acted in the movie *Saaho* ordered by their actor ID.

The following query serves the purpose:

```
select ACTOR.Name from ACTOR, ACTS, MOVIE where  
ACTOR.AID = ACTS.AID and ACTS.MID = MOVIE.MID and  
MOVIE.Name = "Saaho" order by ACTOR.AID;
```

Let us see how the System R optimizer works!!!

# System R style optimization – An example

**Step 1:** Choose the best access paths to each table

- ACTOR: Sequential Scan
- ACTS: Sequential Scan
- MOVIE: Index Look-up on Name

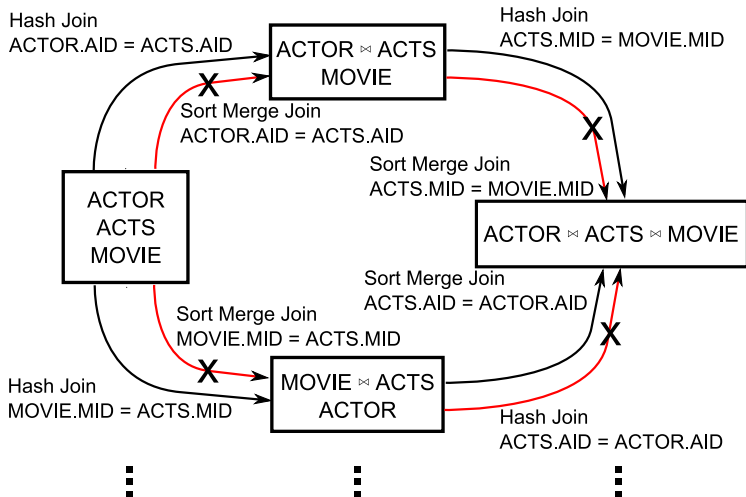
# System R style optimization – An example

**Step 2:** Enumerate all possible join orderings for tables

- ACTOR ⋈ ACTS ⋈ MOVIE
- ACTOR ⋈ MOVIE ⋈ ACTS
- ACTS ⋈ ACTOR ⋈ MOVIE
- ACTS ⋈ MOVIE ⋈ ACTOR
- MOVIE ⋈ ACTOR ⋈ ACTS
- MOVIE ⋈ ACTS ⋈ ACTOR

# System R style optimization – An example

**Step 3:** Determine the join ordering with the lowest cost





# Optimization based on heuristics + cost-based join order search – Pros and cons

## Advantages:

- 1 Usually finds a reasonable plan without having to perform an exhaustive search.

## Disadvantages:

- 1 All the same problems as the heuristic-only approach.
- 2 Left-deep join trees are not always optimal.
- 3 Have to take in consideration the physical properties of data in the cost model (e.g., sort order).

# Optimization based on randomized algorithms

Perform a random walk over a solution space of all possible (valid) plans for a query.

Continue searching until a cost threshold is reached or the optimizer runs for a particular length of time.

**Note:** Postgres genetic algorithm used this.



# Optimization based on stratified search

First rewrite the logical query plan using transformation rules.

- The engine checks whether the transformation is allowed before it can be applied.
- Cost is never considered in this step.

Finally, perform a cost-based search to map the logical plan to a physical plan.

# Optimization based on unified search

Unify the notion of both logical  $\rightarrow$  logical and logical  $\rightarrow$  physical transformations.

– No need for separate stages because everything is transformations.

This approach generates a lot more transformations so it makes heavy use of memorization to reduce redundant work.