



Introduction

- Query optimization is the process the best query execution plan (QEP) among the many possible ones
- Alternative ways to execute a given query
- Equivalent relational algebra expressions

instructor

• Different implementation choices for each relational algebra operation

INSTR(<u>i id</u>, name, dept_name, ...) COURSE(<u>c id</u>, title, ...) TEACHES(<u>i id</u>, <u>c id</u>, ...)

SELECT Lname, C.title FROM INSTRI, COURSE C, TEACHES T WHERE I.i_id = T.i_id AND T.c_id = C.c_id AND dept_name="Music"

teacho te

 $\prod (\sigma(INSTR) \Join (TEACHES \Join COURSE))$

hatz. Korth and Sudarshan

The name of all instructors in the department of Music together with the titles of all courses they teach

System Concents - 6th Fr

 $\prod (\sigma(INSTR | M(TEACHES | MCOURSE)))$



Introduction (Cont.)

A query evaluation plan (QEP) defines exactly what algorithm is used for each operation, and how the execution of the operations is coordinated



Find out how to view query execution plans on your favorite database 1.4

hatz. Korth and Su



Introduction (Cont.)

- Cost difference between query evaluation plans can be enormous . E.g. seconds vs. days in some cases
- Steps in cost-based query optimization
 - 1. Generate logically equivalent expressions using equivalence rules
 - 2. Annotate resulting expressions to get alternative QEP
 - 3. Evaluate/estimate the cost (execution time) of each QEP
 - 4. Choose the cheapest QEP based on estimated cost
- Estimation of QEP cost based on:

se System Concepts - 6th Edition

- Statistical information about relations (stored in the Catalog)
 - number of tuples, number of distinct values for an attribute

1.5

- Statistics estimation for intermediate results
- to compute cost of complex expressions
- Cost formulae for algorithms, computed using statistics



e System Concepts - 6th Edition



©Silberschatz, Korth and Sudarshan See www.db-book.com for conditions on re-use



Transformation of Relational Expressions

- Two relational algebra expressions are said to be equivalent if the two expressions generate the same set of tuples on every legal database instance
 - Note: order of tuples is irrelevant (and also order of attributes)
 - we don't care if they generate different results on databases that violate integrity constraints (e.g., uniqueness of keys)
- In SQL, inputs and outputs are multisets of tuples
 - Two expressions in the multiset version of the relational algebra are said to be equivalent if the two expressions generate the same multiset of tuples on every legal database instance
 - We focus on relational algebra and treat relations as sets
- An equivalence rule states that expressions of two forms are equivalent • One can replace an expression of first form by one of the second form,

ESilberschatz, Korth and Sudarshar

atz. Korth and Sudars

Ex. 13.1(d) *

or vice versa 1.7



Equivalence Rules

1. Conjunctive selection operations can be deconstructed into a sequence of individual selections. (E) = S (S (E))

2. Selection operations are commutative.

$$S_{q_{1}}(G_{q_{2}}(E)) = S_{q_{1}}(S_{q_{1}}(E))$$

$$S_{q_{1}}(S_{q_{2}}(E)) = S_{q_{2}}(S_{q_{1}}(E))$$

3. Only the last in a sequence of projection operations is needed, the others can be omitted

$$\Pi_{L_1}(\Pi_{L_2}(\ldots(\Pi_{L_n}(E))\ldots)) = \Pi_{L_1}(E)$$

where
$$L_1 \subseteq L_2 \subseteq \ldots \subseteq L_n$$

4. Selections can be combined with Cartesian products and theta joins.

hatz. Korth and Sudarsha

a.
$$\sigma_{\theta}(E_1 X E_2) = E_1 M_{\theta} E_2$$

se System Concepts - 6th Edition

b. $\sigma_{\theta_1}(\mathsf{E}_1 \bowtie_{\theta_2} \mathsf{E}_2) = \mathsf{E}_1 \bowtie_{\theta_1 \land \theta_2} \mathsf{E}_2$



Database System Concepts - 6th Edition

Equivalence Rules (Cont.)

- 5. Theta-join (and thus natural joins) operations are commutative. $E_1 \Join_{\theta} E_2 = E_2 \Join_{\theta} E_1$
 - (but the order is important for efficiency)
- 6. (a) Natural join operations are associative: $(E_1 \bowtie E_2) \bowtie E_3 = E_1 \bowtie (E_2 \bowtie E_3)$
 - (again, the order is important for efficiency)
 - (b) Theta joins are associative in the following manner:

 $(E_1 \boxtimes_{\theta_1} E_2) \boxtimes_{\theta_2 \land \theta_3} E_3 = E_1 \boxtimes_{\theta_1 \land \theta_3} (E_2 \boxtimes_{\theta_2} E_3)$

where θ_1 involves attributes from only E_1 and E_2

and θ_2 involves attributes from only E_2 and E_3

More equivalences at Ch. 13.2 of the book *

rshan. Database System Concepts, 6° ed. Silberschatz, Korth, and Sur

base System Concepts - 6th Edition 1.9



Pictorial Depiction of Equivalence Rules



Exercise

Create equivalence rules to push selection inside a left outer join Ex. 13.1(c) *

Disprove the equivalence

```
(R \bowtie S) \bowtie T = R \bowtie (S \bowtie T)
```

Definition (left outer join): the result of a left outer join $T = R \bowtie S$ is a super-set of the result of the join $T' = R \bowtie S$ in that all tuples in T' appear in T. In addition, T preserve those tuples that are lost in the join, by creating tuples in T that are filled with null values

STUD TAKES	stud_id name 1 gino 2 filippo 3 mario stud_id course 1 Algebra 2 Progr. 2 Math 2 Logic	name gino filippo mario	surname bianchi neri rossi	STUD] stud_id 1	I TAKES name gino	surname bianchi	course Math	grade 30
		grade 30 26 22 28 30	1 2 2 3	gino filippo filippo filippo mario	bianchi neri neri neri rossi	Algebra Progr. Math Logic null	26 22 28 30 <i>null</i>	
* Silberscha	tz, Korth, and S	udarshan, Data	abase System Conce	epts, 6* ed.		CSIII	erschatz Kort	h and Sudars



Solutions

Create equivalence rules involving left outer join and selection σ_{θ} (R \supset S) = σ_{θ} (R) \supset S

where θ uses only attributes of R

m Concepts - 6th Edition

■ Disprove the equivalence (R → S) → T = R → (S → T)



1.12

atz. Korth and Sudarshar





Selection Size Estimation

σ_{A=ν}(r)

base System Concepts - 6th Edition

- n_r / V(A,r) : number of records that will satisfy the selection (uniform distribution)
- Field Equality condition on a key attribute: size estimate = 1
- σ_{A ≤ V}(r) (case of σ_{A ≥ V}(r) is symmetric)
 - n: estimated number of tuples satisfying the condition is computed assuming that min(A,r) and max(A,r) are available in catalog
 n = 0 if v < min(A,r)
 - $\rightarrow n = n_r \cdot \frac{v \min(A, r)}{\max(A, r) \min(A, r)}$

otherwise (uniform distribution)

Silberschatz, Korth and Sudarsha

Silberschatz, Korth and Sudarshan

hatz. Korth and Sudarshan

 In absence of statistical information or when v is unknown at time of cost estimation (e.g., v is computed at run-time by the application using the DB) n is assumed to be n_r/2

1.19

If histograms are available, we can refine above estimate by using values for restricted ranges instead of values referring to the entire domain $(n_r, V(A, r), min(A, r), max(A, r))$



Database System Concepts, 6th Ed. ©Silberschatz, Korth and Sudarshan See www.db-book.com for conditions on re-use



Choice of Evaluation Plans

- Must consider the interaction of evaluation techniques when choosing evaluation plans
 - choosing the cheapest algorithm for each operation independently may not yield best overall algorithm. E.g.
 - merge-join may be costlier than hash-join, but may provide a sorted output which reduces the cost for an outer level aggregation
 - nested-loop join may provide opportunity for pipelining
- Practical query optimizers incorporate elements of the following two broad approaches:
 - 1. Search all the plans and choose the best plan in a cost-based fashion

1.21

2. Uses heuristics to choose a plan



Cost-Based Optimization

- Consider finding the best join-order for $r_1 \bowtie r_2 \bowtie \dots r_n$.
- There are (2(n 1))!/(n 1)! different join orders for above expression. With n = 7, the number is 665280, with n = 10, the number is greater than 17.6 billion!
- No need to generate all the join orders. Exploiting some monotonicity (optimal substructure property), the least-cost join order for any subset of {r₁, r₂, ..., r_n} is computed only once.

Cost-Based Optimization: An example

- Consider finding the best join-order for r₁ ⋈ r₂ ⋈ r₃ ⋈ r₄ ⋈ r₅
- Number of possible different join orderings: $\frac{(2(n-1))!}{(n-1)!} = \frac{8!}{4!} = 1680$
- The least-cost join order for any subset of { r₁, r₂, r₃, r₄, r₅} is computed only once
- Assume we want to compute N_{12345} : number of possible different join orderings where r_1, r_2, r_5 as a grouped together, e.g.,

 $(r_1 \bowtie r_2 \bowtie r_3) \bowtie r_4 \bowtie r_5$

- $r_5 \qquad (r_2 \bowtie r_3 \bowtie r_1) \bowtie (r_{\mathcal{B}} \bowtie r_4) \qquad r_4 \bowtie (r_{\mathcal{B}} \bowtie (r_1 \bowtie (r_2 \bowtie r_3)))$
- The naïve approach

e System Concepts - 6th Edition

ase System Concepts - 6th Edition

- $N_{123/45} = N_{123} * N_{45}$
- **N**₁₂₃ = $\frac{4!}{2!}$ = 12 (*N*₁₂₃ : # ways of arranging r_1 , r_2 , and r_3)
- **•** $N_{45} = N_{123} = 12$ (N_{45} : # ways of arranging r_4 and r_5 wrt. block of r_1 , r_2 , and r_3)

1.23

- **N**_{123/45} = 12 * 12 = 144
- Exploiting optimal substructure property:
 - compute only once best ordering for r₁ r₂ r₂ r₃ : 12 possibilities (N₁₂₃)
 - compute best ordering for R₁₂₃ r₄ ⋈ r₅: 12 possibilities (N₄₅)
 - Therefore, **N**₁₂₃₄₅ = 12 + 12 = 24



ase System Concepts - 6th Edition

Dynamic Programming in Optimization

1.22

©Silberschatz, Korth and Sudarshar

hatz. Korth and Sudarshar

To find best join tree (equivalently, best join order) for a set of *n* relations:
 To find best plan for a set *S* of *n* relations, consider all possible plans of the form:

 $S' \Join (S \setminus S')$

- for every non-empty subset S' of S
 Recursively compute costs of best join orders for subsets S' and S \ S'
- to find the cost of each plan. Choose the cheapest of the $2^n 2$ alternatives
- Base case for recursion: single relation access plan
 Apply all selections on R_i using best choice of indices on R_i
- When a plan for a subset is computed, store it and reuse it when it is required again, instead of re-computing it

1.24

Dynamic programming

tabase System Concepts - 6th Edition



Join Order Optimization Algorithm

procedure findbestplan(S)
if (bestplan(S).cost ≠ ∞)
return bestplan[S]
// else bestplan[S] has not been computed earlier, compute it now
if (S contains only 1 relation)
set bestplan[S].plan and bestplan[S].costbased on the best way
of accessing S /* Using selections on S and indices on S */
else for each non-empty subset S1 of S such that S1 ≠ S
P1= findbestplan(S-S1)
A = best algorithm for joining results of P1 and P2
cost = P1.cost+ P2.cost+ cost of A
if cost < bestplan[S].cost
bestplan[S].cost
bestplan[S].cost
cost bestplan[S].cost
bestplan[S].cost
bestplan[S].plan = "execute P1.plan; execute P2.plan;
join results of P1 and P2 using A"
return bestplan[S]</pre>



Cost of Optimization

- With dynamic programming time complexity of optimization is O(3ⁿ).
 With n = 10, this number is 59000 instead of 17.6 billion!
- Space complexity is O(2ⁿ)

tem Concents - 6th Edition

se System Concepts - 6th Edition

 Better time performance when considering only left-deep tree O(n 2ⁿ) Space complexity remains at O(2ⁿ) (heuristic approach)



 Cost-based optimization is expensive, but worthwhile for queries on large datasets (typical queries have small n, generally < 10)

1.26

chatz, Korth and Sudarsha

@Silberschatz, Korth and Sudarshar



abase System Concepts - 6th Edition

Cost Based Optimization with Equivalence Rules

1.25

- Physical equivalence rules equates logical operations (e.g., join) to physical ones (i.e., implementations e.g., nested-loop join, merge join)
 Relational algebra expression are converted into QEP with implementation details
- Efficient optimizer based on equivalence rules depends on
 - A space efficient representation of expressions which avoids making multiple copies of sub-expressions
 - Efficient techniques for detecting duplicate derivations of expressions
 - A form of dynamic programming, which stores the best plan for a subexpression the first time it is optimized, and reuses in on repeated
 - optimization calls on same sub-expression
 - Cost-based pruning techniques that avoid generating all plans (greedy, heuristics, dynamic programming/optimal substructure property)

1.27



erschatz, Korth and Sudarshar

rschatz, Korth and Sudarshar

Heuristic Optimization

- Cost-based optimization is expensive, even with dynamic programming
- Systems may use *heuristics* to reduce the number of choices that must be made in a cost-based fashion
- Heuristic optimization transforms the query-tree by using a set of rules that typically (but not in all cases) improve execution performance:
 - Perform selection early (reduces the number of tuples)
 - Perform projection early (reduces the number of attributes)Perform most restrictive selection and join operations (i.e. with
 - smallest result size) before other similar operations
 Only consider left-deep join orders (particularly suited for pipelining

1.28

as only one input has to be pipelined, the other is a relation)



se System Concepts - 6th Edition

e System Concepts - 6th Edition

Structure of Query Optimizers

- Some systems use only heuristics, others combine heuristics with partial cost-based optimization.
- Many optimizers considers only left-deep join orders.
 - Plus heuristics to push selections and projections down the query tree
 - Reduces optimization complexity and generates plans amenable to pipelined evaluation.
- Heuristic optimization used in some versions of Oracles
 - Repeatedly pick "best" relation to join next
 - > Starting from each of n starting points. Pick best among these



Database System Concepts, 6th Ed. ©Silberschatz, Korth and Sudarshan See www.db-book.com for conditions on re-use