

CSE 444: Database Internals

Lectures 11-12 Query Optimization (part 2)

Magda Balazinska - CSE 444, Spring 2012

1

Query Optimization Algorithm

- Enumerate alternative plans (logical & physical)
- Compute estimated cost of each plan
 - Compute number of I/Os
 - Compute CPU cost
- Choose plan with lowest cost
 - This is called cost-based optimization

Magda Balazinska - CSE 444, Spring 2012

2

Lessons

- Need to consider several physical plans
 - Even for one, simple logical plan
- No magic “best” plan: depends on the data
- In order to make the right choice
 - Need to have **statistics** over the data
 - The B's, the T's, the V's

Magda Balazinska - CSE 444, Spring 2012

3

Outline

- Search space
- Algorithm for enumerating query plans

Magda Balazinska - CSE 444, Spring 2012

4


Relational Algebra Equivalences

- Selections
 - Commutative: $\sigma_{c_1}(\sigma_{c_2}(R))$ same as $\sigma_{c_2}(\sigma_{c_1}(R))$
 - Cascading: $\sigma_{c_1 \wedge c_2}(R)$ same as $\sigma_{c_2}(\sigma_{c_1}(R))$
- Projections
 - Cascading
- Joins
 - Commutative: $R \bowtie S$ same as $S \bowtie R$
 - Associative: $R \bowtie (S \bowtie T)$ same as $(R \bowtie S) \bowtie T$

Magda Balazinska - CSE 444, Spring 2012

5

Left-Deep Plans and Bushy Plans



Magda Balazinska - CSE 444, Spring 2012

6

Commutativity, Associativity, Distributivity

$$R \cup S = S \cup R, R \cup (S \cap T) = (R \cup S) \cap T$$

$$R \bowtie S = S \bowtie R, R \bowtie (S \bowtie T) = (R \bowtie S) \bowtie T$$

$$R \bowtie (S \cup T) = (R \bowtie S) \cup (R \bowtie T)$$

Magda Balazinska - CSE 444, Spring 2012

7

Laws Involving Selection

$$\sigma_{C \text{ AND } C'}(R) = \sigma_C(\sigma_{C'}(R)) = \sigma_C(R) \cap \sigma_{C'}(R)$$

$$\sigma_{C \text{ OR } C'}(R) = \sigma_C(R) \cup \sigma_{C'}(R)$$

$$\sigma_C(R \bowtie S) = \sigma_C(R) \bowtie S$$

$$\sigma_C(R - S) = \sigma_C(R) - S$$

$$\sigma_C(R \cup S) = \sigma_C(R) \cup \sigma_C(S)$$

$$\sigma_C(R \bowtie S) = \sigma_C(R) \bowtie S$$

Assuming C on attributes of R

Magda Balazinska - CSE 444, Spring 2012

8

Example: Simple Algebraic Laws

- Example: $R(A, B, C, D), S(E, F, G)$
- $\sigma_{F=3}(R \bowtie_{D=E} S) = ?$
- $\sigma_{A=5 \text{ AND } G=9}(R \bowtie_{D=E} S) = ?$

Magda Balazinska - CSE 444, Spring 2012

9

Laws Involving Projections

$$\Pi_M(R \bowtie S) = \Pi_M(\Pi_P(R) \bowtie \Pi_Q(S))$$

$$\Pi_M(\Pi_N(R)) = \Pi_M(R)$$

/* note that $M \subseteq N$ */

- Example $R(A,B,C,D), S(E, F, G)$
- $\Pi_{A,B,G}(R \bowtie_{D=E} S) = \Pi_{A,B,G}(\Pi_{A,B,C,D}(R) \bowtie_{D=E} \Pi_{E,F,G}(S))$

Magda Balazinska - CSE 444, Spring 2012

10

Laws involving grouping and aggregation

$$\delta(\gamma_{A, \text{agg}(B)}(R)) = \gamma_{A, \text{agg}(B)}(R)$$

$$\gamma_{A, \text{agg}(B)}(\delta(R)) = \gamma_{A, \text{agg}(B)}(R)$$

if agg is "duplicate insensitive"

Which of the following are "duplicate insensitive"?
sum, count, avg, min, max

$$\gamma_{A, \text{agg}(D)}(R(A,B) \bowtie_{B=C} S(C,D)) = \gamma_{A, \text{agg}(D)}(R(A,B) \bowtie_{B=C} (\gamma_{C, \text{agg}(D)}(S(C,D))))$$

Magda Balazinska - CSE 444, Spring 2012

11

Laws Involving Constraints

Product(pid, pname, price, cid)
Company(cid, cname, city, state)

Foreign key

$$\Pi_{\text{pid, price}}(\text{Product} \bowtie_{\text{cid=cid}} \text{Company}) = \Pi_{\text{pid, price}}(\text{Product})$$

Magda Balazinska - CSE 444, Spring 2012

12

Search Space Challenges

- **Search space is huge!**
 - Many possible equivalent trees
 - Many implementations for each operator
 - Many access paths for each relation
 - File scan or index + matching selection condition
- Cannot consider ALL plans
 - Heuristics: only partial plans with “low” cost

Magda Balazinska - CSE 444, Spring 2012

13

Outline

- Search space
- **Algorithm for enumerating query plans**

Magda Balazinska - CSE 444, Spring 2012

14

Key Decisions

Logical plan

- What logical plans do we consider (left-deep, bushy ?); *Search Space*
- Which algebraic laws do we apply, and in which context(s) ?; *Optimization rules*
- In what order do we explore the search space ?; *Optimization algorithm*

Magda Balazinska - CSE 444, Spring 2012

15

Key Decisions

Physical plan

- What physical operators to use?
- What access paths to use (file scan or index)?
- Pipeline or materialize intermediate results?

These decisions also affect the *search space*

Magda Balazinska - CSE 444, Spring 2012

16

Two Types of Optimizers

- **Heuristic-based optimizers:**
 - Apply greedily rules that always improve plan
 - Typically: push selections down
 - Very limited: no longer used today
- **Cost-based optimizers:**
 - Use a cost model to estimate the cost of each plan
 - Select the “cheapest” plan
 - We focus on cost-based optimizers

Magda Balazinska - CSE 444, Spring 2012

17

Three Approaches to Search Space Enumeration

- Complete plans
- Bottom-up plans
- Top-down plans

Magda Balazinska - CSE 444, Spring 2012

18

Complete Plans

R(A,B)
S(B,C)
T(C,D)

```
SELECT *
FROM R, S, T
WHERE R.B=S.B and S.C=T.C and R.A<40
```

Magda Balazinska - CSE 444, Spring 2012 19

Bottom-up Partial Plans

R(A,B)
S(B,C)
T(C,D)

```
SELECT *
FROM R, S, T
WHERE R.B=S.B and S.C=T.C and R.A<40
```

Why is this better ?

Magda Balazinska - CSE 444, Spring 2012 20

Top-down Partial Plans

R(A,B)
S(B,C)
T(C,D)

```
SELECT *
FROM R, S, T
WHERE R.B=S.B and S.C=T.C and R.A<40
```

Magda Balazinska - CSE 444, Spring 2012 21

Two Types of Plan Enumeration Algorithms

- Dynamic programming (in class)
 - Based on System R (aka Selinger) style optimizer[1979]
 - Limited to joins: *join reordering algorithm*
 - Bottom-up
- Rule-based algorithm (will not discuss)
 - Database of rules (=algebraic laws)
 - Usually: dynamic programming
 - Usually: top-down

Magda Balazinska - CSE 444, Spring 2012 22

Dynamic Programming

Originally proposed in System R [1979]

- Only handles single block queries:

```
SELECT list
FROM R1, ..., Rn
WHERE cond1 AND cond2 AND ... AND cond_k
```

- Some heuristics for search space enumeration:
 - Selections down
 - Projections up
 - Avoid cartesian products

Magda Balazinska - CSE 444, Spring 2012 23

Dynamic Programming

- Search space = join trees
- Algebraic laws = commutativity, associativity
- Algorithm = dynamic programming ☺

Magda Balazinska - CSE 444, Spring 2012 24

Selinger Optimizer Algorithm


- Original Selinger optimizer enumerates different logical and physical plans at the same time
- To simplify the discussion, we will first study the approach considering only logical plans
- We come back to the actual Selinger enumeration algorithm at the end of the lecture

Magda Balazinska - CSE 444, Spring 2012

25

Join Trees

- $R1 \bowtie R2 \bowtie \dots \bowtie Rn$
- Join tree:




- A plan = a join tree
- A partial plan = a subtree of a join tree

Magda Balazinska - CSE 444, Spring 2012

26

Types of Join Trees

- Bushy:




Magda Balazinska - CSE 444, Spring 2012

27

Types of Join Trees

- Right deep:



Magda Balazinska - CSE 444, Spring 2012

28

Types of Join Trees

- Left deep:
 - Work well with existing join algos
 - Nested-loop and hash-join
 - Facilitate pipelining
-
- Selinger algorithm considers only those trees
 - Dynamic programming can be used with all trees

Magda Balazinska - CSE 444, Spring 2012

29

Dynamic Programming

Join ordering:

- Given: a query $R1 \bowtie R2 \bowtie \dots \bowtie Rn$
- Find optimal order
- Assume we have a function $cost()$ that gives us the cost of every join tree

```
SELECT list
FROM R1, ..., Rn
WHERE cond1 AND cond2 AND ... AND condk
```

Magda Balazinska - CSE 444, Spring 2012

30

Dynamic Programming

- For each subquery $Q \subseteq \{R_1, \dots, R_n\}$ compute the following:
 - Size(Q) = the estimated size of Q
 - I.e., the cardinality of the result of Q
 - Plan(Q) = a best plan for Q
 - Cost(Q) = the estimated cost of that plan
 - Note: we focus first on logical plans so we will use as cost estimate the *sum of cardinalities of intermediate relations*

```
SELECT list
FROM R1, ..., Rn
WHERE cond1 AND cond2 AND ... AND condk
```

Magda Balazinska - CSE 444, Spring 2012

31

Dynamic Programming

- Step 1:** For each $\{R_i\}$, set:
 - Size($\{R_i\}$) = $T(R_i)$
 - Plan($\{R_i\}$) = R_i
 - That's the only alternative for a logical plan
 - Cost($\{R_i\}$) = 0
 - Remember that we are computing costs of logical plans

```
SELECT list
FROM R1, ..., Rn
WHERE cond1 AND cond2 AND ... AND condk
```

Magda Balazinska - CSE 444, Spring 2012

32

Dynamic Programming

- Step 2:** For each $Q \subseteq \{R_1, \dots, R_n\}$ involving i relations:
 - Size(Q) = estimate it recursively
 - For every pair of subqueries Q', Q'' s.t. $Q = Q' \cup Q''$ compute $\text{cost}(\text{Plan}(Q') \bowtie \text{Plan}(Q''))$
 - Cost(Q) = the smallest such cost
 - Plan(Q) = the least-cost plan

```
SELECT list
FROM R1, ..., Rn
WHERE cond1 AND cond2 AND ... AND condk
```

Magda Balazinska - CSE 444, Spring 2012

33

Dynamic Programming

- Step 3:** Return Plan($\{R_1, \dots, R_n\}$)

```
SELECT list
FROM R1, ..., Rn
WHERE cond1 AND cond2 AND ... AND condk
```

Magda Balazinska - CSE 444, Spring 2012

34

Example

We use cost model for logical plans:

- $\text{Cost}(P_1 \bowtie P_2) = \text{Cost}(P_1) + \text{Cost}(P_2) + \text{size}(\text{intermediate results for } P_1, P_2)$
- Cost of a scan = 0

Magda Balazinska - CSE 444, Spring 2012

35

Example

- $R \bowtie S \bowtie T \bowtie U$
- Assumptions:

```
SELECT *
FROM R, S, T, U
WHERE cond1 AND cond2 AND ...
```

All join selectivities = 1%

```
T(R) = 2000
T(S) = 5000
T(T) = 3000
T(U) = 1000
```

```
T(R ⋈ S) = 0.01 * T(R) * T(S)
T(S ⋈ T) = 0.01 * T(S) * T(T)
etc.
```

Magda Balazinska - CSE 444, Spring 2012

36

Subquery	Size	Cost	Plan
RS			
RT			
RU			
ST			
SU			
TU			
RST			
RSU			
RTU			
STU			
RSTU			

T(R) = 2000
T(S) = 5000
T(T) = 3000
T(U) = 1000

37

Subquery	Size	Cost	Plan
RS	100k	0	RS
RT	60k	0	RT
RU	20k	0	UR
ST	150k	0	TS
SU	50k	0	US
TU	30k	0	UT
RST	3M	60k	(RT)S
RSU	1M	20k	(UR)S
RTU	600K	20k	(UR)T
STU	1.5M	30k	(UT)S
RSTU	30M	60k +50k=110k	(RT)(SU)

T(R) = 2K
T(S) = 5K
T(T) = 3K
T(U) = 1K

38

Reducing the Search Space

- Restriction 1: only left linear trees (no bushy)
- Why?**
- Restriction 2: no trees with cartesian product

$R(A,B) \bowtie S(B,C) \bowtie T(C,D)$

Plan: $(R(A,B) \bowtie T(C,D)) \bowtie S(B,C)$
has a cartesian product.
Most query optimizers will not consider it

Magda Balazinska - CSE 444, Spring 2012 39

Dynamic Programming: Summary

- Handles only join queries:
 - Selections are pushed down (i.e. early)
 - Projections are pulled up (i.e. late)
- Takes exponential time in general, BUT:
 - Left linear joins may reduce time
 - Non-cartesian products may reduce time further

Magda Balazinska - CSE 444, Spring 2012 40

Completing the Physical Query Plan

- Choose algorithm for each operator
 - How much memory do we have ?
 - Are the input operand(s) sorted ?
- Access path selection for base tables
- Decide for each intermediate result:
 - To materialize
 - To pipeline

Magda Balazinska - CSE 444, Spring 2012 41

More about the Selinger Algorithm

Selinger enumeration algorithm considers

- Different logical and physical plans *at the same time*
- Cost of a plan is IO + CPU
- Concept of *interesting order* during plan enumeration
 - Same order as that requested by ORDER BY or GROUP BY
 - Attributes that appear in equi-join predicates
 - They can speed-up a sort-merge join later

Magda Balazinska - CSE 444, Spring 2012 42

More about the Selinger Algorithm

- Step 1: Enumerate all access paths for a single relation
 - File scan or index scan
 - Keep the cheapest for each *interesting order*
- Step 2: Consider all ways to join two relations
 - Use result from step 1 as the outer relation
 - Consider every other possible relation as inner relation
 - Estimate cost when using sort-merge or nested-loop join
 - Keep the cheapest for each *interesting order*
- Steps 3 and later: Repeat for three relations, etc.

Magda Balazinska - CSE 444, Spring 2012

43

Selinger Algorithm Example

- On the white board

Magda Balazinska - CSE 444, Spring 2012

44