

# CSE 444: Database Internals

## Lecture 11 Query Optimization (part 2)

CSE 444 - Winter 2019

1

## Query Optimizer Overview

- **Input:** A logical query plan
- **Output:** A good physical query plan
- **Basic query optimization algorithm**
  - Enumerate alternative plans (logical and physical)
  - Compute estimated cost of each plan
    - Compute number of I/Os
    - Optionally take into account other resources
  - Choose plan with lowest cost
  - This is called cost-based optimization

CSE 444 - Winter 2019

3

## The Three Parts of an Optimizer

- Cost estimation
  - Based on cardinality estimation
- Search space
  - Algebraic laws, restricted types of join trees
- Search algorithm
  - Will discuss next

CSE 444 - Winter 2019

4

## Search Algorithm

- 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

CSE 444 - Winter 2019

5

## Dynamic Programming

Originally proposed in System R [1979]

- Only handles single block queries:

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

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

CSE 444 - Winter 2019

6

## Dynamic Programming

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

- For each subquery  $Q \subseteq \{R1, \dots, Rn\}$  compute the following:
  - $T(Q)$  = the estimated size of  $Q$
  - $Plan(Q)$  = a best plan for  $Q$
  - $Cost(Q)$  = the estimated cost of that plan

CSE 444 - Winter 2019

7

SELECT list  
FROM R1, ..., Rn  
WHERE cond1 AND cond2 AND ... AND cond.

## Dynamic Programming

- **Step 1:** For each  $\{R_i\}$  do:
  - $T(\{R_i\}) = T(R_i)$
  - $\text{Plan}(\{R_i\}) = \text{access method for } R_i$
  - $\text{Cost}(\{R_i\}) = \text{cost of access method for } R_i$

CSE 444 - Winter 2019 8

SELECT list  
FROM R1, ..., Rn  
WHERE cond1 AND cond2 AND ... AND cond.

## Dynamic Programming

- **Step 2:** For each  $Q \subseteq \{R_1, \dots, R_n\}$  of size  $k$  do:
  - $T(Q) = \text{use estimator}$
  - Consider all partitions  $Q = Q' \cup Q''$   
compute  $\text{cost}(\text{Plan}(Q') \bowtie \text{Plan}(Q''))$
  - $\text{Cost}(Q) = \text{the smallest such cost}$
  - $\text{Plan}(Q) = \text{the corresponding plan}$
- **Note**
  - If we restrict to left-linear trees:  $Q' = \text{single relation}$
  - May want to avoid cartesian products

CSE 444 - Winter 2019 9

SELECT list  
FROM R1, ..., Rn  
WHERE cond1 AND cond2 AND ... AND cond.

## Dynamic Programming

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

CSE 444 - Winter 2019 10

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

## Example

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

$T(R) = 2000$   
 $T(S) = 5000$   
 $T(T) = 3000$   
 $T(U) = 1000$
- Every join selectivity is 0.001

CSE 444 - Winter 2019 11

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

Assume  
 $B(\cdot) = T(\cdot)/10$

Join selectivity  
is 0.001

Subquery	T	Plan	Cost
R	2000		
S	5000		
T	3000		
U	1000		
RS			
RT			
RU			
ST			
SU			
TU			
RST			
RSU			
RTU			
STU			
RSTU			

CSE 444 - Winter 2019 12

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

Assume  
 $B(\cdot) = T(\cdot)/10$

Join selectivity  
is 0.001

Subquery	T	Plan	Cost
R	2000		
S	5000		
T	3000		
U	1000		
RS	10000		
RT	6000		
RU	2000		
ST	15000		
SU	5000		
TU	3000		
RST	30000		
RSU	10000		
RTU	6000		
STU	15000		
RSTU	30000		

CSE 444 - Winter 2019 13

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

Assume  $B(.) = T(.)/10$

Join selectivity is 0.001

Subquery	T	Plan	Cost
R	2000	Clustered index scan R.A	200
S	5000		
T	3000		
U	1000		
RS	10000		
RT	6000		
RU	2000		
ST	15000		
SU	5000		
TU	3000		
RST	30000		
RSU	10000		
RTU	6000		
STU	15000		
RSTU	30000		

14

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

Assume  $B(.) = T(.)/10$

Join selectivity is 0.001

Subquery	T	Plan	Cost
R	2000	Clustered index scan R.A	200
S	5000	Table scan	500
T	3000		
U	1000		
RS	10000		
RT	6000		
RU	2000		
ST	15000		
SU	5000		
TU	3000		
RST	30000		
RSU	10000		
RTU	6000		
STU	15000		
RSTU	30000		

15

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

Assume  $B(.) = T(.)/10$

Join selectivity is 0.001

Subquery	T	Plan	Cost
R	2000	Clustered index scan R.A	200
S	5000	Table scan	500
T	3000	Table scan	300
U	1000	Unclustered index scan U.F	1000
RS	10000		
RT	6000		
RU	2000		
ST	15000		
SU	5000		
TU	3000		
RST	30000		
RSU	10000		
RTU	6000		
STU	15000		
RSTU	30000		

16

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

Assume  $B(.) = T(.)/10$

Join selectivity is 0.001

Subquery	T	Plan	Cost
R	2000	Clustered index scan R.A	200
S	5000	Table scan	500
T	3000	Table scan	300
U	1000	Unclustered index scan U.F	1000
RS	10000	R $\bowtie$ S nested loop join	...
RT	6000		
RU	2000		
ST	15000		
SU	5000		
TU	3000		
RST	30000		
RSU	10000		
RTU	6000		
STU	15000		
RSTU	30000		

17

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

Assume  $B(.) = T(.)/10$

Join selectivity is 0.001

Subquery	T	Plan	Cost
R	2000	Clustered index scan R.A	200
S	5000	Table scan	500
T	3000	Table scan	300
U	1000	Unclustered index scan U.F	1000
RS	10000	R $\bowtie$ S nested loop join	...
RT	6000	R $\bowtie$ T index join	...
RU	2000		
ST	15000		
SU	5000		
TU	3000		
RST	30000		
RSU	10000		
RTU	6000		
STU	15000		
RSTU	30000		

18

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

Assume  $B(.) = T(.)/10$

Join selectivity is 0.001

Subquery	T	Plan	Cost
R	2000	Clustered index scan R.A	200
S	5000	Table scan	500
T	3000	Table scan	300
U	1000	Unclustered index scan U.F	1000
RS	10000	R $\bowtie$ S nested loop join	...
RT	6000	R $\bowtie$ T index join	...
RU	2000	R $\bowtie$ U index join	...
ST	15000	S $\bowtie$ T hash join	...
SU	5000	...	...
TU	3000	...	...
RST	30000		
RSU	10000		
RTU	6000		
STU	15000		
RSTU	30000		

19

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

Subquery	T	Plan	Cost
R	2000	Clustered index scan R.A	200
S	5000	Table scan	500
T	3000	Table scan	300
U	1000	Unclustered index scan U.F	1000
RS	10000	R $\bowtie$ S nested loop join	...
RT	6000	R $\bowtie$ T index join	...
RU	2000	R $\bowtie$ U index join	...
ST	15000	S $\bowtie$ T hash join	...
SU	5000	...	...
TU	3000	...	...
RST	30000	(RT) $\bowtie$ S hash join	...
RSU	10000	(SU) $\bowtie$ R merge join	...
RTU	6000	...	...
STU	15000	...	...
RSTU	30000	(RT) $\bowtie$ (SU) hash join	...

Assume  $B(.) = T./10$   
Join selectivity is 0.001

CSE 444 - Winter 2019 20

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

Subquery	T	Plan	Cost
R	2000	Clustered index scan R.A	200
S	5000	Table scan	500
T	3000	Table scan	300
U	1000	Unclustered index scan U.F	1000
RS	10000	R $\bowtie$ S nested loop join	...
RT	6000	R $\bowtie$ T index join	...
RU	2000	R $\bowtie$ U index join	...
ST	15000	S $\bowtie$ T hash join	...
SU	5000	...	...
TU	3000	...	...
RST	30000	(RT) $\bowtie$ S hash join	...
RSU	10000	(SU) $\bowtie$ R merge join	...
RTU	6000	...	...
STU	15000	...	...
RSTU	30000	(RT) $\bowtie$ (SU) hash join	...

Assume  $B(.) = T./10$   
Join selectivity is 0.001

CSE 444 - Winter 2019 21

## Discussion

- Need to consider both  $R \bowtie S$  and  $S \bowtie R$ 
  - Because the cost may be different! (indexes etc)
- When computing the cheapest plan for
- $(Q) \bowtie R$ , we may consider new access methods for R, e.g. an index look-up that makes sense only in the context of the join

CSE 444 - Winter 2019 22

SELECT list  
FROM R1, ..., Rn  
WHERE cond1 AND cond2 ... AND cond.

## Discussion

Given a query with n relations R1, ..., Rn

- How many entries do we have in the dynamic programming table?
- For each entry, how many alternative plans do we need to inspect?

CSE 444 - Winter 2019 23

SELECT list  
FROM R1, ..., Rn  
WHERE cond1 AND cond2 ... AND cond.

## Discussion

Given a query with n relations R1, ..., Rn

- How many entries do we have in the dynamic programming table?
  - A:  $2^n - 1$
- For each entry, how many alternative plans do we need to inspect?
  - A: for each entry with k tables, examine  $2^k - 2$  plans

CSE 444 - Winter 2019 24


## Reducing the Search Space

- Left-linear trees
- No cartesian products

CSE 444 - Winter 2019 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

CSE 444 - Winter 2019

26

## Types of Join Trees

- Bushy:




CSE 444 - Winter 2019

27

## Types of Join Trees

- Linear :




CSE 444 - Winter 2019

28

## Types of Join Trees

- Right deep:



CSE 444 - Winter 2019

29

## Types of Join Trees

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

CSE 444 - Winter 2019

30