

# Lecture 19: Query Optimization (1)

May 17, 2010

# Announcements

- Homework 3 due on Wednesday in class
  - How is it going?
- Project 4 posted
  - Due on June 2<sup>nd</sup>
  - Start early !

# Where We Are

- We are learning how a DBMS executes a query
- What we learned so far
  - How data is stored and indexed
  - Logical query plans and physical operators
- This week:
  - How to select logical & physical query plans

# Query Optimization Goal

- For a query
  - There exists many logical and physical query plans
  - Query optimizer needs to pick a good one

# Query Optimization Algorithm

- Enumerate alternative plans
- 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

# Example

```
Supplier(sid, sname, scity, sstate)  
Supply(sid, pno, quantity)
```

- Some statistics
  - T(Supplier) = 1000 records
  - T(Supply) = 10,000 records
  - B(Supplier) = 100 pages
  - B(Supply) = 100 pages
  - V(Supplier,scity) = 20, V(Supplier,state) = 10
  - V(Supply,pno) = 2,500
  - Both relations are clustered
- M = 10

```
SELECT sname  
FROM Supplier x, Supply y  
WHERE x.sid = y.sid  
and y.pno = 2  
and x.scity = 'Seattle'  
and x.sstate = 'WA'
```

T(Supplier) = 1000  
T(Supply) = 10,000

B(Supplier) = 100  
B(Supply) = 100

V(Supplier,scity) = 20  
V(Supplier,state) = 10  
V(Supply,pno) = 2,500

M = 10

# Physical Query Plan 1

(On the fly)

$\pi$  sname

Selection and project on-the-fly  
-> No additional cost.

(On the fly)

$\sigma$  scity='Seattle'  $\wedge$  sstate='WA'  $\wedge$  pno=2

(Block-nested loop)

sid = sid

Total cost of plan is thus cost of join:  
=  $B(\text{Supplier}) + B(\text{Supplier}) * B(\text{Supply}) / M$   
=  $100 + 10 * 100$   
= **1,100 I/Os**

Supplier  
(File scan)

Supply  
(File scan)

T(Supplier) = 1000  
T(Supply) = 10,000

B(Supplier) = 100  
B(Supply) = 100

V(Supplier,scity) = 20  
V(Supplier,state) = 10  
V(Supply,pno) = 2,500

M = 10

# Physical Query Plan 2

(On the fly)

$\pi_{\text{sname}}$  (4)

(Sort-merge join)

$\bowtie_{\text{sid} = \text{sid}}$  (3)

(Scan  
write to T1)

(1)  $\sigma_{\text{scity}='Seattle' \wedge \text{sstate}='WA'}$

Supplier  
(File scan)

(Scan  
write to T2)

(2)  $\sigma_{\text{pno}=2}$

Supply  
(File scan)

Total cost

= 100 + 100 \* 1/20 \* 1/10 (1)

+ 100 + 100 \* 1/2500 (2)

+ 2 (3)

+ 0 (4)

Total cost  $\approx$  **204 I/Os**

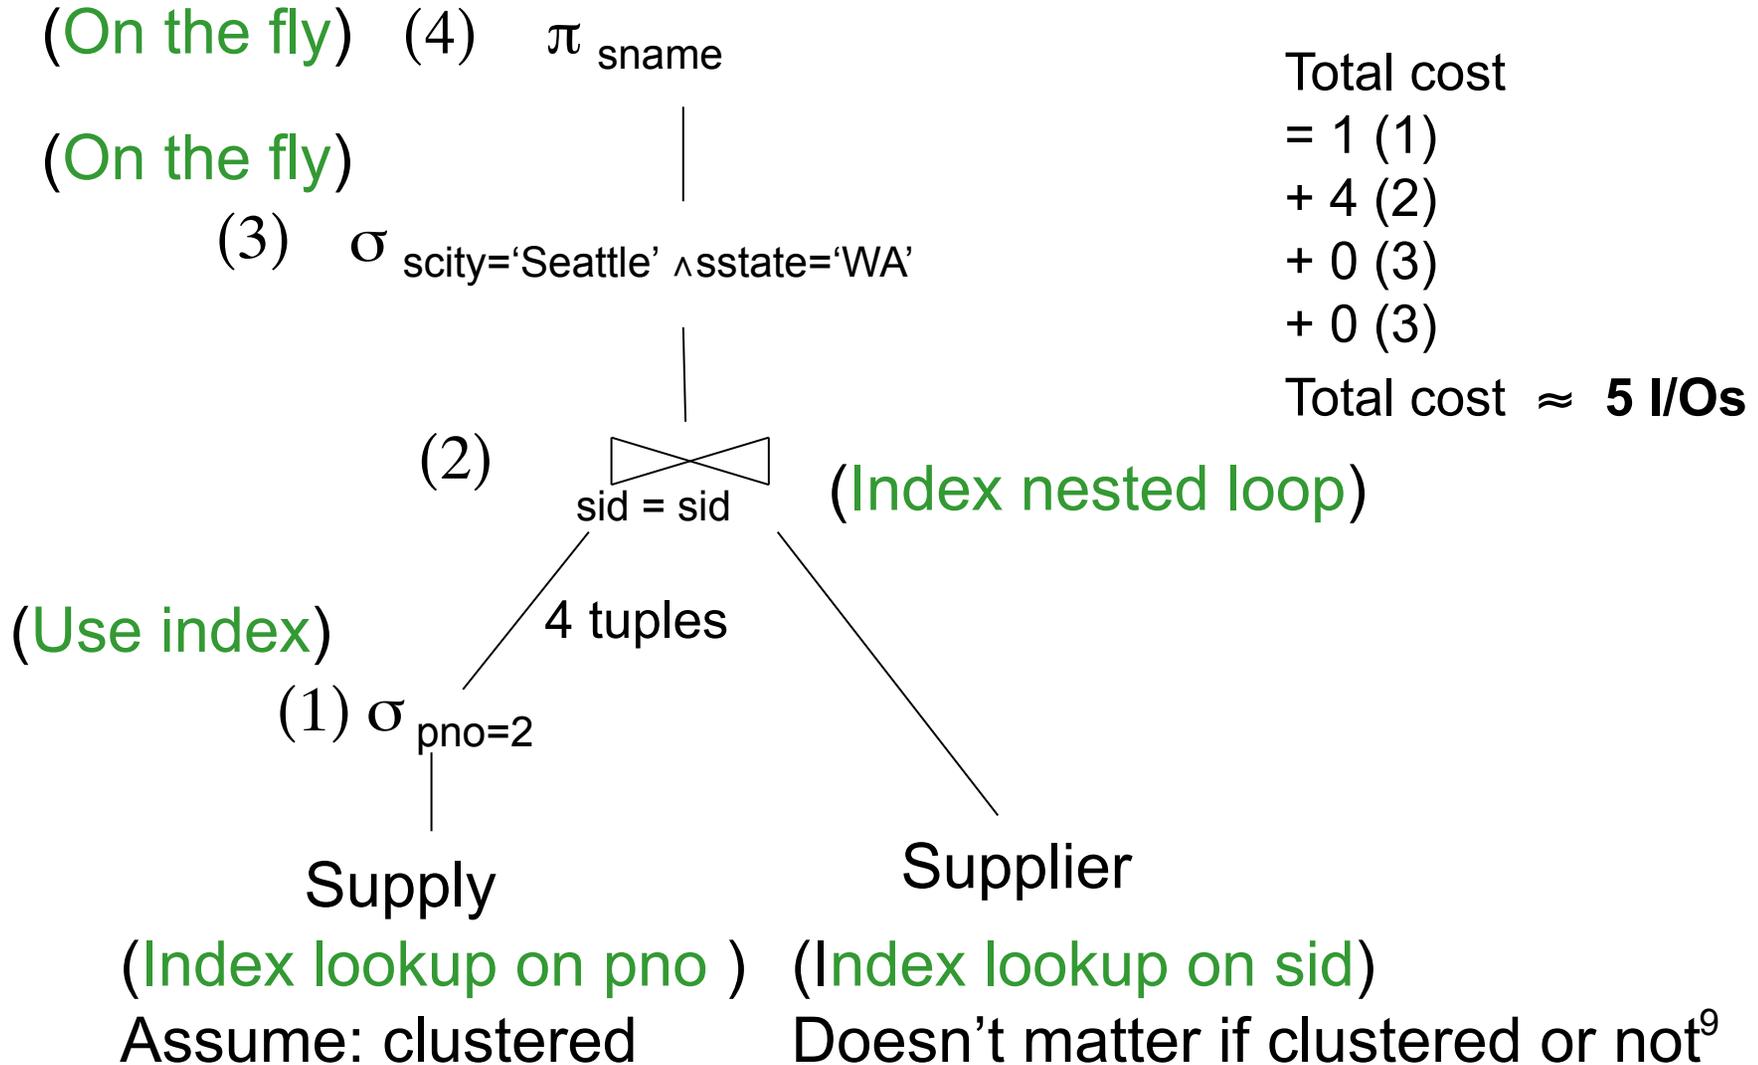
T(Supplier) = 1000  
T(Supply) = 10,000

B(Supplier) = 100  
B(Supply) = 100

V(Supplier,scity) = 20  
V(Supplier,state) = 10  
V(Supply,pno) = 2,500

M = 10

# Physical Query Plan 3



# Simplifications

- In the previous examples, we assumed that all index pages were in memory
- When this is not the case, we need to add the cost of fetching index pages from disk

# Lessons

- Need to consider several physical plan
  - 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

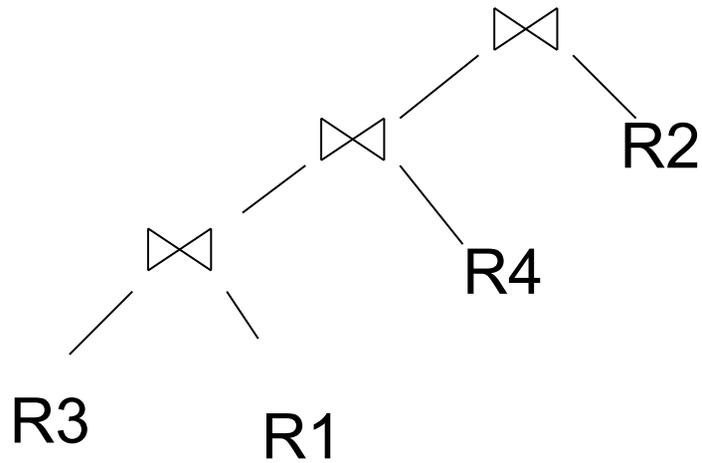
# Outline

- Search space (Today)
- Algorithm for enumerating query plans
- Estimating the cost of a query plan

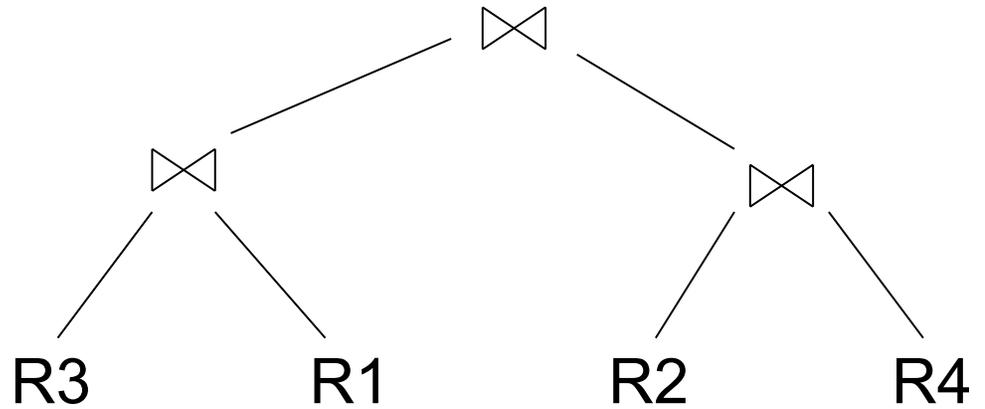
# Relational Algebra Equivalences

- 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
- Joins
  - Commutative :  $R \bowtie S$  same as  $S \bowtie R$
  - Associative:  $R \bowtie (S \bowtie T)$  same as  $(R \bowtie S) \bowtie T$

# Left-Deep Plans and Bushy Plans



Left-deep plan



Bushy plan

# Commutativity, Associativity, Distributivity

$$\begin{aligned}R \cup S &= S \cup R, & R \cup (S \cup T) &= (R \cup S) \cup T \\R \times S &= S \times R, & R \times (S \times T) &= (R \times S) \times T \\R \times S &= S \times R, & R \times (S \times T) &= (R \times S) \times T\end{aligned}$$

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

# Example

Which plan is more efficient ?

$R \bowtie (S \bowtie T)$  or  $(R \bowtie S) \bowtie T$  ?

- Assumptions:
  - Every join selectivity is 10%
    - That is:  $T(R \bowtie S) = 0.1 * T(R) * T(S)$  etc.
  - $B(R)=100$ ,  $B(S) = 50$ ,  $B(T)=500$
  - All joins are main memory joins
  - All intermediate results are materialized

# Laws involving selection:

$$\begin{aligned}\sigma_{C \text{ AND } C'}(R) &= \sigma_C(\sigma_{C'}(R)) = \sigma_C(R) \cap \sigma_{C'}(R) \\ \sigma_{C \text{ OR } C'}(R) &= \sigma_C(R) \cup \sigma_{C'}(R) \\ \sigma_C(R \bowtie S) &= \sigma_C(R) \bowtie S\end{aligned}$$

$$\begin{aligned}\sigma_C(R - S) &= \sigma_C(R) - S \\ \sigma_C(R \cup S) &= \sigma_C(R) \cup \sigma_C(S) \\ \sigma_C(R \bowtie S) &= \sigma_C(R) \bowtie S\end{aligned}$$

When C involves  
only attributes of R

# Example:

## Simple Algebraic Laws

- Example:  $R(A, B, C, D), S(E, F, G)$

$$\sigma_{F=3} (R \bowtie_{D=E} S) = \quad ?$$

$$\sigma_{A=5 \text{ AND } G=9} (R \bowtie_{D=E} S) = \quad ?$$

# Laws Involving Projections

$$\begin{aligned}\Pi_M(R \bowtie S) &= \Pi_M(\Pi_P(R) \bowtie \Pi_Q(S)) \\ \Pi_M(\Pi_N(R)) &= \Pi_M(R) \quad /* \text{ note that } M \subseteq N */\end{aligned}$$

- Example  $R(A,B,C,D)$ ,  $S(E, F, G)$

$$\Pi_{A,B,G}(R \bowtie_{D=E} S) = \Pi_{?}(\Pi_{?}(R) \bowtie_{D=E} \Pi_{?}(S))$$

# Laws involving grouping and aggregation

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

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

if agg is “duplicate insensitive”

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

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

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

# Laws Involving Constraints

Foreign key

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

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

Need a second constraint for this law to hold. Which one ?

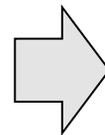
# Example

Foreign key

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

```
CREATE VIEW CheapProductCompany
SELECT *
FROM Product x, Company y
WHERE x.cid = y.cid and x.price < 100
```

```
SELECT pname, price
FROM CheapProductCompany
```



```
SELECT pname, price
FROM Product
```

# Laws with Semijoins

Recall the definition of a semijoin:

- $R \bowtie S = \Pi_{A_1, \dots, A_n} (R \Join S)$
- Where the schemas are:
  - Input:  $R(A_1, \dots, A_n)$ ,  $S(B_1, \dots, B_m)$
  - Output:  $T(A_1, \dots, A_n)$

# Laws with Semijoins

Semijoins: a bit of theory (see *Database Theory*, AHV)

- Given a query:

$$Q = R_1 \bowtie R_2 \bowtie \dots \bowtie R_n$$

- A semijoin reducer for Q is

$$\begin{aligned} R_{i1} &= R_{i1} \times R_{j1} \\ R_{i2} &= R_{i2} \times R_{j2} \\ &\dots \\ R_{ip} &= R_{ip} \times R_{jp} \end{aligned}$$

such that the query is equivalent to:

$$Q = R_{k1} \bowtie R_{k2} \bowtie \dots \bowtie R_{kn}$$

- A full reducer is such that no dangling tuples remain

# Laws with Semijoins

- Example:

$$Q = R(A,B) \bowtie S(B,C)$$

- A reducer is:

$$R_1(A,B) = R(A,B) \bowtie S(B,C)$$

- The rewritten query is:

$$Q = R_1(A,B) \bowtie S(B,C)$$

Why would we do this ?

# Why Would We Do This ?

- Large attributes:

$$Q = R(A,B, D, E, F, \dots) \bowtie S(B,C, M, K, L, \dots)$$

- Expensive side computations

$$Q = \gamma_{A,B, \text{count}(*)} R(A,B,D) \bowtie \sigma_{C=\text{value}}(S(B,C))$$

$$\begin{aligned} R_1(A,B,D) &= R(A,B,D) \bowtie \sigma_{C=\text{value}}(S(B,C)) \\ Q &= \gamma_{A,B, \text{count}(*)} R_1(A,B,D) \bowtie \sigma_{C=\text{value}}(S(B,C)) \end{aligned}$$

# Laws with Semijoins

- Example:

$$Q = R(A,B) \bowtie S(B,C)$$

- A reducer is:

$$R_1(A,B) = R(A,B) \bowtie S(B,C)$$

- The rewritten query is:

$$Q = R_1(A,B) \bowtie S(B,C)$$

Are there dangling tuples ?

# Laws with Semijoins

- Example:

$$Q = R(A,B) \bowtie S(B,C)$$

- A full reducer is:

$$\begin{aligned} R_1(A,B) &= R(A,B) \bowtie S(B,C) \\ S_1(B,C) &= S(B,C) \bowtie R_1(A,B) \end{aligned}$$

- The rewritten query is:

$$Q :- R_1(A,B) \bowtie S_1(B,C)$$

No more dangling tuples

# Laws with Semijoins

- More complex example:

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

- A full reducer is:

$$\begin{aligned} S'(B,C) &:= S(B,C) \bowtie R(A,B) \\ T'(C,D,E) &:= T(C,D,E) \bowtie S(B,C) \\ S''(B,C) &:= S'(B,C) \bowtie T'(C,D,E) \\ R'(A,B) &:= R(A,B) \bowtie S''(B,C) \end{aligned}$$

$$Q = R'(A,B) \bowtie S''(B,C) \bowtie T'(C,D,E)$$

# Laws with Semijoins

- Example:

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

- Doesn't have a full reducer (we can reduce forever)

**Theorem** a query has a full reducer iff it is “acyclic”  
[*Database Theory*, by Abiteboul, Hull, Vianu]

# Example with Semijoins

Emp(eid, ename, sal, did)

Dept(did, dname, budget)

DeptAvgSal(did, avgsal) /\* view \*/

[Chaudhuri'98]

View:

```
CREATE VIEW DepAvgSal As (  
    SELECT E.did, Avg(E.Sal) AS avgsal  
    FROM Emp E  
    GROUP BY E.did)
```

Query:

```
SELECT E.eid, E.sal  
FROM Emp E, Dept D, DepAvgSal V  
WHERE E.did = D.did AND E.did = V.did  
    AND E.age < 30 AND D.budget > 100k  
    AND E.sal > V.avgsal
```

Goal: compute only the necessary part of the view

# Example with Semijoins

Emp(eid, ename, sal, did)

Dept(did, dname, budget)

DeptAvgSal(did, avgsal) /\* view \*/

[Chaudhuri'98]

New view  
uses a reducer:

```
CREATE VIEW LimitedAvgSal As (  
    SELECT E.did, Avg(E.Sal) AS avgsal  
    FROM Emp E, Dept D  
    WHERE E.did = D.did AND D.budget > 100k  
    GROUP BY E.did)
```

New query:

```
SELECT E.eid, E.sal  
FROM Emp E, Dept D, LimitedAvgSal V  
WHERE E.did = D.did AND E.did = V.did  
    AND E.age < 30 AND D.budget > 100k  
    AND E.sal > V.avgsal
```

# Example with Semijoins

Emp(eid, ename, sal, did)

Dept(did, dname, budget)

[Chaudhuri'98]

DeptAvgSal(did, avgsal) /\* view \*/

Full reducer:

```
CREATE VIEW PartialResult AS
  (SELECT E.eid, E.sal, E.did
   FROM Emp E, Dept D
   WHERE E.did=D.did AND E.age < 30
   AND D.budget > 100k)

CREATE VIEW Filter AS
  (SELECT DISTINCT P.did FROM PartialResult P)

CREATE VIEW LimitedAvgSal AS
  (SELECT E.did, Avg(E.Sal) AS avgsal
   FROM Emp E, Filter F
   WHERE E.did = F.did GROUP BY E.did)
```

# Example with Semijoins

New query:

```
SELECT P.eid, P.sal  
FROM PartialResult P, LimitedDepAvgSal V  
WHERE P.did = V.did AND P.sal > V.avgсал
```

# 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