

## Compilation

A query is first parsed into a parse tree based on the grammar for SQL.

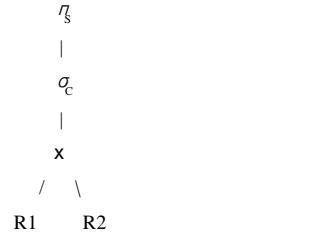
The parse tree is then transformed into a query plan, which is a tree with leaves labeled as relations and non-leaves labeled as relational algebra operations (with subscripts); the children of a node are the arguments of the operation labeling that node.

The query plan is then optimized by transforming it into one that is more efficient.

1

## Query Plans

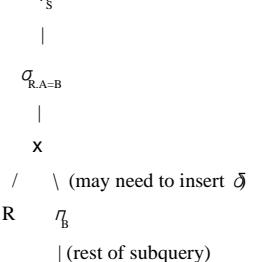
A simple query like SELECT S FROM R1, R2 WHERE C; is translated into a query plan like this:



2

## Subqueries

SELECT S FROM R WHERE A IN SELECT B ...; if inner select is uncorrelated:



3

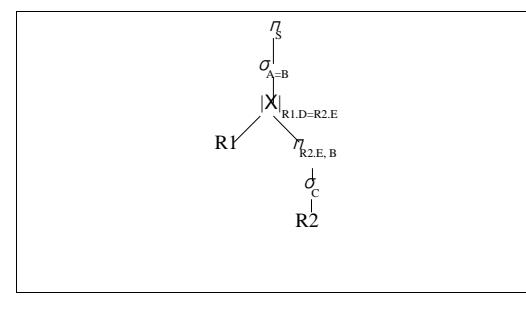
## Correlated Subqueries

Correlated subqueries must be modified pulling out the part of the condition containing the correlation variables, extend the tuples returned by the subqueries with extra attributes, and selecting on the basis of the pulled-out condition after joining the relation(s) with the results of the subqueries.

Example: SELECT S FROM R1  
WHERE A IN SELECT B FROM R2  
WHERE C AND R1.D=R2.E

4

## Query Plan for Example



5

## Transformations

One way to improve query plans is to reduce the size and number of tuples being processed as soon as possible. This can be done with transformations based on equivalences.

Some transformations are also used to break a query down into a combination of simpler queries that can be more efficiently computed. Example: a selection involving required values of two attributes ( $A1=v1$  AND  $A2=v2$ ) can be restructured to do an efficient selection using only one of the attributes, followed by a selection based on the other attribute.

6

## Selection Transformations

$$\sigma_{C_1 \text{ AND } C_2}(R) = \sigma_{C_1}(\sigma_{C_2}(R)) = \sigma_{C_2}(\sigma_{C_1}(R))$$

$$\sigma_{C_1 \text{ OR } C_2}(R) = \sigma_{C_1}(R) \cup \sigma_{C_2}(R)$$

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

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

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

If C depends only on attributes in R, [similar equations if

$$\sigma_C(R \cap S) = \sigma_C(R) \cap S \quad C \text{ depends only on } S$$

$$\sigma_C(R \times S) = \sigma_C(R) \times S \quad (\text{similarly for joins})$$

7

## More Selection

If R has the attributes in C and S has the attributes in C,

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

$$\sigma_C(R | X| S) = \sigma_C(R) | X| \sigma_C(S)$$

Note: sometimes several equations can be combined to get an improvement. Example:

$$\sigma_C(R | X| S) = \sigma_C(R | X| S) = \sigma_C(R) | X| \sigma_C(S)$$

if R has the attributes in C and S has the attributes in C.

8

## Projection

In subscript  $\underset{\text{expr} \rightarrow x}{\text{expr}}$  the attributes in  $\underset{\text{expr}}{\text{expr}}$  are input attributes and  $x$  is an output attribute. In subscript  $\underset{x}{x}$  x is both an input and output attribute.

If L is a projection list, M is the join attributes common to R and S and the input attributes of L that are in R, and N is the join attributes common to R and S and the input attributes of L that are in S, then

$$\pi_L(R | X| S) = \pi_L(\pi_M(R) | X| \pi_N(S))$$

Similarly for other joins, and for product.

9

## More Projection

$$\pi_L(R \cup_B S) = \pi_L(R) \cup_B \pi_L(S)$$

(Won't work for set union, for intersection or difference)

The following is sometimes useful

$$\pi_L(\sigma_C(R)) = \pi_L(\sigma_C(\pi_M(R)))$$

where M lists the input attributes of L and the attributes in C.

10

## Joins, Products

$$R | X|_C S = \sigma_C(R \times S)$$

$R | X|_C S = \pi_L(\sigma_C(R \times S))$  where C equates the attributes that are common to R and S, and L lists all the attributes of R and S but without duplication of the common attributes.

The above equations are usually used to transform products into joins, which are more efficient to compute.

11

## Duplication Elimination

$$\delta R = R$$

Useful when we can prove that R will have no duplicates (a stored relation with a primary key, the result of a grouping operation, or a set operation).

$$\delta R \times S = \delta R \times \delta S$$

Similarly for joins.

$$\delta \sigma_C(R) = \sigma_C(\delta R)$$

$$\delta R \cap_B S = \delta R \cap_B S = R \cap_B \delta S = \delta R \cap_B \delta S$$

12

## Grouping

$$\delta \chi_L(R) = \chi_L(R)$$

$\chi_L(R) = \chi_L(\delta R)$  provided L involves only MIN and MAX aggregate operators.

$\chi_L(R) = \chi_L(\chi_M(R))$  where M contains at least the attributes of R that are in L.

13

## Other Transformations

Operators that are associative and commutative can be rearranged for increased efficiency.

If natural joins and theta joins are combined, the natural joins can be converted to theta joins so that the joins can be rearranged.

14

## Estimating Size of a Selection

$$S = \sigma_{a=v}(R) \text{ Assume } T(S) = T(R)/V(R,a)$$

$$S = \sigma_{a<v}(R) \text{ Assume } T(S) = T(R)/3 \text{ (a heuristic)}$$

$$S = \sigma_{a_v}(R) \text{ Assume } T(S) = T(R) \text{ or assume } T(S) = T(R)*(V(R,a) - 1)/V(R,a)$$

$$S = \sigma_{\text{not } C}(R) \text{ Assume } T(S) = T(R) - T(\sigma_C(R))$$

$$S = \sigma_{C_1 \text{ or } C_2}(R) \text{ Assume } T(S) = T(\sigma_{\text{not(not } C_1 \text{ and not } C_2)}(R))$$

$$\text{Remember that } \sigma_{C_1 \text{ and } C_2}(R) = \sigma_{C_1}(\sigma_{C_2}(R))$$

15

## Estimating Size of a Join

Assume  $T(R | X| S) = T(R)*T(S)/\max(V(R,a), V(S,a))$  when R and S have one attribute (a) in common. To see this, assume uniform distribution of values for attribute a. For a given value of a, there are  $T(R)/V(R,a)$  tuples in R with that value, and  $T(S)/V(S,a)$  tuples in S with that value, so the number of join tuples with that value is  $(T(R)/V(R,a))*(T(S)/V(S,a))$ . There are  $V(R,a)$  distinct values of a in R, so the total number of join tuples is  $V(R,a)*(T(R)/V(R,a))*(T(S)/V(S,a)) = T(R)*T(S)/V(S,a)$ . Alternatively, multiply by  $V(S,a)$  to get  $T(R)*T(S)/V(R,a)$ . Use the smaller estimate (divide  $T(R)*T(S)$  by  $\max(V(R,a), V(S,a))$ ).

16

## Alternative Argument

For each of the  $T(R)$  tuples in R, there are  $T(S)/V(S,a)$  tuples in S with the same value for a, so the total number of join tuples is estimated to be  $T(R)*T(S)/V(S,a)$ . Or, starting with S, for each of the  $T(S)$  tuples in S, there are  $T(R)/V(R,a)$  tuples in R with the same value for a, so the total number of join tuples is estimated to be  $T(R)*T(S)/V(R,a)$ . Take the smaller estimate.

17

## Multiple Join Attributes

Suppose that the attributes common to R and S are  $a_1, \dots, a_k$ . Estimate  $T(R | X| S)$  this way:

$$T = T(R)*T(S)$$

$$\text{for } i \text{ from 1 to } k, T = T/\max(V(R,a_i), V(S,a_i))$$

$$T(R | X| S) = T$$

This assumes that the distributions of the attribute values are uniform and independent of each other.

18

## Multiple Joins

$T(R_1|X| \dots |X| R_n)$  can be estimated this way:

$$T = T(R_1) * \dots * T(R_n)$$

for each attribute  $a$  that is common to two or more of the relations, if  $a$  is common to  $k$  relations,

$T = T /$  the product of the  $k-1$  largest values in  $\{V(R_1, a), \dots, V(R_n, a)\}$

$$T(R_1|X| \dots |X| R_n) = T$$

19

## Heuristics for Other Operations

Estimate  $T(R \cup_B S) = T(R) + T(S)$

Estimate  $T(R \cup_S S) =$  The larger +  $\frac{1}{2}$  of the smaller

Estimate  $T(R \cap S) = \frac{1}{2}$  of the smaller

Estimate  $T(R - S) = T(R) - T(S)/2$

Estimate  $T(\delta R(a_1, \dots, a_n)) =$  smaller of  $T(R)/2$  and  $V(R, a_1) * \dots * V(R, a_n)$

Estimate  $T(\chi_L(R)) =$  smaller of  $T(R)/2$  and  $V(R, a_1) * \dots * V(R, a_n)$

where the  $a_i$  are the grouping attributes.

20

## Enumeration of Physical Query Plans

1. **Exhaustive enumeration** – Try all possible transformations on the original query plan, assign all possible algorithms for computing the individual operations, compute the disk access estimates.
2. **Heuristic selection** – Use heuristics to choose which algorithm to use for each operation.

21

## Examples of Heuristics

- To compute  $\sigma_{a=c}(R)$ , if  $R$  has an index on  $a$ , use the index to look up the tuples with  $a = c$ .
- To compute  $\sigma_{a=c \text{ and } C}(R)$ , if  $R$  has an index on  $a$ , use the index to look up the tuples with  $a = c$ , and immediately test whether  $C$  is true on each tuple as it is retrieved.
- To compute a join, if one relation has an index on the join attributes, use an index-join with that relation in inner loop.
- To compute a join, if one relation is sorted on the join attributes, choose a sort-join over a hash-join.
- For union or intersection of 3 or more relations, combine the smallest relations first.

22

## More Enumerations

3. **Branch-and-bound** – Abandon a physical plan as soon as part of it costs more than the best plan found so far. If the whole plan costs less than the best plan so far, it becomes the new best plan found so far.
4. **Hill climbing** – Start from one plan. Consider several changes to it, choose the change that has the least overall cost. Common changes are choice of algorithm for an operator, or order of doing associative operations.
5. **Dynamic programming** – Working from leaves to root, choose the best way to compute the operation at a node given the previously computed best way for computing the operations at its child nodes.

23

## More Enumerations cont.

6. **Selinger-style optimization** – Keep several plans for computing the operation at a node, the best as in dynamic programming, and some that produce tuples in an order that might be useful higher up in the plan.

24

## Heuristics for Joins

- For the join of two relations, let the smaller relation be the **build** relation and the larger relation be the **probe** relation.
- For the join of three or more relations, use only **left-deep** (left-branching) join trees. There are fewer of them and they tend to be more efficient.
- Use a **greedy algorithm** to do join of three or more relations
  - Start with the smallest relation, join it with the next smallest relation, join the result with the next smallest relation, etc.

25

## Choose Selection Methods

- For  $\sigma_c(R)$  and no useful index on R, scan all of R, testing for C. Cost is  $B(R)$  or  $T(R)$ .
- For  $\sigma_{a=c \text{ and } C}(R)$  and an index on a, retrieve tuples for  $a = c$  and filter them with C test. Cost is  $B(R)/V(R,a)$  or  $T(R)/V(R,a)$ .
- For  $\sigma_{a < c \text{ and } C}(R)$  and other inequalities on a, and an index on a, retrieve the tuples satisfying the inequality and filter them with C test. Cost is  $B(R)/3$  or  $T(R)/3$ .

26

## Choose Join Methods

- Choose among sort-joins, hash-joins, indexed joins if we have enough info about the size of the tables and the space available in memory for buffers.
- Choose a one-pass join (all of the smaller table is put in memory and accessed by the join attributes), or if the smaller table is not much larger than memory, use a nested-loop join, with the smaller table in the outer loop.
- Choose a sort-join if one or more tables are already sorted on the join attributes or two or more joins use the same join attributes.
- If one table is very small and the other table is indexed on the join attributes, use an indexed join.

27

## Choose Join Methods cont.

- If tables are not already sorted or indexed and a multi-pass method is needed, choose a hash-join.

28

## Pipelining, Materialization

Writing the output of an operation to disc is called **materialization**. It is more efficient to **pipeline**; use the output buffer block as an input buffer block for the next operation. Iterators facilitate the necessary interleaving of code for the two operations.

Split the physical query plan into subtrees at the boundaries where intermediate tables get materialized. Generate code for each subtree using iterators. Put the blocks of code in a bottom-up, left-to-right order.

29