

# THE EFFECT OF JOIN SELECTIVITIES ON OPTIMAL NESTING ORDER

*Akhil Kumar and Michael Stonebraker*

*EECS Department  
University of California  
Berkeley, Ca., 94720*

## Abstract

A heuristic query optimizer must choose the best way to process an incoming query. This choice is based on comparing the expected cost of many (or all) of the ways that a command might be processed. This expected cost calculation is determined by statistics on the sizes of the relations involved and the selectivities of the operations being performed. Of course, such estimates are subject to error, and in this paper we investigate the sensitivity of the best query plan to errors in the selectivity estimates. We treat the common case of join queries and show that the optimal plan for most queries is very insensitive to selectivity inaccuracies. Hence, there is little reason for a data manager to spend a lot of effort making accurate estimates of join selectivities.

## 1. INTRODUCTION

In a relational database system, relations are accessed by specifying queries in a query language such as SQL [DATE84] or QUEL [HELD75]. Each user command is processed by a heuristic query optimizer at compile time or at run time to select the minimum expected cost query execution plan. This plan is then executed to deliver the result to the user who submitted the command.

Consider, for example, the following selection query on the EMP relation:

```
retrieve (EMP.name)
where EMP.age > 50 and EMP.salary < 1000
```

If there are secondary indexes on both age and salary, then the optimizer must choose between 3 processing plans, namely scan the whole EMP relation, scan a portion of the age index and then access qualifying EMP records, or scan a portion of the salary index. Which plan has least cost is clearly very sensitive to the expected number of employees who have low salaries and the expected number who have high ages. Hence, data managers accumulate statistics from which such estimates can be made. Current systems keep track of some combination of:

```
the number of tuples in each relation
the minimum and maximum key values for any indexed field
```

---

This research was sponsored by the National Science Foundation under Grant DMC-8504633 and by the Navy Electronics Systems Command under contract N00039-84-C-0039.

a histogram of the number of tuples within various key ranges

Work on optimizing the collection of statistics is reported in [PIAT84, CHRI83, ROWE83].

In the case that a query involves one or more joins, the query optimizer must estimate the number of tuples that will appear in any intermediate join. For example, one might run the following three-way join among the EMP, DEPT and PROJECT relations to find a list of the names of employees along with their department and the project that they work on.

```
retrieve (EMP.name, DEPT.dname, PROJECT.name)
where EMP.dept = DEPT.dname and EMP.name = PROJECT.ename
```

If the optimizer chooses to join EMP and DEPT first, then it must guess the expected size of the partial answer found from:

```
retrieve (EMP.name, DEPT.dname)
where EMP.dept = DEPT.dname
```

Of course, the cost of the remaining computation, that of joining the partial result with the PROJECT relation must also be estimated. Similarly, the query optimizer must guess the size of any other 2-way joins that are constructed by a query plan which it investigates. One way to compute the size of the join of two relations,  $R_i$  and  $R_j$  is:

$$n_{ij} = s_{ij} \times (n_i \times n_j)$$

where

$n_{ij}$  : number of tuples in the result of the above join  
 $s_{ij}$  : Selectivity of a join between relation  $R_i$  and  $R_j$   
 $n_i$  : number of tuples in relation  $R_i$   
 $n_j$  : number of tuples in relation  $R_j$

Of course,  $s_{ij}$  must be estimated, and current optimizers (e.g. [SELI79], [KOOI82]) utilize very crude techniques.

Whenever the actual selectivity differs from the estimated value, then the size of the actual intermediate relation will differ from its estimated value and the actual cost of the corresponding query plan will differ from its estimated value. Consequently, selectivity errors can lead to choosing a plan with minimum estimated cost but whose actual cost may be greater than the actual cost of other possible plans. The purpose of this paper is to study the relationship between inaccuracies in estimates of join selectivities and the actual quality of the minimum cost plan selected using inaccurate statistics. This actual quality is found by comparing the selected plan with the actual optimal plan. In general, we have found that plan quality is very insensitive to inaccurate join selectivities. Hence, there is little reason to spend a lot of effort in constructing accurate statistics in this area. This result should be contrasted with [EPST80] where a different conclusion was reached for a distributed data base environment.

The remainder of this paper is organized as follows. In Section 2, we describe a simulated query planner similar to the one in [SELI79] which was our experimental tool. Section 3 turns to the design of our experiments and introduces the concept of a Q-factor which can be used to measure the sensitivity of a join order to variations in selectivity. Then, in Section 4 the results of our experiments are presented and analyzed in detail.

## 2. THE QUERY PLANNER

To perform this study we implemented a query planner which accepts 4-variable and 5-variable queries and generates the optimal plan for executing them. The planner uses the same heuristics to determine the optimal plan as the System R planner described in [SELI79]. This section discusses the assumptions made by our planner.

First, our planner accepts only join queries and assumes that select and project operations have already been performed. This allows us to concentrate on join selectivities exclusively.

The planner accepts the following inputs:

- 1) Number of relations involved
- 2) Cardinality of each relation
- 3) Tuple size of each relation
- 4) The relations which are sorted and the field on which they are sorted
- 5) The relations having secondary indexes and the fields on which they exist
- 6) Number of joins
- 7) The relations involved in each join and the joining fields from each relation
- 8) The selectivity of each join

Based on the above inputs the planner generates a plan that gives the following:

- 1) The sequence in which the joins should be performed
- 2) The kind of join to be performed -- merge-sort join or nested loop join
- 3) Whether any secondary indexes are used

Our query planner accepts both chain queries and non-chain queries. An example 4-variable chain query is the following:

R1 JOIN1 R2 JOIN2 R3 JOIN3 R4

An example non-chain query is:

R1 JOIN1 R2 AND R1 JOIN2 R3 AND R1 JOIN3 R4

Both chain and non-chain queries may be represented graphically as shown in Figure 1. The join clauses in a query containing n joins are numbered:

$c_1, c_2, \dots, c_n$

For example, the chain query noted above has three join clauses,  $c_1$ ,  $c_2$ , and  $c_3$ .

The other major features of the planner are listed below.


1. The planner works by enumerating certain possible join sequences (also called nesting orders). As in [SELI79], we do not consider joining two relations unless they have a join clause present in the query. Hence, the number of join sequences considered is the number of permutations of the join clauses.
2. A query with n joins is executed by performing n 2-way joins --  $J_1, J_2, \dots, J_i, \dots, J_n$ . The cost of a 2 way join is calculated as follows:

$$CJ_i = (\# \text{ of disk I/O's})_i + M \times (\text{size of answer})_i$$

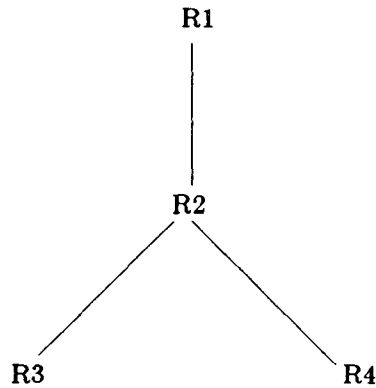
where

$CJ_i$  : cost of performing the  $i^{\text{th}}$  2-way join,  $J_i$

$(\# \text{ of disk I/O's})_i$  : number of disk I/O's required for join  $J_i$



(a) Chain Query



(b) Non-Chain Query

Figure 1: Query Graphs for Chain and Non-Chain Queries

---

(size of answer)<sub>i</sub> : size of the answer of the  $i^{th}$  join,  $J_i$

The first term in the expression above represents the disk cost of processing the query while the expected size of the answer represents a measure of the CPU cost. The constant  $M$  is used to relate the different kinds of costs, and in our experiments we set  $M$  to 0.5. The total cost,  $CJ$  of a plan is the sum of the costs of the individual joins as follows:

$$CJ = \sum_i CJ_i$$

3. The planner considers two types of joins, nested loop joins and merge sort joins ([SELI79]). The nested loop join is performed by accessing all tuples of one relation (the outer relation,  $R_O$ ) and matching each tuple against all the tuples of the inner relation,  $R_I$ . A fixed number of buffer pool pages are assigned to each relation. The correct number of pages from the outer relation are brought into memory. The tuples on each of these pages are compared against all the tuples of  $R_I$  which are loaded into the buffer in groups of pages. The pages allocated to the inner relation are then replaced by a new collection, and the above process is repeated. Thus, each tuple of  $R_O$  is brought into the buffer just once, while the

tuples of  $R_i$  are brought in repeatedly till the join is completed. In our experiments, a 1-page buffer (4096 bytes) was assigned to each relation in memory.

In case a secondary index exists on the joining field of the inner relation, the alternative of accessing the tuples of the inner relation via this index is also considered. Performing a merge-sort join on already sorted relations requires loading each relation iteratively into memory and then writing the result back. In case one or both relations are not already sorted, the cost of sorting the relations is added in determining the cost of the join.

4. For every join sequence being considered, the planner computes the cost of the individual joins using the cheapest method and adds up the cost of all joins in the sequence to determine the total cost. After computing the cost for each join order, the one with the lowest cost is selected as the optimal plan.
5. We assume that all the join clauses are independent. For example, in the chain query above, if the join between R1 and R2 was done first, the next join between R12 and R3 would have the same selectivity as the join between the original R2 and R3.
6. The page size was set to 4096 bytes.
7. Secondary indexes are implemented as B-trees. In order to determine the number of levels within a B-tree index, we assume an index page contains 100 entries. Hence a  $10^6$ -tuple relation requires a 3-level index while a  $10^7$ -tuple relation would require a 4-level index. If the tuples of the inner relation are accessed via an index while performing a nested loop join, the root page of the index is loaded into memory once and kept in the buffer for subsequent use while the lower-level index pages are fetched from disk every time a new tuple has to be accessed.

### 3. EXPERIMENTAL SETUP

#### 3.1. Design of Experiments

The purpose of the experiments was to investigate the effect of varying selectivities on the optimal nesting order. We have performed several experiments on two collections of 4-variable and 5-variable queries. Figure 2 gives a list of 10 4-variable queries that we chose for the experiments. Figure 3 gives a list of 10 5-variable queries while Figure 4 gives the sizes of the relations considered. Each relation is assumed to be 100 bytes wide. In each group we have included an equal number of chain and non-chain queries. They represent a good cross-section of queries that are likely to arise in practice. Each group is composed of queries with varying number of secondary indexes and relations that are sorted on the joining fields. Hence, cases optimizable by both merge-sort and nested loops iteration are included.

In all of our experiments we have varied the selectivities of one or more join clauses over a range of values from 0.33 to 3 times an initial value. This corresponds to an order of magnitude difference between the smallest and the largest selectivity values used in a range. The initial values of selectivities were chosen such that the size of the result of a multi-way join was of the order of the size of the smallest relation in the join. We have observed that a very large number of real world queries conform to this rule. Thus, in a two way join between R1 having  $10^3$  tuples and R2 having  $10^5$  tuples, we would expect an answer of  $10^3$  tuples. This rule served as a guideline in determining the initial

- 
1. R1joinR2joinR3joinR4 (no indexes, no relation sorted)
  2. R1joinR2joinR3joinR4 (indexes on R1,R2,R3,R4)
  3. R1joinR2joinR3joinR4 (indexes on R1,R2,R3; R4 sorted)
  4. R1joinR2joinR3joinR4 (indexes on R2,R4; R1,R3 sorted)
  5. R1joinR2joinR3joinR4 (R2,R3,R4 sorted)
  6. R1joinR2,R2joinR3,R2joinR4 (indexes on R2,and R3; R1,R2,R3,R4 sorted)
  7. R1joinR2,R2joinR3,R2joinR4 (no indexes, no relations sorted)
  8. R1joinR3,R2joinR3,R4joinR3 (R2,R3,R4 have indexes; R1,R2 sorted)
  9. R1joinR4,R2joinR4,R4joinR3 (R1,R3,R4 have indexes)
  10. R1joinR2,R2joinR3,R4joinR2 (R1,R2,R3,R4 sorted)

Figure 2: List of 4-variable queries

---

1. R1joinR2joinR3joinR4joinR5(R1,R2,R3,R5 have indexes; all relations sorted)
2. R1joinR2joinR3joinR4joinR5 (indexes on all relations)
3. R1joinR2joinR3joinR4joinR5 (no indexes, no relations sorted)
4. R1joinR2joinR3joinR4joinR5 (indexes on R1,R2,R3; R4,R5 sorted)
5. R1joinR2joinR3joinR4joinR5 (R2,R3,R4,R5 sorted)
6. R1joinR2,R2joinR3,R2joinR4,R4joinR5 (indexes on all relations; R2,R3 sorted)
7. R1joinR2,R2joinR3,R3joinR4,R2joinR5 (indexes on R2,R3,R4,R5; R1,R2,R3 sorted)
8. R1joinR2,R2joinR4,R2joinR4,R2joinR5 (no indexes, no relations sorted)
9. R1joinR3,R2joinR3,R3joinR4,R4joinR5 (indexes on R1,R2,R3,R4)
10. R1joinR4,R2joinR4,R3joinR4,R3joinR5 (all relations sorted)

Figure 3: List of 5-variable queries

---

selectivity values which are also indicated in Figure 4.

Before discussing details of the experiments we introduce the concept of Q-factor which is our metric for assessing the sensitivity of a nesting order.

### 3.2. Q-factor

The optimal query plan,  $p_{opt}$  for a query is the minimum cost sequence of joins. The performance of any other plan,  $p$  for a combination of selectivities can be judged by the following ratio:

$$r(p,j) = c(p,j) / c(p_{opt},j)$$

where

$r(p,j)$  : ratio of the cost of plan  $p$  to the cost of the optimal plan,  $p_{opt}$  for selectivity combination  $j$

$c(p,j)$  : cost of plan  $p$  for selectivity combination  $j$

By definition, a value of  $r(p,j)$  close to 1 indicates a good plan. A plan which is optimal

---

#### 4-Variable Queries

Relation sizes : R1=10<sup>7</sup>, R2=10<sup>3</sup>, R3=10<sup>6</sup>, R4=10<sup>5</sup>

Initial Selectivities of join clauses :  $s_{c_1}=10^{-7}$ ,  $s_{c_2}=10^{-6}$ ,  $s_{c_3}=10^{-5}$

#### 5-Variable Queries

Relation sizes : R1=10<sup>7</sup>, R2=10<sup>3</sup>, R3=10<sup>5</sup>, R4=10<sup>6</sup>, R5=10<sup>4</sup>

Initial Selectivities of join clauses :  $s_{c_1}=10^{-7}$ ,  $s_{c_2}=10^{-5}$ ,  $s_{c_3}=10^{-6}$ ,  $s_{c_4}=10^{-4}$

Figure 4: Relation Sizes and Selectivities for the sample queries

---

for one set of selectivity values may not continue to be optimal as the selectivity values are varied. Hence, we define the Quality factor (or Q-factor), Q for a query as:

$$Q = \min_{p, j \in J} (\text{avg } r(p,j))$$

where J : set of selectivity values j in the range over which selectivity is varied

The plan,  $p^*$  for which the above expression is minimized has the minimum average cost ratio  $r(p,j)$  over the range of selectivity values. Q is our metric for defining sensitivity of a query to variations in selectivity. The minimum value of Q is 1, and a value of Q close to 1 means that the query is insensitive to variations in selectivity. Thus, the cost of processing the query using  $p^*$  is likely to be very close to the optimal cost at all points over the range J. A higher value of Q denotes greater sensitivity to variations in selectivity and means that the cost of processing the query using  $p^*$  is likely to be considerably sub-optimal for some selectivity values.

A second definition for Quality factor, Q1 is the following:

$$Q1 = \min_p (\max_{j \in J} r(p,j))$$

Again, the plan  $\hat{p}$  for which the above expression is at its minimum is the best overall plan.

The difference between the two definitions for Quality factor is that the first chooses the plan with the best average performance over the range of selectivity values while the second selects the one with the best worst-case performance. Although the best plan determined using the alternative definitions of Quality factor will not always be the same, it will be seen that both metrics usually produce the same plan.

We use  $Q$  as our criterion for selecting the best plan in all our experiments. However, in the first set of experiments in Section 4.2 we also illustrate how  $Q1$  is determined.

## 4. RESULTS OF EXPERIMENTS

### 4.1. Introduction

Four sets of experiments were carried out on a set of 20 queries. In the first set of experiments described in Section 4.2, we varied the selectivity of one join clause over a wide range while the selectivity of another join clause was varied in inverse proportion so as to keep the size of the answer fixed. Section 4.3 discusses experiments in which the selectivity of only one clause was varied over the same range as before and the size of the answer was allowed to vary in the same proportion. In the third set of experiments reported in Section 4.4, the size of the answer was allowed to increase in intervals by increasing each selectivity by a constant factor. Section 4.5 describes experiments in which the size of one of the relations was allowed to vary over a wide range while a corresponding selectivity value was also varied so as to keep the size of the result equal to the size of the smallest relation. Finally Section 4.6 summarizes the conclusions that can be drawn from the four sets of experiments.

### 4.2. Fixed Answer Size

This experiment was conducted on the set of 4-variable and 5-variable queries explained in the previous section. We varied the selectivity of clause  $c_1$  from 0.33 of its initial value to 3 times the initial value. The selectivity of clause  $c_3$  was changed in inverse proportion in order to keep the size of the answer fixed. At 8 different points over the selectivity range, we determined the cost of executing each query by the alternative plans discussed in Section 2 and computed  $Q$ ,  $Q1$ ,  $p^*$ ,  $\hat{p}$ ,  $r(p^*,j)$ , and  $r(\hat{p},j)$ . For all 20 queries  $\hat{p} = p^*$  and Tables 1 and 2 list  $r(p^*,j)$  for the various  $j \in J$  along with  $Q$  and  $Q1$ . The last row in Tables 1 and 2 gives the average values for the 10 queries in each set, and represents the expected performance of a query picked at random from our collection of queries.

The  $Q$ -factor for all 4-variable queries is less than 1.10 which clearly illustrates that on the average, the cost of  $p^*$  is within 10% of the cost of the optimal over the entire range of selectivity values. The  $Q1$  values in the last column of Table 1 show that in only 2 of the 10 queries (queries 2 and 3), the worst-case cost of  $\hat{p}$  is 50% more than the cost of the optimal plan. For the remaining 8 queries the worst-case cost of  $\hat{p}$  is considerably less.

The results of the experiments with 5-variable queries (Table 2) confirm the above observations. For 9 out of 10 queries, the value of  $Q$  is again less than 1.10. For 1 query (query 6) the value of  $Q$  is 1.47. Moreover, the value of  $Q1$  is less than 1.50 for 9 out of 10 5-variable queries.

The above experiments show that even when selectivities are varied over a wide range, the cost of  $\hat{p}$  and  $p^*$  is very close to the cost of the optimal plan for most queries.

### 4.3. Varying One Selectivity

The experiments of the previous section were modified slightly by allowing the selectivity of only one join clause to vary. We varied the selectivity of the first clause,  $c_1$  over



---

Query#	Selectivity Multiplier								Q	Q1
	0.33	0.5	0.75	1.0	1.5	2.0	2.5	3.0		
1	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.02	1.00	1.02
2	1.51	1.00	1.00	1.00	1.00	1.01	1.01	1.02	1.07	1.51
3	1.53	1.01	1.00	1.00	1.00	1.00	1.00	1.00	1.07	1.53
4	1.32	1.00	1.00	1.00	1.00	1.00	1.01	1.01	1.04	1.32
5	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.02	1.00	1.02
6	1.02	1.01	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.02
7	1.00	1.00	1.00	1.00	1.00	1.00	1.21	1.42	1.08	1.42
8	1.01	1.01	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.01
9	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
10	1.02	1.01	1.00	1.00	1.00	1.01	1.01	1.01	1.00	1.02
Avg	1.14	1.00	1.00	1.00	1.00	1.00	1.02	1.05	1.03	1.14

Table 1:  $r(p^*,j)$  ratios and Q-factors over a range of selectivity values (4-variable queries)

---

Query#	Selectivity Multiplier								Q	Q1
	0.33	0.5	0.75	1.0	1.5	2.0	2.5	3.0		
1	1.29	1.15	1.05	1.00	1.00	1.00	1.00	1.00	1.06	1.29
2	1.40	1.03	1.01	1.00	1.04	1.07	1.10	1.13	1.09	1.40
3	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
4	1.00	1.00	1.00	1.01	1.05	1.08	1.12	1.15	1.05	1.15
5	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
6	2.92	1.95	1.30	1.00	1.07	1.13	1.19	1.22	1.47	2.92
7	1.00	1.00	1.00	1.00	1.02	1.04	1.06	1.07	1.02	1.07
8	1.00	1.00	1.00	1.00	1.00	1.00	1.22	1.42	1.08	1.42
9	1.16	1.09	1.04	1.00	1.00	1.00	1.00	1.00	1.03	1.16
10	1.02	1.01	1.00	1.00	1.01	1.01	1.01	1.01	1.00	1.02
Avg	1.28	1.12	1.04	1.00	1.02	1.03	1.06	1.10	1.08	1.28

Table 2:  $r(p^*,j)$  ratios and Q-factors over a range of selectivity values (5-variable queries)

---

the same range as before. The rest of the parameters were the same as in the previous section. The experiments of the previous section were then repeated for the 4-variable and 5-variable queries. Notice that in this case the size of the answer was allowed to vary as the selectivity of the first clause was changed. We again computed  $p^*$ ,  $\hat{p}$ ,  $r(p^*,j)$ , and

$r(p, j)$ ,  $Q$  and  $Q1$ . Tables 3 and 4 give the  $Q$ -factors for the 4-variable and 5-variable queries, respectively.

The second set of experiments reinforce the results of the first set. In 7 of the 10 4-variable queries the  $Q$ -factor is 1.00. The remaining 3 queries have a  $Q$ -factor of less than 1.10. Moreover, 9 out of our 10 5-variable queries from Table 4 have a  $Q$ -factor within 1.06. Query 6 has a  $Q$ -factor of 1.20.

#### 4.4. Incrementing The Answer Size

For the experiments discussed in Section 4.2, the size of the answer was of the same order of magnitude as the smallest relation in the join. In this section we scaled all selectivities by a multiplier in order to increase the size of the answer. Starting with an answer of 1000 tuples, we let the answer size of our sample queries increase up to  $10^6$  in multiples of 10 by increasing the selectivity of each clause by appropriate amounts.

Tables 5 and 6 show the results for 4-variable and 5-variable queries, respectively. For each query the tables give the values of  $Q$ -factor as the size of the answer varies over the range from  $10^3$  to  $10^6$  tuples. The last row in each table gives the average  $Q$ -factor values for the 10 queries for specific answer sizes.

The largest  $Q$ -factor observed in Table 5 is 1.10 (query 4,  $10^6$  tuples). The largest average value across all 4-variable queries (see last row of Table 5) is 1.03. Among the 5-variable queries the average value of  $Q$ -factor across all queries (see last row of Table 6) is less than 1.10. The highest  $Q$ -factor in Table 6 is 1.47 (query 6,  $10^3$  tuples).

We observe that even though some queries (for example, queries 7 and 10 in Table 6) exhibit a  $Q$ -factor that increases with the answer size, the pattern is not maintained in the average values.

---

Query #	Q factor
1	1.00
2	1.07
3	1.06
4	1.04
5	1.00
6	1.00
7	1.00
8	1.00
9	1.00
10	1.00

Table 3:  $Q$ -factors for Best Nesting Orders -- Varying only one Selectivity (4-variable queries)

---

---

Query #	Q factor
1	1.01
2	1.06
3	1.00
4	1.03
5	1.00
6	1.20
7	1.01
8	1.00
9	1.01
10	1.00

Table 4: Q-factors for Best Nesting Orders -- Varying only one Selectivity  
(5-variable queries)

---

Query #	# of Tuples in Answer			
	$10^3$	$10^4$	$10^5$	$10^6$
1	1.00	1.00	1.00	1.00
2	1.07	1.02	1.05	1.01
3	1.07	1.01	1.03	1.01
4	1.04	1.01	1.02	1.10
5	1.00	1.00	1.00	1.00
6	1.00	1.01	1.02	1.02
7	1.08	1.00	1.00	1.02
8	1.00	1.01	1.03	1.10
9	1.00	1.00	1.00	1.00
10	1.00	1.01	1.01	1.02
Avg	1.03	1.01	1.02	1.03

Table 5: Q-factors for Best Nesting Orders -- Varying size of Answer  
(4-variable queries)

---

---

Query #	# of Tuples in Answer			
	$10^3$	$10^4$	$10^5$	$10^6$
1	1.06	1.09	1.00	1.01
2	1.09	1.11	1.05	1.00
3	1.00	1.00	1.00	1.00
4	1.05	1.05	1.00	1.00
5	1.00	1.00	1.00	1.00
6	1.47	1.32	1.11	1.04
7	1.02	1.06	1.19	1.33
8	1.08	1.00	1.00	1.02
9	1.03	1.12	1.07	1.00
10	1.00	1.03	1.15	1.32
Avg	1.08	1.08	1.06	1.07

Table 6: Q-factors for Best Nesting Orders -- Varying size of Answer (5-variable queries)

---

# of tuples	Q-factor	
	Query # 2	Query # 7
$10^2$	1.01	1.15
$10^3$	1.07	1.08
$10^4$	1.03	1.00
$10^5$	1.15	1.02
$10^6$	1.00	1.01
$10^7$	1.00	1.03

Table 7: Q-factors for Varying Size of Relation 2 (4-variable queries)

---

---

# of tuples	Q-factor	
	Query # 2	Query # 6
$10^2$	1.09	1.44
$10^3$	1.09	1.47
$10^4$	1.20	1.48
$10^5$	1.22	1.26
$10^6$	1.22	1.07
$10^7$	1.21	1.06

Table 8: Q-factors for Varying Size of Relation 2 (5-variable queries)

---

#### 4.5. Varying The Size of One Relation

The final set of experiments was done to study the effect of extreme variations in selectivities. We allowed the size of one relation to vary over a wide range and also allowed one of the selectivities to change so that the size of the result remained constant. In order to examine highly sensitive queries we selected the two queries with the highest Q-factor from Tables 1 and 2. For each query we increased the size of the second relation R2 from 100 tuples to  $10^7$  in multiples of 10. We adjusted the selectivities in such a way so as to get an answer equal to the size of the smallest relation. Tables 7 and 8 show the Q-factor values for various sizes of R2 for 4-variable and 5-variable queries respectively.

The experiments in this section were intentionally designed to represent a worst case. In fact, the highest Q-factor values for the 4-variable queries in Table 7 are 1.15 (for both queries 2 and 7). Other Q-factor values in the table are substantially lower. The highest Q-factor values for the 5-variable queries in Table 8 are 1.22 (query 2) and 1.48 (query 6).

#### 4.6. Analysis of The Results

In the four sets of experiments above, we have varied several important parameters for 20 different queries in order to get representative results. We have also identified the 4 most sensitive among the 20 queries and presented worst-case results for them. Among all the 20 queries, there was only one query (query 6 in Figure 3) which exhibited markedly greater sensitivity to changes in selectivity compared to others. It is the only query in which it is conceivable that better a priori estimation of selectivity could result in reasonable savings. For the remaining 19 out of 20 queries, it is highly unlikely that savings in execution cost would offset the additional cost incurred by constructing accurate estimates of selectivities.

Our results show an excellent average behavior of the best plan for both 4-variable and 5-variable queries. We find that when the size of the result is of the same order as the size of the smallest relation, most queries exhibit low sensitivity to variations in selectivity as indicated by the Q-factor values. Even when the size of the result was

increased by multiplying all selectivities by a constant factor, the Q-factor values did not deteriorate.

These results suggest a simple scheme for optimizing multi-join queries. One should simply choose a "middle of the road" value for each join selectivity. This crude estimate is likely to yield good performance. A more sophisticated scheme would be to vary the selectivity about an approximate value. This will allow the determination of the best plan,  $p^*$  and the construction of the corresponding Q-factor. If the Q-factor is low, say less than 1.20,  $p^*$  should be chosen for processing the query. Our experiments show that this will be true for a very large number queries that arise in practical situations. A more accurate determination of selectivity using detailed statistics is warranted only for those queries for which no insensitive plans can be found.

## 5. CONCLUSIONS

This paper has presented a study of the sensitivity of query plans to errors in the estimates for join selectivities. Since the cost of making exact determination of selectivities is prohibitive, estimates of selectivity are at best approximate. Current query planners take these estimates as exact and determine the optimal processing plan based on them. To study sensitivity of query plans to selectivity errors, we introduced a Q-factor for measuring sensitivity of queries to variations in selectivity. The value of this metric was measured for 20 different queries under different parameter values. Surprisingly low Q values were observed, and we can conclude that data managers should make crude selectivity estimates and not spend a great deal of time attempting to provide accurate selectivity values.

## REFERENCES

- [CHRI83] Christodoulakis, S., "Estimating Block Transfers and Join Sizes," Proc. 1983 ACM-SIGMOD Conference on Management of Data, San Jose, CA. May 1983
- [DATE84] Date, C.J., A Guide to DB2, Addison-Wesley Publishing Co., 1986.
- [EPST80] Epstein, R. and Stonebraker, M., "Analysis of Distributed Database Processing Strategies," Proc. VLDB Conference, Montreal, Canada. Sept. 1980.
- [HELD75] Held, G., et. al., "INGRES: A Relational Database System," Proc. 1975 National Computer Conference, Anaheim, CA. June 1975.
- [KOOI82] Kooi, R. and Frankfurth, D., "Query Optimization in INGRES," IEEE Database Engineering, September 1982.
- [PIAT84] Piatetsky-Shapiro, G. and Connell, C., "Accurate Estimation of the Number of Tuples Satisfying a Condition," Proc. 1984 ACM-SIGMOD Conference on Management of Data, Boston, Mass. June 1984.
- [ROWE83] Rowe, N., "Top Down Statistical Estimation on a Database," Proc. 1983 ACM-SIGMOD Conference on Management of Data, San Jose, CA. May 1983.
- [SELI79] Selinger, P. et. al., "Access Path Selection in a Relational Database Management System," Proc. 1979 ACM-SIGMOD Conference on Management of Data, Boston, Mass., June 1979.