

Copyright © 1975, by the author(s).  
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

EFFICIENT EVALUATION OF EXPRESSIONS IN A RELATIONAL ALGEBRA

by

Robert M. Pecherer

Memorandum No. ERL-M510

19 February 1975

ELECTRONICS RESEARCH LABORATORY

College of Engineering  
University of California, Berkeley  
94720

# EFFICIENT EVALUATION OF EXPRESSIONS IN A RELATIONAL ALGEBRA

Robert M. Pecherer

Department of Electrical Engineering and Computer Sciences  
and the Electronics Research Laboratory  
University of California, Berkeley, California 94720  
(415) 454-5427

The retrieval of information from a relational data base is treated as the evaluation of relation-defining expressions in an algebra over the data base relations. An algebra is introduced that offers computational advantages over previously defined relational algebras. A very general storage model is assumed which admits a wide variety of storage structures. An algorithm for  $O(n \log_2 n)$  division is described, and a rapid evaluation schema for a large and important class of expressions is derived. Finally, techniques are explored for translating an algebraic expression into an equivalent expression that is inherently "quicker" to evaluate.

## I. INTRODUCTION

In a relational data base system, the user logically views large quantities of formatted data as stored in time-varying, finite relations of assorted degrees. Updates to and retrievals from his collection of data are often specified in a descriptive, non-procedural language such as DSL-ALPHA [1], SEQUEL [2], SQUARE [3], DAMAS [4], or QUEL [5]. The physical representation of user relations and procedural solutions to non-procedural user requests are the main problem of the system designer, and as yet, efficient implementation techniques are largely unexplored. The presence of multiple users, multiple views and security and integrity constraints complicate the problem and are not discussed. Here, we consider only the formation of response relations (retrievals), so that the time-varying nature of the data base is ignored.

At present, there are three main approaches to the decomposition of non-procedural retrievals:

- (1) Iterative decomposition by tuple substitution [5]. In the retrieval statement, the tuples of one relation are substituted to produce one or more simpler statements with data values replacing variables. By repeated application, all variables are removed and the statements can be directly processed. This approach is further discussed in [4, 5, 6].
- (2) Translation to an algebraic expression employing relational operators on the data base relations. By implementing the operators, the expression is evaluated to produce the response relation. This suggestion is pursued in [7, 8].
- (3) Compilation to a set of procedures in a hierarchical or network-oriented language. This approach is followed in DIAM [9].

The first method has been implemented in the INGRES [5, 6] project at Berkeley. The third approach has been proposed by Senko [9] and by Tschritzis [10]. The feasibility of the second approach has been demonstrated with a translator from the non-procedural language ALPHA to a relational algebra by Codd [7], and an APL implementation of this algebra by Palermo [8]. This paper investigates efficiency considerations for the second approach, building on Palermo [8]. Results are applicable to algebraic expressions generated by the Codd Reduction Algorithm (CRA) [7] and to algebraic expressions in general. The reader is assumed to be familiar with [7], especially the CRA and the retrieval language ALPHA.

We are concerned here with 2 factors:

- (1) efficient implementation procedures for the operators of a relational algebra, and
- (11) translation of algebraic expressions to equivalent algebraic expressions which are inherently "quicker" to evaluate.

The paper is divided into 4 sections. The first introduces definitions and terminology, and describes the relational algebra whose implementation we study here. The second discusses techniques for the efficient implementation of this algebra; results are presented for arbitrary algebraic expressions and for the class of expressions produced by the CRA. The third describes a number of efficiency results obtainable by translation to equivalent expressions, and the last section offers a summary and conclusion.

## II. DEFINITIONS AND TERMINOLOGY FOR A RELATIONAL ALGEBRA

The definitions and terminology are similar to Codd [1, 7, 11, 12]. The relational algebra defined here is derived from [7], but employs fewer, more powerful operators; the differences will be noted.

### Relations

A domain  $D$  is set; a domain is simple if its members are not themselves sets. Let  $D_1, \dots, D_n$  be domains.

A relation  $R$  over  $D_1, \dots, D_n$  is a subset of the Cartesian Product  $D_1 \times \dots \times D_n$  and is said to be an n-ary relation or relation of degree n. The members of  $R$  are called n-tuples or simply tuples. If  $r$  is an  $n$ -tuple,  $r[i]$  designates the  $i^{\text{th}}$  domain value or  $i^{\text{th}}$  attribute of  $r$ , for  $1 \leq i \leq n$ . If  $A = a_1, \dots, a_k$

is a list of integers such that  $1 \leq a_j \leq n$  for  $j = 1, \dots, k$ , then  $A$  is a domain identifying list for any  $n$ -ary relation, and  $r[A]$  designates the  $k$ -tuple  $(r[a_1], \dots, r[a_k])$ . If  $r$  is an  $n$ -tuple and  $s$  an  $m$ -tuple,  $\hat{r}$  designates the  $n+m$ -tuple  $(r[1], \dots, r[n], s[1], \dots, s[m])$ . A relation is first-normal or normal if each of its domains is simple. All relations considered here are assumed to be normal. In the sequel,  $\text{deg}(R)$  denotes the degree of relation  $R$ ; the size of a relation is the number of tuples it contains.

### A Relational Algebra

Any relational algebra consists of operators which map relations or pairs of relations to relations, and a set of relations closed under the operators. The following definitions apply only to normal relations. Note that tuple attributes are identified by position.

Definition (2.1) Let L be a domain identifying list for relation R. The projection  $\pi_L(R)$  is defined by:

$$\pi_L(R) = \{r[L]: r \in R\}$$

Definition (2.2) Let R and S be relations of degree n and m respectively. The product  $R * S$  is a set of  $n+m$ -tuples defined by:

$$R * S = \{r s: r \in R \wedge s \in S\}$$

Definition (2.3) Let R and S be relations, and r and s tuple variables for R and S respectively. Let f be a 0-1 function specified as a Boolean combination of terms of the form

$$x \theta y$$

where  $\theta$  is one of  $\{=, \neq, <, \leq, >, \geq\}$ , and each x and y is an arithmetic expression involving tuple attributes  $r[i]$  and  $s[j]$ , simple functions (such as log, sin, etc.) of attributes, and constants. The join  $R[f]S$  is the set of tuples in  $R * S$  for which the function f is equal to "1." e.g.,  $R[f]S = \{r s: r \in R \wedge s \in S \wedge f(r, s) = 1\}$ .

In (2.3), the symbols "r" and "s" always refer to the left and right operands respectively in the binary join; this is simply a syntactic convenience.

Definition (2.4) Let R be a relation, r a tuple variable for R, and g a Boolean function such as "f" in join, but referencing only attributes of r. The restriction  $R[g]$  is the set of tuples in R for which g is equal to "1."  
 $R[g] = \{r: r \in R \wedge g(r) = 1\}$ .

Definition (2.5) Let R and S be relations of degree n and m respectively; let A and B be domain identifying lists for R and S respectively, both of length  $k < n$  and both without repetition. Let  $\bar{A}$  be a domain identifying list for R complementary to A and in increasing order (viz.  $n=5, A=1,4,3$  then  $\bar{A}=2,5$ ). The division  $R[A \div B]$  designates a subset of  $n-k$ -tuples in  $\pi_{\bar{A}}(R)$  as follows:

$$R[A \div B] = \{r[\bar{A}]: r \in R \wedge \forall s \in S \exists r' \in R \{r[\bar{A}] = r'[\bar{A}] \wedge r'[A] = s[B]\}$$

A more revealing definition and numerous examples of division are given in [7]; a minor difference is that in our definition, when S is vacuous,  $R[A \div B] = \pi_{\bar{A}}(R)$ , whereas in Codd's,  $R[A \div B]$  is the empty relation. A third possibility is to make division by a vacuous relation undefined. Our definition is more consistent with the intended correspondence between division in the algebra and testing for universal quantification in a first-order predicate calculus (see [7, 8] for a discussion). For our purposes, algebraic expressions are assumed to be over non-empty relations, so that the difference in definition is of no concern.

Comparing the above definitions to Codd's Relational Algebra [7], projection and product are identical, as is division when the divisor relation is non-empty. For join and restriction, Codd's notation is such that the only Boolean functions allowed are conjunctions of  $\theta$ -comparisons of attribute values, with all  $\theta$ 's the same. Our definitions are motivated by the belief that the relations of an implemented system will reside in a slow secondary memory, and the input of tuples from R and S to form  $R[f]S$  (or from just R to form  $R[g]$ ) is expected to be more time-consuming than application of the function f (or g). So that the retrieval of (say)

$$\{r s: r \in R \wedge s \in S \wedge (r[1]=s[1] \vee r[2] \neq s[2])\}$$

may be expected to proceed more quickly by evaluating

$$R[r[1] = s[1] \vee r[2] \neq s[2]]S$$

than by evaluating

$$R[r[1] = s[1]]S \text{ and } R[r[2] \neq s[2]]S$$

and forming their union. A similar argument applies for restriction.

Codd's algebra includes the set operations of union, intersection and relative complement; the following identities indicate that the latter 2 are not necessary:

$$(2.6) R \cap S = \pi_{1, \dots, n}(R[r[1] = s[1]] \wedge \dots \wedge r[n] = s[n])S$$

$$(2.7) R \setminus S = (R[r[1] \neq s[1]] \vee \dots \vee r[n] \neq s[n])S$$

where  $\deg(R) = \deg(S) = n$  and S is non-empty

Union cannot be obtained from the other operators since it produces relations with domains which are supersets of the operand relation domains, and none of the other operators can achieve this. Allowing disjunctions in the definitions of the Boolean functions "f" and "g" (for join and restriction) eliminates the primary need for union. If the union of two compatible relations must be formed, we presume this to take place external to the algebra; implementation of union will not be discussed.

The five operators defined here are not a minimal set since

$$(2.8) R * S = R[r[1] = s[1] \vee r[1] \neq s[1]]S, \text{ and}$$

$$(2.9) R[g] = \pi_{1, \dots, n}(R[r[1] = s[1]] \wedge \dots \wedge r[n] = s[n] \wedge g)R$$

where  $\deg(R) = n$ .

No computational advantage is indicated by eliminating product and restriction operators, and in fact, implementation of restriction of a relation by a projection of a join of the relation with itself is probably a poor idea.

In the next section, a storage framework for relations and time approximations for evaluation of the operators are introduced. A fast way to perform division is derived, and a very fast evaluation technique for an important class of algebraic expressions is demonstrated.

### III. EFFICIENCY IN THE IMPLEMENTATION OF RELATIONAL OPERATORS

In this section, we proceed as if our only goal is the evaluation of syntactically correct expressions in the relational algebra over a fixed set of stored data base relations  $R_1, \dots, R_p$  of size  $n_1, \dots, n_p$ .

Correct evaluation of each operator requires examination of every tuple of each operand relation, however, since a relation is a set with no specified order properties, the order in which tuples are examined is of no logical consequence to the result. For the evaluation of certain operators, the retrieval order of tuples in the operand relation(s) can affect the time required. To demonstrate this, we make the following assumptions:

(3.1) All relations (data base, intermediate, result) are maintained in a slow secondary memory as tables of tuples; the order in which they are stored is the only order in which they can be retrieved.

(3.2) All tuples of a relation are encoded in a fixed-length field. The time required to retrieve every tuple of a single relation is proportional to its size.

(3.3) Application of Boolean functions to tuples is performed in primary memory by a single processor in an amount of time which is insignificant compared to the input time for the operand relations or the output time for the result.

We are concerned with the relative efficiency of evaluating the operators for different retrieval orders. We approximate the amount of time required as a function of the size of the operands. When the time is specified as "O(f(n))," this indicates that the actual time T(n) is proportional to f(n), or equivalently,

$$(3.4) \lim_{n \rightarrow \infty} T(n)/f(n) = K$$

where K is a proportionality constant (nonzero).

The evaluation of  $R_1[g]$  requires time  $O(n_1)$  since every tuple of  $R_1$  must be retrieved for testing with g. The evaluation of  $R_1 * R_j$  requires time  $O(n_1 * n_j)$  since  $n_1 * n_j$  tuples must be output. We assume that evaluation of  $R_1[f]R_j$  also requires time  $O(n_1 * n_j)$  since every tuple of  $R_1 * R_j$  must be formed and tested by f. These values are independent of the retrieval order of the operand relations except for certain simple Boolean functions which do not concern us here. We are concerned with projection and division, for which different (known) storage orders require different evaluation times.

The evaluation of  $\pi_L(R_1)$  (when L is not a permutation of the sequence "1,...,deg( $R_1$ )") is complicated by the fact that the result is a set without duplicate tuples. If it is known that  $R_1$  is sorted on the domains of L, all tuples r in  $R_1$  with the same value of r[L] appear sequentially when  $R_1$  is retrieved.

The storage of duplicate tuples can be avoided with comparisons between consecutively retrieved tuples (not by scanning all previously stored tuples), so in this case,  $\pi_L(R_1)$  can be evaluated in time  $O(n_1)$ .

The best method we have found for evaluating  $\pi_L(R_1)$  when  $R_1$  is unsorted is by a simultaneous sort-and-project procedure that eliminates domains not in L and duplicate tuple values. This method is dominated by the sorting time which we take to be  $O(n_1 \log_2 n_1)$

(see Knuth [13], pp. 361-376). Under the given assumptions, we have that:

(3.5)  $\pi_L(R_1)$  can be evaluated in time  $O(n_1)$  if each subset of  $R_1$  with the same projection r[L] is consecutively retrievable.

The amount of time to evaluate  $R_1[A \div B]R_j$  depends not only on  $n_1$  and  $n_j$ , but also on the size of the result; In the following derivation, we assume  $n_1 \geq n_j$ , and that s[B] assumes  $n_j$  distinct values as s ranges over  $R_j$ .

From the definition, we have that for each r in  $R_1$ :

$$(3.6) r[\bar{A}] \in R_1[A \div B]R_j \Leftrightarrow \forall s \in R_j \exists r' \in R_1 [r[\bar{A}] = r'[\bar{A}] \wedge r'[A] = s[B]]$$

$$(3.7) r[\bar{A}] \notin R_1[A \div B]R_j \Leftrightarrow \exists s \in R_j \forall r' \in R_1 [r[\bar{A}] \neq r'[\bar{A}] \vee r'[A] \neq s[B]]$$

To verify the condition on the right hand side (rhs) of (3.6) requires time  $n_j * O(n_1)$  since there are  $n_j$  tuples in  $R_j$ , and it requires time  $O(n_1)$  to find r' in  $R_1$ . To verify the condition on the rhs of (3.7) requires time  $O(n_1)$  since every tuple of  $R_1$  must be checked. If m is the size of the result, m is bounded from below by 0 and above by  $n_1/n_j$  since for each t in the result, there must be  $n_j$  tuples  $z_1, \dots, z_{n_j}$  in  $R_1$  with  $z_1[\bar{A}] = \dots = z_{n_j}[\bar{A}] = t$ , and these tuples cannot satisfy the necessary condition for any other t' in  $R_1[A \div B]R_j$ . This yields the time approximation:

$$(3.9) m * n_j * O(n_1) + (n_1 - m) * O(n_1)$$

for the total evaluation, which is  $O(n_1^2)$  over the range of m. Clever programming can produce smaller values for K in (3.4) for division, but the  $O(n_1^2)$  approximation is valid.

When it is known that  $R_1$  is sorted on  $\bar{A}$ , A and  $R_j$  on B, then for each t in  $R_1[A \div B]R_j$ , (at least)  $n_j$  consecutively retrievable tuples in  $R_1$  will have  $z[\bar{A}] = t$ , and within this subset, the  $n_j$  tuples that satisfy the role of r' in the rhs of (3.6) will appear in the same order as the  $n_j$  tuples of  $R_j$ . The following example illustrates this:

(3.9) Example:  $R_1[2, 3 \div 1, 2]R_j$   
(unsorted  $R_1, R_j$ )

$R_1$	<u>1</u>	<u>2</u>	<u>3</u>	$R_j$	<u>1</u>	<u>2</u>
A	X	1		Z	3	
C	Z	3		X	1	
B	Z	3				
C	X	1				
A	Z	1				
A	Z	3				
B	X	2				
B	Z	1				

( $R_1$  sorted on 1,2,3;  $R_j$  sorted on 1,2)

$R_1$	<u>1</u>	<u>2</u>	<u>3</u>	$R_j$	<u>1</u>	<u>2</u>
A	X	1		X	1	
A	Z	1		Z	3	
A	Z	3				
B	X	2				
B	Z	1				
B	Z	3				
C	X	1				
C	Z	3				

Having achieved this arrangement of the operand relations, evaluation of  $R_1[A \div B]R_j$  can be performed with a sequential scan of  $R_1$  since no tuple of  $R_1$  has to be compared to more than one tuple of  $R_j$ . The cost of sorting both  $R_1$  and  $R_j$  is  $O(n_1 \log_2 n_1)$ .

since  $n_1 \geq n_j$ , and since the division can then be performed in  $O(n_1)$ , the total time is  $O(n_1 \log_2 n_1)$ , which is superior to  $O(n_1^2)$  for unordered  $R_1$  and  $R_j$ . For an arbitrary algebraic expression, this technique and the sort-project technique for projection guarantee that any projection or division can be evaluated in time  $O(n \log_2 n)$  where  $n$  is the size of the projected or divided relation. And under the given assumptions, we also have that:

(3.10)  $R_1[A \div B]R_j$  can be evaluated in  $O(n_1)$  if each subset of  $R_1$  with the same projection  $r[\bar{A}]$  is consecutively retrievable, and the tuples of each subset are retrievable in the same order as the tuples of  $R_j$ .

For an important class of algebraic expressions involving projections and divisions, the conditions in (3.5) and (3.10) can be achieved in the intermediate results of evaluation at no extra cost. The Codd Reduction Algorithm (CRA) generates the class from a restricted set of ALPHA expressions (see [7] for details of the CRA and ALPHA); the general form is:

(3.11)  $\pi_L(G_{k+1}(\dots G_{k+q}((R_1 * \dots * R_{k+q})[g])\dots))$

in which  $G_{k+j}(X)$  is either the projection

$$\pi_{1, \dots, m_{k+j}}(X)$$

or the division

$$X[m_{k+j}+1, \dots, m_{k+j+1} \div 1, \dots, \text{deg}(R_{k+j})]R_{k+j}$$

where  $m_1 = \text{deg}(R_1) + \dots + \text{deg}(R_{i-1})$ .

We refer to such expressions as "CRA-expressions." Suppose the evaluation is performed by producing  $T_0, \dots, T_q, W$  where:

$$T_0 = (R_1 * \dots * R_k * \dots * R_{k+q})[g]$$

$$T_1 = G_{k+q}(T_0)$$

$$\vdots$$

$$T_q = G_{k+1}(T_{q-1})$$

$$W := \pi_L(T_q)$$

When  $g$  is identically true,  $T_i = R_1 * \dots * R_{k+q-i}$  for  $i = 0, \dots, q$ ; each  $G_{k+j}$  produces  $T_{q-j+1} = R_1 * \dots * R_{k+j-1}$  from  $T_{q-j} = R_1 * \dots * R_{k+j}$ . If  $T_0$  is generated as  $(R_1 * \dots * R_k) * (R_{k+1} * (R_{k+2} * (\dots (R_{k+q}))))$ , then for every  $t \in T_1 = R_1 * \dots * R_{k+q-1}$ , the subset  $t * R_{k+q}$  appears sequentially in  $T_0$ , so the conditions in (3.5) are met; if  $R_{k+q}$  is unaltered, the conditions in (3.10) are also met.  $G_{k+q}$  can then be applied to  $T_0$  in time proportional to its size. If  $R_{k+q-1}$  is unaltered and the tuples of  $T_1$  are stored in the order that the tuples of  $T_0$  are sequentially scanned,  $T_1$  meets the conditions also, so again,  $G_{k+q-1}$  can be applied in time proportional to the size of  $T_1$ . Repeating the procedure for  $G_{k+q-1}, \dots, G_{k+1}$ , each such projection or division can proceed in time proportional to the size of the operand. When  $g$  is

any Boolean function, the same argument applies to subsets of  $R_1 * \dots * R_{k+q}, \dots, R_1 * \dots * R_{k+1}$ , so that again, the projections and divisions proceed quickly. Only for the last projection " $W = \pi_L(T_q)$ " can the condition in (3.5) fail.

The importance of this technique is that  $O(n)$  procedures for every division and projection (except the last) are available without any rearrangement of the data base relations. The CRA-expressions can represent complex queries over many relations; rapid evaluation is critical to the performance of the entire system.

#### IV. EFFICIENCY BY TRANSLATION TO AN EQUIVALENT EXPRESSION

Complex queries to an information retrieval system are so time-consuming that even if they represent a small fraction of the total queries, their effect could seriously degrade the response time for simpler requests. Examination of the general CRA-expression

(4.1)  $\pi_L(G_{k+1}(\dots G_{k+q}((R_1 * \dots * R_{k+q})[g])\dots))$

indicates 2 reasons why complex queries require so much time:

- (i) The product relation  $R_1 * \dots * R_{k+q}$  is exceedingly large, perhaps too large to be stored in secondary memory.
- (ii) The restriction " $g$ " and each projection and division require at least a sequential scan of every tuple left by the application of the previously applied operators.

The previous section discussed operator implementation; in this section, properties of the operators are exploited to produce equivalent expressions which reduce the size of the product space and eliminate certain projections and divisions. The translation of an expression to a more easily evaluated expression is expected to proceed mechanically, but guided by accumulated statistics on the data base relations and previous evaluations, and possibly with user interaction. This work was stimulated by Palermo [8].

Palermo [8] has shown that in (4.1),

- (4.2) For  $1 \leq i \leq k$ , any domain of  $R_i$  not referenced in " $L$ " or " $g$ " can be projected out prior to evaluation without affecting the result.
- (4.3) For  $k < i \leq k+q$ , any domain of  $R_i$  not referenced in " $g$ " can similarly be projected out provided  $G_i$  is a projection and not a division.

This allows us to rewrite (4.1) as

(4.4)  $\pi_L(G'_{k+1}(\dots G'_{k+q}((\pi_{L_1}(R_1) * \dots * \pi_{L_{k+q}}(R_{k+q})) [g']\dots))$

where each projection  $\pi_{L_i}(R_i)$  eliminates domains in  $R_i$  in accordance with (4.2) and (4.3), and  $L'$ ,  $G'_{k+j}$  and  $g'$  are derived from  $L$ ,  $G_{k+j}$  and  $g$  in (4.1) to reflect the altered positions of the relevant domains. The new expression (4.4) requires less space to evaluate, but the time to perform projections indicates that this transformation should be used cautiously.

A more immediate result is obtained from the following identities in which  $\text{deg}(R_1) = n$ ,  $\text{deg}(R_2) = m$  and  $\text{deg}(R_3) = q$ :

$$(4.5) \quad \{r_1: r_1 \in R_1 \wedge \exists r_2 \in R_2 \exists r_3 \in R_3 [g]\}$$

$$= \pi_{1, \dots, n} (\pi_{1, \dots, n+m} ((R_1 * R_2 * R_3)[g]))$$

$$= \pi_{1, \dots, n} ((R_1 * R_2 * R_3)[g]).$$

$$(4.6) \quad \{r_1: r_1 \in R_1 \wedge \forall r_2 \in R_2 \forall r_3 \in R_3 [g]\}$$

$$= ((R_1 * R_2 * R_3)[g])[n+m+1, \dots, n+m+q$$

$$+ 1, \dots, q] R_3 [n+1, \dots, n+m + 1, \dots, m] R_2$$

$$= ((R_1 * R_2 * R_3)[g])[n+1, \dots, n+m+q$$

$$+ 1, \dots, m+q] (R_2 * R_3).$$

For CRA-expressions involving sequences of 2 or more consecutive projections (divisions), each such sequence can be replaced with a single projection (division). Although the divisor relation is larger, the size of the dividend relation is the controlling factor; this translation should be performed whenever possible.

The remainder of this section is concerned with replacing products with joins and relations with restricted relations. The results are directed towards limiting the growth of the restricted product which appears in every CRA-expression:

$$(4.7) \quad R_1 * \dots * R_{k+q} [g]$$

However, they also apply to arbitrary algebraic expressions.

Consider the following expression in which  $\text{deg}(R_i) = 3$  for  $i = 1, 2, 3$ :

$$(4.8) \quad (R_1 * R_2 * R_3)[r[1] = 5 \wedge r[3] = r[4] \wedge r[6] \leq r[7]$$

$$\wedge r[5] \neq r[6] \wedge r[3] + r[6] = r[9]].$$

The terms " $r[1] = 5$ " and " $r[5] \neq r[6]$ " each involve attributes from a single operand relation, hence we can rewrite (4.8) as

$$(4.9) \quad ((R_1[r[1] = 5]) * (R_2[r[2] \neq r[3]]) * R_3)$$

$$[r[3] = r[4] \wedge r[6] \leq r[7]$$

$$\wedge r[3] + r[6] = r[9]].$$

In (4.7), the terms " $r[3] = r[4]$ " and " $r[6] \leq r[7]$ " can be eliminated by replacing the product operator with joins. Notice that to avoid ambiguity, extra parenthesization is required:

$$(4.10) \quad (((R_1[r[1] = 5])[r[3] = s[1]]$$

$$(R_2[r[2] \neq r[3]])) [r[6] \leq s[1]] R_3)$$

$$[r[3] + r[6] = r[9]].$$

The final term " $r[3] + r[6] = r[9]$ " can be incorporated into the second join operator. To illustrate the magnitude of time saved, suppose that equality joins and restrictions reduce the size of the result by 90%, inequalities by 10% and all others by 50%. If  $n_1$  is the size of  $R_1$ , the following table gives expected values for the number of times a Boolean function is applied to a tuple or tuple pair in the above expressions:

Expression	Expected # of tuple tests
(4.8)	$n_1 * n_2 * n_3$ .
(4.9)	$n_1 + n_2 + .09(n_1 * n_2 * n_3)$
(4.10)	$n_1 + n_2 + .09(n_1 * n_2) + .009(n_1 * n_2 * n_3)$

If  $n_1 = n_2 = n_3 = 100$ , (4.9) offers a 90% reduction; for (4.10) it is almost 99%.

A further set of improvements is obtainable with rearrangements of the operands. For example, under the assumptions made, the expression

$$(4.11) \quad (R_1 * R_2 * R_3)[r[1] \neq r[4] \wedge r[6] \neq r[7]$$

$$\wedge r[2] = r[8]].$$

should be evaluated as

$$(4.12) \quad \pi_{1,2,3,7,8,9,4,5,6} ((R_1[r[2] = s[2]] R_3)$$

$$[r[1] \neq s[1] \wedge s[3] \neq r[4]] R_2)$$

rather than

$$(4.13) \quad (R_1[r[1] \neq s[1]] R_2)[r[6] \neq s[1]$$

$$\wedge r[2] = s[2]] R_3$$

since the equality join " $r[2] = s[2]$ " in (4.12) is presumed to be more restrictive than the inequality join " $r[1] \neq s[1]$ " in (4.13), and the projection is merely a permutation.

An algorithm has been developed which converts an arbitrary restricted product to an equivalent expression using products, joins and restrictions that is optimal with respect to expected amount of work. Details will appear in a forthcoming report.

## V. CONCLUSION

This paper has addressed the problem of evaluating expressions in a relational algebra over relations in a data base. Procedural considerations for the operators produced a more rapid evaluation schema for division than might otherwise be achieved. For an important class of expressions (the CRA-expressions), it was shown that  $O(n)$  procedures for projections and divisions were achievable when the intermediate results were maintained in certain orders in their actual stored representations. Finally, techniques were demonstrated for translating certain commonly occurring algebraic expressions to equivalent expressions that are in some sense "easier" to evaluate.

When the retrieval of response relations in a non-procedural, descriptive language is performed by translating the retrieval statements into a relational algebra for evaluation, properties of the algebra can be exploited to produce the response relations more rapidly than might be expected. It is hoped that this paper sheds some light on this approach.

## ACKNOWLEDGEMENT

Research sponsored by Naval Electronic Systems Command Contract N00039-75-C-0034.

## BIBLIOGRAPHY

- [1] Codd, E. F., "A Data Base Sublanguage Founded on the Relational Calculus," Proc. 1971 ACM SIGFIDET Workshop on Data Description, Access and Control, San Diego, Nov. 1971.

- [2] Chamberlin, D. D. and R. F. Boyce, "SEQUEL: A Structural English Query Language," IBM Research Laboratory, San Jose, CA
- [3] Boyce, R. F., D. D. Chamberlin, M. M. Hammer and W. F. King, "Specifying Queries as Relational Expressions: SQUARE," IBM Technical Report RJ 1291, October 1973.
- [4] Rothnie, J. B., "The Design of Generalized Data Management Systems," Ph.D. Dissertation, Dept. of Civil Engineering, MIT, September 1972.
- [5] McDonald, N., M. Stonebraker and E. Wong, "Preliminary Design of INGRES," ERL, Univ. Calif., Berkeley, Memo #ERL-M435, April 1974.
- [6] Stonebraker, M. and E. Wong, "INGRES-A Relational Data Base System," ERL, Univ. of Calif., Berkeley, Memo #ERL-M472, November 1974.
- [7] Codd, E. F., "Relational Completeness of Data Base Sublanguages," Courant Computer Science Symposium 6, May 1972.
- [8] Palermo, F. P., "A Data Base Search Problem," IBM Technical Report RJ 1072, July 1972.
- [9] Senko, M. E., E. B. Altman, M. M. Astrahan and P. L. Fehder, "Data Structures and Accessing in Data-Base Systems," IBM Systems Journal, Vol. 12, No. 1, 1973.
- [10] Tsichritzis, D., "A Network Framework for Relation Implementation," Proc. of the IFIPS TC-2 Working Conference on Data Definition Language, January 1975.
- [11] Codd, E. F., "A Relational Model of Data for Large Shared Data Banks," CACM, Vol: 13, No. 6, June 1970.
- [12] Codd, E. F., "Seven Steps to Rendezvous with the Casual User," IBM Research, San Jose, CA, RJ 1333, January 1974.
- [13] Knuth, D. E., The Art of Computer Programming, Volume 3, Sorting and Searching, Addison-Wesley, 1973.