DFT TECHNIQUES FOR SIZE ESTIMATION OF DATABASE JOIN OPERATIONS

KAMİL SARAÇ*, ÖMER EĞECİOĞLU, AMR EL ABBADI

Computer Science Department, University of California Santa Barbara,
Santa Barbara, CA 93106, USA

ABSTRACT

Novel algorithms based on the Discrete Fourier Transform (DFT) are proposed to estimate the size of relations resulting from join operations. We start with an approach in which the frequency distribution values are transformed using the DFT and the Fourier coefficients are used to construct histograms. Our main contribution is a direct approach which uses the amplitudes of the DFT coefficients iteratively. The proposed algorithm gives the exact join size using logarithmic space for the special case of self join. A generalization to compute the join of arbitrary relations is then used to develop two tree-based techniques that provide a spectrum of algorithms which interpolate storage requirements versus accuracy of the estimation obtained. Finally, we present experimental results to exhibit the effectiveness of our approach.

Keywords: Database join, histogram, query optimization, Discrete Fourier Transform.

1. Introduction

Query optimization is an important issue in Database Management Systems (DBMSs). There may be several ways for executing a given query. Finding all these different access plans and deciding the best one may not be practical for a user. As a result, DBMSs need to perform this complex optimization process by themselves. For this reason, most DBMSs use a component called the query optimizer. The aim of query optimizers in relational database systems is to select the most efficient way among all possible ways of executing a query. In practice this process requires approximating the cost of possible execution sequences and selecting the cheapest one in terms of a cost metric, which is usually the cumulative size of intermediate operations. Even though these possible ways result in the same answer, their execution sequence may be different, and hence they may result in different execution costs. A necessary procedure in estimating the execution cost of a query is to estimate the result size of each operation in each possible execution sequence. This information depends on the frequency distribution of the underlying data. In practice, most database systems store summarized information about the frequency distribution (or find it dynamically when needed) and use this information

---

*Computer Science Department, University of California Santa Barbara, Santa Barbara, CA 93106, USA

1

to estimate the result size of various operations and then the cost of the executions. The accuracy of this approximation is crucial to the performance of the database systems. A small error introduced in the execution of a portion of the query can grow exponentially and affect the overall performance considerably [7].

In this paper, we propose novel techniques based on the *Discrete Fourier Transform (DFT)* to estimate the size of relations resulting from join operations. Traditionally, the frequency distributions of the underlying data is stored in the form of histograms. As a first approach, we also use histograms for approximation, but instead of constructing the histogram directly, we first apply DFT, and then construct the histogram based on the amplitudes of the Fourier coefficients. If we consider the frequency vector of an attribute as a time domain signal then the DFT coefficients represent the frequency characteristics of this signal. We base our histogram construction approach on the *end-biased* approach [8], where the most frequent values are stored in the histogram as the first few values, and the rest are averaged. For many time domain signal distributions, the first few DFT coefficients dominate the representation of the signal. Thus by storing the first few DFT coefficients precisely and by averaging the rest of the coefficients in the histogram, the DFT based approach promises to provide a good approximation of the original data frequency distribution.

The second approach is the iterated application of DFT to the frequency distribution values modulo the phase information which allows a logarithmic size representation of a given vector. The algorithms based on this method are our main contribution. The first version of this method called *Self Join using Absolute Value (SJAV)* is exact for self join operations. For arbitrary join operations it provides an upper bound on the size of the join, we refer to this as *Approximation by Absolute Value (AAV)*. The generalizations of *AAV*, *Tree Approximation Algorithms (TAA)* and *Tree Approximation Algorithms with Truncation (TAAT)* are built upon a binary tree representation of the vectors obtained through the iterated application of the DFT. *TAA* uses *AAV* as a subprocedure at the lower levels of the tree, whereas *TAAT* is obtained by truncating the tree at an appropriate level. Both *TAA* and *TAAT* provide a spectrum of algorithms that interpolate storage requirements versus accuracy of the estimates obtained.

This paper is organized as follows. In Section 2 we summarize the previous work on the problem. In section 3 we give the problem statement. In Section 4 we present fundamental properties of the DFT that are used in our algorithms. The algorithms are developed in Sections 5, 6 and 7. We evaluate the performance of the algorithms in Section 8 and conclude the paper in Section 9.

2. Related Work

Several techniques (histogram based, parametric and sampling based) have been proposed in the literature for estimating query result sizes. The survey by Mannino, Chu and Sager is a good reference on the early work on database size estimation [11]. These techniques present various trade-offs in terms of storage, precision, run-time overhead, and make varying assumptions about the distribution of the underlying

data.

2.1. Histogram Techniques

Histogram based techniques [12,6] depend on the underlying data distribution. They do not make any assumptions about the characteristics of the data distribution. They scan the database to extract summary information about the data distribution and store this information into a histogram. Since the nature of this technique is static, as the underlying data changes, the histograms should be reconstructed to reflect these changes. Histograms are the most commonly used techniques in practice [6] and a significant amount of work is done on histogram based estimation techniques. Ioannidis and Poosala made use of types of histograms called *serial* and *end-biased* histograms [8]. Later on they identified several key properties that characterize histograms and provided a taxonomy of histograms based on these properties [16]. Most recent work on histogram based techniques are on histogram maintenance [4] and multi-dimensional histogram construction [15].

2.2. Parametric Techniques

Parametric techniques use parameterized mathematical distributions, e.g. uniform, Zipf, to estimate query result sizes [18,2]. The assumption is that the underlying data follows the a priori selected distribution model. The main advantage of these methods is that they do not require much storage and they have a very small run time overhead. On the other hand, in case of large updates the precomputed distribution parameters may not reflect the distribution characteristics anymore and need to be re-computed. In addition, actual data may or may not follow the selected distribution model.

2.3. Sampling Techniques

Sampling based techniques compute their estimations at run time (query optimization time) by sampling and processing data from the database [5,10]. They do not store any summary information in the database. Thus they do not require any storage space for estimation and do not incur any preprocessing overhead for estimation. In addition, they reflect the most up to date information about the state of the database. The estimation error can be bound by selecting an appropriate sample size using the well-known theory of statistics. On the other hand, the estimation is done per query and no information is stored for future reference. Thus they incur the cost repeatedly.

3. Problem Formulation

We first describe the problem of join size estimation in databases. The development and motivation closely follows that of Ioannidis and Poosala [6,8]. Consider a database with relations $R_0, R_1, \ldots, R_n$. Let a_i, a_{i+1} be attributes of relation R_i.

A tree function-free equality-join query Q is defined as

$$Q = (R_0.a_1 = R_1.a_1 \ and \ R_1.a_2 = R_2.a_2 \ and \ldots and R_{n-1}.a_n = R_n.a_n).$$

The attributes $a_1, a_2, \ldots, a_n$ are called *join attributes*. Relations R_i, $2 \leq i \leq n-1$ participate in Q with two attributes a_i and a_{i+1} and relations R_0 and R_n with one attribute, a_1 and a_n respectively.

In order to optimize the execution of query Q, the size S of the resulting relation must be estimated. Frequency matrices T_i keep track of the joint-frequency distribution of attributes a_i and a_{i+1} of R_i. An exact expression for S is the matrix product

$$S = T_0 T_1 \ldots T_n, \tag{1}$$

but the storage of the frequency matrices and exact computation of this product is too costly.

We follow the common assumption of *attribute independence* formulated in [2] and adopted by most database systems. According to this assumption, the distribution of attribute values are independent of each other. Therefore, if v_i and h_{i+1} are the frequency distributions of individual attributes a_i and a_{i+1} in R_i, then the frequency matrix T_i can be written as the outer product

$$T_i = \frac{1}{|R_i|}(v_i^T \times h_{i+1})$$

where v_i^T is transpose of v_i and $|R_i|$ represents the number of tuples in R_i. This assumption and associativity enables the calculation of S from

$$(h_1 v_1^T)(h_2 v_2^T) \cdots (h_n v_n^T). \tag{2}$$

Note that $h_i v_i^T = \langle h_i, v_i \rangle$ is the standard inner product of the vectors $X = h_i$ and $Y = v_i$.

Example 1 *Let* $T_i = \begin{pmatrix} 40 & 24 & 12 \\ 30 & 18 & 9 \\ 20 & 12 & 6 \end{pmatrix}$ *be the frequency matrix of a relation* R_i. *Then* T_i *can be written as*

$$T_i = \frac{1}{171} \begin{pmatrix} 76, \\ 57, \\ 38 \end{pmatrix} \begin{pmatrix} 90, & 54, & 27 \end{pmatrix}$$

where $|R_i| = 171$, $v_i = (76, \ 57, \ 38)$ *and* $h_{i+1} = (90, \ 54, \ 27)$.

4. Preliminaries

The main motivation for our approach is the use of DFT based techniques to estimate each of the factors in (2), and therefore (1). DFT has been used in various estimation problems in database research such as dimension reduction for searching as well as indexing high dimensional data [1,17].

4

Consider an N-dimensional real vector $X = (x_1, x_2, \ldots, x_N)$. The *Discrete Fourier Transform (DFT)* of X is the N-dimensional complex vector $\hat{X} = (\hat{x}_1, \hat{x}_2, \ldots, \hat{x}_N)$ given by $\hat{X} = FX$ where $F = (1/\sqrt{N}) \| \omega^{(i-1)(j-1)} \|$ is the $N \times N$ Fourier matrix with $\omega = \cos(2\pi/N) + I \sin(2\pi/N)$, and $I = \sqrt{-1}$. It is well known that the DFT of X can be computed using the Fast Fourier Transform in $O(N \log N)$ arithmetic operations. The fundamental properties of the DFT that we make use of are summarized below [3,13]. If we start with an N-dimensional nonnegative real vector X with transform $\hat{X}$, then

- $\hat{x}_1 = (x_1 + x_2 + \cdots + x_N)/\sqrt{N}$ is a nonnegative real number.
- $\hat{x}_i$ and $\hat{x}_{N-i+2}$ are conjugate complex numbers for $i = 2, 3, \ldots, \lceil \frac{N}{2} \rceil$.

In our presentation, we assume that N is odd for simplicity, although our techniques extend to the case of arbitrary N. If we write $N = 2m + 1$, then the coefficient $\hat{x}_1$ is nonnegative real, and the lists $(\hat{x}_2, \hat{x}_3 \ldots, \hat{x}_m)$ and $(\hat{x}_N, \hat{x}_{N-1}, \ldots, \hat{x}_{m+1})$ pair up exactly in conjugate pairs. The most important property of the DFT we use is Parseval's identity [3]:

Theorem 1 *Suppose X and Y are two N-dimensional real vectors. Then $< X, Y > = < \hat{X}, \hat{Y} >$ where $< , >$ denotes the standard complex inner product.*

We use Parseval's identity in the following expanded form for real vectors X and Y. Suppose $\hat{X} = \alpha + I\beta$ is the decomposition of $\hat{X}$ into its real and imaginary components. Similarly, write $\hat{Y} = \gamma + I\delta$. Then Parseval's identity implies that

$$
\begin{aligned}
< X, Y > &= < \hat{X}, \hat{Y} > = \sum_{j=1}^{N} (\alpha_j + I\beta_j)\overline{(\gamma_j + I\delta_j)} \\
&= \sum_{j=1}^{N} (\alpha_j \gamma_j + \beta_j \delta_j) + I \sum_{j=1}^{N} (\beta_j \gamma_j - \alpha_j \delta_j) \\
&= < \alpha, \gamma > + < \beta, \delta > + I(< \beta, \gamma > - < \alpha, \delta >) \qquad (3)
\end{aligned}
$$

Since X and Y are real, the imaginary part on the right hand side of the above equation is zero and

$$
< X, Y > = < \alpha, \gamma > + < \beta, \delta > . \qquad (4)
$$

In other words,

$$
\sum_{i=1}^{N} x_i y_i = \sum_{i=1}^{N} \alpha_i \gamma_i + \sum_{i=1}^{N} \beta_i \delta_i. \qquad (5)
$$

5. DFT Based Histograms

As mentioned earlier, histograms are the most commonly used techniques in query result size estimation. In this section, we present DFT based histogram techniques for the query result size estimation problem. We start with a brief review of database histograms based on Poosala et al. [16].

5.1. Database Histograms

A *histogram* on an attribute X is constructed by partitioning the data distribution T into β mutually disjoint subsets called *buckets* and approximating the frequencies and values in each bucket in some common fashion. In order to construct a histogram on an attribute X, the frequency information for each attribute value is computed (either by scanning the whole relation or by sampling). Then the $< value, frequency >$ tuples of attribute X are used in the construction. Histograms are characterized by the following properties (the terminology is from Poosala et al. [16]):

- Partition Class: The class of histograms considered by the partitioning rule. For example a histogram is *serial* in the sense that histogram buckets correspond to groups of elements of T that are contiguous in the order of the sort parameter.
- Partition Constraint: The mathematical constraint that uniquely identifies the histogram within its partition class. For example *V-Optimal* histogram is the one in which weighted variance of the source values is minimized.
- Sort and Source Parameters: These are parameters derived from T. In the construction of histograms, $< value, frequency >$ tuples are sorted based on the sort parameters and then bucket boundaries are determined according to partition constraint applied on the source parameters. Possible parameters are attribute values (V), frequencies (F), and area (A), which is defined as frequency times spread of an attribute value ($a_i = f_i * s_i$) [16].

Example 2 *According to above characterization, some of the important histograms are defined as below:*

- *V-Optimal(F,F) is a histogram in which a weighted variance of the source values is minimized (partition constraint) and the $< value, frequency >$ tuples are sorted based on frequencies (F) and bucket boundaries are decided according to partition constraint using frequencies (F) as being source parameter.*
- *V-Optimal-End-Biased(F,F) is a histogram which places some of the highest and some of the lowest source parameter values in individual buckets and average out the remaining ones in a bucket (partition constraint) and $< value, frequency >$ tuples are sorted based on frequencies (F) and bucketization is applied using again frequencies (F) as source parameter.*
- *MaxDiff(V,F) is a histogram which places a bucket boundary between two source parameter values that are adjacent (in sort parameter order) if the difference between these values is one of the $\beta - 1$ largest such differences where β is the number of buckets. It sorts $< value, frequency >$ tuples based on attribute values (V), and uses frequencies (F) as source parameter.*

An important characteristics of V-Optimal(F,F) histograms is that they have been proven optimal for tree equality-join queries [8]. On the other hand finding a V-Optimal(F,F) histogram is computationally expensive. Instead Ioannidis and Poosala use V-Optimal-End-Biased(F,F) histograms which are practically easier to construct and are close to optimal histograms in most cases [8].

Below is an example on construction of End-Biased(F,F) and MaxDiff(V,F) histograms:

Example 3 *Let X={(b, 133), (a-, 97), (b+, 89), (b-, 62), (c+, 52), (c, 43), (a, 39), (c-, 37), (d, 12)} be the frequency vector of an attribute in a relation. The optimal bucketization with $\beta = 3$ for End-Biased(F,F) is shown in Figure 1 and that of MaxDiff(V,F) is shown in Figure 2. The corresponding histograms are shown in Figure 3 and Figure 4.*

b	$a-$	$b+, b-, c+, c, a, c-, d$
133	97	89, 62, 52, 43, 39, 37, 12

Fig. 1. Optimal bucketization for End-Biased(F,F) histogram

a	$a-, b+, b$	$b-, c+, c, c-, d$
39	97, 89, 133	62, 52, 43, 37, 12

Fig. 2. Optimal bucketization for MaxDiff(V,F) histogram

133	b
97	$a-$
47.7	$others$

Fig. 3. End-Biased(F,F) histogram

39	a
106.3	$a-, \ b+, \ b$
41.2	$others$

Fig. 4. MaxDiff(V,F) histogram

5.2. Construction of DFT Based End-Biased Histograms

We now develop end-biased histograms using the absolute values of DFT coefficients. Other classes of DFT based histograms could similarly be constructed, e.g., DFT based MaxDiff. In the case of End-Biased(F,F) histograms, in order to find the optimal histogram, the $< value, frequency >$ vector is sorted based on the frequency order and then bucketization is applied on this sorted vector. In the case of DFT based end-biased histogram development, the $< value, frequency >$ vector is kept in the attribute value order. The reason is that in the end-biased histograms, when the frequency values are sorted, the correspondence between attribute values and the frequencies are preserved. But in the DFT domain, the individual DFT coefficients do not have a direct relation with the individual attribute values. For example, consider an attribute a with a frequency vector X. Then the frequency value x_i corresponds to the attribute value a_i, but when we consider $\hat{X}$, the DFT of X, then the entry $\hat{x}_i$ does not have a direct relation with the attribute value a_i. That implies that in finding a result size S by using DFT based histograms, the

188.7	54.3, 54.3	30.7, 30.7, 26.6, 26.6, 30.8, 30.8

Fig. 5. Optimal bucketization for DFT-based End-Biased histogram

188.7	54.3	29.4

Fig. 6. DFT based End-Biased histogram

frequency vectors of two participating relations should be in the same order. For this reason, in order to eliminate the correspondence problem, the frequency vector X is kept in attribute value order and then the DFT transform is applied on it. After obtaining $\hat{X}$, histogram buckets are constructed based on absolute values of the DFT coefficients.

Consider the frequency vector X in the previous example. Then the DFT based end-biased histogram is built by using the absolute values $|\hat{X}|$ of the DFT coefficients of X:

$$|\hat{X}| = \{188.7, 54.3, 30.7, 26.6, 30.8, 30.8, 26.6, 30.7, 54.3\}.$$

Recall that the DFT coefficients $\hat{x}_2, \hat{x}_3, \ldots, \hat{x}_{\lceil \frac{N}{2} \rceil}$ and the coefficients $\hat{x}_N$, $\hat{x}_{N-1}, \ldots, \hat{x}_{\lceil \frac{N}{2} \rceil + 1}$ are complex conjugates. By considering this property we can re-arrange the vector $|\hat{X}|$ as

$$|\hat{X}| = \{188.7, 54.3, 54.3, 30.7, 30.7, 26.6, 26.6, 30.8, 30.8\}.$$

At this point, since we do not deal with individual attribute value correspondences, applying this operation does not cause any error in the process. Furthermore, in deciding the optimal end-biased histogram, we ignore duplications in the $|\hat{X}|$ and find the optimal bucketization by treating duplicates as one number. This increases the accuracy of the approximation by introducing less error due to averaging within the multi-valued bucket. The optimal bucketization for DFT based end-biased histogram is shown in Figure 5. Based on this bucketization, we construct our DFT based end-biased histogram as in Figure 6.

Notice that, unlike the End-Biased(F,F) histograms discussed previously, we do not show any correspondence between bucket values and the actual attribute values. This implies that by using DFT based end-biased histograms we do not store this correspondence information, thus we save in storage compared to the original End-Biased(F,F) histograms.

6. Development of Iterated DFT based Algorithms

6.1. Self Join using Absolute Values of DFT Coefficients

In this subsection we present the iterated DFT based approximation approach. We start by motivating our approach by using the simple case of a self join, i.e. when a relation is joined with itself on the same attribute. We start with the self join case because it is simple and motivates our general approach. Furthermore,

the self join is of particular interest for any application where a dataset is joined with itself on an attribute to find items that share a particular characteristic. This arises in several scientific dataset applications. For example consider a hydrological set of data samples taken over a time period of ocean temperatures. A scientist may be interested in identifying various datasets that share the same temperature. Our approach is particularly efficient for such queries.

Consider a relation R with an attribute a and let X be the frequency vector of a. We will assume for ease of representation that X is N-dimensional where N is of the form of $N = 2^k - 1$ for some integer k. The size S of the result for self join is the inner product of the frequency vector X by itself, i.e. $< X, X >$. Using Parseval's identity in the form (4) with $Y = X$ we have

$$
\begin{aligned}
< X, X > \ &= \ \sum_{i=1}^{N} x_i^2 \ = < \alpha, \alpha > + < \beta, \beta > \\
&= \ \sum_{i=1}^{N} (\alpha_i^2 + \beta_i^2) = \sum_{i=1}^{N} |\hat{x}_i|^2 \ = < \hat{X}, \hat{X} >,
\end{aligned}
\tag{6}
$$

where $|\hat{x}_i| = \sqrt{\alpha_i^2 + \beta_i^2}$ is the absolute value of $\hat{x}_i$. Hence the inner product of a real vector by itself can be calculated either by summing the squares of its values or by summing the squares of the absolute values of its complex DFT coefficients. However the DFT coefficients $\hat{x}_2, \hat{x}_3, \ldots, \hat{x}_N$ of a real vector occur in complex conjugate pairs for odd N. In particular, let a complex vector $\hat{X} = (\hat{x}_1, \hat{x}_2, \hat{x}_3, \ldots, \hat{x}_{N-1}, \hat{x}_N)$ be the DFT of a real vector X with $N = 2^k - 1$. Then the complex numbers $\hat{x}_2$ and $\hat{x}_N$ are complex conjugate pairs with $x_2 = \alpha_2 + I\beta_2$ and $x_N = \alpha_N + I\beta_N$ where $\beta_2 = -\beta_N$. Similarly $\hat{x}_3$ and $\hat{x}_{N-1}$ and so on are complex conjugate pairs.

Since conjugate numbers have the same absolute value, $|\hat{x}_i|^2 = |\hat{x}_{N-i+2}|^2$ for $2 \leq i \leq (N+1)/2$. Therefore combining with (6),

$$
\sum_{i=1}^{N} x_i^2 = \hat{x}_1^2 + 2 \sum_{i=2}^{\frac{N+1}{2}} |\hat{x}_i|^2.
\tag{7}
$$

This reduces the calculation of the desired sum on the left of (7) to the calculation of $\sum_{i=2}^{\frac{N+1}{2}} |\hat{x}_i|^2$. This latter sum is the inner product of the $(N-1)/2$-dimensional vector $X_1 = (|\hat{x}_2|, |\hat{x}_3|, \ldots, |\hat{x}_{(N+1)/2}|)$ by itself. Since this vector is also real, we can take its DFT and iterate this process by using Parseval's identity and (7) at each step. When $N = 2^k - 1$, the next vector to be considered has $(N-1)/2 = 2^{k-1} - 1$ elements, the one after that $2^{k-2} - 1$, and so on. Therefore this process ends in $k = \log(N+1) - 1$ iterations. Let $X_0 = X$ and define X_j to be the real vector $X_j = (x_{j1}, x_{j2}, \ldots, x_{jN_j})$ of length $N_j = 2^{k-j} - 1$ obtained after the j-th iteration, where the j-th iteration consists of taking the DFT of the N_{j-1}-st dimensional real vector X_{j-1}, and consequently setting $X_j = (|\hat{x}_{j-1,2}|, |\hat{x}_{j-1,3}|, \ldots, |\hat{x}_{j-1,N_j+1}|)$.

Let $\hat{x}_{01} = \hat{x}_1$, and let $\hat{x}_{j1}$ be the first element of the DFT $\hat{X}_j$ of the vector X_j. By iterating (7), we obtain the sequence of nonnegative real numbers

$\hat{x}_{01}, \hat{x}_{11}, \ldots, \hat{x}_{k-1,1}$ with the property that

$$< X, X > = \hat{x}_{01}^2 + 2 \left[\hat{x}_{11}^2 + 2 \left[\hat{x}_{21}^2 + \cdots \right] \cdots \right] \right] = \sum_{j=0}^{k-1} 2^j \hat{x}_{j1}^2 . \tag{8}$$

Hence the self join of a relation on an attribute a with N values can be calculated using the $k = \log(N+1)$ representative values computed. We call this algorithm *Self Join using Absolute Value (SJAV)*, and denote the sequence of k values obtained by this method from X by $SJAV(X)$.

Example 4 *Consider the vector $X = (54.34,\ 79.7,\ 25.88,\ 97.13,\ 10.74,\ 37.52,\ 66.94)$ with $< X, X > = 25413.9$. The application of SJAV on X is given in Figure 7. The resulting values $SJAV(X)$ that are stored are $\hat{x}_{01} = 140.7$, $\hat{x}_{11} = 49.33$, $\hat{x}_{21} = 13.69$.*
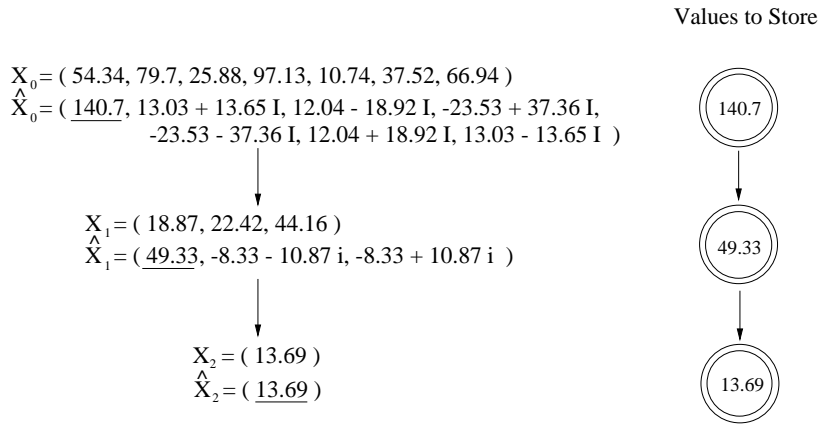
Values to Store

X_0 = (54.34, 79.7, 25.88, 97.13, 10.74, 37.52, 66.94)
$\hat{X}_0$ = (140.7, 13.03 + 13.65 I, 12.04 - 18.92 I, -23.53 + 37.36 I,
        -23.53 - 37.36 I, 12.04 + 18.92 I, 13.03 - 13.65 I)

140.7

X_1 = (18.87, 22.42, 44.16)
$\hat{X}_1$ = (49.33, -8.33 - 10.87 i, -8.33 + 10.87 i)

49.33

X_2 = (13.69)
$\hat{X}_2$ = (13.69)

13.69

Fig. 7. Example 1: The tree structure of *SJAV*.

By using these representative values for X, and using (8), we calculate

$$\begin{aligned} < X, X > &= \hat{x}_{01}^2 + 2\hat{x}_{11}^2 + 2^2\hat{x}_{21}^2 \\ &= 140.7^2 + 2*49.33^2 + 4*13.69^2 = 25413.9 \end{aligned}$$

The self join algorithm using absolute values calculates the exact size of the self join. This is quite significant since all prior algorithms can only provide an approximation to the size of the resulting set. Consider again our hydrology example with ocean temperature. If we assume that the range of water temperature is say from -10^0C to 110^0C, then only seven values need to be maintained by the query optimizer to calculate the exact self join sizes. If the domain of the self join attribute is 10^6, only 20 values need to be maintained by the query optimizer.

6.2. Approximating the Size of Arbitrary Joins

Unfortunately there is no direct equivalent to equation (6) as a way of representing the inner product of two different vectors. Suppose X and Y are two real vectors with DFTs $\hat{X} = \alpha + I\beta$ and $\hat{Y} = \gamma + I\delta$. Thus $\hat{x}_i = \alpha_i + I\beta_i$ and $\hat{y}_i = \gamma_i + I\delta_i$

for $i = 1, 2, \ldots, N$. Even though $< X, Y > \; = \; < \hat{X}, \hat{Y} >$ by Parseval's theorem, the expression (5) for the inner product can be expanded in the frequency domain as

$$< \hat{X}, \hat{Y} > = \sum_{i=1}^{N} \alpha_i \gamma_i + \sum_{i=1}^{N} \beta_i \delta_i,$$

which is not necessarily equal to

$$\sum_{i=1}^{N} |\hat{x}_i||\hat{y}_i| = \sum_{i=1}^{N} \sqrt{\alpha_i^2 + \beta_i^2} \sqrt{\gamma_i^2 + \delta_i^2}. \tag{9}$$

However by the Cauchy-Schwarz inequality the sum in (9) is an upper bound to the actual inner product $< \hat{X}, \hat{Y} >$ and therefore

$$\sum_{i=1}^{N} x_i y_i \leq \sum_{i=1}^{N} |\hat{x}_i||\hat{y}_i|.$$

At this point we will use the right hand side of the above inequality as an approximation to the left hand side and call this as *Approximation by Absolute Values (AAV)*. Now we can use the logarithmic representation obtained above by *SJAV* as follows: Suppose *SJAV* produces the sequences $SJAV(X) = \hat{x}_{01}, \hat{x}_{11}, \ldots, \hat{x}_{k-1,1}$ for X and $SJAV(Y) = \hat{y}_{01}, \hat{y}_{11}, \ldots, \hat{y}_{k-1,1}$ for Y. Then

$$\sum_{i=1}^{N} |\hat{x}_i||\hat{y}_i| = \sum_{j=0}^{k-1} 2^j \hat{x}_{j1} \hat{y}_{j1}.$$

This approximation to $< X, Y >$ using *SJAV* is exact when $X = Y$. When Y is close in the N-dimensional space to a constant multiple of X, we expect the error in the approximation to be small as a consequence of the equality condition in the Cauchy-Schwarz inequality. As the two vectors become more distant in this sense, the error will increase. This crude approximation to $< X, Y >$ has some remarkable asymptotic properties which we mention in Section 6 where we present our experimental results. What we refer to as the *Approximation by Absolute Value (AAV)* algorithm is outlined in Figure 8.

AAV Algorithm:

1. Calculate $SJAV(X)$ for all N-dimensional frequency vectors X by iterated DFT. Each $SJAV(X)$ is a sequence of $\log(N+1)$ real numbers.

2. Approximate $<X,Y>$ by

$$\sum_{j=0}^{k-1} 2^j \hat{x}_{j1} \hat{y}_{j1}$$

   where $SJAV(X) = \hat{x}_{01}, \hat{x}_{11}, \ldots, \hat{x}_{k-1,1}$ and $SJAV(Y) = \hat{y}_{01}, \hat{y}_{11}, \ldots, \hat{y}_{k-1,1}$.

Fig. 8. Outline of the AAV algorithm.

7. Improving AAV: Tree Based Algorithms

AAV is a technique that gives us an upper bound for the inner product of two real vectors, which is exact for self join operations. Now we use AAV as a subroutine to develop more accurate approximation by means of a tree based structure. This results in a sequence of methods called *Tree Approximation Algorithms* (TAA_l), one for every l ranging from 0 to $k-1$. This approach exploits the form of Parseval's identity given in equation (4) for N-dimensional real vectors X and Y. The main idea is to keep the real and the imaginary parts of the transformed vector $\hat{X}$ *exactly* (i.e. without losing the phase information by taking the absolute value as in AAV). This results in two real vectors α and β, each N-dimensional where $\hat{X} = \alpha + I\beta$. Since $\hat{x}_i$ and $\hat{x}_{N-i+2}$ are conjugate complex numbers for $i = 2, 3, \ldots, (N+1)/2$, the numbers in the list $\alpha_2, \alpha_3, \ldots, \alpha_N$ appear in pairs. Similarly the numbers in $\beta_2, \beta_3, \ldots, \beta_N$ appear in pairs. Consider the vectors $h(\alpha) = (\alpha_2, \alpha_3, \ldots, \alpha_{(N+1)/2})$ and $h(\beta) = (\beta_2, \beta_3, \ldots, \beta_{(N+1)/2})$. We invoke two instances of the AAV algorithm: one on $h(\alpha)$ corresponding to the real part of $\hat{X}$ (left subtree), and one on $h(\beta)$ corresponding to the imaginary part of $\hat{X}$ (right subtree). Thus these two real vectors $h(\alpha)$ and $h(\beta)$ are treated as raw inputs to the AAV algorithm, producing two lists $AAV(h(\alpha))$ and $AAV(h(\beta))$ consisting of $k-1$ nonnegative real numbers each, where $N = 2^k - 1$. We demonstrate this process on an example.

Example 5 *Consider the vector X of the previous example, $N = 2^3 - 1 = 7$. The application of TAA_1 in which we now use AAV at level $l = 1$ results in the tree structure given in Figure 9.*

For this example we need to store two lists $AAV(h(\alpha))$ and $AAV(h(\beta))$ of length 2 each, plus the number $\hat{x}_{01}$ that is computed at the root of the tree. The structure of the total list to be stored for X is itself a tree which is depicted on the right side of Figure 9. Since the calculation at the root of the tree in Figure 9 is exact and the application of the approximate AAV algorithm starts at level $l = 1$ of the tree, this algorithm is denoted by TAA_1. Note that AAV can be viewed as the special case TAA_0, starting the application at level 0. We see that this immediately generalizes

12

Fig. 9. Example 2: The tree structure of TAA_1, corresponding to level $l = 1$.

to arbitrary level l, and results in an algorithm TAA_l for $l = 0, 1, \ldots, k-1$ in which the data is exact up to level l and the AAV algorithm is applied to the vectors at the next level. As a boundary case, we set $AAV(X)=X$ if X is a vector of length 1.

Example 6 *Continuing with the vector X of the previous examples, the application of TAA_2 in which we use AAV after level $l = 2$ gives the tree structure of Figure 10.*
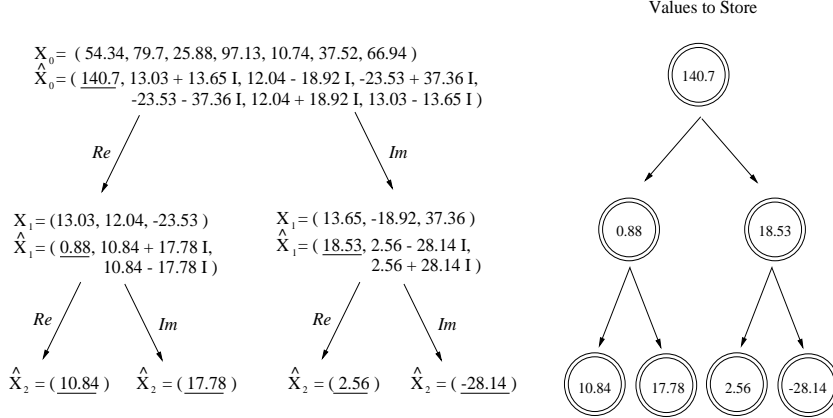
Fig. 10. Example 3: The tree structure of TAA_2, corresponding to level $l = 2$.

The generation of the representative data in the form of a tree for a real vector X using algorithm TAA_l can be described recursively in Figure 11.

7.1. Tree Approximation Algorithms and the TAA Product

In the above, we presented the TAA_l technique to extract a new set of representative data for a vector. In this process, by delaying the application of AAV i.e. increasing the value of l, we will be storing more exact data and hence the approximation based on this data will result in more accurate estimations. In this

---
TAA data generation Algorithm

1. Suppose $N = 2^k - 1$, X is an N-dimensional real vector, and $\hat{X} = (\hat{x}_1, \hat{x}_2, \ldots, \hat{x}_N) = \alpha + I\beta$.

2. If $l = 0$, then $TAA_l = AAV$, and the tree T of values computed as $AAV(X)$ is a linear list of length $\log(N + 1)$, where N is the length of X.

3. If $l > 0$, then T has two principal subtrees T_{Re} and T_{Im}. The value at the root of T is $\hat{x}_1$. The subtrees T_{Re} and T_{Im} are constructed by applying algorithm TAA_{l-1} to the vectors $h(\alpha)$ and $h(\beta)$, respectively.

Fig. 11. TAA representative tree generation algorithm.

subsection we present the approximation algorithm that uses the representative data values found in the above process.

Definition 1 *Suppose T_1 and T_2 are two trees of height H and the same shape where the data field in each node stores a real numerical value. Let $Vec_l(T_1)$ denote the vector of values in the nodes at level l of T_1 ordered from left to right. Similarly, define the vector $Vec_l(T_2)$ for T_2. The TAA product of T_1 and T_2 is defined as*

$$TAA(T_1, T_2) = \sum_{l=0}^{H-1} 2^l < Vec_l(T_1), Vec_l(T_2) > \tag{10}$$

In other words we first imagine T_1 and T_2 lined up so that corresponding nodes are on top of one another. We then multiply the numerical values in the paired nodes, further multiply this number by 2**(level of the node), and then add up the resulting numbers from each node.

Remark 1 *In terms of the TAA product, the approximation given by AAV to $< X, Y >$ is simply $TAA(T_X, T_Y)$ where T_X and T_Y are the trees (chains) of values $AAV(X)$ and $AAV(Y)$ respectively, produced by algorithm AAV.*

TAA algorithm at level l for the approximation of result sizes of join operations can be described succinctly in terms of the TAA product. This is shown in Figure 12.

Theorem 2 *Suppose X and Y are N-dimensional real vectors where $N = 2^k - 1$. Let T_X^l and T_Y^l denote the trees that are generated by TAA_l for the representation of X and Y respectively, $l = 0, 1, \ldots, k - 1$. Then*

- *Each approximation $TAA(T_X^l, T_Y^l)$ computed as given in (10) is an upper bound to the inner product $< X, Y >$.*
  Furthermore

$$TAA(T_X^0, T_Y^0) \geq TAA(T_X^1, T_Y^1) \geq \cdots \geq TAA(T_X^{k-1}, T_Y^{k-1}),$$

  and $TAA(T_X^{k-1}, T_Y^{k-1}) = < X, Y >$ is exact.
- *The number of elements stored in the representative tree T_X^l generated by TAA_l is $2^l(k - l + 1) - 1$.*

14

Fig. 12. Outline of the TAA algorithm.

We give an example for the calculation of the approximation to $< X, Y >$ using the TAA product of the trees T_X and T_Y for $l = 2$.

Example 7 *Take X as in the previous examples. The tree T_X produced by algorithm TAA_2 for X is as given on the right hand side of Figure 10. Take $Y = (69.97, 82.28, 49.67, 36.22, 29.81, 95.85, 51.74)$. The application of TAA_2 on Y produces the tree T_Y given in Figure 13. The approximation to the inner product $< X, Y >$*


Fig. 13. The tree T_Y: Application of TAA_2 to vector Y.

using the formulation in (10) on T_X and T_Y is calculated as follows:

$l = 2$: *$Vec_2(T_X) = (10.84, 17.78, 2.56, -28.14)$,*
$Vec_2(T_Y) = (16.13, -14.18, -14.96, -2.04)$,
contribution to (10): $2^2 * (10.84 * 16.13 - 17.78 * 14.18 - 2.56 * 14.96 + 28.14 * 2.04)$

$l = 1$: *$Vec_1(T_X) = (0.88, 18.53)$, $Vec_1(T_Y) = (8.1, 17.9)$,*
contribution to (10): $2^1 * (0.88 * 8.1 + 18.53 * 17.9)$

$l = 0$: *$Vec_0(T_X) = (140.7)$, $Vec_0(T_Y) = (157.06)$,*
contribution to (10): $2^0 * 140.7 * 157.06$

Summing the contributions from each level, we find $TAA(T_X, T_Y) = 22544$. Since we are using TAA_2 and $2 = k - 1$, this approximation is exact, i.e. $< X, Y > =$

15