

Adaptive and Incremental Processing for Distance Join Queries

Hyoseop Shin, Bongki Moon, and Sukho Lee

Abstract—A spatial distance join is a relatively new type of operation introduced for spatial and multimedia database applications. Additional requirements for ranking and stopping cardinality are often combined with the spatial distance join in online query processing or Internet search environments. These requirements pose new challenges as well as opportunities for more efficient processing of spatial distance join queries. In this paper, we first present an efficient k -distance join algorithm that uses spatial indexes such as R-trees. Bidirectional node expansion and plane-sweeping techniques are used for fast pruning of distant pairs, and the plane-sweeping is further optimized by novel strategies for selecting a sweeping axis and direction. Furthermore, we propose adaptive multistage algorithms for k -distance join and incremental distance join operations. Our performance study shows that the proposed adaptive multistage algorithms outperform previous work by up to an order of magnitude for both k -distance join and incremental distance join queries, under various operational conditions.

Index Terms—Spatial databases, k -distance join, incremental distance join, adaptive query processing, multistage query processing, plane sweeping, sweeping index, estimating cutoff distance.

1 INTRODUCTION

A spatial distance join operation was recently introduced to spatial databases to associate one or more sets of spatial data by distances between them [16]. A distance is usually defined in terms of spatial attributes, but it can be defined in many different ways according to various application specific requirements. In multimedia and image database applications, for example, other metrics such as a *similarity distance function* can be used to measure a distance between two objects in a feature space.

In online decision support and Internet search environments, it is quite common to pose a query that finds the best k matches or reports the results incrementally in the decreasing order of well-matchedness. These type of operations allow users to interact with database systems more effectively and focus on the “best” answers. Since users can say “it is enough already” at any time after obtaining the best answers [9], the waste of system resources can be reduced and, thereby, delivering the results to users more quickly.

This ranking requirement is often combined with a spatial distance join query, and the ranking requirement provides a new opportunity of optimization for spatial distance join processing [10], [12]. For example, consider a query that retrieves the top k pairs (i.e., the nearest pairs) of hotels and restaurants:

```
SELECT h.name, r.name
FROM Hotel h, Restaurant r ORDER BY
distance(h.location, r.location)
STOP AFTER k;
```

For a relatively small stopping cardinality k , the processing time can be reduced significantly by sorting only a fraction of intermediate results enough to produce the k -nearest pairs, instead of sorting an entire set of intermediate results (i.e., a Cartesian product of hotels and restaurants).

A spatial distance join query with a stopping cardinality can be formulated as follows:

$$\sigma_{\text{dist}(r,s) < \mathcal{D}_{\max}}(R \bowtie S),$$

where $\text{dist}(r, s)$ is a distance between two spatial objects $r \in R$ and $s \in S$, and \mathcal{D}_{\max} is a cutoff distance that is determined by a stopping cardinality k and the spatial attribute values of two data sets R and S . It may then be argued that a spatial distance join query can be processed by a spatial join operation [1], [7], [8], [18], [19], [23] followed by a sort operation. Specifically, if a \mathcal{D}_{\max} value can be predicted precisely for a given stopping cardinality k , we can use a spatial join algorithm with a *within* predicate instead of an *intersect* predicate to find the k -nearest pairs of objects. Then, a sort operation will be performed only on the k pairs of objects.

In practice, however, it is almost impossible to estimate an accurate \mathcal{D}_{\max} value for a given stopping cardinality k and, to the best of our knowledge, no method for estimating such a cutoff value has been reported in the literature. If the \mathcal{D}_{\max} value is overestimated, then the results from a spatial join operation may contain too many candidate pairs, which may cause a long delay in a subsequent stage to sort all the candidate pairs. On the other hand, if the \mathcal{D}_{\max} value is underestimated, a spatial join operation may not return a sufficient number of object pairs. Then, the spatial join

- H. Shin is with the Software Center, Samsung Electronics, Apkujong Bldg. 599, Shinsa-Dong, Kangnam-Ku, Seoul, Korea 135-120. E-mail: hyoseop@samsung.com.
- B. Moon is with the Department of Computer Science, University of Arizona, Gould-Simpson Bldg., Rm. 715, 1040 E. 4th St., Tucson, AZ 85721-0077. E-mail: bkmooon@cs.arizona.edu.
- S. Lee is with the School of Computer Science and Engineering, Seoul National University, Shillim-Dong, Kwanak-Ku, Seoul, Korea 151-742. E-mail: shlee@cse.snu.ac.kr.

Manuscript received 1 Nov. 2000; revised 1 Mar. 2002; accepted 6 Aug. 2002. For information on obtaining reprints of this article, please send e-mail to: tkde@computer.org, and reference IEEECS Log Number 113092.




Fig. 1. Hierarchy of spatial containment of R-tree nodes. (a) Tree-structured spatial index. (b) Spatial containment.

operation should be repeated with a new estimate of \mathcal{D}_{max} until k or more pairs are returned. This may cause a significant amount of waste in processing time and resources.

There is another reason that makes it impractical to apply a spatial join algorithm to spatial distance join queries. A spatial join query is typically processed in two steps, *filter* and *refinement*, as proposed in [21]. In a filter step, MBR approximations are used to find pairs of potentially intersected spatial objects. Then, in a refinement step, it is guaranteed that all the qualified (i.e., actually intersected) pairs can be produced from the results returned from the filter step.

In contrast, it is completely unreasonable to process a spatial distance join query in two separate filter and refinement steps, because of the fact that a filtering process is based on MBR approximations. A set of object pairs sorted by distances measured by MBR approximations does not reflect a true order based on actual representations. This is because, for any two pairs of spatial objects $\langle r_1, s_1 \rangle$ and $\langle r_2, s_2 \rangle$, the fact that $dist(MBR(r_1), MBR(s_1)) < dist(MBR(r_2), MBR(s_2))$ does not necessarily imply that $dist(r_1, s_1) < dist(r_2, s_2)$. Consequently, any processing done in the filter step will be of no use for finding the k -nearest object pairs.

In this paper, we propose new strategies for efficiently processing spatial distance join queries combined with ranking requirements. The main contributions of the proposed solutions are:

- New efficient methods are proposed to process distance join queries using spatial index structures such as R-trees. *Bidirectional node expansion* and *optimized plane-sweep* techniques are used for fast pruning of distant pairs, and the plane-sweep is further optimized by novel strategies for selecting a sweeping axis and direction and by using maximum distance for breaking tied pairs.
- Adaptive multistage algorithms are proposed to process distance join queries in a way that the k -nearest pairs are returned *incrementally*. When a stopping cardinality is not known a priori (e.g., in online query processing environments or a complex query containing a distance join as a subquery whose results need to be pipelined to the next stage of the complex query), the adaptive multistage algorithms can produce pairs of objects in a stepwise manner.
- We provide a systematic approach for *estimating the maximum distance* for a distance join query with a stopping cardinality. This estimated distance allows the adaptive multistage algorithms to avoid a *slow start* problem, which may cause a substantial delay

in the query processing. This approach for estimating the maximum distance also allows the size of memory to be parameterized into a queue management scheme, so that data movement between memory and disk can be minimized.

The proposed algorithms achieve up to an order of magnitude performance improvement over previous work for both k -distance join and incremental distance join queries, under various operational conditions.

The rest of this paper is organized as follows: Section 2 surveys the background and related work on processing spatial distance join queries. Major limitations of previous work are also discussed in this section. In Section 3, we present a new improved algorithm based on bidirectional expansion and optimized plane-sweep techniques for k -distance join queries. In Section 4, adaptive multistage algorithms are presented for k -distance join and incremental distance join queries. A queue management scheme parameterized by memory capacity is also presented. Section 5 presents the results of experimental evaluation of the proposed solutions. Finally, Section 6 summarizes the contributions of this paper.

2 BACKGROUND AND PREVIOUS WORK

A spatial index structure R-tree and its variants [3], [6], [14] have been widely used to efficiently access multidimensional data—either spatial or point. Like other tree-structured index methods, an R-tree index partitions a multidimensional space by grouping objects in a hierarchical manner. A subspace occupied by a tree node is always contained in the subspace of its parent node. This hierarchy of spatial containment between R-tree nodes is readily used by spatial distance join algorithms as well as spatial join algorithms.

Suppose r and s are nonleaf nodes of two R-tree indexes R and S , respectively, as in Fig. 1. Then, the minimum distance between r and s is always less than or equal to the minimum distance between one of the entries of r and one of the entries of s . Likewise, the maximum distance between r and s is always greater than or equal to the maximum distance between one of the entries of r and one of the entries of s . This observation leads to the following lemma.

Lemma 1. For two R-tree indexes R and S , if neither $r \in R$ nor $s \in S$ is a root node, then

$$\begin{aligned} dist(r, s) &\geq dist(parent(r), parent(s)), \\ dist(r, s) &\geq dist(r, parent(s)), \\ dist(r, s) &\geq dist(parent(r), s), \end{aligned} \quad (1)$$

where $dist(r, s)$ is the minimum distance between the MBR representations of r and s .

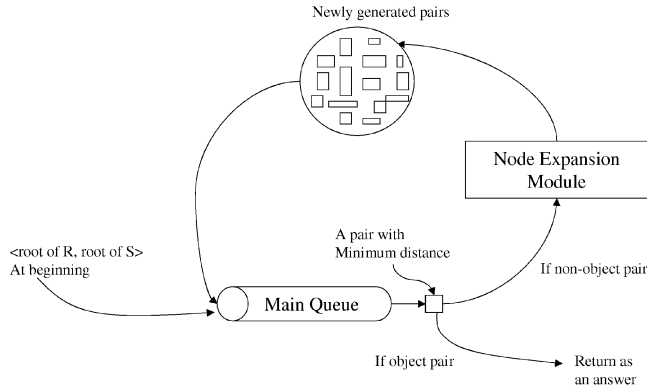


Fig. 2. Framework of incremental distance join (IDJ) processing.

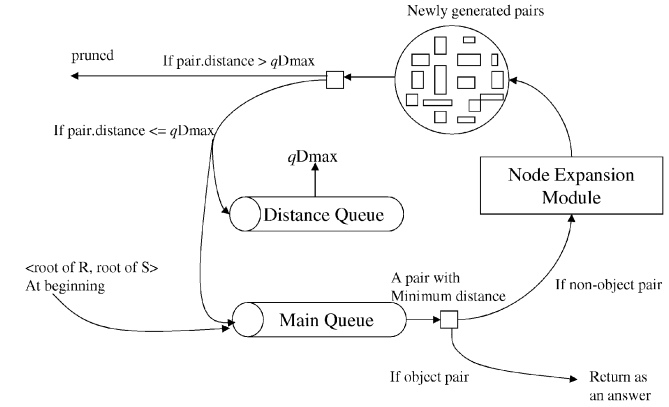
Proof. From the observation above. \square

Lemma 1 allows us to limit the search space, while R-tree indexes are traversed in a top-down manner to process a spatial distance join query. For example, if a pair of nonleaf nodes $\langle r, s \rangle$ turn out to be too far from each other (or their distance is over a certain threshold), then it is not necessary to traverse further down the tree indexes below the nodes r and s . Thus, this lemma provides the key leverage to processing distance join queries efficiently using R-tree indexes.

2.1 Incremental Distance Join and k -Distance Joins

During top-down traversals of R-tree indexes, it is desirable to store examined node pairs in a priority queue, where the node pairs are kept in an increasing order of distances. We call it a *main queue* as opposed to a *distance queue* we will describe later. The main queue initially contains a pair of the root nodes of two R-tree indexes. Each time a pair of nonobject nodes are retrieved from the main queue, the entries of one node are paired up with the entries of the other to generate a new set of node pairs, which are then inserted into the main queue. This process that we call *node expansion* is repeated until the main queue becomes empty or until stopped by an interactive user. If an element retrieved from the main queue is a pair of two objects, the pair is returned immediately to the user as a query result. This is how a spatial distance join query is processed *incrementally*. Fig. 2 depicts a typical framework of processing an incremental distance join (IDJ) query using R-tree indexes.

A distance join query is often given with a stopping cardinality k as in the “stop after” clause of the sample query in Section 1. Since it is known a priori how many object pairs need to be produced for a distance join query, this knowledge can be exploited to improve the performance of the query processing. Suppose a maximum of k -nearest pairs of objects are to be retrieved by a query. One plausible approach is to maintain k candidate pairs of objects during the entire course of query processing. As they are the k -nearest object pairs known at each stage of query processing, any pair of nodes (and any pair of their entries) whose distance is greater than *all* of the k candidate pairs cannot be qualified as a query result. Thus, we can use another priority queue to store the k minimum

Fig. 3. Framework of k -distance join (KDJ) processing.

distances and use the queue to avoid having to insert unqualified pairs into the main queue during the node expansions. We call the priority queue a *distance queue*. Fig. 3 depicts a typical framework of processing a k -distance join (KDJ) query using R-tree indexes and both main and distance queues.

Both main and distance queues can be implemented by heap structures. A main queue is normally implemented as a *min-heap* because the query results are produced in an increasing order of distances. In contrast, a distance queue should be implemented as a *max-heap* that can store at most k distance values. The cutoff distance is determined by the maximum value among the k distances stored in the distance queue. (When the distance queue contains less than k distances, the cutoff distance is set to an infinity.) Pruning node pairs by the distance queue was shown to be very efficient from our experiments, especially when k was rather small. In the rest of the paper, we use qD_{max} to denote the cutoff distance from the distance queue.

2.2 Previous Work

In [16], the authors present both unidirectional and bidirectional node expansion, but conclude based on their experiments that the former provides better performance due to fewer node pairs being produced by their algorithm. When a pair of nodes $\langle r, s \rangle$ are retrieved from a main queue, either node r is paired up with the entries of s , or node s is paired up with the entries of r . None of the pairs are generated from an entry of r and an entry of s . The advantage of the unidirectional expansion is that the number of pairs generated at each expansion step is limited to the fanout of an R-tree index being traversed, and an explosion of the main queue can be avoided. As is acknowledged by the authors of the algorithms, however, the main disadvantage of this approach is that the unidirectional expansion may lead to each node being accessed from disk more than necessary. Also, the unidirectional expansion requires pairing up node r exhaustively with all the entries of node s or vice versa.

For a spatial distance join query with a relatively small stopping cardinality k , the use of a distance queue is an effective means to prevent distant pairs from entering a main queue. For a large k value, however, the distance queue may not work well as an effective pruning tool

```

1: set  $AnswerSet \leftarrow$  an empty set;
2: set  $Q_M, Q_D \leftarrow$  empty main and distance queues;
3: insert a pair  $\langle R.root, S.root \rangle$  into the main queue  $Q_M$ ;
4: while  $|AnswerSet| < k$  and  $Q_M \neq \emptyset$  do
5:   set  $c \leftarrow$  dequeue( $Q_M$ );
6:   if  $c$  is an  $\langle object, object \rangle$  then  $AnswerSet \leftarrow \{c\} \cup AnswerSet$ ;
7:   else  $PlaneSweep(c)$ ;
   end

   procedure  $PlaneSweep(\langle l, r \rangle)$ 
8:   set  $L \leftarrow$  sort_axis( $\{entries\ of\ l\}$ ); // Sort the entries of  $l$  by axis values.
9:   set  $R \leftarrow$  sort_axis( $\{entries\ of\ r\}$ ); // Sort the entries of  $r$  by axis values.
10:  while  $L \neq \emptyset$  and  $R \neq \emptyset$  do
11:     $n \leftarrow$  a node with the min axis value  $\in L \cup R$ ; //  $n$  becomes an anchor.
12:    if  $n \in L$  then
13:       $L \leftarrow L - \{n\}$ ;  $SweepPruning(n, R)$ ;
    else
14:       $R \leftarrow R - \{n\}$ ;  $SweepPruning(n, L)$ ;
    end
  end

  procedure  $SweepPruning(n, List)$ 
15:  for each node  $m \in List$  in an increasing order of axis value do
16:    if  $axis\_distance(n, m) > qD_{max}$  then return; // No more candidates.
17:    if  $real\_distance(n, m) \leq qD_{max}$  then
18:      insert  $\langle n, m \rangle$  into  $Q_M$ ;
19:      if  $\langle n, m \rangle$  is an  $\langle object, object \rangle$  then insert  $real\_distance(n, m)$  into  $Q_D$ ; //
         $qD_{max}$  modified.
    end
  end

```

Fig. 4. Algorithm 1 (\mathcal{B} -KDJ): K-distance join algorithm with bidirectional expansion and plane sweep.

because the cutoff value stored in the distance queue may remain too high for a long duration. This may in turn lead to a long delay particularly in the early stage of query processing. For these reasons, the previous algorithms suffer from poor performance for a k -distance join query with a large k and an incremental distance join query, for which k is unknown in advance.

Moreover, there is an important issue that was not fully addressed in [16]. A hybrid memory/disk technique was proposed as a queue management scheme, which partitions a queue based on the distance range. This technique keeps a partition in the shortest distance range in memory, while the rest of partitions are stored on disk. However, no mechanism was provided to determine a boundary distance value between the partition in memory and the rest, which may have a crucial impact on the performance of queue management.

Recently, a few recursive and iterative algorithms have been proposed [11]. These algorithms make use of various distance metrics such as $MinMax$, $MinMin$, $MaxMin$, and $MaxMax$ to find k closest pairs. Without using a main queue, the recursive algorithms access R-tree nodes recursively following priorities given to the entry pairs within a pair of the parent nodes. The iterative algorithm (called *heap* algorithm) is fairly similar to Hjaltason and Samet's distance join algorithm [16] in that both the algorithms use a distance queue to maintain k candidate pairs during node expansion. One notable difference is that the *heap* algorithm does not store object pairs in the main queue to minimize the size of a main queue. Instead, the *heap* algorithm uses a distance (or candidate) queue to store the k closest pairs of objects. Although this does not guarantee that the main queue always fits in memory, the performance gain by not

storing object pairs in the main queue could be nontrivial, given the potentially large number of object pairs produced by node expansion. Since the *heap* algorithm maintains only the k candidates, the stopping cardinality k must be known a priori. In other words, the *heap* algorithm cannot be used for incremental distance join queries.

Several closely related studies for nearest neighbor queries have been reported in the literature. Among those are nearest neighbor search algorithms based on Voronoi cells [2], [5] and branch and bound techniques [26], [27], a nearest neighbor search algorithm for ranking requirement [15], and multistep k -nearest neighbor search algorithms [17], [28].

Another closely related issue is estimating spatial join selectivity. Some estimation techniques proposed to use supplementary structures such as histograms [24] and wavelets [29]; other estimation techniques were based on uniformity assumption [22] and fractal dimensions for self joins [4]. Recently, Faloutsos et al. [13] proposed a power law to predict the selectivity of spatial join and to estimate the distance of the k th closest pair. This power law will be used to estimate cutoff distances for the adaptive distance join algorithm proposed in this paper.

3 OPTIMIZED PLANE-SWEEP FOR FAST PRUNING

In this section, we propose a new distance join algorithm \mathcal{B} -KDJ (Bidirectional expanding k -Distance Join) using a *bidirectional* node expansion in an attempt to avoid redundant accesses to R-tree nodes (Fig. 4). As is pointed out in Section 2, distance join algorithms based on a unidirectional expansion require accessing an R-tree node more than those based on bidirectional expansions. Under




Fig. 5. Bidirectional node expansion with plane sweeping.

the bidirectional node expansion, for a pair $\langle r, s \rangle$, each of the entries of r is paired up with each of the entries of s . This is essentially a Cartesian product, which may generate more redundant pairs than the unidirectional expansion does. Nonetheless, we will show that the \mathcal{B} -KDJ algorithm can effectively avoid generating redundant pairs by a plane-sweeping technique [25] and novel strategies for choosing an axis and a direction for sweeping. The \mathcal{B} -KDJ algorithm is described in Algorithm 1.

3.1 Bidirectional Node Expansion

Like the distance join algorithms proposed in [16], the \mathcal{B} -KDJ algorithm uses qD_{max} from the distance queue Q_D as a cutoff value to examine node pairs. If a pair of nodes $\langle r, s \rangle$ removed from the main queue are a pair of objects, then the object pair is returned as a query result. Otherwise, the pair is expanded by the *PlaneSweep* procedure for further processing.

Assume that a sweeping axis (i.e., x or y dimensional axis) and a sweeping direction (i.e., forward or backward) are determined, as we will describe in Sections 3.2 and 3.3. Then, the entries of r and s are sorted by x or y coordinates of one of the corners of their MBRs in an increasing or decreasing order, depending on the choice of sweeping axis and sweeping direction. Each node encountered during a plane sweep is selected as an *anchor*, and it is paired up with entries in the other group. For example, in Fig. 5, an entry r_1 of r is selected as an anchor, and the entries s_1 , s_2 , s_3 , and s_4 of s are examined for pairing, as they are within qD_{max} distance from r_1 along the sweeping axis (lines 11-14 and line 16).

Since an axis distance between any pair $\langle r, s \rangle$ is always smaller than or equal to their real distance (i.e., $axis_distance(r, s) \leq real_distance(r, s)$), real distances are computed only for nodes whose axis distances from the anchor are within the current qD_{max} value (line 17). Given that a real distance is more expensive to compute than an axis distance, it may yield nontrivial performance gain. Then, each pair whose real distance is within qD_{max} is inserted into the main queue Q_M (line 18). If it is a pair of objects, then update the current qD_{max} value by inserting the real distance of the object pair into the distance queue Q_D (line 19).

There are alternatives as to what pairs are to be inserted into a distance queue: 1) any pairs encountered during node




Fig. 6. Effect of right selection of the sweeping axis.

expansions or 2) pairs of objects only. If a pair of nonobject R-tree nodes is inserted into a distance queue, then its distance value should be the *maximum* distance (instead of minimum distance), and the minimum number of object pairs that can possibly be generated from the node pair should be maintained in the distance queue, as pointed out in [16]. The minimum number of object pairs can be estimated based on the minimum node occupancy. Since the maximum distance tends to be larger than those of pairs of objects, most nonobject pairs are inserted into a distance queue only to be removed from the distance queue without reducing qD_{max} value. Consequently, the potential benefit from inserting nonobject pairs is expected to be insignificant. More often than not in our experiments, the query processing slowed down slightly due to the overhead of inserting nonobject pairs. Thus, we decide to follow the second option in this paper.

For a relatively small qD_{max} value and two sets of evenly distributed spatial objects, the number of pairs for which the \mathcal{B} -KDJ algorithm computes real distances and performs queue management operations is expected to be roughly $\mathcal{O}(|r| + |s|)$. This justifies the additional cost of sorting entries for plane-sweeping, because the overall cost of the \mathcal{B} -KDJ algorithm would otherwise be $\mathcal{O}(|r| \times |s|)$ by Cartesian products.

3.2 Sweeping Axis

We can improve the \mathcal{B} -KDJ algorithm one step further by deciding the sweeping axis and direction on an individual pair basis. Intuitively, if entries (or data objects) are spread more widely along one dimension (say, x) than the other dimensions, then the bidirectional node expansion is likely to generate a smaller number of node pairs to compute the real distances by plane-sweeping along the dimension x . This is because, when the nodes are more widely spread along a sweeping axis, the chance that a pair of nodes are within a qD_{max} distance along the sweeping axis is lower. For a pair of parent nodes shown in Fig. 6, as an example, it would be better to choose the y -axis as a sweeping axis, as the entries are more widely spread along the y -dimension. On the other hand, if the x -axis is chosen as a sweeping axis, no pair of the entries will be pruned by x -axis distance comparison with qD_{max} because the x -axis distance between any pair of the entries is shorter than the qD_{max} value.

Formally, we define a new metric **sweeping index** as follows, and we use the metric to determine which axis a plane-sweep will be performed on. For a given pair $\langle r, s \rangle$ of R-tree nodes and a given qD_{max} value, we can compute a




Fig. 7. Sweeping index.

sweeping index for each dimension. Conceptually, a sweeping index is a normalized estimation of the number of node pairs we need to compute the real distances for based on the assumption that data objects are uniformly distributed.¹

$$\text{Sweeping Index}_x = \int_0^{|r|_x} \text{Overlap}(qD_{max}, r, t) dt + \int_0^{|s|_x} \text{Overlap}(qD_{max}, s, t) dt. \quad (2)$$

In the first integral term of the equation above, $|r|_x$ is the side length of node r along the dimension x . The function $\text{Overlap}(qD_{max}, r, t)$ is a portion of the side length of s along the dimension x , overlapped with a window of length qD_{max} whose left end point is located at a point t within $|r|_x$ (i.e., $0 \leq t \leq |r|_x$). (See the left diagram in Fig. 7.) Thus, $\text{Overlap}(qD_{max}, r, t)/|s|_x$ represents a fraction of s 's entries intersected with a window $[t, t + qD_{max}]$. The value of the function varies as the window moves along the dimension x from $[0, qD_{max}]$ to $[|r|_x, |r|_x + qD_{max}]$. Therefore, the first integral term represents a relative estimation of the number of s 's entries encountered during the plane-sweeps performed for all the entries of r . The second integral term is symmetric with the first integral, and an identical description can be offered by exchanging r and s .

A smaller sweeping index indicates that the bidirectional expansion needs to compute real distances for a smaller number of nodes pairs. For this reason, the \mathcal{B} -KDJ algorithm chooses a dimension with the minimum sweeping index as a sweeping axis.

One thing we may be concerned about is the cost of computing a sweeping index for each dimension. The sweeping index may appear expensive to compute, as it includes two integral terms. For given $|r|_x$ and $|s|_x$ values and the current qD_{max} value available from the distance queue, however, the sweeping index is reduced to a formula that involves only a few simple arithmetic operations. Suppose nodes r and s are not intersected along a dimension x , the minimum x -axis distance between them is β and node r appears before node s in the plane-sweep direction along x -axis. (Again, see the left diagram in Fig. 7.) Then, the second integral term of (2) becomes zero because all the entries of r have already been swept when the first entry of s is encountered. The first integral term varies

1. An actual number of node pairs for which we need to compute the real distances would be computed by counting the number of s 's entries within qD_{max} axis distance from each entry of r , counting the number of r 's entries within qD_{max} axis distance from each entry of s , and then adding all the counts and dividing the count sum by two. However, this process will be very expensive.

depending on the conditions among qD_{max} , $|r|_x$, and $|s|_x$ values and the proximity (i.e., β) of nodes r and s along a chosen dimension. The right diagram in Fig. 7 illustrates how we can compute the first integral term and obtain a simple expression when a condition $\beta \leq qD_{max} \leq \beta + \min\{|r|_x, |s|_x\}$ is satisfied.

If nodes r and s are not separated, both the integral terms of (2) become nonzero. By similar reasoning, each integral term is also transformed into a formula with only a few simple arithmetic operations. Table 1 summarizes the formulae of the sweeping index for nodes r and s that are in three different spatial relationships: s is separated from, intersected with, or contained in r . The values of α , β , and δ in Table 1 are determined by the side lengths of r and s and their spatial relationship as illustrated in Fig. 8.

Considering that each R-tree node may contain hundreds of entries, it will be a trivial cost to compute a sweeping index for each dimension, while the performance gain by the sweeping axis selection is expected to be significant. This is empirically corroborated by our experiments in Section 5.

3.3 Sweeping Direction

Once a sweeping axis is determined, a sweeping direction can be chosen to be either a *forward* sweep or a *backward* sweep. For a pair of nodes r and s , we can define the forward and backward sweeps as follows:

- A forward plane-sweep scans the entries of r and s in an increasing order of coordinates along the chosen sweeping axis.
- A backward plane-sweep scans the entries of r and s in a decreasing order of coordinates along the chosen sweeping axis.

Consider nodes r and s projected on a sweeping axis. The projected images generate three consecutive closed intervals on the sweeping axis, unless the projected images are completely overlapped. For example, if nodes r and s are intersected as in Fig. 8b, an interval in the left is projected from r , one in the middle from both r and s , and one in the right from s . The interval in the middle may be projected from none of r and s if r and s are separate as in Fig. 8a. Both the intervals in the left and right may be projected from the same node if one node is contained in the other as in Fig. 8c.

However, it does not matter which node an interval is projected from because a sweeping direction is determined solely on the relative length of the intervals in the left and right (i.e., α and δ). A sweeping direction is determined by comparing the length of the left and right intervals: *If the left projected interval is shorter than the right one ($\alpha < \delta$), then a forward direction is chosen. Otherwise, a backward direction is chosen.* By this strategy of choosing a sweeping direction, a pair of nodes closer to each other are likely to be examined earlier than those farther from each other. This in turn allows a pair of closer nodes to be inserted into the main queue (and the distance queue as well if they are an object pair) and helps reduce the qD_{max} value more rapidly.

In summary, the sweeping axis selection improves the bidirectional node expansion step by pruning more pairs of entries whose axis distances are larger than the qD_{max} value,

TABLE 1
The First and Second Integral Terms of the Sweeping Index for r and s

r and s	The first integral term of Equation (2)
Separated	$\begin{cases} 0 & \text{if } q\mathcal{D}_{max} < \beta, \\ \frac{(q\mathcal{D}_{max}-\beta)^2}{2 s _x} & \text{if } \beta \leq q\mathcal{D}_{max} < \beta + \min\{ r _x, s _x\}, \\ \frac{2 r _x(q\mathcal{D}_{max}-\beta)- r _x^2}{2 s _x} & \text{if } r _x + \beta \leq q\mathcal{D}_{max} < s _x + \beta, \\ q\mathcal{D}_{max} - \beta - \frac{ s _x}{2} & \text{if } s _x + \beta \leq q\mathcal{D}_{max} < r _x + \beta, \\ r _x - \frac{(\max\{ r _x + s _x + \beta - q\mathcal{D}_{max}, 0\})^2}{2 s _x} & \text{if } \max\{ r _x, s _x\} + \beta \leq q\mathcal{D}_{max}. \end{cases}$
Intersected	$\begin{cases} q\mathcal{D}_{max} + \frac{q\mathcal{D}_{max}(q\mathcal{D}_{max}-2\delta)}{2 s _x} & \text{if } q\mathcal{D}_{max} < \min\{\alpha, \delta\}, \\ \frac{2 r _x q\mathcal{D}_{max} - \alpha^2}{2 s _x} & \text{if } \alpha \leq q\mathcal{D}_{max} < \delta, \\ q\mathcal{D}_{max} - \frac{\delta^2 + (\max\{q\mathcal{D}_{max} - \alpha, 0\})^2}{2 s _x} & \text{if } \delta \leq q\mathcal{D}_{max} < s _x + \alpha, \\ \alpha + \frac{ s _x^2 - \delta^2}{2 s _x} & \text{if } s _x + \alpha \leq q\mathcal{D}_{max}. \end{cases}$
Contained	$\begin{cases} q\mathcal{D}_{max} & \text{if } q\mathcal{D}_{max} < \alpha, \\ q\mathcal{D}_{max} - \frac{(q\mathcal{D}_{max}-\alpha)^2}{2 s _x} & \text{if } \alpha \leq q\mathcal{D}_{max} < s _x + \alpha, \\ \alpha + \frac{ s _x}{2} & \text{if } s _x + \alpha \leq q\mathcal{D}_{max}. \end{cases}$
r and s	The second integral term of Equation (2)
Separated	0
Intersected	$\begin{cases} q\mathcal{D}_{max} - \frac{q\mathcal{D}_{max}(2\alpha - q\mathcal{D}_{max})}{2 r _x} & \text{if } q\mathcal{D}_{max} < r _x - \alpha, \\ \frac{(r _x - \alpha)^2}{2 r _x} & \text{if } r _x - \alpha \leq q\mathcal{D}_{max}. \end{cases}$
Contained	$\begin{cases} \frac{q\mathcal{D}_{max} s _x}{ r _x} & \text{if } q\mathcal{D}_{max} < \delta, \\ \frac{2q\mathcal{D}_{max} s _x - (q\mathcal{D}_{max}-\delta)^2}{2 r _x} & \text{if } \delta \leq q\mathcal{D}_{max} < s _x + \delta, \\ \frac{ s _x(s _x + 2\delta)}{2 r _x} & \text{if } s _x + \delta \leq q\mathcal{D}_{max}. \end{cases}$

The values of α , β , and δ are determined by the relative positions of r and s as illustrated in Fig. 8.




Fig. 8. Spatial relationships between nodes r and s and their projected intervals. (a) Separated. (b) Intersected. (c) Contained.

while the sweeping direction selection does so by reducing the $q\mathcal{D}_{max}$ value more rapidly.

3.4 Maximum Distance as a Secondary Priority

The main queue maintains node pairs generated by node expansion in an increasing order of their distances. An issue we have not addressed is how we order node pairs of an equal distance in the main queue. This is a nontrivial issue particularly because a pair of intersected nodes are considered to have zero distance between them. Since there may be many pairs of intersected nodes in the main queue,

the performance impact by the way of breaking ties is potentially high.

Hjaltason and Samet used the depths of R-tree nodes to break tie [16]. For two given pairs of nodes of an equal distance, they proposed to give preference to a pair that contains a node with the maximum tree depth among the four tree nodes. This approach may assist their distance join algorithms in getting to leaf nodes and data objects as quickly as possible. However, it is not always beneficial to process node pairs in depth-first order because it does not always accelerate the reduction of $q\mathcal{D}_{max}$ value.




Fig. 9. Breaking ties for node pairs of an equal distance.

Consider two node pairs A and B in Fig. 9 as an example. If we follow the depth-first approach, node pair B will be placed before node pair A in the main queue because the former contains a node whose depth is deeper than those of the nodes in the latter. However, it is very likely that node pair A contains more pairs of entries with shorter distances than node pair B does because the maximum distance of node pair A is shorter than that of node pair B. Based on this observation, we propose to use the maximum distance of a node pair to break ties. By choosing the maximum distance of a node pair as a secondary priority of heap (i.e., the main queue), node pairs with shorter distances can be processed in the earlier stage of distance join. Consequently, qD_{max} value can be reduced more rapidly and the number of distance computations and queue insertions can be reduced.

4 ADAPTIVE MULTISTAGE k -distance JOIN

In the β -KDJ algorithm, qD_{max} value is initially set to an infinity and becomes smaller as the algorithm proceeds. The adaptation of the qD_{max} value has a crucial impact on the performance of the β -KDJ algorithm, as qD_{max} is used as a cutoff to prevent pairs of distant nodes from entering the main queue. If the qD_{max} value approaches to the real D_{max} value slowly, the early stage of the β -KDJ algorithm will be delayed considerably for handling too many pairs of distant nodes. Consequently, at the end of the algorithm processing, the main queue may end up with a large number of distant pairs whose insertions to the main queue were not necessary. The performance effect of *slow start* is more pronounced for a larger k , as the main queue and distance queue tend to grow large for a large k and, thereby, increasing the qD_{max} value. From our experiments with k as high as 100,000, we observed that more than 90 percent of execution time of k -distance join algorithms was spent to produce the first one percent (i.e., 1,000 pairs) of final query results.

In this section, we propose new adaptive multistage distance join algorithms \mathcal{AM} -KDJ and \mathcal{AM} -IDJ that mitigate the slow start problem by *aggressive pruning* and *compensation*.

4.1 Adaptive Multistage k -distance Join

The slow start problem is essentially caused by a pruning strategy using qD_{max} , whose value is dynamically updated as tree indexes are traversed and, therefore, not under direct control of the distance join algorithms. To circumvent this problem, we introduce a new pruning measure eD_{max} , which is an estimated D_{max} value for a given k . The eD_{max} value is set to an initial estimation at the beginning and

adaptively corrected during the algorithm processing. We will discuss techniques for initial estimation and adaptive correction in Section 4.3.

The \mathcal{AM} -KDJ algorithm is similar to the β -KDJ algorithm in that both the algorithms use a bidirectional node expansion. However, unlike the single-stage β -KDJ algorithm, where only qD_{max} is used for pruning, both qD_{max} and eD_{max} are used as cutoff values for pruning distant pairs in the \mathcal{AM} -KDJ algorithm. Specifically, in the *aggressive pruning* stage (described in Algorithm 2 (Fig. 10)),

- eD_{max} is used for pruning based on *axis distances* for aggressive pruning and thereby limiting the size of main and distance queues (line 22).
- qD_{max} is used for further pruning on *real distances* for nodes whose axis distances are within eD_{max} in the same way as β -KDJ.

With a properly estimated eD_{max} value, the \mathcal{AM} -KDJ algorithm can prune a large number of distant pairs in the first stage and avoid a significant portion of delay due to the slow start. However, if the \mathcal{AM} -KDJ algorithm becomes too aggressive by choosing an underestimated eD_{max} value, even close enough pairs may be discarded incorrectly. To avoid any false dismissals, we introduce another queue called *compensation queue* (\mathcal{Q}_C). The compensation queue stores every nonobject node pair retrieved from the main queue (line 11), except those for whom all entries have been examined. Note that qD_{max} , but not eD_{max} , is used for nodes whose axis distances are within eD_{max} . If eD_{max} values are used instead, this algorithm does not guarantee the correctness due to potential false dismissals. Using qD_{max} values also makes the performance of \mathcal{AM} -KDJ fairly insensitive to estimated eD_{max} values.

For example, in Fig. 11 (drawn from Fig. 5), an anchor node r_1 is paired up with nodes s_1 and s_2 , but not with s_3 and s_4 in the aggressive pruning stage because only s_1 and s_2 are within eD_{max} from the anchor node r_1 . Thus, the \mathcal{AM} -KDJ algorithm inserts only two pairs ($\langle r_1, s_1 \rangle$, $\langle r_1, s_2 \rangle$) into a main queue, instead of all four pairs ($\langle r_1, s_1 \rangle$, $\langle r_1, s_2 \rangle$, $\langle r_1, s_3 \rangle$, $\langle r_1, s_4 \rangle$) that would be enqueued by the β -KDJ algorithm. Then, the pair $\langle r, s \rangle$ currently being expanded is inserted into a compensation queue.

The aggressive pruning stage ends when one of the following conditions is satisfied: 1) the main queue becomes empty (line 5), 2) k or more query results have been returned (line 5), or 3) the distance of a node pair retrieved from the main queue becomes greater than eD_{max} (line 9). When condition 2 is satisfied, obviously it is not necessary to execute the compensation stage of the \mathcal{AM} -KDJ algorithm. (An overestimated eD_{max} can also be detected by comparing with qD_{max} value (line 8). In this case, instead of terminating the first stage, \mathcal{AM} -KDJ behaves exactly the same as the β -KDJ algorithm by using qD_{max} alone as a cutoff value.) When condition 3 is satisfied, eD_{max} must have been underestimated because all the object pairs returned after this point will have a greater distance than eD_{max} . Since an object pair with the k th largest distance has not been obtained by the time when the aggressive pruning stage comes to an end, the compensation stage (described in Algorithm 3 (Fig. 12)) begins its processing by inserting all the pairs stored in the compensation queue to the main queue.

```

1: set  $AnswerSet \leftarrow$  an empty set;
2: set  $Q_M, Q_D, Q_C \leftarrow$  empty main, distance and compensation queues ;
3: set  $eD_{max} \leftarrow$  an initial estimated  $D_{max}$ ;
4: insert a pair  $\langle R.root, S.root \rangle$  to the main queue  $Q_M$ ;
5: while  $|AnswerSet| < k$  and  $Q_M \neq \emptyset$  do
6:   set  $c \leftarrow$  dequeue( $Q_M$ );
7:   if  $c$  is an  $\langle object, object \rangle$  then  $AnswerSet \leftarrow \{c\} \cup AnswerSet$ ;
8:   else
9:     if  $qD_{max} \leq eD_{max}$  then  $eD_{max} \leftarrow qD_{max}$ ;           // overestimated  $eD_{max}$ 
10:    if  $c.distance > eD_{max}$  then
11:      reinsert  $c$  back into  $Q_M$ ;
12:      break;           // Terminate the Aggressive Pruning stage.
13:    end
14:    AggressivePlaneSweep( $c$ );
15:    enqueue( $Q_C, c$ );
16:  end
17: end
18: if  $|AnswerSet| < k$  then execute Algorithm 3;
19: procedure AggressivePlaneSweep( $\langle l, r \rangle$ )
20:   set  $L \leftarrow$  sort_axis( $\{entries\ of\ l\}$ );           // Sort the entries of  $l$  by axis values.
21:   set  $R \leftarrow$  sort_axis( $\{entries\ of\ r\}$ );           // Sort the entries of  $r$  by axis values.
22:   while  $L \neq \emptyset$  and  $R \neq \emptyset$  do
23:      $n \leftarrow$  a node with the min axis value  $\in L \cup R$ ;           //  $n$  becomes an anchor.
24:     if  $n \in L$  then
25:        $L \leftarrow L - \{n\}$ ; AggressiveSweepPruning( $n, R$ );
26:        $n.compensate \leftarrow$  a node in  $R$  with the min axis value and not yet paired with  $n$ ;
27:     else
28:        $R \leftarrow R - \{n\}$ ; AggressiveSweepPruning( $n, L$ );
29:        $n.compensate \leftarrow$  a node in  $L$  with the min axis value and not yet paired with  $n$ ;
30:     end
31:   end
32: end
33: procedure AggressiveSweepPruning( $n, List$ )
34:   Same as the SweepPruning procedure in Algorithm 1
35:   except line 16 replaced with the following:
36: 22: if  $axis\_distance(n, m) > eD_{max}$  then return;

```

Fig. 10. Algorithm 2 (\mathcal{AM} -KDJ): adaptive multistage k -distance join algorithm (aggressive pruning).

In the compensation stage, the pairs in the main queue are processed in a similar way as the \mathcal{B} -KDJ algorithm, but there are two notable differences from the \mathcal{B} -KDJ algorithm. First, the entries are not sorted again, if they have

already been sorted in the first stage. Second, for the pairs already expanded once in the first stage, only child pairs not examined in the first stage are processed by plane sweeping. This is feasible by bookkeeping done in the first stage (lines 19 and 21), which stores the information in an additional field ($n.compensate$) attached to a pair being inserted into the compensation queue. For these reasons, the cost of the compensation stage is not considerable compared with the cost of restarting the algorithm. In summary, the \mathcal{AM} -KDJ algorithm uses eD_{max} to avoid the slow start problem in the aggressive pruning stage and speeds up the query processing.

4.2 Adaptive Multistage Incremental Distance Join

Consider online query processing and Internet database search environments, where users interact with database systems in a way the number of required matches can be determined interactively or changed at any point of query processing. Also, consider a complex query that pipelines the results from a spatial distance join to a filter stage. Under these circumstances, the number of pairs (k) that should be returned from a distance join is not known a priori and, hence, a k -distance join algorithm proposed in [16] and the \mathcal{B} -KDJ algorithm presented in Section 3 cannot be used directly.

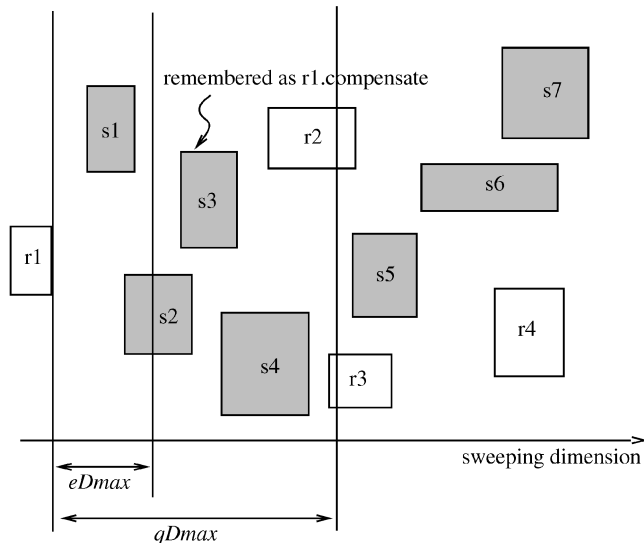


Fig. 11. Aggressive pruning with qD_{max} and eD_{max} .

```

1: insert all elements in  $Q_C$  into  $Q_M$ ;
2: while  $|AnswerSet| < k$  and  $Q_M \neq \emptyset$  do
3:   set  $c \leftarrow \text{dequeue}(Q_M)$ ;
4:   if  $c$  is an  $\langle \text{object}, \text{object} \rangle$  then  $AnswerSet \leftarrow \{c\} \cup AnswerSet$ ;
5:   else  $\text{CompensatePlaneSweep}(c)$ ;
   end
   procedure  $\text{CompensatePlaneSweep}(\langle l, r \rangle)$ 
6:    $L \leftarrow \{ \text{entries of } l \text{ sorted in Stage One} \}$ ;           //  $\{L[1], L[2], \dots, L[|L|]\}$ 
7:    $R \leftarrow \{ \text{entries of } r \text{ sorted in Stage One} \}$ ;       //  $\{R[1], R[2], \dots, R[|R|]\}$ 
8:   while  $L \neq \emptyset$  and  $R \neq \emptyset$  do
9:      $n \leftarrow$  a node with the min axis value  $\in L \cup R$ ;       //  $n$  becomes an anchor.
10:    if  $n \in L$  then
11:       $L \leftarrow L - \{n\}$ ;  $R' \leftarrow \{ \text{node list in } R \text{ not paired with } n \text{ in the Stage One} \}$ ;
12:      //  $\{R[n.\text{compensate}], R[n.\text{compensate} + 1], \dots, R[|R|]\}$ 
13:       $\text{SweepPruning}(n, R')$ ;
14:    else
15:       $R \leftarrow R - \{n\}$ ;  $L' \leftarrow \{ \text{node list in } L \text{ not paired with } n \text{ in the Stage One} \}$ ;
16:      //  $\{L[n.\text{compensate}], L[n.\text{compensate} + 1], \dots, L[|L|]\}$ 
17:       $\text{SweepPruning}(n, L')$ ;
18:    end
19:  end

```

Fig. 12. Algorithm 3 ($\mathcal{AM}\text{-KDJ}$): adaptive multistage k-distance join algorithm (compensation stage).

An important advantage of the $\mathcal{AM}\text{-KDJ}$ algorithm proposed in the previous section is that the $\mathcal{AM}\text{-KDJ}$ algorithm can be extended to an incremental algorithm (we call $\mathcal{AM}\text{-KDJ}$) to support the interactive applications described above. The main difference between the $\mathcal{AM}\text{-KDJ}$ and $\mathcal{AM}\text{-IDJ}$ algorithms is that $\mathcal{AM}\text{-IDJ}$ does not maintain a distance queue. Thus, the $\mathcal{AM}\text{-IDJ}$ algorithm uses eD_{max} alone as a cutoff value for pruning distant pairs because qD_{max} would be drawn only from a distance queue.

Without qD_{max} , $\mathcal{AM}\text{-IDJ}$ works as a stepwise incremental algorithm. First, $\mathcal{AM}\text{-IDJ}$ starts by determining an initial value k_1 and estimating an initial eD_{max1} for k_1 . Then, it performs the same way as the first stage of the $\mathcal{AM}\text{-KDJ}$ algorithm without qD_{max} . However, the first stage may terminate before producing enough object pairs (i.e., less than k_1) if eD_{max} is underestimated. If that happens, the $\mathcal{AM}\text{-IDJ}$ algorithm estimates eD_{max2} value for k_2 ($k_2 > k_1$) and initiates a compensation stage.

Even when a sufficient number of object pairs have been returned from the first stage, users may request more answers. Then, $\mathcal{AM}\text{-IDJ}$ initiates a compensation stage by determining k_2 and estimating a new eD_{max2} accordingly. As shown in Fig. 13 (drawn from Fig. 5), the compensation stage can initiate another compensation stage at the end of its processing, by choosing k_3 and eD_{max3} . This process continues until users stop requesting more answers. In this way, the $\mathcal{AM}\text{-IDJ}$ algorithm can be used to produce query results incrementally without limiting the maximum number of pairs in advance. Except the first stage of the $\mathcal{AM}\text{-IDJ}$ algorithm where the $\text{AggressivePlaneSweep}$ procedure (in Algorithm 2) is used, the $\text{CompensatePlaneSweep}$ procedure (in Algorithm 3) is used to prune distant pairs in the rest of the compensation stages.

4.3 Estimating the Maximum Distance (eD_{max})

Both the $\mathcal{AM}\text{-KDJ}$ and $\mathcal{AM}\text{-IDJ}$ algorithms process a distance join query based on an estimated cutoff value

eD_{max} . Thus, there should be a way to obtain an initial estimate and correct the estimate adaptively as the algorithms proceed. Assuming data sets are uniformly distributed, we provide mechanisms to choose an initial estimate of eD_{max} and to adaptively correct it.

If the distribution of a data set is skewed, then a larger number of close pairs can be found in a smaller dense region of the data space. We expect that the formulae given in this section tend to overestimate eD_{max} value for nonuniformly distributed data sets, especially when a stopping cardinality k is far smaller than the number of all pairs of objects (i.e., $k \ll |R| \times |S|$). This was corroborated by our experiments as described in Section 5.4.

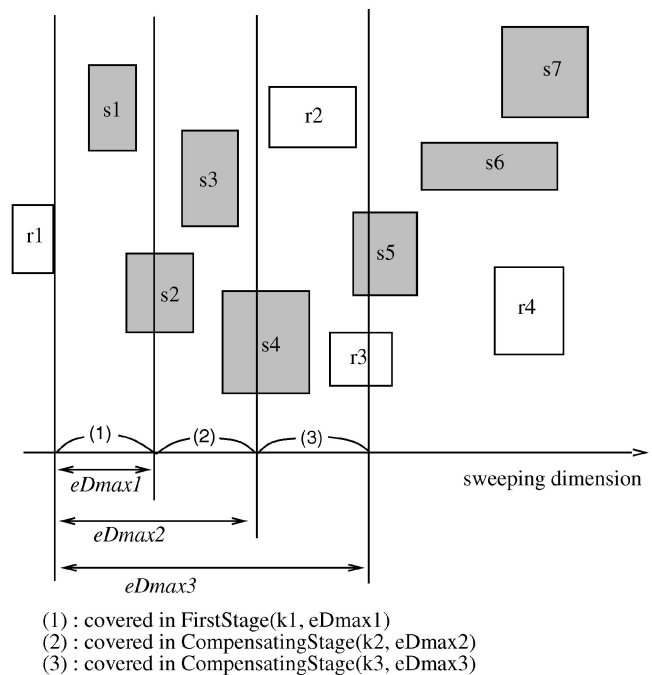


Fig. 13. Step-wise incremental distance join.

4.3.1 Initial Estimation

Let $|R|$ and $|S|$ be the number of data objects in MBRs R and S , respectively. Suppose that most regions of R and S overlap. Then, for a data object r in R contained in the region shared by R and S , the expected number of objects in S within distance d from r is approximated by $|S| \times \frac{\pi \times d^2}{\text{area}(R \cap S)}$, assuming the circle centered at r of radius d is fully contained in the shared region (i.e., $R \cap S$). Thus, by considering all data objects in R , the total number of object pairs within distance d can be approximated by $|R| \times |S| \times \frac{\pi \times d^2}{\text{area}(R \cap S)}$.

When the target number of object pairs, k , is given with a query, we can obtain the initial estimation of \mathcal{D}_{max} by setting k to the above formula as follows:

$$k = |R| \times |S| \times \frac{\pi \times d^2}{\text{area}(R \cap S)},$$

and then by replacing d with $e\mathcal{D}_{max}$. Therefore, for a given stopping cardinality k , the initial estimation of \mathcal{D}_{max} can be obtained by the following equation.

$$e\mathcal{D}_{max} = \sqrt{k \times \rho} \quad \left(\text{where } \rho = \frac{\text{area}(R \cap S)}{\pi \times |R| \times |S|} \right). \quad (3)$$

Evidently, this equation can be applied only when R and S overlap. Nonetheless, it is unlikely this will be a serious limitation because overlapping node pairs always come before nonoverlapping pairs in the main queue \mathcal{Q}_M . For two sets of data objects to be joined by distance, the root nodes of two corresponding R-trees are commonly expected to overlap each other under most practical circumstances. We can then make an initial estimation of \mathcal{D}_{max} from the pair of root nodes.

4.3.2 Adaptive Correction of Estimated Distance $e\mathcal{D}_{max}$

The performance of the $\mathcal{AM}\text{-KDJ}$ and $\mathcal{AM}\text{-IDJ}$ algorithms can be further improved by adaptively adjusting the value of $e\mathcal{D}_{max}$ at runtime. Adaptive correction of $e\mathcal{D}_{max}$ can be done at any point of query processing by estimating a new $e\mathcal{D}_{max}$ from the number of object pairs k_0 ($k_0 < k$) obtained up to the point and the real distance of the k_0 th object pair, $\mathcal{D}_{max(k_0)}$. Specifically, the new estimate $e\mathcal{D}_{max}'$ can be computed from (3) as

$$e\mathcal{D}_{max}' = \sqrt{\mathcal{D}_{max(k_0)}^2 + (k - k_0)\rho} \quad (4)$$

by arithmetic correction, or as

$$e\mathcal{D}_{max}' = \mathcal{D}_{max(k_0)} \times \sqrt{k/k_0} \quad (5)$$

by geometric correction if $\mathcal{D}_{max(k_0)} \neq 0$. In practice, we propose computing $e\mathcal{D}_{max}'$ in both ways, and then choose the minimum if the query processing should be on the aggressive side. Otherwise, the maximum is chosen as $e\mathcal{D}_{max}'$.

Note that the new estimate $e\mathcal{D}_{max}'$ can sometimes grow beyond the previous estimate. If this happens, some pairs whose distances are larger than the previous estimate but smaller than the new estimate could have already been pruned and will never be examined in the current processing stage under the new estimate. Thus, to guarantee

the correctness of the distance join, the algorithm should initiate a compensation stage, as soon as a pair whose distance is smaller than the smallest $e\mathcal{D}_{max}$ is dequeued from the main queue.

4.4 Queue Management

Efficient queue management is one of the key components of the distance join algorithms proposed in this paper. Each of the $\mathcal{B}\text{-KDJ}$, $\mathcal{AM}\text{-KDJ}$, and $\mathcal{AM}\text{-IDJ}$ algorithms relies on the use of one or more priority queues for query processing. In particular, the main queue (\mathcal{Q}_M) is heavily used by all of the proposed algorithms, and its performance impact is significant. In the worst case, the main queue can grow as large as the product of *all* objects of two R-tree indexes. That is, the size of \mathcal{Q}_M is in $\mathcal{O}(|R_{obj}| \times |S_{obj}|)$, where $|R_{obj}|$ and $|S_{obj}|$ are the number of all objects in R and S , respectively. Thus, it is not always feasible to store the main queue in memory.

It was reported in [16] that a simple memory-based implementation might slow down query processing severely, due to excessive virtual memory thrashing. A hybrid memory/disk scheme [16] and a technique based on range partitioning [10] have been proposed to improve queue management and to avoid wasted sorting I/O operations. We adopt a similar scheme for queue management, which partitions a queue by range based on distances of pairs. A partition in the shortest distance range is kept in memory as a heap structure, while the rest of partitions are stored on disk as merely unsorted piles.

When the in-memory heap becomes full, it is *split* into two parts, and then one in the longer distance range is moved to disk as a new segment. When the in-memory heap becomes empty, a disk-resident segment in the shortest distance range or a part of the segment is *swapped in* to memory to fill up the in-memory heap. Each of the **split** and **swap-in** operations requires $\mathcal{O}n \log n$ computational cost for a heap of n elements as well as I/O cost for reading and writing a segment. Thus, it is important to minimize the required number of those operations, which largely depends on the partition boundary values between the in-memory heap and the first disk-resident segment and between those consecutive segments. However, as it is impossible to predict an exact \mathcal{D}_{max} value for a given k , so is it difficult to determine optimal distance values as segment boundaries.

To address this issue, we use (3) to determine the boundary distance values. Suppose n is the number of elements that can be stored in an in-memory heap. Then, the boundary value between the in-memory heap and the first disk-resident segment is given by $\sqrt{n \times \rho}$, and the boundary value between the first and second segments is given by $\sqrt{(2 \times n) \times \rho}$, and so on.

In addition to a main queue, multistage algorithms $\mathcal{AM}\text{-KDJ}$ and $\mathcal{AM}\text{-IDJ}$ use a compensation queue (\mathcal{Q}_C) in the compensation stage. Unlike the main queue, a compensation queue does not store any pair of objects. In other words, a compensation queue can store pairs of nonobject R-tree nodes only. Thus, the size of \mathcal{Q}_C is in $\mathcal{O}(|R_{node}| \times |S_{node}|)$, where $|R_{node}|$ and $|S_{node}|$ are the number

of nodes (both internal and leaf nodes) in R and S , respectively. This is a significantly lower upper-bound than a main queue has. We also observed from our experiments that compensation queues were several orders of magnitude smaller than main queues. As for a distance queue used by \mathcal{B} -KDJ and \mathcal{AM} -KDJ algorithms, its size is always bounded by a given k value. For these reasons, under most circumstances, we assume either a compensation queue or a distance queue fits in memory. If any of these queues outgrows memory, the same partitioning technique used for a main queue will be applied.

5 PERFORMANCE EVALUATION

In this section, we evaluate the proposed algorithms empirically and compare with previous work. In particular, the proposed \mathcal{B} -KDJ, \mathcal{AM} -KDJ, and \mathcal{AM} -IDJ algorithms were compared with Hjaltason and Samet's k -distance and incremental distance join algorithms (hereinafter denoted as HS -KDJ and HS -IDJ, respectively) for k -distance join (KDJ) and incremental distance join (IDJ) queries. We also include the performance of an R-tree-based spatial join algorithm [8] combined with a sort operation (denoted as SJ -SORT) in most of the experiments. For each distance join query, a spatial join operation was performed with a real \mathcal{D}_{max} value to generate the k -nearest pairs. Then, an external sort operation was performed to return the query results in an increasing order of distances. Note that SJ -SORT cannot be applied without knowing a real \mathcal{D}_{max} value, and we made a favorable assumption for SJ -SORT that the real \mathcal{D}_{max} value was known to SJ -SORT a priori. Thus, we conjecture that SJ -SORT followed by an external sort yields the best known lower bound performance for distance join processing.

5.1 Experimental Settings

Experiments were performed on a Sun Ultrasparc-II workstation running on Solaris 2.7. This workstation has 256 MBytes of memory and 9 GBytes of disk storage (Seagate ST39140A) with Ultra 10 EIDE interface. The disk is locally attached to the workstation and used to store databases, queues, and any temporary results. We used the direct I/O feature of Solaris for all the experiments to avoid operating system's cache effects, and the average disk access bandwidth was about 0.5 MBytes/sec for random accesses and about 5 MBytes/sec for sequential accesses.

Data sets. To evaluate distance join algorithms, we used real-world data sets in TIGER/Line97 from the US Bureau of Census [20]. The particular data sets we used were 633,461 streets and 189,642 hydrographic objects from the Arizona state. Throughout the entire set of experiments, the same page size of 4 KBytes was used for disk I/O and R*-tree [3] nodes.

Metrics. We measured the performance of various algorithms based on the following metrics to compare the algorithms in different aspects such as computational cost and I/O cost.

1. *Number of distance computations:* The cost of computing distances between pairs of nodes (or objects) constitutes a significant portion of the computational

cost of a distance join operation. Thus, the total number of distance computations required by a distance join algorithm provides a direct indication of its computational performance.

2. *Number of queue insertions:* The task of managing a main queue is largely I/O intensive as well as CPU intensive. Inserting a node pair into the in-memory portion of the queue is CPU intensive, while inserting into the disk resident portion is I/O intensive. We measured the CPU and I/O cost separately for the two different queue insertions.
3. *Number of R-tree node accesses:* The number of R-tree nodes accessed during distance join processing is another I/O intensive metric. We measured actual number of nodes fetched from disk with varying R-tree buffer sizes.
4. *Response time:* Actual query response times were measured for overall performance of distance join algorithms. CPU and I/O costs were considered separately in measuring the response times.

5.2 Evaluation of k -Distance Joins

In this set of experiments, we varied a stopping cardinality k from 10 to 100,000 to compare the performance of HS -KDJ, \mathcal{B} -KDJ, and \mathcal{AM} -KDJ algorithms. The sizes of in-memory portion of a main queue and R-tree buffer were fixed to 512 KBytes each. For the \mathcal{AM} -KDJ algorithm, we used (3) to estimate $e\mathcal{D}_{max}$ values, and we observed a tendency for $e\mathcal{D}_{max}$ values to be overestimated with respect to real \mathcal{D}_{max} values. For example, for $k = 100,000$, $e\mathcal{D}_{max}$ was about 2.3 times larger than a real \mathcal{D}_{max} .

Fig. 14a shows that both \mathcal{B} -KDJ and \mathcal{AM} -KDJ reduced the number of distance computations significantly. The numbers of distance computations required by the two algorithms were smaller than those required by the HS -KDJ algorithm by up to two orders of magnitude. \mathcal{AM} -KDJ was almost identical to SJ -SORT by this metric. This demonstrates that the optimized plane-sweep method was very effective in pruning pairs generated by bidirectional expansions. On the other hand, the HS -KDJ algorithm examines all possible pairs exhaustively in unidirectional expansions.

In Fig. 14b, both \mathcal{B} -KDJ and \mathcal{AM} -KDJ achieved significant reductions in queue insertions for all k values. \mathcal{AM} -KDJ was always better than \mathcal{B} -KDJ particularly for large k values. This result confirms our conjecture that the optimized plane-sweep method can prevent an explosion of a main queue that would be caused by bidirectional node expansions without the optimized plane-sweep.

Fig. 14c shows the number of R-tree nodes fetched from disk for distance join processing by each algorithm. For large k values, the proposed \mathcal{B} -KDJ and \mathcal{AM} -KDJ algorithms required a far smaller number of R-tree node accesses than HS -KDJ algorithm. For small k values, on the other hand, the HS -KDJ algorithm was slightly better than the other algorithms, due to its more localized node access patterns for small k . Table 2 compares the number of R-tree nodes that would be fetched from disk with R-tree buffer size set to zero. Apparently, the bidirectional node expansion used by the \mathcal{B} -KDJ and \mathcal{AM} -KDJ algorithms

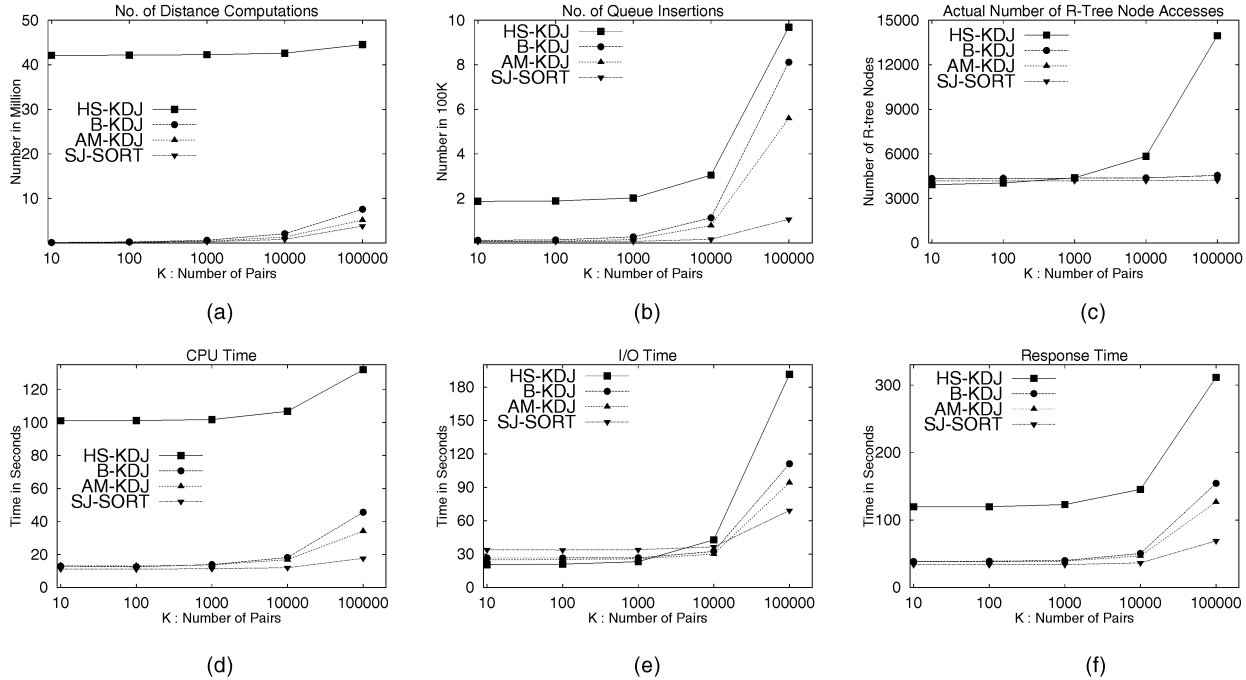


Fig. 14. Performance of k -distance joins. (a) Distance computations. (b) Queue insertions. (c) R-tree node accesses. (d) CPU time. (e) I/O time. (f) Response time.

requires much less number of R-tree node accesses than unidirectional node expansion used by the *HS-KDJ* algorithm. It should be noted that the number of R-tree node accesses for *B-KDJ*, *AM-KDJ*, and *SJ-SORT* algorithms are all identical in Table 2. This is because these algorithms use the same bidirectional node expansion and access the same collection of R-tree nodes, though they may traverse an R-tree index in different orders.

The total CPU time spent on executing each algorithm is shown in Fig. 14d. The *B-KDJ* and *AM-KDJ* algorithms consistently outperformed *HS-KDJ* up to an order of magnitude. This significant improvement in computational cost is due mainly to the reduced number of distance computations. Recall that the unidirectional expansion requires distance computations for an exhaustive set of node pairs, while bidirectional node expansion with plane sweeping requires distance computations only for node pairs whose axis distances are smaller than qD_{max} value at the top of the distance queue. Additionally, the proposed algorithms are further optimized by techniques for selecting

sweeping axis and direction and by using maximum distance as a secondary priority for the main queue.

The total I/O time shown in Fig. 14e reflects mostly the combined effects of queue insertions and R-tree node accesses in Fig. 14b and Fig. 14c, respectively. Fig. 14f shows the response time of each algorithm with the CPU and I/O times combined together. Both the *B-KDJ* and *AM-KDJ* algorithms outperformed the *HS-KDJ* algorithm by a factor of two or three in response times. *AM-KDJ* performed better than *B-KDJ* for large k values, demonstrating that *AM-KDJ* deals with the slow start problem better than *B-KDJ* does. For small k values, both *B-KDJ* and *AM-KDJ* were comparable with *SJ-SORT*. Even for large k values, the response time of *AM-KDJ* was within about 80 percent above that of *SJ-SORT*, which we conjecture yields the best known lower bound performance.

5.3 Impact of Optimized Plane-Sweep and Secondary Priority

We have proposed optimization techniques for *B-KDJ* in Section 3. One is for selecting sweeping axis and direction, which is mainly aimed at reducing the number of distance computations. The other is using the maximum distance between node pairs as a secondary priority of the main queues, which is mainly aimed at reducing the number of queue insertions. To further analyze the performance impacts of the optimization techniques, we measured the performance of *B-KDJ*.

1. with both optimizations turned on,
2. with the sweeping index only,
3. with the secondary priority only,
4. with both optimizations turned off.

TABLE 2
Number of R-Tree Node Accesses for k -distance Joins

KDJ Algorithms	Stopping cardinality k				
	10	100	1,000	10,000	100,000
<i>HS-KDJ</i>	186,184	186,403	186,801	188,354	197,113
<i>B-KDJ</i>	12,652	12,660	12,672	12,688	12,916
<i>AM-KDJ</i>	12,652	12,660	12,672	12,688	12,916
<i>SJ-SORT</i>	12,652	12,660	12,672	12,688	12,916

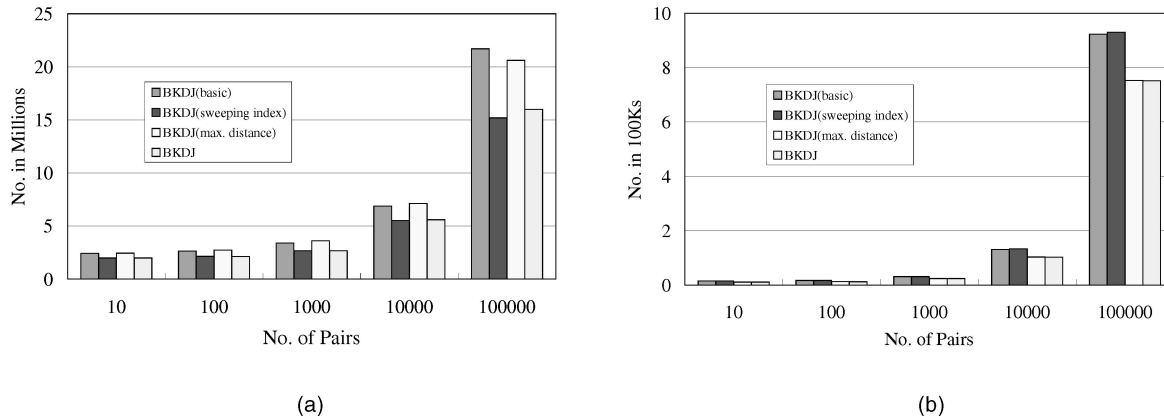


Fig. 15. Improvements by Optimized Plane Sweep for \mathcal{B} -KDJ. (a) Distance computations. (b) Queue insertions.

For the cases with sweeping index turned off, the sweeping index and direction were fixed to x -axis and forward direction.

The sweeping index method alone reduced the number of distance computations by up to 20 percent as shown in Fig. 15a. The use of the maximum distance as a secondary priority alone reduced the number of queue insertions by up to 15 percent as shown in Fig. 15b. The use of the maximum distance also helped decrease the $q\mathcal{D}_{max}$ value more quickly and reduce the number of distance computations slightly as shown in Fig. 15a. However, the synergistic effect of the two optimization techniques was rather insignificant. As they improve the performance of distance join processing largely independently in two different aspects, we recommend that both the optimization techniques be used together.

5.4 Evaluation of Incremental Distance Joins

As in the previous section, we varied a stopping cardinality k from 10 to 100,000 to compare the performance of incremental distance join algorithms HS -IDJ and AM -IDJ. Like the previous experiments for k -distance joins, the sizes of in-memory portion of a main queue and R-tree buffer were fixed to 512 KBytes.

In Figs. 16a and 16b, the AM -IDJ algorithm required 75 to 98 percent less distance computations and queue insertions than the HS -IDJ algorithm did. For large k values, as shown in Fig. 16c, the AM -IDJ algorithm required a much smaller number of disk accesses than the HS -IDJ algorithm. This is because AM -IDJ accesses R-tree nodes using bidirectional node expansion, in the same way as AM -KDJ does. The significant improvement in these three metrics in turn led to improvement in response time by an order of magnitude in Fig. 16f. Specifically, the improvement in CPU time (Fig. 16d) is attributed to the reduction in distance computations and queue insertions, and the improvement in I/O time (Fig. 16e) is attributed to the reduction in queue insertions and R-tree node accesses. Like AM -KDJ algorithm, (3) in Section 4.3.1 was used to estimate $e\mathcal{D}_{max}$ values for the AM -IDJ algorithm.

Now, it is well worthwhile investigating the performance impact of the stopping cardinality k . Generally, KDJ algorithms make use of the a priori knowledge of the

k value to minimize the distance computations and the number of queue insertions. Thus, KDJ algorithms are expected to be much faster than IDJ algorithms. From our experiments, however, HS -KDJ required almost as many distance computations as HS -IDJ did. This indicates that HS -KDJ does not take advantage of the stopping cardinality enough to achieve performance gain in the distance computations.

In contrast, AM -KDJ required only about 70 percent of distance computations that AM -IDJ did (Fig. 16a and Fig. 14a) and required only 8 percent of queue insertions that AM -IDJ did (Fig. 16b and Fig. 14b). This is because KDJ algorithms need not insert a node pair into main queue if its distance is greater than $q\mathcal{D}_{max}$ value. The number of queue insertions has direct impact on both CPU and I/O times. The response time of AM -KDJ algorithm was about 60 percent less than that of AM -IDJ algorithm (see Fig. 16f and Fig. 14f).

5.5 Impact of Memory Size

In this set of experiments, we examined the performance impact of memory constraint on queue management and R-tree access. The sizes of in-memory portion of a main queue and R-tree buffer were varied from 64 KBytes to 1,024 KBytes. We measured the response time of HS -KDJ, \mathcal{B} -KDJ, and AM -KDJ algorithms for a fixed stopping cardinality $k = 100,000$.

5.5.1 Buffer Size for Main Queue

No measurement for SJ -SORT algorithm appears in Figs. 17a through 17c because SJ -SORT algorithm need not use the main queue for distance join processing. As we expected, in Figs. 17a and 17b, the cost of queue management decreased in terms of both the number of required write operations and time spent on the write operations. More noticeable improvement was observed in handling the overflow and underflow of the in-memory portion of queue, by *split* and *swap-in* operations, respectively. (The *split* and *swap-in* operations are described in Section 4.4.) The time spent on the *split* and *swap-in* operations was improved substantially for all three algorithms in Fig. 17c.

It should be noted that the cost of queue management can be further reduced by not storing object pairs in the main queue, as proposed in the recent work by Corral et al.

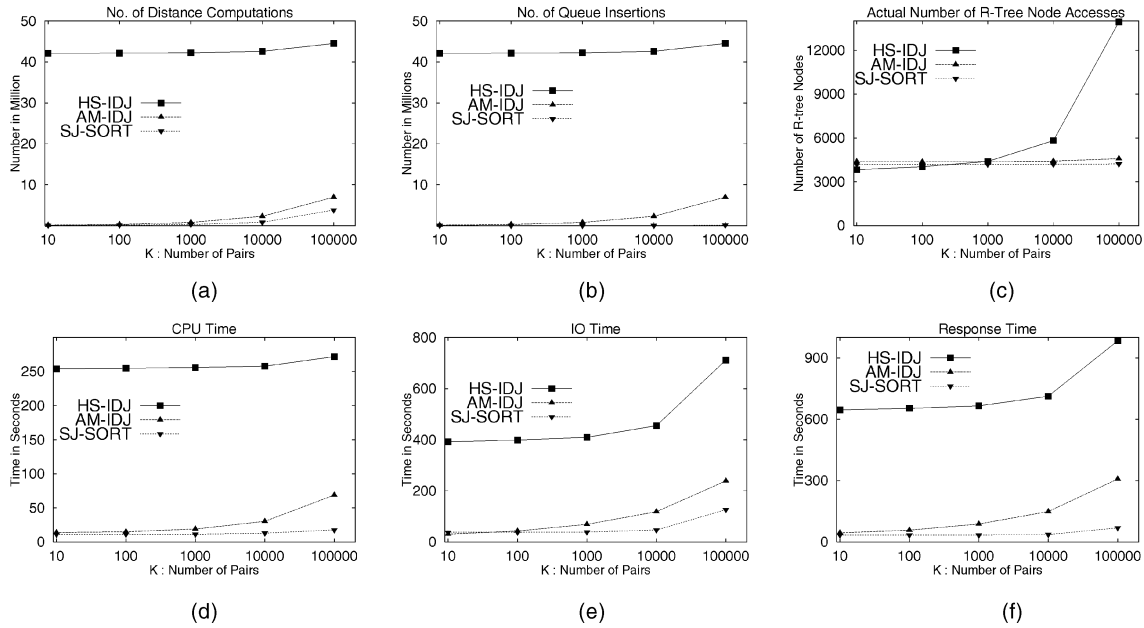


Fig. 16. Performance of incremental distance joins. (a) Distance computations. (b) Queue insertions. (c) R-tree node accesses. (d) CPU time. (e) I/O time. (f) Response time.

[11]. It is straightforward to modify the distance queue to store k object pairs for the \mathcal{B} -KDJ and \mathcal{AM} -KDJ algorithms. The reason we did not use the optimization was that the optimization cannot be applied to incremental distance join queries and it was desired to evaluate the performance of **KDJ** and **IDJ** algorithms on the same basis.

While the CPU time remained almost unchanged in Fig. 17d, the I/O time was improved with more memory for all the algorithms shown in Fig. 17e. The improved response time was mainly attributed to the improved

I/O time. The proposed \mathcal{B} -KDJ and \mathcal{AM} -KDJ algorithms showed consistently better performance in queue management than \mathcal{HS} -KDJ all over the examined range of memory size. This is because \mathcal{B} -KDJ and \mathcal{AM} -KDJ algorithms reduced the number of required queue insertions and queue write operations.

5.5.2 Buffer Size for R-Tree

As shown in Figs. 18a and 18b, a considerable amount of improvement in R-tree accesses was observed by increasing

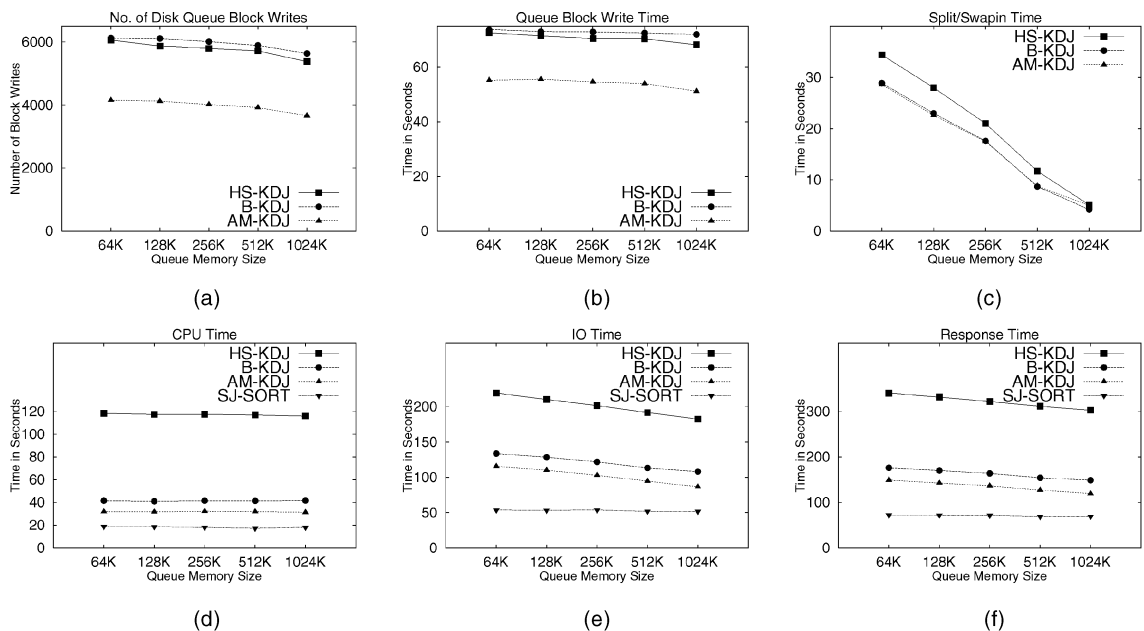


Fig. 17. Impact of queue buffer size. (a) Queue block writes. (b) Queue block write time. (c) Split/swapin time. (d) CPU time. (e) I/O time. (f) Response time.

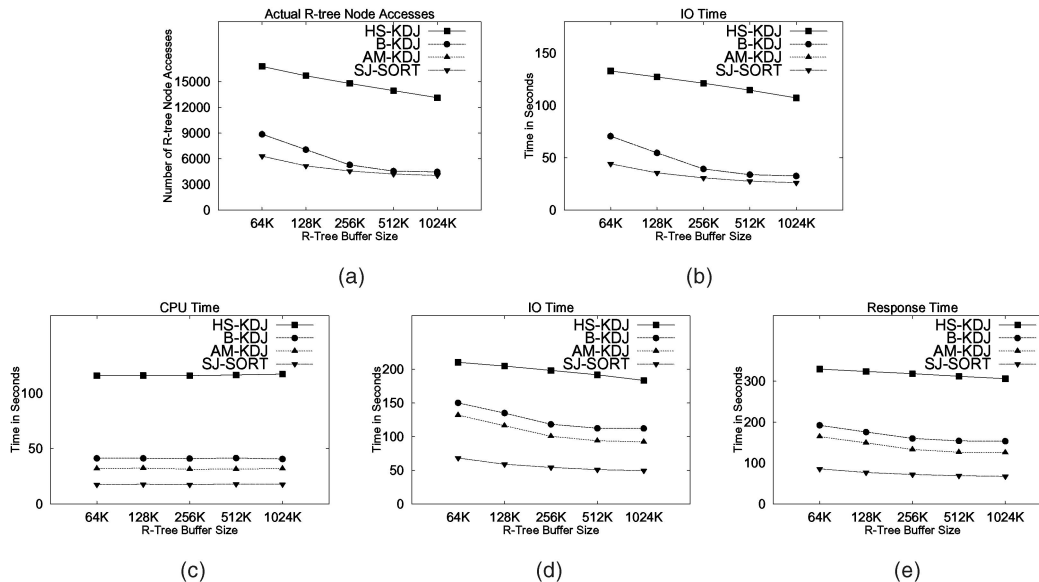


Fig. 18. Impact of R-tree buffer size. (a) R-tree node accesses. (b) R-tree node access time. (c) CPU time. (d) I/O time. (e) Response time.

the size of buffer for R-tree. For example, by increasing the buffer size from 64 KBytes to 1,024 KBytes, the R-tree access time was reduced by 46 percent for the \mathcal{B} -KDJ and \mathcal{AM} -KDJ algorithms. Recall that the \mathcal{B} -KDJ and \mathcal{AM} -KDJ algorithms, which are based on bidirectional node expansion, show the same behavior in R-tree access.

Like the queue management in the previous section, the CPU time spent on R-tree accesses remained almost unchanged, as shown in Fig. 18c. It was again the I/O time that affected the response time most in Figs. 18d and 18e.

5.6 Impact of $e\mathcal{D}_{max}$ Estimation on \mathcal{AM} -KDJ Performance

We designed two sets of experiments to characterize the performance of \mathcal{AM} -KDJ algorithm with respect to the accuracy of estimated $e\mathcal{D}_{max}$ values. In Section 5.6.1, instead of using (3) to estimate $e\mathcal{D}_{max}$, we varied the $e\mathcal{D}_{max}$ value from $0.1 \times \mathcal{D}_{max}$ to $10 \times \mathcal{D}_{max}$. Recall that \mathcal{D}_{max} is a real distance between the k th nearest pair of objects.

5.6.1 Robustness of \mathcal{AM} -KDJ

While fixing a stopping cardinality k to 100,000, we varied the $e\mathcal{D}_{max}$ value from $0.1 \times \mathcal{D}_{max}$ to $10 \times \mathcal{D}_{max}$. When $e\mathcal{D}_{max}$ is overestimated ($e\mathcal{D}_{max} > \mathcal{D}_{max}$), the compensation stage of \mathcal{AM} -KDJ algorithm is not necessary because all the k -nearest pairs will be produced in the first (aggressive pruning) stage. Even when $e\mathcal{D}_{max}$ is overestimated, \mathcal{AM} -KDJ guarantees that $e\mathcal{D}_{max}$ is always smaller than or equal to $q\mathcal{D}_{max}$ (obtained from a distance queue) throughout the first stage. Thus, \mathcal{AM} -KDJ always requires no more distance computation and queue insertion operations than the \mathcal{B} -KDJ algorithm does.

On the other hand, if $e\mathcal{D}_{max}$ is underestimated ($e\mathcal{D}_{max} < \mathcal{D}_{max}$), the node pairs in the compensation queue will be revisited in the compensation stage. Thus, the cost of tree traversals will increase, but it will be bounded by twice the cost of the \mathcal{B} -KDJ algorithm. Although there is no such a bound on the cost of queue management, we observed in most of our experiments that the cost of queue management

was lower than that of the \mathcal{B} -KDJ algorithm. This is because a large number of insertions to a compensation queue were prevented by aggressive pruning, and the compensation queue was several orders of magnitude smaller than the main queue. As discussed in Section 4.1, for a pair already expanded once in the first stage, only child pairs not examined in the first stage are paired up in the compensation stage and thereby wasting no time for redundant work. The value of $q\mathcal{D}_{max}$ is likely to have become quite close to a real \mathcal{D}_{max} value in the compensation stage. So, the \mathcal{AM} -KDJ algorithm usually prunes distant pairs much more efficiently in the compensation stage than the \mathcal{B} -KDJ algorithm would do in a single stage. Therefore, \mathcal{AM} -KDJ outperforms the k -distance join algorithms \mathcal{HS} -KDJ and \mathcal{B} -KDJ, despite the additional cost of compensation stage.

Fig. 19 shows that, as $e\mathcal{D}_{max}$ approaches to a real \mathcal{D}_{max} value, the performance of \mathcal{AM} -KDJ improves consistently in all three metrics. When $e\mathcal{D}_{max}$ increases far beyond the real \mathcal{D}_{max} value, the performance of \mathcal{AM} -KDJ converges to that of the \mathcal{B} -KDJ algorithm. More importantly, however, \mathcal{AM} -KDJ always outperformed \mathcal{B} -KDJ, not to mention \mathcal{HS} -KDJ, with $e\mathcal{D}_{max}$ in a wide spectrum of estimated value range.

We have not measured the cost of compensation queue management. A compensation queue contains pairs of nonobject R-tree nodes. During the first (aggressive pruning) stage of the \mathcal{AM} -KDJ algorithm, the number of pruned pairs is far larger than the number of nonobject pairs inserted into a compensation queue. In most of our experiments, the size of a compensation queue was less than 0.5 percent of the size of a main queue. Thus, the additional cost required for the compensation queue was almost negligible. This is one of the reasons why the \mathcal{AM} -KDJ algorithm always outperformed \mathcal{B} -KDJ, which does not need a compensation queue.

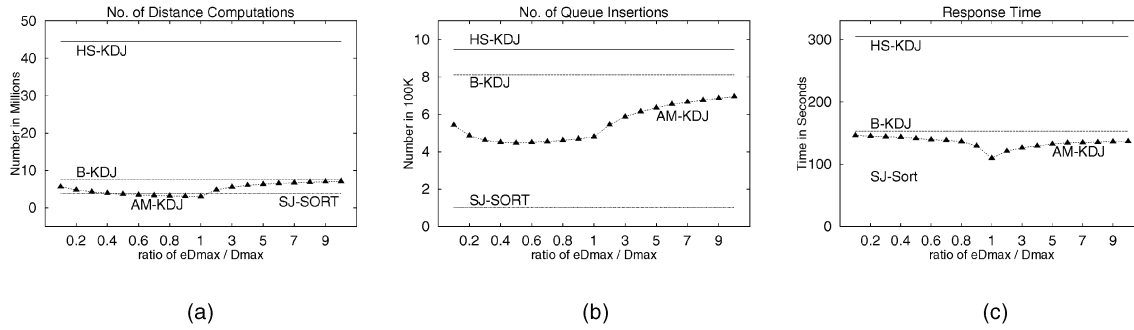


Fig. 19. Performance impact of eD_{max} . (a) Distance computations. (b) Queue insertions. (c) Response time.

5.7 Stepwise Incremental Execution of $\mathcal{AM-IDJ}$

Incremental distance join algorithms do not require a preset stopping cardinality k . Thus, in this set of experiments, we simulated a situation where users repeatedly requested a set of 10,000 nearest pairs at a time until a total of 100,000 nearest pairs were generated. Incremental algorithms $HS-IDJ$ and $AM-IDJ$ each were executed once in a single experiment run, until a total of 100,000 nearest pairs were generated. The sizes of in-memory portion of a main queue and R-tree buffer were fixed to 512 KBytes both for $HS-IDJ$ and $AM-IDJ$.

For $SJ-SORT$, which is not an incremental algorithm, we restarted its processing each time $i \times 10,000$ nearest pairs were generated for i ($1 \leq i \leq 9$). Thus, the performance measurements of $SJ-SORT$ presented in Fig. 20 are cumulative. For example, the response time of $SJ-SORT$ for $k = 20,000$ includes the times spent on executing $SJ-SORT$ twice, once for $k = 10,000$ and another for $k = 20,000$. For each run of $SJ-SORT$, we used a real D_{max} value for each of different stopping cardinalities.

In Fig. 20, we measured the response time of the $\mathcal{AM-IDJ}$ algorithm in two different ways: 1) with eD_{max} values estimated by (3) and 2) with real D_{max} values provided for 10 different k values. When estimated eD_{max} values were provided, $\mathcal{AM-IDJ}$ needed compensation processing only after generating 30,000 pairs and 90,000 pairs, due to overestimated eD_{max} values. In the second case (denoted by $AM-IDJ(D_{max})$ in Fig. 20), a real D_{max} value was provided for each k value from 10,000 through 100,000, to simulate a situation where the next set of 10,000 pairs of objects were repeatedly requested by a user. Consequently, $\mathcal{AM-IDJ}$ was forced to initiate a compensation stage, each time the next set was requested. This overhead slowed down the processing due mainly to redundant R-tree node accesses. Overall, $\mathcal{AM-IDJ}$ showed a fairly consistent performance over varying eD_{max} estimates, as $AM-KDJ$ did in Section 5.6. For all the k values, $\mathcal{AM-IDJ}$ with estimated eD_{max} improved the response time by a factor of two to four, when compared with $HS-IDJ$.

6 CONCLUSIONS

We have developed new distance join algorithms for spatial databases. The proposed algorithms provide significant performance improvement over previous work. The plane-sweep technique optimized by novel strategies for selecting a sweeping axis and direction minimizes the computational

overhead incurred by bidirectional node expansions. The node expansions are further optimized by using maximum distance for breaking tied pairs. We have shown that this optimized plane-sweep technique alone improves processing of a k -distance join query considerably.

The adaptive multistage algorithms employ aggressive pruning and compensation methods to further optimize the distance join processing. These algorithms address a slow start problem by using estimated maximum distances as cutoff values for pruning distant pairs. Our experimental study shows that the proposed algorithms outperformed previous work significantly and consistently for all the stopping cardinalities over a wide spectrum of estimated maximum distances. Ample evidence was observed that the adaptive algorithm yielded significant improvement in query processing time regardless of the techniques used for maximum distance estimations. For a relatively small stopping cardinality, the proposed algorithms achieved up to an order of magnitude improvement over previous work. Assuming data objects are uniformly distributed, we have developed strategies to choose an initial estimate and to correct the estimate adaptively during the query processing.

When the stopping cardinality of a distance join query is unknown (as in online query processing environments or a complex query that contains a distance join as a subquery), the adaptive multistage algorithms process the query in a stepwise manner so that the query results can be returned incrementally.

ACKNOWLEDGMENTS

This work was sponsored in part by the US National Science Foundation CAREER Award (IIS-9876037), US National Science Foundation Grant No. IIS-0100436, and Research Infrastructure program EIA-0080123. It was also





Fig. 20. Step-wise incremental execution.


supported by the Korean Science and Engineering Foundation under Exchange Student Program. The authors assume all responsibility for the contents of the paper.

REFERENCES

- [1] L. Arge, O. Procopiuc, S. Ramaswamy, T. Suel, and J.S. Vitter, "Scalable Sweeping-Based Spatial Join," *Proc. 24th Very Large Databases Conf.*, pp. 259-270, June 1998.
- [2] S. Arya, D.M. Mount, and O. Narayan, "Accounting for Boundary Effects in Nearest Neighbor Searching," *Proc. 11th Ann. Symp. Computational Geometry*, pp. 336-344, 1995.
- [3] N. Beckmann, H.-P. Kriegel, R. Schneider, and B. Seeger, "The R^* -Tree: An Efficient and Robust Access Method for Points and Rectangles," *Proc. 1990 ACM-SIGMOD Conf.*, pp. 322-331, May 1990.
- [4] A. Belussi and C. Faloutsos, "Estimating the Selectivity of Spatial Queries Using the Correlation Fractal Dimension," *Proc. 21st Very Large Databases Conf.*, pp. 299-310, Sept. 1995.
- [5] S. Berchtold, B. Ertl, D. Keim, H.-P. Kriegel, and T. Seidl, "Fast Nearest Neighbor Search in High-Dimensional Spaces," *Proc. 14th Int'l Conf. Data Eng.*, Sept. 1998.
- [6] S. Berchtold, D.A. Keim, and H.-P. Kriegel, "The X-Tree: An Index Structure for High-Dimensional Data," *Proc. 22nd Very Large Databases Conf.*, Sept. 1996.
- [7] T. Brinkhoff, H.-P. Kriegel, R. Schneider, and B. Seeger, "Multistep Processing of Spatial Joins," *Proc. 1994 ACM-SIGMOD Conf.*, pp. 197-208, May 1994.
- [8] T. Brinkhoff, H.-P. Kriegel, and B. Seeger, "Efficient Processing of Spatial Joins Using R-Trees," *Proc. 1993 ACM-SIGMOD Conf.*, pp. 237-246, May 1993.
- [9] M.J. Carey and D. Kossmann, "On Saying 'Enough Already!' in SQL," *Proc. 1997 ACM-SIGMOD Conf.*, pp. 219-230, May 1997.
- [10] M.J. Carey and D. Kossmann, "Reducing the Braking Distance of an SQL Query Engine," *Proc. 24th Very Large Databases Conf.*, pp. 158-169, Aug. 1998.
- [11] A. Corral, Y. Manolopoulos, Y. Theodoridis, and M. Vassilakopoulos, "Closest Pair Queries in Spatial Databases," *Proc. 2000 ACM-SIGMOD Conf.*, pp. 189-200, May 2000.
- [12] D. Donjerkovic and R. Ramakrishnan, "Probabilistic Optimization of Top N Queries," *Proc. 25th Very Large Databases Conf.*, Sept. 1999.
- [13] C. Faloutsos, B. Seeger, A. Traina, and C. Traina Jr., "Spatial Join Selectivity Using Power Laws," *Proc. 2000 ACM-SIGMOD Conf.*, pp. 177-188, May 2000.
- [14] A. Guttman, "R-Trees: A Dynamic Index Structure for Spatial Searching," *Proc. 1984 ACM-SIGMOD Conf.*, pp. 47-57, June 1984.
- [15] G.R. Hjaltason and H. Samet, "Ranking in Spatial Databases," *Proc. Fourth Int'l Symp. Large Spatial Databases (SSD'95)*, pp. 83-95, Sept. 1995.
- [16] G.R. Hjaltason and H. Samet, "Incremental Distance Join Algorithms for Spatial Databases," *Proc. 1998 ACM-SIGMOD Conf.*, pp. 237-248, June 1998.
- [17] F. Korn, N. Sidiropoulos, C. Faloutsos, E. Siegel, and Z. Protopoulos, "Fast Nearest Neighbor Search in Medical Image Databases," *Proc. 22nd Very Large Databases Conf.*, pp. 215-226, June 1996.
- [18] M.-L. Lo and C.V. Ravishankar, "Spatial Joins Using Seeded Trees," *Proc. 1994 ACM-SIGMOD Conf.*, pp. 209-220, May 1994.
- [19] M.-L. Lo and C.V. Ravishankar, "Spatial Hash-Join," *Proc. 1996 ACM-SIGMOD Conf.*, pp. 247-258, June 1996.
- [20] Bureau of the Census, *Tiger/Line Precensus Files: 1997 Technical Documentation*, Washington, DC, 1997.
- [21] J.A. Orenstein, "A Comparison of Spatial Query Processing Techniques for Native and Parameter Spaces," *Proc. 1990 ACM-SIGMOD Conf.*, pp. 343-352, May 1990.
- [22] D. Papadias, N. Mamoulis, and Y. Theodoridis, "Processing and Optimization of Multiway Spatial Joins Using R-Trees," *Proc. 1999 ACM SIGACT-SIGMOD-SIGART Symp. Principles of Database Systems*, pp. 44-55, June 1999.
- [23] J.M. Patel and D.J. DeWitt, "Partition Based Spatial-Merge Join," *Proc. 1996 ACM-SIGMOD Conf.*, pp. 259-270, June 1996.
- [24] V. Poosala, "Histogram-Based Estimation Techniques in Databases," PhD thesis, Univ. of Wisconsin-Madison, 1997.
- [25] F.P. Preparata and M.I. Shamos, *Computational Geometry: An Introduction*. New York: Springer-Verlag, 1985.
- [26] V. Ramasubramanian and K.K. Paliwal, "Fast k-Dimensional Tree Algorithms for Nearest Neighbor Search with Application to Vector Quantization Encoding," *IEEE Trans. Signal Processing*, vol. 40, no. 3, pp. 518-531, Mar. 1992.
- [27] N. Roussopoulos, S. Kelley, and F. Vincent, "Nearest Neighbor Queries," *Proc. 1995 ACM-SIGMOD Conf.*, pp. 71-79, May 1995.
- [28] T. Seidl and H.-P. Kriegel, "Optimal Multistep k-Nearest Neighbor Search," *Proc. 1998 ACM-SIGMOD Conf.*, pp. 154-165, June 1998.
- [29] J.S. Vitter and M. Wang, "Approximate Computation of Multi-dimensional Aggregates of Sparse Data Using Wavelets," *Proc. 1999 ACM-SIGMOD Conf.*, pp. 193-204, June 1999.



Hyoseop Shin received the PhD, MS, and BS degrees at the School of Computer Science and Engineering, Seoul National University, Seoul, Korea, in 2002, 1996, and 1994, respectively. He is a senior engineer in the Software Center, Samsung Electronics, Seoul, Korea. His current research interests include XML, XML database systems, digital tv technologies, and TV-Anytime Forum. During 1999 and 2001, he studied at the Department of Computer Science, University of Arizona, as a visiting student scholar with the financial support of the Korea Science and Engineering Foundation (KOSEF) and Brain Korea 21 (BK21) under Exchange Student Program.



Bongki Moon received the PhD degree in computer science from University of Maryland, College Park, in 1996, and the MS and BS degrees in computer engineering from Seoul National University, Korea, in 1985 and 1983, respectively. He is an assistant professor in the Department of Computer Science, University of Arizona. His current research interests include XML indexing and query processing, high performance spatial and temporal databases, scalable web servers, data mining and warehousing, and parallel and distributed processing. He was a member of the research staff at Communication Systems Division, Samsung Electronics Corp., Korea, from 1985 to 1990. He received the US National Science Foundation CAREER Award in 1999.



Sukho Lee received the BA degree in political science and diplomacy from Yonsei University, Seoul, Korea, in 1964, and the MS and PhD degrees in computer sciences from the University of Texas at Austin in 1975 and 1979, respectively. He is a professor in the School of Computer Science and Engineering, Seoul National University, Seoul, Korea, where he has been leading the Database Research Laboratory. He has served as the president of Korea Information Science Society. His current research interests include database management systems, data warehousing, and multimedia database systems.

► For more information on this or any computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.