Discussion Paper No. 379

An $O(n \log^2 n)$ Algorithm for the kth Longest Path in a Tree with Applications to Location Problems

by

N. Megiddo
A. Tamir
E. Zemel
R. Chandrasekaran

April 1979 - Revised March 1980

Northwestern University
Graduate School of Management
Evanston, Illinois 60201

An $O(n \log^2 n)$ Algorithm for the kth Longest Path in a Tree with Applications to Location Problems

N. Megiddo
A. Tamir
E. Zemel
R. Chandrasekaran

Abstract. Many known algorithms are based on selection in a set whose cardinality is super-linear in terms of the input length. It is desirable in these cases to have selection algorithms that run in sub-linear time in terms of the cardinality of the set. This paper presents a successful development in this direction. The methods developed here are applied to improve the previously known upper-bounds for the time-complexity of various location problems.

1. Introduction

It is now well-known that the kth largest element of an ordered set S can be found in linear time in the cardinality of S, [1]. Since the discovery of that fact, it has been observed by several authors that in some structured sets, the kth largest element may be found even faster. For example, if $S = X + Y$ (where both X and Y consist of n numbers) then the kth largest element of S can be found in $O(n \log n)$ time, even though $|S| = n^2$. This was first achieved by Jefferson, Shamos and Tarjan [15] and by Johnson and Mizoguchi [11], and later generalized and improved by Frederickson and Johnson [6]. A more general case is the following: Suppose that the set S is partitioned into m sorted subsets, such that the kth largest element in each subset can be found in constant time. Fox [5] finds the kth largest element of S in $O(m + k \log m)$ time. Galil and Megiddo [7] solve the problem in $O(m \log^2(|S|/m))$ time. The basic idea of [15] can be used to solve this problem in $O(m \log(|S|/m))$ steps. This was improved by Frederickson and Johnson, [6], who solve the same problem in $O(\max\{m, c \log(k/c)\}$ time, where $c = \min(k,m)$. This is also proved to be an asymptotically optimal bound [6,10].

The structure of S in this latter example is quite abstract. It remains an open question how should other structured sets be handled. For example, suppose that S is the set of all pairs of nodes of a graph, ordered according to the distance (along a shortest path) between the members of the pair. How can we exploit this structure on S for finding the kth largest element? Another interesting example is when S is the set of maximum-flows between

pairs of source-sink in a capacitated network.

In this paper we develop an algorithm for the kth largest element in the set of all simple paths in a tree with edge-lengths. The cardinality of this set is $O(n^2)$ (n is the number of nodes in the tree and each simple path is characterized by its two endpoints). However, our algorithm runs in $O(n \log^2 n)$ time. This fast method of selecting an inter-nodal distance is shown to be very useful in the solution of different combinatorial location problems.

The organization of the paper is as follows. In section 2 we review the two basic approaches to selection in an ordered set with sorted subsets. In section 3 we discuss a decomposition scheme for trees, on which the partition of the set of paths is based. The partitioning itself is developed in section 4 and the solution of the selection problem in the set of paths is summarized in section 5. A brief survey of four different location problems is given in section 6. In section 7 we apply the methods developed in this paper, to obtain improved algorithms for the location problems defined in section 6. In section 8 we briefly discuss the more general case of weighted location problems.

2. An Overview of Selection Algorithms

Suppose that an ordered set S is partitioned into m subsets $S_1, S_2, \ldots, S_m$, such that the kth largest element in each subset can be found in constant time. We distinguish between two methods of selection in S. The first one, which we like to call "trimming," is attributed to Jefferson, Shamos and Tarjan [15] and is also used

in Frederickson and Johnson [6]. The second, which we call "splitting," is a generalization of the linear-time median-finding [1] and is used in Johnson and Mizoguchi [11] and in Galil and Megiddo [7]. For simplicity of exposition, we assume in this section that all members of S are distinct. Handling the general case of S being a multi-set is similar (see the above references). At a given iteration, let $S'_i \subseteq S_i$ be the set of elements still under consideration with $S' = \bigcup_{i=1}^{m} S'_i$.

A. Trimming. Suppose that we are looking for the kth largest element x of S' and assume without loss of generality that $k \leq \frac{1}{2}|S'|$. We first find the lower quartile, y_i, in each S'_i. Next, we consider the set $Y = \{y_1, \ldots, y_m\}$ where each y_i is weighted by $|S'_i|$, and we find the (weighted) lower-quartile y of Y. Obviously, at least one half of the elements of S' are greater than y, and hence $y \leq x$. We can now reduce the set S' by discarding the lower quarter of each subset S'_i for which $y_i \leq y$. This amounts to discarding 1/16 of the set S', and the problem now reduces to finding the kth largest element of the remaining set.

B. Splitting. In this method we first find the median z_i in each S'_i. Next, we find the weighted median z of $Z = \{z_1, \ldots, z_m\}$ (relative to the weights $|S'_i|$). Obviously, z is between the lower and upper quartiles of S'. By computing the rank of z in S', we can tell whether $z \geq x$ or $z < x$. In the latter case, the lower half of every S'_i such that $z_i \leq z$, can be discarded and we look for the kth largest element in the remaining set. Otherwise, the upper half of every S'_i such that $z_i \geq z$ is discarded and we look for the $(k - \frac{1}{4}|S'|)$-th largest element of the remaining set.

It is interesting to compare the logic and overall efficiency of splitting and trimming. One difference between the two methods lies in the position of elements they eliminate. At any given iteration, splitting may eliminate elements from the upper or lower quartiles of S' depending on the outcome of a logical test. In contrast, the elimination process of trimming is not based on any test and the elements eliminated always come from lower parts of S' if $k \leq \frac{1}{2} |S'|$ and from its upper part if the reverse condition holds. As will be pointed out in section 7, a procedure similar to splitting (i.e., based on a logical test) turns out to be preferable for solving various location problems on a tree. As for the efficiency of identifying the kth element of S we note that, in the worst case, splitting eliminates at each iteration more variables than trimming. (Four times as many in the formulation given above although the difference can be reduced by a slight modification of the trimming procedure.) However, the corresponding reduction in the number of iterations enjoyed by the splitting method is more then offset by the effort involved in identifying the rank of z in S' which is necessary to support the logical test. Thus, while the overall complexity achieved by the splitting procedure is $0(m \log^2(|S|/m))$ the corresponding complexity for trimming is only $0(m \log(|S|/m))$.

Our algorithm for the kth longest simple path in a tree exploits the structure of the set S of paths in the following way. We partition S into $m = O(n \log n)$ subsets $S_1, \ldots, S_m$ with the following properties:

(i) The kth largest element in any S_i, as well as its length, can be computed in constant time.

(ii) All the elements of each S_i are paths leading from the same node v_i to other nodes of the tree.

(iii) The k largest and the k smallest elements of any S_i can be discarded in constant time.

(iv) The partitioning process is carried out in $O(n \log^2 n)$ time and $O(n \log n)$ space.

Once this partition is obtained, one can employ the trimming algorithm for finding the kth longest path. This amounts to a time bound of $O(n \log^2 n)$. Details are worked out in the following sections. The partitioning is carried out by a divide-and-conquer algorithm on the tree T. The first step in this direction is an efficient decomposition of the tree which we describe in the next section.

3. Decomposition of Trees

In this section we show how to decompose a tree T into three (or less) subtrees, such that precisely one node of T belongs to more than one subtree, and such that each subtree has no more than $\frac{n}{2} + 1$ nodes (where the set of nodes of T is $N = \{1, \ldots, n\}$).

Suppose that the tree T is given in the form of lists N(i) of all the neighbors of a node i (i=1,...,n). If i and j are

neighbors, then by removing the edge (i,j) two subtrees of T are induced. We denote by $K(i,j)$ the number of nodes in that subtree which contains node i. (Note that K is defined on <u>ordered</u> pairs of neighboring nodes.) It is easy to verify the following:

(1) If i is a leaf where $N(i) = \{j\}$ then $K(i,j) = 1$.

(2) $K(i,j) + K(j,i) = n$ for all pairs of neighbors.

(3) For all j, $\sum_{i \in N(j)} K(i,j) = n-1$.

(4) If $j,k \in N(i)$ $(j \neq k)$ then $K(j,i) < K(i,k)$.

In order to decompose T in the manner described above, we need to find a node x such that for all $i \in N(x)$, $K(i,x) \leq \frac{n}{2}$. The existance of such a node, referred to as the <u>centroid</u> of T, was observed by Jordan in 1869 [12]. Linear time algorithms for finding the centroid appear in Goldman [8], and Kariv and Hakimi, [13]. For the sake of completeness we provide such an algorithm below.

We first note that the computation of all the $K(i,j)$'s can be carried out in $O(n)$ time. This is done as follows. Fix one of the nodes r as the "root," so that every other node i has a "father" $f(i)$ relative to r (i.e. $f(i)$ is the node following i on the path from i to r).

The quantities $K(i,f(i))$ $(i \neq r)$ can be computed recursively by $K(i,f(i)) = 1 + \sum_{j:f(j)=i} K(j,f(j))$ and the computation of all $K(i,j)$'s can be completed by (2). The whole process takes $O(n)$ time.

Once all the $K(i,j)$'s are known, the following process can be used to find a node x such that $K(i,x) \leq \frac{n}{2}$ for all $i \in N(x)$.

```
        x ← 1
1       if K(i,x) ≦ n/2 for all i∈N(x) then stop
        else (there is precisely one i∈N(x) such
                that K(i,x) > n/2).    x ← i
        go to 1
```

This procedure generates a path $1 = x_1, \ldots, x_k = x$ such that $K(x_{j+1}, x_j) > \frac{n}{2}$ $(j=1,\ldots,k-1)$. By (4) and (2), the function $m(x_j) \equiv \max_{i \in N(x_j)} K(i,x_j)$ is monotone decreasing along that path and hence an $x_k = x$ is reached for which $m(x) \leqq \frac{n}{2}$. Obviously, this procedure takes $O(n)$ time.

We now claim that the set $N(x)$ can be partitioned into three or less subsets N_1, N_2, N_3 such that $\sum_{i \in N_j} K(i,x) \leqq \frac{n}{2}$. This is easily proved as follows. Assume $N(x) = \{v_1, \ldots, v_p\}$. By (3) there is s $(1 \leqq s \leqq p)$ such that

$\sum_{i=1}^{s-1} K(v_i,x) \leqq \frac{n-1}{2}$ and $\sum_{i=s+1}^{p} K(v_i,x) \leqq \frac{n-1}{2}$. The desired subsets are $N_1 = \{v_1, \ldots, {}_{s-1}\}$, $N_2 = \{v_s\}$, $N_3 = \{v_{s+1}, \ldots, v_p\}$.

Finally, the partition of $N(x)$ induces a decomposition of T into three or less subtrees T_1, T_2, T_3, namely, T_j is the subtree consisting of x and all the nodes accessible from x via a member of N_j $(j=1,2,3)$. Obviously, x is the only node of T that belongs to more than one such subtree, and also in each T_j there are no more than $\frac{n}{2} + 1$ nodes. The decomposition is carried out in $O(n)$ time.

4. Partition of the Set of Paths in a Tree

In the preceding section we described a decomposition of a tree into three subtrees with a single node x common to the three of them. We refer in this section to that node x as the "decomposer." In this section, S is the set of all simple paths in a tree T. Since there is a one-to-one correspondence between pairs of nodes and simple paths in a tree, we also consider S as the set of pairs of nodes, ordered according to the inter-nodal distances. We partition S into subsets such that the kth largest element in any subset can be found in constant time.

The essence of the partitioning algorithm is as follows. First, we find a decomposer x (see Section 3) and we look at the three subtrees, T_1, T_2, T_3 in the corresponding decomposition. For each T_i (i=1,2,3), we compute all the distances from the node x to all other nodes of T_i, and we sort the set S_i of all simple paths leading from x into T_i, according to these distances. Thus, the node x contributes three sorted subsets to our partition of S. Next, for each node $j \neq x$ in T_1 we can easily compute the sorted set of distances from j to all nodes of T_2, since this is obtained by adding a constant (namely, the distance between j and x) to all elements of S_2. Analogously, for each $j \neq x$ in either T_1 or T_2, we compute the sorted set of distances from j to all nodes of T_3, by adding the distance between j and x to all elements of S_3. Thus, each node $j \neq x$ of T_1 contributes at this stage two sorted subsets and each $j \neq x$ in T_2 contributes one sorted subset to our partition of S. We proceed by decomposing the subtres T_1, T_2, T_3, each along the same lines described above until all the paths (or equivalently, pairs) are enumerated. Throughout this process,

we skip paths leading to or from nodes that have previously served as decomposers, to make sure that each pair of nodes is taken into account precisely once.

The number of subsets created during the partitioning process is estimated as follows. Let $M(n)$ denote the maximum number of subsets in such a partition of S for a tree with n vertices. The tree is decomposed into three subtrees. If n_1, n_2, n_3 are the numbers of nodes in these subtrees, then $n_1 + n_2 + n_3 = n+2$ and $n_i \leqq \frac{n}{2} + 1$. Each node contributes no more than three subsets to the partition of S, and we proceed, recursively, with the subtrees. Hence

$$M(n) \leqq 3n + M(n_1) + M(n_2) + M(n_3)$$

and it follows that $M(n) = 0(n \log n)$.

We now estimate the running-time, $T(n)$ of the partitioning process. It is very essential to note here what is meant by "creating" subsets. The creation of the subsets contributed by the first decomposer requires $0(n \log n)$ time, since we need to compute all distances from the decomposer to all other nodes, and then sort them. However, the creation of other subsets (i.e. subsets contributed by non-decomposers) requires only a few pointers as discussed later in this section. Thus, the general step in the partitioning process consists of (i) Tree decomposition: $0(n)$. (ii) Computing all distances from a single node: $0(n)$. (iii) Sorting the set of all these distances and discarding those associated with previous decomposers $0(n \log n)$. (iv) Creating the subsets pointers and constants: $0(n)$. Thus, the recursive relation is

$$T(n) \leqq C n \log n + T(n_1) + T(n_2) + T(n_3)$$

and therefore $T(n) = 0(n \log^2 n)$.

Next, we turn to discuss the storage aspects of the partitioning algorithm. Whenever a node x serves as a decomposer for a subtree T^1, three sorted sets R_1, R_2, R_3 of distances from x into T^1 are generated. We distinguish between the sets R_i and the subsets S_i that actually constitute our partition. Each set is stored as an array, and the total space for storing these arrays is $O(n \log n)$. (This can be proved by induction.) The partition of S into subsets $S_1, \ldots, S_m$, as well as the reduced forms of S, that are processed by the "trimming" or "splitting" procedures (see Section 2), are handled as follows. Each S_i is characterized by four items. First is a pointer to the corresponding R_j from which S_i is created. Second, is a constant number that should be added to an element of R_j in order to get an element of S_i. Finally, we need two pointers that specify the boundary of that portion of R_j from which S_i is generated. (These two pointers are at the start the same for all the S_i's that rely on the same R_j, but during the "trimming" or "splitting" process, they may become different.) Thus, the total amount of storage that we need is $O(n \log n)$. In addition, at most $O(n \log n)$ storage is required in order to properly maintain the set of trees, T_i, which are generated throughout the algorithm.

We conclude this section with a pigin algol description of the partitioning process. It receives as input a tree T with a set of nodes $N = \{1, \ldots, n\}$ and produces as output a partition $S_1 \ldots S_m$ of the set of inter-nodal distances of

T, where $m = 0(n \log n)$. The sets $S_1 \ldots S_m$ satisfy the properties (i)-(ii) of section 2. The overall complexity bounds for the algorithm are $0(n \log^2 n)$ time and $0(n \log n)$ space. The procedure uses the following terminology:

Q -	current set of subtrees not yet subdivided
B -	current set of nodes which have not as yet served as decomposers.
R_j -	j^{th} sorted set of distances between a decomposer and the nodes of a subtree.
k -	index for set S_k used in the partition
$\gamma(k)$ -	a label identifying the index of subset R_j used to create S_k
$\beta(k)$ -	the constant increment which must be added to each element of R_j to get the corresponding element of S_k

In addition, the procedure uses the following subroutines in the course of its execution:

CENTROID (T) -	Given a tree T return its centroid
SUBTREE (T,x, i) -	Given a tree T, its centroid x, and an index i=1...3, returns the subtree T_i (see last paragraph of section 3)
DISTANCE (T,A,B) -	Given a tree T and two sets of nodes A, and B returns a vector of all the distances d(i,j) $i \in A$, $j \in B$, $i \neq j$.
SORT (D) -	Given a vector D, returns the entries in a sorted way.

```
Procedure  DECOMPOSE (T)

begin

    Q ← T
    B ← N
    j ← 0
    k ← 0

    while   Q ≠ Ø do
        begin
            choose T' from Q
            x ← CENTROID (T')
            for  i = 1,,3   do
                begin
                    T_i ← SUBTREE (T,x,i)
                    N_i ← Nodes of T_i
                    N'_i ← N_i ∩ B \ {x}
                    D_i ← DISTANCE (T_i, x, N'_i)
                    j ← j + i
                    v_i ← j
                    R_j ← SORT (D_i)
                end

            for  each j ε N'_1   do
                begin
                    k ← k + 1
                    γ(k) ← v_2
                    β(k) ← d(j,x)
                    k ← k + 1
                    γ(k) = v_3
                    β(k) ← d(j,x)
                end
```

```
for each      j ε N'_2   do

    begin
        k ← k + 1
        γ(k) ← v_3
        β(k) ← d(j,x)
    end

if   x ε B   do
    begin
        for  i = 1,...3   do
            begin
                k ← k+1
                γ(k) ← v_i
                β(k) ← 0
            end
        B ← B \ {x}
    end
for  i = 1,...3   do
    begin
        if   |N_i| ≥ 3  then   Q ← QUT_i
        else
                D_i ← DISTANCE (T, N'_i, N'_i)
                j ← j + 1
                v_i ← j
                R_j ← SORT (D_i)
                k ← k + 1
                γ(k) ← v_i
                β(k) ← 0
    end
end
```

5. The kth Longest Path in a Tree

Once the partition of the set S of all paths into $m=0(n \log n)$ subsets is established, one can use the techniques introduced in section 2 to find the kth longest path. This amounts to an effort of $0(m \log n/m) = 0(n \log^2 n)$ if one uses trimming and an inferior bound, of $0(n \log^3 n)$, if splitting is used. As the effort involved in generating the partition of S is also $0(n \log^2 n)$ we can conclude that the overall effort for finding the k longest path in T is $0(n \log^2 n)$.

Can this bound be further beaten down? Possibly, but the margin for improvement is slim. An $0(n \log n)$ lower bound on the complexity of the problem can be obtained in a number of ways. The following simple reduction was offered to us by one of the referees. Consider the tree of figure 1 where the heavy line in the center is chosen long enough to ensure that the longest paths in T include one element from X and one from Y. Thus, the well known $0(n \log n)$ bound on selection in X + Y is valid for our problem as well.

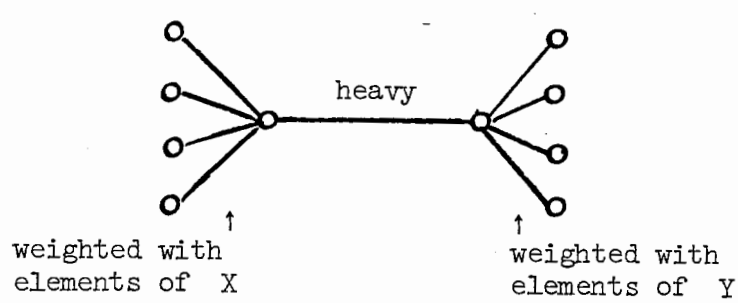

Figure 1

6. Location Problems

We consider here the following different problems of location. First, we assume that a tree T is embedded in the Euclidean plane, so that the edges are line segments whose endpoints are the nodes and edges intersect one another only at nodes. Moreover, each edge has a positive length. (Any tree with positive edge-weights can be so embedded in R^2.) This embedding enables us to talk about points, not necessarily nodes, on the edges. We then denote by $d(x,y)$ the distance, measured along the edges of the tree, between any two points x,y of the tree.

In a typical location problem, one has to select p points of the tree under different assumptions, depending on the particular model considered. In each model, we distinguish between the "supply" set, Σ, (this is the set from which we select the p points) and the "demand" set Δ, with reference to which the objective function is defined. The p-center problem seeks to choose p points $x_1,\ldots,x_p$ from Σ so as to minimize $\sup_{y\in\Delta} \min_{1\leq i\leq p} d(x_i,y)$. The four special cases, where the sets Σ and Δ are either the set of all nodes or the set of all points of the tree, have been discussed and given different algorithms in [2,3,4,9,13].

Following Handler, [9], we use the categorization scheme $\{{N \atop A}\}/\{{N \atop A}\}/p$, interpreted as follows. The first cell describes the supply set Σ, which could be either the set N of all nodes, or the set A of all points. The second cell describes the demand set Δ which could also be either N or A. The third cell indicates the

number of points that we have to select from Σ. For example, N/A/2 refers to selecting two nodes, so as to minimize the maximum (over all points of the tree) of a distance between a point of the tree and the selected node that is nearest to that point. Kariv and Hakimi [13] provide $O(n^2 \log n)$ algorithms for A/N/p and N/N/p. Chandrasekaran and Tamir [3], using a unified approach, solve A/N/p, N/N/p and N/A/p in $O(n^2 \log n)$ time. The A/A/p problem is solved in [4] in $O(n^2 \log^2 p)$ time.

All the algorithms mentioned above are based on the same principle. First, a finite set R of real numbers, which is known to contain the optimal objective function value, is identified. Next, we search R for the minimum value which is feasible in the following sense. A value $r > 0$ is feasible if there exists a set of p points $x_1,\ldots,x_p$ of Σ, such that the distance between any demand point y and its nearest x_i, is not greater than r. Efficient algorithms are known for deciding whether a given r is feasible, and hence the location problem can be solved by a binary search of R, using such a feasibility test. For all four problems this test runs in O(n) time. (See [13] for N/N/p and A/N/p and [4] for N/A/p and A/A/p.) The set R of relevant values in the four different problems is (see [3,4,13]):

Model	The set R
N/N/p	$\{d(i,j)\}_{i,j\in N}$
A/N/p	$\{\frac{1}{2}d(i,j)\}_{i,j\in N}$
N/A/p	$\{d(i,j),\ \frac{1}{2}d(i,j)\}_{i,j\in N}$
A/A/p	$\{\frac{1}{2k}\ d(i,j)\}_{i,j\in N,\ k=1,\ldots,p}$.

Along the lines discussed above, each one of these problems can be solved by computing the set R, and then searching R by repeatedly using linear-time median-finding [1]. This amounts to $O(|R| + n \log |R|)$ time where $|R|$ is the dominant term. Thus, in order to improve this upper-bound, one has to bypass the computation of the set R and still be able to search in that set. This is essentially where we apply the techniques developed in the previous sections.

7. Improved Algorithms for Location Problems

The sets R of relevant values for the various versions of the p- center problem bear a close resemblance to the set S of inter-nodal distances on T. Thus, we can use any algorithm for finding the kth longest element in S to support a binary search over R. Such search involves at each iteration identifying the median element of R, performing the feasibility test and finally discarding one half of the elements. However, we note that identifying the median element at each iteration may be more than one needs. In fact, one can do better by applying a search strategy similar to that of splitting.

Assume that the set R is partitioned into m subsets $R_1, \ldots, R_m$ such that the kth largest element in each subset can be found in constant time. We can employ the following procedure. First, we find the median element, Z_i, in each subset R_i. Next, we find the median element, z, in the set $Z = (Z_1 \ldots Z_m)$, relative to the weights $|R_i|$. Thus, z is between the lower and

upper quartiles of R. This value z can now be tested for feasibility. The test takes $O(n)$ time and determines whether the optimal value v is greater than z (this is when z is not feasible) or not. If $v > z$, we discard the lower half (including z_i) from each R_i such that $z_i \leqq z$. If $v \leqq z$, we discard the upper half (including z_i) from each R_i such that $z_i \geqq z$, with the exception that z itself is not discarded. The search then proceeds with the reduced set R until the optimal value is singled out.

Since each reduction eliminates one quarter of the remaining set, the number of such stages is $O(\log|R|)$. Following is a more detailed analysis for the particular cases.

A. N/N/p

Here R is the set of the inter-nodal distances. It follows from Section 4 that R can be appropriately partitioned into $O(n \log n)$ subsets where the partitioning process takes $O(n \log^2 n)$ time. During the searching process, in each iteration we need $O(n \log n)$ time for identifying the element z and $O(n)$ time for the feasibility test. Thus, the searching stage takes $O((n + n \log n)\log|R|)$ time and hence the location problem is solved in $O(n \log^2 n)$ time.

B. A/N/p

Since $R = \{\frac{1}{2} d(i,j) : i,j \in N\}$ in this case, the location problem is solved by the same partition which is used in N/N/P. Hence, the time bound for this case is $O(n \log^2 n)$.

C. N/A/p

The relevant set here is $R = \{d(i,j), \frac{1}{2} d(i,j) : i,j \in N\}$. Thus, we use essentially the same partition as the one for N/N/p and A/N/p, but in terms of pairs of nodes, each pair is counted twice: once for the distance d(i,j) and once for the number $\frac{1}{2} d(i,j)$. This implies the same time bound of $O(n \log^2 n)$ for this case too.

D. A/A/p

This last case is slightly more complicated than the previous ones. Since $R = \{\frac{1}{2k} d(i,j) : i,j \in N, k=1,\ldots,p\}$ in this case, one way of partitioning R is by using the partition of Section 4 for the set of pairs of nodes, and replicating each subset p times, so that each d(i,j) is multiplied by all the numbers $\frac{1}{2}, \frac{1}{4}, \frac{1}{6}, \ldots, \frac{1}{2p}$. Thus, R is partitioned into $m = O(p\, n \log n)$ subsets while $|R| = O(p\, n^2)$. By applying the searching method, R is successively reduced by a factor of one quarter. Let the set of remaining variables at a given iteration be R', and denote by $T(|R'|)$ the time required by the algorithm to handle this set. We now have

$$T(R') \leq C_1 m + c_2 n + T(3/4|R'|)$$

When the cardinality of R' reaches the level $|R'| = O(m)$ we can search over R' directly using the method of linear time median finding. This involves, at each iteration, finding the median element and the performing the test. The total effort involved in the identification of <u>all</u> the median elements is clearly $O(m)$. Also, since the number of iterations is $O(\log m)$, and since each test requires an effort of $O(n)$, the total effort associated with handling a set of cardinality $O(m)$ is $O(n \log m) = O(m)$. Solving the recursion relation with the initial condition $T(m) = O(m)$

we then get that $T(|R'|) = 0(m \log(|R'|/m))$ and hence the location problem is solved in this approach in time

$$0(p\, n \log n \log(\frac{n}{\log n})) = 0(p\, n \log^2 n).$$

There is an alternative partition that in some cases leads to a better upper-bound. We first compute the $\frac{1}{2}\, n(n-1)$ inter-nodal distances (in $0(n^2)$ time). Then we partition R into $m = \frac{1}{2}\, n(n-1)$ subsets of the form $R_{ij} = \{\frac{1}{2k} d(i,j) : k=1,\ldots,p\}$, where computing the kth largest element in each set is trivial. Applying the searching procedure we obtain the following bound:

$$0(m \log(|R|/m)) = 0(n^2 \log p).$$

To summarize, the different cases are solved with the following upper-bounds

Model	Upper-bound
N/N/p	$0(n \log^2 n)$
A/N/p	$0(n \log^2 n)$
N/A/p	$0(n \log^2 n)$
A/A/p	$0(n \min\{p \log^2 n,\ n \log p\})$

8. Location Problem with Weighted Demands

A more general type of location problem is where the demand is weighted. Specifically, when $\Delta = N$ we may have weights $w_i > 0$ $(i \in N)$ and seek to select $x_1,\ldots,x_p \in \Sigma$ so as to minimize

$$\max_{i \in N}\{w_i \cdot \min_{1 \le j \le p} d(x_j, i)\}.$$

It is shown in [3,13] that the relevant set R generalizes to $\{w_i d(i,j)\}_{i,j \in N}$ in the N/N/p case and to $\{\frac{w_i w_j}{w_i + w_j} d(i,j)\}_{i,j \in N}$ in the A/N/p case. Both cases are solvable in $O(n^2)$ time [4,13], but based on our method an $O(n \log^2 n)$ algorithm for the weighted N/N/p case can be constructed as follows.

Essentially, we consider the set S' of all <u>ordered</u> pairs of nodes together with the linear order induced by the weighted distances $w_i d(i,j)$. The set S' is partitioned, along lines similar to those of Section 4, into $O(n \log n)$ subsets. All the pairs (i,j) in any subset are with the same i, hence the restriction of the order to each subset is independent of the weight w_i. All we have to do during the algorithm, is to multiply the kth largest distance in a set corresponding to i by the weight w_i. Thus the partition satisfies all the properties that are required to obtain a bound of $O(n \log^2 n)$.

References

[1] M. Blum, R.W. Floyd, V.R. Pratt, R.L. Rivest and R.E. Tarjan, "Time Bounds for Selection," J. Com. Sys. Sci. 7 (1972). 448-461.

[2] R. Chandrasekaran and A. Daughety, "Problems of Location on Trees," Discussion Paper No. 357, Center for Mathematical Studies in Economics and Management Science, Northwestern University, 1978.

[3] R. Chandrasekaran and A. Tamir, "Polynomially Bounded Algorithms for Locating P-Centers on a Tree," Discussion Paper No. 358, Center for Mathematical Studies in Economics and Management Science, Northwestern University, 1978.

[4] R. Chandrasekaran and A. Tamir, "An $O((n \log P)^2)$ Algorithm for the Continuous P-Center Problem on a Tree," Discussion Paper No. 367, Center for Mathematical Studies in Economics and Mangement Science, Northwestern University, 1978.

[5] B.L. Fox, "Discrete Optimization via Marginal Analysis," Manag. Sci., 13 (1966) 210-216.

[6] G. N. Frederickson and D.B. Johnson, "Optimal Algorithms for Generating Quantile Information in X+Y and Matrices with Sorted Columns," Proceedings of the 1979 Conference on Information Sciences and Systems, The Johns Hopkins University (to appear).

[7] Z. Galil and N. Megiddo, "A Fast Selection Algorithm and the Problem of Optimum Distribution of Effort," J. ACM 26, (1979), 58-64.

[8] A.J. Goldman, "Optimal Center Location in Simple Networks," Transportation Science, 5, (1971), 212-221.

[9] G.Y. Handler, "Finding Two-Centers of a Tree: The Continuous Case," Transportation Science, 12 (1978) 93-106.

[10] D.B. Johnson and D. S. Kashdan, "Lower Bounds for Selection in X+Y and Other Multisets," J. ACM 25 (1978) 556-570.

[11] D.B. Johnson and T. Mizoguchi, "Selecting the Kth Element in X+Y and $X_1+X_2+\ldots+X_m$," SIAM J. Comput. 7 (1978) 147-153.

[12] C. Jordan, "Sur les Assemblayes des lignes," J. Reine Angew, Math, 70, 185-190 (1869)

[13] O. Kariv and S.L. Hakimi, "An Algorithmic Approach to Network Location Problems; Part I: P-Centers," SIAM J. Appl. Math 37 (1979) 513-538.

[14] _____, "An Algorithmic Approach to Network Location Problems, Part II: P-Medians" SIAM J. Appl. Math 37 (1979) 539-560.

[15] M.I. Shamos, "Geometry and Statistics: Problems at the Interface," In Algorithms and Complexity: New Directions and Recent Results, J.F. Traub, ed., Academic Press, (1976) 251-280.