Discussion Paper No. 397

LOCATING CENTERS ON A TREE WITH DISCONTINUOUS SUPPLY AND DEMAND REGIONS

by

Arie Tamir

and

Eitan Zemel*

Kellogg Graduate School of Management
Northwestern University
Evanston, Illinois 60201

August, 1979

Revised April 1980

*The research of this author was supported in part by the National Science Foundation Grant Number ENG-7902506.

Abstract

Consider a tree $T = (N,E)$ with "supply" and "demand" regions Σ and Δ, each composed of a finite number of disjoint, closed and connected subregions of T, some of which may possibly consist of just one point. Given an integer p, we seek a collection of p "centers" $x_1,\ldots,x_p \in \Sigma$, which minimize the expression $\max_{y\in\Delta} \min_{i=1,\ldots,p} d(y,x_i)$. We present a polynomial algorithm for this problem. Its running time is bounded by $O(n\log^2 n)$ if either Δ or Σ is discrete, and by $O(n \min\{p\log^2 n, n\log p\})$ if both sets contain at least one full edge.

I. Introduction

In this paper we present an efficient algorithm for a general version of the p center problem on an undirected tree network.

To formulate our problem precisely, we assume that an undirected tree $T=T(N,E)$ is embedded in the Euclidean plane, so that edges are line segments whose endpoints are the nodes and edges intersect one another only at nodes. Moreover, each edge of T has a positive length. This embedding enables us to talk about points, not necessarily nodes, on the edges. We denote by A the (infinite) set of points of T. For any two points $x,y \in A$, let $d(x,y)$ denote the distance between x and y, measured along the edges of T. The general p-center problem can be formulated as follows:

(P) Given

(i) A set of points $\Sigma \subseteq A$ (The "supply" set)

(ii) A set of points $\Delta \subseteq A$ (The "demand" set)

(iii) A positive integer p.

Find a collection of p points, $x_1,\ldots,x_p \in \Sigma$ which minimizes the expression

$$\sup_{y\in\Delta} \min_{i=1\ldots p} d(y,x_i) \tag{1}$$

We refer to the particular version of (P) defined on T by a specific choice of Σ, Δ, and p as $\Sigma/\Delta/p$. We use r^* to denote the optimal value of (1).

Following Hakimi [7], different versions of minimax location problems on networks have been studied quite extensively, with emphasis given to the algorithmic aspects. The main results in this respect appear in [2, 3, 4, 5, 6, 7, 8, 9]. To date, polynomially bounded algorithms exist for the four versions of (P) obtained by choosing the pair Σ, Δ as the different combinations of the sets N and A [3, 4, 8, 9]. In addition, a polynomial algorithm is given in

[3] for the case where each of these sets is composed of a finite number of points located anywhere on T. The best worst case time bounds for the above special cases are given in Table 1.

Table 1

model	running time	reference
N/N/p	$O(n\log^2 n)$	9
N/A/p	$O(n\log^2 n)$	9
A/N/p	$O(n\log^2 n)$	9
A/A/p	$O(n\cdot\min\{p\log^2 n,\ n\log p\})$	9
S/D/p where S,D are finite sets of discrete points with $\lvert D\rvert = m$, $\lvert S\rvert = k$	$O(m\cdot n + m^2 k\log(k\cdot m))$	3

In this paper we consider the case of more complicated, yet more realistic sets Σ & Δ. We allow each of these sets to consist of a finite number of disjoint, closed and connected regions, some of which may possibly consist of just one point. Such sets Σ and Δ arise in practice from a variety of reasons. For instance, several regions of T can often be excluded from consideration as potential facility location sites due to the inexistence of appropriate ammenities, restrictive zoning laws, prohibitively high price of property, the desire to maintain a certain distance between facilities and major population centers, etc.

The algorithm presented in this paper can handle such problems quite efficiently. Let n denote the number of nodes of T, including all the tips of the different supply and demand regions. The running time of our algorithm is $O(n\log^2 n)$ if $\Delta \subseteq N$ or

$\Sigma \subseteq N$, and $O(n \min\{p\log^2 n, n\log p\})$ if both Σ and Δ contain at least one full edge. It can be viewed as a generalization of the algorithms of [9] in the sense that its restriction to each of the four basic versions of (P) yields back the algorithm [9] for the relevant version.

The organization of the paper is as follows. In section II we introduce some definitions and preliminaries. In section III we give an overview of existing algorithms for the different versions of (P) and briefly discuss the broad outlines of the algorithm presented here. In section IV we present a set of real numbers, R, which is known to contain the optimal solution value for (P), r^*, and which plays a crucial role in our algorithm. Section V describes the feasibility test which determines whether a given element $r \in R$ satisfies $r < r^*$ or $r \geq r^*$. This test forms a basis for a binary search for r^* on R. Finally, in section VI we sum our algorithm up and discuss its computational complexity.

II. Definitions and Preliminaries

Let $\{\Sigma_i\}$, $i \in I$, and $\{\Delta_j\}$ $j \in J$ be two finite collections of nonempty, closed and connected subsets of T, and assume that $\Sigma_i \cap \Sigma_k = \emptyset$ $(\Delta_j \cap \Delta_\ell = \emptyset)$ for all $i, k \in I$, $i \neq k$ $(j, \ell \in J, j \neq \ell)$. Let, also $\Sigma = \bigcup_{i \in I} \Sigma_i$ and $\Delta = \bigcup_{j \in J} \Delta_j$. We say that a point $x \in \Sigma_i$ $(x \in \Delta_j)$ is <u>extreme</u> in that set if x is a tip of T, or if it is a boundary point of $T\backslash\Sigma_i$ $(T\backslash\Delta_j)$. We denote the set of extreme points of Σ_i (Δ_j) by S_i (D_j) with $S = \bigcup_{i \in I} S_i$, $D = \bigcup_{j \in J} D_j$. We use the convention $N \supseteq S \cup D$ or else we enlarge the set N accordingly. (This may require subdividing some edges into smaller segments.)

The convention $N \supseteq D \cup S$ allows one to partition the open edges of T and its nodes into four types, depending on their intersection with the sets Σ and Δ. An edge is of the supply only (demand only) type if it is included in Σ but not in

Δ (in Δ but not in Σ). It is of the supply and demand type, if it is included in $\Sigma \cap \Delta$ and it is neutral if $\Sigma \cup \Delta$ does not include it. A similar classification is used for the nodes of the tree.

We note that under the above assumptions on Σ and Δ, the supremum in (1) can be replaced by maximum. Furthermore, noting that the maximum in (1) is zero if and only if $\Delta \subseteq \Sigma$ and $|\Delta| \leq p$, we will assume without loss of generality that that maximum is positive.

III. Overview of Algorithms for (P)

The algorithms for the different versions of (P) mentioned in the introduction are based on the same principle. First, we identify a finite set, R, of real numbers which is known to contain the optimal objective function value. Next, we search R for the minimum value which is feasible in the following sense. A value $r > 0$ is called 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 the four models mentioned earlier, this test runs in $O(n)$ time. The set R of relevant values for these four different models is given in Table 2 (see [4, 8]).

Table 2

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

Thus, 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. 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 approach is taken in [9]. It is shown there how the k^{th} longest inter-nodal distance in a tree can be found in an effort which is bounded by $O(n\log^2 n)$. This enables one to find the k-th element in the set R, which in all four cases is closely related to the set of internodal distances on the tree, in time which is less than linear in the cardinality of R. This observation is the basis for the efficient algorithms for the various models of (P) achieved in [9].

The algorithm below uses the same general strategy. First, the set R is derived for the general model considered in this paper. Although this set is slightly more complicated than those applicable for the four simpler models, it still preserves enough resemblence to the set of internodal distances on T to allow one to use the construction of [9] in order to identify the k-th element in this set in sublinear time. We then show how the feasibility tests of [4], [8], can be generalized to deal with the problem treated here, while still running in linear time.

IV. The Set R

For any two points x and y on T, let $P(x,y)$ denote the set of points on the path connecting x and y. A chain X is a finite sequence of distinct points of T (but not necessarily of N) $X = (x_1,\ldots x_m)$, $m \geq 2$, which form a simple path on T, i.e., such that for every $1 \leq i < j < k \leq m$, $x_j \in P(x_i,x_k)$. A chain X is called regular if for some r $d(x_i,x_{i+1}) = r$ $i=1,\ldots,m-1$. It follows from the definition that if X is regular, $r = \frac{d(x_1,x_m)}{m-1}$. A chain X is called alternating if the points $x_1,\ldots,x_m$ are drawn alternatingly from the sets Σ and Δ. If $x_i \in$

$\Sigma \cap \Delta$ for all $i = 1,\ldots,m$, there may be two possible alternating assignments of the points of X to the sets Σ and Δ. We will use the term alternating chain with the understanding that a unique assignment has been specified.

Let $X = (x_1,\ldots,x_m)$ be an alternating chain. We say that the point x_1 is a <u>limit</u> point in X if one of the following two conditions is satisfied:

(i) x_1 is chosen from Σ and any "small" movement from x_1 in the direction of x_m takes one out of the set Σ, i.e., for any $\varepsilon > 0$ there exists a point $u = u(\varepsilon)$ on $P(x_1,x_m)$ with $d(x_1,u) < \varepsilon$ and $u \notin \Sigma$.

(ii) x_1 is chosen from Δ and any small movement from x in any direction away from x_m takes one out of the set Δ, i.e., there exists $\varepsilon > 0$ such that for all points $u \notin P(x_1,x_m)$, $d(x_1,u) < \varepsilon \Rightarrow u \notin \Delta$.

By symmetry, this definition is applied to x_m as well. A point x_i, $1 < i < m$ is an <u>internal</u> point of X if it is neither a limit point of the chain $(x_1,\ldots,x_i)$ nor of the chain $(x_i,x_{i+1},\ldots,x_m)$.

<u>Proposition 1</u>

Let $r^* > 0$ be the optimal solution for (1). Then there exists a regular alternating chain $X = (x_1,\ldots,x_m)$ such that

1. $m \leq 2p+1$
2. $r^* = \dfrac{d(x_1,x_m)}{m-1}$
3. x_1 and x_m are limit points of X.
4. For all $1 < i < m$, x_i is an internal point of X.

<u>Proof</u>.

Let $\bar{X} = \{x_1,\ldots,x_m\}$, $x_i \in \Sigma$, $i=1,\ldots,m$ be an optimal solution for (P) with

$$\max_{y \in \Delta} \min_{i=1,\ldots,p} d(x_i,y) = r^* > 0.$$

Let also

$$Y = \{y \in \Delta \mid \min_{i=1,\ldots,p} d(x_i,y) = r^*\}$$

and

$$\tilde{X} = \{x \in \overline{X} \mid \min_{y \in Y} d(x,y) = r^*\}$$

Note that for $x \in \overline{X}$ $\min_{y \in Y} d(x,y) \geq r^*$.

Obviously, $Y \neq \emptyset$, $\tilde{X} \neq \emptyset$. If there exist more than one optimal solution for (P), let $\overline{X}$ be one which minimizes the cardinality of $\tilde{X}$.

Now let $Z = (z_1 \ldots z_m)$ be a maximal regular alternating chain such that its Σ elements are chosen from $\overline{X}$, its Δ elements are chosen from Y, and $d(z_i,z_{i+1}) = r^*$, $i=1\ldots m-1$. Such a chain obviously exists since $\tilde{X}$ and Y are nonempty. Further, the Σ elements of Z are all members of $\tilde{X}$.

First we show that both z_1 and z_m are limit points of Z. Assume, on the contrary that one of z_1 and z_m is not a limit point, and without loss of generality let that point be z_1. We distinguish between two cases.

1) z_1 represents Σ

This situation is depicted in Figure 1, i.e., there exists $v \neq z_1$, $v \in P(z_1,z_m)$ and $P(z_1,v) \subseteq \Sigma$.

Figure 1

Since $Z = (z_1,\ldots,z_m)$ is maximal, every demand point $u \in \Delta$ to the "left" of z_1 (i.e., $u \in \Delta$, $u \neq z_1$ and $z_1 \in P(u,z_m)$), is either served by a supply center which is not a member of the chain Z, i.e. $d(u,x_i) \leq r^*$ for some $x_i \in \overline{X}$ and $x_i \notin Z$, or else $d(u,z_1) < r^*$. Thus, we can move the supply center, currently at z_1, by a small distance to the right (i.e. along $P(z_1,v)$) without changing the objective value. But, then, such a movement induces a solution to (P) with a smaller value for $|\tilde{X}|$ - contrary to the minimality of the original set $\tilde{X}$.

2) z_1 represents Δ

This situation is depicted in Figure 2, i.e., there exists $v \neq z_1$, $z_1 \in P(v,z_m)$ and $P(v,z_1) \subseteq \Delta$.

Figure 2

Also there exists z_2 in Z, such that z_2 represents Σ and $d(z_1,z_2) = r^*$. Since z_2 is the closest supply center serving z_1, no demand point $u \in P(v,z_1)$, $u \neq z_1$, is served by points on $P(z_1,z_m)$, i.e. $d(u,x) > r^*$ for all $x \in P(z_1,z_m)$, $x \in \bar{X}$.

Let x_k be the closest facility to z_1, among all the facilities serving points on $P(v,z_1) \setminus \{z_1\}$. Since x_k is the closest to z_1, we cannot have $d(x_k,z_1) > r^*$. Also, $d(x_k,z_1) = r^*$ would contradict the maximality of Z, since that would imply $z_1 \in P(x_k,z_m)$ and one could add x_k to the chain. Hence, $d(x_k,z_1) < r^*$, which in turn contradicts the assumption that z_1 belongs to Y.

We have thus shown the existence of a chain satisfying all but condition (4) of the proposition.

Let $Z' = \{z'_1,\ldots,z'_q\}$ be a minimal chain satisfying conditions (1), (2), (3) of the proposition. Suppose that Z' does not satisfy condition (4), and assume that for some $1 < i < q$, z'_i is a limit point of at least one of the two proper subchains $(z'_1,\ldots,z'_i)$, $(z'_i,\ldots,z'_q)$. In either case, we have a contradiction to the minimality of Z'.

$\blacksquare$

Proposition 1 forms a basis for generating a suitable set R for (P). The following simple proposition can be used to bound the cardinality of this set by asserting that the end points of each of the chains referred to in proposition 1 are actually extreme points of Σ or Δ of a special type.

<u>Proposition 2</u>. Let $X = (x_1,\ldots,x_m)$ be an alternating chain in T and let x_1 be a limit point of X. Then:

(1) If x_1 represents Σ, we have

$$x_1 \in S' = \{x \in S \mid \text{there exists } \underline{\text{at least}} \text{ one edge } (x,v), \text{ incident to } x, \text{ such that } (x,v) \notin \Sigma\}$$

(2) If x_1 represents Δ we have

$$x_1 \in D' = \{x \in D \mid \text{there exists at most one edge } (x,v), \text{ incident to } x, \text{ such that } (x,v) \in \Delta\}$$

The proof of Proposition 2 is immediate from the definition of a limit point. Note, that by symmetry, an analogous claim can be made with respect to x_m. $\blacksquare$

It is interesting to compare the sets R mentioned earlier with respect to the four basic models of (P), with the set R of Proposition 1. As all four models are special cases of the model discussed here, we get that the four corresponding sets can be obtained as special cases of Proposition 1. We now turn to examine each of those four cases.

In the paragraphs below, we refer to the set of <u>tips</u> of T by N' ($N' \subseteq N$).

<u>N/N/p</u> here we have

$\Sigma = S = S' = N$

$\Delta = D = D' = N$

and there are no points of Σ or Δ which can serve as internal points. Thus, for each chain, we have $m = 2$ and $R = \{d(i,j)\}i,j \in N$

<u>A/N/p</u> $\Sigma = A,\ S = N' \quad S' = \emptyset$

$\Delta = D = D' = N$

and any point of $A \setminus S$, and no point of Δ, can serve as an internal point. Thus, we have $m = 3$, and $R = \{1/2\ d(i,j)\} i,j \in N$.

N/A/p

$$\Sigma = S = S' = N$$

$$\Delta = A,\ D = D' = N'$$

Any point of $A \setminus D$, but no point of Σ, can serve as internal point. Thus, we have chains with $m = 2,3$ yielding the set $R = \{d(i,j\}_{i \in N, j \in N'} \cup \{\frac{1}{2} d(i,j\}_{i,j \in N}$.

A/A/p

$$\Sigma = A,\ S = N',\ S' = \emptyset$$

$$\Delta = A \quad D = D' = N'$$

any point of $\Sigma \setminus S$ and $\Delta \setminus D$ can serve as an internal point. Since the two end-points of each chain represent Δ we have $m = 2k + 1$, with exactly k points representing Σ. Hence, we must have that $k \leq p$. Therefore,

$$R = \{\frac{1}{2k}\ d(i,j)\}_{i \in N',\ j \in N'}\ k = 1,\ldots,p .$$

We now return to the set R for the general case. As mentioned earlier, one can use the set implied by Proposition 1. However, this set may be quite complex and its determination rather time consuming. Thus, from an algorithmic point of view, it is sometimes more convenient to work with a different set R, which may be larger than the set of Proposition 1 yet is more readily available. The algorithm proposed in this paper is based on the set

$$R = \{\frac{d(x,y)}{j}\}\ x,\ y \in N,\ j \in M = \{1,\ldots m\}$$

where

$$m = \begin{cases} 2 \text{ if either } \Sigma \subseteq N \text{ or } \Delta \subseteq N \\ 2p \text{ if at least one of the pair } \Sigma,\Delta \text{ contains a full edge.} \end{cases}$$

Thus, the cardinality of the set R is $O(n^2)$ in the first case, $O(n^2p)$ in the second. To search over this set, we need a routine which can decide for a given $v \in R$, whether or not v is feasible. Next, we describe such a test, of computational complexity $O(n)$.

V. The Feasibility Test

Like the tests of [4, 8], the test proposed here scans the tree from tips inwards, examining each edge exactly once. However, in contrast with the relatively simple cases treated in those references, one is faced here with a much more complex set of possible combinations of types of edges. Luckily, it turns out that the scheme required to handle the situation is rather simple. However, the arguments needed to support the validity of the test are quite involved and consist of numerous special cases. We give below the essentials of the test procedure. Elaborate proofs and some of the missing details are left for the appendix.

As the test procedure progresses, edges of T are eliminated and the supply and demand regions are modified continuously. At a given stage, let $\overline{T} = (\overline{N}, \overline{E})$ be the current tree under consideration. Similarly, define the sets $\overline{N}'$, $\overline{A}$, $\overline{\Sigma}$ and $\overline{\Delta}$.

We start by rooting the original tree at a given node. The algorithm works by successively eliminating clusters of $\overline{T}$. By a cluster, we mean here a maximal set of edges of the form $\{(x,i): i \in \overline{N}'\}$, where x is fixed, and there is no $j \notin \overline{N}' \cup \{x\}$ such that x is on the (unique) path connecting j with the root. The above x is then called the base of the cluster. A cluster whose base is x is denoted by C(x). If the cluster C(x) does not exhaust the tree $\overline{T}$ (i.e., $C(x) \neq \overline{E}$), there exists a unique edge (x,t) such that $t \notin \overline{N}'$. We call this edge the stem of C(x). A cluster, its base x and its stem (x,t) are depicted in Figure 3.

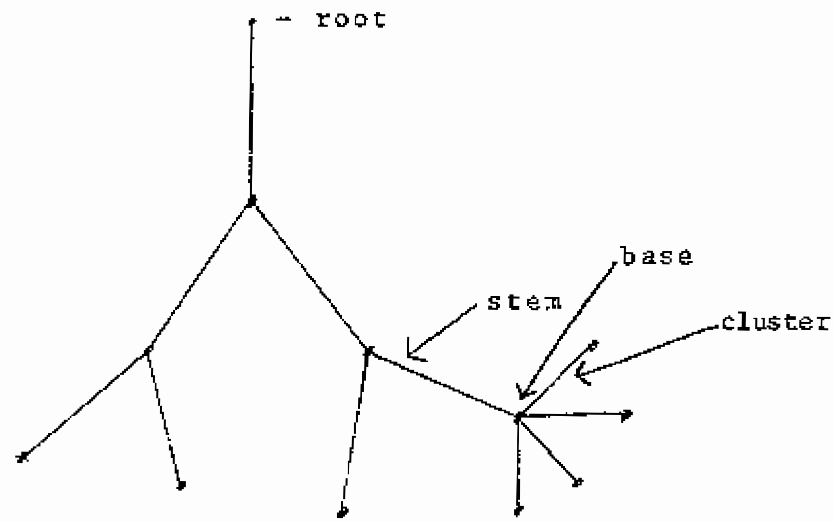

Figure 3 - A rooted tree, with a cluster, base and stem

The process of eliminating a cluster can be devided into three phases. We now briefly describe the operations performed in each.

Phase I - Labeling the Spokes of a Cluster

Let C(x) be a cluster of the current tree. The essense of this phase lies in realizing how the supply-demand nature of a given spoke e = (x,y) of C(x) is fully captured by assuming that all but almost two points of the spoke are neutral. The spoke e = (x,y) is assigned one or two labels: S(e), D(e). The interpretation of the labels is as follows. The label S(e) is associated with a supply point located on e at a distance S(e) away from the base x. A positive label D(e) corresponds to a demand point located on e at a distance D(e) from x. A negative label D(e) corresponds to a demand point located on e at a distance D(e) from x. A negative label D(e) corresponds to a facility already located on e, at a distance -D(e) units away from x. Except for the (at most two) points indicated by the labels, and possibly the base x, all other points of the spoke e are assumed neutral. Further, the spoke length can be reduced to yield

$$d(x,y) = \max\{S(e), |D(e)|\}.$$

The labels assigned to a spoke e = (x,y) satisfy conditions 1-5 below. These conditions are enforced throughout the execution of the feasibility test.

1) At least one of the labels S(e), D(e) is actually assigned

2) $-d(x,y) \leq D(e) \leq d(x,y)$, $D(e) \neq 0$ if D(e) is assigned

$0 < S(e) \leq d(x,y)$ if S(e) is assigned

3) Either $|D(e)| = d(x,y)$

or $S(e) = d(x,y)$

4) $|D(e)| \leq r$ if D(e) is assigned

5) If both D(e) and S(e) are assigned, then

$$r \geq D(e) > 0$$

and

$$S(e) - D(e) \leq r$$

Condition 1 asserts that strictly neutral spokes are eliminated as soon as they are detected. Condition 2 means that the points indicated by the labels (whether supply or demand) lie within the semi open interval (x,y]. Condition 3 implies that at least one of those points lie on the tip y. Condition 4 asserts that there is no need to carry information about a demand point or a facility located at a distance of more than r from x, for otherwise, we have either infeasibility (+D case) or the -D label can be ignored as it cannot cover any demand point of $\overline{T}$. Condition 5 refers to a spoke labeled as D and S simultaneously. Its first part asserts that in such a case, the D label is positive, i.e., corresponds to a demand point and not to a facility. For consider a spoke which contains both a facility and a supply point. If the facility is closer to the base than the supply point, then the latter can be obviously eliminated. In the opposite case, we can replace the spoke with two new spokes one containing only the supply point and the other containing just the facility. The second part of the condition, together with condition 4, implies that if a spoke contains both a demand and a supply point, then the latter must cover the former. Otherwise, the supply point can be ignored, as in the first part of this condition. The reader can note that conditions 4 and 5 imply that

$$\max \{|D(e)|, S(e)\} \leq 2r$$

Thus, at no point in the test will we have to deal with spokes whose length is larger than 2r.

We show in the appendix how a spoke which is an original edge of T can be brought to the form indicated by conditions 1-5. In the next two phases, we indicate how these conditions can be maintained as the algorithm progresses.

Phase II - Labeling the Base of a Cluster

Given a cluster C(x) all of whose spokes have been labeled, we proceed to handle the cluster as a whole. We call this operation labeling the base x of C(x). As will be demonstrated shortly, during this process we eliminate the vast majority of the spokes of this cluster.

It is useful to distinguish between different types of spokes depending on their labels. We say that a spoke is S-labeled if it is assigned an S label. It is S-only labeled if it is S-labeled but no D label is assigned to it. A spoke is +D-labeled (-D-labeled) if its D label is assigned a positive (negative) label, and it is +D only labeled (-D only labeled) if it is +D labeled (-D labeled) but not S-labeled. Finally, a spoke is S-D labeled if it is both S-labeled and +D-labeled. (Condition 5 ensures that a spoke cannot be both S-labeled and -D labeled. Hence, a -D labeled spoke is also -D only labeled.)

Propositions 3 and 4 below allow us to eliminate from C(x) all but at most one spoke of each type. The proof of proposition 3 is given in the appendix. Proposition 4 is immediate and its proof is omitted.

For $x \in N$, denote by $\alpha(x)$ a point in Σ which is closest to x. In particular, if $x \in \Sigma$, $\alpha(x) = x$.

In propositions 3-6 below, ties can be broken arbitrarily.

Proposition 3. Consider a cluster C(x) whose spokes are labeled. We can eliminate without affecting the validity of the test, all the S labels on spokes of C(x) (whether on S only or on S-D labeled spokes) accept perhaps on one spoke which contains $\alpha(x)$. The D only labeled spokes which result from deleting the S labels of S-D labeled spokes may be either of the + or -D type.

Proposition 4. Consider a cluster C(x) whose spokes are labeled. We can eliminate without affecting the validity of the test, all the +D only labeled spokes except for one, whose D label is maximal. Also, we can eliminate all the -D only labeled spokes except the one whose absolute value is minimal and all the S only labeled spokes except the one whose S label is minimal.

Applying propositions 1 and 2 to a given cluster, we are left with at most one spoke of each type. The following two propositions allow us to eliminate additional spokes from C(x). Proof of these propositions are left for the appendix.

Proposition 5. Let $x \in \Sigma$. Then we can eliminate, without affecting the validity of the test, all spokes of C(x) but possibly one. Further, this spoke is either +D or -D labeled only.

Proposition 6. Let $x \notin \Sigma$. Then we can eliminate, without affecting the outcome of the test, all spokes of C(x) but possibly at most two. Furthermore, these two spokes contain (together) at most two labels; one S and one + or -D.

Phase III Expanding the Cluster

After the propositions of the second phase are applied, the given cluster consists of at most two spokes. In the third phase we eliminate the cluster C(x) by expanding it in the direction of its stem, say (x,t). Generally speaking, this is done by replacing C(x) and its stem (x,t) by a new cluster whose spokes are obtained by attaching a neutral copy of the stem (x,t) to each of the spokes of C(x). (Recall that C(x) has at most two spokes.) This operation is demonstrated graphically in Figure 4.

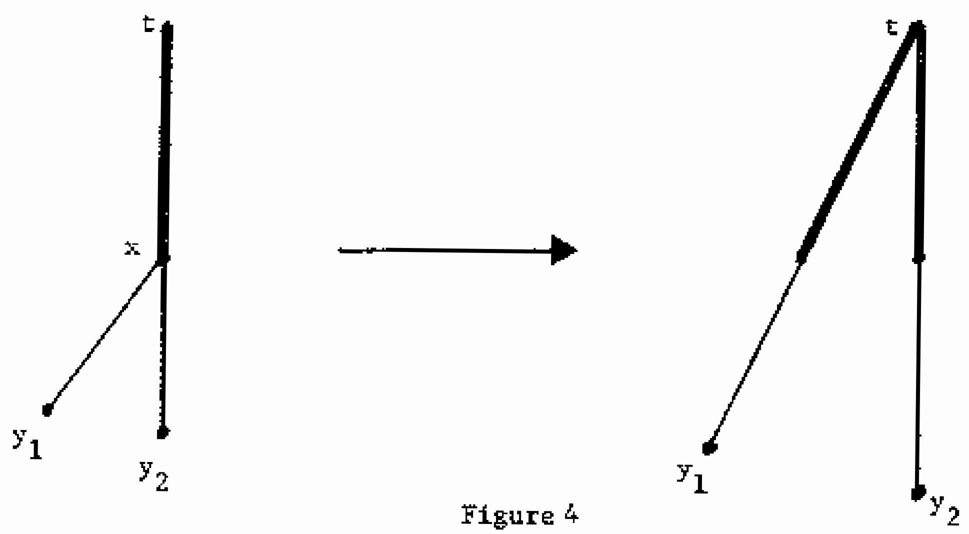

Figure 4

Obviously, such an operation may change some distances between pairs of points of $\tilde{T}$. Specifically, if C(x) has two spokes, then distances between points located on the different spokes are increased by 2d(x,t). Clearly, such an increase is of no importance unless the points in question belong to $\bar{\Sigma}$ and $\bar{\Delta}$ respectively. However, it is shown in the appendix (Remark to Proposition 6) that in this case the two points in question were not within covering distance to begin with. Thus, increasing the distance between them does not affect the validity of the test.

The operation of expanding C(x) can be implemented in practice by updating the labels of the expanded spokes which are now connected to the node t.

The details of the updating routine differ somewhat from case to case depending on the supply demand nature of the original stem (x,t). A precise description of this operation is given in the appendix.

Applying successively the three phases of the test to the clusters of $\overline{T}$, we eventually exhaust all the edges of T and can stop. By keeping a record of the number of facilities established, and comparing this number to p we can decide the feasibility of r. (Of course, we can stop the procedure with a negative answer as soon as the number of facilities established exceed p). The running time of the procedure is established in the following:

Proposition 7.

The test run in time $O(n)$ (including calculating the set $\{\alpha(x),\ x \in N\}$)

Proof: The computation of the sequence $\{\alpha(x)\}$ can be performed in $O(n)$ time by a double scanning of the original tree T. First scan the tree from its tips to the root and find for each node x the closest supply point to x among the descendents of x on the rooted tree T. Then, by scanning back from the root to the tips the sequence $\{\alpha(x)\}$ is found.

To show that the rest of the procedure can be implemented also within the $o(n)$ bound we first note that the time needed to handle any given edge of T (either as a spoke or as a stem of cluster) can be bounded by a constant. We emphasise that this constant is independent of the number of facilities which get established on this edge. This follows the fact that if several facilities are assigned to a certain edge they are located in a periodic pattern at a distance of 2r from each other. Thus, all which has to be specified is the number of these facilities, and the location of the last (i.e. the one closest to the base). Finally we note that the total number of "new edges" created by the algorithm is bounded by 2n, since whenever we expand a cluster it contains at most 2 edges. Thus the total number of edges treated by the test is no more than 3n.

VI. An Algorithm for the p Center Problem

The set R described in section IV. and the feasibility test of section V, form a basis for an algorithm for the p-center problem. As pointed out in section III, a straight forward binary search algorithm will require $O(n^2)$ steps if any one of the sets Σ or Δ contains only discrete points, and $O(n^2p)$ if both contain at least one edge. The bulk of the effort involved corresponds to generating the set R. We now briefly discuss a data structure introduced in [9] which enables us to reduce the computational requirements.

Given a tree T, we show in [9] how the set of inter nodal distances of T can be partitioned into $O(n\log n)$ subsets such that the kth longest element in each subset can be found in constant time. The effort involved in this partition is $O(n\log^2 n)$.

Let $R = \left\{\frac{d(x,y)}{j}\right\}$, $x,y \in N$, $j \in M = \{1,\ldots m\}$ be the set described in section IV. We can use the partition of intermodal distances of T to create a similar partition of the set R by replicating each subset of the former set m times. The number of subsets thus created is $O(n\log n)$ if m=2 and $o(np\log n)$ if m=2p.

Let the subsets which make the partition of R be $R_1 \ldots R_k$. At each itteration, let $R_i^{\sim} \subseteq R_i$ be the set of remaining variables with $R^{\sim} = \bigcup_{i=1}^{k} R_i^{\sim}$. Let $r_i \in R_i^{\sim}$ be the median element in this set and let r be the weighted median element of the set $\{r_i\}_{i=1\ldots k}$, where the weight of r_i is $|R_i^{\sim}|$. The effort needed to find r is $o(k)$. Furthermore, r is known to lie somewhere within the two central quartiles of $R^{\sim}$. Thus, once the feasibility test is performed on r, we can eliminate from $R^{\sim}$ at least one quarter of its elements. Running the test requires $o(n)$ steps and updating the set R requires $o(k)$ steps. Thus, the dominating effort at each iteration is of order $o(k)$.

Let $T(\ell)$ be the effort required to search a set containing ℓ elements. We then have

$$T(\ell) \leq C \cdot k + T(3/4\ \ell)$$

Thus, this approach yields an algorithm of complexity $O(n\log^2 n)$ for the case where $m=2$ (i.e. $k=O(n\log n)$) and $o(np\log^2 n)$ for the case $m=2p$, ($k=O(np\log n)$). For the second case, i.e. when both Σ and Δ contain full edges of T, we can sometime do better following a different approach. This approach is particularly advantageous for large values of p. Here one computes explicitly the set of internodal distances on T and sort this set entirely. The set R can be then described as $m=2p$ duplications of this set. Using the same search strategy on the set R organized in this manner yields an algorithm whose overall complexity is $O(n^2\log p)$. Thus, the "continuous" version of the p- center problem (i.e. the case where $\Sigma \not\subseteq N$, $\Delta \not\subseteq N$), can be solved in $O(n \min\{n\log p, p\log^2 n\}$.

Appendix

We give below the missing details needed to implement the test procedure and supply the missing proofs.

I. Labeling Spokes which are edges of the Original Tree.

Consider a cluster $C(x)$ and let $e=(x,i) \varepsilon C(x)$ be an edge of T. Then exactly one of the following conditions is met:

(i) Each point on e, but possibly x or i, is neither a supply point nor a demand point.

(ii) Each point of e including x and i is a supply and demand point.

(iii) Each point of e is a demand point, but none (with a possible exception of x and i), is a supply point.

(iv) Each point of e is a supply point, but none (with a possible exception of x and i), is a demand point.

Starting with type (i) spoke $e = (x,i)$ we note that if the tip i is neither a supply nor a demand point, we can eliminate e from the tree. The same conclusion follows if i is a supply, but not demand point with $d(x,i) > r$. By eliminating a spoke $e = (x,i)$, we mean here, as in the following paragraphs, eliminating the semi-open interval (x,i). Thus, the base x remains unaffected by such a deletion. Consider now a tip i which is a supply, but not a demand point, and such that $d(x,i) \leq r$. In this case, we leave e as is. We indicate such a spoke configuration by assigning a supply label $S(e) = d(x,i)$ to this spoke. The situation with respect to a demand only tip i is similar. If $d(x,i) > r$, the test is infeasible. If $d(x,i) \leq r$, we leave e as is and indicate this situation by assigning a demand label D9e) = $d(x,i)$. Finally, suppose that i is a supply and demand point. If $d(x,i) > r$, we set a facility at i and delete the spoke (x,i). If $d(x,i) \leq r$, we leave the arc as is. This will require assigning both a supply and a demand label to the arc e, with $S(e) = D(e) = d(x,i)$.

Next we consider a type (iv) arc. If i is not a demand point we can eliminate e as x clearly dominates any point on the semi open interval $(x,i]$ as a potential supply location. Consider, then, a spoke $e = (x,i)$ with $i \in \Delta$. If $d(x,i) \leq r$, then again the supply region in the semi open interval $(x,i]$ can be ignored as x clearly dominates any point in this interval. Thus, we can treat such a spoke as if it were a neutral spoke with a demand only tip. As shown in the previous paragraph such a spoke is assigned a demand label $D(e) = d(x,i)$ but no supply label. In the case $d(x,i) > r$, we must position a facility on e in order to cover the demand point i. It should be evident that the best location for such a facility is at a distance r from i. The demand point on e has been now covered and can be ignored. If $d(x,i) > 2r$, then, clearly, the facility on e can be also ignored and the spoke can be eliminated. Otherwise, the spoke can be shortened up to the location of its facility, i.e., up to a length of $0 < d(x,i) - r \leq r$. We indicate a spoke of this type (i.e., a neutral spoke with a facility at its tip) by a negative-demand label $D(e) = -(d(x,i) - r)$.

Now, let $e = (x,i)$ be of type (iii). Assume first that i is not a supply point. If $d(x,i) > r$, then the test is obviously infeasible. If $d(x,i) \leq r$, we can ignore the demand points in the semi open interval $[x,i)$, as any facility assigned to cover i will cover this interval as well. We label such an arc by setting $D(e) = d(x,i)$. Next, consider the case $i \in \Sigma$. If $d(x,i) \leq r$, we can again ignore all the demand points on e except for i. Thus, e is treated as a neutral spoke with a tip which is both a supply and a demand point. We label such a spoke by $D(e) = S(e) = d(x,i)$. If $d(x,i) > 2r$, then there is no way to cover the demand region on e in its entirety and the problem is infeasible. We are left with the case $r < d(x,i) \leq 2r$. In this case, we must set a facility at i. This facility will cover some of the demand on e and will therefore leave us with a demand only arc of length $d(x,i) - r$. As mentioned previously, we label wuch a spoke by setting $D(e) = d(x,i) - r$.

Finally, we consider a type (ii) arc. If $d(x,i) > 2r$, then at least one facility must be established on e. Obviously, the best position for the first (from the direction of i) such facility is at a distance r from i. Thus, we can reduce the length of e by 2r. In general, let $d(x,i) = k_i(2r) + b_i$ with $0 < b_i \leq 2r$, $k_i \geq 0$, integer. We then set k_i facilities on e at a distance of 2r from each other and reduce the length of e to b_i. Thus, consider a spoke with $d(x,i) \leq 2r$. If $d(x,i) \leq r$, then the supply region in the semi open interval $(x,i]$ can be ignored. We thus set $D(e) = d(x,i)$. If $r < d(x,i) \leq 2r$, we must place a facility on e. The best position of this facility is at distance r from i. Thus, we label the spoke e by $D(e) = -(d(x,i) - r)$.

II. <u>Proof of Proposition 3</u>.

Let $y \neq \alpha(x)$ any supply point located on an S-D labeled spoke of C(x), say $e = (x,i)$. Let z be any demand point on $\bar{T}$ but not on the semi open interval $(x,i]$. The n

$$d(\alpha(x),z) \leq d(\alpha(x), x) + d(x,z) \leq d(y,x) + d(x,z) = d(y,z)$$

Thus, $\alpha(x)$ is (weakly) closer to any demand point accept perhaps to the demand point on e. There are two cases to consider:

(i) $d(\alpha(x),x) + D(e) \leq r$

i.e. $\alpha(x)$ covers the demand point on e as well. In this case x dominates y as a possibly location and we can eliminate y (i.e. drop the label s(e)). This will convert e into D-only labeled spoke.

(ii) $d(\alpha(x),x) + D(e) > r$

In this case there is no supply point, save for y, which can cover the demand point on e. Thus we must establish a facility on y. The spoke e is then transformed into a -D only labeled spoke.

If $y \neq \alpha(x)$ is a supply point on an S-only labled spoke then the same dominance arguments hold and spoke can be eliminated.

III. Proof of Proposition 5.

If $x \in \Sigma$, then $x=\alpha(x)$. Thus, by proposition 3, we can eliminate all the supply labels from spokes of $C(x)$. Applying proposition 4 leaves us with at most two spokes one +D only and the other -D only. Call these spokes e and e' respectively. There are two cases to consider:

(i) $D(e) + |D(e')| \leq r$

i.e. the facility covers the demand point. Thus the letter can be ignored.

(ii) $D(e) + |D(e')| > r$

In this case the demand point of e' must be covered from another facility which is closer to x than $D(e')$. Thus, any (potential) demand point which is covered by the facility on e' will be also covered by this other facility. Consequently, the edge e' can be ignored.

IV. Proof of Proposition 6.

By propositions 3-4 we are left with at most one each of +D only, -D only, S-D, and S only labeled spokes. Furthermore, we have at most one S label.

(i) By the same arguments involved in the proof of proposition 3, we may assume that we have at most one out of the +D only and the -D only labeled spokes.

(ii) Assume a cluster which contains both an S-D and a +D only labled spokes, say e and e' respectively. Clearly we must have

$$D(e') + S(e) = D(e\ ) + d(\alpha(x),x) \leq r$$

or else we have detected infeasibility. We can thus replace the pair e, e' by a unique S-D labled spoke e'' with $S(e'') = S(e)$, $D(e'')= \max \{D(e'), D(e)\}$.

(iii) Assume finally a cluster containing both an S-D and a -D only labeled spokes, e and e`. Obviously if

$$|D(e`)| + D(e) \leq r$$

Then the demand label on e can be dropped. In the opposite case the -D label of e` can be dropped by the same reasoning used in the proof of proposition 5.

Remark

We note that if C(x) contains an S only and a D only labled spokes say e and e`, then

$$S(e) + D(e`) > r$$

Otherwise, we can replace e and e` by one spoke, e``, with D(e``) = D(e`), S(e``) = S(e).

V. The Expansion Routine.

Consider first the case $x \notin \overline{\Sigma}$. This implies that the stem (x,t) is either neutral or is a demand type arc.

The case of a neutral stem is the simplest to dispose of: (i) Assume first that $x \notin \overline{\Lambda}$. Then each spoke of C(x) is enlarged by a distance d(x,t) and the base is transformed from x to t. Enlarging the spoke can be done by incrementing each of the labels on the spokes of C(x) by d(x,t). The only exception is a -D label where the absolute value, rather than the label itself, is increased. We note that the operation of expansion may result in labels which are too long and thus violate some of the conditions 1-5. These violations can be taken care of as follows: If the new label of an S-only labeled or -D only labeled spoke is now larger than r, that spoke can be removed. If the new label of a +D only labeled spoke is larger than r the test yields infeasibility. Now suppose the spoke is S-D labeled. If the new D label is greater than r, a facility is established at the supply point corresponding to the S-label, the D label becomes negative with $D(e) = -S(e)$, and the S label is omitted. If also $|D(e)| > r$ the entire spoke is removed.

(ii) Assume now that $x \in \bar{\Delta}$. If $C(x)$ contains a D only labeled spoke (positive or negative), we can assume with no loss of generality that $x \notin \bar{\Delta}$ and proceed as in (i). If $C(x)$ contains an S-D labeled spoke e, then $S(e) + D(e) \leq r$ (Property 5). In particular, $S(e) \leq r$, which implies that x is covered by any supply point covering the demand point on e corresponding to the D label. Thus, again we may assume $x \notin \bar{\Delta}$, and proceed as in (i). Thus, suppose $C(x)$ contains no D-labeled spokes, and let e be the (unique) spoke of $C(x)$. e is S only labeled. We enlarge this spoke by a distance of $d(x,t)$ and transform the base from x to t. Then we set $S(e) \leftarrow S(e) + d(x,t)$, $D(e) \leftarrow d(x,t)$ and continue as in (i) above.

We now turn to examine a demand type stem (x,t) (still assuming $x \notin \bar{\Sigma}$). This case can be treated essentially like the previous one with some extra care given to making sure that the demand on (x,t) is fully covered.

We first note that if any of the demand points within $C(x)$ will be covered by a facility outside $C(x)$ then this facility will cover (x,t) in its entirety as well. Thus, if there is a +D only labeled spoke in $C(x)$, the demand on (x,t) need not be further considered. Also, if $C(x)$ contains an S-D labeled spoke e with $S(e) + d(x,t) \leq r$, then any facility assigned to cover the demand point on e, whether this point is on e, or some other point outside $C(x)$, will cover (x,t) as well. Finally we note that if there exists a -D labeled spoke, e, which covers t, (i.e. such that $|D(e)| + d(x,t) \leq r$)), then the demand on (x,t) is already covered by this facility, e. In all those cases we can proceed as if the stem (x,t) is completely neutral.

Thus suppose that C(x) does not contain a +D only labeled spoke, or an S-D labeled spoke covering node t, or a -D only labeled spoke covering t.

Let

$$\delta_1(x) = \begin{cases} |D(e)|, & \text{if } C(x) \text{ contains a -D only labeled spoke } e. \\ r & \text{otherwise} \end{cases}$$

$$\delta_2(x) = \begin{cases} S(e) & \text{if } C(x) \text{ contains an S-D labeled spoke } e \\ r & \text{otherwise} \end{cases}$$

Note that $\delta_2(x)$ is well defined since C(x) has at most one S-D labeled spoke (Property 5). Let $\gamma(x) = \min\{\delta_1(x), \delta_2(x)\}$. Then, $\gamma(x) \leq r$. We observe that all demand points on (x,t) which are at a distance less than or equal to $r - \gamma(x)$ from x (except for the point x itself in the case where C(x) has no D labeled spoke), can be ignored since they are either covered by an existing facility or by a facility that will eventually cover some demand point of C(x). Observing that under our assumptions $\gamma(x) + d(x,t) > r$, the demand points on the upper segments of (x,t) of length $\delta = d(x,t) - (r - \gamma(x)) > 0$ will either be covered by a facility not on C(x) or by a facility on an S only labeled spoke of C(x), if one exists. If C(x) does not contain an S only labeled spoke we append a +D only labeled spoke to C(t), whose +D label is set equal to δ. (If $\delta >$ r the test yields infeasibility.) Thus except for the addition of the artificial +D only labeled spoke to C(t), we can treat the arc (x,t) as if it were neutral.

Now suppose that C(x) contains an S-only labeled spoke, e. If $S(e) + d(x,t) \leq r$, set $D(e) = \delta$, and $S(e) \leftarrow S(e) + d(x,t)$. The other spokes of C(x) are treated as if the arc (x,t) were neutral. If $S(e) + d(x,t) > r$, then we append a +D only spoke to C(t), whose +D label is equal to $d(x,t) + S(e) - r$. Also e becomes S-D labeled with $D(e) = \delta$, $S(e) \leftarrow S(e) + d(x,t)$, and the other spokes of $C(x)$ are treated as if (x,t) were neutral. This completes the discussion for a cluster whose base x satisfies $x \notin \bar{\Sigma}$.

Consider now a cluster $C(x)$ such that $x \in \bar{\Sigma}$. We have already seen that such a cluster contains exactly one spoke and that this spoke is either +D only or -D labeled. Let the spoke in question be $e = (x,y)$.

(i) Suppose first that the stem (x,t) is neutral. If also $d(x,t) > r$, delete the spoke if e is -D only labeled; set a facility at x, and delete the spoke, if e is +D only labeled. Let $d(x,t) \leq r$. If e is +D labeled set $S(e) = d(x,t)$ and $D(e) \leftarrow D(e) + d(x,t)$. If e is -D labeled and $D(e) + d(x,t) > r$, delete the D label and set $S(e) = d(x,t)$. Finally if e is -D labeled and $D(e) + d(x,t) \leq r$, the spoke e is now replaced by two spokes of $C(t)$. The first, e_1, is a -D only labeled spoke with $D(e_1) = D(e) + d(x,t)$. The second e_2, is S only labeled with $S(e) = d(x,t)$.

(ii) Next we assume that the stem (x,t) is a demand only arc. Without loss of generality we also assume $d(x,t) \leq 2r$, since otherwise the test yields infeasibility. Suppose that e is -D only labeled. If $d(x,t) + |D(e)| \leq r$ the demand points of (x,t) are covered by the facility located on e and we treat this case as if (x,y) were neutral. Hence suppose $d(x,t) + |D(e)| > r$. The spoke e of $C(x)$ now becomes a spoke of $C(t)$ with $S(e) = d(x,t)$ and $D(e) \leftarrow d(x,t) + |D(e)| - r$.

(a) If $D(e) > 2r$ the test yields infeasibility.

(b) If $r < D(e) \leq 2r$, set a facility at x, delete $S(e)$ and set $D(e) = d(x,t) - r$. (If this new label is greater than r the test yields infeasibility.)

(c) Let $D(e) \leq r$. If also $S(e) \leq r$, we leave the labels as are. Otherwise, ($D(e) \leq r$ but $S(e) > r$), we replace e by two spokes of C(t). The first, e_1, is +D only labeled with $D(e_1) = S(e) - r$. The second, e_2, is S-D labeled with $S(e_2) = S(e)$ and $D(e_2) = D(e)$.

Now let e be +D only labeled, while still assuming (x,t) to be a demand only arc. If $d(x,t) > r$, set a facility at x, and replace the spoke e by a +D only labeled spoke of C(t) with $D(e) = d(x,t) - r$. (If $D(e) > r$ the test yields infeasibility). Let $d(x,t) \leq r$. Then replace e by an S-D spoke of C(t) with $D(e) \leftarrow D(e) + d(x,t)$ and $S(e)=d(x,t)$.

(iii) Now we assume that the stem (x,t) is a supply only arc. First let e be a +D only labeled spoke. Then if $D(e) + d(x,t) \leq r$, replace e by a +D only labeled spoke of C(t) with $D(e) \leftarrow D(e) + d(x,t)$. If $D(e) + d(x,t) > r$, establish a facility at a distance r from y and replace e by -D only labeled spoke of C(t) with $D(e) \leftarrow -(d(x,t) + D(e) - r)$. (Delete this new spoke if $|D(e)| > r$.)

Assume that e is a -D only labeled spoke. If $|D(e)| + d(x,t) \leq r$ replace e by a -D only labeled spoke of C(t) with $D(e) \leftarrow D(e) - d(x,t)$. Otherwise, delete the spoke e and arc (x,t) from the tree.

(iv) Finally we treat the case where the stem (x,t) is both a supply and demand arc. First let e be -D only labeled. If $|D(e)| + d(x,t) \leq r$ replace e by a -D only labeled spoke of C(t) with $D(e) \leftarrow D(e) - d(x,t)$. Otherwise, ($|D(e)| + d(x,t) > r$), replace e by a supply and demand spoke of C(t) whose length is $d(x,t) + |D(e)| - r$. This spoke can be now labeled as was described above for the spokes of the original tree.

Now let e be +D only labeled. If $D(e) + d(x,t)$ r this spoke becomes a +D only labeled spoke, with $D(e)$ $D(e) + d(x,t)$. Otherwise replace e by a supply and demand spoke of $C(t)$ whose length is $d(x,t) + D(e)$. Again we can label this spoke as was indicated previously.

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 Management Science, Northwestern University, 1978.

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

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

[7] S.L. Hakimi, "Optimal Location of Switching Centers and the Absolute Centers and Medians on a Graph," Operations Research, Vol. 12, pp. 450-459 (1964).

[8] O. Kariv and S.L. Hakimi, "An Algorithmic Approach to Network Location Problems, Part I - p-centers" SIAM J. Appl. Math. 37 (1979) pp. 513-538.

[9] N. Megiddo, A. Tamir, E. Zemel and R. Chandrasekaran, "An $O(n \log^2 n)$ Algorithm for the k-th Longest Path in a Tree with Applications to Location Problems," Discussion Paper No. 379, Center for Mathematical Studies in Economics and Management Science, Northwestern University, 1979.