Partitioning of Biweighted Trees

Alessandro Agnetis,[1] Pitu B. Mirchandani,[2] Andrea Pacifici[3]

[1] *Dipartimento di Ingegneria dell'Informazione, Università di Siena, Siena, Italy*

[2] *Systems and Industrial Engineering Department, The University of Arizona, Tucson, Arizona 85721*

[3] *Dipartimento di Informatica, Sistemi e Produzione, Università di Roma "Tor Vergata," Roma, Italy*

Received September 2000; revised July 2001; accepted 23 August 2001

Abstract: A districting problem is formulated as a network partitioning model where each link has one weight to denote travel time and another weight to denote workload. The objective of the problem is to minimize the maximum diameter of the districts while equalizing the workload among the districts. The case of tree networks is addressed and efficient algorithms are developed when the network is to be partitioned into two or three districts. © 2002 Wiley Periodicals, Inc. Naval Research Logistics 49: 143–158, 2002; DOI 10.1002/nav.10003

Keywords: location; network partitioning; bicriteria optimization

1. INTRODUCTION

This paper addresses the problem of districting a network for mobile response units that respond to continuous demands on the links, which we refer to as Highway Patrol Districting (HPD) problem. Given a network, the HPD problem is to determine a partition into p districts, each district being served by a mobile unit. Each link of the network has two weights: One indicates the length of the link (how far apart are the end points of the link), and the second weight is proportional to the workload on that link. The maximum response time of a unit in a district is proportional to the diameter of the district, that is, the maximum distance between two points of the district. Our problem is to minimize the maximum response time while balancing the workloads among the districts.

The motivation for the problem arises from the need to allocate response units patrolling on a highway, responding to incidents such as traffic accidents, helping stalled motorists, catching speed-law violators, and assisting in alleviating nonrecurring unusual congestion. In particular, the Arizona Department of Public Safety has been interested in designing response districts and allocating highway patrol response units to the districts. Arizona divides its highways into fourteen "districts." Each district is patrolled by a number of patrol units. Location of traffic incidents are distributed along the links of the network. Based on historical data on traffic volumes and traffic

Correspondence to: P.B. Mirchandani

accidents, it is reasonable to model the network into link segments so that the probability of an accident is uniformly distributed along each link. This results in the second weight on each link that denotes its workload.

Handler and Mirchandani [8] were the first to introduce location models for continuous link demands, where the objective is to either minimize average travel time (Median Problems) or minimize the maximum travel time (Center Problems). For tree networks they give an exact method, and for general networks they consider an approximate discretized version of continuous demands on links. For general networks Chiu [3] provides exact and heuristic approaches for the 1-median location problem with continuous link demands. Several works, by Sherali et al., deal with special cases of (both capacitated and uncapacitated) median problem on chain networks and trees [1, 14, 15, 16] with a continuum of demand placed on the links. Since "supply" of a facility relates to its maximum workload, the capacitated version of these problems closely relate to our problem. Sherali [14] and Sherali and Nordai [15, 16] consider the case when the total capacity of all facilities equals total demand and the problem is to minimize average travel time. Sherali and Rizzo [17] also investigated the *unbalanced* case, when the total supply does not equal total demand. Kim, Sherali, and Park [9] consider the scenario of an emergency/patrol car traveling along a simple path of a road network while maintaining surveillance of a set of facilities for possible service response; demands may arise discretely on the nodes and/or continuously along the links of the network. They address the minimum objective such as the problem of finding a path that minimizes the weighted sum of distances and the minimax objective of minimizing the farthest weighted distance between the mobile facility and demands during the travel period of the facility. We note, however, that the problem of balancing workloads does not play a major role in these papers.

As the p-Center problem (in particular the "absolute" p-center problem, see [7]) can be stated as the minimization of the maximum diameter in a partition, it relates to our problem at hand. The p-Center problem has been well-studied in the literature. It is NP-complete for p not fixed. However, polynomial algorithms exist that solve the same problem on trees for general values of p [8, 2, 12]. Related center problems of locating paths [13] and subtrees [10] on trees to minimize the maximum distance to demands have also been investigated. In particular, in a recent work, Halman and Tamir [5] study a general class of min–max problems of continuous tree partitioning problems into components (subtrees) where the size of a component relates to its diameter or its length.

We consider an undirected tree network $T = (V, E, w, l)$ embedded on a Euclidean plane (in the following, simply, tree). V is the set of vertices or nodes, E is the set of links or arcs, and w and l are two integer vectors whose components are the workload w_{uv} and the length l_{uv} for all the links $uv \in E$, respectively. The ratio $\rho_{uv} = w_{uv}/l_{uv}$ will be called *density* of the link. Since we assume the workload to be uniformly distributed along each link, the workload of any connected portion of a link uv is given by ρ_{uv} times the length of the portion. We use the term *point of* T to denote either a vertex or an intermediate point of a link of T.

DEFINITION 1: A *subforest* F of T is any set, possibly nonconnected, of vertices and links or fractions of links of T.

For instance, F_2, in Figure 2(a), is a subforest. Note that we are using a continuous extension of the concept of subforest F, allowing the links to be divided in any point, so that one link of T may not entirely belong to subforest F; i.e., F is formed by a collection of intervals of points in the continuum set of points of the edges of T. By a *leaf* we mean an extreme point of a subforest, either it be a node or an intermediate point. We denote by $W(F)$ the total workload of subforest

F. Given two points x and y in F, we indicate by $P(x, y)$ and $d(x, y)$ the path that connects those two points and its total length, respectively. It may happen that $P(x, y)$ is not entirely contained in F.

DEFINITION 2: Given a subforest F of a tree T, let $P(t_1, t_2), t_1, t_2 \in F$, be a path such that $d(t_1, t_2) = \max\{d(x, y) : x \in F, y \in F\}$. We refer to

- value $d(t_1, t_2)$ as the *diameter* of F and denote this value by $D(F)$,
- path $P(t_1, t_2)$ as a *diametrical path* of F, and
- pair of nodes (t_1, t_2) as a pair of *diametrical endpoints* for F.

Note that t_1 and t_2 are necessarily leaves of T.

DEFINITION 3: Given a diametrical path $P(t_1, t_2)$ of a tree T, let $r_1, \ldots, r_s$ be its internal nodes, and let $S_1, \ldots, S_s$ be the s disjoint subtrees rooted at $r_1, \ldots, r_s$, respectively, obtained when the links of the diametrical path $P(t_1, t_2)$ are removed from T. We call S_i a *diametrical subtree* and r_i its *root*, for all $i = 1, 2, \ldots, s$.

Figure 1 illustrates Definition 3. In this paper we always use this definition with reference to a diametrical path of the whole tree T.

DEFINITION 4: A *p-partition* $\pi = \{F_1, F_2, \ldots, F_p\}$ of T is a collection of p disjoint (possibly nonconnected) subforests whose union is the whole tree. Each subforest is called *district*.

Note that, given π, a link may be partially included in a district F_i, the rest of the link being assigned to other districts in π. [See Fig. 2(a): There, e.g., district F_2 is formed by the two intervals between point Q and node 7 of link $\{6, 7\}$ and between point R and node 8 of link $\{6, 8\}$. In this case the district is not connected.] Let $D_\pi := \max_{i=1,\ldots,p}\{D(F_i)\}$. The problem of finding the minimum diameter $D_p^*(T)$ of a p-partition of T:

$$D_p^*(T) = \min\{D_\pi : \pi \text{ is a } p\text{-partition of } T\}$$

is the well-known p-center problem, on trees. As we already mentioned, it is possible to find $D_p^*(T)$ in polynomial time as long as T is a tree (e.g., Frederikson and Johnson [4] provided an $O(|V| \log |V|)$ algorithm.) We use D_p^* instead of $D_p^*(T)$ whenever this does not generate confusion.

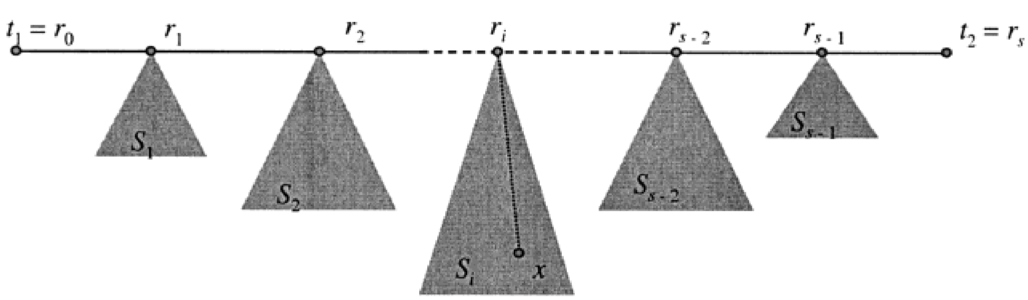

Figure 1. Diametrical path of T and diametrical subtrees.

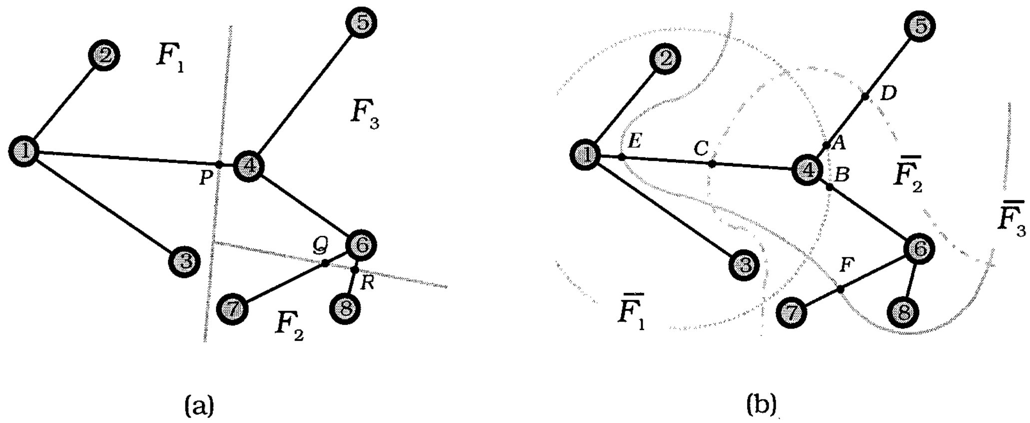

Figure 2. A 3-partition and a 3-cover

DEFINITION 5: A p-partition $\pi = \{F_1, F_2, \ldots, F_p\}$ of T is *balanced* if, for all $i = 1, 2, \ldots, p, W(F_i) = W(T)/p$.

Clearly, there always exists a balanced p-partition of a network (possibly having nonconnected districts). The problem addressed in this paper is the following:

PROBLEM 6: Given a tree network $T = (V, E, w, l)$, find a p-partition π such that $W(F_i) = \frac{1}{p}W(T)$, for all $i = 1, \ldots, p$ and D_π is minimized.

It is useful to introduce the related decision problem:

PROBLEM 7: Given a tree network $T = (V, E, w, l)$ and a rational $\bar{D}$ such that $D_p^*(T) \leq \bar{D} \leq D(T)$, find, if it exists, a p-partition π such that $W(F_i) = \frac{1}{p}W(T)$, for all $i = 1, \ldots, p$ and $D_\pi = \bar{D}$.

Clearly, if we are able to solve Problem 7, we can solve also Problem 6 by performing a binary search over $\bar{D}$. However, for $p = 2$ we will see that this is not necessary.

In order to solve Problem 7, it is convenient to introduce the concept of p-cover which is formally defined below.

DEFINITION 8: A *p-cover* $\chi = \{\bar{F}_1, \bar{F}_2, \ldots, \bar{F}_p\}$ of T is a collection of p connected subforests, named *superdistricts*, such that: (i) $\cup_{i=1}^{p} \bar{F}_i = T$; (ii) each superdistrict $\bar{F}_i$ has diameter $D(\bar{F}_i) = D_\chi$, for all $i = 1, \ldots, p$; (iii) each $\bar{F}_i$ is maximal, i.e., it is not strictly contained in any other subforest having the same diameter D_χ.

Figure 2(b) illustrates Definition 8: Here $\bar{F}_1$ contains links $\{1, 2\}$, $\{1, 3\}$, and $\{1, 4\}$ plus the intervals $\{4, A\}$ and $\{4, B\}$ of links $\{4, 5\}$ and $\{4, 6\}$, respectively. The basic idea of our approach is that we first find a suitable p-cover, and then we carve the p districts out of the p-cover.

The above definition implies that superdistricts are connected. Possibly, we may have intervals of points belonging to more than one superdistrict [e.g., in Fig. 2(b), intervals of points between node 4 and points A, B, and C belong to all the three superdistricts; point E to both $\bar{F}_1$ and $\bar{F}_3$, but not to F_2].

Given a p-partition $\pi = \{F_1, F_2, \ldots, F_p\}$ and a p-cover $\chi = \{\bar{F}_1, \bar{F}_2, \ldots, \bar{F}_p\}$ such that $F_i \subseteq \bar{F}_i$ for all $i = 1, 2, \ldots, p$, we say that π is *contained* in χ, indicated as $\pi \subseteq \chi$. This is the case for the partition and the cover of Figure 2(a) and (b), respectively.

DEFINITION 9: A p-cover χ of T is *balanced* if it contains a balanced partition.

The paper is organized as follows. In Section 2, we characterize feasible solutions to Problem 7. In Section 3, we analyze Problems 6 and 7 when $p = 2$. In Section 4, we analyze problem 7 when $p = 3$. In the Appendix a notation table is given.

2. EXISTENCE OF A BALANCED SOLUTION

In this section we present a characterization of balanced p-covers. In the subsequent sections, efficient algorithms to actually find such a cover for $p = 2$ and $p = 3$ will be described.

The following Theorem 10 gives necessary and sufficient conditions for a p-cover with fixed diameter $D_\chi = \bar{D}$ to be balanced and consequently for a p-partition with diameter $D_\pi \leq \bar{D}$ to be balanced.

THEOREM 10: Given a tree $T = (V, E, w, l)$, a rational $\bar{D}$ such that $D_p^* \leq \bar{D} < D(T)$, and a p-cover $\chi = \{\bar{F}_1, \bar{F}_2, \ldots, \bar{F}_p\}$ with diameter $D_\chi \leq \bar{D}$, then χ is balanced if and only if, for any subset I of the integer numbers $\{1, 2, \ldots, p\}$, the following condition holds:

$$\frac{|I|}{p} W(T) \leq W\left(\bigcup_{i \in I} \bar{F}_i\right). \tag{1}$$

PROOF: Let $I = \{i_1, \ldots, i_r\}, 1 \leq r \leq p$ be a nonempty subset of the first p integers and $\bar{I}$ be its complement $\bar{I} = \{1, \ldots, p\} \backslash I$.

(Only if.) By Definition 9, if there is a balanced p-cover $\chi = \{\bar{F}_1, \ldots, \bar{F}_p\}$ with diameter $D_\chi \leq \bar{D}$ then there is also a balanced p-partition $\pi = \{F_1, \ldots, F_p\} \subseteq \chi$ having diameter $D_\pi \leq \bar{D}$. For any $I \subseteq \{1, 2, \ldots, p\}$, from Definition 5 of balanced partition, as districts are disjoint and each one has weight $\frac{1}{p} W(T)$, we have $W(\cup_{i \in I} F_i) = \frac{|I|}{p} W(T)$. Moreover, as $W(\cup_{i \in I} \bar{F}_i) \geq W(\cup_{i \in I} F_i)$, necessity is proved.

(If.) We next show that if a cover χ exists, having diameter $D_\chi \leq \bar{D}$, and satisfying Eq. (1), then it is possible to find a balanced partition $\pi \subseteq \chi$, with diameter $D_\pi \leq \bar{D}$. Given a p-cover $\chi = \{\bar{F}_1, \bar{F}_2, \ldots, \bar{F}_p\}$ of T, it is possible to determine $2^p - 1$ disjoint, some possibly empty, subforests of T (see Definition 1) in one-to-one correspondence with all the possible nonempty subsets of $\{1, 2, \ldots, p\}$. In particular, letting subset $I = \{i_1, \ldots, i_r\} \subseteq \{1, 2, \ldots, p\}$, we associate with I a subforest H_I defined as all the points that belong exclusively to all the superdistricts in I; i.e.,

$$H_I = \left(\bigcap_{i \in I} \bar{F}_i\right) \Bigg\backslash \left(\bigcup_{i \in \bar{I}} \bar{F}_i\right)$$

and we set $w_I = w(H_I)$. (Figure 3 shows a Venn diagram representing the case $p = 3$.) Any district F_i of any p-partition π contained in χ satisfies the following condition: $F_i \subseteq \bar{F}_i = \cup_{I \in i} H_I$. Therefore, points in subforest $H_I, I = \{i_1, \ldots, i_r\}$, may only belong to r districts of partition $\pi \subseteq \chi$, namely, $F_{i_1}, \ldots, F_{i_r}$. Hence, to design a balanced partition $\pi \subseteq \chi$ we need to "share"

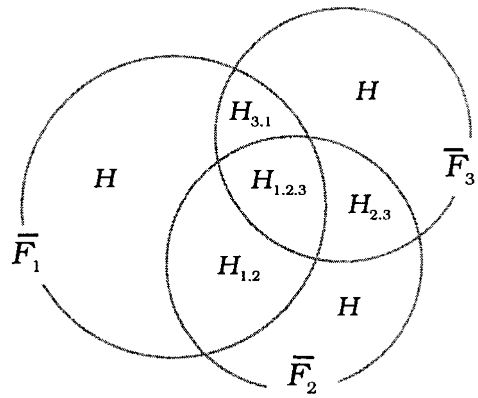

Figure 3. Definition of the seven possible H_I and w_I when $p = 3$.

the workload of H_I, for all I, among districts $F_{i_1}, \ldots, F_{i_r}$ of π, in such a way that the workloads of the resulting districts are equal.

For this purpose, we define the following supply-demand problem. We associate districts $F_1, \ldots, F_p$, to a set of p consumers and the $2^p - 1$ distinct subforests H_I's to a set of $2^p - 1$ suppliers. Supplier/subforest H_I has capacity w_I and supplies district/consumer F_i if and only if $i \in I$. Demands for all the p consumers are set equal to $\frac{1}{p}W(T)$.

Clearly, χ and π are balanced (see Definitions 5 and 9) if and only if the above problem has a feasible solution. In fact, district F_i is obtained by augmenting subforest $H_{\{i\}} = \bar{F}_i \backslash \{\cup_{j \neq i} \bar{F}_j\}$ by means of a suitable portion of $\bar{F}_i \cap \{\cup_{j \neq i} \bar{F}_j\}$. A flow δ from supplier H_I to consumer F_i $(i \in I)$ in the feasible solution of the supply–demand problem indicates that district F_i draws δ units of workload from subforest H_I. A well-known result in network flow theory [11] states that a supply–demand problem has a feasible solution if and only if, for any subset I of consumers, the total demand $(\frac{|I|}{p}W(T))$ does not exceed the total supply that can be sent to I. In our case the total amount of commodity which can be sent to I is given by

$$\sum_{\substack{J \subseteq \{1,\ldots,p\} \\ J \cap I \neq \emptyset}} w_J = W\left(\bigcup_{j \in I} \bar{F}_j\right) \tag{2}$$

This completes the proof. □

3. TWO DISTRICTS

In this section, we consider the case in which T must be partitioned into two districts. We first address Problem 7, and then turn to Problem 6.

3.1. Balanced 2-Covers

In this section we address Problem 7 for $p = 2$.

LEMMA 11: Given a tree $T = (V, E, w, l)$ and a rational $\bar{D}$ such that $D_2^* \leq \bar{D} < D(T)$, let t_1 be a diametrical endpoint of T and let u and v be such that $d(t_1, u) \leq \bar{D}$ and $d(t_1, v) \leq \bar{D}$. Then, $d(u, v) \leq \bar{D}$.

PROOF: Consider a diametrical path, as in Definition 3, $P(t_1, t_2)$ of T (see Fig. 1). For any diametrical subtree S_i of $P(t_1, t_2)$ rooted in r_i, consider a point $x \in S_i$. It is easy to see that

$$d(r_i, x) \leq \min\{d(t_1, r_i), d(t_2, r_i)\}. \tag{3}$$

Consider the set $\bar{F}_1$ of points at distance $\bar{D}$ or less from t_1. We show that the distance between any two points in $\bar{F}_1$ does not exceed $\bar{D}$. Let u and v be any two points in $\bar{F}_1$. Suppose first that u and v belong to two diametrical subtrees, rooted in r_u and r_v, respectively (possibly $r_u = r_v$). Without loss of generality let $d(t_1, r_u) \leq d(t_1, r_v)$. From Eq. (3), $d(u, r_u) \leq d(t_1, r_u)$, which implies $d(u, v) \leq d(u, r_u) + d(r_u, v) \leq d(t_1, v) \leq \bar{D}$. □

LEMMA 12: Given a tree $T = (V, E, w, l)$, and a rational $\bar{D}$ such that $D_2^* \leq \bar{D} < D(T)$, there is a unique cover $\chi = \{\bar{F}_1, \bar{F}_2\}$ with $D_\chi = \bar{D}$.

PROOF: Consider the set $\bar{F}_1$ $(\bar{F}_2)$ of points at distance $\bar{D}$ or less from t_1 (t_2). Since, for $i = 1, 2$, the set $\bar{F}_i$ contains any superdistrict with diameter $\bar{D}$ including t_i (otherwise there would a pair of points (t_i, y) of the superdistrict with $d(t_i, y) > \bar{D}$), from Lemma 11, $\bar{F}_1$ and $\bar{F}_2$ are superdistricts including t_1 and t_2 respectively, both having diameter $\bar{D}$ (recall Definition 8).

We next show that $\chi = \{\bar{F}_1, \bar{F}_2\}$ is the only 2-cover having diameter $\bar{D}$. By contradiction, suppose there is another cover $\chi' = \{\bar{F}_1', \bar{F}_2'\} \neq \chi$, having diameter $\bar{D}$, where $\bar{F}_1'$ and $\bar{F}_2'$ are the superdistricts obtained by starting from a different pair of diametrical endpoints $(t_1', t_2') \neq (t_1, t_2)$ where, with no loss of generality, we let $t_1 \neq t_1'$. If $t_2 = t_2'$, by definition of superdistrict, we have $\bar{F}_2 = \bar{F}_2'$ and therefore $\bar{F}_1 \neq \bar{F}_1'$. If $t_2 \neq t_2'$ and $\bar{F}_1 = \bar{F}_1'$, then it must be $\bar{F}_2 \neq \bar{F}_2'$. By exchanging the role of 1 and 2 we get $t_1 \neq t_1'$ and $\bar{F}_1' \neq \bar{F}_1$. We may henceforth reduce to the case where $t_1 \neq t_1'$ and $\bar{F}_1' \neq \bar{F}_1$.

Note that $t_2 \notin \bar{F}_1$, since $\bar{D} < D(T)$. Also, either t_1 or t_2 is in $\bar{F}_1'$. Without loss of generality, let $t_1 \in \bar{F}_1'$.

It cannot be that $\bar{F}_1 \subseteq \bar{F}_1'$ for otherwise $\bar{F}_1$ would not be maximal or $D(\bar{F}_1') > \bar{D}$ (and analogously, $\bar{F}_1' \not\subseteq \bar{F}_1$). Therefore, suppose that a point $x \in \bar{F}_1 \backslash \bar{F}_1'$ exists. (We can handle the case $x \in \bar{F}_1' \backslash \bar{F}_1$ similarly.) As $x \notin \bar{F}_1'$, it must be $d(t_1', x) > \bar{D}$ and hence, since $d(t_1, x) \leq \bar{D}$, Lemma 11 implies $d(t_1, t_1') > \bar{D}$, which contradicts the fact that $t_1 \in \bar{F}_1'$. □

For $p = 2$ Theorem 10 becomes:

THEOREM 13: Given a tree $T = (V, E, w, l)$, and a rational $\bar{D}$ such that $D_2^* \leq \bar{D} < D(T)$, the 2-cover $\chi = \{\bar{F}_1, \bar{F}_2\}$ with diameter $\bar{D}_\chi = \bar{D}$ is balanced if and only if

$$\min\{W(\bar{F}_1), W(\bar{F}_2)\} \geq \frac{1}{2}W(T). \tag{4}$$

The previous Theorem 13 suggests a straightforward algorithm for Problem 7. Given a tree network $T = (V, E, w, l)$ and a rational $\bar{D}$, the following algorithm finds a balanced 2-partition $\pi = \{F_1, F_2\}$ of T, with diameter $D_\pi = \bar{D}$, or it concludes it does not exist.

1. Determine $\bar{F}_1, \bar{F}_2$ (by marking the points at distance $\bar{D}$ from the diametrical endpoints t_1 and t_2, respectively), and compute $W(\bar{F}_1)$ and $W(\bar{F}_2)$.
2. If Condition (4) does not hold, there is no balanced solution, else
3. Let $F_1 := (\bar{F}_1 \backslash \bar{F}_2) \cup F$, where F is any subforest such that $F \subseteq \bar{F}_1 \cap \bar{F}_2$ and $W(F) = \frac{1}{2}W(T) - W(\bar{F}_1 \backslash \bar{F}_2)$.

4. $F_2 := T \backslash F_1$.

Note that all the steps can be carried out in time $O(|V|)$ by easily adapting any tree-visit algorithm. The latter algorithm can be naturally employed for solving Problem 6 using a binary search for the minimum $\bar{D}$. If we consider that $D_2^* \le \bar{D} \le D(T)$, the overall complexity becomes $O(|V| \log(D(T) - D_2^*))$. In the following section we present an alternative procedure which may turn out to be more efficient.

3.2. Finding the Minimum Diameter

Let $P(t_1, t_2)$ be a diametrical path of T and $\bar{D}$ be a rational such that $D_2^*(T) \le \bar{D} < D(T)$. Consider the set $F_1(\bar{D})$ of points at distance at most $\bar{D}$ from t_1. We define the *bordering set* $B_1(\bar{D})$ as a special set of nodes of T:

$$B_1(\bar{D}) = \{v \in V \text{ such that for all } uv \in E, d(t_1, u) \le \bar{D}, d(t_1, v) > \bar{D}\}.$$

In other words, the nodes in $B_1(\bar{D})$ are those immediately out of subforest $\bar{F}_1(\bar{D})$ as we move away from t_1.

For a given $\bar{D}$, we say that a link $uv \in E$ *crosses the border*, if $d(t_1, u) \le \bar{D}$ and $d(t_1, v) > \bar{D}$. We denote by $\rho_1(\bar{D})$ the sum of the densities of the links crossing the border. As $\bar{D}$ increases, $\rho_1(\bar{D})$ represents the current marginal increment of the workload $W(\bar{F}_1)$.

The algorithm 2_DISTRICTING (illustrated in Table 1) builds superdistrict F_1 including at each step the nearest $v \in B_1$ until $W(F_1) = \frac{1}{2}w(T)$.

THEOREM 14: The algorithm 2_DISTRICTING solves Problem 6 for $p = 2$, in time $O(|V| \log |V|)$.

Table 1. Algorithm 2_DISTRICTING.

Procedure 2_DISTRICTING
Input Network $T = (V, E, w, l)$;
Output A minimum diameter balanced 2-partition $\pi = \{F_1, F_2\}$ of T, and D_π;

1. Compute:
   1.1. A pair of diametrical endpoints of T, t_1 and t_2
   1.2. The value D_2^*, e.g., as in [6]
   1.3. The subforests F_1 and F_2 as the loci of points at distances at most D_2^* from t_1 and t_2, respectively
   1.4. Values $W(F_1)$ and $W(F_2)$, $B_1(\bar{D})$, and $\rho_1(\bar{D})$.
2. Initialize $\bar{D} = D_2^*$ and $\delta := 0$.
3. If $W(F_1) \le \frac{1}{2}W(T)$ goto step 4;
   else if $W(F_2) \le \frac{1}{2}W(T)$ exchange the role of t_1 and t_2 and goto step 4;
   else set $D_\pi = D_2^*$ and adjust F_1 and F_2 by suitably assigning points in $F_1 \cap F_2$ to only one of the two district and goto step 7.
4. Repeat the following:
   4.1. $\bar{D} := \bar{D} + \delta, W(\bar{F}_1) := W(\bar{F}_1) + \rho_1(\bar{D})\delta$;
   4.2. $\delta := \min_{v \in B_1(\bar{D})}\{d(t_1, v) - \bar{D}\}$;
   until $W(\bar{F}_1) + \rho_1(\bar{D})\delta \ge \frac{1}{2}W(T)$
5. $\delta := \frac{1}{\rho_1(\bar{D})}(\frac{1}{2}W(T) - W(\bar{F}_1)); \bar{D} := \bar{D} + \delta$. [$\delta$ is the adjustment required to $\bar{D}$ in order to reach a workload of $\frac{1}{2}W(T)$: note that, due to step 3, δ cannot be negative.]
6. Set F_1 as the locus of points at distance at most $\bar{D}$ from t_1, $F_2 := T \backslash F_1$, and $D_\pi = \bar{D}$.
7. Return $\bar{F}_1$ and $\bar{F}_2$.

PROOF: Let us consider Step 1 first. Any value not greater than D_2^* would be a suitable initial value for $\bar{D}$. In particular, D_2^* may be computed in $O(|V|)$ time (for instance, as in [6]) and this choice does not increase the overall procedure complexity. By performing a tree visit, we can compute all the distances $d(t_1, v)$ in $O(|V|)$. Thereafter, we can rank the nodes in nondecreasing order of $d(t_1, v)$ in $O(|V| \log |V|)$. The set $B_1(\bar{D})$ can be computed simply by checking the distance $d(t_1, v)$ of all nodes v adjacent to nodes u such that $d(t_1, u) \leq \bar{D}$. If $d(t_1, v) > \bar{D}$, we put v in $B_1(\bar{D})$. Concurrently, the workload $W(F_1)$ [where F_1 is the locus of points u of T such that $d(t_1, u) \leq \bar{D}$], and the initial value of $\rho_1(\bar{D})$ can be computed. The latter computations cost $O(|V|)$.

Let us evaluate the computation cost of updating $B_1(\bar{D})$ and $\rho_1(\bar{D})$ in Step 4. As $\bar{D}$ increases, $B_1(\bar{D})$ varies. Every time this step is executed, at least one node leaves $B_1(\bar{D})$. Consider the set of nodes s such that $\delta = d(t_1, s) - \bar{D}$. Let *us* be the arc incident s such that $d(t_1, u) < d(t_1, s)$. All the nodes v such that $sv \in E$ enter $B_1(\bar{D})$. We maintain $B_1(\bar{D})$ ordered with respect to distances from t_1, and hence each insertion requires $O(\log |V|)$ time. Notice that if we consider all the executions of Step 4.2, no node ever enters $B_1(\cdot)$ twice, and since $B_1(\cdot)$ is ordered, δ can be computed in constant time in Step 4.2, so that the overall complexity of Step 2 is $O(|V| \log |V|)$ and the thesis follows. □

4. THREE DISTRICTS

In this section, we consider the case in which T must be partitioned into three districts. We limit ourselves to addressing Problem 7. As discussed in Section 3.1 for the case $p = 2$, the latter algorithm can be exploited, combined with a binary search over the possible values of superdistricts diameter $\bar{D}$, in order to get an algorithm solving Problem 7.

4.1. Balanced 3-Covers

From Theorem 10, we derive the following conditions:

THEOREM 15: Given a tree $T = (V, E, w, l)$, a rational $\bar{D}$ such that $D_3^* \leq \bar{D} < D(T)$, the 3-cover $\chi = \{\bar{F}_1, \bar{F}_2, \bar{F}_3\}$ having diameter $D_\chi = \bar{D}$, is balanced if and only if

$$W(\bar{F}_i) \geq \frac{1}{3}W(T), \qquad i = 1, 2, 3, \tag{5}$$

$$W(\bar{F}_i \cup \bar{F}_j) \geq \frac{2}{3}W(T), \qquad i, j = 1, 2, 3, \quad i \neq j. \tag{6}$$

Hereafter, we show how to efficiently compute a balanced 3-cover $\chi = \{\bar{F}_1, \bar{F}_2, \bar{F}_3\}$ of T with $D_\chi = \bar{D}$, if it exists. First of all notice that if $P(t_1, t_2)$ is a diametrical path of T (recall Definition 3), and a 3-cover $\chi = \{\bar{F}_1, \bar{F}_2, \bar{F}_3\}$ exists having diameter $\bar{D}$, with $\bar{D} < D(T)$, then t_1 and t_2 belong to two different superdistricts, say $t_1 \in \bar{F}_1$ and $t_2 \in \bar{F}_2$. As shown in Lemma 12, for $i = 1, 2$, $\bar{F}_i$ is the set of points of T whose distance from t_i does not exceed $\bar{D}$, then $\bar{F}_1$ and $\bar{F}_2$ are unique. Moreover, it is clear that $T \backslash \{\bar{F}_1 \cup \bar{F}_2\} \subseteq \bar{F}_3$. Let $\bar{T}$ be the tree obtained connecting the components of $T \backslash \{\bar{F}_1 \cup \bar{F}_2\}$ using the links of T. More formally, $\bar{T} = \{\cup P(x, y) : x, y \in T \backslash \{\bar{F}_1 \cup \bar{F}_2\}\}$. Hence, $\bar{F}_3$ must also contain $\bar{T}$.

In the next section we describe a procedure for finding a balanced 3-cover (recall Definition 9) of diameter $\bar{D}$. Basically, once $\bar{F}_1$ and $\bar{F}_2$ are determined, $\bar{F}_3$ must be such that $\bar{T} \subseteq \bar{F}_3$ and

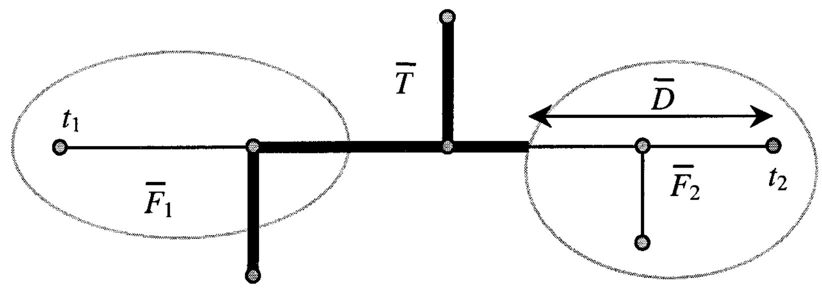

Figure 4. The tree $\bar{T}$.

the conditions of Theorem 15 are satisfied. By Definition 8 of p-cover, $\bar{F}_3$ is a maximal subforest having diameter $D(\bar{F}_3) = \bar{D}$ and therefore there is a single point at distance at most $\bar{D}/2$ from all the points of $\bar{F}_3$, which is its 1-center. Thus, $\bar{F}_3$ is completely defined once we locate its 1-center x.

Our algorithm first identifies the candidate set C of points where the 1-center x of $\bar{F}_3$ can be placed in order to have $\bar{T} \subseteq \bar{F}_3$. The algorithm efficiently enumerates the solutions obtained by locating the 1-center of $\bar{F}_3$ in a point $x \in C$ and evaluates whether the conditions of Theorem 15 hold. The procedure stops as soon as those conditions are satisfied or it concludes that no balanced 3-cover having diameter $D_\chi \leq \bar{D}$ exists.

4.2. Locating the Center of $\bar{F}_3$

Observe that if e is a leaf of $\bar{T}$, then either e is a leaf of $T\backslash\{\bar{F}_1 \cup \bar{F}_2\}$ or e is a point at distance $\bar{D}$ from t_1 or t_2 on a diametrical path $P(t_1, t_2)$. (See Fig. 4.) We denote by C the set of points of T at distance at most $\bar{D}/2$ from all the points in $\bar{T}$. Clearly, the 1-center of $\bar{F}_3$ must belong to C.

LEMMA 16: C is connected.

PROOF: If $C \neq \emptyset$, let x and y be two points in C and $z \in P(x, y)$. For all $t \in \bar{T}$ we have $d(z, t) \leq \max\{d(x, t), d(y, t)\} \leq \bar{D}/2$. Hence z belongs to C. □

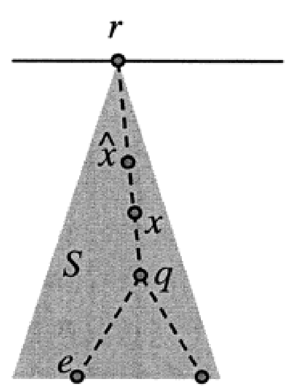

Figure 5. Illustration of Theorem 17.

Consider the set $K = P(t_1, t_2) \cap C$ [where $P(t_1, t_2)$ is a diametrical path of T]. In what follows, if $K = \emptyset$, we denote by $\hat{x}$ the point of C that is closest to $P(t_1, t_2)$. Given a point $x \in T$, we denote by $\bar{F}_3(x)$ the superdistrict having 1-center at x, i.e., $\bar{F}_3(x)$ is the maximal subforest of T with diameter $D(\bar{F}_3) = \bar{D}$ having its 1-center at x.

THEOREM 17: If a balanced 3-cover $\chi = \{\bar{F}_1, \bar{F}_2, \bar{F}_3\}$ of T exists having diameter $\bar{D}$, then there is one such that $t_1 \in \bar{F}_1, t_2 \in \bar{F}_2$, and

- if $K \neq \emptyset$, the 1-center of $\bar{F}_3$ is located on the diametrical path $P(t_1, t_2)$,
- if $K = \emptyset$, the 1-center of $\bar{F}_3$ is located at $\hat{x}$.

PROOF: Suppose $\chi = \{\bar{F}_1, \bar{F}_2, \bar{F}_3\}$ is a balanced 3-cover of T having diameter $\bar{D}$, such that $t_i \in \bar{F}_i, i = 1, 2$. Suppose that the 1-center of $\bar{F}_3$ is a point q not lying on $P(t_1, t_2)$. Let r be the root of the diametrical subtree S containing q, and let e be a deepest leaf of S, i.e., a node such that $d(r, e) = \max_{x \in S}\{d(r, x)\}$. Let x be any point of $P(q, r)$ such that $x \in C$ (possibly $x = q$). Since C is connected, $P(q, x) \subseteq C$ (see Fig. 5). Again, $\bar{F}_3(x)$ is defined as the maximal subforest of T with diameter $D(\bar{F}_3) = \bar{D}$ having its 1-center in x (while the center of $\bar{F}_3$ is q). In order to show that $\bar{F}_3 \subseteq \bar{F}_3(x)$, we consider an *arbitrary* point $y \in \bar{F}_3$ and show that $y \in \bar{F}_3(x)$. We consider two subcases:

(i) $d(y, q) < d(y, x)$. This implies that $y \in S$ and $d(r, y) > d(r, x)$. Suppose first that $S \cap \bar{T} \neq \emptyset$. In this case, e belongs to $\bar{T}$. In fact, if $e \in \bar{F}_1 \backslash \bar{T}$ (or $e \in \bar{F}_2 \backslash \bar{T}$), then the whole subtree S belongs to $\bar{F}_1$ (or $\bar{F}_2$) and therefore $S \cap \bar{T} = \emptyset$. Since $e \in \bar{T}$ and since $x \in C, \bar{F}_3(x)$ covers e and therefore point y. Suppose now $S \cap \bar{T} = \emptyset$. This means that S is completely covered by either $\bar{F}_1$ or $\bar{F}_2$. Without loss of generality, let us suppose S is covered by $\bar{F}_1$. If $\bar{F}_3(x)$ does not cover point y, then $d(r, e) \geq d(x, y) > \bar{D}/2$. On the other hand, since $P(t_1, t_2)$ is a diametrical path, $d(r, t_1) \geq d(e, r)$. But this implies that $d(t_1, e) > \bar{D}$, contradicting the hypothesis that $S \subseteq \bar{F}_1$.

(ii) $d(y, q) \geq d(y, x)$. In this case, since $y \in \bar{F}_3$, then $d(y, q) \leq \frac{1}{2}\bar{D}$ and hence $d(y, x) \leq \frac{1}{2}\bar{D}$, i.e., $y \in \bar{F}_3(x)$.

Since $\bar{F}_3 \subseteq \bar{F}_3(x), W(\bar{F}_3) \leq W(\bar{F}_3(x))$, and $W(\bar{F}_3 \cup \bar{F}_i) \leq W(\bar{F}_3(x) \cup \bar{F}_i)$ for $i = 1, 2$. This means that moving the 1-center of $\bar{F}_3$ from q towards the diametrical path still yields a balanced cover. If $K \neq \emptyset$, we can choose $x = r$. If $K = \emptyset$, note that, since C is connected, $P(\hat{x}, q) \subseteq C$, and, due to the definition of $\hat{x}, \hat{x} \in P(q, r)$. Hence, we can choose $x = \hat{x}$ and this completes the proof. □

4. 3. An Algorithm for Finding a Balanced 3-Cover

The set K can be characterized as the set of points of $P(t_1, t_2)$ whose distance from any leaf of $\bar{T}$ does not exceed $\frac{1}{2}\bar{D}$. If $K \neq \emptyset$, it is easy to pinpoint the two extremes of K along the diametrical path: Call them k_1 and k_2. In a similar fashion to what was done for $p = 2$, we will move the 1-center of $\bar{F}_3$ from k_1 to k_2 along K until either the conditions of Theorem 15 are met or k_2 is reached.

Given a point $x \in K$, we need again to define the bordering set of points of the superdistrict $\bar{F}_3(x)$ centered at x. This time it is convenient to distinguish the bordering sets of the superdistrict $\bar{F}_3(x)$ lying on the two opposite sides of the tree with respect to x. To this aim, consider the two subtrees obtained by cutting T at x. Let T_1 (T_2) be the subtree containing k_1 (k_2). We define

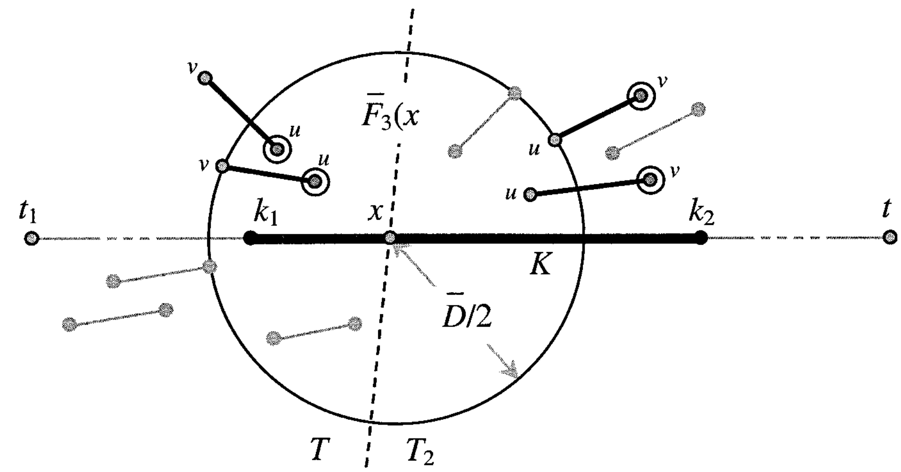

Figure 6. Illustration of bordering sets. The circled u and v nodes are the left and right bordering sets, respectively.

the *left bordering set* $B_3^L(x)$ as the set of nodes lying immediately *inside* $\bar{F}_3(x)$ on the T_1 side, i.e., $B_3^L(x) = \{u \in T_1 | uv \in E, d(x,u) < \frac{1}{2}\bar{D}, d(x,v) \geq \frac{1}{2}\bar{D}\}$. We say that a link $uv \in T_1$ *crosses the left border* of $\bar{F}_3(x)$ if $d(x,u) < \frac{1}{2}\bar{D}$ and $d(x,v) \geq \frac{1}{2}\bar{D}$. Analogously, we define the *right bordering set* $B_3^R(x)$ as the set of nodes lying immediately *outside* $\bar{F}_3(x)$ on the T_2 side, i.e., $B_3^R(x) = \{v \in T_2 | uv \in E, d(x,u) \leq \frac{1}{2}\bar{D}, d(x,v) > \frac{1}{2}\bar{D}\}$, and we say that a link $uv \in T_2$ *crosses the right border* of $\bar{F}_3(x)$ if $d(x,u) \leq \frac{1}{2}D$ and $d(x,v) > \frac{1}{2}\bar{D}$. (See Fig. 6.) Finally the sum of the densities of the links of $B_3^L(x)$ is denoted by $\rho_3^L(x)$, and of $B_3^R(x)$ by $\rho_3^R(x)$. The algorithm returns a balanced cover of diameter $\bar{D}$, if it exists. Starting from k_1, as the 1-center x of $\bar{F}_3(x)$ moves towards k_2, we check the current values of $W(\bar{F}_3(x)), W(\bar{F}_1 \cup \bar{F}_3(x)), W(\bar{F}_2 \cup \bar{F}_3(x))$. Actually, the algorithm takes advantage of the following straightforward fact:

LEMMA 18: If all the links have positive workloads, then as x moves from k_1 to $k_2, W(\bar{F}_1 \cup \bar{F}_3(x))$ is increasing and $W(\bar{F}_2 \cup \bar{F}_3(x))$ is decreasing.

Unlike $W(\bar{F}_1 \cup \bar{F}_3(x))$ and $W(\bar{F}_2 \cup \bar{F}_3(x))$, as x moves from k_1 to $k_2, W(\bar{F}_3(x))$ increases or decreases depending on the relative weights of the links which $\bar{F}_3(x)$ reaches (on the side of T_2) or leaves (on the side of T_1). See Figure 7.

We can define two points $x_1 \in K$ and $x_2 \in K$ such that $W(\bar{F}_1 \cup \bar{F}_3(x_1)) = \frac{2}{3}W(T)$ and $W(\bar{F}_2 \cup \bar{F}_3(x_2)) = \frac{2}{3}W(T)$, respectively. Due to Lemma 18, if $d(t_1, x_1) \leq d(t_1, x_2)$, then only the points of $P(x_1, x_2) \cap K$ are such that conditions (6) are satisfied. If $d(t_1, x_1) > d(t_1, x_2)$, then no feasible solution exists. As illustrated in Figure 7, the slopes of the lines representing the weights $W(\bar{F}_1 \cup \bar{F}_3(x)), W(\bar{F}_2 \cup \bar{F}_3(x))$, and $W(\bar{F}_3(x))$ as x varies are given by $\rho_3^R(x), -\rho_3^L(x)$, and $\rho_3^R(x) - \rho_3^L(x)$, respectively. In particular, $B_3^L(x)$ may change only in a discrete set of points $X_L = \{x \in K | d(u,x) = \frac{1}{2}\bar{D}, u \text{ is a node of } T_1\}$ and, analogously, $B_3^R(x)$ may change only in a discrete set of points $X_R = \{x \in K | d(x,v) = \frac{1}{2}\bar{D}, u \text{ is a node of } T_2\}$.

It is easy to see that, in order to find a point x yielding a balanced cover, we only need to consider a discrete set of values, namely $X = X_L \cup X_R \cup \{x_1, x_2\}$. The algorithm 3_COVER does this in polynomial time. The algorithm is summarized in Table 2. Here, for notation simplicity, we indicate by $x + \delta$ the point obtained moving x by δ towards k_2 along K, and by $x - \delta$ towards k_1 along K.

THEOREM 19: The algorithm 3_COVER solves Problem 7 for $p = 3$, in time $O(|V| \log |V|)$.

PROOF: By similar considerations to those in the proof of Lemma 11, the computation of points k_1 and k_2, the superdistricts $\bar{F}_1$ and $\bar{F}_2$ as well as the workloads $W(\bar{F}_1 \cup \bar{F}_3(k_1))$, $W(\bar{F}_2 \cup \bar{F}_3(k_1))$, and $W(\bar{F}_3(k_1))$ can be computed in $O(n)$ by a tree visit algorithm. Steps 2–3 locate point x_1. We move along $P(t_1, t_2)$ until $W(\bar{F}_1 \cup \bar{F}_3(k_1)) = 2/3W(T)$. Due to Lemma 18, we can proceed from k_1 towards k_2 without ever visiting the same node twice, and hence we can do it in $O(|V| \log |V|)$, since this is identical to what done for Lemma 11 for $p = 2$. Similarly, in Steps 4–5 point x_2 is located, with the same complexity. Steps 6–7 check if we can already conclude our search, and can be done in constant time. At this point, the 1-center can only be located between x_1 and x_2. Step 9 searches the interval $[x_1, x_2]$ for a point x such that the condition (5) $W(\bar{F}_3(x)) \geq \frac{1}{3}W(T)$ is met. In Step 8(b), we advance x to the next point in X. This can still be done in $O(|V| \log |V|)$. If we get to x_2 without ever satisfying condition (5), then no feasible solution exists. □

REMARK 20: The algorithm 3_COVER, combined with a binary search over possible values of $\bar{D}$, implies that Problem 6, for $p = 3$, is solvable in time $O(|V| \log |V| \log(D(T) - D_3^*))$.

5. CONCLUSIONS

Motivated by a districting problem for highway patrol units, we introduced a new network partitioning problem where the network links have two weights. We addressed the special case where the network is a tree and developed efficient algorithms to partition the network into two or three districts. Two possible extensions of the model considered in this paper are general networks and tree networks with more than three districts.

Getting efficient algorithms for general networks is much more complicated, because even without the "workload" weight, when we have the classical p-center model, the problem is NP-complete when p is not fixed.

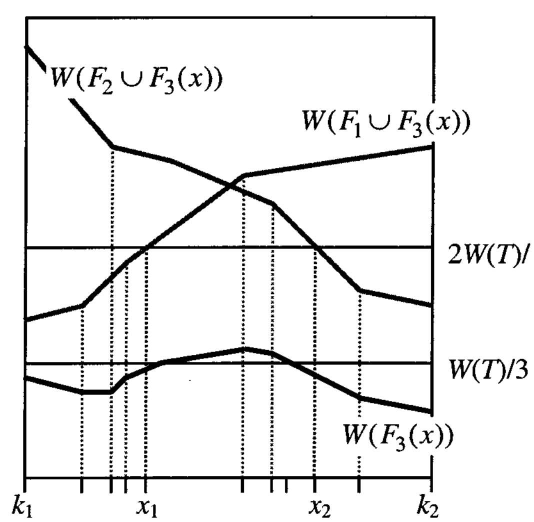

Figure 7. Workloads $W(\cdot)$ as x varies. (The dotted lines denote the discrete set of points where the slopes change.

Table 2. Algorithm 3_COVER.

Procedure FIND 3-COVER
Input Network $T = (V, E, w, l)$, diameter $\bar{D}$ such that $K \neq \emptyset$;
Output A balanced 3-cover $\chi = \{\bar{F}_1, \bar{F}_2, \bar{F}_3\}$ of T with diameter $D_\chi = \bar{D}$, if it exists;

1. Find points k_1 and k_2 along the diametrical path $P(t_1, t_2)$;
2. Compute superdistricts $\bar{F}_1, \bar{F}_2$ and the workloads $W(\bar{F}_1 \cup \bar{F}_3(k_1)), W(\bar{F}_2 \cup \bar{F}_3(k_1)), W(\bar{F}_3(k_1))$;
3. $\delta := 0, x := k_1$;
4. While $W(\bar{F}_1 \cup \bar{F}_3(x + \delta)) < \frac{2}{3}W(T)$ do
   (i) $x := x + \delta$;
   (ii) $\delta := \min_{v \in B_3^R(x)}\{d(x, v) - \frac{1}{2}\bar{D}\}$.
5. $\delta := \max\{0, \frac{1}{\rho_3^R}(\frac{2}{3}W(T) - W(\bar{F}_1 \cup \bar{F}_3(x)))\}; x_1 := x + \delta; \delta := 0; x := k_2$;
6. While $W(\bar{F}_2 \cup \bar{F}_3(x - \delta)) < \frac{2}{3}W(T)$ do
   (i) $x := x - \delta$;
   (ii) $\delta := \min_{u \in B_3^L(x)}\{\frac{1}{2}\bar{D} - d(u, x)\}$.
7. $\delta := \max\{0, \frac{1}{\rho_3^L}(\frac{2}{3}W(T) - W(\bar{F}_2 \cup \bar{F}_3(x)))\}; x_2 := x - \delta$;
8. If $d(t_1, x_1) > d(t_1, x_2)$, then STOP (no feasible solution exists);
9. If $W(\bar{F}_3(x_2) \geq \frac{1}{3}W(T)$ then STOP ($\{\bar{F}_1, \bar{F}_2, \bar{F}_3(x_2)\}$ is a balanced cover);
10. $\delta := 0; x := x_1;$ *found* := FALSE;
11. While $(x < x_2)$ and *found* = FALSE do if $W(\bar{F}_3(x)) \geq \frac{1}{3}W(T)$ then *found* := TRUE else
   (i) $\delta := \min\{\min_{v \in B_3^R(x)}\{d(x, v) - \frac{1}{2}\bar{D}\}, \min_{u \in B_3^L(x)}\{\frac{1}{2}\bar{D} - d(u, x)\}\}$;
   (ii) $x := x + \delta$;
12. If *found* = TRUE then $\{\bar{F}_1, \bar{F}_2, \bar{F}_3(x)\}$ is a balanced cover else, no feasible solution exists;
13. STOP

Restricting to tree networks with fixed p, the problem appears significantly more complex for $p > 3$ than the $p = 3$ case addressed in Section 4. In fact, the reason why we efficiently solved the $p = 3$ case is because the first two centers can be easily located, and then the third center requires only a search over a line. If $p = 4$, we can still locate the first two centers. However, the candidate set for the other two centers is not necessarily a line, but rather a subforest where *two* points are to be found. Hence, the weights w_I depend now on two parameters. It is not clear whether such 2-dimensional set can be searched in polynomial time. One might think of further decomposing this problem, and finding an effective way of obtaining exact 4-partitions. This is a challenging topic for future research. Also, another area of future research is to develop heuristics with good guaranteed worst-case bounds.

Another question concerns how to improve complexity of an algorithm for Problem 6 when $p = 3$. In Section 4 we solve Problem 7 in time $O(|V| \log |V|)$ and noticed (Remark 20) that Problem 6 can be solved in time $O(|V| \log |V| \log(D(T) - D_3^*))$.

APPENDIX: NOTATION

$T = (V, E, w, l)$: tree network, nodes, links, workload and length vectors (Sect. 1.).
$\rho_{uv} = \frac{w_{uv}}{l_{uv}}$: density of link $uv \in E$ (Sect. 1.).
$P(x, y), d(x, y)$: path and distance between two points x and y of the network (Sect. 1.).

$W(F), D(F)$: total workload and diameter of subforest F (Def. 3, Sect. 1.).
t_1, t_2 : diametrical endpoints (Def. 3, Sect. 1.).
$S_i, r_i, i = 1, 2, \ldots s$: diametrical subtrees and their roots (Def. 3, Sect. 1.).
$\pi = \{F_1, \ldots, F_p\}$: p-partition, districts (Def. 4, Sect. 1.).
$D_\pi = \max_{i=1,\ldots,p} D(F_i)$: partition diameter (Sect. 1.).
D_p^* : minimum diameter of a (possibly not balanced) p-partition, i.e., the minimum objective value of p-center problem (Sect. 1.).
$\chi = \{\bar{F}_1, \ldots, \bar{F}_p\}$: p-cover, superdistricts (Def. 8, Sect. 1.).
$D_\chi = D(\bar{F}_i), i = 1, \ldots, p$: superdistrict diameter (Def. 8, Sect. 1.).
$w_{i_1 i_2 \ldots i_k}$: weight of the portion of the tree that can be entirely shared *only* among districts $F_{i_1}, F_{i_2}, \ldots F_{i_k}$ (Sect. 2.).
$B_j, \rho_j, i = 1, 2$: bordering set of superdistrict $\bar{F}_i$, sum of densities of links crossing the border (Sect. 3.2.).
$\bar{T}$: subtree of T obtained connecting the components of $T \setminus \{\bar{F}_1 \cup \bar{F}_2\}$ (Sect. 4.1.).
C : set of points of T at distance at most $\frac{1}{2}\bar{D}$ from all the points of $\bar{T}$ (Sect. 4.2.).
$K : P(t_1, t_2) \cap C$ (Sect. 4.2.).
$\hat{x}$: closest point of C to $P(t_1, t_2)$ when $K = \emptyset$ (Sect. 4.2.).
$\bar{F}_3(x)$: set of points of T whose distance from $x \in C$ is at most $\frac{1}{2}\bar{D}$ (Sect. 4.3.).
k_1, k_2 : extremes of K (Sect. 4.3.).
T_1, T_2 : subtrees of T obtained by cutting T at $x \in K$ (Sect. 4.3.).
$B_3^L(x), \rho_3^L(x), x \in K$: left bordering set of superdistrict $\bar{F}_3(x)$, sum of densities of links crossing the left border (Sect. 4.3.).
$B_3^R(x), \rho_3^R(x), x \in K$: right bordering set of superdistrict $\bar{F}_3(x)$, sum of densities of links crossing the left border (Sect. 4.3.).
$x_1, x_2 \in K$: points of K such that $W(\bar{F}_1 \cup \bar{F}_3(x_1)) = \frac{2}{3}W(T)$ and $W(\bar{F}_2 \cup \bar{F}_3(x_2)) = \frac{2}{3}W(T)$, respectively (sect. 4.3.).

ACKNOWLEDGMENTS

The authors thank the anonymous referees for their comments and suggestions that improved the presentation of our research results. The second author acknowledges the support of the Arizona Department of Public Safety in sponsoring part of this research.

REFERENCES

[1] T.M. Cavalier and H.D. Sherali, Network location problems with continuous link demands: p-medians on a chain and 2-medians on a tree, Eur J Oper Res 23 (1986), 246–245.
[2] R. Chandrasekaran and A. Tamir, An $O((n \log p)^2)$ algorithm for the continuous p-center problem on a tree, SIAM J Algebraic Discrete Methods 1 (1980), 370–375.
[3] S.S. Chiu, The minisum location problem on an undirected network with continuous link demands, presented at the Joint National ORSA/TIMS Meeting, 1982, San Diego, CA.
[4] G.N. Frederickson and D.B. Johnson, Finding k-th paths and p-centers by generating and searching good data structures, J Alg 4 (1983), 61–80.
[5] N. Halman and A. Tamir, Continuous bottleneck tree partitioning problems, Tech. Rep., December 6, 1999, Department of Statistics and Operations Research, School of Mathematical Sciences, Tel-Aviv University, Israel.
[6] G.Y. Handler, Finding two-centers of a tree: The continuous case, Transportation Sci 12 (1978), 93–106.
[7] G.Y. Handler, "p-Center problems," Discrete Location Theory, P.B. Mirchandani and R.L. Francis (Editors), Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley, New York, 1990, pp. 305–347.
[8] G.Y. Handler and P.B. Mirchandani, Location on networks: Theory and Algorithms, MIT Press, Cambridge, MA, 1979.
[9] S. Kim, H.D. Sherali, and H. Park, Minisum and minimax of a moving facility on a network, Comput Oper Res 19(2) (1992), 123–131.
[10] R.K. Kincaid, T.J. Lowe, and T.L. Morin, The location of central structures in trees, Comput Oper Res 15(2) (1988), 103–113.

[11] L. Lovász and M.D. Plumer, Matching theory, Ann. Discrete Math. 29, North-Holland, Amsterdam, 1986, pp. 45–46.
[12] N. Megiddo and A. Tamir, New results on the complexity of p-center problems, SIAM J Comput 12 (1983), 751–758.
[13] J.G. Morris and J.P. Norback, A simple approach to linear facility location, Transportation Sci 14 (1980), 1–8.
[14] H.D. Sherali, Capacitated, balanced, sequential location-allocation problems on chains and trees, Math Program 49 (1991), 381–396.
[15] H.D. Sherali and F.L. Nordai, A capacitated, balanced, 2-median problems on a tree network with a continuum of link demands, Transportation Sci 22(1) (1988), 70–73.
[16] H.D. Sherali and F.L. Nordai, NP-hard capacitated, balanced, p-median problems on a chain graph with a continuum of link demands, Math Oper Res 13(1) (1988), 32–49.
[17] H.D. Sherali and T.P. Rizzo, Unbalanced, capacitated p-median problems on a chain graph with a continuum of link demands, Math Oper Res 13(1) (1991), 32–49.