Trade-offs on the Location of the Core Node in a Network

Jean-François Macq
CORE, Université catholique de Louvain, 34 Voie du Roman Pays, 1348 Louvain-La-Neuve, Belgium

Michel X. Goemans
M.I.T., Department of Mathematics, 77 Massachusetts Ave, Cambridge, Massachusetts 02139

We consider the problem of selecting a core node in a network under two potentially competing criteria, one being the sum of the distances to a set of terminals, the other being the cost of connecting this core node and the terminals with a Steiner tree. We characterize the worst-case trade-off between approximation ratios for the two objectives. Our results, for example, show the existence of a core node in which both objectives are simultaneously within 1.37 times their optimum value (if we were to disregard the other objective). We also consider the problem of minimizing a weighted sum of the two criteria and perform a worst-case analysis of a simple and fast heuristic, which does not need to enumerate possible core locations. This study was motivated by multimedia applications such as videoconferences or multiplayer games in which user-dependent information has to be sent from the users to a core node to be chosen (at a cost proportional to the sum of the distances from the core node), and then global information has to be multicast back from the core node to all users (at a cost proportional to the Steiner tree cost). **© 2004 Wiley Periodicals, Inc.** **NETWORKS, Vol. 44(3), 179–186 2004**

Keywords: multicast communication; core node; bicriteria; Steiner tree; 1-median

1. INTRODUCTION

Many new multimedia applications on the Internet ask for the selection of a meeting point in the network. Such applications are, for instance, videoconferences or multiplayer games. In both cases, each user sends his personal data to the meeting point. The role of this particular point is to gather the information received to create a single composite data flow, which is multicast back to the users. More generally, there exists a class of multicast routing protocols, called *center-based* multicast protocols, which require an administrative center for each multicast group. In the *Core Based Tree* (CBT) protocol ([2]), for instance, if a source node wants to reach a multicast group, the data flow is first sent to a *core* node and then distributed to the group via a *shared tree.* Note that, in general, the source does not need to be a member of the multicast group. However, in this article we focus on a particular case, called *All Receivers Sources* in [3], where the set of senders is (or can be approximated as) the set of receivers. As mentioned above, this case is relevant for several multimedia applications.

Empirical analysis has been carried out to measure the relationship between the choice of the core location and the performance of the routing scheme (see, i.e., [3], [4], [13]). The performance of such a scheme is usually evaluated in terms of delay and bandwidth consumption. For the sake of simulation, the network is modeled as an undirected graph, the communication delay between two nodes of the network is approximated as the number of links on the path used and the bandwidth consumption is evaluated as the sum of the data flow on each link.

If the objective is to minimize the average delay between every pair of users in the group, the optimal routing is achieved on a shortest path tree rooted at the core. In that case, each user-core shortest path is used in both directions, and the optimal location of the core is such as to minimize the sum of the shortest path lengths between the core and every user. In Location Theory terminology, we seek the *1-median* of the users. However, to minimize the bandwidth consumption, we must distinguish two different problems. The first one concerns the optimization of the unicast forward paths between the users and the core. Because each corresponding data flow is particular to its sender, we must optimize each user-core path. Therefore, the total bandwidth consumption of these paths is minimized when the core is located at the 1-median of the users. The second problem is

Received August 2003; accepted April 2004
Correspondence to: M. Goemans; e-mail: goemans@math.mit.edu
Contract grant sponsor: Research Fellowship of the Belgian American Education Foundation (BAEF, to J.-F.M.)
Contract grant sponsor: NSF; contract grant numbers ITR-0121495 and CCR-0098018 (to M.X.G.)
DOI 10.1002/net.20028
Published online in Wiley InterScience (www.interscience.wiley.com).
© 2004 Wiley Periodicals, Inc.

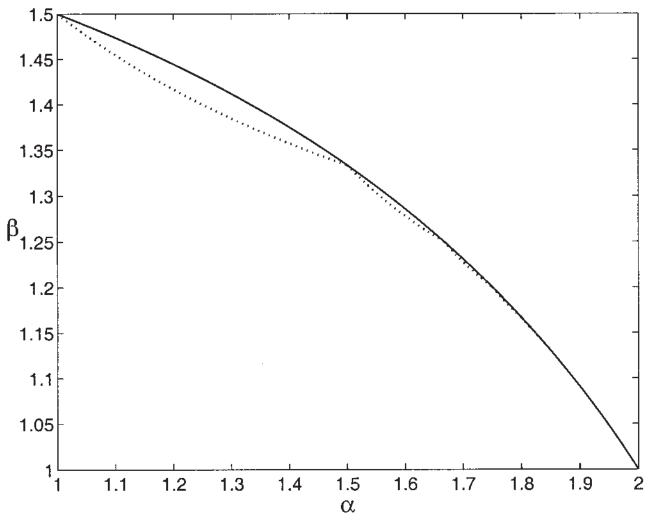

FIG. 1. Trade-off functions.

the optimization of the multicast flow from the core to all the users. Because the core has to send the same data to all the users, minimizing the bandwidth consumption is equivalent to minimizing the number of edges of the shared tree spanning the users and the core. In other words, we must find a *minimum Steiner tree* spanning the users and the core. For this problem, an optimal location for the core is any node on the optimum Steiner tree spanning the users only.

In this article, we consider the problem of locating the core node v under these two criteria, one being the sum of the distances between v and the users or *terminals* (the "median" criterion), the other being the cost of the Steiner tree linking v to the terminals (the "Steiner tree" criterion). Because the 1-median might not be located on an optimum Steiner tree connecting the terminals, there is some trade-off between the two criteria. Our main contribution is to give bounds on this trade-off. To state our result informally, we show that either the optimum 1-median can be connected to one of the terminals without increasing much the cost of the Steiner tree connecting the terminals, or there exists a terminal for which the median criterion is not much more than for the optimum 1-median. More precisely, let α be the ratio for the median criterion between the best terminal and the optimum 1-median, and let β be the ratio of the Steiner tree cost for connecting the terminals and the optimum 1-median to the Steiner tree cost for connecting just the terminals. We can easily show that $\alpha \le 2$, $\beta \le \frac{3}{2}$, and that these values can be attained. More interestingly, we show that for any $(\bar{\alpha}, \bar{\beta})$ $(1 \le \bar{\alpha} \le 2$ and $1 \le \bar{\beta} \le \frac{3}{2})$ such that $\bar{\alpha} = 3 - \frac{1}{2-\bar{\beta}}$ (the *trade-off function,* see Fig. 1), either $\alpha \le \bar{\alpha}$ or $\beta \le \bar{\beta}$. Restating this, if connecting the optimum 1-median to a Steiner tree on the terminals increases its cost by a factor greater than $\bar{\beta}$ then we can select one of the terminals as the core, and the median criterion is no more than $\bar{\alpha}$ times the optimum 1-median value. We actually give a slightly better trade-off function between α and β in Theorem 1; this is indicated by a dotted line on Figure 1. Theorem 1, for example, implies that either α or β is at most $(1 + \sqrt{3})/2 < 1.37$. We would like to mention that even simply showing that β is close to 1 whenever α is close to 2 does not appear to be trivial.

In terms of bicriteria guarantees, our result can be interpreted as saying that, for any $(\bar{\alpha}, \bar{\beta})$ on the trade-off function, there exists a core node such that its median criterion is within $\bar{\alpha}$ of the optimum 1-median and the cost of connecting it to the terminals is at most $\bar{\beta}$ times the cost of connecting only the terminals. In fact, the above discussion says that there exists a core node with either a $(1, \bar{\beta})$ guarantee (if we end up selecting the 1-median), or with a $(\bar{\alpha}, 1)$ guarantee (if we end up selecting one of the terminals). Bicriteria results have been given for many combinatorial optimization problems (see, e.g., [6], [10], [12]). In particular, some problems dealing with trade-offs between shortest paths and spanning tree costs have been studied in [1] and [7].

The study of this trade-off function was motivated by the problem of choosing a core location to minimize the total bandwidth consumption in a multicast communication on a shared tree, under the assumption that all receivers are also sources [3]. Consider, for example, a video-conferencing session in which every terminal sends a data flow to the core consuming say one unit of bandwidth, and the core then multicasts a common virtual scene of bandwidth λ back to the terminals. Typically, λ is an increasing function of the number of terminals and might not even be known in advance. In the *Minimum Bandwidth Cost Core Location Problem,* the goal is to choose a core v so as to minimize the sum of the distances between v and the terminals plus λ times the Steiner tree cost connecting the core v and the terminals. This problem was first studied in [8]. The authors propose an enumeration algorithm that essentially evaluates the objective function at every potential core location. Observe that evaluating this objective function is an NP-hard problem. Computational results are given in [9]. By using a dual-ascent procedure ([15]) to approximate the Steiner trees, they show that it is usually sufficient to explore a small fraction of the graph vertices to obtain good solutions.

If we use a ρ-approximation algorithm to compute a Steiner tree in a core enumeration algorithm, we clearly obtain a ρ-approximation algorithm for the core location problem as well. The currently best known approximation ratio for the Steiner tree problem approaches $1 + \frac{\ln 3}{2} \approx 1.55$ ([11]). However, we use our trade-off results to show that we can approximate the core location problem with a ratio of $1 + \frac{\rho^2}{4}$ (slightly worse than ρ) by approximating only one instance of the minimum cost Steiner tree problem. In particular, for $\rho \approx 1.55$, this gives 1.6, while for $\rho = 2$ corresponding to a minimum spanning tree heuristic for which there exist distributed algorithms, we also obtain a bound of 2.

The remainder of the article is organized as follows. Section 2 gives the basic definitions. Our trade-off function is derived in Section 3. We first prove in Section 3.1 the existence of a decomposition of any tree into a small number of subtrees, each of cost bounded by a fraction of the original cost of the tree; this is stated precisely in Proposition 1. In Section 3.2, the next step is to bound the sum of the distances from any terminal t to the other terminals by going through the optimum Steiner tree for those terminals within the same member of the decomposition as t is and going through the optimum 1-median for the other terminals; this is the essence of Proposition 2. We then apply our trade-off results to the Minimum Bandwidth Cost Core Location Problem and prove that approximating only one Steiner tree is sufficient to obtain a good approximation ratio. These results are presented in Section 4.

2. BASIC DEFINITIONS AND NOTATIONS

Let $G = (V, E, c)$ be a connected, undirected graph with edge costs $c: E \to \mathbb{R}^+$ and $T \subseteq V$ be a particular nonempty subset of vertices called *terminals.* Although we have considered unit costs in the introduction, we carry out our analysis with arbitrary positive edge costs.

The *cost* of a subgraph H of G is the sum of the costs of all edges in H:

$$c(H) = \sum_{e \in E(H)} c_e.$$

The shortest path distance according to the edge costs c_e defines a metric on the vertex set V. The distance between vertices i and j is denoted by $d(i, j)$. For a subgraph H of G, $d_H(i, j)$ denotes the distance between i and j in H.

For any $v \in V$ and $S \subseteq V$, we denote the sum of the distances between v and the vertices in S by

$$f_S(v) = \sum_{t \in S} d(v, t).$$

A *1-median* for S is a vertex minimizing f_S. Whenever $S = T$, we simply write $f(v)$ for $f_T(v)$. Let v^* denote a 1-median for the terminal set T.

For any $S \subseteq V$, a *Steiner tree* for S is a tree spanning S. The cost of a minimum cost Steiner tree for S is denoted by $ST(S)$.

We define $\alpha = \frac{f(v_T)}{f(v^*)}$, where $v_T = \arg\min_{v \in T} f(v)$. In the introduction, β was defined to be $\frac{ST(T \cup \{v^*\})}{ST(T)}$. To make this quantity more tractable, we observe that $ST(T \cup \{v^*\}) \le ST(T) + d(v^*, T)$ where $d(v^*, T) = \min_{v \in T} d(v^*, v)$, and we actually define β to be $1 + \frac{d(v^*, T)}{ST(T)}$. Thus, β is an upper bound on $\frac{ST(T \cup \{v^*\})}{ST(T)}$.

3. RELATIONSHIP BETWEEN α AND β

Here we derive our main inequalities between the parameters α and β.

3.1. Basic Inequalities

We first start by deriving a series of very simple inequalities that in particular show that $\alpha \le 2$ and $\beta \le \frac{3}{2}$.

Lemma 1. *If $\mathcal{T}$ is a tree spanning $S \subseteq V$ then $\sum_{t \in S} f_S(t) \le \frac{|S|^2}{2} c(\mathcal{T})$.*

Proof. We have

$$\sum_{t \in S} f_S(t) = \sum_{t,t' \in S} d(t, t') \le \sum_{t,t' \in S} d_{\mathcal{T}}(t, t')$$
$$= 2 \sum_{e \in E(\mathcal{T})} n_e(|S| - n_e)c_e \le \frac{|S|^2}{2} c(\mathcal{T})$$

where n_e is the number of vertices of S in a connected component of $\mathcal{T} \backslash e$. ■

Applying Lemma 1 to the optimum Steiner tree for T, we derive the following corollary.

Corollary 1. $f(v_T) \le \frac{|T|}{2} ST(T)$.

Lemma 2. $\beta - 1 \le \frac{f(v^*)}{|T| ST(T)}$.

Proof. This follows simply from the definition of $\beta = 1 + d(v^*, T)/ST(T)$ and the fact that $d(v^*, T) \le \frac{1}{|T|} f(v^*)$. ■

Corollary 2. $\beta \le 1 + \frac{1}{2\alpha}$. *In particular,* $\beta \le \frac{3}{2}$.

Proof. From Lemma 2, we have

$$\beta \le 1 + \frac{f(v^*)}{|T| ST(T)} = 1 + \frac{f(v_T)}{\alpha |T| ST(T)} \le 1 + \frac{1}{2\alpha},$$

the equality following from the definition of α and the last inequality from Corollary 1. ■

Lemma 3. $\alpha \le 2 - \frac{2}{|T|}$.

Proof. Let $t^* = \arg\min_{t \in T} d(v^*, t)$. Hence, $f(v^*) \ge |T| d(v^*, t^*)$ and

$$f(v_T) \le f(t^*) = \sum_{t \in T \backslash \{t^*\}} d(t^*, t)$$
$$\le \sum_{t \in T \backslash \{t^*\}} (d(t^*, v^*) + d(v^*, t))$$
$$= (|T| - 1) d(t^*, v^*) + \sum_{t \in T \backslash \{t^*\}} d(v^*, t)$$
$$= (|T| - 2) d(t^*, v^*) + f(v^*) \le f(v^*)\left(2 - \frac{2}{|T|}\right).$$

■

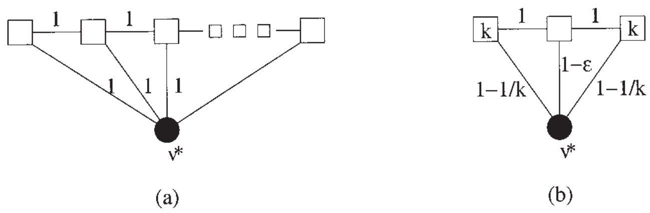

FIG. 2. Tight examples for $\alpha \leq 2$ and $\beta \leq \frac{3}{2}$.

The following examples show that inequalities $\alpha \leq 2$ and $\beta \leq \frac{3}{2}$ are asymptotically tight. We adopt the following convention in Figures 2 and 4: terminals are denoted by squares and the value k inside some indicates that k terminals lie at the same position. The black node represents the 1-median v^* of the terminals. The value close to an edge represents its cost.

Example 1. *Consider the instance depicted in Figure* 2(*a*) *where the terminals lie uniformly on a path of length* $|T| - 1$ *and are at distance* 1 *from* v^*. *Hence*, $f(v^*) = |T|$, $f(v_T) = 2|T| - 4$ (*where* v_T *can be any terminal except the leftmost and rightmost ones*), $ST(T) = |T| - 1$ *and* $ST(T \cup \{v^*\}) = |T|$. *Therefore* $(\alpha, \beta) = (2 - \frac{4}{|T|}, 1 + \frac{1}{|T|-1})$, *which tends to* (2, 1) *when* $|T|$ *tends to* ∞.

Example 2. *In Figure* 2(*b*), $|T| = 2k + 1$ *and the k leftmost and k rightmost terminals lie on a same node. One can verify that* $f(v^*) = 2k - 1 - \epsilon$, $f(v_T) = 2k - 1$ (*where* v_T *can be any leftmost or rightmost terminal*), $ST(T) = 2$ *and* $ST(T \cup \{v^*\}) = 3 - \frac{2}{k} - \epsilon$. *When k tends to* ∞ *and* ϵ *tends to* 0, (α, β) *tends to* $(1, \frac{3}{2})$.

Corollary 2 shows that α must tend to 1 when β tends to $\frac{3}{2}$. Conversely, it is shown in next section that β must tend to 1 when α tends to 2.

3.2. Better Inequalities Using Tree Partitions

For our derivation of the trade-off function, we first need to show the existence of partitions of any tree with a small number of trees and each of a given bounded size. This is stated in Proposition 1.

Definition 1. *A* (δ, M)*-partition of a weighted tree* $\mathcal{T}$ *is a collection of M subtrees* $\mathcal{T}_i$ *of* $\mathcal{T}$ *whose edge sets* $E(\mathcal{T}_i)$ *form a partition of* $E(\mathcal{T})$ *and which satisfy* $c(\mathcal{T}_i) \leq \delta c(\mathcal{T})$.

Definition 2. *A graph H is a* subdivision *of a graph G if H is obtained after iterating the following operation on G*: *Introduce a new vertex w, pick an edge* $e = \{u, v\}$ *and replace it by two edges* $f = \{u, w\}$ *and* $g = \{w, v\}$. *For a weighted graph G, we impose that the previous operation be cost preserving, that is, the cost of the new edges must satisfy* $c_f + c_g = c_e$.

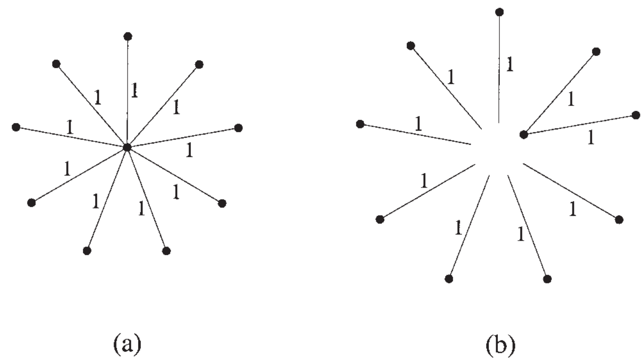

FIG. 3. Tight example for Proposition 1.

Proposition 1. *Let* $\mathcal{T}$ *be a tree. For all* $0 < \delta \leq 1$, *there exists a* $(\delta, \left\lceil \frac{2}{\delta} \right\rceil - 1)$*-partition of a subdivision of* $\mathcal{T}$.

This result is tight in the sense that, for any positive integer M, the smallest value one can take for δ is indeed $\frac{2}{M+1}$ as comes out of Proposition 1. This is shown in the following example.

Example 3. *Consider a star with* $M + 1$ *unit length spokes, as depicted in Figure* 3(*a*). *By the pigeonhole principle, in any* (δ, M)*-partition of a subdivision of the star, there exist two leaves lying in the same subtree. By connectivity, that subtree contains the path of length* 2 *between the two leaves. Hence,* $\delta \geq \frac{2}{M+1}$. *Figure* 3(*b*) *illustrates a* $(\frac{2}{M+1}, M)$ *partition.*

Proof. We first prove the following claim by induction on n, the number of leaves of $\mathcal{T}$.

Claim 1. *For any value* $C > 0$ *and any leaf vertex v of* $\mathcal{T}$, *there exists a partition of a subdivision of* $\mathcal{T}$ *such that the subtree containing v has a cost at most* $\frac{C}{2}$ *and all other subtrees have a cost in* $(\frac{C}{2}, c]$.

Proof of the claim. The proof is by induction. The base case is $n = 2$. Then $\mathcal{T}$ is a path that we partition as follows: starting from the leaf opposite to v, cut pieces of cost C and remove them from the path until the remaining part has a cost less than C. If this cost is at most $\frac{C}{2}$, then the partition is done. Otherwise, do the last cut by considering $\{v\}$ as a subtree (of cost 0). This partition fulfills the claim.

Suppose now that the claim is true for all trees with less than n leaves, where $n \geq 3$. Let w be the closest vertex to v of degree at least 3. If we cut $\mathcal{T}$ in w, we obtain the following connected components: the (v, w)-path and some subtrees $\mathcal{T}_i$ with less than n leaves. For each i, we construct by induction a partition of $\mathcal{T}_i$ that fulfills the condition of the claim with w chosen as the particular leaf. Because the subtrees of the $\mathcal{T}_i$ containing w have cost at most $\frac{C}{2}$, we can merge them into bigger trees such that they all have a cost in $(\frac{C}{2}, c]$ except for one of them, which has a cost at most $\frac{C}{2}$. If needed, consider that $\{w\}$ is that tree. Finally, we extend the low cost tree on the (w, v)-path and finish the construction on this path as in the base case. ∎

Now consider the partition given by the previous claim with $C = \delta c(\mathcal{T})$. Consider the subtree of cost at most $\frac{\delta}{2} c(\mathcal{T})$ and another subtree adjacent to it. If the sum of their cost is less than $C = \delta c(\mathcal{T})$, merge them. In that case all the subtrees have cost in $(\frac{C}{2}, c]$. Otherwise, remark that the average of their cost is in $(\frac{C}{2}, c]$. Let M be the number of subtrees obtained by the previous construction. We must have $M \frac{\delta}{2} c(\mathcal{T}) < c(\mathcal{T})$. Because M is integer, $M \leq \left\lceil \frac{2}{\delta} \right\rceil - 1$, as desired. ∎

From any (δ, M)-partition of the optimum Steiner tree, we can derive an inequality linking α and β as shown in the Proposition below. The idea of the proof is to bound the sum of the distances from any terminal t to the other terminals by going through the optimum Steiner tree for those terminals within the same member of the (δ, M)-partition as t is and going through v^* for the other terminals.

Proposition 2. *If $\beta > 1$ and if there exists a (δ, M)-partition of a subdivision of the optimum Steiner tree $\mathcal{ST}$ on T for some δ and M, then*

$$\alpha \leq 2 - \frac{2}{M} + \frac{\delta}{2(\beta - 1)M}.$$

Proof. If $M = 1$, then $\delta \geq 1$ and it suffices to prove $\alpha \leq \frac{1}{2(\beta-1)}$, which is equivalent to Corollary 2. If $\delta \geq 4(\beta - 1)$, it suffices to prove $\alpha \leq 2$, which is true because of Lemma 3.

So we can suppose w.l.o.g. that $M \geq 2$ and $\delta < 4(\beta - 1)$. Let us define $\delta = \mu(\beta - 1)$ with $0 \leq \mu < 4$.

Let $(\mathcal{T}_i)_{i=1 \ldots M}$ be the subtrees of $\mathcal{T}$ in the (δ, M)-partition. These subtrees naturally induces a partition of T in M blocks $(B_i)_{i=1 \ldots M}$ such that $B_i \subseteq V(\mathcal{T}_i)$. If a terminal $t \in T$ belongs to several subtrees, we arbitrarily choose its block.

For all i, we have

$$|B_i| f(v_T) \leq \sum_{t \in B_i} f(t) = \sum_{t \in B_i} f_{T \setminus B_i}(t) + \sum_{t \in B_i} f_{B_i}(t)$$

$$\leq \sum_{t \in B_i} \sum_{t' \in T \setminus B_i} d(t, t') + \frac{|B_i|^2}{2} c(\mathcal{T}_i) \text{ by Lemma 1}$$

$$\leq \sum_{t \in B_i} \sum_{t' \in T \setminus B_i} (d(t, v^*) + d(v^*, t')) + \frac{|B_i|^2}{2} \delta ST(T)$$

$$= |T \setminus B_i| f_{B_i}(v^*) + |B_i| f_{T \setminus B_i}(v^*) + \frac{|B_i|^2}{2} \delta ST(T)$$

$$= (|T| - 2|B_i|) f_{B_i}(v^*) + |B_i| f(v^*) + \frac{|B_i|^2}{2} \delta ST(T).$$

We now consider two cases.

Case 1. If $|B_i| < \frac{|T|}{2}$ for all i, we use Lemma 2 and rewrite the last inequality as:

$$|B_i| f(v_T) \leq (|T| - 2|B_i|) f_{B_i}(v^*) + |B_i| f(v^*) + \frac{|B_i|^2}{2} \frac{\delta}{(\beta - 1)|T|} f(v^*).$$

A weighted sum over all i yields

$$f(v_T) \sum_{i=1}^{M} \frac{|B_i|}{|T| - 2|B_i|} \leq f(v^*) \left(1 + \sum_{i=1}^{M} \frac{|B_i|}{|T| - 2|B_i|} + \frac{\mu}{2|T|} \sum_{i=1}^{M} \frac{|B_i|^2}{|T| - 2|B_i|} \right),$$

where $\mu = \frac{\delta}{\beta - 1}$. Note that $0 < \mu < 4$. Let us do the following variable substitution: $z_i = \frac{|T| - 2|B_i|}{|T|}$ or, equivalently, $|B_i| = |T| \frac{1 - z_i}{2}$. Thus, the z_i's satisfy

$$0 < z_i \leq 1 \text{ and } \sum_{i=1}^{M} z_i = M - 2.$$

Therefore,

$$\alpha = \frac{f(v_T)}{f(v^*)} \leq \frac{1 + \sum_{i=1}^{M} \frac{1 - z_i}{2z_i} + \frac{\mu}{8} \sum_{i=1}^{M} \frac{(1 - z_i)^2}{z_i}}{\sum_{i=1}^{M} \frac{1 - z_i}{2z_i}}$$

$$= 1 + \frac{1 + \frac{\mu}{8} \left(\sum_{i=1}^{M} \frac{1}{z_i} - 2M + M - 2 \right)}{\frac{1}{2} \sum_{i=1}^{M} \frac{1}{z_i} - \frac{1}{2} M}$$

$$= 1 + \frac{8 + \mu(Z - M - 2)}{4(Z - M)} \text{ where } Z = \sum_{i=1}^{M} \frac{1}{z_i}$$

$$= 1 + \frac{\mu}{4} + \frac{4 - \mu}{2(Z - M)}.$$

Because $\mu < 4$, the last term is maximized when Z is minimum. By Cauchy-Schwarz inequality,

$$M^2 \leq \left(\sum_{i=1}^{M} z_i \right) \left(\sum_{i=1}^{M} \frac{1}{z_i} \right) = (M-2)Z.$$

Hence,

$$\alpha \leq 1 + \frac{\mu}{4} + \frac{(4-\mu)(M-2)}{4M} = 2 - \frac{2-\mu/2}{M},$$

proving the desired bound.

CASE 2. If $|B_i| \geq \frac{|T|}{2}$ for some i, then for this i, we have

$$|B_i| f(v_T) \leq (|T| - 2|B_i|) f_{B_i}(v^*) + |B_i| f(v^*) + \frac{|B_i|^2}{2} \delta ST(T)$$
$$\leq (|T| - 2|B_i|)|B_i|(\beta - 1)ST(T) + |B_i| f(v^*) + \frac{|B_i|^2}{2} \mu(\beta - 1)ST(T)$$

using, as in the proof of Lemma 2, $\frac{f_{B_i}}{|B_i|} \geq d(v^*, T) = (\beta - 1)ST(T)$. Hence,

$$f(v_T) \leq f(v^*) + \left(1 - 2\frac{|B_i|}{|T|} + \frac{\mu}{2}\frac{|B_i|}{|T|}\right)|T|(\beta - 1)ST(T).$$

Because $f(v_T) \geq f(v^*)$, $1 - (2 - \frac{\mu}{2})\frac{|B_i|}{|T|} \geq 0$. Using Lemma 2, $\mu < 4$ and $M \geq 2$, we have

$$\alpha \leq 2 - \left(2 - \frac{\mu}{2}\right)\frac{|B_i|}{|T|} \leq 2 - \left(2 - \frac{\mu}{2}\right)\frac{1}{2}$$
$$\leq 2 - \frac{2 - \mu/2}{M},$$

completing the proof. ∎

We can now prove the following inequalities:

Theorem 1. *If $\beta > 1$ then, for all positive integer M,*

$$\alpha \leq 2 - \frac{2}{M} + \frac{1}{M(M+1)(\beta - 1)}. \tag{1}$$

In particular,

$$\alpha \leq 3 - \frac{1}{2 - \beta}. \tag{2}$$

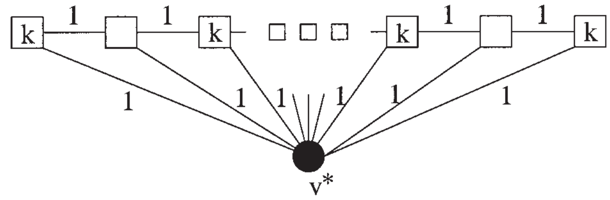

FIG. 4. Worst known examples.

The lower envelope of all inequalities (1) corresponds to the dotted line in Figure 1, while (2) corresponds to the continuous line in the figure.

Proof. For each positive integer M, there exists a (δ, M)-partition of the optimum Steiner tree $\mathcal{ST}$ with $\delta = \frac{2}{M+1}$ by Proposition 1. Proposition 2 immediately yields inequality (1).

Consider two inequalities (1) with $M = n$ and $M = n + 1$, where n is a positive integer. The first inequality is stronger than the second one if and only if:

$$2 - \frac{2}{n} + \frac{1}{n(n+1)(\beta - 1)} \leq 2 - \frac{2}{n+1} + \frac{1}{(n+1)(n+2)(\beta - 1)},$$

or equivalently $\beta \geq 1 + \frac{1}{n+2}$. Therefore, each inequality (1) is the strongest one when $\beta \in [1 + \frac{1}{M+2}, 1 + \frac{1}{M+1}]$, which corresponds to $\alpha \in [2 - \frac{1}{M}, 2 - \frac{1}{M+1}]$. By eliminating the parameter M, one can easily see that all points $(\alpha, \beta) = (2 - \frac{1}{M}, 1 + \frac{1}{M+1})$ lie on the curve $\alpha = 3 - \frac{1}{2-\beta}$. Observe, finally, that the right-hand side of inequality (1) is convex in β and the right-hand side of inequality (2) is concave in β. This proves the validity of inequality (2). ∎

The inequalities of Theorem 1 are not likely to be tight. The worst examples we know are the following.

Example 4. *In the instance depicted in Figure 4, there are M groups of k terminals plus another $M - 1$ terminals between the groups, for a total of $kM + M - 1$ terminals. When k tends to ∞, this can be seen to correspond to $(\alpha, \beta) = (2 - \frac{2}{M}, 1 + \frac{1}{2M-2})$. For example, for $M = 3$, this gives $(\alpha, \beta) = (\frac{4}{3}, \frac{5}{4})$.*

4. MINIMIZING THE TOTAL BANDWIDTH CONSUMPTION

As discussed in the introduction, the *Minimum Bandwidth Cost Core Location Problem* is the problem of finding the best vertex v, which minimizes $cost(v) = f(v) + \lambda ST(T \cup \{v\})$. This problem, first defined in [8], models the situation in which the terminals in T independently send one unit of bandwidth to the core, which sends back λ units of bandwidth along a Steiner tree connecting v and the terminals. The analysis that we provide below does not make any assumption on λ.

Leung and Yum [8] propose an enumeration algorithm that essentially evaluates the objective function at every potential core location. If we use a particular ρ-approximation algorithm to compute the Steiner tree costs in a core enumeration algorithm, we clearly obtain a ρ-approximation algorithm for the core location problem as well, independently of the value of λ. The currently best-known approximation ratio for the Steiner Tree problem approaches $1 + \frac{\ln 3}{2} \approx 1.55$ ([11]). However, we prove in this section that we can approximate the Minimum Bandwidth Cost Core Location Problem with a ratio close to ρ by computing only one ρ-approximation of a minimum cost Steiner tree. Namely, we consider the following core location algorithm:

CoreLocation

input: A graph $G = (V, E, c)$ and a set of terminals $T \subseteq V$.
output: A core vertex.

1. Compute a Steiner tree $\mathcal{ST}_\rho$ spanning T with a ρ-approximation algorithm.
2. Compute $v^* = \arg\min_{v \in V} f(v)$ and $\overline{cost}(v^*) = f(v^*) + \lambda(c(\mathcal{ST}_\rho) + d(v^*, T))$.
3. Compute $v_T = \arg\min_{t \in T} f(t)$ and $\overline{cost}(v_T) = f(v_T) + \lambda c(\mathcal{ST}_\rho)$.
4. **If** $\overline{cost}(v^*) \leq \overline{cost}(v_T)$, **then return** v^*, **else return** v_T.

The main advantage of the algorithm is that we need to solve only one Steiner tree instance. Indeed, in a real setting, finding a good Steiner tree is typically the bottleneck in terms of computation and communication time, whereas the evaluation of f is fast because the shortest path distances are usually precomputed and stored in routing tables.

Proposition 3. *CoreLocation is a* $(1 + \frac{\rho^2}{4})$*-approximation algorithm of the Minimum Bandwidth Cost Core Location Problem.*

Proof. *CoreLocation* returns a node v^{CL} with

$$\begin{aligned}\overline{cost}(v^{CL}) &= \min\{f(v^*) + \lambda(c(\mathcal{ST}_\rho) + d(v^*, T)), f(v_T) \\ &\quad + \lambda c(\mathcal{ST}_\rho)\} \\ &\leq \min\{f(v^*) + \lambda(\rho ST(T) + d(v^*, T)), f(v_T) \\ &\quad + \lambda \rho ST(T))\} \\ &= \min\{f(v^*) + \lambda(\rho + \beta - 1)ST(T), \alpha f(v^*) \\ &\quad + \lambda \rho ST(T)\}.\end{aligned}$$

Clearly, a lower bound on the minimum cost is given by:

$$\min_{v \in V} cost(v) \geq f(v^*) + \lambda ST(T).$$

Therefore, if we set

$$\mu = \frac{f(v^*)}{f(v^*) + \lambda ST(T)},$$

for fixed value of α and β, the approximation ratio of *CoreLocation* is upper bounded as follows:

$$\frac{\overline{cost}(v^{CL})}{\min_{v \in V} cost(v)} \leq \max_{0 \leq \mu \leq 1} \min\{\mu + (1-\mu)(\rho + \beta - 1), \mu\alpha + (1-\mu)\rho\}. \tag{3}$$

Remark that, if $\rho \geq \alpha$, then the second argument of the minimum in (3) is at most $\rho \leq 1 + \frac{\rho^2}{4}$. Thus, we can suppose w.l.o.g. that $\rho < \alpha$. Hence, the first (resp., second) argument is a nonincreasing (resp. nondecreasing) linear function in μ. Also, the first argument is at least (resp. at most) the second one in $\mu = 0$ (resp. $\mu = 1$). This implies that the maximum is attained when μ satisfies

$$\mu + (1-\mu)(\rho + \beta - 1) = \mu\alpha + (1-\mu)\rho \Leftrightarrow \mu = \frac{\beta - 1}{\alpha + \beta - 2}.$$

The upper bound (3) becomes

$$\frac{\overline{cost}(v^{CL})}{\min_{v \in V} cost(v)} \leq \frac{\alpha\beta + \alpha(\rho - 1) - \rho}{\alpha + \beta - 2} = \alpha + \frac{(\alpha - 1)(\rho - \alpha)}{\alpha + \beta - 2}. \tag{4}$$

Observe that the right-hand side of inequality (4) is a nondecreasing function in β (for $\beta \geq 1$) because $\alpha \geq 1$ and $\rho < \alpha$. Therefore, by Theorem 1, we can upper bound the approximation ratio by maximizing the right-hand side of inequality (4) along the curve $\alpha = 3 - \frac{1}{2-\beta}$. However, to make this maximization easier, we will first derive a new inequality:

$$\alpha \leq 3 - \frac{1}{2 - \beta} = 3 - \frac{1}{1 - (\beta - 1)} \leq 3 - (1 + (\beta - 1)) = 3 - \beta.$$

If we maximize the right-hand side of inequality (4) along the line $\alpha + \beta = 3$, the upper bound becomes:

$$\frac{\overline{cost}(v^{CL})}{\min_{v\in V} cost(v)} \leq \max_{1\leq\alpha\leq 2} \alpha + (\alpha - 1)(\rho - \alpha)$$

$$= \max_{1\leq\alpha\leq 2} \alpha(2 + \rho - \alpha) - \rho = \left(1 + \frac{\rho}{2}\right)^2 - \rho = 1 + \frac{\rho^2}{4},$$

where the maximum is attained for $\alpha = 1 + \frac{\rho}{2}$. ■

Let us look at some particular approximation algorithms for the Minimum Cost Steiner Tree problem. If we use a minimum spanning tree algorithm, $\rho = 2$ (see [14]). Hence, by Proposition 3, we obtain a 2-approximation algorithm as well for the core location problem. So in this case it is impossible to improve the approximation ratio of *CoreLocation* by computing other Steiner trees. The minimum spanning tree problem has also the main advantage to have distributed algorithms (see [5] for instance), which are therefore easier to implement in a real network. If we use the best-known approximation algorithm for the Steiner tree problem ($\rho \approx 1.55$), then we obtain a ratio of 1.60. Finally, if we could always compute the optimal Steiner tree spanning the terminals, *CoreLocation* would have an approximation ratio of 1.25. Note that, by using the inequalities of Theorem 1 instead of $\alpha + \beta \leq 3$, one can actually prove that *CoreLocation* is a 1.2-approximation algorithm when $\rho = 1$ (this is attained for $\alpha = \frac{3}{2}$ and $\beta = \frac{4}{3}$ in inequality (4)).

Acknowledgment

J.-F. Macq is a Research Fellow of the Belgian National Funds for Scientific Research (F.N.R.S.). This work was done while visiting M.I.T.

REFERENCES

[1] B. Awerbuch, A. Baratz, and D. Peleg, "Cost-sensitive analysis of communication protocols," Proc 9th ACM Symposium on Principles of Distributed Computing (PODC), Quebec City, Quebec, Canada, 1990, pp. 177–187.

[2] T. Ballardie, P. Francis, and J. Crowcroft, "Core-based trees (CBT): An architecture for scalable inter-domain multicast routing," Proc ACM SIGCOMM, 1993, pp. 85–95.

[3] K.L. Calvert, E.W. Zegura, and M.J. Donahoo, "Core selection methods for multicast routing," Proc 4th Int Conference on Computer Communications and Networks (ICCN'95), Las Vegas, Nevada, 1995, pp. 638–642.

[4] E. Fleury, Y. Huang, and P.K. McKinley, On the performance and feasibility of multicast core selection heuristics, Networks 35 (2000), 145–156.

[5] R.G. Gallager, P.M. Humblet, and P.M. Spira, A distributed algorithm for minimum-weight spanning trees, ACM Transact Program Languages Syst 5 (1983), 66–77.

[6] S.O. Krumke, H. Noltemeier, S.S. Ravi, and M.V. Marathe, Bicriteria compact location problems, Stud Locat Anal 10 (1996), 37–51.

[7] S. Khuller, B. Raghavachari, and N.E. Young, Balancing minimum spanning trees and shortest-path trees, Algorithmica, vol. 14 (1995), 305–321.

[8] Y.W. Leung and T.S. Yum, Connection optimization for two types of videoconferences, IEE Proc Commun 143 (1996), 133–140.

[9] J.-F. Macq, L.A. Wolsey, and B. Macq, "A rendez-vous point selection algorithm for videoconferencing applications," IEEE Int Conf Multimedia Expo (ICME2002), Lausanne, Switzerland, 2002, vol 2, pp. 689–692.

[10] M.V. Marathe, R. Ravi, R. Sundaram, S.S. Ravi, D.J. Rosenkrantz, and H.B. Hunt, Bicriteria network design problems, J Algorithms 28 (1998), 142–171.

[11] G. Robins and A. Zelikovsky, "Improved Steiner tree approximation in graphs," Proc of ACM-SIAM Symp Discrete Algorithms (SODA2000), 2000, pp. 770–779.

[12] C. Stein and J. Wein, On the existence of schedules that are near-optimal for both makespan and total weighted completion time, Operat Res Lett 21 (1997), 115–122.

[13] D. Thaler and C.V. Ravishankar, Distributed center-location algorithms, IEEE J Selected Areas Commun 15 (1997), 291–303.

[14] P. Winter, Steiner problem in networks: A survey, Networks 17 (1987), 129–167.

[15] R.T. Wong, A dual ascent approach for Steiner tree problems on a directed graph, Math Program 28 (1984), 271–287.