Sharing the Cost of Multicast Transmissions in Wireless Networks*

Paolo Penna and Carmine Ventre

Dipartimento di Informatica ed Applicazioni "R.M. Capocelli",
Università di Salerno, via S. Allende 2, I-84081 Baronissi (SA), Italy
{penna,ventre}@dia.unisa.it

Abstract. We investigate the problem of sharing the cost of a multicast transmission in a wireless network where each node (radio station) of the network corresponds to (a set of) user(s) potentially interested in receiving the transmission. As in the model considered by Feigenbaum *et al* [2001], users may act *selfishly* and report a false "level of interest" in receiving the transmission trying to be charged less by the system. We consider the issue of designing a so called *truthful mechanisms* for the problem of maximizing the *net worth* (i.e., the overall "happiness" of the users minus the cost of the transmission) for the case of *wireless* networks. Intuitively, truthful mechanism guarantee that no user has an incentive in reporting a false valuation of the transmission. Unlike the "wired" network case, here the cost of a set of connections implementing a multicast tree is *not* the sum of the single edge costs, thus introducing a complicating factor in the problem. We provide both positive and negative results on the existence of optimal algorithms for the problem and their use to obtain VCG truthful mechanisms achieving the same performances.

1 Introduction

One of the main benefits of ad-hoc wireless networks relies in the possibility of communicating without any fixed infrastructure. Indeed, each station is a radio transmitter/receiver and communication between two stations that are not within their respective transmission ranges can be achieved by *multi-hop* transmissions: a set of intermediate stations forwards the message till its destination.

In this work, we consider the problem of sharing the cost of a multicast transmission in such wireless networks. A set of radio stations implements a *directed communication graph* which can be used to broadcast a (set of) messages from a given source node s to any subset of *users*. In particular, each user j is sitting close to some station i and she can receive the transmission only if i does. In addition, user j benefits from receiving the transmission some amount specified by a value v_j (say, how much j valuates the transmission).

* Work supported by the European Project IST-2001-33135, Critical Resource Sharing for Cooperation in Complex Systems (CRESCCO).

R. Královič and O. Sýkora (Eds.): SIROCCO 2004, LNCS 3104, pp. 255–266, 2004.

Since transmissions along the edges (i.e., links) of the communication graph require some costs (i.e., the power consumption of the station forwarding the messages), one would like to select a suitable set of nodes, to which the transmission is sent to, so that (i) the users share the cost of the transmission, and (ii) the overall *net worth* is maximized: the net worth is defined as the sum of the v_js of all users receiving the transmission minus the overall cost due to the used links.

Although the costs of the links are a property of the network (thus known to the "protocol"), the valuation v_j is clearly a property of user j. Thus, each v_j is a *private* piece of information (a part of the input) held by user j only. So, a user may act *selfishly* and report a different value b_j trying to receive the transmission at a lower price. We thus need to design so called *truthful mechanism* for our problem, that is, a suitable combination of an algorithm A and payments to the users which guarantee that (i) no user j has an incentive in reporting $b_j \neq v_j$, and (ii) the algorithm A, once provided with the correct v_js, returns an optimal solution.

This problem has been previously considered in the context of "classic" wired networks in [15]. In this work we consider a different model, that is, the wireless network one. The main difference between the two models relies on the different cost functions, which turns out to be a key point for solving the problem above.

Cost of Wireless Connections. Consider a *directed weighted communication graph* $\mathcal{G} = (\mathcal{S}, \mathcal{E}, w)$ defined as follows: $\mathcal{S}$ is the set of stations, and $\mathcal{G}$ contains a *directed* edge $(i, j) \in \mathcal{E}$ if and only if the direct transmission from i to j is feasible; in this case the weight $w(i, j)$ is the minimum power required for station i to directly transmit to station j. For instance, in the empty space $w(i, j) = \mathsf{d}(i, j)^2$, where $\mathsf{d}(i, j)$ is the Euclidean distance between i and j.

Each station is a radio transmitter/receiver and a station i is able to *directly* transmit a message to station j if and only if the power P_i used by station i satisfies $P_i \geq w(i, j)$. Stations use *omnidirectional antennas*, and a message sent by station i to j can be also received by every other station j' for which $w(i, j') \leq w(i, j)$. In order to reduce the power consumption, every station i can adjust its transmission power P_i, thus implementing the set of connections $\{(i, j)|\ w(i, j) \leq P_i\}$. Hence, given a set of connections $C \subseteq \mathcal{E}$, its cost is defined as follows:

$$\mathsf{Cost}(C) = \sum_{i \in \mathcal{S}} \max_{j:(i,j) \in C} w(i, j), \tag{1}$$

that is, the *overall power consumption* required to implement all these connections.

Net Worth of a Multicast Transmission. We are interested in sets $C \subseteq \mathcal{E}$ which guarantee that, given a distinguished *source node* s, the set C connects s to a suitable set $D(C) \subseteq \mathcal{S}$ of *destination nodes.* Consider a set U of users, each of them located close to some of the nodes in $\mathcal{S}$. The source s can send some kind of transmission (say a movie or a sport event) to a user j only if j is close to some of the destination nodes $D(C)$. In addition, every user j has a *valuation* v_j of the transmission representing how much she would benefit from receiving it (i.e., how much she would pay for it). As in the model of [15], we consider the situation in which each user j is sitting close to one station, say i; the latter represents the router of the network at distance one hop from user j. So, user j can receive the transmission only if node i does. Observe that, we can always reduce the case of several users located close to the same node to the case of (at most) one user close to one node (consider each user as a node with no outgoing edges and one ingoing edge of cost 0). Given a solution $T \subseteq \mathcal{E}$, its *net worth* is defined as

$$\mathsf{NW}(T) = \mathsf{Worth}(T) - \mathsf{Cost}(T), \tag{2}$$

where $\mathsf{Worth}(T) = \sum_{i \in D(T)} v_i$.

The *cost sharing problem* asks for a $T \subseteq \mathcal{E}$ that, for a given source s, maximizes the net worth function above.

Selfish Users and Economical Constraints. Associated to each node there is a *selfish agent* reporting some (not necessarily true) valuation b_i; the true value v_i is *privately known* to agent i. Based on the reported values $b = (b_1, b_2, \ldots, b_n)$ a *mechanism* $M = (A, P)$ constructs a multicast tree T using algorithm A (i.e. $T = A(b)$) and charges, to each agent i, an amount of money to pay for receiving the transmission equal to $P^i(b)$, with $P = (P^1, P^2, \ldots, P^n)$.

There is a number of natural constraints/goals that we would like a mechanism $M = (A, P)$ to satisfy/meet:

1. Truthfulness (or Strategyproofness)[1]. For every i, let $b_{-i} := (b_1, b_2, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$ and $(b_i, b_{-i}) := b$. The *utility* of agent i when she reports b_i, and the other agents report b_{-i}, is equal to

$$u_i(b_i, b_{-i}) := \begin{cases} v_i - P^i(b_i, b_{-i}) & \text{if } T = A(b_i, b_{-i}) \text{ and } i \in D(T), \\ 0 & \text{otherwise.} \end{cases}$$

We require that, for every i, for every b_{-i}, and for every $b_i \neq v_i$, it holds that $u_i(v_i, b_{-i}) \geq u_i(b_i, b_{-i})$. In other words, whatever strategy the other agents follow, agent i has no incentive to lie about her true valuation v_i. A mechanism satisfying this property is called *truthful.*

2. Efficiency. The *net worth* $\mathsf{NW}(T)$ yielded by the computed solution T is maximum, that is, $\mathsf{NW}(T) = \max_{C \subseteq \mathcal{E}} \{\mathsf{NW}(C)\}$.

3. No Positive Transfer (NPT). No user receives money from the mechanism, i.e., $P^i(\cdot) \geq 0$.

4. Voluntary Participation (VP). We never charge a user an amount of money grater than her *reported* valuation, that is, $\forall b_i, \forall b_{-i} \quad b_i \geq P^i(b_i, b_{-i})$. In particular, a user has always the option to not paying for a transmission for which she is not interested.

5. Consumer Sovereignty (CS). Every user is guaranteed to receive the transmission if she reports a high enough valuation.

6. Budget Balance (BB). $\sum_i P^i(b) = \mathsf{Cost}(A(b))$.

7. Cost Optimality (CO). The set of connections T is optimal w.r.t. the set of receivers $D(T)$, that is, $\mathsf{Cost}(T) = \min_{C \subseteq \mathcal{E}, D(T) = D(C)} \{\mathsf{Cost}(C)\}$.

[1] In Sect. 2 we provide a more general definition of truthfulness which applies to a wide class of problems involving selfish agents.

Clearly, the Efficiency requirement implies the Cost Optimality. Unfortunately, in some cases it is impossible to achieve efficiency, so we will relax it to *r-efficiency*, that is, $r \cdot \mathsf{NW}(T) \geq \max_{C \subseteq \mathcal{E}}\{\mathsf{NW}(C)\}$. In these cases, we will take into account CO and r-CO, i.e., $\mathsf{Cost}(T) \leq r \cdot \min_{\substack{C \subseteq \mathcal{E}, \\ D(T)=D(C)}} \{\mathsf{Cost}(C)\}$.

1.1 Previous Work

Power Consumption and Range Assignment Problems. The problem of computing a broadcast tree of minimal cost for wireless networks has been investigated in [19, 13, 9, 30]. In particular, in [19] the authors proved that the problem is NP-hard to approximate within logarithmic factors, while it remains NP-hard even when considering geometric 2-dimensional networks [9]. Several variants of this problem have been considered in [23, 12, 13, 9, 30, 5, 6, 1] (see also [10] for a survey). However, to our knowledge, no algorithmic solution for optimizing the net worth has been given so far.

Recently, the design of truthful mechanisms for the range assignment problem in presence of "selfish transmitters" (i.e., selfish agents that want to minimize the energy their station has to use) has been investigated in [2] for the strongly connectivity problem, and in [3] for point-to-point transmissions, respectively.

Mechanism Design and Cost-Sharing Mechanisms in Wired Networks. The theory of mechanism design dates back to the seminal papers by Vickrey [29], Clarke [8] and Groves [18], and recently found a natural application to (algorithmic) questions related to the Internet [25] (see also [16] and [26]). VCG mechanisms guarantee the truthfulness under the hypothesis that the mechanism is able to compute the optimum and the optimization function is *utilitarian* (see Sect. 2 for a formal definition of utilitarian problem).

This technique is employed in [15] (and in this work) where the authors consider the wired networks case. They indeed provide a *distributed* optimal algorithm for the case in which the communication graph is a directed tree. This yields a *distributed mechanism*[2] which, for this problem version, satisfies all requirements mentioned above (truthfulness, efficiency, etc.) except for budget balance.

Noticeably, a classical result in game theory [17, 28] implies that, for this model, budget balance and efficiency are mutually exclusive. Additionally, in [14] (see also Theorem 5 in [7]) it is shown that no α-efficiency and β-efficiency can be guaranteed simultaneously, for any two $\alpha, \beta > 1$. So, the choice is to either optimize the efficiency (as in [15]) or to meet budget balance (as in [20, 21, 7]). In the latter case, it is also possible to obtain so called *group strategyproofness*, a stronger notion of truthfulness which can also deal with *coalitions*. On the other hand, if we insist on efficiency, NPT, VP, and CS, then there is essentially only one such mechanism: the marginal-cost mechanism [24], which belongs to the VCG family.

All such negative results apply also to our problem (i.e., wireless networks). Indeed, a simple observation is that every instance of the "wired" case can be

[2] The mechanism is able to compute both the solution and the payments in distributed fashion.

reduced to the wireless one using the following trick: replace every edge (i, j), with two edges $(i, x(i, j))$ and $(x(i, j), j)$ with $w(i, x(i, j)) = 0$ and $w(x(i, j), j) = w(i, j)$. So, also for our problem we have to choose between either budget balance of efficiency.

1.2 Our Results

We consider the problem of designing mechanisms that satisfy truthfulness, efficiency, NPT, VP, CS, and CO in the case of wireless networks. We first show that, even though the problem is not utilitarian, it is possible to adapt the VCG technique so as to obtain truthful mechanisms based on exact algorithms (Sect. 2).

Unfortunately, the problem is NP-hard, thus preventing from a straightforward use of the VCG result to obtain *polynomial-time* truthful mechanisms. Motivated by this negative result, we first consider the problem restricted to communication graphs that are trees (this is the analogous of the result for wired networks in [15]). We prove that, in this case, the optimal net worth can be computed via a polynomial time *distributed* algorithm (Sect. 3.1). The importance of this result[3] is threefold:

- It shows that the hardness of the problem is confined in the choice of a "good" multicast tree, and not in its use: if an "oracle" provides us with a tree containing an optimal multicast tree, then we can compute the optimum in polynomial-time.
- It is used to obtain a truthful *distributed* polynomial-time mechanism satisfying NPT, VP, CS, and efficiency when the given communication graph is a tree. In this case, both the solution and the payments can be computed in distributed fashion using $O(1)$ messages per link.
- It can be used to approximate the problem in some situations for which a good "universal" tree exists, i.e., a tree containing a set of connections of cost not much larger than the optimal solution and reaching the same set of nodes. This approach is similar to that of several wireless multicast protocols[4] which construct a multicast tree by pruning a broadcast tree $\mathcal{T}$ (e.g., MIP, MLU and MLiMST in [13]). In all such cases, one can assume that the communication graph $\mathcal{G}$ is the tree $\mathcal{T}$.

Moreover, we show that a shortest-path tree can be used as universal tree so to obtain a polynomial-time mechanism satisfying truthfulness, NPT, VP, CS, and $O(n)$-CO, for the case of any communication graph $\mathcal{G}$. In addition, our mechanism guarantees $|D(T^*)|$-efficiency, for all instances that admit an optimal solution T^* satisfying $|D(T^*)| \leq \gamma \frac{\mathsf{Worth}(T^*)}{\mathsf{Cost}(T^*)}$, for some constant $\gamma < 1$ (Sect. 3.2). We also prove that, in general, for any $R > 0$, no polynomial-time algorithm can guarantee R-efficiency, unless $\mathsf{P} = \mathsf{NP}$. Notice that, this result rules out the possibility of having polynomial-time mechanisms satisfying $O(n)$-efficiency.

[3] Independently from this work, Biló *et al* [4] also provide polynomial-time truthful mechanisms for trees in the case of wireless networks.

[4] In this case the set of destination nodes (termed multicast group) is given in input.

We then extend our positive result to a class of graphs denoted as trees with *metric free edges* (see Sect. 3.3). Our technical contribution here is a non-trivial algorithm extending the technique and the results for trees.

Finally, we turn our attention to the Euclidean versions of the problem, that is, the case in which points are located on a d-dimensional Euclidean space and $w(i, j) = \mathsf{d}(i, j)^\alpha$, for a fixed constant $\alpha \geq 1$. We first show that the problem remains NP-hard even when $d = 2$ and for any $\alpha > 1$ (Sect. 3.4). For the case $d = 1$ we provide a polynomial-time mechanism satisfying truthfulness, efficiency, NPT, VP, CS and CO, with the additional property of ensuring multicast trees of depth at most h, for any $1 \leq h \leq n-1$ given in input. This result exploits the broadcasting algorithm in [11]. For the case $d = 2$, we present a solution based on the construction of so called *Light Approximate Shortest-path Trees* (LASTs) given in [22]. This achieves a better performance w.r.t. MST-based solutions in several cases.

Due to lack of space some of the proofs are omitted. We refer the interested reader to the full version of this work [27].

2 Optimal Algorithms Yield Truthful Mechanisms

For the sake of completeness, we first recall the classical technique to obtain truthful mechanisms for utilitarian problems known as VCG-mechanism [29, 8, 18]. We then show how to adapt this technique to our (non-utilitarian) problem.

Let us first consider a more general situation in which each agent i has a certain type t_i. The *valuation* of agent i of a solution X is represented by a function $\mathsf{VAL}_i(X, t_i)$, where the function $\mathsf{VAL}_i(\cdot, \cdot)$ is known to the mechanism. A *maximization* problem is *utilitarian* if its objective function $g(\cdot)$, which depends on the agents type vector $t = (t_1, t_2, \ldots, t_n)$, satisfies

$$g(X, t) = \sum_{j=1}^{n} \mathsf{VAL}_j(X, t_j), \qquad (3)$$

for any solution X. Each agent i can declare a different type b_i to the mechanism. We have the following result on the VCG mechanism $M = (\mathsf{ALG}, P_{VCG})$:

Theorem 1. *[18] If* ALG *is an optimal algorithm for a utilitarian problem Π, then the mechanism $M = (\mathsf{ALG}, P_{VCG})$ is truthful for Π.*

Let $\sigma_i(T) = 1$ if $i \in D(T)$, and 0 otherwise. Notice that, simply setting $\mathsf{VAL}_i(T, v_i) := \sigma_i(T) \cdot v_i$ does not satisfy the definition of utilitarian problem since $\mathsf{NW}(T) \neq \sum_{i=1}^{n} \mathsf{VAL}_i(T, v_i) = \mathsf{Worth}(T)$. Nevertheless, the next results states that the VCG technique can be adapted to our problem. The main idea is to initially charge each node by the cost of its ingoing edge in the tree (computed as in the wireless network case) so to "reduce" our problem to a utilitarian one (see [27] for the details).

Theorem 2. *Let A be a (polynomial-time) exact algorithm for maximizing the* $\mathsf{NW}(\cdot)$ *function. Then the cost sharing problem on wireless networks admits a (polynomial-time) mechanism $M = (A, P_A)$ satisfying truthfulness, efficiency, NPT, VP, CS and CO.*

3 Special Communication Graphs

Motivated by the result of the previous section, we focus on the existence of *polynomial-time exact* algorithms for the cost sharing problem on wireless networks. Since the problem is NP-hard, also for Euclidean 2-dimensional instances (Sect. 3.4), in the following we will focus on restrictions for which such algorithms exist.

3.1 Trees

We proceed similarly to [15] and assume that the communication graph is a directed tree $\mathcal{T} = (\mathcal{S}, \mathcal{E})$.

Definition 1. *For every $i \in \mathcal{T}$, let $p(i)$ denote its parent in $\mathcal{T}$. Also let $c_i = w(p(i), i)$ and $c_s = 0$. We denote by $\mathcal{T}_i$ the subtree rooted at i, and by* $\mathsf{ch}(i)$ *the set of i's children. Finally,* $\mathsf{NW}_{\mathsf{opt}}(i)$ *denotes the* optimal net worth *of $\mathcal{T}_i$, that is, the optimum for the instance in which i is the source node and the universal tree is $\mathcal{T}_i$.*

Algorithm Wireless_Trees at node i

1. After receiving a message μ^j from each child $j \in \mathsf{ch}(i)$ do
   (a) $Add(j) := -c_j + \sum_{k \in \mathsf{ch}(i), c_k \leq c_j} \mu^j$;
   (b) $T(i) := \emptyset$;
   (c) if $\max_{j \in \mathsf{ch}(i)} Add(j) < 0$ then $\mu^i := v_i$
   (d) else do
      i. $Add := \max_{j \in \mathsf{ch}(i)} Add(j)$;
      ii. $\mu^i := v_i + Add$;
      iii. $j^* := \arg\max\{j \in \mathsf{ch}(i) | Add(j) = Add\}$;
      iv. $T(i) := \{(i,j) | w(i,j) \leq w(i,j^*)\}$;
2. send μ^i to parent $p(i)$;

Fig. 1. The distributed algorithm for trees computing an optimal solution in bottom-up fashion.

Lemma 1. *For every node i, it holds that*

$$\mathsf{NW}_{\mathsf{opt}}(i) = v_i + \max\{0, \max_{j \in \mathsf{ch}(i)} \{-c_j + \sum_{k \in \mathsf{ch}(i), c_k \leq c_j} \mathsf{NW}_{\mathsf{opt}}(k)\}\}. \qquad (4)$$

Proof. The proof is by induction on the hight h of $\mathcal{T}_i$. Obviously, for $h = 1$, since i is a leaf node, then $\mathsf{NW}_{\mathsf{opt}}(i) = v_i$. Let us now assume that the lemma holds for any $h' \leq h - 1$, and let us prove it for h. Let T_i^* denote an optimal solution for $\mathcal{T}_i$. Let us first observe that, if $\mathsf{NW}_{\mathsf{opt}}(i) = \mathsf{NW}(T_i^*) > v_i$, then T_i^* must contain at least one outgoing edge from node i. Let (i, j) be the longest such edge in

T_i^*. For any node $k \in \mathcal{T}_i$, let $T_{i,k}^*$ denote the subtree of T_i^* rooted at k. Since $(i,j) \in T_i^*$, then $k \in D(T_i^*)$, for all $k \in \mathsf{ch}(i)$ such that $c_k \leq c_j$. So,

$$\mathsf{Worth}(T_i^*) = v_i + \sum_{k \in \mathsf{ch}(i), c_k \leq c_j} \mathsf{Worth}(T_{i,k}^*), \tag{5}$$

$$\mathsf{Cost}(T_i^*) = c_j + \sum_{k \in \mathsf{ch}(i), c_k \leq c_j} \mathsf{Cost}(T_{i,k}^*). \tag{6}$$

Let us now suppose, by contradiction, that there exists a $l \in \mathsf{ch}(i)$, with $c_l \leq c_j$ and $\mathsf{NW}(T_{i,l}^*) < \mathsf{NW}_{\mathsf{opt}}(l)$. Let T_l' denote the subtree of $\mathcal{T}_l$ yielding optimal net worth w.r.t. $\mathcal{T}_l$, that is, $\mathsf{NW}(T_l') = \mathsf{NW}_{\mathsf{opt}}(l)$. Let T_i' be the solution obtained by replacing, in T_i^*, $T_{i,l}^*$ with T_l'. Since T_i' still contains all edges (i,k), with $w(i,k) \leq w(i,j)$, we have that

$$\mathsf{Worth}(T_i') \geq \mathsf{Worth}(T_l') + v_i + \sum_{k \in \mathsf{ch}(i), c_k \leq c_j, k \neq l} \mathsf{Worth}(T_{i,k}^*). \tag{7}$$

Clearly, $\mathsf{Cost}(T_i') = \mathsf{Cost}(T_i^*) - \mathsf{Cost}(T_{i,l}^*) + \mathsf{Cost}(T_l')$. This, combined with Eq. 7, yields $\mathsf{NW}(T_i') \geq \mathsf{NW}(T_i^*) - \mathsf{NW}(T_{i,l}^*) + \mathsf{NW}(T_l')$. From the hypothesis $\mathsf{NW}(T_{i,l}^*) < \mathsf{NW}_{\mathsf{opt}}(l) = \mathsf{NW}(T_l')$, we obtain $\mathsf{NW}(T_i') > \mathsf{NW}(T_i^*)$, thus contradicting the optimality of T_i^*. So, for every $k \in \mathsf{ch}(i)$ with $c_k \leq c_j$, it must hold $\mathsf{NW}(T_{i,k}^*) = \mathsf{NW}_{\mathsf{opt}}(k)$. From Eq.s 6-5 we obtain

$$\begin{aligned} \mathsf{NW}_{\mathsf{opt}}(i) = \mathsf{NW}(T_i^*) &= v_i - c_j + \sum_{k \in \mathsf{ch}(i), c_k \leq c_j} \mathsf{NW}(T_{i,k}^*) \\ &= v_i - c_j + \sum_{k \in \mathsf{ch}(i), c_k \leq c_j} \mathsf{NW}_{\mathsf{opt}}(k). \end{aligned} \tag{8}$$

The optimality of T_i^* implies that, if $\mathsf{NW}_{\mathsf{opt}}(i) > v_i$, then j must be the node in $\mathsf{ch}(i)$ maximizing the right quantity in Eq. 8. The lemma thus follows from the fact that $\mathsf{NW}_{\mathsf{opt}}(i) \geq v_i$: taking no edges in $\mathcal{T}_i$ yield a net worth equal to v_i. This completes the proof.

Theorem 3. *For any communication graph $\mathcal{T}$ which is a tree, algorithm* Wireless_Trees *computes the optimal net worth in polynomial time, using $O(n)$ total messages and sending $O(1)$-messages per link.*

In Fig. 2 we show a distributed top-down algorithm for computing $P_A^j(\cdot)$ of Theorem 2 (see the proof in [27]). The code refers to a non-source node; the computation is initialized by s which, at the end of the bottom-up phase of Algorithm Wireless_Trees, has computed the value $\mathsf{NW}_{\mathsf{opt}} = \mu^s$ and it executes the instructions 1a-1e for every $j \in \mathsf{ch}(s)$.

Using the algorithm in Fig. 2, from Theorem 3 and from Theorem 2 we obtain the following:

Corollary 1. *If the communication graph is a tree, then the cost sharing problem on wireless networks admits a distributed polynomial-time truthful mechanism $M = (A, P_A)$ satisfying efficiency, NPT, VP, CS and CO.*

Algorithm Wireless_Trees_Pay at node $i = p(j)$

1. After receiving the message $(\mathsf{NW}_{\mathsf{opt}}, \lambda^i)$ from parent $p(i)$, for each child $j \in \mathsf{ch}(i)$ do
   (a) $\lambda^{j=0} := \mathsf{NW}_{\mathsf{opt}} - b_j$;
   (b) $x := \max_{k\in \mathsf{ch}(i), k\neq j}\{-c_k + \sum_{l\in \mathsf{ch}(i), c_l \leq c_k} \mu^l\}$;
   (c) $\lambda^{-j} := \mathsf{NW}_{\mathsf{opt}} - \mu^i + v_i + \max\{0, x\}$;
   (d) $\lambda^j := \lambda^{-j} - \lambda^{j=0}$;
   (e) send $(\mathsf{NW}_{\mathsf{opt}}, \lambda^j)$ to child j;

Fig. 2. The distributed algorithm for computing the payments in top-down fashion.

3.2 Do Good Universal Muticast Trees Exist?

In this section we propose an application of our optimal algorithm given in Sect. 3.1. In particular, we consider mechanisms that pre-compute *some* broadcast tree $\mathcal{T}$ (i.e., $D(\mathcal{T}) = \mathcal{S}$) and then solve the problem by computing an optimal subtree T of $\mathcal{T}$. Let ALG_{un} denote the resulting algorithm. The following result is a simple generalization of Corollary 1.

Theorem 4. *There exists a payment function P_{un} such that $M_{un} = (\mathsf{ALG}_{un}, P_{un})$ satisfies truthfulness, NPT, VP and CS. Moreover, if ALG_{un} runs in polynomial time, then the payment functions P_{un} are computable in polynomial time as well.*

The next result provides an upper bound on the approximability of the cost sharing problem when the input communication graph is *not* a tree. The idea is to use a shortest-path tree as universal tree.

Theorem 5. *There exists a polynomial-time mechanism $M_{un} = (\mathsf{ALG}_{un}, P_{un})$ satisfying truthfulness, NPT, VP, CS and $O(l)$-CO, where $l = |D(T)|$ and T is the computed solution. Additionally, for any $\gamma < 1$ and for all instances that admit an optimal solution T^* satisfying $|D(T^*)| \leq \gamma\frac{\mathsf{Worth}(T^*)}{\mathsf{Cost}(T^*)}$, M_{un} guarantees also $(\frac{k}{1-\gamma})$-efficiency, with $k = |D(T^*)|$. Thus, in this case, M_{un} guarantees $O(n)$-efficiency.*

The above result guarantees $O(|D(T^*)|)$-efficiency only in some cases. The next theorem rules out the possibility of obtaining polynomial-time $O(|D(T^*)|)$-efficiency in general. Its proof is a simple adaptation of an analogous result for the wired case in [15].

Theorem 6. *For any $R > 0$, no polynomial-time R-approximation algorithm (mechanism) exists, unless $\mathsf{P} = \mathsf{NP}$.*

3.3 Trees with Metric Free Edges

We consider the case in which the set of edges of the communication graph $\mathcal{G} = (\mathcal{S}, \mathcal{E})$ is partitioned into two sets $\mathcal{T} \cup \mathcal{F}$. Similarly to the case considered in Sect. 3.1, $\mathcal{T}$ induces a spanning tree of $\mathcal{G}$, i.e., $D(\mathcal{T}) = \mathcal{S}$, and, for every node i, we must select a set of outgoing edges in the set $\mathcal{T}$. However, each of these edge also induces a set of "additional connections for free" in the set $\mathcal{F}$, that is, for every $(i,k) \in \mathcal{F}$, $w(i,k) = 0$. So, adding all such connections to the solution does not increase its cost. These connections are specified as follows. Let $w^*(i,j)$ denote the weight of the path in $\mathcal{T}$ connecting the node i to one of its descendants j. Let us define $\mathsf{FREE}(i,j) := \{(i,k) | \ w^*(i,k) \leq w(i,j) \wedge (i,j) \in \mathcal{T} \wedge (i,k) \notin \mathcal{T}\}$. Moreover, for every $T \subseteq \mathcal{T}$, let $\mathcal{F}(T) := \bigcup_{(i,j)\in T} \mathsf{FREE}(i,j)$, and $\mathcal{F} := \mathcal{F}(\mathcal{T})$. We consider the restriction of the problem in which every feasible solution $C \subseteq \mathcal{T} \cup \mathcal{F}$ must fulfill the following property: C contains an edge $(i,k) \in \mathcal{F}$ if and only if it also contains an edge $(i,j) \in \mathcal{T}$ with $w^*(i,k) \leq w(i,j)$. In other words, no edge in $\mathcal{F}$ can appear as the longest outgoing edge of a node in any solution C, thus implying $\mathsf{Cost}(C) = \mathsf{Cost}(C \cap \mathcal{T})$. Observe that this definition captures some restrictions of the 2-dimensional Euclidean case (see [27] for a discussion).

The following result allows us to apply Theorem 2 and obtain a truthful mechanism.

Theorem 7. *The optimal net worth, in the case of metric free edge, of any given tree $\mathcal{T}$ can be computed in polynomial time.*

Corollary 2. *The cost sharing problem on wireless networks, in the case of trees with metric free edges, admits a polynomial-time mechanism $M = (A, P_A)$ satisfying truthfulness, efficiency, NPT, VP, CS, CO.*

3.4 Geometric Euclidean Graphs

In this section we consider so called geometric communication graphs, that is, stations are located on the d-dimensional Euclidean space, the communication graph is a *complete graph* with $w(i,j) := \mathsf{d}(i,j)^\alpha$, where $\alpha \geq 1$ is a fixed constant and $\mathsf{d} : \mathcal{R}^d \to \mathcal{R}^+$ is the Euclidean distance.

Theorem 8. *The problem of maximizing* $\mathsf{NW}(\cdot)$ *is* NP*-hard, even when restricted to geometric wireless networks, for any $d \geq 2$ and any $\alpha > 1$.*

Definition 2. *Let* $\mathsf{MST}(\mathcal{S})$ *denote the minimum spanning tree of a set of points $\mathcal{S} \subseteq \mathcal{R}^d$. Given a source node $s \in \mathcal{S}$,* $\mathsf{MST}_{\mathsf{brd}}(\mathcal{S}, s)$ *denotes the directed spanning tree obtained by considering all edges of* $\mathsf{MST}(\mathcal{S})$ *downward directed from s. Let* $\mathsf{opt}_{\mathsf{brd}}(\mathcal{S}, s)$ *denote the minimum cost among all $T \subseteq \mathcal{S} \times \mathcal{S}$ such that $D(T) = \mathcal{S}$.*

The following result concerns the problem of constructing a tree T minimizing the cost for transmitting to *all* nodes in $\mathcal{S}$:

Theorem 9. *[9, 30] For any $d \geq 1$ and for every $\alpha \geq d$, there exists a constant c_α^d such that, for any $\mathcal{S} \in \mathcal{R}^d$, and for every $s \in \mathcal{S}$,* $\mathsf{Cost}(\mathsf{MST}(\mathcal{S}, s)) \leq c_\alpha^d \cdot \mathsf{opt}_{\mathsf{brd}}(\mathcal{S}, s)$. *In particular, $c_2^2 = 12$ [30].*

Unfortunately, the same approximability result does not hold if the set of destinations is required to be a *subset* of $\mathcal{S}$ (e.g., stations that form a grid of size $\sqrt{n} \times \sqrt{n}$ and only one node with a strictly positive valuation).

A slightly better result can be obtained by using so called *Light Approximate Shortest-path Trees* (LASTs) introduced in [22]:

Theorem 10. *For any k, for any $\beta > 1$, for any $d \geq 2$, and for any $\alpha \geq d$, there exists a polynomial-time mechanism M satisfying truthfulness, NPT, VP, and CS. Additionally, M satisfies $O(1)$-CO whenever the computed solution T' satisfies $D(T') = \mathcal{S}$ or $|D(T')| \leq k$.*

We now consider the problem restricted to *linear* networks, i.e., the Euclidean case with $d = 1$. By using the result in [11] it is possible to prove the following:

Theorem 11. *The optimal h net worth*[5] *of any given communication graph $\mathcal{G}$ corresponding to a linear network can be computed in polynomial time.*

Corollary 3. *The cost sharing problem on linear wireless networks admits a polynomial time mechanism $M = (A, P_A)$ satisfying truthfulness, efficiency, NPT, VP, CS, and CO. This holds also with the additional constraint of computing multicast trees of depth at most h.*

Very recently and independently form this work, Biló *et al* [4] also consider the Euclidean case and provide $O(1)$-BB and $O(1)$-CO mechanisms.

References

1. C. Ambuehl, A. Clementi, M. Di Ianni, N. Lev-Tov, A. Monti, D. Peleg, G. Rossi, and R. Silvestri. Efficient algorithms for low-energy bounded-hop broadcast in ad-hoc wireless networks. In Proc. of STACS, LNCS 2996, pages 418–427, 2004.
2. C. Ambuehl, A. Clementi, P. Penna, G. Rossi, and R. Silvestri. Energy Consumption in Radio Networks: Selfish Agents and Rewarding Mechanisms. In Proc. of SIROCCO, pages 1–16, 2003.
3. L. Anderegg and S. Eidenbenz. Ad hoc-VCG: A Truthful and Cost-Efficient Routing Protocol for Mobile Ad Hoc Networks with Selfish Agents. In Proc. of ACM MOBICOM, pages 245 – 259, 2003.
4. V. Biló, C. Di Francescomarino, M. Flammini, and G. Melideo. Sharing the Cost of Multicast Transmissions in Wireless Networks. In Proc. of SPAA, June 2004.
5. D.M. Blough, M. Leoncini, G. Resta, and P. Santi. On the symmetric range assignment problem in wireless ad hoc networks. In Proc. of IFIP-TCS, pages 71–82. Kluwer vol. 223, 2002.
6. G. Calinescu, I.I. Mandoiu, and A. Zelikovsky. Symmetric connectivity with minimum power consumption in radio networks. In Proc. of IFIP-TCS, pages 119–130. Kluwer vol. 223, 2002.
7. S. Chawla, D. Kitchin, U. Rajan, R. Ravi, and A. Sinha. Profit maximization mechanisms for the extended multicast game. Technical report, School of Computer Science, Carnegie Mellon University, 2002.
8. E.H. Clarke. Multipart Pricing of Public Goods. Public Choice, pages 17–33, 1971.

[5] With h net worth we denote the net worth that we obtain connecting, in at most h hops, the source s with a set of nodes X

9. A. Clementi, P. Crescenzi, P. Penna, G. Rossi, and P. Vocca. On the complexity of computing minimum energy consumption broadcast subgraphs. In Proc. of STACS, LNCS 2010, pages 121–131, 2001.
10. A. Clementi, G. Huiban, P. Penna, G. Rossi, and Y.C. Verhoeven. Some recent theoretical advances and open questions on energy consumption in ad-hoc wireless networks. In Proc. of Workshop on Approximation and Randomization Algorithms in Communication Networks (ARACNE), pages 23–38, 2001.
11. A. Clementi, M. Di Ianni, and R. Silvestri. The minimum broadcast range assignment problem on linear multi-hop wireless networks. Theoretical Computer Science, 299(1-3):751–761, 2003.
12. A. Clementi, P. Penna, and R. Silvestri. Hardness results for the power range assignment problem in packet radio networks. In Proc. of APPROX, LNCS 1671, pages 197–208, 1999.
13. A. Ephremides, G.D. Nguyen, and J.E. Wieselthier. On the Construction of Energy-Effficient Broadcast and Multicast Trees in Wireless Networks. In Proc. of INFOCOM, pages 585–594, 2000.
14. J. Feigenbaum, K. Krishnamurthy, R. Sami, and S. Shenker. Hardness results for multicast cost sharing. Theoretical Computer Science, 304:215–236, 2003.
15. J. Feigenbaum, C.H. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. Journal of Computer and System Sciences, 63(1):21–41, 2001.
16. J. Feigenbaum and S. Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In Proc. of ACM DIALM, pages 1–13, 2002.
17. J. Green, E. Kohlberg, and J.J. Laffont. Partial equilibrium approach to the free rider problem. Journal of Public Econmics, 6:375–394, 1976.
18. T. Groves. Incentive in Teams. Econometrica, 41:617–631, 1973.
19. S. Guha and S. Khuller. Approximation algorithms for connected dominating sets. Algorithmica, 20(4):374–387, 1998.
20. K. Jain and V.V. Vazirani. Applications of approximation algorithms to cooperative games. In Proc. of STOC, 2001.
21. K. Jain and V.V. Vazirani. Equitable cost allocations via primal-dual-type algorithms. In Proc. of STOC, 2002.
22. S. Khuller, B. Raghavachari, and N. Young. Balancing minimum spanning trees and shortest-path trees. Algorithmica, 14(4):305–321, 1995.
23. L.M. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc. Power consumption in packet radio networks. Theoretical Computer Science, 243:289–305, 2000.
24. H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: Budget balance versus efficiency. Economic Theory, 1997.
25. N. Nisan and A. Ronen. Algorithmic Mechanism Design. Games and Economic Behavior, 35:166–196, 2001. Extended abstract in STOC'99.
26. C. H. Papadimitriou. Algorithms, Games, and the Internet. In Proc. of STOC, 2001.
27. P. Penna and C. Ventre. Sharing the Cost of Multicast Transmissions in Wireless Networks. Technical report, CRESCCO, www.ceid.upatras.gr/crescco/, 2003.
28. K. Roberts. The characterization of implementable choice rules. Aggregation and Revelation of Preferences, pages 321–348, 1979.
29. W. Vickrey. Counterspeculation, Auctions and Competitive Sealed Tenders. Journal of Finance, pages 8–37, 1961.
30. P.-J. Wan, G. Calinescu, X.-Y. Li, and O. Frieder. Minimum-energy broadcasting in static ad hoc wireless networks. Proc. of INFOCOM, 2001. Journal version in Wireless Networks, 8(6): 607-617, 2002.