Free-Riders in Steiner Tree Cost-Sharing Games*

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 consider cost-sharing mechanisms for the Steiner tree game. In this well-studied *cooperative* game, each *selfish* user expresses his/her willingness to pay for being connected to a source node s in an underlying graph. A *mechanism* decides which users will be connected and divides the cost of the corresponding (optimal) Steiner tree among these users (*budget balance* condition). Since users can form *coalitions* and misreport their willingness to pay, the mechanism must be *group strategyproof*: even coalitions of users cannot benefit from lying to the mechanism.

We present new *polynomial-time* mechanisms which satisfy a standard set of axioms considered in the literature (i.e., budget balance, group strategyproofness, voluntary participation, consumer sovereignty, no positive transfer, cost optimality) and consider the *free riders* issue recently raised by Immorlica *et al.* [SODA 2005]: it would be desirable to avoid users that are connected for free. We also provide a number of negative results on the existence of *polynomial-time* mechanisms with certain guarantee on the number of free riders. Finally, we extend our technique and results to a variant considered by Biló *et al.* [SPAA 2004] with applications to *wireless* multicast cost sharing.

1 Introduction

Consider the typical scenario in which a set U of n users wishes to *jointly* buy a certain service from some service providing company. Each user $i \in U$ has a *private* value v_i representing his/her willingness to pay for the service offered: v_i is the maximum amount of money that user i is willing to pay for the service or how much he/she would benefit from getting the service. The service provider must then develop a so called *mechanism*, that is, a policy for deciding (i) which users should be serviced and (ii) the price that each of them should pay for getting the service. Developing a fair and economically viable cost-sharing mechanisms is a central problem in *cooperative* game theory with many practical applications (see e.g [18]). In particular, due to its application to multicast and

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

A. Pelc and M. Raynal (Eds.): SIROCCO 2005, LNCS 3499, pp. 231–245, 2005.

bandwidth allocation, *Steiner tree games* (and some variants) received a lot of attention [3, 7, 2, 15]. In this game(s), a network $G = (U \cup \{s\}, E, c)$ is given, where $U = \{1, 2, \ldots, n\}$ corresponds to the set of users and s is a source node. The weight c_e of an edge $e = (i, j) \in E$ denotes the cost of connecting i to j. The (minimum) cost $C(S)$ required to connect a subset S of users to the source is the cost of a (optimal) Steiner tree containing s and S.

An important class of mechanisms is the class of *budget-balanced* mechanisms, that is, the sum of the prices charged to all users is equal to the overall cost $C(S)$ of providing the service to the subset S of users that are selected for being serviced. Observe that, given a subset of users $S \subset U$, computing the optimal cost $C(S)$ is NP-hard [5]. Also, practical applications require the mechanism to output the optimal tree connecting S to s.

In addition, users cannot be assumed to be altruistic nor obedient. Therefore, each user is considered a *selfish agent* reporting some bid value b_i (possibly different from v_i); the true value v_i is *privately known* to agent i. Based on the *reported* values $\mathbf{b} = (b_1, b_2, \ldots, b_n)$ a *mechanism* $M = (A, P)$ uses algorithm A to compute (i) a subset $S(\mathbf{b})$ of users and (ii) a Steiner tree $T(\mathbf{b})$ connecting s to $S(\mathbf{b})$ in G. Then, according to the payment vector $P = (P^1, P^2, \ldots, P^n)$, each user $i \in S(\mathbf{b})$ is charged an amount of money equal to $P^i(\mathbf{b})$. Selfish agents are assumed to be rational and thus, knowing the mechanism M, an user i could report $b_i \neq v_i$ whenever this increases his/her *utility*: given the bids $\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$ of the other agents, the utility function of user i is defined as

$$u_i^M(b_i, \mathbf{b}_{-i}) := \begin{cases} v_i - P^i(\mathbf{b}) & \text{if } i \in S(\mathbf{b}), \\ 0 & \text{otherwise.} \end{cases}$$

In [15] we provided the first mechanism for the Steiner tree game which meets all of the following axioms (previously considered in e.g. [13, 7]):

Cost Optimality (CO). Let $C_{\mathsf{opt}}(S(\mathbf{b}))$ denote the minimum cost required to service all users in $S(\mathbf{b})$. We require that the computed solution is an optimal Steiner tree for connecting the set $S(\mathbf{b})$ to the source s, that is, $\mathsf{COST}(T(\mathbf{b})) = C_{\mathsf{opt}}(S(\mathbf{b}))$.

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

Voluntary Participation (VP). We never charge an user an amount of money greater than her *reported* valuation, that is, $\forall b_i, \forall \mathbf{b}_{-i}\ b_i \geq P^i(b_i, \mathbf{b}_{-i})$. In particular, a user has always the option to not pay for being connected if he/she is not interested. Moreover, $P^i(\mathbf{b}) = 0$, for all $i \notin S(\mathbf{b})$, i.e., only the users getting connected will pay.

Consumer Sovereignty (CS). Every user is guaranteed to get the service if she reports a high enough valuation, that is, $\forall \mathbf{b}_{-i}$, $\exists \overline{b}_i$ such that $i \in S(\overline{b}_i, \mathbf{b}_{-i}) = 1$.

Budget Balance (BB).

1. **Cost recovery.** The cost of the computed solution is recovered from all the users being serviced, that is,

$$\sum_{i \in S(\mathbf{b})} P^i(\mathbf{b}) \geq \mathsf{COST}(T(\mathbf{b})).$$

2. **Competitiveness.** No surplus is created, that is,

$$\sum_{i \in (\mathbf{b})} P^i(\mathbf{b}) \not> \mathsf{COST}(T(\mathbf{b})).$$

If some surplus were created, then a competitor may offer the same service at a better price.

Group Strategyproofness (GSP). We require that a user $i \in U$ that misreport her valuation (i.e., $b_i \neq v_i$) cannot improve her utility nor improve the utility of other users without worsening her own utility (otherwise, a coalition C containing i would secede). Consider a coalition $C \subseteq U$ of users. Let $b_j = v_j$ for all $j \notin C$. Let $\mathbf{b}_C$ and $\mathbf{b}_{-C}$ denote the bid vectors of those users in C and in $U \setminus C$, respectively. The group strategyproofness requires that, if the inequality

$$u_i^M(\mathbf{b}_C, \mathbf{v}_{-C}) \geq u_i^M(\mathbf{v}_C, \mathbf{v}_{-C}) \quad (1)$$

holds for all $i \in C$, then it must hold with equality for all $i \in C$ as well.
It is easy to see that, since the above property must be fulfilled for every possible $\mathbf{v}_{-C}$, then the GSP condition can be restated by replacing $\mathbf{v}_{-C}$ by $\mathbf{b}_{-C}$. In particular, the special case of $C = \{i\}$ yields the weaker notion of *strategyproofness* (or *truthfulness with dominant strategies*): for every user i, $\forall b_i$ and $\forall \mathbf{b}_{-i}$, it holds that $u_i^M(v_i, \mathbf{b}_{-i}) \geq u_i^M(b_i, \mathbf{b}_{-i})$.

It is worth observing that the mechanism in [15] runs in *polynomial time*, in spite of the NP-hardness of the problem of computing an optimal tree for a *given* set of terminals [5]. Intuitively, the mechanism in [15] is always able to pick a set $S(\mathbf{b})$ for which the optimal Steiner tree can be computed in polynomial time, thus ensuring the CO property.

Recently, Immorlica *et al* [6] considered cost-sharing games under the additional constraint of no *free riders*: no user should be serviced for free. This work, among other results, contains a general scheme for obtaining mechanisms with no free riders and satisfying all axioms above. For the Steiner tree game, their approach *cannot* lead to polynomial time mechanism, unless $\mathsf{P} = \mathsf{NP}$. By contrast, our polynomial-time mechanism in [15] is far from the no-free-riders condition: a single user will pay for the cost of the whole tree servicing the selected users S and thus all other users are free riders.

In this work we investigate the existence of mechanisms for the Steiner tree game (and some of its variants) that possibly reduce the number of free riders, still running is polynomial time.

1.1 Previous Work and Our Contribution

The breakthrough paper by Moulin and Shenker [13, 12] provided a natural and powerful technique for building group-strategyproof mechanisms based on the following tool: a *cost-sharing* function $\xi(\cdot)$ which specifies, for each subset S of users, how the cost $C(S)$ is shared among them, that is, $\forall S \subseteq U$, $\sum_{i \in S} \xi(S, i) = C(S)$. They indeed considered a natural scheme for building a mechanism $M(\xi)$ depending on the cost-sharing function $\xi(\cdot)$ (see Fig. 1), and proved that mechanism $M(\xi)$ is group strategyproof if the function $\xi(\cdot)$ is *cross-monotonic*, that is, for all $S' \subset S \subseteq U$, and for all $i \in S'$, if holds that $\xi(S', i) \geq \xi(S, i)$.

Mechanism $M(\xi)$

1. S is initialized to U;
2. If there exists an user i in S with $v_i < \xi(S, i)$ then drop i from S. Keep repeating this step, in arbitrary order, until for every user i in S, $v_i \geq \xi(S, i)$;
3. Charge each user i an amount equal to $P^i(\mathbf{b}) := \xi(S, i)$.

Fig. 1. A general scheme to build a mechanism starting from a cost-sharing function $\xi(\cdot)$ [13]

The existence of such functions was known for the MST game in a work by Kent and Skorin-Kapov [8]. Jain and Vazirani [7] provided a more general technique for building cross-monotonic cost-sharing methods for the MST game and proved that this technique yields a polynomial-time 2-approximate budget-balanced (BB) mechanism for the Steiner tree game. [1] The technique by Jain and Vazirani [7] is based on a non-trivial use of primal-dual algorithms and inspired several works which obtain polynomial-time α-approximate BB mechanisms for other cost-sharing games (namely, metric TSP [7], facility location [14], single-source rent-or-buy [14, 10], wireless multicast [2], Steiner forest [9]).

In a recent work, Immorlica *et al* [6] provided a number of lower bounds on the approximation factor α that cross-monotonic functions can achieve for some cost-sharing games. These lower bounds do not apply in general since mechanisms *not* using cross-monotonic cost-sharing functions may exist [15, 6]. In particular, in [15] we provide a polynomial-time mechanism for the Steiner tree game which achieves all axioms above. Unfortunately, this mechanism is far from the no-free-rider condition since it always charges the cost to a single user. [2]

[1] An α-approximate BB mechanism guarantees (only) that $\mathsf{COST}(T(\mathbf{b})) \leq \sum_i P^i(\mathbf{b}) \leq \alpha \cdot C_{\mathsf{opt}}(S(\mathbf{b}))$.

[2] A similar mechanism for general cost-sharing games has been also presented in [6], though their work does not investigate computational issues for the Steiner tree game.

In this work we provide a new mechanism which guarantees that, given its computed multicast tree T and the subset S_T of users connected to s, there are at most $|S_T| - |\mathsf{leaves}(T)|$ free riders, with $\mathsf{leaves}(T)$ being the set of leaves of tree T. This mechanism still runs in polynomial time and satisfies all axioms above. We also achieve similar results for the wireless multicast game considered in [2, 15]: for this game we obtain α-approximate BB mechanism with the same bound on the number of free riders. The factor α is the same of a mechanism in [15] which, however, has $|S_T| - 1$ free riders. To the best of our knowledge, this factor α is the best known for this game.

Since, in the worst case, our mechanisms yield $|S_T| - 1$ free riders, we investigate the existence of polynomial-time mechanisms having a better guarantee. We first prove that the scheme in [6] for budget-balanced mechanism does not apply to either of our problems and, in the worst case, has the same bad guarantee (i.e., $|S_T| - 1$ free riders). This negative result applies also to exponential-time mechanisms and it is due to the fact that the optimal cost function of our games are not subadditive[3], as required in [6]. Further, we show that *any* budget-balanced mechanism which has no free riders must solve an NP-hard problem. The same negative result holds for α-approximate BB mechanisms for the wireless multicast game, for some $\alpha > 0$. In particular, the $(1+\epsilon)$-approximate BB mechanism given in [6] turns out to be intractable for this game, for small $\epsilon > 0$.

It is worth observing that, under certain hypothesis, α-approximate BB mechanisms with no free riders do not exist for the Steiner tree game, for $\alpha < 2$. This follows from previous results by Immorlica *et al* [6] and van Zwam [17]. The mechanism proposed here satisfies these additional hypothesis and, therefore, free riders cannot be avoided.

1.2 Preliminaries and Notation

Consider a graph $G = (U \cup \{s\}, E, c)$ where the set of terminals coincides with the set of users U. Given a terminal set $S \subseteq U$, we let $ST^*(S)$ denote the minimum-cost tree connecting s to the set S. The tree $MST(S)$ is (any of) the minimum spanning tree(s) over the subgraph of G induced by all and only the vertices in $S \cup \{s\}$. We consider any tree T connecting S to s as rooted at s and we denote by $\mathsf{leaves}(T)$ the set of leaves resulting from this orientation. We consider every such a tree as the set of its directed edges (a, b), where a is the parent of b. For any node a of tree T, we define $\mathsf{leaves}(T, a)$ as the set of all leaf-nodes that are descendent of a.

A *cost-sharing method* for a cost function $C(\cdot)$ is a function $\xi(\cdot)$ which distributes this cost to the users that get the service. The function $\xi(\cdot)$ takes two arguments: a set of users S and a user i and returns a *nonnegative* real number satisfying the following:

$$\text{if } i \notin S \text{ then } \xi(S, i) = 0 \text{ and} \tag{2}$$

[3] A cost function $C(\cdot)$ is subadditive if $C(S \cup \{i\}) > C(S)$, for every $S \subset U$, $i \in U \setminus S$.

$$\sum_{i \in S} \xi(S,i) = C(S). \tag{3}$$

A *β-cost-sharing method* $\xi(\cdot)$ satisfies Eq. 2 and the following relaxation of Eq. 3: $C_{\mathsf{A}}(S) \leq \sum_{i \in S} \xi(S,i) \leq \beta \cdot C_{\mathsf{A}}(S)$. Given a function $\xi : 2^U \times U \rightarrow \mathbf{R}^+ \cup \{0\}$, we define $\mathcal{P}_0^\xi := U$, and $\mathcal{P}_j^\xi := \{S \setminus \{i\} | \ S \in \mathcal{P}_{j-1}^\xi \wedge \xi(S,i) > 0\}$. Intuitively, $\mathcal{P}_j^\xi$ contains the family of all possible sets S that the scheme in Fig. 1 can generate after j users have been dropped. Thus, the set $\mathcal{P}^\xi := \bigcup_{j \geq 0} \mathcal{P}_j^\xi$ contains all possible output sets $S(\mathbf{b})$ of $M(\xi)$. A function $\xi : 2^U \times U \rightarrow \mathbf{R}^+ \cup \{0\}$ is *self cross-monotonic* if, for every $S, S' \in \mathcal{P}^\xi$ with $S' \subset S$, it holds that $\xi(S',i) \geq \xi(S,i)$, for every $i \in S'$.

A mechanism M is *upper continous* if, fixed a vector $\mathbf{b}_{-i}$, if user i is serviced for every $b_i > b$, then it is also serviced for $b_i = b$. Clearly, the mechanism $M(\xi)$ is upper continous. Let A denote an algorithm that, given a set $S \subseteq U$, returns a tree connecting S to s. We plug this algorithm into the scheme in Fig. 1 by adding a final step in which a tree $T(\mathbf{b}) := \mathsf{A}(S(\mathbf{b}))$ is output, thus obtaining a mechanism $M_{\mathsf{A}}(\xi)$ for the Steiner tree game. This defines a cost function $C_{\mathsf{A}}(\cdot)$ satisfying $C_{\mathsf{A}}(S(\mathbf{b})) = \mathsf{COST}(\mathsf{A}(S(\mathbf{b})))$, for every subset $S(\mathbf{b}) \in \mathcal{P}^\xi$.

The following result provides an useful tool for building polynomial-time (approximate) budget-balanced group strategyproof mechanisms:

Theorem 1 ([15]). *For any optimal (respectively, α-approximation) algorithm* A *and any self cross-monotonic β-cost-sharing method* $\xi(\cdot)$ *for* $C_{\mathsf{A}}(\cdot)$*, the mechanism* $M_{\mathsf{A}}(\xi)$ *is group strategyproof, β-approximate BB (respectively, αβ-approximate BB) and satisfies NPT, VP and CS. Moreover,* $M_{\mathsf{A}}(\xi)$ *runs in polynomial time if* A *and* $\xi(\cdot)$ *are polynomial time.*

Observe that, if A is optimal in $\mathcal{P}^\xi$ only (i.e., $C_{\mathsf{A}}(S) = C_{\mathsf{opt}}(S)$, for every $S \in \mathcal{P}^\xi$) then mechanisms $M_{\mathsf{A}}(\xi)$ and $M_{\mathsf{opt}}(\xi)$ will output exactly the same solutions, for any optimal algorithm opt. Hence, the following holds:

Theorem 2. *Let* $\xi(\cdot)$ *be a self-cross monotonic cost-sharing* $\xi(\cdot)$ *for* $C_{\mathsf{A}}(\cdot)$ *and let* A *be optimal in* $\mathcal{P}^\xi$*. Then, the mechanism* $M_{\mathsf{A}}(\xi)$ *is group strategyproof, budget-balanced and satisfies CO, NPT, VP and CS. Moreover,* $M_{\mathsf{A}}(\xi)$ *runs in polynomial time if* A *and* $\xi(\cdot)$ *are polynomial time. Finally,* $M_{\mathsf{A}}(\xi)$ *is upper continous.*

2 Steiner Tree Game

We will develop a cost-sharing method which charges all and only the leaf nodes of the computed solution T. In particular, each time one node is dropped, a new tree is generated by removing the corresponding leaf. The possible trees that this process can possibly generate starting from $MST(U)$ is defined as follows:

Definition 1. *We let* $T_n = MST(U)$ *and* $\mathcal{T}_n = \{T_n\}$*. Then, given* $\mathcal{T}_{j+1}$*, we define inductively* $\mathcal{T}_j := \{T | \ T \cup \{(a,b)\} \in \mathcal{T}_{j+1} \wedge b \in \mathsf{leaves}(T \cup \{(a,b)\})\}$*. Moreover, for any* $T \in \mathcal{T}_j$*, we let* S_T *denote the set of all nodes of tree* T *other than* s*.*

In the next section we prove that any of these trees is optimal, that is, its cost equals the cost of the optimal Steiner tree connecting s to the same set of nodes of the tree.

2.1 Cost Optimality

Observe that, by definition, a minimum spanning tree $MST(U)$ is also an optimal Steiner tree for the terminal set U, that is, $\mathsf{COST}(ST^*(U)) = \mathsf{COST}(MST(U))$. We next show that, starting from this tree, if we repeatedly remove any leaf node, then the tree T that we obtain remains an optimal Steiner tree for the set of nodes it contains:

Theorem 3. *For any* $T \in \mathcal{T}_j$, $\mathsf{COST}(ST^*(S_T)) = \mathsf{COST}(MST(T))$, *with* $0 \leq j \leq n$.

Proof. The proof is by induction on j, starting from $j = n$ down to $j = 0$.

Base step ($j = n$). By Def. 1 we obtain $T = T_n$ and $S_{T_n} = U$. In this case, since there are no Steiner points, then the theorem clearly holds.

Inductive step (from $j + 1$ to j). Let $T \in \mathcal{T}_j$, thus implying that there exists an edge (a, b) such that $T^b := T \cup \{(a, b)\} \in \mathcal{T}_{j+1}$ and b is a leaf in T^b. Hence, $a \in S_T$ and $S_{T^b} = S_T \cup \{b\}$.
By contradiction assume $\mathsf{COST}(T) > \mathsf{COST}(ST^*(S_T))$. We will show that there exists a tree T' spanning S_{T^b} and whose cost is lower than the cost of $ST^*(S_{T^b})$, thus contradicting the optimality of $ST^*(S_{T^b})$. The new tree T' is constructed as follows:

$$T' := \begin{cases} ST^*(S_T) \cup \{(a, b)\} & \text{if } b \notin ST^*(S_T), \\ ST^*(S_T) & \text{otherwise.} \end{cases}$$

We thus obtain

$$\mathsf{COST}(T') \leq \mathsf{COST}(ST^*(S_T)) + c_{(a,b)} < \mathsf{COST}(T) + c_{(a,b)} = \mathsf{COST}(T^b).$$

Since $a \in S_T$, then T' is a tree spanning $S_T \cup \{b\} = S_{T^b}$. By inductive hypothesis, we have $\mathsf{COST}(T^b) = \mathsf{COST}(ST^*(S_{T^b}))$, and the above inequality yields $\mathsf{COST}(T') < \mathsf{COST}(ST^*(S_{T^b}))$. This contradicts the optimality of $ST^*(S_{T^b})$.

This completes the proof.

2.2 Cost-Sharing Function

We first define a function $\xi_T(\cdot)$ that shares the cost of T among its leaves:

Definition 2. *Given a tree* T *and any* $a \in \mathsf{leaves}(T)$, *we let*

$$\xi_T(a) := \sum_{(u,v) \in T : a \in \mathsf{leaves}(T,u)} \frac{c_{(u,v)}}{|\mathsf{leaves}(T,u)|}. \tag{4}$$

Lemma 1. *For any tree T, it holds that $\sum_{a \in \mathsf{leaves}(T)} \xi_T(a) = \mathsf{COST}(T)$.*

Proof. Since for every $a \in \mathsf{leaves}(T)$ there exists $(u,v) \in T$ such that $a \in \mathsf{leaves}(T,u)$, we simply observe that

$$\sum_{(u,v)\in T} \sum_{a \in \mathsf{leaves}(T,u)} \frac{c_{(u,v)}}{|\mathsf{leaves}(T,u)|} = \sum_{a \in \mathsf{leaves}(T)} \sum_{\substack{(u,v)\in T: \\ a \in \mathsf{leaves}(T,u)}} \frac{c_{(u,v)}}{|\mathsf{leaves}(T,u)|}$$

The left hand side is $\mathsf{COST}(T)$, while the right hand side is $\sum_{a \in \mathsf{leaves}(T)} \xi_T(a)$. The lemma thus follows.

We use the trees $\mathcal{T}_j$ and the associated functions $\xi_T(\cdot)$ for defining a self cross-monotonic cost-sharing function $\xi_{\mathsf{leaves}}(\cdot)$ as follows:

$$\xi_{\mathsf{leaves}}(S_T, a) = \begin{cases} \xi_T(a) \text{ if } a \in \mathsf{leaves}(T), \\ 0 \qquad\text{otherwise,} \end{cases} \tag{5}$$

for every $T \in \mathcal{T}_j$, with $0 \le j \le n$.

Lemma 2. *Let $\xi = \xi_{\mathsf{leaves}}$. Then $\mathcal{P}_j^{\xi} = \bigcup_{T \in \mathcal{T}_j} S_T$, for every $0 \le j \le n$.*

Theorem 4. *The function $\xi_{\mathsf{leaves}}(\cdot)$ is self cross-monotonic.*

Proof. It is easy to see that, since $\xi_{\mathsf{leaves}}(\cdot)$ is non-zero only for the leaf nodes, then starting from $T_n = MST(U)$, at each step the mechanism $M(\xi_{\mathsf{leaves}})$ will consider a tree $T \in \mathcal{T}_j$ and remove some leaf b. Let $T' = T \setminus (a,b)$ denote the new tree obtained in this way. We prove that $\xi_{\mathsf{leaves}}(S_T, i) \le \xi_{\mathsf{leaves}}(S_{T'}, i)$, for every $i \in S_{T'}$. Since $S_{T'} = S_T \setminus \{b\}$, it holds that $i \ne b$. For every $(u,v) \in T'$, $|\mathsf{leaves}(T,u)| \ge |\mathsf{leaves}(T',u)|$. Moreover, by definition of T', if $i \in \mathsf{leaves}(T,u)$ for some edge (u,v), then $i \in \mathsf{leaves}(T',u)$. Therefore, if $i \in \mathsf{leaves}(T)$, we obtain

$$\begin{aligned} \xi_{\mathsf{leaves}}(S_T, i) &= \sum_{\substack{(u,v)\in T: \\ i \in \mathsf{leaves}(T,u)}} \frac{c_{(u,v)}}{|\mathsf{leaves}(T,u)|} \le \sum_{\substack{(u,v)\in T: \\ i \in \mathsf{leaves}(T,u)}} \frac{c_{(u,v)}}{|\mathsf{leaves}(T',u)|} \\ &= \sum_{\substack{(u,v)\in T': \\ i \in \mathsf{leaves}(T',u)}} \frac{c_{(u,v)}}{|\mathsf{leaves}(T',u)|} = \xi_{\mathsf{leaves}}(S_{T'}, i). \end{aligned}$$

Otherwise, that is $i \notin \mathsf{leaves}(T)$, it simply holds $\xi_{\mathsf{leaves}}(S_T, i) = 0 \le \xi_{\mathsf{leaves}}(S_{T'}, i)$.

Now consider any two trees $T \in \mathcal{T}_j$ and $T' \in \mathcal{T}_{j-k}$ with $S_{T'} \subset S_T$. By definition, there exists a sequence of trees $T_1, T_2, \ldots, T_{k+1}$, with $T_1 = T$ and $T_k = T'$, which is obtained by repeatedly removing some leaf node of T_l, with $1 \le l \le k-1$. The above argument thus yields $\xi_{\mathsf{leaves}}(S_{T_1}, i) \le \xi_{\mathsf{leaves}}(S_{T_2}, i) \le \cdots \le \xi_{\mathsf{leaves}}(S_{T_k}, i)$. The theorem thus follows from Lemma 2.

2.3 Analysis

We first observe that Lemma 2 and Theorem 3 imply that MST is optimal on $\mathcal{P}^{\xi_{\mathsf{leaves}}}$. Hence, by applying Theorem 2 with $\xi = \xi_{\mathsf{leaves}}$, we obtain the following result:

Corollary 1. *The Steiner tree game admits a mechanism $M_{MST}(\xi)$ which is polynomial-time, group strategyproof, budget-balanced and satisfies CO, NPT, VP and CS. The mechanism is upper continous and garantees that, if T is the tree given in output, then there are at most $|S_T| - |\mathsf{leaves}(T)|$ free riders.*

3 Wireless Multicast Game

A variant of the Steiner tree problem considered in [2, 15] is the *wireless multicast* which is defined as follows. Each node of the graph G corresponds to a station of an ad-hoc network. Stations are located on a c-dimensional Euclidean plane and, given the distance $d(i, j)$ between i and j, the cost of connecting i to j is $c_{(i,j)} := d(i, j)^\gamma$, for some $\gamma > 1$. This quantity represents the power that station i must spend to transmit directly to j. Thus, given a multicast tree T, its cost is the *overall power consumption*, that is,

$$\mathsf{POW}(T) := \sum_{i \in U \cup \{s\}} \max_{(i,j) \in T} \{c_{(i,j)}\},$$

that is, every node contributes as the cost of its longest outgoing edge in T. The game is thus defined as the Steiner tree game with the only difference that the cost function $\mathsf{COST}(\cdot)$ is replaced by the function $\mathsf{POW}(\cdot)$ above. As in [2, 15], we will consider $\gamma \geq c$, since for $\gamma < c$ no approximation algorithm in known; moreover, in many applications $\gamma \geq 2$ and stations are located on the plane (i.e., $c = 2$).

Since $\mathsf{POW}(T) < \mathsf{COST}(T)$ whenever T as at least two leaves, if we apply the payment scheme for the Steiner tree game, the mechanism will create some *surplus*, that is, users will pay more than the cost. We next modify the payment scheme so to avoid this.

Given a tree T, and any node i, let $e_1(i), e_2(i), \ldots, e_k(i)$ denote the list of nodes in T outgoing from i and satisfying $c_{e_j(i)} \leq c_{e_{j+1}(i)}$ (ties are broken arbitrarily). We define a function $w(\cdot)$ which weights the edges of T according to their contribution to $\mathsf{POW}(T)$. In particular, let $\mathsf{mc}(e_j(i))$ denote the *marginal contribution*[4] of edge $e_j(i)$, that is,

$$\mathsf{mc}(e_j(i)) := \begin{cases} c_{e_j(i)} - \max\{c_{e_l(i)} | \ c_{e_l(i)} < c_{e_j(i)}\} & \text{if } j > 1, \\ c_{e_j(i)} & \text{otherwise.} \end{cases} \quad (6)$$

[4] A similar concept is employed in [2] for defining the Shapley value of a wireless multicast tree.

Finally, defined

$$\mathsf{equal}(e_j(i)) := |\{e_l(i) | c_{e_l(i)} = c_{e_j(i)}\}|$$

we let

$$w_{e_j(i)} := \mathsf{mc}(e_j(i))/\mathsf{equal}(e_j(i)). \tag{7}$$

By definition, $\sum_{j=1}^{k} w_{e_j(i)} = c_{e_k(i)}$, thus implying $\mathsf{POW}(T) = \sum_{e \in T} w_e$. We can share this cost by considering the graph G with edge costs c_e replaced by w_e, for every $e \in T$. This idea leads to the following result:

Theorem 5. *Let w_e be defined as in Eq. 7 with respect to $T = MST(U)$. Also let $\xi_{\mathsf{wireless}}(\cdot)$ be defined as in Eq.s 4-5 by replacing c_e with w_e, for every $e \in T$. Then, the function $\xi_{\mathsf{wireless}}(\cdot)$ is a self-cross monotonic cost-sharing method for $C_{\mathsf{A}}(S) := \mathsf{POW}(MST(S))$.*

We can modify the mechanism for the Steiner tree game by replacing $\xi_{\mathsf{leaves}}(\cdot)$ by $\xi_{\mathsf{wireless}}(\cdot)$, thus obtaining a polynomial-time mechanism $M_{MST}(\xi_{\mathsf{wireless}})$ which outputs a tree $T \in \mathcal{T}$ (see Def. 1) and that recovers the corresponding cost $\mathsf{POW}(T)$. In [4] the authors proved that, for any $\gamma \geq \delta$, the total weight of a MST over a set S of points is at most $(3^c - 1)$ times the cost of the optimal wireless broadcast tree for S. (For $c = 2 \leq \gamma$ the constant has been improved down to 7.5 [4].) Thus, using an argument similar to [2, 15], Theorem 3 yields the following result:

Corollary 2. *The wireless multicast game admits a polynomial-time mechanism $M_{MST}(\xi)$ which is group strategyproof, $(3^c - 1)$-approximate BB and satisfies NPT, VP and CS. The mechanism is upper continous and guarantees that, if T is the tree given in output, then there are at most $|S_T| - |\mathsf{leaves}(T)|$ free riders. Additionally, for $c = 2$, mechanism $M_{MST}(\xi_{\mathsf{wireless}})$ is 7.5-approximate BB.*

The above result improves over the mechanism in [15] since in this mechanism there are always $|S_T| - 1$ free riders. In the next section we compare our mechanism with other mechanisms.

4 Negative Results and Open Questions

Immorlica *et al* [6] proposed two mechanisms for avoiding free riders in a cost-sharing game with cost function $C(\cdot)$. The first mechanism is budget-balanced and works for *subadditive* functions $C(\cdot)$:

Mechanism IMM-BUDGET-BALANCE

1. Initialize the set S of serviced users to the empty set and the amount of money m that is already charged to 0;
2. For i from 0 to n, do the following: if $b_i \geq \min\{C(\{i\}), C(S \cup \{i, \ldots, n\}) - m\}$, then include i in S, and charge him/her $\min\{C(\{i\}), C(S \cup \{i, \ldots, n\}) - m\}$ (therefore, m will be increased by this quantity).

The following fact shows that, in our games, the optimal cost function is *not* subadditive and, more importantly, mechanism IMM-BUDGET-BALANCE does *not* guarantee any bound on the number of free riders:

Fact 6. *There exists an instance of the Steiner tree game for which the mechanism* IMM-BUDGET-BALANCE *yields* $n-1$ *free riders. The same result holds for the wireless multicast game.*

Proof. Consider a clique K_n of n nodes and let any of its edges have cost 0. We build a graph G by connecting a new node s to every node of K_n with an edge of cost 1. Then, for every $S \subseteq U$, with $S \neq \emptyset$, $C(S) = 1$. Hence, mechanism IMM-BUDGET-BALANCE charges 0 to all but one user in the final set S.

A similar argument applies to the following instance of the wireless multicast game: consider the star G connecting s to $n \geq 2$ nodes at distance 1 from s. In this case, there will be, in the worst case, $n-1$ free riders.

We also mention that, for the Steiner tree game, a known result by Megiddo [11], combined with a characterization of upper continous budget-balanced mechanisms without free riders [6–Th. 4.3], implies that the above (upper continous) mechanism cannot guarantee the no-free-rider condition even if edges have nonzero weight.

We next show that, unless $\mathsf{P} = \mathsf{NP}$, there exist no polynomial-time mechanism satisfying all axioms above plus the no-free-rider condition.

Theorem 7. *Let* $C_{\mathsf{opt}}(\cdot)$ *be a cost function which is* NP*-hard (to approximate within a factor* α*). Then, no polynomial-time strategyproof mechanism* M *guarantees no free riders and satisfies NPT, VP, CS and (*α*-approximate) BB, unless* $\mathsf{P} = \mathsf{NP}$.

Proof Sketch. The strategyproofness, NPT, BB (no surplus part) and CS conditions imply that, for every $S \subseteq U$, there exists a bid vector $\mathbf{b}^S = (\mathbf{b}_S, \mathbf{0}_{-S})$ such that $S(\mathbf{b}^S) \supseteq S$. The VP and the no-free-rider condition imply that $S(\mathbf{b}^S) = S$. Hence, the (α-approximate) BB condition implies that M is able to compute (an α-approximation of) the optimum $C_{\mathsf{opt}}(S)$ in polynomial time, thus implying $\mathsf{P} = \mathsf{NP}$. (See Appendix 5 for the full proof of this theorem.) □

Since the Steiner tree problem cannot be approximated within some factor $\alpha > 1$, then it is not possible to have polynomial time α-approximate BB mechanisms and no free riders. Hence, the following mechanism for $(1+\epsilon)$-approximate BB mechanisms [6] cannot run in polynomial time for small $\epsilon > 0$:

Mechanism IMM-$(1+\epsilon)$-APPROXIMATE-BB

1. Drop all users i such that $b_i \leq \delta$ and let $R = \{r_1, r_2, \ldots, r_{|S|}\}$ be the (arbitrarily) ordered set of remaining users; /* the value of δ depends on $\epsilon > 0$ */
2. Find the first user $r_i \in R$ such that $b_{r_i} \geq C(\{r_i, r_{i+1}, \ldots, r_{|S|}\}) - n\delta$; The set $S := \{r_i, r_{i+1}, \ldots, r_{|S|}\}$ is serviced, user i pays $C(S) - n\delta$, and every body else in S pays δ.

Observe that, for the wireless multicast game, the mechanism above would be polynomial-time only if a polynomial-time $(1+\epsilon)$-approximation algorithm for the wireless *multicast* optimization problem exists. To the best of our knowledge, the best factor is achieved by a combination of the best known Steiner tree algorithm [16] with the results in [4]. This combination (see [1]) yields an upper bound which is *larger* than the approximation factor $(3^c - 1)$ for the wireless broadcast [4]. Whether such bounds are tight is an interesting open problem. Below we mention other questions that remain open.

4.1 Open Questions

In view of these negative results, one may try to improve our mechanisms along two directions: (i) decrease the number of free riders and/or (ii) improve the approximation factor of the mechanism for wireless multicast game.

As already mentioned, a result by Immorlica *et al* [6] characterizes the class of upper continous α-approximate budget-balanced mechanisms with no free riders as the mechanisms which can be obtained using α-cross-monotonic cost-sharing methods. Hence, the lower bounds on cross-monotonic cost-sharing methods [11, 2] imply that, for the two games considered here, mechanisms without free riders, if any, cannot be obtained directly from the scheme in Fig. 1 or from mechanism IMM-BUDGET-BALANCE (both of them being upper continous). For the Steiner tree game, a tight result by van Zwam [17], implies that the polynomial-time 2-approximate BB mechanisms by [7] are the best possible in the class of upper continous one.

Concerning polynomial-time mechanisms and the free rider issue, the following question is interesting to us: Is there any such mechanism which guarantees that a constant fraction of the serviced set S_T are not free riders? First of all, it is well-known that every node in G of degree 2 cannot be a Steiner point if G is metric. Hence, these nodes could also be charged and participate to the payments. Unfortunately, nodes that in tree $T_n = MST(U)$ have degree 2 *can* be Steiner points and their removal could lead to suboptimal solutions.

Generally speaking, the main question left open is to find a trade off between the number of free-riders and the approximation factor of the budget balance for polynomial-time mechanisms.

References

1. E. Althaus, G. Calinescu, I. Mandoiu, S. Prasad, N. Tchervenski, and A. Zelikovsky. Power efficient range assignment in ad-hoc wireless networks," submitted journal version. *Preliminary results in WCNC '03 and TCS '03*, 2003. Submitted journal version available at http://www.cs.iit.edu/∼calinesc/.
2. V. Biló, C. Di Francescomarino, M. Flammini, and G. Melideo. Sharing the cost of multicast transmissions in wireless networks. In *Proceedings of SPAA*, pages 180–187, 2004.
3. 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.

4. M. Flammini, A. Navarra, R. Klasing, and S. Perennes. Improved approximation results for the minimum energy broadcasting problem. In *Proc. of the joint workshop on Foundations of mobile computing (DIALM-POMC)*, pages 85–91. ACM Press, 2004.
5. M.R. Garey and D.S. Johnson. *Computers and intractability: a guide to the theory of NP-completeness*. Freeman, 1979.
6. N. Immorlica, M. Mahdian, and V. Mirrokni. Limitations of cross-monotonic cost-sharing schemes. In *Proc. of the 16th Annual ACM Symposium on Discrete Algorithms (SODA)*, January 2005.
7. K. Jain and V.V. Vazirani. Applications of approximation algorithms to cooperative games. In *Annual ACM Symposium on Theory of Computing (STOC)*, pages 364–372, 2001.
8. K. Kent and D. Skorin-Kapov. Population monotonic cost allocation on MST's. In *Operational Research Proceedings KOI*, volume 43-48, 1996.
9. J. Könemann, S. Leonardi, and G. Schäfer. A group-strategyproof mechanism for Steiner forests. In *Proc. of the 16th Annual ACM Symposium on Discrete Algorithms (SODA)*. ACM Press, January 2005. To appear.
10. S. Leonardi and G. Schäfer. Cross-monotonic cost-sharing methods for connected facility location games. Technical Report MPI-I-2003-1-017, Max-Plankt-Institut für Informatik, September 2003. Also in Proc. of 5th ACM Conference on Electronic Commerce (EC), pages 242-243, ACM Press, 2004.
11. N. Megiddo. Cost allocation for Steiner trees. *Networks*, 8:1–6, 1978.
12. H. Moulin. Incremental cost sharing: characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16:279–320, 1999.
13. H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. 1997. http://www.aciri.org/shenker/cost.ps.
14. M. Pàl and É. Tardos. Strategy proof mechanisms via primal-dual algorithms. In *Proc. of the 44th FOCS*, 2003.
15. P. Penna and C. Ventre. More Powerful and Simpler Cost-Sharing Methods. In *Proceedings of the 2nd Workshop on Approximation and Online Algorithms (WAOA)*, number 3351 in LNCS, September 2004. To appear in the LNCS series.
16. A. Robins and Zelikovsky. Improved approximation Steiner tree in graphs. In *Proc. of the 11th Annual ACM Symposium on Discrete Algorithms (SODA)*, pages 770–779, 2000.
17. S. van Zwam. A lower bound on the cost recovery of the Steiner tree game with cross-monotonic cost shares. Technical Report 18-04, Dipartimento di Informatica e Sistemistica, Università di Roma "La Sapienza", 2004.
18. H.P. Young. *Cost Allocation*, volume 2, chapter 34 of Handbook of Game Theory, volume 2, edited by R.J. Aumann and S. Hart, pages 1193–1235. Elsevier Science, 1994.

5 Proof of Theorem 7

We next show two useful lemmata which state some basic properties of strategyproof mechanisms. Basically, these results concern how changing one bid in the vector $\mathbf{b}$ can affect the outcome of the mechanism (i.e., the serviced set and the computed payments). Given a mechanism $M = (A, P)$, we let $\sigma_i^A(\mathbf{b}) = 1$

if and only if $i \in S(\mathbf{b})$, with $S(\mathbf{b})$ being the subset of users that algorithm A decides to service on input $\mathbf{b}$.

The following lemma states "steady" conditions on strategyproof mechanisms:

Lemma 3 (keep). *Let $M = (A, P)$ be a strategyproof mechanism and let $S(\mathbf{b})$ denote the set of users serviced on input $\mathbf{b}$. Then, the following conditions must hold:*

$$\sigma_i^A(b_i, \mathbf{b}_{-i}) \Rightarrow \forall b_i' > b_i,\ i \in S(b_i', \mathbf{b}_{-i}); \tag{8}$$

$$\sigma_i^A(b_i, \mathbf{b}_{-i}) = \sigma_i^A(b_i', \mathbf{b}_{-i}) \Rightarrow\ P^i(b_i, \mathbf{b}_{-i}) = P^i(b_i', \mathbf{b}_{-i}). \tag{9}$$

Proof. (Eq. 8). By contradiction, if $\sigma_i^A(b_i', \mathbf{b}_{-i}) = 0$ and $\sigma_i^A(b_i, \mathbf{b}_{-i}) = 1$, for some $b_i' > b_i$ and for some $\mathbf{b}_{-i}$, then consider the situation in which $v_i = b_i'$ and $\mathbf{v}_{-i} = \mathbf{b}_{-i}$. In this case, we obtain

$$v_i \cdot \sigma_i^A(b_i, \mathbf{v}_{-i}) - P^i(b_i, \mathbf{v}_{-i}) \geq v_i - b_i = b_i' - b_i > 0 \geq v_i \cdot \sigma_i^A(v_i, \mathbf{v}_{-i}) - P^i(v_i, \mathbf{v}_{-i}),$$

where the last inequality follows from the fact that $\sigma_i^A(v_i, \mathbf{v}_{-i}) = \sigma_i^A(b_i', \mathbf{v}_{-i}) = 0$ and from the NPT condition. This contradicts the fact that M is strategyproof.

(Eq. 9). By contradiction, let us assume that $\sigma_i^A(b_i, \mathbf{b}_{-i}) = \sigma_i^A(b_i', \mathbf{b}_{-i})$ and $P^i(b_i, \mathbf{b}_{-i}) < P^i(b_i', \mathbf{b}_{-i})$, for some b_i, b_i' and $\mathbf{b}_{-i}$. Consider the situation in which $v_i = b_i'$, thus implying

$$v_i \cdot \sigma_i^A(b_i, \mathbf{v}_{-i}) - P^i(b_i, \mathbf{v}_{-i}) > v_i \cdot \sigma_i^A(v_i, \mathbf{v}_{-i}) - P^i(v_i, \mathbf{v}_{-i}),$$

thus contradicting the fact that M is strategyproof.

The following lemma states that, if users "compete" with each other for being serviced, then the prices cannot be bounded from above by any constant:

Lemma 4 (drop). *Let $M = (A, P)$ a strategyproof mechanism satisfying NPT and CS. For every $\mathbf{b} = (b_1, \ldots, b_n)$, if there exists $i, j \in U$ and $\overline{b}_j$, such that*

$$\sigma_i^A(\mathbf{b}) = 1; \tag{10}$$

$$\sigma_i^A(\overline{\mathbf{b}}) = 0, \ \ \textit{with}\ \overline{\mathbf{b}} = (\overline{b}_j, \mathbf{b}_{-j}). \tag{11}$$

Then there exists a $\overline{b}_i$ such that $P^i(\mathbf{b}') \geq b_i$, where $\mathbf{b}' = (\overline{b}_i, \overline{\mathbf{b}}_{-i})$.

Proof. From the CS condition, there exists $\overline{b}_i$ such that $\sigma_i^A(\overline{b}_i, \overline{\mathbf{b}}_{-i}) = 1$, where $\overline{\mathbf{b}} = (\overline{b}_j, \mathbf{b}_{-j})$. By contradiction, assume that $P^i(\overline{b}_i, \overline{\mathbf{b}}_{-i}) < b_i$. Consider the case in which $v_i = b_i$ and $\mathbf{v}_{-i} = \overline{\mathbf{b}}_{-i}$. Because of the NPT and Eq. 11, it holds that $u_i^M(v_i, \mathbf{v}_{-i}) = u_i^M(b_i, \overline{\mathbf{b}}_{-i}) \leq 0$. Moreover, $u_i^M(\overline{b}_i, \mathbf{v}_{-i}) = u_i^M(\overline{b}_i, \overline{\mathbf{b}}_{-i}) = v_i - P^i(\overline{b}_i, \overline{\mathbf{b}}_{-i}) > 0$, thus contradicting the fact that $M = (A, P)$ is strategyproof.

The above lemma easily implies the following result.

Lemma 5. *Let $M = (A, P)$ a strategyproof mechanism satisfying NPT and CS. For every $S \subseteq U$ and for every $\overline{b} > 0$, there exists a vector $\mathbf{b}^S = (\mathbf{b}_S, \mathbf{0}_{-S})$ such that, if mechanism M returns a set $S(\mathbf{b}^S)$ of users with $S \not\subseteq S(\mathbf{b}^S)$, then there exists $i \in U$ such that $P^i(\mathbf{b}) > \overline{b}$.*

Proof. Simply increase the bids of users in S one by one, from 0 to a value $\overline{b}_i \geq \overline{b}$ such that the current user must be serviced. (This value exists because of the CS condition.) If at some point, a user $i \in S$ previously considered is dropped, because of Lemma 4, there exists a bid vector $\mathbf{b}'$ for which $P^i(\mathbf{b}') > \overline{b}$.

Let $M = (A, P)$ be a polynomial-time mechanism satisfying the hypothesis of Theorem 7. By contradiction, let M run in polynomial time. The above lemma and the (α-approximate) BB imply $S \subseteq S(\mathbf{b})$. The no-free-rider and VP conditions thus yield $S(\mathbf{b}) = S$. Hence, the (α-approximate) BB condition implies that M is able to compute (an α-approximation of) the optimum $C_{\mathsf{opt}}(S)$ in polynomial time, thus implying $\mathsf{P} = \mathsf{NP}$.