Group Strategyproofness and No Subsidy via LP-Duality

Kamal Jain[*] *Vijay V. Vazirani*[*]

Abstract

We make two contributions to co-operative game theory, both of which rely heavily on linear programming duality theory. First, we introduce a fairness criterion on service providers, which we call *No Subsidy*. This condition ensures that a subset of the users is not overcharged in order to subsidize another set of users. We completely characterize cost functions that allow cost sharing mechanisms satisfying this condition. This class of cost functions contains the submodular functions, and can serve the purpose of formally capturing our intuitive notion of functions that exhibit economies of scale.

Besides the well-studied class of non-decreasing, submodular cost functions, the only other class known to admit a budget balanced and group strategyproof cost sharing mechanism is the class of supermodular functions. In our second result, we give a budget balanced and group strategyproof cost sharing method for a cost function that is not in either of these classes. Our cost function arises naturally in the context of correcting an inefficiency of the scheme of multicast routing. Our cost sharing method utilizes the dual solution produced by the primal-dual algorithm for the minimum spanning tree problem, based on the bidirected cut relaxation for this problem.

1 Introduction

Consider a service providing company whose set of possible costumers, also called *users*, is U. For each set $S \subseteq U$ $C(S)$ denotes the cost incurred by the company to serve the users in S. The function C is known as the *cost function*. For concreteness, assume that the company broadcasts news of common interest, such as financial news, on the net. Each user, i, has a utility, u_i', for receiving the news. This utility u_i' is known only to user i. User i enjoys a benefit of $u_i' - x_i$ if she gets the news at the price x_i. If she does not get the news then her benefit is 0. Each user is assumed to be selfish, and hence in order to maximize benefit, may misreport her utility as some other number, say u_i. For the rest of the discussion, *the utility of user i* will mean the number u_i.

A *cost sharing mechanism* determines which users receive the broadcast and at what price. The mechanism is *strategyproof* if the dominant strategy of each user is to reveal the true value of her utility. It is said to be *group strategyproof* if this holds for coalition as well. A cost sharing mechanism is *budget balanced* if the total amount it charges from the receivers is same as the cost incurred by the service provider, $C(S)$. It is *efficient* if it maximizes, over all subsets, S, the sum of the utilities of users in S minus $C(S)$. Ideally, one seeks an efficient, budget balanced and group strategyproof cost-sharing mechanism. A classical result in game theory shows that such

[*]Address: College of Computing, Georgia Institute of Technology, Atlanta, GA 30332–0280, E-mail: {kjain,vazirani}@cc.gatech.edu. Research supported by NSF Grant CCR-9820896.

a mechanism does not exist even for a submodular cost function [6, 14]. Such a mechanism does not exist even after relaxing the condition of group strategyproofness to just strategyproofness.

In view of this limitation, the area of co-operative game theory studies desirable properties that can be achieved by cost sharing mechanisms for specific cost functions. The two best understood cost functions are the submodular and the supermodular cost functions. For these cases, a complete characterization of budget balanced and group strategyproof cost sharing mechanism is known [10, 12]. A cost function is *submodular* if

1. $C(\emptyset) = 0$,
2. for any Q_1 and Q_2, $C(Q_1) + C(Q_2) \geq C(Q_1 \cup Q_2) + C(Q_1 \cap Q_2)$.

It is *supermodular* if the inequality in the second condition is reversed. A cost function is *non-decreasing* if for $Q_1 \subseteq Q_2$, $C(Q_1) \leq C(Q_2)$.

The emphasis in the above mentioned studies is to ensure that users do not misreport their utility values. On the other hand, the issue of ensuring fairness from the service provider seems to have been ignored. In this paper, we rectify this deficiency by introducing the condition of *No Subsidy* on a cost sharing mechanism. This condition ensures that the service provider does not overcharge a subset of the users in order to subsidize another set of users. We characterize completely the class of cost functions that admit a No Subsidy mechanism. This is the class of functions that exhibit the *covering property*. To our knowledge, this is the first theorem of its kind, characterizing all cost functions that allow a certain type of cost sharing mechanism. The proof relies on a non-trivial use of the strong duality theorem of linear programming. By solving a linear program optimally, one can also find a budget balanced No Subsidy mechanism in polynomial time. Traditionally, the latter theorem has played a central role in the study of two person games.

The class of functions that exhibit the covering property includes all submodular functions. Many authors formally capture the notion of *economies of scale* by submodular functions. It is interesting to note that our larger class can also serve this purpose. Hence we open the question of characterizing all budget balanced and group strategyproof mechanism for this class of function.

A fundamental theorem of Moulin and Shenker [12, 10] shows that a cross-monotonic cost sharing method (see Section 2 for a definition) gives rise to a budget balanced and group strategyproof mechanism (we provide a simpler proof of this result). Moreover, if the cost function is submodular, the converse also holds. Two important cost sharing methods proven to be cross-monotonic, if the cost function is non-decreasing and submodular, are:

1. Shapley Value [15]: In the context of multicasting, this method distributes the cost of each edge equally among all the users located downstream of the edge. Moulin and Shenker show that this method achieves the lowest worst case loss of efficiency over all utility profiles [12].
2. Egalitarian [3]: This method, due to Dutta and Ray, seeks equality among all the receivers. Mutuswami has shown that assuming all users draw their utility values from the same probability distribution (with some technical restrictions) this method maximizes the expected size of receivers [11].

The only other class of cost functions known to admit a budget balanced and group strategyproof cost sharing method is the class of supermodular cost functions [10]. The second contribution of our paper is to give a budget balanced and group strategyproof cost sharing method for a cost function that is not in either of these classes. Our cost function arises naturally in the context of correcting an inefficiency of the scheme of multicast routing.

The simplest way of broadcasting a message to many receivers is *unicasting*, under which the message is individually sent to each receiver. This results in several copies of the same message traversing the network links. This waste of bandwidth could be avoided by using a different form of routing, called *multicast routing* [2]. Multicast routing uses a tree connecting all the receivers to the source. The source sends one copy of the message to each neighboring vertex in the tree. These vertices further act as sources for the downstream receivers. In this manner, a source can reach many receivers without sending multiple copies of a message over any link. Traditionally, a tree containing the source and all possible receivers is fixed. The network is then restricted to this tree only. Whenever a message needs to be broadcast to a subset of receivers, multicast routing picks a subtree of this tree. It therefore ignores the high connectivity offered by today's Internet. This is not desirable when the receiver set varies with the message, since the subtree to certain a subset of receivers may be arbitrarily more costly than the cheapest tree to them. On the positive side, Feigenbaum, Papadimitriou and Shenker [5] have recently shown that the marginal cost cost-sharing mechanism can be implemented on this model with the overhead of only two messages per link. The marginal cost cost-sharing mechanism is efficient and strategyproof. However, it is not budget balanced, and may create a deficit.

When the network is not restricted to a tree, the problem of finding the cheapest multicast tree for a given set of receivers is the well-known Steiner tree problem, which is **NP**-hard [8]. Although constant factor approximation algorithms are known for the Steiner tree problem, we do not know of any algorithm which gives either a non-decreasing or a submodular cost function. In fact, the optimal solution also does not give a submodular cost function (for example, consider a four cycle with unit cost edges whose vertices are three receivers and the source). (The fractional optimal solution to the two well-known linear programming relaxations of the Steiner tree problem, undirected cut formulation and bidirected cut formulation, also does not give a submodular cost function.)

In this paper, we give a budget balanced and group strategyproof cost sharing method using a factor 2 approximation algorithm for the optimal Steiner tree problem. If the capacity of each link is infinite then without loss of generality we can assume that the link costs satisfy the triangle inequality. Under this assumption one can show that the minimum spanning tree over a set of receivers and the source is within twice the cost of the cheapest multicast tree [9]. As stated above, the cost function produced by this algorithm is neither non-decreasing nor submodular e.g., a wheel on 5 vertices, where the source is one of the peripheral vertex.

We use Edmonds' exact linear programming relaxation for the minimum spanning tree problem i.e., a relaxation that always has optimal integral solutions [4]. The primal-dual algorithm [4, 13] produces a minimum spanning tree and a dual solution such that the cost of the minimum spanning tree is equal to the value of the dual solution. Our cost sharing method simply distributes each dual equally among all the receivers defining this dual. Clearly, this cost sharing method is budget balanced. We show that it is cross-monotonic as well. In addition, we show that the resulting cost sharing mechanism satisfies the No Subsidy condition as well.

Recently, Goemans and Skutella [7] considered the facility location problem in which we are given a set of cities and a set of potential sites for opening facilities. For each site, we are given the cost of opening a facility there. For each city and site pair, we are given the cost of connecting the city to a facility opened at that site. The aim is to open a set of facilities and connect each city to one of the open facilities so that the sum of the opening and connection costs is minimized. Goemans and Skutella consider the problem of allocating the cost of serving all users in a *fair* manner. Notice that this is not a game theoretic setting since *all* users are being served. Their definition of fairness is that no subset of users be charged more than the cost of serving this subset *optimally*. They give a budget balanced and fair cost allocation only for instances for which the LP-relaxation (due to Balinski [1]) of this problem has an integral optimal solution. The latter restriction is quite severe. Our definition of No Subsidy is inspired by their work.

2 The model

For convenience, let us consider the setting of multicasting for giving the basic definitions and theorems; they are of course valid in the general setting.

Consider an undirected graph $G = (V, E)$ with edge weights c_e's and a marked node *root*. All other nodes are *users*. We denote the set of all users by U. The root is the broadcasting source. It broadcasts a message to a selected set of users. Each user, i, has a non-negative utility, u'_i, for receiving the message. This utility u'_i is known only to user i. User i enjoys a benefit of $u'_i - x_i$ if she gets the message at the price x_i. If she does not get the message then her benefit is 0. Each user is assumed to be selfish and hence in order to maximize the benefit a user can misreport her utility as some other number, say u_i.

In this model we assume that a message can be duplicated at any node at no cost. Edge e charges the price c_e for transporting a message from one end to other. The cost of broadcasting a message is the total price charged by all the edges. We assume that each edge is of infinite capacity, so, a message can be sent from node i to node j through the shortest path between them. Hence we assume that c_e's satisfy the triangle inequality.

Now, suppose that the root has a message to broadcast and every user i has reported utility u_i to the root. The root's job is to come up with the following three things:

1. a set Q of users, selected to receive the message,
2. a tree T containing Q to broadcast the message, and
3. for each user i, the price x_i, to be charged as the cost of delivering the message.

For notational convenience we will also represent Q by its indicator function q, i.e., $q_i = 1$ if $i \in Q$ and $q_i = 0$ otherwise.

Now there are computational and economic constraints to be satisfied. The only computational constraint we are considering in this paper is that only polynomial time is available for computation. The economic constraints are listed below:

1. **Optimum**. T is an optimum tree connecting all the users in Q with the root.
2. **No Positive Transfers** (NPT). For each user i, $x_i \geq 0$, i.e., users will not be payed for receiving a message.

3. **Voluntary Participation** (VP). $q_i u_i - x_i \geq 0$ i.e., if $i \notin Q$ then $x_i = 0$ and if $i \in Q$ then $x_i \leq u_i$ i.e., only those users will pay who will receive the message. Moreover, they will never be asked to pay more than their reported utilities. In other words each user has the option to not receive the message, and if so, derives a benefit of 0.

4. **Consumer Sovereignty** (CS). Every user is guaranteed to receive the message if she reports a high enough utility value.

5. **Budget-balance** (BB). $\sum_{i \in Q} x_i = \sum_{e \in T} c_e$, i.e., total cost of T is payed by all the users receiving the movie.

6. **Efficiency**. $\sum_{i \in Q} u_i - \sum_{e \in T} c_e$ is maximized, i.e., as much worth is created as possible. Ideally we want Pareto optimal, i.e., $\sum_{i \in Q} u'_i - \sum_{e \in T} c_e$ to be maximized. But the next condition reduces the latter to the former.

7. **Group Strategyproof.** Even if a set of users collude, the dominant strategy for each player is to report her utility value truthfully. To be precise, consider a coalition C of users. Let $u_j = u'_j$ for all $j \notin C$. Let (q, x) and (q', x') denote the users served and costs at u and u' respectively. Now, group strategyproofness requires that if the inequality

$$u'_i q_i - x_i \geq u'_i q'_i - x'_i$$

holds for all $i \in C$ then it must hold with equality for all $i \in C$ as well, i.e., if no member of C is made worse off by misreporting their utility values then no member of C is made better off either.

Satisfying the first condition is not computationally possible, assuming $\mathbf{P} \neq \mathbf{NP}$. So we relax it: the cost of T is within twice the cost of the optimum tree connecting users in Q with the root. A classical result in game theory shows that there is no strategyproof mechanisms that is both budget balance and efficient. So condition 6 needs to be relaxed. Feigenbaum, Papadimitriou and Shenker proved that it is NP-hard to compute a constant factor approximation to condition 6. So, we are going to put this condition aside. Moulin and Shenker showed that all other economic constraints mentioned above can be captured by a cross monotonic cost sharing method.

A *cost sharing method* is a function, ξ, which distributes the cost of broadcasting a message to the recipients. More formally, ξ takes two arguments, a set of users Q and a user i, and returns a non-negative real number satisfying the following:

1. if $i \notin Q$ then $\xi(Q, i) = 0$ and

2. $\sum_{i \in Q} \xi(Q, i) = C(Q)$, where $C(Q)$ represents the cost of broadcasting a message to Q. Note that $C(Q)$ is not the optimum cost but it is the cost of the tree T computed for broadcasting the message.

A cost sharing method is *cross-monotonic* if for $Q \subseteq R$, $\xi(Q, i) \geq \xi(R, i)$ for every user $i \in Q$. For every cross monotonic cross sharing function, ξ, Moulin and Shenker give a mechanism $M(\xi)$ which computes Q and $x_i = \xi(Q, i)$.

Mechanism $M(\xi)$

1. Q is initialized to U.

2. if there is a user i in Q with $u_i < \xi(Q, i)$ then drop i from Q. Keep repeating this step, in arbitrary order, until for every user i in Q, $u_i \geq \xi(Q, i)$.

3. set $x_i = \xi(Q, i)$.

This mechanism starts with an attempt to send message to all users, and uses ξ to determine the cost share of each user. It drops, in arbitrary order, any user who cannot pay for receiving the message. The mechanism iterates until it a set of users who can together pay for the cost of the tree needed to serve them. Because of cross-monotonicity of ξ, the eventual set of users left does not depend on the order in which individual users are dropped. Since in every step at least one user is dropped, the mechanism will stop in a linear number of iteration. Such a mechanism clearly achieves BB, NPT, VP and CS. The following theorem shows that it is also group strategyproof.

Theorem 1 *[12, 10] For any cross monotonic cost sharing method ξ, the mechanism $M(\xi)$ is budget balanced, meets NPT, VP, CS and is group strategyproof.*

We provide a simpler proof of group strategyproofness in the Appendix; all other conditions are easy to verify. As shown in [12, 10], the converse of Theorem 1 holds for any submodular function.

Theorem 2 *[12, 10] Suppose U is the set of all users and $C : 2^U \rightarrow \mathcal{R}^+$ is a submodular cost function. If a cost sharing mechanism is budget balanced, meets NPT, VP, CS and is group strategyproof then the mechanism is equivalent to the mechanism $M(\xi)$ for some cross monotonic cost sharing method, ξ.*

3 The No Subsidy condition

Economic constraints ECO1-ECO7 discussed in the last section forbid users from misreporting their utility values. These constraints do not prohibit the service provider from overcharging a set of users [10]. In a competitive market, overcharging a subset of users may mean losing that set of users. Even in the case of a monopoly (like the government), it is not desirable to overcharge a set of users in order to subsidize another set of users. So we add the following economic constraint to the list ECO1-ECO7.

ECO8 No Subsidy If Q is the set of users selected to receive the service and x_i is the cost of service asked from user i, then for any subset R of Q, $\sum_{i \in R} x_i \leq C(R)$, where $C(R)$ is the cost of serving R.

A cost sharing mechanism which is budget balanced, meets NPT, VP, CS and is group strategyproof need not meet the No Subsidy condition. Consider the case when we only have two users a and b. Suppose the cost function is $C() = 0$, $C(a) = 1$, $C(b) = 1$, and $C(a, b) = 3$. Note that this is a supermodular function, so it admits a group strategy proof mechanism [10]. In particular, consider the following mechanism in which u_a and u_b denote the reported utility values for a and b respectively:

- If $u_a > C(a)$ then a receives the service, else not.

- b receives the service if $u_a > C(a)$ and $u_b \geq C(a,b) - C(a)$, or if $u_a \leq C(a)$ and $u_b \geq C(b)$. In all other cases, b does not receive the service.

One can easily check by case analysis that the above mechanism is budget balanced, meets NPT, VP, CS and is group strategyproof but does not meet the No Subsidy condition. In fact there is no mechanism which meets the No Subsidy condition because if both a and b are selected to receive the service (which they can be, by CS) then at least one of a and b is overcharged.

We characterize completely the class of cost functions that admit a No Subsidy mechanism. This is the class of functions that exhibit the *covering property*, defined below. This class includes all submodular functions. Many authors formally capture the notion of *economies of scale* by submodular functions. It is interesting to note that our larger class can also serve this purpose.

A *fractional set*, S, is a set in which an element can appear partially i.e., with each element, e, there is a number $f_S(e) \in [0,1]$, which tells the extent of appearance of e in set S. The union of two fractional sets, S_1 and S_2, is denoted by $S_1 \cup S_2$ and is defined by the function $f_{S_1 \cup S_2} = \min\{f_{S_1} + f_{S_2}, 1\}$. If S is a set and $f \in [0,1]$ then $f \cdot S$ is a fractional set where each element in S appears to the extent of f in $f \cdot S$. Fractional sets $S_1, S_2, \ldots, S_n$ *cover* S if $S_1 \cup S_2 \cup \ldots \cup S_n = S$.

Suppose U is the set of all users and $C : 2^U \to \mathcal{R}^+$ is a cost function. C is said to exhibit the *covering property* if for any set S of users and any covering of S of the form $S = \bigcup_j f_j \cdot S_j$, we have $C(S) \leq \sum_j f_j \cdot C(S_j)$, where each S_j is a set of users.

Theorem 3 *Suppose U is the set of all users and $C : 2^U \to \mathcal{R}^+$ is a cost function. There is a cost sharing mechanism which satisfies budget balance, meets NPT, VP, CS, and also satisfies the No Subsidy condition if and only if C exhibits the covering property.*

Definition 4 A cost sharing method ξ is *weakly cross-monotonic* if for $Q \subseteq R$, $\sum_{i \in Q} \xi(R,i) \leq \sum_{i \in Q} \xi(Q,i) = C(Q)$.

Lemma 5 *Suppose U is the set of all users and $C : 2^U \to \mathcal{R}^+$ is a cost function. There is a budget balanced, NPT, VP, CS and No Subsidy cost sharing mechanism if and only if there is a weakly cross-monotonic cost sharing method.*

Corollary 6 *Suppose U is the set of all users and $C : 2^U \to \mathcal{R}^+$ is a submodular cost function. If a cost sharing mechanism satisfies budget balance, meets NPT, VP, CS and is group strategyproof, then it satisfies the No Subsidy condition also.*

Lemma 7 *Suppose U is the set of all users and $C : 2^U \to \mathcal{R}^+$ is a cost function. There exists a weakly cross-monotonic cost sharing method for C if and only if C exhibits the covering property.*

Proof : **Forward direction**: Let S be a set and $\bigcup_j f_j \cdot S_j$ be a covering for it. Because of VP and CS, S could be selected as the set to receive the service. Suppose each user $i \in S$ is charged x_i as the cost of receiving the service. Then,

$$C(S) = \sum_{i \in S} \xi(S,i) \leq \sum_j (f_j \cdot \sum_{i \in S_j} \xi(S,i)) \leq \sum_j (f_j \cdot C(S_j)).$$

Reverse direction: C allows a weakly cross-monotonic cost sharing method if for every set R of users the following linear program (LP) has a feasible solution.

$$\sum_{i \in R} \xi(R,i) = C(R)$$
$$\forall Q \subseteq R: \sum_{i \in Q} \xi(R,i) \leq C(Q)$$
$$\forall i \in R: \xi(R,i) \geq 0$$

The above LP has a feasible solution if and only if the optimum value of the following LP is at least $C(R)$,

$$\begin{aligned} \text{maximize} \quad & \sum_{i \in R} \xi(R,i) \\ \text{subject to} \quad & \forall Q \subseteq R: \sum i \in Q \xi(R,i) \leq C(Q) \\ & \forall i \in R: \xi(R,i) \geq 0 \end{aligned} \tag{1}$$

Suppose C does not allow a weakly cross-monotonic cost sharing method. Then there will be a set R of users such that the above LP has strictly smaller optimum value than $C(R)$. Hence by the strong duality theorem of linear programming there will be a dual solution of value strictly smaller than $C(R)$. So let us consider the dual of the above LP.

$$\begin{aligned} \text{minimize} \quad & \sum_{Q \subseteq R} C(Q) Y_Q \\ \text{subject to} \quad & \forall i \in R: \sum Q: i \in Q Y_Q \geq 1 \\ & \forall Q \subseteq R: Y_Q \geq 0 \end{aligned} \tag{2}$$

Consider the feasible dual solution Y with value strictly smaller than $C(R)$. We have $\sum_{Q \subseteq R} C(Q) Y_Q < C(R)$. Because of the dual constraints, $\bigcup_{Q \subseteq R} Y_Q \cdot Q = R$, contradicting the covering property. $\square$

Corollary 8 *Submodular functions form a subclass of functions which exhibit the covering property.*

4 The cost-sharing method

In this section, we will give a cost-sharing method for multicasting that achieves budget balance and is cross-monotone, and for any set Q of users chosen for distribution, finds a tree of cost at most twice the optimal Steiner tree containing the root and users Q. Our method utilizes the following two well known facts.

- If the edge costs satisfy the triangle inequality, the cost of a minimum spanning tree on the set of required vertices is within twice the cost of an optimal Steiner tree containing all required vertices.

- There is an exact linear programming relaxation for the minimum spanning tree problem, i.e., a relaxation that always has optimal integral solutions.

The first fact is due to [9]. The second follows from a more general fact due to Edmonds: that there is an exact relaxation for the minimum branching problem. In this problem, we are given a directed graph with non-negative costs on the directed edges, and one of the vertices is marked as root. The problem is to find a minimum cost tree containing all vertices and directed into the root. The transformation from the minimum spanning tree problem in an undirected graph to the minimum branching problem is straightforward. Simply replace each undirected edge $e = (u, v)$ of G by two directed edges $(u \to v)$ and $(v \to u)$ each of cost c_e, and ask for a minimum cost branching directed into the root. Let us denote this directed graph by $H = (V, \vec{E})$. Once the branching is found, we ignore directions on edges to obtain a minimum spanning tree in G.

Let us say that a set $S \subset V$ is *valid* if it is non-empty and does not contain the root. For any set $S \subset V$ and $F \subset \vec{E}$ let $\delta(S) = \{(u \to v) \in \vec{E} \mid u \in S \text{ and } v \in \overline{S}\}$, and $\delta_F(S) = \{(u \to v) \in F \mid u \in S \text{ and } v \in \overline{S}\}$. We state below the LP-relaxation and dual.

$$\begin{array}{lll} \text{minimize} & \displaystyle\sum_{e \in \vec{E}} c_e x_e & \\ \text{subject to} & \displaystyle\sum_{e:\ e \in \delta(S)} x_e \geq 1, & \forall \text{ valid set } S \\ & x_e \geq 0, & e \in \vec{E} \end{array} \tag{3}$$

$$\begin{array}{lll} \text{maximize} & \displaystyle\sum_{\text{valid set } S} y_S & \\ \text{subject to} & \displaystyle\sum_{S:\ e \in \delta(S)} y_S \leq c_e & e \in \vec{E} \\ & y_S \geq 0 & \forall \text{ valid set } S \end{array} \tag{4}$$

Let us say that edge e *feels* dual y_S if $y_S > 0$ and $e \in \delta(S)$. Say that edge e is *tight* if the total amount of dual it feels equals its cost. The dual program is trying to maximize the sum of the dual variables y_S subject to the condition that no edge feels more dual than its cost, i.e., no edge is *over-tight*.

We present Edmonds' algorithm below, which is based on the primal-dual schema. Starting with the trivial primal and dual solutions, it iteratively improves the feasibility of the primal and the optimality of the dual. When an edge becomes tight, it is included in the ordered list F. In any iteration, a set $S \subset V$ is said to be *unsatisfied* if it is valid and $\delta_F(S) = \emptyset$. Any minimal unsatisfied set is said to be *active*. It is easy to see that active sets must be disjoint and must be strongly connected w.r.t. F.

For the purpose of proving properties of this algorithm, we will associate a notion of *time* with the algorithm. In unit time, the algorithm grows duals a unit amount. In an iteration, the algorithm finds all active sets, and raises their dual variables until a new edge, say e, goes tight. At this point, the current iteration ends, and edge e is appended to the list F. The algorithm continues until there are no more unsatisfied sets. At this point, the algorithm prunes F using the procedure of reverse delete. The edges that remain in F form a branching directed into the root.

Algorithm 9 (Minimum branching)

1. **(Initialization)** $F \leftarrow \emptyset$; for each $S \subseteq V$, $y_S \leftarrow 0$.
2. **(Edge augmentation)** While there exists an unsatisfied set do:

   Find all active sets w.r.t. F. For each such set S, raise its dual variable y_S until some edge e goes tight;

   $F \leftarrow F \cup \{e\}$.
3. Let $e_1, e_2, \ldots, e_l$ be the ordered list of edges in F.
4. **(Reverse delete)** For $j = l$ downto 1 do:
   If there are no unsatisfied sets w.r.t. $F - \{e_j\}$, then $F \leftarrow F - \{e_j\}$.
5. Return F.

Because the algorithm raises dual variables only for strongly connected sets, and does a reverse delete in the end, it is possible to show that every valid dual S such that $y_S > 0$, must satisfy $|\delta_F(S)| = 1$. This, and the fact that only tight edges are picked, lead to showing that the cost of the branching picked is precisely equal to the total dual raised, i.e.,

$$\sum_{e \in F} c_e = \sum_{\text{valid set } S} y_S.$$

This important fact shows that the branching found by the algorithm is optimal. As a consequence, we get that LP (3) always has an integral optimal solution. This fact will also enable us to show that our cost-sharing method is budget balanced.

Given a set Q of recipients, we use Algorithm 9 to find a minimum spanning tree containing Q and the root. Our cost-sharing method, ξ, simply distributes each raised dual equally among all vertices contained in this dual. So, the amount charged to user i is

$$\xi(Q, i) = \sum_{S : i \in S} y_S.$$

By the above stated fact, our cost-sharing method is budget balanced.

Theorem 10 *The cost-sharing method, ξ, is cross-monotonic.*

Proof : Consider a run of Algorithm 9 on the set of users Q. Let v be an arbitrary user. At any time t during the run of the algorithm, let S_t be the set of vertices reachable from v by means of tight edges. If S_t does not contain the root, then dual y_{S_t} is growing at time t. Otherwise, v has already reached the root, and there is no dual containing v that is growing at time t. Let t_v denote the time at which v reaches the root. Then,

$$\xi(Q, v) = \int_0^{t_v} \frac{1}{|S_t|} dt.$$

Now, let Q be a set of users, and $w \notin Q$ be another user. Let R and R' be the runs of Algorithm 9 on input $Q \cup \{w\}$ and Q, respectively. Let S_t and S'_t be the reachability of v at time t in runs

R and R'. If v can reach w in R, then $S_t \supset S'_t$, and otherwise $S_t = S'_t$. Therefore, $|S_t| \geq |S'_t|$, and furthermore, v reaches the root at the same time or earlier in run R than in run R'. Hence, $\xi(Q \cup \{w\}, v) \leq \xi(Q, v)$. The lemma follows. $\square$

Corollary 11 *For the cost sharing method ξ given above, the mechanism $M(\xi)$ is budget balanced, meets NPT, VP, CS, is group strategyproof, and satisfies No Subsidy.*

5 Acknowledgements

We wish to thank Hervé Moulin and Scott Shenker for providing valuable pointers into the literature.

References

[1] M. L. Balinski. On finding integer solutions to linear programs. *Proc. IBM Scientific Computing Symp. on Combinatorial Problems*, 225-248. IBM, 1966.

[2] S. Deering and D. Cheriton. Multicast routing in datagram internetworks and extended LANs. In *ACM Transactions on Computer Systems*, 8:85-110, 1990.

[3] B. Dutta and D. Ray. A concept of egalitarianism under participation constraints, *econometrica*, 57, 615-635.

[4] J. Edmonds. Optimum branchings. *J. Res. Nat. Bur. Standards*, 71:233-240, 1967.

[5] J. Feignbaum, C. Papadimitriou and S. Shenker. Sharing the cost of multicast transmission. To appear in *STOC*, 2000.

[6] J. Green, E. Kohlberg, and J. J. Laffont. Partial equilibirium approach to the free rider problem. In *Journal of Public Economics*, 6:375-394, 1976.

[7] M. X. Goemans and M. Skutella. Cooperative Facility Location Games. *Proc. of Soda 2000*, 76-85.

[8] R. M. Karp. Reducibility among combinatorial problems, in R. E. Miller and J. W. Thatcher (eds.), *Complexity of Computer Computation*, Plenum Press, New York, 85-103.

[9] L. D. Kou, G. Markowsky, and L. Berman. A fast algorithm for Steiner trees, *Acta Informatica*, 1981, 15:141-145.

[10] H. Moulin. Incremental cost sharing: characterization by coalition strategy-proofness, *Social Choice and Welfare*, 1999, 16, 279-320.

[11] S. Mutuswami. Strategyproof mechanisms for cost sharing, working paper, Indian Statistical Institute.

[12] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. preprint 1997. Available at *http://www.aciri.org/shenker/cost.ps*

[13] C.H. Papadimitriou and K. Steiglitz. *Combinatorial optimization: algorithms and complexity*, Prentice-Hall, Englewood Cliffs, NJ, 1982.

[14] K. Roberts. The characterization of implementable choice rules. In J. J. Laffont (ed.) *Aggregation and Revelation of Preferences*, North Holland, 1979, pp. 321-348.

[15] L. S. Shapley. A value of n-person games. In H. Kuhn and W. Tucker (eds.) *Contributions to the Theory of Games*, Princeton University Press, 1953, pp. 31-40.

A Appendix

Theorem 12 *[12, 10] For any cross monotonic cost sharing method ξ, the mechanism $M(\xi)$ is budget balanced, meets NPT, VP, CS and is group strategyproof.*

Proof : We will show that $M(\xi)$ is group strategyproof, for a cross monotonic cost sharing method ξ. Let u' be the true utility profile. Assume a coalition C manipulates at u' by u (where $u'_i = u_i$ for all $i \notin C$). Consider two runs, $R(u')$ and $R(u)$ of the above mechanism with true and false utility values respectively. Let Q' and Q be the output sets produced by the two runs, and q' and q be the vectors representing these sets. Assume further that for each $i \in C$, the inequality

$$u'_i q_i - x_i \geq u'_i q'_i - x'_i \tag{5}$$

holds. We will show that for each $i \in C$, this inequality must hold with equality, thereby showing that $M(\xi)$ is group strategyproof.

In an iteration of $M(\xi)$, there could be several users, i, satisfying $u_i < \xi(Q, i)$. Since ξ is cross monotonic, the final outcome does not depend on the manner in which a user is picked to be dropped. Each of the runs $R(u')$ and $R(u)$ makes these choices in an arbitrary (and uncorrelated) manner. Let $s_1, s_2, \ldots, s_k$ be the order in which users were dropped in run $R(u')$.

We will first prove, by contradiction, that $Q \subseteq Q'$. Let s_i be the first user in the list $s_1, s_2, \ldots, s_k$ such that $s_i \in Q$. By the choice of i, each of the users $s_1, \ldots, s_{i-1}$ has been dropped in run $R(u)$, and so $U - \{s_1, \ldots, s_{i-1}\} \supseteq Q$. Now, by cross monotonicity of ξ, $\xi(U - \{s_1, \ldots, s_{i-1}\}, s_i) \leq \xi(Q, s_i)$. Since s_i is dropped in run $R(u')$ and retained in run $R(u)$, $u'_{s_i} < \xi(U - \{s_1, \ldots, s_{i-1}\}, s_i)$, and $\xi(Q, s_i) \leq u_{s_i}$.

Now, there are two cases. If $s_i \notin C$, by assumption, $u'_{s_i} = u_{s_i}$. But the above inequalities give $u'_{s_i} < u_{s_i}$, thereby leading to a contradiction. The second case is that $s_i \in C$. Since $q'_{s_i} = 0$ and $q_{s_i} = 1$, by (5), $u'_{s_i} \geq \xi(Q, s_i)$. But the above inequalities give $u'_{s_i} < \xi(Q, s_i)$, again leading to a contradiction. Hence, $Q \subseteq Q'$.

Now, by cross monotonicity of ξ, the price paid by any member of C in the second run must be at least the price paid in the first run, i.e., $\forall i \in C, \xi(Q, i) \geq \xi(Q', i)$. Hence, for each $i \in C$, (5) must hold with equality. $\square$

Lemma 13 *Suppose U is the set of all users and $C : 2^U \to \mathcal{R}^+$ is a cost function. There is a budget balanced, NPT, VP, CS and No Subsidy cost sharing mechanism if and only if there is a weakly cross-monotonic cost sharing method.*

Proof : **Forward direction**: By VP and CS, any set Q can be selected as a set of service receivers. Suppose x_i is the amount that user i needs to pay if Q is selected as the set of service receivers. Now, assign $\xi(Q, i) = x_i$. By the No Subsidy condition, the cost sharing method defined by ξ is weakly cross-monotonic.

Reverse direction: Consider the mechanism $M(\xi)$ with the modification that dropping of unqualified users happens in some fixed order instead of arbitrary order. One can see that it satisfies budget balance, meets NPT, VP, CS and also satisfies No Subsidy. $\square$