# Applications of Approximation Algorithms to Cooperative Games 

Kamal Jain*

Vijay V. Vazirani ${ }^{\dagger}$

## 1. INTRODUCTION

The Internet, which is intrinsically a common playground for a large number of players with varying degrees of collaborative and selfish motives, naturally gives rise to numerous new game theoretic issues. Computational problems underlying solutions to these issues, achieving desirable economic criteria, often turn out to be NP-hard. It is therefore natural to apply notions from the area of approximation algorithms to these problems. The connection is made more meaningful by the fact that the two areas of game theory and approximation algorithms share common methodology - both heavily use machinery from the theory of linear programming. Various aspects of this connection have been explored recently by researchers $[8,10,15,20,21,26,27$, 29].

In this paper we will consider the problem of sharing the cost of a jointly utilized facility in a "fair" manner. Consider a service providing company whose set of possible customers, 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, $u_{i}^{\prime}$, for receiving the news. This utility $u_{i}^{\prime}$ is known only to user $i$. User $i$ enjoys a benefit of $u_{i}^{\prime}-x_{i}$ 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

[^0][^1]true value of her utility. It is said to be group strategyproof if this holds for coalitions as well. A cost sharing mechanism is budget balancedif 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 [13, 28] shows that such a mechanism does not exist even for a submodular cost function (see Section 2 for formal definitions of these notions). Such a mechanism does not exist even after relaxing the condition of group strategyproofness to just strategyproofness.

In view of this limitation, there are two options: sacrifice either budget balance or efficiency. In the first case, one can show that if the cost function is nondecreasing and submodular, then there is only one way of maximizing efficiency [24]. This is called the marginal cost mechanism. This mechanism is strategyproof, though not group strategyproof. It never creates a budget surplus but can run a deficit, and in many cases raises no revenue at all [24].

In the second case, a fundamental theorem of Moulin and Shenker [23, 24] shows that a cross-monotonic cost sharing method (also known as population monotonic allocation scheme, see Section 2 for a definition) gives rise to a budget balance and group strategyproofness cost sharing mechanism. Moreover, if the cost function is submodular, the converse holds as well.

A nondecreasing, submodular cost function supports an entire class of cross-monotonic cost sharing methods. It is useful to characterize methods from this class possessing special properties so that the service provider may pick one judiciously. Two well known methods are:

1. Shapley Value [31]: In the context of multicasting (see Section 2), 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 [24].
2. Egalitarian [5]: This method, due to Dutta and Ray, seeks to distribute the cost equally 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 the set served [25].

A cost sharing method is said to satisfy the coalition participation constraint if the total cost share of any subset $S$
of the users is no more than $C(S)$, i.e., coalition $S$ has no incentive to break off from the grand coalition. only the users in $S$. Such a cost sharing method is also called weakly cross-monotone. The class of all weakly cross-monotone cost sharing methods corresponding to a given cost function is said to form the core. The classic Bondareva-Shapley Theorem states that the core is nonempty iff the cost function satisfies the covering property (see Section 5 for definition). This allows us to compute an element of $\alpha$-approximate core (see sections 4 and 5).

In this paper we first consider the problem of multicast routing. A tree $T$ containing the source and all possible users is fixed. To serve a subset $S$ of the users, information is routed over the subtree of $T$ containing $S$ and the source. Feigenbaum, Papadimitriou and Shenker [8] 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, thus leading to a linear number of total message, whereas Shapley value requires a quadratic number of messages. They also show that the welfare value of an optimal subtree is NP-hard to approximate within any constant factor.

Although this model for multicast routing is widely used, it suffers from the drawback that the subtree connecting a certain subset $S$ of receivers may be arbitrarily more costly than the cheapest tree to them. The latter of course is an optimal Steiner tree containing $S$ and the source. Such a tree is NP-hard to compute. In addition, Megiddo [22] has shown that for this game, the core is empty and so there does not even exist a weak cross-monotonic cost sharing method for the optimal Steiner tree.

We get around these difficulties by turning to approximation algorithms. A well known factor 2 approximation algorithm for Steiner tree is to find a minimum spanning tree on the required vertices [19]. The minimum spanning tree game has been studied extensively in the literature [2, 11, 12, 17, 18, 16]. Kent and Skorin-Kapov [17] show how to construct a whole class of weak cross-monotonic methods as well as one cross-monotone method for this game. For the latter result, they use matroid properties of spanning trees. Using the primal-dual minimum spanning tree algorithm of Edmonds [6], we show how to construct a class of cross-monotone methods. These methods are parameterized by $n$ mappings, $f_{i}: \mathcal{R}^{+} \rightarrow \mathcal{R}^{+}$, one for each user, $i$. The service provider now has a choice of methods. For instance, he could use his estimates of the probability distribution functions of users' utilities as the mappings.

In a forthcoming paper [14] we build on the above idea of running a primal-dual-type algorithm and distributing costs according to $n$ mappings for constructing a class of cross-monotonic cost sharing methods for an arbitrary nondecreasing submodular function. We also present a broader notion of egalitarianism motivated by the following: If individual users have widely varying financial resources then the egalitarian method of Dutta and Ray does not equalize their opportunity of receiving the service. We introduce the notion of opportunity egalitarianism which attempts to do this. We show how a service provider can attempt to equalize, among the users, the opportunity of their getting the service, based on his estimate of their individual resources. The basic technique underlying our method also enables us to generalize Mutuswami's result: we give an algorithm that maximizes the expected size of the set served even if indi-
vidual utilities are drawn from different probability distributions.

In Section 4 we point out a shortcoming of the condition of budget balance, and propose the $\alpha$-approximate budget balance condition. We use the cost sharing methods derived for Steiner tree to give a 2-approximate budget balanced group strategyproof cost sharing mechanism for the metric TSP game (see Section 4), and leave the open problem of deriving a factor 1.5 solution, using Christofedes' algorithm. Core allocation functions for the metric and Euclidean TSP games have been given by Faigle, Fekete, Hochstattler, and Kern [7] and Fekete and Pulleyblank [9], respectively.

Several NP-hard minimization problems admit a constant factor approximation algorithm based on an
LP-relaxation that is a covering program (e.g., see [32]). In Section 5 we study general approaches for analyzing games based on such problems. We show that the cost function given by an optimal solution to a covering linear program satisfies the covering property (see Section 5 for definitions). We further show that it admits an efficiently computable $\alpha$ approximate weakly cross-monotonic cost sharing method. Perhaps a stronger result holds - that in fact it admits a cross-monotonic cost sharing method. We leave this as an (important) open problem. If this open problem is settled positively, it will lead to a general scheme for obtaining a polynomial time computable $\alpha$-approximate budget balanced group strategyproof cost sharing mechanism for such games. Finding such a method for the facility location game seems particularly interesting. Recently, Goemans and Skutella [10] have studied this game and have a core allocation in case the well known LP-relaxation, due to Balinski [1], has an integral solution.

Finally, we consider the purely combinatorial question of characterizing the space of cross-monotone methods corresponding to a given nondecreasing submodular function. We show that this is a polyhedron and characterize a subset of its corner points. We leave the open problem of characterizing the rest of the corner points.

## 2. THE MODEL

Let us describe the model in the context 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 [4]. 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.

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. Although constant factor approximation algorithms are known for the Steiner tree problem, we do not know of any algorithm which gives either a nondecreasing 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.

Let $G=(V, E)$ be an undirected graph with edge weights $c_{e}$ 's and a marked node root, which is the broadcasting source. All other nodes are users. We denote the set of all users by $U$. 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:

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. Optimality. $T$ 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 paid 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).
(a) Cost Recovery $\sum_{i \in Q} x_{i} \geq \sum_{e \in T} c_{e}$, i.e., the cost of broadcasting the message is recovered from all the users.
(b) Competitiveness $\sum_{i \in Q} x_{i} \nsupseteq \sum_{e \in T} c_{e}$, i.e., no surplus is created. Because if any surplus is created then a competitor can deliver the message at a cheaper cost by reducing the surplus.

The condition of budget balance consists of satisfying both, cost recovery and competitiveness, i.e., $\sum_{i \in Q} x_{i}=$ $\sum_{e \in T} c_{e}$ (the set of users receiving the message pay exactly the total cost of $T$ ).
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., $\sum_{i \in Q} u_{i}^{\prime}-\sum_{e \in T} c_{e}$ to be maximized. But the next condition reduces the latter to the former.
7. Group Strategyproofness. 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 $u_{j}=u_{j}^{\prime}$ for all $j \notin C$. Let $(q, x)$ and $\left(q^{\prime}, x^{\prime}\right)$ denote the users served and costs at $u$ and $u^{\prime}$ respectively. Now, group strategyproofness requires that if the inequality

$$
u_{i}^{\prime} q_{i}-x_{i} \geq u_{i}^{\prime} q_{i}^{\prime}-x_{i}^{\prime}
$$

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 $\mathbf{P} \neq \mathbf{N P}$. We relax it to: 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. Furthermore, Feigenbaum, Papadimitriou and Shenker have shown that it is NP-hard to compute a constant factor approximation to condition 6 as well. 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 function is submodular if

1. $C(\emptyset)=0$,
2. for any $Q_{1}$ and $Q_{2}, C\left(Q_{1}\right)+C\left(Q_{2}\right) \geq C\left(Q_{1} \cup Q_{2}\right)+$ $C\left(Q_{1} \cap Q_{2}\right)$.

It is supermodular if the inequality in the second condition is reversed. A cost function is nondecreasing if for $Q_{1} \subseteq Q_{2}$, $C\left(Q_{1}\right) \leq C\left(Q_{2}\right)$.

A cost allocation function $f$ distributes the cost of sending the message to the entire set of users $U$, i.e., $\forall i \in U, f(i) \geq 0$ and $\sum_{i \in U} f(i)=C(U)$. The core consists of all cost allocation functions $f$ such that $\forall S \subseteq U, \sum_{i \in S} f(i) \leq C(S)$, i.e., no subset of the users have an incentive to secede.

A cost sharing method is a function, $\xi$, which distributes the cost of broadcasting the message to the recipients. More formally, $\xi$ takes two arguments, a set of users $Q$ and a user $i$, and returns a nonnegative 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 $i \in Q$. It is weakly crossmonotonic if for $Q \subseteq R, \sum_{i \in Q} \xi(Q, i) \geq \sum_{i \in Q} \xi(R, i)$. As the name suggests, a cross-monotonic method is weakly crossmonotonic as well. Observe that a weakly cross-monotonic method provides us with a core allocation for each subset of the users.

For every cross-monotonic cross sharing method, $\xi$, Moulin and Shenker give the following 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 $\xi$ 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 $\xi$, 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.

Theorem 1. [24, 23] For any cross monotonic cost sharing method $\xi$, the mechanism $M(\xi)$ is budget balanced, meets $N P T, V P, C S$ and is group strategyproof.

As shown in [23, 24], the converse of Theorem 1 holds for any submodular function. For completeness, we provide a (slightly simpler) proof of group strategyproofness. All other conditions are easy to verify.

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

$$
\begin{equation*}
u_{i}^{\prime} q_{i}-x_{i} \geq u_{i}^{\prime} q_{i}^{\prime}-x_{i}^{\prime} \tag{1}
\end{equation*}
$$

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 $\xi$ 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 $R\left(u^{\prime}\right)$ 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 $R\left(u^{\prime}\right)$.

We will first prove, by contradiction, that $Q \subseteq Q^{\prime}$. 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 $U-\left\{s_{1}, \ldots, s_{i-1}\right\} \supseteq Q$. Now, by cross monotonicity of $\xi, \xi\left(U-\left\{s_{1}, \ldots, s_{i-1}\right\}, s_{i}\right) \leq$ $\xi\left(Q, s_{i}\right)$. Since $s_{i}$ is dropped in rum $R\left(u^{\prime}\right)$ and retained in rum $R(u), u_{s_{i}}^{\prime}<\xi\left(U-\left\{s_{1}, \ldots, s_{i-1}\right\}, s_{i}\right)$, and $\xi\left(Q, s_{i}\right) \leq u_{s_{i}}$.

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

Now, by cross monotonicity of $\xi$, the price paid by any member of $C$ in the second rum must be at least the price paid in the first run, i.e., $\forall i \in C, \xi(Q, i) \geq \xi\left(Q^{\prime}, i\right)$. Hence, for each $i \in C$, (1) must hold with equality.

## 3. THE STEINER TREE GAME

In this section, we will give a class of cost-sharing methods for multicasting that achieves budget balance and is crossmonotone, 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 [19]. 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 nonnegative 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 $(u \rightarrow v)$ and $(v \rightarrow u)$ 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 nonempty and does not contain the root. For any set $S \subset V$ and $F \subset \vec{E}$ let $\delta(S)=\{(u \rightarrow v) \in \vec{E} \mid u \in \dot{S}$ and $v \in \bar{S}\}$, and $\delta_{F}(S)=\{(u \rightarrow v) \in F \mid u \in S$ and $v \in \bar{S}\}$. We state below the LP-relaxation and dual.

$$
\begin{array}{lll}
\operatorname{minimize} & \sum_{e \in \vec{E}} c_{e} x_{e} &  \tag{2}\\
\text { subject to } & \sum_{e: \in \delta(S)} x_{e} \geq 1, \quad \forall \text { valid set } S \\
& x_{e} \geq 0, & e \in \vec{E}
\end{array}
$$

$$
\begin{array}{lll}
\operatorname{maximize} & \sum_{\text {valid set } S} y_{S} &  \tag{3}\\
\text { subject to } & \sum_{S:} y_{S} \leq c_{e} \quad e \in \vec{E} \\
& y_{S} \geq 0 & \forall \text { valid set } S
\end{array}
$$

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 prumes $F$ using the procedure of reverse delete. The edges that remain in $F$ form a branching directed into the root.

## Algorithm 2 (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 $\epsilon_{1}, e_{2}, \ldots, e_{l}$ 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. $F-\left\{e_{j}\right\}$, then $F \leftarrow F-\left\{e_{j}\right\}$.
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 $\left|\delta_{F}(S)\right|=1$. 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 (2) always has an integral optimal solution. This fact will also enable us to show that our cost-sharing method is budget balanced.

Let $Q$ be the set of recipients. We are provided with functions $f_{i}: \mathcal{R}^{+} \rightarrow \mathcal{R}^{+}$, one for each user $i$. Use Algorithm 2 to find a minimum spanning tree containing $Q$ and the root.

We now define the cost-sharing method $\xi$. The cost share for user $i \in Q$ is computed as follows. Let $T$ denote the first time at which there is a path from $i$ to the root consisting of tight edges. At time $t \leq T$, let $S(t)$ denote the set of vertices reachable from $i$ using tight edges. At each time $t<T$, the algorithm grows the dual variable $y_{S(t)}$ at unit rate. Let

$$
F(t)=\sum_{j \in S(t)} f_{j}(t)
$$

Define the cost share for user $i$,

$$
\xi(Q, i)=\int_{0}^{T} \frac{f_{i}(t)}{F(t)} d t
$$

Theorem 3. The cost-sharing method, $\xi$, is cross-monotonic.

Proof. Observe that at each time, the cost-sharing method is simply distributing the growing dual $y_{S(t)}$ among the users in $S(t)$. Since the total dual constructed equals the cost of the tree found,

$$
\sum_{i \in Q} \xi(Q, i)=C(Q)
$$

Consider a run of Algorithm 2 on the set of users $Q$ and let $w \notin Q$ be another user. Let $i \in Q$ be an arbitrary user. Let $R$ and $R^{\prime}$ be the runs of Algorithm 2 on input $Q$ and $Q \cup\{w\}$, respectively. Let $S(t)$ and $S^{\prime}(t)$ be the sets of vertices reachable from $i$ via tight edges at time $t$ in runs $R$ and $R^{\prime}$, respectively. Define

$$
F(t)=\sum_{j \in S(t)} f_{j}(t) \text { and } F^{\prime}(t)=\sum_{j \in S^{\prime}(t)} f_{j}(t)
$$

Let $T$ and $T^{\prime}$ be the first times at which there is a tight path from $i$ to the root in run $R$ and $R^{\prime}$, respectively. Then,

$$
\xi(Q, i)=\int_{0}^{T} \frac{f_{i}(t)}{F(t)} d t
$$

and

$$
\xi(Q \cup\{w\}, i)=\int_{0}^{T^{\prime}} \frac{f_{i}(t)}{F^{\prime}(t)} d t
$$

If $i$ can reach $w$ in run $R^{\prime}$ at time $t$, then $S^{\prime}(t) \supset S(t)$, and otherwise $S(t)=S^{\prime}(t)$. Therefore, for all $t, F^{\prime}(t) \geq F(t)$. Furthermore, $i$ reaches the root at the same time or earlier in run $R^{\prime}$ than in run $R$, i.e., $T^{\prime} \leq T$. Hence, $\xi(Q \cup\{w\}, i) \leq$ $\xi(Q, i)$.

For purposes of efficiency, Algorithm 2 is first run to determine the duals grown, and for each dual, the time at which it started and stopped growing. Since the duals grown form a laminar family, their number is bounded by $2 n$. This information is sufficient for cost allocation. If the functions $f_{i}$ are simple, the cost shares can be computed in closed form; otherwise, one will have to use numerical methods.

Corollary 4. For the cost sharing method $\xi$ given above, the mechanism $M(\xi)$ is budget balanced, meets $N P T$, VP, CS, and is group strategyproof.

If all the functions $f_{i}$ are the same, we get the cost-sharing method of Kent and Skorin-Kapov. This corresponds to distributing each dual $y_{S}$ equally among the users in the set $S$.

Are all possible cross-monotone cost sharing methods for minimum spanning trees captured by the algorithm given above? We give an example to show that the answer to this question is "No". Consider a network on 3 vertices, root, $u$ and $v$. Let the distances be (root, $u)=1,(u, v)=1$ and (root, v) $=2$. Consider the cross-monotone cost-sharing $\operatorname{method} \xi(\{u\}, u)=1, \xi(\{v\}, v)=2, \xi(\{u, v\}, u)=0$ and $\xi(\{u, v\}, v)=2$. When run with vertices $u$ and $v$, our algorithm will assign a nonnegative cost to vertex $u$. Hence, it will never generate this cost-sharing method.

## 4. RELAXING THE BUDGET BALANCE CONDITION

Clearly, the cost sharing method should be such that the service provider does not rum into a deficit. In the presence of competition, it should not create a large surplus either. Budget balance ensures both these conditions, in a mathematically clean manner. However, this is a difficult condition to ensure. Moreover, even when satisfied, it can suffer from the following flaw. Let us consider our Steiner tree solution. We assumed that we were provided with a complete graph, $G$, with edge costs satisfying the triangle inequality, and found a tree $T$ in it. However, the original network, $H$, may not have links connecting all pairs of nodes. Graph $G$ is obtained by taking the closure of $H$, so that edges in $G$ correspond to shortest paths in $H$. Mapping tree $T$ back to $H$ involves replacing edges by paths. This gives rise to a spanning graph, say $H^{\prime}$, which in general contains cycles and multiple edges. Now, budget balance requires the broadcasting company to send a message on each link of $H^{\prime}$, which is clearly wasteful. Any spanning tree that is a subgraph of $H^{\prime}$, say $T^{\prime}$, suffices.

This flaw and the general difficulty of ensuring budget balance motivate the following definition. Let OPT $(Q)$ denote the optimal cost function for serving users $Q$ and let $\alpha \geq 1$ be a constant. (In general, $\alpha$ may be a function of $n-$ for this paper, let us assume it is a constant.) $\xi$ is an $\alpha$-approximate cost sharing method if it satisfies the following:

1. $\alpha$-approximate Competitiveness
$\forall Q \subseteq U: \sum_{i \in Q} \xi(Q, i) \leq \alpha \cdot \mathrm{OPT}(Q)$, i.e., any set $Q$ of users together pays at most $\alpha$ times the optimum cost of serving $Q$.
2. Cost Recovery $\forall Q \subseteq U: \sum_{i \in Q} \xi(Q, i) \geq C(Q)$, for some feasible solution of cost $C(Q)$, i.e., the cost incurred by the service provider is recovered from the users served.
3. if $i \notin Q$ then $\xi(Q, i)=0$.
$\xi$ is cross-monotonic if for $Q \subseteq R, \xi(Q, i) \geq \xi(R, i)$ for every user $i \in Q$. It is weakly cross-monotonic if for $Q \subseteq R$, $\sum_{i \in Q} \xi(Q, i) \geq \sum_{i \in Q} \xi(R, i)$. It is said to be efficiently computable if for any $Q$, both $C$ and $\xi$ are polynomial time computable.

A cost sharing mechanism is $\alpha$-approximate budget balanced if it is $\alpha$-approximate competitive and cost recovering. One can easily obtain the following along the lines of Moulin and Shenker's Theorem 1.

Theorem 5. For any $\alpha$-approximate cross-monotonic cost sharing method $\xi$, mechanism $M(\xi)$ is $\alpha$-approximate budget balanced, meets NPT, VP, CS and is group strategyproof. Furthermore $M(\xi)$ is efficiently computable if $\xi$ is.

Under these definitions, tree $T^{\prime}$ is a 2 -approximate budget balanced solution. Let us illustrate another use of this notion. Consider the metric TSP game, in which edge costs between nodes satisfy the triangle inequality, and a traveling salesman starts at node 1 and executes a tour, visiting a set of users that are chosen by the cost sharing mechanism. Recall that doubling an MST, finding an Eulerian tour and short cutting gives a factor 2 approximation algorithm for metric TSP. Using this fact and Theorem 3 we get.

Theorem 6. There is a 2-approximate budget balanced group strategyproof cost sharing mechanism for the metric TSP game.

An $\alpha$-approximate cost allocation function $f$ is an $\alpha$-approximate competitive and cost recovering way of serving the entire set of users $U$, i.e., $\forall i \in U, f(i) \geq 0$ and $\alpha \mathrm{OPT}(U) \geq \sum_{i \in U} f(i) \geq C(U)$. The $\alpha$-core consists of all $\alpha$-approximate cost allocation functions $f$ such that $\forall S \subseteq$ $U, \sum_{i \in S} f(i) \leq \alpha \mathrm{OPT}(S)$. Clearly, an $\alpha$-approximate weakly cross-monotonic cost sharing method yields an $\alpha$-approximate cost allocation function for each subset of users. A shortcoming of the notion of core is that it turns out to be empty for many games. We hope that the relaxed notion introduced above will alleviate this difficulty.
Remark: In the definition of $\alpha$-approximate weakly crossmonotonic cost sharing method, replacing the condition $\sum_{i \in Q} \xi(Q, i) \geq \sum_{i \in Q} \xi(R, i)$ by $\sum_{i \in Q} \xi(R, i) \leq \alpha \mathrm{OPT}(Q)$ would have lead to a weaker, though still potentially useful, definition.

## 5. A GENERAL APPROACH

Next, we present an open problem, whose positive resolution would lead to a general technique for obtaining an $\alpha$ approximate budget balanced group strategyproof cost sharing method for several games based on NP-hard problems. For the time being, we can construct $\alpha$-approximate core allocations for these games. Let us first present some definitions.

A fractional set, $S$, is a set in which an element can appear partially i.e., with each element, $\epsilon$, 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 $f_{S_{1} \cup S_{2}}=\min \left\{f_{S_{1}}+f_{S_{2}}, 1\right\}$. 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 $C: 2^{U} \rightarrow \mathcal{R}^{+}$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 $C(S) \leq \sum_{j} f_{j} \cdot C\left(S_{j}\right)$, where each $S_{j}$ is a set of users.

The classic Bondareva-Shapley Theorem [3, 30] shows that a necessary and sufficient condition for the existence of a weakly cross-monotonic cost sharing method is that the underlying cost function exhibit the covering property.

Does a cost function satisfying the covering property always admit a cross-monotonic cost sharing method? A positive resolution will lead to the following general scheme.

A covering linear program is a minimization linear program in which all coefficients in the constraint matrix and the objective function are nonnegative. Let $\mathcal{L}$ be a covering linear program. The feasible region of homogenous linear inequalities is called a cone. A cone that lies in the nonnegative orthant will be called a nonnegative cone. Consider the feasible solutions of $\mathcal{L}$ that lie in a nonnegative cone $\mathcal{C}$. Corresponding to each user is a set of constraints of the linear program. For instance, in case of the minimum spanning tree LP, the constraints corresponding to a user $i$ corresponds to sets $S \subset V$ containing $i$ but not the root. Each user has a utility value to get the corresponding constraints satisfied. This utility value can be misreported. Let $\mathcal{C}_{i}$ be the set of constraints corresponding to a user $i$. The set of constraints corresponding to the users in $Q$ is $\bigcup_{i \in Q} \mathcal{C}_{i}$. The cost, $C(Q)$, for serving $Q$ is the optimum objective function value of a solution that satisfies all constraints in $\bigcup_{i \in Q} \mathcal{C}_{i}$ and lies in $\mathcal{C}$. Without loss of generality we assume that each set of constraints corresponds to a user.

Lemma 7. For any covering LP $\mathcal{L}$ and nonnegative cone $\mathcal{C}$, the cost function as defined above exhibits the covering property.

Proof. Suppose $S=\bigcup_{j} f_{j} \cdot S_{j}$, where $S$ and $S_{j}$ 's are sets of users. We want to establish that $C(S) \leq \sum_{j} f_{j}$. $C\left(S_{j}\right)$. Consider the right hand side. Let $\mathrm{x}_{\mathrm{j}}$ be an optimal solution for satisfying all the constraints in $S_{j}$ and lying in cone $\mathcal{C}$. Clearly, $\sum_{j} f_{j} \mathrm{x}_{\mathrm{j}}$ lies in the cone $\mathcal{C}$ as well. Hence, the lemma is proven if we show that $\sum_{j} f_{j} \mathrm{x}_{\mathrm{j}}$ serves all the users in $S$. Consider a user $i$ in $S$. User $i$ is served if all the constraints in $\mathcal{C}_{i}$ are satisfied. Consider a constraint, say $\mathbf{a}^{T} \cdot \mathbf{x} \geq b$, in $\mathcal{C}_{i}$. Note that there is no negative coefficient in a or $b$. Also we have restricted x to nonnegative cone $\mathcal{C}$. Therefore, $\mathbf{a}^{T} \cdot \sum_{j} f_{j} \mathrm{x}_{\mathrm{j}} \geq \mathbf{a}^{T} \cdot \sum_{j: z \in S_{j}} f_{j} \mathrm{x}_{\mathrm{j}}=\sum_{j: i \in S_{j}} f_{j} \mathbf{a}^{T}$. $\mathrm{x}_{\mathrm{j}} \geq \sum_{j: i \in S_{j}} f_{j} b \geq b$, where the last inequality follows from the fact that $i$ is covered in $\bigcup_{j} f_{j} \cdot S_{j}$. This shows that $\sum_{j} f_{j} \mathrm{x}_{\mathrm{j}}$ satisfies all the users in $S$. The cost of the optimum way to satisfy all users in $S, C(S)$, can only be smaller.

Consider an NP-hard minimization problem $\Pi$ for which an $\alpha$ factor approximation algorithm is obtained using an LP-relaxation as a lower bound, i.e., the cost of the solution found is at most $\alpha$ times an optimal solution to the LP. Furthermore, assume that this LP, $\mathcal{P}$, is a covering LP, $\mathcal{L}$, intersected with a nonnegative cone $\mathcal{C}$. (In most cases $\mathcal{C}$ will simply be the nonnegative orthant. As shown below, for facility location, a different cone is required.) If our open problem resolves positively, the optimal cost function for LP
$\mathcal{P}$ admits a cross-monotonic cost sharing method. Multiplying by $\alpha$ gives us an $\alpha$-approximate budget balanced group strategyproof cost sharing mechanism.

Using Lemma 7 and the Bondareva-Shapley Theorem [3, 30] one can show that there exists an $\alpha$-approximate weakly cross monotonic cost sharing method for $\Pi$. However, their theorem uses an exponential sized LP which may be solvable in polynomial time in particular cases, though not in general. The following theorem gives a way of finding one such method efficiently.

Theorem 8. There is an efficiently computable
$\alpha$-approximate weakly cross-monotonic cost sharing method for $\Pi$.

Proof. We may assume w.l.o.g. that there is a unique inequality in $\mathcal{L}$ corresponding to each user and that the constant in this inequality is 1 . The latter is easily ensured by scaling. Suppose the former is not satisfied for user $i$. Consider the inequalities corresponding to user $i$. Pick a new variable $x_{i}$ and replace each inequality $\mathbf{a}^{T} \cdot \mathrm{x} \geq b$ by the homogenous inequality $\mathbf{a}^{T} \cdot \mathbf{x} \geq b x_{i}$. These homogenous inequalities can be pushed into the cone. Further, add the inequality $x_{i} \geq 1$ in $\mathcal{L}$. Clearly the new LP is equivalent to the old one.

Let $\mathcal{D}$ be the dual linear program for $\mathcal{P}$. For subset $Q$ of users, let $\mathcal{P}_{Q}$ be the restriction of $\mathcal{P}$ to inequalities corresponding to $Q$ only (of course retaining all of $\mathcal{C}$ ). Let $\mathcal{D}_{Q}$ denote the dual of $\mathcal{P}_{Q}$.

Consider the following cost sharing method. Let $\mathbf{y}_{Q}$ denote an optimal solution to LP $\mathcal{D}_{Q}$. The cost share of user $i \in Q$ is $\xi(Q, i)=\alpha \mathbf{y}_{Q}(i)$. If $i \notin Q$, then $\xi(Q, i)=0$. Let $\operatorname{OPT}_{f}(Q)$ denote the objective function value of solution $\mathbf{y}_{Q}$.

This cost sharing method is $\alpha$-approximate competitive because $\sum_{i \in Q} \xi(Q, i)=\alpha \mathrm{OPT}_{f}(Q) \leq \alpha \mathrm{OPT}(Q)$. It satisfies cost recovery because the cost of the solution produced by the $\alpha$ factor approximation algorithm for $\Pi$ is $C(Q) \leq \alpha \operatorname{OPT}_{f}(Q)=\sum_{i \in Q} \xi(Q, i)$.

Finally, let us show that it satisfies weak cross-monotonicity. Let $Q \subseteq R$. Let $\mathbf{y}_{R}$ be an optimal solution to LP $\mathcal{D}_{R}$. Let $\mathrm{y}^{\prime}$ denote the restriction of $\mathrm{y}_{R}$ to coordinates corresponding to users in $Q$. The important observation is that $\mathrm{y}^{\prime}$ is a feasible solution to $\mathcal{D}_{Q}$. Since the optimal solution to $\mathcal{D}_{Q}$ can have only a higher objective function value, weak cross-monotonicity follows.

Let us show how Lemma 7 helps overcome some of the difficulties in obtaining a group strategyproof cost sharing method for the facility location game. In this game, 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 we are given the cost of connecting the city to a facility opened at that site. These connection costs satisfy the triangle inequality. In this setting cities are users who want themselves to be connected to an open facility. Each user reports a utility for being connected to an open facility. The cost of serving a set of users is the total cost of opening facilities and connecting each city in the set to one of the open facilities.

One difficulty is that the optimal solution does not admit a weakly cross-monotonic cost sharing method (and hence no cross-monotonic method either). The reason is that the
optimal cost function does not exhibit the covering property. For an instance, consider a cycle on 6 vertices, with 3 cities and 3 facilities alternating. The cost of each edge is 1 and the cost of opening each facility is 2 .

A well known LP-relaxation for this problem, due to Balinski [1], is given below. Several constant factor approximation algorithms are based on this relaxation. Suppose $C$ is the set of cities and $F$ is the set of facilities. Suppose $f_{i}$ is the cost of opening facility $i$ and $c_{i j}$ is the cost of connecting city $j$ to facility $i$. In the corresponding integer program, $y_{i}$ is a $0 / 1$ variable which is 1 iff facility $i$ is open, and $x_{i j}$ be a $0 / 1$ variable which is 1 iff city $j$ is connected to facility $i$.

$$
\begin{array}{lll}
\operatorname{minimize} & \sum_{j \in C, i \in F} c_{i j} x_{i j}+\sum_{i \in F} f_{i} y_{i} &  \tag{4}\\
\text { subject to } & \sum_{i \in F} x_{i j} \geq 1, & \forall j \in C \\
& y_{i}-x_{i j} \geq 0, & \forall j \in C, \forall i \in F \\
& x_{i j} \geq 0, & i \in F, j \in C \\
& y_{i} \geq 0, & i \in F
\end{array}
$$

This value is attained for the incremental cost-sharing method based on any permutation $\sigma$ such that $\sigma(n)=i$.

This leads to the following question regarding the minimum spanning tree game: what is the complexity of computing the minimum and maximum, over all cross-monotone cost-sharing methods, of the cost-share of an individual user in a given coalition? This problem appears to be NP-hard.
Let $\xi$ be a cross-monotone cost-sharing method for $C$. Represent $\xi$ as a point in $n 2^{n-1}$ dimensional real space whose coordinates are indexed by pairs $(Q, i)$, where $Q \subseteq U$ and $i \in Q$. Consider in this space the set of all crossmonotone cost-sharing methods for $C$. This set is a polytope, since it can be described by the following linear system on $n 2^{n-1}$ variables $x(Q, i)$ where $Q \subseteq U$ and $i \in Q$.

$$
\begin{array}{ll}
\sum_{i \in Q} x(Q, i)=C(Q), & \forall Q \subseteq U \\
x(Q, i) \geq x\left(Q^{\prime}, i\right), & \forall Q^{\prime} \subseteq U, Q \subseteq Q^{\prime}, i \in Q \\
x(Q, i) \geq 0, & \forall Q \subseteq U, i \in Q
\end{array}
$$

Proposition 10. Each incremental cost sharing method $\xi_{\sigma}$ for $C$ is a corner points of the polytope described above.

Proof. Suppose $\xi_{\sigma}$ is not a corner point of the polytope. In this case it can be written as the convex combination of two other cross-monotonic cost-sharing methods, $\xi_{a}$ and $\xi_{b}$. We claim that $\xi_{a}, \xi_{b}$ and $\xi_{\sigma}$ are the same. Suppose not. Then there exists $Q \subseteq U$ and $\sigma(i) \in Q$ such that $\xi_{a}(Q, \sigma(i)) \neq \xi_{\sigma}(Q, \sigma(i)) \neq \xi_{b}(Q, \sigma(i))$. Let us pick the smallest such $i$. Without loss of generality assume that $\xi_{a}(Q, \sigma(i))<\xi_{\sigma}(Q, \sigma(i))<\xi_{b}(Q, \sigma(i))$. Let $Q^{\prime}=\{\sigma(j): j \leq i\} \cap Q$. Since $\xi_{\sigma}$ is an incremental costsharing method,

$$
\sum_{\sigma(j) \in Q^{\prime}} \xi_{\sigma}(Q, \sigma(j))=C\left(Q^{\prime}\right)
$$

Therefore,

$$
\sum_{\sigma(j) \in Q^{\prime}} \xi_{b}(Q, \sigma(j))>C\left(Q^{\prime}\right) .
$$

This contradicts the cross-monotonicity of $\xi_{b}$.
Do all the corner points of the polytope correspond to incremental cost-sharing methods? The following cross-monotone cost-sharing method cannot be written as a convex combination of incremental cost-sharing methods, thereby showing
that the answer to this question is "No". $\xi(\{a, b, c\}, a)=$ $2, \xi(\{a, b, c\}, b)=3, \xi(\{a, b, c\}, c)=4, \xi(\{a, b\}, a)=4$, $\xi(\{a, b\}, b)=3, \xi(\{a, c\}, a)=3, \xi(\{a, c\}, c)=4, \xi(\{b, c\}, b)=$ $3, \xi(\{b, c\}, c)=4, \xi(\{a\}, a)=4, \xi(\{b\}, b)=4, \xi(\{c\}, c)=4$. Observe that the cost function of this example is particularly simple: it assigns costs to sets based only on their cardinality. Let us call such a cost function a cardinality cost function.

We leave the open problem of characterizing the rest of the corner points of this polytope. The special case of a cardinality cost function is also interesting.

## 7. ACKNOWLEDGEMENTS

We wish to thank Daniel Granot, Hervé Moulin, Suresh Mutuswami, Christos Papadimitriou, and Scott Shenker for valuable comments and for providing pointers into the literature.

## 8. REFERENCES

[1] M. L. Balinski. On finding integer solutions to linear programs. In Proc. IBM Scientific Computing Symp. on Combinatorial Problems, pages 225-248, 1966.
[2] C. Bird. On cost allocation for a spanning tree: a game theoretic approach. Networks, 6:335-350, 1976.
[3] O. N. Bondareva. Some applications of linear programming to cooperative games. Problemy Kibernetiki, 10:119-139, 1963.
[4] S. Deering and D. Cheriton. Multicast routing in datagram internetworks and extended lans. $A C M$ Transactions on Computer Systems, 8:85-110, 1990.
[5] B. Dutta and D. Ray. A concept of egalitarianism under participation constraints. Econometrica, 57:615-635.
[6] J. Edmonds. Optimum branchings. J. Res. Nat. Bur. Standards, B71:233-240, 1967.
[7] U. Faigle, S. Fekete, W. Hochstattler, and W. Kern. On approximation fair cost allocation in euclidean tsp games. In Electronic Colloquium on Computational Complexity, pages TR95-016, 1995.
[8] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmission. In Proc. 32th Ann. ACM Symp. on Theory of Comp., 2000.
[9] S. P. Fekete and W. R. Pulleyblank. A note on the traveling preacher problem. submitted to, Operations Research Letter.
[10] M. Goemans and M. Sketella. Cooperative facility location games. In Proc. 11th annual Symposium on Discrete Algorithm, pages 76-85, 2000.
[11] D. Granot and G. Huberman. On minimum cost spanning tree games. Mathematical Programming, 21:1-18, 1981.
[12] D. Granot and G. Huberman. On the core and nucleolus of minimum cost spanning tree games. Mathematical Programming, 29:323-347, 1984.
[13] J. Green, E. Kohlberg, and J. J. Laffont. Partial equilibirium approach to the free rider problem. Journal of Public Economics, 6:375-394, 1976.
[14] K. Jain and V. V. Vazirani. Group strategyproof cost allocation under a broader notion of a egalitarianism. In preparation.
[15] P. Karp, Koutsoupias and Shenker. Optimization problems in congestion control. In Proceedings of the 41st FOCS, 2000.
[16] K. Kent and D. Skorin-Kapov. On the core of the minimum steiner tree game in networks. In Annals of Operation Research, volume 57, pages 233-249, 1995.
[17] K. Kent and D. Skorin-Kapov. Population monotonic cost allocation on mst's. In Operational Reearch Proceedings KOI, pages 43-48, 1996.
[18] K. Kent and D. Skorin-Kapov. Distance monotonic stable cost allocation scheme for the minimum cost spanning tree network. In Proceedings of the 19th Conference Information Technology Interfaces ITI, pages 463-468, 1997.
[19] L. Kou, G. Markowsky, and L. Berman. A fast algorithm for Steiner trees. Acta Informatica, 15:141-145, 1981.
[20] Koutsoupias and Papadimitriou. Worst-case equilibria. In Proceedings of STACS, 1998.
[21] Kumar and Kleinberg. Fairness measures for resource allocation. In Proceedings of the 41st FOCS, 2000.
[22] N. Megiddo. Cost allocation for steiner trees. Networks, 8:1-6, 1978.
[23] H. Moulin. Incremental cost sharing: characterization by coalition strategy-proofness. Social Choice and Welfare, 16:279-320, 1999.
[24] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. http://www.aciri.org/shenker/cost.ps, 1997.
[25] S. Mutuswami. Strategyproof mechanism for cost sharing. working paper, Indian statistical Institute.
[26] Nisan. Algorithms for selfish agents - mechanism design for distributed computation. In Proceedings of STACS, 1999.
[27] Nisan and Ronen. Algorithmic mechanism design. In Proceedings of STOC, 1999.
[28] K. Roberts. The characterization of implementable choice rules. In J. J. Laffont, editor, Aggregation and Revealation of Preferences, pages 321-348, North Holland, 1979.
[29] Roughgarden and Tardos. How bad is selfish routing. In Proceedings of the 41st FOCS, 2000.
[30] L. S. Shapley. On balanced sets and cores. Naval Research Logistics Quarterly, 14:453-460.
[31] L. S. Shapley. A value of $n$-person games. In H. Kuhn and W. Tucker, editors, Contributions to the theory of games, pages 31-40. Princeton University Press, 1953.
[32] V. V. Vazirani. Approximation Algorithms. Springer Verlag, Berlin, To appear, 2001.


[^0]:    * Address: One Microsoft Way, Redmond, WA 98052, Email: kamalj@microsoft.com.
    ${ }^{\dagger}$ Address: College of Computing, Georgia Institute of Technology, Atlanta, GA 30332, Email: vazirani@cc.gatech.edu. Research supported by NSF Grant CCR-9820896.

[^1]:    Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
    STOC'01, July 6-8, 2001, Hersonissos, Crete, Greece.
    Copyright 2001 ACM 1-58113-349-9/01/0007 ...\$5.00.

