Cross-Monotonic Cost-Sharing Schemes for Combinatorial Optimization Games: A Survey

Siamak Tazari
University of British Columbia
Department of Computer Science
Vancouver, British Columbia
siamakt@cs.ubc.ca

ABSTRACT

We consider the problem of sharing the cost of receiving a service among a set of users. This problem has been studied extensively in the coalitional game theory literature. In recent years, a certain class of cost-sharing methods, called cross-monotonic cost-sharing schemes, has received special attention, since it has been shown that they lead to group-strategyproof mechanisms. We study the properties, characterization and limitations of these schemes in the domain of combinatorial optimization games.

1. INTRODUCTION

A service is available for a set of agents N. For every subset of agents $S \subseteq N$, there is a certain cost $C(S)$ associated with providing them with this service. A *cost-sharing scheme* specifies for every agent i and every subset S, how much agent i should pay for the service, if the subset S is receiving the service. For concreteness, consider the following example: a service-provider is providing financial news over a network. For every piece of news, a certain subset of users is interested in receiving it. The cost of broadcasting the news to a set of users is equal to the cost of routing the news through the network to them. Assuming that for every link in the network, there is a certain price for using it, it could be cheaper to send the news through a tree to the set of interested users instead of sending it to each user individually. Determining the minimum cost of broadcasting a message in this setting is known as the *multicast routing* problem. Here we consider the problem of sharing the cost of multicast routing: how much should every individual be charged for receiving the news?

Cost-sharing schemes are closely related to *cost-sharing mechanisms*: suppose every agent i has a privately known utility u'_i for receiving the service. If he receives the service for a price x_i, he enjoys a benefit of $u'_i - x_i$; if he does not receive the service his benefit is 0. A cost-sharing mechanism asks every agent to declare their utility for the service and decides which subset of users to serve and how much to charge every agent. But agents act selfishly and might misreport their true utility in order to receive the service for a lower price. Hence, we would like our mechanism to be *strategyproof*[1]: an agent should not be able to benefit by lying. A stronger property is that of being *group-strategyproof*: no coalition of agents should be able to benefit by jointly misreporting their values.

There are a number of desirable and sometimes conflicting properties that we would like to have for cost-sharing schemes. One of them is *budget-balance*: the amount of money collected from the agents should be equal to the actual cost of serving them. Another property, introduced in [22, 21] is that of *cross-monotonicity*[2]: the amount charged to an agent should not increase if the subset of agents being served grows. This encourages the agents to promote the service and is in tune with intuitions about economies of scale. Cross-monotonic cost-sharing schemes have received special attention in the literature, since Moulin and Shenker [16] proved the fundamental theorem that cross-monotonic cost-sharing schemes lead to group-strategyproof cost-sharing mechanisms. We will study the existence and limitations of such schemes with regard to other properties - the most important one being budget-balance - for various cost-functions, especially cost-functions based on combinatorial optimization problems.

2. BASIC DEFINITIONS

Agents. In the following, $N = \{1, \ldots, n\}$ is the set of agents. Every agent i has a utility value u'_i for receiving the service. We assume that agents' utilities are quasi-linear, i.e. are of the form $u'_i q_i - x_i$, where $q_i \in \{0, 1\}$ is an indicator variable specifying if i receives the service or not and x_i is the amount that i has to pay.

Cost function. A cost function is a function C that assigns to every subset of agents, a cost of serving them. We assume it is non-negative and $C(\emptyset) = 0$. It is said to be *non-decreasing* if $S \subseteq T \Rightarrow C(S) \leq C(T)$. A cost function is *submodular* if $C(S \cap T) + C(S \cup T) \leq C(S) + C(T)$. Submodularity is the economies of scale condition and is a property widely studied in the literature [16, 10]. It says that the marginal cost $C(S \cup \{i\}) - C(S)$ of adding an agent i to a set S does not increase when the set S expands. As an

[1] i.e. dominant-strategy truthful

[2] in some of the literature, this property is referred to as population monotonicity.

example, consider the multicast routing problem where the cost of broadcasting the message is calculated as follows: a universal spanning tree T, i.e. a tree containing the service provider and all the users, is fixed; when a subset S of users requests service, the message is sent along the subtree T_S of T that contains S and the service provider. One can check that this cost function is submodular and non-decreasing.

Cost sharing mechanism. A cost sharing mechanism $\mathscr{M}$ is a function that given any vector u of agents' utilities, returns: (i) a subset Q of agents to be served (note that Q can also be represented by its indicator variables $q_i \in \{0,1\}$) and (ii) an amount x_i that is charged to agent i, for every i. There are a number of properties that are desirable for $\mathscr{M}$. We consider the following[3]:

No Positive Transfers (NPT): $\forall i. x_i \geq 0$, no agent is being paid for receiving the service.

Voluntary Participation (VP): If an agent i is not served, then $x_i = 0$ and otherwise $x_i \leq u_i$, i.e. every agent has the option not to participate and receive a utility of 0. In the game-theory literature, this property is also referred to as ex-interim individual rationality.

Consumer Sovereignty (CS): Every agent is guaranteed service if he reports a high enough utility value.

Cost Recovery : $\sum_{i \in Q} x_i \geq C(Q)$, the cost of serving the participants is recovered. This is also known as weak budget-balance.

Competitiveness : $\sum_{i \in Q} x_i \leq C(Q)$, no surplus is created. Otherwise a competitor could offer the service for a cheaper cost.

Budget-Balance (BB): Both cost recovery and competitiveness are satisfied, i.e. $\sum_{iinQ} x_i = C(Q)$.

Efficiency : $\sum_{i \in Q} u_i - C(Q)$ is maximized, i.e. as much worth as possible is created.

Group-Strategyproofness : If C is a coalition of users and no member of C is worse off by misreporting their utility value, then no member is better off either. More formally, let $u_i = u'_i$ for all $i \notin C$ and let (q, x) and (q', x') be the respective outcomes of the mechanism for u and u'. Then if $u'_i q_i - x_i \geq u'_i q'_i - x'_i$ for all $i \in C$, it must hold with equality for all $i \in C$. That is, even if a group of agents collude, they can not benefit by jointly misreporting their utilities.

Cost allocation. A cost allocation for a subset $Q \subseteq U$ is a function $\psi : Q \to \mathbb{R}^+ \cup \{0\}$ that specifies for each agent in Q how much he has to pay for receiving service. It is budget-balanced if $\sum_{i \in Q} \psi(i) = C(Q)$. A cost allocation function is said to be *in the core* if it is budget-balanced and furthermore $\forall S \subseteq Q$, $\sum_{i \in S} \psi(i) \leq C(S)$, i.e. no subset of agents has an incentive to secede.

Cost sharing scheme. A cost sharing scheme is collection of cost allocations for every subset of U. More formally, a cost sharing scheme is a function $\xi : 2^U \times U \to \mathbb{R}^+ \cup \{0\}$ such that if $i \notin Q$, then $\xi(Q, i) = 0$. It is said to be budget-balanced (resp. in the core) if $\xi(Q, \cdot)$ is budget-balanced (resp. in the core) for every $Q \subseteq U$. A cost sharing scheme is *cross-monotonic* if for $R \supseteq Q$, $\xi(R, i) \leq \xi(Q, i)$ for every agent i, i.e. the cost share of any agent does not increase as the subset being served expands. It is not hard to see that a budget-balanced cross-monotonic cost sharing scheme is also in the core but the converse need not hold.

[3]These definitions are largely taken from [9].

3. THE MOULIN-SHENKER THEOREM

In [16], Moulin and Shenker prove the following fundamental theorem:

Theorem 1. For any budget-balanced cross-monotonic cost sharing scheme ξ, there is a mechanism $\mathscr{M}(\xi)$ that meets BB, NPT, VP, CS and is group-strategyproof. Furthermore, if the cost function is submodular, the converse is also true.

The mechanism $\mathscr{M}(\xi)$ is simple: initialize $Q = U$; offer service to the agents in Q at the price $\xi(Q, i)$ dropping anyone who refuses (i.e. his utility is less than the offered price); repeat until $u_i \geq \xi(Q, i)$ for all $i \in Q$. Note that since ξ is cross-monotonic, the order in which the agents are offered service in the above algorithm does not matter, as the offered price can only increase in a later stage.

They study the case of non-decreasing submodular cost functions more closely. Consider the following *incremental cost-sharing method*: let σ be an arbitrary fixed ordering of the agents, let $S \subseteq N$ be a given subset of k agents and $i_1, \ldots, i_k$ be the elements of S, ordered according to σ. Set $\xi_\sigma(S, i_1) = C(\{i_1\})$ and for $2 \leq j \leq k$, $\xi_\sigma(S, i_j) = C(\{i_1, \ldots, i_j\}) - C(\{i_1, \ldots, i_{j-1}\})$. If C is submodular and non-decreasing, this cost-sharing scheme is BB and cross-monotonic. The classical Shapley value [20] is the arithmetic mean of these schemes over all possible σ. It is a classical result in game theory that no strategyproof mechanism is both budget-balanced and efficient, even for submodular cost functions. Moulin and Shenker show that the Shapley value has the lowest worst-case loss of efficiency over all utility profiles [16].

Whereas the Shapley value tries to charge more to agents that cause more cost, Dutta and Ray introduced the *egalitarian cost-sharing method* [3] that attempts to equalize the cost-shares among the agents as much as possible while preserving the core property. Moulin and Shenker [16] show that the egalitarian method is also a BB-cross-monotonic cost-sharing scheme if the underlying cost function is submodular (but not necessarily non-decreasing).

4. MULTICAST ROUTING AND MINIMUM SPANNING TREES

For the multicast routing problem, we can calculate the cost of service to a set of users S, as explained in Section 2, by fixing a universal spanning tree T and considering the cost of a subtree T_S containing the source (i.e. the service provider) and S. This cost function is submodular and non-decreasing [4], so the Shapley value and the egalitarian method both provide us with a BB-cross-monotonic cost-sharing scheme.

There is one problem with this approach: the cost of the *optimal* tree T^*_S connecting the source to the set of serviced agents, could be arbitrarily smaller than the cost of T_S. Notice that T^*_S could even include vertices not in S. T^*_S is known as the *minimum Steiner tree* containing the source and the set S. Finding a minimum Steiner tree is well-known to be NP-hard. Furthermore, if we define the cost of servicing a set S to be the cost of T^*_S, then this cost function is neither submodular nor non-decreasing. But in the case of metric Steiner tree, when the costs of the links of the network define a metric, there exists a 2-approximation algorithm for this problem based on the minimum cost spanning tree (MCST) problem. So, we turn our attention to this problem now and return later to the multicast routing problem.

The well-known MCST problem is given by a graph G with weights on edges and the goal is to find a tree of minimum weight connecting all the vertices. In the MCST cost-sharing game, the cost of providing service to a set of users is given by the MCST that contains them. This problem is well-studied in the literature [11, 9, 15, 2]. Kent and Skorin-Kapov [11] present a class of cost-sharing schemes in the core of this game and one BB-cross-monotonic cost-sharing scheme. In [9], Jain and Vazirani present a class of BB-cross-monotonic cost-sharing schemes parameterized by n *equalizing functions* $f_i : \mathbb{R}^+ \rightarrow \mathbb{R}^+$. These functions can decode some fairness criteria and could for example, relate to the probability distribution functions of user's utilities. They use linear-programming duality in order to achieve this result. Moretti et al. [15] present a BB-cross-monotonic cost-sharing scheme based on Kruskal's greedy MCST algorithm. In [2], Branzei et al. define the *P-value* for MCST, a BB-cross-monotonic cost-sharing scheme derived from the one presented in [15] in a similar fashion as the Shapley-value is derived from incremental cost-sharing schemes.

5. ON APPROXIMATION

It is not always possible to achieve budget-balance together with other properties; or it might be computationally hard to compute the cost shares. For example, in the case of Steiner tree's, we saw that even computing the cost of an optimal Steiner tree is NP-hard. Therefore, Jain and Vazirani [9] introduced the concept of *α-budget-balance*: a cost allocation function ψ for a subset Q of U is α-budget-balanced if $\alpha C(Q) \leq \sum_{i \in Q} \psi(i) \leq C(Q)$. That is, only a fraction α of the cost is recovered[4]. A cost sharing scheme ξ is α-budget-balanced if $\xi(Q, .)$ is α-budget-balanced for every Q in U. The definition of the *α-core* is the same as the definition of the core except that we require α-budget-balance instead of budget-balance. Again, one can show that an α-budget-balanced cross-monotonic cost sharing scheme lies in the α-core but the converse need not hold. A mechanism is said to be α-budget-balanced if its cost allocation has this property.

Jain and Vazirani [9] show that Theorem 1 still holds if BB is replaced by α-BB. As mentioned earlier, the metric Steiner tree problem can be 2-approximated by the MCST problem. Hence, any one of the BB-cross-monotonic cost-sharing schemes for MCST presented in the last section result immediately in a $\frac{1}{2}$-BB cross-monotonic cost-sharing scheme for the Steiner tree game, i.e. also for multicast routing.

For the metric traveling salesman problem (TSP), there exists a 2-approximation based on doubling an MCST; thus, the above mentioned cost-sharing methods also result in a $\frac{1}{2}$-BB cross-monotonic cost-sharing scheme for the TSP-game [9].

Further in the paper [9], Jain and Vazirani also turn to a more general approach: they consider a class of NP-hard minimization problems for which there exists an α factor approximation algorithm based on a linear program (LP). They show that if this LP has a certain property, called the covering property, then there exists an efficiently computable cost sharing scheme in the α-core of the corresponding game. They leave the question of whether there is also a cross-monotonic cost sharing scheme as an important open problem.

Biló et al. [1] consider the problem of multicast routing in wireless networks. In this problem, we assume that the nodes of the network are in d-dimensional euclidean space and every node i has to be assigned a range r_i for transmitting messages. In doing so, it consumes some power, namely $\gamma \cdot r_i^\alpha$, where γ and α are parameters and $\alpha \geq d$ is assumed. Again, a message is to be sent from a source node to a given subset of nodes and the goal is to find a range-assignment, such that the total energy consumption is minimized. The question is again how to distribute this cost among the receivers of the message. Biló et al. [1] argue that if a universal spanning tree T is fixed and costs are calculated based on it, then the associated cost-function is non-decreasing and sub-modular and so there exist cross-monotonic cost-sharing schemes. However, if we consider the optimal range-assignment as our cost-function, then this result only applies if $d = 1$ or $\alpha = 1$. For the case where $\alpha > 1$ or $d > 1$, they study the relationship between wireless multicast routing and the (standard) multicast routing and show that the budget-balanced allocations presented earlier for MCST yield a $\frac{1}{2(3^d-1)}$-BB cross-monotonic cost-sharing scheme for this problem.

6. EQUITABLE COST ALLOCATION

In [10], Jain and Vazirani study the case of submodular cost functions more closely. By generalizing their method of [9], they present a class of (efficiently computable) BB-cross-monotonic cost sharing schemes, derived from a primal-dual type algorithm, parameterized by *equalizing functions* $f_i : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ for every agent i. As in [9], these functions can encode some fairness criterion. They could be used to provide some kind of price discrimination - a property that is crucial to the survival of many industries (e.g. consider airlines: business travelers should be charged more than casual travelers). If these functions are all set to identity, then their method results exactly in the egalitarian method presented by Dutta and Ray [3]; so, it includes this method as a special case. They define the notion of *max-min fairness* and *min-max fairness* for a given set of equalizing functions; max-min fairness aims to maximize the minimum cost-share and min-max fairness aims to minimize the maximum cost-share, thus these criteria ensure that no one underpays / overpays. They denote this property of satisfying both of these fairness criteria at the same time together with being in the core of the game by *opportunity egalitarianism* and show that their cost sharing method is the only opportunity egalitarian method for any given set of equalizing functions. By the end of the paper, they present some special functions that could be used as the equalizing functions such that certain properties are achieved, e.g. maximizing acceptance probability.

[4]One can divide the given inequality by α and thus relax competitiveness instead of cost-recovery. This is, in fact, the way that Jain and Vazirani did it in [9]. The results remain the same.

They also note the Hokari [7] independently generalized Dutta-Ray solutions to give a class of cost-sharing methods that turns out to be identical to these equitable cost allocation methods. Hokari called these methods *monotone path cost allocations*, but it is interesting to note that his formalization and point of view are quite different from Jain and Vazirani - and so are his definitions and algorithms (he does not address issues of algorithmic efficiency, though).

7. MINIMUM ARBORESCENCE, SHORTEST PATH AND NETWORK DESIGN

One can consider a problem similar to MCST in directed graphs, namely finding the minimum arborescence: a set of edges connecting every vertex via a directed path to the source. In [15], Moretti et al. present an example that shows that the corresponding cost-sharing game does not have a budget-balanced cross-monotonic cost-sharing scheme. However, in [14], Moretti et al. consider a simpler problem: finding a minimum arborescence in a directed acyclic graph, where every vertex is guaranteed to have at least a direct link to the source (but this link could be more expensive than using another path). Their study is motivated by the problem of connecting houses on a mountain with a purifier - this problem results exactly in the given setting. They present a BB-cross-monotonic cost-sharing method for the corresponding cost-sharing problem.

In the shortest-path problem, a graph G together with a source s and destination t are given. Every agent is an edge of the graph. In our terminology, he receives service, if he is part of the shortest path connecting s to t. For a given subset of agents, the cost of serving them is equal to the shortest path connecting s to t using only edges of the given subset. This setting makes more sense when regarded in the context of the corresponding cost-sharing mechanism: suppose every agent has a utility for being part of the shortest path. Then the mechanism has to select a shortest path based on the agents' declared utilities and decide how much to charge each agent. Voorneveld and Grahn [23] present a BB-cross-monotonic cost-sharing scheme for a more general setting of this game, where agents are allowed to own a set of edges instead of owning only one edge.

Network design is a term that encompasses a wide range of combinatorial optimization problems. We first consider metric facility-location: we are given a set of facilities F and a set of users N. We want to open a subset of these facilities and connect every user to some facility at minimum cost. Opening a facility p involves a cost f_p and there is also a cost associated with using each link in the network. We assume that the costs of the links obey the triangle inequality, i.e. define a metric. The corresponding cost-sharing problem is given by the cost function where the cost of serving a subset of users is the cost the minimum facility-location connecting only the given subset of users to the facilities.

Pal and Tardos [17] present a general technique for turning a primal-dual algorithm into an α-budget-balanced cross-monotonic cost sharing scheme and thus, by Theorem 1, into a group strategyproof mechanism. Achieving cross-monotonicity can be at the cost of getting a weaker budget-balance-factor α than the approximation factor of the given primal-dual algorithm. Their idea is to use a "ghost process" that virtually keeps paying for the costs during the course of their algorithm and thus ensures cross-monotonicity; the challenge is to show that after these virtual payments are removed, a constant factor of the budget is still being recovered. By applying their method to the metric facility-location game, they achieve a $\frac{1}{3}$-BB cross-monotonic cost-sharing scheme for this problem.

They also consider the single-source rent-or-buy (SSRB) problem: given is a graph G with source s and edge-weights c_e (for every edge e) and a number M. The goal is to find a tree connecting every node to the source, where every edge can either be *bought* at cost $M \cdot c_e$ or *rented* at cost c_e times the number of users that use the edge on their way to the sink. By applying their primal-dual scheme to this problem, they present a $\frac{1}{15}$-BB cross-monotonic cost-sharing scheme for it.

In a follow-up work, Gupta et al. [6] further investigate the SSRB problem by building on the idea of sharing the expected cost of a randomized algorithm given by Gupta et al. in [5]. By using derandomization techniques on that algorithm, they achieve a $\frac{1}{4.6}$-BB cross-monotonic cost-sharing scheme (computable in polynomial time) for this problem - a large improvement over the previous $\frac{1}{15}$ factor.

In the connected facility-location problem, in addition to the requirements of the metric facility location, we also require that the open facilities must be connected via a Steiner tree. With some small alterations, it is also possible to regard SSRB as a special case of this problem. Leonardi and Schäfer [13] present a $\frac{1}{30}$-BB cross-monotonic cost-sharing mechanism for this problem based on the methods suggested by Pal and Tardos in [17].

Finally, we consider the Steiner Forest game: an undirected graph G with edge weights is given, together with a set of k terminal pairs (a terminal pair is a pair of vertices). The goal is to find a subgraph of minimum cost that connects each pair (s, t) of the given terminal pairs (such a subgraph will of course always be a forest). Könemann et al. [12] recently derived a $\frac{1}{2}$-BB cross-monotonic cost-sharing scheme for this network design problem.

8. LIMITATIONS OF CROSS-MONOTONICITY

In a recent paper, Immorlica, Mahdian and Mirrokni [8] study the limitations of cross-monotonic cost sharing schemes. They use a novel technique based on the probabilistic method to derive upper bounds for the budget-balance factor of cross-monotonic cost sharing schemes for several combinatorial optimization games including edge cover, vertex cover, set cover, metric facility location, maximum flow, arborescence packing and maximum matching. The cost function of these games simply assigns to every set of agents the cost of a minimum edge cover, vertex cover, etc. of the subproblem induced by that set of agents. In maximization games we can think of profit sharing instead of cost sharing. For metric facility location, they show an upper bound of $\frac{1}{3}$ - equal to the lower bound achieved by Pal and Tardos [17] - thus closing the gap. For the set cover game they show that no more than an $O(\frac{1}{n})$-fraction of the cost can be recovered. So, basically all hope is lost for this important universal problem. For vertex cover they show an upper bound of $O(n^{-1/3})$, demonstrating that cross-monotonicity is strictly harder to achieve than allocation in the core, since the vertex cover game has a cost-sharing-scheme in the $\frac{1}{2}$-core.

Finally, they consider the implications of their results on the existence of α-budget-balanced group-strategyproof mechanisms. Since the cost functions of these combinatorial optimization games are not submodular, the back-direction of Theorem 1 does not hold. In fact, they give examples of trivial group-strategyproof mechanisms that recover all the cost but intuitively do not seem fair nor efficient (e.g. put the whole burden on only one agent). To resolve this problem, they state two additional properties: *no free riders*, requiring that no one receives the service for free; and *upper continuity*, stating that for every agent i, if he receives service for every bid greater than u_i while holding other bids fixed, then he also gets service if he bids u_i. They show the following theorem:

Theorem 2. The cost function C has an upper-continuous α-budget-balanced group-strategyproof mechanism with no free riders if and only if it has an α-budget-balanced cross-monotonic cost-sharing scheme.

Thus, their negative results also apply to group-strategyproof mechanisms satisfying these two additional properties.

9. BEYOND CROSS-MONOTONICITY

If the underlying cost function is not submodular, the back direction of Theorem 1 does not hold. Hence, there exist group-strategyproof mechanisms that are not derived from cross-monotonic cost-sharing schemes. In [18], Penna and Ventre define the notion of *self-cross-monotonicity*. This notion is closely related to the mechanism given in Theorem 1. Essentially it means that the cross-monotonicity property need only hold for subsets of users that could possibly be considered by that mechanism - not for all subsets[5]. They show that if an approximation algorithm bears certain “reasonable” properties, it will always admit a self-cross-monotonic cost-sharing scheme satisfying NPT, VP and CS with the same BB-factor as the approximation ratio of the algorithm. Based on this idea they derive a group-strategyproof mechanism for the *optimal* Steiner tree game that is *budget-balanced*. This might seem surprising at first, since the Steiner tree problem is NP-hard to begin with; but their trick is that they grow and only consider a subset of users for which they know that their optimal Steiner tree is equal to their MCST. This is accomplished by using Prim’s greedy MCST algorithm. This same mechanism is $\frac{1}{3^d-1}$-BB for the wireless network multicast routing game.

Their suggested mechanisms suffer, though, from a major flaw: they fall under the mechanisms that are categorized as “trivial” in [8], since they put the whole burden of paying the cost on only one agent. That is, if a set of m agents is being served, then $m-1$ of them are always free riders and only one agent has to pay for them all. They tackled this problem in a very recent follow-up paper [19], where they study the free-rider issue for Steiner-tree games. They present a BB group-strategyproof mechanism with at most $|S| - |leaves(T)|$ free riders, when S is the subset being served and T is the Steiner tree being used. A similar result holds for the wireless multicast routing game, where BB is replaced by $\frac{1}{3^d-1}$-BB. In general, they prove the following negative result:

Theorem 3. Let $C_{OPT}(\cdot)$ be a cost function which is NP-hard (to approximate within a factor α). Then, no polynomial-time strategyproof mechanism guarantees no free riders and satisfies NPT, VP, CS and (α-approximate) BB, unless $P = NP$.

[5]the exact definition is rather complicated and out of the scope of this text.

10. CONCLUSIONS AND FUTURE WORK

Since Moulin and Shenker’s paper [16], a lot of attention has been payed to cross-monotonic cost-sharing schemes in the theoretical computer science community. For submodular cost-functions, a large class of such schemes exist and is characterized. But for many other non-submodular cost-functions - such as the functions resulting from most combinatorial optimization games - it seems to be particularly difficult to derive such schemes. In the literature this is often done via linear-programming duality or other often complicated methods.

The negative results of Immorlica et al. [8] show that for many interesting combinatorial optimization games, it is not even possible to achieve good approximations on the budget-balance factor when cross-monotonicity is required. A possible future direction would be to try to introduce randomness into cross-monotonicity and group-strategyproofness; in many areas of computer science, it has been possible to overcome limitations by using randomness. Hence, this direction might be a promising area.

11. REFERENCES

[1] V. Biló, C. D. Francescomarino, M. Flammini, and G. Melideo. Sharing the cost of multicast transmissions in wireless networks. In *Proceedings of SPAA*, pages 180–187, 2004.

[2] R. Branzei, S. Moretti, H. Norde, and S. Tijs. The p-value for cost sharing in minimum cost spanning tree situations. Technical Report 129, Tilburg University, Center for Economic Research, 2003.

[3] B. Dutta and D. Ray. A concept of egalitarianism under participation constraints. *Econometrica*, 57(3):615–635, May 1989.

[4] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of mulicast transmissions. In *Proceedings of STOC*, pages 218–227, 2000.

[5] A. Gupta, A. Kumar, and T. Roughgarden. Simpler and better approximation algorithms for network design. In *Proceedings of STOC*, pages 365–372, 2003.

[6] A. Gupta, A. Srinivasan, and É. Tardos. Cost-sharing mechanisms for network design. In *Proceedings of APPROX-RANDOM*, pages 139–150, 2004.

[7] T. Hokari. Monotone-path dutta-ray solutions on convex games. *Social Choice and Welfare*, 19(4):825–844, October 2002.

[8] N. Immorlica, M. Mahdian, and V. Mirrokni. Limitations of cross-monotonic cost sharing schemes. In *Proceedings of SODA*, pages 602–611, 2005.

[9] K. Jain and V. Vazirani. Applications of approximation algorithms to cooperative games. In *Proceedings of STOC*, pages 364–372, 2001.

[10] K. Jain and V. Vazirani. Equitable cost allocations via primal-dual-type algorithms. In *Proceedings of STOC*, 2002.

[11] K. Kent and D. Skorin-Kapov. Population monotonic cost allocation on mst’s. In *Operational Research Proceedings KOI*, pages 43–48, 1996.

[12] J. Könemann, S. Leonardi, and G. Schäfer. A group-strategyproof mechanism for steiner forests. In *Proceedings of SODA*, pages 612–619, 2005.

[13] S. Leonardi and G. Schäfer. Cross-monotonic cost-sharing methods for connected facility location

games. In *ACM Conference on Electronic Commerce*, pages 242–243, 2004.

[14] S. Moretti, H. Norde, K. H. P. Do, and S. Tijs. Connection problems in mountains and monotonic allocation schemes. Technical Report 12, Tilburg University, Center for Economic Research, 2001.

[15] S. Moretti, H. Norde, and S. Tijs. Minimum cost spanning tree games and population monotonic allocation schemes. Technical Report 18, Tilburg University, Center for Economic Research, 2001.

[16] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18:511–533, 2001.

[17] M. Pal and E. Tardos. Strategy proof mechanisms via primal-dual algorithms. In *Proceedings of FOCS*, 2003.

[18] P. Penna and C. Ventre. More powerful and simpler cost-sharing methods. In *Proceedings of WAOA*, pages 97–110, 2004.

[19] P. Penna and C. Ventre. Free-riders in steiner tree cost-sharing games. In *Proceedings of SIROCCO, to appear*, 2005.

[20] L. S. Shapley. A value of n-person games. *Contributions to the theory of games*, pages 31–40, 1953.

[21] Y. Sprumont. Population monotonic allocation schemes for cooperative games with transferable utility. *Games and Economic Behavior*, 2(4):378–394, December 1990.

[22] W. Thomson. The fair division of a fixed supply among a growing population. *Mathematics of Operations Research*, 8(3):319–326, August 1983.

[23] M. Voorneveld and S. Grahn. Cost allocation in shortest path games. *Mathematical Methods of Operations Research*, 56(2):323–340, November 2002.