This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2008 proceedings.

Cross-Monotonic Multicast

Zongpeng Li
Department of Computer Science
University of Calgary
zongpeng@ucalgary.ca

Abstract*—In the routing and cost sharing of multicast towards a group of potential receivers, cross-monotonicity is a property that states a user's payment can only be smaller when serviced in a larger set. Being cross-monotonic has been shown to be the key in achieving group-strategyproofness. We study multicast schemes that target optimal flow routing, cross-monotonic cost sharing, and budget balance. We show that no multicast scheme can satisfy these three properties simultaneously, and resort to approximate budget balance instead. We derive both positive and negative results that complement each other for directed and undirected networks. We show that in directed networks, no cross-monotonic scheme can recover a constant fraction of optimal multicast cost. We provide a simple scheme that does achieve $\frac{1}{k}$-budget-balance, where k is the number of receivers. Using a probabilistic method rooted in random graph theory, we prove an upper-bound of $\frac{2}{\sqrt{k}}$ for the budget balance ratio. For undirected networks, we derive a constant upper-bound of $\frac{1}{2}$ instead. We further apply a smooth dual growing technique to design a cross-monotonic scheme that recovers $\frac{k+1}{2k\zeta}$ of optimal multicast cost in undirected networks, where ζ is a network-dependent parameter close to 1. This is almost tight against the upper-bound $\frac{1}{2}$. We finally present a two-stage linear optimization model that pursues maximum budget balance in any given specific network, with trade-off in complexity. Optimization results in various network configurations confirm the theoretically established bounds.

Index Terms: Game Theory; Graph Theory; Mathematical Programming/Optimization; Multicast.

I. Introduction

Multicast models one-to-many data dissemination in communication networks, such as the streaming of live video to a large group of Internet users. Traditional multicast algorithm design has been exclusively based on multicast trees or meshes that can be decomposed into a set of trees [1]–[4]. A fundamental limit of the tree-based approach is that optimizing multicast throughput and cost are equivalent to the combinatorial problems of Steiner tree packing and minimum Steiner tree respectively [2], [5], and are NP-hard [6], [7]. However, a recent breakthrough in information theory dramatically changed the picture. A new multicast feasibility characterization is developed by exploiting the encodable as well as replicable properties of information flows: a multicast rate d is feasible in a directed network if and only if it is feasible as a unicast to each receiver separately [8], [9]. Efficient multicast algorithm design were soon spawned from this new network flow based multicast routing structure, in directed, undirected, and wireless network models [10]–[14].

The aforementioned multicast algorithms are all designed for cooperative network environments and focus on optimality and performance. Nonetheless, it has been a growing trend in networking research to realize that it is not always safe to assume that network agents are honest, cooperative and altruistic, and that it is desirable to design network algorithms that are efficient and robust against rational but strategic and selfish agents [3], [15]–[18]. The presence of non-cooperative behaviors introduces a new dimension of constraint into the multicast problem, and makes optimal algorithm design even more challenging. In previous work [18], we have studied multicast routing where information flows are selfish and always select routes that minimize their respective cost. Using shadow price based cost sharing and taxing schemes, we showed that every optimal multicast flow can be successfully enforced, while maintaining fairness and budget balance.

The problem studied in this paper arises from selfish multicast receivers instead. In the market of a multicast network, payments are collected from flows or receivers and paid to links. A multicast solution contains a flow routing scheme and a cost sharing scheme that can be applied to any set of potential receivers. Cost shares collected from receivers are used to cover flow costs at links across the network. Each user has a (potentially different) valuation of the multicast service, known to herself only. The key challenge is to induce receivers to report their true valuations, for better deciding whom to serve, and how much to charge [17], [19], [20]. A *strategyproof* mechanism is one in which each user's dominant strategy [21] is to tell the truth, *i.e.*, lying will not provide any benefit to her own interest. A *group-strategyproof* mechanism is one that is robust against collusion.

It has been shown that the key towards group-strategyproofness is to have *cross-monotonic* cost sharing, where each user's payment smoothly decreases as the service set expands [19], [20], [22]. Beside cross-monotonicity, we would like to also pursue optimality in flow routing, *i.e.*, keeping the routing cost to be shared as low as possible, and a balanced budget, *i.e.*, recovering the flow cost from user payments. Multicast schemes we develop are also *in-core* [23], *i.e.*, no users in the service set have incentive to secede and build their own solution, and satisfy the *no-positive-transfer* property [1], *i.e.*, cost shares are never negative.

We study cross-monotonic multicast in both directed and undirected network models, and provide both positive and negative results that complement each other in each model. We depart from the previous *de facto* standard in cross-monotonic mechanism design, where a pair of corresponding primal and dual solutions are simultaneously manipulated and used in the solution [20], [22], [24], [25]. We compute the primal solution (flow routing) and the dual (cost sharing) independently. The advantage is that we are not obliged to resolve the conflict between primal optimality and dual smoothness. The trade-off is that we need to bound the size of the dual, from both above and below, using the cost of a primal to which it is not directly coupled.

It was observed that the integrality gap in the linear optimization model of a game usually implies an upper-bound on the cost recovery ratio of any cross-monotonic scheme for that game [24], [25]. The presence of exact LP models for multicast [12], [18] suggests some hope on exact budget balance. However, we show that, unfortunately, optimality, cross-monotonicity and budget balance are not simultaneously feasible in general, and relaxing the exact budget balance requirement is therefore a necessary compromise. For directed networks, we design a simple multicast scheme that is optimal, cross-monotonic, and $\frac{1}{k}$-budget-balanced, where k is the number of potential receivers. We further show that no cross-monotonic scheme can always recover a constant fraction of the optimal multicast cost. We design a network pattern based on a complete l-partite hypergraph, with each hyperedge replaced by a gadget tailored for the multicast game. We then apply a probabilistic method [24], [26] to show that the budget balance ratio of any optimal and cross-monotonic scheme is at most $\frac{2}{\sqrt{k}}$ in this network.

We adapt the network pattern and proof technique to undirected networks, and derive a constant upper-bound $\frac{1}{2}$ on the budget balance ratio for optimal and cross-monotonic multicast schemes, contrasting the diminishing bound in directed settings. We provide insights on such a dramatical difference. Assuming sufficient link bandwidth supply, we further design a cross-monotonic multicast scheme that recovers $\frac{k+1}{2k\zeta}$ fraction of the optimal cost, for a parameter ζ close to 1. This pair of negative and positive results are almost tight against each other, since $\frac{k+1}{2k\zeta}$ is close to $\frac{1}{2}$. The multicast scheme design here is based on smooth dual variable growing, which has its root in primal-dual algorithm design [20], [22], [25]. The idea is to trade-off the ratio of cost recovery for smoothness in cost sharing. In particular, we make use of the spanning tree game that is cross-monotonic, optimal, and exactly budget balanced [22]. When bandwidth provisioning in the undirected network is frugal instead, we show that the achievable budget balance ratio is $O(\frac{1}{\beta})$, where β is a parameter of the network that measures how unbalanced the costs are on inter-terminal paths.

We also present a two-stage linear optimization model that pursues the absolute maximum budget balance in specific networks, under optimal routing and cross-monotonic cost sharing constraints. The solutions found vary case by case and do not admit a general characterization, and the associated computational complexity is much higher than that of the general solutions. Experiments in both contrived and randomly generated networks based on this model confirm the theoretically established bounds.

The rest of the paper is organized as follows. We review previous research in Sec. II and explain the network model and notations in Sec. III. Achievable budget balance ratio is studied in Sec. IV for directed networks and in Sec. V for undirected networks. Sec. VI discusses optimizing budget balance in specific network instances, and Sec. VII concludes the paper.

II. Previous Research

Ahlswede *et al.* [8] initiated the study of network coding, which allows information to be encoded/decoded at any node across the network. The most celebrated result in network coding research characterizes multicast feasibility in directed networks: *if a multicast rate d is feasible to each receiver independently, then it is feasible to all receivers simultaneously as a multicast* [8], [9]. Based on this new characterization, efficient algorithms have been designed for optimizing multicast throughput and cost in both directed and undirected networks [10]–[12]. In contrast to many well-studied problems such as set cover and Steiner tree, which are NP-hard, the multicast cost can be modelled as a linear program and efficiently computed. We therefore enforce flow optimality in the multicast schemes discussed throughout this paper.

Moulin and Shenker [19] studied strategyproof sharing mechanisms for costs that are submodular. It was shown that no mechanism can be simultaneously strategyproof, budget balanced, and efficient. An *efficient* mechanism is one that maximizes net social utility [19]. They described a universal transformation from a cross-monotonic cost sharing into a group-strategyproof mechanism, using a simultaneous Cournot tatonnement [21]. Basically, for each set of potential users (start with the full set), compare each user's computed cross-monotonic share with her reported valuation. Exclude a user if the latter is smaller and repeat the procedure until the set stabilizes. In such a game, no user has incentive to lie about her true valuation, even in the presence of collusion. For submodular costs, it was shown that the Shapley value [27] method always leads to cross-monotonic sharing. However, multicast cost is not submodular either with or without network coding; it is not even sub-additive in one of the network models we study, as shown later.

Jain and Vazirani [22] studied cross-monotonic sharing of spanning tree and Steiner tree costs in undirected networks. Based on the dual growing technique rooted in primal-dual algorithm design, they provide an elegant solution for spanning tree that is optimal, cross-monotonic, and exactly budget balanced. This represents the only such scheme known for non-submodular costs so far. The primal (tree) and dual (cost shares) solutions together also imply a $\frac{1}{2}$-optimal, cross-monotonic, and $\frac{1}{2}$-budget-balanced mechanism for building a Steiner tree and sharing its cost. P*á*l and Tardos [20] extended the smooth dual growing technique to facility location and single source rent-or-buy games, resulting in cross-monotonic sharing methods recovering $\frac{1}{3}$ and $\frac{1}{15}$ of the solution cost, respectively. Since both problems are NP-hard, the primal solutions are necessarily approximate ones.

Erd*ö*s and R*é*nyi [26] pioneered probabilistic graph theory research and designed a new method for showing a graph of a certain property exists: instead of constructing a specific instance, which may be hard depending on the property, one may describe a procedure for randomly building graphs and argue that the desired property is satisfied in the output with probability strictly larger than zero [26], [28]. Immorlica *et al.* [24] successfully adapted the probabilistic proof technique in analyzing the limitations of cross-monotonic cost sharing. They describe random constructions of game instances where the expected cost recovery ratio is low, and therefore establish bounds on the achievable budget balance ratio for games such as set cover, vertex cover, and facility location. We also apply the probabilistic technique in this paper, to derive negative results on cross-monotonic multicast schemes.

Wang *et al.* [17] studied the Steiner tree game where both links and nodes can be selfish. Since a strategyproof mechanism against selfish links computes payments to links, it is then natural to share the payments instead of the true multicast cost among the nodes. Both positive and negative results on the budget balance ratio are derived for sharing the payments of a link-strategyproof mechanism that is $\frac{1}{2}$-optimal. We consider in this paper multicast with network coding, where the multicast flow is not necessarily equivalent to a tree or a mesh composed of a set of trees. We also focus on selfish receivers only, and share the exact optimal multicast cost among them.

III. Network Model and Preliminaries

We use a graph $G = (V, E)$ to denote a (directed or undirected) network topology, and vectors $w, c \in Q_+^E$ for link cost and link capacities, respectively. Here Q_+ is the set of positive rational numbers. In a multicast network, $S \in V$ denotes the sender, $T \subseteq V$ denotes the set of potential receivers, and nodes in $V - S - T$ are relay nodes. $k = |T|$ denotes the total number of receivers. In graphical illustrations throughout the paper, *terminal nodes* (S, T) are in black and relay nodes are in white. We use interchangeably the terms *cost shares* and *payments*, *a user* and *a receiver*.

We use vector $f_A \in Q_+^E$ to denote a multicast flow from S to a set of receivers $A \subseteq T$. When A is a singular set $\{u\}$ we simply write f_u. Under the classic linear edge cost model [29], [30], the optimal flow f_A^* can be efficiently computed by solving a linear program [12], [18], using either general methods such as the interior-point algorithm, or tailored sub-gradient algorithms with better running times [11], [12], [18]. We present the LP for directed networks below; the LP for undirected networks is similar [5].

$$\text{Minimize} \qquad \sum_{\vec{uv}} w(\vec{uv}) f(\vec{uv}) \tag{3.1}$$

Subject to:

$$\begin{cases} \sum_{v\in N_\downarrow(u)} f_i(\vec{uv}) = \sum_{v\in N_\uparrow(u)} f_i(\vec{vu}) & \forall T_i \in A, \forall u \\ f_i(\vec{T_iS}) = d & \forall T_i \in A \\ f_i(\vec{uv}) \le f(\vec{uv}) \le c(\vec{uv}) & \forall T_i \in A, \forall\, \vec{uv} \end{cases}$$

$$f_i(\vec{uv}), f(\vec{uv}) \ge 0 \qquad \forall T_i \in A, \forall\, \vec{uv}$$

The LP above assumes that $A = \{T_1, \dots, T_{|A|}\}$. The objective function models the total link flow costs to be minimized. The constraints, in order of appearance, model the conservation of conceptual flows, multicast rate requirement, the merging of conceptual flows, and link capacity bounds, respectively [18]. $\vec{T_iS}$ is a conceptual feedback link introduced to make the LP compact.

We use $y_A(u)$ to represent the cost share of a receiver $u \in A \subseteq T$. $|f_A| = \sum_e w(e)f(e)$ denotes the cost of a multicast flow f_A to receiver set A. $|A|$ denotes the cardinality of a receiver set A. Budget balance can be expressed as $\sum_{u\in A} y_A(u) = |f_A|, \forall A \subseteq T$, for a multicast scheme (f, y). For $x \in [0, 1]$, we say a multicast scheme (f, y) is *x-budget-balanced* if $x|f_A| \le \sum_{u\in A} y_A(u) \le |f_A|, \forall A \subseteq T$, *i.e.*, y always recovers at least x fraction of $|f_A|$ but never more than $|f_A|$. y is *cross-monotonic* if $y_A(u) \le y_B(u)$, $\forall u \in B \subseteq A \subseteq T$. A cost sharing scheme y is *in-core* if $\sum_{u\in B} y_A(u) \le |f_B^*|, \forall B \subseteq A \subseteq T$. If a cross-monotonic y is x-budget-balanced with regard to an optimal flow f^*, then y is automatically in-core.

We say a multicast network has *sufficient* bandwidth provisioning, if (3.1) has the same optimal objective value with or without the constraint in $f \le c$; otherwise we say the network has *frugal* bandwidth provisioning. A set function $g : 2^U \to Q$ is *sub-additive* if $g(A\cup B) \le g(A)+g(B), \forall A, B \subseteq U$; and is *submodular* if $g(A\cup B) \le g(A)+g(B)-g(A\cap B), \forall A, B \subseteq U$. Finally, $E_x(g(x))$ denotes the expectation of a function $g(x)$ over a random variable x.

IV. Directed Networks

We start our study with directed networks, where each link transmits flows in a predetermined direction only. We show that local cost sharing and direct LP dual based sharing are not cross-monotonic, and provide a simple scheme that is optimal, cross-monotonic, and $\frac{1}{k}$-budget-balanced. We prove that no cross-monotonic scheme can recover exactly the optimal multicast cost or even any constant fraction of it, by establishing an upper-bound of $\frac{2}{\sqrt{k}}$ on the budget balance ratio, using the probabilistic method. This is in spite of the fact that min-cost multicast has exact LP formulations.

A. Existing Schemes

Local Sharing is not in-Core. A *local* cost sharing method allocates the multicast cost among users in an edge-wise fashion, and the allocation at each edge depends only on the flow rates of users on that edge. For instance, equal sharing [3], flow proportional sharing, and the Shapely value method [1], [27] all belong to this category. We use an example in Fig. 1 to show that applying local sharing on f^* directly is not in-core, and hence cannot be cross-monotonic (note that local sharing may be cross-monotonic if routing is not required to be optimal).

Here a directed multicast network is shown on the left, each edge is labelled with its unit flow cost. The optimal multicast flow is given on the right, where T_1 has a unit flow on both

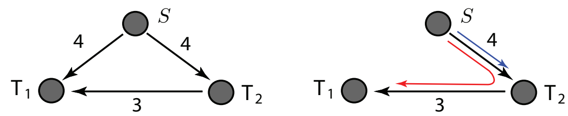

Fig. 1. A multicast network and optimal routing f^*. Required multicast rate $d = 1$ and each link has capacity larger than 1.

links and T_2 has a unit flow on $\overrightarrow{ST_2}$ only. In any local sharing, $y_1 = \frac{4}{2} + 3 = 5$ and $y_2 = \frac{4}{2} = 2$. Now T_1 has incentive not to participate in f^*, since its share is larger than its own optimal solution $|f^*_{T_1}| = 4$. Therefore the solution is not in-core.

Direct LP Dual is Stabilizing but not Smooth. Since min-cost multicast has exact LP formulations, it is natural to consider variables in the dual LP for cost sharing. We have used the dual variables for the flow merging constraints in LP (3.1) as multicast cost shares [18], to successfully enforce selfish multicast flows to stabilize at the min-cost state. Unlike local sharing, LP dual takes into consideration the global network topology and alternate solutions each receiver has. However, in optimal LP solutions, both f^*_A and y^*_A may jump dramatically as the service set $A \subseteq T$ varies, leading to a non-smooth sharing as A grows. In fact, as we will show in Theorem 2, an optimal, cross-monotonic and budget balanced scheme does not exist; since (f^*, y^*) is optimal and exactly budget balanced, it cannot be cross-monotonic.

B. A $\frac{1}{k}$-Budget-Balanced Multicast Scheme

We describe a simple multicast scheme in Table 3.1, which can be applied generally in directed networks with any topology and any link capacity configurations. It uses the optimal multicast flow f^*_A for routing, while charging nodes based on their respective optimal unicast costs in f^*_u.

Table 3.1. Cross-Monotonic Multicast in Directed Networks

For any $A \subseteq T$:

(1) *Solve multicast LP (3.1), let f^*_A be the optimal solution*

(2) for each receiver u in A:
Solve min-cost $S \to u$ unicast flow f^*_u

(3) *Route multicast flows as specified in f^*_A*

(4) *Let each $u \in A$ pay $y_A(u) = \frac{|f^*_u|}{|A|}$*

Theorem 1. The multicast scheme in Table 3.1 is optimal, cross-monotonic, and $\frac{1}{k}$-budget-balanced.

Proof: The optimality of the multicast routing follows directly from the fact that f^*_A is the optimal solution of LP (3.1). To show cross-monotonicity, let A and B be two sets of potential receivers such that $u \in A \subset B \subseteq T$. Since $A \subset B$ implies $|A| < |B|$, we have:

$$y_B(u) = \frac{|f^*_u|}{|B|} < \frac{|f^*_u|}{|A|} = y_A(u)$$

Therefore a user's cost share can get only smaller in a larger user set, and cross-monotonicity is maintained in the sharing.

We next show $\frac{1}{k}$-budget-balance. For any $A \subseteq T$:

$$\sum_{u \in A} y_A(u) = \frac{\sum_{u \in A} |f^*_u|}{|A|} \geq \frac{|f^*_A|}{|A|} \geq \frac{1}{k}|f^*_A|$$

In the derivations above, $\frac{\sum_{u \in A} |f^*_u|}{|A|} \geq \frac{|f^*_A|}{|A|}$ is based on the celebrated feasibility condition for multicast with network coding [8], [9], which decomposes multicast rate feasibility into unicast rate feasibility. More specifically, let f' be the union flow of the f^*_u's, *i.e.*, $f'(\overrightarrow{uv}) = \max_u f^*_u(\overrightarrow{uv}), \forall\, \overrightarrow{uv}$. By the new multicast feasibility condition, f' constitutes a feasible multicast flow of rate d to user set A. Since f' is a union of the f^*_u's, we have (a) $\sum_{u \in A} |f^*_u| \geq |\, f'|$; furthermore, by the optimality of f^*_A we have (b) $|f^*_A| \leq |\, f'|$. Combining (a) and (b) leads to $\frac{\sum_{u \in A} |f^*_u|}{|A|} \geq \frac{|f^*_A|}{|A|}$.

The final step is to show that $\sum_{u \in A} y_A(u) \leq |f^*_A|$:

$$|f^*_u| \leq |\, f^*_A|, \forall u \in A \Longrightarrow \sum_{u \in A} |f^*_u| \leq |\, A||f^*_A|$$

$$\Longrightarrow \sum_{u \in A} \frac{|f^*_u|}{|A|} \leq |f^*_A| \Longrightarrow \sum_{u \in A} y_A(u) \leq |f^*_A|$$

□

The positive result above may appear unsatisfying in that the fraction of cost recovered in the cross-monotonic sharing diminishes as the user group expands. We next present a negative result that shows, unfortunately, constant-fraction cost recovery is impossible.

C. An Upper-Bound on Budget Balance Ratio: $\frac{2}{\sqrt{k}}$

We now establish an upper-bound on the budget balance ratio of optimal cross-monotonic multicast schemes. We do so by first constructing a network pattern based on a complete l-partite hypergraph, and replacing each hyperedge by a tailored multicast gadget. We then select a user set A to show its cost recovery ratio is at most $\frac{2}{\sqrt{k}}$. The selection here is naturally probabilistic, and the recovered cost is an expected value. The rationale is that since each (multicast) scheme comes with its own weakness and forte, it is hard to construct a deterministic scenario where *every* scheme performs 'badly' [24]. Nonetheless, the low expected cost recovery ratio manifests the existence of a specific (scheme-dependent) instance of A where the real ratio is no better.

Theorem 2. For any $\epsilon > 0$, there does not exist a multicast scheme that simultaneously guarantees optimality, cross-monotonicity and $(\frac{2}{\sqrt{k}} + \epsilon)$-budget-balance in general directed networks.

Proof: We construct a directed multicast network as shown in Fig. 2. Terminal nodes in the network include a multicast source S, and $k = lh$ potential multicast receivers grouped into l partites each of size h. For every combination of l nodes each from a different partite, we connect S to them using the gadget shown on the right. The target multicast rate $d = 1$,

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2008 proceedings.

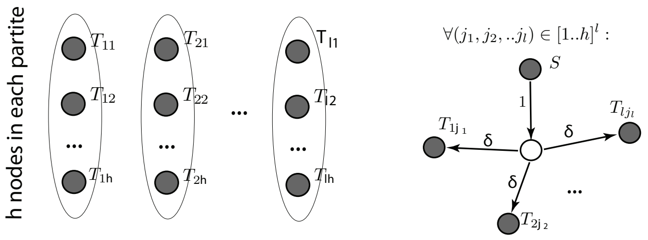

Fig. 2. A directed network pattern for which the achievable budget balance ratio is bounded by $\frac{2}{\sqrt{k}}$.

every edge has capacity larger than 1. Each edge is labelled with its cost; here δ is a small number.

We compose the set of multicast receivers A by constructing two subsets A_1 and A_2, and let $A = A_1 \cup A_2$. For A_1: uniformly randomly pick a partite i, include all nodes there; for A_2: in every other partite, uniformly randomly pick a node and include it in A_2. We next show that due to limitations imposed by cross-monotonicity, the expected cost recovered from such a randomly chosen set A is at most $\frac{2}{\sqrt{k}}$ of the cost incurred in f_A^*.

$$\begin{aligned} & E_A(\textstyle\sum_{T_{ij}\in A} y_A(T_{ij})) \\ =_1 \quad & E_A(\textstyle\sum_{T_{ij}\in A_1} y_A(T_{ij})) + E_A(\textstyle\sum_{T_{ij}\in A_2} y_A(T_{ij})) \\ \leq_2 \quad & hE_{A_2,T_{ij}\in A_1}(y_{A_2+T_{ij}}(T_{ij})) \\ & +E_{A_2,T_{ij}\in A_1}(\textstyle\sum_{T_{ij}\in A_2} y_{A_2+T_{ij}}(T_{ij})) \\ =_3 \quad & h(\tfrac{1}{l}+\delta) + (l-1)(\tfrac{1}{l}+\delta) \\ =_4 \quad & (h+l-1)(\tfrac{1}{l}+\delta) \end{aligned}$$

In the derivations above, $\leq_2$ is based on cross-monotonicity of y. Here T_{ij} is uniformly randomly selected from A_1 and $E_{A_2,T_{ij}\in A_1}(y_{A_2+T_{ij}}(T_{ij}))$ is the expected cost share of T_{ij} within the user set $A_2 + T_{ij}$. Since $A_2 + T_{ij} \subset A$, $y_A(T_{ij}) \leq y_{A_2+T_{ij}}(T_{ij})$ by cross-monotonicity. The following fact is critical in the proof: uniformly randomly picking a user from every partite $j \neq i$ and then uniformly randomly picking a user from partite i is equivalent to uniformly randomly picking a gadget γ from the entire network. Each user in the randomly picked gadget will have the same expected cost share: $\frac{|f_\gamma^*|}{|\gamma|} = \frac{1+l\delta}{l} = \frac{1}{l}+\delta$, and that leads to $=_3$.

Observe that the real cost $|f_A^*| = h+(h+l-1)\delta$. Therefore the expected cost recovery ratio is:

$$E_A(x_A) = \frac{(h+l-1)(\frac{1}{l}+\delta)}{h+(h+l-1)\delta} = \frac{\frac{1}{l}+\delta}{\frac{h}{h+l-1}+\delta}$$

By way of contradiction, assume that for some $0 < \epsilon < 1$, there exists an optimal, cross-monotonic and $(\frac{2}{\sqrt{k}}+\epsilon)$-budget balanced multicast scheme. Choose the values for l, h, and δ based on ϵ, such that:

$$l = h = \sqrt{k} > \frac{2}{1-\epsilon}, \qquad \delta < \frac{h^2\epsilon+1}{(h-h\epsilon-2)(2h-1)},$$

we then have:

$$\frac{\frac{1}{l}+\delta}{\frac{h}{h+l-1}+\delta} = \frac{\frac{1}{h}+\delta}{\frac{h}{2h-1}+\delta} < \frac{2}{h}+\epsilon = \frac{2}{\sqrt{k}}+\epsilon$$

Therefore $E_A(x_A) < \frac{2}{\sqrt{k}}+\epsilon$, contradiction. □

D. Discussions

The upper-bound $\frac{2}{\sqrt{k}}$ suggests that any optimal and cross-monotonic multicast scheme may suffer from poor budget balance in directed networks, since no constant ratio cost recovery is possible in general. This is fundamentally related to the fact that flow transmission on each link is one direction only. Note that in the network in Fig. 2, links to relay nodes can be shared, and links to receivers are private. At private links, the cost recovery is essentially exempted from limitations by cross-monotonicity — no choice of A would make a node u assume that its private links can be shared, and exact budget balance is always achievable at these links. The global budget balance ratio therefore depends on two factors: (a) how small the ratio of budget balance can be forced on shared links by cross-monotonicity, and (b) what is the ratio of shared link cost over private link cost in f^*. While directed and undirected networks share a similar fate in (a), they differ dramatically in (b). The ratio in (b) can be arbitrarily high for directed networks, leading to a rather poor budget balance ratio. In Sec. V, we reveal a different picture for undirected networks.

There is a gap between our positive result (the feasibility of $\frac{1}{k}$-budget balance) and negative result (the upper-bound of $\frac{2}{\sqrt{k}}$ on budget balance). To close this gap, one needs to either design a multicast scheme with asymptotically better cost recovery, or prove a tighter upper-bound. We leave this as future research. The positive and negative results for undirected networks, as shown next, are almost tight against each other, though. The network pattern in Fig. 2 is also empirically studied in Sec.VI, where the limitation of cross-monotonicity is illustrated using numerical results.

V. Undirected Networks

We now switch our focus to undirected networks. We modify the network pattern in Fig. 2 for the undirected setting, and apply the probabilistic method to derive a negative result again: a cross-monotonic multicast scheme can recover no more than half of the optimal multicast cost in general. Assuming sufficient bandwidth supply, we then apply the dual growing technique [20], [22], [25] to derive a smooth cost allocation, which recovers at least $\frac{k+1}{2k\zeta}$ of the optimal cost, and is almost tight against the $\frac{1}{2}$ upper-bound. For frugal bandwidth provisioning, we show that the achievable budget balance ratio is $O(\frac{1}{\beta})$, where β is a network parameter reflecting cost variance among inter-terminal paths.

A. An Upper-Bound on Budget Balance Ratio: $\frac{1}{2}$

Theorem 3. For any $\epsilon > 0$, there does not exist a multicast scheme that simultaneously guarantees optimality, cross-

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2008 proceedings.

monotonicity, and $(\frac{1}{2}+\epsilon)$-budget-balance in general undirected networks.

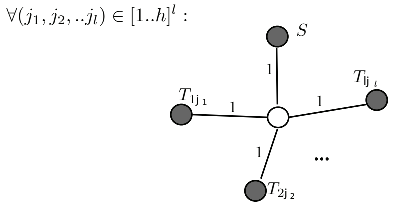

Fig. 3. The undirected multicast gadget.

Proof: we construct a network model similar to the one used in the directed case in Fig. 2. We still start from a complete l-partite hypergraph. However, each hyperedge is replaced with a different gadget, which is tailored for the undirected case and is shown in Fig. 3. The new gadget consists of undirected links each of the same unit cost. The target multicast rate is still $d = 1$ and $c(e) \geq 1, \forall e$.

By randomly picking a subset of receivers $A = A_1 \cup A_2$ in the same way as in the proof of Theorem 2, we have that the total expected pay in A:

$$\begin{aligned} & E_A(\textstyle\sum_{T_{ij}\in A} y_A(T_{ij})) \\ = \; & E_A(\textstyle\sum_{T_{ij}\in A_1} y_A(T_{ij})) + E_A(\textstyle\sum_{T_{ij}\in A_2} y_A(T_{ij})) \\ \leq \; & hE_{A_2, T_{ij}\in A_1}(y_{A_2+T_{ij}}(T_{ij})) \\ & +E_{A_2, T_{ij}\in A_1}(\textstyle\sum_{T_{ij}\in A_2} y_{A_2+T_{ij}}(T_{ij})) \\ = \; & h(1+\tfrac{1}{l}) + (l-1)(1+\tfrac{1}{l}) \\ = \; & (h+l-1)(1+\tfrac{1}{l}) \end{aligned}$$

Furthermore, the optimal multicast cost $|f_A^*| = 2h + l - 1$. Therefore the expected cost recovery ratio is:

$$\begin{aligned} E_A(x_A) &= \frac{E_A(\sum_{T_{ij}\in A} y_A(T_{ij}))}{|f_A^*|} \\ &= \frac{(h+l-1)(1+\frac{1}{l})}{2h+l-1} \\ &= \frac{1}{2} + \frac{l^2+l+2h-2}{2l(2h+l-1)} \end{aligned}$$

Now, by way of contradiction, assume that there does exist an optimal, cross-monotonic multicast scheme recovering $\frac{1}{2}+\epsilon$ of its cost, for some constant $0 < \epsilon \leq \frac{1}{2}$. Pick h and l such that:

$$l > \frac{1}{2\epsilon}, \qquad h > \frac{(l-1)(l+2-2l\epsilon)}{2(2l\epsilon-1)}$$

We then have $(l^2+l+2h-2)/2l(2h+l-1) < \epsilon$, contradiction. □

Note that unlike in the directed case, if links adjacent to receivers have very small cost instead of 1, then arguments in the proof of Theorem 3 would not hold. In particular, receivers in A would interconnect among themselves through relay nodes to get the multicast flow, leading to a smaller $|f_A^*|$ and a larger budget balance ratio achieved. Private links in f^* cannot have arbitrarily small costs, otherwise they may be utilized by other receivers as well, since there is no restriction in link direction in undirected networks.

B. A $\frac{k+1}{2k\zeta}$-Budget-Balanced Multicast Scheme

We next present a positive result on the achievable budget balance ratio, assuming network bandwidth is sufficiently provisioned. Recall that this means LP (3.1) has the same optimal objective value with or without constraint $f \leq c$; practically, having $c(e) \geq d$ at every edge e is sufficient (but not necessary). The case where bandwidth supply is tight will be discussed afterwards.

The multicast scheme presented below is optimal, cross-monotonic, and $\frac{k+1}{2k\zeta}$-budget-balanced. Here ζ denotes the *coding advantage*, *i.e.*, the ratio of achievable multicast throughput with network coding over that without network coding, in the specific network [5], [31]. ζ is proven to be between 1 and 2 and is believed to be close to 1 [31], [32]. In reality, the observed value for ζ is at most $\frac{8}{7}$ for networks with unbounded sizes [32], is at most $\frac{9}{8}$ for small contrived networks [5], and is essentially always 1 for random networks [5]. Therefore $\frac{k+1}{2k\zeta}$ should be very close to $\frac{1}{2}$, and this scheme is almost tight given the $\frac{1}{2}$ lower bound on budget balance established in Theorem 3.

The cross-monotonic multicast scheme here makes use of the classic spanning tree game [22] that is optimal, cross-monotonic, and exactly budget balanced. We apply this game in the closure graph of G and show that it leads to a cross-monotonic multicast cost sharing in G with approximate budget balance.

For any $A \subseteq T$, we first construct a closure graph $G' = (V', E')$ from $G = (V, E)$, such that $V' = A$ and $E' = \{(uv)|\forall u, v \in A\}$, and let $w(uv) = d_G(uv), \forall u, v \in V'$ be the link cost vector for G'. Here $d_G(uv)$ denotes the shortest path length between u and v in G under metric w.

We next run the spanning tree game in G'. In this game, each node u is associated with a *potential* $p(u)$ that is initially zero and grows over time. If at time t, for two nodes u and v in different components, $p(u) + p(v) = w(uv)$, then edge uv becomes *tight* and we merge the two components that contain u and v into one. The rate of potential growth at a node u at time t is $\frac{dp_u(t)}{dt} = \frac{1}{|P_u(t)|}$, where $P_u(t)$ denotes nodes in the same component as u at time t. The total growing rate of a component is always 1. The game stops when all nodes merge into one component. A spanning tree can be constructed from edges that were ever tight. In the original spanning tree game, each node $u \in A$ pays $\int_{t=0}^{t_u} \frac{2}{|P_u(t)|}dt$; we scale the payment in our game to $y_A(u) = \frac{2(k+1)d}{2k\zeta}\int_{t=0}^{t_u} \frac{1}{|P_u(t)|}dt$ instead. Here t_u is the time when u first became a member of the source component. For flow routing, we still solve (3.1) for the optimal multicast flow f_A^*.

Theorem 4. The multicast scheme based on (f^*, y) above is optimal, cross-monotonic, and $\frac{k+1}{2k\zeta}$-budget-balanced.

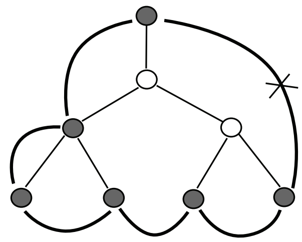

Fig. 4. A Steiner tree in G and a corresponding spanning tree in G'. Take any planar drawing of the Steiner tree, connect all terminal nodes in order using links in G', then connect back to the first terminal. Break the most expensive link in the resulting cycle to form a spanning tree in G'. In each face of the graph, the spanning tree edge cost is no more than the cost of the rest of the face boundary, hence the cost of the cycle is at most twice the cost of the Steiner tree; the spanning tree cost is at most $\frac{k}{k+1}$ fraction of the cycle cost.

Proof: Routing optimality of the scheme follows from the definition of f^*. Cross-monotonicity can be established by observing the following fact: in the spanning tree game, having an extra node in A can only make a node u grow its potential slower, join the source component earlier, and therefore reduce its cost share. We next show $\frac{k+1}{2k\zeta}$-budget-balance by bounding the ratio between the cost of the spanning tree in G' and $|f^*|$.

Let (G, c') be the capacitied network derived from scaling f^*: $c' = \zeta f^*/d$. By the definition of ζ, we should be able to route a multicast flow of 1 in (G, c') without coding, *i.e.*, we can pack at least 1.0 Steiner trees in (G, c'). This manifests the existence of a specific Steiner tree Δ in G, such that $|\Delta| \le \sum_{e\in E} c'(e)$.

Let MST be the tree built in G' during the spanning tree game — MST must be a *minimum* spanning tree in G', since the way it is built is in line with Kruskal's algorithm. The following relation must hold between MST and Δ, a Steiner tree in G connecting terminals in A [22]:

$$|MST| \le \frac{2k}{k+1}|\Delta| \qquad (5.1)$$

A brief illustration of the validity of (5.1) is given in Fig. 4.

Therefore we have:

$$\frac{1}{d}|f_A^*| \le_5 |MST| \le_6 \frac{2k}{k+1}|\Delta| \le_7 \frac{2k\zeta}{(k+1)d}|f_A^*| \qquad (5.2)$$

Here $\le_5$ is due to the optimality of f_A^*, the fact that MST can be used to route a multicast flow of 1, and the fact that c represents a sufficient bandwidth supply. $\le_6$ is from (5.1). $\le_7$ is by the definition of ζ and Δ. (5.2) further implies:

$$\frac{k+1}{2k\zeta}|f_A^*| \le \sum_{u\in A} y_A(u) = \frac{(k+1)d}{2k\zeta}|MST| \le |f_A^*|.$$

□

C. The Case of Frugal Bandwidth Provisioning

In the previous section we have assumed that link bandwidth are over-provisioned in the network. We now further discuss the scenario where bandwidth supply in the network is tight to support the desired throughput d. We modify the multicast scheme in Sec. V-B to obtain an optimal, cross-monotonic, and $O(\frac{1}{\beta})$-budget-balanced multicast scheme. Here β is the *path balance factor* of the network, *i.e.*, the maximum cost difference ratio between two paths connecting terminal nodes. We also show that better than $O(\frac{1}{\beta})$-budget-balance is impossible through a counter example.

Given the capacitied network $N = (G, w, c)$, let $N' = (G, w)$ be its uncapacitied version. For any $A \subseteq T$, we can run the $\frac{k+1}{2k\zeta} = O(1)$ budget balanced cost sharing algorithm in N', and use f_A^*, the optimal multicast flow to A in N, for routing. This multicast scheme is optimal, due to the way f_A^* is defined. It is cross-monotonic, due to the cross-monotonicity of the uncapacitied cost sharing method proved in Theorem 4. Now, let f' be the optimal multicast flow to A with rate d in N', and W be the reduced cost of f_A^* by making every inter-terminal path in f_A^* to have a cost within $O(1)$ of its shortest inter-terminal path. $O(\frac{1}{\beta})$-budget-balance follows from:

$$O(\frac{1}{\beta})|f_A^*| =_8 O(1)W =_9 O(1)|f'| =_{10} \sum_{u\in A} y_A(u)$$
$$\le_{11} |f'| \le_{12} |f_A^*|$$

Here $=_8$ is due to the fact that by relaxing the budget balance ratio from $O(1)$ to β, each inter-terminal path is at most $O(\beta)$ times more expensive. $=_9$ is based on the definition of W and the fact that the total inter-terminal flow rates in f_A^* is within $O(1)$ of that in f'. $=_{10}$ and $\le_{11}$ are from Theorem 4. $\le_{12}$ is due to the fact that both f' and f_A^* are optimal solutions of (3.1), while f' is obtained with one less group of constraints.

The example in Fig. 5 further shows that, $O(\frac{1}{\beta})$-budget-balance is indeed the best we can do in general. In this example, every link has a unit capacity of 1. The two links adjacent to R has cost α and the other three links have cost 1. The target multicast throughput is $d = 2$.

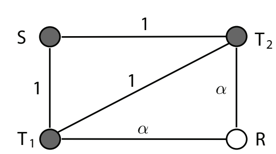

Fig. 5. When bandwidth supply is tight, multicast cost is not sub-additive in undirected networks.

Therefore we have:

$$|f_{T_1}^*| = |f_{T_2}^*| = 3, \;\; |f_T^*| = 2\alpha + 3, \;\; \beta = 2\alpha,$$

and as α grows, the budget balance ratio is limited by cross-monotonicity at $\frac{3}{2\alpha+3}$, which is $O(\frac{1}{\beta})$.

VI. Optimal Budget Balance in Specific Networks

Results from Sec. IV and Sec. V are about guarantees on the budget balance ratio that general and efficiently computable multicast schemes can provide, while maintaining cross-monotonicity and routing optimality. From a different perspective, it also natural to ask: what is the absolute maximum budget balance ratio that can be achieved in a specific network? What is the range of such ratios for real networks? How do they compare with the theoretically established bounds? Towards this direction, we show that in a specific network instance, the problem of pursuing optimal budget balance while maintaining cross-monotonicity and routing optimality naturally admits a two-stage linear optimization model. Based on this model, we study the optimal ratio in various directed and undirected network instances.

A. A Two-Stage Linear Optimization Model

By definition, pursuing maximum budget balance while maintaining optimality in multicast flow routing and cross-monotonicity in cost sharing can be modelled as a two-stage linear optimization. In stage 1, we compute optimal multicast cost $|f^*_A|$ for each $A \subseteq T$, by solving LP (3.1). In stage 2, we solve the following linear program:

Maximize $\quad x$

Subject to:

$$\begin{cases} x|f^*_A| \le \sum_{u\in A} y_A(u) \le |f^*_A| & \forall A \subseteq T \quad (6.1) \\ y_A(u) \le y_B(u) & \forall u \in B \subset A \subseteq T \quad (6.2) \end{cases}$$

$$x, y_A(u) \ge 0 \qquad \forall u \in A \subseteq T$$

Here the objective function is simply x, the budget balance ratio of (f^*, y). Constraint (6.1) requires the solution in y be x-budget-balanced. Constraint (6.2) establishes cross-monotonicity of y over different service sets A. (6.1) and (6.2) together also imply that y is in-core. The numerical results below are obtained by solving the two-stage linear optimization using the interior-point algorithm as implemented in glpk 4.13 [33].

B. x^* for Network Patterns in Sec. IV and Sec. V

The networks in Fig. 6(a) and Fig. 6(b) are isomorphic to network patterns in Fig. 2 and Fig. 3, respectively, with $l = h = 2$. The maximum budget balance ratio x^* found by solving the two-stage model is 0.7537 in the directed case, and is 0.9 in the undirected case. This is in line with our previous results showing that better budget balance can be achieved in undirected networks than in arbitrary directed networks, in similar network topologies.

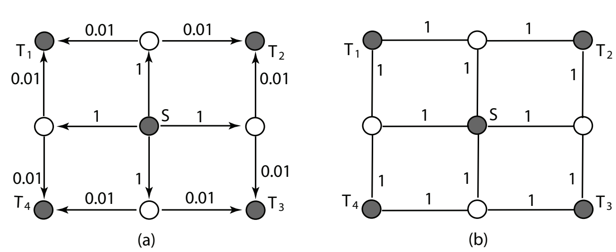

Fig. 6. (a) A directed network where $x^* = 0.7537$. (b) An undirected network where $x^* = 0.9$.

The table below provides details in the solution (f^*, y^*), using the directed case as an example. As we can see, the bottleneck sets that limits x^* from going larger are the ones with cardinality three, each corresponding to the randomly built set A in the proof of Theorem 2. In y^*, each user in A assumes full cost sharing in some gadget and pays only 0.51, while the total cost in A is 2.03 — exactly in line with the probabilistic argument used in Theorem 2.

A (indices of T_i)	$\|f^*_A\|$	$y^*_A(T_i)$	x
1234	2.04	0.51	1.0
123..234	2.03	0.51	0.7537
13, 24	2.02	0.7612	0.7537
12, 14, 23, 34	1.02	0.51	1.0
1..4	1.01	1.01	1.0

C. x^* in Random Networks

We have computed the maximum budget balance ratio x^* for various random networks generated using BRITE [34], a network generator developed at Boston University for simulating Internet-like topologies. Beside undirected networks, we also tested *balanced-bidirected networks* – networks where each pair of neighbors are connected with two directed links, with identical capacity and cost in inverse directions. In the table below, case (A) is for undirected networks with $|V| = 50$, and each entry of x^* is the average value from 5 random networks. Cases (B) and (C) are for undirected networks with $|V| = 30$; the average x^* values from 5 networks are given in (B) and the minimum are given in (C). Cases (D) and (E) are for balanced bidirected networks with $|V| = 30$; (D) lists average values and (E) lists minimum. Bandwidth is sufficiently supplied in each network.

	k=3	5	6	7	8	9	10
(A)	1.0	0.730	0.664	0.642	0.638	0.625	0.601
(B)	0.983	0.772	0.785	0.714	0.723	0.696	0.670
(C)	0.953	0.640	0.686	0.583	0.572	0.648	0.554
(D)	1.0	0.902	0.909	0.932	0.829	0.771	0.709
(E)	1.0	0.716	0.745	0.755	0.736	0.445	0.526

As we can see, x^* decreases as the number of potential users increases in all five cases. We always have $0.5 \le x \le 1.0$ in undirected networks; the only entry where x^* drops below 0.5 is in case (E), the bidirected setting. This confirms the almost tight bound on x^* that we established in Theorem 3 and Theorem 4. x^* is high only when $|T|$ is small; as an extreme case, the reader can verify the fact that perfect budget balance is always feasible for optimal and cross-monotonic multicast to two potential receivers in directed networks. We can also see that, due to the bidirected connection nature, balanced-bidirected networks on average do not suffer from a lower budget balance ratio than undirected networks.

D. Discussions

The multicast schemes we studied in Sec. IV and Sec. V are efficiently computable and are general in that they do not assume any specific information on the network configuration. The two-stage model in this section instead tailors a multicast scheme for each specific input. Such tailored schemes do not have a general characterization, and are extremely expensive to compute — the first stage of the optimization computes exponentially many multicast flows, and the second stage solves an LP with exponentially many constraints. Therefore schemes in Sec. IV and Sec. V are more feasible solutions in practice, except for small multicast groups (*e.g.*, $|T| \leq 15$).

VII. Conclusion

We studied in this paper multicast routing and cost sharing schemes that are optimal, cross-monotonic and budget balanced. The main motivation of cross-monotonicity is to achieve group-strategyproofness. We show that exact budget balance is infeasible in general given the other two requirements. We provide both positive and negative results for achievable budget balance in both directed and undirected networks. Our results in the undirected case is almost tight, and is also verified by simulation results in random networks. We show that constant cost recovery is infeasible in general directed networks, while the best guarantee we can have in undirected networks is close to $\frac{1}{2}$-budget-balance. Furthermore, when bandwidth provisioning is tight in undirected networks, the achievable budget balance ratio is also sensitive to the path cost balance factor β. We finally use a two-stage linear optimization model to compute the maximum budget balance ratio x^* in various random network instances, and verify that x^* can be close to 1 only for small multicast groups, and approaches the theoretically proven range soon as the user group grows large. We leave it as future research to close the gap on the upper and lower bounds proven for directed networks.

Acknowledgement The author would like to thank the anonymous referees for their comments that helped improve this paper, and the Natural Science and Engineering Research Council of Canada (NSERC) and Alberta Inenuity Fund (AIF) for their financial support to his research.

References

[1] J. Feigenbaum, C. Papadimitriou, and S. Shenker, "Sharing the Cost of Multicast Transmissions," *Journal of Computer and System Sciences*, vol. 63, pp. 21–41, 2001.
[2] S. Chen, O. Günlük, and B. Yener, "The Multicast Packing Problem," *IEEE/ACM Transactions on Networking*, vol. 8, no. 3, pp. 311–318, 2000.
[3] C. Chekuri, J. Chuzhoy, L. Lewin-Eytan, J. Naor, and A. Orda, "Non-Cooperative Multicast and Facility Location Games," in *Proceedings of ACM Conference on Electronic Commerce (EC)*, 2006.
[4] M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh, "SplitStream: High-Bandwidth Multicast in Cooperative Environments," in *Proceedings of ACM SOSP*, October 2003.
[5] Z. Li, B. Li, and L. C. Lau, "On Achieving Optimal Multicast Throughput in Undirected Networks," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2467–2485, June 2006.
[6] K. Jain, M. Mahdian, and M. R. Salavatipour, "Packing Steiner Trees," in *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2003.
[7] M. Thimm, "On The Approximability Of The Steiner Tree Problem," in *Mathematical Foundations of Computer Science 2001, Springer LNCS 2136, 678-689*, 2001.
[8] R. Ahlswede, N. Cai, S. R. Li, and R. W. Yeung, "Network Information Flow," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1204–1216, July 2000.
[9] R. Koetter and M. Médard, "An Algebraic Approach to Network Coding," *IEEE/ACM Transactions on Networking*, vol. 11, no. 5, pp. 782–795, October 2003.
[10] Z. Li, B. Li, D. Jiang, and L. C. Lau, "On Achieving Optimal Throughput with Network Coding," in *Proceedings of IEEE INFOCOM*, 2005.
[11] Z. Li and B. Li, "Efficient and Distributed Computation of Maximum Multicast Rates," in *Proceedings of IEEE INFOCOM*, 2005.
[12] D. S. Lun, N. Ratnakar, R. Koetter, M. Médard, E. Ahmed, and H. Lee, "Achieving Minimum-cost Multicast: a Decentralized Approach Based on Network Coding," in *Proceedings of IEEE INFOCOM*, 2005.
[13] Y. Wu, P. A. Chou, Q. Zhang, K. Jain, W. Zhu, and S. Y. Kung, "Network Planning in Wireless Ad Hoc Networks: A Cross-layer Approach," *Journal on Selected Areas in Communications (JSAC)*, vol. 23, no. 1, pp. 136–150, January 2005.
[14] J. Yuan, Z. Li, W. Yu, and B. Li, "A Cross-Layer Optimization Framework for Multihop Multicast in Wireless Mesh Networks," *Journal on Selected Areas in Communications (JSAC)*, vol. 24, no. 11, pp. 2092–2103, November 2006.
[15] C. Papadimitriou, "Algorithms, Games, and the Internet," in *Proc. of the 33rd ACM Symposium on the Theory of Computing (STOC)*, 2001.
[16] W. Wang, X.-Y. Li, and Y. Wang, "Truthful Multicast in Selfish Wireless Networks," in *Proceedings of ACM MobiCom*, 2004.
[17] W. Wang, X.-Y. Li, and Z. Sun, "Sharing the Multicast Payment Fairly," in *Proceedings of the 11th International Computing and Combinatorics Conference (COCOON)*, 2005.
[18] Z. Li, "Min-Cost Multicast of Selfish Information Flows," in *Proceedings of IEEE INFOCOM*, 2007.
[19] H. Moulin and S. Shenker, "Strategyproof Sharing of Submodular Costs: Budget Balance versus Efficiency," *Economic Theory*, vol. 18, no. 3, pp. 511–533, 2001.
[20] M. Pál and E. Tardos, "Group Strategyproof Mechanisms via Primal-Dual Algorithms," in *Proc. of the 44th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2003.
[21] H. Moulin, *Game Theory for the Social Sciences*. New York: New York University Press, 1982.
[22] K. Jain and V. V. Vazirani, "Applications of Approximation Algorithms to Cooperative Games," in *Proc. of the 33rd ACM Symposium on the Theory of Computing (STOC)*, 2001.
[23] H. Moulin, *Cooperative Microeconomics: A Game Theoretic Introduction*. Princeton, NJ: Princeton University Press, 1982.
[24] N. Immorlica, M. Mahdian, and V. S. Mirrokni, "Limitations of Cross-Monotonic Cost Sharing Schemes," in *Proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2005.
[25] J. Könemann, S. Leonardi, and G. Schäfer, "A Group-Strategyproof Mechanism for Steiner Forests," in *Proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2005.
[26] P. Erdös and A. Rényi, "On the Existence of a Factor of Degree One of a Connected Random Graph," *Acta Mathematica Hungarica*, vol. 17, no. 3-4, pp. 359–368, 1966.
[27] L. S. Shapley, "A Value for n-Person Games," *Contributions to the Theory of Games*, pp. 31–40, 1953, Princeton University Press.
[28] D. B. West, *Introduction To Graph Theory, 2nd Edition*. Upper Saddle River, New Jersey: Prentice Hall, 2001.
[29] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Upper Saddle River, New Jersey: Prentice Hall, 1993.
[30] D. P. Bertsekas, *Network Optimization: Continuous and Discrete Models*. Athena Scientific, 1998.
[31] Z. Li and B. Li, "Network Coding in Undirected Networks," in *Proc. of the 38th Annual Conference on Information Sciences and Systems (CISS)*, 2004.
[32] A. Agarwal and M. Charikar, "On the Advantage of Network Coding for Improving Network Throughput," in *Proceedings of the IEEE Information Theory Workshop*, 2004.
[33] *GNU Linear Programming Kit*, http://www.gnu.org/software/glpk/glpk.html.
[34] A. Medina, A. Lakhina, I. Matta, and J. Byers, *BRITE: Boston University Representative Internet Topology Generator*, http://www.cs.bu.edu/brite.