Enforcing Minimum-Cost Multicast Routing against Selfish Information Flows

Zongpeng Li, *Member*, *IEEE*, and Carey Williamson, *Member*, *IEEE*

Abstract—We study multicast in a noncooperative environment where information flows selfishly route themselves through the cheapest paths available. The main challenge is to enforce such selfish multicast flows to stabilize at a socially optimal operating point incurring minimum total edge cost, through appropriate cost allocation and other economic measures, with replicable and encodable properties of information flows considered. We show that known cost allocation schemes are not sufficient. We provide a shadow-price-based cost allocation for networks without capacity limits and show that it enforces minimum-cost multicast. This improves previous result where a 2-approximate multicast flow is enforced. For capacitated networks, computing cost allocation by ignoring edge capacities will not yield correct results. We show that an edge tax scheme can be combined with a cost allocation to strictly enforce optimal multicast flows in this more realistic case. If taxes are not desirable, they can be returned to flows while maintaining weak enforcement of the optimal flow. We relate the taxes to VCG payment schemes and discuss an efficient primal-dual algorithm that simultaneously computes the taxes, the cost allocation, and the optimal multicast flow, with potential of fully distributed implementations.

Index Terms—Communication/networking, multicast, graph algorithms.

1 INTRODUCTION

The classic min-cost flow problem [1], [2] studies the min-cost transmission of commodity flows across a flow network, where each unit edge capacity utilization incurs a cost, and the goal is to minimize the total edge cost while sustaining a target end-to-end flow rate. We consider in this paper the min-cost transmission of information flows in a data network and focus on multicast transmissions where common data is disseminated from a source to multiple destinations. Multicast models real-world applications such as media streaming or the dissemination of popular files. The cost of a communication link is an abstraction of real-world cost, including delay latency, power consumption, and monetary charges.

While both commodity flows and information flows need to confine to the network topology and respect link capacity limits, information flows are unique in that they are *replicable* and *encodable*. Replication and encoding are in general necessary in achieving the full capacity of a data network, and such coding operations applied at potentially any node in the network are referred to as *network coding* [3], [4]. Traditional models of multicast are usually based on Steiner trees, in which either maximizing multicast rate or minimizing multicast cost is computationally intractable [5], [6]. Recent research in network coding reveals a dramatically different structure of multicast: in a directed network, a multicast rate d is feasible if and only if (iff) it is feasible as an independent unicast from the sender to each receiver [3], [4].

- *The authors are with the Department of Computer Science, University of Calgary, 2500 University Drive NW, Calgary, AB T2N 1N4, Canada. E-mail: zongpeng@ucalgary.ca, carey@cpsc.ucalgary.ca.*

Manuscript received 13 Feb. 2008; accepted 26 Sept. 2008; published online 9 Oct. 2008.
Recommended for acceptance by C. Qiao.
For information on obtaining reprints of this article, please send e-mail to: tpds@computer.org, and reference IEEECS Log Number TPDS-2008-02-0058.
Digital Object Identifier no. 10.1109/TPDS.2008.229.

Based on the above result, efficient optimization algorithms for both multicast rate and multicast cost have been successfully designed in the cooperative environment, by exploiting the underlying network flow structure [7], [8]. We consider in this paper a noncooperative environment where flows selfishly minimize their own cost by routing themselves through the cheapest available paths. Such selfish behavior is a well-studied phenomenon in game theory, and the stable state of such selfish routing games is characterized as a *Nash Equilibrium*, where no flow has incentive to deviate from its current route alone, assuming the rest of the flows stick to their routes. It is known that such Nash solutions can lead to very bad performances [9], [10] in general. Therefore, it is necessary to impose calculated economic regulations in the network such that individual decisions of selfish flows jointly lead to a socially desirable operation state of the network.

Previous research has considered the enforcement of optimal multicommodity flows [11], [12], [13]. While they also employ economic measures to regulate the routing of flows among potential paths, it is unique and vital in the multicast problem to encourage flow cost sharing. It is well known that cost sharing can be achieved through combining individual unicast paths into multicast trees. However, two problems are evident along the traditional multicast tree direction. First, the cost may not be really minimum, since a different multicast flow with lower cost may be found by also exploiting the encodable property of information flows [7], [6] (an example is shown in Fig. 1). Second, since optimal routing based on multicast trees is NP-hard to compute, it is unlikely that it would be *exactly* enforced by any efficiently computable regulation scheme in general network topologies [14]. The best result prior to this work is a bicriteria analysis by Bhadra et al. [15], which shows that there always exist cost sharing schemes to enforce a 2-approximate multicast flow.

1045-9219/09/$25.00 © 2009 IEEE Published by the IEEE Computer Society

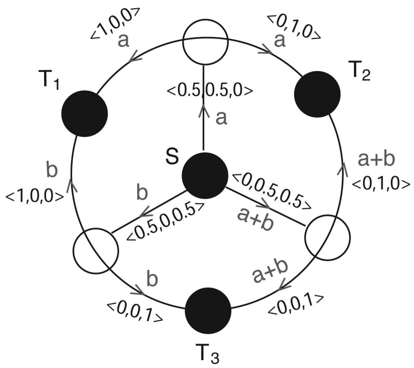

Fig. 1. An example of enforcing min-cost multicast flows. Each link has cost 1 and is labeled with the information flow it carries and its cost share vector $<y_1, y_2, y_3>$. Each edge flow has rate 0.5. The total cost is $0.5 \times 9 = 4.5$ and is optimal. A minimum multicast tree has cost 5.

In this paper, we study the enforcement of optimal multicast flows in general network topologies, through proper edge cost sharing and other economic measures. We first consider the simple case where edges do not have capacity limits and show that the well-known Shapley value [16] method cannot enforce optimal multicast flows. Instead, we formulate the min-cost multicast problem into a pair of primal and dual linear programs, based on the aforementioned multicast feasibility result with network coding due to Ahlswede et al. [3], and propose to allocate edge costs based on the shadow prices of flow merging constraints. We show that a Nash Equilibrium exists and that any optimal multicast flow has a corresponding cost allocation which makes it a Nash flow. The flow-cost pair at Nash Equilibrium also achieves a balanced budget, i.e., the total charges to flows exactly equal the total cost incurred at edges across the network.

For the more realistic case where edges have finite capacity limits, we show that ignoring edge capacities and taking the solution above will not enforce the optimal multicast flow. We instead propose to further establish a nonnegative tax at each edge. We show that each optimal multicast flow can be *strictly* enforced by a pair of edge tax and cost allocation vectors such that the solution remains stable even if the capacity limits on a subset of edges are relaxed. In cases where charging edge taxes is infeasible or undesirable, we prove that there always exists a tax return procedure, after which flows pay true edge costs only and the optimal multicast flow is still Nash. However, the optimal solution is now *weakly* enforced and is sensitive to edge capacity fluctuations. We discuss connections of the edge taxes to the "added value" concept in the Vickrey-Clarke-Groves (VCG) edge payment scheme and to strategyproof multicast [17].

The existential proofs mentioned above rely on path-based linear programming models of the min-cost multicast problem. These LPs, while convenient for analysis purposes, are not suitable for computing the solutions in practice, since they in general contain exponentially many variables/constraints. We finally reformulate the min-cost multicast LPs to reduce their sizes and design an efficient primal-dual algorithm that simultaneously computes the edge taxes, the cost shares, and the optimal multicast flow they enforce. This is achieved by using Lagrange relaxation to remove coupling constraints among flows to distinct receivers and decomposing the entire optimization into a series of shortest path computations through the subgradient method [2]. Our algorithm outperforms general linear programming solution methods in runtime and allows distributed implementations.

The rest of the paper is organized as follows: We review previous research in Section 2 and discuss the network model together with preliminaries in linear programming in Section 3. We study uncapacitated networks in Section 4, capacitated networks in Section 5, generalize the results in Section 6, and design solution algorithms in Section 7. The paper is concluded in Section 8.

2 Previous Research

The seminal work of Ahlswede et al. [3] initiated the research on network coding and showed its necessity in achieving maximum network capacity. An important result they proved is that *in a directed network, a multicast rate d is feasible iff it is feasible to each receiver independently*. Koetter and Médard also derived this result in an algebraic framework [4]. This new characterization of multicast rate feasibility (with network coding) dramatically changed the underlying structure of the multicast problem, namely, from multicast trees to a union of conceptual network flows [18]. Consequently, breakthroughs were made in efficient multicast algorithm design in directed networks [7], undirected networks [18], [8], and wireless ad hoc networks [19], [20], assuming a cooperative environment. In this paper, we instead study how to achieve min-cost multicast when information flows are selfish.

Edge pricing schemes that enforce minimum-delay multicommodity flows are studied in [13] and [12]. Both compute taxes/tolls on edges to guide the selfish routing process. Reference [13] computes the taxes by solving a nonlinear complementary problem, while [12] takes the linear programming approach. Our work in this paper was partly inspired by them and is similar in that edge taxes are also considered as part of the flow regulation measures. The edge taxes we introduce allow more intuitive interpretations and can be eventually returned to the multicast flows without jeopardizing their stability. Another important difference is that we need to also design appropriate cost sharing schemes that induce the desired path sharing among multicast flows.

Feigenbaum et al. studied the sharing of multicast cost in the context of designing strategyproof multicast for selfish receivers with private utility information [14]. Along the multicast tree direction, they show that optimal welfare is NP-hard to approximate within any constant ratio in general networks. They consider acyclic network topologies instead, where the multicast cost is submodular, and both the Shapley value [16] and marginal cost sharing lead to strategyproofness. Lower bounds on communication overhead are derived for these two sharing schemes. In this paper, we perform multicast cost sharing with network coding explicitly considered. As a result, optimal sharing schemes can be computed efficiently in general network topologies.

Bhadra et al. also studied the multicast of selfish information flows in [15]. They use monomial cost functions to approximate edge capacity limits and differentiable relaxations to approximate the $\max$ function in edge flow computation. After these approximations, the min-cost multicast problem is modeled as a nonlinear optimization problem with a differentiable objective function and constraints, and cost shares for enforcing the optimal multicast flow are derived based on Karush-Kuhn-Tucker (KKT) optimality conditions for nonlinear programs [21]. This result applies to power-law edge cost functions only. For the linear cost model and other convex cost functions, they employ a bicriteria approach introduced by Roughgarden and Tardos [10] and design cost shares that enforce a *suboptimal* multicast flow, which has a cost lower than any optimal multicast flow achieving twice the throughput. In this paper, we use exact models for edge capacity limits and edge flow computation and devise economic mechanisms that enforce optimal multicast flows.

Wang et al. [17] studied min-cost multicast in networks with selfish edges, where the edge cost is private information, and each edge reports a cost value at its own choice to maximize its utility. They show that the celebrated VCG payment scheme fails to induce strategyproofness if optimal multicast routing, NP-hard without network coding, is approximated with schemes such as pruning min-spanning tree or link weighted Steiner tree. We establish underlying connections between edge taxes introduced in this paper and the "added value" concept in VCG payments. We argue that paying each edge its declared cost plus such edge taxes according to VCG successfully induces strategyproofness at selfish edges. Therefore, the failure of VCG identified in [17] is due to the fact that approximate multicast algorithms are employed rather than being inherent in the multicast problem itself. Wang et al. [22] further presented a general framework for deciding whether an existing multicast protocol can be transformed into a truthful one and, if possible, how the payments to relay agents should be fairly shared among the receivers.

In the routing and cost sharing of multicast toward a group of potential receivers, each with private utility information for being serviced in the multicast group, the key toward group strategyproofness is to have a cross-monotonic cost sharing scheme [23], [24]. Informally, within cross-monotonic cost sharing, a user's payment can only be smaller when serviced in a larger set. Li [25] studied multicast schemes that target optimal flow routing, cross-monotonic cost sharing, and budget balance. It was shown that these three conditions cannot be satisfied simultaneously in general. Complementing positive and negative results are given for both directed and undirected networks on the achievable approximate budget balance ratio for an optimal and cross-monotonic multicast scheme. The above research is complementary to this work, in that the former studies incentives for strategic receivers to cooperate and the latter studies incentives for selfish traffic to exhibit a socially desirable behavior.

3 Network Model

The data network is modeled as a directed graph $G = (V, E)$. Let $S \in V$ be the multicast sender and $\mathcal{T} = \{T_1, \dots, T_k\} \subseteq V$ be the set of receivers. Here, k is the number of multicast receivers. We use vectors $c, w \in Q_+^E$ to store the capacity limit and cost of edges, respectively. Here, Q_+ is the set of nonnegative rational numbers. Note that $w(e)$ denotes the cost of a unit flow on e, and the total cost of flow $f(e)$ on e is $w(e)f(e)$. Let $\mathcal{P}_i$ be the set of distinct paths available from S to T_i. For $p \in \mathcal{P}_i$, $f(p)$ represents the amount of $S \to T_i$ flow carried by p. The total flow amount on an edge e is denoted as $f(e)$. A scalar d stores the target multicast rate. In this noncooperative multicast routing game, each agent corresponds to an infinitesimally small amount of some $S \to T_i$ flow, the change of whose routing decision alone does not lead to observable variations in any edge state. This models the situation where each data packet makes its own routing decision.

Fig. 1 illustrates some of the concepts we discussed using a multicast network with one sender and three receivers. The target multicast rate is 1, and each edge has cost 1 per unit flow. We ship on each edge an information flow of rate 0.5, of either a, b, or $a + b$ as labeled. Here, $+$ is performed over a finite field and is equivalent to bitwise exclusive or. Each receiver may successfully recover the two original flows a and b from what it receives. A cost allocation is given in the figure, under which the multicast flow shown is Nash. Here, $y_i(e)$ is the cost of a unit flow to T_i on e. For example, the $S \to T_1$ flow has a total rate of 1, with two paths each carrying 0.5. Its cost on each path is $0.5 \times 0.5 + 1 \times 0.5 = 0.75$, and the total $S \to T_1$ flow cost is 1.5. From S to each T_i, there are two different paths with an identical cost of 1.5. Flows do not have incentive to switch between paths, and the routing and cost allocation together constitute a Nash Equilibrium. Constructing a detailed coding scheme upon a given multicast flow is known as *code construction* [26] and is orthogonal to our current topic. In the rest of the paper, we assume that code construction can be efficiently done and omit its further details.

The analysis and algorithm design presented in this paper depend heavily on linear programming duality results, including primal and dual LPs and their correspondence, complementary slackness conditions, and Lagrange duality. These results are discussed in many linear programming textbooks, e.g., [27]. We provide here only a sketched description due to space limitations. Every (primal) LP has an equivalent dual LP. Each variable in the primal maps to a constraint in the dual and vice versa. Equality constraints map to free variables, and inequality constraints map to nonnegative variables. A variable-constraint pair is *complementary* if either the variable is zero or the constraint is tight, i.e., satisfied at equality. If both primal and dual LPs are feasible, then they have pairwise complementary optimal solutions. Conversely, if a pair of primal and dual solutions is feasible and complementary, then they are both optimal. A *Lagrange relaxation* of an LP removes a group of its constraints (let $\vec{x}$ be the variable vector they map to) while adding a corresponding penalty term into the objective function. Viewing the resulting (smaller) LP as a function of $\vec{x}$ and optimizing it over the feasible domain of $\vec{x}$ yields the same optimal solutions (if existent) as in the original LP.

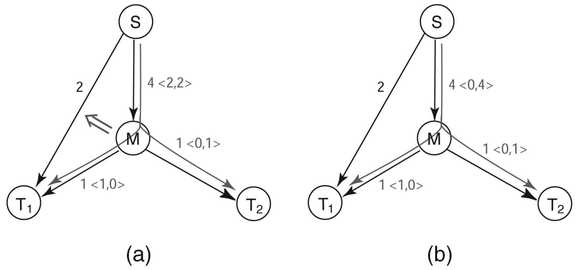

Fig. 2. An example multicast network with its optimal multicast flow. (a) The optimal flow with Shapley value cost sharing is not Nash. (b) The cost allocation under which the optimal flow is Nash. Each nonzero edge flow has rate d. Each edge is labeled with its cost and sharing $w(e) < y_1(e), y_2(e) >$. The double arrow indicates the direction of potential path switches.

In the rest of the paper, we use interchangeably the terms *cost allocation* and *cost sharing* and *edge* and *link*. A multicast flow at Nash Equilibrium is called a *Nash flow*. A network is referred to as an *uncapacitated network* if it has unlimited edge capacities or a *capacitated network* otherwise.

4 Uncapacitated Networks

We begin our study with the simpler case where edges in the network G are assumed to always have sufficient bandwidth supply for the purpose of multicast routing. This allows us to focus on mechanisms that foster appropriate path sharing among flows. We handle realistic network scenarios with finite edge capacity limits in the subsequent section.

4.1 Cost Sharing Based on Shapley Value

The Shapley value is a well-known "fair" value allocation method proposed by Shapley [16]. It was introduced originally in the context of value sharing in coalitional games [28], where each player receives a share depending on how "important" he or she is in the coalitions. It is the only scheme that simultaneously achieves *budget balance*, *anonymity* (player labeling is irrelevant), and *additivity* (a player's shares in two coalitional games sum up to its share in the combination of these two games). In multicast tree cost sharing, it is known that according to the Shapley value, each edge is paid for equally by all its downstream receivers [14].

For cost sharing of multicast with network coding, an accurate definition of the Shapley value can be formulated as

$$y_i(e) = \frac{w(e)}{f_i(e)} \sum_{T_i \notin \mathcal{H} \subseteq \mathcal{T}} \frac{h!(k-h-1)!}{k!} \left(\max_{j \in \mathcal{H} \cup i} f_j(e) - \max_{j \in \mathcal{H}} f_j(e) \right),$$

where $y_i(e)$ is the cost share of a unit $S \to T_i$ flow on e, and $f_j(e) = \sum_{e \in p \in \mathcal{P}_j} f(p)$ is the total amount of a $S \to T_j$ flow on e. $\mathcal{H}$ is the variable set to which we append T_i and compute the additional cost, and h is the cardinality of $\mathcal{H}$. Intuitively, the Shapley value method enumerates all possible orders that flows to different receivers may occur on an edge e, calculates the added cost of each flow upon its appearance, and takes the average as the final share. It satisfies $\sum_i y_i(e) f_i(e) = w(e) f(e)$, i.e., the total flow cost on e is shared exactly, achieving a balanced budget.

TABLE 1
Path-Based Min-Cost Multicast LPs: Primal and Dual

Minimize $\sum_e w(e) f(e)$

Subject to:

$$\begin{cases} \sum_{p \in \mathcal{P}_i} f(p) = d & \forall i \quad \leftrightarrow x_i \\ f(e) \geq \sum_{p \in \mathcal{P}_i : e \in p} f(p) & \forall i, \forall e \quad \leftrightarrow y_i(e) \end{cases}$$

$f(e), f(p) \geq 0 \quad \forall e, \forall p$

Maximize $\sum_i x_i d$

Subject to:

$$\begin{cases} \sum_i y_i(e) \leq w(e) & \forall e \quad \leftrightarrow f(e) \\ x_i \leq \sum_{e \in p} y_i(e) & \forall i, \forall p \in \mathcal{P}_i \quad \leftrightarrow f(p) \end{cases}$$

x_i free; $y_i(e) \geq 0 \quad \forall i, \forall e$

Unfortunately, the Shapley-value-based cost sharing does not successfully enforce min-cost multicast flows. As an example, the multicast session shown in Fig. 2 has a min-cost of 6 and is not Nash under Shapley cost sharing because for $S \to M \to T_1$ flows, there is a cheaper direct path $S \to T_1$ to switch onto. A Nash flow after such switches has cost 7 and is not optimal.

The reason why the Shapley value method is no longer appropriate lies in the fact that it is inherently a local cost sharing method. While it may work fine for a single edge or for a simple tree topology, global attention needs to be paid for multicast flows in general network topologies, to handle the competition among potential flow paths. We now proceed to derive cost shares within the linear programming framework and show that it is always possible to use shadow-price-based cost shares to enforce optimal multicast flows in any network topology.

4.2 LP Formulations

We first formulate the min-cost multicast problem into a pair of path-based LPs in Table 1.

The formulation of the above LPs are based on the fact that a multicast rate d can be achieved by setting up a network flow of rate d from sender S to each receiver T_i; these network flows are *conceptual* in that they share instead of compete for available bandwidth on an edge [18]. In other words, if the $S \to T_i$ and $S \to T_j$ network flows have rates $f_i(e)$ and $f_j(e)$ on an edge e, respectively, then their combined effective flow rate is $\max(f_i(e), f_j(e))$ instead of $f_i(e) + f_j(e)$. This is possible due to the unique encodable and replicable properties of information flows.

4.3 LP Interpretations

In the primal LP, the first constraint requires the total network flow rate to each receiver be the desired multicast throughput d, the second constraint requires the total edge flow $f(e)$ to be no less than each conceptual flow rate $f_i(e)$. Note that the smallest $f(e)$ satisfying this constraint is $f(e) = \max_i f_i(e)$. The objective is to minimize the total multicast flow cost $\sum_e w(e) f(e)$. Following each constraint, we also list its corresponding dual variable. In the dual LP, the first constraint allocates the total cost of an edge $w(e)$ to the k conceptual flows. The second constraint implies that the cost x_i from S to T_i should be no more than the total cost of any path p from S to T_i. Note that the largest value x_i satisfying this constraint is $x_i = \min_{p \in \mathcal{P}_i} \sum_{e \in p} y_i(e)$. The dual objective is to maximize $\sum_i x_i d$.

The interpretation of the primal LP is relatively straightforward: the multicast service provider strives to route a multicast flow of rate d toward each of his or her customer (primal constraints), while attempting to minimize the total edge flow cost incurred (primal objective). The dual LP can be interpreted as a revenue-maximization problem faced by the network operator instead. The network operator wishes to allocate each link cost ($w(e)$) among the receivers ($y_i(e)$) such that when flows are selfish and pick their respective cheapest paths (path cost: x_i), the total revenue collected from flow charges across the network is maximized (dual objective).

4.4 Cost Sharing Based on Shadow Prices

We say that a cost allocation y *enforces* a multicast flow f if f and y together satisfy 1) *stability*, i.e., f is a Nash flow under y, 2) *budget balance*, i.e., $\forall e$, $w(e)f(e) = \sum_i y_i(e) f_i(e)$, and 3) *fairness*, i.e., no flow has a cost share $y_i(e)$ of more than the total edge cost $w(e)$ at any edge e. Requirements 2) and 3) preclude trivial solutions such as assigning a cost $y_i(e) = a$ wherever $f_i(e) > 0$ and $y_i(e) = b$ otherwise for some constants $a < b$.

We now prove that the optimal solutions in the dual LP, also known as *shadow prices*, may serve as cost allocations that enforce optimal multicast flows.

Theorem 1. *If f^* and (x^*, y^*) are a pair of optimal solutions to the primal and dual min-cost multicast LPs, respectively, then y^* enforces f^*.*

Proof. We first show stability. By linear programming duality, f^* and (x^*, y^*) together satisfy the complementary slackness conditions. By dual complementary slackness, we have

$$p \in \mathcal{P}_i, f^*(p) > 0 \longrightarrow x_i^* = \sum_{e \in p} y_i^*(e) \quad \forall i.$$

Since x^* maximizes $\sum_i x_i^* d$, we can conclude that 1) for any receiver T_i, $x_i^* = \min_{p \in \mathcal{P}_i} \sum_{e \in p} y_i^*(e)$ and 2) if p_1, $p_2 \in \mathcal{P}_i$, and $f^*(p_1) > 0$, then $\sum_{e \in p_1} y_i^*(e) \le \sum_{e \in p_2} y_i^*(e)$. In other words, if we use y^* as the cost allocation, then every path p with a nonzero flow is a shortest path, and flows on that path have no incentive of switching to a different path, which makes f^* a Nash flow.

We next show budget balance. By primal complementary slackness, $y_i^*(e) > 0 \longrightarrow f_i^*(e) = f^*(e)$, $\forall i \forall e$, i.e., a $S \rightarrow T_i$ flow has a positive cost on e only if $f_i^*(e) = \max_j f_j^*(e)$. Further, dual complementary slackness implies

$$f^*(e) > 0 \longrightarrow \sum_i y_i^*(e) = w(e) \quad \forall e.$$

Let $\mathcal{I}_e = \{i | f_i^*(e) > 0\}$; then, we have $\sum_i y_i^*(e) f_i^*(e) = \sum_{i \in \mathcal{I}_e} y_i^*(e) f_i^*(e) = \sum_{i \in \mathcal{I}_e} y_i^*(e) f^*(e) = w(e) f(e)$. We finally show fairness: dual feasibility requires that $\forall e$, $\sum_i y_i^*(e) \le w(e)$; since entries in y^* are nonnegative, we have $\forall i \forall e$, $y_i^*(e) \le w(e)$. □

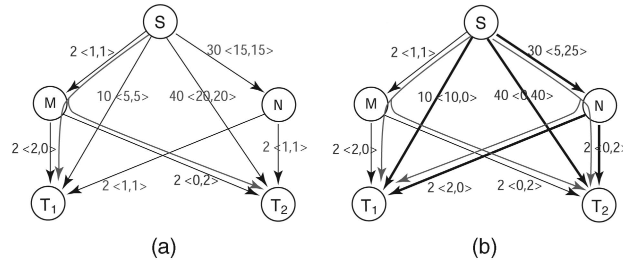

Fig. 3. The same cost allocation y that enforces an optimal multicast flow in the uncapacitated network may fail when edges have limited capacities. (a) y and its enforced optimal multicast flow f when the network is uncapacitated. (b) The capacitated case. Multicast rate $d = 2$. Thin edges have capacity 1, and bold edges have 3. Each nonzero edge flow is of rate 1. The new cost allocation y' that enforces the optimal flow f' is shown. f' is not enforced by y.

Corollary 1. *Every optimal multicast flow is enforceable.*

Proof. If a primal LP is feasible and has an optimal solution, then so does its dual [27]. Therefore, every optimal primal solution f^* of the min-cost multicast LP has a corresponding optimal dual solution (x^*, y^*). This combined with Theorem 1 implies the corollary. □

4.5 Discussions

We proved that every optimal multicast flow f^* can be enforced by a shadow price vector y^*. However, we point out that there may exist under y^* a different Nash flow f' that is not optimal. An example can be derived in Fig. 1 by rearranging flow rates f_1 on the two different $S \rightarrow T_1$ paths from (0.5, 0.5) to, say, (0.2, 0.8). Therefore, the shadow price method guarantees the *price of stability* (ratio of the *best* Nash flow cost to the optimal flow cost) [29] to be 1, but the *price of anarchy* (ratio of the *worst* Nash flow cost to the optimal flow cost) [30] may be higher than 1. We leave it as future research to decide whether there exists a cost allocation scheme that guarantees a price of anarchy of 1 in general.

5 Capacitated Networks

In practice, network links have bounded bandwidth, constituting a new dimension of constraint in multicast routing. We now study how to regulate the selfish multicast routing game in this more realistic scenario. Before presenting the solution, we note that simply ignoring edge capacity limits and applying the shadow-price-based cost allocation described in Section 4 will not yield correct results.

Fig. 3 shows a counter example: Fig. 3a depicts an uncapacitated network, and Fig. 3b depicts a capacitated one, over the same topology. An optimal multicast flow f and its enforcing cost allocation vector y are shown in Fig. 3a. The only optimal multicast flow f' with the specified edge capacity configuration is given in Fig. 3b. We can see that f' is not a Nash flow under y because flows on $S \rightarrow N \rightarrow T_1$ have a better path $S \rightarrow T_1$ to switch onto. The problem is that in the uncapacitated case, the cost allocation y may focus only on the atomic multicast topology with the minimum cost and provide little

TABLE 2
Capacitated Min-Cost Multicast LPs: Primal and Dual

$$\text{Minimize} \quad \sum_e w(e)f(e)$$
Subject to:
$$\begin{cases} \sum_{p\in\mathcal{P}_i} f(p) = d & \forall i & \leftrightarrow x_i \\ f(e) \geq \sum_{p\in\mathcal{P}_i: e\in p} f(p) & \forall i, \forall e & \leftrightarrow y_i(e) \\ f(e) \leq c(e) & \forall e & \leftrightarrow t(e) \end{cases}$$
$$f(e), f(p) \geq 0 \quad \forall e, \forall p$$

$$\text{Maximize} \quad \sum_i x_i d - \sum_e c(e)t(e)$$
Subject to:
$$\begin{cases} \sum_i y_i(e) \leq w(e) + t(e) & \forall e & \leftrightarrow f(e) \\ x_i \leq \sum_{e\in p} y_i(e) & \forall i, \forall p \in \mathcal{P}_i & \leftrightarrow f(p) \end{cases}$$
$$x_i \text{ free; } y_i(e), t(e) \geq 0 \quad \forall i, \forall e$$

guarantee on the rest of the edges. We now proceed to remodel the min-cost multicast problem with finite edge capacities introduced into the picture.

5.1 LP Formulations

The min-cost multicast problem where each edge has a finite constant bandwidth limit can be formulated into a pair of primal and dual LPs, as shown in Table 2.

The primal LP has an extra constraint $f(e) \leq c(e)$ than in the uncapacitated case to model finite edge capacity limits. Recall that c is the edge capacity vector. The dual LP has a corresponding extra variable vector t. From an economic point of view, $t(e)$ reflects how critical bandwidth supply is at edge e. By complementary slackness, $t(e)$ is larger than zero only when e is saturated, i.e., when $f(e) = c(e)$.

5.2 LP Interpretations

Similar to the uncapacitated case, the primal LP can still be interpreted as the cost-minimization problem faced by the multicast service provider. The difference is that routing solutions are now subject to capacity constraints on edges. For the dual LP, the network operator still tries to maximize her revenue that can be collected from selfish multicast flows. However, an extra dimension of flexibility is now available for manipulating routing decisions of the flows: the network operator may decide to set a tax on each edge ($t(e)$). Flow taxes collected will be eventually transferred to a third-party tax bureau (as seen in the dual objective). The objective of the network operator is to maximize his or her net revenue without taxes. Note that although the network operator is not able to keep the taxes collected, he or she in general still wishes to strategically set taxes on edges so that flows are steered toward paths that lead to higher overall revenue.

5.3 Cost Allocation with Edge Taxes

We say that an edge tax and cost sharing pair (t, y) *enforces* a multicast flow f if y enforces f with $w' = w + t$ as the new edge cost vector and that (t, y) *strictly enforces* f if (t, y) enforces f with the edge capacity vector c replaced by any new vector $c' \geq c$. Intuitively, a strictly enforced multicast flow remains stable even if some edge capacity limits are relaxed. We now prove that charging each edge a tax according to t may lead to strict enforcement of optimal multicast flows in capacitated networks.

Theorem 2. *If f^* and (x^*, y^*, t^*) are optimal solutions to the primal and dual min-cost multicast LPs, respectively, then (y^*, t^*) strictly enforces f^*.*

Proof. Since f^* and (x^*, y^*, t^*) are corresponding primal and dual optimal solutions, dual complementary slackness requires that

$$p \in \mathcal{P}_i, f^*(p) > 0 \longrightarrow x_i^* = \sum_{e\in p} y_i^*(e) \quad \forall i.$$

The dual LP maximizes $\sum_i x_i d - \sum_e c(e)t(e)$. Given a fixed vector t, it strives to maximize each x_i in order to maximize $\sum_i x_i d$. Therefore, we know that $x_i^* = \min_{p\in\mathcal{P}_i} \sum_{e\in p} y_i^*(e)$, and

$$p_1, p_2 \in \mathcal{P}_i, f^*(p_1) > 0 \longrightarrow \sum_{e\in p_1} y_i^*(e) \leq \sum_{e\in p_2} y_i^*(e) \quad \forall i.$$

Therefore, toward each receiver T_i, every path with nonzero flow has the minimum after-tax cost, and no flow has incentive to deviate from its current path. In other words, the tax t^* and the cost allocation y^* together makes f^* a Nash flow. The proof of budget balance (with $w + t$ as edge cost) and fairness are essentially the same as that in the uncapacitated case and is omitted.

If the capacities of a subset of edges are increased, all flows in f^* are still feasible, and all paths carrying nonzero flows are still shortest after taxes; therefore, f remains stable. We conclude that y^* and t^* strictly enforces f^*. □

Corollary 2. *In a capacitated network, every optimal multicast flow can be strictly enforced by an edge tax scheme combined with a cost allocation scheme.*

Proof. This follows from the result in Theorem 2 and the fact that every primal optimal solution has a corresponding optimal dual solution. □

5.4 Tax Return

Charging extra taxes from flows may not always be feasible or desirable. We now show that it is possible to return taxes to the multicast flows such that the optimal multicast flow is still enforced by the resulting cost allocation. The challenge is to return taxes to flows that paid them only and to ensure that the multicast flow is still Nash and the budget is still balanced after the return. Our overall solution for tax-free cost sharing in capacitated networks is a two-stage one: compute with-tax cost shares first and then return the taxes appropriately.

Theorem 3. *In a capacitated network, every optimal multicast flow f^* can be enforced by a cost allocation scheme y'.*

Proof. The main idea is to return edge taxes to flows, while maintaining the condition that no path flow has a lower cost path to switch onto. This is possible partly due to the fact that edges have finite capacities. In particular, if a conceptual network flow f_i is already using the full capacity on an edge e, then even if the cost $y_i(e)$ is reduced due to a tax return, it is still infeasible for the rest of the f_i flows to switch onto e since there is no residual edge capacity available.

By Theorem 2, we know that every optimal multicast flow f^* can be strictly enforced by edge taxes t^* and cost allocation y^*, which are optimal dual solutions. Now, consider returning edge taxes as follows:

$$y_i'(e) = \frac{w(e)}{w(e)+t(e)} y_i^*(e) \quad \forall e, \forall i.$$

By primal complementary slackness, we have $t^*(e) > 0 \longrightarrow f^*(e) = c(e)$, $\forall e$, and $y_i^*(e) > 0 \longrightarrow \sum_{p \in \mathcal{P}_i : e \in p} f^*(p) = f^*(e)$, $\forall i$, $\forall e$. In other words, every edge e assigned a nonzero tax $t^*(e)$ has zero residual capacity under flow f^*, and every destination i that has a nonzero cost share $y_i^*(e)$ on edge e has the largest flow rate on e. After the tax return, if a nonzero path flow $f^*(p)$, $p \in \mathcal{P}_i$, wishes to switch to another path $p' \in \mathcal{P}_i$, it must be the case that $\exists e \in p'$, $y_i'(e) < y_i^*(e)$, which implies that $y_i^*(e) > 0$ and $t^*(e) > 0$. Then, by the primal complementary slackness conditions presented above, we have $\sum_{p \in \mathcal{P}_i : e \in p} f^*(p) = f^*(e) = c(e)$, i.e., the edge e would have no residual capacity for a $S \rightarrow T_i$ path flow, and a $p \rightarrow p'$ switch is impossible. Therefore, f^* is a Nash flow under y'. Budget balance and fairness follow from the fact that y^* enforces f^* with edge cost vector $w + t$ and that $y' = \frac{w}{w+t} y^*$. We conclude that the cost allocation y' enforces the optimal multicast flow f^*. □

5.5 Discussions

The complementary slackness conditions imply that taxes in t^* make all utilized $S \rightarrow T_i$ paths equally expensive. We plan to show in a forthcoming work that for an infinite set of network scenarios, the tax vector t^* can be efficiently translated to the added value of an edge e, i.e., the increment of total multicast cost that would occur if e is removed from the network; furthermore, if edge cost $w(e)$ is private information known to e only and edges are selfish, then paying each edge $w'(e) + t^*(e)$ ($w'(e)$ is the declared cost of e) according to the VCG scheme makes e have no incentive to lie. It is a dominant strategy for each edge to report its true cost $w(e)$, leading to a *strategyproof* multicast system. In such cases, the algorithm we present in the next section improves the time complexity of payment computation by a factor of $|E|$, compared to a straightforward algorithm using the definition of VCG. For general scenarios, we will prove that $c(e)t^*(e)$ always constitutes a lower bound on the added value of e.

6 Generalization of Results

In this section, we show that the shadow-price-based enforcement schemes presented in Sections 4 and 5 can be further applied to induce optimal solutions in 1) min-cost multicast with *selfish receivers* and 2) min-cost rooted connectivity in undirected networks.

6.1 Multicast toward Selfish Receivers

So far, we have considered infinitesimally small flows as agents. It is equally interesting to consider the case where each receiver T_i is a selfish agent instead. The fundamental difference between the two models is that while the change of decision of one agent does not noticeably affect the global routing scheme in the former, it does so in the latter. In some literature, the stable system state is referred to as a *Wardrop equilibrium* in the former and a *Nash equilibrium* in the latter [31]. It turns out that the shadow-price-based schemes are strong enough for controlling even selfish receivers, who can reroute a large amount of flow in one step of action. Note that the following fact is true about the LP-based routing and pricing scheme.

Fact. *If f^* and y^* are corresponding optimal primal and dual solutions to the multicast LP in Table 1 (Table 2), then for each i, f_i^* constitutes a min-cost flow of rate d from S to T_i under cost vector y_i^*.*

The above fact is implied by complementary slackness conditions used in Sections 4 and 5, which state that every utilized path is equally expensive and every other path can only be more expensive. Together with Theorems 1 and 2, it leads to the following theorem.

Theorem 4. *For min-cost multicast toward a set of selfish receivers T, each of which minimizes its own routing cost, the optimal global routing is always enforceable.*

6.2 Rooted Connectivity in Undirected Networks

In this section, we switch our focus to undirected networks and consider the following edge-connectivity game. In an undirected network $G = (V, E)$, each node in the terminal set $T \subseteq V$ wishes to purchase edge capacities in the network to make itself d-edge-connected to a certain root node S. If every terminal node does so, then d-edge-connectivity is also realized between every pair of terminal nodes. Therefore, a unicast rate of d can be realized over the purchased link capacities between every pair of terminals. Purchased edge capacities are shared among the terminal nodes; hence, similar to the multicast problem, we also need an appropriate cost sharing scheme to foster collaboration among selfish receivers, which leads to the minimum total purchase payment. We make the following critical observation:

Fact. *The undirected rooted-connectivity game has exactly the same primal and dual LP formulations as in Tables 1 and 2, for the uncapacitated and capacitated cases, respectively.*

The above fact together with Theorems 1 and 2 imply the following theorem.

Theorem 5. *Every optimal solution of the undirected rooted connectivity game is enforceable.*

7 Algorithm Design

We have shown that shadow-price-based cost shares and taxes may enforce optimal multicast flows. We now proceed to discuss how these optimal solutions can be efficiently computed. First, note that the number of distinct paths between two nodes in a general graph may be exponential to the graph size. Consequently, the path-based LP has exponentially many variables (primal) or constraints (dual) and is impractical for computing purposes.

In this section, we first present reformulated link-based min-cost multicast LPs with reduced polynomial sizes. We argue that they are equivalent to the path-based LPs, and in particular, a cost allocation y^* or cost-tax vector pair (y^*, t^*)

TABLE 3
Link-Based Min-Cost Multicast LP: Primal and Dual

Minimize $\sum_{\vec{uv}} w(\vec{uv}) f(\vec{uv})$

Subject to:

$$\begin{cases} \sum_{v\in N_\downarrow(u)} f_i(\vec{uv}) = \sum_{v\in N_\uparrow(u)} f_i(\vec{vu}) & \forall i, u & \leftrightarrow p_i(u) \\ f_i(\vec{T_iS}) = d & \forall i & \leftrightarrow x_i \\ f_i(\vec{uv}) \le f(\vec{uv}) & \forall i, \vec{uv} & \leftrightarrow y_i(\vec{uv}) \end{cases}$$

$f_i(\vec{uv}), f(\vec{uv}) \ge 0 \quad \forall i, \forall \vec{uv}$

Maximize $\sum_i x_i d$

Subject to:

$$\begin{cases} p_i(u) - p_i(v) \le y_i(\vec{uv}) & \forall i, \vec{uv} & \leftrightarrow f_i(\vec{uv}) \\ p_i(T_i) - p_i(S) + x_i \le y_i(\vec{T_iS}) & \forall i & \leftrightarrow f_i(\vec{T_iS}) \\ \sum_i y_i(\vec{uv}) \le w(\vec{uv}) & \forall \vec{uv} & \leftrightarrow f(\vec{uv}) \end{cases}$$

$p_i(u), x_i$ free; $y_i(\vec{uv}) \ge 0 \quad \forall i, \forall u, \forall \vec{uv}$

TABLE 4
Link-Based Min-Cost Multicast LP (Capacitated)

Minimize $\sum_{\vec{uv}} w(\vec{uv}) f(\vec{uv})$

Subject to:

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

$f_i(\vec{uv}), f(\vec{uv}) \ge 0 \quad \forall i, \forall \vec{uv}$

is an optimal solution to the link-based dual LP iff it is an optimal solution to the path-based dual LP. We then discuss how the classic approach of Lagrange relaxation followed by subgradient optimization can be applied to derive effective solution algorithms in both uncapacitated and capacitated networks. These algorithms are combinatorial in nature, consisting of mostly shortest path computations, and are therefore amenable to distributed implementations.

7.1 Equivalence between Path-Based and Link-Based LPs

Again, we start with the uncapacitated case. Presented in Table 3 is the reformulated min-cost multicast linear program based on edge-flow variables, along with its dual. Here, $\vec{uv}$ denotes an edge from node u to node v, and $N_\downarrow(u)$ and $N_\uparrow(u)$ denote the downstream and upstream neighbor set of u in G, respectively. We use $\vec{uv}$ instead of e to represent an edge here because it is helpful to have explicit connections between nodes and their adjacent edges. We assume that there is a conceptual link from each T_i back to S with unlimited capacity, for succinct expression of flow conservation constraints.

The link-based LPs and the path-based LPs model the same min-cost multicast problem and are equivalent to each other. In particular, a pair of primal and dual solutions (f, y) is feasible/optimal in the link-based formulation iff it is feasible/optimal in the path-based formulation. In the primal problem, both LPs establish end-to-end network flows of rate d from S to each T_i. Flow conservation of each conceptual flow is implicit in the path-based formulation but explicit in the link-based formulation. In the dual program, both allocate edge cost w to y_i's and use x_i to compute the shortest $S \to T_i$ path under distance vector y_i. The fact that x_i is the shortest path length in any optimal dual solution is less explicit in the link-based LP, but can be deduced based on the following facts. Variables in p can be interpreted as the *altitude* of nodes [1]. The first dual constraint bounds the altitude difference of two neighboring nodes in the network with y_i, and the second dual constraint bounds x_i with the altitude difference between S and T_i. The equivalence of the two different LP formulations in the capacitated case is similar.

Since the link-based LPs have polynomial sizes, they are more practical to solve using general linear programming solution methods such as the simplex algorithm. However, past experiences suggest that much better scalability can be achieved by exploiting the specific structure of the multicast problem, if possible [6], [7]. Our experiments in large-scale network topologies [6] show that a well-designed subgradient algorithm may outperform both the simplex algorithm and the primal-dual interior-point algorithm by more than an order of magnitude, in terms of runtime. Furthermore, general linear programming solution methods are inherently centralized, while distributed subgradient algorithm design are often possible. We now proceed to describe tailored subgradient algorithms for computing optimal cost sharing and taxes from the link-based LPs.

7.2 Subgradient Algorithm Design

Lagrange relaxation coupled with subgradient optimization has proven effective in designing tailored efficient algorithms for classic combinatorial problems (including the network flow problem) with side constraints [6], [8], [7]. The side constraints can be relaxed so that efficient combinatorial algorithms can be applied to the remaining smaller problem. The price associated with the relaxation is that *a series of*, instead of *one*, smaller problems need to be solved. Lun et al. [7] studied a similar problem of min-cost multicast in cooperative networks and presented a distributed optimization algorithm based on Lagrange relaxation. In a previous work [6], we also applied similar techniques in computing the optimal orientation of an undirected multicast network. The goal in this paper is to compute optimal shadow prices y^* and t^* instead of optimal primal solutions. This leads to subtle but important differences in dualization strategies and subgradient algorithm design. Since the algorithm design in the uncapacitated case is similar to but simpler than that in the capacitated case, we focus on the latter.

We wish to have the subgradient algorithm converge at both y^* and t^* simultaneously. In order to do so, we relax two groups of constraints $f_i \le f \le c$ from the primal LP in Table 4 and introduce both y and t into the new objective function:

$$\begin{aligned}&\sum_{\vec{uv}} w(\vec{uv})f(\vec{uv}) + \sum_i \sum_{\vec{uv}} y_i(\vec{uv})\Big(f_i(\vec{uv}) - f(\vec{uv})\Big)\\&\quad + \sum_i t_i(\vec{uv})\Big(f(\vec{uv}) - c(\vec{uv})\Big)\\&= \sum_{\vec{uv}} f(\vec{uv})\left(w(\vec{uv}) + t(\vec{uv}) - \sum_i y_i(\vec{uv})\right)\\&\quad + \sum_i \sum_{\vec{uv}} y_i(\vec{uv})f_i(\vec{uv}) - \sum_i c(\vec{uv})t(\vec{uv}).\end{aligned}$$

With $f(\vec{uv})$ freely chosen from $[0,\infty)$, $\sum_{\vec{uv}} f(\vec{uv})(w(\vec{uv}) + t(\vec{uv}) - \sum_i y_i(\vec{uv}))$ is unbounded from below if $\sum_i y_i(\vec{uv}) > w(\vec{uv}) + t(\vec{uv})$. Therefore, primal feasibility requires that $\sum_i y_i(\vec{uv}) \le w(\vec{uv}) + t(\vec{uv})$, and we obtain the following Lagrange dual:

Maximize $L(y,t)$
Subject to:

$$\begin{cases} \sum_i y_i(\vec{uv}) \le w(\vec{uv}) + t(\vec{uv}) & \forall\, \vec{uv}\\ y_i(\vec{uv}), t(\vec{uv}) \ge 0 & \forall i, \forall\, \vec{uv}\end{cases}$$

where

$$L(y,t) = \mathrm{Min}_P \sum_{\vec{uv}} \Big(\sum_i f_i(\vec{uv})y_i(\vec{uv}) - c(\vec{uv})t(\vec{uv})\Big)$$

$$P: \begin{cases} \sum_{v\in N_\downarrow(u)} f_i(\vec{uv}) = \sum_{v\in N_\uparrow(u)} f_i(\vec{vu}) & \forall i, \forall u\\ f_i(\vec{T_iS}) = d & \forall i\\ f_i(\vec{uv}) \ge 0 & \forall i, \forall\, \vec{uv}\end{cases}$$

The Lagrange duality theorem [27] assures that the above problem has the same optimal solutions as in the original LP. Due to the relaxation of the interflow coupling constraints, the optimization of $L(y,t)$ is *separable*, i.e., given a pair of fixed dual vectors (y,t)

$$\begin{aligned}&\mathrm{Min}_P \sum_{\vec{uv}} \left(\sum_i f_i(\vec{uv})y_i(\vec{uv}) - c(\vec{uv})t(\vec{uv})\right)\\&= \sum_i \mathrm{Min}_P \sum_{\vec{uv}} f_i(\vec{uv})y_i(\vec{uv}) - \sum_{\vec{uv}} c(\vec{uv})t(\vec{uv}).\end{aligned}$$

During a primal update in each iteration of the subgradient algorithm, the second term $\sum_{\vec{uv}} c(\vec{uv})t(\vec{uv})$ is irrelevant since both $c(\vec{uv})$ and $t(\vec{uv})$ are constant there. The first term $\sum_i \mathrm{Min}_P \sum_{\vec{uv}} f_i(\vec{uv})y_i(\vec{uv})$ can be solved through k shortest path computations: for each i, find a shortest path from S to T_i and assign a flow rate $f_i(\vec{uv}) = d$ to every link $\vec{uv}$ on that path.

During a dual update, given a fixed multicast flow f, we compute new values for y and t with the aid of two prescribed step-size sequences θ_y and θ_t and project them to the positive orthant:

$$y_i' = y_i[k] + \theta_y[k]f_i[k] \quad \forall i, \quad t' = \max(0, t[k] - \theta_t[k]c).$$

Here, $y_i[k]$ denotes the value of y_i at round k of the subgradient iterations, and similarly for $f_i[k]$, $t[k]$, and $\theta[k]$. Values in y' and t' may not be dual feasible in general. In case dual constraints are violated, we need to project (y',t') into the feasibility polytope. We discuss three alternative approaches here for practical implementations. First, we may choose the geometrically closest feasible dual solution [7], $(y[k+1], t[k+1]) = \mathrm{argmin}_{y,t\ge 0;\sum_i y_i \le w+t} \|(y,t),(y',t')\|$. Second, we may proportionally scale the dual variables [6]: $y[k+1] = \alpha y'$ and $t[k+1] = \alpha\, t'$, where $\alpha = w/(\sum_i y_i' - t')$. Third, observe that when $\sum_i y_i(\vec{uv}) > w(\vec{uv}) + t(\vec{uv})$, the subproblem $L(y,t)$ could assign an arbitrarily large value to $f(\vec{uv})$ for the minimization purpose. This violates primal constraint $f(\vec{uv}) \le c(\vec{uv})$ and should incur a high penalty $t(\vec{uv})$. Therefore, another approach is to increase t to satisfy dual constraints: $y[k+1] = y'$ and $t[k+1] = \sum_i y_i - w$.

The subgradient iterations may start at any feasible dual variables, e.g., $y_i = w/k$, $\forall i$, and $t = 0$. Each iteration consists of a primal update followed by a dual update. For convergence, the step sizes should satisfy $\theta \ge 0$, $\lim_{k\to\infty} \theta[k] = 0$, and $\sum_{k=1}^{\infty} \theta[k] = \infty$. Example sequences include $\theta[k] = a/\sqrt{k}$ or $\theta[k] = b/(ck+d)$ for some positive constants a, b, c, and d. Upon convergence, we obtain optimal cost allocation vector y^* and optimal edge tax vector t^*. Then, we may apply the tax return procedure introduced in Section 5.4 to obtain a tax-free solution y' or apply primal recovery techniques [32] to obtain f^*.

8 Conclusion

We studied in this paper how to regulate selfish multicast flows to achieve minimum total cost. We consider explicitly the encodable property of information flows and adopt the conceptual flow structure of multicast routing accordingly. We show that encouraging cost sharing is critical in enforcing min-cost multicast and that traditional cost sharing fails to achieve this goal. We then prove that shadow-price-based cost sharing may enforce optimal multicast flows. Prior to this work, cost shares exist only to enforce suboptimal multicast flows. With further complications of finite link capacity bounds, we propose to enforce the optimal multicast flow with edge taxing and cost sharing combined. We also show that it is possible to return taxes to multicast flows to obtain a tax-free solution. Finally, we present efficient algorithms to compute the necessary cost shares and taxes.

Acknowledgments

This work has been financially supported by the Alberta Ingenuity Fund and by the Natural Sciences and Engineering Council of Canada.

References

[1] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin, *Network Flows: Theory, Algorithms, and Applications.* Prentice Hall, 1993.

[2] D.P. Bertsekas, *Network Optimization: Continuous and Discrete Models.* Athena Scientific, 1998.

[3] R. Ahlswede, N. Cai, S.R. Li, and R.W. Yeung, "Network Information Flow," *IEEE Trans. Information Theory,* vol. 46, no. 4, pp. 1204-1216, July 2000.

[4] R. Koetter and M. Médard, "An Algebraic Approach to Network Coding," *IEEE/ACM Trans. Networking,* vol. 11, no. 5, pp. 782-795, Oct. 2003.

[5] K. Jain, M. Mahdian, and M.R. Salavatipour, "Packing Steiner Trees," *Proc. 10th Ann. ACM-SIAM Symp. Discrete Algorithms (SODA),* 2003.

[6] Z. Li, B. Li, and L.C. Lau, "On Achieving Optimal Multicast Throughput in Undirected Networks," *IEEE Trans. Information Theory,* vol. 52, no. 6, pp. 2467-2485, June 2006.

[7] 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," *Proc. IEEE INFOCOM,* 2005.

[8] Z. Li and B. Li, "Efficient and Distributed Computation of Maximum Multicast Rates," *Proc. IEEE INFOCOM,* 2005.

[9] E. Koutsoupias and C. Papadimitriou, "Worst-Case Equilibria," *Lecture Notes in Computer Science,* vol. 1563, pp. 404-413, 1999.

[10] T. Roughgarden and E. Tardos, "How Bad is Selfish Routing?" *J. ACM,* vol. 49, no. 2, pp. 236-259, 2002.

[11] H. Yang and H. Huang, "The Multi-Class, Multi-Criteria Traffic Network Equilibrium and Systems Optimum Problem," *Transportation Research, Part B,* vol. 38, no. 1, pp. 1-15, Jan. 2004.

[12] L. Fleischer, K. Jain, and M. Mahdian, "Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games," *Proc. 45th IEEE Symp. Foundations of Computer Science (FOCS),* 2004.

[13] G. Karakostas and S.G. Kolliopoulos, "Edge Pricing of Multicommodity Networks for Heterogeneous Selfish Users," *Proc. 45th IEEE Symp. Foundations of Computer Science (FOCS),* 2004.

[14] J. Feigenbaum, C. Papadimitriou, and S. Shenker, "Sharing the Cost of Multicast Transmissions," *J. Computer and System Sciences,* vol. 63, pp. 21-41, 2001.

[15] S. Bhadra, S. Shakkottai, and P. Gupta, "Min-Cost Selfish Multicast with Network Coding," *IEEE Trans. Information Theory,* vol. 52, no. 11, pp. 5077-5087, Nov. 2006.

[16] L.S. Shapley, "A Value for n-Person Games," *Contributions to the Theory of Games.* Princeton Univ. Press, pp. 31-40, 1953.

[17] W. Wang, X. Li, and Y. Wang, "Truthful Multicast in Selfish Wireless Networks," *Proc. ACM MobiCom,* 2004.

[18] Z. Li, B. Li, D. Jiang, and L.C. Lau, "On Achieving Optimal Throughput with Network Coding," *Proc. IEEE INFOCOM,* 2005.

[19] 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," *J. Selected Areas in Comm.,* vol. 23, no. 1, pp. 136-150, Jan. 2005.

[20] J. Yuan, Z. Li, W. Yu, and B. Li, "A Cross-Layer Optimization Framework for Multihop Multicast in Wireless Mesh Networks," *J. Selected Areas in Comm.,* special issue on multi-hop wireless mesh networks, 2006.

[21] H.W. Kuhn and A.W. Tucker, "Nonlinear Programming," *Proc. Second Berkeley Symp. Math. Statistics and Probability,* 1951.

[22] W. Wang, X.-Y. Li, Y. Wang, and Z. Sun, "Designing Multicast Protocols for Non-Cooperative Networks," *IEEE J. Selected Areas in Comm.,* vol. 26, no. 7, pp. 1238-1249, Sept. 2008.

[23] K. Jain and V.V. Vazirani, "Applications of Approximation Algorithms to Cooperative Games," *Proc. 33rd ACM Symp. Theory of Computing (STOC),* 2001.

[24] N. Immorlica, D. Karger, E. Nikolova, and R. Sami, "First-Price Path Auctions," *Proc. Sixth ACM Conf. Electronic Commerce (EC),* 2005.

[25] Z. Li, "Cross-Monotonic Multicast," *Proc. IEEE INFOCOM,* 2008.

[26] S. Jaggi, P. Sanders, P.A. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen, "Polynomial Time Algorithms for Multicast Network Code Construction," *IEEE Trans. Information Theory,* vol. 51, no. 6, pp. 1973-1982, June 2005.

[27] G. Dantzig, *Linear Programming and Extensions.* Princeton Univ. Press, 1998.

[28] M.J. Osborne and A. Rubinstein, *A Course in Game Theory.* MIT Press, 1994.

[29] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden, "The Price of Stability for Network Design with Fair Cost Allocation," *Proc. 45th IEEE Symp. Foundations of Computer Science (FOCS),* 2004.

[30] C. Papadimitriou, "Algorithms, Games, and the Internet," *Proc. 33rd ACM Symp. the Theory of Computing (STOC),* 2001.

[31] E. Altman and L. Wynter, "Euilibrium, Games, and Pricing in Transportation and Telecommunication Networks," *Networks and Spatial Economics,* vol. 4, no. 1, pp. 7-21, Mar. 2004.

[32] H.D. Sherali and G. Choi, "Recovery of Primal Solutions When Using Subgradient Optimization Methods to Solve Lagrangian Duals of Linear Programs," *Operations Research Letters,* vol. 19, 1996.

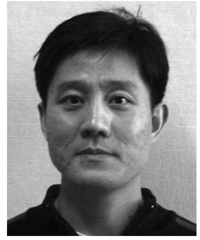

Zongpeng Li received the BE degree in computer science and technology from Tsinghua University, Beijing, in 1999, the MS degree in computer science from the University of Toronto in 2001, and the PhD degree in electrical and computer engineering from the University of Toronto in 2005. He has been working as an assistant professor in the Department of Computer Science, University of Calgary, Calgary, Alberta, Canada, since August 2005. He was named an Edward S. Rogers Sr. Scholar in 2004, received the Alberta Ingenuity New Faculty Award in 2007, was nominated for the Alfred P. Sloan Research Fellow in 2007, and received the Best Paper Award at the Ninth Passive and Active Measurement Conference (PAM) in 2008. His research interests are in computer networks, particularly in network optimization, multicast algorithm design, network game theory, and network coding. He is a member of the IEEE.

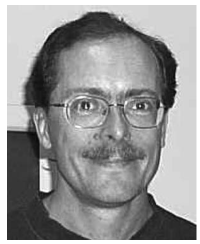

Carey Williamson received the BSc(Hons) degree in computer science from the University of Saskatchewan in 1985 and the PhD degree in computer science from Stanford University in 1992. He holds the Informatics Circle of Research Excellence (iCORE) Chair in the Department of Computer Science, University of Calgary, Calgary, Alberta, Canada, specializing in "Broadband Wireless Networks, Protocols, Applications, and Performance." He also holds an NSERC/iCORE/TELUS Mobility Industrial Research Chair in "Wireless Internet Traffic Modeling." His research interests include Internet protocols, wireless networks, network traffic measurement, workload characterization, network simulation, and Web server performance. He is a member of the IEEE and the IEEE Computer Society.