Min-cost Selfish Multicast with Network Coding

Sandeep Bhadra, Sanjay Shakkottai
Wireless Networking and Communication Group
Department of ECE
The University of Texas at Austin, Austin, TX
E-mail : {bhadra, shakkott}@ece.utexas.edu

Piyush Gupta
Bell Laboratories, Lucent Technologies
600 Mountain Ave.
Murray Hill, NJ 07974
Email: pgupta@research.bell-labs.com

***Abstract*— The single-source min-cost multicast problem is considered, which can be framed as a convex optimization problem with the use of network codes and convex increasing edge costs. A decentralized approach to this problem is presented by Lun, Ratnakar, et.al. for the case where all users cooperate to reach the global minimum. This paper analyzes the cost for the scenario where each of the multicast receivers greedily routes its flows and shows that a Nash equilibrium exists for such a scenario. An allocation rule by which the edge-cost at each edge is allocated to flows through that edge is presented. Under this rule, it is shown that for any (monomial) power-law cost function on each edge, there is a pricing rule such that the flow cost at user equilibrium is the same as the min-cost. This leads to the construction of an autonomous flow-steering algorithm for each receiver, which is also globally optimal.**

I. Introduction

The single-source multicast problem for network coding has received much attention in recent years due to the tractability of designing optimal linear network codes for this case. Ahlswede, Cai and Yeung in [1] prove that for networks where the min-cut max-flow rate cannot be achieved by simple forwarding of packets, coding incoming packets at intermediate routers (network-coding) can help achieve the max-flow min-cut rate for such networks. Further, Ho et. al. [2], [3] suggest the use of Random Linear Codes (RLCs) that can achieve the above linear network code rate asymptotically in the size of the symbol alphabet used for encoding/decoding. Since the intermediate routers can code randomly independent of other routers in the network, RLCs offer the means for decentralized design of network codes and form the basis for practical network coding schemes [4].

The problem of finding the minimum-cost multicast tree for networks has been studied extensively. For a general directed graph with a cost function at each edge, a specified root (source) and a subset of the nodes (receivers), the problem of finding a minimum-cost arborescence rooted at the source and spanning all the receivers is called the Directed Steiner Tree (DST) problem. Approximation algorithms for the DST, which is known to be NP-hard, has received considerable attention in recent years. Charikar et. al. [5] present a non-trivial $O(i(i-1)k^{1/i})$ algorithm in $O(n^i k^2 i)$ time where k is the number of receivers. An LP-relaxation of the problem leads Zosin and Khuller [6] to a $O(D+1)$-approximation for the special case when the subgraph induced by the non-terminal nodes is a tree of depth D.

A decentralized but cooperative scheme has been suggested by Lun et. al. in [7] where the authors solve the network-coding min-cost optimization from [8] using primal-dual methods by message passing between intermediate routers. However, this scheme requires a separate (differential-equation based) controller at each intermediate router for every flow passing through it. Further, many current models of heterogeneous network service provisioning assume that selfish routing decisions are made by end-users based on the price of the links [9], [10]. Such scenarios are likely to emerge with ad-hoc networking where each end-user is attached to a single multicast sink and therefore seeks to minimize her own cost. The dual problem of maximizing utility in a congestion game over a packet-forwarding network is considered in [9], [11]. Recently, the authors in [12] have framed this congestion control problem for network coding for single- and multiple-source multicasts as a generalization of the Eisenberg-Gale convex program to compute market equilibrium in the presence of economies of scale. Further, the primal-dual algorithm in [7] requires computationally intensive steps to be performed at each intermediate router.

Recent work in [13] presents a cooperative node-based primal gradient projection algorithm that is shown to converge to the minimum cost solution (joint routing and congestion control). In this scheme, each intermediate node (router) polls its neighbors for their marginal costs for each session and adjusts routing variables (maintained at each node) appropriately. The source adjusts the admission rate of the multicast session (using a gradient projection algorithm).

In this paper, we seek to design a min-cost flow-allocation algorithm when users are non-cooperative and minimize computation performed at each intermediate router. The users are assumed to be selfish agents that play a non-cooperative game to minimize personal costs selfishly without regard to the global or social optimal, and the expectation is that these users reach a Nash equilibrium if one exists. It is well-known that Nash equilibria do not optimize social welfare in general - a classical example of such an inefficient equilibrium is the ‘Prisoner’s Dilemma’ [14]. Thus it immediately becomes important to quantify the inefficiency inherent in a selfish solution - dubbed the ‘price of anarchy’ [15], [16]. The unicast selfish-agent min-cost routing problem is a classical problem in transportation literature and has been discussed by the authors in [17], [18], which corresponds to the uncoded packet forwarding scenario. The authors in [15], [19], [20] recently calculated the price of anarchy for this problem for a variety of convex cost functions for the capacitated and uncapacitated links. However, the optimization problem for the multicast min-cost flow with network coding as shown in the following section departs significantly from the min-cost unicast flow problem for uncoded packets and thus motivates independent analysis.

A. Main Contributions

In this paper, we consider the min-cost flow routing problem with network coding for the selfish-agent case. We first consider the case with a single source and T multicast sinks (receivers), with each sink requiring a total rate R. We study the case where the network supports multi-path routing. A flow (along a particular path) from the source to a sink accumulates a cost that depends on the flow rate as well as the congestion on each of the links the flow traverses. Each sink t "steers" the flow rate allocation among its paths (i.e., among all paths from the source to the selected sink t such that the sum rate across paths is R) such that *its total cost* is minimized (in other-words, a "greedy" setup for each sink). The main contributions are as follows.

(i) We present the min-cost optimization problem for the single-source multicast with network coding and derive an asymptotically accurate approximation to that problem in Section II. The selfish routing scenario is presented in Section III where a market is defined for bandwidth, being sold by links (sellers), that is utilized by flows to individual sinks (buyers). We develop a mechanism for links (sellers) to allocate the link-costs among users of the link and demonstrate that for *monomial edge cost functions* (see section III), a Nash equilibrium exists, and that the flow allocation at Nash equilibrium corresponds to the min-cost flow. In other-words, we show that the mechanism that we develop for link pricing leads to a rate allocation among users such that "greedy" flow rate allocation by each sink leads to the globally optimal flow rate allocation that minimizes the *total cost* in the network. In terms of algorithmic game theoretic literature, this means that the 'price of anarchy' [15], [16] for the considered "greedy" system is 1.

(ii) Further, inspired by the treatment of Roughgarden and Tardos in [15], we present a bicriteria result in Section III-D. We show that with *any general positive convex link cost functions*, the (network) Nash cost (i.e., the total cost with greedy sinks) when each sink requires a total of rate R is less than the *optimum* network cost (where all sinks cooperate to minimize the total network cost) under a higher rate requirement $2R$.
In the context of a network multicast, a bicriteria result is a good indicator of optimal over-provisioning. While it is often easy to design networks for the optimal case, analyzing sub-optimal cases can be difficult. However, via the bicriteria argument, it follows that a sub-optimal cost is not worse than the optimal cost *under a larger demand*. Hence, a network designer with a view to limit service cost can simply design for the optimal case with the larger demand.

(iii) Next, in Section IV we present *UESSM* – User Equilibrium with Single Source Multicast, a non-cooperative decentralized flow-steering algorithm that provably converges to the min-cost flow allocation for the class of convex, monomial edge cost functions defined in Section III. At each receiver, UESSM "steers" flows across the paths leading to it in order to greedily minimize its own cost. This allows us to achieve the min-cost flow with network coding, without having to maintain state or perform per-flow primal-dual type calculations at every intermediate router. All that links have to do in UESSM is to allocate link costs according to the rule developed in subsection III-A.

(iv) We finally present simulation results for UESSM to illustrate convergence properties.

II. Global Equilibrium

Consider a directed graph $G = (N, A)$ that models the network with the set of nodes N and the set of directed edges between them A. Let $c_e()$ be the cost function corresponding to edge $e \in A$. We assume that $c_e(x)$ is convex, positive, and monotonically increasing in x.

Flows along the set of paths $\mathcal{P}_t$ from s to t are indexed as $f_P \in \mathbb{R}$ for all $P \in \mathcal{P}_t$; $\mathcal{P} = \cup_{t\in T}\mathcal{P}_t$ is the set of all possible paths. Then the optimal cost for a rate R multicast connection from a single source $s \in N$ to sink nodes $T \subset N$ is given by the solution to the following optimization problem similar to [7], [8], GLOBAL(G, c, R):

$$\begin{aligned}
\text{minimize } C(f) &= \sum_{e\in A} C_e(z_e) \\
\text{subject to } z_e &= \max_t \left\{ \sum_{P\in\mathcal{P}_t : e\in P} f_P \right\} \quad \forall e \in A \\
f_P &\geq 0 \quad \forall P \in \mathcal{P} \\
\sum_{P\in\mathcal{P}_t} f_P &= R^{(t)} \quad \forall t \in T.
\end{aligned}$$

However, since $\max\{\ldots\}$ is not differentiable everywhere, we use the $\mathcal{L}_n$ approximation

$$\max\{x_1, x_2, \ldots x_k\} = \lim_{n\to\infty} \left(\sum_{i=1}^{k} x_i^n \right)^{1/n}$$

for analysis, thereby avoiding sub-gradient methods. Following the approximation of the $\max()$ above the $\mathcal{L}_n$-relaxed cost function of GLOBAL(G, c, R),

$$C_n(f) = \sum_{e\in A} c_e \left(\left[\sum_{t\in T} (\sum_{P\in\mathcal{P}_t : e\in P} f_P)^n \right]^{1/n} \right)$$

is differentiable everywhere. We note that a similar approximation has been used for the multicast congestion problem in [21], [7].

Since the cost functions are convex and the constraints form a convex set, the necessary and sufficient first-order Karush-Kuhn-Tucker conditions to solve $\mathcal{L}_n$-GLOBAL(G, c, R) are obtained by forming the Lagrangian

$$L(f,\lambda,\mu) = C_n(f) + \sum_{t\in T}\lambda_t(\sum_{P\in\mathcal{P}_t} f_P - R^{(t)}) - \sum_{P\in\mathcal{P}}\mu_P f_p$$

and differentiating with respect to each flow f_P, and the Lagrangian parameters λ_t and μ_P to yield the unique minimizing solution f^*, λ^*, μ^*. Note that for all $P \in \mathcal{P}$, f^*_P and μ^*_P are complimentary slack, i.e. $f^*_P = 0$ if $\mu^*_P \neq 0$ and vice-versa. Hence for paths with strictly positive flow, differentiating $L(f, \lambda, \mu)$ with respect to a particular flow f_{P_1}, for $P_1 \in \mathcal{P}_j$, gives

$$\sum_{e\in P_1} c'_e(z_e^{(n)})\left(\frac{\sum_{P\in\mathcal{P}_j:e\in P} f_P}{z_e^{(n)}}\right)^{n-1} + \lambda_j = 0,$$

where $c'_e(x) = \frac{\partial c_e(x)}{\partial x}$ and

$$z_e^{(n)} \triangleq \left(\sum_{t\in T}(\sum_{P\in\mathcal{P}_t:e\in P} f_P)^n\right)^{1/n}$$

is the corresponding $\mathcal{L}_n$ relaxation of z_e. This implies that $\forall P_1, P_2 \in \mathcal{P}_j$,

$$\sum_{e\in P_1} c'_e(z_e^{(n)})\left(\frac{\sum_{P\in\mathcal{P}_j:e\in P} f_P}{z_e^{(n)}}\right)^{n-1} \leq \sum_{e\in P_2} c'_e(z_e^{(n)})\left(\frac{\sum_{P\in\mathcal{P}_j:e\in P} f_P}{z_e^{(n)}}\right)^{n-1}. \quad (1)$$

Hence the KKT conditions in Equation (1), together with the uniqueness of solution to the strictly convex $\mathcal{L}_n$-GLOBAL(G, c, R) optimization problem can be summarized in the following lemma.

Lemma 2.1: A network coded multicast flow f is optimal for $\mathcal{L}_n$-GLOBAL(G, c, R) if and only if for all $t \in T$, and any strictly positive $P_1, P_2 \in \mathcal{P}_t$

$$\sum_{e\in P_1} c'_e(z_e^{(n)})\alpha_{e,j}^{(n)} = \sum_{e\in P_2} c'_e(z_e^{(n)})\alpha_{e,j}^{(n)}, \quad (2)$$

where,

$$\alpha_{e,j}^{(n)} \triangleq \frac{z_e^{(n)}}{x_{e,j}}.\frac{1}{\sum_{t\in T}(\frac{x_{e,t}}{x_{e,j}})^n}, \quad (3)$$

and

$$x_{e,j} \triangleq \sum_{e\in P:P\in\mathcal{P}_j} f_P$$

is the total flow of type j through the edge e.

Observe that as $n \to \infty$, $C_n(f) \to C(f)$. Hence the corresponding solutions to the set of conditions in Equation (2) will tend to the limiting conditions as $n \to \infty$.

III. Selfish routing and equilibrium

The solution to GLOBAL finds the optimum flow that minimizes routing cost in the overall network cost. This section deals with the system under the condition that each receiver minimizes its own cost to achieve user equilibrium under a defined bandwidth market to model selfish behavior as shown below. The ultimate goal of this section (and the next, respectively), is to show that under certain conditions on c_e, the user equilibrium corresponds to the global equilibrium (is comparable to the global equilibrium of a related optimization, respectively). These results will motivate a user-equilibrium based distributed optimization algorithm, discussed in Section IV.

A. The bandwidth market and link price-allocation

Each edge $e \in A$ sells bandwidth to the receivers (sinks) which are the users. Note that in the solution to the global problem we were merely concerned with the sum cost $c_e(z_e^{(n)})$ and did not need to consider how the cost of an edge in the network is divided among the flows through that network, while this sharing of costs (price allocation) needs to be defined for the user costs.

Hence, we propose a price allocation rule at each link and subsequently show that under this protocol, the sum cost under user equilibrium is equal to the min-cost for a wide range of cost functions c_e. Our price allocation rule is as follows – for each edge e the total cost of the flows $c_e(z_e^{(n)})$ is divided among flows of all type $t \in T$ so that $\frac{x^n_{e,j}}{\sum_{t\in T} x^n_{e,t}}$ fraction of the edge cost is borne by the flows in $f_P, P \in \mathcal{P}_j$ of type j. In turn $x_{e,j}$ is divided among all flows of type j through edge e in the ratio $f_P/x_{e,j}$ for all $P \in \mathcal{P}_j$. Thus the marginal cost of a flow f_P through a path $P \in \mathcal{P}_j$, $j \in T$

$$d_P^{(n)}(f) \triangleq \sum_{e\in P} c_e(z_e^{(n)})\frac{1}{x_{e,j}}\frac{x^n_{e,j}}{\sum_{t\in T} x^n_{e,t}}. \quad (4)$$

Observe that by simply multiplying and dividing by $z_e^{(n)}$, Equation (4) can be written as

$$d_P^{(n)}(f) = \sum_{e\in P}\frac{c_e(z_e^{(n)})}{z_e^{(n)}}\alpha_{e,j}^{(n)},$$

where $\alpha_{e,j}^{(n)}$ is as defined in Equation (3).

B. User costs and equilibrium

Under the selfish condition, each flow from source s to destination j tries to minimize its marginal cost. This corresponds to each receiver minimizing its own total cost selfishly.

Assuming cost functions are continuous and differentiable everywhere, we define *user equilibrium* as follows,

Definition 3.1: A user equilibrium *is an allocation f feasible in $\mathcal{L}_n$-GLOBAL(G, c, R) such that for any $P_1, P_2 \in \mathcal{P}_t$ where $f_{P_1} > 0$,*

$$d_{P_1}^{(n)}(f) \leq d_{P_2}^{(n)}(f). \quad (5)$$

Note that this version of user equilibrium is also referred to as a *local Nash equilibrium* or *Wardrop equilibrium* in existing literature [17], [20].

Corresponding to this equilibrium, the total system cost for the flow f at Nash equilibrium is then

$$C_n(f) \triangleq \sum_{P \in \mathcal{P}} d_P^{(n)}(f) f_P.$$

In other words, any small $\epsilon \rightarrow 0$ change to the flow allocations from path P_1 to P_2 will only increase the sum cost along the paths in $\mathcal{P}_t$ for sink t. The notion of a local Nash equilibrium can be practically justified in scenarios where end users are in a distributed setting, with no or partial knowledge of the system, and try to reach their own local selfish optima by making small modifications to the flow allocations across paths in $\mathcal{P}_t$, where the flow steering proceeds only if that provides the selfish agent with immediate cost reduction.

C. User equilibrium vs. Global optimum

The similarity between the conditions in Lemma 2.1 and Definition 3.1 have been noticed for the case of costs depending on sum flows through an edge by Dafermos and Sparrow [17] and Beckman [18] and is cited by [15]. An important difference in our case is that while the edge cost in [15], [17], [18] is equally divided among all the flows through it, here, the cost is borne only by the *critical* flows through the edge. Lemma 2.1 and Definition 3.1 then lead us to the following lemma that allows us to formulate the Nash equilibrium condition for a particular set of edge cost functions in terms of a global optimum for the same graph over *a different set of edge cost functions*.

Lemma 3.1 ([17], [18]): A single source multicast flow f solves $\mathcal{L}_n$-GLOBAL(G, c, R) if and only if it is in local Nash equilibrium for $\mathcal{L}_n$-GLOBAL(G, c', R). Also, if $c_e(x)/x$ is continuous and monotonically increasing in x for all $e \in A$, then a local Nash equilibrium flow f exists for $\mathcal{L}_n$-GLOBAL(G, c, R). Moreover, if f and $\tilde{f}$ are feasible flows at Nash equilibrium, then $C_n(f) = C_n(\tilde{f})$.
Proof: See [22]. ∎

This ensures that there exists a flow allocation that satisfies the user equilibrium (5).

We can now present the analog of the main result in Roughgarden and Tardos [15] for the min-cost multicast problem with network coding in the following theorem.

Theorem 3.1: If for an instance (G, c, R) the cost function at each edge e is of the monomial form $c_e(z_e) = a_e z_e^k$ for any fixed $k \in \mathbb{R},\ k > 1$, then for all $n \in \mathbb{N}$, the cost of flow f at local Nash equilibrium $C^{(n)}(f)$ equals the cost $C(n)(f^*)$ of the global min-cost flow f^*.
Proof: Follows immediately by differentiating c_e and using Lemma 3.1. ∎

We note that notwithstanding the simplicity of the proof, the above result is significant due to it's application in Section IV. The result above implies that for an large class of edge cost functions, a global min-cost multicast with network coding can be achieved by merely steering flows across edges to achieve user equilibrium corresponding to each sink t.

Note that in general, the global min-cost flow can be achieved if each link charges the "Lagrangian cost" $h_e(x) = \int_0^x c_e(t)/t\ dt$ instead of the true cost $c_e(x)$. However, this would imply that the seller (link) earns an amount disproportionate to the true value of the goods or services(bandwidth) sold. The link-price allocation scheme detailed in subsection III-A ensures that the seller receives the 'fair' cost $c_e(x)$ but charges the selfish users differently so as to ensure that user equilibrium coincides with the socially-optimal flow allocation.

Observe that at $n = 1$, the $\mathcal{L}_1$-GLOBAL(G, c, R) problem is the same as the classical min-cost flow allocation problem. Also, correspondingly, our price allocation reduces to the allocation of link cost to a sink in linear proportion to the magnitude of flow to that sink through the particular link – thereby making the marginal cost of every flow through a link the same. This is exactly the same as the anarchic scenario in [15] where each flow through a particular edge e has the same marginal cost(edge delay) l_e and the net cost of that edge $c_e = l_e \sum_{e \in P, P \in \mathcal{P}} f_P$.

In general, the results herein define a differentiated pricing scheme for a shared service whose cost depends not on the sum of the demands but on the max demand. At the limit $n \rightarrow \infty$, we observe that only the set of users $T' = \arg\max_{t \in T} \sum_{P \in \mathcal{P}_t : e \in P} f_P$ pay for the cost of the link. Our price allocation rule automatically induces separate selfish agents to collaborate to benefit from this economy of scale.

D. A Bicriteria result for the min-cost flow problem

In this section, we determine a bicriteria result to compare the local Nash cost under a particular rate requirement R with the optimum routing cost under a higher rate requirement $2R$ for any general positive convex increasing cost function.

It is instructive to note that our bicriteria result is the same as that achieved by Roughgarden and Tardos [15] for the uncoded unicast flow case. However, the proof technique in [15] does not apply directly owing to the fact that the edge cost there is a function of the sum flow through an edge, whereas, network coding allows the edge cost to be a function only of the maximum flow of a particular type through that edge.

Theorem 3.2: If f is the flow assignment under local Nash equilibrium for the instance (G, c, R) and f^* is any *flow assignment feasible for the instance $(G, c, 2R)$, then $C(f) \leq C(f^*)$.*
Proof: See [22]. ∎

IV. Distributed Algorithms for Min-cost flow

Section III demonstrates that the sum-cost of the edges with any uniform power-law edge cost function under user equilibrium is the same as the min-cost. This result lends itself readily to the construction of a simple non-cooperative optimal min-cost flow routing algorithm for a single-source multicast with network coding. The following section deals with the single-source multicast for sake of simplicity. It is easy to show that due to the separable and additive nature of the costs for the multiple-source multicast, we can run the same algorithm independently over each session to reach the user equilibrium in this case too.

In this section, we develop UESSM, a non-cooperative decentralized flow-steering algorithm that provably converges to the min-cost flow allocation for the class of convex, monomial edge cost functions defined in Section III. At each receiver, UESSM "flow-steers" among the paths leading to it in order to greedily minimize its cost. This allows us to achieve the min-cost flow with network coding, without having to perform per-flow primal-dual type calculations at every intermediate router.

The implementation of UESSM, assumes *flow routing* between the source and destination, where the source router encodes downstream hop-by-hop routing information into the IP-header, as can be implemented in IPv6. The intermediate routers in the network between the source and sink do not need to maintain state-information locally. All that the intermediate routers need to do is route packets along the outgoing edges corresponding to the hop-by-hop information embedded in each packet and network code across packets of the same type at each instant of time using a random linear code.

Also, each downstream packet aggregates the cost that it has paid along each edge on a particular flow path. For efficiency, this information need not be carried by every downstream packet, but only by representative packets at each iteration of the algorithm.

We initialize with any arbitrary flow allocation among the paths $P \in \mathcal{P}$ that satisfy the source and sink rates of R. The flow allocations in our implementation are elements from a lattice $\mathcal{L} = \{\ldots, -2\Delta, -\Delta, 0, \Delta, 2\Delta, \ldots\}$, for some fixed $\Delta > 0$. Now, one of the sinks $t \in T$ is chosen at random. Sink t now picks two paths $P_1, P_2 \in \mathcal{P}_t$ at random and compares the values of path costs $d_{P_1}(f)$ and $d_{P_2}(f)$. If $d_{P_1}(f) > d_{P_2}(f)$ and $f_{P_1} > \Delta$, then $f_{P_1} \leftarrow f_{P_1} - \Delta$ and $f_{P_2} \leftarrow f_{P_2} + \Delta$ or vice versa.

A. Asynchronous implementation

The implementation of the algorithm above does *not* require synchronous timing between the clocks at the various sink nodes but only requires that the clocks have the same cycle frequency. We assume that the path-delay timescale along the network (for the update of the path costs etc.) is negligible compared to the time-steps in which the algorithm proceeds. Each sink $j \in T$ picks a random delay that is exponentially distributed before adjusting it's flows. Since the exponential distribution is a continuous time-distribution, the collision probability is small. Further, since all flow steering is implemented at the source, the source can be designed to sequentially adjust flows of each sink. This ensures that only one sink adjusts flows at a time in the asynchronous algorithm, thereby retaining the same features as the synchronous implementation.

B. Convergence of UESSM to the min-cost flow

In the following, we will prove that after a sufficiently long period of time UESSM settles with a flow allocation that is within a Δ-*neighborhood* defined below of the user-equilibrium flow allocation in Definition 3.1.

Definition 4.1: **Δ-neighborhood**: A flow allocation f feasible on (G, c, R) is said to be within a Δ-neighborhood of another feasible flow f' if and only if for each $P \in \mathcal{P}$, $|f_P - f'_P| < \Delta$.

Theorem 4.1: UESSM eventually converges to a Δ-neighborhood of f^0.

Proof: See [22]. ∎

C. Simulation results

We simulate UESSM over the classic 7-node butterfly network in [1], [7] with the edge costs shown in Figure 1 for a rate 1 multicast session from source S_1 to destinations D_1 and D_2. The links are marked with the edge cost functions $c_e(x)$. In this example, $\mathcal{P}_1 = \{f_1, f_2, f_3\}$ and $\mathcal{P}_2 = \{F_1, F_2, F_3\}$.

We first study how $C_n(f)$ changes with increasing values of n in the $\mathcal{L}_n$-approximation to the $\max$ function. The trajectories for 100 representative UESSM runs with $\Delta = 0.01$ with varying values of n are plotted in Figure 2. The $n = 1$ case corresponds to multicast without network coding and has a much higher sum-cost than that achieved by the $\mathcal{L}_{100}$ approximation, which is very close to the cost with using the non-differentiable $\max$ function in GLOBAL(G, c, R). However, we note that there is not much gain in going from $n = 10$ to $n = 100$. This suggests that the $\mathcal{L}_n$ approximation works well for even small values of n. This leads us to the open problem of bounding the error in the $\mathcal{L}_n$ approximation for small values of n.

We have also shown error bars corresponding to one standard deviation about the mean, with random initial conditions. We observe that, irrespective of initial conditions, the simulation sum-cost trajectories converges to the mean with progressively small variance. Typical trajectories of flow rates through various paths for the Butterfly network with a step-size of $\Delta = 0.01$ are presented in Figure 3.

ACKNOWLEDGMENTS

The research of S. Bhadra and S. Shakkottai was supported by NSF Grants ACI-0305644, CNS-0325788 and CNS-0347400, and that of P. Gupta was supported in part by NSF Grant CCR-0325673. Early results leading to this paper were presented as an invited talk at the Workshop on Mathematical Modeling and Analysis of Computer Networks, Waterloo, Canada, May 6, 2005. This workshop does not have any printed/electronic proceedings.

REFERENCES

[1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inform. Theory*, vol. 46, pp. 1204–1216, 2000.

[2] T. Ho, R. Koetter, M. Médard, D. Karger, and M. Effros, "The benefits of coding over routing in a randomized setting," in *Proc. 2003 International Symposium on Information Theory*. IEEE, 2003.

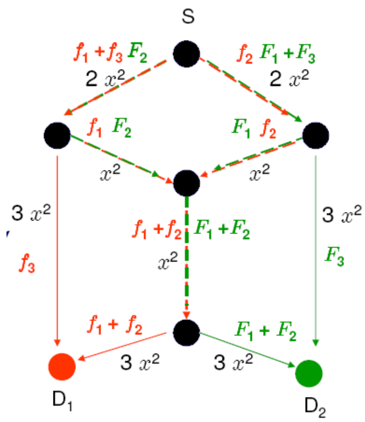

Fig. 1. 7-node Butterfly network.

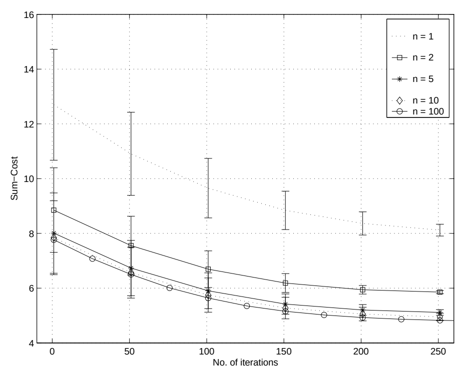

Fig. 2. UESSM Algorithm: Comparing $\mathcal{L}_n$ approximation for $n = 1, 2, 5, 10, 100$.

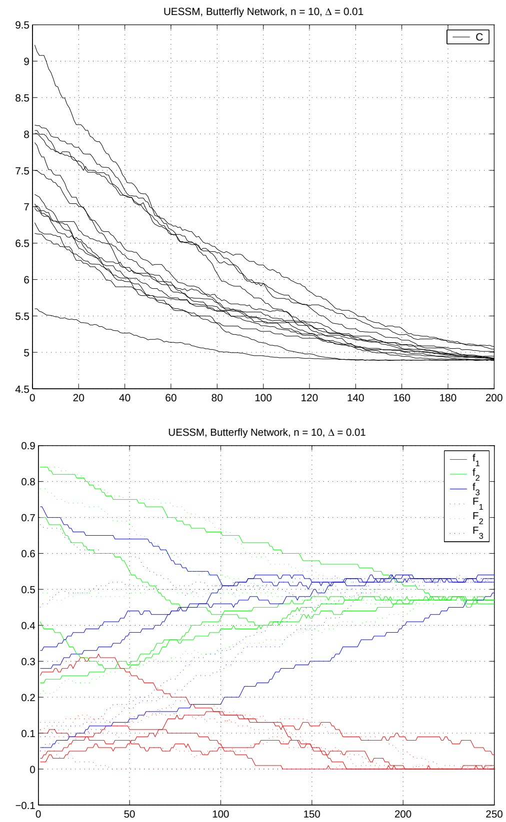

Fig. 3. UESSM Algorithm trajectories: Sum costs and flows for the Butterfly network, $\mathcal{L}_{10}$-GLOBAL(G, c, R), $\Delta = 0.01$.

[3] T. Ho, M. Médard, J. Shi, M. Effros, and D. R. Karger, "On randomized network coding," in *41st Allerton Annual Conference on Communication, Control, and Computing*, Monticello, IL, October 2003.

[4] P. A. Chou, Y. Wu, and K. Jain, "Practical network coding," in *41st Allerton Annual Conference on Communication, Control, and Computing*, Monticello, IL, October 2003.

[5] M. Charikar, C. Chekuri, T. y. Cheung, Z. Dai, A. Goel, S. Guha, and M. Li, "Approximation algorithms for directed steiner problems," in *Proc. Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 1998)*, 1998, pp. 192–200.

[6] L. Zosin and S. Khuller, "On directed steiner trees," in *Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2002)*, 2002, pp. 59–63.

[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," in *Proc. IEEE Infocom 2005*, March 2005.

[8] D. S. Lun, M. Médard, T. Ho, and R. Koetter, "Network coding with a cost criterion," in *Proc. 2004 International Symposium on Information Theory and its Applications (ISITA 2004)*, October 2004.

[9] R. Johari and J. Tsitsiklis, "Efficiency loss in a network resource allocation game," *Mathematics of Operations Research*, vol. 29, no. 3, pp. 407–435, 2004.

[10] D. Acemoglu and A. Ozdaglar, "Flow control, routing and performance from service provider viewpoint," *MIT-LIDS-WP-1696 Technical Report, MIT-LIDS*, December 2003.

[11] F. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operations Research Society*, vol. 49, pp. 237–252, 1998.

[12] K. Jain, V. V. Vazirani, and Y. Yu, "Market equilibria for homothetic, quasi-concave utilities and economies of scale in production," in *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*, January 2005.

[13] Y. Xi and E. M. Yeh, "Distributed algorithms for minimum cost multicast with network coding," in *Proceedings of the Allerton Conference on Communication, Control, and Computing*, Monticello, IL, September 2005.

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

[15] T. Roughgarden and E. Tardos, "How bad is selfish routing?" *Journal of the ACM*, vol. 49, no. 2, pp. 236–259, March 2002.

[16] C. Papadimitriou, "Algorithms, games, and the internet," in *Proc. 33rd Annual ACM Symposium on the Theory of Computing*, 2001, pp. 749–753.

[17] S. C. Dafermos and F. T. Sparrow, "The traffic assignment problem for a general network," *Journal of Research of the National Bureau of Standards, Series B*, vol. 73B, no. 2, pp. 91–118, 1969.

[18] M. Beckmann, C. B. McGuire, and C. B. Winsten, *Studies in the Economics of Transportation*. Yale University Press, 1956.

[19] T. Roughgarden, "The price of anarchy is independent of the network topology," in *Annual ACM Symposium on the Theory of Computing*, 2002, pp. 428–437.

[20] J. Correa, A. Schulz, and N. Stier-Moses, "Selfish routing in capacitated networks," *Mathematics of Operations Research*, vol. 29, no. 4, pp. 961–976, November 2004.

[21] S. Deb and R. Srikant, "Congestion control for fair resource allocation in networks with multicast flows," *IEEE/ACM Transactions on Networking*, pp. 274–285, April 2004.

[22] S. Bhadra, S. Shakkottai, and P. Gupta, "Min-cost selfish multicast with network coding," Bell Labs Technical Report, Tech. Rep., August 2005.