

Network Coding with a Cost Criterion

Desmond S. Lun[†], Muriel Médard[†], Tracey Ho[†], and Ralf Koetter[‡]

[†] Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
E-mail: {dslun, medard, trace}@mit.edu

[‡] Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana, IL 61801, USA
E-mail: koetter@uiuc.edu

Abstract

We consider applying network coding in settings where there is a cost associated with network use. We show that, while minimum-cost multicast problems without network coding are very difficult except in the special cases of unicast and broadcast, finding minimum-cost subgraphs for single multicast connections with network coding can be posed as a linear optimization problem. In particular, we apply our approach to the problem of minimum-energy multicast in wireless networks with omnidirectional antennas and show that it can be handled by a linear optimization problem when network coding is used. For the case of multiple multicast connections, we give a partial solution: We specify a linear optimization problem that yields a solution of no greater cost than any solution without network coding and that we suspect can potentially be substantially better.

1. Introduction

The selection of routes is an issue of utmost importance in data networks that has so far received scant attention in the literature on network coding. Indeed, the standard framework in which network coding is cast, that of network information flow problems [2], assumes that we have a network with limited-capacity links and considers whether or not a given set of connections can be simultaneously established, but gives no consideration to the resources that are consumed as a result of communicating on the links. In addition, such a framework implicitly assumes a certain homogeneity in network traffic — the goal is to ensure that connections are established as long as the network has the capacity to accommodate them, regardless of the type or purpose of the connections — which is frequently

not the case. The most notable example is today's internet, which not only carries different types of traffic, but is also used by a vastly heterogeneous group of end users with differing valuations of network service and performance. It has been variously proposed that such heterogeneous networks be priced [20], with some models allowing for selfish routing decisions based on the price of the links [13, 1].

In the present paper, we consider applying network coding in settings where there is a cost associated with network use, our natural objective being to select subgraphs for coding that minimize the cost incurred.

We commence by considering single multicast connections (which include single unicast and broadcast connections as special cases) in the following section. In Section 3, we study a particular instance of minimum-cost single multicast connections that has attracted much recent interest, that of minimum-energy multicast in wireless networks. In Section 4, we treat the case of multiple multicast connections.

2. Single multicast connections

Whenever the members of a multicast group have a selfish cost objective, or when the network sets link weights to meet its objective or enforce certain policies and each multicast group is subject to a minimum-weight objective, we wish to set up single multicast connections at minimum cost. Network coding for single multicast connections is relatively simple as we have a simple characterization of feasibility in networks with limited-capacity links [2, Theorem 1] and, moreover, it is known that it suffices to consider linear operations over a sufficiently large finite field on a sufficiently long vector created from the source process [15, Theorem 3.3], [14, Theorem 4].

We model the network with a directed graph $G = (N, A)$. For each link $(i, j) \in A$, we associate non-negative numbers a_{ij} and c_{ij} , which are the cost per unit flow and the capacity of the link, respectively.

This work was supported by the National Science Foundation under grant nos. ANI-0335256 and CCR-0325673, and by Hewlett-Packard under grant no. 008542-008.

Suppose we have a source node s producing data at a positive, real rate R that it wishes to transmit to a non-empty set of terminal nodes T . Consider the following linear optimization problem:

$$\begin{aligned}
& \text{minimize} && \sum_{(i,j) \in A} a_{ij} z_{ij} \\
& \text{subject to} && z_{ij} \geq x_{ij}^{(t)}, \quad \forall (i,j) \in A, t \in T, \\
& && \sum_{\{j|(i,j) \in A\}} x_{ij}^{(t)} - \sum_{\{j|(j,i) \in A\}} x_{ji}^{(t)} = \sigma_i^{(t)}, \quad (1) \\
& && \forall i \in N, t \in T, \\
& && c_{ij} \geq x_{ij}^{(t)} \geq 0, \quad \forall (i,j) \in A, t \in T,
\end{aligned}$$

where

$$\sigma_i^{(t)} = \begin{cases} R & \text{if } i = s, \\ -R & \text{if } i = t, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Theorem 1. *The vector z is part of a feasible solution for the linear optimization problem (1) if and only if there exists a network code that sets up a multicast connection in the graph G at rate arbitrarily close to R from source s to terminals in the set T and that puts a flow arbitrarily close to z_{ij} on each link (i,j) .*

Proof. First suppose that z is part of a feasible solution for the problem. Then, for any t in T , we see that the maximum flow from s to t in the network where each link (i,j) has capacity z_{ij} is at least R . So, by Theorem 1 of [2], a network coding solution with flow arbitrarily close to z_{ij} on each link (i,j) exists. Conversely, suppose that we have a network coding solution with flow arbitrarily close to z_{ij} on each link (i,j) . Then the capacity of each link must be at least z_{ij} and, moreover, flows of size R exist from s to t for each t in T (again by Theorem 1 of [2]). Therefore the vector z is part of a feasible solution for the optimization problem. $\square$

It follows straightforwardly from Theorem 1 that the linear optimization problem (1) finds the optimal cost for a rate R multicast connection from s to T in graph G that can be asymptotically achieved with a network code.

To establish minimum-cost multicast with network coding, therefore, it suffices to solve problem (1) and then compute a code that achieves the optimal cost within an arbitrary factor, which can be done systematically in time polynomial in $|N|$, $|A|$, and the block-length of the code [12] or, alternatively, in a random, decentralized fashion [10, 11, 6]. On the other hand, the standard approach for establishing minimum-cost multicast without network coding requires solving the Steiner tree problem on directed graphs, which is

known to be NP-complete (and which, moreover, only really applies when the links are of unlimited capacity). Although tractable approximation algorithms exist for the Steiner tree problem on directed graphs (for example [23, 5]), the multicast routing solutions thus obtained are suboptimal relative to the minimum-cost multicast without network coding, which in turn is suboptimal relative to when network coding is used. Hence network coding promises to provide significant cost improvements for practical multicast routing.

Another advantage offered by network coding is that problem (1) can be easily modified to accommodate convex cost functions, yielding a monotropic programming problem, or, if the cost functions are also piecewise-linear, reformulated into another linear optimization problem. When network coding is not used, it is not at all clear how any non-linear cost functions could be handled and, indeed, solving the Steiner tree problem on directed graphs no longer suffices to find the optimal solution.

Note that, in the special case of unicast, problem (1) reduces to a minimum-cost flow problem, whose solution leads to a fractional routing of flow over a number of paths, which is referred to as bifurcated routing. We see, then, that network coding essentially facilitates the same natural extension to general multicast for bifurcated routing as that for single-path routing. Hence it would appear that, in a cost-efficient network that uses bifurcated routing and that services both unicast and multicast connections, network coding is a *sine qua non*.

We have thus far assumed that the cost is a function (in fact, a separable function) of the vector z , which reflects the flow on every link. There are, however, scenarios where this is not the case. One such scenario is where we are routing selfishly to minimize latency. In this case, we are generally interested in minimizing the latency of each member of the multicast group — we have what Bharath-Kumar and Jaffe [4] term a *destination cost* criterion and, just as the solution for single-path routing is to compute the shortest path from the source to each terminal and route over the resulting tree, the solution for bifurcated routing is to compute the minimum-cost flow from the source to each terminal and code over the resulting “union” of flows. It is not hard to see that, whilst optimizing for destination cost criteria is very easy, it can potentially be very wasteful of network resources. Another scenario is that of energy-limited wireless networks, which we consider in the next section.

3. Minimum-energy multicast in wireless networks

In wireless networks, computing the energy cost is complicated by of the omnidirectionality of the antennas; so when transmitting from node i to node j , we get transmission to all nodes whose distance from i is less than that from i to j “for free” — a phenomenon referred to as the “wireless multicast advantage” in [22]. Under this phenomenon, even the problem of minimum-energy broadcast in wireless networks without network coding is NP-complete [3].

In our formulation of minimum-cost routing with network coding, modifying the cost function to reflect the wireless multicast advantage poses no serious complication, as we now proceed to show.

Let i be a node in N . We impose an ordering $\preceq$ on the set of outgoing links from i , such that $(i, j) \preceq (i, k)$ if and only if $a_{ij} \leq a_{ik}$. Typically, the set of outgoing links from i will be the set of all nodes within a certain, fixed radius of i and the cost a_{ij} of the link between nodes i and j will be proportional to their distance raised to some power α , where $\alpha \geq 2$.

Recall that a network coding solution with flow arbitrarily close to z_{ij} on each link (i, j) exists if and only if we can accommodate flows $x^{(t)}$ for all terminals t in T in the network where each link (i, j) has capacity z_{ij} . Consider a particular link (i, j) . Owing to the omnidirectionality of the antennas, flow can be pushed from i to j by pushing it to any node k such that $(i, k) \in A$ and $(i, k) \succeq (i, j)$, and it follows that the maximum flow $x_{ij}^{(t)}$ that can be pushed for a given t in T is

$$z_{ij} + \sum_{\{k|(i,k) \in A, (i,k) \succeq (i,j)\} \setminus \{j\}} (z_{ik} - x_{ik}^{(t)}). \quad (3)$$

Hence we have

$$\sum_{\{k|(i,k) \in A, (i,k) \succeq (i,j)\}} (z_{ik} - x_{ik}^{(t)}) \geq 0 \quad (4)$$

for all $t \in T$.

Thus, the relevant linear optimization problem that needs to be solved is the following.

$$\begin{aligned} & \text{minimize} \quad \sum_{(i,j) \in A} a_{ij} z_{ij} \\ & \text{subject to} \quad \sum_{\{k|(i,k) \in A, (i,k) \succeq (i,j)\}} (z_{ik} - x_{ik}^{(t)}) \geq 0, \\ & \quad \quad \quad \forall (i, j) \in A', t \in T, \\ & \quad \quad \quad \sum_{\{j|(i,j) \in A\}} x_{ij}^{(t)} - \sum_{\{j|(j,i) \in A\}} x_{ji}^{(t)} = \sigma_i^{(t)}, \\ & \quad \quad \quad \forall i \in N, t \in T, \\ & \quad \quad \quad x_{ij}^{(t)} \geq 0, \quad \forall (i, j) \in A, t \in T, \end{aligned} \quad (5)$$

where A' is a subset of A with the property that the constraint (4) is unique for all $(i, j) \in A'$ (for example,

if (i, j_1) and (i, j_2) are unique members of A such that $a_{ij_1} = a_{ij_2}$, then only one of the two is in A').

Note that we have not included any capacity constraints in problem (5). This is the correct formulation if energy is the most significant constraint and we are interested in optimizing only for energy usage, without regard to rate or spectral efficiency. If, however, bandwidth is a significant constraint and overall throughput a concern, then it is necessary to add the constraint $z \in Z$. Unlike the wireline scenario, the constraint set Z will in general not be separable (i.e. the constraints on separate links will in general be coupled) because of the effect of interference.

4. Multiple multicast connections

We turn our attention to multiple multicast connections in this section. So rather than one source process at a single node, we suppose instead that there are M source processes $X_1, \dots, X_M$ with rates $R_1, \dots, R_M$, respectively, which are generated at (possibly different) nodes $s_1, \dots, s_M$ in N . Each terminal $t \in T$ demands a subset of the source process that are generated in the network, which we specify with the set $D(t) \subset \{1, \dots, M\}$. While the nodes $s_1, \dots, s_M$ can be different, they do not need to be, and an important example where they are not is when a data source has been compressed by a multiresolution or successive refinement source code (see, for example, [9]) and is to be transmitted to users of the network with varying demands of quality.

Given a network with limited-capacity links, the problem of determining whether or not a set of multicast connections is feasible with network coding is considerably more difficult than the equivalent problem when there is only a single multicast connection. All that we currently have are rather cumbersome bounds on the feasible region [21]. In addition, it is known that it is not sufficient to consider linear operations over a sufficiently large finite field on sufficiently long vectors created from the source process — non-linear functions may be necessary in general [7]. Thus it appears that we have little hope of finding minimum-cost solutions. We are not, however, precluded from finding non-trivial cost improvements that network coding can provide. We therefore propose a linear optimization problem whose minimum cost is no greater than the minimum cost of any solution without network coding and show, with a constructive proof, that feasible solutions correspond to network codes that perform linear operations on vectors created from the source processes.

We first introduce some additional notation. For

any node i , let $T(i)$ denote the terminals that are accessible from i , i.e.

$$T(i) = \{t \in T \mid \exists \text{ a forward path to } t \text{ from } i \text{ or } t = i\}, \quad (6)$$

and let $\mathcal{C}(i)$ denote the set of atoms of the algebra generated by $\{D(t)\}_{t \in T(i)}$ (for the reader unfamiliar with set algebras and atoms, see, for example, [8, Section 4.1]), i.e.

$$\mathcal{C}(i) = \left\{ \bigcap_{t \in T(i)} C(t) \mid C(t) = D(t) \text{ or } C(t) = \{1, \dots, M\} \setminus D(t) \setminus \{\emptyset\} \right\}. \quad (7)$$

In essence, what $\mathcal{C}(i)$ gives is a set partition of $\{1, \dots, M\}$ that represents the sources that can be mixed (combined linearly) on links going into i . For a given $C \in \mathcal{C}(i)$, the terminals that receive a source process in C by way of link (j, i) either receive all the source processes in C or none at all. Hence source processes in C can be mixed on link (j, i) as the terminals that receive the mixture will also receive the source processes (or mixtures thereof) necessary for decoding.

Consider the following linear optimization problem:

$$\begin{aligned} & \text{minimize} \quad \sum_{(i,j) \in A} a_{ij} z_{ij} \\ & \text{subject to} \quad c_{ij} \geq z_{ij} = \sum_{C \in \mathcal{C}(j)} y_{ij}^{(C)}, \quad \forall (i, j) \in A, \\ & y_{ij}^{(C)} \geq \sum_{m \in C} x_{ij}^{(t,m)}, \quad \forall (i, j) \in A, t \in T, C \in \mathcal{C}(j), \\ & \sum_{\{j \mid (i,j) \in A\}} x_{ij}^{(t,m)} - \sum_{\{j \mid (j,i) \in A\}} x_{ji}^{(t,m)} = \sigma_i^{(t,m)}, \\ & \quad \quad \quad \forall i \in N, t \in T, m = 1, \dots, M, \\ & x_{ij}^{(t,m)} \geq 0, \quad \forall (i, j) \in A, t \in T, m = 1, \dots, M, \end{aligned} \quad (8)$$

where

$$\sigma_i^{(t,m)} = \begin{cases} R_m & \text{if } v = s_m \text{ and } m \in D(t), \\ -R_m & \text{if } m \in D(i), \\ 0 & \text{otherwise,} \end{cases} \quad (9)$$

and we define $D(i) := \emptyset$ for i in $N \setminus T$. Again, the optimization problem can be easily modified to accommodate convex cost functions.

Theorem 2. *If the vector z is part of a feasible solution for the linear optimization problem (8), then there exists a network code that sets up multicast connections for $m = 1, \dots, M$ at rate arbitrarily close to R_m from*

source s_m to terminals in the set $\{t \in T \mid m \in D(t)\}$ and that puts a flow arbitrarily close to z_{ij} on each link (i, j) .

Proof. Let z be part of a feasible solution. We first consider the case where $R_1 = \dots = R_M = 1$ and the underlying multicommodity flows $\{x^{(t,m)}\}$ are forward path flows of size one. In this case, the codes we use are linear with symbols from a finite field $\mathbb{F}$.

We introduce M nodes $i_1, \dots, i_M$ that are connected to nodes $s_1, \dots, s_M$, respectively. Each link (i_m, s_m) carries one unit of flow of commodity m for each of the terminals in the set $\{t \in T \mid m \in D(t)\}$. We associate with each link (i, j) a set of *global coding vectors* $B(i, j) \subset \mathbb{F}^M$. The set $B(i, j)$ represents the symbols that are transmitted on link (i, j) as a linear function of the original source processes; thus we have $B(i_m, s_m) = \{[0^{m-1}, 1, 0^{M-m}]\}$ for $m = 1, \dots, M$ and the global coding vectors that are put out by a node on its outgoing links must be linear combinations of the ones it receives on its incoming links. Moreover, it is not difficult to see that the terminal t can recover its demands $D(t)$ if and only if

$$\begin{aligned} & \text{span} \left(\bigcup_{\{i \mid (i,t) \in A\}} B(i, t) \right) \\ & \supset \text{span} \left(\bigcup_{m \in D(t)} B(i_m, s_m) \right). \end{aligned} \quad (10)$$

Our proof now follows a development similar to the proof of the main result in [12]. We step through the nodes in topological order, examining the outgoing links and defining global coding vectors on them. On each link (i, j) , we write $B(i, j) = \bigcup_{C \in \mathcal{C}(j)} B_C(i, j)$, where the $B_C(i, j)$ are disjoint, and, if the flow variable $x_{ij}^{(t,m)} = 1$, then it is associated with global coding vectors in the set $B_C(i, j)$ for the unique $C \in \mathcal{C}(j)$ such that $m \in C$.

Every time we define a new global coding vector, we maintain the following invariants:

1. For every terminal $t \in T$, the set of most recently defined global coding vectors associated with each flow path $\{x^{(t,m)}\}_{m \in D(t)}$ forms a set of $|D(t)|$ global coding vectors B_t with the property that $\text{span}(B_t) = \text{span}(\bigcup_{m \in D(t)} B(i_m, s_m))$.
2. The set of global coding vectors $B_C(i, j)$ has the property that $B_C(i, j) \subset \text{span}(\bigcup_{m \in C} B(i_m, s_m))$.

The invariants are initially established by the sets of global coding vectors $B(i_1, s_1), \dots, B(i_M, s_M)$. Now, consider node i and link (i, j) and suppose that the

invariants have been thus far satisfied. Let $C \in \mathcal{C}(j)$, and define $x_{ij}^{(t,C)} := \sum_{m \in C} x_{ij}^{(t,m)}$ and $t^* := \arg \max_{t \in T} x_{ij}^{(t,C)}$. For all terminals in the set $S := \{t \in T \mid x_{ij}^{(t,C)} > 0\}$, there are global coding vectors associated with incoming flows of commodities $m \in C$ that must be replaced by global coding vectors in the set $B_C(i, j)$. First, note that as long as $|\mathbb{F}| \geq |T|$, we can find $x_{ij}^{(t^*, C)}$ valid global coding vectors for $B_C(i, j)$ such that $\dim(\text{span}(B_t)) = |D(t)|$ for all $t \in S$ [12]. Secondly, we have $T(i) \supset T(j)$; and, given $m \in C$, if C' is the unique element of $\mathcal{C}(i)$ such that $m \in C'$, then it is not hard to see that $C' \subset C$. Hence all the global coding vectors associated with incoming flows of commodities $m \in C$ are elements of $\text{span}(\bigcup_{m \in C} B(i_m, s_m))$, so it follows that $B_C(i, j)$, whose elements are linear combinations of these global coding vectors, is a subset of $\text{span}(\bigcup_{m \in C} B(i_m, s_m))$. But, for $t \in S$, $\text{span}(\bigcup_{m \in C} B(i_m, s_m)) \subset \text{span}(\bigcup_{m \in D(t)} B(i_m, s_m))$ since $C \subset D(t)$. Therefore, $\text{span}(B_t) \subset \text{span}(\bigcup_{m \in D(t)} B(i_m, s_m))$ and, because of the dimensionality of $\text{span}(B_t)$, it follows that $\text{span}(B_t) = \text{span}(\bigcup_{m \in D(t)} B(i_m, s_m))$.

We define sets of global coding vectors for all $C \in \mathcal{C}(j)$ as we did above and we see that

$$\begin{aligned} |B(i, j)| &= \sum_{C \in \mathcal{C}(j)} |B_C(i, j)| \\ &= \sum_{C \in \mathcal{C}(j)} \max_{t \in T} \left\{ \sum_{m \in C} x_{ij}^{(t,m)} \right\}. \end{aligned} \quad (11)$$

It is evident that, upon stepping through the entire graph, condition (10) is satisfied and, since z forms part of the feasible solution, we have $z_{ij} \geq |B(i, j)|$, so a flow arbitrarily close to z_{ij} can be placed on each link (i, j) .

In the general case, we code over time $n \geq 1$. We convert the rate- R_m source process X_m into $\lfloor nR_m \rfloor$ unit rate source processes $X_{m,1}, \dots, X_{m,\lfloor nR_m \rfloor}$. Now, for a given $t \in T$ and $m \in \{1, \dots, M\}$, consider the graph G with link capacities $\lceil nx_{ij}^{(t,m)} \rceil$. Since the minimum cut between s_m and t in this graph must be at least $\lfloor nR_m \rfloor$, there exists an integer flow $\chi^{(t,m)}$ of size $\lfloor nR_m \rfloor$ from s_m to t that satisfies $\chi^{(t,m)} \leq \lceil nx^{(t,m)} \rceil$. Using a conformal decomposition, the flow $\chi^{(t,m)}$ can be decomposed into $\lfloor nR_m \rfloor$ forward path flows of size one for each of the source processes $X_{m,1}, \dots, X_{m,\lfloor nR_m \rfloor}$. We have now reduced the general case to the special case where all the source processes are all of unit rate and the underlying multicommodity flows are forward path flows of size one. Therefore, using linear coding with symbols from a finite field, we have sets of global coding vectors on each link (i, j)

with size satisfying

$$|B(i, j)| \leq \sum_{C \in \mathcal{C}(j)} \max_{t \in T} \left\{ \sum_{k \in C} \lceil nx_{ij}^{(t,m)} \rceil \right\}. \quad (12)$$

The rate achieved by such coding is $\lfloor nR_m \rfloor/n$, which differs from R_m by no more than $1/n$, and the flow placed on each link (i, j) can be made as low as $|B(i, j)|/n$, which exceeds z_{ij} by no more than M/n . We obtain the desired result by taking n arbitrarily large. $\square$

Note that because our formulation implicitly assumes that coding delay can be made arbitrarily large, we have avoided the pathologies that arise when the coding delay is fixed, for example, those illustrated by examples in [19, 17, 18]. Indeed, the resolutions to the examples given in the former two papers that use linear coding over longer blocks fall nicely into our formulation and can be obtained as solutions of problem (8).

5. Conclusion

This paper has considered the problem of finding minimum-cost subgraphs for network coding. An important point to note is that this problem is essentially decoupled from the coding problem; that is, we can first determine the amount of flow that must be placed on each edge, then determine the content of the flows. Setting up network connections without network coding can be thought of as consisting of the same pair of problems, except that the coding problem is trivial — it is obvious what the content of the flows should be.

One of our main results is that, while minimum-cost multicast problems without network coding are very difficult except in the special cases of unicast and broadcast, finding minimum-cost subgraphs for single multicast connections with network coding can be posed as a linear optimization problem. For the case of multiple multicast connections, we are only able to give a partial solution: We specify a linear optimization problem that yields a solution of no greater cost than any solution without network coding and that we suspect can potentially be substantially better.

In a separate paper [16], we show that the optimization problem for finding minimum-cost subgraphs for single multicast connections in fact admits a decentralized solution for any convex cost function, which, when coupled with decentralized schemes for constructing network codes [10, 11, 6], forms a fully decentralized approach for achieving minimum-cost multicast. We have not yet developed a distributed algorithm for solving the optimization problem for multiple multicast connections, though it is a clear avenue for future work.

References

- [1] D. Acemoglu and A. Ozdaglar. Flow control, routing, and performance from service provider viewpoint. Technical Report LIDS-WP-1696, MIT LIDS, Dec. 2003.
- [2] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung. Network information flow. *IEEE Trans. Inform. Theory*, 46(4):1204–1216, July 2000.
- [3] A. Ahluwalia, E. Modiano, and L. Shu. On the complexity and distributed construction of energy-efficient broadcast trees in static ad hoc wireless networks. In *Proc. 2002 Conference on Information Sciences and Systems (CISS 2002)*, Mar. 2002.
- [4] K. Bharath-Kumar and J. M. Jaffe. Routing to multiple destinations in computer networks. *IEEE Trans. Commun.*, 31(3):343–351, Mar. 1983.
- [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)*, pages 192–200, 1998.
- [6] P. A. Chou, Y. Wu, and K. Jain. Practical network coding. In *Proc. 41st Annual Allerton Conference on Communication, Control and Computing*, Oct. 2003.
- [7] R. Dougherty, C. Freiling, and K. Zeger. Insufficiency of linear coding in network information flow. Submitted to *IEEE Trans. Inform. Theory*, Feb. 2004.
- [8] R. M. Dudley. *Real Analysis and Probability*, volume 74 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2002.
- [9] M. Effros. Distortion-rate bounds for fixed- and variable-rate multiresolution source codes. *IEEE Trans. Inform. Theory*, 45(6):1887–1910, Sept. 1999.
- [10] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong. Toward a random operation of networks. Submitted to *IEEE Trans. Inform. Theory*.
- [11] T. Ho, M. Médard, J. Shi, M. Effros, and D. R. Karger. On randomized network coding. In *Proc. 41st Annual Allerton Conference on Communication, Control and Computing*, Oct. 2003.
- [12] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen. Polynomial time algorithms for multicast network code construction. Submitted to *IEEE Trans. Inform. Theory*, July 2003.
- [13] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. To appear in *Math. Oper. Res.*
- [14] R. Koetter and M. Médard. An algebraic approach to network coding. *IEEE/ACM Trans. Networking*, 11(5):782–795, Oct. 2003.
- [15] S.-Y. R. Li, R. W. Yeung, and N. Cai. Linear network coding. *IEEE Trans. Inform. Theory*, 49(2):371–381, Feb. 2003.
- [16] 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. Submitted to *IEEE Infocom 2005*.
- [17] M. Médard, M. Effros, D. Karger, and T. Ho. On coding for non-multicast networks. In *Proc. 41st Annual Allerton Conference on Communication, Control and Computing*, Oct. 2003.
- [18] A. Rasala Lehman and E. Lehman. Complexity classification of network information flow problems. In *Proc. 41st Annual Allerton Conference on Communication, Control and Computing*, Oct. 2003.
- [19] S. Riis. Linear versus non-linear Boolean functions in network flow. In *Proc. 2004 Conference on Information Sciences and Systems (CISS 2004)*, Mar. 2004.
- [20] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. *Telecomm. Policy*, 20(3):183–201, 1996.
- [21] L. Song, R. W. Yeung, and N. Cai. Zero-error network coding for acyclic networks. *IEEE Trans. Inform. Theory*, 49(12):3129–3139, Dec. 2003.
- [22] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. Energy-efficient broadcast and multicast trees in wireless networks. *Mobile Networks and Applications*, 7:481–492, 2002.
- [23] L. Zosin and S. Khuller. On directed Steiner trees. In *Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2002)*, pages 59–63, 2002.