

Minimum-Energy Multicast in Mobile Ad Hoc Networks Using Network Coding

Yunnan Wu, *Student Member, IEEE*, Philip A. Chou, *Fellow, IEEE*, and Sun-Yuan Kung, *Fellow, IEEE*

Abstract—The minimum energy required to transmit one bit of information through a network characterizes the most economical way to communicate in a network. In this paper, we show that, under a layered model of wireless networks, the minimum energy-per-bit for multicasting in a mobile ad hoc network can be found by a linear program; the minimum energy-per-bit can be attained by performing network coding. Compared with conventional routing solutions, network coding not only allows a potentially lower energy-per-bit to be achieved, but also enables the optimal solution to be found in polynomial time, in sharp contrast with the NP-hardness of constructing the minimum-energy multicast tree as the optimal routing solution. We further show that the minimum energy multicast formulation is equivalent to a cost minimization with linear edge-based pricing, where the edge prices are the energy-per-bits of the corresponding physical broadcast links. This paper also investigates minimum energy multicasting with routing. Due to the linearity of the pricing scheme, the minimum energy-per-bit for routing is achievable by using a single distribution tree. A characterization of the admissible rate region for routing with a single tree is presented. The minimum energy-per-bit for multicasting with routing is found by an integer linear program. We show that the relaxation of this integer linear program, studied earlier in the Steiner tree literature, can now be interpreted as the optimization for minimum energy multicasting with network coding. In short, this paper presents a unifying study of minimum energy multicasting with network coding and routing.

Index Terms—Energy efficiency, mobility, multicast, network coding, routing, Steiner tree, wireless ad hoc networks.

I. INTRODUCTION

IN THIS paper, we consider the problem of minimum-energy information multicast, namely transmitting common information from a source node s to a set of destination nodes T with the minimum amount of total consumed energy per information bit, in a mobile ad hoc network (MANET).

Several previous works, e.g., [1]–[3], aim at finding the *minimum-energy multicast tree*. Suppose a sender s wants to transmit a message to destinations T . Assuming nodes in the network can only *route*, i.e., replicate and forward, received messages, the multicasting will take place in a sequence of steps. In each step, a node, having received the message so far, forwards the message to some neighbors at a certain power

Paper approved by R. Fantacci, the Editor for Wireless Networks and Systems of the IEEE Communications Society. Manuscript received March 17, 2004; revised November 20, 2004. This paper was presented in part at the IEEE Information Theory Workshop, San Antonio, TX, October 2004.

Y. Wu and P. A. Chou are with Microsoft Research, Redmond, WA 98052-6399 USA (e-mail: yunnanwu@microsoft.com; pachou@microsoft.com).

S.-Y. Kung is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: kung@princeton.edu).

Digital Object Identifier 10.1109/TCOMM.2005.857148




Fig. 1. (a) Example wireless ad hoc network. The locations of the nodes have been marked with dots. Assume each node is equipped with a transmitter operating at a fixed transmission range, which is just sufficient to reach its lateral neighbors, but not the diagonal ones. Under this setting, each physical-layer transmission consumes a unit amount of energy. It is easy to see that the minimum amount of energy required to deliver one message from s to $\{t_1, t_2\}$ is five (transmissions) using the conventional routing approach. One such solution is as follows. First, s broadcasts the message to $\{a, b\}$ using one transmission. Next, a forwards it to c , c forwards to t_1 , b forwards to d , and d forwards to t_2 . (b) Minimum-energy multicast with network coding on this example network. Suppose s has two messages, x_1 and x_2 . First, x_1 is delivered to t_1 with three transmissions and x_2 is delivered to t_2 with three transmissions. Next, t_1 transmits x_1 to e and t_2 transmits x_2 to e . The critical step occurs at e , which broadcasts the XOR result of two messages $x_1 \oplus x_2$ to t_1 and t_2 , consuming only one transmission. Each of the two destinations can recover both x_1 and x_2 by solving a simple linear system of equations.

level. Note that, due to the broadcast nature of radio transmissions, a single transmission by a certain transmitter may successfully reach multiple (neighboring) nodes; this physical-layer broadcast property was called the *wireless multicast advantage* in [1]. Then, the problem is to find a set of relaying nodes and their respective power levels such that all nodes in T receive the message, whereby the total energy expenditure for the task is minimized. Under this formulation, it can be easily concluded that the optimal forwarding scheme should be based on a tree structure. However, the problem of constructing a minimum-energy multicast tree in a wireless ad hoc network is NP-hard; see, e.g., [1]–[3].

We now use a simple example to illustrate that there is room to improve upon the conventional formulation to achieve a potentially lower energy-per-bit. In Fig. 1, the minimum amount of energy required to deliver one message from s to $\{t_1, t_2\}$ is five (transmissions) using the conventional routing approach, whereas two messages can be delivered with nine transmissions using network coding. For details, please refer to Fig. 1.

The lesson drawn from this example is that it is, in general, suboptimal to require the nodes in the network to only replicate and forward, i.e., route, information received. In other words, *network coding* could bring in unique advantages over *routing*. Network coding generalizes the traditional routing

paradigm by allowing nodes to perform arbitrary operations on the information received to generate the output. Historically, throughput gain has been the primary motivation for network coding. In their pioneering work [5], Ahlswede *et al.* gave a simple example network to show that network coding can potentially achieve a higher throughput than routing solutions. More generally, they showed that network coding can achieve the maximum multicast rate.

Network coding is highly applicable to real packet networks. For example, Chou *et al.* [6] presented a prototype system for practical network coding in packet networks, using distributed random linear network coding with buffering. The system achieves throughput close to capacity with low delay and is robust to random packet loss and delay as well as to changes to network topology or capacity. Distributed random network coding were also investigated by Ho *et al.* [7]–[9]. Empirical comparisons of the throughput achievable by routing and network coding have been reported in [10] for static graphs and in [11] for dynamically varying graphs.

An increasingly important application domain of network coding is MANETs. By having random mixture packets self-orchestrate multiple paths, network coding offers built-in error protection and adaptivity to topology changes due to joins, leaves, node or link failures, or congestion; by employing a flooding-type delivery, network coding can be implemented in a distributed fashion easily, whereas the creation and maintenance of distribution trees incurs notable signaling overhead, given the dynamic environment. These properties render network coding potentially useful for unicasting and multicasting in MANETs.

In this study, we explore the advantage of network coding in *economically using network resources*, more specifically, the energy consumption in a MANET. The example in Fig. 1 demonstrates that network coding can lead to solutions that are more economic in using network resources than routing solutions. More generally, we show that, under a layered model of wireless networks, the minimum energy-per-bit for multicasting can be found in polynomial time via a linear program. The linear program outputs an optimal allocation of bit-rate resources on the links; using the allocated link bit rates, the minimum energy-per-bit can be attained by performing network coding, but not routing in general.

The discussions in this paper are based on a popular layered model of wireless networks (see, e.g., [12]), which is a mathematical abstraction of the well-known layered network architecture. The basic assumptions of the layered model can be explained as follows. The lower and upper layers are, respectively, abstracted as supply and demand of communication resources. The interface between the supply and demand is a network of noiseless channels with rate limits, which can be described as $G = (V, E, \mathbf{c})$, where V and E are sets of vertices and edges, respectively, and \mathbf{c} is a length- $|E|$ vector assigning to each edge $e \in E$ a bit-rate limit $c(e)$. Via communication mechanisms such as scheduling and power control, the lower layers can supply a set \mathcal{G} of *realizable graphs*. Using a realizable graph $G \in \mathcal{G}$ as the available communication resources, the network layer coordinates the information flow from the sources to the destinations such that certain end-to-end throughput is

achieved. From an information-theoretic perspective, such decoupling of supply and demand is essentially the separation between channel coding and network coding. Although it may lead to suboptimality, its simplicity facilitates the analysis and invites further engineering insights.

Each realizable graph (V, E, \mathbf{c}) provides certain end-to-end throughput and has an associated power consumption that represents the cost in supplying the bit-rate resources. Thus finding the minimum energy-per-bit boils down to identifying a realizable graph such that the ratio of the power consumption achieved to the multicast throughput achieved is minimized. We shall show this can be carried out via a linear program that jointly optimizes the supply side and the demand side.

II. REALIZABLE GRAPHS FOR MANETS

In this section, we discuss the structure of realizable graphs for MANETs. It is nontrivial to identify the realizable graphs for a wireless network, due to the unique lower layer characteristics of wireless communications. The topology of a wireless ad hoc network, as a composition of individual transmission links, is dynamic. Dynamism arises as a result of control actions such as arranging different transmitters at different times and adjusting the transmission power, as well as random variations in the channel strengths of the wireless links due to fading and node mobility.

We shall start by discussing a static wireless ad hoc network where the channel conditions stay the same during a communication session and then move on to discuss a MANET where the channel conditions evolve over time. This set of realizable graphs represents all possible supplies of bit-rate resources arising from power control and scheduling in the physical and medium access layers (under certain simplifying assumptions about the physical layer). We show that, for the problem of minimum energy multicast, only a small subset of realizable graphs needs to be considered.

A. Static Wireless Ad Hoc Network

A wireless ad hoc network can operate in many different *physical states*, where each physical state represents a “snapshot” of all nodes in the physical layer, such as which nodes are transmitting, what transmitting powers are being used, and what the channel conditions are. A physical state may support a collection of concurrent *links*, which are assumed to be point-to-multipoint in general. Let V_0 denote the set of nodes in the network. A *link* can be described as $u \xrightarrow{c} Y_u$, where $u \in V_0$ is the transmitter, $Y_u \subseteq V_0$ is its receiver set, and c is the associated bit rate in a reliable communication.

Each collection of links supported by a certain physical state corresponds to an *elementary (realizable) graph*. Loosely speaking, an elementary graph refers to a graph that can be directly realized in the physical layer. For example, Fig. 2 shows two physical states and the corresponding elementary graphs. Fig. 2 corresponds to the special case of omnidirectional transmissions. In a first state shown in Fig. 2, only node s is transmitting, and the transmission power is just enough to reach a and b . This supports the link $s \xrightarrow{1.0} \{a, b\}$. The corresponding elementary graph is shown as G_1 . In a second state, only node




Fig. 2. Two example elementary graphs G_1 and G_2 . G_1 corresponds to the physical state where s is transmitting at a power just enough to reach a and b . G_2 corresponds to the physical state where s is transmitting at a power just enough to reach a, b, c , and d .

s is transmitting, and the transmission power is just enough to reach a, b, c , and d . This supports the link $s \xrightarrow{1,0} \{a, b, c, d\}$. The corresponding elementary graph is shown as G_2 .

By timesharing among different physical states, it is possible to achieve any convex combination of the elementary graphs. That is, if λ_k is the relative share of time for the elementary graph $G_k = (V_k, E_k, c_k)$, then it is possible to achieve on average the graph $G = (\cup_k V_k, \cup_k E_k, \sum_k \lambda_k c_k)$. Here the edge capacities c_k are each extended to a length- $|\cup_k E_k|$ vector c'_k in the obvious way. Denote such combinations by $G = \sum_k \lambda_k G_k$.

Let the set of all elementary graphs be \mathcal{B}_0 . The number of elementary graphs $|\mathcal{B}_0|$ generally grows exponentially with the number of network nodes. Fortunately, for the minimum-energy multicast problem, the number of *pertinent* elementary graphs or corresponding physical states is polynomial in the number of nodes, assuming that the achievable bit rate on a link is determined by the signal-to-interference-and-noise ratio (SINR) at the receiver. This is because separating interfering transmissions into different time slots improves the energy efficiency. Since we are focusing on the minimum energy-per-bit, we can restrict our attention to those physical states involving only a single transmitter. It is this fact that results in the polynomial solvability of the minimum-energy multicast problem for a MANET.

With a finite set of elementary graphs $\mathcal{B} \subseteq \mathcal{B}_0$, the set of realizable graphs is

$$\mathcal{G}(\mathcal{B}) = \left\{ G \mid G = \sum_k \lambda_k G_k, \sum_k \lambda_k \leq 1, \lambda_k \geq 0 \forall k, G_k \in \mathcal{B} \right\}$$

where the dependence on \mathcal{B} is explicitly shown. The power consumption of a composite graph $G = \sum_k \lambda_k G_k$ is $P(\sum_k \lambda_k G_k) = \sum_k \lambda_k p(G_k)$, where $p(G_k)$ is the power consumption of elementary graph G_k . The power consumption reflects one possible metric that measures the cost of providing the bit-rate resources at the physical and medium access layers.

1) *Elementary Graphs*: The modeling with elementary graphs is very general; for example, it could potentially represent directional transmissions using multiple antennas. For concreteness, we now present a specific model of elementary graphs, assuming a single-antenna wireless communication system. As discussed above, we only need to consider the physical states with a single transmitter. Since wireless networks operate in an inherent broadcast medium, we consider the modeling of physical-layer broadcast links with elementary

graphs. For simplicity, we consider broadcasting of common information from the transmitter to a set of receivers.

Let u be the transmitter and p_u be its transmitting power. Let the path gain from node $i \in V_0$ to node $j \in V_0$ be a_{ij} . Assume that all receivers have the same noise level σ^2 . The signal-to-noise ratio (SNR) at a node j receiving information from $u \in V_0$ admits the following expression:

$$\text{SNR}_{uj} = \frac{a_{uj} p_u}{\sigma^2}. \quad (1)$$

We adopt a simplistic physical-layer model: as long as the received SNR of a node exceeds a threshold γ , the transmission rate is set to be a common “unit” capacity. Since the associated rate c is always 1, we drop it from the notation $u \xrightarrow{c} Y_u$ and write $u \rightarrow Y_u$ instead. There may be more than one node j whose receiving SNR level exceeds the threshold. Thus, each broadcast link is capable of providing common information to all of these qualifying receivers while consuming some power. To reflect the broadcasting of common information, we propose the following graph model for broadcast links. For each link $u \rightarrow Y_u$, we add to the associated elementary graph a *distinct*¹ virtual vertex (e.g., u') and unit capacity edges uu' , $u'v$, $v \in Y_u$. Two examples are given in Fig. 2 as G_1 and G_2 , in which s' and s'' are the virtual vertices introduced. The virtual vertex plays the role of an artificial bottleneck that constrains the rate of new information going out of the transmitter. Since these virtual vertices do not physically exist, they can only perform routing, instead of arbitrary network coding. In other words, for a virtual vertex u' , information flowing on its outgoing edges are not allowed to be general functions of information flowing on the incoming edges. Fortunately, such a restriction does not compromise the throughput at all.²

In short, an SNR-based approach is used to model elementary graphs. We consider elementary graphs comprising a single broadcast link each. Let $u \rightarrow Y_u = \{v_1, \dots, v_n\}$ be the broadcast link in an elementary graph G_k . This broadcast link is represented by a tree-like structure with a virtual vertex u' . In G_k , we have

$$c_k(uu') = c_k(u'v_1) = \dots = c_k(u'v_n) = 1 \quad (2)$$

and other edges have zero capacity.

2) *Pertinent Elementary Graphs*: By adjusting the transmitting power, the “reach” of a transmitter, i.e., the set of receivers that fall in the transmission range $\{j \in V_0 \mid \text{SNR}_{uj} \geq \gamma\}$, can be accordingly adjusted. Thus, for the problem of minimum-energy multicast, the pertinent elementary graphs will be those with a single transmitter operating at different reaches. The set of reaches of a node will depend on whether the node can select its power from a finite set of values or from a continuous interval. We call the former *discrete power control* and the

¹To see why it can be problematic to treat the virtual vertices associated with a same transmitter as the same in different elementary graphs, please refer to [13].

²To see this, consider an arbitrary network coding solution \mathcal{S} . Modify this solution \mathcal{S} by letting each virtual vertex u' forward whatever information flowing on uu' to its outgoing edges. Then, the receiver nodes in Y_u can produce the information flowing on its outgoing edges as prescribed in \mathcal{S} . This yields another valid coding solution with the same throughput. For a stronger version of this statement, please see Appendix B.

latter *continuous power control*. With the discrete power control model, the total number of pertinent elementary graphs is less than or equal to $Q|V_0|$, where $|V_0|$ is the number of nodes and Q is the number of available transmission power levels at each node. With the continuous power control model, the total number of pertinent elementary graphs is less than or equal to $|V_0| \cdot (|V_0| - 1)$, since the pertinent transmitting power just needs to assume values in a set comprising the minimum power to reach every other node.

B. Mobile Ad Hoc Network

A MANET may experience variations in the channel strength of individual wireless links during a communication session. On a short time scale, e.g., on the order of milliseconds or less, the communication links undergo fading due to the constructive and destructive superpositions of multiple signal paths. On a long time scale, e.g., on the order of seconds or more, user mobility introduces dramatic changes to channel characteristics by varying the relative geographic distances of the communication devices. In this subsection, we specifically consider node mobility as one type of such physically induced dynamism, for the convenience of illustration.

With mobility, the locations of the nodes evolve over time. Suppose the locations of nodes evolve over time as L discrete maps M_l , $l = 1, \dots, L$, where each map M_l refers to a configuration of node locations, lasting for a duration of t_l s. In other words, we assume that the nodes stay at the locations M_l for t_l s and then instantaneously jump to new locations M_{l+1} at the beginning of next interval. In the following, we assume an exact knowledge (or reasonable prediction) of future locations of the nodes. Consequently, a good “motion estimator” is needed to realize the potential advantage offered by the time diversity.

We model mobility by introducing an additional dimension (time) into the graph. The graph for a mobile network is then comprised of several sequentially concatenated layers, with each layer corresponding to a graph for one map. We call such a graph a *time-lined graph*. Each node $v \in V_0$ is now “expanded” into L vertices, v^1, \dots, v^L , one in each layer. The edges in a time-lined graph consist of two type of edges: intralayer edges and interlayer edges. The intralayer edges are edges which originally belong to the graph for each layer. Each of the interlayer edges go from v^l to v^{l+1} and has infinite capacity, modeling information buffering at node $v \in V_0$. These edges are unidirectional because of causality. Hence, we represent a (realizable) time-lined graph as

$$G = G^1 \otimes G^2 \dots \otimes G^L \quad (3)$$

$$G^l = \sum_k \lambda_k^l G_k^l, \quad G_k^l \in \mathcal{B}_l \quad (4)$$

$$\sum_k \lambda_k^l \leq \frac{t_l}{\sum_{l=1}^L t_l}, \quad l = 1, \dots, L \quad (5)$$

$$\lambda_k^l \geq 0 \quad \forall l, k \quad (6)$$

where G^l represents a graph in the l th layer, the notation $G^l \otimes G^{l+1}$ characterizes the interlayer edges introduced, and λ_k^l characterizes the time-sharing proportion for G_k^l , the k th

elementary graph out of a collection \mathcal{B}_l for the l th layer. The power consumption of a time-lined graph is

$$P \left(G^1 \otimes G^2 \dots \otimes G^L \right) = \sum_{l=1}^L \sum_k \lambda_k^l P(G_k^l). \quad (7)$$

III. MINIMUM-ENERGY MULTICAST

In this section, let us focus on a fixed network for clarity of presentation. The extension to the case of a MANET is conceptually the same. Let us denote the set of pertinent elementary graphs by \mathcal{B} . Each elementary graph $G_k \in \mathcal{B}$ contains only a single broadcast link from the discussions in Section II-A. Recall that V_0 denotes the set of physical nodes. Let V_1 denote the set of virtual vertices introduced to model the broadcast links in \mathcal{B} . We use V to denote the enlarged vertex set V_0 and V_1 , and E to denote the union of the edges for $\mathcal{G}(\mathcal{B})$.

Each realizable graph $G \in \mathcal{G}(\mathcal{B}_0)$ represents a feasible supply of bit-rate resources on the links. It has an associated power consumption $P(G)$ watts. Using the bit-rate resources G , certain multicast throughput from s to T can be achieved. Let $C(G)$ denote the *multicast capacity* (from s to T) in G , which is defined as the maximum achievable multicast rate (b/s). Then the minimum energy-per-bit for multicasting from s to T based on a given graph G is

$$\frac{P(G)}{C(G)} \quad (\text{J/b}). \quad (8)$$

Therefore, the minimum energy-per-bit under the layered model is equal to

$$\mathcal{E}^* \equiv \min_{G \in \mathcal{G}(\mathcal{B}_0)} \frac{P(G)}{C(G)} = \min_{G \in \mathcal{G}(\mathcal{B})} \frac{P(G)}{C(G)}. \quad (9)$$

The last equality in (9) follows from the discussions in Section II.

This optimization can be done by a linear program because of two key observations. On the supply side, the set of \mathbf{c} such that $(V, E, \mathbf{c}) \in \mathcal{G}(\mathcal{B})$ can be characterized by a set of linear inequalities. On the demand side, the set of \mathbf{c} that can provide a given multicast rate r can be characterized by a set of linear inequalities.

A. Flow

Given (V, E) , a source node $s \in V$, and a destination node $t \in V$, an s - t -flow is a length- $|E|$ vector \mathbf{f} satisfying the following constraints:

$$\begin{aligned} f(vw) &\geq 0 \quad \forall vw \in E \\ \sum_{w \in V: vw \in E} f(vw) - \sum_{u \in V: uw \in E} f(uw) &= 0 \quad \forall v \in V \setminus \{s, t\} \\ \sum_{w \in V: sw \in E} f(sw) - \sum_{u \in V: us \in E} f(us) &\geq 0. \end{aligned}$$

Here $f(vw)$ denotes the flow on edge vw . Let $\mathcal{F}_{s,t}$ denote the set of s - t flows \mathbf{f} satisfying the above constraints. The *value of the flow* is

$$|\mathbf{f}| \equiv \sum_{w \in V: sw \in E} f(sw) - \sum_{u \in V: us \in E} f(us). \quad (10)$$

A flow from a source to a destination can be decomposed into a sum of several *path-flows* and *cycle-flows*. It turns out that the

$$\mathcal{A}(r) = \left\{ \mathbf{c} \in \mathbb{R}^{|E|} \mid \exists \mathbf{f}_t \in \mathcal{F}_{s,t}, |\mathbf{f}_t| = r, \text{ for each } t \in T, \text{ such that } \mathbf{c} \geq \max_{t \in T} \mathbf{f}_t \right\} \quad (\text{a})$$

$$\mathbf{c} \in \mathcal{A}(r) \iff \min_{t \in T} \rho_{s,t}(\mathbf{c}) \geq r \iff \exists \mathbf{f}_t, \text{ for each } t \in T \text{ satisfying}$$

$$\begin{aligned} 0 &\leq f_t(vw) \leq c(vw), \quad \forall vw \in E, \forall t \in T \\ \sum_{w \in V: vw \in E} f_t(vw) - \sum_{u \in V: uv \in E} f_t(uv) &= 0, \quad \forall v \in V \setminus \{s, t\}, \forall t \in T \\ \sum_{w \in V: sw \in E} f_t(sw) - \sum_{u \in V: us \in E} f_t(us) &= r, \quad \forall t \in T. \end{aligned} \quad (\text{b})$$

Fig. 3. (a) Admissible rate region for multicasting from s to T at a rate r . The notations “ \geq ” and “max” are in the element-wise sense. (b) System of linear inequalities, which is feasible if and only if $\mathbf{c} \in \mathcal{A}(r)$.

cycle-flows can be eliminated without affecting the value of the flow. Each path-flow corresponds to a path from the source with the destination with an associated rate. The value of the flow is the sum of the rates carried by the individual path-flows. Thus, a flow prescribes a way for information to be routed from the source to the destination along parallel paths; the communication rate achieved by such a routing scheme is the value of the flow.

In order for a flow to correspond to a feasible routing arrangement, we often need to enforce that the assigned flow on an edge fits in the available bit-rate resource

$$f(vw) \leq c(vw) \quad \forall vw \in E. \quad (11)$$

The Max-Flow-Min-Cut Theorem, a fundamental result in graph theory, states that, in a graph (V, E) with edge capacities \mathbf{c} , the maximum value of an s - t flow equals the minimum capacity of an s - t cut

$$\rho_{s,t}(\mathbf{c}) \equiv \min_{U:s \in U, t \in \bar{U}} \sum_{vw \in E: v \in U, w \in \bar{U}} c_{vw}. \quad (12)$$

An s - t cut (U, \bar{U}) refers to a partition of the nodes $V = U + \bar{U}$ with $s \in U, t \in \bar{U}$. The capacity of the cut (U, \bar{U}) refers to the sum of the edge capacities for edges going from U to \bar{U} . By an information theoretic argument, the capacity of any s - t cut is an upper bound on the maximum achievable rate from s to t . The Max-Flow-Min-Cut Theorem says that the maximum communication rate from s to t is $\rho_{s,t}(\mathbf{c})$, since it can be achieved by routing along a maximum s - t flow.

B. MAX of Flows

In a graph (V, E) with edge capacities \mathbf{c} , since the capacity of any s - t cut is an upper bound on the achievable transmission rate from s to t , the quantity

$$\min_{t \in T} \rho_{s,t}(\mathbf{c}) \quad (13)$$

is an upper bound on the multicast capacity. Ahlswede *et al.* [5] established that the multicast capacity is (13) by showing that it can be achieved via network coding. This result admits an alternative interpretation in terms of required resources to provide a certain multicast rate. Specifically, consider the end-to-end communication demand of multicasting information from s to T at a rate r in a given (V, E) . The *admissible rate region for multicasting* (from s to T at a rate r) is defined as the set of edge capacity vectors \mathbf{c} such that the rate r can be fulfilled in (V, E, \mathbf{c}) ; denote this by $\mathcal{A}(r)$.




Fig. 4. (a) s - t_1 flow \mathbf{f}_1 on graph (V, E) . (b) s - t_2 flow \mathbf{f}_2 on graph (V, E) . (c) MAX of flows $\max\{\mathbf{f}_1, \mathbf{f}_2\}$, which can provide a multicast rate of 2.0 from s to $\{t_1, t_2\}$.

Lemma 1 (Admissible Rate Region of Multicasting): Given (V, E) , a source s , and destinations T , the admissible rate region for multicasting from s to T at rate r , $\mathcal{A}(r)$, is given in Fig. 3(a). Furthermore, $\mathbf{c} \in \mathcal{A}(r)$ if and only if $\exists \mathbf{f}_t$ for each $t \in T$, satisfying the system of linear inequalities in Fig. 3(b).

In addition

$$\mathcal{A}(r) = r\mathcal{A}(1) \equiv \{r\mathbf{c} \mid \mathbf{c} \in \mathcal{A}(1)\}. \quad (14)$$

Proof: This result follows from [5] readily

$$\mathbf{c} \in \mathcal{A}(r) \quad (15)$$

$$\iff \min_{t \in T} \rho_{s,t}(\mathbf{c}) \geq r \quad (16)$$

$$\iff \rho_{s,t}(\mathbf{c}) \geq r \quad \forall t \in T \quad (17)$$

$$\iff \exists \mathbf{f}_t \in \mathcal{F}_{s,t}, |\mathbf{f}_t| = r, \mathbf{c} \geq \mathbf{f}_t. \quad (18)$$

We use the name *MAX of flows* to refer to an edge-wise maximum of s - t flows providing rate r , i.e., $\max_{t \in T} \mathbf{f}_t$ as in Fig. 3. Just as a flow is the critical structure for communicating from a source to a destination, a MAX of flows plays a fundamental role for multicasting using network coding. Fig. 4 illustrates the structure of a MAX of flows. This example comes from the classical example for network coding introduced in [5]. Fig. 4(a) shows an s - t_1 flow, which prescribes two parallel paths from s to t_1 ; similarly, Fig. 4(b) shows an s - t_2 flow. Fig. 4(c) shows the edge-wise maximum of these two flows, which is sufficient to provide a multicast rate of 2.0.

C. Minimum-Energy Multicast With Network Coding

The demand side of the system is characterized by the condition $\mathbf{c} \in \mathcal{A}(r)$, which can be expressed as a system of

linear inequalities given in Fig. 3(b). Regarding the supply side, recall from Section II that a pertinent realizable graph $G = (V, E, \mathbf{c}) \in \mathcal{G}(\mathcal{B})$ admits the following structure:

$$G = \sum_{k:G_k \in \mathcal{B}} \lambda_k G_k, \quad \sum_k \lambda_k \leq 1, \quad \lambda_k \geq 0 \quad \forall k. \quad (19)$$

Thus, the supply side of the network system, $G = (V, E, \mathbf{c}) \in \mathcal{G}(\mathcal{B})$, can also be expressed as a system of linear inequalities.

Integrating the supply side and the demand side, the minimum-energy multicast problem can be formulated as the following optimization, where $r, \mathbf{c}', \lambda'_k$, are treated as variables:

$$\mathcal{E}^* = \min \frac{\sum_{k:G_k \in \mathcal{B}} \lambda'_k p(G_k)}{r}$$

subject to

$$\begin{aligned} \mathbf{c}' &\in \mathcal{A}(r) = r\mathcal{A}(1) \\ \mathbf{c}' &= \sum_k \lambda'_k \mathbf{c}_k \\ \sum_k \lambda'_k &\leq 1, \lambda'_k \geq 0 \quad \forall k \\ r &> 0. \end{aligned} \quad (20)$$

Here G_k is represented as (V, E, \mathbf{c}_k) , where \mathbf{c}_k is a length- $|E|$ vector.

At first glance, the objective function of the above optimization is nonlinear in the variables. However, we can renormalize the above optimization to arrive at a linear program. Specifically, by a variable change $\lambda_k = \lambda'_k/r, \mathbf{c} = \mathbf{c}'/r$, we have the following.

Theorem 1 (Minimum-Energy Multicast): In a static wireless ad hoc network, let $\mathcal{B} = \{G_k\}$ denote the set of pertinent elementary graphs for the problem of minimum energy multicast, according to a (discrete or continuous) power control model. The minimum energy-per-bit for multicasting from s to T , \mathcal{E}^* , is given by the following optimization:

$$\mathcal{E}^* = \min \sum_{k:G_k \in \mathcal{B}} \lambda_k p(G_k)$$

subject to

$$\begin{aligned} \mathbf{c} &\in \mathcal{A}(1) \\ \mathbf{c} &= \sum_k \lambda_k \mathbf{c}_k \\ \lambda_k &\geq 0 \quad \forall k. \end{aligned} \quad (21)$$

This optimization becomes an explicit linear program upon expanding $\mathbf{c} \in \mathcal{A}(1)$ as in Fig. 3(b). The minimum energy-per-bit can be attained by performing network coding on a realizable graph

$$G^* \equiv \left(V, E, \frac{\mathbf{c}^*}{\sum_k \lambda_k^*} \right) \quad (22)$$

where $\{\lambda_k^*, \mathbf{c}^*\}$ is an optimal solution of (21). The realizable graph G^* provides a multicast rate $1/\sum_k \lambda_k^*$.

Remark: The main difference between (20) and (21) is the absence of the constraint on $\sum_k \lambda_k$. This is due to the observation that the constraint $\sum_k \lambda_k \leq 1/r$ can be satisfied by letting r be sufficiently small. ■

Note that there might be more than one solution achieving \mathcal{E}^* , but with different rates $1/\sum_k \lambda_k^*$. *Theorem 2* examines the rate that can be supported at the minimum energy-per-bit.

Theorem 2 (Max Rate at Min Energy-per-Bit):

- 1) Let r^* denote the reciprocal of the minimum $\sum_k \lambda_k$ value of an optimal solution of (21). More explicitly, $1/r^*$ can be found by the following linear program, whose feasible solutions are the optimal solutions of (21):

$$\frac{1}{r^*} = \min \sum_k \lambda_k$$

subject to

$$\begin{aligned} \mathbf{c} &\in \mathcal{A}(1) \\ \mathbf{c} &= \sum_k \lambda_k \mathbf{c}_k \\ \lambda_k &\geq 0 \quad \forall k \\ \sum_{k:G_k \in \mathcal{B}} \lambda_k p(G_k) &= \mathcal{E}^*. \end{aligned} \quad (23)$$

The quantity r^* bears the following interpretation:

$$r^* = \max_{G \in \mathcal{G}(\mathcal{B})} C(G), \quad \text{subject to } \frac{P(G)}{C(G)} = \mathcal{E}^*. \quad (24)$$

- 2) Note that, in general, (24) may not always be equal to the maximum multicast rate at the minimum energy-per-bit \mathcal{E}^*

$$\max_{G \in \mathcal{G}(\mathcal{B}_0)} C(G), \quad \text{subject to } \frac{P(G)}{C(G)} = \mathcal{E}^*. \quad (25)$$

However, if we assume positive path gains and continuous power control model, r^* is equal to (25), which is the maximum multicast rate at the minimum energy-per-bit \mathcal{E}^* .

Remark: The reason that (24) may be less than (25) is because it may be possible to “pack” multiple links in one elementary graph without incurring extra power consumption, e.g., when the discrete power control model is used. Then, to find the maximum rate at the minimum energy-per-bit, we need to also consider elementary graphs comprising more than one link. However, under the assumptions of part 2) above, any such packing will result in an increase in the energy-per-bit.

Proof:

- 1) Using (20), the right-hand side of (24) is equal to

$$\max r$$

subject to

$$\begin{aligned} \mathbf{c}' &\in \mathcal{A}(r) = r\mathcal{A}(1) \\ \mathbf{c}' &= \sum_k \lambda'_k \mathbf{c}_k \\ \sum_k \lambda'_k &\leq 1, \lambda'_k \geq 0 \quad \forall k \\ \sum_k \lambda'_k p(G_k) &= r\mathcal{E}^*. \end{aligned} \quad (26)$$

Then, (23) follows after a variable change $\lambda_k = \lambda'_k/r, \mathbf{c} = \mathbf{c}'/r$.

- 2) The proof hinges upon the establishment of the following result:

$$G \in \mathcal{G}(\mathcal{B}_0), \frac{P(G)}{C(G)} = \mathcal{E}^* \Rightarrow G \in \mathcal{G}(\mathcal{B}) \quad (27)$$

which is proven below by contradiction. Consider any $G = \sum_k \lambda_k G_k$, $P(G)/C(G) = \mathcal{E}^*$. Assume $\lambda_k > 0 \forall k$. Suppose the elementary graph G_1 comprises more than one link. Then, we can separate G_1 into different elementary graphs $\{G_{1i}\}$, each comprising only a single link. Since, by assumption, the path gains are all positive and the continuous power control model is used, this leads to a power reduction, i.e., $\sum_i p(G_{1i}) < p(G_1)$, which further leads to a reduction in the achievable energy-per-bit. This contradicts the minimality of \mathcal{E}^* . ■

D. Equivalence With Linear Edge-Based Pricing

In the minimum-energy multicasting formulation (21), the system cost is determined from the lower layers as $\sum_k \lambda_k p(G_k)$. In contrast, in graph theory, cost-minimization formulations on a graph are typically based on a more abstract and simpler linear edge-based pricing model. In such a model, the optimization objective assumes the form

$$\min \boldsymbol{\varepsilon}^T \mathbf{c} \quad (28)$$

where $\boldsymbol{\varepsilon}$ is a length- $|E|$ nonnegative vector specifying the per unit cost of the edges, and \mathbf{c} is another length- $|E|$ nonnegative vector specifying the usage on the edges. We now show it is in fact possible to cast (21) into a cost minimization with linear edge-based pricing.

Recall that V_1 denotes the set of virtual vertices introduced to model broadcast links. Thus, $V = V_0 + V_1$. Accordingly, the edge set E can be partitioned into two subsets, comprising, respectively, the edges originating from V_0 (uu' -type) and the edges originating from V_1 ($u'v_j$ -type).

Theorem 3 (Linear Edge-Based Pricing Formulation): The characterization of minimum-energy multicasting in *Theorem 1* is equivalent to a cost minimization with linear edge-based pricing. More exactly, (21) can be simplified into

$$\mathcal{E}^* = \mathcal{E}_{\text{coding}}^*(V, E, \boldsymbol{\varepsilon}, s, T) \equiv \min \boldsymbol{\varepsilon}^T \mathbf{c} \quad (29)$$

subject to $\mathbf{c} \in \mathcal{A}(1)$.

Here, the nonzero entries of the price vector $\boldsymbol{\varepsilon}$ are

$$\varepsilon(uu') \equiv p(G_k), \quad \text{for } uu' \in E \text{ with } u' \in V_1 \quad (30)$$

where $G_k \in \mathcal{B}$ is the elementary graph whose associated broadcast link is represented by a tree structure with virtual nose u' .

Remark: The edge price $\varepsilon(uu')$ can be interpreted as the energy-per-bit for the corresponding broadcast link, since the rate is one under the physical-layer model. Thus, the theorem says that the minimum energy-per-bit for the multicasting session is obtained by optimizing over all relaying arrangements with the link energy-per-bits being the edge prices. To gain some insights, consider the case where there is only a single destination t . Then the minimum energy-per-bit for the session is simply the total energy-per-bit along the most energy efficient s - t path. The multicasting formulation in (29) generalizes this result with a MAX of flows in lieu of an s - t path.

Proof: See Appendix A. ■

Section II-A described an SNR-based model of the elementary graphs in the physical layer. Therein, the capacity of each




Fig. 5. Elementary graph for the AWGN broadcast channel.

edge is either 1 or 0; the power of a broadcast link is such that the received SNR exceeds the threshold. We now reexamine the broadcast links with an information theoretic model. Indeed, for such a model, the minimum-energy multicast formulation can again be reduced to a linear edge-based pricing formulation where the edge prices are the energy-per-bits of the corresponding physical broadcast links.

Let the transmitter be u . Sort the remaining nodes in the network according to the decreasing order of the path gain; suppose they are labeled as $1, \dots, N$ and $a_{u1} > a_{u2} > \dots > a_{uN}$. Assume a discrete time real-valued additive white Gaussian noise (AWGN) reception model at each node k , $Y_i = \sqrt{a_{ui}}X + W_i$, where W_i is Gaussian with variance σ^2 , $\mathcal{N}(0, \sigma^2)$. The input X is constrained to satisfy $E[|X|^2] \leq p_u$. Then, from known results on broadcast channels [14]–[17], we observe that the associated elementary graphs can be characterized by Fig. 5 with edge capacities

$$c_u(uu'_i) = c_u(u'_i1) = \dots = c_u(u'_i i) = \frac{1}{2} \log_2 \left(1 + \frac{\pi_{ui} p_u}{\sum_{j < i} \pi_{uj} p_u + \frac{\sigma^2}{a_{ui}}} \right), \quad (31)$$

$i = 1, \dots, N$

where $\pi_{ui} \geq 0$, $\sum_{i=1}^N \pi_{ui} = 1$.

Letting \mathcal{B} be the family of elementary graphs parameterized by variables $\{\pi_{uk}\}$ and p_u , the minimum (infimum) energy-per-bit \mathcal{E}^* admits a nonlinear optimization formulation that extends (20). It can be shown that

$$\mathcal{E}^* = \mathcal{E}_{\text{coding}}^*(V, E, \boldsymbol{\varepsilon}, s, T) \quad (32)$$

where the nonzero entries of the price vector $\boldsymbol{\varepsilon}$ are

$$\varepsilon(uu'_i) \equiv \frac{2\sigma^2 \ln 2}{a_{ui}} \quad \forall u'_i \in V_1 \quad (33)$$

where V_1 denotes the set of virtual vertices u'_i , as illustrated in Fig. 5. Note that the edge price $\varepsilon(uu'_i)$ again bears the interpretation as the link energy-per-bit; it is the minimum energy-per-bit for the broadcast channel where u is broadcasting common information to nodes $1, \dots, i$, i.e.,

$$\varepsilon(uu'_i) = \inf_{p > 0} \frac{p}{\frac{1}{2} \log_2 \left(1 + \frac{p}{\frac{\sigma^2}{a_{ui}}} \right)}. \quad (34)$$

The derivations are based on the first-order Taylor expansion of (31). Due to space limitation, the details are deferred to a future publication; see also [13].

$$\begin{aligned} \mathcal{A}_{\text{routing}}(r) &= \left\{ \mathbf{c} \in \mathbb{R}^{|E|} \mid \mathbf{c} \geq \sum_j \mu_j \mathbf{g}_j, \sum_j \mu_j = r, \mu_j \geq 0, \forall j, \text{ where } \mathbf{g}_j \in \{0, 1\}^{|E|} \text{ corresponds to a Steiner tree} \right\} \\ \mathcal{A}_{\Sigma}(r) &= \left\{ \mathbf{c} \in \mathbb{R}^{|E|} \mid \exists \mathbf{f}_t \in \mathcal{F}_{s,t}, |\mathbf{f}_t| = r, \text{ for each } t \in T, \text{ such that } \mathbf{c} \geq \sum_{t \in T} \mathbf{f}_t \right\} \\ \mathcal{A}_{\text{singletree}}(r) &= \left\{ \mathbf{c} \in \mathbb{R}^{|E|} \mid \mathbf{c} \geq r \mathbf{g}, \text{ for some } \mathbf{g} \in \mathcal{A}(1) \cap \{0, 1\}^{|E|} \right\} \end{aligned}$$

Fig. 6. Admissible rate region for multicasting with routing $\mathcal{A}_{\text{routing}}(r)$ and its two inner bounds $\mathcal{A}_{\Sigma}(r)$, $\mathcal{A}_{\text{singletree}}(r)$.

IV. MINIMUM-ENERGY MULTICAST WITH ROUTING

In order to fully understand the advantage of network coding over traditional routing, we now examine minimum-energy multicasting via routing-based approaches.

Define the admissible rate region with routing (for multicasting from s to T at a rate r) as the set of \mathbf{c} such that a multicast rate r can be achieved on (V, E, \mathbf{c}) using routing. Denote this region by $\mathcal{A}_{\text{routing}}(r)$. Clearly, $\mathcal{A}_{\text{routing}}(r) \subseteq \mathcal{A}(r)$. Previous work [5] gave an example showing that, in general, $\mathcal{A}_{\text{routing}}(r) \neq \mathcal{A}(r)$.

A general routing solution makes use of multiple distribution trees, each connecting s to T ; such a distribution tree is also called a *Steiner tree* in graph theory. As a result, $\mathcal{A}_{\text{routing}}(r)$ can be characterized as in Fig. 6; this “tree-packing” characterization follows from [18]. However, this formulation involves the set of all Steiner trees, whose size can grow exponentially with the network size. In the following, we first discuss an inner bound of $\mathcal{A}_{\text{routing}}(r)$, $\mathcal{A}_{\Sigma}(r)$, which can be characterized as a system of linear inequalities. Then, we show that, for the minimum-energy multicast problem, we only need to optimize over a subset $\mathcal{A}_{\text{singletree}}(r)$, which can be characterized by a system of linear inequalities plus some integer constraints. This enables the minimum energy-per-bit for routing to be found by an integer linear program.

A. SUM of $|T|$ Flows Inner-Bound

Replacing the maximum operation in the expression of $\mathcal{A}(r)$ by the sum operation, we get an inner bound of $\mathcal{A}_{\text{routing}}(r)$, $\mathcal{A}_{\Sigma}(r)$, in Fig. 6. The region $\mathcal{A}_{\Sigma}(r)$ can be characterized by a system of linear inequalities. However, $\mathcal{A}_{\Sigma}(r)$ is only a subset of $\mathcal{A}_{\text{routing}}(r)$. The reason is as follows. The SUM of flows formulation $\mathcal{A}_{\Sigma}(r)$ allows dedicated resources \mathbf{f}_t to be allocated to the destinations, which can be used to stream distinct data to the destinations. In other words, the SUM of flows solution, also called *multicommodity flow* in the literature (see, e.g., [19]), treats the multicast session as multiple separate unicast sessions. This treatment could be unnecessarily wasteful as it has not taken into account the possibility of sharing resources among different destinations.

B. Admissible Rate Region With a Single Tree and Minimum Energy-per-Bit With Routing

Given a directed graph (V, E) with nonnegative edge prices $\boldsymbol{\varepsilon}$, a source s , and destinations T , the minimum cost to provide a unit multicast rate by routing is

$$\mathcal{E}_{\text{routing}}^*(V, E, \boldsymbol{\varepsilon}, s, T) = \min \boldsymbol{\varepsilon}^T \mathbf{c}, \text{ subject to } \mathbf{c} \in \mathcal{A}_{\text{routing}}(1). \quad (35)$$

Theorem 4 (Minimum-Cost Multicasting With Routing):

- 1) Let $\mathcal{A}_{\text{singletree}}(r)$ be the set of \mathbf{c} such that, in (V, E, \mathbf{c}) , rate r is achievable by routing with a single Steiner tree. Then

$$\mathcal{E}_{\text{routing}}^*(V, E, \boldsymbol{\varepsilon}, s, T) = \min \boldsymbol{\varepsilon}^T \mathbf{c}, \text{ subject to } \mathbf{c} \in \mathcal{A}_{\text{singletree}}(1). \quad (36)$$

- 2) The region $\mathcal{A}_{\text{singletree}}(r)$ can be characterized as

$$\left\{ \mathbf{c} \in \mathbb{R}^{|E|} \mid \mathbf{c} \geq r \mathbf{g}, \text{ for some } \mathbf{g} \in \mathcal{A}(1) \cap \{0, 1\}^{|E|} \right\}. \quad (37)$$

- 3) The minimum cost-per-bit for multicasting with routing is given by the following integer linear program:

$$\begin{aligned} \mathcal{E}_{\text{routing}}^*(V, E, \boldsymbol{\varepsilon}, s, T) &= \min \boldsymbol{\varepsilon}^T \mathbf{g}, \\ &\text{subject to } \mathbf{g} \in \mathcal{A}(1) \cap \{0, 1\}^{|E|}. \end{aligned} \quad (38)$$

A minimum-cost Steiner tree can be constructed from an optimal solution of this integer linear program.

Remark: Note that (36) does not generally hold for non-linear costs. Part 1) says that for the minimum cost problem with linear nonnegative prices, it suffices to restrict attention to routing schemes that use a single Steiner tree – a *minimum cost Steiner tree*. The minimum (cost) Steiner tree problem has been studied extensively in the literature; a survey can be found in the book [20]. This result essentially comes from the Steiner tree literature; see, e.g., [21]–[23].

Proof:

- 1) Let Ω denote the set of all length- $|E|$ binary vectors that correspond to a Steiner tree. Index these vectors as $\mathbf{g}_1, \dots, \mathbf{g}_{|\Omega|}$. Using the expression of $\mathcal{A}_{\text{routing}}(r)$ in Fig. 6, we see that (35) becomes

$$\min \sum_{j=1}^{|\Omega|} \mu_j \boldsymbol{\varepsilon}^T \mathbf{g}_j, \text{ subject to } \sum_{j=1}^{|\Omega|} \mu_j = 1, \mu_j \geq 0 \quad \forall j. \quad (39)$$

Since each Steiner tree has a nonnegative cost, the optimum value in the above optimization is attained by using only a single tree.

- 2) A Steiner tree from s to T is a union of paths from s to each $t \in T$. On the one hand, suppose in (V, E, \mathbf{c}) that rate r can be achieved by routing with a single Steiner tree. By assigning 1 to edges in the tree and 0 to others, we obtain a vector $\mathbf{g} \in \mathcal{A}(1) \cap \{0, 1\}^{|E|}$. It follows that $\mathbf{c} \in (37)$.

On the other hand, for $\mathbf{g} \in \mathcal{A}(1) \cap \{0, 1\}^{|E|}$, since (V, E, \mathbf{g}) provides a multicast rate of 1 from s to T , the set of edges with nonzero $g(vw)$ must contain a path from s to each $t \in T$ and hence a Steiner tree from s to T . Thus,

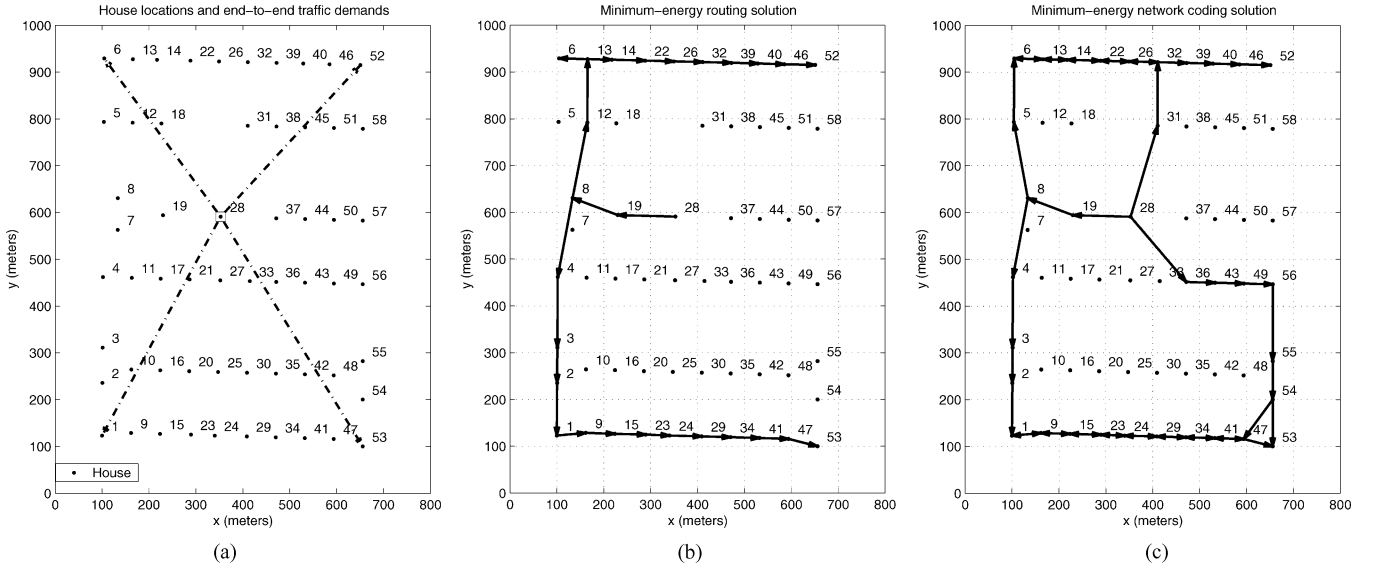


Fig. 7. (a) Example community wireless network. (b) Connectivity graph for the traffic assignment given by the minimum-energy routing solution. (c) Connectivity graph for the traffic assignment given by the minimum-energy multicast solution.

in (V, E, c) or any $c \in (37)$, rate r can be achieved by routing with a single Steiner tree.

- 3) As a result of 1) and 2), the minimum Steiner tree problem can be written as the integer linear program in (38), which was introduced in [21]–[23]. A vector $\mathbf{g} \in \mathcal{A}(1) \cap \{0, 1\}^{|E|}$ corresponds to a Steiner tree plus possibly some extra edges. Given $\mathbf{g} \in \mathcal{A}(1) \cap \{0, 1\}^{|E|}$, a Steiner tree can be explicitly reconstructed by backtracking from the destinations to s . ■

As one might expect, it was observed in the Steiner tree literature that a lower bound of the minimum cost $\mathcal{E}_{\text{routing}}^*(V, E, \boldsymbol{\varepsilon}, s, T)$ can be obtained by relaxing (discarding) the constraint $\mathbf{g} \in \{0, 1\}^{|E|}$. The relaxed linear program coincides with (29). Whereas this linear program was used as a relaxation of the minimum Steiner tree problem, it now admits a more explicit interpretation—it outputs $\mathcal{E}_{\text{coding}}^*(V, E, \boldsymbol{\varepsilon}, s, T)$, which is the minimum cost achievable with network coding for a graph (V, E) with edge prices $\boldsymbol{\varepsilon}$.

From the characterization of $\mathcal{A}_{\text{singletree}}(r)$, the minimum energy-per-bit with routing can be obtained by using $\mathbf{c} \in \mathcal{A}_{\text{singletree}}(1)$ to replace the constraint $\mathbf{c} \in \mathcal{A}(1)$ in (21). Alternatively, it can be obtained as (38) with $\boldsymbol{\varepsilon}$ defined in *Theorem 3*.

C. Integrality Gap

Whereas the minimum energy-per-bit achievable with network coding is given by the linear program in (29), the minimum energy-per-bit achievable with routing is given by the same linear program with additional integrality constraints

$$\mathbf{g} \in \{0, 1\}^{|E|}. \quad (40)$$

Let

$$\beta(V, E, \boldsymbol{\varepsilon}, s, T) \equiv \frac{\mathcal{E}_{\text{routing}}^*(V, E, \boldsymbol{\varepsilon}, s, T)}{\mathcal{E}_{\text{coding}}^*(V, E, \boldsymbol{\varepsilon}, s, T)} \quad (41)$$

denote the gain in economically using resources, offered by network coding over routing. The Steiner tree literature interpreted

this quantity as the *integrality gap*, since it is the ratio of the minimum cost of a Steiner tree (38) and that of its linear program relaxation (29).

A useful metric is the average gain (or its asymptotic growth rate with $|V|$), under appropriate probabilistic measures defined for the problem space $\{(V, E, p, s, T)\}$. Characterizing the average gain is, in our opinion, an interesting future topic.

V. SIMULATIONS

We conduct simulations on an example community wireless network consisting of 58 houses spaced roughly in six rows, shown in Fig. 7(a). The locations of the houses are marked with dots in Fig. 7(a). We consider a multicast session from node 28 to a set of destination nodes $\{1, 6, 52, 53\}$, as illustrated in Fig. 7(a) by four lines.

Simulation parameters are set up as follows. The path gain coefficients are set as $a_{ij} = d(i, j)^{-3}$, where $d(i, j)$ is the distance between node i and j . The SNR threshold γ is set to 4 dB. The noise level is chosen to be 1, as a normalization without essential loss of generality. We adopt the continuous power control model and assume that each node can set its transmission power at any value less than or equal to p_{\max} . In the simulations, we set p_{\max} at a value corresponding to a maximum transmission range of 300 m.

The minimum-energy routing solution and network coding solution are shown in Fig. 7(b) and (c), respectively, which are the connectivity graphs for the optimal traffic assignment. The minimum energy-per-bit with network coding is 98.08% of the minimum energy-per-bit with routing; computing the optimal network coding solution (linear program) takes less than 1/30 the time for computing the optimal routing solution (integer linear program).

The minimum-energy network coding solution of this problem admits an easy interpretation. Suppose the source stream is partitioned into two distinct substreams, $x_1(n)$ and $x_2(n)$, $n = 1, 2, \dots$, which we can think of as odd and even




Fig. 8. Efficient information exchange between nodes 6 and 32. Node 6 has learned substream $x_1(n)$ and node 32 has learned substream $x_2(n)$. The information exchange between nodes 6 and 32 can be done efficiently by exploiting network coding together with physical-layer broadcast. There is a path from node 6 to node 32 and a path from node 32 to node 6. A critical observation is that, with network coding, the union of the two paths is sufficient to provide the required routing capability, i.e., unit rate from node 6 to node 32 and unit rate from node 32 to node 6.

subsequences, for example. Then, substream $x_1(n)$ is delivered from source node 28 to nodes 1 and 6 through a sequence of transmissions: $28 \rightarrow 19$, $19 \rightarrow 8$, $8 \rightarrow \{4, 5\}$, $5 \rightarrow 6$, $4 \rightarrow 3$, $3 \rightarrow 2$, and $2 \rightarrow 1$. Substream x_2 is delivered to nodes 32, 47, and 53 through a sequence of transmissions $28 \rightarrow \{31, 36\}$, $31 \rightarrow 32$, $36 \rightarrow 43$, $43 \rightarrow 49$, $49 \rightarrow 56$, $56 \rightarrow 55$, $55 \rightarrow 54$, and $54 \rightarrow \{47, 53\}$. Now node 6 has substream $x_1(n)$ and node 32 has substream $x_2(n)$ and, thus, nodes 6 and 32 need to exchange information. The information exchange between nodes 6 and 32 can be done efficiently by exploiting network coding together with physical layer broadcast. This is illustrated by the graph representation in Fig. 8. All edges in Fig. 8 have the same capacity. There is a path from node 6 to node 32

$$6 \rightarrow 6' \rightarrow 13 \rightarrow 13' \rightarrow 14 \rightarrow 14' \\ \rightarrow 22 \rightarrow 22' \rightarrow 26 \rightarrow 26' \rightarrow 32 \rightarrow 32'$$

and a path from node 32 to node 6

$$32 \rightarrow 32' \rightarrow 26 \rightarrow 26' \rightarrow 22 \rightarrow 22' \\ \rightarrow 14 \rightarrow 14' \rightarrow 13 \rightarrow 13' \rightarrow 6 \rightarrow 6.$$

A critical observation is that, with network coding, the union of the two paths is sufficient to provide the required routing capability, i.e., unit rate from node 6 to node 32 and unit rate from node 32 to node 6. The structure in Fig. 8 coincides with the simple example of Fig. 1. A similar mutual exchange phenomenon also occurs at the bottom of Fig. 7(b) between nodes 1 and 47. For a detailed discussion of information exchange in a wireless network using network coding, please refer to [24], where we also propose a simple distributed communication scheme using only XOR for network coding.

VI. CONCLUSION

This paper unifies the study of network coding and routing for the problem of minimum-energy multicasting in a wireless ad hoc network. We model a wireless ad hoc network as a graph composed of tree structures representing physical broadcast links. Each edge of the graph has a price, which is the energy-per-bit for the corresponding broadcast link. The minimum-energy multicasting formulation amounts to minimizing the total cost of the consumed bit rates on the edges while providing a unit multicast rate.

For a system that is allowed to use network coding, the admissible rate region, viz. the bit-rate demand on the edges, is characterized by a MAX of flows, which is the edge-wise maximum of flows from the source to each destination. Since the region can be characterized by linear inequalities, the minimum energy-per-bit is found by a linear program.

For a more traditional routing-based system, the admissible rate region is characterized by a sum of trees. Due to linearity of the cost criterion, we only need to use a single tree, the most economic tree. The admissible rate region with a single tree is characterized by a MAX of flows with integer constraints. Thus, the minimum energy-per-bit is found by an integer linear program, whose relaxation becomes the minimum energy multicast formulation with network coding.

In short, we observe that network coding offers advantages over routing in terms of both energy and computation efficiencies. More precisely, the energy advantage is characterized by the integrality gap. The difference in computation is more significant: While optimal routing solution is NP-hard, optimal coding solution is polynomial.

APPENDIX A PROOF OF THEOREM 3

The proof consists of two steps. 1) We show the cost function of (21) is equivalent to $\epsilon^T \mathbf{c}$. 2) We then reduce the constraints of (21) to $\mathbf{c} \in \mathcal{A}(1)$.

- 1) Consider an elementary graph $G_k \in \mathcal{B}$. Let $u \rightarrow \{v_1, \dots, v_n\}$ be the broadcast link in G_k . This broadcast link is represented by a tree-like structure with a virtual vertex $u' \in V_1$. In G_k

$$c_k(uu') = c_k(u'v_1) = \dots = c_k(u'v_n) = 1$$

and other edges have zero capacity. Furthermore, the edges $\{uu', u'v_1, \dots, u'v_n\}$ have zero capacity in all other elementary graphs. This implies that the constraints $\mathbf{c} = \sum_k \lambda_k \mathbf{c}_k$ amount to $|\mathcal{B}|$ sets of constraints in the following form:

$$c(uu') = c(u'v_1) = \dots = c(u'v_n) = \lambda_k. \quad (42)$$

Thus, (21) is equivalent to

$$\begin{aligned} \mathcal{E}^* &= \min \epsilon^T \mathbf{c}, \\ &\text{subject to } \mathbf{c} \in \mathcal{A}(1) \\ c(uu') &= c(u'v_1) = \dots = c(u'v_n) = \lambda_k \quad \forall u' \in V_1 \\ \lambda_k &\geq 0 \quad \forall k. \end{aligned}$$

It follows that the variables λ_k can be eliminated, resulting in

$$\begin{aligned} \mathcal{E}^* &= \min \epsilon^T \mathbf{c}, \\ &\text{subject to } \mathbf{c} \in \mathcal{A}(1) \\ c(uu') &= c(u'v_1) = \dots = c(u'v_n) \quad \forall u' \in V_1. \end{aligned} \quad (43)$$

2) The flow conservation constraints imply

$$f_t(u'v_i) \leq f_t(uu'), \quad i = 1, \dots, n \quad \forall t \in T$$

and hence

$$\max_{t \in T} f_t(u'v_i) \leq \max_{t \in T} f_t(uu'), \quad i = 1, \dots, n. \quad (44)$$

This shows that the demand on $u'v_i$ is always less than or equal to the demand on uu' .

The relations (42) and (44) together imply that, in (43), among the $|E|$ component constraints of

$$\mathbf{c} \geq \max_{t \in T} \mathbf{f}_t$$

the $|E| - |\mathcal{B}|$ constraints corresponding to the $u'v_j$ -type ($u' \in V_1$) edges are redundant. Note further that the $u'v_j$ -type edges do not contribute to the objective function $\boldsymbol{\varepsilon}^T \mathbf{c}$. Thus, relaxing the constraints

$$c(uu') = c(u'v_1) = \dots = c(u'v_n)$$

does not affect the optimal value. As a result, the optimal value of (21) is equivalent to that of (29)

$$\mathcal{E}^* = \mathcal{E}_{\text{coding}}^*(V, E, \boldsymbol{\varepsilon}, s, T). \quad (45)$$

Furthermore, any optimal solution of (29), say \mathbf{c}^* , can be converted into an optimal solution of (21) $\hat{\mathbf{c}}^*$ by setting

$$\hat{c}^*(uu') = \hat{c}^*(u'v_1) = \dots = \hat{c}^*(u'v_n) = c^*(uu'). \quad (46)$$

APPENDIX B

HYBRID CODING/ROUTING RESULT

Lemma 2: Suppose in a graph $G = (V, E, \mathbf{c})$ that the vertex set V is composed of two disjoint sets, V_0 and V_1 , with $s \in V_0$, $T \subseteq V_0$. Suppose also that each vertex $u' \in V_1$ has only one incoming edge $uu' \in E$ and multiple outgoing edges $u'v_j$, $j = 1, \dots, n$, with $c(uu') = c(u'v_j)$, $j = 1, \dots, n$. The maximum rate for multicasting from source s to destinations T can still be achieved even if each vertex $u' \in V_1$ can only perform routing instead of arbitrary network coding.

Proof: Let G' denote the subgraph of G containing the edges $\{uu' \in E | u' \in V_1\}$. Since, in G' , each vertex has in-degree 0 or 1, G' can be decomposed into several edge-disjoint components, where each component is either a tree or a tree plus a "back" edge pointing from a leaf node to the root. For the latter case, none of the vertex in the component can be reached by s . Thus, these vertices can be deleted without affecting the multicast capacity, i.e., the maximum achievable multicast rate.

Hence, without loss of generality, we can assume G' is composed of several edge-disjoint trees. Each tree has a *root edge* through which information can flow in. In G , each such tree is connected to several *egress edges* through which information




Fig. 9. Condition $c(uu') = c(u'v_j)$, $j = 1, \dots, n$, is critical for Lemma 2. If u' can perform network coding, then rate 2.0 is achievable for multicasting from u to $\{t_1, t_2, t_3\}$; if u' can only perform routing, the rate 2.0 cannot be achieved.

can flow out to nodes in V_0 . Note that the information flowing on any egress edge $u'v$, $u' \in V_1$, $v \in V_0$ has to be derived from the information flowing on its corresponding root edge. Consider an arbitrary network coding solution \mathcal{S} . Modify this solution \mathcal{S} by letting each vertex $u' \in V_1$ forward whatever information flowing on uu' to its outgoing edges. Then, the information flowing on an egress edge $u'v$, $u' \in V_1$, $v \in V_0$ is the same as the information flowing on its corresponding root edge; hence, the corresponding node $v \in V_0$ is indeed capable of producing the information flowing on its outgoing edges as prescribed in \mathcal{S} . This yields another valid coding solution satisfying the required restrictions, while achieving the same throughput as \mathcal{S} does. ■

Note that the condition $c(uu') = c(u'v_j)$, $j = 1, \dots, n$, is critical for Lemma 2, as illustrated in Fig. 9 [5]. If u' can perform network coding, then a rate of two is achievable for multicasting from u to $\{t_1, t_2, t_3\}$; if u' can only perform routing, the multicast rate of two cannot be achieved.

APPENDIX C

MULTIPLE MULTICAST SESSIONS

Earlier we have considered finding the minimum energy-per-bit for a single multicast session, under the layered model. A more general problem is to find the minimum energy required to achieve traffic demands formulated as multiple multicast sessions.

Suppose the end-to-end communication demands are comprised of multiple multicast sessions, which we denote by $\langle s_m, T_m, r_m \rangle$, $m = 1, \dots, M$. In the m th multicast session, a source node s_m wants to transmit information to destination nodes T_m at rate r_m . Assume the messages in different sessions are independent. Given (V, E) , let $\mathcal{A}(r_1, \dots, r_m)$ denote the associated admissible rate region. Characterizing $\mathcal{A}(r_1, \dots, r_m)$ remains an open problem in information theory; see [26]. We now briefly discuss an inner bound and an outer bound of $\mathcal{A}(r_1, \dots, r_m)$, both of which can be characterized by a system of linear inequalities.

An inner bound can be established by partitioning the available resources into M shares as $\mathbf{c} = \mathbf{c}_1 + \dots + \mathbf{c}_M$ and use \mathbf{c}_m to provide the communication rate for session m . The corresponding rate region $\underline{\mathcal{A}}(r_1, \dots, r_M)$ is given in Fig. 10. The characterization of $\underline{\mathcal{A}}(r_1, \dots, r_M)$ hinges upon the vector $\sum_{m=1}^M \max_{t \in T^m} \mathbf{f}_t^m$, which was called a *SUM of MAX of flows* in [12]. The set $\underline{\mathcal{A}}(r_1, \dots, r_M)$ can be characterized by a system of linear inequalities [12]. This inner bound is, in

$$\begin{aligned} \underline{\mathcal{A}}(r_1, \dots, r_M) &\equiv \left\{ \mathbf{c} \in \mathbb{R}^{|E|} \mid \exists \mathbf{f}_t^m \in \mathcal{F}_{s_m, t}, |\mathbf{f}_t^m| = r_m, \forall t \in T_m, \forall m \in \{1, \dots, M\}, \text{ such that } \mathbf{c} \geq \sum_{m=1}^M \max_{t \in T} \mathbf{f}_t \right\} \\ \overline{\mathcal{A}}(r_1, \dots, r_M) &\equiv \left\{ \mathbf{c} \in \mathbb{R}^{|E|} \mid \exists \mathbf{f}_j \in \mathcal{F}_{\mathcal{S}_j, \mathcal{T}_j}, |\mathbf{f}_j| = \sum_{m: s_m \in \mathcal{S}_j} r_m, \text{ for each cut type } (\mathcal{S}_j, \mathcal{T}_j), \text{ such that } \mathbf{c} \geq \max_{j=1, \dots, \Delta} \mathbf{f}_j \right\} \end{aligned}$$

Fig. 10. Inner and outer bounds of the admissible rate region for multisource multicasting. Both can be characterized by a linear system of inequalities.

general, loose. To see this, consider Fig. 8. There are two unicast sessions with independent messages. A unit rate for both sessions can be achieved by performing inter-session network coding, while it cannot be achieved by separately processing the two sessions.

A well-known technique for establishing an outer bound is to consider the cuts; see, e.g., Cover and Thomas [27] and Yeung [26]. For any cut (U, \overline{U}) of G , define

$$\mathcal{S}(U) \equiv \{s_m \mid s_m \in U, \exists t \in T_m, t \in \overline{U}\} \quad (47)$$

which is the set of sources for sessions that need to communicate across the cut. The sum rate of all sessions that need to communicate across the cut has to be less than or equal to the total capacity of the cut, viz.,

$$\sum_{m: s_m \in \mathcal{S}(U)} r_m \leq \sum_{vw \in E: v \in U, w \in \overline{U}} c(vw). \quad (48)$$

An inequality of the form (48) is said to be a *cut condition*. The number of cut conditions grows exponentially with $|V|$. We now show how to use a set of flow constraints to efficiently represent the cut conditions.

Given a cut (U, \overline{U}) , construct a set $\mathcal{T}(U)$ that includes one (arbitrarily selected) destination $t \in T_m$ for each m with $s_m \in \mathcal{S}(U)$. Then, all cuts with $\mathcal{S}(U) \neq \emptyset$ can be classified into (at most) Δ types according to the value of $(\mathcal{S}, \mathcal{T})$, where Δ is defined as follows:

$$\begin{aligned} \Delta &\equiv \sum_{\{m_1, \dots, m_k\} \subseteq \{1, \dots, M\}} |T_{m_1}| \cdot |T_{m_2}| \cdots |T_{m_k}| \\ &= (|T_1| + 1)(|T_2| + 1) \cdots (|T_M| + 1) - 1. \end{aligned} \quad (49)$$

We call a cut of type $(\mathcal{S}, \mathcal{T})$ an \mathcal{S} - \mathcal{T} cut. Then, the cut conditions for \mathcal{S} - \mathcal{T} cuts amount to requiring that the capacity of any \mathcal{S} - \mathcal{T} cut is not less than $\sum_{m: s_m \in \mathcal{S}} r_m$. These conditions can be translated into an equivalent flow constraint, due to the Max-Flow-Min-Cut Theorem. Contract the vertices \mathcal{S} into one super-node \hat{s} and contract the vertices \mathcal{T} into one super-node \hat{t} . Contracting a vertex set A into a vertex a means replacing A by a new node a , whose incoming edges are the edges going from $V - A$ to A and whose outgoing edges are the edges going from A to $V - A$. Then, there exists an \hat{s} - \hat{t} flow of rate r if and only if the capacity of any \mathcal{S} - \mathcal{T} cut is greater than or equal to r .

From the above discussions, we obtain an outer bound of $\overline{\mathcal{A}}(r_1, \dots, r_M)$ in Fig. 10. In Fig. 10, $\mathcal{F}_{\mathcal{S}_j, \mathcal{T}_j}$ stands for the linear system of inequalities characterizing an \mathcal{S}_j - \mathcal{T}_j flow (the \hat{s} - \hat{t} flow after contracting \mathcal{S}_j and \mathcal{T}_j into \hat{s} and \hat{t} , respectively).

The characterization of $\overline{\mathcal{A}}(r_1, \dots, r_M)$ hinges upon the vector $\sum_j \mathbf{f}_j$, which is the SUM of Δ flows. Consequently, this outer bound can also be characterized by a set of linear inequalities. Since an s - t flow can be characterized by about $|V| + |E|$ linear inequalities, $\overline{\mathcal{A}}(r_1, \dots, r_M)$ can be characterized by around $(|V| + |E|)\Delta$ linear inequalities.

Combining the linear inequalities for $\underline{\mathcal{A}}(r_1, \dots, r_M)$ (respectively $\overline{\mathcal{A}}(r_1, \dots, r_M)$) with the linear inequalities representing $G \in \mathcal{G}(\mathcal{B})$ leads to upper (lower) bounds of the minimum energy-per-bit.


ACKNOWLEDGMENT

The authors are very grateful for the reviewers for their constructive comments, which have helped improve the clarity of this paper. In particular, one reviewer's suggestion led to the current proof of *Lemma 2*, which is simpler than the proof in an earlier draft. The authors would like to thank Dr. K. Jain for the helpful discussions, Dr. J. Padhye for providing the data for the community wireless networks, and Prof. H. Kobayashi for his comments on the paper. After the initial submission of the manuscript, the authors learned that Lun *et al.* [25] independently considered a similar problem, with the emphasis placed on distributed algorithms.

REFERENCES

- [1] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides, "On construction of energy-efficient broadcast and multicast trees in wireless networks," in *Proc. IEEE Conf. Computer Commun.*, Tel Aviv, Israel, 2000, pp. 586–594.
- [2] W. Liang, "Constructing minimum-energy broadcast trees in wireless ad hoc networks," in *Proc. ACM Int. Symp. Mobile Ad hoc Netw. Computing*, Lausanne, Switzerland, Jun. 2002, pp. 112–122.
- [3] M. Cagalj, J.-P. Hubaux, and C. Enz, "Minimum-energy broadcast in all-wireless networks: NP-completeness and distribution issues," in *Proc. Int. Conf. Mobile Computing Netw.*, 2002, pp. 172–182.
- [4] Y. Wu, P. A. Chou, and S.-Y. Kung, "Minimum-energy multicast in mobile ad hoc networks using network coding," in *Proc. IEEE Inf. Theory Workshop*, San Antonio, TX, Oct. 2004, [CD-ROM].
- [5] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inf. Theory*, vol. 46, no. 4, pp. 1204–1216, Jul. 2000.
- [6] P. A. Chou, Y. Wu, and K. Jain, "Practical network coding," in *Proc. 41st Allerton Conf. Commun., Control, Computing*, Oct. 2003, [CD-ROM].
- [7] T. Ho, R. Koetter, M. Médard, M. Effros, J. Shi, and D. Karger, "Toward a random operation of networks," *IEEE Trans. Inf. Theory*, submitted for publication.
- [8] T. Ho, M. Médard, J. Shi, M. Effros, and D. R. Karger, "On randomized network coding," in *Proc. 41st Allerton Conf. Commun., Control, Computing*, Oct. 2003, [CD-ROM].
- [9] T. Ho, R. Koetter, M. Médard, D. R. Karger, and M. Effros, "The benefits of coding over routing in a randomized setting," in *Proc. Int. Symp. Inf. Theory*, 2003, p. 442.
- [10] Y. Wu, P. A. Chou, and K. Jain, "A comparison of network coding and tree packing," in *Proc. IEEE Int. Symp. Inf. Theory*, Chicago, IL, Jun. 2004, p. 143.

- [11] T. Ho, B. Leong, M. Médard, R. Koetter, Y. Chang, and M. Effros, "On the utility of network coding in dynamic environments," in *Proc. Int. Workshop Wireless Ad-hoc Netw.*, Jun. 2004, [CD-ROM].
- [12] 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," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 1, pp. 136–150, Jan. 2005.
- [13] Y. Wu, "Network coding for multicasting," Ph.D. dissertation, Dept. Elect. Eng., Princeton Univ., Princeton, NJ, 2005.
- [14] T. M. Cover, "Broadcast channels," *IEEE Trans. Inf. Theory*, vol. IT-18, no. 1, pp. 2–14, Jan. 1972.
- [15] P. P. Bergmans, "Random coding theorem for broadcast channels with degraded components," *IEEE Trans. Inf. Theory*, vol. IT-19, no. 2, pp. 197–207, Mar. 1973.
- [16] —, "A simple converse for broadcast channels with additive white Gaussian noise," *IEEE Trans. Inf. Theory*, vol. IT-20, no. 2, pp. 279–280, Mar. 1974.
- [17] R. G. Gallager, "Capacity and coding for degraded broadcast channels," *Probl. Inf. Transm.*, pp. 185–193, Jul.-Sep. 1974.
- [18] K. Jain, M. Mahdian, and M. R. Salavatipour, "Packing Steiner trees," in *Proc. 14th ACM-SIAM Symp. Discrete Algorithms (SODA)*, 2003, pp. 266–274.
- [19] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*. Berlin, Germany: Springer-Verlag, 2003.
- [20] F. K. Hwang, D. S. Richards, and P. Winter, *The Steiner Tree Problem*. Amsterdam, The Netherlands: North-Holland, 1992.
- [21] A. Claus and N. Maculan, "Une nouvelle formulation du problème de Steiner sur un graphe," Univ. Montréal, Centre de Recherche sur les Transports, Montreal, QC, Canada, Tech. Rep. 280, 1983.
- [22] R. T. Wong, "A dual ascent approach for Steiner tree problem in a directed graph," *Math. Program.*, vol. 28, pp. 271–287, 1984.
- [23] J. E. Beasley, "An algorithm for the Steiner tree problem in graphs," *Networks*, vol. 14, pp. 147–159, 1984.
- [24] Y. Wu, P. A. Chou, and S.-Y. Kung, "Information exchange in wireless networks with network coding and physical-layer broadcast," in *Proc. 39th Annu. Conf. Inf. Sci. Syst.*, Mar. 2005, [CD-ROM].
- [25] D. S. Lun, M. Médard, T. Ho, and R. Koetter, "Network coding with a cost criterion," Mass. Inst. Technol., Cambridge, MA, Tech. Rep. LIDS-P-2584, 2004.
- [26] R. W. Yeung, *A First Course in Information Theory*. New York: Kluwer/Plenum, 2002.
- [27] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.




Yunnan Wu (S'02) received the B.S.E. degree in computer science from the University of Science and Technology of China, Hefei, China, in 2000, and the M.A. degree in electrical engineering from Princeton University, Princeton, NJ, in 2002, where he is currently a Ph.D. degree candidate.

He was with Microsoft Research Asia, Beijing, China, from 1999 to 2001 as a Research Assistant, with Bell Laboratory, Lucent Technologies, as a summer intern in 2002, and with Microsoft Research, Redmond, WA, as a summer intern in 2003. Since

2005, he has been a Researcher with Microsoft Research. His research interests include networking, graph theory, information theory, wireless communications, signal processing, and multimedia communications.

Mr. Wu was the recipient of the Best Student Paper Award at the 2000 SPIE and IS&T Visual Communication and Image Processing Conference and the Student Paper Award at the 2005 IEEE International Conference on Acoustics, Speech, and Signal Processing. He is also a recipient of a Microsoft Research Graduate Fellowship for 2003–2005.



Philip A. Chou (S'82–M'87–SM'00–F'04) received the B.S.E. degree in electrical engineering and computer science from Princeton University, Princeton, NJ, in 1980, the M.S. degree in electrical engineering and computer science from the University of California, Berkeley, in 1983, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 1988.

From 1988 to 1990, he was Member of Technical Staff with AT&T Bell Laboratories, Murray Hill, NJ. From 1990 to 1996, he was a Member of Research

Staff with the Xerox Palo Alto Research Center, Palo Alto, CA. In 1997, he was Manager of the Compression Group, VXtreme, Mountain View, CA, before it was acquired by Microsoft in 1997. Since 1998 to the present, he has been a Senior Researcher with Microsoft Research, Redmond, WA. He also served as a Consulting Associate Professor at Stanford University, Stanford, CA, from 1994 to 1995 and has been an Affiliate Professor with the University of Washington, Seattle, since 1998. His research interests include data compression, information theory, communications, and pattern recognition, with applications to video, images, audio, speech, and documents.

Dr. Chou was the recipient of the 2002 ICME Best Paper Award with A. Seghal and the 1993 Signal Processing Society Paper Award with T. Lookabaugh. He served as an Associate Editor in source coding for the IEEE TRANSACTIONS ON INFORMATION THEORY from 1998 to 2001 and served as a Guest Associate Editor for special issues in the IEEE TRANSACTIONS ON IMAGE PROCESSING and the IEEE TRANSACTIONS ON MULTIMEDIA in 1996 and 2004, respectively. From 1998 to 2004, he was a member of the IEEE Signal Processing Society's Image and Multidimensional Signal Processing Technical Committee (IMDSP TC). He is a member of Phi Beta Kappa, Tau Beta Pi, Sigma Xi, and the IEEE Computer, Information Theory, Signal Processing, and Communications Societies and was an active member of the MPEG Committee.



Sun-Yuan Kung (S'74–M'78–SM'84–F'88) received the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 1977.

In 1974, he was an Associate Engineer with Amdahl Corporation, Sunnyvale, CA. From 1977 to 1987, he was a Professor of electrical engineering systems with the University of Southern California, Los Angeles. Since 1987, he has been a Professor of electrical engineering at Princeton University, Princeton, NJ. He held a Visiting Professorship at Stanford University and the Delft University of

Technology, Delft, The Netherlands, in 1984, a Toshiba Chair Professorship at Waseda University, Japan, in 1984, an Honorary Professorship at Central China University of Science and Technology, in 1994, and a Distinguished Chair Professorship at the Hong Kong Polytechnic University since 2001. He has coauthored more than 400 technical publications and numerous textbooks, including *VLSI Array Processors* (Englewood Cliffs, NJ: Prentice-Hall, 1988), *Digital Neural Networks* (Englewood Cliffs, NJ: Prentice-Hall, 1993), *Principal Component Neural Networks* (New York: Wiley, 1996), and *Biometric Authentication: A Machine Learning and Neural Network Approach* (Englewood Cliffs, NJ: Prentice-Hall, 2004). His research interests include VLSI array processors, system identification, neural networks, wireless communication, multimedia signal processing, bioinformatic data mining, and biometric authentication. Since 1990, he has been the Editor-in-Chief of the *Journal of VLSI Signal Processing Systems*.

Prof. Kung was a recipient of the IEEE Signal Processing Society's Technical Achievement Award for his contributions on parallel processing and neural network algorithms for signal processing in 1992, the IEEE Signal Processing Society's Best Paper Award for his publication on principal component neural networks in 1996, and the IEEE Third Millennium Medal in 2000. He was a Distinguished Lecturer of the IEEE Signal Processing Society in 1994.