Minimum-Energy Multicast in Mobile Ad hoc Networks using Network Coding

Yunnan Wu[1]
Princeton University
yunnanwu@princeton.edu

Philip A. Chou
Microsoft Research
pachou@microsoft.com

Sun-Yuan Kung
Princeton University
kung@princeton.edu

Abstract — **The minimum energy required to transmit a bit of information through a network characterizes the most economical way to communicate in a network. In this paper, we show that under a simplified layered model of wireless networks, the minimum-energy multicast problem in mobile ad hoc networks is solvable as a linear program, assuming network coding. Compared with conventional routing solutions, network coding not only promises a potentially lower energy-per-bit, but also enables finding the optimal solution in polynomial time, in sharp contrast with the NP-hardness of constructing the minimum-energy multicast tree as the optimal routing solution.**

I. INTRODUCTION

Consider the problem of minimum-energy information multicast, namely transmitting common information (say, one bit) from a sender s to a set of receivers T with the minimum amount of total consumed energy, in a mobile ad hoc network.

Several previous works, e.g., [1–3], aim at finding the *minimum-energy multicast tree.* Suppose sender s wants to transmit a message to receivers T. Assuming that the intermediate nodes can only *route*, i.e., replicate and forward, messages received, 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 level. Note that due to the broadcast nature of radio transmissions, a single transmission by a certain transmitter may result in multiple nodes recovering the transmitted signals; 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, taking advantage of the physical-layer broadcast property, is NP-complete. Thus, related previous works give heuristic algorithms and investigate their respective approximation ratios.

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. Figure 1 gives an example wireless ad hoc network. The locations of the nodes have been marked with dots. Suppose each node is equipped with a transmitter operating at a fixed transmission range 110m, which is

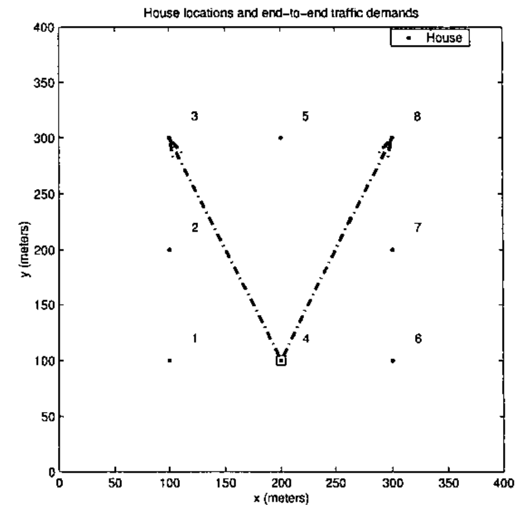

Figure 1: An example wireless ad hoc network, used to illustrate that there is room to improve upon the conventional routing-based formulation to achieve potentially better energy efficiency.

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. Consider multicasting from node 4 to node 3 and node 8. Under the conventional formulation as presented in the previous paragraph, it is easy to see that 5 transmissions is the minimum amount of energy required to deliver 1 message from 4 to $\{3, 8\}$. One such solution is as follows: 4 transmits to $\{1, 6\}$, 1 to 2, 2 to 3, 6 to 7, and 7 to 8.

We now show a more energy-efficient solution. The solution is given as a graph in Figure 2. In Figure 2, we introduce for each node u a virtual vertex u' to model physical layer broadcast. A broadcast link from u to node v_1, v_2, which fall in the radio range of u, is represented in the graph by an edge from u to u', an edge from u' to v_1, and an edge from u' to v_2. Consider multicasting two messages x_1, x_2 from 4 to $\{3, 8\}$. The solution is given in Figure 2 by specifying for each edge the carried messages. First, x_1 is delivered to node 3 with 3 transmissions and x_2 is delivered to node 8 with 3 transmissions. Next, node 3 transmits x_1 to node 5 and node 8 transmits x_2 to node 5. The critical step occurs at node 5, which broadcasts the XOR-ed result of two messages, $x_1 \oplus x_2$, to node 3 and node 8, consuming only one transmission. Each of the two receivers can recover both x_1 and x_2 by solving a simple linear system of equations. According to this solution, two messages are delivered with 9 transmissions, improving upon the 10 transmissions

[1]Contact author: Yunnan Wu. Dept. of Electrical Engineering, Princeton University, Princeton, NJ 08544. Tel.:(609)258-0102. Fax:(609)258-3745.

0-7803-8720-1/04/$20.00 ©2004 IEEE

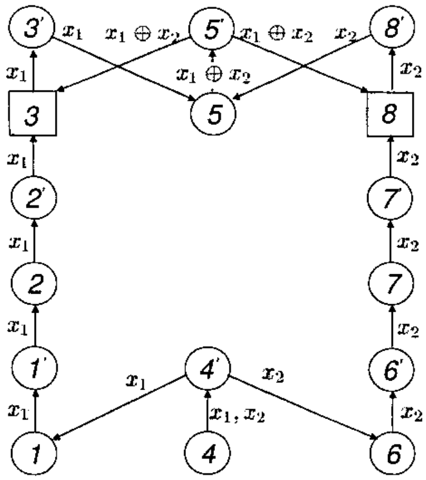

Figure 2: Minimum-energy multicast with network coding on the example network.

with the conventional routing solution. It can be shown that 4.5 transmissions is indeed the minimum to multicast one message.

The lesson of this example is that it is in general suboptimal to require the interior 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 [4], 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 in [4] that the *multicast capacity*, which is defined as the maximum rate at which a sender can communicate common information to a set of receivers, is equal to the minimum of maximum flow values between the sender and each receiver and can be achieved with network coding. Subsequently, Li, Yeung, and Cai [5] showed that it is sufficient for the encoding functions at the interior nodes to be linear. Koetter and Médard [6] gave an algebraic characterization of linear encoding schemes and proved existence of linear time-invariant codes achieving the multicast capacity. Empirical comparisons of the throughput achievable by routing and network coding have been reported in [7] for static graphs and in [8] for dynamically varying graphs.

In this work, we explore the advantage of network coding in economically using network resources, more specifically the energy consumption. The example in Figure 1 and Figure 2, which has a similar structure as the example in [4], demonstrates the network coding can lead to solutions that are more economic in using network resources than routing solutions. More generally, we show that under a simplified layered model of wireless networks, the minimum-energy multicast problem is solvable in polynomial time as a linear program, assuming network coding. Thus, in this context, advantages of network coding over routing are two-fold: network coding can be used to attain the lowest energy-per-bit; at the same time, minimum-energy network coding can be solved in polynomial time, in sharp contrast with the NP-hardness of constructing the minimum-energy multicast tree.

Note that the "minimum energy-per-bit" investigated in this paper is certainly not the fundamental limit in network information theory, where lots of problems remain open. The minimum energy-per-bit of a wireless network in information theoretic sense is still unknown, even for a small network with 3 nodes; see, e.g., [9]. The discussions in this paper are based on a simplified layered model of wireless networks. The basic assumptions of the layered model are as follows. The physical and link layers of the wireless network provide communication resources in the form of a collection of "lossless bit-pipes", each capable of transferring information between (neighboring) nodes at a certain rate. Given a collection of lossless bit-pipes, information can be routed from the sources to the destinations at certain rates in the network layer. The layered model essentially separates the allocation and consumption of bit-rate resources. Such a separation is certainly sub-optimal from an information theoretic perspective. Nonetheless, its simplicity facilitates the analysis and enables some engineering insights to be obtained.

To explain the layered model in more details, we adopt the terminologies and notations in [10]. A collection of lossless bit-pipes mentioned earlier is called a *capacity graph* in [10]. A capacity graph can be described as a triple $G = (V, E, c)$ where V and E are sets of vertices and edges respectively and c is a function assigning to each edge $vw \in E$ a positive edge capacity $c(vw)$. At the physical layer, the network operates in many different *physical states*, each corresponding to an arrangement of concurrent links among neighbors. Each physical state corresponds to an *elementary capacity graph*. At the link layer, by timesharing different physical states, convex combinations of the elementary capacity graphs can be achieved, hence presenting to the upper layers a set of *supported (composite) capacity graphs*.

Each supported capacity graph provides certain communication rate and has an associated power consumption. The minimum energy-per-bit can be found by a linear program that identifies a supported capacity graph and a flow assignment, such that the ratio of the power consumption achieved to the multicast rate achieved, that is, the energy-per-bit, is minimized.

A more general problem, cost (e.g., power) minimization for multiple multicast sessions subject to rate constraints, was considered in [10]; the problem is in general very difficult because of the combinatorial nature of the problem that arises due to the effects of interference. The present paper observes that if the required rate is sufficiently small then the effect of interference becomes immaterial, enabling one to find the minimum energy-per-bit in polynomial time. More details can be found in [11].

Recently, a similar problem was considered in an independent work by Lun et al [12].

II. Capacity Graphs for Wireless Ad hoc Networks

The set of supported capacity graphs represents possible allocations of communication resources arising from scheduling and power control in the physical and link layers. In this section, we discuss the structure of the set of supported capacity graphs, first for a fixed wireless ad hoc network where the channel conditions stay the same during a communication session, and then for a mobile ad hoc network where the channel conditions evolve over time.

A. Fixed 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 the set of associated receivers, 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 capacity graph* (ECG). By time-sharing among different physical states, it is possible to achieve any convex combination of the ECGs. That is, if λ_k is the relative share of time for the ECG $G_k = (V_k, E_k, c_k)$, then it is possible to achieve on average the capacity graph $G = (\cup_k V_k, \cup_k E_k, \sum_k \lambda_k c_k)$, where the edge capacities c_k are each extended to $\cup_k E_k$ in the obvious way. Denote such combinations by $G = \sum_k \lambda_k G_k$. The capacity graphs resulting from timesharing the ECGs will be referred to as *supported (composite) capacity graphs*. A wireless network can be characterized by the set of capacity graphs equal to convex combinations of ECGs. This can be expressed mathematically as

$$\mathcal{G}_0 = \left\{ G \,\middle|\, G = \sum_k \lambda_k G_k,\ \sum_k \lambda_k \leq 1,\ \lambda_k \geq 0\,\forall k, G_k \in \mathcal{B}_0 \right\},$$

which is the set of capacity graphs supported/spanned by the set of all feasible ECGs $\mathcal{B}_0$.

The number of feasible ECGs $|\mathcal{B}_0|$ generally grows exponentially with the number of network nodes. Fortunately, for the minimum-energy multicast problem, the number of *pertinent* elementary capacity graphs, or corresponding physical states, is polynomial in the number of nodes. This is because of the following important albeit simple observation: *each Joule of energy is spent most efficiently when used on a transmission that is free of interference.* Consequently, we can restrict our attention to those physical states involving only a single transmitter, since we are focusing on the minimum energy-per-bit. It is this fact that results in the polynomial solvability of the minimum-energy multicast problem for a mobile ad hoc network.

With a finite set of ECGs $\mathcal{B} \subseteq \mathcal{B}_0$, the convex set of supported capacity graphs becomes

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

where the dependence on $\mathcal{B}$ is explicitly shown. The power consumption of a composite capacity graph $G = \sum_k \lambda_k G_k$ is $P\left(\sum_k \lambda_k G_k\right) = \sum_k \lambda_k P(G_k)$, where $P(G_k)$ is the power consumption of ECG G_k.

Elementary capacity graphs: We now show how to obtain the elementary capacity graph for a given physical state. As discussed above, we only need to consider the physical states with a single transmitter broadcasting to one or multiple receivers.

Let the transmitter be u and let p_u be its transmitting power. Let the path loss ratio from node $i \in V_0$ to node $j \in V_0$ be a_{ij}. Denote the noise variance by σ^2. The SNR at a node j receiving information from $u \in V_0$ admits the following expression: $\mathsf{SNR}_{uj} = a_{uj} p_u / \sigma^2$. We adopt a simplistic physical layer model: as long as the received SNR (single-to-noise ratio) 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.

We now model physical layer broadcast links, which are important since wireless networks operate in an inherent broadcast medium. Note there may be more than one node j whose receiving SNR level exceeds the threshold. In this case, it would be wasteful to require that only one node can receive from the transmitter. On the other hand, it would be overly optimistic to assume that all of them can afford a separate point-to-point unit-capacity link with the transmitter, since distinct information may be loaded on these links in the latter scenario.

Hence, we propose modeling broadcast links by introducing for each broadcast transmitter in each elementary capacity graph a virtual vertex that has point-to-point links with the receivers. For link $u \rightarrow Y_u$, we add to the associated ECG a distinct virtual vertex (e.g., u'), and unit capacity edges $uu', u'v, v \in Y_u$. Examples can be found in Figure 2. The virtual vertex plays the role of an artificial bottleneck which 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. Fortunately, the multicast capacity from s to T can still be achieved even if each such virtual vertex can only perform routing; the details are given in [10, 11].

Pertinent elementary capacity graphs: By adjusting the transmitting power, the "reach" of a transmitter, i.e., the set of receivers that fall in the transmission range, can be accordingly adjusted. Thus for the problem of minimum-energy multicast, the pertinent ECGs 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 collection or from an arbitrarily interval. We call the former *discrete power control* and the latter *continuous power control*. With the discrete power control model, the total number of pertinent ECGs 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 ECGs 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 mobile ad hoc network may experience variations in the channel strength of individual wireless links during a communication session. Suppose the locations of nodes evolve over time as L discrete *maps*, $M_l, l = 1, \ldots, L$, where each map M_l refers to a configuration of node locations, lasting for a duration of t_l (seconds). In other words, we assume that the nodes stay at the locations M_l for t_l seconds, then instantaneously jump to new locations M_{l+1} at the beginning of next interval.

We model mobility by introducing an additional dimension, time, into the capacity graph. The capacity graph for a mobile network is then comprised of several sequentially concatenated layers, each layer corresponding to a capacity graph for one map. We call such a capacity graph a *time-lined capacity graph*. Each node $v \in V_0$ is now "expanded" into L vertices, $v^1, \ldots, v^L$, one in each layer. The edges in a time-lined capacity graph consist of two type of edges: intra-layer edges and inter-layer edges. The intra-layer edges are edges which originally belong to the capacity graph for each layer. Each of the inter-layer 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 uni-directional because of causality. Hence, we represent a (supported) time-lined capacity graph as

$$G = G^1 \overset{\infty}{\Rrightarrow} G^2 \ldots \overset{\infty}{\Rrightarrow} G^L, \tag{1}$$

$$G^l = \sum_k \lambda_k^l G_k^l, \quad G_k^l \in \mathcal{B}_l \tag{2}$$

$$\sum_k \lambda_k^l \leq \frac{t_l}{\sum_{l=1}^L t_l}, \quad l = 1, \ldots, L, \tag{3}$$

$$\lambda_k^l \geq 0, \quad \forall l, k. \tag{4}$$

where G^l represents a capacity graph in the l-th layer, the notation $G^l \overset{\infty}{\Rrightarrow} G^{l+1}$ characterizes the inter-layer edges introduced, and λ_k^l characterizes the time-sharing proportion for G_k^l, the k-th ECG out of a collection $\mathcal{B}_l$ for the l-th layer. The power consumption of a time-lined capacity graph is

$$P\left(G^1 \overset{\infty}{\Rrightarrow} G^2 \ldots \overset{\infty}{\Rrightarrow} G^L\right) = \sum_{l=1}^{L} \sum_k \lambda_k^l P(G_k^l). \tag{5}$$

III. Minimum-Energy Multicast

In Section II, we discussed the structure of the set of capacity graphs supported by the physical and link layers, in both fixed and mobile ad hoc wireless networks. In this section, we assume a fixed network for clarity of presentation. The extension to the case of a mobile ad hoc network is conceptually the same as the case of a fixed wireless network but is more complicated in terms of notation. Let us denote the set of pertinent capacity graphs as $\mathcal{G}$. We use V to denote the enlarged vertex set and E to denote the union of the edges for $\mathcal{G}$. Each supported (composite) capacity graph, $G \in \mathcal{G}$, consumes power $P(G)$ (Watts). If G can be used to provide a multicast rate of r (bits/second), then the associated energy-per-bit of multicasting is

$$\frac{P(G)}{r} \quad \text{(Joules/bit)}. \tag{6}$$

This section addresses the problem of minimum-energy information multicast by finding a supported capacity graph, $G \in \mathcal{G}$, and a flow assignment on it, achieving the minimum energy-per-bit of multicasting. This can be done by a linear program, because of two key observations: first, the structure of $G \in \mathcal{G}$ can be expressed as a set of linear equalities and inequalities; second, the structure of flow assignments supporting a given multicast rate on a given capacity graph G can be characterized by a set of linear equalities and inequalities.

A. Union of flows

Given a capacity graph $G = (V, E, c)$ and two vertices $s, t \in V$, an *s-t-flow* is a nonnegative-valued function f on edges satisfying the following constraints:

$$0 \leq f(vw) \leq c(vw), \quad \forall vw \in E,$$

$$\sum_{w \in V: vw \in E} f(vw) - \sum_{w \in V: wv \in E} f(wv) = 0, \quad \forall v \in V \backslash \{s, t\},$$

and the *flow value* is:

$$|f| \equiv \sum_{w \in V: sw \in E} f(sw) - \sum_{w \in V: ws \in E} f(ws).$$

We use the notation

$$G' = (V', E', c') \preceq G = (V, E, c) \tag{7}$$

to indicate that (V', E') is a subgraph of (V, E), and $c'(vw) \leq c(vw), \forall vw \in E' \subseteq E$. Given a set of flows $\{f_t, t \in T\}$ on G with f_t being an s-t-flow, the *union of flows* is a subgraph $G_g = (V, E, g) \preceq G$, where the function g is

$$g(vw) \equiv \max_{t \in T} f_t(vw), \quad \forall vw \in E. \tag{8}$$

The subgraph G_g has a multicast capacity greater than or equal to $\min_{t \in T} |f_t|$ because the multicast capacity is equal to the minimum of the maximum flow values from s to each $t \in T$ [4]. Thus *any* union of flows (assuming $|f_t| \geq r, \forall t \in T$) is sufficient to provide a rate r with network coding. Conversely, if a subgraph $G' = (V, E, g') \preceq G$ can provide rate r, then it is easy to see that $\exists G'' = (V, E, g'') \preceq G'$ where G'' is a union of flows providing the same rate r. This shows the necessity of *a* union of flows.

B. Minimum-Energy Multicast

With the necessary and sufficient properties of a union of flows, the minimum-energy multicast problem can now be formulated as the following optimization, where r, $f_t(vw)$, $g(vw)$, λ_k, are treated as variables

$$\min \frac{\sum_k \lambda_k P(G_k)}{r} \quad \text{subject to:} \tag{9}$$

$$g(vw) \leq \sum_k \lambda_k c_k(vw), \forall vw \in E \tag{10}$$

$$\sum_k \lambda_k \leq 1, \tag{11}$$

$$\lambda_k \geq 0, \forall k, \tag{12}$$

$$0 \leq f_t(vw) \leq g(vw), \forall vw \in E, \ \forall t \in T \tag{13}$$

$$\sum_{w \in V: vw \in E} f_t(vw) - \sum_{w \in V: wv \in E} f_t(wv) = 0, \quad \forall v \in V \backslash \{s,t\}, \ \forall t \in T \tag{14}$$

$$\sum_{w \in V: sw \in E} f_t(sw) - \sum_{w \in V: ws \in E} f_t(ws) = r, \forall t \in T \tag{15}$$

At first glance, the objective function of the above optimization is nonlinear in the variables. However, it should be easy to see that the minimum value of the above optimization can be achieved at a sufficiently small rate r such that constraint (11) become loose. With some simple mathematical manipulations, we can re-normalize the above optimization to arrive at a linear program. This is formally stated in the following Theorem 1.

Theorem 1 (Minimum-Energy Multicast)
In a fixed wireless ad hoc network, let $\mathcal{B} = \{G_k\}$ denote the set of pertinent ECGs for the problem of minimum energy multicast, according to a (discrete or continuous) power control model. The set of capacity graphs spanned by $\mathcal{B}$ is

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

The minimum energy-per-bit for multicasting from s to T, β^*, is given by the following linear program

$$\beta^* = \min \sum_k \lambda_k P(G_k) \quad \textit{subject to:} \tag{16}$$

$$g(vw) \leq \sum_k \lambda_k c_k(vw), \forall vw \in E \tag{17}$$

$$\lambda_k \geq 0, \forall k, \tag{18}$$

$$0 \leq f_t(vw) \leq g(vw), \forall vw \in E, \ \forall t \in T \tag{19}$$

$$\sum_{w \in V: vw \in E} f_t(vw) - \sum_{w \in V: wv \in E} f_t(wv) = 0, \quad \forall v \in V \backslash \{s,t\}, \ \forall t \in T \tag{20}$$

$$\sum_{w \in V: sw \in E} f_t(sw) - \sum_{w \in V: ws \in E} f_t(ws) = 1, \forall t \in T. \tag{21}$$

The minimum energy-per-bit can be attained by performing network coding on a subgraph $G^* = (V, E, g^*) \preceq G = (V, E, c)$, with

$$g^*(vw) = \frac{1}{\sum_k \lambda_k} g(vw), \forall vw \in E,$$

where $\{g(vw), \forall vw \in E\}$, $\{\lambda_k\}$ are the solutions returned by the linear program in (16)-(21).

IV. Simulations

We conduct simulations on an example community wireless network with 58 houses spaced roughly in 3 rows, shown in Figure 3. The locations of the houses are marked with dots. We consider the multicast session from node 28 to the set of receiver nodes $\{1, 6, 52, 53\}$, as illustrated in Figure 3 by four lines.

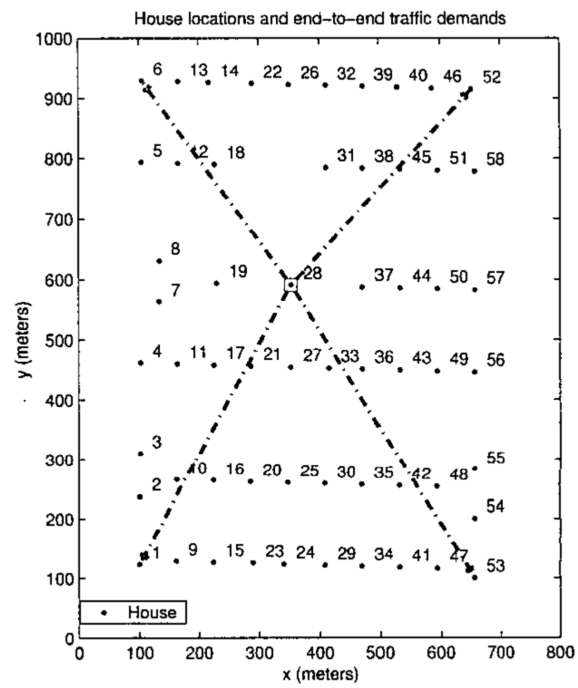

Figure 3: An example community wireless network.

Some simulation parameters are set up as follows. The SNR threshold γ is set to be 4dB. The path loss coefficients are set as $a_{ij} = d(i,j)^{-3}$ where $d(i,j)$ is the distance between node i and j. 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 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 300m.

The connectivity graph for the traffic assignment given by the minimum-energy multicast solution is shown in Figure 4. It happens that the minimum-energy solution of this problem admits an easy interpretation. Suppose the source stream is partitioned

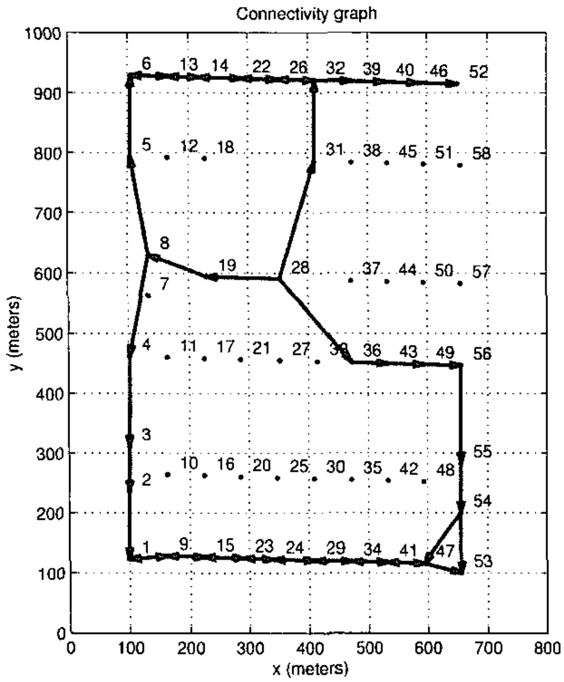

Figure 4: The connectivity graph for the traffic assignment given by the minimum-energy multicast solution.

into two distinct substreams, $x_1(n)$ and $x_2(n)$, $n = 1 \ldots, N$, which we can think of as odd subsequence and even subsequence, for example. Then, substream $x_1(n)$ is delivered from source node 28 to node 1 and node 6 through a sequence of links: $28 \to 19$, $19 \to 8$, $8 \to \{4, 5\}$, $5 \to 6$, $4 \to 3$, $3 \to 2$, $2 \to 1$. Substream x_2 is delivered to node 32, node 47, and node 53, through a sequence of links: $28 \to \{31, 36\}$, $31 \to 32$, $36 \to 43$, $43 \to 49$, $49 \to 56$, $56 \to 55$, $55 \to 54$, $54 \to \{47, 53\}$. Now node 6 has substream $x_1(n)$ and node 32 has substream $x_2(n)$ and thus node 6 and node 32 need to exchange information. The information exchange between node 6 and node 32 can be done efficiently by exploiting network coding together with physical layer broadcast. This is illustrated by the graph representation in Figure 5. All of the edges in Figure 5 have the same capacity. There is a path from node 6 to node 32:

$$6, 6', 13, 13', 14, 14', 22, 22', 26, 26', 32, 32'$$

and a path from node 32 to node 6:

$$32, 32', 26, 26', 22, 22', 14, 14', 13, 13', 6, 6'.$$

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 Figure 5 coincides with the simple example in Figure 2. Similar mutual exchange phenomenon also occurs at the bottom of Figure 4 between node 1 and node 47. It can be generalized from this case study that network coding, combined with physical layer broadcast, can facilitate mutual exchange of information in a wireless ad hoc network.

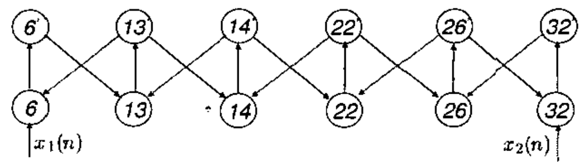

Figure 5: Efficient information exchange between node 6 and node 32.

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. on Computer Comm. (INFOCOM)*, pp. 586-594, 2000.

[2] W. Liang, "Constructing minimum-energy broadcast trees in wireless ad hoc networks," In *Proc. ACM Int'l Symp. on Mobile Ad hoc Networking and Computing (MOBIHOC)*, Lausanne, Switzerland, June 9-11, 2002.

[3] M. Cagalj, J.-P. Hubaux and C. Enz, "Minimum-energy broadcast in all-wireless networks: NP-completeness and distribution issues," In *Proc. Int'l Conf. on Mobile Computing and Networks (MOBICOM)*, pp.172-182, 2002.

[4] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Information Theory*, IT-46(4): 1204-1216, Jul. 2000.

[5] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Trans. Information Theory*, IT-49(2):371-381, Feb. 2003.

[6] R. Koetter, and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Trans. Networking*, vol. 11, no. 5, pp. 782-795, Oct. 2003.

[7] Y. Wu, P. A. Chou, K. Jain, "A comparison of network coding and tree packing," in *Proc. IEEE Int'l Symp. Information Theory (ISIT)*, Chicago, June 2004.

[8] T. Ho, B. Leong, M. Médard, R. Koetter, Y. Chang, M. Effros, "On the utility of network coding in dynamic environments", in *Proc. Int'l Workshop on Wireless Ad-hoc Networks (IWWAN)*, June 2004.

[9] A. El Gamal, S. Zahedi, "Minimum energy communication over a relay channel," in *Proc. Int'l Symp. Information Theory*, June, 2003, Yokohama, Japan, p. 344.

[10] 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," to appear in *IEEE J. Selected Areas in Comm.*, 2005.

[11] Y. Wu, P. A. Chou, and S.-Y. Kung, "Minimum-energy multicast in mobile ad hoc networks using network coding," submitted to *IEEE Trans. Communications*, Mar. 2004.

[12] D. S. Lun, M. Médard, T. Ho, and R. Koetter, "Network coding with a cost criterion," Technical Report LIDS-P-2584, MIT LIDS, Apr. 2004. Also in *Proc. Int'l Symp. Information Theory and its Applications (ISITA2004)*, Oct. 2004.