Efficient Operation of Wireless Packet Networks Using Network Coding

Desmond S. Lun, Muriel Médard, and Ralf Koetter

(Invited Paper)

***Abstract*— The marriage of network coding and wireless packet networks is a natural and attractive one. So, despite network coding being a still nascent field, a considerable body of work on this subject already exists. In this paper, we give a brief overview of this work and hope, thereby, to provide the reader with a firm theoretical basis from which practical implementations and theoretical extensions can be developed.**

***Index Terms*— Ad hoc networks, energy efficient, multicast, network coding, wireless networks**

I. INTRODUCTION

THE notion of coding at the packet level—commonly called network coding—has attracted significant interest since the publication of [1], which showed the utility of network coding for multicast in wireline packet networks. But the utility of network coding reaches much further. And, in particular, it reaches to include both unicast and multicast in wireless packet networks. In fact, wireless packet networks are a most natural setting for network coding because the very characteristics of wireless links that complicate routing, namely, their unreliability and broadcast nature, are the very characteristics for which coding is a natural solution. Couple this with the fact that we are not nearly as constrained in our protocol design choices in the wireless case as we are in the wireline one, and applying network coding to wireless packet networks seems an ideal way of achieving appreciable efficiency gains.

The attractiveness of marrying network coding and wireless packet networks has not escaped the notice of researchers, and numerous papers (e.g., [2]–[13]) have appeared on the subject. The aim of this paper is to synthesize this prior work, with emphasis on the work presented in [5], [7], [9]. We hope, thereby, to provide the reader with a firm theoretical basis for tackling the problem of efficient operation of coded wireless packet networks, from which practical implementations and theoretical extensions can be developed.

We begin with lossless networks, i.e. networks with broadcast links that are effectively lossless, presumably owing to some underlying retransmission scheme. In this case, the main utility of network coding is for multicast—for single unicast connections, coding provides no advantage over routing. In Section III, we move onto lossy networks, where coding is advantageous for single unicast and multicast connections.

This work was supported by the National Science Foundation under grant nos. CCR-0093349, CCR-0325496, and CCR-0325673, by Hewlett-Packard under grant no. 008542-008, and by the Air Force Office of Scientific Research under grant no. F49620-01-1-0365.

D. S. Lun and M. Médard are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, USA (e-mail: dslun@mit.edu, medard@mit.edu).

R. Koetter is with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA (e-mail: koetter@uiuc.edu).

II. LOSSLESS NETWORKS

The main aspect by which lossless wireless networks differ from the lossless wireline networks considered in the original work on network coding (e.g., [1], [14], [15]) is the broadcast nature of the links, i.e. wireless links are either omnidirectional or directed over a large area, meaning that transmissions are often received by more than one node and that transmissions originating from separate nodes may interfere. This characteristic has a significant effect on many network problems, particularly multicast, where it has been given the name, the "wireless multicast advantage", in [16].

For routing, the multicast advantage makes performing energy-efficient multicast more difficult; so much so, in fact, that the problem of minimum-energy broadcast, solved easily in the wireline case by various minimum-weight spanning tree algorithms, becomes NP-complete [17], [18]. Since broadcast is a special case of multicast, the problem of minimum-energy multicast is at least as complex, which forces us to employ heuristics, such as the Multicast Incremental Power (MIP) algorithm described in [16], to address the problem.

For coding, however, the situation is much better: The problem of minimum-energy multicast can be solved exactly in polynomial time. And, since coding subsumes routing, the optimal energy in the former case is generally less than that in the latter. Therefore, coding promises to significantly outperform routing for practical multicast. In simulations, we observed reductions ranging from 13% to 49% in the average total energy of random multicast connections in random wireless networks of varying size as a result of coding as opposed to routing with the MIP algorithm (see Table I).

But a polynomial-time solution is not enough. For many applications, a solution not requiring centralized computation or knowledge is desired. As we shall see, a decentralized solution exists for the coding case, and it comes from coupling decentralized optimization and decentralized, random coding.

A. Model

We model the network with a directed hypergraph $\mathcal{H} = (\mathcal{N}, \mathcal{A})$, where $\mathcal{N}$ is the set of nodes and $\mathcal{A}$ is the set of hyperarcs. A hypergraph is a generalization of a graph, where, rather than arcs, we have hyperarcs. A hyperarc is a pair (i, J), where i, the start node, is an element of $\mathcal{N}$ and J, the set of end nodes, is a non-empty subset of $\mathcal{N}$.

Network size	Approach	Average multicast energy			
		2 sinks	4 sinks	8 sinks	16 sinks
20 nodes	MIP algorithm	30.6	33.8	41.6	47.4
	Network coding	15.5	23.3	29.9	38.1
30 nodes	MIP algorithm	26.8	31.9	37.7	43.3
	Network coding	15.4	21.7	28.3	37.8
40 nodes	MIP algorithm	24.4	29.3	35.1	42.3
	Network coding	14.5	20.6	25.6	30.5
50 nodes	MIP algorithm	22.6	27.3	32.8	37.3
	Network coding	12.8	17.7	25.3	30.3

TABLE I

AVERAGE ENERGY OF RANDOM MULTICAST CONNECTIONS OF UNIT RATE FOR VARIOUS APPROACHES IN RANDOM WIRELESS NETWORKS OF VARYING SIZE. NODES WERE PLACED RANDOMLY WITHIN A 10×10 SQUARE WITH A RADIUS OF CONNECTIVITY OF 3. THE ENERGY REQUIRED TO TRANSMIT AT UNIT RATE TO A DISTANCE d WAS TAKEN TO BE d^2. SOURCE AND SINK NODES WERE SELECTED ACCORDING TO AN UNIFORM DISTRIBUTION OVER ALL POSSIBLE SELECTIONS.

Each hyperarc (i, J) represents a lossless broadcast link from node i to nodes in the non-empty set J. We denote by z_{iJ} the average rate at which packets are injected and received on hyperarc (i, J). The rate vector z is called the coding subgraph and can be varied within a constraint set Z dictated to us by lower layers. We reasonably assume that Z is a convex subset of the positive orthant containing the origin.

To reflect the notion of efficiency, we associate with the network a cost function f that maps feasible rate vectors to real numbers and that we seek to minimize. For wireless networks, it is common for the cost function to reflect energy consumption, but it could also represent, for example, average latency, monetary cost, or a combination of these considerations.

B. Single multicast connections

Given a coding subgraph z, it is shown in [5] that a multicast of rate arbitrarily close to R is achievable with coding from source node s to sink nodes in the set T if and only if there exists, for all $t \in T$, a flow vector $x^{(t)}$ satisfying

$$\sum_{\{J|(i,J)\in\mathcal{A}\}} \sum_{j\in J} x_{iJj}^{(t)} - \sum_{\{j|(j,I)\in\mathcal{A}, i\in I\}} x_{jIi}^{(t)} = \sigma_i^{(t)} \qquad (1)$$

for all $i \in \mathcal{N}$, and

$$\sum_{j\in J} x_{iJj}^{(t)} \leq z_{iJ} \qquad (2)$$

for all $(i, J) \in \mathcal{A}$, where

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

From such a coding subgraph, the connection can be straightforwardly achieved using the decentralized, random coding schemes in [19], [20], [5] or by modifying the deterministic coding schemes in [21]–[23].

Therefore, to establish a minimum-cost multicast connection in a lossless wireless packet network, we simply need to find a coding subgraph z of minimum cost satisfying (1) and (2), then apply a coding scheme to it. These two problems, the subgraph selection problem and the coding problem, are essentially decoupled and can be handled separately.

The coding problem we have already mentioned. The subgraph selection problem equates to solving the following optimization problem.

$$\begin{aligned} &\text{minimize } f(z) \\ &\text{subject to } z \in Z, \\ &z_{iJ} \geq \sum_{j\in J} x_{iJj}^{(t)}, \qquad \forall\ (i,J) \in \mathcal{A},\ t \in T, \\ &\sum_{\{J|(i,J)\in\mathcal{A}\}} \sum_{j\in J} x_{iJj}^{(t)} - \sum_{\{j|(j,I)\in\mathcal{A}, i\in I\}} x_{jIi}^{(t)} = \sigma_i^{(t)}, \\ &\qquad\qquad \forall\ i \in \mathcal{N},\ t \in T, \\ &x_{iJj}^{(t)} \geq 0, \qquad \forall\ (i,J) \in \mathcal{A},\ j \in J,\ t \in T. \end{aligned} \qquad (3)$$

As an example, consider the wireless network depicted in Figure 1(a). We wish to achieve multicast of unit rate to two sinks, t_1 and t_2. We have $Z = [0,1]^{|\mathcal{A}|}$ and $f(z) = \sum_{(i,J)\in\mathcal{A}} a_{iJ} z_{iJ}$, where a_{iJ} is the cost per unit rate shown beside each hyperarc. An optimal solution to problem (3) for this example is shown in Figure 1(b). We have flows $x^{(1)}$ and $x^{(2)}$ of unit size from s to t_1 and t_2, respectively, and, for each hyperarc (i, J), $z_{iJ} = \max(\sum_{j\in J} x_{iJj}^{(1)}, \sum_{j\in J} x_{iJj}^{(2)})$, as we expect from the optimization. To achieve the optimal cost, we code over the subgraph defined by z. A code of length 2 for the subgraph is shown in Figure 1(c). In the figure, X_1 and X_2 refer to the two packets in a coding block. The coding that is performed is that one of the interior nodes receives both X_1 and X_2 and forms the binary sum of the two, outputting the packet $X_1 + X_2$. The code allows both t_1 and t_2 to recover both X_1 and X_2, and it achieves the optimal cost of $5/2$. For routing, even if we were able to solve the NP-complete problem required to find an optimal multicast tree, the optimal cost would be 3.

Problem (3) as it stands is not easy to solve. Its complexity improves if we assume that the cost function is separable and convex, or even linear; i.e. if we suppose $f(z) = \sum_{(i,J)\in A} f_{iJ}(z_{iJ})$, where f_{iJ} is a convex or linear function. This is a very reasonable assumption in many situations.

But the constraint set Z still makes the optimization difficult. This constraint set not only necessarily induces a degree of coupling among separate links because of contention for the wireless medium, it is also usually difficult to describe. One simple, heuristic approach is to find a set of feasible

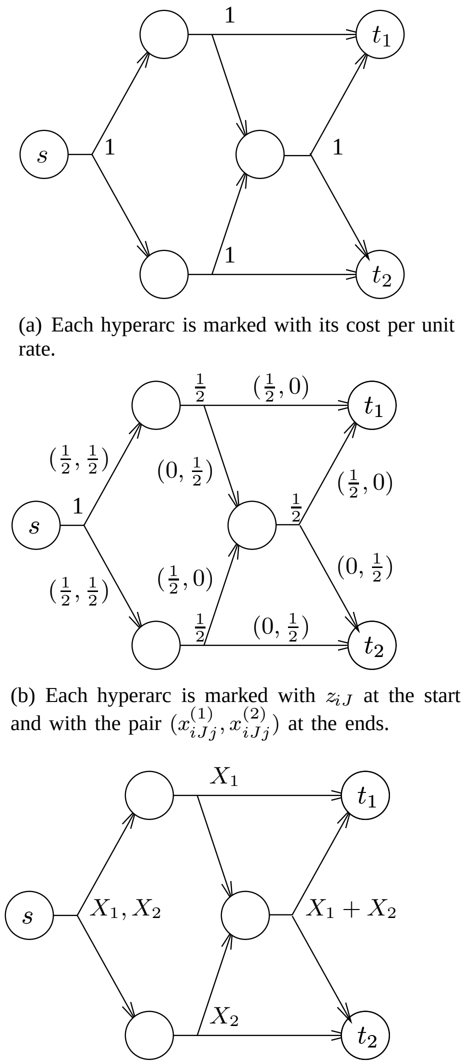

(a) Each hyperarc is marked with its cost per unit rate.

(b) Each hyperarc is marked with z_{iJ} at the start and with the pair $(x_{iJj}^{(1)}, x_{iJj}^{(2)})$ at the ends.

(c) Each hyperarc is marked with its code.

Fig. 1. A wireless network with multicast from s to $T = \{t_1, t_2\}$.

constraint sets $\{Z_1, Z_2, \ldots, Z_N\}$ (for example, each Z_n could correspond to the links established by some set of non-interfering transmitters), then take

$$Z = \sum_{n=1}^{N} \alpha_n Z_n,$$

where $\sum_{n=1}^{N} \alpha_n = 1$ and $\alpha \geq 0$, and include α in the variables to be optimized. The value of α obtained from the optimization can then be used to compute a schedule. This approach is taken in [13], [10], where simple heuristics for finding $\{Z_1, Z_2, \ldots, Z_N\}$ are given.

The constraint set Z poses less of a problem in low-energy systems because nodes seldom transmit and there is less contention. In fact, if energy is the most limiting constraint and we wish to achieve minimum-energy multicast without regard for throughput or bandwidth, which is the problem considered in [16], [17], [18], then Z can be dropped altogether. Hence problem (3) can be rewritten as the following linear optimization problem.

$$\begin{aligned}
&\text{minimize } \sum_{(i,J)\in A} a_{iJ} z_{iJ} \\
&\text{subject to} \\
&z_{iJ} \geq \sum_{j\in J} x_{iJj}^{(t)}, \qquad \forall\ (i,J) \in \mathcal{A},\ t \in T, \\
&\sum_{\{J|(i,J)\in\mathcal{A}\}} \sum_{j\in J} x_{iJj}^{(t)} - \sum_{\{j|(j,I)\in\mathcal{A}, i\in I\}} x_{jIi}^{(t)} = \sigma_i^{(t)}, \\
&\qquad\qquad \forall\ i \in \mathcal{N},\ t \in T, \\
&x_{iJj}^{(t)} \geq 0, \qquad \forall\ (i,J) \in \mathcal{A},\ j \in J,\ t \in T,
\end{aligned} \tag{4}$$

where a_{iJ} represents the energy required to transmit to nodes in J from node i for some fixed time interval. Thus the minimum-energy multicast problem in coded networks can be solved in polynomial time.

Problem (4) can, however, still be simplified further. An assumption that is commonly made in minimum-energy multicast is that, when nodes transmit, they reach all nodes within a certain radius that is a function of their transmission power. So, if node i has M_i power levels that induce broadcast links represented by the hyperarcs $(i, J_1), (i, J_2), \ldots, (i, J_{M_i})$ with increasing power, then $J_1 \subsetneq J_2 \subsetneq \cdots \subsetneq J_{M_i}$. To simplify the optimization, we make the assumption that $J_1 \subsetneq J_2 \subsetneq \cdots \subsetneq J_{M_i}$. This is a more general assumption than the assumption that transmissions are necessarily received over a circular region.

Let $J_1^{(i)}, J_2^{(i)}, \ldots, J_{M_i}^{(i)}$ be the M_i sets of nodes reached by node i with its M_i power levels. Then, for $(i,j) \in \mathcal{A}' := \{(i,j)|(i,J) \in A, J \ni j\}$, we introduce the variables

$$\tilde{x}_{ij}^{(t)} := \sum_{m=m(i,j)}^{M_i} x_{iJ_m^{(i)}j}^{(t)},$$

where $m(i,j)$ is the unique m such that $j \in J_m^{(i)} \setminus J_{m-1}^{(i)}$ (we define $J_0^{(i)} := \emptyset$ for all $i \in \mathcal{N}$ for convenience). Then, provided that $a_{iJ_1^{(i)}} < a_{iJ_2^{(i)}} < \cdots < a_{iJ_{M_i}^{(i)}}$ for all nodes i, problem (4) can be reformulated as the following optimization problem, which was used to obtain the results in Table I.

$$\begin{aligned}
&\text{minimize } \sum_{(i,J)\in A} a_{iJ} z_{iJ} \\
&\text{subject to} \\
&\sum_{n=m}^{M_i} z_{iJ_n^{(i)}} \geq \sum_{k\in J_{M_i}^{(i)} \setminus J_{m-1}^{(i)}} \tilde{x}_{ik}^{(t)}, \\
&\qquad\qquad \forall\ i \in \mathcal{N},\ m = 1, \ldots, M_i,\ t \in T, \\
&\sum_{\{j|(i,j)\in\mathcal{A}'\}} \tilde{x}_{ij}^{(t)} - \sum_{\{j|(j,i)\in\mathcal{A}'\}} \tilde{x}_{ji}^{(t)} = \sigma_i^{(t)}, \\
&\qquad\qquad \forall\ i \in \mathcal{N},\ t \in T, \\
&\tilde{x}_{ij}^{(t)} \geq 0, \qquad \forall\ (i,j) \in \mathcal{A}',\ t \in T.
\end{aligned} \tag{5}$$

We omit the justification of the reformulation for the sake of brevity.

Problem (5) has significantly fewer variables than problem (4) and can be solved using a decentralized algorithm, which is described in full in [8], [9]. This decentralized algorithm requires each node only to know the costs of its incoming and outgoing hyperarcs and to communicate the results of its computations with its neighbors. When coupled with the decentralized coding schemes that we mentioned, the algorithm yields a fully decentralized approach to minimum-energy multicast in lossless wireless packet networks.

III. Lossy networks

The argument for coding is even stronger in lossy networks than it is in lossless networks. Recall that, in lossless networks, coding provides no advantage over routing for single unicast connections—it does only for multicast ones. But in lossy networks, coding provides an efficiency advantage even for single unicast connections. And it also provides robustness. Thus, reliability can be provided without feedback, obviating the need to schedule retransmission requests and to manage the complicated interactions between retransmission protocols at different layers [24]. Such a feedforward approach to reliability is reminiscent of the digital fountain approach [25]. In the digital fountain approach, however, coding is performed only at the source, whereas, in a network coding approach, coding is performed by all nodes, achieving higher throughput in general [5].

A. Model

The model for lossy networks differs from that for lossless networks in only one aspect: the interpretation of the coding subgraph z. Let A_{iJK} be the counting process describing the arrival of packets that are injected on hyperarc (i, J) and received by the set of nodes $K \subset J$, i.e. for $\tau \geq 0$, $A_{iJK}(\tau)$ is the total number of packets received between time 0 and time τ by all nodes in K due to (i, J). Then, we assume

$$\lim_{\tau\to\infty} \frac{A_{iJK}(\tau)}{\tau} = z_{iJK} \tag{6}$$

almost surely, i.e. a long-run average rate exists and is equal to z_{iJK} with probability 1.

B. Single multicast connections

Since single unicast connections are a special case of single multicast connections, we discuss only the latter. Given a coding subgraph z, it is shown in [5], [7] that a multicast of rate arbitrarily close to R is achievable with coding from source node s to sink nodes in the set T if and only if there exists, for all $t \in T$, a flow vector $x^{(t)}$ satisfying

$$\sum_{\{J|(i,J)\in\mathcal{A}\}} \sum_{j\in J} x_{iJj}^{(t)} - \sum_{\{j|(j,I)\in\mathcal{A}, i\in I\}} x_{jIi}^{(t)} = \sigma_i^{(t)}$$

for all $i \in \mathcal{N}$, and

$$\sum_{j\in K} x_{iJj}^{(t)} \leq \sum_{\{L\subset J|L\cap K\neq\emptyset\}} z_{iJL}$$

for all $(i, J) \in \mathcal{A}$ and $K \subset J$.[1] From such a coding subgraph, the only practical scheme we currently have for achieving the connection is the one described in [5], [7]. (For a more restricted model, schemes for demonstrating achievability and for some special networks are given in [2] and [3], [4], respectively.)

[1] Achievability in [5], [7] is shown under more restrictive conditions on A_{iJK} than the ones we currently assume. We have recently extended the result to the current conditions.

We give a brief description of the scheme. We suppose that, at the source node s, we have K message packets $w_1, w_2, \ldots, w_K$, which are vectors of length ρ over the finite field $\mathbb{F}_q$. (If the packet length is b bits, then we take $\rho = \lceil b/\log_2 q\rceil$.) The message packets are initially present in the memory of node s.

The coding operation performed by each node is simple to describe and is the same for every node: Received packets are stored into the node's memory, and packets are formed for injection with random linear combinations of its memory contents whenever a packet injection occurs on an outgoing hyperarc. The coefficients of the combination are drawn uniformly from $\mathbb{F}_q$. Since all coding is linear, we can write any packet x in the network as a linear combination of $w_1, w_2, \ldots, w_K$, namely, $x = \sum_{k=1}^{K} \gamma_k w_k$. We call γ the *global encoding vector* of x, and we assume that it is sent along with x, in its header. The overhead this incurs (namely, $K \log_2 q$ bits) is negligible if packets are sufficiently large.

A sink node collects packets and, if it has K packets with linearly-independent global encoding vectors, it is able to recover the message packets. Decoding can be done by Gaussian elimination. In addition, the scheme can be operated ratelessly, i.e. it can be run indefinitely until all sink nodes in T can decode (at which stage that fact is signaled to all nodes, requiring only a small amount of feedback).

The scheme is remarkably robust: If run for a sufficiently long period of time, it achieves the maximum feasible rate of a given subgraph, with the only assumption on the arrival of received packets on a link being (6). Assumption (6) makes no claims on burstiness or lack thereof—all we require is that a long-run average rate exists. This fact is particularly important in wireless packet networks, where slow fading and collisions often cause packets not be received in a steady stream.

We now turn to finding the coding subgraph. An optimal coding subgraph for minimum-cost multicast in lossy wireless packet networks is found by solving the following optimization problem.

$$\begin{aligned}
&\text{minimize } f(z) \\
&\text{subject to } z \in Z, \\
&\sum_{j\in K} x_{iJj}^{(t)} \leq \sum_{\{L\subset J|L\cap K\neq\emptyset\}} z_{iJL}, \\
&\qquad\qquad \forall\ (i,J)\in\mathcal{A},\ K\subset J,\ t\in T, \\
&\sum_{\{J|(i,J)\in\mathcal{A}\}} \sum_{j\in J} x_{iJj}^{(t)} - \sum_{\{j|(j,I)\in\mathcal{A}, i\in I\}} x_{jIi}^{(t)} = \sigma_i^{(t)}, \\
&\qquad\qquad \forall\ i\in\mathcal{N},\ t\in T, \\
&x_{iJj}^{(t)} \geq 0, \qquad \forall\ (i,J)\in\mathcal{A},\ j\in J,\ t\in T.
\end{aligned} \tag{7}$$

Problem (7) is very similar to problem (3)—the analogous problem in the lossless case. But the difference, in the constraints relating x and z, is an important one: In problem (7), there is, for each $(i, J) \in \mathcal{A}$ and $t \in T$, a constraint for each $K \subset J$. Therefore, the number of constraints grows exponentially with $|J|$. If the number of other nodes reached by each node is small, then solving problem (7) may be practicable; otherwise, we must search for special cases in which problem (7) simplifies or for heuristics to find approximate solutions. Finding such simplifications and heuristics is a clear avenue for further work.

IV. CONCLUSION

We have sought to provide in this paper a brief overview of recent developments in network coding, in particular as they relate to wireless settings. That network coding affords energy savings, which are considerable in certain settings, may be its most immediate benefit over traditional wireless routing approaches. That it also affords the ability to operate networks in a decentralized fashion, while not as easily quantifiable as energy savings, may, however, hold the most promise for improving the workings of wireless networks. Such decentralized approaches may altogether change the types of protocols used for wireless networking.

We have addressed issues of wireless transmission in general terms, and, owing to space and scope considerations, have not explicitly considered interference, mobility, nor issues of efficient information distribution. Interference depends on various factors of the overall system, particularly the medium access protocol, and, while simple collisions may be adequately handled by the lossy network model, network coding and multiple access issues at the physical layer will in general interact in a manner that is not yet understood.

Mobility induces changes in topology and, possibly, traffic patterns. A well-constructed system should not only be able to respond to changes in a robust manner but, when provided with some predictive information regarding movement or traffic trends, be able to account in balanced fashion for both current and future network states and demands. Some initial work [26], [9] indicates that using network coding in dynamic environments may lead to dynamic programming approaches that naturally extend the static optimization techniques discussed in this paper. Whether such dynamic programming is amenable to decentralized implementations and how it may be translated effectively into protocols has not been investigated.

In sensor networks, our goal may not merely be to provide efficient transmission of the data emanating from sensors to some set of target nodes, but rather to convey the total information. Thus, when the data in nearby sensors is appreciably correlated, efficient transmission may require intelligent compression that actively takes into account information redundancy across sensors. There exists an intimate link between coding for distributed compression, i.e. the classical Slepian-Wolf problem, and network coding. Indeed, decentralized, random network coding may be used in an optimal fashion to provide distributed compression in a network, providing results which naturally generalize Slepian-Wolf approaches to the networked setting [27], [19]. Combining such distributed compression with energy efficiency remains to be investigated.

REFERENCES

[1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inform. Theory*, vol. 46, no. 4, pp. 1204–1216, July 2000.

[2] A. F. Dana, R. Gowaikar, R. Palanki, B. Hassibi, and M. Effros, "Capacity of wireless erasure networks," submitted to *IEEE Trans. Inform. Theory*.

[3] R. Khalili and K. Salamatian, "On the achievability of cut-set bound for a class of erasure relay channels: The non degraded case," in *Proc. 2004 International Symposium on Information Theory and its Applications (ISITA 2004)*, Oct. 2004, pp. 269–274.

[4] ——, "On the capacity of multiple input erasure relay channels," in *Proc. WINMEE, RAWNET and NETCOD 2005 Workshops*, Apr. 2005.

[5] D. S. Lun, M. Médard, and M. Effros, "On coding for reliable communication over packet networks," in *Proc. 42nd Annual Allerton Conference on Communication, Control, and Computing*, Sept.–Oct. 2004, invited paper.

[6] D. S. Lun, M. Médard, T. Ho, and R. Koetter, "Network coding with a cost criterion," in *Proc. 2004 International Symposium on Information Theory and its Applications (ISITA 2004)*, Oct. 2004, pp. 1232–1237.

[7] D. S. Lun, M. Médard, R. Koetter, and M. Effros, "Further results on coding for reliable communication over packet networks," submitted to *2005 IEEE International Symposium on Information Theory (ISIT 2005)*.

[8] D. S. Lun, N. Ratnakar, R. Koetter, M. Médard, E. Ahmed, and H. Lee, "Achieving minimum-cost multicast: A decentralized approach based on network coding," in *Proc. IEEE Infocom 2005*, Miami, FL, Mar. 2005.

[9] D. S. Lun, N. Ratnakar, M. Médard, R. Koetter, D. R. Karger, T. Ho, and E. Ahmed, "Minimum-cost multicast over coded packet networks," submitted to *IEEE Trans. Inform. Theory*.

[10] Y. E. Sagduyu and A. Ephremides, "Joint scheduling and wireless network coding," in *Proc. WINMEE, RAWNET and NETCOD 2005 Workshops*, Apr. 2005.

[11] J. Widmer, C. Fragouli, and J.-Y. Le Boudec, "Low-complexity energy-efficient broadcasting in wireless ad-hoc networks using network coding," in *Proc. WINMEE, RAWNET and NETCOD 2005 Workshops*, Apr. 2005.

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

[13] 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. Select. Areas Commun.*, vol. 23, no. 1, pp. 136–150, Jan. 2005.

[14] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Trans. Inform. Theory*, vol. 49, no. 2, pp. 371–381, Feb. 2003.

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

[16] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides, "Energy-efficient broadcast and multicast trees in wireless networks," *Mobile Networks and Applications*, vol. 7, pp. 481–492, 2002.

[17] A. Ahluwalia, E. Modiano, and L. Shu, "On the complexity and distributed construction of energy-efficient broadcast trees in static ad hoc wireless networks," in *Proc. 2002 Conference on Information Sciences and Systems (CISS 2002)*, Mar. 2002.

[18] W. Liang, "Constructing minimum-energy broadcast trees in wireless ad hoc networks," in *Proc. 3rd ACM International Symposium on Mobile Ad Hoc Networking & Computing (MOBIHOC '02)*, 2002, pp. 112–122.

[19] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong, "Toward a random operation of networks," submitted to *IEEE Trans. Inform. Theory*.

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

[21] N. J. A. Harvey, D. R. Karger, and K. Murota, "Deterministic network coding by matrix completion," in *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*, Jan. 2005.

[22] C. Fragouli and E. Soljanin, "Information flow decomposition for network coding," submitted to *IEEE Trans. Inform. Theory*.

[23] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen, "Polynomial time algorithms for multicast network code construction," submitted to *IEEE Trans. Inform. Theory*.

[24] C. Liu and E. Modiano, "On the interaction between layered protocols: The case of window flow control and ARQ," in *Proc. 2002 Conference on Information Sciences and Systems (CISS 2002)*, Mar. 2002.

[25] M. Luby, "LT codes," in *Proc. 43rd Annual IEEE Symposium on Foundations of Computer Science*, Nov. 2002, pp. 271–280.

[26] D. S. Lun, M. Médard, and D. R. Karger, "On the dynamic multicast problem for coded networks," in *Proc. WINMEE, RAWNET and NETCOD 2005 Workshops*, Apr. 2005.

[27] T. C. Ho, M. Médard, and M. Effros, "Network coding for correlated sources," in *Proc. 2004 Conference on Information Sciences and Systems (CISS 2004)*, Mar. 2004, invited paper.