

On the Benefits of Network Coding for Wireless Applications

Christina Fragouli

EPFL

Lausanne, Switzerland

Email: christina.fragouli@epfl.ch

Jörg Widmer

DoCoMo Euro-Labs

Munich, Germany

Email: widmer@docomolab-euro.com

Jean-Yves Le Boudec

EPFL

Lausanne, Switzerland

Email: jean-yves.leboudec@epfl.ch

Abstract— We argue that the main benefits of network coding in a wireless environment might manifest in situations where the topology dynamically changes, and operation is restricted to distributed algorithms that do not employ knowledge about the network environment. We consider several problem instances in this set-up, that include broadcasting information to all nodes of the network and collecting sensor measurements. We show that in many such cases, under some simplifying assumptions, the problem is theoretically equivalent to simple variations of the coupon collector problem. Thus network coding can offer benefits of a factor of $\log n$, where n is the number of nodes and the benefits are in terms of energy efficiency, as was proven in [3]. We present simulation results under more realistic conditions that support this claim.

I. INTRODUCTION

Deployment of network coding [1], [11] in wireless networks can help to better exploit shared resources such as wireless bandwidth, and to conserve scarce resources, such as battery life. In this context, we examined in [8], [16] the benefits in terms of energy efficiency that network coding may offer for the problem of broadcasting in wireless ad-hoc networks. We measure energy efficiency as energy expenditure per successfully transmitted information bit. By broadcasting we mean that every node is a source that wants to transmit information to all other nodes. The problem of broadcasting is interesting not only because it abstracts diverse practical applications, but also because this is a situation where information mixing is clearly beneficial and where we thus expect network coding to offer most benefits.

In [8], [16] we showed that network coding reduces the required energy by a factor of $1/2$ and $3/4$ for the circular network and the square grid network respectively. Since large random networks tend to perform like a square grid, one would expect that for general wireless networks, network coding offers similar benefits. In fact, it is straightforward to show that for the problem of broadcasting over fixed networks, network coding can only offer constant benefits in terms of energy efficiency [7].

In this paper, we argue that the main benefits network coding offers might be in terms of operational complexity in dynamically changing environments. We consider a wireless environment, where the network configuration constantly changes, because nodes move, turn on and off, roam out of range, etc. We focus our attention to very simple decentralized

distributed algorithms, where nodes do not know the identity of their neighbors. Our motivation is that, in a dynamically changing environment, such updates are costly. For a number of examples we reduce the problem of energy efficiency to simple variations of the coupon collector problem (see for example [13]). This problem was examined in conjunction with network coding in [3], and it was shown network coding can offer benefits that increase as $\log n$ where n is the number of nodes.

We additionally look at sensor network applications, where sensor nodes would like to communicate information to one (or more) collectors, and either the sensor nodes or the collectors are mobile. In such scenarios again a benefit of $\log n$ is possible. One such application was also examined in [5], where each node is allowed a unit memory element and the problem was reduced to the coupon's collector problem. Our work differs in that, we allow multiple memory elements, and we allow network coding and routing of the same complexity in terms of transmissions.

This paper is organized as follows. In Section II we state the wireless broadcasting problem, review literature results, and position our paper in this framework. In Section III we present the analysis for several instances of the broadcasting problem. In Section IV we discuss sensor network applications. Finally, in Section V we present simulation results.

II. THE PROBLEM OF WIRELESS BROADCASTING

A. Problem Statement

We consider a wireless network with n nodes, where all nodes want to receive the same information, and where the network configuration dynamically changes. The information might initially be distributed among all nodes, or might be concentrated in a few nodes, as specified in Section III according to the application. We assume that the transmission range is the same for all transmitting nodes, and that each transmission conveys one unit of information to all nodes within its range.

Optimizing for energy efficiency amounts to minimizing the number of transmissions required to convey a unit of information to all receivers. More precisely, let T_{nc} denote the total number of transmissions required to broadcast one information unit to all nodes when we use network coding. Similarly, let T_w denote the required number of transmissions when we do

not use network coding. We are interested in calculating $\frac{T_{nc}}{T_w}$, in a dynamically changing environment, where the network topology constantly changes. We will consider two sources of variance:

- Nodes turn on and off independently at random with probability p . This for example captures nodes roaming out of range in a cellular environment, users in a peer-to-peer network turning off their devices to save energy, or more simply channel erasures due to noise and interference.
- Nodes move at a high speed with respect to the broadcast session duration. In Section III we will use a uniform at random mobility model to derive theoretical results, while in Section V we will examine the performance under a more realistic mobility model.

We will also assume that nodes have no information about the network topology. In particular, a transmitter does not know the identity of the nodes within its transmission range.

B. Related Work

There exist two main bodies of theoretical work related to this paper. In both cases, the emphasis is in minimizing the *speed* of information dissemination, which is expressed in terms of rounds of transmissions. We present the results as related to our specific problem. We then briefly review experimental results in the network coding realm.

1. Epidemic algorithms for rumor spreading

This work focuses on networks represented as graphs, and distributed algorithms, where, similarly to our case, nodes do not have information about the nodes they are communicating with. At each round, each node randomly chooses a communication partner among the nodes that are connected to it through an edge, and either “pushes” or “pulls” information from it (see for example [4], [10]). Results in the literature establish that $O(n \log n)$ rounds are required to disseminate the messages. Recently, work in [3] showed that using a network coding over a complete graph requires $O(n)$ rounds, which is a significant reduction.

2. Broadcasting in radio communication networks

In this body of work the wireless environment is modeled as a graph, where, when a node transmits a message, it is received by all its neighbors, and where a node successfully receives information if and only if exactly one of its neighbors is transmitting. Again transmissions are divided into rounds, where in each round a subset of the nodes transmits, in a way scheduled to minimize conflicts and maximize information spreading. The goal is to disseminate the information in the smallest number of rounds. Both centralized and decentralized algorithms are presented. Indicative results include that, the problem is NP-hard, there exist static networks where the number of required rounds is $\Theta(\log^2 n)$ while there exist mobile networks where the number of required rounds is $\Omega(n)$ [6], [2], [14]. Using a similar model, the problem of minimizing energy consumption over a static wireless network was recently studied in [9].

3. Algorithms

Within the scope of the network coding literature, a number of papers have proposed algorithms that employ network coding over a dynamically changing wireless environment and evaluated their performance through simulation results. Closest to our particular broadcasting problem is [17] which shows that, from the viewpoint of packet delivery ratio and overhead, network coding compares very favorably to flooding. Minimum cost multicasting using network coding was examined in [18] for mobile networks and in [12] for fixed networks. Our work differs in that, rather than solving the routing problem, we focus on assessing the benefits network coding may offer.

C. This paper

With respect to the previous work, our work is positioned as follows. We are interested in wireless networks, where a broadcast message is received by all neighbors within a certain radius. While most work in the broadcasting literature looks at the “speed” of dissemination, which is measured in terms of the required rounds, our measure of performance is energy efficiency, which translates in number of transmissions. Thus, we could, for example, have one transmission per round, and have no conflicts.

Work in [9] also considers optimization for energy efficiency, but over wireless networks modeled as arbitrary graphs. Although this approach has its merit and is interesting, it is not clear how well it applies in practical wireless networks, where the existence of “edges” connecting nodes reflects the positioning of the nodes on the plane and is not arbitrary.

Moreover, our interest is not in worst case bounds, as in [14], but average performance. In this sense, our work is closer to rumor-spreading using network coding [3]. We will make this connection precise in Section III. Finally, we are not interested in “static” networks, but in dynamically changing topologies, where nodes do not have information about the network topology.

III. WIRELESS BROADCASTING APPLICATIONS

A. Ad-hoc wireless network

We consider n nodes, where each node is a source that wants to transmit information to all other nodes. We divide time into rounds. The transmission policy is that during each round each active node transmits once and transmissions occur successively, without conflict. We can employ this simplified policy without loss of optimality, because we do not care about delay, but only about energy efficiency.

We assume that at the beginning of each round nodes are placed uniformly at random on a unit-area disc of radius $1/\sqrt{\pi}$. This corresponds to having a uniform at random mobility pattern, where the rounds are far enough in time to allow a node to move anywhere on the disc with equal probability between rounds. We use this generous mobility model to simplify the analysis, but examine more realistic mobility models through simulation results in Section V.

We also assume that each node can successfully broadcast information within a radius of r with

$$r = \Theta\left(\frac{1}{\sqrt{n}}\right)$$

fixed for all nodes. Thus at each round each node will have on the average a constant number of $k = n\pi r^2$ neighbors, of which $k(1-p)$ will be active.

We are going to compare the energy efficiency in the case where we use forwarding and where we use network coding. We underline our assumption that nodes do not know which are their neighbors, or what information they already have. Thus, in the case of forwarding, without loss of generality, we can assume that during each round, and at each (possibly new) position, node i always broadcasts x_i . In the case of network coding, each node will transmit a random linear combination over some finite field F_q of the symbols it has previously received.

Theorem 1: Broadcasting to all receivers can be achieved using on the average

–without network coding: $\frac{\Theta(n \log n)}{(1-p)^2}$ rounds,

–with network coding: $\frac{\Theta(n)}{(1-p)^2}$ rounds,

where at each round occur on the average $(1-p)n$ transmissions. Thus on the average

$$\frac{T_{nc}}{T_w} = \Theta\left(\frac{1}{\log n}\right).$$

Proof Outline

Consider first the case of forwarding, and a given node j that would like to transmit its message x_j to all other $n-1$ nodes.

Construct a bipartite graph as follows. The left part consists of the $n-1$ nodes. The right part consists of M nodes v_i , where node v_i corresponds to round i , and is connected to the neighbors of node j during this round. Thus the degree of node v_i is a random variable with average $k(1-p)$. We are asking, how many right hand side nodes do we need, i.e., what number of rounds, so that node j transmits its message in all other nodes. This simple analysis has been performed in the context of LT and Raptor codes (see for example [15]-Proposition 1) where it was shown that M should scale as $\Theta(n \log n)$. Since node j will be active with probability $(1-p)$, it is easy to see that the average number of rounds we will need will equal

$$\frac{\Theta(n \log n)}{(1-p)^2}.$$

This problem can also be thought as a variation of the coupon collector's problem. The coupon collector's problem in its standard form is described as buying boxes of some product, and in each box there exists one coupon, chosen uniformly at random from a collection of n coupons. We are asking what is the average number of boxes we need to buy to collect all n coupons (see for example [13]). It is well known that in this case the answer is $O(n \log n)$ coupons. Our case is a simple variation, where now each box contains on the average $k(1-p)$ different coupons.

In [3] it was shown that use of network coding with the standard coupon collector problem reduces the number of

required rounds to n . In our case as well, node j will be active on the average $(1-p)m$ out of m rounds. Given that it is active, it will receive on the average $k(1-p)$ transmissions from its active neighbors. Using standard arguments in the network coding literature, and provided that the field F_q is large enough, each received transmission will bring new information to the node j . Thus, it will be able to decode all n information units on the average after $\frac{\Theta(n)}{(1-p)^2}$ rounds. \square

Note that the performance of network coding is not affected by node mobility. In contrast, mobility has a significant effect on forwarding. Initially, as nodes randomly move, the information is disseminated faster than in the case of a static network. However, because of the assumption that *nodes do not know what information their neighbors have*, as approximately half the nodes collect the information, more and more often transmissions do not bring new information to the receiving nodes. This is a well known observation in the context of rumor spreading algorithms over networks represented as graphs.

We underline again that our results do not hold, if we assume that nodes have some information about other nodes in their transmission range. For example, if we assume that a node knows *how many* active nodes are in its transmission range, it can wait (a possibly infinite time) until *all* nodes are simultaneously in its transmission range and are active, and then broadcast its message using just one transmission.

B. Broadcasting in Cellular Networks

In the cellular network model, we have m base-stations and n mobile phone receivers. The base-stations have K information units, that they want to transmit to all mobiles. We assume that the transmission range is the same for all base-stations, each transmission conveys one unit of information, and that the coverage areas of the base-stations do not overlap.

In this model, base-stations are always active, while nodes are mobile and may turn on and off. A node will be active and successfully receive information approximately $M(1-p)$ out of every M rounds. Thus, if base-stations broadcast using an erasure correcting code of rate $(1-p)$, then each transmission brings useful information at each node. For a node to receive K messages we will need $\frac{K}{(1-p)}$ rounds.

In the case of forwarding, assume that base-stations randomly select and transmit one of the K messages. Thus each node at each round observes one of the messages uniformly at random. We can think of this problem as a balls-in-bins experiment, where the bins are the K messages the node wants to collect, and the balls correspond to the rounds. Using standard results [13] we will again need on the average $\frac{n \log n}{(1-p)}$ rounds. Thus, we again realize a $\log n$ benefit.

IV. DATA COLLECTION IN SENSOR NETWORKS

We consider a sensor network with n nodes, where each node i has an independent observation x_i . There also exist a set of k collector nodes. We want the union of information that these k nodes collect to be sufficient to retrieve all x_i . We consider two models. In the first model, the sensor nodes themselves are mobile, while the collector nodes are static.

We call this the *mobile nodes model*. In the second model, we have k collectors that move randomly among the nodes, and collect the information. We call this the *moving collector model*.

A. Mobile Node Model

This model corresponds to applications where sensor nodes are placed on mobile objects such as cars or wildlife, that measure statistics to be communicated to base-stations. Assume that sensor nodes transmit at a constant range, to other sensor nodes as well as to the base-stations.

In the case of forwarding, we have one more variation of the coupon collector problem, where now we have k collectors, and we are asking how many boxes should the collectors buy so that the union of their coupons covers all n possibilities. For k constant with respect to n , which is the most realistic case, it is easy to see that the same order arguments apply. That is, we need $\Theta(n \log n)$ transmissions, while use of network coding results to $\Theta(n)$ transmissions. This application is very similar to the work described in [5], the only essential difference being that we allow k collectors.

B. Mobile Collector Model

In this model nodes are static, while collectors are mobile. In particular, we will assume for simplicity that nodes are placed on a square grid, and that we have a total of n^2 nodes. We consider two phases.

In the first phase, the sensor nodes exchange information according for example to the decentralized protocol described in [8] during $m < n^2$ rounds, where at each round all nodes of the square grid transmit once. We assume that each node transmission is successfully received by its four closest neighbors. As a result, in the case of network coding, after m rounds each node will have $4m + 1$ observations which will depend on the information of the $\Theta(m^2)$ nodes that are within a radius of m . Note that information collected by neighbor nodes may have a significant intersection.

In the case of forwarding, we will assume that at each round each node forwards with equal probability one of the four messages it has received in the previous round. Then after m rounds each node will have collected $4m + 1$ data x_i , from nodes within a radius of m . In particular, given our transmission model, each bit x_i will perform a random walk with m steps, and thus on the average we expect it to have reached nodes within distance $\Theta(\sqrt{m})$. We will make here the simplifying assumption that each node will receive all $\Theta(m)$ information symbols x_i within a radius of $\Theta(\sqrt{m})$.

In the second phase the a mobile collector samples k nodes uniformly at random from the square grid. We are asking what is the minimum number of nodes to sample to collect all information. Obviously $k \geq n^2 / (4m + 1)$.

In the case of network coding, we can think of our problem as randomly covering the square grid with disks of radius m . Consider a particular node i . The probability that the node is not covered by a disc during a given round equals the probability that the center of the disk is not within distance m

from node i . That is, if we uniformly at random choose the center to be any of the n^2 points of the grid, then the center is not one of the $\Theta(m^2)$ points that are in distance m from node i . So this probability equals

$$1 - \frac{m^2}{n^2}.$$

Repeating the experiment for k rounds, the probability that a node is not covered equals

$$\left(1 - \frac{m^2}{n^2}\right)^k.$$

Assume that we choose $m = \sqrt{n}$ and that we require that the probability a node is not covered decays at least as $1/n$. Then we will need at least $k = n \log n$ rounds. In fact, since

$$k \geq n^2 / (4m + 1) \geq n \log n$$

we will need $k = \Theta(n\sqrt{n})$ rounds, i.e. the optimal number of rounds. (Each round corresponds to sampling a node as the collector moves).

In the case of forwarding, we will now have that the probability a node is not cover (from disks of radius $\sqrt{m} = n^{1/4}$) will equal

$$\left(1 - \frac{1}{n\sqrt{n}}\right)^k$$

and thus we will need

$$\Theta(n\sqrt{n} \log n)$$

rounds. As a conclusion, our approximate analysis again indicates a loss of a factor of $\log n$.

V. SIMULATION RESULTS

A. Simulations for broadcast in ad-hoc networks

As a simulation environment, we use a random topology with a fixed node density. To avoid edge effects, we let the simulation area wrap around at the edges. Nodes have a nominal radio range of 250m and each node has on average four neighbors.

We will first present simulation results of the case where at each round each node is placed uniformly at random in the rectangular simulation area. The left graph of Fig. 1 shows a log-log plot of the total number of packet transmissions required such that all nodes receive all packets for different network sizes. In general, network coding achieves a 100% delivery ratio using much fewer transmissions. We further observe that this gap widens as the network size increases. The semi-log plot in the right graph of Figure 1 confirms the $\log n$ factor in the ratio of overhead of flooding and network coding from the theoretical analysis. A similar overhead ratio can be observed in simulations with a more dense or more sparse network (not shown here).

The mobility model used in Fig. 1 corresponds to the generic mobility model used in the theoretical analysis in Section III-A. This model essentially implies that the composition of the neighborhood of a node is completely uncorrelated from round to round.

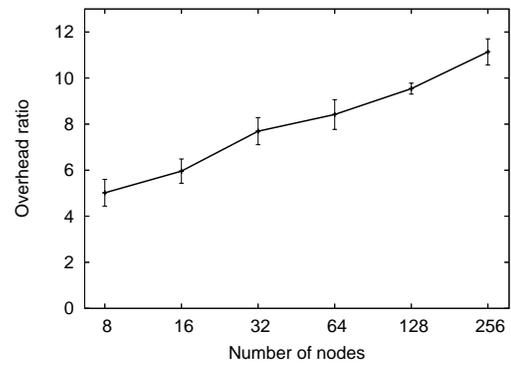
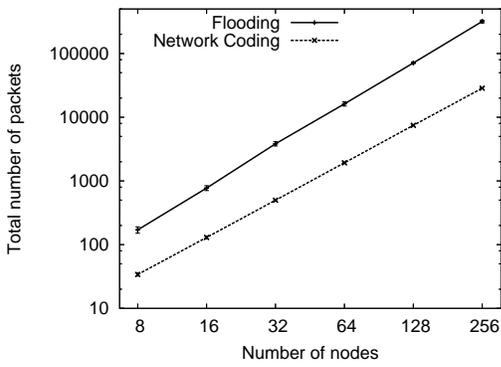


Fig. 1. Overhead in terms of required number of packet transmissions for each node to receive all packets (left) and ratio of number flooding overhead to network coding overhead (right) for uniformly random mobility.

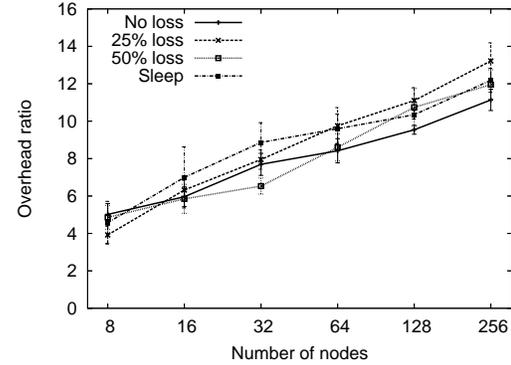
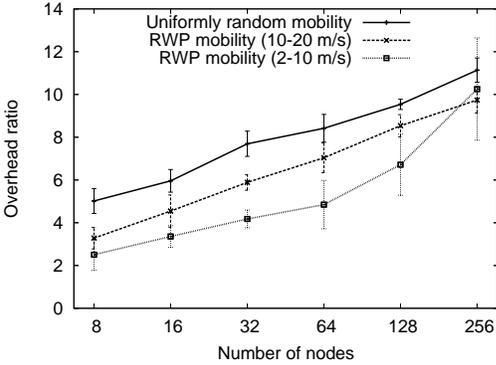


Fig. 2. Ratio of number flooding overhead to network coding overhead for random waypoint mobility.

Fig. 3. Ratio of number flooding overhead to network coding overhead for sleeping nodes.

In practice, this is true only when the node speed is very high or the packet transmission rate is very low. A less generous mobility implies that less data is transported through the network by node mobility and has instead to be forwarded via intermediate nodes. In Fig. 2 we present simulation results for a more realistic mobility model. We use the same simulation parameters as before, but also show the overhead ratio for mobility according to the random-waypoint mobility model with no pause time and movement speeds uniformly distributed between 2 m/s and 10 m/s as well as 10 m/s and 20 m/s, respectively. With the random-waypoint mobility model, nodes pick a random destination whose location is uniformly distributed in the simulation area as well as a movement speed with which they travel until the destination is reached.

We can see that in this case, although network coding still offers significant benefits, the performance gap with routing is smaller. This agrees with our intuition that, when mobility is more restricted, network coding performance will decrease, because how well the data is “mixed” plays a crucial role for the network coding analysis.

To model sleep in the same way as in the theoretical analysis, we drop packets upon packet transmission *and* packet reception with a certain probability (to emulate sleep at the sender and the receivers). We also compare this simple model to a more realistic implementation of node sleep. Here, all

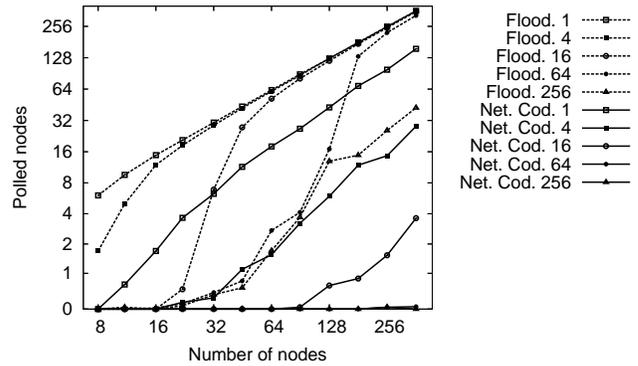


Fig. 4. Number of polled nodes for different numbers of transmissions per node for flooding and network coding.

nodes are within radio range and the sleep ratio (wake time over total time) is set such that we have on average the same number of active neighbors as in the previous settings. Sleep duration follows an exponential distribution. From Fig. 3 we see that there is no clear difference between the different sleep models and the overhead ratio remains basically unaffected. Of course the overhead itself increases with the sleep ratio.

B. Simulations for data collection

We simulate data dissemination and collection for sensor networks using a forwarding algorithm that sends out a packet with a certain probability, whenever a packet is received for

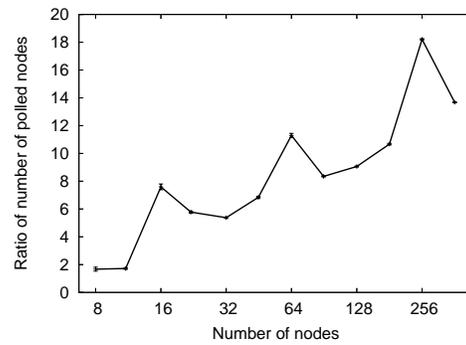
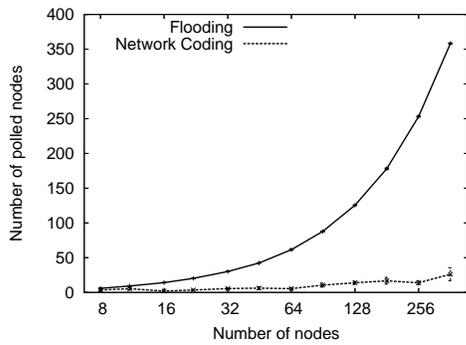


Fig. 5. Number of nodes that have to be polled by the data collector to obtain all packets (left graph) and the ratio of polled nodes for flooding to polled nodes for network coding (right graph).

the first time (flooding) or an innovative packet is received (network coding). When no data is disseminated and each node holds only its own packets, network coding and flooding obviously have the same performance. In case data is disseminated such that all nodes can decode all packets and a data collection phase is unnecessary, network coding offers a constant benefit in terms of required number of transmissions [8]. Between these two extreme cases, network coding provides higher performance gains in terms of number of polled nodes.

We use a static scenario without node mobility. We further vary the number of nodes in the network, as well as the network size, to have on average 12 neighbors per node. The total number of transmissions per node is limited and we count, how many nodes have to be polled by the data collector (sink) after the data dissemination phase, to obtain all data. Even if nodes are allowed to transmit only once, network coding benefits a lot in very small scenarios but has little impact in larger networks, as shown in Fig. 4. As long as the allowed number of transmissions is significantly below the number of nodes in the network, flooding needs to poll almost all nodes. Here, flooding with 256 transmissions per node achieves roughly the same performance as network coding with 4 transmissions per node. For more than 16 transmissions per node, network coding delivers all packets to the sink, such that no polling is necessary.

For the next set of simulations, we set the number of packets each node is allowed to transmit to $\lfloor 0.25\sqrt{n} \rfloor$. In this setting, flooding takes almost no advantage of the proactive data dissemination (left graph of Fig. 5). The number of polled nodes is reduced by slightly more than the number of packets that each node transmits. In contrast, with network coding the sink node needs to poll only a few other nodes to be able to decode all original packets. Looking at the ratio of polled nodes for flooding to polled nodes for network coding (right graph of Fig. 5), we again see an approximately logarithmic increase. The spikes that are evident for network sizes of 16, 64, and 256 nodes stem from the rounding of the allowed number of transmissions.

VI. CONCLUSIONS

In this paper we examined several wireless applications that include broadcasting and data collection from sensor nodes.

We saw that with simple decentralized algorithms, network coding may offer logarithmic benefits in terms of energy efficiency over highly variable environments.

REFERENCES

- [1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung. Network information flow. *IEEE Trans. Inform. Theory*, pages 1204–1216, July 2000.
- [2] N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. On the complexity of radion communication. *SODA*, 1989.
- [3] S. Deb and M. Medard. Algebraic gossip: A network coding approach to optimal multiple rumor mongering. In *Proc. Allerton*, Oct. 2004.
- [4] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. *ACM SIGOPS Operating Systems Review*, January 1988.
- [5] A. G. Dimakis, V. Prabhakaran, and K. Ramchandran. Ubiquitous access to distributed data in large-scale sensor networks through decentralized erasure codes. In *Symposium on Information Processing in Sensor Networks (IPSN '05)*, Apr. 2005.
- [6] M. Elkin and G. Kortsarz. Logarithmic inapproximability of the radio broadcast problem. *Journal of Algorithms*, pages 8–25, July 2004.
- [7] C. Fragouli, J. Widmer, and J.-Y. L. Boudec. Energy-efficient broadcasting in wireless ad-hoc networks using network coding. In *In preparation for ACM Trans. in Networking*, 2006.
- [8] C. Fragouli, J. Widmer, and J. Y. LeBoudec. A network coding approach to energy efficient broadcasting: from theory to practice. *Infocom 2006*, March 2006.
- [9] K. Jain and K. Talwar. On the power saving of network coding. *Allerton*, Oct 2005.
- [10] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking. Randomized rumor spreading. *SODA*, 2000.
- [11] S.-Y. R. Li, R. W. Yeung, and N. Cai. Linear network coding. *IEEE Trans. Inform. Theory*, 49:371–381, Feb. 2003.
- [12] D. Lun, M. Medard, T. Ho, and R. Koetter. Network coding with a cost criterion. In *ISITA*, Oct. 2004.
- [13] M. Mitzenmacher and E. Upfal. Probability and computing : Randomized algorithms and probabilistic analysis. *Cambridge Press*, 2005.
- [14] R. Prakash, A. Schiper, M. Mohsin, D. Cavin, and Y. Sasson. A lower bound for broadcasting in mobile ad hoc networks. *Technical Report (IC/2004/37)*, June 2004.
- [15] A. Shokrollahi. Raptor codes. *Submitted to IEEE Trans. Information Theory*, 2004.
- [16] J. Widmer, C. Fragouli, and J.-Y. Le Boudec. Energy efficient broadcasting in wireless ad hoc networks. *First Workshop on Network Coding*, March 2005.
- [17] J. Widmer and J.-Y. Le Boudec. Network coding for efficient communication in extreme networks. In *Workshop on delay tolerant networking and related networks (WDTN-05)*, Philadelphia, PA, Aug. 2005.
- [18] Y. Wu, P. A. Chou, and S.-Y. Kung. Minimum-energy multicast in mobile ad hoc networks using network coding. In *Proc. IEEE Information Theory Workshop*, Oct. 2004.