On the Network Coding Advantage for Wireless Multicast in Euclidean Space

Ashish Goel*
Stanford University
Stanford, CA 94305
ashishg@stanford.edu

Sanjeev Khanna†
University of Pennsylvania
Philadelphia, PA 19104
sanjeev@cis.upenn.edu

Abstract

Multicast is a fundamental communication operation in wireless sensor networks whereby a source sensor transmits its information to a relevant subset of sensors in the network. Motivated by this, we study the advantage of network coding for minimizing the total power needed for multicast in wireless networks. We show that there is an absolute constant, depending only on the power gradient and the dimension of the underlying Euclidean space, that bounds the maximum advantage of network coding. An interesting aspect of our result is that it shows that the advantage of coding remains bounded by a constant even when compared to a multicast scheme without coding that is restricted to do only point-to-point transmissions.

1 Introduction

A basic operation in communication networks is multicasting whereby a node, called the *source*, sends some data to a subset of nodes in the network, referred to as the *sinks*. The maximum rate at which a source can send data to the sinks (throughput) is clearly bounded by the minimum capacity cut that separates the source from one or more of the sinks. When there is only one sink, the celebrated Max Flow-Min Cut theorem [7] shows that the maximum throughput achievable matches the minimum capacity cut. However, once we allow even two sinks, it is easy to construct examples where the maximum throughput is strictly smaller than minimum capacity cut that separates the source from one of the sinks. Ahlswede *et al.* [3] recently introduced the notion of *network coding* where the network nodes are allowed additional operations besides merely storing and forwarding the data received by them. In this model, intermediate nodes have the ability to decode the arriving data and forward some encoding of the data received by them. Ahlswede *et al.* [3] showed that in this more general model, there always exists a scheme that allows the source to send the data at a rate that matches the minimum capacity cut separating the source from at least one sink. The worst-case gap between the throughput achievable with network coding and without network coding is referred to as the *coding advantage for throughput maximization*. We will denote it by Θ_t in this paper. For the multicast problem, Li *et al.* [9] showed that there exists undirected instances with $\Theta_t \geq 9/8$. Sanders *et al.* [13] strengthened this to $\Theta_t = \Omega(\log n)$ for multicast in directed networks. Agarwal and Charikar [1], improved these bounds to $8/7$ for undirected multicast and to $\Omega((\frac{\log n}{\log \log n})^2)$ for the directed case. More importantly, they establish an elegant connection, namely, Θ_t is equal to the worst-case integrality gap of a natural linear programming formulation for the Steiner tree problem.

Network coding in wireless networks has also received considerable attention in recent years, and has in large part been motivated by application to sensor networks which typically use wireless transmission and are often highly resource constrained (see [6] for an example). In [14, 8], the problem of broadcasting in a wireless network with a minimum number of transmissions has been studied. It is assumed that each node is a source for the broadcast and that every node broadcasts at a fixed range (hence all transmission costs are identical). The authors show that, asymptotically speaking, network coding reduces the number of transmissions needed by a factor of 2 in case of circular networks (nodes are placed around a circle with equal spacing), and by a factor of $4/3$ in case of grid networks. Moreover, these bounds are shown to be tight and distributed algorithms are given for realizing the advantage of network coding. The problem of finding a minimum energy multicast tree with network coding is known to be polynomial-time solvable [11, 15]. This is in sharp contrast to the setting where only forwarding of packets is allowed at each node. In the latter setting, even the problem of computing a minimum energy broadcast tree is known to be NP-

*Research supported by an NSF CAREER award 0339262, and an Alfred P. Sloan Faculty Fellowship.
†Research supported by a Guggenheim Fellowship and NSF Award CCF-0635084.

978-0-7695-3157-1/08 $25.00 © 2008 IEEE
DOI 10.1109/IPSN.2008.29

hard [5, 4, 10]. However, the following question was left open: *What is the maximum network coding advantage for multicast in wireless networks, where the cost of a transmission is some function of the distance?*

1.1 Our Results

In this paper, we provide bounds on the advantage (in terms of power consumed) of using network coding for multicasting in wireless networks. We assume that we are given a set of n points in d-dimensional space where the distance between any pair of points is the Euclidean distance between them. We are also given a parameter α, the *power gradient*; a node can broadcast a unit size message to all nodes within Euclidean distance r from it, by transmitting the message with power $\pi(r) = \pi_0 r^\alpha$ for some constant π_0 and $\alpha \geq 2$. This is a standard and reasonable assumption on the form of the power function, see [12], for instance.

Given an instance of the multicast problem in d-dimensional space with power gradient α, we analyze the quantity Θ_p, defined as the worst-case ratio of the minimum power needed to multicast a unit of data from the source to the sinks without network coding to the minimum power needed when network coding is used. We refer to Θ_p as the *coding advantage for power minimization*. Our main result is as follows.

Theorem 1 *For any $\alpha > d \geq 2$, there exists a constant $c_1(\alpha, d) \leq \frac{2\alpha \cdot 3^d}{\alpha - d}$ depending only on α and d such that $\Theta_p \leq c_1(\alpha, d)$ for wireless multicast.*

In other words, there is a constant which depends only on the power gradient and the dimension of the space, that bounds the coding advantage for minimizing the power needed for multicast. Moreover, we show that this constant upper bound on coding advantage holds even when compared to data forwarding schemes that are allowed only point-to-point transmissions.

1.2 Overview of the Proof Technique

We will distinguish between two kinds of wireless transmission in proving our main result. The first is a *simultaneous transmission*: when a source transmits a message, all receivers that are within its range (the range depends on the transmission power) receive the message. This is the kind of transmission we are primarily interested in. The other kind of transmission is called a *point to point transmission*, where only one receiver can receive the transmitted message; this kind of transmission is an artifact in our proof.

The proof of our main result is presented in section 4 and is based on the following technical steps:

1. We first use a cut-based linear program (LP) to lower bound the optimal power required for wireless multicast in the presence of network coding.
2. Then we show that a feasible solution to this cut-based LP formulation can be transformed into a feasible solution for an LP that computes the optimum fractional multicast tree which does not use network coding and *which only performs point to point transmissions*.
3. We next use the geometry of Euclidean spaces to show that the cost of the fractional tree generated by the above transformation is not much larger than the optimum solution to the cut LP (and hence the optimal power needed with network coding). This requires a technical lemma that we discuss separately in Section 3 which relates the cost of point to point vs simultaneous transmission for complete (i.e. integral) trees.
4. Finally, we use the fact that the optimum fractional multicast tree (assuming only point to point transmissions) is a good approximation to the optimum integer multicast tree (again assuming only point to point transmissions). Since multicast using simultaneous transmissions can not consume any more power than multicast using point to point transmissions, it follows that the network coding advantage in reducing the power needed for multicast is bounded by a constant factor.

We note that our result also holds for the case $\alpha = d$ albeit with slightly different constants, as explained at the end of section 3.

2 Preliminaries

The input to the wireless multicast problem is a set V of n wireless nodes (sensors) situated in a d-dimensional space, a set of *sinks* $X \subseteq V$, and a distinguished vertex s, called the *source*. Let α be the power gradient, that is, to transmit a message of unit size to all recipients within Euclidean distance r, a node needs to transmit the message with power $\pi(r) = \pi_0 r^\alpha$ for some constant π_0. We will assume, without loss of generality, that $\pi_0 = 1$. The goal is to send a unit of information from the source s to the set X of sink nodes, minimizing the total power consumed.

Let $\mathsf{dist}(u, v)$ denote the Euclidean distance between points $u, v \in V$. Further, let $R = \{\mathsf{dist}(u, v) : u, v \in V\}$ denote the set of possible inter-point distances, and hence the set of possible transmission radii of interest. For $r \in R, u \in V$, let $N_r(u)$ denote the set $\{v \in V : \mathsf{dist}(u, v) \leq r\}$. This is the set of all nodes in a Euclidean ball of radius r centered at u; in particular this always includes u. Let $\tau_r(u)$ denote a minimum spanning tree (rooted at u) of the set $N_r(u)$.

3 Point-to-Point vs. Simultaneous Transmission

Consider any rooted Steiner tree T connecting the source s to vertices in X (i.e. T is a tree that spans s and X but may span other vertices in the graph as well). We say that an edge (v, w) belongs to T if w is a child of v in T. We define two different transmission costs associated with a given tree T. The first cost, referred to as the *simultaneous transmission cost* and denoted by $C(T)$, measures the total amount of power consumed if each node transmits a single copy of the message with just enough power for the message to be received by each of its children in the tree. The second cost, referred to as the *point-to-point transmission cost* and denoted by $D(T)$, is the total amount of power consumed if each node sends a separate message to each of its children, with each copy being transmitted with just enough power to reach that particular child. Let C_{ST} denote the minimum value of $C(T)$ over all Steiner trees connecting s to X. Similarly, let D_{ST} denote the minimum value of $D(T)$ over all Steiner trees connecting s to X. Clearly, $C_{\mathsf{ST}} \leq D_{\mathsf{ST}}$ since for any Steiner tree T, we have $C(T) \leq D(T)$.

The next lemma bounds the worst-case gap between the power needed to transmit a unit of information to all input points within some radius r of a given point using a single simultaneous transmission versus a point-to-point transmission scheme. A similar lemma was established by Clementi *et al.* [5]. We present a proof here for completeness. Our proof is more elementary and may be of independent interest.

Given any node u in the network and a radius r, let $\tau_r(u)$ denote a spanning tree of nodes in $N_r(u)$ that minimizes $D(\tau_r(u))$.

Lemma 2 *If $\alpha > d \geq 2$, then for any point $u \in V$ and a radius r_0, there exists a constant $c(\alpha, d) \leq \frac{\alpha 3^d}{\alpha - d}$ such that $D(\tau_{r_0}(u)) \leq c(\alpha, d)\pi(r_0)$.*

Proof: Let $b(x)$ denote the maximum possible number of disjoint balls of radius $0 < x \leq r_0$ with their centers inside a ball of radius r_0 centered at u. Since all the radius x balls must be contained inside a ball of radius $r_0 + x$ centered at u, we have $b(x) < ((r_0 + x)/x)^d$. Let S denote the minimum spanning tree, under the Euclidean metric, of the set of vertices $N_{r_0}(u)$. Observe that no edge in this tree can be longer than r_0. Let $M(r)$ denote the number of edges in this tree which are of length greater than r, where $0 < r < r_0$. Deleting these edges creates $M(r)+1$ components in S. Pick an arbitrary node from each of these components and draw a ball of radius $r/2$ around it. The resulting $M(r) + 1$ balls must be disjoint; if not, we would have a spanning tree that is cheaper than the minimum spanning tree S, a contradiction. Thus, $M(r) + 1 \leq b(r/2)$, or

$$M(r) < \left(\frac{(r_0 + r/2)}{(r/2)}\right)^d.$$

Since $r < r_0$, we obtain

$$M(r) < 3^d \left(\frac{r_0}{r}\right)^d.$$

Recall that $\pi(r) = r^\alpha$ is the power required to transmit a message over a distance r. The point to point transmission cost $D(S)$of the spanning tree S is at most

$$\int_0^{r_0} 3^d \left(\frac{r_0}{r}\right)^d \mathbf{d}\pi(\mathbf{r}) = \left(\frac{\alpha}{\alpha - d}\right) 3^d r_0^\alpha.$$

Recall that S was chosen to be a minimum spanning tree of $N_{r_0}(u)$ under the Euclidean metric, not under the point to point transmission cost. Since $\tau_{r_0}(u)$ is the tree which minimizes the point to point transmission cost over all the nodes in $N_{r_0}(u)$, we can conclude that $D(\tau_{r_0}(u)) \leq D(S)$ which concludes the proof of this lemma. ■

The results of [5] also hold for $\alpha = d$ (with a worse exponential dependence on α, but no dependence on $(\alpha - d)^{-1}$) and hence the results in this paper can also be extended to the case $\alpha = d$ with different constants.

4 Wireless Multicast: Coding Advantage for Power Minimization

We will now prove our main result, using the high level structure outlined in Section 1.2.

Let C_{CODE} represent the minimum total power required to multicast a message of unit size from s to all the sinks in X, where the multicast can use an arbitrary topology, and intermediate nodes are allowed to divide the message into arbitrarily small pieces and use network coding. The minimum total power required to multicast a message of unit size from s to nodes in X, where the multicast occurs over a Steiner tree rooted at s, and no coding is allowed is precisely the quantity C_{ST} defined earlier. Our goal is to prove that $C_{\mathsf{ST}} \leq \Theta_p C_{\mathsf{CODE}}$ where Θ_p depends only on the dimension d and the gradient α, and not on the size of V or the location of points in V.

4.1 A Lower Bound on C_{CODE}

Consider a set $S \subset V$ such that $s \in S$. In order to multicast a unit-sized message to all nodes in X, the sum of the sizes of all the messages transmitted from S to $V - S$ must be at least 1 unless $X \subseteq S$. We call this the "cut constraint". Let $y_r(u)$ denote the amount of information transmitted by u over a distance r. Also, let $\mathsf{dist}(u, V - S)$

denote the minimum distance from u to a node in $V - S$. Then, the cut constraint for S corresponds to the condition

$$\sum_{u \in S} \left(\sum_{r \in R, r \geq \mathsf{dist}(u, V-S)} y_r(u) \right) \geq 1.$$

The quantity C_{CUT} is the minimum power required for a set of transmissions to satisfy all the cut constraints. More formally, C_{CUT} is the value of the following linear program:

$$C_{\mathsf{CUT}} = \text{minimize} \sum_u \sum_r \pi(r) y_r(u) \tag{1}$$

subject to
$\forall S \subset V, s \in S, X \not\subseteq S$:

$$\sum_{u \in S} \left(\sum_{r \in R, r \geq \mathsf{dist}(u, V-S)} y_r(u) \right) \geq 1 \tag{2}$$

$$\forall u \in V, r \in R: \;\; y_r(u) \geq 0. \tag{3}$$

Since any network coding solution must satisfy all the cut constraints, we obtain

$$C_{\mathsf{CODE}} \geq C_{\mathsf{CUT}}. \tag{4}$$

4.2 An Upper Bound on C_{ST}

Let D_{FRAC} denote the solution to the following linear program:

$$D_{\mathsf{FRAC}} = \text{minimize} \sum_{u,v} \pi(\mathsf{dist}(u,v)) z(u,v) \tag{5}$$

subject to

$$\forall S \subset V, s \in S, X \not\subseteq S: \sum_{u \in S, v \notin S} z(u,v) \geq 1 \tag{6}$$

$$\forall u, v \in V: \;\; z(u,v) \geq 0. \tag{7}$$

If we require $z(u,v)$ to be integral, we would obtain an integer program for computing D_{ST}. Thus D_{FRAC} can be viewed as a *fractional relaxation* for computing D_{ST}, and hence gives a lower bound on its cost. It is also known [2] that

$$D_{\mathsf{ST}} \leq 2 D_{\mathsf{FRAC}}. \tag{8}$$

Further, since $D(T) \geq C(T)$ for all Steiner trees T rooted at s, it follows that

$$C_{\mathsf{ST}} \leq D_{\mathsf{ST}} \leq 2 D_{\mathsf{FRAC}}. \tag{9}$$

Even for the special case of $X = V$, the problem of finding the optimal tree (i.e. the one that requires the least total power) for wireless multicast is itself NP-hard [5, 4, 10]. So we will settle for an upper bound on the quantity C_{ST}.

4.3 Bounding the Power Gain

We will now show that an optimal Steiner tree rooted at s consumes at most a constant factor more power than the optimum network coding solution, for any fixed $\alpha > d \geq 2$. Thus our proof is constructive. In particular, we will show that $D_{\mathsf{FRAC}} \leq c(\alpha, d) \cdot C_{\mathsf{CUT}}$, completing the proof of Theorem 1. Surprisingly, we will show this result where the network coding solution is allowed to use simultaneous transmission, whereas the Steiner tree is restricted to the more expensive point-to-point transmissions.

Consider the optimum solution y to the linear program for computing C_{CUT} (LP 1-3). Define variables $z(v,w)$ to denote $\sum_{u,r:(v,w) \in \tau_r(u)} y_r(u)$. We establish two useful properties of these variables in the next two lemmas.

Lemma 3 *The variables $z(v,w)$ satisfy the constraints in the linear program for computing D_{FRAC} (LP 5-7).*

Proof: The variables $z(v,w)$ are non-negative combinations of the variables $y_r(u)$. Since the variables $y_r(u)$ are feasible for the LP 1-3, the variables $z(v,w)$ must be non-negative and hence satisfy constraints 7.

Now consider any set $S \subset V$ such that $s \in S$ and $X \not\subseteq S$. Since y is a feasible solution to the LP 1-3, we must have $\sum_{u \in S} \sum_{r \in R, r \geq \mathsf{dist}(u, V-S)} y_r(u) \geq 1$. If $r \geq \mathsf{dist}(u, V - S)$ then $N_r(u) \cap (V-S) \neq \emptyset$. Also, if $u \in S$ then $N_r(u) \cap S \neq \emptyset$. Hence, if $u \in S$ and $r \geq \mathsf{dist}(u, V-S)$ then $\tau_r(u) \cap (S \times (V-S)) \neq \emptyset$. Now,

$$\begin{aligned}
& \sum_{v \in S, w \notin S} z(v,w) \\
= & \sum_{v \in S, w \notin S} \left(\sum_{u,r:(v,w) \in \tau_r(u)} y_r(u) \right) \\
\geq & \sum_{v \in S, w \notin S} \left(\sum_{u,r:u \in S, r \geq \mathsf{dist}(u,V-S), (v,w) \in \tau_r(u)} y_r(u) \right) \\
= & \sum_{u,r:u \in S, r \geq \mathsf{dist}(u,V-S)} y_r(u) \cdot |\tau_r(u) \cap (S \times (V-S))| \\
\geq & \sum_{u,r:u \in S, r \geq \mathsf{dist}(u,V-S)} y_r(u) \\
& \qquad [\text{Since } \tau_r(u) \cap (S \times (V-S)) \neq \emptyset] \\
\geq & \; 1 \\
& \qquad [\text{Since } y_r(u) \text{ satisfy constraints 2}].
\end{aligned}$$

Hence, the variables z also satisfy constraints 6 in LP 5-7. ∎

Lemma 4 $\sum_{v,w \in V} \pi(\mathsf{dist}(v,w)) z(v,w) \leq c(\alpha, d) C_{\mathsf{CUT}}$.

Proof: By definition of $D(T)$, Lemma 2 implies

$$\sum_{(v,w)\in\tau_r(u)} \pi(\mathsf{dist}(v,w)) \le c(\alpha,d)\pi(r).$$

We multiply both sides by $y_r(u)$ and then sum over all r and u to obtain

$$\sum_{r\in R,u\in V} y_r(u) \sum_{(v,w)\in\tau_r(u)} \pi(\mathsf{dist}(v,w))$$

$$\le c(\alpha,d) \sum_{r\in R,u\in V} \pi(r)y_r(u).$$

We exchange the order of summation on the LHS to obtain

$$\sum_{v,w\in V} \pi(\mathsf{dist}(v,w)) \sum_{r,u:(v,w)\in\tau_r(u)} y_r(u)$$

$$\le c(\alpha,d) \sum_{r\in R,u\in V} \pi(r)y_r(u).$$

The RHS is now just $c(\alpha,d)C_{\mathsf{CUT}}$. The inner sum on the LHS is just $z(v,w)$. Hence, we obtain

$$\sum_{v,w\in V} \pi(\mathsf{dist}(v,w))z(v,w) \le c(\alpha,d)C_{\mathsf{CUT}}.$$

■

Observe that the LHS is Lemma 4 is exactly the objective function of the LP 5-7 for computing D_{FRAC}. Our main technical lemma is now immediate:

Lemma 5 $D_{\mathsf{FRAC}} \le c(\alpha,d)C_{\mathsf{CUT}}$.

Proof: Since z is feasible for the LP 5-7 (Lemma 3), we have $D_{\mathsf{FRAC}} \le \sum_{v,w\in V} \pi(\mathsf{dist}(v,w))z(v,w)$. Combining this with Lemma 4 implies $D_{\mathsf{FRAC}} \le c((\alpha,d)C_{\mathsf{CUT}}$. ■

We now prove the bound on the amount of power saving that can be obtained using network coding.

Theorem 6 $C_{\mathsf{ST}} \le 2c(\alpha,d)C_{\mathsf{CUT}}$.

Proof: From Lemma 5 we have $D_{\mathsf{FRAC}} \le c((\alpha,d)C_{\mathsf{CUT}}$. Using equation 8, we obtain $D_{\mathsf{ST}} \le 2c(\alpha,d)C_{\mathsf{CUT}}$. Combining this with equations 9 and 4, we obtain $C_{\mathsf{ST}} \le 2c(\alpha,d)C_{\mathsf{CODE}}$. ■

Thus the network coding advantage in reducing the power needed is at most a constant that depends solely on the power gradient and the dimension of the space. It is somewhat surprising that the power consumed by a minimum Steiner tree is only a constant factor more than the best network coding solution, *even if we only use point-to-point transmissions over the Steiner tree*.

5 Open problems

One important open problem raised by our work is whether the advantage of network coding is bounded by a constant if we consider the total network throughput that can be obtained before the network runs out of power (as opposed to the power consumed by an individual multicast as we have done in this paper).

Another open problem is to obtain tighter bounds on the constant $c(\alpha,d)$.

References

[1] A. Agarwal and M. Charikar. On the Advantage of Network Coding for Improving Network Throughput. *IEEE Information Theory Workshop*, San Antonio, Oct 25-29, 2004.

[2] A. Agrawal, P. N. Klein, and R. Ravi. When Trees Collide: An Approximation Algorithm for the Generalized Steiner Problem on Networks. *SIAM J. Comput. (SICOMP)*, 24(3), pp. 440-456, 1995.

[3] R. Ahlswede, N. Cai, S.-Y. R. Li and R. W. Yeung. Network information flow. *IEEE Trans. on Information Theory*, vol. 46, pp. 1204-1216, 2000.

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

[5] A. Clementi, P. Crescenzi, P. Penna, G. Rossi, and P. Vocca. A Worst-case Analysis of a MST-based Heuristic to Construct Energy-efficient Broadcast Subtrees in Wireless Networks. *Technical Report 010 of the Univ. of Rome "Tor Vergata,"* 2001 Preliminary version is STACS'01.

[6] A. G. Dimakis, D. Petrovic, and K. Ramchandran. From Dumb Wireless Sensors to Smart Networks using Network Coding. *Information Theory and Applications (ITA)*, 2006.

[7] L. R. Jr. Ford and D. R. Fulkerson. Maximal Flow Through a Network. *Canadian Journal of Mathematics*, pp. 99–404, 1956.

[8] C. Fragouli, J. Widmer, and J.-Y. Le Boudec. A Network Coding Approach to Energy Efficient Broadcasting: From Theory to Practice. *IEEE INFOCOM*, 2006.

[9] Z. Li, B. Li, D. Jiang, and L. C. Lau. On Achieving Optimal End-to-End Throughput in Data Networks: Theoretical and Empirical Studies. *Technical Report, University of Toronto*, May 2004

[10] W. Liang. Constructing minimum energy broadcast trees in wireless ad hoc networks. *ACM MOBIHOC*, pp. 112–122, 2002.

[11] D.S. Lun, N. Ratnakar, R. Koetter, M. Medard, M., E. Ahmed, and H. Lee. Achieving minimum-cost multicast: A decentralized approach based on network coding. *IEEE INFOCOM*, pp. 1607-1617, 2005.

[12] K. Pahlavan and A. Levesque. *Wireless information networks*. Wiley-Interscience, 1995.

[13] P. Sanders, S. Egner, and L. M. G. M. Tolhuizen. Polynomial time algorithms for network information flow. *SPAA*, pp. 286-294, 2003.

[14] J. Widmer, C. Fragouli, and J.-Y. Le Boudec. Energy-efficient broadcasting in wireless ad-hoc networks. *Netcod* 2005.

[15] Y. Wu, P. A. Chou, and S. Kung. Minimum-energy multicast in mobile ad hoc networks using network coding. *IEEE Transactions on Communications*, 53 (11), pp. 1906-1918, 2005.