

On Average Throughput and Alphabet Size in Network Coding

Chandra Chekuri
Bell Labs, Lucent
Murray Hill NJ 07974, USA
Email: chekuri@lucent.com

Christina Fragouli
School of Comput. and Commun. Sciences
EPFL, Lausanne
Emial: christina.fragouli@epfl.ch

Emina Soljanin
Bell Labs, Lucent
Murray Hill NJ 07974, USA
Email: emina@lucent.com

Abstract—We analyze a special class of configurations with h sources and N receivers to demonstrate the throughput benefits of network coding and deterministic code design. We show that the throughput benefits network coding offers can increase proportionally to \sqrt{N} , with respect to the *average* as well as the minimum throughput. We also show that while for this class of configurations there exists a deterministic coding scheme that realizes these benefits using a binary alphabet, randomized coding may require an exponentially large alphabet size.

I. INTRODUCTION

Consider a communication network represented as a directed graph $G = (V, E)$ with unit capacity edges, and h unit rate information sources S_1, \dots, S_h that simultaneously transmit information to N receivers R_1, \dots, R_N located at distinct nodes. Assume that the min-cut between the sources and each receiver node is h . The max-flow, min-cut theorem states that, if receiver R_i could utilize the network resources by himself, he would be able to receive information at rate h .

Recently it has been realized that allowing nodes in communication networks to re-encode the information they receive in addition to re-routing, increases the capacity of the network. This type of coding is termed network coding [1], [2]. In fact it was shown that by liner re-encoding, the min-cut rate can be achieved in multicasting to several sinks [1], [2]. That is, using network coding, allows each receiver to retrieve information at rate h , even when N receivers share the network resources. This is generally not the case when we use routing, *i.e.*, when we allow intermediate nodes only to forward and not to code. Thus network coding can offer throughput benefits as compared to routing.

A central question in this area is whether we can quantify how large these benefits are. Let T_{nc}^i denote the rate that receiver R_i experiences when network coding

is used, and T_u^i the rate when only uncoded transmission on is allowed. We are interested in calculating $\frac{T_u^i}{T_{nc}^i}$.

In [2] it was shown that, for undirected graphs, the throughput benefit for all receivers is bounded by a factor of two, *i.e.*, $\frac{T_u^i}{T_{nc}^i} \leq 2$. This result does not apply to directed graphs. In fact, the authors in [3] provided an example network where the throughput benefits scale proportionally to the number of sources, namely, there exists a receiver R_i such that $\frac{T_u^i}{T_{nc}^i} = \frac{1}{h}$. In other words, if we compare the minimum rate guaranteed to all receivers under routing with the rate that network coding can offer, the benefits network coding offers are proportional to the number of sources h . In [4] it was shown that these benefits equal the integrality gap of a standard linear programming formulation for the directed steiner tree problem. Known lower bounds on the integrality gap are $\Omega(\sqrt{N})$ [5] and $\Omega((\log n / \log \log n)^2)$ [6] where n is the number of nodes in the underlying graph.

For applications that are robust to loss of packets, such as real time audio and video, we might be interested in comparing the average throughput we can achieve with and without network coding, where the averaging is performed over the rate that each individual receiver experiences. This is especially true when the number of receivers is large and the throughput they experience tends to concentrate around a much larger value than the minimum. A recent paper [7], examined the average throughput achieved with routing and calculated the ratio $\frac{\sum_i T_u^i}{\sum_i T_{nc}^i}$, where $\sum_i T_{nc}^i$ and $\sum_i T_u^i$ is the sum rate for network coding and routing respectively. Using this performance measure, we showed that for a large class of configurations, including the example in [3], network coding can offer only a constant factor benefit as compared to routing. In [8], we show that the average throughput benefits of network coding can also be related to the integrality gap of a standard linear programming

formulation for the directed steiner tree problem.

The purpose of this paper is to describe and analyze a class of directed graph configurations with N receivers. These configurations were originally constructed in [5] to obtain a lower bound on the integrality gap for the directed steiner tree problem. Our observation is that these graphs can also be used to illustrate two interesting points related to network coding. First, we show that employing network coding over this class of directed graphs can offer throughput benefits proportional to \sqrt{N} , where N is the number of receivers, with regard to the average (and as a result to the minimum) throughput. Second, we show that using randomizing coding over these configurations may require an alphabet size exponentially larger than the minimum alphabet size required.

The idea in randomized network coding [3], [9] is to randomly combine over a finite field the incoming information flows and show that the probability of error can become arbitrarily small as the size of the finite field increases. We show that for this class of configurations, to guarantee a small probability of error, we may need to use an exponentially large alphabet size. In contrast, we prove that a binary alphabet size is in fact sufficient for network coding. We construct a deterministic algorithm that has linear complexity, can be used to perform network coding over this class of configurations, and requires binary alphabet.

The paper is organized as follows. In Section II we describe the structure of these configurations, and prove an upper bound on the average throughput. In Section III we construct deterministic coding schemes and in Section IV we compare them against randomized coding schemes.

II. THE NETWORK

Let N and p , $p \leq N$, be two integers and $\mathcal{I} = \{1, 2, \dots, N\}$ be an index set. We define two more index sets: \mathcal{A} as the set of all $(p-1)$ -element subsets of \mathcal{I} and \mathcal{B} as the set of all p -element subsets of \mathcal{I} . We consider a class of networks, illustrated in Fig. 1, and defined by two parameters N and p as follows: Source S transmits information to N receiver nodes $R_1 \dots R_N$ through a network of three sets of nodes A , B and C . A -nodes are indexed by the elements of \mathcal{A} , and B and C -nodes, by the elements of \mathcal{B} . An A node is connected to a B node if the index of A is a subset of the index of B . A receiver node is connected to the C nodes whose indices contain the index of the receiver. All edges in the graph have unit capacity. The out-degree of the source node is




Fig. 1. Source transmits information to N receivers.

$\binom{N}{p-1}$. Two specific members of this family of networks are shown in Fig. 2 and Fig. 3.

We can compute the degrees of the nodes in the network by simple combinatorics:

Proposition 1:

- the out-degree of A nodes is $N - (p - 1)$,
- the in-degree of B nodes is p ,
- the out-degree of C nodes is p ,
- the in-degree of the receiver nodes is $\binom{N-1}{p-1}$.

We next compute the value of the min-cut between the source node and each receiver node, or equivalently, the number of edge disjoint paths between the source and each receiver.

Theorem 1: There are exactly $\binom{N-1}{p-1}$ edge disjoint paths between the source and each receiver.

Proof: Consider receiver i . It is connected to the $\binom{N-1}{p-1}$ distinct C -nodes indexed by the elements of \mathcal{B} containing i . Each of the C -nodes is connected to the B -node with the same index. All paths between the source and the receiver i have to go through these B and C -nodes. Therefore the number of edge disjoint paths between the source and the receiver can not be larger than $\binom{N-1}{p-1}$. To show that there that many of edge disjoint paths, we proceed as follows: After removing i from the indices of the receiver i B -nodes, we are left with $\binom{k-1}{p-1}$ distinct sets of size $p - 1$, *i.e.* distinct elements of \mathcal{A} . We use the A -nodes indexed by these elements of \mathcal{A} to connect the receiver i B -nodes to the source. ■

Therefore, the sum rate with network coding T_{nc} is equal to $N \binom{N-1}{p-1}$. We next find an upper bound to the sum rate without network coding T_u and the to the ratio T_u/T_{nc} .

Theorem 2: In a network in Fig. 1, when $h = \binom{N-1}{p-1}$

$$\frac{T_u}{T_{nc}} \leq \frac{p-1}{k-p+1} + \frac{1}{p}. \quad (1)$$

Proof: If only routing is permitted, the information is transmitted from the source node to the receiver through a number of trees, each carrying a different information source. Let a_t be the number of A -nodes in tree t , and c_t , the number of B and C -nodes. Note that $b_t \geq a_t$, and that the c_t C -nodes are all descendants of the a_t A -nodes. Therefore, we can count the number of the receivers spanned by the tree as follows: Let $n_t(A(j))$ be the number of C -nodes connected to the j -th A -node in the tree. Note that

$$\sum_{j=1}^{a_t} n_t(A(j)) = c_t.$$

The maximum number of receivers the tree can reach through this A -node is $n_t(A(j)) + p - 1$. Consequently, the maximum number of receivers the tree can reach is

$$\sum_{j=1}^{a_t} [n_t(A(j)) + p - 1] = a_t(p - 1) + c_t.$$

To find an upper bound to the routing throughput, we need to find the number of receivers that can be reached by a set of disjoint trees. Note that for any set of disjoint trees we have

$$\sum_t a_t = \binom{k}{p-1} \text{ and } \sum_t c_t = \binom{k}{p}.$$

Therefore, T_u can be upper-bounded as

$$\begin{aligned} T_u &= \sum_t (a_t(p-1) + c_t) \\ &= (p-1) \sum_t a_t + \sum_t c_t \leq (p-1) \binom{N}{p-1} + \binom{k}{p}. \end{aligned}$$

The sum rate with network coding T_{nc} is equal to $N \binom{N-1}{p-1}$. Thus we get that

$$\frac{T_u}{T_{nc}} \leq \frac{p-1}{N-p+1} + \frac{1}{p}.$$

For a fixed N , the LHS of the above inequality is minimized for

$$p = \frac{N+1}{\sqrt{N}+1} \cong \sqrt{N},$$

and for this value of p ,

$$\frac{T_u}{T_{nc}} = 2 \frac{\sqrt{N}}{1+N} \lesssim \frac{2}{\sqrt{N}}. \quad (2)$$

III. DETERMINISTIC CODING

We show that network coding can be done by using the binary alphabet. We first explain how the coding is done for two special cases of p : case when $p = 2$ and case when $p = N - 1$, and then proceed with the general case.

A. $p = 2$

Consider the case when $p = 2$ for arbitrary N . An example for $N = 4$ is shown in Fig. 2. In this case the




Fig. 2. $N=4, p=2$

number of information sources is $h = N - 1$. We can code over the binary field as follows: Since the number of edges going out of S into A nodes is N , we can send the $N - 1$ sources over the first $N - 1$ of these edges and not use the N -th edge. In other words, the coding vector of the i -th of these edges is the i -th basis vector e_i for $i = 1, 2, \dots, N - 1$. The B -nodes merely sum their inputs over \mathbb{F}_2^h , and forward the result to the C -nodes. Consequently, the coding vectors on the branches going to receiver N are the $N - 1$ basis vectors, and the coding vectors on the branches going to receiver i for $i = 1, 2, \dots, N - 1$ are e_i and $e_j + e_i$ for $j = 1, \dots, N - 1$ and $j \neq i$.

B. $p = N - 1$

Consider the case when $p = 2$ for arbitrary N . An example for $N = 5$ is shown in Fig. 2. In this case the number of information sources is $h = N - 1$. The number of C -nodes is N . Each subset of $N - 1$ C -nodes is observed by a receiver. Therefore coding vectors of the edges between the B and C -nodes must belong to an arc, namely any $N - 1$ of them must be linearly independent.




Fig. 3. $N=5, p=4$

The following maximal arc over $\mathbb{PG}(N-2, 2)$ has N points:

$$\begin{pmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & 1 \\ 1 & 1 & \dots & 1 \end{pmatrix} \quad (3)$$

We can obtain this arc by coding as follows: To the $N-1$ edges going from the source to the A nodes whose label does not contain N , we assign $N-1$ basis vectors of over $\mathbb{F}_2^{(N-1)}$. We remove all other edges outgoing of the source, and then all A -nodes which lost their connection with the source, and the edges coming out of the removed A nodes. Consequently, the first of the B -nodes has $N-1$ inputs. By addition, of these inputs the coding vector between this B and its corresponding C node becomes $(11\dots 1)$. The rest of the B -nodes have only one input. Thus we get the binary arc 3) at the last set of edges.

C. The General Case

For arbitrary values of p and N , network coding can be done as follows: We first remove the edges going out of S into those A -nodes whose labels contain N . There are $\binom{N-1}{p-2}$ such edges. Since the number of edges going out of S into A -nodes is $\binom{N}{p-1}$, the number of remaining edges is $\binom{N}{p-1} - \binom{N-1}{p-2} = \binom{N-1}{p-1}$. We label these edges by the $h = \binom{N-1}{p-1}$ different basis elements of \mathbb{F}_2^h . We further remove all A -nodes which have lost their connection with the source S , as well as their outgoing edges. The B -nodes merely sum their inputs over \mathbb{F}_2^h , and forward the result to the C -nodes.

Consider a C -node that the N -th receiver is connected to. Its label, say ω , is a p -element subset of \mathcal{I} containing N . Because of our edge removal, the only A -node that this C -node is connected to is the one with the label $\omega \setminus \{N\}$. Therefore, all C -nodes that the N -th receiver is connected to have a single input, and all those inputs are different. Consequently, the N -th receiver observes all the sources directly.

Each of the receivers $1, 2, \dots, N-1$ will have to solve a system of equations. Consider one of these receivers, say j . Some of the C -nodes that the j -th receiver is connected to have a single input: those are the nodes whose label contains N . There are $\binom{N-2}{p-2}$ such nodes, and they all have different labels. For the rest of the proof, it is important to note that each of these labels contains j , and the $\binom{N-2}{p-2}$ labels are all $(p-1)$ -element subsets of \mathcal{I} which contain j and do not contain N . Let us now consider the remaining $\binom{N-1}{p-1} - \binom{N-2}{p-2} = \binom{N-2}{p-1}$ C -nodes that the j -th receiver is connected to. Each of these nodes is connected to p A -nodes. The labels of $p-1$ of these A -nodes contain j , and only one does not. That label is different for all C -nodes that the receiver j is connected to. Consequently, the j -th receiver gets $\binom{N-2}{p-2}$ sources directly, and each source of the remaining $\binom{N-2}{p-1}$ as a sum of that source and some $p-1$ of the sources received directly.

D. A Lower Bound to the Throughput T_u

We can derive a lower bound to the sum rate when only routing is used as follows: Assume that we route h sources into the A -nodes as we did for the coding described in the previous section and forward them to through only those B -nodes which have a single input. Then one receiver will receive all $\binom{N-1}{p-1}$ sources and $N-1$ receiver will receive $\binom{N-2}{p-2}$ sources. We have

$$\begin{aligned} T_u &\geq \binom{N-1}{p-1} + (N-1) \binom{N-2}{p-2} \\ &= \binom{N-1}{p-1} + (p-1) \binom{N-1}{p-1} \\ &= p \binom{N-1}{p-1}, \end{aligned}$$

and consequently

$$\frac{T_u}{T_{nc}} \geq \frac{p}{N}.$$

IV. RANDOM CODING

A. General Networks

For a general network with N receivers in which coding is performed by random assignment of coding vectors over the alphabet \mathbb{F}_q , a lower bound to the probability P_N^d that all N receivers will be able to decode is derived in [9] to be

$$P_N^d \geq \left(1 - \frac{N}{q}\right)^n,$$

where n is defined in [9] to be the number of edges where coding is performed. In our case, $n \geq \binom{N}{p}$, and the lower bound becomes

$$P_N^d \geq \left(1 - \frac{N}{q}\right)^{\binom{N}{p}} \approx e^{-\frac{N\binom{N}{p}}{q}}.$$

We next look into randomized coding for the class of networks under investigation. We first consider the case when randomized coding is used at all nodes with multiple inputs, namely the source node and all the B nodes, and then the case when the coding at the source node is done deterministically as in Sec. III-C, and randomized coding is done at the B nodes with multiple inputs after the removal of edges as in Sec. III-C.

B. Random Coding for the Special Class of Networks

First, randomized coding is used at the source node to decide which linear combination goes to each A -node. Then:

Pr (receiver j has a full rank set of equations)=
 Pr (each node C receiver j observes increases his rank)=
 $\prod_{i=2}^h \Pr$ (node C_i that receiver j observes increases his rank)=

$\prod \Pr$ (node C_i increases receiver j rank | $\{A_i\}$ inputs of C_i do not lie in the span of $\{C_1 \dots C_{i-1}\}$) Pr($\{A_i\}$ inputs of C_i do not lie in the span of $\{C_1 \dots C_{i-1}\}$) =
 $\geq \prod (1 - \frac{1}{q})^2 = (1 - \frac{1}{q})^{2(h-1)} = (1 - \frac{1}{q})^{2\binom{N-1}{p-1}}$.

C. Random Coding at B nodes

Assume that we choose the coding vectors for the edges going into the A -nodes as we did for the deterministic coding described in Sec. III-C, but now the B -nodes randomly combine their inputs instead of summing them.

Consider receiver j . As before $\binom{N-2}{p-2}$ or its C -nodes are connected to a single input. Consider one of the remaining $\binom{N-2}{p-1}$ C -nodes that the receiver j is connected to. The corresponding B node will form a random linear combination of the $p-1$ sources that are directly received and of an additional source. Therefore, if the random linear combining is performed over \mathbb{F}_q , the C

will observe a linear combination of only the $p-1$ sources directly received with probability $1/q$, namely only if the coefficient zero is chosen for the additional source. Thus the receiver j receives an independent linear combination from a C node with p inputs with probability $1-1/q$. Since the linear combining at each multi-input B node is performed independently, receiver j will be able to decode all h sources with probability

$$\Pr\{\text{single receiver decodes}\} = \left(1 - \frac{1}{q}\right)^{\binom{N-2}{p-1}}.$$

We can also compute the probability that all receivers be able to decode all sources. Note that this happens when all multi-input B nodes use a nonzero coefficient for the not-directly received source. Since there are $\binom{N}{p} - \binom{N-1}{p-1} = \binom{N}{p-1}$ such nodes, we obtain

$$\Pr\{\text{all receiver decode}\} = \left(1 - \frac{1}{q}\right)^{\binom{N}{p-1}}.$$

Thus similarly with before, if we want this probability to be greater than e^{-1} , we need to choose $q \geq \binom{N}{p-1}$.

REFERENCES

- [1] R. Ahlswede, N. Cai, S-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Transac. Inform. Theory*, pp. 1204–1216, July 2000.
- [2] S-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Transac. Inform. Theory*, vol. 49, pp. 371–381, Feb. 2003.
- [3] P. Sanders, S. Egner, and L. Tolhuizen, "Polynomial time algorithms for network information flow," *Proc. 15th ACM Symposium on Parallel Algorithms and Architectures*, 2003.
- [4] A. Agarwal and M. Charikar, "On the advantage of network coding for improving network throughput," in *Proc. of 2004 IEEE Information Theory Workshop*, San Antonio, Texas, Oct. 2004.
- [5] L. Zosin and S. Khuller, "On directed steiner trees," in *Proc. 13th Annual ACM/SIAM Symposium on Discrete Algorithms (SODA '02)*, 2002, pp. 59–63.
- [6] E. Halperin, G. Kortsarz, R. Krauthgamer, A. Srinivasan, and N. Wang, "Integrality ratio for group steiner trees and directed steiner trees," in *Proc. of 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Jan. 2003.
- [7] C. Fragouli and E. Soljanin, "On average throughput benefits for network coding," *Proc. 42nd Annual Allerton Conference*, Monticello, IL, Oct. 2004.
- [8] C. Chekuri, C. Fragouli, and E. Soljanin, "On average throughput and alphabet size in network coding," in preparation for *IEEE Trans. Inform. Theory*.
- [9] T. Ho, M. Médard, J. Shi, M. Effros, and D. R. Karger, "On randomized network coding," *Proc. 41st Annual Allerton Conference*, Monticello, IL, Oct. 2003.

ACKNOWLEDGMENT

This work was supported the NSF under award No. CCR-0325673 and FNS under award No. 200021-103836/1.