Network-Coding in Interference Networks

Brian Smith and Sriram Vishwanath
Department of Electrical and Computer Engineering
University of Texas at Austin
1, University Station C0803
Austin, Texas 78712
e-mail: {bsmith, sriram}@ece.utexas.edu

Abstract — **This paper investigates the gains of network coding in Gaussian interference networks. The term "interference networks" indicates networks where different users' transmissions effect each other's reception in some way (i.e. interfere). Following the well-known results for non-interfering, noise-free networks and for erasure networks, it is shown using examples that network coding in interference networks can also result in substantial gains. Intriguingly, network coding is found to benefit interference networks for configurations where there is no obvious advantage in the equivalent non-interference network. However, it is found that unlike non-interfering networks, non-cooperative network coding is a suboptimal coding strategy.**

I. Introduction

Recent results on routing information over networks have shown that by allowing nodes which can perform computations, rather than limiting the nodes solely to forwarding the data, can in some cases increase the total throughput, [1, 9, 7, 11, 6, 5, 4, 3, 10, 8]. Consider routing over a network as shown in Figure 1. Data from a distant set of sources S is being routed to nodes A and B which is then relayed to the ultimate set of receivers C. Conventional wisdom suggests that plain routing of packets be performed by either picking one of the paths (either A or B), or by sending sending unrelated information to A and B. A fast emerging area in networks is to relate information along these two paths. Network coding, one method of producing such a relationship, is a distributed coding technique where each node transmits linear combinations of the data available to that node. Another method of use for relay networks is block-Markov superposition encoding [2], which also relates the transmitted codewords by introducing a dependence over multiple pieces of data [12].

We will use illustrative examples to demonstrate the use of both of these techniques, first individually and then in concert, over various Gaussian interference networks. For any non-interference network, define the equivalent interference network to be the network with the same topological interconnections between nodes. The difference between the two equivalent networks is that in the non-interference case, signals on different the edges entering or leaving a particular node will have no impact on each other, while in the interference network the signals will interact. The examples in this paper will show that for some network configurations, the same network coding schemes that are used in the non-interfering network will be appropriate for use in an equivalent interference network. However, other examples show that there also exist configurations for which network coding is of benefit only to the interference network, and not to the equivalent non-interference network.

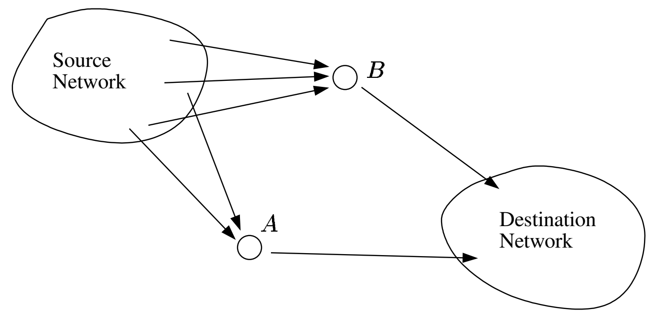

Figure 1: Correlations Can be Induced on a Network

II. Classical Network Coding

The introduction and primary analysis of network coding has been for bit-pipe networks, i.e. those where each channel from transmitter to receiver is assumed ideal and where transmissions are assumed not to interfere with each other. Other work has studied erasure networks [13], but with a non-interfering multiple access receiving model. In these scenarios, significant gains from employing network-coding as opposed to relaying have been established. In fact, network-coding has been shown to be the optimal transmission scheme when multicasting information over bit-pipe networks and over some erasure networks [1, 13].

The classic network-coding example from [1], illustrated in Figure 2, is a great means of illustrating the gains of network-coding over routing. In this example, source S desires to communicate information A and B to destinations D_1 and D_2, at the maximum rate possible. Each link is assumed to be an ideal channel supporting a data rate of 1 bit/transmission. As is well known, the link $X \rightarrow Y$ in Figure 2 acts as a bottleneck: if either A's data or B's data is sent over this link, one of the receivers D_1 or D_2 cannot receive data at the full max-flow min-cut rate of 2 bits/transmission that it could if it were the only intended receiver. The surprising result of network coding is that a multicasting rate of R is achievable from a single source to all receivers simultaneously as long as each

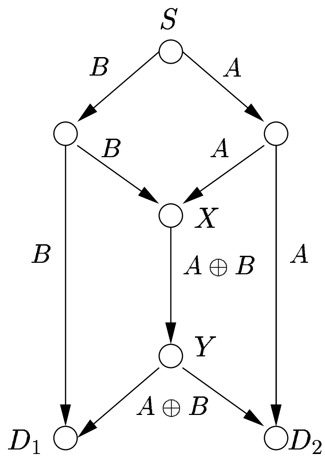

Figure 2: The Network Coding Example

of their min-cut max-flow values are greater than or equal to R. This of course implies that there is a method to achieve 2 bits/transmission to D_1 and D_2, and the appropriate strategy is to send the the exclusive-OR of A and B through link $X \to Y$. Each destination, knowing $A \oplus B$ and either the bit A or B, can decode the other bit.

We can also consider replacing this non-interfering network with an equivalent Gaussian interference network which has the same nodes and edge connections, as in Figure 3. The received signal at each node shall be the sum of all the transmitted signals (each perhaps multiplied by a constant gain) intended for that node, with the addition of unit power Gaussian noise. Each node will transmit complex signals limited to a given maximum power. Now, each node is broadcasting information to neighboring nodes in the sense of being an information theoretic *broadcast channel*, and receiving information on a noisy uplink (a *multiple access channel*). This model could easily be replaced by an alternate interference channel model (such as a binary symmetric channel), but all the examples analyzed here use a Gaussian network.

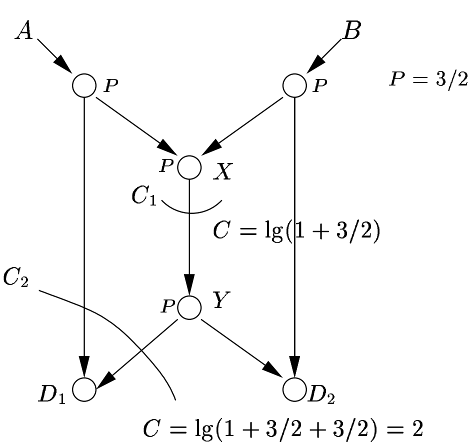

Figure 3: A Gaussian Interference Network

Allow the power available at the source to be arbitrarily large, so that we can view the model as having two independent sources A and B and focus on the more interesting portions of the network. All other nodes will be allotted a power of $P = 3/2$ and all channel gains will be unity. Thus, the capacity of a single link, such as across the cut C_1 in Figure 3 is $\lg(1 + P) = \lg(5/2) < 2$ bits/transmission. The sum capacity of a multiple access channel with two independent transmitters, for example across cut C_2 in Figure 3, will be $\lg(1 + P + P) = 2$ bits/transmission.

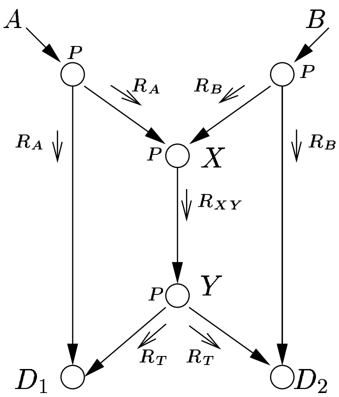

Figure 4: Rate Constraints for the Gaussian Network

Referring to Figure 4, if each broadcast node is used as a multicast (sending the same information down each link), the rate constraints are:

$$R_A \leq \lg(1 + 3/2) \tag{1}$$
$$R_B \leq \lg(1 + 3/2) \tag{2}$$
$$R_A + R_B \leq \lg(1 + 3/2 + 3/2) \tag{3}$$
$$R_{XY} \leq \lg(1 + 3/2) \tag{4}$$
$$R_T \leq \lg(1 + 3/2) \tag{5}$$
$$R_T + R_A \leq \lg(1 + 3/2 + 3/2) \tag{6}$$
$$R_T + R_B \leq \lg(1 + 3/2 + 3/2) \tag{7}$$

Thus, a pure routing strategy on this interference network will suffer from the same limitations as a pure routing strategy in the bit-pipe model of Section II. The link $X \to Y$ across cut C_1 in Figure 3 has a maximum capacity of $\lg(5/2)$ bits/transmission (from Equation 4). However, data can arrive at node X at a rate of 2 bits/transmission (Equation 3), so with a pure routing strategy, some of the data will be lost. This is the analogue to the situation in the equivalent non-interference network in Figure 2, where the network must decide whether to route either the bit from A or the bit from B.

It is straightforward to see that an average rate of 1.5 bits/transmission can be achieved in this interference network[Fig 3]: Assign a value of 1 to all of the rates R_A, R_B, R_{XY}, and R_T. Let the information sent by node X to Y and multicast by node Y to D_1 and D_2 be the bits from source A. Node D_1 receives A's data, and node D_2 receives data from both A and B. Time-sharing can be used so that on average, 1.5 bits/transmission arrive at each destination node. It can be shown (using linear programming with appropriate constraints) that no better pure routing strategy, with the broadcast nodes acting as multicast, exists.

However, using the network coding scheme of the equivalent classic non-interference network from Figure 2, the links $X \to Y$, $Y \to D_1$, and $Y \to D_2$ in Figure 4 can all carry the data $A \oplus B$. Thus, 2 bits/transmission are received at each terminal. The rate of 2 bits/transmission is known to be the maximum rate achievable to each terminal using independent codebooks. The maximum rate is bounded by each cut across the network, and in particular observe the cut C_2 in Figure 3. This cut is across the multiple access channel and therefore has a sum rate of $\lg(1+3/2+3/2) = 2$. With 2 bits/transmission to each terminal, network coding has achieved the upper bound. Exceeding this rate across a multiple access channel by correlating the codewords is the subject of Section IV.

III. Example: Benefit of Network Coding on an Interference Network

Recall that in the network coding scheme described for the interference network of Figure 4 all of the broadcast channel nodes are used as multicast. Specifically, each node sends the same information to all of the receivers that are connected to it. Intuitively, multicasting is a good idea for this configuration because the sum rate obtained by broadcasting different information to each receiver will in general be less than the multicasting rate. This occurs because the portion of the broadcast signal that a particular receiver does not decode is seen by that receiver as Gaussian noise. The decrease in rate is illustrated in Figure 5. A broadcast channel with power constraint P and unit noise can either 1) multicast at a rate of $\lg(1+P)$ bits/transmission to both nodes or 2) broadcast at a rate of $\lg(1+\alpha P)$ to one node and $\lg(1+\frac{(1-\alpha)P}{1+\alpha P})$ to the other (where $0 \le \alpha \le 1$) for a total of $\lg(1+P)$ bits/transmission. From one point of view, the sum rate of information transfer is the same in either case, since the same amount of information crosses cut C_1; however, in the multicast scheme, a greater amount of information crosses each of the cuts C_2 and C_3 than in the broadcast scheme.

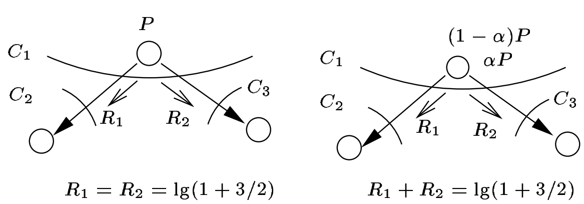

Figure 5: Multicasting and Broadcasting

To lend further credence to the argument that the multicasting scheme is often superior to broadcasting, consider the network in Figure 6. Three independent sources with power P and P', as indicated, desire to transmit to two receiver nodes. When the center node is multicasting, a sum rate of

$$\lg \frac{(1+P'+P)\,(1+P'+P)}{(1+P)} \tag{8}$$

to the two multiple-access nodes achieved. (The denominator term is the upper bound on the rate of duplicate information being sent to the receivers). Under a broadcasting scheme for the center node, the rate can be only

$$\begin{aligned}\lg\left(\left(1+\tfrac{P'+\alpha P}{1+(1-\alpha)P}\right)\left(1+\tfrac{P'+(1-\alpha)P}{1+\alpha P}\right)\right)\\ = \lg \tfrac{\left(1+P'+P\right)\left(1+P'+P\right)}{(1+P+\alpha(1-\alpha)P^2)}\end{aligned} \tag{9}$$

which is less than the rate achieved in multicast mode, even when accounting for the duplicate information being sent. The point of this example is to demonstrate that, in at least the simple examples of a interference networks studied here, multicasting is a reasonable strategy to employ.

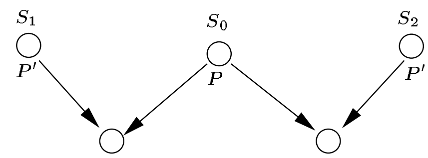

Figure 6: An Example: Multicasting Outperforms Broadcasting

If a broadcast node in an interference network is constrained to send the same information down each link, then it is acting in a similar manner to a configuration shaped like a "T" in a non-interference network, as in Figure 7. In this "T" sub-network, if each pipe has equal capacity, then there is no advantage in not sending the full content available at the central node out each of its branches. This simile will be utilized in the following example.

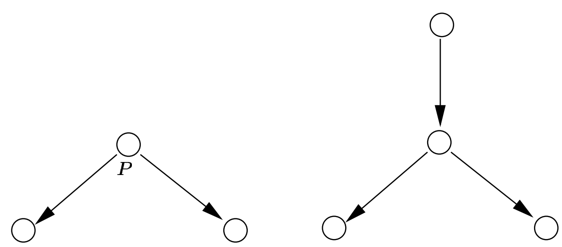

Figure 7: Multicasting in an Interference Network is Like a "T" in a Bit-Pipe Network

When an interference network contains a broadcast node, which acts similarly to a "T" in a non-interference network [Fig 7], it opens up the possibility of a network coding benefit. The network in Figure 8, with a "bow-tie" shape, is just such a network. In the equivalent non-interference network of Figure 9, there is no benefit to network coding: the full rate of 2 bits/transmission to each receiver is clearly achievable by routing alone. We shall demonstrate that the network coding is an appropriate technique, however, for the equivalent interference network of Figure 10.

In the "bow-tie" Gaussian interference network of Figure 10, $\lg(5/2)$ bits/transmission can be obtained by "turning off" the central node X. In that case, there is no interference and data is routing from each source to a single corresponding destination node. Even better, if each of the three broadcast nodes

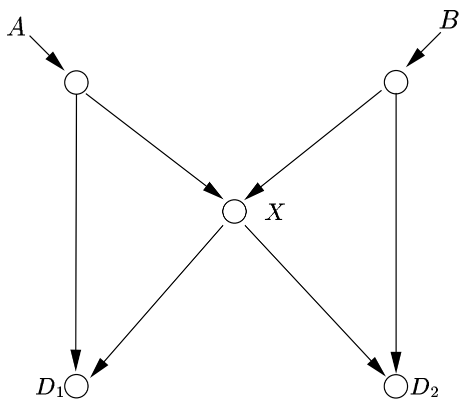

Figure 8: Bow-Tie Network

multicasts at full power $P = 3/2$ (with the central node alternately choosing A or B to transmit), then an average rate of 1.5 bits/transmission is achieved. This procedure is effectively identical in structure to the routing without network coding performed in the classic non-interference example of Figure 2.

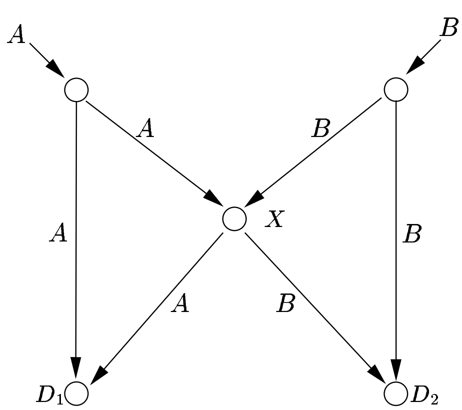

Figure 9: Network Coding Cannot Benefit Non-Interference Bow-Tie Network

But, in accordance with the classic example, the central node X in Figure 10 can perform the exclusive-OR operation on the data A and B and multicast the result to the receivers D_1 and D_2. This scheme will allow the maximum rate of 2 bits/transmission (which is still bounded above by the capacity of the cut C_1 across a multiple access channel) to be achieved. Some interference networks can actually benefit from network coding, as shown in this example, even when the equivalent non-interference network cannot. Intuitively, this phenomenon occurs because of the interaction between data intended for two difference receivers by the same transmitter in the channel models for interference networks. In Figure 10, broadcast node X, by multicasting the same data to nodes D_1 and D_2, is acting as if it were in the non-interference configuration of Figure 7.

Further, it is tempting to propose a parallel hypothesis to Ahlswede's network coding theorem, namely that "If R is less than the minimum of the rates achievable from a single source

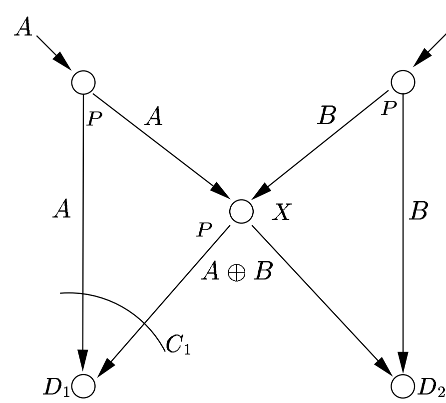

Figure 10: Network Coding Increases Throughput for Interference Model

node S to each of several terminal nodes D_i individually in a Gaussian interference network by using independent codewords, then a rate R is achievable to the nodes D_i simultaneously in the same Gaussian network." However, much future work remains to prove or to find counter-example to such a proposition.

IV. Example: Node-Cooperation and Network Coding in Concert

All of the previous examples of network coding on interference networks in this paper assume that each node uses independently generated codebooks. Since the bulk of the work in network coding theory has been performed concerning non-interference networks, non-cooperative network coding is the logical first step to take in studying network coding on interference networks. It is known, however, that node-cooperative strategies, i.e. correlating codewords, can increase the rate of point-to-point transmission over some networks (the relay channel, or the multiple access channel with correlated sources, for example). The next example will show that using node-cooperation and network-coding strategies in tandem can be a useful tool in coding schemes for interference networks.

Consider a portion of a Gaussian interference network as detailed in Figure 11. Through the previous actions of the network, the node T_1 has access to the source U at an arbitrarily high rate, the node T_2 has V, and the center node T_0 has obtained both U and V. For clarity, T_0 is assumed to have power constraint P_0 while the other nodes are each limited to P. Allow $P = P_0 = 3/2$. By routing alone, the arguments in Section III show that a maximum of 1.5 bits/transmission, on average, can reach the destination nodes. The obvious network coding (center node T_0 multicasts $A \oplus B$) scheme again increases this rate to 2 bits/transmission. However, an even greater rate can be obtained by allowing the nodes to cooperate and send correlated signals.

Let γ and γ_0 be powers bounded by $0 \leq \gamma \leq P$ and $0 \leq \gamma_0 \leq P_0$. We will use timesharing and define two modes of operation for the system, each to take place for 50% of the overall operation: For the first mode, generate three codebooks

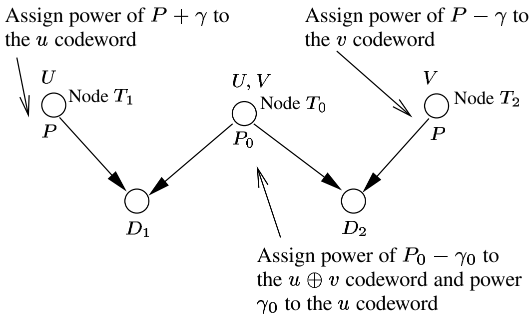

Figure 11: Network Coding and Node Cooperation

of size 2^{nR_u}, 2^{nR_v}, and 2^{nR_*}, each as $p(x) \sim \mathcal{N}(0,1)$, so that there exist codewords $x_u^n(u)$, $x_v^n(v)$ and $x_*^n(t)$ for $u \in \{1..2^{nR_u}\}$, $v \in \{1..2^{nR_v}\}$, $t \in \{1..2^{nR_*}\}$. The index t shall represent the value of $u \oplus v$.

In each time-slot of this mode, the transmitter node T_1 sends the codeword $x_1^n(u) = \sqrt{P+\gamma}x_u^n(u)$ which has a power of $P+\gamma$. The transmitter node T_2 sends the codeword $x_2^n(v) = \sqrt{P-\gamma}x_v^n(v)$, which has a power of $P-\gamma$. The central node T_0 sends the codeword $x_0^n(u,v) = \sqrt{\gamma_0}x_u^n(u) + \sqrt{P_0-\gamma_0}x_*^n(t)$, which has a total power of P_0.

The appropriate indices u,v, and t are recovered correctly (by jointly typical decoding, for example) at both multiple-access destination nodes when n is large and the following conditions are met. From the destination node D_1:

$$R_* < \lg\left(1 + \frac{P_0-\gamma_0}{1}\right) \tag{10}$$

$$R_u < \lg\left(1 + \frac{\left(\sqrt{P+\gamma}+\sqrt{\gamma_0}\right)^2}{1}\right) \tag{11}$$

$$R_u + R_* < \lg\left(1 + \frac{\left(\sqrt{P+\gamma}+\sqrt{\gamma_0}\right)^2 + P_0 - \gamma_0}{1}\right) \tag{12}$$

and from the multiple access channel at destination node D_2:

$$R_* < \lg\left(1 + \frac{P_0-\gamma_0}{1+\gamma_0}\right) \tag{13}$$

$$R_v < \lg\left(1 + \frac{P-\gamma}{1+\gamma_0}\right) \tag{14}$$

$$R_v + R_* < \lg\left(1 + \frac{P_0-\gamma_0+P-\gamma}{1+\gamma_0}\right) \tag{15}$$

In the second mode, the roles of T_1 and T_2 are reversed, so transmitted power at these nodes will each ultimately average to P. The codebooks of X_U and X_V can simply be swapped and reused for this mode.

On average, each destination node receives $(R_u + R_v)/2$ bits/transmission directly. From that information, and from the exclusive-OR information it receives at a rate R_*, the destination nodes can then decode the opposite sources. This decoding will be at a rate which is the minimum of R_* and $(R_u + R_v)/2$ bits/transmission. Numerical calculations show that each destination can receive at a rate of approximately 2.06 by assigning γ and γ_0 both values of 0.1 and performing appropriate rate-splitting at the multiple-access receivers.

V. Conclusions

The salient features of this paper are:

1. The "bit-pipe" model, where each channel between two nodes is modeled as a non-interfering pipe with a threshold capacity c, is often a useful abstraction. However, it suffers from the following issues:
   - It neglects interactions between channels such as interference and other constraints on the system.
   - In a wireless scenario, it imposes an orthogonality on adjacent channels, which in many cases may not be the best strategy.
   - Any gains due to node-cooperation are often lost with this model.
2. A noisy channel model captures many of these effects. It is however, more complex to handle. Results such as network coding gain, that are known for the "bit-pipe" model can be replicated for the noisy channel model.
3. Simple network configurations exist where the bit-pipe model promises no network-coding gain, but the equivalent interference model shows network-coding gain by interference mitigation.
4. Node cooperation strategies, which by definition have no effect in non-interference networks, can increase the throughput of interference networks above the rate which is possible when the network is limited to using independent codebooks at each node.

References

[1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung. Network information flow. *IEEE Trans. Inform. Theory*, 46:1204–1216, 2000.

[2] T. M. Cover and A. El Gamal. Capacity theorems for the relay channel. *IEEE Trans. Inform. Theory*, 25:572–584, 1979.

[3] R. Dougherty, C. Freiling, and K. Zeger. Linearity and solvability in multicast networks. *Proc. CISS*, 2004.

[4] C. Fragouli, E. Soljanin, and A. Shokrollahi. Network coding and a coloring problem. *Proc. CISS*, 2004.

[5] T. Ho, M. Médard, M. Effros, and R. Koetter. Network coding for correlated sources. *Proc. CISS*, 2004.

[6] T. Ho, M. Médard, J. Shi, M. Effros, and D. R. Karger. On randomized network coding. *Proc. 41st Annual Allerton Conf. on Commun. Control and Computing*, 2003.

[7] R. Koetter and M. Médard. An algebraic approach to network coding. *IEEE/ACM Trans. Networking*, 11:782–795, 2003.

[8] S. Y. R. Li and R. W. Yeung. Network multicast flow via linear coding. *Proc. Symp. Operations research and Appl.*, pages 197–211, 1998.

[9] S.-Y.R. Li, R. W. Yeung, and N. Cai. Linear network coding. *IEEE Trans. Inform. Theory*, 49:371–381, 2003.

[10] Z. Li and B. Li. Network coding in undirected networks. *Proc. CISS*, 2004.

[11] L. Song and R. W. Yeung. Network information flow - multiple sources. *Proc. IEEE Intl. Symp. Inform. Theory*, page 102, 2001.

[12] G. Kramer, M. Gastpar and P. Gupta, "Capacity theorems for wireless relay channels", *Proc. 41st Allerton Conf. on Commun., Control and Computing*, Oct 1-3, 2003.

[13] R. Gowaikar, A. F. Dana, R. Palanki, B. Hassibi, M. Effros, "On the Capacity of Erasure Relay Networks", *Proceedings of ISIT 2004*, 2004.