On Randomized Network Coding[*]

Tracey Ho[†], Muriel Médard[†], Jun Shi[§], Michelle Effros[‡] and David R. Karger[†]

[†]Massachusetts Institute of Technology, [§]University of California, Los Angeles,
[‡]California Institute of Technology

[†]{trace@, medard@, karger@csail.}mit.edu, [§]junshi@ee.ucla.edu, [‡]effros@caltech.edu

Abstract

We consider a randomized network coding approach for multicasting from several sources over a network, in which nodes independently and randomly select linear mappings from inputs onto output links over some field. This approach was first described in [3], which gave, for acyclic delay-free networks, a bound on error probability, in terms of the number of receivers and random coding output links, that decreases exponentially with code length. The proof was based on a result in [2] relating algebraic network coding to network flows. In this paper, we generalize these results to networks with cycles and delay. We also show, for any given acyclic network, a tighter bound in terms of the probability of connection feasibility in a related network problem with unreliable links. From this we obtain a success probability bound for randomized network coding in link-redundant networks with unreliable links, in terms of link failure probability and amount of redundancy.

1   Introduction

We consider a randomized network coding approach for distributed transmission and compression of information in multi-input multicast networks. This family of problems includes traditional single-source multicast for content delivery, and the reachback problem for sensor networks, in which several, possibly correlated, sources transmit to a single receiver.

In this approach, first described in [3], each link carries a linear combination of signals from incident incoming links. The linear coefficients for each link are independently and randomly chosen from some finite field. The receivers need only know the overall linear combination of source processes in each of their incoming signals. This information can be sent through the network as a vector, for each signal, of coefficients corresponding to each of the source processes, updated at each coding node by applying the same linear mappings to the coefficient vectors as to the information signals.

Reference [3] considered independent or linearly correlated sources on acyclic delay-free networks, and showed an upper bound on error probability that decreases exponentially with the length of the codes. The proof was based on results in [2] linking multicast network coding to network flows/bipartite matching. It was noted that this approach

---

achieves robust routing and compression in combination within the network, differing from traditional approaches which first do source/diversity coding followed by routing of coded information. Any available network capacity can be fully exploited for robustness to link failures and coding error, while not hampering accommodation of new multicast sources.

In this paper, we generalize these results to networks with cycles and delay. We also show, for acyclic networks, a relation between the probability of randomized coding success and the probability of connection feasibility in a related network problem with unreliable links. This result is useful for obtaining tighter bounds on randomized coding success probability that are based on more specific network characteristics. From it we obtain a tighter bound for link-redundant networks with unreliable links, in terms of link failure probability and amount of redundancy, showing how these factors affect randomized coding success probability.

Going from the acyclic delay-free case to the case with cycles and delay, the scalar coefficients of the linear combinations become polynomials in a delay variable. The number of terms of these polynomials that must be sent, and the memory required at the receivers, depend on the number of links involved in cycles (memory registers) in the network. For less frequently changing networks, instead of sending coefficient vectors through the network, there can be a phase in which the sources take turns to each send a unit impulse through the network.

1.1 Overview

A brief overview of related work is given in Section 1.2. In Section 2, we provide the algebraic model we consider for our networks. Our main results are given in 3, and our proofs and ancillary results in Section 4. We present our conclusions and some directions for further work in Section 5.

1.2 Related Work

Ahlswede et al. [1] showed that with network coding, as symbol size approaches infinity, a source can multicast information at a rate approaching the smallest minimum cut between the source and any receiver. Li et al. [6] showed that linear coding with finite symbol size is sufficient for multicast. Koetter and Médard [5] presented an algebraic framework for network coding that recaptured previous results and gave an algebraic condition for checking the validity of a given linear multicast code. Sanders et al. [7] and Jaggi et al. [4] proposed centralized algorithms for single source multicast using a subgraph consisting of flow solutions to individual receivers, and showed that randomization with centralized testing could yield computational advantage.

2 Model

We adopt the model of [5], which represents a network as a directed graph $\mathcal{G}$. Discrete independent random processes $X_1, \ldots, X_r$ are observable at one or more source nodes, and there are $d \geq 1$ receiver nodes. The output processes at a receiver node β are denoted $Z(\beta, i)$. The *multicast* connection problem is to transmit all the source processes to each of the receiver nodes.

There are ν links in the network. Link l is an *incident outgoing link* of node v if $v = \text{tail}(l)$, and an *incident incoming link* of v if $v = \text{head}(l)$. We call an incident outgoing link of a source node a *source link* and an incident incoming link of a receiver node a *terminal link*. Edge l carries the random process $Y(l)$.

The time unit is chosen such that the capacity of each link is one bit per unit time, and the random processes X_i have a constant entropy rate of one bit per unit time. Edges with larger capacities are modelled as parallel edges, and sources of larger entropy rate are modelled as multiple sources at the same node.

The processes X_i, $Y(l)$, $Z(\beta, i)$ generate binary sequences. We assume that information is transmitted as vectors of bits which are of equal length u, represented as elements in the finite field $\mathbb{F}_{2^u}$. The length of the vectors is equal in all transmissions and all links are assumed to be synchronized with respect to the symbol timing. In this paper we consider linear coding[1]. For a linear code, the signal $Y(j)$ on a link j is a linear combination of processes X_i generated at node $v = \text{tail}(j)$ and signals $Y(l)$ on incident incoming links l. For the delay-free case, this is represented by the equation

$$Y(j) = \sum_{\{i\,:\,X_i \text{ generated at } v\}} a_{i,j} X_i + \sum_{\{l\ :\ \text{head}(l)\,=\,v\}} f_{l,j} Y(l)$$

and an output process $Z(\beta, i)$ at receiver node β is a linear combination of signals on its terminal links, represented as

$$Z(\beta, i) = \sum_{\{l\,:\,\text{head}(l)=\beta\}} b_{\beta_{i,l}} Y(l)$$

For multicast on a network with link delays, memory is needed at the receiver nodes, but memoryless operation suffices at all other nodes [5]. We consider unit delay links, modeling links with longer delay as links in series. The corresponding linear coding equations are

$$Y_{t+1}(j) \;=\; \sum_{\{i\,:\,X_i \text{ generated at } v\}} a_{i,j} X_{it} + \sum_{\{l\ :\ \text{head}(l)\,=\,v\}} f_{l,j} Y_t(l) \tag{1}$$

$$Z_{t+1}(\beta, i) \;=\; \sum_{\{l\,:\,\text{head}(l)=\beta\}} \sum_{u=t-\mu}^{t} b_{\beta_{i,l_{t-u}}} Y_u(l) \tag{2}$$

where μ represents the memory required.

The coefficients $\{a_{i,j}, f_{l,j}, b_{\beta_{i,l}} \in \mathbb{F}_{2^u}\}$ can be collected into $r \times \nu$ matrices $A = (a_{i,j})$ and $B_\beta = (b_{\beta_{i,j}})$, and the $\nu \times \nu$ matrix $F = (f_{l,j})$, whose structure is constrained by the network. For acyclic graphs, we number the links ancestrally, i.e. lower-numbered links upstream of higher-numbered links, so matrix F is upper triangular with zeros on the diagonal. A triple (A, F, B), where

$$B = \left[\begin{array}{c} B_1 \\ \hline \vdots \\ \hline B_d \end{array} \right]$$

specifies the behavior of the network, and represents a *linear network code*. We use the following notation:

---

[1]which is sufficient for multicast [6]

$Y(3) = a_{1,3}X_1 + f_{1,3}Y(1)$
$\qquad\quad + f_{2,3}Y(2)$

Figure 1: Illustration of linear coding at a node.

- $G = \begin{cases} (I - F)^{-1} \text{ in the acyclic delay-free case}^2 \\ (I - DF)^{-1} \text{ in the case with delay}^3 \end{cases}$

- $G_{\mathcal{H}}$ is the submatrix consisting of columns of G corresponding to links in set $\mathcal{H}$

- $\underline{a}_j$, $\underline{c}_j$ and $\underline{b}_j$ denote column j of A, AG and B respectively

Matrix AG gives the transfer matrix from input processes to signals on each link; the connection problem is feasible if and only if AGB_β^T has full rank for each receiver β [5].

3 Main Results

We generalize results presented in [2, 3] to arbitrary graphs which may have cycles and delay.

Theorem 1 *For a feasible multicast connection problem on a (possibly cyclic) network with unit delay links, independent or linearly correlated sources, and a network code in which some or all code coefficients are chosen independently and uniformly over all elements of a finite field $\mathbb{F}_q$ (some coefficients can take fixed values as long as these values preserve feasibility[4]), the probability that all the receivers can decode the source processes is at least $(1 - d/q)^\eta$ for $q > d$, where d is the number of receivers and η is the number of links carrying random combinations of source processes and/or incoming signals.*

The proof of the above theorem relies on the following result, which allows us to easily characterize the dependence of the transfer matrix determinant on the random coefficients.

Theorem 2 *For an arbitrary (possibly cyclic) network with unit delay links, the transfer matrix $A(I - DF)^{-1}B_\beta^T$ for receiver β in a network code (A, F, B) is nonsingular if and only if the corresponding Edmonds matrix $\begin{bmatrix} A & 0 \\ I - DF & B_\beta^T \end{bmatrix}$ is nonsingular.*

The bound in Theorem 1 is a very general one, applying across all networks with the same number of receivers and the same number of links with independently chosen random linear mappings. Our next goal is to find tighter bounds by taking into account more specific network characteristics. To this end, we establish a connection between randomized coding success probability and network connection feasibility when links are

---

[2]The inverse exists since F is nilpotent.

[3]The inverse exists since the determinant is a nonzero polynomial in D.

[4]i.e. the result holds for networks where not all nodes perform random coding, or where signals add by superposition on some channels

unreliable, for acyclic networks with or without link delays. This is useful for cases where analysis of connection feasibility is easier than direct analysis of randomized coding, for example in the case of networks with unreliable links and excess capacity.

Theorem 3 *For a d-receiver multicast problem on an acyclic network, the success probability of a random network code in the field of size q is greater than or equal to the probability that the network connections remain feasible after deleting each link of the original graph with probability d/q.*

Theorem 4 *For a connection problem with r sources and links with failure probability p, let y be the minimum redundancy, i.e. deletion of any y links in the network preserves feasibility. A lower bound on the probability that a particular receiver receives all processes is*

$$\sum_{x=r}^{r+y} \binom{r+y}{x} \left(1 - p - \frac{1-p}{q}\right)^{Lx} \left(1 - \left(1 - p - \frac{1-p}{q}\right)^{L}\right)^{r+y-x}$$

where L is the longest source-receiver path in the network.

4 Proofs and Ancillary Results

4.1 Randomized Network Coding on Arbitrary Graphs with Delay

Proof of Theorem 2: Note that

$$\begin{bmatrix} I & -A(I-DF)^{-1} \\ 0 & I \end{bmatrix} \begin{bmatrix} A & 0 \\ I-DF & B_\beta^T \end{bmatrix} = \begin{bmatrix} 0 & -A(I-DF)^{-1}B_\beta^T \\ I-DF & B_\beta^T \end{bmatrix}$$

The first matrix, $\begin{bmatrix} I & -A(I-DF)^{-1} \\ 0 & I \end{bmatrix}$, has determinant 1. So $\det\left(\begin{bmatrix} A & 0 \\ I-DF & B_\beta^T \end{bmatrix}\right)$ equals $\det\left(\begin{bmatrix} 0 & -A(I-DF)^{-1}B_\beta^T \\ I-DF & B_\beta^T \end{bmatrix}\right)$, which can be expanded as follows:

$$\begin{aligned} &\det\left(\begin{bmatrix} 0 & -A(I-DF)^{-1}B_\beta^T \\ I-DF & B_\beta^T \end{bmatrix}\right) \\ =\ & (-1)^{r\nu}\det\left(\begin{bmatrix} -A(I-DF)^{-1}B_\beta^T & 0 \\ B_\beta & I-DF \end{bmatrix}\right) \\ =\ & (-1)^{r\nu}\det(-A(I-DF)^{-1}B_\beta^T)\det(I-DF) \\ =\ & (-1)^{r(\nu+1)}\det(A(I-DF)^{-1}B_\beta^T)\det(I-DF) \end{aligned}$$

Since $\det(I-DF)$ is nonzero, the result follows. ∎

Lemma 1 *The determinant polynomial of the Edmonds matrix $\begin{bmatrix} A & 0 \\ I-DF & B_\beta^T \end{bmatrix}$ associated with a network code (A, F, B) in a network with delay is a polynomial in delay variable D, whose coefficients have maximum degree ν in variables $\{a_{x,j}, f_{i,j}\}$, and are linear in each variable $\{a_{x,j}, f_{i,j}\}$.*

Proof: Each variable $\{a_{x,j}, f_{i,j}, b_{x,j}\}$ appears in only one entry of the Edmonds matrix. The determinant can be written as the sum of products of $r + \nu$ entries, one from each row and column. Each such product is linear in each variable $\{a_{x,j}, f_{i,j}, b_{x,j}\}$, has degree at most $r + \nu$ in variables $\{a_{x,j}, f_{i,j}, b_{x,j}\}$, and has degree r in variables $\{b_{x,j}\}$. ∎

Lemma 2 *Let P be a polynomial in $\mathbb{F}[\xi_1, \xi_2, \ldots]$ of degree less than or equal to $d\eta$, in which the largest exponent of any variable ξ_i is at most d. Values for $\xi_1, \xi_2, \ldots$ are chosen independently and uniformly at random from $\mathbb{F}_q \subseteq \mathbb{F}$. The probability that P equals zero is at most $1 - (1 - d/q)^\eta$ for $d < q$.*

Proof: For any variable ξ_1 in P, let d_1 be the largest exponent of ξ_1 in P. Express P in the form $P = \xi_1^{d_1} P_1 + R_1$, where P_1 is a polynomial of degree at most $d\eta - d_1$ that does not contain variable ξ_1, and R_1 is a polynomial in which the largest exponent of ξ_1 is less than d_1. By the Principle of Deferred Decisions, the probability $\Pr[P = 0]$ is unaffected if we set the value of ξ_1 last after all the other coefficients have been set. If, for some choice of the other coefficients, $P_1 \neq 0$, then P becomes a polynomial in $\mathbb{F}[\xi_1]$ of degree d_1. By the Schwartz-Zippel Theorem, this probability $\Pr[P = 0 | P_1 \neq 0]$ is upper bounded by d_1/q. So

$$
\begin{aligned}
\Pr[P = 0] &\leq \Pr[P_1 \neq 0]\frac{d_1}{q} + \Pr[P_1 = 0] \\
&= \Pr[P_1 = 0]\left(1 - \frac{d_1}{q}\right) + \frac{d_1}{q}
\end{aligned}
\tag{3}
$$

Next we consider $\Pr[P_1 = 0]$, choosing any variable ξ_2 in P_1 and letting d_2 be the largest exponent of ξ_2 in P_1. We express P_1 in the form $P_1 = \xi_2^{d_2} P_2 + R_2$, where P_2 is a polynomial of degree at most $d\eta - d_1 - d_2$ that does not contain variables ξ_1 or ξ_2, and R_2 is a polynomial in which the largest exponent of ξ_2 is less than d_2. Proceeding similarly, we assign variables ξ_i and define d_i and P_i for $i = 3, 4, \ldots$ until we reach $i = k$ where P_k is a constant and $\Pr[P_k = 0] = 0$. Note that $1 \leq d_i \leq d < q \ \forall \ i$ and $\sum_{i=1}^{k} d_i \leq d\eta$, so $k \leq d\eta$. Applying Schwartz-Zippel as before, we have for $k' = 1, 2, \ldots, k$

$$
Pr[P_{k'} = 0] \leq Pr[P_{k'+1} = 0]\left(1 - \frac{d_{k'+1}}{q}\right) + \frac{d_{k'+1}}{q}
\tag{4}
$$

Combining all the inequalities recursively, we can show by induction that

$$
\Pr[P = 0] \leq \frac{\sum_{i=1}^{k} d_i}{q} - \frac{\sum_{i \neq j} d_i d_j}{q^2} + \ldots + (-1)^{k-1}\frac{\prod_{i=1}^{k} d_i}{q^k}
$$

where $0 \leq d\eta - \sum_{i=1}^{k} d_i$.

Now consider the integer optimization problem

$$
\begin{aligned}
\text{Maximize} \quad & f = \frac{\sum_{i=1}^{d\eta} d_i}{q} - \frac{\sum_{i \neq j} d_i d_j}{q^2} + \ldots + (-1)^{d\eta-1}\frac{\prod_{i=1}^{d\eta} d_i}{q^{d\eta}} \\
\text{subject to} \quad & 0 \leq d_i \leq d < q \ \forall \ i \in [1, d\eta], \\
& \sum_{i=1}^{d\eta} d_i \leq d\eta, \quad \text{and} \quad d_i \text{ integer}
\end{aligned}
\tag{5}
$$

whose maximum is an upper bound on $\Pr[P = 0]$.

We first consider the non-integer relaxation of the problem. Let $\underline{d}^* = \{d_1^*, \ldots, d_{d\eta}^*\}$ be an optimal solution.

For any set S_h of h distinct integers from $[1, d\eta]$, let $f_{S_h} = 1 - \frac{\sum_{i \in S_h} d_i}{q} + \frac{\sum_{i,j \in S_h, i \neq j} d_i d_j}{q^2} - \ldots + (-1)^h \frac{\prod_{i \in S_h} d_i}{q^h}$. We can show by induction on h that $0 < f_{S_h} < 1$ for any set S_h of h distinct integers in $[1, d\eta]$.

If $\sum_{i=1}^{d\eta} d_i^* < d\eta$, then there is some $d_i^* < d$, and there exists a feasible solution $\underline{d}$ such that $d_i = d_i^* + \epsilon$, $\epsilon > 0$, and $d_h = d_h^*$ for $h \neq i$, which satisfies

$$f(\underline{d}) - f(\underline{d}^*) \quad = \quad \frac{\epsilon}{q} \left(1 - \frac{\sum_{h \neq i} d_h^*}{q} + \ldots + (-1)^{d\eta - 1} \frac{\prod_{h \neq i} d_h^*}{q^{d\eta - 1}} \right)$$

This is positive, contradicting the optimality of $\underline{d}^*$.

Next suppose $0 < d_i^* < d$ for some d_i^*. Then there exists some d_j^* such that $0 < d_j^* < d$, since if $d_j^* = 0$ or d for all other j, then $\sum_{i=1}^{d\eta} d_i^* \neq d\eta$. Assume without loss of generality that $0 < d_i^* \leq d_j^* < d$. Then there exists a feasible vector $\underline{d}$ such that $d_i = d_i^* - \epsilon$, $d_j = d_j^* + \epsilon$, $\epsilon > 0$, and $d_h = d_h^* \ \forall \ h \neq i, j$, which satisfies

$$f(\underline{d}) - f(\underline{d}^*) = - \left(\frac{(d_i^* - d_j^*)\epsilon - \epsilon^2}{q^2} \right) \left(1 - \frac{\sum_{h \neq i,j} d_h^*}{q} - \ldots + (-1)^{d\eta - 2} \frac{\prod_{h \neq i,j} d_h^*}{q^{d\eta - 2}} \right)$$

This is again positive, contradicting the optimality of $\underline{d}^*$.

Thus, $\sum_{i=1}^{d\eta} d_i^* = d\eta$, and $d_i^* = 0$ or d. So exactly η of the variables d_i^* are equal to d. Since the optimal solution is an integer solution, it is also optimal for the integer program (5). The corresponding optimal $f = \eta \frac{d}{q} - \binom{\eta}{2} \frac{d^2}{q^2} + \ldots + (-1)^{\eta - 1} \frac{d^\eta}{q^\eta} = 1 - \left(1 - \frac{d}{q} \right)^\eta$. ∎

Proof of Theorem 1: To check if a network code (A, F, B) transmits all source processes to receiver β, it suffices to check that the determinant of the corresponding Edmonds matrix is nonzero (Theorem 2). This determinant, which we denote by P_β, is a polynomial in delay variable D, whose coefficients are linear in each variable $\{a_{x,j}, f_{i,j}\}$ and have degree at most ν in these variables (Lemma 1). Each column corresponds to a link in the network; the number of columns containing variable terms equals η, the number of links carrying random combinations of source processes and/or incoming signals. The product $\prod_\beta P_\beta$ for d receivers is, accordingly, a polynomial in delay variable D, whose coefficients are polynomials in $\{a_{x,j}, f_{i,j}\}$ of degree at most $d\eta$, and in which the largest exponent of each of these variables is at most d. These properties still hold if some variables are set to deterministic values which do not make the product identically zero.

Linearly correlated sources can be viewed as pre-specified linear combinations of underlying independent processes. Unlike the independent sources case where each nonzero entry of the A matrix can be set independently, in this case there are linear dependencies among the entries. The columns $\underline{a}_j$ of the A matrix are linear functions $\underline{a}_j = \sum_k \alpha_j^k \underline{v}_j^k$ of column vectors $\underline{v}_j^k$ that represent the composition of the source processes at tail(j) in terms of the underlying independent processes. Variables α_j^k in column $\underline{a}_j$ can be set independently of variables $\alpha_{j'}^k$ in other columns $\underline{a}_{j'}$. It can be seen from Lemma 1 that for any particular j, each product term in the polynomial P_β for any receiver β contains at most one variable $a_{i,j} = \sum_k \alpha_j^k v_{i,j}^k$. P_β is thus linear in the variables α_j^k, and also in variables $f_{i,j}$, which are unaffected by the source correlations. So any variable in the product of d such polynomials has maximum exponent d.

Applying Lemma 2 gives us the required bound.

For the single-receiver case, the bound is attained for a network consisting only of links forming a single set of r disjoint source-receiver paths. ∎

4.2   Connections with Link Reliability

Proof of Theorem 3: Consider any link j, and a set $\mathcal{S}$ of d' arbitrary $(r \times r - 1)$ rank-$(r-1)$ matrices in $(\mathbb{F}_q(D))^{r \times r-1}$, such that, for each matrix in $\mathcal{S}$, link j has among its inputs a signal whose associated vector is not in the column space of the matrix. Let $\underline{v}_i \in (\mathbb{F}_q(D))^r$ be the vector associated with the i^{th} input to link j. Let $Y(j) = \sum_i D f_i \underline{v}_i$ be the vector associated with link j.

Each entry of $Y(j)$ is a polynomial in $\mathbb{F}_q(D, f_1, f_2, \ldots)$ that is linear in coefficients f_i. The determinant of an $r \times r$ matrix which has $Y(j)$ as one of its columns, and whose $r-1$ other columns are independent of coefficients f_i, is thus linear in coefficients f_i. The product of d' such determinants has maximum degree d' in coefficients f_i.

If coefficients f_i are chosen uniformly and independently from $\mathbb{F}_q$, by the Schwartz-Zippel Theorem, this product is nonzero with probability at least $1 - d'/q$. Denoting by $E_{\mathcal{S},j}$ the event that adding $Y(j)$ as an additional column to each of the matrices in $\mathcal{S}$ gives a full rank matrix, we have $\Pr(E_{\mathcal{S},j}) \geq 1 - d'/q$.

Next consider a number of sets $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_n$ each consisting of d' arbitrary $(r \times r - 1)$ rank-$(r-1)$ matrices, such that for each matrix in $\mathcal{S}_k$, $1 \geq k \geq n$, link j has among its inputs a signal whose associated vector is not in the column space of the matrix. Then $\Pr(\bigcup_{k=1}^n E_{\mathcal{S}_k,j}) \geq 1 - d'/q$.

Each receiver receives all processes successfully if the submatrix of AG corresponding to r of its incident incoming links, or terminal links, has full rank. The connection problem is feasible if and only if each receiver has a set of r link-disjoint paths, one from each source.

Let j be the highest-indexed link in an ancestral ordering, where lower-indexed links feed into higher-indexed links. Consider any given signals on all other links. There are three cases:

Case 1: Regardless of the code coefficients for j, there cannot exist full rank sets of r terminal links for each receiver.

Case 2: Regardless of the code coefficients for j, each receiver has a full rank set of r terminal links.

Case 3: For some choice of code coefficients for link j, each receiver has a full rank set of r terminal links, i.e. link j has among its inputs signals whose associated vectors are not in the column space of the submatrices of AG corresponding to the terminal links of one or more other receivers. By our earlier arguments, such a choice is made with probability at least $1 - d'/q$, where d' is the number of receivers downstream of link j.

In all three cases, the probability that each receiver has a set of r terminal links with a full rank set of inputs when code coefficients for link j are chosen randomly is greater than or equal to that in the case where link j is deleted with probability $d/q \geq d'/q$.

We next consider the problem where link j is deleted with probability d/q, and random code coefficients are chosen for all other links. From our earlier arguments, the probability that any set of r undeleted paths to each receiver has a full rank set of inputs is less than or equal to the probability of success in the original network coding problem.

We continue in this fashion, at each stage considering a new problem in which we delete with probability d/q the next highest-indexed link as well as each previously con-

sidered link. Random code coefficients are chosen for all other links. At each stage, for any choice of surviving links among the set of randomly deleted links, the problem is either infeasible, or there exist one or more sets of random coding links incident to undeleted paths to each receiver which, if full rank, preserve feasibility of the problem. The probability that any set of r undeleted paths to each receiver has a full rank set of inputs is less than or equal to the probability of success in the original network coding problem. ∎

Proof of Theorem 4: For a given network of non-failed links, we can find a lower bound by considering the more general case where a source process can be available at one or more source node locations, and by analyzing the probability that the connections remain feasible when links fail with probability $1/q$, which by Theorem 3 gives us a lower bound on network coding success probability. The success probability for a network whose links fail with probability p is thus lower bounded by the probability that the connections remain feasible when links fail with probability $1 - (1 - p)(1 - 1/q)$.

We show by induction on y that a network consisting of $r + y$ disjoint source-receiver paths, any r of which can transmit all processes, has a success probability that is less than or equal to that for any y-redundant network.

Consider a network $\mathcal{G}_1$ consisting of $r + y$ disjoint source-receiver paths any r of which can transmit all processes. Let $\mathcal{G}_2$ be any other y-redundant network.

For $i = 1, 2$, we consider a set $\mathcal{P}_i$ of links forming r disjoint paths from each source to the receiver on graph $\mathcal{G}_i$. We distinguish two cases:

Case 1: None of the links in $\mathcal{P}_i$ fail. In this case the connections are feasible.

Case 2: There exists some link $j_i \in \mathcal{P}_i$ that fails.

The probability of either case occurring is the same for $i = 1, 2$. Since

$$\Pr(\text{success}) = \Pr(\text{case 1}) + \Pr(\text{case 2}) \Pr(\text{success}|\text{case 2})$$

$\Pr(\text{success}|i = 1) \leq \Pr(\text{success}|i = 2)$ iff $\Pr(\text{success}|\text{case 2}, i = 1) \leq \Pr(\text{success}|\text{case 2}, i = 2)$.

For $y = 0$, the hypothesis is true since $\Pr(\text{success}|\text{case 2}) = 0$ for $i = 1, 2$. For $y > 0$, in case 2 we can remove link j_i leaving a $(y - 1)$-redundant graph $\mathcal{G}'_i$. By the induction hypothesis, the probability of success for $\mathcal{G}'_1$ is less than or equal to that for $\mathcal{G}'_2$.

Thus, $\mathcal{G}_1$ gives a lower bound on success probability, which is the probability that all links on at least r of $r + y$ length-L paths do not fail. The result follows from observing that each path does not fail with probability $\left((1 - p)(1 - \frac{1}{q}) \right)^L$. ∎

5   Conclusion

We have presented bounds for the success probability of distributed randomized network coding for multi-source multicast in networks. The first is a very general bound for arbitrary networks, which may have cycles or delay, in terms of the number of receivers and the number of links with independently chosen linear mappings. We have also shown an approach for obtaining tighter results for more specific networks. For any given acyclic network, we can bound randomized coding success probability by the probability of connection feasibility in a related network problem with unreliable links. From this we obtain a success probability bound for randomized network coding in networks with unreliable links and excess capacity, in terms of link failure probability and amount of redundancy.

Further work includes extensions to different applications, such as non-multicast. It would also be of interest to consider various protocols for different communication scenarios and evaluate the associated overhead, comparing this with traditional routing based approaches.

Acknowledgments

References

[1] R. Ahlswede, N. Cai, S.-Y.R. Li and R.W. Yeung, "Network Information Flow", IEEE-IT, vol. 46, pp. 1204-1216, 2000.

[2] T. Ho, D. R. Karger, M. Médard and R. Koetter, "Network Coding from a Network Flow Perspective ", Proceedings of the 2003 IEEE International Symposium on Information Theory.

[3] T. Ho, R. Koetter, M. Médard, D. R. Karger and M. Effros, "The Benefits of Coding over Routing in a Randomized Setting", Proceedings of the 2003 IEEE International Symposium on Information Theory.

[4] S. Jaggi, P.A. Chou and K. Jain, "Low Complexity Algebraic Network Codes", Proceedings of the 2003 IEEE International Symposium on Information Theory.

[5] R. Koetter and M. Médard, "An Algebraic Approach to Network Coding", IEEE/ACM Transactions on Networking, to appear.

[6] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding", IEEE Transactions on Information Theory, vol. 49, pp. 371-381, 2003.

[7] P. Sanders, S. Egner, and L. Tolhuizen, "Polynomial Time Algorithms For Network Information Flow", In 15th ACM Symposium on Parallel Algorithms and Architectures, pages 286-294, 2003.