

Random Linear Network Coding: A free cipher?

Luísa Lima Muriel Médard João Barros

Abstract—We consider the level of information security provided by random linear network coding in network scenarios in which all nodes comply with the communication protocols yet are assumed to be potential eavesdroppers (i.e. “nice but curious”). For this setup, which differs from wiretapping scenarios considered previously, we develop a natural *algebraic security* criterion, and prove several of its key properties. A preliminary analysis of the impact of network topology on the overall network coding security, in particular for complete directed acyclic graphs, is also included.

Index Terms—security, information theory, graph theory, network coding.

I. INTRODUCTION

Under the classical networking paradigm, in which intermediate nodes are only allowed to store and forward packets, information security is usually viewed as an independent feature with little or no relation to other communication tasks. In fact, since intermediate nodes receive exact copies of the sent packets, data confidentiality is commonly ensured by cryptographic means at higher layers of the protocol stack. Breaking with the ruling paradigm, network coding allows intermediate nodes to mix information from different data flows [1], [2] and thus provides an intrinsic level of data security — arguably one of the least well understood benefits of network coding.

Previous work on this issue has been mostly concerned with constructing codes capable of splitting the data among different links, such that reconstruction by a wiretapper is either very difficult or impossible. In [3], the authors present a secure linear network code that achieves perfect secrecy against an attacker with access to a limited number of links. A similar problem is considered in [4], featuring a random coding approach in which only the input vector is modified. [5] introduces a different information-theoretic security model, in which a system is deemed to be secure if an eavesdropper is unable to get any decoded or decodable (also called *meaningful*) source data. Still focusing on wiretapping attacks, [6] provides a simple security protocol exploiting the network topology: an attacker is shown to be unable to get any

L. Lima (luisalima@ieee.org) and J. Barros (barros@dcc.fc.up.pt) are with the Instituto de Telecomunicações (IT) and the Department of Computer Science, Faculdade de Ciências da Universidade do Porto, Portugal. M. Médard (medard@mit.edu) is with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. This work was partly supported by the Fundação para a Ciência e Tecnologia (Portuguese Foundation for Science and Technology) under grant SFRH/BD/24718/2005 and by AFOSR under grant “Robust Self-Authenticating Network Coding” AFOSR 000106. Part of this work was done while the first author was a visiting student at the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology.




Fig. 1. Canonical Network Coding Example. In this image, intermediate nodes are represented with squares. With this code, node 4 is a vulnerability for the network since it can decode all the information sent through it. Note that the complete opposite happens for node 5, that receives no meaningful information whatsoever.

meaningful information unless it can access those links that are necessary for the communication between the legitimate sender and the receiver, who are assumed to be using network coding. As a distributed capacity-achieving approach for the multicast case, randomized network coding [7], [8] has been shown to extend naturally to packet networks with losses [9] and Byzantine modifications (both detection and correction [10], [11], [12], [13]). [14] adds a cost criterion to the secure network coding problem, providing heuristic solutions for a coding scheme that minimizes both the network cost and the probability that the wiretapper is able to retrieve all the messages of interest.

In this work, we approach network coding security from a different angle: our focus is *not* on the threat posed by external wiretappers but on the more general threat posed by intermediate nodes. We assume that the network consists entirely of “nice but curious” nodes, i.e. they comply with the communication protocols (in that sense, they are well-behaved) but may try to acquire as much information as possible from the data that passes through them (in which case, they are potentially malicious). This notion is highlighted in the following example.

Example 1: Consider the canonical network coding example with 7 nodes, shown in *Figure 1*. Node 1 sends a flow to sinks 6 and 7 through intermediate nodes 2, 3, 4 and 5. From the point of security, we can distinguish between three types of intermediate nodes in this setting: (1) those that only get a non-meaningful part of the information, such as node 5; (2) those that obtain all of the information, such as node 4; and (3) those that get partial yet meaningful information, such as nodes 2 and 3. Although this network code could be considered *secure* against single-edge external wiretapping — i.e., the wiretapper is not able to retrieve the whole data simply

by eavesdropping on a single edge — it is clearly insecure against internal eavesdropping by an intermediate node.

Motivated by this example, we set out to investigate the security potential of network coding. Our main contributions are as follows:

- *Problem Formulation:* We formulate a secure network coding problem, in which all intermediate nodes are viewed as potential eavesdroppers and the goal is to characterize the intrinsic level of security provided by random linear network coding.
- *Algebraic Security Criterion:* Based on the notion that the number of decodable bits available to each intermediate node is limited by the degrees of freedom it receives, we are able to provide a natural secrecy constraint for network coding and to prove some of its most fundamental properties.
- *Security Analysis for Complete Directed Acyclic Graphs:* As a preliminary step towards understanding the interplay between network topology and security against eavesdropping nodes, we present a rigorous characterization of the achievable level of algebraic security for this class of complete graphs.

The remainder of this paper is organized as follows. First, a formal problem statement is in *Section II*, followed by a detailed analysis of the algebraic security of Randomized Linear Network Coding in *Section III*. In *Section IV*, this analysis is carried out specifically for complete directed acyclic graphs. The paper concludes with *Section V*.

II. PROBLEM SETUP

We adopt the network model of [2]: we represent the network as an acyclic directed graph $G = (V, E)$, where V is the set of nodes and E is the set of edges. Edges are denoted by round brackets $e = (v, v') \in E$, in which $v = \text{head}(e)$ and $v' = \text{tail}(e)$. The set of edges that end at a vertex $v \in V$ is denoted by $\Gamma_I(v) = \{e \in E : \text{head}(e) = v\}$, and the in-degree of the vertex is $\delta_I(v) = |\Gamma_I(v)|$; similarly, the set of edges originating at a vertex $v \in V$ is denoted by $\Gamma_O(v) = \{e \in E : \text{tail}(e) = v\}$, the out-degree being represented by $\delta_O(v) = |\Gamma_O(v)|$.

Discrete random processes X_1, \dots, X_K are observable at one or more source nodes. To simplify the analysis, we shall consider that each network link is free of delays and that there are no losses. Moreover, 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. We shall consider multicast connections as it is the most general type of single connection; there are $d \geq 1$ receiver nodes. The objective is to transmit all the source processes to each of the receiver nodes.

In linear network coding, edge $e = (v, u)$ carries the process $Y(e)$, which is defined below:

$$Y(e) = \sum_{l: X_l \text{ generated at } v} \alpha_{l,e} X(v, l) + \sum_{e': \text{head}(e') = \text{tail}(e)} \beta_{e',e} Y(e')$$

The *transfer matrix* M describes the relationship between an input vector \underline{x} and an output vector \underline{z} , $\underline{z} = \underline{x}M$; $M = A(I - F)^{-1}B^T$, where A and B represent, respectively, the linear mixings of the input vector and of the output vector, and have sizes $K \times |E|$ and $\nu \times |E|$. F is the adjacency matrix of the directed labelled line graph corresponding to the graph G . In this paper we shall not consider matrix B , which only refers to the decoding at the receivers. Thus, we shall mainly analyse parts of the matrix AG , such that $G = (I - F)^{-1}$; \underline{a}_i and \underline{c}_i denote column i of A and AG , respectively. We define the *partial transfer matrix* $M'_{\Gamma_I(v)}$ (also called *auxiliary encoding vector* [9]) as the observable matrix at a given node v , i.e. the observed matrix formed by the symbols received at a node v . This is equivalent to the fraction of the data that an intermediate node has access to in a multicast transmission.

Regarding the coding scheme, we consider the random linear network coding scheme introduced in [7]: and thus each coefficient of the matrices described above is chosen independently and uniformly over all elements of a finite field \mathbb{F}_q , $q = 2^m$.

Our goal is to evaluate the *intrinsic security* of random linear network coding, in multicast scenarios where all the intermediate nodes in the network are potentially malicious eavesdroppers. Specifically our threat model assumes that intermediate nodes perform the coding operations as outlined above, and will try to decode as much data as possible.

III. ALGEBRAIC SECURITY OF RANDOM LINEAR NETWORK CODING

A. Algebraic security

The Shannon criterion for information-theoretic security [15] corresponds in general terms to a zero mutual information between the cypher-text (C) and the original message (M), i.e. $I(M; C) = 0$. This condition implies that an attacker must guess $\leq H(M)$ symbols to be able to compromise the data. With network coding, on the other hand, if the attacker is capable of guessing M symbols, $K - M$ additional observed symbols are required for decoding — by noting that each received symbol is a linear combination of the K message symbols from the source, we can see that a receiver must receive K coded symbols in order to recover one message symbol. Thus, as will be shown later, restricted rank sets of individual symbols do not translate into immediately decodable data with high probability. This notion is illustrated in *Figure 2*. In the scheme shown on top, each intermediate node can recover half of the transmitted symbols, whereas in the bottom scheme none of the nodes can recover any portion of the sent data.

Definition 1 (Algebraic Security Criterion): The level of security provided by random linear network coding is measured by the number of symbols that an intermediate node v has to guess in order to decode *one* of the transmitted symbols. From a formal point of view,

$$\Delta_S(v) = \frac{K - (\text{rank}(M'_{\Gamma_I(v)}) + l_d)}{K},$$




Fig. 2. Example of algebraic security. In the upper scheme data is not protected, whereas in the lower scheme nodes 2 and 3 are unable to recover any data symbols.

where l_d represents the number of partially diagonalizable lines of the matrix (i.e. the number of message symbols that can be recovered by Gaussian elimination).

Notice that the previous definition is equivalent to computing the difference between the global rank of the code and the local rank in each intermediate node v . Moreover, as more and more symbols become compromised of security criteria, the level of security tends to 0, since as we shall show in this section, with high probability the number of individually decodable symbols l_d goes to zero as the size of the field goes to infinity.

B. Security Characterization

We are now ready to solve the problem of characterizing the algebraic security of random linear network coding. The key to our proofs is to analyze the properties of the partial transfer matrix at each intermediate node. Recall that there are two cases in which the intermediate node can gain access to relevant information: (1) when the partial transfer matrix has full rank and (2) when the partial transfer matrix has diagonalizable parts. Thus, we shall carry out independent analyzes in terms of rank and in terms of partially diagonalizable matrices. The following lemmas will be useful.

Lemma 1: In the random linear network coding scheme,

$$P(\Delta_S > 0) \leq P(\exists v : \delta_I(v) > K).$$

Proof: See the Appendix. ■

It follows from this lemma that it is only necessary to consider the case in which $K \leq \delta_I(v)$.

Lemma 2: The probability that a linear combination of independent and uniformly distributed values in \mathbb{F}_q yields the zero result is bounded by

$$P(X_{lin} = 0) \leq \frac{2q + h(q)}{q^2},$$

where $h(q)$ is a function such that $O(h(q)) < O(q^2)$. Moreover, $P(X_{lin} = 0)$ tends to 0 when $q \rightarrow \infty$.

Proof: See the Appendix. ■

Lemma 3: The probability of obtaining y zeros in one line of the $\xi \times \xi$ transfer matrix M is bounded by

$$P(Y = y) \leq \binom{\xi}{\xi - y} \left(\frac{2q + h(q)}{q^2} \right)^y \left(1 - \frac{2q + h(q)}{q^2} \right)^{\xi - y}.$$

Proof: See the Appendix. ■

Theorem 1: Let $P(l_d > 0)$ be the probability of recovering a strictly positive number of symbols l_d at the intermediate nodes with $\delta_I(v) \leq K - 1$ by Gaussian elimination. Then, $P(l_d > 0) \rightarrow 0$ with $q \rightarrow \infty$ and $K \rightarrow \infty$.

Proof: Let M' be the transpose of the partial transfer matrix at some vertex v , $M' = M_{\Gamma_I(v)}^T$. We consider the process of Gaussian elimination of M' . It is unnecessary to consider rank K , since in that case the matrix, w.h.p. is invertible and hence diagonalizable [8]. Thus, M' is a $\delta_I(v) \times K$ matrix, $\delta_I(v) < K$.

We prove the theorem constructively by analysing the probability of having $K - 1$ zeros in one or more lines of M' . Let p be the probability of having $K - 1$ zeros in a line of M' , and let X be a random variable representing the recoverable number of symbols when an intermediate node has $\delta_I(v)$ degrees of freedom. It follows from Lemma 3 that

$$p = \binom{K}{K-1} \left(\frac{2q + h(q)}{q^2} \right)^1 \left(1 - \frac{2q + h(q)}{q^2} \right)^{K-1}.$$

In the base case with $\delta_I(v) = 1$, at most $X = 1$ symbols can be recovered, since there are not enough degrees of freedom to perform Gaussian elimination and the only chance for recovering a symbol is that the line of the matrix M already has $K - 1$ zeros. The probability for this is p .

In the case that $1 < \delta_I(v) < K$, we can obtain directly a number $L = l$ of lines with $K - 1$ zeros, and a number $\delta_I(v) - l$ of lines in the opposite situation. Since we have $\delta_I(v)$ degrees of freedom to perform Gaussian elimination, we can obtain at most $\delta_I(v)$ symbols by successive elimination. At each step the probability of obtaining a line with $K - 1$ zeros is bounded by p .

By analysing the different possibilities of combinations for the lines that already have $K - 1$ zeros and the ones that can be obtained by Gaussian elimination, we get

$$P(X = x) \leq \sum_{l=0}^x \binom{\delta_I(v)}{l} p^l (1-p)^{\delta_I(v)-l} P_l(X = x)$$

$$P_l(X = x) \leq \binom{\delta_I(v) - l}{x - (\delta_I(v) - l)} p^{x - \delta_I(v) + l} (1-p)^{2\delta_I(v) - 2l - x},$$

where $P_l(X = x)$ represents $P(X = x | L = l)$.

Approximating the binomial distribution by a normal distribution yields

$$P_l(X = x) \approx \frac{e'}{\sqrt{2\pi(\delta_I(v) - l)p(1-p)}},$$

where

$$e' = \exp\left(-\frac{1}{2} \frac{(x - (\delta_I(v) - l)p)^2}{(\delta_I(v) - l)p(1-p)}\right)$$

Since $p \rightarrow p^* < 1$, we can state that, when $q \rightarrow \infty$ and $p \rightarrow 0$ is $\approx \exp(x^2)$. When K goes to ∞ , so does x , and hence

$$\exp(x^2)_{x \rightarrow \infty} \rightarrow 0,$$

and

$$P_l(X = K - 1)_{q \rightarrow \infty, K \rightarrow \infty} \rightarrow 0.$$

Since

$$P(X = K-1) = \sum_{l=0}^{K-1} \binom{\delta_I(v)}{l} p^l (1-p)^{\delta_I(v)-l} P_l(X = K-1),$$

and $P_l(X = K-1)$ decreases exponentially, and l only increases linearly,

$$P(X = K-1)_{q \rightarrow 0, K \rightarrow \infty} \rightarrow 0.$$

The probability of obtaining $X < K-1$ symbols is bounded by $P(X = K-1)$; it follows that the probability of decoding X symbols with any $\delta_I(v) < K$ goes to zero as q and K tend to infinity. ■

IV. ALGEBRAIC SECURITY OF THE COMPLETE GRAPH

Notice that, in consequence of the property outlined in Lemma 1, the algebraic security of a graph is topology dependent. A node with $\delta_I(v) \geq K$ will not necessarily receive a full-rank partial transfer matrix. The rank depends on the available paths between sources and each intermediate node. More specifically, depending on the topology of the graph, some nodes may receive only combinations of symbols derived from matrices with restricted rank, i.e. less than K . This includes, for example, trees, where a node connected directly to the source by a link of capacity C can only have children that receive at most rank C .

As a first step towards general network models, we consider the case of complete acyclic directed graphs $G = (V, E)$, $n = |V|$, which can be generated as follows.

- Generate random labels for the n vertices. These have some ordering $\{e_1, e_2, \dots, e_n\}$ associated to them;
- Make an outgoing (directed) edge from the vertex with the minimum label to every vertex with a higher label;
- Continue until we reach a vertex where there are no more possibilities for connections.

This algorithm generates a complete acyclic directed graph with one source, one sink and $|E| = n(n-1)/2$ edges, since the total degree of each vertex is $n-1 = \delta_I(v) + \delta_O(v)$. The source and the sink are naturally determined as those nodes that have only outgoing edges or only incoming edges, respectively. The ordering ensures that this algorithm always generates an acyclic directed graph, conferring the graphs generated in this way specific properties such as the distribution of the in and out-degrees. These properties can be determined directly from the order of the vertex using $\delta_O(v) = n - \text{order}(v)$ and $\delta_I(v) = n - \delta_O(v) - 1 = \text{order}(v) - 1$.

Before proving our next theorem, we introduce the following lemmas.

Lemma 4: In complete acyclic directed graphs, a node that receives R symbols, receives w.h.p. a partial transfer matrix with rank equal to $\min(R, K)$.

Proof: See the Appendix. ■

Lemma 5: For the complete directed acyclic graph, w.h.p.,

$$\Delta_S(v) = \frac{K - \min(K, \text{order}(v))}{K}.$$

Proof: See the Appendix. ■

Theorem 2: Let ϕ_S be the *secure max-flow*, defined as the maximum number of symbols that may be secured in a transmission by using random linear network coding. For a complete acyclic directed graph with n nodes, the secure max-flow equals the max-flow min-cut capacity of the network and is $n-1$. Conversely, the minimum numbers of required symbols for secured transmission is $n-1$ symbols.

Proof:

Suppose, by contradiction, that $K = n-1$ is the max-flow min-cut capacity of the complete directed acyclic graph. The maximum order of an intermediate node v is $n-2$, thus by Lemma 5 we have $\Delta_S(v) = 1/(n-1)$. It follows that the secure max-flow of the complete acyclic directed graph equals the capacity of the graph.

By contradiction, let the minimum number of required symbols for secured transmission be $m_s \leq n-2$. There exists an intermediate node v such that $\text{order}(v) = n-1$, and consequently, $\Delta_S(v) = 0$. Then the minimum number of required symbols for secure transmission is $m_s = n-1$. ■

It follows that the way to secure this class of complete graphs is to transmit at the max-flow min-cut capacity, if necessary by adding “dummy” symbols.

V. CONCLUSIONS

Intrigued by the security potential inherent to random linear network coding, we developed a specific algebraic security criterion, for which we proved a set of key properties. Perhaps one of the most striking conclusions of our analysis is that algebraic security with network coding is very dependent on the topology of the network. As an example, we focused on complete acyclic directed graphs, and determined the secure max-flow, as well as the minimum number of symbols required for algebraic security. As part of our ongoing work, we are extending this analysis to other more general network models. Ultimately, we would like to develop secure communication protocols capable of exploiting random linear network coding as an *almost free* cypher.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge insightful discussions with Rui A. Costa (Univ. of Porto).

REFERENCES

- [1] R. Ahlswede, N. Cai, S.Y.R. Li, and R.W. Yeung, “Network information flow,” *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1204–1216, 2000.
- [2] R. Koetter and M. Medard, “An algebraic approach to network coding,” *IEEE/ACM Transactions on Networking*, vol. 11, no. 5, pp. 782–795, 2003.
- [3] N. Cai and R.W. Yeung, “Secure network coding,” *Proceedings of the IEEE International Symposium on Information Theory*, Lausanne, Switzerland, 2002.
- [4] J. Feldman, T. Malkin, C. Stein, and R.A. Servedio, “On the capacity of secure network coding,” *Proc. of the 42nd Annual Allerton Conference on Communication, Control, and Computing*, 2004.
- [5] K. Bhattad and K.R. Narayanan, “Weakly secure network coding,” *Proc. of the First Workshop on Network Coding, Theory, and Applications (NetCod)*, Riva del Garda, Italy, 2005.
- [6] K. Jain, “Security based on network topology against the wiretapping attack,” *IEEE Wireless Communications*, vol. 11, no. 1, pp. 68–71, 2004.

- [7] T. Ho, R. Koetter, M. Medard, D.R. Karger, and M. Effros, "The benefits of coding over routing in a randomized setting," *Proc. of the IEEE International Symposium on Information Theory (ISIT)*, Yokohama, Japan, June/July 2003.
- [8] T. Ho, M. Medard, J. Shi, M. Effros, and D.R. Karger, "On randomized network coding," *Proceedings of the 41st Annual Allerton Conference on Communication, Control, and Computing*, 2003.
- [9] D.S. Lun, M. Medard, R. Koetter, and M. Effros, "On Coding for Reliable Communication over Packet Networks," *Arxiv preprint cs.IT/0510070*, 2005.
- [10] T. Ho, B. Leong, R. Koetter, M. Medard, M. Effros, and DR Karger, "Byzantine modification detection in multicast networks using randomized network coding," *Proceedings of the International Symposium on Information Theory*, Yokohama, Japan, June/July 2003.
- [11] S. Jaggi, M. Langberg, T. Ho, and M. Effros, "Correction of adversarial errors in networks," *Proceedings of the International Symposium on Information Theory*, Adelaide, Australia, September 2005.
- [12] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, and M. Medard, "Resilient Network Coding In the Presence of Byzantine Adversaries," *IEEE Infocom*, 2006.
- [13] Sidharth Jaggi, *Design and Analysis of Network Codes*, Ph.D. thesis, California Institute of Technology, 2005.
- [14] J. Tan and M. Medard, "Secure Network Coding with a Cost Criterion," *Proc. 4th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt'06)*, Boston MA, April, 2006.
- [15] C.E. Shannon, *Communication Theory of Secrecy Systems*, Bell Systems Technical Journal, Vol. 28, pp. 656-715, October 1949.

APPENDIX

Proof of Lemma 1

We will prove this constructively in terms of the ranks of parts of the transfer matrix. The auxiliary encoding vector in each intermediate node v is given by

$$M'_{\Gamma_I(v)} = (A(I - F)^{-1})_{\Gamma_I(v)},$$

where $M'_{\Gamma_I(v)}$ denotes the columns of the matrix corresponding to the incoming edges of v . The dimension of $M'_{\Gamma_I(v)}$ is $K \times \delta_I(v)$, with $\delta_I(v) < |E|$.

To determine the rank of the partial transfer matrix, we note that the transfer matrix $M = A(I - F)^{-1}B^T$ for the network must be invertible, and hence, $\text{rank}(M) = K$. On the other hand, to determine the rank of $A(I - F)^{-1}$ we use the fact that $(I - F)^{-1}$ is invertible and thus $\text{rank}((I - F)^{-1}) = |E|$. We also have

$$\text{rank}(A(I - F)^{-1}) \leq |E|,$$

because the dimension of $A(I - F)^{-1}$ is $K \times |E|$. But, since

$$\text{rank}(A(I - F)^{-1}B^T) = K = \min(\text{rank}(A(I - F)^{-1}), B)$$

holds and $K < |E|$ (true because K must be less than the minimum cut in the network) we conclude that

$$\text{rank}(A(I - F)^{-1}) = K.$$

We now consider $\Delta_S(v)$ at some vertex v . For that, we can consider two distinct cases: the first one is if $K < \delta_I(v)$. In this case, we cannot assume anything about $\Delta_S(v)$, since the rank of the matrix $M'_{\Gamma_I(v)}$ will be dependent on the topology of the network. As for the second case, $\text{rank}(M'_{\Gamma_I(v)}) < K \Rightarrow \Delta_S(v) < 0$. ■

Proof of Lemma 2

Contrary to the sum, the product of independent and uniformly distributed values in \mathbb{F}_q is *not* independent and uniformly distributed. In fact, there are two ways to obtain a zero in a multiplication in \mathbb{F}_q : (1) by multiplication between an element $a \in \mathbb{F}_q$ and 0, and (2) by multiplication over two elements $a \in \mathbb{F}_q$ and $b \in \mathbb{F}_q$, such that $a \neq 0$ and $b \neq 0$, but $ab = 0$. Now, the total number of entries of the multiplicative table between q elements of \mathbb{F}_q is q^2 , and there are at most $2q$ instances of the first case: q instances of $ab = 0$, $a = 0$ and $b \neq 0$, and q instances of $ab = 0$, $a \neq 0$ and $b = 0$. As for the second case, it is possible to prove by contradiction that the number of zeros obtained this way is strictly less than q^2 : if this was not the case, all products of elements of \mathbb{F}_q would be zero, and that is absurd. Since this is true for any q , the number of zeros grows $O(h(q)) < O(q^2)$. Thus, we have

$$P(X_{lin} = 0) \leq \frac{2q + h(q)}{q^2}.$$

Since for large enough q we have $(2 + h(q))/q < 1$, it follows that

$$P(X_{lin} = 0)_{q \rightarrow \infty} = 0. \quad \blacksquare$$

Proof of Lemma 3

Each position of a line of the transfer matrix M is a linear combination of independently and uniformly chosen values in \mathbb{F}_q , and thus, the probability of obtaining a zero in a position is given by *Lemma 2*. The result follows by considering all the combinations of the possible positions in which the Y zeros may occur. ■

Proof of Lemma 4

Suppose that a given intermediate node receives $R = K + \theta$ symbols, $\theta \geq 0$. It is clear that the maximum possible rank is K and thus there is a way to remove θ columns s.t. the rank of the resulting set will still be at maximum K . Now consider the case in which vertex v receives at most K symbols. If the columns are linearly dependent, the condition

$$\{x_{h_1}c_{h_1} + x_{h_2}c_{h_2} + \dots + x_{h_n}c_{h_n} = (0\dots 0)^T\},$$

such that $x_{h_1}, x_{h_2}, \dots, x_{h_n}$ not all 0, $\in \mathbb{F}_q$ and h_1, h_2, \dots, h_n represent the columns $\in \Gamma_I(v)$, will be satisfied. Since the linear combination of lines of the transfer matrix is again a linear combination of independent and uniformly distributed values in \mathbb{F}_q , it follows from *Lemma 3* that the probability of obtaining $(0\dots 0)^T$ tends to 0 when $q \rightarrow \infty$ and $K \rightarrow \infty$, and thus, the columns $h_1, h_2, \dots, h_n \in \Gamma_I(v)$ are linearly independent w.h.p. ■

Proof of Lemma 5

It follows from *Lemma 4* that w.h.p., the number of symbols received by a vertex is the rank of the partial transfer matrix received (and at most K) and thus

$$\Delta_S(v) = \frac{K - \min(K, \delta_I(v))}{K} = \frac{K - \min(K, \text{order}(v) - 1)}{K}$$

■