On the Capacity of Secure Network Coding

Jon Feldman[*] Dept. of IEOR

Tal Malkin[†] Dept. of CS

Rocco A. Servedio[‡] Dept. of CS

Cliff Stein[§] Dept. of IEOR

Columbia University, New York, NY

{jonfeld@ieor, tal@cs, rocco@cs, cliff@ieor}.columbia.edu

Abstract

We consider the problem of using a multicast network code to transmit information securely in the presence of a "wire-tap" adversary who can eavesdrop on a bounded number of network edges. Cai & Yeung (ISIT, 2002) gave a method to alter any given linear network code into a new code that is secure. However, their construction is in general inefficient, and requires a very large field size; in many cases this is much greater than the field size required by standard network code construction algorithms to achieve the min-cut capacity (without a security guarantee).

In this paper we generalize and simplify the method of Cai & Yeung, and show that the problem of making a linear network code secure is equivalent to the problem of finding a linear code with certain generalized distance properties. We show that if we give up a small amount of overall capacity, then a random code achieves these properties using a much smaller field size — in some cases a field of constant size suffices — than the construction of Cai & Yeung. We add further support to this approach by showing that if we are not willing to give up any capacity, then a large field size may sometimes be required to achieve security.

1 Introduction

1.1 The Network Coding Model. An instance of the *multicast linear network coding* problem consists of a directed acyclic graph $G = (V_G, E_G)$, a source node s_G, a set T_G of sink nodes, a message length n, and a field $\mathbb{F}_q$ of size q. The edges of G are used to transmit information through the graph; each edge carries one element of $\mathbb{F}_q$ per time step.

The goal in network coding is to design a scheme whereby an arbitrary message vector $\mathbf{m} \in \mathbb{F}_q^n$, which originates at the source node s_G, may be communicated over this network so that each sink can recover the entire vector $\mathbf{m}$. Such a scheme, which we refer to as a *solution* to the network coding problem, is said to be *feasible* if two conditions are met: (i) For each edge (u, v) the symbol transmitted over (u, v) is some function of the symbols that are available at node u. If u is the source, the entire message vector $\mathbf{m}$ is available; otherwise the symbols transmitted on edges (w, u) into node u are available. (ii) For each

[*]J. Feldman was supported by an NSF Postdoctoral Research Fellowship DMS-0303407.

[†]T. Malkin was partially supported by NSF Early Career Development (CAREER) Grant CCF-0347839.

[‡]R. Servedio was partially supported by NSF Early Career Development (CAREER) Grant CCF-0347282.

[§]C. Stein was partially supported by NSF Grant DMI-9970063.

sink node in T_G there must be some function which, if applied to the symbols received at that node, yields the original message $\mathbf{m}$.

In a *linear* network code, each of the functions described above is a linear function. Thus a linear solution to the network coding problem is given by a list of vectors $(\mathbf{v}[e])_{e \in E_G}$ describing which linear combination of the messages is transmitted on each edge (the symbol $\mathbf{v}[e] \cdot \mathbf{m}$ is carried on edge e). Feasibility implies that for all edges (u, v), the symbol $\mathbf{v}[u, v] \cdot \mathbf{m}$ may be computed as a linear combination of the symbols $\mathbf{v}[w, u] \cdot \mathbf{m}$ carried on edges (w, u) into node u.

It is now well-known that given an instance of the multicast linear network coding problem with $q \geq |T_G|$, a feasible solution exists if and only if the minimum cut between the source and each sink is of size at least n [1, 13]; moreover, efficient algorithms are known for constructing feasible solutions [9].

1.2 Security Against a Wire-Tap Adversary. In this paper, we consider the following problem: we are given an instance of the multicast linear network coding problem as described above, and a feasible solution. Our task is to make this solution *secure* against an eavesdropping adversary. Throughout the paper we assume a computationally unbounded "wire-tap" adversary against the network code. This adversary has access to the symbols which are transmitted over some unknown set of at most k edges of the network. Additionally, the adversary has full knowledge of the network code itself and of whatever protocol we use for security. We would like to transmit some information x over the network in a way that is information-theoretically secure against this adversary; roughly speaking, information-theoretically secure means that no matter which set of k edges the adversary accesses, she should gain no information about the information x. (We give a precise definition later.)

The study of secure network coding in general, and our model in particular, are well motivated. We first note that by the results of [9], given an instance of the problem, a feasible solution can be efficiently constructed. Thus, it is quite reasonable for us to assume that a feasible solution is already given to us, and we just need to transform it to a secure one. Network coding in general has the appealing advantage of being able to increase the throughput of a given network, and has been suggested as a practical tool for use in content distribution networks over the Internet, as well as ad-hoc wireless networks [3]. Because of the insecure nature of such networks, it is important to consider security issues in network coding. While there are many possible definitions of security, we feel that our framework—in which no assumptions are made about which k edges have been compromised and we strive for information-theoretic, i.e. perfect, security—is a strong and natural one to study.

1.3 Related Work. Network coding was introduced by Ahlswede *et al.* [1], and has since received a lot of attention (e.g., [11, 6, 12]). The problem of making a linear network code secure was first studied by Cai & Yeung [2], who considered a "wire-tap" adversary. We give more details of their work throughout the paper. Jain [10] also considers this model, and gives more precise security conditions for the case of a single sink. Ho *et al.* [7] consider the related problem of network coding in the presence of a *Byzantine attacker* which can modify data sent from a node in the network.

Given a network code with message length n, and a wire-tap adversary that is capable of looking at sets of at most $k < n$ edges, Cai & Yeung [2] suggest using a linear "secret-sharing" method to provide security in the network. Instead of sending n message symbols, the user sends k random symbols and $n - k$ message symbols. Additionally, the code itself undergoes a certain linear transformation. Cai & Yeung give sufficient conditions for this transformation to guarantee security. They also show that as long as the field size $q > \binom{|E_G|}{k}$, a secure linear transformation exists.[1] Unfortunately, this lower bound is much greater than the $q \geq |T_G|$ lower bound required for constructing the code itself without a security guarantee [13]. Also, their construction of the linear transformation takes at least $\binom{|E_G|}{k}$ time steps. This complexity, as well as the required lower bound on the field size q, is quite restrictive when k is large.

1.4 Our Results. Cai & Yeung achieve security by altering the code (using a linear transformation) on each edge. We take a somewhat different approach: we only modify the input, while leaving the code unchanged. Although it is easily shown that our method is equivalent in power to that of Cai & Yeung, it has the advantage that the code itself does not need modification to be used securely.

We adapt and simplify the proof of Cai & Yeung showing that certain independence conditions are sufficient for the linear transformation to yield a secure code; in addition, we show that these same conditions are also necessary.

We then prove that a secure linear transformation exists if and only if there is a solution to a certain generalized (classical) code construction problem. With this new coding-theoretic characterization, we are able to give a positive result based on methods similar to proving the Gilbert-Varshamov bound (see [8]). We show that if one is willing to give up a little bit of capacity—namely, sending $(1+\epsilon)k$ random symbols and $n-(1+\epsilon)k$ message symbols—then a random linear transformation will be secure with high probability, as long as (roughly) $q \geq \Theta(|E_G|^{1/\epsilon})$. This is superior to the bound $q \geq \binom{|E_G|}{k}$ in most cases, and allows a trade-off between capacity and field size. For very large $k = \Theta(|E_G|)$, our lower bound becomes $q > 2^{\Omega(1/\epsilon)}$, a constant independent of $|E_G|$.

We also give a negative result, supporting the need to give up capacity in order to achieve security with a small field size. Using a result [5] on the covering radius [4] which linear codes can achieve, we show that if one insists upon sending $n-k$ message symbols, then there are cases where the field size must be almost as large as $|E_G|^{\sqrt{k}}$. (We give more precise statements of both our positive and negative results later in the paper.)

2 Preliminaries

Throughout the paper all vectors v are row vectors unless otherwise indicated, and we write v^T to denote the corresponding column vector. If v is an n-dimensional row vector and w is an m-dimensional row vector we write (v, w) to denote the $(n+m)$-dimensional row vector obtained by concatenating v and w.

Given $x \in \mathbb{F}_q^N$, the *ball of radius* d around x is the set of all vectors in $\mathbb{F}_q^N$ which differ from x in at most d coordinates. We write $\mathrm{Vol}_q(d, N)$ to denote the number of vectors in this ball.

2.1 Secure Linear Network Coding. As in Section 1, we assume that we have a (not necessarily secure) network code which can be used to transmit a message $\mathbf{m} \in \mathbb{F}_q^n$. Our goal is to transform the code into one allowing secure transmission of some input information, in the presence of a wire-tapping adversary. To do this, we apply an initial linear encoding $E : \mathbb{F}_q^{n-\ell} \times \mathbb{F}_q^{\ell} \to \mathbb{F}_q^n$ to the input information $x \in \mathbb{F}_q^{n-\ell}$ and some randomly

[1] Their lower bound is actually $q > |\mathcal{A}|$, where $\mathcal{A}$ is the set of possible subsets of edges available to the adversary. In our model we assume no knowledge of which set of k edges the adversary may have infiltrated, so $\mathcal{A}$ consists of all sets of k edges and we have $|\mathcal{A}| = \binom{|E_G|}{k}$.

chosen $r \in \mathbb{F}_q^\ell$; the vector $\mathbf{m} = E(x, r)$ is then transmitted as normal, using the original (deterministic) network code. This approach (as opposed to changing the code itself) is natural, and can be used in applications where the sender does not have control over the network code, yet wants to achieve security.

By the feasibility of the network code, the entire message vector $\mathbf{m} = E(x, r)$ can be recovered at the sinks, and so $E(x, r)$ must uniquely define the information x. Formally, we say E is *sound* if, for all x, x' where $x \neq x'$, we have $E(x, r) \neq E(x', r')$ for all r, r'.

The *capacity* of the network code is $n - \ell$, the dimension of the information vector x.

The security condition we impose (which is equivalent to the one given by Cai & Yeung) is that for any set of at most k edges in the network, knowing the values sent across those edges gives no information about the information word x. Formally, for any $P = \{e_1, \ldots, e_{k'}\} \subseteq E_G$ where $|P| = k' \leq k$, for any $\mathbf{a} = (a_1, \ldots, a_{k'}) \in \mathbb{F}_q^{k'}$, and for any $x \in \mathbb{F}_q^{n-\ell}$, let

$$R(x, \mathbf{a}) = \{r \in \mathbb{F}_q^\ell : E(x, r) \cdot \mathbf{v}[e_i] = a_i \text{ for } i = 1, \ldots, k'\}. \tag{1}$$

That is, $R(x, \mathbf{a})$ is the set of all possible vectors r such that when $\mathbf{m} = E(x, r)$ is sent, the observed information on P is $\mathbf{a}$.

The encoding E is *secure* if for all $P, \mathbf{a}$ as above, and for all $x, x' \in \mathbb{F}_q^{n-\ell}$, we have $|R(x, \mathbf{a})| = |R(x', \mathbf{a})|$. Note that the definition implies that when r is chosen uniformly at random, whatever symbols $(a_1, \ldots, a_{k'})$ the adversary sees on the edges she controls gives no information about the transmitted x.

2.2 Connection to Secret Sharing. We note that the problem, as defined above, is a generalization of the cryptographic concept of secret sharing. Indeed, in the degenerate case of a network code consisting of two nodes, a source and a sink, with n parallel edges between them, our problem becomes an instance of secret sharing: We need to find a way to encode a secret into n "shares" so that no k of the shares give any information about the secret, but getting all shares allows recovery of the secret. For general networks, we have $|E_G|$ *linear combinations* of the n shares, and the requirement is that no subset of at most k of these linear combinations gives any information about the secret, while knowing all n shares (which are computable by the feasibility of the given code) allows recovery of the secret.

3 Necessary and Sufficient Conditions for Secure Coding

In this section we describe a general family of linear encoders $E(x, r)$, and show necessary and sufficient independence conditions for such encoders to be secure. One direction (sufficiency) of the proof in this section is implicit in work of Cai & Yeung [2]. We offer a simpler presentation, a generalization to the case $\ell \geq k$, and a proof that the condition given is also necessary.

Let $x = (x_1, \ldots, x_{n-\ell})$ be an information vector, and $r = (r_1, \ldots, r_\ell)$ be a random vector, where $n > \ell \geq k$. (Recall that k is the bound on the adversary). Let $\mathbf{y} = (x, r)$ be the concatenation of the information and the random vectors. Our secure encoding function is $E(x, r) = \mathbf{y}M^{-1}$, where M is an invertible n-by-n matrix. So, each edge of the network carries the symbol $(\mathbf{y}M^{-1}) \cdot \mathbf{v}[e]$, where $\mathbf{v}[e]$ is given by the network code. Since M is invertible, we have that the encoding is sound.

The encoding is secure under the following conditions on M:

Theorem 1 *For an invertible n-by-n matrix M, the encoding $E(x,r) = (x,r)M^{-1}$ is secure if and only if any set consisting of*

(a) at most k linearly independent vectors from $\{\mathbf{v}[e]\}_{e \in E_G}$, and

(b) any number of vectors from the first $n-\ell$ columns of M

is linearly independent.

Proof: We first show that the encoding is secure if the independence condition is met. Suppose the adversary has access to a set P of edges, where $k' = |P| \leq k$. Let $V = V(P)$ be a n-by-k' matrix where the columns of V are the vectors $\{\mathbf{v}^T[e]\}_{e \in P}$. We may assume that V has rank k', since otherwise the adversary could drop an edge and not lose any information. When $\mathbf{y}$ is sent over the network, the adversary sees $(\mathbf{y}M^{-1}\mathbf{v}^T[e])$ for each edge $e \in P$; in other words, the adversary sees the length-k' vector $\mathbf{a} = \mathbf{y}M^{-1}V$.

Given an information vector $\hat{x} \in \mathbb{F}_q^{n-\ell}$, the set $R(\hat{x}, \mathbf{a})$ (as defined in (1)) has one member for every solution to the system of $n - \ell + k'$ equations in n unknowns described by $\hat{\mathbf{y}}V' = (\hat{x}, \mathbf{a})$, where $\hat{\mathbf{y}} \in \mathbb{F}_q^n$ is unknown, and V' is the following n by $(n - \ell + k')$ coefficient matrix:

$$V' = \left[\begin{array}{c|c} I_{n-\ell} & M^{-1}V \\ 0 & \end{array}\right].$$

Now suppose that V' has full rank; then, for all $(\mathbf{a}, \hat{x})$ pairs on the right hand side, the system of equations has exactly the same number of solutions. It follows that for all $\mathbf{a}$, we have $R(x, \mathbf{a}) = R(x', \mathbf{a})$ for all distinct information vectors x, x', and thus the encoding is secure.

To prove that V' has full rank, we consider the matrix MV' (recall that M has full rank by definition). This matrix MV' has its first $n - \ell$ columns matching the first $n - \ell$ columns of M, and the last k columns are the matrix $MM^{-1}V = V$. Thus, if the conditions on M in the theorem hold, the matrix MV' has full rank for all possible choices of V, and thus the encoding is secure.

For the other direction, suppose the independence condition is not met. This means that there is some nontrivial linear combination of some l.i. set $\{\mathbf{v}[e]\}_{e \in P}$ of at most k edge vectors which equals some nontrivial linear combination of the first $n - \ell$ (l.i.) columns of M. If we define V and V' in terms of P as above, then this is equivalent to saying that MV' is not full rank, and thus V' is not full rank, since M is full rank by assumption. We may conclude that $V'(z_1, z_2)^T = 0$ for some $z_1 \in \mathbb{F}_q^{n-\ell}$, $z_2 \in \mathbb{F}_q^{k'}$. Also, we know that $z_1 \neq 0$ and $z_2 \neq 0$ by looking at the structure of V' (using the fact that $I_{n-\ell}$, V and M are all full rank). So, we have $\begin{bmatrix} I_{n-\ell} \\ 0 \end{bmatrix} z_1^T + M^{-1}Vz_2^T = 0$.

Fix some information vector $x \in \mathbb{F}_q^{n-\ell}$ and random vector $r \in \mathbb{F}_q^{\ell}$. Let $\mathbf{a} = (x,r)M^{-1}V$ be the vector of observed symbols on the edges P. Since $\mathbf{a}$ is the result of a possible choice of r, we have that $R(x, \mathbf{a}) > 0$. Note that

$$\mathbf{a} \cdot z_2 = (x,r)M^{-1}Vz_2^T = -(x,r)\begin{bmatrix} I_{n-\ell} \\ 0 \end{bmatrix} z_1^T = -x \cdot z_1.$$

In fact, the relation $\mathbf{a} \cdot z_2 = -x \cdot z_1$ holds for any pair of possible information vectors x and observed vectors $\mathbf{a}$ which satisfy $R(x, \mathbf{a}) > 0$.

Let $x' = x + e_i$ where i is an index such that $(z_1)_i \neq 0$. Thus we have that $x' \cdot z_1 \neq x \cdot z_1 = -\mathbf{a} \cdot z_2$. We conclude that $\mathbf{a}$ is not a possible observed vector for x', and thus $R(x', \mathbf{a}) = 0$. We have demonstrated an $\mathbf{a}, x, x'$ where $R(x, \mathbf{a}) \neq R(x', a)$, so the encoding is not secure. ■

3.1 The Existence of a Secure Matrix.

We say that a matrix M which meets the conditions of Theorem 1 is *secure*. Implicit in the work of Cai & Yeung [2] is a proof that a secure matrix M exists with $\ell = k$, as long as the field size q satisfies $q > \binom{|E_G|}{k}$. However, the algorithm they give in their proof for finding such a matrix M takes at least $\binom{|E_G|}{k}$ time steps. Also, having an alphabet of this size may well be prohibitive for certain networks.

4 Finding a Secure Linear Network Code is a Coding Problem

In this section we give our result showing that finding a secure matrix M meeting the independence conditions of Theorem 1 is equivalent to finding a linear code with certain generalized distance properties. Roughly speaking, the code we are looking for must have all its codewords far away from any word in a linear subspace defined by the vectors $\{\mathbf{v}[e]\}_{e\in E_G}$. In Sections 5 and 6 we will use this equivalence to establish upper and lower bounds on the field size required for secure network coding.

4.1 Preliminaries.

We henceforth write N for $|E_G|$. Let Z be an $n \times N$ matrix whose columns are the vectors $\{v[e]\}_{e\in E_G}$ in some fixed (arbitrary) order. By construction of the network code, since there must be n linearly independent vectors $\mathbf{v}[e]$ in every source-sink cut, the matrix Z must have rank n. Let A be an $(N-n) \times N$ generator matrix for the null space of Z. (Equivalently, A is the parity-check matrix if Z is regarded as a generator for a code.)

We will henceforth use the notation M to mean the first $n-\ell$ columns of the invertible square matrix "M" in Theorem 1. With this new notation, a matrix M is secure[2] iff

$$Mx^T + Zw^T \neq 0 \text{ for all } x \in \mathbb{F}_q^{n-\ell}, w \in \mathbb{F}_q^N \text{ s.t. } x \neq 0, |w| \leq k. \qquad (2)$$

We define a notion of "distance" between two matrices that is (roughly) the minimum distance between two vectors in the span of their rows. More precisely, for an $\alpha \times n$ matrix P and a $\beta \times n$ matrix Q, we define

$$\delta(P,Q) \equiv \min_{x\in\mathbb{F}_q^\alpha, y\in\mathbb{F}_q^\beta, y\neq 0} \Delta(xP, yQ),$$

where Δ is the Hamming distance. Note the slight asymmetry in the treatment of P and Q, namely that x can be 0^α but y cannot be 0^β; this makes the minimum distance of the code generated by Q an upper bound on $\delta(P,Q)$.

4.2 Main Theorem.

Now we present our main theorem relating the above notion of distance to the existence of a secure matrix M:

Theorem 2 *Given the matrix A as defined above, there exists a secure $n \times (n-\ell)$ matrix M if and only if there exists an $(n-\ell) \times N$ matrix B with $\delta(A,B) > k$.*

Proof: Suppose there is some $n-\ell \times N$ matrix B with $\delta(A,B) > k$. Let $M = ZB^T$. Note that B must have rank $n-\ell$, since otherwise it could not have $\delta(A,B) > 0$.

Because $\delta(A,B) > k$, and A generates the null space of Z, we have that $\Delta(y^T, B^T x^T) > k$ for all $x \in \mathbb{F}_q^{n-\ell}$, $x \neq 0$ and $y : Zy^T = 0$. Therefore, $Z(B^T x^T + w^T) \neq 0$ for all $x \in \mathbb{F}_q^{n-\ell}$,

[2] Since the security of "M" (as in Theorem 1) depends only on its first $n-\ell$ columns, we may extend a matrix M from (2) to be square and invertible using an arbitrary extension of M to a basis, as in [2].

$x \neq 0$ and $w \in \mathbb{F}_q^N$ where $|w| \leq k$. This implies $ZB^T x^T + Zw^T = Mx^T + Zw^T \neq 0$ for all such x, w, which are exactly the security conditions in (2).

For the other direction, if we suppose there is a secure M, we construct B as follows. For each column M_i of M, let the ith column of B^T be an arbitrary member of the coset $\{y \in \mathbb{F}_q^N : Zy = M_i\}$. Note that we again have $M = ZB^T$.

Since M is secure, we have $Mx + Zw \neq 0$ for all $x \neq 0$, $|w| \leq k$ (from (2)), and so $Z(B^T x + w) \neq 0$ for all such x, w. Thus for all $y : Zy = 0$, and $x \neq 0$, we have $\Delta(B^T x, y) > k$. This implies $\delta(A, B) > k$. ∎

Thus the question of whether there exists a secure M is equivalent to the question of whether there exists a matrix B with good distance from A, the generator for the null space of Z.

4.3 A Generalized Coding Problem. Having proved Theorem 2, our interest is now in the following problem:

> **Span Distance Problem:** Given an α-by-N matrix A with rank α, whose entries belong to $\mathbb{F}_q$, find a β-by-N matrix B over $\mathbb{F}_q$ such that $\delta(A, B) > k$.

We can regard this question as a generalization of the classical code construction problem: if $\{xB\}_{x \in \mathbb{F}_q^\beta}$ is regarded as a code, then every non-zero codeword must have good distance not only from the all-zeros codeword, but also from every other word generated by A.

In the following sections, we consider the *Span Distance Problem* abstractly. When we apply this problem to Theorem 2, we have $\alpha = N - n$ and $\beta = n - \ell$. Setting ϵ such that $\ell = (1 + \epsilon)k$, we are now interested in the case of the span distance problem where $k = \frac{N - \alpha - \beta}{1 + \epsilon}$. The parameter ϵ represents the amount of capacity we are willing to give up in order to reduce the field size necessary to achieve security.

5 A Positive Result: giving up capacity to save on field size

The main theorem in this section is the following:

Theorem 3 *Let A be an arbitrary α-by-N matrix with rank α over $\mathbb{F}_q$, and let B be a random β-by-N matrix over $\mathbb{F}_q$. Let k, ϵ be such that $k = \frac{N - \alpha - \beta}{1 + \epsilon}$. Then we have $\delta(A, B) > k$ with probability at least $1 - P_{BAD}$, where*

$$P_{BAD} = q^{-(1+\epsilon)k} \mathrm{Vol}_q(k, N). \tag{3}$$

Proof: The argument follows along the same lines as the classical argument that random linear codes meet the Gilbert-Varshamov bound. Let BAD be the set of words in $\mathbb{F}_q^N$ with distance at most k from some linear combination of the rows of A. Using the bound $|\mathrm{BAD}| \leq q^\alpha \mathrm{Vol}_q(k, N)$, we have that for a particular $x_1 \in \mathbb{F}_q^\beta$, the probability (over choices of B) that $x_1 B \in \mathrm{BAD}$ is at most $\frac{q^\alpha \mathrm{Vol}_q(k,N)}{q^N}$. Applying a union bound over $x_1 \in \mathbb{F}_q^\beta$, we have the probability of some $x_1 B$ being in BAD is at most $P_{BAD} \leq q^{\alpha + \beta - N} \mathrm{Vol}_q(k, N) = q^{-(1+\epsilon)k} \mathrm{Vol}_q(k, N)$. ∎

5.1 Applying Theorem 3. Here we show that Theorem 3 allows us to use fields of quite modest size and still achieve a good probability bound in (3). We have

$$\mathrm{Vol}_q(k, N) = \sum_{i=0}^{k} (q-1)^i \binom{N}{i}. \tag{4}$$

We consider two different ranges of values for k (these are $k = o(N)$ and $k = \Theta(N)$) and use different upper bounds on $\mathrm{Vol}_q(k, N)$ in these two cases. We use the following facts in the bounds.

Fact 4 *[8] For any $q \geq 2$, if $0 < k < (1 - 1/q)N$, then*

(a) the largest term in the sum (4) is the $i = k$ term,

(b) $\frac{\log \mathrm{Vol}_q(k,N)}{\log q} = (H_q(k/N) \pm o(1))N$, where

$$H_q(\delta) \quad = \quad \delta \log_q(q-1) - \delta \log_q \delta - (1-\delta)\log_q(1-\delta). \tag{5}$$

We first consider the case $k = o(N)$. In this case, from Fact 4(a), it follows that $\mathrm{Vol}_q(k, N) \leq (k+1)q^k N^k$, and so (3) implies that the probability $1 - P_{BAD}$ is positive as long as $q > (k+1)^{1/(\epsilon k)} \cdot N^{1/\epsilon}$. To obtain a high probability result such as $P_{BAD} < N^{-c}$ for some constant $c > 0$, it suffices to take $q > (k+1)^{1/(\epsilon k)} \cdot N^{(1+c/k)/\epsilon}$. These lower bounds on q are easily seen to be much less restrictive than the $q > \binom{N}{k}$ lower bound of Cai & Yeung, at the cost of only a small loss in capacity (number of information symbols we can transmit). As one example, if we take $\epsilon = 1$ then we achieve capacity $n - 2k$ (as opposed to Cai & Yeung's $n - k$), but we require only that the field size q be (roughly) at least $N^{1+c/k}$ which is close to N for moderate k and small constant c. (Of course, even smaller lower bounds on q can be achieved by taking $\epsilon > 1$.) Thus, if k satisfies both $k = \omega(1)$ and $k = o(N)$, we lose only a $(1 - o(1))$ factor in capacity while obtaining a superpolynomial savings in field size.

We now consider the case $k = \Theta(N)$, and show that here we can achieve even more dramatic savings in field size. Taking $k = \delta N$ where $\delta = \Theta(1)$ is some constant and plugging Fact 4(b) into (3), we have that $P_{BAD} \leq q^{-(1+\epsilon)\delta N} q^{N(H_q(\delta)+o(1))}$. It is easy to see from (5) that $H_q(\delta) < \delta + \frac{1}{\log q}$, and thus (5) implies that $P_{BAD} < q^{(-\delta\epsilon/2+o(1))N}$ provided that $\delta\epsilon/2 > 1/\log q$, i.e. $q > 2^{2/(\delta\epsilon)}$. Since δ is a fixed constant independent of N, the field size lower bound is $2^{\Omega(1/\epsilon)}$ which is independent of N. (We remind the reader that in order to construct a multicast linear network code efficiently [9], it must be the case that $q > |T_G|$.)

6 A Negative Result: communicating securely at capacity can require large field size

The main result in this section is the following theorem:

Theorem 5 *Let $\alpha = N - \frac{\log N}{\log q} - \frac{\log \mathrm{Vol}_q(k,N)}{\log q} + 2\log N + \log q + \log \ln q$. If α, β satisfy*

$$k + \beta < N - \alpha = \frac{\log N}{\log q} + \frac{\log \mathrm{Vol}_q(k, N)}{\log q} - 2\log N - \log q - \log \ln q \tag{6}$$

then there is an $\alpha \times N$ matrix A over $\mathbb{F}_q$ such that there is no $\beta \times N$ matrix B over $\mathbb{F}_q$ for which $\delta(A, B) > k$.

In words, this theorem says that for certain values of α and β, if q is too small then there exists an $\alpha \times N$ matrix A over $\mathbb{F}_q$ for which the span distance problem cannot be solved if we take $k = N - \alpha - \beta$. Since $k = N - \alpha - \beta$ corresponds to taking $\epsilon = 0$, this means that if we do not give up some capacity then there need not exist a secure matrix M unless the field size q is quite large.

6.1 Using a Code with Good Covering Radius. To establish Theorem 5, we need to find a full-rank α-by-N matrix A which is such that for all full-rank β-by-N matrices B, there is a point $x_1 B$, $x_1 \neq 0$ and a point $x_2 A$ where $\Delta(x_1 B, x_2 A) \leq k = N - \alpha - \beta$.

For the case $\beta = 1$, this is exactly a question of constructing a code A with small *covering radius*. The covering radius [4] of a code is the minimum value d such that the union of the spheres of radius d around the points in the code cover the entire space $\mathbb{F}_q^n$. Suppose A had covering radius of at most $N - \alpha - \beta$. Then, no matter what B is (B is a single vector, since $\beta = 1$), it has distance at most $N - \alpha - \beta$ to some point $x_2 A$. Now suppose A has covering radius $d > N - \alpha - \beta$. Then there is some vector B where $\Delta(B, x_2 A) > N - \alpha - \beta$ for all x_2. Furthermore, any scalar multiple of B will also have distance at least d from any $x_2 A$ (to see this, note that if $\Delta(aB, x_2 A) < d$, then $\Delta((1/a)aB, (1/a)(x_2 A)) < d$, and so $\Delta(B, ((1/a)x_2)A) < d$, a contradiction).

Thus for $\beta = 1$, a construction of A with covering radius of at most $N - \alpha - \beta$ is necessary and sufficient for a negative result. Additionally, for $\beta > 1$, showing that there exists an $\alpha \times N$ matrix A with covering radius at most $N - \alpha - \beta$ is sufficient for a negative result.

Cohen and Frankl [5] gave upper bounds on the covering radius of linear codes over $\mathbb{F}_q$. Their analysis can be used to obtain the following result:[3]

Theorem 6 *For any value $1 \leq d \leq N$, there is a D-dimensional linear code over $\mathbb{F}_q$ with block length N (i.e. a vector subspace of $\mathbb{F}_q^N$) which has covering radius at most d, where*

$$D \equiv N - \frac{\log N}{\log q} - \frac{\log \mathrm{Vol}_q(d, N)}{\log q} + 2\log N + \log q + \log \ln q.$$

Combining Theorem 6 with the discussion at the beginning of this section gives Theorem 5.

6.2 Applying Theorem 5. We give one example here of how Theorem 5 can be applied. Other interesting examples are possible, but we omit them for space reasons.

Let τ be any constant satisfying $0 < \tau < 1/2$. Let $c = N^\tau$, let $k = \sigma c^2 \log N$ (we will specify σ shortly) and let $q = N^c$. By Fact 4(a), we can get a fairly good lower bound on $\mathrm{Vol}_q(k, N)$ just by considering the last term of the sum. We have $\mathrm{Vol}_q(k, N) \geq (q-1)^k \binom{N}{k} \geq (q/2)^k \left(\frac{N}{k}\right)^k$. We thus have that $\frac{\log \mathrm{Vol}_q(k,N)}{\log q} \geq k + \frac{-1 + k\log N - k\log k}{\log q}$. Plugging the above parameter settings for k and q into Equation (6) (but not substituting in yet for c), we have that Equation (6) is satisfied if

$$\begin{aligned}\sigma c^2 \log N + \beta \quad < \quad & \frac{1}{c} + \sigma c^2 \log N + \frac{-1 + \sigma c^2 \log^2 N - \sigma c^2 \log N \log(\sigma c^2 \log N)}{c \log N} \\ & -2\log N - c\log N - \log(c \ln N).\end{aligned}$$

[3] The expression for D in Theorem 6 is slightly different from the result as stated in [5]. Their analysis implicitly assumes that the field size q is independent of N; however this assumption need not hold for us. We give a complete derivation of Theorem 6 in the full version of the paper.

This inequality is equivalent to

$$\beta < \frac{1}{c} - \frac{1}{c\log N} + \sigma c \log N - \sigma c \log(\sigma c^2 \log N) - (c+2)\log N - \log(c \ln N).$$

Now since $c = N^\tau$, it can be verified that taking $\sigma = \frac{2}{1-2\tau}$ (a fixed constant since $0 < \tau < 1/2$ is a fixed constant) makes the right-hand side of this last inequality at least $(c-3)\log N$ (for sufficiently large N), so β can be any value smaller than this bound. This example shows that for a wide range of values of k the lower bound on field size required for this method of secure multicast linear network coding, if no capacity is given up, can be as large as $N^{\Omega(\sqrt{k/\log k})}$. It is interesting to contrast this lower bound with the upper bound of $\binom{N}{k}$ of Cai & Yeung.

7 Future Work

Several interesting directions for future research suggest themselves. Can quantitative improvements on our results for secure network coding be achieved, perhaps by studying nonlinear network coding schemes? Can network codes for information transmission problems other than multicast be made secure using our techniques? Another natural direction is to consider secure network coding in a framework where only statistical security or security against computationally bounded adversaries is required, as opposed to the information-theoretic security criterion we studied.

References

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

[2] N. Cai and R. W. Yeung. Secure network coding. In *International Symposium on Information Theory (ISIT '02)*, June 2002.

[3] P. A. Chou, Y. Wu, and K. Jain. Practical network coding. In *Proc. 41st Annual Allerton Conference on Communication, Control, and Computing*, October 2003.

[4] G. Cohen, I. Honkala, S. Litsyn, and A. Lobstein. *Covering Codes*. Elsevier, Amsterdam, 1997.

[5] G. D. Cohen and P. Frankl. Good coverings of Hamming spaces with spheres. *Discrete Mathematics*, 56:125–131, 1985.

[6] T. Ho, D. Karger, M. Médard, and R. Koetter. Network coding from a network flow perspective. In *International Symposium on Information Theory (ISIT '03)*, June 2003.

[7] T. Ho, B. Leong, R. Koetter, M. Médard, M. Effros, and D. R. Karger. Byzantine modification detection in multicast networks using randomized network coding. In *IEEE International Symposium on Information Theory (ISIT 2004)*, June 2004.

[8] W. C. Huffman and V. Pless. *Fundamentals of Error Correcting Codes*. Cambridge University Press, 2003.

[9] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen. Polynomial time algorithms for multicast network code construction. Submitted to IEEE Transactions on Information Theory, July 2003.

[10] K. Jain. Security based on network topology against the wiretapping attack. *IEEE Wireless Communications*, pages 68–71, February 2004.

[11] R. Koetter and M. Médard. An algebraic approach to network coding. *Transactions on Networking*, October 2003.

[12] A. R. Lehman and E. Lehman. Complexity classification of network information flow problems. In *Symposium on Discrete Algorithms (SODA '04)*, January 2004.

[13] S.-Y. R. Li, R. W. Yeung, and N. Cai. Linear network coding. *IEEE Transactions on Information Theory*, 49:371–381, February 2003.