On Field Size and Success Probability in Network Coding

Olav Geil[1], Ryutaroh Matsumoto[2], and Casper Thomsen[1]

[1] Department of Mathematical Sciences, Aalborg University, Denmark
olav@math.aau.dk,
caspert@math.aau.dk
[2] Department of Communications and Integrated Systems,
Tokyo Institute of Technology, Japan
ryutaroh@rmatsumoto.org

Abstract. Using tools from algebraic geometry and Gröbner basis theory we solve two problems in network coding. First we present a method to determine the smallest field size for which linear network coding is feasible. Second we derive improved estimates on the success probability of random linear network coding. These estimates take into account which monomials occur in the support of the determinant of the product of Edmonds matrices. Therefore we finally investigate which monomials can occur in the determinant of the Edmonds matrix.

Keywords: Distributed networking, linear network coding, multicast, network coding, random network coding.

1 Introduction

In a traditional data network, an intermediate node only forwards data and never modifies them. Ahlswede et al. [1] showed that if we allow intermediate nodes to process their incoming data and output modified versions of them then maximum throughput can increase, and they also showed that the maximum throughput is given by the minimum of maxflows between the source node and a sink node for single source multicast on an acyclic directional network. Such processing is called network coding. Li et al. [10] showed that computation of linear combinations over a finite field by intermediate nodes is enough for achieving the maximum throughput. Network coding only involving linear combinations is called linear network coding. The acyclic assumption was later removed by Koetter and Médard [9].

In this paper we shall concentrate on the error-free, delay-free multisource multicast network connection problem where the sources are uncorrelated. However, the proposed methods described can be generalized to deal with delays as in [7]. The only exception is the description in Section 7.

Considering multicast, it is important to decide whether or not all receivers (called sinks) can recover all the transmitted information from the senders (called

J. von zur Gathen, J.L. Imaña, and Ç.K. Koç (Eds.): WAIFI 2008, LNCS 5130, pp. 157–173, 2008.

sources). It is also important to decide the minimum size q of the finite field $\mathbf{F}_q$ required for linear network coding.

Before using linear network coding we have to decide coefficients in linear combinations computed by intermediate nodes. When the size q of a finite field is large, it is shown that random choice of coefficients allows all sinks to recover the original transmitted information with high probability [7]. Such a method is called random linear network coding and the probability is called success probability. As to random linear network coding the estimation or determination of the success probability is very important. Ho et al. [7] gave a lower bound on the success probability.

In their paper [9], Koetter and Médard introduced an algebraic geometric point view on network coding. As explained in [3], computational problems in algebraic geometry can often be solved by Gröbner bases. In this paper, we shall show that the exact computation of the minimum q can be made by applying the division algorithm for multivariate polynomials, and we will show that improved estimates for the success probability can be found by applying the footprint bound from Gröbner basis theory. These results introduce a new approach to network coding study. As the improved estimates take into account which monomials occur in the support of the determinant of a certain matrix [7] we study this matrix in details at the end of the paper.

2 Preliminary

We can determine whether or not all sinks can recover all the transmitted information by the determinant of some matrix [7]. We shall review the definition of such determinant. Let $G = (V, E)$ be an directed acyclic graph with possible parallel edges that represents the network topology. The set of source and sink nodes is denoted by S and T respectively. Assume that the source nodes S together get h symbols in $\mathbf{F}_q$ per unit time and try to send them.

Identify the edges in E with the integers $1, \ldots, |E|$. For an edge $j = (u, v)$ we write $\mathrm{head}(j) = v$ and $\mathrm{tail}(j) = u$. We define the $|E| \times |E|$ matrix $F = (f_{i,j})$ where $f_{i,j}$ is a variable if $\mathrm{head}(i) = \mathrm{tail}(j)$ and $f_{i,j} = 0$ otherwise. The variable $f_{i,j}$ is the coding coefficient from i to j.

Index h symbols in $\mathbf{F}_q$ sent by S by $1, \ldots, h$. We also define an $h \times |E|$ matrix $A = (a_{i,j})$ where $a_{i,j}$ is a variable if the edge j is an outgoing edge from the source $s \in S$ sending the i-th symbol and $a_{i,j} = 0$ otherwise. Variables $a_{i,j}$ represent how the source nodes send information to their outgoing edges.

Let $X(l)$ denote the l-th symbol generated by the sources S, and let $Y(j)$ denote the information sent along edge j. The model is described by the following relation

$$Y(j) = \sum_{i=1}^{h} a_{i,j} X(i) + \sum_{i:\mathrm{head}(i)=\mathrm{tail}(j)} f_{i,j} Y(i) \ .$$

For each sink $t \in T$ define an $h \times |E|$ matrix B_t whose (i, j) entry $b_{t,i,j}$ is a variable if $\mathrm{head}(j) = t$ and equals 0 otherwise. The index i refers to the i-th

symbol sent by one of the sources. Thereby variables $b_{t,i,j}$ represent how the sink t process the received data from its incoming edges.

The sink t records the vector

$$\boldsymbol{b^{(t)}} = \left(b_1^{(t)}, \ldots, b_h^{(t)}\right)$$

where

$$b_i^{(t)} = \sum_{j:\mathrm{head}(j)=t} b_{t,i,j} Y(j) \; .$$

We now recall from [7] under which conditions all informations sent by the sources can always be recovered at all sinks. As in [7] we define the Edmonds matrix M_t for $t \in T$ by

$$M_t = \begin{pmatrix} A & 0 \\ I - F & B_t^T \end{pmatrix} \; . \tag{1}$$

Define the polynomial P by

$$P = \prod_{t \in T} |M_t| \; . \tag{2}$$

P is a multivariate polynomial in variables $f_{i,j}$, $a_{i,j}$ and $b_{t,i,j}$. Assigning a value in $\mathbf{F}_q$ to each variable corresponds to choosing a coding scheme. Plugging the assigned values into P gives an element $k \in \mathbf{F}_q$. The following theorem from [7] tells us when the coding scheme can be used to always recover the information generated at the sources S at all sinks in T.

Theorem 1. *Let the notation and the network coding model be as above. Assume a coding scheme has been chosen by assigning values to the variables $f_{i,j}$, $a_{i,j}$ and $b_{t,i,j}$. Let k be the value found by plugging the assigned values into P. Every sink $t \in T$ can recover from $\boldsymbol{b^{(t)}}$ the informations $X(1), \ldots, X(h)$ no matter what they are, if and only if $k \neq 0$ holds.*

Proof. See [7]. □

3 Computation of the Minimum Field Size

We shall study computation of the minimum symbol size q. For this purpose we will need the division algorithm for multivariate polynomials [3, Sec. 2.3] to produce the remainder of a polynomial $F(X_1, \ldots, X_n)$ modulo $(X_1^q - X_1, \ldots, X_n^q - X_n)$ (this remainder is independent of the choice of monomial ordering). We adapt the standard notation for the above remainder which is

$$F(X_1, \ldots, X_n) \text{ rem } (X_1^q - X_1, \ldots, X_n^q - X_n) \; .$$

The reader unfamiliar with the division algorithm can think of the above remainder of $F(X_1, \ldots, X_n)$ as the polynomial produced by the following procedure. As long as we can find an X_i such that X_i^q divides some term in the polynomial under consideration we replace the factor X_i^q with X_i wherever it occurs. The process continues until the X_i-degree is less than q for all $i = 1, \ldots, n$. It is clear that the above procedure can be efficiently implemented.

Proposition 1. *Let $F(X_1, \ldots, X_n)$ be an n-variate polynomial over $\mathbf{F}_q$. There exists an n-tuple $(x_1, \ldots, x_n) \in \mathbf{F}_q^n$ such that $F(x_1, \ldots, x_n) \neq 0$ if and only if*

$$F(X_1, \ldots, X_n) \text{ rem } (X_1^q - X_1, \ldots, X_n^q - X_n) \neq 0 \, .$$

Proof. As $a^q = a$ for all $a \in \mathbf{F}_q$ it holds that $F(X_1, \ldots, X_n)$ evaluates to the same as $R(X_1, \ldots, X_n) := F(X_1, \ldots, X_n)$ rem $(X_1^q - X_1, \ldots, X_n^q - X_n)$ in every $(x_1, \ldots, x_n) \in \mathbf{F}_q^n$. If $R(X_1, \ldots, X_n) = 0$ therefore $F(X_1, \ldots, X_n)$ evaluates to zero for every choice of $(x_1, \ldots, x_n) \in \mathbf{F}_q^n$. If $R(X_1, \ldots, X_n)$ is nonzero we consider it first as a polynomial in $\mathbf{F}_q(X_1, \ldots, X_{n-1})[X_n]$ (that is, a polynomial in one variable over the quotient field $\mathbf{F}_q(X_1, \ldots, X_{n-1})$). But the X_n-degree is at most $q-1$ and therefore it has at most $q-1$ zeros. We conclude that there exists an $x_n \in \mathbf{F}_q$ such that $R(X_1, \ldots, X_{n-1}, x_n) \in \mathbf{F}_q[X_1, \ldots, X_{n-1}]$ is nonzero. Continuing this way we find $(x_1, \ldots, x_n)$ such that $R(x_1, \ldots, x_n)$ and therefore also $F(x_1, \ldots, x_n)$ is nonzero. □

From [7, Th. 2] we know that for all prime powers q greater than $|T|$ linear network coding is possible. It is now straightforward to describe an algorithm that finds the smallest field $\mathbf{F}_q$ of prescribed characteristic p for which linear network coding is feasible. We first reduce the polynomial P from (2) modulo the prime p. We observe that although P is a polynomial in all the variables $a_{i,j}$, $b_{t,i,j}$, $f_{i,j}$ the variable $b_{t,i,j}$ appears at most in powers of 1. This is so as it appears at most in a single entry in M_t and does not appear elsewhere. Therefore $\mathbf{F}_q$ can be used for network coding if P rem p does not reduce to zero modulo the polynomials $a_{i,j}^q - a_{i,j}$, $f_{i,j}^q - f_{i,j}$. To decide the smallest field $\mathbf{F}_q$ of characteristic p for which network coding is feasible we try first $\mathbf{F}_q = \mathbf{F}_p$. If this does not work we then try $\mathbf{F}_{p^2}$ and so on. To find an $\mathbf{F}_q$ that works we need at most to try $\lfloor \log_p(|T|) \rfloor$ different fields as we know that linear network coding is possible whenever $q > |T|$.

Note that once a field $\mathbf{F}_q$ is found such that the network connection problem is feasible the last part of the proof of Proposition 1 describes a simple way of deciding coefficients $(x_1, \ldots, x_n) \in \mathbf{F}_q^n$ that can be used for network coding.

From [4, Sec. 7.1.3] we know that it is an NP-hard problem to find the minimum field size for linear network coding. Our findings imply that it is NP-hard to find the polynomial P in (2).

4 Computation of the Success Probability of Random Linear Network Coding

In random linear network coding we from the beginning fix for a collection

$$K \subseteq \{1, \ldots, h\} \times \{1, \ldots, |E|\}$$

the $a_{i,j}$'s with $(i,j) \in K$ and also we fix for a collection

$$J \subseteq \{1, \ldots, |E|\} \times \{1, \ldots, |E|\}$$

the $f_{i,j}$'s with $(i,j) \in J$. This is done in a way such that there exists a solution to the network connection problem with the same values for these fixed coefficients. A priori of course we let $a_{i,j} = 0$ if the edge j is not emerging from the source sending information i, and also a priori we of course let $f_{i,j} = 0$ if j is not an adjacent downstream edge of i. Besides these a priori fixed values there may be good reasons for also fixing other coefficients $a_{i,j}$ and $f_{i,j}$ [7]. If for example there is only one upstream edge i adjacent to j we may assume $f_{i,j} = 1$. All the $a_{i,j}$'s and $f_{i,j}$'s which have not been fixed at this point are then chosen randomly and independently. All coefficients are to be elements in $\mathbf{F}_q$. If a solution to the network connection problem exists with the $a_{i,j}$'s and the $f_{i,j}$'s specified, it is possible to determine values of $b_{t,i,j}$ at the sinks such that a solution to the network connection problem is given. Let μ be the number of variables $a_{i,j}$ and $f_{i,j}$ chosen randomly. Call these variables $X_1, \ldots, X_\mu$. Consider the polynomial P in (2) and let $\widetilde{P}$ be the polynomial made from P by plugging in the fixed values of the $a_{i,j}$'s and the fixed values of the $f_{i,j}$'s (calculations taking place in $\mathbf{F}_q$). Then $\widetilde{P}$ is a polynomial in $X_1, \ldots, X_\mu$. The coefficients of $\widetilde{P}$ are polynomials in the $b_{t,i,j}$'s over $\mathbf{F}_q$. Finally, define

$$\widehat{P} := \widetilde{P} \text{ rem } (X_1^q - X_1, \ldots, X_\mu^q - X_\mu) \, .$$

The success probability of random linear network coding is the probability that the random choice of coefficients will lead to a solution of the network connection problem[1] as in Section 2. That is, the probability is the number

$$\begin{aligned} &\frac{|\{(x_1, \ldots, x_\mu) \in \mathbf{F}_q^\mu | \widetilde{P}(x_1, \ldots, x_\mu) \neq 0\}|}{q^\mu} \\ &= \frac{|\{(x_1, \ldots, x_\mu) \in \mathbf{F}_q^\mu | \widehat{P}(x_1, \ldots, x_\mu) \neq 0\}|}{q^\mu} \, . \end{aligned} \tag{3}$$

To see the first result observe that for fixed $(x_1, \ldots, x_\mu) \in \mathbf{F}_q^\mu$, $\widetilde{P}(x_1, \ldots, x_\mu)$ can be viewed as a polynomial in the variables $b_{t,i,j}$'s with coefficients in $\mathbf{F}_q$ and recall that the $b_{t,i,j}$'s occur in powers of at most 1. Therefore, if $\widetilde{P}(x_1, \ldots, x_\mu) \neq 0$, then by Proposition 1 it is possible to choose the $b_{t,i,j}$'s such that if we plug them into $\widetilde{P}(x_1, \ldots, x_\mu)$ then we get nonzero. The last result follows from the fact that $\widetilde{P}(x_1, \ldots, x_\mu) = \widehat{P}(x_1, \ldots, x_\mu)$ for all $(x_1, \ldots, x_\mu) \in \mathbf{F}_q^\mu$. In this section we shall present a method to estimate the success probability using Gröbner basis theoretical methods.

We briefly review some basic definitions and results of Gröbner bases. See [3] for a more detailed exposition. Let $\mathcal{M}(X_1, \ldots, X_n)$ be the set of monomials in the variables $X_1, \ldots, X_n$. A monomial ordering $\prec$ is a total ordering on $\mathcal{M}(X_1, \ldots, X_n)$ such that

$$L \prec M \Longrightarrow LN \prec MN$$

[1] This corresponds to saying that each sink can recover the data at the maximum rate promised by network coding.

holds for all monomials $L, M, N \in \mathcal{M}(X_1, \ldots, X_n)$ and such that every nonempty subset of $\mathcal{M}(X_1, \ldots, X_n)$ has a unique smallest element with respect to $\prec$. The leading monomial of a polynomial F with respect to $\prec$, denoted by $\mathrm{LM}(F)$, is the largest monomial in the support of F. Given a polynomial ideal I and a monomial ordering the footprint $\Delta_\prec(I)$ is the set of monomials that cannot be found as leading monomials of any polynomial in I. The following proposition explains our interest in the footprint (for a proof of the proposition see [2, Pro. 8.32]).

Proposition 2. *Let* $\mathbf{F}$ *be a field and consider the polynomials* $F_1, \ldots, F_s \in \mathbf{F}[X_1, \ldots, X_n]$. *Let* $I = \langle F_1, \ldots, F_s \rangle \subseteq \mathbf{F}[X_1, \ldots, X_n]$ *be the ideal generated by* $F_1, \ldots, F_s$. *If* $\Delta_\prec(I)$ *is finite then the number of common zeros of* $F_1, \ldots, F_s$ *in the algebraic closure of* $\mathbf{F}$ *is at most equal to* $|\Delta_\prec(I)|$.

Proposition 2 is known as the footprint bound. It has the following corollary.

Corollary 1. *Let* $F \in \mathbf{F}[X_1, \ldots, X_n]$ *where* $\mathbf{F}$ *is a field containing* $\mathbf{F}_q$. *Fix a monomial ordering and let*

$$X_1^{j_1} \cdots X_n^{j_n} = \mathrm{LM}\big(F \text{ rem } (X_1^q - X_1, \ldots, X_n^q - X_n)\big) \, .$$

The number of zeros of F *over* $\mathbf{F}_q$ *is at most equal to*

$$q^n - \prod_{v=1}^{n} (q - j_v) \, . \tag{4}$$

Proof. We have

$$\begin{aligned} &\Delta_\prec(\langle F, X_1^q - X_1, \ldots, X_n^q - X_n \rangle) \\ &\qquad \subseteq \Delta_\prec(\langle \mathrm{LM}(F \text{ rem } (X_1^q - X_1, \ldots, X_n^q - X_n)), X_1^q, \ldots, X_n^q \rangle) \end{aligned}$$

and the size of the latter set equals (4). The result now follows immediately from Proposition 2. □

Theorem 2. *Let as above* $\widetilde{P}$ *be found by plugging into* P *some fixed values for the variables* $a_{i,j}$, $(i,j) \in K$, *and by plugging into* P *some fixed values for the variables* $f_{i,j}$, $(i,j) \in J$, *and by leaving the remaining* μ *variables flexible. Assume as above that there exists a solution to the network connection problem with the same values for these fixed coefficients. Denote by* $X_1, \ldots, X_\mu$ *the variables to be chosen by random and define* $\widehat{P} := \widetilde{P} \text{ rem } (X_1^q - X_1, \ldots, X_\mu^q - X_\mu)$. *(Note that if* $q > |T|$ *then* $\widehat{P} = \widetilde{P}$*). Consider* $\widehat{P}$ *as a polynomial in the variables* $X_1, \ldots, X_\mu$ *and let* $\prec$ *be any fixed monomial ordering. Writing* $X_1^{j_1} \cdots X_\mu^{j_\mu} = \mathrm{LM}(\widehat{P})$ *the success probability is at least*

$$q^{-\mu} \prod_{v=1}^{\mu} (q - j_v) \, . \tag{5}$$

As a consequence the success probability is in particular at least

$$q^{-\mu} \min \left\{ \prod_{i=1}^{\mu} (q - s_i) \middle| X_1^{s_1} \cdots X_\mu^{s_\mu} \text{ is a monomial in the support of } \widehat{P} \right\} . \tag{6}$$

Proof. Let $\mathbf{F}$ be the quotient field $\mathbf{F}_q(X_1, \ldots, X_\mu)$. The result in (5) now follows by applying Corollary 1 and (3). As the leading monomial of $\widetilde{P}$ is of course a monomial in the support of $\widetilde{P}$ (6) is smaller or equal to (5). □

Remark 1. The condition in Theorem 2 that there exists a solution to the network connection problem with the coefficients corresponding to K and J being as specified is equivalent to the condition that $\widehat{P} \neq 0$.

We conclude this section by mentioning without a proof that Gröbner basis theory tells us that the true success probability can be calculated as

$$q^{-\mu}\left(q^\mu - |\Delta_\prec(\langle \widetilde{P}, X_1^q - X_1, \ldots, X_\mu^q - X_\mu\rangle)|\right) .$$

This observation is however of little value as it seems very difficult to compute the footprint

$$\Delta_\prec(\langle \widetilde{P}, X_1^q - X_1, \ldots, X_\mu^q - X_\mu\rangle)$$

due to the fact that μ is typically a very high number.

5 The Bound by Ho et al.

In [7] Ho et al. gave a lower bound on the success probability in terms of the number of edges j with associated random coefficients[2] $\{a_{i,j}, f_{l,j}\}$. Letting η be the number of such edges [7, Th. 2] tells us that if $q > |T|$ and if there exists a solution to the network connection problem with the same values for the fixed coefficients, then the success probability is at least

$$p_{\text{Ho}} = \left(\frac{q - |T|}{q}\right)^\eta . \tag{7}$$

The proof in [7] of (7) relies on two lemmas of which we only state the first one.

Lemma 1. *Let η be defined as above. The determinant polynomial of M_t has maximum degree η in the random variables $\{a_{i,j}, f_{l,j}\}$ and is linear in each of these variables.*

Proof. See [7, Lem. 3]. Alternatively the proof can be derived as a consequence of Theorem 3 in Section 7. □

Recall, that the polynomial P in (2) is the product of the determinants $|M_t|$, $t \in T$. Lemma 1 therefore implies that the polynomial $\widetilde{P}$ has at most total degree equal to $|T|\eta$ and that no variable appears in powers of more than $|T|$. The assumption $q > |T|$ implies $\widehat{P} = \widetilde{P}$ which makes it particular easy to see that the same of course holds for $\widehat{P}$. Combining this observation with the following lemma shows that the numbers in (5) and (6) are both at least as large as the number (7).

[2] We state Ho et al.'s bound only in the case of delay-free acyclic networks.

Lemma 2. *Let $\eta, |T|, q \in \mathbf{N}$, $|T| < q$ be some fixed numbers. Let $\mu, x_1, \ldots, x_\mu \in \mathbf{N}_0$ satisfy*

$$0 \leq x_1 \leq |T|, \ldots, 0 \leq x_\mu \leq |T|$$

and $x_1 + \cdots + x_\mu \leq |T|\eta$. The minimal value of

$$\prod_{i=1}^{\mu} \left(\frac{q - x_i}{q} \right)$$

(taken over all possible values of $\mu, x_1, \ldots, x_\mu$) is

$$\left(\frac{q - |T|}{q} \right)^{\eta}.$$

Proof. Assume μ and $x_1, \ldots, x_\mu$ are chosen such that the expression attains its minimal value. Without loss of generality we may assume that

$$x_1 \geq x_2 \geq \cdots \geq x_\mu$$

holds. Clearly, $x_1 + \cdots + x_\mu = |T|\eta$ must hold. If $x_i < |T|$ and $x_{i+1} > 0$ then

$$(q - x_i)(q - x_{i+1}) > (q - (x_i + 1))(q - (x_{i+1} - 1))$$

which cannot be the case. So $x_1 = \cdots = x_\eta = |T|$. The remaining x_j's if any all equal zero. □

6 Examples

In this section we apply the methods from the previous sections to two concrete networks. We will see that the estimate on the success probability of random linear network coding that was described in Theorem 2 can be considerably better than the estimate described in [7, Th. 2]. Also we will apply the method from Section 3 to determine the smallest field of characteristic two for which network coding can be successful.

As random linear network coding is assumed to take place at the nodes in a decentralized manner, one natural choice is to set $f_{i,j} = 1$ whenever the indegree of the end node of edge i is one and j is the downstream edge adjacent to i. Clearly, if j is not a downstream edge adjacent to i we set $f_{i,j} = 0$. Whenever none of the above is the case we may choose $f_{i,j}$ randomly. Also if there is only one source and the outdegree of the source is equal to the number of symbols to be send we may enumerate the edges from the source by the numbers $1, \ldots, h$ and set $a_{i,j} = 1$ if $1 \leq i = j \leq h$ and set $a_{i,j} = 0$ otherwise. This strategy can be generalized also to deal with the case of more sources. In the following two examples we will choose the variables in the manner just described. The network in the first example is taken from [4, Ex. 3.1] whereas the network in the second example is new.

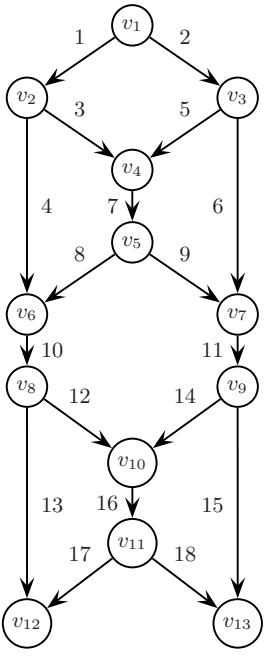

Fig. 1. The network from Example 1

Example 1. Consider the delay-free and acyclic network in Figure 1. There is one sender v_1 and two receivers v_{12} and v_{13}. The min-cut max-flow number is two for both receivers so we assume that two independent random processes emerge from sender v_1. We consider in this example only fields of characteristic 2. Following the description preceding the example we set $a_{1,1} = a_{2,2} = 1$ and $a_{i,j} = 0$ in all other cases. Also we let $f_{i,j} = 1$ except

$$f_{3,7}, f_{5,7}, f_{4,10}, f_{8,10}, f_{9,11}, f_{6,11}, f_{12,16}, f_{14,16}$$

which we choose by random. As in the previous sections we consider $b_{t,i,j}$ as fixed but unknown to us. The determinant polynomial becomes

$$\widetilde{P} = (b^2c^2e^2gh + c^2f^2gh + a^2d^2f^2gh)Q \ ,$$

where

$$\begin{array}{llll} a = f_{3,7} & b = f_{5,7} & c = f_{4,10} & d = f_{8,10} \\ e = f_{9,11} & f = f_{6,11} & g = f_{12,16} & h = f_{14,16} \end{array}$$

and $Q = |B'_{v_{12}}||B'_{v_{13}}|$. Here, $B'_{v_{12}}$ respectively $B'_{v_{13}}$ is the matrix consisting of the nonzero columns of $B_{v_{12}}$ respectively the nonzero columns of $B_{v_{14}}$. Restricting to fields $\mathbf{F}_q$ of size at least 4 we have $\widehat{P} = \widetilde{P}$ and we can therefore immediately

apply the bounds in Theorem 2. Applying (6) we get the following lower bound on the success probability

$$P_{\text{new }2}(q) = \frac{(q-2)^3(q-1)^2}{q^5} \ .$$

Choosing as monomial ordering the lexicographic ordering $\prec_{\text{lex}}$ with

$$a \prec_{\text{lex}} b \prec_{\text{lex}} d \prec_{\text{lex}} e \prec_{\text{lex}} g \prec_{\text{lex}} h \prec_{\text{lex}} f \prec_{\text{lex}} c$$

the leading monomial of $\widetilde{P}$ becomes c^2f^2gh and therefore from (5) we get the following lower bound on the success probability

$$P_{\text{new }1}(q) = \frac{(q-2)^2(q-1)^2}{q^4} \ .$$

For comparison the bound (7) from [7] states that the success probability is at least

$$P_{\text{Ho}}(q) = \frac{(q-2)^4}{q^4} \ .$$

We see that $P_{\text{new }1}$ exceeds P_{Ho} with a factor $(q-1)^2/(q-2)^2$, which is larger than 1. Also $P_{\text{new }2}$ exceeds P_{Ho}. In Table 1 we list values of $P_{\text{new }1}(q)$, $P_{\text{new }2}(q)$ and $P_{\text{Ho}}(q)$ for various choices of q.

Table 1. From Example 1: Estimates on the success probability

q	4	8	16	32	64
$P_{\text{new }1}(q)$	0.140	0.430	0.672	0.893	0.909
$P_{\text{new }2}(q)$	0.703×10^{-1}	0.322	0.588	0.773	0.880
$P_{\text{Ho}}(q)$	0.625×10^{-1}	0.316	0.586	0.772	0.880

We next consider the field $\mathbf{F}_2$. We reduce $\widetilde{P}$ modulo $(a^2-a, \ldots, h^2-h)$ to get

$$\widehat{P} = (bcegh + cfgh + adfgh)Q \ .$$

From (6) we see that the success probability of random network coding is at least 2^{-5}. Choosing as monomial ordering the lexicographic ordering described above (5) tells us that the success probability is at least 2^{-4}. For comparison the bound (7) does not apply as we do not have $q > |T|$. It should be mentioned that for delay-free acyclic networks the network coding problem is solvable for all choices of $q \geq |T|$ [8] and [11]. From this fact one can only conclude that the success probability is at least 2^{-8} (8 being the number of coefficients to be chosen by random).

Example 2. Consider the network in Figure 2. The sender v_1 generates 3 independent random processes. The vertices v_{11}, v_{12} and v_{13} are the receivers.

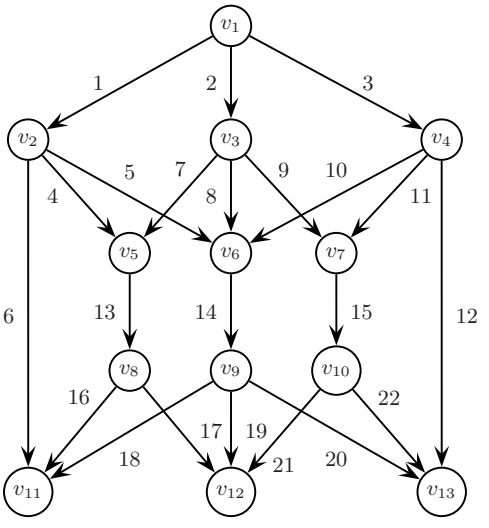

Fig. 2. The network from Example 2

We will apply network coding over various fields of characteristic two. We start by considering random linear network coding over fields of size at least 4. As $4 > |T| = 3$ we know that this can be done successfully.

We set $a_{1,1} = a_{2,2} = a_{3,3} = 1$ and $a_{i,j} = 0$ in all other cases. We let $f_{i,j} = 1$ except $f_{4,13}, f_{7,13}, f_{5,14}, f_{8,14}, f_{10,14}, f_{9,15}, f_{11,15}$, which we choose by random. As in the last section we consider $b_{t,i,j}$ as fixed but unknown to us. Therefore $\widetilde{P} = \widehat{P}$ is a polynomial in the seven variables $f_{4,13}, f_{7,13}, f_{5,14}, f_{8,14}, f_{10,14}, f_{9,15}, f_{11,15}$. The determinant polynomial becomes

$$\widehat{P} = (abcdefg + abce^2f^2 + b^2c^2efg)Q \ ,$$

where

$$a = f_{4,13} \qquad b = f_{5,14} \qquad c = f_{7,13} \qquad d = f_{8,14}$$
$$e = f_{9,15} \qquad f = f_{10,14} \qquad g = f_{11,15}$$

and $Q = |B'_{v_{11}}|\,|B'_{v_{12}}|\,|B'_{v_{13}}|$. Here, $B'_{v_{11}}$ respectively $B'_{v_{12}}$ respectively $B'_{v_{13}}$ is the matrix consisting of the nonzero columns of $B_{v_{11}}$ respectively the nonzero columns of $B_{v_{12}}$ respectively the nonzero columns of $B_{v_{13}}$. Choosing a lexicographic ordering with d being larger than the other variables and applying (5) we get that the success probability is at least

$$P_{\text{new }1}(q) = \frac{(q-1)^7}{q^7} \ .$$

Applying (6) we see that the success probability is at least

$$P_{\text{new }2}(q) = \frac{(q-1)^3(q-2)^2}{q^5}.$$

For comparison (7) tells us that success probability is at least

$$P_{\mathrm{Ho}}(q) = \frac{(q-3)^3}{q^3} .$$

Both bound (5) and bound (6) exceed (7) for all values of $q \geq 4$. In Table 2 we list $P_{\mathrm{new}\ 1}(q)$, $P_{\mathrm{new}\ 2}(q)$ and $P_{\mathrm{Ho}}(q)$ for various values of q.

Table 2. From Example 2: Estimates on the success probability

q	4	8	16	32	64
$P_{\mathrm{new}\ 1}(q)$	0.133	0.392	0.636	0.800	0.895
$P_{\mathrm{new}\ 2}(q)$	0.105	0.376	0.630	0.799	0.895
$P_{\mathrm{Ho}}(q)$	0.156×10^{-1}	0.244	0.536	0.744	0.865

We next consider the field $\mathbf{F}_2$. We reduce $\widetilde{P}$ modulo $(a^2-a, \ldots, g^2-g)$ to get

$$\widehat{P} = (abcdefg + abcef + bcefg)Q .$$

From (6) we see that the success probability of random network coding is at least 2^{-7}. Choosing a proper monomial ordering we get from (5) that the success probability is at least 2^{-5}. For comparison neither [7], [8], nor [11] tells us that linear network coding is possible.

7 The Topological Meaning of $|M_t|$

Recall from Section 5 that Ho et al.'s bound (7) relies on the rather rough Lemma 1. The following theorem gives a much more precise description of which monomials can occur in the support of P and $\widetilde{P}$ by explaining exactly which monomials can occur in $|M_t|$. Thereby the theorem gives some insight into when the bounds (5) and (6) are much better than the bound (7). The theorem states that if K is a monomial in the support of $|M_t|$ then it is the product of $a_{i,j}$'s, $f_{i,j}$'s and $b_{t,i,j}$'s related to h edge disjoint paths $P_1, \ldots, P_h$ that originate in the senders and end in receiver t.

Theorem 3. *Consider a delay-free acyclic network. If K is a monomial in the support of the determinant of M_t then it is of the form $K_1 \cdots K_h$ where*

$$K_u = a_{u,l_1^{(u)}} f_{l_1^{(u)},l_2^{(u)}} f_{l_2^{(u)},l_3^{(u)}} \cdots f_{l_{s_u-1}^{(u)},l_{s_u}^{(u)}} b_{t,v_u,l_{s_u}^{(u)}}$$

for $u = 1, \ldots, h$. Here, $\{v_1, \ldots, v_h\} = \{1, \ldots, h\}$ holds and $l_1^{(1)}, \ldots, l_h^{(h)}$ respectively $l_{s_1}^{(1)}, \ldots, l_{s_h}^{(h)}$ are pairwise different. Further

$$f_{l_i^{(u_1)},l_{i+1}^{(u_1)}} \neq f_{l_j^{(u_2)},l_{j+1}^{(u_2)}}$$

unless $u_1 = u_2$ and $i = j$ hold. In other words K corresponds to a product of h edge disjoint paths.

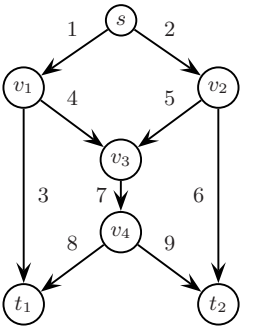

Fig. 3. The butterfly network

Proof. A proof can be found in the appendix. □

We illustrate the theorem with an example.

Example 3. Consider the butterfly network in Figure 3. A monomial K is in the support of $|M_{t_1}|$ if and only if it is in the support of the determinant of

$$N_{t_1} = (n_{i,j}) = \begin{bmatrix} I + F & B_{t_1}^T \\ A & 0 \end{bmatrix}$$

$$= \begin{bmatrix} 1 & 0 & f_{1,3} & f_{1,4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & f_{2,5} & f_{2,6} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & b_{t_1,1,3} & b_{t_1,2,3} \\ 0 & 0 & 0 & 1 & 0 & 0 & f_{4,7} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & f_{5,7} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & f_{7,8} & f_{7,9} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & b_{t_1,1,8} & b_{t_1,2,8} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ a_{1,1} & a_{1,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ a_{2,1} & a_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

By inspection we see that the monomial

$$K = a_{1,1}a_{2,2}b_{t_1,1,3}b_{t_1,2,8}f_{7,8}f_{5,7}f_{2,5}f_{1,3}$$

is in the support of $|N_{t_1}|$. We can write $K = K_1K_2$ where

$$K_1 = a_{1,1}f_{1,3}b_{t_1,1,3} \quad \text{and} \quad K_2 = a_{2,2}f_{2,5}f_{5,7}f_{7,8}b_{t_1,2,8}\ .$$

This is the description guaranteed by Theorem 3. To make it easier for the reader to follow the proof of Theorem 3 in the appendix we now introduce some of the notations to be used there. By inspection the monomial K can be written

$$K = \prod_{i=1}^{11} n_{i,p(i)}$$

where the permutation p is given by

$$p(1) = 3 \quad p(2) = 5 \quad p(3) = 10 \quad p(4) = 4 \quad p(5) = 7 \quad p(6) = 6$$
$$p(7) = 8 \quad p(8) = 11 \quad p(9) = 9 \quad p(10) = 1 \quad p(11) = 2$$

Therefore if we index the elements in $\{1, \ldots, 11\}$ by

$$i_1 = 10 \quad i_2 = 1 \quad i_3 = 3 \quad i_4 = 11 \quad i_5 = 2 \quad i_6 = 5$$
$$i_7 = 7 \quad i_8 = 8 \quad i_9 = 4 \quad i_{10} = 6 \quad i_{11} = 9$$

then we can write

$$K_1 = n_{i_1,p(i_1)} n_{i_2,p(i_2)} n_{i_3,p(i_3)}$$
$$K_2 = n_{i_4,p(i_4)} n_{i_5,p(i_5)} n_{i_6,p(i_6)} n_{i_7,p(i_7)} n_{i_8,p(i_8)}$$

and we have

$$n_{i_9,p(i_9)} = n_{i_{10},p(i_{10})} = n_{i_{11},p(i_{11})} = 1$$

corresponding to the fact $p(i_9) = i_9$, $p(i_{10}) = i_{10}$ and $p(i_{11}) = i_{11}$.

Remark 2. The procedures described in the proof of Theorem 3 can be reversed. This implies that there is a bijective map between the set of edge disjoint paths $P_1, \ldots, P_h$ in Theorem 3 and the set of monomials in $|M_t|$.

Theorem 3 immediately applies to the situation of random network coding if we plug into the $a_{i,j}$'s and into the $f_{t,i,j}$'s on the paths $P_1, \ldots, P_h$ the fixed values wherever such are given. Let as in Lemma 1 η be the number of edges for which some coefficients $a_{i,j}, f_{i,j}$ are to be chosen by random. Considering the determinant as a polynomial in the variables to be chosen by random with coefficients in the field of rational expressions in the $b_{t,i,j}$'s we see that no monomial can contain more than η variables and that no variable occurs more than once. This is because the paths $P_1, \ldots, P_h$ are edge disjoint. Hence, Lemma 1 is a consequence of Theorem 3.

Acknowledgments

The authors would like to thank the anonymous referees for their helpful suggestions.

References

1. Ahlswede, R., Cai, N., Li, S.-Y.R., Yeung, R.W.: Network information flow. IEEE Transactions on Information Theory 46(4), 1204–1206 (2000)
2. Becker, T., Weispfenning, V.: Gröbner Bases - A Computational Approach to Commutative Algebra. Springer, Berlin (1993)
3. Cox, D., Little, J., O'Shea, D.: Ideals, Varieties, and Algorithms, 2nd edn. Springer, Berlin (1996)

4. Fragouli, C., Soljanin, E.: Network Coding Fundamentals. In: Foundations and Trends in Networking, vol. 2(1). now Publishers Inc., Hanover (2007)
5. Geil, O.: On codes from norm-trace curves. Finite Fields and their Applications 9(3), 351–371 (2003)
6. Ho, T., Karger, D.R., Médard, M., Koetter, R.: Network Coding from a Network Flow Perspective. In: Proceedings. IEEE International Symposium on Information Theory, Yokohama, Japan, July 2003, p. 441 (2003)
7. Ho, T., Médard, M., Koetter, R., Karger, D.R., Effros, M., Shi, J., Leong, B.: A Random Linear Network Coding Approach to Multicast. IEEE Transactions on Information Theory 52(10), 4413–4430 (2006)
8. Jaggi, S., Chou, P.A., Jain, K.: Low Complexity Algebraic Multicast Network Codes. In: Proceedings. IEEE International Symposium on Information Theory, Yokohama, Japan, July 2003, p. 368 (2003)
9. Koetter, R., Médard, M.: An Algebraic Approach to Network Coding. IEEE/ACM Transactions on Networking 11(5), 782–795 (2003)
10. Li, S.-Y.R., Yeung, R.W., Cai, N.: Linear Network Coding. IEEE Transactions on Information Theory 49(2), 371–381 (2003)
11. Sanders, P., Egner, S., Tolhuizen, L.: Polynomial Time Algorithms for Network Information Flow. In: Proceedings of the 15th ACM Symposium on Parallel Algorithms, San Diego, USA, June 2003, pp. 286–294 (2003)

A Proof of Theorem 3

The proof of Theorem 3 calls for the following technical lemma.

Lemma 3. *Consider a delay-free acyclic network with corresponding matrix F as in Section 2. Let I be the $|E| \times |E|$ identity matrix and define*

$$\Gamma = (\gamma_{i,j}) = I + F \, .$$

Given a permutation p on $\{1, \ldots, |E|\}$ write

$$p^{(i)}(\lambda) = \overbrace{p(p(\cdots(\lambda)\cdots))}^{i \text{ times}}$$

If for some $\lambda \in \{1, \ldots, |E|\}$ the following hold

(1) $\lambda, p(\lambda), \ldots, p^{(x)}(\lambda)$ *are pairwise different*
(2) $p^{(x+1)}(\lambda) \in \{\lambda, p(\lambda), \ldots, p^{(x)}(\lambda)\}$
(3) $\gamma_{\lambda,p(\lambda)}, \gamma_{p(\lambda),p(p(\lambda))}, \ldots, \gamma_{p^{(x)}(\lambda),p^{(x+1)}(\lambda)}$ *are all nonzero*

then $x = 0$.

Proof. Let p be a permutation and let x and λ be numbers such that (1), (2) and (3) hold. As p is a permutation then (1) and (2) implies that $p(p^{(x)}(\lambda)) = \lambda$. Aiming for a contradiction assume $x > 0$. As $p(\eta) = \eta$ does not hold for any $\eta \in \{\lambda, p(\lambda), \ldots, p^{(x)}(\lambda)\}$,

$$\gamma_{\lambda,p(\lambda)}, \gamma_{p(\lambda),p^{(2)}(\lambda)}, \ldots, \gamma_{p^{(x)}(\lambda),p^{(x+1)}(\lambda)}$$

are all non-diagonal elements in $I + F$. By (3) we therefore have constructed a cycle in a cycle-free graph and the assumption $x > 0$ cannot be true. □

Proof (of Theorem 3). A monomial is in the support of the determinant of M_t if and only if it is in the support of the determinant of

$$N_t = \begin{pmatrix} I+F & B_t^T \\ A & 0 \end{pmatrix} = (n_{i,j}) \, .$$

To ease the notation in the present proof we consider the latter matrix. Let p be a permutation on $\{1, \ldots, |E|+h\}$ such that

$$\prod_{s=1}^{|E|+h} n_{s,p(s)} \neq 0 \, . \tag{8}$$

Below we order the elements in $\{1, \ldots, |E|+h\}$ in a particular way by indexing them $i_1, \ldots, i_{|E|+h}$ according to the following set of procedures.

Let $i_1 = |E|+1$ and define recursively

$$i_s = p(i_{s-1})$$

until $|E| < p(i_s) \leq |E|+h$. Note that this must eventually happen due to Lemma 3. Let s_1 be the (smallest) number such that $|E| < p(i_{s_1}) \leq |E|+h$ holds. This corresponds to saying that $n_{i_1,p(i_1)}$ is an entry in A, that $n_{i_2,p(i_2)}$, $\ldots, n_{i_{s_1-1},p(i_{s_1-1})}$ are entries in $I+F$, and that $n_{i_{s_1},p(i_{s_1})}$ is an entry in B_t^T. Observe, that $p(i_r) = i_r$ cannot happen for $2 \leq r \leq s_1$ as already $p(i_{r-1}) = i_r$ holds. As $n_{i_r,p(i_r)}$ is non-zero by (8) we therefore must have

$$n_{i_r,p(i_r)} = f_{i_r,p(i_r)} = f_{i_r,i_{r+1}}$$

for $2 \leq r < s_1$. Hence,

$$(n_{i_1,p(i_1)}, \ldots, n_{i_{s_1},p(i_{s_1})}) = (a_{1,i_2}, f_{i_2,i_3}, \ldots, f_{i_{s_1-1},i_{s_1}}, b_{t,v_1,i_{s_1}})$$

for some v_1. Denote this sequence by P_1. Clearly, P_1 corresponds to the polynomial K_1 in the theorem.

We next apply the same procedure as above starting with $i_{s_1+1} = |E|+2$ to get a sequence P_2 of length s_2. Then we do the same with $i_{s_1+s_2+1} = |E|+3, \ldots, i_{s_1+\cdots s_{h-1}+1} = |E|+h$ to get the sequences $P_3, \ldots, P_h$. For $u = 2, \ldots, h$ we have

$$\begin{aligned} P_u &= \left(n_{i_{s_1+\cdots+s_{u-1}+1},p(i_{s_1+\cdots+s_{u-1}+1})}, \ldots, n_{i_{s_1+\cdots+s_u},p(i_{s_1+\cdots+s_u})}\right) \\ &= \Big(a_{u,i_{s_1+\cdots+s_{u-1}+2}}, f_{i_{s_1+\cdots+s_{u-1}+2},i_{s_1+\cdots+s_{u-1}+3}}, \ldots, \\ &\qquad f_{i_{s_1+\cdots+s_u-1},i_{s_1+\cdots+s_u}}, b_{t,v_u,i_{s_1+\cdots+s_u}}\Big) \, . \end{aligned}$$

Clearly, P_u corresponds to K_u in the theorem. Note that the sequences $P_1, \ldots, P_h$ by the very definition of a permutation are edge disjoint in the sense that

(1) $n_{i,j}$ occurs at most once in $P_1, \ldots, P_h$,
(2) if n_{j,l_1}, n_{j,l_2} occur in $P_1, \ldots, P_h$ then $l_1 = l_2$,
(3) if $n_{j_1,l}, n_{j_2,l}$ occur in $P_1, \ldots, P_h$ then $j_1 = j_2$.

Having indexed $s_1 + \cdots + s_h$ of the integers in $\{1, \ldots, |E| + h\}$ we consider what is left, namely

$$\Lambda = \{1, \ldots, |E| + h\} \setminus \{i_1, \ldots, i_{s_1+\ldots+s_h}\} .$$

By construction we have $i_1 = |E| + 1, \ldots, i_{s_1+\cdots+s_{h-1}+1} = |E| + h$ and therefore $\Lambda \subseteq \{1, \ldots, |E|\}$. Also by construction for every

$$\delta \in \{1, \ldots, |E|\} \cap \{i_1, \ldots, i_{s_1+\cdots+s_h}\}$$

we have $\delta = p(\epsilon)$ for some $\epsilon \in \{i_1, \ldots, i_{s_1+\cdots+s_h}\}$. Therefore $p(\lambda) \in \Lambda$ for all $\lambda \in \Lambda$ holds. In particular $p^{(x)}(\lambda) \in \{1, \ldots, |E|\}$ for all x. From Lemma 3 we conclude that $p(\lambda) = \lambda$ for all $\lambda \in \Lambda$. □