Variable-Rate Linear Network Coding

Silas L. Fong and Raymond W. Yeung

Abstract

We introduce variable-rate linear network coding for single-source finite acyclic network. In this problem, the source of a network transmits messages at different rates in different time sessions and every non-source node in the network decodes the messages if possible. We propose two efficient algorithms for implementing variable-rate linear network coding under different circumstances.

Index Terms

Network coding, variable-rate, linear broadcast, static network codes.

1 INTRODUCTION

Network coding, first studied by Yeung and Zhang [1] and Ahlswede et al. [2], reveals that if coding is applied at the nodes in a network, rather than routing alone, the network capacity can be increased. Li et al. [3] and subsequently Koetter and Medard [4] proved that linear network coding is sufficient to achieve the maximum capacity in a single-source finite acyclic network. Consequently, linear network coding for single-source finite acyclic networks has been a subject of much research interest. We refer the reader to [5] (see also [6]) for a tutorial on the subject. In this work, they classify linear network codes for single-source finite acyclic networks into four classes: (a) generic; (b) linear dispersion; (c) linear broadcast; (d) linear multicast. These four classes of linear network codes possess properties of decreasing strength.

Although there has been much investigation into various properties of linear network codes with a fixed rate, little research has been undertaken to investigate into the possible relationships among codes with different rates. In our previous work [7], the linkage among linear broadcasts of different rates was studied and the concept has been adopted by J. Goseling and J. H. Weber [8] for minimum-cost multicasting. This paper is an extension of [7], and variable-rate linear network coding with link failure is studied for the first time.

Silas L. Fong and Raymond W. Yeung are with the Department of Information Engineering, The Chinese University of Hong Kong, N.T., Hong Kong (e-mail: lhfong5@ie.cuhk.edu.hk; whyeung@ie.cuhk.edu.hk).

This paper is organized as follows. Section 2 presents various classes of linear network code, including linear broadcast. Section 3 presents the concept of variable-rate linear network coding and provides algorithms for efficient implementations of variable-rate linear network coding. In Section 4, the results in Section 3 are extended to the scenario with link failure. Section 5 concludes this paper.

2 PRELIMINARIES

2.1 Linear Network Code

A *network* is represented by a finite directed graph $G = (E, V)$ consisting of node set V and edge set E. *Nodes* are denoted by upper case letters (X, Y, etc). *Edges* are denoted by lower case letters (e, i, etc) on which a symbol from a finite field F, called the base field, can be transmitted. For simplicity, we assume every transmission on a channel and every internal processing of any node incur no delay. The *source* node is denoted by S which generates a message every unit time. The maximum flow from the source S to a non-source node T is denoted by $maxflow(T)$. The set of incoming edges and outgoing edges of node U are denoted by $In(U)$ and $Out(U)$ respectively. Let a pair of edges (d, e) be called an *adjacent pair* when there exists a node T with $d \in In(T)$ and $e \in Out(T)$.

In a linear network code, all the information symbols are regarded as elements of a base field F. These symbols include the symbols that comprise the information source as well as the symbols transmitted on the channels. For example, F is taken to be the field $GF(2)$ when the information unit is the bit. Furthermore, encoding and decoding are based on linear algebra defined on the base field, so that efficient algorithms for encoding and decoding as well as for code construction can be obtained. The global description of a linear network code described in [5] is used in this paper.

Definition 1: Let F be a finite field and ω be a positive integer. An *ω-dimensional F-valued linear network code* on an acyclic communication network consists of a scalar $k_{d,e}$ for every adjacent pair (d, e) in the network as well as an ω-dimensional column vector f_e for every edge e in the network such that:

(i) $f_e = \sum_{d \in In(T)} k_{d,e} f_d$, where $e \in Out(T)$;

(ii) The vectors f_e for the ω imaginary channel $e \in In(S)$ form the natural basis of the vector space F^ω.

The vector f_e is called *the global encoding kernel for edge e*. The local encoding kernel at the node T refers to the $|In(T)| \times |Out(T)|$ matrix $K_T = [k_{d,e}]_{d \in In(T), e \in Out(T)}$.

Let the source generate a message $\vec{x}$ in the form of an ω-dimensional row vector. A node T receives the symbols $\vec{x} \cdot f_d$, $d \in In(T)$, from which it calculates the symbol $\vec{x} \cdot f_e$ for sending onto each edge $e \in Out(T)$ via the linear formula

$$\vec{x} \cdot f_e = \vec{x} \cdot \sum_{d \in In(T)} k_{d,e} f_d = \sum_{d \in In(T)} k_{d,e} (\vec{x} \cdot f_d),$$

where the first equality follows from (i).

Given the local encoding kernels at all the nodes in an acyclic network, the global encoding kernels can be calculated recursively in any upstream-to-downstream order by (i), while (ii) provides the boundary conditions. An ω-dimensional F-valued linear network code can be viewed as an F-valued linear network code that enables the source to transmit a message consisting of ω data units.

Linear multicast and linear broadcast are described in [5] and their definitions are stated as follows:

Definition 2: Let vectors f_e denote the global encoding kernels in an ω-dimensional F-valued linear network code on a single-source finite acyclic network. Let

$$V_T = span\{f_d : d \in In(T)\}.$$

Then, the linear network code qualifies as a *linear multicast* and a *linear broadcast* respectively if the following statements hold:

(i) $dim(V_T) = \omega$ for every non-source node T with $maxflow(T) \geq \omega$;

(ii) $dim(V_T) = \min\{\omega, maxflow(T)\}$ for every non-source node T.

Clearly, $(ii) \Rightarrow (i)$. Thus, every linear broadcast is a linear multicast. Let p be the number of non-source node T with $maxflow(T) \geq \omega$ in an acyclic network. Using the algorithm proposed in [9], we can construct an ω-dimensional linear multicast on the network if the size of the base field is larger than p. A slight modification of this algorithm proves the following theorem.

Theorem 1: Given a single-source finite acyclic network with n non-source nodes and a finite field F, an ω-dimensional F-valued linear broadcast can be constructed if $|F| > n$.

Proof: It is similar to the proof in [9] and therefore omitted. □

Generally, a larger base field is required for constructing a linear broadcast than a linear multicast in the same network because the algorithms for constructing a linear broadcast need to consider more nodes compared with the algorithms for constructing a linear multicast.

2.2 Linear Network Code with Link Failure

In the discussion so far, a linear network code has been defined on a network with a fixed topology, where all the channels are assumed to be available at all times. In real life, a communication network often suffers from link failures or traffic congestions from time to time. In other words, the effective configuration of a communication network may vary from time to time. Link failures need to be handled efficiently because otherwise a large amount of data can be lost, especially when the data rate is high. Consider the use of, for instance, an ω-dimensional multicast on an acyclic network for multicasting a sequence of messages generated at the source node. When no channel failure occurs, a non-source node

T with $maxflow(T)$ at least equal to ω would be able to decode the sequence of messages. In case of link failures, if $maxflow(T)$ in the resulting network is at least ω, the sequence of messages in principle can still be received at that node. However, the deployment of a network code for the new network topology is involved, which not only is cumbersome but also may cause a significant loss of data during the switchover. In order to develop an efficient scheme for handling link failures, a class of linear network code called static network code described in [5] is studied in this paper, which can provide the network with maximum robustness in case of channel failures. The configuration formally defined in [5] and the global description of static network code in [5] are stated as follows:

Definition 3: A configuration ε of a network is a mapping from the set of channels in the network to the set $\{0, 1\}$. Channels in $\varepsilon^{-1}(0)$ are idle channels with respect to this configuration, and the subnetwork resulting from the deletion of idle channels will be called the ε-subnetwork. The maximum flow from the source S to a non-source node T over the ε-subnetwork is denoted as $maxflow_{\varepsilon}(T)$.

Definition 4: Let F be a finite field and ω be a positive integer. Let $k_{d,e}$ be the local encoding kernel for every adjacent pair (d, e) in an ω-dimensional F-valued linear network code on an acyclic communication network. The ε-global encoding kernel for the channel e, denoted by $f_{e,\varepsilon}$, is the ω-dimensional column vector calculated recursively in an upstream-to-downstream order by:

(i) $f_{e,\varepsilon} = \varepsilon(e) \sum_{d \in In(T)} k_{d,e} f_{d,\varepsilon}$, where $e \in Out(T)$.

(ii) The ε-global encoding kernel for the ω imaginary channels are independent of ε and form the natural basis of the vector space F^{ω}.

In the above definition, the local encoding kernels $k_{d,e}$ remain unchanged with ε. Let the source generate a message $\vec{x}$ in the form of an ω-dimensional row vector. A node T receives the symbols $\vec{x} \cdot f_{d,\varepsilon}$, $d \in In(T)$, from which it calculates the symbol $\vec{x} \cdot f_{e,\varepsilon}$ for sending onto each edge $e \in Out(T)$ via the linear formula

$$\vec{x} \cdot f_{e,\varepsilon} = \varepsilon(e) \sum_{d \in In(T)} k_{d,e} (\vec{x} \cdot f_{d,\varepsilon}).$$

In particular, a channel e with $\varepsilon(e) = 0$ has $f_{e,\varepsilon} = \vec{0}$ according to (i) and transmits the symbol $\vec{x} \cdot f_{e,\varepsilon} = 0$. In a real network, whenever a symbol is not received on an input channel due to channel failures, the symbol is regarded as being 0.

Static linear multicast and static linear broadcast are described in [5] and their definitions are stated as follows:

Definition 5: Following the notation of Definition 4 and letting

$$V_{T,\varepsilon} = span\{f_{d,\varepsilon} : d \in In(T)\},$$

an ω-dimensional F-valued linear network code on a single-source finite acyclic network qualifies as a static linear multicast and a static linear broadcast respectively if the following statements hold:

(i) $dim(V_{T,\varepsilon}) = \omega$ for every configuration ε and every non-source node T with $maxflow_{\varepsilon}(T) \geq \omega$;

(ii) $dim(V_{T,\varepsilon}) = \min\{\omega, maxflow_{\varepsilon}(T)\}$ for every configuration ε and every non-source node T.

While the configuration ε varies, the local encoding kernels remain unchanged. Therefore, the advantage of using a static linear broadcast in case of link failures is that the local operation at any node in the network is affected only at the minimal level. Each receiving node in the network, however, needs to know the configuration ε before decoding the source message correctly.

Let p be the number of non-source node T with $maxflow(T) \geq \omega$ and m be the number of configurations in an acyclic network. Using the algorithm proposed in [4], we can construct an ω-dimensional static linear multicast on the network if the size of the base field is larger than mp. A slight modification of this algorithm proves the following theorem.

Theorem 2: Given a single-source finite acyclic network with n non-source nodes, m configurations and a finite field F, an ω-dimensional F-valued static linear broadcast can be constructed if $|F| > mn$.

Proof: It is similar to the proof of constructing a static linear multicast in [4] and therefore omitted. □

3 VARIABLE-RATE LINEAR NETWORK CODING

In a single-source finite acyclic network, suppose the source wants to transmit messages at one of q possible rates within a session. Let $\bar{q}$ be the highest among the q rates. To avoid triviality, assume $\bar{q} \leq maxflow(T)$ for at least one non-source node T. We are now required to design a linear network coding system which enables every non-source node T to decode the message if $maxflow(T)$ is at least equal to the transmission rate in that session. In this section, we assume that link failures do not occur in the network; networks with link failure will be treated in Section 4.

The most effective solution based on existing results for the scenario described above is to use the algorithm proposed by Jaggi et al. [9] to obtain q linear multicasts of different dimensions for the same network. Consequently, every node is required to store q different copies of local encoding kernels in order to be able to apply the suitable local encoding kernel for that session. This increases the complexity of the system considerably if the system is implemented in hardware. Besides, changing the local encoding kernels at the nodes consumes resources in the network.

As an attempt to alleviate the shortcomings in the solution above, we propose in this section a new scheme based on linear broadcast for more efficient implementation of variable-rate linear network coding. Throughout this paper, all the networks concerned are single-source finite acyclic networks and we let F^{ω} denote the vector space of all ω-dimensional column vectors.

Lemma 1: An ω-dimensional F-valued linear network code is given on an acyclic network where $\omega \geq 2$. Let f_e be the global encoding kernel for all edge $e \in E$. Let $I_{\omega-1}$ denote the $(\omega-1) \times (\omega-1)$ identity matrix and let $\vec{b} \in F^{\omega-1}$ be any arbitrary $(\omega-1)$-dimensional column vector. Let

$$f_e^{\omega-1} = \left[\begin{array}{cc} I_{\omega-1} & \vec{b} \end{array} \right] f_e \tag{1}$$

for all non-imaginary channel e. Then, $f_e^{\omega-1}, e \in E$ constitute the global encoding kernels of an $(\omega-1)$-dimensional F-valued linear network code in the same base field F. In particular, the local encoding kernel of this $(\omega-1)$-dimensional linear network code at every non-source node is the same as that of the original ω-dimensional linear network code.

Proof: Let $k_{d,e}$ be the local encoding kernel for every adjacent pair (d,e) of the given ω-dimensional F-valued linear network code. We will show that $f_e^{\omega-1}, e \in E$ constitute the global encoding kernels of an $(\omega-1)$-dimensional F-valued linear network code by demonstrating the existence of the corresponding local encoding kernel $k_{d,e}^{\omega-1}$ for every adjacent pair (d,e).

By convention, we assume that the global encoding kernel for the $\omega-1$ imaginary channels form the standard basis of $F^{\omega-1}$. For any channel $e \in Out(S)$, since $f_e^{\omega-1}$ as specified in (1) is in $F^{\omega-1}$, $k_{d,e}^{\omega-1}, d \in In(S)$ can always be chosen.

For all non-imaginary channel $e \notin Out(S)$, let $k_{d,e}^{\omega-1} = k_{d,e}$. We now verify the relation

$$f_e^{\omega-1} = \sum_{d \in In(T)} k_{d,e}^{\omega-1} f_d^{\omega-1} \tag{2}$$

by considering

$$f_e = \sum_{d \in In(T)} k_{d,e} f_d.$$

Multiplying both sides by $\left[\begin{array}{cc} I_{\omega-1} & \vec{b} \end{array} \right]$, we obtain

$$\left[\begin{array}{cc} I_{\omega-1} & \vec{b} \end{array} \right] f_e = \sum_{d \in In(T)} k_{d,e} \left[\begin{array}{cc} I_{\omega-1} & \vec{b} \end{array} \right] f_d.$$

Then (2) immediately follows from (1), since $k_{d,e}^{\omega-1} = k_{d,e}$ for all non-imaginary channel $e \notin Out(S)$. This shows that $f_e^{\omega-1}, e \in E$ constitute the global encoding kernels of an $(\omega-1)$-dimensional F-valued linear network code with the local encoding kernels $k_{d,e}^{\omega-1}$. In particular, $k_{d,e}^{\omega-1} = k_{d,e}$ for every adjacent pair (d,e) for $e \notin Out(S)$. In other words, the local encoding kernel at every non-source node of the $(\omega-1)$-dimensional linear network code specified by $f_e^{\omega-1}, e \in E$ is the same as that of the original ω-dimensional linear network code. □

Definition 6: Let an ω-dimensional F-valued linear broadcast on an acyclic network where $\omega \geq 2$ and $\vec{b} \in F^{\omega-1}$, an $(\omega-1)$-dimensional column vector, be given. Define

$$f_e^{\omega-1} = \left[\begin{array}{cc} I_{\omega-1} & \vec{b} \end{array} \right] f_e$$

for all non-imaginary channel e, where f_e is the global encoding kernel for channel e. Then, $\vec{b}$ is called a reduction vector for the given linear broadcast if $f_e^{\omega-1}, e \in E$ specify an $(\omega-1)$-dimensional F-valued linear broadcast.

Lemma 2: Let F be a finite field, and ω and m be integers such that $\omega \geq 2$ and $1 \leq m \leq \omega - 1$. Let $\vec{c}_1, \vec{c}_2, \ldots, \vec{c}_m \in F^\omega$ be m linearly independent vectors, and let

$$\vec{d}_i = \left[\begin{array}{cc} I_{\omega-1} & \vec{b} \end{array} \right] \vec{c}_i \tag{3}$$

for $i = 1, 2, \ldots, m$, where

$$\vec{b} = \left[\begin{array}{cccc} b_1 & b_2 & \cdots & b_{\omega-1} \end{array} \right]^T$$

and $b_1, b_2, \ldots, b_{\omega-1}$ are indeterminates in F. Then, there exists a nonzero polynomial

$$p(b_1, b_2, \ldots, b_{\omega-1}) = a_0 + a_1 b_1 + a_2 b_2 + \ldots + a_{\omega-1} b_{\omega-1}$$

where a_j's are constants in F such that $\vec{d}_1, \vec{d}_2, \ldots, \vec{d}_m$ are linearly independent whenever

$$p(b_1, b_2, \ldots, b_{\omega-1}) \neq 0.$$

Proof: Construct the matrix

$$D_m = \left[\begin{array}{cccc} \vec{d}_1 & \vec{d}_2 & \cdots & \vec{d}_m \end{array} \right].$$

We will show that there exists an $m \times m$ submatrix A of D_m whose determinant is equal to a nonzero polynomial in $b_1, b_2, \ldots, b_{\omega-1}$. We will further show that $det(A)$ has the form

$$a_0 + a_1 b_1 + a_2 b_2 + \ldots + a_{\omega-1} b_{\omega-1}$$

where a_j's are constants in F. Then by letting

$$p(b_1, b_2, \ldots, b_{\omega-1}) = det(A),$$

since A is a submatrix of D_m, it follows that $\vec{d}_1, \vec{d}_2, \ldots, \vec{d}_m$ are linearly independent whenever $p(b_1, b_2, \ldots, b_{\omega-1})$ is evaluated to a nonzero value in F.

To facilitate our discussion, we write

$$\vec{c}_i = \left[\begin{array}{c} \vec{h}_i \\ k_i \end{array} \right], \tag{4}$$

where $\vec{h}_i \in F^{\omega-1}$ and $k_i \in F$ for $i = 1, 2, \ldots, m$. It is readily seen from (3) that

$$\vec{d}_i = \vec{h}_i + k_i \vec{b}$$

for $i = 1, 2, \ldots, m$, which implies

$$D_m = \left[\begin{array}{cccc} \vec{h}_1 + k_1 \vec{b} & \vec{h}_2 + k_2 \vec{b} & \cdots & \vec{h}_m + k_m \vec{b} \end{array} \right]. \tag{5}$$

We first show that there exists some $\vec{b} \in F^{\omega-1}$ such that $\vec{d}_1, \vec{d}_2, \ldots, \vec{d}_m$ are linearly independent. Assume the contrary, i.e., $\vec{d}_1, \vec{d}_2, \ldots, \vec{d}_m$ are linearly dependent for all $\vec{b}$. We will show that this leads to a contradiction.

$Case\ 1 : |span\{\vec{h}_1, \vec{h}_2, \ldots, \vec{h}_m\}| < \omega - 1$

Since $|span\{\vec{h}_1, \vec{h}_2, \ldots, \vec{h}_m\}|$ is at most $\omega - 2$, a vector $\vec{z} \in F^{\omega-1}$ can always be found such that $\vec{z} \notin span\{\vec{h}_1, \vec{h}_2, \ldots, \vec{h}_m\}$. Then, by our assumption, $\{\vec{d}_i\}$ are linearly dependent for all $\vec{b}$, in particular for $\vec{b}$ equals $\vec{z}$. In other words, $\{\vec{d}_i + k_i\vec{z}\}$ are linearly dependent, i.e.,

$$\begin{aligned} & t_1(\vec{h}_1 + k_1\vec{z}) + t_2(\vec{h}_2 + k_2\vec{z}) \\ + \quad & \ldots + t_m(\vec{h}_m + k_m\vec{z}) = \vec{0} \end{aligned}$$

for some $t_1, t_2, \ldots t_m \in F$ where not all t_i's are equal to 0. Regrouping the terms, we have

$$\begin{aligned} & (t_1\vec{h}_1 + t_2\vec{h}_2 + \ldots + t_m\vec{h}_m) \\ + \quad & (t_1k_1 + t_2k_2 + \ldots + t_mk_m)\vec{z} = \vec{0}. \end{aligned}$$

Since $\vec{z} \notin span\{\vec{h}_1, \vec{h}_2, \ldots, \vec{h}_m\}$, this implies

$$\begin{cases} t_1k_1 + t_2k_2 + \ldots + t_mk_m = 0 \\ t_1\vec{h}_1 + t_2\vec{h}_2 + \ldots + t_m\vec{h}_m = \vec{0}. \end{cases}$$

Consequently,

$$t_1\vec{c}_1 + t_2\vec{c}_2 + \ldots t_m\vec{c}_m = \vec{0}$$

(cf.(4)), which contradicts the linear independence among $\vec{c}_1, \vec{c}_2, \ldots, \vec{c}_m$.

$Case\ 2 : |span\{\vec{h}_1, \vec{h}_2, \ldots, \vec{h}_m\}| = \omega - 1$

Since m is at most $\omega - 1$ and $|span\{\vec{h}_1, \vec{h}_2, \ldots, \vec{h}_m\}|$ equals $\omega - 1$, m equals $\omega - 1$ and $\vec{h}_1, \vec{h}_2, \ldots, \vec{h}_m$ are linearly independent. However, for $\vec{b}$ equals $\vec{0}$,

$$\vec{d}_i = \vec{h}_i$$

for $i = 1, 2, \ldots, m$. Then, $\{\vec{d}_i\}$ are linearly independent for $\vec{b}$ equals $\vec{0}$, which contradicts our assumption.

Combining the two cases, we have shown that $\vec{d}_1, \vec{d}_2, \ldots, \vec{d}_m$ are linearly independent for some $\vec{b}$. For this choice of $\vec{b}$, there exists a submatrix A of D_m such that $det(A)$ is evaluated to a nonzero value. Since $det(A)$ is a polynomial in the indeterminates $b_1, b_2, \ldots, b_{\omega-1}$, this implies that $det(A)$ is a nonzero polynomial in these indeterminates. Since D_m is $(\omega - 1) \times m$ and A is an $m \times m$ submatrix of D_m, we see from (5) that

$$A = \left[\begin{array}{ccc} \vec{r}_1 + k_1\vec{b}' & \vec{r}_2 + k_2\vec{b}' & \cdots \quad \vec{r}_m + k_m\vec{b}' \end{array} \right],$$

where $\vec{r}_1, \vec{r}_2, \ldots, \vec{r}_m, \vec{b}' \in F^m$ are the corresponding subvectors of $\vec{h}_1, \vec{h}_2, \ldots, \vec{h}_m$ and $\vec{b}$ respectively. If $k_1 = \ldots = k_m = 0$, then $det(A) = a_0$ where $a_0 \in F$. Otherwise, assume without loss of generality that $k_1 \neq 0$. Then, by means of column operations on A, we see that $det(A)$ can be expressed in the form

$$C \left| \left[\begin{array}{cccc} \vec{l}_1 + \tau \vec{b}' & \vec{l}_2 & \cdots & \vec{l}_m \end{array} \right] \right|$$

where $C, \tau \in F$ and $\vec{l}_i \in F^m$. It then follows that in $det(A)$, the power of each component of $\vec{b}$ is at most one. Therefore,

$$det(A) = a_0 + a_1 b_1 + a_2 b_2 + \ldots + a_{\omega-1} b_{\omega-1}$$

where a_j's are constants in F for $j = 0, 1, \ldots, \omega - 1$. Let

$$p(b_1, b_2, \ldots, b_{\omega-1}) = det(A)$$

and this completes the proof of the lemma. □

Lemma 3: Let n be the total number of non-source nodes in an acyclic network, and an ω-dimensional F-valued linear broadcast be given, where $\omega \geq 2$. Then a reduction vector can be found if $|F| > n$.

Proof: Let f_e be the global encoding kernel of the given linear broadcast for all edge $e \in E$. Let

$$\vec{b} = \left[\begin{array}{cccc} b_1 & b_2 & \cdots & b_{\omega-1} \end{array} \right]^T$$

be an $(\omega - 1)$-dimensional column vector where all b_ξ's are indeterminates in F, and let

$$f_e^{\omega-1} = \left[\begin{array}{cc} I_{\omega-1} & \vec{b} \end{array} \right] f_e$$

for all non-imaginary channel e. The existence of a reduction vector is proved by showing that by suitably choosing $\vec{b}$, $f_e^{\omega-1}, e \in E$ specify an $(\omega - 1)$-dimensional F-valued linear broadcast.

For each non-source node T, let

$$m = \min\{\omega - 1, maxflow(T)\}.$$

Then, m linearly independent vectors f_e can always be chosen from the set of incoming edge $e \in In(T)$ since the given linear network code is a linear broadcast. Denote these m vectors by $\vec{c}_1, \vec{c}_2, \ldots, \vec{c}_m$ and let

$$\vec{c}_i{}^{\omega-1} = \left[\begin{array}{cc} I_{\omega-1} & \vec{b} \end{array} \right] \vec{c}_i$$

for $i = 1, 2, \ldots, m$, where $\vec{c}_i{}^{\omega-1}$ is an $(\omega - 1)$-dimensional column vectors. Note that m as well as the vectors $\vec{c}_1, \vec{c}_2, \ldots, \vec{c}_m$ and $\vec{c}_1{}^{\omega-1}, \vec{c}_2{}^{\omega-1}, \ldots, \vec{c}_m{}^{\omega-1}$ depend on node T although this is not explicitly indicated in order to keep the notation simple. Let g_T be the nonzero polynomial $p(b_1, b_2, \ldots, b_{\omega-1})$ in Lemma 2, which exists because $\vec{c}_1, \vec{c}_2, \ldots, \vec{c}_m$ are linearly independent. Let N_T denote the solution space of

$$g_T(b_1, b_2, \ldots, b_{\omega-1}) = 0.$$

Since g_T is a nonzero polynomial in $\omega-1$ variables, $|N_T| \leq |F|^{\omega-2}$. We now consider

$$\left| F^{\omega-1} \cap (\bigcup_T N_T) \right|$$

in order to find a reduction vector

$$\vec{v} = \left[\begin{array}{cccc} v_1 & v_2 & \cdots & v_{\omega-1} \end{array} \right]^T.$$

By the union bound,

$$\left| F^{\omega-1} \cap (\bigcup_T N_T) \right| \leq \sum_T \left| (F^{\omega-1} \cap N_T) \right|.$$

Since $|N_T| \leq |F|^{\omega-2}$ and $n < |F|$, this implies

$$\begin{aligned} \sum_T \left| (F^{\omega-1} \cap N_T) \right| &\leq \sum_T |F|^{\omega-2} \\ &= n\,|F|^{\omega-2} \\ &< |F|^{\omega-1}. \end{aligned}$$

Therefore,

$$\left| F^{\omega-1} \cap (\bigcup_T N_T) \right| < \left| F^{\omega-1} \right|$$

and we can find $\vec{v} \in F^{\omega-1}$ such that $\vec{v} \notin \bigcup_T N_T$. In other words, $\vec{v}$ can be obtained such that

$$g_T(v_1, v_2, \ldots, v_{\omega-1}) \neq 0$$

for each non-source node T, which implies $\vec{c}_1^{\,\omega-1}, \vec{c}_2^{\,\omega-1}, \ldots, \vec{c}_m^{\,\omega-1}$ are linearly independent for each non-source node T when $\vec{b} = \vec{v}$ by Lemma 2. Consequently, $f_e^{\omega-1}, e \in E$ specify an $(\omega-1)$-dimensional F-valued linear network code that

$$\begin{aligned} dim(V_T) &= m \\ &= \min\{\omega-1, maxflow(T)\} \end{aligned}$$

for each non-source node T when $\vec{b} = \vec{v}$. Therefore, $f_e^{\omega-1}, e \in E$ specify an $(\omega-1)$-dimensional F-valued linear broadcast for $\vec{b} = \vec{v}$. It then follows from Definition 6 that $\vec{v}$ is a reduction vector for the given linear broadcast. □

Lemma 3 provides an algorithm to find a reduction vector and an application of Lemma 3 is illustrated by the following simple example.

Example 1: An acyclic network with 7 non-source nodes and a 3-dimensional $GF(11)$ linear broadcast on the network are shown in Fig. 1. The local encoding kernels at the non-source nodes of the linear broadcast are shown in Fig. 2. Since $|GF(11)| > 7$, a reduction vector can be found by Lemma 3 and $\left[\begin{array}{cc} 1 & 2 \end{array} \right]^T$ is found to be a reduction vector. The corresponding 2-dimensional $GF(11)$ linear broadcast

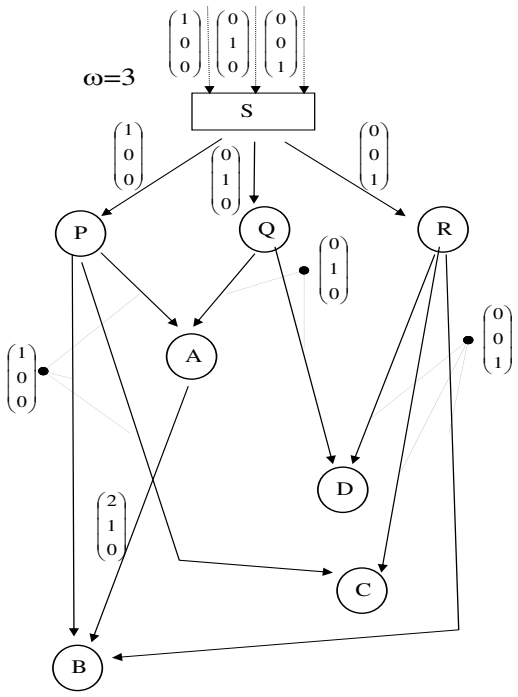

Fig. 1. A 3-dimensional $GF(11)$ linear broadcast

K_P	K_Q	K_R	K_A
$[1\ \ 1\ \ 1]$	$[1\ \ 1]$	$[1\ \ 1\ \ 1]$	$\begin{bmatrix}2\\1\end{bmatrix}$

Fig. 2. The local encoding kernels at the non-source nodes

constructed by the reduction vector is shown in Fig. 3. It can be easily observed that the two linear broadcasts have the same local encoding kernels at all the non-source nodes.

Theorem 3: Let n be the total number of non-source nodes in an acyclic network. An ω-dimensional F-valued linear broadcast is given on the network where $\omega \geq 2$ and $|F| > n$. Then, for every $h = 1, 2, \ldots, \omega - 1$, an h-dimensional F-valued linear broadcast can be constructed such that these linear broadcasts have the same local encoding kernels at all the non-source nodes.

Proof: Using Lemma 3, a reduction vector for the given linear broadcast can be found and an $(\omega - 1)$-dimensional linear broadcast is obtained. By Lemma 1, the local encoding kernel of this $(\omega - 1)$-dimensional linear broadcast at every non-source node is the same as that of the original ω-dimensional linear broadcast. By repeating this procedure, each time reducing the dimension of the linear broadcast

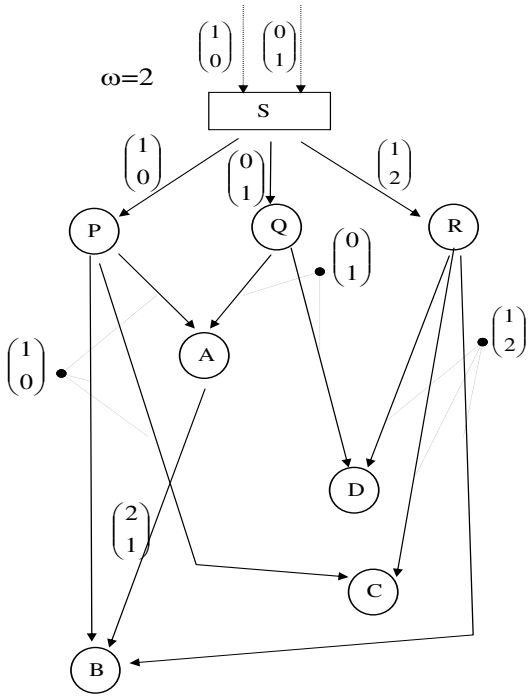

Fig. 3. A 2-dimensional $GF(11)$ linear broadcast

by one, the desired set of linear broadcasts can be obtained. □

Theorem 1 and Theorem 3 together render an efficient implementation of linear broadcasts of different dimensions on the same network. In particular, they provide a solution to the scenario described at the beginning of this section. This solution, which has the advantage that each non-source node is required to store only one copy of the local encoding kernel, is summarized by the following two steps:

Step 1 : Let n be the number of non-source node in the network and a $\bar{q}$-dimensional F-valued linear broadcast where $|F| > n$ is constructed by Theorem 1.

Step 2 : Lower-dimension linear broadcasts are obtained from the $\bar{q}$-dimensional broadcast by Theorem 3.

4 VARIABLE-RATE LINEAR NETWORK CODING WITH LINK FAILURE

In a single-source finite acyclic network with $2^{|E|}$ possible configurations, suppose the source wants to transmit messages at one of q possible rates within a session. Let $\bar{q}$ be the highest among the q rates. Let ε_Ω denote the configuration (cf. Definition 3) with no link failure, i.e., $\varepsilon_\Omega(e) = 1$ for all non-imaginary channel $e \in E$. To avoid triviality, assume $\bar{q} \leq maxflow_{\varepsilon_\Omega}(T)$ for at least one non-source node T. We are now required to design a linear network coding system which enables every non-source node T to decode the message if $maxflow_{\varepsilon}(T)$ is at least equal to the transmission rate in that session with ε

being the configuration. If we want to minimize the complexity of the local operation at all the nodes, an effective solution based on existing results is to use the algorithm proposed in [4] to construct q static linear multicasts of different dimensions for the same network. Consequently, every node is required to store q different copies of local encoding kernels in order to be able to apply the suitable local encoding kernel for that session.

As an attempt to alleviate the shortcomings in the solution above, we extend the scheme developed in the previous section in order to implement variable-rate linear network coding with link failure more efficiently.

Lemma 4: An ω-dimensional F-valued linear network code is given on an acyclic network. Let $f_{e,\varepsilon}$ be the ε-global encoding kernel for all edge $e \in E$ and every configuration ε. Let $I_{\omega-1}$ denote the $(\omega-1)\times(\omega-1)$ identity matrix and let $\vec{b} \in F^{\omega-1}$ be any arbitrary $(\omega-1)$-dimensional column vector. Let

$$f_{e,\varepsilon}^{\omega-1} = \left[\begin{array}{cc} I_{\omega-1} & \vec{b} \end{array} \right] f_{e,\varepsilon} \tag{6}$$

for all non-imaginary channel e and every configuration ε. Then, $f_{e,\varepsilon}^{\omega-1}, e \in E$ constitute the ε-global encoding kernels of an $(\omega-1)$-dimensional F-valued linear network code in the same base field F. In particular, the local encoding kernel of this $(\omega-1)$-dimensional linear network code at every non-source node is the same as that of the original ω-dimensional linear network code.

Proof: It is similar to the proof in Lemma 1 and therefore omitted. □

Definition 7: Let an ω-dimensional F-valued static linear broadcast on an acyclic network where $\omega \geq 2$, and $\vec{b} \in F^{\omega-1}$, an $(\omega-1)$-dimensional column vector, be given. Define

$$f_{e,\varepsilon}^{\omega-1} = \left[\begin{array}{cc} I_{\omega-1} & \vec{b} \end{array} \right] f_{e,\varepsilon}$$

for all non-imaginary channel e and every configuration ε, where $f_{e,\varepsilon}$ is the global encoding kernel for channel e under configuration ε. Then, $\vec{b}$ is called a static reduction vector for the given static linear broadcast if $f_{e,\varepsilon}^{\omega-1}, e \in E$ specify an $(\omega-1)$-dimensional F-valued linear broadcast for every configuration ε.

Lemma 5: Let n be the total number of non-source nodes in an acyclic network and m be the total number of configurations ε in the network. For any ω-dimensional F-valued static linear broadcast where $\omega \geq 2$, a static reduction vector can be found if $|F| > mn$.

Proof: Let $f_{e,\varepsilon}$ be the global encoding kernel of the given static linear broadcast for all edge $e \in E$ and every possible configuration ε. Let

$$\vec{b} = \left[\begin{array}{cccc} b_1 & b_2 & \cdots & b_{\omega-1} \end{array} \right]^T$$

be an $(\omega-1)$-dimensional column vector where all b_ξ's are indeterminates in F, and let

$$f_{e,\varepsilon}^{\omega-1} = \left[\begin{array}{cc} I_{\omega-1} & \vec{b} \end{array}\right] f_{e,\varepsilon}$$

for all non-imaginary channel e and every configuration ε. The existence of a static reduction vector is proved by showing that by suitably choosing $\vec{b}$, $f_{e,\varepsilon}^{\omega-1}, e \in E$ specify an $(\omega-1)$-dimensional F-valued linear broadcast for every configuration ε.

For each configuration ε, the network code on the ε-subnetwork is a linear broadcast. Therefore, we let a nonzero polynomial $g_{T,\varepsilon}(b_1, b_2, \ldots, b_{\omega-1})$ be g_T in the proof of Lemma 3 for each non-source node T under each ε. Let $N_{T,\varepsilon}$ denote the solution space of

$$g_{T,\varepsilon}(b_1, b_2, \ldots, b_{\omega-1}) = 0.$$

Since $g_{T,\varepsilon}$ is a nonzero polynomial in $\omega-1$ variables, $|N_{T,\varepsilon}| \le |F|^{\omega-2}$. We now consider

$$\left| F^{\omega-1} \cap \left(\bigcup_{\varepsilon} \left(\bigcup_{T} N_{T,\varepsilon} \right) \right) \right|$$

in order to find a static reduction vector

$$\vec{v} = \left[\begin{array}{cccc} v_1 & v_2 & \cdots & v_{\omega-1} \end{array}\right]^T.$$

By the union bound,

$$\left| F^{\omega-1} \cap \left(\bigcup_{\varepsilon} \left(\bigcup_{T} N_{T,\varepsilon} \right) \right) \right| \le \sum_{\varepsilon} \left(\sum_{T} \left| \left(F^{\omega-1} \cap N_{T,\varepsilon} \right) \right| \right).$$

Since $|N_T| \le |F|^{\omega-2}$ and $mn < |F|$, this implies

$$\begin{aligned} \sum_{\varepsilon} \left(\sum_{T} \left| \left(F^{\omega-1} \cap N_{T,\varepsilon} \right) \right| \right) &\le \sum_{\varepsilon} \left(\sum_{T} |F|^{\omega-2} \right) \\ &= mn\,|F|^{\omega-2} \\ &< |F|^{\omega-1}. \end{aligned}$$

Therefore,

$$\left| F^{\omega-1} \cap \left(\bigcup_{\varepsilon} \left(\bigcup_{T} N_{T,\varepsilon} \right) \right) \right| < \left| F^{\omega-1} \right|$$

and we can find $\vec{v} \in F^{\omega-1}$ such that $\vec{v} \notin \bigcup_\varepsilon (\bigcup_T N_{T,\varepsilon})$. In other words, $\vec{v}$ can be obtained such that

$$g_{T,\varepsilon}(v_1, v_2, \ldots, v_{\omega-1}) \neq 0$$

for each non-source node T under every configuration ε. By the similar arguments as the proof in Lemma 3, $f_{e,\varepsilon}^{\omega-1}, e \in E$ specify an $(\omega-1)$-dimensional F-valued linear broadcast for $\vec{b} = \vec{v}$ under every configuration ε. It then follows from Definition 7 that $\vec{v}$ is a static reduction vector. □

Theorem 4: Let n be the total number of non-source nodes in an acyclic network and m be the total number of configurations. An ω-dimensional F-valued static linear broadcast is given on the network

where $\omega \geq 2$ and $|F| > mn$. Then, for every $h = 1, 2, \ldots, \omega - 1$, an h-dimensional F-valued static linear broadcast can be constructed such that these static linear broadcasts have the same local encoding kernels at all the non-source nodes.

Proof: Using Lemma 5, a static reduction vector for the given static linear broadcast can be found and an $(\omega - 1)$-dimensional static linear broadcast is obtained. By Lemma 4, the local encoding kernel of this $(\omega - 1)$-dimensional static linear broadcast at every non-source node is the same as that of the original ω-dimensional static linear broadcast. By repeating this procedure, each time reducing the dimension of the static linear broadcast by one, the desired set of static linear broadcasts can be obtained. □

Theorem 2 and Theorem 4 together render an efficient implementation of static linear broadcasts of different dimensions on the same network. In particular, they provide a solution to the scenario described at the beginning of this section. This solution, which has the advantage that each non-source node is required to store only one copy of the local encoding kernel, is summarized by the following two steps:

Step 1 : Let n be the number of non-source nodes in the network and a $\bar{q}$-dimensional F-valued static linear broadcast where $|F| > 2^{|E|}n$ is constructed by Theorem 2.

Step 2 : Lower-dimension static linear broadcasts are obtained from the $\bar{q}$-dimensional static linear broadcast by Theorem 4.

We assume that all the $2^{|E|}$ possible configurations may occur in the network at the beginning of this section. Conceivably, a practical application may deal with only a certain collection $\{\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_k\}$ of configurations in order to provide link contingency, network security, network expandability, transmission redundancy, alternate routing upon congestion, etc. If the possible configurations in the network are reduced to $\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_k$, Theorem 2 and Theorem 4 can still be applied to construct static linear broadcasts that have the same local encoding kernels at all the non-source nodes. However, the number of configurations considered in Theorem 2 and Theorem 4 is reduced from $2^{|E|}$ to k. Consequently, the threshold on the sufficient size of the base field as well as the computational complexity in Theorem 2 and Theorem 4 will be lower.

5 CONCLUSION

A scheme that enables efficient implementation of variable-rate linear network coding in a single-source finite acyclic network is developed. In our scheme, the same local encoding kernel at every non-source node can be used for different transmission rates. In addition, two efficient algorithms are proposed for implementing variable-rate linear network coding in different situations. Compared with solutions based on existing results, our algorithms are simpler and require less storage space.

Further research includes the complexity analysis of our algorithms that enable efficient implementation of variable-rate linear network coding. The performance analysis of randomly designed codes for variable-rate linear network coding is also interesting for future research.

REFERENCES

[1] R. W. Yeung and Z. Zhang, "Distributed source coding for satellite communications," *IEEE Trans. Inform. Theory*, vol. IT-45, pp. 1111–1120, 1999.

[2] R. Ahlswede, N. Cai, S.-Y. R. Li and R. W. Yeung, "Network information flow," *IEEE Transactions on Information Theory*, vol. vol. 46, pp. pp. 1204–1216, 2000.

[3] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Transactions on Information Theory*, Feb. 2003.

[4] R. Koetter, M. Medard, "An algebraic approach to network coding," *Transactions on Networking*, Oct. 2003.

[5] R. W. Yeung, S.-Y. R. Li, N. Cai, and Z. Zhang, "Network coding theory," *Foundations and Trends in Communications and Information Theory*, 2006.

[6] R. W. Yeung, *Information theory and network coding*. Springer, 2008.

[7] S. L. Fong and R. W. Yeung, "Variable-rate linear network coding," in *Proc. IEEE Information Theory Workshop, Cheungdu, China*, Oct. 2006, pp. 409–412.

[8] J. Goseling and J. H. Weber, "Multi-rate network coding for minimum-cost multicasting," in *Proc. IEEE ISIT'08*, 2008.

[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," *IEEE Transactions on Information Theory*, vol. IT 51, no. 1973–1982, Mar. 2005.