On Coding for Non-Multicast Networks*

Muriel Médard Michelle Effros David Karger Tracey Ho

Abstract

We consider the issue of coding for non-multicast networks. For multicast networks, it is known that linear operations over a field no larger than the number of receivers are sufficient to achieve all feasible connections. In the case of non-multicast networks, necessary and sufficient conditions are known, if we restrict ourselves to linear codes over a finite field [1]. However, no linearity sufficiency results exist for non-multicast networks. Indeed, [2] shows that linearity over a field is not sufficient in general. We present a coding theorem that provides necessary and sufficient conditions, in terms of receiver entropies, for an arbitrary set of connections to be achievable on any network. We conjecture that linearity is sufficient to satisfy the coding theorem, when linear operations are performed over vectors rather than scalars in a field. We illustrate the intuition of this conjecture with an example. This work is part of an ongoing cooperation with R. Koetter.

1 Introduction

The beautiful results of [3] and [2] suggest a fundamental difference between the network coding problems for multicast and non-multicast networks. In [3], Li, Yeung, and Cai prove that linear coding is sufficient for multicast network coding problems. In contrast, Lehman and Lehman [2] demonstrate that the same linear coding model is not sufficient to achieve the optimal network coding performance in networks with arbitrary demands.

We begin with a simple new example that teases out the behavior that breaks the linear coding definition in [2]. The given example suggests that the problem discovered by Lehman and Lehman is not that linearity fails in non-multicast networks, but rather that prior definitions of linearity are too restrictive to admit known solutions to some simple non-multicast networks. Roughly, the prior models allow linear operations on a symbol-by-symbol basis but prohibit linear operations on vectors of symbols. While these schemes group symbols into vectors prior to coding, each resulting vector is considered as a single symbol from a fixed finite field. Necessary and sufficient conditions for achievability of connections for such linear schemes over a finite field have been established in [1]. However, allowing only linear operations on symbols over some finite field does not achieve all outcomes afforded by linear operations on vectors created from the original source symbols.

*M. Médard (medard@mit.edu) and T. Ho (trace@mit.edu) are with the Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology, Cambridge, MA 02139. M. Effros (effros@caltech.edu) is with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125. D. Karger (karger@csail.mit.edu) is with the Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology, Cambridge, MA 02139. This research is supported in part by NSF Grants CCR-0325324 and CCR-0220039, University of Illinois subaward #03-25673, Hewlett-Packard 008542-008, and Caltech's Lee Center for Advanced Networking.

The example we give in this paper leads to our conjecture that linear coding (according to its most general definition) is sufficient for network coding on non-multicast networks.

We pursue the conjecture of linearity by seeking an understanding of the space of achievable network coding problems. Towards this end, we first consider a collection of network simplifications, demonstrating that the most general network coding problem can be solved by solving a considerably smaller class of possible network problems. Using the given simplifications, we finally prove a coding theorem that gives necessary and sufficient conditions for the achievability of a collection of fixed demands in a general network. This coding theorem is only in terms of entropies at the receivers.

2 Example

We begin by providing a non-multicast example on a $K_{3,4}$ graph shown in Figure 1. This example is due to R. Koetter. The sources and demands are as shown in the figure.

Notice that a pure routing solution fails. Using routing, only one of the two sources sent to node 4 can be forwarded to any receiver. Therefore, the receiver that requires both of the sources that are sent to node 4 cannot be satisfied.

The following argument demonstrates that no code that uses only linear operations on symbols from a fixed finite field can solve the given network. We begin by examining node 4. If the output on any of the edges leaving node 4 is a linear combination of the node's inputs, then either

- edges (1,3) and (1,4) must carry the same linear combination or
- edges (2,4) and (2,5) must carry the same linear combination.

If neither of these conditions is satisfied, then the receivers cannot separate A and B at the decoders. Suppose that (1,3) and (1,4) carry the same mixture, then either

- sources A and A' cannot be unmixed or
- only one of A and A' gets through to the receivers.

In either case, the code fails to meet the given demands.

Now consider mixing at nodes 1 and 2. If node 1 mixes A and A' along either of its edges, say edge (1,3), then any node receiving information that passed through (1,3) needs the information that passed through (1,4) as well. Since 4 does no mixing, then the receiver in question can get only one stream originating at node 2. This leaves that receiver the ability to get only one of B or B'. Since all demands are represented in the network, one receiver cannot reconstruct its required sources.

While the above argument shows that linear coding on scalars is insufficient to meet the given demands, linear coding on vectors achieves the desired goal. Figure 2 gives a solution. This solution requires operations over two time steps. Note that the capacity of each link is now two bits per two time steps. In effect, we are renormalizing over time. In this 2-dimensional vector framework, we can meet all of the demands at all of the receivers using a pure routing solution. This solution can also be posed as a linear coding solution on a vector space of dimension 2.

Note that the reduction of 3-SAT in [2] also provides an example in which a vector solution is necessary. Thus, while the example in [2] violates linearity if we do not allow vector linear solutions, it bears a vector linear solution.

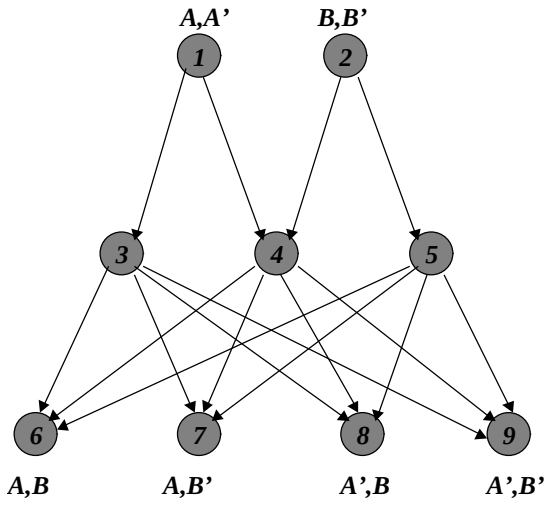

Figure 1: Linear operations on symbols are insufficient to achieve the capacity of this network.

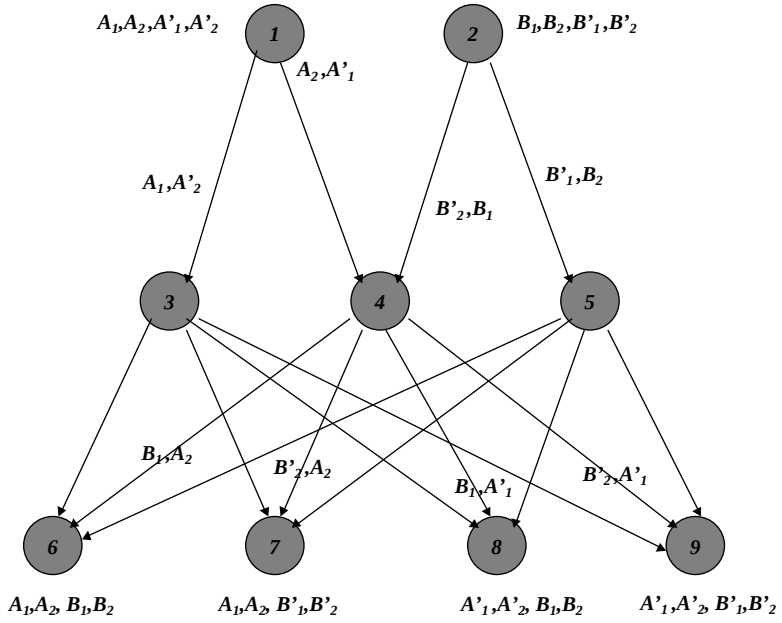

Figure 2: Linear operations on vectors are sufficient to achieve the capacity of this network.

3 Conjecture

We conjecture that linear coding – under its most general definition – is sufficient for network coding in systems with arbitrary demands. We next derive a collection of necessary and sufficient conditions, in terms of receiver entropies, for meeting a collection of demands on an arbitrary, fixed network. This argument suggests a partial proof of the sufficiency of linearity by showing that a collection of demands can be met with a code that uses only linear operations at the internal nodes (along with possibly non-linear operations at the system encoders and decoders) if and only if the entropies achievable by a non-linear operation at any step in the network are identical to the entropies achievable by a linear operation (perhaps of a different vector dimension) at the same location.

4 Coding Theorem

We define a network coding problem by describing a graph $\mathcal{G}$, a source matrix $\mathcal{S}$, and a demand matrix $\mathcal{D}$. The graph, which is by assumption directed and acyclic, is characterized by its vertex set $\mathcal{V}$ (representing the network nodes), its edge set $\mathcal{E}$ (representing the network links), and its link capacities $\mathcal{R} = [r(e)]_{e\in\mathcal{E}}$;[1] we therefore write $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{R})$. The source matrix is a $K \times |\mathcal{V}|$ binary matrices describing which, if any, of K possible sources nodes enter the network at nodes $v \in \mathcal{V}$. Similarly, the demand matrix is a $K\times|\mathcal{V}|$ binary matrix describing which, if any, of K possible sources are required at nodes $v \in \mathcal{V}$.More precisely, let $\mathbf{B} = (B_1, \ldots, B_K)$ be a collection of K input processes. Since this paper considers block coding, we use $\mathbf{B}_1, \mathbf{B}_2, \ldots$ to describe samples of source vector $\mathbf{B}$, with $\mathbf{B}_t = (B_{1,t}, B_{2,t}, \ldots, B_{K,t})$ representing the vector of source samples at integer time t and $B_k^n = (B_{k,1}, \ldots, B_{k,n})$ representing the first n samples of the kth random process $B_{k,1}, B_{k,2}, \ldots$ for any $k \in \mathcal{K} = \{1, \ldots, K\}$. The source matrix $\mathcal{S} = [s(k, v)]_{(k,v)\in\mathcal{K}\times\mathcal{V}}$ has entries

$$s(k, v) = \begin{cases} 1 & \text{if source } B_k \text{ is available at node } v \\ 0 & \text{otherwise.} \end{cases}$$

The demand matrix $\mathcal{D} = [d(k, v)]_{(k,v)\in\mathcal{K}\times\mathcal{V}}$ has entries

$$d(k, v) = \begin{cases} 1 & \text{if node } v \text{ requires source } B_k \\ 0 & \text{otherwise.} \end{cases}$$

We assume $s(k, v) = 1$ implies $d(k, v) = 0$ and $\sum_{v\in\mathcal{V}} d(k, v) > 0$ for all $k \in \mathcal{K}$. Thus a single node of the network can be both a transmitter and a receiver, but not for the same source; further each source entering the network has a destination distinct from all network entry points for that source.

Given an arbitrary graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{R})$, define a network code (f^n, g^n) for graph $\mathcal{G}$ to be a set of edge encoders and node decoders. Each encoder f_e^n in the set $f^n = \{f_e^n\}_{e\in\mathcal{E}}$ is a deterministic mapping describing the output on edge $e \in \mathcal{E}$ as a function of the available inputs. The decoder set $g^n = \{g_{k,v}^n\}_{(k,v)\in\mathcal{K}\times\mathcal{V}}$ defines K node decoders $(g_{1,v}^n, \ldots, g_{k,v}^n)$ for each node $v \in \mathcal{V}$; each decoder $g_{k,v}^n$ is a deterministic mapping giving the reconstruction of source B_k^n at node v. More notation is necessary to make these definitions precise.

For any edge $e \in \mathcal{E}$, let $t(e)$ and $h(e)$ denote the node from which edge e originates and to which it goes, respectively. For any node $v \in \mathcal{V}$, let $\mathcal{I}(v) \subseteq \mathcal{E}$ and $\mathcal{O}(v) \subseteq \mathcal{E}$

[1] We here assume that each link is lossless when used at any rate below its capacity and disallow the use of a link at any rate above its capacity.

denote the set of edges coming into and going out from node v, respectively. Finally, assuming a block coding strategy with blocklength n, let $w_e^n \in \mathcal{W}_e^n$ and $\mathbf{x}_v^n \in \mathcal{X}_v^n$ be the message transmitted across edge $e \in \mathcal{E}$ and all information available to node $v \in \mathcal{V}$, respectively, in a single use of a fixed block coding strategy. (Capital W_e^n and $\mathbf{X}_v^n$ denote the values corresponding to a random network input.) Message w_e^n is the binary message that traverses edge e in n uses of the network; thus

$$\mathcal{W}_e^n = \{0,1\}^{nr(e)}.$$

Source $\mathbf{x}_v^n$ is the concatenation of all sources available to node v and all messages incoming from the edges in $\mathcal{I}(v)$; thus

$$\mathcal{X}_v^n = \{0,1\}^{n(\sum_{k=1}^K s(k,v) + \sum_{e \in \mathcal{I}(v)} r(e))}.$$

For each edge $e \in \mathcal{E}$, any deterministic mapping

$$f_e^n : \mathcal{X}_{t(e)}^n \to \mathcal{W}_e^n$$

represents a legitimate encoder for mapping the space of available input information into the space of messages that can be transmitted in n uses of channel e. Similarly, for each $(k,v) \in \mathcal{K} \times \mathcal{V}$, any deterministic mapping

$$g_{k,v}^n : \mathcal{X}_v^n \to \{0,1\}^n$$

gives a legitimate decoder for reconstructing input source vector B_k^n from output source vector $\mathbf{X}_v^n$ at node v. Sequentially applying the above definitions (visiting the edges and nodes in an order that is consistent with the partial ordering imposed by directed graph $\mathcal{G}$) gives

$$\begin{aligned} \mathbf{x}_v^n &= \left((b_k^n)_{k \in \mathcal{K}: s(k,v)=1}, (w_e^n)_{e \in \mathcal{I}(v)} \right) \\ w_e^n &= f_e^n(\mathbf{x}_{t(e)}^n) \end{aligned}$$

when the values of the input processes (sample values of $B_1^n, \ldots, B_K^n$) are $b_1^n, \ldots, b_K^n$.

We say that $(\mathcal{G}, \mathcal{S}, \mathcal{D})$ is *achievable* if and only if there exists some $n \geq 1$ and a code (f^n, g^n) such that, if $\{\mathbf{x}_v^n\}$ are the output source vectors resulting from input source vectors $(b_k^n)_{k=1}^K$, then

$$\left(g_{k,v}^n(\mathbf{x}_v^n) \right)_{(k,v) \in \mathcal{K} \times \mathcal{V}} = (d(k,v) \cdot b_k^n)_{(k,v) \in \mathcal{K} \times \mathcal{V}}$$

for all $(b_k^n)_{k=1}^K \in \{0,1\}^{nK}$. Notice that for any (k,v) such that $d(k,v) = 0$, the decoding rule $g_{k,v}^n(\mathbf{x}_v^n) = 0^n$ for all $\mathbf{x}_v^n \in \mathcal{X}_v^n$ meets the above constraint; thus node v need only reconstruct the samples b_k^n for values of k where $d(k,v) = 1$.

Theorem 1 *For any $(\mathcal{G}, \mathcal{S}, \mathcal{D})$, there exists a corresponding $(\tilde{\mathcal{G}}, \tilde{\mathcal{S}}, \tilde{\mathcal{D}})$ such that only one node in $\tilde{\mathcal{G}}$ serves as a transmitter ($|\{v \in \mathcal{V} : \sum_{k=1}^K \tilde{s}(k,v) > 0\}| = 1$), each receiver in $\tilde{\mathcal{G}}$ demands only one source ($\sum_{k=1}^K \tilde{d}(k,v) \leq 1$ for all $v \in \mathcal{V}$), each receiver is connected to only one other node and has incoming capacity 1 (for each $v \in \mathcal{V}$ with $\sum_{k=1}^K \tilde{d}(k,v) = 1$, $|\mathcal{I}(v)| = 1$, $|\mathcal{O}(v)| = 0$, and $\sum_{e \in \mathcal{I}(v)} r(e) = 1$), and $(\mathcal{G}, \mathcal{S}, \mathcal{D})$ is achievable if and only if $(\tilde{\mathcal{G}}, \tilde{\mathcal{S}}, \tilde{\mathcal{D}})$ is achievable.*

Proof: Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{R})$, $\mathcal{S}$, and $\mathcal{D}$ be the given graph, source matrix, and demand matrix. We build $\tilde{\mathcal{G}} = (\tilde{\mathcal{V}}, \tilde{\mathcal{E}}, \tilde{\mathcal{R}})$, $\tilde{\mathcal{S}}$, and $\tilde{\mathcal{D}}$ by modifying $\mathcal{G}$, $\mathcal{S}$, and $\mathcal{D}$ as follows. First we add a single new transmitter node v_0 to serve as the single initiation point for all K sources. For each $v \in \mathcal{V}$ such that $\sum_{k=1}^{K} \tilde{s}(k,v) > 0$, we add a directed edge $e_{v_0,v}$ from v_0 to v. Next, for each $(k,v) \in \mathcal{K} \times \mathcal{V}$ that satisfies $s(k,v) = 1$ or $d(k,v) = 1$, we add a new node $\tilde{v}_{v,k}$ and a new edge $\tilde{e}_{v,k}$ from node v to node $\tilde{v}_{v,k}$. Let $\mathcal{V}' = \{v_0\}$, $\mathcal{E}' = \{e_{v_0,v} : \sum_{k=1}^{K} s(k,v) > 0\}$, $\mathcal{V}'' = \{\tilde{v}_{v,k} : (k,v) \in \mathcal{K} \times \mathcal{V} \wedge s(k,v) + d(k,v) = 1\}$, and $\mathcal{E}'' = \{\tilde{e}_{v,k} : (k,v) \in \mathcal{K} \times \mathcal{V} \wedge s(k,v) + d(k,v) = 1\}$ be the sets of added nodes and edges. Set $\tilde{\mathcal{V}} = \mathcal{V} \cup \mathcal{V}' \cup \mathcal{V}''$, $\tilde{\mathcal{E}} = \mathcal{E} \cup \mathcal{E}' \cup \mathcal{E}''$, $\tilde{\mathcal{R}} = [\tilde{r}(e)]_{e \in \tilde{\mathcal{E}}}$, $\tilde{\mathcal{S}} = [\tilde{s}_{k,v}]_{(k,v) \in \mathcal{K} \times \tilde{\mathcal{V}}}$, and $\tilde{\mathcal{D}} = [\tilde{d}_{k,v}]_{(k,v) \in \mathcal{K} \times \tilde{\mathcal{V}}}$, where

$$\begin{aligned}
\tilde{r}(e) &= \begin{cases} \sum_{k=1}^{K} s(k,v) & \text{if } e \in \mathcal{E}' \\ 1 & \text{if } e \in \mathcal{E}'' \\ r(e) & \text{otherwise} \end{cases} \\
\tilde{s}(k,v) &= \begin{cases} 1 & \text{if } v \in \mathcal{V}' \\ 0 & \text{otherwise} \end{cases} \\
\tilde{d}(k,v) &= \begin{cases} 1 & \text{if } v \in \mathcal{V}'' \\ 0 & \text{otherwise.} \end{cases}
\end{aligned}$$

If $(\mathcal{G}, \mathcal{S}, \mathcal{D})$ is achievable then there exists an $n \geq 1$ and a network code (f^n, g^n) such that for any input $(b_k^n)_{k=1}^{K}$,

$$\left(g_{k,v}^n(\mathbf{x}_v^n)\right)_{(k,v) \in \mathcal{K} \times \mathcal{V}} = \left(d(k,v) \cdot b_k^n\right)_{(k,v) \in \mathcal{K} \times \mathcal{V}}.$$

Define $(\tilde{f}^n, \tilde{g}^n)$

$$\begin{aligned}
\tilde{f}_e^n(\mathbf{x}_{t(e)}^n) &= \begin{cases} (b_k^n)_{k:s(k,h(e))=1} & \text{if } e \in \mathcal{E}' \\ g_{k,v}^n(\mathbf{x}_{t(e)}^n) & \text{if } e = e_{v,k} \in \mathcal{E}'' \\ f_e(\mathbf{x}_{t(e)}^n) & \text{otherwise} \end{cases} \\
\tilde{g}_{k,v}^n(\mathbf{x}_v^n) &= \begin{cases} 0^n & \text{if } v \in \mathcal{V}' \\ \mathbf{x}_v^n & \text{if } v \in \mathcal{V}'' \\ g_{k,v}^n(\mathbf{x}_v^n) & \text{otherwise.} \end{cases}
\end{aligned}$$

In words, the new transmitter distributes the sources to the transmitters from $\mathcal{G}$; each receiver from $\mathcal{G}$ decodes the desired sources and transmits them to its new children. Since the rates are sufficient for these tasks by construction,

$$\left(\tilde{g}_{k,v}^n(\mathbf{x}_v^n)\right)_{(k,v) \in \mathcal{K} \times \tilde{\mathcal{V}}} = \left(\tilde{d}(k,v) \cdot b_k^n\right)_{(k,v) \in \mathcal{K} \times \tilde{\mathcal{V}}}.$$

If $(\tilde{\mathcal{G}}, \tilde{\mathcal{S}}, \tilde{\mathcal{D}})$ is achievable, then there exists an $n \geq 1$ and a network code $(\tilde{f}^n, \tilde{g}^n)$ such that for all input processes $(b_k^n)_{k \in \mathcal{K}}$

$$\left(\tilde{g}_{k,v}^n(\mathbf{x}_v^n)\right)_{(k,v) \in \mathcal{K} \times \tilde{\mathcal{V}}} = \left(\tilde{d}(k,v) \cdot b_k^n\right)_{(k,v) \in \mathcal{K} \times \tilde{\mathcal{V}}}.$$

For any e such that $\sum_{k=1}^{K} s(k,t(e)) > 0$, let $\mathbf{x}_{t(e)}^n = (\mathbf{x}_{t(e),1}^n, \mathbf{x}_{t(e),2}^n)$, where $\mathbf{x}_{t(e),1}^n = (b_k^n)_{k:s(k,v)=1}$ is the source portion of $\mathbf{x}_{t(e)}^n$ and $\mathbf{x}_{t(e),2}^n$ is the portion of $\mathbf{x}_{t(e)}^n$ comprised of messages forwarded from other nodes in the network. Then define $\{(f^n, g^n)\}$ as

$$\begin{aligned}
f_e^n(\mathbf{x}_{t(e)}^n) &= \begin{cases} \tilde{f}_e^n(\tilde{f}_{e_{k_0,t(e)}}(\mathbf{x}_{t(e),1}^n), \mathbf{x}_{t(e),2}^n) & \text{if } \sum_{k=1}^{K} s_{k,t(e)} > 0 \\ \tilde{f}_e^n(\mathbf{x}_{t(e)}^n) & \text{otherwise} \end{cases} \\
g_{k,v}^n(\mathbf{x}_v^n) &= \begin{cases} \tilde{g}_{k,\tilde{v}_{k,v}}^n(\tilde{f}_{e_{k,v}}^n(\mathbf{x}_v^n)) & \text{if } \sum_{k=1}^{K} d(k,v) > 0 \\ \tilde{g}_{k,v}^n(\mathbf{x}_v^n) & \text{otherwise.} \end{cases}
\end{aligned}$$

The network code on the smaller graph $\mathcal{G}$ emulates the behavior of the network code on the larger graph $\mathcal{G}'$. Since the latter allows reconstruction of the sources at the receivers that demand them, the former does as well, giving

$$\left(g^n_{k,v}(\mathbf{x}^n_v)\right)_{(k,v)\in\mathcal{K}\times\mathcal{V}} = (d(k,v)\cdot b^n_k)_{(k,v)\in\mathcal{K}\times\mathcal{V}}$$

when the inputs are $(b^n_k)^n_{k=1}$. $\square$

The preceding theorem demonstrates that there is no loss in generality associated with restricting our attention to network coding problems with a single transmitter and a collection of receivers such that each receiver is connected to exactly one other node, desires exactly one source, and has incoming capacity exactly equal to 1. The remainder of this paper focuses on network coding problems (henceforth called *single-transmitter, single-demand* or STSD systems) that satisfy these conditions.

The following theorem demonstrates a form of equivalence between a collection of random variables and a family of functions of those random variables that meet a matroidal condition defined there. The result is useful for characterizing the set of achievable STSD problems and is also potentially interesting in its own right as a property of random variables.

Theorem 2 *Let $B^K = (B_1, \ldots, B_K)$ be independent, uniformly distributed bits. Suppose that there exists a collection $\{f_v\}_{v=1}^V$ of deterministic functions $f_v : \{0,1\}^K \to \{0,1\}$, and let $X_v = f_v(B^K)$ for each v. Then the following statements are equivalent.*

- *There exist deterministic mappings $h : \{0,1\}^K \to \{0,1\}^K$ and $g_v : \{0,1\} \to \{0,1\}$ such that*

  $$g_v(f_v(h(B^K))) = B_{k_v}$$

  for each $v \in \{1,\ldots,V\}$, where $\{k_v\}_{v\in\{1,\ldots,V\}}$ is a collection of integers such that the set $\{v : k_v = k\}$ is non-empty for every k.

- *There exists a $K \times V$ matrix $D = [d_{k,v}]$ such that $\sum_{k=1}^K d_{k,v} = 1$ for all v,*

  $$\sum_{k=1}^K d_{k,\{1,\ldots,V\}} = K,$$

  and for any $\mathcal{A} \subseteq \{1,\ldots,V\}$

  $$H(X_\mathcal{A}) = \sum_{k=1}^K d_{k,\mathcal{A}},$$

  where $X_\mathcal{A} = (X_v)_{v\in\mathcal{A}}$ and $d_{k,\mathcal{A}} = 1$ if $\sum_{v\in\mathcal{A}} d_{k,v} > 0$ and 0 otherwise.

Proof: Given the first condition, we satisfy the second condition by setting $d_{k,v} = 1$ if $k_v = k$ and 0 otherwise.

Given the second condition, fix an initial $\mathcal{A}$ to be a smallest possible subset of $\{1,\ldots,V\}$ such that $d(k,\mathcal{A}) = K$. Let $\{v_1,\ldots,v_K\}$ denote the K members of $\mathcal{A}$, where $d_{k,v_j} = 1$ if $k = j$ and 0 otherwise. For the given set,

$$K = H(X_\mathcal{A}) = H(B^K) = H(X_\mathcal{A}, B^K)$$

since the functions f_v are deterministic. Thus $H(B^K|X_{\mathcal{A}}) = H(X_{\mathcal{A}}|B^K) = 0$, and there is a one-to-one mapping $h : \{0,1\}^K \to \{0,1\}^K$ such that $X_{\mathcal{A}} = x_{\mathcal{A}}$ if and only if $B^K = h(x_{\mathcal{A}})$. Given this mapping,

$$X_{\mathcal{A}} = h^{-1}(B^K) = (f_{v_1}(B^K), \ldots, f_{v_K}(B^K)).$$

In this case, for any $b^K \in \{0,1\}^K$, $X_{\mathcal{A}} = b^K$ is achieved by $B^K = h(b^K)$, giving

$$f_{v_k}(h(b^K)) = b_k$$

for all $k \in \{1, \ldots, K\}$. Thus setting $g_v(x) = x$ for all $v \in \mathcal{A}$ meets the desired constraint on the given subset. Now for any $v \notin \mathcal{A}$ for which $d_{k,v} = 1$, $H(X_v) = H(X_{v_k}) = H(X_v, X_{v_k}) = 1$, where X_{v_k} is again the unique element of $\mathcal{A}$ that satisfies $d_{k,v_k} = 1$. Again $H(X_v|X_{v_k}) = H(X_{v_k}|X_v) = 0$ implies that there exists a one-to-one mapping between X_{v_k} and X_v. Let g_v denote the mapping that gives $X_{v_k} = g_v(X_v)$. The given $\{g_v\}_{v=1}^V$ and h together satisfy the first condition. □

Corollary 1 applies Theorem 2 to characterize the family of achievable STSD problems. The corollary demonstrates that $(\mathcal{G}, \mathcal{S}, \mathcal{D})$ is achievable if and only if there exists a collection of deterministic edge codes that achieves the desired entropies at the receivers in the sense described in Theorem 2.

Following the above notational conventions, Corollary 1 uses $d(k, \mathcal{A})$, where $d(k, \mathcal{A}) = 1$ if $\sum_{v \in \mathcal{A}} d(k, v) > 0$ and $d(k, \mathcal{A}) = 0$ otherwise, to denote the combined demand for source k among a collection of nodes $\mathcal{A} \subseteq \mathcal{V}$. Similarly, $\mathbf{X}_{\mathcal{A}}^n = (\mathbf{X}_v^n)_{v \in \mathcal{A}}$ describes the combined source output for all nodes in any $\mathcal{A} \subseteq \mathcal{V}$ resulting from the application of the blocklength-n code f^n. Finally, $\tilde{\mathcal{V}} = \{v \in \mathcal{V} : \sum_{k=1}^K d(k, v) = 1\}$ to denote the set of terminal vertices in an STSD graph.

Corollary 1 *Given an STSD problem $(\mathcal{G}, \mathcal{S}, \mathcal{D})$ that satisfies $d(k, \mathcal{V}) = 1$ for every $k \in \mathcal{K}$, $(\mathcal{G}, \mathcal{S}, \mathcal{D})$ is achievable if and only if there exist deterministic encoders f^n such that if the input $\mathbf{B} = (B_1, \ldots, B_K)$ is a collection of K independent binary sources drawn from the Bernoulli(1/2) distribution, then*

$$H(\mathbf{X}_{\mathcal{A}}^n) = \sum_{k=1}^K d(k, \mathcal{A})$$

for all $\mathcal{A} \subseteq \mathcal{V}$.

Proof: To apply the previous theorem directly, we consider, for any fixed n, a representation of the problem $(\mathcal{G}, \mathcal{S}, \mathcal{D})$ in which the sources, demands, and capacities associated with all time-instances $t \in \{1, \ldots, n\}$ are separately and explicitly represented. That is, each source input B_k^n is represented by n individual source inputs $B_k(1), \ldots, B_k(n)$, the demands for the n samples from each source are treated as n separate demands, and the capacity of edge e is given by $nr(e)$ to denote the total capacity over n channel uses. Applying the simplification from Theorem 1 to this time-separated graph gives a network with single-bit inputs and single-bit demands. Theorem 2 then yields the desired result; in this case, the functions under consideration describe the cumulative effect of the codes at all nodes through which the source bits have passed enroute to a given receiver. □

5 Conclusions

This paper poses a conjecture that linear coding suffices for network coding on arbitrary networks. The conjecture is supported by an example demonstrating that previously noted problems with linear coding can be addressed using linear network codes on vector inputs. We introduce a coding theorem for networks, which may be viewed as a coding theorem for functions of random variables. The coding theorem demonstrates that a functional input can be reconstructed from its mappings if and only if the entropies of all subsets of functional outputs have entropies identical to the entropies of the inputs they attempt to reconstruct. Applying this theorem to the outputs of a network code gives necessary and sufficient conditions for the achievability of a collection of demands on a fixed network in terms of the entropies of the received outputs. The coding theorem gives a new tool for investigating the sufficiency of linear coding in networks with non-multicast demands.

References

[1] R. Koetter and M. Médard. Beyond routing: an algebraic approach to network coding. In *Proceedings of INFOCOM 2002*, volume 1, pages 122–130, 2002.

[2] A. R. Lehman and E. Lehman. Complexity classification of network information flow problems. In *Proceedings of the Forty-first Allerton Conference on Control, Communication, and Computing*, Allerton, IL, October 2003.

[3] S.-Y.R. Li, R.W. Yeung, and N. Cai. Linear network coding. *IEEE Transactions on Information Theory*, IT-49(2):371–381, February 2003.