

On the Impact of a Single Edge on the Network Coding Capacity

Shirin Jalali

Center for Mathematics of Information
California Institute of Technology
shirin@caltech.edu

Michelle Effros

Departments of Electrical Engineering
California Institute of Technology
effros@caltech.edu

Tracey Ho

Department of Electrical Engineering
California Institute of Technology
tho@caltech.edu

Abstract—In this paper, we study the effect of a single link on the capacity of a network of error-free bit pipes. More precisely, we study the change in network capacity that results when we remove a single link of capacity δ . In a recent result, we proved that if all the sources are directly available to a single super-source node, then removing a link of capacity δ cannot change the capacity region of the network by more than δ in each dimension. In this paper, we extend this result to the case of multi-source, multi-sink networks for some special network topologies.

I. PROBLEM STATEMENT

Consider a communication problem defined by a network, a collection of sources, and a collection of sinks. The network is a directed graph with nodes representing communication devices and edges representing error-free, point-to-point communication channels with finite capacities. The sources are independent data streams, and each is available to precisely one node in the network. Each sink is a node in the network that desires some subset of the data streams; the desired subset may differ from one sink to the next. The capacity of the network, also called the “network coding capacity,” describes the set of achievable rates for every possible combination of sources and sinks. Solving for the capacity is a challenging open problem. In this paper, we investigate a simpler question: what is the effect of a single link on the network coding capacity of such a network? Specifically, we wish to understand whether decreasing the capacity of a single edge e from $C_e \geq \delta$ to $C_e - \delta$ can change the capacity region of the network by more than δ in each dimension.

In [1], we posed this question and proved that if all sources are available at one node, then changing the capacity of a single link by δ reduces each achievable rate vector by at most δ in each dimension. In this paper, we extend this result to a family of multi-source, multi-sink networks.

II. NOTATION

Throughout the paper, finite sets are denoted by script letters such as \mathcal{X} and \mathcal{Y} . The size of a finite set \mathcal{A} is denoted by $|\mathcal{A}|$. Random variables are denoted by upper case letters such as X and Y . We represent the alphabet of random variable X by \mathcal{X} . Bold letters, for example $\mathbf{X} = (X_1, \dots, X_n)$ and $\mathbf{x} = (x_1, \dots, x_n)$ represent vectors. The length of a vector is implied in the context, and its ℓ^{th} element is denoted by X_ℓ . For a set $\mathcal{F} \subseteq \{1, 2, \dots, n\}$, $\mathbf{x}_{\mathcal{F}} = (x_i)_{i \in \mathcal{F}}$, where the

elements are sorted in ascending order of their indices. For a vector $\mathbf{X} \in \mathbb{R}^n$, let $\mathbf{X}^+ = \max(\mathbf{0}, \mathbf{X})$, where $\mathbf{0}$ is a zero-valued vector of length n , and the max operator is applied component-wise.

III. SYSTEM MODEL

Consider an acyclic error-free network \mathcal{N} denoted by a directed graph $G = (\mathcal{V}, \mathcal{E})$ with nodes \mathcal{V} and edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. Each edge $e = (v_1, v_2) \in \mathcal{E}$ represents an error-free channel from Node v_1 to Node v_2 . We use $C_e > 0$ to denote that channel’s capacity. For each node $v \in \mathcal{V}$, $\text{In}(v) = \{(v_1, v) : (v_1, v) \in \mathcal{E}\}$ and $\text{Out}(v) = \{(v, v_1) : (v, v_1) \in \mathcal{E}\}$ denote the set of incoming and outgoing edges for Node v respectively.

Let $\mathcal{S} = \{1, 2, \dots, k\}$ denote the set of sources available in the network, and let

$$\alpha : \mathcal{S} \rightarrow \mathcal{V},$$

specify the source availability. Thus for each $s \in \mathcal{S}$, $\alpha(s)$ describes the unique node where source s is available. Likewise, for each $v \in \mathcal{V}$, let $\sigma(v) \subseteq \mathcal{S}$ denote the set of sources observed by Node v , i.e.,

$$\sigma(v) = \{s : \alpha(s) = v\}.$$

Finally, for each $v \in \mathcal{V}$, let $\beta(v) \subseteq \mathcal{S}$ denote the set of sources that Node v is interested in recovering.

A network code of block length n and rate $\mathbf{R} = (R_s)_{s \in \mathcal{S}}$ over such a network is described as follows. Each source $s \in \mathcal{S}$ generates some message $M_s \in \mathcal{M}_s = \{1, 2, \dots, 2^{nR_s}\}$. For each $e \in \mathcal{E}$, let $\mathcal{W}_e = \{1, 2, \dots, 2^{nC_e}\}$. The coding operations performed by each node can be categorized as follows

- 1) Encoding functions:

For each $v \in \mathcal{V}$ and $e \in \text{Out}(v)$, the encoding function corresponding to Edge e is a mapping

$$g_e : \prod_{s \in \sigma(v)} \mathcal{M}_s \times \prod_{e' \in \text{In}(v)} \mathcal{W}_{e'} \rightarrow \mathcal{W}_e.$$

- 2) Decoding functions:

For each $v \in \mathcal{V}$ and $s \in \beta(v)$, the decoding function for source s at Node v is a mapping

$$g_v^s : \prod_{s' \in \sigma(v)} \mathcal{M}_{s'} \times \prod_{e \in \text{In}(v)} \mathcal{W}_e \rightarrow \mathcal{M}_s.$$

A rate vector $\mathbf{R} = (R_s)_{s \in \mathcal{S}}$ is said to be achievable on network \mathcal{N} , if for any $\epsilon > 0$, there exists a block length n large enough and a coding scheme of block length n operating at rate \mathbf{R} such that for all $v \in \mathcal{V}$ and $s \in \beta(v)$

$$\mathbb{P}(\hat{M}_s^{(v)} \neq M_s) \leq \epsilon,$$

where $\hat{M}_s^{(v)}$ denotes the reconstruction of message M_s at Node v . For sources \mathcal{S} , availability mapping $\alpha(\cdot)$, and demand mapping $\beta(\cdot)$, let $\mathcal{R}(\mathcal{N}, \mathcal{S}, \alpha, \beta)$ denote the set of achievable rates on Network \mathcal{N} .

In the discussion that follows, we use \mathcal{N} to describe the original network and \mathcal{N}' to describe the new network that results when we reduce the capacity of a single, fixed edge $e \in \mathcal{E}$ from $C_e \geq \delta$ to $C'_e = C_e - \delta$. If $C_e = \delta$, then edge e is removed from \mathcal{N} to obtain \mathcal{N}' .

IV. PRIOR WORK

Network codes are communication schemes in which every node is allowed to perform arbitrary functions on its inputs in creating its outputs. The idea was first proposed by Ahlswede, Cai, Li, and Yeung in 2000 [2]. They proved that Ford and Fulkerson's famous max-flow min-cut theorem for unicast networks [3], also holds in multicast networks. (Here a "unicast network" refers to a network with a single source and a single sink node, while a "multicast network" refers to a network with one source and multiple sink nodes, each requiring all data available at the source.) While it is always possible to achieve the capacity in a unicast network using only routing at the relay nodes, Ahlswede et al. showed that there exist networks where coding is required to achieve the multicast capacity. Linear coding operations suffice for achieving the capacity of a multicast network by [4]. While both the capacity region and the structure of capacity-achieving codes are known for multicast demands, neither the capacity nor a low-complexity family of codes sufficient for achieving the capacity is known for most demand types. Linear codes are insufficient for achieving the capacity under general demands by [5].

Computing the capacity region of an error-free network can be cast as a convex optimization problem with a linear cost function over the space of *normalized entropic vectors* with some other linear constraints [6][7]. This characterization reveals that network information theory problems over noiseless networks could be solved if we could explicitly characterize the set of entropy vectors. While there has been a lot of effort in recent years geared towards developing a better understanding of the set of entropy vectors (c.f. [8], [9], [10], [11], [12], [13]), to date the problem remains largely unsolved.

In this paper, we study the problem from a different perspective. Instead of trying to find the capacity region of a network, we focus on the effect of a single link on that capacity region. Precisely, we try to understand the effect on network capacity of changing the capacity on a single edge $e \in \mathcal{E}$ from $C_e \geq \delta$ to $C'_e = C_e - \delta$, which effectively changes just one linear constraint in the problem as described above.

V. RESULTS

Before stating our main result in Section V-E, we briefly review some cases where the impact, in terms of network capacity, of reducing C_e is already known or straightforward to characterize.

A. Demand Types with Tight Cut-Set Bounds

For a variety of demand types, including multicast, multi-source multicast, single-source with non-overlapping demands, and single-source with non-overlapping demands and a multicast demand, network coding capacity can be fully characterized by the corresponding cut-set bounds [14]. Reducing C_e to $C_e - \delta$ for a single edge $e \in \mathcal{E}$ reduces the capacity of every cut by at most δ . Therefore, if $(\mathcal{S}, \alpha, \beta)$ describes any such demand type, and $\mathbf{R} \in \mathcal{R}(\mathcal{N}, \mathcal{S}, \alpha, \beta)$, then $(\mathbf{R} - \delta \cdot \mathbf{1})^+ \in \mathcal{R}(\mathcal{N}', \mathcal{S}, \alpha, \beta)$, where \mathcal{N}' is the modified network, as described in Section III, and $\mathbf{1}$ is the all-ones vector.

B. Links Connected to Terminal Nodes

Consider a terminal node $v_o \in \mathcal{V}$; then Node v_o has no outgoing edges ($\text{Out}(v_o) = \emptyset$). Let $p = |\text{In}(v_o)|$ denote the number of edges incoming to v_o , and let W_1, W_2, \dots, W_p denote the messages carried by these links. Further, assume that the link corresponding to the message W_1 has capacity δ . For any $s \in \beta(v_o)$,

$$\begin{aligned} I(M_s; W_2, \dots, W_p) &= I(M_s; W_1, W_2, \dots, W_p) - I(M_s; W_1 | W_2, \dots, W_p) \\ &\geq I(M_s; W_1, W_2, \dots, W_p) - H(W_1) \\ &\geq I(M_s; W_1, W_2, \dots, W_p) - n\delta. \end{aligned}$$

This proves that removing this link reduces the capacity from source s to node v by at most δ . Since Node v has only incoming edges, this change does not affect the capacities at any other nodes in the network. As a result, applying, for each $s \in \sigma(v)$, an outer code with rate $R_s - \delta$ and codewords drawn uniformly at random yields expected error probability approaching 0 as the coding dimension grows without bound. This proves the existence of a good collection of codes. Therefore, $\mathbf{R} = (R_s : s \in \mathcal{S}) \in \mathcal{R}(\mathcal{N}, \mathcal{S}, \alpha, \beta)$, implies $\mathbf{R}' = (R'_s : s \in \mathcal{S}) \in \mathcal{R}(\mathcal{N}', \mathcal{S}, \alpha, \beta)$, where $R'_s = R_s$ for all $s \in \mathcal{S} \setminus \sigma(v)$ and $R'_s = (R_s - \delta)^+$ for all $s \in \sigma(v)$.

C. Super Source Node

For the case where all the sources are available to a *super source node* ($\sigma(v_o) = \mathcal{S}$ for some $v_o \in \mathcal{V}$, as shown in Fig. 1), we showed in [1] that changing the capacity of any link $e \in \mathcal{E}$ from $C_e \geq \delta$ to $C'_e = C_e - \delta$ changes the network capacity region by at most δ in each dimension (i.e., $\mathbf{R} \in \mathcal{R}(\mathcal{N}, \mathcal{S}, \alpha, \beta)$ implies $(\mathbf{R} - \delta \cdot \mathbf{1})^+ \in \mathcal{R}(\mathcal{N}', \mathcal{S}, \alpha, \beta)$).




Fig. 1. All sources available directly at a super source node v_o

D. Linear Network Coding

Consider a linear network code of block length n and rate $\mathbf{R} = (R_s)_{s \in \mathcal{S}}$ operating on network \mathcal{N} . Let $e \in \mathcal{E}$ be a fixed link of capacity $C_e = \delta$ inside this network. In this case, we treat both source messages and the messages traversing each link in the network as binary vectors. Since the code is linear, the message W_e sent across link e can be written as a linear combination of the source messages $\{M_s\}_{s \in \mathcal{S}}$. Precisely,

$$W_e = \sum_{s \in \mathcal{S}} A_{s,e} M_s, \quad (1)$$

where for each $s \in \mathcal{S}$, $A_{s,e}$ denotes a binary matrix of dimension $nC_e \times nR_s$ and all additions in (1) are binary additions. Let \mathcal{M}_0 denote the set of messages that yield message $W_e = \mathbf{0}$ on link e using the given linear code, i.e.,

$$\mathcal{M}_0 \triangleq \{(M_s)_{s \in \mathcal{S}} : \sum_{s \in \mathcal{S}} A_{s,e} M_s = \mathbf{0}\}.$$

If we restrict our attention to this subset of messages, then we can run the given linear code in the absence of edge e since the value of W_e for all such messages is fixed and known. Unfortunately, choosing messages from \mathcal{M}_0 may require coordination among the source nodes. We therefore choose messages from a subset of \mathcal{M}_0 that requires no such coordination. Namely, we transmit only messages from \mathcal{M}_{00} , where \mathcal{M}_{00} is defined as

$$\mathcal{M}_{00} \triangleq \{(M_s)_{s \in \mathcal{S}} : A_{s,e} M_s = \mathbf{0} \text{ for all } s \in \mathcal{S}\}.$$

By sending only messages $(M_s)_{s \in \mathcal{S}} \in \mathcal{M}_{00}$, we guarantee that $W_e = \mathbf{0}$; since $\mathcal{M}_{00} = \prod_{s \in \mathcal{S}} \{M_s : A_{s,e} M_s = \mathbf{0}\}$, the source nodes can transmit only messages from \mathcal{M}_{00} without coordination. The resulting rate is $(1/n) \log |\{M_s : A_{s,e} M_s = \mathbf{0}\}| \geq (R_s - \delta)^+$ for each $s \in \mathcal{S}$. Thus we can apply the code from \mathcal{N} on the network \mathcal{N}' to achieve reliable communication at rate $(\mathbf{R} - \delta \cdot \mathbf{1})^+$.

The given argument demonstrates that removing a single link of capacity $C_e = \delta$ changes the rate achievable with linear coding by at most δ in each dimension. The same argument can be used to show that reducing the capacity of some edge e with $C_e > \delta$ to $C'_e = C_e - \delta$ reduces the rate achievable with linear coding by at most δ in each dimension. This can be seen by treating a link of capacity $C_e > \delta$ as a pair of parallel links of capacities $C_e - \delta$ and δ , respectively, and applying the previous argument.

Unfortunately, as noted in Section IV, linear network codes are not sufficient for achieving the capacity of general error-free networks. Thus, the given strategy proves only that reducing the capacity of a link by δ changes the set of rates achievable using linear coding by at most δ in each dimension. If rate \mathbf{R} is achievable using linear coding on \mathcal{N} , then rate $(\mathbf{R} - \delta \cdot \mathbf{1})^+$ is achievable using linear coding on \mathcal{N}' .

E. Main Result

Consider the k -unicast network \mathcal{N} shown in Fig. 2(a). Here, $\alpha(s) = v_s$ and $\beta(v_{k+s}) = \{s\}$ for all $s \in \mathcal{S}$; that is, each message $s \in \mathcal{S}$ is a unicast from node v_s to node v_{k+s} . In a blocklength- n code, $\mathcal{M}_s \in \{1, 2, \dots, 2^{nR_s}\}$ denotes the source message for Source s , and \hat{M}_s represents the reconstruction of M_s at sink node v_{k+s} . When we remove the link e of capacity $C_e = \delta$ from \mathcal{N} , we obtain the network \mathcal{N}' shown in Fig. 2(b).




Fig. 2. A multiple unicast network with special structure

Theorem 1: For any $\mathbf{R} \in \mathcal{R}(\mathcal{N}, \mathcal{S}, \alpha, \beta)$,

$$(\mathbf{R} - \delta \cdot \mathbf{1})^+ \in \mathcal{R}(\mathcal{N}', \mathcal{S}, \alpha, \beta).$$

Proof: Fix $\mathbf{R} = (R_1, R_2, \dots, R_k) \in \mathcal{R}(\mathcal{N}, \mathcal{S}, \alpha, \beta)$. We first consider the case where $\min\{R_1, \dots, R_k\} \geq \delta$. Given a code of blocklength n , for each $s \in \mathcal{S}$, let $P_{e,s}^{(n)} \triangleq \mathbb{P}(M_s \neq \hat{M}_s)$ denote the error probability in reconstructing source s at sink v_{k+s} . For any $p \in [0, 1]$, let $h(p) = -p \log(p) - (1-p) \log(1-p)$ be the binary entropy function. Since \mathbf{R} is achievable on \mathcal{N} , for any $\epsilon > 0$ and n large enough there exists a rate- \mathbf{R} code of blocklength n such that $\max\{P_{e,s}^{(n)} : s \in \mathcal{S}\} \leq \epsilon$ and $\max\{h(P_{e,s}^{(n)}) : s \in \mathcal{S}\} \leq \epsilon$. Given any $\epsilon > 0$, fix such a code. We next use this family of codes to prove the existence of a multiple access code for communicating the sources from nodes v_1, \dots, v_k to node a and a broadcast code for transmitting all sources $s \in \mathcal{S}$ from Node a to nodes v_{k+1}, \dots, v_{2k} , respectively, both at rates $\mathbf{R} - \delta \cdot \mathbf{1}$. In the arguments that follows, we use W_e , \mathbf{W}_i , and \mathbf{W}_o to denote the message sent through the link e of capacity $C_e = \delta$, the inputs of Node a , and the outputs of Node a , respectively (see Fig. 2(a)).

Consider the k -user multiple access channel with inputs $\mathbf{M} = (M_1, M_2, \dots, M_k)$ and output \mathbf{W}_i . The capacity region of this k -user MAC is the set of rate vectors $\mathbf{r} = (r_1, r_2, \dots, r_k)$ satisfying

$$\sum_{s \in \mathcal{A}} r_s \leq I(\mathbf{M}_{\mathcal{A}}; \mathbf{W}_i | \mathbf{M}_{\mathcal{A}^c}, Q),$$

for all $\mathcal{A} \subseteq \mathcal{S}$ and some

$$p(q)p(m_1|q)p(m_2|q) \dots p(m_k|q).$$

Define

$$\mathbf{r}_{\text{mac}} \triangleq (I(M_1; \mathbf{W}_i), I(M_2, \mathbf{W}_i), \dots, I(M_k; \mathbf{W}_i)),$$

under the distribution imposed by the code fixed above. In the argument that follows, we first show that \mathbf{r}_{mac} falls in the capacity region of the MAC and then prove that \mathbf{r}_{mac} satisfies the desired rate constraint.

Since the messages M_1, \dots, M_k are independent, for any sets $\mathcal{A} \subseteq \mathcal{S}$ and $\mathcal{A}^c = \mathcal{S} \setminus \mathcal{A}$,

$$\begin{aligned} \sum_{s \in \mathcal{A}} r_{\text{mac},s} &= \sum_{s \in \mathcal{A}} I(M_s; \mathbf{W}_i) \\ &= \sum_{s \in \mathcal{A}} [H(M_s) - H(M_s | \mathbf{W}_i)] \\ &= H(M_{\mathcal{A}}) - \sum_{s \in \mathcal{A}} H(M_s | \mathbf{W}_i) \\ &\leq H(M_{\mathcal{A}}) - H(M_{\mathcal{A}} | \mathbf{W}_i) \\ &\leq H(M_{\mathcal{A}}) - H(M_{\mathcal{A}} | \mathbf{W}_i, M_{\mathcal{A}^c}) \\ &= I(M_{\mathcal{A}}; \mathbf{W}_i | M_{\mathcal{A}^c}). \end{aligned}$$

Thus, \mathbf{r}_{mac} falls in the capacity region of the MAC.

We next bound each term in \mathbf{r}_{mac} . For each $s \in \mathcal{S}$,

$$\begin{aligned} H(M_s | \mathbf{W}_i) &\leq H(M_s, W_e | \mathbf{W}_i) \\ &= H(M_s | W_e, \mathbf{W}_i) + H(W_e | \mathbf{W}_i) \\ &\leq nR_s P_{e,s}^{(n)} + h(P_{e,s}^{(n)}) + n\delta \end{aligned}$$

by Fano's inequality [15]. Hence,

$$\begin{aligned} I(M_s; \mathbf{W}_i) &= H(M_s) - H(M_s | \mathbf{W}_i) \\ &\geq n(R_s - \delta) - nR_s\epsilon - \epsilon, \end{aligned} \quad (2)$$

since $\max\{P_{e,s}^{(n)}, h(P_{e,s}^{(n)})\} \leq \epsilon$ by assumption. Recall that $\epsilon > 0$ is arbitrary; thus (2) implies that $(\mathbf{R} - \delta \cdot \mathbf{1})$ is achievable on the described MAC.

We next deliver these messages to their intended receivers using the broadcast channel (BC) from Node a to the sinks v_{k+1}, \dots, v_{2k} . Again, we apply the previously chosen code, operating the code in the absence of edge e by sending only source messages for which the message across edge e is a fixed value w_e to be chosen next.

By the data processing inequality and Fano's inequality [15],

$$\begin{aligned} H(\mathbf{W}_o | W_e) &= H(\mathbf{W}_o, W_e) - H(W_e) \\ &\geq I(\mathbf{W}_o, W_e; M_1, \dots, M_k) - H(W_e) \\ &\geq I(\hat{M}_1, \dots, \hat{M}_k; M_1, \dots, M_k) - n\delta \\ &\geq n \sum_{s=1}^k R_s - H(M_1, \dots, M_k | \hat{M}_1, \dots, \hat{M}_k) - n\delta \\ &\geq n \sum_{s=1}^k R_s - n \sum_{s=1}^k R_s \epsilon - k\epsilon - n\delta. \end{aligned}$$

Since $H(\mathbf{W}_o | W_e) = \sum_{w_e \in \mathcal{W}_e} H(\mathbf{W}_o | W_e = w_e) p(w_e)$, there exists some $w_e \in \mathcal{W}_e$ such that

$$H(\mathbf{W}_o | W_e = w_e) \geq n \sum_{s=1}^k R_s - n \sum_{s=1}^k R_s \epsilon - k\epsilon - n\delta. \quad (3)$$

Moreover, for any $w_e \in \mathcal{W}_e$

$$H(M_1, \dots, M_k | W_e = w_e) \geq H(\mathbf{W}_o | W_e = w_e) \quad (4)$$

since \mathbf{W}_o is a deterministic function of M_1, \dots, M_k .

Fixing the message W_e to a value of w_e that satisfies (3), we get a k -user deterministic broadcast channel (BC) [15] with input \mathbf{W}_o and outputs $(\hat{M}_1, \dots, \hat{M}_k)$. Appendix A summarizes prior results on the capacity region for this BC, which achieves reliable transmission at all rates $\mathbf{r} = (r_1, r_2, \dots, r_k)$ for which

$$\sum_{s \in \mathcal{A}} r_s \leq H(\hat{M}_{\mathcal{A}} | W_e = w_e),$$

for all $\mathcal{A} \subseteq \mathcal{S}$. We now prove that this set of rates includes the rate $\mathbf{r}_{\text{bc}} = n(\mathbf{R} - \delta \cdot \mathbf{1})$. Note from (3) and (4) that

$$\begin{aligned} &H(\hat{M}_1, \dots, \hat{M}_k | W_e = w_e) \\ &\geq I(M_1, \dots, M_k; \hat{M}_1, \dots, \hat{M}_k | W_e = w_e) \\ &\geq n \sum_{s=1}^k R_s - 2n \sum_{s=1}^k R_s \epsilon - 2k\epsilon - n\delta. \end{aligned} \quad (5)$$

On the other hand, for any $\mathcal{A} \subseteq \mathcal{S}$,

$$\begin{aligned} &H(\hat{M}_{\mathcal{A}} | W_e = w_e) + H(\hat{M}_{\mathcal{A}^c} | W_e = w_e) \\ &\geq H(\hat{M}_1, \dots, \hat{M}_k | W_e = w_e), \end{aligned} \quad (6)$$

where $H(\hat{M}_{\mathcal{A}^c} | W_e = w_e) \leq \sum_{s \in \mathcal{A}^c} nR_s$. Hence, combining (5) and (6),

$$H(\hat{M}_{\mathcal{A}} | W_e = w_e) \geq n \sum_{s \in \mathcal{A}} R_s - 2n \sum_{s=1}^k R_s \epsilon - 2k\epsilon - n\delta.$$

Thus, since ϵ is arbitrary, $n(\mathbf{R} - \delta \cdot \mathbf{1})$ is achievable on the given BC. This implies that the messages received by node a at rate \mathbf{r}_{mac} can be delivered to their intended receivers, which concludes the proof for the case where $R_s > \delta$ for all $s \in \mathcal{S}$.

Finally, note that if there are some sources with $R_s \leq \delta$, then we can use the same argument by sending constant messages for all such sources in both the MAC and the BC. \blacksquare




Fig. 3. A special case of the network shown in Fig. 2(a)

A special case of the network shown in Fig. 2(a) is shown in Fig. 3. Theorem 1 immediately applies.

Note that Theorem 1 can also be used to derive an outer bound on the capacity region of the k -unicast network \mathcal{N} shown in Fig. 2(a). Let $\mathcal{R}_1 \triangleq \mathcal{R}(\mathcal{N}_1, \mathcal{S}, \alpha_1, \beta_1)$ and $\mathcal{R}_2 \triangleq \mathcal{R}(\mathcal{N}_2, \mathcal{S}, \alpha_2, \beta_2)$ denote the capacity regions of the networks \mathcal{N}_1 and \mathcal{N}_2 shown in Fig. 2(b), with $\alpha_1(s) = v_s$, $\alpha_2(s) = a$ and $\beta_2(v_{s+k}) = s$, for $s \in \mathcal{S}$. Moreover, $\beta_1(a) = \mathcal{S}$, $\beta_1(v) = \emptyset$ for $v \in \mathcal{V} \setminus a$, and $\beta_2(v) = \emptyset$ for $v \in \mathcal{V} \setminus \{v_{k+1}, \dots, v_{2k}\}$. Note that \mathcal{R}_1 and \mathcal{R}_2 correspond to a multicast network and a single source network with non-overlapping demands, respectively. Hence, as mentioned before, in both cases the capacity regions are computable and are fully characterized by the cut-set bounds [14].

Corollary 1: Let $\mathcal{R}_o \triangleq \{R + \delta \cdot \mathbf{1} : R \in \mathcal{R}_1 \cap \mathcal{R}_2\}$. Then,

$$\mathcal{R}(\mathcal{N}, \mathcal{S}, \alpha, \beta) \subseteq \mathcal{R}_o.$$

VI. CONCLUSION

In this paper we study the effect of a single link on the network coding capacity of a network of error-free bit pipes. For some special topologies of multi-source multi-sink networks, we prove that our result from [1] continues to hold; that is, reducing the capacity of a link by δ changes the capacity region by at most δ in each dimension. The question of whether or not this result holds for all networks remains an open area for future research.

APPENDIX A

DETERMINISTIC BROADCAST CHANNEL

A k -user deterministic broadcast channels (DBC) with input $x \in \mathcal{X}$ and outputs $\{Y_s \in \mathcal{Y}_s\}_{s \in \mathcal{S}}$ is a k -user broadcast channel such that for any $x \in \mathcal{X}$ and $(y_1, \dots, y_k) \in \mathcal{Y}_1 \times \mathcal{Y}_2 \times \dots \times \mathcal{Y}_k$,

$$P((Y_1, \dots, Y_k) = (y_1, \dots, y_k) | X = x) \in \{0, 1\}. \quad (\text{A-1})$$

Since the capacity region of a BC depends only on the receivers' conditional marginal distributions [15], (A-1) implies that a K -user DBC can be described by k functions (f_1, \dots, f_k) ,

$$f_s : \mathcal{X} \rightarrow \mathcal{Y}_s,$$

such that $Y_s = f_s(X)$ for $s \in \mathcal{S}$.

While the capacity region for general BCs remains unsolved, the capacity region of a k -user DBC is known and can

be described by the union of the set of rates (R_1, R_2, \dots, R_k) satisfying

$$\sum_{s \in \mathcal{A}} R_s \leq H(Y_{\mathcal{A}}),$$

for any $\mathcal{A} \subseteq \{1, \dots, k\}$, for some $P(X)$ [16], [17].

ACKNOWLEDGMENTS

This work was supported in part by Caltech's Center for the Mathematics of Information (CMI), DARPA ITMANET grant W911NF-07-1-0029, the Air Force Office of Scientific Research under grant FA9550-10-1-0166, and Caltech's Lee Center for Advanced Networking.

REFERENCES

- [1] H. Tracey, M. Effros, and S. Jalali. On equivalence between network topologies. In *48th Annu. Allerton Conf. Communication, Control, and Computing*, Sep 2010.
- [2] R. Ahlswede, Ning Cai, S.-Y.R. Li, and R.W. Yeung. Network information flow. *IEEE Trans. Inform. Theory*, 46(4):1204–1216, July 2000.
- [3] L. R. Ford and D. R. Fulkerson. Maximal ow through a network. *Canad. J. Math.*, 8:399–404, 1956.
- [4] S.-Y.R. Li, R.W. Yeung, and Ning Cai. Linear network coding. *IEEE Trans. Inform. Theory*, 49(2):371–381, February 2003.
- [5] R. Dougherty, C. Freiling, and K. Zeger. Insufficiency of linear coding in network information flow. In *Proc. IEEE Int. Symp. Inform. Theory*, pages 264–267, September 2005.
- [6] B. Hassibi and S. Shadbakht. Normalized entropy vectors, network information theory and convex optimization. In *IEEE Workshop on Information Theory*, pages 1–5, July 2007.
- [7] X. Yan, R. W. Yeung, and Z. Zhang. The capacity region for multi-source multi-sink network coding. In *Proc. IEEE Int. Symp. Inform. Theory*, pages 116–120, June 2007.
- [8] R.W. Yeung. A framework for linear information inequalities. *Information Theory, IEEE Transactions on*, 43(6):1924–1934, November 1997.
- [9] Z. Zhang and R.W. Yeung. A non-shannon-type conditional inequality of information quantities. *Information Theory, IEEE Transactions on*, 43(6):1982–1986, November 1997.
- [10] Ho-Leung Chan and R.W. Yeung. A combinatorial approach to information inequalities. In *Information Theory and Networking Workshop, 1999*, page 63, 1999.
- [11] T.H. Chan and R.W. Yeung. On a relation between information inequalities and group theory. *Information Theory, IEEE Transactions on*, 48(7):1992–1995, July 2002.
- [12] R. Dougherty, C. Freiling, and K. Zeger. Six new non-shannon information inequalities. In *Proc. IEEE Int. Symp. Inform. Theory*, pages 233–236, July 2006.
- [13] F. Matus. Infinitely many information inequalities. In *Proc. IEEE Int. Symp. Inform. Theory*, pages 41–44, June 2007.
- [14] R. Koetter and M. Medard. An algebraic approach to network coding. *IEEE/ACM Trans. Networking*, 11(5):782795, 2003.
- [15] T. Cover and J. Thomas. *Elements of Information Theory*. Wiley, New York, 2nd edition, 2006.
- [16] K. Marton. The capacity region of deterministic broadcast channels. In *Proc. IEEE Int. Symp. Inform. Theory*, Paris-Cachan, France, 1977.
- [17] M. S. Pinsker. Capacity of noiseless broadcast channels. *Probl. Inform. Transm.*, pages 92–102, 1978.