

Capacity Bounds for Networks of Broadcast Channels

Michelle Effros

Department of Electrical Engineering
California Institute of Technology

Abstract—This paper derives new outer bounds on the capacities of networks comprised of broadcast and point-to-point channels. The results are tight in some cases, and methods for bounding their error in general are discussed. The results given demonstrate the simplicity of generalizing network coding results to networks of noisy channels using the family of network equivalence tools. The approach taken is not inherently a cut-set approach and frequently yields tighter bounds than those achieved by traditional cut-sets.

I. INTRODUCTION

To date, the family of tools available for finding outer bounds on capacity regions of general networks under general demands is quite limited. Central among these tools are the cut-set bounds of [1, Theorem 15.10.1], which bound the rate that can flow from the nodes in any set S to all other nodes S^c in the network by the capacity that would be obtained if S and S^c could each act as a single supernode. More analytical tools (see, for example, [2], [3], [4], [5]) and computational tools (for example, [6], [4]) are available for finding general capacity outer bounds in networks of point-to-point noiseless links. While some of these tools also extend to point-to-point noisy channels [7], [8], [5], the wireline network coding literature focuses primarily on networks of noiseless links.

Until recently, only the cut-set bounds of [1, Theorem 15.10.1] applied for upper bounding the capacities of wireless networks. Cut-set bounds are known to be tight for a small number of wireless network models (e.g., networks of deterministic broadcast channels under multicast demands [9]). The gap between cut-set rates and capacities is non-zero even for many three-node networks [1, Section 15.10] and can be arbitrarily large in networks of noiseless links [3]. More recently, some of the more sophisticated generalizations on cut-set bounds have been generalized to a family of wireless network models with broadcast transmissions, interference receptions, and independent vertex noise [10].

In [11], [12], [13], the author and her collaborators bound the capacity of arbitrary networks of noisy point-to-point, broadcast, multiple access, and interference channels by the capacities of corresponding networks of noiseless, point-to-point bit pipes. The central result for point-to-point channels [11, Theorem 3] proves the optimality of the separation of channel and network coding in wireline networks. Precisely, it proves that the capacity of a network \mathcal{N} of noisy point-to-point channels equals the capacity of a network $\hat{\mathcal{N}}$ of noiseless bit pipes, where each noisy channel in \mathcal{N} is replaced

in $\hat{\mathcal{N}}$ by a bit pipe of the corresponding capacity. While achievability follows from the channel capacity, the converse does not since obtaining the optimal network performance sometimes requires operating individual channels above their capacities [11]. The separation result in [11] generalizes [7], [8] from multicast to general demands.

Applying network equivalence techniques to a network \mathcal{N} of noisy broadcast, multiple access, and interference channels gives outer bounds on the capacity of \mathcal{N} in terms of the capacity of a (family) of networks $\{\mathcal{N}(p)\}_p$ of noiseless, point-to-point bit pipes [12]. The resulting bounds are not always achievable, so the use of noiseless bit pipes does not imply separation. In fact, separation is known not to hold for networks of noisy broadcast channels [14] or even networks of noiseless broadcast channels [9] in which each channel output is a deterministic function of the channel input.

This work demonstrates the simplicity of deriving capacity outer bounds for wireless networks under general demands using the equivalence tools of [11, Theorem 3] and [12, Theorem 4]. Like previous works (e.g., [15], [9], [16]), the model given focuses on the broadcast nature of wireless channels assuming no interference. This choice is made only for simplicity since the multiple access and interference channel models of [12] are easy to apply. (Their evaluation may be more difficult since they rely on auxiliary random variables.)

Results for a variety of multicast and non-multicast demand problems are included. These results are tight in some cases (including networks where separation is known not to be optimal). Bounds on their accuracy are discussed for others. The bounds given differ fundamentally from cut-set bounds even for problems like multicast demands where cut-set bounds on the corresponding noiseless network would be tight.

Section II gives the problem set-up and background. Section III includes results and examples.

II. PRELIMINARIES

Consider an m -node network \mathcal{N} . At each time t , node v transmits a random variable $X_t^{(v)}$ and receives a random variable $Y_t^{(v)}$. The network is memoryless, so it is characterized by a conditional probability distribution

$$p(\mathbf{y}|\mathbf{x}) = p(y^{(1)}, \dots, y^{(m)} | x^{(1)}, \dots, x^{(m)}).$$

The network structure is described by a hypergraph with node set $V = \{1, \dots, m\}$ and hyperedge set $E = \{e_1, \dots, e_k\}$. Each hyperedge $e = [V_1, V_2]$ has a collection of input nodes

$V_1 \subset V$ and a set of output nodes $V_2 \subseteq V \setminus V_1$. This work focuses on networks of point-to-point ($|V_1| = |V_2| = 1$) and two-receiver broadcast channels ($|V_1| = 1, |V_2| = 2$), though Theorems 2.1 and 2.2 apply using the general definition. The indegree and outdegree of node v are

$$d_{\text{in}}(v) \stackrel{\text{def}}{=} |\{[V_1, V_2] \in E | v \in V_2\}|$$

$$d_{\text{out}}(v) \stackrel{\text{def}}{=} |\{[V_1, V_2] \in E | v \in V_1\}|.$$

When $d_{\text{out}}(v)$ and $d_{\text{in}}(v)$ exceed 1, $\mathcal{X}^{(v)} = \prod_{d=1}^{d_{\text{out}}(v)} \mathcal{X}^{(v,d)}$, $\mathcal{Y}^{(v)} = \prod_{d=1}^{d_{\text{in}}(v)} \mathcal{Y}^{(v,d)}$, $X_t^{(v)} = (X_t^{(v,1)}, \dots, X_t^{(v,d_{\text{out}}(v))})$, and $Y_t^{(v)} = (Y_t^{(v,1)}, \dots, Y_t^{(v,d_{\text{in}}(v))})$, where the outputs and inputs of v are described in the (fixed but arbitrary) order imposed by the hyperedge indices. Let $V_1(e)$ and $V_2(e)$ denote the input and output ports of edge e . For example, if channel $e = [\{u\}, \{v_1, v_2\}]$ has input port (u, s) for some $s \in \{1, \dots, d_{\text{out}}(u)\}$ and output ports (v_1, r_1) and (v_2, r_2) for some $r_i \in \{1, \dots, d_{\text{out}}(v_i)\}$, $i \in \{1, 2\}$, then $V_1(e) = \{(u, s)\}$ and $V_2(e) = \{(v_1, r_1), (v_2, r_2)\}$. The channels described by the elements of E are independent, so

$$p(\mathbf{y}|\mathbf{x}) = \prod_{e \in E} p(y^{V_2(e)} | x^{V_1(e)}),$$

where for any set $S \subseteq S$, $X^S \stackrel{\text{def}}{=} (X^{(v)} : v \in S)$ and $Y^S \stackrel{\text{def}}{=} (Y^{(v)} : v \in S)$ are the vectors of channel inputs and outputs, respectively, corresponding to S and x^S and y^S are particular instances of those random variables. (When S is empty, X^S and Y^S are treated as fixed but arbitrary constants.)

A blocklength- n code operates the network over n time steps with the goal of communicating message $W^{(v \rightarrow v')} \in \mathcal{W}^{(v \rightarrow v')} \stackrel{\text{def}}{=} \{1, \dots, 2^{nR^{(v \rightarrow v')}}\}$ from node v to node v' . (Since nodes need not send messages to themselves, $R^{(v \rightarrow v)} \equiv 0$ for all $v \in V$.) The vector of rates $R^{(v \rightarrow v')}$ is denoted by \mathcal{R} . While the given notation suggests multiple unicast demands, the results apply to general demands using the argument from [17]. (While [17] treats networks of noiseless links, the argument applies immediately here as well.)

Using the given definitions, a network is written as a triple

$$\left(\prod_{v=1}^m \mathcal{X}^{(v)}, \prod_{e \in E} p(y^{V_2(e)} | x^{V_1(e)}), \prod_{v=1}^m \mathcal{Y}^{(v)} \right)$$

with the added constraint that $X_t^{(v)}$ is a function of $\{Y_1^{(v)}, \dots, Y_{t-1}^{(v)}, W^{(v \rightarrow 1)}, \dots, W^{(v \rightarrow m)}\}$ alone.

Remark 1: Including a “no transmission” symbol in the input and output alphabets of each channel allows arbitrary scheduling within this time-step model.

Definition 1: Let a network

$$\mathcal{N} \stackrel{\text{def}}{=} \left(\prod_{v=1}^m \mathcal{X}^{(v)}, \prod_{e \in E} p(y^{V_2(e)} | x^{V_1(e)}), \prod_{v=1}^m \mathcal{Y}^{(v)} \right)$$

be given corresponding to a graph $G = (V, E)$. A blocklength- n solution $\mathcal{S}(\mathcal{N})$ to this network is a set of encoding and

decoding functions:

$$X_t^{(v)} : (\mathcal{Y}^{(v)})^{t-1} \times \prod_{v''=1}^m \mathcal{W}^{(v \rightarrow v'')} \rightarrow \mathcal{X}^{(v)}$$

$$\hat{W}^{(v' \rightarrow v)} : (\mathcal{Y}^{(v)})^n \times \prod_{v''=1}^m \mathcal{W}^{(v \rightarrow v'')} \rightarrow \mathcal{W}^{(v' \rightarrow v)}$$

mapping $(Y_1^{(v)}, \dots, Y_{t-1}^{(v)}, W^{(v \rightarrow 1)}, \dots, W^{(v \rightarrow m)})$ to $X_t^{(v)}$ and $(Y_1^{(v)}, \dots, Y_n^{(v)}, W^{(v \rightarrow 1)}, \dots, W^{(v \rightarrow m)})$ to $\hat{W}^{(v' \rightarrow v)}$. Solution $\mathcal{S}(\mathcal{N})$ is called a (λ, \mathcal{R}) -solution, denoted $(\lambda, \mathcal{R}) - \mathcal{S}(\mathcal{N})$, if the encoding and decoding functions imply $\Pr(W^{(v \rightarrow v')} \neq \hat{W}^{(v \rightarrow v')}) < \lambda$ for all v, v' .

Definition 2: The rate region $\mathcal{R}(\mathcal{N}) \subset \mathbb{R}_+^{m^2}$ of a network \mathcal{N} is the closure of the set of rate vectors \mathcal{R} such that for any $\lambda > 0$ a $(\lambda, \mathcal{R}) - \mathcal{S}(\mathcal{N})$ solution exists for all blocklengths n sufficiently large.

Theorem 2.1 ([11, Theorem 3]): Consider networks

$$\mathcal{N} = \left(\mathcal{X}^{(1,1)} \times \dots \times \mathcal{X}^{(u,s)} \times \dots \times \mathcal{X}^{(m,d_{\text{out}}(m))}, \right.$$

$$p(y^{(v,r)} | x^{(u,s)}) \prod_{e \in E \setminus \{\bar{e}\}} p(y^{V_2(e)} | x^{V_1(e)}),$$

$$\mathcal{Y}^{(1,1)} \times \dots \times \mathcal{Y}^{(v,r)} \times \dots \times \mathcal{Y}^{(m,d_{\text{in}}(m))} \left. \right)$$

$$\hat{\mathcal{N}} = \left(\mathcal{X}^{(1,1)} \times \dots \times \hat{\mathcal{X}}^{(u,s)} \times \dots \times \mathcal{X}^{(m,d_{\text{out}}(m))}, \right.$$

$$\delta(\hat{x}^{(u,s)} - \hat{y}^{(v,r)}) \prod_{e \in E \setminus \{\bar{e}\}} p(y^{V_2(e)} | x^{V_1(e)}),$$

$$\mathcal{Y}^{(1,1)} \times \dots \times \hat{\mathcal{Y}}^{(v,r)} \times \dots \times \mathcal{Y}^{(m,d_{\text{in}}(m))} \left. \right),$$

where $(\hat{\mathcal{X}}^{(u,s)}, \delta(\hat{x}^{(u,s)} - \hat{y}^{(v,r)}), \hat{\mathcal{Y}}^{(v,r)})$ is a bit pipe that noiselessly maps $\max_{p(x^{u,s})} I(X^{(u,s)}; Y^{(v,r)})$ bits from its input to its output at each time step. Then

$$\mathcal{R}(\mathcal{N}) = \mathcal{R}(\hat{\mathcal{N}}).$$

Theorem 2.2 ([12, Theorem 4]): Given a network

$$\mathcal{N} = \left(\mathcal{X}^{(1,1)} \times \dots \times \mathcal{X}^{(u,s)} \times \dots \times \mathcal{X}^{(m,d_{\text{out}}(m))}, \right.$$

$$p(y^{(v_1,r_1)}, y^{(v_2,r_2)} | x^{(u,s)}) \prod_{e \in E \setminus \{\bar{e}\}} p(y^{V_2(e)} | x^{V_1(e)}),$$

$$\mathcal{Y}^{(1,1)} \times \dots \times \mathcal{Y}^{(v_1,r_1)} \times \dots \times \mathcal{Y}^{(v_2,r_2)} \times \dots$$

$$\left. \times \mathcal{Y}^{(m,d_{\text{in}}(m))} \right),$$

let $d_u = d_{\text{out}}(u) + 1$, $d_1 = d_{\text{out}}(v_1) + 1$, and $d_2 = d_{\text{in}}(v_2) + 1$. For each marginal $p = \{p(x^{(u,s)})\}_{x^{(u,s)}}$, define network $\hat{\mathcal{N}}(p)$

$$\hat{\mathcal{N}}(p) = \left(\mathcal{X}^{(1,1)} \times \dots \times \hat{\mathcal{X}}^{(u,s)} \times \dots \times \hat{\mathcal{X}}^{(u,d_u)} \times \dots \right.$$

$$\times \hat{\mathcal{X}}^{(v_1,d_1)} \times \dots \times \mathcal{X}^{(m,d_{\text{out}}(m))},$$

$$\delta(\hat{x}^{(u,s)} - \hat{y}^{(v_1,r_1)}) \delta(\hat{x}^{(u,d_u)} - \hat{y}^{(v_2,r_2)})$$

$$\cdot \delta(\hat{x}^{(v_1,d_1)} - \hat{y}^{(v_2,d_2)}) \prod_{e \in E \setminus \{\bar{e}\}} p(y^{V_2(e)} | x^{V_1(e)}),$$

$$\mathcal{Y}^{(1,1)} \times \dots \times \hat{\mathcal{Y}}^{(v_1,r_1)} \times \dots \times \hat{\mathcal{Y}}^{(v_2,r_2)} \times \dots$$

$$\left. \times \hat{\mathcal{Y}}^{(v_2,d_2)} \times \dots \times \mathcal{Y}^{(m,d_{\text{in}}(m))} \right),$$

where

$$(\hat{\mathcal{X}}^{(u,s)}, \delta(\hat{x}^{(u,s)} - \hat{y}^{(v_1,r_1)}), \hat{\mathcal{Y}}^{(v_1,r_1)})$$

$$(\hat{\mathcal{X}}^{(v_1,d_1)}, \delta(\hat{x}^{(v_1,d_1)} - \hat{y}^{(v_2,d_2)}), \hat{\mathcal{Y}}^{(v_2,d_2)})$$

$$(\hat{\mathcal{X}}^{(u,d_u)}, \delta(\hat{x}^{(u,d_u)} - \hat{y}^{(v_2,r_2)}), \hat{\mathcal{Y}}^{(v_2,r_2)})$$




Fig. 1. Two deterministic broadcast channel models from [12].




Fig. 2. An alternative formulation of the models from Fig. 1.

are noiseless bit pipes of capacities $I(X^{(u,s)}; Y^{(v1,r1)})$, $I(X^{(u,s)}; Y^{(v1,r1)})$, and $I(X^{(u,s)}; Y^{(v2,r2)} | Y^{(v1,r1)})$ for distribution $p(x^{(u,s)})p(y^{(v1,r1)}, y^{(v2,r2)} | x^{(u,s)})$. Then

$$\mathcal{R}(\mathcal{N}) \subseteq \bigcup_p \mathcal{R}(\hat{\mathcal{N}}(p)).$$

Theorem 2.1 shows that the capacity of any network \mathcal{N} containing noisy channel $(\mathcal{X}^{(u,s)}, p(y^{(v,r)} | x^{(u,s)}), \mathcal{Y}^{(v,r)})$ from node u to node v equals the capacity of a distinct network $\hat{\mathcal{N}}$ that is identical to \mathcal{N} except that the channel described above is replaced by a rate- $(\max_{p(x)} I(X^{(u,s)}; Y^{(v,r)}))$ noiseless bit pipe $(\hat{\mathcal{X}}^{(u,s)}, \delta(x^{(u,s)} - y^{(v,r)}), \hat{\mathcal{Y}}^{(v,r)})$. Theorem 2.2 bounds the capacity of any network \mathcal{N} containing a noisy broadcast channel $(\mathcal{X}^{(u,s)}, p(y^{(v1,r1)}, y^{(v2,r2)} | x^{(u,s)}), \mathcal{Y}^{(v1,r1)} \times \mathcal{Y}^{(v2,r2)})$ from node u to nodes v_1 and v_2 by the capacity of a distinct network $\hat{\mathcal{N}}$ that is identical to \mathcal{N} except that the given broadcast channel is replaced by three noiseless bit pipes. The given model is asymmetrical in broadcast receivers v_1 and v_2 . Since the labeling is arbitrary, this result effectively offers two models for each broadcast channel – both of which guarantee outer bounds on the capacity region of the original network. An illustration of both models appears in Figure 1. The same results apply immediately to the models in Figure 2.

III. RESULTS

Repeated application of Theorems 2.1 and 2.2 allows us to bound the capacity of any network of memoryless point-to-point and broadcast channels by the capacity of a network of noiseless, point-to-point bit pipes. Since the broadcast channel model is asymmetrical in receivers 1 and 2 and the index labeling is arbitrary, for each network containing k broadcast channels, the given results provide 2^k possible network models – each of which gives a (possibly distinct) outer bound on the network capacity region.

Theorem 3.1 derives general bounds for noisy broadcast networks. The results illustrate the power of Theorems 2.1 and 2.2 by showing the simplicity of deriving new results for networks of noisy wireless and wireline channels using prior results from the traditional network coding literature.

For any sets $A, T \subset V$ with $A \cap T = \emptyset$,

$$\mathcal{S}_{A,T} \stackrel{\text{def}}{=} \{S \subset V : A \subseteq S \wedge T \setminus S \neq \emptyset\}.$$

Let $\mathcal{S}_{a,T} = \mathcal{S}_{A,T}$ when $A = \{a\}$.

Theorem 3.1: Consider a network

$$\mathcal{N} \stackrel{\text{def}}{=} \left(\prod_{v=1}^m \mathcal{X}^{(v)}, \prod_{e \in E} p(y^{V_2(e)} | x^{V_1(e)}), \prod_{v=1}^m \mathcal{Y}^{(v)} \right)$$

defined on graph $G = (V, E)$ with $|V_1(e)| = 1$ and $|V_2(e)| \leq 2$ for all $e \in E$.

1) The capacity C for the multicast problem with source $s \in V$ and terminals $T \subset V \setminus \{s\}$ is bounded as

$$C \leq \max_{\{p(x^{V_1(e)}): e \in E\}} \min_{S \in \mathcal{S}_{s,T}} \sum_{e: V_1(e) \subseteq S} I(X^{V_1(e)}; Y^{V_2(e) \setminus S})$$

2) Given a multi-source multicast problem with source nodes $\{s_1, \dots, s_k\}$ and demand nodes $T \subset V \setminus \{s\}$ such that each node $v \in T$ requires rate R_i from source node s_i for all $i \in \{1, \dots, k\}$, then

$$\sum_{i \in A} R_i \leq \max_{\{p(x^{V_1(e)}): e \in E\}} \min_{S \in \mathcal{S}_{s,T}} \sum_{e: V_1(e) \subseteq S} I(X^{V_1(e)}; Y^{V_2(e) \setminus S})$$

$$\forall A \subseteq \{1, \dots, k\}.$$

3) Let $s \in V$ be a single source node, and let $T = \{t_1, \dots, t_k\} \subset V \setminus \{s\}$ be a collection of terminal nodes; node t_i wishes to collect message $W^{(s,i)}$ from s at rate R_i . Messages $\{W^{(s,i)} : i \in \{1, \dots, k\}\}$ are independent. Then

$$\sum_{i \in A} R_i \leq \max_{\{p(x^{V_1(e)}): e \in E\}} \min_{\substack{S \subset V \setminus A: \\ s \in S}} \sum_{e: V_1(e) \subseteq S} I(X^{V_1(e)}; Y^{V_2(e) \setminus S})$$

$$\forall A \subseteq \{1, \dots, k\}.$$

4) Let $s \in V$ be a single source node, and let $T_1 = \{t_1, \dots, t_k\} \subset V \setminus \{s\}$ be a collection of terminal nodes; t_i wishes to collect message $W^{(s,i)}$ from s at rate R_i . Let $T_2 \subseteq V \setminus (\{s\} \cup T_1)$ be a distinct set of terminal nodes for which each node $t \in T_2$ demands all k messages. Messages $\{W^{(s,i)} : i \in \{1, \dots, k\}\}$ are independent. Then

$$\sum_{i=1}^k R_i \leq \max_{\{p(x^{V_1(e)}): e \in E\}} \min_{S \in \mathcal{S}_{s,T_2}} \sum_{e: V_1(e) \subseteq S} I(X^{V_1(e)}; Y^{V_2(e) \setminus S})$$

$$\sum_{i \in A} R_i \leq \max_{\{p(x^{V_1(e)}): e \in E\}} \min_{\substack{S \subset V \setminus A: \\ s \in S}} \sum_{e: V_1(e) \subseteq S} I(X^{V_1(e)}; Y^{V_2(e) \setminus S})$$

$$\forall A \subseteq \{1, \dots, k\}.$$

Proof. 1) Let $\hat{\mathcal{N}}_{p,\ell}$ denote the network model associated with a given choice of input distributions $p = \{p(x^{V_1(e)}): e \in E\}$ and model choices $\ell = \{\ell_e : e \in E\}$, where $\ell_e : V_2(e) \rightarrow \{1, \dots, |V_2(e)|\}$ specifies which receiver of each broadcast channel serves as receiver 1. (The model choice $\ell(e)$ for a

point-to-point link e is unique since in that case $|V_2(e)| = 1$ and the channel model is tight.) If the optimal multicast coding strategy on \mathcal{N} , and therefore its optimal collection of input distributions $p^* = \{p(x^{V_1(e)}) : e \in E\}$, were known, then the proofs of Theorems 2.1 and 2.2 show that the multicast capacity of \mathcal{N} is no greater than the multicast capacity of every network $\mathcal{N}_{p^*,\ell}$. For any fixed $\mathcal{N}_{p,\ell}$, the multicast capacity of $\mathcal{N}_{p,\ell}$ equals the network's min-cut by [2]. Thus

$$C \leq \min_{\ell} \min_{p^*} \text{mincut}(\mathcal{N}_{p^*,\ell}) \leq \max_p \min_{\ell} \text{mincut}(\mathcal{N}_{p,\ell}).$$

For any cut S , hyperedge e , and model choice $\ell(e)$ for e , let $\mathcal{V}(S, e, \ell(e))$ denote the total capacity of all edges of the model for channel e corresponding to model choice $\ell(e)$ passing from nodes in S to nodes in S^c . Then

$$\begin{aligned} \min_{\ell} \text{mincut}(\mathcal{N}_{p,\ell}) &= \min_{\ell} \min_{S \in \mathcal{S}_{s,T}} \sum_{e \in E} \mathcal{V}(S, e, \ell(e)) \\ &= \min_{S \in \mathcal{S}_{s,T}} \sum_{e \in E} \min_{\ell(e)} \mathcal{V}(S, e, \ell(e)) \end{aligned}$$

The following analysis finds $\min_{\ell(e)} \mathcal{V}(S, e, \ell(e))$ for each S , e , and $\ell(e)$. For any cut S and hyperedge e such that $V_1(e) \not\subseteq S$, there exists a model choice $\ell(e)$ for which $\mathcal{V}(S, e, \ell(e)) = 0$. For any cut S and hyperedge e for which $V_1(e) \subseteq S$ and $|V_2(e) \setminus S| = 1$, $\mathcal{V}(S, e, \ell(e))$ equals $I(X^{V_1(e)}; Y^{V_2(e) \setminus S})$ for one model of e and $I(X^{V_1(e)}; Y^{V_2(e)})$ for the other; thus $\min_{\ell(e)} \mathcal{V}(S, e, \ell(e)) = I(X^{V_1(e)}; Y^{V_2(e) \setminus S})$. Finally, for any cut S and hyperedge e for which $V_1(e) \subseteq S$ and $|V_2(e) \setminus S| = 2$, $\mathcal{V}(S, e, \ell(e))$ equals $I(X^{V_1(e)}; Y^{V_2(e)}) = I(X^{V_1(e)}; Y^{V_2(e) \setminus S})$ for both choices of $\ell(e)$. Thus

$$\min_{\ell(e)} \mathcal{V}(S, e, \ell(e)) = \begin{cases} I(X^{V_1(e)}; Y^{V_2(e) \setminus S}) & \text{if } V_1(e) \subseteq S \\ 0 & \text{otherwise,} \end{cases}$$

which gives the desired result.

The proofs for 2, 3, and 4 apply [18, Theorems 8, 9, and 10] in arguments similar to those used for 1. ■

Theorem 3.1 applies network coding results for demand types where cut-set bounds are tight in networks of noiseless, point-to-point links. This does not, however, imply tightness for the noisy networks. Theorem 3.1 is tight for multicast across deterministic broadcast channels by [9]. It also matches the multicast capacity bounds of [15], [16]; these bounds may or may not be tight since existing achievability results employ side information about the erasure locations. We can bound the tightness of Theorem 3.1 in general by comparing the given upper bounds with the corresponding capacities of the network of lower bounding models from [12]. For example, when \mathcal{N} is a network of memoryless Gaussian channels with independent noise at the receivers, Theorem 3.1 bounds the rate region of \mathcal{N} using a family of bounds of the form $\sum_i R_i \leq \max_p \min_S \sum_e I(X^{V_1(e)}; Y^{V_2(e) \setminus S})$. For each such upper bound, the following lower bound also applies

$$\sum_i R_i \geq \max_p \min_S \sum_e [I(X^{V_1(e)}; Y^{V_2(e) \setminus S}) - E(p, S, e)]$$

by [19, Theorem 4.2], where $E(p, S, e) = 0$ if $V_2(e) \setminus S = \emptyset$ or $|V_2(e)| = 1$ and $1/2$ otherwise. (The same theorem gives




Fig. 3. The multiple unicast network evaluated in Lemma 3.2.

corresponding bounds for networks containing multiple access channels and broadcast channels with dependent noise at their receivers.) Thus for Gaussian channels with independent noise, each bound in Theorem 3.1 differs from the true capacity by at most $k_o/2$, where $k_o \leq |E|$ is the number of broadcast channels contributing to each worst-case cut.

Note that application of cut-set bounds directly to the noisy network generally gives looser bounds than those found in Theorem 3.1. Given a cut S , the looseness in the cut-set bound for S on \mathcal{N} arises in part from the potential dependence between the inputs to distinct channels traversing from S to S^c . This dependence may result from statistical dependence in the noise between distinct receivers of a broadcast channel or from shared information received from other nodes.

A variety of techniques are available for bounding capacities of networks of noiseless links when cut-set bounds are not tight (e.g., [3], [5]). Since the bounds are complex, Lemma 3.2 bounds capacity for a specific multiple unicast problem rather than giving a general form bound. The given example, which modifies an example from [3] from a network of noiseless, point-to-point links to a network of broadcast channels with dependent receiver noise, shows that the difference between the resulting bounds and the traditional cut-set bounds can be arbitrarily large. We here employ the noiseless bit-pipe models of Theorems 2.1 and 2.2 and then bound the capacities of the underlying network using an argument like that in [3].¹ For the sake of concreteness, we pose the problem as a network of Gaussian broadcast channels. We note, however, that the technique applies equally well to other symmetric broadcast channels and does not rely on the specific channel model.

Lemma 3.2: Consider the network of noisy broadcast channels with $V = \{1, \dots, 3n+1\}$ and $E = \{e_1, \dots, e_{2n}\}$ shown in Figure 3. Each hyperedge $e \in E$ is a two-receiver Gaussian broadcast channel with power constraint P . The noise observed at the two receivers has variance $N_1 = N_2 = N$ and correlation coefficient $\rho \neq 0$. If each node $u \in \{1, \dots, n\}$ transmits an independent unicast to its terminal $v = 2n + u$ at

¹A bounding tool similar to that of [3] that applies to some networks with broadcast transmissions, interference receptions, and independent vertex noise appears in [10]. Lemma 3.2 violates that paper's independent vertex noise assumption. Nonetheless, in his (explicitly self-identifying) review of this paper, the first author of [5] claims that the method of [5] can also be used to derive the upper bound in Lemma 3.2. The author has not verified that claim.




Fig. 4. A model for the network in Figure 3.

rate R , then the optimal achievable rate satisfies

$$R = C/n,$$

where $C = (1/2) \log(1+P/N)$. Applying the cut-set approach of [1] on network \mathcal{N} gives minimal outer bound $R \leq C/2$.

Proof. Consider the model shown in Figure 4. Channel e_n is replaced by a model from Figure 2 while all other channels are replaced by models from Figure 1. The model for channel e_n introduces an extra node to the network, here labeled $3n+2$.

Fix the network size parameter n . For any achievable rate R , there exists a multiple unicast code for the given network that reconstructs each message $W^{(i,2n+i)}$ with error probability that can be made arbitrarily small for code blocklength sufficiently large. Fix such a code. Let U_e specify the information traversing edge e using this code, and let $U \stackrel{\text{def}}{=} U_{[3n+2, n+1]}$. We first show that for every edge e in the bounding network, we can reconstruct U_e from U with probability approaching 1.

We visit nodes $n+i$, $2n+i$, and $3n+i$ for i increasing from 1 to n , sequentially showing that the inputs – and therefore the outputs – of each node can be reconstructed from U with probability approaching 1. The only input to node $n+1$ is U , thus $U_{[n+1, n+2]}$ and $U_{[n+1, 2n+1]}$ are deterministic functions of U . Node $2n+1$ has only one input, thus $\hat{W}^{(1, 2n+1)}$ and $U_{[2n+1, n+2]}$ are deterministic functions of $U_{[n+1, 2n+1]}$ and therefore of U . At node 1, $U_{[1, 2]}$ and $U_{[1, n+2]}$ are deterministic functions of $W^{(1, 2n+1)}$, which equals $\hat{W}^{(1, 2n+1)}$ with probability approaching 1; thus we can reconstruct $U_{[1, 2]}$ and $U_{[1, n+2]}$ from U with probability approaching 1. Continuing this argument for nodes $n+i$, $2n+i$, i , for each subsequent i shows that all n messages $W^{(i, 2n+i)}$ can be reconstructed with probability approaching 1 using only U . Thus nR is bounded by the capacity of edge $[3n+3, n+1]$, which is at most $C = (1/2) \log(1+P/N)$ by Theorem 2.2. Operating broadcast channels e_1, \dots, e_n at rate point $(C, 0)$ and e_{n+1}, \dots, e_{2n} at rate point $(\frac{n-1}{n}C, \frac{1}{n}C)$ (which is achievable by a time-sharing argument) and routing over the resulting asymptotically lossless network achieves the desired bound.

For the cut-set bound, fix any cut S and $k \geq 2$. Let $1 \leq v_1 < v_2 < \dots < v_k \leq n$ be the sources separated from their sinks by S (i.e., $(v_i, 2n+v_i) \in S \times S^c$). For each $i < k$, the path $(v_i, n+v_i+1, n+v_i+2, \dots, n+v_{i+1}, 2n+v_{i+1})$ must cross at least once from S to S^c . Since these paths are disjoint, the cut across S bounds kR by a value no less than $(k-1)C$. The resulting bound is minimized when $k=2$. ■

IV. CONCLUSION

This work derives a collection of network capacity bounds using point-to-point noiseless models derived in prior work. The results demonstrate the usefulness of equivalence tools, which facilitate bound derivation for noisy networks. The results are tight for some examples. Methods for quantifying bound accuracy are available for others. Finding the precise error in general may be difficult since even the capacity of a single broadcast channel remains unsolved. When the bounds are not known to be tight, bounds on their accuracy can be derived.

ACKNOWLEDGMENT

This work was supported by DARPA ITMANET under grant number W911NF-07-1-0029.

REFERENCES

- [1] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley, second ed., 2006.
- [2] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. on Information Theory*, vol. IT-46, pp. 1204–1216, July 2000.
- [3] N. Harvey and R. Kleinberg, "Tighter cut-set bounds for k -pairs communication problems," in *Proc. of the Allerton Conference on Communication, Control, and Computing*, (Monticello, IL), Sept. 2005.
- [4] N. Harvey, R. Kleinberg, and A. R. Lehman, "On the capacity of information networks," *IEEE Trans. on Information Theory*, vol. 52, pp. 2345–2364, June 2006.
- [5] G. Kramer and S. Savari, "Edge-cut bounds on network coding rates," *Journal of Network and Systems Mngmnt*, vol. 14, pp. 49–67, Mar. 2006.
- [6] L. Song, R. W. Yeung, and N. Cai, "Zero-error network coding for acyclic networks," *IEEE Trans. on Information Theory*, vol. 49, pp. 3129–3139, July 2003.
- [7] S. Borade, "Network information flow: limits and achievability," in *Proc. of the IEEE Int. Symp. on Information Theory*, p. 139, July 2002.
- [8] L. Song, R. W. Yeung, and N. Cai, "A separation theorem for single-source network coding," *IEEE Trans. on Information Theory*, vol. 52, pp. 1861–1871, May 2006.
- [9] N. Ratnakar and G. Kramer, "The multicast capacity of deterministic relay networks with no interference," *IEEE Trans. on Information Theory*, vol. 52, pp. 2425–2432, June 2006.
- [10] G. Kramer and S. Savari, "Capacity bounds for relay networks," in *Information Theory and App.s Workshop*, (San Diego) IEEE, Jan. 2006.
- [11] R. Koetter, M. Effros, and M. Médard, "On a theory of network equivalence," in *Proc. of the IEEE Information Theory Workshop*, (Volos, Greece), pp. 326–330, June 2009.
- [12] R. Koetter, M. Effros, and M. Médard, "Beyond network equivalence," in *Allerton Annual Conference on Communications, Control, and Computing*, (Monticello, IL), Sept. 2009.
- [13] R. Koetter, M. Effros, and M. Médard, "On a theory of network equivalence," *IEEE Trans. on Information Theory*, 2010.
- [14] R. Gowaikar, A. F. Dana, B. Hassibi, and M. Effros, "A practical scheme for wireless network operation," *IEEE Trans. on Communications*, vol. 55, pp. 463–476, Mar. 2007.
- [15] A. Dana, R. Gowaikar, R. Palanki, B. Hassibi, and M. Effros, "Capacity of wireless erasure networks," *IEEE Trans. on Information Theory*, vol. 52, pp. 789–804, Mar. 2006.
- [16] D. S. Lun, M. Médard, R. Koetter, and M. Effros, "On coding for reliable communication over packet networks," *Physical Communication*, vol. 1, pp. 3–20, Mar. 2008.
- [17] R. Dougherty and K. Zeger, "Nonreversibility and equivalent constructions of multiple-unicast networks," *IEEE Trans. on Information Theory*, vol. 52, no. 11, pp. 5067–5077, 2006.
- [18] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Trans. on Networking*, vol. 11, pp. 782–795, Oct. 2003.
- [19] M. Effros, "On capacity outer bounds for a simple family of wireless networks," in *Information Theory and App.s Workshop*, (San Diego), IEEE, Jan. 2010. Invited paper.