Tighter Cut-based Bounds for k-pairs Communication Problems

Nicholas J. A. Harvey* Robert Kleinberg†

Abstract

We study the extent to which combinatorial cut conditions determine the maximum network coding rate of k-pairs communication problems. We seek a combinatorial parameter of directed networks which constitutes a valid upper bound on the network coding rate but exceeds this rate by only a small factor in the worst case. (This worst-case ratio is called the ***gap*** of the parameter.) We begin by considering vertex-sparsity and meagerness, showing that both of these parameters have a gap which is linear in the network size. Due to the weakness of these bounds, we propose a new bound called ***informational meagerness***. This bound generalizes both vertex-sparsity and meagerness and is potentially the first known combinatorial cut condition with a sublinear gap. However, we prove that informational meagerness does not tightly characterize the network coding rate: its gap can be super-constant.

1 Introduction

A k-pairs communication problem is a special type of network coding problem in which each message has a single source node and a single sink node. Such problems warrant study for both practical and theoretical reasons. On the practical side, k-pairs communication problems model the vast majority of contemporary communication sessions, which typically involve a single sender and single receiver. On the theoretical side, characterizing the capacity of k-pairs communication problems has implications for long-standing open problems in computational complexity [1]. The term "multiple unicast sessions" has also been used to refer to k-pairs communication problems [7, 11].

A key goal of the existing work on k-pairs communication problems is to compare the network coding rate to the corresponding multicommodity flow rate, i.e., the maximum transmittion rate achievable when data is treated as a fluid that cannot be copied or coded. On the lower-bound side, there are examples [5, 6, 11] showing that the use of coding can increase the rate by an unbounded factor in directed graphs (as the size of the graph tends to infinity). In contrast, no undirected graph is known where the network coding rate exceeds the multicommodity flow rate. The ***undirected k-pairs conjecture*** claims that actually these two rates are always equal [1, 5, 6, 11].

On the upper-bound side, various approaches have been proposed to bound the network coding rate. A general outer bound on the capacity of multi-commodity information networks can be found in standard textbooks [4]. This bound was used by Borade [3]

*MIT Computer Science and Artificial Intelligence Laboratory. nickh@mit.edu. Supported by a Natural Sciences and Engineering Research Council of Canada Fellowship.

†MIT Department of Mathematics. rdk@math.mit.edu. Supported by a Fannie and John Hertz Foundation Fellowship.

to obtain a bound on the capacity of directed information networks. Building on this result, Kramer and Savari [10] give an general outer bound, called the bidirected cut-set bound, on the capacity of (directed or undirected) information networks. Their work shows how bounds based on cuts can be applied to information networks, even when undirected channels have complex rate regions. Kramer and Savari [9] also proposed a stronger outer bound based on a technique from Bayesian networks called d-separation. Song et al. [12] define a linear program whose optimum value bounds the network coding rate in directed acyclic graphs. Harvey et al. [1, 5] define a similar linear program which bounds the network coding rate in undirected and cyclic graphs.

This paper focuses on cut-based techniques for proving upper bounds on the rate of directed network coding problems. The skeptical reader might wonder why we restrict our attention to cut-based approaches, rather than more elaborate bounds (e.g., LP-based bounds) which are potentially stronger. One reason is that cut-based methods are a very natural approach for inferring properties of information flows. Another reason is that cut-based bounds are computationally attractive, since they can be efficiently verified: the cut itself acts as a succinct certificate of the upper bound.

This paper considers three cut-based bounds, which we refer to as ***edge-sparsity***, ***vertex-sparsity***, and ***meagerness***. Sparsity is a classical notion from the multicommodity flow literature which bounds the rate of a concurrent[1] multicommodity flow solution [13]. Sparsity can be defined using either sets of edges or sets of vertices, hence our distinction between edge-sparsity and vertex-sparsity. Meagerness is a notion introduced by Harvey et al. [6] to address some shortcomings of sparsity in directed information networks. These bounds are defined more formally below.

1.1 Definitions

We now introduce our formal definitions and notation. An instance of the k-pairs communication problem consists of: (1) A directed graph $G = (V, E)$. (2) A capacity $c(e) \in \mathbb{R}^+$ for each edge e. (3) A set $\mathcal{I}$ of "commodities" of size k. (4) For each $i \in \mathcal{I}$, a source vertex $\sigma(i) \in V$, a sink vertex $\tau(i) \in V$, and a demanded communication rate $d_i \in \mathbb{R}^+$. It is convenient to discuss the flow of information by considering only sets of edges, so we will assume without loss of generality that each source has a single out-edge $S(i) = (\sigma(i), s(i))$ and no in-edges, and that each sink has a single in-edge $T(i) = (t(i), \tau(i))$ and no out-edges. For $I \subseteq \mathcal{I}$, define $\mathcal{S}(I) = \{ S(i) : i \in I \}$. For convenience, define $\mathcal{S} = \mathcal{S}(\mathcal{I})$. The sets $\mathcal{T}(I)$ and $\mathcal{T}$ are defined analogously.

A ***network code*** specifies a function for each $e \in E$ which determines the message sent on e as a function of the messages received at e's tail vertex. We assume that the sources are mutually independent random processes, where $S(i)$ transmits a uniformly distributed message in some alphabet Σ_i. A ***network coding solution*** is a network code in which $H(S(i)) = H(T(i))$ and $H(T(i)|S(i)) = 0$ for each $i \in \mathcal{I}$. A network coding solution has ***rate*** r if there exists a constant α such that $H(e) \leq \alpha \cdot c(e)\ \forall e \in E$ and $H(S(i)) \geq \alpha \cdot r \cdot d_i\ \forall i \in \mathcal{I}$. The ***network coding rate*** of an instance G is the supremum of the rates of all network coding solutions.

We now present some tools for bounding the network coding rate of an instance. A

[1]Here, concurrent means that the demanded rate of all sources is scaled by the same factor. Hence, the feasible rate region is a subset of $\mathbb{R}$, not of $\mathbb{R}^k$.

set $B \subseteq E$ is said to be ***downstream*** of $A \subseteq E$ if there is no path from $\mathcal{S}$ to B in $G \setminus A$. This relation is denoted $A \leadsto B$. A basic property is $A \leadsto B \implies H(A) = H(A, B)$, where $H(A)$ denotes the joint entropy of the messages transmitted on the edges in A. A strictly stronger notion than downstreamness is ***informational dominance*** [5]. We say edge set A informationally dominates edge set B if the information transmitted on edges in A determines the information transmitted on edges in B in *all* network coding solutions with positive rate. This relation is denoted $A \overset{\mathrm{i}}{\leadsto} B$, and the set of all edges dominated by A is denoted $\mathrm{Dom}(A)$. We will discuss informational dominance further in Section 3.

We now define and elaborate on the cut-based bounds mentioned above.

- *Edge-sparsity.* Given $A \subseteq E$, we say that A ***separates*** commodity i if there is no path from $S(i)$ to $T(i)$ in $G \setminus A$. Let $C^e(A) = \sum_{e \in A} c(A)$ and $D^e(A)$ denote the total demand of commodities separated by A. The sparsity of edge set A is $\mathscr{S}^e(A) := C^e(A)/D^e(A)$. The edge-sparsity of G is $\mathscr{S}^e_G := \min_{A \subseteq E} \mathscr{S}^e(A)$. The edge-sparsity is an upper bound for the maximum flow rate in both undirected and directed graphs and for the network coding rate in undirected graphs. However, the well-known butterfly graph [2] shows that the edge-sparsity is *not* an upper bound for the network coding rate in directed graphs.
- *Vertex-sparsity.* Given $U \subset V$, we say that U ***separates*** commodity i if $\sigma(i) \in U$ and $\tau(i) \in \overline{U}$ or vice-versa. Let $C^v(U)$ denote the total capacity of all edges with one endpoint in U and the other in $\overline{U}$. Let $D^v(U)$ denote the total demand of commodities separated by U. The sparsity of U is $\mathscr{S}^v(U) := C^v(U)/D^v(U)$. The vertex-sparsity of G is $\mathscr{S}^v_G := \min_{U \subset V} \mathscr{S}^v(U)$. This quantity is an upper bound for the maximum flow rate and the network coding rate in both undirected and directed graphs (see, e.g., [3, 5, 7, 10]). However, edge-sparsity gives a stronger bound on the flow rate.
- *Meagerness.* Given an edge set $A \subseteq E$ and a commodity set $P \subseteq \mathcal{I}$, we say that A ***isolates*** P if for all pairs $i, j \in P$, every path from $S(i)$ to $T(j)$ intersects A. The meagerness of A and G are respectively

  $$\mathcal{M}(A) := \min_{P\,:\,A \text{ isolates } P} \frac{C(A)}{\sum_{i \in P} d_i} \qquad \text{and} \qquad \mathcal{M}_G := \min_{A \subseteq E} \mathcal{M}(A).$$

  This quantity is an upper bound for the network coding rate in directed and undirected graphs [5, 6]. However, for undirected graphs, edge-sparsity gives a stronger bound.

For any bound mentioned above, we define its ***gap*** to be the ratio between the value of the bound and the network coding rate of the instance.

1.2 Our Contributions

We analyze the quality of these cut-based bounds and show that both vertex-sparsity and meagerness can have a gap that is linear in the size of the instance. Since these bounds can be weak, we propose a new combinatorial cut condition called ***informational meagerness***. This bound generalizes both vertex-sparsity and meagerness and is potentially the first known combinatorial cut condition with a sublinear gap. However, we prove that informational meagerness does not tightly characterize the network coding rate: its gap can be logarithmic in the size of the instance.

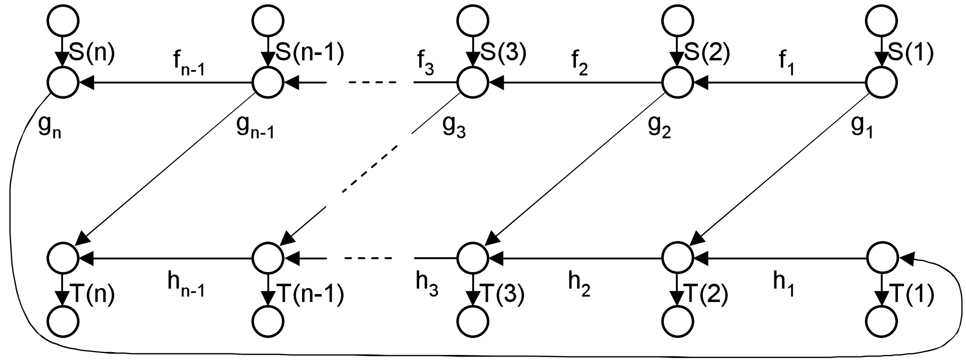

Figure 1: The graph $\mathcal{G}_n$ has meagerness gap and vertex-sparsity gap at least $\Omega(n)$.

2 Vertex-sparsity and Meagerness

As mentioned above, edge-sparsity is not an upper bound for the network coding rate in directed k-pairs instances. On the other hand, vertex-sparsity and meagerness are valid upper bounds, but these bounds are not always tight. It was shown in [5, 6], via an example called the split butterfly, that the meagerness bound can have a gap of at least 3/2. Formally, we exhibit a family of graphs $\mathcal{G}_n$ for which the vertex-sparsity gap and meagerness gap are $\Omega(n)$.

2.1 A Linear Gap

For $n \geq 2$, we define a directed graph $\mathcal{G}_n$ with n commodities and $4n$ vertices. The edges of $\mathcal{G}_n$ are $f_i := (s(i), s(i+1))$, $g_i := (s(i), t(i+1))$, and $h_i := (t(i), t(i+1))$ (for $1 \leq i < n$) and $g_n := (s(n), t(1))$. All edges have unit capacity and all commodities have unit demand. The graph $\mathcal{G}_n$ is illustrated in Figure 1.

Lemma 1. The network coding rate of $\mathcal{G}_n$ is $1/n$.

Proof. Consider any network coding solution in $\mathcal{G}_n$. We have the following chain of downstreamness and informational dominance relations:

$$
\begin{aligned}
\{g_n\} &\rightsquigarrow \{g_n, T(1), h_1\} && (g_n \text{ is the only inbound edge to } \{T(1), h_1\}.)\\
&\overset{\text{i}}{\rightsquigarrow} \{S(1), f_1, g_1, h_1\} && (S(1) \text{ and } T(1) \text{ are equal RVs and } \{S(1)\} \rightsquigarrow \{f_1, g_1\}.)\\
&\rightsquigarrow \{S(1), f_1, T(2), h_2\} && (\{g_1, h_1\} \text{ are the inbound edges to } \{T(2), h_2\}.)\\
&\overset{\text{i}}{\rightsquigarrow} \{S(1), S(2), f_2, g_2, h_2\} && (S(2) \text{ and } T(2) \text{ are equal RVs and } \{S(2), f_1\} \rightsquigarrow \{f_2, g_2\}.)\\
&\rightsquigarrow \{S(1), S(2), f_2, T(3), h_3\} && (\{g_2, h_2\} \text{ are the inbound edges to } \{T(3), h_3\}.)
\end{aligned}
$$

Using an inductive argument along these lines, we obtain the informational dominance relation $\{g_n\} \overset{\text{i}}{\rightsquigarrow} \{S(1), \ldots, S(n)\}$. Since g_n has capacity 1, we obtain that $\alpha \geq H(g_n) \geq H(S(1), \ldots, S(n)) \geq \alpha \cdot n \cdot r$, where r is the rate of our network coding solution and α is the unspecified constant appearing in the definition of "rate." Thus $r \leq 1/n$. Moreover, this rate is achievable by a trivial multiplexing solution. □

Lemma 2. $\mathcal{M}(\mathcal{G}_n) = 1$.

Proof. We consider a set of commodities P and show that any set which isolates P

must be large. If P contains a single commodity then we must remove at least one edge to isolate this commodity, so we obtain meagerness of at least 1 for this case. So suppose that $P = \{i_1, i_2, \ldots, i_k\}$ for $k \geq 2$ where $i_1 < \cdots < i_k$. We claim that we must remove at least k edges to isolate P. To prove this, we will exhibit k edge-disjoint paths connecting the sources and (non-corresponding) sinks for commodities in P. For $1 \leq a < k$, we have the path $\sigma_{i_a}, s_{i_a}, t_{i_a+1}, t_{i_a+2}, \ldots, t_{i_{a+1}}, \tau_{i_{a+1}}$, and the final path is $\sigma_{i_k}, s_{i_k}, s_{i_k+1}, s_{i_k+2}, \ldots, t_1, t_2, \ldots, t_{i_1}, \tau_{i_1}$. $\square$

Lemma 3. $\mathscr{S}^v_{\mathcal{G}_n} = \Omega(1)$.

Proof. Without loss of generality, we may contract all source and sink edges. Next, let $U \subset V(\mathcal{G}_n)$ be arbitrary. We will show that any commodity $i > 1$ separated by U can be charged uniquely to an edge that is cut by U. If $\{s(i), s(i-1)\} \subseteq U$ and $t(i) \in \overline{U}$, or vice versa, then we charge to the edge g_{i-1}, which is cut by U. Otherwise, if $\{s(i-1), t(i)\} \subseteq U$ and $s(i) \in \overline{U}$, or vice versa, then we charge to the edge f_{i-1}, which is cut by U. Hence $D^v(U) \leq C^v(U) + 1$, since we have neglected commodity $i = 1$. $\square$

Lemmas 1-3 show that $\mathcal{G}_n$ has meagerness gap and vertex-sparsity gap at least $\Omega(n)$.

For completeness, we mention another example which can be used to show the same result.[2] Consider a directed n-cycle where each node is a source and the corresponding sink is at distance $n-1$ from the source. All edges have unit capacity and all commodities have unit demand. A moment's thought shows that the maximum flow rate and the edge-sparsity are both $1/(n-1)$, but the vertex-sparsity and meagerness are both equal to 1. In fact, the network coding rate in this instance is also $1/(n-1)$ since it has been proven that the network coding rate equals the maximum flow rate in any instance of the k-pairs communication problem on a directed cycle [5]. (As an exercise, the reader may wish to use our techniques to prove this claim.) Therefore this example also has meagerness gap and vertex-sparsity gap at least $\Omega(n)$.

3 Informational Meagerness

Before defining informational meagerness, we first define informational isolation. First observe that the definition of isolation from Section 1 has the following equivalent statement: a set $A \subseteq E$ isolates P if and only if $A \cup \mathcal{S}(\overline{P}) \rightsquigarrow \mathcal{T}(P)$. By analogy, we say that A ***informationally isolates*** P if $A \cup \mathcal{S}(\overline{P}) \overset{\mathrm{i}}{\rightsquigarrow} \mathcal{T}(P)$. Note that the set of commodities isolated by A is a subset of the commodities informationally isolated by A. The informational meagerness (or, more whimsically, ***iMeagerness***) of A and G are:

$$i\mathcal{M}(A) := \min_{P\,:\,A \text{ informationally isolates } P} \frac{C(A)}{\sum_{i \in P} d_i} \qquad \text{and} \qquad i\mathcal{M}_G := \min_{A \subseteq E} i\mathcal{M}(A).$$

Since $A \overset{\mathrm{i}}{\rightsquigarrow} B \implies H(A) \geq H(A, B)$, it follows that $i\mathcal{M}_G$ bounds the network coding rate in G.

The next section will show that $i\mathcal{M}_G$ is not a tight bound on the network coding rate. We will construct a family of graphs, called the ***iterated split butterfly***, which proves that the iMeagerness gap of a directed graph with n vertices can be at least $\Omega(\log n)$.

[2]This example was communicated to us by J. Kleinberg [8] as an instance in which the vertex-sparsity greatly exceeds the maximum flow value. We learned of this example only after we had discovered $\mathcal{G}_n$.

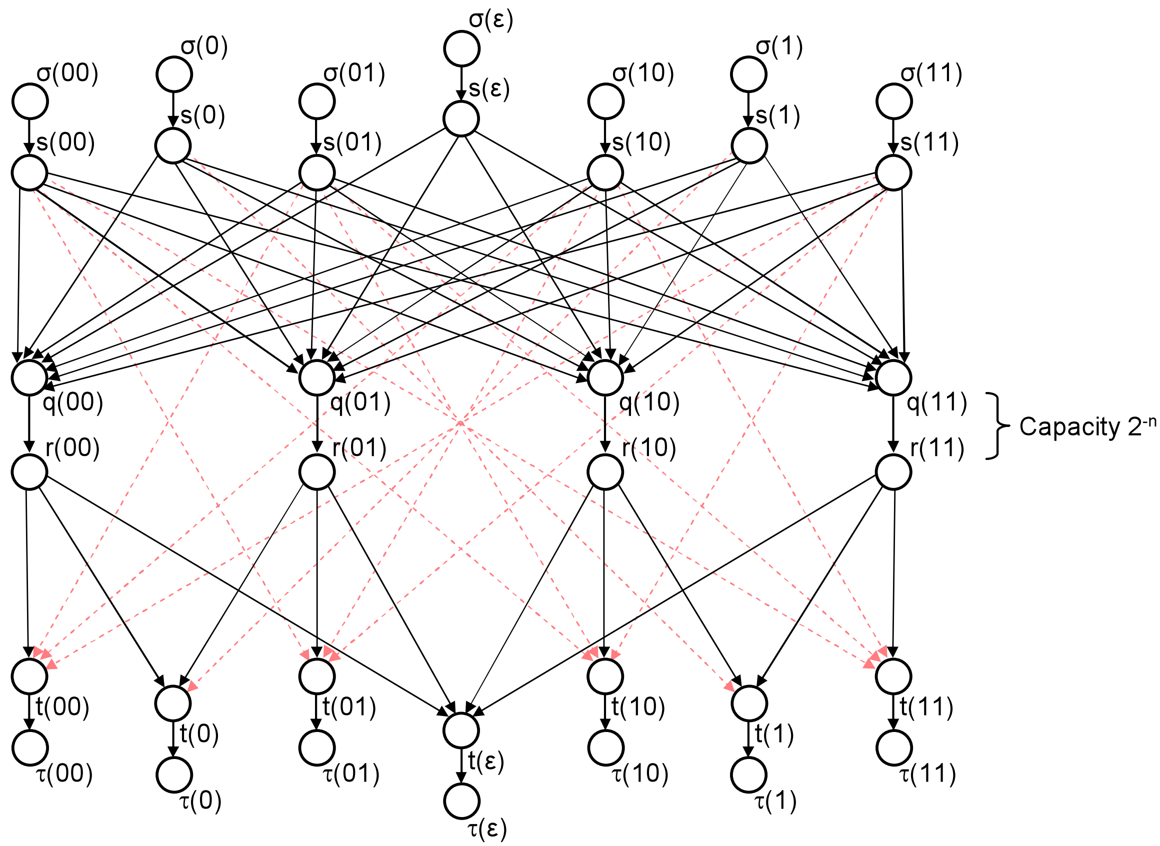

Figure 2: The graph $\mathcal{ISB}(2)$. All edges except $(q(\lambda), r(\lambda))$ have infinite capacity. The edges between $(s(a), t(b))$ for cousins a, b are indicated with dotted red lines.

3.1 Definition of $\mathcal{ISB}(n)$

For a positive integer n, let $\mathbb{T}(n)$ denote a rooted complete binary tree of height n, i.e., $\mathbb{T}(n)$ has $2^{n+1} - 1$ nodes indexed by all the binary strings of length less than or equal to n (including the empty string), and node a is an ancestor of node b if and only if a is indexed by a prefix of the binary string corresponding to b. (We will sometimes refer to the same relation by saying that b is a descendant of a. Note, in particular, that every node is considered to be both an ancestor and a descendant of itself.) The leaves of $\mathbb{T}(n)$ are indexed by the 2^n binary strings of length n. We will refer to leaves of $\mathbb{T}(n)$ using Greek letters $\lambda, \mu, \nu, \ldots$ and we will refer to nodes of $\mathbb{T}(n)$ (including leaves) using Roman letters $a, b, c, \ldots$. For a node a of $\mathbb{T}(n)$ we define the ***depth*** of a to be the length of the corresponding binary string. We define two nodes a, b to be ***cousins*** if they have the same depth but $a \neq b$.

We now define a directed graph $\mathcal{ISB}(n)$ whose vertex and edge sets are based on the structure of $\mathbb{T}(n)$. The vertex set of $\mathcal{ISB}(n)$ consists of the following vertices:

- For every node a of $\mathbb{T}(n)$ (including the leaves) there are four vertices $\sigma(a)$, $s(a)$, $\tau(a)$, $t(a)$ of $\mathcal{ISB}(n)$.
- For every leaf λ of $\mathbb{T}(n)$ there are two vertices $q(\lambda), r(\lambda)$ of $\mathcal{ISB}(n)$.

The edge set of $\mathcal{ISB}(n)$ consists of the following directed edges:

- Edges $(\sigma(a), s(a))$ and $(t(a), \tau(a))$ for every node a of $\mathbb{T}(n)$.

- Edges $(q(\lambda), r(\lambda))$ for every leaf λ of $\mathbb{T}(n)$.
- Edges $(s(a), q(\lambda))$ for every node a and leaf λ.
- Edges $(r(\lambda), t(a))$ for every node a and every leaf λ which is a descendant of a.
- Edges $(s(a), t(b))$ for every pair of cousins a, b.

All edges of $\mathcal{ISB}(n)$ have infinite capacity except for the edges $(q(\lambda), r(\lambda))$, which each have capacity 2^{-n}.

We now define a network coding instance in $\mathcal{ISB}(n)$. The commodities are indexed by the nodes of $\mathbb{T}(n)$. The source edge of commodity a is $(\sigma(a), s(a))$ and the sink edge is $(t(a), \tau(a))$. The demand of commodity a is $2^{-\mathrm{depth}(a)}$.

3.2 Calculating the iMeagerness of $\mathcal{ISB}(n)$

Before we can compute the iMeagerness of $\mathcal{ISB}(n)$, we must describe the combinatorial characterization of informational dominance. Let $G(B, c)$ to be the graph constructed from G by removing all edges that cannot reach $T(c)$, all edges in B, and then all edges not reachable from a source. An ***indirect walk*** is a sequence of directed paths $Q_1, P_1, Q_2, P_2, \ldots, Q_\ell$ where all paths start at a source, Q_1 starts at $S(c)$, Q_j ends at $T(c)$, and Q_i and P_i end at the same vertex for all $1 \le i < \ell$.

Theorem 4 (Harvey et al. [5]). $\mathrm{Dom}(A)$ is the unique minimal set of edges B satisfying the following four criteria:

1. $A \subseteq B$.
2. For every commodity c, the source edge $S(c) = (\sigma(c), s(c))$ belongs to B if and only if the sink edge $T(c) = (t(c), \tau(c))$ belongs to B.
3. Every edge which is not in B is reachable in $G \setminus B$ from a source.
4. For every source edge $S(c)$ which is not in B, there is an indirect walk for commodity c in $G(B, c)$.

Lemma 5. Suppose a, b are cousins in $\mathbb{T}(n)$. Let $T(a, b)$ denote the set of vertices incident to a source or sink edge for commodity a or b, and let

$$Q(a, b) = \{\, q(\lambda) \,:\, \lambda \text{ is a descendant of } a \text{ or } b \,\}.$$

Let $R(a, b)$ be the set of all edges of $\mathcal{ISB}(n)$ with at least one endpoint lying outside the set $Q(a, b) \cup T(a, b)$. Then $\mathrm{Dom}(R(a, b)) = R(a, b)$.

Proof. We simply verify the criteria of Theorem 4 with $A = B = R(a, b)$. Criteria (1) and (2) are trivially satisfied. To verify criterion (3), note that the edges not in $R(a, b)$ are the source/sink edges for commodities a and b, the edges from $\{s(a), s(b)\}$ to $Q(a, b)$, and the edges $(s(a), t(b))$ and $(s(b), t(a))$ (since we assume that a and b are cousins). Clearly all of these edges are reachable from a source edge.

Criterion (4) is verified by arbitrarily choosing leaves λ, μ which are descendants of a, b, respectively, and using the two indirect walks:

$$Q_1 = \sigma(a), s(a), q(\lambda)\ ;\ P_1 = \sigma(b), s(b), q(\lambda)\ ;\ Q_2 = \sigma(b), t(a), \tau(a)$$
$$Q_1 = \sigma(b), s(b), q(\mu)\ ;\ P_1 = \sigma(a), s(a), q(\mu)\ ;\ Q_2 = \sigma(a), t(b), \tau(b).$$

Since all four criteria are satisfied, we have $\mathrm{Dom}(R(a, b)) = R(a, b)$. $\square$

Lemma 6. The iMeagerness of $\mathcal{ISB}(n)$ is at least $1/2$.

Proof. Suppose A is a set of edges of $\mathcal{ISB}(n)$ and P is a non-empty set of commodities informationally isolated by A. (By abuse of notation, we will also use P to denote the corresponding set of nodes of $\mathbb{T}(n)$.) We must prove that $C(A)$ is at least $D(P)/2$. This is trivial if A contains any infinite-capacity edges. Suppose now that A is contained in the set of finite-capacity edges of $\mathcal{ISB}(n)$.

Our key claim is that P can not contain two nodes which are cousins. To see this, suppose for the sake of contradiction that P contains a pair of cousins a, b. Let $\mathcal{S}(P)$ denote the set of source edges for commodities in P and let $\mathcal{S}(\overline{P})$ denote the set of all other source edges. Note that every edge $S(x) \in \mathcal{S}(\overline{P})$ has neither endpoint in $T(a,b) \cup Q(a,b)$, and hence $S(x) \in R(a,b)$. Additionally, every edge in $e \in A$ has finite capacity, so $e = (q(\lambda), r(\lambda))$ for some λ. Since $r(\lambda) \notin T(a,b) \cup Q(a,b)$, it follows that $e \in R(a,b)$. We have shown that $A \cup \mathcal{S}(\overline{P}) \subseteq R(a,b)$. A basic property of Dom is monotonicity: $X \subseteq Y \implies \mathrm{Dom}(X) \subseteq \mathrm{Dom}(Y)$. This property implies that $\mathrm{Dom}(A \cup \mathcal{S}(\overline{P})) \subseteq \mathrm{Dom}(R(a,b)) = R(a,b)$. On the other hand, $S(a) \notin R(a,b)$ by definition, so clearly $\mathcal{S}(P) \not\subseteq R(a,b)$. Thus $\mathcal{S}(P) \not\subseteq \mathrm{Dom}(A \cup \mathcal{S}(\overline{P}))$, i.e., $A \cup \mathcal{S}(\overline{P}) \stackrel{\mathrm{i}}{\not\rightsquigarrow} \mathcal{S}(P)$. This contradicts our assumption that A informationally isolates P.

Our key claim implies that no two distinct nodes in P have the same depth. Let a be the node of minimum depth in P, i.e., a is the closest to the root. Recalling that a node at depth d corresponds to a commodity whose demand is 2^{-d}, we see that the combined demand of all commodities in P is at most $2 \cdot 2^{-\mathrm{depth}(a)}$. So to finish, it suffices to prove that the combined capacity of edges in A is at least $2^{-\mathrm{depth}(a)}$, i.e., that $|A| \geq 2^{n-\mathrm{depth}(a)}$, which is the number of leaves in $\mathbb{T}(n)$ that are descendants of a. In fact, A must contain each edge $(q(\lambda), r(\lambda))$ where λ is a descendant of a, because otherwise the path $\sigma(a), s(a), q(\lambda), r(\lambda), t(a), \tau(a)$ would lie in the complement of $A \cup \mathcal{S}(\overline{P})$, contradicting our hypothesis that $S(a)$ is informationally dominated by $A \cup \mathcal{S}(\overline{P})$. $\square$

3.3 Calculating the network coding rate for $\mathcal{ISB}(n)$

For every node $a \in \mathbb{T}(n)$, define edge sets $E(a), \widetilde{E}(a), F(a)$ in $\mathcal{ISB}(n)$ as follows: $F(a)$ contains the source edge $S(b) = (\sigma(b), s(b))$ for every node $b \in \mathbb{T}(n)$ which is not an ancestor of a; $E(a)$ contains the set $F(a)$, and it also contains the edge $(q(\lambda), r(\lambda))$ for every leaf λ which is a descendant of a. Finally let $\widetilde{E}(a) = E(a) \cup \{S(a)\}$.

Lemma 7. The sets $E(a), \widetilde{E}(a), F(a)$ satisfy the following:

- **Fact 1:** If a is a node whose two children are a_L, a_R then $\widetilde{E}(a_L) \cup \widetilde{E}(a_R) = E(a)$ and $\widetilde{E}(a_L) \cap \widetilde{E}(a_R) = F(a)$.
- **Fact 2:** $E(a) \rightsquigarrow T(a)$, and $\widetilde{E}(a) \subseteq \mathrm{Dom}(E(a))$.

Proof. The first fact follows straightforwardly from the definitions of $E(a), \widetilde{E}(a), F(a)$. The second fact is verified as follows. Let Π be a path from any source to the sink edge $T(a) = (t(a), \tau(a))$. The edge preceding $t(a)$ on the path Π is either of the form $(s(b), t(a))$ for some cousin b of a, or it is of the form $(r(\lambda), t(a))$ for some descendant λ of a. In the former case, Π must contain the edge $S(b) = (\sigma(b), s(b))$ since this is the only incoming edge to $s(b)$. In the latter case, Π must contain the edge $(q(\lambda), r(\lambda))$ since this is the only incoming edge to $r(\lambda)$. Thus, in either case, Π intersects $E(a)$ and therefore

$E(a) \rightsquigarrow T(a)$. Since $S(a)$ and $T(a)$ are equal random variables under any network coding solution, we immediately obtain $\widetilde{E}(a) \subseteq \mathrm{Dom}(E(a))$. $\square$

Theorem 8. The network coding rate for $\mathcal{ISB}(n)$ is at most $1/(n+1)$.

Proof. Our first claim in the proof is the following fact: for $k = 1, 2, \ldots, n$,

$$\sum_{\mathrm{depth}(a)=k} H(E(a)) - \sum_{\mathrm{depth}(a)=k-1} H(E(a)) \geq \sum_{\mathrm{depth}(a)=k-1} H(F(a)). \tag{1}$$

This is proved as follows. Note that for any two nodes a_L, a_R of depth k with a common parent a, we have $H(\widetilde{E}(a_L)) + H(\widetilde{E}(a_R)) \geq H(\widetilde{E}(a_L) \cup \widetilde{E}(a_R)) + H(\widetilde{E}(a_L) \cap \widetilde{E}(a_R)) = H(E(a)) + H(F(a))$, by submodularity of entropy and fact 1. Thus, grouping all nodes of depth k into pairs with a common parent, we obtain

$$\sum_{\mathrm{depth}(a)=k} H(\widetilde{E}(a)) \geq \sum_{\mathrm{depth}(a)=k-1} \big(H(E(a)) + H(F(a))\big). \tag{2}$$

Next, fact 2 implies that $H(E(a)) \geq H(\widetilde{E}(a))$. Hence equation (1) follows.

Now summing equation (1) for $k = 1, 2, \ldots, n$, we obtain

$$\sum_{\mathrm{depth}(a)=n} H(E(a)) - \sum_{\mathrm{depth}(a)=0} H(E(a)) \geq \sum_{\mathrm{depth}(a)<n} H(F(a)). \tag{3}$$

Let r denote the root of $\mathbb{T}(n)$, i.e., the unique node at depth 0. From fact 2, we have $H(E(r)) \geq H(\widetilde{E}(r))$. Next, since $\mathcal{S} = F(r) \cup \{S(r)\} \subseteq \widetilde{E}(r)$ and all sources are independent, we obtain $H(\widetilde{E}(r)) \geq \sum_{a \in \mathbb{T}(n)} H(S(a))$. Hence, substituting into equation (3) and writing a node a at depth n as a leaf λ, we obtain

$$\sum_{\text{leaf } \lambda} H(E(\lambda)) \geq \sum_{\mathrm{depth}(a)<n} H(F(a)) + \sum_{a \in \mathbb{T}(n)} H(S(a)). \tag{4}$$

For a leaf λ let $B(\lambda)$ denote the edge $(q(\lambda), r(\lambda))$ and note that $E(\lambda)$ is the disjoint union of the sets $F(\lambda)$ and $\{B(\lambda)\}$. Hence, $H(F(\lambda)) + H(B(\lambda)) \geq H(E(\lambda))$ (by submodularity). Combining this with (4) we obtain

$$\sum_{\text{leaf } \lambda} H(B(\lambda)) + \sum_{\text{leaf } \lambda} H(F(\lambda)) \geq \sum_{a \in \mathbb{T}(n)} H(S(a)) + \sum_{\mathrm{depth}(a)<n} H(F(a)). \tag{5}$$

We now claim that the following identity holds:

$$\sum_{\text{leaf } \lambda} H(F(\lambda)) = \sum_{\mathrm{depth}(a)<n} H(F(a)). \tag{6}$$

To prove this, let $b \not\preceq a$ denote the relation "b is not an ancestor of a." Using the definition of $F(\lambda)$ and independence of sources, we see that the left side of (6) is equal to

$$\begin{aligned} \sum_{\text{leaf } \lambda} \sum_{b \not\preceq \lambda} H(S(b)) &= \sum_{b \in \mathbb{T}(n)} |\{\, \lambda : \lambda \text{ is a leaf } \wedge\ b \not\preceq \lambda \,\}| \cdot H(S(b)) \\ &= \sum_{b \in \mathbb{T}(n)} (2^n - 2^{n-\mathrm{depth}(b)}) \cdot H(S(b)). \end{aligned}$$

By the same reasoning, the right side of (6) is equal to

$$\begin{aligned}\sum_{\mathrm{depth}(a)<n}\sum_{b\not\preceq a} H(S(b)) &= \sum_{b\in\mathbb{T}(n)} |\{\ a\ :\ \mathrm{depth}(a)<n\ \wedge\ b\not\preceq a\ \}|\cdot H(S(b)) \\ &= \sum_{b\in\mathbb{T}(n)} \big((2^n-1)-(2^{n-\mathrm{depth}(b)}-1)\big)\cdot H(S(b)),\end{aligned}$$

thus establishing the identity (6). Subtracting (6) from (5) we obtain

$$\sum_{\text{leaf }\lambda} H(B(\lambda)) \geq \sum_{a\in\mathbb{T}(n)} H(S(a)).$$

If there exists a network coding solution of rate r, then there exists a constant α such that $H(B(\lambda)) \leq 2^{-n}\cdot\alpha$ for all leaves λ, and $H(S(a)) \geq r\cdot 2^{-\mathrm{depth}(a)}\cdot\alpha$ for all nodes a. Hence

$$\sum_{\text{leaf }\lambda} 2^{-n}\cdot\alpha \geq \sum_{a\in\mathbb{T}(n)} r\cdot 2^{-\mathrm{depth}(a)}\cdot\alpha.$$

This simplifies to $\alpha \geq r(n+1)\alpha$, since $\mathbb{T}(n)$ has 2^n leaves and $\sum_{a\in\mathbb{T}(n)} 2^{-\mathrm{depth}(a)} = n+1$. Hence $r \leq 1/(n+1)$ as claimed. $\square$

Acknowledgements

We thank Varsha Dani for her careful reading of this paper, and for pointing out an error in our original proof of Lemma 2.

References

[1] M. Adler, N. J. A. Harvey, K. Jain, R. Kleinberg, and A. R. Lehman. On the capacity of information networks. To appear in the ACM-SIAM Symposium on Discrete Algorithms, 2006.

[2] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung. Network information flow. *IEEE Transactions on Information Theory*, 46(4):1204–1216, July 2000.

[3] S. P. Borade. Network information flow: Limits and achievability. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, 2002.

[4] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, 1991.

[5] N. J. A. Harvey, R. Kleinberg, and A. R. Lehman. On the capacity of information networks. Submitted to the the special issue of the IEEE Transactions on Information Theory and the IEEE/ACM Transactions on Networking on Networking and Information Theory.

[6] N. J. A. Harvey, R. D. Kleinberg, and A. R. Lehman. Comparing network coding with multicommodity flow for the k-pairs communication problem. Technical Report 964, MIT LCS, Sept. 2004.

[7] K. Jain, V. Vazirani, R. W. Yeung, and G. Yuval. On the capacity of multiple unicast sessions in undirected graphs. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Sept. 2005.

[8] J. Kleinberg. Personal communication, May 2005.

[9] G. Kramer and S. Savari. Progressive d-separating edge set bounds on network coding rates. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Sept. 2005.

[10] G. Kramer and S. A. Savari. On networks of two-way channels. In *Proceedings of the DIMACS Workshop on Algebraic Coding Theory and Information Theory*, Dec. 2003. To appear.

[11] Z. Li and B. Li. Network coding: The case of multiple unicast sessions. In *Proceedings of the 42nd Allerton Annual Conference on Communication, Control, and Computing*, 2004.

[12] L. Song, R. Yeung, and N. Cai. Zero-error network coding for acyclic networks. *IEEE Transactions on Information Theory*, 49:3129–3139, Dec. 2003.

[13] V. Vazirani. *Approximation Algorithms*. Springer, 2001.