Conservative Network Coding

Nicholas J. A. Harvey Kamal Jain Lap Chi Lau Chandra Nair Yunnan Wu

***Abstract*—** Motivated by practical networking scenarios, we introduce a notion of restricted communication called *conservative networking*. Consider a network of lossless links and a number of independent sources. Each node needs to recover a certain subset of the sources. However, each node is *conservative* in that all information it receives can only be a function of the sources it will ultimately recover. For acyclic networks, we show that conservative networking admits a clean characterization: (i) the rates achievable by integer routing, factional routing, and network coding are equal, and (ii) this rate is determined by a simple cut bound. However, this clean characterization does not extend to cyclic networks. We present cyclic examples showing that (i) fractional routing can be strictly better than integer routing, and (ii) network coding can be strictly better than fractional routing. This work underscores the difficulties generally encountered in cyclic networks.

I. Introduction

In their pioneering work on network coding, Ahslwede et al. [1] determined the capacity (maximum achievable rate) for multicasting information in a network of lossless channels with bit-rate constraints. Ahslwede et al. showed that the multicast capacity is equal to the minimum capacity of a cut separating the source from a receiver. Furthermore, an example is given [1], which shows that the traditional routing scheme, where nodes only store and forward data, cannot achieve the multicast capacity in general. Instead, to achieve the capacity, we need to allow nodes to perform *network coding*, i.e., generating output data by encoding (i.e., computing certain functions of) previously received input data.

By now, the problem of single session multicasting is well understood, from both theoretical and practical perspectives. However, the multi-session network coding problem, where multiple multicast sessions with independent data share a network, remains an open challenge.

In this paper, we focus on a restricted form of multi-session communication called *conservative networking*. By being conservative, we require each node to reject any received message that involves information that this node does not need. In other words, all information a node is allowed to receive can only be a function of the sources it will ultimately recover.

There are at least two motivations for considering the conservative networking model. First, in practical networks such as a peer-to-peer network, nodes may be conservative because (i) they have no incentive to forward to others data in which they have no vested interest, and (ii) there are many security issues that arise when data is forwarded through nodes which are not the intended recipients. Second, conservative networking problems involving a single session are theoretically elegant, and optimal solutions can easily be characterized.

Conservative networking with a single session is essentially a broadcasting problem, where the source node wants to transfer the same information to all other nodes (see, e.g., Edmonds [2] or Wu et al. [3]). One of Edmonds' fundamental results in graph theory establishes that the broadcast capacity can be achieved by routing. Specifically, given a directed graph and a source node s, the maximum number of edge disjoint spanning trees rooted at s is equal to the minimum capacity of cuts separating the source from another node. Thus, for conservative networking with a single session, routing is optimal and network coding is not needed to achieve capacity.

This motivates us to investigate whether Edmonds' theorem can be generalized in the multi-session setting. For acyclic networks, we indeed obtain a clean characterization: (i) the rates achievable by integer routing, factional routing, and network coding are equal, and (ii) this rate is determined by a simple cut bound.

Somehow surprisingly, this clean characterization does not extend to cyclic networks. We present cyclic examples showing that (i) fractional routing can be strictly better than integer routing, and (ii) network coding can be strictly better than fractional routing. Our construction is based on a graphical reduction, which transforms any multi-session unicast problem G into a multi-session conservative networking problem G'. We show the graphical reduction preserves the routing capacity and does not decrease the coding capacity. Then by applying the reduction to known multiple unicast examples, we obtain the desired counter examples.

II. Preliminaries

A. Communication Problems

We begin with informal definitions. A *communication problem* is modeled as a directed graph in the following way. Nodes (computers, routers, etc.) participating in the problem are modeled as vertices of the graph. Communication channels between nodes are modeled as edges. If a node has a stream of information that it wishes to share with some other entities, then the information is called a *commodity*, the node which has the information is called a *source*, and the nodes which wish to receive the information are called *receivers*. In general, a communication problem will involve several commodities. More formally, we have

Research performed while all of the authors were at Microsoft Research, Redmond, WA. N. Harvey is currently with the MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA, and is supported by a Natural Sciences and Engineering Research Council of Canada PGS Scholarship, by NSF contract CCF-0515221 and by ONR grant N00014-05-1-0148. L. C. Lau is currently with the Department of Computer Science, University of Toronto, Canada, and is supported by a Microsoft Fellowship.

Definition 2.1: A communication problem is a tuple

$$\big(G = (V, E),\, k,\, (s_1, T_1),\, \ldots,\, (s_k, T_k) \big)$$

where $G = (V, E)$ is a directed graph, $k \in \mathbb{N}$ is an integer specifying the number of commodities, $s_i \in V$ is the source for commodity i, and $T_i \subseteq V - s_i$ is the set of receivers for commodity i.

The channels of a communication problem are used to transmit information. This is usually formalized by defining an alphabet of symbols which can be sent on a channel, and asserting that, at each time step, an edge may send a *message*, which consists of a single symbol from this alphabet. This can be formalized (see, e.g., Harvey et al. [4]), although we do not do so here. Rather, we assume that the reader has previously seen a formalization, or is willing to believe that a reasonable formalization exists. The reader should also note that handling cyclic graphs involves several subtleties.

In general, the channels of a communication problem may have different bandwidths. For simplicity, we ignore this effect and assume that all channels have the same capacity. This assumption amounts to assuming that there is a single alphabet from which any edge may transmit any symbol.

For communication problems arising in practice, it is common that, for each commodity, there is a single receiver node which wishes to receive that commodity. More formally, we have

Definition 2.2: A *k-pairs communication problem* is one in which each $|T_i| = 1$. In the literature, this is also known as a *multiple-session unicast problem.*

Given a communication problem, the natural question to study is: *How much information can be sent from the source of a commodity to the receivers?* To state this question in more detail, we need additional definitions.

B. Solutions

A *solution* is a scheme for transmitting a commodity from its source node to its receiver nodes, via the communication channels. Informally, we require that, for each receiver node, it must be able to reconstruct the information for all commodities that it wishes to receive, given the messages that arrived on its inbound edges. One can define different classes of solutions, based on various schemes for constructing messages.

A *routing solution* is one in which the information of the commodity is regarded as immutable. It cannot be manipulated or coded with other commodities.

The basic sort of routing solution is an *integer routing solution*, in which every edge is associated with a single commodity, and at every time step, the message transmitted on an edge consists of a single symbol from its associated commodity's information stream. In the language of combinatorial optimization, an integer routing solution is a packing of k directed Steiner trees; the i^{th} tree is rooted at s_i, and spans the nodes $T_i + s_i$. For k-pairs communication problems, an integer routing solution is simply a packing of k edge-disjoint directed s_i-t_i paths.

One can also define *fractional routing solutions* where, at each time step, an edge is associated with a single commodity, but at different time steps, the edge can be associated with different commodities. This amounts to multiplexing several integer routing solutions over several time steps. In the language of combinatorial optimization, a fractional routing solution amounts to a fractional packing of Steiner trees. For k-pairs communication problems, a fractional routing solution is simply a (fractional) multicommodity flow.

Finally, the most general sort of solution is a *network coding solution*. In such solutions, the message transmitted on an edge can be an arbitrary function of the edges already received by the edge's tail node. Network coding solutions have no counterpart in combinatorial optimization.

C. Rate

The quality of a solution is determined by the amount of information that it can transmit. This notion is captured by the (concurrent) *rate* of the solution, defined as follows.

- *Integer routing solution*: Let p be an integer. Suppose that a solution packs p Steiner trees for each commodity i, where all of these trees are disjoint. Then we say that the solution has rate p.
- *Fractional routing solution*: Let p be a rational number. Suppose that the solution fractionally packs Steiner trees and the total weight of all trees for commodity i is at least p. Then we say that the solution has rate p.
- *Network coding solution*: Suppose that the solution delivers a symbols of the information stream to the receivers over b time steps. Then we say that the solution has rate $p = a/b$. This is only a vague, intuitive definition; for a formal definition see, e.g., Harvey et al. [4].

Given a communication problem, it is interesting to compare the best rate that can be achieved by solutions of the various classes. We define:

$$\begin{aligned} \mathcal{R}_{\text{int}}(G) &= \text{max rate of an integer routing solution} \\ \mathcal{R}_{\text{frac}}(G) &= \text{max rate of a fractional routing solution} \\ \mathcal{N}(G) &= \text{supremum rate of a network coding solution} \end{aligned}$$

It is very interesting to understand these quantities further. A natural question is: *Are there simple upper and lower bounds on these quantities?*

At the very least, we have the following inequalities

$$\mathcal{R}_{\text{int}}(G) \;\leq\; \mathcal{R}_{\text{frac}}(G) \;\leq\; \mathcal{N}(G), \tag{1}$$

since any integer routing solution is trivially a fractional one, and any fractional routing solution is trivially a coding solution. Some natural questions include: *When are these inequalities strict? When does equality hold?*

D. Conservative Solutions

As stated earlier, this paper focuses on a particular class of solutions called *conservative solutions*. Informally, a solution is conservative if commodity i is never transmitted outside the set $T_i + s_i$. More formally, we have

Definition 2.3: The set of commodities that vertex v wishes to send/receive is denoted $c(v)$. The set of commodities that *both* endpoints of an edge $e = (u, v)$ wish to send/receive is denoted $c(e)$. Thus,

$$\begin{aligned} c(v) &:= \{\, i \,:\, v \in T_i + s_i \,\} \\ c(e) &:= c(u) \cap c(v). \end{aligned}$$

A solution is called conservative if every message transmitted on edge e is only a function of the commodities in $c(e)$.

The notation $\mathcal{R}^{\text{cons}}_{\text{int}}(G)$ is defined to be the maximum rate of a conservative integer routing solution in G. The notations $\mathcal{R}^{\text{cons}}_{\text{frac}}(G)$ and $\mathcal{N}^{\text{cons}}(G)$ are defined similarly. As in Eq. (1), we have

$$\mathcal{R}^{\text{cons}}_{\text{int}}(G) \;\leq\; \mathcal{R}^{\text{cons}}_{\text{frac}}(G) \;\leq\; \mathcal{N}^{\text{cons}}(G). \tag{2}$$

Note that conservative solutions for k-pairs communication problems are rather trivial.

E. LP characterization of $\mathcal{R}^{\text{cons}}_{\text{frac}}(G)$

In general, integer routing rates and network coding rates are quite difficult to analyze. However, fractional routing solutions can be characterized using linear programs. Conservative fractional routing solutions have a particularly clean formulation, which we describe next.

Definition 2.4 (Flow): An *s–t flow* is a nonnegative vector $\boldsymbol{f}$ of length $|E|$ satisfying the *flow conservation constraint*

$$\text{excess}_v(\boldsymbol{f}) = 0 \qquad \forall v \in V - \{s, t\}, \tag{3}$$

where

$$\text{excess}_v(\boldsymbol{f}) \;:=\; \sum_{\substack{e \in E \text{ s.t.} \\ \text{head}(e) = v}} f_e \;-\; \sum_{\substack{e \in E \text{ s.t.} \\ \text{tail}(e) = v}} f_e, \tag{4}$$

is the *flow excess of* v, i.e., the amount of incoming flow minus the amount of outgoing flow for node v. Rather, the flow excess is not required to be zero at s and t. The flow excess at the destination node t (i.e., $\text{excess}_t(\boldsymbol{f})$) is an important quantity called the *value* of the flow.

Let $\mathcal{F}_{s,t}(r)$ denote the set of s–t flows in G, each with flow value r. Then $\boldsymbol{f} \in \mathcal{F}_{s,t}(r)$ if and only if

$$\boldsymbol{f} \geq \boldsymbol{0}, \tag{5}$$
$$\text{excess}_s(\boldsymbol{f}) = -r, \tag{6}$$
$$\text{excess}_t(\boldsymbol{f}) = r, \tag{7}$$
$$\text{excess}_v(\boldsymbol{f}) = 0, \quad \forall v \in V - \{s, t\}\,. \tag{8}$$

Constraint (6) is in fact redundant as it can be derived from (7) and (8). Note that the above inequalities are linear in $\boldsymbol{f}$ and r; for this reason, $\mathcal{F}_{s,t}(r)$ is called the *s–t flow polytope*. A useful property of $\mathcal{F}_{s,t}(r)$ is its linearity in r, i.e.,

$$\mathcal{F}_{s,t}(r) \;=\; r \cdot \mathcal{F}_{s,t}(1) \;:=\; \{\, r\boldsymbol{f} \,:\, \boldsymbol{f} \in \mathcal{F}_{s,t}(1) \,\}\,. \tag{9}$$

The maximum rate of fractional routing, $\mathcal{R}^{\text{cons}}_{\text{frac}}(G)$, is given by the following linear program.

$$\begin{array}{lllll} \max & r & & & \\ \text{s.t.} & \boldsymbol{g}_1 + \cdots + \boldsymbol{g}_k & \leq & \boldsymbol{c} & \\ & g_i(e) & = & 0 & \forall e \notin G_i \\ & \boldsymbol{f}^i_t & \leq & \boldsymbol{g}_i & \forall i \in [k],\ \forall t \in T_i \\ & \boldsymbol{f}^i_t & \in & \mathcal{F}_{s_i,t}(r) & \forall i \in [k],\ \forall t \in T_i \end{array}$$

Here $[k]$ denotes the set $\{1, \ldots, k\}$ and G_i denotes the subgraph induced by the node set $T_i{+}s_i$. This can be explained as follows. We split the total available capacity $\boldsymbol{c}$ into k parts (or subgraphs) $\boldsymbol{g}_1$, …, $\boldsymbol{g}_k$, one for each commodity. Due to the conservative networking constraint, we can assume that the flows for commodity i only use edges in G_i. Each subgraph $\boldsymbol{g}_i$ must support a broadcast rate of r. Due to Edmonds' theorem and the Max-Flow-Min-Cut Theorem, this holds if and only if $\boldsymbol{g}_i$ contains an s_i–t flow $\boldsymbol{f}^i_t$ with rate r for each $t \in T_i$.

F. A Cut Condition

We now derive a "cut condition" (upper bound) on $\mathcal{N}^{\text{cons}}(G)$. First, define the notation

$$E[A, B] \;:=\; \{\, (a, b) \in E \,:\, a \in A,\, b \in B \,\}\,.$$

A *cut* is an arbitrary subset $U \subseteq V$ of the vertices. The edges crossing the cut (from U to $\overline{U}$) are precisely $E[U, \overline{U}]$. Let $C \subseteq [k]$ be an arbitrary subset of the commodities. Let us now consider the edges crossing the cut which, in a conservative solution, can transmit commodities in C. Denote this set by

$$E[U, \overline{U}]_C \;:=\; \left\{\, e \in E[U, \overline{U}] \,:\, C \cap c(e) \neq \emptyset \,\right\}.$$

The commodities which must transmit information across the cut are

$$C_U \;:=\; \left\{\, i \in C \,:\, s_i \in U \ \text{and}\ T_i \cap \overline{U} \neq \emptyset \,\right\}.$$

A simple counting argument yields

Proposition 2.5: For every $U \subseteq V$, we have

$$\mathcal{N}^{\text{cons}}(G) \;\leq\; \frac{|E[U, \overline{U}]_C|}{|C_U|}.$$

Let us define the *conservative cut value* on G to be

$$\mathcal{C}^{\text{cons}}(G) \;:=\; \min_{C \subseteq [k]} \;\min_{U \subseteq V}\; \frac{|E[U, \overline{U}]_C|}{|C_U|}.$$

Then Proposition 2.5 asserts that

$$\mathcal{N}^{\text{cons}}(G) \;\leq\; \mathcal{C}^{\text{cons}}(G). \tag{10}$$

A natural question is: *Is this bound tight?*

III. Acyclic Graphs

We begin our investigation of conservative solutions by considering acyclic graphs. Our main result is as follows.

Theorem 3.1: Consider a communication problem on an acyclic graph G. Then

$$\mathcal{R}^{\text{cons}}_{\text{frac}}(G) \;=\; \mathcal{N}^{\text{cons}}(G) \;=\; \mathcal{C}^{\text{cons}}(G),$$

and

$$\mathcal{R}^{\text{cons}}_{\text{int}}(G) \;=\; \lfloor \mathcal{N}^{\text{cons}}(G) \rfloor$$

(cf. Eq. (2) and Eq. (10)).

Proof. Let $r = \mathcal{C}^{\text{cons}}(G)$, and assume that this quantity is positive. We will construct a conservative fractional routing solution also of rate r, and a conservative integer routing solution of rate $\lfloor r \rfloor$. The argument is iterative, and it examines the vertices in the topological order $v_1, v_2, \ldots$. Suppose that we have already considered vertices $V_i = \{v_1, \ldots, v_i\}$ and that the messages transmitted between those vertices have been adjusted to conform to a conservative fractional routing solution.

Consider now the vertex v_{i+1}, and assume for convenience that it is not a source for any commodity. The edges inbound to v_{i+1} are precisely

$$E[V_i, \{v_{i+1}\}] \;=\; E[V - v_{i+1}, \{v_{i+1}\}],$$

since we consider the vertices in topological order.

Let us now construct a bipartite graph H_{i+1} with vertex set $X \cup Y$ where

$$\begin{aligned} X &= \{\, x_j \,:\, j \in c(v_{i+1}) \,\} \\ Y &= \{\, y_e \,:\, e \in E[\,V - v_{i+1},\, \{v_{i+1}\}\,] \,\}, \end{aligned}$$

and the x_j's and y_e's are new vertices. The graph H_{i+1} has an edge (x_j, y_e) iff $j \in c(e)$. Now fix a set $U \subseteq X$, and let $C = \{\, j : x_j \in U \,\}$. Let $\delta(U) \subseteq Y$ denote the neighbors of U in H_{i+1}. Then clearly

$$\delta(U) \;=\; \{\, y_e \,:\, e \in E[\,V - v_{i+1},\, \{v_{i+1}\}\,]_C \,\}.$$

Therefore

$$\begin{aligned} |\delta(U)| &= |E[\,V - v_{i+1},\, \{v_{i+1}\}\,]_C| \\ &\geq r \cdot |C_{V - v_{i+1}}| \quad \text{(by Proposition 2.5)} \\ &= r \cdot |c(v_{i+1}) \cap C| \\ &= r \cdot |U|. \end{aligned} \qquad (11)$$

By Hall's Theorem (see the Appendix), Eq. (11) implies that there exists a fractional matching M where each node in X has (fractional) degree at least r. In other words, there exist $w_{j,e} \geq 0$ such that

$$\begin{aligned} \sum_e w_{j,e} &\geq r \quad \forall j \\ \sum_j w_{j,e} &\leq 1 \quad \forall e. \end{aligned}$$

Since G is acyclic, the nodes in V_i must have fully decoded all commodities for which they are a receiver, using the routing solution constructed so far. This means that, if $w_{j,e} > 0$ then the tail of edge e has already fully decoded commodity j, so it can indeed provide a $w_{j,e}$ fraction of commodity j to v_{i+1} on edge e. Using this fractional matching, we augment the conservative fractional routing solution on V_i to one on V_{i+1}.

The same argument applies to the case of integral routing solutions. Eq. (11) also implies, via Hall's Theorem, that there exists an integral matching M for which each node in X has (fractional) degree at least $\lfloor r \rfloor$. Therefore a conservative integer routing solution on V_i can be augmented to one on V_{i+1} with rate at least $\lfloor r \rfloor$. ■

IV. Cyclic Graphs

In this section we investigate whether the results of the previous section can be extended to graphs with cycles. Surprisingly, the answer is no!

Theorem 4.1: There exist graphs G_1 and G_2 such that

$$\mathcal{R}^{\text{cons}}_{\text{int}}(G_1) \;\leq\; \mathcal{R}^{\text{cons}}_{\text{frac}}(G_1) - 1 \qquad (12)$$

$$\mathcal{R}^{\text{cons}}_{\text{frac}}(G_2) \;<\; \mathcal{N}^{\text{cons}}(G_2) \qquad (13)$$

Additionally, computing $\mathcal{R}^{\text{cons}}_{\text{int}}(G)$ is NP-hard.

A. The Reduction

The argument depends on a certain reduction from k-pairs communication problems to general communication problems. Let G denote a k-pairs problem. From this, we will construct another communication problem G'.

The reduction from G to G'.

The new instance G' is constructed as follows. For each commodity, we create a new dummy node u_i. The vertices are

$$V' \;=\; V \cup \{\, u_i \,:\, 1 \leq i \leq k \,\}.$$

The edges are

$$\begin{aligned} E' \;=\; & E \,\cup\, \{\, (t_i, u_i) \,:\, 1 \leq i \leq k \,\} \\ & \cup\, \{\, (u_i, v) \,:\, 1 \leq i \leq k,\, v \in V \,\}. \end{aligned}$$

The receivers for commodity i are

$$T_i \;=\; V + u_i.$$

An example of this reduction is shown in Fig. 1.

Lemma 4.2: Any rate-r integer routing solution for G corresponds to a rate-r conservative integer routing solution for G', and vice-versa.

Proof. $\Rightarrow$: Given the integer routing solution for G, we additionally transmit commodity i on the edges (t_i, u_i) and (u_i, v) for each $v \in V$. Every vertex in $V + u_i = T_i$ thereby receives commodity i, so this is an integer routing solution for G'. It is clearly conservative.

$\Leftarrow$: Suppose that we have a conservative integer routing solution for G', i.e., an integer packing of directed Steiner trees. The Steiner tree for commodity i necessarily contains

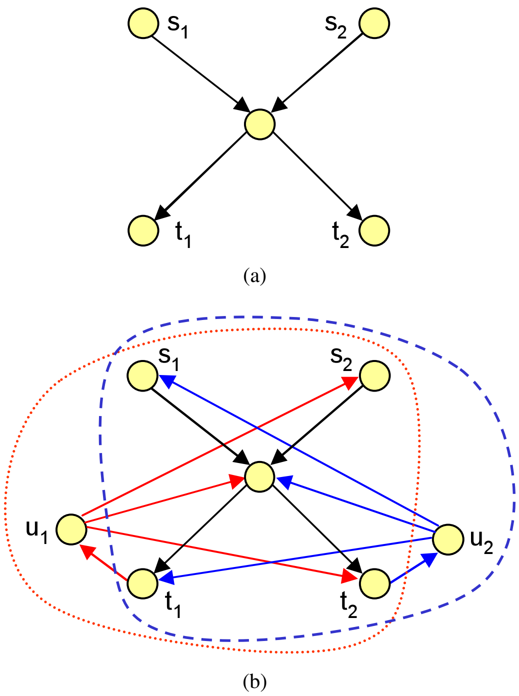

Fig. 1. (a) A graph G. (b) The graph G' obtained by applying the reduction.

a path from s_i to u_i, since u_i is a receiver for commodity i (i.e., $u_i \in T_i$). This path does not traverse any u_j where $j \neq i$ since $u_j \notin T_i$ and since we assume that the given solution is conservative. The penultimate vertex on this path must be t_i, so we obtain an s_i-t_i path P_i which traverses only vertices in V. The paths $P_1, \ldots, P_k$ are disjoint (since they came from an integer Steiner packing) and therefore they form an integer routing solution for G. ∎

The same argument extends easily to the fractional case.

Lemma 4.3: Any rate-r fractional routing solution for G corresponds to a rate-r conservative fractional routing solution for G', and vice-versa.

For the case of network coding, the easier direction of the argument continues to hold.

Lemma 4.4: Any rate-r network coding solution for G yields a rate-r conservative network coding solution for G'.

Thus we have shown that, for any k-pairs communication problem G,

$$\begin{aligned} \mathcal{R}_{\text{int}}(G) &= \mathcal{R}^{\text{cons}}_{\text{int}}(G') \\ \mathcal{R}_{\text{frac}}(G) &= \mathcal{R}^{\text{cons}}_{\text{frac}}(G') \\ \mathcal{N}(G) &\leq \mathcal{N}^{\text{cons}}(G'). \end{aligned}$$

One might be tempted to conjecture that $\mathcal{N}(G) = \mathcal{N}(G')$ should also hold, but surprisingly this is false! This result, proven in Section IV-D, gives yet another example of the counter-intuitive nature of network coding.

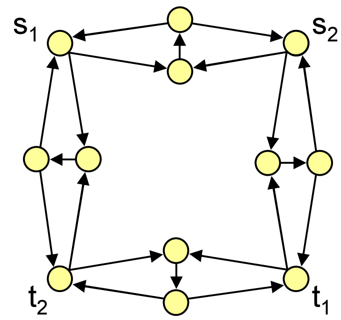

Fig. 2. The communication problem proving Proposition 4.5.

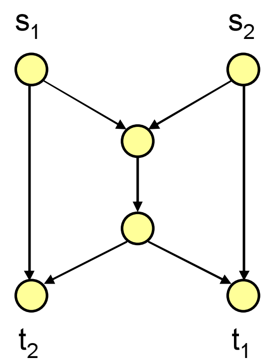

Fig. 3. The communication problem proving Proposition 4.6.

B. Existing Results

Before proving our main theorem, we recall some well-known results from the literature.

Proposition 4.5 (Folklore): There exists a k-pairs communication problem G for which $\mathcal{R}_{\text{int}}(G) \leq \mathcal{R}_{\text{frac}}(G) - 1$.

Proof. Let G be the communication problem shown in Fig. 2. It has $\mathcal{R}_{\text{int}}(G) = 0$ but $\mathcal{R}_{\text{frac}}(G) = 1$. ∎

Proposition 4.6: There exists a k-pairs communication problem G in which $\mathcal{R}_{\text{frac}}(G) < \mathcal{N}(G)$.

Proof. The "butterfly graph" (Ahlswede et al. [1]) has a k-pairs variant G shown in Fig. 3. This graph has $\mathcal{R}_{\text{frac}}(G) = 1/2$ but $\mathcal{N}(G) \geq 1$. ∎

Theorem 4.7 (Fortune et al. [5]): Given a k-pairs communication problem G, the problem of deciding whether $\mathcal{R}_{\text{int}}(G) = \mathcal{R}_{\text{frac}}(G)$ is NP-complete.

In fact, a much stronger result can be shown.

Theorem 4.8 (Guruswami et al. [6]): It is NP-hard to approximate the value of $\mathcal{R}_{\text{int}}(G)$ to within a factor of $|E|^{1/2-\epsilon}$ for any $\epsilon > 0$.

C. Proof of Main Theorem

We now assemble the results mentioned earlier to proof the main theorem of this section.

Proof (of Theorem 4.1). Apply the reduction of Section IV-A to the graph of G of Proposition 4.5, obtaining a graph G_1. This graph has $\mathcal{R}^{\text{cons}}_{\text{int}}(G_1) = 0$ and $\mathcal{R}^{\text{cons}}_{\text{frac}}(G_1) = 1$, thereby proving Eq. (12).

Apply the reduction of Section IV-A to the graph of G of Proposition 4.6, obtaining a graph G_2. This graph has $\mathcal{R}^{\text{cons}}_{\text{frac}}(G_2) = 1/2$ and $\mathcal{N}^{\text{cons}}(G_2) = 1$, thereby proving Eq. (13).

Apply the reduction of Section IV-A to the family of instances used in Theorem 4.7. This shows that computing the value $\mathcal{R}^{\text{cons}}_{\text{int}}(G)$ is NP-hard. ∎

For the sake of clarity, we now give a second, more direct proof of Eq. (13).

Proof. Consider the communication problem G in Fig. 4 (a). We claim that $\mathcal{R}^{\text{cons}}_{\text{frac}}(G) < \mathcal{N}^{\text{cons}}(G)$. It is important to note that G is (just barely) cyclic; if it were acyclic, then we would have $\mathcal{R}^{\text{cons}}_{\text{frac}}(G) = \mathcal{N}^{\text{cons}}(G)$ by Theorem 3.1.

The first step is to show that $\mathcal{R}^{\text{cons}}_{\text{frac}}(G) \leq 1/2$. (In fact, this rate is obviously achievable.) Suppose that there exists a conservative fractional routing solution of rate r.

- First, consider the set $U = \{v_4, v_5, v_7\}$. Both commodities must enter U, via the two edges e_1 and e_2. Since our solution is conservative, e_1 can only send commodity 1. So r units of e_2's capacity must be used commodity 2.
- Next, consider the set $U = \{v_4\}$. Both commodities must enter U, via the two edges e_2 and e_3. Since our solution is conservative, e_3 can only send commodity 2. So r units of e_2's capacity must be used for commodity 1.

Combining these observations, we see that e_2 must send r units of both commodity 1 and commodity 2, so $r \leq 1/2$.

The second step is to show that $\mathcal{N}^{\text{cons}}(G) \geq 1$. (In fact, this rate is obviously optimal.) This follows from the rate-1 network coding solution shown in Fig. 4 (b). Note that the timing of message transmissions in the solution is critical — commodity B must arrive at v_4 on e_3 before v_4 can transmit commodity A on e_4. ∎

D. A Graph G with $\mathcal{N}(G) < \mathcal{N}^{\text{cons}}(G')$

Let G be the k-pairs communication problem shown in Fig. 5 (a).

Lemma 4.9: $\mathcal{N}(G) = 2/3$.

Proof. First we show that $\mathcal{N}(G) \geq 2/3$. Imagine divide each commodity into two parts, each of size $1/3$ unit. For example, we write commodity A as A_1, A_2, and the combined size of A_1 and A_2 is $2/3$. The network coding solution shown in Fig. 5 (b) achieves rate $2/3$.

Now we show that $\mathcal{N}(G) \leq 2/3$. We treat each commodity as a random variable, then use a sequence of entropy arguments and Markov chain arguments to establish the bound. For a detailed description of this technique see, Harvey et al. [4], Jain et al. [7], Kramer and Savari [8], etc.

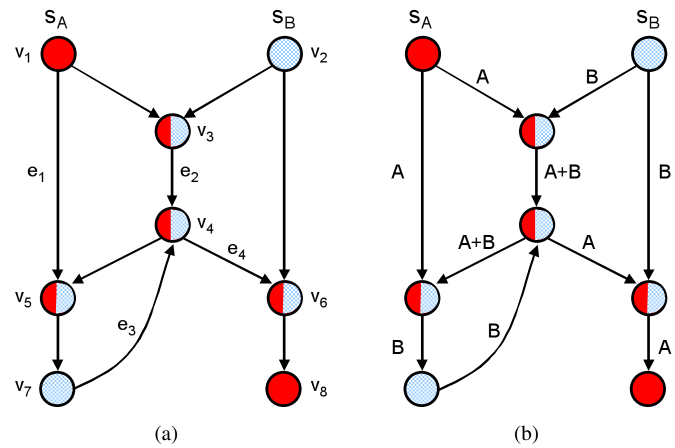

Fig. 4. (a) The communication problem to prove Eq. (13). The source for commodity 1 is v_1, and the receivers are $T_1 = \{v_3, v_4, v_5, v_6, v_8\}$, shown in red. The source for commodity 2 is v_2, and the receivers are $T_2 = \{v_3, v_4, v_5, v_6, v_7\}$, shown in light blue. (b) A conservative network coding solution of rate 1.

$$
\begin{aligned}
&H(s_D) + H(e_4) + H(e_{15}) \\
&\geq H(s_D, e_4, e_{13}, e_{15}, e_{17}) \\
&\geq H(s_D, e_4, e_{13}, e_{15}, e_{17}, e_{10}, e_{18}, e_{19}) \\
&\geq H(s_D, e_4, e_{13}, e_{15}, e_{17}, e_{10}, e_{18}, e_{19}, s_A) \\
&\geq H(s_D, e_4, e_{13}, e_{15}, e_{17}, e_{10}, e_{18}, e_{19}, s_A, e_{14}) \\
&\geq H(s_D, e_4, e_{13}, e_{15}, e_{17}, e_{10}, e_{18}, e_{19}, s_A, e_{14}, s_C) \\
&\geq H(s_D, e_4, e_{13}, e_{15}, e_{17}, e_{10}, e_{18}, e_{19}, s_A, ..., s_C, e_5, e_7) \\
&\geq H(s_D, e_4, e_{13}, e_{15}, e_{17}, ..., s_A, ..., s_C, e_5, e_7, e_{12}) \\
&\geq H(s_D, e_4, e_{13}, e_{15}, e_{17}, ..., s_A, ..., s_C, e_5, e_7, e_{12}, e_{16}) \\
&\geq H(s_D, ..., e_{17}, ..., s_A, ..., s_C, ..., e_{16}, e_{21}) \\
&\geq H(s_D, ..., e_{17}, ..., s_A, ..., s_C, ..., e_{21}, s_B) \\
&\geq H(s_A, s_B, s_C, s_D) \\
&= H(s_A) + H(s_B) + H(s_C) + H(s_D)
\end{aligned}
$$

Rewriting, $H(s_A) + H(s_B) + H(s_C) \leq H(e_4) + H(e_{15}) \leq 2$, implying that the network coding rate is at most $2/3$. ∎

Lemma 4.10: $\mathcal{N}(G') = 1$.

Proof. The graph G' is as shown in Fig. 5 (c). This figure also describes a conservative network coding solution of rate 1. The crucial detail is the new edges from u_B and u_C to the central vertex with out-degree 4. These edges can respectively transmit the unencoded commodities B and C to the central vertex. Therefore the unencoded commodities B and C can be transmitted to t_A and t_D, allowing them to decode A and D respectively. As long as t_A, t_B, t_C, t_D have decoded their desired commodities, all receivers can decode their desired commodities due to the edges (not shown) leaving the u_i vertices. ∎

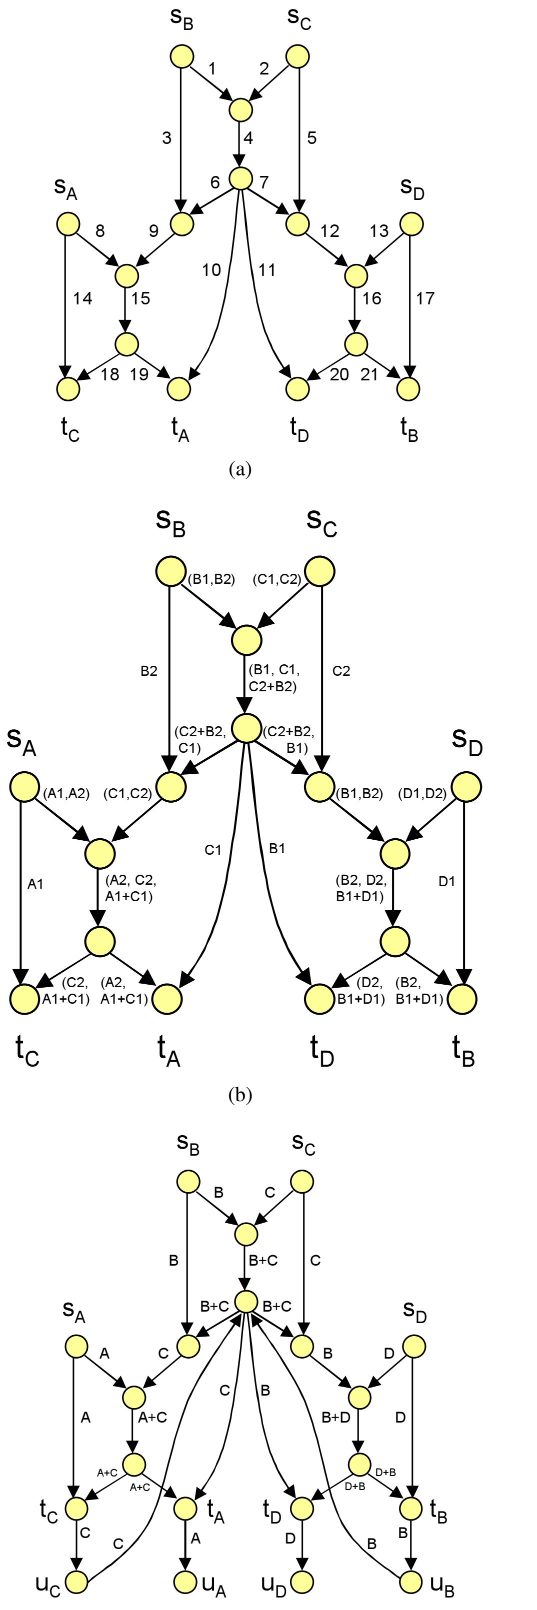

Fig. 5. (a) The k-pairs communication problem for proving Lemma 4.9 and Lemma 4.10. (b) An optimal network coding solution of rate $2/3$. (c) The graph, after being transformed by the reduction of Section IV-A, has a network coding solution of rate 1.

References

[1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1204–1216, July 2000.

[2] J. Edmonds, "Edge-disjoint branchings," in *Combinatorial Algorithms*, ser. Courant Computer Science Symposium, R. Rustin, Ed. Algorithmics Press, 1973, vol. 9, pp. 91–96.

[3] Y. Wu, K. Jain, and S.-Y. Kung, "A unification of network coding and tree-packing (routing) theorems," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2398–2409, June 2006.

[4] N. J. A. Harvey, R. Kleinberg, and A. R. Lehman, "On the capacity of information networks," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2345–2364, June 2006.

[5] S. Fortune, J. Hopcroft, and J. Wyllie, "The directed subgraph homeomorphism problem," *Theoretical Computer Science*, vol. 10, pp. 111–121, 1980.

[6] V. Guruswami, S. Khanna, R. Rajaraman, F. B. Shepherd, and M. Yannakakis, "Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems," in *Proceedings of the ACM Symposium on Theory of Computing*, 1999, pp. 19–28.

[7] K. Jain, V. V. Vazirani, and G. Yuval, "On the capacity of multiple unicast sessions in undirected graphs," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2805–2809, June 2006.

[8] G. Kramer and S. A. Savari, "Edge-cut bounds on network coding rates," *Journal of Network and Systems Management*, vol. 14, no. 1, Mar. 2006.

[9] P. Hall, "On representatives of subsets," *The Journal of the London Mathematical Society*, vol. 10, pp. 26–30, 1935.

[10] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*. Springer-Verlag, 2003.

Appendix
Hall's Theorem

In this appendix, we state the generalization of Hall's theorem used in Section III. Let $H = (X \cup Y, E)$ be a bipartite graph. For $U \subseteq X$, let $\delta(U)$ denote the set of neighbors of vertices in U.

Theorem 1.1 (Hall [9]): If $|\delta(U)| \geq |U|$ for all $U \subseteq X$, then H has a perfect matching.

Proof. See Schrijver [10, Section 22]. ■

Hall's theorem has the following "packing generalization".

Corollary 1.2: Let k be a positive integer. If $|\delta(U)| \geq k \cdot |U|$ for all $U \subseteq X$, then H has k disjoint perfect matchings.

Proof. See Schrijver [10, Theorem 22.10]. The idea is to split each X-vertex into k copies (with identical neighbors), obtaining a graph H'. By Theorem 1.1, H' has a perfect matching, which can be split into k disjoint perfect matchings of H. ■

The generalization that we require is as follows.

Corollary 1.3: Let $r = p/q$ be a rational number, where p and q are positive integers. If $|\delta(U)| \geq r \cdot |U|$ for all $U \subseteq X$, then H contains p matchings $M_1, \ldots, M_p$ such that:

- For every $i \in [p]$, each X-vertex is incident with exactly one edge in M_i, and
- Each Y-vertex has at most q incident edges amongst all p of the matchings.

Proof. The proof is similar to Corollary 1.2 except that each X-vertex is split into p copies, and each Y-vertex is split into q copies. ■