

On the Capacity of Information Networks

Nicholas J. A. Harvey, Robert Kleinberg and April Rasala Lehman

Abstract—An outer bound on the rate region of noise-free information networks is given. This outer bound combines properties of entropy with a strong information inequality derived from the structure of the network. This blend of information theoretic and graph theoretic arguments generates many interesting results. For example, the capacity of directed cycles is characterized. Also, a gap between the sparsity of an undirected graph and its capacity is shown. Extending this result, it is shown that multicommodity flow solutions achieve the capacity in an infinite class of undirected graphs, thereby making progress on a conjecture of Li and Li. This result is in sharp contrast to the situation with directed graphs, where a family of graphs are presented in which the gap between the capacity and the rate achievable using multicommodity flows is linear in the size of the graph.

Index Terms—Capacity, Informational Dominance, k -pairs Communication Problems, Multicommodity Flow, Multiple Unicast Sessions, Network Coding, Sparsity

I. INTRODUCTION

WE study the capacity of information networks, focusing on techniques for proving outer bounds. Understanding the capacity of information networks is a long-standing open problem [1]. In this paper, we consider only networks that are noise-free, interference-free, and where any two-way channels have simple linear capacity regions. To completely characterize the capacity of such networks, one would desire computable outer bounds on capacity and algorithms for finding solutions that achieve those bounds. Recently the term *network coding* has been used to describe algorithms for communication problems in information networks. Proving outer bounds on the capacity of a network is an important ingredient in its complete characterization. From an algorithmic perspective, outer bounds are useful for proving performance guarantees: one can prove that a solution is nearly optimal by comparing it to an outer bound.

Two different fields have sought to understand the capacity of a network: combinatorial optimization and information theory. In combinatorial optimization, the capacity of a network is typically determined by analyzing packing problems (e.g., flows) that are constrained by the graph structure. This

theory grew out of a need to understand shipment of cargo in transportation networks [2] and does not capture the subtleties of information transmission. On the other hand, information theory provides a deep understanding of complex communication problems over structurally simple channels but does not yet fully extend to arbitrary graph structures. Combining ideas from both of these theories allows us to make progress on understanding the capacity of information networks.

Our work is motivated by the following questions, all of which have been considered in prior work, e.g., [1], [3]–[12]. Further discussion of this prior work is given in Section I-B.

- What structural properties of a graph (e.g., cuts) can be used to prove outer bounds on the feasible rate region? Are these bounds tight?
- Can we obtain improved bounds by combining these ideas with information theoretic techniques?
- Do the answers to the foregoing questions differ for directed graphs and undirected graphs?
- When do information networks behave like transportation networks, from the perspective of achievable rates?

Our investigation of these questions leads to several interesting results, which we summarize in the next section.

A. Our Results

We first give a brief summary of our contributions and elaborate on them below. Our key contributions are:

- a new computable outer bound on the rate region for network coding problems on all graphs, including cyclic and undirected graphs.
- a proof of a gap between sparsity and achievable rate in undirected graphs.¹
- the concept of *informational dominance*, which allows us to derive strong information theoretic inequalities from the structure of the graph.
- a proof that the network coding rate equals the multicommodity flow rate for any directed cycle.
- an example showing a tight gap between the network coding rate and multicommodity flow rate in directed acyclic graphs.

Some of these statements concern the *k*-pairs communication problem, in which the network must support k point-to-point connections. The k -pairs communication problem is closely related to the classical multicommodity flow problem. A motivation for our work is understanding the differences

Manuscript received March 14, 2005; revised February 1, 2006. A preliminary version of this work was presented at the DIMACS Working Group on Network Coding, Piscataway, NJ, January 2005.

N. Harvey is with the MIT Computer Science and Artificial Intelligence Laboratory. Portions of the first author's work were supported by a Natural Sciences and Engineering Research Council of Canada PGS Scholarship. R. Kleinberg is with the Cornell University Department of Computer Science. Portions of this work were done while the second author was with the MIT Department of Mathematics and the Berkeley Computer Science Division, supported by a Fannie and John Hertz Foundation Fellowship and an NSF Mathematical Sciences Postdoctoral Research Fellowship, respectively. A. Lehman is with Google Inc. Portions of this work were done while the third author was with the MIT Computer Science and Artificial Intelligence Laboratory.

¹The sparsity is the minimum, over all cuts, of the capacity of the edges in the cut divided by the demand of all commodities that must cross the cut; for more formal definitions, see Section III. To the best of our knowledge, no such gap between sparsity and achievable rate was previously known. The same gap was proven independently and concurrently by Jain et al. [13] and Kramer and Savari [14].

and similarities between the k -pairs communication problem, the corresponding multicommodity flow problem and their respective combinatorial characterizations. Thus many of our results are described in terms of the k -pairs communication problem, but later generalized to arbitrary communication problems in Section XI.

We lay the groundwork for the paper in Section II with formal definitions for network coding in directed acyclic graphs (DAGs). We then consider two cut conditions both of which are upper bounds on the maximum multicommodity flow rate. The first condition, *sparsity*, is not an upper bound on the network coding rate; this follows from an example of Ahlswede et al. [8]. We illustrate an extreme form of this phenomenon in Section III: a gap between the maximum network coding rate and the maximum multicommodity flow rate that is tight in terms of the number of vertices, edges, and commodities in the problem instance.

The shortcomings of sparsity motivate our second cut condition, defined in Section IV, which we call *meagerness*. We prove that meagerness is an upper bound on the network coding rate but give an example where this bound is not tight. This proof illustrates an important synergy between combinatorics and information theory: by combining bounds based on the graph structure with information theoretic inequalities (e.g., submodularity of entropy) one can obtain bounds that are tighter than those proven using cuts alone. This same principle has been employed in previous and concurrent work, e.g., [3], [7], [13], [14].

In Section V, we generalize the idea of obtaining information theoretic bounds based on the graph structure. We seek to understand precisely when the information in one part of the graph necessarily determines the information in another part of the graph. More formally, 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 . The relation is required to hold for all network coding solutions, regardless of the rate of the solution. (In defining the informational dominance relation we implicitly assume that sources are uncorrelated, or more generally that every k -tuple of messages has a positive probability of being transmitted by the k sources.) If this relation holds then the capacity of the edges in A is an upper bound on the entropy of the information transmitted on edges in $A \cup B$. For this reason, informational dominance plays a key role in our investigation of network capacity. We also present a complete combinatorial characterization of the informational dominance relation. This yields a polynomial time algorithm for the following problem: Given an instance of the network coding problem on a graph G and an edge set A , find all edges informationally dominated by A .

We then define network coding in graphs with cycles and in particular undirected graphs. In discussing the capacity of cyclic networks, Song et al. state that “existing tools are not powerful enough, and we need to develop new tools for this class of problems” [3]. One difficulty is in ensuring that network codes in such graphs have an operational meaning, because naïve definitions of network coding in graphs with cycles can lead to cyclic dependencies among the coding

functions, resulting in codes whose interpretation is either ambiguous or self-contradictory. Our approach is to define an infinite directed acyclic graph G^* that represents the flow of information in G over time; essentially the same transformation was used by other authors, e.g. [8], [15], [16]. Although G^* resolves the operational issues, one would prefer to reason about rate bounds in G itself. We achieve this by extending the information dominance characterization from G^* to G . This allows us to state an outer bound for general graphs based on the structure of G itself, without reference to the infinite graph G^* .

We employ these techniques to establish two bounds. First, we show that the network coding rate equals the multicommodity flow rate for all k -pairs communication problems on directed cycles. This question was previously investigated by Kramer and Savari [7]. Second, we consider the conjecture of Li and Li [17] which states that the network coding rate equals the multicommodity flow rate for all k -pairs communication problems on undirected networks. In undirected graphs, the sparsity is an upper bound on both the maximum multicommodity flow rate and the maximum network coding rate. The interesting cases of the conjecture arise when there is a gap between the sparsity and the maximum multicommodity flow rate. We consider one such example, the Okamura-Seymour graph. We prove that the network coding rate equals the maximum multicommodity flow rate in this example. To our knowledge, this is the first demonstration of a provable gap between the sparsity bound and the network coding rate in undirected networks.² Our proof extends to show that the network coding rate equals the maximum concurrent flow rate in an infinite family of instances generalizing the Okamura-Seymour graph; this family includes infinitely many instances in which the sparsity bound is not tight.

B. Related Work

Previous work in the network coding literature has investigated how the capacity of an information network can increase when coding operations on the information are permitted. The seminal work of Ahlswede et al. [8] gives a simple example, shown in Figure 1(a), where the use of coding can increase the capacity. Much of the subsequent work focused primarily on *multicast problems*, which are problems with a single information source and multiple sinks. For multicast problems, it has been shown that coding can increase capacity by a polylogarithmic factor [9], [18] in directed graphs and a constant factor in undirected graphs [9], [10].

Recently, much attention has been given to the k -pairs communication problem [4], [5], [7], [10], [13], [14], [17], which is also known as a problem of “multiple unicast sessions”. A key goal of this previous work is to understand

²This claim relates our bound to prior work. Upper bound techniques similar to ours have been proposed concurrently in [13] and [14]. In Section V-B and Section VII, we formally compare our techniques to the bound of [13]. A complete draft of [14] was not available to the authors at the time that this paper was submitted, and therefore it was not known whether informational dominance is at least as strong as their bound. A rigorous comparison with the bound of [14] is left to future work. For the benefit of the reader, we include references to both [13] and [14] in the subsequent text to indicate results that have been concurrently and independently derived elsewhere.

the network coding rate for k -pairs communication problems and to compare this rate to the corresponding multicommodity flow rate. An example due to Li and Li [4] and Harvey et al. [5] shows that coding can increase the capacity by an unbounded factor in directed graphs, as the size of the graph tends to infinity. This paper improves on the prior work by giving the strongest known increase in capacity with coding for directed graphs. Also, we give a new upper bound for the network coding rate in undirected graphs; to the best of our knowledge, this is the strongest known such bound.²

A general outer bound on the capacity of information networks (including those with noise and interference) can be found in standard textbooks [1, Section 14.10]. This bound was used by Borade [6] to give a tight outer bound on the capacity of multicast networks. Extending this result to networks with multiple commodities, Borade obtains a bound on the capacity of directed information networks which is essentially the sparsity bound. Kramer and Savari [7] give an general outer bound on the capacity of (directed or undirected) information networks called the *bidirected cut-set bound*. This work shows how bounds based on cuts can be applied to information networks, even when undirected channels have complex rate regions. Kramer and Savari demonstrate how to apply their bound to various examples, including a directed 3-cycle. Their analysis demonstrates the effectiveness of combining combinatorial and information-theoretic techniques.

Previous work on DAGs has considered relations similar to informational dominance. Song et al. [3] define a linear program containing the constraint that every edge must be a function of the edges inbound to its tail vertex. Jain et al. [13] independently and concurrently derived related inequalities known as the “input-output inequality” and “crypto inequality”. These conditions are all restricted forms of informational dominance in that their transitive closure does not necessarily determine all informational dominance relations. The notion of informational dominance is closely related to conditional independence of random variables in a functional dependence graph. Kramer [11], [12] has explored this direction in depth, and showed that a condition from Bayesian networks called *d-separation* is sufficient to establish certain informational dominance relationships. Techniques from all of the aforementioned prior work² can be used to prove *sufficient* conditions for informational dominance. Our work is the first to give *necessary and sufficient* conditions for informational dominance. It is not known whether our necessary and sufficient conditions, when combined with Shannon’s basic inequalities for entropy, lead to strictly tighter bounds than those obtainable using other existing sufficient conditions for informational dominance, e.g., [11]–[13]. It should be stressed that our bounds apply only to noise-free networks, whereas some of the other existing approaches such as [11], [12] may be applied to networks with noisy channels as well.

Jain et al. [13] and Kramer and Savari [14] also examined the capacity of the Okamura-Seymour example and proved that the network coding rate equals the multicommodity flow rate. Jain et al. extend their bound to an infinite class of bipartite networks, as do we.

The outer bound on the capacity of DAGs that we describe

in Section IV-C defines a rate region which is contained in the rate region $\mathfrak{R}_{\text{out}}$ defined by the outer bound of Song et al. [3]. Previous work, e.g. [7], [11], [12], has derived inner and outer bounds for undirected and cyclic networks. We also obtain an outer bound for such networks, through the use of informational dominance.

II. DEFINITIONS FOR DAGS

In this section we define a model for network coding in directed acyclic graphs. There are two parts to our model. First, we define the type of communication problem that we consider. Then we define the concept of a network coding solution. We conclude the section by relating our definitions to the information theoretic formulation which is more prevalent in the literature.

Our definition of a network coding instance involves graphs which may be countably infinite. Allowing infinite graphs facilitates the definition of undirected instances in Section VI. However, the word “path” should be interpreted throughout this paper as meaning “finite-length path”, and the word “reachable” should be interpreted as meaning “reachable by a finite-length path”.

Definition 1 (Instance of the Network Coding Problem):

An instance of the network coding problem is specified by the following:

- A directed acyclic graph $G = (V, E)$, which may be countably infinite.
- A nonnegative real capacity $c(e)$ for each edge e .
- A finite set \mathcal{I} that indexes the “commodities”. By convention, $k = |\mathcal{I}|$.
- For each commodity $i \in \mathcal{I}$, a source node $\sigma(i) \in V$ and a set of sinks $\mathbf{T}(i) \subseteq V$.
- For each commodity $i \in \mathcal{I}$, a positive real number d_i which is the demanded communication rate.

This definition, along with all results in this paper, can be generalized to allow for commodities with multiple source nodes, each transmitting an identical message; the details are relatively straightforward and hence omitted. We describe the remaining definitions and results in terms of the k -pairs communication problem, where each commodity only has a single sink node, i.e., $\mathbf{T}(i) = \{\tau(i)\}$. Without loss of generality, one may assume that each source $\sigma(i)$ has a single out-edge $S(i) = (\sigma(i), s(i))$ and no in-edges, and that each sink $\tau(i)$ has a single in-edge $T(i) = (t(i), \tau(i))$ and no out-edges. This convenient transformation is used throughout the paper since it allows one to reason about both sources and communication links by considering only sets of edges. We refer to $S(i)$ as the *source edge* and $T(i)$ as the *sink edge* for commodity i . These edges will be assumed to have infinite capacity. Let \mathcal{S} be the set of all source edges and \mathcal{T} be the set of all sink edges.

We define a network coding solution for the k -pairs communication problem. A natural extension of this definition applies to the general network coding problem.

Definition 2 (Network Coding Solution): Given a network coding instance with an underlying acyclic graph G , a *network coding solution* specifies for each edge e an alphabet

$\Sigma(e)$ and a function $f_e : \prod_{i \in \mathcal{I}} \Sigma(S(i)) \rightarrow \Sigma(e)$ such that three conditions are satisfied. Let $M = (M_1, M_2, \dots, M_k)$ be the k -tuple of symbols transmitted on source edges $(S(1), S(2), \dots, S(k))$.

- If $e \in E$ is not a source-edge, then the function $f_e : \prod_{i \in \mathcal{I}} \Sigma(S(i)) \rightarrow \Sigma(e)$ is computable from the functions on edges in $\text{In}(e)$, which we define to be the set of incoming edges of the tail of e .
- For every sink edge $T(i) \in \mathcal{T}$, $f_{T(i)}(M) = M_i$.
- All but finitely many of the alphabets $\Sigma(e)$ are singleton sets.

We refer to M_i as the *message* for commodity i . For a specified network coding solution and an arbitrary set $A = \{e_a : a \in \mathcal{A}\}$ of edges, it is convenient to introduce the following notation: first $\Sigma(A) = \prod_{e \in A} \Sigma(e)$, second $f_A : \Sigma(S) \rightarrow \Sigma(A)$ is the function where for every k -tuple of source messages M ,

$$f_A(M) = \{f_{e_a}(M) : a \in \mathcal{A}\},$$

and third $\text{In}(A) = \cup_{e \in A} \text{In}(e)$.

Although an instance of the network coding problem specifies a capacity for each edge we do not restrict the size of the alphabet used in a solution. However, when the alphabets are finite, we define the rate of a network coding solution in terms of the capacities of the edges.

Definition 3 (Rate): The *rate* of a network coding solution is the maximum value of r such that there exists a constant b with $\log_b |\Sigma(e)| \leq c(e)$ for each $e \in E$, and $\log_b |\Sigma(S(i))| \geq rd_i$ for each commodity i . For a given network coding problem, the *network coding rate* is defined to be the supremum of the rates of all network coding solutions.

Our definition of rate is a single real number that scales the amount of information sent by each source by the same factor. More generally, we could define for each commodity i a rate r_i and consider the region of feasible rate vectors. The use of a single scalar rate simplifies our notation and effectively poses no restriction since the source rates can be scaled individually using the vector of demands. Scalar definitions of rate are commonly used in the multicommodity flow literature. The quantity $\min_i r_i$ is called the *concurrent flow rate* and the quantity $\sum_i r_i$ is called the *total flow rate*. Thus our definition of network coding rate is comparable to the concurrent multicommodity flow rate. We discuss the extension of our results to multidimensional rate regions in Section XI.

We conclude by briefly relating our definitions to the information theoretic formulation of network coding which is more prevalent in the prior literature. While our formulation is combinatorial in nature and defines notions such as ‘‘capacity’’ and ‘‘rate’’ without explicitly assuming a probability distribution on the messages transmitted, it is possible to relate our definitions to information theory by assigning a probability distribution on the messages transmitted in the network, and examining the entropy of different random variables under this distribution. Such entropy-based arguments will be vital in the second half of this paper.

Suppose for each source edge $S(i)$ we choose a message M_i independently and uniformly at random from $\Sigma(S(i))$.

If we let M be the random variable (M_1, \dots, M_k) , then we can define for each edge e the random variable $Y_e = f_e(M)$. The Shannon entropy of Y_e is denoted $H(Y_e)$. For convenience, we define the following notational shorthands. For a given network code and a given set A of edges, we let $Y_A = (Y_{e_1}, Y_{e_2}, \dots, Y_{e_a})$ where $\{e_1, e_2, \dots, e_a\}$ is the subset of A consisting of edges which have non-singleton alphabets. Additionally, let $H(A)$ denote $H(Y_A)$.

The definition of a network coding solution can be given in information theoretic terms. For example, the requirement that $f_{T(i)}(M) = M_i$ is essentially equivalent to the conditions $H(S(i)) = H(T(i))$ and $H(Y_{T(i)}|Y_{S(i)}) = 0$. Similarly, the condition that the function f_e be computable from information on in-edges to e can be expressed as $H(Y_e|Y_{\text{In}(e)}) = 0$. This last condition implies that $H(S, A) = H(S)$ for any set of edges A . Lastly, we recast our definition of rate in terms of the entropy of the information transmitted over each edge.

Lemma 1: Given a network coding solution defined on a directed acyclic graph G of rate r , there exists a constant b such that the following hold. For all edges $e \in E$,

$$H(e) \leq c(e) \log_2 b$$

and for all commodities $i \in \mathcal{I}$,

$$H(S(i)) \geq rd_i \log_2 b.$$

A key property of entropy that we will frequently use is *submodularity*: for any sets U and V of random variables, $H(U) + H(V) \geq H(U \cup V) + H(U \cap V)$. This is easily proven by the non-negativity of the mutual information $I(U \setminus V; V \setminus U | U \cap V)$. The following generalization of submodularity is also useful, and will be applied in Theorem 27. Let A_1, \dots, A_n be a collection of sets. Define

$$B_{i,n} = \bigcup_{\{\alpha_1, \dots, \alpha_i\} \subseteq \{1, \dots, n\}} A_{\alpha_1} \cap \dots \cap A_{\alpha_i}$$

for $1 \leq i \leq n$. The proof of this theorem is given in the appendix.

Theorem 2 (n-way submodularity):

$$\sum_{i=1}^n H(A_i) \geq \sum_{i=1}^n H(B_{i,n})$$

III. SPARSITY

We consider the k -pairs communication problem in finite directed acyclic graphs. Our goal is to understand cut-based upper bounds on the maximum achievable rate. A natural place to begin is by considering cut conditions which supply upper bounds on the maximum multicommodity flow. The *sparsity* of a graph is well-known to be an upper bound on the maximum multicommodity flow (see, e.g., [19]). However, the network coding rate can exceed the sparsity in directed acyclic graphs, as is shown by the well-known example in Figure 1(a), due to Ahlswede et al. [8]. Understanding why sparsity is not an upper bound on the maximum network coding rate sheds light on the type of cut condition we should be considering. In addition, using sparsity as an upper bound on the maximum multicommodity flow rate, we give a tight characterization of the gap between the maximum network coding rate and the maximum multicommodity flow in terms of the number of vertices, edges, and commodities in the problem instance.

A. Preliminaries

To define sparsity, we must first define cuts. Given a graph $G = (V, E)$ a *cut* is a subset A of the edge set E . (Such a cut is also known as a cut-set or a disconnecting edge set.) We define the capacity of a cut as $C(A) = \sum_{e \in A} c(e)$. Suppose there is a k -pairs communication problem defined on G . If every path from $S(i)$ to $T(i)$ in G intersects A then we say that A *separates* commodity i . Let

$$D(A) = \sum_{i: A \text{ separates } i} d_i.$$

Definition 4 (Sparsity): The sparsity of a cut A and a graph G are respectively defined as

$$\mathcal{S}(A) = \frac{C(A)}{D(A)} \quad \text{and} \quad \mathcal{S}_G = \min_{A \subseteq E} \mathcal{S}(A).$$

The sparsity is an upper bound on the maximum multicommodity flow rate in both directed and undirected graphs (see, e.g., [19]–[22]). Intriguingly, in a directed graph the sparsity of the graph is not an upper bound on the maximum network coding rate [8]. In the remainder of this section we present an example of this phenomenon.

B. A large gap

One measure of the benefit of network coding solutions over flow solutions is the ratio of the network coding rate and the maximum concurrent multicommodity flow rate, which we call the network coding *gap*. We now describe a family of graphs which have a large network coding gap.

The graphs are defined recursively, where $\mathcal{G}(1)$ is the well-known example shown in Figure 1(a). Note that the capacity of the middle edge e in this example bounds the *total* multicommodity flow rate and the *concurrent* network coding rate. More formally, suppose edge e has capacity c and that the rates achievable for the two commodities are f_1, f_2 (for flow) and r_1, r_2 (for network coding). Then we must have $f_1 + f_2 \leq c$ and $\max\{r_1, r_2\} \leq c$, implying a gap of 2.

The approach to obtaining a larger gap is to add more instances of $\mathcal{G}(1)$ that somehow share the capacity of their middle edges. To do this, we construct a larger graph $\mathcal{G}(n)$ by building up from $\mathcal{G}(n-1)$. We add several new instances of $\mathcal{G}(1)$ in which the middle edge e has been removed. The idea is that communication through $\mathcal{G}(n-1)$ will play the role of these missing middle edges. More formally, for each commodity i in $\mathcal{G}(n-1)$, we add a new instance of $\mathcal{G}(1)$ where $s(i)$ is identified with a , $t(i)$ is identified with b but there is no edge from $s(i)$ to $t(i)$. Note that the pair $(s(i), t(i))$ are no longer a commodity in the new instance. This construction is illustrated in Figure 1(b).

We now analyze the total multicommodity flow rate in $\mathcal{G}(n)$. The total multicommodity flow rate through an instance of $\mathcal{G}(1)$ is bounded by the capacity of edge e . In our recursive construction, edge e is omitted but its effective capacity is determined by the communication rate achievable between the corresponding commodities in $\mathcal{G}(n-1)$. By induction, the total communication rate of these commodities is at most 1, and hence the combined effective capacity of the missing middle edges is 1. Thus the total multicommodity flow rate in $\mathcal{G}(n)$

is at most 1. The concurrent multicommodity flow rate is at most $1/2^n$ since there are 2^n commodities.

We now analyze the network coding rate. In $\mathcal{G}(1)$, it is well known that both commodities can communicate at rate 1: all source and edge alphabets are $\Sigma = \{0, 1\}$ and edge e sends the exclusive or (XOR) of the source messages. In order for all commodities in $\mathcal{G}(n)$ to communicate at rate 1, each instance of $\mathcal{G}(1)$ must have effective capacity 1 for its missing middle edge. But by induction this capacity is available in $\mathcal{G}(n-1)$, and hence the network coding rate in $\mathcal{G}(n)$ is 1. One may verify that this network coding solution is obtained when each node simply sends the XOR of its inputs on all of its outputs.

It is easy to verify that the graph $\mathcal{G}(n)$ has 2^n commodities, $\Theta(2^n)$ vertices and $\Theta(2^n)$ edges. Thus we have proven the following theorem.

Theorem 3: For any $k \geq 2$, there exists an instance of the k -pairs communication problem on a directed acyclic graph $G = (V, E)$ with $|V| = \Theta(k)$ and $|E| = \Theta(k)$ for which the network coding gap is $\Theta(k)$.

The graph $\mathcal{G}(n)$ clearly illustrates the distinction between sparsity and network coding rate. Let A be a cut consisting of the single edge e in the instance of $\mathcal{G}(1)$ at the base of the recursive construction. Then $\mathcal{S}(A) = 1/2^n$ since this cut separates all commodities. However, the network coding rate is 1, which greatly exceeds the sparsity.

C. \mathcal{G} is optimal

We now argue that the instance \mathcal{G} is optimal in terms of the gap between the maximum multicommodity flow rate and the network coding rate. More formally, we provide upper bounds on the network coding gap that asymptotically match Theorem 3. In this subsection of the paper, unlike all other sections, we do not assume that each commodity has its own source and sink edges which are not shared by other commodities. This change allows us to provide a stronger bound for instances in which the number of commodities, k , greatly exceeds the number of vertices or edges in the network. The following theorem makes four assumptions: (1) edges have unit capacity, (2) all commodities have unit demands, (3) there exists a path from each source to its sink, and (4) no vertex is both the source and the sink for the same commodity.

Theorem 4: Let $G = (V, E)$ be a directed graph. For any instance of the k -pairs communication problem satisfying the above assumptions, the network coding gap is at most $\min\{|E|, |V|, k\}$.

Proof: First we show that the gap is bounded by $|E|$. A simple pigeon-hole argument shows that the maximum network coding rate is at most $|E|/k$. (Using the more formal tools of Section V, this claim follows since the set E informationally dominates all sources.) Next we consider the maximum concurrent multicommodity flow rate. Since there exists a path from each source to its sink, we can pack these paths fractionally by scaling their flow by $1/k$. Hence the maximum concurrent flow rate is at least $1/k$ and the gap is at most $|E|$.

Now we bound the gap by $|V|$ and k simultaneously. Consider a network coding solution achieving rate r . Let




Fig. 1: (a) The instance $\mathcal{G}(1)$, with the sources and sinks for its two commodities indicated. The edge e has unit capacity and all others have infinite capacity. This instance is due to Ahlswede et al. [8]. (b) The instance $\mathcal{G}(2)$ is constructed from $\mathcal{G}(1)$. The sources and sinks for its four commodities are indicated. All new edges have infinite capacity. (c) The instance $\mathcal{G}(3)$ is constructed from $\mathcal{G}(2)$. The sources and sinks for its eight commodities are indicated.

V_S be the set of vertices that are sources for at least one commodity. For a vertex $v \in V_S$, let $\mathcal{I}_v = \{i : S(i) = v\}$ be the set of commodities with v as the source, and let $T_v = \{T(i) : i \in \mathcal{I}_v\}$ be the set of corresponding sinks. Create a new vertex v' and for each $i \in \mathcal{I}_v$ add an edge $(T(i), v')$ of capacity r , allowing parallel edges. Then our network coding solution provides communication at rate $r \cdot |\mathcal{I}_v|$ between vertices v and v' . Thus the minimum cut separating v and v' must have capacity at least $r \cdot |\mathcal{I}_v|$. Hence, ignoring the commodities not in \mathcal{I}_v , there exists a v - v' flow f_v of value $r \cdot |\mathcal{I}_v|$. Moreover, for every $w \in T_v$, the flow to vertex v' is at most r times the number of commodities in \mathcal{I}_v whose sink vertex is w . Thus the flow from v to w must be exactly $r \cdot |\{i \in \mathcal{I}_v : T(i) = w\}|$. This shows that the flow f_v provides disjoint flows of rate r for all commodities in \mathcal{I}_v . We may fractionally pack the collection of flows $\{f_v : v \in V_S\}$ by reducing their rates by a factor of $1/|V_S|$. This yields a multicommodity flow solution of rate $r/|V_S|$, and hence the gap is at most $|V_S|$. Since $|V_S| \leq \min\{|V|, k\}$, the proof is complete. ■

IV. MEAGERNESS

Sparsity does not necessarily provide an upper bound on the network coding rate in directed acyclic graphs; this is shown by the example of Ahlswede et al. [8] in Figure 1(a). We now define a new combinatorial cut condition called “meagerness” which does provide an upper bound. We also present an example showing that it is not a tight upper bound.

A. Preliminaries

Let $G = (V, E)$ be a finite directed acyclic graph.

Definition 5 (Isolation): Given an edge set $A \subseteq E$ and a subset of commodities $P \subseteq \mathcal{I}$, we say A isolates P if for all $i, j \in P$, every path from $S(i)$ to $T(j)$ intersects A .

A cut A that isolates a set of commodities P must disconnect each source for a commodity in P from the sink for every commodity in P . For a set of commodities P , $d(P) = \sum_{i \in P} d_i$.

Definition 6 (Meagerness): The *meagerness* of a cut A and a graph G are respectively defined as

$$\mathcal{M}(A) = \min_{P: A \text{ isolates } P} \left\{ \frac{C(A)}{d(P)} \right\}$$

$$\mathcal{M}_G = \min_{A \subseteq E} \mathcal{M}(A).$$

To understand isolating cuts and meagerness, it is useful to consider again the graph $\mathcal{G}(1)$ shown in Figure 1(a). The edge e in that example separates the source for each commodity from the corresponding sink. Therefore, if the set P is a single commodity, then the middle edge isolates P . On the other hand, if P is the set of both commodities then the middle edge does not isolate P .

Lemma 5: For any instance of the k -pairs communication problem on a directed graph $G = (V, E)$, the network coding rate is at most \mathcal{M}_G .

This lemma can be proven by a simple counting argument. We omit this proof since this lemma also follows from our stronger results on informational dominance in Section V.




Fig. 2: The split butterfly instance, which has meagerness 1 but network coding rate $2/3$.

B. The Split Butterfly

We now describe an instance of the k -pairs communication problem where the network coding rate is strictly smaller than the meagerness of the graph. The instance, called the *split butterfly*, is depicted in Figure 2. Let $G = (V, E)$ be the graph corresponding to this instance. The set of commodities is $\mathcal{I} = \{a, b, c\}$ and the demand for each commodity is 1. Each edge in G has capacity 1. It is easy to verify that the meagerness of split butterfly is 1. We will show that the network coding rate is $2/3$.

First we show that rate $2/3$ is achievable. Let $\Sigma(S(i)) = \{0, 1\}^2$ for all i and $\Sigma(e) = \{0, 1\}^3$ for all e . The message M_a is transmitted on the only path from $s(a)$ to $t(a)$. Similarly, the message M_b is transmitted on the only path from $s(b)$ to $t(b)$. Each of the bits for message M_c are transmitted on one of the two paths from $s(c)$ to $t(c)$. Therefore, edges e_1 and e_2 must transmit 3 bits each. Every other edge transmits fewer than 3 bits.

Our next step is to prove an upper bound of $2/3$ on the network coding rate for the split butterfly. We use the information theoretic view of network coding described in Section II.

Lemma 6: The network coding rate of the split butterfly instance is $2/3$.

Proof: From the structure of the graph G , the symbol transmitted on edge $T(b)$ must be computable from the information on $S(a)$ and e_1 . Similarly, the symbol transmitted on $T(a)$ must be a function of $S(b)$ and e_2 , and the symbol transmitted on $T(c)$ must be a function of e_1 and e_2 . These observations imply the following entropy equalities.

$$H(S(a), e_1) = H(S(a), T(b), e_1) \quad (1)$$

$$H(S(b), e_2) = H(T(a), S(b), e_2) \quad (2)$$

$$H(e_1, e_2) = H(T(c), e_1, e_2) \quad (3)$$

$$H(S(a), e_1) + H(S(b), e_2) \quad (4)$$

$$= H(S(a), S(b), e_1) + H(S(a), S(b), e_2)$$

$$H(S(a)) + H(S(b)) + H(e_1) + H(e_2) \quad (5)$$

$$\geq H(S(a), S(b), e_1, e_2) + H(S(a), S(b))$$

$$H(e_1) + H(e_2) \geq H(S(a), S(b), e_1, e_2) \quad (6)$$

$$H(S(a), S(b), e_1, e_2) = H(S(a), S(b), S(c), e_1, e_2) \quad (7)$$

$$H(e_1) + H(e_2) \geq H(S(a), S(b), S(c), e_1, e_2) \quad (8)$$

$$H(e_1) + H(e_2) \geq H(S(a)) + H(S(b)) + H(S(c)) \quad (9)$$

$$(c(e_1) + c(e_2)) \geq (d_a + d_b + d_c)r \quad (10)$$

(1)-(3) follow from the discussion of G . (4) follows from (1) and (2) and from the fact that $Y_{T(a)} = Y_{S(a)}$ and $Y_{T(b)} = Y_{S(b)}$. (5) follows from (4) by submodularity. (6) follows from (5) because the sources are independent. (7) follows from (3). (8) follows from (6) and (7). (9) follows from (8) since entropy is non-decreasing and the sources are independent. (10) follows from (9) and Lemma 1, assuming rate r is achievable. Thus we obtain $2/3 \geq r$ since all edges have capacity 1 and all commodities have demand 1. ■

This proof shows that the upper bound on rate given by the meagerness of the graph is sometimes loose.

C. Information Theoretic Constraints

Examining the proof of Lemma 6 carefully, it uses certain characteristics of entropy, such as submodularity, and certain inequalities implied by the structure of the graph. These are the key techniques that are used throughout the rest of the paper. Let us now formalize the latter type of inequalities. If the graph structure implies that the information transmitted on a set of edges B is computable from the information transmitted on a set of edges A , we say A *informationally dominates* B . In the preceding proof, we derived informational dominance relationships by ad hoc arguments. The next section gives a combinatorial approach for deriving such inequalities in arbitrary DAGs.

By combining all informational dominance constraints with certain properties of entropy, one can obtain an outer bound on the capacity of arbitrary DAGs. The details of this bound are deferred to Section VIII, where we give an outer bound that also holds for cyclic and undirected graphs. Before this general outer bound can be presented, we must first extend our basic definitions and informational dominance to general graphs. This is done in sections VI and VII.

V. INFORMATIONAL DOMINANCE

In this section, we formally define informational dominance and present a condition based on the structure of the graph that entirely characterizes the informational dominance relationship. This result provides a complete answer to the natural question “When does the information on one set of edges necessarily determine the information on another set of edges?”, in the case when every k -tuple of messages has a positive probability of being transmitted by the k sources. This characterization also leads to a polynomial time algorithm for finding the set of all edges informationally dominated by a given set.

A. Preliminaries

Assume we are given a k -pairs communication problem defined on a directed acyclic graph G . (Recall that G may be countably infinite, but the terms “path” and “reachable” refer to paths of finite length.) Without loss of generality, we assume that every edge is reachable in G from some source and that for every commodity i there is a path from $S(i)$ to $T(i)$ in G .

Definition 7 (Informational Dominance): An edge set A *informationally dominates* edge set B if for all network coding solutions and k -tuples of messages x and y , $f_A(x) = f_A(y)$ implies $f_B(x) = f_B(y)$.

Definition 8 (Dom(A)): For an edge set A ,

$$\text{Dom}(A) = \{ e : A \text{ informationally dominates } e \}.$$

It is clear from the definition that the informational dominance relation is a pre-order (i.e., it is reflexive and transitive). However, for finite G , it is not clear if determining whether a set of edges A informationally dominates a set of edges B is even recursively decidable³. The difficulty arises because the informational dominance relationship must hold for *all* network coding solutions, of which there are infinitely many. In Section V-C we present a characterization of the informational dominance relationship that is entirely based on the structure of the graph. The characterization is somewhat cumbersome, so Section V-B formulates some simple combinatorial conditions which are sufficient but not necessary for informational dominance. These simple conditions are much easier to work with than the full necessary and sufficient condition developed in Section V-C, and they play a role in the outer bound proofs for specific networks considered in this paper.

B. Simple sufficient conditions

Definition 9 (Downstream): We say that an edge set B is *downstream* of edge set A if every path from S to B in G must intersect A . We denote the relation “ B is downstream of A ” using the notation $A \rightsquigarrow B$.

The following lemma relates downstreamness to informational dominance. Its proof follows from the more general Theorem 10.

Lemma 7: If $A \rightsquigarrow B$ then $B \subseteq \text{Dom}(A)$.

For example, the proof of Lemma 6 relies on the following three downstream relationships:

$$\{S(a), e_1\} \rightsquigarrow \{S(a), T(b), e_1\}$$

$$\{S(b), e_2\} \rightsquigarrow \{S(b), T(a), e_2\}$$

$$\{e_1, e_2\} \rightsquigarrow \{T(c), e_1, e_2\}.$$

A useful special case of downstreamness is the following relation identified by Jain et al. [13].

Input-Output Inequality. The edges entering a vertex set U informationally dominate all edges adjacent to U .

³Informally, “recursively decidable” means that the problem can be solved by a Turing machine which can run for an arbitrarily large amount of time but must necessarily terminate. For a formal definition, see Sipser [23].




Fig. 3: An example of an indirect walk from $s(i)$ to $t(i)$.

More precisely, the input-output inequality can be expressed by the following corollary of Lemma 7.

Corollary 8 (Jain et al. [13]): Given a vertex set $U \subseteq V(G)$, let $E_{\text{in}}(U)$ denote the set of edges from $V(G) \setminus U$ to U , let $E_{\text{adj}}(U)$ denote the set of edges adjacent to U , and let $\mathcal{S}(U)$ denote the set $\mathcal{S} \cap E_{\text{adj}}(U)$. Then $E_{\text{adj}}(U) \subseteq \text{Dom}(E_{\text{in}}(U) \cup \mathcal{S}(U))$.

Another useful sufficient condition for informational dominance involves the weakly connected components of the graph $G \setminus A$. Recall that two vertices u, v of a directed graph with edge set E are said to belong to the same weakly connected component if there exists a path from u to v which ignores the directions of the edges.

Lemma 9: If $\sigma(i)$ and $\tau(i)$ are in different weakly connected components of $G \setminus A$ then $\{S(i), T(i)\}$ is informationally dominated by A .

Again, the proof is omitted because the lemma follows immediately from Theorem 10. A similar lemma, called the ‘‘crypto inequality,’’ was identified by Jain et al. [13]; we will discuss this lemma in the context of undirected graphs in Section VII.

C. Combinatorial Characterization of $\text{Dom}(A)$

Recall that for an edge set A , $\text{Dom}(A)$ is defined to be the set of all edges informationally dominated by A . The difficult part of characterizing $\text{Dom}(A)$ is determining when the source and sink edges for a commodity are informationally dominated by A . Roughly, we need to determine when the edge $T(i)$ can determine any information about the message M_i without using information on edges in A . The notion of an *indirect walk* is key to making this determination. An example of an indirect walk is shown in Figure 3.

Definition 10 (Indirect walk): Let H be a graph and let $\mathcal{S}(H)$ be the set of source edges. For a source edge $S(i)$ and an edge e , an *indirect walk* from $S(i)$ to e is a sequence of (finite length) directed paths $Q_1, P_1, Q_2, P_2, \dots, Q_j$ in H where

- The first node in Q_1 is the head of $S(i)$ and the last node in Q_j is the tail of e .
- The last node in Q_ℓ is the same as the last node in P_ℓ for all $\ell < j$.
- The first node in each path is the head of a source edge in $\mathcal{S}(H)$.

An indirect walk from $S(i)$ to $T(i)$ is called an *indirect walk for commodity i* .

We may assume that each path Q_i has a distinct source as its initial vertex, and hence the number of paths in an indirect walk is indeed finite. Note that the paths P_j essentially traverse backwards from the tail of Q_i to a source. Finally, we need to define a subgraph $G(A, i)$ of G . We will show that if there is no indirect walk for commodity i in the subgraph $G(A, i)$, then $S(i)$ and $T(i)$ are informationally dominated by A . An example of the graph $G(A, i)$ is shown in Figure 4 (b).

Definition 11 ($G(A, i)$): Given a graph G , an edge set A and a commodity i , let $G(A, i)$ be the graph obtained from G by the following process:

- Remove any edge or vertex that does not have a path to $T(i)$ in G .
- Remove all edges of A .
- Remove any edge or vertex that is not reachable from a source edge in the remaining graph.

The remainder of this section is devoted to proving the following combinatorial characterization of $\text{Dom}(A)$.

Theorem 10: For an edge set A , the set $\text{Dom}(A)$ satisfies the following conditions.

Condition 1: $A \subseteq \text{Dom}(A)$.

Condition 2: $S(i) \in \text{Dom}(A)$ if and only if $T(i) \in \text{Dom}(A)$.

Condition 3: Every edge in $E \setminus \text{Dom}(A)$ is reachable in $G \setminus \text{Dom}(A)$ from a source.

Condition 4: For every source edge $S(i)$ in $G \setminus \text{Dom}(A)$, there is an indirect walk for commodity i in $G(\text{Dom}(A), i)$.

Furthermore, any set B satisfying these conditions contains $\text{Dom}(A)$.

Note that Condition 3 is a statement about downstreamness: it asserts that for any edge e , if $A \rightsquigarrow \{e\}$ then $e \in \text{Dom}(A)$. We will prove the two directions of Theorem 10 separately. Section V-D proves that $\text{Dom}(A)$ necessarily satisfies these conditions. Section V-E proves these conditions are sufficient: $\text{Dom}(A)$ is the unique minimal superset of A meeting these conditions.

D. Necessary Conditions

Condition 1 is trivially satisfied by reflexivity of informational dominance. Condition 2 follows from the definition of a network coding solution, which requires that the messages on sink edges equal the messages on the corresponding source edges. For Conditions 3–4, we will show that an edge set B that violates one of these two conditions informationally dominates an edge in $E \setminus B$. This fact then implies that $\text{Dom}(A)$ must satisfy these conditions. To see this, suppose otherwise. Then there is an edge e in $E \setminus \text{Dom}(A)$ that is informationally dominated by $\text{Dom}(A)$. By transitivity, e is also informationally dominated by A . Thus $e \in \text{Dom}(A)$, contradicting the choice of e .

1) *Condition 3 and Downstreamness:* The third condition says that every edge in $E \setminus \text{Dom}(A)$ is reachable in $G \setminus \text{Dom}(A)$ by some source edge. Let us consider this situation in the context of a general edge set B . The following lemma proves that downstreamness is a sufficient condition for informational dominance.

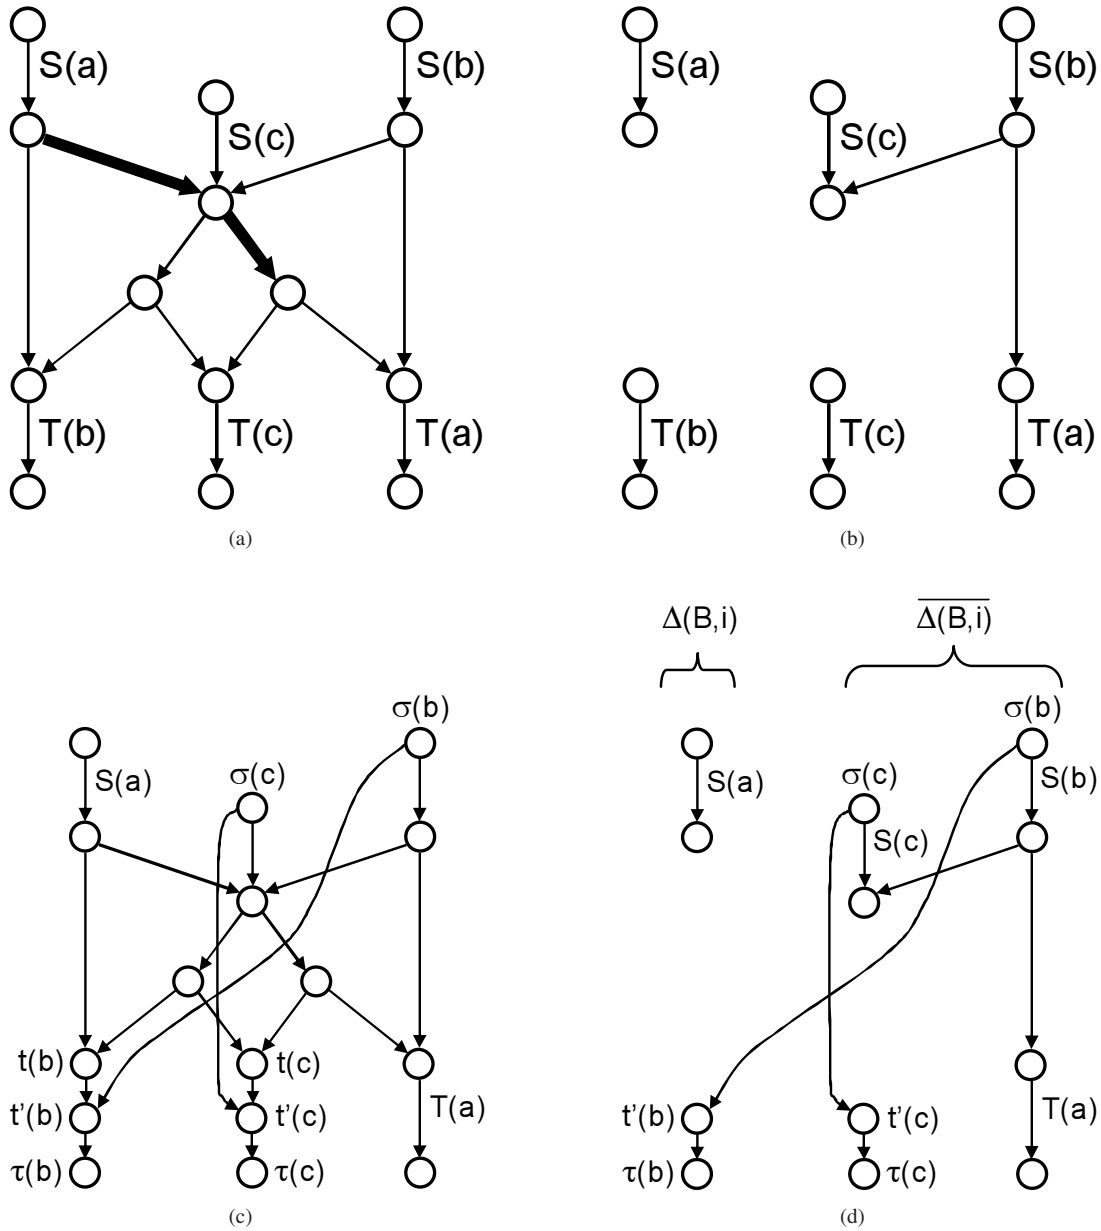


Fig. 4: (a) The split butterfly example is depicted here as graph G . The bold edges denote a cut B . (b) The graph $G(B, i)$, where i is chosen to be commodity a . Note that an edge that is inbound to $t(a)$ is deleted because it is not reachable from a source edge. (c) The graph G^i introduces “shortcut edges” for all commodities except commodity i . (d) The graph $G^i(B, i)$ is shown, along with the partition of S into $\Delta(B, i)$ and $\overline{\Delta(B, i)}$.

Lemma 11: For any network code and for any sets B and B' of edges, if B' is downstream of B , then there exists a function $h : \Sigma(B) \rightarrow \Sigma(B')$ such that $h \circ f_B = f_{B'}$.

Proof: Let D be the set of all edges that are downstream of B and have non-singleton alphabets. Note that $B \subseteq D$ and $B' \subseteq D$. Order the edges of D in topological order so that for any edge $e \in D$, every edge in $\text{In}(e)$ is either not in D or comes before e in the ordering. Let the ordering of D be $(e_1, e_2, \dots, e_{|D|})$.

We prove by induction on k that $f_{e_k} = h_{e_k} \circ f_B$ for some function h_{e_k} . Consider an arbitrary edge e_k and let $D_{k-1} = \{e_1, \dots, e_{k-1}\}$. Let e be an arbitrary edge in $\text{In}(e_k)$. We claim

that there is a function h_e such that $f_e = h_e \circ f_B$. If $e \in B$ then the claim is trivial. If $e \in D_{k-1}$ then the claim holds by induction. If $e \notin B \cup D_{k-1}$ then, by definition of D and choice of the ordering, there is a path from a source to e in $G \setminus B$. This contradicts the assumption that $e_k \in D$ and hence the claim is proven.

By the definition of a network coding solution, there exists a function $g_{e_k} : \Sigma(\text{In}(e_k)) \rightarrow \Sigma(e_k)$ which maps the functions on edges in $\text{In}(e_k)$ to f_{e_k} in the sense that $f_{e_k} = g_{e_k} \circ f_{\text{In}(e_k)}$. Combining g_{e_k} with the functions h_e for $e \in \text{In}(e_k)$ gives the desired function h_{e_k} . ■

Lemma 11 implies that a set of edges B' that is downstream of B is informationally dominated by B . Therefore, $\text{Dom}(A)$ must satisfy Condition 3.

2) *Condition 4*: We now consider the last condition. Without loss of generality we assume that Conditions 2 and 3 hold for a set of edges B but there exists a commodity i with no indirect walk from $S(i)$ to $T(i)$ in $G(B, i)$. We prove that $T(i)$ is informationally dominated by B .

We start by simplifying the instance G by adding “shortcut edges” for all commodities except commodity i . This will allow us to focus on communication from $S(i)$ to $T(i)$ and ignore whether other commodities are communicated correctly. We construct the simplified graph G^i from G by the following process, which is illustrated in Figure 4.

- For each commodity $j \neq i$, add a node $t'(j)$ and replace sink edge $T(j) = (t(j), \tau(j))$ with two edges $T'(j) = (t(j), t'(j))$ and $T(j) = (t'(j), \tau(j))$. The latter edge is the new sink edge for commodity j .
- For each commodity $j \neq i$, add an edge $\epsilon_j = (\sigma_j, t'(j))$.

The graph $G^i(B, i)$ is obtained from G^i as in Definition 11. Notice that an edge ϵ_j cannot be contained in an indirect walk from $S(i)$ to $T(i)$ in $G^i(B, i)$ because the edges $(t(j), t'(j))$ have no path to $T(i)$ and are therefore removed. Thus, if there does not exist an indirect walk for commodity i in $G(B, i)$, then there does not exist an indirect walk for commodity i in $G^i(B, i)$.

Given a network coding solution N for G , the network coding solution N^i is defined as follows.

- Source σ_i sends M_j on edge ϵ_j for all $j \neq i$. That is, $\Sigma^i(\epsilon_j) = \Sigma(S(j))$ and $f_{\epsilon_j}^i(M) = M_j$.
- Let e be an edge in G^i such that e is not a source or sink edge and $e \neq \epsilon_j$ for all j . If there exists a path from e to $T(i)$, then $\Sigma^i(e) = \Sigma(e)$ and $f_e^i = f_e$. Otherwise, $\Sigma(e) = \{0\}$ and $f_e^i(M) = 0$.

Again, by construction, if N is a network coding solution for the original instance then N^i is a network coding solution for the instance defined on G^i . The following lemma proves another useful property of N^i .

Lemma 12: Given two k -tuples of messages x and y and a set of edges B , if $f_B(x) = f_B(y)$ under network coding solution N , then $f_B^i(x) = f_B^i(y)$ under network coding solution N^i .

Proof: An edge $e \in B$ transmits the same function under N and N^i if there is a path from e to $T(i)$. Therefore $f_e^i(x) = f_e^i(y)$. If there is no path from e to $T(i)$, then e always transmits 0 under N^i . Therefore, $f_B^i(x) = f_B^i(y)$. ■

We are now ready to prove that Condition 4 is necessary.

Lemma 13: If there does not exist an indirect walk for commodity i in $G(B, i)$, then B informationally dominates $T(i)$.

Proof: We argue by contradiction. Suppose $T(i)$, and therefore $S(i)$, is not informationally dominated by B . Then there exists a network coding solution N and two k -tuples of messages x and y such that $x_i \neq y_i$ but $f_B(x) = f_B(y)$. Let G^i and N^i be defined as above. By Lemma 12, this implies that in graph G^i under network coding solution N^i , $f_B^i(x) = f_B^i(y)$ and $x_i \neq y_i$. By assumption, there is no

indirect walk from $S(i)$ to $T(i)$ in $G(B, i)$ and hence, as observed previously, there is none in $G^i(B, i)$ either.

We partition the set of source edges into two sets. Source edge $S(j) \in \Delta(B, i)$ if there is no indirect walk from $S(j)$ to $T(i)$ in $G^i(B, i)$. Let $\overline{\Delta(B, i)} = S \setminus \Delta(B, i)$. Let \underline{z} be a k -tuple of messages such that $z_j = y_j$ for all $S(j) \in \overline{\Delta(B, i)}$ and $z_j = x_j$ otherwise. We prove three claims which lead to a contradiction, as explained in the final paragraph of the proof.

Claim 1: $z_i \neq y_i$. By assumption, there is no indirect walk from $S(i)$ to $T(i)$, so $S(i) \in \Delta(B, i)$. Therefore $z_i = x_i$ and by assumption, $x_i \neq y_i$. This implies that $z_i \neq y_i$.

Claim 2: $f_B^i(z) = f_B^i(y)$. Of the edges in B , only finitely many have alphabets which are not singletons. Sort these edges topologically and label them $\{e_1, e_2, \dots, e_b\}$. For every edge $e \in B \setminus \{e_1, e_2, \dots, e_b\}$, the equation $f_e^i(z) = f_e^i(y)$ holds trivially since $|\Sigma(e)| = 1$. Now suppose that there is an edge $e_j \in \{e_1, e_2, \dots, e_b\}$ such that $f_{e_j}^i(z) \neq f_{e_j}^i(y)$; let e_j be the first such edge. There are four cases to consider:

- 1) There is no path from e_j to $T(i)$ in G .
- 2) Every path from a source to the tail of e_j intersects B .
- 3) There is a path in $G^i(B, i)$ from a source in $\overline{\Delta(B, i)}$ to the tail of e_j .
- 4) There is a path in $G^i(B, i)$ from a source in $\Delta(B, i)$ to the tail of e_j .

For the first case, if there is no path from e_j to $T(i)$ in G , then $f_{e_j}^i(z) = 0 = f_{e_j}^i(y)$ which contradicts the assumption that e_j transmits a different message under the two k -tuples of messages. If we are in the second case, then e_j is downstream of the set B_j of edges in B that precede e_j in the ordering. Therefore, e_j can only transmit two different symbols if some edge that precedes it in the ordering also transmits two different symbols. This contradicts the assumption that e_j was the first such edge.

Therefore we are in case three or four. It is crucial to note that these two cases are mutually exclusive since, if we were in both cases, then there would exist an indirect walk from a source in $\Delta(B, i)$ to $T(i)$, contradicting the definition of $\Delta(B, i)$. Thus, if we are in case three then $\overline{\Delta(B, i)} \cup B_j \rightsquigarrow e_j$. By choice of e_j , every edge in B_j transmits the same symbol under the two k -tuples of messages z and y . In addition, every $S(j) \in \overline{\Delta(B, i)}$ has $z_j = y_j$ by the definition of z . Therefore, $f_{e_j}^i$ is determined by values which are identical under message vectors z and y , contradicting our assumption that $f_{e_j}^i(z) \neq f_{e_j}^i(y)$. Hence we cannot be in case three.

Finally, suppose that we are in case four. We will argue that $f_{e_j}^i(z) = f_{e_j}^i(x)$. In case four edge e_j is downstream of B_j and the sources $\Delta(B, i)$. Let \mathcal{E} be this set of edges and let h be the function mapping symbols on the edges of \mathcal{E} to the symbol transmitted on e_j . Consider an edge $e \in \mathcal{E}$. If $e \in B_j$, then $f_e^i(z) = f_e^i(y) = f_e^i(x)$, by our choice of e_j and choice of x and y . Otherwise, $e = S(j)$ for some $S(j) \in \Delta(B, i)$. By the definition of z , $z_j = x_j$. Therefore, for every edge $e \in \mathcal{E}$, $f_e^i(z) = f_e^i(x)$. This implies that $h(f_{\mathcal{E}}^i(z)) = h(f_{\mathcal{E}}^i(x))$. Since $e_j \in B$, we assumed that it transmits the same symbol under the two k -tuples of messages x and y . Therefore

$$h(f_{\mathcal{E}}^i(y)) = h(f_{\mathcal{E}}^i(x)) = h(f_{\mathcal{E}}^i(z))$$

This contradicts our assumption that $f_{e_j}^i(z) \neq f_{e_j}^i(y)$. Therefore, $f_B^i(z) = f_B^i(y)$.

Claim 3: $f_{T(i)}^i(z) = f_{T(i)}^i(y)$. By the definition of an indirect walk and the set $\Delta(B, i)$, the edge $T(i)$ is downstream of B and the source edges in $\Delta(B, i)$. We have shown that the set of edges B transmit the same symbols under the two k -tuples of messages z and y . By construction, z and y assign the same message to any commodity j with $S(j) \in \Delta(B, i)$. Therefore, the tail of $T(i)$ receives the same information under the two k -tuples of messages.

The three preceding claims show that $T(i)$ transmits the same symbol under z and y but the message for commodity i is not the same, contradicting the fact that N^i is a network coding solution. This implies that N is not a network coding solution and therefore the set B informationally dominates $S(i)$. ■

E. Sufficiency of Conditions

We have shown that $\text{Dom}(A)$ must satisfy all four conditions of Theorem 10. We now show that if all four conditions are satisfied by a set of edges B , then B does not informationally dominate any edge in $E \setminus B$.

Lemma 14: If all four conditions of Theorem 10 are satisfied by a set of edges B , then B does not informationally dominate any edge in $E \setminus B$.

Proof: We need to show two things: that B does not informationally dominate any source or sink edge in $G \setminus B$ and that B does not informationally dominate an edge e which is not a source or sink edge. We handle these two cases separately.

For the first case, by Conditions 2 and 4, an edge $S(i)$ is in $G \setminus B$ if and only if $T(i)$ is also in $G \setminus B$ and there exists an indirect walk from $S(i)$ to $T(i)$ in $G(B, i)$. Let $Q_1, P_1, \dots, Q_{j-1}, P_{j-1}, Q_j$ be the indirect walk from $S(i)$ to $T(i)$ in $G(B, i)$. For each path segment Q_ℓ let v_ℓ be the last node in Q_ℓ . By the definition of $G(B, i)$, for $1 \leq \ell < j$ there exists a path R_ℓ from v_ℓ to the tail of $T(i)$ in G .

Relabel the source edges as follows: for all $1 \leq \ell \leq j$, $S(\ell)$ now refers to the edge whose head is the first node in Q_ℓ . (Thus, $S(i)$ is now called $S(1)$.) We now sketch the portion of our network coding solution that transmits M_1 from $S(1)$ to $T(1)$ as follows. (The alphabets will be vector spaces over the field $\{0, 1\}$, and the symbol \oplus should be interpreted as the addition operation in one of these vector spaces.)

- For $1 \leq \ell \leq j$, the path Q_ℓ is used to send M_ℓ from $S(\ell)$ to v_ℓ .
- For $1 \leq \ell < j$, the path P_ℓ is used to send $M_{\ell+1}$ from $S(\ell+1)$ to v_ℓ .
- For $1 \leq \ell < j$, the path R_ℓ is used to send $M_\ell \oplus M_{\ell+1}$ to $T(1)$.

We have not yet specified the alphabets to be used by this network coding solution. Effectively, all source and edge alphabets are set to $\{0, 1\}$. However, if any of the Q , P or R paths intersect then we must accommodate all overlapping transmissions by increasing the alphabets of the edges accordingly. So for an edge e which appears in p of the paths $\{P_\ell, Q_\ell, R_\ell : 1 \leq \ell \leq j\}$, we put $\Sigma(e) = \{0, 1\}^p$. This

increase in alphabet size reduces the rate of the solution but the rate is irrelevant for this proof.

We now explain how the portion of the network coding solution described above transmits M_1 from $S(1)$ to $T(1)$. For $1 \leq \ell < j$, path R_ℓ sends $M_\ell \oplus M_{\ell+1}$ on path R_ℓ to the tail of $T(1)$. In addition, path Q_j sends M_j from $S(j)$ to the tail of $T(1)$. Therefore, the tail of $T(1)$ receives j linearly independent combinations of the j messages M_1, M_2, \dots, M_j . The message M_1 can be computed from the values of these j symbols and therefore M_1 can be transmitted on edge $T(1)$.

To show that B does not informationally dominate any source or sink edge in $G \setminus B$, we must exhibit two k -tuples of messages x and y , such that $f_B(x) = f_B(y)$ and $x_1 \neq y_1$. The messages x and y are constructed as follows. For $1 \leq \ell \leq j$, let $x_\ell = 0$ and $y_\ell = 1$. For $\ell > j$, let $x_\ell = y_\ell = 0$. By construction, $x_1 \neq y_1$ but for all $1 \leq \ell < j$, $x_\ell \oplus x_{\ell+1} = y_\ell \oplus y_{\ell+1}$. Therefore, for $1 \leq \ell < j$, every edge in the path in R_ℓ transmits the same symbol for the two k -tuples of messages. Since these are the only paths that might intersect B on which any information is transmitted, $f_B(x) = f_B(y)$.

Repeating the above argument for all $S(i)$ in $G \setminus B$ in turn proves that every source and sink in $G \setminus B$ is not informationally dominated by B .

Now consider an edge e in $G \setminus B$ that is not a source or sink edge. By condition 3, edge e is reachable from some source $S(j)$ in $G \setminus B$. Since $S(j)$ is not informationally dominated by B , there are two k -tuples of messages x and y such that $f_B(x) = f_B(y)$ but $x_j \neq y_j$. Given this network coding solution, we augment it by also sending M_j from $S(j)$ to e along the path in $G \setminus B$. Edge e then transmits M_j in addition to whatever information it was already transmitting. Since $x_j \neq y_j$ but $f_B(x) = f_B(y)$, edge e is not informationally dominated by B . ■

By Lemma 14, any superset B of A satisfying the four conditions of Theorem 10 does not informationally dominate any edges in $E \setminus B$. Since $\text{Dom}(A)$ is defined to be the set of edges informationally dominated by A , any such set B must contain $\text{Dom}(A)$.

Corollary 15: Any set of edges B that satisfies the four conditions of Theorem 10 contains $\text{Dom}(A)$.

Lemmas 11, 13, 14 and Corollary 15 prove Theorem 10

F. Efficiently Computing $\text{Dom}(A)$

Let $G = (V, E)$ be finite. If we know $\text{Dom}(A)$, then we can easily check if A informationally dominates a set B of edges by checking if $B \subseteq \text{Dom}(A)$. Therefore, we turn our attention to efficiently finding $\text{Dom}(A)$.

Initially we let $\tilde{A} = A$ and check if any of the conditions of Theorem 10 are violated by \tilde{A} . If so, then there exists an edge in $E \setminus \tilde{A}$ that is informationally dominated by \tilde{A} . If such an edge exists, then we add it to \tilde{A} and continue. Our procedure terminates when there are no edges in $E \setminus \tilde{A}$ that \tilde{A} informationally dominates. A straightforward induction argument shows that the final set is $\text{Dom}(A)$.

At each iteration of the algorithm, we either find a new edge that is informationally dominated by the set \tilde{A} or we stop. Finding a new edge that is informationally dominated

by \tilde{A} can be done by checking each of the four conditions of Theorem 10. Since we begin with $\tilde{A} = A$ and only add edges, Condition 1 will always be satisfied. Conditions 2 and 3 are straightforward to check in time linear in $|G|$. We show how to check Condition 4 efficiently.

Lemma 16: Given a finite graph G , an edge set \tilde{A} and a commodity i , there exists an indirect walk for commodity i in $G(\tilde{A}, i)$ if and only if $S(i)$ and $T(i)$ are in the same weakly connected component of $G(\tilde{A}, i)$.

Proof: Let $U(\tilde{A}, i)$ be the graph obtained from $G(\tilde{A}, i)$ by undirecting all edges. The edges $S(i)$ and $T(i)$ are in the same weakly connected component of $G(\tilde{A}, i)$ if and only if there is a path from $S(i)$ to $T(i)$ in $U(\tilde{A}, i)$.

Suppose there exists an indirect walk for commodity i in $G(\tilde{A}, i)$. If we undirect the edges in this walk and remove repeated edges, we get a path from $S(i)$ to $T(i)$ in $U(\tilde{A}, i)$.

Now suppose there is a path $p(A, i)$ from $S(i)$ to $T(i)$ in $U(\tilde{A}, i)$. We construct an indirect walk for commodity i in $G(\tilde{A}, i)$ from $p(A, i)$. Let p be the set of edges and nodes in $G(\tilde{A}, i)$ corresponding to edges and nodes of $p(A, i)$. The set p can be decomposed into directed subpaths $q_1, p_1, q_2, p_2, \dots, q_j$ where:

- The first node in q_1 is the head of $S(i)$.
- The last node in q_j is the tail of $T(i)$.
- The last node, v_ℓ , in q_ℓ is the same as the last node in p_ℓ for $1 \leq \ell < j$.

In order to turn p into an indirect walk $P = Q_1, P_1, Q_2, P_2, \dots, P_j$, we need to extend the subpaths so that they all start at sources. Since the first node in q_1 is $S(i)$, we let $Q_1 = q_1$. For $1 \leq \ell < j$, let u_ℓ be the first node in p_ℓ . Since u_ℓ corresponds to a node in $G(\tilde{A}, i)$ there exists a source, $S(\ell)$, with a path r_ℓ to u_ℓ in $G(\tilde{A}, i)$. Let P_ℓ be the path obtained by first following r_ℓ and then following p_ℓ . Note that $q_{\ell+1}$ starts at the same node as p_ℓ . Let $Q_{\ell+1}$ be the path consisting of r_ℓ followed by $q_{\ell+1}$. Defining P_ℓ and $Q_{\ell+1}$ in this way for $1 \leq \ell < j$ results in an indirect walk for commodity i in $G(\tilde{A}, i)$. ■

The graphs $G(\tilde{A}, i)$ and $U(\tilde{A}, i)$ can be constructed in time linear in $|G|$. Furthermore, finding a path from $S(i)$ to $T(i)$ in $U(\tilde{A}, i)$ can also be done in time linear in $|G|$. This implies the following corollary.

Corollary 17: Given a directed acyclic graph G and a set A of edges, the set $\text{Dom}(A)$ can be found in time $O(|G|^3)$.

This corollary also shows that we can efficiently test whether a set of edges A informationally dominates a set of edges B . Since B is informationally dominated by A if and only if $B \subseteq \text{Dom}(A)$, the corollary implies that we can test this informational dominance relationship in time $O(|G|^3)$.

VI. DEFINITIONS FOR GENERAL GRAPHS

In this section we extend the definitions of Section II to finite undirected graphs and finite graphs with cycles. Given a k -pairs communication problem on an undirected graph $G = (V, E)$, we first define a corresponding directed graph $\hat{G} = (\hat{V}, \hat{E})$. Each undirected edge $\{u, v\} \in E$ is represented in \hat{G} as two oppositely directed edges (u, v) and (v, u) . For directed cyclic graphs, we simply set $\hat{G} = G$. As was the case with

DAGs, one may assume that each source $\sigma(i)$ in \hat{G} has a single out-edge $S(i) = (\sigma(i), s(i))$ and no in-edges, and that each sink $\tau(i)$ has a single in-edge $T(i) = (t(i), \tau(i))$ and no out-edges. Next, we must define a network code for the graph \hat{G} . The definition of Section II cannot immediately be used since it applies only to directed acyclic graphs. The following subsection addresses this issue.

A. Time-Expanded Graph

When defining a network coding solution in a graph with cycles, it is important to ensure that the solution has an operational meaning. Intuitively, we must ensure a solution could actually be implemented in the network over a sequence of time steps. This ensures that there exists an appropriate causality relationship between the random variables on the network edges. This notion is made precise below by specifying a leveled directed acyclic graph G^* representing the flow of information in G over a sequence of time steps, and requiring that a network coding solution in G should come from a network coding solution in G^* . As mentioned earlier, time-expanded graphs have been used in prior work (see, e.g., [8], [15], [16]) for the same purpose of modeling information flow in a graph with cycles, although our construction differs in some technical details such as having an infinite number of levels. A similar approach was advocated in [5], but that paper proposed a model of network coding based on a slightly different definition of G^* where nodes are memoryless. Here, as in [8], [15], [16], we advocate a model which implicitly treats each node of G as having infinite memory.

Definition 12 (Time-Expanded Graph G^):* Given an instance of the network coding problem in a directed graph G , the *time-expanded graph* $G^* = (V^*, E^*)$ is a directed acyclic graph obtained using the following transformation.

- The vertex set V^* includes the set $V \times \mathbb{Z}$. The vertex of V^* represented by the ordered pair (v, t) will be denoted by v_t .
- For each edge $e = (u, v) \in E$ there are edges $e_t = (u_{t-1}, v_t)$ in E^* for $t \in \mathbb{Z}$.
- For each $v \in V$ and each $t \in \mathbb{Z}$, there is an edge $(v_{t-1}, v_t) \in E^*$. These edges are referred to as *memory edges*.
- For each commodity $i \in \mathcal{I}$, we add two new vertices $\sigma(i), \sigma'(i)$ connected by a directed edge $S(i)$, and we add directed edges $S_t(i) = (\sigma'(i), s(i)_t)$ to E^* , for each $t \in \mathbb{Z}$.
- Similarly, for each $i \in \mathcal{I}$, we add two new vertices $\tau(i), \tau'(i)$ connected by a directed edge $T(i)$, and we add directed edges $T_t(i) = (t(i)_t, \tau(i))$ to E^* for each $t \in \mathbb{Z}$.

Given an edge set $A \subseteq \hat{E} = E(\hat{G})$, the time-expansion $A^* \subseteq E^*$ is defined as follows.

- For each source edge $S(i) \in A$, A^* contains the edges $S(i), \{S_t(i) : t \in \mathbb{Z}\}$ of G^* .
- For each sink edge $T(i) \in A$, A^* contains the edges $T(i), \{T_t(i) : t \in \mathbb{Z}\}$ of G^* .
- For every other edge $e \in A$, A^* contains the edges $\{e_t : t \in \mathbb{Z}\}$.

- For every vertex $v \in \hat{V} = V(\hat{G})$ such that A contains every incoming edge of v , A^* contains the memory edges $\{(v_t, v_{t+1}) : t \in \mathbb{Z}\}$.

If a network coding instance is defined on an undirected graph G , then the graph G^* is obtained from G by first replacing each undirected edge of G by two oppositely directed edges and then performing the above transformation. Since G^* is a directed acyclic graph, the definition of a network code and network coding solution from Section II are still applicable. A useful property of network coding solutions in G^* is that they may be adjusted by “time translations.” In other words, given a network coding solution N and an integer k , we may obtain another network coding solution $N_{(k)}$ by using the same edge alphabets and coding functions, but reassigning them so that $\Sigma(e_t)$ and f_{e_t} are assigned to edge e_{t+k} rather than edge e .

This definition resolves the foundational issues of defining a network coding solution in a graph with cycles. However, it is often easier to reason about the graph \hat{G} rather than the time-expanded graph G^* . For this purpose it is convenient to have a notion of a network coding solution defined on \hat{G} .

Definition 13 (Concise Network Coding Solution): Given an instance of the network coding problem and a solution defined in the time-expanded graph model, we define a concise representation of this solution as follows. For $e \in \hat{E}$:

- The *edge alphabet* is defined as $\Sigma(e) = \prod_{j=-\infty}^{\infty} \Sigma(e_j)$.
- The *edge function* is the function $f_e : \Sigma(S) \rightarrow \Sigma(e)$ where

$$f_e(M) = (\dots, f_{e_{-1}}(M), f_{e_0}(M), f_{e_1}(M), \dots).$$

B. Rate of a Network Coding Solution

Using the concise representation of a network coding solution, we define the rate of a solution.

Definition 14 (Rate (directed)): Given an instance of the network coding problem with underlying directed graph G , we say a network coding solution achieves rate r if all edge alphabets in the concise representation are finite, and there exists a constant b such that $\log_b |\Sigma(e)| \leq c(e)$ for each $e \in E$ and $\log_b |\Sigma(S(i))| \geq rd_i$ for each commodity i .

Definition 15 (Rate (undirected)): Given an instance of the network coding problem with underlying undirected graph G , we say a network coding solution achieves rate r if all edge alphabets in the concise representation are finite, and there exists a constant b such that for each edge $e \in E$ (represented by oppositely directed edges \vec{e} and \bar{e} in \hat{G}),

$$\log_b \left(|\Sigma(\vec{e})| \cdot |\Sigma(\bar{e})| \right) \leq c(e)$$

and for each commodity i ,

$$\log_b |\Sigma(S(i))| \geq rd_i.$$

Our definition of a rate r code allows the code designer to split the use of an undirected edge in any possible way. There could be other interpretations for the capacity region of an undirected edge; this issue was explored by Kramer and Savari [7]. For example, one could assume that an undirected edge (u, v) with capacity c can be used, at each time step, as either a capacity c channel from u to v or as a capacity c channel from v to u . Another possibility would be to interpret

this edge as allowing a channel of capacity c from u to v and a channel of capacity c from v to u . This at most doubles the maximum achievable rate and would be more restrictive than just using our definition with doubled edge capacities. We feel our definition is justified by the following reasons. First, it is consistent with the well-known push-to-talk two-way channel. Second, it is consistent with the usual view of network flow on undirected graphs in the combinatorial optimization literature.

The information theoretic view of network coding on directed acyclic graphs described in Section II extends to graphs with cycles. The definition is based on concise network coding solutions. For edge $e \in \hat{E}$, define the random variable $Y_e = f_e(M)$, where M is the random variable for the source messages. As before, we use $H(e)$ as shorthand for $H(Y_e)$. The analogues to Lemma 1 are as follows.

Lemma 18: Let G be a directed graph. Given a concise network coding solution of rate r , there exists a constant b such that the following hold. For all edges $e \in E$,

$$H(e) \leq c(e) \log_2 b$$

and for all commodities $i \in \mathcal{I}$,

$$H(S(i)) \geq rd_i \log_2 b.$$

Lemma 19: Let G be an undirected graph and let \hat{G} be the corresponding directed graph. Given a concise network coding solution of rate r , there exists a constant b such that the following hold. For any pair \vec{e} and \bar{e} of edges corresponding to an undirected edge e ,

$$H(\vec{e}) + H(\bar{e}) \leq c(e) \log_2 b$$

and for any edge $S(i)$

$$H(S(i)) \geq rd_i \log_2 b.$$

VII. INFORMATIONAL DOMINANCE IN GENERAL GRAPHS

In Section V, we introduced combinatorial criteria and an algorithm for determining when an edge set A informationally dominates another edge set B in a DAG. In this section we generalize these characterizations of informational dominance to graphs with cycles and undirected graphs. Recall that for an edge set $A \subseteq \hat{E}$ there is a corresponding edge set $A^* \subseteq E^*$ (the *time-expansion* of A) in the time-expanded graph.

Lemma 20: $\text{Dom}(A^*) = \text{Dom}(A^*)$.

Proof: We begin by proving that an edge e_t is informationally dominated by A^* if and only if e is informationally dominated by A . One direction is easy: if e is informationally dominated by A , then by definition this means that for any concise network code in \hat{G} , the information on A determines the signal transmitted on e , or equivalently the information on A^* determines the signals transmitted on e_t for all $t \in \mathbb{Z}$. To prove the other direction, if A doesn't informationally dominate e , then there is at least one network code N in G^* and at least one edge e_t such that the signal on e_t is not determined by A^* . But the property “ e_t is informationally dominated by A^* ” is invariant under time translations: for any integer k the network code $N_{(k)}$ has the property that e_{t+k} is not determined by A^* . So we find that none of the edges $\{e_t : t \in \mathbb{Z}\}$ are in $\text{Dom}(A^*)$. A similar argument applies to

the edges $S_t(i), S(i), T_t(i), T(i)$: each such edge belongs to $\text{Dom}(A)^*$ iff it belongs to $\text{Dom}(A^*)$.

Finally we prove that a memory edge (v_t, v_{t+1}) is in $\text{Dom}(A)^*$ if and only if it is in $\text{Dom}(A^*)$. Let $\mu_v = \{(v_t, v_{t+1}) : t \in \mathbb{Z}\}$ denote the set of all memory edges at v , and let $B = \text{In}(\mu_v)$. Note that μ_v is downstream of $B \setminus \mu_v$. If all incoming edges of v in G belong to $\text{Dom}(A)$ then (by the preceding paragraph) $\text{Dom}(A^*)$ contains $B \setminus \mu_v$, so $\text{Dom}(A^*)$ contains μ_v as well. Conversely, if $(v_t, v_{t+1}) \in \text{Dom}(A^*)$ then every incoming edge of v_t must be in $\text{Dom}(A^*)$. (Otherwise, a network code in which A^* fails to determine the signal on $e_{t-1} = (u_{t-1}, v_t)$ may easily be modified to a network code in which A^* fails to determine the signal on $e = (v_t, v_{t+1})$, by copying the signal from e_{t-1} to e .) Again using the preceding paragraph, this implies that every incoming edge of v is in $\text{Dom}(A)$, hence $(v_t, v_{t+1}) \in \text{Dom}(A)^*$. ■

Next, we show that the conditions of Theorem 10 are preserved under time-expansion.

Lemma 21: For an edge set $B \subseteq \hat{E}$, B satisfies conditions 2–4 in \hat{G} if and only if B^* satisfies 2–4 in G^* .

Proof: From the definition of the time-expansion B^* , we know that for each commodity i , $S(i) \in B^*$ iff $S(i) \in B$ and $T(i) \in B^*$ iff $T(i) \in B$. This establishes condition 2.

For conditions 3 and 4, it is useful to make the following definition. If $P = (v^{(1)}, v^{(2)}, \dots, v^{(j)})$ is a simple path in G , then P^* is the path $(v_1^{(1)}, v_2^{(2)}, \dots, v_j^{(j)}, v_{j+1}^{(j)}, v_{j+2}^{(j)}, \dots, v_n^{(j)})$ in G^* . (Note that $j \leq n$ because P is a simple path.) If P is a path in \hat{G} which traverses a source or sink edge $S(i)$ or $T(i)$, then P^* is the path in G^* formed by first removing the source or sink edges from P , then applying the above transformation to obtain a path in G^* , then prepending the path $(\sigma(i), \sigma'(i), v_1^{(1)})$ or appending the path $(v_n^{(j)}, \tau(i), \tau'(i))$.

If a vertex $v \in \hat{V}$ is reachable from a source in $\hat{G} \setminus B$ using a path P , then P^* is a path from a source to v_n in $G^* \setminus B^*$. Using time translation, we obtain a path from a source to v_t in G^* , for all t . Conversely, if there is a path P from a source to a vertex v_t in G^* , then by deleting the memory edges of P and dropping the time coordinates of all other edges, we obtain a path from a source to v in $\hat{G} \setminus B$. This establishes condition 3.

Finally, if $Q_1, P_1, Q_2, P_2, \dots, Q_j$ is an indirect path in $\hat{G}(B, i)$, then $Q_1^*, P_1^*, Q_2^*, P_2^*, \dots, Q_j^*$ is an indirect path in $G^*(B^*, i)$. Conversely, given an indirect path in $G^*(B^*, i)$ we may obtain an indirect path in $\hat{G}(B, i)$ by deleting all memory edges and dropping the time coordinates of all other edges. This establishes condition 4. ■

Corollary 22: For any edge set $A \subseteq \hat{E}$, $\text{Dom}(A)$ satisfies conditions 1–4. Any other set satisfying conditions 1–4 contains $\text{Dom}(A)$.

Corollary 23: There is an algorithm to compute $\text{Dom}(A)$ in time $O(|G|^3)$.

Proof: The algorithm described in Section V-F doesn't rely on the fact that G is a directed acyclic graph, it only relies on the combinatorial characterization in terms of properties 1–4. ■

Corollary 24: Lemma 7 and Lemma 9 remain valid when the underlying graph G is any directed graph.

A particularly useful case of Lemma 9 is the following relation identified by Jain et al. [13].

Crypto Inequality. If a graph G is partitioned into two parts, the set of edges crossing between the parts informationally dominates all commodities whose source and sink are in different parts.

More precisely, the crypto inequality can be expressed as follows.

Corollary 25 (Jain et al. [13]): Let G be an undirected graph, whose vertex set is partitioned into two subsets U and \bar{U} . Let $E(U, \bar{U})$ be the set of all edges in \hat{G} with one endpoint in U and the other in \bar{U} , and let $S(U, \bar{U})$ be the set of all edges $S(i)$ where i is a commodity whose source and sink belong to different connected components of $G \setminus E(U, \bar{U})$. Then $S(U, \bar{U}) \subseteq \text{Dom}(A)$.

VIII. OUTER BOUND FOR GENERAL GRAPHS

In this section we present an outer bound on the maximum achievable rate for general graphs. This bound is based on the following constraints. The constraints are applied in the graph G itself and represent constraints on *concise* network coding solutions.

- *Polymatroid inequalities:* The entropy function H is a non-negative, non-decreasing, submodular set function. In other words, for all $A, B \subseteq E$,

$$H(B) - H(A \cap B) \geq H(A \cup B) - H(A) \geq 0.$$

- *Informational Dominance:* For $A, B \subseteq E$, if A informationally dominates B , then $H(A) \geq H(B)$.
- *Independence of sources:* For any set $S(i_1), S(i_2), \dots, S(i_j)$ of sources,

$$H(S(i_1), \dots, S(i_j)) = H(S(i_1)) + \dots + H(S(i_j)).$$

- *Correctness:* For every commodity i , the edges $S(i)$ and $T(i)$ transmit the same symbol. Consequently, for any set U of edges, $H(U \cup \{S(i)\}) = H(U \cup \{T(i)\})$.
- *Rate:* A solution of rate r exists if and only if there exists a constant b such that $H(S(i)) \geq rd_i$ for all i and:

Directed Graphs: For every edge e ,

$$H(e) \leq c(e) \log_2 b.$$

Undirected Graphs: For every edge e ,

$$H(\vec{e}) + H(\overleftarrow{e}) \leq c(e) \log_2 b.$$

The constraints listed above constitute a powerful technique for proving upper bounds on the rate achievable by network coding in general graphs. These constraints can be used to generate a linear program. (Unfortunately the size of the linear program is exponential in the problem size.) Song, Yeung and Cai [3] presented a similar LP-bound for directed acyclic graphs, leaving open the problem of generalizing this bound to all graphs. (Their bound used a constraint that can be viewed as a restricted version of informational dominance.) A way to extend their bound to general graphs is to transform G into the time-expansion of G^* . Since this is a directed acyclic graph, the LP-bound of Song, Yeung and Cai then applies. There are many reasons why this solution is not entirely satisfactory. First, G^* is an infinite graph. The bound given by

considering solutions on any finite truncation of G^* could be strictly smaller than the actual maximum network coding rate. Secondly, working directly with edge sets in the much smaller graph \tilde{G} leads to substantially simpler and more condensed proofs of outer bounds, as well as faster algorithms for finding such outer bounds.

A. A Bound in Terms of $\overline{\Gamma^*}$

The first constraint says that entropy is a submodular, non-negative and non-decreasing set function. These characteristics imply what are known as *Shannon-type information inequalities* or the *basic inequalities*. Suppose we have n random variables. There are $2^n - 1$ joint entropies associated with these n random variables (one for every non-empty subset). For any n , these inequalities define a region in \mathbb{R}^N where $N = 2^n - 1$. It is known that there are points in this region that cannot be interpreted as the joint entropies associated with a collection of n random variables (see, e.g., [16]). This implies that the $2^n - 1$ joint entropies associated with a collection of n random variables must satisfy some additional constraints. Yeung defines Γ^* as the set of points in \mathbb{R}^N that correspond to the joint entropies associated with some n random variables [16]. Recently, a few *non-Shannon-type information inequalities* have been found [24], [25] but a complete characterization of the region Γ^* is still unknown.

An stronger outer bound could be obtained by replacing the first constraint with the set of all constraints that define Γ^* , or possibly the convex closure of Γ^* . Defining either of these regions in a computable way is a major open question. For now, we present our bound in terms of the basic inequalities but note that a better understanding of Γ^* or its closure could result in a tighter upper bound on the achievable rate for network coding.

IX. THE CAPACITY OF DIRECTED CYCLES

We now illustrate the use of informational dominance on an example problem: characterizing the capacity of directed cycles. In previous work, Kramer and Savari [7] specifically considered a 3-cycle with two commodities. This example is well-known since the cut-set bound obtained by partitioning the vertices into two sets does not tightly bound the multicommodity flow rate [22, Chapter 70]. Kramer and Savari proved a bound on the capacity that is tighter than the cut-set bound via Fano's inequality and properties of mutual information. They also posed the problem of characterizing the capacity of arbitrary directed cycles.

We prove that the network coding rate equals the maximum concurrent multicommodity flow rate for any k -pairs communication problem on a directed cycle. Let r be the rate of the maximum concurrent multicommodity flow. We begin by finding the sparsest cut. For a directed cycle, the sparsest cut consists of a single edge. This is a consequence of the fact that for each commodity i , there is only one directed path from $s(i)$ to $t(i)$. For a single edge u , the sparsity of the edge is given by the number of commodities whose path from $s(i)$ to $t(i)$ on the cycle traverses u . Let $\mathcal{I}(u)$ be the set of such

commodities and let e be an edge such that

$$\frac{c(e)}{\sum_{i \in \mathcal{I}(e)} d_i} = r.$$

In other words, the set $\{e\}$ is a sparsest cut.

Lemma 26: \mathcal{S} is informationally dominated by the set $A = \{e\} \cup \{S(i) : i \notin \mathcal{I}(e)\}$.

Proof: Let $\mathcal{J} \subseteq \mathcal{I}$ be the set of commodities i such that $S(i)$ is not informationally dominated by A . If $\mathcal{J} \neq \emptyset$, then we may label the commodities in \mathcal{J} as i_1, i_2, \dots, i_q according to the order in which their sink is encountered on the directed cycle starting at the head of edge e . Since $S(i_1) \notin \text{Dom}(A)$, Theorem 10 ensures that there is an indirect walk $Q_1, P_1, Q_2, P_2, \dots, Q_j$ for commodity i_1 in $G(\text{Dom}(A), i_1)$. Let i_j be the commodity whose source $s(i_j)$ is the initial vertex of path Q_j . Note that $S(i_j) \notin \text{Dom}(A)$ since $S(i_j)$ belongs to an indirect walk in $G(\text{Dom}(A), i_1)$; hence $i_j \in \mathcal{J}$. As we traverse the directed cycle starting at the head of e , we do not encounter $t(i_j)$ before $t(i_1)$ (by the defining property of i_1) and we do not encounter $t(i_1)$ before $s(i_j)$ (because Q_j is a path from $s(i_j)$ to $t(i_1)$ in $G \setminus \{e\}$). Hence we do not encounter $t(i_j)$ before $s(i_j)$, which means that the path from $s(i_j)$ to $t(i_j)$ lies in $G \setminus \{e\}$ and hence $i_j \notin \mathcal{I}(e)$. But this means that $S(i_j) \in A$, contradicting the fact that $S(i_j)$ does not belong to $\text{Dom}(A)$. ■

Lemma 26 implies that $H(A) \geq H(\mathcal{S})$. Using the submodularity of the entropy function and the independence of sources,

$$\begin{aligned} H(e) + \sum_{i \notin \mathcal{I}(e)} H(S(i)) &\geq \sum_{i \in \mathcal{I}} H(S(i)) \\ H(e) &\geq \sum_{i \in \mathcal{I}(e)} H(S(i)). \end{aligned}$$

If r^* is the rate of this concise solution, then by Lemma 18 there exists a constant b such that

$$c(e) \log b \geq H(e) \geq \sum_{i \in \mathcal{I}(e)} H(S(i)) \geq \sum_{i \in \mathcal{I}(e)} r^* d_i \log b.$$

We conclude that

$$r^* \leq \frac{c(e)}{\sum_{i \in \mathcal{I}(e)} d_i}.$$

This shows that the network coding rate equals the multicommodity flow rate.

X. THE CAPACITY OF UNDIRECTED GRAPHS

In Section VIII, we presented an outer bound for the rate of network coding problems on undirected graphs. In this section we focus on k -pairs communication problems in undirected graphs. Using the constraints that comprise our outer bound, we prove a gap between the sparsity of the graph and the network coding rate. This is the first known proof of such a gap in an undirected graph.⁴ An extension of this proof shows that the network coding rate equals the maximum concurrent multicommodity flow rate for an infinite class of interesting graphs.

⁴As mentioned earlier, the same gap was proven independently and concurrently by Jain et al. [13] and by Kramer and Savari [14].

A. Preliminaries

Our objective is to understand the relationship between multicommodity flow and network coding in undirected graphs. Network coding in undirected graphs is very different from network coding in directed graphs. To understand this difference we must review the definition of sparsity of a graph.

Definition 16 (Sparsity): For a k -pairs communication problem on an undirected graph G , let $\mathcal{I}(A)$ be the set of commodities whose source and sink are in different connected components of $G' = (V, E \setminus A)$. The *sparsity* of a set of edges A is given by

$$\mathcal{S}(A) = \frac{\sum_{e \in A} c(e)}{\sum_{i \in \mathcal{I}(A)} d_i}$$

and the sparsity of G is equal to $\min_{A \subseteq E} \mathcal{S}(A)$.

Sparsity illustrates a key difference between undirected and directed graphs: for undirected graphs, sparsity is an upper bound on the network coding rate, whereas for directed graphs it is not. For the directed case, this claim follows from our results in Section III. For the undirected case, the claim can be verified by observing that $S(i) \in \text{Dom}(A)$ whenever A is an undirected cut which separates the source and sink for commodity i , and then applying the informational dominance constraint from Section VIII. Thus, for undirected graphs, we have the following inequalities:

$$\begin{aligned} & \text{maximum concurrent flow rate} \\ & \leq \text{network coding rate} \leq \text{sparsity}. \end{aligned}$$

Therefore, in undirected instances in which the maximum multicommodity flow rate equals the sparsity of the graph, network coding provides no improvement to the rate. As observed by Kramer and Savari [7] and Li and Li [4], this implies that the network coding rate equals the flow rate for instances with $k \in \{1, 2\}$ and instances on planar graphs where the sources and sinks are on the boundary of the outer face.

However, for some graphs the sparsity can be greater than the maximum multicommodity flow rate. For example, Leighton and Rao [21] showed that the sparsity can exceed the maximum multicommodity flow rate by a factor of $\Omega(\log n)$, when the network is a constant-degree expander graph with n vertices. The interesting open question is whether the network coding rate can be larger than the maximum multicommodity flow rate for instances of the k -pairs communication problem on undirected graphs. Li and Li [4], [17] and Harvey et al. [5] independently conjectured that these two rates are always equal. While the only evidence supporting this belief is for restricted classes of graphs, our justification for making the conjecture is the same as for Edmonds' conjecture [26] that $P \neq NP$:

My reasons are the same as for any mathematical conjecture: (1) It is a legitimate mathematical possibility, and (2) I do not know.

One obstacle to proving the conjecture is the difficulty in proving upper bounds on the network coding rate which are stronger than those obtained by considering the sparsity of individual cuts. However, the techniques developed in this

paper allow us to prove the conjecture for a specific class of graphs where there is a gap between the sparsity and the maximum multicommodity flow rate. This class of graphs is described in the following sections. We remark that other contemporary work, e.g. [13], [14], has also proven bounds stronger than the sparsity bound.

B. The Okamura-Seymour Example

We consider a small example, shown in Figure 5, due to Okamura and Seymour [27]. This example is a 4-commodity flow problem in an undirected graph with 5 vertices. Each source has demand 1 and each edge has capacity 1. The maximum concurrent flow in this graph is $3/4$, achieved when each of the commodities a, b, c sends half of its flow on each of the two-hop paths from its source to its sink, and commodity d sends one-third of its flow on each of the two-hop paths from its source to its sink.




Fig. 5: The Okamura-Seymour Example.

The constraints from Section VIII suffice to prove the following theorem.

Theorem 27: The network coding rate of the Okamura-Seymour example is $3/4$.

Proof: We use \vec{uv} to refer to the directed edge from u to v and uv to refer to the pair $\{\vec{uv}, \vec{vu}\}$. We define $H(\vec{uv})$ to be the entropy of the random variable associated with edge \vec{uv} .

We start with the following downstreamness relationships.

$$\{S(a), \vec{da}, \vec{ea}\} \rightsquigarrow \{S(a), T(c), da, ea\}$$

$$\{S(b), \vec{db}, \vec{eb}\} \rightsquigarrow \{S(b), T(a), db, eb\}$$

$$\{S(c), \vec{dc}, \vec{ec}\} \rightsquigarrow \{S(c), T(b), dc, ec\}$$

Using the informational dominance and correctness constraints, we obtain the following entropy equalities.

$$H(S(a), \vec{da}, \vec{ea}) = H(S(a), S(c), da, ea)$$

$$H(S(b), \vec{db}, \vec{eb}) = H(S(a), S(b), db, eb)$$

$$H(S(c), \vec{dc}, \vec{ec}) = H(S(b), S(c), dc, ec)$$

Considering the right-hand side of these equations, note that each edge appears exactly once and the sources

$S(a), S(b), S(c)$ appear exactly twice. Adding these equations together and applying Theorem 2 we obtain

$$\begin{aligned} & H(S(a), \vec{da}, \vec{ea}) + H(S(b), \vec{db}, \vec{eb}) + H(S(c), \vec{dc}, \vec{ec}) \\ & \geq H(S(a), S(b), S(c), da, ea, db, eb, dc, ec) \\ & \quad + H(S(a), S(b), S(c)). \end{aligned}$$

Using submodularity on the left, independence of sources on the right, and cancelling terms:

$$\begin{aligned} & H(\vec{da}, \vec{ea}) + H(\vec{db}, \vec{eb}) + H(\vec{dc}, \vec{ec}) \\ & \geq H(S(a), S(b), S(c), da, ea, db, eb, dc, ec). \end{aligned} \quad (11)$$

We now use a new informational dominance relationship:

$$\begin{aligned} & \text{Dom}(\{S(a), S(b), S(c), da, ea, db, eb, dc, ec\}) \\ & \supseteq \{S(a), S(b), S(c), S(d)\} \end{aligned}$$

which follows from Lemma 9. Thus we can replace the right-hand side of Equation (11) with $H(S(a), S(b), S(c), S(d))$. Since the sources are independent,

$$\begin{aligned} & H(\vec{da}, \vec{ea}) + H(\vec{db}, \vec{eb}) + H(\vec{dc}, \vec{ec}) \\ & \geq H(S(a)) + H(S(b)) + H(S(c)) + H(S(d)). \end{aligned} \quad (12)$$

Next, starting with the following two downstreamness relationships,

$$\begin{aligned} \{S(d), \vec{ad}, \vec{bd}, \vec{cd}\} & \rightsquigarrow \{S(d), ad, bd, cd\} \\ \{\vec{ae}, \vec{be}, \vec{ce}\} & \rightsquigarrow \{T(d), ae, be, ce\} \end{aligned}$$

a similar sequence of arguments shows

$$\begin{aligned} & H(\vec{ad}, \vec{bd}, \vec{cd}) + H(\vec{ae}, \vec{be}, \vec{ce}) \\ & \geq H(S(a)) + H(S(b)) + H(S(c)) + H(S(d)). \end{aligned} \quad (13) \quad (14)$$

Adding (12) and (13) together implies

$$\begin{aligned} & H(\vec{da}, \vec{ea}) + H(\vec{db}, \vec{eb}) + H(\vec{dc}, \vec{ec}) \\ & + H(\vec{ad}, \vec{bd}, \vec{cd}) + H(\vec{ae}, \vec{be}, \vec{ce}) \\ & \geq 2(H(S(a)) + H(S(b)) + H(S(c)) + H(S(d))). \end{aligned}$$

By submodularity, we can replace each of the joint entropies on the left with the sum of the entropy of the individual random variables. Applying Lemma 19, if there exists a rate r solution, then

$$\begin{aligned} & c(da) + c(db) + c(dc) + c(ea) + c(eb) + c(ec) \\ & \geq 2r(d_a + d_b + d_c + d_d). \end{aligned}$$

In this example, each edge has capacity one and the demand for each commodity is one. This implies that $r \leq \frac{3}{4}$. ■

C. Special Bipartite Graphs

The analysis of the Okamura-Seymour example generalizes to an infinite family of undirected graphs which we call *special bipartite graphs*. This class of graphs contains infinitely many instances in which there is a gap between the sparsity and the

maximum multicommodity flow rate.⁵ Here we prove that the maximum achievable rate with coding equals the maximum multicommodity flow rate for special bipartite graphs.

Suppose G is a bipartite graph whose vertex set is partitioned into two independent sets V and W . Consider an instance of the k -pairs communication problem in G with the property that for every commodity, the source and sink both belong to V or they both belong to W . Let $E(W, V)$ denote the set of directed edges from W to V in \hat{G} . Let $\mathcal{S}(V)$ denote the set of sources in V , and $\mathcal{S}(W)$ the set of sources in W . For a vertex v , let $\text{In}(v)$ denote the set of incoming edges of v in \hat{G} , and let $\mathcal{T}(v)$ denote the set of edges in \mathcal{T} whose tail is v . Note that $\mathcal{S}(v) \in \text{In}(v)$ and $\text{In}(V) = E(W, V) \cup \mathcal{S}(V)$.

The relation $\text{In}(v) \rightsquigarrow (\text{In}(v) \cup \mathcal{T}(v))$ shows that $\text{Dom}(\text{In}(v)) \supseteq \text{In}(v) \cup \mathcal{T}(v)$, implying

$$H(\text{In}(v)) \geq H(\text{In}(v) \cup \mathcal{T}(v)).$$

By the correctness property, we may replace the sink random variables with source random variables. Consider the right-hand side of these equations for all $v \in V$. Each edge in $E(W, V)$ appears exactly once and each source in $\mathcal{S}(V)$ appears exactly twice. Adding these equations together and applying Theorem 2 we obtain

$$\sum_{v \in V} H(\text{In}(v)) \geq H(\text{In}(V)) + H(\mathcal{S}(V)).$$

Applying submodularity to the left-hand side, we obtain

$$\sum_{e \in E(W, V)} H(e) + H(\mathcal{S}(V)) \geq H(\text{In}(V)) + H(\mathcal{S}(V)). \quad (15)$$

We now claim that $\mathcal{S} \subseteq \text{Dom}(\text{In}(V))$. This holds because for any $S(i) \in \mathcal{S}(W)$, there is no indirect path from $S(i)$ to $T(i)$ in $G \setminus \text{In}(V)$. This yields the inequality

$$H(\text{In}(V)) \geq H(\mathcal{S}). \quad (16)$$

Combining Equation (15) and Equation (16) and canceling terms, we obtain

$$\sum_{e \in E(W, V)} H(e) \geq H(\mathcal{S}).$$

A symmetric argument shows that

$$\sum_{e \in E(V, W)} H(e) \geq H(\mathcal{S}).$$

Summing these two inequalities, we obtain an entropy inequality which implies

$$\begin{aligned} \sum_{e \in E} c(e) & \geq 2 \sum_{i \in \mathcal{I}} r d_i \\ r & \leq \frac{\sum_e c(e)}{2 \sum_i d_i}. \end{aligned}$$

This inequality is tight in instances where each source-sink pair is joined by a 2-hop path, and a dual-optimal length function assigns length 1 to every edge of G ; we call such instances *special bipartite graphs*. The Okamura-Seymour

⁵However, the gap is never greater than a factor of 2 for special bipartite graphs.

example is a special bipartite graph. It is not hard to come up with infinitely many other bipartite graphs satisfying this property.

Theorem 28: Multicommodity flow achieves the same maximum rate as network coding in any instance of the k -pairs communication problem defined on an underlying undirected bipartite graph G where a dual-optimal length function assigns 1 to every edge and every source-sink pair is distance 2 apart.

XI. GENERALIZATION

In this section we describe three generalizations of our results. First, we show how to extend our outer bound to general network coding problems. This requires extending our combinatorial characterization of informational dominance. We then describe how our outer bound can be modified to define an outer bound on the region of feasible rate vectors.

A. Preliminaries

In an instance of the general network coding problem, each commodity i can have multiple sinks $\mathcal{T}(i) = \{\tau(i)_1, \tau(i)_2, \dots, \tau(i)_{|\mathcal{T}(i)|}\}$.⁶ The notions *network coding solution*, *rate*, and *informational dominance* are defined exactly as before. For our purposes it is still convenient to represent sources and sinks with directed edges even when the network itself is represented as an undirected graph. Therefore, for every commodity i and every sink node $\tau(i)_j$, we assume $\tau(i)_j$ has no out-edges and $T(i)_j = (t(i)_j, \tau(i)_j)$ is the single in-edge to $\tau(i)_j$.

B. Outer Bound for General Network Coding Problems

We extend our outer bound from Section VIII to general network coding problems. Recall that the bound in Section VIII is derived from a set of conditions that any network coding solution must satisfy. Some of these conditions apply unchanged to general network coding problems, such as the basic properties of entropy. Others, such as the correctness condition, need to be adjusted for general network coding problems. We restate that bound with the necessary adjustments.

- *Polymatroid inequalities:* H is a non-negative, non-decreasing, submodular set function.
- *Informational Dominance:* For $A, B \subseteq E$, if A informationally dominates B , then $H(A) \geq H(B)$.
- *Independence of commodities:* For any set $\{S(i_1), S(i_2), \dots, S(i_\ell)\}$ of source edges,

$$H(S(i_1), \dots, S(i_\ell)) = H(S(i_1)) + \dots + H(S(i_\ell)).$$

- *Correctness:* For every commodity i , the source edge $S(i)$ and the sink edges $T(i)_1, \dots, T(i)_{|\mathcal{T}(i)|}$ transmit the same symbol. Consequently, for any set $U \subseteq E$ and $1 \leq j \leq |\mathcal{T}(i)|$, $H(S(i), U) = H(T(i)_j, U)$.
- *Rate:* A solution of rate r exists if and only if there exists a constant b such that $H(S(i)) \geq r d_i$ for all i and:

⁶It is also possible to allow multiple sources for a commodity, each transmitting an identical message. As remarked earlier, we will omit the details since they are relatively straightforward and since most realistic scenarios are covered by the model in which each commodity has only one source.

Directed Graphs: For every edge e ,

$$H(e) \leq c(e) \log_2 b.$$

Undirected Graphs: For every edge e ,

$$H(\vec{e}) + H(\overleftarrow{e}) \leq c(e) \log b.$$

Our outer bound in Section VIII is computable because each of the conditions can be used to create a set of linear constraints. Taken together these constraints form a large linear program whose optimal value is an upper bound on the maximum network coding rate. In order to fully extend our outer bound to general network coding problems we need to show that the informational dominance condition is decidable.

C. Informational Dominance for General Network Coding

We extend the combinatorial characterization of the informational dominance relationship to general network coding problems. The extension is straightforward and implies an equally straightforward adaptation of our polynomial time algorithm for information dominance to the general case. These extensions allow us to restate Theorem 10 for general network coding problems. Let $G(\text{Dom}(A), T(i)_j)$ be the graph $G(\text{Dom}(A), i)$ constructed with respect to the sink edge $T(i)_j$ for commodity i .

Theorem 29: For an edge set A , the set $\text{Dom}(A)$ satisfies the following conditions.

Condition 1: $A \subseteq \text{Dom}(A)$.

Condition 2: For every commodity i , $\text{Dom}(A)$ either contains the source edge and all sink edges for commodity i or it does not contain any of these edges.

Condition 3: Every edge in $E \setminus \text{Dom}(A)$ is reachable in $G \setminus \text{Dom}(A)$ from a source.

Condition 4: For every sink edge $T(i)_j$ in $G \setminus \text{Dom}(A)$, there is an indirect walk from $S(i)$ to $T(i)_j$ in $G(\text{Dom}(A), T(i)_j)$.

Furthermore, any set B satisfying these conditions contains $\text{Dom}(A)$.

D. Rate Regions

Thus far we have described our outer bound in terms of a single rate. We did this to facilitate our comparison of network coding and multicommodity flow. Suppose instead we wished to consider a vector of feasible rates $\mathbf{r} = [r_1, r_2, \dots, r_k]$ where the interpretation is that commodity i can send at a rate of $r_i d_i$. Our outer bound can easily be modified to handle rate vectors by replacing the rate condition with the following more general condition.

- *Rate:* A solution with rate vector $\mathbf{r} = [r_1, r_2, \dots, r_k]$ exists if and only if there exists a constant b such that $H(S(i)) \geq r_i d_i$ for all i and:

Directed Graphs: For every edge e ,

$$H(e) \leq c(e) \log_2 b.$$

Undirected Graphs: For every edge e ,

$$H(\vec{e}) + H(\overleftarrow{e}) \leq c(e) \log b.$$

All the other conditions from our outer bound remain the same.

XII. CONCLUSIONS

We have presented several results concerning the capacity of information networks; similar results have been concurrently and independently derived by Jain et al. [13] and Kramer and Savari [14]. First, we have derived an outer bound on the rate region of general information networks by blending graph theoretic and information theoretic techniques. Our bound applies to undirected and cyclic graphs as well as directed acyclic graphs. We have used our outer bound to precisely characterize the achievable rates for two infinite classes of graphs. We have also shown that the capacity is much larger than the maximum multicommodity flow rate in certain directed networks.

There are many important open questions related to the work in this paper.

- *Does our outer bound exactly characterize the network coding rate?* Our outer bound relies on properties of the entropy function and the concept of informational dominance. Are these the essential ingredients in determining the capacity of a network? For that matter, do other concurrently proposed bounds, e.g. [14], give a tighter bound on the network capacity? Our conditions can be specified as a linear program, which is extremely large even for small instances. Does the optimal value of this linear program equal the network coding rate? We suspect that the answer is no. For example, perhaps one can construct an example where non-Shannon-type information inequalities are needed to obtain a tight outer bound. In fact, Song, Yeung and Cai [3] stated a bound for directed acyclic graphs in which they optimized over the convex closure of Γ^* ; as mentioned in Section VIII-A, Γ^* is the set of all vectors p in $\mathbb{R}^{2^n - 1}$ for which there exists a set of n random variables such that p represents the $2^n - 1$ joint entropies defined by the set of n random variables. This region is defined by Shannon's basic inequalities, as well as certain non-Shannon-type information inequalities which are not fully understood. What bound for general graphs do we obtain if we optimize over the same set of vectors that Song, Yeung and Cai do?
- *For instances of the k -pairs communication problem on undirected graphs, does the multicommodity flow rate equal the network coding rate?* This conjecture remains open. The answer is yes for all instances in which the maximum multicommodity flow value equals the sparsity. In addition, we've found an infinite class of instances for which the answer is yes even though there is a gap between the maximum multicommodity flow value and the sparsity. Can this result be extended to all instances of the k -pairs communication problem on undirected graphs? A related question is to understand the relationship between the polytope of feasible multicommodity flows and the polytope characterizing the entropies associated with network coding solutions for a given instance of the k -pairs communication problem.
- *Are there efficient algorithms to compute or approximate the network coding rate?* Our outer bound consists of an extremely large linear program where any feasible point

gives an outer bound on the network coding rate. Solving this linear program exactly seems very difficult. Can one approximate the optimal value to this linear program in polynomial time? This would give an efficient means for computing an outer bound on the rate region for general information networks.

- *Are there efficient algorithms to compute solutions with optimal rate for general network coding problems?* Suppose one could exactly solve the linear program corresponding to our outer bound. Does the solution give valid joint entropies for some set of random variables? Can one construct a network coding solution by finding coding functions for the edges of the network that match these joint entropies? If it is computationally hard to find solutions with optimal rate, one might be content with an approximately optimal rate. Can the LP of our outer bound be used to prove the quality of an approximate solution?

ACKNOWLEDGMENTS

The authors would like to thank Eric Lehman, Michel Goemans, David Karger, Tom Leighton, Madhu Sudan, Varsha Dani, and László Babai for useful discussions related to this work. We would also like to thank the anonymous reviewers for suggesting many useful improvements to the exposition.

APPENDIX

n -WAY SUBMODULARITY

We now prove Theorem 2. For notational convenience, let $B_{0,\ell} = A_{\ell+1}$ and $B_{\ell+1,\ell} = \emptyset$ (where $1 \leq \ell < n$). It is straightforward to verify that the following two properties hold for $1 \leq j \leq \ell + 1$.

$$\begin{aligned} B_{j,\ell} \cup (B_{j-1,\ell} \cap A_{\ell+1}) &= B_{j,\ell+1} \\ B_{j,\ell} \cap (B_{j-1,\ell} \cap A_{\ell+1}) &= B_{j,\ell} \cap A_{\ell+1} \end{aligned} \quad (17)$$

The proof of Theorem 2 relies on the following lemma.

Lemma 30: Fix a value ℓ where $1 \leq \ell < n$. For $0 \leq j \leq \ell$,

$$\begin{aligned} &\sum_{i=1}^{\ell} H(B_{i,\ell}) + H(A_{\ell+1}) \\ &\geq \sum_{i=1}^j H(B_{i,\ell+1}) + \sum_{i=j+1}^{\ell} H(B_{i,\ell}) + H(B_{j,\ell} \cap A_{\ell+1}). \end{aligned}$$

Proof: By induction on j . For $j = 0$, the statement is trivial. For $j \geq 1$, we have

$$\begin{aligned} &\sum_{i=1}^{\ell} H(B_{i,\ell}) + H(A_{\ell+1}) \\ &\geq \sum_{i=1}^{j-1} H(B_{i,\ell+1}) + \sum_{i=j}^{\ell} H(B_{i,\ell}) + H(B_{j-1,\ell} \cap A_{\ell+1}) \\ &\geq \sum_{i=1}^{j-1} H(B_{i,\ell+1}) + \sum_{i=j+1}^{\ell} H(B_{i,\ell}) \\ &\quad + H(B_{j,\ell} \cup (B_{j-1,\ell} \cap A_{\ell+1})) \\ &\quad + H(B_{j,\ell} \cap (B_{j-1,\ell} \cap A_{\ell+1})) \end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^{j-1} H(B_{i,\ell+1}) + \sum_{i=j+1}^{\ell} H(B_{i,\ell}) \\
&\quad + H(B_{j,\ell+1}) + H(B_{j,\ell} \cap A_{\ell+1})
\end{aligned}$$

The first inequality holds by induction. The second inequality follows by applying submodularity to the j^{th} term of the second sum. The last equality follows from (17). ■

Proof (of Theorem 2). The proof is by induction on n . For $n = 2$, the theorem simply states that H is submodular. Assuming that the theorem holds for $n - 1$, we prove it for n .

$$\begin{aligned}
\sum_{i=1}^{n-1} H(A_i) + H(A_n) &\geq \sum_{i=1}^{n-1} H(B_{i,n-1}) + H(A_n) \\
&\geq \sum_{i=1}^n H(B_{i,n}).
\end{aligned}$$

The first inequality follows by induction. The second inequality follows by applying Lemma 30 with $\ell = n - 1$ and $j = \ell$. ■

REFERENCES

- [1] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley, 1991.
- [2] A. Schrijver, "On the history of combinatorial optimization (till 1960)." [Online]. Available: <http://homepages.cwi.nl/~lex/files/histco.pdf>
- [3] L. Song, R. Yeung, and N. Cai, "Zero-error network coding for acyclic networks," *IEEE Transactions on Information Theory*, vol. 49, pp. 3129–3139, Dec. 2003.
- [4] 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*, Sept. 2004.
- [5] N. J. A. Harvey, R. D. Kleinberg, and A. R. Lehman, "Comparing network coding with multicommodity flow for the k -pairs communication problem," MIT LCS, Tech. Rep. 964, September 2004.
- [6] S. P. Borade, "Network information flow: Limits and achievability," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, July 2002.
- [7] G. Kramer and S. A. Savari, "On networks of two-way channels," in *DIMACS Workshop on Algebraic Coding Theory and Information Theory (December 15-18, 2003, Rutgers University, Piscataway, NJ)*, ser. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, A. Ashikhmin and A. Barg, Eds. Rhode Island: American Mathematical Society, 2005, vol. 68, pp. 133–143.
- [8] 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.
- [9] A. Agarwal and M. Charikar, "On the advantage of network coding for improving network throughput," in *Proceedings of the IEEE Information Theory Workshop*, Oct. 2004.
- [10] Z. Li, B. Li, D. Jiang, and L. C. Lau, "On achieving throughput with network coding," in *Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, Mar. 2005.
- [11] G. Kramer, "Capacity results for the discrete memoryless network," *IEEE Transactions on Information Theory*, vol. 49, no. 1, pp. 4–21, Jan. 2003.
- [12] —, "Directed information channels with feedback," Ph.D. dissertation, Swiss Federal Institute of Technology (ETH), Zurich, 1998.
- [13] 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.
- [14] 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, to appear.
- [15] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Transactions on Information Theory*, vol. 49, no. 2, pp. 371–381, 2003.
- [16] R. W. Yeung, *A First Course in Information Theory*. Kluwer, 2002.
- [17] Z. Li and B. Li, "Network coding in undirected networks," in *Proceedings of the 38th Annual Conference on Information Sciences and Systems (CISS)*, Mar. 2004.
- [18] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen, "Polynomial time algorithms for multicast network code construction," *IEEE Transactions on Information Theory*, vol. 51, no. 6, pp. 1973–1982, June 2005.
- [19] V. Vazirani, *Approximation Algorithms*. Springer, 2001.
- [20] J. Cheriyan, H. J. Karloff, and Y. Rabani, "Approximating directed multicuts," in *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, Oct. 2001, pp. 348–356.
- [21] T. Leighton and S. Rao, "Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms," *Journal of the ACM*, vol. 46, no. 6, pp. 787–832, Nov. 1999.
- [22] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*. Springer-Verlag, 2003.
- [23] M. Sipser, *Introduction to the Theory of Computation*. Course Technology, 1996.
- [24] Z. Zhang and R. W. Yeung, "On characterization of entropy function via information inequalities," *IEEE Transactions on Information Theory*, vol. 44, no. 4, pp. 1440–1452, July 1998.
- [25] K. Makarychev, Y. Makarychev, A. Romashchenko, and N. Vereshchagin, "A new class of non Shannon type inequalities for entropies," *Communications in Information and Systems*, vol. 2, no. 2, pp. 147–166, Dec. 2002.
- [26] J. Edmonds, "Optimum branchings," *J. Research of the National Bureau of Standards*, vol. 71B, pp. 233–240, 1967.
- [27] H. Okamura and P. Seymour, "Multicommodity flows in planar graphs," *Journal of Combinatorial Theory, Series B*, vol. 31, no. 1, pp. 75–81, 1981.