On the Multiple Unicast Network Coding Conjecture

Michael Langberg
Computer Science Division
Open University of Israel
Raanana 43107, Israel
mikel@openu.ac.il

Muriel Médard
Research Laboratory of Electronics
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
medard@mit.edu

Abstract

In this paper, we study the multiple unicast network communication problem on undirected graphs. It has been conjectured by Li and Li [CISS 2004] that, for the problem at hand, the use of network coding does not allow any advantage over standard routing. Loosely speaking, we show that under certain (strong) connectivity requirements the advantage of network coding is indeed bounded by 3.

I. Introduction

In the network coding paradigm, internal nodes of the network may mix the information content of the received packets before forwarding them. This mixing (or encoding) of information has been extensively studied over the last decade, e.g., [2], [11], [9], [7], [6]. While the advantages of network coding in the *multicast* setting are currently well understood, this is far from being the case in the context of *general* network coding. In particular, determining the capacity of a general network coding instance is a long standing open problem, e.g., [3], [14].

In the general network coding problem, a set of source nodes $\{s_i\} \subseteq V$, wishes to transmit information to a set of terminal nodes $\{t_j\} \subseteq V$, according to a set of source/terminal requirements $\{(s_i, t_j)\}$ (implying that terminal t_j is interested in the information available at source s_i). For directed networks it was shown in [4] that any general network coding problem can be reduced to a multiple-unicast network coding instance in which there are k source/terminal pairs (s_i, t_i) and the objective is to design a coding scheme which allows t_i to recover the information present at s_i. Unlike the multicast scenario, determining the capacity of a k-unicast network coding instance is a long standing open problem. Specifically, it is currently not known whether this problem is solvable in polynomial time, is NP-hard, or maybe it is even undecidable [10] (the undecidability assumes that the alphabet size can be arbitrary and unbounded). Nevertheless, it is known that there is an unbounded gap between the capacity of the k-unicast problem in the directed network coding setting as opposed to the traditional setting of routing (were no encoding at internal nodes is allowed), e.g., [1]. This gap even holds for the simpler multicast scenario. The advantage of using network coding over traditional routing is the central theme discussed in this paper and is denoted throughout as the *coding advantage*.

Network coding in *undirected* networks has received considerably less attention from the research community. In such settings, the network is modeled by an undirected graph $G = (V, E)$. Each link $(v, u) \in E$ can transmit the information in both directions, i.e., from v to u and from u to v, subject to the restriction that the total amount of information transmitted over link (v, u) does not exceed its capacity.

The problems of unicast, broadcast, and multicast in undirected networks were studied by Li and Li in [12]. It was shown that for unicast and broadcast there is no advantage in the use of network coding over traditional routing. For the case of multicast, the coding advantage was shown to be at most 2, which complements the result of [1] stating that this advantage may be at least $8/7$. Little is known regarding the coding advantage for the more general k-unicast setting. To this day, the possibility that the advantage be unbounded (i.e., a function of the size of the network) has not been ruled out in the literature. In [12], [13] it is conjectured that for undirected graphs there is no coding advantage at all. This fact was verified on several special cases such as bipartite graphs [8], [5] and planar graphs [13] however is still open in general.

Loosely speaking, the Li and Li conjecture states that an undirected graph allowing a k-unicast connection using network coding also allows the same connection using routing. In this work we address a relaxed version of this

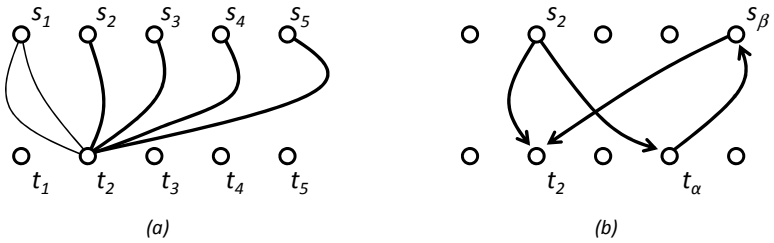

Fig. 1. *(a)* The flow F_2 equals the (disjoint) union of the flows $f_{i,2}$ for $i = 1, \ldots, k$ (here $k = 5$). Notice that each flow of type $f_{i,2}$ may consist of multiple paths connecting s_i and t_2. For example, in the figure, the flow $f_{1,2}$ consists of 2 disjoint paths (each of half the capacity) and the flow $f_{2,2}$ consists of a single path. *(b)* The flows corresponding to f_2^* in the proof of Lemma 1. Some of the flows connect s_2 and t_2 "directly", while others are routed via relay nodes t_α and s_β. Notice, that in both *(a)* and *(b)*, we represent our graph G schematically by drawing the source and terminal nodes only.

conjecture. Our relaxation has the following flavor. We show that an undirected graph allowing *more* than a k-unicast connection using network coding will *almost* allow a k-unicast connection using routing. The question here is how exactly do we define "more" and "almost".

Recall that in the k-unicast problem, there are k sources $\{s_1, \ldots, s_k\}$, k terminals $\{t_i, \ldots, t_k\}$, and one is required to design an information flow allowing each source s_i to transmit information to its corresponding terminal t_i. In the *k-multicast* problem, one is required to design an information flow allowing each source s_i to transmit information to all the terminals $\{t_1, \ldots, t_k\}$. Clearly, requiring that a network allows a k-multicast connection implies the corresponding k-unicast connection. In this work we show that an undirected graph allowing a k-*multicast* connection at rate r using network coding will allow the corresponding k-unicast connection at rate *close* to r, namely at rate $r/3$. The proof of our result is very simple in nature, and is based on a certain *flow decomposition* of the graph at hand.

We would like to stress that in our comparison of coding verses routing we are considering multiple-multicast coding rate on one hand and multiple unicast routing rate on the other. Considering the multiple unicast problem on both ends (as in the k-unicast conjecture) remains an intriguing open problem. In addition, we note that (although strongly related) the term of "k-multicast" and the standard notion of "multicast" (in which there is a single source) have subtle differences that seem central to the work at hand. We elaborate on these differences in Section I-C.

We further observe that our result implies an appealing *qualitative* statement regarding the use of network coding. Let $G = (V, E)$ be a directed graph. We denote by $\bar{G} = (V, \bar{E})$ the *undirected* graph obtained from G in which each directed edge $e = (u, v)$ in E appears as an undirected edge (u, v) in $\bar{E}$. Our result outlined above now implies the following statement. Given a directed graph G which allows k-multicast communication at rate r on k source/terminal pairs (s_i, t_i), by undirecting the edges of G (to obtain $\bar{G}$) one can obtain a feasible k-unicast routing solution of rate at least $r/3$. Namely, this implies the following informal statement:

Statement 1: In the setting in which one is guaranteed k-multicast communication, but requires only k-unicast: undirecting the edges of G is *as strong* as allowing network coding (up to a factor of 3).

A. Preliminaries

Let $G = (V, E)$ be a directed graph. We denote by $\bar{G} = (V, \bar{E})$ the *undirected* graph obtained from G in which each directed edge $e = (u, v)$ of capacity c_e in E appears as an undirected edge (u, v) of capacity c_e in $\bar{E}$. For simplicity we will consider $c_e = 1$ throughout this work. Our results extend naturally to arbitrary edge capacities.

We consider the k-unicast problem on directed and undirected networks. For directed networks, an instance of the problem is a graph G and k source terminal pairs (s_i, t_i) in G. The objective is to trasmitt information generated at source s_i to terminal t_i at maximal rate. The information generated at different sources is assumed to be independent. For a rigorous and detailed definition of the transmission rate and capacity of network coding instances see for example [14].

We define the *multiple unicast routing rate* of G, denoted as $\mathtt{U}_{\mathtt{RR}}(G)$, as the maximum value r such that there exists a routing scheme which enables communication between every source s_i and its corresponding terminal t_i at rate r. Equivalently, $\mathtt{U}_{\mathtt{RR}}(G)$ equals the value of the multicommodity flow on the instance at hand. We define the

multiple unicast coding rate of G, denoted as $\mathtt{U}_{\mathtt{CR}}(G)$, as the maximum value r such that there exists a network coding scheme which enables communication between every source s_i and its corresponding terminal t_i at rate r.

The k-unicast problem on undirected graphs $\bar{G}$ is defined similarly. Roughly speaking, a routing scheme (network coding scheme) on $\bar{G}$ is said to satisfy the capacity requirements of $\bar{G}$ if one can direct the edges of $\bar{G}$ to obtain a directed graph H for which the scheme still satisfies the capacity requirements of H. In this process, an undirected edge (u, v) of capacity c can be turned into two directed edges (u, v) and (v, u) of capacities c_1 and c_2 respectively, for any $c_1 + c_2 = c$. More formally, we seek a directed graph $H = (V, E_H)$ for which $\bar{H} = \bar{G}$ and the scheme at hand satisfies the capacity requirements of H. Here, we consider $\bar{H} = \bar{G}$ iff the graphs have the same edge set and edge capacitites. H is sometimes referred to as an *orientation* of $\bar{G}$. We define the routing rate $\mathtt{U}_{\mathtt{RR}}(\bar{G})$ and coding rate $\mathtt{U}_{\mathtt{CR}}(\bar{G})$ accordingly.

Multiple unicast conjecture [12], [13]: In [12], [13], Li and Li conjectured that in an undirected network $\bar{G}$ with multiple unicast sessions, network coding does not lead to any coding advantage. Namely, that

Conjecture 1 ([12], [13]): For an undirected graph $\bar{G}$ it holds that $\mathtt{U}_{\mathtt{CR}}(\bar{G}) = \mathtt{U}_{\mathtt{RR}}(\bar{G})$.

An equivalent way to phrase this conjecture is:

Conjecture 2: For any directed graph G it holds that $\mathtt{U}_{\mathtt{CR}}(G) \leq \mathtt{U}_{\mathtt{RR}}(\bar{G})$.

This can be interpreted as follows. Given a directed graph G which allows coding rate $\mathtt{U}_{\mathtt{CR}}(G)$, by undirecting the edges of G one can obtain a feasible routing solution of rate at least $\mathtt{U}_{\mathtt{CR}}(G)$. Namely, this implies the following informal statement: **"Undirecting the edges of G is *as strong* as allowing network coding"**. For completeness, we prove the equivalence between the two conjectures.

Proof: Assume that Conjecture 1 holds. Namely that $\mathtt{U}_{\mathtt{CR}}(\bar{G}) = \mathtt{U}_{\mathtt{RR}}(\bar{G})$. As $\mathtt{U}_{\mathtt{CR}}(G) \leq \mathtt{U}_{\mathtt{CR}}(\bar{G})$, we conclude that $\mathtt{U}_{\mathtt{CR}}(G) \leq \mathtt{U}_{\mathtt{CR}}(\bar{G}) = \mathtt{U}_{\mathtt{RR}}(\bar{G})$. Assume now that Conjecture 2 holds. Let H be the directed graph for which $\bar{H} = \bar{G}$ and $\mathtt{U}_{\mathtt{CR}}(H) = \mathtt{U}_{\mathtt{CR}}(\bar{G})$. Namely, H is the directed graph that *realizes* $\mathtt{U}_{\mathtt{CR}}(\bar{G})$. Now by Conjecture 2, $\mathtt{U}_{\mathtt{CR}}(\bar{G}) = \mathtt{U}_{\mathtt{CR}}(H) \leq \mathtt{U}_{\mathtt{RR}}(\bar{H}) = \mathtt{U}_{\mathtt{RR}}(\bar{G})$. As it always holds that $\mathtt{U}_{\mathtt{CR}}(\bar{G}) \geq \mathtt{U}_{\mathtt{RR}}(\bar{G})$, we conclude Conjecture 1. ■

B. Our result

In this work we prove a relaxed version of Conjecture 1 and 2. We start by some definitions. Let G be a directed graph and $\{(s_i, t_i)\}_{i=1}^{k}$ be a set of k source/terminal pairs. We say that G allows k-multicast communication of rate r (or that $\mathtt{M}_{\mathtt{CR}}(G) = r$) between the sources $\{s_1, \ldots, s_k\}$ and the terminals $\{t_1, \ldots, t_k\}$ if there is a network coding scheme which allows each terminal t_j to recover the information (of entropy r) present at each one of the sources. Again, for a rigorous and detailed definition of the transmission rate and capacity of network coding instances see for example [14].

It is not hard to verify that $\mathtt{M}_{\mathtt{CR}}(G) > 0$ only if G contains a path between each source s_i and terminal t_j. We refer to such graphs G as *strongly connected*. We say that a directed graph G has strong connectivity $\mathtt{SC}(G) = r$ if for every terminal t_j, $j = 1, \ldots, k$, there exists a *valid* multicommodity flow F_j consisting of k disjoint flows $\{f_{i,j}\}_{i=1}^{k}$ where each $f_{i,j}$ connects s_i and t_j with capacity r. See Figure 1(a). Here and throughout, a multicommodity flow F is valid if in taking all (disjoint) flows $f \in F$ together, one does not exceed the given edge capacities. It is easy to verify that the strong connectivity of G equals the capacity $\mathtt{M}_{\mathtt{CR}}(G)$. For completeness, the proof is given below.

Claim 1: $\mathtt{M}_{\mathtt{CR}}(G) = \mathtt{SC}(G)$.

Proof: In what follows, we assume that $\mathtt{SC}(G) = 1$ or $\mathtt{M}_{\mathtt{CR}}(G) = 1$ (the proof extends naturally to the general case of $r > 1$ as well). Assume that $\mathtt{SC}(G) = 1$. The first direction of our assertion now follows since G satisfies the so-called *multicast requirements*. Namely, enhance G by adding a new node s connected by an edge of unit capacity to all sources s_i. Denote the enhanced graph by G_s. Consider the multicast of k units of information over G_s from s to all terminals t_j. As $\mathtt{SC}(G) = 1$, the minimum cut between s and each t_j in G_s is at least k. This implies the existence of a network code over G_s which allows the required multicast, e.g., [11]. As s has exactly k outgoing edges of unit capacity, we may assume w.l.o.g. that in this network coding scheme no encoding is performed on the edges leaving s. It is now not hard to verify that the exact same coding scheme when applied on the original graph G will allow each terminal t_j to recover the information of all sources s_i, implying that $\mathtt{M}_{\mathtt{CR}}(G) \geq 1$.

Now assume that $\mathtt{M}_{\mathtt{CR}}(G) = 1$. As before consider the graph G_s. The coding scheme of G directly implies a multicast coding scheme for G_s, which in turn imply for each terminal j a flow from s to t_j of capacity k, e.g.,

Communication task	Notation with coding	Notation without coding	Coding advantage
Multiple multicast	$\mathtt{M}_{\mathtt{CR}}(\bar{G}) = \mathtt{SC}(\bar{G})$	Not referred to in this work	At least $8/7$ [1]
Multiple unicast	$\mathtt{U}_{\mathtt{CR}}(\bar{G})$	$\mathtt{U}_{\mathtt{RR}}(\bar{G})$	Unknown, see conjecture of [12], [13]

TABLE I

A SUMMARY OF THE NOTATION USED THROUGHOUT THIS WORK. LET $\bar{G}$ BE AN UNDIRECTED GRAPH. FOR THE MULTIPLE MULTICAST SCENARIO, IT HOLDS THAT $\mathtt{M}_{\mathtt{CR}}(\bar{G}) = \mathtt{SC}(\bar{G})$ (A SLIGHT VARIANT OF CLAIM 1). THE GAP BETWEEN ROUTING AND CODING IN THIS CASE IS AT LEAST THAT OF THE (SINGLE SOURCE) MULTICAST CASE, WHICH IS AT LEAST $8/7$ [1]. FOR THE MULTIPLE-UNICAST SCENARIO IT HOLDS THAT $\mathtt{U}_{\mathtt{CR}}(\bar{G}) \geq \mathtt{U}_{\mathtt{RR}}(\bar{G})$, HOWEVER NO GAP IS KNOWN TO EXIST. IN A SENSE, IN THIS WORK WE COMPARE *apples and oranges*. NAMELY, WE COMPARE $\mathtt{M}_{\mathtt{CR}}(\bar{G})$ AND $\mathtt{U}_{\mathtt{RR}}(\bar{G})$. WE NOTE THAT THE TABLE ABOVE REFERS TO UNDIRECTED GRAPHS $\bar{G}$. THE RATE OBTAINABLE FOR ANY DIRECTED ORIENTATION G OF $\bar{G}$ IS AT MOST THAT OF $\bar{G}$ (IN ALL THE SETTINGS ABOVE).

[14]. As the edges leaving s are all unit capacity, that latter implies a set of disjoint flows $\{f_{i,j}\}_i$, where $f_{i,j}$ connects s_i and t_j with unit capacity. This in turn implies that $\mathtt{SC}(G) \geq 1$. ■

The result of this work can be summarized in the following theorem. Before we state our theorem, we refer the reader to a summary of our notation in Table I.

Theorem 1: Let G be a directed graph and $\{(s_i, t_i)\}_{i=1}^k$ be a set of k source terminal pairs. Then $3\mathtt{U}_{\mathtt{RR}}(\bar{G}) \geq \mathtt{M}_{\mathtt{CR}}(G)$.

We note that one may phrase Theorem 1 in the following equivalent manner: Let $\bar{G}$ be an undirected graph and $\{(s_i, t_i)\}_{i=1}^k$ be a set of k source terminal pairs. Then $3\mathtt{U}_{\mathtt{RR}}(\bar{G}) \geq \mathtt{M}_{\mathtt{CR}}(\bar{G})$. To put our result in perspective, we further elaborate on the results and proof techniques appearing in [12].

C. Comparison to techniques of [12]

Let G be a directed graph. In the network coding *multicast* scenario, there is a single source s which wants to transmit the exact same information to a subset T of terminals in G. In the work of [12] the task of multicasting over undirected graphs $\bar{G}$ was studied. Using our notation, it was shown in [12] that $2\pi(\bar{G}) \geq \mathtt{SC}(\bar{G}) \geq \mathtt{SC}(G) = \mathtt{M}_{\mathtt{CR}}(G)$. Here $\pi(\bar{G})$ is the multicast routing rate, and $\mathtt{SC}(\bar{G})$ is the (minimum over the) Min-Cut between s and terminals $t_j \in T$ denoted by λ in [12]. The result of [12] is similar in nature to our main result. In fact, the constant of 2 in the work of [12] beats the constant 3 appearing in our result. However, the multicast scenario differs from that of k-multicast studied in this work in the sense that there is no single source node s but rather k source nodes $s_1, \ldots, s_k$.

In a nutshell, the crux of the proof of [12] includes a reduction in which the multicast instance $\bar{G}$ undergoes several *splitting* modifications, until it is turned into an instance $\bar{G}'$ to the *broadcast* problem (in which the terminal set includes the entire vertex set of $\bar{G}'$). Roughly speaking, this reduction preserves the (relationship between the) values of $\pi(\bar{G}')$ and $\mathtt{SC}(\bar{G}')$ when compared to that of $\pi(\bar{G})$ and $\mathtt{SC}(\bar{G})$. Once turned into a broadcast instance, it is proven that $2\pi(\bar{G}') \geq \mathtt{SC}(\bar{G}')$. This implies that $2\pi(\bar{G}) \geq \mathtt{SC}(\bar{G}) \geq \mathtt{SC}(G) = \mathtt{M}_{\mathtt{CR}}(G)$.

To the best of our judgment, the reduction used in [12] does not adapt to the k-multicast scenario addressed in this work. The main reason being the lack of a single source s governing the multicast connection. One may attempt to use the reduction of [12] combined with the ideas of Claim 1 in which we transform a multi-source instance into a single source instance. However, in such attempts, the reduced graph will have diverse connectivity and will no longer match the broadcast scenario of [12] and its analysis.

II. PROOF OF THEOREM 1

As before we assume that $\mathtt{M}_{\mathtt{CR}}(G) = 1$, and use the fact that this implies $\mathtt{SC}(G) = 1$ (Claim 1). The proof extends naturally to the general case as well (in which $\mathtt{M}_{\mathtt{CR}}(G) = r$). Consider the graph $\bar{G}$. Clearly, as $\mathtt{SC}(G) = 1$ it holds that $\mathtt{SC}(\bar{G}) = 1$ also. We now prove the following Lemma which implies that $\mathtt{U}_{\mathtt{RR}}(\bar{G}) \geq 1/3$. This will conclude our proof.

Lemma 1: Let $k \geq 2$. If for every $j = 1, \ldots, k$ there exists a valid multicommodity flow F_j consisting of k disjoint flows $\{f_{i,j}\}_{i=1}^k$ where each $f_{i,j}$ connects s_i and t_j with unit capacity; then there exists a valid multicommodity flow F^* consisting of k disjoint flows f_i^* connecting s_i to t_i, each of capacity $1/3$.

Proof: Consider the family $\mathcal{F}$ of unit capacity flows $\cup_j F_j = \{f_{ij} | i,j \in [k]\}$. Here, and throughout, $[k] = \{1,2,\ldots,k\}$. The family $\mathcal{F}$ is not necessarily a valid multicommodity flow in $\bar{G}$, in the sense that taking all flows in $\mathcal{F}$ one may exceed certain edge capacities. We first start by defining a variant of $\mathcal{F}$ that is indeed a multicommodity flow in $\bar{G}$. Recall, that each $F_j = \{f_{ij} | i \in [k]\}$ is a valid multicommodity flow in $\bar{G}$. Moreover, $\cup_j F_j = \mathcal{F}$. Thus, it holds that reducing the capacity of flows in $\mathcal{F}$ from unit value to a value of $\frac{1}{k}$ will result in a valid multicommodity flow. Let $\mathcal{F}_{\frac{1}{k}}$ be the set of flows appearing in $\mathcal{F}$ after their capacity has been reduced to $\frac{1}{k}$.

We now, refine the family $\mathcal{F}_{\frac{1}{k}}$ as follows: for each flow $f_{ij} \in \mathcal{F}_{\frac{1}{k}}$ of capacity $\frac{1}{k}$ we define $3k-4$ identical flows $\{f_{ij}^{\ell}\}_{\ell=1}^{3k-4}$, each of capacity $\frac{1}{k(3k-4)}$. Denote the new collection of flows by $\mathcal{F}_{\frac{1}{k(3k-4)}} = \{f_{ij}^{\ell} | i,j \in [k],\ \ell \in [3k-4]\}$. It is not hard to verify that $\mathcal{F}_{\frac{1}{k(3k-4)}}$ is a valid multicommodity flow.

Finally, we turn the valid multicommodity flow $\mathcal{F}_{\frac{1}{k(3k-4)}}$ into a flow $F^* = \{f_i^* | i \in [k]\}$ as asserted. It suffices to define f_i^* for each $i \in [k]$. The flow f_i^* will consist of two types of flows. The first type of flows will connect s_i and t_i directly. Namely, we add to f_i^*, $2k-3$ flows f_{ii}^{ℓ} from the set $\mathcal{F}_{\frac{1}{k(3k-4)}}$. The second flow type will connect s_i and t_i via two "relays" t_α and s_β. Namely, for each $\alpha, \beta \in [k] \setminus \{i\}$ we will add to f_i^* a flow $f_{i,\alpha}^{\ell}$ from s_i to t_α; the *reverse* of a flow $f_{\beta,\alpha}^{\ell}$ from t_α to s_β; and a flow $f_{\beta,i}^{\ell}$ from s_β to t_i. See Figure 1(b). These three flows together will connect s_i and t_i. All in all, to construct f_i^* we use the following flows of $\mathcal{F}_{\frac{1}{k(3k-4)}}$:

- $2k-3$ copies of f_{ii}.
- For each $\alpha \in [k] \setminus \{i\}$: $k-1$ copies of $f_{i\alpha}$.
- For each $\beta \in [k] \setminus \{i\}$: $k-1$ copies of $f_{\beta i}$.
- For each $\alpha, \beta \in [k] \setminus \{i\}$: 1 copy of $f_{\beta\alpha}$.

The total amount of flow from s_i to t_i will be

$$\frac{(2k-3)+(k-1)^2}{k(3k-4)} = \frac{k^2-2}{k(3k-4)} > \frac{1}{3}$$

It remains to show that $F^* = \{f_i^* | i \in [k]\}$ is indeed a valid multicommodity flow. Namely, that it uses exactly the flows of $\mathcal{F}_{\frac{1}{k(3k-4)}}$ as its building blocks. Let $i \neq j$. In the process of constructing the flows in F^* we use flows of type f_{ij}^{ℓ} exactly $3k-4$ times: $(k-1)$ times when constructing f_i^*; $(k-1)$ times when constructing f_j^*; and once for each f_α^* when $\alpha \neq i$ and $\alpha \neq j$. The same goes for the flow f_{ii}^{ℓ}: $(2k-3)$ times when constructing f_i^*; and once for each f_α^* when $\alpha \neq i$. ■

III. Conclusions

We have shown that, in undirected graphs that are r-strongly connected, the use of network coding for k-multicast is comparable (within a factor of 3) to the routing rate of an arbitrary set of k unicast connections. Our results address a relaxed version of the Li and Li conjecture, using a different approach to that used by Li and Li, which does not extend gracefully to our setting. We would like to stress that in our comparison of coding verses routing we are considering k-multicast coding rate on one hand and k-unicast routing rate on the other. Considering the multiple unicast problem on both ends (as in the k-unicast conjecture) remains an intriguing open problem. An interesting consequence of our result is that the bulk of the advantage (in our setting) of coding versus not coding in directed graphs may be obtained through considering an undirected version of the graph. This may have interesting consequences for wireless networks, since they are generally undirected. While it may at first blush seem that our results imply a bound of a factor of 3 for the advantage of k-multicast coding versus k-unicast non-coding in wireless networks, such a conclusion would misinterpret our results. Indeed, broadcast and half-duplex constraints do not in general allow us to operate a wireless network as an arbitrary undirected network.

There are several interesting directions for future work. We choose to mention the one that motivated this work: Can one prove the Li and Li conjecture when restricted to graphs that are r-strongly connected? Namely, given such graphs, can one show no (or a limited) advantage for coding in the multiple unicast setting, thus proving the Li and Li conjecture at least for r-strongly connected networks.

REFERENCES

[1] A. Agarwal and M. Charikar. On the Advantage of Network Coding for Improving Network Throughput. *Information Theory Workshop*, pages 247– 249, 2004.

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

[3] T. Cover and J. Thomas. *Elements of Information Theory, second edition*. John Wiley and Sons, 2006.

[4] R. Dougherty and K. Zeger. Nonreversibility and Equivalent Constructions of Multiple-Unicast Networks. *IEEE Transactions on Information Theory*, 52(11):5067–5077, 2006.

[5] N.J.A. Harvey, R. Kleinberg, and A. R. Lehman. On the Capacity of Information Networks. *IEEE Transactions on Information Theory*, 52(6):2345– 2364, 2006.

[6] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong. A Random Linear Network Coding Approach to Multicast. *IEEE Transactions on Information Theory*, 52(10):4413–4430, 2006.

[7] 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*, 51(6):1973– 1982, 2005.

[8] K. Jain, V. V. Vazirani, and G. Yuval. On the Capacity of Multiple Unicast Sessions in Undirected Graphs. *IEEE/ACM Transactions on Networking*, 14:2805 – 2809, 2006.

[9] R. Koetter and M. Médard. An Algebraic Approach to Network Coding. *IEEE/ACM Transactions on Networking*, 11(5):782 – 795, 2003.

[10] A. Rasala Lehman. *Network Coding*. Ph.d. thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, 2005.

[11] S.-Y. R. Li, R. W. Yeung, and N. Cai. Linear Network Coding. *IEEE Transactions on Information Theory*, 49(2):371 – 381, 2003.

[12] Z. Li and B. Li. Network Coding in Undirected Networks. In *Proceedings of the 38th Annual Conference on Information Sciences and Systems (CISS 2004)*, March 2004.

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

[14] R. W. Yeung, S.Y.R. Li, N. Cai, and Z. Zhang. *Network Coding Theory*. Now Publishers Inc, 2006.