

A Constant Bound on Throughput Improvement of Multicast Network Coding in Undirected Networks

Zongpeng Li, *Member, IEEE*
 Department of Computer Science
 University of Calgary
 zongpeng@ucalgary.ca

Baochun Li, *Senior Member, IEEE*
 Dept. of Electrical & Computer Engineering
 University of Toronto
 bli@eecg.toronto.edu

Lap Chi Lau
 Dept. of Computer Science & Engineering
 Chinese University of Hong Kong
 chi@cse.cuhk.edu.hk

Abstract—Recent research in network coding shows that, joint consideration of both coding and routing strategies may lead to higher information transmission rates than routing only. A fundamental question in the field of network coding is: how large can the throughput improvement due to network coding be? In this paper, we prove that in undirected networks, the ratio of achievable multicast throughput with network coding to that without network coding is bounded by a constant factor of 2, *i.e.*, network coding can at most double the throughput. This result holds for any undirected network topology, any link capacity configuration, any multicast group size, and any source information rate. This constant bound 2 represents the tightest bound that has been proved so far in general undirected settings, and is to be contrasted with the unbounded potential of network coding in improving multicast throughput in directed networks.

Index Terms—Network Coding, Multicast, Undirected Networks, Routing, Complexity

I. INTRODUCTION

THE throughput of information transmission within a data network is constrained by the network topology and link capacities. Traditional techniques in improving transmission throughput focus on strategically routing information flows along high bandwidth or multiple paths from the source to the destinations. Recently, it is shown that such routing strategies alone may not be sufficient. Rather, it is necessary to consider encoding/decoding data on nodes in the network, in order to achieve the optimal throughput [1], [2]. Since these coding operations are not restricted to source or destination nodes, they are referred to as *network coding*. A classic example that illustrates the power of network coding is shown in Fig. 1, where each link has unit capacity. With network coding, the achievable throughput is two. Without coding, the achievable throughput is only one, if integral routing is required, *i.e.*, if all link flow rates are integers (0 or 1 in this example).

Similar to source erasure codes, encoding and decoding operations in network coding are also defined over finite fields, which have fixed length representation of symbols. Therefore, information flows do not increase in size after being encoded. The introduction of network coding has essentially expanded the available strategies to achieve optimal transmission throughput: rather than only relying on routing strategies, an optimal transmission strategy to achieve the maximum throughput includes both a routing scheme and a corresponding coding scheme. Optimal throughput achieved




Fig. 1. A classic example from [1]: network coding helps achieve a multicast throughput of 2, in a topology where each link has unit capacity.

with coding is always lower bounded by that achieved without coding.

An important direction in network coding research is to understand and quantify the potential of network coding in improving end-to-end throughput of information transmission, in general multi-hop communication networks [3], [4], [5], [6], [7], [8], [9], [10], [11]. In this paper, we consider undirected networks with bidirectional links. We compare the achievable throughput with coding both a routing scheme and a corresponding coding scheme. Optimal throughput achieved with coding is always lower bounded by that achieved without coding.

An important direction in network coding research is to understand and quantify the potential of network coding in improving end-to-end throughput of information transmission, in general multi-hop communication networks [3], [4], [5], [6], [7], [8], [9], [10], [11]. In this paper, we consider undirected networks with bidirectional links. We compare the achievable throughput with coding to other parameters that have been previously defined to reflect a communication network's connectivity or capacity. Such parameters include the packing number (which is also the achievable throughput without coding), strength, and connectivity. We consider three types of communication sessions: unicast (one-to-one), broadcast (one-to-all) and multicast (one-to-many). We examine the relative order among the above four quantities, from which we derive upper bounds for the *coding advantage*, *i.e.*, the ratio of throughput improvement due to network coding. In contrast to previous work, which shows the coding advantage is not finitely bounded in directed networks [5], [9], [12], we show that the coding advantage is always bounded by a constant factor of two in undirected networks. Our proof holds for either fractional routing, where information flows can be split

and merged at arbitrarily fine scales, or half-integer routing, where each information flow being transmitted has either an integer or half-integer rate.

We extend the result to an Internet-like bidirected network model, where the bound for the coding advantage depends on how imbalanced link capacities can be on two opposite links between the same pair of nodes. We also provide brief discussions on the integral flow model and the hypergraph model. In addition, we prove that the achievable throughput, and therefore the coding advantage, are independent of the location of the information source within the communication group, which is a unique property that is only valid in undirected settings. Finally, we show that in many cases, including in both directed and undirected networks, with both integral and fractional routing, optimal throughput with network coding is much easier to compute than optimal throughput without coding. In particular, the small bounds on the coding advantage may lead to nice approximation algorithms.

The remainder of the paper is organized as follows. We introduce related work on network coding in Sec. II, prove bounds on the coding advantage for unicast, broadcast, and multicast in Sec. III, extend the result to group communications in Sec. IV, and to other network models in Sec. V. We then discuss the coding advantage and complexity issues in Sec. VI and finally conclude the paper in Sec. VII.

II. PREVIOUS RESEARCH

Ahlsweede *et al.* [1] initiated the study of network coding. They show examples that demonstrate the benefit of network coding, in terms of throughput improvement. They also prove the fundamental result that, for a multicast transmission in a directed network, if a rate x can be achieved for each receiver independently, it can also be achieved for the entire session. Koetter *et al.* [2] also derived this result for directed acyclic networks within an algebraic framework. They further extend the discussion to multiple transmissions, and examined the benefit of network coding in terms of robust networking.

Li *et al.* [13] show that linear codes suffice in achieving optimal throughput for a multicast transmission. The bound on the necessary base field size is first given by Koetter *et al.* [2]. They show that for a multicast session with throughput r and number of receivers k , there exists a solution based on a finite field $GF(2^m)$, for some $m \leq \lceil \log_2(kr + 1) \rceil$. This bound is then improved by Ho *et al.* to $m \leq \lceil \log_2(r + 1) \rceil$ [14].

Li *et al.* [13] also proposed the first code assignment algorithm, which performs an exponential number of vector independence tests. Jaggi *et al.* [9] observed that, exploiting flow information in the routing strategy dramatically simplifies the task, and designed a polynomial time code assignment algorithm accordingly. They also show that, in directed networks with integral routing, the coding advantage may increase proportionally as $\Omega(\log |V|)$, and therefore may be arbitrarily high.

Agarwal and Charikar [4] show that the potential of network coding to improve throughput is equivalent to the integrality gap of the bidirected cut relaxation of the minimum Steiner tree problem. This result is related to the bound 2 on coding

advantage proved in this paper in the following two aspects. First, it provides an alternative proof method for the bound 2 in the fractional flow case, since the integrality gap of the bidirected cut relaxation is bounded by the gap of the undirected relaxation of the minimum Steiner tree problem, which is known to be bounded by 2 [15], [16]. Second, it shows that the advantage of network coding in decreasing multicast cost is also bounded by 2 (assuming the linear link cost model in the classic min-cost flow literature [17]), since the minimum Steiner tree IP and its bidirected cut relaxation yield the optimal cost without and with network coding, respectively. The proof of bound 2 in this paper is combinatorial in nature and does not rely on linear programming techniques; it works for not only fractional routing but also half-integer routing.

Chekuri *et al.* studied the coding advantage in directed networks, without the symmetrical throughput requirement on multicast receivers [5]. They demonstrate classes of networks where the coding advantage is bounded by 2, including networks with two unit sources, or networks with two receivers. They also construct a network pattern where the coding advantage grows at rate $\Theta(\sqrt{|V|})$. In this paper, we focus on undirected networks with symmetrical multicast throughput. The bound we prove holds for any multicast throughput and any multicast group size.

Cannons *et al.* [18] and Dougherty *et al.* [19] also performed comparison studies between network capacity with network coding (coding capacity) and network capacity with routing only (routing capacity). In particular, they show that the network capacity is independent of the coding alphabet, and that while routing capacity of a network is always achievable, coding capacity is not.

Empirical comparisons of multicast throughput with and without network coding can be found in [3] and [20]. Both comparisons suggest that for random network topologies, the coding advantage is marginal at best. Other advantages are suggested instead, such as ease of management, robustness and ease for algorithm design.

The coding advantage in the case of multiple unicast sessions has also been examined in recent literature. While it is known that the coding advantage is larger than 1 for either directed networks or integral flows, it was conjectured by Harvey *et al.* [21] and by us [22] that network coding does not make a difference in undirected networks with fractional flows. This conjecture remains unsettled except for some special network cases [6], [7], [8].

Early network coding research often focus on directed acyclic networks, where the lack of cycles facilitates the definition and computation of the global coding vectors at each network node. For converting an undirected network into a directed one, our approach in this paper is the same as that used in classic network flow research and that adopted by Harvey *et al.* [6] and Kramer *et al.* [23]: each undirected link uv with capacity $C(uv)$ is viewed as a pair of directed links \vec{uv} and \vec{vu} , whose capacities $C(\vec{uv})$ and $C(\vec{vu})$ can be flexibly chosen as long as $C(\vec{uv}) + C(\vec{vu}) \leq C(uv)$. For extending the definition of network coding from directed acyclic networks to general directed networks, we refer the readers to the technique

of time-parameterized layered graphs by Ahlswede *et al.* [1] and by Yeung [24].

III. FINITE BOUNDS ON CODING ADVANTAGE IN UNDIRECTED NETWORKS

Network coding introduces a new dimension to the information transmission problem. Traditionally, only the routing dimension is considered in a transmission strategy; with network coding, a transmission strategy includes both the routing scheme and the coding scheme. Considering both dimensions together is necessary to achieve the maximum information transmission rate.

We use a simple graph $G = (V, E)$ to represent the topology of a network, and use a function $C : E \rightarrow Z^+$ to denote link capacities. However, when we study fractional routing models, having integer values in C is neither important nor necessary, and we assume C can take rational values too. The communication group is $M = \{S, T_1, \dots, T_k\} \subseteq V$, with S being the sender of the unicast, broadcast, or multicast session, by default. In our graphical illustrations, terminal nodes (nodes in the communication group) are black, and relay nodes (nodes not in the communication group) are white. We focus on fractional routing in this section, and will extend our discussion to integral routing later in the paper. For integral routing, all link capacities and flow rates have integer values. For fractional routing, we assume link capacities may be shared fractionally in both directions, and flows can be split and merged at arbitrarily fine scales.

We use $\chi(N)$ to denote the maximum throughput of a network N containing a single transmission session. We compare $\chi(N)$ with other parameters that are defined to characterize the connectivity or capacity of a communication network, including the *packing number*, *strength* and *connectivity*. We study and compare the four parameters for unicast, broadcast, and multicast transmissions, respectively. Based on results obtained from such comparison studies, we derive a bound on the coding advantage in each case. We refer the readers to our previous work [20], [25] for precise linear programming formulations of $\chi(N)$.

Packing refers to the procedure of finding pairwise edge-disjoint sub-trees of G , in each of which the communication group remains connected. The packing number of a communication network N is denoted as $\pi(N)$, and is equal to the maximum throughput without coding. The reason is that, each tree can be used to transmit one unit information flow from the sender to all receivers, therefore the packing number gives the maximum number of unit information flows that can be transmitted. Formally, let \mathcal{T} denote the set of all possible trees connecting the communication group, then the packing number can be defined using the following liner program:

$$\begin{aligned} & \text{Maximize} && \sum_{t \in \mathcal{T}} f(t) \\ & \text{Subject to:} && \\ & \left\{ \begin{array}{ll} \sum_{uv \in t} f(t) \leq C(uv) & \forall uv \in E \\ f(t) \geq 0 & \forall t \in \mathcal{T} \end{array} \right. \end{aligned}$$

In the LP above, $f(t)$ is a variable representing the amount of information flow one ships along tree t . One can either

require $f(t)$ to take integer values only (assuming all inputs in $C(uv)$ are integers), leading to integral packing, or allow $f(t)$ to take rational values, leading to fractional packing.

Strength is a kind of partition connectivity of the network [26], and is denoted as $\eta(N)$. Let \mathcal{P} denote all possible ways of partitioning N into a number of disconnected components, such that each component contains at least one terminal node. Then:

$$\eta(N) = \min_{p \in \mathcal{P}} \frac{|E_c|}{|p| - 1}$$

Here E_c is the set of inter-component links, and $|E_c| = \sum_{uv \in E_c} C(uv)$; $|p|$ is the number of components that the network is separated into, in p .

Connectivity refers to the minimum edge connectivity between a pair of nodes in the communication group, and is denoted as $\lambda(N)$. It is also the minimum size of a cut that separates the communication group. Fig. 2 illustrates the concept of these four parameters using an example multicast network.

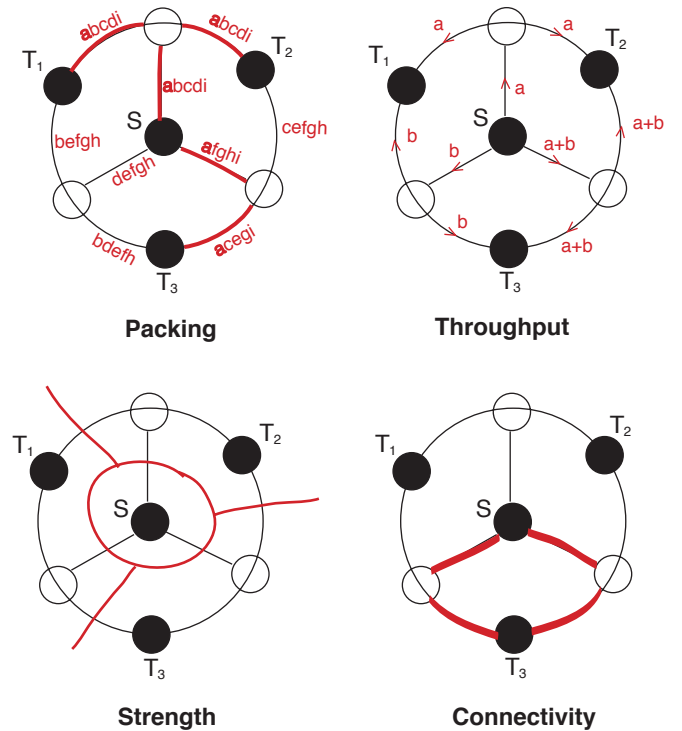


Fig. 2. The four parameters of a communication network N : packing number $\pi(N)$, throughput $\chi(N)$, strength $\eta(N)$ and connectivity $\lambda(N)$. In this particular network with unit capacity at each undirected link: $\pi(N) = 1.8$, nine trees (each labeled with a letter between ‘a’ and ‘i’) each of rate 0.2 can be packed; $\chi(N) = 2$, two unit information flows a and b can be delivered to all receivers simultaneously; $\eta(N) = 2$, in the illustrated optimal partition, $\frac{|E_c|}{p-1} = \frac{6}{4-1} = 2$; $\lambda(N) = 2$, each pair of terminal nodes is 2-edge-connected.

A. Unicast

In an undirected network with a unicast session $N = \{(G(V, E), C: E \rightarrow Z^+, M = \{S, T\} \subseteq V)\}$, the packing number $\pi(N)$ becomes the number of edge-disjoint S - T paths. The throughput $\chi(N)$ is the maximum information rate that

can be achieved in the $S \rightarrow T$ transmission. The strength $\eta(N)$ is now minimized over all cuts separating S and T — no valid partition with more than two components exists, since each valid component contains at least one terminal node and there are two terminal nodes only in the network. The connectivity becomes the edge-connectivity between S and T , *i.e.*, the minimum amount of edge capacity one needs to remove from the network, in order to separate S and T into disconnected components.

Based on previous results, we can show that the four quantities turn out to be all equal for a unicast transmission:

Theorem 1. For a unicast transmission in an undirected network, N ,

$$\pi(N) = \chi(N) = \eta(N) = \lambda(N).$$

Proof: Due to the fact that $\eta(N)$ can be minimized over simple cuts (cuts which separates the network into exactly two components) only, it becomes identical to $\lambda(N)$ by definition; both are equal to the size of the min-cut between S and T . Furthermore, observe that uncoded throughput is always bounded by coded throughput — network coding allows instead of requires coding, and any valid transmission scheme with routing only is still valid in the paradigm of network coding. Therefore $\pi(N) \leq \chi(N)$. Next, the $S \rightarrow T$ information rate is bounded by the S - T min-cut, *i.e.*, $\chi(N) \leq \eta(N)$. In order to finish the proof, it is sufficient to show $\pi(N) = \lambda(N)$, which is implied by Menger's Theorem [27]: *Let u, v be two vertices of a graph G . The maximum number of pairwise edge-disjoint u - v paths equals to the minimum number of edges whose removal separates u from v in G .* \square

It follows from Theorem 1 that network coding is not necessary in order to achieve the maximum throughput for a unicast session:

Corollary 1. The coding advantage for a unicast session is always 1.

B. Broadcast

Let $N = \{G(V, E), C : E \rightarrow Z^+, M = V = \{S, T_1, \dots, T_k\}\}$ be an undirected network containing a broadcast session, with S being the broadcast sender, and all other nodes in the network being receivers. The packing number $\pi(N)$ becomes the spanning tree packing number, *i.e.*, the maximum number of pair-wise edge-disjoint spanning trees that can be identified in the network. A spanning tree is a tree that connects every node in the network. The throughput $\chi(N)$ is the maximum information rate from S to every other node in the network, simultaneously. The strength $\eta(N)$ is still as defined; just note that for a broadcast network, every partition is valid, since each component of a partition always contains some node from the communication group. Connectivity $\lambda(N)$ becomes the size of the minimum simple cut of the network.

The fact that all nodes in the network request the same information leads to the following nice property for broadcast transmissions:

Theorem 2. For a broadcast transmission in an undirected

network, N ,

$$\frac{1}{2}\lambda(N) \leq \pi(N) = \chi(N) = \eta(N) \leq \lambda(N)$$

Proof: We first show that $\pi(N) = \chi(N) = \eta(N)$. Tutte-Nash-Williams Theorem characterizes the relationship between integral spanning tree packing and network strength [27], [28], [29]: *A graph G has x pairwise edge-disjoint spanning trees if and only if, for every vertex partition, there are at least $(p-1)x$ edges with endpoints in different components, where p is the number of components in the partition.* Tutte-Nash-Williams Theorem shows that, for the integral spanning tree packing problem, $\pi(N) = \lfloor \eta(N) \rfloor$. In the fractional flow model, one can apply the technique of scaling edge capacities up, and then scaling the solution down accordingly, to derive $\pi(N) = \eta(N)$ from the integral packing result. Furthermore, since the spanning tree packing number $\pi(N)$ is equal to the uncoded throughput, it can not exceed the coded throughput $\chi(N)$, *i.e.* $\pi(N) \leq \chi(N)$. Next, we observe that, if the network is partitioned into p components, each component not containing the source needs a total incoming edge capacity x in order to achieve throughput x ; therefore $(p-1)x$ inter-component edge capacity is required in total. This leads to $\chi(N) \leq \eta(N)$. Combining the above results, we have $\pi(N) = \chi(N) = \eta(N)$.

We next show that $\frac{1}{2}\lambda(N) \leq \pi(N)$ and $\eta(N) \leq \lambda(N)$. By definition, $\eta(N) \leq \lambda(N)$ holds since $\lambda(N)$ can be viewed as a special case of $\eta(N)$, where only partitions containing two components are considered. We now prove that $\frac{1}{2}\lambda(N) \leq \chi(N)$, using Nash-Williams' Weak Graph Orientation Theorem [30]: *a graph G has an x -edge-connected orientation if and only if it is $2x$ -edge-connected.* The weak orientation theorem implies that, in the fractional model, a broadcast network N always has a $\frac{1}{2}\lambda(N)$ -edge-connected orientation, *i.e.*, an orientation where the max-flow between each pair of nodes is at least $\frac{1}{2}\lambda(N)$. Combined with the result in directed networks, *a transmission rate that can be independently achieved for each receiver can be achieved for the entire session*, this implies $\frac{1}{2}\lambda(N) \leq \chi(N)$. \square

From Theorem 2, we can see that network coding has no potential in improving broadcast throughput either:

Corollary 2. The coding advantage for a broadcast session is always 1.

C. Multicast

Multicast is a more general form of communication than both unicast and broadcast. A unicast session can be viewed as a special case of multicast, where exactly two nodes in V are in the multicast group M . A broadcast session can also be viewed as a special case of multicast, where all nodes in V are in the multicast group M . In general, the multicast group M can be any subset of V that has size two or larger, and the packing problem becomes Steiner tree packing. A Steiner tree is a subtree in the network connecting all the terminal nodes, each relay node may or may not appear in the tree.

Theorem 3. For a multicast transmission in an undirected network, $N = \{G(V, E), C : E \rightarrow Z^+, M = \{S, T_1, \dots, T_k\} \subseteq$

$V\}$,

$$\frac{1}{2}\lambda(N) \leq \pi(N) \leq \chi(N) \leq \eta(N) \leq \lambda(N).$$

Proof: The fact that uncoded throughput is bounded by coded throughput again leads to $\pi(N) \leq \chi(N)$. Furthermore, the partition connectivity condition is necessary for a certain throughput to be feasible in multicast networks as well, therefore $\chi(N) \leq \eta(N)$. Next, $\eta(N) \leq \lambda(N)$ still holds due to the same argument as in the broadcast case — $\lambda(N)$ can be considered as a special case of $\eta(N)$ where only partitions containing two components are allowed. We now have $\pi(N) \leq \chi(N) \leq \eta(N) \leq \lambda(N)$, and need only to focus on the validity of $\frac{1}{2}\lambda(N) \leq \pi(N)$ in the rest of the proof, which is more complex in the case of multicast.

We build the proof of $\frac{1}{2}\lambda(N) \leq \pi(N)$ in the multicast case upon the same result proven in the broadcast case. More specifically, we transform the multicast network into a broadcast network, by eliminating the existence of relay nodes, while guaranteeing that $\frac{1}{2}\lambda(N) \leq \pi(N)$ holds after the transformation only if it holds before it.

Before describing the transformation in more detail, we first introduce Mader’s Undirected Splitting Theorem [30]: *Let $G(V + z, E)$ be an undirected graph so that (V, E) is connected and the degree $d(z)$ is even. Then there exists a complete splitting at z preserving the edge-connectivity between all pairs of nodes in V .*

A split-off operation at node z refers to the replacement of a 2-hop path $u-z-v$ by a direct edge between u and v , as illustrated in Fig. 3. A complete splitting at z is the procedure of repeatedly applying split-off operations at z until z is isolated.




Fig. 3. A split-off at node z .

The Undirected Splitting Theorem says that, if a graph has an even-degree non-cut node, then there exists a split-off operation at that node, after which the pairwise connectivities among the other nodes remain unchanged; and by repeatedly applying such split-off operations at this node, one can eventually isolate it from the rest of the graph, without affecting the edge-connectivity of any pair of nodes in it.

Now, consider repeatedly applying one of the following two operations on a multicast network: (1) apply a complete splitting at a non-cut relay node, preserving pairwise edge connectivities among terminal nodes in M ; or (2) add a relay node that is a M-cut node into the multicast group M , *i.e.*, change its role from a relay node to a receiver node. Here a M-cut node is one whose removal separates the multicast group into more than one disconnected components. Fig. 4 illustrates these two operations with a concrete example.

In order to meet the even node degree requirement in the Undirected Splitting Theorem, we first double each link




Fig. 4. Transforming a multicast network into a broadcast network, where the validity of $\frac{1}{2}\lambda(N) \leq \pi(N)$ can be traced back. (a) The original multicast network, with unit capacity on each link. (b) The network after applying operation (2), moving the M-cut node A into the multicast group. A becomes receiver T_3 . (c) The network after applying operation (1). A split-off was done at relay node B . A broadcast network is obtained.

capacity in the input network, then we scale the solution down by a factor of $\frac{1}{2}$ at the end. Note that, assuming the input network has integer link capacities, then each node has an even degree after doubling link capacities. Furthermore, a split-off operation does not affect the parity of the degree of any node in the network. Therefore the Undirected Splitting Theorem guarantees that as long as there are relay nodes that are not cut nodes, operation (1) is possible. Furthermore, operation (1) does not increase $\pi(N)$. Therefore, if $\frac{1}{2}\lambda(N) \leq \pi(N)$ holds after applying operation (1), it holds before applying operation (1) as well. Operation (2), applied to M-cut nodes, does not affect either $\pi(N)$ or $\lambda(N)$. So, again we can claim that for operation (2), if $\frac{1}{2}\lambda(N) \leq \pi(N)$ holds after applying the operation, it holds before applying the operation as well.

As long as there are relay nodes in the multicast network, at least one of the two operations can be applied. If both operations are possible, operation (1) takes priority. Since each operation reduces the number of relay nodes by one, eventually we obtain a broadcast network with terminal nodes only. By Theorem 2, $\frac{1}{2}\lambda(N) \leq \pi(N)$ holds.

Finally, note that we obtained an integral transmission strategy after doubling each link capacity. Therefore, after we scale the solution back by a factor of $\frac{1}{2}$, the transmission strategy is half-integral. \square

Corollary 3. For a multicast transmission in an undirected network, the coding advantage is upper-bounded by a constant factor of two, as long as half-integer routing is allowed.

Proof: By Theorem 3, $\frac{1}{2}\lambda(N) \leq \pi(N)$ and $\chi(N) \leq \lambda(N)$ as long as half integer routing is allowed. Therefore we conclude $\frac{1}{2}\chi(N) \leq \pi(N)$, *i.e.*, the coding advantage $\chi(N)/\pi(N) \leq 2$. \square

Fig. 5 shows the optimal throughput without coding of the multicast session given in Fig. 1, assuming half-integral routing and arbitrarily fractional routing respectively, with the network being undirected. Links labeled with the same letter or number form a Steiner tree. For example, the tree labeled with ‘a’ is highlighted in bold edges. In (a), each tree has capacity 0.5; in (b), trees labeled with a letter each has capacity 0.25, and trees labeled with a number have capacity 0.125. As a result, uncoded throughput achieved is 1.5 in (a) and 1.875 in (b) respectively, by transmitting a flow along each Steiner tree, with the flow rate equal to the tree capacity. Since optimal throughput with coding is 2, the corresponding coding advantages are 1.333 and 1.067, respectively.




Fig. 5. Throughput without network coding, for the example shown in Fig.1. (a) With half-integer routing, optimal throughput is 1.5. Each flow is of rate 0.5. (b) With arbitrary fractional routing, optimal throughput is 1.875. Each flow labeled with a letter and a number has rate 0.25 and 0.125, respectively.

We point out that none of the inequalities in Theorem 3 can be replaced with equality in general multicast networks. In particular, the network in Fig. 1 is an example where $\pi(N) < \chi(N)$, showing that packing number is not always equal to multicast throughput; the network in Fig. 6 is an example where $\chi(N) < \eta(N)$, showing that strength is not always equal to multicast throughput either. The value of $\chi(N)$ was computed using the linear optimization approach in [25]; the value of $\eta(N)$ was computed by optimizing over all valid network partitions. In the same network, $\pi(N) = 13.5$ can be obtained by enumerating all different multicast trees and then solving the tree-packing LP, and $\lambda(N) = 16$ can be obtained by solving two max-flow LPs. Therefore, we also have $\frac{1}{2}\lambda(N) < \pi(N)$ and $\eta(N) < \lambda(N)$ here. The construction of the network topology was inspired by the “gadget” used in the proof by Dahlhaus *et al.* [31] that shows the multiterminal cut problem is NP-Complete.




Fig. 6. A multicast network where packing number $\pi(N) = 13.5$, maximum throughput $\chi(N) = 13.5$, strength $\eta(N) = 14$ and connectivity $\lambda(N) = 16$. Unlabeled links each has capacity 1. This example shows that multicast throughput is in general not equivalent to strength in undirected networks.

IV. SOURCE INDEPENDENCE AND CODING ADVANTAGE FOR GROUP COMMUNICATIONS

In this section, we show that the achievable throughput for a multicast transmission does not depend on which node in the multicast group acts as the sender. In other words, if we move the information source from one node in the multicast group onto another, the optimal coding throughput remains unchanged. First, note that such a property does not hold in directed networks, where the connectivity between two nodes

can be arbitrarily different in the two directions. Second, it is rather obvious that this property holds for multicast without coding. The uncoded multicast problem is equivalent to the Steiner tree packing problem, and the packing number is defined upon the network topology and the terminal set, regardless of which node in the terminal set is the “sender”.

However, with network coding considered, it is less obvious whether the source independence property still holds. In Theorem 4, we provide an affirmative answer, based on which we then extend the bound 2 on coding advantage for multicast to the case of group communication.

Theorem 4. The optimal throughput of a multicast transmission in an undirected network is completely determined by the network topology, the link capacities, and the multicast group; it is not dependent on the selection of the sender within the multicast group.




Fig. 7. Two scenarios of reversing the $A \rightarrow B$ flow. Thicker links are being reversed.

Proof: The proof we present below is based on the following fact: a directed multicast transmission is feasible if and only if it satisfies all the simple cut conditions [2], i.e., if every simple cut between the source and any receiver has size no less than the multicast rate.

Suppose we exchange the sender and receiver roles between two terminal nodes A and B , and the optimal throughput before the exchange is f . Then there must exist a network flow of rate f from the sender to every receiver (including, in particular, from A to B). Consider reversing the $A \rightarrow B$ network flow, which has rate f . We show that after the reversal, simple cut conditions are still satisfied. Let C be another multicast node. Consider a cut that separates B and C . There are two cases, either A is in the same partition as B , or A is in the same partition as C , as shown in Fig. 7. In the first case, we have net flow of rate f traversing the cut from the AB component to the C component before the reversal, and an equal amount of flow in both directions will be reversed; therefore after the reversal, we still have the same amount of flow going from the AB component towards the C component. In the second case, similarly, the total flow going from the AC component towards the B component is f before the reversal, and all flows crossing the cut will be reversed. Therefore, after the flow reversal, we have flows of total rate f going from the B component towards the AC component. \square

Our proof also shows that, after the information source is moved, the same multicast throughput can be achieved with exactly the same bandwidth consumption on each link. Therefore, we can derive the following corollary:

Corollary 4.a A multicast rate is feasible if and only if it is

feasible with the information source separated into independent sub-sources and redistributed among the multicast group.

Fig. 8 shows an example containing the same network as in Fig. 1, with the two unit information sources at the top multicast node moved onto the two bottom multicast nodes respectively. Each information source can still be transmitted to all three multicast nodes, after the movement.




Fig. 8. Optimal transmission strategy after splitting and moving the information source, for the network shown in Fig. 1. Here P_1 is the source for a and P_2 is the source for b .

Corollary 4.a is relevant to video conferencing, where each participant multicasts his/her local audio/video data to every other participant, and receives audio/video data from them as well. By Corollary 4.a, a video conferencing session is feasible with a certain sending throughput requirement at each participant, if and only if the multicast transmission obtained by congregating all throughput requirement at one of the participants is feasible.

Since both throughput with and without network coding are source independent in undirected networks, we can extend the bound of 2 proved in Theorem 3 to the case of group communication:

Corollary 4.b The coding advantage for a group communication session in an undirected network is upper-bounded by a constant factor of 2.

As a final note of this section, we point out that the ‘flow reversing’ technique used in the proof of Theorem 4 is not always applicable in general networks. For example, Dougherty and Zeger [32] demonstrated a specific (directed) network N_7 consisting of multiple unicast sessions, such that N_7 has a network coding solution while its *reverse network* does not. A reverse network is defined as the network obtained from reversing every link direction and exchanging the sender/receiver role of every unicast session.

V. CODING ADVANTAGE IN RELATED NETWORK MODELS

A. Internet-like Bidirectional Networks

Given the fact that the coding advantage is finitely bounded in undirected networks but not in directed ones, it is natural to ask which model is closer to real-world networks, and whether the coding advantage is bounded in such networks. A real-world computer network, such as the current generation Internet, is usually bidirectional but not undirected. If u and v are two neighboring routers in the Internet, the amount of bandwidth available from u to v and that from v to u are both fixed and independent of each other. At a certain

moment, if the $u \rightarrow v$ link is congested and the $v \rightarrow u$ link is idling, it is not feasible to ‘borrow’ bandwidth from the $v \rightarrow u$ direction to the $u \rightarrow v$ direction, due to the lack of a dynamic bandwidth allocation module. Therefore, the Internet resembles an undirected network in that communication is bidirectional, and resembles a directed network in that each link is directed with a fixed amount of bandwidth.

The Internet can be better modeled with a *balanced directed* network. In a balanced directed network, each link has a fixed direction. However, a pair of neighboring nodes u and v are always mutually reachable through a direct link, and the ratio between $c(\vec{uv})$ and $c(\vec{vu})$ is upper-bounded by a constant ratio α . In cases where $\alpha = 1$, we have an absolutely balanced directed network. This is rather close to the reality in the Internet backbone, although last-hop connections to the Internet exhibit a higher level of asymmetry in upstream/downstream capacities. Based on the constant bound developed in the previous section, we now show that the coding advantage in such an α -balanced network is also finitely bounded.

Theorem 5. For a multicast session in an α -balanced bidirectional network, the coding advantage is upper-bounded by $2(\alpha + 1)$.

Proof: We first define a few notations. Let $N_{1:\alpha}$ be the α -balanced network; let N_1 be an undirected network with the same topology as $N_{1:\alpha}$, where $c(uv)$ in N_1 is equal to the smaller one of $c(\vec{uv})$ and $c(\vec{vu})$ in $N_{1:\alpha}$; let $N_{\alpha+1}$ be the undirected network obtained by multiplying every link capacity in N_1 with $\alpha + 1$. Then we have:

$$\begin{aligned} \pi(N_{1:\alpha}) &\geq \pi(N_1) \geq \frac{1}{\alpha+1} \pi(N_{\alpha+1}) \geq \frac{1}{\alpha+1} \frac{1}{2} \chi(N_{\alpha+1}) \\ &\geq \frac{1}{2(\alpha+1)} \chi(N_{1:\alpha}) \end{aligned}$$

In the derivations above, the third inequality is an application of Theorem 3; the other inequalities are based on definitions. \square

From Theorem 5, we can see that the more ‘balanced’ a directed network is, a smaller bound on the coding advantage can be claimed. In the case of an absolutely balanced network, the bound is 4. In arbitrary directed networks, α may approach ∞ , and correspondingly a finite bound on the coding advantage does not exist.

B. Integral Flows and Hypergraph Models

We now briefly discuss the cases of integral flow rates and hypergraph network models. The motivation of having integral flows lies behind the fact that in packet-switched networks, each packet constitutes an atomic data unit and data flow rates should be discrete. We showed in [33] that in the model of undirected networks with integral flows, $\chi(N) \leq 26\pi(N)$, *i.e.*, the coding advantage is bounded by a constant factor of 26. A hypergraph is similar to a graph except that each edge may connect more than two vertices. It naturally models data transmission in wireless networks equipped with omnidirectional antennae, where each transmission is heard by all users within a certain effective reception range. Király and

Lau [34] recently showed that in the hypergraph model, if a communication group is $2x$ -hyperedge-connected, then there is an orientation within which it is x -hyperedge-connected. This demonstrates that the multicast throughput can achieve at least half hyperedge connectivity. Furthermore, the hypertree packing number is upper-bounded by the hyperedge connectivity. In conclusion, the coding advantage is bounded by a factor of 2 in hypergraph models.

C. Non-Uniform Demand Networks and Average Throughput

Internet clients often exhibit significant heterogeneity in their network connections. Dial-up, cable modem, ADSL and campus networks each may provide a different connection capacity. Consequently, the max-flows between the multicast sender and different receivers may be drastically different. For real-world multicast applications such as media streaming, a common technique for accommodating heterogeneous clients in the same session is *layered coding* [35]. With layered coding, the source media is encoded into a base layer and several enhancement layers. Depending on the receiving capacity, a receiver may choose to receive and decode a subset of all the layers only, and accept a suboptimal playback quality.

Such real-world heterogeneity is naturally modeled by the *non-uniform demand network* model [36], [37]. A non-uniform demand network is very similar to a multicast network, with one node acting as the information source and a set of nodes acting as receivers. However, the requirement that all receivers receive information at the same rate is relaxed. Different users may receive information at a rate tailored according to her own capacity. The goal is to maximize the total throughput of all the receivers, or equivalent, the average throughput among all the receivers.

Unfortunately, non-uniform multicast is computationally a much harder problem than multicast. Cassuto and Bruck [37] proved that determining whether a rate vector χ_1, \dots, χ_k (χ_i denotes the desired throughput for receiver T_i) is feasible is NP-hard; consequently, maximizing the average throughput $\frac{1}{k} \sum_i \chi_i$ is also NP-hard.

A practical heuristic to circumvent the complexity of non-uniform multicast is to use successive layering [35]. The idea is to first compute the multicast rate that can be achieved for all receivers, and the corresponding multicast flow; then remove the utilized network capacities from N , remove the bottleneck receivers who can not achieve higher rates, and compute the maximum multicast rate in the residual network. These two steps are repeated until there is no more receivers in the residual network. By Theorem 3, within each layer of normal multicast, we can still achieve at least half of the coding throughput. Therefore the upper bound of 2 for the coding advantage holds if the average throughput with network coding is computed using this practical solution.

Characterizing the coding advantage using the true maximum average throughput with network coding is theoretically a more interesting direction. A similar framework as in the proofs of Theorem 2 and Theorem 3 can be adopted, *i.e.*, utilize the fact that coding throughput is upper-bounded by the connectivity, and then bound the ratio between throughput

without coding and connectivity. Based on a partial Steiner tree packing result by Bang-Jensen *et al.* [38], Chekuri *et al.* [36] proved that without network coding, a rate vector χ_1, \dots, χ_k can be achieved, such that for each T_i , χ_i is no less than half of the max-flow from S to T_i . This implies that the upper-bound of 2 for the coding advantage remains valid, for average throughput of non-uniform multicast.

VI. CODING ADVANTAGE AND COMPLEXITY

By establishing constant bounds on the coding advantage, so far in the paper we have been arguing in the direction that the benefit of network coding in terms of improving throughput is limited. Empirical studies on the coding advantage reveal a similar picture. For multicast in undirected networks, the largest coding advantage value observed is only slightly larger than one. In [25], we showed a small network where the value is $\frac{9}{8}$; in [4], the authors there pointed out a network pattern in which the coding advantage approaches $\frac{8}{7}$ as the network size grows towards infinity. It was also observed in [25] that the coding advantage value is almost always 1 in randomly constructed network topologies.

On the other hand, however, network coding may dramatically reduce the complexity of many optimization problems that arise in information dissemination. In some cases, the underlying network flow structure of coded multicast transmission leads to efficient linear optimization algorithms; in some other cases, the small bound on coding advantage leads to nice polynomial-time approximation algorithms. We provide some of such examples in the rest of this section.

A. Information Exchange

In an *information exchange* session, two nodes A and B need to transmit information to each other, with throughput requirement f_{AB} and f_{BA} respectively. This form of communication arises in scenarios such as: two sensor nodes exchange sensed data with each other [39], [40], two receivers in an asynchronous file downloading session reconcile received data with each other [41], or two online messaging applications stream multimedia data to each other concurrently. An information exchange session can also be viewed as two simultaneous unicast sessions between a pair of nodes, in opposite directions.

If network coding is ignored, then even a problem as simple as determining the feasibility of a single information exchange session is NP-hard, in directed networks with integral routing. One may derive this NP-hardness result from the proof given by Fortune *et al.* [42] that shows the edge-disjoint path problem is NP-hard for two opposite commodities. On the other hand, when network coding is taken into consideration, the information exchange problem becomes nicely tractable. As shown in Fig. 9, we can transform the coded information exchange problem into a coded multicast problem, which requires just two max-flow computations [1]. In the transformation, we add an extra source node to be the multicast sender, then assume the two unicast nodes are multicast receivers. Connect the sender S with A and B with an edge of capacity f_{AB} and f_{BA} respectively. Then we can verify that the original

information exchange session is feasible if and only if the resulting multicast session can achieve throughput $f_{AB} + f_{BA}$ with coding. The latter requirement is equivalent to have both the $S \rightarrow A$ max-flow and the $S \rightarrow B$ max-flow to be at least $f_{AB} + f_{BA}$ [1].




Fig. 9. Transforming the coded information exchange problem to the coded multicast problem.

B. Multicast with Fractional Routing

We now switch back to the undirected network model, and consider again multicast with fractional routing. Without network coding, the optimal multicast throughput problem is equivalent to the fractional Steiner tree packing problem, which is NP-Complete and APX-hard [43]. However, once network coding is supported, the optimal throughput can then be computed efficiently, due to the nice network flow structure underlying the coded information dissemination problem. We pointed out in [44] that optimal multicast with network coding in the fractional flow model can be solved using the linear programming approach. Various efficient and distributed algorithms have been designed thereafter within the linear optimization framework, for achieving maximum multicast throughput or minimum multicast cost [45], [46], [25], [47], [48], [49].

To conclude, in all the above three examples we have shown, the optimal transmission throughput problem is much more tractable with network coding considered. In the first and the third example, the problem is NP-Complete without network coding, and is P with network coding. In the second example, the problem is NP-hard either with or without network coding, but with network coding the optimal solution is much easier to approximate using polynomial-time algorithms.

C. Multicast with Integral Routing

Continuing from the previous subsection, we now further restrict our solutions to integral flows only. Without network coding, the achievable multicast throughput equals to the (integral) Steiner tree packing number. It has been shown that the Steiner tree packing problem is NP-Complete [43], [50]; it is even worse in the integral case: there did not exist any known polynomial time algorithm that can approximate the problem to any constant ratio, until the result $26\pi(N) \geq \lambda(N)$ [33] was recently proven in the integral flow model.

On the other hand, by taking network coding into consideration, we are led to a 2-approximation for the optimal multicast throughput. We show this claim by examining the relation between connectivity and throughput in the integral

model. We have shown that $\frac{1}{2}\lambda(N) \leq \pi(N) \leq \chi(N)$ holds in the fractional model; more accurately, it holds as long as half-integer flows are allowed. In the integral model, it is not known whether $\frac{1}{2}\lambda(N) \leq \pi(N)$ still holds or not. In fact, it is a well known open problem in graph theory. The best result known so far is the aforementioned ratio $\frac{1}{26}$. On the other hand, we show that throughput with network coding can still achieve half connectivity in the integral routing model.

Theorem 5. For a multicast transmission in an undirected network, N , $\lfloor \frac{1}{2}\lambda(N) \rfloor \leq \chi(N)$ holds under the integral routing requirement.

Proof: Our proof is based on Nash-Williams' Strong Graph Orientation Theorem [30]: every undirected graph $G(V, E)$ has an orientation $G' = (V, D)$ for which $\lambda_{G'}(u, v) \geq \lfloor \frac{1}{2}\lambda_G(u, v) \rfloor$, for all $u, v \in V$. This theorem guarantees the existence of a "well-balanced" orientation of an undirected network, within which the connectivity from any node u to any other node v is at least half of the connectivity between u and v in the original undirected network.

From the theorem above, we know that if $\lambda(N) = x$, then there is an integral orientation of the network, such that the directed connectivity among the multicast group M is at least $\lfloor \frac{1}{2}x \rfloor$. This implies that the integral max-flow from S to each receiver T_i is at least $\lfloor \frac{1}{2}x \rfloor$. Then by the feasibility condition of directed multicast [1], [2], there is an integral routing scheme to achieve $\chi(N) \geq \lfloor \frac{1}{2}x \rfloor$. \square

Corollary 5. The optimal multicast throughput problem in undirected networks with integral routing can be approximated within a factor of two in polynomial time.

Proof: The claim $\chi(N) \leq \lambda(N)$ in Theorem 3 still holds in the integral case. Combined with Theorem 5, we have $\lfloor \frac{1}{2}\lambda(N) \rfloor \leq \chi(N) \leq \lambda(N)$. Therefore computing $\lambda(N)$ gives a 2-approximation for $\chi(N)$. Note that $\lambda(N)$ is obviously computable in polynomial time — in the worst case, one can compute the max-flow between each pair of multicast nodes, and take the minimum value among them. It is also possible to find the detailed transmission strategy that achieves the approximated throughput value, since polynomial time algorithms exist for both the orientation [30] and code assignment [9], [12]. \square

To conclude, in all the above three examples we have shown, the optimal transmission throughput problem is much more tractable with network coding considered. In the first and the third example, the problem is NP-Complete without network coding, and is P with network coding. In the second example, the problem is NP-hard either with or without network coding, but with network coding the optimal solution is much easier to approximate using polynomial-time algorithms.

VII. CONCLUDING REMARKS

In this paper, we have compared the coded transmission throughput with the packing number, strength, and connectivity, in an undirected network with a unicast, broadcast, and multicast transmission, respectively. Our results lead to small constant bounds on the coding advantage in these cases: the coding advantage is always 1 for unicast or broadcast, and

is at most 2 for multicast. The bound 2 for multicast is then extended to the case of group communication, by showing that achievable throughput is source independent in undirected networks, either with or without network coding. We also make the observation that applying network coding makes it possible to design efficient algorithms that compute and achieve the optimal transmission throughput. Finite bounds on the coding advantage discussed throughout this paper are summarized in Table I.

The following questions on the coding advantage are still open. First, for multiple unicast sessions that concurrently co-exist in the same network, is the coding advantage always 1, assuming undirected networks with fractional routing [21], [22]? Second, with arbitrary fractional routing, can the bound of two for coding advantage be further tightened? Third, is the bound of two still valid if we replace the half-integer routing requirement with integral routing? Finally, is the bound of two still valid if we have multiple concurrent communication sessions?

REFERENCES

- [1] R. Ahlswede, N. Cai, S. R. Li, and R. W. Yeung, "Network Information Flow," *IEEE Trans. on Inform. Theory*, vol. 46, no. 4, pp. 1204–1216, July 2000.
- [2] R. Koetter and M. Médard, "An Algebraic Approach to Network Coding," *IEEE/ACM Trans. on Networking*, vol. 11, no. 5, pp. 782–795, October 2003.
- [3] Y. Wu, P. A. Chou, and K. Jain, "A Comparison of Network Coding and Tree Packing," in *Proc. of ISIT*, 2004.
- [4] A. Agarwal and M. Charikar, "On the Advantage of Network Coding for Improving Network Throughput," in *IEEE Inform. Theory Workshop*, October 2004.
- [5] C. Chekuri, C. Fragouli, and E. Soljanin, "On Average Throughput and Alphabet Size in Network Coding," *IEEE Trans. on Inform. Theory*, vol. 52, no. 6, pp. 2410–2424, June 2006.
- [6] N. J. A. Harvey, R. Kleinberg, and A. R. Lehman, "On the Capacity of Information Networks," *IEEE Trans. on Inform. Theory*, vol. 52, no. 6, pp. 2345–2364, June 2006.
- [7] N. J. A. Harvey and R. Kleinberg, "Tighter Cut-Based Bounds for k -pairs Communication Problems," in *43rd Annual Allerton Conference on Communication, Control, and Computing*, 2005.
- [8] K. Jain, V. V. Vazirani, and G. Yuval, "On the Capacity of Multiple Unicast Sessions in Undirected Graphs," *IEEE Trans. on Inform. Theory*, vol. 52, no. 6, pp. 2805–2809, June 2006.
- [9] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Eger, K. Jain, and L. M. G. Tolhuizen, "Polynomial Time Algorithms for Multicast Network Code Construction," *IEEE Trans. on Inform. Theory*, vol. 51, no. 6, pp. 1973–1982, June 2005.
- [10] D. Tuninetti and C. Fragouli, "On the Throughput Improvement due to Limited Complexity Processing at Relay Nodes," in *Proc. of ISIT*, 2005.
- [11] J. Liu, D. Goeckel, and D. Towsley, "Bounds on the Gain of Network Coding and Broadcasting in Wireless Networks," Department of Computer Science, University of Massachusetts, UM-CS-2006-050, Tech. Rep., October 2006.
- [12] P. Sanders, S. Eger, and L. Tolhuizen, "Polynomial Time Algorithm for Network Information Flow," in *Proc. of the 15th ACM Symposium on Parallelism in Algorithms and Architectures*, 2003.
- [13] S. Y. R. Li, R. W. Yeung, and N. Cai, "Linear Network Coding," *IEEE Trans. on Inform. Theory*, vol. 49, no. 2, pp. 371–381, February 2003.
- [14] T. Ho, D. R. Karger, M. Médard, and R. Koetter, "Network Coding From a Network Flow Perspective," in *Proc. of ISIT*, 2003.
- [15] A. Agrawal, P. Klein, and R. Ravi, "When Trees Collide: An Approximation Algorithm for the Generalized Steiner Tree Problem on Networks," *SIAM Journal on Computing*, vol. 23, no. 3, pp. 440–456, June 1995.
- [16] M. X. Goemans and D. P. Williamson, "A General Approximation Technique for Constrained Forest Problems," *SIAM Journal on Computing*, vol. 24, no. 2, pp. 296–317, April 1995.
- [17] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Upper Saddle River, New Jersey: Prentice Hall, 1993.
- [18] J. Cannons, R. Dougherty, C. Freiling, and K. Zeger, "Network Routing Capacity," in *Proc. of ISIT*, 2005.
- [19] R. Dougherty, C. Freiling, and K. Zeger, "Unachievability of Network Coding Capacity," *IEEE Trans. on Inform. Theory*, vol. 52, no. 6, pp. 2365–2372, June 2006.
- [20] Z. Li, B. Li, D. Jiang, and L. C. Lau, "On Achieving Optimal Throughput with Network Coding," in *Proc. of IEEE INFOCOM*, 2005.
- [21] N. J. A. Harvey, R. D. Kleinberg, and A. R. Lehman, "Comparing Network Coding with Multicommodity Flow for the k -pairs Communication Problem," Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, <http://theory.csail.mit.edu/nickh/Publications/Gaps/TR964.pdf>, Tech. Rep., November 2004.
- [22] Z. Li and B. Li, "Network Coding: The Case of Multiple Unicast Sessions," in *42nd Annual Allerton Conference on Communication, Control, and Computing*, 2004.
- [23] G. Kramer and S. A. Savari, "Edge-Cut Bounds on Network Coding Rates," *Journal of Network and Systems Management*, vol. 14, no. 1, pp. 49–67, 2006.
- [24] R. W. Yeung, *A First Course in Information Theory*. New York, NY: Springer, 2002.
- [25] Z. Li, B. Li, and L. C. Lau, "On Achieving Optimal Multicast Throughput in Undirected Networks," *IEEE Trans. on Inform. Theory*, vol. 52, no. 6, pp. 2467–2485, June 2006.
- [26] W. H. Cunningham, "Optimal Attack and Reinforcement of a Network," *Journal of the ACM*, vol. 32, pp. 549–561, 1985.
- [27] D. B. West, *Introduction To Graph Theory, 2nd Edition*. Upper Saddle River, New Jersey: Prentice Hall, 2001.
- [28] W. T. Tutte, "On the Problem of Decomposing a Graph Into n Connected Factors," *J. London Math. Soc.*, vol. 36, pp. 221–230, 1961.
- [29] C. S. J. A. Nash-Williams, "Edge-disjoint Spanning Trees of Finite Graphs," *J. London Math. Soc.*, vol. 36, pp. 445–450, 1961.
- [30] A. Frank, "Edge-connection of Graphs, Digraphs, and Hypergraphs," Enevry Research Group on Combinatorial Optimization, Budapest, Hungary, <http://www.cs.elte.hu/egres>, Tech. Rep., September 2000.
- [31] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis, "The Complexity of Multiterminal Cuts," *SIAM Journal on Computing*, vol. 23, no. 1, pp. 864–894, 1994.
- [32] R. Dougherty and K. Zeger, "Nonreversibility and Equivalent Construction of Multiple-Unicast Networks," *IEEE Trans. on Inform. Theory*, vol. 52, no. 11, pp. 5067–5077, 2006.
- [33] L. C. Lau, "An Approximate Max-Steiner-Tree-Packing Min-Steiner-Cut Theorem," in *Proc. of the 45th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2004.
- [34] T. Király and L. C. Lau, "Approximate Min-Max Theorems for Steiner Rooted-Oriented Hypergraphs," in *Proc. of the 47th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.
- [35] J. Zhao, F. Yang, Q. Zhang, Z. Zhang, and F. Zhang, "LION: Layered Overlay Multicast with Network Coding," *IEEE Trans. on Multimedia*, vol. 8, no. 5, October 2006.
- [36] C. Chekuri, C. Fragouli, and E. Soljanin, "On Achievable Information Rates in Single-Source Non-Uniform Demand Networks," in *Proc. of ISIT*, 2006.
- [37] Y. Cassuto and J. Bruck, "Network Coding for Non-Uniform Demands," in *Proc. of ISIT*, 2005.
- [38] J. Bang-Jensen, A. Frank, and B. Jackson, "Preserving and Increasing Local Edge-Connectivity in Mixed Graphs," *SIAM Journal on Discrete Mathematics*, vol. 8, no. 2, pp. 155–178, May 1995.
- [39] J. M. Kahn, R. H. Katz, and K. S. J. Pister, "Next Century Challenges: Mobile Networking for "Smart Dust,"" in *Proc. of the 5th Annual ACM International Conference on Mobile Computing and Networking (MobiCom)*, 1999.
- [40] D. Estrin, L. Girod, G. Pottie, and M. Srivastava, "Instrumenting the World With Wireless Sensor Networks," in *Proc. of the International Conference on Acoustics, Speech and Signal Processing*, 2001.
- [41] J. Byers, J. Considine, M. Mitzenmacher, and S. Rost, "Informed Content Delivery Across Adaptive Overlay Networks," in *Proc. of ACM SIGCOMM*, 2002.
- [42] S. Fortune, J. Hopcroft, and J. Wyllie, "The Directed Subgraph Homeomorphism Problem," *Theoretical Computer Science*, vol. 10, no. 2, pp. 111–121, 1980.
- [43] K. Jain, M. Mahdian, and M. R. Salavatipour, "Packing Steiner Trees," in *Proc. of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2003.
- [44] Z. Li and B. Li, "Network Coding in Undirected Networks," in *Proc. of the 38th Annual Conference on Information Sciences and Systems (CISS)*, 2004.

TABLE I
FINITE BOUNDS ON THE CODING ADVANTAGE

Configuration	unicast	broadcast	multicast	group com- munication	bidirectional multicast	integral multicast	hypergraph
$\frac{\chi(N)}{\pi(N)} \leq$	1	1	2	2	$2(\alpha + 1)$	26	2

- [45] D. S. Lun, N. Ratnakar, M. Médard, R. Koetter, D. R. Karger, T. Ho, E. Ahmed, and F. Zhao, "Minimum-cost Multicast over Coded Packet Networks," *IEEE Trans. on Inform. Theory*, vol. 52, no. 6, pp. 2608–2623, 2006.
- [46] Y. Wu, P. A. Chou, Q. Zhang, K. Jain, W. Zhu, and S. Kung, "Network Planning in Wireless Ad Hoc Networks: A Cross-Layer Approach," *IEEE J. on Sel. Areas in Comm.*, vol. 23, no. 1, January 2005.
- [47] J. Yuan, Z. Li, W. Yu, and B. Li, "A Cross-Layer Optimizatin Framework for Multihop Multicast in Wireless Mesh Networks ," *IEEE J. on Sel. Areas in Comm.*, vol. 24, no. 11, November 2006.
- [48] Y. Cui, Y. Xue, and K. Nahrstedt, "Optimal Distributed Multicast Routing Using Network Coding," *SIGMETRICS Performance Evaluation Review*, vol. 32, no. 2, pp. 47–49, 2004.
- [49] M. Wang, Z. Li, and B. Li, "A High Throughput Overlay Multicast Infrastructure with Network Coding," in *Proc. of IEEE International Workshop on Quality of Service (IWQoS)*, Passau, Germany, June 2005.
- [50] M. Thimm, "On The Approximability Of The Steiner Tree Problem," in *Mathematical Foundations of Computer Science 2001*, Springer LNCS 2136, 678-689, 2001.

Lap Chi Lau is an Assistant Professor in the Department of Computer Science and Engineering at the Chinese University of Hong Kong. He received his B.Sc. degree in computer science from the Chinese University of Hong Kong, his M.Sc. degree in computer science from the University of Toronto under the supervision of Derek Corneil, and his Ph.D. degree in computer science from the University of Toronto under the supervision of Michael Molloy. His research interests are in combinatorial optimization, approximation algorithms and graph theory. For his work on Steiner tree packing, he received Machtey award for best student paper in IEEE conference FOCS 2004. For his thesis work, he received the Doctoral Prize 2007 from the Canadian Mathematical Society.

Zongpeng Li received his B.E. degree in Computer Science and Technology from Tsinghua University (Beijing) in 1999, his M.S. degree in Computer Science from University of Toronto in 2001, and his Ph.D. degree in Electrical and Computer Engineering from University of Toronto in 2005. He has been working as an Assistant Professor at the Department of Computer Science in the University of Calgary since August 2005. Zongpeng was named an Edward S. Rogers Sr. Scholar in 2004, won the Alberta Ingenuity New Faculty Award in 2007, was nominated for the Alfred P. Sloan Research Fellow in 2007, and received the Best Paper Award at the Ninth Passive and Active Measurement Conference (PAM) in 2008. His research interests are in computer networks, particularly in network optimization, multicast algorithm design, network game theory and network coding.

Baochun Li received his B.Engr. degree in 1995 from Department of Computer Science and Technology, Tsinghua University, Beijing, China, and his M.S. and Ph.D. degrees in 1997 and 2000 from the Department of Computer Science, University of Illinois at Urbana-Champaign. Since 2000, he has been with the Department of Electrical and Computer Engineering at the University of Toronto, where he is currently a Full Professor. He holds the Nortel Networks Junior Chair in Network Architecture and Services since October 2003, and the Bell University Laboratories Endowed Chair in Computer Engineering since August 2005. In 2000, he was the recipient of the IEEE Communications Society Leonard G. Abraham Award in the Field of Communications Systems. In 2005, he won the Best Paper Award at the Thirteenth IEEE International Workshop on Quality of Service (IWQoS). His research interests include large-scale multimedia systems, peer-to-peer networks, applications of network coding, and wireless networks.