

Network Coding in Undirected Networks

Zongpeng Li, Baochun Li

Department of Electrical and Computer Engineering

University of Toronto

{arcane, bli}@eecg.toronto.edu

Abstract—

Recent work in network coding shows that, it is necessary to consider both the *routing* and *coding* strategies to achieve optimal throughput of information transmission in data networks. So far, most research on network coding has focused on the model of *directed* networks, where each communication link has a fixed direction. In this paper, we study the benefits of network coding in *undirected* networks, where each communication link is bidirectional. Our theoretical results show that, for a single unicast or broadcast session, there are no improvements with respect to throughput due to network coding. In the case of a single multicast session, such an improvement is bounded by a factor of two, as long as half integer routing is permitted. This is dramatically different from previous results obtained in directed networks. We also show that multicast throughput in an undirected network is independent of the selection of the sender within the multicast group. We finally show that, rather than improving the optimal achievable throughput, the benefit of network coding is to significantly facilitate the design of efficient algorithms to compute and achieve such optimal throughput.

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.

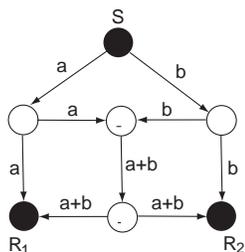


Fig. 1. With network coding, the achievable multicast throughput is 2.

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

While previous studies of network coding focus on directed networks with unidirectional links, 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 [3], we show that the coding advantage is always bounded by a constant factor of two in undirected networks. Our proof holds for half-integer routing, where each information flow being transmitted has either an integer or half-integer rate.

In addition, we prove that the achievable throughput is 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 more amenable to compute than optimal throughput without coding.

The remainder of the paper is organized as follows. We introduce related work on network coding in Sec. II, compare throughput with coding to other network parameters in Sec. III, discuss the source independence property of coded transmission in Sec. IV, and investigate the benefit of applying network coding from the perspective of computational complexity in Sec. V. Finally we conclude the paper and point out open problems in Sec. VI.

II. RELATED WORK

Ahlsvede *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.* [4] 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 $m \leq \lceil \log_2(r + 1) \rceil$ [5].

Li *et al.* [4] proposed the first code assignment algorithm, which performs an exponential number of vector independence tests. Sanders *et al.* [3] 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.

Zhu *et al.* [6] utilize network coding in designing their multicast scheme in overlay networks. In their empirical studies, throughput improvement over existing overlay multicast protocols with routing only are observed.

III. OPTIMAL THROUGHPUT IN UNDIRECTED NETWORKS: A COMPARISON STUDY

Network coding introduces a new dimension into 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 $\chi(N)$ to denote the maximum throughput of a network N containing a single transmission session. We compare $\chi(N)$ with other parameters that have been 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.

Packing refers to the computation of 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. *Strength* is a kind of partition connectivity of the network [7], and is denoted as $\eta(N)$. It is defined as the minimum ratio of $|E_c|/(p - 1)$, where p is the number of components the communication group is separated into, E_c is the set of inter-component

links, and the partition is required to have at least one source or destination node in each component. *Connectivity* refers to the minimum edge connectivity between a pair of nodes in the communication group, and is denoted as $\lambda(N)$.

We use a simple graph $G = (V, E)$ to represent the topology of a network, and use a rational function $C : E \rightarrow Q^+$ to denote link capacities. 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, nodes in the communication group are black, and relay nodes are white. We focus on fractional routing in this section, and will discuss integral routing in the following sections. 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.

A. Unicast

In an undirected network with a unicast session $N = \{(G(V, E), C: E \rightarrow Q, M = \{S, T\} \subseteq V)\}$, the packing number $\pi(N)$ becomes the number of edge-disjoint S - T paths. Throughput $\chi(N)$ is the maximum information rate that can be achieved in the $S \rightarrow T$ transmission. Strength $\eta(N)$ is now minimized over all simple cuts separating S and T , since no valid partition with more than two components exists. 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 .

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 only, it becomes identical to $\lambda(N)$, and both are equal to the min-cut between S and T . Furthermore, observe that uncoded throughput is always bounded by coded throughput, therefore $\pi(N) \leq \chi(N)$. Then, 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 [8]: *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 $N = \{(G(V, E), C : E \rightarrow Q^+, M = V = \{S, T_1, \dots, T_k\})\}$ 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. Throughput $\chi(N)$ is the maximum information rate from S to every other node in the network, simultaneously. Strength $\eta(N)$ is still as defined; just note that for a broadcast network, every partition is valid, since each component within 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 [8], [9], [10]: *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 $\pi(N) = \chi(N) = \eta(N)$.

By definition, $\eta(N) \leq \lambda(N)$, 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 [11]: *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 $\frac{1}{2}\lambda(N)$. Then by the results 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 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.

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

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

Proof: The fact that uncoded throughput is bounded by coded throughput again leads to $\pi(N) \leq \chi(N)$. And the partition 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.

The proof of $\frac{1}{2}\lambda(N) \leq \pi(N)$ contains one more step than in the broadcast case. We first need to remove relay nodes in $V - M$. In order to do so, we introduce Mader's Undirected Splitting Theorem [11]: *Let $G(V + z, E)$ be an undirected graph so that there is no node-cut incident to z and the degree $d(z)$ is even. Then there exists a complete splitting at z preserving the local edge-connectivities of all pairs of nodes in V .*

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

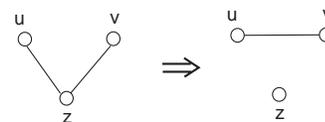


Fig. 2. A split off at node z .

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 connectivity 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 pairwise edge-connectivity 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 multicast nodes in M ; and (2) add a relay node that is a cut node into the multicast group M .

In order to meet the even node degree requirement in the undirected splitting theorem, we first double each link 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) 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. By Theorem 2, $\frac{1}{2}\lambda(N) \leq \pi(N)$ holds.

Finally, note that we obtained an integral transmission strategy achieving throughput $\frac{1}{2}\lambda(N)$, 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

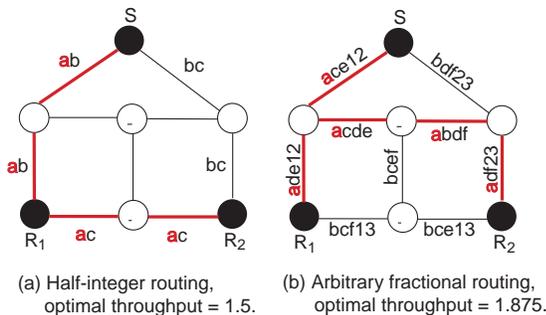


Fig. 3. Throughput without coding, for the example shown in Fig. 1.

Fig. 3 shows the optimal throughput without coding of the multicast session given in Fig. 1, assuming half-integer 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’ has been highlighted in bold edges. In (a), each tree has capacity 0.5; in (b), trees labeled with a letter have 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.

IV. SOURCE INDEPENDENCE IN UNDIRECTED NETWORKS

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 coded 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 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 Steiner set, regardless of which node in the Steiner set is the “sender”.

However, with network coding considered, it is less obvious whether the source independence property still holds. In Theorem 4, we show that it is the case.

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.

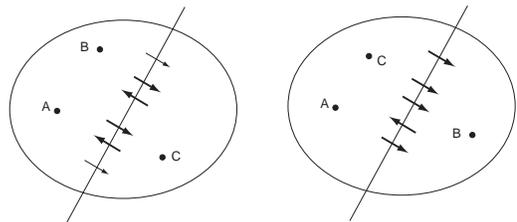


Fig. 4. Two scenarios of reversing the $A \rightarrow B$ flow. Darker links are being reversed.

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].

Suppose we exchange the sender and receiver roles between two multicast nodes A and B , and the optimal throughput before the exchange is f . Consider reversing the $A \rightarrow B$ 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. 4. 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 strength 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 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. 5 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.

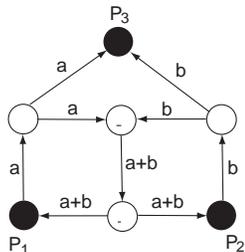


Fig. 5. Optimal transmission strategy after splitting and moving the information source, for the network shown in Fig. 1.

Corollary 4 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 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.

V. BENEFIT OF NETWORK CODING: THE COMPLEXITY PERSPECTIVE

In Sec. III, we have shown that the benefit of network coding is rather limited, from the perspective of throughput improvement. The coding advantage is 1 for both unicast and broadcast, and is at most 2 for multicast. Our empirical experiences show that even for multicast, the coding advantage is usually much smaller than the theoretical bound 2: for contrived networks in which network coding are necessary to achieve optimal throughput, the largest coding advantage value we observed is 1.125; for large scale random networks, the coding advantage we observed is always 1 [12].

Although its potential of improving throughput is limited, we have found that applying network coding has another major advantage: *reducing the computational complexity for computing and achieving the optimal throughput*. In this section, we illustrate this advantage with three examples, with both directed and undirected networks, fractional and integral routing considered.

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 [13][14], two receivers in an asynchronous file downloading session reconcile received data with each other [15], 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.* [16] 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. 6, 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].

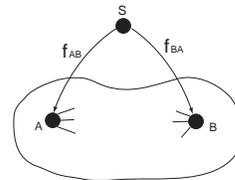


Fig. 6. Transforming the coded information exchange problem to the coded multicast problem.

B. Multicast with Integral Routing

We now switch back to undirected networks, and consider a multicast session with integral routing requirement. As discussed previously, the achievable multicast throughput equals to the steiner tree packing number. It has been shown that the steiner tree packing problem is NP-Complete [17], [18]; it is even worse in the integral case: there does not exist any known polynomial time algorithm that can approximate the problem to any constant ratio.

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. Even if the ratio $\frac{1}{2}$ is replaced with any other small constant ratio, the answer is still unknown. In network terminologies, this means that even if we know the multicast group has connectivity cx , for some arbitrarily large constant c , we still do not know whether throughput without coding can achieve x or not, with the integral routing requirement. On the other hand, we show that throughput with 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 [11]: 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 $u, v \in V$.

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 Directed Multicast Theorem, 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 [11] and code assignment [3].

C. Multicast with Fractional Routing

If we remove the integral routing requirement, the optimal multicast throughput problem without coding is equivalent to the fractional steiner tree packing problem, which is still NP-Complete and APX-hard [17]. However, once network coding is supported, the optimal throughput can then be computed efficiently. We have been able to formulate the computation of optimal multicast throughput with coding as a problem of computing rates of conceptual flows, which can be solved as a linear optimization problem. The interested reader is referred to our other work [12] for more details.

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 does not have known constant ratio approximations within polynomial time, while it has a 2-approximation in polynomial time with network coding. Therefore, although network coding may not lead to much higher values for the optimal throughput, it leads to more efficient algorithms that computes such optimal throughput.

VI. CONCLUSION AND FUTURE DIRECTIONS

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. We also show that achievable throughput with coding is independent of the selection of sender

within the communication group. Finally, we make the observation that applying network coding makes it possible to design efficient algorithms that compute and achieve the optimal transmission throughput.

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? 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? We intend to study these important questions as part of our future work.

REFERENCES

- [1] R. Ahlswede, N. Cai, S. R. Li, and R. W. Yeung, "Network Information Flow," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1204–1216, July 2000.
- [2] R. Koetter and M. Médard, "An Algebraic Approach to Network Coding," *IEEE/ACM Transactions on Networking*, vol. 11, no. 5, pp. 782–795, October 2003.
- [3] P. Sanders, S. Egner, and L. Tolhuizen, "Polynomial Time Algorithm for Network Information Flow," in *Proceedings of the 15th ACM Symposium on Parallelism in Algorithms and Architectures*, 2003.
- [4] S. Y. R. Li, R. W. Yeung, and N. Cai, "Linear Network Coding," *IEEE Transactions on Information Theory*, vol. 49, pp. 371, 2003.
- [5] T. Ho, D. R. Karger, M. Médard, and R. Koetter, "Network Coding From a Network Flow Perspective," in *Proceedings of ISIT*, 2003.
- [6] Y. Zhu, B. Li, and J. Guo, "Multicast with Network Coding in Application-Layer Overlay Networks," *IEEE Journal on Selected Areas in Communications, Special Issue on Recent Advances in Service Overlay Networks*, pp. 107–120, January 2004.
- [7] W. H. Cunningham, "Optimal Attack and Reinforcement of a Network," *Journal of the ACM*, vol. 32, pp. 549–561, 1985.
- [8] D. B. West, *Introduction To Graph Theory, 2nd Edition*, Prentice Hall, Upper Saddle River, New Jersey, 2001.
- [9] W. T. Tutte, "On the Problem of Decomposing a Graph Into n Connected Factors," *J. London Math. Soc.*, vol. 36, pp. 221–230, 1961.
- [10] C. St. J. A. Nash-Williams, "Edge-disjoint Spanning Trees of Finite Graphs," *J. London Math. Soc.*, vol. 36, pp. 445–450, 1961.
- [11] A. Frank, "Edge-connection of Graphs, Digraphs, and Hypergraphs," Tech. Rep., Enervary Research Group on Combinatorial Optimization, Budapest, Hungary, <http://www.cs.elte.hu/egres>, September 2000.
- [12] Z. Li, B. Li, D. Jiang, and L. C. Lau, "On Achieving Optimal End-to-End Throughput in Data Networks: Theoretical and Empirical Studies," Tech. Rep., Department of Electrical and Computer Engineering, University of Toronto, 2004.
- [13] J. M. Kahn, R. H. Katz, and K. S. J. Pister, "Next Century Challenges: Mobile Networking for "Smart Dust"," in *Proceedings of the 5th Annual ACM International Conference on Mobile Computing and Networking (MobiCom)*, 1999.
- [14] D. Estrin, L. Girod, G. Pottie, and M. Srivastava, "Instrumenting the World With Wireless Sensor Networks," in *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, 2001.
- [15] J. Byers, J. Considine, M. Mitzenmacher, and S. Rost, "Informed Content Delivery Across Adaptive Overlay Networks," in *Proc. of ACM SIGCOMM*, 2002.
- [16] S. Fortune, J. Hopcroft, and J. Wyllie, "The Directed Subgraph Homeomorphism Problem," *Theoretical Computer Science*, vol. 10, no. 2, pp. 111–121, 1980.
- [17] K. Jain, M. Mahdian, and M. R. Salavatipour, "Packing Steiner Trees," in *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2003.
- [18] M. Thimm, "On The Approximability Of The Steiner Tree Problem," in *Mathematical Foundations of Computer Science 2001, Springer LNCS 2136, 678-689*, 2001.