On the Integral Coding Advantage in Unit Combination Networks

David Holcomb[*], Eric D. Manley[†], Jason Grout[‡], Alex Hoyer[§]

Abstract

Network coding is a networking paradigm which allows network nodes to combine different pieces of data at various steps in the transmission rather than simply copying and forwarding the data. Network coding has various applications, and can be used to increase throughput, routing efficiency, robustness, and security. The original benefit that was demonstrated was improving the allowable transmission rate for a multicast session, and this application has been the focus of much research. One important parameter, the *coding advantage*, is the ratio of throughput with network coding to that without. The multicast networks that have a nontrivial coding advantage (i.e., coding advantage greater than 1) all seem to contain a substructure called the combination network which has a source, n relay nodes, and $\binom{n}{k}$ receivers in which each receiver is adjacent to a unique subset of k relay nodes. The coding advantage in combination networks has previously been determined for networks with fractional routing. In this paper, we address integral routing, which is more appropriate for networks (like optical wavelength-division-multiplexing networks) which allow only coarse-grained subdivision of the available bandwidth on any given channel. We give exact formulas for the integral coding advantage in both directed and undirected networks. For directed networks, we show that the coding advantage is $k/\left\lfloor \frac{n}{n-k+1} \right\rfloor$. For undirected networks, we show that the coding advantage is $k/(k-1)$. The latter result fits with conjectures that the integral coding advantage in any undirected network is bounded above by 2.

1 Introduction

Network coding is a relatively recent advancement in information theory proposed by Ahlswede *et al.* [2] that has the potential to increase throughput, robustness, routing efficiency, and security in data networks. Without network coding, nodes on a transmission route copy data on incoming channels and

[*]dholcom2@illinois.edu, University of Illinois, Urbana-Champaign, IL

[†]eric.manley@drake.edu, Drake University, Des Moines, IA

[‡]jason.grout@drake.edu, Drake University, Des Moines, IA

[§]ahoyer3@math.gatech.edu, Georgia Institute of Technology, Atlanta, GA

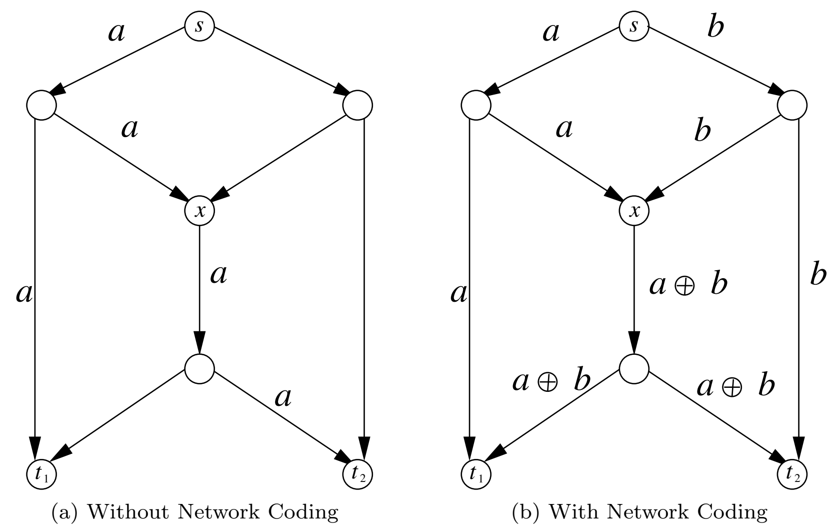

Figure 1: Network coding with multicast communication. Node s is the source, t_1 and t_2 are receivers, and each edge has capacity 1.

forward it along outgoing channels. With network coding, the nodes instead transmit functions of the symbols received on incoming channels. In this paper, we focus on the single-source multicast problem in which a single source transmits a message to many receivers. Ahlswede *et al.* proved that if a transmission rate is achievable for each receiver as a unicast, then network coding can guarantee that same rate to the entire group as a multicast transmission [2]. With linear network coding, the coding functions are limited to linear combinations of incoming symbols, which Li *et al.* showed is sufficient to achieve the maximum multicast flow rate [6]. Each receiver decodes the original symbols by solving the resulting system of linear equations received on its incoming channels.

Figure 1 presents an example of how network coding works on the butterfly network, a canonical structure first presented by Ahlswede *et al.* [2]. In this example, we have a 7-node directed network with one source node s and two receiver nodes t_1 and t_2. Each edge has capacity 1, and the message to be sent is of size 2 (represented as symbols a and b). Thus, in any given single time interval, to transmit both symbols with traditional routing, we need two edge-disjoint directed Steiner trees. Note that if we set up any single directed Steiner tree for transmitting symbol a (an example is shown in Figure 1a), any possible directed Steiner tree for b is blocked. Without network coding, the only way to transmit both symbols requires twice as much time (i.e. send a and then send b) or necessitates a network upgrade. However, if network coding is allowed, node x can compute $a \oplus b$ (with $\oplus$ denoting XOR, bitwise logical exclusive-OR) and transmit the messages as shown in Figure 1b. The symbol $a \oplus b$ has the

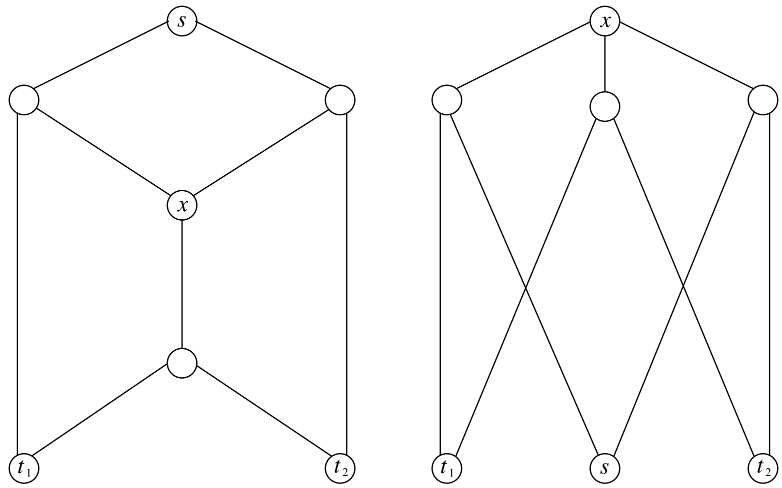

(a) Undirected Butterfly Network (b) $C_{3,2}$

Figure 2: The butterfly network redrawn in the style of a combination network.

same size as both a and b, and so the whole message can be transmitted in a single time interval. Receiver t_1 receives a and $(a \oplus b)$ on its incoming channels and may recover b by computing $a \oplus (a \oplus b) = b$. Similarly, t_2 receives both b and $(a \oplus b)$ and can recover a. Thus, network coding allows us to double our throughput in this network.

Network coding research is often focused on exploring the potential benefits of network coding. One way this is measured is by the *coding advantage* of a network: the ratio of maximum throughput with network coding to maximum throughput without network coding. Previous research has suggested that only certain kinds of networks experience an increase in throughput due to network coding. In fact, some experimental studies have found that very few randomly-generated networks benefit from network coding [8, 10]. The multicast networks that have a nontrivial coding advantage (i.e., coding advantage greater than one) all seem to contain a substructure called the combination network [9, 13], which is the class of networks studied in this paper. We will define combination networks formally in Section 2; however, we note that combination networks are simply a generalization of the butterfly network (see Figure 2). That is, it appears that the network shown in Figure 1 contains the primary structure which allows an improvement to single-source multicast with network coding. Thus, understanding butterfly and combination networks is an important step in coming to a greater understanding of network coding in general.

The coding advantage depends in part on how the capacity of an edge can be split between signals. For instance, with *integral routing*, an edge's capacity is an integer and can only be allocated in integral portions. With *fractional routing*, an edge's capacity can be divided in arbitrary, even non-integral proportions. Li and Li showed that in the fractional case, general undirected networks with a single unicast or broadcast have a trivial coding advantage (i.e. the coding advantage is 1). They also showed that the coding advantage for fractional routing on an arbitrary undirected network is bounded above by 2 [7]. Agarwal and Charikar provided Steiner tree linear programs whose integrality gap is equivalent to the fractional coding advantage [1]. Noting the apparent exceptional status of combination networks, Maheshwar, Li, and Li considered the reduction in multicast *cost* (which is important when communication channels have associated costs). They found that *cost advantage* in uniform-cost combination networks was $(\binom{n}{k} + n - k + 1)/(\binom{n}{k} + \frac{n}{k})$, and showed that this is an upper bound on the coding advantage in a related capacitated network [9].

However, not much attention has been given to integral coding advantage. In high-bandwidth, coarse-grained networks like optical wavelength-division-multiplexing (WDM) networks, each network channel has large bandwidth, and it may not be possible to split its capacity among many messages in an arbitrary way, so integral routing is a more appropriate model than fractional routing. The application of network coding to optical WDM networks has been the subject of recent research [4, 5, 11, 12], so it is valuable to examine the coding advantage problem in the integral routing case.

In this paper, we examine the coding advantage with integer routing for unit combination networks, i.e., combination networks where all edges have capacity one. We find closed forms for the coding advantage in both the directed and undirected cases. These closed forms help us determine exactly the conditions for one of these networks to have nontrivial coding advantage. In the directed case, the coding advantage is not bounded above. This result complements Jaggi *et al.*'s result that the coding advantage is unbounded from above in combination networks and increases with $\Omega(\log |V|)$, where V is the vertex set [3]. We also find that, in the undirected case, the coding advantage in unit combination networks is bounded above by 2, with the upper bound being tight. This complements Li and Li's result that, allowing fractional routing (or, at least half-integer routing), the coding advantage in undirected networks is bounded above by 2.

The rest of this paper is organized as follows. In Section 2, we discuss unit combination networks and provide some useful lemmas and notation. The coding advantage for directed networks is considered in Section 3, and undirected networks are considered in Section 4. We conclude in Section 5.

2 Unit Combination Networks

In this section, we formally define the class of networks which we are investigating. We also provide some observations on how the relevant throughput parameters are determined and some lemmas that will be useful in proving our main results. The umbrella class of networks we are studying is combination networks which are defined as follows.

Definition 1. *Let n, k be integers such that $n > k \geq 1$. The* combination network $C_{n,k} = (V, E, s \in V, T \subseteq V)$*, with nodes V, edges E, source s, receivers T, and relay nodes $V \setminus (T \cup \{s\})$, is constructed with n relay nodes and $\binom{n}{k}$ receiver nodes. The edges set E contains an edge from s to each relay node as well as edges between the relays and T such that each of the $\binom{n}{k}$ nodes in T is adjacent to a unique subset of the relay nodes of size k.*

In directed combination networks, all edges incident on the source are directed away from the source and all edges incident on a receiver are directed towards the receiver. For examples of three combination networks, see Figure 3. In this paper we focus on combination networks which are *unit networks*.

Definition 2. *A* unit network $G = (V, E, s, T)$ *is a network with vertex set V, edge set E with all edges having capacity one, source s, and receiver set T.*

Before defining the integral coding advantage for unit networks, we make the following observations about throughput in unit networks.

Observation 1. *The maximum throughput with network coding in a directed unit network $G = (V, E, s, T)$ is*

$$\chi(G) = \min\{p_t | t \in T\}$$

where p_t is the maximum number of edge-disjoint directed paths from s to t.

Proof. This is a special case of Ahlswede *et al.*'s theorem on the maximum throughput with network coding [2]. □

If the graph is undirected, $\chi(G)$ is the maximum throughput over all possible orientations of G.

Observation 2. *The maximum integral throughput without network coding in a directed unit network $G = (V, E, s, T)$ is $\pi(G)$, the maximum number of edge-disjoint directed integral trees containing directed paths from s to T that can be packed into G.*

If the graph is undirected, $\pi(G)$ is the maximum throughput over all possible orientations of G. In the undirected case, without loss of generality, the maximum can be taken over all Steiner trees. In the directed case, we can likewise take the maximum over all the analogous directed Steiner trees. When it is clear from context, we use the term Steiner tree in both cases.

We now define the *integral coding advantage* in unit networks and give an observation that network coding generalizes traditional routing.

Definition 3. *The* integral coding advantage *for unit network $G = (V, E, s, T)$ is $\chi(G)/\pi(G)$, the ratio of the maximum throughput with network coding to the maximum throughput without network coding.*

Observation 3. *Because each tree that can be packed into a graph also contains a path from the source to each receiver, the integral coding advantage of a unit network is always at least one.*

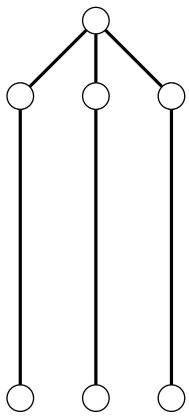

(a) $C_{3,1}$

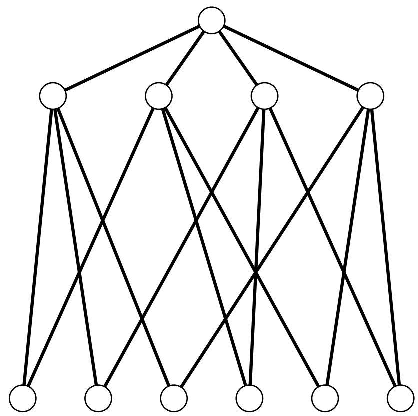

(b) $C_{4,2}$

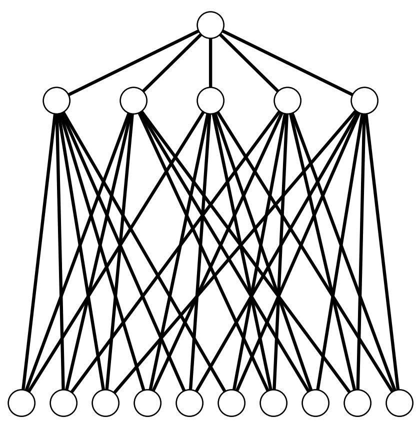

(c) $C_{5,3}$

Figure 3: Combination networks. The source is at the top, the relays are in the middle-top, and the receivers are at the bottom.

Now that we have given the necessary definitions, we provide a series of lemmas which will be used in the proofs of our main results.

Lemma 1. *In $C_{n,k}$, any relay node is adjacent to $\binom{n-1}{k-1}$ receiver nodes.*

Proof. There are $\binom{n-1}{k-1}$ ways to choose k of the relay nodes where one of those chosen is some specific relay node. $\square$

Lemma 2. *In any unit combination network $\chi(C_{n,k}) = k$.*

Proof. For any receiver node v, there exists a set of k relay nodes adjacent to v, and every relay node in this set is adjacent to the source. Therefore, there exist k edge-disjoint paths from the source to v. By Observation 1, the maximum throughput is k. $\square$

Lemma 3. *The coding advantage for $C_{n,1}$ is one.*

Proof. $C_{n,1}$ is itself the only possible Steiner tree (e.g. Figure 3a). By Lemma 2, the coding advantage is one. $\square$

Lemma 4 (Smith *et al.* [13])**.** *Any Steiner tree in a combination network must reach at least $n - k + 1$ relays.*

Proof. Consider for any subset of $i \in \{1, 2, ..., n\}$ relay nodes, there are $\binom{n-i}{k}$ receivers adjacent to none of the relays in the subset. We need $\binom{n-i}{k} = 0$ in order to cover all receivers, which only happens when $n - i < k$, which means $i \geq n - k + 1$. Then, the minimum number of relays that must be in a tree for that tree to contain all receivers is $n - k + 1$. $\square$

3 Directed Combination Networks

We now give one of our main results, an exact formula for the integral coding advantage in unit directed combination networks.

Theorem 1. *In the case of unit directed combination networks $C_{n,k}$, the integral coding advantage is $\frac{k}{\left\lfloor \frac{n}{n-k+1} \right\rfloor}$.*

Proof. Since all relay-receiver edges are directed away from the relays, the only way any tree can reach a relay is through the source-relay edges, and there are only n such edges. From Lemma 4, any completed Steiner tree needs to reach at least $n - k + 1$ relays. Then, in order to give each tree at least $n - k + 1$ source-relay edges, there can be at most $\left\lfloor \frac{n}{n-k+1} \right\rfloor$ trees. If we construct the $\left\lfloor \frac{n}{n-k+1} \right\rfloor$ trees by giving them disjoint subsets of $n - k + 1$ source-relay edges, and all the relay-receiver edges incident on their newly connected relays, all trees can be connected to all receivers, so $\pi(C_{n,k}) = \left\lfloor \frac{n}{n-k+1} \right\rfloor$. With Lemma 2, this is sufficient to show that the coding advantage in a unit directed combination network $C_{n,k}$ is $\frac{k}{\left\lfloor \frac{n}{n-k+1} \right\rfloor}$. $\square$

We will now show that the integral coding advantage for unit directed combination networks is unbounded from above. Jaggi *et al.* also gave this result in [3], although they did not give the exact coding advantage.

Corollary 1. *Integral coding advantage is unbounded from above in the case of unit directed combination networks.*

Proof. Consider the unit combination networks $C_{2k-1,k}$ for $k \geq 1$. The coding advantage for such a network is

$$\frac{k}{\left\lfloor \frac{n}{n-k+1} \right\rfloor} = \frac{k}{\left\lfloor \frac{2k-1}{2k-1-k+1} \right\rfloor} = \frac{k}{\left\lfloor \frac{2k-1}{k} \right\rfloor}$$
$$\geq \frac{k}{\frac{2k-1}{k}} = \frac{k^2}{2k-1} > \frac{k^2}{2k} = \frac{k}{2}$$

Since $\lim_{k\to\infty} \frac{k}{2} = \infty$, there exists $C_{2k-1,k}$ with arbitrarily large coding advantage. □

This confirms that unit directed combination networks are interesting because they lead to a nontrivial coding advantage. In fact, there is a nontrivial coding advantage for all $k \geq 2$.

Corollary 2. *In any unit directed combination network $C_{n,k}$ with $k \geq 2$, $\pi(C_{n,k}) < k$, so the coding advantage is nontrivial.*

Proof. Let $C_{n,k}$ be a unit directed combination network with $k \geq 2$. From the proof of Theorem 1, $\pi(C_{n,k}) = \left\lfloor \frac{n}{n-k+1} \right\rfloor$. Assume that $\left\lfloor \frac{n}{n-k+1} \right\rfloor \geq k$. Then

$$\frac{n}{n-k+1} \geq k$$

Because $n-k+1 > 0$ (since we require $n > k$):

$$n \geq kn - k^2 + k$$

$$k^2 - k \geq n(k-1)$$

$$k(k-1) \geq n(k-1)$$

Because $k \geq 2$, we have $k \geq n$. Thus $C_{n,k}$ is not a combination network, which is a contradiction (Definition 1). Therefore $\pi(C_{n,k}) < k$. □

4 Undirected Combination Networks

We now address the case of unit undirected combination networks. Note that, in an undirected unit combination network, a tree can be oriented to send a signal along a relay-receiver edge "backwards" from the receiver to the relay. Thus, there is a greater potential for packing in Steiner trees than in the directed case. However, we will show that combination networks still have a nontrivial coding advantage in this case, but that the coding advantage is bounded above by 2.

Even though the edges can be oriented in either direction, referring to a relay-receiver edge as being used in the backward or forward direction is sometimes helpful, so we define the use of these terms as follows:

Definition 4. *A* backward edge*, when referring to a combination network, is an edge incident on a receiver that is used by a tree to send a signal away from that receiver.*

Definition 5. *A* forward edge*, when referring to a combination network, is an edge incident on a receiver that is used by a tree to send a signal toward that receiver.*

In the following two lemmas, we show that $\pi(C_{n,k}) = k-1$ when $k \geq 2$ for unit undirected combination networks, which will give us a formula for the coding advantage.

Lemma 5. *In any unit undirected combination network $C_{n,k}$ where $k \geq 2$, $\pi(C_{n,k}) < k$.*

Proof. Assume to contradiction that k edge-disjoint Steiner trees have been packed into some unit undirected combination network $C_{n,k}$. Each of the k trees has an edge incident on each receiver. However, each receiver is only incident on k edges, therefore all edges incident on any receiver node belong to distinct trees. Thus, no tree may, after reaching a receiver node from a relay-receiver edge, use another relay-receiver edge backwards to reach a relay node, as this would deprive some other tree of the edge, preventing it from reaching the receiver. Because no tree can use any relay-receiver edge to go backwards to a relay, the only way a tree can reach a relay is through the source-relay edges. The graph then degenerates into the directed case, and $\pi(C_{n,k}) = \left\lfloor \frac{n}{n-k+1} \right\rfloor$ (Theorem 1). From Corollary 2, $\pi(C_{n,k}) < k$. □

Lemma 6. *In any unit undirected combination network $C_{n,k}$ where $k \geq 2$, $\pi(C_{n,k}) \geq k-1$.*

Proof. We consider four cases.

Case 1: $k = 2$. Consider the unit undirected combination network $C_{n,2}$. It is possible to pack $k-1 = 2-1 = 1$ edge-disjoint Steiner tree into this graph because the graph is connected. Any spanning tree will also be a Steiner tree. Thus $\pi(C_{n,2}) \geq 1$.

Case 2: $k = n-1$ and $k \neq 2$. Consider the unit undirected combination network $C_{n,n-1}$ (that is, where $k = n-1$ and $k \neq 2$). We can pack in $k-1 = n-2$ Steiner trees as follows: Give each tree one distinct source-relay edge. At this point, each tree contains the source and one distinct relay. Then, give each tree all the relay-receiver edges incident on its included relay node. At this point, there are $n-(n-2)=2$ relay nodes not yet included in any tree. Call these relays x and y. We then have

- $\binom{n-(n-2)}{n-1-(n-2)} = \binom{2}{1} = 2$ receivers included in all trees and adjacent to exactly one of x and y.
- $\binom{n-2}{n-1-2} = \binom{n-2}{n-3} = n-2$ receivers adjacent to both x and y. These receivers are adjacent to $(n-1)-2 = n-3$ of the relays other than x and y, and are thus not included in exactly one tree each. Give the trees backward edges incident on these $n-2$ receivers so that each of x and y can be connected to $\lfloor (n-2)/2 \rfloor$ of the trees. Note that this leaves one unused edge incident on each of these receivers and incident to the other of x and y. Where n is odd, one tree contains neither x nor y, so give this tree the source-relay edges to x and y. Finally, propagate the trees on unused edges from x and y to receivers that are not yet included in all $n-2$ trees.

All receivers are now included in all $n-2$ trees showing that $\pi(C_{n,n-1}) \geq n-2$

Case 3: $k = n-2$ and $k \neq 2$. Consider the unit undirected combination network $C_{n,n-2}$ (that is, where $k = n-2$ and $k \neq 2$). We show that we can pack in $k-1 = n-3$ Steiner trees. Furthermore, we will also ensure the following property:

Property 1. *The $k-1$ trees are configured in such a way that if we select any $k-1$ relay nodes, we can pick one tree from each of the $k-1$ relay nodes so that all trees are represented.*

This property will serve as a necessary invariant in the inductive proof used in Case 4 which will prove the lemma in all other cases.

Begin as in Case 2 by giving each tree one distinct source-relay edge. At this point, each tree includes the source and one distinct relay. Then, give each tree all the relay-receiver edges incident on its included relay. At this point, there are $n-(n-3) = 3$ relay nodes not yet included in any tree. Call these relays x, y, and z. We then have

- $\binom{n-(n-3)}{n-2-(n-3)} = \binom{3}{1} = 3$ receivers included in all trees.
- $\binom{3}{2}\binom{n-3}{n-2-2} = 3\binom{n-3}{n-4} = 3(n-3)$ receivers adjacent to exactly two of x, y, and z; and therefore, adjacent to $n-4$ of the initial relay nodes (note that since $k > 2$ and $n = k+2$ we have $n \geq 5$). These receivers are each not included in exactly one tree and have two unused incident edges. Divide these receivers into three disjoint sets based on which pair of x, y, and z they're adjacent to: A is the set adjacent to x and y, B is the set adjacent to x and z, and C is the set adjacent to y and z. Note that each of A, B, and C contains $n-3$ receivers. Because each receiver is adjacent to a different subset of $n-2$ of the relays, within each of A, B, and C, all receivers are adjacent to different sets of relay nodes. Then, we can give each tree a backward edge from one of the the receivers in A to x, from

B to y, and from C to z. Then, all of x, y, and z are included in all trees. Furthermore, each receiver in A, B, and C still has one unused incident edge also incident on one of x, y, or z. Use these edges to connect all the receivers in A, B, and C to all trees.

- $\binom{n-3}{n-2-3} = \binom{n-3}{n-5}$ receivers adjacent to all three of x, y, and z. These receivers are adjacent to $n-5$ of the initial relays, and thus are not included in exactly two of the $n-3$ trees. Because each of x, y, and z are now included in all trees, we can use the unused edges incident on x, y, and z, and on these receivers to connect all trees to these receivers.

Thus, all receivers are now included in all $n-3$ trees showing that $\pi(C_{n,n-2}) \geq n-3$. Note also that $n-3$ of the relays are (uniquely) included in exactly one of the trees and the other three relays are included in all trees. Thus, selecting any $n-3$ relays, we can find a representative tree for each relay that allows all $n-3$ trees to be simultaneously represented. This maintains Property 1, so we can use Case 3 as a base case for an inductive proof for all remaining cases.

Case 4: $k < n-2$. Now, suppose a unit undirected combination network $C_{n,k}$ for $k \leq n-2$ has a Steiner tree packing (where each tree contains the source and all receivers) with $k-1$ trees satisfying Property 1. Note that by Case 3, this holds for $k = n-2$.

Then, consider the graph $C_{n+1,k}$. For such a graph, $k \leq n-2 \Rightarrow k \leq (n+1)-2$, and moreover, there exists some induced subgraph isomorphic to $C_{n,k}$. Call this subgraph G, and include in it the source node, and some combination of n of the relays and their $\binom{n}{k}$ adjacent receivers. Call the $(n+1)-n = 1$ relay not in this subgraph $r_{\bar{G}}$. Call the receivers not in this subgraph $T_{\bar{G}}$. By the inductive hypothesis, $k-1$ edge-disjoint Steiner trees can be packed into G such that Property 1 holds. Because any relay is adjacent to $\binom{n}{k-1}$ of the receivers, and because $\binom{n}{k-1} = \frac{n!}{(k-1)!(n-k+1)!} = \frac{n(n-1)!}{(n-k+1)(k-1)!(n-k)!} = \frac{n}{n-k+1}\binom{n-1}{k-1} > \binom{n-1}{k-1}$, each of the relays in G is adjacent to at least one receiver in $T_{\bar{G}}$. Furthermore, all receivers in $T_{\bar{G}}$ are adjacent to $r_{\bar{G}}$, and to $k-1$ of the relays in G. Then, because Property 1 holds, each receiver in $T_{\bar{G}}$ can use its edges incident on the $k-1$ relays in G to connect it to all trees. We can use the edges connecting $r_{\bar{G}}$ with the receivers in $T_{\bar{G}}$ to connect all trees to $r_{\bar{G}}$. Now, all receivers are included in all trees, and Property 1 has been maintained because $r_{\bar{G}}$ is included in all trees. Therefore, our inductive hypothesis is maintained.

Putting all four cases together, we conclude that $\pi(C_{n,k}) \geq k-1$ for all $n > k \geq 2$. $\square$

We now have a formula for the coding advantage in unit undirected combination networks.

Theorem 2. *The coding advantage for a unit undirected combination network* $C_{n,k}$ *is 1 if* $k = 1$ *and* $\frac{k}{k-1}$ *otherwise.*

Proof. By Lemma 3, the coding advantage is 1 when $k = 1$. By Lemma 5, $\pi(C_{n,k}) < k$, and by Lemma 6, $\pi(C_{n,k}) \geq k-1$, so $\pi(C_{n,k}) = k-1$. Then, by Lemma 2, $\chi(C_{n,k}) = k$. Therefore, $\chi(C_{n,k})/\pi(C_{n,k}) = k/(k-1)$ when $k \geq 2$. $\square$

The coding advantage is largest when k is small and is bounded above by 2 which is achieved only when $k = 2$. Li and Li showed that in the case of general undirected graphs allowing half-integer routing that the coding advantage is bounded above by two [7]. Since combination networks represent the known class of graphs with nontrivial coding advantages, this fits with the conjecture that the integral coding advantage is bounded above by 2 for general graphs, although, to the best of our knowledge, this question remains open.

5 Conclusions

The flow in a combination network structure is the only known flow structure to produce a nontrivial coding advantage. We derived formulas to describe integral coding advantage for unit combination networks in both the directed (Theorem 1) and undirected (Theorem 2) cases. In the case of undirected combination networks, the coding advantage is bounded above by two. In the directed case, the coding advantage is not bounded above. Furthermore, as corollaries, we have shown that, except for the case where $k = 1$, all unit combination networks have a nontrivial integral coding advantage. This represents an important step in understanding when network coding is beneficial in coarse-grained networks and contributes to a better understanding of network coding in general.

References

[1] A. Agarwal and M. Charikar. On the advantage of network coding for improving network throughput. In *IEEE Information Theory Workshop*, pages 247–249, October 2004.

[2] R. Ahlswede, Ning Cai, S.-Y.R. Li, and R.W. Yeung. Network information flow. *Information Theory, IEEE Transactions on*, 46, July 2000.

[3] S. Jaggi, P. Sanders, P.A. Chou, M. Effros, S. Egner, K. Jain, and L.M.G.M. Tolhuizen. Polynomial time algorithms for multicast network code construction. *Information Theory, IEEE Transactions on*, 51, June 2005.

[4] A. E. Kamal. A generalized strategy for 1+N protection. In *Proceedings of the IEEE International Conference on Communications (ICC2008)*, 2008.

[5] Minkyu Kim, Muriel Médard, and Una-May O'Reilly. Network coding and its implications on optical networking. In *Optical Fiber Communication Conference*, page OThO3. Optical Society of America, 2009.

[6] S.-Y. R. Li, R. W. Yeung, and N. Cai. Linear network coding. *IEEE Trans. Inf. Theory*, 49(2):371 – 381, February 2003.

[7] Zongpeng Li and Baochun Li. Network coding in undirected networks. In *38th Annual Conference on Information Sciences and Systems (CISS)*, March 2004.

[8] Zongpeng Li, Baochun Li, Dan Jiang, and Lap Chi Lau. On achieving optimal throughput with network coding. In *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, March 2005.

[9] Shreya Maheshwar, Zongpeng Li, and Baochun Li. Bounding the power of combination network coding in undirected networks. *Information Theory, IEEE Transactions on*, Preprint.

[10] Eric D. Manley. *Network Coding for WDM All-Optical Networks*. PhD thesis, University of Nebraska, 2009.

[11] Eric D. Manley, Jitender S. Deogun, Lisong Xu, and Dennis R. Alexander. All-optical network coding. *IEEE/OSA Journal of Optical Communications and Networking*, 2(4):175 – 191, April 2010.

[12] R.C. Menendez and J.W. Gannet. Efficient, fault-tolerant all-optical multicast networks via network coding. In *Conference on Optical Fiber communication/National Fiber Optic Engineers Conference (OFC/NFOEC)*, pages 1–3, February 2008.

[13] A. Smith, B. Evans, Zongpeng Li, and Baochun Li. The cost advantage of network coding in uniform combinatorial networks. In *Sensor, Mesh and Ad Hoc Communications and Networks Workshops, 2008. SECON Workshops '08. 5th IEEE Annual Communications Society Conference on*, 2008.