# On the Advantage of Network Coding for Improving Network Throughput

**(extended abstract)**

Amit Agarwal  Moses Charikar[1]
Dept. of Computer Science
Princeton University
Princeton, NJ 08544
{aagarwal,moses}@cs.princeton.edu

*Abstract* — **Given a data network with link capacities, we consider the throughput of the network for a multicast session involving a source node and a given set of terminals. It is known that network coding can improve the throughput of the network. We study the coding advantage, i.e. the ratio of the throughput using network coding to that without using network coding. We show that the maximum coding advantage for a given network is equal to the integrality gap of certain linear programming (LP) formulations for Steiner tree. This holds for both directed as well as undirected networks. For directed networks, the coding advantage is equal to the integrality gap of the directed Steiner tree LP formulation; for undirected networks, the coding advantage is equal to the integrality gap of the bidirected cut LP formulation for Steiner tree. This relates the coding advantage to well studied notions in combinatorial optimization. Further, this connection improves the known bounds on the coding advantage for both undirected as well directed networks.**

## I. Introduction

Given a network with capacities on links, a fundamental problem is to compute the maximum multicast throughput possible for communication between a source node and a set of receivers. Traditionally, intermediate nodes in the network are allowed to merely store and forward information packets. In this model, the multicast throughput problem is computationally hard to solve optimally. In fact, this problem is equivalent to the problem of maximum fractional steiner tree packing, which is NP-hard [3]. In a seminal paper, Ahlswede etal [1] introduced the notion of *network coding*, where intermediate nodes are allowed to encode and decode messages they receive. In this model, they gave a simple characterization of the maximum throughput possible in a directed network for a multicast session between a source and a given set of receivers, and demonstrated that network coding could increase the throughput. Sanders etal [7] demonstrated examples where the gap between the throughput using network coding to that without using coding was $\Omega(\log n)$, where $n$ is the number of receivers. Recently, Li etal [5] studied the multicast throughput of undirected networks and gave a linear programming formulation to compute the throughput in the undirected case. Similar techniques were also used recently by Kramer and Savari [4] to bound the multicast capacity of undirected networks with network coding. [5] also studied the *coding advantage*, i.e. the ratio of the throughput using network coding to that without using coding. For undirected networks, [5] gave various examples with coding advantage $> 1$, the best value being $9/8$ and proved that it is no more than 2.

In this paper we show an interesting connection between the coding advantage and the integrality gap of linear programming formulations for minimum weight Steiner tree. Since the minimum weight Steiner tree problem is NP-hard for both undirected and directed networks, polynomial time solvable LP relaxations of Steiner tree are commonly used to obtain a lower bound on the optimal Steiner tree weight. The quality of the bound provided by the LP relaxation is measured by its integrality gap, i.e. the ratio between the optimum Steiner tree weight and the optimum solution to the LP relaxation. We show that for undirected networks, the maximum coding advantage is equal to the integrality gap of the bidirected cut relaxation for the undirected Steiner tree problem [6]. For directed networks, we show that the coding advantage is equal to the integrality gap of a natural LP formulation of directed Steiner tree.

Our results improve on the best bounds known for the maximum coding advantage in both directed and undirected networks. Using an integrality gap example due to Goemans (presented in [8]), we show a family of undirected networks with coding advantage approaching $8/7$, improving on the $9/8$ bound from [5]. For directed networks, using a recent integrality gap example for directed Steiner tree given by Halperin etal. [2], we show that the coding advantage can be $\Omega((\log n/\log\log n)^2)$ where $n$ is the number of receivers. It is interesting to note that determining the integrality gaps of these LP formulations for undirected and directed Steiner tree are well studied and interesting open problems in the computer science optimization literature. In the undirected case, it is known that the gap is at most 2, and an upper bound of $3/2$ is known for a special class of graphs caled quasi-bipartite graphs [6]. Determining whether the gap is strictly lower than 2 is a major open problem. In the directed case, establishing whether the gap is upper bounded by a polygarithmic function is a major open problem.

There are interesting parallels between the coding advantage examples obtained in the network coding community and the integrality gap examples produced in the computer science optimization community. Some of the coding advantage examples given in [5] appear in [6] as integrality gap examples. The integrality gap example of Goemans (described in [8]) relies on a gadget which is exactly the same as a simple, well known example to illustrate the advantage of network coding,

[1]This work was supported by NSF ITR grant CCR-0205594, DOE Early Career Principal Investigator award DE-FG02-02ER25540, NSF CAREER award CCR-0237113 and an Alfred P. Sloan Fellowship.

first given in [1]. Finally, the example given to show the coding advantage in directed networks in [7] is exactly the same as a well known integrality gap example for set cover, a special case of directed Steiner tree.

In this extended abstract, we present our results for the undirected case. We omit a detailed exposition of the results for the directed case; the details are similar to those in the undirected setting.

## II. Preliminaries

We represent the input network as an undirected graph $G = (V, E)$. Let $c : E \to R^+$ be an assignment of non-negative capacities to the edges. We use $c_e$ to denote the capacity of edge $e \in E$. For given edge capacities $c$, we will investigate the network throughput for a multicast session between source $m_0$ and receivers $m_1, \ldots, m_k$. The amount of information that can be sent with network coding is denoted by $\chi(G, c)$. The amount of information that can be sent without network coding is denoted by $\Pi(G, c)$ and is the same as the fractional steiner tree packing number for $G$. Define the coding advantage to be the ratio $\frac{\chi(G,c)}{\Pi(G,c)}$. Let $\tau$ be the set of all possible Steiner trees for $G$ connecting the given set of terminals $\{m_0, \ldots, m_k\}$.

The Steiner packing number is given by the following linear program: (there is a variable $x_t$ for every possible Steiner tree $t \in \tau$)

$$\begin{aligned} \max \quad & \textstyle\sum_{t\in\tau} x_t \\ \sum_{t\in\tau : e\in t} x_t \le c_e, \quad & \forall e \in E \\ x_t \ge 0, \quad & \forall t \in \tau \end{aligned}$$

The dual of the above linear program is as follows:

$$\begin{aligned} \min \quad & \textstyle\sum_{e\in E} c_e y_e \\ \sum_{e\in t} y_e \ge 1, \quad & \forall t \in \tau \\ y_e \ge 0, \quad & \forall e \in E \end{aligned}$$

Note that $y_e$ are the variables in the above LP.

As given in [5], the following LP (called the cFlow LP in [5]) computes the optimal throughput achievable with network coding. To give the LP, we use similar notation as in [5] (modified for consistency).

$$\begin{aligned} \max \quad & f^* \\ c_a \ge 0 \quad & \forall a \in D \\ c_{e_1} + c_{e_2} = c_e \quad & \forall e \in E \\ f^i(a) \le C(a) \quad & \forall i \in [1 \ldots k], \forall a \in D \\ f^i_{in}(v) = f^i_{out}(v) \quad & \forall i \in [1 \ldots k], \forall v \in V - \{m_0, m_i\} \\ f^i_{in}(m_0) = 0 \quad & \forall i \in [i \ldots k] \\ f^i_{out}(m_i) = 0 \quad & \forall i \in [1 \ldots k] \\ f^* = f^i_{in}(m_i) \quad & \forall i \in [1 \ldots k] \end{aligned}$$

The optimum of this LP is $\chi(G, c)$.

Let $w : E \to R^+$ be an assignment of non-negative weights to the edges. We use $w_e$ to denote the weight of edge $e \in E$. For given edge weights $w$, $OPT(G, w)$ denotes the weight of the minimum weight Steiner tree on $G$. Finding $OPT(G, w)$ is NP-hard. One can formulate the following bidirected cut integer program to find $OPT(G, w)$. For each undirected edge $e \in E$, we have two directed edges $e_1$ and $e_2$ which represent the two orientations of $e$; both have the same weight as the undirected edge. $D = \{e_1, e_2, \forall e \in E\}$. A valid set $C$ is a subset of vertices such that $C$ contains the source $m_0$ and $\bar{C}$ contains at least one terminal. $\delta(C) = \{(u, v) \in D : u \in C, v \notin C\}$. The following integer program computes the minimum weight Steiner tree:

$$\begin{aligned} \min \quad & \textstyle\sum_{a\in D} w_a c_a \\ \sum_{a\in\delta(C)} c_a \ge 1 \quad & \forall \text{ valid sets } C \\ c_a \in \{0, 1\} \quad & \forall a \in D \end{aligned}$$

Note that $c_a$ are the variables in the above integer program; the optimum value is $OPT(G, w)$. In the bidirected cut LP relaxation, we replace the last constraint by the following

$$c_a \ge 0, \forall a \in D$$

Let the optimum value of the bidirected cut relaxation for $G$ with given weights $w$ be denoted by $B(G, w)$. The maximum value of the ratio $OPT(G, w)/B(G, w)$ is called the integrality gap of the LP relaxation.

Note that the coding advantage is invariant under multiplicative scaling of capacities; similarly, the integrality gap is invariant under multiplicative scaling of weights.

## III. Analysis for the Undirected Case

In this section we present the arguments to show that maximum coding advantage for a given network $G$ is the same as the integrality gap of the bidirected cut relaxation for minimum weight Steiner tree. We establish this by proving inequality in both directions between these two quantities.

**Theorem III.1.**

$$\max_c \frac{\chi(G, c)}{\Pi(G, c)} \le \max_w \frac{OPT(G, w)}{B(G, w)}.$$

*Proof.* Let us consider the Steiner packing number for a graph with capacities $c_e$. Assume, without loss of generalization, that the capacities are scaled so that the value of the objective of the cFlow LP for this graph (i.e. $\chi(G, c)$) is exactly 1.

Consider the dual to the Steiner packing LP; by strong duality, the optimum value is equal to $\Pi(G, c)$. We claim that the optimum solution of this dual gives us a gap example for the bidirected cut relaxation, with integrality gap at least as large as the coding advantage. To see why this is true, notice that we can view the $y_e$'s as edge costs. The first constraint of the dual gives us the condition that every Steiner tree under these edge costs should have cost at least 1.

**Claim III.2.** *$\sum_{e\in E} c_e y_e$ is an upper bound on the value of the bidirected cut relaxation for the Steiner tree instance with edge costs given by the $y_e$'s.*

*Proof.* Since the graph $G$ with capacities $c_e$ has a cFlow value of 1, every edge can be bidirected and the capacity distributed between the two oppositely directed copies so that a flow of 1 can be routed separately from the source to every terminal according to these capacities. But then these bidirected capacities give a valid solution to the bidirected cut relaxation. (Note that the directed capacity across any cut separating the source and at least one terminal must be at least 1). □

Hence, from the optimum solution of the dual, there exists a setting of weights $w$ such that $OPT(G,w) \geq 1$ and $B(G,w) \leq \Pi(G,c)$. ☐

**Theorem III.3.**

$$\max_c \frac{\chi(G,c)}{\Pi(G,c)} \geq \max_w \frac{OPT(G,w)}{B(G,w)}.$$

*Proof.* Consider an instance $G$ with weights $w$ for which the bidirected cut relaxation has gap $g$. From $G$, we will construct an instance for which the coding advantage is atleast $g$.

Consider the capacities $c_a$ (one for each orientation of each edge) returned by the bidirected cut relaxation on $G$. For each undirected edge $e$, let $c_e$ to be the sum of the edge capacities for both orientations of edge $e$. Now run the cFlow LP on $G$ with these capacities.

**Claim III.4.** *The value of the cFlow LP (i.e. $\chi(G,c)$) for the graph $G$ with capacities $c_e$ is atleast* 1.

*Proof.* The bidirected cut relaxation shows how to distribute $c_e$ amongst forward and back edges such that for every cut separating the source and a terminal, the directed capacity is atleast 1. This means that the directed graph can support a flow of at least 1 from the source to every terminal. ☐

The Steiner tree packing number for this instance is $\Pi(G,c)$. This means for any setting of weights $w$, there is a Steiner tree of cost at most $\frac{\sum w_e c_e}{\Pi(G,c)}$. Recall that we started with a gap instance for the bidirected cut relaxation with $B(G,c) = \sum w_e c_e$ and $\frac{OPT(G,c)}{B(G,c)} = g$. But the above statement implies that there is a Steiner tree of cost at most $\frac{B(G,c)}{\Pi(G,c)}$. Hence $\Pi(G,c) \leq 1/g$, which implies that the coding advantage $\frac{\chi(G,c)}{\Pi(G,c)} \geq g$, the integrality gap. ☐

## Gap example for the undirected case

For completeness, we reproduce the best known gap example for the bidirected cut formulation for Steiner tree (due to Goemans, as described in [8]). The example consists of $n+1$ terminals, labeled $a_0, \ldots, a_n$. In addition we have $2\binom{n}{2}$ vertices, which are labeled $c_{ij}$ and $d_{ij}$, for $1 \leq i,j \leq n$. The vertices are connected as shown in Figure 1. For every $1 \leq k \leq n$, we have an additional vertex $b_k$. We have edges $(a_i, b_i)$ and $(b_i, a_0)$ for all $i$ of cost 2. For every $1 \leq i,j \leq n$, $i \neq j$, we also have edges $(a_i, c_{ij})$ of cost 2, and edges $(b_i, d_{ij})$ of cost 1. Finally, for every $1 \leq i,j \leq n$, $i \neq j$ we have edges $(c_{ij}, d_{ij})$ of cost 1.

The integrality gap of this example approaches $\frac{8}{7}$ as $n \to \infty$. A feasible solution for the linear relaxation is the following: each edge is selected to extent $\frac{1}{n}$. Using these values as edge capacities, we get an example with coding advantage approaching $\frac{8}{7}$. The example is similar to the gap example described previously; a portion of the graph is depicted in Figure 2.

## 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, 2000.

[2] E. Halperin, G. Kortsarz, R. Krauthgamer, A. Srinivasan and N. Wang, “Integrality ratio for Group Steiner Trees and Directed Steiner Trees,” *Proc. 14th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 275–284, 2003.


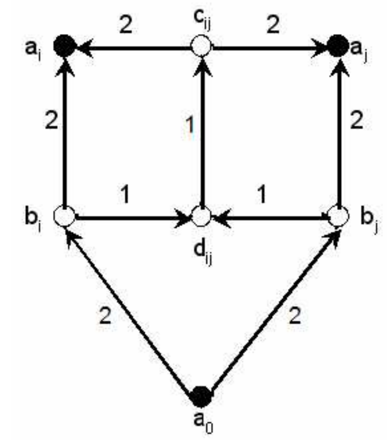


Fig. 1: Gadget used in construction of gap example for bidirected cut relaxation.

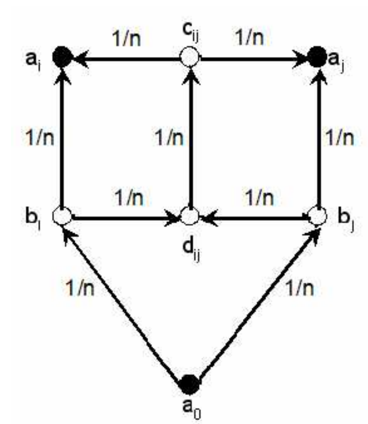


Fig. 2: Portion of the example with high coding advantage.


[3] K. Jain, M. Mahdian, and M. R. Salavatipour, “Packing Steiner trees,” *Proc. 14th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 266–274, 2003.

[4] G. Kramer, and S. A. Savari, “Cut sets and information flow in networks of two-way channels,” *Proc. IEEE International Symposium on Information Theory (ISIT)*, p. 33, 2004

[5] Z. Li, B. Li, D. Jiang, and L. C. Lau, “On achieving optimal end-to-end throughput in data networks: theoretical and empirical studies,” *ECE Technical Report*, University of Toronto, 2004.

[6] S. Rajagopalan, and V. Vazirani, “On the bidirected cut relaxation for the metric Steiner tree problem,” *Proc. 11th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 742–751, 1999.

[7] P. Sanders, S. Egner, and L. Tolhuizen, “Polynomial Time Algorithms for Network Information Flow,” *Proc. 15th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pp. 286-294, 2003.

[8] A. Sinha, “Steiner trees: first summer paper for the PhD program at GSIA,” *manuscript*, 2000.