Space Information Flow: Multiple Unicast

Zongpeng Li
Dept. of Computer Science, University of Calgary
and Institute of Network Coding, CUHK
zongpeng@ucalgary.ca

Chuan Wu
Department of Computer Science
The University of Hong Kong
cwu@cs.hku.hk

***Abstract*—The multiple unicast network coding conjecture states that for multiple unicast in an undirected network, network coding is equivalent to routing. Simple and intuitive as it is, the conjecture has remained open since its proposal in 2004 [1], [2], and is now a well-known unsolved problem in the field of network coding. In this work, we provide a proof to the conjecture in its space/geometric version. Space information flow is a new paradigm being proposed [3], [4]. It studies the transmission of information in a geometric space, where information flows are free to propagate along any trajectories, and may be encoded wherever they meet. The goal is to minimize a natural bandwidth-distance sum-product (network volume), while sustaining end-to-end unicast and multicast communication demands among terminals at known coordinates. The conjecture is true in networks only if it is true in space. Our main result is that network coding is indeed equivalent to routing in the space model. Besides its own merit, this partially verifies the original conjecture, and further leads to a geometric framework [5] for a hopeful proof to the conjecture.**

I. Background and Introduction

A. The Multiple Unicast Network Coding Conjecture

Departing from the classic store-and-forward principle of data networking, network coding encourages information flows to be "mixed" in the middle of a network, via means of coding [6], [7]. While network coding for a single communication session (unicast, broadcast or multicast) is well understood by now, the case of multiple independent sessions (multi-source, multi-sink) is much harder, with less results known [8]. A basic scenario in the latter is the multiple unicast setting, where multiple independent one-to-one communication demands co-exist in a network. With routing, the optimal solution can be computed by solving a multicommodity flow (MCF) linear program; with network coding, the structure and the computational complexity of the optimal solution are largely unknown.

In directed networks, network coding can augment the capacity region of multiple unicast. For example, Fig. 1(A) shows a network coding solution for two unicast sessions in a directed network, where each session has a throughput of 1. Without network coding, it is not hard to verify that achieving a throughput of 1 and 1 for both sessions concurrently is infeasible, given the pre-defined link directions. In general, the *coding advantage*, the ratio of the maximum throughput achievable with network coding over that with routing, may grow linearly with the network size [9].

However, no >1 coding advantage has been observed for multiple unicast in the undirected setting. For example, Fig. 1(B) shows a MCF with end-to-end flow rate of 1 for each of the two unicast sessions. Li and Li [1] and Harvey *et al*. [2] conjectured that network coding is equivalent to routing for multiple unicast in undirected networks.

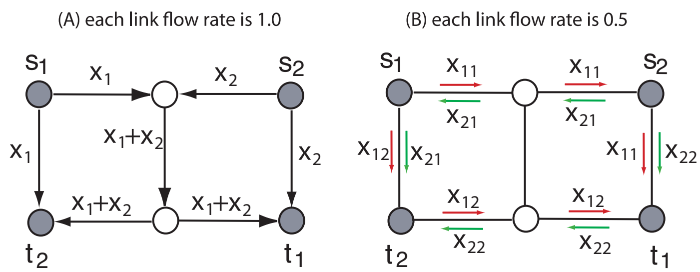

Fig. 1. (Example from [1].) Two unicast sessions, from s_1 to t_1 and from s_2 to t_2, each with target rate 1. All link capacities are 1. (A). Solution with network coding. (B) Solution without network coding.

Despite a series of research effort devoted [10]–[12], rather limited progresses have been made towards settling this fundamental problem in network coding. Besides "easy" cases where the cut set bounds can be achieved without network coding [1], [2], the conjecture has been verified only in small, fixed networks and their variations, such as the Okamura-Seymour network [10], [11]. It is worth noting that such verification already involves new tools in network information theory such as information dominance [10], input-output equality and crypto equality [11]. A growing agreement is that new tools beyond a "simple blend" of graph theory and information theory are required for eventually settling the conjecture. In this work, we prove the geometric version of the multiple unicast conjecture, by further incorporating mature techniques in geometry into the picture.

In 2007, Mitzenmacher *et al*. compiled a list of seven open problems in network coding [13], where the multiple unicast conjecture appears as problem number 1. Chekuri commented that claiming an equivalence between network coding and routing for *all* undirected networks is a "bold conjecture", and that the problem of fully understanding network coding for multiple unicast sessions is still "wild open" ([14], p51-55).

B. Space Information Flow

Space information flow is a new subject of study being proposed [3], [4]. It considers terminals at known locations in a geometric space, with unicast, broadcast or multicast communication demands among them. Information flows can be transmitted along any trajectories in the space, and may be replicated wherever desired, or encoded wherever they meet. The goal is to minimize the total bandwidth-distance sum-product, while sustaining given end-to-end communication rates. Besides being a conceivable theoretical problem of "network coding in space", space information flow models the min-cost design of a blueprint of a communication network, which deserves renewed research attention given network coding [3]. As we will discuss later, space information flow also opens the door to geometric approaches for studying network information flow problems, including in particular the multiple unicast network coding conjecture in graphs.

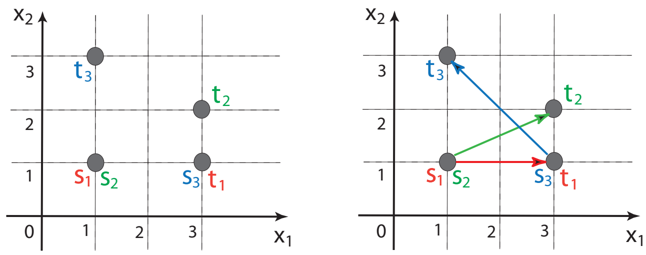

Fig. 2. A 2-D example of space information flow: meeting communication demands among nodes in space. A min-cost solution is to be computed, for three unit-demand unicast sessions from s_1 to t_1, from s_2 to t_2 and from s_3 to t_3, respectively (left). Given network coding, is there a solution better than MCF (right)?

For a quick feel of space information flow, consider three unicast sessions each with unit demand, from s_1 to t_1, from s_2 to t_2 and from s_3 to t_3, respectively, in a 2-D Euclidean space as shown in Fig.1. We can route an information flow along any path in space, insert relay nodes wherever desired, and replicate or encode information flows wherever desired. We aim to minimize the volume of the solution network induced, $\sum_e ||e||\mathbf{f}(e)$. Here e is a link employed for flow transmission, $||e||$ is the length of e in space, and $\mathbf{f}(e)$ is the rate of information flow routed across e. What is the optimal solution for satisfying the three unicast demands? Can network coding lead to better solutions than routing (MCF)? Recent examples show that network coding can outperform routing when the demand in space is multicast [3], [4]. What about multiple unicast?

C. Summary of Results

Our main result in this work is that network coding is equivalent to routing, for multiple unicast in the new space information flow model. We restrict our attention to Euclidean spaces; the case of non-Euclidean spaces, such as n-dimensional spaces under Chebyshev distance, is still being investigated [5].

We first analyze the simple case of a 1-D space, where a single point constitutes a cut. A natural requirement on a valid multiple unicast solution here is that, at any given point A, the total amount of flows at A, aggregated from both directions, should be at least the total demand of unicast sessions whose terminals reside on different sides of A. We take integration on both sides of this inequality along all points in the 1-D space, and prove that network coding can not improve upon an optimal solution based on routing (MCF).

For the general case of a h-D space, $h \geq 2$, our approach is to reduce the problem into 1-D, by applying the mature tool of projection in geometry. We prove that, if network coding can outperform MCF in h-D, then it can do so in 1-D, thereby leading to a contradiction. More specifically, we show that in a given case where a network coding based solution has a smaller cost than that of MCF, there must exist a 1-D subspace, onto which the projection of the network coding solution is still cheaper than the projection of the MCF solution. The challenge here is that such a "good" candidate subspace for projection is hard to find. It is problem dependent and no fixed subspace always works. We prove the existence of such an elusive subspace without explicitly identifying it, through an argument of integrating the projected network coding and MCF solutions over all possible 1-D rays from origin.

D. Relevance and Discussions

In Sec. II-C, we prove that the *cost advantage*, the potential advantage of network coding over routing in terms of reducing data transmission cost, is always at least as high in networks than in space. Therefore, our result in this paper partially verifies the original multiple unicast network coding conjecture in networks.

Perhaps more interesting is that the new space information flow perspective provides a promising direction for attacking the original conjecture itself. In a sibling work [5], we describe a geometric framework that is hopeful for eventually resolving the original conjecture. We briefly preview this geometric framework, as well its connection to this work, in Sec. III-D.

Given that network coding is equivalent to routing for multiple unicast in a Euclidean space, it is interesting to ask whether the same holds for multicast. In two sibling work [3], [4], we study the multicast problem in space, with network coding explicitly considered. There we present examples that show network coding and routing are indeed different in space, prove upper-bounds on the cost advantage, analyze the achievability of optimality with finite solutions, and discuss the complexity of optimal multicast in a geometric space.

II. Problem Models

A. Network Information Flow

We represent an existing network, directed or undirected, using a graph $G = (V, E)$. The vector $\mathbf{c} \in \mathbf{Z}_+^E$ stores capacities of links in E. Here $\mathbf{Z}_+$ is the set of positive integers. Another vector $\mathbf{w} \in \mathbf{Q}_+^E$ represents the *distance* or *cost* of links in E, and $\mathbf{w}_e$ can be interpreted as the cost of routing a unit flow through that link. Here $\mathbf{Q}_+$ represents the set of positive rational numbers.

For the min-cost multiple unicast problem, we consider k unicast sessions co-existing in network G, and let s_i and t_i be the sender and receiver of session $i \in \{1, \ldots, k\}$. We use $\mathbf{r}$ to denote the target throughput vector of the k sessions, and $\mathbf{r}_i$ is the required throughput of session i. Without network coding, a solution to the multiple unicast problem is a multicommodity flow (MCF), which can be represented using a link flow vector $\mathbf{f} \in \mathbf{Q}_+^E$. The min-cost MCF can be computed by solving a linear program [1], [2].

A network coding solution to the multiple unicast problem has two components: (A) a flow component, for how much flow to transmit over each link, and (B) a coding component, for where and how to encode and decode the information flows. We denote the underlying link flow vector in (A) using $\mathbf{f} \in \mathbf{Q}_+^E$ too. In undirected networks, a network coding scheme may be dynamic in that the transmission scheme is a time-slotted one (a convolutional code), and a different flow routing and coding scheme is adopted in each different time slot [6], [10]. In this case, we simply let $\mathbf{f_e}$ be the time-average flow rate at link e. There is no known linear program of polynomial size that computes the min-cost network coding solution.

B. Space Information Flow

In the space information flow problem, we are given a set of terminal nodes, with (multiple) unicast or multicast communication demand. The space we consider in this work is a h-D Euclidean space, $h \geq 1$. A node u has coordinate $(x_{1,u}, x_{2,u}, \ldots, x_{h,u})$. The Euclidean distance between two nodes u and v is

$$||uv||_h = \left(\sum_{i=1}^{h} (x_{i,u} - x_{i,v})^2 \right)^{1/2}$$

Given a space information flow vector $\mathbf{f}$, a network can be induced, over the same nodes and links as in $\mathbf{f}$, by viewing $\mathbf{f}_e$ as the capacity of e. The distance of e is denoted as $||e||_h$. The cost of $\mathbf{f}$ is then $\sum_e ||e||_h \mathbf{f}_e$. This reflects the general rule that the longer and the wider a communication cable, the more expensive it is. For the sake of cost minimization, apparently, only straight line segments need to be considered in $\mathbf{f}$.

Given two vectors $\vec{p}$ and $\vec{q}$, $\vec{p} \cdot \vec{q} = |||p||||q|||\cos\theta|$ is the absolute value of the inner product of $\vec{p}$ and $\vec{q}$, where θ is the angle between $\vec{p}$ and $\vec{q}$.

C. Paradigm Comparison

We can establish a connection between the cost advantage in space and that in graphs. Given a problem instance, in the form of either multiple unicast or multicast, let β_d, β_u and β_s be the max cost advantage possible in directed networks, undirected networks, and space, respectively. Then we have the following relation among the three:

Theorem 2.1. $\beta_d \geq \beta_u \geq \beta_s$.

Proof: We first show that $\beta_d \geq \beta_u$. Given the max cost advantage β_u in undirected networks, let Δ_u be a problem instance where this cost advantage is achieved, and let $\mathbf{f}$ be the underlying flow of the optimal network coding solution. We can create a corresponding problem instance Δ_d for the directed setting, by viewing $\mathbf{f}$ as the directed network, while keeping the terminal nodes, link costs and target throughput intact. With network coding, the cost of the optimal solution is the same in Δ_d and in Δ_u. Without network coding, the cost of the optimal solution can only increase from Δ_u to Δ_d, since the latter is more restrictive. Therefore $\beta_d \geq \beta_u$.

The proof to $\beta_u \geq \beta_s$ is similar, by viewing the underlying flow $\mathbf{f}$ of the optimal network coding solution for Δ_s as an undirected network. The directions in $\mathbf{f}$ are ignored. The cost of a link e is taken as $||e||_h$. □

Given Theorem 2.1, we know that all upper-bounds on the cost advantage proven for the undirected model are still valid in the space model. Conversely, all lower-bounds that we can prove for the space model will also be valid for the undirected model. For example, an upper-bound of 2 is known for cost advantage in undirected multicast networks [15]–[17]. This bound automatically holds for multicast in a space of any dimension. In Sec. III, we prove that the cost advantage for multiple unicast is always 1 in space. Unfortunately, this does not directly imply the multiple unicast conjecture in undirected networks. We discuss how this bound is connected to the conjecture in Sec. III-D.

III. Space Information Flow: multiple unicast

A. The Multiple Unicast Conjecture for Network Information Flow

In their original work where the multiple unicast conjecture was proposed [1], Li and Li first formulated the conjecture in the throughput domain, and then applied linear programming duality to translate it into an equivalent version in the cost domain.

The Multiple Unicast Conjecture [1], [2]
Throughput domain: *For k independent unicast sessions in a capacitated undirected network $(G, \mathbf{c})$, a throughput vector $\mathbf{r}$ is feasible with network coding if and only if it is feasible with routing.*
Cost domain: *Let $\mathbf{f}$ be the underlying flow vector of a network coding solution for k independent unicast sessions with throughput vector $\mathbf{r}$, in a cost-weighted undirected network $(G, \mathbf{w})$. Then $\sum_e \mathbf{w}_e \mathbf{f}_e \geq \sum_i \mathbf{d}_i \mathbf{r}_i$, where $\mathbf{d}_i$ is the shortest path distance between the sender and receiver of session i under metric $\mathbf{w}$.*

Intuitively, the throughput version of the conjecture claims that network coding cannot help improve throughput, while the cost version claims that network coding cannot help reduce transmission cost. In the rest of this section, we prove the cost version of the multiple unicast conjecture for space information flow, where the cost $\mathbf{w}_e$ becomes, naturally, the Euclidean length $||e||$ of link e.

B. Multiple Unicast in 1-D Space

In a 1-D space, each line segment (or edge) e between two neighboring vertices forms a cut of the network. The amount of flow $\mathbf{f^1}_e$ over e has to be at least the total throughput requirement of terminal pairs separated by the removal of e. We next prove that this implies the multiple unicast conjecture in 1-D space.

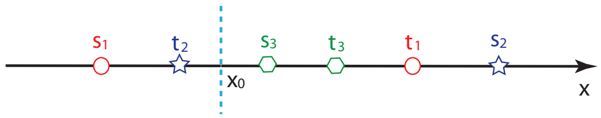

Fig. 3. Three unicast sessions in 1-D. Total flow crossing point x_0, $\mathbf{f^1}_{x_0}$, is lower-bounded by $Demand((-\infty, x_0); (x_0, \infty)) = \mathbf{r_1} + \mathbf{r_2}$.

Theorem 3.1. *Given k independent unicast sessions in 1-D space, let $\mathbf{f^1}$ be the underlying flow vector of a network coding solution achieving a rate vector $\mathbf{r}$. Then $\sum_e(||e||_1\mathbf{f^1}_e) \geq \sum_i(||s_it_i||_1\mathbf{r}_i)$.*

Proof: For a given point x in the 1-D space, let $\mathbf{f^1}_x$ be the total amount of flow crossing x, in both directions. Note that the point x constitutes a cut of the 1-D space, and therefore $\mathbf{f^1}_x$ is lower-bounded by the flow demand between the left sub-space $(-\infty, x)$ and the right sub-space (x, ∞), denoted as $Demand((-\infty, x); (x, \infty))$. We integrate both sizes over the entire 1-D space, and obtain:

$$\int_{x=-\infty}^{\infty} \mathbf{f^1}_x dx \geq \int_{x=-\infty}^{\infty} Demand((-\infty, x); (x, \infty))dx$$

$$= \sum_i ||s_it_i||_1 \mathbf{r}_i$$

Furthermore, note that $\sum_e(||e||_1\mathbf{f^1}_e) = \int_{x=-\infty}^{\infty} \mathbf{f^1}_x dx$. We conclude that $\sum_e(||e||_1\mathbf{f^1}_e) \geq \sum_i(||s_it_i||_1\mathbf{r}_i)$. □

C. Multiple Unicast in h-D Space

We now consider multiple unicast demands in a h-D space, for $h \geq 2$. While only the cases of $h = 2$ and $h = 3$ allow intuitive interpretations, the problem is as well-defined for higher dimensions, which is helpful in connecting to the original multiple unicast conjecture in graphs, because embedding a graph metric into a geometric space often requires a high dimension space [5].

We prove the multiple unicast conjecture by projecting the problem from h-D to 1-D, and then apply Theorem 3.1. The requirement on the projection is: a coding solution has total cost less than the specified bound in the conjecture, only if it does so after the projection. The main difficulty of the proof is that an optimal or "good" direction for projection is actually hard to find. In particular, it is not sufficient to always project onto one of the axes. We show indirect evidence instead, for the existence of such a good direction, by taking an integration over all possible rays at origin for projection.

Theorem 3.2. *For k independent unicast sessions in a h-D space, $h \geq 2$, assume there is a network coding solution with underlying flow vector $\mathbf{f^h}$, we have $\sum_e(\mathbf{f^h}_e||e||_h) \geq \sum_i(||s_it_i||_h\mathbf{r}_i)$.*

Proof: Assume, by way of contradiction, that $\sum_e(\mathbf{f^h}_e||e||_h) < \sum_i(||s_it_i||_h\mathbf{r}_i)$. We construct the k pairs unicast instance and its network coding solution in 1-D by projecting their counter parts from h-D. Our goal is to show that there exists a 1-D sub-space/direction in the h-D space, onto which the projection satisfies $\sum_e(\mathbf{f^1}_e||e||_1) < \sum_i(||s_it_i||_1\mathbf{r}_i)$, and therefore obtain a contradiction to Theorem 3.1.

As shown in Fig. 4, let Φ be the surface of the h-D unit hyper-sphere at the origin. We can enumerate all possible directions in h-D by traversing all points on Φ, and connecting to there from the origin. Let $\overrightarrow{p}$ be the vector from origin to the corresponding point on Φ, let $\overrightarrow{1}$ be the unit vector $(1, 0, 0, \ldots, 0)$.

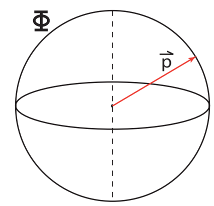

Fig. 4. Project and Integrate over all possible diretions $\overrightarrow{p}$.

The integration over the closed surface Φ for all the projections of $\mathbf{f^h}$ is:

$$\oiint_\Phi \sum_e(\mathbf{f^h}_e(e \cdot \overrightarrow{p}))d\Phi =_1 \sum_e \oiint_\Phi \mathbf{f^h}_e(e \cdot \overrightarrow{p})d\Phi$$

$$=_2 \sum_e \oiint_\Phi \mathbf{f^h}_e||e||_h(\overrightarrow{1} \cdot \overrightarrow{p})d\Phi$$

$$=_3 \sum_e(\mathbf{f^h}_e||e||_h) \oiint_\Phi (\overrightarrow{1} \cdot \overrightarrow{p})d\Phi$$

The nice property of this integration is that it is separable, in the sense that we can perform integration for each link flow segment first, and then take the summation ($=_1$). Furthermore, we observe that when we integrate for each line segment, the orientation of that line segment does not matter, since we vary the projection direction to take all possible values ($=_2$).

The integration over the closed surface Φ for all the projections of $\{\overrightarrow{s_it_i} \,|i = 1, \ldots, k\}$ is:

$$\oiint_\Phi \sum_i(\overrightarrow{s_it_i} \cdot \overrightarrow{p})d\Phi = \sum_i \oiint_\Phi (\overrightarrow{s_it_i} \cdot \overrightarrow{p})d\Phi$$

$$= \sum_e \oiint (||s_it_i||_h(\overrightarrow{1} \cdot \overrightarrow{p}))d\Phi = \sum_i ||s_it_i||_h \oiint_\Phi (\overrightarrow{1} \cdot \overrightarrow{p}))d\Phi$$

Since $\sum_e(\mathbf{f^h}_e||e||_h) < \sum_i ||s_it_i||_h$ by assumption, we claim that:

$$\oiint_\Phi \sum_e(\mathbf{f^h}_e(e \cdot \overrightarrow{p}))d\Phi < \oiint_\Phi \sum_i(\overrightarrow{s_it_i} \cdot \overrightarrow{p})d\Phi$$

Since the terms being integrated on both sides are non-negative, we claim that, there must exist a particular direction $\overrightarrow{p^*}$, for which

$$\sum_e(\mathbf{f^h}_e(e \cdot \overrightarrow{p^*})) < \sum_i(\overrightarrow{s_it_i} \cdot \overrightarrow{p^*})$$

□

D. Connection to the Multiple Unicast Conjecture

Comparing Theorem 3.2 with the original multiple unicast conjecture in undirected networks, we note that the two statements are rather similar. The only difference lies in the fact that Theorem 3.2 is based on Euclidean distances, whereas the conjecture is based on a cost metric induced from a graph. A natural direction for settling the conjecture is then to embed the graph metric into a geometric space, and then utilize Theorem 3.2.

An *isometric* embedding of a graph G into a space is one that preserves pairwise node distances in G. The distance between two nodes u and v in G is the shortest path length between u and v in G. By Theorem 3.2, we can see that, if there exists a certain space to which an isometric embedding of the network is feasible, and our projection based proof technique for dimension reduction can be adapted to carry through, then we can prove the multiple unicast conjecture. While no Euclidean space always permits isometric embedding of graphs, there do exist non-Euclidean spaces that satisfy this property [5]. For embedding into Euclidean spaces, if we relax the isometric requirement and allow a distance distortion ratio up to α in the embedding, we can prove the cost advantage is upper-bounded by α in the original graph.

Following this direction, we present a geometric framework for studying the multiple unicast network coding conjecture in a sibling work [5]. The framework consists of four major steps: (i) translating the conjecture from throughput domain to cost domain, (ii) embed the network into an Euclidean or non-Euclidean space, in an isometric or low-distortion manner, (iii) reduce the problem from high-dimension space to low dimension space, and (iv) prove that there cannot be a counter example to the conjecture in low dimension.

Based on this framework, we have been able to formulate a unified proof to a number of results, including (1) the conjecture holds for two unicast sessions, (2) the gap between a network coding solution and a routing solutions is at most $O(\log k)$, (3) the conjecture holds for uniform complete graphs, (4) the gap between a network coding solution and a routing solutions is at most $\sqrt{2}$ in uniform grid networks, (5) the conjecture is true in star networks [18], and (6) the conjecture is true in a class of infinitely many layered or bipartite networks. Among these, (1) and (2) were known before, but were proved using different techniques. Results (3)-(6) are new and not previously known. The proofs to results (2), (3) and (4) resort to embedding into a Euclidean space, and directly build upon the main result in this work. The proofs to results (1), (5) and (6) resort to embedding into a non-Euclidean space instead, where isometric embedding of a graph metric is feasible.

IV. Conclusion and Future Directions

This work is among the first that studies the problem of space information flow, with a focus on the case of multiple unicast sessions. We proved that for multiple unicast in a Euclidean space, network coding is equivalent to routing. This partially verifies the well-known multiple unicast conjecture in network coding. This result, together with the new space information flow perspective, leads to a new geometric framework for studying the original multiple unicast conjecture. For the paradigm of space information flows, the multicast case appears even more interesting, where basic problems such as the computational complexity of the optimal multicast solution in space are yet to be investigated.

References

[1] Z. Li and B. Li, "Network Coding: The Case of Multiple Unicast Sessions," in *Proc. of The 42nd Annual Allerton Conference on Communication, Control, and Computing*, 2004.

[2] N. J. A. Harvey, R. D. Kleinberg, and A. R. Lehman, "Comparing Network Coding with Multicommodity Flow for the k-pairs Communication Problem," Tech. Rep., CSAIL, MIT, http://theory.csail.mit.edu/ nickh/Publications/Gaps/TR964.pdf, November 2004.

[3] Z. Li and C. Wu, "Space Information Flow," submitted to NetCod 2012.

[4] X. Yin, Y. Wang, X. Wang, X. Xue, and Z. Li, "Min-Cost Multicast Network in Euclidean Space," submitted to ISIT 2012.

[5] T. Xiahou, C. Wu, J. Huang, and Z. Li, "A Geometric Framework for Studying the Multiple Unicast Network Coding Conjecture," submitted to NetCod 2012.

[6] 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.

[7] 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.

[8] X. Yan, R. W. Yeung, and Z. Zhang, "The Capacity Region for Multi-source Multi-sink Network Coding," in *Proc. of IEEE International Symposium on Information Theory (ISIT)*, 2007.

[9] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. M. G. Tolhuizen, "Polynomial Time Algorithms for Multicast Network Code Construction," *IEEE Transactions on Information Theory*, vol. 51, no. 6, pp. 1973–1982, June 2005.

[10] N. J. A. Harvey, R. Kleinberg, and A. R. Lehman, "On the Capacity of Information Networks," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2345–2364, June 2006.

[11] K. Jain, V. V. Vazirani, R. W. H. Yeung, and G. Yuval, "On the Capacity of Multiple Unicast Sessions in Undirected Graphs," in *Proc. of IEEE International Symposium on Information Theory (ISIT)*, 2005.

[12] M. Langberg and M. Médard, "On the Multiple Unicast Network Coding Conjecture," in *47th Annual Allerton Conference on Communication, Control, and Computing*, 2009.

[13] M. Mitzenmacher, "Network Coding: Open Problems," Tech. Rep., http://mybiasedcoin.blogspot.com/2007/08/network-coding-open-problems.html, August 2007.

[14] C. Chekuri, "Routing vs. Network Coding," Tech. Rep., University of Illinois, Urbana-Champaign, temple.birs.ca/∼09w5103/chekuri_09w5103_talk.pptx, 2009.

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

[16] A. Agarwal and M. Charikar, "On the Advantage of Network Coding for Improving Network Throughput," in *IEEE Information Theory Workshop*, October 2004.

[17] C. Fragouli and E. Soljanin, *Network Coding Fundamentals*, Now Publishers, 2007.

[18] S. M. S. T. Yazdi, S. A. Savari, and G. Kramer, "Network Coding in Star Networks," in *Proc. of IEEE International Symposium on Information Theory (ISIT)*, 2008.