A Geometric Framework for Investigating the Multiple Unicast Network Coding Conjecture

Tang Xiahou†, Chuan Wu‡, Jiaqing Huang§, Zongpeng Li†¶

† Department of Computer Science, University of Calgary, Canada

‡ Department of Computer Science, The University of Hong Kong

§ Dept. of Electronics & Information Engineering, Huazhong University of Science and Technology, China

¶ Institute of Network Coding, The Chinese University of Hong Kong

***Abstract*—The multiple unicast network coding conjecture states that for multiple unicast sessions in an undirected network, network coding is equivalent to routing. Simple and intuitive as it appears, 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. Based on a recently proposed tool of *space information flow* [3]–[5], we present a geometric framework for analyzing the multiple unicast conjecture. The framework consists of four major steps, in which the conjecture is transformed from its throughput version to cost version, from the graph domain to the space domain, and then from high dimension to 1-D, where it is to be eventually proved. We apply the geometric framework to derive unified proofs to known results of the conjecture, as well as new results previously unknown. A possible proof to the conjecture based on this framework is outlined.**

I. Introduction

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 sessions (multi-source, multi-sink) is much harder, with fewer results known [8]. The case of multiple independent one-to-one unicast sessions is probably the most basic scenario of the multi-source multi-sink setting. With routing, multiple unicast is equivalent to the combinatorial problem of multicommodity flows (MCF) [9], which is polynomial time computable. With network coding, the structure and the computational complexity of the optimal solution are largely unknown.

If the network is directed, network coding can outperform routing for multiple unicast sessions. Fig. 1(A) shows a network coding solution for two unicast sessions, each with an end-to-end throughput of 1. If each link direction is fixed, then we can verify that achieving a throughput of 1 and 1 concurrently is infeasible without network coding. The potential of throughput improvement due to network coding is unbounded, for multiple unicast in a directed network [1].

Interestingly, the picture is drastically different in undirected networks, where the capacity of a link is flexibly sharable in two opposite directions. No example is known where network coding makes a difference from routing. Fig. 1(B) shows a MCF with end-to-end flow rate of 1 and 1, which is feasible if the underlying network in Fig. 1(A) is undirected. Harvey

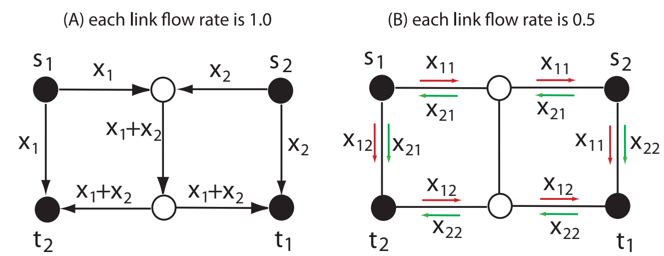

Fig. 1. (Example from [1].) Two unicast sessions, from s_1 to t_1 and from s_2 to t_2, each with target rate 1. All link capacities are 1. (A) Solution with network coding. (B) Solution without network coding.

et al. [2] and Li and Li [1] conjectured that network coding is equivalent to routing for multiple unicast in undirected networks.

Despite a series of research effort devoted [10]–[12] to it, this fundamental problem in network coding has witnessed rather limited progresses towards its resolution. 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 such as information dominance [10], input-output equality and crypto equality [11].

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). 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 apply a recently proposed tool, *space information flow* [3]–[5], to develop a geometric framework for studying the multiple unicast network coding conjecture. The framework consists of four main steps. In Step 1, LP duality is applied for translating the conjecture from its throughput version to an equivalent cost version. In Step 2, graph embedding is performed, for translating the cost version from the network domain to the space domain. Step 3 aims at dimension reduction that brings the problem from a high dimension space to 1-D. Step 4 contains a direct proof in 1-D, where the cut condition on information flow transmission is readily applicable. Based on the geometric framework, we derive unified proofs to a number of known results on the conjecture, as well as new results unknown before.

Step 1 of the framework borrows an existing result from previous work [1]. Step 2 is inspired by the very recent work on space information flow, where the optimal transmission of information flows, in a geometric space instead of in a fixed network topology, is studied. The embedding techniques applied include both classic and new ones. Step 3 exploits recent results developed in the space information flow paradigm and new results developed in this work. Step 4 is relatively simple, where the proof is done by taking an integration over the 1-D space on both sides of the cut condition inequality [4].

We hope that this framework will shed light onto the original multiple unicast conjecture in network coding, and possibly other problems in network information flow as well.

II. Model and Preliminaries

We use $G = (V, E)$ to represent an undirected network, with $|V| = n$ nodes. Let $c \in \mathcal{Q}_+^E$ be a link capacity vector, and $w \in \mathcal{Q}_+^E$ be a link cost vector. Here $\mathcal{Q}_+$ is the set of positive rational numbers. For the multiple unicast problem, the set V contains in particular k sender-receiver pairs, s_i and t_i, $1 \le i \le k$. The k unicast sessions are independent, and have a desired throughput vector $\mathbf{r} = (r_1, \ldots, r_k)$.

In the max-throughput version of the multiple unicast problem, we are given a capacitated network (G, c), and wish to maximize a ratio $\alpha \ge 0$, such that the throughput vector $\alpha\mathbf{r}$ can be achieved. Let α_{NC} and α_{MCF} be the maximum values of α possible, under network coding and routing (MCF), respectively, then the *coding advantage* is defined as the ratio α_{NC}/α_{MCF}.

In the min-cost version of the multiple unicast problem, we are given a link-weighted network (G, w), with each link having unlimited capacity. Under routing (MCF), the minimum cost for achieving a throughput vector $\mathbf{r}$ is $\sum_i (d_i r_i)$, where d_i is the shortest path length between s_i and t_i in G, under cost vector w. Under network coding, we wish to minimize the total solution cost $\sum_e (w(e)f(e))$, such that vector f together with some code assignment forms a valid network coding solution for achieving throughput vector $\mathbf{r}$. Assume f^* is the underlying flow vector of an optimal network coding solution, we define the *cost advantage* of network coding as the ratio $\sum_i (d_i r_i)/\sum_e (w(e)f^*(e))$.

A h-D space with p-norm distance is denoted as l_p^h. For two nodes u and v in l_p^h with coordinates $(x_{u1}, \ldots, x_{uh})$ and $(x_{v1}, \ldots, x_{vh})$, respectively, the distance between u and v is:

$$||u, v||_p = \left(\sum_{i=1}^{h} |x_{ui} - x_{vi}|^p \right)^{\frac{1}{p}}$$

For the multiple unicast version of the space information flow problem, we are given k pairs of terminals, $(s_i, t_i), 1 \le i \le k$, in a space l_p^h. We seek the min-cost solution that can achieve a throughput vector $\mathbf{r}$, under the rule that relay nodes can be inserted anywhere for free, and the cost of a one-hop transmission is proportional to both its flow rate and its geometric distance. Under routing (MCF), the optimal cost is $\sum_i (||s_i t_i||_p r_i)$. Under network coding, let f^* be the underlying flow vector of the optimal solution. The cost is then $\sum_{e \in f^*} ||e||_p f(e)$. The cost advantage of network coding here is $\sum_i (||s_i t_i||_p r_i)/\sum_{e \in f^*} (||e||_p f^*(e))$.

III. The Geometric Framework

In this section, we describe the geometric framework for studying the multiple unicast conjecture, including its four major steps.

A. Step 1. From Throughput to Cost: LP Duality

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 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 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 $\sum_e (w(e)f(e)) \ge \sum_i (d_i \mathbf{r}_i)$.*

Li and Li proved that the throughput version of the conjecture is equivalent to the cost version, by applying LP duality in the form of the Japanese Theorem. In particular, their proof leads to the following result that will be used in this work:

Theorem 3.1. (Li and Li, 2004 [1]) Given an undirected network G with k pairs of unicast terminals specified, and any desired throughput vector $\mathbf{r}$, the maximum coding advantage in (G, c) over all $c \in \mathcal{Q}_+^E$, equals the maximum cost advantage in (G, w) over all $w \in \mathcal{Q}_+^E$.

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 Step 1 of the framework, we apply Theorem 3.1 to translate the statement to be proven from its throughput version to cost version.

B. Step 2. From Network to Space: Graph Embedding

An *embedding* of a link-weighted graph $(G = (V, E), w)$ into a space l_p^h involves assigning a h-D coordinate to each node $u \in V$. In the multiple unicast problem, we embed either the closure or the partial closure of G. The *closure network* G' is a complete network defined over the same set of vertices as in G, such that the cost of a link $e = (uv)$ equals d_{uv}, the shortest path length between u and v in G. The *partial closure* of G is G with direct links added between each pair of s_i and t_i, with cost d_i.

A closure embedding has distortion β if $||uv||_p \leq d(uv) \leq \beta \cdot ||uv||_p, \forall u, v \in V$. A partial closure embedding has distortion β if $||s_i t_i||_p \leq d_i \leq \beta ||s_i t_i||_p$, and $||e||_p \leq w(e), \forall e \in E$. In both cases, the embedding is *isometric* if $\beta = 1$.

Theorem 3.2. For k pairs of unicast sessions in an undirected network (G, w), with desired throughput vector $\mathbf{r}$, assume G has a β-distortion **closure embedding** in a space l_p^h. If the cost advantage is 1 after the embedding, then it is upper-bounded by β before the embedding.

Proof: If there is a network coding solution in G, with an underlying flow vector f satisfying $\sum_e (w(e)f(e)) < \sum_i (d_i r_i)$, then there is such a f' in G', by the definition of a closure network. The embedding of f' leads to a solution in l_p^h, where $\sum_e (||e||_p f'(e)) < \beta \cdot \sum_i (||s_i t_i||_p r_i)$ due to the β-distortion property of the embedding. □

A similar result holds for partial embedding as well.

Theorem 3.3. For k pairs of unicast sessions in an undirected network (G, w), assume there is a β-distortion **partial closure embedding** of G in a space l_p^h. If the cost advantage is 1 after the embedding, then it is upper-bounded by β before the embedding.

The proof of Theorem 3.3 is similar to that of Theorem 3.2, and is omitted. Informally, when the original link cost vector is 'nice', *e.g.*, satisfying the triangular inequality, partial closure embedding may be preferred. Otherwise, closure embedding is likely to be more helpful. A special case of Theorem 3.2 and Theorem 3.3 is when $\beta = 1$, then cost advantage is 1 after the embedding only if it is 1 before the embedding.

C. Step 3. From High Dimension to 1-D: Projection

Step 3 of the framework aims to simplify the statement to be proven from high dimension to 1-D. We introduce a few results useful for such dimension reduction.

Theorem 3.4. If there exists a configuration of k unicast sessions in l_∞^n, $n > k$, where $\sum_e (||e||_\infty f_e) < \sum_i (||s_i t_i||_\infty r_i)$, then there exists a configuration of k unicast sessions in l_∞^k, where the same inequality holds.

Proof: For each session i of the k unicast sessions in the l_∞^n space, let's define the *primary coordinate* of i as $argmax_j |x_{s_i j} - x_{t_i j}|$. We project the original multiple unicast instance from l_∞^n to l_∞^k by truncating the coordinate of each node in the following way: keep k coordinates including all the primary coordinates, dropping other coordinates.

After the projection from l_∞^n to l_∞^k above, the distance $||s_i, t_i||_\infty$ remains unchanged, for each session i. The distance between any two nodes u and v cannot increase. Therefore, $\sum_e (||e||_\infty f_e)$ does not increase due to the projection, while $\sum_i (d_i r_i)$ remains unchanged due to the projection, and hence the theorem is true. □

By definition, the normed spaces are all equivalent in 1-D. In particular, there is no difference between l_2^1 and l_∞^1. Therefore we drop the norm p from l_p^1, and simply write l^1.

Theorem 3.5. If there exists a configuration of k unicast sessions in l_∞^2, where $\sum_e (||e||_\infty f_e) < \sum_i (||s_i t_i||_\infty r_i)$, then there exists a configuration of k unicast sessions in l^1, where the same inequality holds.

Proof: Let $\vec{x}$ and $\vec{y}$ be two vectors in a space l_p^h. We define the projection of $\vec{x}$ onto $\vec{y}$ as $proj(\vec{x}, \vec{y}) = \frac{|\vec{x} \cdot \vec{y}|}{||\vec{y}||_p}$, where $\cdot$ is the inner product operation.

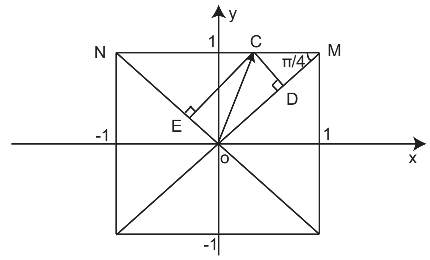

Fig. 2. Projecting a unit vector in l_∞^2 to the two diagonal lines. While unit vectors form a circle in l_2^2, they form a square in l_∞^2. The total Euclidean length of the two projected vectors is constant, and is $\sqrt{2}$, since $||OD||_2 + ||OE||_2 = ||OM||_2$.

As shown in Fig. 2, given a unit length vector ($\overrightarrow{OC}$) in l_∞^2, the total Euclidean length of the two projected line segments onto the two diagonal lines ($\overrightarrow{OM}$ and $\overrightarrow{ON}$) is constant ($\sqrt{2}$).

Since $\sum_e (||e||_\infty f_e) < \sum_i (d_i r_i)$ by assumption, we have:

$$\sum_e (f_e (proj(e, \overrightarrow{OM}) + proj(e, \overrightarrow{ON})))$$

$$< \sum_i (r_i (proj(\overrightarrow{s_i t_i}, \overrightarrow{OM}) + proj(\overrightarrow{s_i t_i}, \overrightarrow{ON})))$$

From the inequality above, we can conclude that for at least one of $\overrightarrow{OM}$ and $\overrightarrow{ON}$, the projected network coding solution still has a smaller total cost than the cost of the projected MCF solution. □

Theorem 3.6. (Li and Wu, 2012 [4].) If there exists a configuration of k unicast sessions in l_2^h, for any $h \geq 2$, where $\sum_e (||e||_2 f_e) < \sum_i (||s_i t_i||_2 r_i)$, then there exists a configuration of k unicast sessions in l^1, where the same inequality holds.

Below we formulate a conjecture that generalizes Theorem 3.5 from l_∞^2 to l_∞^h for $h \geq 2$. It can also be viewed as the transformation of Theorem 3.6 from l_2^h to l_∞^h. Later we show that this conjecture implies the original multiple unicast network coding conjecture.

Conjecture 3.1. If there exists a configuration of k unicast sessions in l_∞^h for some $h \geq 2$, where $\sum_e (||e||_\infty f_e) < \sum_i (||s_i t_i||_\infty r_i)$, then there exists a configuration of k unicast sessions in l^1, where the same inequality holds.

D. Step 4. Prove Conjecture in 1-D: Integrating Cut Inequality

In a recent work, Li and Wu prove the following equivalence between network coding and routing in 1-D spaces. Their approach is to write an inequality for the cut condition at a point in l^1, and then integrate both sides of the inequality over all points in l_1.

Theorem 3.7. (Li and Wu, 2012 [4].) For any configuration of k unicast sessions in l^1, we always have $\sum_e(||e||_1 f_e) \geq \sum_i(||s_i t_i||r_i)$.

IV. Unified Proofs to Previous Results

In this section, we demonstrate the application of the geometric framework designed in Sec. III, by providing unified proofs to three known results of the multiple unicast conjecture. Existing proofs to these results are rather different from each other.

A. The Case of Two Unicast Sessions

Theorem 4.1. For two unicast sessions in an undirected network (G, c), network coding is equivalent to routing (MCF), *i.e.*, a throughput vector (r_1, r_2) is feasible with network coding if and only if it is feasible with routing.

Proof:

Step 1. Transformation: Apply Theorem 3.1 to all network configurations with $k = 2$, to translate the statement from its throughput version to cost version.

Step 2. Embedding: Apply Theorem 3.2, to translate the statement to be proven from the network information flow domain to the space information flow domain, from G to l_∞^n. A network (G, w) with n nodes has an isometric closure embedding into l_∞^n, as reviewed below.

Let u and v be two nodes in l_∞^n, at location $(x_{u1}, \ldots, x_{un})$ and $(x_{v1}, \ldots, x_{vn})$, respectively. The ∞-norm distance, or Chebyshev distance, between u and v is:

$$||u, v||_\infty = \lim_{p\to\infty} \left(\sum_{i=1}^{n} |x_{ui} - x_{vi}|^p\right)^{\frac{1}{p}} = \max_i |x_{ui} - x_{vi}|$$

We number the nodes in G and hence G' as u_1, u_2, ..., u_n. We can embed each node u_i, $1 \leq i \leq n$ by assigning it the coordinates $(x_{i1} = d_{i1}, x_{i2} = d_{i2}, \ldots, x_{ii} = d_{ii} = 0, \ldots, x_{i,n} = d_{i,n})$, where d_{ij} is the shortest path length between u_i and u_j in G. After such an embedding, we can verify that for any $1 \leq k \leq n$, $|d_{ik} - d_{jk}| \leq |d_{ij}|$ due to the triangular inequality satisfied by cost metric d in G', and hence $||u_i, u_j||_\infty = d_{i,j}$ by the definition of ∞-norm distance above.

Step 3. Projection: Apply Theorem 3.4 to reduce the space from l_∞^n to l_∞^k ($k = 2$), and then apply Theorem 3.5 to further reduce to l^1.

Step 4. 1-D Proof: Apply Theorem 3.7 to prove the statement in l^1, concluding the proof to Theorem 4.1. □

B. The $O(\log k)$ Upper-Bound in The General Case

We first introduce the following result on embedding a link-weighted graph into an Euclidean space:

Theorem 4.2. (Bourgain, 1985) The closure of an edge-weighted graph (G, w) with n nodes can be embedded into l_p for any $1 \leq p \leq \infty$, with distortion $O(\log n)$.

We now prove the following upper-bound on the coding advantage for multiple unicast in a general undirected network:

Theorem 4.3. For k unicast sessions in a undirected capacitated network (G, c) with n vertices, the coding advantage is upper-bounded by $O(\log k)$.

Proof: **Step 1. Transformation:** Apply Theorem 3.1, to translate the statement from throughput version to its cost version.

Step 2. Embedding: We apply Theorem 3.3 to translate the problem from G to a Euclidean space. We need a partial closure embedding of G with distortion $O(\log k)$.

We first apply Theorem 4.2 to the $2k$ terminal nodes, and obtain an embedding of these $2k$ nodes into a Euclidean space, with distortion $O(\log k)$. We then add the other $n - 2k$ nodes into the Euclidean space one by one. At each step, we make sure that the distance between the new node u and every existing node v is at most d_{uv}. Then we can conclude that if the cost advantage is 1 in the target Euclidean space, then it is upper-bounded by $O(\log k)$ before the embedding in G.

Step 3. Projection: Apply Theorem 3.6 to reduce the space from l_2^h to l^1. Here h is the dimension required for the embedding in Step 2 to be feasible.

Step 4. 1-D Proof: Apply Theorem 3.7 to prove the statement in l^1, concluding the proof to Theorem 4.3. □

C. Multiple Unicast in Star Networks

A network G is a *star network*, if there is a (center) node u in G, such that every other node is directly connected to u only. It has been previously studied in the literature of network coding for multiple unicast sessions [15].

Theorem 4.4. For k unicast sessions in an undirected network (G, c) with a star topology that satisfies the following property, network coding is equivalent to routing: for each session i, at least one of s_i or t_i locates at a node that is a source or destination of at most three sessions.

Proof:

Step 1. Transformation: Apply Theorem 3.1 to undirected star networks, to translate the statement from throughput version to its cost version.

Step 2. Embedding: We apply Theorem 3.3 to transform the problem from G to l_∞^2. We show a partial closure embedding of the star network (G, w) into l_∞^2, with $\beta = 1$, guaranteeing the distance between every pair of s_i and t_i remains unchanged during the embedding

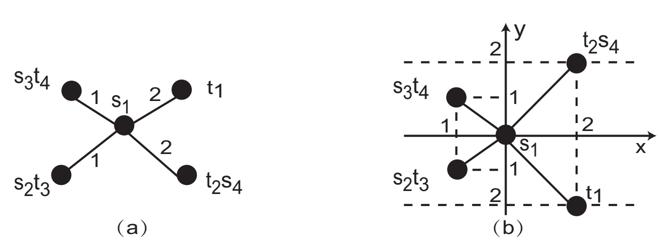

Fig. 3. Embedding a star network with heterogeneous cost into l_∞^2. (a) Original network G. (b) Embedding in l_∞^2.

As shown in Fig. 3, we first embed the center node to the origin O in l_∞^2. The other nodes are distributed on the four quadrants. For each pair of s_i and t_i, if neither is at center O, then embed the pair to different quadrants. The distance between this pair is $(d_{s_iO} + d_{t_iO})$, thus the pairwise weights remain unchanged. Given the condition in the theorem, such an isometric embedding into l_∞^2 is always possible.

Step 3. Projection: Apply Theorem 3.5 to reduce the space from l_∞^2 to l^1.

Step 4. 1-D Proof: Apply Theorem 3.7 to prove the statement in l^1, concluding the proof to Theorem 4.4. □

V. New Results in Cost Domain

In this section, we further apply the geometric framework from Sec. III to prove a number of new results.

A. Complete Networks

We prove that in a complete network with uniform cost, network coding can not outperform routing, for multiple unicast sessions.

Theorem 5.1. For k unicast sessions in a network (G, w), if G is a complete graph and w is a uniform cost vector, then the cost advantage is 1.

Proof:

Step 1. Transformation: In this case, we are proving network coding is equivalent to coding in the cost domain only. Step 1 in the framework does not apply.

Step 2. Embedding: We describe an isometric closure embedding of the uniform complete network G into l_2^n. For each vertex i, $i = 1, 2, \cdots, n$, let all the coordinates of i be zero, except that the ith coordinate is $\frac{\sqrt{2}}{2}$. Consequently, the distance between any two points is 1 in the target space, and we obtain an isometric embedding of G. We can then apply Theorem 3.2 to transform the problem from G to l_2^n.

Step 3. Projection: Apply Theorem 3.6 to reduce the space from l_2^n to l^1.

Step 4. 1-D Proof: Apply Theorem 3.7 to prove the statement in l^1, concluding the proof to Theorem 5.1. □

If the complete network (G, w) does not have a uniform cost in w, then we can similarly prove that the cost advantage is upper-bounded by the maximum heterogeneity in list costs $\max_{e_1, e_2 \in G} w(e_1)/w(e_2)$. The difference in the proof lies in Step 3, where we resort to an embedding with a distortion, instead of an isometric one.

B. Grid Networks

Theorem 5.2. For k pairs of unicast sessions in a 2-D square grid network (G, w) with uniform cost in w, the cost advantage is at most $\sqrt{2}$ (Fig. 4). If each pair of s_i and t_i is further on the same row or column, then cost advantage is 1.

Proof: **Step 1. Transformation:** Not applicable.

Step 2. Embedding: We perform a partial closure embedding of the grid network into l_2^2 in the straightforward way. The distortion is upper-bounded by $\sqrt{2}$. If each pair of s_i and t_i is on the same horizontal or vertical line, then we obtain an isometric partial closure embedding. Then we can apply Theorem 3.3 to transform the problem from G to l_2^2.

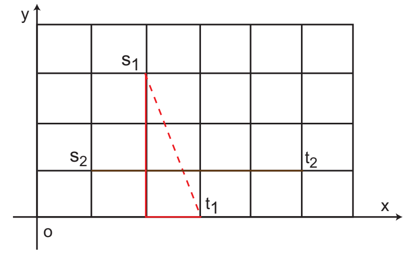

Fig. 4. The straightforward embedding of a 2-D square grid network into l_2^2.

Step 3. Projection: Apply Theorem 3.6 to reduce the space from l_2^2 to l^1.

Step 4. 1-D Proof: Apply Theorem 3.7 to prove the statement in l^1, concluding the proof to Theorem 5.2. □

The result in Theorem 5.2 can be enhanced and generalized in a number of directions. For instance, if the original network G is a uniform h-D grid instead of a 2-D grid, for some $h \geq 2$, then we can embed G into l_2^h with distortion $\sqrt{h}$, leading to an upper-bound of $\sqrt{h}$ on the cost advantage of network coding for multiple unicast sessions.

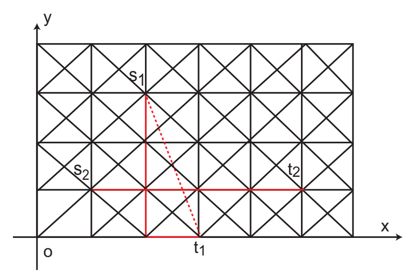

Fig. 5. Embedding a grid network with diagonal links into l_∞^2.

Furthermore, consider a 2-D uniform grid network G that further includes diagonal lines within all minimal squares, also with unit cost. We can embed the partial closure of G into l_∞^2 in an isometric fashion, as shown in Fig. 5. Here the isometric embedding is obtained by applying the most straightforward way of embedding G into a plane. Applying this as Step 2 in the framework, we can prove that network coding is equivalent to routing in G.

C. Layered Networks

A layered network is a generalization of a bipartite network into multi-partite, such that edges exist between neighboring partite/layers only, as shown in Fig. 6. We prove that, if links from each layer have uniform cost, then the cost advantage for multiple unicast is 1. If links from each layer have heterogeneous costs, the cost advantage can still be bounded by the degree of intra-layer cost heterogeneity.

In a layered network (G, w), let L be a layer of links. Define *cost heterogeneity* of layer L as $\max_{e, e' \in L} \frac{w(e)}{w(e')}$.

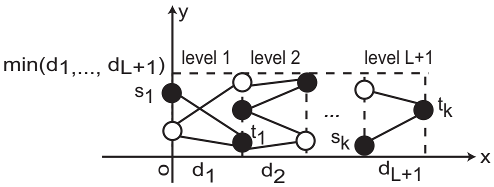

Fig. 6. Embedding a layered network (can be viewed as generalization of both a bipartite network and a combination network $C_{n,k}$) into l_∞^2. d_i is the uniform cost of links in layer i. If link costs in layer i is not uniform, we scale them to uniform before embedding, losing a factor equivalent to the cost heterogeneity.

Theorem 5.3. For k pairs of unicast sessions in an undirected layered network (G, w), with each pair of s_i and t_i lying at different layers, the cost advantage of network coding is 1 if each layer has uniform link costs. Otherwise, the cost advantage is upper-bounded by the maximum cost heterogeneity over all layers.

Proof: **Step 1. Transformation:** Not applicable.

Step 2. Embedding: We embed the partial closure of the layered network G into l_∞^2, as shown in Fig. 6. If each layer has uniform link cost, then in the embedding, the distance between any pair of s_i and t_i remains unchanged. The distance between any other pair of nodes can only decrease. If each layer has heterogeneous link costs, then the distortion of the embedding can be upper-bounded by the maximum cost heterogeneity over all layers. We can then apply Theorem 3.3 to transform the problem from G to l_∞^2.

Step 3. Projection: Apply Theorem 3.5 to reduce the space from l_∞^2 to l^1.

Step 4. 1-D Proof: Apply Theorem 3.7 to prove the statement in l^1, concluding the proof to Theorem 5.2. □

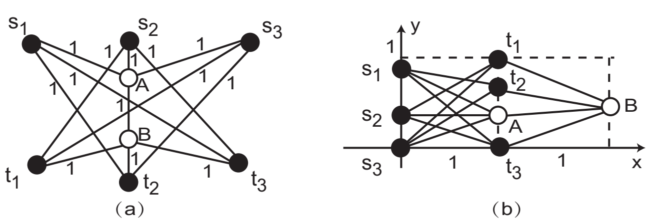

Fig. 7. A specific layered network and its embedding.

A special case of a layered network, as shown in Fig. 7(a), was used to demonstrate that network coding can have an arbitrarily large coding advantage for multiple unicast sessions [1]. There are k pairs of unicast sessions. Each s_i is connected to node A and all the receivers except t_i. A is connected to B. Each receiver t_i is connected to B as well as all the senders except s_i. If we assume each link has a unit cost (instead of a unit capacity [1]), Fig. 7(b) depicts the embedding of this network into l_∞^2. From Theorem 5.3, we know that network coding does not make a difference here, contrasting the arbitrarily large coding advantage under uniform link capacities.

VI. Conclusion

We applied a recently proposed tool, space information flow, to design a geometric framework for analyzing the multiple unicast conjecture, a well-known open problem in network coding. Based on the framework, we obtain unified proofs to a number of new results as well as existing results on the multiple unicast conjecture. We conclude by suggesting the following direction for proving the conjecture itself, based on the framework:

A possible proof to the multiple unicast conjecture
Step 1. Transformation: Apply Theorem 3.1 to translate the conjecture from its throughput version to cost version.
Step 2. Embedding: Based on the isometric closure embedding of G into l_∞^n, apply Theorem 3.3 to transform the problem from G to l_∞^n.
Step 3. Projection: Prove and then apply Conjecture 3.1, to reduce the problem from l_∞^n to l^1.
Step 4. 1-D Proof: Apply Theorem 3.7 to prove the statement in l^1, concluding the proof to the conjecture.

Acknowledgement

The last author would like to thank Michael Langberg, Sidharth Jaggi and Lap Chi Lau for helpful discussions.

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," CSAIL, MIT, MIT-LCS-TR-964, Tech. Rep., November 2004.
[3] Z. Li and C. Wu, "Space Information Flow," submitted to NetCod 2012.
[4] ——, "Space Information Flow: Multiple Unicast," submitted to ISIT 2012.
[5] X. Yin, Y. Wang, X. Wang, X. Xue, and Z. Li, "Min-Cost Multicast Network in Euclidean Space," submitted to ISIT 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] T. C. Hu, "Multi-commodity Network Flows," *Operations Research*, vol. 11, pp. 344–360, 1963.
[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," http://mybiasedcoin.blogspot.com, August 06 2007 post, Tech. Rep., August 2007.
[14] C. Chekuri, "Routing vs. Network Coding," University of Illinois, Urbana-Champaign, temple.birs.ca/~09w5103/chekuri_09w5103_talk.pptx, Tech. Rep., 2009.
[15] 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.