

A Geometric Perspective to Multiple-Unicast Network Coding

Tang Xiahou, Zongpeng Li, *Senior Member, IEEE*, Chuan Wu, *Member, IEEE*,
and Jiaqing Huang, *Member, IEEE*

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, and is now a well-known unsolved problem in the field of network coding. Based on a recently proposed tool of space information flow, 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.

Index Terms—Network coding, multiple-unicast, space information flow, geometric information flow, multicommodity flow.

I. INTRODUCTION

DEPARTING from the classic *store-and-forward* paradigm of data routing, 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 (assuming fractional flows are allowed). With network coding, the structure and the computational complexity of the optimal solution are largely unknown.

Manuscript received October 14, 2012; revised June 6, 2013; accepted January 31, 2014. Date of publication February 27, 2014; date of current version April 17, 2014. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada, in part by the National Natural Science Foundation of China under Grant 61271227, and in part by the University Grants Committee, Hong Kong, under Grant AoE/E-02/08. This paper was presented at the 2012 IEEE International Symposium on Information Theory and the 2012 IEEE International Symposium on Network Coding.

T. Xiahou and Z. Li are with the Department of Computer Science, University of Calgary, Calgary, AB T2N 1N4, Canada (e-mail: txiahou@ucalgary.ca; zongpeng@ucalgary.ca).

C. Wu is with the Department of Computer Science, University of Hong Kong, Hong Kong (e-mail: cwu@cs.hku.hk).

J. Huang is with the Huazhong University of Science and Technology, Wuhan 430074, China (e-mail: jqhuang@hust.edu.cn).

Communicated by T. Uyematsu, Associate Editor for Shannon Theory.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2014.2308998




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.




Fig. 2. (Coding solution on the left is an example from [10].) Three unicast sessions, each with target rate 1. All link capacities are 1. (A) Solution with network coding. (B) Solution without network coding - each of a_1, a_2, \dots, c_2 is an information flow of rate 0.5.

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 demand of 1. If each link has a unit capacity and a predefined direction as shown, 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. Fig. 2 shows another example of three unicast sessions, where a routing based solution (right) can achieve the same maximal throughput vector of $\{1, 1, 1\}$ under network coding (left). Harvey *et al.* [2] and Li and Li [1] conjectured that network coding is equivalent to routing for multiple-unicast in undirected networks, in the sense that any throughput vector

feasible under routing is feasible with network coding, and vice versa.

Despite a series of research effort devoted to it (*e.g.*, [11]–[13]), 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 [11], [12]. It is worth noting that such verification already involves new tools such as information dominance [11], input-output inequality and crypto inequality [12].

In 2007, Mitzenmacher *et al.* compiled a list of seven open problems in network coding [14], 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” [15, pp. 51–55]. 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.

Step 1 of the framework borrows an existing result from previous work [1]. Step 2 builds upon 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. 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].

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. In the throughput domain, it is known that the multiple-unicast conjecture is true under the following special cases: (a) when the number of unicast sessions is 2, (b) when all the unicast sessions have their sources or receivers co-located on the same node, and (c) when the network has a star topology. For the general case, it is further known that (d) the gap between network coding throughput and routing throughput is upper-bounded by a ratio of $O(\log k)$, where k is the number of unicast sessions. While existing proofs to results in (a)–(d) are rather different from one another, we utilize the proposed geometric framework to design unified proofs to all of them. The difference among our unified proofs lies in the choice of the target geometric space for network embedding,

and the technique in reducing the dimension of the geometric space. In the cost domain, the comparison between network coding and routing for multiple unicast sessions is relatively new. We examine a class of networks that include uniform complete networks, uniform or non-uniform grid networks and layered networks, and prove upper-bounds on the ratio by which network coding can save cost over routing. 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.

The rest of the paper is organized as follows. We review related literature in Sec. II, and go through preliminary material in Sec. III. Sec. IV describes the geometric framework in detail, including a number of results in each step of the framework, in preparation for Sec. V and Sec. VI, where existing results on the multiple-unicast conjecture in the throughput domain and new results in the cost domain are proved, respectively, in a unified way. Finally, Sec. VII concludes the paper, while suggesting a possible proof to the multiple-unicast conjecture based on the proposed geometric framework.

II. PREVIOUS RESEARCH

Harvey and Kleinberg [16] studied how various types of combinatorial cut conditions can approach the network coding throughput in an undirected network. They defined and examined three types of bounds, based on edge-sparsity, vertex-sparsity and meagerness, respectively. However, it is found that none of these bounds exactly or closely upper-bounds the achievable throughput of network coding.

Jain *et al.* [12] studied the multiple-unicast Conjecture in the Okamura-Seymour graph, a well-known example with four unicast sessions, where the cut bound is 1 (for each session) and routing can achieve a throughput of $3/4$ only. Applying tools in entropy calculus, including input-output equality and crypto inequality to the network at hand, they show that the max throughput with network coding is also $3/4$, and hence the conjecture is true there. Independently, Harvey *et al.* [11] also studied the Okamura-Seymour example, by blending graph theory and information theory techniques. Techniques from both Jain *et al.* and Harvey *et al.* can be applied to prove the multiple-unicast Conjecture in a special type of bipartite networks [11], [12].

Al-Bashabsheh and Yongacoglu [17] and Kramer and Savari [18] further studied the multiple-unicast Conjecture in another specific network, Hu’s 3-commodity network, where a gap exists between the routing throughput and the cut bound. The network has 6 vertices, 8 links and 3 unicast sessions. Based a close examination of the network structure, they apply input-output equality and submodularity of information entropy [17], and the d -separation technique from Bayesian networks [18], respectively, for proving that the max throughput both with and without network coding is $8/7$ for each unicast session.

Yazdi *et al.* [19], [20] studied the multiple-unicast problem in undirected ring networks. Based on an extension of the Japanese Theorem [21], [22], they show that the multiple-unicast Conjecture is true when all nodes in the network lie on a cycle.

Langberg and Médard [13] studied the multiple-unicast Conjecture, aiming at proving a relaxed version of it. They prove that under a certain connectivity assumption, throughput of network coding is at most three times throughout of routing. More precisely, if the network is well connected such that it is possible to concurrently multicast from every sender s_i to all receivers at rate r , then routing can achieve a unicast rate of $r/3$ for each of the unicast sessions.

Traskov *et al.* [23] studied the multiple-unicast problem in directed networks. They design two sub-optimal but practical code construction techniques for the network coding solution, based on linear programming and integer programming, respectively.

Wang and Shroff [24] propose and study *pairwise network coding* in directed networks, cyclic or not. They present a graph theoretic characterization of network coding restricted to two information flows, and discuss the sufficiency of linear network coding, the complexity of identifying coding opportunities, and the bandwidth efficiency of the solutions. Erez and Feder [25] studied a similar problem of two unicast sessions in directed acyclic networks. They construct a backpressure style distributed solution based on network coding, for improving upon multicommodity flow based solutions.

In a recent work, Huang and Ramamoorthy [26] studied 3 unicast sessions in directed acyclic networks. Combining entropy arguments with network topology constraints, they characterize a number of cases where rates are achievable or not achievable by network coding.

III. 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 , $1 \leq i \leq k$. The k unicast sessions are independent, and have a desired throughput vector $r = (r_1, \dots, r_k)$. An orientation (V, A) of G is specified by a set of nodes V that is the same as in G , and a set of directed links $A = \{\vec{uv}, \vec{vu} \mid \forall (u, v) \in E\}$. A flow vector $f \in \mathcal{Q}_+^A$ specifies the rate of information flow transmitted at each directed link in A . If $e \in A$ is a directed link, then $f(e)$ is the scalar component in f that specifies that flow rate at e . If $e = (u, v) \in E$ is an undirected link then $f(e) = f(\vec{uv}) + f(\vec{vu})$ is the total flow rate between u and v .

A. Multiple-Unicast: Network Information Flow

In the max-throughput version of the multiple-unicast problem, we are given a capacitated network (G, c) , and wish to maximize a ratio $\alpha \geq 0$, such that the throughput vector αr can be achieved. Let α_{NC} and α_R 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_R .

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), each unicast session can be routed separately since there is no

inter-session coding. The minimum cost for each session i is $d_i r_i$, where d_i is the length of the shortest path between s_i and t_i in G , assuming each link e has cost $w(e)$. The minimum total cost under routing for achieving a throughput vector r is therefore $\sum_i (d_i r_i)$. 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 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 $\sum_i (d_i r_i) / \sum_e (w(e) f^*(e))$.

B. Multiple-Unicast: Space Information Flow

Space information flow (geometric information flow) is a new subject of study being proposed [3], [5], which can be viewed as a generalization of the geometric Steiner tree problem by introducing network coding. 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 geometric 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]. 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.

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}, \dots, x_{uh}) and (x_{v1}, \dots, x_{vh}) , respectively, the distance between u and v is:

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

The superscript h , the subscript p , or both, may be omitted when their specific values are not important or clear from the context.

For the multiple-unicast version of the space information flow problem, we are given k pairs of terminals, (s_i, t_i) , $1 \leq i \leq k$, in a space l_p^h . We seek the min-cost solution that can achieve a throughput vector 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 $\sum_i (\|s_i, t_i\|_p r_i)$, due to the triangular inequality of distances in l_p^h , known as Minkowski inequality in the literature. Under network coding, let f^* be the underlying flow vector of the optimal solution. For an edge e between two nodes u and v , let $\|e\|_p = \|u, v\|_p$ be the length of e . The cost is then $\sum_e \|e\|_p f^*(e)$. The cost advantage of network coding here is $\sum_i (\|s_i, t_i\|_p r_i) / \sum_e (\|e\|_p f^*(e))$.

Fig. 3(a) shows a multicast version of the space information flow problem in a 2-D Euclidean space [27]. Among the




Fig. 3. A multicast space information flow problem in a 2-D Euclidean space (example taken from [27]), where network coding outperforms routing, *i.e.*, the cost advantage is strictly larger than 1. The target multicast throughput is 1 bps. (a) Six multicast terminals. Radius of circle is 1. (b) Minimum Euclidean Steiner tree, cost is 4.64. (c) Minimum multicast flow with network coding, cost is 4.57.




Fig. 4. 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)?

six multicast terminals (A-F), five (A-E) are evenly distributed along a circle centered at F. Any of the six node can be chosen as the multicast source, with the other five being receivers. Fig. 3(b) shows the optimal solution with routing, which has cost 4.64. Three relay nodes are inserted. Fig. 3(c) shows the optimal solution with network coding, which has cost 4.57. Five relay nodes are inserted. Despite its small value, the gap between the two optimal costs reveals that multicast with network coding is a fundamentally different problem from geometric Steiner trees, with a different problem structure, and perhaps a different computational complexity.

For the multiple-unicast version 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. 4. 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, $\sum_e ||e||f(e)$. Here e is a link employed for flow transmission, $||e||$ is the length of e in space, and $f(e)$ 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], [5]. What about multiple-unicast?

C. Paradigm Comparison

Given a space information flow vector f , a network can be induced, over the same nodes and links as in f , by viewing $f(e)$ as the capacity of e . The distance of e is denoted as $||e||$. The cost of f is then $\sum_e ||e||f(e)$.

This reflects the general rule that the longer and the wider a communication cable is, the more expensive it is. For the sake of cost minimization, apparently, only straight line segments (or geodiscs in manifolds without a global coordinate system) need to be considered in f .

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 ζ_d , ζ_u and ζ_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 3.1. $\zeta_d \geq \zeta_u \geq \zeta_s$.

Proof: We first show that $\zeta_d \geq \zeta_u$. Given the maximum cost advantage ζ_u^* in undirected networks, let Δ_u be a problem instance where this cost advantage is achieved, and let 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 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 , since f^* constitutes an optimal solution for both. Without network coding, the cost of the optimal solution can only increase from Δ_u to Δ_d , since the latter is more restrictive — a routing solution faces extra constraints in the form of predefined link directions. Therefore $\zeta_d \geq \zeta_u$.

We next show that $\zeta_u \geq \zeta_s$. Given the maximum cost advantage ζ_s^* in geometric spaces, let Δ_s be a problem instance where this cost advantage is achieved, and let f^* be the underlying flow of the optimal network coding solution. We can create a corresponding problem instance Δ_u for the undirected network setting, by viewing f^* as the underlying network topology. The orientation of links in f^* is cancelled. If two links appear between the same pair of nodes as a result, they are merged into a single link with the sum capacity. The set of terminal nodes and the target throughput remain unchanged. The cost of a link between two nodes u and v , $w(uv)$, is taken as the geometric distance $||u, v||_p^h$. With network coding, the cost of the optimal solution is the same in Δ_s and in Δ_u , since f^* constitutes an optimal solution for both. Without network coding, the cost of the optimal solution can only increase from Δ_s to Δ_u , since the latter is more restrictive — a routing solution faces extra constraints in the form of a predefined network topology (including the number of relay nodes, and the interconnection among the terminal nodes and relay nodes). Therefore $\zeta_u \geq \zeta_s$. \square

Given Theorem 3.1, 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 [28]–[30]. This bound automatically holds for multicast in a geometric space of any dimension, with normed distances.

D. Discussions

At a first glance, it may appear that the space information flow problem is less *rich* than the network information flow

problem, in that only terminal nodes present in the problem input, and a detailed network topology including relay nodes and link connections are not specified. Furthermore, one may naturally worry that transforming the multiple-unicast problem from the network domain into the space domain (as we describe in Sec. IV) makes the problem trivial and less interesting, because the interconnection information essential for a network is lost; and hence, the geometric perspective may not be helpful in proving the original multiple-unicast network coding conjecture. Surprisingly, results we prove suggest a different picture. In particular, we show the following relations between the network version and the geometric version of the multiple-unicast conjecture:

1. If the conjecture is true in networks, then it is true in any geometric space with a properly defined notion of distance, satisfying non-negativity, symmetry, and triangular inequality (implied from Theorem 3.1).
2. In particular, if the conjecture is true in networks, then it is true in an Euclidean space of any dimension. If the conjecture is true in all Euclidean spaces, then it is *approximately true* in networks (the gap between network coding throughput and routing throughput cannot be too large).
3. The conjecture is true in all networks **if and only if** it is true in all geometric spaces with Chebyshev distance.

Item 3. above shows that the multiple-unicast problem carries over its full flavor from the network domain to the space domain.

In the space information problem, for either multiple-unicast or multicast, one may generalize the flow cost function from linear to nonlinear, including sub-linear and super-linear versions. However, we restrict our attention to linear cost only in this paper. The linear cost model is a natural and basic model to start with, and it also appears sufficient for our purpose of translating the multiple-unicast conjecture into its geometric version.

IV. 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 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 r , in a cost-weighted undirected network (G, \mathbf{w}) . Then $\sum_e (w(e)f(e)) \geq \sum_i (d_i r_i)$.

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 4.1. (Li and Li, 2004 [1]) Given an undirected network G with k pairs of unicast terminals specified, and any desired throughput vector 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 4.1 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 metric closure of the link-weighted network G . It is a complete network defined over the same set of vertices as in G , such that the cost of a link $e = (u, v)$ 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 $\|u, v\|_p \leq d(uv) \leq \beta \cdot \|u, v\|_p, \forall u, v \in V$. 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 4.2. For k pairs of unicast sessions in an undirected network (G, w) , with desired throughput vector r , assume G has a β -distortion **closure embedding** in a space l_p^h . If the cost advantage is $\zeta \geq 1$ after the embedding, then it is upper-bounded by $\beta\zeta$ before the embedding.

Proof: Assume, by way of contradiction, that there is a network coding solution in G , with an underlying flow vector f satisfying $\beta\zeta \sum_e (w(e)f(e)) < \sum_i (d_i r_i)$, 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 $\beta\zeta \cdot \sum_e (\|e\|_p f'(e)) < \beta \cdot \sum_i (\|s_i, t_i\|_p r_i)$ due to the β -distortion property of the embedding, contradicting the assumption that the cost advantage after the embedding is ζ . \square

A similar result holds for partial embedding as well.

Theorem 4.3. 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 ζ after the embedding, then it is upper-bounded by $\beta\zeta$ before the embedding.

The proof of Theorem 4.3 is similar to that of Theorem 4.2, 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 4.2 and Theorem 4.3 is when $\beta = 1$, then cost




Fig. 5. Projection from l_∞^n to l_∞^k : (1) for each session $i \in \{1, \dots, k\}$, identify primary coordinate for session i ; (2) for each node u , remove its $(n - k)$ non-primary coordinates and keep its k primary coordinates.

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 4.4. 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 $\text{argmax}_j |x_{s_i, j} - x_{t_i, j}|$. We project the original multiple-unicast instance from l_∞^n to l_∞^k by truncating the coordinate of each point in the following way: keep k coordinates including all the primary coordinates, dropping other coordinates. More specifically, let \mathcal{J} be the set of all primary coordinates. The size of \mathcal{J} is at most k . Each point $(x_j, j = 1, \dots, n)$ in the original n -D space is mapped to a point in the $|\mathcal{J}|$ -D space of coordinates $(x_j, j \in \mathcal{J})$. Fig. 5 illustrates such a projection from l_∞^n to l_∞^k with an example of five nodes and two unicast sessions.

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. \square

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 4.5.** 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{p} and \vec{q} be two vectors in a space l_p^h . We define the projection of \vec{p} onto \vec{q} as the Euclidean projection $\text{proj}(\vec{p}, \vec{q}) = \frac{\vec{p} \cdot \vec{q}}{\|\vec{q}\|_2^2}$, where \cdot is the inner product operation and $\vec{p} \cdot \vec{q} = \|\vec{p}\|_2^h \|\vec{q}\|_2^h \cos \theta$, with θ being the angle between \vec{p} and \vec{q} .

As shown in Fig. 6, given a unit length vector (\vec{OC}) in l_∞^2 , the total Euclidean length of the two projected line segments onto the two diagonal lines $(\vec{OM}$ and $\vec{ON})$ is constant ($\sqrt{2}$).




Fig. 6. 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 $\|OD\|_2^2 + \|OE\|_2^2 = \|OM\|_2^2$.




Fig. 7. Project and Integrate over all possible directions \vec{p} .

Therefore we know that for any link e in l_∞^2 ,

$$\frac{\text{proj}(e, \vec{OM}) + \text{proj}(e, \vec{ON})}{\|e\|_\infty^2} = \sqrt{2}$$

Since $\sum_e (\|e\|_\infty^2 f(e)) < \sum_i (d_i r_i)$ by assumption, we have:

$$\begin{aligned} & \sum_e (f(e) (\text{proj}(e, \vec{OM}) + \text{proj}(e, \vec{ON}))) \\ & < \sum_i (r_i (\text{proj}(\vec{s}_i \vec{t}_i, \vec{OM}) + \text{proj}(\vec{s}_i \vec{t}_i, \vec{ON}))) \end{aligned}$$

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

Theorem 4.6. If there exists a configuration of k unicast sessions in l_2^h , for any $h \geq 2$, where $\sum_e (\|e\|_2^h f(e)) < \sum_i (\|s_i, t_i\|_2^h r_i)$, then there exists a configuration of k unicast sessions in l^1 , where the same inequality holds.

Proof: Given the $h - D$ problem instance, such that $\sum_e (f(e) \|e\|_2^h) < \sum_i (\|s_i, t_i\|_2^h r_i)$, we construct a $1 - D$ k pairs unicast instance and its network coding solution by projecting their counterparts 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 $\sum_e (f(e) \|e\|_2^1) < \sum_i (\|s_i, t_i\|_2^1 r_i)$.

As shown in Fig. 7, 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 \vec{p} be the vector from origin to the corresponding point on Φ , let $\vec{1}$ be the h -D unit vector $(1, 0, 0, \dots, 0)$.




Fig. 8. Integrated projection of a link e over all possible directions is constant, regardless of the orientation of e — an example of projecting from l_2^2 to l^1 . The integration equals the area of the shaded region between 0 and 2π . The orientation of e affects the phase of the cosine curve but not the area. Here x -axis is the orientation of \vec{p} in the 2-D space, and the absolute value of y -axis is $e \cdot \vec{p}$.

The integration over the closed surface Φ for all the projections of f is:

$$\begin{aligned} \iint_{\Phi} \sum_e (f(e)(e \cdot \vec{p})) d\Phi &= \sum_e \iint_{\Phi} f(e)(e \cdot \vec{p}) d\Phi \\ &= \sum_e \iint_{\Phi} f(e) \|e\|_2^h (\vec{1} \cdot \vec{p}) d\Phi \\ &= \sum_e (f(e) \|e\|_2^h) \iint_{\Phi} (\vec{1} \cdot \vec{p}) d\Phi \end{aligned}$$

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 ($=_2$), as illustrated in Fig. 8.

The integration over the closed surface Φ for all the projections of $\{s_i t_i r_i | i = 1, \dots, k\}$ is:

$$\begin{aligned} \iint_{\Phi} \sum_i (s_i t_i \cdot \vec{p}) r_i d\Phi &= \sum_i \iint_{\Phi} (s_i t_i \cdot \vec{p}) r_i d\Phi \\ &= \sum_e \iint_{\Phi} (\|s_i, t_i\|^h (\vec{1} \cdot \vec{p})) r_i d\Phi \\ &= \sum_i \|s_i, t_i\|^h \iint_{\Phi} (\vec{1} \cdot \vec{p}) r_i d\Phi \end{aligned}$$

Since $\sum_e (f(e) \|e\|^h) < \sum_i \|s_i, t_i\|^h$ by assumption, we claim that:

$$\iint_{\Phi} \sum_e (f(e)(e \cdot \vec{p})) d\Phi < \iint_{\Phi} \sum_i (s_i t_i \cdot \vec{p}) r_i d\Phi$$

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

$$\sum_e (f(e)(e \cdot \vec{p}^*)) < \sum_i (s_i t_i \cdot \vec{p}^*) r_i \quad \square$$

Below we formulate a conjecture that generalizes Theorem 4.5 from l_{∞}^2 to l_{∞}^h for $h \geq 2$. It can also be viewed as the transformation of Theorem 4.6 from l_2^h to l_{∞}^h . Later we show that this conjecture implies the original multiple-unicast network coding conjecture.

Conjecture 4.1. If there exists a configuration of k unicast sessions in l_{∞}^h for some $h \geq 2$, where $\sum_e (\|e\|_{\infty}^h f(e)) <$




Fig. 9. Three unicast sessions in 1-D. Total flow crossing point x_0 , $f_{x_0}^1$, is lower-bounded by $Demand((-\infty, x_0); (x_0, \infty)) = \mathbf{r}_1 + \mathbf{r}_2$.

$\sum_i (\|s_i, t_i\|_{\infty}^h r_i)$, then there exists a configuration of k unicast sessions in l^1 , where the same inequality, with dimension h replaced by 1, holds.

Theorem 4.6 and Theorem 4.7 in the next subsection show that the multiple-unicast network coding conjecture is true for Euclidean spaces of any dimension (in the literature, an Euclidean space of dimension higher than 3 is sometimes referred to as a Hilbert space). We recall that the celebrated Nash Embedding Theorem states that every Riemannian manifold can be isometrically embedded into some Euclidean space [31]. However, this does not directly imply a proof of the multiple-unicast network coding conjecture in Riemannian manifolds, since: the isometry claimed in Nash Embedding Theorem is between geodesic distances before and after embedding, not between geodesic distances before embedding and straight-line distances after embedding. In fact, it may be interesting to study (a) the possibility of isometric closure embedding of an undirected graph into a Riemannian manifold, and (b) the correctness of the multiple-unicast conjecture in Riemannian manifolds.

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

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 $f(e)$ 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 geometric spaces, confirming that any solution, with network coding or not, can not break through the throughput-distance sum-product barrier in a 1-D space.

Theorem 4.7. Given k independent unicast sessions in 1-D space, let f^1 be the underlying flow vector of a network coding solution achieving a rate vector r . Then $\sum_e (\|e\|^1 f(e)) \geq \sum_i (\|s_i, t_i\|^1 r_i)$.

Proof: For a given point x in the 1-D space, let f_x^1 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 f_x^1 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 quantities over the entire 1-D space, and obtain:

$$\begin{aligned} \int_{x=-\infty}^{\infty} f_x^1 dx &\geq \int_{x=-\infty}^{\infty} Demand((-\infty, x); (x, \infty)) dx \\ &= \sum_i \|s_i, t_i\|^1 r_i \end{aligned}$$

Furthermore, note that $\sum_e (\|e\|^1 f(e)) = \int_{x=-\infty}^{\infty} f_x^1 dx$. We conclude that $\sum_e (\|e\|^1 f(e)) \geq \sum_i (\|s_i, t_i\|^1 r_i)$. \square

V. UNIFIED PROOFS TO PREVIOUS RESULTS

In this section, we demonstrate the application of the geometric framework designed in Sec. IV, by providing unified proofs to a few known results of the multiple-unicast conjecture in the throughput domain.

A. The Cases of Two Unicast Sessions and Co-Located Terminals

A set of vertices $\Gamma \subseteq V$ covers the k unicast sessions if $\forall i \in \{1, \dots, k\}, s_i \in \Gamma$ or $t_i \in \Gamma$. In other words, a cover of the k unicast sessions is a set of nodes that includes at least one of the two terminal nodes from each of the k unicast sessions. Such a cover is a vertex cover of the demand graph [32] for the multiple-unicast problem.

We next prove that, if there exists a two-node cover Γ , *i.e.*, $|\Gamma| = 2$, then network coding is equivalent to routing. From this result, we derive as corollaries that the multiple-unicast conjecture is true in the following two cases, which are known in the literature: (a) if the number of unicast sessions k is 2, and (b) if all the k unicast sessions have their sources (or receivers) co-located at a common node.

Theorem 5.1. For k unicast sessions in an undirected network (G, c) , if there exists a two-node cover $\Gamma = \{u, v\} \subseteq V$, then network coding is equivalent to routing (MCF), *i.e.*, a throughput vector r is feasible with network coding if and only if it is feasible with routing.

Proof: **Step 1. Transformation:** Apply Theorem 4.1 to all network configurations with a two-node cover, to translate the statement from its throughput version to cost version.

Step 2. Embedding: Apply Theorem 4.2, 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}, \dots, x_{un}) and (x_{v1}, \dots, x_{vn}) , respectively. The ∞ -norm distance, or Chebyshev distance, between u and v is:

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

We number the nodes in G and hence G' as u_1, u_2, \dots, u_n . We can embed each node u_i , $1 \leq i \leq n$ by assigning it the coordinates $(x_{i1} = d_{i1}, x_{i2} = d_{i2}, \dots, x_{ii} = d_{ii} = 0, \dots, x_{in} = d_{in})$, 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: We can apply a primary coordinate based dimension reduction technique, similar to that used in the proof of Theorem 4.4, for reducing the geometric space from l_∞^n to l_∞^2 . Given the two-node cover $\Gamma = \{u, v\}$, we can select as two primary coordinates the two coordinates in l_∞^n that correspond to distances to nodes u and v in G respectively (recall that each coordinate in l_∞^n denotes distances from some

node to a given node in G , by the way G is embedded into l_∞^n). Then, we can project each node u_i from l_∞^n to l_∞^2 by truncating its $(n - 2)$ non-primary coordinates, resulting in a 2-D coordinate (x_{iu}, x_{iv}) .

Next, we apply Theorem 4.5 to further reduce the geometric space from l_∞^2 to l^1 .

Step 4. 1-D Proof: Apply Theorem 4.7 to prove the statement in l^1 , concluding the proof to Theorem 5.1. \square

Corollary 5.1. 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: This can be proved by applying Theorem 5.1, while selecting $\Gamma = \{s_1, s_2\}$ to be the cover. \square

Corollary 5.2. If all the k unicast sessions have their sources (or receivers) co-located at a common nodes, then network coding is equivalent to routing (MCF), *i.e.*, a throughput vector r is feasible with network coding if and only if it is feasible with routing.

Proof: Let u be the node in G where the k sources (or receivers) are co-located. Choose another node $v \neq u$ from G arbitrarily. Let $\Gamma = \{u, v\}$. By the definition of a node cover, Γ is a valid node cover for the k unicast sessions that have co-located sources (receivers). Applying Theorem 5.1, we can conclude that network coding is equivalent to routing in this case. \square

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

Theorem 5.2. 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 4.1, to translate the statement from throughput version to its cost version.

Step 2. Embedding: We translate the problem from G to a geometric space with l_1 distance, where:

$$\|u, v\|_1 = \left(\sum_i |x_{ui} - x_{vi}| \right)^{\frac{1}{1}} = \sum_i |x_{ui} - x_{vi}|$$

Recall that $d(u, v)$ denotes the shortest path distance between two nodes u and v in G . For each node u in G , map u to $(d(u, A_i) | i = 1, \dots, O(\log^2 k))$, where each A_i is a randomly chosen subset of the terminals $T = \{s_i, t_i | 1 \leq i \leq k\}$ generated in the following way: for each $h < k$ that is a power of 2, randomly pick $O(\log h)$ sets $A \subseteq T$ of cardinality h ; map u to $(d(u, A_i))$ where $d(u, A) = \min_{v \in A} d(u, v)$. It is proved that the mapping satisfies *almost surely* both (i) $\|u, v\|_1 \leq d(u, v)$, $\forall u, v \in V$, and (ii) $\frac{d(u, v)}{O(\log^2 k)} \leq \|u, v\|_1 \leq d(u, v)$, $\forall u, v \in T$ ([33], Corollary 3.4). If the embedding fails to satisfy both (i) and (ii), one can repeat the randomized embedding until (i) and (ii) are satisfied, leading to a $O(\log k)$ -distortion partial embedding of G into $l_1^{O(\log^2 k)}$.

Step 3. Projection: In this step we aim at reducing the space from $l_1^{O(\log^2 k)}$ to l^1 . If the total size of the network coding solution f is smaller than that of the routing solution in

$l_1^{O(\log^2 k)}$ to l^1 , i.e.:

$$\sum_{i=1}^{O(\log^2 k)} \sum_{e=(u,v) \in f} (|x_{ui} - x_{vi}| f(e)) < \sum_{i=1}^{O(\log^2 k)} \sum_{j=1}^k (|x_{s_j i} - x_{t_j i}| r_i),$$

there must be at least one particular i , such that $1 \leq i \leq O(\log^2 k)$, and

$$\sum_{e=(u,v) \in f} (|x_{ui} - x_{vi}| f(e)) < \sum_{j=1}^k (|x_{s_j i} - x_{t_j i}| r_i),$$

meaning that there is a problem instance in l^1 , where the network coding cost is smaller than the routing cost.

Step 4. 1-D Proof: Apply Theorem 4.7 to prove the statement in l^1 , concluding the proof to Theorem 5.2. \square

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 [34].

Theorem 5.3. 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 4.1 to undirected star networks, to translate the statement from throughput version to its cost version.

Step 2. Embedding: We apply Theorem 4.3 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 (a) the distance between every pair of s_i and t_i remains unchanged during the embedding, and (b) the length of every edge e is upper bounded by $w(e)$ in the network.

In the first step of the embedding, we map the center node O of the star network to the origin of l_∞^2 .

In the second step, we randomly place non-center nodes into one of the four quadrants of l_∞^2 . For each of the non-center node u , select one of the four quadrants randomly and place u along the diagonal line of that quadrant, with distance to each of the two axes being $w(uO)$ in G . For instance, if the second quadrant is picked, then the coordinates assigned to u are $(-w(uO), w(uO))$.

The third step resolves potential conflicts that arise when s_i and t_i are mapped into the same quadrant, for some unicast session $i \in \{1, \dots, k\}$. If such a session i exists, then by the assumption in the theorem, either s_i or t_i is located at a node that is a terminal for at most three unicast sessions. Without loss of generality, let's assume it's s_i , which locates at a node v in G . v hosts at most two other unicast sessions, i' and i'' . We relocate node v to a quadrant that is different from both (i) its current quadrant, and (ii) the quadrant that hosts the other terminal of unicast sessions i' and i'' . Fig. 10 shows an example of such embedding.

For each pair of terminals s_i and t_i , if one of them resides at the center O , then $e = (s_i, t_i)$ is an edge in G and the shortest




Fig. 10. Embedding a star network with heterogeneous cost into l_∞^2 . (a) Original network G . (b) Embedding in l_∞^2 .

s_i - t_i path in G is the link (s_i, t_i) because G is a star network and there are no alternative paths. Furthermore, $d_i = w(e)$, and $\|s_i, t_i\|_\infty^2 = \max\{w(e), w(e)\} = w(e) = d_i$. If neither s_i or t_i is at the center O , then they are mapped to two different quadrants in l_∞^2 , and it can be verified that $\|s_i, t_i\|_\infty^2 = \|s_i, O\|_\infty^2 + \|t_i, O\|_\infty^2 = w(s_i, O) + w(O, t_i) = d_i$.

The last step of verifying the validity of the partial closure embedding is to show that for every edge (u, v) in G , $\|u, v\|_\infty^2 \leq d_{uv}$. If one of u and v is the center node, then $\|u, v\|_\infty^2 = w(u, v) = d_{uv}$ by the way the embedding is defined. If neither u or v is the center node, then $\|u, v\|_\infty^2 = d_{u,v}$ if u and v are mapped to different quadrants, and $\|u, v\|_\infty^2 < d_{u,v}$ if u and v are mapped to the same quadrant.

Step 3. Projection: Apply Theorem 4.5 to reduce the space from l_∞^2 to l^1 .

Step 4. 1-D Proof: Apply Theorem 4.7 to prove the statement in l^1 , concluding the proof to Theorem 5.3. \square

VI. NEW RESULTS IN COST DOMAIN

In this section, we further apply the geometric framework from Sec. IV 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 6.1. 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 , $i = 1, 2, \dots, n$, let all the coordinates of i be zero, except that the i th 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 4.2 to transform the problem from G to l_2^n .

Step 3. Projection: Apply Theorem 4.6 to reduce the space from l_2^n to l^1 .

Step 4. 1-D Proof: Apply Theorem 4.7 to prove the statement in l^1 , concluding the proof to Theorem 6.1. \square

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




Fig. 11. The straightforward embedding of a 2-D square grid network into l_2^2 .




Fig. 12. Embedding a grid network with diagonal links into l_∞^2 .

$\max_{e_1, e_2 \in G} w(e_1)/w(e_2)$. 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 6.2. 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. 11). 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 4.3 to transform the problem from G to l_2^2 .

Step 3. Projection: Apply Theorem 4.6 to reduce the space from l_2^2 to l^1 .

Step 4. 1-D Proof: Apply Theorem 4.7 to prove the statement in l^1 , concluding the proof to Theorem 6.2. \square

The result in Theorem 6.2 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.

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. 12. 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 .




Fig. 13. 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 are not uniform, we scale them to uniform before embedding, losing a factor equivalent to the cost heterogeneity.

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. Specifically, nodes in a layered network G can be partitioned into $V = V_1 \oplus \dots \oplus V_{L+1}$, and links in G can be partitioned into $E = E_1 \oplus \dots \oplus E_L$, such that for each link $e = (uv) \in E_l$, for any $1 \leq l \leq L$, $u \in V_l$ and $v \in V_{l+1}$. A bipartite network is a special case where $L = 2$. A combination network is a special case where $L = 3$ and the first layer has one node only.

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 — the *cost heterogeneity* of layer l is $\rho_l = \max_{e, e' \in E_l} \frac{w(e)}{w(e')}$.

Theorem 6.3. For k unicast sessions in a layered network (G, w) where each pair of s_i and t_i are located at different layers, the cost advantage of network coding is upper-bounded by $\rho = \max_l \rho_l$, the maximum intra-layer link cost heterogeneity.

Proof: Step 1. Transformation: Not applicable.

Step 2. Embedding: We embed the layered network G into l_∞^2 , as shown in Fig. 13. First, map each layer of nodes V_i to a vertical line in l_∞^2 , such that all nodes in V_i share the same x -coordinate x_i , and $x_{i+1} - x_i = \min_{e \in E_i} w(e)$. For the y -coordinate, we place each node in V_i randomly in the range $[0, \min_i(x_{i+1} - x_i)]$, making sure that for each link $e = (uv) \in E$, the Chebyshev length of e is measured as the difference between the x -coordinates of u and v . We can verify that such an embedding satisfies both requirements of a ρ -distortion partial closure embedding: (i) $\frac{d(s_i, t_i)}{\rho} \leq \|s_i, t_i\|_\infty^2 \leq d(s_i, t_i)$, $\forall 1 \leq i \leq k$, and (ii) $\|e\|_\infty^2 \leq w(e)$, $\forall e \in E$. Then Theorem 4.3 can be applied to transform the problem from G to l_∞^2 .

Step 3. Projection: Apply Theorem 4.5 to reduce the space from l_∞^2 to l^1 .

Step 4. 1-D Proof: Apply Theorem 4.7 to prove the statement in l^1 , concluding the proof to Theorem 6.2. \square

A special case of a layered network, as shown in Fig. 14(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 source s_i is connected to node A and every receiver t_j for $1 \leq j \leq k$ and $j \neq i$. Each receiver t_i is connected to B and every source s_j for $1 \leq j \leq k$ and $j \neq i$. If we assume each link has a unit cost (instead of a unit capacity [1]), Fig. 14(b) depicts the embedding of this network into l_∞^2 . From Theorem 6.3,




Fig. 14. A specific layered network (a) and its embedding (b).

we know that network coding does not make a difference here, contrasting the arbitrarily large coding advantage under uniform link capacities.

VII. 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. Our studies show that the cost version of the multiple-unicast conjecture is true in Euclidean/Hilbert spaces, in Riemannian manifolds and in l_1 spaces. We conclude by suggesting the following direction for proving the conjecture itself, through transforming the original conjecture in networks into l_∞ spaces:

A possible proof to the multiple-unicast conjecture

Step 1. Transformation: Apply Theorem 4.1 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 4.3 to transform the problem from G to l_∞^n .

Step 3. Projection: Prove and then apply Conjecture 4.1, to reduce the problem from l_∞^n to l^1 .

Step 4. 1-D Proof: Apply Theorem 4.7 to prove the statement in l^1 , concluding the proof to the conjecture.

ACKNOWLEDGMENT

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

REFERENCES

- [1] Z. Li and B. Li, "Network coding: The case of multiple unicast sessions," in *Proc. 42nd Annu. Allerton Conf. Commun., Control, Comput.*, 2004, pp. 1–9.
- [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, Cambridge, MA, USA, Tech. Rep. MIT-LCS-TR-964, Nov. 2004.
- [3] Z. Li and C. Wu, "Space information flow," Dept. Comput. Sci., Univ. Calgary, Alberta, AB, Canada, Tech. Rep. 2012-0728-sif, 2012.
- [4] Z. Li and C. Wu, "Space information flow: Multiple unicast," in *Proc. IEEE ISIT*, Jul. 2012, pp. 1897–1901.
- [5] X. Yin, Y. Wang, X. Wang, X. Xue, and Z. Li, "Min-cost multi-cast network in Euclidean space," in *Proc. IEEE ISIT*, Jul. 2012, pp. 1316–1320.
- [6] R. Ahlswede, N. Cai, S. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inf. Theory*, vol. 46, no. 4, pp. 1204–1216, Jul. 2000.
- [7] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Trans. Netw.*, vol. 11, no. 5, pp. 782–795, Oct. 2003.
- [8] X. Yan, R. W. Yeung, and Z. Zhang, "The capacity region for multi-source multi-sink network coding," in *Proc. IEEE ISIT*, Jun. 2007, pp. 116–120.
- [9] T. C. Hu, "Multi-commodity network flows," *Oper. Res.*, vol. 11, no. 3, pp. 344–360, 1963.
- [10] R. Dougherty and K. Zeger, "Nonreversibility and equivalent construction of multiple-unicast networks," *IEEE Trans. Inf. Theory*, vol. 52, no. 11, pp. 5067–5077, Nov. 2006.
- [11] N. J. A. Harvey, R. Kleinberg, and A. R. Lehman, "On the capacity of information networks," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2345–2364, Jun. 2006.
- [12] K. Jain, V. V. Vazirani, R. W. H. Yeung, and G. Yuval, "On the capacity of multiple unicast sessions in undirected graphs," in *Proc. IEEE ISIT*, Sep. 2005, pp. 563–567.
- [13] M. Langberg and M. Médard, "On the multiple unicast network coding conjecture," in *Proc. 47th Annu. Allerton Conf. Commun., Control, Computing*, 2009, pp. 222–227.
- [14] M. Mitzenmacher. (2007, Aug.). *Network coding: Open problems* [Online]. Available: <http://mybiasedcoin.blogspot.ca/2007/08/network-coding-open-problems.html>
- [15] C. Chekuri, "Routing vs. Network coding," Dept. Comput. Sci., Univ. Illinois at Urbana-Champaign, Champaign, IL, USA, Tech. Rep. 09w5103, 2009.
- [16] N. J. A. Harvey and R. Kleinberg, "Tighter cut-based bounds for k-pairs communication problems," in *Proc. 43rd Annu. Allerton Conf. Commun., Control, Computing*, 2005, pp. 1–10.
- [17] A. Al-Bashabsheh and A. Yongacoglu, "On the k-pairs problem," in *Proc. IEEE ISIT*, Jul. 2008, pp. 1828–1832.
- [18] G. Kramer and S. A. Savari, "Edge-cut bounds on network coding rates," *J. Netw. Syst. Manag.*, vol. 14, no. 1, pp. 49–67, 2006.
- [19] S. M. S. Yazdi, S. A. Savari, F. Farnoud, and G. Kramer, "A multimesage capacity region for undirected ring networks," in *Proc. IEEE ISIT*, Jun. 2007, pp. 1091–1095.
- [20] S. M. S. Yazdi, S. A. Savari, G. Kramer, K. Carlson, and F. Farnoud, "On the multimesage capacity region for undirected ring networks," *IEEE Trans. Inf. Theory*, vol. 56, no. 4, pp. 1930–1947, Apr. 2010.
- [21] M. Iri, "On an extension of the maximum-flow minimum-cut theorem to multicommodity flows," *J. Oper. Res. Soc. Jpn.*, vol. 13, no. 3, pp. 129–135, 1971.
- [22] K. Onaga and O. Kakusho, "On feasibility conditions of multicommodity flows in networks," *IEEE Trans. Circuit Theory*, vol. 18, no. 4, pp. 425–429, Jul. 1971.
- [23] D. Traskov, N. Ratnakar, D. S. Lun, R. Koetter, and M. Médard, "Network coding for multiple unicasts: An approach based on linear optimization," in *Proc. IEEE ISIT*, Jul. 2006, pp. 1758–1762.
- [24] C.-C. Wang and N. B. Shroff, "Pairwise intersession network coding on directed networks," *IEEE Trans. Inf. Theory*, vol. 56, no. 8, pp. 3879–3900, Aug. 2010.
- [25] E. Erez and M. Feder, "Improving the multicommodity flow rates with network coding for two sources," *IEEE J. Sel. Areas Commun.*, vol. 27, no. 5, pp. 814–824, Jun. 2009.
- [26] S. Huang and A. Ramamoorthy, "A note on the multiple unicast capacity of directed acyclic networks," in *Proc. IEEE ICC*, Jun. 2011, pp. 1–6.
- [27] J. Huang, X. Yin, X. Zhang, X. Du, and Z. Li, "On space information flow: Single multicast," in *Proc. IEEE Symp. Netw. Coding*, Jun. 2013, pp. 1–6.
- [28] Z. Li and B. Li, "Network coding in undirected networks," in *Proc. 38th Annu. CISS*, 2004, pp. 1–6.
- [29] A. Agarwal and M. Charikar, "On the advantage of network coding for improving network throughput," in *Proc. IEEE Inf. Theory Workshop*, Oct. 2004, pp. 247–249.
- [30] C. Fragouli and E. Soljanin, *Network Coding Fundamentals*. Norwell, MA, USA: Now Publishers, 2007.
- [31] J. Nash, "The imbedding problem for Riemannian manifolds," *Ann. Math.*, vol. 63, no. 1, pp. 20–63, Jan. 1956.
- [32] C. Chekuri, M. Mydlarz, and F. B. Shepherd, "Multicommodity demand flow in a tree and packing integer programs," *ACM Trans. Algorithms*, vol. 3, no. 3, pp. 5067–5077, Aug. 2007.
- [33] N. Linial, E. London, and Y. Rabinovich, "The geometry of graphs and some of its algorithmic applications," *Combinatorica*, vol. 15, no. 2, pp. 215–245, Feb. 1995.
- [34] S. M. S. T. Yazdi, S. A. Savari, and G. Kramer, "Network coding in star networks," in *Proc. IEEE ISIT*, Jul. 2008, pp. 325–329.

Tang Xiahou obtained her bachelor's degree from University of Science and Technology of China in 2010. After that, she studied at University of Calgary under the supervision of Dr. Zongpeng Li and obtained her master's degree in 2012. Currently she is working as a developer at Microsoft in Beijing.

Chuan Wu (M'08) received her B.Engr. and M.Engr. degrees in 2000 and 2002 from the Department of Computer Science and Technology, Tsinghua University, China, and her Ph.D. degree in 2008 from the Department of Electrical and Computer Engineering, University of Toronto, Canada. Between 2002 and 2004, she worked in the Information Technology industry in Singapore. Since September 2008, Chuan Wu has been an Assistant Professor in the Department of Computer Science at the University of Hong Kong. Her research is in the areas of cloud computing, online and mobile social networks, peer-to-peer networks and wireless networks. She is a member ACM.

Zongpeng Li (M'05–SM'12) received his B.E. degree in Computer Science and Technology from Tsinghua University (Beijing) in 1999, his M.S. degree in Computer Science from University of Toronto in 2001, and his Ph.D. degree in Electrical and Computer Engineering from University of Toronto in 2005. Since 2005, he has been with the Department of Computer Science in the University of Calgary. In 2011-2012, Zongpeng was a visitor at the Institute of Network Coding, Chinese University of Hong Kong. His research interests are in computer networks and network coding.

Jiaqing Huang (M'08) received his B.Sc. degree and M.S. degree in Electronic Science and Technology from East China Normal University (Shanghai) in 1994 and 1997, respectively, and his Ph.D. degree in Electronics and Information Engineering from Huazhong University of Science and Technology (Wuhan) in 2003. He has been working at the Department of Electronics and Information Engineering in the Huazhong University of Science and Technology since July 1997, where he is currently an Associate Professor. His research interests are in communications, particularly network coding.