

Coarse differentiation and multi-flows in planar graphs

James R. Lee*

Prasad Raghavendra†

Abstract

We show that the multi-commodity max-flow/min-cut gap for series-parallel graphs can be as bad as 2, matching a recent upper bound [8] for this class, and resolving one side of a conjecture of Gupta, Newman, Rabinovich, and Sinclair.

This also improves the largest known gap for planar graphs from $\frac{3}{2}$ to 2, yielding the first lower bound that doesn't follow from elementary calculations. Our approach uses the *coarse differentiation* method of Eskin, Fischer, and Whyte in order to lower bound the distortion for embedding a particular family of shortest-path metrics into L_1 .

1 Introduction

Since the appearance of [24] and [4], low-distortion metric embeddings have become an increasingly powerful tool in studying the relationship between cuts and multicommodity flows in graphs. For background on the field of metric embeddings and their applications in theoretical computer science, we refer to Matoušek's book [26, Ch. 15], the surveys [19, 23], and the compendium of open problems [25].

One of the central connections lies in the correspondence between low-distortion L_1 embeddings, on the one hand, and the *Sparsest Cut problem* (see, e.g. [24, 4, 3, 2]) and *concurrent multi-commodity flows* (see, e.g. [17, 12]) on the other. This relationship allows one to bring sophisticated geometric and analytic techniques to bear on classical problems in graph partitioning and in the theory of network flows. In the present paper, we show how techniques developed initially in geometric group theory can be used to shed new light on the connections between sparse cuts and multi-commodity flows in planar graphs.

Multi-commodity flows and sparse cuts. Let $G = (V, E)$ be an undirected graph, with a *capacity* $C(e) \geq 0$ associated to every edge $e \in E$. Assume that we are given k pairs of vertices $(s_1, t_1), \dots, (s_k, t_k) \in V \times V$ and $D_1, \dots, D_k \geq 1$. We think of the s_i as *sources*, the t_i as *targets*, and the value D_i as the *demand* of the *terminal pair* (s_i, t_i) for *commodity* i .

In the *MaxFlow* problem the objective is to maximize the *fraction* λ of the demand that can be shipped simultaneously for all the commodities, subject to the capacity constraints. Denote this maximum by λ^* . A straightforward upper bound on λ^* is the *sparsest cut ratio*. Given any subset $S \subseteq V$, we write

$$\Phi(S) = \frac{\sum_{uv \in E} C(uv) \cdot |\mathbf{1}_S(u) - \mathbf{1}_S(v)|}{\sum_{i=1}^k D_i \cdot |\mathbf{1}_S(s_i) - \mathbf{1}_S(t_i)|},$$

where $\mathbf{1}_S$ is the characteristic function of S . The value $\Phi^* = \min_{S \subseteq V} \Phi(S)$ is the minimum over all cuts (partitions) of V , of the ratio between the total capacity crossing the cut and the total demand

*University of Washington, Seattle. Research partially supported by NSF CAREER award CCF-0644037. Part of this research was completed while the author was a postdoctoral fellow at the Institute for Advanced Study, Princeton.

†University of Washington, Seattle. Research supported in part by NSF grant CCF-0343672.

crossing the cut. In the case of a single commodity (i.e. $k = 1$) the classical MaxFlow-MinCut theorem states that $\lambda^* = \Phi^*$, but in general this is no longer the case. It is known [24, 4] that $\Phi^* = O(\log k)\lambda^*$. This result is perhaps the first striking application of metric embeddings in combinatorial optimization (specifically, it uses Bourgain’s embedding theorem [6]).

Indeed, the connection between L_1 embeddings and multi-commodity flow/cut gaps can be made quite precise. For a graph G , let $c_1(G)$ represent the largest distortion necessary to embed any shortest-path metric on G into L_1 (i.e. the maximum over all possible assignments of non-negative lengths to the edges of G). Then $c_1(G)$ gives an upper bound on the ratio between the sparsest cut ratio and the maximum flow for any multi-commodity flow instance on G (i.e. with any choices of capacities and demands) [24, 4]. Furthermore, this connection is tight in the sense that there is always a multi-commodity flow instance on G that achieves a gap of $c_1(G)$ [17].

Despite significant progress [28, 17, 12, 8, 7], some fundamental questions are still left unanswered. As a prime example, consider the well-known *planar embedding conjecture* [17, 19, 23, 25]:

There exists a constant C such that every planar graph metric embeds into L_1 with distortion at most C .

In initiating a systematic study of L_1 embeddings [17] for minor-closed families, Gupta, Newman, Rabinovich, and Sinclair put forth the following vast generalization of this conjecture (we refer to [14] for the relevant graph theory).

Conjecture 1 (Minor-closed embedding conjecture). *If \mathcal{F} is any non-trivial minor-closed family, then $\sup_{G \in \mathcal{F}} c_1(G) < \infty$.*

Lower bounds on the multi-commodity max-flow/min-cut ratio in planar graphs. While techniques for proving upper bounds on the L_1 -distortion required to embed such families has steadily improved, progress on lower bounds has been significantly slower, and recent breakthroughs in lower bounds for L_1 embeddings of discrete metric spaces that rely on discrete Fourier analysis [21, 20] do not apply to excluded-minor families.

The best previous lower bound on $c_1(G)$ when G is a planar graph occurred for $G = K_{2,n}$, i.e. the complete $2 \times n$ bipartite graph. By a straightforward generalization of the lower bound of Okamura and Seymour [28], it is possible to show that $c_1(K_{2,n}) \rightarrow \frac{3}{2}$ as $n \rightarrow \infty$ (see also [1] for a simple proof of this fact in the dual setting).

We show that, in fact there is an infinite family of series-parallel (and hence, planar) graphs $\{G_n\}$ such that $\lim_{n \rightarrow \infty} c_1(G_n) = 2$; this is not only a new lower bound for planar graphs, but yields an *optimal* lower bound on the L_1 -distortion (and hence the flow/cut gap) for series-parallel graphs. The matching upper bound was recently proved in [8].

1.1 Results and techniques

Our lower bound approach is based on exhibiting *local rigidity* for pieces of metric spaces under low-distortion embeddings into L_1 (which we take to mean $L_1([0, 1])$ throughout). This circle of ideas, and the relationship to theory of metric differentiation are a long-studied phenomena in geometric analysis (see e.g. [18, 29, 5, 9]). More recently, they have been applied to the study of L_1 embeddings [10, 11] based on local rigidity results for sets of finite perimeter in the Heisenberg group [16]; see [22] for the relevance to integrality gaps for the Sparsest Cut problem.

Our basic approach is simple; we know that $c_1(K_{2,n}) \leq \frac{3}{2}$ for every $n \geq 1$. But consider $s, t \in V(K_{2,n})$ which constitute the partition of size 2. Say that a cut $S \subseteq V(K_{2,n})$ is *monotone with respect to s and t* if every simple s - t path in $K_{2,n}$ has at most one edge crossing the cut (S, \bar{S}) . It is not difficult




Figure 1: A single edge H , $H \circledast K_{2,3}$, and $H \circledast K_{2,3} \circledast K_{2,2}$.

to show that if an L_1 embedding is composed entirely of cuts which are monotone with respect to s and t , then that embedding must have distortion at least $2 - \frac{2}{n}$.

Consider now the recursively defined family of graphs $K_{2,n}^{\circledast k}$, where $K_{2,n}^{\circledast 1} = K_{2,n}$ and $K_{2,n}^{\circledast k}$ arises by replacing every edge of $K_{2,n}^{\circledast k-1}$ with a copy of $K_{2,n}$. The family $\{K_{2,2}^{\circledast k}\}_{k \geq 1}$ are the well-known *diamond graphs* of [27, 17]. We show that in any low-distortion embedding of $K_{2,n}^{\circledast k}$ into L_1 , for $k \geq 1$ large enough, it is possible to find a (metric) copy of $K_{2,n}$ for which the induced embedding is composed almost entirely of monotone cuts. The claimed distortion bound follows, i.e. $\lim_{n,k \rightarrow \infty} c_1(K_{2,n}^{\circledast k}) = 2$. In Section 5, we exhibit embeddings which show that for every fixed n , $\lim_{k \rightarrow \infty} c_1(K_{2,n}^{\circledast k}) < 2$, thus it is necessary to have the base graphs grow in size.

The ability to find these monotone copies of $K_{2,n}$ inside a low-distortion L_1 embedding of $K_{2,n}^{\circledast k}$ arises from two sources. The first is the *coarse differentiation* technique of Eskin, Fischer, and Whyte [15] which gives a discrete approach to finding local regularity in distorted paths; this is carried out in Section 3. The second aspect is the relationship between regularity and monotonicity for L_1 embeddings which is expounded upon in Section 3.2, and relies on the well-known fact that every L_1 embedding decomposes in a certain way into a distribution over cuts.

2 Preliminaries

For a graph G , we will use $V(G), E(G)$ to denote the sets of vertices and edges of G , respectively. Sometimes we will equip G with a non-negative length function $\text{len} : E(G) \rightarrow \mathbb{R}_+$, and we let d_{len} denote the shortest-path (semi-)metric on G . We say that len is a *reduced length* if $d_{\text{len}}(u, v) = \text{len}(u, v)$ for every $(u, v) \in E(G)$. All length functions considered in the present paper will be reduced. We will write d_G for the path metric on G if the length function is implicit. For an integer n , let $K_{2,n}$ denote the complete bipartite graph with 2 vertices on one side, and n on the other.

2.1 s - t graphs and \circledast -products

An s - t graph G is a graph which has two distinguished vertices $s, t \in V(G)$. For an s - t graph, we use $s(G)$ and $t(G)$ to denote the vertices labeled s and t , respectively. Throughout this article, the graphs $K_{2,n}$ are considered s - t graphs in the natural way (the two vertices forming one side of the partition are labeled s and t).

Definition 2.1 (Composition of s - t graphs). *Given two s - t graphs H and G , define $H \circledast G$ to be the s - t graph obtained by replacing each edge $(u, v) \in E(H)$ by a copy of G (see Figure 1). Formally,*

- $V(H \circledast G) = V(H) \cup (E(H) \times V(G) \setminus \{s(G), t(G)\})$.

- For every edge $e = (u, v) \in E(H)$, there are $|E(G)|$ edges,

$$\left\{ \left((e, v_1), (e, v_2) \right) \mid (v_1, v_2) \in E(G) \text{ and } v_1, v_2 \notin \{s(G), t(G)\} \right\} \cup \\ \left\{ \left(u, (e, w) \right) \mid (s(G), w) \in E(G) \right\} \cup \left\{ \left((e, w), v \right) \mid (w, t(G)) \in E(G) \right\}$$

- $s(H \circlearrowleft G) = s(H)$ and $t(H \circlearrowleft G) = t(H)$.

If H and G are equipped with length functions $\text{len}_H, \text{len}_G$, respectively, we define $\text{len}_{H \circlearrowleft G}$ as follows. Using the preceding notation, for every edge $e = (u, v) \in E(H)$,

$$\begin{aligned} \text{len}((e, v_1), (e, v_2)) &= \frac{\text{len}_H(e)}{d_{\text{len}_G}(s(G), t(G))} \text{len}_G(v_1, v_2) \\ \text{len}(u, (e, w)) &= \frac{\text{len}_H(e)}{d_{\text{len}_G}(s(G), t(G))} \text{len}_G(s(G), w) \\ \text{len}((e, w), v) &= \frac{\text{len}_H(e)}{d_{\text{len}_G}(s(G), t(G))} \text{len}_G(w, t(G)). \end{aligned}$$

This choice implies that $H \circlearrowleft G$ contains an isometric copy of $(V(H), d_{\text{len}_H})$.

Observe that there is some ambiguity in the definition above, as there are two ways to substitute an edge of H with a copy of G , thus we assume that there exists some arbitrary orientation of the edges of H . However, for our purposes the graph G will be symmetric, and thus the orientations are irrelevant.

Definition 2.2 (Recursive composition). For an s - t graph G and a number $k \in \mathbb{N}$, we define $G^{\circlearrowleft k}$ inductively by letting $G^{\circlearrowleft 0}$ be a single edge of unit length, and setting $G^{\circlearrowleft k} = G^{\circlearrowleft k-1} \circlearrowleft G$.

The following result is straightforward.

Lemma 2.3 (Associativity of \circlearrowleft). For any three graphs A, B, C , we have $(A \circlearrowleft B) \circlearrowleft C = A \circlearrowleft (B \circlearrowleft C)$, both graph-theoretically and as metric spaces.

Definition 2.4. For two graphs G, H , a subset of vertices $X \subseteq V(H)$ is said to be a copy of G if there exists a bijection $f : V(G) \rightarrow X$ with distortion 1, i.e. $d_H(f(u), f(v)) = C \cdot d_G(u, v)$ for some constant $C > 0$.

Now we make the following two simple observations about copies of H and G in $H \circlearrowleft G$.

Observation 2.5. The graph $H \circlearrowleft G$ contains $|E(H)|$ distinguished copies of the graph G , one copy corresponding to each edge in H .

Observation 2.6. The subset of vertices $V(H) \subseteq V(H \circlearrowleft G)$ form an isometric copy of H .

For any graph G , we can write $G^{\circlearrowleft N} = G^{\circlearrowleft k-1} \circlearrowleft G \circlearrowleft G^{N-k}$. By observation 2.5, there are $|E(G^{\circlearrowleft k-1})| = |E(G)|^{k-1}$ copies of G in $G^{\circlearrowleft k-1} \circlearrowleft G$. Now using observation 2.6, we obtain $|E(G)|^{k-1}$ copies of G in $G^{\circlearrowleft N}$. We refer to these as the *level- k copies* of G , and their vertices as *level- k vertices*.

In the case of $K_{2,n}^{\circlearrowleft N}$, we will use a compact notation to refer to the copies of $K_{2,n}$. For two level- k vertices $x, y \in V(K_{2,n}^{\circlearrowleft N})$, we will use $K_{2,n}^{(x,y)}$ to denote the copy of $K_{2,n}$ for which x and y are the s - t points. Note that such a copy does not exist between all pairs of level- k vertices.

2.2 Cuts and L_1 embeddings

Cuts. A cut of a graph is a partition of V into (S, \bar{S}) —we sometimes refer to a subset $S \subseteq V$ as a cut as well. A cut gives rise to a semi-metric; using indicator functions, we can write the cut semi-metric as $\rho_S(x, y) = |\mathbf{1}_S(x) - \mathbf{1}_S(y)|$. A fact central to our proof is that embeddings of finite metric spaces into L_1 are equivalent to sums of positively weighted cut metrics over that set (for a simple proof of this see [13]).

A *cut measure* on G is a function $\mu : 2^V \rightarrow \mathbb{R}_+$ for which $\mu(S) = \mu(\bar{S})$ for every $S \subseteq V$. Every cut measure gives rise to an embedding $f : V \rightarrow L_1$ for which

$$\|f(u) - f(v)\|_1 = \int |\mathbf{1}_S(u) - \mathbf{1}_S(v)| d\mu(S), \quad (1)$$

where the integral is over all cuts (S, \bar{S}) . Conversely, to every embedding $f : V \rightarrow L_1$, we can associate a cut measure μ such that (1) holds. We will use this correspondence freely in what follows. When V is a finite set (as it will be throughout), for $A \subseteq 2^V$, we define $\mu(A) = \sum_{S \in A} \mu(S)$.

Embeddings and distortion. If $(X, d_X), (Y, d_Y)$ are metric spaces, and $f : X \rightarrow Y$, then we write

$$\|f\|_{\text{Lip}} = \sup_{x \neq y \in X} \frac{d_Y(f(x), f(y))}{d_X(x, y)}.$$

If f is injective, then the *distortion* of f is defined by $\text{dist}(f) = \|f\|_{\text{Lip}} \cdot \|f^{-1}\|_{\text{Lip}}$. A map with distortion D will sometimes be referred to as *D-bi-lipschitz*. If $d_Y(f(x), f(y)) \leq d_X(x, y)$ for every $x, y \in X$, we say that f is *non-expansive*. For a metric space X , we use $c_1(X)$ to denote the least distortion required to embed X into L_1 .

3 Coarse differentiation

In the present section, we study the regularity of paths under bi-lipschitz mappings into L_1 . Our main tool is based on differentiation [15]. First, we need a discrete analog of bounded variation.

Definition 3.1. A sequence $\{x_1, x_2, \dots, x_k\} \subseteq X$ in a metric space (X, d) is said to be ϵ -efficient if

$$d(x_1, x_k) \leq \sum_{i=1}^{k-1} d(x_i, x_{i+1}) \leq (1 + \epsilon) d(x_1, x_k)$$

Of course the left inequality follows trivially from the triangle inequality.

Definition 3.2. A function $f : Y \rightarrow X$ between two metric spaces (X, d) and (Y, d') , is said to be ϵ -efficient on $P = \{y_1, y_2, \dots, y_k\} \subseteq Y$ if the sequence $f(P) = \{f(y_1), f(y_2), \dots, f(y_k)\}$ is ϵ -efficient in X .

For the sake of simplicity, we first present the coarse differentiation argument for a function on $[0, 1]$. Let $f : [0, 1] \rightarrow X$ be a non-expansive map into a metric space (X, d) . Let $M \in \mathbb{N}$ be given, and for each $k \in \mathbb{N}$, let $L_k = \{jM^{-k}\}_{j=0}^{M^k} \subseteq [0, 1]$ be the set of *level- k points*, and let $S_k = \{(jM^{-k}, (j+1)M^{-k}) : j \in \{1, \dots, M^k - 1\}\}$ be the set of *level- k pairs*.

For an interval $I = [a, b]$, $f|_I$ denotes the restriction of f to the interval I . Now we say that $f|_I$ is ϵ -efficient at granularity M if

$$\sum_{j=0}^{M-1} d\left(f\left(a + \frac{(b-a)j}{M}\right), f\left(a + \frac{(b-a)(j+1)}{M}\right)\right) \leq (1 + \epsilon) d(f(a), f(b)).$$

Further, we say that a function f is (ε, δ) -inefficient at level k if

$$|\{(a, b) \in S_k : f|_{[a,b]} \text{ is not } \varepsilon\text{-efficient at granularity } M\}| \geq \delta M^k.$$

In other words, the probability that a randomly chosen level k restriction $f|_{[a,b]}$ is not ε -efficient is at least δ . Otherwise, we say that f is (ε, δ) -efficient at level k . The main theorem of this section follows.

Theorem 3.3 (Coarse differentiation). *If a non-expansive map $f : [0, 1] \rightarrow X$ is (ε, δ) -inefficient at an α -fraction of levels $k = 1, 2, \dots, N$, then $\text{dist}(f|_{L_{N+1}}) \geq \frac{1}{2}\varepsilon\alpha\delta N$.*

Proof. Let $D = \text{dist}(f|_{L_{N+1}})$, and let $1 \leq k_1 < \dots < k_h \leq N$ be the $h \geq \lfloor \alpha N \rfloor$ levels at which f is (ε, δ) -inefficient.

Let us consider the first level k_1 . Let $S'_{k_1} \subseteq S_{k_1}$ be a subset of size $|S'_{k_1}| \geq \lfloor \delta |S_{k_1}| \rfloor$ for which

$$(a, b) \in S'_{k_1} \implies f|_{[a,b]} \text{ is not } \varepsilon\text{-efficient at granularity } M$$

For any such $(a, b) \in S'_{k_1}$, we know that

$$\begin{aligned} \sum_{j=0}^{M-1} d\left(f\left(a + jM^{-k_1-1}\right), f\left(a + (j+1)M^{-k_1-1}\right)\right) &> (1 + \varepsilon)d(f(a), f(b)) \\ &\geq d(f(a), f(b)) + \varepsilon \frac{M^{-k_1}}{D}. \end{aligned}$$

by the definition of (not being) ε -efficient, and the fact that $d(f(a), f(b)) \geq |a - b|/D$. For all segments $(a, b) \in S_{k_1} - S'_{k_1}$, the triangle inequality yields

$$\sum_{j=0}^{M-1} d\left(f\left(a + jM^{-k_1-1}\right), f\left(a + (j+1)M^{-k_1-1}\right)\right) \geq d(f(a), f(b))$$

By summing the above inequalities over all the segments in S_{k_1} , we get

$$\sum_{(u,v) \in S_{k_1+1}} d(f(u), f(v)) \geq \sum_{(a,b) \in S_{k_1}} d(f(a), f(b)) + \frac{\varepsilon\delta}{2D},$$

where the extra factor 2 in the denominator on the RHS just comes from removing the floor from $|S'_{k_1}| \geq \lfloor \delta |S_{k_1}| \rfloor$. Similarly, for each of the levels k_2, \dots, k_h , we will pick up an excess term of $\varepsilon\delta/(2D)$. We conclude that

$$1 \geq \sum_{(u,v) \in S_{N+1}} d(f(u), f(v)) \geq \frac{\varepsilon\delta h}{2D},$$

where the LHS comes from the fact that f is non-expansive. Simplifying achieves the desired conclusion. \square

3.1 Differentiation for families of geodesics

Let $G = (V, E)$ be an unweighted graph, and let \mathcal{P} denote a family of geodesics (i.e. shortest-paths) in G . Furthermore, assume that every $\gamma \in \mathcal{P}$ has length M^r for some $M, r \in \mathbb{N}$. Let $f : (V, d_G) \rightarrow X$ be a non-expansive map into some metric space (X, d) .

For the sake of convenience, we will index the vertices along the paths using numbers from $[0, 1]$. Specifically, we will refer to the i^{th} vertex along the path $\gamma \in \mathcal{P}$ by $\gamma\left(\frac{i}{M^r}\right)$. For indices a, b , we will use $\gamma[a, b]$ to denote the sub path starting at $\gamma(a)$ and ending at $\gamma(b)$. We will also use $f|_{\gamma[a, b]}$ to denote the restriction of f to the path $\gamma[a, b]$. As earlier, the function $f|_{\gamma[a, b]}$ is said to be ε -efficient at granularity M if

$$\sum_{j=0}^{M-1} d\left(f\left(\gamma\left(a + M^{-1}(b-a)j\right)\right), f\left(\gamma\left(a + M^{-1}(b-a)(j+1)\right)\right)\right) \leq (1 + \varepsilon) d(f(a), f(b)).$$

Let the sets L_k and S_k be defined as before. Thus a level- k segment of a path $\gamma \in \mathcal{P}$ is $\gamma[a, b]$ for some $(a, b) \in S_k$. We say that f is (ε, δ) inefficient at level k for the family of paths \mathcal{P} if the following holds:

$$\left| \{(a, b) \in S_k, \gamma \in \mathcal{P} : f|_{\gamma[a, b]} \text{ is not } \varepsilon\text{-efficient at granularity } M\} \right| \geq \delta M^k |\mathcal{P}|.$$

A straightforward variation of the proof of Theorem 3.3 yields the following.

Theorem 3.4. *If a non-expansive map $f : V \rightarrow X$ is (ε, δ) -inefficient at an α -fraction of levels $k = 1, 2, \dots, N$, then $\text{dist}(f) \geq \frac{1}{2}\varepsilon\alpha\delta N$.*

Proof. Let $D = \text{dist}(f)$, and let $1 \leq k_1 < \dots < k_h \leq N$ be the $h \geq \lfloor \alpha N \rfloor$ levels for which f is (ε, δ) -inefficient at level k_i .

Let us consider the first level k_1 . Let $S'_{k_1} \subseteq \mathcal{P} \times S_{k_1}$ be a subset of size $|S'_{k_1}| \geq \lfloor \delta |S_{k_1}| |\mathcal{P}| \rfloor$ for which

$$(\gamma, (a, b)) \in S'_{k_1} \implies f|_{\gamma[a, b]} \text{ is not } \varepsilon\text{-efficient at granularity } M.$$

For any such $(\gamma, (a, b)) \in S'_{k_1}$, we know that

$$\begin{aligned} \sum_{j=0}^{M-1} d\left(f\left(\gamma\left(a + jM^{-k_1-1}\right)\right), f\left(\gamma\left(a + (j+1)M^{-k_1-1}\right)\right)\right) &> (1 + \varepsilon) d(f(\gamma(a)), f(\gamma(b))) \\ &\geq d(f(\gamma(a)), f(\gamma(b))) + \varepsilon \frac{M^{r-k_1}}{D}. \end{aligned}$$

by the definition of (not being) ε -efficient, and the fact that $d(f(\gamma(a)), f(\gamma(b))) \geq M^r |a - b| / D$. In particular, summing both sides over all the segments $\gamma[a, b]$ over all paths γ and segments $[a, b] \in S_{k_1}$ (and replacing the preceding inequality by the triangle inequality if $(a, b) \notin S'_{k_1}$), we get

$$\sum_{\gamma \in \mathcal{P}} \sum_{(u, v) \in S_{k_1+1}} d(f(\gamma(u)), f(\gamma(v))) \geq \sum_{\gamma \in \mathcal{P}} \sum_{(a, b) \in S_{k_1}} d(f(\gamma(a)), f(\gamma(b))) + \frac{\varepsilon \delta M^r |\mathcal{P}|}{2D},$$

Similarly, for each of the levels k_2, \dots, k_h , we will pick up an excess term of $\varepsilon \delta M^r |\mathcal{P}| / (2D)$. We conclude that

$$M^r |\mathcal{P}| \geq \sum_{\gamma \in \mathcal{P}} \sum_{(u, v) \in S_{N+1}} d(f(\gamma(u)), f(\gamma(v))) \geq \frac{\varepsilon \delta h M^r |\mathcal{P}|}{2D},$$

The desired conclusion follows. \square

3.2 Efficient L_1 -valued maps and monotone cuts

Finally, we relate monotonicity of L_1 -valued mappings to properties of their cut decompositions.

Definition 3.5. A sequence $P = \{x_1, x_2, \dots, x_k\} \subseteq X$ is said to be monotone with respect to a cut (S, \bar{S}) (where $X = S \uplus \bar{S}$) if $S \cap P = \{x_1, x_2, \dots, x_i\}$ or $\bar{S} \cap P = \{x_1, x_2, \dots, x_i\}$ for some $1 \leq i \leq k$.

If μ is a cut measure on a finite set X and $x, y \in X$, we define the separation measure $\mu^{x|y}$ as follows: For every $S \subseteq X$, let $\mu^{x|y}(S) = \mu(S)|\mathbf{1}_S(x) - \mathbf{1}_S(y)|$.

Lemma 3.6. Let (X, d) be a finite metric space, and let $P = \{x_1, x_2, \dots, x_k\} \subseteq X$ be a finite sequence. Given a mapping $f : X \rightarrow L_1$, let μ be the corresponding cut measure (see (1)). If f is ϵ -efficient on P , then

$$\mu^{x_1|x_k}(\{S : P \text{ is monotone with respect to } (S, \bar{S})\}) \geq (1 - \epsilon)\|f(x_1) - f(x_k)\|_1.$$

Proof. If the sequence P is not monotone with respect to a cut (S, \bar{S}) , then

$$\sum_{i=1}^{k-1} |\mathbf{1}_S(x_i) - \mathbf{1}_S(x_{i+1})| \geq 2|\mathbf{1}_S(x_1) - \mathbf{1}_S(x_k)|.$$

Now, let $\mathcal{E} = \{S : P \text{ is not monotone with respect to } (S, \bar{S})\}$, and for the sake of contradiction, assume that $\mu^{x_1|x_k}(\mathcal{E}) > \epsilon\|f(x_1) - f(x_k)\|_1$, then

$$\begin{aligned} \sum_{i=1}^{k-1} \|f(x_i) - f(x_{i+1})\|_1 &= \sum_{i=1}^{k-1} \left[\int_{\mathcal{E}} |\mathbf{1}_S(x_i) - \mathbf{1}_S(x_{i+1})| d\mu(S) + \int_{\bar{\mathcal{E}}} |\mathbf{1}_S(x_i) - \mathbf{1}_S(x_{i+1})| d\mu(S) \right] \\ &\geq 2 \int_{\mathcal{E}} |\mathbf{1}_S(x_1) - \mathbf{1}_S(x_k)| d\mu(S) + \int_{\bar{\mathcal{E}}} |\mathbf{1}_S(x_1) - \mathbf{1}_S(x_k)| d\mu(S) \\ &= 2\mu^{x_1|x_k}(\mathcal{E}) + \mu^{x_1|x_k}(\bar{\mathcal{E}}) \\ &> (1 + \epsilon)\|f(x_1) - f(x_k)\|_1, \end{aligned}$$

where we observe that $\|f(x_1) - f(x_k)\|_1 = \mu^{x_1|x_k}(\mathcal{E}) + \mu^{x_1|x_k}(\bar{\mathcal{E}})$. This is a contradiction, since f is assumed to be ϵ -efficient on P . \square

4 The distortion lower bound

Our lower bound examples are the recursively defined family of graphs $\{K_{2,n}^{\otimes k}\}_{k=1}^{\infty}$. We recall that the graphs $K_{2,2}^{\otimes k}$ are known as diamond graphs [17, 27].

Lemma 4.1. Let G be an s - t graph with a uniform length function, i.e. $\text{len}(e) = 1$ for every $e \in E(G)$. Then for every $\epsilon, D > 0$, there exists an integer $N = N(G, \epsilon, D)$ such that the following holds: For any non-expansive map $f : G^{\otimes N} \rightarrow X$ with $\text{dist}(f) \leq D$, there exists a copy G' of G in $G^{\otimes N}$ such that f is ϵ -efficient on all s - t geodesics in G' .

Proof. Let $M = d_G(s, t)$, and let \mathcal{P}_G denote the family of s - t geodesics in G . Fix $\delta = \frac{1}{|\mathcal{P}_G|}$, $\alpha = \frac{1}{2}$ and $N = \frac{8D}{\epsilon\delta}$.

Let \mathcal{P} denote the family of all s - t geodesics in $G^{\otimes N}$. Each path in \mathcal{P} is of length M and consists of M^N edges. From the choice of parameters, observe that $\frac{1}{2}\epsilon\alpha\delta N > D$. Applying Theorem 3.4 to the

family \mathcal{P} , any non-expansive map f with $\text{dist}(f) \leq D$ is (ϵ, δ) -efficient at an $\alpha = \frac{1}{2}$ -fraction of levels $k = 1, 2, \dots, N$. Specifically, there exists a level k such that f is (ϵ, δ) -efficient at level k .

For a uniformly random choice of path $\gamma \in \mathcal{P}$, and level- k segment (a, b) of γ , $f|_{\gamma[a, b]}$ is not ϵ -efficient at granularity M with probability at most δ . In case of the family \mathcal{P} , each of the level- k segments is nothing more than an s - t geodesic in a level- k copy of G .

If, for at least one of the level- k copies of G , f is ϵ -efficient on all the s - t geodesics in that copy, the proof is complete. On the contrary, suppose each level- k copy has an s - t geodesic on which f is not ϵ -efficient. Then in each level- k copy at least a $\delta = \frac{1}{\mathcal{P}_G}$ -fraction of the s - t geodesics are ϵ -inefficient. As the level- k copies partition the set of all level- k segments, this implies that at least a δ -fraction of the segments are ϵ -inefficient. This contradicts the fact that f is (ϵ, δ) -efficient at level k . \square

Although we will not need it, the same type of argument proves the following generalization to weighted graphs G . The idea is that in $G^{\otimes N}$ for N large enough, there exists a copy of a subdivision of G with each edge finitely subdivided. Paying small distortion, we can approximate G (up to uniform scaling) by this subdivided copy, where the latter is equipped with uniform edge lengths.

Lemma 4.2. *Let G be an s - t graph with arbitrary non-negative edge lengths $\text{len} : E(G) \rightarrow \mathbb{R}_+$. Then for every $\epsilon, D > 0$, there exists an integer $N = N(G, \epsilon, D, \text{len})$ such that the following holds: For any non-expansive map $f : G^{\otimes N} \rightarrow X$ with $\text{dist}(f) \leq D$, there exists a copy G' of G in $G^{\otimes N}$ such that f is ϵ -efficient on all s - t geodesics in G' .*

In the graph $K_{2,n}$, we will refer to the n vertices other than s, t by $M = \{m_i\}_{i=1}^n$.

Lemma 4.3. *For $\epsilon < \frac{1}{2}$ and any function $f : V(K_{2,n}) \rightarrow L_1$ that is ϵ/n -efficient with respect to each of the geodesics s - m_i - t , for $1 \leq i \leq n$, we have $\text{dist}(f) \geq 2 - \frac{2}{n} - 2\epsilon$.*

Proof. Let μ be the cut measure corresponding to f . By scaling, we may assume that

$$\|f(s) - f(t)\|_1 = \mu\{S : \mathbf{1}_S(s) \neq \mathbf{1}_S(t)\} = 1.$$

Let $V = V(K_{2,n})$. Without loss of generality, we assume that μ is supported on $2^V \setminus \{\emptyset, V\}$. Let γ_i be the geodesic s - m_i - t for $i \in \{1, 2, \dots, n\}$. Define

$$\mathcal{E} = \{S : (S, \bar{S}) \text{ is not monotone with respect to } \gamma_i \text{ for some } i \in [n]\}.$$

Applying Lemma 3.6, by a union bound and the fact that f is ϵ/n efficient on every γ_i , we see that $\mu(\mathcal{E}) \leq \epsilon$.

Consider a cut (S, \bar{S}) that is monotone with respect to all the γ_i geodesics, and such that $\mu(S) > 0$. Let us refer to these cuts as *good* cuts. By monotonicity, and the fact that $S \notin \{\emptyset, V\}$, we know that $|\mathbf{1}_S(s) - \mathbf{1}_S(t)| = 1$. Thus for a good cut (S, \bar{S}) , we have

$$\sum_{i, j \in [n]} |\mathbf{1}_S(m_i) - \mathbf{1}_S(m_j)| = 2(|S| - 1)(n - |S| - 1) \leq \frac{n^2}{2}. \quad (2)$$

It follows that,

$$\begin{aligned} \sum_{i, j \in [n]} \|f(m_i) - f(m_j)\|_1 &= \int_{\mathcal{E}} \sum_{i, j \in [n]} |\mathbf{1}_S(m_i) - \mathbf{1}_S(m_j)| d\mu(S) + \int_{\bar{\mathcal{E}}} \sum_{i, j \in [n]} |\mathbf{1}_S(m_i) - \mathbf{1}_S(m_j)| d\mu(S) \\ &\leq \mu(\bar{\mathcal{E}}) \frac{n^2}{2} + \mu(\mathcal{E}) n^2 \\ &\leq (1 - \epsilon) \frac{n^2}{2} + \epsilon n^2 \\ &\leq \frac{(1 + \epsilon)n^2}{2} \|f(s) - f(t)\|_1. \end{aligned}$$

where in the first inequality, we have used (2), and we recall that $\|f(s) - f(t)\|_1 = 1$.

Contrasting this with the fact that

$$\sum_{i,j \in [n]} d_{K_{2,n}}(m_i, m_j) = n(n-1) d_{K_{2,n}}(s, t)$$

yields

$$\text{dist}(f) \geq \frac{n(n-1)}{\frac{(1+\epsilon)n^2}{2}} = \frac{2}{1+\epsilon} \left(1 - \frac{1}{n}\right) \geq 2 - \frac{2}{n} - 2\epsilon.$$

□

Theorem 4.4. *For any $n \geq 2$, $\lim_{k \rightarrow \infty} c_1(K_{2,n}^{\otimes k}) \geq 2 - \frac{2}{n}$.*

Proof. For any $\epsilon' > 0$, let N be the integer obtained by applying Lemma 4.1 to $K_{2,n}$ with $\epsilon = \epsilon'/n$, $D = 2$ and $G = K_{2,n}$. We will show that for any map $f : K_{2,n}^{\otimes N} \rightarrow L_1$, $\text{dist}(f) \geq 2 - \frac{2}{n} - 2\epsilon'$. Without loss of generality, assume that f is non-expansive. If $\text{dist}(f) \leq 2$, then from Lemma 4.1 there exists a copy of $K_{2,n}$ in which f is $\frac{\epsilon'}{n}$ on all the s - t geodesics. Using Lemma 4.3, we see that on *this* copy of $K_{2,n}$ we get $\text{dist}(f|_{K_{2,n}}) \geq 2 - \frac{2}{n} - 2\epsilon'$. The result follows by taking $\epsilon' \rightarrow 0$. □

5 Embeddings of $K_{2,n}^{\otimes k}$

In this section, we show that for every fixed n , $\lim_{k \rightarrow \infty} c_1(K_{2,n}^{\otimes k}) < 2$.

A next-embedding operator. Let T be a random variable ranging over subsets of $V(K_{2,n}^{\otimes k})$, and let S be a random variable ranging over subsets of $V(K_{2,n})$. We define a random subset $P_S(T) \subseteq V(K_{2,n}^{\otimes k+1})$ as follows. One moves from $K_{2,n}^{\otimes k}$ to $K_{2,n}^{\otimes k+1}$ by replacing every edge $(x, y) \in E(K_{2,n}^{\otimes k})$ with a copy of $K_{2,n}$ which we will call $K_{2,n}^{(x,y)}$. For every edge $(x, y) \in K_{2,n}^{\otimes k}$, let $S^{(x,y)}$ be an independent copy of the cut S (which ranges over subsets of $V(K_{2,n})$). We form the cut $P_S(T) \subseteq V(K_{2,n}^{\otimes k+1})$ as follows. If $(x, y) \in E(K_{2,n}^{\otimes k})$, then for $v \in V(K_{2,n}^{(x,y)})$, we put

$$\mathbf{1}_{P_S(T)}(v) = \begin{cases} \mathbf{1}_{P_S(T)}\left(s(K_{2,n}^{(x,y)})\right) & \text{if } \mathbf{1}_{S^{(x,y)}}(v) = \mathbf{1}_{S^{(x,y)}}\left(s(K_{2,n}^{(x,y)})\right) \\ \mathbf{1}_{P_S(T)}\left(t(K_{2,n}^{(x,y)})\right) & \text{otherwise} \end{cases}$$

We note that, strictly speaking, the operator P_S depends on n and k , but we allow these to be implicit parameters.

5.1 Embeddings for small n

Consider the graph $K_{2,n}$ with vertex set $V = \{s, t\} \cup M$. An embedding in the style of [17] would define a random subset $S \subseteq V$ by selecting $M' \subseteq M$ to contain each vertex from M independently with probability $\frac{1}{2}$, and then setting $S = \{s\} \cup M'$. The resulting embedding has distortion 2 since, for every pair $x, y \in M$, we have $\Pr[\mathbf{1}_S(x) \neq \mathbf{1}_S(y)] = \frac{1}{2}$. To do slightly better, we choose a uniformly random subset $M' \subseteq M$ of size $\lfloor \frac{n}{2} \rfloor$ and set $S = \{s\} \cup M'$ or $S = \{s\} \cup (M \setminus M')$ each with probability half. In this case, we have

$$\Pr[\mathbf{1}_S(x) \neq \mathbf{1}_S(y)] = \frac{\lfloor \frac{n}{2} \rfloor \cdot \lfloor \frac{n+1}{2} \rfloor}{\binom{n}{2}} > \frac{1}{2},$$




Figure 2: The two cases of Theorem 5.1

resulting in a distortion slightly better than 2. A recursive application of these ideas results in $\lim_{k \rightarrow \infty} c_1(K_{2,n}^{\otimes k}) < 2$ for every $n \geq 1$, though the calculation is complicated by the fact that the worst distortion is incurred for a pair $\{x, y\}$ with $x \in M(H)$ and $y \in M(G)$ where H is a copy of $K_{2,n}^{\otimes k_1}$ and G is a copy of $K_{2,n}^{\otimes k_2}$, and the relationship between k_1 and k_2 depends on n . (For instance, $c_1(K_{2,2}) = 1$ while $\lim_{k \rightarrow \infty} c_1(K_{2,2}^{\otimes k}) = \frac{4}{3}$.)

Theorem 5.1. *For any $n, k \in \mathbb{N}$, we have $c_1(K_{2,n}^{\otimes k}) \leq 2 - \frac{2}{2^{\lfloor \frac{k}{2} \rfloor + 1}}$.*

Proof. For simplicity, we prove the bound for $K_{2,2n}^{\otimes k}$. A similar analysis holds for $K_{2,2n+1}^{\otimes k}$. We define a random cut $S_k \subseteq V(K_{2,2n}^{\otimes k})$ inductively. For $k = 1$, choose a uniformly random partition $M(K_{2,2n}^{\otimes 1}) = M_s \cup M_t$ with $|M_s| = |M_t| = n$, and let $S_1 = \{s(K_{2,2n}^{\otimes 1})\} \cup \{M_s\}$. The key fact which causes the distortion to be less than 2 is the following: For any $x, y \in M(K_{2,2n}^{\otimes 1})$, we have

$$\Pr[\mathbf{1}_{S_1}(x) \neq \mathbf{1}_{S_1}(y)] = \frac{n^2}{\binom{2n}{2}} = \frac{n}{2n-1} > \frac{1}{2}. \quad (3)$$

This follows because there are $\binom{2n}{2}$ pairs $\{x, y\} \in M(K_{2,2n}^{\otimes 1})$ and n^2 are separated by S_1 .

Assume now that we have a random subset $S_k \subseteq V(K_{2,2n}^{\otimes k})$. We set $S_{k+1} = P_{S_1}(S_k)$ where P_{S_1} is the operator defined above, which maps random subsets of $V(K_{2,2n}^{\otimes k})$ to random subsets of $V(K_{2,2n}^{\otimes k+1})$. In other words $S_k = P_{S_1}^{k-1}(S_1)$.

Let $s_0 = s(K_{2,2n}^{\otimes k})$ and $t_0 = t(K_{2,2n}^{\otimes k})$. It is easy to see that the cut $S = S_k$ defined above is always monotone with respect to every s_0 - t_0 shortest path in $K_{2,2n}^{\otimes k}$, thus every such path has exactly one edge cut by S_k , and furthermore the cut edge is uniformly chosen from along the path, i.e. $\Pr[\mathbf{1}_S(x) \neq \mathbf{1}_S(y)] = 2^{-k}$ for every $(x, y) \in E(K_{2,2n}^{\otimes k})$. In particular, it follows that if $u, v \in V(K_{2,2n}^{\otimes k})$ lie along the same simple s_0 - t_0 path, then $\Pr[\mathbf{1}_S(u) \neq \mathbf{1}_S(v)] = 2^{-k}d(u, v)$.

Now consider any $u, v \in V(K_{2,2n}^{\otimes k})$. Fix some shortest path P from u to v . By symmetry, we may assume that P goes left (toward s_0) and then right (toward t_0). Let s be the left-most point of P . In this case, $s = s(H)$ for some subgraph H which is a copy of $K_{2,2n}^{\otimes k'}$ with $k' \leq k$, and such that $u, v \in V(H)$; we let $t = t(H)$. We also have $d(u, v) = d(u, s) + d(s, v)$. Let $M = M(H)$, and fix $x, y \in M$ which lie along the s - u - t and s - v - t shortest-paths, respectively. Without loss of generality, we may assume that $d(s, v) \leq d(s, y)$. We need to consider two cases (see Figure 5.1).

Case I: $d(u, s) \leq d(x, s)$.

For any pair $a, b \in V(K_{2,2n}^{\otimes k})$, we let $\mathcal{E}_{a,b}$ be the event $\{\mathbf{1}_S(a) \neq \mathbf{1}_S(b)\}$. In this case, we have $\Pr[\mathcal{E}_{u,v}] = \Pr[\mathcal{E}_{s,t}] \cdot \Pr[\mathcal{E}_{u,v} | \mathcal{E}_{s,t}]$. Since s, t clearly lie on a shortest s_0 - t_0 path, we have $\Pr[\mathcal{E}_{s,t}] =$

$2^{-k}d(s, t)$. For any event \mathcal{E} , we let $\mu[\mathcal{E}] = \Pr[\mathcal{E} \mid \mathcal{E}_{s,t}]$. Now we calculate using (3),

$$\begin{aligned} \mu[\mathcal{E}_{u,v}] &\geq \mu[\mathcal{E}_{x,y}] (\mu[\mathcal{E}_{x,s} \mid \mathcal{E}_{x,y}] \mu[\mathcal{E}_{u,s} \mid \mathcal{E}_{x,s}, \mathcal{E}_{x,y}] + \mu[\mathcal{E}_{x,t} \mid \mathcal{E}_{x,y}] \mu[\mathcal{E}_{v,s} \mid \mathcal{E}_{x,t}, \mathcal{E}_{x,y}]) \\ &= \frac{n}{2n-1} \left(\frac{1}{2} \cdot \frac{d(u, s)}{d(x, s)} + \frac{1}{2} \cdot \frac{d(v, s)}{d(y, s)} \right) \\ &= \frac{n}{2n-1} \frac{d(u, v)}{d(s, t)}. \end{aligned}$$

Hence in this case, $\Pr[\mathbf{1}_S(u) \neq \mathbf{1}_S(v)] \geq \frac{n}{2n-1} \cdot 2^{-k} d(u, v)$.

Case II: $d(u, s) \geq d(x, s)$.

Here, we need to be more careful about bounding $\mu[\mathcal{E}_{u,v}]$. It will be helpful to introduce the notation $a \mapsto b$ to represent the event $\{\mathbf{1}_S(a) = \mathbf{1}_S(b)\}$. We have,

$$\begin{aligned} \mu[\mathcal{E}_{u,v}] &= \mu[x \mapsto t, y \mapsto s] + \mu[x \mapsto t, y \mapsto t, v \mapsto s] + \mu[x \mapsto s, y \mapsto s, u \mapsto t] \\ &\quad + \mu[x \mapsto s, y \mapsto t, u \mapsto t, v \mapsto s] + \mu[x \mapsto s, y \mapsto t, u \mapsto s, v \mapsto t] \\ &= \frac{1}{2} \frac{n}{2n-1} + \frac{n-1}{2n-1} \frac{d(v, y) + d(u, x)}{d(s, t)} + \frac{1}{2} \frac{n}{2n-1} \left(\frac{d(u, x)d(v, y) + d(u, t)d(v, s)}{d(x, t)d(y, s)} \right) \end{aligned}$$

If we set $A = \frac{d(v, s)}{d(s, t)}$ and $B = \frac{d(u, x)}{d(s, t)}$, then $\frac{d(u, v)}{d(s, t)} = \frac{1}{2} + A + B$ and simplifying the expression above, we have

$$\mu[\mathcal{E}_{u,v}] = \frac{1}{2} + B + \frac{A}{2n-1} - \frac{4n}{2n-1} AB$$

Since the shortest path from u to v goes through s by assumption, we must have $A + B \leq \frac{1}{2}$. Thus we are interested in the minimum of $\mu[\mathcal{E}_{u,v}] / (\frac{1}{2} + A + B)$ subject to the constraint $A + B \leq \frac{1}{2}$. It is easy to see that the minimum is achieved at $A + B = \frac{1}{2}$, thus setting $B = \frac{1}{2} - A$, we are left to find

$$\min_{0 \leq A \leq \frac{1}{2}} \left\{ 1 - 2A + \frac{4nA^2}{2n-1} \right\} = \frac{2n+1}{4n}.$$

(The minimum occurs at $A = \frac{1}{2} - \frac{1}{4n}$.) So in this case, $\Pr[\mathbf{1}_S(u) \neq \mathbf{1}_S(v)] \geq \frac{2n+1}{4n} 2^{-k} d(u, v)$.

Combining the above two cases, we conclude that the distribution $S = S_k$ induces an L_1 embedding of $K_{2,2n}^{\otimes k}$ with distortion at most $\max\{\frac{2n-1}{n}, \frac{4n}{2n+1}\} = 2 - \frac{2}{2n+1}$. A similar calculation yields

$$c_1(K_{2,2n+1}^{\otimes k}) \leq \left(\min_{0 \leq A \leq \frac{1}{2}} \left\{ 1 - 2A + \frac{4(n+1)A^2}{2n+1} \right\} \right)^{-1} = 2 - \frac{2}{2n+3}.$$

□

Acknowledgments

We thank Alex Eskin for explaining coarse differentiation to us, Yuri Rabinovich for relating the distortion 2 conjecture, and Kostya and Yury Makayrchev for a number of fruitful discussions.

References

- [1] A. Andoni, M. Deza, A. Gupta, P. Indyk, and S. Raskhodnikova. Lower bounds for embedding of edit distance into normed spaces. In *Proceedings of the 14th annual ACM-SIAM Symposium on Discrete Algorithms*, 2003.
- [2] S. Arora, J. R. Lee, and A. Naor. Euclidean distortion and the Sparsest Cut. *J. Amer. Math. Soc.*, 21(1):1–21, 2008.
- [3] S. Arora, S. Rao, and U. Vazirani. Expander flows, geometric embeddings, and graph partitionings. In *36th Annual Symposium on the Theory of Computing*, pages 222–231, 2004. To appear, *J. ACM*.
- [4] Y. Aumann and Y. Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM J. Comput.*, 27(1):291–301 (electronic), 1998.
- [5] Y. Benyamini and J. Lindenstrauss. *Geometric nonlinear functional analysis. Vol. 1*, volume 48 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2000.
- [6] J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel J. Math.*, 52(1-2):46–52, 1985.
- [7] B. Brinkman, A. Karagiozova, and J. R. Lee. Vertex cuts, random walks, and dimension reduction in series-parallel graphs. In *39th Annual Symposium on the Theory of Computing*, pages 621–630, 2007.
- [8] A. Chakrabarti, A. Jaffe, J. R. Lee, and J. Vincent. Embeddings, flows, and cuts in 2-sums of graphs. Submitted, 2008.
- [9] J. Cheeger. Differentiability of Lipschitz functions on metric measure spaces. *Geom. Funct. Anal.*, 9(3):428–517, 1999.
- [10] J. Cheeger and B. Kleiner. Differentiating maps into L^1 and the geometry of BV functions. Preprint, 2006.
- [11] J. Cheeger and B. Kleiner. Generalized differential and bi-Lipschitz nonembedding in L^1 . *C. R. Math. Acad. Sci. Paris*, 343(5):297–301, 2006.
- [12] C. Chekuri, A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair. Embedding k -outerplanar graphs into l_1 . *SIAM J. Discrete Math.*, 20(1):119–136, 2006.
- [13] M. M. Deza and M. Laurent. *Geometry of cuts and metrics*, volume 15 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1997.
- [14] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, third edition, 2005.
- [15] A. Eskin, D. Fisher, and K. Whyte. Quasi-isometries and rigidity of solvable groups. Preprint, 2006.
- [16] B. Franchi, R. Serapioni, and F. Serra Cassano. Rectifiability and perimeter in the Heisenberg group. *Math. Ann.*, 321(3):479–531, 2001.

- [17] A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair. Cuts, trees and l_1 -embeddings of graphs. *Combinatorica*, 24(2):233–269, 2004.
- [18] J. Heinonen. *Lectures on analysis on metric spaces*. Universitext. Springer-Verlag, New York, 2001.
- [19] P. Indyk. Algorithmic applications of low-distortion geometric embeddings. In *42nd Annual Symposium on Foundations of Computer Science*, pages 10–33. IEEE Computer Society, 2001.
- [20] S. Khot and A. Naor. Nonembeddability theorems via Fourier analysis. *Math. Ann.*, 334(4):821–852, 2006.
- [21] S. Khot and N. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1 . In *46th Annual Symposium on Foundations of Computer Science*, pages 53–62. IEEE Computer Soc., Los Alamitos, CA, 2005.
- [22] J. R. Lee and A. Naor. l_p metrics on the Heisenberg group and the Goemans-Linial conjecture. In *47th Annual Symposium on Foundations of Computer Science*. IEEE Computer Soc., Los Alamitos, CA, 2006.
- [23] N. Linial. Finite metric-spaces—combinatorics, geometry and algorithms. In *Proceedings of the International Congress of Mathematicians, Vol. III (Beijing, 2002)*, pages 573–586, Beijing, 2002. Higher Ed. Press.
- [24] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- [25] J. Matoušek. Open problems on low-distortion embeddings of finite metric spaces. Online: <http://kam.mff.cuni.cz/~matousek/metrop.ps>.
- [26] J. Matoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.
- [27] I. Newman and Y. Rabinovich. A lower bound on the distortion of embedding planar metrics into Euclidean space. *Discrete Comput. Geom.*, 29(1):77–81, 2003.
- [28] H. Okamura and P. D. Seymour. Multicommodity flows in planar graphs. *J. Combin. Theory Ser. B*, 31(1):75–81, 1981.
- [29] P. Pansu. Métriques de Carnot-Carathéodory et quasiisométries des espaces symétriques de rang un. *Ann. of Math. (2)*, 129(1):1–60, 1989.