

Embeddings of topological graphs: Lossy invariants, linearization, and 2-sums

Amit Chakrabarti*

Alexander Jaffe[†]

James R. Lee^{‡†}

Justin Vincent[†]

Abstract

We study the properties of embeddings, multicommodity flows, and sparse cuts in minor-closed families of graphs which are also closed under 2-sums; this includes planar graphs, graphs of bounded treewidth, and constructions based on recursive edge replacement. In particular, we show the following.

- Every graph which excludes K_4 as a minor (in particular, series-parallel graphs) admits an embedding into L_1 with distortion at most 2, confirming a conjecture of Gupta, Newman, Rabinovich, and Sinclair, and improving over their upper bound of 14. This shows that in every multi-commodity flow instance on such a graph, one can route a maximum concurrent flow whose value is at least half the cut bound. Our upper bound is optimal, as it matches a recent lower bound of Lee and Raghavendra.
- We move beyond K_4 -minor-free graphs by showing that every W_4 -minor-free-graph embeds into L_1 with $O(1)$ distortion, where W_4 is the 4-wheel. By a characterization of Seymour, these graphs are precisely subgraphs of 2-sums of K_4 's.
- We prove that if \mathcal{G} and \mathcal{H} are two minor-closed families and \mathcal{G} is closed under taking 2-sums, then members of \mathcal{G} embed non-trivially into non-contracting distributions over members of \mathcal{H} if and only if $\mathcal{G} \subseteq \mathcal{H}$. This significantly generalizes a result of Gupta, et al. where \mathcal{G} and \mathcal{H} are the families of K_4 -minor-free graphs and trees, respectively.

It implies, for instance, that W_4 -free graphs do not embed into distributions over K_4 -free graphs, and that treewidth- $(k+1)$ planar graphs cannot be non-trivially embedded into distributions over arbitrary graphs of treewidth k , improving a result of Carroll and Goel which involves graphs of treewidth $k+3$ and k , respectively. This contrasts rather vividly with the results of Indyk and Sidiropoulos which show that, for fixed $g \geq 1$, genus- g graphs do embed into distributions over planar graphs.

1 Introduction

Since the appearance of [19] and [2], low-distortion metric embeddings have become an increasingly powerful tool in the study of multi-commodity flows and sparse cuts in graphs. For background on the field of metric embeddings and their applications, we refer to Matoušek's book [21, Ch. 15], the surveys [11, 18], and the compendium of open problems [20]. One of the most intriguing and well-studied lines of embedding research concerns the relationship between the topology of a graph and the embeddability of its shortest-path metric into L_1 . As a prime example, consider following the well-known conjecture.

Conjecture 1 (Planar embedding conjecture). *There exists a constant C such that every planar graph metric embeds into L_1 with distortion at most C .*

This conjecture was first posed in published form by Gupta, Newman, Rabinovich, and Sinclair [10], but has been well-known since the seminal paper of Linial, London, and Rabinovich [19]. The conjecture has been emphasized in treatments by Indyk [11], Linial [18], and Matoušek [21].

The driving force behind such questions lies in the intimate relationship between low-distortion L_1 embeddings, on the one hand, and *sparse cuts* and *concurrent multi-commodity flows*, on the other. More specifically,

*Dartmouth College, Department of Computer Science, Hanover, NH. Part of this work completed while the author was visiting the University of Washington.

[†]University of Washington, Department of Computer Science and Engineering, Seattle, WA. Research supported in part by NSF CCF-0644037.

[‡]Part of this work done while the author was at the Institute for Pure and Applied Mathematics, UCLA.

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 and the maximum concurrent flow for any multi-commodity flow instance on G (i.e. with any choices of capacities and demands) [19, 2]. Furthermore, this connection is tight in the sense that there is always a multi-commodity flow instance on G that achieves a gap of $c_1(G)$ [10]. One derives from these facts a $c_1(G)$ -approximation to the sparsest cut problem on G .

The present paper continues the study of L_1 embeddings of topologically restricted families of graphs by introducing a number of new techniques, for both upper and lower bounds, that shed light on the corresponding metric and cut structures. In some instances, we are able to produce optimal upper bounds on distortion, and hence optimal bounds on the multi-commodity flow/cut gap in well-studied families of graphs. We remark that 2-summing is a crucial operation in the setting of this study; planar graphs are obviously closed under 2-sums, and the best-known lower bounds [16] on the multi-commodity max-flow/min-cut gap in planar graphs are formed by recursive 2-summing of graphs.

Topologically restricted families of graphs. We will be concerned with families of graphs \mathcal{F} which are closed under taking minors, i.e. such that if $G \in \mathcal{F}$, and G' is a graph which results from G by removing an edge or vertex, or by contracting one of its edges, then $G' \in \mathcal{F}$ as well; we refer to [6] for the relevant graph theory. Examples of such families include the set of all planar graphs, or the set of all graphs which can be drawn without edge crossings on a surface of genus g , for any fixed $g \in \mathbb{N}$.

The seminal work of Robertson and Seymour [26] shows that every minor-closed family \mathcal{F} is actually characterized by a finite list of forbidden minors. We say that a family of graphs is *non-trivial* if \mathcal{F} does not contain all graphs, and we write $c_1(\mathcal{F}) = \sup_{G \in \mathcal{F}} c_1(G)$. In initiating a systematic study of L_1 embeddings [10] for minor-closed families, Gupta, Newman, Rabinovich, and Sinclair put forth the following vast generalization of the planar embedding conjecture.

Conjecture 2 (Minor-closed embedding conjecture). *If \mathcal{F} is any non-trivial minor-closed family, then $c_1(\mathcal{F}) < \infty$.*

So far, progress on the preceding conjecture has been limited. If L is a finite list of graphs, we write $\mathcal{F}L$ for the (minor-closed) family of graphs which do not contain

any member of L as a minor. From Kuratowski's theorem, we know that the planar embedding conjecture is equivalent to the assertion that $c_1(\mathcal{F}\{K_{3,3}, K_5\}) < \infty$.

The classical theory. The work of Okamura and Seymour [23] implies that $c_1(\mathcal{F}\{K_{2,3}\}) = 1$, where we will use the standard notation $K_{m,n}$ and K_m for the complete $m \times n$ bipartite graph and the complete graph on m vertices, respectively. We recall that $\mathcal{F}\{K_{2,3}\}$ is precisely the family of graphs whose 2-connected components are either outerplanar or a copy of K_4 .

In fact, much of the classical theory of multi-commodity flows and cuts, when translated to the language of L_1 embeddings, concerns non-expansive embeddings of graph metrics which are isometric when restricted to various subsets of the vertices. The previously mentioned theorem of Okamura and Seymour actually follows from a stronger result: Let (G, d) be the shortest-path metric on a weighted planar graph G , and let $F \subseteq V(G)$ be any face of G . Then there exists a non-expansive mapping $f : V(G) \rightarrow L_1$ such that $f|_F$ is an isometry.

This was further strengthened by Okamura [22]: If $F' \subseteq V(G)$ is another face of G , then in fact one can construct a non-expansive map such that both $f|_F$ and $f|_{F'}$ are isometries. Another result of Okamura states that if we fix a node $v \in F$, then there exists a non-expansive map f where $f|_F$ is an isometry and $\|f(v) - f(u)\|_1 = d_G(u, v)$ for every $u \in V(G)$. In fact, stronger results are known, but these require the language of edge-disjoint paths [23, 22] or cut packings [27].

In [23], Okamura and Seymour also showed that $c_1(K_{2,3}) \geq 4/3$, which exhibits the need to move away from isometric embeddings in order to study L_1 embeddings (and, equivalently, multi-commodity flows) for richer families of graphs.

Embeddings with distortion. In [10], Gupta, et al. made the first step in this direction by showing that $c_1(\mathcal{F}\{K_4\}) < 14$, and a well-known generalization of [23] shows that $c_1(K_{2,n}) \rightarrow \frac{3}{2}$ as $n \rightarrow \infty$, hence $c_1(\mathcal{F}\{K_4\}) \geq \frac{3}{2}$ as well (see [1] for a proof that the unweighted metric on $K_{2,n}$ yields this bound). We recall that $\mathcal{F}\{K_4\}$ is precisely the family of graphs whose 2-connected components are series-parallel. One of the contributions of this paper is a proof that $c_1(\mathcal{F}\{K_4\}) \leq 2$, matching a lower bound of Lee and Raghavendra [16], and yielding the first exact bound for a minor-closed family of graphs whose distortion is not exactly 1. Independent of the present work, we have learned that Evans and Safari [7] achieve a tighter analysis of the Gupta, et al. [10] embedding, showing that the actual distortion

incurred is at most 6. Furthermore, they showed that the embedding of [10] incurs distortion at least 3.

We mention a few additional results in this realm. In [4], it is shown that if \mathcal{O}_k is the family of k -outerplanar graphs, then $c_1(\mathcal{O}_k) \leq c^k$ for some constant $c \geq 1$. Rao [25], based on the fundamental work [13], shows that for any $r \geq 1$, any n -point graph $G \in \mathcal{F}\{K_{r,r}\}$ satisfies $c_1(G) = O(r^3 \sqrt{\log n})$. We remark that this bound has been improved to $c_1(G) = O(r \sqrt{\log n})$ by the papers [9, 14]. Indyk and Sidiropoulos [12] showed that every genus g graph embeds into a distribution over planar graphs (see Section 1.2 for a definition) with distortion $\exp(O(g))$. In particular, it follows that $c_1(\{\text{genus } g \text{ graphs}\}) \leq \exp(O(g)) \cdot c_1(\mathcal{F}\{K_{3,3}, K_5\})$.

Finally, we mention a result of Lee and Naor [15] which solves the Lipschitz extension problem for shortest-path metrics on topologically restricted graph metrics: For any finite family of graphs L , there exists a constant c_L such that if (G, d) is a shortest-path metric with $G \in \mathcal{F}L$, $S \subseteq V(G)$, and $f : S \rightarrow L_1$, then there exists an extension $\tilde{f} : V(G) \rightarrow L_1$ with $\tilde{f}|_S = f$ and $\|\tilde{f}\|_{\text{Lip}} \leq c_L \|f\|_{\text{Lip}}$.

1.1 Results and techniques

In what follows, we assume a general knowledge of concurrent multi-commodity flows and the sparsest cut problem; we refer the reader to [19, 10] for the necessary background. For notions related to cut distributions on graphs and embeddings of discrete metric spaces, see Section 1.2.

Optimal embeddings for K_4 -free graphs. In Section 2, we prove that $c_1(\mathcal{F}\{K_4\}) = 2$, resolving a conjecture of Gupta, Newman, Rabinovich, and Sinclair (communicated to us by Y. Rabinovich), and improving over their previous bound of 14 [10]. This provides the first optimal upper bound on the L_1 -distortion of an excluded-minor family that does not embed into L_1 isometrically, and hence the first optimal *approximate* multicommodity max-flow/min-cut theorem for a natural family of graphs (as opposed to the previous *exact* max-flow/min-cut theorems). In addition, this yields a 2-approximation for the Sparsest Cut problem in K_4 -free graphs.

When one considers an isometric L_1 embedding of a graph $G = (V, E)$, each cut supported by the embedding has a very strong form: If $S \subseteq V$ is such a cut, and $G/\partial S$ represents the graph arising from contracting the edges across S , then for every $u, v \in V$, we must have

$$d_G(u, v) - d_{G/\partial S}(u, v) = |\mathbf{1}_S(u) - \mathbf{1}_S(v)|, \quad (1)$$

where $\mathbf{1}_S$ is the characteristic function of S . This is a *lossless* property, in the sense that one can now pass to an embedding of $G/\partial S$ inductively; an isometry of $G/\partial S$ extends to an isometry of G .

Now, we recall that K_4 -free graphs are precisely those whose 2-connected components are series-parallel. Unfortunately, it is not difficult to see that the shortest-path metric of $K_{2,3}$ (a series-parallel graph) does not embed isometrically into L_1 ; in fact, $c_1(K_{2,3}) = 4/3$. There is no non-trivial cut in $K_{2,3}$ satisfying (1), and thus an inductive embedding procedure must necessarily be more delicate.

The key to achieving the optimal upper bound $c_1(\mathcal{F}\{K_4\}) \leq 2$ lies in designing the proper *lossy* inductive invariant, which obviously can no longer involve simply a bound on the distortion. After defining the proper invariant, the embedding proceeds by induction via two types of *local moves*. In one type of move, we make a cut that satisfies (1). The existence of such cuts depends on the fact that our graphs are unweighted and *bipartite* (we first perform a general reduction to this case). Similar considerations arise in the classical theory (see, e.g. [27]), and in the dual setting this corresponds to requiring the graph to be Eulerian as in [23]. The second type of move involves removing an ear from the graph (an ear is a subpath all of whose internal vertices have degree 2), inducting on the remainder, and then randomly extending the embedding to the removed ear.

The inductive invariant is somewhat complicated (see Section 2.2), but one key property is suggested by the paper [16] which contains the lower bound $c_1(\mathcal{F}\{K_4\}) \geq 2$. The lower bound technique in that paper is based on a differentiation argument which shows that “locally,” the cuts occurring in any low-distortion embedding of a path must have a certain rigid form. Reflecting this property, we ensure that every geodesic ear is embedded isometrically. This is possible for series-parallel graphs because, almost by definition, there is a “global direction” along which all geodesic ears travel.

Lower bounds on embeddings into distributions over simpler graphs. A sometimes useful technique for constructing L_1 embeddings for graphs is to first embed them into a distribution over simpler graphs with small expected distortion. (More specifically, we refer to distributions over non-contractive embeddings.) This was employed by Gupta, et al. [10] for outerplanar graphs, and by Chekuri, et al. [4] for k -outerplanar graphs.

On the other hand, [10] shows that series-parallel graphs cannot be embedded into distributions over trees, and [3] shows that treewidth- $(k + 3)$ graphs cannot be

embedded into distributions over treewidth- k graphs. This still leaves open some interesting possibilities, e.g. can *planar* treewidth- $(k + 1)$ graphs be embedded into treewidth- k graphs? Can W_4 -free graphs be embedded into distributions over series-parallel graphs?

In Section 3, we show that the answer to all these questions is negative, by the following general lower bound. Let \mathcal{G} and \mathcal{H} be two minor-closed families of graphs, and suppose that \mathcal{G} contains some non-tree graph, and is closed under the operation of taking 2-sums. If n -vertex graphs from \mathcal{G} can be embedded into distributions over graphs in \mathcal{H} with $o(\log n)$ distortion, then $\mathcal{G} \subseteq \mathcal{H}$. The preceding questions are all answered by this result. This contrasts rather vividly with the results of Indyk and Sidiropoulos [12] which show that genus- g graphs *do embed* into distributions over planar graphs. Of course, the family of genus- g graphs is not closed under 2-sums.

Our construction is based on “linearizing” an arbitrary base graph so that all edges are pointed in the same “direction,” and then recursively 2-summing the graph with itself. The analysis then follows the form of Rabinovich and Raz [24] and Gupta, et al. [10].

1.2 Preliminaries

Graphs and metrics. We deal exclusively with finite graphs $G = (V, E)$ which are free of loops and parallel edges. We will also write $V(G)$ and $E(G)$ for the vertex and edge sets of G , respectively. In Section 2, we deal exclusively with unweighted graphs, while in other sections we sometimes equip G with a length function $\text{len} : E \rightarrow \mathbb{R}_+$ on edges. We will denote the metric space associated with a graph G as (V, d_G) , where d_G is the shortest path metric according to the edge lengths. We say that a length function len is *reduced* if G is $\text{len}(u, v) = d_G(u, v)$ for every $(u, v) \in E$.

2-sums and $\bar{2}$ -sums. A standard construction in topological graph theory takes two disjoint graphs $G = (V, E)$ and $H = (W, F)$ and constructs the *2-sum* $G \oplus_2 H$, which arises by first taking the disjoint union of G and H , and then choosing edges $e \in E$ and $f \in F$, identifying them, together with their endpoints, and removing the resulting joined edge. The notation \oplus_2 is ambiguous, as it doesn’t specify how the graphs are summed together, but will always be clear from context. If the joined edge is not removed, we refer to this as a $\bar{2}$ -sum. Note that the 2-sum is always a subgraph of the $\bar{2}$ -sum.

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 pseudometric; using indicator functions, we can write the cut pseudometric 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 [5]).

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

$$\|f(u) - f(v)\|_1 = \int |\mathbf{1}_S(u) - \mathbf{1}_S(v)| d\mu(S), \quad (2)$$

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 (2) holds. We will use this correspondence freely in what follows.

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 $\text{dist}(f) = \|f\|_{\text{Lip}} \cdot \|f^{-1}\|_{\text{Lip}}$. If $d_Y(f(x), f(y)) \leq d(x, y)$ for every $x, y \in X$, we say that f is *non-expansive*.

If $f : X \rightarrow L_1$ and $g : X \rightarrow L_1$ are two embeddings, we write $f \oplus g : X \rightarrow L_1$ for a mapping that satisfies $\|(f \oplus g)(x) - (f \oplus g)(y)\|_1 = \|f(x) - f(y)\|_1 + \|g(x) - g(y)\|_1$ for every $x, y \in X$.

2 Embeddings for series-parallel graphs

In this section, we prove that $c_1(V, d_G) \leq 2$ whenever $G = (V, E)$ is a series-parallel graph with shortest-path metric d_G . Crucially, our proofs will be based on the assumptions that G is unweighted and bipartite. Suppose that (G, d_G) is a weighted graph metric. We may approximate the weights by rationals, scale them to be integers, and then subdivide every edge of weight $W \in \mathbb{N}$ into a path of length W to obtain an unweighted graph metric $(G', d_{G'})$ into which G embeds with distortion arbitrarily close to 1. If G is K_4 -free, then so is G' . Thus we may assume that such graphs are unweighted without loss of generality. Also, we may replace G' by a bipartite graph G'' into which G' embeds with distortion 1 by replacing every edge of G' with a path of length 2.

2.1 Convex distortion and random extension

We begin with some definitions. Let $G = (V, E)$ be an arbitrary unweighted graph. A path $\pi \in \mathcal{P}_{xy}$ with endpoints $x, y \in V$ is called an *ear* if every vertex of $\text{int}(\pi)$ has degree exactly 2 in G , and the length of π is at least 2. We say that π is *geodesic* if $\text{len}(\pi) = d_G(x, y)$ and *slack* otherwise.

If G is a series-parallel graph with endpoints $s, t \in V$, we say that a path π is *direct* if it is a subpath of an s - t path. We define π to be *minimal* if it is a direct, geodesic ear. We say that $(x, y) \in V^2$ form a *minimal pair* if either $x = y$ or \mathcal{P}_{xy} contains a minimal path. We say that minimal pairs (x, y) and (a, b) are *independent* if, for all geodesic ears $\pi \in \mathcal{P}_{xy}$ and $\pi' \in \mathcal{P}_{ab}$, π and π' are edge-disjoint.

Lemma 2.1. *Either G is a path, or G contains a pair of vertices $x, y \in V$ and two distinct direct ears P, Q both of which have endpoints x, y .*

Proof. Assume that G is not a path. If G was formed by the serial composition of H, H' then one of H or H' is not a path, and we are done by induction. If G was formed by the parallel composition of H, H' then either both H, H' are paths, in which case we may take $\{x, y\} = \{s, t\}$ and $\{P, Q\} = \{H, H'\}$, else one of H or H' is not a path and again we are done by induction. \square

Next, we discuss the canonical way of extending a cut measure over the rest of the graph to a geodesic ear.

Random extension. Suppose that $x, y \in V$ and $x \neq y$. Let $\pi \in \mathcal{P}_{xy}$ be a geodesic ear. Let $\hat{G} = (\hat{V}, \hat{E})$ be the graph G with the vertices of $\text{int}(\pi)$ removed. If we have an embedding $\hat{f} : \hat{V} \rightarrow L_1$ with corresponding cut measure $\hat{\mu}$, we define the *standard extension* $f : V \rightarrow L_1$ implicitly, by defining its corresponding cut measure μ as follows. Order the vertices of π naturally from x to y (which we will refer to as left to right). Let e be a uniformly random edge of π . For every $A \subseteq \hat{V}$, we define a random extension $\text{ext}(A) \subseteq V$ by its characteristic function,

$$\mathbf{1}_{\text{ext}(A)}(u) = \begin{cases} \mathbf{1}_A(u) & \text{if } u \in \hat{V} \\ \mathbf{1}_A(x) & \text{if } u \in \pi \text{ is to the left of } e \\ \mathbf{1}_A(y) & \text{if } u \in \pi \text{ is to the right of } e. \end{cases}$$

Finally, we define μ as follows.

$$\mu(S) := \Pr[S = \text{ext}(S \cap \hat{V})] \cdot \hat{\mu}(S \cap \hat{V}),$$

for any $S \subseteq V$. The following lemma is easy to verify.

Lemma 2.2. *Given f, \hat{f} as above, we have*

1. For any $u \in \pi, v \in V$,

$$\|f(u) - f(v)\|_1 = \frac{d_G(u, y)}{d_G(x, y)} \|f(x) - f(v)\|_1 + \frac{d_G(u, x)}{d_G(x, y)} \|f(y) - f(v)\|_1. \quad (3)$$

2. For $u, v \notin \pi$, $\|f(u) - f(v)\|_1 = \|\hat{f}(u) - \hat{f}(v)\|_1$.

3. $\|f\|_{\text{Lip}} = \|\hat{f}\|_{\text{Lip}}$.

Now we introduce the notion of “imaginary points,” which can be thought of as points that could exist on a new geodesic ear with endpoints $x, y \in V$.

Imaginary points. Given a pair $x, y \in V$, we will use the notation $\alpha\bar{x}\bar{y}$ to denote an “imaginary vertex” lying on the “line” between x and y at distance $\alpha d_G(x, y)$ from y . Thus $0\bar{x}\bar{y} = y$ and $1\bar{x}\bar{y} = x$. We then define, for any $\alpha, \beta \in (0, 1)$, and any pairs of (not necessarily distinct) vertices $x, y \in V$ and $a, b \in V$,

$$d_f(\alpha\bar{x}\bar{y}, \beta\bar{a}\bar{b}) := \alpha\beta \|f(x) - f(a)\|_1 + \alpha(1-\beta) \|f(x) - f(b)\|_1 + \beta(1-\alpha) \|f(y) - f(a)\|_1 + (1-\alpha)(1-\beta) \|f(y) - f(b)\|_1.$$

We also extend $d_G(\cdot, \cdot)$ to imaginary points via

$$d_G(\alpha\bar{x}\bar{y}, \beta\bar{a}\bar{b}) := \min \left\{ \begin{aligned} &(1-\alpha)d_G(x, y) + d_G(x, a) + (1-\beta)d_G(a, b), \\ &(1-\alpha)d_G(x, y) + d_G(x, b) + \beta d_G(a, b), \\ &\alpha d_G(x, y) + d_G(y, a) + (1-\beta)d_G(a, b), \\ &\alpha d_G(x, y) + d_G(y, b) + \beta d_G(a, b) \end{aligned} \right\}$$

Convex embeddings of 4-tuples. We say that $f : V \rightarrow L_1$ has *convex contraction* D on $\{(x, y), (a, b)\} \subseteq V^2$ if the following holds. For any $\alpha, \beta \in [0, 1]$, we have

$$d_f(\alpha\bar{x}\bar{y}, \beta\bar{a}\bar{b}) \geq \frac{d_G(\alpha\bar{x}\bar{y}, \beta\bar{a}\bar{b})}{D}. \quad (4)$$

Now we record two results about convex embeddings.

Lemma 2.3. *If P is a path metric and $f : P \rightarrow L_1$ is an isometry, then f has convex contraction 2 on $\{(x, y), (a, b)\}$ whenever (x, y) and (a, b) are independent in P .*

Proof. It is clear that every isometric embedding of the path extends to an isometric embedding of \mathbb{R} , thus we

will prove the lemma for \mathbb{R} . Let $x, y, a, b \in \mathbb{R}$ be any four points such that $x \leq y$ and $a \leq b$. We now consider two cases.

First, suppose $[x, y] \cap [a, b] = \emptyset$. Then, without loss of generality, we may assume $x \leq y \leq a \leq b$. Let $u = x + (1 - \alpha)|x - y|$ and $v = a + (1 - \beta)|a - b|$. Using $\alpha|x - y| = |u - y|$ and $(1 - \beta)|a - b| = |a - v|$, it is easy to verify that the left hand side of (4) is precisely $|u - v|$. Also, we have $d_G(\alpha\overline{xy}, \beta\overline{ab}) = |u - v|$ because the minimum on the right hand side is basically enumerating possible shortest paths from u to v . So in this case, (4) holds with $D = 1$.

Next, suppose $[x, y] \cap [a, b] \neq \emptyset$. Since the x - y and a - b paths are required to be edge-disjoint, we may assume, without loss of generality, that $x \leq a = b \leq y$. Now (4) with $D = 2$ reduces to

$$2\alpha|x - a| + 2(1 - \alpha)|y - a| \geq \min \{(1 - \alpha)|x - y| + |x - a|, \alpha|x - y| + |y - a|\}.$$

This is worst when $|x - a| = |y - a|$ and $\alpha = \frac{1}{2}$, where we get equality. \square

The next lemma will be used heavily in our inductive argument.

Lemma 2.4. *Suppose that $G = (V, E)$ is a graph and $x, y \in V$ are the endpoints of a geodesic ear P . Let $\hat{V} = V \setminus \text{int}(P)$, consider $\hat{f} : (\hat{V}, d_G) \rightarrow L_1$, and let $f : V \rightarrow L_1$ be the standard extension of \hat{f} . If \hat{f} has convex contraction D on $\{(x, y), (a, b)\}$ for $a, b \in \hat{V}$ then f has convex contraction D on $\{(u, v), (a, b)\}$ for any $u, v \in P$.*

Proof. Consider $u, v \in P$ and $a, b \in \hat{V}$. Let $\delta, \gamma > 0$ be such that $\gamma d_G(x, y) = d_G(u, x)$ and $\delta d_G(x, y) = d_G(v, y)$. For any $\alpha, \beta \in (0, 1)$, we have

$$\begin{aligned} d_f(\alpha\overline{uv}, \beta\overline{ab}) &= \alpha\beta\|f(u) - f(a)\|_1 + \alpha(1 - \beta)\|f(u) - f(b)\|_1 \\ &\quad + \beta(1 - \alpha)\|f(v) - f(a)\|_1 \\ &\quad + (1 - \alpha)(1 - \beta)\|f(v) - f(b)\|_1 \quad (5) \\ &= d_f([\alpha(1 - \gamma) + (1 - \alpha)\delta]\overline{xy}, \beta\overline{ab}), \quad (6) \end{aligned}$$

where, in passing from (5) to (6), we use Lemma 2.2(1) to expand each of the four L_1 distances in the preceding line. Setting $A = \alpha(1 - \gamma) + (1 - \alpha)\delta$, and using the fact that \hat{f} has convex contraction D on $\{(x, y), (a, b)\}$ yields

$$d_f(\alpha\overline{uv}, \beta\overline{ab}) = d_{\hat{f}}(A\overline{xy}, \beta\overline{ab}) \geq \frac{d_G(A\overline{xy}, \beta\overline{ab})}{D}.$$

In order to finish, we need to show that $d_G(A\overline{xy}, \beta\overline{ab}) \geq d_G(\alpha\overline{uv}, \beta\overline{ab})$, which follows from a number of applications of the triangle inequality; details are deferred to the full version. \square

Finally, we have a lemma to deal with slack ears.

Lemma 2.5. *Let $G = (V, E)$ be any bipartite graph with a slack ear P . Then there exists a non-trivial cut $R \subseteq \text{int}(P)$ such that if $\hat{G} = G/\partial R$, then the following holds. For any embedding $\hat{f} : V(\hat{G}) \rightarrow L_1$, there exists an embedding $f : V \rightarrow L_1$ such that*

1. $\text{dist}(f) = \text{dist}(\hat{f})$ and $\|f\|_{\text{Lip}} = \|\hat{f}\|_{\text{Lip}}$
2. For $u, v \in V$, we have

$$\|f(u) - f(v)\|_1 = \|\hat{f}(u) - \hat{f}(v)\|_1 + |\mathbf{1}_R(u) - \mathbf{1}_R(v)|.$$

Proof. Let P be a slack ear in G . Since G is bipartite, we have $\text{len}(P) = d_G(x, y) + 2k$ for some $k \geq 1$. Let $R \subseteq P$ be the subpath of P with the leftmost $\lfloor \frac{k+1}{2} \rfloor$ and the rightmost $\lceil \frac{k+1}{2} \rceil$ vertices removed, so that R is a direct path of length $d_G(x, y) + k - 1$. Let $\hat{G} = G/\partial R$. We claim that for every pair $u, v \in V$, we have

$$d_G(u, v) = d_{\hat{G}}(u, v) + |\mathbf{1}_R(u) - \mathbf{1}_R(v)|. \quad (7)$$

This is easy to see for $u, v \notin P$, since for such a pair $d_{\hat{G}}(u, v) = d_G(u, v)$. Next, the endpoints of R are distance $d_G(x, y) + k - 1$ apart, while any other path between these points has length at least $d_G(x, y) + k + 1$, thus $d_{\hat{G}}(u, v) = d_G(u, v)$ for $u, v \in R$ as well. Similarly, for any $u, v \in P \setminus R$, we have $d_G(u, v) \leq d_G(x, y) + k - 1 \leq \text{len}(R)$, hence $d_{\hat{G}}(u, v) = d_G(u, v)$ for $u, v \notin R$. Finally, for $u \in R$ and $v \in P \setminus R$, it is obvious that $d_{\hat{G}}(u, v) = d_G(u, v) - 1$, completing the proof of (7).

Consider now any embedding $\hat{f} : V(\hat{G}) \rightarrow L_1$. To obtain a mapping $f : V \rightarrow L_1$, we simply define $f = \hat{f} \oplus \mathbf{1}_R$, where \hat{f} can be naturally thought of as a map on G . By (7), we have $\text{dist}(f) = \text{dist}(\hat{f})$ and $\|f\|_{\text{Lip}} = \|\hat{f}\|_{\text{Lip}}$. Claim (2) of the lemma follows immediately. \square

2.2 The embedding

We now show that every series-parallel graph embeds into L_1 with distortion at most 2. The embedding is constructed inductively; at every step, we find either a taut ear or a slack ear. A taut ear is removed from the graph; an embedding of the resulting graph is constructed inductively, and is extended to the taut ear via random




Figure 1. In this orientation, no slack ears occur in the induction.




Figure 2. Embedding of $K_{2,3}$ where a slack ear occurs in the second step.

extension and Lemma 2.2. Upon encountering a slack ear, we make a cut according to Lemma 2.5, contract the two edges crossing this cut, and induct on the remaining graph. To analyze the complex interleaving of these two operations, our inductive hypothesis maintains small convex contraction on certain pairs in the graph.

Figures 1 and 2 show two possible induction sequences for $K_{2,3}$, depending on the choice of endpoints. Initial cuts are denoted by dotted lines, while cuts that were formed by random extension are denoted by dashed lines. The inductive sequence of graphs should be read left to right, while the cuts are formed from right to left.

Theorem 2.6. *Let $G = (V, E)$ be a bipartite series-parallel graph with endpoints $s, t \in V$. Let the notions of minimal pairs and minimal paths be defined with respect to (s, t) . Then there is a mapping $f : V \rightarrow L_1$ with $\|f\|_{\text{Lip}} \leq 1$ and $\text{dist}(f) \leq 2$ such that the following properties hold.*

- (P1) For any minimal pair $(x, y) \in V^2$, we have $\|f(x) - f(y)\|_1 = d_G(x, y)$.
- (P2) If (x, y) and (a, b) are independent minimal pairs, then f has convex contraction 2 on $\{(x, y), (a, b)\}$.

Proof. Our proof is by induction on $|V|$. First, assume that G is a path P (in particular, this holds if $|V| \leq 2$). Then obviously G possesses an isometric embedding into L_1 , and by Lemma 2.3, we are done.

If G is not a path, then by Lemma 2.1, there exists a pair $x, y \in V$ connected by a pair of distinct direct ears P, Q . The analysis now splits into two cases.

Case I: Both P and Q are geodesic.

Let \hat{G} be the graph obtained by removing the vertices of $P \setminus \{x, y\}$ from G . It is straightforward to verify that \hat{G} is a bipartite series-parallel graph with endpoints s, t . By induction, \hat{G} has an embedding $\hat{f} : V(\hat{G}) \rightarrow L_1$ with $\text{dist}(\hat{f}) \leq 2$ and which satisfies properties (P1) and (P2). Let $f : V \rightarrow L_1$ be the standard extension of \hat{f} .

The fact that $\|f\|_{\text{Lip}} \leq 1$ follows from part (3) of Lemma 2.2. From parts (1) and (2) of Lemma 2.2, we only need to check the condition $\text{dist}(f) \leq 2$ for pairs $u \in P$ and $v \notin P$. To this end, let $\alpha = \frac{d_G(u, y)}{d_G(x, y)}$ and use Lemma 2.2(1) to write

$$\begin{aligned}
& \|f(u) - f(v)\|_1 \\
&= \alpha \|f(x) - f(v)\|_1 + (1 - \alpha) \|f(y) - f(v)\|_1 \\
&= \alpha \|\hat{f}(x) - \hat{f}(v)\|_1 + (1 - \alpha) \|\hat{f}(y) - \hat{f}(v)\|_1 \\
&= d_f(\alpha \overline{xy}, v) \\
&\geq \frac{d_G(\alpha \overline{xy}, v)}{2} \\
&\geq \frac{d_G(u, v)}{2}.
\end{aligned}$$

where the first inequality uses (4) with $\beta = 1$ and $v = a = b$, along with the induction hypothesis (P2) applied to the minimal pairs (x, y) and (v, v) . Note that (x, y) is minimal in \hat{G} because it is connected by the minimal path Q . The final inequality is simply the triangle inequality.

Now, suppose the pair (u, v) is minimal in G , but not minimal (or existent) in \hat{G} . We claim that this can only occur for $u, v \in P$. This is because neither x nor y can be an internal vertex for a direct ear in G , since in that case they would have degree at least 3 (being connected to Q, P , and some other vertex).

Thus, verifying (P1) for f is easy. Using the fact that $\|f(x) - f(y)\|_1 = \|\hat{f}(x) - \hat{f}(y)\|_1 = d_G(x, y)$ by induction, Lemma 2.2 implies that f is isometric on P . In particular, (P2) is verified for pairs $(u, v), (a, b)$ with $u, v, a, b \in P$. Hence we need only verify (P2) for $u, v \in P$ and a minimal pair $a, b \in V(\hat{G})$. But this follows from Lemma 2.4.

Case II: P is slack.

Let $R \subseteq \text{int}(P)$ be the subset guaranteed by Lemma 2.5. Obviously $\hat{G} = G/\partial R$ is a bipartite series-parallel graph with endpoints s and t , so by induction we have a distortion 2 embedding $\hat{f} : V(\hat{G}) \rightarrow L_1$ for which properties (P1) and (P2) hold with respect to \hat{G} . Applying Lemma 2.5, we arrive at a mapping $f : V \rightarrow L_1$. Since every minimal pair in G is also a minimal pair in \hat{G} , f satisfies (P1) and (P2) with respect to G . \square

3 Lower bounds on probabilistic embeddings into simpler graphs

We now present a general lower bound for embedding one class of graphs probabilistically into another. We rely heavily on the approach of Gupta et al. [10]. Our main theorem follows. We recall that, for $p \in \mathbb{N}$, a p -subdivision of a graph H refers to the graph which results from replacing every edge of H by an ear of length p . If H is weighted, then every edge of an ear in the k -subdivision is given weight equal to that of the edge it replaces.

Theorem 3.1. *Let G be an unweighted graph with path metric d_G , and let \mathcal{X} be any family of metric spaces. Suppose there exists a constant $c \geq 1$ such that embedding the p -subdivision of (G, d_G) into any member of \mathcal{X} requires distortion at least $c \cdot p$. Then there exists an infinite family of graphs $\{H_k\}_{k=1}^\infty$ with $|H_k| \rightarrow \infty$ such that the following holds: Each H_k can be constructed by 2-sums starting with copies of G and a collection of disjoint cycles, and any probabilistic embedding of H_k into a distribution over spaces in \mathcal{X} incurs distortion $\Omega(\log |V(H_k)|)$.*

For simplicity, we mention the following corollary.

Corollary 3.2. *Let \mathcal{G} and \mathcal{H} be two minor-closed families of graphs, such that \mathcal{G} has at least one member which is not a tree, and \mathcal{G} is closed under taking 2-sums. If every n -vertex member of \mathcal{G} embeds into a distribution over graphs from \mathcal{H} with $o(\log n)$ distortion, then $\mathcal{G} \subseteq \mathcal{H}$.*

Some special cases of this theorem are relevant to the problems in this paper, and close the possibility of a certain strategy for L_1 embeddings. The proofs of these corollaries is given at the end of the present section.

Corollary 3.3. *There exists an infinite family $\{H_k\}$ of W_4 -free graphs for which any probabilistic embedding into a distribution over K_4 -free graphs has distortion $\Omega(\log |V(H_k)|)$.*

Corollary 3.4. *There exists an infinite family $\{H_k\}$ of $(K_5 \setminus e)$ -free graphs for which any probabilistic embedding into a distribution over planar graphs has distortion $\Omega(\log |V(H_k)|)$.*

Corollary 3.5. *For any integer $w \geq 2$, there exists an infinite family $\{H_k\}$ of planar treewidth- w graphs for which any probabilistic embedding into a distribution over treewidth- $(w - 1)$ graphs has distortion $\Omega(\log |V(H_k)|)$.*

These results are obviously tight, since [8] gives an $O(\log n)$ -distortion embedding of any n -point metric spaces into a distribution over trees.

The proof of Theorem 3.1 is a generalization of the lower bound from [10] for embedding series-parallel (treewidth 2) graphs into distributions over dominating trees. Their diamond graph is a special case of our construction with the 2×2 grid acting as the base graph. However, extending their results to more arbitrary graphs relies on some additional insights.

Proof of Theorem 3.1. We begin by defining the following graph property.

Definition 3.6. *A weighted, undirected graph $G = (V, E)$ is said to be (s, t) -aligned if $(s, t) \in E$, and for all $e \in E$, there exists an s - t geodesic containing e .*

Lemma 3.7. *For every unweighted 2-connected graph $G = (V, E)$, and $(s, t) \in E$, there exists an assignment of non-negative integral weights to E under which G is (s, t) -aligned.*

The proof of Lemma 3.7 is based on the following simple lemma from [17].

Lemma 3.8 ([17]). *A graph $G = (V, E)$ is 2-connected if and only if for every $(s, t) \in E$, there exists an ordering $s = v_0, v_1, \dots, v_{n-1} = t$ of V such that for every $j \in \{1, 2, \dots, n - 2\}$, there are i, k with $i < j < k$ and $(v_i, v_j), (v_j, v_k) \in E$.*

Next, we perform a further modification to a graph G , allowing it to serve as the base level of our construction.

Lemma 3.9. *For any (s, t) -aligned, integrally weighted graph $G = (V, E)$, there exists an unweighted, undirected graph $G' = (V', E')$ with $V \subseteq V'$, such that G embeds into G' with distortion 1, and such that E' can be partitioned into disjoint s - t geodesics in G' . Furthermore, the graph $G'' = (V', E' \cup \{(s, t)\})$ is constructible by taking the 2-sum of G with a number of cycles.*

Proof. We construct G' through a series of alterations to G . First let $G_0 = (V_0, E_0)$ be G with each edge replaced by an unweighted ear of the same length. This is equivalent to 2-summing an edge of length L with a cycle of length $L + 1$.

We then create $G' = (V', E')$ as follows. We further subdivide each edge of E_0 once, so that it becomes an ear of length 2. Now, for each edge $e \in E_0$, pick an arbitrary s - t geodesic P of G_0 that passes through e (such a path exists because G is (s, t) -aligned). For each edge $f \neq e$ in P , add an additional length-2 ear to G' , parallel

to the length-2 ears already corresponding to f . This operation can clearly be performed by 2-summing with a number of disjoint cycles. Assign to e , not the geodesic P , but the geodesic P' which uses the newly added ears. Each of the s - t paths formed in this way is a geodesic, and they form a partition of the edges of the final graph G' .

For technical reasons, we need to produce the graph G'' which is G' with an additional (s, t) edge. Clearly G'' can also be constructed by 2-sums with disjoint cycles. \square

We are now ready to describe our recursive construction. Given an initial unweighted graph G , let H_1 be the unweighted graph G' resulting from applying Lemma 3.7 to G , and then Lemma 3.9 to the resulting weighted version of G . For $k \geq 2$, we construct H_k by replacing every edge of H_{k-1} by a copy of H_1 , identifying the endpoints of the edge with $s, t \in V(H_1)$ arbitrarily. This operation can be performed by 2-summing copies of G'' onto H_{k-1} , and thus each H_k is constructible via 2-sums of G and series-parallel graphs. There is a natural way to think about $V(H_1) \subseteq V(H_2) \subseteq \dots \subseteq V(H_{k-1}) \subseteq V(H_k)$, and also of $E(H_i)$ as a set of pairs in $V(H_k)$ for $i \leq k$. This will be useful in what follows.

We now turn to lower bounding the distortion of probabilistically embedding H_j into a distribution over \mathcal{X} . As in [10], it suffices to prove that for any $(X, d) \in \mathcal{X}$, and any non-contracting mapping $\varphi : V(H_k) \rightarrow X$, we have

$$\sum_{(u,v) \in E(H_k)} d_X(\varphi(u), \varphi(v)) \geq \Omega(\log |H_k|) \cdot |E(H_k)|.$$

Let $m = |E(H_1)|$.

Proposition 3.10. *For all $1 \leq i \leq k$, H_k contains m^{k-i} edge-disjoint copies of H_i .*

Proof. This can easily be observed by induction over $k \geq i$. It is obviously for $i = k$. For $k > i$, each ‘‘edge’’ of $E(H_1)$ in H_k is really a copy of H_{k-1} , which contains m^{k-1-i} copies of H_i ; hence H_k contains m^{k-i} copies. \square

Let q refer to the number of disjoint geodesic s - t geodesics in H_1 , and let r refer to the length of these paths. As shown earlier, $m = q \cdot r$.

Proposition 3.11. *H_i contains q^{i-1} edge-disjoint r^{i-1} -subdivisions of H_1 .*

Proof. This is trivial for $i = 1$. For larger i , notice that each edge of the initial H_1 structure has been replaced by q^{i-1} edge-disjoint paths of length r^{i-1} . This too can be easily seen by induction. \square

Now, because G embeds into H_1 with distortion 1, every p -subdivision of H_1 must incur distortion at least $c \cdot p$ under any embedding into a member of \mathcal{X} . Hence, in any such subdivision, there must be at least one edge (u, v) with $d_X(\varphi(u), \varphi(v)) \geq c \cdot p$ (recalling that φ is non-contractive).

For any edge $(u, v) \in E(H_k)$, assign the color $i \in \mathbb{N}$ to (u, v) if $d_X(\varphi(u), \varphi(v)) \geq c \cdot r^{i-1}$. Let $C_e \subseteq \mathbb{N}$ be the set of colors assigned to the edge $e \in E(H_k)$. Then, combining Propositions 3.10 and 3.11, we see that for $1 \leq i \leq k$, we have $|\{e : i \in C_e\}| \geq m^{k-i} \cdot q^{i-1}$.

It follows that

$$\begin{aligned} \sum_{(u,v) \in E(H_k)} d_X(\varphi(u), \varphi(v)) &\geq \sum_{(u,v) \in E} \max_{i \in C(u,v)} c \cdot r^{i-1} \\ &\geq \sum_{(u,v) \in E} \frac{c}{r} \sum_{i \in C(u,v)} r^{i-1} \\ &= \frac{c}{r} \sum_{i=1}^k |\{e : i \in C_e\}| \cdot r^{i-1} \\ &\geq \frac{c}{r} \sum_{i=1}^k m^{k-i} \cdot q^{i-1} \cdot r^{i-1} \\ &= \frac{c}{r} \sum_{i=1}^k m^{k-i} \cdot m^{i-1} \\ &= \frac{c}{r} \sum_{i=1}^k m^{k-1} \\ &= \frac{c \cdot k \cdot m^{k-1}}{r} \\ &= k \cdot \frac{c}{mr} |E(H_k)|, \end{aligned}$$

where we have used the fact that $|E(H_k)| = m^k$. Since, $c, m, r = O(1)$, we have bounded the average edge stretch from below by $\Omega(k) = \Omega(\log |V(H_k)|)$. \square

We now move onto proving Corollaries 3.2–3.5.

Proof of Corollary 3.2. We begin with the following lemma, which follows from the work of Rabinovich and Raz [24].

Lemma 3.12 (Carroll and Goel [3]). *Let G be unweighted graph, and H a weighted graph, such that H does not contain G as a minor. Then any embedding of the p -subdivision of G into H incurs distortion at least $k/6 - 1/2$.*

Let \mathcal{X} be the set of metrics induced on graphs from \mathcal{H} . Suppose, for the sake of contradiction, that $\mathcal{G} \not\subseteq \mathcal{H}$, and let $G \in \mathcal{G} \setminus \mathcal{H}$. Now apply Theorem 3.1 to G and \mathcal{X} to get a family of graphs $\{H_k\}$. Since \mathcal{G} is minor-closed and contains a non-tree member, it also contains a cycle of length at least 3. Since \mathcal{G} is closed under 2-sums and minor-closed, it therefore contains a cycle of every length. Thus we have $H_k \in \mathcal{G}$ for every $k \geq 1$, hence n -vertex members of \mathcal{G} do not embed into distributions over members of \mathcal{H} with $o(\log n)$ distortion, completing the proof. \square

References

- [1] A. Andoni, M. Deza, A. Gupta, P. Indyk, and S. Raskhodnikova. Lower bounds for embedding edit distance into normed spaces. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, MD, 2003)*, pages 523–526, New York, 2003. ACM.
- [2] 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.
- [3] D. Carroll and A. Goel. Lower bounds for embedding into distributions over excluded minor graph families. In *Proceedings of the 12th European Symposium on Algorithms*, 2004.
- [4] 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 (electronic), 2006.
- [5] M. M. Deza and M. Laurent. *Geometry of cuts and metrics*, volume 15 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1997.
- [6] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, third edition, 2005.
- [7] W. Evans and M. A. Safari. On the ℓ_1 embedding of series-parallel graphs. Manuscript, 2007.
- [8] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 448–455, 2003.
- [9] J. Fakcharoenphol and K. Talwar. An improved decomposition theorem for graphs excluding a fixed minor. In *Proceedings of 6th Workshop on Approximation, Randomization, and Combinatorial Optimization*, volume 2764 of *Lecture Notes in Computer Science*, pages 36–46. Springer, 2003.
- [10] A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair. Cuts, trees and l_1 -embeddings of graphs. *Combinatorica*, 24(2):233–269, 2004.
- [11] 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.
- [12] P. Indyk and A. Sidiropoulos. Probabilistic embeddings of bounded genus graphs into planar graphs. In *Proceedings of the 23rd Annual Symposium on Computational Geometry*. ACM, 2007.
- [13] P. N. Klein, S. A. Plotkin, and S. Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pages 682–690, 1993.
- [14] R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: a new embedding method for finite metrics. *Geom. Funct. Anal.*, 15(4):839–858, 2005.
- [15] J. R. Lee and A. Naor. Extending Lipschitz functions via random metric partitions. *Invent. Math.*, 160(1):59–95, 2005.
- [16] J. R. Lee and P. Raghavendra. Coarse differentiation and multi-flows in planar graphs. 2007.
- [17] A. Lempel, S. Even, and I. Cederbaum. An algorithm for planarity testing of graphs. In *Theory of Graphs, International Symposium (Rome, 1966)*, pages 215–232, 1966.
- [18] 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.
- [19] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- [20] J. Matoušek. Open problems on low-distortion embeddings of finite metric spaces. Online: <http://kam.mff.cuni.cz/~matousek/metrop.ps>.
- [21] J. Matoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.
- [22] H. Okamura. Multicommodity flows in graphs. *Discrete Appl. Math.*, 6(1):55–62, 1983.
- [23] H. Okamura and P. D. Seymour. Multicommodity flows in planar graphs. *J. Combin. Theory Ser. B*, 31(1):75–81, 1981.
- [24] Y. Rabinovich and R. Raz. Lower bounds on the distortion of embedding finite metric spaces in graphs. *Discrete Comput. Geom.*, 19(1):79–94, 1998.
- [25] S. Rao. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In *Proceedings of the 15th Annual Symposium on Computational Geometry*, pages 300–306, New York, 1999. ACM.
- [26] N. Robertson and P. D. Seymour. Graph minors. XX. Wagner’s conjecture. *J. Combin. Theory Ser. B*, 92(2):325–357, 2004.
- [27] A. Schrijver. Distances and cuts in planar graphs. *J. Combin. Theory Ser. B*, 46(1):46–57, 1989.