A face cover perspective to ℓ_1 embeddings of planar graphs

Arnold Filtser *

Abstract

It was conjectured by Gupta et al. [Combinatorica04] that every planar graph can be embedded into ℓ_1 with constant distortion. However, given an n-vertex weighted planar graph, the best upper bound on the distortion is only $O(\sqrt{\log n})$, by Rao [SoCG99]. In this paper we study the case where there is a set K of terminals, and the goal is to embed only the terminals into ℓ_1 with low distortion. In a seminal paper, Okamura and Seymour [J.Comb.Theory81] showed that if all the terminals lie on a single face, they can be embedded isometrically into ℓ_1. The more general case, where the set of terminals can be covered by γ faces, was studied by Lee and Sidiropoulos [STOC09] and Chekuri et al. [J.Comb.Theory13]. The state of the art is an upper bound of $O(\log \gamma)$ by Krauthgamer, Lee and Rika [SODA19]. Our contribution is a further improvement on the upper bound to $O(\sqrt{\log \gamma})$. Since every planar graph has at most $O(n)$ faces, any further improvement on this result, will be a major breakthrough, directly improving upon Rao's long standing upper bound. Moreover, it is well known that the flow-cut gap equals to the distortion of the best embedding into ℓ_1. Therefore, our result provides a polynomial time $O(\sqrt{\log \gamma})$-approximation to the sparsest cut problem on planar graphs, for the case where all the demand pairs can be covered by γ faces.

1 Introduction

Metric embeddings is a widely used algorithmic technique that have numerous applications, notably in approximation, online and distributed algorithms. In particular, embeddings into ℓ_1 have implications to graph partitioning problems. Specifically, the ratio between the Sparsest Cut and the maximum multicommodity flow (also called flow cut gap) is upper bounded by the distortion of the optimal embedding into ℓ_1 (see [LLR95, GNRS04]).

Given a weighted graph $G = (V, E, w)$ with the shortest path metric d_G, and embedding $f : V \to \ell_1$, the *contraction* and *expansion* of f are the smallest τ, ρ, respectively, such that for every pair $u, v \in V$,

$$\frac{1}{\tau} \cdot d_G(u, v) \leq \|f(u) - f(v)\|_1 \leq \rho \cdot d_G(u, v) \quad .$$

The *distortion* of the embedding is $\tau \cdot \rho$. If $\tau = 1$ (resp. $\rho = 1$) we say that the embedding is non-contractive (expansive). If $\rho = O(1)$, we say that the embedding is Lipschitz.

In this paper we focus on embeddings of planar graphs into ℓ_1. Rao [Rao99] showed that every n-vertex planar graph can be embedded into ℓ_1 with distortion $O(\sqrt{\log n})$. The best known lower bound is 2 by Lee and Raghavendra [LR10]. A long standing conjecture by Gupta *et al.* [GNRS04] states that every graph family excluding a fixed minor, and in particular planar graphs, can be embedded into ℓ_1 with constant distortion.

Consider the case where there is a set $K \subseteq V$ of terminals, and we are only interested in embedding the terminals into ℓ_1. This version is sufficient for the flow cut-gap equivalence, where the terminals are the vertices with demands. Better embeddings might be constructed when K has a special structure. A *face cover* of G is a set of faces such that every terminal belongs to some face from the set (see Figure 1 for an illustration). Given a drawing of G in the plane, denote by $\gamma(G, K)$ the minimal size of a face cover. It was

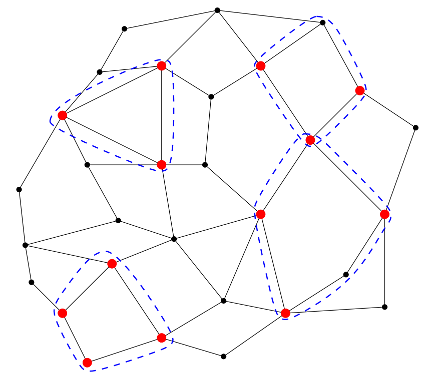

Figure 1: The terminal vertices colored in red. The size of the face cover is 4. The faces in the cover are encircled by a blue dashed lines.

*Columbia University. Email: arnold273@gmail.com. Most of the work was conducted while the author was affiliated with Ben-Gurion University of the Negev. Work supported in part by Simons Foundation, ISF grant 1817/17, and by BSF Grant 2015813.

shown by Hurkens, Schrijver and Tardos [HST86], that the result of Okamura and Seymour [OS81] implies that if $\gamma(G, K) = 1$, that is, all the terminals lie on a single face, then K embeds isometrically into ℓ_1 (a special case is when G is outerplanar). For the general case, where $\gamma(G, K) = \gamma \geq 1$, the methods of Lee and Sidiropoulos [LS09] imply distortion of $2^{O(\gamma)}$. Chekuri, Shepherd and Weibel [CSW13] constructed an embedding with distortion of 3γ. Recently, Krauthgamer, Lee and Rika [KLR19] managed to construct an embedding into ℓ_1 with $O(\log \gamma)$ distortion by first applying a stochastic embedding into trees. This method has benefits, since trees are very simple to work with. Additionally, the result of [KLR19] is tight w.r.t stochastic embedding into trees. We improve upon [KLR19] by embedding directly into ℓ_1.

THEOREM 1.1. *Let $G = (V, E, w)$ be a weighted planar graph with a given drawing in the plane and $K \subseteq V$ a set of terminals. There is an embedding of K into ℓ_1 with distortion $O(\sqrt{\log \gamma(G, K)})$. Moreover, this embedding can be constructed in polynomial time.*

Since every n-vertex graph has $O(n)$ faces, by setting $K = V$, Theorem 1.1 re-proves the celebrated result of Rao [Rao99]. Moreover, any improvement upon Theorem 1.1 will be a major breakthrough.

In addition, Theorem 1.1 has implication on the *sparsest cut* problem. Let $c : E \to \mathbb{R}_+$ be an assignment of capacities to the edges, and $d : \binom{K}{2} \to \mathbb{R}_+$ assignment of demands to terminal pairs. The sparsity of a cut S is the ratio between the capacity of the edges crossing the cut to the demands crossing the cut. The sparsest cut is the cut with minimal sparsity. Theorem 1.1 implies:

COROLLARY 1.1. *Let $G = (V, E)$ be a weighted planar graph with a given drawing in the plane, $K \subseteq V$ a set of terminals, capacities $c : E \to \mathbb{R}_+$ and demands $d : \binom{K}{2} \to \mathbb{R}_+$. Let $\gamma(G, K) = \gamma$. Then there is a polynomial time $O(\sqrt{\log \gamma})$-approximation algorithm for the sparsest-cut problem.*

See [LLR95, GNRS04, KLR19] for further details.

1.1 Technical Ideas In a recent paper of Abraham *et al.* [AFGN18], among other results, the authors constructed an $O(\sqrt{\log n})$-distortion embedding of planar graphs into ℓ_1. This embedding is based on *shortest path decompositions* (SPD). Even though the distortion is similar, the new embedding is very different from the classic embedding of Rao [Rao99]. An SPD is a hierarchical decomposition of a graph using shortest paths. The first level of the partition is simply V. In level i, all the clusters are connected. To construct level $i + 1$, we remove a single shortest path from every cluster of level i. Level $i + 1$ consists of the remaining connected components. This process is repeated until all the vertices are removed. The SPDdepth is the depth of the hierarchy. Using cycle separators [Mil86] it is possible to create an SPD of depth $O(\log n)$ for every planar graph. [AFGN18] showed that every graph which admits an SPD of depth k, can be embedded into ℓ_1 with distortion $O(\sqrt{k})$. In particular $O(\sqrt{\log n})$ for planar graphs.

In this paper we generalize the notion of SPD by defining *partial* SPD (PSPD). The difference is that in PSPD we do not need all the vertices to be removed. That is, in PSPD the last level of the hierarchy is allowed to be non-empty. Given a planar graph G with a terminal set K and a face cover of size $\gamma(G, K) = \gamma$, using cycle separators [Mil86] we create a PSPD of depth $O(\log \gamma)$, such that for every cluster C in the lower level of the hierarchy, all the remaining terminals $K \cap C$ lie on a single face. In other words, each such cluster is an Okamura-Seymour (O-S) graph.

We invoke the embedding of [AFGN18] on our PSPD , as a result we get an embedding with expansion $O(\sqrt{\log \gamma})$, where every pair of terminals v, u that either was separated by the PSPD , or lie close enough to some removed shortest path, has constant contraction. All is left to do is take care of terminal pairs that remained in the same cluster, and lie far from the cluster boundary. As each such cluster is O-S graph, it embeds isometrically to ℓ_1. However, we cannot simply embed each cluster independently of the entire graph. Such an oblivious embedding will create an unbounded expansion, as close-by pairs might belong to different clusters.

Our solution, and the main technical part of the paper, is to create a truncated embedding [1]. Specifically, consider a cluster C where all the terminals lie on a single face F. Let $\mathcal{B} = V \setminus C$ be the boundary of C, which is the set of vertices outside C. We construct a Lipschitz embedding f of F into ℓ_1 such that the norm $\|f(v)\|_1$ of every vertex $v \in F$ is bounded by its distance to the boundary $d_G(v, \mathcal{B})$, while f has constant contraction for pairs far enough from the boundary. Our final embedding is defined as a concatenating of the embedding for the PSPD with a truncated embedding for every cluster, providing a constant contraction on all pairs and $O(\sqrt{\log \gamma})$ expansion.

Our truncated embedding does not uses the embedding of [OS81]. As a middle step, given a parameter $t > 0$, we provide a uniformly truncated embedding [2]

[1] An embedding $f : V \to \ell_1$ is truncated if for every vertex v, $\|f(v)\|_1$ is bounded by some formerly specified number.

[2] In uniformly truncated embedding $\|f(v)\|_1$ is bounded for all the vertices by a global single parameter.

f_t such that f_t is Lipschitz, the norm $\|f_t(v)\|_1$ of every vertex $v \in F$ is exactly t, and f_t provides constant contraction for pairs at distance at most t. The construction of the uniformly truncated embedding goes through a stochastic embedding into trees. To create the non-uniformly truncated embedding we combine uniformly truncated embeddings for all possible truncation scales.

1.2 Related Work

The notion of face cover $\gamma(G, K)$ was extensively studied in the context of Steiner tree problem [EMJ87, Ber90, KNvL19], cuts and (multicommodity) flows [MNS85, CW04], all pairs shortest path [Fre91, Fre95, CX00] and cut sparsifiers [KR17, KPZ17]. Given a drawing and a terminal set K, $\gamma(G, K)$ can be found in $2^{O(\gamma(G,K))} \cdot \text{poly}(n)$ time, but generally it is known to be NP-hard [BM88]. Frederickson [Fre91] (Lemma 7.1) presented a polynomial-time approximation scheme (PTAS) for the problem of finding a face cover of minimum size. Specifically, given a planar graph with a drawing, Frederickson's algorithm finds a face cover of size at most $(1+\epsilon)\cdot\gamma(G, K)$ in $O(2^{\frac{3}{\epsilon}} \cdot n)$ time. Denote by $\gamma^*(G, K)$ the minimal size of a face cover over all planar drawings of G. It is known that computing $\gamma^*(G, K)$ is NP-hard [BM88]. Frederickson [Fre91] presented a 4-approximation for $\gamma^*(G, K)$ in the special case where $K = V$, i.e. the terminals are the entire set V. However, for general $K \subseteq V$, to the best of the author's knowledge, no approximation is known.

It is well known that Euclidean metrics, as well as distributions over trees, embed isometrically into ℓ_1 (See [Mat02]). Therefore, in order to construct a bounded distortion embedding to ℓ_1, it is enough to embed into either ℓ_2 or a distribution over trees.

Outerplanar graphs are 1-outerplanar. A graph is called k-outerplanar, if by removing all the vertices on the outer face, the graph becomes $k-1$-outerplanar. Chekuri *et al.* [CGN$^+$06] proved that k-outerplanar graphs embed into distribution over trees with $2^{O(k)}$ distortion.

Next consider minor-closed graph families. Following [GNRS04], Chakrabarti *et al.* [CJLV08] showed that every graph with treewidth-2 (which excludes K_4 as a minor) embeds into ℓ_1 with distortion 2 (which is tight, as shown by [LR10]). Already for treewidth-3 graphs, it is unknown whether they embed into ℓ_1 with a constant distortion. Abraham *et al.* [AFGN18] showed that every graph with pathwidth k embeds into ℓ_1 with distortion $O(\sqrt{k})$ (through ℓ_2), improving a previous result of Lee and Sidiropoulos [LS13] who showed a $(4k)^{k^3+1}$ distortion (via embedding into trees). Graphs with treewidth k are embeddable into ℓ_2 with distortion $O(\sqrt{k\log n})$ [KLMN05]. For genus g graphs, [LS10] showed an embedding into Euclidean space with distortion $O(\log g + \sqrt{\log n})$. Finally, for H-minor-free graphs, combining the results of [AGG$^+$14, KLMN05] provides Euclidean embeddings with $O(\sqrt{|H|\log n})$ distortion.

For other notions of distortion, Abraham *et al.* [ABN11] showed that β-decomposable metrics (which include planar graphs as well as all other families mentioned in this section), for fixed β, embed into ℓ_2 with scaling distortion $O(\sqrt{\log\frac{1}{\epsilon}})$. This means that for every $\epsilon \in (0,1)$ all but an ϵ fraction of the pairs in $\binom{V}{2}$ have distortion at most $O(\sqrt{\log\frac{1}{\epsilon}})$. Bartal *et al.* [BFN16] proved that β-decomposable metrics (for fixed β) embed into ℓ_2 with prioritized distortion $O(\sqrt{\log j})$. In more detail, given a priority order $\{v_1, \dots, v_n\}$ over the vertices, the pair $\{v_i, v_j\}$ for $j \le i$, will have distortion at most $O(\sqrt{\log j})$

2 Preliminaries

Graphs. We consider connected undirected graphs $G = (V, E)$ with edge weights $w : E \to \mathbb{R}_{\ge 0}$. Let d_G denote the shortest path metric in G. For a vertex $x \in V$ and a set $A \subseteq V$, let $d_G(x, A) := \min_{a\in A} d(x, a)$, where $d_G(x, \emptyset) := \infty$. For a subset of vertices $A \subseteq V$, let $G[A]$ denote the induced graph on A. Let $G \setminus A := G[V \setminus A]$ be the graph after deleting the vertex set A from G.

See Section 1 for definitions of *embedding, distortion, contraction*, *expansion* and *Lipschitz*. We say that an embedding is *dominating* if it is non-contractive. Given a graph family $\mathcal{F}$, a *stochastic embedding* of G into $\mathcal{F}$ is a distribution $\mathcal{D}$ over pairs (H, f_H) where $H \in \mathcal{F}$ and f_H is an embedding of G into H. We say that $\mathcal{D}$ is dominating if for every $(H, f_H) \in \text{supp}(\mathcal{D})$, f_H is dominating. We say that a dominating stochastic embedding $\mathcal{D}$ has expected distortion t, if for every pair $u, v \in V$ it holds that

$$\mathbb{E}_{(H,f_H)\sim\mathcal{D}}\left[d_H\left(f_H(u), f_H(v)\right)\right] \le t \cdot d_G(u, v)\ .$$

A *terminated* planar graph $G = (V, E, w, K)$ is a planar graph (V, E, w), with a subset of terminals $K \subseteq V$. A graph G is outerplanar if there is a drawing of G in the plane such that all the vertices lie on the unbounded face. A *face cover* is a set of faces such that every terminal lies on at least one face from the cover. Given a graph G with a drawing in the plane, denote by $\gamma(G, K)$ the minimal size of a face cover. In the special case where all the terminals are covered by a single face, i.e. $\gamma(G, K) = 1$, we say that G is an Okamura-Seymour graph, or O-S graph for short.

A *tree decomposition* of a graph $G = (V, E)$ is a tree T with nodes $B_1, \dots, B_s$ (called *bags*) where each B_i is a subset of V such that: (1) For every edge $\{u, v\} \in E$, there is a bag B_i containing both u and v. (2) For every vertex $v \in V$, the set of bags containing

v form a connected subtree of T. The *width* of a tree decomposition is $\max_i\{|B_i|-1\}$. The *treewidth* of G is the minimal width of a tree decomposition of G. It is straightforward to verify that every tree graph has treewidth 1.

Given a set of s embeddings $f_i : V \to \mathbb{R}^{d_i}$ for $i \in \{1,\dots,s\}$, the *concatenation* of $f_1,\dots,f_s$, denoted by $\oplus_{i=1}^{s} f_i$, is a function $f : V \to \mathbb{R}^{\sum_i d_i}$, where the coordinates from 1 to d_1 correspond to f_1, the coordinates from d_1+1 to d_1+d_2 correspond to f_2, etc.

3 Partial Shortest Path Decomposition

Abraham *et al.* [AFGN18] defined *shortest path decompositions* (SPD s) of "low depth". Every (weighted) path graph has an SPDdepth 1. A graph G has an SPDdepth k if there exist a *shortest path* P, such that every connected component in $G\backslash P$ has an SPDdepth $k-1$. In other words, given a graph, in SPD we hierarchically delete shortest paths from each connected component, until no vertices remain. In this paper we define a generalization called *partial shortest path decomposition* (PSPD), where we remove the requirement that all the vertices will be deleted. See the formal definition below. In Section 6 we will argue that every terminated planar graph with face cover of size γ has a PSPD of depth $O(\log\gamma)$ such that in each connected component in the lower level of the hierarchy, all the terminals lie on a single face.

A *partial partition* $\mathcal{X}$ of a set X, is a disjoint set of subsets of X. In other words, for every $A \in \mathcal{X}$, $A \subseteq X$, and for every different subsets $A, B \in \mathcal{X}$, $A \cap B = \emptyset$.

DEFINITION 3.1. (PSPD) *Given a weighted graph $G = (V,E,w)$, a PSPD of depth k is a pair $\{\mathcal{X},\mathcal{P}\}$, where $\mathcal{X}$ is a collection $\mathcal{X}_1,\dots,\mathcal{X}_{k+1}$ of partial partitions of V, $\mathcal{P}$ is a collection of sets of paths $\mathcal{P}_1,\dots,\mathcal{P}_k$, and:*

1. *$\mathcal{X}_1 = \{V\}$.*
2. *For every $1 \le i \le k$ and every cluster $X \in \mathcal{X}_i$, there exist a unique path $P_X \in \mathcal{P}_i$ such that P_X is a shortest path in $G[X]$.*
3. *For every $2 \le i \le k+1$, $\mathcal{X}_i$ consists of all connected components of $G[X \setminus P_X]$ over all $X \in \mathcal{X}_{i-1}$.*

The remainder of the PSPD $\{\mathcal{X},\mathcal{P}\}$ is a pair $\{\mathcal{C},\mathcal{B}\}$ where $\mathcal{C} = \mathcal{X}_{k+1}$ is the set of connected components in the final level of the PSPD , and $\mathcal{B} = \bigcup_{i=1}^{k} \cup\mathcal{P}_i$ is the set of all the vertices in the removed paths. $\mathcal{B}$ is also called the boundary.

Under Definition 3.1 SPD is a special case of PSPD where $\mathcal{C} = \emptyset$ (and $\mathcal{B} = V$). The main theorem in [AFGN18] states that if a graph G has SPD of depth k, then it is embeddable into ℓ_1 with distortion $O(\sqrt{k})$ [3]. [AFGN18] construct a different embedding for each level of the decomposition. Each such embedding is Lipschitz, while for every pair of vertices $u,v \in V$ there is some level i such that the embedding for this level has constant contraction w.r.t. u,v. Specifically, the level i with the bounded contraction guarantee is the first level in which either u,v are separated or the distance between $\{u,v\}$ to a deleted path is at most $\frac{d_G(u,v)}{12}$. In particular, given a PSPD , by using the exact same embedding from [AFGN18] (w.r.t. the existing levels in the decomposition), we get the following theorem.

THEOREM 3.1. (EMBEDDING USING PSPD) *Let $G = (V,E,w)$ be a weighted graph, and let $\{\mathcal{X},\mathcal{P}\}$ be a PSPD of depth k with remainder $\{\mathcal{C},\mathcal{B}\}$. There is an embedding $f : V \to \ell_1$ with the following properties:*

1. *For every $u,v \in V$, $\|f(v)-f(u)\|_1 \le O(\sqrt{k}) \cdot d_G(u,v)$.*
2. *For every $u,v \in V$ which are either separated by $\mathcal{C}$ (that is u,v do not belong to the same cluster in $\mathcal{C}$), or such that $\min\{d_G(v,\mathcal{B}), d_G(u,\mathcal{B})\} \le \frac{d_G(u,v)}{12}$, it holds that $\|f(v)-f(u)\|_1 \ge d_G(u,v)$.*

4 Uniformly Truncated Embedding

In this section we construct a uniformly truncated embedding for O-S graphs into ℓ_1. Specifically, given a truncation parameter t, we show how to embed O-S graphs into ℓ_1 via a Lipschitz map such that the norm of all the vectors is exactly t, and it is non-contractive for terminals at distance at most t. We will use two previous results on stochastic embeddings. The following theorem was proven by Englert *et al.* [EGK+14] (Thm. 12) in a broader sense. Lee *et al.* [LMM15] (Thm. 4.4) observed that it implies embedding of O-S graphs into outerplanar graphs.

THEOREM 4.1. *Consider a weighted planar graph $G = (V,E,w)$ with $F \subseteq V$ being a face. There is a stochastic embedding of F into dominating outerplanar graphs with expected distortion $O(1)$.*

The following theorem was proven by Gupta *et al.* [GNRS04] (Thm. 5.4).

[3] In fact [AFGN18] proved a more general result, stating that G is embeddable into ℓ_p for $p \in [1,\infty]$, with distortion $O(k^{\min\{\frac{1}{2},\frac{1}{p}\}})$. Similarly, in Theorem 3.1 can replace ℓ_1 with ℓ_p and the expansion $\sqrt{k}$ with $k^{\min\{\frac{1}{2},\frac{1}{p}\}}$. The contraction condition and values remains the same.

THEOREM 4.2. *Consider a weighted outerplanar graph $G = (V, E, w)$. There is a stochastic embedding of G into dominating trees with expected distortion $O(1)$.*

As it was already observed in [KLR19], we conclude (the proof is to Appendix A):

COROLLARY 4.1. *Consider a planar graph $G = (V, E, w)$ with a face F. There is a stochastic embedding of F into dominating trees with expected distortion $O(1)$.*

A first step towards truncated embedding of O-S graphs will be a truncated embedding of trees.

LEMMA 4.1. *Let $T = (V, E, w)$ be some tree and let $t > 0$ be a truncation parameter. There exists an embedding $f : T \to \ell_1$ such that the following holds:*

1. *Sphere surface: for every $v \in V$, $\|f(v)\|_1 = t$.*
2. *Lipschitz: for every $u, v \in V$,* $\|f(v) - f(u)\|_1 \le 4 \cdot d_T(v, u)$.
3. *Bounded contraction: for every $u, v \in V$, $\|f(v) - f(u)\|_1 \ge \min\{d_T(v, u), t\}$.*

Proof. Add a new vertex v_t to T with edges of weight $\frac{t}{2}$ to all the other vertices. Call the new graph T_t. Notice that T_t has treewidth 2. According to Chakrabarti *et al.* [CJLV08], there is an embedding f_{tw} of T_t into ℓ_1 with distortion 2. By rescaling, we can assume that the contraction is 1 (and the expansion is at most 2). Additionally, by shifting, we can assume that $f_{\text{tw}}(v_t) = \vec{0}$. Let f be the embedding f_{tw} with an additional coordinate. The value of every $v \in V$ in the new coordinate equals to $t - \|f_{\text{tw}}(v)\|_1$. We argue that f has the desired properties.

The first property follows as for every vertex $v \in V$, the distance in T to v_t is exactly $\frac{t}{2}$, therefore $\|f_{\text{tw}}(v)\|_1 = \|f_{\text{tw}}(v) - f_{\text{tw}}(v_t)\|_1 \le 2 \cdot d_{T_t}(v, v_t) = t$. Therefor $\|f(v)\|_1 = \|f_{\text{tw}}(v)\|_1 + |t - \|f_{\text{tw}}(v)\|_1| = t$.

The second property follows as f_{tw} has expansion 2, and distances in T_t can only decrease w.r.t. distances in T. Thus for every $u, v \in V$, $\|f_{\text{tw}}(v) - f_{\text{tw}}(u)\|_1 \le 2 \cdot d_{T_t}(v, u) \le 2 \cdot d_T(v, u)$. By the triangle inequality,

$$\begin{aligned}\|f(v) - f(u)\|_1 &= \|f_{\text{tw}}(v) - f_{\text{tw}}(u)\|_1 \\ &\quad + |(t - \|f_{\text{tw}}(v)\|_1) - (t - \|f_{\text{tw}}(u)\|_1)| \\ &\le 2 \cdot \|f_{\text{tw}}(v) - f_{\text{tw}}(u)\| \\ &\le 4 \cdot d_T(v, u) \ .\end{aligned}$$

For the third property, consider some pair $u, v \in V$. As every shortest path containing the new vertex v_t will be of weight at least t, it holds that $d_{T_t}(v, u) = \min\{d_T(v, u), t\}$. We conclude that $\|f(v) - f(u)\|_1 \ge \|f_{\text{tw}}(v) - f_{\text{tw}}(u)\|_1 \ge d_{T_t}(v, u) = \min\{d_T(v, u), t\}$. ∎

Next, we construct an embedding of O-S graphs into the sphere of radius t in ℓ_1.

COROLLARY 4.2. *Let $G = (V, E, w)$ be a planar graph, F a face, and $t > 0$ a truncation parameter. There exists an embedding $f : F \to \ell_1$ such that the following holds:*

1. *Sphere surface: for every $v \in F$, $\|f(v)\|_1 = t$.*
2. *Lipschitz: for every $u, v \in F$, $\|f(v) - f(u)\|_1 \le O(d(v, u))$.*
3. *Bounded Contraction: for every $u, v \in F$, $\|f(v) - f(u)\|_1 \ge \min\{d_G(v, u), t\}$.*

Proof. Let $\mathcal{D}$ be the distribution over dominating trees guaranteed in Corollary 4.1. For every $T \in \text{supp}(\mathcal{D})$, let f_T be the embedding of T into ℓ_1 from Lemma 4.1 with parameter t. Our embedding is constructed by concatenating all f_T, scaled by their probabilities. That is, $f = \oplus \{\Pr[T] \cdot f_T \mid T \in \text{supp}(\mathcal{D})\}$.

The first property follows as for every $v \in X$ and f_T, $\|f(v)\|_1 = t$. Similarly, the third property follows as for every $u, v \in F$ and $T \in \text{supp}(\mathcal{D})$, $\|f_T(v) - f_T(u)\|_1 \ge \min\{d_T(v, u), t\} \ge \min\{d(v, u), t\}$.

The second property follows as for every $v, u \in V$,

$$\begin{aligned}\|f(v) - f(u)\|_1 &= \sum_T \Pr[T] \cdot \|f_T(v) - f_T(u)\|_1 \\ &\le \sum_T \Pr[T] \cdot 4 \cdot d_T(v, u) \\ &= 4 \cdot \mathbb{E}_{T \sim \mathcal{D}}[d_T(v, u)] \\ &= O(d_G(v, u)) \ .\end{aligned}$$

∎

REMARK 4.2. *Efficient construction: the support of the distribution $\mathcal{D}$ might be of exponential size. Nevertheless, we can bypass this barrier by carefully sampling polynomially many trees.*
Denote by m the number of vertices on F. For a pair v, u, by Markov inequality, the probability that a sampled tree has distortion larger than m^3 on u, v is $O(m^{-3})$. We say that a tree is bad if it has distortion m^3 on some pair, otherwise it is good. By the union bound, the probability for sampling a bad tree is $O(1/m)$. Let $\mathcal{D}'$ be the distribution $\mathcal{D}$ restricted to good trees only. $\mathcal{D}'$ is a distribution over dominating trees with constant expected distortion, and worst case distortion m^3. Sample m^6 trees $T_1, \dots, T_{m^6}$ from $\mathcal{D}'$. By Hoeffding [4] and union inequalities, w.h.p. the average

[4] See https://sarielhp.org/misc/blog/15/09/03/chernoff.pdf, Theorem 7.4.3 .

distortion of all pairs will be constant. Define the embedding $f = \circ \left\{ m^{-6} \cdot f_{T_i} \right\}_{i=1}^{m^6}$. The proof above still goes through.

5 Non-Uniformly Truncated Embedding

In this section we generalize Corollary 4.2. Instead of a uniform truncation parameter t for all the vertices, we will allow a somewhat customized truncation .

LEMMA 5.1. *Let $G = (V, E, w)$ be a planar graph with a given drawing on the plane. Let $F, \mathcal{I}, \mathcal{B} \subseteq V$ such that $F \subseteq \mathcal{I}$, $\mathcal{I} \cup \mathcal{B} = V$, $\mathcal{I} \cap \mathcal{B} = \emptyset$, and F is a face in $G[\mathcal{I}]$. Then there is an embedding $f : F \to \ell_1$ such that the following holds:*

1. *For every $v \in F$, $\|f(v)\|_1 = d_G(v, \mathcal{B})$.*
2. *Lipschitz: for every $u, v \in F$,* $$\|f(v) - f(u)\|_1 \le O(d_G(v, u)).$$
3. *Bounded Contraction: for every $c \ge 1$, and every $u, v \in F$ such that $\min\{d_G(v, \mathcal{B}), d_G(u, \mathcal{B})\} \ge \frac{d_G(u,v)}{c}$ it holds that $\|f(v) - f(u)\|_1 \ge \frac{d_G(v,u)}{12 \cdot c}$.*

Proof. We will construct the non-uniformly truncated embedding by a smooth combination of uniformly truncated embeddings for all possible truncation scales. A similar approach was applied in [AFGN18]. Assume (by scaling) w.l.o.g. that the minimal weight of an edge in G is 1. Let $M \in \mathbb{N}$ be minimal integer such that the diameter of G is strictly bounded by 2^M.

Consider the graph $G[\mathcal{I}]$ induced by $\mathcal{I}$. Note that $G[\mathcal{I}]$ is an O-S graph w.r.t. F. For every distance scale $t \in \{0, 1, \dots, M\}$ let f_t be the embedding of F w.r.t. the shortest path metric induced by $G[\mathcal{I}]$ from Corollary 4.2 with truncation parameter 2^t. For a vertex $v \in \mathcal{I}$, let $t_v \in \mathbb{N}$ be such that $d_G(v, \mathcal{B}) \in [2^{t_v}, 2^{t_v+1})$. Set $\lambda_v = \frac{d_G(v,\mathcal{B}) - 2^{t_v}}{2^{t_v}}$. Note that $0 \le \lambda_v < 1$. For $t \in \{0, \dots, M\}$ we define a function $\tilde{f}_t : F \to \ell_1$,

$$\tilde{f}_t(v) \doteq \begin{cases} \lambda_v \cdot f_t(v) & \text{if } t = t_v + 1, \\ (1 - \lambda_v) \cdot f_t(v) & \text{if } t = t_v, \\ \vec{0} & \text{otherwise.} \end{cases}$$

Define f to be the concatenation of $\tilde{f}_0(v), \dots, \tilde{f}_M(v)$.

For every $v \in F$, according to Corollary 4.2 it holds that

(5.1)
$$\begin{aligned} &\|f(v)\|_1 \\ &= (1 - \lambda_v) \cdot \left\| \tilde{f}_{t_v}(v) \right\|_1 + \lambda_v \cdot \left\| \tilde{f}_{t_v+1}(v) \right\|_1 \\ &= (1 - \lambda_v) \cdot 2^{t_v} + \lambda_v \cdot 2^{t_v+1} \\ &= \frac{2^{t_v+1} - d_G(v, \mathcal{B})}{2^{t_v}} \cdot 2^{t_v} + \frac{d_G(v, \mathcal{B}) - 2^{t_v}}{2^{t_v}} \cdot 2^{t_v+1} \\ &= d_G(v, \mathcal{B}) \ . \end{aligned}$$

Next we prove that f is Lipschitz. Consider a pair of vertices $u, v \in \mathcal{F}$. If $d_G(u, v) < d_{G[\mathcal{I}]}(u, v)$, then the shortest path from u to v in G has to go through the boundary $\mathcal{B}$. It follows that $d_G(u, \mathcal{B}) + d_G(v, \mathcal{B}) \le d_G(u, v)$. We conclude

$$\begin{aligned} \|f(v) - f(u)\|_1 &\le \|f(v)\|_1 + \|f(u)\|_1 \\ &\overset{(5.1)}{=} d_G(v, \mathcal{B}) + d_G(u, \mathcal{B}) \\ &\le d_G(v, u) \ . \end{aligned}$$

Otherwise, $d_G(u, v) = d_{G[\mathcal{I}]}(u, v)$. It follows from Corollary 4.2 that for every scale parameter t, $\|f_t(v) - f_t(u)\| \le O\left(d_{G[\mathcal{I}]}(u, v)\right) = O\left(d_G(u, v)\right)$. We will prove a similar inequality for $\tilde{f}_t$. As $\tilde{f}_t(u)$ and $\tilde{f}_t(v)$ combined might be nonzero in at most 4 different scales, the bound on expansion will follow.

Denote by p_t the scaling factor of v in $\tilde{f}_t(v)$. That is, $p_{t_v+1} = \lambda_v$, $p_{t_v} = 1 - \lambda_v$, and $p_t = 0$ for $t \notin \{t_v, t_v + 1\}$. Similarly, define q_t for u. First, observe that for every t,

$$\begin{aligned} \left\| \tilde{f}_t(v) - \tilde{f}_t(u) \right\|_1 &= \|p_t \cdot f_t(v) - q_t \cdot f_t(u)\|_1 \\ &\le \min\{p_t, q_t\} \cdot \|f_t(v) - f_t(u)\|_1 \\ &\quad + |p_t - q_t| \cdot \max\{\|f_t(v)\|_1, \|f_t(u)\|_1\} \\ &\le O(d_G(u, v)) + |p_t - q_t| \cdot 2^t \ . \end{aligned}$$

It suffice to show that $|p_t - q_t| = O(d_G(u, v)/2^t)$. Indeed, for indices $t \notin \{t_u, t_u + 1, t_v, t_v + 1\}$, $p_t = q_t = 0$, and in particular $|p_t - q_t| = 0$. Let us consider the other cases. W.l.o.g. assume that $d_G(v, \mathcal{B}) \ge d_G(u, \mathcal{B})$ and hence $t_v \ge t_u$. We proceed by case analysis.

- $\boldsymbol{t_u = t_v}$ **:** In this case, $|p_{t_v} - q_{t_v}| = |(1 - \lambda_v) - (1 - \lambda_u)| = |\lambda_v - \lambda_u| = |p_{t_v+1} - q_{t_v+1}|$. The value of this quantity is bounded by

$$\begin{aligned} \lambda_v - \lambda_u &= \frac{d_G(v, \mathcal{B}) - 2^{t_v}}{2^{t_v}} - \frac{d_G(u, \mathcal{B}) - 2^{t_v}}{2^{t_v}} \\ &= \frac{d_G(v, \mathcal{B}) - d_G(u, \mathcal{B})}{2^{t_v}} \\ &\le \frac{d_G(u, v)}{2^{t_v}} \ . \end{aligned}$$

Hence, we get that $|p_t - q_t| = O(d_G(u,v)/2^t)$ for all $t \in \{t_v, t_v+1\}$.

- $\boldsymbol{t_u = t_v - 1}$: It holds that

$$\begin{aligned}&\lambda_v + (1-\lambda_u)\\ &\leq\ 2\cdot\frac{d_G(v,\mathcal{B})-2^{t_v}}{2^{t_v}}+\frac{2^{t_u+1}-d_G(u,\mathcal{B})}{2^{t_u}}\\ &=\ \frac{\left(d_G(v,\mathcal{B})-2^{t_v}\right)-\left(2^{t_u+1}-d_G(u,\mathcal{B})\right)}{2^{t_u}}\\ &\leq\ \frac{d_G(u,v)}{2^{t_u}}\ .\end{aligned}$$

We conclude:

$$\begin{aligned}|p_{t_v+1}-q_{t_v+1}| &= \lambda_v = O(d_G(u,v)/2^{t_v+1})\\ |p_{t_v}-q_{t_v}| &= |1-\lambda_v-\lambda_u| = O(d_G(u,v)/2^{t_v})\\ |p_{t_u}-q_{t_u}| &= 1-\lambda_u = O(d_G(u,v)/2^{t_u})\end{aligned}$$

- $\boldsymbol{t_u < t_v - 1}$: By the definition of t_v and t_u, $d_G(v,u) \geq d_G(v,\mathcal{B}) - d_G(u,\mathcal{B}) \geq 2^{t_v} - 2^{t_u+1} \geq 2^{t_v-1}$. It follows that for every $t \leq t_v+1$, $|p_t - q_t| \leq 1 \leq \frac{d_G(u,v)}{2^{t_v-1}} = O\left(\frac{d_G(u,v)}{2^t}\right)$.

Next we argue that f has small contraction for pairs far enough form the boundary. Consider a pair of vertices $u,v \in F$, and let $c \geq 1$ such that $\min\{d_G(v,\mathcal{B}), d_G(u,\mathcal{B})\} \geq \frac{d_G(u,v)}{c}$. It holds that $2^{t_v} > \frac{1}{2}\cdot d_G(v,\mathcal{B}) \geq \frac{1}{2c}\cdot d_G(u,v)$. For every $t \geq t_v$, by the contraction property of Corollary 4.2, it holds that

(5.2)
$$\|f_t(v) - f_t(u)\|_1 \geq \min\left\{d_{G[\mathcal{I}]}(v,u), 2^t\right\} \geq \frac{d_G(v,u)}{2c}\ ,$$

where the last inequality follows as $c \geq 1$. Assume w.l.o.g. that $d_G(v,\mathcal{B}) \geq d_G(u,\mathcal{B})$, thus $t_v \geq t_u$ and let $t \in \{t_v, t_v+1\}$ such that $p_t \geq \frac{1}{2}$. Set $S = 12\cdot c$. We consider two cases:

- If $|p_t - q_t| > \frac{d_G(v,u)}{2^t\cdot S}$ then

$$\begin{aligned}\|f(v)-f(u)\|_1 &\geq \|p_t\cdot f_t(v) - q_t\cdot f_t(u)\|_1\\ &\geq \left|\|p_t\cdot f_t(v)\|_1 - \|q_t\cdot f_t(u)\|_1\right|\\ &= |p_t-q_t|\cdot 2^t > \frac{d_G(v,u)}{S}\ .\end{aligned}$$

- Else $|p_t - q_t| \leq \frac{d_G(v,u)}{2^t\cdot S}$. First, assume that $p_t \geq q_t$. As $\frac{d_G(v,u)}{2^t\cdot S} \leq \frac{2c}{S} = \frac{1}{6}$, it holds that $q_t \geq \frac{1}{3}$. We conclude

$$\begin{aligned}&\|f(v)-f(u)\|_1\\ &\geq\ \|p_t\cdot f_t(v) - q_t\cdot f_t(u)\|_1\\ &\geq\ q_t\cdot\|f_t(v)-f_t(u)\|_1 - |p_t-q_t|\cdot\|f_t(v)\|\\ &\overset{(5.2)}{\geq}\ \frac{1}{3}\cdot\frac{d_G(v,u)}{2c} - \frac{d_G(v,u)}{S}\ =\ \frac{d_G(v,u)}{S}\ .\end{aligned}$$

The case where $q_t > p_t$ is symmetric.

∎

REMARK 5.2. *In Lemma 5.1 we used uniformly truncated embeddings in order to create a non-uniformly truncated embedding. Such a transformation might be relevant in other contexts as well. Consider a case where each vertex v has some truncation parameter s_v, and there exists a uniformly truncated embedding for every parameter t. As long as for every v,u, $|s_v - s_u| = O(d_G(u,v))$, following the same construction as above, one can create a similar non-uniformly truncated embedding where $\|f(v)\| = s_v$.*

6 Embedding Parametrized by Face Cover: Proof of Theorem 1.1

We will use the following separator theorem in order to create a PSPD .

THEOREM 6.1. *Let $G = (V,E,w,K)$ be a weighted terminated planar graph. Suppose that $\gamma(G,K) = \gamma$. Then there are two shortest paths P_1, P_2 in G, such that for every connected component C in $G\backslash\{P_1\cup P_2\}$ it holds $\gamma(G[C], K\cap C) \leq \frac{2}{3}\gamma + 1$.*

We defer the proof of Theorem 6.1 to Appendix B. If the size of the face cover is 2, we can reduce this size to 1 by removing a single shortest path containing vertices from both faces, the remaining graph will be O-S. Similarly, if the size of the face cover is 3 we can reduce to 1 by removing a pair of shortest paths. We can invoke Theorem 6.1 repeatedly in order to hierarchically partition G, reducing the size of the face cover in each iteration. After $O(\log\gamma)$ iterations, the size of the face cover in each connected component will be at most 1.

COROLLARY 6.1. *Let $G = (V,E,w,K)$ be a weighted terminated planar graph such that $\gamma(G,K) = \gamma$. Then there is an PSPD $\{\mathcal{X},\mathcal{P}\}$ of depth $O(\log\gamma)$ and remainder $\{\mathcal{C},\mathcal{B}\}$, such that for every cluster $C \in \mathcal{C}$, all the terminals in C lie on a single face (i.e. $\forall C \in \mathcal{C}$, $\gamma(C, K\cap C) \leq 1$).*

Given the PSPD $\{\mathcal{X},\mathcal{P}\}$ above, with remainder $\{\mathcal{C},\mathcal{B}\}$, we are ready to define the embedding of Theorem 1.1. Let f_{PSPD} be the embedding of G into ℓ_1 from Theorem 3.1, restricted to K. For every cluster $C \in \mathcal{C}$ let $\mathcal{B}_C = V\setminus C$ and F_C be the outer face of C. Let f_C be the embedding from Lemma 5.1 with parameters $F_C, C, \mathcal{B}_C$. Let $\hat{f}_C$ be the embedding f_C restricted to $K\cap C$, and extended to K by sending every $v \in K\setminus C$ to $\vec{0}$. The final embedding f will be a concatenation of f_{PSPD} with $\hat{f}_C$ for all $C \in \mathcal{C}$.

Expansion Consider a pair of vertices $v,u \in V$. By Theorem 3.1, $\|f_{\mathsf{PSPD}}(v) - f_{\mathsf{PSPD}}(u)\|_1 = O\left(\sqrt{\log\gamma}\right)\cdot d_G(v,u)$. On the other hand, for every $C \in \mathcal{C}$, using Lemma 5.1, if $u,v \in C$ then

$\|f_C(v) - f_C(u)\|_1 = O(1) \cdot d_G(v,u)$. Otherwise if $v \in C, u \notin C$ $\|f_C(v) - f_C(u)\| = d_G(v,\mathcal{B}) \leq d_G(v,u)$ (similarly for $v \notin C, u \in C$). As each vertex is nonzero only in a single function f_C, the $O(\sqrt{\log \gamma})$ bound on the expansion follows.

Contraction Consider a pair of terminal vertices v,u. If either u,v are separated by $\mathcal{C}$ or $\min\{d_G(v,\mathcal{B}), d_G(u,\mathcal{B})\} \leq \frac{d_G(u,v)}{12}$, it holds that $\|f_{\mathsf{PSPD}}(v) - f_{\mathsf{PSPD}}(u)\|_1 \geq d_G(u,v)$ and we are done. Otherwise, there must exist a cluster $C \in \mathcal{C}$ such that $u,v \in F_C$ and $\min\{d_G(v,\mathcal{B}_C), d_G(u,\mathcal{B}_C)\} \geq \frac{d_G(u,v)}{12}$. By Lemma 5.1 $\left\|\hat{f}_C(v) - \hat{f}_C(u)\right\|_1 = \|f_C(v) - f_C(u)\|_1 \geq \frac{d_G(u,v)}{12 \cdot 12}$.

Theorem 1.1 now follows. Bellow we discuss the implementation details of the embedding.

6.1 Polynomial Implementation Given a planar graph with a drawing in the plane, using the PTAS of Frederickson [Fre91] we can find a face cover of size $2 \cdot \gamma(G,K)$ in linear time (see Section 1.2). Note that using a cover of size $2 \cdot \gamma(G,K)$ instead of $\gamma(G,K)$ is insignificant for our $O(\sqrt{\log(\gamma(G,K))})$ upper bound. Next, construct a PSPD for this face cover using cycle separators. Since we construct at most n separators, the construction of the PSPD also takes polynomial time. Given a PSPD , the embedding of [AFGN18] is efficiently computed.

After the creation of f_{PSPD} we are left with a remainder $\{\mathcal{C},\mathcal{B}\}$. For every $C \in \mathcal{C}$ we start by computing uniformly truncated embeddings (see Remark 4.2). Given an O-S graph with k terminals, there are at most $2k$ truncation scales (as each terminal participates in two scales only). Thus in polynomial time we can compute the embedding for all truncation scales, and thus compute the non-uniformly truncated embedding.

Acknowledgments

The author would like to thank Ofer Neiman for helpful discussions.

References

[ABN11] Ittai Abraham, Yair Bartal, and Ofer Neiman. Advances in metric embedding theory. *Advances in Mathematics*, 228(6):3026 – 3126, 2011. 3

[AFGN18] Ittai Abraham, Arnold Filtser, Anupam Gupta, and Ofer Neiman. Metric embedding via shortest path decompositions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 952–963, 2018. 2, 3, 4, 6, 8

[AGG+14] Ittai Abraham, Cyril Gavoille, Anupam Gupta, Ofer Neiman, and Kunal Talwar. Cops, robbers, and threatening skeletons: padded decomposition for minor-free graphs. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 79–88, 2014. 3

[Ber90] Marshall W. Bern. Faster exact algorithms for steiner trees in planar networks. *Networks*, 20(1):109–120, 1990. 3

[BFN16] Yair Bartal, Arnold Filtser, and Ofer Neiman. On notions of distortion and an almost minimum spanning tree with constant average distortion. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 873–882, 2016. The embedding of β-decomposable metrics appears in the full version http://arxiv.org/abs/1609.08801. 3

[BM88] Daniel Bienstock and Clyde L. Monma. On the complexity of covering vertices by faces in a planar graph. *SIAM J. Comput.*, 17(1):53–76, 1988. 3

[CGN+06] Chandra Chekuri, Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Embedding k-outerplanar graphs into ℓ_1. *SIAM J. Discrete Math.*, 20(1):119–136, 2006. 3

[CJLV08] Amit Chakrabarti, Alexander Jaffe, James R. Lee, and Justin Vincent. Embeddings of topological graphs: Lossy invariants, linearization, and 2-sums. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 761–770, 2008. 3, 5

[CSW13] Chandra Chekuri, F. Bruce Shepherd, and Christophe Weibel. Flow-cut gaps for integer and fractional multiflows. *J. Comb. Theory, Ser. B*, 103(2):248–273, 2013. 2

[CW04] Danny Z. Chen and Xiaodong Wu. Efficient algorithms for k-terminal cuts on planar graphs. *Algorithmica*, 38(2):299–316, 2004. 3

[CX00] Danny Z. Chen and Jinhui Xu. Shortest path queries in planar graphs. In *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing*, STOC '00, pages 469–478, New York, NY, USA, 2000. ACM. 3

[EGK+14] Matthias Englert, Anupam Gupta, Robert Krauthgamer, Harald Räcke, Inbal Talgam-Cohen, and Kunal Talwar. Vertex sparsifiers: New results from old techniques. *SIAM J. Comput.*, 43(4):1239–1262, 2014. 4

[EMJ87] Ranel E. Erickson, Clyde L. Monma, and Arthur F. Veinott Jr. Send-and-split method for minimum-concave-cost network flows. *Math. Oper. Res.*, 12(4):634–664, 1987. 3

[Fre91] Greg N. Frederickson. Planar graph decomposition and all pairs shortest paths. *J. ACM*, 38(1):162–204, 1991. 3, 8

[Fre95] Greg N. Frederickson. Using cellular graph embeddings in solving all pairs shortest paths problems. *J. Algorithms*, 19(1):45–85, 1995. 3

[GNRS04] Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Cuts, trees and ℓ_1s-embeddings of graphs. *Combinatorica*, 24(2):233–269, 2004. 1, 2, 3, 4

[HST86] Cor A. J. Hurkens, Alexander Schrijver, and Éva Tardos. On fractional multicommodity flows and distance functions. *Discrete Mathematics*, 73:99–109, 1986. 2

[KLMN05] Robert Krauthgamer, James R. Lee, Manor Mendel, and Assaf Naor. Measured descent: a new embedding method for finite metrics. *Geometric and Functional Analysis*, 15(4):839–858, 2005. 3

[KLR19] Robert Krauthgamer, James R. Lee, and Havana Rika. Flow-cut gaps and face covers in planar graphs. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 525–534, 2019. 2, 5

[KNvL19] Sándor Kisfaludi-Bak, Jesper Nederlof, and Erik Jan van Leeuwen. Nearly eth-tight algorithms for planar steiner tree with terminals on few faces. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1015–1034, 2019. 3

[KPZ17] Nikolai Karpov, Marcin Pilipczuk, and Anna Zych-Pawlewicz. An exponential lower bound for cut sparsifiers in planar graphs. *CoRR*, abs/1706.06086, 2017. 3

[KR17] Robert Krauthgamer and Inbal Rika. Refined vertex sparsifiers of planar graphs. *CoRR*, abs/1702.05951, 2017. 3

[LLR95] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995. 1, 2

[LMM15] James R Lee, Manor Mendel, and Mohammad Moharrami. A node-capacitated okamura–seymour theorem. *Mathematical Programming*, 153(2):381–415, 2015. 4

[LR10] James R. Lee and Prasad Raghavendra. Coarse differentiation and multi-flows in planar graphs. *Discrete & Computational Geometry*, 43(2):346–362, 2010. 1, 3

[LS09] James R. Lee and Anastasios Sidiropoulos. On the geometry of graphs with a forbidden minor. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 245–254, 2009. 2

[LS10] James R. Lee and Anastasios Sidiropoulos. Genus and the geometry of the cut graph. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 193–201, 2010. 3

[LS13] James R. Lee and Anastasios Sidiropoulos. Pathwidth, trees, and random embeddings. *Combinatorica*, 33(3):349–374, 2013. 3

[Mat02] Jiri Matoušek. *Lectures on discrete geometry*. Springer-Verlag, New York, 2002. 3

[Mil86] Gary L. Miller. Finding small simple cycle separators for 2-connected planar graphs. *J. Comput. Syst. Sci.*, 32(3):265–279, 1986. 2, 10

[MNS85] Kazuhiko Matsumoto, Takao Nishizeki, and Nobuji Saito. An efficient algorithm for finding multicommodity flows in planar networks. *SIAM J. Comput.*, 14(2):289–302, 1985. 3

[OS81] Haruko Okamura and P.D. Seymour. Multicommodity flows in planar graphs. *Journal of Combinatorial Theory, Series B*, 31(1):75 – 81, 1981. 2

[Rao99] Satish Rao. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In *Proceedings of the Fifteenth Annual Symposium on Computational Geometry, Miami Beach, Florida, USA, June 13-16, 1999*, pages 300–306, 1999. 1, 2

[Tho04] Mikkel Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *J. ACM*, 51(6):993–1024, November 2004. 10

A O-S into trees: Proof of Corollary 4.1

Corollary 4.1. *Consider a planar graph $G = (V, E, w)$ with a face F. There is a stochastic embedding of F into dominating trees with expected distortion $O(1)$.*

Proof. Let $\mathcal{D}_{OP}$ be the distribution over outerplanar graphs from Theorem 4.1. For every $G' \in \text{supp}(\mathcal{D}_{OP})$, let $\mathcal{D}_{G'}$ be the distribution over trees from Theorem 4.2 (w.r.t. G'). We define a distribution $\mathcal{D}$ of embeddings of F into trees as follows. First sample an outerplanar graph G' using $\mathcal{D}_{OP}$. Then sample a tree T using $\mathcal{D}_{G'}$. We argue that the distribution $\mathcal{D}$ has the desired properties.

All the embeddings in the support of $\mathcal{D}$ are dominating metrics as both $\mathcal{D}_{OP}$ and $\mathcal{D}_{G'}$ (for all G') have this property. Consider $v, u \in V$, it holds that

$$
\begin{aligned}
&\mathbb{E}_{T'\sim\mathcal{D}}\left[d_T(v,u)\right] \\
&= \sum_{G'} \Pr_{\mathcal{D}_{OP}}\left[G'\right]\cdot\left(\sum_{T}\Pr_{\mathcal{D}_G}\left[T \mid G'\right]\cdot d_T(v,u)\right) \\
&= \sum_{G} \Pr_{\mathcal{D}_{OP}}\left[G'\right]\cdot O\left(d_{G'}(v,u)\right) \quad = \quad O\left(d_G(v,u)\right) .
\end{aligned}
$$

□

B Faces Separator: Proof of Theorem 6.1

Theorem B.1. *Let $G = (V, E, w, K)$ be a weighted terminated planar graph. Suppose that $\gamma(G, K) = \gamma$. Then there are two shortest paths P_1, P_2 in G, such that for every connected component C in $G \backslash \{P_1 \cup P_2\}$ it holds $\gamma(G[C], K \cap C) \leq \frac{2}{3}\gamma + 1$.*

Proof. Even though similar statements to Theorem 6.1 already appeared in the literature, we provide a proof for completeness. A planar cycle separator theorem for vertices have already became folklore [Mil86, Tho04]. Specifically, given a weight function $\omega : V \to \mathbb{R}_+$ over the vertices, and a root vertex $v \in V$, one can efficiently find a cycle S, that consists of two shortest paths rooted at v, such that the total weight of the vertices in each connected component of $G[V \setminus S]$ is at most $\frac{2}{3}\sum_{v\in V}\omega(v)$.

We start by defining a weight function ω. Let $\mathcal{F}$ be a face cover of size γ. For every face $F \in \mathcal{F}$, let $v_F \in F$ be an arbitrary vertex (not necessarily unique). Initially the weight of all the vertices is 0. For every $F \in \mathcal{F}$, add a single unit of weight to v_F. Note that the total weight of all the vertices is γ, while for every $F \in \mathcal{F}$, the total weight of the vertices in F is at least 1. See Figure 2 for an illustration.

Let v be an arbitrary vertex on the outer face. We use the planar cycle separator theorem w.r.t. the weight function ω and the root vertex v. As a result, we get a pair of shortest paths P_1, P_2 rooted in v. Let C be a connected component in $G\setminus\{P_1 \cup P_2\}$. The total weight of all the vertices in C is bounded by $\frac{2}{3}\gamma$. Consider the drawing of C obtained by removing all other vertices from the drawing of G. Next we define a face cover $\mathcal{F}_C$. For every $F \in \mathcal{F}$, if $v_F \in C$ then add F to $\mathcal{F}_C$ (or the new face containing the remainder of F). Additionally, add the outer face in the drawing of C to $\mathcal{F}_C$.

It is straightforward that $|\mathcal{F}_C| \le \frac{2}{3}\gamma + 1$. We argue that $\mathcal{F}_C$ is a face cover for $K \cap C$. Indeed, let $u \in K \cap C$. Let $F \in \mathcal{F}$ be some face s.t. $u \in F$. If $v_F \in C$, then F (or its remainder) is in $\mathcal{F}_C$, and therefore u is covered. Otherwise, $v_F \notin C$. Therefore v_F and u were separated by the deletion of P_1, P_2. Necessarily some vertex of F belongs to $P_1 \cup P_2$. We conclude that u is now part of the outer face, and therefore covered. □

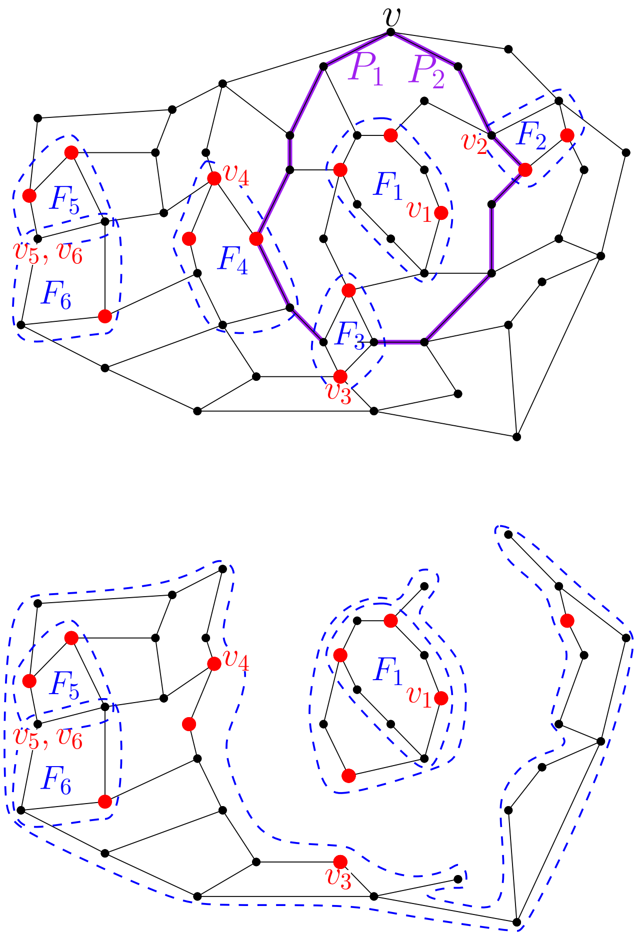

Figure 2: *On the top displayed a graph G. The terminals are colored red. The face cover consist of the faces $F_1, \dots, F_6$, surrounded by blue dashed lines. For each face F_i let v_{F_i} (denoted v_i) be some vertex on F_i. Define a weight function ω by adding a unit of weight to every v_i. The separator consists of shortest paths P_1, P_2 colored purple.*
On the bottom displayed the graph after removing all separator vertices. In each connected component C, a new face cover is defined by taking the outer face and adding a single face for every $v_i \in C$.