

Flow-cut gaps and face covers in planar graphs

Robert Krauthgamer*[†]
Weizmann Institute

James R. Lee[‡]
University of Washington

Havana (Inbal) Rika[§]
Weizmann Institute

Abstract

The relationship between the sparsest cut and the maximum concurrent multi-flow in graphs has been studied extensively. For general graphs, the worst-case gap between these two quantities is now settled: When there are k terminal pairs, the flow-cut gap is $O(\log k)$, and this is tight. But when topological restrictions are placed on the flow network, the situation is far less clear. In particular, it has been conjectured that the flow-cut gap in planar networks is $O(1)$, while the known bounds place the gap somewhere between 2 (Lee and Raghavendra, 2003) and $O(\sqrt{\log k})$ (Rao, 1999).

A seminal result of Okamura and Seymour (1981) shows that when all the terminals of a planar network lie on a single face, the flow-cut gap is exactly 1. This setting can be generalized by considering planar networks where the terminals lie on one of $\gamma > 1$ faces in some fixed planar drawing. Lee and Sidiropoulos (2009) proved that the flow-cut gap is bounded by a function of γ , and Chekuri, Shepherd, and Weibel (2013) showed that the gap is at most 3γ . We significantly improve these asymptotics by establishing that the flow-cut gap is $O(\log \gamma)$. This is achieved by showing that the edge-weighted shortest-path metric induced on the terminals admits a stochastic embedding into trees with distortion $O(\log \gamma)$. The latter result is tight, e.g., for a square planar lattice on $\Theta(\gamma)$ vertices.

The preceding results refer to the setting of *edge-capacitated* networks. For *vertex-capacitated networks*, it can be significantly more challenging to control flow-cut gaps. While there is no exact vertex-capacitated version of the Okamura-Seymour Theorem, an approximate version holds; Lee, Mendel, and Moharrami (2015) showed that the vertex-capacitated flow-cut gap is $O(1)$ on planar networks whose terminals lie on a single face. We prove that the flow-cut gap is $O(\gamma)$ for vertex-capacitated instances when the terminals lie on at most γ faces. In fact, this result holds in the more general setting of submodular vertex capacities.

*Weizmann Institute of Science, Israel. Work partially supported by ONR Award N00014-18-1-2364, the Israel Science Foundation grant #1086/18, and a Minerva Foundation grant. Email: robert.krauthgamer@weizmann.ac.il

[†]Part of this work was done while the author was visiting the Simons Institute for the Theory of Computing.

[‡]University of Washington, Seattle, WA. Work supported by NSF grants CCF-1616297 and CCF-1407779 and a Simons Investigator Award. Email: jrl@cs.washington.edu

[§]Weizmann Institute of Science, Israel. Email: havana.rika@weizmann.ac.il

Contents

1	Introduction	2
1.1	The flow-cut gap	3
1.2	The vertex-capacitated flow-cut gap	3
1.3	Stochastic embeddings	4
2	Approximation by random trees	5
2.1	Random partitions, embeddings, and peeling	5
2.2	Approximation by OS-instances	7
2.3	From OS-instances to random trees	10
3	Polymatroid flow-cut gaps	10
3.1	Embeddings into thin trees	11

1 Introduction

We present some new upper bounds on the gap between the concurrent flow and sparsest cut in planar graphs in terms of the topology of the terminal set. Our proof employs low-distortion metric embeddings into ℓ_1 , which are known to have a tight connection to the flow-cut gap (see, e.g., [LLR95, GNRS04]). We now review the relevant terminology.

Consider an undirected graph G equipped with nonnegative edge lengths $\ell : E(G) \rightarrow \mathbb{R}_+$ and a subset $\mathsf{T} = \mathsf{T}(G) \subseteq V(G)$ of *terminal vertices*. We use $d_{G,\ell}$ to denote the shortest-path distance in G , where the length of paths is computed using the edge lengths ℓ . We use $c_1^+(G, \ell; \mathsf{T})$ to denote the minimal number $D \geq 1$ for which there exists 1-Lipschitz mapping $F : V(G) \rightarrow \ell_1$ such that $F|_{\mathsf{T}(G)}$ has bilipschitz distortion D . In other words,

$$\forall u, v \in V(G) : \quad \|f(u) - f(v)\|_1 \leq d_{G,\ell}(u, v), \quad (1)$$

$$\forall s, t \in \mathsf{T}(G) : \quad \|f(s) - f(t)\|_1 \geq \frac{1}{D} \cdot d_{G,\ell}(s, t). \quad (2)$$

For an undirected graph G , we define $c_1^+(G; \mathsf{T}) := \sup_{\ell} c_1^+(G, \ell; \mathsf{T})$, where ℓ ranges over all nonnegative lengths $\ell : E(G) \rightarrow \mathbb{R}_+$. We also define $c_1^+(G, \ell) := c_1^+(G, \ell; V(G))$ and $c_1^+(G) := c_1^+(G; V(G))$. Finally, for a family \mathcal{F} of finite graphs, we denote $c_1^+(\mathcal{F}) := \sup\{c_1^+(G) : G \in \mathcal{F}\}$, and for $k \in \mathbb{N}$, we denote

$$c_1^+(\mathcal{F}; k) := \sup\{c_1^+(G; \mathsf{T}) : G \in \mathcal{F}, \mathsf{T} \subseteq V(G), |\mathsf{T}| = k\}.$$

Let \mathcal{F}_{fin} denote the family of all finite graphs, and $\mathcal{F}_{\text{plan}}$ the family of all planar graphs. It is known that $c_1^+(\mathcal{F}_{\text{fin}}; k) = \Theta(\log k)$ [AR98, LLR95] for all $k \geq 1$. For planar graphs, one has $c_1^+(\mathcal{F}_{\text{plan}}; k) \leq O(\sqrt{\log k})$ [Rao99] and $c_1^+(\mathcal{F}_{\text{plan}}) \geq 2$ [LR10].

Fix a plane graph G (this is a planar graph G together with a drawing in the plane). For $\mathsf{T} \subseteq V(G)$, we define the quantity $\gamma(G; \mathsf{T})$ to be the smallest number of faces in G that together cover all the vertices of T , and $\gamma(G) := \gamma(G; V(G))$. We say that the pair (G, T) is an *Okamura-Seymour instance*, or in short an *OS-instance*, if it can be drawn in the plane with all its terminal on the same face, i.e., if there is a planar representation for which $\gamma(G; \mathsf{T}) = 1$. A seminal result of Okamura and Seymour [OS81] implies that $c_1^+(G; \mathsf{T}) = 1$ whenever (G, T) is an OS-instance.

The methods of [LS09] show that $c_1^+(G; \mathsf{T}) \leq 2^{O(\gamma(G; \mathsf{T}))}$, and the authors of [CSW13, Theorem 4.13] give a direct proof that $c_1^+(G; \mathsf{T}) \leq 3\gamma(G; \mathsf{T})$. Our main result is the following improvement.

Theorem 1.1. *For every plane graph G and terminal set $\mathsf{T} \subseteq V(G)$, it holds that*

$$c_1^+(G; \mathsf{T}) \leq O(\log \gamma(G; \mathsf{T})).$$

A long-standing conjecture [GNRS04] asks whether $c_1^+(\mathcal{F}) < \infty$ for every family \mathcal{F} of finite graphs that is closed under taking minors, and such that \mathcal{F} does not contain all finite graphs. If true, this conjecture would of course imply that one can replace the bound of Theorem 1.1 with a universal constant.

It is known that a plane graph G has treewidth $O(\sqrt{\gamma(G)})$ [KLL02]. If we use $\mathcal{F}_{\text{tw}}(w)$ and $\mathcal{F}_{\text{pw}}(w)$ to denote the families of graphs of treewidth w and pathwidth w , respectively, then it is known that $c_1^+(\mathcal{F}_{\text{tw}}(2))$ is finite [GNRS04], but this remains open for $c_1^+(\mathcal{F}_{\text{tw}}(3))$. (On the other hand, $c_1^+(\mathcal{F}_{\text{pw}}(w)) < \infty$ for every $w \geq 1$ [LS13], and the best-known quantitative bound is $c_1^+(\mathcal{F}_{\text{pw}}(w)) \leq O(\sqrt{w})$ [AFGN18].)

1.1 The flow-cut gap

We now define the flow-cut gap, and briefly explain its connection to c_1^+ . Consider an undirected graph G with terminals $\mathbb{T} = \mathbb{T}(G)$. Let $c : E(G) \rightarrow \mathbb{R}_+$ denote an assignment of *capacities* to edges, and $d : \binom{\mathbb{T}}{2} \rightarrow \mathbb{R}_+$ an assignment of *demands*. The triple (G, c, d) is called an (undirected) *network*. The *concurrent flow* value of the network is the maximum value $\lambda > 0$, such that $\lambda \cdot d(\{s, t\})$ units of flow can be routed between every demand pair $\{s, t\} \in \binom{\mathbb{T}}{2}$, simultaneously but as separate commodities, without exceeding edge capacities.

Given the network (G, c, d) and a subset $S \subset V$, let $\text{cap}(S)$ be denote the total capacity of edges crossing the cut $(S, V \setminus S)$, and let $\text{dem}(S)$ denote the sum of $d(\{s, t\})$ over all pairs $\{s, t\} \in \binom{\mathbb{T}}{2}$ that cross the same cut. The *sparsity of a cut* $(S, V \setminus S)$ is defined as $\text{cap}(S)/\text{dem}(S)$, and the *sparsest-cut value* of (G, c, d) is the minimum sparsity over all cuts in G . Finally, the *flow-cut gap* in the network (G, c, d) is defined as the ratio

$$\text{gap}(G, c, d) := \frac{\text{sparsest-cut}(G, c, d)}{\text{concurrent-flow}(G, c, d)} \geq 1,$$

where the inequality is a basic exercise.

For a graph G (without capacities and demands), denote $\text{gap}(G) := \sup_{c, d} \text{gap}(G, c, d)$, where $c : E(G) \rightarrow \mathbb{R}_+$ and $d : \binom{\mathbb{T}}{2} \rightarrow \mathbb{R}_+$ range over assignments of capacities and demands. The following theorem presents the fundamental duality between flow-cut gaps and ℓ_1 distortion.

Theorem 1.2 ([AR98, LLR95, GNRS04]). *For every finite graph G , $\text{gap}(G) = c_1^+(G)$.*

Thus the following theorem states our main result in terms of flow-cut gaps.

Theorem 1.3. *Every for every plane graph G and terminal set $\mathbb{T} \subseteq V(G)$, it holds that*

$$\text{gap}(G) \leq O(\log \gamma(G; \mathbb{T})).$$

1.2 The vertex-capacitated flow-cut gap

One can consider the analogous problems in more general networks; for instance, those which are *vertex-capacitated* (instead of edge-capacitated). In that setting, bounding the flow-cut gap appears to be more significantly more challenging than for edge capacities. The authors of [FHL05] establish that the vertex-capacitated flow-cut gap is $O(\log k)$ for general networks with k terminals, and this bound is known to be tight [LR99].

For planar networks, Lee, Mendel, and Moharrami [LMM15] sought a vertex-capacitated version of the Okamura-Seymour Theorem [OS81], and proved that the vertex-capacitated flow-cut gap is $O(1)$ for instances (G, \mathbb{T}) satisfying $\gamma(G; \mathbb{T}) = 1$.

However, it was not previously known whether the gap is bounded even for $\gamma(G; \mathsf{T}) = 2$. We prove that in planar vertex-capacitated networks (G, T) with $\gamma = \gamma(G; \mathsf{T})$, the flow-cut gap is $O(\gamma)$; see Theorem 3.1. In fact, we prove this result in the more general setting of submodular vertex capacities, also known as *polymatroid networks*. This model was introduced in [CKRV15] as a generalization of vertex capacities, and the papers [CKRV15, LMM15] showed that more refined methods in metric embedding theory are able to establish upper bounds on the flow-cut gap even in this general setting.

1.3 Stochastic embeddings

Instead of embedding plane graphs with $\gamma(G; \mathsf{T}) < \infty$ directly into ℓ_1 , we will establish the stronger result that such instances can be randomly approximated by trees in a suitable sense.

If (X, d_X) is a finite metric space and \mathcal{F} is a family of finite metric spaces, then a *stochastic embedding of (X, d_X) into \mathcal{F}* is a probability distribution μ on pairs $(\varphi, (Y, d_Y))$ such that $\varphi : X \rightarrow Y$, $(Y, d_Y) \in \mathcal{F}$, and $d_Y(\varphi(x), \varphi(y)) \geq d_X(x, y)$ for all $x, y \in X$. The *expected stretch of μ* is defined by

$$\text{str}(\mu) := \max \left\{ \frac{\mathbb{E}_{(\varphi, (Y, d_Y)) \sim \mu} [d_Y(\varphi(x), \varphi(y))]}{d_X(x, y)} : x, y \in X \right\}.$$

We will refer to an undirected graph G equipped with edge lengths $\ell_G : E(G) \rightarrow \mathbb{R}_+$ as a *metric graph*, and we use d_G to denote the corresponding shortest-path distance. If G is equipped implicitly with a set $\mathsf{T}(G) \subseteq V(G)$ of terminals, we refer to it as a *terminated graph*. A graph equipped with both lengths and terminals will be called a *terminated metric graph*. We will consider any graph or metric graph G as terminated with $\mathsf{T}(G) = V(G)$ if terminals are not otherwise specified.

Given a terminated metric graph G , a *stochastic terminal embedding of G into a family \mathcal{F} of terminated metric graphs* is a distribution μ over pairs (φ, F) such that $\varphi : V(G) \rightarrow V(F)$ and $F \in \mathcal{F}$, and it holds that terminals map to terminals:

$$\forall t \in \mathsf{T}(G), \quad \mathbb{P}[\varphi(t) \in \mathsf{T}(F)] = 1,$$

and the embedding is non-contracting on terminals:

$$\forall s, t \in \mathsf{T}(G), \quad \mathbb{P}_{(\varphi, F) \sim \mu} [d_F(\varphi(s), \varphi(t)) \geq d_G(s, t)] = 1. \quad (3)$$

The *expected stretch* of this embedding, again denoted $\text{str}(\mu)$, is defined just as for general metric spaces:

$$\text{str}(\mu) := \max \left\{ \frac{\mathbb{E}_{(\varphi, F) \sim \mu} [d_F(\varphi(u), \varphi(v))]}{d_G(u, v)} : u, v \in V(G) \right\}. \quad (4)$$

Theorem 1.4. *Consider a terminated metric plane graph G with $\gamma = \gamma(G; \mathsf{T}(G))$. Then G admits a stochastic terminal embedding into the family of metric trees with expected stretch $O(\log \gamma)$.*

Theorem 1.4 immediately yields Theorem 1.1 using the fact that every finite tree metric embeds isometrically into ℓ_1 (see, e.g., [GNRS04] for further details). The bound $O(\log \gamma)$ is optimal up to the implied constant, as it is known that for an $m \times m$ planar grid equipped with uniform edge lengths, the expected stretch of any stochastic embedding into metric trees is at least $\Omega(\log m)$ [KRS01]. (A similar lower bound holds for the diamond graphs [GNRS04].)

Theorem 1.4 may also be of independent interest (including when $\mathsf{T}(G) = V(G)$) as embedding into dominating trees has many applications, including to competitive algorithms for online problems such as buy-at-bulk network design [AA97], and to approximation algorithms for combinatorial optimization, e.g., for the group Steiner tree problem [GKR00]. We remark that stochastic

terminal embeddings into metric trees were employed by [GNR10] in the context of approximation algorithms, and were later used in [EGK⁺14] to design flow sparsifiers.

2 Approximation by random trees

Before introducing our primary technical tools, we will motivate their introduction with a high-level overview of the proof of Theorem 1.4. Fix a terminated metric plane graph G with $\gamma = \gamma(G; \mathbb{T}(G)) > 1$. Our plan is to approximate G by an OS-instance (where all terminals lie on a single face) by uniting the γ faces covering $\mathbb{T}(G)$, while approximately preserving the shortest-path metric on G . The use of stochastic embeddings will come from our need to perform this approximation randomly, preserving distances only in expectation. Using the known result that OS-instances admit stochastic terminal embeddings into metric trees, this will complete the proof.

A powerful tool for randomly “simplifying” a graph is the Peeling Lemma [LS09], which informally “peels off” any subset $A \subset V(G)$ from G , by providing a stochastic embedding of G into graphs obtained by “gluing” $G \setminus A$ and the induced graph $G[A]$. The expected stretch of the embedding depends on how “nice” A is; for example, it is $O(1)$ when A is a shortest path in a planar G . The Peeling Lemma can be used to stochastically embed G into dominating OS-instances with expected stretch $2^{O(\gamma)}$ [CSW13, Section 4.5], by iteratively peeling off a shortest path A between two special faces (which has the effect of uniting them into a single face).

In contrast, our argument applies the Peeling Lemma only once. We pick A to form a connected subgraph in G that spans the γ distinguished faces. By cutting along A , one effectively merges all γ faces into a single face in a suitably chosen drawing of $G \setminus A$. The Peeling Lemma then provides a stochastic terminal embedding of G into a family of OS-instances that are constructed from copies of A and $G \setminus A$.

The expected stretch we obtain via the Peeling Lemma is controlled by how well the (induced) terminated metric graph on A can be stochastically embedded into a distribution over metric trees. For this purpose, we choose the set A to be a shortest-path tree in G that spans the γ distinguished faces, and then use a result of Sidiropoulos [Sid10] to stochastically embed A into metric trees with expected stretch that is logarithmic in the number of *leaves* (rather than logarithmic in the number of vertices, as in stochastic embeddings for general finite metric spaces [FRT04]). We remark that this is non-trivial because, while A will be (topologically) a tree spanning γ faces, the relevant metric on A is d_G (which is not a path metric on $G[A]$).

2.1 Random partitions, embeddings, and peeling

For a finite set S , we use $\text{Trees}(S)$ to denote the set of all metric spaces (S, d) that are isometric to $(V(T), d_T)$ for some metric tree T .

Theorem 2.1 (Theorem 4.4 in [Sid10]). *Let G be a metric graph, and let P_1, \dots, P_m be shortest paths in G sharing a common endpoint. Then the metric space $(\cup_{i=1}^m V(P_i), d_G)$ admits a stochastic embedding into $\text{Trees}(\cup_{i=1}^m V(P_i))$ with expected stretch at most $O(\log m)$.*

Let (X, d) be a finite metric space. A distribution ν over partitions of X is called (β, Δ) -Lipschitz if every partition P in the support of ν satisfies $S \in P \implies \text{diam}_X(S) \leq \Delta$, and moreover, for every $x, y \in X$,

$$\mathbb{P}_{P \sim \nu} [P(x) \neq P(y)] \leq \beta \cdot \frac{d(x, y)}{\Delta},$$

where for $x \in X$, we use $P(x)$ to denote the unique set in P containing x .

We denote by $\beta_{(X,d)}$ the infimal $\beta \geq 0$ such that for every $\Delta > 0$, the metric (X, d) admits a (β, Δ) -Lipschitz random partition. The following theorem is due to Klein, Plotkin, and Rao [KPR93] and Rao [Rao99].

Theorem 2.2. *For every planar graph G , we have $\beta_{(V(G), d_G)} \leq O(1)$.*

Let G be a metric graph, and consider $A \subseteq V(G)$. The *dilation of A inside G* is defined to be

$$\text{dil}_G(A) := \max_{u, v \in A} \frac{d_{G[A]}(u, v)}{d_G(u, v)},$$

where $d_{G[A]}$ denotes the induced shortest-path distance on the metric graph $G[A]$.

For two metric graphs G, G' , a *1-sum of G with G'* is a graph obtained by taking two disjoint copies of G and G' , and identifying a vertex $v \in V(G)$ with a vertex $v' \in V(G')$. This definition naturally extends to a 1-sum of any number of graphs. Note that the 1-sum inherits a canonical length function from G and G' .

Consider a subset $A \subseteq V(G)$. For $a \in A$, let G_A^a denote the graph $G[(V(G) \setminus A) \cup \{a\}]$. We define the graph \widehat{G}_A as the 1-sum of $G[A]$ with $\{G_A^a : a \in A\}$, where $G[A]$ is glued to G_A^a at their common copy of $a \in A$.

Let us write the vertex set of \widehat{G}_A as the disjoint union:

$$V(\widehat{G}_A) = \hat{A} \sqcup \bigsqcup_{a \in A} \{(a, v) : v \in V(G) \setminus A\},$$

where $\hat{A} := \{\hat{a} : a \in A\}$, and $\hat{a} \in V(\widehat{G}_A)$ denotes the copy of a in the canonical image of $G[A]$ in \widehat{G}_A , and (a, v) corresponds to the image of $v \in V(G) \setminus A$ in G_A^a . Say that a mapping $\psi : V(G) \rightarrow V(\widehat{G}_A)$ is a *selector map* if it satisfies:

1. For each $a \in A$, $\psi(a) = \hat{a}$.
2. For each $v \in V(G) \setminus A$, $\psi(v) \in \{(a, v) : a \in A\}$.

In other words, a selector maps each $a \in A$ to its unique copy in \widehat{G}_A , and maps each $v \in V(G) \setminus A$ to one of its $|A|$ copies in \widehat{G}_A .

Lemma 2.3 (The Peeling Lemma [LS09]). *Let $G = (V, E)$ be a metric graph and fix a subset $A \subseteq V$. Let G' be obtained by removing all the edges inside A :*

$$G' := (V, E') \quad \text{with} \quad E' = E \setminus E(G[A]),$$

and denote $\beta = \beta_{(V, d_{G'})}$. Then there is a stochastic embedding μ of G into the metric graph \widehat{G}_A such that μ is supported on selector maps and, moreover, $\text{str}(\mu) \leq O(\beta \cdot \text{dil}_G(A))$.

Remark 2.4. *The statement of the Peeling Lemma in [LS09] (see also [BLS10]) does not specify all the above details about the stochastic mapping from G to \widehat{G}_A , but they can be easily verified by inspecting the proof.*

2.1.1 Composition

Consider now some metric tree $T \in \text{Trees}(A)$. Via the identification between A and $\hat{A} \subseteq V(\widehat{G}_A)$, we may consider the associated metric tree $\hat{T} \in \text{Trees}(\hat{A})$. Define the metric graph $\widehat{G}_A[[T]]$ with vertex set $V(\widehat{G}_A)$ and edge set

$$E(\widehat{G}_A[[T]]) := (E(\widehat{G}_A) \setminus E(\widehat{G}_A[\hat{A}])) \cup E(\hat{T}),$$

where the edge lengths are inherited from \widehat{G}_A and \widehat{T} , respectively. In other words, we replace the edges of $\widehat{G}_A[\widehat{A}]$ with those coming from \widehat{T} .

Finally, denote by

$$\mathcal{F}_{G,A} := \left\{ \widehat{G}_A[\![T]\!] : T \in \text{Trees}(A) \right\}$$

the family of all metric graphs arising in this manner.

Lemma 2.5. *Every graph in $\mathcal{F}_{G,A}$ is a 1-sum of some $T \in \text{Trees}(A)$ with the graphs $\{G_A^a : a \in A\}$.*

Suppose that μ is a stochastic embedding of G into \widehat{G}_A that is supported on pairs (ψ, \widehat{G}_A) , where ψ is a selector map. Let ν denote a stochastic embedding of (A, d_G) into $\text{Trees}(A)$. By relabeling vertices, we may assume that ν is supported on pairs (id, T) where $\text{id} : A \rightarrow A$ is the identity map.

Consider finally the stochastic embedding of G into $\mathcal{F}_{G,A}$ which we denote $\nu \circ \mu$, and define by

$$(\nu \circ \mu)(\psi, \widehat{G}_A[\![T]\!]) := \mu(\psi, \widehat{G}_A) \cdot \nu(\text{id}, T) \quad \forall T \in \text{Trees}(A).$$

While notationally cumbersome, the following claim is now straightforward.

Lemma 2.6 (Composition Lemma). *It holds that*

$$\text{str}(\nu \circ \mu) \leq \text{str}(\nu) \cdot \text{str}(\mu).$$

2.2 Approximation by OS-instances

Let us now show that every terminated metric plane graph G with $\gamma = \gamma(G; \mathbb{T}(G))$ admits a stochastic terminal embedding into OS-instances. In Section 2.3, we recall how OS-instances can be stochastically embedded into metric trees, thereby completing the proof of Theorem 1.4.

Let $F_1, F_2, \dots, F_\gamma$ be faces of G that cover $\mathbb{T}(G)$, and denote $T_i := V(F_i) \cap \mathbb{T}(G)$. For each $i \geq 1$, fix an arbitrary vertex $v_i \in V(F_i)$. Denote $r := v_1$, and for each $i \geq 2$, let P_i be the shortest path from v_i to r . Finally, let P be the tree obtained as the union of the graphs $\{G[P_i] : i \geq 2\}$.

We present now Klein’s Tree-Cut operation [Kle06]. It takes as input a plane graph G and a tree T in G , and “cuts open” the tree to create a new face F_{new} . More specifically, consider walking “around” the tree and creating a new copy of each vertex and edge of T encountered along the way. This operation maintains planarity while replacing the tree T with a simple cycle C_T that bounds the new face. It is easy to verify that C_T has two copies of every edge of T , and $\deg_T(v)$ copies of every vertex of T , where $\deg_T(v)$ stands for the degree of v in T . This Tree-Cut operation can also be found in [Bor04, BKK07, BKM09].

We apply Klein’s Tree-Cut operation to G and the tree P , and let G_1 be the resulting metric plane graph with the new face F_{new} , after we replace P with a simple cycle C_P ; see Figure 1 for illustration. Since P shares at least one vertex with each face F_1, \dots, F_γ in G (namely, v_i), the cycle C_P shares at least one vertex with each face F_i in G_1 .

We now construct G_2 by applying two operations on G_1 . First, for every face F_i that shares exactly one vertex with C_P , namely only u_i (which is a copy of v_i), we add a new edge of length 0 incident to both F_i and F_{new} as follows. Let $N_{G_1}^1(v_i)$ be all the neighbors of v_i in G_1 embedded between the face F_i and F_{new} on one side, and $N_{G_1}^2(v_i)$ be all its neighbors on the other side. We split v_i into two vertices v'_i, v''_i in G_1 and connect all the vertices in $N_{G_1}^1(v_i)$ to v'_i and all the vertices in $N_{G_1}^2(v_i)$ to v''_i . We finish this operation by connecting v'_i to v''_i with an edge of length 0; see Figure 2.

Note that this new edge $\{v'_i, v''_i\}$ is incident to both F_i and F_{new} , and this operation maintains the planarity, along with the distance metric of G_1 (in the straightforward sense, where one takes




Figure 1: In G , the tree P (in blue) is incident to all $\gamma = 4$ distinguished special faces (drawn in green). G_1 is obtained by applying the Tree-Cut operation on G and P , which creates a new face F_{new} . Finally, G_2 is obtained by duplicating some vertices on F_{new} and connecting copies of the same vertex by length-zero edges (the dashed red edges).

a quotient by vertices at distance 0 from each other). The second operation adds between all the copies of the same $v \in V(P)$ a star with edge length 0 drawn inside F_{new} . Note that adding the stars inside F_{new} does not violate the planarity since all the copies of the vertices in C_P are ordered by the walk around P . It is easy to verify that if we identify each $v \in V(P)$ with one of its copies in G_2 arbitrarily then

$$\forall x, y \in V(G), \quad d_G(x, y) = d_{G_2}(x, y). \quad (5)$$

Lemma 2.7. $(V(P), d_G)$ admits a stochastic embedding into $\text{Trees}(V(P))$ with expected stretch at most $O(\log \gamma)$.

Proof. Apply Theorem 2.1 on the shortest-paths P_2, \dots, P_γ in G , with shared vertex $v_1 = r$. \square

Let $A \subseteq V(G_2)$ denote all the vertices on the boundary of F_{new} in G_2 . To every $T \in \text{Trees}(V(P))$, we can associate a tree $T' \in \text{Trees}(A)$ by identifying $x \in V(P)$ with one of its copies in A , and attaching the rest of its copies to x with an edge of length 0. Using (5) in conjunction with Lemma 2.7 yields the following.

Corollary 2.8. (A, d_{G_2}) admits a stochastic embedding into $\text{Trees}(A)$ with expected stretch at most $O(\log \gamma)$.

Let H be the graph obtained from G_2 by adding an edge $\{u, v\}$ of length $d_G(u, v)$ between every pair of vertices $u, v \in A$. By construction, we have $\text{dil}_H(A) = 1$. Let $E' := E(H) \setminus E(H[A])$, and $H' = (V(H), E')$. Since H' and H_A^a for $a \in A$ are subgraphs of the planar graph G_2 , they are planar as well. Thus, by Theorem 2.2, we have $\beta_{(V(H), d_{H'})} \leq O(1)$.

By applying the Peeling Lemma (Lemma 2.3) to H and $A \subseteq V(H)$, we obtain a stochastic embedding μ of H into \widehat{H}_A such that μ is supported on selector maps and $\text{str}(\mu) \leq O(1)$. Using Corollary 2.8 and the fact that (A, d_{G_3}) is the same as (A, d_{G_2}) , we obtain a stochastic embedding ν of (A, d_{G_3}) into $\text{Trees}(A)$ with $\text{str}(\nu) \leq O(\log \gamma)$.




Figure 2: The neighbors of v_i are partitioned into two sets (colored red and blue) by going around v_i in the plane and watching for the location of faces F_i and F_{new} , to eventually split v_i into two.

Define $\mathsf{T}(H)$ to be the set $\mathsf{T}(G)$, together with all the copies of nodes in $\mathsf{T}(G)$ created in the construction of H , and

$$\mathsf{T}(\widehat{H}_A) := \{\hat{a} : a \in \mathsf{T}(H)\} \cup \{(v, a) : v \in \mathsf{T}(H), a \in A\}.$$

Applying Lemma 2.6 to the pair μ, ν (in conjunction with Lemma 2.5) yields a stochastic embedding $\pi := \nu \circ \mu$ satisfying the following:

Lemma 2.9. *$(V(G), d_G)$ admits a stochastic embedding π into the family of metric graphs that are 1-sums of a metric tree with the graphs $\{H_A^a : a \in A\}$, and $\text{str}(\pi) \leq O(\log \gamma)$. Moreover, every $(\varphi, W) \in \text{supp}(\pi)$ satisfies $\varphi(\mathsf{T}(G)) \subseteq \mathsf{T}(W)$.*

For any subgraph W of H , we define $\mathsf{T}(W) := V(W) \cap \mathsf{T}(H)$. We now show that all the pairs $\{(H_A^a, \mathsf{T}(H_A^a)) : a \in A\}$ are OS-instances.

Lemma 2.10. *For every $a \in A$, there is a face F_a in H_A^a such that $V(F_a) \supseteq \mathsf{T}(H_A^a)$.*

Proof. Fix $a \in A$. The graph G_2 is planar, and while H need not be planar, the subgraphs $G_2[(V(G_2) \setminus A) \cup \{a\}]$ and H_A^a are identical for each $a \in A$. Thus, it suffices to prove the lemma for the subgraphs $G_2[(V(G_2) \setminus A) \cup \{a\}]$.

Observe that if we remove from G_2 a vertex $v \in V(G_2)$, then all the faces incident to v in G_2 become one new face in the graph $G_2 \setminus \{v\}$. Moreover, if we remove from G_2 both endpoints of an edge $\{u, v\}$, then all the faces incident to either u or v become one new face in $G_2 \setminus \{u, v\}$. Recall that $G_2[A]$ is a simple cycle (bounding F_{new}), thus $G_2[A \setminus \{a\}] = G_2[A] \setminus \{a\}$ is connected, and all the faces incident to at least one vertex in $A \setminus \{a\}$ become one new face in $G_2[(V(G_2) \setminus A) \cup \{a\}]$, which we denote F_{new}^a .

By construction of G_2 (which splits a vertex of G_1 if it is the only vertex incident to both F_i and F_{new}), every face F_i is incident to at least two vertices in A , and thus to at least one in $A \setminus \{a\}$. It follows that all the terminals in $G_2[(V(G_2) \setminus A) \cup \{a\}]$ are on the same face F_{new}^a . In addition, since a has at least one neighboring vertex $b \in A$, at least one face is incident to both a and b in G_2 , and it becomes part of the face F_{new}^a in $G_2[(V(G_2) \setminus A) \cup \{a\}]$. Therefore, $a \in V(F_{new}^a)$ as well, and the lemma follows. \square

Since the 1-sum of trees with OS-instances is again an OS-instance, together with Lemma 2.9, this yields:

Corollary 2.11. *G admits a stochastic terminal embedding into a family \mathcal{F} of terminated metric plane graphs, where each $W \in \mathcal{F}$ satisfies $\gamma(W; \mathbb{T}(W)) = 1$.*

2.3 From OS-instances to random trees

We need a couple of known embedding theorems.

Theorem 2.12 ([GNRS04, Thm. 5.4]). *Every metric outerplanar graph admits a stochastic embedding into metric trees with expected stretch $O(1)$.*

The next result is proved in [LMM15, Thm. 4.4] (which is essentially a restatement of [EGK⁺14, Thm. 12])

Theorem 2.13. *If G is a terminated metric plane graph and $\gamma(G; \mathbb{T}(G)) = 1$, then G admits a stochastic terminal embedding into metric outerplanar graphs with expected stretch $O(1)$.*

In conjunction with Theorem 2.12, this shows that every OS-instance admits a stochastic terminal embedding into metric trees with expected stretch $O(1)$. Combined with Corollary 2.11, this finishes the proof of Theorem 1.4.

3 Polymatroid flow-cut gaps

We now discuss a network model introduced in [CKRV15] that generalizes edge and vertex capacities. Recall that if S is a finite set, then a function $f : 2^S \rightarrow \mathbb{R}$ is called *submodular* if $f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$ for all subsets $A, B \subseteq S$. For an undirected graph $G = (V, E)$, we let $E(v)$ denote the set of edges incident to v . A collection $\vec{\rho} = \{\rho_v : 2^{E(v)} \rightarrow \mathbb{R}_+\}_{v \in V}$ of monotone, submodular functions are called *polymatroid capacities on G* .

Say that a function $\varphi : E \rightarrow \mathbb{R}_+$ is *feasible with respect to $\vec{\rho}$* if it holds that for every $v \in V$ and subset $S \subseteq E(v)$, it holds that $\sum_{e \in S} \varphi(e) \leq \rho_v(S)$. Given demands $\mathbf{dem} : V \times V \rightarrow \mathbb{R}_+$, one defines the *maximum concurrent flow value* of the polymatroid network $(G, \vec{\rho}, \mathbf{dem})$, denoted $\text{mcf}_G(\vec{\rho}, \mathbf{dem})$, as the maximum value $\epsilon > 0$ such that one can route an ϵ -fraction of all demands simultaneously using a flow that is feasible with respect to $\vec{\rho}$.

For every subset $S \subseteq E$, define the cut semimetric $\sigma_S : V \times V \rightarrow \{0, 1\}$ by $\sigma_S(u, v) := 0$ if and only if there is a path from u to v in the graph $G(V, E \setminus S)$. Say that a map $g : S \rightarrow V$ is *valid* if it maps every edge in S to one of its two endpoints in V . One then defines the *capacity of a set $S \subseteq E$* by

$$\nu_{\vec{\rho}}(S) := \min_{\substack{g: S \rightarrow V \\ \text{valid}}} \sum_{v \in V} \rho_v(g^{-1}(v)).$$

The *sparsity* of S is given by

$$\Phi_G(S; \vec{\rho}, \mathbf{dem}) := \frac{\nu_{\vec{\rho}}(S)}{\sum_{u, v \in V} \mathbf{dem}(u, v) \sigma_S(u, v)}.$$

We also define $\Phi_G(\vec{\rho}, \mathbf{dem}) := \min_{\emptyset \neq S \subseteq E} \Phi(S; \vec{\rho}, \mathbf{dem})$. Our goal in this section is to prove the following theorem.

Theorem 3.1. *There is a constant $C \geq 1$ such that the following holds. Suppose that $G = (V, E)$ is a planar graph and $D \subseteq F_1 \cup F_2 \cup \dots \cup F_\gamma$, where each F_i is a face of G . Then for every collection $\vec{\rho}$ of polymatroid capacities on G and every set of demands $\mathbf{dem} : D \times D \rightarrow \mathbb{R}_+$ supported on D , it holds that*

$$\text{mcf}_G(\vec{\rho}, \mathbf{dem}) \leq \Phi_G(\vec{\rho}, \mathbf{dem}) \leq C\gamma \cdot \text{mcf}_G(\vec{\rho}, \mathbf{dem}).$$

3.1 Embeddings into thin trees

In order to prove this, we need two results from [LMM15]. Suppose G is an undirected graph, T is a connected tree, and $f : V(G) \rightarrow V(T)$. For every distinct pair $u, v \in V(G)$, let P_{uv}^T denote the unique simple path from $f(u)$ to $f(v)$ in T . Say that the map f is Δ -thin if, for every $u \in V(G)$, the induced subgraph on $\bigcup_{v:\{u,v\} \in E(G)} P_{uv}^T$ can be covered by Δ simple paths in T emanating from $f(u)$.

Suppose further that G is equipped with edge lengths $\ell : E(G) \rightarrow \mathbb{R}_+$. If (X, d_X) is a metric space and $f : V(G) \rightarrow X$, we make the following definition. For $\tau > 0$ and any $u \in V(G)$:

$$|\nabla_\tau f(u)|_\infty := \max \left\{ \frac{d_X(f(u), f(v))}{\ell(u, v)} : \{u, v\} \in E \text{ and } \ell(u, v) \in [\tau, 2\tau] \right\}$$

Fact 3.2. *Suppose that $f : V(G) \rightarrow \mathbb{R}$ is 1-Lipschitz, where $V(G)$ is equipped with the path metric $d_{G,\ell}$. Then f is 2-thin and*

$$\max \{ |\nabla_\tau f(u)|_\infty : u \in V(G), \tau > 0 \} \leq 1.$$

Theorem 3.3 (Rounding theorem [LMM15]). *Consider a graph $G = (V, E)$ and a subset $D \subseteq V$. Suppose that for every length $\ell : E \rightarrow \mathbb{R}_+$, there is a random Δ -thin mapping $\Psi : V \rightarrow V(T)$ into some random tree T that satisfies:*

1. For every $v \in V$ and $\tau > 0$: $\mathbb{E} |\nabla_\tau \Psi(v)|_\infty \leq L$
2. For every $u, v \in D$:

$$\mathbb{E} [d_T(\Psi(u), \Psi(v))] \geq \frac{d_{G,\ell}(u, v)}{K}.$$

Then for every collection $\vec{\rho}$ of polymatroid capacities on G and every set of demands $\text{dem} : D \times D \rightarrow \mathbb{R}_+$ supported on D , it holds that

$$\text{mcf}_G(\vec{\rho}, \text{dem}) \leq \Phi_G(\vec{\rho}, \text{dem}) \leq O(\Delta K L) \cdot \text{mcf}_G(\vec{\rho}, \text{dem}).$$

Theorem 3.4 (Face embedding theorem [LMM15]). *Suppose that $G = (V, E)$ is a planar graph and $D \subseteq V$ is a subset of vertices contained in a single face of G . Then for every $\ell : E \rightarrow \mathbb{R}_+$, there is a random 4-thin mapping $\Psi : V \rightarrow V(T)$ into a random tree metric that satisfies the assumptions of Theorem 3.3 with $K, L \leq O(1)$.*

We now use this to prove the following multi-face embedding theorem; combined with Theorem 3.3, it yields Theorem 3.1.

Theorem 3.5 (Multi-face embedding theorem). *Suppose that $G = (V, E)$ is a planar graph and $D \subseteq F_1 \cup F_2 \cup \dots \cup F_\gamma$, where each F_i is a face of G . Then for every $\ell : E \rightarrow \mathbb{R}_+$, there is a random 4-thin mapping $\Psi : V \rightarrow V(T)$ into a random tree metric that satisfies the assumptions of Theorem 3.3 with $L \leq O(1)$ and $K \leq O(\gamma)$.*

Proof. For each $i = 1, 2, \dots, \gamma$, let $\Psi_i : V \rightarrow V(T_i)$ be the random 4-thin mapping guaranteed by Theorem 3.4 with constants $1 \leq K_0, L_0 \leq O(1)$, and let $\Psi'_i : V \rightarrow \mathbb{R}$ be the 2-thin mapping given by $\Psi'_i(v) = d_{G,\ell}(v, F_i)$ (recall Fact 3.2). Now let $\Psi : V \rightarrow V(T)$ be the random map that arises from choosing one of $\{\Psi_1, \dots, \Psi_\gamma, \Psi'_1, \dots, \Psi'_\gamma\}$ uniformly at random. Then Ψ is a random 4-thin mapping satisfying (1) in Theorem 3.3 for some $L \leq O(1)$.

Consider now some $u \in F_i$ and $v \in V$. Let $u' \in F_i$ be such that $d_{G,\ell}(v, u') = d_{G,\ell}(v, F_i)$. If $d_{G,\ell}(u', v) \geq \frac{d_{G,\ell}(u,v)}{4K_0L_0}$, then

$$\mathbb{E} [d_T(\Psi(u), \Psi(v))] \geq \frac{1}{2\gamma} |\Psi'_i(u) - \Psi'_i(v)| = \frac{d_{G,\ell}(u', v)}{2\gamma} \geq \frac{d_{G,\ell}(u, v)}{8\gamma K_0 L_0}.$$

If, on the other hand, $d_{G,\ell}(u', v) < \frac{d_{G,\ell}(u,v)}{4K_0L_0}$, then

$$\begin{aligned} \mathbb{E} [d_T(\Psi(u), \Psi(v))] &\geq \frac{1}{2\gamma} \mathbb{E} [d_{T_i}(\Psi_i(u), \Psi_i(v))] \\ &\geq \frac{1}{2\gamma} \mathbb{E} [d_{T_i}(\Psi_i(u), \Psi_i(u')) - d_{T_i}(\Psi_i(u'), \Psi_i(v))] \\ &\geq \frac{1}{2\gamma} \left(\frac{d_{G,\ell}(u, u')}{K_0} - L_0 d_{G,\ell}(u', v) \right) \\ &\geq \frac{1}{2\gamma} \left(\frac{d_{G,\ell}(u, v) - d_{G,\ell}(u', v)}{K_0} - \frac{d_{G,\ell}(u, v)}{4K_0} \right) \\ &\geq \frac{1}{2\gamma} \left(\frac{3}{4} \frac{d_{G,\ell}(u, v)}{K_0} - \frac{d_{G,\ell}(u', v)}{K_0} \right) \\ &\geq \frac{d_{G,\ell}(u, v)}{4\gamma K_0}. \end{aligned}$$

Thus Ψ also satisfies (2) in Theorem 3.3 with $K \leq O(\gamma)$, completing the proof. \square

Acknowledgments

We thank Anupam Gupta for useful discussions at an early stage of the research.

References

- [AA97] B. Awerbuch and Y. Azar. Buy-at-bulk network design. In *38th Annual Symposium on Foundations of Computer Science, FOCS '97*, pages 542–547. IEEE Computer Society, 1997.
- [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*, pages 952–963, 2018.
- [AR98] Yonatan Aumann and Yuval Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM J. Comput.*, 27(1):291–301, 1998.
- [BKK07] Glencora Borradaile, Claire Kenyon-Mathieu, and Philip N. Klein. A polynomial-time approximation scheme for Steiner tree in planar graphs. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007*, pages 1285–1294, 2007.
- [BKM09] Glencora Borradaile, Philip N. Klein, and Claire Mathieu. An $o(n \log n)$ approximation scheme for Steiner tree in planar graphs. *ACM Trans. Algorithms*, 5(3):31:1–31:31, 2009.
- [BLS10] Glencora Borradaile, James R. Lee, and Anastasios Sidiropoulos. Randomly removing g handles at once. *Comput. Geom.*, 43(8):655–662, 2010.
- [Bor04] G. Borradaile. *Exploiting Planarity for Network Flow and Connectivity Problems*. PhD thesis, Brown University, 2004.
- [CKRV15] Chandra Chekuri, Sreeram Kannan, Adnan Raja, and Pramod Viswanath. Multicommodity flows and cuts in polymatroidal networks. *SIAM J. Comput.*, 44(4):912–943, 2015.

- [CSW13] Chandra Chekuri, F. Bruce Shepherd, and Christophe Weibel. Flow-cut gaps for integer and fractional multiflows. *Journal of Combinatorial Theory, Series B*, 103(2):248 – 273, 2013.
- [EGK⁺14] M. Englert, A. Gupta, R. Krauthgamer, H. Räcke, I. Talgam-Cohen, and K. Talwar. Vertex sparsifiers: New results from old techniques. *SIAM Journal on Computing*, 43(4):1239–1262, 2014.
- [FHL05] Uriel Feige, Mohammad Taghi Hajiaghayi, and James R. Lee. Improved approximation algorithms for minimum-weight vertex separators. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 563–572, 2005.
- [FRT04] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004.
- [GKR00] N. Garg, G. Konjevod, and R. Ravi. A polylogarithmic approximation algorithm for the group Steiner tree problem. *J. Algorithms*, 37(1):66–84, 2000.
- [GNR10] Anupam Gupta, Viswanath Nagarajan, and R. Ravi. An improved approximation algorithm for requirement cut. *Oper. Res. Lett.*, 38(4):322–325, 2010.
- [GNRS04] Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Cuts, trees and l_1 -embeddings of graphs. *Combinatorica*, 24(2):233–269, 2004.
- [Kle06] Philip N. Klein. A subset spanner for planar graphs, with application to subset TSP. In *38th Annual ACM Symposium on Theory of Computing*, pages 749–756, 2006.
- [KLL02] Ton Kloks, Chuan-Min Lee, and Jiping Liu. New algorithms for k-face cover, k-feedback vertex set, and k-disjoint cycles on plane and planar graphs. In *Graph-Theoretic Concepts in Computer Science, 28th International Workshop, WG 2002, Revised Papers*, pages 282–295, 2002.
- [KPR93] Philip Klein, Serge A. Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the Twenty-fifth Annual ACM Symposium on Theory of Computing*, STOC '93, pages 682–690, 1993.
- [KRS01] Goran Konjevod, R. Ravi, and F. Sibel Salman. On approximating planar metrics by tree metrics. *Inform. Process. Lett.*, 80(4):213–219, 2001.
- [LLR95] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- [LMM15] James R. Lee, Manor Mendel, and Mohammad Moharrami. A node-capacitated Okamura-Seymour theorem. *Math. Program.*, 153(2, Ser. A):381–415, 2015.
- [LR99] T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.
- [LR10] James R. Lee and Prasad Raghavendra. Coarse differentiation and multi-flows in planar graphs. *Discrete & Computational Geometry*, 43(2):346–362, 2010.
- [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*, pages 245–254, 2009.
- [LS13] James R. Lee and Anastasios Sidiropoulos. Pathwidth, trees, and random embeddings. *Combinatorica*, 33(3):349–374, 2013.
- [OS81] Haruko Okamura and P.D. Seymour. Multicommodity flows in planar graphs. *Journal of Combinatorial Theory, Series B*, 31(1):75 – 81, 1981.
- [Rao99] Satish Rao. Small distortion and volume preserving embeddings for planar and euclidean metrics. In *Proceedings of the Fifteenth Annual Symposium on Computational Geometry*, SCG '99, pages 300–306. ACM, 1999.
- [Sid10] Anastasios Sidiropoulos. Optimal stochastic planarization. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pages 163–170, 2010.