A Node-capacitated Okamura-Seymour Theorem

James R. Lee*
Computer Science & Engineering
University of Washington
jrl@cs.washington.edu

Manor Mendel†
Mathematics and Computer Science
Open University of Israel
mendelma@gmail.com

Mohammad Moharrami*
Computer Science & Engineering
University of Washington
mohammad@cs.washington.edu

ABSTRACT

The classical Okamura-Seymour theorem states that for an edge-capacitated, multi-commodity flow instance in which all terminals lie on a single face of a planar graph, there exists a feasible concurrent flow if and only if the cut conditions are satisfied. Simple examples show that a similar theorem is impossible in the node-capacitated setting. Nevertheless, we prove that an approximate flow/cut theorem does hold: For some universal $\varepsilon > 0$, if the node cut conditions are satisfied, then one can simultaneously route an ε-fraction of all the demands. This answers an open question of Chekuri and Kawarabayashi. More generally, we show that this holds in the setting of multi-commodity polymatroid networks introduced by Chekuri, et. al. Our approach employs a new type of random metric embedding in order to round the convex programs corresponding to these more general flow problems.

Categories and Subject Descriptors

G.2.2 [**Mathematics of Computing**]: Discrete Mathematics—*Graph Thoery*

General Terms

Theory

Keywords

Metric embeddings; Multi-commodity flows

1. INTRODUCTION

The relationship between flows and cuts in graphs has played a fundamental role in combinatorial optimization. A seminal result in the study of multi-commodity flows is the classical Okamura-Seymour theorem [23] which we now recall.

An *undirected flow network* is an undirected graph $G = (V, E)$ together with a capacity function on edges $\mathsf{cap} : E \to [0, \infty)$. A set of *demands* is specified by a symmetric mapping $\mathsf{dem} : V \times V \to [0, \infty)$. For $u, v \in V$, denote by $\varphi_{uv} : E \to [0, \infty)$ the undirected u-v flow. The (edge) capacity constrains require that for every $e \in E$, $\sum_{u,v \in V} \varphi_{uv}(e) \leq \mathsf{cap}(e)$. Given such an instance, let $\mathsf{mcf}_G(\mathsf{cap}, \mathsf{dem})$ be the largest value ε such that one can simultaneous route $\varepsilon \cdot \mathsf{dem}(u, v)$ units of flow between u and v for every $u, v \in V$ while not violating any of the edge capacities. This optimization describes the *maximum concurrent flow problem.*

For two subsets $S, T \subseteq V$, let $\mathsf{cap}(S, T)$ denote the total capacity of all edges with one endpoint in S and one in T. Similarly, let $\mathsf{dem}(S, T) = \sum_{u \in S} \sum_{v \in T} \mathsf{dem}(u, v)$. To give an upper bound on mcf, we can consider cuts in G, described by subsets $S \subseteq V$. To every such subset we assign a value called the *sparsity* of the cut:

$$\Phi_G(S; \mathsf{cap}, \mathsf{dem}) = \frac{\mathsf{cap}(S, \bar{S})}{\mathsf{dem}(S, \bar{S})} .$$

It is straightforward to see that for any $S \subseteq V$, we have $\mathsf{mcf}_G(\mathsf{cap}, \mathsf{dem}) \leq \Phi_G(S; \mathsf{cap}, \mathsf{dem})$. The *sparsest cut* is the one which gives the best upper bound on mcf. In this vein, we define

$$\Phi_G(\mathsf{cap}, \mathsf{dem}) = \min_{S \subseteq V} \Phi_G(S; \mathsf{cap}, \mathsf{dem}) .$$

Thus we have the relationship $\mathsf{mcf}_G(\mathsf{cap}, \mathsf{dem}) \leq \Phi_G(\mathsf{cap}, \mathsf{dem})$ and the flow/cut gap question asks how close this upper bound is to the truth.

To state the Okamura-Seymour theorem, we need one final piece of notation. We say that the demand function dem is *supported on a subset* $D \subseteq V$ if $\mathsf{dem}(u, v) > 0$ only when $u, v \in D$. The classical Max-flow Min-cut Theorem [10] implies that if the demand dem is supported on a two-element subset $\{s, t\} \subseteq V$, then for any capacities cap, we have $\mathsf{mcf}_G(\mathsf{cap}, \mathsf{dem}) = \Phi_G(\mathsf{cap}, \mathsf{dem})$. An extension of Hu [14] shows that if dem is supported on a 4-element subset $D \subseteq V$, the same equality holds. The Okamura-Seymour theorem states that whenever G is a planar graph and the demand is supported on a single face, there is likewise no flow/cut gap.

*Research supported by NSF CCF-1217256 and a Sloan Research Fellowship.

†Mathematics and Computer Science Department, Open University of Israel. Research partially supported by ISF grants 221/07, 93/11, and BSF grants 2006009, 2010021. This work was carried out during visits of the author to Microsoft Research and the University of Washington.

THEOREM 1.1 ([23]). *Let $G = (V, E)$ be a planar graph, and let $F \subseteq V$ be any face of G. Then for any capacities* $\mathsf{cap} : E \to [0, \infty)$ *and any demands* $\mathsf{dem} : V \times V \to [0, \infty)$ *supported on F, we have*

$$\mathsf{mcf}_G(\mathsf{cap}, \mathsf{dem}) = \Phi_G(\mathsf{cap}, \mathsf{dem}) .$$

We remark that this theorem has applications beyond flow/cut gaps. For instance, in [7] the authors use it as a fundamental step in solving the edge-disjoint paths problem in planar graphs with constant congestion. A significant motivation of the present paper is to serve as a step in extending their work to the more difficult vertex-disjoint paths problem.

Indeed, one can consider generalizations of edge-capacitated networks. A prominent example is to consider capacities on vertices (one motivation is that this a more relevant type of constraint in wireless networks [6]). While one can simulate edge capacities by introducing a new vertex in the middle of an edge, it does not seem that any reduction is known by which one can simulate vertex capacities with edge capacities.

Vertex-capacitated flows. Formally, we define a *vertex-capacitated flow network* by considering a capacity function $\mathsf{cap} : V \to [0, \infty)$ assigning capacities to vertices instead of edges. It seems that the most elegant way to think about capacities in this setting is as follows: If a flow of value α is sent along a path P from s to t, then it consumes $\alpha/2$ capacity at s and t and α capacity at each of the intermediate nodes of P.[1] Formally, in the multi-commodity setting, the vertex capacity constrains require that for every $w \in V$, $\sum_{e \ni w} \sum_{u,v \in V} \varphi_{uv}(e) \leq 2\,\mathsf{cap}(w)$. The corresponding definition of the maximum concurrent flow follows immediately; we use the notation mcf^v_G for the vertex-capacitated version. For the definition of Φ^v_G, we have to be slightly more careful. For a subset $S \subseteq V$ of the vertices, denote by $G[S]$ the induces subgraph of G on S. We define a function $\rho_S : V \times V \to \{0, \frac{1}{2}, 1\}$ by

$$\rho_S(u,v) = \begin{cases} \frac{1}{2} & |\{u,v\} \cap S| = 1 \\ 1 & u,v \in S \\ 1 & u,v \in \bar{S} \text{ and } u,v \text{ are in distinct connected} \\ & \qquad\qquad \text{components of } G[\bar{S}] \\ 0 & \text{otherwise.} \end{cases}$$

In other words, we are only given half-credit for separating u and v if exactly one of them is in the separator. Then we define

$$\Phi^v_G(S; \mathsf{cap}, \mathsf{dem}) = \frac{\sum_{v \in S} \mathsf{cap}(v)}{\sum_{u,v \in V} \mathsf{dem}(u,v)\rho_S(u,v)} ,$$

and $\Phi^v_G(\mathsf{cap}, \mathsf{dem}) = \min_{S \subseteq V} \Phi^v_G(S; \mathsf{cap}, \mathsf{dem})$. It is straightforward to verify that $\mathsf{mcf}^v_G(\mathsf{cap}, \mathsf{dem}) \leq \Phi^v_G(\mathsf{cap}, \mathsf{dem})$.

These precise definitions ensure that a classical Max-flow Min-cut theorem holds when the demand is supported on a single pair (this follows from Menger's theorem [22]). They also allow other natural properties in the multi-commodity setting; it is an exercise to show that for any tree T, we have $\mathsf{mcf}^v_T(\mathsf{cap}, \mathsf{dem}) = \Phi^v_T(\mathsf{cap}, \mathsf{dem})$ for any choice of capacities and demands. In some sense, this property will be crucial for our approach later.

[1] This particular choice does not materially affect any theorem in the paper which deals with approximate flow/cut gaps.

Unfortunately, there is no exact vertex-capacitated analog of the Okamura-Seymour Theorem. The planar graph in Figure 1 has all vertices on the outer face. The capacities are specified on the vertices and the demands are given by dotted edges in the figure; all demands have value 1. It is straightforward to check that one has $\Phi^v_G(\mathsf{cap}, \mathsf{dem}) = 1$ and yet $\mathsf{mcf}^v_G(\mathsf{cap}, \mathsf{dem}) = 5/7$. The shaded nodes form a vertex cut of sparsity 1.

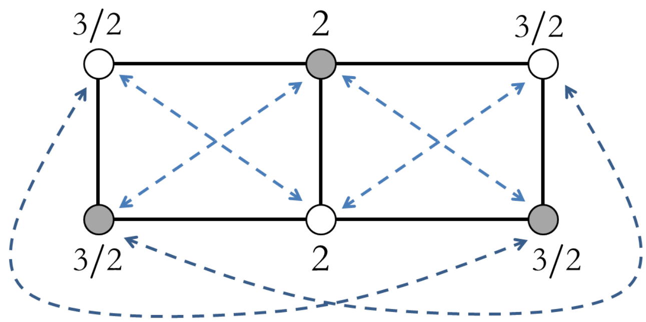

Figure 1: A counterexample to an exact node-capacitated Okamura-Seymour Theorem.

However, a main result of the present paper is that an approximate version does hold in the node-capacitated setting, answering a question posed by Chekuri and Kawarabayashi.

THEOREM 1.2 (APPROXIMATE O-S THEOREM). *There exists a constant $\varepsilon > 0$ such that the following holds. Let $G = (V, E)$ be a planar graph and let $F \subseteq V$ be any face of G. Then for any vertex capacities* $\mathsf{cap} : V \to [0, \infty)$ *and any demands* $\mathsf{dem} : V \times V \to [0, \infty)$ *supported on F, we have*

$$\mathsf{mcf}^v_G(\mathsf{cap}, \mathsf{dem}) \geq \varepsilon \cdot \Phi^v_G(\mathsf{cap}, \mathsf{dem}) .$$

In fact, our result holds in the more general setting of undirected polymatroid networks which we discuss next.

1.1 Polymatroid networks

Motivated by applications to information flow in wireless networks, Chekuri et. al. [6] introduced a generalization of vertex capacities by putting a submodular capacity function at every vertex. Recall that a function $f : 2^S \to \mathbb{R}$ over a finite set S is called *submodular* if $f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$ for all $A, B \subseteq S$. Let $G = (V, E)$ be a graph and suppose that for a multi-commodity flow $\varphi = \{\varphi_{st}\}_{s,t \in V}$ in G, we use $\varphi(e) = \sum_{s,t} \varphi_{st}(e)$ to denote the total flow through the edge e. For a vertex $v \in V$, we use $E(v)$ to denote the edges incident to v. Let $\vec{\rho} = \{\rho_v : 2^{E(v)} \to [0, \infty)\}_{v \in V}$ be a collection of monotone, submodular functions called *polymatroid capacities.* A flow φ is *feasible with respect to* $\vec{\rho}$ if for every $v \in V$ and every subset $S \subseteq E(v)$, we have $\sum_{e \in S} \varphi(e) \leq \rho_v(S)$.

Given a demand function $\mathsf{dem} : V \times V \to [0, \infty)$, we can define the *maximum concurrent flow value* of a polymatroid network by $\mathsf{mcf}_G(\vec{\rho}, \mathsf{dem})$ as the maximum $\varepsilon > 0$ such that one can route an ε-fraction of all demands simultaneously using a flow that is feasible with respect to $\vec{\rho}$.

The corresponding notion of a sparse cut is now a little trickier. For every subset of edges $S \subseteq E$, we can define the cut semi-metric $\sigma_S : V \times V \to \{0,1\}$ on V by $\sigma_S(x,y) = 0$ if and only if there exists a path from x to y in the graph $G(V, E \setminus S)$. Following [6], we call a map $g : S \to V$ *valid* if it maps every edge in S to one of its two endpoints in V. We can then define the *capacity of a set* $S \subseteq E$ by

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

Finally, we define the *sparsity of* S by

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

and define $\Phi_G(\vec{\rho}, \mathsf{dem}) = \min_{S\subseteq V} \Phi_G(S; \vec{\rho}, \mathsf{dem})$. It is not too difficult to see that, again,

$$\mathsf{mcf}_G(\vec{\rho}, \mathsf{dem}) \le \Phi_G(\vec{\rho}, \mathsf{dem}) \,.$$

In [6], it is proved that when dem is supported on a single pair, we have

$$\Phi_G(\vec{\rho}, \mathsf{dem}) \le 2 \cdot \mathsf{mcf}_G(\vec{\rho}, \mathsf{dem}) \,.$$

Unfortunately, the factor 2 is necessary, and owes itself to a slight defect in the notion of undirected polymatroid networks. If one were to say that a flow only consumes half the capacity of an edge if it originates at an endpoint (as in the vertex-capacitated case described above), then we would obtain an exact single-commodity max-flow/min-cut theorem in this setting. Indeed, for directed polymatroid networks, such a result is classical [13, 17]. Since we are concerned here with approximate flow/cut gaps, this will not be an issue, and we follow [6]. We obtain an Okamura-Seymour theorem for polymatroid networks as well, answering a question posed to us by Chandra Chekuri.

THEOREM 1.3 (POLYMATROID O-S THEOREM). *There exists a constant $\varepsilon > 0$ such that the following holds. Let $G = (V, E)$ be a planar graph and let $F \subseteq V$ be any face of G. Then for any polymatroid capacities $\vec{\rho}$ and any demands $\mathsf{dem} : V \times V \to [0, \infty)$ supported on F, we have*

$$\mathsf{mcf}_G(\vec{\rho}, \mathsf{dem}) \ge \varepsilon \cdot \Phi_G(\vec{\rho}, \mathsf{dem}) \,.$$

Theorem 1.2 is a special case of Theorem 1.3. Indeed, vertex capacity $\mathsf{cap} : V \to [0, \infty)$, is (up to a factor of 2) equivalent to vertex polymatroid capacity $\vec{\rho}_v(\emptyset) = 0$ and $\vec{\rho}_v(S) = \mathsf{cap}(v)$ for $\emptyset \ne S \subseteq E(v)$. With this definition of $\vec{\rho}$, it is immediate to check that $\mathsf{mcf}_G(\vec{\rho}, \mathsf{dem}) \le \mathsf{mcf}^v_G(\mathsf{cap}, \mathsf{dem}) \le 2\mathsf{mcf}_G(\vec{\rho}, \mathsf{dem})$ and $\Phi_G(\vec{\rho}, \mathsf{dem}) \le \Phi^v_G(\mathsf{cap}, \mathsf{dem}) \le 2\Phi_G(\vec{\rho}, \mathsf{dem})$.

1.2 Embeddings and flow/cut gaps

Our main tools in proving Theorems 1.2 and 1.3 are various embeddings of metric spaces. To this end, we first recall known results in the edge and vertex-capacitated settings. In the next section, we discuss the new types of embeddings we need to handle vertex-capacitated and polymatroid networks.

A *metric graph* $G = (V, E, \mathsf{len})$ is an undirected graph equipped with a non-negative length function on edges $\mathsf{len} : E \to [0, \infty)$. We extend the length function to paths $P \subseteq E$ by setting $\mathsf{len}(P) = \sum_{e\in P} \mathsf{len}(e)$. Associated to every such length is the shortest-path pseudo-metric on G defined by $d_{\mathsf{len}}(u, v) = \min_P \mathsf{len}(P)$ where the minimum is over all u-v paths P in G. We say that a pseudo-metric d on V is *supported on the graph* G if $d = d_{\mathsf{len}}$ for some length function on E. In many situations we will only be considering a single length function on G at a time, and then we write d_G instead of d_{len}.

We will consider embeddings of such graph metrics into various other spaces. Given two metric spaces (X, d_X) and (Y, d_Y) and a function $f : X \to Y$, we define the *Lipschitz constant of* f by

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

If $\|f\|_{\mathrm{Lip}} \le L$, we say that f is *L-Lipschitz.*

We define the *distortion* of the map f by $\mathsf{dist}(f) = \|f\|_{\mathrm{Lip}} \cdot \|f^{-1}\|_{\mathrm{Lip}}$. The L_1 *distortion of a metric space* (X, d_X), written $c_1(X, d_X)$, denotes the infimum of $\mathsf{dist}(f)$ over all maps $f : X \to L_1$. The next theorem gives a tight relationship between flow/cut gaps in graphs and L_1 embeddings of the metric supported on them. It follows from [20] and [12].

THEOREM 1.4. *Consider any graph $G = (V, E)$ and any subset $D \subseteq V$. Let*

$$K_1(G, D) = \sup_{\mathsf{cap},\mathsf{dem}} \frac{\Phi_G(\mathsf{cap}, \mathsf{dem})}{\mathsf{mcf}_G(\mathsf{cap}, \mathsf{dem})} \,,$$

where the supremum is over all capacity functions $\mathsf{cap} : E \to [0, \infty)$ and all demand functions supported on D. Let

$$K_2(G, D) = \sup_d \left[\inf_{f:(V,d)\to L_1} \mathsf{dist}(f|_D) \right] ,$$

where the supremum is over all metrics d supported on G and the infimum is over all 1-Lipschitz mappings $f : V \to L_1$. Then $K_1(G, D) = K_2(G, D)$.

In particular, the Okamura-Seymour Theorem (Thm. 1.1) can be restated as the following fact about embeddings of planar graphs: For any metric planar graph $G = (V, E)$ and any face $F \subseteq V$, there exists a 1-Lipschitz mapping $f : V \to L_1$ such that $\mathsf{dist}(f|_F) = 1$.

Vertex-capacitated flows and ℓ_1^{dom} embeddings. Unfortunately, L_1 embeddings are not sufficient for the study of vertex-capacitated flow/cut gaps; we refer to [9] for some examples. Instead, [9] uses a stronger notion of embedding. For simplicity, we discuss such embeddings only for finite metric spaces. An *ℓ_1^{dom} embedding of a finite pseudometric space* (X, d) is a random 1-Lipschitz mapping $\Lambda : X \to \mathbb{R}$. One then defines

$$\mathsf{dist}(\Lambda) = \max_{x,y\in X} \frac{\mathbb{E}\,|\Lambda(x) - \Lambda(y)|}{d(x, y)} \,,$$

and writes $c_1^{\mathrm{dom}}(X, d)$ for the infimum of $\mathsf{dist}(\Lambda)$ over all such random mappings $\Lambda : X \to \mathbb{R}$. It is straightforward to verify that $c_1(X, d) \le c_1^{\mathrm{dom}}(X)$ and there are many interesting cases when this inequality is strict (see [9, 3]). Such embeddings were initially studied by Matousek and Rabinovich [21]. It was shown in [9] that they can be used to bound vertex-capacitated flow/cut caps, and [6] extended this to undirected polymatroid networks.

THEOREM 1.5 ([6]). *Consider a graph $G = (V, E)$ and a subset $D \subseteq V$. Suppose there is a constant $K \ge 1$ such that for every metric d supported on G, we have $c_1^{\mathrm{dom}}(D, d) \le K$. Then for every set of polymatroid capacities $\vec{\rho}$ on G and every $\mathsf{dem} : V \times V \to [0, \infty)$ supported on D, we have*

$$\mathsf{mcf}_G(\vec{\rho}, \mathsf{dem}) \ge \frac{1}{2K} \Phi_G(\vec{\rho}, \mathsf{dem}) \,.$$

Despite the power of the preceding theorem, it is insufficient for proving our main results. Since $c_1^{\mathrm{dom}}(X,d)$ is at least the Euclidean distortion of (X,d), Bourgain's lower bound on the Euclidean distortion of trees [2] implies that there are n-point tree metrics (T_n,d_n) with $c_1^{\mathrm{dom}}(T_n,d_n) = \Omega(\sqrt{\log\log n})$. In the next section, we introduce a new notion of embedding that is sufficient for proving vertex-capacitiated and polymatroid versions of the Okamura-Seymour theorem.

1.3 Length functions, star-shaped embeddings, and single-scale gradients

We first setup a polymatroid embedding problem which follows from the duality theorem of [6]. Fix a finite ground set S. Given a function $\rho : \{0,1\}^S \to \{0,1\}$, we define its *Lovász extension* $\hat\rho : [0,\infty)^S \to [0,\infty)$ by

$$\hat\rho(z) = \int_0^\infty \rho(z^\theta)d\theta\,,$$

where $z^\theta \in \{0,1\}^S$ has $(z^\theta)_i = 1$ whenever $z_i \geq \theta$. Observe that for a constant $\alpha > 0$, we have $\hat\rho(\alpha\cdot z) = \alpha\cdot\hat\rho(z)$. We will associate 2^S and $\{0,1\}^S$ via the mapping which sends a subset $A \subseteq S$ to its characteristic function $\mathbf{1}_A \in \{0,1\}^S$. Likewise, we will associate functions $S \to [0,\infty)$ with elements of $[0,\infty)^S$.

In the rest of this section, we will consider families of functions $\mathcal{F} = \{\ell_v : E(v) \to [0,\infty)\}_{v\in V}$ associated to a graph $G = (V,E)$. Given a length function $\mathsf{len} : E \to [0,\infty)$, we say that $\mathcal{F}$ *is adapted to* len if for every edge $e = \{u,v\} \in E$, we have

$$\mathsf{len}(e) \leq \ell_u(e) + \ell_v(e)\,.$$

The next result is from [6].

Theorem 1.6 (Duality Theorem). *For any graph $G = (V,E)$ the following holds. For any polymatroid capacities $\vec\rho = \{\rho_v : v \in V\}$ and any demands $\mathsf{dem} : V\times V \to [0,\infty)$,*

$$\mathsf{mcf}_G(\vec\rho,\mathsf{dem}) = \min_{\mathsf{len},\{\ell_v\}}\left[\frac{\sum_{v\in V}\hat\rho_v(\ell_v)}{\sum_{u,v\in V}\mathsf{dem}(u,v)d_{\mathsf{len}}(u,v)}\right], \quad (1)$$

where the minimum is over all length functions $\mathsf{len} : E \to [0,\infty)$ on G and all families $\{\ell_v : E(v) \to [0,\infty)\}_{v\in V}$ adapted to len.

The preceding theorem shows that to prove flow/cut gaps, it suffices to find for every given length function len and any $\{\ell_v\}_{v\in V}$ adapted to len, a set $S \subseteq E$ for which

$$\Phi_G(S;\vec\rho,\mathsf{dem}) \leq C\cdot\frac{\sum_{v\in V}\hat\rho_v(\ell_v)}{\sum_{u,v\in V}\mathsf{dem}(u,v)d_{\mathsf{len}}(u,v)}$$

for some constant $C > 0$. This gives rise to an embedding problem with differs from the classical one in a way which we now describe informally.

In the case of edge-capacitated flows and L_1 embeddings, to satisfy the Lipschitz property, it suffices to consider the stretch of each edge separately. For vertex-capacitated flows, and more generally polymatroid networks, we must *coordinate* the stretch of the edges adjacent to a vertex. In essence, a vertex has to "pay" in the corresponding "Lipschitz constant" if *any* of its adjacent edges is stretched. Thus we should try as much as possible to stretch the edges adjacent to a vertex simultaneously.

This makes some standard techniques (e.g. random embeddings into trees as in [12]) inappropriate for our study (although some of the principles in [12] will prove invaluable). Certainly ℓ_1^{dom} embeddings achieve this coordination because they are (by definition) Lipschitz in every coordinate, but as we mentioned earlier, they are insufficient for proving our main theorems.

To satisfy this goal, we must pay careful attention to the image of the edges in our embeddings. On the other hand, to overcome the limitations of ℓ_1^{dom}, we will increase our target spaces to include general metric trees.

Star-shaped mappings. Say that a graph H is *star-shaped* if H is the subdivision of some star graph. Suppose that $G = (V,E)$ is a graph, T is a tree, and $\lambda : V \to V(T)$ is an arbitrary map. For every $u,v \in V(T)$, let $P_{uv} \subseteq V(T)$ be the unique simple path between u and v in T. We say that λ is a *star-shaped mapping* if, for every $u \in V(T)$, the induced graph on

$$\{P_{uv} : v \in V(T), E(\lambda^{-1}(u),\lambda^{-1}(v)) \neq \emptyset\}$$

is star-shaped. In other words, if we consider the paths in T which correspond to edges in G, then all such paths emanating from the same vertex in T should form a star-shaped subgraph.

In addition to controlling the *shape* of a mapping, we need to control the lengths of the "arms" of the star simultaneously. Fortunately (and this property will be crucial to the approach of Section 4), we will only need to bound the stretch over single scales.

Single-scale ℓ_∞ gradients. If we are given a metric graph $G = (V,E,\mathsf{len})$ and a mapping $f : V \to (X,d_X)$ into a metric space (X,d_X), we make the following definition: For any $\tau > 0$,

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

In Section 2, we prove the following theorem which shows how such mappings can be used for polymatroid flow/cut gaps.

Theorem 1.7 (Main rounding theorem). *Let $G = (V,E,\mathsf{len})$ be a metric graph and suppose there exists a random metric tree T and a random star-shaped mapping $F : V \to V(T)$ such that for some $K \geq 1$,*

$$\max_{v\in V}\sup_{\tau>0}\mathbb{E}\,|\nabla_\tau F(v)|_\infty \leq K\,. \quad (2)$$

Then for any family of functions $\{\ell_v : E(v)\to[0,\infty)\}_{v\in V}$ adapted to len, and for any polymatroid capacities $\vec\rho = \{\rho_v\}_{v\in V}$ and demands $\mathsf{dem} : V\times V\to[0,\infty)$, we have

$$\Phi_G(\vec\rho,\mathsf{dem}) \leq \frac{64K\sum_{v\in V}\hat\rho_v(\ell_v)}{\sum_{u,v\in V}\mathsf{dem}(u,v)\cdot\mathbb{E}\,[d_T(F(u),F(v))]}\,. \quad (3)$$

1.4 The embedding theorem

In light of Theorem 1.7, we are able to prove Theorem 1.3 by constructing appropriate random embeddings into trees. In the present section we state our main embedding theorem and give an outline of its proof.

THEOREM 1.8. *There exist constants $K, L \geq 1$ such that the following holds. If $G = (V, E)$ is a metric planar graph, and $F \subseteq V$ is any face of G, then there exists a random tree T and random star-shaped mapping $\Lambda : V \to V(T)$ such that the following conditions hold.*

i) For every $u \in V$ and $\tau > 0$, we have $\mathbb{E}\,|\nabla_\tau \Lambda(u)|_\infty \leq K$.

ii) For every $u, v \in F$,

$$\mathbb{E}\left[d_T(\Lambda(u), \Lambda(v))\right] \geq \frac{d_G(u,v)}{L}\,. \tag{4}$$

Combined with the rounding theorem (Theorem 1.7) and duality (Theorem 1.6), this immediately yields Theorem 1.3 and, in particular, a vertex-capacitated Okamura-Seymour theorem (Theorem 1.2).

PROOF OF THEOREM 1.3. Fix a planar graph $G = (V, E)$, a face $F \subseteq V$ of G, demands $\mathsf{dem} : V \times V \to [0, \infty)$ supported on F, and polymatroid capacities $\vec{\rho}$. By Theorem 1.7, there exists a length function $\mathsf{len} : E \to [0, \infty)$ and a family $\{\ell_v : E(v) \to [0, \infty)\}_{v \in V}$ adapted to len such that

$$\mathsf{mcf}_G(\vec{\rho}, \mathsf{dem}) = \frac{\sum_{v \in V} \hat{\rho}_v(\ell_v)}{\sum_{u,v \in V} \mathsf{dem}(u,v) d_{\mathsf{len}}(u,v)}\,.$$

Consider the metric planar graph $G = (V, E, \mathsf{len})$. By Theorem 1.8 there exist a random tree T and a random star-shaped embedding $\Lambda : V \to T$ satisfying (2) with $K = 1$, and (4) with some universal constant $L > 0$. Applying Theorem 1.7 with Λ, we conclude

$$\begin{aligned}\Phi_G(\vec{\rho}, \mathsf{dem}) &\overset{(3)\wedge(4)}{\leq} \frac{64KL \sum_{v \in V} \hat{\rho}(\ell_v)}{\sum_{u,v \in V} \mathsf{dem}(u,v) d_G(u,v)} \\ &\overset{(1)}{=} 64KL \cdot \mathsf{mcf}_G(\vec{\rho}, \mathsf{dem}).\end{aligned}$$

□

We now give a brief outline of the proof of Theorem 1.8.

First step: Outerplanar graphs into random trees. Theorem 1.8 is proved in two main steps. First, in Section 3, we prove it for the special case of outerplanar graphs; this is precisely the situation where the face F satisfies $F = V$ in Theorem 1.8. It is known that outerplanar graph metrics embed into distributions over dominating trees [12], but this is not sufficient for our purposes; these maps are not star-shaped and do not satisfy the gradient conditions. Instead our proof is inspired by the result of Charikar and Sahai [5] stating that every outerplanar graph metric can be embedded into the product of two trees with $O(1)$ distortion. In particular, each of these two embeddings must be $O(1)$-Lipschitz, so one hopes that the star-shaped and gradient properties might be achievable with their techniques.

Indeed, by following their basic induction and using a heavily modified variant of their embedding, we are able to obtain the desired result. Unfortunately, for this purpose we are not able to obtain a product of two trees; instead we need an entire distribution, but this suffices in light of Theorem 1.7.

Second step: Retracting onto a face. The second step follows the approach of [8] for proving that face metrics (i.e. those metrics arising from taking the shortest-path metric on a planar graph restricted to a face) embed into distributions over dominating trees; this result was originally proved in [19] via a different method. In [8], the authors randomly retract a planar graph $G = (V, E)$ onto a prescribed face $F \subseteq V$ in such a way that edges are not stretched too much in expectation.

Their embedding has the rather convenient property (not shared by previous random retractions) that stars are mapped to stars, satisfying our star-shaped ambitions. Thus we are left to wrestle with the ℓ_∞ gradient issue. By using stronger properties of known random partitioning schemes for planar graphs [15]—specifically the fact that such partitions are "padded" in the language of [11, 16]—we are able to show that all single-scale ℓ_∞ gradients are $O(1)$ in expectation under the random retraction. We remark that this mapping does *not* preserve global ℓ_∞ gradients in expectation, and this is the main reason we have introduced the single-scale definition. This pushes some non-trivial work to the rounding theorem in Section 2 which must now show that all the scales can be rounded simultaneously.

1.5 Preliminaries

Here we review some additional definitions before diving into the proofs. We deal exclusively with finite graphs $G = (V, E)$ which are free of loops and parallel edges. We will also write $V(G)$ and $E(G)$ for the vertex and edge sets of G, respectively. A *metric graph* is a graph G equipped with a non-negative length function on edges $\mathsf{len} : E \to [0, \infty)$. We will denote the pseudometric space associated with a metric graph G as (V, d_G), where d_G is the shortest path metric according to the edge lengths. Note that $d_G(x, y) = 0$ may occur even when $x \neq y$, and if G is disconnected, there will be pairs $x, y \in V$ with $d_G(x, y) = \infty$. We allow both possibilities throughout the paper. An important point is that *all length functions in the paper are assumed to be reduced,* i.e. they satisfy the property that for every $e = (u, v) \in E$, $\mathsf{len}(e) = d_G(u, v)$. For $v \in V$ and $R \geq 0$, we write $B_G(v, R) = \{u \in V : d_G(u, v) \leq R\}$.

In the present paper, paths in graphs are always simple, i.e., no vertex appears twice. Given a metric graph G, we extend the length function to paths $P \subseteq E$ by setting $\mathsf{len}(P) = \sum_{e \in P} \mathsf{len}(e)$. We recall that for a subset $S \subseteq V$, $G[S]$ represents the induced graph on S. For a pair of subsets $S, T \subseteq V$, we use the notations $E(S, T) = \{(u, v) \in E : u \in S, v \in T\}$ and $E(S) = E(S, S)$, and if $v \in V$, we write $E(v) = E(\{v\}, V \setminus \{v\})$.

Given a set X, a *random map* $F : X \to Y$ is shorthand for some probability space (Ω, μ) and a distribution over mappings $\{F_\omega : X \to Y_\omega\}_{\omega \in \Omega}$. Note that both F and Y are random variables. In all our constructions, X and Y_ω are finite sets. When no confusion arises, probabilistic expressions containing F and Y should be understood as been taken over the probability space (Ω, μ). When we refer to a property of Y or F, it should be understood that this property holds *for all* Y_ω and $F_\omega : X \to Y_\omega$, $\omega \in \Omega$.

2. POLYMATROID NETWORKS AND EMBEDDINGS

Our primary goal in the present section is to prove Theorem 1.7 which shows that random tree embeddings can be used to bound flow/cut gaps in polymatroid networks. We start in Section 2.1 by showing that a fixed "thin" mapping into a tree can be use for rounding. In Section 2.2, we prove the crucial property that every star-shaped mapping into a tree can be converted to a random thin map. Finally in Section 2.3, we combine these results with a multi-scale analysis to show that a suitable distribution over star-shaped mappings into random trees suffices for rounding.

2.1 Thin-star tree rounding

Consider a graph G, a connected tree T, and a map $f : V(G) \to V(T)$. For every pair $u, v \in V$, let P_{uv} denote the unique simple path from $f(u)$ to $f(v)$ in T. We say that f is Δ-*thin* if, for every $u \in V(G)$, the induced graph on $\bigcup_{v:\{u,v\}\in E(G)} P_{uv}$ can be covered by Δ simple paths in T emanating from $f(u)$. The next lemma gives a generalization of line-embedding rounding [9, 6] to arbitrary thin maps into trees.

Lemma 2.1. *Let $G = (V, E)$ be a graph, T a connected metric tree, and let $f : V \to V(T)$ be a Δ-thin map. Suppose that the set of functions $\{\ell_v : E(v) \to [0, \infty)\}_{v\in V}$ is such that $d_T(f(u), f(v)) \le \ell_u(e) + \ell_v(e)$ for every edge $e = \{u, v\} \in E$.*

Then for any polymatroid capacities $\vec{\rho} = \{\rho_v\}_{v\in V}$ and demands $\mathsf{dem} : V \times V \to [0, \infty)$, there exists a subset of edges $S \subseteq E$ such that

$$\Phi_G(S; \vec{\rho}, \mathsf{dem}) \le \frac{\Delta \sum_{v\in V} \hat{\rho}_v(\ell_v)}{\sum_{u,v\in V} \mathsf{dem}(u, v) \cdot d_T(f(u), f(v))}\,.$$

Proof. For every edge $\{u, v\} \in E$, let P_{uv} denote the unique simple path between $f(u)$ and $f(v)$ in T. For every $a \in E(T)$, we define the subset $S(a) \subseteq E$ by

$$S(a) = \Big\{\{u, v\} \in E : a \in E(P_{uv})\Big\}\,.$$

Observe that if $a \in E(P_{xy})$ for some $x, y \in V$, then $\sigma_{S(a)}(x, y) = 1$. Thus we have, for any $x, y \in V$,

$$\sum_{a\in E(T)} \mathsf{len}_T(a) \cdot \sigma_{S(a)}(x, y) \ge d_T(f(x), f(y))\,. \tag{5}$$

Next, we give an upper bound on $\nu_{\vec{\rho}}(S(a))$ for every $a \in E(T)$. First, arbitrarily orient the edges of $E(T)$. Fix $a = (x, y) \in E(T)$ according to this orientation. Consider any $\lambda \in [0, \mathsf{len}_T(a)]$. For an edge $e \in S(a)$, choose the orientation $e = (u, v)$ such that P_{uv} traverses a in the order (x, y). We will assign the edge e to the vertex u if

$$d_T(f(u), x) + \lambda \le \ell_u(e)\,, \tag{6}$$

and otherwise assign e to the vertex v. This gives, for every $\lambda \in [0, \mathsf{len}_T(a)]$, a valid assignment $g_{a,\lambda} : S(a) \to V$. Integrating yields

$$\mathsf{len}_T(a) \cdot \nu_{\vec{\rho}}(S(a)) \le \int_0^{\mathsf{len}_T(a)} \left(\sum_{v\in V} \rho_v(g_{a,\lambda}^{-1}(v))\right) d\lambda\,. \tag{7}$$

Our next goal is to show that, for every $v \in V$, we have

$$\sum_{a\in E(T)} \int_0^{\mathsf{len}_T(a)} \rho_v(g_{a,\lambda}^{-1}(v))\, d\lambda \le \Delta\hat{\rho}_v(\ell_v)\,. \tag{8}$$

To this end, fix $v \in V$. Since f is Δ-thin, there are $k \le \Delta$ paths $P_1, P_2, \ldots, P_k$ in T emanating from $f(v) \in V(T)$ such that the following holds: If $S(a)$ contains an edge with endpoint v, then $a \in E(P_i)$ for some $i \in \{1, 2, \ldots, k\}$. Thus we can write

$$\sum_{a\in E(T)} \int_0^{\mathsf{len}_T(a)} \rho_v(g_{a,\lambda}^{-1}(v))\, d\lambda \le \sum_{i=1}^{k} \sum_{a\in E(P_i)} \int_0^{\mathsf{len}_T(a)} \rho_v(g_{a,\lambda}^{-1}(v))\, d\lambda\,, \tag{9}$$

and it suffices to bound each term of the latter sum separately.

To this end, fix $i \in \{1, 2, \ldots, k\}$. For $\theta \in [0, \mathsf{len}(P_i)]$, let

$$S_v(\theta) = \Big\{\{u, v\} \in E : f(u) \in V(P_i) \text{ and } \ell_v(\{u, v\}) \ge \theta\Big\}\,.$$

By the assignment rule (6), the fact that $d_T(f(u), f(v)) \le \ell_u(\{u, v\}) + \ell_v(\{u, v\})$ for every $\{u, v\} \in E$, and monotonicity of ρ_v, we have

$$\begin{aligned}\sum_{a\in E(P_i)} \int_0^{\mathsf{len}_T(a)} \rho_v(g_{a,\lambda}^{-1}(v))\, d\lambda &\le \int_0^{\infty} \rho_v(S_v(\theta))\, d\theta \\ &\le \int_0^{\infty} \rho_v(\ell_v^{\theta})\, d\theta \\ &= \hat{\rho}_v(\ell_v)\,,\end{aligned}$$

where in the final line we have used the definition of the Lovász extension $\hat{\rho}_v$ and the notation: $\ell_v^{\theta}(\{u, v\}) = 1$ if $\ell_v(\{u, v\}) \ge \theta$ and $\ell_v^{\theta}(\{u, v\}) = 0$ otherwise. Combining this with (9) yields (8).

Now interchanging sums and integrals in (8) and summing (7) over $a \in E(T)$ yields

$$\sum_{a\in E(T)} \mathsf{len}_T(a) \cdot \nu_{\vec{\rho}}(S(a)) \le \Delta \sum_{v\in V} \hat{\rho}_v(\ell_v)\,.$$

Using this in conjunction with (5), we have

$$\begin{aligned}&\min_{S\subseteq E} \Phi_G(S; \vec{\rho}, \mathsf{dem}) \\ &\le \min_{a\in E(T)} \frac{\nu_{\vec{\rho}}(S(a))}{\sum_{u,v\in V} \mathsf{dem}(u, v)\sigma_{S(a)}(u, v)} \\ &\le \frac{\sum_{a\in E(T)} \mathsf{len}_T(a) \cdot \nu_{\vec{\rho}}(S(a))}{\sum_{a\in E(T)} \mathsf{len}_T(a) \sum_{u,v\in V} \mathsf{dem}(u, v)\sigma_{S(a)}(u, v)} \\ &\le \frac{\Delta \sum_{v\in V} \hat{\rho}_v(\ell_v)}{\sum_{u,v\in V} \mathsf{dem}(u, v) d_T(f(u), f(v))}\,,\end{aligned}$$

completing the proof. □

2.2 Random thinning

Next we show how an arbitrary star-shaped map into a tree can be converted into a random 4-thin map. Due to lack of space, the proof is deferred to the full version.

Lemma 2.2. *Let $G = (V, E)$ be a graph, T a connected metric tree, and let $f : V \to V(T)$ be a 1-Lipschitz star-shaped map. Then there exists a random connected metric tree T' and a random 4-thin map $F : V \to V(T')$ satisfying the following conditions:*

i) F is 1-Lipschitz with probability one.

ii) For every $u, v \in V$, we have

$$\mathbb{E}\, d_{T'}(F(u), F(v)) \ge \frac{1}{2}\, d_T(f(u), f(v))\,.$$

The next result follows from Lemma 2.2 and Lemma 2.1.

COROLLARY 2.3. *Let $G = (V, E)$ be a graph, T a connected metric tree, and $f : V \to V(T)$ a star-shaped mapping. Suppose that the set of functions $\{\ell_v : E(v) \to [0, \infty)\}_{v \in V}$ is such that $d_T(f(u), f(v)) \le \ell_u(e) + \ell_v(e)$ for every edge $e = \{u, v\} \in E$.*

Then for any polymatroid capacities $\vec{\rho} = \{\rho_v\}_{v \in V}$ and demands $\mathsf{dem} : V \times V \to [0, \infty)$, there exists a subset of edges $S \subseteq E$ such that

$$\Phi_G(S; \vec{\rho}, \mathsf{dem}) \le \frac{8 \sum_{v \in V} \hat{\rho}_v(\ell_v)}{\sum_{u,v \in V} \mathsf{dem}(u, v) \cdot d_T(f(u), f(v))} .$$

2.3 Rounding random star-shaped embeddings

Finally, we are ready prove the main result of this section connecting embeddings to polymatroid flow/cut gaps. We restate Theorem 1.7 here for the sake of the reader

THEOREM 2.4. *Let $G = (V, E, \mathsf{len})$ be a metric graph and suppose there exists a random connected metric tree T and a random star-shaped mapping $F : V \to V(T)$ such that for some $K \ge 1$,*

$$\max_{v \in V} \sup_{\tau > 0} \mathbb{E} \, |\nabla_\tau F(v)|_\infty \le K .$$

Then for any set of functions $\{\ell_v : E(v) \to [0, \infty)\}_{v \in V}$ that is adapted to len, and for any polymatroid capacities $\vec{\rho} = \{\rho_v\}_{v \in V}$ and demands $\mathsf{dem} : V \times V \to [0, \infty)$, there exists a subset of edges $S \subseteq E$ such that

$$\Phi_{\vec{\rho}, \mathsf{dem}}(S) \le \frac{64K \sum_{v \in V} \hat{\rho}_v(\ell_v)}{\sum_{u,v \in V} \mathsf{dem}(u, v) \cdot \mathbb{E} \left[d_T(F(u), F(v)) \right]} .$$

PROOF. Using the fact that $\hat{\rho}_v$ is monotone, we may first scale $\{\ell_v(e) : v \in V, e \in E(v)\}$ down and assume that for $\{u, v\} \in E$, we have $\mathsf{len}(\{u, v\}) = \ell_u(\{u, v\}) + \ell_v(\{u, v\})$. Next, by rounding all the length functions up, we may assume that $\{\ell_v(e) : v \in V, e \in E(v)\}$ are dyadic:

$$\{\ell_v(e) : v \in V, e \in E(v)\} \subseteq \{2^h : h \in \mathbb{Z}\}, \tag{10}$$

and that for $\{u, v\} \in E$, we have

$$\mathsf{len}(\{u, v\}) \ge \frac{1}{2} \Big(\ell_u(\{u, v\}) + \ell_v(\{u, v\}) \Big) . \tag{11}$$

Now define the random functions $\{\tilde{\ell}_v : E(v) \to [0, \infty)\}_{v \in V}$ by

$$\tilde{\ell}_v(\{u, v\}) = \begin{cases} 0 & \text{if } \ell_v(\{u, v\}) < \ell_u(\{u, v\}) \\ 2\ell_v(\{u, v\}) \cdot \frac{d_T(F(u), F(v))}{\mathsf{len}(u, v)} & \text{otherwise.} \end{cases}$$

Then, by definition, we have $d_T(F(u), F(v)) \le \tilde{\ell}_u(\{u, v\}) + \tilde{\ell}_v(\{u, v\})$ for every $\{u, v\} \in E$ since $\{\ell_v\}$ is adapted to len.

We define a new family $\{\hat{\ell}_v\}$ by $\hat{\ell}_v(e) = \sup\{\tilde{\ell}_v(e') : \ell_v(e') \le \ell_v(e)\}$. Observe that $\hat{\ell}_v \ge \tilde{\ell}_v$ pointwise, thus by monotonicity, $\hat{\rho}_v(\hat{\ell}_v) \ge \hat{\rho}_v(\tilde{\ell}_v)$. Additionally, we have $\ell_v(e) \le \ell_v(e')$ if and only if $\hat{\ell}_v(e) \le \hat{\ell}_v(e')$. Thus the collections of edge sets $\{\ell_v^\theta : \theta \in [0, \infty)\}$ and $\{\hat{\ell}_v^\theta : \theta \in [0, \infty)\}$ are identical.

Enumerate the set of values $\{\ell_v(e) : e \in E(v)\} \cup \{0\}$ by $0 = \tau_0 < \tau_1 < \tau_2 < \cdots < \tau_k$ so that

$$\hat{\rho}_v(\ell_v) = \sum_{i=0}^{k-1} (\tau_{i+1} - \tau_i) \rho_v(\ell_v^{\tau_i}) \ge \frac{1}{2} \sum_{i=0}^{k-1} \tau_{i+1} \rho_v(\ell_v^{\tau_i}) , \tag{12}$$

where the latter inequality holds since $\tau_{i+1} \ge 2\tau_i$ by (10).

For $i = 1, 2, \ldots, k$, we can likewise set $\hat{\tau}_i = \max\{\hat{\ell}_v(e) : \ell_v(e) = \tau_i\}$. By construction, we have $0 = \hat{\tau}_0 \le \hat{\tau}_1 \le \hat{\tau}_2 \le \cdots \le \hat{\tau}_k$, and

$$\hat{\rho}_v(\hat{\ell}_v) = \sum_{i=0}^{k-1} (\hat{\tau}_{i+1} - \hat{\tau}_i) \rho_v(\ell_v^{\tau_i}) .$$

Finally, define $\tilde{\tau}_i = \max\{\tilde{\ell}_v(e) : \ell_v(e) = \tau_i\}$. Observe that if $\tilde{\tau}_{i+1} \ne \hat{\tau}_{i+1}$, then $\hat{\tau}_{i+1} = \hat{\tau}_i$, thus we can write

$$\hat{\rho}_v(\hat{\ell}_v) \le \sum_{i=0}^{k-1} \tilde{\tau}_{i+1} \rho_v(\ell_v^{\tau_i}) . \tag{13}$$

Using the definition of $\tilde{\ell}_v$, we have

$$\tilde{\tau}_i = \max \left(\{0\} \cup \left\{ \tilde{\ell}_v(e) : \ell_v(e) = \tau_i \text{ and } \ell_v(e) \ge \frac{1}{2} \mathsf{len}(e) \right\} \right) .$$

Furthermore, by (11), if $\ell_v(e) = \tau_i$, then $\mathsf{len}(e) \ge \frac{1}{2}\tau_i$. Thus,

$$\mathbb{E}[\tilde{\tau}_i] \le \left(\mathbb{E}|\nabla_{\tau_i/2} F(v)|_\infty + \mathbb{E}|\nabla_{\tau_i} F(v)|_\infty \right) \tau_i \le 2K\tau_i .$$

Using (13) and (12), this implies

$$\begin{aligned} \mathbb{E}[\hat{\rho}_v(\hat{\ell}_v)] &\le \sum_{i=0}^{k-1} \mathbb{E}[\tilde{\tau}_{i+1}] \rho_v(\ell_v^{\tau_i}) \\ &\le 2K \sum_{i=0}^{k-1} \tau_{i+1} \rho_v(\ell_v^{\tau_i}) \le 4K \hat{\rho}_v(\ell_v) . \end{aligned}$$

Applying Corollary 2.3 completes the proof. $\square$

3. STAR-SHAPED EMBEDDINGS OF OUTERPLANAR GRAPHS INTO TREES

Our goal is now to prove that every metric outerplanar graph admits a random Lipschitz, star-shaped embedding into a random tree.

THEOREM 3.1. *There is a constant $K \ge 1$ such that the following holds. Let $G = (V, E)$ be a metric outerplanar graph. Then there is a random metric tree T and a random 1-Lipschitz, star-shaped mapping $F : V \to V(T)$ such that for every $u, v \in V$, $\mathbb{E}[d_T(F(u), F(v))] \ge d_G(u, v)/K$.*

We begin by setting up the notations and definitions needed to prove Theorem 3.1.

3.1 Notation and definitions

For a graph $G = (V, E)$, and $v \in V$, we use the notation $N_G(v) = \{u : (u, v) \in E\}$ to denote the set neighbors of v in the graph G. For a path P, we define the cycle $C(P, \ell)$ as the cycle obtained by connecting the endpoints of P with an edge of length ℓ. The length of the cycle C, is given by $\mathsf{len}(C) = \mathsf{len}(P) + \ell$. In this section, it is helpful to think of cycles as continuous cycles and $V(P) \subseteq C$ as points on the cycle.

For a cycle C and a point p on the cycle we define $\mathsf{flat}(C, p)$ to be the path where p is one end point and $x \in C$ is mapped to the point at distance $d_C(p, x)$ from p on the path. Moreover for points $x, y \in C$ we use $d_{\mathsf{flat}(C,p)}(x, y) = |d_C(x, p) - d_C(y, p)|$ to denote the distance between x and y on the path $\mathsf{flat}(C, p)$. See Figure 2 for an example.

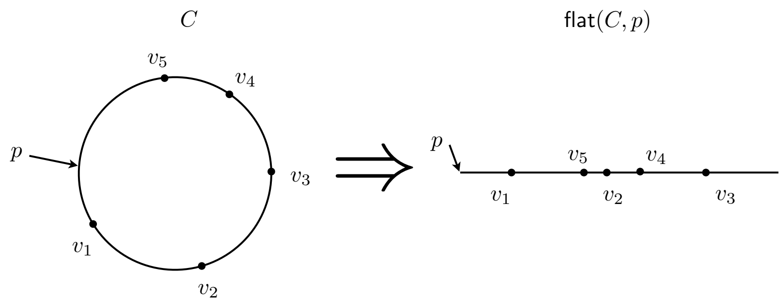

Figure 2: The flattening of the cycle C.

For two paths $P = (u_1, \ldots, u_m)$ and $Q = (v_1, \ldots, v_n)$ with the same length, we define the *glue*ing of P and Q as follows. We first identifying the end points u_1 with v_1, and u_m with v_n to specify the end points of the resulting path. Then we map each point $x \in V(P) \cup V(Q)$ so that the distance between x and the end points of the path is preserved.

Finally, for two given trees, T_1 and T_2, and pairs of vertices $u_1, v_1 \in V(T_1)$ and $u_2, v_2 \in V(T_2)$ such that $d_{T_1}(u_1, v_1) = d_{T_2}(u_2, v_2)$, we define $\mathsf{glue}(T_1, T_2; u_1, v_1; u_2, v_2)$ as the tree resulting from gluing the trees T_1 and T_2 on the unique path between u_1 and v_1 in T_1 and u_2 and v_2 in T_2. See Figure 3.

3.2 Framework

Our approach to Theorem 3.1 employs the framework of Charikar and Sahai (see Theorem 4 in [5]). Any outerplanar graph can be constructed by considering a sequence of paths P_i, and then doing the following: Start with $G_1 = P_1$. At step i, we consider some edge $e_i = (u_i, v_i)$ on the outer face of G_i, and obtain G_{i+1} by either attaching the endpoints of P_i to u_i and v_i, or by attaching only one endpoint of P_i to either u_i or v_i.

In this section we only consider biconnected outerplanar graphs (so the endpoints of P_i are always attached to u_i and v_i), since we can simply take the embedding of biconnected components of a graph that are connected by a single vertex into trees, and glue the trees on the image of the common vertex to obtain an embedding for the whole graph.

We also use the concept of a slack structure [12]. We say that an outerplanar graph has an α-slack structure if it can be built out of paths P_i such that the length of any path P_i which attaches to both endpoints of an edge e_i is at least α times the length of e_i. The following lemma is a straightforward generalization of a fact from [12], where it is proved for $\alpha = 2$.

Lemma 3.2 ([12]). *Consider any $\alpha \geq 1$. Given an outerplanar metric graph $G = (V, E, \mathsf{len}_G)$, there is an outerplanar metric graph $H = (V, E', \mathsf{len}_H)$ with $E' \subseteq E$, and such that H has an α-slack structure. Furthermore, $d_G \geq d_H \geq (1/\alpha) d_G$, and for every $(u, v) \in E'$,*

$$d_H(u, v) = \mathsf{len}_H(u, v). \qquad (14)$$

Thus, by incurring distortion at most α, we may assume that the outerplanar graph G has an α-slack structure. We will build our embedding inductively based on the sequence of the paths $P_1 \ldots, P_m$ provided by Lemma 3.2.

Random extension. Given an embedding of a metric graph G into a random metric tree T, $F : V(G) \to V(T)$ and a new path P attached to the points $u, v \in V(G)$, we extend the embedding of G to an embedding for $\hat{G} = G \cup P$ into a random tree $\hat{T}$, using the following operation. Let $C = C(P, d_T(F(u), F(v)))$. To extend the tree T, we choose two *anchor* points $p, q \in C$, and map the vertices of C onto two paths $L = \mathsf{flat}(C, p)$ and $R = \mathsf{flat}(C, q)$.

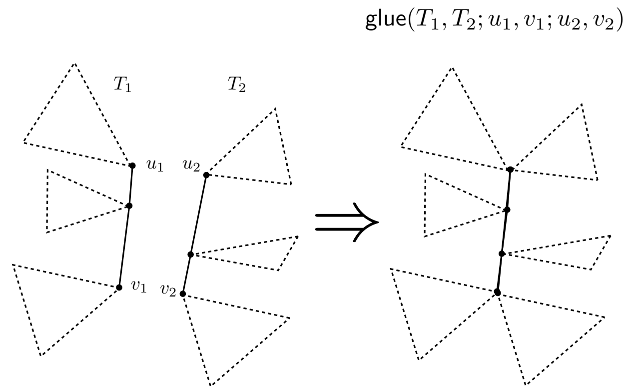

Figure 3: Gluing of the trees T_1 and T_2.

We put $\hat{T} = \mathsf{glue}(T, L; F(u), F(v); u, v)$ with probability 1/2 and $\hat{T} = \mathsf{glue}(T, R; F(u), F(v); u, v)$ with probability 1/2. This specifies a random mapping $\hat{F} : V(G) \to V(\hat{T})$. Since it will be clear from context which vertices we are gluing onto, we will use the notations $\mathsf{glue}(T, L)$ and $\mathsf{glue}(T, R)$ without specifying the vertices. Note that the gluing can be done if and only if $d_L(u, v) = d_R(u, v) = d_C(u, v) = d_T(F(u), F(v))$. Moreover, If the map $F : V(G) \to V(T)$ is 1-Lipschitz, then so is the extension $\hat{F}$.

A significant difference between our construction and that of [5] is in the way we choose the anchor points. For our purposes, it is not enough to simply look at the α-slack graph; we need to use the structure of the original graph when we choose the anchor points in order to maintain the star-shaped property. The algorithm of [5] is able to construct an embedding using only two trees, while we embed the graph into a distribution over trees. In the next section, we present a distortion bound for this embedding based on the distance between the anchor points in the cycle.

3.3 Bounding the distortion

Before we can state the main lemma of this section, we need the following definition. For a cycle C, and points

$u, v \in C$ we say that a pair of points $p, q \in C$ is (α, β)*-apart with respect to another pair* $\{u, v\}$ if $d_C(p, q) = \alpha \mathsf{len}(C)$ and for $a \in \{u, v\}$ and $b \in \{p, q\}$:

$$\beta \, \mathsf{len}(C) \le d_C(a, b) \le \left(\frac{1}{2} - \beta\right) \mathsf{len}(C).$$

We now state a lemma capturing our main inductive step. Due to lack of space, we defer its proof to the full version.

Lemma 3.3. *Let G be a graph, T be random metric tree, and let $F : V(G) \to V(T)$ be a random 1-Lipschitz map such that $\mathbb{E}[d_T(F(x), F(y)] \ge d_G(x, y)/6$ for every $x, y \in V(G)$. Let $\hat{G}$ be a graph constructed by attaching a path P with*

$$\mathsf{len}(P) \ge 160 \cdot d_G(u, v) \tag{15}$$

onto a pair of vertices $u, v \in V(G)$. Let $C = C(P, d_T(F(u), F(v)))$, and $p, q \in C$ be any pair of points that are $(1/6, 1/16)$-apart with respect to $\{u, v\}$ in C, and let $\hat{T}$ be the random extension of T by C with respect to the anchor points p and q. The embedding $\hat{F} : V(\hat{G}) \to V(\hat{T})$ is also 1-Lipschitz and such that for all $x, y \in V(\hat{G})$, $\mathbb{E}[d_{\hat{T}}(\hat{F}(x), \hat{F}(y)] \ge d_{\hat{G}}(x, y)/6$.

3.4 The star-shaped property

To prove Theorem 3.1, we need to choose the anchor points such that the resulting map is star-shaped. Due to lack of space, the argument is deferred to the full version.

4. CONNECTED RANDOM RETRACTIONS

Our goal now is to complete the proof of Theorem 1.8 by showing that every planar graph can be randomly retracted onto a specified face in such a way that the face can itself be endowed with an outerplanar metric. Combining this with our embedding of outerplanar graphs into random trees from Section 3, we will be able to prove Theorem 1.8; this is done in Section 4.3.

In the next section, we review the notion of "padded partitions" of metric spaces. The existence of such partitions for planar graphs (due to [15]) will be one of our two central ingredients here. The other ingredient is the method of [8] for the construction of random *connected* retractions. They work with a weaker notion of random partitions, so their results (as stated in [8]) are not strong enough for us. In Section 4.2, we follow their proof closely but use padded partitions, allowing us to obtain the stronger conclusion we require.

4.1 Padded partitions of graphs

Random partitions are a powerful tool in the theory of embeddings of finite metric spaces; see, e.g., [1, 24, 16, 18]. A particularly powerful notion is that of a "padded" partition. We review the relevant definitions in the special setting of finite metric spaces.

Consider a metric space (X, d). We will sometimes think of a partition P of X as a map $P : X \to 2^X$ sending each $x \in X$ to the unique set in P containing it. We say that P is τ*-bounded* if $\mathrm{diam}(S) \le \tau$ for every $S \in P$. We say that a random partition $\mathcal{P}$ is τ-bounded if this holds almost surely. A random partition $\mathcal{P}$ is (α, τ)*-padded* if it is τ-bounded and, additionally, for every $x \in X$ and $R \ge 0$, we have

$$\mathbb{P}[B_X(x, R) \not\subseteq \mathcal{P}(x)] \le \alpha \cdot \frac{R}{\tau},$$

where $B_X(x, R) = \{y \in X : d(x, y) \le R\}$.

The main random partitioning result we require is from [15], though it first appeared in this form later (see [24, 16, 18]).

Theorem 4.1 ([15]). *There exists a constant $\alpha > 0$ such that if $G = (V, E)$ is a metric planar graph, then for every $\tau > 0$, (V, d_G) admits an (α, τ)-padded random partition. Furthermore, the distribution of the partition can be sampled from in polynomial-time in the size of G.*

4.2 Random retractions

We now use random partitions to construct random retractions. This was first done in [4] in the context of the 0-extension problem on graphs. Further work includes [18], which concerns the Lipschitz extension problem, and [19], where the authors are primarily concerned with randomly simplifying the topology of metric graphs. The proof of the next theorem follows from the techniques of [8] for constructing a connected retraction. We are able to obtain a stronger conclusion by using a stronger assumption about the random partitions. The proof itself is deferred to the full version.

Theorem 4.2. *Let $G = (V, E)$ be a metric graph and suppose that for some $\alpha \ge 2$ and every $\tau \ge 0$, (V, d_G) admits an (α, τ)-padded random partition. Then for any subset $S \subseteq V$, there exists a random mapping $F : V \to S$ such that the following properties hold.*

i) For every $x \in S$, $F(x) = x$.

ii) For every $x \in V$ and $\tau > 0$, $\mathbb{E}\, |\nabla_\tau F(x)|_\infty \le O(\alpha \log \alpha)$.

iii) For every $x \in S$, the set $F^{-1}(x)$ is a connected subset of G.

4.3 Retracting to an outerplanar graph

Finally, we use the random retractions of the preceding section to randomly embed every metric on the face of a planar graph into an outerplanar graph in a suitable way. This technique is also taken from [8], although again we require some stronger properties of the embedding.

Theorem 4.3. *There is a constant $K > 1$ such that for any metric planar graph $G = (V, E)$ and face $V_0 \subseteq V$, there is a random outerplanar metric graph H and a random mapping $F : V \to V(H)$ satisfying the following:*

i) For every edge $\{u, v\} \in E$, either $F(u) = F(v)$ or $\{F(u), F(v)\} \in E(H)$.

ii) For every $u, v \in V_0$, $d_H(F(u), F(v)) \ge d_G(u, v)$.

iii) For every $u \in V$ and $\tau \ge 0$, we have $\mathbb{E}\, |\nabla_\tau F(u)|_\infty \le K$.

Proof. By Theorem 4.1, we can apply Theorem 4.2 to the metric graph G with $S = V_0$. Let $F : V \to V_0$ be the random mapping guaranteed by Theorem 4.2.

We construct a metric graph H with vertex set V_0 and an edge $\{u, v\}$ of length $d_G(u, v)$ whenever there is an edge between the sets $F^{-1}(u)$ and $F^{-1}(v)$ in G. Since the sets $\{F^{-1}(u) : u \in V_0\}$ are connected, the resulting graph H is outerplanar. Also, property (i) is immediate.

Property (ii) follows because F is the identity on V_0 and the edges in H have length equal to the distance between their endpoints in (V_0, d_G). Property (iii) follows from Theorem 4.2(i). □

We can now complete the proof of Theorem 1.8.

Proof of Theorem 1.8. Let $\Lambda_1 : V \to V(H)$ be the random mapping from V onto the vertices of an outerplanar metric graph H guaranteed from Theorem 4.3. Let $\Lambda_2 : V(H) \to V(T)$ be the random mapping of H into trees from Theorem 3.1. The mapping $\Lambda = \Lambda_2 \circ \Lambda_1 : V \to V(T)$ is mapping guaranteed by the theorem. Combining the star-shaped property of Theorem 3.1 with Theorem 4.3(i) implies that Λ is star-shaped. Property (ii) of Theorem 1.8 is a consequence of Theorem 4.3(ii) and Theorem 3.1. Finally, property (i) follows from the Lipschitz condition of Theorem 3.1 and property (iii) of Theorem 4.3. $\square$

Acknowledgements

We are grateful to Chandra Chekuri and Bruce Shepherd for bringing the node-capacitated problem to our attention and for their encouragement during the course of the project.

5. REFERENCES

[1] Y. Bartal. On approximating arbitrary metrices by tree metrics. In *STOC '98 (Dallas, TX)*, pages 161–168. ACM, New York, 1998.

[2] J. Bourgain. The metrical interpretation of superreflexivity in Banach spaces. *Israel J. Math.*, 56(2):222–230, 1986.

[3] B. Brinkman, A. Karagiozova, and J. R. Lee. Vertex cuts, random walks, and dimension reduction in series-parallel graphs. In *STOC'07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 621–630. ACM, New York, 2007.

[4] G. Calinescu, H. Karloff, and Y. Rabani. Approximation algorithms for the 0-extension problem. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 8–16, Philadelphia, PA, 2001. SIAM.

[5] M. Charikar and A. Sahai. Dimension reduction in the ℓ_1 norm. In *FOCS*, pages 551–560, 2002.

[6] C. Chekuri, S. Kannan, A. Raja, and P. Viswanath. Multicommodity flows and cuts in polymatroidal networks. In *ITCS*, pages 399–408, 2012.

[7] C. Chekuri, S. Khanna, and F. B. Shepherd. Edge-disjoint paths in planar graphs with constant congestion. *SIAM J. Comput.*, 39(1):281–301, 2009.

[8] M. Englert, A. Gupta, R. Krauthgamer, H. Räcke, I. Talgam-Cohen, and K. Talwar. Vertex sparsifiers: New results from old techniques. In *APPROX-RANDOM*, pages 152–165, 2010.

[9] U. Feige, M. Hajiaghayi, and J. R. Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM J. Comput.*, 38(2):629–657, 2008. Prelim. version in *STOC 2005*.

[10] L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.

[11] A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *FOCS*, pages 534–543, 2003.

[12] A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair. Cuts, trees and l_1-embeddings of graphs. *Combinatorica*, 24(2):233–269, 2004.

[13] R. Hassin. Minimum cost flow with set-constraints. *Networks*, 12(1):1–21, 1982.

[14] T. C. Hu. Multi-commodity network flows. *J. ORSA*, 11:344–360, 1963.

[15] P. N. Klein, S. A. Plotkin, and S. Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pages 682–690, 1993.

[16] R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: a new embedding method for finite metrics. *Geom. Funct. Anal.*, 15(4):839–858, 2005.

[17] E. L. Lawler and C. U. Martel. Computing maximal "polymatroidal" network flows. *Math. Oper. Res.*, 7(3):334–347, 1982.

[18] J. R. Lee and A. Naor. Extending Lipschitz functions via random metric partitions. *Invent. Math.*, 160(1):59–95, 2005.

[19] J. R. Lee and A. Sidiropoulos. On the geometry of graphs with a forbidden minor. In *STOC*, pages 245–254, 2009.

[20] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.

[21] J. Matoušek and Y. Rabinovich. On dominated l_1 metrics. *Israel J. Math.*, 123:285–301, 2001.

[22] K. Menger. Zur allgemeinen kurventheorie. *Fund. Math.*, 10:96–116, 1927.

[23] H. Okamura and P. D. Seymour. Multicommodity flows in planar graphs. *J. Combin. Theory Ser. B*, 31(1):75–81, 1981.

[24] S. Rao. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In *Proceedings of the 15th Annual Symposium on Computational Geometry*, pages 300–306, New York, 1999. ACM.