Genus and the geometry of the cut graph

[extended abstract]

James R. Lee* Anastasios Sidiropoulos†

Abstract

We study the quantitative geometry of graphs in terms of their genus, using the structure of certain "cut graphs," i.e. subgraphs whose removal leaves a planar graph. In particular, we give optimal bounds for random partitioning schemes, as well as various types of embeddings. Using these geometric primitives, we present exponentially improved dependence on genus for a number of problems like approximate max-flow/min-cut theorems, approximations for uniform and non-uniform Sparsest Cut, treewidth approximation, Laplacian eigenvalue bounds, and Lipschitz extension theorems and related metric labeling problems.

We list here a sample of these improvements. All the following statements refer to graphs of genus g, unless otherwise noted.

- We show that such graphs admit an $O(\log g)$-approximate multi-commodity max-flow/min-cut theorem for the case of uniform demands. This bound is optimal, and improves over the previous bound of $O(g)$ [KPR93, FT03]. For general demands, we show that the worst possible gap is $O(\log g + C_P)$, where C_P is the gap for planar graphs. This dependence is optimal, and already yields a bound of $O(\log g + \sqrt{\log n})$, improving over the previous bound of $O(\sqrt{g \log n})$ [KLMN04].
- We give an $O(\sqrt{\log g})$-approximation for the uniform Sparsest Cut, balanced vertex separator, and treewidth problems, improving over the previous bound of $O(g)$ [FHL05].
- If a graph G has genus g and maximum degree D, we show that the kth Laplacian eigenvalue of G is $(\log g)^2 \cdot O(kgD/n)$, improving over the previous bound of $g^2 \cdot O(kgD/n)$ [KLPT09]. There is a lower bound of $\Omega(kgD/n)$, making this result almost tight.
- We show that if (X, d) is the shortest-path metric on a graph of genus g and $S \subseteq X$, then every L-Lipschitz map $f : S \to Z$ into a Banach space Z admits an $O(L \log g)$-Lipschitz extension $\tilde{f} : X \to Z$. This improves over the previous bound of $O(Lg)$ [LN05], and compares to a lower bound of $\Omega(L\sqrt{\log g})$. In a related way, we show that there is an $O(\log g)$-approximation for the 0-extension problem on such graphs, improving over the previous $O(g)$ bound.
- We show that every n-vertex shortest-path metric on a graph of genus g embeds into L_2 with distortion $O(\log g + \sqrt{\log n})$, improving over the previous bound of $O(\sqrt{g \log n})$. Our result is asymptotically optimal for every dependence $g = g(n)$.

1 Introduction

The geometry of finite metric spaces plays a fundamental role in a number of areas of graph theory, and in the study of approximation algorithms for NP-hard problems. Unlike classical metric spaces (e.g. the Euclidean space $\mathbb{R}^d$), these spaces often do not come equipped with a natural measure that interacts nicely with the geometry. To combat this lack of structure, the use of randomness becomes a fundamental tool.

One of the most prominent techniques involves *random low-diameter partitions.* For instance, one considers a random partition of some metric space (X, d) into pieces of diameter at most $\Delta > 0$, with the property that nearby points are rarely placed into different sets, e.g. with probability at most $\kappa \cdot d(x, y)/\Delta$ for some parameter $\kappa > 0$ that measures the quality of the random partition.

Such constructs have been fundamental in areas like probabilistic embeddings into trees [Bar96, Bar98, FRT03] (with a host of applications to approximation and online algorithms), Sparsest Cut and its variants [KPR93, Rao99, ARV04, CGR05, FHL05], geometric representations of graphs, e.g. [KL03, GKL03, Rab03, KLMN04, CDG+09, Lee09], the Lipschitz extension [LN05] and 0-extension problems [CKR01, FHRT03, AFH+04, LN04], proximity data structures

*Computer Science & Engineering, University of Washington. Research partially supported by NSF grant CCF-0644037 and a Sloan Research Fellowship. E-mail: jrl@cs.washington.edu.

†Toyota Technological Institute at Chicago. E-mail: tasos@tti-c.org.

[CGMZ05, MN07, AGMW07], eigenvalue bounds for graphs [BLR08, KLPT09], and a host of other approximation algorithms (e.g. for TSP [Tal04] and Unique Games [CMM06]).

One of the most widely used partition schemes arises from work of Klein, Plotkin, and Rao [KPR93], as interpreted by [Rao99], and quantitatively improved in [FT03]. It is shown that if (X, d) is the shortest-path metric on a graph which excludes K_h (the complete graph on h vertices) as a minor, then one can take $\kappa = O(h^2)$ above. In particular, for graphs of genus g, one has $\kappa = O(g)$. In the present work, we give a construction that achieves $\kappa = O(\log g)$, which is the optimal dependence. Speaking to the power of such random partitions, this has a number of applications to, e.g. approximate max-flow/min-cut theorems, approximations for uniform and non-uniform Sparsest Cut, treewidth approximation, Laplacian eigenvalue bounds, Lipschitz extension theorems, and various embeddings into normed spaces. These applications are discussed in detail in Section 4.1.

In this extended abstract, we focus mainly on random partitions schemes. We discuss some additional geometric consequences of our work in Section 4, including a variety of optimal embedding results. The proofs of these theorems, as well as a full account of their consequences are deferred to final paper.

We now present some preliminary definitions on metric spaces. The reader comfortable with this terminology can skip to the technical overview in Section 1.2.

1.1 Preliminaries

Throughout the paper, all metric spaces are assumed to be finite unless otherwise stated.

Graphs and metrics. We deal exclusively with finite graphs $G = (V, E)$ which are free of loops and parallel edges. We will also write $V(G)$ and $E(G)$ for the vertex and edge sets of G, respectively. A *metric graph* is a graph G equipped with a non-negative length function on edges $\mathsf{len} : E \to \mathbb{R}_+$. We will denote the metric space associated with a graph G as (V, d_G), where d_G is the shortest path metric according to the edge lengths. Note that $d_G(x, y) = 0$ may occur even when $x \neq y$, and also 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 $e = (u, v) \in E$, $\mathsf{len}(e) = d_G(u, v)$.

Given a metric graph G, we extend the length function to paths $P \subseteq E$ by setting $\mathsf{len}(P) = \sum_{e\in P} \mathsf{len}(e)$. For a pair of vertices $a, b \in P$, we use the notation $P[a, b]$ to denote the sub-path of P from a to b. 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 $E(S) = E(S, S)$. For a vertex $u \in V$, we write $N(u) = \{v \in V : (u, v) \in E\}$.

For a metric space (X, d), a number $R \geq 0$, and a point $x \in X$, we use

$$B_{(X,d)}(x, R) = \{y \in X : d(x, y) \leq R\}$$

to denote the closed ball of radius R about x in X. If the metric space is clear from context, we sometimes just write $B(x, R)$. For a subset $S \subseteq X$, we write $\mathrm{diam}_{(X,d)}(S) = \max_{x,y\in S} d(x, y)$. Again, we omit the subscript if the metric and/or space are clear from context. Finally, for sets $S, T \subseteq X$ and a point $x \in X$, we use the notations $d(x, S) = \min_{y\in S} d(x, y)$ and $d(T, S) = \min_{x\in T} d(x, S)$.

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

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

If f is injective, then the *distortion of* f is defined by $\mathsf{dist}(f) = \|f\|_{\mathrm{Lip}} \cdot \|f^{-1}\|_{\mathrm{Lip}}$. If $d_Y(f(x), f(y)) \leq d_X(x, y)$ for every $x, y \in X$, we say that f is *non-expansive.* For $p \geq 1$, we write $c_p(X, d_X)$ for the infimal distortion over all maps from X into some L_p space.

1.2 Technical overview

Let $G = (V, E)$ be a metric graph of genus g (i.e. which can be drawn on an orientable surface of genus g without edge crossings). Given a parameter $\Delta > 0$, our goal is to produce a random partition P of V into sets of diameter at most Δ, with the property that for every $x \in V$, the probability that $B(x, \Delta/\alpha)$ is all mapped to the same set of P is $\Omega(1)$. Here, α is some parameter; our goal is to achieve $\alpha = O(\log g)$. We refer to Section 2 for a formal discussion of random partitions.

In Section 2.1, we develop some generic primitives for combining and manipulating these kinds of partitions. In particular, we show that if there is a subgraph H of G such that if both $(V(H), d_G)$ and the *induced path metric* on $G \setminus H$ admit good random partitions, then so does (V, d_G). Note that the induced path metric on $G\setminus H$ is *not* the same as d_G restricted to $V(G\setminus H)$.

This immediately suggests an approach: Find a subgraph H so that $G \setminus H$ is planar, and $(V(H), d_G)$ admits good random partitions. Then the planar portion $G \setminus H$ can be dealt with using [KPR93]. We do this in Section 3 by letting H consist of a system of cycles in G, each of which is a shortest representative from its homotopy class. It is easy to deduce that H itself is a union of $O(g)$ shortest paths in G. Since each shortest path, being isometric to a subset of the real line, admits good random partitions, it suffices to show that we can achieve $\alpha = O(\log g)$ for the union of such objects. This is done in Lemma 2.7 using a modification of the random partitioning algorithm of [CKR01]. The final result appears in Theorem 4.1.

In Section 4.1, we discuss various applications of this scheme, as well as applications of the additional embedding theorems stated in 4.

2 Random partitions

In the present section, we will provide some basic primitives for random partitions of metric spaces. First, we introduce the three types of random partitions that we will be most interested in. Fix a metric space (X, d). If P is a partition of X, we will often think of it is as a mapping $P : X \to 2^X$ which sends a point $x \in X$ to the set $P(x) \in P$ with $x \in P(x)$.

Padded partitions. Let P be a random partition of X. We say that P is (ρ, δ, Δ)-*padded* if the following conditions hold.

1. For all $Q \in \mathrm{supp}(P)$, for all $S \in P$, we have $\mathrm{diam}(S) \leq \Delta$.
2. For every $x \in X$,

$$\Pr[B(x, \rho) \subseteq P(x)] \geq \delta.$$

We write $\alpha(X, d; \delta, \Delta)$ for the infimal constant $\alpha \geq 1$ such that (X, d) admits a $(\Delta/\alpha, \delta, \Delta)$-padded random partition.

Lipschitz partitions. Let P be a random partition of X. We say that P is (L, Δ)-*Lipschitz* if the following conditions hold.

1. For all $Q \in \mathrm{supp}(P)$, for all $S \in P$, we have $\mathrm{diam}(S) \leq \Delta$.
2. For every $x, y \in X$,

$$\Pr[P(x) \neq P(y)] \leq L \cdot d(x, y).$$

We write $\beta(X, d; \Delta)$ for the infimal constant $\beta \geq 1$ such that (X, d) admits a $(\beta/\Delta, \Delta)$-Lipschitz random partition.

Spreading partitions. We say that (X, d) admits a (ρ, δ, Δ)-*spreading partition* if there is a random partition $X = Z \cup \bar{Z}$ such that for all $x, y \in X$ satisfying $d(x, y) > \Delta$, we have

$$\Pr[x \in Z \text{ and } d(y, Z) \geq \rho] \geq \delta.$$

We write $\zeta(X, d; \delta, \Delta)$ for the infimal constant $\zeta \geq 1$ such that (X, d) admits a $(\Delta/\zeta, \delta, \Delta)$-spreading partition.

We recall some known theorems on the relationships between such partitions. The next two lemmas are from [LN04]. The first shows that padded partitions yield Lipschitz partitions of similar parameters.

Lemma 2.1. (Padded to Lipschitz) *If (X, d) admits a (ρ, δ, Δ)-padded partition, then it also admits an $(4\delta\rho, 2\Delta)$-Lipschitz partition.*

Lemma 2.2. *For every $\delta > 0, \Delta > 0$, there exists a constant $C \geq 1$ and a $d \geq 1$ such that $(\mathbb{R}^d, \|\cdot\|_2)$ does not admit a $(C\Delta/d, \delta, \Delta)$-padded partition.*

On the other hand, we have the following result from [CCGG98] which, when combined with the preceding lemma, shows that padded partitions are, in some sense, strictly stronger than Lipschitz partitions.

Lemma 2.3. *For every $\Delta > 0$ and $d \geq 1$, there exists a $(\Delta/O(\sqrt{d}), \Delta)$-Lipschitz partition of $(\mathbb{R}^d, \|\cdot\|_2)$.*

Lemma 2.4. (Padded to spreading) *If (X, d) admits a (ρ, δ, Δ)-padded partition, then it also admits a $(\rho, \delta/4, \Delta)$-spreading partition.*

Proof. Let $P : X \to 2^X$ be a (ρ, δ, Δ)-padded partition. For every $C \in P$, let $\sigma_C \in \{0, 1\}$ be an independent, uniformly distributed random variable. Define a random subset $Z \subseteq X$ by

$$Z = \{x \in X : \sigma_{P(x)} = 0\}.$$

We claim that $Z \cup \bar{Z}$ is a $(\rho, \delta/4, \Delta)$-spreading partition. To see this, consider $x, y \in X$ with $d(x, y) > \Delta$. Then if the two independent events $\{\sigma_{P(x)} = 0\}$ and $\{\sigma_{P(y)} = 1\}$ occur as well as the event $\{B(y, \rho) \subseteq P(y)\}$, then $x \in Z$ and $d(y, Z) \geq \rho$. Clearly the probability of these three events occurring simultaneously is at least $\delta \cdot (1/2)^2$, completing the proof.

We now list some known results about the existence of such partitions. The first is an elementary observation.

Lemma 2.5. *If $X \subseteq \mathbb{R}$, then for any $\delta \in (0, 1), \Delta > 0$, $\alpha(X, |\cdot|; \delta, \Delta) \leq 2/(1 - \delta)$.*

The next theorem follows from the work of Klein, Plotkin, and Rao [KPR93] as envisioned in [Rao99], with the listed quantitative bound due to [FT03].

Theorem 2.1. *If $G = (V, E)$ is a graph metric and G excludes K_r as a minor, then for every $\Delta > 0$, $\alpha(V, d_G; \frac{1}{2}, \Delta) = O(r^2)$. In particular, if G has genus g, then $\alpha(V, d_G; \frac{1}{2}, \Delta) = O(g)$.*

The next theorem is due to Bartal [Bar96] (see also [CKR01] for an alternate proof).

THEOREM 2.2. *For any finite metric space* (X,d) *and any* $\Delta > 0$, $\alpha(X, d; \frac{1}{2}, \Delta) = O(\log |X|)$.

Finally, we state the following theorem of Lee [Lee05], which relies heavily on the work of Arora, Rao, and Vazirani [ARV04].

THEOREM 2.3. *There exists a constant* $\delta > 0$ *such that the following holds for every* $\Delta > 0$. *If* (X,d) *is a finite metric space for which* $(X, \sqrt{d})$ *admits a* C*-bi-Lipschitz embedding into a Hilbert space, then* $\zeta(X, d; \delta, \Delta) = O(C^2 \sqrt{\log |X|})$.

2.1 Padded partitions The following lemmas refer to a metric space (X,d). For a subset $S \subseteq X$, we use $N_\beta(S) = \{x \in X : d(x,S) \le \beta\}$ to denote the β-neighborhood of S in X.

LEMMA 2.6. (EXTENSION TO NEIGHBORHOODS) *For every* $\delta, \Delta > 0$ *and* $S \subseteq X$, *the following holds. If* (S,d) *admits a* (ρ, δ, Δ)*-padded partition, then* $(N_{\rho/4}(S), d)$ *admits a* $(\rho/2, \delta, \Delta + \rho/2)$*-padded partition.*

Proof. Let $\widehat{S} = N_{\rho/4}(S)$. Let P be a (ρ, δ, Δ)-padded partition of (S,d). Let $\Gamma : X \to S$ be chosen so that $d(x, \Gamma(x)) = d(x, S)$ for every $x \in X$. Construct a random partition $\widehat{P}$ of $\widehat{S}$ by defining, for every $C \in P$, $\widehat{C} = \{x \in \widehat{S} : \Gamma(x) \in C\}$. It is clear that $\{\widehat{C}\}_{C \in P}$ is a partition of $\widehat{S}$, and for any $C \in P$,

$$\mathrm{diam}(\widehat{C}) \le \mathrm{diam}(C) + 2(\rho/4) \le \Delta + \rho/2.$$

Now, consider any $x \in \widehat{S}$. If $B_S(\Gamma(x), \rho) \subseteq P(x)$, we claim that $B_{\widehat{S}}(x, \rho/2) \subseteq \widehat{P}(x)$. This follows because if $y \in \widehat{S}$ is such that $d(x,y) \le \rho/2$, then

$$d(\Gamma(x), \Gamma(y)) \le d(x, \Gamma(x)) + d(x,y) + d(y, \Gamma(y)) \le \rho,$$

implying that $\widehat{P}(x) = \widehat{P}(y)$ under our assumption. Thus,

$$\Pr[B_{\widehat{S}}(x, \rho/2) \subseteq \widehat{P}(x)] \ge \Pr[B_S(\Gamma(x), \rho) \subseteq P(x)] \ge \delta.$$

LEMMA 2.7. (COMPOSITION FOR FINITE COVERINGS) *Let* $C_1, C_2, \ldots, C_k \subseteq X$ *and* $\delta, \Delta > 0$ *be arbitrary. Suppose that for each* $i \in [k]$, (C_i, d) *admits a* (ρ, δ, Δ)*-padded random partition* P_i. *If* $S = C_1 \cup C_2 \cup \cdots \cup C_k$, *then* (S,d) *admits a* $(\rho/O(\log k), \delta/2, \Delta + \rho/2)$*-padded partition.*

Proof. Let π be a random permutation of $\{1, 2, \ldots, k\}$, and let $\alpha \in [\frac{1}{2}, 1]$ be chosen uniformly at random. For $i = 1, 2, \ldots, k$ define

$$S_i = \left\{x : d(x, C_{\pi(i)}) \le \alpha\rho/4\right\} \setminus \{S_1 \cup \cdots \cup S_{i-1}\}$$

It is clear that $\{S_1, \ldots, S_k\}$ forms a partition of S, which we call P_0.

Now, by definition $S_i \subseteq N_{\rho/4}(C_{\pi(i)})$, hence applying Lemma 2.6 to $P_{\pi(i)}$ yields a $(\rho/2, \delta, \Delta + \rho/2)$-padded partition $\widehat{P}_i$ of S_i. Our final partition of S is $P = \bigcup_{i=1}^k \widehat{P}_i$. We proceed to prove that P is $(\rho/O(\log k), \frac{1}{2}\delta, \Delta + \rho/2)$-padded. The diameter bound is clear.

Fixing $x \in S$ and $R \le \rho/2$, we have

$$\begin{aligned}
\Pr \quad & [B_S(x,R) \subseteq P(x)] \\
& \ge \Pr[B_S(x,R) \subseteq P_0(x)] \\
& \quad \cdot \Pr[B_S(x,R) \subseteq P(x) \mid B_S(x,R) \subseteq P_0(x)] \\
& \ge \delta \cdot \Pr[B_S(x,R) \subseteq P_0(x)],
\end{aligned}$$

using the fact that each $\widehat{P}_i$ is $(\rho/2, \delta, \Delta + \rho/2)$-padded. Thus it suffices to prove that $\Pr[B_S(x, \rho/(C \log k)) \subseteq P_0(x)] \ge \frac{1}{2}$ for some sufficiently large constant $C \ge 1$. The analysis follows [CKR01].

Order the sets $C_{i_1}, C_{i_2}, \ldots, C_{i_k}$ in increasing order of their distance from x, i.e. so that $d(x, C_{i_j}) \le d(x, C_{i_{j+1}})$ for $j = 1, 2, \ldots, k-1$. Let $I_j = [d(x, C_{i_j}) - R, d(x, C_{i_j}) + R]$. Write $\mathcal{E}_j$ for the event that $\alpha\rho/4 \le d(x, C_{i_j}) + R$ and i_j is the minimal element, according to the total order on $\{1, 2, \ldots, k\}$ induced by π, for which $\alpha\rho/4 \ge d(x, C_{i_j}) - R$. It is not difficult to see that the event $\{B_S(x,t) \not\subseteq P_0(x)\}$ is contained in the events $\bigcup_{j=1}^k \mathcal{E}_j$, hence

$$\begin{aligned}
\Pr[B_S(x,t) \not\subseteq P_0(x)] \quad & \le \quad \sum_{j=1}^k \Pr[\mathcal{E}_j] \\
& = \quad \sum_{j=1}^k \Pr[\alpha\rho/4 \in I_j] \cdot \Pr[\mathcal{E}_j \mid \alpha\rho/4 \in I_j] \\
& \le \quad \sum_{j=1}^k \frac{2R}{\rho/8} \cdot \frac{1}{j} \le O(\log k)\frac{R}{\rho},
\end{aligned}$$

where we have used the fact that, conditioned on $\alpha\rho/4 \in I_j$, $\mathcal{E}_j$ only occurs if i_j is the minimal index (according to π) among $\{i_1, \ldots, i_j\}$. It is now clear that we can choose $R \ge \rho/(C \log k)$ (for a suitable constant C) such that $\Pr[B_S(x,t) \not\subseteq P_0(x)] \le \frac{1}{2}$, completing the proof.

The next lemma is specific to metrics with a graph structure.

LEMMA 2.8. *Let* $G = (V,E)$ *be a metric graph, and* $S \subseteq V$ *an arbitrary subset. Suppose that*

1. (S,d) *admits a* (ρ, δ, Δ)*-padded partition.*

2. *For every $R \subseteq V$ with $R \cap S = \emptyset$, $(R, d_{G[R]})$ admits a (ρ, δ, Δ)-padded partition.*

Then, (V, d_G) admits an $(\rho/32, \delta/2, \Delta + \rho/2)$-padded partition.

Proof. Let $P : S \to 2^S$ be the partition from condition (1) above and for every $R \subseteq V$ with $R \cap S = \emptyset$, let $P_R : R \to 2^R$ be the partition promised by condition (2) above. Let $\Gamma : V \to S$ be such that $d_G(x, \Gamma(x)) = d_G(x, S)$ for all $x \in V$.

Now, let $\beta \in [0, 1]$ be chosen uniformly at random, and put $R = V \setminus N_{\beta\rho/4}(S)$. Let $\widehat{P}$ be the $(\rho/2, \delta, \Delta + \rho/2)$-padded partition of $N_{\beta\rho/4}(S)$ guaranteed by applying Lemma 2.6 to P. Define a random partition $P^* : V \to 2^V$ by

$$P^*(x) = \begin{cases} P_R(x) & x \in R \\ \widehat{P}(x) & \text{otherwise} \end{cases}$$

For every $x \in V \setminus R$, we have $\mathrm{diam}_{(V,d_G)}(P(x)) \leq \Delta + \rho/2$ and for every $x \in R$,

$$\mathrm{diam}_{(V,d_G)}(P_R(x)) \leq \mathrm{diam}_{(R,d_{G[R]})}(P_R(x)) \leq \Delta,$$

where the latter inequality follows because $d_G(x, y) \leq d_{G[R]}(x, y)$ for all $x, y \in R$.

Consider now any point $x \in V$. Let $\mathcal{E}$ be the event

$$\left\{B_{(V,d_G)}(x, \rho/16) \subseteq R\right\} \cup \left\{B_{(V,d_G)}(x, \rho/16) \subseteq V \setminus R\right\}.$$

Then, by our random choice of $\beta \in [0, 1]$, we have

$$\Pr[\mathcal{E}] \geq 1 - \frac{2 \cdot (\rho/16)}{\rho/4} = \frac{1}{2}.$$

Now, observe that if $B_{(V,d_G)}(x, \rho/16) \subseteq V \setminus R$, then

$$\begin{aligned}&\Pr[B_{(V,d_G)}(x, \rho/16) \subseteq P^*(x)] \\ &\quad = \Pr[B_{(V,d_G)}(x, \rho/16) \subseteq \widehat{P}(x)] \geq \delta.\end{aligned}$$

On the other hand, if $B_{(V,d_G)}(x, \rho/16) \subseteq R$, then

$$B_{(V,d_G)}(x, \rho/32) = B_{(R,d_{G[R]})}(x, \rho/32),$$

hence

$$\begin{aligned}&\Pr[B_{(V,d_G)}(x, \rho/32) \subseteq P^*(x)] \\ &\quad = \Pr[B_{(R,d_{G[R]})}(x, \rho/32) \subseteq P_R(x)] \geq \delta.\end{aligned}$$

It follows that

$$\Pr[B_{(V,d_G)}(x, \rho/32) \subseteq P^*(x)] \geq \Pr[\mathcal{E}] \cdot \delta \geq \delta/2,$$

completing the proof.

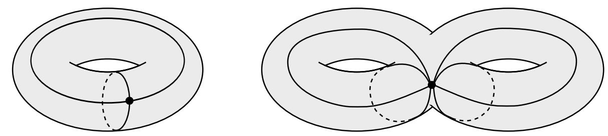

Figure 1: Example of systems of loops for surfaces of genus 1 and 2.

3 Homotopy generators and the cut graph

Let G be a genus-g graph embedded into an orientable genus-g surface S, and let x be a vertex of G. A *system of loops with basepoint* x is a collection of $2g$ cycles $C_1, \ldots, C_{2g}$ containing x such that the complement of $\bigcup_{i=1}^{2g} C_i$ in S is homeomorphic to a disk. Examples of systems of loops are depicted in figure 1 (see also [EW05]). A system of loops is called *optimal* if every C_i is the shortest cycle in its homotopy class. We remark that the set of cycles in a system of loops is a set of generators for the fundamental group $\pi_1(S, x)$.

Algorithms for computing optimal systems of loops have been given by Colin de Verdière and Lazarus [dVL02] and by Erickson and Whittlesey [EW05]. The later algorithm works as follows: Let T be a shortest-path tree in G with root the basepoint x. For every edge $e \in G \backslash T$ let $\sigma(e)$ be the loop obtained by concatenating e with the two paths in T between x and the end-points of e. Let also J be the dual of $G \setminus T$ in S. For every edge $e \in G \setminus T$, we set the weight of its dual $e^* \in J$ to be equal to $\mathsf{len}(\sigma(e))$. Let T' be a maximum spanning tree in J. Finally, let A be the set of the duals of the edges of J, that are not in T'. The resulting system of loops is $X = \{\sigma(e)\}_{e \in A}$.

It now easily follows that the union of all the cycles in X can by decomposed into $O(g)$ shortest paths with disjoint interiors. To see that, let R be the subtree of T induced by all the paths between x and the end-points of the edges in A. Since $|A| = 2g$, it follows that R has at most $4g$ leaves. Moreover, every branch of R is a shortest path, and therefore R can be decomposed into at most $8g - 1$ shortest paths. Since every edge in A is trivially also a shortest path, we obtain the following lemma.

Lemma 3.1. *Let G be a graph embedded into an orientable surface S of genus g. Then, there exists a subgraph H of G such that the complement of H in S is homeomorphic to a disk, satisfying the following properties:*

(a) There exists a collection of $k \leq 12g - 1$ shortest paths $P_1, \ldots, P_k$ in G with disjoint interiors, such that $H = \bigcup_{i \in [k]} P_i$.

(b) There exists a collection of $4g$ shortest-paths $Q_1, \ldots, Q_{4g}$ in G, having x as a common end-point, such that $V(H) = \bigcup_{i \in [4g]} V(Q_i)$.

3.1 Properties of the cut graph We now prove some properties of the metric $(V(H), d_G)$ from Lemma 3.1. Our first lemma gives a family of padded partitions.

LEMMA 3.2. *Let (X, d) be any metric space with $X \subseteq C_1 \cup \cdots \cup C_k$, where each C_i is isometric to a subset of the real line. Then for any $\Delta > 0$, $\alpha(X, d; \frac{1}{4}, \Delta) = O(\log k)$.*

Proof. By Lemma 2.5, for every $\Delta' > 0$, there exists a $(\Delta'/4, \frac{1}{2}, \Delta')$-padded partition of each C_i. Now applying Lemma 2.7 yields a $(\Delta'/O(\log k), \frac{1}{4}, 5\Delta'/4)$-padded partition of (X, d). Setting $\Delta = 5\Delta'/4$ yields a $(\Delta/O(\log k), \frac{1}{4}, \Delta)$-padded partition for every $\Delta > 0$, completing the proof.

The next lemma shows that $(V(H), d_G)$ is close to a metric of bounded pathwidth.

LEMMA 3.3. *The metric $(V(H), d_G)$ from Lemma 3.1 embeds into a graph of pathwidth $O(g)$ with distortion $O(1)$.*

Proof. We will define a graph J with $V(H) = V(J)$ of pathwidth $O(g)$. By Lemma 3.1(b). We have that there exists $x \in V(H)$, and shortest-paths $Q_1, \ldots, Q_{4g}$ in G having x as a common end-point, such that

$$V(H) = \bigcup_{i \in [4g]} Q_i.$$

In order to simplify the argument, we will assume that each path Q_i consists of unit-length edges. This is without loss of generality because after scaling we can assume that the distances are integers, and we can replace every edge $\{u, v\}$ by a path of $\mathsf{len}_G(u, v)$ unit-length edges, without changing the distances in the graph.

Suppose that for each $i \in 4g$ we have $Q_i = q_{i,1}, \ldots, q_{i,t_i}$, where $q_{i,1} = x$. Let $t = \max_{i \in [4g]} t_i$. First, we define a path-decomposition $\mathcal{X}$ for J, with $\mathcal{X} = \{X_i\}_{i=1}^{t}$. We set

$$X_1 = \{x\},$$

and for any $i > 1$, we set

$$X_i = \bigcup_{j \in [4g]} \{q_{\max\{t_i, i-1\}, j}, q_{\max\{t_i, i\}, j}\}.$$

Finally, we define the edge-set of J to be

$$E(J) = \bigcup_{X \in \mathcal{X}} \bigcup_{u \neq v \in X} \{u, v\},$$

and we set the length of each edge $\{u, v\} \in E(J)$ to be $d_G(u, v)$. It is clear that $\mathcal{X}$ is a path-decomposition for J, of width $8g - 1$.

It remains to show that the induced embedding of $(V(H), d_G)$ into $(V(J), d_J)$ has small distortion. It is clear that the embedding is non-contracting, so it suffices to bound the expansion. Consider $u, v \in V(H)$, and let $u \in V(Q_i)$, $v \in V(Q_j)$. Fix a shortest-path P between u and v in G. Assume w.l.o.g. that $u = q_{i,l}$, $v = q_{j,r}$, for some $l \in [t_i]$, $r \in [t_j]$. We have

$$\begin{aligned} d_J(u, v) &\leq d_J(q_{i,l}, q_{j,l}) + d_J(q_{j,l}, q_{j,r}) \\ &= d_G(q_{i,l}, q_{j,l}) + d_G(q_{j,l}, q_{j,r}) \\ &= d_G(q_{i,l}, q_{j,r}) - 2|l - r| \\ &= d_G(u, v) + 2|d_G(x, u) - d_G(x, v)| \\ &\leq 3 d_G(u, v) \end{aligned}$$

Therefore, the distortion is at most 3.

4 Partitions, embeddings, and applications

Our first theorem concerns padded partitions of bounded genus graphs.

THEOREM 4.1. *Let $G = (V, E)$ be a metric graph of orientable genus g. Then for every $\Delta > 0$, $\alpha(V, d_G; \frac{1}{8}, \Delta) = O(\log g)$.*

Proof. Let H be the subgraph guaranteed by Lemma 3.1. By Lemma 3.1(a) and Lemma 3.2, for every $\Delta > 0$, $(V(H), d_G)$ admits a $(\Delta/O(\log g), \frac{1}{4}, \Delta)$-padded partition. Now, upon removing $V(H)$ from V, the induced graph $G[V \setminus V(H)]$ is planar, hence for any $R \subseteq V$ with $R \cap V(H) = \emptyset$, $(R, d_{G[R]})$ is the shortest-path metric on a planar graph, hence by Theorem 2.1, for every $\Delta > 0$, this metric admits a $(\Delta/O(1), \frac{1}{4}, \Delta)$-padded partition. Applying Lemma 2.8 with $S = V(H)$ yields a $(\Delta/O(\log g), \frac{1}{8}, O(\Delta))$-padded partition of (V, d_G), yielding the statement of the theorem.

Now, consider a finite metric space (X, d), and a non-negative weight function $\omega : X \times X \to [0, \infty)$ which is symmetric, i.e. such that $\omega(x, y) = \omega(y, x)$ for all $x, y \in X$. For any $p \in (0, \infty)$, to any mapping $f : X \to \mathbb{R}$ we associate the quantity

$$\mathsf{avd}_{\omega,p}(f) = \|f\|_{\mathrm{Lip}} \left(\frac{\sum_{x,y \in X} \omega(x, y) |f(x) - f(y)|^p}{\sum_{x,y \in X} \omega(x, y)\, d(x, y)^p} \right)^{1/p}.$$

This is notion of “average distortion” of a mapping was studied by Rabinovich [Rab03], and is closely related to concepts like concentration of measure [MS86] and the observable diameter of a metric space [Gro07]. There are a number of applications of such mappings; see

[FHL05] for applications of the $p = 1$ case, and [BLR08] for $p = 2$. We will be particularly interested in the case where ω can be written as $\omega(x, y) = \pi(x)\pi(y)$ for some $\pi : X \to [0, \infty)$. We refer to such an ω as a *product weight.*

The following theorem is useful in analyzing various semi-definite programs.

THEOREM 4.2. *Let $G = (V, E)$ be a graph of orientable genus g. Let (V, d) be a metric space with the property that $(V, \sqrt{d})$ embeds isometrically into a Hilbert space, and let d_G be the path metric arising from an edge (u, v) having length $d(u, v)$. Then for every $p \geq 1$ and product weight ω, there exist a mapping $f : (V, d_G) \to \mathbb{R}$ with*

$$\mathsf{avd}_{\omega,p}(f) \lesssim C(p)\sqrt{\log g},$$

where $C(p)$ is some constant depending only on p.

The proof of this theorem is deferred to the full version. Using known results on average distortion embeddings in planar graphs [Rab03, KPR99], through a sequence of reductions, it suffices to find such an embedding for the cut graph. For this purpose, we use Lemma 3.3, in conjunction with the technique of Rabinovich [Rab03] for graphs of small treewidth.

Next, we have the following theorem on Euclidean embeddings.

THEOREM 4.3. *Let $G = (V, E)$ be a metric graph of orientable genus g. Then there is an embedding $f : (V, d_G) \to L_2$ with $\mathsf{dist}(f) \lesssim \log g + \sqrt{\log |V|}$.*

In the full version, we prove this theorem by combining the planar embedding theorem of Rao [Rao99] with an embedding of the cut graph that requires only $O(\log g)$ distortion. The latter embedding is constructed in a novel way using the measured descent technique [KLMN04].

Finally, we have the following two conditional embedding theorems.

THEOREM 4.4. *Let $C(n)$ be the supremum of $c_1(X, d)$ over all n-point planar graph metrics. Then for every metric graph $G = (V, E)$ of orientable genus g, we have*

$$c_1(V, d_G) \lesssim \log g + C(n).$$

THEOREM 4.5. *If every metric supported on a graph of pathwidth k probabilistically embeds into a distribution over trees with distortion $f(k)$, then genus-g graphs probabilistically embed into a distribution over planar graphs with distortion $f(O(g)) \cdot O(\log g)$.*

4.1 Applications

We now give some sample applications. We defer a more complete list to the full version.

Sparsest Cut and multi-commodity flows. It is well-known that Theorem 4.1 implies an $O(\log g)$-approximate max-flow/min-cut theorem for product multi-commodity flow instances in genus-g graphs (see [LR99] and [Rab03]). This improves over the bound of $O(g)$ from [KPR93, FT03], and is tight as n-vertex expander graphs yield an $\Omega(\log n)$ gap [LR99]. Since every n-vertex graph has genus $O(n^2)$, this yields the desired lower bound. This also yields an $O(\log g)$-approximation to Sparsest Cut with uniform demands [Rab03],

In the case of general demands, Theorem 4.3 implies an $O(\log g + \sqrt{\log n})$-approximate max-flow/min-cut theorem, and Theorem 4.4 implies that in general the gap is at most $O(\log g + C_P(n))$, where $C_P(n)$ is the maximum gap for an n-point planar network.

Vertex cuts and treewidth approximation. Using Theorem 4.2 in conjunction with [FHL05] yields an $O(\sqrt{\log g})$-approximation for the edge and vertex versions of uniform Sparsest Cut, as well as an $O(\sqrt{\log g})$-approximation for approximately optimal tree decompositions and approximating treewidth. This improves over the previous bounds of $O(g)$.

Laplacian eigenvalue bounds. In a direct application of Theorem 4.1 in the paper [KLPT09], it follows that if a graph G has genus g and maximum degree D, then the kth Laplacian eigenvalue is at most $O(kDg/n)\cdot(\log g)^2$. Note that this bound is almost tight as there are examples (see [GHT84]) which have kth eigenvalue $\Omega(kDg/n)$. The previous bound had $(\log g)^2$ replaced by g^2.

Lipschitz extension and 0-extension. When combined with [LN05], Theorem 4.1 yields the following. Let Z be a Banach space. For any metric graph $G = (V, E)$ of genus-g, any $S \subseteq V$, and any $f : S \to Z$, there exists a mapping $\tilde{f} : V \to Z$ with $\tilde{f}|_S = f$ and $\|\tilde{f}\|_{\mathrm{Lip}} \leq O(\log g)\|f\|_{\mathrm{Lip}}$. This improves over the previous bound of $O(g)$ and compares to the known lower bound of $\Omega(\sqrt{\log g})$ (coming from an n-point metric space with an $\Omega(\sqrt{\log n})$ lower bound [JL84]).

When combined with either [AFH$^+$04] or [LN04] (the former paper uses the earthmover relaxation, while the latter uses the standard 0-extension LP from [CKR01]), this yields an $O(\log g)$-approximation to the 0-extension problems on graphs of genus g.

References

[AFH$^+$04] A. Archer, J. Fakcharoenphol, C. Harrelson, R. Krauthgamer, K. Talwar, and E. Tardos. Approximate classification via earthmover metrics. In *15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1072–1080, January 2004.

[AGMW07] I. Abraham, C. Gavoille, D. Malkhi, and U. Wieder. Strong-diameter decompositions of minor free graphs. In *19th ACM Symposium on Parallelism in Algorithms and Architectures*, 2007.

[ARV04] Sanjeev Arora, Satish Rao, and Umesh Vazirani. Expander flows, geometric embeddings, and graph partitionings. In *36th Annual Symposium on the Theory of Computing*, pages 222–231, 2004.

[Bar96] Yair Bartal. Probabilistic approximations of metric space and its algorithmic application. In *37th Annual Symposium on Foundations of Computer Science*, pages 183–193, October 1996.

[Bar98] Yair Bartal. On approximating arbitrary metrics by tree metrics. In *Annual ACM Symposium on Theory of Computing*, pages 183–193, 1998.

[BLR08] P. Biswal, J. R. Lee, and S. Rao. Eigenvalue bounds, spectral partitioning, and metrical deformations via flows. In *49th Annual Symposium on Foundations of Computer Science*. IEEE Computer Soc., Los Alamitos, CA, 2008.

[CCGG98] M. Charikar, C. Chekuri, A. Goel, and S. Guha. Rounding via trees: deterministic approximation algorithms for group Steiner trees and k-median. In *30th Annual ACM Symposium on Theory of Computing*, pages 114–123. ACM, New York, 1998.

[CDG$^+$09] T.-H. Hubert Chan, Kedar Dhamdhere, Anupam Gupta, Jon Kleinberg, and Aleksandrs Slivkins. Metric embeddings with relaxed guarantees. *SIAM J. Comput.*, 38(6):2303–2329, 2009.

[CGMZ05] H. T-H. Chan, A. Gupta, B. M. Maggs, and S. Zhou. On hierarchical routing in doubling metrics. In *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2005.

[CGR05] S. Chawla, A. Gupta, and H. Räcke. An improved approximation to sparsest cut. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, Vancouver, 2005. ACM.

[CKR01] Gruia Calinescu, Howard Karloff, and Yuval 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.

[CMM06] Eden Chlamtac, Konstantin Makayrchev, and Yury Makarychev. How to play unique game using embeddings. In *47th Annual Syposium on Foundations of Computer Science*, 2006.

[dVL02] È. Colin de Verdière and F. Lazarus. Optimal system of loops on an orientable surface. In *Annual Symposium on Foundations of Computer Science*, 2002.

[EW05] Jeff Erickson and Kim Whittlesey. Greedy optimal homotopy and homology generators. In *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1038–1046, 2005.

[FHL05] U. Feige, M. T. Hajiaghayi, and J. R. Lee. Improved approximation algorithms for minimum-weight vertex separators. In *37th Annual ACM Symposium on Theory of Computing*. ACM, 2005. To appear, *SIAM J. Comput.*

[FHRT03] Jittat Fakcharoenphol, Chris Harrelson, Satish Rao, and Kunal Talwar. An improved approximation algorithm for the 0-extension problem. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 257–265, New York, 2003. ACM.

[FRT03] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 448–455, 2003.

[FT03] J. Fakcharoenphol and K. Talwar. An improved decomposition theorem for graphs excluding a fixed minor. In *Proceedings of 6th Workshop on Approximation, Randomization, and Combinatorial Optimization*, volume 2764 of *Lecture Notes in Computer Science*, pages 36–46. Springer, 2003.

[GHT84] John R. Gilbert, Joan P. Hutchinson, and Robert Endre Tarjan. A separator theorem for graphs of bounded genus. *J. Algorithms*, 5(3):391–407, 1984.

[GKL03] Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *44th Symposium on Foundations of Computer Science*, pages 534–543, 2003.

[Gro07] Misha Gromov. *Metric structures for Riemannian and non-Riemannian spaces*. Modern Birkhäuser Classics. Birkhäuser Boston Inc., Boston, MA, english edition, 2007. Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates.

[JL84] William B. Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability (New Haven, Conn., 1982)*, volume 26 of *Contemp. Math.*, pages 189–206. Amer. Math. Soc., Providence, RI, 1984.

[KL03] R. Krauthgamer and J. R. Lee. The intrinsic dimensionality of graphs. In *35th Annual ACM Symposium on Theory of Computing*. ACM, 2003.

[KLMN04] R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: A new embedding method for finite metrics. *Geom. Funct. Anal.*, 2004. To appear.

[KLPT09] J. Kelner, J. R. Lee, G. Price, and S.-H. Teng. Higher eigenvalues of graphs. To appear, *FOCS*, 2009.

[KPR93] Philip N. Klein, Serge A. Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pages 682–690, 1993.

[KPR99] M. R. Korupolu, C. G. Plaxton, and Rajmohan Rajaraman. Placement algorithms for hierarchical cooperative caching. *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 586–595, 1999.

[Lee05] J. R. Lee. On distance scales, embeddings, and efficient relaxations of the cut cone. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 92–101, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[Lee09] J. R. Lee. Volume distortion for subsets of Euclidean spaces. *Discrete Comput. Geom.*, 41(4):590–615, 2009.

[LN04] James R. Lee and Assaf Naor. Metric decomposition, smooth measures, and clustering. *Preprint*, 2004.

[LN05] James R. Lee and Assaf Naor. Extending Lipschitz functions via random metric partitions. *Invent. Math.*, 160(1):59–95, 2005.

[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.

[MN07] Manor Mendel and Assaf Naor. Ramsey partitions and proximity data structures. *J. Eur. Math. Soc. (JEMS)*, 9(2):253–275, 2007.

[MS86] Vitali D. Milman and Gideon Schechtman. *Asymptotic theory of finite-dimensional normed spaces*, volume 1200 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1986. With an appendix by M. Gromov.

[Rab03] Y. Rabinovich. On average distortion of embedding metrics into the line and into L_1. In *35th Annual ACM Symposium on Theory of Computing*. ACM, 2003.

[Rao99] Satish 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.

[Tal04] K. Talwar. Bypassing the embedding: Approximation schemes and distance labeling schemes for growth restricted metrics. In *36th Annual ACM Symposium on Theory of Computing*, 2004.