

Optimal stochastic planarization

Anastasios Sidiropoulos

Toyota Technological Institute at Chicago

Email: tasos@ttic.edu

Abstract—It has been shown by Indyk and Sidiropoulos [1] that any graph of genus $g > 0$ can be stochastically embedded into a distribution over planar graphs with distortion $2^{O(g)}$. This bound was later improved to $O(g^2)$ by Borradaile, Lee and Sidiropoulos [2]. We give an embedding with distortion $O(\log g)$, which is asymptotically optimal.

Apart from the improved distortion, another advantage of our embedding is that it can be computed in polynomial time. In contrast, the algorithm of [2] requires solving an NP-hard problem.

Our result implies in particular a reduction for a large class of geometric optimization problems from instances on genus- g graphs, to corresponding ones on planar graphs, with a $O(\log g)$ loss factor in the approximation guarantee.

Keywords—embeddings; genus; planar graphs.

I. INTRODUCTION

Planar graphs constitute an important class of combinatorial structures, since they can be used to model a wide variety of natural objects. At the same time, they have properties that give rise to improved algorithmic solutions for numerous graph problems, if one restricts the set of possible inputs to planar graphs (see, for example [3]).

One natural generalization of planarity involves the genus of a graph. Informally, a graph has genus g , for some $g \geq 0$, if it can be drawn without any crossings on the surface of a sphere with g additional handles (see Section I-C). For example, a planar graph has genus 0, and a graph that can be drawn on a torus has genus at most 1.

In a way, the genus of a graph quantifies how far it is from being planar. Because of their similarities to planar graphs, graphs of small genus usually exhibit nice algorithmic properties. More precisely, algorithms for planar graphs can usually be extended to graphs of bounded genus, with a small loss in efficiency or quality of the solution (e.g. [4]). Unfortunately, many such extensions are complicated and based on ad-hoc techniques.

Inspired by Bartal’s stochastic embedding of general metrics into trees [5], Indyk and Sidiropoulos [1] showed that every metric on a graph of genus g can be stochastically embedded into a planar graph with distortion at most exponential in g (see Section I-C for a formal definition of stochastic embeddings). Since the distortion measures the ability of the probabilistic mapping to preserve metric properties of the original space, it is desirable to make this quantity as small as possible. The above bound was later improved by Borradaile, Lee, and Sidiropoulos [2], who

obtained an embedding with distortion polynomial in g . In the present paper, we give an embedding with distortion $O(\log g)$, which matches the $\Omega(\log g)$ lower bound from [2]. The statement of our main result follows.

Theorem I.1 (Stochastic planarization). *Any graph G of genus g , admits a stochastic embedding into a distribution over planar graphs, with distortion $O(\log g)$. Moreover, given a drawing of G into a genus- g surface, the embedding can be computed in polynomial time.*

We note that Theorem I.1 can be equivalently stated for compact 2-dimensional simplicial manifolds, i.e. continuous spaces obtained by glueing together finitely many triangles, with every point having a neighborhood homeomorphic to a disk. The result for these spaces can be obtained via a careful affine extension of our embedding over simplices. Since our focus is on algorithmic applications, we omit the details, and restrict our discussion to the discrete case (i.e. finite graphs).

A. Our techniques

In [1] it was shown that a graph of genus g can be stochastically embedded into a distribution over graphs of genus $g - 1$, with constant distortion. Repeating this g times results in a planar graph, but yields distortion exponential in g . The improvement of [2] was obtained by giving an algorithm that removes all handles at once. The main technical tool used to achieve this was the Peeling Lemma from [6]. The idea is that given a graph G of genus g , one can find a subgraph $H \subset G$, which we refer to as the *cut graph*, such that (i) $G \setminus H$ is planar, (ii) H has dilation $g^{O(1)}$, and (iii) H can be stochastically embedded into a planar graph. The resulting distortion of the embedding produced via the Peeling Lemma is proportional to the dilation of H , and therefore polynomial in g .

It was further shown in [2] that any cut graph has dilation $\Omega(g)$, imposing a limitation on their technique. We overcome this barrier as follows. We first find a cut graph consisting of $O(g)$ shortest paths with a common end-point. These paths are obtained from the generators of the fundamental group of the underlying surface, due to Erickson and Whittlesey [7]. In the heart of our analysis, we show how to embed a collection of shortest paths with a common end-point, into a random tree with distortion $O(\log g)$. This result can be viewed as a generalization of the tree-embedding theorem due to Fakcharoenphol, Rao, and Talwar [8], who showed

that any n -point metric space admits a stochastic embedding into a tree with distortion $O(\log n)$.

This connection with tree embeddings seems surprising, since planar graphs appear to be significantly more complicated topologically. For instance, even embedding the $n \times n$ grid into a random tree, requires distortion $\Omega(\log n)$, due to a lower bound of Alon, Karp, Peleg, and West [9]. Gupta, Newman, Rabinovich, and Sinclair [10] have shown that the same lower bound of $\Omega(\log n)$ holds even for embedding very simple classes of planar graphs into trees, such as series-parallel graphs (i.e. even for planar graphs of treewidth 2).

Our tree-embedding result is obtained by combining the approach from [8] with the algorithm of Lee and Sidiropoulos [11] for computing random partitions for graphs of small genus. We remark however that the algorithm of [8] computes an embedding into an *ultrametric*¹, and it can be shown that even a single shortest path cannot be embedded into a random ultrametric with distortion better than $\Omega(\log n)$. We therefore need new ideas to obtain distortion $O(\log g)$. One key ingredient towards this is a new random decomposition scheme, which we refer to as *alternating partitions*, and which takes into account the topology of the paths that we wish to partition. These techniques might be of independent interest.

B. Applications

Optimization: As in the case of stochastic embeddings of arbitrary metrics into trees [5], we obtain a general reduction from a class of optimization problems on genus- g graphs, to their restriction on planar graphs. We now state precisely the reduction. Let V be a set, $\mathcal{I} \subset \mathbb{R}_+^{V \times V}$ a set of non-negative vectors corresponding all feasible solutions for a minimization problem, and $c \in \mathbb{R}_+^{V \times V}$. Then, we define the *linear minimization problem* (\mathcal{I}, c) to be the computational problem where we are given a graph $G = (V, E)$, and we are asked to find $s \in \mathcal{I}$, minimizing

$$\sum_{\{u,v\} \in V \times V} c_{u,v} \cdot s_{u,v} \cdot d(u,v)$$

Observe that this definition captures a very general class of problems. For example, MST can be encoded by letting \mathcal{I} be the set of indicator vectors of the edges of all spanning trees on V , and c the all-ones vector. Similarly, one can easily encode problems such as TSP, Facility-Location, k -Server, Bi-Chromatic Matching, etc.

We can now state an immediate Corollary of our embedding result.

Corollary I.2. *Let $\Pi = (\mathcal{I}, c)$ be a linear minimization problem. If there exists a polynomial-time α -approximation algorithm for Π on planar graphs, then there exists a*

¹A metric space (X, d) where for every $x, y, z \in X$, $d(x, y) \leq \max\{d(x, z), d(z, y)\}$.

randomized polynomial-time $O(\alpha \cdot \log g)$ -approximation algorithm for Π on graphs of genus g .

We remark that in the above Corollary, we can allow the approximation guarantee α to be a function of n , using the following argument. It has been shown by Englert et al. [12] that for any planar graph $G = (V, E)$, and any $U \subseteq V$, there exists a stochastic embedding of (U, d_G) into a distribution over planar graphs with $|U|$ vertices. Composing our stochastic embedding with this result, gives an embedding into a random planar graph with the same number of vertices as the input graph.

Metric embeddings: One of the most intriguing open problems in the theory of metric embeddings is determining the optimal distortion for embedding planar graphs, and more generally graphs that exclude a fixed minor, into L_1 (see e.g. [6], [10], [13], [14]). We remark that by the work of Linial, London, and Rabinovich [13], this distortion equals precisely the maximum multi-commodity max-flow/min-cut gap on these graphs, and is therefore of central importance in divide-and-conquer algorithms that are based on Sparsest-Cut [15], [16]. Our embedding result immediately implies the following corollary. The first proof of this statement was given in [11], where it was derived via a fairly complicated argument.

Corollary I.3. *If all planar graphs embed into L_1 with distortion at most α , then all graphs of genus g embed into L_1 with distortion $O(\alpha \cdot \log g)$.*

C. Preliminaries

Throughout the paper, we consider graphs $G = (V, E)$ with a non-negative length function $\text{len} : E \rightarrow \mathbb{R}$. For a pair $u, v \in V(G)$, we denote the length of the shortest path between u and v in G , with the lengths of edges given by len , by $d_G(u, v)$. By scaling all edge lengths, we may assume that the minimum distance is 1. Let also Δ denote the diameter of G . Unless otherwise stated, we restrict our attention to finite graphs.

Graphs on surfaces: Let us recall some notions from topological graph theory (an in-depth exposition can be found in [17]). A *surface* is a compact connected 2-dimensional manifold, without boundary. For a graph G we can define a one-dimensional simplicial complex C associated with G as follows: The 0-cells of C are the vertices of G , and for each edge $\{u, v\}$ of G , there is a 1-cell in C connecting u and v . A *drawing* of G on a surface S is a continuous injection $f : C \rightarrow S$. The *orientable genus* of a graph G is the smallest integer $g \geq 0$ such that C can be drawn into a sphere with g handles. Note that a graph of genus 0 is a planar graph.

Metric embeddings: A mapping $f : X \rightarrow Y$ between two metric spaces (X, d) and (Y, d') is *non-contracting* if $d'(f(x), f(y)) \geq d(x, y)$ for all $x, y \in X$. If (X, d) is any finite metric space, and \mathcal{Y} is a family of finite metric spaces,

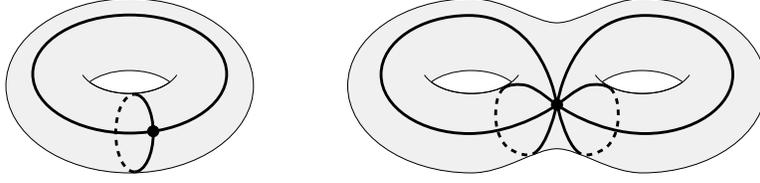


Figure 1. Example of systems of loops for surfaces of genus one and two.

we say that (X, d) admits a stochastic D -embedding into \mathcal{Y} if there exists a random metric space $(Y, d') \in \mathcal{Y}$ and a random non-contracting mapping $f : X \rightarrow Y$ such that for every $x, y \in X$,

$$\mathbb{E} \left[d'(f(x), f(y)) \right] \leq D \cdot d(x, y). \quad (1)$$

The infimal D such that (1) holds is the *distortion of the stochastic embedding*. A detailed exposition of results on metric embeddings can be found in [18] and [19].

D. Organization

The rest of the paper is organized as follows. In Section II we show that in any graph of genus g , we can find a collection of $O(g)$ shortest paths with a common end-point, whose removal leaves a planar graph. In Section III we define alternating partitions for the metric space M induced on these paths. Using these partitions, we show in Section IV how to embed M into a random tree, with distortion $O(\log g)$. Finally, in Section V we combine this tree embedding with the Peeling Lemma, to obtain our main result.

II. HOMOTOPY GENERATORS

Let G be a genus- g graph embedded into an orientable genus- g surface S , and let r be a vertex of G . A *system of loops with basepoint r* is a collection of $2g$ cycles C_1, \dots, C_{2g} containing r , 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 [7] for a detailed exposition). The set of cycles in a system of loops generate the fundamental group $\pi_1(S, r)$.

A system of loops is called *optimal* if every C_i is the shortest cycle in its homotopy class. Algorithms for computing optimal systems of loops have been given by Colin de Verdière and Lazarus [20] and by Erickson and Whittlesey [7]. The later algorithm has the property that each cycle C_i can be decomposed into either two shortest paths with common end-point r , or two such shortest paths, and an edge between the other two end-points. We therefore have the following.

Lemma II.1 (Greedy homotopy generators [7]). *Let G be a graph embedded into an orientable surface S of genus g . Then, there exists a subgraph H of G satisfying the following properties:*

- (i) *The complement of H in S is homeomorphic to a disk.*
- (ii) *There exists $r \in V(G)$, and a collection of $4g$ shortest-paths Q_1, \dots, Q_{4g} in G , having r as a common end-point, such that $V(H) = \bigcup_{i \in [4g]} V(Q_i)$.*

III. ALTERNATING PARTITIONS

Let G be a graph. By rescaling the edge-lengths we may assume w.l.o.g. that the minimum distance in G is one. Let $\mathcal{P} = \{P_1, \dots, P_k\}$ be a collection of shortest paths in G , with a common end-point $r \in V(G)$. Let $X = \bigcup_{i=1}^k V(P_i)$. We consider the metric space (X, d_G) . We define a collection $\{\mathcal{C}_i\}_{i=0}^{2+\log \Delta}$, where each \mathcal{C}_i is a random partition of X into sets of diameter less than 2^i , and such that for any $i \in \{1, \dots, 1+\log \Delta\}$, \mathcal{C}_i is a refinement of \mathcal{C}_{i+1} . We refer to the resulting collection $\{\mathcal{C}_i\}_i$ as *alternating partitions* for (X, d_G) .

Pick a permutation $\sigma \in S_g$, and reals $\alpha \in [0, 1)$, $\beta \in [1, 2)$, uniformly, and independently at random. This is all the randomness that will be used in the construction.

We set $\mathcal{C}_{2+\log \Delta} = \{X\}$, i.e. the trivial partition that places all points into the same cluster. For $i = 1 + \log \Delta, \dots, 1$, given \mathcal{C}_{i+1} we define \mathcal{C}_i by performing two partitioning steps that we describe below (see also Figure 2).

Horizontal partitioning step: Let $A \in \mathcal{C}_{i+1}$. We partition A into clusters $\{A_s\}_{s=1}^k$. We consider the paths in \mathcal{P} in the order $P_{\sigma(1)}, \dots, P_{\sigma(k)}$. For each $s \in \{1, \dots, k\}$, we form the cluster

$$A_s = \left(A \cap N_G(P_{\sigma(s)}, \beta \cdot 2^{i-3}) \right) \setminus \bigcup_{t=1}^{s-1} A_t,$$

where $N_G(P_l, \delta) = \{x \in X : d_G(x, P_l) \leq \delta\}$. We say that the path $P_{\sigma(s)}$ is the *trunk* of A_s . For notational convenience, we also refer to $P_{\sigma(1)}$ as the trunk of the unique cluster $\{X\}$ in the partition $\mathcal{C}_{2+\log \Delta}$. We refer to the clusters $\{A_s\}_{s=1}^k$ as the *horizontal children* of A .

Vertical partitioning step: Next, we proceed to partition each horizontal child A_s of A into a set of clusters $\{A_{s,j}\}_{j \in \mathbb{N}_0}$, so that for any integer $j \geq 0$,

$$A_{s,j} = \{x \in A_s : (j-1+\alpha) \cdot 2^{i-2} \leq d_G(r, x) < (j+\alpha) \cdot 2^{i-2}\}$$

We refer to the clusters $\{A_{s,j}\}_{j \in \mathbb{N}_0}$ as the *vertical children* of A_s . We also say that $P_{\sigma(s)}$ is the trunk of $A_{s,j}$. Finally, we add all non-empty clusters $A_{s,j}$ to \mathcal{C}_i .

²It suffices to choose α and β within $O(\log n)$ bits of precision.

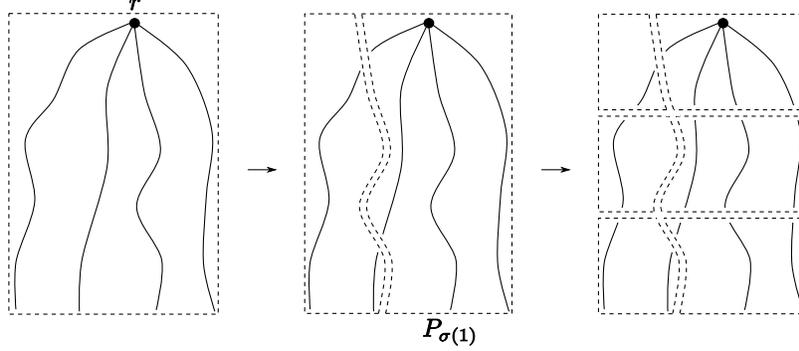


Figure 2. Alternating horizontal and vertical partitioning steps.

This concludes the description of the construction of the alternating partitions $\{\mathcal{C}_i\}_i$ for (X, d_G) .

Lemma III.1. *For any $i \in \{0, \dots, 2 + \log \Delta\}$, and $A \in \mathcal{C}_i$, we have $\text{diam}_G(A) < 2^i$.*

Proof: Let P_s be the trunk of A . Let Q be the subpath of P_s that is contained in A . By the construction of the vertical children we have $\text{len}(Q) < 2^{i-2}$. Moreover, by the construction of the horizontal children we have that for any $x \in A$, $d_G(x, Q) < 2^{i-2}$. Therefore, for any $x, y \in A$ we have $d_G(x, y) \leq d_G(x, C) + d_G(y, C) + \text{len}(Q) < 2^i$. ■

IV. EMBEDDING THE CUT GRAPH INTO A RANDOM TREE

As in the previous section, let G be a graph, and $r \in V(G)$. Let P_1, \dots, P_k be shortest paths in G with common end-point r , and define $X = \bigcup_{i=1}^k V(P_k)$. We will use the alternating partitions $\{\mathcal{C}_i\}_i$ constructed in the previous section to obtain a stochastic embedding of (X, d_G) into a distribution over trees, with distortion $O(\log k)$.

For any $Y \subseteq V(G)$, let $\text{top}(Y) = \min_{v \in Y} d_G(r, v)$, and $\text{bottom}(Y) = \max_{v \in Y} d_G(r, v)$.

We proceed by induction on the partitions $\{\mathcal{C}_i\}_{i \in \mathbb{N}_0}$, starting from \mathcal{C}_0 . For every cluster $A \in \mathcal{C}_i$ we construct a tree T_A and an injection $f_A : A \rightarrow V(T_A)$. We inductively maintain the following invariant:

- (I) For every cluster A with trunk P_s , there exists in T_A a copy of the subpath of P_s containing all vertices $v \in V(P_s)$ with $\text{top}(A) \leq d_G(r, v) \leq \text{bottom}(A)$. We refer to this path as the *stem* of A . We denote by r_A the vertex in the stem of A which is closest to r in G . We refer to r_A as the *root* of A .

By Lemma III.1 we have that every cluster $A \in \mathcal{C}_0$ has diameter less than the minimum distance in G , and therefore contains a single vertex. We set T_A to be the trivial tree containing that vertex. The map f_A sends the unique vertex in A to its copy in T_A .

Suppose now that we have constructed a tree for every cluster in \mathcal{C}_{i-1} , for some $i \geq 1$. We will show how to

obtain a tree for every cluster in \mathcal{C}_i . Let $A \in \mathcal{C}_i$, and let $\{A_s\}_{s=1}^k$ be the horizontal children of A . For a horizontal child A_s , let $\{A_{s,j}\}_{j \in \mathbb{N}_0}$ be its vertical children. Recall that each such $A_{s,j}$ is a cluster in \mathcal{C}_{i-1} . Therefore, by the induction hypothesis we have already computed a tree $T_{A_{s,j}}$ for every $A_{s,j}$, and an injection $f_{A_{s,j}} : A_{s,j} \rightarrow V(T_{A_{s,j}})$. We construct the tree T_A in two steps:

Vertical composition step: We first combine the graphs of the vertical children of each A_s , to obtain an intermediate tree T_{A_s} . This is done as follows. Recall that P_s is the trunk of A . By the inductive invariant (I) we have that for every vertical child $A_{s,j}$, its stem $Q_{s,j}$ is a path in $T_{A_{s,j}}$, and each such $Q_{s,j}$ is a copy of a subpath of P_s . In particular, since for all i we have $\text{bottom}(A_{s,i}) < \text{top}(A_{s,i+1})$, it follows that the stems of distinct vertical children correspond to disjoint subpaths of P_s . Let Q_s be the subpath containing all vertices $v \in P_s$ with $\text{top}(A_s) \leq d_G(r, v) \leq \text{bottom}(A_s)$. Since $A_{s,j} \subseteq A_s$ it follows that $Q_{s,j} \subseteq Q_s$. We form the tree T_{A_s} by taking a copy of Q_s and identifying for every vertical child $A_{s,j}$, the stem $Q_{s,j}$ in $T_{A_{s,j}}$ with its copy in Q_s . The path Q_s becomes the stem of A_s . The mapping f_{A_s} is defined by composing each $f_{A_{s,j}}$ with the natural inclusion $V(T_{A_{s,j}}) \rightarrow V(T_{A_s})$. See Figure 3(a).

Horizontal composition step: Next, we combine the trees T_{A_s} for all horizontal children $\{A_s\}_{s=1}^k$ of A , to obtain T_A . Let P_t be the trunk of A . Observe that there exists a non-empty horizontal child A_t of A . For any $l \in \{1, \dots, k\}$, with $l \neq t$, we connect r_{A_l} with r_{A_t} via an edge of length 2^i . Let T_A be the resulting tree, and f_A be the induced injection from $\bigcup_{s=1}^k A_s$ to $V(T_A)$. Note that the root of A is $r_A = r_{A_t}$. See Figure 3(b).

It is easy to verify that the inductive invariant (I) is maintained. Finally, we set $T = T_{\{X\}}$, and $f = f_{\{X\}}$. This concludes the description of the embedding $f : X \rightarrow V(T)$.

A. Bounding the distortion

It is straight-forward to verify that the mapping $f : X \rightarrow V(T)$ is non-contracting, so it remains to bound the expected

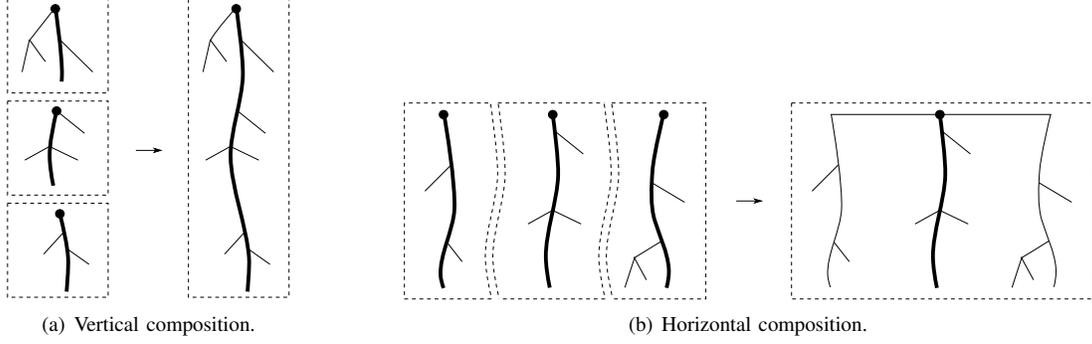


Figure 3. Vertical and horizontal composition steps.

expansion for every pair of vertices. We begin with a useful Lemma.

Lemma IV.1. *Let $i \in \{0, \dots, 2 + \log \Delta\}$, let $A \in \mathcal{C}_i$, and let Q be the stem of A . Then, for any $v \in A$, we have $d_T(f(v), Q) \leq 2^{i+2}$.*

Proof: For any $j \in \{0, \dots, i\}$, let $A_j = \mathcal{C}_j(v)$. We have that T_{A_0} is the tree containing only $f(v)$. For any $j \in \{0, \dots, i\}$, let Q_j be the stem of A_j . We have $\text{len}(Q_j) < 2^{j-2}$, and $d_T(r_{A_j}, Q_{j+1}) \leq d_T(r_{A_j}, r_{A_{j+1}}) \leq 2^{j+1}$. Thus $d_T(f(v), Q) \leq \sum_{j=0}^{i-1} (\text{len}(Q_j) + d_T(r_{A_j}, Q_{j+1})) < \sum_{j=0}^{i-1} (2^{j-2} + 2^{j+1}) < 2^{i+2}$ ■

For the remaining of the analysis, we fix two vertices $u, v \in X$. We wish to bound $\mathbb{E}[d_T(f(u), f(v))]$, where the expectation is taken over the randomness used in constructing the alternating partitions $\{\mathcal{C}_i\}_i$ (i.e. α , β , and σ).

We begin by introducing some notation. We say that a path $P_s \in \mathcal{P}$ settles $\{u, v\}$ at level i if u and v are in the same cluster in \mathcal{C}_{i+1} , and P_s is the first path w.r.to the ordering σ such that P_s is the trunk of at least one of the clusters $\mathcal{C}_i(u)$, $\mathcal{C}_i(v)$. Moreover, we say that P_s cuts horizontally $\{u, v\}$ at level i if it settles $\{u, v\}$ at level i , and exactly one of the clusters $\mathcal{C}_i(u)$, $\mathcal{C}_i(v)$ has P_s as its trunk (i.e. u and v get separated during the horizontal partitioning step, while constructing \mathcal{C}_i).

Similarly, we say that P_s saves $\{u, v\}$ at level i if u and v are in the same cluster in \mathcal{C}_{i+1} , and P_s is the trunk of a horizontal child of a cluster in \mathcal{C}_{i+1} containing both u and v . We say that P_s cuts vertically $\{u, v\}$ at level i if it saves $\{u, v\}$ at level i , and $\mathcal{C}_i(u) \neq \mathcal{C}_i(v)$ (i.e. u and v get separated during the vertical partitioning step, while constructing \mathcal{C}_i).

Let γ_s^i , resp. δ_s^i , be the supremum of $d_T(f(u), f(v))$ when s cuts $\{u, v\}$ at level i horizontally, resp. vertically, taken over all possible random choices of the algorithm. That is,

$$\gamma_s^i = \sup_{\alpha, \beta, \sigma} \{d_T(f(u), f(v)) : s \text{ cuts horizontally } \{u, v\} \text{ at level } i\}$$

$$\delta_s^i = \sup_{\alpha, \beta, \sigma} \{d_T(f(u), f(v)) : s \text{ cuts vertically } \{u, v\} \text{ at level } i\}$$

Then, we have

$$\mathbb{E}[d_T(f(u), f(v))] \leq \Phi_1 + \Phi_2, \quad (2)$$

where

$$\Phi_1 = \sum_{s=1}^k \sum_{i=0}^{2+\log \Delta} \gamma_s^i \cdot \Pr[P_s \text{ cuts horizontally } \{u, v\} \text{ at level } i]$$

$$\Phi_2 = \sum_{s=1}^k \sum_{i=0}^{2+\log \Delta} \delta_s^i \cdot \Pr[P_s \text{ cuts vertically } \{u, v\} \text{ at level } i]$$

We will bound each one of these quantities separately.

Lemma IV.2. $\Phi_1 \leq O(\log k) \cdot d_G(u, v)$.

Proof: Define the interval

$$I_s = [\min\{d_G(u, P_s), d_G(v, P_s)\}, \max\{d_G(u, P_s), d_G(v, P_s)\}]$$

In order for P_s to cut horizontally $\{u, v\}$ at level i it must be the case that $\beta \cdot 2^{i-3} \in I_s$. Since β is chosen from $[1, 2)$ uniformly at random, it follows by the triangle inequality that this happens with probability at most

$$\Pr[\beta \cdot 2^{i-3} \in I_s] \leq |I_s|/2^{i-3} \leq d_G(u, v)/2^{i-3} \quad (3)$$

Assume w.l.o.g. that $d_G(P_1, \{u, v\}) \leq \dots \leq d_G(P_k, \{u, v\})$. Conditioned on the event that $\beta \cdot 2^{i-3} \in I_s$, any of the paths P_1, \dots, P_s can settle $\{u, v\}$. Therefore,

$$\Pr[P_s \text{ settles } \{u, v\} \mid \beta \cdot 2^{i-3} \in I_s] \leq 1/s \quad (4)$$

Next we bound γ_s^i . Suppose that a path P_s cuts horizontally $\{u, v\}$ at level i . Let Q be the stem of the cluster in \mathcal{C}_{i+1} containing both u and v . By Lemma IV.1 we conclude that

$$\begin{aligned} \gamma_s^i &\leq d_T(f(u), f(v)) \\ &\leq d_T(f(u), Q) + \text{len}(Q) + d_T(f(v), Q) \\ &\leq 2^{i+5} \end{aligned} \quad (5)$$

Observe that since $\beta \in [1, 2)$, it follows that for every $s \in \{1, \dots, k\}$, there exists an integer i_s such that the path P_s can cut horizontally $\{u, v\}$ only at level i_s , or $i_s + 1$. Therefore

$$\begin{aligned}
\Phi_1 &\leq \sum_{s=1}^k \sum_{i=0}^1 2^{i_s+j+5} \\
&\quad \cdot \Pr[P_s \text{ settles } \{u, v\} \text{ at level } i_s \text{ and } \beta \cdot 2^{i_s+j-3} \in I_s] \\
&\leq \sum_{i=1}^k \sum_{j=0}^1 2^{i_s+j+5} \\
&\quad \cdot \Pr[P_s \text{ settles } \{u, v\} \text{ at level } i_s + j \mid \beta \cdot 2^{i_s+j-3} \in I_s] \\
&\quad \cdot \Pr[\beta \cdot 2^{i_s+j-3} \in I_s] \\
&\leq O\left(\sum_{s=1}^k 2^{i_s} \cdot \frac{1}{s} \cdot \frac{d_G(u, v)}{2^{i_s}}\right) \\
&= O(\log k) \cdot d_G(u, v)
\end{aligned}$$

Lemma IV.3. $\Phi_2 \leq O(\log k) \cdot d_G(u, v)$.

Proof: Define

$$J_s = [\max\{d_G(u, P_s), d_G(v, P_s)\}, \infty)$$

$$\begin{aligned}
R_s^i &= \bigcup_{j=0}^{\infty} [j \cdot 2^{i-2} + \min\{d_G(r, u), d_G(r, v)\}, \\
&\quad j \cdot 2^{i-2} + \max\{d_G(r, u), d_G(r, v)\}]
\end{aligned}$$

Denote by \mathcal{E}_1^i the event $\beta \cdot 2^{i-3} \in J_s$, and by \mathcal{E}_2^i the event $\alpha \cdot 2^{i-2} \in R_s^i$. In order for P_s to cut vertically $\{u, v\}$ at level i , both \mathcal{E}_1^i and \mathcal{E}_2^i must hold.

Assume w.l.o.g. that $d_G(P_1, \{u, v\}) \leq \dots \leq d_G(P_k, \{u, v\})$. Conditioned on \mathcal{E}_1^i , any of P_1, \dots, P_s can save $\{u, v\}$. Therefore,

$$\Pr[P_s \text{ saves } \{u, v\} \mid \mathcal{E}_1^i] \leq 1/s \quad (6)$$

By the triangle inequality we have

$$\Pr[\mathcal{E}_2^i] \leq d_G(u, v)/2^{i-2} \quad (7)$$

We next upper bound δ_s^i . Suppose that P_s cuts vertically $\{u, v\}$ at level i . Let $L_s = \max\{d_G(u, P_s), d_G(v, P_s)\}$, and $j_s = 2 + \lceil \log L_s \rceil$. Assume w.l.o.g. that $d_G(r, u) \leq d_G(r, v)$. It follows by the construction of T that P_s is the trunk of $\mathcal{C}_{j_s+1}(u)$ and $\mathcal{C}_{j_s+1}(v)$. Therefore, there exist clusters $U_1, \dots, U_t \in \mathcal{C}_{j_s+1}$ with $u \in U_1, v \in U_t$, such that P_s is the trunk of every U_i , and such that the stems of U_1, \dots, U_t are consecutive subpaths of P_s . Let Q_i be the stem of U_i . For any $j \in \{1, \dots, t-1\}$, the path Q_j is connected to the path Q_{j+1} via a path W_j , such that $\text{len}(Q_j) + \text{len}(W_j) \leq 2^{j_s-1}$.

By lemma IV.1 we have

$$\begin{aligned}
\delta_s^i &\leq d_T(f(u), Q_1) + d_T(f(v), Q_t) \\
&\quad + \sum_{j=1}^{t-1} (\text{len}(Q_j) + \text{len}(W_j)) + \text{len}(Q_t) \\
&\leq 2^{j_s+4} + |d_G(r, u) - d_G(r, v)| + 2 \cdot 2^{j_s-1} \\
&\leq 2^8 \cdot L_s + d_G(u, v) \quad (8)
\end{aligned}$$

Let $\tau_s = \max\{j_s, \lceil \log d_G(u, v) \rceil - 1\}$. Observe that for any $i < j_s$, we have $\Pr[\mathcal{E}_1^i] = 0$. Moreover, if P_s saves $\{u, v\}$ at level i , it must be the case that there exists $A \in \mathcal{C}_{i+1}$ containing both u and v , and thus with $\text{diam}_G(A) \geq d_G(u, v)$. By Lemma III.1 we have $i \geq \lceil \log d_G(u, v) \rceil - 1$. Therefore, for any $i < \tau_s$ we have $\Pr[P_s \text{ cuts } \{u, v\} \text{ at level } i] = 0$. Combining with (6), (7), and (8), we have

$$\begin{aligned}
\Phi_2 &\leq \sum_{s=1}^k \sum_{i=\tau_s}^{2+\log \Delta} \delta_s^i \\
&\quad \cdot \Pr[P_s \text{ saves } \{u, v\} \text{ at level } i \mid \mathcal{E}_1^i \text{ and } \mathcal{E}_2^i] \\
&\quad \cdot \Pr[\mathcal{E}_1^i \text{ and } \mathcal{E}_2^i] \\
&\leq \sum_{s=1}^k \frac{1}{s} \sum_{i=\tau_s}^{2+\log \Delta} \delta_s^i \cdot \frac{d_G(u, v)}{2^{i-2}} \\
&\leq d_G(u, v) \cdot \sum_{s=1}^k \frac{1}{s} \sum_{i=\tau_s}^{2+\log \Delta} \frac{2^8 \cdot L_s + d_G(u, v)}{2^{i-2}} \\
&< d_G(u, v) \cdot \sum_{s=1}^k \frac{1}{s} \cdot \frac{2^8 \cdot L_s + d_G(u, v)}{2^{\tau_s-3}} \\
&\leq O(\log k) \cdot d_G(u, v)
\end{aligned}$$

Combining (2) with lemmas IV.2 and IV.3 we obtain the main result of this section. ■

Theorem IV.4. *Let G be a graph, and let $\mathcal{P} = \{P_1, \dots, P_k\}$ be a collection of shortest-paths in G , sharing a common end-point. Then, the metric space $(\bigcup_{i=1}^k V(P_i), d_G)$ admits a stochastic embedding into a distribution over trees with distortion $O(\log k)$.*

V. PLANARIZATION

Let (X, d) be a metric space. A distribution \mathcal{F} over partitions of X is called (β, Δ) -Lipschitz if every partition in the support of \mathcal{F} has only clusters of diameter at most Δ , and for every $x, y \in X$,

$$\Pr_{C \in \mathcal{F}} [C(x) \neq C(y)] \leq \beta \cdot \frac{d(x, y)}{\Delta}.$$

We denote by $\beta_{(X, d)}$ the infimum β such that for any $\Delta > 0$, the metric (X, d) admits a (Δ, β) -Lipschitz random partition, and we refer to $\beta_{(X, d)}$ as the *modulus of decomposability* of (X, d) . The following theorem is due to Klein, Plotkin, and Rao [21], and Rao [22].

Theorem V.1 ([21], [22]). *For any planar graph G , we have $\beta_{(V(G), d_G)} = O(1)$.*

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

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

For a graph G and a graph family \mathcal{F} we write $G \stackrel{D}{\rightsquigarrow} \mathcal{F}$ to denote the fact that G stochastically embeds into a distribution over graphs in \mathcal{F} , with distortion D .

For two graphs G, G' , a *1-sum* of G with G' is a graph obtained by taking two disjoint copies of G and G' , and identifying a vertex $v \in V(G)$ with a vertex $v' \in V(G')$. For a graph family \mathcal{X} , we denote by $\oplus_1 \mathcal{X}$ the closure of \mathcal{X} under 1-sums.

Lemma V.2 (Peeling Lemma [6]). *Let G be a graph, and $A \subseteq V(G)$. Let $G' = (V(G), E')$ be a graph with $E' = E(G) \setminus E(G[A])$, and let $\beta = \beta_{(V, d_{G'})}$ be the corresponding modulus of decomposability. Then, there exists a graph family \mathcal{F} such that $G \stackrel{D}{\rightsquigarrow} \mathcal{F}$, where $D = O(\beta \cdot \text{dil}_G(A))$, and every graph in \mathcal{F} is a 1-sum of isometric copies of the graphs $G[A]$ and $\{G[V \setminus A \cup \{a\}]\}_{a \in A}$.*

We will use the following auxiliary Lemma.

Lemma V.3 (Composition Lemma). *Let G be a graph, and let $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ be graph families. If $G \stackrel{\alpha}{\rightsquigarrow} \oplus_1(\mathcal{X} \cup \mathcal{Y})$, and for any $G' \in \mathcal{X}$, $G' \stackrel{\beta}{\rightsquigarrow} \mathcal{Z}$, then $G \stackrel{\alpha \cdot \beta}{\rightsquigarrow} \oplus_1(\mathcal{Z} \cup \mathcal{Y})$.*

Proof sketch: It follows by direct composition of the two embeddings. ■

Proof of Theorem I.1: Let G be a genus- g graph, drawn on a genus- g surface S . Let H be the subgraph of G given by Lemma II.1. Recall that there exists $r \in V(G)$ and shortest paths P_1, \dots, P_{4g} in G , having r as a common end-point, such that $V(H) = \bigcup_{i=1}^{4g} V(P_i)$.

Let us write $G_1 = G$, and $X = V(H)$. By Theorem IV.4, we have

$$(X, d_{G_1}) \stackrel{O(\log g)}{\rightsquigarrow} \text{Trees} \quad (9)$$

After scaling the lengths of the edges in G_1 , we may assume that the minimum distance is one. Let G_2 be the graph obtained from G_1 as follows. For every edge $\{u, v\} \in E(G_1)$ with $u \in X_1$ and $v \notin X_1$, we replace $\{u, v\}$ with two edges $\{u, w\}$ and $\{w, v\}$, with $\text{len}(\{u, w\}) = 1/2$, and $\text{len}(\{w, v\}) = \text{len}(\{u, v\}) - 1/2$. Let $Y = X \cup N_{G_2}(X)$, i.e. the set X , together with all new vertices w introduced above.

Observe that for any $x, y \in X$, we have $d_{G_1}(x, y) = d_{G_2}(x, y)$. Therefore by (9),

$$(X, d_{G_2}) \stackrel{O(\log g)}{\rightsquigarrow} \text{Trees}$$

This embedding can be extended to Y as follows. For every tree T in the support of the distribution, and for every vertex

$w \in Y \setminus X$, we attach w to T by adding an edge of length $1/2$ between w and the unique neighbor of w in X . Since we only add leaves to T , the new graph is still a tree. It is straight-forward to verify that the resulting stochastic embedding has distortion $O(\log g)$, and thus

$$(Y, d_{G_2}) \stackrel{O(\log g)}{\rightsquigarrow} \text{Trees} \quad (10)$$

Let G_3 be the graph obtained from G_2 by adding an edge $\{u, v\}$ of length $d_G(u, v)$, between every pair of vertices $u, v \in Y$. Let $E'_3 = E(G_3) \setminus E(G_3[Y])$, and $G'_3 = (V(G_3), E'_3)$. By cutting the surface S along H we obtain a drawing of $G \setminus H$ into the interior of a disk. Since every vertex in $Y \setminus X$ is attached to $G \setminus H$ via a single edge, it follows that this planar drawing can be extended to G'_3 . Thus, the graph G'_3 is planar, and by theorem V.1 we have $\beta_{(V(G'_3), d_{G'_3})} = O(1)$.

Similarly, we have that for any $y \in Y$, the graph $G_3[V(G_3) \setminus Y \cup \{y\}]$ is planar.

Observe that $\text{dil}_{G_3}(Y) = 1$. By the Peeling Lemma (Lemma V.2) we have that G_3 can be stochastically embedded with distortion $O(\text{dil}_{G_3}(Y) \cdot \beta_{G'_3}) = O(1)$ into a distribution over graphs J , where J is obtained by 1-sums of isometric copies of $G_3[Y]$ and $\{G_3[V(G_3) \setminus Y \cup \{y\}]\}_{y \in Y}$. Each graph in $\{G_3[V(G_3) \setminus Y \cup \{y\}]\}_{y \in Y}$ is planar, and therefore

$$G_3 \stackrel{O(1)}{\rightsquigarrow} \oplus_1(\text{Planar}, \{G_3[Y]\}) \quad (11)$$

By (10), and since the metric (Y, d_{G_3}) is the same as the metric (Y, d_{G_2}) , we have that

$$G_3[Y] \stackrel{O(\log g)}{\rightsquigarrow} \text{Trees} \quad (12)$$

Note that the 1-sum of two planar graphs is also planar. Therefore, combining (11), (12) and Lemma V.3, we obtain

$$G_3 \stackrel{O(\log g)}{\rightsquigarrow} \oplus_1(\text{Trees} \cup \text{Planar}) = \oplus_1 \text{Planar} = \text{Planar}$$

Since G_3 contains an isometric copy of G , this implies

$$G \stackrel{O(\log g)}{\rightsquigarrow} \text{Planar},$$

concluding the proof. ■

ACKNOWLEDGEMENTS

The author wishes to thank Glencora Borradaile, Piotr Indyk, and James R. Lee for their invaluable contributions to the development of the techniques used in this paper. He also thanks Dimitrios Thilikos for a question which motivated the problem studied here, as well as Jeff Erickson, and Sariel Har-Peled for several enlightening discussions.

REFERENCES

- [1] P. Indyk and A. Sidiropoulos, “Probabilistic embeddings of bounded genus graphs into planar graphs,” in *Proc. 23rd Annual ACM Symposium on Computational Geometry*, 2007.
- [2] G. Borradaile, J. R. Lee, and A. Sidiropoulos, “Randomly removing g handles at once,” in *Proc. 25th Annual ACM Symposium on Computational Geometry*, 2009.
- [3] B. S. Baker, “Approximation algorithms for np-complete problems on planar graphs,” *J. ACM*, vol. 41, no. 1, pp. 153–180, 1994.
- [4] E. W. Chambers, J. Erickson, and A. Nayyeri, “Homology flows, cohomology cuts,” in *Proc. 41st Annual ACM Symposium on Theory of Computing*, 2009.
- [5] Y. Bartal, “Probabilistic approximation of metric spaces and its algorithmic applications,” in *37th Annual Symposium on Foundations of Computer Science (Burlington, VT, 1996)*. Los Alamitos, CA: IEEE Comput. Soc. Press, 1996, pp. 184–193.
- [6] J. R. Lee and A. Sidiropoulos, “On the geometry of graphs with a forbidden minor,” *Proceedings of the 41st STOC*, Jan 2009.
- [7] J. Erickson and K. Whittlesey, “Greedy optimal homotopy and homology generators,” in *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2005, pp. 1038–1046.
- [8] J. Fakcharoenphol, S. Rao, and K. Talwar, “A tight bound on approximating arbitrary metrics by tree metrics,” in *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, 2003, pp. 448–455.
- [9] N. Alon, R. Karp, D. Peleg, and D. B. West, “Graph-theoretic game and its applications to the k -server problem,” *SIAM Journal on Computing*, 1991.
- [10] A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair, “Cuts, trees and ℓ_1 embeddings,” *Annual Symposium on Foundations of Computer Science*, 1999.
- [11] J. R. Lee and A. Sidiropoulos, “Genus and the geometry of the cut graph,” in *Proc. 21st ACM-SIAM Symposium on Discrete Algorithms*, 2010.
- [12] M. Englert, A. Gupta, R. Krauthgamer, H. Räcke, I. Talgam, and K. Talwar, “Vertex sparsifiers: New results from old techniques,” *CoRR*, vol. abs/1006.4586, 2010.
- [13] N. Linial, E. London, and Y. Rabinovich, “The geometry of graphs and some of its algorithmic applications,” *Proceedings of 35th Annual IEEE Symposium on Foundations of Computer Science*, pp. 577–591, 1994.
- [14] C. Chekuri, A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair, “Embedding k -outerplanar graphs into ℓ_1 ,” in *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, 2003, pp. 527–536.
- [15] F. Leighton and S. Rao, “Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms,” *Journal of the ACM*, vol. 46, pp. 787–832, 1999.
- [16] S. Arora, S. Rao, and U. Vazirani, “Expander flows, geometric embeddings, and graph partitionings,” in *36th Annual Symposium on the Theory of Computing*, 2004, pp. 222–231.
- [17] B. Mohar and C. Thomassen, *Graphs on Surfaces*. John Hopkins, 2001.
- [18] P. Indyk, “Tutorial: Algorithmic applications of low-distortion geometric embeddings,” *Annual Symposium on Foundations of Computer Science*, 2001.
- [19] J. Matousek, *Lectures on Discrete Geometry*. Springer, 2002.
- [20] E. C. de Verdière and F. Lazarus, “Optimal system of loops on an orientable surface,” in *Annual Symposium on Foundations of Computer Science*, 2002.
- [21] 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*, 1993, pp. 682–690.
- [22] S. Rao, “Small distortion and volume preserving embeddings for planar and Euclidean metrics,” in *Proceedings of the 15th Annual Symposium on Computational Geometry*. New York: ACM, 1999, pp. 300–306.