Probabilistic Embeddings of Bounded Genus Graphs Into Planar Graphs

Piotr Indyk
MIT CSAIL
indyk@theory.csail.mit.edu

Anastasios Sidiropoulos
MIT CSAIL
tasos@theory.csail.mit.edu

ABSTRACT

A probabilistic C-embedding of a (guest) metric M into a collection of (host) metrics $M'_1, \ldots, M'_k$ is a randomized mapping F of M into one of the $M'_1, \ldots, M'_k$ such that, for any two points p, q in the guest metric:

1. The distance between $F(p)$ and $F(q)$ in any M'_i is not smaller than the original distance between p and q.

2. The expected distance between $F(p)$ and $F(q)$ in (random) M'_i is not greater than some constant C times the original distance, for $C \geq 1$.

The constant C is called the *distortion* of the embedding. Low-distortion probabilistic embeddings enable reducing algorithmic problems over "hard" guest metrics into "easy" host metrics.

We show that every metric induced by a graph of bounded genus can be probabilistically embedded into planar graphs, with constant distortion. The embedding can be computed efficiently, given a drawing of the graph on a genus-g surface.

Categories and Subject Descriptors

F.2 [**Analysis of Algorithms and Problem Complexity**]: General

General Terms

Algorithms Theory

Keywords

Embeddings, Probabilistic Approximation, Bounded Genus Graphs, Planar Graphs

1. INTRODUCTION

Planar graphs constitute an important class of combinatorial structures, since they can often 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 e.g. [3]).

A natural generalization of planarity can be obtained using the notion of 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 later in this section for a formal definition). For example, a planar graph has genus 0, and a graph that can be drawn on a torus has genus 1.

The genus of a graph can be interpreted as a quantity expressing how far a graph is from being planar. To that extend, graphs of small genus usually exhibit nice algorithmic properties, mainly due to their inherent similarities with planar graphs. More precisely, algorithms for planar graphs can usually be extended to graphs of bounded genus, with a small loss in efficiency, or in the quality of the solution. However, such extensions can in some cases be complicated, and based on ad-hoc techniques.

In this paper we give a general method for solving problems on graphs of bounded genus, by reducing them to corresponding problems on planar graphs. Our approach is inspired by Bartal's probabilistic approximation of general metrics by trees [4]. We show that for any graph G of bounded genus, there exists a distribution $\mathcal{F}$ over planar graphs, such that for any pair of vertices $u, v \in V(G)$, if we pick a graph H from $\mathcal{F}$, then the expected distance between u and v in H is distorted by at most a constant factor (see later in this section for a precise definition). This in turn implies that a general class of problems on graphs of bounded genus involving optimizing a combination of distances, can be reduced to corresponding problems on planar graphs.

1.1 Our results

We show that every graph of genus g can be $O(1)$-probabilistically approximated by a distribution over graphs of genus $g-1$. By repeatedly applying this procedure g times, we obtain that every graph of genus g can be $2^{O(g)}$-probabilistically approximated by a distribution over planar graphs. In particular, this implies that for graphs of bounded genus, the expected stretch in the above approximation is $O(1)$.

We complement this result by two lower bounds. First, we show that for any $n > 0$, there exist an n-vertex graph that cannot be $o(\log n / \log \log n)$-probabilistically approximated by planar graphs. Since the genus of a graph is at most polynomial in n, this implies that for any $g > 0$, we cannot always have a probabilistic approximation of a graph

SCG'07, June 6–8, 2007, Gyeongju, South Korea.

of genus g by a distribution over planar graphs, with expected expansion $o(\log g/\log\log g)$.

Furthermore, we show that for any $g \geq 1$, and for any $n_0 > 0$, there exists $n \geq n_0$, and an n-vertex graph of genus g, for which any (deterministic) embedding into a graph of genus $g-1$, has distortion $\Omega(n/g)$. This lower bound motivates our use of probabilistic approximations, since in the worst case, any *single* embedding cannot have small distortion.

1.2 Our techniques

Prior work on probabilistic approximation (see later in this section for details), is based on constructing a probability distribution over partitions of the input metric space. The partitions are chosen so that they can later be combined into a simple graph structure.

Our approach deviates from this general framework, since our arguments are inherently topological. Informally, our algorithm for constructing a probabilistic approximation works as follows. We pick a small non-separating cycle C of the input graph G, and compute a *randomly shifted* version of C. Intuitively, this can be seen as shifting C randomly along a handle. We then remove C from G, obtaining a graph of smaller genus. By choosing an appropriate C, we can show that the expected expansion of the distance between any pair of vertices is bounded by some constant.

1.3 Applications

Graph optimization.

As in the case of probabilistically approximating metrics by distributions over tree metrics [4], we obtain a general reduction from a class of graph optimization problems over bounded genus graphs, to a restriction over planar graphs. More precisely, we show the following.

COROLLARY 1. *Let Π be a graph optimization problem in which the objective function depends linearly on the distances of the input graph. If there exists an α-approximation algorithm for Π on planar graphs, then there exists a $O(\alpha)$-approximation algorithm for graphs of bounded genus.*

Embedding into ℓ_1.

Gupta et. al. [14] conjectured that planar graphs embed into ℓ_1 with distortion $O(1)$. The conjecture has been proven for the special case of series-parallel [14], and $O(1)$-outerplanar graphs [8]. Our result implies a strengthening of this conjecture for graphs of bounded genus that was suggested by Thilikos [22]. This strengthening is obtained as follows. Assume that every planar graph embeds into ℓ_1 with distortion c. Then, given a graph G of bounded genus, we can first $O(1)$-probabilistically approximate G by a distribution $\mathcal{F}$ over planar graphs. For each $H \in \mathcal{F}$ there exists an embedding into ℓ_1 with distortion c, and so by concatenating all the embeddings (weighted by their corresponding probabilities), we obtain an embedding of G into ℓ_1 with distortion $O(c)$. We thus show the following result.

COROLLARY 2. *If every planar graph embeds into ℓ_1 with distortion c, then every graph of bounded genus embeds into ℓ_1 with distortion $O(c)$.*

1.4 Related Work

Probabilistic embeddings were first considered by Karp [16], where it was shown that the n-cycle can be $O(1)$-probabilistically approximated by a distribution over spanning subtrees. The first result for general graphs was obtained in [2], where it was shown that any graph can be $2^{O(\log n \log\log n)}$-probabilistically approximated by a distribution over spanning trees.

The notion of probabilistic embeddings was introduced by Bartal in [4], where it was shown that every metric can be $O(\log^2 n)$-probabilistically approximated by a distribution over trees. The same paper gave a $\Omega(\log n)$ lower bound for the same quantity. The upper bound was later improved to $O(\log n \log\log n)$ in [5], and subsequently to $O(\log n)$ in [12]. For the case of approximating by spanning subgraphs, it was shown in [9] that every graph can be $O(\log^2 n \log\log n)$-probabilistically approximated by a distribution over spanning subtrees. When the input is a series-parallel graph, the same bound was improved to $O(\log n)$ in [10].

Probabilistic approximation by distributions over more complex classes of graphs has also been considered. In particular, [21] shows that every metric of constant doubling dimension can be $(1+\epsilon)$-probabilistically approximated by a distribution over graphs of bounded treewidth. In [7] it was shown that there exist graphs of treewidth k that cannot be $o(\log n)$-probabilistically approximated by a distribution over graphs of treewidth $k-3$. In the same paper it was also shown that for any k, there exist planar graphs that cannot be $o(\log n)$-probabilistically approximated by a distribution over graphs of treewidth k.

The later bound implies in particular that in order to obtain a $o(\log n)$-probabilistic approximation for graphs of bounded genus, one needs to consider distributions over more complex families of graphs.

Moreover, the results in [7] imply that there exists a fixed minor H, such that the family of H-minor free graphs cannot be $o(\log n)$-probabilistically approximated by a distribution over planar graphs. Therefore, it is impossible to generalize our positive results to arbitrary minor-free families of graphs.

1.5 Notation and Definitions

Topological graph theory.

Let us recall some notions from topological graph theory (an in-depth exposition can be found in [19]). A *surface* is a two-dimensional manifold. Let S_g be a compact connected orientable surface without boundary, and of genus g. 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. An embedding of G on a surface S is a continuous injection $f : C \to V$. The genus of a graph G is the smallest integer g such that C can be embedded on S_g.

Graphs and metric embeddings.

Unless stated otherwise, all the graphs that we consider in this paper, are assumed to be finite, simple, undirected, weighted, and without loops. For a graph $G=(V,E)$, and for $u,v \in V(G)$, we denote by $D_G(u,v)$ the length of the shortest-path between u and v in G. By scaling the edge weights of G, we may assume that the minimum distance between any pair of points in G is 1.

An *embedding* of a graph G into a graph H is a function $f: V(G) \to V(H)$. The *distortion* of such an embedding is equal to the minimum c, such that there exists $r > 0$, such that for any $u, v \in V(G)$,

$$r \cdot D_G(u,v) \leq D_H(u,v) \leq c \cdot r \cdot D_H(u,v).$$

The embedding f is called an *isometry* if for any $u, v \in V(G)$, $D_H(u,v) = D_G(u,v)$, and it is called *non-contracting* if for any $u, v \in V(G)$, $D_H(u,v) \geq D_G(u,v)$.

We use the notion of probabilistic approximation introduced in [4]. We say that a graph H *dominates* a graph G, if $V(G) \subseteq V(H)$, and for any $u, v \in V(G)$, $D_G(u,v) \leq D_H(u,v)$. For a graph G, a parameter $\alpha \geq 1$, and a probability distribution $\mathcal{F}$ over graphs H that dominate G, we say that $\mathcal{F}$, *α-probabilistically approximates* G, if for any $u, v \in V(G)$,

$$\mathbf{E}_{H \in \mathcal{F}}[D_H(u,v)] \leq \alpha \cdot D_G(u,v).$$

A detailed exposition of combinatorial and algorithmic results concerning metric embeddings, can be found in [18], and [15].

2. PRELIMINARIES

Let G be a graph embedded on S_g, and let C be a cycle of G. Then, C is called *non-separating*, if the corresponding cycle in S_g does not separate S_g.

Given a graph embedded on a surface of bounded genus, it is essential for our algorithm to be able to compute the shortest non-separating cycle of the graph. The first polynomial-time algorithm for this problem was given by Thomassen in [23]. The next lemma states the best currently known time bound, obtained by Cabello and Chambers [6].

Lemma 1. *[[6]] Let G be an n-vertex graph of genus g. Given an embedding of G on S_g, we can compute the shortest non-separating cycle of G, in time $O(g^3 n \log n)$.*

We will now introduce a basic operation that reduces the genus of a graph. Let G be a graph of genus g, embedded on S_g. Let C be a non-separating cycle of G. We say that *the cut along C in G induces a graph G'*, if G' can be obtained from G by cutting S_g along C, and by attaching two copies of C on the boundary of each resulting disk. The following is a standard fact from combibatorian topology (cf. lemma 4.2.4, page 106 of [19]).

Lemma 2. *Let G be a graph of genus g, embedded on S_g, and let C be a non-separating cycle of G. If a cut along C induces a graph G', then the genus of G' is less than g.*

3. CONSTRUCTING THE DISTRIBUTION

Let G be a graph of genus g, embedded on S_g. We will describe an algorithm that computes a distribution $\mathcal{F}$ over graphs of genus at most $g-1$, which $O(1)$-probabilistically approximates G.

By triangulating each face of G, we can assume w.l.o.g. that G is a triangulation of S_g. Note that when triangulating G, for each new edge that we add, we can set its length to be equal to the shortest-path distance between its end-points. This way, the shortest-path metric of the resulting graph remains the same.

We begin by computing a shortest non-separating cycle C of G, using the algorithm from lemma 1. Let k be the length of C. Let G' be the graph induced by the cut along C in G. Let also C_1, and C_2 be the two copies of C in G'.

Claim 1. $D_{G'}(C_1, C_2) \geq k/2$.

Proof. Assume that the assertion is not true. Pick a path p from $v_1 \in C_1$ to $v_2 \in C_2$, with length less than $k/2$. Let v_2' be the copy of v_2 in C_1. Observe that either the clock-wise arc of C_1 from v_1 to v_2', of the clock-wise arc of C_1 from v_2' to v_1 has length at most $k/2$. Let A be this arc. The concatenation of p and A gives a non-separating cycle in G of length strictly less than k, contradicting the minimality of C. $\square$

Let $L = D_{G'}(C_1, C_2)$. For each $\alpha \geq 0$, let V_α be the set of vertices of G' that are at distance α from C_1. Formally, $V_\alpha = \{v \in V(G') | D_{G'}(C_1, v) \leq \alpha\}$.

Let $I = [0, L/2)$. For each $\alpha \in I$, let E_α be the set of edges of G' that are between the sets V_α, and $V(G') \setminus V_\alpha$. That is, $E_\alpha = \{\{u,v\} \in E(G') | u \in V_\alpha, v \notin V_\alpha\}$. By removing E_α from G' we obtain a set of connected components. Moreover, there exists a connected component H_α, such that $C_2 \subseteq H_\alpha$. Define E'_α to be the set of edges between H_α and the rest of G'. Formally, $E'_\alpha = \{\{u,v\} \in E(G') | u \notin H_\alpha, v \in H_\alpha\}$.

Note that E'_α is also a subset of the edge set of G. For each $\alpha \in I$, let G_α be the graph obtained from G as follows. Fix a shortest path P of length L between C_1 and C_2 in G'. First, we remove E'_α from G. Starting from C_1, let z_α be the first vertex of $V(G) \setminus V_\alpha$, visited by P. For each edge $e \in E'_\alpha$, let u be the end-point of e in $V(G) \setminus V_\alpha$. We add an edge $\{u, z_\alpha\}$, with length $D_G(u, z_\alpha)$.

We will use the following result from [1].

Lemma 3 ([1]). *Let $U \subseteq V(G')$ be such that the removal of U from G', disconnects C_1 from C_2. Then, there exists $Z \subseteq U$, which induces a non-separating cycle in G.*

The described construction is depicted in Figure 1. We define $\mathcal{F}$ to be the uniform distribution over $\{G_\alpha\}_{\alpha \in I}$. This is the final distribution that we construct, so it remains to show that it has the claimed properties. First, we need to show that for each $\alpha \in I$, G_α is indeed a connected graph of genus less than g.

Lemma 4. *For each $\alpha \in I$, G_α is a connected graph of genus at most $g-1$.*

Proof. We first argue that G_α is connected. Observe that for each $v \in V(G)$, there exists a path from v either to C_1, or to C_2 in G_α. It follows that there exists a path from v to C in G_α, and since C is a connected subgraph, G_α is connected.

Next, we show that the genus of G_α is strictly less than that of G. Let U be the set of vertices of V_α that are end-points of edges in E'_α. Clearly, U separates C_1 from C_2 in G'. Thus, by lemma 3 there exists $Z \in U$ that induces a non-separating cycle in G. By lemma 2, it follows that G_α has genus at most $g-1$. $\square$

Next, we need to show that every graph in the support of $\mathcal{F}$ dominates the input graph G.

Lemma 5. *For each $u, v \in V(G)$, for each $\alpha \in I$,*

$$D_{G_\alpha}(u,v) \geq D_G(u,v).$$

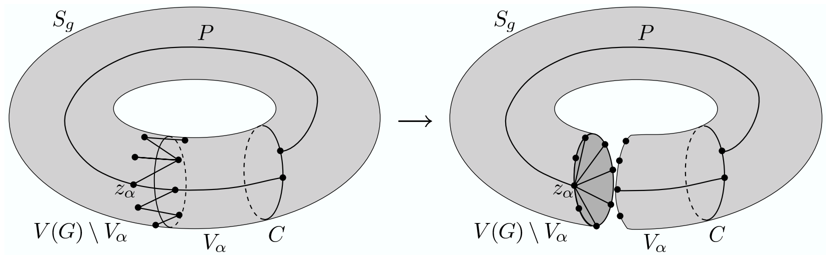

Figure 1: Obtaining the graph G_α from G.

PROOF. Observe that G_α if obtained from G be removing the edges in E'_α, and by adding some new edges. Clearly, by removing the edges in E'_α we do not decrease any distances. Moreover, for each edge that we add, we set its length to be equal to the shortest-path distance between its end-points in G. Therefore, the distance between any two vertices in G_α is at least their distance in G. $\square$

Finally, we need to bound the expected expansion of $\mathcal{F}$.

LEMMA 6. *For each* $\{u, v\} \in E(G)$, $\mathbf{E}_{\Gamma\in\mathcal{F}}[D_\Gamma(u, v)] \leq 8 \cdot D_G(u, v)$.

PROOF. Let $\{u, v\} \in E(G)$, and assume w.l.o.g. that $D_{G'}(C_1, u) \leq D_{G'}(C_1, v)$. Observe that an edge e is in the set E'_α, only if $\alpha \in [D_{G'}(C_1, u), D_{G'}(C_1, v))$. Thus,

$$\begin{aligned}\mathbf{Pr}_{\Gamma\in\mathcal{F}}[e \notin E(\Gamma)] &\leq (D_{G'}(C_1, v) - D_{G'}(C_1, u))\frac{2}{L} \\ &\leq 2 \cdot D_G(u, v)/L\end{aligned}$$

Consider now $\alpha \in I$, such that $e \notin E(G_\alpha)$. We will show that the distance between u and v in G_α is not too large. Since $e \notin E(G_\alpha)$, it follows that e is one of the edges that are removed from G while constructing G_α. That is, $e \in E'_\alpha$. We have that $u \in V_\alpha$, and $v \notin V_\alpha$.

Observe that the shortest path between C_1 and u does not pass through $V(G) \setminus V_\alpha$, and thus it is contained in G_α. Thus, $D_{G_\alpha}(u, C) = D_G(u, C_1) = \alpha \leq L/2$.

Moreover,

$$\begin{aligned}D_{G_\alpha}(v, C_2) &\leq D_{G_\alpha}(v, z_\alpha) + D_{G_\alpha}(z_\alpha, C_2) \\ &\leq D_G(v, z_\alpha) + L \\ &\leq D_G(v, u) + D_G(u, C) + \mathrm{diam}_G(C) \\ &\quad + D_G(C, z_\alpha) + L \\ &\leq D_G(v, u) + \alpha + k/2 + L + L \\ &\leq D_G(v, u) + 7L/2\end{aligned}$$

Putting everything together, we obtain that for any $\{u, v\} \in E(G)$,

$$\begin{aligned}\mathbf{E}_{\Gamma\in\mathcal{F}}[D_\Gamma(u, v)] &\leq D_G(u, v) \cdot \mathbf{Pr}_{\Gamma\in\mathcal{F}}[e \in E(\Gamma)] \\ &\quad + (D_G(v, u) + 7L/2) \cdot \mathbf{Pr}_{\Gamma\in\mathcal{F}}[e \notin E(\Gamma)] \\ &\leq 8 \cdot D_G(u, v)\end{aligned}$$

$\square$

Combining lemmata 5, and 6, we can show the main result of this section.

THEOREM 1. *Every graph of genus g can be $O(1)$-probabilistically approximated by a distribution over graphs of genus at most $g - 1$.*

PROOF. Consider the distribution $\mathcal{F}$ over graphs of genus at most $g - 1$ described above. Let $u, v \in V(G)$. Since by lemma 5 the distances in G_i are non-contracting, it suffices to show that $\mathbf{E}_{\Gamma\in\mathcal{F}}[D_\Gamma(u, v)] = O(D_G(u, v))$. Consider a shortest path $q = x_1, x_2, \ldots, x_t$ between u and v in G. By the linearity of expectation, and by lemma 6,

$$\begin{aligned}\mathbf{E}_{\Gamma\in\mathcal{F}}[D_\Gamma(u, v)] &\leq \sum_{j=1}^{t-1} \mathbf{E}_{\Gamma\in\mathcal{F}}[D_\Gamma(x_j, x_{j+1})] \\ &\leq 8 \cdot \sum_{j=1}^{t-1} D_G(x_j, x_{j+1}) \\ &= 8 \cdot D_G(u, v)\end{aligned}$$

$\square$

COROLLARY 3. *Every graph of genus g can be $2^{O(g)}$-probabilistically approximated by a distribution over planar graphs.*

PROOF. We repeatedly apply theorem 1, g times. $\square$

4. LOWER BOUND FOR PROBABILISTIC APPROXIMATION BY PLANAR GRAPHS

It has been shown by Erdős and Sachs [11] that there exist dense graphs of high girth.

LEMMA 7. *For every $\gamma \geq 3$, and every $n \geq 3$, there exists a connected graph on n vertices, with at least $\frac{1}{4}n^{1+1/\gamma}$ edges, and girth greater than γ.*

THEOREM 2. *For any $n > 0$, there exists an n-vertex graph G, such that any probabilistic approximation of G by a distribution over planar graphs, has expected expansion $\Omega(\log n/\log\log n)$.*

PROOF. Let $n > 0$, and assume that any graph on n vertices can be α-probabilistically approximated by a distribution over planar graphs.

Let H be a graph of girth γ as given by lemma 7, for some $\gamma > 0$ to be defined later. Fix a spanning tree T of H. For each subset of edges $Y \subseteq E(H) \setminus E(T)$, let G_Y be the graph of $V(G)$, with edge set $Y \cup E(T)$. Let A be the family of all possible subgraph of H that contain T. That is,

$$A = \{G_Y | Y \subseteq E(H)\}$$

For each $G_Y \in A$, there exists a distribution $\mathcal{F}_Y$ over planar graphs that α-probabilistically approximates G_Y. Fix $u, v \in V(G_Y)$. By the Markov inequality, we have

$$\mathbf{Pr}_{G'\in\mathcal{F}_Y}[D_{G'}(u, v) > 2\alpha D_{G_Y}(u, v)] < 1/2$$

Thus, if we pick $k = 2\log n$ planar graphs $G_Y^1, \ldots, G_Y^k \in \mathcal{F}_Y$, then the minimum distance between u, and v in all of these graphs, is in the range $[D_{G_Y}(u,v), 2\alpha D_{G_Y}(u,v)]$, with probability at least $1-n^{-2}$. By the union bound, this holds for all pairs $u,v \in V(G_Y)$, with positive probability.

We next show that we can obtain a succint representation of an approximation of G_Y, using $G_Y^1,\ldots,G_Y^k$. Note that this is not immediate, since each G_Y^i might contain steiner nodes. For each $i \in [k]$, pick a collection P_Y^i of shortest-paths of G_Y^i, satisfying the following properties:

- For each $u,v \in V(H)$, there exists a unique shortest-path between u and v in P_Y^i.
- For each shortest-path $p \in P_Y$, and for each $v_1, v_2 \in p$, the subpath q of p between v_1 and v_2, is also in P_Y.

Let J_Y^i be the graph obtained by taking the union of all the paths in P_Y^i, and by replacing induced subpaths, by single edges. Observe that the number of vertices of J_Y^i with degree greater than 2 is at most $2\binom{\binom{n}{2}}{2} < n^4$. This is because by the choise of P_Y^i, each pair of paths in P_Y^i can contribute at most 2 such vertices. Since J_Y^i does not contain induced paths of length greater than 1, it follows that it has at most $n+n^4 < 2n^4$ vertices.

Since each J_Y^i is planar, it follows by a result of [17] that there exist constants $C_1, C_2 > 0$, such that J_Y^i can be embedded into $C_1 \log n$-dimensional ℓ_∞, with distortion C_2. Observe that the distance between any pair of vertices in H is an integer between 1 and n. It follows that when embedding J_Y^i into ℓ_∞, after appropriate scaling, we can round each coordinate to the closest integer in $\{0,1,\ldots,n^2\}$, and incure distortion at most 2. Thus, the restriction of the embedding on $V(G_Y)$ can be represented by at most $C_1 n \log^2 n$ bits. Since we use $k = 2\log n$ distinct samples G_Y^i, we obtain that we can represent the distances in G_Y up to a factor of $4\alpha C_2$, using at most $C_1 n \log^3 n$ bits.

Thus, the total number of distinct representations of graphs G_Y is at most $2^{C_1 n \log^3 n}$. On the other hand, for each $G_Y, G_Z \in A$, with $G_Y \neq G_Z$, there exist $w_1, w_2 \in V(H)$, such that either $\{w_1,w_2\} \in E(G_Y)$, and $\{w_1,w_2\} \notin E(G_Z)$, or $\{w_1,w_2\} \notin E(G_Y)$, and $\{w_1,w_2\} \in E(G_Z)$. Assume w.l.o.g. that $\{w_1,w_2\} \in E(G_Y)$, and $\{w_1,w_2\} \notin E(G_Z)$. Since H has girth more than γ, we have that $D_{G_Y}(w_1,w_2) = 1$, and $D_{G_Z}(w_1,w_2) > \gamma$. If we pick $\gamma > 4\alpha C_2$, it follows that each $G_Y \neq G_Z$ should have distinct representations. Thus, the number of distinct representations is at least the size of the family A. That is

$$\begin{aligned} 2^{C_1 n \log^3 n} &\geq |A| \\ &= 2^{|E(H)|-n+1} \\ &> 2^{\frac{1}{4}n^{1+1/\gamma}-n+1} \\ &> 2^{\frac{1}{8}n^{1+1/(4\alpha C_2)}} \end{aligned}$$

Thus, $\alpha = \Omega(\log n/\log\log n)$. □

5. LOWER BOUND FOR EMBEDDING GRAPHS OF GENUS G INTO GRAPHS OF GENUS $G-1$

In [20] (see also [13]), it is shown that any embedding of the n-cycle into a tree has distortion $\Omega(n)$. This fact motivates the use of probabilistic approximation by distributions over trees. Similarly, we can motivate the use of probabilistic approximation of bounded genus graphs, by showing that any fixed embedding of a graph of genus g into a graph of genus $g-1$, cannot always result in small distortion.

THEOREM 3. *For any $g \geq 1$, and for any $n > 0$, there exist an n-vertex unweighted graph G of genus g, such that any embedding of G into a graph of genus $g-1$ has distortion $\Omega(n/g)$.*

PROOF. The complete graph K_t on t vertices has genus $g > t^2/20$. Let G be the graph obtained from K_t by replacing every edge by a path of length $n/\binom{t}{2}$. Note that G has n vertices. Let g' be the genus of G. Observe that K_t is a minor of G, thus $g' \geq g$. Moreover, every drawing of G induces a drawing of K_t: for each $u,v \in V(K_t)$, the arc connecting $\{u,v\}$ is the image of the path connecting u and v in G. Thus, $g' = g$.

Consider now an embedding f of G into a graph H of genus $g-1$, with distortion c. We need to show a lower bound for c. After scaling the edge lengths of H, we can assume that f is non-contracting. For a graph A, and for a pair of vertices u,v of A, let $P_{u,v}^A$ denote a shortest path between u and v in A. For any $u,v \in K_t$, with $u \neq v$, define $H_{u,v}$ to be the edge-induced subgraph of H containing all the edges in the shortest paths between the edges in $P_{u,v}^G$. Formally, the edge set of $H_{u,v}$ is defined to be

$$E(H_{u,v}) = \bigcup_{\{x,y\} \in E(P_{u,v}^G)} \bigcup_{\{z,w\} \in E(P_{x,y}^H)} \{\{z,w\}\}$$

Consider now an non-crossing embedding σ of H into S_{g-1}. We can construct a (crossing) embedding σ' of K_t into S_{g-1} as follows. For each $u,v \in V(K_t)$, $\sigma(H_{u,v})$ is a connected one-dimensional simplicial complex, which is a subspace of S_{g-1}. Pick a path $\rho_{u,v}$ in $\sigma(H_{u,v})$ between $\sigma(u)$, and $\sigma(v)$. Then, the union of the paths $\rho_{u,v}$, for $u,v \in V(K_t)$ induce an embedding σ' of K_t into S_{g-1}.

Since the genus of K_t is g, there exist $u_1,v_1,u_2,v_2 \in V(K_t)$ such that $\rho_{u_1,v_1} \cap \rho_{u_2,v_2} \neq \emptyset$. We first show that we can pick u_1,v_1,u_2,v_2, so that their end-points are distinct. Assume otherwise, and let $u_1 = u_2$. We can further assume w.l.o.g. that ρ_{u_1,v_1} and ρ_{u_2,v_2} intersect at a single point. Otherwise, we can find two consecutive intersection points p_1, p_2 that can be removed by replacing the sub-path of ρ_{u_1,v_1} between p_1 and p_2 by a slightly translated copy of the sub-path of ρ_{u_1,v_2} between p_1 and p_2.

Let now p be the distinct intersection point of ρ_{u_1,v_1} and ρ_{u_2,v_2}. By exchanging the parts of the two paths that appear between v_1 and p, and by slightly perturbing the two arcs, we can reduce the number of crossings of a_1 and a_2 by one. By repeating the above process, we obtain a drawing with no crossings, contradicting the fact that the genus of K_t is g.

We have thus shown that we can pick distinct u_1,v_1,u_2,v_2, so that $\rho_{u_1,v_1} \cap \rho_{u_2,v_2} \neq \emptyset$. It follows that $H_{u_1,v_1} \cap H_{u_2,v_2} \neq \emptyset$. That is, there exists an edge $\{u_1',v_1'\}$ in P_{u_1,v_1}^G, and and edge $\{u_2',v_2'\}$ in P_{u_2,v_2}^G, such that the paths $P_{u_1',v_1'}^H$ and $P_{u_2',v_2'}^H$ intersect. It follows that there exists $q \in V(P_{u_1',v_1'}^H) \cap V(P_{u_2',v_2'}^H)$, such that for any $i \in \{1,2\}$,

$$D_H(u_i',v_i') = D_H(u_i',q) + D_H(q,v_i').$$

Since the embedding is non-contracting, we have

$$\begin{aligned} D_H(u'_1, u'_2) &\leq D_H(u'_1, q) + D_H(u'_2, q) \\ &\leq D_H(u'_1, v'_1) + D_H(u'_2, v'_2) \\ &\leq cD_G(u'_1, v'_1) + cD_G(u'_2, v'_2) \\ &= 2c \end{aligned}$$

On the other hand, since the end-points of the edges $\{u_1, v_1\}$, and $\{u_2, v_2\}$ are distinct in K_t, we obtain $D_G(u'_1, u'_2) \geq n/\binom{t}{2}$. By non-contraction, $2c = D_H(u'_1, u'_2) \geq D_G(u'_1, u'_2) \geq n/\binom{t}{2}$. Thus, $c = \Omega(n/g)$. □

Acknowledgements.

The authors wish to thank Dimitrios Thilikos for a discussion that led to the formulation of the problem considered in this paper.

6. REFERENCES

[1] M. O. Albertson and J. P. Hutchinson. On the independence ratio of a graph. *J. Graph Theory*, 2:1–8, 1978.

[2] 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.

[3] B. S. Baker. Approximation algorithms for np-complete problems on planar graphs. *J. ACM*, 41(1):153–180, 1994.

[4] Y. Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. *Annual Symposium on Foundations of Computer Science*, 1996.

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

[6] S. Cabello and E. Chambers. Multiple source shortest paths in a genus g graph. In *SODA*, 2007.

[7] D. E. Carroll and A. Goel. Lower bounds for embedding into distributions over excluded minor graph families. In *ESA*, pages 146–156, 2004.

[8] 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*, pages 527–536, 2003.

[9] M. Elkin, Y. Emek, D. Spielman, and S.-H. Teng. Lower-stretch spanning trees. In *Annual ACM Symposium on Theory of Computing*, pages 494–503, 2005.

[10] Y. Emek and D. Peleg. A tight upper bound on the probabilistic embedding of series-parallel graphs. In *SODA*, pages 1045–1053, 2006.

[11] P. Erdős and H. Sachs. Reguläre graphen gegebener taillenweite mit minimaler knotenzahl. *Wiss. Z. Matin-Luther-Univ. Halle-Wittenberg Math.-Natur. Reihe*, 12:251–257, 1963.

[12] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *Annual ACM Symposium on Theory of Computing*, 2003.

[13] A. Gupta. Steiner points in tree metrics don't (really) help. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 220–227, 2001.

[14] A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair. Cuts, trees and l1 embeddings. *Annual Symposium on Foundations of Computer Science*, 1999.

[15] P. Indyk. Tutorial: Algorithmic applications of low-distortion geometric embeddings. *Annual Symposium on Foundations of Computer Science*, 2001.

[16] R. M. Karp. A $2k$-competitive algorithm for the circle. Manuscript, 1989.

[17] 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.

[18] J. Matousek. *Lectures on Discrete Geometry.* Springer, 2002.

[19] B. Mohar and C. Thomassen. *Graphs on Surfaces.* John Hopkins, 2001.

[20] Y. Rabinovich and R. Raz. Lower bounds on the distortion of embedding finite metric spaces in graphs. *Discrete Computational Geometry*, 1998.

[21] K. Talwar. Bypassing the embedding: algorithms for low dimensional metrics. In *Annual ACM Symposium on Theory of Computing*, pages 281–290, 2004.

[22] D. Thilikos. Personal Communication.

[23] C. Thomassen. Embeddings of graphs with no short noncontractible cycles. *J. Combin. Theory, Ser B*, 48:155–177, 1990.