COMPATIBLE CONNECTIVITY-AUGMENTATION OF PLANAR DISCONNECTED GRAPHS

Greg Aloupis,[*] *Luis Barba,*[†] *Paz Carmi,*[‡] *Vida Dujmović,*[§] *Fabrizio Frati,*[¶] *and Pat Morin*[‖]

August 18, 2014

Abstract. Motivated by applications to graph morphing, we consider the following *compatible connectivity-augmentation problem*: We are given a labelled n-vertex planar graph, $\mathcal{G}$, that has $r \geq 2$ connected components, and $k \geq 2$ isomorphic planar straight-line drawings, $G_1, \ldots, G_k$, of $\mathcal{G}$. We wish to augment $\mathcal{G}$ by adding vertices and edges to make it connected in such a way that these vertices and edges can be added to $G_1, \ldots, G_k$ as points and straight-line segments, respectively, to obtain k planar straight-line drawings isomorphic to the augmentation of $\mathcal{G}$. We show that adding $\Theta(nr^{1-1/k})$ edges and vertices to $\mathcal{G}$ is always sufficient and sometimes necessary to achieve this goal. The upper bound holds for all $r \in \{2, \ldots, n\}$ and $k \geq 2$ and is achievable by an algorithm whose running time is $O(nr^{1-1/k})$ for $k = O(1)$ and whose running time is $O(kn^2)$ for general values of k. The lower bound holds for all $r \in \{2, \ldots, n/4\}$ and $k \geq 2$.

[*]Department of Computer Science, Tufts University, aloupis.greg@gmail.com

[†]School of Computer Science, Carleton University and Département d'Informatique, Université Libre de Bruxelles, lbarbafl@ulb.ac.be

[‡]Department of Computer Science, Ben-Gurion University of the Negev, carmip@cs.bgu.ac.il

[§]School of Computer Science and Electrical Engineering, University of Ottawa, vida.dujmovic@uottawa.ca

[¶]School of Information Technologies, The University of Sydney, fabrizio.frati@sydney.edu.au

[‖]School of Computer Science, Carleton University, morin@scs.carleton.ca

1 Introduction

Consider the following problem, which will be more carefully formalized below. We are given several different labelled planar straight-line drawings (or simply drawings) of the same disconnected labelled graph, $\mathcal{G}$. We wish to make $\mathcal{G}$ connected by adding vertices and edges in such a way that these vertices and edges can also be added to the drawings of $\mathcal{G}$ while preserving planarity. The objective is to do this while minimizing the number of edges and vertices added. As the example in Figure 1 shows, it is not always possible to just add edges to $\mathcal{G}$; sometimes additional vertices are necessary.

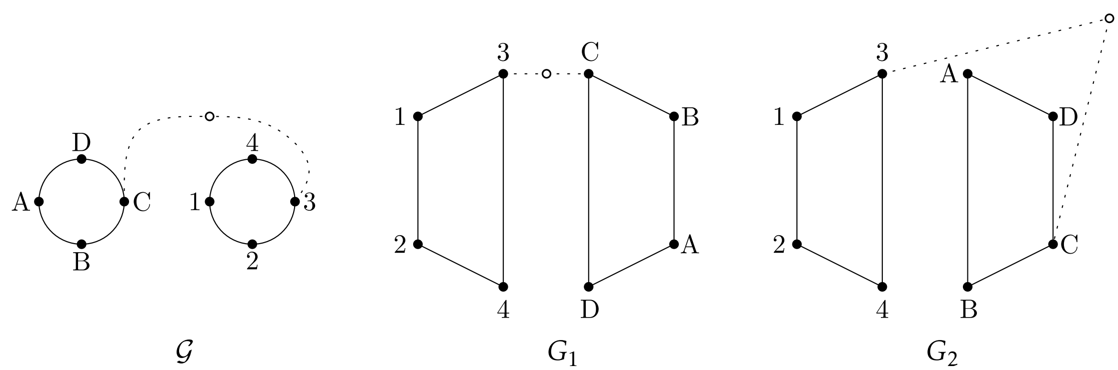

Figure 1: Two drawings, G_1 and G_2, of the same graph, $\mathcal{G}$, where making $\mathcal{G}$ connected requires the addition both of edges and vertices. In this case, $\mathcal{G}$ is made connected by adding the hollow vertex and two dashed edges.

The motivation for this work comes from the problem of morphing planar graphs, which has many applications [**?**, **?**, **?**, **?**, **?**] including computer animation. Imagine an animator who wishes to animate a scene in which a character's expression goes from neutral, to surprised, to happy (see Figure 2). The animator can draw these three faces, but does not want to hand-draw the 30–60 frames required to animate the change of expression. The strokes used to draw the character's features can be converted into paths and these can be merged into components corresponding to the character's eyes, nose, mouth and so on. A correspondence between the same elements in different pictures is also given.[1] Thus, the input is three isomorphic drawings of the same planar graph.

In this setting, animating the face becomes a problem of *morphing* (i.e., continuously deforming) one drawing of a planar graph into another drawing of the same planar graph while maintaining planarity of the drawing throughout the deformation. This morphing problem has been studied since 1944, when Cairns [**?**] showed such a transformation always exists. Since then, a sequence of results has shown that morphs can be done effi-

[1] In many cases, the correspondence is a byproduct of the creation process. For example, in Figure 2, the second two faces were obtained by copying and then editing the first one.

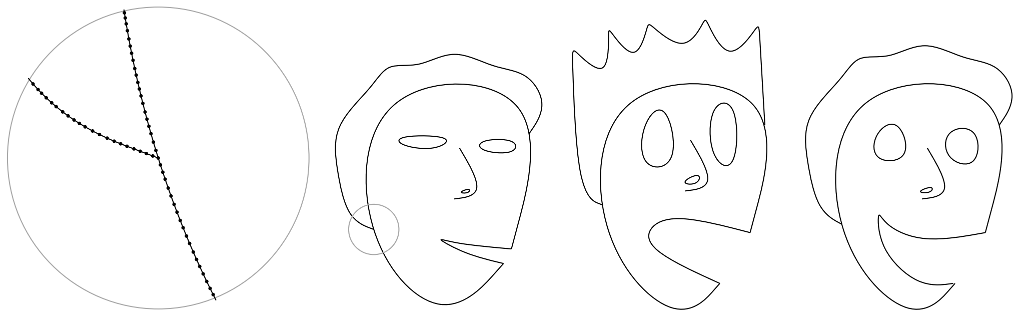

Figure 2: Computer-assisted animation frequently involves morphing between a sequence of drawings of the same planar graph. Zooming in on a section of the image reveals that the artist's strokes are approximated by polygonal paths

ciently, so that the motion can be described concisely [**?**, **?**, **?**, **?**]. The most recent such result [**?**] shows that any planar drawing of an n-vertex *connected* planar graph can be morphed into any isomorphic drawing using a sequence of $O(n)$ *linear morphs*, in which vertices move along linear trajectories at constant speed.

The morphing algorithms discussed above require that the input graph, $\mathcal{G}$, be connected. In many applications of morphing (for example in Figure 2) the input graph is not connected. Before these morphing algorithms can be used, $\mathcal{G}$ must be augmented into a connected graph, $\mathcal{H}$, but this augmentation must be compatible with the drawings of $\mathcal{G}$. At the same time, the complexity of the morph produced by a morphing algorithm depends on the number of vertices of $\mathcal{H}$. Therefore, we want to find an augmentation with the fewest number of vertices. This motivates the theoretical question studied in the current paper.

1.1 Formal Problem Statement and Main Result

A *drawing* of a graph $\mathcal{G} = (V, E)$ is a one-to-one function $\psi\colon V \to \mathbb{R}^2$. A drawing is *planar* if (a) for every pair of edges uw and xy in E, the open line segment with endpoints $\psi(u)$ and $\psi(w)$ is disjoint from the open line segment with endpoints $\psi(x)$ and $\psi(y)$ and (b) for every edge uw and every vertex y, $\psi(y)$ is not contained in the open line segment with endpoints $\psi(u)$ and $\psi(w)$. Two planar drawings, ψ_1 and ψ_2, of $\mathcal{G}$ are *isomorphic* if there exists a continuous family of planar drawings $\{\psi^{(t)}\colon 0 \le t \le 1\}$ of $\mathcal{G}$ such that $\psi^{(0)} = \psi_1$ and $\psi^{(1)} = \psi_2$.[2]

We call a graph, $G = (V, E)$ a *geometric planar graph* if it is the image of some planar drawing of a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. That is, $V(G) = \{\psi(v) : v \in V(\mathcal{G})\}$, $E(G) = \{(\psi(u), \psi(w)) : (u, w) \in E(\mathcal{G})\}$, and ψ is a planar drawing of $\mathcal{G}$. When clear from context, we will sometimes

[2]By Cairn's result, this is equivalent to saying that the two drawings of G have the same rotation schemes, the same cycle-vertex containment relationship, and the same outer face.

treat a geometric planar graph interchangeably with the set of points and line segments defined by its vertices and edges, respectively.

We will avoid repeatedly referencing drawing functions like ψ. Instead, we will talk about a graph $\mathcal{G}$ and k isomorphic drawings $G_1,\ldots,G_k$ of $\mathcal{G}$. This means that each G_i is the geometric graph given by the drawing of $\mathcal{G}$ with some function ψ_i and that $\psi_1,\ldots,\psi_k$ are pairwise isomorphic. When necessary, we may talk about the vertex v in G_i where v is actually a vertex of $\mathcal{G}$; this should be taken to mean the vertex $\psi_i(v)$ in G_i.

We are now ready to state the main problem studied in this paper. Given $k > 1$ planar isomorphic drawings $G_1,\ldots,G_k$ of $\mathcal{G}$, a *compatible augmentation*, $\mathcal{H}$, of $\mathcal{G}$ is a supergraph of $\mathcal{G}$ such that (1) $\mathcal{H}$ is connected, and (2) there exist planar isomorphic drawings, $H_1,\ldots,H_k$, of $\mathcal{H}$ such that $H_i \supset G_i$ for every $i \in \{1,\ldots,k\}$. We prove the following result:

Main Result: *If $\mathcal{G}$ is a graph with n vertices and $r \in \{2,\ldots,n\}$ connected components and $G_1,\ldots,G_k$, $k \geq 2$, are isomorphic planar drawings of $\mathcal{G}$, then there always exists a compatible augmentation of $\mathcal{G}$ whose size is $O(nr^{1-1/k})$.*

Furthermore, this bound is tight; for every $r \in \{2,\ldots,\lfloor n/4 \rfloor\}$ and $k \geq 2$, there exists a graph $\mathcal{G}$ with r components and k isomorphic planar drawings for which any compatible augmentation has size $\Omega(nr^{1-1/k})$.

These results show that the (worst-case) cost of an augmentation is very sensitive to the number, k, of drawings, but only up to a point. For a fixed value of r, our bounds range from $\Theta(nr^{1/2})$ (when $k = 2$) to $\Theta(nr)$ (for $k \geq \log r$). On the other hand, for the common case where $k = 2$, our bounds vary from $\Theta(n)$ (when $r \in O(1)$) up to $\Theta(n^{3/2})$ (when $r \in \Theta(n)$). Neither r nor k causes the complexity of the augmentation to blow up beyond $\Theta(n^2)$.

1.2 Related Work

To the best of our knowledge, there is little work on compatible connectivity-augmentation of planar graphs, though there is work on isomorphic triangulations of polygons. Refer to Figure 3. In this setting, the graph $\mathcal{G}$ is a cycle and one has two non-crossing drawings, P and Q, of $\mathcal{G}$. The goal is to augment $\mathcal{G}$ (and the two drawings P and Q) so that $\mathcal{G}$ becomes a near-triangulation, and P and Q become (geometric) triangulations of the interiors of the polygons whose boundaries are P and Q. Aronov *et al.* [**?**] showed that this can always be accomplished with the addition of $O(n^2)$ vertices and that $\Omega(n^2)$ vertices are sometimes necessary. Kranakis and Urrutia [**?**] showed that this result can be made sensitive to the number of reflex vertices of P and Q, so that the number of triangles required is $O(n+pq)$ where p and q are the number of reflex vertices of p and q, respectively.

Babikov *et al.* [**?**] showed that the result of Aronov *et al.* can be extended to polygons with holes. This work is the most closely related to ours because it encounters (the special case $k = 2$ of) our problem as a subproblem. In their setting, the graph $\mathcal{G}$ is a collection of

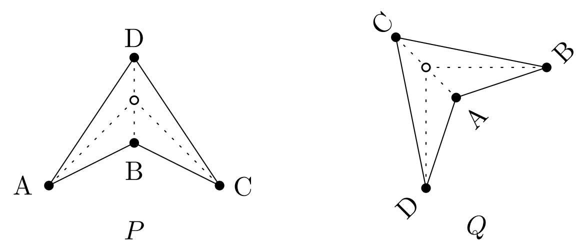

Figure 3: Compatible triangulations of two 4-gons P and Q.

r cycles, the drawings P and Q are such that one cycle, $\mathcal{C}$, of $\mathcal{G}$ contains all the others in its interior and no other pair of cycles is nested in P or Q. In the first stage of their algorithm, they build a connected supergraph $\mathcal{H}'$ of $\mathcal{G}$, but their supergraph has size $\Theta(n^2)$ in the worst case. The main theorem in the current paper shows that this step of their algorithm could be done with a graph $\mathcal{H}'$ having only $O(nr^{1/2})$ edges (but completing this graph to a triangulation may still requires $\Omega(n^2)$ edges in the worst case).

Finally, several papers have dealt with the problem of increasing the connectivity of a (single) geometric planar graph while adding few vertices and edges. Abellenas *et al.* [**?**] consider the problem of adding edges to a planar drawing in order to make it 2-edge connected and showed that $\lfloor(2n-2)/3\rfloor$ edge are sometimes necessary and $6n/7$ edges are always sufficient. Tóth [**?**] later obtained the tight upper-bound of $\lfloor(2n-2)/3\rfloor$ for the same problem. Rutter and Wolff [**?**] show that finding the minimum number of edges required to achieve 2-edge connectivity is NP-hard.

1.3 Outline

To guide the reader, we give a rough sketch of our upper bound proof, which is illustrated in Figure 4. We will assume, for the sake of simplicity, that every component has at least one vertex incident to the outer face.

For each component, $\mathcal{C}_i$, of $\mathcal{G}$ we select a distinguished *corner*, a_i, of G_1. (A corner is the space between two consecutive edges incident to some vertex of the outer face). The corner a_i is called the *attachment corner* for component $\mathcal{C}_i$. Notice that, since $G_1,\ldots,G_k$ are isomorphic, a_i appears as a corner in each of $G_1,\ldots,G_k$. The augmentation that we ultimately create will consist of a path that connects to each component $\mathcal{C}_i$ at its attachment corner a_i.

Next, for each drawing, G_j, we add $r-1$ edges to make G_j into a connected graph, G_j^*; these edges are not, in general, edges that take part in the final augmentation of $G_1,\ldots,G_k$ (in fact the edges of G_j^* not in G_j might be different from the edges of G_h^* not in G_h, if $j \neq h$). We then traverse the boundary of the outer face of G_j^* to obtain a polygonal path, φ_j, of length $O(n)$ that comes close to every corner of G_j. The path φ_j is then used

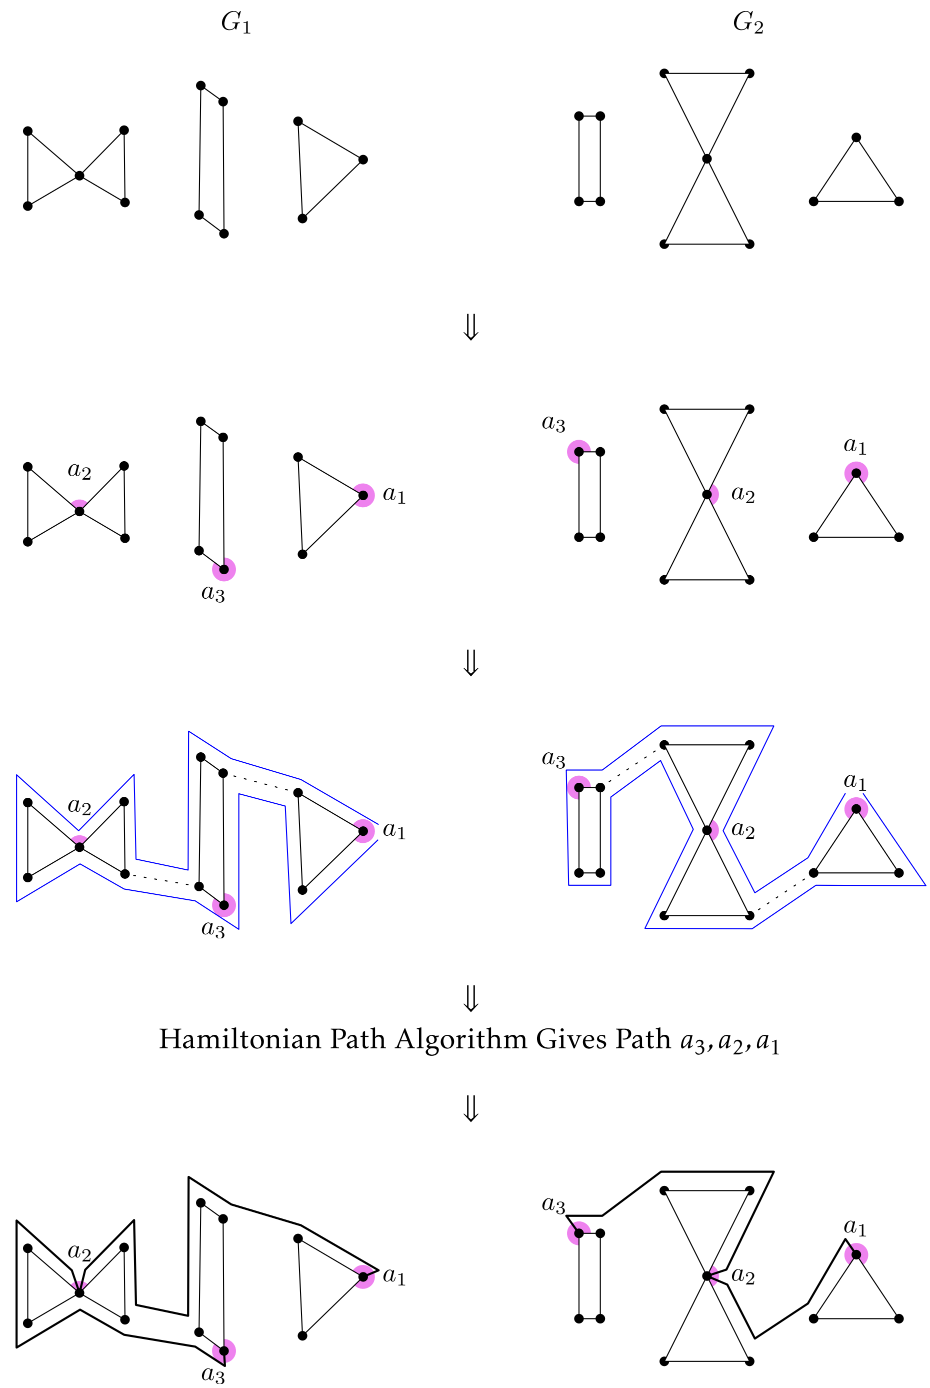

Figure 4: The algorithm for making a compatible augmentation of $\mathcal{G}$ works by defining corners $a_1, \ldots, a_r$, taking a spanning path on each augmented drawing G_i^* that visits all corners, and using this information to compute a permutation of $a_1, \ldots, a_r$ that can be drawn efficiently in each G_i.

to define an integer distance $d_j(a_\ell, a_m)$ for any two attachment corners a_ℓ and a_m. This distance includes information about the number of edges of φ_j between a_ℓ and a_m as well the sizes of the some of the components visited while walking from a_ℓ to a_m along φ_j.

Next, we take a leap into k dimensions by using the distance functions $d_1,\ldots,d_k$ to produce a k-dimensional point set $X = \{x_1,\ldots,x_r\}$ that lives in a hypercube of side-length $O(n)$. This mapping has the property that, by adding a path of length $O(\|x_\ell - x_m\|_\infty)$ to $\mathcal{G}$, the attachment corners a_ℓ and a_m can be joined in each of $G_1,\ldots,G_k$ while preserving planarity.

Now, since the point set X is in $\mathbb{R}^k$, has r points, and lives in a hypercube of side-length $O(n)$, a classic argument about the geometric Travelling Salesman Problem [**?**, **?**] implies that it has a spanning path whose length, measured in the ℓ_∞ norm, is $O(nr^{1-1/k})$.[3] This implies that $\mathcal{G}$ can be made connected with a collection of $r-1$ paths, whose endpoints are the corners $a_1,\ldots,a_r$, having total size $O(nr^{1-1/k})$, and that each of these paths can be drawn in a planar fashion in each of $G_1,\ldots,G_k$.

At this point, all that remains is to show that each of these $r-1$ paths be drawn in each of $G_1,\ldots,G_k$ without crossing each other. This part of the proof involves carefully winding these paths around the components in $G_1,\ldots,G_k$ using paths close to the paths $\varphi_1,\ldots,\varphi_k$ defined above. This part of the proof resembles the first part of the proof of Babikov *et al.* [**?**], but is complicated by the fact that we have to be quite careful that the number of edges in these paths remains in $O(nr^{1-1/k})$.

The remainder of the paper is organized as follows: In Section 2 we start by solving the special case in which the graph G has no edges. This special case is already non-trivial and introduces some of the main ideas used in solving the full problem, which is tackled in Section 3. Section 4 presents a lower bound construction that matches our upper bound.

2 Upper bounds for trivial components

As a warmup, we consider a (trivial) graph containing n vertices and no edges. Before constructing a compatible augmentation, we provide a subroutine that constructs a "short" planar spanning path of a given ordered set of points.

2.1 Spanning paths of point sets

Let S be a set of n points in the plane with distinct x-coordinates. Given a point $v \in S$, let $\mathrm{rank}(v)$ denote the number of points of S that lie to the left of (having smaller x-coordinate than) v.

Given an arbitrary order $(v_1, v_2, \ldots, v_n)$ of the points of S, we want to construct a

[3] The original argument was for the standard ℓ_2 norm and has an extra $\sqrt{k}$ factor. This factor disappears in the ℓ_∞ norm. See Appendix A for details.

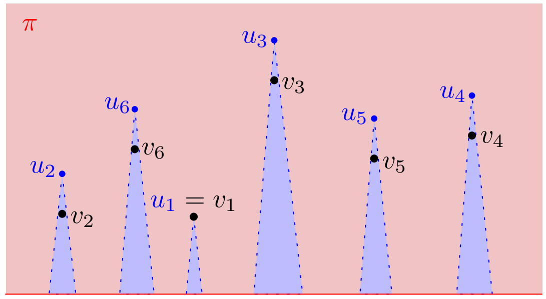

Figure 5: The halfplane π and the cones $\Delta_1,\ldots,\Delta_n$ with apexes at $u_1,\ldots,u_n$.

path R that connects them in this order and such that:

$$|R| = O\left(\sum_{i=1}^{n-1} |\operatorname{rank}(v_i) - \operatorname{rank}(v_{i+1})|\right).$$

This paper uses the $|\cdot|$ operator in several different ways, depending on the type of its argument. For a real number, x, $|x|$ is the absolute value of x. For a walk, $R = (r_0,\ldots,r_k)$, $|R| = k$ denotes the number of edges traversed by R. For a (weakly-)simple polygon, P whose vertices—as encountered during a counterclockwise traversal—are $(a_1,\ldots,a_k)$, $|P| = k$, denotes the number of edges of P.

Consider a horizontal line ℓ below S and let π be the closed halfspace supported by ℓ that contains S. We present an algorithm that constructs R iteratively; during the ith iteration of the algorithm, the path is extended with $O(|\operatorname{rank}(v_i) - \operatorname{rank}(v_{i-1})|)$ vertices to include v_i. For each $i \in \{1,\ldots,n\}$, after the ith iteration of the algorithm, we maintain the invariant that ℓ does not intersect R, and we also maintain the *escape invariant* which is defined as follows: For each $j \in \{i,\ldots,n\}$, there is a closed cone Δ_j with apex u_j such that (1) u_j lies above v_j and has the same x-coordinate as v_j ($u_j = v_j$ if $j = i$), (2) Δ_j contains v_j and no other point of S, (3) Δ_j contains the ray originating at v_j in the direction of the negative y-axis, (4) Δ_j does not intersect R, and (5) Δ_h and Δ_j are disjoint inside π, for every $h \in \{i,\ldots,n\}$ with $h \neq j$.

Initialize R as a path that consists of the single vertex v_1. In order to establish the escape invariant, we define $u_1 = v_1$; also, for each $j \in \{2,\ldots,n\}$, we define u_j as an arbitrary translation up of v_j; further, for each $j \in \{1,\ldots,n\}$, we let Δ_j be a cone with apex on u_j sufficiently narrow so that these cones do not intersect inside π; see Figure 5.

Now assume that R is a path connecting v_1 with v_j, for some $j \in \{1,\ldots,n-1\}$. We extend R by appending a path that connects v_j with v_{j+1}.

First, we translate Δ_{j+1} down until its apex u_{j+1} coincides with v_{j+1}. Let π_j be the closure of the set obtained from π by removing Δ_h, for every $h \in \{j, j+1,\ldots,n\}$; see Figure 6 (right). That is, π_j is a halfspace with dents made by the removal of $n-j+1$ cones. Observe

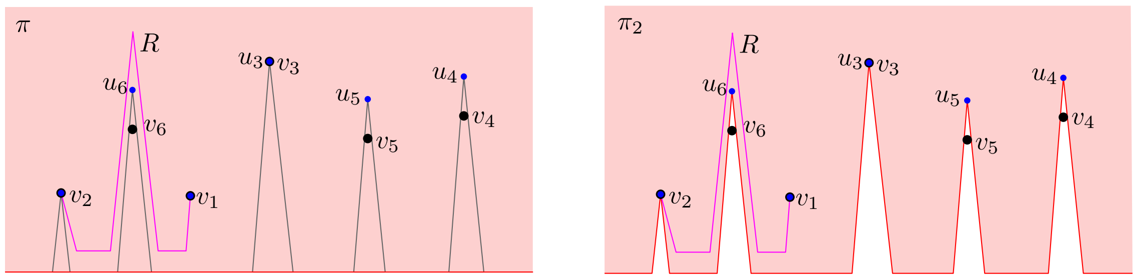

Figure 6: The boundary of π_2 is in the path from v_2 to v_3.

that, for every pair of apexes u_i and u_h, with $i,h \geq j$, the boundary of π_j contains a path from u_i to u_h with $O(|\text{rank}(v_i) - \text{rank}(v_h)|)$ edges. Because ℓ does not intersect R and by the escape invariant, the boundary of π_j intersects R only at v_j. Moreover, again by the escape invariant, for each v_i with $i > j$, v_i lies outside of π_j except for v_{j+1} that lies on its boundary. Because both v_j and v_{j+1} lie on the boundary of π_j, which does not intersect R other than at v_j, we can connect v_j with v_{j+1} via a path contained in the boundary of π_j with length $O(|\text{rank}(v_j) - \text{rank}(v_{j+1})|)$. In this way, we extend R to a planar path that contains $v_1, \ldots, v_{j+1}$.

After connecting v_j with v_{j+1}, for each $h \in \{j+2, \ldots, n\}$, either Δ_h is disjoint from R, or it shares some portion of its boundary with R. However, the interior of Δ_h does not intersect R. To preserve the escape invariant, for each $h \in \{j+2, \ldots, n\}$, we translate π and Δ_h downwards by a sufficiently small amount, ε, and we scale Δ_{j+1} horizontally down, while keeping its apex at v_{j+1}. To conclude, each translated or scaled cone is contained in the previous one, u_h lies above v_h, for each $h \in \{j+2, \ldots, n\}$, and u_{j+1} coincides with v_{j+1}. Therefore, by choosing ε sufficiently small, we maintain the escape invariant and obtain the following result.

Lemma 2.1. *Given an order $(v_1, \ldots, v_n)$ of the vertices of S, there exists a planar path R that connects every point of S in the given order such that the number of vertices of R between v_i and v_{i+1} is $O(|\text{rank}(v_i) - \text{rank}(v_{i+1})|)$, for each $i \in \{1, \ldots, n-1\}$.*

Proof. Recall that in each iteration, the algorithm computes a path connecting v_j with v_{j+1} that does not cross the portion of the path already constructed. Because this invariant is maintained throughout, the resulting path is planar.

Since the path that connects v_j with v_{j+1} follows the boundary of π_j and since this boundary has length $O(|\text{rank}(v_j) - \text{rank}(v_{j+1})|)$ between v_j and v_{j+1}, the path that connects v_j with v_{j+1} has length $O(|\text{rank}(v_j) - \text{rank}(v_{j+1})|)$. Consequently, the total length of R is given by $O\left(\sum_{i=1}^{n-1} |\text{rank}(v_i) - \text{rank}(v_{i+1})|\right)$. □

2.2 Compatible drawings of point sets

Recall that in this section $\mathcal{G}$ is a graph with n trivial components. Let $G_1,\ldots,G_k$ be $k > 1$ isomorphic drawings of $\mathcal{G}$, i.e., G_i is a sequence of n points in the plane. Assume without loss of generality that no two points of G_i share the same x-coordinate. Given a vertex v of $\mathcal{G}$, let $\mathrm{rank}_{G_i}(v)$ denote the number of points of G_i having smaller x-coordinate than v, and let $x_v = (\mathrm{rank}_{G_1}(v),\ldots,\mathrm{rank}_{G_k}(v))$ be a point in the integer grid $[0; n-1]^k$ in $\mathbb{R}^k$. Let $X = \{x_v : v \in V(\mathcal{G})\}$ and let P be the shortest Hamiltonian path of X when distance is measured using the ℓ_∞ norm, so that

$$\|x_v - x_u\|_\infty = \max\{|\mathrm{rank}_{G_i}(v) - \mathrm{rank}_{G_i}(u)| : i \in \{1,\ldots,k\}\} \ .$$

It is known that the length of P is $O(n^{2-1/k})$ (see Corollary A.2 in Appendix A). Note that the order of the points of P induces an order on the vertices of $\mathcal{G}$ and hence, an order on the vertices of each G_i.

Theorem 2.2. *For each $i \in \{1,\ldots,n\}$, we can construct a path R_i of length $O(n^{2-1/k})$ that connects every point of G_i so that $G_i \cup R_i$ is planar. Moreover, for any distinct $i, j \in \{1,\ldots,n\}$, $G_i \cup R_i$ and $G_j \cup R_j$ are isomorphic.*

Proof. By relabelling, let $(v_1,\ldots,v_n)$ denote the order of the vertices of $\mathcal{G}$ induced by the Hamiltonian path, P, described above. The graph $\mathcal{H}$, which is an augmentation of $\mathcal{G}$, is a path that visits the vertices $v_1,\ldots,v_n$ in this order. Letting d_j denote the ℓ_∞ distance between x_{v_j} and $x_{v_{j+1}}$, the path $\mathcal{H}$ includes an additional $O(d_j)$ vertices between v_j and v_{j+1}. It follows that the number of vertices in $\mathcal{H}$ is proportional to the length of P, which is $O(n^{2-1/k})$.

For each G_i, we use Lemma 2.1 to draw $\mathcal{H}$ as a planar path, R_i, that connects the vertices $v_1,\ldots,v_n$ in this order in the drawing G_i. Since $d_j \geq |\mathrm{rank}_{G_i}(v_j) - \mathrm{rank}_{G_i}(v_{j+1})|$, the $O(d_j)$ vertices in $\mathcal{H}$ between v_j and v_{j+1} are enough to draw the $O(|\mathrm{rank}_{G_i}(v_j) - \mathrm{rank}_{G_i}(v_{j+1})|)$ vertices in R_i between v_j and v_{j+1}. Since the vertices of each G_i are connected in the same order, $G_i \cup R_i$ is isomorphic to $G_j \cup R_j$ for each $i, j \in \{1,\ldots,k\}$. □

3 The general problem

In this section, we extend the result presented in Section 2 to graphs with non-trivial components. We follow the same general scheme used in Section 2 for the case of trivial (isolated vertex) components: We define k different orderings of the components of $\mathcal{G}$ and use these orderings (and the sizes of these components) to define an r-point set, X, in $\mathbb{R}^k$. A short path that visits all points in X is then translated back into a short path, R, that visits all components of $\mathcal{G}$. The path R is then added, as a polygonal path, R_i, to each drawing, G_i, of $\mathcal{G}$.

Unlike the case in which components are isolated vertices, there is no natural ordering of the components of G_i, so we must define one. Also, the drawing of path R_i is considerably more complicated. In Section 2, R_i is drawn incrementally, and always passes above components that are not yet included in R_i and below components that are already included in R_i. In this section, we redefine "above" and "below". The number of edges required to go above or below a component depends on its size and structure.

3.1 Preliminaries

Let C be a connected geometric planar graph. Let $v_0, v_1, \ldots, v_k, v_0$ be the sequence of vertices of C visited by a counterclockwise Eulerian tour along the boundary of the outer face of C. Note that v_i may be equal to v_j for some $i \neq j$. A vertex v_i in this sequence is called a *corner* of C. We consider the boundary of C, denoted by ∂C, to be the boundary of the weakly-simple polygon $(v_0, \ldots, v_k, v_0)$ whose vertex set is the set of corners of C.[4]

Let $\varepsilon > 0$. For each corner v_i of ∂C, let ℓ_i be the half-line starting at v_i that bisects the angle between the edges $v_{i-1}v_i$ and v_iv_{i+1} in the outer face of C. Let z_i be the point at distance ε from v_i along ℓ_i. We call z_i the *ε-copy* of v_i. Let $\partial_\varepsilon C$ be the piecewise-linear cycle defined by the sequence $(z_0, z_1, \ldots, z_k, z_0)$. We call $\partial_\varepsilon C$ the *ε-fattening* of C. An ε-fattening $\partial_\varepsilon C$ is *simple* if $\partial_\varepsilon C$ is a simple polygon that contains C. Note that $\partial_\varepsilon C$ is simple, provided that ε is sufficiently small. In this paper, we consider only simple ε-fattenings; see Figure 7. Note that the number of edges between two corners of ∂C along the boundary of C is the same as the number of edges between their ε-copies along $\partial_\varepsilon C$.

3.2 Connected augmentations

Let G be a geometric planar graph with r connected components such that each component is adjacent to the outer face. Two vertices are *visible* if the open segment joining them does not intersect G. Let T_G be a smallest set of edges of the visibility graph of G that need to be added to G to make it connected. As there are always two components containing mutually visible vertices, we can connect them and repeat recursively. Thus, T_G has $r-1$ edges. (Loosely, we can think of T_G as a spanning tree of G's components.) Let $G^* = G \cup T_G$. We say that G^* is a *connected augmentation* of G; see Figure 7.

Let $C_1, \ldots, C_r$ be the components of G. For each $i \in \{1, \ldots, r\}$, let $a_i \in C_i$ be an arbitrary corner of ∂C_i (note that a_i is adjacent to the outer face). We call a_i the *attachment corner* of C_i.

Let φ be the path on the corners of ∂G^* (hence φ is also a walk on the vertices of ∂G^*) obtained by splitting ∂G^* at the corner a_1. That is, φ is a path that visits every corner

[4]More formally, ∂C is the boundary of the unbounded component of $\mathbb{R}^2 \setminus C$, when we treat C as the union of all its edges (line segments) and vertices (points).

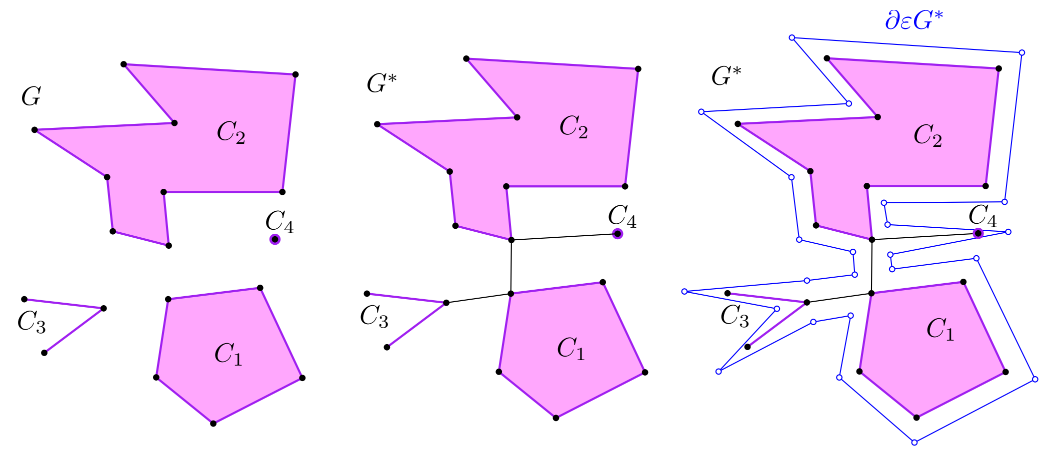

Figure 7: The graph G (left); a connected augmentation, G^*, of G (middle); and the ε-fattening, $\partial_\varepsilon G^*$ (right).

of ∂G^* exactly once except for a_1, which is visited twice. Given two corners u and v in ∂G^*, let $\varphi(u,v)$ denote the unique path in φ that connects u with v. Let $A(u,v)$ be the set of attachment corners of G visited by $\varphi(u,v)$. For two attachment corners $u = a_i$ and $v = a_j$, define

$$\sigma_G(u,v) = |\varphi(u,v)| + 2 \sum_{a_t \in A(u,v)} |\partial C_t| - |\partial C_i| - |\partial C_j|,$$

which we call the *cost* of going from u to v. The definition of $\sigma_G(u,v)$ is designed to capture the fact that, if u and v occur consecutively on the path R we construct, then the portion of R between u and v will have length at least $|\varphi(u,v)|$ since it follows φ, and it may also take a detour around every attachment corner $a_i \in A(u,v)$. If it takes this detour at some $a_i \in A(u,v)$, then it does so by walking around ∂C_i which requires an additional $|\partial C_i|$ edges.

Lemma 3.1. *If $u = a_i$ is an attachment corner of G, then $\sigma_G(a_1,u) < 6n$. Moreover, if $v = a_j$, with $j > i$, is another attachment corner of G, then $\sigma_G(u,v) = \sigma_G(v,u) = \sigma_G(a_1,v) - \sigma_G(a_1,u)$.*

Proof. Recall that G^* is a graph with n vertices, so ∂G^* is a weakly-simple polygon with at most n distinct vertices, so $|\partial G^*| \le 2n-2 < 2n$. Furthermore, $\varphi(a_1,u) \subset \partial G^*$, so $|\varphi(a_1,u)| \le |\partial G^*| < 2n$. Similarly, $2\sum_{a_i \in A(a_1,u)} |\partial C_i| \le 2\sum_{i=1}^{r} |\partial C_i| < 4n$. Therefore, $\sigma_G(a_1,u) < 6n$, which proves the first part of the lemma.

To prove the second part of the lemma, first observe that $\sigma_G(u,v) = \sigma_G(v,u)$ by definition. To prove the second equality, denote the relevant attachment corners of G^* by $a_1, a_2, \ldots, a_i = u, a_{i+1}, \ldots, a_j = v$. Then $A(a_1,u) = \{a_1,\ldots,a_i\}$, $A(a_1,v) = \{a_1,\ldots,a_j\}$, and

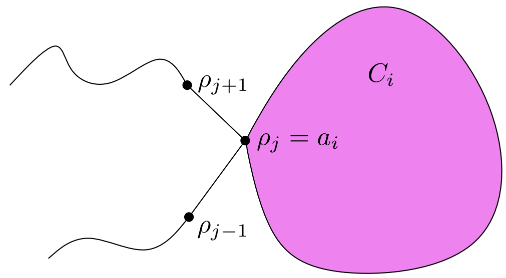

Figure 8: The component C_i is to the right of the path $R = (\rho_1, \ldots, \rho_t)$.

$A(u,v) = \{a_i, \ldots, a_j\}$, so

$$\begin{aligned}\sigma_G(a_1,v) - \sigma_G(a_1,u) &= |\varphi(a_1,v)| - |\varphi(a_1,u)| + 2\sum_{t=1}^{j}|\partial C_t| - 2\sum_{t=1}^{i}|\partial C_t| \\ &\quad - |\partial C_1| + |\partial C_1| - |\partial C_j| + |\partial C_i| \\ &= |\varphi(u,v)| + 2\sum_{t=i+1}^{j}|\partial C_t| - |\partial C_j| + |\partial C_i| \\ &= |\varphi(u,v)| + 2\sum_{t=i}^{j}|\partial C_t| - |\partial C_j| - |\partial C_i| \\ &= |\varphi(u,v)| + 2\sum_{a_t \in A(u,v)}|\partial C_t| - |\partial C_j| - |\partial C_i| \\ &= \sigma_G(u,v)\ .\end{aligned}$$

□

3.3 Spanning paths for connected augmentations

Let $a_1, \ldots, a_r$ be an arbitrary order of the attachment corners of G (we can get the incremental indexing by relabeling the components). Given a path $R = (\rho_1, \rho_2, \ldots, \rho_t)$ that passes through all attachment corners of G, we say that C_i *lies to the right of* R if (1) a_i is the only vertex of C_i that belongs to $V(R)$, and (2) if $a_i = \rho_j$ for some $j \in \{1, \ldots, t\}$, then ρ_{j-1} and ρ_{j+1} appear as consecutive vertices when sorting—in the graph $G \cup R$—the neighbors of a_i in clockwise order around a_i (see Figure 8).

We now show how to construct a path R that connects the attachment corners of G in the given order, i.e., if $i < j$, then a_i is visited before a_j by R. We want to construct R so that each component C_i of G lies to the right of R and so that $G \cup R$ is a planar geometric graph. Moreover, we want the subpath of R between a_j and a_{j+1} to have $O(\sigma_G(a_j, a_{j+1}))$ vertices. We initialize R with the trivial path that contains only a_1, and then extend R iteratively, so that each new corner a_i is included in R. Recall that for any given $\varepsilon > 0$, $\partial_\varepsilon G^*$ denotes the ε-fattening of G^* (see Section 3.1). Let $\mu > 0$ be a small constant to be specified

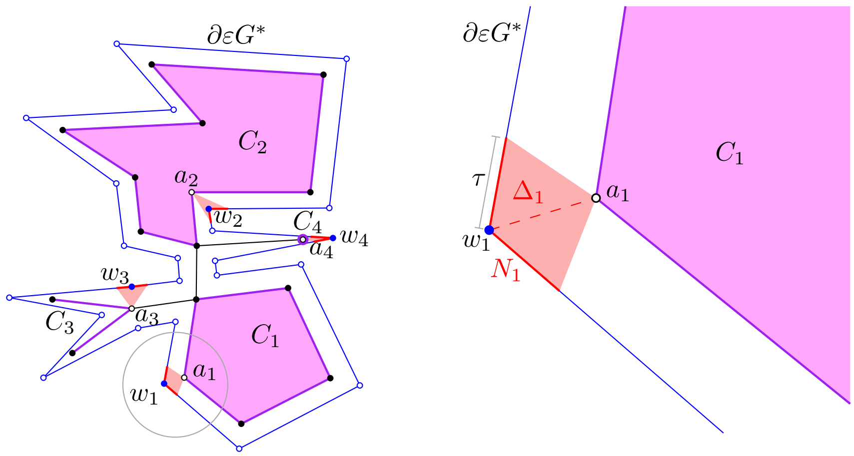

Figure 9: The ϵ-fattening of G^* and the "cones" $\Delta_1, \ldots, \Delta_r$.

later. Initially, let $\varepsilon = 2\mu$ and let $\delta = \mu/2$. Let $\lambda < \mu/2^{r+1}$ be a constant sufficiently small so that $\partial_\lambda C_i \cap \partial_\lambda C_j = \emptyset$ for any distinct $i, j \in \{1, \ldots, r\}$. Throughout, λ remains constant while ε and δ are redefined at each iteration. However, as an invariant we maintain $\lambda < \delta < \varepsilon$.

For each $i \in \{1, \ldots, r\}$, let w_i be the ε-copy of a_i. Split $\partial_\varepsilon G^*$ at w_1, i.e., $\partial_\varepsilon G^*$ is a path with both endpoints equal to w_1. By choosing ε sufficiently small, we guarantee that $\partial_\varepsilon G^*$ is simple, i.e., $\partial_\varepsilon G^*$ is isomorphic to φ. We say that two points in the plane are *R-visible* if the open segment joining them does not intersect R. Let $\tau > 0$. For each $i \in \{1, \ldots, r\}$ such that a_i is not an interior point of R, consider the set of points $N_i \subset \partial_\varepsilon G^*$ that are at distance at most τ from w_i. Let Δ_i be the convex hull of $N_i \cup \{a_i\}$, i.e., Δ_i is a "cone" with apex at a_i; see Figure 9. (We deliberately misuse the word "cone" here because the "cones" $\Delta_1, \ldots, \Delta_r$ in this section play the same roles as the cones $\Delta_1, \ldots, \Delta_n$ in Section 2.)

While constructing R, we also maintain the *escape invariant* which is defined as follows. Assume that R so far connects $a_1, \ldots, a_i$, for some $i \in \{1, \ldots, r-1\}$. Then: (1) R intersects neither $\partial_\varepsilon G^*$ nor its unbounded face; (2) for each $j \in \{i+1, \ldots, r\}$, R intersects neither the simple polygon bounded by $\partial_\delta C_j$ nor the cone Δ_j; (3) $\Delta_h \cap \Delta_j = \emptyset$, for any distinct $h, j \in \{1, \ldots, r\}$; and (4) w_i is R-visible from a_i.

In particular, Conditions (2) and (4) of the escape invariant imply that, for each $j \in \{i, \ldots, r\}$, every point in N_j (including w_j) is R-visible from a_j. The escape invariant holds when $R = \{a_1\}$, provided that τ is sufficiently small.

Assume that we have constructed a path R that connects a_1 with a_j, for some $j \in \{1, \ldots, r-1\}$, and that the escape invariant holds. To extend R, we create a new path that connects a_j with a_{j+1} without crossing R while maintaining the escape invariant. Recall that we consider $\partial_\varepsilon G^*$ to be a path with both endpoints on w_1.

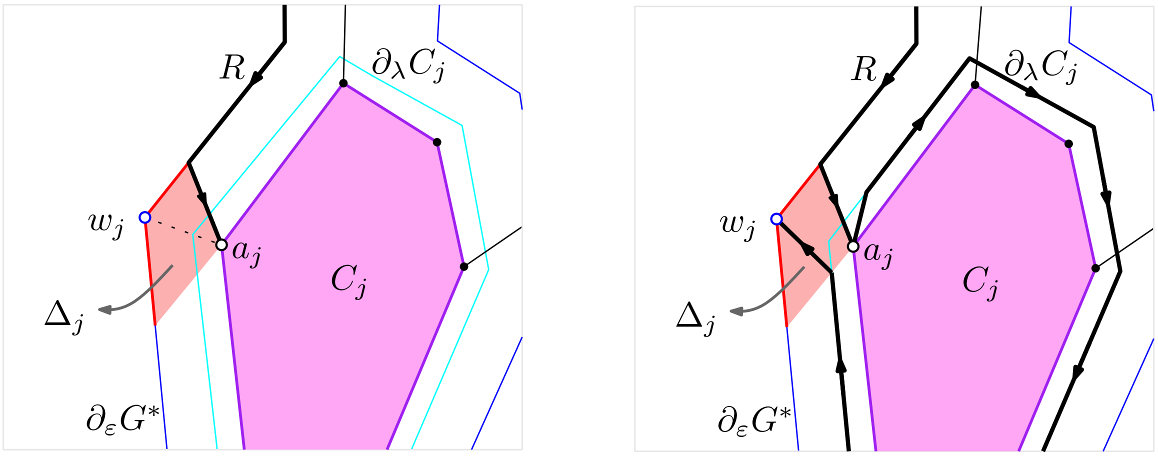

Figure 10: When extending R from a_j to a_{j+1} we have to take care to keep C_j to the right of R.

The first part of the path connecting a_j with a_{j+1} consists of a path connecting a_j with w_j. If $j = 1$, or if $j > 1$ and R together with the edge $a_j w_j$ leaves C_j to its right, then connect a_j with w_j via a straight-line segment; since w_j is R-visible from a_j, this segment does not cross R. Otherwise, connect a_j with $\partial_\lambda C_j$ via a straight-line segment and traverse $\partial_\lambda C_j$ clockwise before moving to w_j on $\partial_\varepsilon G^*$. By the escape invariant, no crossing occur in this drawing. In this way, we guarantee that C_j lies to the right of the constructed path; see Figure 10 for an illustration. Because $\lambda < \delta < \varepsilon$ and since $a_i \in V(R)$, the escape invariant is preserved.

The path from a_j to a_{j+1} continues with a path from w_j to w_{j+1}, which follows the unique path in $\partial_\varepsilon G^*$ from w_j to w_{j+1}. However, whenever we reach an endpoint of N_i for some $i \in \{j+2,\dots,r\}$, we take a *detour* to the other endpoint of N_i while avoiding its interior so that the points in the interior of N_i remain R-visible from a_i; see Figure 11. Formally, we walk from the reached endpoint of N_i to $\partial_\delta C_i \setminus \Delta_i$ along the boundary of Δ_i. Then, we traverse the path $\partial_\delta C_i \setminus \Delta_i$ before moving to the other endpoint of N_i from the endpoint of $\partial_\delta C_i \setminus \Delta_i$. Note that R does not intersect the interior of the simple polygon bounded by $\partial_\delta C_i$ nor the interior of Δ_i. Moreover, R remains inside the simple polygon bounded by $\partial_\varepsilon G^*$.

Once we go around C_i, we are back on $\partial_\varepsilon G^*$ on the other endpoint of N_i. In this way, we continue going towards w_{j+1} along $\partial_\varepsilon G^*$ until reaching an endpoint of N_{j+1}. Once we reach an endpoint of N_{j+1}, we move directly from this endpoint to a_{j+1}.

Because $\partial_\varepsilon G^*$ is isomorphic to φ, the constructed path between a_j and a_{j+1} has length at most $|\varphi(a_j, a_{j+1})|$ plus the length of the boundaries of the components for which the path detoured. Because each component we walked around has its attachment corner

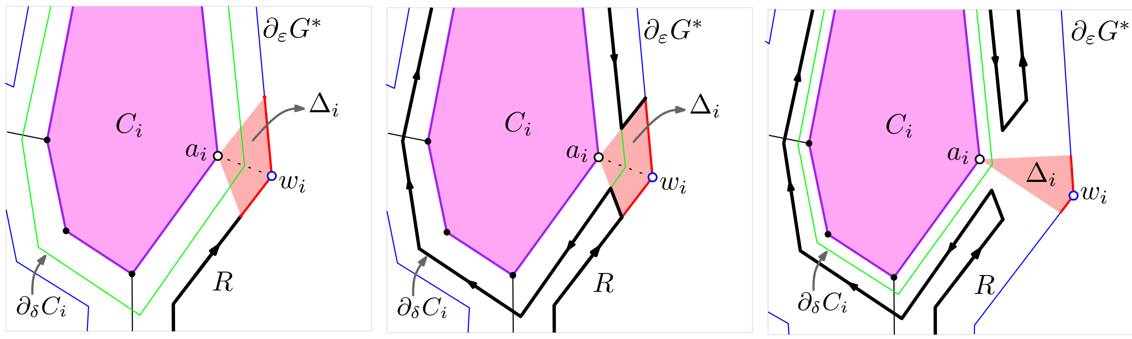

Figure 11: The "detour" taken to avoid crossing the cone Δ_i (left, middle); and the narrowing of the cone Δ_i as well as the redefinition of the ε- and δ-fattenings of G^* and C_i, respectively.

on the path $\varphi(a_j, a_{j+1})$, and thus in $A(a_j, a_{j+1})$, the length of the constructed path between a_j and a_{j+1} is

$$O\left(|\varphi(a_j, a_{j+1})| + \sum_{a_i \in A(a_j, a_{j+1})} |\partial C_i|\right) = O(\sigma_G(a_j, a_{j+1})) \ .$$

After reaching a_{j+1}, we increase ε by a factor of two. Similarly, we decrease the value of δ by a factor of two. That is, after reaching a_{j+1}, $\varepsilon = \mu 2^{j+1}$ while $\delta = \mu/2^{j+1}$ and hence, we guarantee that $\lambda < \delta < \varepsilon$. Also, $\partial_\varepsilon G^*$ is still simple, provided that μ is initially chosen to be sufficiently small. Finally, we reduce τ by a factor of two and update N_i and Δ_i accordingly, for each $i \in \{1, \ldots, n\}$; see Figure 11 (c).

Recall that for each $a_i \notin R$, R intersected neither the interior of Δ_i nor the interior of the polygon bounded by $\partial_\delta C_i$. Moreover, R remained within $\partial_\varepsilon G^*$. Therefore, after increasing (*resp.* reducing) ε (*resp.* δ), we preserve the escape invariant for the next iteration of the algorithm. We iterate until all attachment corners of G are visited by R.

Lemma 3.2. *Given an arbitrary order $a_1, \ldots, a_r$ of the attachment corners of G, there is a path R connecting all attachment corners of G in the given order such that $R \cup G$ is planar, every component C_i of G lies to the right of R when oriented from a_1 to a_r, and the subpath of R between a_j and a_{j+1} has $O(\sigma_G(a_j, a_{j+1}))$ vertices, for each $j \in \{1, \ldots, r-1\}$.*

Proof. By construction, the attachment corners are visited by R in the given order; also, the subpath, γ_j, of R between a_j and a_{j+1} has $O(\sigma_G(a_j, a_{j+1}))$ vertices. For each component C_i, a_i is the only vertex of C_i visited by R. Moreover, the construction guarantees that C_i lies to the right of R when oriented from a_1 to a_r.

To prove that R is planar, recall that in each round we extend R by constructing a path γ_j that connects a_j with a_{j+1}. We claim that at this point, no edge of γ_j crosses the

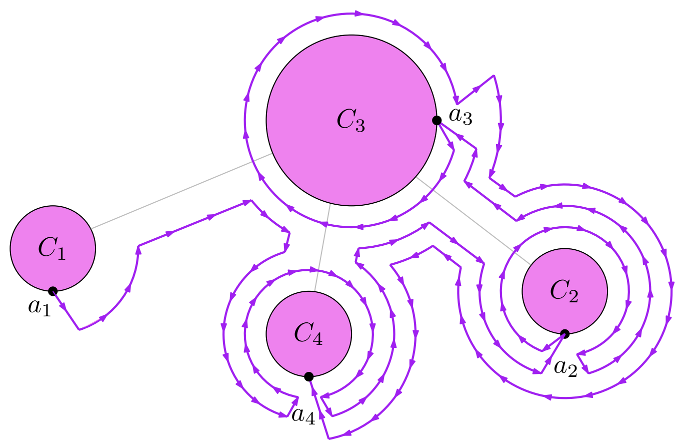

Figure 12: An example of the algorithm for generating a spanning path that connects $a_1, \ldots, a_4$.

portion of R constructed so far. Indeed, because the value of ε (*resp.* δ) increases (*resp.* decreases) in each round, the edges of γ_j that lie on the boundaries of some $\partial_\delta C_i$ or on $\partial_\varepsilon G^*$ cannot cross R by the escape invariant. Moreover, this invariant states that for each $a_i \notin R$, R does not intersect Δ_i. Because each cone Δ_i is narrowed in each round, the edges of γ_j that lie on the boundary of this cone cannot cross R. Finally, because $\lambda < \delta$, the edges of γ_j that lie on $\partial_\lambda C_j$ do not cross R. Therefore, we conclude that by concatenating γ_j and R, we obtain a planar path. □

Figure 12 illustrates the algorithm of Lemma 3.2 on a small example. In this example, the path from a_1 to a_2 passes by a_4, so R detours around C_4 in order to preserve the escape invariant at a_4. After R attaches to a_2 and a_3, it winds around components C_2 and C_3, respectively, in order to ensure that these components attach to the right of R.

3.4 Compatible drawings of planar graphs

Let $\mathcal{G}$ be a planar graph with n vertices and r connected components. Let $G_1, \ldots, G_k$ be k planar isomorphic drawings of $\mathcal{G}$. For now, we will assume that, in these drawings, every component of $\mathcal{G}$ has at least one vertex incident to the outer face. We show how to construct a compatible augmentation of $\mathcal{G}$ of size $O(nr^{1-1/k})$.

Let $\mathcal{C}_1, \ldots, \mathcal{C}_r$ be the connected components of $\mathcal{G}$. Because $G_1, \ldots, G_k$ are isomorphic, we can select one attachment corner from each component in the drawing G_1, and this attachment corner also appears in each of $G_2, \ldots, G_k$. Thus, for each $j \in \{1, \ldots, r\}$, we choose an attachment corner a_j of ∂C_j such that a_j is incident to the outer face of C_j.

For each $i \in \{1, \ldots, k\}$, let G_i^* be a connected augmentation of G_i, as defined in Section 3.2. For each $i \in \{1, \ldots, k\}$ and $j \in \{1, \ldots, r\}$, let $\text{rank}_i(j) = \sigma_{G_i^*}(a_1, a_j)$. For each

$j \in \{1,\ldots,r\}$, let $x_j \in \mathbb{R}^k$ be a point corresponding to the component C_j such that $x_j = (\text{rank}_1(a_j), \text{rank}_2(a_j), \ldots, \text{rank}_k(a_j))$. Let $X = \{x_1,\ldots,x_r\} \subset \mathbb{R}^k$ denote the resulting set of points. Lemma 3.1 implies that X is contained in an integer grid of side length $4n$.

Let P be the shortest Hamiltonian path of X under the ℓ_∞ norm. As before, because X is contained in the k-dimensional integer grid of side-length $4n$ and $|X| = r$, the maximum (ℓ_∞) length of P is $O(nr^{1-1/k})$. Note that the order of the points in P induces an order of the components of $\mathcal{G}$ and hence an order of the attachment corners of each G_i.

Theorem 3.3. *For each $1 \le i \le k$, we can construct a path R_i of length $O(nr^{1-1/k})$ that connects every component of G_i such that $G_i \cup R_i$ is planar. Moreover, for each $1 \le i < j \le k$, $G_i \cup R_i$ is isomorphic to $G_j \cup R_j$.*

Proof. By relabelling, let $(a_1,\ldots,a_r)$ denote the order of the attachment corners of G_i induced by P. Letting d_j denote the ℓ_∞ distance between x_j and x_{j+1}, we denote by $\mathcal{H}$ a path that passes through (the vertices corresponding to corners) $a_1,\ldots,a_r$ in this order, and that includes an additional $O(d_j)$ vertices between a_j and a_{j+1}. Thus, the number of vertices in $\mathcal{H}$ is proportional to the length of P, which is $O(nr^{1-1/k})$.

For each G_i, we use Lemma 3.2 to draw $\mathcal{H}$ as a planar path, R_i, that connects $a_1,\ldots,a_r$ in this order. By construction,

$$\begin{aligned} d_j &= \max\{|\text{rank}_i(j+1) - \text{rank}_i(j)| : i \in \{1,\ldots,k\}\} \\ &\ge |\text{rank}_i(j+1) - \text{rank}_i(j)| \ , \end{aligned}$$

so the $O(d_j)$ vertices in $\mathcal{H}$ between a_j and a_{j+1} are enough to draw the $O(\sigma_{G_i}(a_j, a_{j+1}))$ vertices in R_i between a_j and a_{j+1}, since $\sigma_{G_i}(a_j,a_{j+1}) = \sigma_{G_i}(a_1,a_{j+1}) - \sigma_{G_i}(a_1,a_j) = |\text{rank}_i(j+1) - \text{rank}_i(j)|$.

To conclude, each R_i visits each component only at its attachment corner, the attachment corners of each G_i are connected in the same order, and R_i leaves every component to the right when oriented from a_1 to a_r. Therefore, $G_i \cup R_i$ is isomorphic to $G_j \cup R_j$ for each $1 \le i < j \le k$. □

3.5 Handling interior components

In the preceding section, we assumed that the drawings $G_1,\ldots,G_k$ of G were such that every component was incident to the outer face. To see that the assumption is not necessary, observe that we can first use the preceding algorithm to connect all the components that do appear on the outer face using a polygonal path that is contained on the outer face. The number of vertices used in this path is $O(n'r^{1-1/k})$, where n' is the number of vertices on the outer face.

Next, for each interior face, f, that has multiple components $C_1,\dots,C_t$ on its boundary, we can use (a small modification of) the preceding algorithm to connect $C_1,\dots,C_t$ and the outer boundary of f using a path that is contained in f. This path has length $O(n_f r^{1-1/k})$. We then repeat this step on each face. The result is a connected augmentation of $\mathcal{G}$ whose total size is $O(Nr^{1-1/k})$ where $N = O(n)$ is the total size of all faces.

3.6 An algorithm

We remark that Theorem 3.3 yields an efficient algorithm for constructing the augmentation $\mathcal{H}$, and even the drawings $H_1,\dots,H_k$. The main steps involved are:

1. Finding connected planar supergraphs $G_1^*,\dots,G_k^*$ of the drawings $G_1,\dots,G_k$. For each planar graph G_i, this can easily be done in $O(n\log n)$ time using, for example, a plane sweep algorithm that maintains the invariant that all components with a vertex to the left of the sweep-line are already joined by edges. Thus, this step takes $O(kn\log n)$ time.

2. Constructing the point set X and finding the path P. Constructing X takes $O(kn)$ time, while a path P of length $O(n^{2-1/k})$ can be obtained from an (approximate) minimum spanning tree of X. For constant values of k, an approximate MST can be computed in $O(n\log n)$ time using the algorithm of Calahan and Kosaraju [**?**]. For larger values of k, the actual minimum spanning tree can be computed in $O(kn^2)$ time.

3. Constructing each of the paths $R_1,\dots,R_k$. Each of these paths is easily constructed in $O(n^{2-1/k})$ time once we have determined values of ϵ, δ, τ, and λ that are sufficiently small. A more careful examination of our algorithm reveals that all that is really needed is a value of ε such that $\partial_\epsilon G_i^*$ is simple, for each $i \in \{1,\dots,k\}$. Once we have this value of ε, the values of the remaining variables can taken from the set $\{i\varepsilon/3r : i \in \{1,\dots,3r\}\}$.

   It turns out that a value $\varepsilon \le cm/n$, where m is the minimum non-zero difference between x coordinates or y coordinates in $G_1,\dots,G_k$, and c is a constant, is sufficiently small. Thus, a suitable ε can be computed in $O(kn\log n)$ time by sorting.

This yields the following algorithmic result about connected augmentations:

Theorem 3.4. *An augmentation satisfying the conditions of Theorem 3.3 can be computed in $O(kn^2)$ time for any value of k. If k is constant, then the augmentation can be computed in $O(nr^{1-1/k})$ time.*

The latter result is worst-case optimal since, in the next section we will show that there exists inputs where every augmentation has size $\Omega(nr^{1-1/k})$.

4 Lower Bounds

Our lower bounds are based on the following lemma. It says that we can find k permutations of $\{1,\ldots,r\}$ such that for half the indices $i \in \{1,\ldots,r\}$, and every $j \in \{1,\ldots,r\} \setminus \{i\}$, there is a permutation in which i and j are at distance $\Omega(r^{1-1/k})$.

Lemma 4.1. *Let $t = (1/2)^{1+1/k} \cdot (r-1)^{1-1/k}$. There exists permutations $\pi^{(1)},\ldots,\pi^{(k)}$ of $\{1,\ldots,r\}$ such that for at least half the values of $i \in \{1,\ldots,r\}$ and for every $j \in \{1,\ldots,r\} \setminus \{i\}$,*

$$\max\left\{\left|\pi_i^{(s)} - \pi_j^{(s)}\right| : s \in \{1,\ldots,k\}\right\} \geq t \ . \tag{1}$$

Proof. This proof is an application of the probabilistic method. Select each of $\pi^{(1)},\ldots,\pi^{(k)}$ independently and uniformly from among all $r!$ permutations of $\{1,\ldots,r\}$. Fix a particular index i and a particular index j. For a particular $s \in \{1,\ldots,k\}$, the probability that $|\pi_i^{(s)} - \pi_j^{(s)}| \leq t$ is at most $2t/(r-1)$ since the set $\{\hat{\jmath} \in \{1,\ldots,r\} : |\pi_i^{(s)} - \pi_{\hat{\jmath}}^{(s)}| \leq t\}$ is a random subset of at most $2t$ elements drawn without replacement from $\{1,\ldots,r\} \setminus \{i\}$.

Therefore, since $\pi^{(1)},\ldots,\pi^{(k)}$ are chosen independently,

$$\Pr\left\{\max\left\{\left|\pi_i^{(s)} - \pi_j^{(s)}\right| : s \in \{1,\ldots,k\}\right\} \leq t\right\} \leq (2t/(r-1))^k = \frac{1}{2(r-1)} \ .$$

In particular, the expected number of such $j \in \{1,\ldots,r\} \setminus \{i\}$ is at most $1/2$ so, by Markov's Inequality, the probability that there exists at least one such j is at most $1/2$. Thus, with probability at least $1/2$, the index i satisfies (1) and therefore the expected number of indices $i \in \{1,\ldots,r\}$ that satisfy (1) is $r/2$. We conclude that there must exist some permutations $\pi^{(1)},\ldots,\pi^{(k)}$ that satisfy (1) for at least half the indices $i \in \{1,\ldots,r\}$. □

Using Lemma 4.1, we can prove a lower bound that matches the upper bound obtained in our general construction.

Theorem 4.2. *For every positive integer n and every $r \in \{2,\ldots,\lfloor n/4 \rfloor\}$, there exists a graph $\mathcal{G}$ having n vertices, r connected components, and k isomorphic drawings $G_1,\ldots,G_k$ such that any compatible augmentation of $\mathcal{G}$ has size $\Omega(nr^{1-1/k})$.*

Proof. Since the lemma only claims an asymptotic result, we may assume without loss of generality that r is even and that $2r$ divides n.

The graph $\mathcal{G}$ consists of r disjoint paths, $\mathcal{C}_1,\ldots,\mathcal{C}_r$, each of length n/r. Each of the drawings $G_1,\ldots,G_k$ has the vertices of $\mathcal{G}$ on the same point and edge set. The point set consists of the vertices of r nested regular n/r-gons, $P_1,\ldots,P_r$, each centered at the origin and having nearly the same size. Refer to Figure 13 (left). More precisely, $P_1 \subset P_2 \subset \cdots \subset P_r$ and the sizes are chosen so that any segment joining two non-consecutive vertices of P_i

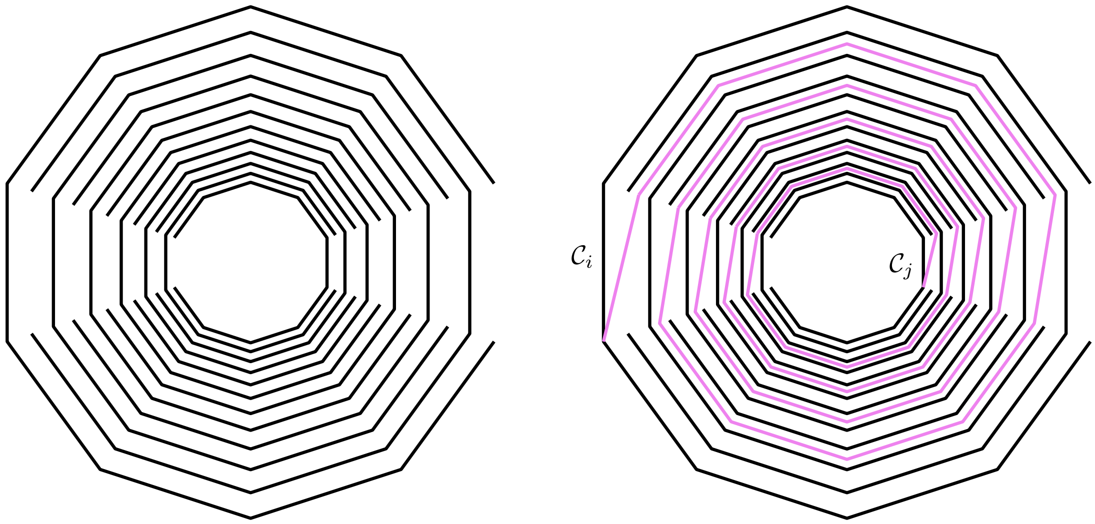

Figure 13: In the construction in Theorem 4.2, all drawings use the same set of vertices and line segments and the drawing of a path that joins $\mathcal{C}_i$ to $\mathcal{C}_j$ must travel around all paths drawn between the drawing of $\mathcal{C}_i$ and $\mathcal{C}_j$.

intersects the interior of P_{i-1}. The drawings $G_1, \ldots, G_k$ are obtained from the permutations $\pi^{(1)}, \ldots, \pi^{(k)}$ given by Lemma 4.1. In the drawing G_x, the path $\mathcal{C}_i$ is drawn on the vertices of $P_{\pi_i^{(x)}}$. If $y = \pi_i^{(x)}$ is even, the drawing uses all the edges of P_y except the left-most edge. If y is odd, the drawing uses all the edges of P_y except the right-most edge.

Now, without loss of generality, consider some edge-minimal compatible augmentation $\mathcal{H}$ of $\mathcal{G}$. For each component $\mathcal{C}_i$ of G, let T_i be any path in $\mathcal{H}$ that has one endpoint on $\mathcal{C}_i$, one endpoint on some other component $\mathcal{C}_j$, $j \neq i$, and no vertices of $\mathcal{G}$ in its interior. Now, for each of the $r/2$ indices $i \in \{1, \ldots, r\}$ that satisfy (1), the path T_i joins a vertex of $P_{\pi_i^{(s)}}$ to a vertex of $P_{\pi_j^{(s)}}$, $j \neq i$, and $|\pi_i^{(s)} - \pi_i^{(s)}| \geq t$. This path must have length $\Omega(tn/r)$ since it has to "go around" the paths between $P_{\pi_i^{(s)}}$ and $P_{\pi_j^{(s)}}$; see Figure 13 (right).

Thus far, we have shown that for at least $r/2$ values of $i \in \{1, \ldots, r\}$, the component C_i is the endpoint of a path, T_i, of length at least $\Omega(tn/r) = \Omega(nr^{-1/k})$. It is tempting to claim the result at this point, since $(r/2) \cdot \Omega(nr^{-1/k}) = \Omega(nr^{1-1/k})$. Unfortunately, there is a little more work that needs to be done, since two such paths T_i and T_j may not be disjoint, so summing their lengths double-counts the contribution of the shared portion.

To finish up we note that, since the augmentation $\mathcal{H}$ is minimal, it is a tree; $\mathcal{G}$ contains no cycles, so any cycle in $\mathcal{H}$ contains an edge not in $\mathcal{G}$ that could be removed. Now, observe that if we traverse the outer face of (any planar drawing of) $\mathcal{H}$ then we obtain a non-simple path, P, that traverses each edge of $\mathcal{H}$ exactly twice. If we consider the set of maximal subpaths of P with no vertex of $\mathcal{G}$ in their interior, we obtain a set of r edge-disjoint paths, $Q_1,\ldots,Q_r$ and, for every component $\mathcal{C}_i$ of $\mathcal{G}$, there is a vertex of $\mathcal{C}_i$ that is an endpoint of at least one such path. Therefore, from the preceding discussion, the total length of $Q_1\ldots,Q_r$ is $\Omega(nr^{1-1/k})$. But since each edge of $\mathcal{H}$ appears at most twice in these subpaths, we conclude that $\mathcal{H}$ has $\Omega(nr^{1-1/k})$ edges. Since $\mathcal{H}$ is a tree, it has $\Omega(nr^{1-1/k})$ vertices. □

Acknowledgement

This work was initiated at the *Second Workshop on Geometry and Graphs*, held at the Bellairs Research Institute, March 9-14, 2014. We are grateful to the other workshop participants for providing a stimulating research environment.

A Shortest Tour in the Uniform Norm

Under the ℓ_∞ *metric*, the distance between two points $x = (x_1, \ldots, x_k)$ and $y = (y_1, \ldots, y_k)$ is $\|x - y\|_\infty = \max\{|x_i - y_i| : i \in \{1, \ldots, k\}\}$.

Lemma A.1. *Let P be a set of $r \geq 2$ points contained in the k-dimensional cube $[0,1]^k$, for $k \geq 2$. Then there exists a spanning path of P whose length under the ℓ_∞ metric is at most $cr^{1-1/k}$, where c is a universal constant. In particular, c does not depend on k or r.*

Proof. The following proof is a rehashing of an argument used by Moran [**?**]. An argument of Few [**?**], which begins by stabbing $[0,1]^k$ with a grid of $\Theta(r^{1-1/k})$ lines, could also be used to establish the same asymptotic result.

First, we note that it is sufficient to upper-bound the ℓ_∞-length of the minimum spanning-tree of P, since this can be transformed into a path of at most twice its length [**?**].

For any point $p \in \mathbb{R}^k$, the *uniform ball* of radius d centered at p, defined as

$$B_\infty(p,d) = \{q \in \mathbb{R}^k : \|p - q\|_\infty \leq d\}$$

is a cube of side-length $2d$ and has volume $(2d)^k$. If $d < 1$ and $p \in [0,1]^k$, then $B_\infty(p,d) \cap [0,1]^k$ contains a cube of side-length d, so $B_\infty(p,d) \cap [0,1]^k$ has volume at least d^k.

The preceding implies that the set P contains two points p and q, such that $\|p - q\|_\infty \leq 2r^{-1/k}$; otherwise, one could pack r disjoint cubes, each of volume greater than $1/r$ into $[0,1]^k$.

Now, we can construct a spanning tree of P by repeatedly taking the pair of points $p, q \in P$ that minimize $\|p - q\|_\infty$, adding the edge pq to our spanning tree and then removing q from P. Since, at the ith step of this algorithm, the set P contains $r - i + 1$ points, the total ℓ_∞-length of all the edges added to this tree is

$$\begin{aligned}
\sum_{i=2}^{r} 2i^{-1/k} &\leq 2^{1-1/k} + \int_2^r 2x^{-1/k}\,\mathrm{d}x \\
&= 2^{1-1/k} + (2/(1-1/k))\left(r^{1-1/k} - 2^{1-1/k}\right) \\
&\leq 4r^{1-1/k} + O(1) \\
&\leq Cr^{1-1/k} \quad ,
\end{aligned}$$

for a sufficiently large constant C and any $r, k \geq 2$. Thus, the result holds for $c = 2C$. □

By uniformly scaling the point set P by a factor of n, we obtain the following corollary of Lemma A.1, which is used in our algorithm:

Corollary A.2. *Let P be a set of $r \geq 2$ points contained in the cube $[0, n]^k$, for $k \geq 2$. Then there exists a spanning path of P whose length under ℓ_∞ metric is at most $cnr^{1-1/k}$, where c is a universal constant. In particular, c does not depend on k, r, or n.*