

Computing β -Stretch Paths in Drawings of Graphs

Esther M. Arkin

Stony Brook University, NY, USA

Faryad Darabi Sahneh

University of Arizona, Tucson, AZ, USA

Alon Efrat

University of Arizona, Tucson, AZ, USA

Fabian Frank

University of Arizona, Tucson, AZ, USA

Radoslav Fulek

University of Arizona, Tucson, AZ, USA

Stephen Kobourov

University of Arizona, Tucson, AZ, USA

Joseph S. B. Mitchell

Stony Brook University, NY, USA

Abstract

Let f be a drawing in the Euclidean plane of a graph G , which is understood to be a 1-dimensional simplicial complex. We assume that every edge of G is drawn by f as a curve of constant algebraic complexity, and the ratio of the length of the longest simple path to the length of the shortest edge is $\text{poly}(n)$. In the drawing f , a path P of G , or its image in the drawing $\pi = f(P)$, is β -stretch if π is a simple (non-self-intersecting) curve, and for every pair of distinct points $p \in P$ and $q \in P$, the length of the sub-curve of π connecting $f(p)$ with $f(q)$ is at most $\beta \|f(p) - f(q)\|$, where $\|\cdot\|$ denotes the Euclidean distance. We introduce and study the β -stretch Path Problem (β SP for short), in which we are given a pair of vertices s and t of G , and we are to decide whether in the given drawing of G there exists a β -stretch path P connecting s and t . The β SP also asks that we output P if it exists.

The β SP quantifies a notion of “near straightness” for paths in a graph G , motivated by gerrymandering regions in a map, where edges of G represent natural geographical/political boundaries that may be chosen to bound election districts. The notion of a β -stretch path naturally extends to cycles, and the extension gives a measure of how gerrymandered a district is. Furthermore, we show that the extension is closely related to several studied measures of local fatness of geometric shapes.

We prove that β SP is strongly NP-complete. We complement this result by giving a quasi-polynomial time algorithm, that for a given $\varepsilon > 0$, $\beta \in O(\text{poly}(\log |V(G)|))$, and $s, t \in V(G)$, outputs a β -stretch path between s and t , if a $(1 - \varepsilon)\beta$ -stretch path between s and t exists in the drawing.

2012 ACM Subject Classification Mathematics of computing \rightarrow Paths and connectivity problems; Theory of computation \rightarrow Computational geometry

Keywords and phrases stretch factor, dilation, geometric spanners

Digital Object Identifier 10.4230/LIPIcs.SWAT.2020.7

Funding This work is supported in part by NSF grants CCF-1526406, CCF-1740858, CCF-1712119, and DMS-1839274.

1 Introduction

We study an optimal path problem in planar drawings of graphs, in which we represent edges as curves of constant algebraic complexity. We seek a path in a graph G from a given vertex s to another given vertex t that is, in a precise sense, as close as possible to the straight-line



© Esther M. Arkin, Faryad Darabi Sahneh, Alon Efrat, Fabian Frank, Radoslav Fulek, Stephen Kobourov, and Joseph S. B. Mitchell;

licensed under Creative Commons License CC-BY

17th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2020).

Editor: Susanne Albers; Article No. 7; pp. 7:1–7:20



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

segment from s to t . We formalize this notion by saying that an $s - t$ path is a β -stretch path if the distance between any two points along the path (not only the endpoints) is at most β times the Euclidean distance between them.

The notion of “ β -stretch” in this definition is similar to the notion of stretch in a multiplicative β graph spanner [17], where we want to remove edges from the graph while ensuring that the shortest path distance in the spanner is at most β times the length of a shortest path in the original graph. Thorough reviews of existing results for geometric spanners are available in [4, 9, 16]. In our problem we are not sparsifying the graph; instead, we try to find the most “natural” path connecting two given vertices s and t in a given embedded graph. If we interpret the embedded graph as the road network of a country, such paths can be used as an initial step to partition the country into regions with natural shapes. One of our motivations, in fact, is the problem of computing natural regions that, in a precise sense, avoid gerrymandering. A few definitions have been proposed in the literature to characterize what a “natural” path could entail. For example, a path in a drawing of a graph is defined to be *self-approaching* [1, 12] if for any two points p and q on the path, when moving from p to q along the path, the Euclidean distance to q is decreasing. Icking et al. [12] proved that a self-approaching path is 5.3332-stretch.

The problem of computing β -stretch paths bears similarities to the graph dilation problem, where for every pair of vertices s and t in a geometric graph, we compare the shortest-path distance between s and t to their actual Euclidean distance in the plane, and return the largest ratio of these two values over all pairs (s, t) . In the special case of cycles this problem is known as computing the maximum detour of a polygonal chain [8]. Klein and Kutz show that computing a minimum-dilation graph that connects a given n -point set in the plane with at most m edges is NP-hard [14]. In one direction, if we are given an embedded geometric graph with a dilation ratio that is at most as large as our target stretch factor, a weaker variant of a β -stretch path exists between every pair of vertices $s - t$, in which we consider only pairs of vertices along the path rather than points. However, since the dilation is a global property an $s - t$ path that is β -stretch in the given graph might still exist even if the dilation is more than β . We elaborate on other connections to our problem in Section 1.3.

We naturally extend the notion of β -stretch paths to β -stretch cycles. Interestingly, we show that a β -stretch cycle bounds a locally “fat” shape in the sense as defined by De Berg [7], with the parameter of fatness depending on β . The converse is easily seen not to be true. Our notion of β -stretch cycles may have applications to computing geographic partitions into regions whose shapes are well shaped in a sense that cannot be captured with fatness criteria.

The rest of the paper is organized as the following. We formally define the β -stretch path problem in Section 1.1, followed by key main results and an overview of related results in the literature in Section 1.2 and 1.3, respectively. In Section 2, we prove a relation between β -stretch cycles and locally γ -fat shapes. Section 3 proves that β -stretch path problem is strongly NP-complete. Section 4 develops a quasi-polynomial approximation scheme algorithms for β -stretch path problem and its extension to computing β -stretch cycles. We conclude with open problems and future directions in Section 5. Omitted proofs are in the Appendix (Section 6).

1.1 Problem Statement

Let $G = (V, E)$ be a finite simple graph, with vertex set V and edge set $E \subseteq \binom{V}{2}$. A *drawing* of a graph is a representation of G in the Euclidean plane \mathbb{R}^2 , in which vertices are distinct points and edges are Jordan arcs represented as curves of *constant algebraic complexity*, i.e., described by a constant number of polynomial equations (inequalities), whose maximum

degree is bounded by a fixed constant.

Formally, a drawing of a graph is a continuous map $f : G \rightarrow \mathbb{R}^2$, where we treat G as a 1-dimensional simplicial complex. The representation of a vertex $v \in V$, an edge $e \in E$, and a path $P \subseteq G$ in the drawing f is $f(v)$, $f(e)$, and $f(P)$, respectively. Here, we consider a *generalized path* that can end in a midpoint of an edge.

We will distinguish paths in a graph from paths in a drawing of a graph. The reason is that we will consider “paths” in a drawing that end in relative interiors of edges. Treating G as a 1-dimensional simplicial complex, a *path* in a drawing f of G is $f(P)$, where P is a generalized path in G . We will be denoting paths in a drawing by lower case Greek letters.

Let $\|\cdot\|$ be the Euclidean norm. Let $P \subseteq G$ denote a path between p and $q \in G$. If both p and q are vertices of G then P corresponds to a usual path in G . Let f be a drawing of G . Then $\pi = f(P)$ is the path *between p and q* in f . Let $\pi(p', q')$ denote the sub-path of π between $p', q' \in G$, that is, $\pi(p', q') = f(P(p', q'))$, where $P(p', q') \subseteq P$ is the path between p' and q' . If we want to specify a path π together with its endpoints s and t we denote it by $\pi(s, t) = \pi$. The path π *passes through* all of the vertices and edges of G intersecting P . The *length* of the path π , denoted by $\|\pi\|$, is the usual Euclidean length, which can be computed as $\int_P \|f'(x)\| dx$. The *distance* between $s \in P$ and $t \in P$ along π , denoted by $d_\pi(s, t)$, is the length of the sub-curve of π between $f(s)$ and $f(t)$.

β -stretch path. Let π be a path in f free of self-intersections. For $\beta \geq 1$, path π is a *β -stretch path* if for every $p, q \in P$ we have

$$\frac{d_\pi(p, q)}{\|f(p) - f(q)\|} \leq \beta. \quad (1)$$

β -stretch cycle. Let C be a simple cycle in G so that $\gamma = f(C)$ is free of self-intersections. The cycle γ in f is a *β -stretch cycle* if for every pair of points p and q on C we have

$$\frac{d_\gamma(p, q)}{\|f(p) - f(q)\|} = \frac{\min\{d_\pi(p, q), d_{\pi'}(p, q)\}}{\|f(p) - f(q)\|} \leq \beta, \quad (2)$$

where $\pi = \pi(p, q)$ and $\pi' = \pi'(p, q)$ are the two paths between q and p whose union is γ .

The left hand side of (1) and (2) is the *stretch factor of p and q along π and γ* , respectively. The maximum of the stretch factor of p and q over distinct $p, q \in P$ and $p, q \in C$ is the *stretch factor* of π and γ , respectively. Note that a β -stretch path (cycle) is a β' -stretch path (cycle), for every $\beta' \geq \beta$. If a path π or a cycle γ is self-intersecting, its stretch factor is undefined.

► **Problem 1.** *β -STRETCH PATH PROBLEM (β SP).* We are given a drawing f of a graph G , $\beta \geq 1$, $s \in V(G)$ and $t \in V(G)$. Decide whether there exists a β -stretch path in f between s and t . The instance of the problem is denoted by (G, f, β, s, t) .

A self-intersection-free cycle γ in a drawing f of G *separates* $s \in G \setminus C$ from $t \in G \setminus C$ if $f(s)$ and $f(t)$ are contained in different connected components of the complement of γ in \mathbb{R}^2 .

► **Problem 2.** *β -STRETCH CYCLE PROBLEM (β CP).* We are given a drawing f of a graph G , $\beta \geq 1$, $s \in V(G)$ and $t \in V(G)$. Decide whether there exists a β -stretch cycle in f separating s from t . The instance of the problem is denoted by (G, f, β, s, t) .

1.2 Main Results

Our main results proved in Sections 3, 4.2 and 4.3, respectively, are the following.

► **Theorem 1.** *β SP is strongly NP-complete.*

► **Theorem 2.** *Let (G, f, β, s, t) be an instance for β SP with $\text{poly}(\log n) \geq \beta \geq 1$. Suppose that the shortest edge length in f is 1, and that there exists $c > 0$ such that the longest simple path in f has length at most n^c . Under the above assumptions there exists a QPTAS for β SP. In other words, there exists a quasi-polynomial-time algorithm that for a fixed $\text{poly}(\log n) \geq \beta \geq 1$ and $\varepsilon > 0$ returns a β -stretch path between s and t if a $\beta(1 - \varepsilon)$ -stretch path between s and t exists in f .*

► **Theorem 3.** *Let (G, f, β, s, t) be an instance for β SC with $\text{poly}(\log n) \geq \beta \geq 1$. Suppose that the shortest edge length in f is 1, and that there exists $c > 0$ such that the longest path in f has the length at most n^c . Under the above assumptions there exists a QPTAS for β SC. In other words, there exists a quasi-polynomial-time algorithm that for a fixed $\text{poly}(\log n) \geq \beta \geq 1$ and $\varepsilon > 0$ returns a β -stretch cycle separating s from t if a $\beta(1 - \varepsilon)$ -stretch cycle separating s from t exists in f .*

1.3 Related Work

Dilation or stretch factor [16] is perhaps the most common measure for the quality of a geometric graph. There is a subtle difference between the stretch factor of a path versus the stretch factor of a graph. For a path, the stretch factor only pertains to its endpoints, while for a graph the stretch factor pertains to every pair of the graph vertices. Our definition of β -stretch path falls in the middle as it pertains to all pairs of points belonging to the path.

It is worth mentioning that a line of existing results in the literature is not about designing a geometric graph with desired stretch factor, but about the fast computation of the stretch factor, given the graph. Narasimhan and Smid [15] considered the problem of computing the stretch factor of a Euclidean graph, defined as the maximum ratio of graph distance and Euclidean distance between any two vertices of the graph. Using Callahan and Kosaraju's well-separated pair decomposition, they showed that there exists a EPTAS for computing the stretch factor running in $O(|V|^{3/2})$ time, which is much faster than computing all-pairs-shortest-path distances. For general weighted graphs, Cohen proposed fast algorithms to compute paths with a desired stretch factor [6]. The stretch factor, in this case, is the ratio of the path length to the graph distance. Farshi *et al.* studied the problem of adding an edge to a Euclidean graph that lowers its stretch factor as much as possible [11].

Chen *et al.* [5] recently proposed a new straightness measure for a path. A polygonal chain (p_1, p_2, \dots, p_n) is a c -chain if for all $1 \leq i < j < k \leq n$, we have $\|p_i - p_j\| + \|p_j - p_k\| \leq c\|p_i - p_k\|$. There is a connection between the notion of c -chain and our proposed notion of β -stretch paths. On the one hand, if a chain is β -stretch, it is trivial to show that it is also a β -chain according to the definition in [5]. On the other hand, a c -chain bounds the possible stretch of the chain according to [5, Theorem 1–3]. Even though the analysis is only for the endpoints of the path, the results readily follow for any pair of points on the chain. Hence, it indeed implies the chain has β -stretch (with the difference of only checking pairs of vertices, not the points on the connecting segments).

A closely related notion to our β -stretch path is the notion of quasiconvexity as defined by Azzam and Schul [3]. A connected subset Γ of the Euclidean space is said to be *quasiconvex* if any two points x and y in Γ can be connected via a path in Γ whose length is bounded by a constant times the Euclidean distance between x and y [3]. According to this definition, a

β -stretch path is quasiconvex with constant β . The problem studied by Azzam and Schul is in some sense opposite to ours. Given a connected set Γ and a target set of points K , they compute a superset $\tilde{\Gamma} \supset \Gamma$ that connects the K points, has Hausdorff length comparable to that of Γ , and is quasiconvex. We, instead, look for a path that is a subset of the given connected set (graph) and that is quasiconvex with a constant stretch factor β . While a short quasiconvex set always exists [3, Theorem 1], we show that determining whether a β -path exists is strongly NP-complete.

One measure of “compactness” designed to quantify gerrymandering in political districting is the Polsby-Popper score, based on the ratio of the area of a district to the square of the district’s perimeter [18]. See [19] for a discussion of shape measures used in the study of gerrymandering.

2 β -Stretch Curves and Locally γ -Fat Shapes

In order to model inputs that represent realistic objects, computational geometers introduced the notion of *fat shapes*. The aim of this section is to argue that our notion of β -stretch cycles captures a local variant of fatness.

Roughly speaking, a planar shape, understood as a closed topological disk T , is locally γ -fat if every disk that is centered in T and is not containing the whole T has at least a γ -fraction of its area in T . Let $D \subset \mathbb{R}^2$ denote a disk. Let $D \cap S$, for $S \subseteq \mathbb{R}^2$, denote the path connected component of $D \cap S$ containing the center of D .

Locally γ -fat shape [2, 7]. For $0 \leq \gamma \leq \frac{1}{2}$, a closed topological disk $T \subseteq \mathbb{R}^2$ is locally γ -fat if for every disk D centered in T that does not contain D in its interior, we have $\text{area}(T \cap D) \geq \gamma \cdot \text{area}(D)$.

We remark that there exists a variant of local γ -fatness that considers $\text{area}(T \cap D)$ rather than $\text{area}(T \cap D)$ [20, 21]. The following applies also to this weaker notion of local γ -fatness.

The notion of β -stretch cycles extends to any measurable Jordan curve, in particular, boundaries of “nice” topological disks. In the following theorem, we show that by controlling the stretch factor of the boundary of a topological disk, we also control its local fatness. In particular, lowering the stretch factor increases the fatness. The corresponding lower bound on the local fatness is the inverse of a linear function of the stretch factor with the leading constant factor 2π . We also show that the stretch factor of the boundary cannot be bounded by a function of its local fatness.

► **Theorem 4.** *Every closed topological disk $T \subset \mathbb{R}^2$, whose boundary ∂T is measurable and β -stretch, is locally $\frac{1}{2\pi\beta}$ -fat. For every $\beta > 1$, there exists a locally $\frac{1}{32\pi}$ -fat topological disk whose boundary is not a β -stretch cycle.*

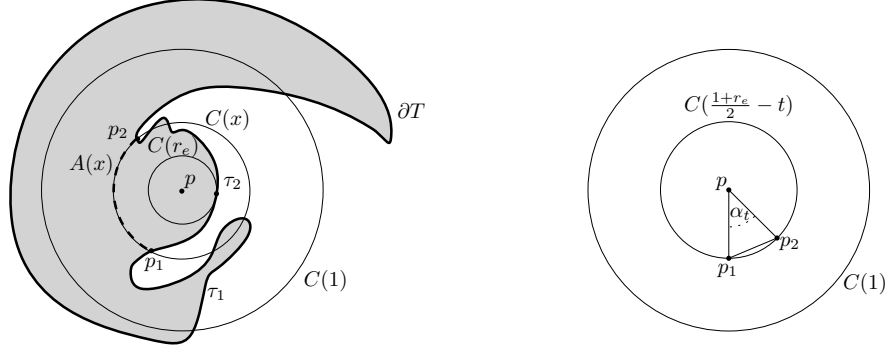
Proof. Let D denote a disk, centered at a point $p \in T$, that does not contain T in its interior. We need to show that $\frac{1}{2\pi\beta} \text{area}(D) \leq \text{area}(T \cap D)$.

Let $D(r)$ and $C(r)$, for $r \geq 0$, denote the disk and circle, respectively, with radius r centered at p . By rescaling, we assume that $D = D(1)$ is a unit disk. Let $r_e = \min\{r \geq 0, (C(r) \cap \partial T) \neq \emptyset\}$. Hence, r_e is the radius of the largest disk $D(r_e)$, whose interior does not intersect ∂T . Since D does not contain T in its interior, we have $r_e \leq 1$.

We will presently show that $\left(r_e^2 + \frac{(1-r_e)^2}{2\pi\beta}\right) \text{area}(D) = \left(r_e^2 + \frac{(1-r_e)^2}{2\pi\beta}\right) \pi \leq \text{area}(T \cap D)$. Then optimizing over the value of r_e , such that $0 \leq r_e \leq 1$, in the previous two inequalities gives the desired lower bound $\frac{1}{2\pi\beta} \text{area}(D)$ on $\text{area}(T \cap D)$. The lower bound is minimized for $r_e = 0$. It remains to show that $\left(r_e^2 + \frac{(1-r_e)^2}{2\pi\beta}\right) \pi \leq \text{area}(T \cap D)$. The first term is due to the fact that $D(r_e) \subseteq T$ since $p \in T$.

7:6 Computing β -Stretch Paths in Drawings of Graphs

To get the second term we consider slices $S(r) = T \cap C(r)$, for $r_e \leq r \leq 1$. First, we treat $r \in [r_e, \frac{1+r_e}{2}]$. We claim that $S(\frac{1+r_e}{2} - t)$, for $0 \leq t \leq \frac{1-r_e}{2}$, contains a circular arc of angular length greater than or equal to $\frac{1}{\beta} \cdot 2 \frac{1-r_e-2t}{1+r_e-2t}$. The claim is proved with the help of the following lemma; see Figure 1 for an illustration.



■ **Figure 1** An illustration of Lemma 5 (left) and inequality (3) (right).

► **Lemma 5.** *The slice $S(x)$, $r_e < x \leq 1$, contains a circular arc $A(x)$, whose relative interior is contained in the interior of $T \cap D$, and whose endpoints $p_1 \in \partial T$ and $p_2 \in \partial T$ split ∂T into two parts τ_1 and τ_2 sharing p_1 and p_2 , such that $\tau_2 \cap C(r_e) \neq \emptyset$ and $\tau_1 \cap C(1) \neq \emptyset$.*

Proof. Refer to Figure 1 (left). First, we perturb ∂T a little bit to eliminate touchings between $C(x)$ and ∂T without increasing the total length of $C(x)$ contained in the interior of T . Let p'_1 and p'_2 denote a point in $\partial T \cap C(r_e)$ and $\partial T \cap C(1)$, respectively. Let τ'_1 and τ'_2 denote the two parts of ∂T connecting p'_1 and p'_2 . We assume that τ'_2 is shortest possible. In particular, τ'_2 is contained in $\partial(T \cap D)$. Note that both τ'_1 and τ'_2 intersect $C(x)$ in an odd number of path connected components.

Let A_1, \dots, A_k denote the path connected components of $T \cap C(x)$. Note that none of A_i 's is a point since we eliminated touchings between ∂T and $C(x)$. It must be that there exists A_j , $1 \leq j \leq k$, such that one endpoint of A_j belongs to τ'_1 and the other to τ'_2 . Indeed, otherwise the number of path connected components in $\tau'_1 \cap C(x)$ and $\tau'_2 \cap C(x)$ would be even.

By the choice of τ'_2 , putting $A(x) = A_j$ concludes the proof. ◀

We show that $A(\frac{1+r_e}{2} - t)$ from Lemma 5 is an arc of the desired angular length, which is at least $\frac{1}{\beta} \cdot 2 \frac{1-r_e-2t}{1+r_e-2t}$. Let τ_1 and τ_2 , and p_1 and p_2 be as in Lemma 5 for $x = \frac{1+r_e}{2} - t$. Note that due to the choice of t and the fact that $C(r_e) \cap \tau_2 \neq \emptyset$, we have $d_{\tau_2}(p_1, p_2) \geq 2(\frac{1-r_e}{2} - t)$. The same inequality holds for $d_{\tau_1}(p_1, p_2)$, since $\tau_1 \cap C(1) \neq \emptyset$. Let α_t denote the smaller angle defined by the rays emanating from p through p_1 and p_2 . Since ∂T is β -stretch, we have, see Figure 1 (right),

$$\beta \geq \frac{2(\frac{1-r_e}{2} - t)}{\|p_1 - p_2\|} = \frac{2(\frac{1-r_e}{2} - t)}{2 \sin \frac{\alpha_t}{2} (\frac{1+r_e}{2} - t)}. \quad (3)$$

The desired lower bound $\frac{1}{2\beta} \cdot \frac{1-r_e-2t}{1+r_e-2t}$ on the angular length of $A(\frac{1+r_e}{2} - t)$ follows since this is lower bounded by $2 \sin \frac{\alpha_t}{2}$.

Similarly we prove that $S(\frac{1+r_e}{2} + t)$, for $0 \leq t \leq \frac{1-r_e}{2}$, contains a circular arc of angular length at least $\frac{1}{\beta} \cdot 2 \frac{1-r_e-2t}{1+r_e+2t}$.

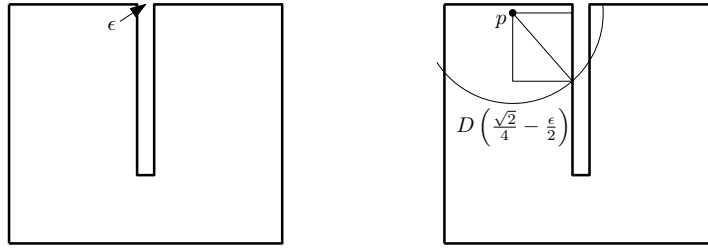
Finally, by summing up infinitesimal thickenings of the slices of width dt we get

$$\begin{aligned} \text{area}(D \cap T) &\geq \frac{1}{2\beta} \int_0^{\frac{1-r_e}{2}} 2 \frac{1-r_e-2t}{1+r_e-2t} \left(\left(\frac{1+r_e}{2} - t \right)^2 - \left(\frac{1+r_e}{2} - t - dt \right)^2 \right) + \\ &+ \frac{1}{2\beta} \int_0^{\frac{1-r_e}{2}} 2 \frac{1-r_e-2t}{1+r_e+2t} \left(\left(\frac{1+r_e}{2} + t \right)^2 - \left(\frac{1+r_e}{2} + t - dt \right)^2 \right), \end{aligned}$$

which simplifies to

$$\text{area}(D \cap T) \geq \frac{2}{\beta} \int_0^{\frac{1-r_e}{2}} (1-r_e-2t) dt.$$

It follows that $\frac{(1-r_e)^2}{2\beta} \leq \text{area}(T \cap D)$, concluding the proof of the first part of the theorem.



■ **Figure 2** A family of topological disks T witnessing that a locally $\frac{1}{32\pi}$ -fat shape can have boundary with an arbitrarily large stretch factor, which is achieved by choosing ϵ arbitrarily small.

Refer to Figure 2. For the second part of the theorem, consider a topological disk T , that is a unit square with an $\epsilon > 0$ wide slit from the middle of an edge to the center as in Figure 2. Clearly, if we choose $\epsilon < \frac{1}{\beta}$ then ∂T is not a β -stretch cycle. However, T stays locally $\frac{1}{32\pi}$ -fat for $\epsilon > 0$. Indeed, it is not hard to see that for $r < \frac{\sqrt{2}}{4} - \frac{\epsilon}{2}$, a disk $D(r)$ centered at a point p in T of radius r has $\text{area}(T \cap D(r)) \geq \left(\frac{r}{\sqrt{2}} \right)^2 > \frac{r^2}{32} = \frac{\text{area}(D(r))}{32\pi}$. For $r \geq \frac{\sqrt{2}}{4} - \frac{\epsilon}{2}$, we have $\text{area}(T \cap D(r)) \geq \frac{1}{16}$, but it is enough to consider $r \leq \sqrt{2}$, since otherwise the whole T is contained in $D(r)$. Hence, $\text{area}(T \cap D(r)) \geq \frac{1}{16} = \frac{2\pi}{32\pi} \geq \frac{\text{area}(D(r))}{32\pi}$. ◀

3 NP-completeness of βSP

The aim of this section is to prove Theorem 1. Let G, f, s and t be as in the statement of the problem βSP . First, we show that we can certify that a given path π in f is a β -stretch path in polynomial time, which follows by the next lemma.

► **Lemma 6.** *Let π be a non-self-intersecting path in f between s and t . There exists a quadratic time algorithm to check if π is a β -stretch path.*

Proof. Note that it is enough to compute the maximum of

$$\max_{s \in e, t \in f} \frac{d_\pi(s, t)}{\|f(s) - f(t)\|}, \tag{4}$$

over pairs of edges e and f on the path P in G such that $\pi = f(P)$. Due to a constant algebraic complexity of edges in f , (4) can be seen as a rational function of two variables whose

maximum can be computed in constant time by the standard calculus and approximated by solving a system of polynomial equations, and therefore the quadratic time complexity follows. \blacktriangleleft

Thus, the problem is in NP, and it remains to argue the NP-hardness. We proceed by a reduction from the graph vertex cover problem, which is one of the first known NP-complete problems from Karp's seminal paper [13], and which we state next. A *vertex cover* in a graph $G = (V, E)$ is a subset V' of its vertex set V such that every edge in E has at least one vertex in V' .

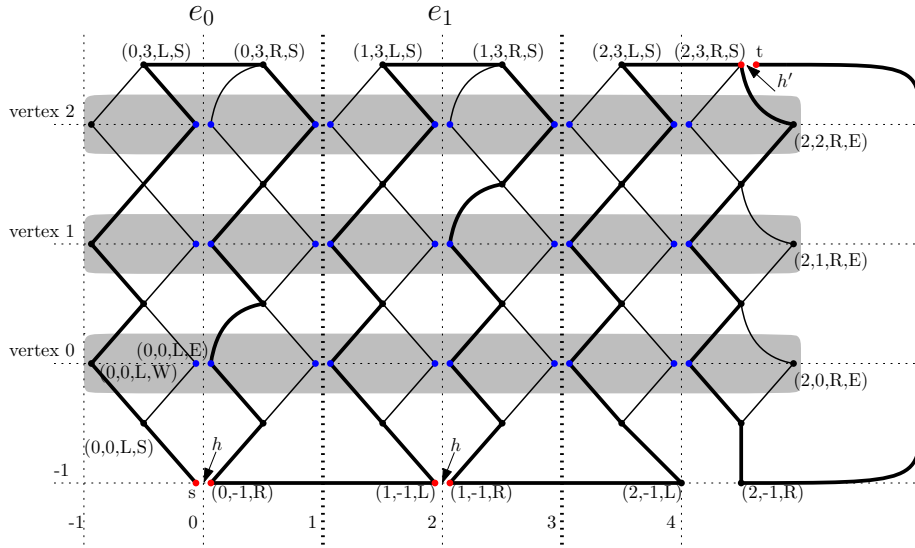
► **Problem 3.** *VERTEX COVER.* We are given a graph G , and a positive integer k . Decide whether there exists a vertex cover in G of size at most k . The instance of the problem is denoted by (G, k) .

For any instance (G, k) of vertex cover we construct an instance (H, f, β, s, t) of β SP that is positive if and only if (G, k) is positive. It will follow from the reduction that β SP is strongly NP-complete, since all of the numerical values in the constructed instance of β SP are bounded by a polynomial in the size of G . The construction follows.

Note that the problem β SP in trees is solvable in quadratic time, by Lemma 6, since in a tree there exists exactly one path between every pair of vertices. Our reduction shows that β SP becomes NP-hard even for graphs whose maximal 2-connected components are cycles.

We put $\beta = n^5$, where n is the number of vertices in G . Let m be the number of edges in G . We identify $V(G)$ with $[n] = \{0, \dots, n-1\}$ and label the edges e_0, \dots, e_{m-1} . The graph $H = (V(H), E(H))$ is constructed as follows; see Figure 3 for an illustration. Roughly, H is composed of chains of 4-cycles arranged in a serial fashion between the distinguished vertices s and t , and drawn as diamonds. Each 4-cycle in a chain (except the two rightmost chains) corresponds to an edge-vertex pair in G , and each pair of consecutive chains except the last one corresponds to an edge of G . Two consecutive chains are joined by an edge or a subdivided edge. The abstract graph H depends only on the number of vertices and edges in G , that is, n and m , and the structure of G is encoded in the drawing of H . Every vertex of H is either a triplet or a 4-tuple: the first element corresponds to an index of an edge of G or is equal to m , the second element corresponds to a vertex of G or is equal to -1 or n , the third element is “L” (for left) or “R” (right), and the fourth element is “E” (for east), “S” (for south) or “W” (for west). Formally, the vertex set is $V(H) = \{s = (0, -1, L), t\} \cup \{(v, e, \alpha, \beta) \mid v \in [n], e \in [m+1], \alpha \in \{L, R\}, \beta \in \{E, S, W\}\} \cup \{(e, n, \alpha, S), (-1, e, \alpha) \mid e \in [m+1], \alpha \in \{L, R\}\}$, and the edge set $E(H) = \{(e, v, \alpha, W)(e, v, \alpha, S), (e, v, \alpha, S)(e, v, \alpha, E), (e, v, \alpha, E)(e, v+1, \alpha, S), (e, v+1, \alpha, S)(e, v, \alpha, W) \mid v \in [n], \alpha \in \{L, R\}, e \in [m+1]\} \cup \{(e, -1, R)(e+1, -1, L), (e, n, L)(e, n, R) \mid e \in [m]\} \cup \{(e, -1, \alpha)(e, 0, \alpha, S) \mid e \in [m+1], \alpha \in \{L, R\}\} \cup \{(m, -1, R)t\}$.

The drawing f represents H in a zig-zag fashion, and has a grid-like structure reminiscent of the edge-vertex incidence matrix of G with rows corresponding to the vertices and columns corresponding to the edges of G . Thus, every chain of 4-cycles of H occupies its own column, and 4-cycles corresponding to the same vertex of G occupy their own row. First, we define $f(v)$ for each $v \in V(H)$. Let $\varepsilon = \beta^{-1} = n^{-5}$. Let $h > 0$ and $h' > 0$ be sufficiently small constants that we specify later. We put $f(t) = (2m + \frac{1}{2} + h', n - \frac{1}{2})$. We put $f((e, -1, L)) = (2e - h, -1)$ and $f((e, -1, R)) = (2e + h, -1)$. We put $f((m, -1, L)) = (2m, -1)$ and $f((m, -1, R)) = (2m + 1, -1)$. We put $f((e, v, L, E)) = (2e - \varepsilon, v)$, $f((e, v, R, E)) = (2e + 1 - \varepsilon, v)$, $f((e, v, L, W)) = (2e - 1 + \varepsilon, v)$, and $f((e, v, R, W)) = (2e + \varepsilon, v)$. We put $f((e, v, L, S)) = (2e - \frac{1}{2}, v - \frac{1}{2})$ and $f((e, v, R, S)) = (2e + \frac{1}{2}, v - \frac{1}{2})$, for $v \in [n]$ and $e \in [m+1]$.



■ **Figure 3** The drawing f of H in the NP-hardness reduction if G is a path on three vertices 0, 1 and 2, with edges $e_0 = 02$ and $e_1 = 21$. Letters in the 3rd and 4th component of a vector representing a vertex stand for Left, Right and East, South, West, respectively. A β -stretch path π between s and t is depicted bold, and corresponds to the minimum vertex cover $VC(\pi)$ of G consisting of the single vertex 2. (A vertex v is contained in $VC(\pi)$ if and only if π passes through $(2, v, R, E)$.)

In f , all of the edges are drawn as straight-line segments except in the following cases. For every $v \in V$ and e_i such that $v \in e_i$, we draw the edge $(i, v, R, W)(i, v + 1, R, S)$ in a close neighborhood of the straight-line segments connecting their end vertices as an xy -monotone curve (that is, a curve that intersects every vertical and horizontal line in at most 1 point) that is longer by more than $20n^{-4}$ in comparison with the straight-line segment $(i, v, R, W)(i, v + 1, R, S)$. We do not care about the shape of the curve and we can think of it as a slightly perturbed line segment. Note that the length of the curve is at most $\sqrt{2}\|f((i, v, R, W)) - f((i, v + 1, R, S))\|$. In the same way, we also draw all of the edges $(m, v, R, E)(m, v + 1, R, S)$, for all $v \in [n]$. Finally, we draw the edge $(m, -1, R)t$ as a concatenation of the horizontal line segment between $f(t)$ and the point $p = f((m, n, R, S)) - (20n^{-4}, 0) \in \mathbb{R}^2$ and a y -monotone curve (that is, every horizontal line intersects the curve at most once) of length $10n$ between $f(m, -1, R)$ and p such that its relative interior does not pass very close to the rest of the drawing.

To finish the drawing $f = f(h, h')$ it remains to choose the values of h and h' . We denote $f_{\text{aux}} = f(0, 0)$ an auxiliary drawing of H with $h = h' = 0$. Let $\pi_e = f_{\text{aux}}(P_e)$ be the 2nd shortest path in f_{aux} between the vertex $(e, -1, L)$ and $(e, -1, R)$, which is independent of the choice of $e \in [m]$. Note that π_e is a path all of whose edges but 1 are drawn as line segments, and its first and last vertex coincide in the drawing. We put $h = \frac{\|\pi_e\|}{2\beta} \leq \frac{20n}{2n^5} = 10n^{-4}$. Let $\pi' = f_{\text{aux}}(P')$ be the $(k + 1)$ -st shortest path in f_{aux} between (m, n, R, S) and t . We put $h' = \frac{\|\pi'\|}{\beta} \leq \frac{20n}{n^5} = 20n^{-4}$. Note that π' is a path with all but $k + 1$ of its edges drawn as line segments, and its first and last vertex t coincide in the drawing.

► **Observation 7.** *The path $f(P_e)$, for $e \in [m]$, and $f(P')$ is shorter than π_e and π' , respectively, and longer than $\|\pi_e\| - 20n^{-4}$ and $\|\pi'\| - 20n^{-4}$.*

For every $v \in [n]$, $e \in [m + 1]$ and $\alpha \in \{L, R\}$, every path in G between s and t must pass either through (e, v, α, W) or (e, v, α, E) . Furthermore, due to the very short distances between blue vertices in the figure we have the following.

7:10 Computing β -Stretch Paths in Drawings of Graphs

► **Lemma 8.** *Let π be a β -stretch path in f between s and t . If π passes through (e, v, L, E) then π passes through (e, v, R, E) and (e', v, α, E) , for all $e' > e$ and $\alpha \in \{L, R\}$. If π passes through (e, v, R, E) then π passes through (e', v, α, E) , for all $e' > e$ and $\alpha \in \{L, R\}$.*

Proof. Suppose that π passes through (e, v, L, E) , and, for the sake of contradiction, let $e' \geq e$ denote the smallest value such that π passes through $(e, v, \alpha, E) \neq (e, v, L, E)$ for some $\alpha \in \{L, R\}$. Suppose that $e = e'$. The other case is treated analogously. By the construction of the drawing f , $\|f((e, v, L, E)) - f((e, v, R, W))\| = 2\epsilon = \frac{2}{\beta}$, and $d_\pi((e, v, L, E), (e, v, R, W)) > 2\epsilon$. Hence, the stretch factor of π is strictly more than β (contradiction). ◀

Proof of Theorem 1. It is easy to verify that the construction of (H, f, β, s, t) can be carried out in polynomial time, and all of the numerical values appearing in the construction of f can be bounded from above by a polynomial function of n , the number of vertices in G . Thus, the strong NP-completeness of β SP follows once we show that (G, k) is a positive instance if and only if (H, f, β, s, t) is a positive instance.

First, if (G, k) is a positive instance, there exists a vertex cover $V' \subseteq V$ of G of size at most k . Let π_{\max} denote the longest path of H in f . Let π be the path in f between s and t passing through (e, v, α, w) if and only if $v \in V'$, for all $e \in [m+1]$ and $\alpha \in \{L, R\}$. We need to show that π is a β -stretch path. Note that π is uniquely determined, and that by the choice of β , the only possible pairs of points that could violate the property of π being a β -stretch path are $(e, -1, L)$ and $(e, -1, R)$, for some $e \in [m]$, and (m, n, R, S) and t . Indeed, it is easy to check that the union of two edges sharing a vertex is always a β -stretch path in f , which follows from the fact that an xy -monotone curve is at most $\sqrt{2}$ -stretch. Hence, in order to violate that π is a β -stretch path, we need to find a pair of points $p \in e_i \in E(H)$ and $q \in e_{i'} \in E(H)$, $e_i \cap e_{i'} = \emptyset$, such that $f(p) \in \pi$, $f(q) \in \pi$, and $\|f(p) - f(q)\| < \frac{\|\pi_{\max}\|}{\beta} < \frac{20n^3}{n^5} = 20n^{-2}$. We can assume that n is sufficiently large such that the pre-image in f of a disk neighborhood of $f(p) \in \mathbb{R}^2$, $p \in H$, with radius $20n^{-2}$ is a single component of H , that does not intersect a pair of edges not sharing a vertex, except when p is very close to $(e, -1, \alpha)$, for some $e \in [m+1]$, $\alpha \in \{L, R\}$, (m, n, R, S) or t , which are colored red in the figure.

Since V' is a vertex cover, we have $d_\pi((i, -1, L), (i, -1, R)) \leq \|\pi_i\|$, for all $i \in [m]$. Indeed, for each $i \in [m]$, the path π misses two non-linear edges incident to $(i, v, R, 0)$ for $v \in e_i$ such that $v \in V'$. Then by Observation 7, $\frac{d_\pi((i, -1, L), (i, -1, R))}{\|f(i, -1, L) - f(i, -1, R)\|} \leq \frac{\|\pi_i\|}{2h} = \beta$. Furthermore, since $|V'| \leq k$, we have $d_\pi((m, n, S, R), t) \leq \|\pi'\|$. Then by Observation 7, $\frac{d_\pi((m, n, S, R), t)}{\|f(m, n, S, R) - f(t)\|} \leq \frac{\|\pi'\|}{h} = \beta$.

Second, if π is a β -stretch path between s and t , let $\text{VC}(\pi) \subseteq V$ be defined as follows. A vertex v is contained in $\text{VC}(\pi)$ if and only if π passes through (m, v, R, E) . Since π is β -stretch, we have $d_\pi((m, n, R, S), t) \leq h'\beta = \frac{\|\pi'\|}{\beta}\beta = \|\pi'\|$. If $|\text{VC}(\pi)| > k$ then by Observation 7 and the length of non-geodesic edges $d_\pi((m, n, R, S), t) > \|\pi'\| - 20n^{-4} + 20n^{-4} = \|\pi'\|$, which is in contradiction with the previous claim. Hence, $|\text{VC}(\pi)| \leq k$. It remains to show that $\text{VC}(\pi)$ is a vertex cover of G .

For the sake of contradiction, suppose that there exists an uncovered edge, that is, an edge $uv = e_i \in E$ such that $e_i \cap \text{VC}(\pi) = \emptyset$. On the one hand, by Lemma 8 and the definition of $\text{VC}(\pi)$, π passes through (i, u, R, W) and (i, v, R, W) . Hence, by Observation 7 and the length of non-geodesic edges, $d_\pi((e, -1, L), (e, -1, R)) > \|\pi_e\| - 20n^{-4} + 20n^{-4} = \|\pi_e\|$. On the other hand, since π is β -stretch, $d_\pi((e, -1, L), (e, -1, R)) \leq 2h\beta = 2\frac{\|\pi_e\|}{2\beta}\beta = \|\pi_e\|$ (contradiction). ◀

Note that our NP-hardness proof involves large stretch values (here, $\beta = n^5$). It would be interesting to show NP-hardness for small stretch values.

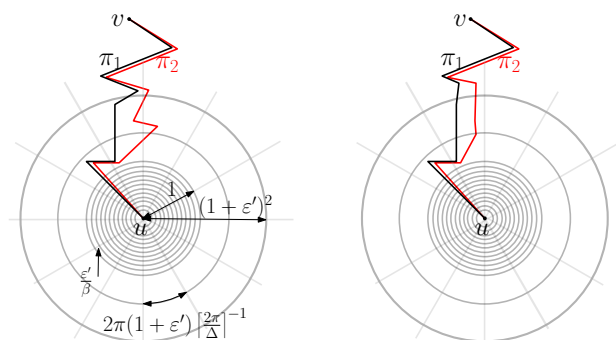
4 Approximation Algorithms

In Section 3, we proved that β SP is strongly NP-complete, which rules out that there exists a FPTAS [22, Section 8] for it, unless $P=NP$; see [22, Corollary 8.6]¹. Let (G, f, β, s, t) be an instance of β SP, and let $\beta^* = \operatorname{argmin}_{\beta}((G, f, \beta, s, t)$ is positive), which is well defined by compactness. In other words, it is highly unlikely that we can approximate β^* within a factor of $(1 + \varepsilon)$, for any $\varepsilon > 0$, in time that is polynomial in both $|V(G)|$ and $\frac{1}{\varepsilon}$.

To complement our hardness result, we show that there exists an algorithm with a quasi-polynomial, that is $O(n^{\operatorname{poly}(\log n)})$, running time that for a given $\varepsilon > 0$ and β , $1 \leq \beta \leq \log^c n$, for some fixed $c \geq 1$, returns a β -stretch path between s and t if a $\beta(1 - \varepsilon)$ -stretch path between s and t exists thereby proving Theorem 2. We assume that ε, c and β satisfy the above properties in the rest of the section. Unless specified otherwise, the base of \log is 2.

4.1 A Path Filtering Scheme

We give a path filtering scheme that we use in Section 4.2 to prove Theorem 2. The main idea behind our algorithm is the following. Since we are aiming only at $\varepsilon > 0$ approximation, we do not need to take into account all of the possible paths between s and t . From a set of paths that are very “similar“ to each other, in the sense that we specify later, we only keep one candidate and delete the rest. Our algorithm proceeds in $\lceil \log n \rceil$ rounds; in the i -th round we compute a set of at most quasi-polynomially many (in terms of n, ε and β) paths of G with at most 2^i edges that are $(1 - \varepsilon_i)\beta$ -stretch in f , for some small ε_i ’s, such that $\varepsilon_0 = \varepsilon, \varepsilon_i > \varepsilon_{i+1}$, and $\varepsilon_{\lceil \log n \rceil} = 0$. In the following, we rigorously define what we mean by “similar”, and how we cluster similar paths. In particular, we cluster paths connecting the same pair of vertices u and v according to their behaviour with respect to stretched radial grids centered at their end vertex u or v ; see Figure 4 for an illustration.



■ **Figure 4** A pair of paths π_1 and π_2 that are not equivalent (on the left) and that are equivalent (on the right) w.r.t. a radial grid centered at u .

Radial grid. Let $\varepsilon > 0$, $\varepsilon' = \varepsilon/\beta$, $r_i = (1 + \varepsilon')^i$ and $\Delta = \frac{\varepsilon'}{1 + \varepsilon'}$. The *radial grid* $F_u(\varepsilon, \beta)$ centered at a point (vertex) $u \in V(G)$ consists of $\lceil \frac{\beta}{\varepsilon'} \rceil$ circles centered at $f(u)$ of radius $i \frac{\varepsilon'}{\beta}$, for $i \in \left[\left\lceil \frac{\beta}{\varepsilon'} \right\rceil \right]$, and circles of radius r_i , for $i \in \lceil [c \log_{1 + \varepsilon'} n] + 1 \rceil$, and $D = \lceil \frac{2\pi}{\Delta} \rceil$ equiangular

¹ Indeed, we can place the vertices in the construction of the reduction on a grid of polynomial size in $n = |V(G)|$ with the unit corresponding to $n^{1/10}$.

7:12 Computing β -Stretch Paths in Drawings of Graphs

spaced rays emanating from $f(u)$. (Recall that we assumed that the shortest edge has length 1 and the largest simple path length is n^c for some constant $c > 0$.) The complement of the radial grid $F_u(\varepsilon, \beta)$ in \mathbb{R}^2 consists of at most $N = D \cdot (\lceil \frac{1}{\varepsilon'} \rceil + \log_{1+\varepsilon'} n^c) = O(\text{poly}(\log n))$ two-dimensional open path connected components, whose closures are *cells* of $F_u(\varepsilon, \beta)$. Note that, ε is treated as a constant and $\beta = O(\text{poly}(\log n))$ by the hypothesis of Theorem 2. In the following, we disregard unbounded cells since they do not intersect $f(G)$. Without loss of generality, we assume that $F_u(\varepsilon, \beta)$ is sufficiently generic with respect to f , that is, $F_u(\varepsilon, \beta) \cap f(G)$ consists of a finite set of points. To this end we might need to slightly perturb the value of ε .

Let $\pi = \pi(u, v)$ be a path in f . Let Σ_π^u denote the subset of cells of $F_u(\varepsilon, \beta)$ that π intersects. We group paths $\pi = \pi(u, v)$ between u and v according to Σ_π^u and approximate distances between u and cells σ in Σ_π^u , which we define next. Let $d_\pi(\sigma, u)$ be the minimum length of the sub-path of π between the point p on π such that $f(p) \in \sigma$ and u . Let r_σ denote the Euclidean distance from u to a furthest point in σ from u . Let $\Xi_\pi^u = \Xi_\pi^u(\varepsilon, \beta) = \left\{ \left(\sigma, \left\lfloor \log_{1+\varepsilon'} \frac{d_\pi(\sigma, u)}{r_\sigma} \right\rfloor \right) \mid \sigma \in \Sigma_\pi^u \right\}$. If π is a β -stretch path, then $\frac{d_\pi(\sigma, u)}{r_\sigma} \leq \beta$. Therefore the second component of each pair in Ξ_π^u is a natural number not bigger than $\lfloor \log_{1+\varepsilon'} \beta \rfloor$.

Path equivalence. Two paths $\pi = \pi(u, v)$ and $\pi' = \pi'(u, v)$ are *equivalent with respect to the radial grid $F_u(\varepsilon, \beta)$* if the first and last edge of π and π' are identical, $\Xi_\pi^u(\varepsilon, \beta) = \Xi_{\pi'}^u(\varepsilon, \beta)$, and the length of π differs from the length of π' by a multiplicative factor of at most $(1 + \varepsilon)$.

Intuitively, equivalent paths pass through the same cells with almost similar distances from u to each intersected cell. Let N be as above, the number of the cells, and $k = \lfloor \log_{1+\varepsilon'} \beta \rfloor + 1$. The crucial aspect of the grid $F_u(\varepsilon, \beta)$ is that there are at most k^N pairwise non-equivalent paths. We have $k^N = (\log_{1+\varepsilon'} \beta)^{cD(\lceil \frac{1}{\varepsilon'} \rceil + \log_{1+\varepsilon'} n)} = O(\text{poly}(\log n)^{\text{poly}(\log n)}) = O(n^{\text{poly}(\log \log n)})$, which is quasi-polynomial in n .

The following lemma (proved in Section 6.1) quantifies the approximation guarantee of our filtering scheme.

► **Lemma 9.** *Let $j \in \mathbb{N}$ such that $j \geq 2$. Let $\pi_1 = \pi_1(u = v_0, v_1), \pi_2 = \pi_2(v_1, v_2) \dots, \pi_j = \pi_j(v_{j-1}, w = v_j)$, and $\pi'_1 = \pi'_1(u = v_0, v_1), \pi'_2 = \pi'_2(v_1, v_2), \dots, \pi'_j = \pi'_j(v_{j-1}, w = v_j)$ be β -stretch paths such that π_i and π'_i , for every $1 \leq i < j$, are equivalent with respect to $F_{v_i}(\varepsilon, \beta_0)$ and $F_{v_{i-1}}(\varepsilon, \beta_0)$, for some $\beta_0 \geq \beta$. Then the following holds.*

If $\pi = \pi_1 \frown \pi_2 \dots \frown \pi_j$ is not a β -stretch path, then $\pi' = \pi'_1 \frown \pi'_2 \dots \frown \pi'_j$ is not a $(1 - 31\varepsilon)\beta$ -stretch path.

4.2 Approximation algorithm for paths

We give an algorithm proving Theorem 2. Refer to the pseudo-code of Algorithm 1. We initialize $\Psi_0 := E(G)$ and $\varepsilon' := \frac{\ln(1-\varepsilon)^{-1}}{32^{\lceil \log n \rceil}}$. The algorithm proceeds in $\lceil \log n \rceil$ many steps, and in the i -th step it computes a set of $\frac{1-\varepsilon}{(1-31\varepsilon')^i} \beta$ -stretch paths Ψ_i in G such that every path in Ψ_i has at most 2^i edges. The set Ψ_{i+1} is computed from $\Psi_{\leq i} = \bigcup_{j \leq i} \Psi_j$ as follows. We pick every pair of distinct paths $\pi_1(u, v) \in \Psi_{\leq i}$ and $\pi_2(v, w) \in \Psi_{\leq i}$ such that the concatenation $\pi = \pi(u, w) = \pi_1(u, v) \frown \pi_2(v, w)$ is a self-intersection free path with at least $2^i + 1$ edges. We put π into Ψ_{i+1} if π is a $\frac{1-\varepsilon}{(1-31\varepsilon')^{i+1}} \beta$ -stretch path. At the end of the $(i + 1)$ -st step, we recursively delete for every pair of vertices u and v of G in Ψ_{i+1} a path $\pi'(u, v)$ if an equivalent path $\pi'(u, v)$ with respect to $F_u(\varepsilon', \beta)$ and $F_v(\varepsilon', \beta)$ still exists in Ψ_{i+1} .

The algorithm outputs a β -stretch path between s and t if $\Psi_{\leq \lceil \log n \rceil}$ contains such a path.

Correctness. Suppose that there exists a $(1 - \varepsilon)\beta$ -stretch path π_0 in f connecting s and t with ℓ edges. We show that the algorithm outputs a β -stretch path connecting s and t . We show by induction on i that after the i -th step of the algorithm, in $\Psi_{\leq i}$ there exists a sequence S_i of $\lceil \frac{\ell}{2^i} \rceil$ paths, whose concatenation is a $\beta \frac{1-\varepsilon}{(1-31\varepsilon')^i}$ -stretch path π_i between s and t . If the claim holds, we are done, since, for a sufficiently large n , we have

$$(1 - 31\varepsilon')^{-\lceil \log n \rceil} (1 - \varepsilon)\beta = \left(1 - \frac{31 \ln(1 - \varepsilon)^{-1}}{32 \lceil \log n \rceil}\right)^{-\lceil \log n \rceil} (1 - \varepsilon)\beta < e^{\ln(1 - \varepsilon)^{-1}} (1 - \varepsilon)\beta = \beta.$$

In the base case the claim holds by the existence of π_0 . By the induction hypothesis, we suppose that the claim holds after the i -th round. We apply Lemma 9 with $\beta_0 := \beta, \varepsilon := \varepsilon'$, and $\beta := \beta \frac{1-\varepsilon}{(1-31\varepsilon')^i}$ to the paths in S_i , whose concatenation π_i in the given order plays the role of π' , and to the equivalent representatives of consecutive pairs of paths in S_i that were not deleted from $\Psi_{\leq i+1}$, whose concatenation plays the role of π . It follows that π is $\beta \frac{1-\varepsilon}{(1-31\varepsilon')^{i+1}}$ -stretch yielding S_{i+1} . Putting $\pi_{i+1} = \pi$ concludes the proof of the correctness of the algorithm.

Running time. The bottleneck of the algorithm is clearly the path filtering scheme that filters all but quasi-polynomially many paths, and therefore the claimed running time follows by the fact that the algorithm ends in $\lceil \log n \rceil$ steps and Lemma 6.

■ **Algorithm 1** Approximation algorithm.

Data: An instance of β SP (G, f, β, s, t) and $\varepsilon > 0$.

Result: A β -stretch path between s and t in f if a $(\beta(1 - \varepsilon))$ -stretch path between s and t exists. (The algorithm can possibly output a β -stretch path even if no $(\beta(1 - \varepsilon))$ -stretch path exists.)

$\varepsilon' := \frac{\ln(1-\varepsilon)^{-1}}{32 \lceil \log n \rceil}$;

$\Psi_0 := E(G), i := 0$; (Ψ_i : the set of candidate β -stretch paths with at most 2^i edges.)

while $\Psi_i \neq \emptyset$ **do**

$\Psi_{i+1} := \emptyset$;

for $\pi_1(u, v), \pi_2(v, w) \in \bigcup_{j \leq i} \Psi_j$ **do**

if $\pi = \pi(u, w) = \pi_1(u, v) \frown \pi_2(v, w)$ has at least $2^i + 1$ edges, and is a $\beta \frac{1-\varepsilon}{(1-31\varepsilon')^{i+1}}$ -stretch path. **then**

 add π to Ψ_{i+1}

while there exists two equivalent paths $\pi(u, v)$ and $\pi'(u, v)$ with respect to

$F_u(\varepsilon', \beta)$ and $F_v(\varepsilon', \beta)$ in Ψ_{i+1} . **do**

 remove π from Ψ_{i+1}

$i \leftarrow i + 1$;

return A β -stretch path between s and t if $\bigcup_i \Psi_i$ contains such path.

4.3 Approximation Algorithm for Cycles

We discuss an extension of the algorithm from Section 4.2 from paths to cycles thereby establishing Theorem 3. Let (G, f, β, s, t) be the input instance for β CP. Let $G_0 = G \setminus \{s, t\}$. We subdivide the edges of G_0 such that every edge has the length at least 1 and at most 2 in f . Let f_0 denote the drawing of G_0 inherited from f . The graph G_0 has polynomially many vertices in terms of the number of vertices of G . We will work with the input instance $(G_0, f_0, \beta, s_0, t_0)$ of β SP, where $s_0, t_0 \in V(G_0)$ and $\varepsilon_0 = 1 - \sqrt{1 - \varepsilon}$. The reason for the

7:14 Computing β -Stretch Paths in Drawings of Graphs

choice of smaller ε_0 is that we will need to work with ε_0 such that $(1 - \varepsilon_0)^2 = (1 - \varepsilon)$. Intuitively, we try to combine all pairs of paths joining the same pair of vertices in $\Psi_{\leq \lceil \log n \rceil}$ constructed by the algorithm from Section 4.2.

A self-intersection free cycle in f_0 separates $f_0(s)$ from $f_0(t)$ if and only if it crosses the line segment between $f_0(s)$ and $f_0(t)$ an odd number of times. In order to keep track of the parity of crossings of paths with the line segment between s and t , we extend the path filtering scheme from Section 4.1 as follows.

Path equivalence. Two paths $\pi = \pi(u, v)$ and $\pi' = \pi'(u, v)$ are *equivalent with respect to the radial grid* $F_u(\varepsilon, \beta)$ in f_0 if the first and last edge of π and π' are identical, $\Xi_\pi^u(\varepsilon, \beta) = \Xi_{\pi'}^u(\varepsilon, \beta)$, the length of π differs from the length of π' by a multiplicative factor of at most $(1 + \varepsilon)$, and additionally the parities of the number of crossings of π' and π with the line segment connecting $f_0(s)$ and $f_0(t)$ are the same.

Algorithm. First, we run a brute-force algorithm to find a β -stretch separating cycle C such that the length of $\gamma = f(C)$ is at least $\frac{4}{\varepsilon_0} + 2$. If we fail to find a β -stretch cycle C , we run the algorithm from Section 4.2 with the input instance $(G_0, f_0, \beta, s_0, t_0)$, for $\varepsilon_0 > 0$, using the previously modified notion of path equivalence with radial grids parametrized by $\varepsilon'(\varepsilon_0) = \frac{\ln(1-\varepsilon_0)^{-1}}{3200 \lceil \log n \rceil}$ and β , that is, $F_u(\varepsilon'/100, \beta)$ rather than $F_u(\varepsilon', \beta)$ in comparison with the original algorithm. The algorithm returns $\Psi_{\leq \lceil \log n \rceil}$. We check if there exists a pair of paths in $\Psi_{\leq \lceil \log n \rceil}$, whose concatenation is a β -stretch cycle C separating s from t . If this is the case we output C .

Correctness. Suppose that there exists a $(1 - \varepsilon)\beta$ -stretch cycle $\gamma = f(C)$ in G_0 separating s from t . Let P_1 and P_2 denote a pair of paths in G between $u \in V(G_0)$ and $v \in V(G_0)$, whose union is C . We choose P_1 and P_2 so that the difference of the length of $\pi_1 = f(P_1)$ and $\pi_2 = f(P_2)$ is minimized. Note that this difference is at most 2. Suppose that π_1 is not shorter than π_2 . We claim that π_1 and π_2 are $\frac{1-\varepsilon}{1-\varepsilon_0}\beta$ -stretch paths. Indeed, for any $p_1, p_2 \in P_1$ $d_\gamma(p_1, p_2) \geq d_{\pi_1}(p_1, p_2) - 2 \geq (1 - \varepsilon_0)d_{\pi_1}(p_1, p_2)$. The first inequality is by the choice of P_1 and P_2 , and the second one by the fact that the length of π_1 is at least $\frac{2}{\varepsilon_0}$, since the length of γ is at least $\frac{4}{\varepsilon_0} + 2$.

Note that $\frac{1-\varepsilon}{1-\varepsilon_0}\beta = (1 - \varepsilon_0)\beta$. Mimicking the proof of the correctness of the algorithm from Section 4.2, we derive that $\Psi_{\leq \lceil \log n \rceil}$ contains a pair of $(1 - \varepsilon_0)\beta$ -stretch paths P'_1 and P'_2 joining the same pair of vertices at P_1 and P_2 such that the concatenation of $\pi'_1 = f_0(P'_1)$ and $\pi'_2 = f_0(P'_2)$ is a β -stretch cycle γ' . To this end we need to adapt Lemma 9 to the case when $u = w$.

► **Lemma 10.** *Let $\varepsilon > 0$ be sufficiently small. Let $j \in \mathbb{N}$ such that $j \geq 2$. Let $\pi_1 = \pi_1(u = v_0, v_1), \pi_2 = \pi_2(v_1, v_2) \dots, \pi_j = \pi_j(v_{j-1}, u = v_j)$, and $\pi'_1 = \pi'_1(u = v_0, v_1), \pi'_2 = \pi'_2(v_1, v_2), \dots, \pi'_j = \pi'_j(v_{j-1}, u = v_j)$ be β -stretch paths such that π_i and π'_i , for every $0 \leq i \leq j$, are equivalent with respect to $F_{v_i}(\varepsilon/100, \beta_0)$ and $F_{v_{i-1}}(\varepsilon/100, \beta_0)$, for some $\beta_0 \geq \beta$. Then the following holds. If $\gamma = \pi_1 \widehat{\cap} \pi_2 \widehat{\cap} \dots \widehat{\cap} \pi_j$ has length at least 20, and is not a β -stretch cycle, then $\gamma' = \pi'_1 \widehat{\cap} \pi'_2 \widehat{\cap} \dots \widehat{\cap} \pi'_j$ is not a $(1 - 31\varepsilon)\beta$ -stretch cycle. Furthermore, γ separates s from t if and only if γ' separates s from t .*

5 Conclusion and Future Work

We proved that β SP is strongly NP-complete, but our reduction seems to work only with large β that is polynomial in the number of vertices n of the input graph. A natural open problem is to determine the complexity of β SP for β constant or logarithmic in n . We proposed a

quasi-polynomial algorithm for β SP that works only for β that is at most logarithmic in n , and that has a quasi-polynomial running already for constant values of β . Therefore we find the problem of devising a PTAS for β SP interesting even when β is a fixed constant.

This leads us to suspect that devising an approximation algorithm for β SP becomes easier if we restrict ourselves to drawings of graphs in which the vertex set is supported by an integer grid of a polynomial size and edges are straight-line segments.

In the future, we intend to extend our work in the following direction, motivated by the computation of districts that avoid gerrymandering. We mark some vertices in a plane graph as “important” and we wish to cut the graph into regions, whose boundaries are β -stretch cycles, such that each region contains exactly one important vertex. A related work by Eppstein et al. [10] describes a method for defining geographic districts in road networks using stable matching. However, their resulting regions might even be disconnected. As we discussed in Section 2, the β -stretch condition is more constraining than local fatness; a locally fat region, whose boundary has a large stretch factor, might look like the shape in Figure 2, which is indicative of a gerrymandered district, with a selective slit removed. We propose that partitioning of geographic regions using β -stretch paths/cycles can lead to districting solutions that may better avoid gerrymandering. We leave this work for future study.


References

- 1 Soroush Alamdari, Timothy M Chan, Elyot Grant, Anna Lubiw, and Vinayak Pathak. Self-approaching graphs. In *International Symposium on Graph Drawing*, pages 260–271. Springer, 2012.
- 2 Boris Aronov, Mark De Berg, Esther Ezra, and Micha Sharir. Improved bounds for the union of locally fat objects in the plane. *SIAM Journal on Computing*, 43(2):543–572, 2014.
- 3 Jonas Azzam and Raanan Schul. How to take shortcuts in Euclidean space: making a given set into a short quasi-convex set. *Proceedings of the London Mathematical Society*, 105(2):367–392, 2012.
- 4 Prosenjit Bose and Michiel Smid. On plane geometric spanners: A survey and open problems. *Computational Geometry*, 46(7):818–830, 2013.
- 5 Ke Chen, Adrian Dumitrescu, Wolfgang Mulzer, and Csaba D Tóth. On the stretch factor of polygonal chains. *arXiv preprint arXiv:1906.10217*, 2019.
- 6 Edith Cohen. Fast algorithms for constructing t -spanners and paths with stretch t . *SIAM Journal on Computing*, 28(1):210–236, 1998.
- 7 Mark de Berg. Improved bounds on the union complexity of fat objects. *Discrete & Computational Geometry*, 40(1):127–140, 2008.
- 8 Annette Ebbers-Baumann, Rolf Klein, Elmar Langetepe, and Andrzej Lingas. A fast algorithm for approximating the detour of a polygonal chain. *Computational Geometry*, 27(2):123–134, 2004.
- 9 David Eppstein. Spanning trees and spanners. *Handbook of computational geometry*, pages 425–461, 1999.
- 10 David Eppstein, Michael T. Goodrich, Doruk Korkmaz, and Nil Mamano. Defining equitable geographic districts in road networks via stable matching. In *Proceedings of the 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, GIS 2017, Redondo Beach, CA, USA, November 7-10, 2017*, pages 52:1–52:4, 2017.
- 11 Mohammad Farshi, Panos Giannopoulos, and Joachim Gudmundsson. Improving the stretch factor of a geometric network by edge augmentation. *SIAM Journal on Computing*, 38(1):226–240, 2008.
- 12 Christian Icking, Rolf Klein, and Elmar Langetepe. Self-approaching curves. *Mathematical Proceedings of the Cambridge Philosophical Society*, 125(3):441–453, 1999.

- 13 Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- 14 Rolf Klein and Martin Kutz. Computing geometric minimum-dilation graphs is np-hard. In *International Symposium on Graph Drawing*, pages 196–207. Springer, 2006.
- 15 Giri Narasimhan and Michiel Smid. Approximating the stretch factor of euclidean graphs. *SIAM Journal on Computing*, 30(3):978–989, 2000.
- 16 Giri Narasimhan and Michiel Smid. *Geometric spanner networks*. Cambridge University Press, 2007.
- 17 David Peleg and Alejandro A Schäffer. Graph spanners. *Journal of graph theory*, 13(1):99–116, 1989.
- 18 Daniel D. Polsby and Robert D. Popper. The third criterion: Compactness as a procedural safeguard against partisan gerrymandering. *Yale Law and Policy Review*, 9(2):301–353, 1991.
- 19 Kristopher Tapp. Measuring political gerrymandering. *The American Mathematical Monthly*, 126(7):593–609, 2019.
- 20 A Frank van der Stappen, Dan Halperin, and Mark H Overmars. The complexity of the free space for a robot moving amidst fat obstacles. *Computational Geometry*, 3(6):353–373, 1993.
- 21 A Frank van der Stappen and Mark H Overmars. Motion planning amidst fat obstacles. In *Proceedings of the tenth annual symposium on Computational geometry*, pages 31–40. ACM, 1994.
- 22 Vijay V Vazirani. *Approximation algorithms*. Springer Science & Business Media, 2013.

6 Appendix

6.1 Proof of Lemma 9



■ **Figure 5** An illustration of Lemma 9 when $j = 2$. A radial grid centered at v_1 , and a pair of paths $\pi = \pi_1 \frown \pi_2$ and $\pi' = \pi'_1 \frown \pi'_2$ that are equivalent with respect to the radial grid centered at v_1 .

► **Lemma 9.** *Let $j \in \mathbb{N}$ such that $j \geq 2$. Let $\pi_1 = \pi_1(u = v_0, v_1), \pi_2 = \pi_2(v_1, v_2) \dots, \pi_j = \pi_j(v_{j-1}, w = v_j)$, and $\pi'_1 = \pi'_1(u = v_0, v_1), \pi'_2 = \pi'_2(v_1, v_2), \dots, \pi'_j = \pi'_j(v_{j-1}, w = v_j)$ be β -stretch paths such that π_i and π'_i , for every $1 \leq i < j$, are equivalent with respect to $F_{v_i}(\varepsilon, \beta_0)$ and $F_{v_{i-1}}(\varepsilon, \beta_0)$, for some $\beta_0 \geq \beta$. Then the following holds.*

If $\pi = \pi_1 \frown \pi_2 \frown \dots \frown \pi_j$ is not a β -stretch path, then $\pi' = \pi'_1 \frown \pi'_2 \frown \dots \frown \pi'_j$ is not a $(1 - 31\varepsilon)\beta$ -stretch path.

Proof. Refer to Figure 5. Assume that π is not a β -stretch path. It follows that either π contains a self-intersection, or there exists two points q and p on π , whose stretch factor is bigger than β . Formally, in either case, there exists a pair of points p and q in G such that

$$\frac{d_\pi(p, q)}{\|f(p) - f(q)\|} > \beta. \quad (5)$$

It is enough to consider the case, in which p is on π_1 and q is on π_j , and p and q are not contained in the union of 2 consecutive edges of π . Indeed, these 2 consecutive edges would be also both on π' by the definition of the equivalent paths.

We show that π' is not a $\beta(1 - 31\varepsilon)$ -stretch path. Consider the cell σ_q and σ_p in the radial grid $F_{v_1}(\varepsilon, \beta_0)$ and $F_{v_{j-1}}(\varepsilon, \beta_0)$, respectively, that contains p and q . Let $q' \in G$ and $q'' \in G$, and $p' \in G$ and $p'' \in G$, respectively, be the points such that $f(q') \in \sigma_q$ and $f(q'') \in \sigma_q$, and $f(p') \in \sigma_p$ and $f(p'') \in \sigma_p$, respectively, minimizing $d_{\pi'}(q', v)$ and $d_{\pi'}(q'', v)$, and $d_{\pi'}(p', v)$ and $d_{\pi'}(p'', v)$. We show that the stretch factor of p' and q' along π' is bigger than $\beta(1 - 16\varepsilon)$, which will conclude the proof. To this end we first derive several simple inequalities.

Since π_1 and π'_1 , and π_j and π'_j are equivalent with respect to $F_{v_1}(\varepsilon, \beta_0)$ and $F_{v_{j-1}}(\varepsilon, \beta_0)$, respectively, the values of $d_{\pi'}(q', v_1)$ and $d_{\pi'}(q'', v_1)$, and $d_{\pi'}(p', v_{j-1})$ and $d_{\pi'}(p'', v_{j-1})$ are within the factor of $(1 + \varepsilon')$ of each other, where $\varepsilon' = \varepsilon/\beta_0$. Since π_1 is a β -stretch paths, $d_\pi(q, q'') \leq \beta L_{\sigma_q}$, where L_{σ_q} is the diameter of σ_q . Therefore

$$d_\pi(q, v_1) = d_\pi(q, q'') + d_\pi(q'', v_1) \leq \beta L_{\sigma_q} + (1 + \varepsilon')d_{\pi'}(q', v_1). \quad (6)$$

The same holds for p, p' and p'' . By the construction of $F_{v_1}(\varepsilon, \beta)$ and $F_{v_{j-1}}(\varepsilon, \beta)$, the diameter of $\sigma \in \{\sigma_p, \sigma_q\}$ such that $r_\sigma = (1 + \varepsilon')^{i+1}$ can be bounded from the above as follows

$$L_\sigma < (1 + \varepsilon')^{i+1} - (1 + \varepsilon')^i + \frac{2\pi\varepsilon'}{1 + \varepsilon'}(1 + \varepsilon')^i \leq (1 + 2\pi)\frac{\varepsilon'}{1 + \varepsilon'}r_\sigma. \quad (7)$$

The upper bound on the diameter of all of the other cells σ contained in the unit disk centered at v_1 and v_{j-1} , respectively, follows if p and q is contained in the annulus between the unit circle and the circle of radius $\frac{1}{\beta_0}$ centered at v_1 and v_{j-1} .

$$L_\sigma < \frac{\varepsilon'}{\beta_0} + \frac{2\pi\varepsilon' \left(r_\sigma - \frac{\varepsilon'}{\beta_0}\right)}{\varepsilon' + 1} < \varepsilon' \left(r_\sigma - \frac{\varepsilon'}{\beta_0}\right) + 2\pi \left(r_\sigma - \frac{\varepsilon'}{\beta_0}\right) \varepsilon' = (1 + 2\pi)\varepsilon' \left(r_\sigma - \frac{\varepsilon'}{\beta_0}\right) \quad (8)$$

By the triangle inequality, $\|f(q) - f(p)\| \geq \|f(q') - f(p')\| - \|f(q) - f(q')\| - \|f(p) - f(p')\| \geq \|f(p') - f(q')\| - L_{\sigma_q} - L_{\sigma_p}$. Therefore

$$\begin{aligned} \beta &\stackrel{(5)}{<} \frac{d_\pi(q, v_1) + d_\pi(v_1, v_2) + \dots + d_\pi(v_{j-1}, p)}{\|f(q) - f(p)\|} \\ &\stackrel{(6)}{\leq} \frac{(1 + \varepsilon')(d_{\pi'}(q', v_1) + \dots + d_{\pi'}(v_{j-1}, p')) + \beta(L_{\sigma_q} + L_{\sigma_p})}{\|f(q') - f(p')\| - L_{\sigma_q} - L_{\sigma_p}} \\ &\leq \frac{d_{\pi'}(q', v_1) + \dots + d_{\pi'}(v_{j-1}, p')}{\|f(q') - f(p')\|} \frac{1 + \varepsilon'}{1 - \frac{L_{\sigma_q} + L_{\sigma_p}}{\|f(q') - f(p')\|}} + \beta \frac{\frac{L_{\sigma_q} + L_{\sigma_p}}{\|f(q') - f(p')\|}}{1 - \frac{L_{\sigma_q} + L_{\sigma_p}}{\|f(q') - f(p')\|}}. \end{aligned} \quad (9)$$

We consider two cases depending on whether π' is a β -stretch path. If π' is not a β -stretch path, then it is also not a $\beta(1 - 16\varepsilon')$ -stretch path and we are done. If π' is a β -stretch path

7:18 Computing β -Stretch Paths in Drawings of Graphs

and both σ_q and σ_p are not contained in the unit disk centered at v_1 and v_{j-1} , respectively, then we must have

$$\|f(p') - f(q')\| \geq \frac{d_{\pi'}(p', q')}{\beta} > \frac{\|f(q') - f(v_1)\| + \|f(v_{j-1}) - f(p')\|}{\beta} \geq \frac{r_{\sigma_q} + r_{\sigma_p}}{(1 + \varepsilon')\beta}. \quad (10)$$

Combining (10) with the upper bound (7) on L_σ from the above yields

$$\frac{L_{\sigma_q} + L_{\sigma_p}}{\|f(q') - f(p')\|} < \frac{(1 + 2\pi)\varepsilon'(r_{\sigma_q} + r_{\sigma_p})}{(r_{\sigma_q} + r_{\sigma_p})/\beta} = (1 + 2\pi)\varepsilon \frac{\beta}{\beta_0} \leq (1 + 2\pi)\varepsilon. \quad (11)$$

If σ_q and σ_p is contained in the annulus between the unit circle and the circle of radius $\frac{1}{\beta_0}$ centered at v_1 and v_{j-1} , respectively, then (10) becomes

$$\|f(p') - f(q')\| > \frac{\|f(q') - f(v_1)\| + \|f(v_{j-1}) - f(p')\|}{\beta} \geq \frac{r_{\sigma_q} - \varepsilon'/\beta_0 + r_{\sigma_p} - \varepsilon'/\beta_0}{\beta}. \quad (12)$$

Then using (8) and (12), we recover the upper bound from (11).

$$\frac{L_{\sigma_q} + L_{\sigma_p}}{\|f(q') - f(p')\|} < \frac{(1 + 2\pi)(r_{\sigma_q} - \varepsilon'/\beta_0 + r_{\sigma_p} - \varepsilon'/\beta_0)\varepsilon'}{\frac{r_{\sigma_q} - \varepsilon'/\beta_0 + r_{\sigma_p} - \varepsilon'/\beta_0}{\beta}} = (1 + 2\pi)\varepsilon \frac{\beta}{\beta_0} \leq (1 + 2\pi)\varepsilon \quad (13)$$

If σ_q is contained in the annulus between the unit circle and the circle of radius $\frac{1}{\beta_0}$ centered at v_1 , and σ_p is not contained in the unit disk centered at v_{j-1} then (10) becomes.

$$\|f(p') - f(q')\| > \frac{\|f(q') - f(v_1)\| + \|f(v_{j-1}) - f(p')\|}{\beta} \geq \frac{\frac{r_{\sigma_p}}{(1 + \varepsilon')} + (r_{\sigma_q} - \frac{\varepsilon'}{\beta_0})}{\beta}. \quad (14)$$

Then using (7),(8) and (10), we again recover the upper bound from (11).

$$\frac{L_{\sigma_q} + L_{\sigma_p}}{\|f(q') - f(p')\|} < \frac{(1 + 2\pi) \left(r_{\sigma_q} - \varepsilon'/\beta_0 + \frac{r_{\sigma_p}}{(1 + \varepsilon')} \right) \varepsilon'}{\frac{\frac{r_{\sigma_p}}{(1 + \varepsilon')} + (r_{\sigma_q} - \varepsilon'/\beta_0)}{\beta}} = (1 + 2\pi)\varepsilon \frac{\beta}{\beta_0} \leq (1 + 2\pi)\varepsilon \quad (15)$$

Finally, if σ_q is contained in the disk of radius $\frac{1}{\beta_0}$ centered at v_1 we distinguish two cases depending on whether σ_p is contained in the unit disk centered at v_{j-1} . If this is the case, q is contained on an edge of π_1 incident to v_j , since π_1 is a β -stretch path, and $\beta_0 \geq \beta$. Hence, as every edge has length at least 1 in f , we have that σ_p is not contained in the unit disk centered at v_{j-1} with diameter $\frac{1}{\beta_0}$. Indeed, q and p are not contained in two consecutive edges of π and therefore they are at distance more than 1 along π , and thus, σ_p is not in the disk of radius $\frac{1}{\beta}$, but $\beta_0 \geq \beta$. Depending on whether σ_p is contained in the unit disk centered at v_{j-1} , we obtain one of the following bounds.

$$\|f(p') - f(q')\| \geq \frac{d_{\pi'}(p', q')}{\beta} > \frac{\|f(v_{j-1}) - f(p')\|}{\beta} \geq \frac{\frac{r_{\sigma_p}}{(1 + \varepsilon')}}{\beta} \quad (16)$$

$$\|f(p') - f(q')\| \geq \frac{d_{\pi'}(p', q')}{\beta} > \frac{\|f(v_{j-1}) - f(p')\|}{\beta} \geq \frac{r_{\sigma_p} - \varepsilon'/\beta_0}{\beta} \quad (17)$$

Then using (7),(8) and (16) and (17), we again recover an upper bound analogous to (11), but worse by a multiplicative factor of 2.

$$\frac{L_{\sigma_q} + L_{\sigma_p}}{\|f(q') - f(p')\|} \leq \frac{2L_{\sigma_p}}{\|f(q') - f(p')\|} \leq 2(1 + 2\pi)\varepsilon \quad (18)$$

Using (11), (13), (15), and (18), (9) can be in every possible case rewritten as follows, which concludes the proof.

$$\begin{aligned} \frac{d_{\pi'}(q', p')}{\|f(q') - f(p')\|} &= \frac{d_{\pi'}(q', v_1) + \dots + d_{\pi'}(v_{j-1}, p')}{\|f(q') - f(p')\|} > \beta \frac{1 - 4(1 + 2\pi)\varepsilon}{1 + \varepsilon/\beta} \\ &> \beta \frac{1 - 4(1 + 2\pi)\varepsilon}{1 + \varepsilon} > \frac{1 - 31\varepsilon}{1 + \varepsilon} \beta > (1 - 31\varepsilon)\beta \end{aligned} \quad \blacktriangleleft$$

6.2 Proof of Lemma 10

► **Lemma 10.** *Let $\varepsilon > 0$ be sufficiently small. Let $j \in \mathbb{N}$ such that $j \geq 2$. Let $\pi_1 = \pi_1(u = v_0, v_1), \pi_2 = \pi_2(v_1, v_2) \dots, \pi_j = \pi_2(v_{j-1}, u = v_j)$, and $\pi'_1 = \pi'_1(u = v_0, v_1), \pi'_2 = \pi'_2(v_1, v_2), \dots, \pi'_j = \pi'_j(v_{j-1}, u = v_j)$ be β -stretch paths such that π_i and π'_i , for every $0 \leq i \leq j$, are equivalent with respect to $F_{v_i}(\varepsilon/100, \beta_0)$ and $F_{v_{i-1}}(\varepsilon/100, \beta_0)$, for some $\beta_0 \geq \beta$. Then the following holds. If $\gamma = \pi_1 \widehat{\cap} \pi_2 \widehat{\cap} \dots \widehat{\cap} \pi_j$ has length at least 20, and is not a β -stretch cycle, then $\gamma' = \pi'_1 \widehat{\cap} \pi'_2 \widehat{\cap} \dots \widehat{\cap} \pi'_j$ is not a $(1 - 31\varepsilon)\beta$ -stretch cycle. Furthermore, γ separates s from t if and only if γ' separates s from t .*

Proof. The proof is analogous to the proof of Lemma 9 except that we consider distances along γ and γ' , which are cycles rather than paths. Due to this reason we slightly weaken some inequalities. The second claim of the lemma is immediate from the definition of the path equivalence. In the following we derive the first claim.

Assume that γ is not a β -stretch cycle. It follows that either γ contains a self-intersection, or there exists two points q and p on π , whose stretch factor is bigger than β . Formally, in either case, there exists a pair of points p and q in G_0 such that

$$\frac{d_\gamma(p, q)}{\|f_0(p) - f_0(q)\|} > \beta. \quad (19)$$

It is enough to consider the case, in which p is on $\pi_{i'}$ and q is on $\pi_{j'}$, and p and q are not contained in the union of 2 consecutive edges of γ . Indeed, these 2 consecutive edges would be also both on γ' by the definition of the equivalent paths, and the edges have length at most 2. Therefore the minimum length curve between p and q in γ is contained in these 2 consecutive edges.

We show that π' is not a $\beta(1 - 31\varepsilon)$ -stretch path. Consider the cell σ_q and σ_p in the radial grid $F_{v_1}(\varepsilon/100, \beta_0)$ and $F_{v_{j-1}}(\varepsilon/100, \beta_0)$, respectively, that contains p and q . We have $\varepsilon' = \frac{\varepsilon}{100\beta_0}$. The rest of the proof differs from the proof of Lemma 9 in the following weaker consequence of a variant of (6), and other inequalities with $d_{\pi'}(q', p')$ that needs to be replaced with $d_{\gamma'}(q', p')$.

$$d_\gamma(q, p) = \beta(L_{\sigma_q} + L_{\sigma_p}) + (1 + 100\varepsilon')d_{\gamma'}(q', p'), \quad (20)$$

where $f_0(q') \in \pi_{i'} \cap \sigma_q$ and $f_0(p') \in \pi_{j'} \cap \sigma_p$.

In the following we derive (20). Let $\pi = \pi(q, p) \subset \gamma$ such that $d_\pi(q, p) = d_\gamma(q, p)$. Let $\pi' = \pi'(q', p') \subset \gamma'$ such that $\pi' \cap \pi'_i \neq \emptyset$ if and only if $\pi \cap \pi_i \neq \emptyset$. Thus, π' is equivalent to π .

Let $\ell(\gamma)$ and $\ell(\gamma')$ denote the length of γ and γ' , respectively. If $d_{\pi'}(q', p') = d_{\gamma'}(q', p')$ then (20) holds by the same argument as in the proof of Lemma 9.

Otherwise, $d_{\gamma'}(q', p') = \ell(\gamma') - d_{\pi'}(q', p')$. Furthermore, $d_{\pi'}(q', p') = \beta(L_{\sigma_q} + L_{\sigma_p}) + (1 + \varepsilon')d_\gamma(q, p) \leq \beta(L_{\sigma_q} + L_{\sigma_p}) + \frac{1}{2}\ell(\gamma) \leq \beta(L_{\sigma_q} + L_{\sigma_p}) + \frac{1}{2}\ell(\gamma')(1 + \varepsilon')$. Combining the previous two (in)equalities we get that $d_{\gamma'}(q', p') \geq \ell(\gamma') - \beta(L_{\sigma_q} + L_{\sigma_p}) - \frac{1}{2}\ell(\gamma')(1 + \varepsilon') = \frac{1}{2}\ell(\gamma')(1 - \varepsilon') - \beta(L_{\sigma_q} + L_{\sigma_p})$.

7:20 Computing β -Stretch Paths in Drawings of Graphs

By the previous paragraph, and (7) and (8),

$$\frac{d_{\pi'}(q', p')}{d_{\gamma'}(q', p')} \leq \frac{\frac{1}{2}\ell(\gamma')(1 + \varepsilon') + \beta(L_{\sigma_q} + L_{\sigma_p})}{\frac{1}{2}\ell(\gamma')(1 - \varepsilon') - \beta(L_{\sigma_q} + L_{\sigma_p})} \leq \frac{\frac{1}{2}\ell(\gamma')(1 + \varepsilon') + 16\varepsilon'\ell(\gamma')}{\frac{1}{2}\ell(\gamma')(1 - \varepsilon') - 16\varepsilon'\ell(\gamma')} \leq \frac{1 + 33\varepsilon'}{1 - 33\varepsilon'} \quad (21)$$

Now, (20) follows from (6) using (21) for sufficiently small ε' . ◀