

Geometric Spanners for Points Inside a Polygonal Domain

Mohammad Ali Abam, Marjan Adeli, Hamid Homapour, and Pooya Zafar Asadollahpoor

Department of Computer Engineering, Sharif University of Technology, Iran
abam@sharif.edu, {madeli, homapour, zafar}@ce.sharif.edu

Abstract

Let \mathcal{P} be a set of n points inside a polygonal domain \mathcal{D} . A polygonal domain with h holes (or obstacles) consists of h disjoint polygonal obstacles surrounded by a simple polygon which itself acts as an obstacle. We first study t -spanners for the set \mathcal{P} with respect to the geodesic distance function π where for any two points p and q , $\pi(p, q)$ is equal to the Euclidean length of the shortest path from p to q that avoids the obstacles interiors. For a case where the polygonal domain is a simple polygon (i.e., $h = 0$), we construct a $(\sqrt{10} + \epsilon)$ -spanner that has $O(n \log^2 n)$ edges. For a case where there are h holes, our construction gives a $(5 + \epsilon)$ -spanner with the size of $O(n\sqrt{h} \log^2 n)$.

Moreover, we study t -spanners for the visibility graph of \mathcal{P} ($VG(\mathcal{P})$, for short) with respect to a hole-free polygonal domain \mathcal{D} . The graph $VG(\mathcal{P})$ is not necessarily a complete graph or even connected. In this case, we propose an algorithm that constructs a $(3 + \epsilon)$ -spanner of size $O(n^{4/3+\delta})$. In addition, we show that there is a set \mathcal{P} of n points such that any $(3 - \epsilon)$ -spanner of $VG(\mathcal{P})$ must contain $\Omega(n^2)$ edges.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Geometric Spanners, Polygonal Domain, Visibility Graph

Digital Object Identifier 10.4230/LIPIcs.SOCG.2015.186

1 Introduction

Background. Let $\mathcal{G} = (V, E)$ be an undirected edge-weighted graph and let $d_{\mathcal{G}}(p, q)$ be the length of the weighted shortest path from p to q in \mathcal{G} . Let $t \geq 1$ be a real number. The subgraph $\mathcal{S} = (V, E_{\mathcal{S}})$ of \mathcal{G} is called a t -spanner if for any two vertices $p, q \in V$, we have $d_{\mathcal{S}}(p, q) \leq t \cdot d_{\mathcal{G}}(p, q)$. Any path from p to q in \mathcal{S} whose weight is at most $t \cdot d_{\mathcal{G}}(p, q)$ is called a t -path. The dilation or stretch factor of \mathcal{S} is the minimum t for which \mathcal{S} is a t -spanner of \mathcal{G} . The size of \mathcal{S} is defined as the number of edges in $E_{\mathcal{S}}$.

t -spanners have been mostly studied on complete graphs coming from finite metric spaces. Let (\mathcal{P}, d) be a finite metric space where \mathcal{P} is a set of n points. Consider the complete graph \mathcal{G}_c over \mathcal{P} where $wt(p, q) = d(p, q)$ (wt denotes weight) for any two points $p, q \in \mathcal{P}$. For any t -spanner \mathcal{S} of \mathcal{G}_c , we have $d_{\mathcal{S}}(p, q) \leq t \cdot d(p, q)$. Indeed, the spanner \mathcal{S} approximates distances in the metric space up to a factor of t . The t -spanner \mathcal{S} is usually called the t -spanner of the metric space (\mathcal{P}, d) . In this paper, we are interested in spanners in a geometric context, i.e., the metric space comes from a geometric space like the Euclidean space. Here, the graph \mathcal{G}_c is the complete Euclidean graph on \mathcal{P} (i.e., weights are the Euclidean distances). A geometric t -spanner is an Euclidean graph \mathcal{S} on \mathcal{P} such that $d_{\mathcal{S}}(p, q) \leq t \cdot |pq|$ for all points $p, q \in \mathcal{P}$ where $|pq|$ denotes the Euclidean distance between p and q .

In some applications like road networks, when constructing spanners, the main goal is to obtain a small dilation while not using too many edges. However, one may want to obtain



© Mohammad Ali Abam, Marjan Adeli, Hamid Homapour, and Pooya Zafar Asadollahpoor; licensed under Creative Commons License CC-BY

31st International Symposium on Computational Geometry (SoCG'15).

Editors: Lars Arge and János Pach; pp. 186–197



Leibniz International Proceedings in Informatics

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

spanners with a number of additional properties such as small weight – weight proportional to the weight of a Minimum Spanning Tree (MST) – and bounded degrees.

Previous work. Althöfer *et al.* [7] were first to study sparse spanners on edge-weighted graphs that have the triangle-inequality property. They showed that for any integer number $t > 0$, there is a $(2t + 1)$ -spanner with $\mathcal{O}(n^{1+1/t})$ edges where n is the number of vertices. This can be applied to any metric space (\mathcal{P}, d) as the complete graph over \mathcal{P} in the metric space has the triangle-inequality property. Geometric spanners have been attracted a lot of attention over the past two decades. It has been shown that for any set of n points in \mathbb{R}^d and any $\varepsilon > 0$, there is a $(1 + \varepsilon)$ -spanner with $\mathcal{O}(n/\varepsilon^{d-1})$ edges – see the recent book by Narasimhan and Smid [13] and references therein for this and many other results on geometric spanners. When the doubling dimension of a metric space is bounded, similar results to the Euclidean setting are possible [12, 14].

Let the points of \mathcal{P} be in a surface $\mathcal{M} \in \mathbb{R}^3$ and let $d_{\mathcal{M}}(p, q)$ be the weight of the shortest (i.e., the minimum weight) path from p to q on \mathcal{M} for any two points $p, q \in \mathcal{P}$. Obviously, $(\mathcal{P}, d_{\mathcal{M}})$ is a metric space and its doubling dimension is not necessarily bounded. Therefore, results to metric spaces with bounded doubling dimension cannot be applied to the metric space $(\mathcal{P}, d_{\mathcal{M}})$ and now the main question is: is it possible to obtain a spanner with a constant stretch factor and a near-linear number of edges for the metric space $(\mathcal{P}, d_{\mathcal{M}})$? Abam *et al.* [3] considered a special case where the surface \mathcal{M} is a plane containing several pillars with width and length of zero but with non-negative height. They assume the points of \mathcal{P} lie at the top of the pillars. This variant can be seen as a set of n weighted points in a plane in which for any two points p and q , their distance is defined to be $wt(p) + |pq| + wt(q)$ where $wt(x)$ is the weight of point x and $|pq|$ is the Euclidean distance of p and q . They presented a $(5 + \varepsilon)$ -spanner with a linear number of edges for any given $\varepsilon > 0$. They also showed that when \mathcal{M} is the boundary of a convex object, it is possible to obtain $(1 + \varepsilon)$ -spanner with a linear number of edges.

Problem statement. Suppose a set \mathcal{P} of n points are given inside a polygonal domain \mathcal{D} which consists of a simple polygon containing h disjoint polygonal holes. The holes and the simple-polygon boundary can be seen as obstacles. Consider the metric space (\mathcal{P}, π) where $\pi(p, q)$ for any points $p, q \in \mathcal{P}$ is equal to the Euclidean length of the shortest path from p to q that avoids the obstacles interiors; the so-called the geodesic distance of p and q . Moreover, let $VG(\mathcal{P})$ be the visibility graph of \mathcal{P} with respect to the polygonal domain \mathcal{D} , i.e., $p, q \in \mathcal{P}$ are connected in $VG(\mathcal{P})$ iff the segment pq avoids the obstacles interiors. Note that $VG(\mathcal{P})$ is not necessarily a complete graph or even a connected graph. In this paper, we investigate the existence of t -spanners with few edges for both the metric space (\mathcal{P}, π) and $VG(\mathcal{P})$. Note that the polygonal domain \mathcal{D} can be seen as a surface. Indeed, obstacles can be seen as walls, tall enough such that any shortest path between two points p and q avoids the walls. Therefore, the metric space (\mathcal{P}, π) is a special case of the metric space $(\mathcal{P}, d_{\mathcal{M}})$ where \mathcal{M} is a surface in \mathbb{R}^3 .

Our results. The first part of our work as explained in Section 2 is devoted to the metric space (\mathcal{P}, π) . For a case where the polygonal domain \mathcal{D} is a simple polygon (i.e., $h = 0$), we construct a $(\sqrt{10} + \varepsilon)$ -spanner that has $\mathcal{O}(n \log^2 n)$ edges. We extend this result to the case where there are h holes. We show that our construction gives a $(5 + \varepsilon)$ -spanner with the size of $\mathcal{O}(n\sqrt{h} \log^2 n)$ for any given $\varepsilon > 0$. The diameter of both spanners is 2. As the second part of our work, in Section 3 we study t -spanners for $VG(\mathcal{P})$. We first show how to obtain

a $(3 + \epsilon)$ -spanner for any given $\epsilon > 0$ of size $O(n^{4/3+\delta})$ for some $\delta > 0$ and then we show that there is a set \mathcal{P} of n points such that any $(3 - \epsilon)$ -spanner of \mathcal{P} must have $\Omega(n^2)$ edges.

2 Spanners for the metric space (\mathcal{P}, π)

Let \mathcal{P} be a set of n points inside a polygonal domain \mathcal{D} which is a simple polygon containing h polygonal disjoint obstacles. Let $\pi(p, q)$ for any two points $p, q \in \mathcal{P}$ be the geodesic distance of p and q with respect to \mathcal{D} . We first present our spanner construction when $h = 0$ in Section 2.1 and then give our general spanner construction in Section 2.2.

2.1 Spanners for points inside a simple polygon


Our spanner construction is based on the SSPD [2, 4] as defined next. For a set Q of n points in \mathbb{R}^d (i.e., the d -dimensional Euclidean space), a pair decomposition of Q is a set of pairs of subsets of Q , such that for every pair of points of $p, q \in Q$ there exists a pair (A, B) in the decomposition such that $p \in A$ and $q \in B$ or vice versa. For a point set A , let $\text{radius}(A)$ be the radius of the minimum enclosing disc of A . An s -Semi-Separated Pair Decomposition (s -SSPD) of Q is a pair decomposition of Q such that for every pair (A, B) , the distance between A and B (i.e. the distance of their minimum enclosing discs) is larger than s times the minimum of the $\text{radius}(A)$ and $\text{radius}(B)$. For a point set Q and a constant $s > 0$, we know there exists an s -SSPD whose weight, $\sum |A| + |B|$ over all pairs, is $O(n \log n)$. The SSPD was introduced to overcome the obesity problem of the Well-Separated Pair Decomposition (WSPD) [9, 15]: there is a set of n points, such that for any WSPD of it, $\sum |A| + |B|$ over all pairs in the WSPD is $\Omega(n^2)$.

Spanner construction. For the given $\epsilon > 0$, we first explain our spanner construction and then prove that the resulting spanner \mathcal{S} is a $(\sqrt{10} + \epsilon)$ -spanner. Our construction is as follows. We partition the simple polygon \mathcal{D} into two simple sub-polygons using a vertical segment ℓ (called the splitting segment) in such a way that each sub-polygon contains at most two-thirds of the points of \mathcal{P} – see [8] for details. For each point $p \in \mathcal{P}$, we compute the point $p_\ell \in \ell$ which has the minimum geodesic distance to p among all points on ℓ . We call p_ℓ the projection of p into ℓ and for a subset A of \mathcal{P} , we define $C_\ell(A)$ to be a point of A whose geodesic distance to ℓ is the smallest. We then compute an s -SSPD for projected points p_ℓ where $s = 4/\epsilon$. Note that some points may have the same projection on ℓ . In this case we treat them as different points while constructing the SSPD. For each pair (A, B) in the SSPD where $\text{radius}(A) \leq \text{radius}(B)$, we add edge $(p, C_\ell(\mathcal{P}(A)))$ to the spanner \mathcal{S} for all points p whose $p_\ell \in A \cup B$ where $\mathcal{P}(A) = \{p \in \mathcal{P} | p_\ell \in A\}$ – recall that an edge (p, q) corresponds to the shortest geodesic path between p and q . We recursively process both simple sub-polygons.

Spanner size. Let $T(n)$ be the size of the resulting spanner \mathcal{S} . Clearly, $T(n) = \sum (|A| + |B|) + T(n_1) + T(n_2)$ where $n_1 + n_2 = n$ and $n_1, n_2 \geq n/3$. Since $\sum (|A| + |B|) = O(n \log n)$ by the SSPD property, we can simply conclude that the spanner size is $O(n \log^2 n)$.

It remains to show that the resulting spanner \mathcal{S} is a $(\sqrt{10} + \epsilon)$ -spanner. We first state the following lemma which plays a key role in our proof showing \mathcal{S} is a $(\sqrt{10} + \epsilon)$ -spanner.

► **Lemma 1.** *Suppose ABC is a right triangle with $\angle CAB = 90$. Let H be a y -monotone path between B and D such that the region bounded by AB , AD , and H is convex where D is some point on edge AC . We have $3|H| + |DC| \leq \sqrt{10}|BC|$ where $|\cdot|$ denotes the Euclidean length.*



■ **Figure 1** A right triangle and a y -monotone convex chain inside it.

Proof. We claim $|H|^2 + |DC|^2 \leq |BC|^2$ and will prove it later. For any two real numbers x and y , we know $(x^2 + y^2)(3^2 + 1^2) \geq (3x + y)^2$. By setting $x = |H|$ and $y = |DC|$, we get $3|H| + |DC| \leq \sqrt{10}|BC|$ as desired.

To prove $|H|^2 + |DC|^2 \leq |BC|^2$, let D' be the point on BC with the same y -coordinate with D . Since H is a convex chain inside triangle $DD'B$ with endpoints D and B , we know

$$|H| \leq |BD'| + |D'D|.$$

Using the above well-known geometric inequality, we have

$$\begin{aligned} |H|^2 + |DC|^2 &\leq (|BD'| + |D'D|)^2 + |DC|^2 \\ &= |BD'|^2 + 2|BD'| \cdot |D'D| + |D'D|^2 + |DC|^2 \\ &= |BD'|^2 + 2|BD'| \cdot |D'D| + |D'C|^2 \\ &\leq |BD'|^2 + 2|BD'| \cdot |D'C| + |D'C|^2 \\ &= (|BD'| + |D'C|)^2 = |BC|^2 \end{aligned}$$




Now, we are ready to prove the main result of this section.

► **Lemma 2.** *The resulting spanner \mathcal{S} of the above construction is a $(\sqrt{10} + \varepsilon)$ -spanner with diameter 2.*

Proof. Any two points $p, q \in \mathcal{S}$ lie at different sides of the splitting segment ℓ at one step of the recursive construction. At this step, there is a semi-separated pair (A, B) that $p_\ell \in A$ and $q_\ell \in B$ or vice versa. WLOG assume $p_\ell \in A$ and $q_\ell \in B$ and moreover assume $\text{radius}(A) \leq \text{radius}(B)$. Let $c = C_\ell(\mathcal{P}(A))$ which of course is a point of \mathcal{P} – see Fig. 2. We recall that among all points whose projections are in A , point c has the minimum geodesic distance to ℓ .

According to our construction at this step of the recursion, edges (p, c) and (q, c) are added to \mathcal{S} . Thus, the length of the shortest path between p and q in \mathcal{S} is at most $\pi(p, c) + \pi(c, q)$. We next show that $\pi(p, c) + \pi(c, q) \leq (\sqrt{10} + \varepsilon)\pi(p, q)$. This implicitly shows the diameter of \mathcal{S} is 2.

Let $\text{SP}(x, y)$ be the shortest path from point x to y with respect to \mathcal{D} for any two points x and y . By the definition of π , the Euclidean length of $\text{SP}(x, y)$ is $\pi(x, y)$. $\text{SP}(p, q)$ definitely intersects ℓ at some point, say r . Let p' (q') be the point at which $\text{SP}(p, q)$ and $\text{SP}(p, p_\ell)$ ($\text{SP}(q, q_\ell)$) get separated – see Fig. 2 to get insight to our notations. It is clear both $\text{SP}(p', p_\ell)$



■ **Figure 2** The splitting segment ℓ partitions the simple polygon into two simple sub-polygons such that each part has at most two-thirds of the points. The projections of points into ℓ are depicted with subscript ℓ .

$SP(q', q_\ell)$ are y -monotone convex chains. $SP(p, q)$ consists of $SP(p, p')$, $SP(p', r)$, $SP(r, q')$ and $SP(q', q)$. We know $\pi(p', r) \geq |p'r|$ and $\pi(q', r) \geq |q'r|$. If we let $h_{p'}$ and $h_{q'}$ be the perpendicular projections of p' and q' on ℓ , in both triangles $q'h_{q'}r$ and $p'h_{p'}r$, the conditions of Lemma 1 hold. All these observations help us to prove the lemma as follows.

Since the distance function π has the triangle-inequality property, we have:

$$\begin{aligned}\pi(p, c) &\leq \pi(p, p_\ell) + |p_\ell c_\ell| + \pi(c_\ell, c) \\ \pi(c, q) &\leq \pi(c, c_\ell) + |c_\ell q_\ell| + \pi(q_\ell, q).\end{aligned}$$

Considering $|c_\ell q_\ell| \leq |c_\ell p_\ell| + |p_\ell r| + |r q_\ell|$ and $\pi(c, c_\ell) \leq \pi(p, p_\ell)$, therefore:

$$\begin{aligned}\pi(p, c) + \pi(c, q) &\leq 3\pi(p, p_\ell) + 2|p_\ell c_\ell| + |p_\ell r| + |r q_\ell| + \pi(q_\ell, q) \\ &= 3\pi(p, p') + 3\pi(p', p_\ell) + 2|p_\ell c_\ell| + |p_\ell r| + |r q_\ell| + \pi(q_\ell, q') + \pi(q', q).\end{aligned}$$

We can apply Lemma 1 to both triangles $q'h_{q'}r$ and $p'h_{p'}r$ and get the following inequalities

$$\begin{aligned}3\pi(p', p_\ell) + |p_\ell r| &\leq \sqrt{10}\pi(p', r) \\ |r q_\ell| + \pi(q_\ell, q') &\leq \sqrt{10}\pi(r, q').\end{aligned}$$

These together yield:

$$3\pi(p', p_\ell) + |p_\ell r| + |r q_\ell| + \pi(q_\ell, q') \leq \sqrt{10}\pi(p', q').$$

Finally, since in the semi-separated pair (A, B) the distance between each two points in A is at most $\frac{2}{s}$ times of the distance between each two points of A and B , we can get:

$$|p_\ell c_\ell| \leq \frac{2}{s}|p_\ell q_\ell| \leq \frac{2}{s}\pi(p, q).$$

If we set $s = \frac{4}{\varepsilon}$, the following inequality holds:




Figure 3 (a) Tight example for the given algorithm in Section 2.1, (b) Any $(2 - \epsilon)$ -spanner in a simple polygonal domain must contain $\Omega(n^2)$ edges, (c) The key property in Lemma 1 does not hold anymore for a polygonal domain with holes.

$$\begin{aligned} \pi(p, c) + \pi(c, q) &\leq 3\pi(p, p') + \sqrt{10}\pi(p', q') + 2|p_\ell c_\ell| + \pi(q', q) \\ &\leq \left(\sqrt{10} + \frac{4}{s}\right)\pi(p, q). \end{aligned}$$



Tight example. As a tight example for our construction, consider the simple polygon in Fig. 3(a) in which $\pi(p, q)$ equals $\sqrt{10}y$ while the shortest path in \mathcal{S} is $10y$.

Lower bound. Consider the simple polygon in Fig. 3(b). When d gets close to 0, $\pi(p, q)$ gets close to $2l$ for any two points p and q . If there is no edge between p and q , the shortest path in \mathcal{S} must go through at least one intermediate vertex, say t . Therefore, the length of the shortest path from p to q , which is at least $\pi(p, t) + \pi(t, q)$, becomes greater than $(2 - \epsilon)\pi(p, q)$ if d is chosen small enough. This implies that to get a $(2 - \epsilon)$ -spanner, we need all edges.

Putting all this together, we get the following theorem.


► **Theorem 3.** Let $\epsilon > 0$ be a given real number. Suppose a set \mathcal{P} of n points is given inside a simple polygon \mathcal{D} . There is a $(\sqrt{10} + \epsilon)$ -spanner with diameter 2 of size $\mathcal{O}(n \log^2 n)$ for the metric space (\mathcal{P}, π) . Moreover, there is a set \mathcal{P} of n points such that any $(2 - \epsilon)$ -spanner of the metric space (\mathcal{P}, π) must contain $\Omega(n^2)$ edges.

2.2 Spanners for points inside a polygonal domain with h holes

Suppose the polygonal domain \mathcal{D} contains h disjoint polygonal holes. Our spanner construction is based on the following decomposition.

► **Lemma 4.** The polygonal domain \mathcal{D} with h holes can be decomposed into $\mathcal{O}(h)$ simple polygons using $\mathcal{O}(h)$ vertical segments (called splitting segments) avoiding the holes interiors such that each simple polygon has at most 3 splitting segments on its boundary.

Proof. As the first step, from the leftmost and rightmost points of each obstacle, we draw two vertical extensions; one going downward until an obstacle is hit and one going upward until an obstacle is hit – see Fig. 4(a). This clearly decomposes the polygonal domain into



■ **Figure 4** (a) Planar decomposition of the polygonal domain \mathcal{D} (first step), (b) Decomposing regions with more than three vertical extensions (second step).

$\mathcal{O}(h)$ simple polygons. But one simple polygon may have $m > 3$ vertical extensions on its boundary. In this case, as the second step, we draw $\mathcal{O}(m)$ vertical extensions inside the simple polygon and decompose it into $\mathcal{O}(m)$ simple polygons such that each new simple polygon has at most three vertical extensions on its boundary. To do that, we first draw a vertical extension such that on each of its side there are roughly half of the vertical extensions. We continue recursively on both sides – see Fig. 4(b). The number of the extra vertical extensions satisfies this recursion: $T(m) = 2T(m/2) + 1$, $T(3) = 0$. Therefore, $T(m) = \mathcal{O}(m)$. As each vertical extension of the first step of the construction is adjacent to at most two simple polygons, the total number of the extra extensions is $\mathcal{O}(h)$. ◀

Suppose the decomposition described in Lemma 4 is available to us. We construct a vertex-weighted graph $\mathcal{G}_{\mathcal{D}}$ as follows. We assign a vertex to each simple polygon and associate it with the number of points in \mathcal{P} that are contained in that simple polygon as its weight. We connect two vertices if their corresponding simple polygons are adjacent. Obviously, $\mathcal{G}_{\mathcal{D}}$ is a planar graph with $\mathcal{O}(h)$ vertices. Our divide-and-conquer construction algorithm uses the following well-known theorem for planar graphs.

► **Theorem 5** ([6]). *Suppose $\mathcal{G} = (V, E)$ is a planar vertex-weighted graph with $|V| = m$. Then, there is a (\sqrt{m}) -separator for \mathcal{G} , i.e., V can be partitioned into three sets A , B and C such that (i) $|C| = \mathcal{O}(\sqrt{m})$, (ii) there is no edge between A and B and (iii) $wt(A), wt(B) \leq 2/3wt(V)$, where $wt(X)$ is the weight summation over all vertices in X .*

Theorem 5 can be applied to the graph $\mathcal{G}_{\mathcal{D}}$ as it is a planar graph. In the following, we explain in details how to construct a spanner \mathcal{S} for the metric space (\mathcal{P}, π) .

1. We first construct $\mathcal{G}_{\mathcal{D}}$ and compute its $\mathcal{O}(\sqrt{h})$ -separator. Let A , B and C be the three sets defined in Theorem 5.
2. We collect $\mathcal{O}(\sqrt{h})$ splitting segments into a set H . More precisely, for each vertex of C (we know $|C| = \mathcal{O}(\sqrt{h})$), we add at most the three splitting segments that appear on the boundary of the simple polygon corresponding to the vertex.
3. For each splitting segment ℓ in H , we apply one recursive step of the given algorithm in Section 2.1.
4. We recursively process the induced subgraphs on A and B until one vertex is left. Each vertex at the last level of the recursion corresponds to a simple polygon in the decomposition of Lemma 4. For each such simple polygon, we apply the whole algorithm given in Section 2.1.

Spanner size. Like the argument given in Section 2.1, at each step of the recursion, for each splitting segment, we add $\mathcal{O}(n \log n)$ edges, and in total for $\mathcal{O}(\sqrt{h})$ splitting segments we add $\mathcal{O}(\sqrt{hn} \log n)$ edges. The whole recursive algorithm except at the leaves of the recursion tree, adds $\mathcal{O}(\sqrt{hn} \log^2 n)$ edges. At the leaf v , we add $\mathcal{O}(n_v \log^2 n_v)$ edges where n_v is the number of points inside the corresponding simple polygon. We know $\sum n_v = n$ and therefore, the total added edges at the leaves is $\mathcal{O}(n \log^2 n)$. All this together state that the spanner size is $\mathcal{O}(\sqrt{hn} \log^2 n)$.

Stretch factor. It is tempting to believe that using the argument of Section 2.1, we can show that the spanner \mathcal{S} is a $(\sqrt{10} + \varepsilon)$ -spanner. But unfortunately, a key property that Lemma 1 relies on, does not hold anymore for a polygon domain with holes. This key property is: $\text{SP}(p, r)$ (i.e., the shortest path from p to r) and $\text{SP}(p, p_\ell)$ topologically are the same. When there are holes, this may not happen as depicted in Fig. 3(c). In the figure, $\text{SP}(p, r)$ goes above the specified hole and $\text{SP}(p, p_\ell)$ goes below that hole. Fortunately, we still can show that the spanner \mathcal{S} has a constant stretch factor.

► **Lemma 6.** *The resulting spanner \mathcal{S} of the above construction is a $(5 + \varepsilon)$ -spanner of the metric space (\mathcal{P}, π) .*

Proof. Consider the top level of our recursive construction. The polygonal domain \mathcal{D} is partitioned into three components, one of which is the separator – see Fig. 5. For any two points p and q which are (i) not in the same component or (ii) in the same separator component but in different simple polygons, the shortest paths from p to q intersects at least one of $\mathcal{O}(\sqrt{h})$ splitting segments collected from the separator. Let ℓ be such a splitting segment. Consider the step of the algorithm working on ℓ . There is a semi-separated pair (A, B) such that $p_\ell \in A$ and $q_\ell \in B$ or vice versa. WLOG assume $p_\ell \in A$ and $q_\ell \in B$ and assume $\text{radius}(A) \leq \text{radius}(B)$. If we let $c = C_\ell(\mathcal{P}(A))$, we know edges (p, c) and (q, c) exist in spanner \mathcal{S} . Hence, the shortest path between p and q in \mathcal{S} is at most $\pi(p, c) + \pi(c, q)$. According to the triangle inequality of π , we have:


$$\begin{aligned} \pi(p, c) &\leq \pi(p, p_\ell) + |p_\ell c_\ell| + \pi(c_\ell, c) \\ \pi(c, q) &\leq \pi(c, c_\ell) + |c_\ell q_\ell| + \pi(q_\ell, q). \end{aligned}$$

We know:

- $\pi(c, c_\ell)$:
 $\pi(c, c_\ell) \leq \pi(p, p_\ell) \leq \pi(p, q)$
- $\pi(p, p_\ell) + \pi(q_\ell, q)$:
 $\pi(p, p_\ell) + \pi(q_\ell, q) \leq \pi(p, r) + \pi(r, q) = \pi(p, q)$
- $|c_\ell q_\ell|$:
 since $|p_\ell r| \leq \pi(p_\ell, p) + \pi(p, r)$ and $\pi(p, p_\ell) \leq \pi(p, r)$ (the same holds for q and q_ℓ), then :

$$\begin{aligned} |c_\ell q_\ell| &\leq |c_\ell p_\ell| + |p_\ell q_\ell| \\ &\leq \frac{2}{s} \pi |p_\ell q_\ell| + |p_\ell q_\ell| \\ &\leq \left(\frac{2}{s} + 1\right) (|p_\ell r| + |r q_\ell|) \\ &\leq \left(\frac{2}{s} + 1\right) (2\pi(p, q)) \end{aligned}$$
- $|p_\ell c_\ell|$:
 From $c_\ell, p_\ell \in A, q_\ell \in B$ and the SSPD property, we have:

$$|p_\ell c_\ell| \leq \frac{2}{s} |p_\ell q_\ell| \leq \frac{4}{s} \pi(p, q)$$



■ **Figure 5** Any path from one component to another one must intersect the separator's boundaries.

All this together give us:

$$\pi(p, q) \leq \pi(p, c) + \pi(c, q) \leq \left(5 + \frac{8}{s}\right)\pi(p, q)$$

We just need to set $s = \frac{8}{\varepsilon}$. Considering the top level of the recursive construction in the proof can be adjusted to the level at which properties (i) or (ii) are satisfied or both points p and q lie in a simple polygon and their shortest path does not intersect any splitting segments of the separators. In the latter, since we run the whole algorithm of Section 2.1, certainly there is a $(\sqrt{10} + \varepsilon)$ -path between p and q . ◀

To summarize, we get the following theorem.

► **Theorem 7.** *Let $\varepsilon > 0$ be a given real number. Suppose a set \mathcal{P} of n points is given inside a simple polygon \mathcal{D} containing h holes. There is a $(5 + \varepsilon)$ -spanner with diameter 2 of size $\mathcal{O}(n\sqrt{h}\log^2 n)$ for the metric space (\mathcal{P}, π) .*

3 Spanners for the visibility graph


Let \mathcal{P} be a set of n points inside a simple polygon \mathcal{D} (i.e., a polygonal domain without hole). Let $VG(\mathcal{P})$ be the visibility graph of \mathcal{P} , which is not necessarily connected. Here, the goal is to find a t -spanner \mathcal{S} with few edges of $VG(\mathcal{P})$, that is, for any two points $p, q \in \mathcal{P}$, their shortest distance in \mathcal{S} is at most t times their shortest distance in $VG(\mathcal{P})$.

Since $VG(\mathcal{P})$ is a special case of weighted graphs holding triangle-inequality property, by applying the algorithm given in [7] we can get the following spanner.

► **Theorem 8.** *For any integer $t > 0$, there is a $(2t + 1)$ -spanner \mathcal{S} such that the number of edges in \mathcal{S} is $\mathcal{O}(n^{1+1/t})$.*

If we set $t = 1$, the above theorem gives us a 3-spanner of size $\mathcal{O}(n^2)$. We next show that it is possible to get $(3 + \varepsilon)$ -spanner of size $\mathcal{O}(n^{4/3+\delta})$ for any $\varepsilon > 0$.

Spanner construction. We first decompose \mathcal{D} using a splitting segment ℓ into two simple polygons \mathcal{D}_L and \mathcal{D}_R each containing at most $2/3n$ points of \mathcal{P} . Let $VG_\ell(\mathcal{P})$ be the subgraph of $VG(\mathcal{P})$ containing every edge of $VG(\mathcal{P})$ that intersects ℓ . We next explain how to find a




■ **Figure 6** (a) The visibility cone of a point. (b) Any $(3 - \varepsilon)$ -spanner of the visibility graph has size of $\Omega(n^2)$.

$(3 + \varepsilon)$ -spanner of $VG_\ell(\mathcal{P})$ with $O(n^{4/3+\delta})$ edges. By recursing on \mathcal{D}_L and \mathcal{D}_R , we can get a $(3 + \varepsilon)$ -spanner of $VG(\mathcal{P})$ with $O(n^{4/3+\delta})$ edges.

The main idea is to represent $VG_\ell(\mathcal{P})$ which is a bipartite graph, as the union of some complete bipartite graphs and find a spanner for each complete bipartite graph. Let $\sigma(p)$ be the visibility cone of p , that is, all half-lines originating from p and intersecting ℓ – see Fig. 6(a). (p, q) is an edge of $VG_\ell(\mathcal{P})$ if and only if $q \in \sigma(p)$ and $p \in \sigma(q)$. For ease of presentation, we call points in \mathcal{D}_L and \mathcal{D}_R red points and blue points, respectively. We map each $\sigma(p)$ to a segment in the dual plane by the standard transformation [11] where a point (a, b) is mapped to the line $y = ax + b$ and vice versa. It is easy to see that (p, q) is an edge of $VG_\ell(\mathcal{P})$ if and only if the segments corresponding to $\sigma(p)$ and $\sigma(q)$ intersect each other. Therefore, the edges in $VG_\ell(\mathcal{P})$ correspond to the intersection of two segments sets and vice versa. Let us call them red segments (corresponding to the red points) and blue segments (corresponding to the blue points). We then construct a segment-intersection-searching data structure [5] for the red segments, which is a multilevel partition tree, each of whose nodes is associated with a canonical subset of red segments. The total size of canonical subsets is $O(n^{4/3+\delta})$. For every blue segment, all red segments intersecting it can be reported as a union of $O(n^{1/3+\delta})$ pairwise disjoint canonical subsets which is useful to construct a clique cover of VG_ℓ without computing all intersections. Therefore, we can represent $VG_\ell(\mathcal{P})$ as the union of some complete bipartite graphs with the total size $O(n^{4/3+\delta})$. We then compute a $(3 + \varepsilon)$ -spanner of size $O(m \log m)$ for each complete bipartite graph with m vertices as described in [1].

Lower bound. Consider a set of $n/2$ points on a segment whose endpoints are $(0, 0)$ and $(0, \alpha)$ and a set of $n/2$ points on a segment whose endpoints are $(1, 0)$ and $(1, \alpha)$. We can put all n points in a simple polygon as depicted in Fig. 6(b) such that every point on each segment can see any point on the other segment and any two points on a segment cannot see each other. Let p and q be two points on the different segments. For an spanner \mathcal{S} of the visibility graph, if the edge (p, q) does not exist in the spanner, any path between p and q in \mathcal{S} must have at least three edges and since the length of each edge is almost the length of (p, q) – we can choose α small enough depending on ε – the spanner cannot be a $(3 - \varepsilon)$ -spanner. Therefore, the spanner must have every edge of the visibility graph which implies the spanner size is $\Omega(n^2)$.

Putting all together, we get the following result.



■ **Figure 7** Lower bound construction.

► **Theorem 9.** For any given $\varepsilon > 0$, there is a $(3 + \varepsilon)$ -spanner of $VG(\mathcal{P})$ that contains $O(n^{4/3+\delta})$ edges for some $\delta > 0$. Moreover, there is a set \mathcal{P} such that any $(3 - \varepsilon)$ -spanner of the visibility graph $VG(\mathcal{P})$ has size of $\Omega(n^2)$.

► **Remark.** If the polygonal domain \mathcal{D} has h holes, we can apply the technique of Section 2.2 to get a $(3 + \varepsilon)$ -spanner of size $O(\sqrt{h}n^{4/3+\delta})$. Moreover, it is possible to find a set \mathcal{P} of n points such that any $(5 - \varepsilon)$ -spanner must have $\Omega(n^{4/3})$ edges. An instance of the line-point incidence problem [10] with $\Omega(n^{4/3})$ incidences can be used to construct the desired instance. To sketch the overall plan, we introduce two sets A (red points) and B (blue points) inside a polygon domain with holes such that (i) for any $p, p' \in A$ and $q, q' \in B$, $|pq|$ is almost $|p'q'|$ and (ii) two points from A cannot see each other and the same holds for B , and (iii) there is no cycle of length 4 in the bipartite visibility graph and (iv) the number of edges in the visibility graph is $\Omega(n^{4/3})$. All this together mean the girth is at least 6 and all edges have almost the same weight. Therefore, any $(5 - \varepsilon)$ -spanner must contain $\Omega(n^{4/3})$ edges. To get the desired point set, consider a $\sqrt{n} \times \sqrt{n}$ grid as depicted in Fig. 7. The number of grid points (p, q) inside the black square where $GCD(p, q) = 1$ is $\Omega(n^{1/3})$. Look at each of these points as a vector. For each vector, we draw a line parallel to the vector from each grid point. The number of different lines for each vector is $O(n^{2/3})$ and the number of incidences is obviously n . In total we have $O(n)$ lines and $\Omega(n^{4/3})$ incidences. We can look at the lines as blue segments. We also put n red parallel segments in the grid points with the negative slope α and very small length. Now, we dualize the segments to cones with the standard transformation. Let A and B be the dual of red and blue segments respectively – note that points in A and B are apexes of the cones. We can put some obstacles such that for every point in A or B , the dual of the visibility cone is exactly the corresponding segment in our incidence construction. It is easy to see that A and B satisfy the required properties by making α and the scale of the grid smaller.

Acknowledgements. The first author would like to thank Pankaj Agarwal and Mark de Berg who initiated the problem and gave valuable suggestions. Moreover, we would like to thank the anonymous reviewers for their valuable comments.

References

- 1 Mohammad Ali Abam, Paz Carmi, Mohammad Farshi, and Michiel Smid. On the power of the semi-separated pair decomposition. *Computational Geometry*, 46(6):631–639, 2013.
- 2 Mohammad Ali Abam, Mark De Berg, Mohammad Farshi, and Joachim Gudmundsson. Region-fault tolerant geometric spanners. *Discrete & Computational Geometry*, 41(4):556–582, 2009.

- 3 Mohammad Ali Abam, Mark De Berg, Mohammad Farshi, Joachim Gudmundsson, and Michiel Smid. Geometric spanners for weighted point sets. *Algorithmica*, 61(1):207–225, 2011.
- 4 Mohammad Ali Abam and Sarel Har-Peled. New constructions of sspds and their applications. *Computational Geometry*, 45(5):200–214, 2012.
- 5 Pankaj K Agarwal and Micha Sharir. Applications of a new space-partitioning technique. *Discrete & Computational Geometry*, 9(1):11–38, 1993.
- 6 Noga Alon, Paul Seymour, and Robin Thomas. Planar separators. *SIAM Journal on Discrete Mathematics*, 7(2):184–193, 1994.
- 7 Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993.
- 8 Prosenjit Bose, Jurek Czyzowicz, Evangelos Kranakis, Danny Krizanc, and Anil Maheshwari. Polygon cutting: Revisited. In *Proceedings of Japanese Conference on Discrete & Computational Geometry (JCDCG'98)*, volume 1763 of *LNCS*, pages 81–92. Springer, 1998.
- 9 Paul B Callahan and S Rao Kosaraju. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *Journal of the ACM*, 42(1):67–90, 1995.
- 10 Kenneth L Clarkson, Herbert Edelsbrunner, Leonidas J Guibas, Micha Sharir, and Emo Welzl. Combinatorial complexity bounds for arrangements of curves and spheres. *Discrete & Computational Geometry*, 5(1):99–160, 1990.
- 11 Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 2008.
- 12 Sarel Har-Peled and Manor Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM Journal on Computing*, 35(5):1148–1184, 2006.
- 13 Giri Narasimhan and Michiel Smid. *Geometric spanner networks*. Cambridge University Press, 2007.
- 14 Kunal Talwar. Bypassing the embedding: algorithms for low dimensional metrics. In *Proceedings of Annual ACM symposium on Theory of computing*, pages 281–290, 2004.
- 15 Kasturi R Varadarajan. A divide-and-conquer algorithm for min-cost perfect matching in the plane. In *Proceedings of Annual Symposium on Foundations of Computer Science*, pages 320–331, 1998.