

THE COMPLEXITY OF GEODESIC SPANNERS

Sarita de Berg,* Marc van Kreveld,* and Frank Staals*

ABSTRACT. A *geometric t -spanner* for a set S of n point sites is an edge-weighted graph for which the (weighted) distance between any two sites $p, q \in S$ is at most t times the original distance between p and q . We study geometric t -spanners for point sets in a constrained two-dimensional environment P . In such cases, the edges of the spanner may have non-constant complexity. Hence, we introduce a novel spanner property: the spanner *complexity*, that is, the total complexity of all edges in the spanner. Let S be a set of n point sites in a simple polygon P with m vertices. We present an algorithm to construct, for any fixed integer $k \geq 1$, a $2\sqrt{2}k$ -spanner with complexity $O(mn^{1/k} + n \log^2 n)$ in $O(n \log^2 n + m \log n + K)$ time, where K denotes the output complexity. When we relax the restriction that the edges in the spanner are shortest paths, such that an edge in the spanner can be any path between two sites, we obtain for any constant $\varepsilon \in (0, 2k)$ a *relaxed geodesic* $(2k + \varepsilon)$ -spanner of the same complexity, where the constant is dependent on ε . When we consider sites in a polygonal domain P with holes, we can construct a relaxed geodesic $6k$ -spanner of complexity $O(mn^{1/k} + n \log^2 n)$ in $O((n + m) \log^2 n \log m + K)$ time. Additionally, for any constant $\varepsilon \in (0, 1)$ and integer constant $t \geq 2$, we show a lower bound for the complexity of any $(t - \varepsilon)$ -spanner of $\Omega(mn^{1/(t-1)} + n)$.

1 Introduction

In the design of networks on a set of nodes, we often consider two criteria: few connections between the nodes, and small distances. Spanners are geometric networks on point sites that replace the small distance criterion by a small detour criterion. Formally, a *geometric t -spanner* for a set S of n point sites is an edge-weighted graph $\mathcal{G} = (S, E)$ for which the (weighted) distance $d_{\mathcal{G}}(p, q)$ between any two sites $p, q \in S$ is at most $t \cdot d(p, q)$, where $d(p, q)$ denotes the distance between p and q in the distance metric we consider [34]. The smallest value t for which a graph \mathcal{G} is a t -spanner is called the *spanning ratio* of \mathcal{G} . The number of edges in the spanner is called the *size* of the spanner.

In the real world, spanners are often constructed in some sort of environment. For example, we might want to connect cities by a railway network, where the tracks should avoid obstacles such as mountains or lakes. One way to model such an environment is by a polygonal domain. In this paper, we study the case where the sites lie in a polygonal domain P with m vertices and h holes, and we measure the distance between two points p, q by their *geodesic distance*: the length of the shortest path between p and q fully contained within P . An example of such a spanner is provided in Figure 1.

* *Utrecht University*, {s.deberg,m.j.vankreveld,f.staals}@uu.nl




Figure 1: A spanner on a set of point sites in a polygonal domain. Because of the orange edges, the spanner has a relatively high complexity.

The spanning ratio and the size of spanners are not the only properties of interest. Many different properties have been studied, such as total weight (or lightness), maximum degree, (hop) diameter, and fault-tolerance [4, 9, 11, 14, 22, 31, 32, 37]. When we consider distance metrics for which the edges in the spanner no longer have constant complexity, another interesting property of spanners arises: the spanner *complexity*, i.e. the total complexity of all edges in the spanner. In our railway example, this corresponds to the total number of bends in the tracks. A spanner with a low number of bends may be desired, as trains can drive faster on straight tracks, and it makes construction cheaper. In this paper, we study this novel property for point sites in a polygonal domain, where the complexity of an edge is simply the number of line segments in the path. In this setting, a single edge may have complexity $\Theta(m)$. Naively, a spanner of size E could thus have complexity $\Theta(mE)$. Our goal is to compute an $O(1)$ -spanner of size $O(n \text{ polylog } n)$ with small complexity, preferably near linear in both n and m .

When studying spanning trees of points, two variants exist: with or without Steiner points. The same is true for spanners, where Steiner points can be used to obtain lighter and sparser spanners [6, 31]. In this paper we focus on the variant where Steiner points are *not* allowed. In subsequent work, de Berg et al. [19] study how the introduction of Steiner points affects the complexity lower bounds, and where to place the Steiner points to obtain a lower complexity spanner. They show that Steiner points have limited utility, in particular most lower bounds are unaffected by introducing a constant number of Steiner points.

Related work. For the Euclidean distance in \mathbb{R}^d , and any fixed $\varepsilon > 0$, there is a $(1 + \varepsilon)$ -spanner of size $O(n/\varepsilon^{d-1})$ [37]. For the more general case of metric spaces of bounded doubling dimension we can also construct a $(1 + \varepsilon)$ -spanner of size $O(n/\varepsilon^{O(d)})$ [13, 24, 27]. These results do not apply when the sites lie in a polygon, and we measure their distances using the geodesic distance. Abam et al. [1] show there is a set of n sites in a simple polygon P for which any geodesic $(2 - \varepsilon)$ -spanner has $\Omega(n^2)$ edges. They also construct a geodesic $(\sqrt{10} + \varepsilon)$ -spanner of size $O(n \log^2 n)$ for sites in a simple polygon, and a geodesic $(5 + \varepsilon)$ -spanner of size $O(n\sqrt{h} \log^2 n)$ for sites in a polygonal domain. Recently, Abam, de Berg,

and Seraji [3] showed that a geodesic $(2 + \varepsilon)$ -spanner with $O(n \log n)$ edges exists for points on a polyhedral terrain, thereby almost closing the gap between the upper and lower bound on the spanning ratio. However, they show only the existence of such a spanner, and leave constructing one open. Moreover, all of these spanners can have high, $\Omega(nm)$, complexity.

Abam, de Berg, and Seraji [3] make use of spanners on an *additively weighted* point set in \mathbb{R}^d . In this setting, the distance between two sites p, q is $w(p) + |pq| + w(q)$ for $p \neq q$, where $w(p)$ is the non-negative weight of a site $p \in S$ and $|pq|$ denotes the Euclidean distance, and 0 for $p = q$. Such additively weighted spanners were studied before by Abam et al. [2], who obtain a $(5 + \varepsilon)$ -spanner of linear size, and a $(2 + \varepsilon)$ -spanner of size $O(n \log n)$. They also provide a lower bound of $\Omega(n^2)$ on the size of any $(2 - \varepsilon)$ -spanner. Abam et al. [3] improve these results and obtain a nearly optimal additively weighted $(2 + \varepsilon)$ -spanner of size $O(n)$.

The other key ingredient for the geodesic $(2 + \varepsilon)$ -spanner of Abam et al. [3] is a *balanced shortest-path separator*. Such a separator consists of either a single shortest path between two points on the boundary of the terrain, or three shortest paths that form a *shortest-path triangle*. This separator partitions the terrain into two subterrains, and we call it balanced when each of these terrains contains roughly half of the sites in S . In their constructive proof for the existence of such a balanced separator, they assume that the three shortest paths in a shortest-path triangle are disjoint, except for their mutual endpoints. However, during their construction it can actually happen that these paths are *not* disjoint. When this happens, it is unclear exactly how to proceed. In Appendix A, we correct this technical issue. Just like for the $(2 + \varepsilon)$ -spanner, the computation of a balanced separator is left for future research. We show how to compute such a separator efficiently in a polygonal domain, using techniques proposed by Thorup [40], and Lipton and Tarjan [33].

Next to spanners on the complete Euclidean geometric graph, spanners under line segment constraints were studied [8, 10, 12, 16, 17]. In this setting, a set C of line segment constraints is provided, where each line segment connects two sites in S and no two line segments properly intersect. The goal is to construct a spanner on the *visibility graph* of S with respect to C . Clarkson [16] showed how to construct a linear sized $(1 + \varepsilon)$ -spanner for this graph. Later, (constrained) Yao- and Θ -graphs were also considered in this setting [8, 12]. If the segments in C form a polygonal domain P , this setting is similar to ours, except that *all* vertices of P are included as sites in S . Thus the complexity of each edge is constant, and additionally it is required that there are short paths between the vertices of P .

Low complexity paths are studied in the *minimum-link path* problem. In this problem, the goal is to find a path that uses a minimal number of links (edges) between two sites in a domain, for example a simple polygon [23, 28, 36, 38]. This problem focuses only on the complexity of the path, with no restriction on the length of the path. Mitchell et al. [35] consider the related problem of finding the *shortest* path with at most k edges between two points p, q in a simple polygon. They give an algorithm to compute a k -link path with length at most $(1 + \varepsilon)$ times the length of the shortest k -link path, for any $\varepsilon > 0$. This result can not be applied to our setting, as the length of our paths should be bounded in terms of $d(p, q)$, i.e. the shortest m -link path, instead of the shortest k -link path.

Our results. We first consider the simple setting where the sites lie in a simple polygon, i.e. a polygonal domain without holes. We show that in this setting any $(3 - \varepsilon)$ -spanner has complexity $\Omega(mn)$ in the worst-case, thus implying that the $(2 + \varepsilon)$ -spanner of Abam, de Berg, and Seraji [3] has also complexity $\Omega(mn)$ in the worst-case, despite having $O(n \log n)$ edges.

To improve this complexity, we first introduce a simple 2-spanner with $O(n \log n)$ edges for an additively weighted point set in a 1-dimensional Euclidean space; see Section 2. In Section 3, we use this result to obtain a geodesic $2\sqrt{2}$ -spanner with $O(n \log^2 n)$ edges for a point set in a simple polygon. We recursively split the polygon by a chord λ such that each subpolygon contains roughly half of the sites, and build a 1-dimensional spanner on the sites projected to λ . We then extend this spanner into one that also has bounded complexity. For any fixed integer $k \geq 1$, we obtain a $2\sqrt{2}k$ -spanner with complexity $O(mn^{1/k} + n \log^2 n)$. Furthermore, we provide an algorithm to compute such a spanner that runs in $O(n \log^2 n + m \log n + K)$ time, where K denotes the output complexity. When we output each edge explicitly, K is equal to the spanner complexity. However, as each edge is a shortest path, we can also output an edge implicitly by only stating the two sites it connects. In this case K is equal to the size of the spanner.

In Sections 4 and 5, we extend our results for a simple polygon to a polygonal domain. There are two significant difficulties in this transition: (i) we can no longer partition the polygon by a line segment such that each subpolygon contains roughly half of the sites, and (ii) the shortest path between two sites p, q may not be homotopic to the path from p to q via another site c .

We solve problem (i) by using a shortest-path separator similar to Abam, de Berg, and Seraji [3]. To apply the shortest-path separator in a polygonal domain, we need new additional ideas, which we discuss in Section 4. In particular, we allow one additional type of separator in our version of a shortest-path separator: two shortest paths from a point in P to the boundary of a single hole. We show that this way there indeed always exists such a separator in a polygonal domain, and provide an $O(m \log m + n \log n)$ time algorithm to compute one, using ideas of Thorup [40].

To overcome problem (ii), we allow an edge (p, q) to be any path from p to q . In networks, the connections between two nodes are often not necessarily optimal paths, the only requirement being that the distance between two hubs does not become too large. Thus allowing other paths between two sites seems a reasonable relaxation. We call such a spanner a *relaxed geodesic spanner*. This way, we obtain a relaxed geodesic $6k$ -spanner of size $O(n \log^2 n)$ and complexity $O(mn^{1/k} + n \log^2 n)$ that can be computed in $O((n + m) \log^2 n \log m + K)$ time. Because our edges always consist of at most three shortest paths, we can again output the edges implicitly in $O(n \log^2 n)$ time.

In Section 6, we use a relaxed geodesic spanner in a simple polygon to improve the spanning ratio of the low complexity spanner from $2\sqrt{2}k$ to $2k + \varepsilon$, for any constant $\varepsilon \in (0, 2k)$. To achieve this, we adapt the refinement suggested by Abam, de Berg, and Seraji [3] in such a way that the spanning ratio is reduced while the complexity increases only by a constant factor (dependent on ε). In the preliminary version of this paper [20] we claimed that application of this refinement was straightforward and that it could also




Figure 2: Construction of the additively weighted 1-dimensional spanner. The green triangle represents all points that are at distance at most $d_w(c, O)$ from O .

be applied in a polygonal domain. However, this no longer guarantees bounded complexity. Our new approach does provide the desired complexity bound, but only in case the domain is a simple polygon. We leave obtaining a $(2k + \varepsilon)$ -spanner of the same complexity in a polygonal domain as an open problem.

Finally, in Section 7, we provide lower bounds on the complexity of (relaxed) geodesic spanners. For any constant $\varepsilon \in (0, 1)$ and integer constant $t \geq 2$, we show a lower bound for the complexity of a $(t - \varepsilon)$ -spanner in a simple polygon of $\Omega(mn^{1/(t-1)} + n)$. Therefore, the spanning ratio of $2k + \varepsilon$ of our $O(mn^{1/k} + n \log^2 n)$ complexity spanner is about a factor two off optimal. For the case of a $(3 - \varepsilon)$ -spanner, we prove an even stronger lower bound of $\Omega(nm)$.

Throughout the paper, we make the general position assumption that all vertices of P and sites in S have distinct x - and y -coordinates. Symbolic perturbation, in particular a shear transformation, can be used to remove this assumption [18].

2 A 1-dimensional additively weighted 2-spanner

We consider how to compute an additively weighted spanner \mathcal{G} in 1-dimensional Euclidean space, where each site $p \in S$ has a non-negative weight $w(p)$. The distance $d_w(p, q)$ between two sites $p, q \in S$ is given by $d_w(p, q) = w(p) + |pq| + w(q)$, where $|pq|$ denotes the Euclidean distance. Without loss of generality, we can map \mathbb{R}^1 to the x -axis, and the weights to the y -axis, see Figure 2. This allows us to speak of the sites left (or right) of some site p .

To construct a spanner \mathcal{G} , we first partition the sites into two sets S_ℓ and S_r of roughly equal size by a point O with $w(O) = 0$. The set S_ℓ contains all sites left of O , and S_r all sites right of O . Sites that lie on the vertical line through O are not included in either of the sets. We then find a site $c \in S$ for which $d_w(c, O)$ is minimal. For all $p \in S$, $p \neq c$, we add the edge (p, c) to \mathcal{G} . Finally, we handle the sets S_ℓ and S_r , excluding the site c , recursively.

Lemma 1. *The graph \mathcal{G} is a 2-spanner of size $O(n \log n)$ and can be constructed in $O(n \log n)$ time.*

Proof. As we add $O(n)$ edges in each level of the recursion, the total number of edges in \mathcal{G} is $O(n \log n)$. Consider two sites $p, q \in S$. Let c be the chosen center at the level of

the recursion where p and q are assigned to different subsets S_ℓ and S_r . Assume without loss of generality that $p \in S_\ell$ and $q \in S_r$. Note that, because $p \in S_\ell$ and $q \in S_r$ we have $d_w(p, q) = d_w(p, O) + d_w(q, O)$. Furthermore, $d_w(c, O) \leq d_w(p, O)$ and $d_w(c, O) \leq d_w(q, O)$, by the choice of c . Because both edges (p, c) and (q, c) are in \mathcal{G} , we get for $d_{\mathcal{G}}(p, q)$:

$$d_{\mathcal{G}}(p, q) \leq d_w(p, O) + 2d_w(c, O) + d_w(q, O) \leq 2d_w(p, O) + 2d_w(q, O) = 2d_w(p, q). \quad (1)$$

When there is no such center point, so p or q lies on the vertical line through O at some level, then it still holds for this level that $d_w(p, q) = d_w(p, O) + d_w(q, O)$. Equation (1) again gives that $d_{\mathcal{G}}(p, q) \leq 2d_w(p, q)$.

We can find a point O that separates the points into two sets S_ℓ and S_r of equal size at each level of the recursion in linear time. Additionally, a linear number of edges is added to the spanner at each level. The running time is thus $O(n \log n)$. \square

3 Spanners in a simple polygon

3.1 A simple geodesic spanner

Just like Abam, de Berg, and Seraji [3], we use our 1-dimensional spanner to construct a geodesic spanner. We are more interested in the simplicity of the spanner than its spanning ratio, as we base our low complexity spanners, to be discussed in Section 3.2, on this simple geodesic spanner. Let P be a simple polygon, and let ∂P denote the polygon boundary. We denote by $d(p, q)$ the geodesic distance between $p, q \in P$, and by $\pi(p, q)$ the shortest (geodesic) path from p to q . We analyze the simple construction using any 1-dimensional additively weighted t -spanner of size $O(n \log n)$. We show that restricting the domain from a polygonal terrain (or a polygonal domain) to a simple polygon improves the spanning ratio from $3t$ to $\sqrt{2}t$. The construction can be refined to achieve a spanning ratio of $t + \varepsilon$, see Section 3.1.1.

As in [1] and [3], we first partition P into two subpolygons P_ℓ and P_r by a line segment λ , such that each subpolygon contains at most two thirds of the sites in S [7]. We assume, without loss of generality, that λ is a vertical line segment and P_ℓ is left of λ . Let S_ℓ be the sites in the closed region P_ℓ , and $S_r := S \setminus S_\ell$. For each site $p \in S$, we then find the point p_λ on λ closest to p . We call p_λ the projection of p . Note that this point is unique, because the shortest path to a line segment is unique in a simple polygon. We denote by S_λ the set of all projected sites. As λ is a line segment, we can define a weighted 1-dimensional Euclidean space on λ , where $w(p_\lambda) := d(p, p_\lambda)$ for each $p_\lambda \in S_\lambda$. We compute a t -spanner $\mathcal{G}_\lambda = (S_\lambda, E_\lambda)$ for this set. For each pair $(p_\lambda, q_\lambda) \in E_\lambda$, we add the edge (p, q) , which is $\pi(p, q)$, to our spanner \mathcal{G} . Finally, we recursively compute spanners for S_ℓ and S_r , and add their edges to \mathcal{G} as well.

Lemma 2. *The graph \mathcal{G} is a geodesic $\sqrt{2}t$ -spanner of size $O(n \log^2 n)$.*

Proof. As \mathcal{G}_λ has $O(n \log n)$ edges (Lemma 1) that directly correspond to edges in \mathcal{G} , and the recursion has $O(\log n)$ levels, we have $O(n \log^2 n)$ edges in total. Let p, q be two sites in S . If both are in S_ℓ (or S_r), then there is a path of length $\sqrt{2}td(p, q)$ by induction. So,




Figure 3: The shortest path $\pi(p, q)$ crosses λ at r . The difference in length between the direct path from z to r and the path through p_λ can be bounded by considering the triangle $\mathcal{T} = (z, z', r)$.

we assume w.l.o.g. that $p \in S_\ell$ and $q \in S_r$. Let r be the intersection point of $\pi(p, q)$ and λ . Observe that p_λ and q_λ must be on opposite sides of r , otherwise r cannot be on the shortest path. We assume, without loss of generality, that p_λ is above r and q_λ below r . Because \mathcal{G}_λ is a t -spanner, we know that there is a weighted path from p_λ to q_λ of length at most $td_w(p_\lambda, q_\lambda)$. As $w(p_\lambda) = d(p, p_\lambda)$, this directly corresponds to a path in the polygon. So,

$$d_{\mathcal{G}}(p, q) \leq d_{\mathcal{G}_\lambda}(p_\lambda, q_\lambda) \leq td_w(p_\lambda, q_\lambda) = t(d(p, p_\lambda) + |p_\lambda r| + |r q_\lambda| + d(q_\lambda, q)). \quad (2)$$

Let z be the point where the shortest paths from p to p_λ and r separate. See Figure 3 for an illustration. Consider the right triangle $\mathcal{T} = (z, z', r)$, where z' is the intersection point of the line perpendicular to λ through z and the line containing λ . Note that z' does not necessarily lie within P . For this triangle we have that

$$|zr| \geq \frac{\sqrt{2}}{2}(|zz'| + |z'r|). \quad (3)$$

Next, we show that the path from z to p_λ is a y -monotone convex polygonal chain ending at or below z' . Consider the vertical ray through z upwards to the polygon boundary. We call the part of ∂P between where the ray hits ∂P and the top endpoint of λ the *top* part of ∂P . Similarly, for a downwards ray, we define the *bottom* part of ∂P . There are no vertices on $\pi(z, p_\lambda)$ from the bottom part of ∂P , because such a vertex would then also occur on the shortest path to r . This is in contradiction with the definition of z . If z sees z' , then $p_\lambda = z'$, otherwise the chain must bend at one or more vertices of the top part of ∂P , and thus lie below z' . It follows that $\pi(z, p_\lambda)$ is contained within \mathcal{T} . Similarly, we conclude that $\pi(z, r)$ is contained within \mathcal{T} . Additionally, this gives us that $d(z, p_\lambda) \leq |zz'| + |z'p_\lambda|$, and $d(z, r) \geq |zr|$. Together with Equation (3) this yields $d(z, p_\lambda) + |p_\lambda r| \leq |zz'| + |z'r| \leq \sqrt{2}|zr| \leq \sqrt{2}d(z, r)$. And thus

$$d(p, p_\lambda) + |p_\lambda r| = d(p, z) + d(z, p_\lambda) + |p_\lambda r| \leq d(p, z) + \sqrt{2}d(z, r) \leq \sqrt{2}d(p, r).$$

Symmetrically, we find for q that $d(q, q_\lambda) + |q_\lambda r| \leq \sqrt{2}d(q, r)$. From this, together with Equation (2), we conclude that $d_{\mathcal{G}}(p, q) \leq t(\sqrt{2}d(p, r) + \sqrt{2}d(r, q)) = \sqrt{2}td(p, q)$. \square

Applying Lemma 2 to the spanner of Section 2 yields a $2\sqrt{2}$ -spanner of size $O(n \log^2 n)$.

3.1.1 A refinement to obtain a $(t + \varepsilon)$ -spanner

Abam et al. [3] refine their spanner construction to obtain a $(2 + \varepsilon)$ -spanner for any constant $\varepsilon > 0$. In the following lemma, we apply their refinement to the construction of the spanner proposed in Section 3.1 and obtain a $(t + \varepsilon)$ -spanner.

Lemma 3. *Using any 1-dimensional additively weighted t -spanner of size $O(n \log n)$, we can construct a $(t + \varepsilon)$ -spanner for a set S of n point sites that lie in a simple polygon P of size $O(c_{\varepsilon, t} n \log^2 n)$, where $c_{\varepsilon, t}$ is a constant depending on ε and t . The number of sites used to construct the 1-dimensional spanner is $O(c_{\varepsilon, t} n)$.*

Proof. In the refined construction, instead of adding only a single point p_λ to S_λ for each site p , we additionally add a collection of $O(1/\delta^2)$ points on λ “close” to p_λ to S_λ , where δ is a constant depending on ε . These additional points all lie within distance $(1 + 2/\delta) \cdot d(p, p_\lambda)$ of p_λ . The points are roughly equally spaced on the line segment within this distance. To be precise, the segment is partitioned into $O(1/\delta^2)$ pieces of length $\delta \cdot d(p, p_\lambda)$, and for each piece i the point $p_\lambda^{(i)}$ closest to p is added to S_λ . The weight of each point is again chosen as the geodesic distance to p . The 1-dimensional spanner(s) \mathcal{G}_λ is then computed on this larger set S_λ . For each edge $(p_\lambda^{(i)}, q_\lambda^{(j)})$ in \mathcal{G}_λ , we again add the edge (p, q) to the final spanner \mathcal{G} .

Abam, de Berg, and Seraji prove that for each $p, q \in S$, there are points $p_\lambda^{(i)}, q_\lambda^{(j)} \in S_\lambda$ such that $d_w(p_\lambda^{(i)}, q_\lambda^{(j)}) \leq (1 + \delta) \cdot d(p, q)$. As \mathcal{G}_λ is a t -spanner for S_λ , choosing $\delta = \varepsilon/t$ implies that $d_{\mathcal{G}}(p, q) \leq t \cdot d_w(p_\lambda^{(i)}, q_\lambda^{(j)}) \leq t(1 + \delta)d(p, q) = (t + \varepsilon)d(p, q)$. In other words, \mathcal{G} is a $(t + \varepsilon)$ -spanner of S . Note that the number of edges in the spanner has increased because we use $O(n/\delta^2)$ instead of $O(n)$ points to compute the 1-dimensional spanner. This results in a spanner of size $O(c_{\varepsilon, t} n \log^2 n)$, where $c_{\varepsilon, t} = O(t^2/\varepsilon^2)$ is a constant depending on ε . \square

3.2 Low complexity geodesic spanners

In general, a geodesic spanner $\mathcal{G} = (S, E)$ in a simple polygon with m vertices may have complexity $O(m|E|)$. It is easy to see that the $2\sqrt{2}$ -spanner of Section 3.1 can have complexity $\Omega(nm)$, just like the spanners in [3]. As one of the sites, c , is connected to all other sites, the polygon in Figure 4 provides this lower bound. The construction in Figure 4 even shows that the same lower bound holds for the complexity of any $(3 - \varepsilon)$ -spanner. A description of the construction and the proof of the lower bound are deferred to Section 7.1. Additionally, the following theorem implies a trade-off between the spanning ratio and the spanner complexity.

Theorem 1. *For any constant $\varepsilon \in (0, 1)$ and integer constant $t \geq 2$, there exists a set of n point sites in a simple polygon P with $m = \Omega(n)$ vertices for which any (relaxed) geodesic $(t - \varepsilon)$ -spanner has complexity $\Omega(mn^{1/(t-1)})$.*




Figure 4: A simple polygon with a set of sites for which any $(3 - \varepsilon)$ -spanner has complexity $\Omega(nm)$.

The proofs of these lower bounds are in Section 7. Next, we present a spanner that almost matches this bound. We first present a $4\sqrt{2}$ -spanner of bounded complexity, and then generalize the approach to obtain a $2\sqrt{2}k$ -spanner of complexity $O(mn^{1/k} + n \log^2 n)$, for any integer $k \geq 2$. In Section 6, we show how to improve the spanning ratio of this spanner to $2k + \varepsilon$, for any $\varepsilon \in (0, 2k)$.

3.2.1 A $4\sqrt{2}$ -spanner of complexity $O(m\sqrt{n} + n \log^2 n)$

To improve the complexity of the geodesic spanner, we adapt our construction for the additively weighted spanner \mathcal{G}_λ as follows. After finding the site $c_\lambda \in S_\lambda$ for which $d_w(c_\lambda, O)$ is minimal, we do not add all edges (p_λ, c_λ) , $p_\lambda \in S_\lambda$, to \mathcal{G}_λ . Instead, we form groups of sites whose original sites (before projection) are “close” to each other in P . We thus consider the location of a site p corresponding to p_λ when constructing the 1-dimensional spanner \mathcal{G}_λ . For each group S_i , we add all edges $(p_\lambda, c_{i,\lambda})$, $p_\lambda \in S_i$, to \mathcal{G}_λ , where $c_{i,\lambda}$ is the site in S_i for which $d_w(c_{i,\lambda}, O)$ is minimal. Finally, we add all edges $(c_{i,\lambda}, c_\lambda)$ to \mathcal{G}_λ .

To make sure the complexity of our spanner does not become too large, we must choose the groups in such a way that the edges in our spanner do not cross “bad” parts of the polygon too often. We achieve this by making groups of roughly equal size, where shortest paths within each group S_i are contained within a region R_i that is (almost) disjoint from the regions R_j of other groups. We first show how to solve a recursion that is later used to bound the complexity of the spanner. The subsequent lemma formally states the properties that we desire of our groups, and bounds the complexity of such a spanner. Finally, we discuss how to form groups that satisfy these properties. As the edges in \mathcal{G}_λ are in one-to-one correspondence with edges in \mathcal{G} , we sometimes slightly abuse terminology by referring to the complexity of the edges in \mathcal{G} that correspond to the edges in \mathcal{G}_λ simply as the complexity of \mathcal{G}_λ .

Lemma 4. *For any integer constant $k \geq 1$ and integer constants $c_1, c_2 \geq 0$, the recursion $T(n, m) = T(n/2, m_1) + T(n/2, m_2) + O(mn^{1/k} + n \log^{c_1} n)$ with $T(1, m) = O(m)$, where $m_1 + m_2 = m + c_2$, solves to $T(n) = O(mn^{1/k} + n \log^{c_1+1} n)$.*

Proof. We will write the recursion as a sum over all levels i and all subproblems j at each level. Because n is halved in each level, the recursion has $O(\log n)$ levels. There are 2^i subproblems at level i of the recursion.

Let $m_{i,j}$ denote the m value used in subproblem j of level i . We consider the sum of the m values over all subproblems at level i which we denote by M_i , so $M_i = \sum_{j=0}^{2^i} m_{i,j}$. We prove by induction that $M_i = m + c_2 \cdot (2^i - 1)$ for any $i \geq 1$. For $i = 1$ this states that $M_1 = m + c_2 \cdot (2^1 - 1) = m + c_2$, which is equivalent to $m_1 + m_2 = m + c_2$. Suppose that the hypothesis $M_i = m + c_2 \cdot (2^i - 1)$ holds for $i = k$. The 2^{k+1} subproblems at level $k+1$ consist of $2^{k+1}/2 = 2^k$ pairs of subproblems (j_1, j_2) , each generated by a subproblem j at level k , for which $m_{k+1,j_1} + m_{k+1,j_2} = m_{k,j} + c_2$. We can thus find M_{k+1} by summing over these pairs. So, $M_{k+1} = \sum_{(j_1,j_2)} (m_{k,j} + c_2) = M_k + c_2 \cdot 2^k = m + c_2 \cdot (2^k - 1) + c_2 \cdot 2^k = m + c_2 \cdot (2^{k+1} - 1)$.

We are now ready to formulate the recursion as a summation. For simplicity we use that $M_i \leq m + c_2 2^i$.

$$\begin{aligned}
T(n, m) &= \sum_{i=0}^{O(\log n)} \sum_{j=0}^{2^i} \left(m_{i,j} \left(\frac{n}{2^i} \right)^{1/k} + \frac{n}{2^i} \log^{c_1} \left(\frac{n}{2^i} \right) \right) \\
&= \sum_{i=0}^{O(\log n)} \sum_{j=0}^{2^i} m_{i,j} \left(\frac{n}{2^i} \right)^{1/k} + \sum_{i=0}^{O(\log n)} \sum_{j=0}^{2^i} \frac{n}{2^i} \log^{c_1} \left(\frac{n}{2^i} \right) \\
&= \sum_{i=0}^{O(\log n)} \left(\frac{n}{2^i} \right)^{1/k} \sum_{j=0}^{2^i} m_{i,j} + O(n \log^{c_1+1} n) \\
&\leq \sum_{i=0}^{O(\log n)} \left(\frac{n}{2^i} \right)^{1/k} (m + c_2 2^i) + O(n \log^{c_1+1} n) \\
&= mn^{1/k} \sum_{i=0}^{O(\log n)} \frac{1}{2^{i/k}} + c_2 n^{1/k} \sum_{i=0}^{O(\log n)} 2^{(1-1/k)i} + O(n \log^{c_1+1} n) \\
&= O(mn^{1/k}) + c_2 n^{1/k} \cdot O(n^{1-1/k}) + O(n \log^{c_1+1} n) \\
&= O(mn^{1/k} + n \log^{c_1+1} n). \quad \square
\end{aligned}$$

Lemma 5. *If the groups satisfy the following properties, then \mathcal{G} has $O(m\sqrt{n} + n \log^2 n)$ complexity:*

1. each group contains $\Theta(\sqrt{n})$ sites, and
2. each vertex of P appears on the shortest path between a pair of sites in the same group for only $O(1)$ groups.

Proof. We will first prove the complexity in P of the edges added in one level of \mathcal{G}_λ is $O(m\sqrt{n} + n)$. Two types of edges are added to the spanner: (a) edges from some c_i to c , and (b) edges from some $p \in S_i$ to c_i . Observe that we drop the subscript λ here, i.e. c

corresponds to c_λ , because we refer to the location of the sites in the polygon rather than in the 1-dimensional space. According to property 1, there are $\Theta(\sqrt{n})$ groups, and thus $\Theta(\sqrt{n})$ type (a) edges, that each have a complexity of $O(m)$. Thus the total complexity of these edges is $O(m\sqrt{n})$. Let r_i be the maximum complexity of a shortest path between any two sites in S_i and let V_i be the set of vertices this path visits. Property 2 states that for any $v \in V_i$ it holds that $|\{j \mid v \in V_j\}| = O(1)$, which implies that $\sum_i r_i = O(m)$. The complexity of all type (b) edges is thus $O(n) + \sum_i r_i O(\sqrt{n}) = O(m\sqrt{n} + n)$.

Next, we show that in both recursions, the recursion for \mathcal{G}_λ and the recursion on P_ℓ and P_r , not only the number of sites, but also the complexity of the polygon is divided over the two subproblems. Recall that the 1-dimensional space corresponds to the *vertical* line segment λ . So, splitting the sites into left and right of O in the 1-dimensional space corresponds to splitting the polygon *horizontally* at O : all sites left (right) of O in the 1-dimensional space lie in the part of the polygon below (above) this horizontal line segment. Thus, shortest paths between sites left of O use part of the polygon that is disjoint from the shortest paths between the sites right of O . This means that for two subproblems we have that $m_1 + m_2 = m$, where m_i denotes the maximum complexity of a path in subproblem i . The recursion for the complexity is now given by

$$T(n, m) = T(n/2, m_1) + T(n/2, m_2) + O(m\sqrt{n} + n), \text{ with } m_1 + m_2 = m.$$

According to Lemma 4 this solves to $T(n) = O(m\sqrt{n} + n \log n)$.

Similarly, the split by λ divides the polygon into two subpolygons, while adding at most two new vertices. As all vertices, except for the endpoints of λ , are in P_ℓ or P_r (not both), the total complexity of both subpolygons is at most $m + 4$. We obtain the following recursion

$$T(n, m) = T(n/2, m_1) + T(n/2, m_2) + O(m\sqrt{n} + n \log n), \text{ with } m_1 + m_2 = m + 4.$$

Lemma 4 again states this solves to $T(n) = O(m\sqrt{n} + n \log^2 n)$. \square

Forming the groups. To form groups that satisfy these two properties, we consider the shortest path tree SPT_c of c : the union of all shortest paths from c to the vertices of P . We include the sites $p \in S \setminus \{c\}$ as leaves in SPT_c as children of their apex, i.e., the last vertex on $\pi(c, p)$. This gives rise to an ordering of the sites in S , and thus of the weighted sites in S_λ , based on the in-order traversal of the tree. We assign the first $\lceil \sqrt{n} \rceil$ sites to S_1 , the second $\lceil \sqrt{n} \rceil$ to S_2 , etc. See Figure 5.

Clearly these groups satisfy property 1. Proving that they also satisfy property 2 is more involved. For each group S_i , consider the minimal subtree \mathcal{T}_i of SPT_c containing all $p \in S_i$. \mathcal{T}_i defines a polygonal region R_i in P as follows. Refer to Figure 5 for an illustration. Let v_i be the root of \mathcal{T}_i . Consider the shortest path $\pi(v_i, a)$, where a is the first site of S_i in \mathcal{T}_i by the ordering used before. Let π_a be the path obtained from $\pi(v_i, a)$ by extending the last segment of $\pi(v_i, a)$ to the boundary of P . Similarly, let π_b be such a path for the last site of S_i in \mathcal{T}_i . We take R_i to be the region in P rooted at v_i and bounded by π_a , π_b , and the boundary of P , that contains the sites in S_i . In case v_i is c , we split R_i into two regions R_j and R_k , such that the angle of each of these regions at c is at most π . The set




Figure 5: The shortest path tree of c . The polygon vertices are grey in the tree. Each group S_i has an associated polygonal region R_i in P .

S_i is then also split into two sets S_j and S_k accordingly. The following three lemmas on R_i and \mathcal{T}_i together imply that the groups satisfy property 2.

Lemma 6. *Only vertices of P that are in \mathcal{T}_i can occur in R_i .*

Proof. Let v be a vertex in R_i . The paths π_a and π_b are shortest paths from the polygon boundary to v_i . Additionally, we can extend these paths with $\pi(v_i, c)$ to obtain shortest paths to the site c . The shortest path from v to c cannot intersect either of these shortest paths twice, thus v is in \mathcal{T}_i . \square

Lemma 7. *All shortest paths between sites in S_i are contained within R_i .*

Proof. Suppose there is a shortest path $\pi(p, q)$, $p_\lambda, q_\lambda \in S_i$, that is not contained within R_i . Then $\pi(p, q)$ must exit R_i through π_a and enter again through π_b , or the other way around. The path thus goes around v_i . Consider a line through v_i that does not pass through the interior of R_i . Note that this exists because v_i is not a reflex vertex of R_i . As P is a simple polygon, the line must intersect $\pi(p, q)$ twice. This provides a shortcut of the shortest path by following along the line, which is a contradiction. \square

Lemma 8. *Any vertex $v \in SPT_c$ occurs in at most two trees \mathcal{T}_i and \mathcal{T}_j as a non-root node.*

Proof. Let $v \in \mathcal{T}$, and $\mathcal{T}(v)$ be the subtree rooted at v . Suppose v is a non-root node of three trees $\mathcal{T}_i, \mathcal{T}_j$, and \mathcal{T}_k , $i < j < k$. Then there must be three sites $p_i \in S_i, p_j \in S_j$, and $p_k \in S_k$ in $\mathcal{T}(v)$. As our groups occur in order, we know p_i is before p_j , and p_j is before p_k in the in-order traversal of SPT_c . Let p be the parent of v . As v is a non-root node of each of the subtrees, we know that p is in $\mathcal{T}_i, \mathcal{T}_j$, and \mathcal{T}_k as well. This implies that there must be a site $p'_j \in S_j$ in $\mathcal{T} \setminus \mathcal{T}(v)$. If p'_j is before p_j , then p'_j is also before p_i , because $p'_j \notin \mathcal{T}(v)$. This implies that p_i is in S_j , because it lies between $p'_j \in S_j$ and $p_j \in S_j$, which is a contradiction. If p'_j is after p_j , then the same reasoning implies $p_k \in S_j$, which is also a contradiction. \square

Note that the root r of \mathcal{T}_i is never used in a shortest path between sites in S_i , because r cannot be a reflex vertex of R_i . Consequently, Lemma 5 states that the spanner has complexity $O(m\sqrt{n} + n \log^2 n)$.

Lemma 9. *The graph \mathcal{G} is a geodesic $4\sqrt{2}$ -spanner of size $O(n \log^2 n)$.*

Proof. We prove the 1-dimensional spanner \mathcal{G}_λ is a 4-spanner with $O(n \log n)$ edges. Together with Lemma 2, this directly implies \mathcal{G} is a $4\sqrt{2}$ -spanner with $O(n \log^2 n)$ edges.

To construct \mathcal{G}_λ , we still add only a single edge for each site in each level of the recursion. Thus, the total number of edges is $O(n \log n)$. Again, consider two sites $p_\lambda, q_\lambda \in S_\lambda$, and let c_λ be the chosen center point at the level where p_λ and q_λ are separated by O . Let S_i be the group of p_λ and S_j the group of q_λ . Both the edges $(p_\lambda, c_{i,\lambda})$ and $(c_{i,\lambda}, c_\lambda)$ are in \mathcal{G}_λ , similarly for q_λ . We thus have a path $p_\lambda \rightarrow c_{i,\lambda} \rightarrow c_\lambda \rightarrow c_{j,\lambda} \rightarrow q_\lambda$ in \mathcal{G}_λ . Using that $d_w(p_\lambda, c_{i,\lambda}) \leq d_w(p_\lambda, O) + d_w(c_{i,\lambda}, O)$, because of the triangle inequality, and $d_w(c_{i,\lambda}, O) \leq d_w(p_\lambda, O)$, we find:

$$\begin{aligned} d_{\mathcal{G}_\lambda}(p_\lambda, q_\lambda) &= d_w(p_\lambda, c_{i,\lambda}) + d_w(c_{i,\lambda}, c_\lambda) + d_w(c_\lambda, c_{j,\lambda}) + d_w(c_{j,\lambda}, q_\lambda) \\ &\leq d_w(p_\lambda, O) + 2d_w(c_{i,\lambda}, O) + 2d_w(c_\lambda, O) + 2d_w(c_{j,\lambda}, O) + d_w(q_\lambda, O) \\ &\leq 4d_w(p_\lambda, O) + 4d_w(q_\lambda, O) \\ &= 4d_w(p_\lambda, q_\lambda). \end{aligned} \quad \square$$

3.2.2 A $2\sqrt{2}k$ -spanner of complexity $O(mn^{1/k} + n \log^2 n)$

We first sketch how to generalize the approach of Section 3.2.1 to obtain a spanner with a trade-off between the (constant) spanning ratio and complexity, and then formally prove the result in Lemma 10. Fix $N = n^{1/k}$, for some integer constant $k \geq 1$. In the construction of \mathcal{G}_λ , instead of $\Theta(\sqrt{n})$ groups, we create $\Theta(N)$ groups. For each of these groups we select a center, and then partition the groups further recursively. In each level of the recursion the same fixed value N is used to create the groups. By connecting each center to its parent center, we obtain a tree of height k . This results in a spanning ratio of $2\sqrt{2}k$.

Lemma 10. *For any integer constant $k \geq 1$, there exists a geodesic $2\sqrt{2}k$ -spanner of size $O(n \log^2 n)$ and complexity $O(mn^{1/k} + n \log^2 n)$.*

Proof. We prove that the spanner \mathcal{G}_λ we just sketched is a $2k$ -spanner with $O(n \log n)$ edges whose complexity is $O(mn^{1/k} + n \log n)$. Lemma 2, together with Lemma 4, then implies that \mathcal{G} is a geodesic $2\sqrt{2}k$ -spanner of complexity $O(mn^{1/k} + n \log^2 n)$ and size $O(n \log^2 n)$. We first describe the 1-dimensional spanner \mathcal{G}_λ in more detail, then analyze the complexity, and finally analyze its size and spanning ratio.

Fix $N = n^{1/k}$. When building a single level of the 1-dimensional spanner \mathcal{G}_λ , instead of $\Theta(\sqrt{n})$ groups of size $\Theta(\sqrt{n})$, we create $\Theta(N)$ groups of size $\Theta(n^{1-1/k})$, based on the shortest path tree of c , for some integer constant $k \geq 1$. After selecting a center c_i for each group S_i , we recursively split the groups further based on the shortest path trees of the new

centers c_i , until we reach groups of size one. For each group S_i at level j of this recursion, a center $c_i^{(j)}$ is selected as the site in S_i for which $d_w(c_i^{(j)}, O)$ is minimal. We add an edge from each $c_i^{(j)}$ to its parent $c_i^{(j-1)}$. The final tree T obtained this way has height k .

Let $r_i^{(j)}$ be the number of vertices in $R_i^{(j)}$ of a group S_i at level j , where $R_i^{(j)}$ is defined as in Section 3.2.1. In other words, $r_i^{(j)}$ is the maximum complexity of a shortest path between two sites in S_i at level j , as in the proof of Lemma 5. Let $S_1^{(j)}, \dots, S_{O(N)}^{(j)}$ be the subgroups of some group $S_i^{(j-1)}$. Then the corresponding regions $R_1^{(j)}, \dots, R_{O(N)}^{(j)}$ partition $R_i^{(j-1)}$, and Lemma 8 implies that a vertex of $R_{O(N)}^{(j)}$ can be in at most two of the smaller regions. Thus, we still have that $\sum_i r_i^{(j)} = O(m)$. This implies that the complexity all edges from level j to $j - 1$ is $O(mn^{1/k} + n)$. As T has constant height $O(k)$, the total complexity of the edges in T is also $O(mn^{1/k} + n)$. Lemma 4 implies that the complexity of \mathcal{G}_λ is $O(mn^{1/k} + n \log n)$.

In a single level of the recursion to build the 1-dimensional spanner \mathcal{G}_λ , only one edge is added for each site, namely to its parent in the tree. We thus still add $O(n)$ edges in each level of the recursion, and $O(n \log n)$ edges in total.

Consider two sites $p, q \in S_\lambda$, and let c be the chosen center point at the level where p and q are separated by O . The spanning ratio of the 1-dimensional spanner is determined by the number of sites we visit on a path from $p \in S_\ell$ to $q \in S_r$. Observe that the height of the tree T is k . We assume w.l.o.g. that the path from p to q is as long as possible, i.e. visits $2k - 1$ vertices. In a slight abuse of notation, we denote by $c_p^{(j)}$ the centers on the path in T from p to the root c . Then there is a path $p \rightarrow c_p^{(k-1)} \rightarrow c_p^{(k-2)} \rightarrow \dots \rightarrow c^{(0)} (= c) \rightarrow \dots \rightarrow c_q^{(k-2)} \rightarrow c_q^{(k-1)} \rightarrow q$. Using that $d_w(c_p^{(j)}, c_p^{(j+1)}) \leq d_w(c_p^{(j)}, O) + d_w(c_p^{(j+1)}, O)$ and that $d_w(c_p^{(j)}, O) \leq d_w(p, O)$, we find:

$$\begin{aligned} d_{\mathcal{G}_\lambda}(p, q) &= d_w(p, c_p^{(k-1)}) + \sum_{j=0}^{k-2} d_w(c_p^{(j+1)}, c_p^{(j)}) + \sum_{j=0}^{k-2} d_w(c_q^{(j)}, c_q^{(j+1)}) + d_w(c_q^{(k-1)}, q) \\ &\leq d_w(p, O) + 2 \sum_{j=1}^{k-1} d_w(c_p^{(j)}, O) + 2d_w(c, O) + 2 \sum_{j=1}^{k-1} d_w(c_q^{(j)}, O) + d_w(q, O) \\ &\leq d_w(p, O) + 2(k - 1)d_w(p, O) + 2d_w(c, O) + 2(k - 1)d_w(q, O) + d_w(q, O) \\ &\leq 2k(d_w(p, O) + d_w(q, O)) \\ &= 2kd_w(p, q). \end{aligned}$$

Thus the spanning ratio of the 1-dimensional spanner \mathcal{G}_λ is $2k$. □

3.3 Construction algorithm

In this section we propose an algorithm to construct the spanners of Section 3.2. The following gives a general overview of the algorithm, which computes a $2\sqrt{2}k$ -spanner in

$O(n \log^2 n + m \log n)$ time. In the rest of this section we will discuss the steps in more detail.

1. Preprocess P for efficient shortest path queries and build both the vertical decomposition \mathcal{VD} and horizontal decomposition \mathcal{HD} of P .
2. For each $p \in S$, find the trapezoid in \mathcal{VD} and \mathcal{HD} that contains p . For each trapezoid $\nabla \in \mathcal{VD}$, store the number of sites of S that lies in ∇ and sort these sites on their x -coordinates.
3. Recursively compute a spanner on the sites S in P :
 - (a) Find a vertical chord λ of P such that λ partitions P into two polygons P_ℓ and P_r , and each subpolygon contains at most $2n/3$ sites using the algorithm of Lemma 11.
 - (b) For each $p \in S$, find the point p_λ on λ and its weight using the algorithm of Lemma 12, and add this point to S_λ .
 - (c) Compute an additively weighted 1-dimensional spanner \mathcal{G}_λ on the set S_λ using the algorithm of Lemma 13 or Lemma 14.
 - (d) For every edge $(p_\lambda, q_\lambda) \in E_\lambda$ add the edge (p, q) to \mathcal{G} .
 - (e) Recursively compute spanners for S_ℓ in P_ℓ and S_r in P_r .

In step 1, we preprocess the polygon in $O(m)$ time such that the distance between any two points $p, q \in P$ can be computed in $O(\log m)$ time [15, 25]. We also build the horizontal and vertical decompositions of P , and a corresponding point location data structure, as a preprocessing step in $O(m)$ time [15, 30]. We then perform a point location query for each site $p \in S$ in $O(n \log m)$ time in step 2 and sort the sites within each trapezoid in $O(n \log n)$ time in total. The following lemma describes the algorithm to compute a vertical chord that partitions P into two subpolygons such that each of them contains roughly half of the sites in S . It is based on the algorithm of Bose et al. [7] that finds such a chord without the constraint that it should be vertical. Because of this constraint, we use the vertical decomposition of P instead of a triangulation in our algorithm.

Lemma 11. *In $O(n + m)$ time, we can find a vertical chord of P that partitions P into two subpolygons P_ℓ and P_r such that each subpolygon contains at most $2n/3$ sites of S .*

Proof. Consider the dual tree of the vertical decomposition \mathcal{VD} . Because our polygon vertices have distinct x - and y -coordinates, the maximum degree of any node in the tree is four; at most two neighbors to the right and two to the left of the trapezoid. We select an arbitrary node r as root of the tree. For each node v , we compute $c(v)$: the number of sites in S that lie in some trapezoid of the subtree rooted at v . These values can be computed in linear time (in the size of the polygon) using a bottom-up approach, because we already know the number of sites that lie in each trapezoid.

Let v be an arbitrary node in the tree and ∇_v the corresponding trapezoid. We first show that there is a vertical segment contained in ∇_v that partitions the polygon such that each subpolygon contains at most $2/3$ of the sites if




Figure 6: The trapezoid ∇_v has parent $\nabla_{p(v)}$ and relevant children ∇_u and ∇_w .

1. $n/3 \leq c(v) \leq 2n/3$, or
2. ∇_v contains at least $2n/3$ sites, or
3. for each child w of v we have $c(w) < n/3$, and $c(v) > 2n/3$.

In case 1, we choose λ as the vertical segment between the trapezoid of v and its parent. In case 2, we choose λ as a segment for which exactly $n/3$ sites in ∇_v lie left of λ . As ∇_v contains more than $2n/3$ sites, at most $n/3$ sites lie outside of ∇_v , thus P_ℓ contains between $n/3$ and $2n/3$ sites.

In case 3, we choose λ to lie within ∇_v . Note that v must have at least two children, otherwise Δ_v would contain at least $2n/3 - n/3 = n/3$ sites, thus we would be in case 1 or case 2. Assume that $\nabla_{p(v)}$ lies right of ∇_v , where $p(v)$ denotes the parent of v . We consider the two children u, w for which the trapezoids lie left of ∇_v , see Figure 6. When there is only one such site u , we consider $c(w) = 0$. It holds that $c(u) + c(w) < 2n/3$. We choose λ such that $\max(0, n/3 - c(u) - c(w))$ sites in ∇_v lie left of λ . Thus P_ℓ contains between $n/3$ and $2n/3$ sites of S . Similarly, we consider the children on the right when $\nabla_{p(v)}$ lies left of ∇_v .

We now show that if none of the above conditions hold, there is a path in the tree to a node for which one of these conditions holds. If none of the above conditions hold, then either $c(v) < n/3$, or $c(v) > 2n/3$ and there is a child w of v with $c(w) \geq n/3$. In the first case, we consider the node $p(v)$, in the second case we consider the node w with $c(w) \geq n/3$. By continuing like this, a path in the tree is formed. If the path ends up in r , it must hold that $c(r) > 2n/3$, and thus condition 2 holds. If the path ends up in a leaf node, then either condition 1 or 2 must hold for the leaf node.

This proves not only that there exists such a vertical segment, but also provides a way to find such a segment. As the tree contains $O(m)$ nodes, we can find a trapezoid that can contain λ in $O(m)$ time. We can separate the sites in a trapezoid by a vertical line segment such that exactly x sites lie left of the segment in linear time. Thus, the algorithm runs in $O(n + m)$ time. \square

The following lemma states that we can find the projections p_λ efficiently. The algorithm produces not only these projected sites, but also the shortest path tree SPT_λ of λ . The shortest path tree of a segment λ is defined similarly to that of a point: it is the




Figure 7: An example of the coloring of the horizontal decomposition in a simple polygon.

tree given by the union of all shortest paths from the vertices of P to their closest point on λ . In SPT_λ we also include all shortest paths from the sites in S to λ . See Figure 8 for an illustration.

Lemma 12. *We can compute the closest point p_λ on λ and $d(p, p_\lambda)$ for all sites $p \in S$, and the shortest path tree SPT_λ , in $O(m + n \log m)$ time.*

Proof. Consider the horizontal decomposition \mathcal{HD} of P and its dual tree T . Let x_λ denote the x -coordinate of the vertical line segment λ and let t be its top endpoint and b its bottom endpoint. We choose the root r of the tree as the trapezoid ∇_r for which $t \in \nabla_r$ and $\nabla_r \cap \lambda \neq \emptyset$. We color T as follows, see Figure 7 for an example. Every node v for which the corresponding trapezoid ∇_v contains the top or bottom endpoint of λ is colored orange. Note that the root r is thus colored orange. Every node v for which ∇_v is crossed from top to bottom by λ is colored blue. For every other node $v \in T$, consider its lowest ancestor w that is colored either blue or orange. If w is blue, then v is colored green, if w is orange, v is colored purple. Next, we describe how to find p_λ for a site $p = (x_p, y_p) \in \nabla_v$ for each color of v .

Blue: The horizontal line segment connecting p and λ is contained within P . So, $p_\lambda = (x_\lambda, y_p)$.

Orange: The shortest path to λ is again a line segment. If p lies above t or below b , then p_λ is t or b , respectively. Otherwise, $p_\lambda = (x_\lambda, y_p)$.

Green: Let w be the highest green ancestor of v (possibly v itself). Suppose that ∇_w lies above the trapezoid $\nabla_{p(w)}$ of its (blue) parent and ∇_w lies in P_ℓ . Let q be the bottom right corner of ∇_w . Then $q_\lambda = (x_\lambda, y_q)$, as q also lies in $\nabla_{p(w)}$. We will show that $p_\lambda = q_\lambda$. Suppose to the contrary that $p_\lambda \neq q_\lambda$. As v is a descendant of w , any path from p to λ must intersect the bottom of ∇_v . Let q' be the this intersection point. Because q' lies on the bottom boundary of ∇_w , the horizontal line segment $q'q_\lambda$ is contained within P and thus $q'_\lambda = q_\lambda$. Because $d(q', q'_\lambda) = d(q', q_\lambda) < d(q', p_\lambda)$, subpath optimality implies that

$d(p, q_\lambda) < d(p, p_\lambda)$, which is a contradiction. So $p_\lambda = q_\lambda$. Symmetrically, we can show a similar statement when ∇_v lies below $\nabla_{p(v)}$ and/or ∇_v lies in P_r .

Purple: Let w be the highest purple ancestor of v . Assume that $\nabla_{p(w)}$ lies below ∇_w . Note that the bottom segment of ∇_w lies above t . It follows that for all points q on this segment we have $q_\lambda = t$. According to the same argument as for the green trapezoids, we thus have $p_\lambda = t$. Symmetrically, if $\nabla_{p(w)}$ lies above ∇_w , then $p_\lambda = b$.

We can thus find p_λ for all $p \in S$ as follows. We perform a depth first search on T starting at r . When we visit a node v , we first determine its color in $O(1)$ time. Then for all $p \in \nabla_v$ we determine p_λ as described before. Note that this can also be done in constant time, because for a green/purple node we already computed the projected sites for the parent trapezoid. This thus takes $O(m + n)$ time overall.

The shortest path tree SPT_λ can now be computed as follows. All vertices of P in a blue trapezoid are children of λ . For each vertex q on the boundary of the blue and green region, see Figure 7, we include the shortest path tree of q restricted to its respective green region as a subtree of q . We include the shortest path trees of t and b restricted to the top and bottom orange and purple regions as a subtree of λ . Because the regions where we construct a shortest path trees are disjoint, we can compute all of these shortest path trees in $O(m)$ time [26]. Finally, for all $p \in S$ we compute $d(p, p_\lambda)$ in $O(n \log m)$ time, and include each site in SPT_λ as a child of their apex. \square

Lemma 13. *Given SPT_λ , we can construct a 4-spanner \mathcal{G}_λ on the additively weighted points S_λ , where the groups satisfy the properties of Lemma 5, in $O(n \log n + m)$ time.*

Proof. The 4-spanner of Section 3.2.1 requires an additional step at each level of the recursion, namely the formation of $\Theta(\sqrt{n})$ groups. We first discuss the running time to construct a 4-spanner when forming the groups as in Section 3.2.1, and then improve the running time by introducing a more efficient way to form the groups.

In Section 3.2.1, the groups are formed based on the shortest path tree of the site c . Building the shortest path tree, and a corresponding point location data structure, takes $O(m)$ time [26]. Then, we perform a point location query for each site to find its apex in the shortest path tree, and add the sites to the tree. These queries take $O(n \log m)$ time in total. We form groups based on the traversal of the tree. Note that we do not distinguish between sites with the same parent in the tree, as the tree \mathcal{T}_i (and thus the region R_i) obtained contains the same vertices of P regardless of the order of these sites. After obtaining the groups, we again add only $O(n)$ edges to the spanner. The overall running time of the algorithm is thus $O((m + n \log m) \log n)$.

This running time can be improved by using another approach to form the groups. To form groups that satisfy the properties of Lemma 5, and thus result in a spanner of the same complexity, we can use any partition of P into regions R_i , as long R_i as contains $\Theta(\sqrt{n})$ sites and $\sum_i r_i = O(m)$. Next, we describe how to form such groups efficiently using SPT_λ .

We first define an ordering on the sites. This is again based on the traversal of some shortest path tree. Instead of considering the shortest path tree of a point site, we consider the shortest path tree SPT_λ of λ . Again, all sites in S are included in this shortest path tree.




Figure 8: The shortest path tree SPT_λ and associated polygonal region R_i for each group S_i .

Additionally, we split the node corresponding to λ into a node for each distinct projection point on λ (of the vertices and the sites) and add an edge between each pair of adjacent points, see Figure 8. We root the tree at the node corresponding to the bottom endpoint of λ . Whenever a node t on λ has multiple children, in other words, when multiple sites are projected to the same point t , our assumption that all y -coordinates are distinct ensures that all these sites lie either in P_ℓ or P_r .

The groups are formed based on the in-order traversal of this tree, which can be performed in $O(m+n)$ time. As before, the first $\lceil \sqrt{n} \rceil$ are in S_1 , the second in S_2 , etc. The groups thus satisfy the first property. Next, we show they also satisfy the second property.

For each group S_i , we again consider the minimal subtree \mathcal{T}_i of SPT_λ containing all $p \in S_i$. \mathcal{T}_i defines a region R_i in P as follows. Let a be the first site of S_i in \mathcal{T}_i by the ordering used before. Assume that a lies in P_ℓ . We distinguish two cases: $a_\lambda \in \mathcal{T}_i$, or $a_\lambda \notin \mathcal{T}_i$. When $a_\lambda \in \mathcal{T}_i$, then let π_a be the path obtained from $\pi(a_\lambda, a)$ by extending the last segment to the boundary of P . Additionally, we extend π_a into P_r horizontally until we hit the polygon boundary. When $a_\lambda \notin \mathcal{T}_i$, consider the root v_i of \mathcal{T}_i . Let π_a be the path obtained from $\pi(v_i, a)$ by extending the last segment of the path to the boundary of P . Similarly, let π_b be such a path for the rightmost site of S_i in \mathcal{T}_i . We take R_i to be the region in P bounded by π_a , π_b , and some part of the boundary of P that contains the sites in S_i . See Figure 8. Note that, as before, only vertices of P that are in \mathcal{T}_i can occur in R_i . All shortest paths between sites in S_i are contained within R_i . Just as for the shortest path tree of c , Lemma 8 implies that any vertex $v \in SPT_\lambda$ occurs in at most two trees \mathcal{T}_i and \mathcal{T}_j as a non-root vertex. We conclude that any vertex is used by shortest paths within at most two groups.

After splitting λ at a point O , the tree SPT_λ is also split into two trees \mathcal{T}_ℓ and \mathcal{T}_r that contain exactly the sites in S_ℓ and S_r . We can thus reuse the ordering to form groups at each level of the recursion. This way, the total running time at a single level of the recursion is reduced to $O(n)$. The overall running time thus becomes $O(n \log n + m)$. \square

Lemma 14. *Given SPT_λ , we can construct a $2k$ -spanner \mathcal{G}_λ on the additively weighted points S_λ , where groups are formed as in Section 3.2.2, in $O(n \log n + m)$ time.*

Proof. To construct the 1-dimensional $2k$ -spanner of Section 3.2.2, we can use the shortest path tree of λ to form the groups as before. Note that we can select a center for each group after computing the groups, as including the center in the subgroups does not influence spanning ratio or complexity. After ordering the sites based on the in-order traversal of SPT_λ , we can build the tree of groups in linear time using a bottom up approach. As before, fix $N = n^{1/k}$. We first form the $\Theta(N^k)$ lowest level groups, containing only a single site, and select a center for each group. Each group at level i is created by merging $\Theta(N)$ groups at level $i - 1$, based on the same ordering. We do not perform this merging explicitly, but for each group we select the site closest to O of the merged level- $(i - 1)$ centers as the center. Because our center property, being the closest to O , is decomposable, this indeed gives us the center of the entire group. This way, we can compute the edges added in one level of the recursion in linear time, so the running time remains $O(n \log n + m)$. \square

The total running time thus becomes $O((n(\log n + \log m) + m) \log n) = O(n \log^2 n + m \log n)$. Here, we used that $n \log n \log m = O(n \log^2 n)$ for $m < n^2$, and $n \log n \log m = O(m \log n)$ for $m \geq n^2$. By splitting the polygon alternately based on the sites and the polygon vertices, we can replace the final $O(\log n)$ factor by $O(\min(\log n, \log m))$. Together with Lemma 10, we obtain the following theorem.

Theorem 2. *Let S be a set of n point sites in a simple polygon P with m vertices, and let $k \geq 1$ be any integer constant. We can build a geodesic $2\sqrt{2}k$ -spanner of size $O(n \log^2 n)$ and complexity $O(mn^{1/k} + n \log^2 n)$ in $O(n \log^2 n + m \log n + K)$ time, where K is the output complexity.*

4 Balanced shortest-path separators

Let S be a set of n point sites inside a polygonal domain P with m vertices. A straightforward approach to construct a geodesic spanner for a polygonal domain would be to use the same construction we used for a simple polygon in Section 3.1. In a polygonal domain, we cannot simply split the domain into two subdomains by a line segment, as we did to partition a simple polygon, because any line segment that appropriately partitions S might intersect one or more holes. However, to apply our 1-dimensional spanner, we require only that the splitting curve λ is a shortest path in P . Sadly, even if we allow a shortest path between two points on the outer boundary of P as our separator, it is not always possible to split the sites into sets S_ℓ and S_r of roughly equal size. See Figure 9 for an example. We thus need another approach for subdividing the domain.

In this section, we develop an algorithm to partition the polygonal domain P into two subdomains P_ℓ and P_r such that roughly half of the sites in S lie in P_ℓ and half in P_r . Additionally, we require that the curve bounding P_ℓ consists of at most three shortest paths and possibly part of ∂P , where ∂P denotes the boundary the polygonal domain, i.e. the boundary of the outer polygon and the boundary of the holes.

In the version of this paper published at SoCG 2023 [20], we adapted the balanced shortest-path separator of Abam, de Berg, and Seraji [3] for this purpose. However, a




Figure 9: No shortest path between two points on the boundary of P can separate the sites into two groups. The sites can be separated by three shortest paths, for example using the orange triangle.




Figure 10: A 1-separator (blue), a 2-separator (orange), and a 3-separator (green).

much more efficient construction can be achieved by using a different approach based on the separator by Thorup [40], which is the approach we take here.

To partition the polygonal domain into two subdomains P_ℓ and P_r , such that each contains roughly half of the sites in S , we allow three different types of separator. These separators consist of 1, 2, or 3 shortest paths, as seen in Figure 10. Formally, we define a balanced shortest-path separator as follows.

Definition 1. A *balanced shortest-path separator* (sp-separator) partitions the polygonal domain into two subdomains P_ℓ and P_r , such that $2n/9 \leq |S_\ell| \leq 2n/3$. The separator is of one the following three types.

- A *1-separator* consists of a single shortest path $\pi(u, v)$ that connects two points u, v on the outer boundary of P .
- A *2-separator* consists of two shortest paths $\pi(u, v)$ and $\pi(u, w)$, where $u \in P$ and $v, w \in \partial P$ on the boundary of the same hole or the outer polygon.
- A *3-separator* consists of three shortest paths $\pi(u, v)$, $\pi(v, w)$, and $\pi(w, u)$, where $u, v, w \in P$, and each pair of shortest paths overlaps in a single interval, starting at the common endpoint.

Note that one of the subpolygons P_ℓ or P_r might be disconnected for a 2- or 3-separator. We also allow a shortest path as a *degenerate* 3-separator, P_ℓ is then the degenerate polygon that is this shortest path. We call u, v, w the *corners* of the separator. Slightly

abusing our notation, we refer to the closed subset of \mathbb{R}^2 corresponding to the polygonal domain P_ℓ of a separator Δ by Δ itself. In the rest of this section, we prove the following theorem.

Theorem 3. *Let S be a set of n point sites in a polygonal domain P with m vertices. A balanced sp -separator exists and it can be computed in $O(m \log m + n \log n)$ time.*

Thorup [40] proposed a shortest path separator for a polygonal domain P consisting of at most six shortest paths (and possibly part of ∂P) that separates P into three parts that each contain at most half of the triangles of some triangulation of P . His idea is to build the shortest path tree of an arbitrary vertex, and then apply his earlier algorithm to find a separator for a planar graph using a spanning tree of the graph [39]. This approach is in turn based on the separator theorem for planar graphs by Lipton and Tarjan [33]. In our case, we want our separator not to separate the vertices of the polygon, but the sites in S . Next, we give a standalone explanation of how to handle this case, using the ideas of Thorup [40], and Lipton and Tarjan [33]. Our goal is to find a balanced separator consisting of at most three shortest paths. In other words, we want to find a 1-, 2-, or 3-separator.

We start by constructing the shortest path tree T of an arbitrary vertex v on the outer boundary of P . We then triangulate the space both inside and outside the convex hull of P , including the interior of the obstacles, constrained by the edges of ∂P and T . The triangulation within the convex hull of P can be done using straight line segments. However, we want to triangulate the space outside of the convex hull as well, but for this we may use curved edges, as illustrated in Figure 11. We call a triangle *free* when it lies in P , so its interior is not part of a hole or the exterior. We assign each site to the triangle that contains the site. If there are multiple sites in the interior of the boundary of two triangles, then all of these sites are assigned to the same free triangle (chosen arbitrarily). We set the weight of a triangle to the number of sites that is assigned to it. If there is a triangle of weight greater than $n/3$, then we can easily find a balanced separator that is contained within this triangle in $O(n \log n)$ time. When there are many sites on an edge of this triangle, our assignment ensures these are all assigned to the same triangle, which allows us to use a segment on this edge as a degenerate 3-separator. In the rest of this section, we thus assume that there is no such a heavy triangle.

We find a balanced separator by constructing a sequence $\Delta_0 \supset \Delta_1 \supset \dots \supset \Delta_k$ of separators such that $|\Delta_i \cap S| > 2n/3$ for $i = 0, \dots, k-1$ and $n/3 \leq |\Delta_k \cap S| \leq 2n/3$. Each separator Δ_i consists of an edge of the triangulation $(u_i, v_i) \notin T$ and the shortest paths $\pi(v, u_i)$ and $\pi(v, v_i)$. See Figure 11 for an illustration.

Let (u_0, v_0) be an arbitrary non-tree edge. If (u_0, v_0) is in the interior of a hole (or in the exterior), then the shortest paths $\pi(u_0, v)$ and $\pi(v_0, v)$, together with the obstacle boundary (or the boundary of the outer polygon), form a separator. To be precise, this is a 2-separator, unless u_0 or v_0 is equal to v then it is a 1-separator. Otherwise, $\pi(u_0, v)$, $\pi(v_0, v)$, and $\pi(u_0, v_0)$ together form a 3-separator. Let Δ_0 correspond to the side of the separator that has the highest weight. If the weight of Δ_0 is at most $2n/3$, we are done. Otherwise, we find the next separator Δ_1 as follows.

Consider the triangle (u_0, v_0, w) adjacent to (u_0, v_0) contained within Δ_0 . If one of




Figure 11: The shortest path tree T of v (blue edges) and a triangulation (purple edges inside P and grey edges outside P). The separator Δ_i (dashed orange) separates the white part of the polygon from the blue part.

the sides of this triangle is a tree edge, then let the non-tree edge be (u_1, v_1) . Otherwise, consider the separators defined by $(\pi(u_0, w), \pi(u_0, v), \pi(w, v))$ and $(\pi(v_0, w), \pi(v_0, v), \pi(w, v))$. Note that these separators are contained within Δ_0 , and their union is equal to Δ_0 . In this case, let (u_1, v_1) be the edge (u_0, w) or (v_0, w) corresponding to the separator with largest weight. As the triangle (u_0, v_0, w) has weight less than $n/3$, the weight of Δ_1 is at least $n/3$, and Δ_1 contains at least one less triangle than Δ_0 . We continue this process until we find a balanced separator Δ_k . Note that the separator thus always consists of two shortest paths to v from vertices of P and is closed by either a line segment between two vertices of P , or part of the boundary. As the weight of the separator corresponds to the number of sites in P_ℓ , this gives us a separator of the sites.

Building the shortest path tree of v and triangulating it takes $O(m \log m)$ time [15, 26, 30]. We can compute the weight of all triangles in $O(n \log m)$ time using $O(n)$ point location queries. Next, we argue that the search for the separator can then be performed in $O(m)$ time [33]. Computing the weight of Δ_0 can be done in $O(m)$, for example using a breadth-first search approach. The next separator Δ_{i+1} contains at least one triangle less than Δ_i , so it can take at most m steps to find a valid separator. In such a step there are two cases: either one side of the triangle (u_i, v_i, w) is a tree edge, or none of its sides are. In the first case we simply compute the weight of Δ_{i+1} by subtracting the weight of the triangle (u_i, v_i, w) from the weight of Δ_i . In the second case, we must find the weight of the separators defined by $(\pi(u_i, w), \pi(u_i, v), \pi(w, v))$ and $(\pi(v_i, w), \pi(v_i, v), \pi(w, v))$. As they partition Δ_i it is sufficient to compute the weight of just one of these separators. We aggregate the weight of the two separators simultaneously, by using a breadth-first search in both problems and alternately taking one step in each problem. We stop whenever we find the complete weight of one of the separators. The number of step in the BFS is thus

at most twice the size of the smallest of the two separators. As the triangles in one of these separators are not contained in any of the separators $\Delta_{i+1}, \dots, \Delta_k$ (but are in Δ_i), the running time over all separators $\Delta_0, \dots, \Delta_{k-1}$ is $O(m)$. We conclude that the total running time to compute a valid separator is $O(n \log n + m \log m + n \log m) = O(n \log n + m \log m)$.

5 Spanners in a polygonal domain

We consider a set of point sites S that lie in a polygonal domain P with m vertices and h holes. Let ∂P denote the boundary of the outer polygon. In Section 5.1, we first discuss how to obtain a simple geodesic spanner for a polygonal domain, using the separator of Section 4. As before, the complexity of this spanner can be high. In Section 5.2.1, we discuss an adaptation to the spanner construction that achieves lower-complexity spanners, where the edges in the spanner are no longer shortest paths.

5.1 A simple geodesic spanner

To construct a geodesic spanner for a polygonal domain we augment the construction we used for a simple polygon in Section 3.1. As discussed in Section 4, we cannot split the polygon into two subpolygons by a line segment λ . However, to apply our 1-dimensional spanner, we require only that the splitting curve λ is a shortest path in P . Instead of a line segment, we use the balanced sp-separator of Section 4 to split the polygonal domain. There are three types of such a separator: a shortest path between two points on ∂P (1-separator), two shortest paths starting at the same point and ending at the boundary of a single hole (2-separator), three shortest paths $\pi(u, v)$, $\pi(v, w)$, and $\pi(u, w)$ with $u, v, w \in P$ (3-separator). See Figure 10 for an illustration and Definition 1 for a formal definition. Let P_ℓ be the polygonal domain to the left of λ , when λ is a 1-separator, and interior to λ , when λ is a 2- or 3-separator. Symmetrically, P_r is the domain to the right of λ for a 1-separator and exterior to λ for a 2-, or 3-separator. As before, let S_ℓ be the sites in the closed region P_ℓ , and $S_r := S \setminus S_\ell$.

To compute a spanner on the set S , we project the sites to each of the shortest paths defining the separator, and consecutively run the 1-dimensional spanner algorithm to construct \mathcal{G}_λ once on each shortest path. Note that these projections are no longer unique, as there might be two topologically distinct shortest paths to λ . However, we can simply select one such that no two paths properly intersect to obtain the desired spanning ratio and spanner complexity. We then add the edge (p, q) to our spanner \mathcal{G} for each edge (p_λ, q_λ) in the 1-dimensional spanners. Finally, we recursively compute spanners for the sites S_ℓ in P_ℓ and S_r in P_r , just like in the simple polygon case.

Whenever the sp-separator intersects a single hole at two or more different intervals, then part of P_ℓ or P_r becomes disconnected. When this happens, we simply consider each connected polygonal domain as a separate subproblem, and recurse on all of them. Let n_i and m_i denote the number of sites and worst-case complexity of a shortest path in subproblem i . This means that only reflex vertices are counted for m_i , which are the only relevant vertices for the spanner complexity. The sites are partitioned over the subproblems, so we

have $\sum_i n_i = n$. The only new vertices (not of P) that can be included in the subproblems are the at most three corners of the separator. Each vertex can be a reflex vertex in only one of the subproblems, thus $\sum_i m_i \leq m + 3$. In the recursion, an increase in the number of subproblems means that we might have more than $c \cdot 2^i$ vertices not of P at level i , but the depth of the recursion tree is then proportionally decreased. All further proofs on complexity of our spanners are written in term of P_ℓ and P_r , but translate to the case of multiple subproblems.

Next, we analyze the spanner construction using any 1-dimensional additively weighted t -spanner of size $O(n \log n)$.

Lemma 15. *The graph \mathcal{G} is a geodesic $3t$ -spanner of size $O(n \log^2 n)$.*

Proof. As the 1-dimensional spanner has size $O(n \log n)$, there are still $O(n \log^2 n)$ edges in \mathcal{G} . What remains is to argue that \mathcal{G} is a $3t$ -spanner. Let p, q be two sites in S . In contrast to the simple polygon case, a shortest path between two sites p, q in S_ℓ (resp. S_r) is not necessarily contained in P_ℓ (resp. P_r). Therefore, we distinguish two different cases: either $\pi(p, q)$ is fully contained within P_r or P_ℓ , or there is a point $r \in \pi(p, q) \cap \lambda$ for some shortest path λ of the separator. In the first case, there exists a path in \mathcal{G} of length at most $3td(p, q)$ by induction. In the second case, we have

$$d_{\mathcal{G}}(p, q) \leq d_{\mathcal{G}_\lambda}(p, q) \leq td_w(p_\lambda, q_\lambda) = t(d(p, p_\lambda) + d(p_\lambda, q_\lambda) + d(q, q_\lambda)). \quad (4)$$

Additionally, we use that $d(p_\lambda, q_\lambda) \leq d(p_\lambda, r) + d(r, q_\lambda)$ and $d(p_\lambda, r) \leq d(p_\lambda, p) + d(p, r)$, because of the triangle inequality, so

$$d(p_\lambda, q_\lambda) \leq d(p_\lambda, r) + d(q_\lambda, r) \leq d(p_\lambda, p) + d(p, r) + d(q_\lambda, q) + d(q, r) \leq 2d(p, q).$$

It follows that $d_{\mathcal{G}}(p, q) \leq 3td(p, q)$. □

Applying the simple 1-dimensional spanner of Section 2 results in a 6-spanner. However, using the refinement of Lemma 3, we again obtain a $(2 + \varepsilon)$ -spanner.

5.2 Low complexity spanners in a polygonal domain

To obtain spanners of low complexity in a simple polygon, we formed groups of sites such that shortest paths within a group were disjoint from shortest paths of other groups. We proposed two different ways of forming these groups, based on the shortest path tree of the central site c , and based on the shortest path tree of the separator λ . Both of these approaches do not directly lead to a low complexity spanner in a polygonal domain, as we explain next.

Both methods can still be applied in a polygonal domain, as the shortest path tree of both a site and a shortest path is still well-defined. However, it does not give us the property that we want for our groups. In particular, the second property discussed in Lemma 13: each vertex of P is only used by shortest paths within $O(1)$ groups, does not hold. This is because the shortest path between two vertices u, v is not necessarily homotopic to the




Figure 12: Assigning the sites to groups based on the shortest path tree of c , as described in Section 3.2.1, forms these colored groups. The shortest path from each site to c is shown dashed. Each shortest path between two sites of a group contains both vertices v and w .

path $\pi(u, c) \cup \pi(c, v)$. Thus paths within a group can go around a certain hole, while their shortest paths to c (or λ) do not. See Figure 12 for an example. The construction can easily be expanded to ensure there are more sites in each group, by simply adding as many sites very close to the existing ones, or to more than three groups, by adding an additional hole above the construction with two corresponding sites. Consequently, the property that each vertex of P is used only by shortest paths within $O(1)$ groups does not hold.

So far, we assumed that every edge $(p, q) \in E$ is a shortest path between p and q . To obtain a spanner of low complexity, we can also allow an edge between p and q to be any path between the two sites. We call such a spanner a relaxed geodesic spanner. Note that our lower bounds still hold in this case. In Figure 4, every path between $p \in S_\ell$ and $q \in S_r$ has complexity $\Theta(m)$, and in Figure 18 we can easily adapt the top side of the polygon such that any path between two sites p, q has the same complexity as $\pi(p, q)$.

5.2.1 A 12-spanner of complexity $O(m\sqrt{n} + n \log^2 n)$

To obtain a low complexity spanner in a polygonal domain, we adapt our techniques for the simple-polygon $4\sqrt{2}$ -spanner in such a way that we avoid the problems we just sketched. The main difference with the simple polygon approach is that for an edge (p_λ, q_λ) in \mathcal{G}_λ , the edge (p, q) that we add to \mathcal{G} is no longer $\pi(p, q)$. Instead, let (p, q) be the shortest path from p to q via p_λ and q_λ , excluding any overlap of the path. We denote this path by $\pi_\lambda(p, q)$. This path is not unique, for example when $\pi(p, p_\lambda)$ is not unique, but choosing the path that is used to construct SPT_λ will do. Formally, $\pi_\lambda(p, q)$ is defined as follows.

Definition 2. The path $\pi_\lambda(p, q)$ is given by:

- $\pi(p, p_\lambda) \cup \pi(p_\lambda, q_\lambda) \cup \pi(q_\lambda, q)$, where $\pi(p_\lambda, q_\lambda) \subseteq \lambda$, if $\pi(p, p_\lambda)$ and $\pi(q, q_\lambda)$ are disjoint,
- $\pi(p, r) \cup \pi(r, q)$, where r denotes the closest point to p of $\pi(p, p_\lambda) \cap \pi(q, q_\lambda)$, otherwise.

One of the properties that we require of the groups, see Lemma 5, has changed, namely that each vertex of P appears on the shortest path between a pair of sites in the same group for only $O(1)$ groups. Instead of the shortest paths between sites in a group, we consider the paths $\pi_\lambda(p, q)$ of Definition 2. The following lemma shows that we can obtain a spanner with similar complexity as in a simple polygon when groups satisfy this adjusted property.

Lemma 16. *If the groups satisfy the following properties, then \mathcal{G} has $O(m\sqrt{n} + n \log^2 n)$ complexity:*

1. *each group contains $\Theta(\sqrt{n})$ sites, and*
2. *each vertex of P appears on the path π_λ between a pair of sites in the same group for only $O(1)$ groups.*

Proof. Note that the complexity of any path $\pi_\lambda(p, q)$ is $O(m)$, as it can use a vertex of P at most once. Thus the proof of Lemma 5 directly implies that the complexity of the edges in one level of \mathcal{G}_λ is $O(m\sqrt{n} + n)$.

In the recursion to construct \mathcal{G}_λ , splitting the sites by O no longer corresponds to a horizontal split in the polygon. However, the paths $\pi(p, p_\lambda)$, $p \in S_\ell$, are still disjoint from the paths $\pi(q, q_\lambda)$, $q \in S_r$. For the two subproblems generated by the split by O it thus still holds that $m_1 + m_2 = m$, where m_i denotes the maximum complexity of a path in subproblem i . Lemma 4 states that this recursion solves to $O(m\sqrt{n} + n \log n)$.

In the recursion where the domain is partitioned into two subpolygons P_ℓ and P_r , we now add at most three new vertices to the polygonal domain, namely the three corners of the sp-separator. Each vertex of P can only be a reflex vertex in either P_ℓ or P_r , so $m_1 + m_2 \leq m + 3$. Lemma 4 implies that this recursion for the complexity solves to $O(m\sqrt{n} + n \log^2 n)$. \square

As in Lemma 13, we form the groups based on the traversal of the shortest path tree SPT_λ . We again include all sites in S in the shortest path tree. Whenever a node v of λ has multiple children, we let all nodes that correspond to vertices/sites that lie to the left of λ come before nodes that correspond to vertices/sites that lie to the right of λ in the in-order traversal. Within these sets, the vertices/sites are ordered from bottom to top, as seen from v . See Figure 13 for an example. The subtree rooted at the start and end point of λ is simply a part of the shortest path tree of the start/end point. The first $\lceil \sqrt{n} \rceil$ sites in the in-order traversal are in S_1 , the second $\lceil \sqrt{n} \rceil$ in S_2 , etc.

Clearly, these groups satisfy property 1 of Lemma 16. To show these groups satisfy property 2, we again consider for each group S_i the minimal subtree \mathcal{T}_i of SPT_λ that contains all $p \in S_i$. Whenever there is more than one vertex of λ in \mathcal{T}_i , we choose the leftmost of these vertices as the root of \mathcal{T}_i .

Lemma 17. *An edge $\pi_\lambda(p, q)$ with $p_\lambda, q_\lambda \in S_i$ bends only at vertices in \mathcal{T}_i .*




Figure 13: The shortest path tree SPT_λ and the corresponding group assignment.

Proof. The path $\pi_\lambda(p, q)$ is a subpath of $\pi(p, p_\lambda) \cap \pi(p_\lambda, q_\lambda) \cap \pi(q, q_\lambda)$. In particular, it is the path from p to q in \mathcal{T}_i via their lowest common ancestor. This is p_λ or q_λ when $\pi(p, p_\lambda) \cap \pi(q, q_\lambda) = \emptyset$ and the vertex r , as in Definition 2, otherwise. \square

Lemma 18. *Any vertex in SPT_λ occurs in at most two trees \mathcal{T}_i and \mathcal{T}_j as a non-root node.*

Proof. Follows directly from the proof of Lemma 8. \square

Except for the root nodes, property 2 thus holds. As each \mathcal{T}_i has only one root node by definition, the number of groups using a single root node v in their paths may be large, but the sum of the number of groups that use a node over all root nodes is still $O(\sqrt{n})$. From this and Lemma 16 we conclude that \mathcal{G} has complexity $O(m\sqrt{n} + n \log^2 n)$.

Lemma 19. *The graph \mathcal{G} is a geodesic 12-spanner of size $O(n \log^2 n)$.*

Proof. The number of edges is exactly the same as in the $4\sqrt{2}$ -spanner in a simple polygon. The way the groups are formed in the construction of the 1-dimensional spanner \mathcal{G}_λ does not influence its spanning ratio, thus \mathcal{G}_λ is a 4-spanner (see Lemma 9). Note that even for our redefined edges in \mathcal{G} it holds that $d_{\mathcal{G}}(p, q) \leq d_{\mathcal{G}_\lambda}(p_\lambda, q_\lambda)$. Lemma 15 then directly implies that \mathcal{G} is a 12-spanner. \square

5.2.2 A $6k$ -spanner of complexity $O(mn^{1/k} + n \log^2 n)$

The generalization of the construction, discussed in Section 3.2.2, where $\Theta(n^{1/k})$ groups are recursively partitioned into smaller groups, is also applicable in a polygonal domain. As our groups satisfy the required properties, the complexity of this spanner remains $O(mn^{1/k} + n \log^2 n)$, as in the simple polygon, but the spanning ratio increases to $6k$.

Lemma 20. *Let S be a set of n point sites in a polygonal domain P with m vertices, and let $k \geq 1$ be any integer constant. There exists a geodesic $6k$ -spanner of size $O(n \log^2 n)$ and complexity $O(mn^{1/k} + n \log^2 n)$.*

5.3 Construction algorithm

In this section we discuss an algorithm to compute the geodesic spanners of Section 5.2. The following gives an overview of the algorithm that computes a $6k$ -spanner of complexity $O(mn^{1/k} + n \log^2 n)$ in $O(n^2 \log m + nm \log m)$ time.

1. Find an sp-separator such that P is partitioned into two polygons P_ℓ and P_r , and S_ℓ contains at least $2n/9$ and at most $2n/3$ sites using the algorithm of Theorem 3.
2. For each shortest path λ of the separator:
 - (a) For each $p \in S$ find the weighted point p_λ on λ and add this point to S_λ using the algorithm of Lemma 21.
 - (b) Compute an additively weighted 1-dimensional spanner \mathcal{G}_λ on the set S_λ .
 - (c) For every edge $(p_\lambda, q_\lambda) \in E_\lambda$ add the edge $(p, q) = \pi_\lambda(p, q)$ to \mathcal{G} .
3. Recursively compute spanners for S_ℓ in P_ℓ and S_r in P_r .

The algorithm starts by finding a balanced sp-separator. According to Theorem 3 this takes $O(n \log n + m \log m)$ time. We then continue by building a 1-dimensional additively weighted spanner on each of the shortest paths defining λ as follows.

Lemma 21. *We can compute the closest point p_λ on λ and $d(p, p_\lambda)$ for all sites $p \in S$, and the shortest path tree SPT_λ , in $O((m + n) \log m)$ time.*

Proof. Hershberger and Suri [29] show how to build the shortest path map of a point site in a polygonal domain in $O(m \log m)$ time. They also note that this extends to non-point sources, such as line segments, and to $O(m)$ sources, without increasing the running time. We can thus build the shortest path map of λ in P in $O(m \log m)$ time, using that λ has complexity $O(m)$. After essentially triangulating each region of the shortest path map using the vertices of P , we obtain SPT_λ in the same time bound. After building a point location data structure for this augmented shortest path map, we can query it for each site $p \in S$ to find p_λ and $d(p, p_\lambda)$ in $O(\log m)$ time. \square

Lemma 22. *Given SPT_λ , we can construct a 4-spanner \mathcal{G}_λ on the additively weighted points S_λ , where the groups satisfy the properties of Lemma 16, in $O(n \log n + m)$ time.*

Proof. As in Lemma 13, we can reuse SPT_λ to form the groups based on the ordering produced by an in-order traversal of the tree. See Section 5.2.1 for an exact description of this ordering. The ordering allows us to form the groups for a level of the 1-dimensional spanner in $O(n)$ time, thus the total running time is $O(n \log n)$. \square




Figure 14: The site p and the additional points $p_\lambda^{(i)}$ it generates for $i \in \{-3, \dots, 3\}$.

For the general $2k$ -spanner for additively weighted sites, using the ordering of the sites and Lemma 14 to build the tree of groups bottom up implies that we can construct this spanner in $O(n \log n + m)$ time as well.

After computing the additively weighted spanner \mathcal{G}_λ , we add edge $(p, q) = \pi_\lambda(p, q)$ to \mathcal{G} for every edge $(p_\lambda, q_\lambda) \in \mathcal{G}_\lambda$. We can either compute and store these edges explicitly, which would take time and space equal to the complexity of all added edges, or we can store them implicitly by only storing the points p_λ and q_λ , or the point r from Definition 2 when the paths are not disjoint. The point r is the lowest common ancestor of the nodes p and q in SPT_λ . This can be computed in $O(1)$ time after $O((n + m) \log(n + m))$ preprocessing time [5].

As there are at most three shortest path that define the separator, step 2 takes $O(n \log n + m \log m)$ time in total, as all other terms can be subsumed. This means that this is also the dominant term in the construction. The total running time is thus $O(n \log^2 n + m \log m \log n)$.

Theorem 4. *Let S be a set of n point sites in a polygonal domain P with m vertices, and let $k \geq 1$ be any integer constant. We can build a relaxed geodesic $6k$ -spanner of size $O(n \log^2 n)$ and complexity $O(mn^{1/k} + n \log^2 n)$ in $O(n \log^2 n + m \log m \log n + K)$ time, where K is the output complexity.*

6 Improving the spanning ratio in a simple polygon

In this section, we discuss how to improve the spanning ratio of the low complexity spanner in a simple polygon from $2\sqrt{2}k$ to $2k + \varepsilon$, for any constant $\varepsilon \in (0, 2k)$. Observe that this is indeed improves the spanning ratio when $\varepsilon \in (0, (2\sqrt{2} - 2)k)$. To apply the refinement of Lemma 3 in combination with our low complexity spanner we require some additional properties on the collection of points added in the refinement. If we apply the lemma directly,

and thus add $O(1/\delta^2)$ additional points to the 1-dimensional space for each site in S , where $\delta = \varepsilon/(2k)$, then the recursion in the 1-dimensional spanner \mathcal{G}_λ no longer corresponds to a horizontal split in the polygon. In other words, the proof of Lemma 5 no longer holds. To fix this issue we make three changes to this naive application. First, we use the paths as defined in Definition 2 instead of shortest paths as edges in our spanner. This allows us to construct a low complexity spanner on any tree, instead of requiring the tree to be the shortest path tree of λ . Second, we do not choose the additional points on λ uniformly, but in such a way that the angle of the final segment on the path with λ is in a fixed set of angles (of size dependent on δ). Third, instead of constructing a $2k$ -spanner on a single 1-dimensional space, we consider a constant number (dependent on δ) of subsets of points, and construct a spanner on each set separately. Next, we elaborate on the second and third change.

Selecting additional points. For a site $p \in S$, consider the closest point p_λ on λ as before. Let x be a point on λ . As P is a simple polygon, the absolute value of the angle of the final segment on the path $\pi(p, x)$ (with respect to the x -axis) strictly increases as x moves away from p_λ . This implies that for any $\alpha \in (-\pi, \pi)$ there is at most one point p_α on λ such that the final segment of $\pi(p, p_\alpha)$ has angle α . Note that for $\alpha = 0$ the point p_α is equal to p_λ if p_λ is not one of the endpoints of λ . We use this property to select additional points corresponding to the same angles for each site.

As in Lemma 3, we select $O(n/\delta^2)$ points on λ in total, where $\delta = \varepsilon/(2k)$. Instead of adding these points directly to S_λ , we create $O(1/\delta^2)$ sets $S_\lambda^{(i)}$ of size $O(n)$. The new points are again weighted by the distance to their original site. We define $S_\lambda^{(0)} := \{p_\lambda : p \in S\}$. Next, we define the other sets $S_\lambda^{(i)}$.

For ease of description we do not consider the angle of the final segment in radians, but measure the angle differently. Let z be a vertex of P and let z' be the horizontal projection of z on the line through λ (observe that z' might not be on λ). Furthermore, let z^* be a point on λ visible from z . See Figure 14. We then describe the angle of the segment zz^* by the ratio between $|zz'|$ and $|z'z^*|$. Let $p_\lambda^{(i)}$ be the point on λ for which the last segment on $\pi(p, p_\lambda^{(i)})$ has ratio $1/(\delta|i|)$ and lies above p_λ if $i > 0$ or below p_λ if $i < 0$, if it exists. The set $S_\lambda^{(i)}$ is then defined as $S_\lambda^{(i)} := \{p_\lambda^{(i)} : p \in S\}$. We will only consider the sets $S_\lambda^{(i)}$ for $i \in I$, where $I = \{-\lceil 3/\delta^2 \rceil, -\lceil 3/\delta^2 \rceil + 1, \dots, \lceil 3/\delta^2 \rceil - 1, \lceil 3/\delta^2 \rceil\}$. Let S_λ be the union of these sets $S_\lambda^{(i)}$.

To compute the sets $S_\lambda^{(i)}$ we use the approach from Lemma 12 in both P_ℓ and P_r , which computes the set $S_\lambda^{(0)}$. For each set $S_\lambda^{(i)}$, instead of using the horizontal decomposition, we use a decomposition that consists of trapezoids whose sides are parallel to the angle of the last segment of the paths in the set $S_\lambda^{(i)}$. Such a decomposition can for example be computed simply by rotating P and S accordingly, and then computing a horizontal decomposition. We then use the coloring of Lemma 12 to find the site $p_\lambda^{(i)}$ and its distance to p for each $p \in S$. These additional points and the trees that they form can thus be computed in $O(1/\delta^2 \cdot (m + n \log m))$ time.




Figure 15: For each site the path to its projection and the (valid) additional points for $i \in \{-2, -1, 0, 1, 2\}$ are shown. The fat grey tree is the concatenation of the $i = -1$ paths in P_ℓ , the $i = 2$ paths in P_r , and λ .

Constructing a spanner. To obtain a $(2k + \varepsilon)$ -spanner in P , we need that for two sites $p \in S_\ell$ and $q \in S_r$ there are points $p_\lambda^{(i)}, q_\lambda^{(j)} \in S_\lambda$ such that $d_{\mathcal{G}_\lambda}(p_\lambda^{(i)}, q_\lambda^{(j)}) \leq 2k \cdot d_w(p_\lambda^{(i)}, q_\lambda^{(j)})$ and that $d_w(p_\lambda^{(i)}, q_\lambda^{(j)}) \leq (1 + \delta) \cdot d(p, q)$. We do not require these properties for all pairs of sites in S , but only for pairs that lie on opposite sides of λ . This means that we do not need to build a spanner on the entire set S_λ . Let $S_{\lambda,\ell}^{(i)}$ and $S_{\lambda,r}^{(i)}$ denote the subsets of $S_\lambda^{(i)}$ that contain all points generated by sites in S_ℓ and S_r respectively. We build a 1-dimensional $2k$ -spanner $\mathcal{G}_{i,j}$ for each set $S_{i,j} := S_{\lambda,\ell}^{(i)} \cup S_{\lambda,r}^{(j)}$, $i, j \in I$. We thus construct $O(1/\delta^4)$ spanners on sets of size $O(n)$. The edges in \mathcal{G}_λ are then simply the union of the edges in all spanners $\mathcal{G}_{i,j}$.

To construct a spanner $\mathcal{G}_{i,j}$, we apply our algorithm to construct a low complexity spanner on a tree $T_{i,j}$ instead of SPT_λ . The tree $T_{i,j}$ consists of λ together with the paths from each point $p_\lambda^{(i)} \in S_{i,j}$ (or $p_\lambda^{(j)}$) to its original site p . The following lemma implies that this indeed results in a tree.

Lemma 23. *For two sites $p, q \in S$ the paths $\pi(p, p_\lambda^{(i)})$ and $\pi(q, q_\lambda^{(i)})$ do not properly intersect.*

Proof. For contradiction, let x be the first point on $\pi(p, p_\lambda^{(i)}) \cap \pi(q, q_\lambda^{(i)})$ as seen from p . We know that x defines a unique point $x_\lambda^{(i)}$ on λ . As x is on $\pi(p, p_\lambda^{(i)})$ the subpath from x to $p_\lambda^{(i)}$ is also a shortest path. It follows that $x_\lambda^{(i)} = p_\lambda^{(i)}$. Similarly, $x_\lambda^{(i)} = q_\lambda^{(i)}$. We conclude that $p_\lambda^{(i)} = q_\lambda^{(i)}$, and thus the paths from p and q overlap from x on. \square

As in the polygonal domain setting, it follows that the complexity of the edges of $\mathcal{G}_{i,j}$ in P is $O(m\sqrt{n} + n \log^2 n)$. The complexity of all edges in \mathcal{G}_λ in P is thus $O(1/\delta^4 \cdot (m\sqrt{n} + n \log^2 n))$.

What remains is to prove this indeed results in the desired spanning ratio for sites on opposing sides of λ . Without loss of generality, let $p \in S_\ell$ and $q \in S_r$. We want to show




Figure 16: Notation used in Lemma 24.

that there are points $p_\lambda^{(i)}, q_\lambda^{(j)} \in S_\lambda$ such that $d_w(p_\lambda^{(i)}, q_\lambda^{(j)}) \leq (1 + \delta) \cdot d(p, q)$. Let r be the intersection point of $\pi(p, q)$ and λ . The following lemma states that there is a short path in \mathcal{G}_λ between p and r .

Lemma 24. *Let $r \in \lambda$ and $p \in S$. Then there is a point $p_\lambda^{(i)} \in S_\lambda$ such that $d(p, p_\lambda^{(i)}) + d(p_\lambda^{(i)}, r) \leq (1 + \delta)d(p, r)$.*

Proof. Without loss of generality, we assume that r is below p_λ on λ . Let $p^* \in S_\lambda$ be the point generated by p on $\pi(r, p_\lambda)$ that is closest to r . Furthermore, let z be the last vertex on $\pi(p, r) \cap \pi(p, p^*)$, and let z' be the horizontal projection of z on the line through λ . See Figure 16 for an illustration of the notation. Observe that when z is not the last vertex on $\pi(p, p^*)$ then the last vertex on this path must be above $\pi(p, r)$. We consider the following two cases: either $d(r, p^*) > \delta|zz'|$, or $d(r, p^*) \leq \delta|zz'|$.

If $d(r, p^*) > \delta|zz'|$, we claim that $p^* = p_\lambda^{(i')}$ where $i' = -\lceil 3/\delta^2 \rceil$, i.e. p^* is the outermost point generated by p . Consider the point z^* at distance $|zz'| \cdot \delta|i'|$ below z' . If $p^* = z^*$, then $d(r, p^*) > \delta|zz'|$ implies that the point $p_\lambda^{(i'-1)}$ should lie between p^* and r , which is a contradiction. If $p^* \neq z^*$, then p^* must lie below z^* , as it must be generated by a vertex below the segment zz^* . It follows that $d(r, z^*) > d(r, p^*) > \delta|zz'|$. This again implies the contradiction that $p_\lambda^{(i'-1)}$ is between p^* and r . We conclude that the claim holds. It follows that $|z'r| \geq |z'z^*| = |zz'| \cdot \delta \lceil 3/\delta^2 \rceil \geq |zz'| \cdot (1/\delta + 2/\delta) \geq |zz'| \cdot (1 + 2/\delta)$, as $\delta \in (0, 1)$.

Rewriting this equation gives $|zz'| \leq \delta/2 \cdot (|z'r| - |zz'|)$. It then follows that

$$\begin{aligned}
 d(p, p^*) + d(p^*, r) &= d(p, z) + d(z, p^*) + d(p^*, r) \\
 &\leq d(p, z) + |zz'| + |z'p^*| + d(p^*, r) \\
 &= d(p, z) + |zz'| + |z'r| \\
 &\leq d(p, z) + |zz'| + |z'z| + |zr| \\
 &= d(p, z) + 2|zz'| + |zr| \\
 &\leq d(p, z) + \delta \cdot (|z'r| - |zz'|) + |zr| \\
 &\leq d(p, z) + \delta \cdot (|zz'| + |zr| - |zz'|) + |zr| \\
 &\leq d(p, z) + (1 + \delta)|zr| \\
 &\leq (1 + \delta)d(p, r).
 \end{aligned}$$

If $d(r, p^*) \leq \delta|zz'|$, then

$$d(p, p^*) + d(p^*, r) \leq d(p, z) + d(z, p^*) + \delta|zz'| \leq d(p, z) + d(z, r) + \delta d(z, r) \leq (1 + \delta)d(p, r).$$

Here, we use that $d(z, p^*) \leq d(z, r)$, as p^* is closer to p_λ than r , and that $|zz'| \leq d(z, r)$. \square

Lemma 24 also implies that there is a point $q_\lambda^{(j)} \in S_\lambda$ such that $d(q, q_\lambda^{(j)}) + d(q_\lambda^{(j)}, r) \leq (1 + \delta)d(q, r)$. It follows that $d_w(p_\lambda^{(i)}, q_\lambda^{(j)}) \leq d(p, p_\lambda^{(i)}) + d(p_\lambda^{(i)}, r) + d(q, q_\lambda^{(j)}) + d(q_\lambda^{(j)}, r) \leq (1 + \delta)(d(p, r) + d(q, r)) = (1 + \delta)d(p, r)$. As $\mathcal{G}_{i,j}$ is a $2k$ -spanner on $S_{i,j}$, this results in a spanning ratio of $2k(1 + \delta) = 2k + \varepsilon$ for \mathcal{G} when choosing $\delta = \varepsilon/(2k)$.

Combining this result with the spanner construction in Section 3.3, we obtain the following theorem.

Theorem 5. *Let S be a set of n point sites in a simple polygon P with m vertices, and let $k \geq 1$ be any integer constant. For any constant $\varepsilon \in (0, 2k)$, we can build a relaxed geodesic $(2k + \varepsilon)$ -spanner of size $O(c_{\varepsilon,k} n \log^2 n)$ and complexity $O(c_{\varepsilon,k}(mn^{1/k} + n \log^2 n))$ in $O(c_{\varepsilon,k} n \log^2 n + m \log n + K)$ time, where $c_{\varepsilon,k}$ is a constant depending only on ε and k , and K is the output complexity.*

7 Lower bounds for complexity

In this section, we consider lower bounds on the complexity of spanners. We first describe a simple $\Omega(nm)$ lower bound construction for a (relaxed) geodesic $(3 - \varepsilon)$ -spanner, and then prove a (slightly worse) $\Omega(mn^{1/(t-1)})$ lower bound construction for a $(t - \varepsilon)$ -spanner.

7.1 Lower bound for $(3 - \varepsilon)$ -spanners

Theorem 6. *For any constant $\varepsilon \in (0, 1)$, there exists a set of n point sites in a simple polygon P with $m = \Omega(n)$ vertices for which any (relaxed) geodesic $(3 - \varepsilon)$ -spanner has complexity $\Omega(mn)$.*




Figure 17: A simple polygon with a set of sites for which any $(3 - \varepsilon)$ -spanner has complexity $\Omega(nm)$.

Proof. Consider the construction given in Figure 17. We assume that $m = \Omega(n)$. We split the sites into two sets S_ℓ and S_r equally. The sites lie in long ‘spikes’ of length ℓ , either on the left (S_ℓ) or right (S_r) of a central passage of complexity $\Theta(m)$. We show that this construction gives a complexity of $\Omega(nm)$ for any $(3 - \varepsilon)$ -spanner.

When w gets close to 0, the distance between any two sites p, q approaches 2ℓ . To get a $(3 - \varepsilon)$ -spanner, we can thus have at most one intermediate site on the path from p to q . We assume that all possible, constant complexity, edges between vertices on the same side of the construction are present in the spanner. To make sure sites on the left also have (short) paths to sites on the right, we have to add some additional edges that go through the central passage of the polygon, which forces paths to have complexity $\Theta(m)$. We will show that we need $\Theta(n)$ of these edges, each of complexity $\Theta(m)$, to achieve a $(3 - \varepsilon)$ -spanner, thus proving the $\Omega(nm)$ lower bound.

Let $q \in S_r$. For each $p \in S_\ell$, we need a path with at most one intermediate site to q . There are two ways to achieve this: we can go through an intermediate site on the left, or on the right. In the first case, we add an edge from a site $p' \in S_\ell$ to q . In the second case, we need to add an edge from each $p \in S_\ell$ to any site $q' \in S_r$. In the first case we thus add only one edge, while in the second case we add $\Theta(n)$ edges (that go through the central passage). Let $k \geq 1$ be the number of sites in S_r that have a direct edge to some site in S_ℓ . If $k < |S_r|$, then there is some site $q \in S_r$ for which we are in case two, and we thus have $\Theta(n)$ edges of complexity $\Theta(m)$. If $k = |S_r|$, then there is a direct edge to each of the $\Theta(n)$ sites in S_r , and we therefore also end up with a complexity of $\Omega(nm)$. \square

7.2 General lower bounds

Theorem 1. For any constant $\varepsilon \in (0, 1)$ and integer constant $t \geq 2$, there exists a set of n point sites in a simple polygon P with $m = \Omega(n)$ vertices for which any (relaxed) geodesic $(t - \varepsilon)$ -spanner has complexity $\Omega(mn^{1/(t-1)})$.

Proof. Consider the construction of the polygon P shown in Figure 18. Let p_1, \dots, p_n be the




Figure 18: Construction of a simple polygon P for which any $(t - \varepsilon)$ -spanner has complexity $\Omega(n^{1+1/(t-1)})$. The shortest path from the i -th to the j -th site has complexity $\Theta(i - j)$.

sites from left to right. Thus, the complexity of any path from the i -th site p_i to the j -th site p_j is at least $|i - j|$. When w is close to 0, the distance between any two sites approaches 2ℓ . To achieve a spanning ratio of $(t - \varepsilon)$, the path in the spanner from p_i to p_j can visit at most $t - 2$ other vertices. In other words, we can go from p_i to p_j in at most $t - 1$ hops. This is also called the hop-diameter of the spanner.

As the spanning ratio is determined only by the number of hops on the path, we can model the spanner in a much simpler metric space ϑ_n . This is a 1-dimensional Euclidean space with n points v_1, \dots, v_n that lie on the x -axis at coordinates $1, 2, \dots, n$. The edge (v_i, v_j) thus has length (or weight) $|i - j|$. Any spanning subgraph of ϑ_n of hop-diameter t and total weight W (the weight of a graph is the sum of the weights of its edges) is in one-to-one correspondence to an $(t + 1 - \varepsilon)$ -spanner of P of complexity $\Theta(W)$. Dinitz, Elkin, and Solomon [21] prove the following on the relation between the *hop-radius* and weight of any spanning subgraph of ϑ_n . The hop-radius $h(G, r)$ of a graph G with respect to a root r is defined as maximum number of hops that is needed to reach any vertex in the graph from the root. The hop-radius $h(G)$ of G is then defined as $\min_{r \in V} h(G, r)$. Note that the hop-diameter is an upper bound on $h(G)$.

Lemma 25 (Dinitz et al. [21]). *For any sufficiently large integer n and positive integer $t < \log n$, any spanning subgraph of ϑ_n with hop-radius at most t has weight at least $\Omega(t \cdot n^{1+1/t})$.*

The lemma implies that any $(t - \varepsilon)$ -spanner of P , which has hop-diameter $t - 1$, has complexity $\Omega((t - 1) \cdot n^{1+1/(t-1)}) = \Omega(n^{1+1/(t-1)})$, for constant t .

To achieve a lower bound for $m > n$, we slightly adapt our polygon such that the path between two adjacent spikes has complexity $\Theta(m/n)$. This implies that the path between p_i and p_j has a complexity of at least $m/n \cdot |i - j|$ instead. For this adapted polygon P' , any spanning subgraph of ϑ_n of hop-diameter t and total weight W is in one-to-one correspondence to an $(t + 1 - \varepsilon)$ -spanner of P' of complexity $\Theta(m/n \cdot W)$. It follows from Lemma 25 that any $(t - \varepsilon)$ -spanner of P' has complexity $\Omega(m/n \cdot n^{1+1/(t-1)}) = \Omega(mn^{1/(t-1)})$. \square

Corollary 1. For any integer constant $t \geq 2$, there exists a set of n point sites in a simple polygon P with $m = \Omega(n)$ vertices for which any (relaxed) geodesic t -spanner has complexity $\Omega(mn^{1/t})$.

8 Conclusion

In this paper, we have taken a first look at a new measure of compactness of geodesic spanners, namely the spanner complexity. We presented both constructions for low complexity spanners, as well as lower bounds on the complexity for a given spanning ratio. A clear direction for future work lies in closing the gap between the upper and lower bounds. This gap is much larger for a polygonal domain, as our spanner of complexity $O(mn^{1/k} + n \log^2 n)$ has a spanning ratio of $6k$, while the simple polygon spanner of similar complexity has a spanning ratio of only $2k + \varepsilon$. This difference lies in the application of the refinement suggested by Abam, de Berg, and Seraji [3]. For a simple polygon, we managed to adapt their approach to reduce the spanning ratio such that the complexity of the spanner is increased by only a constant factor. However, it does not seem straightforward to apply the same refinement in a polygonal domain while retaining the low complexity property.

For a polygonal domain, Abam et al. [1] suggest a different construction from our approach that results in a spanner of size dependent on the number of holes in P . Our techniques to construct a low complexity spanner can also be applied to this construction. This results in a $6k$ -spanner of size $O(\sqrt{hn} \log^2 n)$ and complexity $O(\sqrt{h}(mn^{1/k} + n \log^2 n))$. As the separators used here are always line segments instead of general shortest paths, it would be interesting to see if it is possibly easier to apply the refinement of [3] in this setting.

References

- [1] Mohammad Ali Abam, Marjan Adeli, Hamid Homapour, and Pooya Zafar Asadollahpoor. Geometric spanners for points inside a polygonal domain. In *31st International Symposium on Computational Geometry, SoCG*, volume 34 of *LIPICs*, pages 186–197, 2015.
- [2] Mohammad Ali Abam, Mark de Berg, Mohammad Farshi, Joachim Gudmundsson, and Michiel H. M. Smid. Geometric spanners for weighted point sets. *Algorithmica*, 61(1):207–225, 2011.
- [3] Mohammad Ali Abam, Mark de Berg, and Mohammad Javad Rezaei Seraji. Geodesic spanners for points on a polyhedral terrain. *SIAM J. Comput.*, 48(6):1796–1810, 2019.
- [4] Sunil Arya, Gautam Das, David M. Mount, Jeffrey S. Salowe, and Michiel H. M. Smid. Euclidean spanners: short, thin, and lanky. In *In Twenty-Seventh Annual ACM Symposium on Theory of Computing, STOC, Proceedings*, pages 489–498. ACM, 1995.
- [5] Michael A. Bender and Martin Farach-Colton. The LCA problem revisited. In *In Theoretical Informatics, 4th Latin American Symposium, LATIN, Proceedings*, volume 1776 of *Lecture Notes in Computer Science*, pages 88–94. Springer, 2000.

- [6] Sujoy Bore and Csaba D. Tóth. Euclidean Steiner spanners: Light and sparse. *SIAM Journal on Discrete Mathematics*, 36(3):2411–2444, 2022.
- [7] Prosenjit Bose, Jurek Czyzowicz, Evangelos Kranakis, Danny Krizanc, and Anil Maheshwari. Polygon cutting: Revisited. In *Discrete and Computational Geometry, Japanese Conference, JCDCG, Revised Papers*, volume 1763 of *LNCS*, pages 81–92, 1998.
- [8] Prosenjit Bose, Rolf Fagerberg, André van Renssen, and Sander Verdonschot. On plane constrained bounded-degree spanners. *Algorithmica*, 81(4):1392–1415, 2019.
- [9] Prosenjit Bose, Joachim Gudmundsson, and Michiel H. M. Smid. Constructing plane spanners of bounded degree and low weight. *Algorithmica*, 42(3-4):249–264, 2005.
- [10] Prosenjit Bose and J. Mark Keil. On the stretch factor of the constrained Delaunay triangulation. In *In the 3rd International Symposium on Voronoi Diagrams in Science and Engineering, ISVD*, pages 25–31. IEEE Computer Society, 2006.
- [11] Prosenjit Bose and Michiel H. M. Smid. On plane geometric spanners: A survey and open problems. *Comput. Geom.*, 46(7):818–830, 2013.
- [12] Prosenjit Bose and André van Renssen. Spanning properties of Yao and Θ -graphs in the presence of constraints. *Int. J. Comput. Geom. Appl.*, 29(2):95–120, 2019.
- [13] T.-H. Hubert Chan, Anupam Gupta, Bruce M. Maggs, and Shuheng Zhou. On hierarchical routing in doubling metrics. *ACM Trans. Algorithms*, 12(4):55:1–55:22, 2016.
- [14] T.-H. Hubert Chan, Mingfei Li, Li Ning, and Shay Solomon. New doubling spanners: Better and simpler. *SIAM J. Comput.*, 44(1):37–53, 2015.
- [15] Bernard Chazelle. Triangulating a simple polygon in linear time. *Discret. Comput. Geom.*, 6:485–524, 1991.
- [16] Kenneth L. Clarkson. Approximation algorithms for shortest path motion planning. In *In the 19th Annual ACM Symposium on Theory of Computing, STOC*, pages 56–65. ACM, 1987.
- [17] Gautam Das. The visibility graph contains a bounded-degree spanner. In *In the 9th Canadian Conference on Computational Geometry, CCCG*, 1997.
- [18] Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. *Computational geometry: Algorithms and applications, 3rd Edition*. Springer, 2008.
- [19] Sarita de Berg, Tim Ophelders, Irene Parada, Frank Staals, and Jules Wulms. The complexity of geodesic spanners using steiner points, 2024.
- [20] Sarita de Berg, Marc J. van Kreveld, and Frank Staals. The complexity of geodesic spanners. In *proc. 39th International Symposium on Computational Geometry, SoCG*, volume 258 of *LIPICs*, pages 16:1–16:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

- [21] Yefim Dinitz, Michael Elkin, and Shay Solomon. Shallow-low-light trees, and tight lower bounds for Euclidean spanners. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 519–528, 2008.
- [22] Michael Elkin and Shay Solomon. Optimal Euclidean spanners: Really short, thin, and lanky. *J. ACM*, 62(5):35:1–35:45, 2015.
- [23] Subir Kumar Ghosh. Computing the visibility polygon from a convex set and related problems. *J. Algorithms*, 12(1):75–95, 1991.
- [24] Lee-Ad Gottlieb and Liam Roditty. An optimal dynamic spanner for doubling metric spaces. In *16th Annual European Symposium on Algorithms, ESA*, volume 5193 of *Lecture Notes in Computer Science*, pages 478–489, 2008.
- [25] Leonidas J. Guibas and John Hershberger. Optimal shortest path queries in a simple polygon. *J. Comput. Syst. Sci.*, 39(2):126–152, 1989.
- [26] Leonidas J. Guibas, John Hershberger, Daniel Leven, Micha Sharir, and Robert Endre Tarjan. Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica*, 2:209–233, 1987.
- [27] Sariel Har-Peled and Manor Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM J. Comput.*, 35(5):1148–1184, 2006.
- [28] John Hershberger and Jack Snoeyink. Computing minimum length paths of a given homotopy class. *Comput. Geom.*, 4:63–97, 1994.
- [29] John Hershberger and Subhash Suri. An optimal algorithm for Euclidean shortest paths in the plane. *SIAM J. Comput.*, 28(6):2215–2256, 1999.
- [30] David G. Kirkpatrick. Optimal search in planar subdivisions. *SIAM J. Comput.*, 12(1):28–35, 1983.
- [31] Hung Le and Shay Solomon. Truly optimal Euclidean spanners. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 1078–1100. IEEE Computer Society, 2019.
- [32] Christos Levkopoulos, Giri Narasimhan, and Michiel H. M. Smid. Efficient algorithms for constructing fault-tolerant geometric spanners. In *In the 13th Annual ACM Symposium on the Theory of Computing, STOC, Proceedings*, pages 186–195. ACM, 1998.
- [33] Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.
- [34] Joseph S. B. Mitchell and Wolfgang Mulzer. Proximity algorithms. In *Handbook of Discrete and Computational Geometry (3rd Edition)*, chapter 32, pages 849–874. Chapman & Hall/CRC, 2017.
- [35] Joseph S. B. Mitchell, Christine D. Piatko, and Esther M. Arkin. Computing a shortest k-link path in a polygon. In *33rd Annual Symposium on Foundations of Computer Science, FOCS*, pages 573–582. IEEE Computer Society, 1992.

- [36] Joseph S. B. Mitchell, Günter Rote, and Gerhard J. Woeginger. Minimum-link paths among obstacles in the plane. *Algorithmica*, 8(5&6):431–459, 1992.
- [37] Giri Narasimhan and Michiel H. M. Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.
- [38] Subhash Suri. A linear time algorithm for minimum link paths inside a simple polygon. *Computer Vision, Graphics, and Image Processing*, 35(1):99–110, 1986.
- [39] Mikkel Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *J. ACM*, 51(6):993–1024, 2004.
- [40] Mikkel Thorup. Compact oracles for approximate distances around obstacles in the plane. In *In the 15th Annual European Symposium, ESA, Proceedings*, volume 4698 of *Lecture Notes in Computer Science*, pages 383–394. Springer, 2007.
- [41] Haitao Wang. A new algorithm for Euclidean shortest paths in the plane. In *53rd Annual ACM SIGACT Symposium on Theory of Computing, STOC*, pages 975–988. ACM, 2021.

A Alternative proof for the existence of a balanced separator

In an earlier version of this paper [20], we gave a constructive proof for the existence of a balanced sp-separator using a similar approach to Abam, de Berg, and Seraji [3]. As we consider a polygonal domain instead of a polyhedral terrain, we need to adapt the approach slightly to fit this setting. Moreover, we correct a technical issue of [3]. In this section we first give our proof for the existence of a separator in a polygonal domain, and then comment how to fix the technical issue in [3] using our approach.

Existence of a separator. Let P be a polygonal domain with h holes. We start by trying to find a 1-separator from an arbitrary fixed point $u \in \partial P$ to a point $v \in \partial P$. The point v is moved clockwise along the boundary of P , starting at u , to find a separator satisfying our constraints, i.e. P_ℓ contains between $2n/9$ and $2n/3$ sites of S . This way, we either find a balanced 1-separator, or jump over at least $2n/3 - 2n/9 = 4n/9$ sites at a point v . In this case, the region bounded by two shortest paths between u and v contains at least $4n/9$ sites. We then try to find a 2- or 3-separator contained within this region.

To find a 2- or 3-separator, we construct a sequence of 3-separators $\Delta_0 \supset \Delta_1 \supset \dots \supset \Delta_k$, where either the final 3-separator Δ_k contains between $2n/9$ and $2n/3$ sites, or we find a balanced 2-separator within Δ_k . During the construction, the invariant that $|\Delta_i \cap S| \geq 2n/9$ is maintained. The first 3-separator, Δ_0 , has as corners the two points u and v on ∂P from before, and a point w , which is an arbitrary point on one of the two shortest paths connecting u and v . So, Δ_0 is exactly the region bounded by the two shortest path connecting u and v that contains at least $4n/9$ sites.

Whenever Δ_i contains at most $2n/3$ sites, we are done. If not, we find Δ_{i+1} as follows. If Δ_i is a degenerate 3-separator, we simply select a subpath of the shortest path that contains $2n/3$ sites. Similarly, if at least $2n/9$ sites lie on any of the bounding shortest paths. If neither of these cases holds, then there are at least $2n/3 + 1 - 3(2n/9 - 1) = 4$ sites in the interior of Δ_i . In the following, we find a 3-separator Δ_{i+1} with either $|\Delta_{i+1} \cap S| < |\Delta_i \cap S|$, or $|\text{Int}(\Delta_{i+1}) \cap S| < |\text{Int}(\Delta_i) \cap S|$, or a valid 2-separator contained within Δ_i . Because each Δ_i contains either fewer sites, or fewer sites in its interior, than its predecessor, we eventually find a 3-separator with the desired number of sites, or we end up in one of the easy degenerate cases. In the following description, we drop the subscript i for ease of description.

We define a *good* path to be a path from a point $p \in \Delta$ to a corner u to be a shortest path $\pi(p, u)$ that is fully contained within Δ . To make sure our definition is also correct when one of the corners lies inside Δ , we do not allow the path to cross $\pi(v, w)$. See Figure 19 for an illustration. Essentially, we see the coinciding part of the shortest paths as having an infinitesimally small separation between them. The following provides a formal definition of a good path.

Definition 3. A shortest path $\pi(p, u)$ from a point p that lies in a 3-separator Δ to a corner u of Δ , is a *good path* if it is fully contained within Δ , and $\pi(p, u)$ is a shortest path in the polygonal domain $P \cap \Delta$, where the outer polygon is Δ .




Figure 19: Of the shortest paths $\pi(p, u)$, $\pi(q, u)$, and $\pi(r, u)$ only $\pi(p, u)$ is a good path.




Figure 20: Any path from z to v that exits Δ through $\pi(u, w)$ must intersect a path from z to u .

We consider the closed region $Z_u \subset \Delta$ such that for every $p \in Z_u$ there is a good path to u . Because for any p' on a good path $\pi(p, u)$ the path $\pi(p', u)$ is also a good path, Z_u is connected. Z_u is bounded by $\pi(u, v)$, $\pi(u, w)$, ∂H , and a curve B_u that connects v to w , see Figure 20. Because we do not consider ∂H to be part of B_u , this is a possibly disconnected curve that consists of edges from the shortest path map of u and Δ . The shortest path map of a point p in a polygonal domain P partitions the free space into maximal regions, such that for any two points in the same region the shortest paths from u to both points use the same vertices of P [41]. We call the curves of the shortest path map for which there are two topologically distinct paths from p to any point on the curve *walls*. We prove the following lemma of Abam et al. [3] for our definition of a good path.

Lemma 26 (Lemma 3.2 of [3]). *For any point $z \in B_u$, there are good paths $\pi(z, u)$, $\pi(z, v)$, and $\pi(z, w)$ to the three corners of Δ .*

Proof. By definition of B_u , there is a good path from z to u . When z lies on $\pi(v, w)$, the subpaths of $\pi(v, w)$ from z to v and to w are good paths. So, assume $z \notin \pi(v, w)$. We will prove by contradiction that there is a good path from z to v , and by symmetry from z




Figure 21: The region Z_u where all points have good paths to u . In blue and green the 2-separator with corners z_1, u, z_2 and z_1, v, z_2 , respectively.


w . Suppose there is no good path from z to v . Because z is on B_u , there is also a shortest path from z to u that is not a good path, so it is not contained within Δ , or it crosses Δ . This path $\pi(z, u)$ must exit (or cross) Δ through $\pi(v, w)$, because otherwise the path could simply continue along $\pi(v, u)$ or $\pi(w, u)$ and stay within Δ . Similarly, $\pi(z, v)$ must exit Δ through $\pi(u, w)$. The path $\pi(z, u)$ either goes around v or w , as shown in Figure 20. In both cases, any path to v that starts at z and exits (or crosses) through $\pi(u, w)$ and does not intersect $\pi(u, w)$ again, must intersect the path $\pi(z, u)$. As these shortest paths start at the same point, this is a contradiction with the fact that two shortest paths can only cross once. \square

This lemma implies that for any point z on B_u , the 3-separator with corners u, v , and z is contained within Δ . We use this observation by moving a point z along B_u from w to v . At the start, the 3-separator defined by u, v , and z , which we denote by T_z , is equal to Δ . At the end, T_z is equal to the degenerate 3-separator $\pi(u, v)$. Note that, in contrast to the situation of Abam et al. [3] for a terrain, B_u is not necessarily continuous, as it can be interrupted by holes. But, if B_u intersects a hole, it intersects this hole exactly twice, because Z_u is connected. A directed walk along B_u , jumping at holes, is thus still well-defined. We walk a point z along B_u until one of the following happens: (1) $|S \cap T_z|$ or $|S \cap \text{Int}(T_z)|$ decreases, or (2) z encounters a hole.


In case (1), either T_z contains at least $2n/9$ sites, then we set $\Delta_{i+1} := T_z$, or we jump over at least $2n/3 - 2n/9 = 4n/9$ sites. This can happen because the shortest path to u , to v , or both jump over a hole. We assume the path to u jumps, the approach for when the path to v jumps is symmetric. The 3-separator u, z, w' , with w' on one of the two $\pi(z, u)$, could contain the same number of sites, in the closure and in its interior, as Δ . Therefore, we select an arbitrary site $s \in S$ that lies in the region bounded by the two shortest paths. We then consider the two 3-separators with u, z , and s as corners. As s now lies on the boundary of the 3-separators, both of these separators contain less sites in their interior than Δ . And, as they partition a region that contains at least $4n/9$ sites, one of them contains at least $2n/9$ sites. We set Δ_{i+1} to be this 3-separator.

In case (2), let z_1 denote the point we encounter the hole H_i , and z_2 be the point where B_u exits H_i , see Figure 21. Whenever $|T_{z_2} \cap S| \geq 2n/9$, we simply continue the walk at z_2 , until we again end up in one of the two cases. If not, then there are at least $2n/3 - 2n/9 = 4n/9$ sites in $T_{z_1} \setminus T_{z_2}$, because $T_{z_2} \subset T_{z_1}$. Note that $T_{z_1} \setminus T_{z_2}$ is exactly the union of the two 2-separators z_1, u, z_2 and z_1, v, z_2 . This means that either the 2-separator z_1, u, z_2 or z_1, v, z_2 contains at least $2n/9$ sites. Suppose that the separator including u contains at least $2n/9$ sites. We then try to find a balanced 2-separator z_1, u, z^* with $z^* \in \partial H_i$ that is contained within the 2-separator z_1, v, z_2 . For each point $z^* \in \partial H_i$ on the “ u -side” of H_i (blue in Figure 21), $\pi(z^*, u)$ is contained within the 2-separator z_1, u, z_2 , as it cannot cross $\pi(z_1, u)$ and $\pi(z_2, u)$. As we did to find a 1-separator, we walk z^* along ∂H_i from z_1 to z_2 on the “ u -side”. As before, we now find a balanced 2-separator, and we are done, or we jump over at least $2n/3 - 2n/9 = 4n/9$ sites. Again, this region does not necessarily contain less sites than Δ , thus we continue as before by selecting a site s in this region as the third corner of Δ_{i+1} . Similarly, when the 2-separator with v as a corner contains more than $2n/9$ sites, we walk z^* along the other side of H_i and consider the 2-separator z_1, v, z^* .

Correction of a technical issue in [3]. Abam, de Berg, and Seraji [3] present an approach to find a balanced sp-separator on a terrain \mathcal{T} . They define an sp-separator as either a 1-separator or a 3-separator, where the three shortest paths are disjoint except for their mutual endpoints. However, on a polyhedral terrain these paths might also not be disjoint. This can, for example, happen when choosing a site s in the area bounded by two shortest paths as a new corner. Figure 22 illustrates this for a polygonal domain, and can be generalized to a terrain by making the holes into very high walls. Their subsequent definition of a good path, which is simply a shortest path contained within Δ , would then imply that the entire interior of the box, including the blue region in Figure 22a, is contained in Z_u . The question is then how we move z from w to v . If we would move z along $\pi(w, v)$ at the start, the site w will immediately enter the interior of the 3-separator u, v, z , see Figure 22b. The number of sites in the interior has thus actually *increased* instead of decreased. In a polyhedral terrain, we can define a good path similar to Definition 3, but then consider the path $\pi(p, u)$ in the polyhedral terrain $\mathcal{T} \cap \Delta$. Using our renewed definition, the proof of Abam et al. [3] also holds in the case of coinciding edges on a terrain.



(a) The region Z_u according to our definition of a good path in orange. The region Z_u according to [3] includes the blue region.



(b) The 3-separator T_z has more sites in its interior than the original 3-separator Δ .

Figure 22: A polygonal domain where the shortest paths in the separator Δ with corners u, v, w are not disjoint.