Spanner for the $0/1/\infty$ Weighted Region Problem

Joachim Gudmundsson
The University of Sydney, Australia

Zijin Huang
The University of Sydney, Australia

André van Renssen
The University of Sydney, Australia

Sampson Wong
University of Copenhagen, Denmark

Abstract

We consider the problem of computing an approximate weighted shortest path in a weighted planar subdivision, with weights assigned from the set $\{0, 1, \infty\}$. The subdivision includes zero-cost regions (0-regions) with weight 0 and obstacles with weight ∞, all embedded in a plane with weight 1. In a polygonal domain, where the 0-regions and obstacles are non-overlapping polygons (not necessarily convex) with in total N vertices, we present an algorithm that computes a $(1+\varepsilon)$-approximate spanner of the input vertices in expected $\widetilde{O}(N/\varepsilon^3)$ time[1], for $0 < \varepsilon < 1$. Using our spanner, we can compute a $(1+\varepsilon)$-approximate weighted shortest path between any two points (not necessarily vertices) in $\widetilde{O}(N/\varepsilon^3)$ time. Furthermore, we prove that our results more generally apply to non-polygonal convex regions. Using this generalisation, one can approximate the weak partial Fréchet similarity [7] between two polygonal curves in expected $\widetilde{O}(n^2/\varepsilon^2)$ time, where n is the total number of vertices of the input curves.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases weighted region problem, approximate shortest path, spanner

Digital Object Identifier 10.4230/LIPIcs.WADS.2025.33

Related Version *Full Version*: https://arxiv.org/abs/2407.01951

Funding *Joachim Gudmundsson*: This research was partially funded by the Australian Government through the Australian Research Council (project number DP240101353).
André van Renssen: This research was partially funded by the Australian Government through the Australian Research Council (project number DP240101353).

1 Introduction

The Weighted Region Problem (WRP) is a generalization of the shortest path problem, considering a planar subdivision E where each face has a non-negative weight associated with it. A path σ in E can be partitioned into a set of subpaths $\{\sigma_1, ..., \sigma_k\}$ based on its intersection with faces in E, where a subpath σ_i starts at the point s_i and ends at the point t_i. Both s_i and t_i must lie on the boundary of the same face F_i. The weight of the subpath σ_i is the Euclidean length of σ_i times the weight assigned to F_i. The total weight of a path is the sum of the weights of its subpaths. The goal of the WRP is to find the weighted shortest path from a source point s to a target point t. When the weights are in the set $\{0, 1, \infty\}$, this problem is referred to as the $0/1/\infty$ Weighted Region Problem [15].

[1] $\widetilde{O}(\cdot)$ notation ignores poly-logarithmic dependencies on N and $1/\varepsilon$.

19th International Symposium on Algorithms and Data Structures (WADS 2025).
Editors: Pat Morin and Eunjin Oh; Article No. 33; pp. 33:1–33:15
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Researchers have conjectured that the WRP is difficult to solve [15], and recent studies confirm this conjecture – the Weighted Region Problem is unsolvable in the algebraic computation model over the rational numbers. De Carufel et al. [14] demonstrated that the WRP cannot be solved exactly even with only three different weights. De Berg et al. [12] confirmed its unsolvability with just two different weights. Mitchell and Papadimitriou [17] illustrated that in two dimensions, a weighted shortest path can intersect at least $\Omega(n^2)$ boundaries even when the regions are convex.

Due to the difficulty of solving the WRP exactly, approximation algorithms have been considered. A common approach is to discretize the problem space, either by assuming the space is a tessellation of convex polygons with exactly one associated weight [5], or by placing Steiner (sample) points on the boundaries of the regions [1, 2, 9, 16, 18]. In these approaches, the number of sample points depends not only on the complexity of the regions but also on geometric parameters such as the maximum integer coordinate of any vertex and the ratio r_w of the maximum weight over the minimum weight. As r_w increases, so does the number of required sample points. As a result, the weights are required to be strictly positive.

1.1 Related work

Our work is closely related to the data structure and algorithm by Gewali et al. [15] to solve the $0/1/\infty$ weighted region problem. Their algorithm takes a polygonal domain with N vertices as input and constructs a *critical graph* $\mathcal{G}^* = (\mathcal{V}^*, \mathcal{E}^*)$ (a type of visibility graph) with $O(N^2)$ edges. Dijkstra's shortest path algorithm can be used on $\mathcal{G}^*$ to compute a weighted shortest path between any pair of vertices in $O(N^2)$ time.

In $0/1/\infty$ weighted regions, a weighted shortest path P^* avoids obstacles and traverses the 0-regions freely, while minimizing its length in the plane (1-region). Consider two (closed) regions A and B, each either a 0-region or an obstacle. The key observation in [15] is that an edge in P^* connecting A and B must be *locally optimal* (see Fact 16). For example, an edge (a, b) connecting two convex 0-regions A and B must be perpendicular to the tangent touching $a \in A$ and the tangent touching $b \in B$. Gewali et al. [15] showed that G^* contains all such locally optimal edges in G^*, which implies that G^* must contain the optimal path between any pair of vertices in G^*.

1.2 Our Contribution

In this paper, we build on the work by Gewali et al. [15] with a focus on the 0-regions as they are not handled well by existing approximation schemes (using sample points or tessellation). In Section 2, we consider the 0/1 weighted region problem where the 0-regions are convex but not necessarily polygonal.

▶ **Problem 1.** *In the planar subdivision induced by the plane with weight* 1 *and a set* $\mathcal{Z}$ *of non-overlapping convex zero-cost regions (0-regions) with weight* 0*, given an approximation error* $0 < \varepsilon < 1$*, find a* $(1 + \varepsilon)$*-approximate weighted shortest path from an arbitrary point* s *to an arbitrary point* t*.*

The high-level idea is that, in order to obtain $(1 + \varepsilon)$-approximate shortest paths, we place $O(1/\varepsilon)$ sample points on the boundary of each 0-region; the number of sample points is independent of other parameters. Using these sample points, we construct a Θ-graph and $O(1/\varepsilon)$ trapezoidal maps, which are part of our data structure $\mathcal{B}$. The trapezoidal maps ensure the existence of good paths between 0-regions that are close[2] to each other, while

[2] A precise definition is provided in Lemma 8 and 10, Section 2.

the Θ-graph ensures the same for 0-regions that are far from each other. To the best of our knowledge, our algorithm is the first near-linear time $(1+\varepsilon)$-approximation algorithm that finds an approximated weighted shortest path in a 0/1 weighted region. Note that both Theorem 1 and 3 apply to polygonal domains with non-convex regions, since they can be triangulated into a linear number of triangles.

▶ **Theorem 1.** *Consider a planar subdivision induced by a plane with weight 1, containing a set $\mathcal{Z}$ of non-overlapping convex 0-regions with weight 0. Let $|\mathcal{Z}| = n$ and N denote the total number of vertices in $\mathcal{Z}$. For any approximation factor $0 < \varepsilon < 1$, a data structure $\mathcal{B}$ can be constructed over $\mathcal{Z}$ in $O(N + (n/\varepsilon^2)(\log(n/\varepsilon) + \log N))$ expected time, with a total size of $O(N + n/\varepsilon^2)$. When queried with points s and t, $\mathcal{B}$ can return a weighted path P from s to t in $O(N + n/\varepsilon^2 + (n/\varepsilon)\log(n/\varepsilon) + (\log N)/\varepsilon)$ time, satisfying $\mathbf{w}(P) \leq (1+\varepsilon)\cdot\mathbf{w}(P^*)$, where P^* is the optimal weighted shortest path from s to t.*

To use our algorithm on an application, we prove the above theorem in a more general setting, where the 0-regions are non-polygonal. In Section 3, we use our algorithm to approximate the partial weak Fréchet similarity of two polygonal curves. This problem was first studied by Buchin et al. [7], and they presented a cubic time algorithm. De Carufel et al. [13] later transformed the problem into a weighted shortest path problem amidst 0/1-regions. Using Theorem 1, our algorithm is the first near-quadratic time $(\sqrt{2}+\varepsilon)$-approximation algorithm for computing the partial weak Fréchet similarity between a pair of polygonal curves.

Buchin et al. [8] showed that there is no strongly subquadratic time algorithm for approximating the weak Fréchet distance within a factor less than 3 unless the strong exponential-time hypothesis fails. Approximating the partial weak Fréchet similarity is at least as hard as approximating the weak Fréchet distance. As a result, it is unlikely that a subquadratic time algorithm exists.

▶ **Theorem 2.** *One can approximate the partial weak Fréchet similarity of two curves within a factor of $(\sqrt{2}+\varepsilon)$ in $O((n^2/\varepsilon^2)\log(n/\varepsilon))$ expected time.*

In Section 4, we generalise our data structure to also allow convex obstacles that cannot be traversed, i.e., obstacles of weight ∞. Let $\mathbf{d}(a,b)$ denote the weight of the weighted shortest path from a to b amidst $0/1/\infty$-weighted region. By introducing additional sample points, we show that if we need to take a detour from a sample point a to a sample point b, there exists a set D (a detour) of edges in $\mathcal{B}$ such that the total length of D approximates the distance $\mathbf{d}(a,b)$ within a factor of $1+\varepsilon$.

In the special case that the 0-regions and obstacles are polygonal, $\mathcal{B}$ is a $(1+\varepsilon)$-spanner of the input vertices. To the best of our knowledge, our algorithm is the first near-linear time $(1+\varepsilon)$-approximation algorithm for the weighted shortest path in a $0/1/\infty$ weighted region.

▶ **Theorem 3.** *Consider a planar subdivision induced by a plane with a weight of 1, consisting of two sets of convex and non-overlapping regions: 0-regions $\mathcal{Z}$ with a weight of 0, and obstacles $\mathcal{O}$ with a weight of ∞. Let $n = |\mathcal{Z}| + |\mathcal{O}|$ and let N denote be the total number of vertices in $\mathcal{Z} \cup \mathcal{O}$. For any approximation factor $0 < \varepsilon < 1$, a data structure $\mathcal{B}$ can be constructed over $\mathcal{Z} \cup \mathcal{O}$ in $O(N + (n/\varepsilon^3)(\log(n/\varepsilon) + \log N))$ expected time, with a total size of $O(N + n/\varepsilon^3)$. When queried with arbitrary points s and t, $\mathcal{B}$ returns a path P from s to t in $O(N + n/\varepsilon^3 + (n/\varepsilon^2)\log(n/\varepsilon) + (\log N)/\varepsilon)$ time, ensuring that $\mathbf{w}(P) \leq (1+\varepsilon)\cdot\mathbf{w}(P^*)$, where P^* is the optimal weighted shortest path from s to t.*

2 Shortest path amidst 0-regions

The exact version of Problem 1 has a brute-force $\Omega(n^2)$ algorithm, by computing the distance between every pair of 0-regions. To compute an approximate solution, the goal is to construct an undirected weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with a near-linear number of edges, such that there exists a path P between two 0-regions in $\mathcal{G}$ with $\mathbf{w}(P) \leq (1+\varepsilon) \cdot \mathbf{w}(P^*)$, for a fixed parameter $\varepsilon > 0$, where P^* is the weighted shortest path.

To this end, we will use two data structures: trapezoidal maps and Θ-graphs. Both data structures are used to determine which pairs of 0-regions are connected. The trapezoidal maps will ensure that there exist good paths between 0-regions that are close to each other, while the Θ-graph ensures the same for 0-regions that are far from each other.

2.1 Construction of the data structure

In order to define our data structure, we first define a set of directions. Let $\theta < \pi/6$ be a fixed positive real number. Let $r(k\theta)$ be the direction with a counter-clockwise angle of $k\theta$ with the positive x-axis. Let $r(p, k\theta)$ be the ray originating from the point p with a counter-clockwise angle of $k\theta$ with the positive x-axis. To simplify the discussion, we will assume that $(\pi/2)/\theta \in \mathbb{Z}$ to guarantee that if $r(k\theta)$ exists, then so does $r(k\theta + \pi/2)$.

For a 0-region A, we define a set $\mathcal{SP}(A)$ of sample points on the boundary of A. Let $sp(A, k\theta)$ be a sample point on the boundary ∂A of A such that $sp(A, k\theta)$ is extreme in the direction $r(k\theta)$ (see Figure 1). When the geometric region A is clear from context, we write $sp(k\theta)$ instead of $sp(A, k\theta)$.

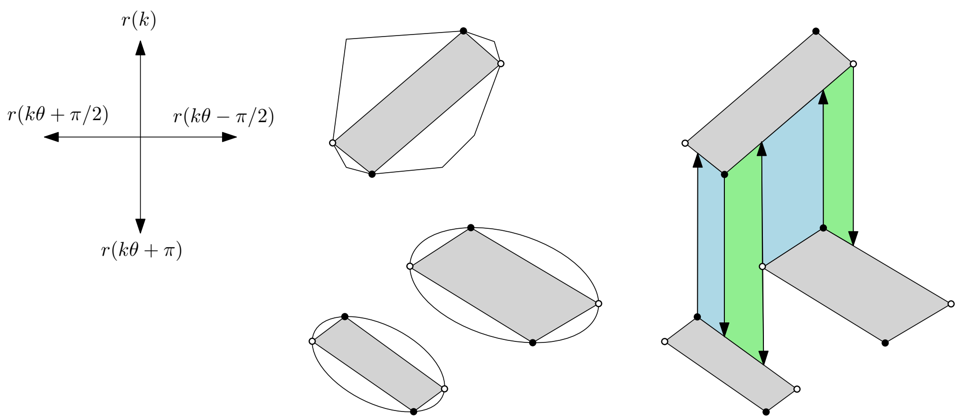

Figure 1 The sample points that are extreme in the directions of $r(k\theta)$ and $r(k\theta+\pi)$ are marked with black dots. The sample points that are extreme in the directions of $r(k\theta \pm \pi/2)$ are marked with circles. Using these sample points, we generate a simplified polygon and construct $\mathcal{M}(k)$; the blue and green regions are examples of faces in $\mathcal{M}(k)$.

Let $\partial A(a, a')$ define the subset of ∂A traversed from point a to point a' in counter-clockwise order, where $a, a' \in \partial A$. We say a line l *overlaps* A if l and A intersect at more than one point. We say two regions, A and B, are non-overlapping if their interiors do not intersect.

A point $p \in \partial A$ can be an extreme point for more than one direction, in which case we call p a vertex of A. There may be more than one extreme point on A for a single direction. If p is the extreme point on A for consecutive directions $\{r(k\theta), r((k+1)\theta), ..., r((k+m)\theta)\}$, we say $p = sp(k\theta)$, $p = sp((k+1)\theta)$, ..., and $p = sp((k+m)\theta)$ simultaneously. If there is more than one extreme point for a single direction $r(k\theta)$, these extreme points must lie on some segment $ab \subseteq \partial A$, and we say both a and b are $sp(k\theta)$. The sample points on the boundary of a convex region can be computed by traversing the boundary.

▶ **Observation 4.** *Given n convex regions with N vertices in total, there are $O(n/\theta)$ sample points, and it takes $O(N + n/\theta)$ time to compute them.*

Using the sample points on a 0-region A, we can generate a simplified 0-region (a convex polygon) `simpl`(A) by connecting adjacent sample points of every 0-region (see Figure 1). Using the set `simpl`$(\mathcal{Z})$ of simplified 0-regions, we will generate a set of trapezoidal maps, and we say `simpl`(A) and `simpl`(B) are adjacent if they are both adjacent to the same face in a trapezoidal map. We construct the query data structure $\mathcal{B}$ using Algorithm 1. We analyse $\mathcal{B}$ in the full version. Let $\|AB\|$ denote the Euclidean distance between two geometric objects A and B.

Algorithm 1 Construct $\mathcal{B}$ with 0-regions.

This algorithm takes as input a set $\mathcal{Z}$ of non-overlapping and convex 0-regions, and constructs a data structure $\mathcal{B}$. The undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is initially empty.

1. Compute the sample points $\mathcal{SP}(\mathcal{Z})$, and add $\mathcal{SP}(\mathcal{Z})$ to $\mathcal{V}$.
2. Pick an arbitrary sample point as the anchor $ak(A)$ for every $A \in \mathcal{Z}$. For each $p \in \mathcal{SP}(A)$, add $e = (p, ak(A))$ to $\mathcal{E}$, and set $\mathbf{w}(e) = 0$.
3. For each direction $r(k\theta)$, generate a trapezoidal map $\mathcal{M}(k)$ using `simpl`$(\mathcal{Z})$, and do the following for each $\mathcal{M}(k)$ (see [11, Theorem 6.3 and 6.8] for trapezoidal map construction).
   - For each face $F \in \mathcal{M}(k)$ adjacent to A and B, $A \neq B$, add $e = (ak(A), ak(B))$ to $\mathcal{E}$, and set $\mathbf{w}(e) = \|AB\|$.
4. With $\mathcal{V}$ as the input, generate a Θ-graph $\mathcal{G}_\Theta = (\mathcal{V}, \mathcal{E}_\Theta)$.
   - For each edge $(p, q) \in \mathcal{E}_\Theta$, if p and q do not belong to the same 0-region, add $e = (p, q)$ to $\mathcal{E}$, and set $\mathbf{w}(e) = \|pq\|$.
5. Return $\mathcal{B} = \{\mathcal{M}(k) \mid \forall k \in [0, 2\pi/\theta), k \in \mathbb{Z}\} \cup \{\mathcal{G}, \mathcal{G}_\Theta\}$ as the data structure.

▶ **Lemma 5.** *Given an approximation factor $0 < \varepsilon < 1$, and n non-overlapping convex 0-regions with total complexity N, one can build the data structure $\mathcal{B}$ in $O(N + (n/\varepsilon^2)(\log(n/\varepsilon) + \log N))$ time, and the total size of $\mathcal{B}$ is $O(N + n/\varepsilon^2)$.*

2.2 Trapezoidal map

Before arguing that there exists a good path using the edges constructed, we start with an observation about the pair of points realising the shortest distance between two 0-regions. For a convex region A and a point $p \in \partial A$, there exists at least one *supporting line* $l = l_t(A, p)$ going through p such that A lies entirely in one of the two halfplanes determined by l [19]. Let $p \in \partial A$, and $q \in \partial B$. We can observe that if pq realises $\|AB\|$, then pq must be perpendicular to a pair of supporting lines $l_t(A, p)$ and $l_t(B, q)$.

▶ **Observation 6.** *Let A and B be two convex regions. Let pq be the line segment realising $\|AB\|$, where $p \in A$ and $q \in B$. The segment pq must be perpendicular to a pair of supporting lines $l_t(A, p)$ and $l_t(B, q)$.*

We will show that if pq realises the distance between two 0-regions, we can transform pq into another segment pq' such that pq' is parallel to some direction $r(k\theta)$, and $\|pq'\|$ approximates $\|pq\|$. To do this, we first need a fact (see the full version for a proof).

▶ **Lemma 7.** *Given a segment pq, let α (resp. β) be the acute angle between pq and $r(p,(k+1)\theta)$ (resp. $r(p,k\theta)$), where $\alpha+\beta=\theta<\pi/6$. Let q' be the intersection of $r(p,(k+1)\theta)$ and $r(q,k\theta+\pi/2)$. We have that $\|pq'\|=(\cos(\beta)/\cos(\theta))\cdot\|pq\|$, and $\|qq'\|=(\sin(\alpha)/\cos(\theta))\cdot\|pq\|$.*

Let pq realise $\|AB\|$. We consider the scenario when $\|pq\|$ is relatively small compared to the horizontal span of (say) B. Using the above lemma, we will show that we have constructed a set of edges in $\mathcal{E}$ connecting two sample points $a\in A$ and $b\in B$, such that the total weight of these edges approximates $\|pq\|$.

▶ **Lemma 8.** *Let pq realise $\mathbf{d}(A,B)$, where $p\in A$ and $q\in B$. Let $p\in\partial A(a,a')$, and $q\in\partial B(b',b)$, where points a and a' (resp. b and b') are adjacent sample points on 0-region A (resp. B). If $\max\{\|pa'\|,\|qb\|\}\geq(\sin(\alpha)/\cos(\theta))\cdot\|pq\|$ or $\max\{\|pa\|,\|qb'\|\}\geq(\sin(\beta)/\cos(\theta))\cdot\|pq\|$, then there exists a path $P\subseteq\mathcal{E}$ from A to B such that $\mathbf{w}(P)\leq(\cos(\beta)/\cos(\theta))\cdot\|pq\|$.*

Proof. Without loss of generality, assume that $\|qb\|\geq(\sin(\alpha)/\cos(\theta))\cdot\|pq\|$. Lemma 7 implies that there exists a point $q'\in\partial B(b,q)$ such that pq' is in some direction $r(k\theta)$, and $\|pq'\|\leq(\cos(\beta)/\cos(\theta))\cdot\|pq\|$. Because

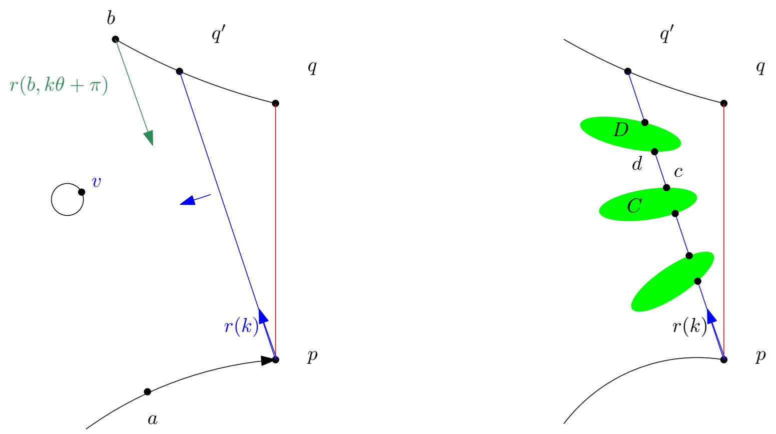

Figure 2 In the left figure, if pq' does not overlap a 0-region, we slide pq' until it touches a sample point. In the right figure, we slide each inter-region segment (cd as an example) the same way.

Observe that pq cannot overlap any 0-region; otherwise, pq does not realise $\mathbf{d}(A,B)$. If pq' does not overlap any 0-region (see Figure 2, left), we fix the orientation of pq', and move p along $\partial A(p,a)$ and q along $\partial B(b,q)$, until pq' touches a sample point.

If pq' touches a (resp. b), then $r(a,k)$ (resp. $r(b,k\theta+\pi)$) hits B (resp. A). If pq' touches a sample point $v\notin A\cup B$, v must be extreme in the direction $r(k\theta-\pi/2)$. As a result, $r(v,k)$ hits B, and $r(v,k\theta+\pi)$ hits A. In either case, A and B are adjacent in some face of $\mathcal{M}(k)$, and the edge $e=(ak(A),ak(B))$ is in $\mathcal{E}$ by construction. As a result, $\mathbf{w}(e)=\|AB\|=\|pq\|$. This is also trivially true when p or q' is already a sample point.

Otherwise, the segment pq' overlaps a set E' of 0-regions, and there exists a path P from A to B through E' (see Figure 2, right). Since the 0-regions do not overlap, the boundaries of the 0-regions in E' partition pq' into a set of intra-region and inter-region segments. Let cd be one among the set S of inter-region segments, where c is on a 0-region C, and d is on a 0-region D. Using the same argument as above, one can slide cd until it touches a sample point, and edge $(ak(C),ak(D))\in\mathcal{E}$ exists by construction.

In total, traveling from A to B via the 0-regions E' must be less costly than $\|pq'\|$, since $\mathbf{w}(ak(C), ak(D)) = \|CD\| \leq \|cd\|$, and the intra-region segments have weight 0. Summing up the cost of P, we have that

$$\mathbf{w}(P) = \sum_{cd \in S} \mathbf{w}(ak(C), ak(D)) < \sum_{cd \in S} \|cd\| < \|pq'\| = \frac{\cos(\beta)}{\cos(\theta)} \cdot \|pq\|. \qquad \blacktriangleleft$$

2.3 Θ-Graph

In Step 4 of Algorithm 1, we constructed a Θ-graph $\mathcal{G}_\Theta = (\mathcal{V}_\Theta, \mathcal{E}_\Theta)$ using the defined sample points. The vertices $\mathcal{V}_\Theta$ are simply all sample points. The edges in $\mathcal{E}_\Theta$ are constructed using the standard Θ-graph construction [11]. Recall that in Algorithm 1, for every edge $(p, q) \in \mathcal{E}_\Theta$, with $p \in A$, $q \in B$, and $A \neq B$, we add an edge (p, q) to $\mathcal{E}$, and set $\mathbf{w}(p, q) = \|pq\|$.

The Θ-graph constructs a set of "good" edges in $\mathcal{E}$ when the distances between p (resp. q) and its adjacent sample points are small compared to $\|pq\|$. In this case, we argue that there exists a pair of sample points $a \in A$ and $b \in B$, such that $\|ab\| \leq (1/\cos(\theta)) \cdot \|pq\|$. Similar to Lemma 7, we prove the following geometric property in the full version.

▶ **Lemma 9.** *Given a segment pq, let α (resp. β) be the acute angle between pq and $r(p, (k+1)\theta)$ (resp. $r(p, k\theta)$), where $\alpha+\beta = \theta < \pi/6$. Let q' be the intersection of $r(p, (k+1)\theta)$ and $r(q, k\theta+\pi/2)$, and let p' be the intersection of $r(q, k\theta+\pi)$ and $r(p, (k+1)\theta+\pi/2)$. Let c be the intersection of pq' and qp'. We have that $\|cp'\| + \|cq'\| = (\sin(\alpha) + \sin(\beta))/(\cos(\theta)\sin(\theta))$, and $\|p'q'\| = (1/\cos(\theta)) \cdot \|pq\|$.*

In the case that both p and q are close to their adjacent sample points, we prove that there exists a pair (a, b) of sample points such that $\|ab\|$ approximates $\|pq\|$ in the full version.

▶ **Lemma 10.** *Let pq realise $\mathbf{d}(A, B)$, where $p \in A$ and $q \in B$. Let $p \in \partial A(a', a)$, and $q \in \partial B(b, b')$, where points a and a' (resp. b and b') are adjacent sample points on A (resp. B). If $\max\{\|pa'\|, \|qb\|\} < (\sin(\alpha)/\cos(\theta)) \cdot \|pq\|$ and $\max\{\|pa\|, \|qb'\|\} < (\sin(\beta)/\cos(\theta)) \cdot \|pq\|$, then $\mathbf{d}(a, b) < (1/\cos(\theta)) \cdot \|pq\|$.*

2.4 The quality of the path

For now, assume that s and t lie in some 0-region. An optimal s-t path P^* consists of a set of segments, where the endpoints of each segment lie on the boundaries of the 0-regions. A segment pq either lies within a 0-region or connects two different 0-region. Since it costs nothing to follow an edge inside a 0-region, the weight of P is the total weight of those edges connecting different 0-regions.

Let pq realise the distance between 0-regions A and B, where p lies on ∂A between sample points a and a', and q lies on ∂B between sample points b and b'. In Lemma 8, we have shown that if $\max\{\|pa'\|, \|qb\|\} \geq (\sin(\alpha)/\cos(\theta)) \cdot \|pq\|$ or $\max\{\|pa\|, \|qb'\|\} \geq (\sin(\beta)/\cos(\theta)) \cdot \|pq\|$, there exists a path $P \subseteq \mathcal{E}$ from a sample point on A to a sample point on B of length at most $(\cos(\beta)/\cos(\theta)) \cdot \|pq\|$.

In Lemma 10, we have shown that if $\max\{\|pa'\|, \|qb\|\} < (\sin(\alpha)/\cos(\theta)) \cdot \|pq\|$ and $\max\{\|pa\|, \|qb'\|\} < (\sin(\beta)/\cos(\theta)) \cdot \|pq\|$, there exist sample points $a \in A$ and $b \in B$, such that $\mathbf{d}(a, b) < (1/\cos(\theta)) \cdot \|pq\|$. To obtain a path between a and b in this case, we rely on the Θ-graph. The tightest bounds on the length of this path are due to Bose et al. [4], who showed that the spanning ratio of a Θ-graph is at most $r_\theta = 1 + 2\sin(\theta/2)/(\cos(\theta/2) - \sin(\theta/2))$. Our final approximation ratio is $(1/\cos(\theta)) \cdot r_\theta$. Given an input approximation factor $1+\varepsilon$, we can compute a desired angle $\theta \in O(\varepsilon)$. The undirected graph $\mathcal{G} \in \mathcal{B}$ is therefore a $(1+\varepsilon)$-spanner of the sample points. See the full version for details.

▶ **Lemma 11.** *The graph $\mathcal{G} \in \mathcal{B}$ contains a path $P \subseteq \mathcal{E}$ from sample point a to sample point b such that $\mathbf{w}(P) \leq (1+\varepsilon) \cdot \mathbf{w}(P^*)$, where P^* is the optimal path from a to b.*

Given a pair of query points s and t, we add s and t to $\mathcal{B}$. We treat both s and t as 0-regions with no interior to enable previous lemmas, yielding the theorem below. See the full version for details.

▶ **Theorem 1.** *Consider a planar subdivision induced by a plane with weight 1, containing a set $\mathcal{Z}$ of non-overlapping convex 0-regions with weight 0. Let $|\mathcal{Z}| = n$ and N denote the total number of vertices in $\mathcal{Z}$. For any approximation factor $0 < \varepsilon < 1$, a data structure $\mathcal{B}$ can be constructed over $\mathcal{Z}$ in $O(N + (n/\varepsilon^2)(\log(n/\varepsilon) + \log N))$ expected time, with a total size of $O(N + n/\varepsilon^2)$. When queried with points s and t, $\mathcal{B}$ can return a weighted path P from s to t in $O(N + n/\varepsilon^2 + (n/\varepsilon)\log(n/\varepsilon) + (\log N)/\varepsilon)$ time, satisfying $\mathbf{w}(P) \leq (1+\varepsilon) \cdot \mathbf{w}(P^*)$, where P^* is the optimal weighted shortest path from s to t.*

3 Partial weak Fréchet similarity

This section highlights one application of our data structure: approximating the partial weak Fréchet similarity. The Fréchet distance is a popular measure of the similarity between two polygonal curves. An *orientation-preserving reparameterisation* is a continuous and bijective function $f : [0,1] \to [0,1]$ such that $f(0) = 0$, and $f(1) = 1$. The $\mathbf{width}_{f,g}(\pi, \sigma)$ between two curves π and σ with respect to the reparameterisations f and g, is defined as follows.

$$\mathbf{width}_{f,g}(\pi, \sigma) = \max_{t \in [0,1]} \|\pi(f(t)) - \sigma(g(t))\|$$

Consider the scenario where a person is walking his dog with a leash connecting them: the person needs to stay on π while walking according to f, and the dog needs to stay on σ while walking according to g. The maximum leash length is the width between π and σ with respect to the reparameterisations f and g. The standard Fréchet distance $\delta_F(\pi, \sigma)$ is the minimum leash length required over all possible walks (defined by reparameterisations f and g).

$$\delta_F(\pi, \sigma) = \inf_{f,g \in [0,1] \to [0,1]} \mathbf{width}_{f,g}(\pi, \sigma)$$

Problems relating to the Fréchet distance are commonly solved in a configuration space called the *freespace diagram.* The *free space* $\mathcal{F}_d(\pi, \sigma)$ with respect to the Fréchet distance d is the union of all pairs of points $x \in \pi$ and $y \in \sigma$ such that the distance between x and y is at most d. As opposed to the free space, we will call $[0, \|\pi\|] \times [0, \|\sigma\|] \setminus \mathcal{F}_d(\pi, \sigma)$ the *forbidden space.*

$$\mathcal{F}_d(\pi, \sigma) = \{(x, y) \in [0, \|\pi\|] \times [0, \|\sigma\|] \mid \|\pi(x) - \sigma(y)\| \leq d\}$$

The *freespace diagram* $\mathcal{D}_d(\pi, \sigma)$ is a data structure that stores the free space $\mathcal{F}_d(\pi, \sigma)$ in n^2 cells. Alt and Godau [3] showed that the intersection of the free space with each cell is the intersection of an ellipse and a rectangle. Therefore the free space in each cell is convex, and its boundary is of constant complexity. For two polygonal curves with complexity n, Alt and Godau [3] proved the following fact.

▶ **Fact 12.** *The freespace diagram contains at most n^2 cells, and it can be constructed in $O(n^2 \log n)$ time. The free space inside each cell is the intersection of an ellipse and a rectangle.*

It is well-known that if one can find an xy-monotone path in $\mathcal{D}_d(\pi, \sigma)$ from the bottom-left corner s to the top-right corner t via the free space, then $\delta_F(\pi, \sigma) \leq d$.

The notion of weak Fréchet distance relaxes the requirement of the reparameterisation f: it still needs to be continuous but not bijective. This means that the person and the dog can walk backward. To determine if the weak Fréchet distance $\delta_{wF}(\pi, \sigma)$ is at most d, we need to find only a (potentially not xy-monotone) path through the free space from s to t in $\mathcal{D}_d(\pi, \sigma)$. Buchin et al. [8] showed that there is no strongly subquadratic time algorithm for approximating the weak Fréchet distance within a factor less than 3 unless the strong exponential-time hypothesis fails.

Buchin et al. [7] proposed the *partial Fréchet similarity* (partial similarity in short) to deal with the Fréchet distance's sensitivity to outliers. Instead of determining whether a leash of length d is enough to complete the walk, partial similarity determines how much can be completed given a leash of length d. The partial similarity is the total length of the portion of two curves that are matched under the Fréchet distance d.

Let $\|xy\|_p$ be the distance between point x and point y under the L_p norm. Let $\|v\|$ be the L_2 norm of the vector v. Under the L_p metric, given the desired Fréchet distance d, the partial similarity $S_{f,g}(\pi, \sigma)$ of curves π and σ with respect to the reparameterisations f and g is formally defined as follows [7].

$$S_{f,g}(\pi, \sigma) = \int_{\|\pi(f(t))\sigma(g(t))\|_p \leq d} (\|\pi(f(t))'\| + \|\sigma(g(t))'\|) dt$$

Naturally, we want to compute a pair of reparameterisations f and g that maximise the partial similarity. To do this, Buchin et al. [7] proposed a cubic time algorithm under L_1. They showed that it is sufficient to find an xy-monotone and *rectilinear* path P from s to t such that P intersects as much free space as possible, where s (resp. t) is the bottom-left (resp. top-right) corner of the freespace diagram.

Under the weak Fréchet distance, the monotonicity requirement is removed. But since a path P can traverse back and forth in the freespace diagram, it is no longer meaningful for P to intersect as much free space as possible. We instead are interested in computing a path that intersects as little forbidden space as possible to minimise the portions of the two curves that are not matched within distance d. Therefore, solving the partial weak Fréchet distance problem under the L_2 metric is equivalent to finding a weighted rectilinear shortest path amidst a set of $O(n^2)$ non-overlapping and convex 0-regions embedded in the plane (the forbidden space) with weight 1. By Fact 12, a 0-region is the free space within a cell, which has constant complexity.

Amidst the 0-regions and measured in L_p metric, let OPT_{L_p} denote the weight of the weighted shortest path from s to t, and let ALG_{L_p} denote the weight of the weighted path computed by Theorem 1. Since $OPT_{L_1} \leq \sqrt{2} \cdot OPT_{L_2}$, we have

$$OPT_{L_2} \leq OPT_{L_1} \leq \sqrt{2} \cdot OPT_{L_2} \leq \sqrt{2} \cdot ALG_{L_2} \leq \sqrt{2} \cdot (1+\varepsilon) OPT_{L_2},$$

which leads to the following theorem.

▶ **Theorem 2.** *One can approximate the partial weak Fréchet similarity of two curves within a factor of* $(\sqrt{2}+\varepsilon)$ *in* $O((n^2/\varepsilon^2)\log(n/\varepsilon))$ *expected time.*

4 Shortest path amidst 0-regions and obstacles

In this section, we generalise our data structure from Section 2 to allow convex obstacles that cannot be traversed, i.e., obstacles with weight ∞. Our problem is finding an approximate shortest path amidst 0-regions and obstacles. More concretely, we consider the following problem.

▶ **Problem 2.** *In the planar-subdivision induced by the plane with weight* 1*, and a set of non-overlapping convex regions consisting of obstacles with weight* ∞*, and* 0*-regions with weight* 0*, given an approximation error* $0 < \varepsilon < 1$*, find a* $(1+\varepsilon)$*-approximate weighted shortest path from point* s *to point* t*.*

In this section, let $\mathbf{d}(A, B)$ denote the minimum distance between two geometric regions A and B in a $0/1/\infty$ weighted setting. The Θ-graph can be constructed in an environment with obstacles. Clarkson [10] described such construction over points and polygonal obstacles, and proved that a path that $(1+\varepsilon)$-approximates $\mathbf{d}(a, b)$ exists in the Θ-graph, where a and b are vertices. We will use this Θ-graph in the rest of the paper.

Like in the previous section, we will first describe the construction of the data structure $\mathcal{B}$ and analyse the time and space complexity. We then show that $\mathcal{G} \in \mathcal{B}$ contains a good path between every pair of sample points. We use this to argue the approximation ratio for arbitrary s and t.

4.1 Construction of the data structure

In order to deal with obstacles, we need to define two new types of sample points. For clarity, we refer to the sample points defined previously as the *original sample points*. In Section 2, a trapezoidal map $\mathcal{M}(k)$ was only used to determine if two 0-regions should be connected, and we did not explicitly compute the intersection of a vertical segment and the boundary of a region. With the introduction of obstacles, we do need such intersections. Consider a sample point a. When constructing $\mathcal{M}(k)$, we shoot two vertical rays from $a \in A$, one upwards and one downwards. Let p be the first intersection of $r(a, k)$ with the boundary of some region that is not A. We call p a *propagated sample point*.

The other type of sample points we need is the *tangent sample points*. Given two disjoint obstacles A and B, and a common tangent l, if l touches A at point a and B at b, we add a and b as tangent sample points. When we say a point a is a sample point, a can be any type of sample point.

Recall that $\texttt{simpl}(A)$ is the simplified region by connecting every pair of adjacent sample points of A. $\texttt{simpl}(\mathcal{Z})$ is the set of simplified 0-regions, and $\texttt{simpl}(\mathcal{O})$ is the set of simplified obstacles. We formally define the construction of our data structure. Given a set $\mathcal{O}$ of convex obstacles and a set $\mathcal{Z}$ of convex 0-regions, we build our data structure using Algorithm 2. See the full version for the analysis.

▶ **Lemma 13.** *Given an approximation error* $0 < \varepsilon < 1$*, and* n *non-overlapping convex regions including* 0*-regions and obstacles with total complexity* N*, one can build the data structure* $\mathcal{B}$ *in* $O(N + (n/\varepsilon^3)(\log(n/\varepsilon) + \log N))$ *expected time, and the total size of* $\mathcal{B}$ *is* $O(N + n/\varepsilon^3)$*.*

The structure of the rest of the section is as follows. By Lemma 14, the distance between two adjacent sample points on the boundary of an obstacle approximates the straight line segment. Therefore, we show that we can "snap" the vertices of an optimal path to our sample points. For every segment $pq \subseteq P^*$, we then argue that either the trapezoidal map or the Θ-graph contains a path approximating $\|pq\|$ to within a factor of $1+\varepsilon$.

Algorithm 2 Construct $\mathcal{B}$ with 0-regions and obstacles.

This algorithm takes as input a set of non-overlapping and convex regions, including 0-regions $\mathcal{Z}$ and obstacles $\mathcal{O}$, and constructs a data structure $\mathcal{B}$. The undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is initially empty.

1) Compute the original sample points $\mathcal{SP}(\mathcal{Z}) \cup \mathcal{SP}(\mathcal{O})$, and add them to $\mathcal{V}$.
2) For each direction $r(k\theta)$, generate a trapezoidal map $\mathcal{M}(k)$ using $\texttt{simpl}(\mathcal{Z}) \cup \texttt{simpl}(\mathcal{O})$.
3) For every $\mathcal{M}(k)$, do the following for every face F adjacent to A and B, $A \neq B$.
   a. Compute the propagated sample points, and add them to $\mathcal{V}$.
   b. If A and B are both 0-regions, add $e = (ak(A), ak(B))$ to $\mathcal{E}$, and set $\mathbf{w}(e) = \|AB\|$.
   c. If at least one of A and B is an obstacle, let ab and $a'b'$ be the vertical segments defining F, where $a, a' \in A$, and $b, b' \in B$. Add edges $e_1 = (a, b)$ and $e_2 = (a', b')$ to $\mathcal{E}$. Set $\mathbf{w}(e_1) = \|ab\|$ and $\mathbf{w}(e_2) = \|a'b'\|$.
   d. If A and B are both obstacles, we compute their common tangents. For each common tangent that touches A at a and B at b, add a and b to $\mathcal{V}$.
4) Redefine $\texttt{simpl}(A)$ as the polygon generated by connecting adjacent sample points of A, original, propagated, and tangent sample points included. With $\mathcal{V}$ and $\texttt{simpl}(\mathcal{O})$ as the input, generate a Θ-graph $\mathcal{G}_\Theta = (\mathcal{V}, \mathcal{E}_\Theta)$.
   - For each edge $(p, q) \in \mathcal{E}_\Theta$, if p and q belongs to different regions A and B, add $e = (p, q)$ to $\mathcal{E}$ and set $\mathbf{w}(e) = \|pq\|$.
5) For every pair of adjacent sample points a, a' on an obstacle, add $e = (a, a')$ to edges, and set $\mathbf{w}(e) = \|aa'\|$.
6) Pick an arbitrary sample point as the anchor $ak(A)$ for every $A \in \mathcal{Z}$. For each sample point a of a 0-region A, add $e = (a, ak(A))$ to $\mathcal{E}$, and set $\mathbf{w}(e) = 0$.
7) Return $\mathcal{B} = \{\mathcal{M}(k) \mid \forall k \in [0, 2\pi/\theta), k \in \mathbb{Z}\} \cup \{\mathcal{G}, \mathcal{G}_\Theta\}$ as the data structure.

4.2 Walking on the boundary is not expensive

We first argue that if a and a' are adjacent sample points of A, then $\|aa'\|$ approximates $\|\partial A(a, a')\|$ within a factor of $\sec(\theta/2)$. Therefore, if we find a path P amidst the simplified obstacles, we only need to pay a small factor to transform P into a path amidst the original obstacles. Then, we prove that if pq is part of the optimal path, we can replace pq with a path $P \subseteq \mathcal{E}$, such that $\mathbf{w}(P)$ approximates $\|pq\|$. See the full version for a formal proof.

▶ **Lemma 14.** *Let a and b be adjacent sample points on ∂A, where a appear after b in a counter-clockwise walk. We have that $\|\partial A(a, b)\| \leq \sec(\theta/2) \cdot \|ab\|$.*

The above lemma implies the following. Let $\mathbf{d_S}(a, b)$ be the distance between point a and point b amidst simplified obstacles and simplified 0-regions, and let P be the path achieving this distance. If we partition P using the sample points, in the worst case, each segment connects adjacent sample points on obstacles. This implies the following corollary.

▶ **Corollary 15.** *Let a and b be two adjacent sample points. We have that $\mathbf{d}(a, b) \leq \sec(\theta/2) \cdot \mathbf{d_S}(a, b)$.*

4.3 Snapping a segment of the optimal path to the sample points

Gewali et al. [15] defined three types of *locally optimal* edges joining two simple polygonal regions, and they proved that the shortest path from s to t must be comprised of these locally optimal edges [15, Lemma 2.5] (ignoring edges in 0-regions). For convex obstacles and 0-regions, we need to consider only four types of segments.

▶ **Fact 16.** *If segment pq is in the optimal weighted path P^* amidst convex and non-overlapping 0-regions and obstacles, there must exists two supporting lines $l_t(p)$ and $l_t(q)$ such that pq belong to one of the following cases (ignoring segments in 0-regions).*

1) *pq connects two 0-regions such that $pq \perp l_t(q)$ and $pq \perp l_t(p)$.*

2) *pq connects the point p on a 0-region A and the point q on an obstacle B such that $pq \subset l_t(q)$, and $pq \perp l_t(p)$.*

3) *pq lies on one of the common tangent of two different obstacles.*

4) *p and q are two points on the same obstacle, and $pq = \partial A(p, q)$ or $pq = \partial A(q, p)$.*

In Section 4.3.1, 4.3.2, and 4.3.3, we handle each type of edge in Case 1, 2, and 3, respectively. For an edge pq in each of these cases, we argue that there exists a good path in $\mathcal{G} \in \mathcal{B}$ that approximates pq. Section 4.4 summarises the approximation ratio using the Case 4 edges. In the following lemmas, we assume without loss of generality that pq lies between the direction $r(k\theta)$ and $r((k+1)\theta)$, and α (resp. β) is the measure of the acute angle between $r(p, k\theta)$ (resp. $r(p, (k+1)\theta)$) and pq. By Corollary 15, we consider only simplified obstacles. We delegate the full proofs to the appendix.

4.3.1 Case 1: pq connects two 0-regions

We observe that, unfortunately, Lemma 8 does not trivially apply. When we rotate pq to pq', if pq' overlaps obstacles, a path generated using the skewed set of obstacles can be much longer than $\|pq'\|$, since the path would have to take a detour around the obstacles. The following lemmas resolve this issue. See the full version for full proofs of the lemmas below.

▶ **Lemma 17.** *Let A and B be two 0-regions. Let $pq \subseteq P^*$, where $p \in \partial A$, and $q \in \partial B$. If $\|qb\| \geq (\sin(\alpha)/\cos(\theta)) \cdot \|pq\|$, where b is a sample point adjacent to q, then there exists a sample point $a \in \partial A$ such that $\mathbf{d_S}(a, b) \leq (\cos(\beta)/\cos(\theta)) \cdot \|pq\|$.*

▶ **Lemma 18.** *Let A and B be two 0-regions. Let $pq \subseteq P^*$, where $p \in \partial A$, and $q \in \partial B$. If $\|qb\| < (\sin(\alpha)/\cos(\theta)) \cdot \|pq\|$ and $\|pa\| < (\sin(\beta)/\cos(\theta)) \cdot \|pq\|$, where a (resp. b) is a sample point adjacent to p (resp. q), then $\mathbf{d_S}(a, b) \leq (\sin(\alpha) + \sin(\beta))/(\cos(\theta)\sin(\theta)) \cdot \|pq\|$.*

We combine Lemma 17 and Lemma 18 to obtain a bound on the path length where p and q both lie on 0-regions. Note that when two points p and a lie on the boundary of the same 0-region, $\mathbf{d_S}(a, p) = 0$.

▶ **Lemma 19.** *Let A and B be two convex 0-regions. Let $pq \subseteq P^*$, where $p \in \partial A$ and $q \in \partial B$. There exists a pair of sample points $a \in A$ and $b \in B$, such that $\mathbf{d_S}(p, a) + \mathbf{d_S}(a, b) + \mathbf{d_S}(b, q) \leq \max\{(\cos(\beta)/\cos(\theta)), (\sin(\alpha) + \sin(\beta))/(\cos(\theta)\sin(\theta))\} \cdot \|pq\|$.*

4.3.2 Case 2: pq connects two obstacles

Using Fact 16, we know that if pq connects two obstacles, then pq must lie on a common tangent of A and B. When two obstacles are close, p (and q) may be very far from its adjacent original sample points. Hence, we need the tangent sample points when two obstacles are close (connected via a trapezoidal map). We now show that if p and q lie on obstacles, an approximate path in $\mathcal{B}$ exists. See the full version for a full proof of the lemma below.

▶ **Lemma 20.** *Let A and B be two convex obstacles. Let $pq \subseteq P^*$, where $p \in \partial A$ and $q \in \partial B$. There exists a pair of sample points $a \in A$ and $b \in B$, such that $\mathbf{d_S}(p, a) + \mathbf{d_S}(a, b) + \mathbf{d_S}(b, q) \leq (1/\cos(\theta)) \cdot \|pq\|$.*

4.3.3 Case 3: pq connects a 0-region and an obstacle

In this section, we prove that if pq connects an obstacle A and a 0-region B, there is a path P that approximate pq. See the full version for a full proof of the lemma below.

▶ **Lemma 21.** *Let A be a convex obstacle and let B be a convex 0-region. Let $pq \subseteq P^*$, where $p \in \partial A$ and $q \in \partial B$. There exists a pair of sample points $a \in A$ and $b \in B$, such that $\mathbf{d_S}(p,a) + \mathbf{d_S}(a,b) + \mathbf{d_S}(b,q) \leq (\cos(\beta)/\cos(\theta)) \cdot \|pq\|$.*

4.4 The quality of the path

In Lemma 19, 20, and 21, we have shown that for every segment pq in Case 1-3 in Fact 16, either there exists a path $P \subseteq \mathcal{E}$ such that $\mathbf{w}(P)$ approximates $\|pq\|$, or there exist two sample points a and b such that $\mathbf{d_S}(p,a) + \mathbf{d_S}(a,b) + \mathbf{d_S}(b,q)$ approximates $\|pq\|$. Taking the maximum ratio in the three lemmas, we have the following.

$$\mathbf{d_S}(p,a) + \mathbf{d_S}(a,b) + \mathbf{d_S}(b,q) \leq \frac{2\sin(\frac{\theta}{2})}{\cos(\theta)\sin(\theta)} \cdot \|pq\|$$

For a Case 4 segment $\partial A(p,q)$, where both p and q lies on the obstacle A, assume without loss of generality that p occurs before q in P^*, and let $pq = \partial A(p,q)$.

If $\partial A(p,q)$ contains no sample point, then assume that the optimal path uses segment $p'p$ to reach A, and qq' to leave A. We argue that there exists an approximate path P that approximates $\|p'p\| + \|\partial A(p,q)\| + \|qq'\|$. Let a (resp. b) be the closest sample point to p (resp. q), such that $\partial A(p,q) \subseteq \partial A(a,b)$. In Lemma 19, 20, and 21, we have payed for a path $P_p \subseteq \mathcal{E}$ from p' to p though a and a path $P_q \subseteq \mathcal{E}$ from q to q' through b. Since there is no sample point on $\partial A(p,q)$, instead of going from a to p and q to b, we take the path ab directly. The unused cost of $\mathbf{d_S}(a,p)$ and $\mathbf{d_S}(q,b)$ pays for $\|ab\|$.

Let $p_\perp$ (resp. $q_\perp$) be the orthogonal projection of p (resp. q) on ab. Clearly, $\|\partial A(a,p)\| \geq \|ap_\perp\|$ and $\|\partial A(q,b)\| \geq \|bq_\perp\|$. By Lemma 14, $\|\partial A(p,q)\| \leq \sec(\theta/2) \cdot \|p_\perp q_\perp\|$. Therefore, we connect P_p and P_q using ab to generate a path P, and we have that

$$\|P\| \leq \frac{2\sin(\frac{\theta}{2})}{\cos(\theta)\sin(\theta)} \cdot (\|p'p\| + \|\partial A(p,q)\| + \|qq'\|).$$

If $\partial A(p,q)$ contains at least one sample point $\{a, ..., b\}$, then by Lemma 14, we have that

$$\mathbf{d_S}(p,a) + \mathbf{d_S}(a,b) + \mathbf{d_S}(b,q) \leq \sec\left(\frac{\theta}{2}\right) \cdot \|pq\|.$$

Bose and van Renssen [6] showed that in an environment with polygonal obstacles, the Θ-graph described by Clarkson [10] has a spanning ratio of at most $r_\theta = 1+2\sin(\theta/2)/(\cos(\theta/2)-\sin(\theta/2))$. We also need to apply the factor to traverse the boundaries of convex obstacles to account for the difference compared to the boundaries of simplified obstacles, as in Lemma 14. For a desired approximation factor $1+\varepsilon$, we can compute a $\theta \in O(\varepsilon)$, and thus $\mathcal{B}$ contains a $(1+\varepsilon)$-spanner of the sample points.

▶ **Lemma 22.** *In $\mathcal{B}$, there exists a path $P \subseteq \mathcal{E}$ between any pair of sample points (a,b) such that $\mathbf{w}(P) \leq (1+\varepsilon) \cdot \mathbf{d}(a,b)$.*

Given a pair of query points s and t, we add s and t to $\mathcal{B}$. We treat both s and t as obstacles with no interior to enable previous lemmas, yielding the theorem below. See the full version for details.

▶ **Theorem 3.** *Consider a planar subdivision induced by a plane with a weight of* 1*, consisting of two sets of convex and non-overlapping regions: 0-regions* $\mathcal{Z}$ *with a weight of* 0*, and obstacles* $\mathcal{O}$ *with a weight of* ∞*. Let* $n = |\mathcal{Z}| + |\mathcal{O}|$ *and let* N *denote be the total number of vertices in* $\mathcal{Z} \cup \mathcal{O}$*. For any approximation factor* $0 < \varepsilon < 1$*, a data structure* $\mathcal{B}$ *can be constructed over* $\mathcal{Z} \cup \mathcal{O}$ *in* $O(N + (n/\varepsilon^3)(\log(n/\varepsilon) + \log N))$ *expected time, with a total size of* $O(N + n/\varepsilon^3)$*. When queried with arbitrary points* s *and* t*,* $\mathcal{B}$ *returns a path* P *from* s *to* t *in* $O(N + n/\varepsilon^3 + (n/\varepsilon^2)\log(n/\varepsilon) + (\log N)/\varepsilon)$ *time, ensuring that* $\mathbf{w}(P) \leq (1+\varepsilon) \cdot \mathbf{w}(P^*)$*, where* P^* *is the optimal weighted shortest path from* s *to* t*.*

References

1. Lyudmil Aleksandrov, Anil Maheshwari, and Jörg-Rüdiger Sack. Approximation algorithms for geometric shortest path problems. In *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing*, pages 286–295, New York, NY, USA, May 2000. Association for Computing Machinery. doi:10.1145/335305.335339.
2. Lyudmil Aleksandrov, Anil Maheshwari, and Jörg-Rüdiger Sack. Determining approximate shortest paths on weighted polyhedral surfaces. *Journal of the ACM*, 52(1):25–53, January 2005. doi:10.1145/1044731.1044733.
3. Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *International Journal of Computational Geometry & Applications*, 05:75–91, March 1995. doi:10.1142/S0218195995000064.
4. Prosenjit Bose, Jean-Lou de Carufel, Pat Morin, André van Renssen, and Sander Verdonschot. Towards tight bounds on theta-graphs: more is not always better. *Theoretical Computer Science*, 616:70–93, February 2016. doi:10.1016/j.tcs.2015.12.017.
5. Prosenjit Bose, Guillermo Esteban, David Orden, and Rodrigo I. Silveira. On approximating shortest paths in weighted triangular tessellations. *Artificial Intelligence*, 318:103898, May 2023. doi:10.1016/j.artint.2023.103898.
6. Prosenjit Bose and André van Renssen. Spanning properties of Yao and theta-graphs in the presence of constraints. *International Journal of Computational Geometry & Applications*, 29(02):95–120, June 2019. doi:10.1142/S021819591950002X.
7. Kevin Buchin, Maike Buchin, and Yusu Wang. Exact algorithms for partial curve matching via the Fréchet distance. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 645–654, USA, January 2009. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611973068.71.
8. Kevin Buchin, Tim Ophelders, and Bettina Speckmann. SETH says: weak Fréchet distance is faster, but only if it is continuous and in one dimension. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2887–2901, USA, January 2019. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611975482.179.
9. Siu-Wing Cheng, Jiongxin Jin, and Antoine Vigneron. Triangulation Refinement and Approximate Shortest Paths in Weighted Regions. In *Proceedings of the 2015 Annual ACM-SIAM Symposium on Discrete Algorithms*, Proceedings, pages 1626–1640. Society for Industrial and Applied Mathematics, December 2014. doi:10.1137/1.9781611973730.108.
10. Kenneth L. Clarkson. Approximation algorithms for shortest path motion planning. In *Proceedings of the nineteenth annual ACM Symposium on Theory of Computing*, pages 56–65, New York, NY, USA, January 1987. Association for Computing Machinery. doi:10.1145/28395.28402.
11. Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*. Springer, Berlin, Heidelberg, 2008. doi:10.1007/978-3-540-77974-2.
12. Sarita de Berg, Guillermo Esteban, Rodrigo I. Silveira, and Frank Staals. Exact solutions to the Weighted Region Problem, February 2024. arXiv:2402.12028 [cs]. doi:10.48550/arXiv.2402.12028.

13 Jean-Lou de Carufel, Amin Gheibi, Anil Maheshwari, Jörg-Rüdiger Sack, and Christian Scheffer. Similarity of polygonal curves in the presence of outliers. *Computational Geometry*, 47(5):625–641, July 2014. doi:10.1016/j.comgeo.2014.01.002.

14 Jean-Lou de Carufel, Carsten Grimm, Anil Maheshwari, Megan Owen, and Michiel Smid. A note on the unsolvability of the weighted region shortest path problem. *Computational Geometry*, 47(7):724–727, August 2014. doi:10.1016/j.comgeo.2014.02.004.

15 Laxmi P. Gewali, Alex C. Meng, Joseph S. B. Mitchell, and Simeon Ntafos. Path planning in $0/1/\infty$ weighted regions with applications. In *Proceedings of the Fourth Annual Symposium on Computational Geometry*, pages 266–278, New York, NY, USA, January 1988. Association for Computing Machinery. doi:10.1145/73393.73421.

16 Mark Lanthier, Anil Maheshwari, and Jörg-Rüdiger Sack. Approximating weighted shortest paths on polyhedral surfaces. In *Proceedings of the Thirteenth Annual Symposium on Computational Geometry*, pages 274–283, New York, NY, USA, August 1997. Association for Computing Machinery. doi:10.1145/262839.262984.

17 Joseph S. B. Mitchell and Christos H. Papadimitriou. The weighted region problem: finding shortest paths through a weighted planar subdivision. *Journal of the ACM*, 38(1):18–73, January 1991. doi:10.1145/102782.102784.

18 Zheng Sun and John H. Reif. On finding approximate optimal paths in weighted regions. *Journal of Algorithms*, 58(1):1–32, January 2006. doi:10.1016/j.jalgor.2004.07.004.

19 Victor Andreevich Toponogov. *Differential Geometry of Curves and Surfaces: A Concise Guide*. Birkhauser, 2006.