Navigating Weighted Regions with Scattered Skinny Tetrahedra

Siu-Wing Cheng[*] Man-Kwun Chiu[†] Jiongxin Jin[‡] Antoine Vigneron[§]

Abstract

We propose an algorithm for finding a $(1+\varepsilon)$-approximate shortest path through a weighted 3D simplicial complex $\mathcal{T}$. The weights are integers from the range $[1, W]$ and the vertices have integral coordinates. Let N be the largest vertex coordinate magnitude, and let n be the number of tetrahedra in $\mathcal{T}$. Let ρ be some arbitrary constant. Let κ be the size of the largest connected component of tetrahedra whose aspect ratios exceed ρ. There exists a constant C dependent on ρ but independent of $\mathcal{T}$ such that if $\kappa \leq \frac{1}{C} \log\log n + O(1)$, the running time of our algorithm is polynomial in n, $1/\varepsilon$ and $\log(NW)$. If $\kappa = O(1)$, the running time reduces to $O(n\varepsilon^{-O(1)}(\log(NW))^{O(1)})$.

[*]Department of Computer Science and Engineering, HKUST, Hong Kong. Research supported by Research Grant Council, Hong Kong, China (project no. 611812).

[†]National Institute of Informatics (NII), Tokyo, Japan and JST, ERATO, Kawarabayashi Large Graph Project.

[‡]Google Inc.

[§]School of Electrical and Computer Engineering, Ulsan National Institute of Science and Technology, Ulsan, Korea.

1 Introduction

Finding shortest paths in a geometric environment is a classical optimization problem in computational geometry (e.g. [7, 8, 14, 15, 18, 19, 20, 21, 25, 26, 28, 29, 31]). In 2D and terrains, researchers have also studied cost models in applications that are non-L_p metrics and anisotropic (e.g.[1, 2, 3, 4, 5, 9, 10, 11, 12, 13, 22, 23, 27, 32, 33]). In 3D, other than motion planning, shortest path is a popular tool for simulating seismic raytracing in ray-based tomography schemes for studying some geological properties, and the time required to traverse different regions may vary (e.g. [17, 24]).

The weighted region problem is a way to model the unequal difficulties in traversing different regions [27]. In 3D, we are given a simplicial complex $\mathcal{T}$ of n tetrahedra. These tetrahedra and their vertices, edges and triangles are called the *simplices* of $\mathcal{T}$. Given two simplices in $\mathcal{T}$, either they are disjoint or their intersection is another simplex in $\mathcal{T}$. Every vertex has integral coordinates and let N denote the largest vertex coordinate magnitude. Each tetrahedron τ is associated with an integral weight $\omega_\tau \in [1, W]$. For every edge or triangle, its weight is equal to the minimum weight among the tetrahedra incident to that edge or triangle. The cost of a path that lies in a simplex σ is equal to the path length multiplied by ω_σ. Given a path P in $\mathcal{T}$, we denote its length by $\|P\|$ and its cost by $\mathrm{cost}(P) = \sum_{\text{simplex } \sigma} \omega_\sigma \|P \cap \sigma\|$. The weighted region problem is to find the least-cost path from a given source vertex to a given destination vertex.

The weighted region problem in 2D has been studied extensively. Fully polynomial time approximation schemes are known [11, 27]. There are also successful discretization schemes whose running time is linear in the input size and dependent on some geometric parameter of the polygonal domain [5, 33]. In contrast, only one algorithm for the weighted region problem in 3D has been proposed (Aleksandrov et al. [6]). The authors present a $(1+\varepsilon)$-approximation algorithm whose running time is $O\left(Kn\varepsilon^{-2.5} \log \frac{n}{\varepsilon} \log^3 \frac{1}{\varepsilon}\right)$, where K is asymptotically at least the cubic power of the maximum aspect ratio of the tetrahedra in the worst case. (Aspect ratio is defined in Section 2.) It is an open problem whether an FPTAS exists for the 3D weighted region problem. Our contribution is an algorithm whose running time depends on combinatorial parameters only and it is in fact polynomial when the tetrahedra with large aspect ratios are scattered.

Let ρ be an arbitrary constant independent of $\mathcal{T}$. We call a tetrahedron *skinny* if its aspect ratio exceeds ρ. Two skinny tetrahedra are *connected* if their boundaries touch, and the transitive closure of this relation gives the connected components of skinny tetrahedra. Let κ be the number of tetrahedra in the largest connected component of skinny tetrahedra.

We present a $(1+\varepsilon)$-approximation algorithm for the 3D weighted region problem. It runs in $O(2^{2^{O(\kappa)}} n\varepsilon^{-7} \log^2 \frac{W}{\varepsilon} \log^2 \frac{NW}{\varepsilon})$ time. The hidden constant in the exponent $O(\kappa)$ is dependent on ρ but independent of $\mathcal{T}$. Thus, there exists a constant C dependent on ρ but independent of $\mathcal{T}$ such that if $\kappa \leq \frac{1}{C} \log\log n + O(1)$, the running time is polynomial in n, $1/\epsilon$ and $\log(NW)$. If $\kappa = O(1)$, the running time is linear in n. In comparison, the running time in [6] has the advantage of being independent from N and W, but K can be arbitrarily large even if there are only $O(1)$ skinny tetrahedra. Putting the result in [6] in our model, K is a function of N and n in the worst case, and K can be $\Omega(\frac{1}{n}N^3 + 1)$.

2 Preliminaries

A path P in $\mathcal{T}$ consists of *links* and *nodes.* A link is a maximal segment that lies in a simplex of $\mathcal{T}$. Nodes are link endpoints. We assume that P does not bend in the interior of any simplex because such a bend can be shortcut. So the nodes of P lie at vertices, edges and triangles.

Given two points x and y in this order in P, we use $P[x, y]$ to denote the subpath between them.

The *simplex sequence* of a path P is the ordered sequence Σ of vertices, edges and triangles that intersect the interior of P from u to v. If P has the minimum cost among all paths from u to v with simplex sequence Σ, we call P a *locally shortest path* (with respect to Σ). The shortest path from u to v is the locally shortest path with the minimum cost among all possible simplex sequences.

Let $B(x, r)$ denote a closed ball centered at a point x with radius r.

The aspect ratio of a tetrahedron τ is the ratio of the radius of the smallest sphere that encloses τ to the radius of the largest sphere inscribed in τ. If the aspect ratio is bounded by a constant, all angles of τ are bounded from below and above by some constants. A tetrahedron is *skinny* if its aspect ratio exceeds some arbitrary constant ρ fixed *a priori*. If a tetrahedron is not skinny, it is *fat*.

Two tetrahedra are *connected* if their boundaries touch. The equivalence classes of the transitive closure of this relation are called *connected components* of tetrahedra. Two tetrahedra are *edge-connected* if they share at least one edge. The equivalence classes of the transitive closure of this relation are called *edge-connected components* of tetrahedra. A *cluster* is a connected component of skinny tetrahedra. By definition, every cluster has at most κ tetrahedra.

For every simplex σ in $\mathcal{T}$, $\mathrm{star}(\sigma)$ denotes the set of tetrahedra that have σ as a boundary simplex. Given a set $\mathcal{K}$ of simplices, $|\mathcal{K}|$ denotes the union of all simplices in $\mathcal{K}$ and $\mathrm{bd}(\mathcal{K})$ denotes the set of simplices in the boundary of $|\mathcal{K}|$.

For simplicity, we will show a $1 + O(\varepsilon)$ approximation ratio, which can be reduced to $1 + \varepsilon$ by tuning some constants. Our algorithm discretizes $\mathcal{T}$ and builds an edge-weighted graph $\mathcal{G}$ so that the shortest path in $\mathcal{G}$ is a $1 + O(\varepsilon)$ approximation. This approach is also taken in [6] in 3D. Unlike in [6], in order to allow for skinny tetrahedra, we discretize the fat tetrahedra only, and the edges in $\mathcal{G}$ represent approximate shortest paths that may not lie within a single tetrahedron. This also leads to a very different analysis for obtaining the approximation ratio of $1 + \varepsilon$.

Let $\{u, v\}$ be a pair of vertices of $\mathcal{G}$. If u and v lie in a cluster, we would ideally connect them by an edge with weight equal to the shortest path cost between u and v within the cluster. However, even if a simplex sequence is given, finding the locally shortest path requires solving a nonlinear system derived using Snell's law. It is unclear how to do this exactly. Instead, we switch to convex distance functions induced by convex polytopes with $O(1/\varepsilon)$ vertices, so that the modified metrics give $1 + O(\varepsilon)$ approximations of the original metrics. Under the modified metrics, the locally shortest path with respect to Σ can be obtained by linear programming. We enumerate all possible simplex sequences to find the shortest path cost within the cluster under the modified metrics.

3 Placement of Steiner points

For every vertex v in $\mathcal{T}$, the fat tetrahedra in $\mathrm{star}(v)$ may form multiple edge-connected components and we call each a *fat substar*. For an edge or triangle σ, there is at most one fat substar in $\mathrm{star}(\sigma)$.

Definition 1 *Let x be a point in the union of vertices, edges and triangles of $\mathcal{T}$. Let σ be the simplex of the lowest dimension containing x. For every fat substar F of σ, define $\delta_F(x)$ to be the minimum distance from x to a simplex in $\mathrm{bd}(F)$ that does not contain x. When σ is an edge or triangle, there is at most one fat substar of σ and so we simplify the notation to $\delta(x)$. Figure 1 shows some examples.*

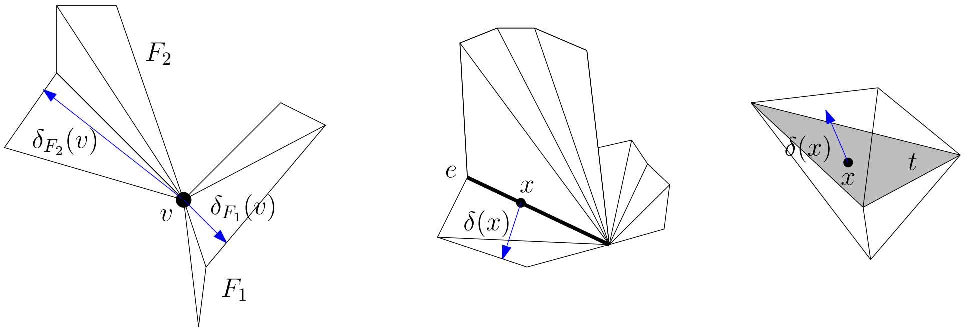

Figure 1: Examples of fat substars and $\delta_F(x)$.

Remark 1: For a vertex v of $\mathcal{T}$, $\delta_F(v)$ is the distance between v and a triangle opposite v in some tetrahedron $\tau \in F$. Since the tetrahedra in F have bounded aspect ratio and there are $O(1)$ of them, $\delta_F(v) = \Theta(\|e\|)$ for every edge $e \in F$.

Remark 2: For a point x in the interior of an edge e, $\delta(x)$ is the distance between x and an edge or triangle σ such that σ bounds a fat tetrahedron incident to e and σ shares only a vertex v with e. Thus, $\delta(x) = \|vx\| \sin\theta$, where θ is the angle between e and σ, which is bounded from below by a constant. So $\delta(x) = \Theta(\|vx\|)$.

Remark 3: Consider the edge-ball B_q placed at the intersection q between uv and the boundary of N_u. By Remark 1, $\|qu\| = \frac{\varepsilon}{3W}\delta_F(u) = \Theta(\frac{\varepsilon}{W}\|uv\|)$, where F is the fat substar of u that contains uv. By definition, $\delta(q) \leq \|qu\| = O(\frac{\varepsilon}{W}\|uv\|)$. By Remark 2, $\delta(q) = \Omega(\min\{\|qu\|, \|qv\|\}) = \Omega(\min\{\frac{\varepsilon}{W}\|uv\|, (1-\frac{\varepsilon}{W})\|uv\|\}$. So $\delta(q) = \Theta(\frac{\varepsilon}{W}\|uv\|)$. The radius of B_q is $\frac{\varepsilon}{3}\delta(q) = \Theta(\frac{\varepsilon^2}{W}\|uv\|)$.

For every vertex v of $\mathcal{T}$ and every fat substar F of v, define a *vertex-ball* $B_{v,F} = B(v, \frac{\varepsilon}{3W}\delta_F(v))$. Let N_v be the union of $B_{v,F} \cap F$ over all fat substars F.

Let uv be an edge of a fat tetrahedron in $\mathcal{T}$. We place Steiner points in uv outside N_u and N_v as follows. Initialize $\mathcal{B}$ to be the union of the interiors of N_u and N_v. Find the point $p \in uv \setminus \mathcal{B}$ such that $\delta(p)$ is maximum. Make p a Steiner point. Define an *edge-ball* $B_p = B(p, \frac{\varepsilon}{3}\delta(p))$. Add the interior of B_p to $\mathcal{B}$. Repeat until $uv \setminus \mathcal{B}$ is empty. Finally, make the intersection point q between uv and the boundary of N_u a Steiner point and introduce an edge-ball $B_q = B(q, \frac{\varepsilon}{3}\delta(q))$. Repeat the same for the intersection point between uv and the boundary of N_v.

As we will see below, the edge-balls centered at two consecutive Steiner points strictly outside N_u and N_v overlap significantly. After placing Steiner points strictly outside N_u and N_v, an extreme edge-ball may have a tiny overlap with N_u or N_v. In this case, if x is a point on some triangle incident to uv such that x lies close to this tiny overlap, then $\delta(x)$ can be arbitrarily small. This will cause a problem in discretizing triangles. Thus, we place two more edge-balls at the intersection points between uv and the boundaries of N_u and N_v. Figure 2 shows an example.

Lemma 3.1 *Let uv be an edge of a fat tetrahedron. The edge uv is covered by the union of N_u, N_v, and the edge-balls centered at the Steiner points in uv. For every consecutive pair of Steiner points $p, q \in uv$ strictly outside N_u and N_v,*

(i) $\|pq\| = \frac{\varepsilon}{3} \cdot \max\{\delta(p), \delta(q)\}$, and

(ii) if $\delta(p) \geq \delta(q)$, then q lies on the boundary of B_p; otherwise, p lies on the boundary of B_q.

There are $O\left(\frac{1}{\varepsilon}\log\frac{W}{\varepsilon}\right)$ Steiner points in uv.

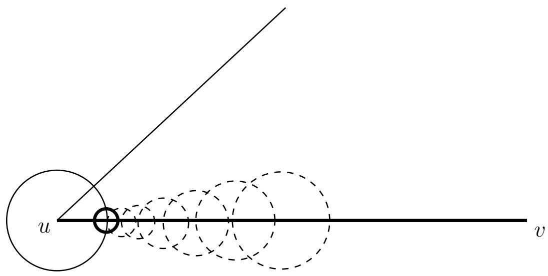

Figure 2: Placement of Steiner points on an edge uv. The solid circle centered at u denotes a vertex ball at u. The dashed circles denote some edge-balls placed with their centers in the interior of uv. The bold circle centered at the intersection of uv and the solid circle is the edge-ball placed at the intersection of uv and the boundary of N_u.

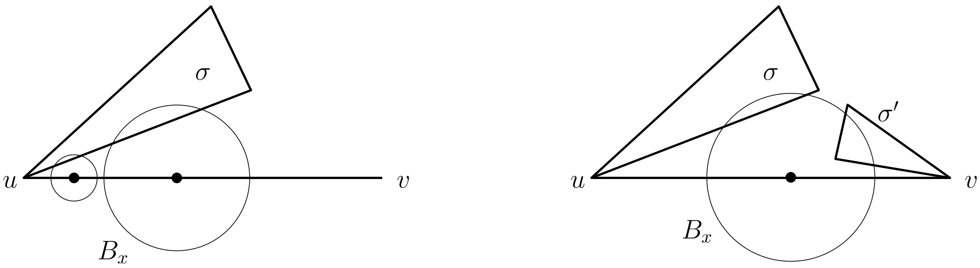

Figure 3: As x moves from u towards v, the radius $\delta(x)$ of the ball B_x is determined by an edge or triangle σ incident to u. So $\delta(x)$ grows linearly in size, and $\delta(x)$ reaches the maximum when it touches another edge or triangle σ' incident to v.

Proof. The construction ensures the coverage of uv. Consider moving a point x from u along the edge uv towards v. When x is arbitrarily close to u, $\delta(x)$ is determined by some edge or triangle σ incident to u, that is, $\delta(x)$ is the distance from x to σ. Therefore, if we draw a ball B_x centered at x with radius $\delta(x)$, B_x touches σ and does not intersect any other simplex in $\text{star}(uv)$ that is not incident to the interior of uv. Figure 3 shows an example. As x moves towards v, σ remains the closest simplex to x among the edges and triangles incident to u. Therefore, the radius $\delta(x)$ of B_x grows linearly, and B_x remains disjoint from any other edge or triangle that is incident to u or v. At some moment during the movement of x, B_x comes into contact with another edge or triangle σ' incident to v. At this moment, $\delta(x)$ reaches its maximum. When x moves further towards v, σ' determines $\delta(x)$ instead and $\delta(x)$ decreases linearly. In summary, $\delta(x)$ increases linearly from a limit of zero at u and then decreases linearly to a limit of zero at v.

The placement of Steiner points strictly outside N_u and N_v begins with the point $p \in uv$ that maximizes $\delta(p)$. Therefore, the point $q \in uv$ that maximizes $\delta(q)$ outside the interiors of N_u, N_v, and B_p must lie on the boundary of B_p. Repeating this argument establishes properties (i) and (ii) in the lemma.

Let $p, q \in uv$ be two consecutive Steiner points strictly outside N_u and N_v such that $\delta(x)$ increases linearly from a limit of zero from u to p and then to q. By Remark 2, $\delta(p) = \Theta(\|pu\|)$. We have shown in property (i) that $\|pq\| \geq \frac{\varepsilon}{3}\delta(p)$. By the linear increase in $\delta(\cdot)$, we get $\delta(q) = (1 + \|pq\|/\|pu\|)\delta(p) \geq (1 + \Theta(\varepsilon))\delta(p)$. The next Steiner point after q is thus at distance at least $\frac{\varepsilon}{3}\delta(q) \geq \frac{\varepsilon}{3}(1 + \Theta(\varepsilon))\,\delta(p)$ from q. In other words, the distance between consecutive Steiner points strictly outside N_u and N_v increases repeatedly by at least a factor $1 + \Theta(\varepsilon)$. By Remark 3, the distance between two consecutive Steiner points that are strictly outside N_u is $\Omega(\frac{\varepsilon^2}{W}\|uv\|)$ near u. Then, it increases to at most $\frac{\varepsilon}{6}\|uv\|$ in the interior of uv. The same holds for the sequence of Steiner points from N_v. Hence, there are $O(\log_{1+\Theta(\varepsilon)} \frac{W}{\varepsilon}) = O\left(\frac{1}{\varepsilon} \log \frac{W}{\varepsilon}\right)$ Steiner points. ◻

Lemma 3.2 *Placing Steiner points on an edge takes $O\left(\frac{1}{\varepsilon} \log \frac{W}{\varepsilon}\right)$ time.*

Proof. The placement of Steiner points on an edge uv begins with finding the interior point p such that $\delta(p)$ is maximum. There are $O(1)$ tetrahedra in the fat substar F of uv, so there are $O(1)$ edges and triangles in F that are incident to u or v. In $O(1)$ time, we can determine the edge or triangle σ that determines $\delta(x)$ for a point $x \in uv$ near u. Similarly, we can determine the edge or triangle σ' that determines $\delta(x)$ for a point $x \in uv$ near v. As discussed in the proof of Lemma 3.1, σ and σ' determine the maximum $\delta(p)$.

Finding the exact location of $p^* \in uv$ that maximizes $\delta(\cdot)$ involves computing the medial axis of σ and σ' and then intersecting this medial axis with uv. Medial axis computation is rather complex. Instead, it suffices to find an approximate location $\tilde{p} \in uv$ such that the edge-ball $B_{\tilde{p}} = B(\tilde{p}, \frac{\varepsilon}{3}\delta(\tilde{p}))$ contains p^*. Then, the function $\delta(\cdot)$ still decreases linearly in the two intervals $uv \setminus B_{\tilde{p}}$ towards u and v, respectively, and so the proof of Lemma 3.1 still applies. We find such a point $\tilde{p}$ by binary search. We begin with the midpoint q of uv. If q is closer to σ than σ', we recurse on the interval closer to v; otherwise, we recurse on the interval closer to u. The recursion stops when the interval length is less than $\frac{\varepsilon}{3}\delta(q)$, q being the interval midpoint, and we place the edge-ball $B_q = B(q, \frac{\varepsilon}{3}\delta(q))$. Since the binary search keeps the point p^* that maximizes $\delta(\cdot)$ in the interval, the last q is the desired $\tilde{p}$ and $B_{\tilde{p}}$ contains p^*.

By Remark 2, $\delta(p^*) = \Theta(\|p^*u\|) = \Theta(\|p^*v\|)$, which implies $\delta(p^*) = \Theta(\|uv\|)$. Since $B_{\tilde{p}}$ touches an edge or triangle incident to u or v, $B_{\tilde{p}}$ cannot be contained inside $B(p^*, \frac{\varepsilon}{3}\delta(p^*))$. So the radius $\frac{\varepsilon}{3}\delta(\tilde{p})$ of $B_{\tilde{p}}$ is at least $\frac{\varepsilon}{3}\delta(p^*) - \|p^*\tilde{p}\|$. Note that $\|p^*\tilde{p}\|$ is at most half of the length of the last interval. Thus, $\frac{\varepsilon}{3}\delta(\tilde{p}) \geq \frac{\varepsilon}{3}\delta(p^*) - \|p^*\tilde{p}\| \geq \frac{\varepsilon}{3}\delta(p^*) - \frac{\varepsilon}{6}\delta(\tilde{p})$. It follows that $\delta(\tilde{p}) \geq \frac{2}{3}\delta(p^*) = \Omega(\|uv\|)$. The binary search takes $O(1)$ time per probe. The initial interval is uv, and the recursion stops when the interval length is less than $\frac{\varepsilon}{3}\delta(\tilde{p}) = \Omega(\varepsilon\|uv\|)$. Therefore, there are $O(\log \frac{1}{\varepsilon})$ probes, meaning that we can place the first edge-ball in $O(\log \frac{1}{\varepsilon})$ time.

Subsequently, since $\delta(x)$ increases and then decreases from one endpoint of an edge to the other endpoint, there are at most two gaps on the edge to be covered during the placement of Steiner points. Therefore, after placing the first edge-ball, it takes $O(1)$ time to place each Steiner point subsequently. ◻

Let t be a triangle of a fat tetrahedron. Since the vertex-balls and edge-balls on an edge overlap significantly, we can show that for every point x in t that lies outside the vertex-balls and edge-balls in the boundary of t, the point x is relatively far from the boundary of t.

Lemma 3.3 *There exists a constant $c > 0$ such that for every triangle t of a fat tetrahedron and for every point $x \in t$ that lies outside vertex-balls and edge-balls in the boundary of t, the distance between x and the boundary of t is at least $c\varepsilon^2\ell/W$, where ℓ is the longest edge length of t.*

Proof. The point x is closest to the boundary of t when x lies at the intersection of $\partial B_p \cap \partial B_q \cap t$, where p and q are two consecutive Steiner points or vertices on some edge e of t. There are three cases depending on the locations of p and q.

Case 1: p and q are two Steiner points, and neither p nor q lies on the boundaries of the vertex-balls at the endpoints of e. Without loss of generality, assume that $\delta(p) \geq \delta(q)$, so q lies on the boundary of B_p by Lemma 3.1. As shown in the left image in Figure 4, the distance between x and the edge e is $\frac{\varepsilon}{3}\delta(q) \cdot \sin\left(\arccos \frac{\delta(q)}{2\delta(p)}\right) = \delta(q) \cdot \frac{\varepsilon}{3}(\delta(p)^2 - \frac{1}{4}\delta(q)^2)^{1/2}/\delta(p) \geq \frac{\varepsilon}{2\sqrt{3}}\delta(q)$.

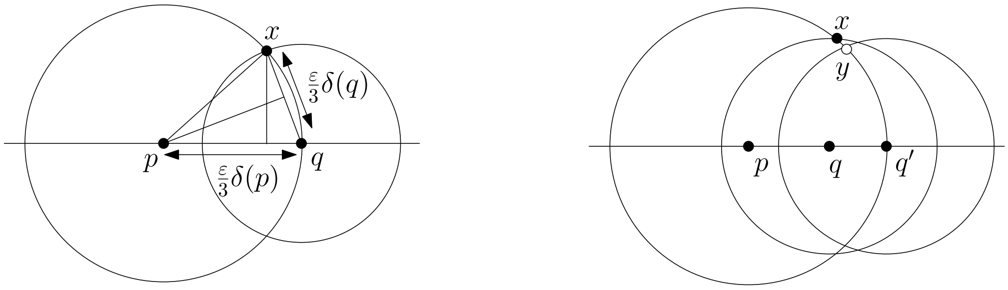

Figure 4: In the left figure, the distance of x from the edge is $\delta(q)\cdot\frac{\varepsilon}{3}(\delta(p)^2-\frac{1}{4}\delta(q)^2)^{1/2}/\delta(p)$. In the right figure, the point x is farther from the edge than the white dot which is the intersection $\partial B_p \cap \partial B' \cap t$.

Since q is outside the vertex-balls at the endpoints of e, by Remarks 1 and 2, we have $\delta(q) = \Omega(\frac{\varepsilon}{W}\|e\|)$, which is $\Omega(\varepsilon\ell/W)$ because because all angles of t are bounded from above and below by some constants. Therefore, the distance between x and e is $\Omega(\varepsilon^2\ell/W)$.

Case 2: p is a vertex-ball and q is a Steiner point in the boundary of B_p. Using the analysis in Case 1, the distance between x and e is at least $\frac{\varepsilon}{2\sqrt{3}}\delta(q)$, which is $\Omega(\varepsilon^2\ell/W)$ as $\delta(q) = \Omega(\frac{\varepsilon}{W}\|e\|)$ by Remark 3.

Case 3: p and q are two Steiner points, and q lies on the boundary of the vertex-ball at an endpoint u of e. By construction, B_p must overlap with the vertex-ball at u, and therefore, $q \in B_p$. Also, $\delta(p) > \delta(q)$ as $\delta(\cdot)$ decreases linearly from p towards u. Let q' be intersection of e and the boundary of B_p that is closer to u. Refer to the right image in Figure 4. Let B' be a copy of B_q with its center at q'. Let x and y be the intersection points $\partial B_p \cap \partial B_q \cap t$ and $\partial B_p \cap \partial B' \cap t$, respectively. Clearly, the distance between from x to e is greater than the distance from y to e. Using the analysis in Case 1, the distance between y and e is at least $\frac{\varepsilon}{2\sqrt{3}}\delta(q)$, which is $\Omega(\varepsilon^2\ell/W)$ as $\delta(q) = \Omega(\frac{\varepsilon}{W}\|e\|)$ by Remark 3. □

The placement of Steiner points in a triangle uvw of a fat tetrahedron is slightly more involved. In the interior of uvw, the value of $\delta(x)$ is determined by the triangles of at most two fat tetrahedra incident to uvw. Consider one triangle t out of these candidates. Orient space so that uvw is horizontal. The graph of the distance function from x to t is a plane that makes an angle arctan$(\sin\theta)$ with the horizontal, where θ is the dihedral angle between t and uvw (which is bounded from below and above by some constants). The graph of $\delta(x)$ is thus a lower envelope of planes. Moreover, this lower envelope H is supported by exactly three planes induced by three triangles that share with uvw the edges uv, vw and uw. Let ℓ denote the longest edge length of uvw. The maximum height of H is $h_{\max} = \Theta(\ell)$ as the tetrahedra defining $\delta(x)$ have bounded aspect ratios. For each point x in the interior of uvw that are close to and outside the vertex-balls and edge-balls at the boundary of uvw, by Lemma 3.3, $\delta(x) \geq c\varepsilon^2\ell/W$ for some constant $c > 0$. Let H^+ denote the portion of H at height $h_{\min} = c\varepsilon^2\ell/W^2$ or above. We will place Steiner points in the projection of H^+ in uvw. By the geometry of H, a cross-section of H bounds a triangle that has the same angles as uvw and projects to the interior of uvw.

Define $h_0 = h_{\max}$ and for $i \geq 1$, $h_i = h_{i-1}/(1+\varepsilon)$. Let $A_i \subset uvw$ be the triangular annulus that the portion of H between heights h_i and h_{i+1} projects to. Both the inner and outer boundaries of this annulus are similar to uvw. The area of A_i is $\Theta((h_i - h_{i+1})(h_i + h_{i+1})) = \Theta(\varepsilon h_i^2)$. We place Steiner points in each A_i as follows. Initialize $\mathcal{B} = \emptyset$. Make an arbitrary point $p \in A_i \setminus \mathcal{B}$ a Steiner point. Define a *triangle-ball* $B_p = B(p, \frac{\varepsilon}{3}\delta(p))$. Add the interior of B_p to $\mathcal{B}$. Repeat until $A_i \setminus \mathcal{B}$ is empty.

Lemma 3.4 *Let uvw be a triangle of a fat tetrahedron. The triangle uvw is covered by* *the union of N_u, N_v, N_w, and edge-balls and triangle-balls with centers in uvw. There are $O\left(\frac{1}{\varepsilon^2}\log\frac{W}{\varepsilon}\right)$ Steiner points in uvw.*

Proof. The construction ensures the coverage of uvw. For every pair of Steiner points p and q in A_i, q lies outside B_p or p lies outside B_q, depending on whether p or q was placed first. Therefore, $\|pq\| \geq \frac{\varepsilon}{3} \cdot \min\{\delta(p), \delta(q)\}$. The value of $\delta(x)$ in A_i is between h_i and h_{i+1}. Therefore, if we place disks of radii $\frac{\varepsilon}{6}h_{i+1}$ centered at the Steiner points in A_i, the disks are disjoint. At least a constant fraction of each such disk lies inside A_i. Therefore, there are $O(\varepsilon h_i^2/(\varepsilon^2 h_{i+1}^2)) = O(1/\varepsilon)$ Steiner points in A_i. As i increases, h_i decreases and approaches $h_{\min} = \Theta(\varepsilon^2 h_{\max}/W^2)$. Observe that $h_i = (1+\varepsilon)^{-i}h_{\max}$. Hence, $(1+\varepsilon)^{-i}h_{\max} \geq h_{\min}$, which implies that $i = O\left(\log_{1+\varepsilon}\frac{W}{\varepsilon}\right) = O\left(\frac{1}{\varepsilon}\log\frac{W}{\varepsilon}\right)$. It follows that there are $O\left(\frac{1}{\varepsilon^2}\log\frac{W}{\varepsilon}\right)$ Steiner points in uvw. ◻

Lemma 3.5 *Placing Steiner points in uvw takes $O(\frac{1}{\varepsilon^4}\log\frac{W}{\varepsilon})$ time.*

Proof. While placing Steiner points in A_i, we punch holes in A_i and the next Steiner point is identified outside the holes. This can be done by constructing an arrangement of disks and computing the depths of the arrangement cells along the way. Any point in a cell with depth zero can be picked as the next Steiner point. As argued in the proof of Lemma 3.4, there are $O(1/\varepsilon)$ holes, so the arrangement has $O(1/\varepsilon^2)$ complexity and can be constructed in $O(1/\varepsilon^2)$ time, meaning that we spend $O(1/\varepsilon^3)$ time per A_i. Hence, the total time is $O(\frac{1}{\varepsilon^4}\log\frac{W}{\varepsilon})$. ◻

It may be possible to improve the time complexity stated in Lemma 3.5 by maintaining the arrangement of disks incrementally, instead of rebuilding from scratch whenever a new Steiner point is placed. However, the time complexity in Lemma 3.5 is not a bottleneck in the entire algorithm.

4 Steiner graph and snapping

The vertices of $\mathcal{T}$ and the Steiner points form the vertices of $\mathcal{G}$. Before defining the edges of $\mathcal{G}$, we first define *extended clusters*. An extended cluster C^* consists of the skinny tetrahedra in a cluster C and the tetrahedra in contact with C. The tetrahedra in $C^* \setminus C$ are fat, and therefore, there are $O(\kappa)$ tetrahedra in C^*. If a boundary simplex σ of C^* is in contact with the boundary of C, then σ must also be a boundary simplex of $\mathcal{T}$.

There are two kinds of edges in $\mathcal{G}$. Each edge of the first kind connects two graph vertices x and y in the same extended cluster C^*. The edge weight is $1+O(\varepsilon)$ times the shortest path cost in C^* from x to y. We will show in Section 5 how to compute such an edge weight. Each edge of the second kind connects two graph vertices in a vertex star free of skinny tetrahedra. The edge weight is $1+O(\varepsilon)$ times the shortest path cost in that vertex star, which can also be computed by the method in Section 5. Notice that $\mathcal{T}$ is covered by the extended clusters and vertex stars free of skinny tetrahedra. Due to the overlap among extended clusters and vertex stars, we may construct multiple edges between two graph vertices, and if so, we keep the edge between them with the lowest weight.

Assuming that $\mathcal{G}$ has been computed, we prove below that a shortest path in $\mathcal{G}$ is a $(1+O(\varepsilon))$-approximate shortest path in $\mathcal{T}$. We need three technical lemmas (Lemmas 4.1, 4.2, and 4.3) that snap a path to vertices and Steiner points.

Lemma 4.1 *Let v be a vertex of a fat tetrahedron. Let F be a fat substar of v. Let x be a point in $|F|$ such that $\|vx\| \geq \delta_F(v)/2$. Let P be a path such that a subpath of P in $|F|$ connects x*

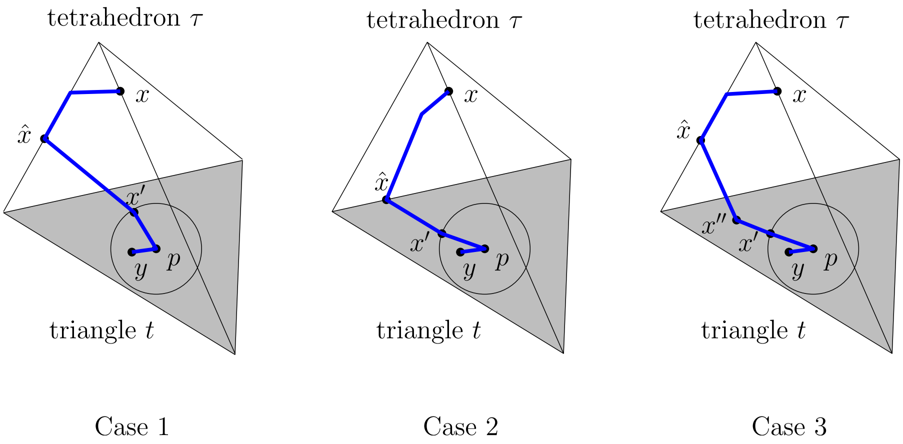

Figure 5: Three cases in the proof of Lemma 4.2.

to a point $y \in B_{v,F}$. We can convert $P[x,y]$ to a path Q from x to y so that $Q \subset |F|$, Q passes through v, and $\text{cost}(Q) \leq (1+O(\varepsilon)) \cdot \text{cost}(P[x,y])$.

Proof. Since $\|vx\| \geq \delta_F(v)/2$, $x \notin B_{v,F}$. Let x' be the first entry point of $P[x,y]$ into $B_{v,F}$. We replace $P[x,y]$ by $P[x,x'] \cup x'v \cup vy$. Observe that $\|xx'\| \geq \|vx\| - \|vx'\| \geq \frac{1}{2}\delta_F(v) - \frac{\varepsilon}{3W}\delta_F(v) \geq \frac{3-2\varepsilon}{6}\delta_F(v)$. Therefore, $\text{cost}(x'v) \leq W\|x'v\| = \frac{\varepsilon}{3}\delta_F(v) \leq \frac{2\varepsilon}{3-2\varepsilon}\|xx'\| \leq O(\varepsilon) \cdot \text{cost}(P[x,x']) \leq O(\varepsilon) \cdot \text{cost}(P[x,y])$. Also, $\text{cost}(vy) \leq W\|vy\| \leq \frac{\varepsilon}{3}\delta_F(v)$, which is at most $O(\varepsilon) \cdot \text{cost}(P[x,y])$ by the analysis above. Therefore, $\text{cost}(P[x,x'] \cup x'v \cup vy) \leq \text{cost}(P[x,x']) + O(\varepsilon) \cdot \text{cost}(P[x,y]) \leq (1+O(\varepsilon)) \cdot \text{cost}(P[x,y])$. ◻

Lemma 4.2 *Let t be a triangle of a fat tetrahedron τ. Let p be a Steiner point in the interior of t, and let B_p denote the triangle-ball centered at p. Let P be a path such that a subpath of P in τ connects a point x in a boundary simplex of τ other than t to a point $y \in B_p \cap t$. We can convert $P[x,y]$ to a path Q from x to y so that $Q \subset \tau$, Q passes through p, and* $\text{cost}(Q) \leq (1+O(\varepsilon)) \cdot \text{cost}(P[x,y])$.

Proof. $P[x,y] \subset \tau$ by assumption. Let x' be the last entry point of $P[x,y]$ into B_p. We replace $P[x,y]$ by $P[x,x'] \cup x'p \cup py$. To analyze the cost of $P[x,x'] \cup x'p \cup py$, retrace $P[x,x']$ from x' towards x until we hit a boundary simplex of τ other than t for the first time at a point $\hat{x}$. Note that $\delta(p) \leq \|p\hat{x}\|$. So $\|\hat{x}x'\| \geq \|p\hat{x}\| - \|px'\| \geq \delta(p) - \frac{\varepsilon}{3}\delta(p) = \frac{3-\varepsilon}{3}\delta(p)$. There are three cases illustrated in Figure 5.

Case 1: $P[\hat{x},x']$ is a segment whose interior lies in the interior of τ. We have $\text{cost}(x'p) \leq \frac{\varepsilon}{3}\omega_\tau\delta(p) \leq \frac{\varepsilon}{3-\varepsilon}\omega_\tau\|\hat{x}x'\| \leq O(\varepsilon) \cdot \text{cost}(P[x,x']) \leq O(\varepsilon) \cdot \text{cost}(P[x,y])$. Similarly, $\text{cost}(py) \leq \frac{\varepsilon}{3}\omega_t\delta(p) \leq \frac{\varepsilon}{3}\omega_\tau\delta(p) \leq O(\varepsilon) \cdot \text{cost}(P[x,y])$.

Case 2: $P[\hat{x},x']$ is a segment whose interior lies in the interior of t. Then the interior of $P[\hat{x},y]$ lies in the interior of t. The analysis is similar to that in Case 1. We have $\text{cost}(x'p) = \frac{\varepsilon}{3}\omega_t\delta(p) \leq \frac{\varepsilon}{3-\varepsilon}\omega_t\|\hat{x}x'\| \leq O(\varepsilon) \cdot \text{cost}(P[x,x']) \leq O(\varepsilon) \cdot \text{cost}(P[x,y])$. Similarly, $\text{cost}(py) \leq \frac{\varepsilon}{3}\omega_t\delta(p) \leq O(\varepsilon) \cdot \text{cost}(P[x,y])$.

Case 3: $P[\hat{x},x']$ consists of two segments $\hat{x}x''$ and $x''x'$ whose interiors lie in the interiors of τ and t, respectively. Then the interior of $P[x'',y]$ lies in the interior of t. If $\|\hat{x}x''\| \geq \frac{1}{2}\|\hat{x}x'\|$, then we adapt the analysis in Case 1 using the relation $\delta(p) \leq \frac{3}{3-\varepsilon}\|\hat{x}x'\| \leq \frac{6}{3-\varepsilon}\|\hat{x}x''\|$. Otherwise, $\|x''x'\| \geq \frac{1}{2}\|\hat{x}x'\|$ and we adapt the analysis in Case 2 using the relation $\delta(p) \leq \frac{3}{3-\varepsilon}\|\hat{x}x'\| \leq \frac{6}{3-\varepsilon}\|x''x'\|$. ◻

Lemma 4.3 *Let e be an edge of a fat tetrahedron. Let F denote the fat substar of e. Let p be a Steiner point in the interior of e, and let B_p denote the edge-ball centered at p. Let x be a point in $|F|$ such that $\|px\| \geq \delta(p)/2$. Let P be a path such that a subpath of P in $|F|$ connects x to a point $y \in B_p \cap t$, where t is a triangle in F incident to e. Suppose that y lies outside every triangle-ball B_q where $q \in t$. Then, we can convert $P[x, y]$ to a path Q from x to y so that $Q \subset |F|$, Q passes through p, and $\mathrm{cost}(Q) \leq (1 + O(\varepsilon)) \cdot \mathrm{cost}(P[x, y])$.*

Proof. Since y lies outside every triangle-ball B_q where $q \in t$, y is at distance $O(\frac{\varepsilon^2}{W^2}\|e\|)$ from e. Let y' be the closest point in e to y. Since $\|px\| \geq \delta(p)/2$, $x \notin B_p$. Let x' be the first entry point of $P[x, y]$ into B_p.

Let σ be the triangle or tetrahedron with the minimum weight among those incident to e and visited by $P[x, x']$. Suppose that $P[x, x']$ enters σ for the first time at a point a. We replace $P[x, y]$ by $P[x, a] \cup ap \cup py' \cup yy'$. Figure 6 illustrates the conversion.

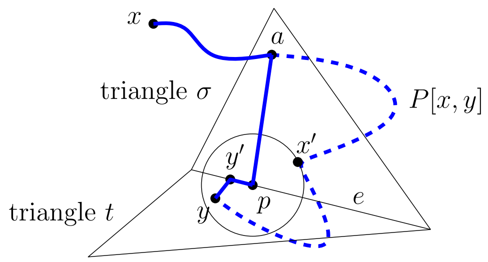

Figure 6: The resulting path Q in the proof of Lemma 4.3.

The subpath $P[a, x']$ is contained in $|F| \subseteq \mathrm{star}(e)$. Since σ has the minimum weight among the simplices incident to e and visited by $P[x, x'] \supseteq P[a, x']$, we conclude that

$$\mathrm{cost}(P[a, x']) \geq \omega_\sigma \|ax'\|, \quad \mathrm{cost}(P[x, x']) \geq \omega_\sigma \|xx'\|.$$

Also,

$$\|xx'\| \geq \|px\| - \|px'\| \geq \frac{1}{2}\delta(p) - \frac{\varepsilon}{3}\delta(p) = \frac{3 - 2\varepsilon}{6}\delta(p).$$

Then, $\mathrm{cost}(ap) = \omega_\sigma \|ap\| \leq \omega_\sigma \|ax'\| + \omega_\sigma \|px'\| = \omega_\sigma \|ax'\| + \frac{\varepsilon}{3}\omega_\sigma \delta(p) \leq \mathrm{cost}(P[a, x']) + \frac{2\varepsilon}{3-2\varepsilon}\omega_\sigma \|xx'\| \leq \mathrm{cost}(P[a, x']) + O(\varepsilon) \cdot \mathrm{cost}(P[x, x'])$. Next, $\mathrm{cost}(py') = \omega_e \|py'\| \leq \frac{\varepsilon}{3}\omega_\sigma \delta(p) \leq \frac{2\varepsilon}{3-2\varepsilon}\omega_\sigma \|xx'\| \leq O(\varepsilon) \cdot \mathrm{cost}(P[x, x'])$. Also, $\mathrm{cost}(yy') \leq W\|yy'\| \leq O(W \cdot \frac{\varepsilon^2}{W^2}\|e\|)$. Recall that p is not inside the vertex-balls at the endpoints of e, and these vertex-balls have radius $\Omega(\varepsilon\|e\|/W)$. Therefore, $\delta(p) = \Omega(\varepsilon\|e\|/W)$ by Remark 2. Hence, $\mathrm{cost}(yy') \leq O(\varepsilon) \cdot \delta(p) \leq O(\varepsilon) \cdot \|xx'\| \leq O(\varepsilon) \cdot \mathrm{cost}(P[x, x'])$. ◻

Let P be a path in $\mathcal{T}$ from a source vertex v_s to a destination vertex v_d. The next result shows that P can be converted to a path from v_s to v_d in the Steiner graph $\mathcal{G}$ such that $\mathrm{cost}(Q) \leq (1 + O(\varepsilon)) \cdot \mathrm{cost}(P)$.

Lemma 4.4 *Let P be a path in $\mathcal{T}$ from v_s to v_d. We can convert P to a path Q in $\mathcal{T}$ from v_s to v_d such that there exists nodes $v_s = z_0, z_1, \ldots, z_m = v_d$ of Q with the following properties.*

- *For all $i \in [0, m]$, z_i is a vertex of the Steiner graph $\mathcal{G}$.*
- *For all $i \in [0, m-1]$, $Q[z_i, z_{i+1}]$ is contained in an extended cluster or a vertex star free of skinny tetrahedron.*

Moreover, $\mathrm{cost}(Q) \leq (1 + O(\varepsilon)) \cdot \mathrm{cost}(P)$. Hence, $\mathcal{G}$ gives a $1 + O(\varepsilon)$ approximation because $\mathcal{G}$ contains the edges $\{z_i, z_{i+1}\}$, $i \in [0, m-1]$, with weight $(1 + O(\varepsilon)) \cdot \mathrm{cost}(Q[z_i, z_{i+1}])$.

Proof. We first determine a sequence of points x_i, simplices σ_i, and vertices u_i, $i \in [0, m]$, in this order along P from v_s to v_d. We set $x_0 = \sigma_0 = u_0 = v_s$. For $i = 1, 2, \ldots$, we determine x_i, σ_i, and u_i as follows.

Suppose that u_{i-1} is disjoint from all clusters. If $P[x_{i-1}, v_d] \subseteq \text{star}(u_{i-1})$, then set $m = i$ and $\sigma_m = x_m = v_d$. Otherwise, let x_i be the first exit point of $P[x_{i-1}, v_d]$ from $\text{star}(u_{i-1})$ and let σ_i be the boundary simplex of $\text{star}(u_{i-1})$ of the lowest dimension that contains x_i. The simplex σ_i must be disjoint from u_{i-1}; otherwise, σ_i would be in the boundary of $\mathcal{T}$, meaning that $P[x_{i-1}, v_d]$ could not exit at $x_i \in \sigma_i$.

The remaining possibility is that u_{i-1} is contained in a cluster C_{i-1}. Let C^*_{i-1} be the corresponding extended cluster. If $P[x_{i-1}, v_d] \subseteq C^*_{i-1}$, then set $m = i$ and $\sigma_m = x_m = v_d$. Otherwise, let x_i be the first exit point of $P[x_{i-1}, v_d]$ from C^*_{i-1} and let σ_i be the boundary simplex of C^*_{i-1} of the lowest dimension that contains x_i. The simplex σ_i must be disjoint from C_{i-1}; otherwise, σ_i would be in the boundary of $\mathcal{T}$ and $P[x_{i-1}, v_d]$ could not exit at $x_i \in \sigma_i$.

After determining σ_i and x_i, we determine u_i as follows. If σ_i is a vertex, then $u_i = \sigma_i$. If σ_i is an edge, then u_i is the endpoint of σ_i nearest to x_i. If σ_i is a triangle, find the edge e of σ_i nearest to x_i and then set u_i to be the endpoint of e nearest to x_{i-1}. After determining u_i, if $x_i \neq v_d$, we increment i and repeat the above.

Throughout the determination of x_i, σ_i and u_i for $i = 1, 2, \ldots, m$, we maintain the following two invariants:

- Invariant 1: For every vertex v of σ_i, if F is the fat substar of v that intersects $P[x_{i-1}, x_i]$ and x is the last entry point of $P[x_{i-1}, x_i]$ into $|F|$, then $\|vx\| \geq \delta_F(v)/2$.

  Proof. Let X be $\text{star}(u_{i-1})$ or $C^*_{i-1} \setminus C_{i-1}$, whichever case is appropriate in determining σ_i, x_i and u_i. Let σ be the simplex in F of the lowest dimension containing x. The simplex σ belongs to $F \cap X$, and σ_i is separated from u_{i-1} by fat tetrahedra in $F \cap X$. If v is not a vertex of σ, then σ is a boundary simplex of F in X and so $\|vx\| \geq \delta_F(v)$ by definition. If v is a vertex of σ, then either σ is contained in the interior of $|F \cap X|$ or σ is also a boundary simplex of $\mathcal{T}$. In either case, since $P[x_{i-1}, x_i]$ enters $|F|$ at $x \in \sigma$, we conclude that $x = x_{i-1}$ and $\sigma = \sigma_{i-1}$. As argued previously, σ_i is disjoint from u_{i-1}, so $v \neq u_{i-1}$, which implies that σ_{i-1} is an edge or triangle incident to both v and u_{i-1}. In either case, our method for determining u_{i-1} guarantees that u_{i-1} is closer to x_{i-1} than v. Therefore, $\|vx\| = \|vx_{i-1}\| \geq \|vu_{i-1}\| - \|u_{i-1}x_{i-1}\| \geq \delta_F(v) - \|u_{i-1}x_{i-1}\| \geq \delta_F(v) - \|vx\|$. Hence, $\|vx\| \geq \delta_F(v)/2$.

- Invariant 2: For every edge e of σ_i and every Steiner point $p \in e$, $\|px\| \geq \delta(p)/2$ where x is the last entry point of $P[x_{i-1}, x_i]$ into the fat substar of e.

  Proof. Let F denote the fat substar of e. Let X be $\text{star}(u_{i-1})$ or $C^*_{i-1} \setminus C_{i-1}$, whichever case is appropriate in determining x_i, σ_i and u_i. Let σ be the simplex in F of the lowest dimension containing x. We use $d(A, B)$ to denote the distance between two point sets A and B. If the Steiner point p is disjoint from σ, then $\|px\| \geq d(p, \sigma) \geq \delta(p)$. Suppose that p is incident to σ. Note that e is separated from u_{i-1} by fat tetrahedra in $F \cap X$. Therefore, $\sigma \neq e$, which implies that σ is a triangle incident to e. Either σ is contained in the interior of $|F \cap X|$ or σ is a boundary simplex of $\mathcal{T}$. In either case, since $P[x_{i-1}, x_i]$ enters $|F|$ at $x \in \sigma$, we conclude that $x = x_{i-1}$ and $\sigma = \sigma_{i-1}$. As argued previously, σ_i is disjoint from u_{i-1}, which implies that u_{i-1} is a vertex of σ_{i-1} opposite to the edge e. Refer to Figure 7. Since u_{i-1} is not an endpoint of e, by our method for determining u_{i-1}, the point x_{i-1} must be closer to some edge e' of σ_{i-1} different from e. Then, $\|px_{i-1}\| \geq (\|px_{i-1}\| + d(x_{i-1}, e'))/2 \geq d(p, e')/2 \geq \delta(p)/2$.

After determining the sequence $x_0, x_1, x_2, \ldots, x_m$, we transform P as follows. For every $i \in [1, m-1]$, we apply Lemma 4.1, Lemma 4.2, or Lemma 4.3 to $P[x_{i-1}, x_i]$ and σ_i. First, if

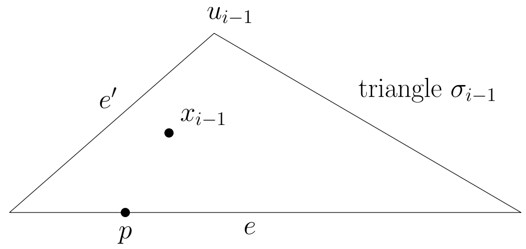

Figure 7: Invariant 2 in the proof of Lemma 4.4.

$x_i \in B_{v,F}$ for some vertex v of σ_i and fat substar F of v, then we apply Lemma 4.1. (Note that invariant 1 is required for invoking Lemma 4.1.) Otherwise, if σ_i is a triangle and $x_i \in B_p$ for some Steiner point p in the interior of t, then we apply Lemma 4.2. If neither Lemma 4.1 nor Lemma 4.2 is applicable, then Lemma 4.3 must be applicable, that is, there is a fat tetrahedron τ incident to σ_i such that:

- σ_i is an edge of τ and $x_i \in B_p$ for some Steiner point $p \in \sigma_i$, or
- σ_i is a triangle of τ, $x_i \in B_p \cap \sigma_i$ for some Steiner point p on an edge of σ_i, and x_i lies outside all triangle-balls with centers in σ_i.

Lemma 4.3 is applicable in either case. (Note that invariant 2 is required for invoking Lemma 4.2.) For $i \in [0, m-1]$, let Y_i denote $\mathrm{star}(u_i)$ or C_i^*, whichever is appropriate in determining x_{i+1}. Irrespective of whether Lemma 4.1, Lemma 4.2, or Lemma 4.3 is invoked on $P[x_i, x_{i+1}]$ and σ_{i+1}, we conclude that for $i \in [0, m-2]$,

(i) $P[x_i, x_{i+1}]$ is converted to a path Q_i from x_i to x_{i+1} that makes a detour to a vertex or Steiner point $z_{i+1} \in \sigma_{i+1}$.

(ii) The path Q_i is contained inside $|Y_i|$.

(iii) $\mathrm{cost}(Q_i) \leq (1 + O(\varepsilon))\,\mathrm{cost}(P[x_i, x_{i+1}])$.

Define $Q_{m-1} = P[x_{m-1}, x_m] = P[x_{m-1}, v_d]$, $z_0 = v_s$, and $z_m = v_d$. Property (iii) above implies that $\sum_{i=0}^{m-1} \mathrm{cost}(Q_i) \leq (1 + O(\varepsilon))\,\mathrm{cost}(P)$. By property (ii) above, for $i \in [0, m-2]$, Q_i is contained in $|Y_i|$. Since $Q_{m-1} = P[x_{m-1}, x_m]$, the subpath Q_{m-1} is contained in $|Y_{m-1}|$. For $i \in [0, m-1]$, observe that Q_i ends at a vertex or Steiner point z_{i+1} in Y_i. It implies that the path $Q = \bigcup_{i=0}^{m-1} Q_i$ can be viewed as $\bigcup_{i=0}^{m-1} Q[z_i, z_{i+1}]$ such that for $i \in [0, m-1]$, the subpath $Q[z_i, z_{i+1}]$ is contained in a vertex star free of skinny tetrahedron or an extended cluster. This completes the proof. ◻

5 Processing extended clusters and vertex stars

Let Γ be a connected set of $O(\kappa)$ tetrahedra. Let p and q be two points in the union of vertices, edges, and triangles in Γ. We present an algorithm to compute a $(1 + O(\varepsilon))$-approximate shortest path in Γ from p to q.

5.1 Locally shortest path

For every triangle $t \in \Gamma$, its *unit disk* is the Euclidean disk D_t that is centered at the origin, lies on a plane parallel to t, and has radius $1/\omega_t$. The travel cost from a point x to a point y in t is λ if changing the radius of $D_t + x$ to λ/ω_t puts y on the boundary of the shrunk or expanded

$$
\begin{array}{lll}
 & \min \quad \sum_{i=0}^{m} z_i + \sum_{i=1}^{m} z'_i & \\
\text{subject to} & x_i = \sum_{j=1}^{3} \alpha_{i,j} v_{i,j} & \forall\, i \in [0, m+1]\ \forall\, j \in [1,3] \\
 & x'_i = \sum_{j=1}^{3} \alpha'_{i,j} v_{i,j} & \forall\, i \in [0, m+1]\ \forall\, j \in [1,3] \\
 & z_i \geq \langle x'_{i+1} - x_i, n_f \rangle / \langle n_f, n_f \rangle & \forall\, i \in [0, m]\ \forall\, \text{facet } f \text{ of } D^*_{\tau_i} \\
 & z'_i \geq \langle x_i - x'_i, n_f \rangle / \langle n_f, n_f \rangle & \forall\, i \in [1, m]\ \forall\, \text{edge } f \text{ of } D^*_{\sigma_i} \\
 & \alpha_{i,j} \geq 0,\ \alpha'_{i,j} \geq 0 & \forall\, i \in [0, m+1]\ \forall\, j \in [1,3] \\
 & \sum_{j=1}^{3} \alpha_{i,j} = 1 = \sum_{j=1}^{3} \alpha'_{i,j} & \forall\, i \in [0, m+1]
\end{array}
$$

Figure 8: Linear programming system.

disk. To approximate D_t, we place $\Theta(1/\sqrt{\varepsilon})$ points roughly uniformly on the boundary of D_t as follows. Enclose D_t by a concentric unit square. Place points on the square boundary at distance $\sqrt{\varepsilon}$ apart. Project these points radially onto the boundary of D_t. Let D^*_t denote the convex hull of the points on the boundary of D_t. One can measure the travel cost from x to y by shrinking or expanding $D^*_t + x$ instead. It is easy to check that D^*_t ensures a $1 + O(\varepsilon)$ approximation of the cost under D_t.

For every tetrahedron $\tau \in \Gamma$, its *unit ball* D_τ is the Euclidean ball centered at the origin with radius $1/\omega_\tau$. The travel cost in τ can be measured by shrinking or expanding D_τ as before. We place $\Theta(1/\varepsilon)$ points roughly uniformly on the boundary of D_τ as follows. Enclose D_τ by a concentric unit cube. Divide the facets of the unit cube into uniform grids so that each grid box has side length $\sqrt{\varepsilon}$. Project the grid vertices radially onto the boundary of D_τ. Let D^*_τ be the convex hull of these points on the boundary of D_τ. It is easy to check that D^*_τ gives a $1 + O(\varepsilon)$ approximation of the cost under D_τ.

Computing D^*_t and D^*_τ for all triangles and tetrahedra takes $O\left(\frac{n}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ time.

Let $\Sigma = (\sigma_1, \sigma_2, \cdots, \sigma_m)$ be a given simplex sequence. Let p and q be two points in some tetrahedra incident to σ_1 and σ_m, respectively. We show how to compute the locally shortest path from p to q with respect to Σ under the modified metric by linear programming. Consider the case that every σ_i is a triangle denoted by $v_{i,1}v_{i,2}v_{i,3}$. The case of some σ_i being vertices or edges can be handled similarly.

Let $x_i x'_{i+1}$ be a possible path link where $x_i \in \sigma_i$ and $x'_{i+1} \in \sigma_{i+1}$. Let τ_i denote the tetrahedron bounded by σ_i and σ_{i+1}. Using barycentric coordinates, the variable $x_i \in \mathbb{R}^3$ satisfies the constraint $x_i = \sum_{j=1}^{3} \alpha_{i,j} v_{i,j}$ for some non-negative variables $\alpha_{i,j} \in \mathbb{R}$ such that $\sum_{j=1}^{3} \alpha_{i,j} = 1$. Similarly, the variable $x'_{i+1} \in \mathbb{R}^3$ satisfies $x'_{i+1} = \sum_{j=1}^{3} \alpha'_{i+1,j} v_{i+1,j}$ for some non-negative variables $\alpha'_{i+1,j} \in \mathbb{R}$ such that $\sum_{j=1}^{3} \alpha'_{i+1,j} = 1$. For convenience, assume that $v_{0,j} = p$ and $v_{m+1,j} = q$ for $j \in [1,3]$. We need the facet g of $D^*_{\tau_i}$ that contains the direction of the vector $x'_{i+1} - x_i$ because the cost of $x_i x'_{i+1}$ is equal to $\langle x'_{i+1} - x_i, n_g \rangle / \langle n_g, n_g \rangle$, where $\langle \cdot, \cdot \rangle$ denotes the inner product operator and n_g denotes the vector that goes from the origin to a point in the support plane of g such that $n_g \perp g$. By the convexity of $D^*_{\tau_i}$, the facet f of $D^*_{\tau_i}$ that gives the largest $\langle x'_{i+1} - x_i, n_f \rangle / \langle n_f, n_f \rangle$ is the correct facet g. Therefore, we introduce a variable $z_i \in \mathbb{R}$ and require $z_i \geq \langle x'_{i+1} - x_i, n_f \rangle / \langle n_f, n_f \rangle$ for every facet f of $D^*_{\tau_i}$. Part of the total path cost is $\sum_{i=0}^{m} z_i$. The minimization ensures that $z_i = \langle x'_{i+1} - x_i, n_g \rangle / \langle n_g, n_g \rangle$ at the end. We also allow for potential critical refraction at σ_{i+1}, i.e., allow for the link $x'_{i+1} x_{i+1} \subset \sigma_{i+1}$. To capture the cost of $x'_{i+1} x_{i+1}$, we introduce another variable z'_{i+1} and require $z'_{i+1} \geq \langle x_{i+1} - x'_{i+1}, n_f \rangle / \langle n_f, n_f \rangle$ for every edge f of $D^*_{\sigma_{i+1}}$. The objective is to minimize $\sum_{i=0}^{m} z_i + \sum_{i=1}^{m} z'_i$. The linear programming system is given in Figure 8.

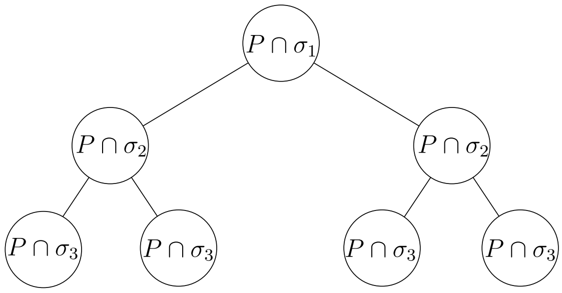

Figure 9: Visualizing a simplex sequence for a shortest path.

There are $\Theta(m\varepsilon^{-1})$ constraints and $\Theta(m)$ variables. The coefficients in the constraints $x_i = \sum_{j=1}^{3} \alpha_{i,j} v_{i,j}$ and $x_i' = \sum_{j=1}^{3} \alpha_{i,j}' v_{i,j}$ have magnitudes N or less because every coordinate of $v_{i,j}$ has magnitude at most N. The vertex coordinates in $D^*_{\tau_i}$ result from multiplying $1/\omega_{\tau_i}$ with the coordinates of the grid vertices on the unit cube. The grid box side length is $\sqrt{\varepsilon}$. Therefore, $O\left(\log \frac{W}{\varepsilon}\right)$ bits suffice for a vertex coordinate in $D^*_{\tau_i}$. For every facet f of $D^*_{\tau_i}$, we first compute an outward normal ν_f of f by taking cross-product using the vertices of f. The coordinates of ν_f thus require $O(\log \frac{W}{\varepsilon})$ bits. Let u be a vertex of f. We solve the linear equation $\langle \frac{1}{\alpha}\nu_f, \frac{1}{\alpha}\nu_f - u \rangle = 0$ for $\alpha \in \mathbb{R}$ such that $\frac{1}{\alpha}\nu_f$ lies on the support plane of f, i.e., $n_f = \frac{1}{\alpha}\nu_f$. Thus, α requires $O(\log \frac{W}{\varepsilon})$ bits and so does n_f. The same conclusion applies to the constraints $\langle x_i - x_i', n_f \rangle / \langle n_f, n_f \rangle$ for every edge f of $D^*_{\sigma_i}$. In summary, the total number of bits to encode the linear program is $O\left(m\varepsilon^{-1} \log \frac{NW}{\varepsilon}\right)$. The ellipsoid method [30] solves the above linear program in $O(m^7\varepsilon^{-3} \log^2 \frac{NW}{\varepsilon} + m^8\varepsilon^{-2} \log^2 \frac{NW}{\varepsilon})$ arithmetic operations.

5.2 Approximate shortest path

To compute the approximate shortest path in Γ from p to q, our strategy is to enumerate all possible simplex sequences from p to q, use the method in Section 5.1 to compute a $1 + O(\varepsilon)$ approximation of the locally shortest path with respect to each simplex sequence, and finally select the shortest one among these locally shortest paths. The remaining questions are how long a simplex sequence and how many simplex sequences we need to consider.

Consider a shortest path P in Γ from p to q. Let $\sigma_1, \sigma_2, \cdots$ be the simplices in Γ in non-decreasing order of weights. We can assume that $P \cap \sigma_1$ is connected. Otherwise, we can shortcut P by joining the two connected components in $P \cap \sigma_1$ by a line segment in σ_1 without increasing the path cost. For a similar reason, we can assume that $P \cap \sigma_2$ has at most two connected components. In general, $P \cap \sigma_i$ has at most 2^{i-1} connected components. This argument is best visualized as arranging the connected components in a full binary tree with $P \cap \sigma_1$ at the root, two nodes of $P \cap \sigma_2$ at the next level, and so on as illustrated in Figure 9. It follows that the simplex sequence is at most $2^{O(\kappa)}$ long. Consequently, there are at most $2^{2^{O(\kappa)}}$ simplex sequences. There are $O(\frac{\kappa^2}{\varepsilon^4} \log^2 \frac{W}{\varepsilon})$ pairs of vertices and Steiner points in an extended cluster or vertex star free of skinny tetrahedra. We repeat the approximate shortest path computation $O(n \cdot \frac{\kappa^2}{\varepsilon^4} \log^2 \frac{W}{\varepsilon})$ times, invoking the result in Section 5.1 at most $2^{2^{O(\kappa)}}$ times with $m = 2^{O(\kappa)}$ for each approximate shortest path computation.

The above processing completes the construction of the Steiner graph $\mathcal{G}$ which has $O(\frac{n}{\varepsilon^2} \log \frac{W}{\varepsilon})$ vertices and $O(\frac{n}{\varepsilon^4} \log^2 \frac{W}{\varepsilon})$ edges by Lemmas 3.1 and 3.4. We can run Dijkstra's algorithm to find the shortest path from v_s to v_d in $\mathcal{G}$ in $O(\frac{n}{\varepsilon^2} \log \frac{W}{\varepsilon}(\log \frac{n}{\varepsilon} + \log\log \frac{W}{\varepsilon}) + \frac{n}{\varepsilon^4} \log^2 \frac{W}{\varepsilon})$ time [16]. But this is dominated by the construction cost of $\mathcal{G}$.

Theorem 5.1 *Let ρ be an arbitrary constant. Let $\mathcal{T}$ be a simplicial complex of n tetrahedra such that vertices have integral coordinates with magnitude at most N and tetrahedra have* *integral weights in the range* $[1, W]$*. Let* κ *be the number of tetrahedra in the largest connected component of tetrahedra whose aspect ratios exceed* ρ*. For all* $\varepsilon \in (0,1)$ *and for every pair of source and destination vertices* v_s *and* v_d *in* $\mathcal{T}$*, we can find a* $(1+\varepsilon)$*-approximate shortest path in* $\mathcal{T}$ *from* v_s *to* v_d *in* $O(2^{2^{O(\kappa)}} n\varepsilon^{-7} \log^2 \frac{W}{\varepsilon} \log^2 \frac{NW}{\varepsilon})$ *time.*

6 Conclusion

We presented a $(1+\varepsilon)$-approximation algorithm for the shortest path problem for weighted regions in three dimensions. The novelty of this algorithm is that the time complexity depends on combinatorial parameters only and it is sensitive to the size of the largest connected component of skinny tetrahedra. There exists a constant $C \geq 1$ such that if this size is at most $\frac{1}{C} \log\log n + O(1)$, then our running time is polynomial in n, $1/\varepsilon$, and $\log(NW)$, where n is the total number of tetrahedra, N is the largest vertex coordinate magnitude, and W is the maximum region weight. When the size of the largest connected component of skinny tetrahedra is $O(1)$, our running time is linear in n and polynomial in $1/\varepsilon$ and $\log(NW)$. Our result holds irrespective of the worst aspect ratio among the tetrahedra. It remains an open problem to find a FPTAS for the shortest path problem for weighted regions in three dimensions.

References

[1] M. Ahmed. *Constrained Shortest Paths in Terrains and Graphs.* PhD Thesis, University of Waterloo, Canada, 2009.

[2] M. Ahmed, S. Das, S. Lodha, A. Lubiw, A. Maheshwari, and S. Roy. Approximation algorithms for shortest descending paths in terrains. *Journal of Discrete Algorithms*, 8 (2010), 214–230.

[3] M. Ahmed and A. Lubiw. Shortest descending paths through given faces. *Computational Geometry: Theory and Applications*, 42 (2009), 464–470.

[4] M. Ahmed and A. Lubiw. Shortest descending paths: towards an exact algorithm. *International Journal of Computational Geometry and Applications*, 21 (2011), 431–466.

[5] L. Aleksandrov, A. Maheshwari, and J.-R. Sack. Determining approximate shortest paths on weighted polyhedral surfaces. *Journal of ACM*, 52 (2005), 25–53.

[6] L. Aleksandrov, H. Djidjev, A. Maheshwari, and J.-R. Sack. An approximation algorithm for computing shortest paths in weighted 3-d domains. *Discrete and Computational Geometry*, 50 (2013), 124–184.

[7] D.Z. Chen and H. Wang. Computing shortest paths among curved obstacles in the plane. *Proceedings of the 29th Annual Symposium on Computational Geometry*, 2013, 369–378.

[8] J. Chen and Y. Han. Shortest paths on a polyhedron. *International Journal of Computational Geometry and Applications*, 6 (1996), 127–144.

[9] S.-W. Cheng and J. Jin. Approximate shortest descending paths. *SIAM Journal on Computing*, 43 (2014), 410–428.

[10] S.-W. Cheng and J. Jin. Shortest paths on polyhedral surfaces and terrains. *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, 2014, 373–382.

[11] S.-W. Cheng, J. Jin, and A. Vigneron. Triangulation refinement and approximate shortest paths in weighted regions. *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2015, 1626–1640.

[12] S.-W. Cheng, H.-S. Na, A. Vigneron, and Y. Wang. Approximate shortest paths in anisotropic regions. *SIAM Journal on Computing*, 38 (2008), 802–824.

[13] S.-W. Cheng, H.-S. Na, A. Vigneron, and Y. Wang. Querying approximate shortest paths in anisotropic regions. *SIAM Journal on Computing*, 39 (2010), 1888–1918.

[14] J. Choi, J. Sellen, and C.-K. Yap. Approximate Euclidean shortest path in 3-space. *Proceedings of the 10th Annual Symposium on Computational Geometry*, 1994, 41–48.

[15] K.L. Clarkson. Approximation algorithms for shortest path motion planning. *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, 1987, 56–65.

[16] M. Fredman and R. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of ACM*, 34 (1987), 596–615.

[17] B. Giroux and B. Larouche. Task-parallel implementation of 3D shortest path raytracing for geophysical applications. *Computers & Geosciences*, 54 (2013), 130–141.

[18] S. Har-Peled. Constructing approximate shortest path maps in three dimensions. *SIAM Journal on Computing*, 28 (1999), 1182–1197.

[19] J. Hershberger and S. Subhash. An optimal algorithm for Euclidean shortest paths in the plane. *SIAM Journal on Computing*, 28 (1999), 2215–2256.

[20] J. Hershberger, S. Subhash, and H. Yildiz. A near-optimal algorithm for shortest paths among curved obstacles in the plane. *Proceedings of the 29th Annual Symposium on Computational Geometry*, 2013, 359–368.

[21] S. Kapoor, S.N. Maheshwari, and J.S.B. Mitchell. An efficient algorithm for Euclidean shortest paths among polygonal obstacles in the plane. *Discrete and Computational Geometry*, 18 (1997), 377–383.

[22] M. Lanthier, A. Maheshwari, and J.-R. Sack. Approximating weighted shortest paths on polyhedral surfaces. *Algorithmica*, 30 (2001), 527–562.

[23] C. Mata and J.S.B. Mitchell. A new algorithm for computing shortest paths in weighted planar subdivisions. *Proceedings of the 13th Annual Symposium on Computational Geometry*, 1997, 264–273.

[24] W. Menke. *Geophysical Data Analysis: Discrete Inverse Theory*, Academic Press, 2012.

[25] J.S.B. Mitchell. Geometric shortest paths and network optimizations. In *Handbook of Computational Geometry*, Elsevier Science, (J.-R. Sack and J. Urrutia, eds), 2000.

[26] J.S.B. Mitchell, D. Mount, and C. Papadimitriou. The discrete geodesic problem. *SIAM Journal on Computing*, 1987, 647–668.

[27] J.S.B. Mitchell and C.H. Papadimitrou. The weighted region problem: finding shortest paths through a weighted planar subdivision. *Journal of ACM*, 8 (1991), 18–73.

[28] J.S.B. Mitchell and M. Sharir. New results on shortest paths in three dimensions. *Proceedings of the 20th Annual Symposium on Computational Geometry*, 2004, 124–133.

[29] C.H. Papadimitriou. An algorithm for shortest-path motion in three dimensions. *Information Processing Letters*, 20(1985), 259–263.

[30] C.H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity.* Dover, 1998.

[31] Y. Schreiber and M. Sharir. An optimal-time algorithm for shortest paths on a convex polytope in three dimensions. *Discrete and Computational Geometry*, 39 (2008), 500–579.

[32] Z. Sun and J. Reif. On finding energy-minimizing paths on terrains. *IEEE Transactions on Robotics*, 21 (2005), 102–114.

[33] Z. Sun and J. Reif. On finding approximate optimal paths in weighted regions. *Journal of Algorithms*, 58 (2006), 1–32.