

Navigating Weighted Regions with Scattered Skinny Tetrahedra

Siu-Wing Cheng^{1(✉)}, Man-Kwun Chiu^{2,3}, Jiongxin Jin⁴,
and Antoine Vigneron⁵

¹ Department of Computer Science and Engineering, HKUST,
Hong Kong, Hong Kong
`scheng@cse.ust.hk`

² National Institute of Informatics (NII), Tokyo, Japan

³ JST, ERATO, Kawarabayashi Large Graph Project, Tokyo, Japan

⁴ Google Inc., Seattle, USA

⁵ Visual Computing Center,

King Abdullah University of Science and Technology (KAUST),
Thuwal, Saudi Arabia

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 $O(n\varepsilon^{-O(1)}(\log(NW))^{O(1)})$.

Keywords: Weighted region · Shortest path · Approximation algorithm

1 Introduction

Finding shortest paths are classical geometric optimization problems (e.g. [4, 10–12, 15]). In 2D, researchers have also studied cost models in applications that are non- L_p metrics and anisotropic (e.g. [1, 2, 5–9, 14, 17]). 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 (e.g. [13]).

The weighted region problem is a way to model the unequal difficulties in traversing different regions [14]. 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

S.-W. Cheng—Supported by Research Grants Council, Hong Kong, China (project no. 611812).

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 $\text{cost}(P) = \sum_{\text{simplex } \sigma} \omega_\sigma \|P \cap \sigma\|$. 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 [7, 14]. 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 [2, 17]. In contrast, only one algorithm for the weighted region problem in 3D has been proposed (Aleksandrov et al. [3]). The authors [3] present a $(1 + \varepsilon)$ -approximation algorithm whose running time is $O(Kn\varepsilon^{-2.5} \log \frac{n}{\varepsilon} \log^3 \frac{1}{\varepsilon})$, 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 Sect. 2.) It is an open problem whether an FPTAS exists for the 3D weighted region problem.

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 , $1/\varepsilon$ and $\log(NW)$. If $\kappa = O(1)$, the running time is linear in n . In comparison, the running time in [3] 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 [3] 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. Recall that every cluster has at most κ tetrahedra.

For every simplex σ in \mathcal{T} , $\text{star}(\sigma)$ 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 $\text{bd}(\mathcal{K})$ 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 [3]. However, 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.

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 $\text{star}(v)$ 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 $\text{star}(\sigma)$.

Definition 1. *Let x be a point in the union of vertices, edges and triangles of \mathcal{T} . Let σ be the simplex of 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 $\text{bd}(F)$ 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 $\delta(x)$.*

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 that bounds a fat tetrahedron incident to e and shares only a vertex v with e . Also, $\delta(x) = \Theta(\|vx\|)$.

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 .

Lemma 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 , $\|pq\| \geq \frac{\varepsilon}{3} \cdot \max\{\delta(p), \delta(q)\}$, and either p lies on the boundary of B_q or q lies on the boundary of B_p . There are $O\left(\frac{1}{\varepsilon} \log \frac{W}{\varepsilon}\right)$ Steiner points in uv .*

Proof. The construction ensures the coverage of uv . Assume that q was placed after p . By construction, q is not inside $B(p, \frac{\varepsilon}{3}\delta(p))$ and so $\|pq\| \geq \frac{\varepsilon}{3}\delta(p)$. As q is placed after p , $\delta(q) \leq \delta(p)$ and so $\|pq\| \geq \frac{\varepsilon}{3}\delta(p) \geq \frac{\varepsilon}{3}\delta(q)$.

In the interior of uv , $\delta(x)$ 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 the third property in the lemma.

Let F be the fat substar of u that contains uv . At the intersection point x between uv and the boundary of N_u , $\delta(x) \leq \|ux\| = \frac{\varepsilon}{3W}\delta_F(u) = \Theta\left(\frac{\varepsilon}{W}\|uv\|\right)$ by Remark 1. By Remark 2, $\delta(x) = \Omega(\min\{\|ux\|, \|vx\|\}) = \Omega(\min\{\frac{\varepsilon}{W}\|uv\|, (1 - \frac{\varepsilon}{W})\|uv\|\})$. So $\delta(x) = \Theta\left(\frac{\varepsilon}{W}\|uv\|\right)$. Similarly, at the intersection point x between uv and the boundary of N_v , $\delta(x) = \Theta\left(\frac{\varepsilon}{W}\|uv\|\right)$. The maximum value of $\delta(x)$ in the interior of uv is at most $\frac{1}{2}\|uv\|$. 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 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 $1 + \Theta(\varepsilon)$ from $\Omega(\frac{\varepsilon^2}{W}\|uv\|)$ at the boundary of N_u to $O(\varepsilon\|uv\|)$ in the interior of uv . The same holds for the sequence of Steiner points from N_v . Hence, there are $O(\log_{1+\Theta(\varepsilon)} \frac{W}{\varepsilon}) = O(\frac{1}{\varepsilon} \log \frac{W}{\varepsilon})$ Steiner points. \square

Lemma 2. *Placing Steiner points on an edge takes $O(\frac{1}{\varepsilon} \log \frac{W}{\varepsilon})$ time.*

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 , $\delta(x) \geq c\varepsilon^2\ell/W$ for some constant $c > 0$.¹ 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. *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 $O(\frac{1}{\varepsilon^2} \log \frac{W}{\varepsilon})$ Steiner points in uvw .*

Proof. The construction ensures the coverage of uvw . We can show as in the proof of Lemma 1 that $\|pq\| \geq \frac{\varepsilon}{3} \max\{\delta(p), \delta(q)\}$ for every pair of Steiner points p and q placed in A_i . 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 $i = O(\log_{1+\varepsilon} \frac{W}{\varepsilon}) = O(\frac{1}{\varepsilon} \log \frac{W}{\varepsilon})$. It follows that there are $O(\frac{1}{\varepsilon^2} \log \frac{W}{\varepsilon})$ Steiner points in uvw . \square

¹ The smallest value of $\delta(x)$ occurs near the edge-ball centered at the intersection point between uv and the boundary of N_u or the boundary of N_v .

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

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 Sect. 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 Sect. 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 5, 6, and 7) that snap a path to vertices and Steiner points.

Lemma 5. *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 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. Let x' be the first entry point of $P[x,y]$ into $B_{v,F}$. We replace $P[x,y]$ by $P[x,x'] \cup x'v \cup vy$. We have $\text{cost}(x'v) \leq W\|x'v\| = \frac{\varepsilon}{3}\delta_F(v) \leq \frac{2\varepsilon}{3-2\varepsilon}\|xx'\| \leq O(\varepsilon) \cdot \text{cost}(P[x,x']) \leq O(\varepsilon) \cdot \text{cost}(P[x,y])$. Similarly, $\text{cost}(vy) \leq O(\varepsilon) \cdot \text{cost}(P[x,y])$. \square

Lemma 6. *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 . 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 $\delta(p) \leq \|p\hat{x}\|$. We replace $P[x,y]$ by $P[x,x'] \cup x'p \cup py$. Figure 1 illustrates the three cases below.




Fig. 1. The path Q in cases 1, 2 and 3 in the proof of Lemma 6 from left to right.

Case 1: $P[\hat{x}, x']$ is a segment whose interior lies in the interior of τ . We have $\text{cost}(x'p) = \frac{\varepsilon}{3}\omega_\tau\delta(p) \leq \frac{\varepsilon}{3-\varepsilon}\omega_\tau\|\hat{x}x'\| \leq O(\varepsilon) \cdot \text{cost}(P[x, x']) \leq O(\varepsilon) \cdot \text{cost}(P[x, y])$. 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 . We analyze the extra cost as in Case 1 with ω_τ replaced by ω_t .

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 $\delta(p) \leq \frac{6}{3-\varepsilon}\|\hat{x}x''\|$. Otherwise, $\|x''x'\| \geq \frac{1}{2}\|\hat{x}x'\|$ and we adapt the analysis in Case 2 using the relation $\delta(p) \leq \frac{6}{3-\varepsilon}\|x''x'\|$. \square

Lemma 7. *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 $\text{cost}(Q) \leq (1 + O(\varepsilon)) \cdot \text{cost}(P[x, y])$.*

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 . Let x' be the first entry




Fig. 2. The resulting path Q in the proof of Lemma 7.

point of $P[x, y]$ into B_p . Retrace $P[x, x']$ from x' towards x until we hit x or a simplex in $\text{bd}(F)$ other than e for the first time. Let \hat{x} be the point where retracing stops. Note that $\delta(p) \leq 2\|p\hat{x}\|$.

Let σ be the triangle or tetrahedron with the minimum weight among those incident to e and visited by $P[\hat{x}, x']$. Suppose that $P[\hat{x}, x']$ 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 2 illustrates the conversion. First, $\text{cost}(ap) = \omega_\sigma \|ap\| \leq \omega_\sigma \|ax'\| + \frac{\varepsilon}{3} \omega_\sigma \delta(p) \leq \text{cost}(P[a, x']) + \frac{2\varepsilon}{3-2\varepsilon} \omega_\sigma \|\hat{x}x'\| \leq \text{cost}(P[a, x']) + O(\varepsilon) \cdot \text{cost}(P[x, x'])$. Next, $\text{cost}(py') = \omega_e \|py'\| \leq \frac{\varepsilon}{3} \omega_\sigma \delta(p) \leq \frac{2\varepsilon}{3-2\varepsilon} \omega_\sigma \|\hat{x}x'\| \leq O(\varepsilon) \cdot \text{cost}(P[x, x'])$. Also, $\text{cost}(yy') \leq W \|yy'\| \leq O(W \cdot \frac{\varepsilon^2}{W^2} \|e\|)$. 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, $\text{cost}(yy') \leq O(\varepsilon) \cdot \delta(p) \leq O(\varepsilon) \cdot \|\hat{x}x'\| \leq O(\varepsilon) \cdot \text{cost}(P[x, x'])$. \square

Next, we convert a path P from v_s to v_d to a path Q such that the nodes $v_s = u_{i_1}, u_{i_2}, u_{i_3}, \dots, u_{i_m} = v_d$ in Q are vertices of \mathcal{G} , and for all $j \geq 1$, $Q[u_{i_j}, u_{i_{j+1}}]$ is contained in an extended cluster or a vertex star free of skinny tetrahedra. Moreover, $\text{cost}(Q) \leq (1 + O(\varepsilon)) \cdot \text{cost}(P)$. Therefore, \mathcal{G} gives a $1 + O(\varepsilon)$ approximation because \mathcal{G} contains the edges $\{u_{i_j}, u_{i_{j+1}}\}$ with weight $(1 + O(\varepsilon)) \cdot \text{cost}(Q[u_{i_j}, u_{i_{j+1}}])$.

Lemma 8. *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 the nodes $v_s = u_{i_1}, u_{i_2}, u_{i_3}, \dots, u_{i_m} = v_d$ in Q are vertices of \mathcal{G} , and for all $j \geq 1$, $Q[u_{i_j}, u_{i_{j+1}}]$ is contained in an extended cluster or a vertex star free of skinny tetrahedra. Moreover, $\text{cost}(Q) \leq (1 + O(\varepsilon)) \cdot \text{cost}(P)$.*

Proof. Let P_0 denote a path from v_s to v_d in \mathcal{T} .

Suppose that v_s is disjoint from all clusters. If P_0 does not leave $\text{star}(v_s)$, then v_d is a vertex in $\text{star}(v_s)$ and the lemma is trivially true. Assume that P_0 leaves $\text{star}(v_s)$ for the first time at a point y . Then y lies in a boundary simplex σ of $\text{star}(v_s)$ disjoint from v_s . We modify $P_0[v_s, y]$ by applying Lemmas 5, 6, or 7 to make a detour to a vertex or Steiner point $p \in \sigma$.

Suppose that v_s is contained in a cluster C . Recall that C^* denotes the extended cluster corresponding to C . If P_0 does not leave C^* , then v_d is a vertex in C^* and there is nothing to prove. Assume that P_0 leaves C^* for the first time at a point y . Let x be the point in the boundary of C that P_0 leaves C for the last time before reaching y . Let σ be the simplex of lowest dimension in $\text{bd}(C^*)$ that contains y . The simplex σ is disjoint from the boundary of C ; otherwise, σ would be a boundary simplex of \mathcal{T} , meaning that P_0 cannot leave C^* at y , a contradiction. We modify $P_0[x, y]$ by applying Lemmas 5, 6, or 7 to make a detour to a vertex or Steiner point $p \in \sigma$.

Let P_1 denote the path resulted from modifying P_0 . The extra cost of $O(\varepsilon) \cdot \text{cost}(P_0[v_s, y])$ can be charged to $P_0[v_s, y]$. Then we work on $P_1[y, v_d]$. Recall that y belongs to the boundary simplex σ of $\text{star}(v_s)$ or an extended cluster containing v_s , whichever case is applicable. We identify a vertex v as follows. If σ is a vertex, let $v = \sigma$. If σ is an edge, let v be the closest endpoint of σ to y . If σ is a triangle, let e be the closest edge of σ to y and then let v be

the closest endpoint of e to y . Then, we repeat the previous analysis on v and $P_1[y, v_d]$. That is, we check the exit of $P_1[y, v_d]$ from $\text{star}(v)$ or an extended cluster containing v , whichever case is applicable. The vertex or Steiner point p to which P_0 was snapped belongs to σ and p is already a vertex of \mathcal{G} . The next application of Lemmas 5, 6, or 7 will convert $P_1[y, v_d]$ to a path $P_2[y, v_d]$ that passes through a vertex or Steiner point q such that $P_2[p, q]$ lies in an extended cluster or a vertex star free of skinny tetrahedra. The extra charge in converting P_1 to P_2 can be charged to a subpath of $P_1[y, v_d]$. Repeating the argument proves the lemma. \square

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 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 $1/\omega_\tau$. Similar to the 2D case, D_τ can be approximated by a convex hull D_τ^* with $O(1/\varepsilon)$ vertices. Computing D_t^* and D_τ^* for all triangles and tetrahedra takes $O(\frac{n}{\varepsilon} \log \frac{1}{\varepsilon})$ time.

Let $\Sigma = (\sigma_1, \sigma_2, \dots, \sigma_m)$ 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 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 $\sum_{i=0}^m z_i + \sum_{i=1}^m z'_i$.

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 . Roughly speaking, 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, $O(\log \frac{W}{\varepsilon})$ 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 $O(m\varepsilon^{-1} \log \frac{NW}{\varepsilon})$. The ellipsoid method [16] 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 Sect. 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 $\sigma_1, \sigma_2, \dots$ 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. 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 Sect. 5.1 at most $2^{2^{O(\kappa)}}$ times with $m = 2^{O(\kappa)}$ for each approximate shortest path computation.

Theorem 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.*

References

1. Ahmed, M.: Constrained Shortest Paths in Terrains and Graphs. Ph.D. Thesis, University of Waterloo, Canada (2009)
2. Aleksandrov, L., Maheshwari, A., Sack, J.-R.: Determining approximate shortest paths on weighted polyhedral surfaces. *J. ACM* **52**, 25–53 (2005)
3. Aleksandrov, L., Djidjev, H., Maheshwari, A., Sack, J.-R.: An approximation algorithm for computing shortest paths in weighted 3-d domains. *Discrete. Comput. Geom.* **50**, 124–184 (2013)
4. Chen, J., Han, Y.: Shortest paths on a polyhedron. *Int. J. Comput. Geom. Appl.* **6**, 127–144 (1996)
5. Cheng, S.-W., Jin, J.: Approximate shortest descending paths. *SIAM J. Comput.* **43**, 410–428 (2014)
6. Cheng, S.-W., Jin, J.: Shortest paths on polyhedral surfaces and terrains. In: *Proceedings of ACM Symposium on Theory of Computing*, pp. 373–382 (2014)
7. Cheng, S.-W., Jin, J., Vigneron, A.: Triangulation refinement and approximate shortest paths in weighted regions. In: *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pp. 1626–1640 (2015)
8. Cheng, S.-W., Na, H.-S., Vigneron, A., Wang, Y.: Approximate shortest paths in anisotropic regions. *SIAM J. Comput.* **38**, 802–824 (2008)
9. Cheng, S.-W., Na, H.-S., Vigneron, A., Wang, Y.: Querying approximate shortest paths in anisotropic regions. *SIAM J. Comput.* **39**, 1888–1918 (2010)
10. Choi, J., Sellen, J., Yap, C.-K.: Approximate Euclidean shortest path in 3-space. In: *Proceedings of the Annual Symposium on Computational Geometry*, pp. 41–48 (1994)
11. Clarkson, K.L.: Approximation algorithms for shortest path motion planning. In: *Proceedings of the ACM Symposium on Theory Computing*, pp. 56–65 (1987)
12. Hershberger, J., Subhash, S.: An optimal algorithm for Euclidean shortest paths in the plane. *SIAM J. Comput.* **28**, 2215–2256 (1999)
13. Menke, W.: *Geophysical Data Analysis: Discrete Inverse Theory*. Academic Press, New York (2012)
14. Mitchell, J.S.B., Papadimitriou, C.H.: The weighted region problem: finding shortest paths through a weighted planar subdivision. *J. ACM* **8**, 18–73 (1991)
15. Papadimitriou, C.H.: An algorithm for shortest-path motion in three dimensions. *Inf. Process. Lett.* **20**, 259–263 (1985)
16. Papadimitriou, C.H., Steiglitz, K.: *Combinatorial Optimization: Algorithms and Complexity*. Dover, New York (1998)
17. Sun, Z., Reif, J.: On finding approximate optimal paths in weighted regions. *J. Alg.* **58**, 1–32 (2006)