Edge Flips in Surface Meshes*

Siu-Wing Cheng[†] Jiongxin Jin[‡]

Abstract

Little theoretical work has been done on edge flips in surface meshes despite their popular usage in graphics and solid modeling to improve mesh equality. We propose the class of (ε, α)-meshes of a surface that satisfy several properties: the vertex set is an ε-sample of the surface, the triangle angles are no smaller than a constant α, some triangle has a good normal, and the mesh is homeomorphic to the surface. We believe that many surface meshes encountered in practice are (ε, α)-meshes or close to being one. We prove that flipping the appropriate edges can smooth a dense (ε, α)-mesh by making the triangle normals better approximations of the surface normals and the dihedral angles closer to π. Moreover, the edge flips can be performed in time linear in the number of vertices. This helps to explain the effectiveness of edge flips as observed in practice and in our experiments. A corollary of our techniques is that, in $\mathbb{R}^2$, every triangulation with a constant lower bound on the angles can be flipped in linear time to the Delaunay triangulation.

1 Introduction

Surface meshes are popular representations of smooth surfaces in computer graphics and solid modeling. The quality and smoothness of surface meshes are often improved by applying edge flips. For example, each candidate edge flip is assigned a score in [11] that measures how the flip can decrease a cost function that reflects the overall discrete curvature of the mesh, and the edge flips are applied in a greedy manner based on the scores; edge flips are used in [15] to improve the aspect ratios of the triangles for flow simulations; edge flips are performed in [1, 21] to reduce the variance in the vertex degrees as well as to improve the aspect ratios of the triangles. Despite the popularity of edge flips, there has been no theoretical study of their impact on the surface mesh quality.

We propose the class of (ε, α)-*meshes* of a closed surface that satisfy several properties: the vertex set is an ε-sample of the surface, the triangle angles are no smaller than a constant α, some triangle has a good normal, and the mesh is homeomorphic to the surface. We believe that many surface meshes encountered in practice are (ε, α)-meshes or close to being one. We prove that flipping the appropriate edges can smooth a dense (ε, α)-mesh by making the triangle normals better approximations of the surface normals and the dihedral angles closer to π. Moreover, the edge flips can be performed in time linear in the number of vertices. This helps to explain the effectiveness of edge flips as observed in practice and in our experiments. We also show that edge flips can be applied locally in a dense (ε, α)-mesh: given a subset V of the vertices, edge flips can be performed in $O(|V|)$ time to improve the mesh smoothness at the vertices in V. (See Theorem 2 in Section 6.) In $\mathbb{R}^3$, our definition of edge flippability is different from the usual empty circumsphere criterion, and it can be checked by a primitive that compares the

*Research supported by Research Grant Council, Hong Kong, China (612109). The work was done while Jin was at the Department of Computer Science and Engineering, HKUST.

[†]Department of Computer Science and Engineering, HKUST, Hong Kong. Email: scheng@cse.ust.hk

[‡]Google Inc., Seattle, USA. Email: jamesjjx@google.com

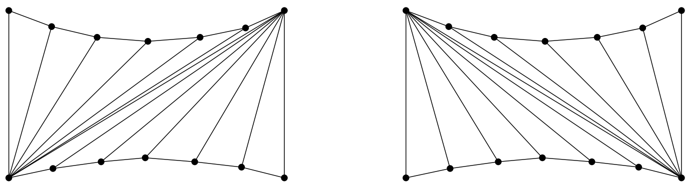

Figure 1: An example from [17] in which $\Omega(n^2)$ edge flips are needed to convert the left triangulation to the right triangulation and vice versa.

inner products of vectors of the form $u - v$, where u and v are appropriate mesh vertices. (See Definition 3 in Section 4.)

Our results are obtained by showing that, upon the termination of edge flips, the circumradius of every triangle τ is at most $\varepsilon + O(\varepsilon^\kappa)$ times the local feature size at any vertex of τ, where κ is any fixed constant in $(1, 1.5)$. (The circumradii of the triangles may be much bigger before the edge flips.) By standard surface sampling results, smaller circumradii make the mesh smoother and a better approximation of the underlying surface. The ability to decrease the triangle circumradii is also useful in maintaining deforming surface meshes [6, 7]. The circumradius bound of $\varepsilon + O(\varepsilon^\kappa)$ times the local feature size is proved by showing that edge flips make the diametric ball of every triangle almost empty of vertices. The proofs require the mesh vertices to form a very dense sample of the underlying surface, however, our experiments show that edge flips work well even if the vertices are not very dense. (See Section 7.)

A corollary of our techniques is that, in $\mathbb{R}^2$, given a triangulation with n vertices and all angles greater than some constant independent of n, it can be converted in $O(n)$ time to the Delaunay triangulation by edge flips (Theorem 1 in Section 4). In the general case where the angles are not bounded from below by some constant, Hurtado et al. [17] proved that $\Omega(n^2)$ edge flips are needed to convert any one of the two triangulations in Figure 1 to the other, which implies that one of the triangulations in Figure 1 needs $\Omega(n^2)$ edge flips to become Delaunay. As shown in Figure 1, some angles can be as small as $O(1/n)$. Since the angles are not greater than some constant, our result (Theorem 1) cannot be applied in this case to deduce a linear bound on the number of edge flips.

We have employed our edge flip procedure in a surface reconstruction protytype. The prototype first extracts a locally uniform subsample from the input sample, invokes Cocone [3] to reconstruct a triangular mesh, and then inserts the remaining sample points to the mesh via edge flips. Good reconstructions and speedups were obtained on several models in the Stanford 3D scanning repository when compared with running Cocone alone. Refer to [8] for more details. Section 7 shows detailed experiments with the edge flip procedure alone on some synethetic models. The experiments show that the edge flips are effective in improving the mesh quality and they can be performed efficiently.

2 Preliminaries

For every pair of points $x, y \in \mathbb{R}^3$, $d(x, y)$ denotes the Euclidean distance between x and y. Given a point set $Y \subseteq \mathbb{R}^3$, $d(x, Y)$ denotes $\inf_{y \in Y} d(x, y)$. Given two vectors $\vec{u}$ and $\vec{v}$, $\angle(\vec{u}, \vec{v})$ denotes the angle between them which lies in the range $[0, \pi]$. Given three points a, b and c, we use $\angle abc$ to denote $\angle(a - b, c - b)$. Given a subset $X \subset \mathbb{R}^3$, $\mathrm{aff}(X)$ denotes the affine subspace of the lowest dimension that contains X. Let h and h' be two linear objects such as vectors, segments, lines, polygons, and planes. We use $\angle_a(h, h')$ to denote the nonobtuse angle between $\mathrm{aff}(h)$ and $\mathrm{aff}(h')$. Let $B(x, r)$ denote the ball in $\mathbb{R}^3$ with center x and radius r. Given a ball B,

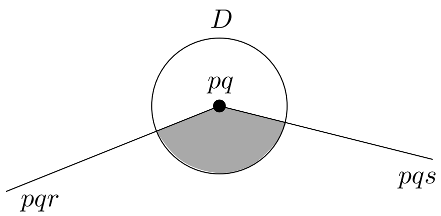

Figure 2: The figure shows a cross-section of pqr and pqs. D is partitioned by $pqr \cup pqs$ into two sectors. The angle of the smaller shaded sector is the dihedral angle at pq.

we use ∂B to denote its boundary. Given a triangle τ, c_τ denotes its circumcenter, γ_τ denotes its circumradius, B_τ denotes the *diametric ball* $B(c_\tau, \gamma_\tau)$ of τ, and $\mathbf{n}_\tau$ denotes a unit vector orthogonal to $\mathrm{aff}(\tau)$. The diametric ball of τ is the smallest circumscribing ball of τ.

A *triangulated polygonal surface* T is a set of vertices, edges and triangles such that the intersection of every pair of elements in T is either empty or an element in T, and for every vertex of T, its incident triangles form a topological disk (i.e., the union of the incident triangles is homeomorphic to a two-dimensional disk). The union of the vertices, edge and triangles form the *underlying space* $|T|$ of T. The *star* of a vertex $p \in T$, denoted $\mathrm{star}(p)$, is the set of edges and triangles in T that are incident to p. When there is ambiguity, we may write $\mathrm{star}(p, T)$ to specify the underlying triangulated polygonal surface. Take two triangles $pqr, pqs \in T$ that share the edge pq. Place an arbitrarily small two-dimensional geometric disk D (embedded in $\mathbb{R}^3$) such that the center of D lies in the interior of pq and D is orthogonal to pq. D is partitioned by pqr and pqs into two sectors. The *dihedral angle* at pq is the angle of the smaller sector. Figure 2 shows an example. Equivalently, the dihedral angle at pq is equal to $\pi - \angle_a(\mathbf{n}_{pqr}, \mathbf{n}_{pqs})$.

Let $\Sigma \subset \mathbb{R}^3$ be a closed connected smooth surface throughout this paper. For every point $x \in \Sigma$, a *medial ball* B at x is a maximal ball tangent to Σ at x such that the interior of B does not intersect Σ. The *medial axis* $\mathcal{M}$ of Σ is the set of centers of medial balls at points in Σ. The *local feature size* of a point $x \in \Sigma$ is $f(x) = d(x, \mathcal{M})$. The local feature size function f is 1-Lipschitz, i.e., $f(x) \leq f(y) + d(x, y)$ [9]. A finite point set $P \subset \Sigma$ is an *ε-sample* of Σ for some $\varepsilon \in (0, 1)$ if $d(x, P) \leq \varepsilon f(x)$ for every point $x \in \Sigma$. The *nearest point map* ν maps a point $x \in \mathbb{R}^3 \setminus \mathcal{M}$ to the point $\nu(x) \in \Sigma$ closest to x. The map ν is continuous and for every point $x \notin \mathcal{M}$, the line through x and $\nu(x)$ is normal to Σ at $\nu(x)$. We use $\mathbf{n}_x$ to denote the outward unit surface normal at a point $x \in \Sigma$.

A *mesh of* Σ is a triangulated polygonal surface T such that the vertices of T are points in Σ and $|T|$ is homeomorphic to Σ. If the triangles in T have small circumradii with respect to the local feature sizes, $|T|$ does not intersect $\mathcal{M}$. In this case, the restriction of ν to $|T|$ is well-defined and we denote it by ν_T. If ν_T is continuous and bijective (the inverse of ν_T is continuous as Σ is compact), then ν_T is a homeomorphism from $|T|$ to Σ.

Some standard surface sampling results in the literature are stated in Lemma 2.1(i–iv) below and Lemma 2.1(v) follows from the 1-Lipschitzness of f.

Lemma 2.1 ([2, 9, 16]) *Let p, q and r be three distinct points on Σ.*

(i) *If $d(p, q) \leq \varepsilon f(p)$ for some $\varepsilon < 1$, then $\angle(\mathbf{n}_p, \mathbf{n}_q) \leq \varepsilon/(1-\varepsilon)$ and $\angle_a(\mathbf{n}_p, pq) \geq \arccos(\varepsilon/2)$. When $\varepsilon \leq 0.1$, $\arccos(\varepsilon/2) > \pi/2 - 0.51\varepsilon$.*

(ii) *Assume that the largest angle of the triangle pqr is at the vertex p. If $\gamma_{pqr} \leq \varepsilon f(p)$ for some $\varepsilon < 0.5$, then*

(a) $\angle_a(\mathbf{n}_p, \mathbf{n}_{pqr}) \leq \theta_\varepsilon = \arcsin(\varepsilon) + \arcsin\left(\frac{2}{\sqrt{3}} \sin(2 \arcsin(\varepsilon))\right)$,

(b) $\angle_a(\mathbf{n}_q, \mathbf{n}_{pqr})$ *and* $\angle_a(\mathbf{n}_r, \mathbf{n}_{pqr})$ *are at most* $\phi_\varepsilon = \theta_\varepsilon + \frac{2\varepsilon}{1-2\varepsilon}$.

When $\varepsilon \leq 0.1$, $\theta_\varepsilon < 3.5\varepsilon$ *and* $\phi_\varepsilon < 5.83\varepsilon$.

(iii) *For every point* z *on the tangent plane of* Σ *at* p, *if* $d(p,z) = \varepsilon f(p)$ *for some* $\varepsilon \leq 0.25$, *then* $d(z,\nu(z)) < 2\varepsilon^2 f(p) = 2\varepsilon\, d(p,z)$.

(iv) *If* p *belongs to an* ε*-sample* P *of* Σ *for some* $\varepsilon < 1$, *then* $d(p, P \setminus \{p\}) \leq \frac{2\varepsilon}{1-\varepsilon} f(p)$.

(v) *If* $d(p,q) \leq \varepsilon_1 f(p) + \varepsilon_2 f(q)$ *for some* $\varepsilon_1 < 1$ *and some* $\varepsilon_2 < 1$, *then* $f(p) \leq \frac{1+\varepsilon_2}{1-\varepsilon_1} f(q)$ *and* $d(p,q) \leq \frac{\varepsilon_1+\varepsilon_2}{1-\varepsilon_1} f(q)$.

Lemma 2.2(i, ii) below state several extensions of Lemma 2.1(iii). They show that given a point $x \in \mathbb{R}^3$ and a point $p \in \Sigma$, if $d(p,x) = O(\varepsilon f(p))$ and px is almost orthogonal to $\mathbf{n}_p$, then x is at distance $O(\varepsilon^2 f(p))$ from Σ.

Lemma 2.2 *Let* p, q *and* r *be any three points on* Σ. *Let* c_0, c_1 *and* c *be any positive values.*

(i) *For every point* $x \in \mathbb{R}^3$, *if* $d(p,x) \leq c_0\varepsilon f(p)$ *and* $\angle_a(\mathbf{n}_p, px) \geq \pi/2 - c_1\varepsilon$ *for some* $\varepsilon < \min\left\{1, \frac{1}{4c_0}, \frac{\pi}{2c_1}\right\}$, *then* $d(x,\nu(x)) \leq (2c_0 + c_1)\varepsilon\, d(p,x) \leq c_0(2c_0+c_1)\varepsilon^2 f(p)$.

(ii) *For every point* $z \in \Sigma$, *if* $d(p,z) \leq c\varepsilon f(p)$ *for some* $\varepsilon \leq \frac{1}{10c}$, *then for every point* $x \in pz$, $d(x,\nu(x)) \leq 2.51c\varepsilon\, d(p,x) \leq 2.51c^2\varepsilon^2 f(p)$.

(iii) *Suppose that* $\gamma_{pqr} \leq c\varepsilon f(p)$ *for some* $\varepsilon < \min\left\{1, \frac{1}{72c}\right\}$. *Then, for every point* $x \in \mathrm{aff}(pqr) \cap B(c_{pqr}, \gamma_{pqr})$, $d(x,\nu(x)) \leq 10c\varepsilon\, d(p,x) \leq 20c^2\varepsilon^2 f(p)$ *and* $d(p,\nu(x)) \leq (2c\varepsilon + 20c^2\varepsilon^2) f(p)$.

Proof. Consider (i). Let x' be the projection of x onto the tangent plane of Σ at p. So $d(p,x') \leq d(p,x) \leq c_0\varepsilon f(p)$. Also, $d(x,x') \leq d(p,x)\sin(c_1\varepsilon) \leq c_1\varepsilon\, d(p,x)$. By Lemma 2.1(iii), $d(x',\nu(x')) \leq 2c_0\varepsilon\, d(p,x') \leq 2c_0\varepsilon\, d(p,x)$. The shortest connection from x to Σ is the segment connecting x and $\nu(x)$, which is not longer than the path that moves linearly from x to x' and then to $\nu(x')$. Therefore, $d(x,\nu(x)) \leq d(x,x') + d(x',\nu(x')) \leq (2c_0+c_1)\varepsilon\, d(p,x) \leq c_0(2c_0 + c_1)\varepsilon^2 f(p)$. This proves (i).

By Lemma 2.1(i), $\angle_a(\mathbf{n}_p, pz) \geq \arccos(c\varepsilon/2) > \pi/2 - 0.51c\varepsilon$. Then, applying (i) with $c_0 = c$ and $c_1 = 0.51c$, we conclude that for every point $x \in pz$, $d(x,\nu(x)) \leq 2.51c\varepsilon\, d(p,x) \leq 2.51c^2\varepsilon^2 f(p)$. This proves (ii).

Without loss of generality, assume that the largest angle of pqr is at q. Since $d(p,q) \leq 2\gamma_{pqr} \leq 2c\varepsilon f(p)$, Lemma 2.1(v) implies that $f(p) \leq \frac{1}{1-2c\varepsilon} f(q)$ and so $\gamma_{pqr} \leq \frac{c\varepsilon}{1-2c\varepsilon} f(q)$. By Lemma 2.1(ii)(b), $\angle_a(\mathbf{n}_p, \mathbf{n}_{pqr}) < \frac{5.83c\varepsilon}{1-2c\varepsilon}$ which is less than $6c\varepsilon$ by the assumption of $\varepsilon < \frac{1}{72c}$. Therefore, $\angle_a(\mathbf{n}_p, px) \geq \angle_a(\mathbf{n}_p, pqr) = \pi/2 - \angle_a(\mathbf{n}_p, \mathbf{n}_{pqr}) > \pi/2 - 6c\varepsilon$. Since $d(p,x) \leq 2\gamma_{pqr} \leq 2c\varepsilon f(p)$, we can apply (i) with $c_0 = 2c$ and $c_1 = 6c$, which gives $d(x,\nu(x)) \leq 10c\varepsilon\, d(p,x) \leq 20c^2\varepsilon^2 f(p)$. Therefore, $d(p,\nu(x)) \leq d(p,x) + d(x,\nu(x)) \leq (2c\varepsilon + 20c^2\varepsilon^2) f(p)$. This proves (iii). $\square$

3 Surface meshes

We propose a class of surface meshes defined as follows.

Definition 1 *For every* $\varepsilon \in (0,1)$ *and every constant* $\alpha \in (0, \pi/3]$, *an* $(\boldsymbol{\varepsilon}, \boldsymbol{\alpha})$***-mesh*** *of* Σ *is a triangulation* T *that satisfies the following conditions.*

- *The vertices of T form an ε-sample of Σ.*
- *The angles of every triangle in T are at least α.*
- *There exist a triangle τ in T and a vertex p of τ such that $\angle_a(\mathbf{n}_p, \mathbf{n}_\tau) \leq \arcsin\left(\frac{0.8}{1+2\csc(\alpha/2)}\right)$.*
- *ν_T is a homeomorphism from $|T|$ to Σ.*

Despite the constant lower bound α on the triangle angles, the vertex set of an (ε, α)-mesh may not be a *locally uniform sample* [8, 10, 14] because the nearest neighbor distance of a vertex p is not necessarily $\Omega(\varepsilon f(p))$ and the number of vertices at distance $O(\varepsilon f(p))$ or less from p may not be bounded from above by a constant. The third condition in Definition 1 requires a triangle τ and a vertex p of τ such that $\mathbf{n}_\tau$ is a reasonable approximation of $\mathbf{n}_p$. For example, if $\gamma_\tau = O(\varepsilon f(p))$, then τ satisfies the third condition by Lemma 2.1(ii), provided that ε is sufficiently small. We also verified the third condition on two models, bunny and armadillo. Figure 3 gives more detailed information. (Since the smooth surfaces are unknown, the surface normals at the mesh vertices are estimated by local principal component analysis.) We conjecture that the third condition is a consequence of the other three conditions in Definition 1 as long as ε is small enough.

23503 vertices
$\alpha = 10.050482°$
$\min \angle_a(\mathbf{n}_p, \mathbf{n}_\tau) = 0.091873°$
Bound in the 3rd condition: $\leq 1.923°$

73706 vertices
$\alpha = 4.138016°$
$\min \angle_a(\mathbf{n}_p, \mathbf{n}_\tau) = 0.082493°$
Bound in the 3rd condition: $\leq 0.8127°$

Figure 3: Statistics on checking the third condition in Definition 1.

We study some geometric and combinatorial properties of (ε, α)-meshes in the rest of this section. We first prove some properties of a vertex star in Lemma 3.1 below, assuming that the vertices are dense and there is a small triangle in the vertex star. These properties give bounds on the number of triangles, triangle circumradii, normal deviations, and dihedral angles within the star. Recall that the definition of (ε, α)-mesh requires ε to be a number from $(0, 1)$ and α to be an angle from $(0, \pi/3]$.

Lemma 3.1 *Let $\mu_1 = \mu_0(\csc\alpha)^{4\pi/\alpha - 1}$ for some $\mu_0 \geq 1$. Let T be an (ε, α)-mesh of Σ such that $\varepsilon \leq \min\left\{\frac{1}{72\mu_1}, \frac{\sin(\alpha/4)}{6\mu_1}\right\}$. If a vertex $p \in T$ is incident to a triangle with circumradius at most $\mu_0 \varepsilon f(p)$, then the following properties are satisfied.*

(i) *For every triangle $\tau \in \mathrm{star}(p)$, $\gamma_\tau \leq \mu_1 \varepsilon f(p)$ and $\angle_a(\mathbf{n}_p, \mathbf{n}_\tau) < \frac{6\gamma_\tau}{f(p)} \leq 6\mu_1\varepsilon$.*

(ii) *There are at most $4\pi/\alpha$ edges in $\mathrm{star}(p)$.*

(iii) *For every triangle $\tau \in \mathrm{star}(p)$ and every point $x \in \tau$, $\angle_a(\mathbf{n}_{\nu(x)}, \mathbf{n}_\tau) < 9\mu_1\varepsilon$.*

(iv) *For every pair of triangles* $\sigma, \tau \in \mathrm{star}(p)$ *that share an edge, the dihedral angle at* $\sigma \cap \tau$ *is greater than* $\pi - \frac{6\gamma_\sigma + 6\gamma_\tau}{f(p)} \geq \pi - 12\mu_1\varepsilon$.

(v) *Let* h *be the distance between* p *and its nearest vertex. There is an edge in* $\mathrm{star}(p)$ *of length less than* $1.1h \csc\alpha < \frac{2.2\varepsilon \csc\alpha}{1-\varepsilon} f(p)$.

Proof. By assumption, there is a triangle $\tau_1 \in \mathrm{star}(p)$ such that $\gamma_{\tau_1} \leq \mu_0 \varepsilon f(p)$. Let τ_2 be a triangle in $\mathrm{star}(p)$ that shares an edge with τ_1. The edge $\tau_1 \cap \tau_2$ has length at most $2\gamma_{\tau_1} \leq 2\mu_0\varepsilon f(p)$. Since the angles of τ_2 are at least α, its circumradius γ_{τ_2} is at most $\mu_0\varepsilon \csc\alpha\, f(p)$. Apply this argument to the triangles in $\mathrm{star}(p)$ in circular order starting from τ_1. Then, for the ith triangle τ_i visited, its circumradius γ_{τ_i} is at most $\mu_0\varepsilon(\csc\alpha)^{i-1} f(p)$, and its edge lengths are at most $2\mu_0\varepsilon(\csc\alpha)^{i-1} f(p)$.

For $i \leq 4\pi/\alpha$, since $\mu_0\varepsilon(\csc\alpha)^{i-1} \leq \mu_1\varepsilon$ by the definition of μ_1, Lemma 2.2(ii) implies that for every edge e of τ_i, if we take a point $x \in e$ arbitrarily close to p, then $d(x, \nu(x)) \leq 2.51 \cdot 2\mu_0\varepsilon(\csc\alpha)^{i-1} d(p,x) < 6\mu_0\varepsilon(\csc\alpha)^{i-1} d(p,x) \leq 6\mu_1\varepsilon\, d(p,x)$. This implies that e makes an angle at most $\arcsin(6\mu_1\varepsilon)$ with the line tangent to the curve $\nu(e)$ at p. Since $\arcsin(6\mu_1\varepsilon) \leq \alpha/4$ by the assumption of the lemma, we conclude that for $i \leq 4\pi/\alpha$, the angle of the curved triangle $\nu(\tau_i)$ at p is at least $\alpha - 2\arcsin(6\mu_1\varepsilon) \geq \alpha/2$. Hence, there are at most $4\pi/\alpha$ triangles in $\mathrm{star}(p)$. This establishes the correctness of (ii).

Take an arbitrary triangle $\tau \in \mathrm{star}(p)$. It follows from the above discussion that $\gamma_\tau \leq \mu_1\varepsilon f(p)$. Let q be the vertex of τ at which the angle is the largest. Since $d(p,q) \leq 2\gamma_\tau \leq 2\mu_1\varepsilon f(p)$, Lemma 2.1(v) implies that $f(p) \leq \frac{1}{1-2\mu_1\varepsilon} f(q)$. Then, Lemma 2.1(ii)(b) implies that $\angle_a(\mathbf{n}_p, \mathbf{n}_\tau) < \frac{5.83\gamma_\tau}{f(q)} \leq \frac{5.83\gamma_\tau}{(1-2\mu_1\varepsilon)f(p)}$. One can verify that $\frac{5.83}{1-2\mu_1\varepsilon} < 6$ by our assumption of $\varepsilon \leq \frac{1}{72\mu_1}$ in the lemma. Therefore, $\angle_a(\mathbf{n}_p, \mathbf{n}_\tau) < \frac{6\gamma_\tau}{f(p)} \leq 6\mu_1\varepsilon$, establishing the correctness of (i).

Consider (iii). Take a point x in any triangle $\tau \in \mathrm{star}(p)$. Since $\gamma_\tau \leq \mu_1\varepsilon f(p)$, Lemma 2.2(iii) implies that

$$
\begin{aligned}
d(x, \nu(x)) &\leq 20\mu_1^2\varepsilon^2 f(p), && (1)\\
d(p, \nu(x)) &\leq 2\mu_1\varepsilon f(p) + 20\mu_1^2\varepsilon^2 f(p) \leq 2.3\mu_1\varepsilon f(p). && (2)
\end{aligned}
$$

In (2), the bound of $2.3\mu_1\varepsilon f(p)$ follows from the assumption of $\mu_1\varepsilon \leq 1/72$ in the lemma. The inequality $\mu_1\varepsilon \leq 1/72$ or simply $\varepsilon \leq 1/72$ will be used often in the rest of the proof, but we will not make it explicit again to simplify the presentation. Since $d(p, \nu(x)) \leq 2.3\mu_1\varepsilon f(p)$, Lemma 2.1(i) implies that

$$
\angle(\mathbf{n}_p, \mathbf{n}_{\nu(x)}) \leq \frac{2.3\mu_1\varepsilon}{1 - 2.3\mu_1\varepsilon} < 3\mu_1\varepsilon. \tag{3}
$$

By (i), $\angle_a(\mathbf{n}_{\nu(x)}, \mathbf{n}_\tau) \leq \angle_a(\mathbf{n}_p, \mathbf{n}_\tau) + \angle(\mathbf{n}_p, \mathbf{n}_{\nu(x)}) < 6\mu_1\varepsilon + 3\mu_1\varepsilon = 9\mu_1\varepsilon$. This proves (iii).

Consider (iv). Let pq be the common edge of two triangles σ and τ in T. By (i), the dihedral angle at pq is either less than $\frac{6\gamma_\sigma + 6\gamma_\tau}{f(p)} \leq 12\mu_1\varepsilon$ or greater than $\pi - \frac{6\gamma_\sigma + 6\gamma_\tau}{f(p)} \geq \pi - 12\mu_1\varepsilon$. We show that the former case is impossible. Let x be a point in the interior of τ that is arbitrarily close to the midpoint of pq. Consider the line L through x and $\nu(x)$, which is normal to Σ at $\nu(x)$. By (i) and (3), $\angle_a(\mathbf{n}_\tau, L) \leq \angle(\mathbf{n}_p, \mathbf{n}_{\nu(x)}) + \angle_a(\mathbf{n}_p, \mathbf{n}_\tau) < 9\mu_1\varepsilon$. Similarly, $\angle_a(\mathbf{n}_\sigma, L) < 9\mu_1\varepsilon$. If the dihedral angle at pq is less than $12\mu_1\varepsilon$, then L would intersect σ at a point y. By (1), we get $d(y, \nu(x)) \leq d(p,x) + d(p,y) + d(x, \nu(x)) \leq (4\mu_1\varepsilon + 20\mu_1^2\varepsilon^2) f(p)$. Relating $f(p)$ to $f(\nu(x))$ using (2) and Lemma 2.1(v), we obtain $d(y, \nu(x)) \leq \frac{4\mu_1\varepsilon + 20\mu_1^2\varepsilon^2}{1-2.3\mu_1\varepsilon} f(\nu(x)) < f(\nu(x))$. Thus, $\nu(x)$ is the closest point in Σ to y. This contradicts the injectivity of ν_T because both x and y are mapped to $\nu(x)$.

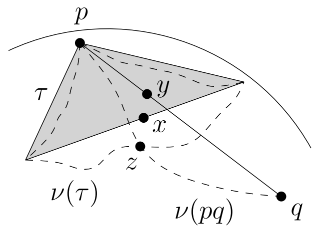

Figure 4: The dashed curved triangle is $\nu(\tau)$ and the dashed curve from p to q is $\nu(pq)$. Both x and y map to z under ν.

Consider (v). Let q be the vertex in T closest to p. By Lemma 2.1(iv), $d(p,q) \leq \frac{2\varepsilon}{1-\varepsilon} f(p)$. If pq is an edge in T, we are done. Assume that pq is not an edge in T. Then, there exists a triangle $\tau \in \text{star}(p)$ such that $\nu(pq)$ leaves $\nu(\tau)$ at a point z. Let x denote the point $\nu^{-1}(z) \cap \tau$, which lies on the edge of τ opposite p. Let y denote the point $\nu^{-1}(z) \cap pq$. Figure 4 shows an example. By Lemma 2.2(ii), $d(p,z) \leq d(p,y) + d(y,z) \leq (1 + 2.51 \cdot 2\varepsilon/(1-\varepsilon))\, d(p,y) < \frac{1+5\varepsilon}{1-\varepsilon}\, d(p,y) \leq \frac{2\varepsilon+10\varepsilon^2}{(1-\varepsilon)^2} f(p)$. Then, Lemma 2.1(i) implies that $\angle(\mathbf{n}_p, \mathbf{n}_z) \leq \frac{2\varepsilon+10\varepsilon^2}{1-4\varepsilon-9\varepsilon^2}$ which is less than 3ε by the condition on ε in the assumption of the lemma. Then, (i) implies that $\angle_a(\mathbf{n}_z, \tau) \geq \angle_a(\mathbf{n}_p, \tau) - \angle(\mathbf{n}_p, \mathbf{n}_z) > \pi/2 - (3 + 6\mu_1)\varepsilon > \pi/2 - 9\mu_1\varepsilon$. Let p' denote the projection of p onto the tangent plane H of Σ at z. Then, $p'z$ is the common projection of px and py onto H. As a result, $d(p,x) = d(p',z) \sec(\angle_a(px, H)) \leq d(p',z) \sec(\pi/2 - \angle_a(\mathbf{n}_z, \tau)) \leq d(p,y) \sec(9\mu_1\varepsilon) < 1.1\, d(p,q)$. The lengths of the two sides of τ in star(p) are at most the height of τ from p divided by $\sin\alpha$, which is at most $d(p,x) \csc\alpha < 1.1 \csc\alpha\, d(p,q) < \frac{2.2\varepsilon \csc\alpha}{1-\varepsilon} f(p)$. This establishes the correctness of (v). $\square$

Next, we prove in Lemma 3.2 below that the properties in Lemma 3.1 in fact hold for every vertex, provided that ε is sufficiently small. This subclass of (ε, α)-meshes is the subject of study in the rest of this paper. We define them formally as follows.

Definition 2 *Let* $\mu_0 = \max\left\{4(\csc\alpha)^{4\pi/\alpha+2},\, 40(\csc\alpha)^2 \csc(\alpha/2)\right\}$. *Let* $\mu_1 = \mu_0 (\csc\alpha)^{4\pi/\alpha-1}$. *An* (ε,α)*-mesh of* Σ *is* ***dense*** *if* $\varepsilon \leq \min\left\{\frac{1}{72\mu_1}, \frac{\sin(\alpha/4)}{6\mu_1}, \frac{(\sin(\alpha/2))^2}{800+420\sin(\alpha/2)}\right\}$.

The above definition of a dense mesh requires ε to be pessimistically small (e.g. $\varepsilon \leq 6.14 \times 10^{-18}$ when $\alpha = \pi/6$) in order to facilitate the subsequent theoretical analysis. However, our analysis is probably not tight. Also, the edge flip procedure is oblivious of the value of ε and produces good experimental results as reported in Section 7.

Lemma 3.2 *Every dense* (ε,α)*-mesh* T *of* Σ *satisfies the following properties, where* μ *denotes the constant* $2(\csc\alpha)^{4\pi/\alpha+1}$.

(i) *For each vertex* $p \in T$ *and every triangle* $\tau \in \text{star}(p)$, $\gamma_\tau \leq \mu\varepsilon f(p)$ *and* $\angle_a(\mathbf{n}_p, \mathbf{n}_\tau) < \frac{6\gamma_\tau}{f(p)} \leq 6\mu\varepsilon$.

(ii) *Every vertex in* T *is incident to at most* $4\pi/\alpha$ *edges.*

(iii) *For each triangle* $\tau \in T$ *and each point* $x \in \tau$, $\angle_a(\mathbf{n}_{\nu(x)}, \mathbf{n}_\tau) < 9\mu\varepsilon$.

(iv) *For every pair of triangles* $\sigma, \tau \in T$ *that share an edge, the dihedral angle at* $\sigma \cap \tau$ *is greater than* $\pi - \frac{6\gamma_\sigma + 6\gamma_\tau}{f(p)} \geq \pi - 12\mu\varepsilon$.

(v) *For every vertex* $p \in T$, *there is an edge in* star(p) *of length less than* $1.1h\csc\alpha < \frac{2.2\varepsilon\csc\alpha}{1-\varepsilon} f(p)$, *where* h *is the distance between* p *and its nearest vertex.*

Proof. By the third condition in the definition of (ε, α)-meshes, there exist a triangle $\tau \in T$ and a vertex p of τ such that

$$\theta = \angle_a(\mathbf{n}_p, \mathbf{n}_\tau) \leq \arcsin\left(\frac{0.8}{1 + 2\csc(\alpha/2)}\right). \tag{4}$$

Our plan is to prove that $\gamma_\tau \leq \mu_0 \varepsilon f(p)$ so that Lemma 3.1 can be applied to p. Then, we tighten the factor μ_1 in Lemma 3.1(i, iii, iv) to μ and extend the argument to other vertices.

Let $\delta = \frac{40\varepsilon}{\sin(\alpha/2) - 20\varepsilon}$. Using the definition of δ and the condition $\varepsilon \leq \frac{(\sin(\alpha/2))^2}{800 + 420\sin(\alpha/2)}$ in the definition of dense (ε, α)-meshes, one can verify that

$$\delta \leq \frac{0.1}{2\csc(\alpha/2) + 1}. \tag{5}$$

We claim that the height of τ from p is at most $\delta f(p)$.

Claim 3.1 *The height of τ from p is at most $\delta f(p)$.*

Proof. Suppose for the sake of contradiction that the height of τ from p is greater than $\delta f(p)$. Let C be the circle in $\mathrm{aff}(\tau)$ such that its center x lies on the bisector of the angle of τ at p, $d(p, x) = 0.25\delta f(p)$, and the radius of C is $d(p, x)\sin(\alpha/2)$. Note that $x \in \tau$ and $C \subset \tau$ by the assumption that the height of τ from p is greater than $\delta f(p)$. We show below that $\nu(C)$ contains a large empty region, which contradicts the property of ε-sampling.

Let x' be the projection of x onto the tangent plane of Σ at p. Then, $d(x, \nu(x)) \leq d(x, x') + d(x', \nu(x'))$, which is at most $d(x, x') + 0.5\delta\, d(p, x')$ by Lemma 2.1(iii) because $d(p, x') \leq d(p, x) = 0.25\delta f(p)$. Also, $d(x, x') \leq d(p, x)\sin\theta$. Therefore,

$$d(x, \nu(x)) \leq d(p, x)\sin\theta + 0.5\delta\, d(p, x') \leq (\sin\theta + 0.5\delta)\, d(p, x). \tag{6}$$

Similarly, for every point y in the circumference of C, we can derive that $d(y, \nu(y)) \leq (\sin\theta + \delta)\, d(p, y) \leq (\sin\theta + \delta)(d(p, x) + d(x, y))$. It follows that

$$\begin{aligned} d(\nu(x), \nu(y)) &\geq d(x, y) - d(x, \nu(x)) - d(y, \nu(y)) \\ &\geq d(x, y) - (\sin\theta + 0.5\delta)\, d(p, x) - (\sin\theta + \delta)(d(p, x) + d(x, y)). \end{aligned}$$

Notice that $d(x, y) = \mathrm{radius}(C) = d(p, x)\sin(\alpha/2)$. Therefore,

$$\begin{aligned} d(\nu(x), \nu(y)) &\geq d(x, y) - (\sin\theta + 0.5\delta)\csc(\alpha/2)\, d(x, y) - \\ &\quad (\sin\theta + \delta)(1 + \csc(\alpha/2))\, d(x, y) \\ &> d(x, y)(1 - (\sin\theta + \delta)(2\csc(\alpha/2) + 1)). \end{aligned} \tag{7}$$

By (7), (4) and (5),

$$d(\nu(x), \nu(y)) > 0.1\, d(x, y) = 0.1\sin(\alpha/2)\, d(p, x) = 0.025\delta\sin(\alpha/2)\, f(p). \tag{8}$$

By (6),

$$d(p, \nu(x)) \leq d(p, x) + d(x, \nu(x)) \leq d(p, x)(1 + \sin\theta + 0.5\delta).$$

Plug (5) and (4) into the inequality above. It gives the relation $d(p, \nu(x)) \leq \left(1 + \frac{0.85}{2\csc(\alpha/2)+1}\right) d(p, x) = 0.25\delta\left(1 + \frac{0.85}{2\csc(\alpha/2)+1}\right) f(p)$, which is less than $0.5\delta f(p)$

because $\frac{0.85}{1+2\csc(\alpha/2)} < 1$. Thus, $f(\nu(x)) \le (1+0.5\delta)f(p)$ by Lemma 2.1(v). Substituting this into (8) gives

$$d(\nu(x),\nu(y)) > \frac{0.025\delta\sin(\alpha/2)}{1+0.5\delta} f(\nu(x)) = \varepsilon f(\nu(x)). \tag{9}$$

The last step follows from the definition of δ. The closed curve $\nu(C)$ encloses a topological disk $D \subset \Sigma$ that contains $\nu(x)$, and D does not contain any sample point because $D \subset \nu(\tau)$. It is also known that $\Sigma \cap B(\nu(x), \varepsilon f(\nu(x)))$ is a topological disk [5]. By (9), $d(\nu(x),\nu(C)) > \varepsilon f(\nu(x))$, which implies that $\Sigma \cap B(\nu(x), \varepsilon f(\nu(x))) = D \cap B(\nu(x), \varepsilon f(\nu(x)))$. But then $\Sigma \cap B(\nu(x), \varepsilon f(\nu(x)))$ contains no sample point as D contains none, contradicting the definition of ε-sampling. This establishes the claim that the height of τ from p is at most $\delta f(p)$. □

By Claim 3.1, the two edges of τ in star(p) are at most $\delta \csc\alpha\, f(p)$ long. Therefore, $\gamma_\tau \le 0.5\delta(\csc\alpha)^2 f(p)$. Using the conditions on μ_0 and ε in the definition of dense (ε,α)-meshes, one can verify that $0.5\delta(\csc\alpha)^2 \le \mu_0\varepsilon$ and, therefore, $\gamma_\tau \le \mu_0\varepsilon f(p)$. This makes Lemma 3.1 applicable to p. We get a sharper bound on the circumradii in star(p) as follows.

By Lemma 3.1(v), there is an edge in star(p) of length less than $4\varepsilon\csc\alpha f(p)$ and, hence, there is a triangle in star(p) with circumradius less $2\varepsilon(\csc\alpha)^2 f(p)$. The longest edge length of this triangle is less than $4\varepsilon(\csc\alpha)^2 f(p)$. The next triangle in cyclic order around p has circumradius less than $2\varepsilon(\csc\alpha)^3 f(p)$ and longest edge length less than $4\varepsilon(\csc\alpha)^3 f(p)$. Continuing with this reasoning and applying Lemma 3.1(ii), we deduce that every triangle in star(p) has circumradius less than $2\varepsilon(\csc\alpha)^{4\pi/\alpha+1} f(p) = \mu\varepsilon f(p)$ and longest edge length less than $2\mu\varepsilon f(p)$. This results in the $\mu\varepsilon f(p)$ bound stated in (i), which is an improvement of the $\mu_1\varepsilon f(p)$ bound stated in Lemma 3.1(i).

We proceed to bound the circumradii of triangles incident to other vertices. Take any neighbor q of p, since $d(p,q) \le 4\varepsilon(\csc\alpha)^{4\pi/\alpha+1} f(p)$, Lemma 2.1(v) implies that $d(p,q) \le 4\varepsilon(\csc\alpha)^{4\pi/\alpha+1} f(q)/(1-4\varepsilon(\csc\alpha)^{4\pi/\alpha+1})$, which is at most $2\mu_0\varepsilon\sin\alpha\, f(q)$ by the conditions on μ_0 and ε in the definition of dense (ε,α)-meshes. Therefore, the circumradius of any triangle incident to pq is at most $\mu_0\varepsilon f(q)$, implying that Lemma 3.1 holds for q. Moreover, we can apply the reasoning in the previous paragraph again to show that every triangle in star(q) has circumradius less than $\mu\varepsilon f(q)$ and longest edge length less than $2\mu\varepsilon f(q)$. As a result, starting from p, we can traverse T in a breadth first manner to deduce that Lemma 3.1 holds for all vertices of T and bound the circumradii in the star of every vertex v by $\mu\varepsilon f(v)$.

Once the factor μ_1 in the circumradius bound in Lemma 3.1(i) is improved to μ, one can verify that every factor μ_1 in Lemma 3.1(i, iii, iv) is also improved to μ as well. □

Lemma 3.2(i, iv) show that a decrease in the triangle circumradii improves the bounds on the normal deviations and dihedral angles. Although Lemma 3.2(i) states that the triangle circumradii are bounded by $\mu\varepsilon f(p)$, the factor μ is large in comparison with the factor $\varepsilon + O(\varepsilon^\kappa)$, where $\kappa \in (1, 1.5)$, that is guaranteed by repeated edge flips as shown in Sections 4–6 later. The value ε is required to be very small in Definition 2, but our experimental results in Section 7 show that edge flips work well even if the vertices are not very dense.

4 Edge flips

There are three results in this section. Lemma 4.1 shows that dense (ε,α)-meshes are closed under the operation of flipping a *flippable* edge (to be defined shortly). Lemmas 4.2 and 4.3

show that a linear number edge flips suffice to ensure that no flippable edge remains. Theorem 1 is a corollary of our techniques and it states that a planar triangulation with all angles greater than some constant can be flipped to the Delaunay triangulation in linear time. Throughout this section, let μ denote the constant $2(\csc\alpha)^{4\pi/\alpha+1}$ as defined in Lemma 3.2.

Flipping the common edge pq of two triangles pqr and pqs means replacing them by two new triangles prs and qrs. We flip pq only if it satisfies the following criterion, which ensures that the minimum angle in the two new triangles produced by the edge flip is greater than the minimum angle in the two triangles before the edge flip.

Definition 3 *Let pqr and pqs be two triangles in a triangulated polygonal surface. The edge pq is* ***flippable*** *if and only if*

(i) $\angle prs > \min\{\pi/2, \angle pqs\}$ *and* $\angle psr > \min\{\pi/2, \angle pqr\}$,

(ii) $\angle qrs > \min\{\pi/2, \angle qps\}$ *and* $\angle qsr > \min\{\pi/2, \angle qpr\}$,

(iii) $\angle rps > \max\{\angle rpq, \angle spq\}$ *and* $\angle rqs > \max\{\angle rqp, \angle sqp\}$,

(iv) $\angle prq > \max\{\angle prs, \angle qrs\}$ *and* $\angle psq > \max\{\angle psr, \angle qsr\}$,

(v) *rs is currently not an edge in the triangulated polygonal surface.*

To test the flippability of an edge according to Definition 3, one can compare the cosines of the angles involved, which can be done by comparing the inner products of some vectors of the form $u - v$, where u and v are appropriate mesh vertices. It is well-known in the plane that if s lies inside the circumcircle of pqr, then r lies inside the circumcircle of pqs and flipping pq increases the minimum angle. However, in three dimensions, it is possible that s lies inside the diametic ball of pqr, while r lies outside the diametric ball of pqs. This motivates us to define flippability to increase the minimum angle directly. This also gives an aggressive strategy to improve the mesh quality.

We show in Lemma 4.1 below that dense (ε, α)-meshes are closed under the flipping of flippable edges.

Lemma 4.1 *Let T be a dense (ε, α)-mesh of Σ. If an edge $pq \in T$ is flippable, then flipping pq produces a dense (ε, α)-mesh.*

Proof. The vertex set remains an ε-sample after the edge flip. The first condition in Definition 1 and the condition on ε in Definition 2 are thus satisfied. Definition 3 implies that all angles are at least α after flipping pq. It remains to check the third and fourth conditions in Definition 1. Let pqr and pqs be the triangles incident to pq.

By Lemma 3.2(v), there is an edge e in star(r) of length less than $4\varepsilon \csc\alpha\, f(r)$. The edge e is not affected by flipping pq. Let σ be any triangle incident to e after flipping pq. The circumradius of σ is at most length$(e)/(2\sin\alpha) < 2\varepsilon(\csc\alpha)^2 f(r)$. Let v be the vertex of σ at which the angle is the largest. Then $d(r, v) \le 2\gamma_\sigma < 4\varepsilon(\csc\alpha)^2 f(r)$. By Lemma 2.1(v), $f(r) \le \frac{1}{1-4\varepsilon(\csc\alpha)^2} f(v)$ which implies that $\gamma_\sigma < \frac{2\varepsilon(\csc\alpha)^2}{1-4\varepsilon(\csc\alpha)^2} f(v)$. By Lemma 2.1(ii)(a), $\angle_a(\mathbf{n}_v, \mathbf{n}_\sigma) < \frac{7\varepsilon(\csc\alpha)^2}{1-4\varepsilon(\csc\alpha)^2}$, which is less than $\arcsin\left(\frac{0.8}{1+2\csc(\alpha/2)}\right)$ by the condition on ε in the definition of dense (ε, α)-meshes. Therefore, the third condition in Definition 1 is satisfied.

Before we establish the fourth condition in Definition 1, we first bound the circumradii of the triangles in star(r). The bound on the number of triangles in star(r) may increase from $4\pi/\alpha$ to $4\pi/\alpha + 1$ after flipping pq. Nevertheless, we have shown in the previous paragraph that the two triangles incident to e (after flipping pq) have circumradii less than $2\varepsilon(\csc\alpha)^2 f(r)$.

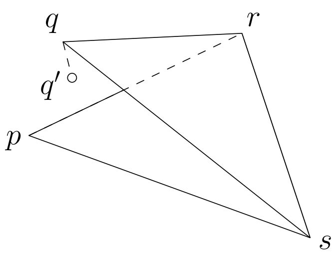

Figure 5: If the dihedral angle at rs is less than $12\mu\varepsilon$, then p and q' lie on the same side of rs as p and $\angle qrq' < 12\mu\varepsilon$.

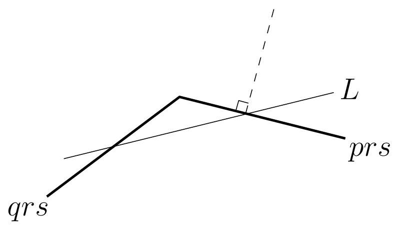

Figure 6: The bold segments represent the cross-sections of the triangles prs and qrs. The dihedral angle between the two triangles is obtuse. For L to intersect both prs and qrs, $\angle_a(L, prs)$ must be less than π minus the dihedral angle between prs and qrs.

Therefore, we can go around the triangles in star(r) and repeat the analysis to show that their circumradii are less than $2\varepsilon(\csc\alpha)^{4\pi/\alpha+1}f(r) = \mu\varepsilon f(r)$.

It remains to verify that the nearest point map ν restricted to the updated mesh is a homeomorphism onto Σ. As shown in the previous paragraph, both γ_{prs} and γ_{qrs} are less than $\mu\varepsilon f(r)$, which is less than $0.5f(r)$ by the condition on ε in the definition of dense (ε, α)-meshes. Therefore, the updated mesh avoids the medial axis of Σ, which means that the restriction of ν to the updated mesh is well-defined. Since $prs \cup qrs$ and $pqr \cup pqs$ share the same boundary, it suffices to prove that the restriction of ν to $prs \cup qrs$ is injective and $\nu(prs \cup qrs) = \nu(pqr \cup pqs)$.

We first establish several properties of prs and qrs. Although we have not proved that the updated mesh is an (ε, α)-mesh, since γ_{prs} and γ_{qrs} are less than $\mu\varepsilon f(r)$, the proofs of Lemma 3.1(i, iii) can be applied to the triangles prs and qrs to give:

$$\angle_a(\mathbf{n}_r, \mathbf{n}_{prs}) < 6\mu\varepsilon, \quad \angle_a(\mathbf{n}_r, \mathbf{n}_{qrs}) < 6\mu\varepsilon \tag{10}$$

$$\forall\, x \in prs, \ \ \angle_a(\mathbf{n}_{\nu(x)}, \mathbf{n}_{prs}) < 9\mu\varepsilon \tag{11}$$

Let ϕ denote the dihedral angle at rs after flipping pq. By (10), either $\phi < 12\mu\varepsilon$ or $\phi > \pi - 12\mu\varepsilon$. We show that the first case is impossible. Assume to the contrary that $\phi < 12\mu\varepsilon$. Refer to Figure 5. Let q' be the projection of q onto aff(prs), which must lie on the same side of rs as p by the assumption of $\phi < 12\varepsilon$. Then, $\angle qrq' \leq \phi < 12\mu\varepsilon$. Therefore, $\angle prq \leq \angle prq' + \angle qrq' = |\angle prs - \angle q'rs| + \angle qrq' \leq |\angle prs - \angle qrs| + 2 \cdot \angle qrq' < |\angle prs - \angle qrs| + 24\mu\varepsilon = \max\{\angle prs, \angle qrs\} - \min\{\angle prs, \angle qrs\} + 24\mu\varepsilon$. Because $\angle prs$ and $\angle qrs$ are at least α, the value of $-\min\{\angle prs, \angle qrs\} + 24\mu\varepsilon$ is negative, which implies that $\angle prq < \max\{\angle prs, \angle qrs\}$. This is a contradiction to the flippability of pq. We conclude that

$$\phi > \pi - 12\mu\varepsilon. \tag{12}$$

Next, we show that ν is injective on $prs \cup qrs$. Suppose to the contrary that there exist two points $x, y \in prs \cup qrs$ such that $\nu(x) = \nu(y)$. Without loss of generality, assume that

$x \in prs$. Let L be the line through $\nu(x)$ and parallel to $\mathbf{n}_{\nu(x)}$. Both x and y are on L, so y must belong to qrs. Consider the plane that contains L and is perpendicular to prs. This plane contains a triangle Δ (shown shaded in Figure 6) that is bounded by L, prs and qrs. The angle of Δ at rs is minimized and equal to the dihedral angle ϕ between prs and qrs when Δ is orthogonal to $\mathrm{aff}(rs)$. Otherwise, the angle of Δ at rs can be bigger than ϕ. Since $\angle_a(L, prs)$ is another angle of Δ, $\angle_a(L, prs)$ must be less than π minus the angle of Δ at rs, implying that $\angle_a(L, prs) < \pi - \phi$. It follows from (12) that $\angle_a(L, \mathbf{n}_{prs}) = \pi/2 - \angle_a(L, prs) > \pi/2 - 12\mu\varepsilon$. On the other hand, $\angle_a(L, \mathbf{n}_{prs}) = \angle_a(\mathbf{n}_{\nu(x)}, \mathbf{n}_{prs}) < 9\mu\varepsilon$ by (11). This is a contradiction because $\pi/2 - 12\mu\varepsilon > 9\mu\varepsilon$. Therefore, ν is injective on $prs \cup qrs$.

Next, we show that $\nu(prs \cup qrs) = \nu(pqr \cup qrs)$. Let $C = pr \cup ps \cup qr \cup qs$ denote the common boundary of $prs \cup qrs$ and $pqr \cup pqs$. A continuous bijective map from a compact space is a homeomorphism onto its image. Therefore, $\nu(prs \cup qrs)$ and $\nu(pqr \cup pqs)$ are topological disks with a common boundary $\nu(C)$. Since $\nu(prs \cup qrs)$ and $\nu(pqr \cup pqs)$ have the same boundary, either $\nu(prs \cup qrs) = \nu(pqr \cup pqs)$ or the interiors of $\nu(prs \cup qrs)$ and $\nu(pqr \cup pqs)$ are disjoint. We prove that the latter is impossible. Take any point $x \in prs \setminus C$. The four triangles prs, qrs, pqr and pqs form the boundary of the tetrahedron $pqrs$. The line through x and $\nu(x)$ must enter $pqrs$ at x and exit $pqrs$ at another point, say $y \in pqr \cup pqs \cup qrs$.

Since $\gamma_{prs} \leq \mu\varepsilon f(r)$, Lemma 2.2(iii) implies that $d(r, \nu(x)) \leq (2\mu\varepsilon + 20\mu^2\varepsilon^2) f(r)$. Then, Lemma 2.1(v) implies that $f(r) \leq \frac{1}{1-2\mu\varepsilon-20\mu^2\varepsilon^2} f(\nu(x))$. If $y \in qrs$, then since $\gamma_{qrs} \leq \mu\varepsilon f(r)$, we obtain $d(y, \nu(x)) \leq d(r, y) + d(r, \nu(x)) \leq 2\gamma_{qrs} + (2\mu\varepsilon + 20\mu^2\varepsilon^2) f(r) \leq (4\mu\varepsilon + 20\mu^2\varepsilon^2) f(r) \leq \frac{4\mu\varepsilon+20\mu^2\varepsilon^2}{1-2\mu\varepsilon-20\mu^2\varepsilon^2} f(\nu(x))$, which is less than $f(\nu(x))$ by the condition on ε in the definition of dense (ε, α)-meshes. But this implies that $\nu(x)$ is the closest point in Σ to y (i.e., $\nu(y) = \nu(x)$), contradicting the injectivity of ν on $prs \cup qrs$. We conclude that $y \in pqr \cup pqs$. Since γ_{pqr} and γ_{pqs} are at most $\mu\varepsilon f(p)$ by Lemma 3.2(i), we can repeat the analysis above with r replaced by p to obtain $d(y, \nu(x)) \leq \frac{4\mu\varepsilon+20\mu^2\varepsilon^2}{1-2\mu\varepsilon-20\mu^2\varepsilon^2} f(\nu(x)) < f(\nu(x))$. It follows that $\nu(x)$ is the closest point in Σ to y, i.e., $\nu(y) = \nu(x)$. As a result, the interiors of $\nu(pqr \cup pqs)$ and $\nu(prs \cup qrs)$ are not disjoint, which implies that $\nu(prs \cup qrs) = \nu(pqr \cup pqs)$. □

Lemma 4.1 shows that edge flips keep the mesh within the class of dense (ε, α)-meshes. But will the edge flips terminate? If so, how long will the process take? The termination of the edge flips can be argued as follows. The *angle vector* of a mesh is the list of all the angles in the triangles sorted in nondecreasing order. After an edge flip, the minimum angle in the two new triangles is greater than the minimum angle in the two old triangles according to Definition 3. Therefore, an edge flip increases the angle vector of the surface mesh lexicographically. It follows that the same mesh cannot be generated more than once by successive edge flips. Hence, the edge flips must terminate. We show in the next two lemmas that the number of such edge flips is in fact linear in the number of vertices.

Lemma 4.2 *Let T be a dense (ε, α)-mesh of Σ. Let $\mathcal{G}$ be the graph such that its nodes are the vertices of T and two nodes of $\mathcal{G}$ are connected if and only if they are connected in some dense (ε, α)-mesh of Σ that has the same vertex set as T. Then, the degree of every node of $\mathcal{G}$ is at most $(1.1\mu^2(\sin\alpha)^2 + 1)^2$.*

Proof. Let p be a node in $\mathcal{G}$. Take an edge $pq \in \mathcal{G}$. Thus, pq is an edge in some dense (ε, α)-mesh T' of Σ. By Lemma 3.2(v), there is an edge $e \in \mathrm{star}(p, T')$ of length less than $1.1h \csc\alpha$, where h is the distance between p and its nearest vertex. For each triangle in T' incident to e, its circumradius is less than $0.55h(\csc\alpha)^2$ and its longest side length is less than $1.1h(\csc\alpha)^2$. By Lemma 3.2(ii), there are at most $4\pi/\alpha$ triangles in $\mathrm{star}(p, T')$ and going around these triangles shows that the longest edge in $\mathrm{star}(p, T')$ has length less than $1.1h(\csc\alpha)^{4\pi/\alpha} = 0.55h\mu\sin\alpha$.

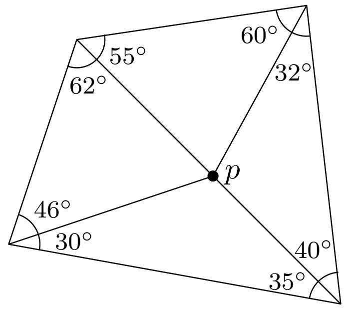

Figure 7: vector$(p) = (30°, 32°, 35°, 40°, 46°, 55°, 60°, 62°)$.

By applying the above reasoning to the vertex q, we can show that $d(p,q)$ is less than $0.55\mu \sin\alpha$ times the nearest neighbor distance of q. Therefore, $d(p,q) \leq 0.55\mu \sin\alpha\, d(q,r)$, which implies that

$$h \leq d(p,q) \leq 0.55\mu \sin\alpha\, d(q,r). \tag{13}$$

The circumradius of any triangle incident to pq in T' is at most $\mu\varepsilon f(p)$ by Lemma 3.2(i). As a result,

$$d(p,q) \leq 2\mu\varepsilon f(p). \tag{14}$$

Take another edge $pr \in \mathcal{G}$. By applying the above argument to the dense (ε, α)-mesh that contains pr, we obtain

$$d(p,r) \leq 2\mu\varepsilon f(p). \tag{15}$$

By (14) and Lemma 2.1(v), $f(p) \leq \frac{1}{1-2\mu\varepsilon} f(q) < 2f(q)$. By (14) and (15), $d(q,r) \leq d(p,q) + d(p,r) \leq 4\mu\varepsilon f(p) \leq 8\mu\varepsilon f(q)$. Lemma 2.1(i) implies that

$$\angle_a(qr, \mathbf{n}_p) \geq \angle_a(qr, \mathbf{n}_q) - \angle(\mathbf{n}_p, \mathbf{n}_q) > \pi/2 - 4.08\mu\varepsilon - \frac{2\mu\varepsilon}{1 - 2\mu\varepsilon} > \pi/2 - 7\mu\varepsilon. \tag{16}$$

Let q' and r' be the projections of q and r onto the tangent plane of Σ at p, respectively.

$$d(q', r') = d(q,r) \sin(\angle_a(qr, \mathbf{n}_p)) \overset{(13)\&(16)}{>} \frac{h\cos(7\mu\varepsilon)}{0.55\mu \sin\alpha} > h/(\mu \sin\alpha).$$

Therefore, for every pair of nodes in $\mathcal{G}$ adjacent to p, their projections onto the tangent plane of Σ at p are separated by a distance $h/(\mu\sin\alpha)$ or more. If we place a disk with radius $h/(2\mu\sin\alpha)$ and center at the projection of each node in $\mathcal{G}$ adjacent to p, such disks are disjoint. Moreover, these disks are contained in a bigger disk with center p and radius $0.55h\mu\sin\alpha + h/(2\mu\sin\alpha)$. Thus, the number of nodes in $\mathcal{G}$ adjacent to p is at most

$$\frac{\pi\left(0.55h\mu\sin\alpha + h/(2\mu\sin\alpha)\right)^2}{\pi\left(h/(2\mu\sin\alpha)\right)^2} = (1.1\mu^2(\sin\alpha)^2 + 1)^2.$$

$\square$

Lemma 4.2 implies the following result that repeated edge flips takes no more than linear time to finish.

Lemma 4.3 *Let T be a dense (ε, α)-mesh of Σ. Edge flips can be applied to T repeatedly until no edge is flippable in time linear in the number of vertices in T.*

Proof. Let n denote the number of vertices in T. By Lemma 3.2(ii), T contains $O(n)$ edges. We first scan all edges in T in $O(n)$ time and put the flippable ones into a queue.

We dequeue edges one by one. For each dequeued edge e, if e is still present and flippable in the current mesh, then we flip it. (Since dense (ε, α)-meshes are closed under edge flips by Lemma 4.1 and every vertex has constant degree by Lemma 3.2(ii), it takes $O(1)$ time to check the existence of e.) If e is flipped, two new triangles are produced and we enqueue the boundary edges of the union of the two new triangles that are flippable. The above procedure repeats until the queue becomes empty. The running time is thus $O(n + \text{number of edge flips})$.

By Lemma 4.1, all intermediate meshes are dense (ε, α)-meshes. To analyze the effect of an edge flip on a vertex star, for every vertex p, we define vector(p) to be the list of angles opposite p in the triangles in star(p), sorted in nondecreasing order. Figure 7 shows an example. If the edge pq between two triangles pqr and pqs is flipped, Definition 3 implies that for every $v \in \{p, q, r, s\}$, the minimum angle opposite v in prs and qrs is greater than the minimum angle opposite v in pqr and pqs. It implies that for every $v \in \{p, q, r, s\}$, vector(v) increases lexicographically after flipping pq. As a result, when the star of a vertex p is changed by an edge flip, the same star of p cannot be reproduced later. By Lemma 4.2, there are $O(1)$ vertices that may become neighbors of p at some point during the edge flips. They form only $O(1)$ different stars of p, i.e., the star of p can be changed only $O(1)$ times. Summing over the n vertices shows that there are $O(n)$ edges flips. $\square$

The proof techniques of Lemmas 4.2 and 4.3 can be used to prove our first main result, namely, any planar triangulation with angles greater than a constant can be flipped to the Delaunay triangulation in linear time. This helps to explain why it is fast in practice to convert a planar triangulation to the Delaunay triangulation by edge flips.

Theorem 1 *Given a triangulation of n points in $\mathbb{R}^2$ with angles greater than a constant independent of n, it can be converted to the Delaunay triangulation by edge flips in $O(n)$ time.*

Proof. In converting a planar triangulation to the Delaunay triangulation, the empty circumcircle criterion is often applied. Let pqr and pqs be two triangles in the triangulation that share the edge pq. By the empty circumcircle criterion, we can flip pq if r is enclosed by the circumcircle of pqs (in which case, s is also enclosed by the circumcircle of pqr).

It is an easy corollary of the Inscribed Angle Theorem that, given two triangles pqr and pqs in $\mathbb{R}^2$, if r is enclosed by the circumcircle of pqs, the following relations hold: $\angle prs > \angle pqs$, $\angle psr > \angle pqr$, $\angle qrs > \angle qps$, and $\angle qsr > \angle qpr$. Therefore, if pq can be flipped by the empty circumcircle criterion, then pq is also flippable according to Definition 3.[1] As a result, we can repeatedly flip edges that are flippable according to Definition 3 or the empty circumcircle criterion.

Assume that all angles in the triangulation are greater than α, a constant independent of n. Let $\mathcal{G}$ be the graph such that its nodes are the triangulation vertices and two nodes are connected if and only they are connected in some intermediate triangulation produced by the edge flips. In every intermediate triangulation, every vertex has degree at most $2\pi/\alpha$ because every incident angle is at least α. We apply the analysis of Lemma 4.2 to show that every vertex in $\mathcal{G}$ has a constant degree as follows. Let p be a vertex. Let pq and pr be two edges in $\mathcal{G}$. Let h be the distance between p and its nearest vertex. In the triangulation that contains pq, p is incident to an edge of length at most $h \csc \alpha$, so the longest edge in star(p) is at most $h(\csc \alpha)^{2\pi/\alpha}$ long. Thus, $d(p, q) \leq h(\csc \alpha)^{2\pi/\alpha}$.

[1] In $\mathbb{R}^2$, if pq is flippable, condition (iv) of Definition 3 forces $pqrs$ to be a convex quadrilateral and conditions (i, ii) of Definition 3 force r to be inside the circumcircle of pqs. That is, Definition 3 and the empty circumcircle criterion are equivalent in $\mathbb{R}^2$.

Similarly, the length of some edge in star(q) is at most $\csc\alpha$ times the nearest neighbor distance of q, and hence at most $d(q,r)\csc\alpha$. Therefore, the longest edge in star(q) is at most $d(q,r)(\csc\alpha)^{2\pi/\alpha}$ long, implying that $d(p,q) \leq d(q,r)(\csc\alpha)^{2\pi/\alpha}$.

It follows that $d(q,r) \geq d(p,q)(\sin\alpha)^{2\pi/\alpha} \geq h(\sin\alpha)^{2\pi/\alpha}$. Since q and r are two arbitrary nodes in $\mathcal{G}$ adjacent to p, we conclude that if we place a disk with radius $0.5h(\sin\alpha)^{2\pi/\alpha}$ and center at each node in $\mathcal{G}$ adjacent to p, such disks are disjoint. Also, these disks lie inside a bigger disk with center p and radius $h(\csc\alpha)^{2\pi/\alpha} + 0.5h(\sin\alpha)^{2\pi/\alpha}$. A packing argument shows that the degree of p in $\mathcal{G}$ is at most $(2(\csc\alpha)^{4\pi/\alpha}+1)^2 = O(1)$.

Then, the proof of Lemma 4.3 shows that the edge flips can be applied repeatedly until the triangulation becomes Delaunay in time linear in the number of vertices. □

5 Vertex neighborhood: flatness and injective projection

We develop two results in this section, Lemmas 5.1 and 5.4, that will be needed when we prove that edge flips can decrease the triangle circumradii in Section 6. Lemma 5.1 is an analog of Lemma 3.2(i) for triangles that are near a vertex p but possibly not in star(p). Lemma 5.4 shows that a neighborhood of p projects injectively onto every plane that makes an angle at least $\pi/3$ with $\mathbf{n}_p$. Throughout this section, $\mu = 2(\csc\alpha)^{4\pi/\alpha+1}$ as defined in Lemma 3.2.

Lemma 5.1 *For all $c > 0$, if T is a dense (ε,α)-mesh of Σ and $\varepsilon \leq \frac{1}{2c\mu+8\mu}$, then for every triangle $\tau \in T$ and every vertex $p \in T$ such that $d(p,\tau) \leq c\mu\varepsilon f(p)$, $\angle_a(\mathbf{n}_p,\mathbf{n}_\tau) < (2c+7.5)\mu\varepsilon$.*

Proof. Let q be the vertex of τ that subtends the largest angle in τ. By Lemma 3.2(i), $d(p,q) \leq c\mu\varepsilon f(p) + 2\gamma_\tau \leq c\mu\varepsilon f(p) + 2\mu\varepsilon f(q)$. Lemma 2.1(v) implies that $d(p,q) \leq \frac{(c+2)\mu\varepsilon}{1-2\mu\varepsilon} f(p)$. Then, by Lemma 2.1(i, ii(a)), $\angle_a(\mathbf{n}_p,\mathbf{n}_\tau) \leq \angle_a(\mathbf{n}_q,\mathbf{n}_\tau) + \angle(\mathbf{n}_p,\mathbf{n}_q) < 3.5\mu\varepsilon + \frac{(c+2)\mu\varepsilon}{1-(c+4)\mu\varepsilon}$, which is at most $(2c+7.5)\mu\varepsilon$ by our assumption that $\varepsilon \leq \frac{1}{2c\mu+8\mu}$. □

We need two technical lemmas in order to prove Lemma 5.4. The first one states that for every vertex p, any line that makes a small angle with $\mathbf{n}_p$ cannot be tangent to $|T|$ locally at any point near p.

Lemma 5.2 *For all $c > 0$, if T is a dense (ε,α)-mesh of Σ and $\varepsilon \leq \frac{1}{2c\mu+8\mu}$, then for every vertex $p \in T$ and every line L such that $\angle_a(\mathbf{n}_p, L) \leq \pi/2 - (4c+21)\mu\varepsilon$, L cannot be tangent to $|T|$ locally at any point in $B(p, c\mu\varepsilon f(p))$.*

Proof. Assume to the contrary that L is tangent to $|T|$ locally at a point $x \in B(p, c\mu\varepsilon f(p))$. There are three possibilities: x lies in the interior of a triangle $\tau \in T$, x lies in the interior of an edge $e \in T$, and x is a vertex of T.

In the first case, $\angle_a(L,\mathbf{n}_\tau) = \pi/2$, which implies that $\angle_a(\mathbf{n}_p,\mathbf{n}_\tau) \geq \angle_a(L,\mathbf{n}_\tau) - \angle_a(\mathbf{n}_p,L) \geq (4c+21)\mu\varepsilon$. This is a contradiction to Lemma 5.1.

In the second case, let σ and τ be the triangles incident to e. By Lemma 3.2(iv), the dihedral angle at e is greater than $\pi - 12\mu\varepsilon$. Lemma 5.1 implies that $\angle_a(\mathbf{n}_p, \tau\cap\sigma) \geq \angle_a(\mathbf{n}_p,\tau) > \pi/2 - (2c+7.5)\mu\varepsilon$. Therefore, $\max\{\angle_a(\mathbf{n}_p,\mathbf{n}_\sigma), \angle_a(\mathbf{n}_p,\mathbf{n}_\tau)\} > \pi/2 - 6\mu\varepsilon - (2c+7.5)\mu\varepsilon - \angle_a(\mathbf{n}_p,L) \geq (2c+7.5)\mu\varepsilon$. This is again a contradiction to Lemma 5.1.

In the third case, there exist triangles $\sigma,\tau \in$ star(x) such that σ and τ share an edge, and if we translate L to any point in $\sigma\cap\tau$, the translated line is tangent to $|T|$ locally at that point. We can then repeat the analysis in the previous paragraph to obtain a contradiction. □

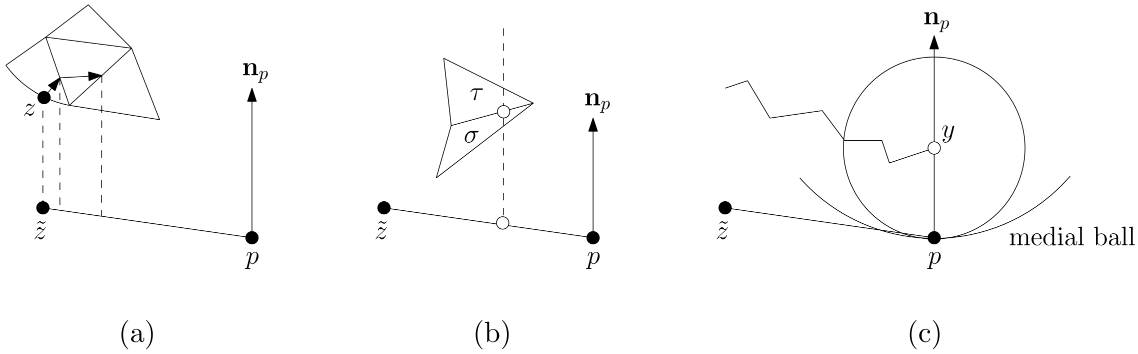

Figure 8: (a) Trace line segments across the triangles and the projection on $p\tilde{z}$ moves monotonically towards p. (b) We cannot turn back as we trace the line segments; otherwise, the normal of some triangle would make too large an angle with $\mathbf{n}_p$. (c) We obtain a polyline with an endpoint y that projects onto p, and if $y \neq p$, then $B(y, d(p, y))$ is contained in a medial ball at p.

The second technical result states that, within a neighborhood of a vertex, $|T|$ is connected and fairly flat.

Lemma 5.3 *For all $c \geq 1$, if T is a dense (ε, α)-mesh of Σ and $\varepsilon \leq \frac{1}{40c\mu}$, then*

(i) *for every vertex $p \in T$, $|T| \cap B(p, c\mu\varepsilon f(p))$ is connected,*

(ii) *for every vertex $p \in T$ and every point $x \in |T|$ such that $d(p, x) \leq c\mu\varepsilon f(p)$, $\angle_a(\mathbf{n}_p, px) > \pi/2 - (2c + 7.5)\mu\varepsilon$.*

Proof. Recall that ν_T denotes the restriction of the nearest point map ν to $|T|$. Let B_0 denote $B(p, 2c\mu\varepsilon f(p))$. Since $c\mu\varepsilon < 1/2$ by the assumption of the lemma, $\Sigma \cap B_0$ is a topological disk [5]. Therefore, $\nu_T^{-1}(\Sigma \cap B_0)$ is also a topological disk.

We first analyze the structure of $\nu_T^{-1}(\Sigma \cap B_0)$. Let x be a point in $\nu_T^{-1}(\Sigma \cap B_0)$. Let τ be a triangle in T that contains x. By Lemmas 3.2(i) and 2.2(iii), for every vertex w of τ, $d(x, \nu(x)) \leq 20\mu^2\varepsilon^2 f(w) \leq \frac{1}{2}c\mu\varepsilon f(w)$ as $c \geq 1$ and $c\mu\varepsilon \leq 1/40$ by the assumption of the lemma. Then, $d(p, w) \leq d(p, \nu(x)) + d(x, \nu(x)) + d(w, x) \leq 2c\mu\varepsilon f(p) + 2.5\mu\varepsilon f(w)$ as $d(w, x) \leq 2\gamma_\tau \leq 2\mu\varepsilon f(w)$ by Lemma 3.2(i). Then, $f(w) \leq \frac{1+2c\mu\varepsilon}{1-2.5\mu\varepsilon} f(p) < 2f(p)$ by Lemma 2.1(v) and the assumption of the lemma that $\mu\varepsilon \leq c\mu\varepsilon \leq 1/40$. Therefore, $d(x, \nu(x)) \leq \frac{1}{2}c\mu\varepsilon f(w) < c\mu\varepsilon f(p)$. This implies that the distance between p and the boundary of $\nu_T^{-1}(\Sigma \cap B_0)$ is greater than $(2c\mu\varepsilon - c\mu\varepsilon)f(p) \geq c\mu\varepsilon f(p)$ and, hence, the boundary of $\nu_T^{-1}(\Sigma \cap B_0)$ lies outside the ball $B_1 = B(p, c\mu\varepsilon f(p))$.

Let C be the connected component in $\nu_T^{-1}(\Sigma \cap B_0) \cap B_1$ that contains p. Since the boundary of $\nu_T^{-1}(\Sigma \cap B_0)$ is outside B_1, the boundary of C is a subset of ∂B_1. We first prove two claims.

Claim 5.1 *Let $\tilde{x}$ denote the projection of a point $x \in C$ onto the tangent plane of Σ at p. If z is a boundary point of C that minimizes $d(p, \tilde{z})$ over all boundary points of C, then $\angle_a(\mathbf{n}_p, pz) > \pi/2 - (2c + 7.5)\mu\varepsilon$.*

Proof. Starting from z, trace line segments across triangles in C that project onto subsets of the segment $p\tilde{z}$. Refer to Figure 8(a).

As we trace the line segments, the projection on $p\tilde{z}$ must move monotonically towards p. Otherwise, we turn back at a point $x \in |T| \cap B_1$, which means that a line parallel to $\mathbf{n}_p$ is tangent to $|T|$ locally at x. But this is a contradiction to Lemma 5.2.

As we trace the line segments, we cannot encounter any boundary point of C because, by our choice of z, no boundary point of C projects to a point in $p\tilde{z}$ closer to p than $\tilde{z}$.

Consequently, by tracing line segments across triangles in C that project onto subsets of the segment $p\tilde{z}$, we obtain a polyline in C from z to some point y such that p is the projection of y on $p\tilde{z}$ and this polyline is monotone with respect to $\mathrm{aff}(p\tilde{z})$. We claim that y is equal to p. If $y \neq p$, then Σ does not intersect the interior of $B(y, d(p, y))$ because $d(p, y) \leq c\mu\varepsilon f(p) < f(p)$ as $y \in B_1$, and therefore, $B(y, d(p, y))$ is strictly contained in a medial ball at p. Refer to Figure 8(c). The point $\nu_T(y)$ is equal to p because $B(y, d(p, y))$ intersects Σ at p only, but this contradicts the injectivity of ν_T. Hence, the endpoints of the polyline must be z and p.

If $\angle_a(\mathbf{n}_p, pz) \leq \pi/2 - (2c + 7.5)\mu\varepsilon$, then some line segment ℓ in this polyline must make an angle at most $\pi/2 - (2c + 7.5)\mu\varepsilon$ with $\mathbf{n}_p$, which implies that the angle between $\mathbf{n}_p$ and the normal of the triangle containing ℓ is at least $(2c+7.5)\mu\varepsilon$. But this contradicts Lemma 5.1. It follows that $\angle_a(\mathbf{n}_p, pz) > \pi/2-(2c+7.5)\mu\varepsilon$. □

Claim 5.2 *For every point $x \in C$, $\angle_a(\mathbf{n}_p, px) > \pi/2 - (2c + 7.5)\mu\varepsilon$.*

Proof. For every point $x \in C$, let $\tilde{x}$ denote the projection of x onto the tangent plane of Σ at p. Let z be a boundary point of C that minimizes $d(p, \tilde{z})$ over all boundary points of C. If $d(p, \tilde{x}) \geq d(p, \tilde{z})$, then B_1 constrains $d(x, \tilde{x})$ to be at most $d(z, \tilde{z})$, which implies that $\angle_a(\mathbf{n}_p, px) \geq \angle_a(\mathbf{n}_p, pz) > \pi/2-(2c+7.5)\mu\varepsilon$. If $d(p, \tilde{x}) < d(p, \tilde{z})$, then we can a trace a polyline from x to p as in the proof of Claim 5.1 to show that $\angle_a(\mathbf{n}_p, px) > \pi/2 - (2c + 7.5)\mu\varepsilon$. □

In the rest of the proof, we show that $|T| \cap B_1$ is equal to C, from which properties (i) and (ii) of the lemma follow. Assume to the contrary that $|T| \cap B_1 \neq C$. Since the boundary of $\nu_T^{-1}(\Sigma \cap B_0)$ is outside B_1, C is a connected component of $|T| \cap B_1$ and $|T| \cap B_1$ has connected component(s) other than C. Let q be a point in $(|T| \cap B_1) \setminus C$ that is closest to p. Let G be the plane tangent to $B(p, d(p, q))$ at q.

Suppose that q is a vertex of T. So $q \in \Sigma$. Because $\angle_a(\mathbf{n}_p, pq) > \pi/2 - 0.51c\mu\varepsilon$ by Lemma 2.1(i), $\angle_a(\mathbf{n}_p, G)$ is less than $0.51c\mu\varepsilon$. So there is a line $L \subset G$ through q such that $\angle_a(\mathbf{n}_p, L) = \angle_a(\mathbf{n}_p, G) < 0.51c\mu\varepsilon$ and L is tangent to $|T|$ locally at q. This is a contradiction to Lemma 5.2.

Suppose that q lies in the interior of an edge or a triangle in T. By our choice of G, the triangles in T that contain q lie on one side of G and p lies on the opposite side of G. By Lemma 5.1, for every triangle τ that contains q, $\angle_a(\mathbf{n}_p, \mathbf{n}_\tau) < (2c + 7.5)\mu\varepsilon$, which implies that

$$\angle_a(\mathbf{n}_p, G) > \pi/2 - (2c + 7.5)\mu\varepsilon \iff \angle_a(\mathbf{n}_p, pq) < (2c + 7.5)\mu\varepsilon. \quad (17)$$

Observe that $B(q, d(p, q)) \subset B_0$, which implies that

$$\nu_T(q) \in \Sigma \cap B_0.$$

Refer to Figure 9(a). Let D denote the closure of the complement of the double cone with apex p, half aperture $\pi/2 - (2c + 7.5)\mu\varepsilon$, and axis parallel to $\mathbf{n}_p$. Let L be the line through q and $\nu_T(q)$, which must be normal to Σ at $\nu_T(q)$.

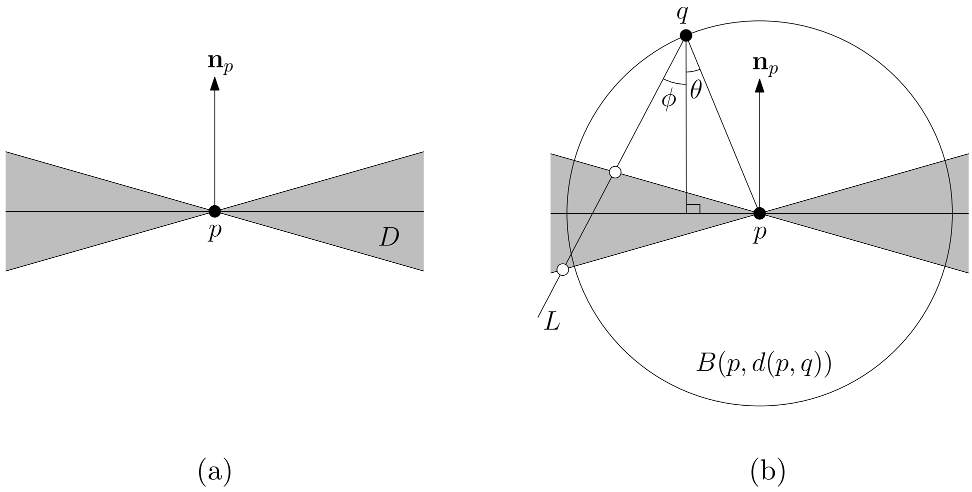

Figure 9: (a) Both $\Sigma \cap B_0$ and C are subsets of D (shown shaded). (b) The angle θ is less than $(2c + 7.5)\mu\varepsilon$ and the angle ϕ is less than $\frac{2c\mu\varepsilon}{1-2c\mu\varepsilon}$. By the sine law, the white dots are at distances at most $d(p,q)\,\frac{\sin(\phi+\theta)}{\cos(\phi+(2c+7.5)\mu\varepsilon)}$ from p.

Claim 5.3 *L intersects C.*

Proof. By Lemma 2.1(i), $\Sigma \cap B_0 \subseteq D$ and for every point $x \in \Sigma \cap B_0$, $\angle(\mathbf{n}_p, \mathbf{n}_x) \leq \frac{2c\mu\varepsilon}{1-2c\mu\varepsilon}$. Thus, $\angle_a(\mathbf{n}_p, L) \leq \frac{2c\mu\varepsilon}{1-2c\mu\varepsilon}$. Refer to Figure 9(b). Let ℓ be the maximum distance from p to the two intersection points between L and the boundary of D. By the sine law, $\ell = d(p,q)\,\frac{\sin(\phi+\theta)}{\cos(\phi+(2c+7.5)\mu\varepsilon)} \leq d(p,q)\,\frac{\phi+\theta}{1-\phi-(2c+7.5)\mu\varepsilon}$. Recall that $\theta \leq (2c+7.5)\mu\varepsilon$ and $\phi < \frac{2c\mu\varepsilon}{1-2c\mu\varepsilon}$. Since $c \geq 1$ and $\varepsilon \leq \frac{1}{40c\mu}$ by the assumption of the lemma, we obtain $\ell \leq d(p,q)\,\frac{(5c+7.5)\mu\varepsilon}{1-(5c+7.5)\mu\varepsilon} \leq d(p,q)\,\frac{13c\mu\varepsilon}{1-13c\mu\varepsilon} \leq d(p,q)$. Thus, L pierces through D completely inside B_1.

We have shown in Claim 5.2 that for every point $x \in C$, $\angle_a(\mathbf{n}_p, px) > \pi/2 - (2c+7.5)\mu\varepsilon$. Therefore, $C \subseteq D \cap B_1$. Recall that the boundary of C is a subset of ∂B_1. If L does not intersect C, we can slide a copy L' of L towards p and L' must eventually be tangent to $|C| \subseteq |T| \cap B_1$ locally at some point. But this is a contradiction to Lemma 5.2. $\square$

Let r denote a point in $L \cap C$. Notice that $q \neq r$ because $\angle_a(\mathbf{n}_p, pq) < (2c+7.5)\mu\varepsilon$ by (17) and, therefore, $q \notin D$. Since $\nu_T(q)$ belongs to B_0, by Lemma 2.1(v), $f(\nu_T(q)) \geq (1-2c\mu\varepsilon)f(p)$, which is greater than $3c\mu\varepsilon f(p) \geq d(p,r) + d(p,\nu_T(q)) \geq d(r,\nu_T(q))$. Therefore, $B(r, d(r,\nu_T(q))$ is contained in a medial ball at $\nu_T(q)$, which implies $\nu_T(q)$ is the closest point in Σ to r. But then $\nu_T(r) = \nu_T(q)$, contradicting the injectivity of ν_T.

In summary, we obtain a contradiction if we assume that $(|T| \cap B_1) \setminus C \neq \emptyset$. Therefore, $|T| \cap B_1 = C$, implying the correctness of properties (i) and (ii) of the lemma. $\square$

We are ready to show that a neighborhood of a vertex p projects injectively onto every plane that makes an angle at least $\pi/3$ with $\mathbf{n}_p$.

Lemma 5.4 *For all $c \geq 1$, if T is a dense (ε, α)-mesh of Σ and $\varepsilon \leq \frac{1}{40c\mu}$, then for every vertex $p \in T$, $|T| \cap B(p, c\mu\varepsilon f(p))$ is connected and it projects injectively onto any plane that makes an angle at least $\pi/3$ with $\mathbf{n}_p$.*

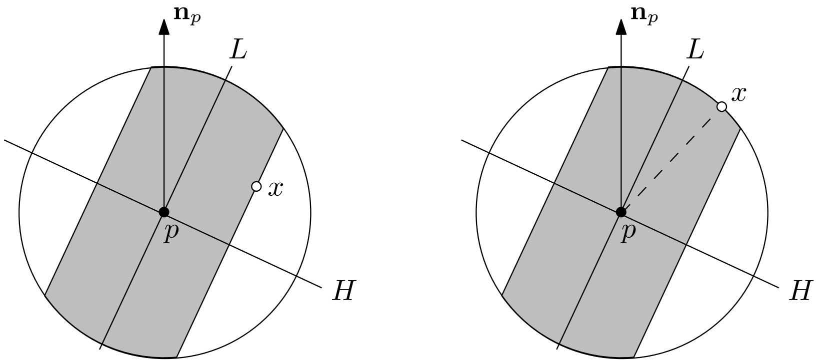

(a) (b)

Figure 10: The ball denotes $B(p, 1.5c\mu\varepsilon f(p))$. The shaded region consists of points in $B(p, 1.5c\mu\varepsilon f(p))$ at distance $c\mu\varepsilon f(p)$ or less from L.

Proof. By Lemma 5.3(i), $|T| \cap B(p, c\mu\varepsilon f(p))$ is connected. Let H be any plane such that $\angle_a(\mathbf{n}_p, H) \geq \pi/3$. Consider the projection of $|T| \cap B(p, c\mu\varepsilon f(p))$ onto H. Let L be the line through p perpendicular to H. So $\angle_a(\mathbf{n}_p, L) \leq \pi/6$.

Suppose for the sake of contradiction that the projection of $|T| \cap B(p, c\mu\varepsilon f(p))$ onto H is not injective. It follows that the projection of $|T| \cap B(p, 1.5c\mu\varepsilon f(p))$ onto H is not injective either. So there exist two points $x, y \in |T| \cap B(p, 1.5c\mu\varepsilon f(p))$ such that xy is parallel to L. We choose x and y so that $d(x, L) = d(y, L)$ is minimized. This choice of x and y means that their distances from L are at most $c\mu\varepsilon f(p)$.

Suppose that neither x nor y belongs to the boundary of $B(p, 1.5c\mu\varepsilon f(p))$. By Lemma 5.3(ii), L intersects $|T| \cap B(p, c\mu\varepsilon f(p))$ at p only. Refer to Figure 10(a). If we translate a copy L' of L linearly towards xy, then our choice of x and y implies that when L' reaches xy, L' is tangent to $|T|$ locally at x or y. This is a contradiction to Lemma 5.2.

Suppose that x or y belongs to the boundary of $B(p, 1.5c\mu\varepsilon f(p))$, say x. See Figure 10(b). Since $\angle_a(px, L) \leq \arcsin(1/1.5)$ and $\angle_a(\mathbf{n}_p, L) = \pi/2 - \angle(\mathbf{n}_p, H) \leq \pi/6$, we obtain $\angle_a(\mathbf{n}_p, px) \leq \angle_a(\mathbf{n}_p, L) + \angle_a(px, L) \leq \pi/6 + \arcsin(1/1.5) < 1.26$. This is impossible because Lemma 5.3(ii) implies that $\angle_a(\mathbf{n}_p, px) > \pi/2 - (2c + 7.5)\mu\varepsilon > 1.26$ as $2c + 7.5 \leq 9.5c$ and $\varepsilon \leq \frac{1}{40\mu\varepsilon}$ by the assumption of the lemma. □

6 Almost empty diametric balls

In this section, we prove that if a dense (ε, α)-mesh does not have any flippable edge, then the diametric ball of every mesh triangle is almost empty in the sense that a concentric ball with a slightly smaller radius is empty of mesh vertices. The "almost emptiness" of the diametric ball implies that the circumradius of a triangle τ is bounded by $\varepsilon + O(\varepsilon^\kappa)$ times the local feature size at any vertex of τ, where κ is any fixed constant in $(1, 1.5)$. The factor $\varepsilon + O(\varepsilon^\kappa)$ is much smaller than the factor $\mu\varepsilon = 2(\csc\alpha)^{4\pi/\alpha+1}\varepsilon$ in Lemma 3.2(i, iv) and, therefore, edge flips make the mesh smoother and a better approximation of Σ.

We first show that the edge flippability is related to a local notion of emptiness of diametric balls of triangles. Let pqr and pqs be two triangles in a surface mesh of Σ. Recall that the diametric ball of pqr is denoted by $B_{pqr} = B(c_{pqr}, \gamma_{pqr})$, where c_{pqr} denotes the circumcenter of pqr. We say that a point x *lies inside* B_{pqr} if x lies in the interior of B_{pqr}. In the case that

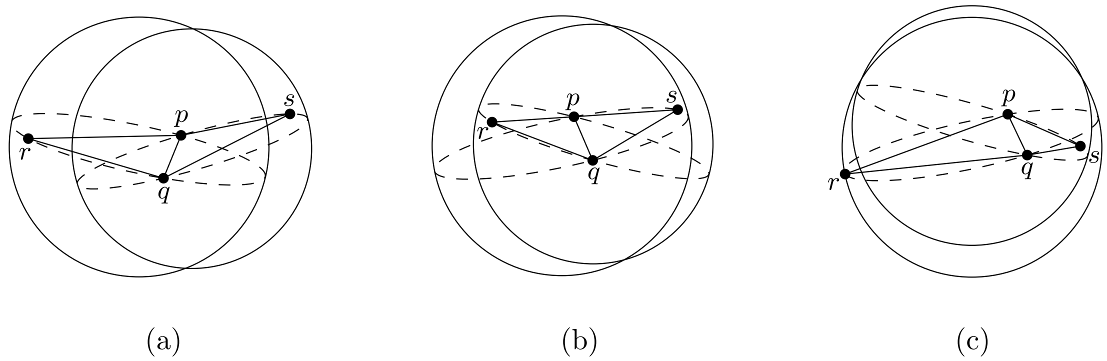

Figure 11: (a) r does not lie inside B_{pqs} and s does not lie inside B_{pqr}. (b) r lies inside B_{pqs} and s lies inside B_{pqr}. (c) s lies inside B_{pqr} and r does not lie inside B_{pqs}.

pqr and pqs are coplanar, r lies inside B_{pqs} if and only if s lies inside B_{pqr}. However, this is not true when the two triangles are not coplanar. Figure 11 shows all three possible configurations.

Definition 4 *Let pqr and pqs be two triangles in a triangulated polygonal surface. The edge pq is **illegal** if r lies inside B_{pqs} and s lies inside B_{pqr}. Otherwise, pq is **legal**.*

The next lemma shows that every illegal edge is flippable.

Lemma 6.1 *If T is a dense (ε, α)-mesh of Σ, then every illegal edge in T is flippable according to Definition 3.*

Proof. Let pqr and pqs be two triangles in T such that the common edge pq is illegal. We first prove that $\angle prs > \min\{\pi/2, \angle pqs\}$. Refer to Figure 12(a). Since r lies inside B_{pqs}, the ray from p through r hits a point $x \in \partial B_{pqs}$, which implies that $\angle prs > \angle pxs$. The circumcircle of pqs, being a great circle of B_{pqs}, is not smaller than the circumcircle of pxs. Therefore, $\frac{d(p,s)}{2\sin\angle pqs} = \gamma_{pqs} \geq \gamma_{pxs} = \frac{d(p,s)}{2\sin\angle pxs}$, which implies that $\sin\angle pxs \geq \sin\angle pqs$. If $\angle pqs \leq \pi/2$, then $\sin\angle pxs \geq \sin\angle pqs \Rightarrow \angle pxs \geq \angle pqs$, and therefore, $\angle prs > \angle pxs \geq \angle pqs$. If $\angle pqs > \pi/2$, we show that $\angle prs > \pi/2$ as follows. The plane H through ps perpendicular to pqs cuts ∂B_{pqs} into two subsets. Let C denote the subset that contains q. Since $\angle pqs > \pi/2$, C is the smaller of the two subsets as illustrated in Figure 12. Since pqx and pqr are coplanar, the dihedral angle between pqx and pqs is the same as that between pqr and pqs, which is obtuse by Lemma 3.2(iv). This forces x to lie in C. Imagine that we fix the locations of p and s and move x on C. When x lies on H, $\angle pxs$ achieves its minimum value of $\pi/2$ (because ps is a diameter of the disk $H \cap B_{pqs}$). Thus, $\angle pxs \geq \pi/2$ in general and $\angle prs > \angle pxs \geq \pi/2$.

Similarly, we can prove that $\angle psr > \min\{\pi/2, \angle pqr\}$, $\angle qrs > \min\{\pi/2, \angle qps\}$, and $\angle qsr > \min\{\pi/2, \angle qpr\}$.

Next, we show that $\angle rps > \max\{\angle rpq, \angle spq\}$. Put a very small sphere S centered at p, which intersects pq, pr and ps at the points q', r' and s', respectively. Refer to Figure 12(b). The points q', r' and s' lie in a quarter of S bounded by the tangent plane of B_{pqr} at p and $\text{aff}(pqr)$. The plane orthogonal to both $\text{aff}(pr)$ and $\text{aff}(pqr)$ bounds a halfspace H^+ that contains r'. $H^+ \cap S$ is the union of great circular arcs of S that are incident to r' and are as long as the great circular arc between q' and r'. Also, the bounding plane of H^+ meets the great circular arc between q' and r' at right angle. Since the dihedral angle at pq is obtuse by Lemma 3.2(iv), s' does not belong to H^+ which implies that the great circular arc between r' and s' is longer than the one between q' and r'. Hence, $\angle rps > \angle rpq$. A similar argument shows that $\angle rps > \angle spq$.

Similarly, we can show that $\angle rqs > \max\{\angle rqp, \angle sqp\}$.

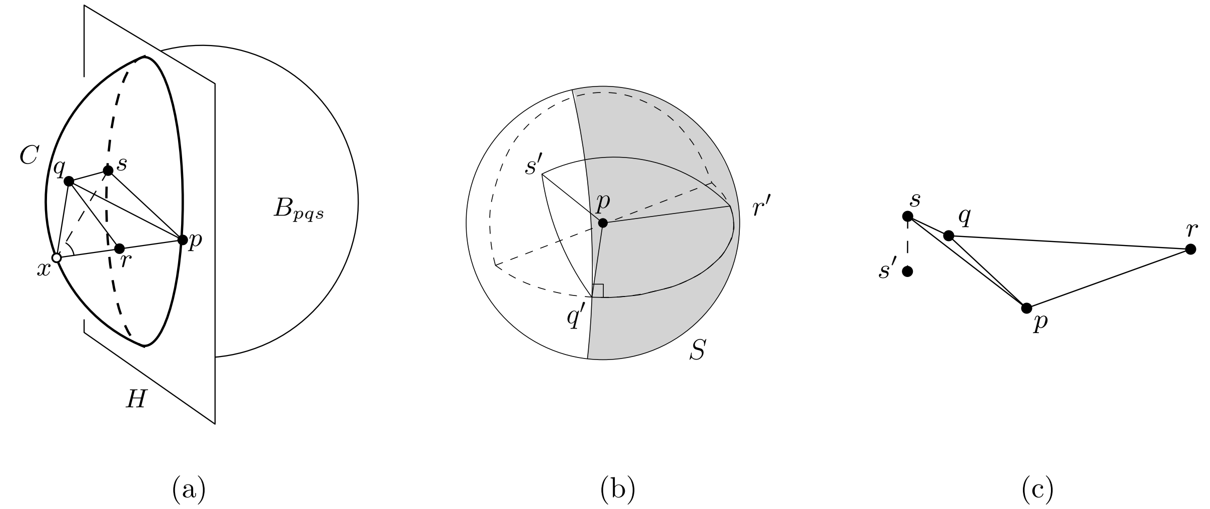

Figure 12: (a) The ray from p through r hits ∂B_{pqs} at x. The plane H through ps perpendicular to pqs cuts ∂B_{pqs} into two subsets, where C is the one containing q. If C is smaller than a hemisphere, then $\angle pxs \geq \pi/2$. (b) A quarter of S contains the points q', r' and s'. The plane orthogonal to both aff(pr) and aff(pqr) bounds a halfspace H^+ containing r'. $H^+ \cap S$ is shown shaded. (c) s' is the projection of s onto aff(pqr).

Next, we show that $\angle prq > \max\{\angle prs, \angle qrs\}$. Since s lies inside B_{pqr}, by Lemma 3.2(i), $d(r,s) < 2\gamma_{pqr} \leq 2\mu\varepsilon f(r)$. Then, Lemma 2.1(i) and Lemma 3.2(i) imply that $\angle_a(pqr, rs) \leq \angle_a(\mathbf{n}_r, \mathbf{n}_{pqr}) + \pi/2 - \angle_a(\mathbf{n}_r, rs) < 8\mu\varepsilon$. Since s lies inside B_{pqr} and the dihedral angle at pq is obtuse by Lemma 3.2(iv), the projection of rs onto aff(pqr) intersects pqr. Refer to Figure 12(c). Let s' be the orthogonal projection of s onto aff(pqr). It follows that $\angle prs' \geq \angle prs - \angle_a(pqr, rs) > \angle prs - 8\mu\varepsilon$. Similarly, $\angle qrs' > \angle qrs - 8\mu\varepsilon$. Therefore, $\angle prq = \angle prs' + \angle qrs' > \angle prs + \angle qrs - 16\mu\varepsilon$. We have shown earlier that $\angle prs > \min\{\pi/2, \angle pqs\}$ and $\angle qrs > \min\{\pi/2, \angle qps\}$. Moreover, $\angle pqs$ and $\angle qps$ are at least α. Therefore, $\angle prq > \max\{\angle prs, \angle qsr\} + \alpha - 16\mu\varepsilon > \max\{\angle prs, \angle qrs\}$ because $\alpha > 16\mu\varepsilon$ by the condition on ε in the definition of (ε, α)-meshes.

Similarly, we can show that $\angle psq > \max\{\angle psr, \angle qsr\}$.

It remains to prove that rs is not an edge in T. Suppose to the contrary that rs is an edge in T. We have argued before that the projection of rs onto aff(pqr) intersects pqr. By Lemma 3.2(i), $\angle_a(\mathbf{n}_r, pqr) > \pi/2 - 6\mu\varepsilon$, which is greater than $\pi/3$ by the condition on ε in the definition of dense (ε, α)-meshes. But then the projection of star(r) onto aff(pqr) should be injective according to Lemma 5.4, a contradiction. □

By Lemmas 4.1, 3.2(iv) and 6.1, there is no illegal edge when there is no flippable edge in a dense (ε, α)-mesh.

6.1 Technical results

We first prove two technical results before showing in the next subsection that the diametric balls of the triangles are almost empty when no edge is flippable.

Given two triangles σ and τ in an (ε, α)-mesh that share an edge, the boundaries of their diametric balls intersect in a circle. Let H be the support plane of this circle. Although the dihedral angle between σ and τ is close to π, the plane H may still be nearly parallel to σ or

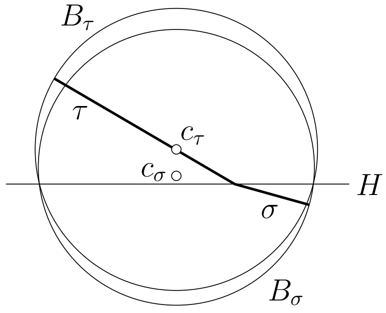

Figure 13: The figure shows a cross-section view of the triangles σ and τ. The plane H can make a small angle with σ or τ even if the dihedral angle between σ and τ is close to π.

τ.[2] Figure 13 shows an example.

The first technical result states that if τ makes a small angle with H, then γ_σ and γ_τ are similar and c_σ and c_τ are close to each other. Recall that $\mu = 2(\csc\alpha)^{4\pi/\alpha+1}$ as defined in Lemma 3.2.

Lemma 6.2 *For all $c > 0$, if T is a dense (ε, α)-mesh of Σ and $\varepsilon \le \frac{1}{2c\mu+24\mu}$, then for every pair of triangles $\sigma, \tau \in T$ such that σ and τ share an edge and $\angle_a(\tau, H) \le c\mu\varepsilon$, where $H = \mathrm{aff}(\partial B_\tau \cap \partial B_\sigma)$, the following relations hold.*

(i) $(1 - 2(6+c)\mu\varepsilon)\gamma_\sigma < \gamma_\tau < (1 + 4(6+c)\mu\varepsilon)\gamma_\sigma$.

(ii) $d(c_\tau, c_\sigma) < 4(6+c)\mu\varepsilon\gamma_\tau$.

Proof. Refer to Figure 13. Since H contains the edge $\sigma \cap \tau$, the assumption of $\angle_a(\tau, H) \le c\mu\varepsilon$ implies that $d(c_\tau, H) \le \gamma_\tau \sin(c\mu\varepsilon)$. By Lemma 3.2(iv), $\angle_a(\tau, \sigma) < 12\mu\varepsilon$. Thus, $\angle_a(\sigma, H) \le \angle_a(\sigma, \tau) + \angle_a(\tau, H) < (12+c)\mu\varepsilon$, which implies that $d(c_\sigma, H) < \gamma_\sigma \sin((12+c)\mu\varepsilon)$. The segment $c_\sigma c_\tau$ is perpendicular to H. Therefore,

$$\begin{aligned} d(c_\sigma, c_\tau) &\le d(c_\sigma, H) + d(c_\tau, H) \\ &< \gamma_\sigma \sin((12+c)\mu\varepsilon) + \gamma_\tau \sin(c\mu\varepsilon) \\ &< (12+c)\mu\varepsilon\gamma_\sigma + c\mu\varepsilon\gamma_\tau. \end{aligned}$$

By the triangle inequality, $\gamma_\tau \le \gamma_\sigma + d(c_\sigma, c_\tau)$. Then, $\gamma_\tau < \gamma_\sigma + c\mu\varepsilon\gamma_\tau + (12+c)\mu\varepsilon\gamma_\sigma$, which gives

$$\gamma_\tau < \frac{1 + (12+c)\mu\varepsilon}{1 - c\mu\varepsilon}\gamma_\sigma = \left(1 + \frac{(12+2c)\mu\varepsilon}{1 - c\mu\varepsilon}\right)\gamma_\sigma < (1 + 4(6+c)\mu\varepsilon)\gamma_\sigma.$$

The last step substitutes $1 - c\mu\varepsilon$ by $1/2$ which follows from the assumption of the lemma that $\varepsilon \le \frac{1}{2c\mu+24\mu}$. Symmetrically, one can derive from the inequality $\gamma_\sigma - d(c_\sigma, c_\tau) \le \gamma_\tau$ that

$$\gamma_\tau > \frac{1 - (12+c)\mu\varepsilon}{1 + c\mu\varepsilon}\gamma_\sigma = \left(1 - \frac{(12+2c)\mu\varepsilon}{1 + c\mu\varepsilon}\right)\gamma_\sigma > (1 - 2(6+c)\mu\varepsilon)\gamma_\sigma.$$

This proves (i). Then,

$$\begin{aligned} d(c_\sigma, c_\tau) &< (12+c)\mu\varepsilon\gamma_\sigma + c\mu\varepsilon\gamma_\tau \\ &\le (12+c)\mu\varepsilon \cdot (\gamma_\tau + d(c_\sigma, c_\tau)) + c\mu\varepsilon\gamma_\tau \\ &= (12+2c)\mu\varepsilon\gamma_\tau + (12+c)\mu\varepsilon\, d(c_\sigma, c_\tau). \end{aligned}$$

[2] If σ and τ are coplanar, then H must be perpendicular to the support plane of σ and τ.

Rearranging terms, we obtain

$$d(c_\sigma, c_\tau) < \frac{(12+2c)\mu\varepsilon\gamma_\tau}{1-(12+c)\mu\varepsilon} \le 4(6+c)\mu\varepsilon\gamma_\tau.$$

In the last step, we substitute $1-(12+c)\mu\varepsilon$ by $1/2$, which follow from the assumption of the lemma that $\varepsilon \le \frac{1}{2c\mu+24\mu}$. This proves (ii). ◻

The second technical result states that, in a dense (ε,α)-mesh, if a vertex v lies inside the diametric ball of a triangle τ, then for every edge pq of τ, the triangle pqv makes a small angle with τ.

Lemma 6.3 *Let T be a dense (ε,α)-mesh of Σ. For every triangle $\tau \in T$, every edge pq of τ and every vertex $v \in T$, if v lies inside B_τ, then $\angle_a(pqv,\tau) < 10\mu\varepsilon\csc(\alpha/2)$.*

Proof. Let τ be a triangle in T. Let pq be an edge of τ. Let v be a vertex of T that lies inside B_τ. Let $\tilde{v}$ be the orthogonal projection of v onto $\mathrm{aff}(pq)$.

By Lemma 3.2(i), $\angle_a(\mathbf{n}_p,\mathbf{n}_\tau) < 6\mu\varepsilon$. It can be verified that $\varepsilon \le 1/(80\mu)$ by the condition on ε in the definition of dense (ε,α)-meshes. Therefore, we can invoke Lemma 5.4 with $c=2$ to project $T \cap B(p, 2\mu\varepsilon f(p))$ injectively onto $\mathrm{aff}(\tau)$. Let v' denote the projection of v. In the projection, either v' is connected to p or v' lies outside the projection of $\mathrm{star}(p)$. Also, v' lies inside the circumcircle of τ because v lies inside B_τ. Since $d(p,v) \le 2\gamma_\tau \le 2\mu\varepsilon f(p)$ by Lemma 3.2(i), we obtain $\angle(\mathbf{n}_p, pv) \ge \pi/2 - 1.02\mu\varepsilon$ by Lemma 2.1(i). Therefore,

$$\angle_a(\tau, pv) \le \angle_a(\mathbf{n}_p,\mathbf{n}_\tau) + \pi/2 - \angle_a(\mathbf{n}_p, pv) < 8\mu\varepsilon. \tag{18}$$

Since $d(q,v) \le 2\gamma_\tau \le 2\mu\varepsilon f(q)$ by Lemma 3.2(i), we can similarly show that

$$\angle_a(\tau, qv) < 8\mu\varepsilon. \tag{19}$$

Therefore,

$$d(v,v') < \min\{d(p,v),\, d(q,v)\} \cdot \sin(8\mu\varepsilon). \tag{20}$$

There are two cases to consider depending on whether $\tilde{v}$ lies in the interior of pq or not.

Suppose that $\tilde{v}$ lies in the interior of pq. Refer to Figure 14(a). Let pqr be the triangle that shares pq with τ. Let r' be the projection of r onto $\mathrm{aff}(\tau)$. Similar to the derivation of (18) and (19), we can show that $\angle_a(\tau, pr) < 8\mu\varepsilon$ and $\angle_a(\tau, qr) < 8\mu\varepsilon$. It follows that $\angle pqr'$ and $\angle qpr'$ are at least $\alpha - 8\mu\varepsilon \ge \alpha/2$ by the condition on ε in the definition of (ε,α)-meshes. Since v' lies outside the interior of τ and pqr', we conclude that $\angle qpv' \ge \alpha/2$ or $\angle pqv' \ge \alpha/2$. Assume without loss of generality that $\angle qpv' \ge \alpha/2$. So

$$\begin{aligned} d(v',\tilde{v}) &\ge d(p,v')\sin(\alpha/2) \ge \big(d(p,v) - d(v,v')\big)\sin(\alpha/2) \\ &> \frac{1-\sin(8\mu\varepsilon)}{\sin(8\mu\varepsilon)}\, d(v,v')\sin(\alpha/2). \qquad (\because (20)) \end{aligned}$$

Therefore,

$$\angle_a(pqv,\tau) = \arctan\left(\frac{d(v,v')}{d(v',\tilde{v})}\right) < \frac{\sin(8\mu\varepsilon)}{1-\sin(8\mu\varepsilon)}\csc(\alpha/2) < 10\mu\varepsilon\csc(\alpha/2).$$

The inequality $\sin(8\mu\varepsilon)/(1-\sin(8\mu\varepsilon)) < 10\mu\varepsilon$ is used in the last step above, which follows from the condition on ε in the definition of dense (ε,α)-meshes.

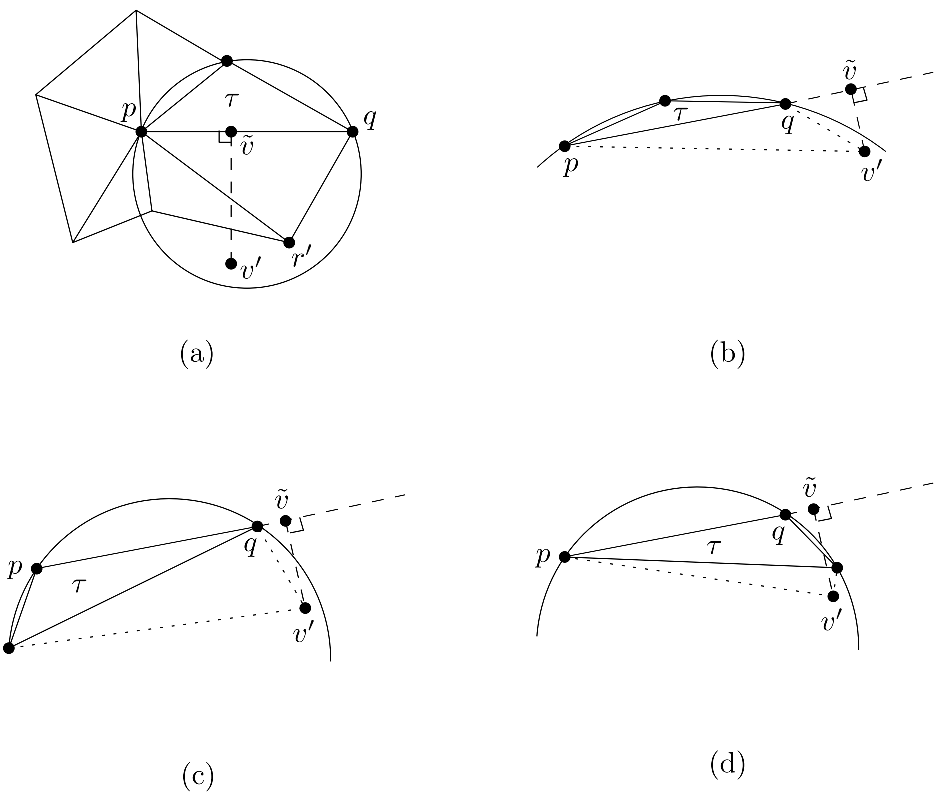

Figure 14: Illustrations for the proof of Lemma 6.3.

Suppose that $\tilde{v}$ does not lie in the interior of pq. Then, either p or q is the nearest point in τ to v, say q. It follows that $\angle pqv' \geq \pi/2$. Figure 14(b,c,d) show the three possible configurations. We claim that $\angle pqv' \in [\pi/2, \pi - \alpha]$ in all three configurations. Consider Figure 14(b, c). Let C be the convex hull of τ and v'. Since v' lies inside the circumcircle of τ, any two opposite angles of C sum to less than π. Then, the angle of C at q is less than $\pi - \alpha$ because the angle of C opposite q is at least α. It follows that $\angle pqv'$ is between $\pi/2$ and $\pi - \alpha$. In Figure 14(d), $\angle pqv'$ is at most the angle of τ at q which is at most $\pi - 2\pi$. This establishes our claim. Since $\angle pqv' \in [\pi/2, \pi - \alpha]$, we get

$$\begin{aligned} d(v', \tilde{v}) &= d(q, v') \sin \angle pqv' \;\geq\; \big(d(q, v) - d(v, v')\big) \sin \alpha \\ &> \frac{1 - \sin(8\mu\varepsilon)}{\sin(8\mu\varepsilon)}\, d(v, v') \sin \alpha. \qquad (\because (20)) \end{aligned}$$

Hence,

$$\angle_a(pqv, \tau) = \arctan\left(\frac{d(v, v')}{d(v', \tilde{v})}\right) < \frac{\sin(8\mu\varepsilon)}{1 - \sin(8\mu\varepsilon)} \csc \alpha < 10\mu\varepsilon \csc \alpha.$$

$\square$

6.2 Main analysis

We first prove in Lemma 6.4 below that if two triangles pqr and pqs share a legal edge pq, then r is either outside B_{pqs} or near ∂B_{pqs}. (The same conclusion can be drawn for s.) This is the base case of the subsequent induction to show that every vertex is either outside B_{pqs} or near ∂B_{pqs} when all edges are legal.

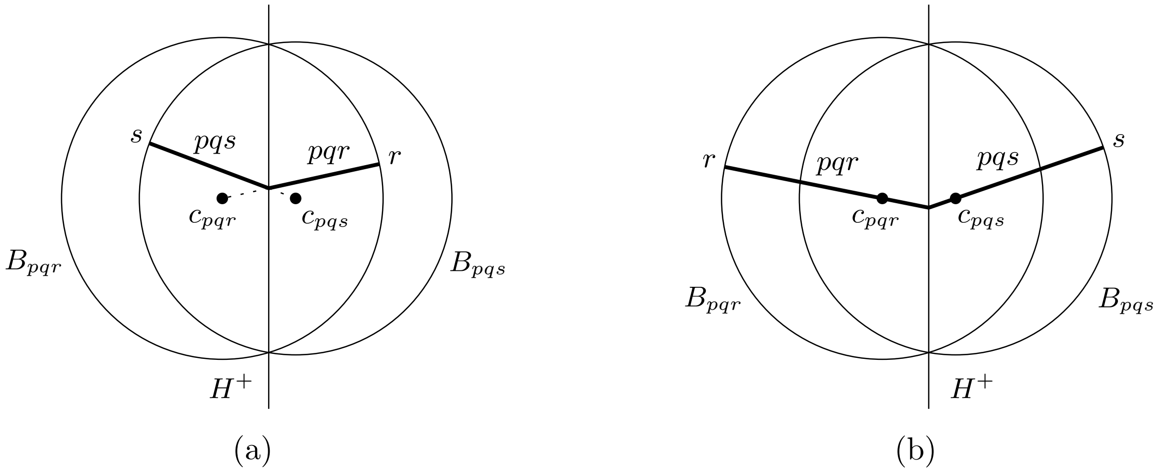

Figure 15: Two side-views taken in the direction of pq. (a) $pqr \subset B_{pqs}$ and $pqs \subset B_{pqr}$, implying that pq is illegal. (b) The edge pq is legal.

Lemma 6.4 *Let T be a dense (ε, α)-mesh. For every pair of triangles $pqr, pqs \in T$ that share a legal edge pq, $d(r, c_{pqs}) > (1 - 108\mu\varepsilon)\gamma_{pqs}$.*

Proof. Let $H = \mathrm{aff}(\partial B_{pqr} \cap \partial B_{pqs})$. Suppose that $\angle_a(pqr, H) \leq 12\mu\varepsilon$. Since T is a dense (ε, α)-mesh, it can be verified that $\varepsilon \leq 1/(48\mu)$ by the condition on ε in the definition of dense (ε, α)-meshes. Therefore, we can apply Lemma 6.2 with $c = 12$ to obtain $d(c_{pqr}, c_{pqs}) < 72\mu\varepsilon\gamma_{pqr}$ and $\gamma_{pqr} > (1 - 36\mu\varepsilon)\gamma_{pqs}$. So $d(r, c_{pqs}) \geq d(r, c_{pqr}) - d(c_{pqr}, c_{pqs}) > \gamma_{pqr} - 72\mu\varepsilon\gamma_{pqr} > (1 - 72\mu\varepsilon)(1 - 36\mu\varepsilon)\gamma_{pqs} > (1 - 108\mu\varepsilon)\gamma_{pqs}$ as stated in the lemma. Suppose that $\angle_a(pqr, H) > 12\mu\varepsilon$. By Lemma 3.2(iv), the dihedral angle at pq is greater than $\pi - 12\mu\varepsilon$. Thus, H separates pqr and pqs. Let H^+ be the half-space bounded by H that contains pqs. If $B_{pqs} \cap H^+ \subset B_{pqr}$, then $B_{pqr} \setminus H^+ \subset B_{pqs}$. Refer to Figure 15(a). Therefore, r lies inside B_{pqs} and s lies inside B_{pqr}, but then pq is illegal, a contradiction. So we are in the case of Figure 15(b), where $B_{pqs} \cap H^+ \not\subset B_{pqr}$. Then, r does not lie inside B_{pqs}, which implies that $d(r, c_{pqs}) \geq \gamma_{pqs}$. ◻

By Lemma 6.4, given a triangle τ with no illegal edges, we only need to worry about whether B_τ contains a vertex p several triangles away. We introduce some notation to facilitate the analysis of such a case.

Definition 5 *Let T be a dense (ε, α)-mesh of Σ. Let p be a vertex in T that lies inside the diametric ball of a triangle $\tau \in T$. We define a* ***triangle sequence* $\mathbf{seq}(p, \tau)$** *from τ to p as follows. By Lemma 5.4, $T \cap B(c_\tau, \gamma_\tau)$ projects injectively to* $\mathrm{aff}(\tau)$*. Let p' denote the projection of p. There is a unique vertex v of τ such that the line segment $p'v$ intersects the interior of τ. Walking linearly from v to p' visits the projections of some triangles on* $\mathrm{aff}(\tau)$ *in order. The sequence* $\mathrm{seq}(p, \tau)$ *is the corresponding ordered list of these triangles. That is, if $(\tau_1', \tau_2', \tau_3', \ldots)$ is the sequence of projections of triangles visited as we walk linearly from v to p', where τ_i' denotes the projection of τ_i, then* $\mathrm{seq}(p, \tau) = (\tau_1, \tau_2, \tau_3, \ldots)$*. Note that $\tau_1 = \tau_1' = \tau$. The last triangle in* $\mathrm{seq}(p, \tau)$ *belongs to* $\mathrm{star}(p)$*.*

If p lies inside B_τ, our plan to show that p is near ∂B_τ goes as follows. Let $\lambda = 1 - \varepsilon^c$ for some fixed constant $c \in (0, 0.5)$. If $d(p, c_\tau) \geq \lambda\gamma_\tau$, then p is already near ∂B_τ. Suppose to the contrary that $d(p, c_\tau) < \lambda\gamma_\tau$. We extract the shortest prefix $(\tau_1, \tau_2, \ldots, \tau_k)$ of $\mathrm{seq}(p, \tau)$ such that $d(p, c_{\tau_k}) \geq (1 - \sqrt{\varepsilon})\gamma_{\tau_k}$. Therefore, $d(p, c_{\tau_i}) < (1 - \sqrt{\varepsilon})\gamma_{\tau_i}$ for $i \in [1, k-1]$. We prove two results below. Lemma 6.5 shows that for $i \in [1,k]$, $d(p,c_{\tau_i})/\gamma_{\tau_i} \geq 1-\rho_i$, where ρ_i increases slowly as i decreases from k to 1. In other words, the lower bound on $d(p,c_{\tau_i})/\gamma_{\tau_i}$ decreases slowly as i decreases from k to 1. Lemma 6.6 shows that k is at most a constant. Since $d(p,c_{\tau_k}) \geq (1-\sqrt{\varepsilon})\gamma_{\tau_k}$, we can work backward from $i=k-1$ to 1 and apply Lemmas 6.5 and 6.6 to obtain a lower bound on $d(p,c_{\tau_1}) = d(p,c_\tau)$. This lower bound will contradict the assumption of $d(p,c_\tau) < \lambda\gamma_\tau$. Hence, $d(p,c_\tau) \geq \lambda\gamma_\tau$ is the only possibility, i.e., p is near ∂B_τ.

Lemma 6.5 *Let T be a dense (ε,α)-mesh of Σ for a sufficiently small ε. Suppose that a vertex $p \in T$ lies inside the diametric ball of a triangle $\tau \in T$. If $(\tau_1,\tau_2,\ldots,\tau_k)$ is a prefix of* $\mathrm{seq}(p,\tau)$ *such that $d(p,c_{\tau_k}) \geq (1-\sqrt{\varepsilon})\gamma_{\tau_k}$ and $\tau_i \cap \tau_{i+1}$ is legal for $i \in [1,k-1]$, then for $i \in [1,k]$, $d(p,c_{\tau_i}) \geq (1-\rho_i)\gamma_{\tau_i}$, where $\rho_i = 2^{k-i}(\csc\alpha)^{2k-2i}\sqrt{\varepsilon}$.*

Proof. We prove the lemma by induction from $i=k$ down to $i=1$. The base case of $i=k$ is true because $\rho_k = \sqrt{\varepsilon}$ and $d(p,c_{\tau_k}) \geq (1-\sqrt{\varepsilon})\gamma_{\tau_k}$ by the assumption of the lemma.

The induction hypothesis is $d(p,c_{\tau_i}) \geq (1-\rho_i)\gamma_{\tau_i}$ for some $i \in [2,k]$. Let H denote $\mathrm{aff}(\partial B_{\tau_i} \cap \partial B_{\tau_{i-1}})$. There are two cases in the analysis of $d(p,c_{\tau_{i-1}})$ depending on $\angle_a(\tau_{i-1},H)$.

Suppose that $\angle_a(\tau_{i-1},H) \leq 10\mu\varepsilon\csc(\alpha/2)$. We apply Lemma 6.2 with $c = 10\csc(\alpha/2)$ to obtain $\gamma_{\tau_{i-1}} < (1+24\mu\varepsilon+40\mu\varepsilon\csc(\alpha/2))\gamma_{\tau_i}$ and $d(c_{\tau_{i-1}},c_{\tau_i}) < (24\mu\varepsilon+40\mu\varepsilon\csc(\alpha/2))\gamma_{\tau_{i-1}}$. Then,

$$\begin{aligned}
d(p,c_{\tau_{i-1}}) &\geq d(p,c_{\tau_i}) - d(c_{\tau_{i-1}},c_{\tau_i}) \\
&> (1-\rho_i)\gamma_{\tau_i} - (24\mu\varepsilon+40\mu\varepsilon\csc(\alpha/2))\gamma_{\tau_{i-1}} \\
&> \frac{1-\rho_i}{1+24\mu\varepsilon+40\mu\varepsilon\csc(\alpha/2)}\gamma_{\tau_{i-1}} - (24\mu\varepsilon+40\mu\varepsilon\csc(\alpha/2))\gamma_{\tau_{i-1}} \\
&> (1-24\mu\varepsilon-40\mu\varepsilon\csc(\alpha/2))(1-\rho_i)\gamma_{\tau_{i-1}} - (24\mu\varepsilon+40\mu\varepsilon\csc(\alpha/2))\gamma_{\tau_{i-1}} \\
&= (1-\rho_i)\gamma_{\tau_{i-1}} - (2-\rho_i)(24\mu\varepsilon+40\mu\varepsilon\csc(\alpha/2))\gamma_{\tau_{i-1}} \\
&> (1-\rho_i)\gamma_{\tau_{i-1}} - 2(24\mu\varepsilon+40\mu\varepsilon\csc(\alpha/2))\gamma_{\tau_{i-1}}.
\end{aligned}$$

The inequality $2(24\mu\varepsilon+40\mu\varepsilon\csc(\alpha/2)) \leq \sqrt{\varepsilon}$ is satisfied for a sufficiently small ε, and $\rho_i \geq \sqrt{\varepsilon}$ by definition. Therefore,

$$d(p,c_{\tau_{i-1}}) > (1-2\rho_i)\gamma_{\tau_{i-1}} > (1-\rho_{i-1})\gamma_{\tau_{i-1}}.$$

Suppose that $\angle_a(\tau_{i-1},H) > 10\mu\varepsilon\csc(\alpha/2)$. If p does not lie inside $B_{\tau_{i-1}}$, then $d(p,c_{\tau_{i-1}}) \geq \gamma_{\tau_{i-1}}$ and we are done. Assume for the rest of the proof that p lies inside $B_{\tau_{i-1}}$. Let G be the plane that contains the edge $\tau_{i-1} \cap \tau_i$ and is perpendicular to τ. Refer to Figure 16(a). Let Δ be the triangle with p and the endpoints of $\tau_i \cap \tau_{i-1}$ as vertices. Lemma 6.3 implies that

$$\angle_a(\Delta,\tau_{i-1}) < 10\mu\varepsilon\csc(\alpha/2). \tag{21}$$

We first show that H separates τ_{i-1} from p and τ_i. The dihedral angle at $\tau_{i-1} \cap \tau_i$ is greater than $\pi - 12\mu\varepsilon$ by Lemma 3.2(iv). Since $10\mu\varepsilon\csc(\alpha/2) > 12\mu\varepsilon$, H must separate τ_{i-1} from τ_i. We show that H separates τ_{i-1} and Δ too. Let q and q_{i-1} be two arbitrary vertices of τ and τ_{i-1}, respectively. Recall that p lies inside B_τ by the assumption of the lemma, and we are in the case of p lying inside $B_{\tau_{i-1}}$. Therefore,

$$\begin{aligned}
d(q,q_{i-1}) &\leq d(q,p) + d(q_{i-1},p) \\
&\leq 2\gamma_\tau + 2\gamma_{\tau_{i-1}} \\
&\leq 2\mu\varepsilon f(q) + 2\mu\varepsilon f(q_{i-1}) && (\because \text{Lemma 3.2(i)}) \\
&\leq \frac{4\mu\varepsilon}{1-2\mu\varepsilon} f(q). && (\because \text{Lemma 2.1(v)})
\end{aligned}$$

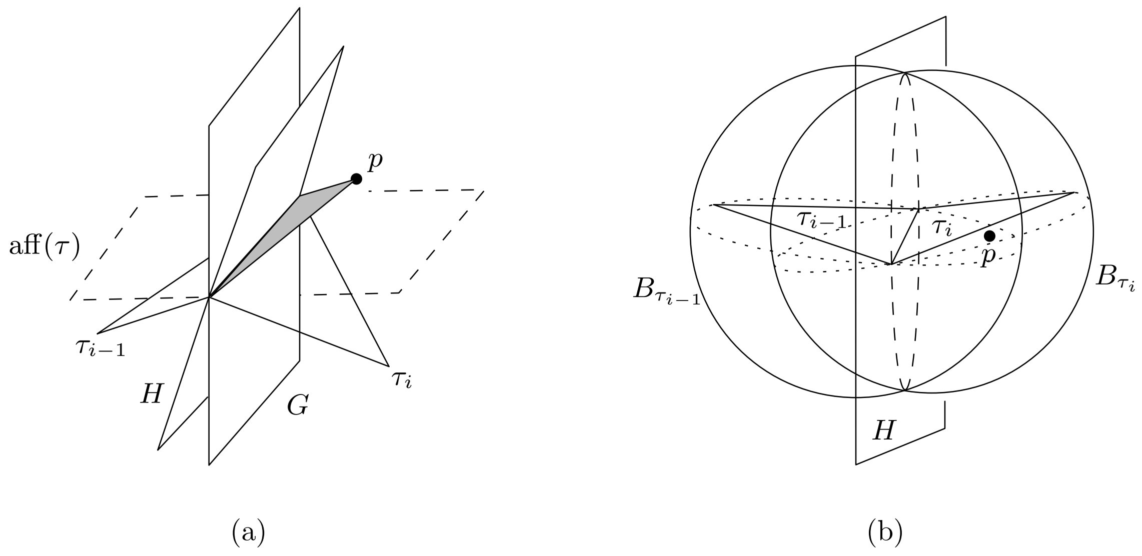

Figure 16: (a) The dashed plane is aff(τ). The shaded triangle is Δ and its vertices are p and the endpoints of $\tau_{i-1} \cap \tau_i$. (b) p is farther from $B_{\tau_{i-1}}$ than B_{τ_i} in terms of the power distance.

It follows that

$$\begin{aligned} \angle_a(\mathbf{n}_\tau, \mathbf{n}_{\tau_{i-1}}) &\leq \angle_a(\mathbf{n}_q, \mathbf{n}_\tau) + \angle_a(\mathbf{n}_{q_{i-1}}, \mathbf{n}_{\tau_{i-1}}) + \angle(\mathbf{n}_q, \mathbf{n}_{q_{i-1}}) \\ &< 12\mu\varepsilon + \frac{4\mu\varepsilon}{1-6\mu\varepsilon} \qquad (\because \text{ Lemmas 2.1(i) and 3.2(i)}) \\ &< 17\mu\varepsilon. \end{aligned}$$

The last step follows from the assumption that ε is sufficiently small. Since the plane G is perpendicular to τ, the inequality $\angle_a(\mathbf{n}_\tau, \mathbf{n}_{\tau_{i-1}}) < 17\mu\varepsilon$ implies that $\angle_a(G, \tau_{i-1}) > \pi/2 - 17\mu\varepsilon$. By the definition of seq(p, τ), G separates τ_{i-1} from both p and τ_i. This forces the dihedral angle between τ_{i-1} and Δ to be greater than $\pi/2 - 17\mu\varepsilon$. Since $\angle_a(\Delta, \tau_{i-1}) < 10\mu\varepsilon \csc(\alpha/2)$ by (21), the dihedral angle between τ_{i-1} and Δ must be greater than $\pi - 10\mu\varepsilon\csc(\alpha/2)$ then. Recall that we are in the case of $\angle_a(\tau_{i-1}, H) > 10\mu\varepsilon\csc(\alpha/2)$, which means that τ_{i-1} and Δ must be separated by H. It follows that H separates τ_{i-1} from both p and τ_i.

Refer to Figure 16(b). The plane H is the bisector of $B_{\tau_{i-1}}$ and B_{τ_i} with respect to the power distance [4]. Since the edge $\tau_{i-1} \cap \tau_i$ is legal by assumption of the lemma, $B_{\tau_{i-1}}$ contains the portion of B_{τ_i} on the side of H opposite to p. This implies that the power distance of p from $B_{\tau_{i-1}}$ is at least the power distance of p from B_{τ_i}. That is,

$$d(p, c_{\tau_{i-1}})^2 - \gamma_{\tau_{i-1}}^2 \geq d(p, c_{\tau_i})^2 - \gamma_{\tau_i}^2.$$

By the induction hypothesis, $d(p, c_{\tau_i}) \geq (1-\rho_i)\gamma_{\tau_i}$. Since τ_{i-1} and τ_i share an edge and every angle in τ_i is at least α, $\gamma_{\tau_{i-1}} \geq \gamma_{\tau_i}\sin\alpha$. Therefore,

$$d(p, c_{\tau_{i-1}})^2 \geq (1-\rho_i)^2\gamma_{\tau_i}^2 - \gamma_{\tau_i}^2 + \gamma_{\tau_{i-1}}^2 \geq \left(1 - \rho_i(2-\rho_i)(\csc\alpha)^2\right)\gamma_{\tau_{i-1}}^2.$$

We obtain the desired result of $d(p, c_{\tau_{i-1}})^2 \geq (1-\rho_{i-1})^2\gamma_{\tau_{i-1}}^2$ if

$$\begin{aligned} & 1 - \rho_i(2-\rho_i)(\csc\alpha)^2 &\geq\; & (1-\rho_{i-1})^2 \\ \Leftrightarrow\; & \rho_i(2-\rho_i)(\csc\alpha)^2 &\leq\; & \rho_{i-1}(2-\rho_{i-1}) \\ \Leftrightarrow\; & 2-\rho_i &\leq\; & 2(2-\rho_{i-1}) \end{aligned}$$

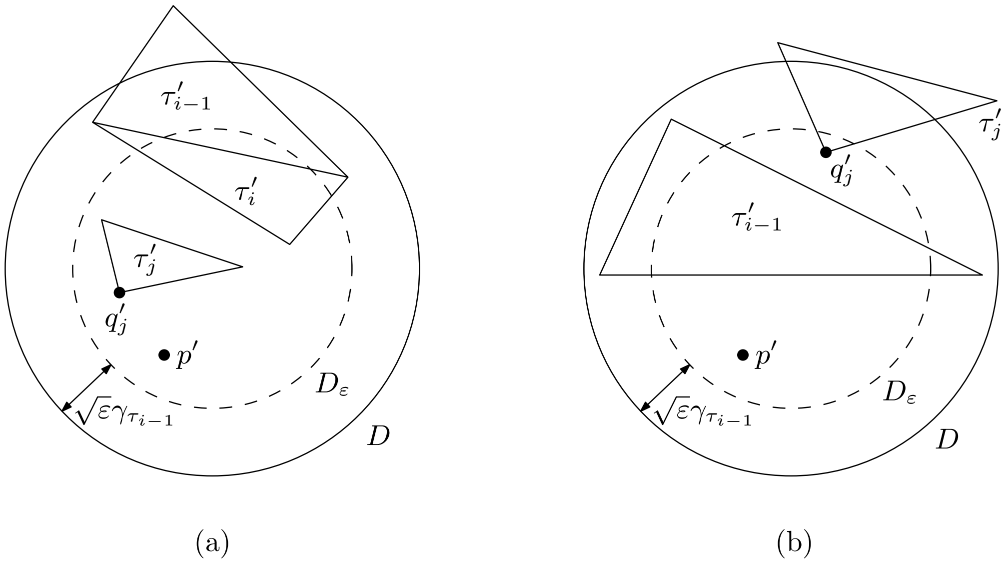

Figure 17: (a) G separates τ_{i-1} from q_j and τ_i. (b) If p' and q'_j are in different regions in D_ε outside τ'_{i-1}, it is impossible for the oriented segment from the vertex v of τ to p' in the definition of $\text{seq}(p,\tau)$ to intersect τ'_{i-1} before τ'_j.

The inequality $2-\rho_i \le 2(2-\rho_{i-1})$ holds when $\rho_i \le (\sin^2\alpha)/2$ because $2-\rho_{i-1} = 2-2\rho_i(\csc\alpha)^2 \ge 1$ then. On the other hand, if $\rho_i > (\sin^2\alpha)/2$, then $\rho_{i-1} = 2(\csc\alpha)^2\rho_i > 1$ and, therefore, the inequality $d(p, c_{\tau_{i-1}}) \ge (1-\rho_{i-1})\gamma_{\tau_{i-1}}$ is trivially true. ◻

Let $\lambda = 1-\varepsilon^c$ for some fixed constant $c \in (0, 0.5)$. The next result shows that if $d(p, c_\tau) < \lambda\gamma_\tau$ and $d(p, c_{\tau_i}) < (1-\sqrt{\varepsilon})\gamma_{\tau_i}$ for $i \in [1, k-1]$, then k is at most some constant.

Lemma 6.6 *Let $\lambda = 1-\varepsilon^c$ for some fixed constant $c \in (0, 0.5)$. Let T be a dense (ε,α)-mesh of Σ for a sufficiently small ε. Suppose that a vertex $p \in T$ and a triangle $\tau \in T$ satisfy $d(p, c_\tau) < \lambda\gamma_\tau$. If $(\tau_1, \tau_2, \ldots, \tau_k)$ is a prefix of $\text{seq}(p,\tau)$ such that $d(p, c_{\tau_i}) < (1-\sqrt{\varepsilon})\gamma_{\tau_i}$ and $\tau_i \cap \tau_{i+1}$ is legal for $i \in [1, k-1]$, then $k \le \kappa_{\text{seq}} = 2\pi/(\alpha - \sin\alpha) + 2$.*

Proof. For $j \in [1, k]$, let q_j denote an arbitrary vertex of τ_j and define $\ell(j)$ to be the smallest index in $[1, k]$ such that $d(q_j, c_{\tau_{\ell(j)}}) \ge (1-\sqrt{\varepsilon})\gamma_{\tau_{\ell(j)}}$. Note that $\ell(j) \le j$ because $d(q_j, c_{\tau_j}) = \gamma_{\tau_j} > (1-\sqrt{\varepsilon})\gamma_{\tau_j}$. The definition of $\ell(j)$ implies that $d(q_j, c_i) < (1-\sqrt{\varepsilon})\gamma_{\tau_i}$ for $i \in [1, \ell(j)-1]$. We first prove a claim. For $i \in [1, k]$, define $\rho_i = 2^{k-i}(\csc\alpha)^{2k-2i}\sqrt{\varepsilon}$ as in Lemma 6.5.

Claim 6.1 *For every $j \in [1, k]$ and every $i \in [1, \ell(j)]$, $d(q_j, c_{\tau_i}) \ge (1-\rho_{k-\ell(j)+i})\gamma_{\tau_i}$.*

Proof. The proof works by induction from $i = \ell(j)$ down to $i = 1$. The base case is true because, by the definition of $\ell(j)$, $d(q_j, c_{\tau_{\ell(j)}}) \ge (1-\sqrt{\varepsilon})\gamma_{\tau_{\ell(j)}} = (1-\rho_k)\gamma_{\tau_{\ell(j)}}$. The induction hypothesis is $d(q_j, c_{\tau_i}) \ge (1-\rho_{k-\ell(j)+i})\gamma_{\tau_i}$ for some $i \in [2, \ell(j)]$.

Consider τ_{i-1}. Since $i-1 < \ell(j)$, the definition of $\ell(j)$ implies that $d(q_j, c_{\tau_{i-1}}) < (1-\sqrt{\varepsilon})\gamma_{\tau_{i-1}}$. We want to apply an analysis similar to the proof of Lemma 6.5.

We first show that τ_{i-1} is separated from both q_j and τ_i by the plane G through $\tau_{i-1} \cap \tau_i$ perpendicular to $\text{aff}(\tau)$. Refer to Figure 17(a). Let p', q'_j, τ'_i, τ'_{i-1}, and $c'_{\tau_{i-1}}$ be the projections of p, q_j, τ_i, τ_{i-1}, and $c_{\tau_{i-1}}$ in $\text{aff}(\tau)$, respectively. Let D denote the disk in $\text{aff}(\tau)$ with center $c'_{\tau_{i-1}}$ and radius $\gamma_{\tau_{i-1}}$. Let D_ε be the concentric disk in $\mathrm{aff}(\tau)$ with center $c'_{\tau_{i-1}}$ and radius $(1-\sqrt{\varepsilon})\gamma_{\tau_{i-1}}$.

Recall that $d(p, c_{\tau_{i-1}}) < (1-\sqrt{\varepsilon})\gamma_{\tau_{i-1}}$ by the assumption of the lemma and $d(q_j, c_{\tau_{i-1}}) < (1-\sqrt{\varepsilon})\gamma_{\tau_{i-1}}$ as $i-1 < \ell(j)$. This implies that q'_j and p' belong to D_ε.

We argue that the vertices of τ'_{i-1} lie outside D_ε as follows. Let q be an arbitrary vertex of τ. By the definition of $\mathrm{seq}(p,\tau)$, every triangle in $\mathrm{seq}(p,\tau)$ intersects $B(c_\tau, \gamma_\tau)$. Therefore, $d(q, \tau_{i-1}) \leq 2\gamma_\tau \leq 2\mu\varepsilon f(q)$ by Lemma 3.2(i), and then $\angle_a(\mathbf{n}_\tau, \mathbf{n}_{\tau_{i-1}}) \leq \angle_a(\mathbf{n}_q, \mathbf{n}_\tau) + \angle_a(\mathbf{n}_q, \mathbf{n}_{\tau_{i-1}})$, which is less than $18\mu\varepsilon$ by Lemmas 3.2(i) and 5.1. The distances from $c'_{\tau_{i-1}}$ to the vertices of τ'_{i-1} are at least $\gamma_{\tau_{i-1}}\cos(\angle_a(\mathbf{n}_\tau, \mathbf{n}_{\tau_{i-1}})) > \gamma_{\tau_{i-1}}\cos(18\mu\varepsilon) > (1-\sqrt{\varepsilon})\gamma_{\tau_{i-1}}$ for a sufficiently small ε. Therefore, the vertices of τ'_{i-1} lie outside D_ε.

The edges of τ'_{i-1} partitions D_ε into no more than four regions, at most three outside τ'_{i-1} and at most one inside τ'_{i-1}. Since the projection of $T \cap B(c_\tau, \mu\varepsilon\gamma_\tau)$ onto $\mathrm{aff}(\tau)$ is injective by Lemma 5.4, p' and q'_j lie in the regions in D_ε outside τ'_{i-1}. Let v be the vertex of τ such that the segment vp' intersects the interior of τ. In other words, $\mathrm{seq}(p,\tau)$ is defined by vp', and therefore, the oriented segment from v to p' intersects τ'_{i-1} before τ'_j. It follows that p' and q'_j are in the same region in D_ε outside τ'_{i-1}; otherwise, it would be impossible for the oriented segment from v to p' to intersect τ'_{i-1} before τ'_j. Figure 17(b) shows such an impossible configuration. Therefore, p', q'_j and τ'_i are on the same side of $\tau'_{i-1} \cap \tau'_i$, which implies that G separates τ_{i-1} from q_j and τ_i.

Given that G separates τ_{i-1} from q_j and τ_i. we can invoke the same inductive proof of Lemma 6.5. The base case is changed from $d(p, c_{\tau_k}) \geq (1-\rho_k)\gamma_{\tau_k}$ to $d(q_j, c_{\tau_{\ell(j)}}) \geq (1-\rho_k)\gamma_{\tau_{\ell(j)}} = (1-\sqrt{\varepsilon})\gamma_{\tau_{\ell(j)}}$. We go through the same steps in the proof of Lemma 6.5 to show that $d(q_j, c_{\tau_i}) \geq (1-\rho_{k-\ell(j)+i})\gamma_{\tau_i}$ for $i \in [1, \ell(j)]$. □

The inequality $\rho_{k-\ell(j)+i} \leq \rho_{k-j+i}$ holds because $\ell(j) \leq j$. By using the relation $\rho_{k-\ell(j)+i} \leq \rho_{k-j+i}$ and setting $i=1$, Claim 6.1 gives

$$\forall\, j \in [1,k], \quad d(q_j, c_\tau) \geq (1-\rho_{k-j+1})\gamma_\tau. \tag{22}$$

Let $\beta = (1+\max\{\cos(2\alpha), \lambda\})/2$. Define k_ε to be the largest integer such that

$$\beta \leq 1 - 2^{k_\varepsilon-1}(\csc\alpha)^{2k_\varepsilon-2}\sqrt{\varepsilon}.$$

There exists such an integer k_ε because ε is sufficiently small and $\sqrt{\varepsilon}$ is asymptotically smaller than ε^c. For $i \in [1, \min\{k_\varepsilon, k\}]$, $1-2^{k_\varepsilon-1}(\csc\alpha)^{2k_\varepsilon-2}\sqrt{\varepsilon} \leq 1-2^{i-1}(\csc\alpha)^{2i-2}\sqrt{\varepsilon} = 1-\rho_{k-i+1}$. It follows that

$$\forall\, i \in [1, \min\{k_\varepsilon, k\}], \quad \beta \leq 1-\rho_{k-i+1}. \tag{23}$$

Next, we prove a claim on the projection of τ_i onto $\mathrm{aff}(\tau)$.

Claim 6.2 *For $i \in [1,k]$, let τ'_i and q'_i denote the projections of τ_i and an arbitrary vertex q_i of τ_i onto $\mathrm{aff}(\tau)$, respectively.*

(i) *For every $i \in [1,k]$, each angle of τ'_i is greater than $\alpha/2$.*

(ii) *For every $i \in [1, \min\{k_\varepsilon, k\}]$ and every vertex q'_i of τ'_i, $d(c_\tau, q'_i)$ is greater than $\gamma_\tau \cdot \max\{\cos(2\alpha), \lambda\}$.*

Proof. Let q be an arbitrary vertex of τ. For every $i \in [1,k]$ and every vertex q_i of τ_i, since τ_i intersects $B(c_\tau, \gamma_\tau)$ by the definition of $\mathrm{seq}(p,\tau)$, Lemma 3.2(i) implies that $d(q, q_i) \leq 2\gamma_\tau + 2\gamma_{\tau_i} \leq 2\mu\varepsilon f(q) + 2\mu\varepsilon f(q_i)$, and then Lemma 2.1(v) implies that

$$\forall\, i \in [1,k],\ \forall \text{vertex } q_i \text{ of } \tau_i, \quad d(q, q_i) \leq \frac{4\mu\varepsilon}{1-2\mu\varepsilon} f(q). \tag{24}$$

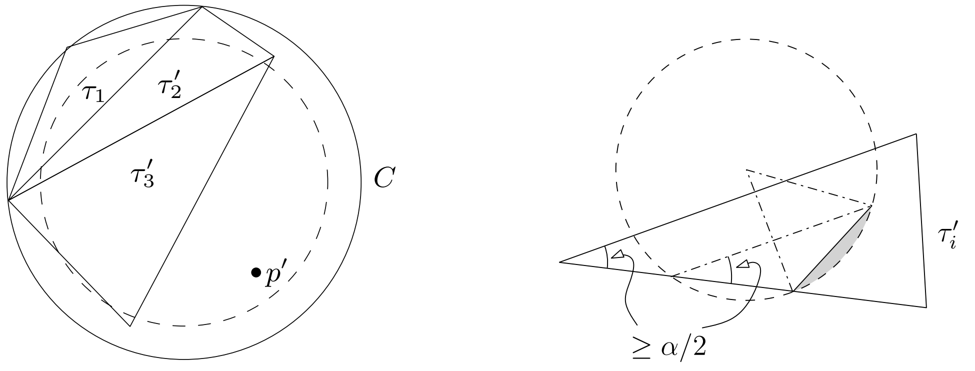

(a) (b)

Figure 18: (a) The dashed circle has radius $\gamma_{\tau_1} \cdot \max\{\cos(2\alpha), \lambda\}$. For every $i \in [1, \min\{k_\varepsilon, k\}]$, the vertices of τ_i' are outside the dashed circle. (b) The area of τ_i' inside the dashed circle is greater than the area of the shaded region, which is $(\alpha - \sin\alpha)/2$ times the square of the radius of the dashed circle.

By Lemma 5.4, $T \cap B(q, 5\mu\varepsilon f(q))$ projects injectively onto $\mathrm{aff}(\tau)$. Since ε is sufficiently small, by (24), τ_i is completely contained in $T \cap B(q, 5\mu\varepsilon f(q))$ for $i \in [1, k]$. Let τ_i' denote the projection of τ_i in $\mathrm{aff}(\tau)$. For $i \in [1, k]$, since the lengths of the edges of τ_i are at most $2\gamma_{\tau_i} \leq 2\mu\varepsilon f(q_i)$, Lemmas 2.1(i) and 3.2(i) and (24) imply that the angle between $\mathrm{aff}(\tau)$ and each edge of τ_i incident to q_i is less than $1.02\mu\varepsilon + \angle(\mathbf{n}_{q_i}, \mathbf{n}_q) + \angle_a(\mathbf{n}_q, \mathbf{n}_\tau) < 1.02\mu\varepsilon + \frac{4\mu\varepsilon}{1-6\mu\varepsilon} + 6\mu\varepsilon$, which is less than $12\mu\varepsilon$ for a sufficiently small ε. Therefore, for $i \in [1, k]$, each angle of τ_i' is at least $\alpha - 24\mu\varepsilon > \alpha/2$, establishing the correctness of (i).

Take any $i \in [1, \min\{k_\varepsilon, k\}]$ and any vertex q_i of τ_i. Let q_i' denote the projection of q_i onto $\mathrm{aff}(\tau)$. By (24) and Lemmas 2.1(i) and 3.2(i), $\angle_a(qq_i, \tau) \leq \pi/2 - \angle_a(qq_i, \mathbf{n}_q) + \angle_a(\mathbf{n}_q, \mathbf{n}_\tau) < 0.51 \cdot \frac{4\mu\varepsilon}{1-2\mu\varepsilon} + 6\mu\varepsilon < 9\mu\varepsilon$. Thus, $d(q, q_i') > d(q, q_i)\cos(9\mu\varepsilon)$ and $d(q_i, q_i') < d(q, q_i)\sin(9\mu\varepsilon)$.

If $d(q, q_i) \geq 2\gamma_\tau$, then $d(c_\tau, q_i') \geq d(q, q_i') - d(c_\tau, q) > 2\gamma_\tau\cos(9\mu\varepsilon) - \gamma_\tau$. Since $\lambda = 1 - \varepsilon^c$ for some fixed constant $c \in (0, 0.5)$, the inequality $2\cos(9\mu\varepsilon) - 1 > \max\{\cos(2\alpha), \lambda\}$ holds for a sufficiently small ε, i.e., $d(c_\tau, q_i') > \gamma_\tau \cdot \max\{\cos(2\alpha), \lambda\}$.

If $d(q, q_i) < 2\gamma_\tau$, then by (22) and (23), $d(c_\tau, q_i) \geq (1 - \rho_{k-i+1})\gamma_\tau \geq \beta\gamma_\tau$. Thus, $d(c_\tau, q_i') \geq d(c_\tau, q_i) - d(q_i, q_i') > \beta\gamma_\tau - 2\gamma_\tau\sin(9\mu\varepsilon)$. By the definition of β, the inequality $\beta - 2\sin(9\mu\varepsilon) > \max\{\cos(2\alpha), \lambda\}$ holds for a sufficiently small ε. Thus, $d(c_\tau, q_i') > \gamma_\tau \cdot \max\{\cos(2\alpha), \lambda\}$. This proves the correctness of (ii). $\square$

Refer to Figure 18(a). The solid circle denotes the circumcircle C of τ and the dashed circle denotes the concentric circle with radius $\gamma_\tau \cdot \max\{\cos(2\alpha), \lambda\}$. By Claim 6.2(i,ii), for $i \in [1, \min\{k_\varepsilon, k\}]$, the angles of τ_i' are at least $\alpha/2$ and the vertices of τ_i' are outside the dashed circle. Since $d(p, c_\tau) < \lambda\gamma_\tau$ by the assumption of the lemma, p' lies inside the dashed circle. Since the radius of the dashed circle is at least $\gamma_\tau\cos(2\alpha)$ and the angles of τ are at least α, one can verify that the dashed circle intersects τ.

Consider $\tau_i' \cap \tau_{i+1}'$ for some $i \in [1, \min\{k_\varepsilon, k\} - 1]$ such that τ_i' intersects the dashed circle. By the definition of $\mathrm{seq}(p, \tau)$, $\mathrm{aff}(\tau_i' \cap \tau_{i+1}')$ separates p' from τ_i' in $\mathrm{aff}(\tau)$. Therefore, $\tau_i' \cap \tau_{i+1}'$ must cross the dashed circle in order that τ_i' intersects the dashed circle and p' lies inside the dashed circle. Inductively, this implies that for $i \in [2, \min\{k_\varepsilon, k\} - 1]$, at least two edges of τ_i' cross the dashed circle. Since the minimum angle of τ_i' is at least $\alpha/2$, one can verify that the area of τ_i' inside the dashed circle is at least $(\alpha - \sin\alpha)/(2\pi)$ times the area of the dashed circle. Refer to Figure 18(b). This implies that the range $[2, \min\{k_\varepsilon, k\} - 1]$ has cardinality at most $2\pi/(\alpha - \sin\alpha)$ and, therefore, $\min\{k_\varepsilon, k\} \leq 2\pi/(\alpha - \sin\alpha) + 2$. We set $\kappa_{\text{seq}} = 2\pi/(\alpha - \sin\alpha) + 2$. By choosing ε small enough so that $k_\varepsilon > \kappa_{\text{seq}}$, we obtain $k = \min\{k_\varepsilon, k\} \leq \kappa_{\text{seq}}$. □

We are ready to prove Theorem 2. Theorem 2(i, ii) state that edge flips make the diametric ball of every triangle almost empty and the circumradius of every triangle at most $\varepsilon + O(\varepsilon^{1+c})$ times the local feature size at every vertex of the triangle, where c is a fixed constant in $(0, 0.5)$. Plugging this circumradius bound into Lemma 3.2(i, iv) shows that edge flips can provably smooth the mesh. Theorem 2(iii) further states that edge flips can be applied to the local neighborhood of any subset of vertices in time linear in the subset size.

Theorem 2 *For every constant $c \in (0, 0.5)$ and every constant $\alpha \in [0, \pi/3]$, there exists $\varepsilon_0 \in (0, 1)$ depending on c and α such that for every $\varepsilon \in (0, \varepsilon_0]$, if T is an (ε, α)-mesh of a connected closed smooth surface, then the following properties are satisfied.*

(i) *We can flip flippable edges in T until no edge is flippable in time linear in the number of vertices in T. An (ε, α)-mesh T' is produced in the end and for every triangle $\tau \in T'$, $B(c_\tau, (1 - \varepsilon^c)\gamma_\tau)$ does not contain any vertex.*

(ii) *For every vertex $p \in T$ and every triangle $\tau \in \text{star}(p)$, if $B(c_\tau, (1 - \varepsilon^c)\gamma_\tau)$ does not contain any vertex, then $\gamma_\tau \leq (\varepsilon + O(\varepsilon^{1+c}))f(p)$.*

(iii) *Given any subset V of vertices of T, we can flip flippable edges in $O(|V|)$ time to produce an (ε, α)-mesh T' so that for every triangle $\tau \in T'$ that is incident to a vertex in V or a neighbor of a vertex in V, $B(c_\tau, (1 - \varepsilon^c)\gamma_\tau)$ does not contain any vertex.*

Proof. We assume that ε_0 is small enough that T is a dense (ε, α)-mesh. Let n denote the number of vertices in T. Consider (i). The edge flips take $O(n)$ time by Lemma 4.3. Suppose for the sake of contradiction that there is a vertex $p \in T'$ and a triangle $\tau \in T'$ such that $d(p, c_\tau) < (1 - \varepsilon^c)\gamma_\tau$. Let $(\tau_1 = \tau, \tau_2, \ldots, \tau_k)$ be the shortest prefix of $\text{seq}(p, \tau)$ such that $d(p, c_{\tau_k}) \geq (1 - \sqrt{\varepsilon})\gamma_{\tau_k}$. Lemmas 6.1 and 6.6 imply that $k \leq \kappa_{\text{seq}}$ and Lemma 6.5 implies that

$$d(p, c_\tau) \geq \left(1 - 2^{k-1}(\csc\alpha)^{2k-2}\sqrt{\varepsilon}\right)\gamma_\tau.$$

Since $c \in (0, 0.5)$, the inequality above implies that $d(p, c_\tau) \geq (1 - \varepsilon^c)\gamma_\tau$ for a sufficiently small ε, contradicting the assumption that $d(p, c_\tau) < (1 - \varepsilon^c)\gamma_\tau$. Hence, every vertex in T' is at distance $(1 - \varepsilon^c)\gamma_\tau$ or more from c_τ.

Consider (ii). Recall that ν denotes the nearest point map of the given closed surface. Take any triangle $\tau \in \text{star}(p)$. By Lemma 3.2(i), $\gamma_\tau \leq \mu\varepsilon f(p)$. We apply Lemma 2.2(iii) with $c = \mu$ to obtain:

$$d(c_\tau, \nu(c_\tau)) \leq 10\mu\varepsilon\gamma_\tau \tag{25}$$

$$d(p, \mu(c_\tau)) \leq (2\mu\varepsilon + 20\mu^2\varepsilon^2)f(p) \tag{26}$$

There is a vertex q at distance $\varepsilon f(\nu(c_\tau))$ or less from $\nu(c_\tau)$ as the vertices of T form an ε-sample of the closed surface. By (i), $d(q, c_\tau) \geq (1 - \varepsilon^c)\gamma_\tau$, which implies that $\varepsilon f(\nu(c_\tau)) \geq d(q, \nu(c_\tau)) \geq d(q, c_\tau) - d(c_\tau, \nu(c_\tau)) \geq (1 - \varepsilon^c - 10\mu\varepsilon)\gamma_\tau$. Therefore,

$$\gamma_\tau \leq \frac{\varepsilon}{1 - \varepsilon^c - 10\mu\varepsilon}\, f(\nu(c_\tau)).$$

By (26) and Lemma 2.1(v), $f(\nu(c_\tau)) \le (1 + 2\mu\varepsilon + 20\mu^2\varepsilon^2)f(p)$. Therefore, for a sufficiently small ε, $\gamma_\tau \le \frac{\varepsilon}{1-\varepsilon^c-10\mu\varepsilon} f(\nu(c_\tau)) = (\varepsilon + O(\varepsilon^{1+c}))f(p)$.

Consider (iii). We flip flippable edges that are within κ_{seq} edges away from any vertex $p \in V$ until no such flippable edge can be found. Lemma 4.2 implies that there are $O(1)$ vertices within κ_{seq} edges away from p in the graph $\mathcal{G}$ defined in Lemma 4.2. Therefore, the edges that are flipped are incident to $O(|V|)$ vertices. The neighborhood of each such vertex can only change $O(1)$ times as argued in the proof of Lemma 4.3. It follows that the edge flips terminate in $O(|V|)$ time. Let T' denote the dense (ε, α)-mesh produced.

Let τ be any triangle in T' that is incident to a vertex in V or a neighbor of a vertex in V. Suppose for the sake of contradiction that there is a vertex $p \in T'$ such that $d(p, c_\tau) < (1-\varepsilon^c)\gamma_\tau$. Extract the shortest prefix $(\tau_1 = \tau, \tau_2, \ldots, \tau_k)$ of seq(p, τ) such that $d(p, c_{\tau_k}) \ge (1 - \sqrt{\varepsilon})\gamma_{\tau_k}$. Since every edge that is at most κ_{seq} edges away from τ is non-flippable, Lemmas 6.1 and 6.6 imply that $k \le \kappa_{\text{seq}}$. Then Lemma 6.5 implies that $d(p, c_\tau) \ge \left(1 - 2^{k-1}(\csc\alpha)^{2k-2}\sqrt{\varepsilon}\right)\gamma_\tau$, which yields the same contradiction of $d(p, c_\tau) \ge (1 - \varepsilon^c)\gamma_\tau$ as in the analysis of (i). Hence, every vertex in T' is at distance $(1 - \varepsilon^c)\gamma_\tau$ or more from c_τ. □

The algorithm in the proof of Theorem 2(iii) is simple, but it may be impractical as κ_{seq} is large. A more practical method works as follows. Let V be a subset of vertices. Let F denote the triangles incident to a vertex in V or a neighbor of a vertex in V. We initialize a variable m to be a small constant, say 3, and then flip flippable edges that are within m edges away from V, and then mark triangles that are within m edges away from V. (The set F is updated correspondingly.) Afterwards, we check if the diametric ball of every triangle in F intersects marked triangles only. If the diametric ball of some triangle in F intersects a non-marked triangle, then we double the value of m and repeat the above. Each iteration runs in $O(|V|)$ time. There are at most $O(\log \kappa_{\text{seq}})$ iterations, but potentially much smaller in practice.

7 Experiments

We performed experiments on three datasets, including the surface meshes of a sphere, a torus, and an object of a boomerang shape. The experiments were run on a machine with an Intel Xeon E5450 3.00GHZ CPU with 16GB RAM. Our program uses CGAL and is compiled with options "-O3 -DNDEBUG".

The spherical and toroidal surface meshes are generated as follows. We uniformly sample 10K points on a unit sphere and a torus with major radius 5 and minor radius 3. We run the Cocone algorithm [3, 9] on the sample points to produce the surface meshes. The output meshes of the Cocone algorithm are usually good, so we need to worsen its quality in order to demonstrate the effect of edge flips. We worsen the mesh quality by performing five rounds of random edge flips. In each round, every edge is flipped with probability 1/2, provided that the flip does not produce a dihedral angle smaller than 120°. Given the worsened mesh, the smoothed mesh is obtained by repeatedly flipping flippable edges.

For the sphere data, Table 1 shows the statistics of dihedral angles, deviation of triangle normals from the surface normals at the vertices of the trinagles, circumradii, and triangle angles. Figures 19–22 show the distributions of these quantities. There were 108953 edge flippability tests and 15148 edge flips. The total running time is 15 milliseconds.

Table 2 shows the statistics for the torus data. Figures 23–26 show the corresponding distributions. There were 107214 edge flippability tests and 14816 edge flips. The total running time is 16 milliseconds.

We also experimented with a boomerang-like object obtained from a snapshot of the deformation of a topological ball [13, 19]. There are 5802 vertices in the mesh and the deformation snapshot is produced by a mesh maintenance algorithm developed by us [7]. Figure 27 shows the snapshot. We worsen the snapshot in the same way as described before for the sphere and torus cases. Then, we smooth the mesh by repeatedly flipping flippable edges. Table 3 shows some statistics on the dihedral angles, circumradii, and triangle angles. (We do not provide statistics on normal deviations because we do not know the surface normals.) Figures 28–30 show the distributions of of these quantities. There were 58901 edge flippability tests and 8002 edge flips. The total running time is 7 milliseconds.

In all three datasets, it is clear that repeated edge flips are efficient and effective in improving the mesh quality and smoothing surface meshes.

8 Conclusion

We proposed the class of (ε, α)-meshes of closed surfaces and study the effect of edge flips on them. In $\mathbb{R}^3$, our definition of edge flippability is different from the usual empty circumsphere criterion. Given a dense (ε, α)-mesh, we prove that repeated edge flips can lower the circumradius of every mesh triangle to $\varepsilon + O(\varepsilon^\kappa)$ times the local feature size at any vertex of the triangle, where κ is any fixed constant in $(1, 1.5)$. Then, standard surface sampling results in the literature show that the normal deviation becomes smaller and the dihedral angles in the mesh become closer to π. That is, the mesh is smoother. This helps to explain why edge flips are effective in improving the mesh quality and smoothing surface meshes as observed in practice. Our experimental results also confirm this observation. The vertex densities in our datasets are not high in contrast to the condition on ε required by our theoretical results. A corollary of our techniques is that, in $\mathbb{R}^2$, every triangulation with a constant lower bound on the angles can be flipped in linear time to the Delaunay triangulation.

Acknowledgment

We would like to thank the anonymous reviewers for their helpful comments. We also thank Man-Kit Lau for helping us to generate Figure 3.

References

[1] P. Alliez, M. Meyer, and M. Desbrun. Interactive geometry remeshing. *ACM Transactions on Graphics*, 21 (2002), 347–354.

[2] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete & Computational Geometry*, 22(4):481–504, 1999.

[3] N. Amenta, S. Choi, T.K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *International Journal of Computational Geometry and Applications*, 12:125–141, 2002.

[4] F. Aurenhammer and R. Klein. Voronoi diagrams. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 201–290. Elsevier, 2000.

[5] J.-D. Boissonnat and F. Cazals. Natural neighbor coordinates of points on a surface. *Computational Geometry: Theory and Applications*, 19 (2001), 87–120.

[6] S.-W. Cheng and T.K. Dey. Maintaining deforming surface meshes. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, 2008, 112–121.

[7] S.-W. Cheng and J. Jin. Edge flips and deforming surface meshes. In *Proceedings of the 27th Annual Symposium on Computational Geometry*, 331–340, 2011.

[8] S.-W. Cheng, J. Jin, and M.-K. Lau. A fast and simple surface reconstruction algorithm. *Proceedings of the 28th Annual Symposium on Computational Geometry*, 2012, 69–78.

[9] T.K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis.* Cambridge University Press, New York, NY, USA, 2006.

[10] T. K. Dey, S. Funke and E.A. Ramos. Surface reconstruction in almost linear time under locally uniform sampling. *European Workshop on Computational Geometry*, Berlin, March 2001.

[11] N. Dyn, K. Hormann, S.-J. Kim, and D. Levin. *Optimizing 3D triangulations using discrete curvature analysis*, pages 135–146. Vanderbilt University, 2001.

[12] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, Springer, 1987.

[13] D. Enright, R. Fedkiw, J. Ferziger, and I. Mitchell. A hybrid particle level set method for improved interface capturing. *Journal of Computational Physics*, 183 (2002), 83–116.

[14] S. Funke and E.A. Ramos. Smooth-surface reconstruction in near-linear time. *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2002, 781–790.

[15] O.S. Galaktionov, P.D. Anderson, G.W.M. Peters, and F.N. Van de Vosse. An adaptive front tracking technique for three-dimensional transient flows. *International Journal for Numerical Methods in Fluids*, 32 (2000), 201–217.

[16] J. Giesen and U. Wagner. Shape dimension and intrinsic metric from samples of manifolds. *Discrete & Computational Geometry*, 32(2):245–267, 2004.

[17] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete & Computational Geometry*, 22(3):333–346, 1999.

[18] C. L. Lawson. Transforming triangulations. *Discrete Mathematics*, 3(4):365 – 372, 1972.

[19] R.J. LeVeque. High-resolution conservative algorithms for advection in incompressible flow. *SIAM Journal on Numerical Analysis*, 33 (1996), 627–665.

[20] G. L. Miller, T. Phillips, and D. Sheehy. Linear-size meshes. In *20th Canadian Conference on Computational Geometry*, 2008.

[21] V. Surazhsky and C. Gotsman. Explicit surface remeshing. *Proceedings of the Eurogrpahics Symposium on Geometry Processing*, 2005, 20–30.

Table 1: Surface meshes of a unit sphere.

	Normal deviation (mean/SD/max)	180°− Dihedral angle (mean/SD/max)	circumradius (mean/SD/max)	triangle angle (mean/SD/min/max)
Cocone output	1.5/0.56/3.8	1.4/0.87/5.3	0.027/0.010/0.066	60/29.4/0.2/170.6
Worsened mesh	5.5/7.68/59	8.8/10.56/60	0.10/0.13/0.87	60/52.4/0.07/179.3
Smoothed mesh	1.5/0.56/3.8	1.4/0.87/5.3	0.027/0.010/0.066	60/29.4/0.2/170.6

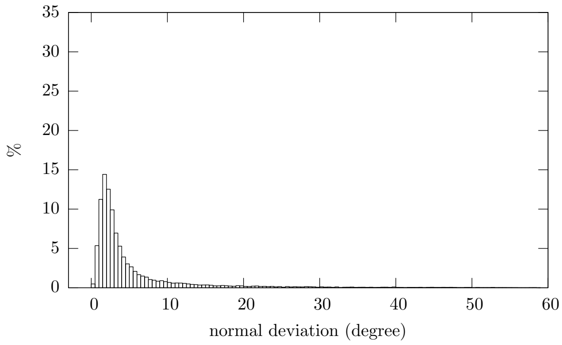

(a) Worsened mesh

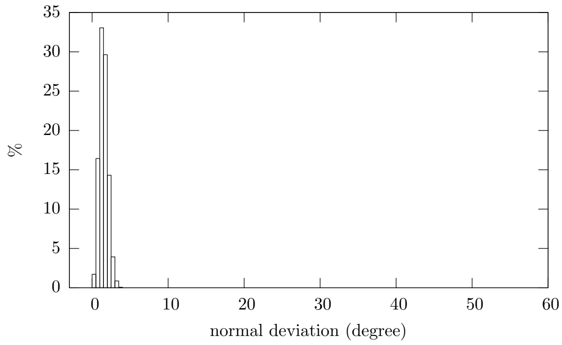

(b) Smoothed mesh

Figure 19: Distributions of the normal deviations in the spherical meshes.

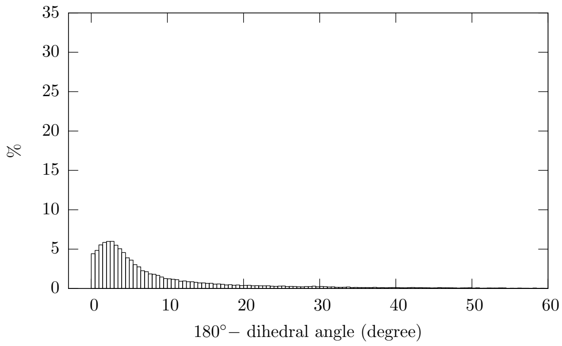

(a) Worsened mesh

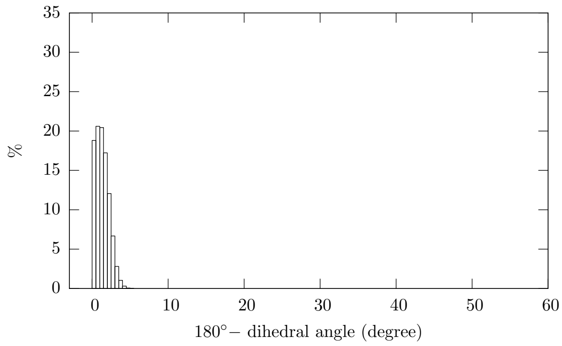

(b) Smoothed mesh

Figure 20: Distributions of the dihedral angles in the spherical meshes.

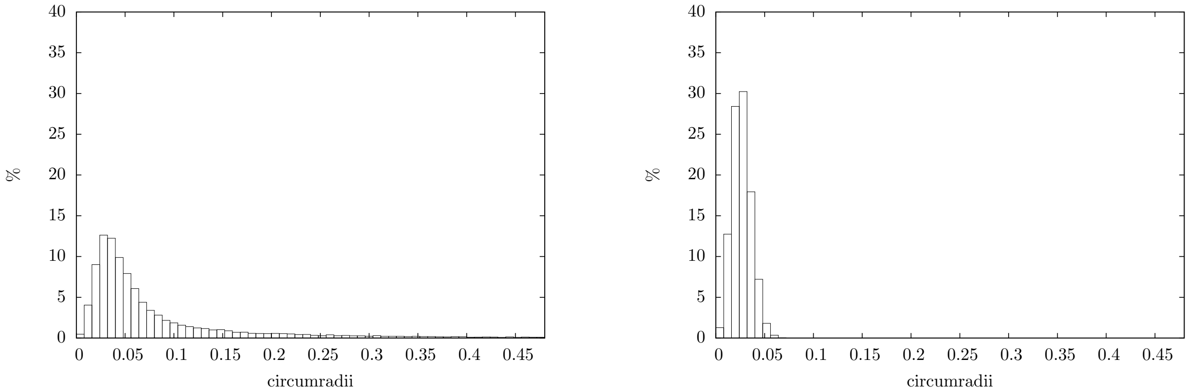

(a) Worsened mesh (b) Smoothed mesh

Figure 21: Distributions of the circumradii in the spherical meshes.

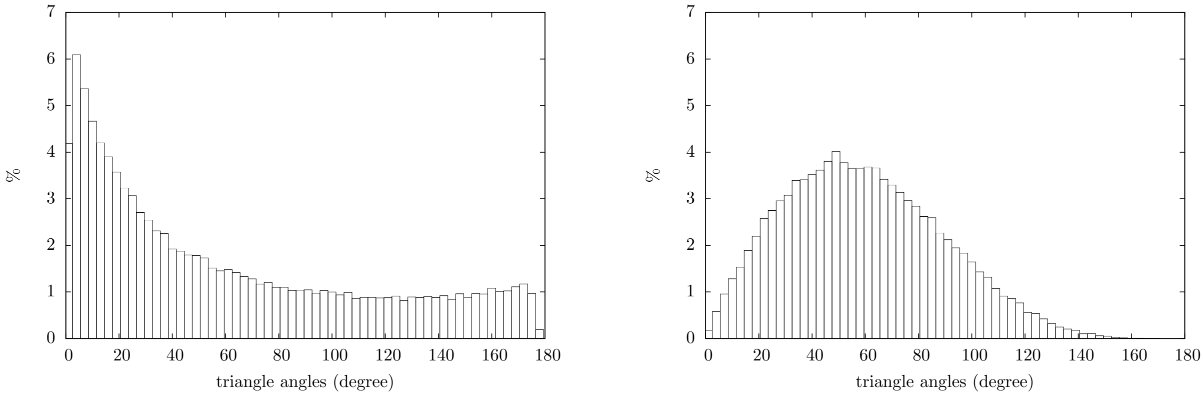

(a) Worsened mesh (b) Smoothed mesh

Figure 22: Distributions of the triangle angles in the spherical meshes.

Table 2: Surface meshes of a torus.

	Normal deviation (mean/SD/max)	$180^\circ-$ Dihedral angle (mean/SD/max)	circumradius (mean/SD/max)	triangle angle (mean/SD/min/max)
Cocone output	2.6/1.50/16.8	2.6/1.93/16.3	0.18/0.07/0.46	60/29.6/0.3.171.6
Worsened mesh	6.8/8.11/63.7	10/11.09/60	0.63/1.32/92.8	60/51.6/0.02/179.9
Smoothed mesh	2.6/1.49/15.8	2.6/1.92/16	0.18/0.07/0.44	60/29.6/0.3/171.6

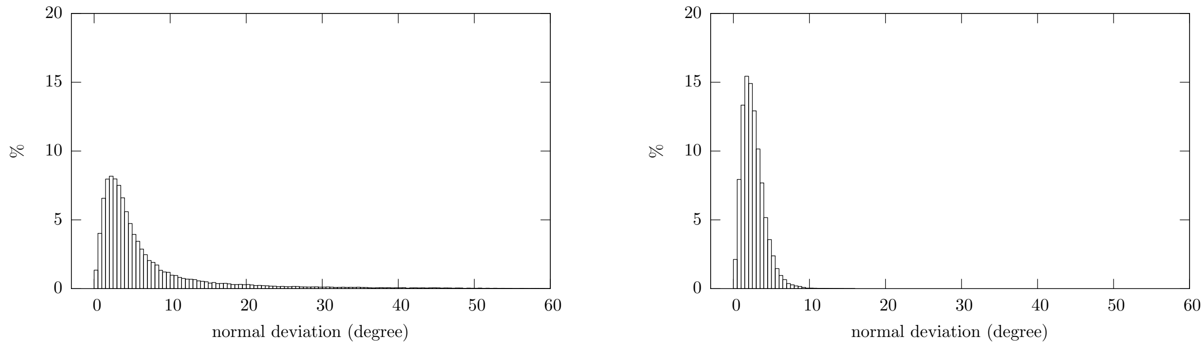

(a) Worsened mesh (b) Smoothed mesh

Figure 23: Distributions of the normal deviations in the toroidal meshes.

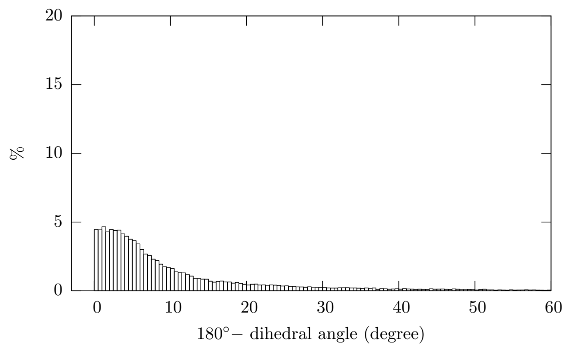

(a) Worsened mesh

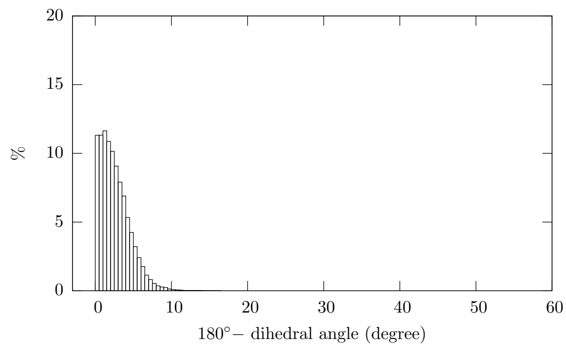

(b) Smoothed mesh

Figure 24: Distributions of the dihedral angles in the toroidal meshes.

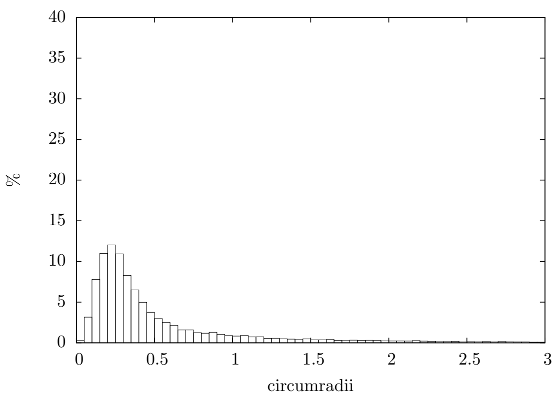

(a) Worsened mesh

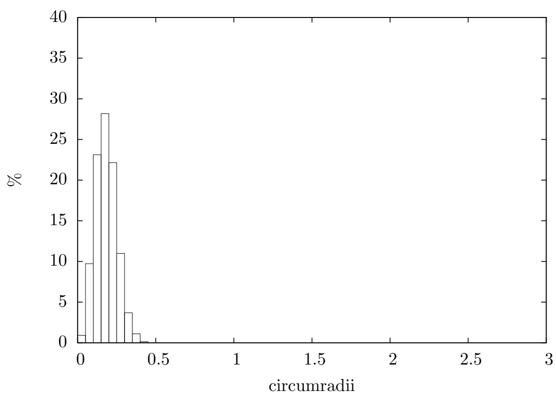

(b) Smoothed mesh

Figure 25: Distributions of the circumradii in the toroidal meshes.

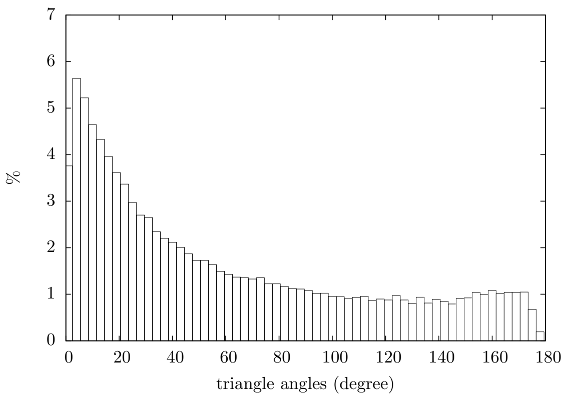

(a) Worsened mesh

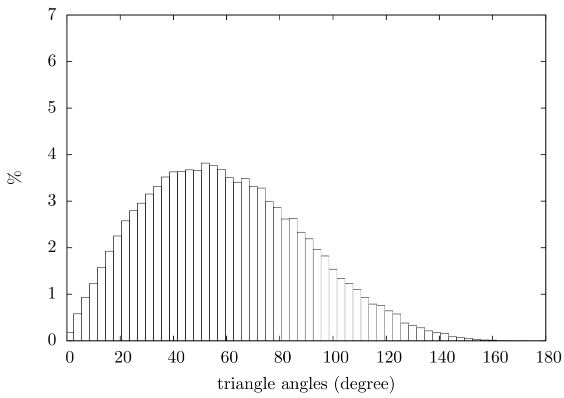

(b) Smoothed mesh

Figure 26: Distributions of the triangle angles in the toroidal meshes.

Table 3: Surface meshes of the boomerang-like object.

	180°− Dihedral angle (mean/SD/max)	circumradius (mean/SD/max)	triangle angle (mean/SD/min/max)
Deformation snapshot	4.3/5.23/50.8	0.006/0.003/0.028	60/23.3/13.9/143.9
Worsened mesh	11.2/12.11/60	0.024/0.078/6.3	60/50.2/0.03/179.9
Smoothed mesh	4.3/5.23/50.9	0.006/0.003/0.028	60/23.3/13.9/143.9

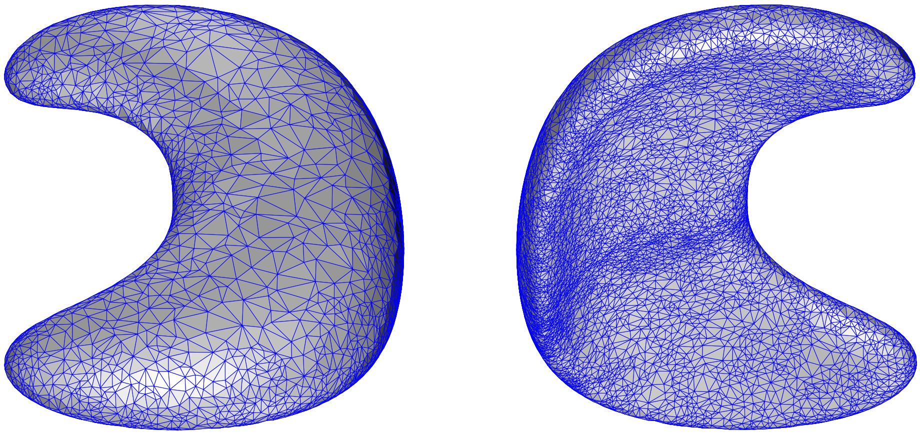

Figure 27: The snapshot produced by a maintenance algorithm developed by us [7].

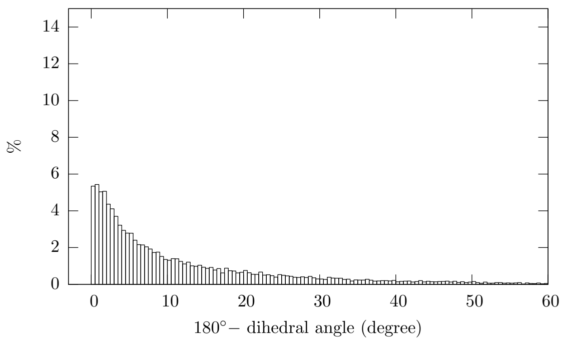

(a) Worsened mesh

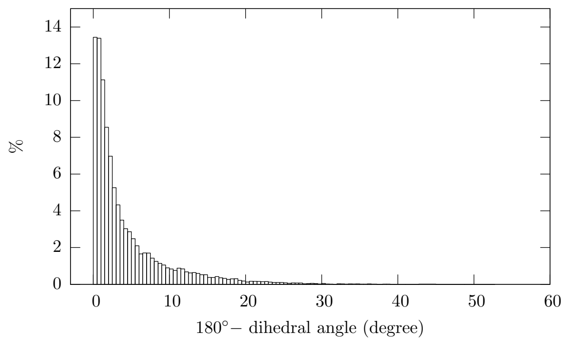

(b) Smoothed mesh

Figure 28: Distributions of the dihedral angles in the meshes of the boomerang-like object.

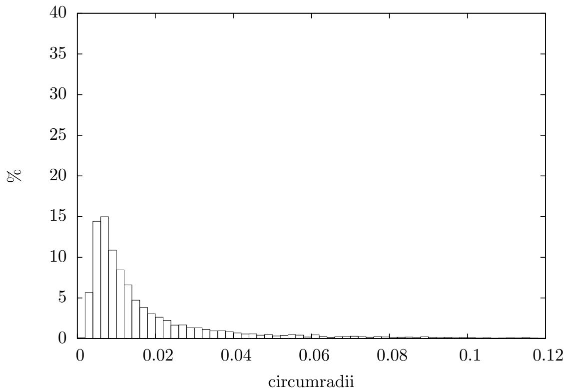

(a) Worsened mesh

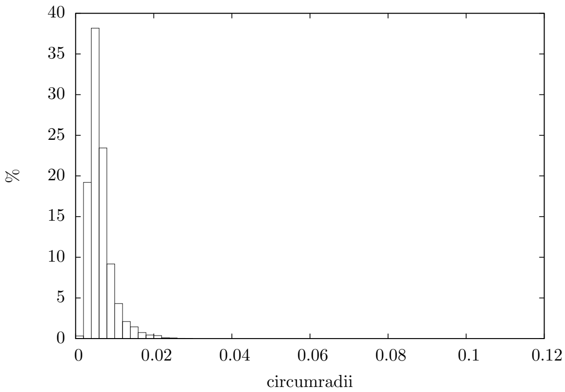

(b) Smoothed mesh

Figure 29: Distributions of the circumradii in the meshes of the boomerang-like object.

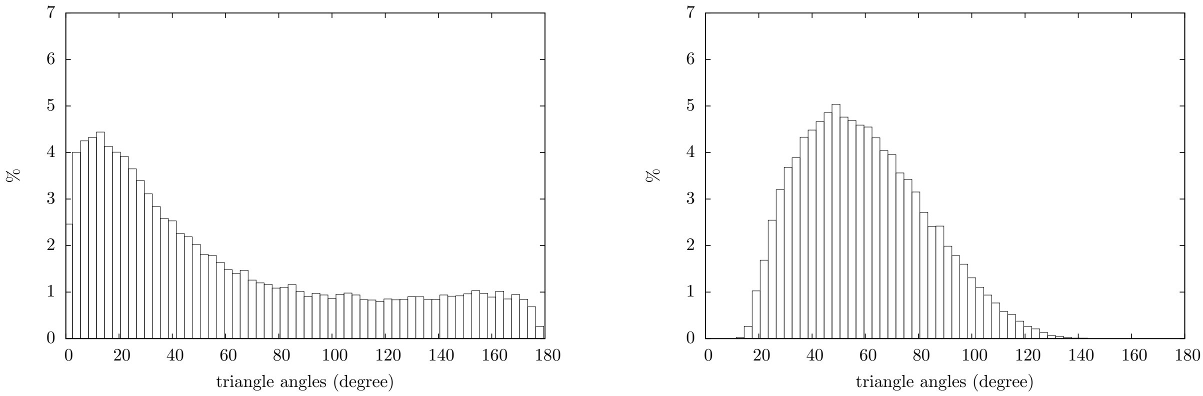

(a) Worsened mesh

(b) Smoothed mesh

Figure 30: Distributions of the triangle angles in the meshes of the boomerang-like object.