Sliver Exudation *

Siu-Wing Cheng†, Tamal K. Dey‡, Herbert Edelsbrunner§,
Michael A. Facello¶ and Shang-Hua Teng‖

Abstract

A sliver is a tetrahedron whose four vertices lie close to a plane and whose projection to that plane is a convex quadrilateral with no short edge. Slivers are notoriously common in 3-dimensional Delaunay triangulations even for well-spaced point sets. We show that if the Delaunay triangulation has the ratio property introduced in [15] then there is an assignment of weights so the weighted Delaunay triangulation contains no slivers. We also give an algorithm to compute such a weight assignment.

Keywords. Mesh generation, computational geometry; tetrahedral meshes, mesh quality, (weighted) Delaunay triangulations, slivers, algorithms.

1 Introduction

This paper studies slivers in 3-dimensional Delaunay triangulations and in particular, the question of how we can rid ourselves of slivers. This section introduces the general context in which this question arises and reviews what is known about it.

*Research by the second, third and fifth authors is partially supported by the U.S. Department of Energy, under Subcontract B341494, and by NSF under grant DMS 98-73945. Research of the third author is also supported by NSF under grant CCR-96-19542, and by ARO under grant DAAG55-98-1-0177. Research of the fifth author is also supported by an Alfred P. Sloan Fellowship.

†Department of Computer Science, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong.

‡Department of Computer Science and Engineering, Indian Institute of Technology, Kharagpur 721302, India.

§Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, and Raindrop Geomagic, Champaign, Illinois 61820.

¶Raindrop Geomagic, Champaign, Illinois 61820.

‖Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801.

Mesh generation. Meshes are cell complexes that decompose spatial domains for the purpose of numerical simulation and analysis. In this paper we exclusively consider meshes made up of tetrahedral cells. We use mathematical terminology whenever reasonable and define a *tetrahedral mesh* as a simplicial complex in $\mathbb{R}^3$. The face-to-face property of the mesh is implicit because a simplicial complex requires that any two simplices are either disjoint or meet in a common triangle, edge, or vertex. We also require that every triangle, edge, and vertex in the mesh is face of a tetrahedron in the mesh.

A spatial domain is typically given in terms of its boundary constructed using a computer-aided design system. The *tetrahedral mesh generation problem* assumes the boundary is piecewise linear and asks for the construction of a tetrahedral mesh that covers the spatial domain with that boundary. The size and shape of the triangles and tetrahedra is important because it relates to the convergence and stability of numerical methods such as the finite element analysis, see Strang and Fix [20].

Probably the most common tetrahedral meshes are Delaunay triangulations, which are named after Boris Delaunay [7] and also known as duals to Voronoi tilings, which are named after Georges Voronoi [22]. They are supported by fast algorithms both for the construction and for the maintenance under local changes. In this paper we make essential use of a somewhat larger class of tetrahedral meshes referred to as weighted Delaunay triangulations. This class has been studied extensively in the geometry literature where its meshes are known as regular triangulations [3] and also as coherent triangulations [12]. The fast algorithms for Delaunay triangulations extend with minor modification to the larger class of weighted Delaunay triangulations [9].

Previous work. The generation of meshes with well-shaped triangles in $\mathbb{R}^2$ is reasonably well understood. Bern, Eppstein and Gilbert prove that quad-tree decompositions can be used to generate meshes free of badly shaped triangles that adapt to the local density of input specifications [2]. Ruppert proves the same for his version of the Delaunay refinement method [18]. Experimental studies suggest the latter method adapts better to input specifications and outperforms the quad-tree approach with smaller and smoother meshes.

The generation of meshes of well-shaped tetrahedra in $\mathbb{R}^3$ seems considerably more difficult. The extension of the quad-tree and the Delaunay refinement methods to $\mathbb{R}^3$ both encounter significant difficulties. Mitchell and Vavasis [16] use oct-trees to generate a tetrahedral mesh for the volume bounded by a specified polyhedral boundary. Dey, Bajaj and Sugihara [8] and Shewchuck [19] extend the Delaunay refinement algorithm to $\mathbb{R}^3$ but fail to address the problem of slivers.

The disturbing presence of slivers in 3-dimensional Delaunay triangulations has been recognized at least since the experimental study of Cavendish, Field and Frey [4]. Talmor [21] notes that even well-spaced vertices do not prevent slivers. Chew [5] sketches an algorithm that eliminates slivers by adding points in a randomized manner.

Results. The main result of this paper is a method that eliminates slivers without adding any new point and without moving any point of the given set. Intuitively, the method applies physical pressure and squeezes the Delaunay triangulation. Most slivers give way to the pressure and disappear. The remaining slivers migrate to the boundary where they can be peeled off. Alternatively, they can be treated with boundary enforcement heuristics described in the mesh generation literature. We suppress the distraction of boundary effects by considering periodic point sets $S \in \mathbb{R}^3$. In other words, we choose a finite set $S_0 \subseteq [0,1)^3$ and duplicate it within each integer unit cube: $S = S_0 + \mathbb{Z}^3$, where $\mathbb{Z}^3$ is the three-dimensional integer grid. The sliver elimination method assumes the points are distributed so each tetrahedron in the Delaunay triangulation has the ratio property introduced by Miller et al. [15]: the radius of the circumsphere is bounded from above by a constant times the length of the shortest edge. If necessary the ratio property can be generated by adding points at circumcenters of violating Delaunay tetrahedra.

We show that under the assumption of the ratio property we can assign small real weights to the points so the weighted Delaunay triangulation contains no sliver. We refer to this result as the Sliver Theorem. Another way to think of the result is that the ratio property for the Delaunay triangulation implies the existence of a sliver-free triangulation of the same set of points. Since the ratio property prevents all other types of undesirable elements, our result implies a triangulation free of any badly shaped tetrahedron. This complements a result of Talmor [21] that a triangulation without badly shaped tetrahedron implies the ratio property for the Delaunay triangulation of the same set of points. In other words, for a periodic set $S \subseteq \mathbb{R}^3$ the ratio property for its Delaunay triangulation is equivalent to the existence of a triangulation without any badly shaped tetrahedron.

Since the sliver-free triangulation is a weighted Delaunay triangulation it can be obtained from the unweighted Delaunay triangulation by a sequence of flips. The algorithm in this paper is thus similar to but also different from Joe's heuristic, which improves tetrahedral shape by flipping [13]. Joe's heuristic is greedy and halts the improvement of a vertex neighborhood at a local optimum. The algorithm in this paper improves a vertex neighborhood by following a more global optimization strategy.

Outline. Section 2 discusses the shape of triangles and tetrahedra. Section 3 introduces Delaunay triangulations for unweighted and for weighted points. Sections 4 and 5 prove geometric results needed in the proof of the Sliver Theorem, which is presented in Section 6. Section 7 turns this theorem into a sliver removing algorithm. Section 8 concludes the paper.

2 Tetrahedral Shape

A triangle or tetrahedron is badly shaped if it has at least one small angle. Some badly shaped tetrahedra have badly shaped bounding triangles, but there are also tetrahedra with small angles none of whose four triangles is badly shaped. This section explains what exactly we mean by good and bad shape and how we talk about it.

Shape measures. The mesh generation literature is rich in measures of simplex quality. A common term is the *aspect ratio*, which is often but not always defined as the radius of the smallest containing sphere over the radius of the largest contained sphere. Related is the *measure of degeneracy* defined as the length of the longest edge over the radius of the largest contained sphere. The latter is motivated by the finite element convergence analysis of Ciarlet [6]. Liu and Joe [14] consider several other measures for tetrahedra and study how they relate to each other. In this paper we use distance, radius, angle and volume to express the quality of triangles and tetrahedra.

To simplify discussions we use fuzzy language for size descriptors. In each case we suppose the existence of a small constant, $\varepsilon > 0$, which can be used to make the statement precise. For example, an angle φ is *small* if $\varphi < \varepsilon$ and *large* if $\varphi > \pi - \varepsilon$. An aspect ratio is *large* if it exceeds $1/\varepsilon$. We also use fuzzy descriptors in a relative sense. For example, the edge pq of a triangle pqr is *short* if $\|p-q\| < \varepsilon \cdot \max\{\|p-r\|, \|q-r\|\}$. Similar relative conventions are adopted for points that are *close* to each other or to a line or plane, etc.

Badly shaped triangles. A triangle with large aspect ratio has at least one small angle and all three vertices close to a line. There are two types: a *dagger* with one short edge and a *blade* with no short edge, see Figure 1.

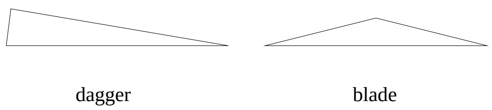

Figure 1: The dagger has one short edge and at least one small angle. The blade has no short edge and therefore one large and two small angles.

Badly shaped tetrahedra. Among the tetrahedra with large aspect ratio we distinguish the ones with at least three badly shaped triangles from the others. A tetrahedron of the former type has four vertices close to a line. The points can be close or far in the direction along this line, and we distinguish the cases 3-1 (*spire*), 1-2-1 (*spear*), 1-1-1-1 (*spindle*), 2-1-1 (*spike*), 2-2 (*splinter*), see Figure 2. The spire has a cycle of three short edges and therefore a cycle of three daggers among its four triangles. The *splinter* has two opposite short edges and therefore four daggers, two in each direction. The spear and the spike both have one short edge and therefore two daggers and two blades as triangles. The spindle has no short edge and therefore four blades as triangles.

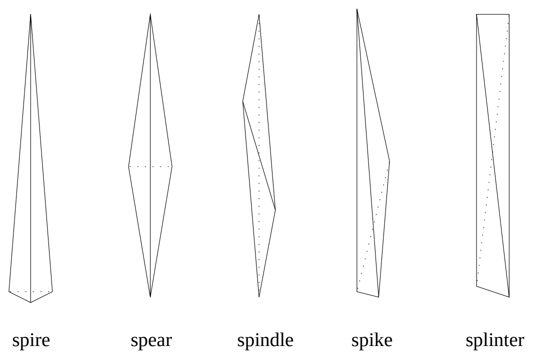

Figure 2: From left to right the number of daggers among the four triangles is at least three for the spire, two for the spear, zero for the spindle, two for the spike, and four for the splinter.

A tetrahedron whose vertices are not close to a line has a large aspect ratio if its vertices are close to a plane, and we distinguish the cases where two points are close to each other (*wedge*), where three points are close to a line (*spade*), where the projection to the plane is a triangle with a point inside (*cap*), and where the projection is a quadrilateral (*sliver*), see Figure 3.

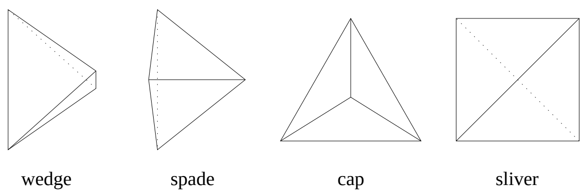

Figure 3: From left to right the number of long edges with small dihedral angle is one for the wedge, two for the spade, three for the cap, and four for the sliver.

A similar but different classification of badly shaped tetrahedra can be found in Bern et. al [1]. Their classification is based on dihedral angles while ours primarily considers face angles.

Radius-edge ratio. Let $pqrs$ be a tetrahedron, X the radius of its circumsphere, and L the length of its shortest edge, see Figure 4. $pqrs$ has *Ratio Property* $[\varrho_0]$ for a constant ϱ_0 if $X/L \leq \varrho_0$. If a tetrahedron has Ratio Property $[\varrho_0]$ then so do all its triangles. A triangulation has *Ratio Property* $[\varrho_0]$ if all its tetrahedra have it.

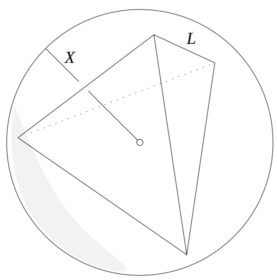

Figure 4: The vertices of the tetrahedron lie on the circumsphere with radius X. The length of the shortest edge is L.

The ratio attains its minimum for the regular tetrahedron where $X/L = \sqrt{6}/4 \approx 0.612$. Space cannot be tiled with copies of the regular tetrahedron alone so triangulations require a larger value of ϱ_0. Ratio Property $[\varrho_0]$ eliminates all badly shaped triangles and all badly shaped tetrahedra other than the slivers. If $\varrho_0 < 1/\sqrt{2} \approx 0.707$ then all face angles are acute so that even slivers cannot exist. However, the ratio property for such a small constant can usually not be obtained and we need a different method to eliminate slivers.

3 Delaunay Triangulations

The proof of the Sliver Theorem uses weighted Delaunay triangulations in an essential manner. This section introduces Delaunay triangulations, weighted points, and weighted Delaunay triangulations.

Delaunay triangulations. Let S be a discrete set of points in $\mathbb{R}^3$. We permit infinite sets but they must be locally finite. For simplicity assume that S is in general position. In particular, for every four points in S there is a sphere that passes through them and for any five points there is no such sphere. A sphere is *empty* if it encloses no point of S, or equivalently, if all points lie either on or outside the sphere. The convex hull of points $p, q, r, s \in S$ is a tetrahedron, denoted as $pqrs$, and a *Delaunay tetrahedron* if the circumsphere is empty. The *Delaunay triangulation* of S, denoted as $\mathrm{Del}\, S$, is the 3-complex consisting of all Delaunay tetrahedra and their triangles, edges, and vertices.

Delaunay triangulations are popular meshes for several reasons. If S is in general position then $\mathrm{Del}\, S$ is unique and can be efficiently constructed [4, 9]. The changes caused by deleting or inserting a point are typically local. $\mathrm{Del}\, S$ contains all edges of a minimum spanning tree, and for each $p \in S$ it contains the edge to the closest point. Delaunay triangulations are optimal with respect to smallest containing spheres of tetrahedra, see [17].

Given a Delaunay triangulation we can generate Ratio Property $[\varrho_0]$ by adding points at circumcenters of violating tetrahedra, see e.g. [19]. If $\varrho_0 \geq 1$ then the minimum distance between a new point and any of the old points is at least the minimum distance between any two of the old points. In the periodic interpretation of $\mathbb{R}^3$ we have a finite amount of volume, which permits only finitely many points if the interpoint distances are bounded by a fixed positive lower bound. The method can therefore not add infinitely many points and halts after a finite amount of time.

Weighted points and distance. A *weighted point*, $\hat{p} = (p, P^2) \in \mathbb{R}^3 \times \mathbb{R}$, is interpreted as a sphere or ball with center p and radius P, see Figure 5. The *weighted distance* between $\hat{p}$ and $\hat{z} = (z, Z^2)$ is

$$\|\hat{p} - \hat{z}\| = \sqrt{\|p - z\|^2 - P^2 - Z^2}.$$

$\hat{p}$ and $\hat{z}$ are *orthogonal* if the weighted distance vanishes: $\hat{p} \perp \hat{z}$ if $\|\hat{p} - \hat{z}\| = 0$. Any four spheres in $\mathbb{R}^3$ have a common orthogonal sphere, called the *orthosphere*. For example, if the four spheres are points then the orthosphere is the unique circumsphere of the tetrahedron they define. Unless the four centers lie in a common plane, the orthosphere is unique and has finite radius.

The corresponding observation one dimension lower is that any three circles in $\mathbb{R}^2$ have a common orthogonal circle, called the *orthocircle*, see Figure 5. Unless the three centers are collinear, the orthocircle has finite radius.

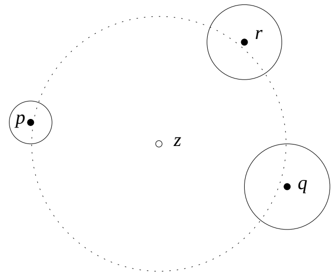

Figure 5: The dotted circle is orthogonal to the three solid circles. Since the radii of the solid circles are positive, their centers all lie outside the dotted circle.

Weighted Delaunay triangulations. A weighted generalization of Delaunay triangulations is obtained by substituting spheres for points and orthospheres for circumspheres. To be specific, let $\mathsf{w} : S \to \mathbb{R}$ be a weight assignment and consider the defined set of spheres:

$$\hat{S} = \{(p, P^2) \mid p \in S \text{ and } P^2 = \mathsf{w}(p)\}.$$

Assume $\hat{S}$ is in general position, which among other things implies that every four spheres have a common orthogonal sphere and no five spheres have one. A sphere $\hat{z}$ is *empty* if $\|\hat{z} - \hat{p}\| \geq 0$ for every $\hat{p} \in \hat{S}$. The convex hull of four sphere centers is a tetrahedron and a *weighted Delaunay tetrahedron* if the common orthosphere is empty. The *weighted Delaunay triangulation* of $\hat{S}$, denoted as $\mathrm{Del}\, \hat{S}$, is the 3-complex consisting of all weighted Delaunay tetrahedra and their triangles, edges, and vertices. If all radii are zero then the weighted Delaunay triangulation of the spheres is the same as the Delaunay triangulation of the centers.

The center p of a sphere $\hat{p} \in \hat{S}$ may or may not belong to the weighted Delaunay triangulation. Specifically, p is a vertex in $\mathrm{Del}\, \hat{S}$ iff there is a sphere not necessarily in $\hat{S}$ that is orthogonal to $\hat{p}$ and has positive weighted distance to all other spheres in $\hat{S}$. In this paper we choose weights in a way that guarantees the existence of such spheres. It follows that the set of centers is also the set of vertices.

Cross-sections. Orthogonality is inherited from spheres to circles if the slicing plane passes through at least one of the two centers.

CLAIM 1. If $(p, P^2) \perp (z, Z^2)$ then any plane through p intersects the two spheres in two orthogonal circles.

PROOF. Let (u, U^2) and (v, V^2) be the circles of intersection between the plane and the two spheres. We have $u = p$, $U^2 = P^2$ and $V^2 = Z^2 - \|z - v\|^2$. Then

$$\begin{aligned} \|u - v\|^2 &= \|u - z\|^2 - \|z - v\|^2 \\ &= P^2 + Z^2 - \|z - v\|^2 \\ &= U^2 + V^2, \end{aligned}$$

which shows that (u, U^2) and (v, V^2) are orthogonal. □

Claim 1 can be extended to dimensions different from 3. Consider for example the two-dimensional case. If (p, P^2) and (z, Z^2) are two orthogonal circles then any line passing through p intersects them in two orthogonal intervals. Given two intervals there is a unique third interval orthogonal to both. It follows that all circles that are orthogonal to two circles (p, P^2) and (q, Q^2) intersect the edge from p to q in the same two points.

4 Linear Relations

This section proves a number of relations between distances, weighted distances, radii, and areas needed for the proof of the Sliver Theorem in Section 6. We begin by introducing notation that simplifies computations and discussions.

Relation. Two quantities X and Y are said to be *linearly related*, denoted as $X \sim Y$, if there are constants $c, C > 0$ with $c \cdot X \leq Y \leq C \cdot X$. Note that $\sim$ satisfies

$$\begin{aligned} X \sim Y &\implies Y \sim X, \\ (X \sim Y) \wedge (Y \sim Z) &\implies X \sim Z, \end{aligned}$$

but it is not an equivalence relation because the constants deteriorate in the repeated application of the second rule: if c', C' are the constant for $Y \sim Z$ then $c \cdot c', C \cdot C'$ are the constants for $X \sim Z$. The relation combines well with arithmetic operations on positive quantities:

$$(X \sim Y) \wedge (U \sim V) \implies \begin{cases} X + U \sim Y + V, \\ X \cdot U \sim Y \cdot V, \\ X/U \sim Y/V. \end{cases}$$

If c'', C'' are the constants for $U \sim V$ then $\min\{c, c''\}$, $\max\{C, C''\}$ are the constants for the sums, $c \cdot c'', C \cdot C''$ for the products, and $c/C'', C/c''$ for the ratios. In this paper we obtain new linear relations from constant length chains of old linear relations.

Weight property. We suppose that the radii of the spheres are not large relative to the distances between their centers. To make this precise we say a pair of spheres $\hat{p} = (p, P^2)$, $\hat{q} = (q, Q^2)$ has *Weight Property* $[\omega_0]$ for a constant $\omega_0 \in (0, 1/2)$ if $0 \leq P, Q \leq \omega_0 \|p - q\|$. A set of spheres has *Weight Property* $[\omega_0]$ if every pair has it. The upper bound on the radii implies the spheres are pairwise disjoint. It also implies that the weighted distance between two spheres is not very different from the Euclidean distance between the two centers:

CLAIM 2. If a pair of spheres $\hat{p}, \hat{q}$ has Weight Property $[\omega_0]$ then $\|p - q\| \sim \|\hat{p} - \hat{q}\|$.

PROOF. We establish $c_2 \cdot \|p - q\| \leq \|\hat{p} - \hat{q}\| \leq C_2 \cdot \|p - q\|$ for constants $c_2 = \sqrt{1 - 2\omega_0^2}$ and $C_2 = 1$. By definition we have $\|\hat{p} - \hat{q}\|^2 = \|p - q\|^2 - P^2 - Q^2$ and $P^2, Q^2 \geq 0$ implies $\|\hat{p} - \hat{q}\| \leq \|p - q\|$. We get the lower bound from $P^2, Q^2 \leq {\omega_0}^2 \|p - q\|^2$, which implies $\sqrt{1 - 2{\omega_0}^2} \cdot \|p - q\| \leq \|\hat{p} - \hat{q}\|$. □

Area and radius. Let pqr be a triangle and X the radius of the circumcircle. X is no smaller than half the length of the longest edge, and if pqr has Ratio Property $[\varrho_0]$ then X is also not much larger than that. This implies that X^2 is not much different from the area of the triangle, which we denoted as $|pqr|$:

CLAIM 3. If pqr has Ratio Property $[\varrho_0]$ then $X^2 \sim |pqr|$.

PROOF. We establish $c_3 \cdot X^2 \leq |pqr| \leq C_3 \cdot X^2$ with

$$c_3 = \min\{\frac{\sqrt{3}}{4{\varrho_0}^2}, \frac{1}{2{\varrho_0}^2} \sin(2 \arccos \frac{1}{2\varrho_0})\}$$

and $C_3 = \pi$. The upper bound is clear because pqr is enclosed by the circumcircle with radius X. For the lower bound we assume pr is the longest of the three edges and we use $\|p - q\|, \|q - r\| \geq X/\varrho_0$ as implied by Ratio Property $[\varrho_0]$. The line from q to the circumcenter, x, cuts the angle ψ at q into two, one inside pqx and the other inside xqr. The cosine of the angle inside pqx is

$$\cos \psi_p = \frac{\|p - q\|}{2X} \geq \frac{1}{2\varrho_0}.$$

pqx is an isosceles triangle with equal angles at p and q. We therefore take the arccosine in the interval from the empty to the right angle and get

$$\psi_p \leq \arccos \frac{1}{2\varrho_0} < \frac{\pi}{2}.$$

The same bounds hold for ψ_r inside xqr. Because pr is the longest side, $\psi = \psi_p + \psi_r$ is at least $\frac{\pi}{3}$ and at most twice the upper bound for ψ_p. The sine of ψ is therefore at least the smaller sine of the two bounds: $\sin \psi \geq \min\{\sin \frac{\pi}{3}, \sin(2 \arccos \frac{1}{2\varrho_0})\}$. The area of pqr is half the base length times the height:

$$\begin{aligned} |pqr| &= \frac{1}{2} \cdot \|p - q\| \cdot \|q - r\| \cdot \sin \psi \\ &\geq \frac{X^2}{2{\varrho_0}^2} \cdot \min\{\frac{\sqrt{3}}{2}, \sin(2 \arccos \frac{1}{2\varrho_0})\}, \end{aligned}$$

which implies the lower bound. $\square$

Note that $\sin(2 \arccos \frac{1}{2\varrho_0})$ is larger than, equal to, less than $\sqrt{3}/2$ if ϱ_0 is less than, equal to, larger than 1.

Radius and radius. Ratio Property $[\varrho_0]$ and Weight Property $[\omega_0]$ together imply that for a triangle the radii of the circumcircle and the orthocircle are not very different. Let $\hat{p}, \hat{q}, \hat{r}$ be three spheres that define an orthocircle with radius Z and whose centers define a circumcircle with radius X.

Claim 4. If $\hat{p}, \hat{q}, \hat{r}$ have Weight Property $[\omega_0]$ and pqr has Ratio Property $[\varrho_0]$ then $X \sim Z$.

Proof. We establish $c_4 \cdot X \leq Z \leq C_4 \cdot X$ with $c_4 = \sqrt{1 - 4{\omega_0}^2}$ and $C_4 = 1 + 2\varrho_0{\omega_0}^2$. The minimum weighted distance of the circumcenter, x, from the three weighted points is a lower bound for the radius of the orthocircle. To bound that minimum note that $2X$ is an upper bound on the length of each edge and ${\omega_0}^2(2X)^2$ is an upper bound on the weight of each point:

$$\begin{aligned} Z^2 &\geq \min\{\|x - \hat{p}\|^2, \|x - \hat{q}\|^2, \|x - \hat{r}\|^2\} \\ &\geq X^2 - 4{\omega_0}^2 X^2. \end{aligned}$$

To obtain an upper bound consider the perpendicular bisectors k and ℓ of edges pq and qr, which intersect at x, see Figure 6. Let $\hat{k}$ be the line of points

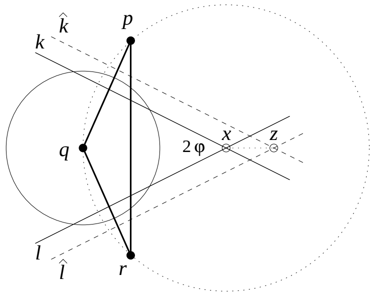

Figure 6: To avoid a tiny parallelogram we draw the strips wider and the circle around q larger than allowed by Weight Property $[\omega_0]$.

with equal weighted distance from $\hat{p}$ and $\hat{q}$, and let $\hat{\ell}$ be the same line for $\hat{q}$ and $\hat{r}$. The width of the strip between k and $\hat{k}$ is a maximum only if the weights of $\hat{p}$ and $\hat{q}$ are as different as possible, for example $P^2 = 0$ and $Q^2 = {\omega_0}^2\|p - q\|^2$. In this case the width is $W = {\omega_0}^2\|p - q\|/2 \leq {\omega_0}^2 X$. The same upper bound holds for the width of the strip between ℓ and $\hat{\ell}$. The centers x of the circumcircle and z of the orthocircle are diagonally opposite vertices of the parallelogram formed by $k, \hat{k}, \ell, \hat{\ell}$, see Figure 6. Let the angle at x inside the parallelogram be 2φ. The edge xz cuts this angle into two and we assume that the angle between xz and k inside the quadrilateral is $\xi \geq \varphi$. The distance between the two centers is therefore

$$\|x - z\| = \frac{W}{\sin \xi} \leq \frac{{\omega_0}^2 \cdot X}{\sin \varphi}.$$

By Ratio Property $[\varrho_0]$, we have $\|p - q\|, \|q - r\| \geq X/\varrho_0$. Hence $X \cdot \sin \varphi \geq X/2\varrho_0$ and therefore $\sin \varphi \geq 1/2\varrho_0$. The radius of the orthocircle is bounded from above by the radius of the circumcircle plus the distance between the centers:

$$Z \leq X + \|x - z\| \leq X + 2\varrho_0{\omega_0}^2 \cdot X,$$

which is the upper bound claimed at the beginning of the proof. $\square$

Fortunately, Claim 4 does not extend to tetrahedra where it fails for slivers with four almost cocircular vertices.

Parametrizing slivers. Let $pqrs$ be a tetrahedron, V the volume, and L the length of the shortest edge. We define $\sigma = \sigma(pqrs) = V/L^3$ and use it as a measure of quality. Assuming Ratio Property $[\varrho_0]$ we call $pqrs$ a *sliver* if σ is less than some threshold $\sigma_0 > 0$ to be specified later. It is useful to relate this measure with a distance-radius ratio. Let D be the Euclidean distance of point p from the plane passing through qrs and let Y be the radius of the circumcircle of qrs, see Figure 7.

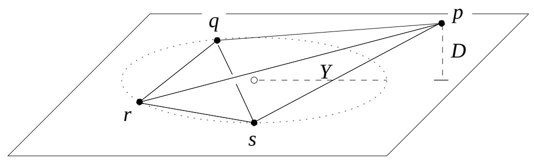

Figure 7: D/Y can be defined for each ordering of the four vertices, but all four ratios are linearly related to σ.

Claim 5. If a tetrahedron $pqrs$ has Ratio Property $[\varrho_0]$ then $D/Y \sim \sigma$.

Proof. We establish $c_5 \cdot D/Y \leq \sigma \leq C_5 \cdot D/Y$ for $c_5 = c_3/24$ and $C_5 = \pi{\varrho_0}^3/3$. The triangle qrs has Ratio Property $[\varrho_0]$, so $Y \sim L$ with constants $c = 1/\varrho_0$ and $C = 2$. By Claim 3 we have $Y^2 \sim |qrs|$ with constants c_3, C_3. The volume of the tetrahedron is $|qrs| \cdot D/3$. Therefore

$$\frac{D}{Y} = \frac{Y^2 \cdot D}{Y^3} \sim \frac{|qrs| \cdot D}{3 \cdot L^3} = \sigma.$$

Following the rules for combining linear relations we get constants $c_5 = c_3/3C^3$ and $C_5 = C_3/3c^3$. $\square$

5 Length and Degree Bounds

This section recalls a result of Talmor [21] which is then used to extend results of Miller et al. [15] from unweighted to weighted Delaunay triangulations.

Weighted ratio property. Let S be a periodic set of points in $\mathbb{R}^3$. In other words, $S = S_0 + \mathbb{Z}^3$ where $S_0 \subseteq [0,1)^3$ is finite and $\mathbb{Z}^3$ is the three-dimensional integer grid. For a point $x \in \mathbb{R}^3$ let $N(x)$ be the distance to the second closest point in S. If $x \in S$ then x itself is closest and $N(x)$ is the distance to the closest point in $S - \{x\}$. The following result is Theorem 3.6.2 in Talmor's thesis [21]:

Claim 6. Assume $\mathrm{Del}\, S$ has Ratio Property $[\varrho_0]$. Then there is a constant c_T depending only on ϱ_0 such that $N(x) \geq c_T \cdot N(z)$ for every empty sphere (z, Z^2) that passes through x.

We extracted $c_T = 64 \cdot {\varrho_0}^2 M^c$ from Talmor's thesis, where $M = \max\{2\varrho_0/\sin\tau, 4{\varrho_0}^2\}$, $c = 2/(1-\cos\frac{\tau}{4})$, $\tau = (\arcsin\frac{1}{2\varrho_0})/2$. We use Claim 6 to derive a property for weighted Delaunay triangulations reminiscent of the ratio property for Delaunay triangulations. A periodic sphere set is defined by a weight assignment $\mathsf{w} : S_0 \to \mathbb{R}$. As usual $\hat{z} = (z, Z^2)$ denotes the orthosphere and L the length of the shortest edge of a tetrahedron.

Claim 7. Assume $\mathrm{Del}\, S$ has Ratio Property $[\varrho_0]$ and $\hat{S}$ has Weight Property $[\omega_0]$. Then there exists a constant ϱ_1 depending only on ϱ_0 and ω_0 such that $Z/L \leq \varrho_1$ for every tetrahedron in $\mathrm{Del}\, \hat{S}$.

Proof. We establish the bound for the constant $\varrho_1 = (1+\omega_0)/c_T$. Let $\hat{z}$ and L be orthosphere and shortest edge length for a tetrahedron $pqrs \in \mathrm{Del}\, \hat{S}$. Assume $L = \|p-q\|$, which implies $N(p) \leq L$. Because all points of S lie on or outside $\hat{z}$ we have $N(z) \geq Z$. Let x be a point on the intersection of the two orthogonal spheres $\hat{p}$ and $\hat{z}$. By Weight Property $[\omega_0]$ we have $\|x-p\| \leq \omega_0 \cdot N(p) \leq \omega_0 \cdot L$. Therefore $N(x) \leq \|x-q\| \leq \|x-p\| + \|p-q\| \leq (1+\omega_0) \cdot L$. Claim 6 implies

$$Z \;\leq\; N(z) \;\leq\; \frac{N(x)}{c_T} \;\leq\; \frac{1+\omega_0}{c_T} \cdot L,$$

as stated. □

Edge-length variation. For a graph G with vertices and straight edges in $\mathbb{R}^3$ we are interested in comparing the length of edges. Specifically, we define the *length variation* at a vertex $p \in G$ as

$$\nu(p, G) \;=\; \max\{\|p-q\|/\|p-u\|\},$$

where the maximum is taken over all edges pq, pu in G. Our first result shows that triangles in the weighted Delaunay triangulation inherit a constant upper bound on the variation of their edge lengths from the Delaunay triangulation.

Claim 8. If $\hat{S}$ has Weight Property $[\omega_0]$, $\mathrm{Del}\, S$ has Ratio Property $[\varrho_0]$, and $pqr \in \mathrm{Del}\, \hat{S}$ then $\|p-q\| \sim \|p-r\|$.

Proof. We establish $c_8 \cdot \|p-q\| \leq \|p-r\| \leq C_8 \cdot \|p-q\|$ for $c_8 = \sqrt{1-4{\omega_0}^2}/2\varrho_1$ and $C_8 = 1/c_8$. The length of an edge is at most twice the radius of the circumcircle, X, and by Claim 4 that radius is linearly related to the radius of the orthocircle:

$$\|p-q\| \;\leq\; 2X \;\leq\; \frac{2Z}{\sqrt{1-4{\omega_0}^2}}.$$

By Claim 7 the length of pq is $\|p-q\| \geq Z/\varrho_1$. The same bounds hold for $\|p-r\|$ which implies the claimed linear relation. □

Edges forming small angles. If two edges pq and pu share a common endpoint we denote the angle at that endpoint as $\angle qpu$. All angles between edges are measured between 0 and π. We show that a small angle implies about equal length, and this is even true if the two edges arise in two different weighted Delaunay triangulations:

Claim 9. Assume $\mathrm{Del}\, S$ has Ratio Property $[\varrho_0]$, $\hat{S}_1$ and $\hat{S}_2$ have Weight Property $[\omega_0]$, and $pq \in \mathrm{Del}\, \hat{S}_1$, $pu \in \mathrm{Del}\, \hat{S}_2$. Then there is a constant $\eta_0 > 0$ such that $\angle qpu < \eta_0$ implies $\|p-q\| \sim \|p-u\|$.

Proof. We establish the implication for the constant

$$\eta_0 \;=\; \frac{1}{2} \cdot \arctan \frac{{\varrho_1}^2 - \sqrt{{\varrho_1}^2 + {\omega_0}^2 - 1/4}}{\sqrt{{\omega_0}^2 + 1/4}}$$

and $c_9 \cdot \|p-q\| \leq \|p-u\| \leq C_9 \cdot \|p-q\|$ for $c_9 = (1-\omega_0)/2$ and $C_9 = 1/c_9$. Let $\hat{z} = (z, Z^2)$ be the orthosphere of a tetrahedron that contains pq as one of its edges. We cut $\hat{z}$ with the plane that passes through p, q, u and let $\hat{y} = (y, Y^2)$ be the circles of intersection, see Figure 8. Let k be the line passing through p and tangent to $\hat{y}$. By Claim 1, $\hat{y}$ is orthogonal to the circles at which the plane intersects $\hat{p}$ and $\hat{q}$. All circles orthogonal to $\hat{p}$ and to $\hat{q}$ meet pq in the same two points, see the comment after Claim 1. The distance between these two points is twice the radius of the smallest orthocircle. By Weight Property $[\omega_0]$ that radius is

$$V \;\geq\; \sqrt{\frac{1}{4} - {\omega_0}^2} \cdot \|p-q\|.$$

Let a be the point on the circle $\hat{y}$ so that ya intersects pq in a right angle, see Figure 8. The intersection point

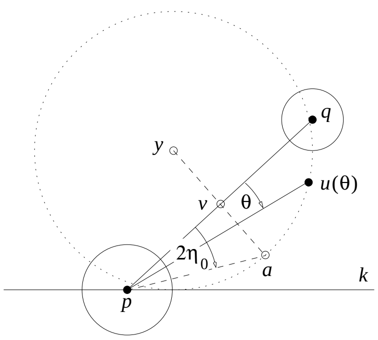

Figure 8: The dotted circle is orthogonal to the two solid ones. Edges with small angle θ cannot be short.

is the center v of the smallest circle orthogonal to $\hat{p}$ and $\hat{q}$. By Claim 7 we have $Y \leq \varrho_1 \cdot \|p-q\|$. The normal distance of a from pq is $A = Y - \sqrt{Y^2 - V^2}$ which assumes its minimum when Y is as large and V is as small as possible. Therefore

$$\begin{aligned} A &\geq \varrho_1\|p-q\| - \sqrt{\varrho_1{}^2\|p-q\|^2 - \frac{1-4\omega_0{}^2}{4}\|p-q\|^2} \\ &= (\varrho_1 - \sqrt{\varrho_1{}^2 + \omega_0{}^2 - 1/4}) \cdot \|p-q\|. \end{aligned}$$

The angle at p between pq and line k is at least the angle between pq and pa, which is

$$\begin{aligned} \angle qpa &= \arctan \frac{A}{\sqrt{P^2+V^2}} \\ &\geq \arctan \frac{\varrho_1 - \sqrt{\varrho_1{}^2 + \omega_0{}^2 - 1/4}}{\sqrt{\omega_0{}^2 + 1/4}} \end{aligned}$$

because $P^2 \leq \omega_0{}^2 \cdot \|p-q\|^2$ and $V^2 \leq \|p-q\|^2/4$. Note that the latter expression is twice the constant η_0.

For each angle $0 \leq \theta \leq 2\eta_0$, let $k(\theta)$ be the line passing through p that forms an angle θ with pq, see Figure 8. $k(\theta)$ intersects $\hat{y}$ in two points and we let $u(\theta)$ be the point further from p. Finally, we define $f(\theta) = \|p - u(\theta)\|$. We have $f(0) \geq (1-\omega_0) \cdot \|p-q\|$ and $f(2\eta_0) \geq 0$. Since f is a concave function it follows that $f(\theta)$ is at least $(1-\omega_0)/2$ times the length of pq for all $\theta \leq \eta_0$. By assumption there is an angle $\theta \leq \eta_0$ so $u(\theta)$ lies on the edge from p to u. The lower bound of the linear relation follows:

$$\|p-u\| \;\geq\; f(\theta) \;\geq\; \frac{1-\omega_0}{2} \cdot \|p-q\|.$$

The upper bound follows by a symmetric argument that exchanges q and u. ◻

Weighted Delaunay edges. Miller et al. prove that if $\mathrm{Del}\, S$ has Ratio Property $[\varrho_0]$ then there is a constant upper bound on the length variation at every vertex [15]. We use Claim 7 to prove the same is true for the graph of all possible weighted Delaunay edges. Let $K = K(S)$ be the union of all weighted Delaunay triangulations defined by weight assignments $\mathsf{w} : S \to \mathbb{R}$ whose sphere sets have Weight Property $[\omega_0]$. Since K contains all edges of the unweighted Delaunay triangulation it connects every point p to the closest point $q \in S$. Let G be the graph of all edges in K.

Claim 10. If $\mathrm{Del}\, S$ has Ratio Property $[\varrho_0]$ then there is a constant $\nu_0 > 0$ such that $\nu(p, G) \leq \nu_0$ for every vertex $p \in S$.

Proof. We establish the upper bound on length variation for the constant

$$\nu_0 \;=\; \left(\frac{2}{1-\omega_0}\right)^m \cdot \left(\frac{2\varrho_1}{\sqrt{1-4\omega_0{}^2}}\right)^{m-1},$$

where $m = 2/(1 - \cos\frac{\eta_0}{4})$. The argument is based on the two elementary geometry facts provided by Claims 8 and 9. Let Σ be the sphere of directions centered at p. We form a maximal packing of circular caps each with angle $\eta_0/4$. This means that if a is the center and b is a point on the boundary of a cap then $\angle apb = \eta_0/4$. The set of caps with the same centers and with radii $\eta_0/2$ covers Σ. Since the area of each cap in the first set is $(1-\cos\frac{\eta_0}{4})/2$ times the area of the sphere, the number of caps is at most some constant $m = 2/(1-\cos\frac{\eta_0}{4})$. The remainder of this proof uses the larger caps, which cover Σ.

For each edge $pq \in K$ let the point $q' \in \Sigma$ be the radial projection of q. Similarly, for each triangle $pqr \in K$ consider the arc on Σ that is the radial projection of qr. The points and arcs form a graph. Let pq be the longest and pu the shortest edge with endpoint p. We walk in the graph from q' to u'. This path leads from cap to cap and we just record the sequence of caps visited. If the path leaves a cap and returns to it later we ignore the detour and record the cap only once. In the end we have a sequence of at most m caps.

When we walk from point to point we track the length of the corresponding edges. As long as we stay within a single cap the length decreases at most by a factor of $(1-\omega_0)/2$, see Claim 9. If we step from one cap to the next the length decreases by at most a factor of $\sqrt{1-4\omega_0{}^2}/2\varrho_1$, see Claim 8. The number of caps is at most m so $\|p-u\| \geq \|p-q\|/\nu_0$. The claim follows because $\nu(p,G) = \|p-q\|/\|p-u\| \leq \nu_0$. ◻

All $\mathrm{Del}\, \hat{S}$ with Weight Property $[\omega_0]$ are subcomplexes of K. Claim 10 thus implies that for all such weighted Delaunay triangulations the length variation at every vertex is bounded from above by the constant ν_0.

Constant degree. The bound on the length variation in Claim 10 implies that each vertex belongs to at most a constant number of edges in K. A straightforward volume argument suffices to establish this fact.

CLAIM 11. If $\mathrm{Del}\, S$ has Ratio Property $[\varrho_0]$ then there is a constant δ_0 such that every vertex $p \in S$ belongs to at most δ_0 edges in K.

PROOF. We prove the claim for the constant $\delta_0 = (2{\nu_0}^2+1)^3$, where ν_0 is the constant in Claim 10. Let pq be the longest and pu the shortest edge with endpoint p. Assume without loss of generality that $\|p-u\| = 1$. Let r be a neighbor of p and s a neighbor of r. We have $\|p-r\| \geq 1$ by assumption $\|r-s\| \geq 1/\nu_0$ by Claim 10. For each neighbor r of p let Γ_r be the open ball with center r and radius $1/2\nu_0$. The balls are pairwise disjoint and fit inside the ball Γ with center p and radius $\|p-q\| + 1/2\nu_0$. The volume of Γ is

$$\begin{aligned} |\Gamma| &= \frac{4\pi}{3}(\|p-q\| + \frac{1}{2\nu_0})^3 \\ &\leq \frac{4\pi}{3}\left(\frac{2{\nu_0}^2+1}{2\nu_0}\right)^3 \\ &= (2{\nu_0}^2+1)^3 \cdot |\Gamma_r|. \end{aligned}$$

In words, at most $\delta_0 = (2{\nu_0}^2+1)^3$ neighbor balls fit into Γ. This implies that δ_0 is an upper bound on the number of neighbors of p. □

6 Sliver Theorem

Sections 4 and 5 provide the technical prerequisites for the proof of the main result of this paper, which is presented in this section.

Weight pumping idea. The main idea in the proof of the Sliver Theorem is to assign a weight P^2 to each point $p \in S$ so the weighted Delaunay triangulation is free of slivers. To get Weight Property $[\omega_0]$ we choose the weight of p in the interval $W(p) = [0, {\omega_0}^2 N^2(p)]$. Given a sliver $pqrs$ we use the pigeonhole principle to show there is a weight P^2 in $W(p)$ so $pqrs$ does not belong to the weighted Delaunay triangulation. While considering the tetrahedra around p we keep the weights of q, r, s unchanged and exclude the sliver $pqrs$ from the triangulation merely by manipulating the weight of p.

To be specific consider the weight interval $W(p)$ and for each sliver $pqrs$ mark the subinterval W_{qrs} of weights $P^2 \in W(p)$ for which $pqrs$ belongs to the weighted Delaunay triangulation, see Figure 9. We prove shortly that the length of the subintervals goes to zero as we lower the threshold for slivers. We also prove that there is only a constant number of subintervals to be considered. Hence we can choose a positive constant threshold

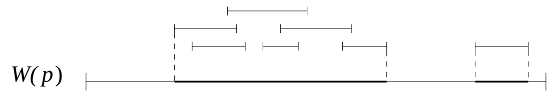

Figure 9: The subintervals cover all weights of p for which the weighted Delaunay triangulation contains a sliver.

small enough so $W(p)$ is not covered by the subintervals. Any weight $P^2 \in W(p)$ outside all subintervals will do.

Existence interval. For a tetrahedron $pqrs$ consider the orthocenter, z, as a function of the weight P^2 of p. We define $H(P)$ as the signed distance of z from the plane passing through qrs, see Figure 10. $H(P)$ is

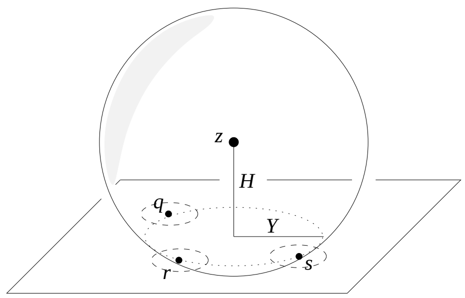

Figure 10: If the weights of q, r, s are fixed and the weight of p varies then the orthocenter of the tetrahedron moves on the normal line passing through the orthocenter of qrs.

positive if z and p lie on the same side and it is negative if they lie on different sides of the plane. As usual, L is the length of the shortest edge of $pqrs$.

CLAIM 12. If $\hat{p}, \hat{q}, \hat{r}, \hat{s}$ have Weight Property $[\omega_0]$ and $pqrs$ has Ratio Property $[\varrho_0]$ then there is a constant c_{12} such that $[-c_{12}L, c_{12}L]$ contains all values of $H(P)$ for which the orthosphere of $pqrs$ is empty.

PROOF. We establish $c_{12} = \sqrt{{\varrho_1}^2 + {\omega_0}^2 - 1/4}$. The square radius of the orthosphere is $Z^2 = H(P)^2 + Y^2$, which by Claim 7 is bounded from above by ${\varrho_1}^2 L^2$. The radius of the circumcircle of qrs is $X \geq L/2$, and by Claim 4 the radius of the orthocircle is $Y \geq \sqrt{1-4{\omega_0}^2} \cdot X$. Putting everything together we get

$$\begin{aligned} H(P)^2 &\leq {\varrho_1}^2 L^2 - (1-4{\omega_0}^2) \cdot X^2 \\ &\leq ({\varrho_1}^2 - \frac{1-4{\omega_0}^2}{4}) \cdot L^2 \end{aligned}$$

as claimed. □

Pumping motion. The bound on $H(P)^2$ is translated into a bound on the weight of p. For this purpose we look at the motion of the orthocenter. Specifically, we relate P^2 to the displacement of z along the line of its motion, which we denote as ℓ. As in Figure 7 the distance of p from the plane passing through qrs is denoted as D.

CLAIM 13. $H(P) = H(0) - \frac{P^2}{2D}$.

PROOF. Let E be the distance of p from ℓ. Then $Z^2 + P^2 = (H(P) - D)^2 + E^2$, see Figure 11. The

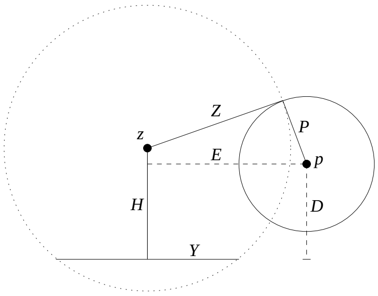

Figure 11: The orthocenter z moves down as the circle around p grows.

square radius of the orthosphere is $Z^2 = H(P)^2 + Y^2$ and therefore $H(P)^2 = (H(P) - D)^2 + E^2 - P^2 - Y^2$. After cancelling $H(P)^2$ we have

$$H(P) = \frac{D^2 + E^2 - Y^2}{2D} - \frac{P^2}{2D}.$$

The first term on the right side is $H(0)$ and the second is the displacement of z if we change the weight of p from 0 to P^2. ◻

Subinterval length. The goal is to show that the subintervals of $W(p)$ can be made as small as necessary. Recall the notation related to the parametrization of slivers introduced in Section 4.

CLAIM 14. The length of the subinterval defined by $pqrs$ is $|W_{qrs}| \leq c \cdot \sigma Y^2$.

PROOF. The constant is $c = 8c_{12}/c_5$. By Claim 12 the tetrahedron $pqrs$ belongs to the weighted Delaunay triangulation only if $-c_{12}L \leq H(P) \leq c_{12}L$. Using Claim 13 we get an interval for the weight of p:

$$2D \cdot (-c_{12}L + H(0)) \leq P^2 \leq 2D \cdot (c_{12}L + H(0)).$$

This interval contains W_{qrs}. To bound its length we use $L \leq 2Y$ and $D/Y \leq \sigma/c_5$ from Claim 5. Then

$$|W_{qrs}| = 4c_{12} \cdot DL \leq \frac{8c_{12}}{c_5} \cdot \sigma Y^2,$$

as claimed. ◻

Finale. We have now all pieces together to state and prove the main result of this paper, namely that there are weight assignments whose weighted Delaunay triangulations are free of slivers.

SLIVER THEOREM. Assume $\mathrm{Del}\, S$ has Ratio Property $[\varrho_0]$. Then there is a constant $\sigma_0 > 0$ and a weight assignment defining a set of spheres $\hat{S}$ with Weight Property $[\omega_0]$ such that $\sigma > \sigma_0$ for all tetrahedra $pqrs \in \mathrm{Del}\, \hat{S}$.

PROOF. We establish the result for the constant

$$\sigma_0 = \frac{c_5 \cdot (1 - 4{\omega_0}^2) \cdot {\omega_0}^2}{8c_{12} \cdot {\varrho_1}^2 {\nu_0}^2 \cdot (2{\nu_0}^2 + 1)^9}.$$

Let $p \in S$ and assume without loss of generality that the distance to its closest neighbor in K is 1. The length of its weight interval is therefore $|W(p)| = {\omega_0}^2$. Let $pqrs \in K$ be a sliver, that is, a tetrahedron with $\sigma < \sigma_0$. By Claim 14, it defines a subinterval of length $|W_{qrs}| \leq (8c_{12}/c_5) \cdot \sigma Y^2$. By Claim 10, the edges pq, pr, ps have length at most ν_0 each. By Claim 7, the radius of the orthosphere of $pqrs$ is $Z \leq \varrho_1 \cdot \nu_0$. The radius of the orthocircle of qrs is at most Z, and by Claim 4, the radius of the circumcircle is

$$Y \leq \frac{Z}{\sqrt{1 - 4{\omega_0}^2}} \leq \frac{\varrho_1 \cdot \nu_0}{\sqrt{1 - 4{\omega_0}^2}}.$$

The number of subintervals is at most the number of tetrahedra in K that share p. By Claim 11 there are at most δ_0 edges sharing p, so there are fewer than ${\delta_0}^3$ such tetrahedra. Let $I(p)$ be the part of $W(p)$ covered by subintervals defined by slivers. The total covered length is

$$\begin{aligned} |I(p)| &< \frac{8c_{12}}{c_5} \cdot \sigma_0 Y^2 \cdot {\delta_0}^3 \\ &\leq \frac{8c_{12} \cdot {\varrho_1}^2 {\nu_0}^2 \cdot (2{\nu_0}^2 + 1)^9}{c_5 \cdot (1 - 4{\omega_0}^2)} \cdot \sigma_0 \\ &\leq |W(p)|. \end{aligned}$$

We can therefore find a weight $P^2 \in W(p) - I(p)$. Every tetrahedron $pqrs \in K$ compatible with the weight P^2 of p has $\sigma \geq \sigma_0$. We repeat the argument for every point $p \in S$ and obtain a weight assignment that satisfies the claim. ◻

We remark that the above proof is not circular although at first sight it may appear that weight assignments for different points can interact in complicated ways. The next section discusses this issue in some detail.

7 Algorithm

This section develops two versions of an algorithm that eliminates slivers by weight assignment. The algorithm assumes the points are distributed so the Delaunay triangulation has Ratio Property $[\varrho_0]$. The first version is sequential, the second is parallel, and they both run in optimal time.

General strategy. The main step of the algorithm assigns a real weight P^2 to every point p in the given set S. This is done by processing the points in an arbitrary sequence. When processing point $p \in S$ we compute subintervals of $W(p)$ and we choose $P^2 \in W(p)$ outside all subintervals. The Sliver Theorem guarantees that for a properly chosen constant σ_0 such a weight exists. After processing all points the weighted Delaunay triangulation contains no tetrahedron with value of σ less than σ_0.

A critical issue is the apparent circularity of the algorithm: the change of the weight of q may alter some of the subintervals for p if pq is an edge in K. If p precedes q in the processing order then q may readmit tetrahedra around p that have been eliminated earlier by choice of P^2. There are two crucial observations that break the circularity. The first is that the quality measure for slivers is symmetric: $\sigma(pqrs) = \sigma(qprs)$. The second is that we only *increase* the weight of q so each newly admitted tetrahedron has q as a vertex. For its own sake, q chooses its weight to avoid all tetrahedra with $\sigma < \sigma_0$. As a consequence, any tetrahedron readmitted around p has value of σ at least σ_0.

The key step is the construction of subintervals of $W(p)$. Recall that each subinterval corresponds to a tetrahedron $pqrs$ with measure $\sigma < \sigma_0$. The Sliver Theorem suggests we use the constant σ_0 for which it proves the subintervals do not cover $W(p)$. In view of the miserably pessimistic estimate of σ_0 we follow a different strategy. Consider all tetrahedra with vertex p, and for each $pqrs$ consider the unbounded rectangle

$$R_{qrs} \quad = \quad W_{qrs} \times [\sigma(pqrs), +\infty)$$

in the $P^2 \times \sigma$ plane. The boundary of the union of rectangles forms the *skyline* over $W(p)$, see Figure 12. The best choice for P^2 is the weight coordinate of a highest point on the skyline.

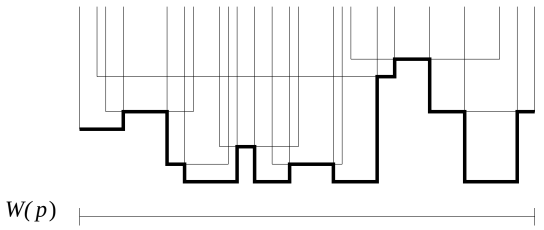

Figure 12: Each rectangle over $W(p)$ covers the values of σ around p that cannot be achieved if the weight is chosen in the P^2-interval of the rectangle.

Searching the skyline. The skyline is constructed by considering all tetrahedra in weighted Delaunay triangulations generated by varying the weight of p and keeping all other weights fixed. At the time the algorithm works on $p \in S$ it has already processed some of the points. Let $\mathsf{w} : S \to \mathbb{R}$ be the current weight assignment, and for every $P^2 \in W(p)$ let $\mathsf{w}_P : S \to \mathbb{R}$ be defined by

$$\mathsf{w}_P(u) \quad = \quad \begin{cases} \mathsf{w}(u) & \text{if } u \in S - \{p\}, \\ P^2 & \text{if } u = p. \end{cases}$$

Note that $\mathsf{w}_0 = \mathsf{w}$. Let $\hat{S}_P$ correspond to w_P. For every $\hat{S}_P$ we are only interested in the set $T_P \subseteq \operatorname{Del} \hat{S}_P$ of tetrahedra that have p as a vertex. The skyline is defined by the union of all T_P for $P^2 \in W(p)$. This union is computed by continuously increasing P^2 from 0 to its maximum, which is ${\omega_0}^2 N^2(p)$. The set T_P changes only at discrete moments:

$$0 = P_0^2 < P_1^2 < \ldots < P_{k-1}^2 < P_k^2 = {\omega_0}^2 N^2(p).$$

Define $T_i = T_P$ for P^2 between P_{i-1}^2 and P_i^2. In the non-degenerate case the step from T_{i-1} to T_i consists of a single flip operation [9, 17]. The sequence of flips is determined using a priority queue storing T_i. Each tetrahedron carries the time or weight when it is destroyed by the weight increase, and P_i^2 is the earliest such time of any tetrahedron in the priority queue. That tetrahedron is removed from the priority queue and from the weighted Delaunay triangulation, and new tetrahedra are inserted. By Claim 11 the union of all sets T_P has only constant cardinality so that processing the changing set costs only constant time in total. At any moment P^2 the minimum σ value of any tetrahedron in T_P is the height of the skyline above $P^2 \in W(p)$.

To summarize, we now have a sequential algorithm that takes time $\mathrm{O}(n \log n)$, where n is the number of points in S_0. Recall that $S = S_0 + \mathbb{Z}^3$ is periodic and so is the Delaunay triangulation. Any one of a number of published algorithms can be adapted to periodic sets in a way so it touches only points in S_0. The adaptation of the algorithm in [9] or in [11] takes time $\mathrm{O}(n \log n)$ because the size of $\operatorname{Del} S$ in a period is $\mathrm{O}(n)$ by Claim 11. After constructing $\operatorname{Del} S$ we assign weights in time $\mathrm{O}(n)$ as explained above. The construction of the corresponding weighted Delaunay triangulation is a side effect of the weight assignment step.

Parallel algorithm. A parallel version of the algorithm can be obtained by taking advantage of Claim 11, which asserts that vertices in K have constant size neighborhoods. Recall that K is the union of all weighted Delaunay triangulations, where the union is taken over all weight assignments $\mathsf{w} : S_0 \to \mathbb{R}$ whose sphere sets have Weight Property $[\omega_0]$. The *degree* of a vertex $p \in K$, denoted as $\delta(p)$, is the number of edges with endpoint p. By Claim 11, $\delta(p)$ is bounded from above by the constant δ_0. Hence K has a vertex coloring with $\delta_0 + 1$ colors. Two vertices of the same color share no tetrahedra in any of the weighted Delaunay triangulations so the weight assignment algorithm can be applied simultaneously. In other words, k processors can assign weights to the points of one color class in parallel and achieve optimal speed-up.

To summarize we now have a parallel algorithm that takes time $\mathrm{O}(n \log n/k)$ for $k = \mathrm{O}(n)$ processors. The first step constructs the graph of all edges in K using a randomized algorithm of Frieze, Miller and Teng [11]. As in the sequential case this algorithm needs to be adapted to periodic point sets, which is not difficult. The graph is then colored with a constant number of colors. The final two steps are the same as for the sequential algorithm except that the coloring is used to parallelize both the construction of the initial Delaunay triangulation and the weight assignment.

8 Discussion

This paper shows that if the Delaunay triangulation of a periodic point set in $\mathbb{R}^3$ has Ratio Property $[\varrho_0]$ then slivers can be removed by assigning small real weights to the points. In other words, Ratio Property $[\varrho_0]$ implies the existence of a weighted Delaunay triangulation without any badly shaped tetrahedron. This complements the implication in the other direction proved by Talmor [21] and thus establishes the equivalence of Ratio Property $[\varrho_0]$ and the existence of triangulations without badly shaped tetrahedra.

Experiments. The technical statement of our result involves a positive constant σ_0 that tells slivers from other tetrahedra. For practical purposes a large σ_0 is desirable. The estimate for σ_0 provided by the Sliver Theorem is miserably tiny, and it will be important to implement the algorithm and collect empirical estimates from computational experiments. The primary goal is to gain insight into how big a constant σ_0 we can expect in practical cases and how σ_0 depends on ϱ_0 and on ω_0.

The sequential algorithm uses an ordering of the vertices and it would be interesting to know whether some orderings perform better than others. The worst-first ordering suggests itself, but it is not clear that it yields higher values of σ_0 than a random ordering.

Boundary effects. Until now we avoided any mention of boundary effects. Applications usually triangulate bounded and non-convex domains $\Omega \subseteq \mathbb{R}^3$ given in terms of their boundary represented by a 2-dimensional complex B. If we choose a finite set $S \subseteq \Omega$ we can construct $\mathrm{Del}\, S$ and remove simplices outside Ω. This works fine as long as $\mathrm{Del}\, S$ *conforms* to B, by which we mean that $\mathrm{Del}\, S$ contains a 2-dimensional subcomplex that subdivides B. However, finding a set S so its Delaunay triangulation conforms to a given 2-dimensional complex is a difficult problem in general. Edelsbrunner and Tan [10] describe a polynomial solution to the 2-dimensional version of the problem, but at this time there is no such solution available in $\mathbb{R}^3$. Heuristic strategies that add points on B dense enough to force boundary conformity seem to work in practice and are described in the applied literature. The algorithm described in Section 7 works fine even for triangulations of non-convex domains, and it can remove slivers inside and outside Ω. Occasionally, the change of a weight will challenge the conformity of the weighted Delaunay triangulation, and additional points will have to placed to reinforce the boundary.

Acknowledgements

The authors thank Jeff Erickson for suggesting some of the names in our categorization of undesirable tetrahedra and Dafna Talmor for helpful discussions, in particular on the material covered in Section 5.

References

[1] M. Bern, L. P. Chew, D. Eppstein and J. Ruppert. Dihedral bounds for mesh generation is high dimensions. *In* "Proc. 6th Ann. ACM-SIAM Sympos. Discrete Algorithms 1995", 189–196.

[2] M. Bern, D. Eppstein and J. Gilbert. Provably good mesh generation. *J. Comput. Syst. Sci.* **48** (1994), 384–409.

[3] L. J. Billera, P. Filliman and B. Sturmfels. Construction and complexity of secondary polytopes. *Adv. Math.* **83** (1990), 155–179.

[4] J. C. Cavendish, D. A. Field and W. H. Frey. An approach to automatic three-dimensional finite element mesh generation. *Internat. J. Numer. Methods Engrg.* **21** (1985), 329–347.

[5] L. P. Chew. Guaranteed-quality Delaunay meshing in 3D. Short paper in "Proc. 13th Ann. Sympos. Comput. Geom., 1997", 391–393.

[6] P. G. Ciarlet. *The Finite Element Method for Elliptic Problems.* North-Holland, 1978.

[7] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk* **7** (1934), 793–800.

[8] T. K. Dey, C. Bajaj and K. Sugihara. On good triangulations in three dimensions. *Internat. J. Comput. Geom. Appl.* **2** (1992), 75–95.

[9] H. Edelsbrunner and N. R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica* **15** (1996), 223–241.

[10] H. Edelsbrunner and T. S. Tan. An upper bound for conforming Delaunay triangulations. *Discrete Comput. Geom.* **10** (1993), 197–213.

[11] A. M. Frieze, G. L. Miller and S.-H. Teng. Separator based divide and conquer in computational geometry. *In* "Proc. 5th ACM Sympos. Parallel Algor. Arch., 1992", 420–430.

[12] I. M. Gelfand, M. M. Kapranov and A. V. Zelevinsky. *Discriminants, Resultants and Multidimensional Determinants.* Birkhäuser, Boston, 1994.

[13] B. Joe. Delaunay versus min-max solid angle triangulations for three-dimensional mesh generation. *Internat. J. Numer. Methods Engrg.* **31** (1991), 987–997.

[14] A. Liu and B. Joe. Relationship between tetrahedron shape measures. *BIT* **34** (1994), 268–287.

[15] G. L. Miller, D. Talmor, S.-H. Teng and N. Walkington. A Delaunay based numerical method for three dimensions: generation, formulation and partition. *In* "Proc. 27th Ann. ACM Sympos. Theory Comput. 1995", 683–692.

[16] S. A. Mitchell and S. A. Vavasis. Quality mesh generation in three dimensions. In "Proc. 8th Ann. Sympos. Comput. Geom., 1992", 212–221.

[17] V. T. Rajan. Optimality of the Delaunay triangulation in $\mathbb{R}^d$. *Discrete Comput. Geom.* **12** (1994), 189–202.

[18] J. Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms* **18** (1995), 548–585.

[19] J. Shewchuck. Tetrahedral mesh generation by Delaunay refinement. *In* "Proc. 14th Ann. Sympos. Comput. Geom. 1998", 86–95.

[20] G. Strang and G. J. Fix. *An Analysis of the Finite Element Method.* Prentice Hall, Englewood Cliffs, New Jersey, 1973.

[21] D. Talmor. Well-spaced points for numerical methods. Report CMU-CS-97-164, Dept. Comput. Sci., Carnegie-Mellon Univ., Pittsburgh, Penn., 1997.

[22] G. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. *J. Reine Angew. Math.* **133** (1907), 97–178, and **134** (1908), 198–287.