Anisotropic Voronoi Diagrams and Guaranteed-Quality Anisotropic Mesh Generation

François Labelle
flab@cs.berkeley.edu

Jonathan Richard Shewchuk
jrs@cs.berkeley.edu

Department of Electrical Engineering and Computer Sciences
University of California at Berkeley
Berkeley, California 94720

Abstract

We introduce *anisotropic Voronoi diagrams*, a generalization of multiplicatively weighted Voronoi diagrams suitable for generating guaranteed-quality meshes of domains in which long, skinny triangles are required, and where the desired anisotropy varies over the domain. We discuss properties of anisotropic Voronoi diagrams of arbitrary dimensionality—most notably circumstances in which a site can see its entire Voronoi cell. In two dimensions, the anisotropic Voronoi diagram dualizes to a triangulation under these same circumstances. We use these properties to develop an algorithm for anisotropic triangular mesh generation in which no triangle has an angle smaller than 20°, as measured from the skewed perspective of any point in the triangle.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems

General Terms

Algorithms, Theory

Keywords

Anisotropic Voronoi diagram, anisotropic mesh generation

1. Introduction

The best-performing triangulations for interpolation and numerical modeling have triangles or tetrahedra whose aspect ratios and orientations are chosen to suit the function they interpolate, or the partial differential equation whose solution they approximate. Triangles, tetrahedra, or most generally d-simplices in E^d (henceforth *elements*) that are nearly equilateral are excellent for some applications; for others, elements that are long and thin, like those depicted in Figure 1, can offer better accuracy with fewer elements [3, 9, 12, 18]. Applications in the latter class are said to exhibit *anisotropic* behavior.

The construction of anisotropic triangulations that meet the needs of these applications is an important problem for which many heuristic solutions are available [5, 7, 11, 19]. However, these algorithms have no guarantee of success. For example, the edge flip algorithm for constructing the Delaunay triangulation is easily modified to take anisotropy into account. Alternatively, George and Borouchaki [7] suggest an anisotropic version of the Bowyer–Watson algorithm [6, 20] for inserting a site into a Delaunay triangulation. But is the final triangulation produced by either of these algorithms unique? What are its properties? Will the flip algorithm terminate or loop forever?

Here we describe an approach that puts anisotropic meshing on firm theoretical ground. In Section 3 we define *anisotropic Voronoi diagrams*, a generalization of multiplicatively weighted Voronoi diagrams [4]. Anisotropic Voronoi diagrams can be defined in any dimensionality. The geometric dual of an anisotropic Voronoi diagram is not generally a triangulation. We describe conditions in which the Voronoi cells are guaranteed to be entirely visible from their generating sites in Section 5. For the special case of two dimensions, the same conditions also guarantee that the planar dual is a geometric triangulation with no inverted triangles.

These results make possible an algorithm that generates high-quality two-dimensional anisotropic meshes by refining an anisotropic Voronoi diagram to enforce the conditions that guarantee that the dual is a triangulation, and to remove any poor-quality elements

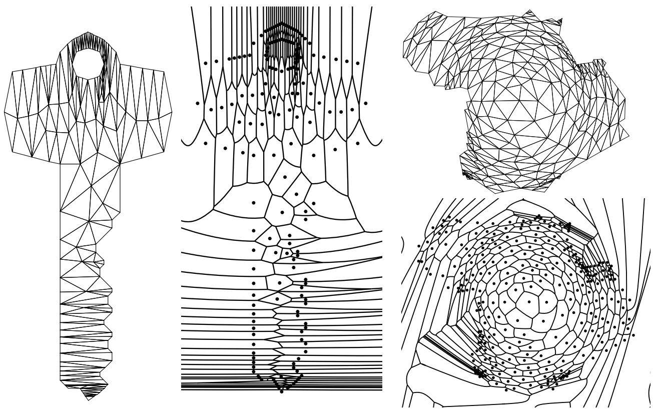

Figure 1: Anisotropic meshes generated by Voronoi refinement, and the anisotropic Voronoi diagrams used to generate them.

Supported in part by the National Science Foundation under Awards ACI-9875170, CMS-9980063, CCR-0204377, and EIA-9802069, and by a gift from the Okawa Foundation.

SoCG'03, June 8–10, 2003, San Diego, California, USA.

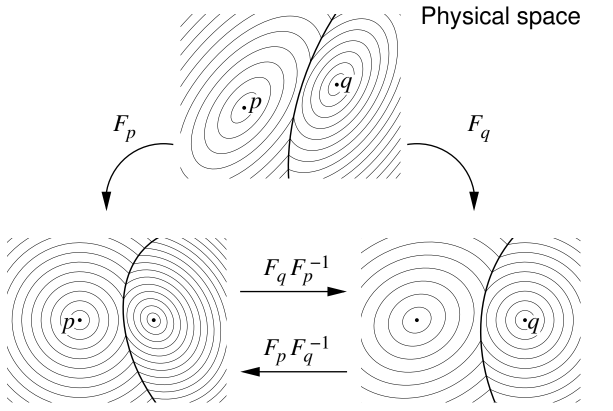

Figure 2: The deformation tensor F_p maps physical space into a space where p's distance metric is isotropic.

from the triangulation. We present the algorithm in Section 9, and prove in Section 10 that it generates an anisotropic mesh whose triangles are all of good quality.

2. Anisotropy: Measures and Goals

Consider a domain $\Omega \subseteq E^d$. Suppose that at each point p in Ω there is a symmetric positive definite *metric tensor* M_p, provided by the user, which dictates how lengths and angles are measured from the perspective of p. The metric tensor is most easily represented as a $d \times d$ matrix. We wish to define a Voronoi diagram over E^d, and perhaps its Delaunay dual restricted to Ω.

Given a metric tensor M_p, define a *deformation tensor* F_p to be any $d \times d$ matrix satisfying

$$M_p = F_p^T F_p \quad \text{and} \quad \det F_p > 0. \tag{1}$$

F_p maps the *physical space* E^d to a rectified space where lengths, areas, and angles as p sees them are measured in the usual way, as Figure 2 illustrates. For example, if q_1 and q_2 are two points in E^d, the distance between q_1 and q_2 as measured by p is

$$d_p(q_1, q_2) = \|F_p q_1 - F_p q_2\|_2 = \sqrt{(q_1 - q_2)^T M_p (q_1 - q_2)}.$$

We will also use the shorthand notations $d_p(q) = d_p(p, q)$ and $d(p, q) = \min\{d_p(q), d_q(p)\}$. Note that $d_p(\cdot, \cdot)$ satisfies the triangle inequality, but $d(\cdot, \cdot)$ does not.

The angle $\theta = \angle q_1 q_2 q_3$ as measured by p is

$$\theta = \arccos \frac{(q_1 - q_2)^T M_p (q_3 - q_2)}{d_p(q_1, q_2) d_p(q_3, q_2)}.$$

In the mesh generation problem, each point in the domain would like to be in an element that is as close to equilateral and equiangular as possible, as measured by that point. For example, if a point q lies in a triangle with vertices u, v, and w, the triangle with vertices $F_q u$, $F_q v$, and $F_q w$ should have no angle close to 0° or 180°.

The deformation tensor F_p is underconstrained. If the $d \times d$ matrix Q represents a proper orthogonal transformation (i.e. a rotation), replacing F_p with QF_p makes no difference. To compute F_p from M_p one can choose the Cholesky decomposition, a symmetric square root, or any other F_p that satisfies (1).

Where does M_p come from? It usually expresses the effects of element shape on interpolation error—or for finite element methods, the effects of shape on discretization error and stiffness matrix conditioning. For interpolation, M_p is often closely related to the Hessian matrix of the function to be interpolated, and may even be the Hessian. Unlike the Hessian, M_p must be positive definite, but its eigenvalues may encode upper bounds on the magnitude of the curvature of the function at p along d principal axes of curvature. For finite element methods, M_p may also take into account the natural anisotropy of the partial differential equation whose solution is sought. See elsewhere [18] for details.

M_p may or may not also encode the ideal size of each element. We have the option of representing the ideal shape and size in the metric tensor, or of representing shape only and handling size separately. With the former option, the ideal element has, say, a unit edge length as measured by p; a larger value of M_p indicates the desire for smaller elements. (But be aware that in two dimensions, the circumradius as measured by p is a better gauge of interpolation error than the edge lengths [18].) We revisit these choices in Section 9.

A central tool in our work is the idea of the *relative deformation* $F_q F_p^{-1}$ which maps a point p's view of the world to another point q's. (See Figure 2.) The *relative distortion* $\tau(p, q) = \max\{\|F_q F_p^{-1}\|_2, \|F_p F_q^{-1}\|_2\}$ gives an upper bound on how differently p and q perceive distances. For any points p, q, and a, $\tau(p, q) \geq 1$, $\tau(p, q) = \tau(q, p)$, and $\tau(p, q) \leq \tau(p, a)\tau(a, q)$.

PROPOSITION 1. *Let p, q, a, and b be points in Ω. Then*

$$\frac{d_p(a, b)}{\|F_p F_q^{-1}\|_2} \leq d_q(a, b) \leq \|F_q F_p^{-1}\|_2 \, d_p(a, b) \quad \textit{and}$$
$$\frac{d_p(a, b)}{\tau(p, q)} \leq d_q(a, b) \leq \tau(p, q) d_p(a, b).$$

PROOF. $d_q(a, b) = \|F_q(a - b)\|_2 = \|F_q F_p^{-1} F_p(a - b)\|_2 \leq \|F_q F_p^{-1}\|_2 \|F_p(a - b)\|_2 = \|F_q F_p^{-1}\|_2 d_p(a, b) \leq \tau(p, q) d_p(a, b)$. Similarly, $d_p(a, b) \leq \|F_p F_q^{-1}\|_2 d_q(a, b) \leq \tau(p, q) d_q(a, b)$. ■

3. Anisotropic Voronoi Diagrams

The domain Ω and the metric tensor field M together are a Riemannian manifold, so the natural way to define an anisotropic Voronoi diagram of a point set is to compute a Voronoi diagram on that manifold. Leibon and Letscher [10] do just that. In Riemannian geometry, the length of a path $\mathbf{s} = \mathbf{s}(t)$, $0 \leq t \leq 1$, connecting two points $p = \mathbf{s}(0)$ and $q = \mathbf{s}(1)$ is calculated by the path integral

$$R(\mathbf{s}) = \int_0^1 \sqrt{\frac{d\mathbf{s}}{dt}^T M_{\mathbf{s}(t)} \frac{d\mathbf{s}}{dt}} \, dt,$$

and the distance between two points p and q is the length of the shortest path connecting them. The Voronoi diagram is defined in the usual manner: the Voronoi cell $\mathrm{Vor}(v)$ of a point site v is the set of all points on the manifold that are at least as close to v as to any other site. Leibon and Letscher show that if there is an upper bound on the sectional curvature of the manifold, and if the sites are spaced densely enough, the Voronoi diagram dualizes to a well-defined Delaunay triangulation.

Unfortunately, finding the shortest Riemannian path between two points is a computationally difficult operation. Local geodesic paths can be approximated numerically, but the algorithms for doing so are slow enough to make generating large anisotropic meshes impractical. These problems can be overcome by using approximate distance computations and approximate geodesics or other heuristics [7], but none of these has been amenable to the kind of analysis that would yield a provably good algorithm.

Here, we propose an approximation to the Riemannian Voronoi diagram that can be generated reasonably efficiently, yet makes it possible to prove that our Delaunay triangulations have favorable properties. Like Leibon and Letscher, we show that if the spacing

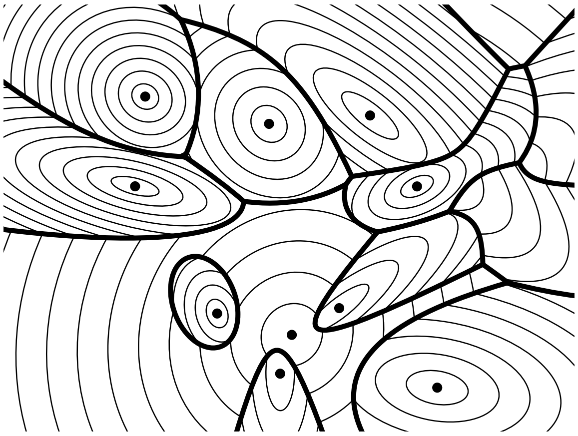

Figure 3: An anisotropic Voronoi diagram. Thin arcs are isocontours of the nearest site's distance metric.

of the sites satisfies the right conditions, we can guarantee a geometrically well-defined two-dimensional Delaunay triangulation; furthermore, we can offer triangles with guaranteed good quality.

DEFINITION 1 (ANISOTROPIC VORONOI DIAGRAM). *Let V be a set of sites. The* Voronoi cell *of a site v in V is*

$$\mathrm{Vor}(v) = \{p \in E^d : d_v(p) \le d_w(p) \text{ for all } w \in V\}.$$

Any subset of sites $W \subseteq V$ induces a Voronoi cell $\mathrm{Vor}(W) = \bigcap_{w \in W} \mathrm{Vor}(w)$ of points equally close to the sites in W and no closer to any others. If it is not empty, such a cell has a dimensionality of $\dim(\mathrm{Vor}(W)) \ge d+1-|W|$, achieving equality if the sites are in general position. Every site in W is said to own $\mathrm{Vor}(W)$. *The* anisotropic Voronoi diagram *of V is the arrangement of the Voronoi cells $\{\mathrm{Vor}(W) : W \subseteq V, W \ne \emptyset, \mathrm{Vor}(W) \ne \emptyset\}$.*

Figure 3 depicts an example. Anisotropic Voronoi diagrams are a generalization of multiplicatively weighted Voronoi diagrams [4] in which the distance metric is anisotropic. If the metric tensor field M is isotropic (i.e. for some scalar field c, $M_p = c_p I$ for all $p \in \Omega$, where I is the identity tensor), the anisotropic Voronoi diagram is the multiplicatively weighted Voronoi diagram.

One odd characteristic that anisotropic Voronoi diagrams share with multiplicatively weighted Voronoi diagrams is that a Voronoi cell (of any dimension) can consist of multiple connected components. Cells are partitioned into *faces* of the arrangement whose relative interiors are connected. 0-faces are *Voronoi vertices*, and 1-faces are *Voronoi arcs*. A d-face that does not contain its generating site is called an *orphan*. (Figure 3 has three orphans.)

PROPOSITION 2. *The boundary between two adjoining d-cells is composed of patches of a quadratic curve or surface.*

PROOF. Any point q in $\mathrm{Vor}(v) \cap \mathrm{Vor}(w)$ satisfies $d_v(q) = d_w(q)$, or $(q-v)^T M_v (q-v) = (q-w)^T M_w (q-w)$. This equation is quadratic in q, therefore every point in $\mathrm{Vor}(v) \cap \mathrm{Vor}(w)$ lies on a common quadratic curve or surface. ■

In three dimensions quadratic surfaces are also called quadrics. In two dimensions, nondegenerate quadratic curves are also called conic sections (circles, ellipses, parabolas, and hyperbolas).

Another way to think of anisotropic Voronoi diagrams is as the lower envelope of the arrangement in E^{d+1} of the paraboloids $z_v(p) = (p-v)^T M_v (p-v)$. By projecting the faces of the lower envelope down to E^d we form the *minimization diagram* [16] of the paraboloids, which is the anisotropic Voronoi diagram.

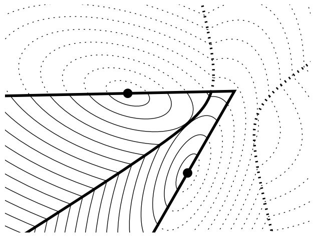

Figure 4: A wedge.

4. Diagram Complexity and Construction

The complexity of a d-dimensional n-site anisotropic Voronoi diagram is in $O(n^{d+\epsilon})$, where ϵ is an arbitrary (small) positive constant, by virtue of Halperin and Sharir's upper bounds on the complexity of lower envelopes [8, 15]. To construct a lower bound example of $\Omega(n^d)$ worst-case complexity, choose paraboloids whose isocontours are frisbee-shaped, axis-aligned, and form a grid. (Our thanks go to Micha Sharir for this example).

Two-dimensional anisotropic Voronoi diagrams can be constructed in $O(n^{2+\epsilon})$ time by a divide-and-conquer algorithm of Agarwal, Schwarzkopf, and Sharir [2, 16] for minimization diagrams. Three-dimensional anisotropic Voronoi diagrams can be constructed in $O(n^{3+\epsilon})$ expected time by a random incremental algorithm of Agarwal, Aronov, and Sharir [1, 16].

For the purpose of Voronoi refinement, we need an incremental site insertion algorithm, but we cannot randomize the order in which sites are inserted. At any rate, any incremental site insertion algorithm is too slow for our needs. We discuss a sneaky alternative, fast enough for practical use, in Sections 7 and 8.

5. Anisotropic Delaunay Triangulations

The dual of the standard Voronoi diagram is the Delaunay triangulation. Our anisotropic Voronoi diagram can be very complicated, and its geometric dual may contain inverted or repeated simplices and other irregularities. In this section we describe conditions under which the dual of the anisotropic Voronoi diagram is a correct triangulation.

DEFINITION 2. *Let v and w be two sites. Define the* wedge *between these two sites as the locus of points q for which the angle $\angle qvw$ as viewed from v is less than $90°$, and the angle $\angle qwv$ as viewed from w is less than $90°$. (See Figure 4.) Mathematically,*

$$\begin{aligned}\mathrm{wedge}(v,w) \quad &= \quad \{q \in E^d : (q-v)^T M_v (w-v) > 0 \\ &\qquad \textit{and } (q-w)^T M_w (v-w) > 0\}.\end{aligned}$$

LEMMA 3 (VISIBILITY LEMMA). *Let v and w be two sites in E^d. If we restrict the two-site Voronoi diagram of $\{v,w\}$ to* $\mathrm{wedge}(v,w)$*, then v can see its entire cell, and w can see its entire cell as well.*

PROOF. The restricted Voronoi cell of v is defined by the following three inequalities.

$$(q-v)^T M_v (w-v) > 0.$$
$$(q-w)^T M_w (v-w) > 0.$$
$$(q-w)^T M_w (q-w) \ge (q-v)^T M_v (q-v).$$

Let q be any point that satisfies these inequalities. The visibility claim is that $q' = \lambda q + (1-\lambda) v$ satisfies these inequalities as well, for $0 < \lambda < 1$. This claim can be verified by substituting q' for q and showing that the inequalities hold for q', given that they hold for q. The result holds for w by symmetry. ∎

Lemma 3 is tight. The visibility property stops precisely where the Voronoi surface exits the wedge.

DEFINITION 3. *A Voronoi k-face $f \subseteq \mathrm{Vor}(W)$, with $0 \leq k < d$, is said to be* wedged *if for every pair of distinct sites $v_1, v_2 \in W$, every point q on f falls inside* $\mathrm{wedge}(v_1, v_2)$. *For example, $f \subseteq \mathrm{Vor}(\{v_1, v_2, v_3\})$ is wedged if $f \subseteq \mathrm{wedge}(v_1, v_2) \cap \mathrm{wedge}(v_2, v_3) \cap \mathrm{wedge}(v_3, v_1)$.*

THEOREM 4 (VISIBILITY THEOREM). *If every lower-dimensional face of a d-face of $\mathrm{Vor}(v)$ is wedged, then the d-face is star-shaped and every point in the d-face is visible from v.*

PROOF. For the sake of contradiction, let p be a point in the d-face that is not visible from v. Let q be the point furthest from p on the line segment pv that is visible from p. Because p is not visible from v, neither is q. Let w be the owner of the first face encountered strictly after q on the ray $\vec{qv}$. Then q lies on a face owned by v and w. By assumption this face is wedged. Imagine a Voronoi diagram with sites v and w only. In this diagram q is in $\mathrm{Vor}(v)$ and in $\mathrm{wedge}(v, w)$, but q is not visible from v, contradicting Lemma 3. ∎

The following lemma implies that if a Voronoi surface is not wedged, we can insert a new site on it that is not close to an existing site—a handy tool for mesh generation.

LEMMA 5. *Let q be a point in $\mathrm{Vor}(v) \cap \mathrm{Vor}(w)$ that lies outside $\mathrm{wedge}(v, w)$ on the side of w. Let $\gamma \geq 1$ be a constant for which $\tau(v, w) \leq \gamma$. Then the proximity of q to v and w is bounded by $d_v(q) = d_w(q) \geq d_w(v)/\sqrt{\gamma^2 - 1}$.*

PROOF. Because q is on w's side of $\mathrm{wedge}(v, w)$, $d_w(q, v)^2 \geq d_w(q)^2 + d_w(v)^2$ by Pythagoras' Theorem. By Proposition 1, $d_w(q, v)^2 \leq \tau(v, w)^2 d_v(q)^2 \leq \gamma^2 d_v(q)^2$. Because $q \in \mathrm{Vor}(v) \cap \mathrm{Vor}(w)$, $d_v(q) = d_w(q)$. The result follows by combining inequalities and rearranging terms. ∎

The rest of this section applies to the two-dimensional case only.

LEMMA 6 (TRIANGLE ORIENTATION LEMMA). *Let q be a Voronoi vertex owned by the sites v_1, v_2, v_3. If q is wedged, then the orientation of the triangle $v_1v_2v_3$ matches the ordering of the cells $\mathrm{Vor}(v_1), \mathrm{Vor}(v_2), \mathrm{Vor}(v_3)$ locally around q. In other words, if at q the cells $\mathrm{Vor}(v_1), \mathrm{Vor}(v_2), \mathrm{Vor}(v_3)$ occur clockwise, then the sites v_1, v_2, v_3 occur clockwise in the plane, and vice versa.*

PROOF. The triangle $v_1v_2v_3$ cannot be degenerate, because if, say, the angle at v_1 is 180°, then $\mathrm{wedge}(v_1, v_2)$ and $\mathrm{wedge}(v_1, v_3)$ are disjoint (because they are defined as open sets), and so q cannot lie in their intersection. Imagine a Voronoi diagram with sites v_1, v_2, and v_3 only. By Lemma 3, $d_{v_1}(q') \leq d_{v_2}(q')$ and $d_{v_1}(q') \leq d_{v_3}(q')$ for any point q' on the line segment v_1q. Therefore v_1 can see q in the three-point Voronoi diagram. Symmetrically, so can v_2 and v_3.

There are three cases. (1) q lies in the triangle $v_1v_2v_3$. (2) q lies on the opposite side of exactly one edge of the triangle. (3) q lies on the opposite side of exactly two edges. In cases (1) and (2), the three-way visibility property implies matching orientations as claimed. Case (3) implies opposite orientations, but case (3) is impossible. If q lies on the opposite side of edges v_1v_2 and v_1v_3, say, then as measured by v_1, $\angle qv_1v_2 + \angle qv_1v_3 > 180^\circ$, which contradicts the wedge properties $\angle qv_1v_2 < 90^\circ$ and $\angle qv_1v_3 < 90^\circ$. The lemma follows. ∎

THEOREM 7 (DUAL TRIANGULATION THEOREM). *Let the domain Ω be a polygonal subset of the plane, let V be a set of sites in Ω which include every vertex of Ω, and let D be the anisotropic Voronoi diagram of V. Let $D|_\Omega$ be the restriction of D to Ω. Suppose that each Voronoi arc cut by the restriction operation is owned by the endpoints of the edge of Ω that cuts it. If all the Voronoi arcs and vertices of $D|_\Omega$ are wedged, then the geometric dual of $D|_\Omega$ is a polygonalization of Ω (with strictly convex polygons), and is a triangulation of Ω if V is in general position. Arbitrarily triangulating each polygon yields what we call an* anisotropic Delaunay triangulation *of (V, Ω).*

PROOF. $D|_\Omega$ has no orphans because an orphan in D that survives the restriction to Ω unchanged is ruled out by Theorem 4, and an orphan that is cut or created by the restriction to Ω would defy the assumption that each cut Voronoi arc is dual to the domain edge that cuts it. Thus every 2-face in $D|_\Omega$ dualizes to the site that generates it.

Every Voronoi vertex has degree three or greater, and thus dualizes to a polygon. The polygon cannot have a repeated site v because that would imply that four Voronoi arcs owned by v meet at the vertex: two of these arcs have points that are not visible from v, contradicting Lemma 3. If the degree of the vertex exceeds three, triangulate the polygon arbitrarily. Every Voronoi vertex of $D|_\Omega$ is wedged, so by Lemma 6, triangles that share an edge have compatible orientations. By transitivity, this means that all triangles are positively oriented regardless of how each polygon is triangulated, so each polygon is strongly convex.

Every uncut Voronoi arc of $D|_\Omega$ is incident to two distinct Voronoi vertices, because the Voronoi arc lies within the wedge of its two generating sites so it cannot form a loop around one of the sites. Every uncut Voronoi arc thus dualizes to an edge between two triangles, and by assumption every cut arc dualizes to an edge on the boundary of Ω. Therefore, the geometric dual of $D|_\Omega$ is a valid triangulation of V covering Ω. ∎

When $M_p \equiv I$, Voronoi edges and vertices are always wedged, and Theorem 7 reiterates what we already know: the dual of the standard Voronoi diagram is a triangulation. However, for arbitrary sets of sites and arbitrary metric tensor fields, the preconditions of Theorem 7 seldom hold. The theorem becomes useful in conjunction with the Voronoi refinement method described in Section 9.

6. Triangle Quality

There is a simple relationship between the minimum angle $\theta_{\min}$ of a triangle and the ratio β of its circumradius to its shortest edge length: $\sin \theta_{\min} = \frac{1}{2\beta}$. Here we revisit this relationship for our anisotropic notions of "circumradius" and "shortest edge."

LEMMA 8. *Consider a triangle with sides a, b, and c and opposite angles A, B, and C such that $c^2 \geq |a^2 - b^2|$. Imagine another triangle with sides a', b', and c' and opposite angles A', B', and C'. If $a' \geq a$, $b' \geq b$ and $c' \leq c$, then $C' \leq C$.*

PROOF. Scaling does not change angles, so we can assume $c' = c$ by scaling every side of the second triangle by $c/c' \geq 1$. Imagine the circle through the vertices of the first triangle. Since $c' = c$, fix the vertices at A and B and generate the second triangle from the first by moving the vertex at C only. $c^2 \geq |a^2 - b^2|$ implies $c^2 \geq a^2 - b^2$ and $c^2 \geq b^2 - a^2$, which implies $A \leq 90^\circ$ and $B \leq 90^\circ$, so when the sides a and b increase, the vertex at C can only move out of the circle, which by circle geometry implies that $C' \leq C$. ∎

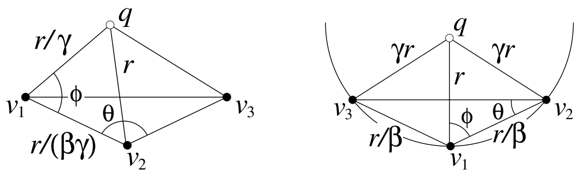

Figure 5: Two cases for bounding $\theta = \angle v_1 v_2 v_3$.

THEOREM 9. *Let q be a Voronoi vertex owned by the sites v_1, v_2, and v_3. Let the* circumradius *associated with q be $r = d_{v_1}(q) = d_{v_2}(q) = d_{v_3}(q)$. Let the* shortest edge length *ℓ of the triangle $t = \triangle v_1 v_2 v_3$ be $\min\{d(v_1, v_2), d(v_2, v_3), d(v_3, v_1)\}$. Let $1 \leq \gamma \leq \sqrt{2}$ and $\beta \geq 1/\sqrt{2}$ be constants such that $\beta \geq r/\ell$ and $\tau(v_i, v_j) \leq \gamma$ for $i, j \in \{1, 2, 3\}$. Let $\chi = 1/(2\beta) - (\gamma^2 - 1)\beta/2$.*

If $\chi > 0$, then q is wedged. Furthermore, each angle θ of t satisfies $\theta \geq \arcsin \chi$, where θ is measured from the perspective of some vertex of t (strangely, not necessarily the same vertex for each angle, and not necessarily the vertex where the angle is). Note that this only implies that $\sin \theta \geq \chi$ if θ is acute.

PROOF. Let $\phi = \angle q v_1 v_2$ as measured by v_1. Note that $d_{v_1}(v_1, q) = r$, $d_{v_1}(v_1, v_2) \geq r/\beta$, and $d_{v_1}(v_2, q) \leq \gamma d_{v_2}(q) = \gamma r$. Because $\gamma \geq 1$ and $\beta \geq 1/\sqrt{2}$, we have $(\gamma r)^2 \geq |r^2 - (r/\beta)^2|$ so Lemma 8 applies and by the law of cosines,

$$\cos \phi \geq \frac{r^2 + (r/\beta)^2 - (\gamma r)^2}{2r(r/\beta)} = \frac{1}{2\beta} - \frac{(\gamma^2 - 1)\beta}{2} = \chi.$$

If $\chi > 0$, then $\cos \phi > 0$, so $\phi < 90^\circ$ and one of the two inequalities required to show that $q \in \mathrm{wedge}(v_1, v_2)$ is satisfied (namely $(q - v_1)^T M_{v_1} (v_2 - v_1) > 0$). By repeating the argument for $\phi = \angle q v_i v_j$, $i, j \in \{1, 2, 3\}, i \neq j$, we obtain all six inequalities required to show that $q \in \mathrm{wedge}(v_1, v_2) \cap \mathrm{wedge}(v_2, v_3) \cap \mathrm{wedge}(v_3, v_1)$; therefore q is wedged.

Next we bound $\theta = \angle v_1 v_2 v_3$, the angle at vertex v_2. Consider the case where q lies on the opposite side of edge $v_1 v_3$ from v_2 (but on the triangle-side of edges $v_1 v_2$ and $v_2 v_3$) as illustrated in Figure 5 (left). As measured by v_2, $\angle v_2 q v_1 + \angle v_2 q v_3 < 180^\circ$. Assume without loss of generality that $\angle v_2 q v_1 < 90^\circ$. If we compute the angle $\phi = \angle q v_1 v_2$ as before, but now as measured by v_2, we have $d_{v_2}(v_1, q) \geq r/\gamma$, $d_{v_2}(v_1, v_2) \geq r/\beta \geq r/(\beta\gamma)$, and $d_{v_2}(v_2, q) = r$. Multiplying these bounds by γ does not change the angles, so Lemma 8 applies as before and $\cos \phi \geq \chi$ as measured by v_2 as well. Because $\angle v_2 q v_1 < 90^\circ$, this implies that $\sin \angle v_1 v_2 q \geq \chi$, which in turn implies that $\theta \geq \angle v_1 v_2 q \geq \arcsin \chi$ as measured by v_2.

The proof of Lemma 6 tells us there are only two other cases: either q lies in t, or q lies on the opposite side of exactly one edge, $v_1 v_2$ or $v_2 v_3$—assume without loss of generality $v_2 v_3$. In either case the line $v_1 q$ cuts the angle $\angle v_3 v_1 v_2$. Consider Figure 5 (right), drawn from v_1's perspective, where t achieves $\sin \theta = \cos \phi = \chi$. In the figure, imagine v_1 and q are fixed. We claim that the illustrated positions of v_2 and v_3 minimize θ.

For the case where q lies on the opposite side of $v_2 v_3$, we choose to measure $\theta = \angle v_1 v_2 v_3$ from v_1's perspective. Because $d_{v_1}(v_2) \geq r/\beta$ and $d_{v_1}(v_2, q) \leq \gamma r$, v_2 is constrained to lie *inside* the illustrated circle (which circumscribes our proposed worst-case t). By circle geometry, moving v_2 subject to this constraint can only increase θ. Once v_2 is fixed, moving v_3 subject to the symmetrical constraints $d_{v_1}(v_3) \geq r/\beta$ and $d_{v_1}(v_3, q) \leq \gamma r$ cannot decrease θ either. Therefore from v_1's perspective, $\theta \geq \arcsin \chi$.

For the case where q lies in t, there are two other configurations of the vertices that sometimes locally minimize θ, but the case just described is always the global minimum for the ranges of γ and β specified in the theorem. Details are omitted.

Rotate the vertex labels and repeat the argument to provide a lower bound for all three angles of t. ■

COROLLARY 10. *Let q, v_1, v_2, v_3, γ, and β satisfy the conditions of Theorem 9. Let p be an arbitrary point such that $\tau(p, v_i) \leq \gamma$ for $i \in \{1, 2, 3\}$. Let $\chi = 1/(2\beta) - (\gamma^2 - 1)\beta/2$. Then each angle θ of the triangle $t = \triangle v_1 v_2 v_3$ satisfies $\arcsin(\chi/\gamma^2) \leq \theta \leq 2 \arccos(\chi/\gamma^2)$, as measured by p.*

PROOF. Let $\theta_p = \angle v_1 v_2 v_3$ as measured by p. Write θ_{v_i} for the same angle as measured by v_i. By Theorem 9, for some $j \in \{1, 2, 3\}$, either $\sin \theta_{v_j} \geq \chi$ or $\theta_{v_j} \geq 90^\circ$. Let A_p be the area of t as measured by p, and A_{v_j} the area of t as measured by v_j. By a well-known formula for the sine of an angle,

$$\begin{aligned}
\sin \theta_p &= \frac{2A_p}{d_p(v_1, v_2) d_p(v_2, v_3)} = \frac{2A_{v_j} \det(F_p F_{v_j}^{-1})}{d_p(v_1, v_2) d_p(v_2, v_3)} \\
&\geq \frac{2A_{v_j} \|F_p F_{v_j}^{-1}\|_2 / \|F_{v_j} F_p^{-1}\|_2}{\|F_p F_{v_j}^{-1}\|_2^2 d_{v_j}(v_1, v_2) d_{v_j}(v_2, v_3)} \\
&= \frac{2A_{v_j}}{\|F_{v_j} F_p^{-1}\|_2 \|F_p F_{v_j}^{-1}\|_2 d_{v_j}(v_1, v_2) d_{v_j}(v_2, v_3)} \\
&\geq \frac{1}{\gamma^2} \sin \theta_{v_j}.
\end{aligned}$$

Therefore $\sin \theta_p \geq \chi/\gamma^2$ if $\theta_{v_j} \leq 90^\circ$. If θ_{v_j} is obtuse but θ_p is acute, divide θ_{v_j} into a 90° piece and a leftover, and apply the bound to the 90° portion, yielding $\sin \theta_p > 1/\gamma^2 > \chi/\gamma^2$. By repeating the argument for each angle of t, we have that from p's perspective, the sine of every acute angle of t is at least χ/γ^2. A triangle with no angle smaller than $\arcsin(\chi/\gamma^2)$ has no angle greater than $180^\circ - 2\arcsin(\chi/\gamma^2) = 2\arccos(\chi/\gamma^2)$. ■

7. Loose Anisotropic Voronoi Diagrams

In an ordinary Voronoi diagram, inserting a site is by nature a local operation. Although one new site can cause extensive changes to most of the diagram, that is not the common case. Randomized incremental construction algorithms enjoy expected constant-time site insertions in theory (not counting the point location time), and Ruppert's Delaunay refinement algorithm for quality mesh generation [13] usually enjoys constant-time site insertions in practice.

Anisotropic Voronoi diagrams do not inherit this virtue, because a Voronoi cell is not necessarily connected. An incremental site insertion algorithm must inspect the entire diagram in search of points that are closer to the new site than to the older sites, thereby creating orphans where necessary. This makes site insertion slow.

However, a goal of our mesh generation algorithm is to create a Voronoi diagram that has no orphans. Our meshing algorithm works correctly even if we use a sloppy incremental site insertion algorithm that forgets to install new orphans.

Let V be a set of sites, and consider again the arrangement $\mathcal{A}$ in E^{d+1} of the paraboloids $z_v(p) = (p - v)^T M_v (p - v)$ for all $v \in V$, whose lower envelope is the anisotropic Voronoi diagram. Each paraboloid in $\mathcal{A}$ is sliced into d-faces where it intersects other paraboloids. We wish to choose a subset of these d-faces that form a structure that is similar to the lower envelope, but easier to construct.

Let F be a subset of the d-faces of $\mathcal{A}$, and let $\bigcup F$ be the union of these faces. Suppose F satisfies the following two conditions. First, F contains every d-face that contains its generating site. Second, $\bigcup F$ is a manifold without boundary that intersects every vertical line (parallel to the x_{d+1}-axis) at exactly one point. In other

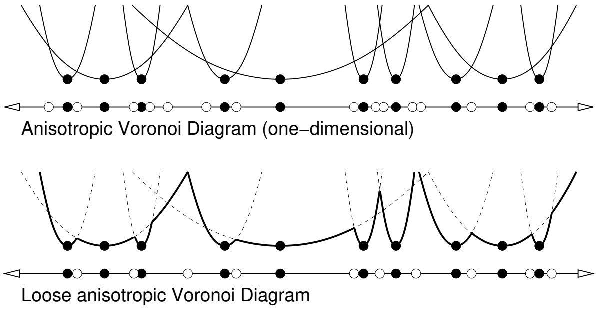

Figure 6: A loose anisotropic Voronoi diagram often has fewer orphans than the true anisotropic Voronoi diagram. Black points are sites; white points are Voronoi vertices.

words, if $f(p)$ is a function that maps each point p in E^d to a real number h such that $(p, h) \in \bigcup F$, then f is single-valued, continuous, and defined everywhere in E^d. See the lower half of Figure 6 for an example. For any F that satisfies these conditions, let L be the arrangement formed by projecting the d-faces in F to E^d; we call L a *loose Voronoi diagram* of V.

For a loose Voronoi diagram L, a site v *owns* a face g of L if g is the projection of a face of F that lies on v's paraboloid.

In general, a set of sites V may have many loose Voronoi diagrams; the true anisotropic Voronoi diagram of V is one of them. They differ from each other by the orphans they contain. For each site v in V, every loose Voronoi diagram of V contains the d-face of $\mathrm{Vor}(v)$ that contains v, or a superset of that face. The other components of $\mathrm{Vor}(v)$ may or may not be present, and v may own territory that is not part of $\mathrm{Vor}(v)$.

Observe that if the true Voronoi diagram D has no orphans, then it is the only loose Voronoi diagram, because E^d is covered by the d-faces that contain their generating sites. Likewise, if $D|_\Omega$ has no orphans, then for any loose diagram L, $L|_\Omega = D|_\Omega$. Therefore, if D satisfies the preconditions of the Dual Triangulation Theorem, then $D|_\Omega$ can be constructed by incrementally inserting its sites into a loose Voronoi diagram.

8. Incremental Voronoi Diagram Construction

This section sketches an incremental site insertion algorithm which, given a two-dimensional loose Voronoi diagram L of a set V of sites and a site $v \notin V$, constructs a loose Voronoi diagram L^v of the set $V \cup \{v\}$. Let F be the set of 2-faces of the paraboloid arrangement $\mathcal{A}$ that project to faces in L. Let z be v's paraboloid, $z(q) = (q - v)^T M_v (q - v)$. Let F^v be the three-dimensional arrangement formed by z and the faces in F. In F^v, z is sliced into 2-faces. One of these 2-faces projects to a 2-face f that contains v.

The main idea is that the algorithm inserts only the 2-face f into L. Therefore, the updated loose Voronoi diagram L^v has only one 2-face owned by v (namely f), and v has no orphans. The update step is thus local in nature and will often run in constant time per site insertion in practice—especially in the Voronoi refinement algorithm of Section 9, which quickly makes the sites regularly spaced. The worst-case running time for a sequence of n site insertions is between $\Omega(n^2)$ and $O(n^3)$ (to narrow this range is an open problem), but the worst-case time is unlikely to be realized during Voronoi refinement.

Observe that the loose Voronoi diagram created by the algorithm is not necessarily orphan-free. The insertion of v does not create any orphans owned by v, but it may split other sites' cells up into multiple faces, thereby creating orphans owned by other sites.

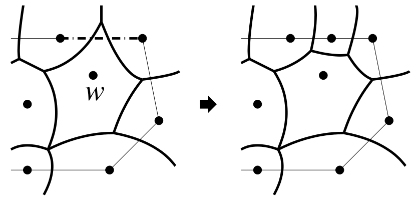

Figure 7: A segment encroached upon by w is split.

The algorithm runs in three steps. The first step finds a 2-face of L that contains the new site v. (In the Voronoi refinement algorithm, this information is already provided.) The second step computes the overlay of f with L—the arrangement in the plane formed by L and the boundary of f. Let G be the set of 2-faces of L that intersect the interior of f. This step exploits locality by using depth-first search to visit only the 2-faces in G. Let g be a face in G with owner w. The algorithm determines whether the boundary of f intersects g—more directly, whether the curve $d_v(p) = d_w(p)$ intersects g—and constructs the overlay of f with g. The details of overlay computation are routine [2, 16] and are omitted.

The third step merges all faces in the interior of f into one 2-face owned by v. The resulting arrangement is L^v.

9. Anisotropic Mesh Generation by Voronoi Refinement

Let X be a planar straight line graph (PSLG)—a set of sites and segments in the plane that are required to appear in the mesh. Let $\Omega \subset E^2$ be a finite domain to triangulate (whose boundary must be the union of some of the segments in X), and let M be a metric tensor field defined over Ω. The following Voronoi refinement algorithm is an anisotropic revision of Ruppert's Delaunay refinement algorithm for guaranteed-quality isotropic mesh generation [13].

To guarantee that his algorithm will terminate, Ruppert demands that any two segments of X that share an endpoint v must be separated by an angle of at least 90°. We impose the condition that the angle is at least 79°, where the angle is measured according to the metric tensor M_v. (The angle constraint can be relaxed a bit; see Theorem 11.)

Voronoi refinement begins with the construction of the Voronoi diagram D of the sites in X under the metric tensor field M. The algorithm works equally well whether D is a true anisotropic Voronoi diagram or a loose anisotropic Voronoi diagram. At this time, the dual of $D|_\Omega$ is not necessarily a triangulation.

Let s be any segment in X, and let a and b be its endpoints. Say that s is *encroached* if it intersects a Voronoi cell belonging to any site w other than a and b (see Figure 7). If s is encroached, *split* it by inserting a new site z in $s \cap \mathrm{Vor}(w)$ (at a location to be discussed shortly). Update the Voronoi diagram D by incrementally inserting z. The splitting operation replaces s with two subsegments az and zb, which may or may not be encroached and need to be split further.

When possible, split a segment so that $d_a(z) = d_b(z)$. Occasionally, however, this is inappropriate because the "midpoint" z satisfying this condition does not lie in $\mathrm{Vor}(w)$ for any encroaching site w. This can happen if $s \cap \mathrm{Vor}(w)$ lies entirely on one side of the "midpoint." In this case, place the new site z in $s \cap \mathrm{Vor}(w)$ as close to the "midpoint" as possible.

The purpose of segment splitting is to ensure that each subsegment is an edge in the geometric dual of $D|_\Omega$. When no encroached segment or subsegment exists, the algorithm attempts to eliminate a poorly shaped triangle, an orphan, or another irregularity by inserting a site. A point $q \in \Omega$ is said to be a *violator* if

- q lies on a Voronoi 1-cell $\mathrm{Vor}(v) \cap \mathrm{Vor}(w)$ but is not wedged (in $\mathrm{wedge}(v, w)$)—call q a *wedge violator*; or
- q is a Voronoi vertex that dualizes to a triangle t that is inverted, is too large, or has an angle less than some constant θ_{bound}, as measured by the metric tensor M_p for any point p in t.

The Voronoi refinement algorithm chooses an arbitrary violator q and attempts to insert a site there, thereby eliminating the violator. It might not be necessary to eliminate every violator; the algorithm may stop as soon as $D|_\Omega$ dualizes to a triangulation of satisfactory quality. However, the proof in Section 10 guarantees termination even when no violator is spared.

A new site is never permitted to encroach upon a segment of X. If q does not encroach upon any segment, the algorithm updates D to reflect the insertion of q, then looks for another violator. If q encroaches upon a segment s, q is not inserted. Instead, s is split as described above. The splitting site z must lie in $s \cap \mathrm{Vor}(q)$, where $\mathrm{Vor}(q)$ is the Voronoi cell that would have been created if q were inserted. Subject to this restriction, z must lie as close to the "midpoint" defined by $d_a(z) = d_b(z)$ as possible.

If the Voronoi refinement algorithm terminates, the dual of $D|_\Omega$ is a triangulation of Ω by Theorem 7, and the triangles are of good quality (because otherwise the algorithm would not stop). But does it terminate? Section 10 shows that, with reasonable restrictions on M and θ_{bound}, it does.

To simplify programming, an implementation only needs to find Voronoi "arcs" that are ellipses (which always have violators), and Voronoi vertices that dualize to inverted or poor-quality triangles. (For this purpose, the "dual" of an orphan is the site that owns it.) An attack on these violators will also eliminate all orphans and "island" Voronoi cells (enclosed in other Voronoi cells, possibly adjoining other islands) without any need to explicitly test for their existence. We omit the proof.

Because the Voronoi refinement algorithm refines the mesh heavily in regions where M varies rapidly, we suggest the following variation. For any point p, let $M'_p = (\det M_p)^{-1/d} M_p$, so the metric tensor M'_p has determinant one. The tensors M_p and M'_p measure angles identically but lengths differently. Use M_p to judge whether a triangle is too large (as before), but define the anisotropic Voronoi diagram using the tensor field M'_p. Then the relative distortion between two points is solely attributable to differences in how they measure angles, and not differences in how they measure areas. In our experiments, the modified algorithm produces meshes with fewer triangles.

10. Proof of Termination

Intuitively, the Voronoi refinement algorithm inserts sites for one of two reasons: either the spacing of sites is not uniform enough to guarantee good-quality triangles, or neighboring sites have metric tensors that strongly disagree with each other. The distances at which these effects occur are called the *local feature size* and the *bounded distortion radius*.

DEFINITION 4. *For a metric tensor field M and a fixed $\gamma > 1$, the* bounded distortion radius $\mathrm{bdr}(p, \gamma)$ *at a point $p \in \Omega$ is the greatest number such that for every point $q \in \Omega$, if $d_p(q) \leq \mathrm{bdr}(p, \gamma)$, then $\tau(p, q) \leq \gamma$. (Note that $\mathrm{bdr}(p, \gamma)$ can be infinite.) The* distortion disk *centered at p is the elliptical disk $\{q \in E^2 : d_p(q) \leq \mathrm{bdr}(p, \gamma)\}$. Let $\mathrm{bdr}_{\min}(\gamma) = \inf_{p \in \Omega} \mathrm{bdr}(p, \gamma)$.*

Every point in p's distortion disk has a similar view of lengths and angles. The Voronoi refinement algorithm is guaranteed to terminate if $\mathrm{bdr}_{\min}(\gamma)$ is positive. This is always true if the metric tensor field M is continuous over Ω with bounded spatial gradients, but is sometimes true even with minor discontinuities in M. The value of γ depends on how strong an angle bound the user demands; see Theorem 11 below.

Let q and q' be two distinct *feature points*—points that lie on input sites or segments in X. Say that q and q' are *intertwined* if they lie on a common segment of X, or if they lie on segments s and s', respectively, where s and s' share a common endpoint b and $\tau(q, b) \leq \gamma$ and $\tau(q', b) \leq \gamma$.

DEFINITION 5. *For a PSLG X and a fixed $\gamma > 1$, the* local feature size $\mathrm{lfs}(p)$ *at a point $p \in \Omega$ is the radius (as measured by p) of the smallest elliptical disk (circular from p's perspective) centered at p that intersects two feature points of X that are not intertwined. In other words, there are non-intertwined feature points q and q' such that $d_p(q) = \mathrm{lfs}(p)$ and $d_p(q') \leq \mathrm{lfs}(p)$, but this is not true for any radius smaller than $\mathrm{lfs}(p)$. Let $\mathrm{lfs}_{\min} = \inf_{p \in \Omega} \mathrm{lfs}(p)$.*

This definition of local feature size is similar to Ruppert's, but it is adjusted to account for the fact that two segments that meet at an angle greater than $90°$, from the perspective of their shared endpoint, might meet at an angle of $1°$, from the perspective of points q and q' on each segment. In this case, q and q' are not intertwined, and the distance between them (as they measure it) influences the local feature size nearby. However, if $\mathrm{bdr}_{\min} > 0$, then $\mathrm{lfs}_{\min} > 0$.

For brevity, the following theorem applies only when the Voronoi refinement algorithm does not refine triangles for being too large. It is straightforward (if tedious) to adapt the proof to cover refinement of oversized triangles as well.

THEOREM 11. *Let $\theta_{\mathrm{bound}} < \arcsin \frac{1}{2\sqrt{2}}$ be a constant. Suppose a triangle t is considered to be* poorly shaped *(thus its circumcenter is a violator) if t has an angle less than θ_{bound}, as measured by an arbitrary point p in t. Let γ be the real root of*

$$\gamma^8 + 2\gamma^4 \sqrt{\gamma^2 + 1} \sin \theta_{\mathrm{bound}} - \gamma^4 - 1 = 0$$

that is greater than one. Suppose any two adjoining segments of X are separated by an angle of at least $2 \arcsin(\gamma^2/2)$, as measured from the perspective of the intersection point. (This angle is less than $79°$, and can be reduced arbitrarily close to $60°$ by choosing γ closer to one.) If $\mathrm{bdr}_{\min}(\gamma) > 0$, then the Voronoi refinement algorithm described in Section 9 generates a triangulation T wherein no triangle t has an angle less than θ_{bound} as measured by any point $p \in t$, and every pair of sites $u \neq v$ in T satisfies

$$d_u(v) \geq \min \left\{ \frac{\mathrm{lfs}_{\min}}{\gamma}, \frac{\mathrm{bdr}_{\min}(\gamma)}{(\gamma^5 + \gamma^3)\sqrt{\gamma^2 + 1}} \right\}.$$

If $\theta_{\mathrm{bound}} = 0°$—the minimum requirement for a geometrically valid triangulation—then $\gamma \doteq 1.1278$. This suggests that if neighboring sites measure lengths more than about 12% differently, that may be enough to trigger refinement to reduce the disparity, even if the sites are uniformly spaced. If $\theta_{\mathrm{bound}} = 10°$, $\gamma \doteq 1.0629$. θ_{bound} may be as high as about $20.7°$ (as it is for Ruppert's algorithm), but as θ_{bound} approaches this upper limit, the algorithm may have to refine the edges to very short lengths to keep the relative distortion locally small enough.

Our proof of Theorem 11 proceeds through many lemmas, in which all references to Voronoi cells refer to cells of the true anisotropic Voronoi diagram, even if the Voronoi refinement algorithm uses a loose Voronoi diagram instead. Lemma 20 shows that it does no harm for the algorithm to work with the latter.

LEMMA 12. *Let v be a site of a Voronoi diagram D, and let q be a point in* $\mathrm{Vor}(v)$. *For some $\gamma \geq 1$, suppose $d_v(q) \geq \mathrm{bdr}(v, \gamma)$. Then for any site w of D, $d(w, q) \geq \mathrm{bdr}_{\min}(\gamma)/\gamma$. Thus, inserting q into D creates no inter-site distance shorter than* $\mathrm{bdr}_{\min}(\gamma)/\gamma$.

PROOF. Suppose for the sake of contradiction that for some site w of D, $d(w, q) < \mathrm{bdr}_{\min}(\gamma)/\gamma$. Then either $d_w(q) < \mathrm{bdr}(w, \gamma)/\gamma$ or $d_q(w) < \mathrm{bdr}(q, \gamma)/\gamma$, so $\tau(w, q) \leq \gamma$. In the former case $d_w(q) < \mathrm{bdr}_{\min}(\gamma)/\gamma$, and in the latter case $d_w(q) \leq \gamma d_q(w) < \mathrm{bdr}_{\min}(\gamma)$.

Because q is in $\mathrm{Vor}(v)$, $d_v(q) \leq d_w(q) < \mathrm{bdr}_{\min}(\gamma)$, which contradicts the assumption that $d_v(q) \geq \mathrm{bdr}(v, \gamma)$. ■

LEMMA 13. *Let D be a Voronoi diagram, and let q be a point where no site of D lies. Let z be a point that would be in $\mathrm{Vor}(q)$ if q were inserted into D. For some $\gamma \geq 1$, suppose $d_q(z) \geq \mathrm{bdr}(q, \gamma)$. Then for any site w of D, $d(w, z) \geq \mathrm{bdr}_{\min}(\gamma)/\gamma$.*

PROOF. Essentially the same as the proof of Lemma 12. (Observe that because z would lie in $\mathrm{Vor}(q)$ if q were inserted, $d_q(z) \leq d_w(z)$ for any site w of D.) ■

LEMMA 14. *Let u and v be sites of a Voronoi diagram D. Let q be a point on the Voronoi 1-cell $\mathrm{Vor}(u) \cap \mathrm{Vor}(v)$. For some $\gamma \geq 1$, suppose $d_u(v) \geq \mathrm{bdr}(u, \gamma)$. Then for any site w of D, $d(w, q) \geq \mathrm{bdr}_{\min}(\gamma)/(\gamma^3 + \gamma)$.*

PROOF. If $d_u(q) \geq \mathrm{bdr}(u, \gamma)$ or $d_v(q) \geq \mathrm{bdr}(v, \gamma)$, the result follows by Lemma 12. Otherwise, by the definition of bounded distortion radius, $\tau(u, q) \leq \gamma$ and $\tau(v, q) \leq \gamma$, so $\tau(u, v) \leq \gamma^2$.

Suppose for the sake of contradiction that for some site w of D, $d(w, q) < \mathrm{bdr}_{\min}(\gamma)/(\gamma^3 + \gamma)$. Then either $d_w(q) < \mathrm{bdr}(w, \gamma)/(\gamma^3 + \gamma)$ or $d_q(w) < \mathrm{bdr}(q, \gamma)/(\gamma^3 + \gamma)$, so $\tau(w, q) \leq \gamma$. In the former case $d_w(q) < \mathrm{bdr}_{\min}(\gamma)/(\gamma^3 + \gamma)$, and in the latter case $d_w(q) \leq \gamma d_q(w) < \mathrm{bdr}_{\min}(\gamma)/(\gamma^2 + 1)$.

Because q is in $\mathrm{Vor}(u) \cap \mathrm{Vor}(v)$, $d_u(q) = d_v(q) \leq d_w(q) < \mathrm{bdr}_{\min}(\gamma)/(\gamma^2 + 1)$. Recall that $\tau(u, v) \leq \gamma^2$, so $d_u(v, q) \leq \gamma^2 d_v(q) < \gamma^2 \mathrm{bdr}_{\min}(\gamma)/(\gamma^2 + 1)$. By the triangle inequality, $d_u(v) \leq d_u(q) + d_u(v, q) < \mathrm{bdr}_{\min}(\gamma)$, which contradicts the assumption that $d_u(v) \geq \mathrm{bdr}(u, \gamma)$. The result follows. ■

LEMMA 15. *Let v be a site of a Voronoi diagram D, and let q be a point where no site of D lies. Let z be a point in $\mathrm{Vor}(v)$ that would be in $\mathrm{Vor}(q)$ if q were inserted into D. For some $\gamma \geq 1$, suppose $d_q(v) \geq \mathrm{bdr}(q, \gamma)$. Then for any site w of D, $d(w, z) \geq \mathrm{bdr}_{\min}(\gamma)/(\gamma^3 + \gamma)$.*

PROOF. Omitted. Similar to the proof of Lemma 14, but it uses Lemma 13 as well as Lemma 12. ■

LEMMA 16. *Let D be a Voronoi diagram, and let q be a point in Ω. For any $\gamma \geq 1$ and for every site w of D, $d_q(w) \geq \min\{d_w(q)/\gamma, \mathrm{bdr}(q, \gamma)\}$.*

PROOF. For each site w, either $d_q(w) \geq \mathrm{bdr}(q, \gamma)$ (satisfying the lemma) or $d_q(w) < \mathrm{bdr}(q, \gamma)$. In the latter case, $\tau(q, w) \leq \gamma$ by the definition of bdr and $d_q(w) \geq d_w(q)/\gamma$. ■

LEMMA 17. *Let X be a PSLG whose intersecting segments satisfy the angle condition specified in Theorem 11. Let s be a segment in X with endpoints a and b. Let D be a Voronoi diagram whose sites include the sites in X (including a and b). Let v be a site of D that encroaches upon s. Let z be a point in $\mathrm{Vor}(v) \cap s$. Let $m = \min\{d_a(v), d_b(v)\}$.*

Then for any site w of D, $d(w, z) \geq \min\{m, \mathrm{lfs}_{\min}/\gamma, \mathrm{bdr}_{\min}(\gamma)\}$. Thus, inserting z into D creates no inter-site distance shorter than $\min\{m, \mathrm{lfs}_{\min}/\gamma, \mathrm{bdr}_{\min}(\gamma)\}$.

PROOF. Either v is a site in X or v lies on a segment of X (otherwise v would not have been inserted). Suppose v and z are not intertwined. Then $d_v(z) \geq \mathrm{lfs}(v)$ by the definition of lfs. Because z is in $\mathrm{Vor}(v)$, $d_w(z) \geq d_v(z) \geq \mathrm{lfs}(v)$ for any site w of D. By Lemma 16, $d_z(w) \geq \min\{\mathrm{lfs}(v)/\gamma, \mathrm{bdr}(z, \gamma)\}$, so the lemma holds.

If v and z are intertwined, then v lies on a segment s' that adjoins s. Assume without loss of generality that the shared endpoint of s and s' is b; then $\tau(v, b) \leq \gamma$ by the definition of intertwined. We claim that $d_b(v, z) > \gamma^2 d_b(v)$. Suppose for the sake of contradiction that $d_b(v, z) \leq \gamma^2 d_b(v)$. Because z is in $\mathrm{Vor}(v)$, $d_b(z) \geq d_v(z) \geq d_b(v, z)/\gamma$, so $d_b(v, z) \leq \gamma d_b(z) < \gamma^2 d_b(z)$.

Let θ be the angle $\angle vbz$ as measured by b. Given the constraints $d_b(v, z) \leq \gamma^2 d_b(v)$ and $d_b(v, z) < \gamma^2 d_b(z)$, θ is maximized at $\theta = 2\arcsin(\gamma^2/2)$ when the constraints achieve equality. But the second constraint cannot achieve equality, and the angle condition is $\theta \geq 2\arcsin(\gamma^2/2)$, so the claim holds by contradiction.

Therefore, for any site w of D, $d_w(z) \geq d_v(z) \geq d_b(v, z)/\gamma > \gamma d_b(v) \geq \gamma m$. By Lemma 16, $d_z(w) \geq \min\{m, \mathrm{bdr}(z, \gamma)\}$. ■

LEMMA 18. *Let q be a Voronoi vertex that dualizes to a poorly shaped triangle $t = \triangle v_1 v_2 v_3$. Define γ as in Theorem 11. Suppose that $d_{v_i}(v_j) < \mathrm{bdr}(v_i, \gamma)$ for every $i, j \in \{1, 2, 3\}$. Let the* shortest edge length *ℓ of t be* $\min\{d(v_1, v_2), d(v_2, v_3), d(v_3, v_1)\}$.

Then for every site w of D, $d_w(q) > \gamma^2 \sqrt{\gamma^2 + 1}\ell$.

PROOF. For each vertex v_i of t, the distortion disk of v_i is large enough to enclose the other two vertices, so it encloses t. Therefore, $\tau(v_i, p) \leq \gamma$ for any vertex v_i of t and point p in t. Because t is poorly shaped, there is a point p in t from whose perspective t has an angle $\theta < \theta_{\mathrm{bound}}$. Let $r = d_{v_1}(q) = d_{v_2}(q) = d_{v_3}(q)$. Let $\beta = \gamma^2 \sqrt{\gamma^2 + 1}$. If $\beta \geq r/\ell$, then by Corollary 10, $2\gamma^2 \sin\theta \geq 1/\beta - (\gamma^2 - 1)\beta$. Substitution of β gives $\gamma^8 + 2\gamma^4 \sqrt{\gamma^2 + 1} \sin\theta - \gamma^4 - 1 \geq 0$. By assumption, $\gamma^8 + 2\gamma^4 \sqrt{\gamma^2 + 1} \sin\theta_{\mathrm{bound}} - \gamma^4 - 1 = 0$, so $\sin\theta \geq \sin\theta_{\mathrm{bound}}$, a contradiction; hence $\beta < r/\ell$.

For every site w of D, $d_w(q) \geq r > \beta\ell = \gamma^2 \sqrt{\gamma^2 + 1}\ell$. ■

LEMMA 19. *Let u and v be sites of a Voronoi diagram D. Let q be a wedge violator in $\mathrm{Vor}(u) \cap \mathrm{Vor}(v)$ that is outside* $\mathrm{wedge}(u, v)$. *Define γ as in Theorem 11. Then for every site w of D, $d_w(q) \geq \min\{\gamma^2 \sqrt{\gamma^2 + 1}\, d(u, v), \mathrm{bdr}_{\min}/(\gamma^3 + \gamma)\}$.*

PROOF. If $\tau(u, v) > \gamma$, then $d_v(q) \geq \mathrm{bdr}_{\min}/(\gamma^3 + \gamma)$ by Lemma 14. If $\tau(u, v) \leq \gamma$, assume q is outside $\mathrm{wedge}(u, v)$ on the side of u. From Lemma 5 we have $d_v(q) \geq d_u(v)/\sqrt{\gamma^2 - 1}$. By assumption, $\gamma^8 - \gamma^4 - 1 \leq 0$, so $d_u(v)/\sqrt{\gamma^2 - 1} \geq \gamma^2 \sqrt{\gamma^2 + 1} d_u(v) \geq \gamma^2 \sqrt{\gamma^2 + 1}\, d(u, v)$.

Because $q \in \mathrm{Vor}(v)$, for every site w of D, $d_w(q) \geq d_v(q) \geq \min\{\gamma^2 \sqrt{\gamma^2 + 1}\, d(u, v), \mathrm{bdr}_{\min}/(\gamma^3 + \gamma)\}$. ■

LEMMA 20. *Suppose the Voronoi refinement algorithm maintains a loose anisotropic Voronoi diagram instead of a true anisotropic Voronoi diagram. Let q be a point that the algorithm identifies as a violator, but is not really a violator. The bound of Lemma 19 applies to q (for some pair of sites u and v).*

PROOF. Let L be the loose Voronoi diagram maintained by the algorithm at the moment q is identified as a violator, and let D be the true anisotropic Voronoi diagram of the same sites. Because q is incorrectly identified as a violator, q lies in a 2-cell c of L owned by some site z whose Voronoi cell (in D) does not contain z.

Consider first the case where q lies in the interior of c. Because $q \notin \mathrm{Vor}(z)$, there must be some other site v for which $q \in \mathrm{Vor}(v)$.

The 2-cell of L owned by v includes every point of $\mathrm{Vor}(v)$ that is visible to v within $\mathrm{Vor}(v)$. Therefore, q is not visible to v within $\mathrm{Vor}(v)$. Let p be the point nearest v on the line segment qv such that the line segment qp lies entirely in $\mathrm{Vor}(v)$. (It is possible that $p = q$.) Let $u \neq v$ be a site such that $p \in \mathrm{Vor}(u)$ and some point of $\mathrm{Vor}(u)$ lies between p and v. By Lemma 3, p is outside $\mathrm{wedge}(u, v)$, so p is a wedge violator. By Lemma 19, $d_v(p) \geq \min\{\gamma^2\sqrt{\gamma^2+1}\, d(u, v), \mathrm{bdr}_{\min}/(\gamma^3 + \gamma)\}$. Because p lies between q and v, $d_v(q) \geq d_v(p)$. Because $q \in \mathrm{Vor}(v)$, for every site w of D, $d_w(q) \geq d_v(q)$. The result follows.

Now consider the case where q lies on the boundary of c. Let $c' = c \setminus \mathrm{Vor}(z)$, and observe that $q \in c'$. Because the bound holds for every point in the interior of c', it holds for every point on the boundary of c' by the continuity of distances. ■

LEMMA 21. *Let q be a site inserted at a violator by the Voronoi refinement algorithm. Let D be the Voronoi diagram just before q is inserted. Let $m = \min_{u\neq v} d_u(v)$ where the sites u and v are chosen from among all sites of D. For every site w of D, $d(q, w) \geq \min\{\gamma\sqrt{\gamma^2+1}m, \mathrm{bdr}_{\min}(\gamma)/(\gamma^4 + \gamma^2)\}$. Thus, the insertion of q into D creates no inter-site distance shorter than that.*

PROOF. The violator q does not encroach upon any subsegment (otherwise it would not be inserted). If q is a wedge violator, the result follows from Lemmas 19 and 16. If q is a violator by mistaken identity, the result follows from Lemmas 20 and 16. Otherwise, q dualizes to a poorly shaped triangle $t = \triangle v_1v_2v_3$. If $d_{v_i}(v_j) \geq \mathrm{bdr}(v_i, \gamma)$ for any $i, j \in \{1, 2, 3\}$, then by Lemma 14, $d(q, w) \geq \mathrm{bdr}_{\min}(\gamma)/(\gamma^3 + \gamma)$ as claimed. Hence, assume that $d_{v_i}(v_j) < \mathrm{bdr}(v_i, \gamma)$ for every $i, j \in \{1, 2, 3\}$. By Lemma 18, $d_w(q) \geq \gamma^2\sqrt{\gamma^2+1}m$. The result follows from Lemma 16. ■

LEMMA 22. *Let s be a segment with endpoints a and b in a PSLG X. Let z be the "midpoint" of s, defined so that $d_a(z) = d_b(z)$. Let D be a Voronoi diagram whose sites include the sites in X (including a and b). Let q be a point such that $\mathrm{Vor}(q)$ would contain z if q were inserted into D (thus q would encroach upon s).*

Suppose $\tau(a, q) \leq \gamma$ and $\tau(b, q) \leq \gamma$ for some $\gamma \geq 1$. Then $d_a(z) = d_b(z) \geq \min\{d_a(q), d_b(q)\}/(\gamma\sqrt{\gamma^2+1})$.

PROOF. Consider a line l through z that is perpendicular to s from q's point of view. Specifically, l is composed of the points p that satisfy $d_q(b, z)^2 + d_q(z, p)^2 = d_q(b, p)^2$ (Pythagoras' Law). Suppose without loss of generality that q lies on the same side of l as b, or q lies on l. Then $d_q(b, z)^2 + d_q(z)^2 \geq d_q(b)^2$. Because $\tau(b, q) \leq \gamma$, it follows that $\gamma^2 d_b(z)^2 + d_q(z)^2 \geq d_b(q)^2/\gamma^2$.

$\mathrm{Vor}(q)$ would contain z if q were inserted into D, so $d_q(z) \leq d_b(z)$. Therefore, $(\gamma^2 + 1)d_b(z)^2 \geq d_b(q)^2/\gamma^2$, from which the result follows. ■

LEMMA 23. *Let z be a site inserted on a segment s because a violator q (not inserted) encroaches upon s. Let D be the Voronoi diagram just before z is inserted. Let $m = \min_{u\neq v} d_u(v)$ where the sites u and v are chosen from among all sites of D. Define γ as in Theorem 11. For every site w of D, $d(z, w) \geq \min\{m, \mathrm{bdr}_{\min}(\gamma)/((\gamma^5 + \gamma^3)\sqrt{\gamma^2+1})\}$.*

PROOF. The segment s intersects only two Voronoi cells, namely $\mathrm{Vor}(a)$ and $\mathrm{Vor}(b)$ where a and b are the endpoints of s. Otherwise, the algorithm would have split s before trying to insert q. There are two cases to consider. In Case A, z lies on the Voronoi curve $\mathrm{Vor}(a) \cap \mathrm{Vor}(b)$. In Case B, z lies in just one Voronoi cell—without loss of generality, say $\mathrm{Vor}(b)$—because if q were inserted, $\mathrm{Vor}(q)$ would not intersect $s \cap \mathrm{Vor}(a)$.

If $d_q(b) \geq \mathrm{bdr}(q, \gamma)$, the result follows by Lemma 15. Hence, assume $\tau(q, b) \leq \gamma$. Similarly, in Case A the result follows if $d_q(a) \geq \mathrm{bdr}(q, \gamma)$, so assume in that case that $\tau(q, a) \leq \gamma$.

Let $v_1, \ldots, v_k$ be the sites whose Voronoi cells contain q. In Case A, by Lemma 22, $d_b(z) \geq \min\{d_a(q), d_b(q)\}/(\gamma\sqrt{\gamma^2+1})$. Because q lies in $\mathrm{Vor}(v_i)$ for each $i \in [1, k]$ and z lies in $\mathrm{Vor}(b)$, for any site w of D we have $d_w(z) \geq d_b(z) \geq \min\{d_a(q), d_b(q)\}/(\gamma\sqrt{\gamma^2+1}) \geq d_{v_i}(q)/(\gamma\sqrt{\gamma^2+1})$.

A similar result holds in Case B. Let D_q be the Voronoi diagram formed by inserting q into D (even though the algorithm does not insert q). In D_q, there are one or two intersection points $s \cap \mathrm{Vor}(b) \cap \mathrm{Vor}(q)$, and z is the one furthest from b. Although z lies in $\mathrm{Vor}(b)$ in D_q, either z is not visible from b within $\mathrm{Vor}(b)$, or z lies where $\mathrm{Vor}(q)$ tangentially intersects s at a single point. By Lemma 3, z lies outside $\mathrm{wedge}(b, q)$ on q's side of the wedge. From Lemma 5 we have $d_q(z) \geq d_q(b)/\sqrt{\gamma^2-1}$. In D, q lies in $\mathrm{Vor}(v_i)$ for each $i \in [1, k]$, and in D_q, z lies in $\mathrm{Vor}(q)$, so for any site w of D we have $d_w(z) \geq d_q(z) \geq d_q(b)/\sqrt{\gamma^2-1} \geq d_b(q)/(\gamma\sqrt{\gamma^2-1}) \geq d_{v_i}(q)/(\gamma\sqrt{\gamma^2-1})$.

In either Case A or B, $d_w(z) \geq d_{v_i}(q)/(\gamma\sqrt{\gamma^2+1})$ for any site w of D and any $i \in [1, k]$. Suppose that $d_{v_i}(q) \geq \mathrm{bdr}(v_i, \gamma)$ for some $i \in [1, k]$. Then $d_w(z) \geq \mathrm{bdr}(v_i, \gamma)/(\gamma\sqrt{\gamma^2+1})$ and, by Lemma 16, $d_z(w) \geq \mathrm{bdr}_{\min}(\gamma)/(\gamma^2\sqrt{\gamma^2+1})$ for any site w of D, so the result follows.

If the supposition is not true for any $i \in [1, k]$, then $d_{v_i}(q) < \mathrm{bdr}(v_i, \gamma)$ and thus $\tau(v_i, q) \leq \gamma$ for every $i \in [1, k]$. This implies that $\tau(v_i, v_j) \leq \gamma^2$ for any $i, j \in [1, k]$.

Suppose that there are two sites v_i and v_j for which $i, j \in [1, k]$ and $d_{v_i}(v_j) \geq \mathrm{bdr}(v_i, \gamma)$. By the triangle inequality, $d_{v_i}(v_j) \leq d_{v_i}(q) + d_{v_i}(q, v_j) \leq d_{v_i}(q) + \gamma^2 d_{v_j}(q) = (\gamma^2 + 1)d_{v_i}(q) \leq (\gamma^3 + \gamma)\sqrt{\gamma^2+1}d_w(z)$ for any site w of D. Thus $d_w(z) \geq \mathrm{bdr}(v_i, \gamma)/((\gamma^3 + \gamma)\sqrt{\gamma^2+1})$ and, by Lemma 16, $d_z(w) \geq \mathrm{bdr}_{\min}(\gamma)/((\gamma^4 + \gamma^2)\sqrt{\gamma^2+1})$ for each site w, so the result follows.

If the supposition is not true, then $d_{v_i}(v_j) < \mathrm{bdr}(v_i, \gamma)$ for every pair of sites v_i and v_j whose Voronoi cells contain q. One of the preconditions for Lemma 18 is thus satisfied.

Depending on what type of violator q is, Lemma 18, 19, or 20 applies, showing that $d_w(q) \geq \min\{\gamma^2\sqrt{\gamma^2+1}m, \mathrm{bdr}_{\min}(\gamma)/(\gamma^3 + \gamma)\}$ for every site w of D. It follows that $d_w(z) \geq d_{v_i}(q)/(\gamma\sqrt{\gamma^2+1}) \geq \min\{\gamma m, \mathrm{bdr}_{\min}(\gamma)/((\gamma^4 + \gamma^2)\sqrt{\gamma^2+1})\}$ and, by Lemma 16, $d_z(w) \geq \min\{m, \mathrm{bdr}_{\min}(\gamma)/((\gamma^5 + \gamma^3)\cdot\sqrt{\gamma^2+1})\}$ for each site w of D as claimed. ■

PROOF OF THEOREM 11. The anisotropic Voronoi diagram of the sites in X has no inter-site distance less than $\mathrm{lfs}_{\min}$, by the definition of $\mathrm{lfs}_{\min}$. By Lemmas 17, 21, and 23, every inter-site distance created by any site the algorithm inserts is at least as great as either $\mathrm{lfs}_{\min}/\gamma$, $\mathrm{bdr}_{\min}(\gamma)/((\gamma^5 + \gamma^3)\sqrt{\gamma^2+1})$, or the shortest inter-site distance existing prior to the site insertion.

By induction on the sequence of site insertions, for any two sites v and w that are ever inserted by the Voronoi refinement algorithm, $d(v, w) \geq \mu = \min\{\mathrm{lfs}_{\min}/\gamma, \mathrm{bdr}_{\min}(\gamma)/((\gamma^5+\gamma^3)\sqrt{\gamma^2+1})\}$.

Let $\lambda_{\max}$ be an upper bound on the largest eigenvalue of M over the domain Ω. For any two sites v and w, $\mu \leq d_v(w) = \|F_v(v - w)\|_2 \leq \|F_v\|_2\|v - w\|_2 = (\max_x x^T M_v x/x^T x)^{1/2}\|v - w\|_2 \leq \sqrt{\lambda_{\max}}\, \|v - w\|_2$. Therefore, measured in physical space, $\|v - w\|_2 \geq \mu/\sqrt{\lambda_{\max}}$. Imagine that a disk of radius $\mu/(2\sqrt{\lambda_{\max}})$ is centered at each site. The interiors of these disks do not intersect. Only a finite number of such nonoverlapping disks can be packed so their centers lie in Ω, therefore the algorithm terminates. ■

11. Conclusions

There are several clear directions for improvement and extension of our results. We suspect it is straightforward (albeit messy) to prove that our Voronoi refinement algorithm generates *graded* meshes in which edge lengths are dictated primarily by local features—the mesh is not refined to very small triangles in one region because of very small features far away. Ruppert [13] proves such a result for isotropic meshing. However, the anisotropic grading result would be weaker than Ruppert's isotropic result because a site whose metric tensor is small sees faraway elements as being close by, and therefore its influence on them is less attenuated. Specifically, a site with a small metric tensor has a wide paraboloid which can create orphans at a great distance, thereby causing refinement in distant locations to remove the orphans. These effects might be mitigated by simply removing the orphans without inserting new sites (exploiting the flexibility of loose Voronoi diagrams), or by resizing the metric tensors (and paraboloids) of problem sites.

Our Voronoi refinement algorithm, like Ruppert's Delaunay refinement algorithm, recovers segments of the PSLG by splitting them. There are many advantages to using a constrained Delaunay triangulation (CDT) instead. But can CDTs be made anisotropic? Seidel [14] proposes the *extended Voronoi diagram*, which dualizes the (isotropic) CDT by extending the plane so that each segment serves as a gateway to another "sheet" of paper, inhabited by some of the faces of the Voronoi diagram. As a natural generalization, we suggest extended anisotropic Voronoi diagrams, which dualize to constrained triangulations under the same conditions for which anisotropic Voronoi diagrams dualize to triangulations.

For practical mesh generation, it is attractive to maintain the restricted Voronoi diagram $D|_\Omega$ rather than D because the rest of D is not needed. We could also redefine loose site insertion so that the depth-first search stops at the boundary of Ω, but the diagram created this way is not necessarily $L|_\Omega$. It is an open problem whether an analog of Lemma 20 holds for this weaker site insertion method. More generally, if loose insertion and extended anisotropic Voronoi diagrams can be made to work together, the combination might enjoy the speed of loose insertion and the benefits of CDTs.

In three dimensions or more, we cannot prove that our Voronoi refinement algorithm creates a triangulation with no geometrically inverted elements, for the same reason that it is difficult to prove that isotropic three-dimensional mesh generation algorithms create no sliver tetrahedra. Delaunay refinement algorithms for tetrahedral mesh generation are nevertheless effective in practice at removing slivers [17], and we believe the same is true of three-dimensional Voronoi refinement in the anisotropic case.

12. References

[1] Pankaj K. Agarwal, Boris Aronov, and Micha Sharir. *Computing Lower Envelopes in Four Dimensions with Applications*. Proceedings of the Tenth Annual Symposium on Computational Geometry, pages 348–358. Association for Computing Machinery, 1994.

[2] Pankaj K. Agarwal, Otfried Schwarzkopf, and Micha Sharir. *Overlay of Lower Envelopes and its Applications*. Technical Report CS-1994-18, Department of Computer Science, Duke University, 1994.

[3] Thomas Apel. *Anisotropic Finite Elements: Local Estimates and Applications*. B. G. Teubner Verlag, Stuttgart, 1999.

[4] Franz Aurenhammer and Herbert Edelsbrunner. *An Optimal Algorithm for Constructing the Weighted Voronoi Diagram*. Pattern Recognition **17**:251–257, 1984.

[5] Frank J. Bossen and Paul S. Heckbert. *A Pliant Method for Anisotropic Mesh Generation*. Fifth International Meshing Roundtable, pages 63–74, October 1996.

[6] Adrian Bowyer. *Computing Dirichlet Tessellations*. Computer Journal **24**(2):162–166, 1981.

[7] Paul-Louis George and Houman Borouchaki. *Delaunay Triangulation and Meshing: Application to Finite Elements*. Hermès, Paris, 1998.

[8] Dan Halperin and Micha Sharir. *New Bounds for Lower Envelopes in Three Dimensions, with Applications to Visibility in Terrains*. Discrete & Computational Geometry **12**:313–326, 1994.

[9] Kenneth E. Jansen, Mark S. Shephard, and Mark W. Beall. *On Anisotropic Mesh Generation and Quality Control in Complex Flow Problems*. Tenth International Meshing Roundtable (Newport Beach, California), pages 341–349. Sandia National Laboratories, October 2001.

[10] Greg Leibon and David Letscher. *Delaunay Triangulations and Voronoi Diagrams for Riemannian Manifolds*. Proceedings of the Sixteenth Annual Symposium on Computational Geometry (Hong Kong), pages 341–349. Association for Computing Machinery, June 2000.

[11] Xiang-Yang Li, Shang-Hua Teng, and Alper Üngör. *Biting Ellipses to Generate Anisotropic Mesh*. Eighth International Meshing Roundtable (Lake Tahoe, California), pages 97–108, October 1999.

[12] Shmuel Rippa. *Long and Thin Triangles Can Be Good for Linear Interpolation*. SIAM Journal on Numerical Analysis **29**(1):257–270, February 1992.

[13] Jim Ruppert. *A Delaunay Refinement Algorithm for Quality 2-Dimensional Mesh Generation*. Journal of Algorithms **18**(3):548–585, May 1995.

[14] Raimund Seidel. *Constrained Delaunay Triangulations and Voronoi Diagrams with Obstacles*. 1978–1988 Ten Years IIG (H. S. Poingratz and W. Schinnerl, editors), pages 178–191. Institute for Information Processing, Graz University of Technology, 1988.

[15] Micha Sharir. *Almost Tight Bounds for Lower Envelopes in Higher Dimensions*. Discrete & Computational Geometry **12**:327–345, 1994.

[16] Micha Sharir and Pankaj K. Agarwal. *Davenport-Schinzel Sequences and their Geometric Applications*. Cambridge University Press, New York, 1995.

[17] Jonathan Richard Shewchuk. *Tetrahedral Mesh Generation by Delaunay Refinement*. Proceedings of the Fourteenth Annual Symposium on Computational Geometry (Minneapolis, Minnesota), pages 86–95. Association for Computing Machinery, June 1998.

[18] ———. *What is a Good Linear Finite Element? Interpolation, Conditioning, Anisotropy, and Quality Measures*. Manuscript, 2002.

[19] Kenji Shimada, Atsushi Yamada, and Takayuki Itoh. *Anisotropic Triangular Meshing of Parametric Surfaces via Close Packing of Ellipsoidal Bubbles*. Sixth International Meshing Roundtable (Park City, Utah), pages 375–390. Sandia National Laboratories, October 1997.

[20] David F. Watson. *Computing the n-dimensional Delaunay Tessellation with Application to Voronoi Polytopes*. Computer Journal **24**(2):167–172, 1981.