Fat Voronoi Diagrams*

Gary L. Miller† Todd Phillips‡ Donald R. Sheehy§

1 Why Fat Voronoi Diagrams?

Voronoi refinement is a powerful tool for efficiently generating meshes for finite element simulation. The classic definition of quality in a mesh can be achieved by bounding the *aspect ratio* of the Voronoi cells measured as the ratio of the circumscribing and inscribing radii as measured from the site. There are tight upper and lower bounds on the number of extra points needed to achieve such a Voronoi diagram. The use of good aspect ratio Voronoi diagrams is central to both quadtree methods [1][1] and Voronoi refinement algorithms [3].

Unfortunately, bounding the aspect ratio in this way is often an overkill, with lower bounds on the size and runtime that depend on the *spread* of the input set, a geometric quantity that may be unbounded in n. In this paper, we give a relaxed definition of Voronoi cell quality called *fatness* that captures many of the nice properties of the old definition without being subject to the lower bounds on the size. We give upper and lower bounds on the complexity of such Voronoi diagrams and provide an algorithm to generate such a Voronoi diagram with only a linear number of extra points. In future work we hope to understand fat Voronoi diagrams well enough to design the next generation meshing algorithm with them.

The first and simplest question that arises in this area is whether or not a cell in a fat Voronoi diagram can have an unbounded number of neighbors. We prove that this is not possible for fat Voronoi diagrams in the plane and conjecture that similar bounds hold in higher dimensions. As Figure 1 demonstrates, this is a property peculiar to fat Voronoi diagrams; it holds neither for general fat complexes nor for weighted Voronoi diagrams.

*This work was partially supported by the National Science Foundation under grant number CCF-0635257.

†Department of Computer Science, Carnegie Mellon University, glmiller@cs.cmu.edu

‡Department of Computer Science, Carnegie Mellon University, tp517@cs.cmu.edu

§Department of Computer Science, Carnegie Mellon University, dsheehy@cs.cmu.edu

[1] The corners of a balanced quadtree form a good aspect ratio Voronoi diagram.

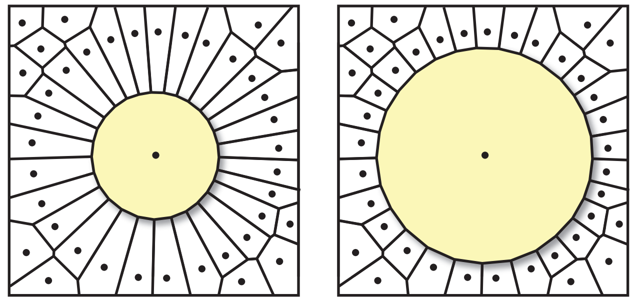

Figure 1: Right: A fat Voronoi cell can have arbitrarily many neighbors but at least some of them must be skinny. **Left:** If weights are allowed on the vertices, the corresponding weighted Voronoi diagram can be fat and yet have unbounded maximum degree.

2 Definitions and Notations

Let M be a finite set of points in $\mathbb{R}^d$ that we call *vertices*. We assume there is some compact, convex bounding region $R \subset \mathbb{R}^d$ that contains all of M. The *Voronoi cell* of a vertex $v \in M$ is the set of points in R for which v is the nearest neighbor in M, and is denoted $\mathrm{Vor}(v)$. The *Voronoi diagram* of M is the cell complex formed by the Voronoi cells of M, the bounding region R, and the corresponding intersections. The cells are Voronoi diagram are called *faces*.

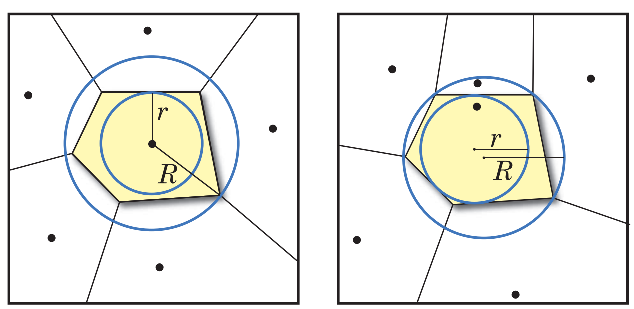

Figure 2: The cell on the left has good aspect ratio. The cell on the right is fat but it does not have good aspect ratio.

The dual complex to the Voronoi diagram is the *Delaunay triangulation*. For points in general position, the Delaunay triangulation is a simplicial complex. The set of simplices incident to a vertex p is called the *star* of p and the boundary of the star of p is called the *link* of p.

The *in-ball* of $\mathrm{Vor}(v)$ is the largest inscribed ball and its radius is r_v. The *out-ball* of $\mathrm{Vor}(v)$ is the smallest circumscribed ball and its radius is R_v. A Voronoi cell, $\mathrm{Vor}(v)$, is τ-fat if $\frac{R_v}{r_v} \le \tau$. The classic definition of a *good aspect ratio* Voronoi cell, in which the in-ball and out-ball are centered at v, is a special case of fatness, but fat cells do not necessarily have good aspect ratio (see Figure 2).

3 Global Upper Bounds

Generic fat complexes, which are not necessarily Voronoi diagrams, are well studied in the literature. There are several definitions of "fatness", but in the case of convex cells, they are equivalent up to constants. The only relevant difference is whether the definition uses a ratio of one-dimensional values (such as radii in our definition), or a d-dimensional ratio of volumes. The former definition makes explicit the exponential dependence on d.

Global upper bounds on complexity exist for general fat complexes. The classic approach to bounding fat complexes is to bound the number of larger neighbors of any cell [5]. This bounds the total number of neighbor relations in the complex, thus bounding above the total number of $d-1$-faces of a τ-fat complex by $(4\tau)^d n$. This bounds the average number of neighbors per vertex by $(4\tau)^d$. In the special case of fat Voronoi diagrams, these arguments can be extended to count the total count of all faces of all dimensions, giving an upper bound of $(4\tau)^{d^2} n$. This approach has been shown to work in [2], although with less attention to the constant depending on dimension.

We wish to strengthen these global bounds by showing that in the special case of a Voronoi complex, the size of a neighborhood is locally bounded.

4 Local Upper Bounds

The upper bounds of the previous section rely only on the fatness of the complex. They do not use any properties intrinsic to Voronoi diagrams. A complex may be fat and yet have unbounded degree. Thus, the guarantees are global rather than local. In the plane, it is possible to prove a local bound, namely that every cell has at most a constant number of neighbors.

Let $\mathrm{Vor}(P)$ be a fat Voronoi diagram. For any $p \in P$, let Q_p be the set of Delaunay neighbors of p. That is, for each $q \in Q_p$, there is a clockwise oriented Voronoi edge (u, v) dual to $\overline{pq}$ in $\mathrm{Vor}(P)$. For input points in general position, there is a unique edge e_u (e_v) in the Voronoi diagram that emanates from u (respectively v) but is not a boundary edge of $\mathrm{Vor}(p)$. Let ℓ_u and ℓ_v be the lines containing e_u and e_v respectively. We define two relevant angles:

$\theta_q =$ the angle between ℓ_u and ℓ_v.
$\sigma_q = \angle upv$.

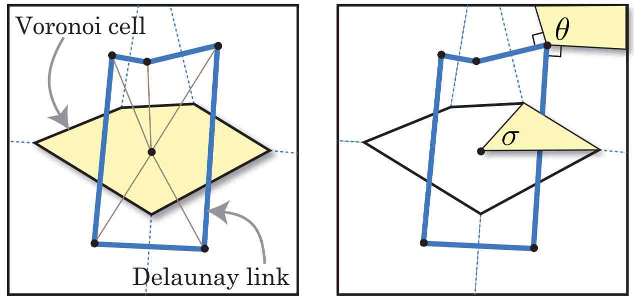

Figure 3: The Voronoi cell and the dual Delaunay link. There is a natural tradeoff between σ and θ. Although, the Voronoi cell is convex, the Delaunay link may not be.

We observe the following fact about these angles.

$$2\pi = \sum_{q \in Q_p} \theta_q = \sum_{q \in Q_p} \sigma_q. \tag{1}$$

This gives us two ways to measure the angles, one from the perspective of the Delaunay triangulation and the other from the perspective of the Voronoi diagram. Ultimately, we will show that at least one of σ_q or θ_q must be large for each neighbor q. This will allow us the bound the size of Q_p.

Let p be any vertex of P and let q be any neighbor in Q_p. To simplify notation, let $\sigma = \sigma_q$ and let $\theta = \theta_q$. The following three lemmas capture the relevant tradeoffs between σ and θ. The proofs are omitted for space.

Lemma 1. *If* $\theta \ge 0$ *then* $\max\{\sigma, \theta\} \ge \arcsin \frac{1}{4\tau}$.

Lemma 2. *If* $\theta < 0$ *then* $\sigma \ge \arcsin \frac{1}{\tau}$.

Lemma 3. *If* $\theta < 0$ *then* $\sigma + \theta \ge 0$.

Theorem 4. *If* $\mathrm{Vor}(M)$ *is* τ*-fat then for all* $p \in M$, $|Q_p| \le \frac{6\pi}{\arcsin \frac{1}{4\tau}}$.

Proof. Partition Q_p based on the magnitude of θ into $Q^- = \{q \in Q_p : \theta_q < 0\}$, $Q^+ = Q_p \backslash Q^-$. We will bound the size of each set individually. Using Lemma 2, we derive the following.

$$2\pi = \sum_{q \in Q_p} \sigma_q \ge \sum_{q \in Q^-} \sigma_q \ge |Q^-| \arcsin \frac{1}{\tau}.$$

This implies that $|Q^-| \le \frac{2\pi}{\arcsin \frac{1}{\tau}}$. We now bound $|Q^+|$

as follows.

$$\begin{aligned}
4\pi &= \sum_{q\in Q_p} \sigma_q + \theta_q && \text{[by Equation (1)]}\\
&\geq \sum_{q\in Q^+} \sigma_q + \theta_q && \text{[by Lemma 3]}\\
&\geq \sum_{q\in Q^+} \max\{\sigma_q, \theta_q\} && \text{[because } \sigma_q, \theta_q \geq 0\text{]}\\
&\geq |Q^+| \arcsin \frac{1}{4\tau} && \text{[by Lemma 1].}
\end{aligned}$$

This implies that $|Q^+| \leq \frac{4\pi}{\arcsin\frac{1}{4\tau}}$. So we see that

$$|Q_p| = |Q^-| + |Q^+| \leq \frac{2\pi}{\arcsin\frac{1}{\tau}} + \frac{4\pi}{\arcsin\frac{1}{4\tau}} \leq \frac{6\pi}{\arcsin\frac{1}{4\tau}}.$$

□

This result for planar Voronoi diagrams leads us to the following conjecture.

Conjecture 1 (The Fat Voronoi Conjecture)**.** *If* Vor(M) *is τ-fat then every vertex p in M has at most $2^{O(d)}$ neighbors, where τ is a constant independent of d.*

The Fat Voronoi Conjecture does not imply an asymptotic improvement on the bound on the total number of faces of the Voronoi diagram. For example, it is easy to see that the total number of faces can still be $2^{O(d^2)}$. However, we further conjecture that the situation is not as bad as that.

Conjecture 2 (The Strong Fat Voronoi Conjecture)**.** *If* Vor(M) *is τ-fat for then for every vertex p in M,* Vor(p) *has at most $2^{O(d \log d)}$ faces, where τ is a constant independent of d.*

The Strong Fat Voronoi Conjecture would imply the lower bound we derive in Theorem 5 is tight for constant values of the fatness parameter, τ.

5 Lower Bounds

In this section, we prove a nontrivial lower bound on the number of faces of a fat Voronoi cell. For constant τ, this bound is $2^{\Omega(d \log d)}$, which is somewhat surprising given that the integer lattice $\mathbb{Z}^d$ has Voronoi cells with only 2^d faces. This discrepancy arises because cubes are not fat as their dimension increases; a d-cube has fatness $\sqrt{d}$. The following theorem captures exactly this tradeoff in the lower bound.

Theorem 5. *If* Vor(M) *is a τ-fat Voronoi diagram, then it has at least $2^{\Omega(d)}\left(\frac{\sqrt{d}}{\tau}\right)^d n$ faces.*

Proof. For any vertex $p \in M$, let $s_1 \ldots s_k$ be the simplices of Del(M) with a vertex at p. For any set of vectors A let cone(A) denote the non-negative linear combinations of vectors of A. We recall that the polar A° of a set $A \subset \mathbb{R}^d$ is defined as.

$$A^\circ = \{y \in R^d : y \cdot a \leq 1 \text{ for all } a \in A\}.$$

Let C_s be cone$(\{p - q : q \in S\})$. For any simplex s there is a dual circumcenter x at a corner of Vor(p). Let C_x be $\{y - x : |y - p| \leq |y - q| \text{ for all } q \in s\}$. Note that $C_s = C_x^\circ$.

The Voronoi cell Vor(p) is fat and thus contains an in-ball b. If we let B_x denote cone$(\{y - x : y \in b\})$, then we have $B_x \subset C_x$. Recall that polarity reverses containment, so $C_x^\circ \subset B^\circ$. Letting B_s be the cone polar to B_x, it follows that

$$C_s = C_x^\circ \subset B_x^\circ = B_s.$$

Let θ_x and θ_s be the half-angles of the circular cones B_x and B_s respectively. Because the cones are polar to each other, $\theta_x + \theta_s = \frac{\pi}{2}$. Because Vor$(p)$ is τ-fat, $\theta_x \geq \arcsin\frac{1}{2\tau}$. Otherwise, b would not fit in C_x. So,

$$\cos\theta_s \geq \cos\left(\frac{\pi}{2} - \arcsin\frac{1}{2\tau}\right) = \frac{1}{2\tau}.$$

Let $\mathbb{B}$ be the unit ball centered at the origin. We can now pack the simplicial cones C_{s_i} restricted to $\mathbb{B}$ and apply Lemma 6.

$$\Gamma_d = \sum_{i=1}^{k} \text{vol}(C_{s_i} \cap \mathbb{B}) \leq \frac{k\Gamma_d}{2^{O(d)}\left(\frac{\sqrt{d}}{\tau}\right)^d}. \tag{2}$$

It follows that k, the number of simplices at p, is at least $2^{\Omega(d)}\left(\frac{\sqrt{d}}{\tau}\right)^d$. □

Lemma 6. *For all simplices s with a vertex at p,*

$$\text{vol}(C_s \cap \mathbb{B}) \leq \frac{\Gamma_d}{2^{O(d)}\left(\frac{\sqrt{d}}{\tau}\right)^d}.$$

Proof. Let H_1 and H_2 be halfspaces normal to the axis of B_s at distances $\frac{1}{2\tau}$ and 1 respectively. Since $\mathbb{B} \subset H_2$, it follows that

$$\text{vol}(C_s \cap \mathbb{B}) \leq \text{vol}(C_s \cap H_2). \tag{3}$$

The two sets $C_s \cap H_2$ and $B_s \cap H_2$ are each the convex closure of a base and a vertex at the origin. Moreover, their bases both lie in the boundary of H_2 so the ratio of their volumes is equal to the ratio of their bases. Since the base of $C_s \cap H_s$ is a $d-1$-simplex contained in the $d-1$-dimensional ball that is the base of $B_s \cap H_2$, the ratio of their volumes is bounded using Stirling's approximation as follows.

$$\frac{\mathrm{vol}(C_s \cap H_2)}{\mathrm{vol}(B_s \cap H_2)} \leq \frac{\mathrm{vol}(\mathbb{S})}{\Gamma_{d-1}} \leq \frac{1}{2^{O(d)} d^{\frac{d}{2}}}, \tag{4}$$

where $\mathbb{S}$ represents the regular $d-1$-simplex inscribed in the $d-1$-dimensional unit ball. By scaling $B_S \cap H_2$ down by a factor of 2τ, we get $B_S \cap H_1$, and therefore

$$\mathrm{vol}(B_S \cap H_2) = (2\tau)^d \mathrm{vol}(B_S \cap H_1). \tag{5}$$

Since $B_S \cap H_1 \subset \mathbb{B}$, we have that

$$\mathrm{vol}(B_S \cap H_1) < \Gamma_d. \tag{6}$$

Together, Equations (3), (4), (5), and (6) imply the statement of the lemma. □

6 A Fat Voronoi Algorithm

In previous work, we showed how to construct a linear size superset of a point set whose Delaunay triangulation has linear size [4]. It is possible to modify that algorithm to produce a superset whose Voronoi diagram is fat. The key to linear size meshing is to decompose the point set into so-called *well-paced* sets. These sets are ordered so that each point is near a face of the Voronoi diagram of its predecessors. A well-paced set of points can be extended to a good aspect ratio Voronoi diagram (in classic sense) using only a linear number of extra points.

Since not all point sets are well-paced, it is necessary to divide up space into a hierarchy of regions so that at any level, the points appear well-paced and a lower level looks like a single point.

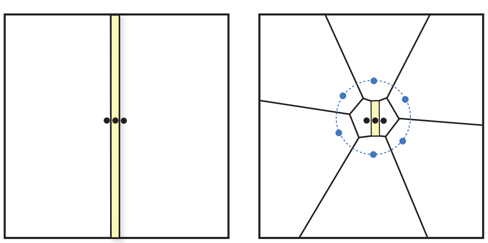

Figure 4: Long, skinny Voronoi cells are truncated by an appropriately size bounding cage.

Consider an ordering $(p_1, \ldots, p_n)$ on the input set P. Let $P_i = \{p_1, \ldots p_i\}$ be the ith prefix. We can define two related distances on the prefixes by taking the distance to the nearest and second nearest neighbor. Formally, $d_i(x) = \min p_j \in P_i |x - p_j|$ and $f_i(x) = \min p_j, p_k \in P_i \max\{|x-p_j|, |x-p_k\}$. The set P is θ-well-paced with respect to the given ordering and the bounding region R if $\frac{d_i(p_{i+1})}{f_i(p_{i+1})} > \theta$ for all $i = 2 \ldots n-1$ and $\frac{|p_1 - p_2|}{\mathrm{radius}(R)} > \theta$.

The FatVoronoi algorithm works by finding a maximal θ-well-paced subset S. A standard Voronoi refinement algorithm will only produce linearly many Steiner points for a θ-well-paced input. All input vertices in $P \setminus S$ must be significantly closer to one point of S than any other, for otherwise they would be θ-well-paced. A bounding *cage* composed of a constant number of points is placed around any vertex of S that is the nearest neighbor of a point in $P \setminus S$. These cages guarantee that non-fat cells of $\mathrm{Vor}(P)$ do not "escape" (see Figure 4). Moreover, the new vertices added for the cage all have Fat Voronoi cells. The cages and the input vertices they contain are handled recursively and the final output is the Voronoi diagram of the union of the outputs for each recursive call. In the end, we achieve the following Theorem.

Theorem 7. *Given n points $P \subset \mathbb{R}^d$, the* FatVoronoi *algorithm can produce a superset M of P of size $2^{O(d)}n$ such that $\mathrm{Vor}(M)$ is $O(1)$-fat in time $O_d(n^2)$ time.*

7 Remarks

There is at least some bound on the number of neighbors of a Fat Voronoi diagram in general dimension. Recall the classic definition of the *aspect-ratio* ρ of a $\mathrm{Vor}(v)$ as the ratio of the radii of the smallest out-ball to largest in-ball, where the balls must be *centered* at v. The packing arguments bounding the number of larger neighbors[5] can be extended to the following:

Lemma 8. *If $\mathrm{Vor}(P)$ is τ-fat and $\mathrm{Vor}(v)$ has aspect ratio ρ, then v has at most $2^{O(d)} \log \rho$ neighbors.*

While this lemma is much weaker than Conjecture 1, it may at least provides an approach to a more generic high dimensional argument.

References

[1] M. Bern, D. Eppstein, and J. R. Gilbert. Provably Good Mesh Generation. *Journal of Computer and System Sciences*, 48(3):384–409, June 1994.

[2] C. Gray. Nearest-neighbor queries with well-spaced points. *International Symposium on Voronoi Diagrams in Science and Engineering*, 0:42–49, 2010.

[3] B. Hudson, G. Miller, and T. Phillips. Sparse Voronoi Refinement. In *Proceedings of the 15th International Meshing Roundtable*, pages 339–356, Birmingham, Alabama, 2006. Long version available as Carnegie Mellon University Technical Report CMU-CS-06-132.

[4] G. L. Miller, T. Phillips, and D. R. Sheehy. Linear-size meshes. In *CCCG: Canadian Conference in Computational Geometry*, 2008.

[5] M. H. Overmars. Point location in fat subdivisions. *Information Processing Letters*, 44:261–265, 1992.