

An Output-Sensitive Algorithm for Computing Weighted α -Complexes*

Donald R. Sheehy[†]

Abstract

An α -complex is a subcomplex of the Delaunay triangulation of a point set $P \subset \mathbb{R}^d$ that is topologically equivalent to the union of balls of radius α centered at the points of P . In this paper, we give an output-sensitive algorithm to compute α -complexes of n -point sets in constant dimensions, whose running time is $O(f \log n \log \frac{\alpha}{s})$, where s is the smallest pairwise distance and f is the number of simplices in the $c\alpha$ -complex for a constant c . The algorithm is based on a refinement of a recent algorithm for computing the full Delaunay triangulation of P . We also extend the algorithm to work with weighted points provided the weights are appropriately bounded. The new analysis, which may be of independent interest, bounds the number of intersections of k -faces of a Voronoi diagram with $(d - k)$ -faces of the Voronoi diagram of a carefully constructed super-set.

1 Introduction

The starting point for many algorithmic problems in computational geometry is the discrete representation of continuous objects. The α -complex gives a topologically faithful representation of a union of balls as a subcomplex of the Delaunay triangulation of the centers [6]. Weighted α -complexes model the case where the radii of the balls are permitted to vary.

As with the Delaunay triangulation, the α -complex has a dramatic difference in the number of simplices in best- and worst-case examples. However, it can be that even though the Delaunay triangulation may be large, say $\Theta(n^{\lceil d/2 \rceil})$ simplices, the α -complex may still be quite small. Thus, our goal is to compute the α -complex without computing the full Delaunay triangulation. Our approach will be to modify a recent output-sensitive algorithm for computing Delaunay triangulations [11] as well as providing a new perspective to the analysis that gives nearly tight bounds on the number of bistellar flips in a restricted case of kinetic Delaunay triangulations, a result of independent interest.

Contributions Our main contributions are the following.

1. We introduce a generalization of the aspect ratio of a Voronoi cell that applies also to the cells of dimension less than d and relate the aspect ratio to the number of flips needed in removing a subset of vertices from a Delaunay triangulation in the output-sensitive Delaunay triangulation algorithm of Miller and Sheehy [11]. This gives a tighter analysis and also leads to the following algorithmic results.
2. We give a generalization of the Miller-Sheehy algorithm to handle weighted points, assuming the weight of any point is less than half the distance to its nearest neighbor.
3. We give a variation of the algorithm that returns the (weighted) α -complex of the point set without computing the full Delaunay triangulation.




Figure 1: The Voronoi diagram and α -complex of a point set in the plane.

Related Work The classic reference for α -complexes is the survey by Edelsbrunner [6]. Several interesting variations of α -complexes have been proposed including conformal α -complexes [8] which use an alternative to weighting to approximate variations in radii and alpha-beta witness complexes [1] which relax the condition that the output be embedded in \mathbb{R}^d .

We put a restriction on the class of weight functions that are permitted. A generalization to arbitrary weights would mean a new output-sensitive algorithm for convex hulls. The restriction is precisely that used by Cheng et al. [4] on sliver exudation, a method that adds weights to Delaunay triangulations to eliminate certain badly shaped simplices. That work is closely related to

*Partially supported by the National Science Foundation under grant number CCF-1464379

[†]University of Connecticut don.r.sheehy@gmail.com

further use of weights in surface reconstruction to account for curvature [3]. Other curve and surface reconstruction algorithms explicitly use α -complexes [2, 13].

2 Background

Voronoi and Power Diagrams The *Euclidean norm* of a point $x \in \mathbb{R}^d$ is denoted $\|x\|$ and the *Euclidean distance* between points $x, y \in \mathbb{R}^2$ is $\|x - y\|$. The distance from a point to a set is defined as $\mathbf{d}(x, P) := \min_{p \in P} \|x - p\|$. A *weighted point set* is a finite set $P \subset \mathbb{R}^d$ and a weight function $w : P \rightarrow \mathbb{R}_{\geq 0}$. An unweighted set may be viewed as a weighted set with all weights 0. The *power* of a weighted point p is the function $\pi_p(x) := \|p - x\|^2 - w(p)^2$. The *power distance* between a point x and a weighted set P is defined as $\pi_P(x) := \min_{p \in P} \pi_p(x)$. The *power diagram* Vor_P of a weighted point set P is the set of nonempty polyhedra called *Voronoi cells* indexed by subsets $S \subseteq P$ as follows.

$$\text{Vor}_P(S) := \{x \in \mathbb{R}^d \mid \forall s \in S : \pi_P(x) = \pi_s(x)\}$$

One can easily check that for unweighted point sets, this yields the standard Euclidean Voronoi diagram.

The dual diagram, Del_P , is the set of convex closures of the sets S such that $\text{Vor}_P(S)$ is nonempty. Duals of power diagrams go by several names including *weighted Delaunay triangulations*, *regular triangulations*, and *coherent triangulations*. These names assume that the points are in sufficiently general position that the duals are triangulations. In general the weighted Delaunay triangulation is the orthogonal projection in \mathbb{R}^d of the lower convex hull of the points $P^+ := \{(p, \|p\|^2 - w(p)^2) \in \mathbb{R}^{d+1} \mid p \in P\}$.

We say that a point set P is *mildly weighted* if for all $p \in P$, we have $w(p) < \frac{1}{2} \min_{q \in P \setminus \{p\}} \|p - q\|$. In particular, this implies that if the points are viewed as balls with radii equal to the weights, then the balls are disjoint. Unweighted points are mildly weighted.

For mildly weighted points P , we define the *weighted feature size* as

$$\mathbf{f}_{P,w}(x) := \sqrt{\min_{(u,v) \in \binom{P}{2}} \max_{p \in \{u,v\}} \pi_p(x)}.$$

This is the square root of the second smallest power distance from x to a point of P . If the points were not mildly weighted, the feature size could be imaginary at some points. If the points are unweighted, then we will simply write \mathbf{f}_P for the feature size, and, in this case, the square root of the power distance is just the Euclidean distance. The function \mathbf{f}_P is sometimes called the Ruppert local feature size and is ubiquitous in the analysis of Delaunay and Voronoi refinement mesh generation [12].

Weighted α -Complexes An *orthoball* of a set of weighted points S is a ball B with center c and radius r such that $\pi_p(c) = r^2$ for all $p \in S$. The minimum radius for an orthoball of S is called the *orthoradius*. For unweighted points, the orthoball is called the *circumball* and the orthoradius is called the *circumradius*.

The α -offsets of a weighted point set are defined as $P^\alpha := \{x \in \mathbb{R}^d \mid \pi_P(x) \leq \alpha^2\}$. A Voronoi cell of a subset $\sigma \subseteq P$ restricted to the offsets is defined as $\text{Vor}_P^\alpha(\sigma) := \text{Vor}_P(\sigma) \cap P^\alpha$ and the corresponding Voronoi diagram is $\text{Vor}_P^\alpha := \{\text{Vor}_P^\alpha(\sigma) \mid \sigma \subseteq P\}$. The α -complex is the subcomplex of the Delaunay triangulation restricted to the α -offsets as follows.

$$\text{Del}_P^\alpha := \{\sigma \in \text{Del}_P \mid \text{Vor}_P^\alpha(\sigma) \neq \emptyset\}.$$

Equivalently the α -complex may be defined as the *nerve* of set of clipped Voronoi cells $\{\text{ball}(p, \alpha) \cap \text{Vor}_P(p) \mid p \in P\}$, i.e. an abstract simplicial complex with a simplex for every subset of P whose corresponding clipped Voronoi cells have a common intersection. The Nerve Theorem, a standard result in algebraic topology guarantees that Del_P^α is homotopy equivalent to P^α . This topological guarantee was extended by Edelsbrunner and Shah [7] and forms the foundation of many of the topological guarantees in surface reconstruction [5].

Aspect Ratios of Voronoi Cells We will assume here and throughout that all Voronoi cells are bounded and convex. There are two different ways this will be enforced. First, we will consider a global bounding ball Ω that contains all the points and restrict our attention to the intersection of the full Voronoi cells with Ω . Second, when considering α -complexes, we will intersect the Voronoi cell of a point p with the ball of radius α centered at P . Having bounded cells allows the following definition (illustrated in Figure 2).

Definition 1 If P is a set of mildly weighted points and $F \in \text{Vor}_P$, the aspect ratio of F is defined as

$$\text{aspect}_P(F) := \frac{\max_{x \in F} \mathbf{f}_{P,w}(x)}{\min_{y \in F} \mathbf{f}_{P,w}(y)}.$$

More generally, we let aspect_P denote the geometric mean of the aspect ratios of all Voronoi cells (of all dimensions) in Vor_P , i.e.

$$\text{aspect}_P := \left(\prod_{F \in \text{Vor}_P} \text{aspect}(F) \right)^{1/f},$$

where $f = |\text{Vor}_P|$. A more useful way to write this definition is the following.

$$f \log(\text{aspect}_P) = \sum_{F \in \text{Vor}_P} \log(\text{aspect}_P(F)). \quad (1)$$




Figure 2: The aspect ratio of two cells of a Voronoi diagram. Left: a 2-dimensional cell. Right: a 1-dimensional cell. In both cases, the distances are measured from the point of P (in the interior of the 2-dimensional cell) to the nearest and farthest points in the cell F .

Well-Spaced Points A set of points M inside a bounding domain Ω is called τ -well-spaced if for all $q \in M$, $\text{aspect}_M(\text{Vor}_M(q)) \leq \tau$. As Voronoi cells of well-spaced points are nearly balls, simple packing arguments imply that there is a constant c_1 such that $\text{Vor}_M(q)$ has at most c_1 faces for all $q \in M$. Given a set of n points P and a bounding ball Ω , there exists a τ -well-spaced superset M of P as long as $\tau > 2$. Asymptotically minimal well-spaced supersets are *graded* in the sense that there is a constant K such that for all $v \in M$, we have $\mathbf{f}_P(v) \leq K\mathbf{f}_M(v)$. The grading condition implies that there is a constant c_2 such that for all $r > 0$, at most c_2 points of M have Voronoi cells intersecting $\text{annulus}(q, r, 2r)$ for any $q \in P$ [14], where $\text{annulus}(q, r, 2r)$ denotes $\text{ball}(q, 2r) \setminus \text{ball}(q, r)$. The constants K , c_1 and c_2 only depend on d and τ . Moreover, such a superset can be found in $O(n \log n + |M|)$ time [10]. Finally, we will use another important fact about graded, well-spaced point sets, namely that there is a constant γ such that $r \leq \gamma\mathbf{f}_M(x)$ for all x in any empty ball of radius r (see [9, Lemma 6.1]).

We will say a point set P is *annulus-free* if there is no point p and radius R such that $\text{ball}(p, r)$ contains more than one point of P and $\text{annulus}(p, r, 10r)$ contains no points of P . The constant 10 here is arbitrary. The size of a τ -well-spaced superset M from an annulus-free set P is known to be $O(n)$, so the running time to compute M is $O(n \log n)$ [14]. For α -complexes, point sets that are not annulus-free are rather uninteresting: if $r > \alpha$ then the ball points in the ball form a separate component; if $r \ll \alpha$ then the points are much closer than the scale and so replacing them with a single point results in (Hausdorff-)close offsets.

A Kinetic View of Refinement Given a set $S \subset \mathbb{R}^d$ of $d + 2$ points in general position, there are precisely two different triangulations of S . A bistellar flip is a local change in a triangulation that swaps between the two

triangulations of such a subset of $d + 2$ points. Given a point set P in a bounding ball B and a constant τ , there exists a τ -well-spaced superset $M \supseteq P$. Starting from the Delaunay triangulation of M , one may obtain the Delaunay triangulation of P , by incremental bistellar flips that ultimately remove the points of $M \setminus P$ except those on the convex hull. This is done by changing the weights linearly and tracking the incremental changes that occur in the weighted Delaunay triangulation. As the change in weights may be viewed as a change in heights for a kinetic convex hull problem, the combinatorial changes can all be computed by replacing the coordinates in the usual Delaunay in-sphere predicate with the linear functions describing the motion. These changes are stored in a heap and are processed one at a time.

3 The Algorithm

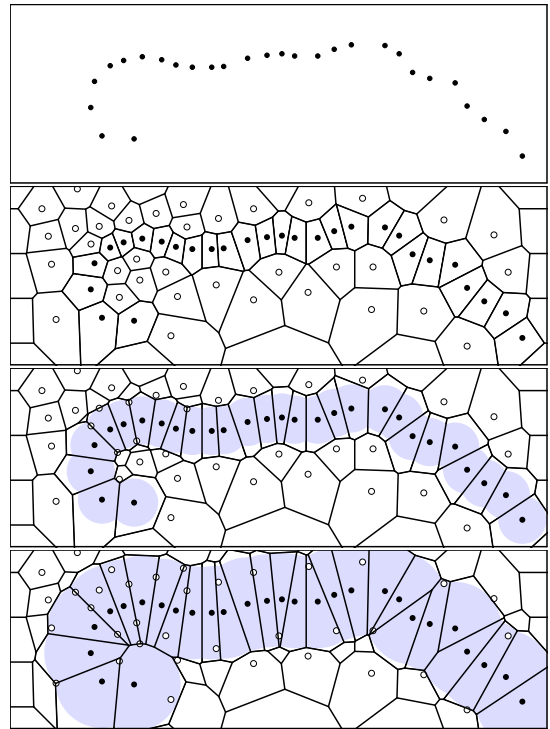


Figure 3: The algorithm is illustrated from top to bottom in terms of the Voronoi diagram. Starting from the input points (black), Steiner points are added (white). Weight is then added to the input points causing local changes to the Voronoi diagram until the weighted Voronoi cells of the input points contain the α -offsets.

In this section, we describe the algorithm for computing the α -complex of a set of mildly weighted points and prove its correctness. The algorithm starts by building a linear-size Delaunay triangulation of a well-spaced superset M of the input points P . The extra points

are called *Steiner points*. Then it adjusts the weights to match the input weights, leaving the weights of the Steiner points as zero. The projective view of weighted Delaunay triangulation assigns a height to a point p equal to $\|p\|^2 - w(p)^2$. Now, this height is treated as a $(d+1)$ st coordinate and the weighted Delaunay triangulation is the orthogonal projection of the lower convex hull back into \mathbb{R}^d . In the algorithm, we treat the weights as a function of time, so, the weight $w(p)$ is specified in the input but the algorithm uses

$$w(p, t) = \begin{cases} \sqrt{w(p)^2 + t} & \text{if } p \in P \\ w(p) & \text{otherwise.} \end{cases}$$

The weighted point set at time t is denoted M_t . The height of a point p at time t is

$$h(p, t) := \|p\|^2 - w(p, t)^2.$$

As t increases, the input points get pulled downward. Generically, each combinatorial change in the weighted Delaunay triangulation is a single bistellar flip. As the input points move downward, the Steiner points are flipped out of the triangulation. Note that the definition of $w(p, t)$ guarantees that the height $h(p, t)$ is either constant or a linear function of t . Weighted points $S = \{p_1, \dots, p_{d+2}\}$ in \mathbb{R}^d have a common orthoball exactly when they lie on a common hyperplane after lifting. Thus, we can check this condition by computing

$$\text{flip}_S(t) = \det \begin{bmatrix} p_{1,1} & \cdots & p_{d+2,1} \\ \vdots & \cdots & \vdots \\ p_{1,d} & \cdots & p_{d+2,d} \\ h(p_1, t) & \cdots & h(p_{d+2}, t) \\ 1 & \cdots & 1 \end{bmatrix}.$$

Note that $\text{flip}_S(t)$ is a linear function of t and so we can compute the *flip time* t_S satisfying $\text{flip}_S(t_S) = 0$. More generally, when the coordinates of the points (and not just the heights) are polynomials in t , $\text{flip}_S(t)$ is some polynomial, and computing the roots of $\text{flip}_S(t)$ gives the changes in the Delaunay triangulation as the points move. This more general setting is the quintessential example in the field of kinetic data structures, a generalization of the line-sweep paradigm.

In our case, we are only modifying the height and so the algebraic computations are much simpler. The main data structure is a heap called the *flip heap* that stores the possible flips ordered by time. We identify each flip with a facet in Del_{M_t} . The steps of the construction given in Algorithm 1.

The following lemma guarantees that stopping the kinetic part of the algorithm at time $t = \alpha^2$, will still allow us to construct the α -complex.

Lemma 2 *If a simplex $\sigma \in \text{Del}_P$ is contained in a d -simplex $\sigma' \in \text{Del}_P$, of orthoradius at most \sqrt{t} then the*

Algorithm 1 Compute the α complex for a mildly weighted point set.

-
- 1: **procedure** ALPHA COMPLEX(P, α)
 - 2: Compute a graded, τ -well-spaced superset M of P in a bounding ball B containing P .
 - 3: For each facet F of Del_M , compute the flip time t_F and insert the key-value pair (t_F, F) into the flip heap. Skip the insertion if $t_F > \alpha^2$.
 - 4: **while** The flip heap is nonempty **do**
 - 5: Pop a facet F off the heap
 - 6: Attempt to flip F , and push any new facets to the heap if their flip time is at most α^2 .
 - 7: Output all simplices containing only points of P that have an orthoradius at most α .
-

flip time when σ first appears in the ALPHA COMPLEX algorithm is at most t .

Proof. First, observe that for sufficiently large α , every simplex of Del_P will appear at some time. Let t_0 be the time when σ first appears and let c be the orthocenter of the corresponding flip. Let p be any vertex of σ , so $t_0 = \pi_p(c)$. This means that c is the orthocenter of σ' , the smallest d -dimensional simplex (by orthoradius) in Del_P containing σ . Suppose for contradiction that $t_0 > t$. At time t , some vertex q of M has a Voronoi cell containing c such that $q \neq p$. So, $\pi_q(c, t) < \pi_p(c, t)$ and so it follows from the definition of the power distance that $0 \leq \pi_q(c) \leq \pi_p(c) - t$. Because $t_0 = \pi_p(c)$, the preceding inequality implies that $t_0 \leq t$, a contradiction. Therefore, we conclude that the flip time t_0 when σ first appears is at most t as claimed. \square

Lemma 2 now implies the following theorem as it guarantees that the algorithm finds all simplices of the α -complex despite stopping at time α^2 .

Theorem 3 *Given a set P of mildly weighted points and a parameter α , the ALPHA COMPLEX algorithm above returns the weighted α -complex of P .*

4 Analysis

The starting point for the analysis of the running time of our output-sensitive algorithm for α -complexes is the following lemma of Miller & Sheehy (proven as a first step in Lemma 6 of [11]) describing the flips in the algorithm; it says that the flips are in one-to-one correspondence with intersections between Vor_M and Vor_P . We extend it to the weighted case.

Lemma 4 *Let $P \subset \mathbb{R}^d$ be a finite point set and let $M \supseteq P$ be any superset of P . A set S of $d+2$ points of M will be involved in a flip in the kinetic refinement reversal if and only if $\text{Vor}_M(S \setminus P) \cap \text{Vor}_P(S \cap P) \neq \emptyset$.*

The preceding lemma gives a static way to count the kinetic changes in the algorithm; it suffices to count intersections between the starting and ending Voronoi diagrams. In previous work, a coarse bound on the number of intersections was given by exploiting the fact that the point set M in the algorithm is well-spaced. That analysis give a bound of $O(\log \Delta)$ flips per simplex. We will now give a more refined analysis that bounds these flips instead in terms of a more local parameter. Incidentally, this will also improve the running-time guarantee for certain known hard instances for Delaunay triangulation, such some that are known to produce $\Theta(n^2)$ simplices in \mathbb{R}^3 .

Lemma 4 implies a natural way to partition the set of flips by assigning the set of points S in the flip to the simplex $S \cap P$ of Del_P . This is clearly possible, because according to the lemma $\text{Vor}_P(S \cap P)$ must be nonempty for the flip to occur and so $S \cap P \in \text{Del}_P$. To count the total flips, it will suffice to bound the number of flips assigned to each simplex of Del_P .

Lemma 5 *Let $P \subset \mathbb{R}^d$ be a set of mildly weighted points. Let M be a τ -well-spaced superset of P . Let $F \in \text{Vor}_P$ be any face. There is a constant c that depends only on d and τ such that*

$$|\{G \in \text{Vor}_M \mid F \cap G \neq \emptyset\}| \leq c \log(\text{aspect}_P(F)).$$

Proof. Let $S \subset P$ be such that $F = \text{Vor}_P(S)$ and let $q \in S$ be any point. Recall from Section 2, for graded supersets M , there is a constant c_2 such that for all $r > 0$, at most c_2 points of M have Voronoi cells intersecting annulus($q, r, 2r$). Moreover, there is a constant c_1 such that each such Voronoi cell has at most c_1 faces. Let $x = \text{argmax}_{x \in F} \mathbf{f}_{P,w}(x)$ and $y = \text{argmin}_{y \in F} \mathbf{f}_{P,w}(y)$. Let $A_i = \text{annulus}(q, 2^i r, 2^{i+1} r)$ for all integers i , where $r = \|q - y\|$. By Lemma 9, $\|x - q\| \leq r \text{aspect}_P(F)$. It follows that $F \subset \bigcup_{i=0}^{\lceil \log(\text{aspect}_P(F)) \rceil - 1} A_i$. As there are at most $c_1 c_2$ intersections in any A_i and thus at most $c_1 c_2 \lceil \log(\text{aspect}_P(F)) \rceil$ faces of Vor_M intersect F . \square

We first give an upper bound on the total number of flips in terms of the aspect ratio of Vor_P . This bound applies independent of the value of α and thus gives a potentially tighter bound on the number of flips in computing the full Delaunay triangulation of P .

Theorem 6 *For a mildly weighted point set P and any constant $\alpha \geq 0$, the total number of flips in the ALPHACOMPLEX algorithm is $O(f \log(\text{aspect}_P))$, where $f = |\text{Vor}_P|$.*

Proof. By Lemma 4 and (1), it will suffice to prove that each k -face F of Vor_P intersects at most $O(\log \text{aspect}(F))$ $(d - k)$ -faces of Vor_M , which is precisely the conclusion of Lemma 5. \square

The following lemma guarantees that the flips that occur in the algorithm, all happen “close” to the input points. That is, none of the intersections causing a flip are farther than a constant times α from the points of P . This is the key to proving an output-sensitive running time for α -complexes as it says that the output simplices are discovered (approximately) in order of their orthoradius.

Lemma 7 *Let P be a set of mildly weighted points, let $p \in P$, let $\alpha \in \mathbb{R}$, and let x be the center of a flip in the ALPHACOMPLEX algorithm that occurs at time t . There is a constant c_3 that depends only on τ and d such that if $\mathbf{f}_P(p) \leq 2\sqrt{2}\alpha$ and $t \leq \alpha^2$, then $\|x - p\| \leq c_3\alpha$.*

Proof. Let $b := w(p, t) = \sqrt{w(p)^2 + t} \leq \sqrt{2}\alpha$, where the last inequality follows from the mildness assumption and the hypothesis that $t \leq \alpha^2$. Let r be the radius of the orthoball of the flip centered at x , so $r = \sqrt{\|x - p\|^2 - b^2}$. Let y be the point on the line segment \overline{xp} such that $\|x - y\| = r$. Now, for τ -well-spaced points M , it is known that there is a constant γ such that any ball of radius r that contains no points of M has the property that $r \leq \gamma \mathbf{f}_P(z)$ for all z in the ball. So, in our case, this implies that $r \leq \gamma \mathbf{f}_P(y)$ because $y \in \text{ball}(x, r)$ and no points M lie in this ball.

We consider two cases. First, if $r < b$, then $\|x - p\|^2 = r^2 + b^2$ implies that $\|x - p\| \leq 2\alpha$. So in that case, it suffices to choose $c_3 \geq 2$. Second, if $r \geq b$, then similarly,

$$\begin{aligned} \|x - p\| &\leq \sqrt{2}r \\ &\leq \sqrt{2}\gamma \mathbf{f}_M(y) \\ &\leq \sqrt{2}\gamma \mathbf{f}_P(y) \\ &\leq \sqrt{2}\gamma (\mathbf{f}_P(p) + \|p - y\|) \\ &\leq \sqrt{2}\gamma (2\sqrt{2}\alpha + b) \\ &\leq 6\gamma\alpha. \end{aligned}$$

So, $c_3 = 6\gamma$ is the desired constant. \square

The preceding lemma provides the main new tool for analyzing the running time of ALPHACOMPLEX and also indicates why the output-sensitive term is not precisely the output size but rather the size of a constant factor larger α -complex.

Theorem 8 *For a mildly weighted, annulus-free point set $P \subset \mathbb{R}^d$, the total running time of ALPHACOMPLEX (P, α) is $O(f \log(n) \log(\alpha/s))$ where $s := \min_{p \in P} \mathbf{f}_P(p)$, $f = |\text{Del}_P^{c_3\alpha}|$, and c_3 is the constant from Lemma 7.*

Proof. The preprocessing phase to compute M and Del_M takes $O(n \log n)$ time [12]. Adjusting Del_M to form Del_{M_t} requires only a constant number of flips per vertex. This fact is used in the work of Cheng et al. on

sliver exudation [4], but also follows from the arguments of Lemma 7.

The main loop processes flips and adds them to a heap. Lemma 7 implies that all of the flips are contained in $P^{c_3\alpha}$. Let f be the number of faces of $\text{Vor}_P^{c_3\alpha}$. The aspect ratio of each such face is at most $\log(c_3\alpha/s)$. So, Theorem 6 implies that the total number of flips is $O(f \log(\alpha/s))$. Not every flip on the heap is performed, but at most a constant number of potential flips are added to the flip heap for every actual flip, so the total number of heap operations is also $O(f \log(\alpha/s))$. Thus, the total running time is $O(f \log(n) \log(\alpha/s))$. \square

5 Conclusion

In this paper, we generalized the output-sensitive algorithm of Miller and Sheehy for Delaunay triangulations to also give guarantees for mildly weighted points and for α -complexes. Along the way, we generalized the notion of the aspect ratio of a Voronoi cell to give a meaningful definition for weighted Voronoi cells of dimension less than d . This new definition gives a sharper analysis for the Miller and Sheehy algorithm and also simplifies the analysis of the modified algorithm.

One future direction is to look at more recent generalizations of α -complexes introduced for topological inference from point cloud data such as the alpha-beta witness complexes of Attali et al. [1]. It may also be useful to apply this approach to Voronoi-based manifold reconstruction as many use the Delaunay triangulation restricted to the manifold which is a subcomplex of the Delaunay triangulation restricted to a union of balls (see for example [2, 13]). We are also interested in relaxing the mild weighting assumption and replacing it with a Lipschitz condition on the weights. In that case, weights could be larger as long as they don't vary too quickly.

References

- [1] D. Attali, H. Edelsbrunner, J. Harer, and Y. Mileyko. Alpha-beta witness complexes. In *WADS*, 2007.
- [2] F. Bernardini and C. L. Bajaj. Sampling and reconstructing manifolds using alpha-shapes. In *In Proc. 9th Canad. Conf. Comput. Geom.*, pages 193–198, 1997.
- [3] J.-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential Delaunay complexes. *Discrete & Computational Geometry*, 51(1):221–267, 2014.
- [4] S.-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S.-H. Teng. Sliver exudation. *JACM: Journal of the ACM*, 47, 2000.
- [5] T. K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, 2007.
- [6] H. Edelsbrunner. The union of balls and its dual shape. *Discrete & Computational Geometry*, 13:415–440, 1995.

- [7] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. In *Proceedings of the 10th ACM Symposium on Computational Geometry*, pages 285–292, 1994.
- [8] J. Giesen, F. Cazals, M. Pauly, and A. Zomorodian. The conformal alpha shape filtration. *The Visual Computer*, 22(8):531–540, 2006.
- [9] 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.
- [10] G. L. Miller, T. Phillips, and D. R. Sheehy. Beating the spread: Time-optimal point meshing. In *Proceedings of the 26th ACM Symposium on Computational Geometry*, pages 321–330, 2011.
- [11] G. L. Miller and D. R. Sheehy. A new approach to output-sensitive construction of Voronoi diagrams and Delaunay triangulations. *Discrete & Computational Geometry*, 52(3):476–491, 2014.
- [12] G. L. Miller, D. R. Sheehy, and A. Velingker. A fast algorithm for well-spaced points and approximate Delaunay graphs. In *Proceedings of the 29th annual Symposium on Computational Geometry*, pages 289–298, 2013.
- [13] S. H. Park, S. S. Lee, and J. H. Kim. A surface reconstruction algorithm using weighted alpha shapes. In L. Wang and Y. Jin, editors, *Fuzzy Systems and Knowledge Discovery*, volume 3613 of *Lecture Notes in Computer Science*, pages 1141–1150. Springer, 2005.
- [14] D. R. Sheehy. New Bounds on the Size of Optimal Meshes. *Comput. Graph. Forum*, 31(5):1627–1635, 2012.

A The Weighted Aspect Ratio

In the following lemma, we show that if we replace the power distance with the Euclidean distance, the aspect ratio cannot go up.

Lemma 9 *Let P be a mildly weighted point set. Let $S \subset P$ and let $F = \text{Vor}_P(S)$, where $|S| \geq 2$. Let $x = \text{argmax}_{x \in F} \mathbf{f}_{P,w}(x)$ and $y = \text{argmin}_{y \in F} \mathbf{f}_{P,w}(y)$. For all $q \in S$, $\|q - x\| \leq \|q - y\| \text{aspect}_P(F)$.*

Proof. First, observe that by the choice of x and y , we know that $\mathbf{f}_{P,w}(x)^2 = \pi_q(x)$ and $\mathbf{f}_{P,w}(y)^2 = \pi_q(y)$. The desired inequality now follows from the definitions of π_q and $\text{aspect}_P(F)$ as follows.

$$\begin{aligned} \|q - x\|^2 &= \pi_q(x) + w(q)^2 \\ &= \pi_q(y) \text{aspect}_P(F)^2 + w(q)^2 \\ &\leq \text{aspect}_P(F)^2 (\pi_q(y) + w(q)^2) \\ &= \|q - y\|^2 \text{aspect}_P(F)^2. \end{aligned}$$

\square