CGL

Computational Geometric Learning

A New Approach to Output-Sensitive Voronoi Diagrams and Delaunay Triangulations

---

Gary L. Miller   Donald R. Sheehy

CGL Technical Report No.: CGL-TR-64

A New Approach to Output-Sensitive Voronoi Diagrams

Gary L. Miller and Donald R. Sheehy

Abstract

We describe a new algorithm for computing the Voronoi diagram of a set of n points in constant-dimensional Euclidean space. The running time of our algorithm is $O(f \log n \log \Delta)$ where f is the output complexity of the Voronoi diagram and Δ is the spread of the input, the ratio of largest to smallest pairwise distances. Despite the simplicity of the algorithm and its analysis, it improves on the state of the art for all inputs with polynomial spread and near-linear output size. The key idea is to first build the Voronoi diagram of a superset of the input points using ideas from Voronoi refinement mesh generation. Then, the extra points are removed in a straightforward way that allows the total work to be bounded in terms of the output complexity, yielding the output sensitive bound. The removal only involves local flips and is inspired by kinetic data structures.

1 Introduction

Voronoi diagrams and their duals, Delaunay triangulations, are ubiquitous in computational geometry, both as a source of interesting theory and a tool for applications [3]. Algorithms for planar Voronoi diagrams abound where optimal algorithms are known. In higher dimensions, the situation is complicated by the large gap between the best-case and worst-case output complexity. A Voronoi diagram of n points in $\mathbb{R}^d$ can have between $\Theta(n)$ and $\Theta(n^{\lceil d/2 \rceil})$ faces [18, 23] (we suppress constant factors that only depend on d). This motivates the search for *output-sensitive* algorithms and analysis, where the time and space guarantees depend on both the input size n and the number of output faces f.

Computing Voronoi diagrams in $\mathbb{R}^d$ reduces to computing convex hulls in $\mathbb{R}^{d+1}$. This relationship is perhaps most clear in the dual, where the Delaunay triangulation is the projection of the lower hull of the input points

lifted into $\mathbb{R}^{d+1}$ by the standard *parabolic lifting*:

$$(x_1, \ldots, x_d) \mapsto \left(x_1, \ldots, x_d, \sum_{i=1}^{d} x_i^2 \right).$$

All of the previous literature on output-sensitive Voronoi diagrams in higher dimensions is directed at solving the more general problem of computing convex hulls.[1]

Seidel gave an algorithm based on polytope shelling that runs in $O(n^2 + f \log n)$ time [22]. The quadratic term is from a preprocess that solves a d-dimensional linear program for each input point. Matousek and Schwartzkopf showed how to exploit the common structure in these linear programs to improve the running time to $O(n^{2-2/(\lceil d/2\rceil+1)} \log^{O(1)} n + f \log n)$ [19].

Another paradigm of algorithms uses the dual notions of gift-wrapping [26] and pivoting [4]. Both approaches can enumerate the facets of a simple polytope in $O(nf)$ time, thus paying approximately linear time per face. Since Voronoi diagrams have at least n faces, these methods do not improve on the LP-based methods except that they avoid the exponential dependence on the dimension inherent in such methods.

Chan gave an algorithm based on gift-wrapping that runs in $O(n \log f + (nf)^{1-1/(\lceil d/2\rceil+1)} \log^{O(1)} n)$ time [6]. A later work by Chan et al. gave an algorithm that runs in $O((n + (nf)^{1-1/\lceil (d+1)/2\rceil} + fn^{1-2/\lceil (d+1)/2\rceil}) \log^{O(1)} n)$ time [7]. Even when $f = \Theta(n)$, the running time is poly(n) per face.

Better bounds are known for 3- and 4-dimensional Voronoi diagrams where truly polylogarithmic time per face algorithms are known. Chan et al. gave an algorithm that achieves $O(f \log^2 n)$ for $\mathbb{R}^3$ [7]. Amato and Ramos gave an $O(f \log^3 n)$-time algorithm in $\mathbb{R}^4$ [2].

Voronoi diagrams and Delaunay triangulations are used in mesh generation (see the recent book by Chan et al. [10]). Extra vertices called Steiner points are added in a way that keeps the complexity down. Perhaps surprisingly, the number of vertices increases, but the total number of faces can decrease. Such a mesh can be constructed in $O(n \log \Delta)$ time using only $O(n \log \Delta)$ vertices [14, 20].

In this paper, we propose a new algorithm for constructing Voronoi diagrams that uses Voronoi refinement mesh generation as a preprocess. Then, it removes all of the Steiner points using a method derived from the field of kinetic data structures. We prove that at most $O(f \log \Delta)$ local changes occur during the removal process. Each local change requires only constant

---

[1] To avoid confusion, when reporting running times of known algorithms, we always give the results as they apply to Voronoi diagrams rather than convex hulls.

time to process. These local changes are ordered via a heap data structure which adds an extra factor of $O(\log(f \log \Delta)) = O(\log n + \log \log \Delta)$ to the running time. We assume an asymptotic floating point model of computation where points are represented by their coordinates with $O(\log n)$-bit floating point numbers [13]. In such a model, $\log \log \Delta = O(\log n)$. Thus, the total running time is $O(f \log n \log \Delta)$.

Unlike previous work on output-sensitive Voronoi diagram construction, our algorithm does not use a reduction to the convex hull problem. Instead, it uses specific properties of the Voronoi diagram to get an improvement.

Contribution We present the MESHVORONOI algorithm, a new, output-sensitive algorithm for computing Voronoi diagrams in $\mathbb{R}^d$. MESHVORONOI runs in $O(f \log n \log \Delta)$ time. When $\Delta = \text{poly}(n)$ such as the case when input points have integer coordinates with $O(\log n)$ bits of precision, the running time is $O(f \log^2 n)$. Even the $O(n^{\lceil d/2 \rceil})$ worst-case inputs for Voronoi diagrams can be represented with polynomial (even linear) spread. We get an improvement over existing algorithms for inputs with polynomial spread in dimension $d > 3$ when $f = O(n^{2-2/\lceil d/2 \rceil})$. Moreover, the analysis in terms of the spread is often quite loose. The running time really depends on the aspect ratios of the individual Voronoi cells of the output, for which the spread stands in as an easy to describe upper bound. The other advantage of the MESHVORONOI algorithm is that it is simple both to describe and to analyze.

Related Work We will make use of the flip-based construction of weighted Delaunay triangulations similar to that presented by Edelsbrunner and Shah [11]. In that paper, the concern was to add a single point to a regular triangulation but when run backwards, it describes the removal of a single point by local flips. We are interested in removing sets of points simultaneously, the Steiner points of the mesh. In this respect, the problem more resembles the Delaunay triangulation splitting problem for which Chazelle et al. gave a linear time algorithm for the plane [8]. The only higher dimensional analogue of this result is the extension by Chazelle and Mulzer to the case of splitting convex polytopes in $\mathbb{R}^3$ [9].

Our algorithm may be viewed as a special case of a kinetic convex hull problem, and indeed, the main tools come directly from the literature on kinetic data structures [12]. In general, the kinetic convex hull problem is much harder than the instances arising in our algorithm and it is only because of the specific geometric structure of these instances that we are

able to prove useful bounds.

Joswig and Ziegler presented a different approach to output sensitive convex hulls using homology calculations [17]. This algorithm is shown to be output-sensitive for simplicial polytopes like those produced for Delaunay triangulations of points in general position. However, they do not improve bounds on the asymptotic running times in terms of n and f as their construction passes through several reductions.

Many low-dimensional algorithms for computing Voronoi diagrams depend on incremental construction, where the points are added one at a time. When the input can be degenerate, Bremner showed that incremental constructions cannot be output-sensitive [5]. Our algorithm does resemble an incremental algorithm. The main difference is that it adds *more* than just the input points in the first phase of the algorithm. The harder work is removing these extra points.

2 Background

Points and Distances in Euclidean Space We will deal exclusively with the case of points in d-dimensional Euclidean space. The Euclidean norm of x is denoted $\|x\|$ and thus the Euclidean distance between two points x and y is $\|x - y\|$. For a point $x \in \mathbb{R}^d$ and a compact set $U \subset \mathbb{R}^d$, define $\mathrm{d}(x, U) := \min_{y \in U} \|x - y\|$. The **spread** Δ of a set of points is the ratio of the largest to smallest pairwise distances.

We make the simplifying assumption that the spread is at most exponential, i.e. $\Delta = 2^{O(n)}$. This is equivalent to assuming an asymptotic floating point notation where coordinates are stored as floating point numbers with $O(\log n)$-bit words (see [13] for a more detailed treatment of this model). In the case of exponential or super-exponential spread, several previous methods are known to beat the $O(f \log n \log \Delta)$ running time of our algorithm. Moreover, it is possible to simulate a $2^{O(n)}$ upper bound on the spread by decomposing the point set into a hierarchy of $O(n)$ point sets, each with such a bound on the spread (see Miller et al. [20] for an application of this approach to mesh generation). The assumption that the spread is at most exponential allows a clearer exposition of the main ideas of our approach without introducing the complexity of hierarchical point sets.

Voronoi Diagrams and Power Diagrams Let $P \subset \mathbb{R}^d$ be a finite set of points. The **Voronoi diagram** of P is a cell complex decomposing $\mathbb{R}^d$ into convex, polyhedral cells such that all points in a cell share a common set

of nearest neighbors among the points of P. Formally, for a subset $S \subset P$, the Voronoi cell $\mathrm{Vor}_P(S)$ is defined as the set of all points x of $\mathbb{R}^d$ such that $\mathrm{d}(x, P) = \|x - y\|$ for all $y \in S$, i.e.

$$\mathrm{Vor}_P(S) := \{x \in \mathbb{R}^d : \mathrm{d}(x, P) = \|x - y\| \text{ for all } y \in S\}$$

When $S = \{v\}$ is a singleton, we will abuse notation and refer to its Voronoi cell as $\mathrm{Vor}_P(v)$ instead of $\mathrm{Vor}_P(\{v\})$. For points in **general position** (no $k + 3$ points lying on a common k-sphere), the affine dimension of $\mathrm{Vor}_P(S)$ is $d - (|S| - 1)$.
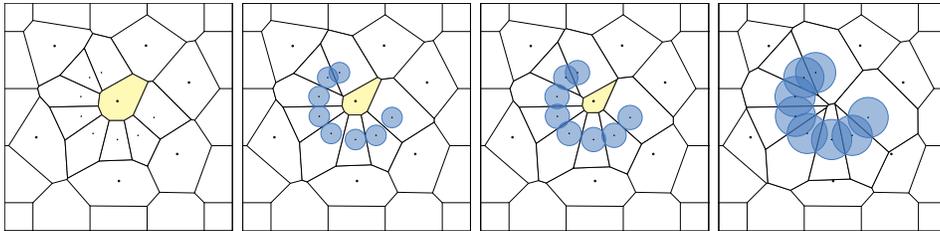

Figure 1: A weighted Voronoi diagram with weights indicated by disks. As the weights of some points increase from left to right, some cells disappear from the diagram altogether.

The Voronoi diagram of P, denoted Vor_P, has a natural dual called the **Delaunay diagram**, denoted Del_P. For point sets in general position, the Delaunay diagram is an embedded simplicial complex called the **Delaunay triangulation**. The Voronoi/Delaunay duality is realized combinatorially by inverting the posets of the corresponding cell complexes, identifying each k-face of the Voronoi diagram with a $(d - k)$-simplex of the Delaunay triangulation.

If the points P have real-valued weights $w : P \to \mathbb{R}$ and the Euclidean distance from $p \in P$ to $x \in \mathbb{R}^d$ is replaced with the power distance $\pi_p(x) = \|x - p\|^2 - w(p)^2$, then the Voronoi diagram becomes a **weighted Voronoi diagram**, also known as a **power diagram**. The dual is still well-defined and is known as the **weighted Delaunay diagram** (or the **weighted Delaunay triangulation** when it is a simplicial complex).

Weighted Delaunay diagrams are the projections of convex polytopes in one dimension higher. In fact, interpreting the power distance to the origin as a height function for the points lifted into $\mathbb{R}^{d+1}$, gives the weighted Delaunay diagram as the projection of the lower convex hull of these lifted points. In particular, setting all weights to 0 gives the (unweighted) Delaunay diagram as a convex hull in $\mathbb{R}^{d+1}$. Thus, there is a strong connection

between the problems of computing convex hulls, Delaunay triangulations, and Voronoi diagrams.

Orthoballs and Encroachment   A weighted point p **encroaches** $B = \text{ball}(c, r)$ if

$$\pi_p(c) < r^2,$$

and it is **orthogonal** to B if

$$\pi_p(c) = r^2.$$

For a collection of $d + 1$ weighted points in $\mathbb{R}^d$ (not all on a hyperplane), the **orthoball** is the unique ball orthogonal to each of the weighted points, and its center and radius are called the **orthocenter** and **orthoradius** respectively. For unweighted points, a point encroaches a ball if its in the interior, and it is orthogonal if it is on the boundary. The orthocenter of unweighted points is called the **circumcenter** and its center and radius are called the **circumcenter** and **circumradius** respectively.

   If a weighted point encroaches the orthoball of a simplex σ then we say that σ is encroached. For points in general position, the weighted Delaunay triangulation is the unique triangulation in which no simplex is encroached by any weighted vertex. In the unweighted case, this corresponds to the property that no vertex lies in the interior of the circumball of any Delaunay simplex.
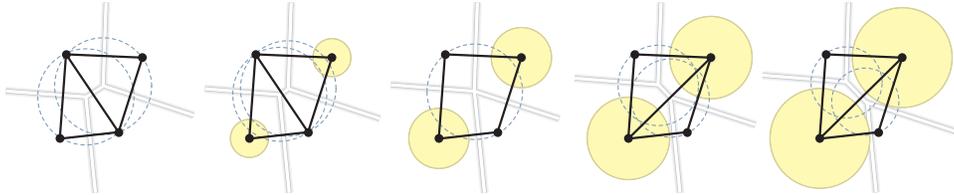

Figure 2:   From left to right, increasing the weight of two points causes a flip in the weighted Delaunay triangulation. The dotted circles indicate the orthoballs of the weighted Delaunay triangles. The center figure indicates the exact weight when the flip occurs as the orthoballs of all four possible triangles coincide.

Flips in Triangulations   Bistellar flips are a useful way to make local changes in triangulations. This operation takes a collection of $d + 2$ vertices S whose convex hull $\text{ch}(S)$ is a subcomplex of the triangulation, and replace

its interior with a new triangulation. In $\mathbb{R}^2$, there are three classes of flips: those that swap the diagonal of a quadrilateral, those that introduce a new vertex in a triangle by splitting it into three, and those that remove a vertex incident to exactly three triangles. These are called respectively $(2,2)$-, $(1,3)$-, and $(3,1)$-flips, where the numbers correspond to the number of d-simplices removed and inserted respectively. In d dimensions, there are similar (i,j)-flips for nonnegative integers i,j such that $i+j = d+2$, where each replaces i d-simplices with j new d-simplices.

Any pair of weighted Delaunay triangulations can be transformed, from one to the other by a sequence of bistellar flips. This process of local changes is at the heart of incremental constructions [11]. Any continuous change in the weights of a set of points causes a change in the corresponding weighted Delaunay triangulation that can be realized by a sequence of bistellar flips. In such a case, the flips happen exactly at those moments when the diagram is no longer a triangulation. For unweighted points, this happens when some set of $d+2$ points lie on a common $(d-1)$-sphere. This can be tested by a linear predicate. That is, $d+2$ points $p_1, \ldots, p_{d+2}$ lie on a d-sphere if and only if

$$\det \begin{bmatrix} p_1 & \cdots & p_{d+2} \\ \|p_1\|^2 & \cdots & \|p_{d+2}\|^2 \\ 1 & \cdots & 1 \end{bmatrix} = 0.$$

These represent the degenerate configurations of points. As before, for weighted points, we replace the norm with the power distance to find that $d+2$ weighted points form a degenerate configuration if and only if

$$\det \begin{bmatrix} p_1 & \cdots & p_{d+2} \\ \|p_1\|^2 - w(p_1)^2 & \cdots & \|p_{d+2}\|^2 - w(p_{d+2})^2 \\ 1 & \cdots & 1 \end{bmatrix} = 0.$$

This determinant test is a consequence of the lifting definition of the weighted Delaunay triangulation.

Voronoi Aspect Ratios and Voronoi Refinement   The **in-radius** of a Voronoi cell $\mathrm{Vor}_P(q)$ is the radius of the largest ball centered at q contained in $\mathrm{Vor}_P(q)$. The **out-radius** of $\mathrm{Vor}_P(q)$ is the radius of the smallest ball centered at q that contains all of the vertices of $\mathrm{Vor}_P(q)$. Such a ball contains all of $\mathrm{Vor}_P(q)$ for bounded Voronoi cells. The **aspect ratio** of the Voronoi cell of q is the ratio the out-radius over the in-radius, denoted $\mathrm{aspect}_P(q)$. A Voronoi diagram has **bounded aspect ratio** if every cell has bounded

aspect ratio. More generally, we say a set of points M is τ-**well-spaced** if for all $v \in M$, $\text{aspect}_M(v) \leq \tau$.

Bounded aspect ratio Voronoi diagrams have many nice properties. The most relevant for our purposes is that no d-dimensional Voronoi cell has more than $2^{O(d)}$ facets (faces of codimension 1) [21].

For any set of n points P, there exists a τ-well-spaced superset M for any constant $\tau > 2$. Moreover, such a superset can be computed in $O(n \log n + |M|)$ time with $|M| = O(n \log \Delta)$ [20]. The process of adding points to improve the Voronoi aspect ratio is called **Voronoi refinement**. It is perhaps more widely known in its dual form, Delaunay refinement. The extra points added are called **Steiner points**.

The analysis of Voronoi refinement depends on a function called the **Ruppert feature size**, defined for all $x \in \mathbb{R}^d$ as the distance to the second nearest input point.

$$\mathbf{f}_P(x) := \max_{p \in P} \mathrm{d}(x, P \setminus \{p\})$$

One defines the feature size with respect to a set M similarly and denote it $\mathbf{f}_M$. Voronoi refinement produces a τ-well-spaced set of points M such that for each vertex $v \in M$,

$$\mathbf{f}_P(v) \leq K\mathbf{f}_M(v),$$

where $K = \frac{2\tau}{\tau - 2}$ (see [14] or [24, Thm. 3.3.2]).

The total number of output points is, up to constant factors, determined by the **feature size measure** of the domain Ω, defined as

$$\mu(\Omega) := \int_\Omega \frac{dx}{\mathbf{f}_P(x)}.$$

For a wide class of input domains, the feature size measure is $O(n)$ [25]. For general inputs, the bound of $O(n \log \Delta)$ mentioned above is well known and can even be derived as a corollary to Lemma 5 below.

Sparse Voronoi Refinement   The meshing preprocess that we use is called Sparse Voronoi Refinement (SVR) and is due to Hudson et al. [14]. SVR is able to avoid the worst case complexity of Voronoi diagrams because it guarantees that the intermediate state is always a well-spaced point set and thus has a Voronoi diagram of linear complexity.

For the case of point sets in a bounding box, the SVR algorithm is easy to describe. It is an incremental construction that proceeds by alternating between two phases called **break** and **clean**. The break phase attempts to

add a Steiner point q at the farthest vertex of a Voronoi cell that contains an input point that has not yet been inserted. If there is an input point p too near to q, then p is added instead. The clean phase repeatedly attempts to add the farthest vertex of any cell with aspect ratio greater than τ until none are left. As in the break phase, if ever there is an input point too close, then it is added instead.

The running time of SVR is $O(n \log \Delta + |M|)$. The $n \log \Delta$ term comes from the point location data structure which associates each point with the Voronoi cell that contains it. When attempting to add a point, the nearby Voronoi cells are checked to see if any input points are nearby to insert instead.

Acar et al. developed an efficient implementation of SVR in 3-d [1]. It can also be efficiently parallelized [15]. Recently, Miller et al. showed that a variation of SVR runs in $O(n \log n + |M|)$ time by using a more complex point location scheme [20].

3   Algorithm

In this section, we describe the MESHVORONOI algorithm. It has three phases: a preprocess where the input points are placed in a bounding box and meshed, a removal phase that eliminates most of the Steiner points by incremental flipping, and a cleanup phase that removes the outer bounding box. The main data structure is a heap that stores the facets that are scheduled for removal by flipping. We call this the **flip heap**.

Throughout the algorithm, there is a global time parameter t. There will be a superset M of the input points P. At time t, M_t denotes the set M with weights, where the weight of a point p at time t is given as

$$w(p,t) := \left\{ \begin{array}{ll} \sqrt{t} & \text{if } p \in P \\ 0 & \text{otherwise} \end{array} \right.$$

Using this weighting scheme, there exists a sufficiently large t such that for all $t' > t$, $\text{Del}_{M_t} = \text{Del}_{M_{t'}}$. We use $M_\star$ to refer to M_t, where t is some such sufficiently large value.

3.1   Checking Potential Flips

A set S of $d+2$ points forms a **potential flip** at time t if $\text{ch}(S) \subset \text{Del}_{M_t}$. A potential flip can be identified with any of its interior facets. For example, in the plane, $(2,2)$-flips are usually associated with the edge in the interior
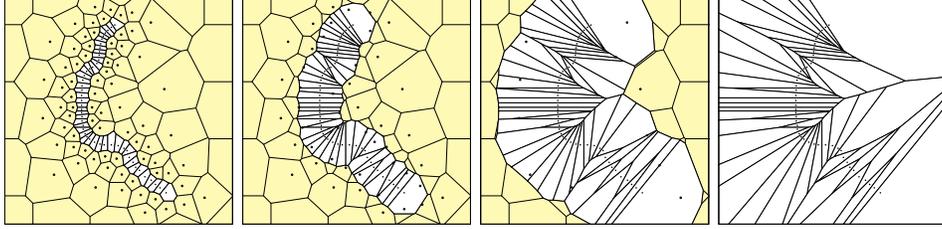
Figure 3: An illustration of the algorithm from left to right. Starting with a point set, it is extended to a well-spaced superset. The cells of the Steiner points are shaded. Then, the weights of the input points are increased until the extra cells disappear.

of a convex quadrilateral that will get replaced during the flip. In general, the facet representing a potential flip stands for the $d+2$ points comprising the two simplices sharing that facet.

The weights of the points of M vary with the time parameter t. For $p \in M$, the power distance from $x \in \mathbb{R}^d$ to p at time t is given as

$$\pi_p(x,t) := \|p - x\|^2 - w(p,t)^2 = \left\{ \begin{array}{ll} \|p - x\|^2 - t & \text{if } p \in P \\ \|p - x\|^2 & \text{otherwise} \end{array} \right.$$

So, for $d+2$ points $p_1, \ldots, p_{d+2}$ comprising a potential flip, the time associated with the flip is the value of t such that

$$\det \left[\begin{array}{ccc} p_1 & \cdots & p_{d+2} \\ \pi_{p_1}(0,t) & \cdots & \pi_{p_{d+2}}(0,t) \\ 1 & \cdots & 1 \end{array} \right] = 0.$$

The determinant on the left hand side is a linear function of t, so it is easy to solve for t. When evaluating a potential flip, the algorithm simply checks if the time associated with the flip is before or after the current time. This approach to viewing linear, geometric predicates as polynomial functions of time is well established in the area of kinetic data structures [12]. In our case, only one row of the matrix is changing with time, and the change is linear, so there is no need to solve higher degree polynomial systems as is common in other kinetic data structures problems.

3.2   Preprocessing

The first step in the algorithm is to add a constant number of points to form a bounding region around the input. This can be done so that the

total complexity of the convex hull of the augmented point set is a constant. Second, the Sparse Voronoi Refinement algorithm adds $O(n \log \Delta)$ Steiner points to produce a τ-well-spaced superset M, for a constant $\tau > 2$ (choosing $\tau = 3$ is reasonable). Third, the facets of the Delaunay triangulation of M are each checked for a potential flip and added to the flip heap accordingly.

3.3 Flipping out Steiner points

We maintain the weighted Delaunay triangulation through a sequence of incremental flips induced by the changes in weights. While the flip heap is nonempty, we pop the next potential flip. The time t is set to be the time of this potential flip. If the flip is still valid, i.e. all of the relevant facets are still present in the weighted Delaunay triangulation at time t, then we perform the flip. Otherwise, we do nothing and continue. If any new facets are introduced, we check them for potential flips and add them to the flip heap accordingly. Then we loop.

3.4 Postprocessing

To complete the construction, the algorithm removes all boundary vertices and all incident Delaunay simplices.

4 Analysis

4.1 Boundary Issues

The weighting will not remove all of the Steiner points. However, as the following lemma shows, only the boundary vertices will remain.

Lemma 1 (Only boundary vertices remain). *If $q \in \mathrm{Del}_{M_t}$ for some Steiner point q and all $t \geq 0$, then q is a boundary vertex.*

Proof. Let p be any input point and let q be any non-boundary Steiner point. Let

$$t_\star = \max_{x \in \mathrm{Vor}_M(q)} \|p - x\|^2.$$

Such a maximum exists because the Voronoi cells of non-boundary vertices are compact. Suppose for contradiction that $q \in \mathrm{Del}_{M_t}$ for some $t > t_\star$. Since $q \in \mathrm{Del}_{M_t}$, there must exist some point $y \in \mathrm{Vor}_{M_t}(q)$. It follows that $y \in \mathrm{Vor}_M(q)$ as well and so $\|p - y\|^2 \leq t_\star$. We now observe that

$$\pi_p(y, t) = \|p - y\|^2 - t \leq t_\star - t < 0 \leq \|q - y\|^2 = \pi_q(y, t),$$

contradicting the assumption that $y \in \mathrm{Vor}_{M_t}(q)$. $\hfill\square$

We need to show that removing the boundary vertices in a naive way as a post-process gives the correct output. For this, it will suffice to show the following lemma.

Lemma 2 (Induced Subcomplex). *There exists $t_\star$ such that for all $t > t_\star$, Del_P is an induced subcomplex of Del_{M_t}.*

Proof. For all $t \geq 0$, and each $p \in P$, we have $p \in \mathrm{Vor}_{M_t}(p)$. That is, input points are always their own nearest neighbors as weights increase. This implies that the vertices of Del_P are all present in Del_{M_t}.

Suppose for contradiction that some simplex σ is in $\mathrm{Del}_P \backslash \mathrm{Del}_{M_t}$ for some $t > r^2$, where r is the circumradius of σ. Then there must be some Steiner point q encroaching the orthoball of σ at time t. Let x be the circumcenter of σ. Since σ is composed only of input points, x is also the orthocenter of σ for all times t. For all $p \in \sigma$, we have

$$\pi_p(x, t) = \|p - x\|^2 - t = r^2 - t < 0 \leq \|q - x\|^2 = \pi_q(x, t).$$

It follows that q does not encroach the orthoball of σ at time t. This contradiction implies that $\mathrm{Del}_P \subset \mathrm{Del}_{M_{t_\star}}$ as long as $t_\star$ is greater than the largest squared circumradius of any simplex in Del_P.

To show that Del_P is an induced subcomplex and complete the proof, we observe that because Del_P and Del_{M_t} are embedded simplicial complexes covering the convex closure of P, there can be no simplices in Del_{M_t} containing only vertices of P that are not already included in Del_P. $\hfill\square$

Finally, the last important fact to check is that there is not too much extra work to put the points in a bounding domain. In meshing, this slack between the bounding box and the input is called scaffolding [16].

Lemma 3 (Bounded Scaffolding). *The number of simplices removed in the final step of the algorithm is $O(f)$. That is, $|\mathrm{Del}_{M_\star} \setminus \mathrm{Del}_P| = O(f)$.*

Proof. Any simplex $\sigma \in \mathrm{Del}_{M_\star} \setminus \mathrm{Del}_P$ can be written as a disjoint union $\sigma = S \sqcup T$ where $S \in \mathrm{ch}(M)$ and $T \in \mathrm{ch}(P)$. By construction, $\mathrm{ch}(M)$ has a constant number of vertices, and since M is well-spaced, $|\mathrm{ch}(M)|$ is also a constant. Since $\mathrm{ch}(P) \subseteq \mathrm{Del}_P$, we have $|\mathrm{ch}(P)| = O(f)$. It follows that there can be at most $|\mathrm{ch}(M)| \cdot |\mathrm{ch}(P)| = O(f)$ such simplices. $\hfill\square$

4.2 Running time analysis

Theorem 4 (Running Time). *Given n points $P \subset \mathbb{R}^d$, the* MeshVoronoi *algorithm constructs the Voronoi diagram of P in $O(f \log \Delta \log n)$ time, where f is the number of faces of* Vor_P.

Before proceeding to the proof of the running time guarantee, we will first bound the number of combinatorial changes of the weighted Voronoi diagram and thus also the number of heap operations. These are the main technical points of the analysis.

We first prove a relatively standard packing bound on the number of Voronoi cells of Vor_M that can intersect the Voronoi cell of an input point.

Lemma 5 (Packing). *Let M be a τ-well-spaced superset of P satisfying the feature size condition that $\mathbf{f}_P(q) \leq K\mathbf{f}_M(q)$ for all $q \in M$ for some constants τ and K. Let p be any point of P. Then, the number of points $q \in M$ such that $\mathrm{Vor}_M(q) \cap \mathrm{Vor}_P(p) \neq \emptyset$ is $O(1 + \log(\mathrm{aspect}_P(p)))$.*

Proof. The proof will be by a volume packing argument. Let M_p denote the set of points $q \in M \setminus \{p\}$ such that $\mathrm{Vor}_M(q) \cap \mathrm{Vor}_P(p) \neq \emptyset$. For any $q \in M_p$ and $x \in \mathrm{Vor}_M(q) \cap \mathrm{Vor}_P(p)$, we derive the following bound on the distance between p and q.

$$
\begin{aligned}
\|p - q\| &\leq \|q - x\| + \|p - x\| && \text{[by the triangle inequality]} \\
&\leq \|q - x\| + \mathbf{f}_P(x) && [x \in \mathrm{Vor}_P(p)] \\
&\leq 2\|q - x\| + \mathbf{f}_P(q) && [\mathbf{f}_P \text{ is 1-Lipschitz}] \\
&\leq 2\tau\mathbf{f}_M(q) + \mathbf{f}_P(q) && [\text{since } M \text{ is } \tau\text{-well-spaced}] \\
&\leq (2\tau + K)\mathbf{f}_M(q) && [\text{by the feature size condition}]
\end{aligned}
$$

Define $\gamma := \frac{1}{4\tau + 2K}$ so that the preceding bound may be written as

$$
\gamma\|p - q\| \leq \frac{1}{2}\mathbf{f}_M(q).
$$

We partition the set M_p into geometrically growing spherical shells A_i, where for each integer i,

$$
A_i := \{q \in M_p \mid (1 + \gamma)^{i-1} \leq \|p - q\| \leq (1 + \gamma)^i\}.
$$

For each $q \in M_p$, define the ball $B_q := \mathrm{ball}(q, \gamma\|p - q\|)$. The definition of B_q and the bound on $\|p - q\|$ imply

$$
B_q \subseteq \mathrm{ball}(q, \frac{1}{2}\mathbf{f}_M(q)) \subset \mathrm{Vor}_M(q),
$$

and so the balls B_q are pairwise disjoint. For $q \in A_i$, we further get that $B_q \subset \text{ball}(p, (1 + \gamma)^{i+1})$. Thus,

$$\text{Vol}(\text{ball}(p, (1 + \gamma)^{i+1})) \geq \text{Vol}\left(\bigsqcup_{q \in A_i} \text{Vol}(B_q)\right)$$

$$\geq |A_i| \text{Vol}(\text{ball}(q, \gamma(1 + \gamma)^{i-1})).$$

It follows that $|A_i| \leq (1 + \gamma)^{2d}/\gamma^d$. The nearest point of M_p to p has distance at least $\mathbf{f}_P(p)/K$ from p by the feature size condition. So, for $i < \lfloor \log_{1+\gamma}(\mathbf{f}_P(p)/K) \rfloor$, A_i is empty. The farthest point of M_p to p has distance at most $2 \, \text{aspect}_P(p)\mathbf{f}_P(p)$ from p by the triangle inequality and the definition of M_p. So, for $i > \lceil \log_{1+\gamma}(2 \, \text{aspect}_P(p)\mathbf{f}_P(p)) \rceil$, A_i is empty. Thus, there are at most

$$\lceil \log_{1+\gamma}(2 \, \text{aspect}_P(p)\mathbf{f}_P(p)) \rceil - \lfloor \log_{1+\gamma}(\mathbf{f}_P(p)/K) \rfloor$$

$$= O(1 + \log(\text{aspect}_P(p)))$$

nonempty sets A_i. This completes the proof as we have shown M_p can be decomposed into $O(1 + \log(\text{aspect}_P(p)))$ sets, each of constant size. $\quad\square$

Counting Flips The main challenge in the analysis is to bound the number of flips that happen in the transformation from the Voronoi diagram of M to the Voronoi diagram of P. The key to bounding this number is to observe that each such flip is witnessed by the intersection of a k-face of Vor_M and a $(d - k)$-face of Vor_P for some k. Thus, we count these intersections instead. This intuition is made precise in the following lemmas. The first bounds the number of flips performed by the algorithm. The latter bounds the number of potential flips considered by the algorithm, since not all potential flips are performed.

Lemma 6 (Flip Bound). *Given n points $P \subset \mathbb{R}^d$, the MESHVORONOI algorithm performs $O(f \log \Delta)$ flips, where $f = |\text{Vor}_P|$.*

Proof. We observe that a flip occurs exactly when the weights cause some pair of adjacent simplices to encroach on each other. That is, the two simplices share a common orthoball. Let U be the $d + 2$ vertices comprising these simplices and let c be the center of their orthoball. Let $k + 1$ be the number Steiner points in U. The weights of the input points in U are all equal and thus these points are all also equidistant from c. So, c is contained in a k-face of Vor_P. The Steiner points of U have weight 0, so

14

they are equidistant from c and closer to c than any of the input points. So, c is contained in a $(d-k)$-face of Vor_M, where M is the bounded aspect ratio superset of P constructed by the algorithm. Thus, the point c is the intersection of a k-face of Vor_P and a $(d-k)$-face Vor_M. It will suffice to bound the number of such intersections to get an upper bound on the number of flips. We will show that for each of the f faces of Vor_P, there are $O(\log \Delta)$ intersections with Vor_M to be counted.

Let F be any k-face of Vor_P that intersects a $(d-k)$-face of Vor_M. Then, F also intersects a d-face of Vor_M. Let $\text{Vor}_P(p)$ be a d-face of Vor_P containing F. Since the d-faces of Vor_M have only a constant number of faces, it will suffice to bound intersections between $\text{Vor}_P(p)$ and the d-faces of Vor_M. From Lemma 5, there are at most $O(\log \Delta)$ such intersections. $\quad\square$

Lemma 7 (Potential Flip Bound)**.** *Given n points $P \subset \mathbb{R}^d$, the* MeshVoronoi *algorithm sees $O(f \log \Delta)$ potential flips, where $f = |\text{Vor}_P|$.*

Proof. At the start of the algorithm, there is one potential flip for each facet of Vor_M. This is at most $O(n \log \Delta)$. During the rest of the algorithm, there are at most $\binom{d+2}{d} = O(d^2)$ new potential flips each time a real flip occurs, one for each new facet that appears. By Lemma 6, this is $O(f \log \Delta)$. $\quad\square$

The main result.　We are now ready to prove the running time guarantee, Theorem 4.

Proof of Theorem 4. It will suffice to bound the running time of each phase of the algorithm. The preprocessing takes $O(n \log \Delta)$ from the running time of Sparse Voronoi Refinement. Seeding the heap also takes $O(n \log \Delta)$ time as each facet is checked in constant time and the amortized cost of heap insertion is $O(1)$. There are at most two heap operations for each potential flip, one insertion and one deletion. The total number of potential flips is $O(f \log \Delta)$ as shown in Lemma 7. Deleting the minimum element from a heap with $O(f \log \Delta)$ elements requires $O(\log(f \log \Delta)) = O(\log n)$ time. Thus, the total time for all heap operations is $O(f \log \Delta \log n)$ as desired. $\quad\square$

5　Conclusion

The algorithm we have presented is a direct combination of Delaunay mesh generation and kinetic data structures. The output-sensitive running time depends on the log of the spread. This is the usual cost of doing a geometric divide and conquer. It remains an interesting question if it is possible to

replace the $\log \Delta$ term with a $\log n$ by some more combinatorial divide and conquer. The other log-term coming from the heap operations may also permit some improvement as it is clear that many flips are geometrically independent and so their ordering is not strict.

Another possible direction of future work is to exploit hierarchical meshes to keep the number of Steiner points linear [20]. However, it is not known how to leverage this into an improvement over the bounds presented here. At best it replaces the $\log \Delta$ with $\max\{\log \Delta, n\}$.

Going forward, we hope to extend the methods here to the convex hull problem. Likely, this will require significant new ideas, but it may be possible to achieve a similar output-sensitive running time of $O(f \log \Delta \log n)$ for computing the convex hull of n points with f faces.

References

[1] Umut A. Acar, Benoît Hudson, Gary L. Miller, and Todd Phillips. SVR: Practical engineering of a fast 3D meshing algorithm. In *Proc. 16th International Meshing Roundtable*, pages 45–62, 2007.

[2] Nancy M. Amato and Edgar A. Ramos. On computing voronoi diagrams by divide-prune-and-conquer. In *Symposium on Computational Geometry*, pages 166–175, 1996.

[3] Franz Aurenhammer. Voronoi diagrams–a survey of a fundamental geometric data structure. *ACM Comput. Surv.*, 23(3):345–405, September 1991.

[4] David Avis and Komei Fukuda. A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra. *Discrete & Computational Geometry*, 8(1):295–313, 1992.

[5] David Bremner. Incremental convex hull algorithms are not output sensitive. *Discrete & Computational Geometry*, 21(1):57–68, 1999.

[6] Timothy M. Chan. Output-sensitive results on convex hulls, extreme points, and related problems. *Discrete & Computational Geometry*, 16(4):369–387, 1996.

[7] Timothy M. Chan, Jack Snoeyink, and Chee-Keng Yap. Primal dividing and dual pruning: Output-sensitive construction of four-dimensional polytopes and three-dimensional voronoi diagrams. *Discrete & Computational Geometry*, 18(4):433–454, 1997.

[8] Bernard Chazelle, Olivier Devillers, Ferran Hurtado, Mercè Mora, Vera Sacristán, and Monique Teillaud. Splitting a Delaunay triangulation in linear time. *Algorithmica*, 34:39–46, 2002.

[9] Bernard Chazelle and Wolfgang Mulzer. Computing hereditary convex structures. *Discrete & Computational Geometry*, 45(4):796–823, 2011.

[10] Siu-Wing Cheng, Tamal K. Dey, and Jonathan Richard Shewchuk. *Delaunay Mesh Generation*. CRC Press, 2012.

[11] Herbert Edelsbrunner and Nimish R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica*, 15, 1996.

[12] Leonidas Guibas. Kinetic data structures. In Dinesh P. Mehta and Sartaj Sahni, editors, *Handbook of Data Structures and Applications*. CRC Press, 2005.

[13] Sariel Har-Peled and Manor Mendel. Fast construction of nets in low dimensional metrics, and their applications. *SIAM Journal on Computing*, 35(5):1148–1184, 2006.

[14] Benoît Hudson, Gary Miller, and Todd 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.

[15] Benoît Hudson, Gary L. Miller, and Todd Phillips. Sparse Parallel Delaunay Refinement. In *19th Annual ACM Symposium on Parallelism in Algorithms and Architectures*, pages 339–347, San Diego, June 2007.

[16] Benoît Hudson, Gary L. Miller, Todd Phillips, and Donald R. Sheehy. Size complexity of volume meshes vs. surface meshes. In *SODA: ACM-SIAM Symposium on Discrete Algorithms*, 2009.

[17] Michael Joswig and Günter M. Ziegler. Convex hulls, oracles, and homology. *J. Symb. Comput*, 38(4):1247–1259, 2004.

[18] V. Klee. On the complexity of d-dimensional voronoi diagrams. *Archiv der Mathematik*, 34:75, 1980.

[19] Jiří Matoušek and Otfried Schwarzkopf. Linear optimization queries. In *Symposium on Computational Geometry*, pages 16–25, 1992.

[20] Gary L. Miller, Todd Phillips, and Donald R. Sheehy. Beating the spread: Time-optimal point meshing. In *SOCG: Proceedings of the 26th ACM Symposium on Computational Geometry*, pages 321–330, 2011.

[21] Gary L. Miller, Dafna Talmor, Shang-Hua Teng, and Noel Walkington. On the radius-edge condition in the control volume method. *SIAM J. on Numerical Analysis*, 36(6):1690–1708, 1999.

[22] Raimund Seidel. Constructing higher-dimensional convex hulls at logarithmic cost per face. In *STOC: ACM Symposium on Theory of Computing*, 1986.

[23] Raimund Seidel. On the number of faces in higher-dimensional Voronoi diagrams. In *Proceedings of the 3rd Annual Symposium on Computational Geometry*, pages 181–185, 1987.

[24] Donald R. Sheehy. *Mesh Generation and Geometric Persistent Homology*. PhD thesis, Carnegie Mellon University, 2011.

[25] Donald R. Sheehy. New Bounds on the Size of Optimal Meshes. *Computer Graphics Forum*, 31(5):1627–1635, 2012.

[26] Garret Swart. Finding the convex hull facet by facet. *Journal of Algorithms*, 6(1):17–48, 1985.