



A step in the Delaunay mosaic of order k

Herbert Edelsbrunner, Anton Nikitenko, and Georg Osang 

Abstract. Given a locally finite set $X \subseteq \mathbb{R}^d$ and an integer $k \geq 0$, we consider the function $\mathbf{w}_k: \text{Del}_k(X) \rightarrow \mathbb{R}$ on the dual of the order- k Voronoi tessellation, whose sublevel sets generalize the notion of alpha shapes from order-1 to order- k (Edelsbrunner et al. in IEEE Trans Inf Theory IT-29:551–559, 1983; Krasnoshchekov and Polishchuk in Inf Process Lett 114:76–83, 2014). While this function is not necessarily generalized discrete Morse, in the sense of Forman (Adv Math 134:90–145, 1998) and Freij (Discrete Math 309:3821–3829, 2009), we prove that it satisfies similar properties so that its increments can be meaningfully classified into critical and non-critical steps. This result extends to the case of weighted points and sheds light on k -fold covers with balls in Euclidean space.

Keywords. Discrete Morse theory, Order- k Voronoi tessellations, Hyperplane arrangements, Order- k Delaunay mosaics, Rhomboid tilings, Weighted points, Power distance.

1. Introduction

Given a locally finite set $X \subseteq \mathbb{R}^d$ and non-negative integer k , the *order- k Voronoi tessellation* decomposes \mathbb{R}^d into closed convex domains in which the k nearest points in X are the same [9, 14]. In other words, knowing the domain that contains a point $x \in \mathbb{R}^d$ is equivalent to knowing which k points of X are nearest to x . This motivates the use of the tessellation as a data structure for the k -nearest neighbor problem [3]. Similar to the ordinary (order-1) Voronoi tessellation, the order- k Voronoi tessellation has a natural dual [1], which we refer to as the *order- k Delaunay mosaic*, denoted $\text{Del}_k(X)$. Its cells represent collections of domains with non-empty common intersection. We are interested

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 78818 Alpha), and by the DFG Collaborative Research Center TRR 109, ‘Discretization in Geometry and Dynamics’, through Grant No. I02979-N35 of the Austrian Science Fund (FWF).

in the function $\mathbf{w}_k: \text{Del}_k(X) \rightarrow \mathbb{R}$ that maps each cell to the minimum radius, r , such that the corresponding intersection of domains contains a point at distance at most r from each one of its k nearest neighbors. The sublevel sets of \mathbf{w}_k generalize the notion of alpha shapes from $k = 1$ to orders $k \geq 1$ [5, 13]. Recently, the stochastic properties of \mathbf{w}_k have been studied [6] and algorithms for computing the persistence have been presented [7]. We shed additional light on these results by establishing that \mathbf{w}_k behaves similar to a discrete Morse function. We hasten to mention that \mathbf{w}_k neither satisfies the requirements of a *discrete Morse function* [10] nor the slightly weaker requirements of a *generalized discrete Morse function* [11]. Nevertheless, we can classify the increments in the sublevel sets into critical and non-critical steps with predictable impact on the homotopy type. To state the result, we note that each cell of the order- k Delaunay mosaic is spanned by several size- k subsets of X , and we shall determine how to characterize incremental steps of \mathbf{w}_k in terms of such subsets. Specifically, taking any set of $\ell + 1 \leq d + 1$ points in X , there is the unique smallest sphere that passes through these points. Assuming general position, the convex hull of these points is an ℓ -simplex. Some of the size- k subsets of the $\ell + 1$ points defining the sphere together with the points of X inside the sphere form cells of the Delaunay mosaic, and some of these cells constitute the unique step of \mathbf{w}_k corresponding to this sphere; see Sect. 2. Our main result is the following classification of the topology types of the steps:

- We call the configuration that defines a step *self-centered* if the center of the corresponding sphere is contained in the simplex spanned by the points on the sphere. Adding the cells in the step changes the Euler characteristic of the sublevel set, which implies that it also changes the homotopy type, so we refer to it as a *critical step* of \mathbf{w}_k .
- We call the configuration *altruistic* if the center of the corresponding sphere is not contained in the simplex spanned by the points on the sphere. Adding the cells in the step preserves the homotopy type, so we refer to it as a *non-critical step* of \mathbf{w}_k .

With minor adjustments, this classification extends to the case in which the points have real weights and the squared Euclidean distance is replaced by the power distance. This is most transparent when we view the Voronoi tessellations in \mathbb{R}^d as projections of the levels in a hyperplane arrangement in \mathbb{R}^{d+1} ; see e.g. [8]. Correspondingly, we view the Delaunay mosaics as horizontal slices of a rhomboid tiling in \mathbb{R}^{d+1} ; see [7] for the construction in the unweighted case, which readily extends to points with weights.

Outline Section 2 introduces background on weighted Voronoi tessellations and hyperplane arrangements, their dual Delaunay mosaics and rhomboid tilings, and discrete Morse functions. Section 3 gives the proof of the main result, thus extending the framework of discrete Morse theory to include squared radius functions on order- k Delaunay mosaics. Section 4 concludes the paper.

2. Background

It will be useful to see Voronoi tessellations as levels in hyperplane arrangements and Delaunay mosaics as slices of rhomboid tilings. We note that the different terminology for the *tilings*, *mosaics* and *tessellations* has roots in the different origins of the established terms, and we take advantage of it to simplify the distinction between the studied objects. In general, all these three notions refer to cell complexes whose cells are (possibly unbounded) convex polyhedra. To avoid redundancies, we explain everything for the more general, weighted case.

Hyperplane arrangements Let X be a locally finite set of points with real weights in \mathbb{R}^d , and let J be the corresponding *index set*. The *power distance* of a point $x \in \mathbb{R}^d$ from a weighted point $(x_j, w_j) \in X \subseteq \mathbb{R}^d \times \mathbb{R}$ is $\pi_j(x) = \|x - x_j\|^2 - w_j$. For $I \subseteq J$, the corresponding (*weighted*) *Voronoi domain* is the set of points that satisfy $\pi_i(x) \leq \pi_j(x)$ for all $i \in I$ and $j \in J \setminus I$. For each non-negative integer k , the (*weighted*) *order- k Voronoi tessellation*, denoted $\text{Vor}_k(X)$, is the collection of Voronoi domains for sets I that satisfy $|I| = k$. Setting $w_j = 0$ for all indices j , we get the unweighted situation as a special case.

To illuminate the structure of the order- k Voronoi tessellation, let $\varpi, f_j: \mathbb{R}^d \rightarrow \mathbb{R}$ be defined by mapping $x \in \mathbb{R}^d$ to $\varpi(x) = \frac{1}{2}\|x\|^2$ and to $f_j(x) = \langle x, x_j \rangle - \frac{1}{2}[\|x_j\|^2 - w_j]$. The graph of ϖ is a paraboloid in \mathbb{R}^{d+1} , and the graph of f_j is the hyperplane that touches the shifted paraboloid defined by $\varpi + \frac{1}{2}w_j$ in the point $(x_j, \varpi(x_j) + \frac{1}{2}w_j)$. We refer to the collection of hyperplanes as the *arrangement* of X . Let now $S = S(x, w)$ be the $(d-1)$ -dimensional *sphere* with center $x \in \mathbb{R}^d$ and squared radius $w \in \mathbb{R}$. For $w > 0$ this is an ordinary sphere, for $w = 0$ it is a point, and for $w < 0$ it is what we call an *imaginary sphere*. There are different ways of visualizing the latter concept, but singularly important in this paper is that its squared radius is negative, giving the correct power distance if plugged into the formula given above. In either case, S partitions the index set into three subsets:

$$\text{In}(S) = \{j \in J \mid \pi_j(x) < w\}, \tag{1}$$

$$\text{On}(S) = \{j \in J \mid \pi_j(x) = w\}, \tag{2}$$

$$\text{Out}(S) = \{j \in J \mid \pi_j(x) > w\}. \tag{3}$$

In the unweighted case, when $w_j = 0$ for all indices j , the three sets correspond to the points inside, on, and outside the sphere. In the weighted case, the condition in (2) can be rewritten as $\|x - x_j\|^2 = w + w_j$, which we geometrically interpret as having two spheres, $S(x, w)$ and $S(x_j, w_j)$, that intersect at a right angle. It is easy to see that $i \in \text{In}(S), \text{On}(S), \text{Out}(S)$ iff $f_j(x) - \frac{1}{2}[\|x\|^2 - w] = \frac{1}{2}[w_j + w - \|x - x_j\|^2]$ is positive, zero, negative or, equivalently, iff the point $y(x) = (x, \frac{1}{2}[\|x\|^2 - w])$ in \mathbb{R}^{d+1} lies below, on, above the graph of f_j . This motivates us to consider the decomposition of \mathbb{R}^{d+1} defined by the hyperplanes. For each partition of J into three sets, we consider the corresponding partition

of the set of the hyperplanes in \mathbb{R}^{d+1} , and define the corresponding *cell* to consist of points that at the same time are *on or under* the hyperplanes in the first set, *on* the hyperplanes in the second set, and *on or above* the hyperplanes in the third set. Important for us is the case when the three-partition is defined by a fixed sphere S . Then the corresponding cell consists of points $(x, z) \in \mathbb{R}^{d+1}$ that satisfy $z \leq f_j(x)$ if $j \in \text{In}(S)$, $z = f_j(x)$ if $j \in \text{On}(S)$, and $z \geq f_j(x)$ if $j \in \text{Out}(S)$. Assuming general position, this cell has dimension $d - p = d + 1 - |\text{On}(S)|$. We refer to the $(d - p)$ -dimensional cells as $(d - p)$ -*cells* and to the $(d + 1)$ -cells as *chambers*. We write $\text{Arr}(X)$ for the set of all cells. The sphere defines a chamber if $\text{On}(S) = \emptyset$, and in this case we call $k = |\text{In}(S)|$ the *depth* of the chamber, because there are precisely k hyperplanes above it. The relation between the chambers and the Voronoi domains should be clear.

Proposition 2.1. (Chambers and Domains [8]) *Let X be a locally finite set of points with real weights in \mathbb{R}^d . For every non-negative integer k , there is a bijection between the domains of $\text{Vor}_k(X)$ and the chambers at depth k of $\text{Arr}(X)$ such that every Voronoi domain is the vertical projection of its corresponding chamber to \mathbb{R}^d . \square*

Rhomboid tiling To dualize the Voronoi tessellations, we generalize the construction of [1] to the unweighted case. Specifically, we map every domain of $\text{Vor}_k(X)$ to the sum of the corresponding k points, for the moment ignoring the weights. For every non-empty common intersection of domains, we collect the images of the domains that contain this intersection, and we add the convex hull of these points as a cell to the dual, which we refer to as the *order- k Delaunay mosaic* of X , denoted $\text{Del}_k(X)$. Note that this definition is not coordinate invariant, but it can be made so by substituting the average for the sum, which is just a rescaling. We refrain from doing this to simplify the notation.

We give an alternate description of these mosaics after dualizing the hyperplane arrangement. To this end, we write $y_j = (x_j, -1) \in \mathbb{R}^{d+1}$, for every $(x_j, w_j) \in X$, and $y_I = \sum_{i \in I} y_i$, for every $I \subseteq J$. The $(d + 1)$ -st coordinate of y_I is $-|I|$, and we call $k = |I|$ the *depth* of the point. For every sphere S in \mathbb{R}^{d+1} — which we recall may be a point or imaginary — we let $\text{rho}(S) = \text{conv}\{y_I \mid \text{In}(S) \subseteq I \subseteq \text{In}(S) \cup \text{On}(S)\}$ be the *rhomboid* of S . It is *anchored* at the vertex with least depth, which is $y_{\text{In}(S)}$, and it is *spanned* by the vectors y_i with $i \in \text{On}(S)$. We say X is in *general position* if for every collection of $d + 2 - p \geq 1$ spheres S there are at most $p + 1$ weighted points that belong to all $d + 2 - p$ sets $\text{On}(S)$. For such a set X , the dimension of the rhomboid defined by a sphere S is the number of vectors that span it, which is $\dim \text{rho}(S) = |\text{On}(S)|$. The *rhomboid tiling* of X , denoted $\text{Rho}(X)$, is the collection of all rhomboids defined by spheres in \mathbb{R}^{d+1} ; see Fig. 1. As argued for the unweighted case in [7], $\text{Rho}(X)$ is a *polytopal complex* in \mathbb{R}^{d+1} , which means its cells are convex polytopes with disjoint interiors such that the boundary of every polytope is the union of other polytopes in the complex [15]. Indeed, we claim the following properties:




FIGURE 1 The rhomboid tiling of five weighted points on the real line. The weights are the squared radii of the colored half-circles. The horizontal line at depth k intersects the tiling in the order- k Delaunay mosaic of the points. Since D has much smaller weight than its two neighbors, it is not part of the order-1 mosaic. Nonetheless, D participates in all mosaics of order higher than one

Proposition 2.2. (Rhomboid Tiling) *Let X be a locally finite set of points with real weights in general position in \mathbb{R}^d . Then*

1. $\text{Rho}(X)$ is dual to $\text{Arr}(X)$;
2. $\text{Rho}(X)$ is a polytopal complex of rhomboids;
3. the horizontal slice at integer depth $k \geq 0$ is the order- k Delaunay mosaic of X .

We refer to the proof of the corresponding result for unweighted points in [7] and note that it readily generalizes to the weighted case. We explain Property 1 of Proposition 2.2 because it is repeatedly used. The duality is established by a bijection $\rho \mapsto \rho^*$ in which ρ is a $(p + 1)$ -dimensional rhomboid in $\text{Rho}(X)$ and ρ^* is a $(d - p)$ -cell in $\text{Arr}(X)$ such that $\rho \subseteq \varrho$ iff $\varrho^* \subseteq \rho^*$. The points inside, on, and outside the sphere that defines ρ correspond to the hyperplanes above, containing, under the cell ρ^* . Note also that the vertical partial order of the cells in the arrangement agrees with the same order on the rhomboids in the tiling.

Discrete Morse theory Next we assign to each rhomboid in the tiling a real value. Before doing so, we give a few definitions. Let K be a polytopal complex. Its *Hasse diagram* is a directed graph whose nodes are the cells in K and whose arcs are the pairs of cells $\sigma \subseteq \tau$ with $\dim \tau = \dim \sigma + 1$. A function $f : K \rightarrow \mathbb{R}$ is *monotonic* if $f(\sigma) \leq f(\tau)$ whenever $\sigma \subseteq \tau$. The *level set* for a value $w \in \mathbb{R}$ is the set of cells $f^{-1}(w) \subseteq K$. It is a set of nodes in the Hasse diagram. A *step* of f is a maximal subset of a level set whose induced subgraph in the Hasse diagram is connected. We note that the steps of f partition K . An *interval* of K is given by cells $\sigma \subseteq v$ and consists of all faces of v that share σ as a face, denoted $[\sigma, v] = \{\tau \in K \mid \sigma \subseteq \tau \subseteq v\}$. We call σ the *lower bound* and v the *upper bound* of $[\sigma, v]$. The interval is *singular* if $\sigma = v$. A monotonic function $f : K \rightarrow \mathbb{R}$ is *generalized discrete Morse* if every step is an interval; see [11].

For comparison, f is *discrete Morse* if every step is an interval of size 1 or 2; see [10] but note that the original definition is in-essentially more general by allowing $f(\sigma) > f(\tau)$ for pairs $\sigma \subseteq \tau$ in a step.

We are interested in the case in which $K = \text{Rho}(X)$ and f is the squared radius function, which we now define. Recall that each rhomboid $\rho \in \text{Rho}(X)$ has a dual cell $\rho^* \in \text{Arr}(X)$. For each point $y = (x, z) \in \rho^*$, we set $w(y) = \|x\|^2 - 2z$, and we define the *squared radius function*, $\mathbf{w}: \text{Rho}(X) \rightarrow \mathbb{R}$, by mapping ρ to $\mathbf{w}(\rho) = \min_{y \in \rho^*} w(y)$. To develop a geometric intuition for this function, we sweep the hyperplane arrangement from top to bottom with a paraboloid. Specifically, for $t \in \mathbb{R}$, the paraboloid is the graph of $\varpi - \frac{1}{2}t$. Then $\mathbf{w}(\rho)$ is the minimum $t \in \mathbb{R}$ such that the corresponding paraboloid has a non-empty intersection with ρ^* . Assuming general position, the sequence in which the paraboloid encounters the cells in $\text{Arr}(X)$ follows a few simple rules. For example, when the paraboloid encounters a vertex in $\text{Arr}(X)$, then it has already encountered $2^{d+1} - 1$ of the chambers incident to the vertex, and it touches the unique last incident chamber for the first time, as well as all faces of this chamber that share the vertex. More generally, when the paraboloid touches the intersection of $p + 1$ hyperplanes for the first time, this happens at an interior point of a $(d - p)$ -cell contained in this intersection. At the same time the paraboloid touches a unique chamber together with all faces of this chamber that share the $(d - p)$ -cell as a common face. We therefore have the following result.

Proposition 2.3. (Squared Radius Function) *Let X be a locally finite set of points with real weights in general position in \mathbb{R}^d . Then $\mathbf{w}: \text{Rho}(X) \rightarrow \mathbb{R}$ is a generalized discrete Morse function. Furthermore, every step of \mathbf{w} has a vertex as a lower bound, and there is only one singular interval, namely the vertex at the origin, which corresponds to the empty index set.*

The same claim restricted to the unweighted case can be found in [7]. Its proof extends with obvious modifications to the proof in the weighted case, which we therefore omit. This proposition says that the paraboloid always enters a chamber together with a subset of its faces while sweeping, thus defining a *shelling* of $\text{Arr}(X)$ (see [4]). This chamber corresponds to the vertex of a rhomboid and the faces of the chamber correspond to the faces of the rhomboid that share this vertex. As we will see shortly, the vertex is not necessarily the lowest vertex of the rhomboid. Note that \mathbf{w} is rather special because it has a limited collection of interval types. This is best seen by constructing $\text{Rho}(X)$ one step at a time. After starting with the vertex at the origin, each step glues a new rhomboid of dimension at least 1 together with all missing faces to the complex. Such a step preserves the homotopy type, which implies that every non-empty sublevel set of \mathbf{w} is contractible.

Horizontal integer slices Let P_k be the horizontal hyperplane at integer depth k in \mathbb{R}^{d+1} . By Proposition 2.2, the intersection of P_k with $\text{Rho}(X)$ gives the (weighted) order- k Delaunay mosaic. We aim at associating each cell in

$\text{Del}_k(X)$ with the rhomboid in $\text{Rho}(X)$ such that the cell is the intersection of P_k with the rhomboid, but there is an ambiguity for the vertices of $\text{Del}_k(X)$, which belong to several intersections. We will associate them to the vertices of $\text{Rho}(X)$, but to avoid special cases, we formulate the definitions for the rhomboids without their boundary. Note that the (relative) interiors of the rhomboids partition the union of the rhomboids. Accordingly, each cell $\sigma \in \text{Del}_k(X)$ is the closure of the intersection of P_k with the interior of a unique rhomboid $\rho = \rho(\sigma)$. Writing $p = \dim \sigma$, we have $\dim \rho(\sigma) = p + 1$ if $p \geq 1$, and $\dim \rho(\sigma) = 0$ if $p = 0$. Similarly, we get the *squared radius function* of $\text{Del}_k(X)$ from that of $\text{Rho}(X)$: $\mathbf{w}_k(\sigma) = \mathbf{w}(\rho(\sigma))$. We are interested in the partition of $\text{Del}_k(X)$ into the steps of \mathbf{w}_k . To this end, let σ and τ be two cells in $\text{Del}_k(X)$ and note that $\sigma \subseteq \tau$ iff $\rho(\sigma) \subseteq \rho(\tau)$. It follows that each step of \mathbf{w}_k is the horizontal slice of a step of \mathbf{w} .

Letting $\rho \in \text{Rho}(X)$, we write $\text{top}(\rho)$ and $\text{btm}(\rho)$ for the vertices with minimum and maximum depth, and we write $\text{last}(\rho)$ for the vertex with maximum value of \mathbf{w} . Proposition 2.3 implies that $\lambda = \text{last}(\rho)$ for every interval $[\lambda, \rho]$ of $\mathbf{w}: \text{Rho}(X) \rightarrow \mathbb{R}$. For some rhomboids, we have $\text{last}(\rho) = \text{btm}(\rho)$, but not necessarily for all. Depending on the shape of the rhomboid, λ can indeed be any vertex of ρ other than $\text{top}(\rho)$. We formally state this as a lemma:

Lemma 2.4. (Last not Top Vertex) *Let X be a locally finite set of points with real weights in general position in \mathbb{R}^d . Then $\lambda \neq \text{top}(\rho)$ for every non-singular interval $[\lambda, \rho]$ of \mathbf{w} .*

Indeed, the chamber in $\text{Arr}(X)$ that is dual to $\text{top}(\rho)$ lies above ρ^* . Since ρ is an upper bound, the point at which the paraboloid first touches ρ^* during the sweep is an interior point. Hence, $\text{top}(\rho)^*$ has a lower value of \mathbf{w} and therefore does not belong to the interval.

3. Topology of a step

This section proves our main result: that critical and non-critical steps of \mathbf{w}_k can be distinguished by whether $\text{last}(\rho)$ is equal to or different from $\text{btm}(\rho)$, with ρ the smallest rhomboid in $\text{Rho}(X)$ that contains the corresponding step of \mathbf{w} . The addition of a critical step changes the homotopy type of the sublevel set and the addition of a non-critical step preserves it. We begin with an enumeration of the types.

Step types By Proposition 2.3, every step of $\mathbf{w}: \text{Rho}(X) \rightarrow \mathbb{R}$ is an interval $[\lambda, \rho]$ in which ρ is the maximum rhomboid that satisfies $\lambda = \text{last}(\rho)$. The interval consists of all faces of ρ that share λ . Assuming $\dim \rho = p + 1 \geq 0$, the rhomboid has vertices at $p + 2$ depth values, and letting k be the depth value of $\text{btm}(\rho)$, these values are $k - g$ for $0 \leq g \leq p + 1$. By Lemma 2.4, λ can assume only $p + 1$ of these depth values. If $\lambda \neq \text{btm}(\rho)$, then P_{k-g} has a non-empty intersection with the interior of at least one rhomboid in the




FIGURE 2 A 3-rhomboid ρ with dashed silhouette separating the (gray) faces that share λ from the (transparent) other faces. Each slice is blue and shown together with the corresponding subgraph of the Hasse diagram of the Delaunay mosaic of the corresponding order. Left: the self-centered configurations whose corresponding critical steps consist of a triangle without its boundary, a triangle with its edges but without vertices, and a vertex. Middle and right: the altruistic configurations whose corresponding non-critical steps consist in both cases of a triangle with one edge, and a triangle with two edges and the shared vertex

interval for $1 \leq g \leq p$, and if $\lambda = \text{btm}(\rho)$, then there is one more, namely for $0 \leq g \leq p$. In total, we count $p^2 + p + 1$ possible types of slices; see Fig. 2 for an illustration of the types for $p = 2$. Some of these types are symmetric. We refer to the $p + 1$ slices in case $\lambda = \text{btm}(\rho)$ as *self-centered* and the p^2 other slices as *altruistic*. The terminology is motivated by the fact that $\text{last}(\rho) = \text{btm}(\rho)$ iff the convex hull of the points whose indices are in $\text{On}(S)$ contain the center of S . There is an ambivalent case, when the center lies on the boundary of the convex hull, but this can be prevented by slightly strengthening the general position assumption to forbid this situation.

Topology type Letting $A \subseteq \text{Rho}(X)$ be a step of \mathbf{w} , we write $P_k \cap A \subseteq \text{Del}_k(X)$ for the corresponding step of \mathbf{w}_k . The *Euler characteristic* of A is $\chi(A) = \sum_{\rho \in A} (-1)^{\dim \rho}$. Since A is necessarily an interval, its Euler characteristic vanishes, unless $A = \{0\}$, in which case it is 1. The Euler characteristic of the slice is $\chi(P_k \cap A) = \sum_{\tau \in P_k \cap A} (-1)^{\dim \tau}$, which may or may not be zero. We write $|A|$ for the union of interiors of the rhomboids in A , and $P_k \cap |A|$ for its slice at depth k . Let $\mathbb{H}^p \subseteq \mathbb{R}^p$ be the set of points with non-negative first coordinate, and note that $\chi(\mathbb{H}^p) = 0$ for all $p \geq 1$. Two topological spaces have the same *topology type* if there is a homeomorphism between them, and in this case they have the same Euler characteristic. For example, the half-open interval, $[0, 1)$, has the same topology type as \mathbb{H}^1 , which we denote as $[0, 1) \approx \mathbb{H}^1$. We represent $[0, 1)$ by an edge together with one of its endpoints, so the Euler characteristic, which is the alternating sum of cells vanishes. We will see that every altruistic configuration has the topology type of \mathbb{H}^p , for some value of p , while every self-centered configuration has non-zero Euler characteristic.




FIGURE 3 Three self-centered configurations in \mathbb{R}^3 . From left to right: a tetrahedron without boundary faces, an octahedron with four of its triangles but no other faces, and a tetrahedron with all of its faces except for the vertices

Theorem 3.1. (Topology of a Step) *Let X be a locally finite set of points with real weights in general position in \mathbb{R}^d , let $A = [\lambda, \rho]$ be a step of \mathbf{w} , set $p + 1 = \dim \rho$, write k for the depth of $\text{btm}(\rho)$, and recall that $P_{k-g} \cap A$ is a step of \mathbf{w}_{k-g} .*

1. *If $\text{last}(\rho) = \text{btm}(\rho)$, then $\chi(P_{k-g} \cap A) \neq 0$ for $0 \leq g \leq p$.*
2. *If $\text{last}(\rho) \neq \text{btm}(\rho)$, then $P_{k-g} \cap |A| \approx \mathbb{H}^p$ and therefore $\chi(P_{k-g} \cap A) = 0$ for $1 \leq g \leq p$.*

All other horizontal integer slices of A are empty.

Proof. We first consider the self-centered configurations, when $\lambda = \text{last}(\rho) = \text{btm}(\rho)$. For $g = 0$, the hyperplane P_{k-g} contains λ and avoids the interiors of all other rhomboids in $A = [\lambda, \rho]$. The Euler characteristic of this slice is one and therefore non-zero, as claimed. For $1 \leq g \leq p$, P_{k-g} has non-empty intersections with the interiors of the rhomboids of dimension larger than g and empty intersections with the interiors of all other rhomboids in the interval; see the left panel of Fig. 2 for the three cases that occur for $p = 2$, and see Fig. 3 for three of the four cases that occur for $p = 3$. Therefore, $\chi(P_{k-g} \cap A) = \sum_{q=g+1}^{p+1} (-1)^q \binom{p+1}{q}$, in which the binomial coefficient is the number of q -dimensional faces of a $(p + 1)$ -dimensional rhomboid that share a common vertex, namely λ . This sum evaluates to $(-1)^{g+1} \binom{p}{g}$, which is non-zero, as claimed.

We second consider the altruistic configurations, when $\lambda \neq \text{btm}(\rho)$. Let λ' be the vertex of ρ opposite to λ , and project ρ orthogonally to the hyperplane normal to $\lambda' - \lambda$. The projection is a p -dimensional convex polytope. Call the preimage of its (relative) boundary the *silhouette* of ρ , and note that it is a $(p - 1)$ -dimensional topological sphere that contains all vertices of ρ other than λ and λ' ; see Fig. 2. None of the rhomboids in the silhouette belong to $A = [\lambda, \rho]$. In fact, the silhouette separates the boundary rhomboids of ρ that are in this interval from the boundary rhomboids that are not in the interval. Since $\text{btm}(\rho)$ and $\text{top}(\rho)$ belong to the silhouette, P_k and $P_{k-(p+1)}$ both have empty *intersection* with the interiors of all rhomboids in $[\lambda, \rho]$, as claimed.

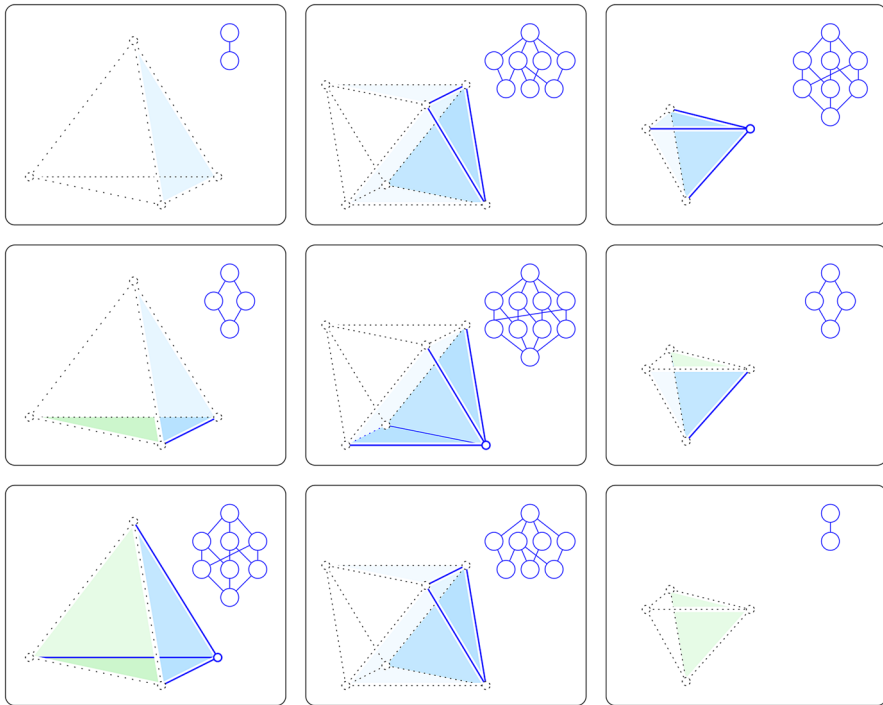


FIGURE 4 The 9 altruistic configurations in \mathbb{R}^3 . Compared to $\text{btm}(\rho)$ at depth k , the vertex $\lambda = \text{last}(\rho)$ has depth $k - 3$, $k - 2$, and $k - 1$ in the *top*, *middle*, and *bottom* row. Similarly, the slice is at depth $k - 3$, $k - 2$, and $k - 1$ in the left, middle, and right column

We thus assume $1 \leq g \leq p$ for the remainder of this proof. At depth $k - g$, the horizontal hyperplane intersects ρ in a convex polytope of dimension p , and it intersects the boundary of ρ on both sides of the silhouette. To go from one side to the other along the boundary of ρ intersected with P_{k-g} , we have to cross the intersection of P_{k-g} with the silhouette, which we will prove is a topological $(p - 2)$ -sphere. We conclude that an open $(p - 1)$ -ball of the boundary belongs to $P_{k-g} \cap |A|$, and the complementary closed $(p - 1)$ -ball does not belong to $P_{k-g} \cap |A|$. It follows that the slice of the interval has the topology type of \mathbb{H}^p , as claimed. The middle and right panels of Fig. 2 illustrate the four altruistic configurations for $p = 2$, and Fig. 4 illustrates the nine altruistic configurations for $p = 3$. Reading the eight outer cases in a circle around the center case, we note that each is symmetric to the diagonally opposite type. In other words, there are really only five altruistic types for $p = 3$.

We return to the intersection of P_{k-g} with the silhouette and the claim that this intersection is a topological sphere of dimension $p - 2$. For $p = 1$, ρ is a convex quadrangle, its silhouette consists of two vertices, $\text{top}(\rho)$ and $\text{btm}(\rho)$, and P_{k-1} passes through the other two vertices thus intersecting the silhouette

in the empty set—the (-1) -sphere—as claimed. Assuming $p \geq 2$, we denote the silhouette by \mathbb{S} , we recall that it is a $(p - 1)$ -sphere, and we write $e: \mathbb{S} \rightarrow \mathbb{R}$ for the depth function on the silhouette. Its extreme values are $e(\text{top}(\rho)) = k - (p + 1)$ and $e(\text{btm}(\rho)) = k$, and $P_{k-g} \cap \mathbb{S} = e^{-1}(k - g)$. To prove that this level set is a $(p - 2)$ -sphere, it suffices to show that e has only two critical points, namely the minimum at $\text{top}(\rho)$ and the maximum at $\text{btm}(\rho)$. The case $p = 2$ is easy. Here we have a 3-rhomboid whose silhouette is a hexagon. The difference between the depths of the endpoints of any edge is 1. We thus need three edges to go from $\text{btm}(\rho)$ at depth k to $\text{top}(\rho)$ at depth $k - 3$ and another three edges to go back. It follows that P_{k-g} meets the silhouette in two points—a 0-sphere—as claimed.

The argument for $p > 2$ is different. Recall that e is a continuous function on a $(p - 1)$ -sphere, this sphere is decomposed into $(p - 1)$ -rhomboids, and e is affine on each of these rhomboids. If e has a critical point (in piecewise-linear sense) in addition to the minimum at $\text{top}(\rho)$ and the maximum at $\text{btm}(\rho)$, then it also has a saddle, and this saddle must be a vertex of some of the rhomboids. To contradict the existence of a saddle, note that the $(p - 1)$ -rhomboids meet in groups of p at a common vertex. Let ν be such a shared vertex and cut each incident $(p - 1)$ -rhomboid with the $(p - 2)$ -dimensional plane that passes through the vertices adjacent to ν . We thus get p $(p - 1)$ -simplices, which can be seen are the facets of a p -simplex. It follows that ν can be a minimum, a maximum, or a regular point, but it cannot be a saddle of e . Hence, every horizontal slice at depth strictly between $k - (p + 1)$ and k is a $(p - 2)$ -sphere, as required. \square

Consequences Our main technical result is Theorem 3.1, which we now turn into a statement about the filtration of order- k Delaunay mosaics. Let $w_0 < w_1 < \dots$ be the sorted values of \mathbf{w}_k and write $K_\ell = \mathbf{w}_k^{-1}[-\infty, w_\ell] \subseteq \text{Del}_k(X)$ for every $\ell \geq 0$. Assuming X is in general position, the difference between any two contiguous mosaics is a collection of steps, and by slightly strengthening the notion of general position, we may assume that each difference is a *single* step: $A_\ell = K_\ell \setminus K_{\ell-1}$. For example, all vertices of the order-1 Delaunay mosaic of unweighted points share the function value, 0, and we can perturb the set by assigning small weights. While it is not necessary, we simplify the following statement by using this stronger notion of general position.

Corollary 3.2. (Filtration of Order- k Delaunay Mosaics) *Let X be a locally finite set of points with real weights in general position in \mathbb{R}^d , and let $0 \leq k$ and $0 \leq u \leq v$ be integers.*

1. *If exactly one of the steps A_u, A_{u+1}, \dots, A_v of \mathbf{w}_k is critical, then K_u and K_v have different Euler characteristics and therefore different homotopy types.*
2. *If A_u, A_{u+1}, \dots, A_v are all non-critical steps of \mathbf{w}_k , then K_u and K_v have the same homotopy type.*

This corollary of Theorem 3.1 is a direct extension of a theorem about discrete Morse functions in [10]. Other results in this theory can be similarly extended.

4. Discussion

The main result of this paper is a topological characterization of the incremental steps of the squared radius function on the order- k Delaunay mosaic of a locally finite set of possibly weighted points in Euclidean space. With this insight, we gain a topological interpretation of the probabilistic analysis of this function for a stationary Poisson point process [6]. While the critical steps do not determine the topology of the sublevel sets, they provide bounds on the ranks of their homology groups. In contrast to the order-1 case studied in [2], the squared radius function in the order- k case is neither discrete Morse nor generalized discrete Morse [10, 11]. Since the function nevertheless behaves similar to a Morse function, it may be considered a geometrically motivated further extension of the framework; see also [12, Chapter 11] for algebraically motivated extensions of discrete Morse theory.

In conclusion, we mention that our result requires the given points be in general position. While this assumption does not imply that the Delaunay mosaics are simplicial, it simplifies the analysis by guaranteeing that the dual of the corresponding hyperplane arrangement is a complex of rhomboids. It would be interesting to generalize the theory to locally finite sets that are not necessarily in general position.

Funding Open access funding provided by Institute of Science and Technology (IST Austria).

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access. This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

- [1] Aurenhammer, F.: *A new duality result concerning Voronoi diagrams*. Discrete Comput. Geom. **5**, 243–254 (1990)
- [2] Bauer, U., Edelsbrunner, H.: *The Morse theory of Čech and Delaunay complexes*. Trans. Am. Math. Soc. **369**, 3741–3762 (2017)
- [3] Bhatia, N., Vandana: *Survey of nearest neighbor techniques*. Int. J. Comput. Sci. Inf. Sec. **8**, 302–305 (2010)
- [4] Björner, A., Wachs, M.L.: *Shellable nonpure complexes and posets. I*. Trans. Am. Math. Soc. **348**, 1299–1327 (1996)
- [5] Edelsbrunner, H., Kirkpatrick, D.G., Seidel, R.: *On the shape of a set of points in the plane*. IEEE Trans. Inf. Theory **IT-29**, 551–559 (1983)
- [6] Edelsbrunner, H., Nikitenko, A.: *Poisson-Delaunay mosaics of order k* . Manuscript, IST Austria, Klosterneuburg, Austria (2017)
- [7] Edelsbrunner, H., Osang, G.: *The multi-cover persistence of Euclidean balls*. In: Proceedings of the 34th Annual Symposium on Computational Geometry, 2018, 34:1–34:14 (2018)
- [8] Edelsbrunner, H., Seidel, R.: *Voronoi diagrams and arrangements*. Discrete Comput. Geom. **1**, 25–44 (1986)
- [9] Fejes Toth, G.: *Multiple packing and covering of the plane with circles*. Acta Math. Acad. Sci. Hung. **27**, 135–140 (1976)
- [10] Forman, R.: *Morse theory for cell complexes*. Adv. Math. **134**, 90–145 (1998)
- [11] Freij, R.: *Equivariant discrete Morse theory*. Discrete Math. **309**, 3821–3829 (2009)
- [12] Koslov, D.: *Combinatorial Algebraic Topology*. Springer, Berlin (2008)
- [13] Krasnoshchekov, D., Polishchuk, V.: *Order- k alpha hulls and alpha shapes*. Inf. Process. Lett. **114**, 76–83 (2014)
- [14] Shamos, M.I., Hoey, D.J.: *Closest-point problems*. In: Proceedings of the 16th Annual IEEE Symposium on Foundations of Computer Science 1975, 151–162 (1975)
- [15] Ziegler, G.M.: *Lectures on Polytopes*. Springer, New York, NY (1995)

Herbert Edelsbrunner, Anton Nikitenko and Georg Osang
IST Austria (Institute of Science and Technology Austria)
Am Campus 1
3400 Klosterneuburg
Austria
e-mail: georg.osang@ist.ac.at

Herbert Edelsbrunner
e-mail: edels@ist.ac.at

Anton Nikitenko
e-mail: anton.nikitenko@ist.ac.at

Received: February 27, 2020.

Revised: February 18, 2021.

Accepted: February 20, 2021.