Continuous and Discrete Radius Functions on Tessellations and Mosaics

Ranita Biswas
IST Austria (Institute of Science and Technology Austria), Klosterneuburg, Austria
ranita.biswas@ist.ac.at

Sebastiano Cultrera di Montesano
IST Austria (Institute of Science and Technology Austria), Klosterneuburg, Austria
sebastiano.cultrera@ist.ac.at

Herbert Edelsbrunner
IST Austria (Institute of Science and Technology Austria), Klosterneuburg, Austria
herbert.edelsbrunner@ist.ac.at

Morteza Saghafian
Department of Mathematical Sciences, Sharif University of Technology, Tehran, Iran
morteza.saghafian65@student.sharif.edu

Abstract

The Voronoi tessellation in $\mathbb{R}^d$ is defined by locally minimizing the power distance to given weighted points. Symmetrically, the Delaunay mosaic can be defined by locally maximizing the negative power distance to other such points. We prove that the average of the two piecewise quadratic functions is piecewise linear, and that all three functions have the same critical points and values. Discretizing the two piecewise quadratic functions, we get the alpha shapes as sublevel sets of the discrete function on the Delaunay mosaic, and analog shapes as superlevel sets of the discrete function on the Voronoi tessellation. For the same non-critical value, the corresponding shapes are disjoint, separated by a narrow channel that contains no critical points but the entire level set of the piecewise linear function.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Voronoi tessellations, Delaunay mosaics, PL functions, alpha shapes, radius functions, continuous and discrete Morse theory.

Funding This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant no. 788183, from the Wittgenstein Prize, Austrian Science Fund (FWF), grant no. Z 342-N31, and from the DFG Collaborative Research Center TRR 109, 'Discretization in Geometry and Dynamics', Austrian Science Fund (FWF), grant no. I 02979-N35.

Lines 780

1 Introduction

The starting point for the work reported in this paper is the role of the general position assumption in the construction of Delaunay mosaics, and more specifically of their radius functions. Without general position assumption, the mosaics are not simplicial and the radius functions are not discrete Morse. How do we relax the theory to allow for non-generic data? Related to this question is the symmetry between Voronoi tessellations and Delaunay mosaics that appears when we introduce weights, and non-generic data is essential to realize this symmetry. In this paper, we weave the two strands of inquiry together by studying the continuous and discrete radius functions that define Voronoi tessellations and Delaunay mosaics for weighted points not necessarily in general position. We prove new results on these tessellations and mosaics by exploiting the structural properties of these functions.

Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The Voronoi tessellation and the dual Delaunay mosaic are classic topics in discrete geometry and go back at least to the seminal papers by Voronoi [19] and by Delaunay [3]. The radius function on the Delaunay mosaic was first introduced in [6], along with its sublevel sets, which are the alpha shapes of the given points. Three-dimensional alpha shapes have found ample applications in shape modeling [8, 11, 14] and in the analysis of biomolecules [7]. The connection to discrete Morse theory, as introduced by Forman [9] and generalized by Freij [10], was exploited for the purpose of surface reconstruction in [5]; see also [18]. We formulate the extension of discrete Morse theory needed to encompass radius functions on non-generic Delaunay mosaics and thus facilitate their application when non-generic position is essential, such as in crystallography.

Non-general position of points with weights is also essential when we interpret a Voronoi tessellation as a Delaunay mosaic and vice versa. By this we do not mean to take the tessellation to its dual mosaic but rather to construct a different set of weighted points whose Delaunay mosaic is essentially identical to the Voronoi tessellation of the first set. Viewing the tessellation and the mosaic as projections of the boundary complexes of convex polytopes, this construction follows by observing that the polar of a convex polyhedron is still a convex polyhedron. Notwithstanding, we get new insights into a much studied subject by looking into the details of this symmetry. We mention four such results, the first of which is combinatorial.

- Let $\mu \neq \nu$ be cells of a Voronoi tessellation, and write μ^*, ν^* for the corresponding cells in the dual Delaunay mosaic. Then $\operatorname{int} \mu \cap \nu^* \neq \emptyset$ implies $\operatorname{int} \nu \cap \mu^* = \emptyset$.

The second result is about the piecewise quadratic functions, $vor, del \colon \mathbb{R}^d \to \mathbb{R}$, whose pieces define the Voronoi tessellation and the dual Delaunay mosaic, respectively. Choosing opposite signs, the average defined by $sd(x) = \frac{1}{2}[vor(x) + del(x)]$ is piecewise linear. We use the above combinatorial insight to prove the following result.

- Extending concepts from smooth Morse theory to piecewise quadratic and piecewise linear functions, we show that $vor, del, sd \colon \mathbb{R}^d \to \mathbb{R}$ have the same critical points and the same critical values.

Discretizing the two piecewise quadratic functions, we get radius functions on the Voronoi tessellation and Delaunay mosaic, $\mathtt{vor} \colon \mathrm{Vor}(X) \to \mathbb{R}$ and $\mathtt{del} \colon \mathrm{Del}(X) \to \mathbb{R}$. For generic collections of weighted points, they are discrete Morse but not so for non-generic collections.

- Extending concepts from discrete Morse theory, we describe the structure of the steps of the radius functions on the Voronoi tessellation and Delaunay mosaic for weighted points in non-general position.

The fourth result sheds light on the relation between the sub- and superlevel sets of these discrete functions.

- We show that the underlying spaces of $\mathtt{del}^{-1}(-\infty, t]$ and $\mathtt{vor}^{-1}[t, \infty)$ are disjoint for all non-critical values t.

In particular, the channel between the two underlying spaces is free of critical points, the level set of the piecewise linear function, $sd^{-1}(t)$, splits it into two halves, and each half deformation retracts to the respective underlying space. Keeping track of the homology of the complementing subcomplexes, we get the basic relation of Alexander duality.

Outline. Section 2 presents background in discrete geometry. Section 3 studies the piecewise quadratic functions that define the Voronoi tessellation and Delaunay mosaic as well as their average, which is piecewise linear. Section 4 considers the corresponding discrete functions and introduces a framework to relate their properties to the standard axioms of discrete Morse theory. Section 5 relates the sublevel sets of one with the superlevel sets of the other. Section 6 concludes the paper.

2 Background

We review Voronoi tessellations and the dual Delaunay mosaics, which we introduce for points with real weights in Euclidean space. In addition, we describe the standard polarity transform and its relation to the tessellation and the mosaic. Finally, we explain how to view tessellations and mosaics as projections of convex polytopes.

2.1 Voronoi Tessellations and Delaunay Mosaics

We refer to $a = (p_a, w_a) \in \mathbb{R}^d \times \mathbb{R}$ as a *weighted point*, with *location* $p_a \in \mathbb{R}^d$ and *weight* $w_a \in \mathbb{R}$. Let $B \subseteq \mathbb{R}^d \times \mathbb{R}$ be a set of weighted points whose projection to $\mathbb{R}^d$ is injective and locally finite. In other words, for every location there is an open neighborhood that separates it from the other locations. It is common to interpret $a = (p_a, w_a)$ as a sphere, with center p_a and squared radius w_a, but for this we have to allow for spheres with non-positive squared radii. The *power distance* of a point $x \in \mathbb{R}^d$ from $a = (p_a, w_a)$ is $\pi_a(x) = \|x - p_a\|^2 - w_a$. It is positive outside the sphere, zero on the sphere, and negative inside the sphere. Of course, for a sphere with negative squared radius, all points are outside. For a subset $A \subseteq B$, consider all points $x \in \mathbb{R}^d$ with equal power distance from the weighted points in A and strictly larger power distance from the other weighted points, and call its closure the *(Voronoi) cell* of A, denoted $\mathrm{cell}(A)$. Each non-empty cell is a convex polyhedron in $\mathbb{R}^d$, and its dimension depends on A. The *(weighted) Voronoi tessellation* of B, denoted $\mathrm{Vor}(B)$, is the collection of non-empty cells. It is a *polyhedral complex* in the sense that every cell is a convex polyhedron, every face of a cell is again a cell, and any two cells are either disjoint or intersect in a common face, which is therefore also a cell in the tessellation. A cell of dimension p has faces of dimension from 0 to p, and we call the faces of dimension $p-1$ its *facets*. Define the *dual cell* of A as the convex hull of the locations in A, denoted $\mathrm{cell}^*(A)$, which is again a convex polyhedron. The dimension of a cell and its dual cell are necessarily complementary: if $p = \dim \mathrm{cell}(A)$ and $q = \dim \mathrm{cell}^*(A)$, then $p + q = d$. The *(weighted) Delaunay mosaic* of B, denoted $\mathrm{Del}(B)$, is the collection of dual cells. Figure 1 illustrates the concepts by drawing a Voronoi tessellation and the corresponding Delaunay mosaic on top of each other.

In $\mathbb{R}^d$, we call a Voronoi tessellation *simple* if every p-dimensional cell is face of exactly $q + 1 = d - p + 1$ top-dimensional cells, and we call a Delaunay mosaic *simplicial* if every q-dimensional dual cell is the convex hull of $q + 1$ points. Clearly, a Voronoi tessellation is simple iff the corresponding Delaunay mosaic is simplicial. We stress that this paper does not assume that $\mathrm{Vor}(B)$ be simple and $\mathrm{Del}(B)$ be simplicial, and we introduce these notions primarily to clarify the difference between the generic and the non-generic situation.

Besides $\mathrm{Vor}(B)$ and $\mathrm{Del}(B)$, we will be interested in subcomplexes and subsets of these complexes. To stress the difference, we note that a *subcomplex* is closed under taking faces, while a *subset* does not necessarily enjoy this property. We call a subset *open* if it is closed under taking cofaces. As an example consider a subset $K \subseteq \mathrm{Vor}(B)$ and let $K^* \subseteq \mathrm{Del}(B)$ contain $\mathrm{cell}^*(A)$ iff $\mathrm{cell}(A) \in K$. Clearly, K is a subcomplex of the Voronoi tessellation iff K^* is an open subset of $\mathrm{Del}(B)$, and vice versa. While the cells in a complex may intersect, their (relative) interiors are disjoint. Indeed, for every $x \in \mathbb{R}^d$ there is a unique cell $\tau \in \mathrm{Vor}(B)$ whose interior contains x. The same is true for the Delaunay mosaic if we restrict ourselves to points x in the convex hull of the locations. As suggested in Figure 1, we will extend the

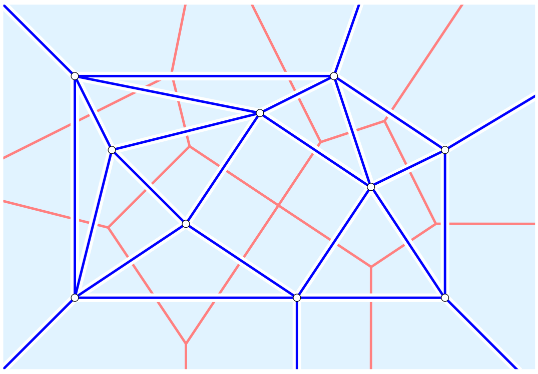

Figure 1: The overlay of a Voronoi tessellation and its dual Delaunay mosaic. The former is not simple because it contains one vertex incident to four edges, and the latter is not simplicial because it contains one region with four edges. We add half-lines to the mosaic to decompose the complement of the convex hull into convex cells.

Delaunay mosaic artificially so that this restriction can be removed. We define the *underlying space* of a subset K of a polyhedral complex as the union of interiors of its cells:

$$|K| = \{x \in \mathbb{R}^d \mid x \in \operatorname{int} \tau \text{ for some } \tau \in K\}. \tag{1}$$

If K is a complex, then this is just the union of cells, but if K is not a complex, then the union of interiors is a strict subset of the union of cells.

2.2 Polarity

We introduce the *paraboloid map*, $\varpi \colon \mathbb{R}^d \to \mathbb{R}$, defined by $\varpi(x) = \frac{1}{2}\|x\|^2$ and we are interested in the most elementary version of polarity with respect to this paraboloid, which relates a point $u = (u_1, u_2, \ldots, u_{d+1})$ in $\mathbb{R}^{d+1}$ with the hyperplane of points $x \in \mathbb{R}^{d+1}$ that satisfy $x_{d+1} = u_1 x_1 + \ldots + u_d x_d - u_{d+1}$. We denote this hyperplane by u^*, we call u^* the *polar hyperplane* of u (with respect to ϖ), and we call $u = (u^*)^*$ the *polar point* of u^* (with respect to ϖ). Importantly, the transform preserves incidences, that is: $u \in v^*$ iff $v \in u^*$ for any two points $u, v \in \mathbb{R}^{d+1}$. The transform also preserves sidedness, which we introduce by saying that u lies *below*, *on*, *above* v^* if u_{d+1} is less than, equal to, greater than $v_1 u_1 + \ldots + v_d u_d - v_{d+1}$. Specifically, u is above v^* iff v is above u^*, and together with the preservation of incidences, this implies u is below v^* iff v is below u^*.

To express the relation between the Voronoi tessellation and the Delaunay mosaic in terms of the polarity transform, we map every weighted point in $\mathbb{R}^d \times \mathbb{R}$ to a lifted point and its polar hyperplane in $\mathbb{R}^{d+1}$. For every weighted point $a = (p_a, w_a)$, we represent the two by a constant map and an affine map, $f_a, g_a \colon \mathbb{R}^d \to \mathbb{R}$:

$$f_a(x) = \tfrac{1}{2}\|p_a\|^2 - \tfrac{1}{2} w_a, \tag{2}$$
$$g_a(x) = \langle p_a, x \rangle - f_a(x), \tag{3}$$

so that $(p_a, f_a(p_a))$ is the lifted point and $\operatorname{img} g_a = g_a(\mathbb{R}^d)$ is its polar hyperplane (with respect to ϖ). It is not difficult to verify that the average of the two maps on p_a gives us the value of ϖ on p_a:

$$\tfrac{1}{2}[f_a(p_a) + g_a(p_a)] = \tfrac{1}{2}\|p_a\|^2 = \varpi(p_a). \tag{4}$$

Returning to the connection with the weighted points, the zero-set of $g_a - \varpi$ consists of the points $x \in \mathbb{R}^d$ for which

$$g_a(x) - \varpi(x) = -\tfrac{1}{2}\|x - p_a\|^2 + \tfrac{1}{2}w_a = -\tfrac{1}{2}\pi_a(x) \tag{5}$$

vanishes. In words, the zero-set of $g_a - \varpi$ is also the zero-set of π_a, namely the sphere with center p_a and squared radius w_a. We call two weighted points $a = (p_a, w_a)$ and $b = (p_b, w_b)$ *orthogonal* if $\|p_a - p_b\|^2 = w_a + w_b$. It is a straightforward exercise to show that this is equivalent to $g_a(p_b) = f_b(p_b)$ or, in words, that the lifted point of b lies on the hyperplane of a. If both weights are positive, Pythagoras' theorem implies that the zero-sets of π_a and π_b — which are spheres with squared radii w_a and w_b — intersect orthogonally.

Next, we generalize the relations between points and hyperplanes to collections $A \subseteq \mathbb{R}^d \times \mathbb{R}$ whose projection to $\mathbb{R}^d$ is injective and locally finite. Write $\mathrm{flat}(A)$ for the affine hull of the locations: $\mathrm{flat}(A) = \mathrm{aff}\,\{p_a | a \in A\}$, and $\mathrm{sol}(A)$ for the set of points $x \in \mathbb{R}^d$ that satisfy $g_a(x) = g_b(x)$ for all $a, b \in A$. For example, if $A = \{a = (p_a, w_a)\}$, then $\mathrm{flat}(A) = p_a$ and $\mathrm{sol}(A) = \mathbb{R}^d$. Assuming the locations of the points in A are affinely independent, we write $q + 1 = \#A$ and $p = d - q$, and observe that

- $\dim \mathrm{flat}(A) = q$ and $\dim \mathrm{sol}(A) = p$,
- $\mathrm{flat}(A)$ and $\mathrm{sol}(A)$ are orthogonal affine subspaces of $\mathbb{R}^d$, and we write $y = y(A)$ for the intersection point.

Indeed, if all weights are zero, then $\mathrm{sol}(A)$ is the set of centers of spheres that pass through all points of A. This set is a p-dimensional affine subspace of $\mathbb{R}^d$ orthogonal to the q-dimensional affine hull of A. When we adjust the weight of $a \in A$, this affine subspace does not change other than by moving parallel to its initial position. So $\mathrm{flat}(A)$ and $\mathrm{sol}(A)$ retain the two properties stated above.

In addition to the two affine subspaces, we introduce two affine functions, $f_A \colon \mathbb{R}^d \to \mathbb{R}$ and $g_A \colon \mathbb{R}^d \to \mathbb{R}$, that generalize f_a and g_a as defined in (2) and (3). Specifically, f_A agrees with f_a at p_a for every $a = (p_a, w_a) \in A$ and its restriction to $\mathrm{sol}(A)$ is constant. Similarly, g_A agrees with g_a within $\mathrm{sol}(A)$ for every $a \in A$ and its restriction to $\mathrm{flat}(A)$ is constant. Recall that $y(A) = \mathrm{sol}(A) \cap \mathrm{flat}(A)$.

▶ **Lemma 2.1** (Common Maximum). *Let $A \subseteq \mathbb{R}^d \times \mathbb{R}$ be a set of weighted points whose locations are affinely independent. Then $y = y(A)$ is the common maximum of*

- *(i) the restriction of $f_A - \varpi$ to $\mathrm{flat}(A)$,*
- *(ii) the restriction of $g_A - \varpi$ to $\mathrm{sol}(A)$,*
- *(iii) the average, $\frac{1}{2}[f_A + g_A] - \varpi$, and in this case the value of the maximum vanishes.*

Proof. We begin by mapping every location $x \in \mathrm{flat}(A)$ to a weighted point $u \in \mathbb{R}^d \times \mathbb{R}$ with $p_u = x$ and $w_u = 2\varpi(x) - 2f_A(x)$, noting that $f_u(x) = f_A(x)$. Similarly, we map every location $x \in \mathrm{sol}(A)$ to $v \in \mathbb{R}^d \times \mathbb{R}$ with $p_v = x$ and $w_v = 2\varpi(x) - 2g_A(x)$, noting that $f_v(x) = g_A(x)$. By construction, $g_u \colon \mathbb{R}^d \to \mathbb{R}$ agrees with g_A on $\mathrm{sol}(A)$ and, symmetrically, $g_v \colon \mathbb{R}^d \to \mathbb{R}$ agrees with f_A on $\mathrm{flat}(A)$. Hence, $\|p_u - p_v\|^2 = w_u + w_v$, which for positive weights is equivalent to the zero-sets of π_u and π_v intersecting orthogonally. Observe that this is true for all pairs $(p_u, p_v) \in \mathrm{flat}(A) \times \mathrm{sol}(A)$, so we have what for two lines in $\mathbb{R}^2$ is sometimes called a coaxal system [17].

If we now fix v with $p_v \in \mathrm{sol}(A)$, we get u with minimum weight by minimizing $\|p_v - p_u\|^2$. This minimum is attained for $p_u = y$, and since $w_u = 2\varpi(p_u) - 2f_A(p_u)$, this implies that y maximizes $f_A - \varpi$, as claimed in (i). The proof of (ii) is symmetric.

While we considered only the restrictions of f_A and g_A to affine subspaces, they are defined on the entire $\mathbb{R}^d$. Hence, the map $f\colon \mathbb{R}^d \to \mathbb{R}$ sending x to $f(x) = \frac{1}{2}[f_A(x) + g_A(x)]$ is well defined. It is affine since f_A and g_A are affine. Letting x' and x'' be the orthogonal projections of $x \in \mathbb{R}^d$ onto flat(A) and sol(A), respectively, we have $f(x) = \frac{1}{2}[f_A(x')+g_A(x'')]$. At the intersection of the two affine subspaces, we have $f(y) - \varpi(y) = 0$ by (4). At every other point $x \in \mathbb{R}^d$, $f(x) - \varpi(x) < 0$, simply because $f_A(x') - \varpi(x') \leq f_A(y) - \varpi(y)$ and $g_A(x'') - \varpi(x'') \leq g_A(y) - \varpi(y)$, with strict inequality at least once. This implies (iii). ◻

We note that (iii) implies that the graph of $\frac{1}{2}[f_A + g_A]$ is the unique hyperplane in $\mathbb{R}^{d+1}$ that touches the graph of ϖ in the point $(y, \varpi(y))$.

2.3 Projection of Envelopes

Since the Voronoi tessellation is defined in terms of minimum power distance, it can equally well be defined in terms of maximum affine function values. Specifically, let $env\colon \mathbb{R}^d \to \mathbb{R}$ be the upper envelope of the affine maps: $env(x) = \max_{a \in B} g_a(x)$, and call the linear pieces of this envelope the *faces* of env. It is not difficult to see that there is a bijection between the faces of env and the cells of $\mathrm{Vor}(B)$ such that every cell is the vertical projection of the corresponding face to $\mathbb{R}^d$. This property was known already to Voronoi [19].

A similar construction exists for Delaunay mosaics, which is usually phrased in terms of the convex hull of the points $(p_a, f_a(p_a))$ in $\mathbb{R}^{d+1}$. We call a face of this convex polytope *lower* if there is a non-vertical hyperplane in $\mathbb{R}^{d+1}$ such that the face lies in the hyperplane and the rest of the polytope lies above it. It is not difficult to see that there is a bijection between the lower faces of this polytope and the cells of $\mathrm{Del}(B)$ such that every cell is the vertical projection of the corresponding lower face to $\mathbb{R}^d$. In this paper, it is convenient to add arbitrarily steep "ramps" around the polytope whose vertical projections decompose the rest of $\mathbb{R}^d$ into convex cells. In other words, we introduce $end\colon \mathbb{R}^d \to \mathbb{R}$ as the upper envelope of all affine maps $g_c\colon \mathbb{R}^d \to \mathbb{R}$ that satisfy $g_c(x) \leq y$ for every point $(x, y) \in \mathbb{R}^d \times \mathbb{R}$ of the polytope. Most of these maps are redundant, except those whose graphs support facets, and the ramps that support $(d-1)$-dimensional faces on the silhouette of the polytope. Then there is a set of weighted points, $C \subseteq \mathbb{R}^d \times \mathbb{R}$, possibly including a point at infinity, whose projection to $\mathbb{R}^d$ is locally finite such that $end(x) = \max_{c \in C} g_c(x)$. Now we have complete symmetry and can write $\mathrm{Del}(B) = \mathrm{Vor}(C)$ as well as $\mathrm{Vor}(B) = \mathrm{Del}(C)$. We call C the *polar set* of B and, symmetrically, B the polar set of C.

3 Continuous Functions

In this section, we consider two piecewise quadratic functions, whose pieces define the Voronoi tessellation and its dual Delaunay mosaic. The main result is that these two functions and their piecewise linear average have the same critical points.

3.1 Piecewise Quadratic and Piecewise Linear Functions

Recall that $env, end\colon \mathbb{R}^d \to \mathbb{R}$ are piecewise linear convex functions. Comparing them with ϖ, we get two piecewise quadratic functions, $vor, del\colon \mathbb{R}^d \to \mathbb{R}$, and one piecewise linear function, $sd\colon \mathbb{R}^d \to \mathbb{R}$, defined by

$$vor(x) = \varpi(x) - env(x), \tag{6}$$
$$del(x) = end(x) - \varpi(x). \tag{7}$$
$$sd(x) = \tfrac{1}{2}[end(x) - env(x)] = \tfrac{1}{2}[del(x) + vor(x)]. \tag{8}$$

As illustrated in Figure 2, *del* dominates *vor*, which implies that their average, *sd*, is sandwiched between them. To prove this formally, we introduce the *common subdivision* of

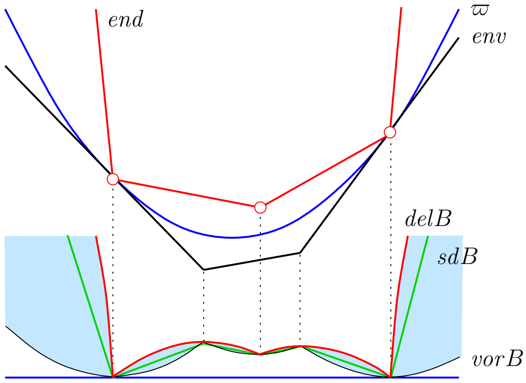

Figure 2: The paraboloid function, the two envelope functions, and their piecewise quadratic and piecewise linear differences.

the tessellation and the mosaic, denoted $\mathrm{Sd}(B)$, which consists of all cells $\gamma = \tau \cap \sigma^*$ with $\tau \in \mathrm{Vor}(B)$ and $\sigma^* \in \mathrm{Del}(B)$. Since τ and σ^* are convex, so is γ. The restrictions of *del* and of *vor* to γ are quadratic, while the restriction of *sd* to γ is linear.

▶ **Lemma 3.1** (Sandwich). *Let $B \subseteq \mathbb{R}^d \times \mathbb{R}$ have an injective and locally finite projection to $\mathbb{R}^d$. Then $del(x) \geq sd(x) \geq vor(x)$ for every $x \in \mathbb{R}^d$.*

Proof. Let $a \in \mathbb{R}^d \times \mathbb{R}$ such that $f_a(p_a) = env(p_a)$. Hence, $f_a(p_a) \geq g_b(p_a)$ for all $b \in B$, with equality at least once. Since the polarity transform preserve sidedness, we have $f_b(p_b) \geq g_a(p_b)$, for all $b \in B$, and therefore $end(y) \geq g_a(y)$ for all $y \in \mathbb{R}^d$, which includes $y = p_a$. Writing $x = p_a$, this implies

$$del(x) - vor(x) = end(x) + env(x) - 2\varpi(x) \geq g_a(x) + f_a(x) - 2\varpi(x), \tag{9}$$

in which the right-hand side vanishes because of (4). This implies the claimed inequalities. ◻

The inequalities in Lemma 3.1 imply that the sublevel sets and the superlevel sets of the three functions are nested:

$$del^{-1}(-\infty, t] \subseteq sd^{-1}(-\infty, t] \subseteq vor^{-1}(-\infty, t], \tag{10}$$
$$del^{-1}[t, \infty) \supseteq sd^{-1}[t, \infty) \quad \supseteq vor^{-1}[t, \infty). \tag{11}$$

The sublevel set of *del* and the superlevel set of *vor*, for a common value t, are illustrated in Figure 3 together with the channel between these two sets. We will see shortly that the three functions share the critical points, at which they all agree.

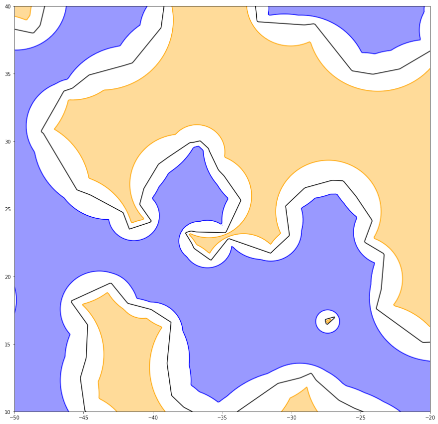

Figure 3: The *black* level set of *sd* splits the *white* channel into two. The corresponding superlevel set of *vor* is *orange* and the sublevel set of *del* is *blue*.

3.2 Two Auxiliary Lemmas

We need three auxiliary results to prove that the functions defined in (6), (7), (8) share the critical points and values, two of which will be presented in this subsection. The first result is a new combinatorial statement about Voronoi tessellations and Delaunay mosaics.

▶ **Lemma 3.2** (Excluded Crossing). *Let $B \subseteq \mathbb{R}^d \times \mathbb{R}$ have an injective and locally finite projection to $\mathbb{R}^d$, let $\mu \neq \nu$ be cells in* $\mathrm{Vor}(B)$ *and recall that μ^*, ν^* are their dual cells in* $\mathrm{Del}(B)$. *If $\mathrm{int}\,\mu \cap \nu^* \neq \emptyset$, then $\mathrm{int}\,\nu \cap \mu^* = \emptyset$.*

Proof. To reach a contradiction, assume that both intersections are non-empty, so we can choose points $x \in \mathrm{int}\,\mu \cap \nu^*$ and $y \in \mathrm{int}\,\nu \cap \mu^*$. Since the interiors of μ and ν are disjoint, we have $x \neq y$. Let $M, N \subseteq B$ be such that $\mu = \mathrm{cell}(M)$ and $\nu = \mathrm{cell}(N)$. By definition of a cell, x has the same power distance from all $a \in M$, and a strictly larger power distance from all $b \in B \setminus M$. Write $R_M = \pi_a(x)$ with $a \in M$, and write $R_N = \pi_c(y)$ with $c \in N$. Assume without loss of generality that $R_N \geq R_M$. Then every weighted point $a \in M$ satisfies $\pi_a(y) \geq R_N \geq R_M = \pi_a(x)$, so $\|y - p_a\| \geq \|x - p_a\|$. Drawing the perpendicular bisector of x and y, this implies that all p_a with $a \in M$ lie in the closed half-space that contains x. Since y lies outside this half-space, it is not contained in the convex hull of the p_a with $a \in M$, but this contradicts $y \in \mu^*$. ◻

We remark that we take the interiors of μ and ν so that the two hypothesized intersection points are different. This detail is a crucial aspect of the proof. Indeed, it is possible to have $\mu \cap \nu^* \neq \emptyset$ and $\nu \cap \mu^* \neq \emptyset$: let ν^* be a right-angled triangle in $\mathbb{R}^2$ and μ^* its longest edge. Then ν is the circumcenter of the triangle, which lies on μ^*, and μ has ν as an endpoint.

Write $\mathbb{S}^{d-1}$ for the unit sphere in $\mathbb{R}^d$. The second result is a geometric statement about the common intersection of *hemispheres*, which are closed subsets of $\mathbb{S}^{d-1}$ that are bounded by great-spheres of dimension $d-2$. Note that a unit vector, $e \in \mathbb{S}^{d-1}$, defines both a point as well as a hemisphere, namely the one whose points $y \in \mathbb{S}^{d-1}$ satisfy $\langle e, y \rangle \leq 0$.

▶ **Lemma 3.3** (Hemispheres). *The common intersection of a collection of hemispheres of $\mathbb{S}^{d-1}$ is either contractible or a $(p-1)$-dimensional great-sphere with $0 \leq p \leq d$.*

Proof. Let $E \subseteq \mathbb{S}^{d-1}$ be the set of vectors defining the hemispheres in the given collection. If $E \neq \emptyset$ and there is a point $x \in \mathbb{S}^{d-1}$ with $\langle e, x \rangle < 0$ for all $e \in E$, then the hemispheres have a non-empty and contractible common intersection. Otherwise, let $x \in \mathbb{S}^{d-1}$ such that $\langle e, x \rangle \leq 0$, for all $e \in E$, with equality for a minimum number of vectors. If x does not exist, then the intersection of hemispheres is empty, which is the case $p = 0$ in the claimed statement. When x exists, it may not be unique, but the vectors e for which the scalar product vanishes are unique. Similarly, the linear span of these vectors is unique, and letting $0 \leq d - p \leq d$ be its dimension, the common intersection of the hemispheres is a $(p-1)$-dimensional great-sphere. The case $p = d$ corresponds to an empty collection of hemispheres so that the common intersection is the entire $\mathbb{S}^{d-1}$. ◻

3.3 In- and Out-Links

The third result is a topological statement about vector fields defined by two convex polytopes, $P, Q \subseteq \mathbb{R}^d$, whose dimensions are complementary, $p = \dim P$ and $q = \dim Q$ with $p + q = d$, and whose affine hulls intersect in a single point. The product, $P \times Q$, is a convex polytope of dimension d. Its boundary is a topological $(d-1)$-sphere that decomposes into a thickened $(p-1)$-sphere and a thickened $(q-1)$-sphere: $\partial(P \times Q) = (\partial P \times Q) \cup (P \times \partial Q)$. Indeed, for every $s \in \partial(P \times Q)$, there are unique points $y \in P$ and $z \in Q$ such that $s = y + z$, and at least one of y and z belongs to the respective boundary. We are interested in $\psi \colon \partial(P \times Q) \to \mathbb{S}^{d-1}$ defined by mapping $s = y + z$ to $\psi(s) = \frac{1}{2}(y - z)$; see Figure 4 for an illustration. To study

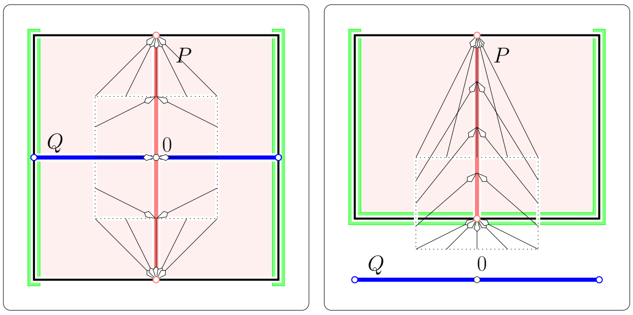

Figure 4: The map $\psi \colon \partial(P \times Q) \to \mathbb{S}^1$ illustrated for two intersecting line segments on the *left* and for two disjoint line segments on the *right*. For better visualization, we anchor the vectors at the boundary points of $\frac{1}{2}(P \times Q)$, and we highlight the in-links in green.

ψ, we introduce the *in-link* and *out-link* of P and Q:

$$
\begin{aligned}
\mathit{in}\mathrm{Lk}(P, Q) &= \{s \in \partial(P \times Q) \mid \langle \psi(s), \mathbf{n}(s) \rangle \leq 0\}, && (12)\\
\mathit{out}\mathrm{Lk}(P, Q) &= \{s \in \partial(P \times Q) \mid \langle \psi(s), \mathbf{n}(s) \rangle \geq 0\}, && (13)
\end{aligned}
$$

in which $\mathbf{n}(s)$ is the unit outward directed normal at s. This normal is unique for every facet, which we recall is a face of dimension $d-1$, but it is not unique for faces of dimension $d-2$ or less. We remedy this difficulty by writing $\mathbf{n}(s)$ for the collection of normals that interpolate between the normals of the incident facets, and by including s in the in- or out-link if the respective inequality is satisfied for at least one vector in $\mathbf{n}(s)$. In the left panel of Figure 4, the in-link consists of the left edge and the right edge of the product, while the out-link consists of the remaining two edges. Both have the homotopy type of the 0-sphere. In the right panel, the in-link consists of three edges, with the out-link containing the remaining, top edge. Both links are contractible. The important difference is that P and Q intersect in the left panel while they are disjoint in the right panel.

▶ **Lemma 3.4** (In- and Out-Link). *Let $P, Q \subseteq \mathbb{R}^d$ be convex polytopes with orthogonal affine hulls of complementary dimensions: $p = \dim P$, $q = \dim Q$, and $p + q = d$. Then*

$$\operatorname{int} P \cap \operatorname{int} Q \neq \emptyset \implies in\mathrm{Lk}(P,Q) \simeq \mathbb{S}^{q-1}, out\mathrm{Lk}(P,Q) \simeq \mathbb{S}^{p-1}, \tag{14}$$

$$P \cap Q = \emptyset \implies in\mathrm{Lk}(P,Q) \text{ and } out\mathrm{Lk}(P,Q) \text{ contractible}, \tag{15}$$

$$\operatorname{int} P \cap \operatorname{int} Q = \emptyset \text{ and } P \cap Q \neq \emptyset \implies in\mathrm{Lk}(P,Q) \text{ or } out\mathrm{Lk}(P,Q) \text{ contractible}. \tag{16}$$

Proof. Assume that the affine hulls of P and Q intersect at $0 \in \mathbb{R}^d$. Every facet E of $R = P \times Q$ is either of the form $F \times Q$ or $P \times G$, in which F and G are facets of P and Q, respectively. Whether or not E belongs to the in-link or the out-link depends on the relative position of E and 0, and the rule is opposite for the two forms. To explain, we call E *visible* (from 0) if $\langle \mathbf{n}(s), s \rangle \leq 0$ for every $s \in E$ and *invisible* (from 0) if $\langle \mathbf{n}(s), s \rangle \geq 0$ for every $s \in E$. We observe that $in\mathrm{Lk}(P,Q)$ contains all visible facets E of the form $E = F \times Q$ and all invisible facets of the form $E = P \times G$, while $out\mathrm{Lk}(P,Q)$ contains all invisible facets of the first type and all visible facets of the second type.

In the first case, when $\operatorname{int} P \cap \operatorname{int} Q \neq \emptyset$, 0 belongs to the interior of R. Hence all facets of R are invisible, which implies that the in-link is $P \times \partial Q$, which has the homotopy type of a $(q-1)$-sphere. Symmetrically, the out-link is $\partial P \times Q$, which has the homotopy type of the $(p-1)$-sphere. This proves (14).

To prepare the second case, consider a q-dimensional convex polytope Q in $\mathbb{R}^q$, and let $0 \in \mathbb{R}^q$ be outside Q and not contained in the affine hull of any of its facets. This partitions the facets into the visible and invisible ones from 0. Letting H be a hyperplane that separates 0 from Q, we can apply a projective transformation that maps H to infinity, 0 to another point $0'$, and Q to another convex polytope Q', all in $\mathbb{R}^q$. We may imagine this transform moves H to infinity, pushing 0 in front of it to disappear to infinity and then return from the other side. Importantly, a facet of Q is visible from 0 iff the corresponding facet of Q' is invisible from $0'$. We will make use of this construction shortly.

In the second case, when $P \cap Q = \emptyset$, not all facets of R are invisible. Since $0 \notin R$, it is outside at least one of P and Q, and we assume without loss of generality $0 \notin Q$. To distinguish the two types of facets of R, we consider P and Q within their respective affine hulls. Specifically, there is a bijection between the visible facets of R on the one side, and the visible facets of P inside $\operatorname{aff} P$ and of Q inside $\operatorname{aff} Q$ on the other side. For the in-link, we need the visible facets of P and the invisible facets of Q, so we apply a projective transformation that maps Q to Q' and 0 to $0'$ — all still in $\operatorname{aff} Q$ — such that a facet of Q is invisible from 0 iff the corresponding facet of Q' is visible from $0'$. This transformation does not affect P. We get a new product, $R' = P \times Q'$ and we are interested in the part of the boundary that is visible from $0'$. Since R' is convex and $0' \notin R'$, this part of $\partial R'$ is contractible, which implies that the corresponding part of ∂R, which is $in\mathrm{Lk}(P,Q)$, is also contractible. Symmetrically, the invisible part of $\partial R'$ is contractible, which implies that $out\mathrm{Lk}(P,Q)$ is also contractible. This proves (15).

In the third case, when $\operatorname{int} P \cap \operatorname{int} Q = \emptyset$ and $P \cap Q \neq \emptyset$, 0 belongs to ∂R. The facets that contain 0 are both visible and invisible (from 0). Assume $0 \in \partial Q$. Then we can move 0 to $0'$, still within $\operatorname{aff} Q$ but slightly outside Q, in such a way that a facet of Q is visible from 0 iff it is visible from $0'$. Now we are in the second case as far as the visible facets of Q are concerned, which implies that the out-link of P and Q is contractible. This proves (16). Note that this construction is not symmetric, as moving 0 to $0''$ inside Q preserves the invisible facets of Q but does not imply a contractible in-link. However, we need only one contractible link, which completes the proof. ◻

3.4 Up- and Down-Links

Since the continuous functions we study are not smooth, it is necessary to define what we mean by a critical point. We need a definition that is general enough to apply to piecewise linear and to piecewise quadratic functions. Letting $f \colon \mathbb{R}^d \to \mathbb{R}$ be such a function and $x \in \mathbb{R}^d$, we write $S_r = S_r(x)$ for the $(d-1)$-sphere with radius $r > 0$ and center x. Letting S_r^- contain all $y \in S_r$ with $f(y) \leq 0$, we note that its homotopy type is the same for all sufficiently small radii. Fixing a sufficiently small $\varepsilon > 0$, we call S_ε^- the *down-link* of x and f, denoted $dn\mathrm{Lk}(x, f)$. Symmetrically, S_r^+ contains all points $y \in S_r$ with $f(y) \geq 0$, and we call S_ε^+ the *up-link* of x and f, denoted $up\mathrm{Lk}(x, f)$. We call x a *non-critical point* of f if at least one of the two links is contractible. All points with topologically more complicated up- and down-links are *critical points* of f, where we note that the empty link is not contractible. See Figure 5 for the local pictures that arise for a 2-dimensional piecewise linear function. In the generic case, the down-link is contractible iff the up-link is contractible. The "at least one" rule is used to classify borderline cases as non-critical. An example is the southern hemisphere as the down-link and the northern hemisphere together with the south-pole as the up-link.

To study the critical points of $f = vor$, we fix $x \in \mathbb{R}^d$ and let $A \subseteq B$ be the subset of weighted points such that $x \in \operatorname{int} \operatorname{cell}(A)$. Setting $h^2 = vor(x)$, x lies on the boundary of $vor^{-1}(-\infty, h^2]$, which is a union of closed balls, namely the balls with centers p_a and squared radii $w_a + h^2$, for $a \in B$. Specifically, x lies on the boundary of such a ball if $a \in A$, and it lies outside the ball if $a \in B \setminus A$. We get the two links by intersecting the union and its closed complement with a sphere of sufficiently small radius ε:

$$dn\mathrm{Lk}(x, vor) = S_\varepsilon(x) \cap vor^{-1}(-\infty, h^2], \tag{17}$$

$$up\mathrm{Lk}(x, vor) = S_\varepsilon(x) \cap vor^{-1}[h^2, \infty). \tag{18}$$

Scaling the small sphere back to unit size, we get a closed cap that approximates the complement of a hemisphere arbitrarily closely for each $a \in A$, and the down-link as the union of these caps. By Lemma 3.3, there are only $d + 2$ possible shapes for $dn\mathrm{Lk}(x, vor)$, namely either contractible or a thickened $(q - 1)$-dimensional great sphere for $0 \leq q \leq d$. Symmetrically, there are only $d+2$ possible shapes for $up\mathrm{Lk}(x, vor)$, namely either contractible or a thickened $(p - 1)$-dimensional great sphere with $p = d - q$. If at least one of the two links is contractible, then x is a non-critical point of vor, and otherwise, it is a critical point with index q. The symmetric argument applies to del, so x can be either a non-critical point of del or a critical point with the same index, q.

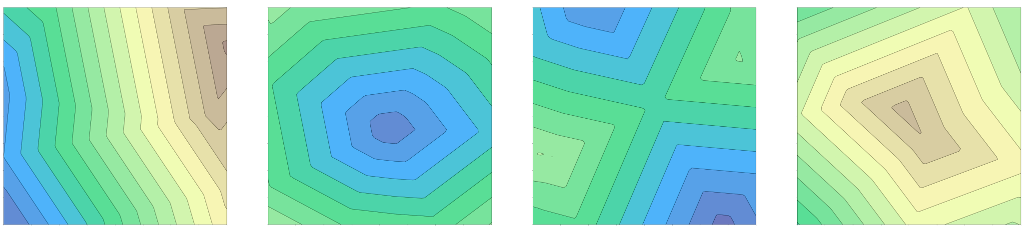

Figure 5: From *left* to *right*: typical patterns of level sets in the neighborhood of a non-critical point, a minimum (index 0), a saddle (index 1), and a maximum (index 2) in two dimensions. The corresponding down-link is a single contractible arc, empty, two disjoint contractible arcs, and the full circle, respectively. The patterns are cut out of the larger context in Figure 9(d), where the middle level set is shown using thin black lines.

3.5 Coincidental Critical Points

Recall that $del(x) \geq sd(x) \geq vor(x)$ by Lemma 3.1. We strengthen this result by proving further connections between the three functions. Specifically, we prove that every point $x \in \mathbb{R}^d$ is of the same type for vor and for del, as well as for their average. Recall that the restriction of the latter to a d-dimensional cell $\gamma = \tau \cap \sigma^*$ satisfies

$$sd(x) = \tfrac{1}{2}[del(x) + vor(x)] = \tfrac{1}{2}[-\tfrac{1}{2}\pi_c(x) + \tfrac{1}{2}\pi_b(x)] = \tfrac{1}{2}\langle x, p_c - p_b \rangle + \text{const}, \tag{19}$$

in which $b \in B$ and $c \in C$ such that $\tau = \text{cell}(b)$ and $\sigma^* = \text{cell}(c)$. Hence, $p_c - p_b$ is twice the gradient of sd at every point in $\text{int}\,\gamma$. We use this insight to prove the main result of this section.

▶ **Theorem 3.5** (Coincidental Critical Points). *Let $B \subseteq \mathbb{R}^d \times \mathbb{R}$ have an injective and locally finite projection to $\mathbb{R}^d$. Then $x \in \mathbb{R}^d$ is a critical point of $vor \colon \mathbb{R}^d \to \mathbb{R}$ iff it is a critical point of $del \colon \mathbb{R}^d \to \mathbb{R}$ iff it is a critical point of $sd \colon \mathbb{R}^d \to \mathbb{R}$, and in this case $del(x) = sd(x) = vor(x)$ and the index of x is the same for all three functions.*

Proof. We prove that $x \in \mathbb{R}^d$ is a critical point (of vor, del, and sd) iff $x = \text{int}\,\nu \cap \text{int}\,\nu^*$ for a cell $\nu \in \text{Vor}(B)$ and its dual cell $\nu^* \in \text{Del}(B)$, and that the index of such a critical point is $q = \dim \nu^*$. Furthermore, $del(x) = sd(x) = vor(x)$ in this case by (4).

We begin with $f = vor$, which maps every $x \in \mathbb{R}^d$ to half the smallest power distance to a weighted point in B. The restriction of vor to a cell ν is also the restriction of a quadratic function on $\text{aff}\,\nu$ to ν. This quadratic function has a unique minimum, namely at $y = \text{aff}\,\nu \cap \text{aff}\,\nu^*$. The only possibility for a point $x \in \text{int}\,\nu$ to be a critical point of vor is therefore $x = y$. This implies that $\text{int}\,\nu \cap \text{aff}\,\nu^* \neq \emptyset$ is necessary for x to be critical. Symmetrically, $\text{aff}\,\nu \cap \text{int}\,\nu^* \neq \emptyset$ is necessary, which implies that $\text{int}\,\nu \cap \text{int}\,\nu^* \neq \emptyset$ is necessary. It is easy to see that the latter condition is also sufficient because vor increases along all directions within $\text{aff}\,\nu$ and it decreases in all directions within $\text{aff}\,\nu^*$. The index is the dimension of the affine subspace within which x is a maximum of f, which is $q = \dim \nu^*$, as claimed. The argument for $f = del$ is symmetric and therefore omitted. The index is still q, and not p as suggested by symmetry, because del maps every $x \in \mathbb{R}^d$ to the *negative* of the smallest power distance to a weighted point in C.

The argument for $f = sd$ is more involved. Since this function is piecewise linear, the only possible critical points are the vertices of $\text{Sd}(B)$. To simplify the argument, we assume that cells ν and μ^* with complementary dimensions have interiors that are either disjoint or intersect in a single point, which is therefore a vertex of $\text{Sd}(B)$. Writing $u = \text{int}\,\nu \cap \text{int}\,\mu^*$, we let $S_\varepsilon(u)$ be a sufficiently small sphere centered at u. It intersects a cell of $\text{Sd}(B)$ iff that cell is incident to u. The intersections of these cells with $S_\varepsilon(u)$ define a cell complex on the sphere. By construction, μ is dual to the collection of cells incident to μ^*, and ν^* is dual to the collection of cells incident to ν. Setting $P = \mu$ and $Q = \nu$, this implies that $P \times Q$ is dual to the collection of cells incident to u, and the boundary complex of $P \times Q$ is dual to the complex on $S_\varepsilon(u)$. Every point $v \in \mathbb{S}^{d-1}$ is a direction, and we write $sd_v(u)$ for the right derivative of sd at u in the direction v. The goal is to prove that the down- and up-links of u and sd are closely related to the in- and out-links of P and Q, namely

$$dn\text{Lk}(u, sd) \simeq in\text{Lk}(P, Q) \quad \text{and} \quad up\text{Lk}(u, sd) \simeq out\text{Lk}(P, Q). \tag{20}$$

By Lemma 3.4, the in- and out-links of P and Q either have the homotopy types of $\mathbb{S}^{q-1}$ and $\mathbb{S}^{p-1}$, if $\text{int}\,P \cap \text{int}\,Q \neq \emptyset$, or at least one link is contractible, if $\text{int}\,P \cap \text{int}\,Q = \emptyset$. Assuming (20), this implies that the down- and up-links of u and sd have the homotopy types of $\mathbb{S}^{q-1}$

and $\mathbb{S}^{p-1}$, if $\nu = \mu$, and at least one is contractible, if $\nu \neq \mu$. Indeed, $\nu \neq \mu$ together with $\operatorname{int} \nu \cap \operatorname{int} \mu^* \neq \emptyset$ implies $\operatorname{int} P \cap \operatorname{int} Q = \emptyset$ by Lemma 3.2.

We finally prove (20). Recall that every vertex of $P \times Q$ corresponds to a d-cell of $\mathrm{Sd}(B)$ incident to u, and every facet corresponds to an edge incident to u. Recall also that the map $\psi \colon \partial(P \times Q) \to \mathbb{S}^{d-1}$ introduced in Section 3.3 sends every vertex $s = y + z$ of $P \times Q$ to $\psi(s) = \frac{1}{2}(y - z)$. In the notation of equation (19), $y = p_c$ and $z = p_b$, so $\psi(s)$ is the gradient of sd restricted to the d-cell in $\mathrm{Sd}(B)$ that corresponds to s. To continue, we assume u is the origin of $\mathbb{R}^d$, we consider a facet E of $P \times Q$, and we let e be the corresponding edge of $\mathrm{Sd}(B)$ emanating from u. Observe that the gradient of the restriction of sd to the edge e is a constant multiple of the unit outer normal of $P \times Q$ at E, $\mathrm{const} \cdot \mathbf{n}_E$.

If a linear function $g \colon \mathbb{R}^d \to \mathbb{R}$ agrees with sd along e, then the projection of ∇g onto the line spanned by $\mathbf{n}_E$ is the gradient of the restriction. It follows that $\langle \nabla g, \mathbf{n}_E \rangle = \mathrm{const}$, and this holds in particular for the linear functions that correspond to the vertices of E. The gradient of any affine combination is the affine combination of the gradients. Hence, there is a unique affine combination of the functions corresponding to the vertices of E whose gradient is shortest, denoted g_E, and this gradient is of course $\mathrm{const} \cdot \mathbf{n}_E$. It follows that $\mathbf{n}_E$ belongs to $dn\mathrm{Lk}(u, sd)$ iff E belongs to $in\mathrm{Lk}(P, Q)$. By the nerve theorem, the full subcomplex of the decomposition of $S_\varepsilon(u)$ defined by the vertices with non-positive $\langle \nabla g_E, \mathbf{n}_E \rangle$ has the same homotopy type as $in\mathrm{Lk}(P, Q)$. The rest of the down-link deformation retracts to this full subcomplex, which implies the left homotopy equivalence in (20). The symmetric argument relating the up-link of u and sd with the out-link of P and Q implies the right homotopy equivalence in (20). This completes the proof. ◻

4 Discrete Functions

Parallel to the continuous functions studied in Section 3, we introduce discrete functions on the Voronoi tessellation, the Delaunay mosaics, and their common subdivision. We then study the structure of their steps, which we classify depending on their effect on the homology of the sublevel set.

4.1 Discrete Morse Theory

Letting K be a polyhedral complex in $\mathbb{R}^d$, we call $f \colon K \to \mathbb{R}$ a *discrete function*. It is *monotonic* if $f(\nu) \leq f(\mu)$ whenever ν is a face of μ in K, and it is *anti-monotonic* if $-f$ is monotonic. For every $t \in \mathbb{R}$, we call $f^{-1}(t)$ a *level set*, $f^{-1}(-\infty, t]$ a *sublevel set*, and $f^{-1}[t, \infty)$ a *superlevel set* of f. For completeness, we start by introducing the terminology of discrete Morse theory, which we adapt to polyhedral complexes.

The *Hasse diagram* of K is the directed graph whose nodes are the cells of K, with an arc from ν to μ if $\nu \subseteq \mu$ and $\dim \nu = \dim \mu - 1$. We note that $f \colon K \to \mathbb{R}$ is monotonic iff the values along every directed path of the Hasse diagram are non-decreasing. A *step* of f is a connected component of the Hasse diagram restricted to a level set of f, and we write ∇f for the collection of steps, which partitions K. We construct the *step graph* by taking the steps in ∇f as nodes and drawing an arc from I to J if there are cells $\nu \in I$ and $\mu \in J$ such that the Hasse diagram has an arc from ν to μ. In other words, the step graph is obtained from the Hasse diagram by contracting every arc whose end-cells share the function value. It follows that the values along every directed path of the step graph are strictly increasing.

A monotonic $f \colon K \to \mathbb{R}$ is a *discrete Morse function* if every step is either a pair or a singleton; see [9] but note that we inessentially simplified the setting by requiring that the cells in a pair share the same value. The singletons contain the *critical cells* and the pairs contain the *non-critical cells* of f. Following the convention in smooth Morse theory [15], where the index of a critical point is indicative of the effect of advancing the sublevel set beyond its value, we call the dimension of a critical cell its *index*. Indeed, adding a critical p-cell gives either birth to a p-cycle or death to a $(p-1)$-cycle, which affects the homology of the complex accordingly. In contrast, removing the two cells of a pair $\{\nu, \mu\}$ — which is allowed only if the result is still closed — has no effect on the homotopy type and therefore on the homology of the complex [9].

To generalize the concept, we call a subset $J \subseteq K$ an *interval* if there are cells $\alpha, \omega \in K$ such that $J = \{\nu \in K \mid \alpha \subseteq \nu \subseteq \omega\}$. In words, the interval has a unique *lower bound*, α, and a unique *upper bound*, ω, and consists of all faces of ω that have α as a face. A monotonic $f \colon K \to \mathbb{R}$ is a *generalized discrete Morse function* if every step is an interval; see [10]. The intervals of size one contain the *critical cells* and all other intervals contain the *non-critical cells* of f. Removing the cells of an interval of size larger than one from K is referred to as a *collapse*, which is allowed only if the result is still closed.

In the simplicial case, the Hasse diagram restricted to an interval is isomorphic to the 1-skeleton of a cube of the appropriate dimension. Choosing a direction, we get a collection of parallel edges of the cube, which corresponds to a partition of the interval into pairs. In the polyhedral case, such a partition is not quite as obvious but it exists. In other words, every collapse can be decomposed into a sequence of elementary collapses. The proof of this claim reduces to the fact that every convex polytope allows for a discrete Morse function with a single critical cell, which is a vertex [2]. This implies that if L can be obtained from K by a sequence of possibly non-elementary collapses, K and L have the same homotopy type.

4.2 Min and Max Functions

Taking the minimum or maximum over all points of a cell, we turn the continuous functions of Section 3 into discrete functions. In particular, we introduce $\mathtt{vor} \colon \mathrm{Vor}(B) \to \mathbb{R}$, $\mathtt{del} \colon \mathrm{Del}(B) \to \mathbb{R}$, and $\mathtt{sdn}, \mathtt{sdx} \colon \mathrm{Sd}(B) \to \mathbb{R}$ defined by

$$\mathtt{vor}(\tau) = \max_{x \in \tau^*} del(x), \tag{21}$$

$$\mathtt{del}(\sigma^*) = \min_{x \in \sigma} vor(x), \tag{22}$$

$$\mathtt{sdn}(\gamma) = \min_{x \in \gamma} sd(x), \tag{23}$$

$$\mathtt{sdx}(\gamma) = \max_{x \in \gamma} sd(x). \tag{24}$$

We note that $\mathtt{vor}$ is defined in terms of *del* and $\mathtt{del}$ in terms of *vor*. This is not a mistake but motivated by our desire to remain consistent with the standard literature on alpha shapes, where $\mathtt{del}$ is the (squared) radius function; see [6, 8]. It is also possible to define $\mathtt{vor}$ in terms of *vor* and $\mathtt{del}$ in terms of *del*, which gives slightly different discrete functions with essentially the same properties. It will often be convenient to apply the discrete Voronoi and Delaunay functions to the common subdivision. Technically, these are different functions, $\mathtt{sdv}, \mathtt{sdd} \colon \mathrm{Sd}(B) \to \mathbb{R}$, defined by $\mathtt{sdv}(\gamma) = \mathtt{vor}(\tau)$ and $\mathtt{sdd}(\gamma) = \mathtt{del}(\sigma^*)$, whenever $\gamma = \tau \cap \sigma^*$.

4.3 Classification with Homology

As introduced in Section 4.1, the step graph of a monotonic function defines a partial order on the steps. We can construct the complex by adding the steps one at a time according to a linear extension of this partial order. To determine the effect of adding a step to a subcomplex, we compute its relative homology, as we now explain.

Let $J_0, J_1, \ldots, J_m$ be a linear extension of the partial order defined by the step graph of $f \colon K \to \mathbb{R}$. This order may or may not be consistent with the sublevel sets of f, in the sense that the corresponding values listed in the same order may or may not be sorted. Write $K_j = \bigcup_{0 \leq i \leq j} J_i$, note that K_j is closed, and get $K_{j+1} = K_j \sqcup J_{j+1}$ by adding the next step. To describe how the addition of $J = J_{j+1}$ affects the homology of the complex, we consider the pair $(\bar{J}, \dot{J})$, in which $\bar{J} = \mathrm{cl}\, J$ is the closure and $\dot{J} = \bar{J} \setminus J$. Since $K_j \sqcup J$ is a complex, we have $\dot{J} = K_j \cap \bar{J}$, which is the intersection of two complexes and therefore a complex itself. We are interested in the relative homology of $(\bar{J}, \dot{J})$, since it will allow us to deduce the homology of K_{j+1} from that of K_j. Fixing a field to compute homology, we classify the steps according to the ranks of the relative homology groups, which we denote as $\beta_p = \mathrm{rank}\, \mathsf{H}_p(\bar{J}, \dot{J})$ for all dimensions p.

▶ **Definition 4.1** (Critical Step). *We call J a* non-critical step *of f if $\beta_p = 0$ for all $p \geq 0$. Otherwise, J is a* critical step. *It is a* simple critical step *of* index *p if all ranks vanish except in a single dimension, p, in which $\beta_p = 1$.*

We now explain how to deduce the homology of a complex from the homology of its predecessor and the relative homology of the step. We get the homology of $K_{j+1} = K_j \sqcup J$ using the long exact sequence of a pair:

$$\cdots \to \mathsf{H}_p(K_j) \to \mathsf{H}_p(K_{j+1}) \to \mathsf{H}_p(K_{j+1}, K_j) \to \mathsf{H}_{p-1}(K_j) \to \ldots. \tag{25}$$

Note that $\mathsf{H}_p(K_{j+1}, K_j)$ is isomorphic to $\mathsf{H}_p(\bar{J}, \dot{J})$ for every dimension p by excision. Assuming the ranks of the homology groups of K_j and of $(\bar{J}, \dot{J})$ are given, there are very few options for the ranks of K_{j+1} that make the sequence exact. For example, if J is a non-critical step, then $\mathrm{rank}\, \mathsf{H}_p(K_{j+1}) = \mathrm{rank}\, \mathsf{H}_p(K_j)$ for every p. If J is a simple critical step with index p, then either $\mathrm{rank}\, \mathsf{H}_p(K_{j+1}) = \mathrm{rank}\, \mathsf{H}_p(K_j) + 1$ or $\mathrm{rank}\, \mathsf{H}_{p-1}(K_{j+1}) = \mathrm{rank}\, \mathsf{H}_{p-1}(K_j) - 1$, with equal ranks in all other dimensions.

4.4 Critical and Non-critical Steps

Note that for a discrete or generalized discrete Morse function, every critical step is simple and indeed consists of only a single cell. In contrast, the discrete version of a generic piecewise linear map can have non-simple critical steps, such as monkey saddles, etc. However, these steps are still special since each has a unique lower bound, which is a vertex.

Similarly, the discrete functions in this paper are special cases within the general framework introduced in the previous subsection. In particular, each step of the Delaunay function, `del`: $\mathrm{Del}(B) \to \mathbb{R}$, has a unique upper bound, as we will prove shortly. To include the discrete Voronoi function in this discussion, we note that `vor`: $\mathrm{Vor}(B) \to \mathbb{R}$ is anti-monotonic, so $-$`vor` is monotonic, the above discussion applies, and every step of `vor` has a unique upper bound as well. Furthermore, the critical steps of `del` and `vor` contain a single cell each and are therefore simple, as we now prove.

▶ **Theorem 4.2** (Step Shape). *Every step of* `vor` *and of* `del` *has a unique upper bound, and if it is critical, then it consists of a single cell whose dimension is equal to the index of the step.*

Proof. We first prove that every step of `del` has a unique upper bound, and we omit the proof for `vor`, which is symmetric. By definition,

$$\mathtt{del}(\sigma^*) = \min_{x \in \sigma}[\varpi(x) - env(x)], \tag{26}$$

in which $env = \varpi - vor$ is piecewise linear and convex. Because ϖ is strictly convex, the minimum on the right-hand side of (26) is attained at a unique point, which we denote $y = y(\sigma)$. The step J of `del` that contains σ^* also contains every $\tau^* \in \mathrm{Del}(B)$ with $y(\tau) = y$. It contains no other cell, else there would be a cell with two points minimizing a strictly convex function. Without loss of generality, assume that σ^* is the unique cell in J such that σ contains y in its interior. It follows that $\sigma \subseteq \tau$ for all $\tau^* \in J$, which is equivalent to $\tau^* \subseteq \sigma^*$ for all $\tau^* \in J$. Hence, σ^* is the unique upper bound of J.

We second prove that every step that contains two or more cells is non-critical. Such a step, J, has a unique upper bound, σ^*. Write $q = \dim \sigma^*$, and let $A \subseteq B$ contain the weighted points such that σ^* is the convex hull of their locations. Let $S_r(x)$ be the smallest sphere such that $\pi_a(x) = r^2$ for every $a \in A$, and recall that this sphere is unique. Because σ^* is an upper bound, we have $\pi_b(x) > r^2$ for all $b \in B \setminus A$. All cells $\tau^* \in J \setminus \{\sigma^*\}$ are faces of σ^* that are *visible* from x. By this we mean that the line segment connecting x and a point $z \in \mathrm{int}\, \tau^*$ is disjoint from $\mathrm{int}\, \sigma^*$, while the line that passes through x and z has a non-empty intersection with $\mathrm{int}\, \sigma^*$. This implies that the union of interiors of the cells in $J \setminus \{\sigma^*\}$ is an open $(q-1)$-ball. As before, we define $\bar{J} = \mathrm{cl}\, J$ and $\dot{J} = \bar{J} \setminus J$. Since $\bar{J}$ is a closed q-ball and $\dot{J}$ is a closed $(q-1)$-ball in its boundary, the rank of $\mathsf{H}_p(\bar{J}, \dot{J}) = 0$ for every dimension p. Hence, J is non-critical, which implies that every critical step consists of a single cell, as claimed. Adding a cell of dimension q to the appropriate sublevel set affects either the q-th or the $(q-1)$-st homology group, which implies that the index of a critical step is the dimension of its cell, again as claimed. ◻

We observe that our definition of a critical step is consistent with that of a critical point. An interesting detail are the borderline non-critical points, which we recall have a contractible down-link and a non-contractible up-link, or the other way round. Correspondingly in the discrete setting, we call $\tau \in \mathrm{Vor}(B)$ a *borderline non-critical cell* if $\tau \cap \tau^* \neq \emptyset$ but $\mathrm{int}\, \tau \cap \mathrm{int}\, \tau^* = \emptyset$. A borderline critical cell is not critical, but there are arbitrarily small perturbations of the weighted points in B that render such a cell critical. Note that τ is a borderline non-critical cell of `vor` iff τ^* is a borderline non-critical cell of `del`. To bring such cases in focus, we introduce a condition that avoids them.

▶ **Definition 4.3** (General Position). *A set $B \subseteq \mathbb{R}^d \times \mathbb{R}$ with injective and locally finite projection to $\mathbb{R}^d$ is in* general position *if* `vor` *has no borderline non-critical cell or, equivalently, if* `del` *has no borderline non-critical cell.*

Note that this notion of general position is independent of the condition that guarantees simple Voronoi tessellations and simplicial Delaunay mosaics.

5 Complementing Subcomplexes

The main new concept in this section, is the channel between complementing subcomplexes of the tessellation and the mosaic. This channel acts like a buffer between the complexes, not unlike the buffer created from the second barycentric subdivision in the standard proof of Alexander duality [16].

5.1 Sub- and Superlevel Sets

Observe that for `del` and `sdx`, the value of a cell is larger than or equal to the values of its faces, and for `vor` and `sdn`, it is less than or equal to the values of its faces. It follows that the following sub- and superlevel sets are complexes:

$$\mathrm{Vor}^t(B) = \mathtt{vor}^{-1}[t, \infty), \tag{27}$$
$$\mathrm{Del}_t(B) = \mathtt{del}^{-1}(-\infty, t], \tag{28}$$
$$\mathrm{Sd}^t(B) = \mathtt{sdn}^{-1}[t, \infty), \tag{29}$$
$$\mathrm{Sd}_t(B) = \mathtt{sdx}^{-1}(-\infty, t]. \tag{30}$$

We extend (10) and (11) from the continuous to the discrete setting.

▶ **Lemma 5.1** (Nested Spaces). *Let $B \subseteq \mathbb{R}^d \times \mathbb{R}$ have an injective and locally finite projection to $\mathbb{R}^d$. Then $|\mathrm{Del}_t(B)| \subseteq |\mathrm{Sd}_t(B)|$ and $|\mathrm{Vor}^t(B)| \subseteq |\mathrm{Sd}^t(B)|$.*

Proof. Recall the functions $\mathtt{sdv}, \mathtt{sdd} \colon \mathrm{Sd}(B) \to \mathbb{R}$ introduced at the end of Section 4.2. By construction, the underlying spaces of their sub- and superlevel sets agree with those of `vor` and `del`. In particular, $|\mathtt{sdd}^{-1}(-\infty, t]| = |\mathrm{Del}_t(B)|$ and $|\mathtt{sdv}^{-1}[t, \infty)| = |\mathrm{Vor}^t(B)|$. By Lemma 3.1, we have

$$\mathtt{sdd}(\gamma) \geq \mathtt{sdx}(\gamma) \geq \mathtt{sdn}(\gamma) \geq \mathtt{sdv}(\gamma), \tag{31}$$

for every $\gamma \in \mathrm{Sd}(B)$. As illustrated in Figure 6, this implies $\mathtt{sdd}^{-1}(-\infty, t] \subseteq \mathtt{sdx}^{-1}(-\infty, t]$ and $\mathtt{sdv}^{-1}[t, \infty) \subseteq \mathtt{sdn}^{-1}[t, \infty)$. The sequence of inequalities in (31) thus imply the two claimed containment relations. ◻

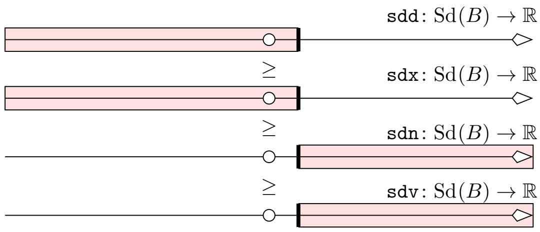

Figure 6: The four discrete functions on the common subdivision, which dominate each other from *top* to *bottom*. All indicated sub- and superlevel sets are for the same value, t.

Let $t \in \mathbb{R}$ be a value different from $sd(x)$ for all vertices x of $\mathrm{Sd}(B)$. Then $\mathrm{Sd}_t(B) \cap \mathrm{Sd}^t(B) = \emptyset$, and similarly their underlying spaces are disjoint. Combining the two relations in Lemma 5.1, we therefore have $|\mathrm{Del}_t(B)| \cap |\mathrm{Vor}^t(B)| = \emptyset$, which we illustrated in Figure 7. On the other hand, if t is the value of a vertex, x, then x belongs to $\mathrm{Sd}_t(B)$ as well as to $\mathrm{Sd}^t(B)$. If x is furthermore a critical point of sd, then x belongs also to $|\mathrm{Del}_t(B)|$ and to $|\mathrm{Vor}^t(B)|$.

5.2 Channel

Since the sub- and superlevel sets of `del` and `vor` considered in Lemma 5.1 have disjoint underlying spaces, it makes sense to study the space in between. For each value $t \in \mathbb{R}$,

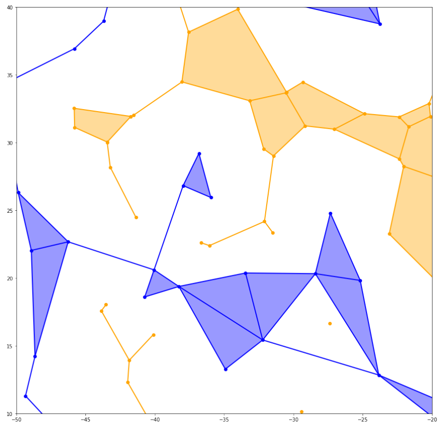

Figure 7: Complementing subcomplexes of the Voronoi tessellation, in *orange*, and the Delaunay mosaic, in *blue*. The complexes are constructed for a non-critical value of t, for which their underlying spaces are disjoint.

this is the underlying space of an open collection of cells in the common subdivision of the tessellation and the mosaic. For each cell $\gamma = \tau \cap \sigma^*$ in $\mathrm{Sd}(B)$, the relevant values are

$$t_0(\gamma) = \mathtt{sdv}(\gamma) = \mathtt{vor}(\tau), \tag{32}$$

$$t_1(\gamma) = \mathtt{sdd}(\gamma) = \mathtt{del}(\sigma^*). \tag{33}$$

Moving from $-\infty$ to ∞ along the real numbers, τ is dropped from $\mathrm{Vor}^t(B)$ at $t = t_0(\gamma)$ and σ^* is added to $\mathrm{Del}_t(B)$ at $t = t_1(\gamma)$. If τ is a critical cell of $\mathtt{vor}$ and $\sigma^* = \tau^*$ is the corresponding critical cell of $\mathtt{del}$, then γ is a point that belongs to both underlying spaces at $t = t_0(\gamma) = t_1(\gamma)$, and to exactly one of these underlying spaces for all other values of t. For all other cells $\gamma = \tau \cap \sigma^*$, Lemma 5.1 implies $t_0(\gamma) < t_1(\gamma)$. In all cases, γ belongs to the space in between $\mathrm{Del}_t(B)$ and $\mathrm{Vor}^t(B)$ for all $t_0(\gamma) < t < t_1(\gamma)$. More formally, we define the *channel* of B at t:

$$\mathrm{Ch}_t(B) = \{\gamma = \tau \cap \sigma^* \mid \tau \notin \mathrm{Vor}^t(B), \sigma^* \notin \mathrm{Del}_t(B)\}; \tag{34}$$

see Figure 8. This is the complement of the union of two subcomplexes of $\mathrm{Sd}(B)$ or, equivalently, the intersection of two open sets:

$$\mathrm{Ch}_t(B) = \mathrm{Sd}(B) \setminus \left[\mathtt{sdd}^{-1}(-\infty, t] \cup \mathtt{sdv}^{-1}[t, \infty)\right] \tag{35}$$

$$= \mathtt{sdd}^{-1}(t, \infty) \cap \mathtt{sdv}^{-1}(-\infty, t). \tag{36}$$

Recall that $\mathtt{sdd}(\gamma)$ is at least the maximum and $\mathtt{sdv}(\gamma)$ is at most the minimum $sd(x)$ over all points $x \in \gamma$. It follows that $sd^{-1}(t)$ is disjoint of the underlying spaces of $\mathtt{sdd}^{-1}(-\infty, t]$ and $\mathtt{sdv}^{-1}[t, \infty)$, unless t is a critical value of sd, in which case the corresponding critical points belong to all three. Hence, $sd^{-1}(t)$ is contained in the underlying space of the channel, unless t is a critical value, in which case the level set passes through the corresponding critical points. We state this insight together with a straightforward related property more formally.

▶ **Theorem 5.2** (Split Channel). *Let $B \subseteq \mathbb{R}^d \times \mathbb{R}$ have an injective and locally finite projection to $\mathbb{R}^d$, and let $t \in \mathbb{R}$ be different from all critical values of sd. Then*
- $sd^{-1}(t) \subseteq |\mathrm{Ch}_t(B)|$,
- $sd^{-1}(t)$ *is an orientable $(d-1)$-manifold.*

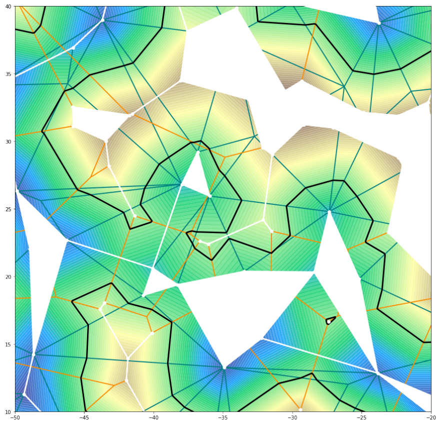

Figure 8: The channel decomposed into cells of the common subdivision of the Voronoi tessellation and the Delaunay mosaic, with the two complementing subcomplexes forming the *white* background. In *black*, we superimpose the level set of sd for the value of t that splits the channel into two.

On the other hand, if t is a critical value of sd, then both of these properties are violated, but only at the corresponding critical points, and at these points both, level set and channel, go through topological reorganization.

5.3 Evolution of Channel

For every non-critical value $t \in \mathbb{R}$, we have a partition of $\mathbb{R}^d$ into the underlying space of the superlevel set of `vor`, of the sublevel set of `del`, and of the channel in between. We are interested in the evolution of this partition as t goes from $-\infty$ to ∞. It is convenient to study the corresponding partition of the common subdivision,

$$\mathrm{Sd}(B) = \mathtt{sdd}^{-1}(-\infty, t] \sqcup \mathrm{Ch}_t(B) \sqcup \mathtt{sdv}^{-1}[t, \infty), \tag{37}$$

as t goes from $-\infty$ to ∞. At the beginning, the only non-empty set in the partition is the superlevel set of `sdv`, and step by step the cells migrate first to the channel and second to the sublevel set of `sdd`, until, at the end, the latter is the only non-empty subset in the partition. Indeed, every change in this process is the migration of a step of `sdv` to the channel or the migration of a step of `sdd` from the channel. We distinguish between non-critical steps and critical steps of index q, with $0 \leq q \leq d$. By Theorem 4.2, the cells of an index q critical step subdivide an open q-cell in $\mathrm{Del}(B)$ or in $\mathrm{Vor}(B)$.

Write J_i and t_i for the steps of `sdv` and `sdd` and their values, for $0 \leq i \leq m$. We assume the indexing satisfies $t_i \leq t_{i+1}$ for $0 \leq i < m$, and in case of a tie, the steps of `sdv` precede those of `sdd`. Write V_i and D_i for the two complexes after processing steps J_0 through J_i, and let $C_i = \mathrm{Sd}(B) \setminus [V_i \sqcup D_i]$ be the third set in the partition. We get the next partition as

$$\begin{array}{lll} V_{i+1} = V_i \setminus J_{i+1}, & C_{i+1} = C_i \sqcup J_{i+1}, & D_{i+1} = D_i, \\ V_{i+1} = V_i, & C_{i+1} = C_i \setminus J_{i+1}, & D_{i+1} = D_i \sqcup J_{i+1}, \end{array}$$

in which the first row describes the change if the step belongs to `sdv` and the second row if the step belongs to `sdd`. To avoid discussing the homology of unbounded spaces, we add a point at infinity to compactify $\mathbb{R}^d$ to $\mathbb{S}^d$.

CASE J_{i+1} is non-critical. Then the p-th homology groups of V_i and V_{i+1} are isomorphic, and so are the p-th homology groups of D_i and D_{i+1}, for every p.

Case J_{i+1} is an index q critical step of `sdv`. Then either $\beta_q(V_{i+1}) = \beta_q(V_i) - 1$ or $\beta_{q-1}(V_{i+1}) = \beta_{q-1}(V_i) + 1$, with equality for the ranks in all other dimensions.

Case J_{i+1} is an index q critical step of `sdd`. Then either $\beta_q(D_{i+1}) = \beta_q(D_i) + 1$ or $\beta_{q-1}(D_{i+1}) = \beta_{q-1}(D_i) - 1$, with equality for the ranks in all other dimensions.

Recall that the critical steps come in pairs of complementary indices $p + q = d$. Assuming J_{i+1}, J_{i+2} is such a pair of critical steps, one of `sdv` and the other of `sdd`, we get either $\beta_p(V_{i+2}) = \beta_p(V_i) - 1$ or $\beta_{p-1}(V_{i+2}) = \beta_{p-1}(V_i) + 1$ for the ranks on one side of the channel, and either $\beta_q(D_{i+2}) = \beta_q(D_i) + 1$ or $\beta_{q-1}(D_{i+2}) = \beta_{q-1}(D_i) - 1$ for the ranks on the other side of the channel. This is consistent with Alexander duality but fails to imply it as we did not yet pair up the events on the two sides.

5.4 Crushing the Channel

This subsection addresses the missing step in the proof of Alexander duality for V_i and D_i. To this end, we show that the channel that separates the two complexes can be deformation retracted. Let $t \in \mathbb{R}$ such that $D_i = \mathrm{Del}_t(B)$ and $V_i = \mathrm{Vor}^t(B)$, and recall that $|D_i| \subseteq vor^{-1}(-\infty, t]$ and $|V_i| \subseteq del^{-1}[t, \infty)$. Since the situation is symmetric, it suffices to talk about D_i. By definition, a *boundary cell* of D_i is contained in $\partial|D_i|$, and by construction, $\sigma^* \in D_i$ is a boundary cell iff the intersection of the corresponding spheres has a non-empty contribution to the boundary of $vor^{-1}(-\infty, t]$. Letting p be the dimension of the dual cell, $\sigma \in \mathrm{Vor}(B)$, and p_a be one of the vertices of σ^*, this contribution is $A_\sigma = \sigma \cap S_r(p_a)$, in which the squared radius of the sphere is $r^2 = w_a + t$. Hence, A_σ is a subset of a $(p-1)$-sphere, which may or may not be connected. An important part of the construction is the *join* of σ^* and A_σ, which is the union of line segments connecting the two sets:

$$\sigma^* * A_\sigma = \{(1-\lambda)y + \lambda z \mid y \in \sigma^*, z \in A_\sigma, 0 \leq \lambda \leq 1\}. \tag{38}$$

Writing $U_t = vor^{-1}(-\infty, t]$ and following [4], we decompose $U_t \setminus |D_i|$ into such joins. The deformation retraction will happen along the *fibers* of this decomposition, which are the line segments in the joins. We therefore need that the fibers cover $U_t \setminus |D_i|$ and that they do not intersect except at shared endpoints. But this is clear because the entire decomposition can be obtained by projecting pieces of a convex surface in $\mathbb{R}^{d+1}$ to $\mathbb{R}^d$. This surface is the boundary of the convex hull of the graphs of *end* and $\varpi + t$. The pieces that belong to the graph of *end* project to cells in D_i, the pieces that bridge the gap between the two graphs project to the joins, and the rest belongs to the graph of ϖ, which we do not project.

We now return to splitting the channel along the middle, by which we mean that we split it along $sd^{-1}(t)$. It is important that each fiber intersect this level set in exactly one point.

▶ **Lemma 5.3** (Fiber Crossing). *Let $B \subseteq \mathbb{R}^d \times \mathbb{R}$ have an injective and locally finite projection to $\mathbb{R}^d$, let $t \in \mathbb{R}$ be a non-critical value, and let y, z be endpoints of a fiber in the decomposition of $U_t \setminus |\mathrm{Del}_t(B)|$. Then there is a unique $0 \leq \lambda \leq 1$ such that $sd((1-\lambda)y + \lambda z) = t$.*

Proof. We have $sd(y) < t < sd(z)$ for the fiber with endpoints $y \in \sigma^*$ and $z \in A_\sigma$. It follows that the fiber intersects $sd^{-1}(t)$ an odd number of times. To show that this number is 1, we recall that the sublevel set of *vor* and the superlevel set of *del* are both unions of balls:

$$vor^{-1}(-\infty, t] = \bigcup_{a \in B} a_t \quad \text{and} \quad del^{-1}[t, \infty) = \bigcup_{c \in C} c_t, \tag{39}$$

in which a_t is the ball with center p_a and squared radius $w_a + t$, for $a \in B$, and c_t is the ball with center p_c and squared radius $w_c - t$, for $c \in C$. By construction, we have

$\|p_a - p_c\|^2 \geq w_a + t + w_c - t$; that is: a_t and c_t are orthogonal or further than orthogonal from each other. Recall that $del(x) \geq sd(x) \geq vor(x)$ for every $x \in \mathbb{R}^d$, by Lemma 3.1. This implies that the two unions of balls cover the entire $\mathbb{R}^d$, and that the level set of sd at t is contained in their intersection; see Figure 3 and equations (10) and (11). We first consider the special case in which y is a vertex of $\mathrm{Del}_t(B)$ and z is a point of the sphere bounding the corresponding ball: assuming $w_a + t > 0$, we set $y = p_a$ and let z be a point on the boundary of a_t. Of course y and z belong to the boundary of their respective sets. Assuming c_t contains z, there is a unique $0 < \lambda_c \leq 1$ such that $x = (1-\lambda)y + \lambda z$ belongs to c_t iff $\lambda_c \leq \lambda$. Setting $\lambda_c = \infty$ if c_t does not contain z, we let $\lambda_{\min} = \min_{c \in C} \lambda_c$. Hence, x belongs to $del^{-1}[t, \infty)$ iff $\lambda_{\min} \leq \lambda$. We prove the claim by first extending this construction to general fibers and second arguing about the overlap of the two unions of balls.

Let $y \in \sigma^*$ and $z \in A_\sigma$ be the endpoints of a fiber, and consider the ball with center y and squared radius $\|z - y\|^2$. It is not necessarily a ball a_t with $a \in B$, but it is contained in the union of balls a_t, with p_a a vertex of σ^*, and its boundary contains the intersection of the boundaries of these balls. It follows that it is orthogonal or further than orthogonal from all balls c_t, with $c \in C$. By construction, $z \in del^{-1}[t, \infty)$, so there is a unique $0 < \lambda_{\min} \leq 1$ such that a point $x = (1-\lambda)y + \lambda z$ of the fiber belongs to $del^{-1}[t, \infty)$ iff $\lambda_{\min} \leq \lambda$. In summary, the points at which the fiber intersects the level set all lie between $y' = (1-\lambda_{\min})y + \lambda_{\min} z$ and z. Write $[y', z]$ for this portion of the fiber, which we orient from y' to z. It is not difficult to see that the restriction of vor to $[y', z]$ is a strictly increasing piecewise quadratic function. Indeed, if there is a vertex p_a of σ^* such that the Voronoi cell of a contains the entire segment from y' to z, then vor restricted to $[y', z]$ is quadratic and its extension along the line attains its minimum outside $[y', z]$, namely at p_a, which lies before y'. If there is no such vertex p_a, then we trace the segment from z back to y', passing through a sequence of Voronoi cells. Each time we pass from one cell to another, the slope of the restriction of vor increases. It follows that also in this case, we reach y' before we reach a minimum. Similarly, the restriction of del to $[y', z]$ is a strictly increasing piecewise quadratic function. It follows that sd restricted to $[y', z]$ is a strictly increasing piecewise linear function, which implies that it crosses t exactly once. Hence, the fiber intersects $sd^{-1}(t)$ in exactly one point, as claimed. ◻

To construct the deformation retraction, we clip every fiber where it intersects $sd^{-1}(t)$ and retract the remaining piece to its endpoint in $|D_i|$. To describe this formally, we write $z' = (1-\lambda')y + \lambda' z$, with λ' the unique solution to $sd((1-\lambda)y + \lambda z) = t$, and we write $M_t = sd^{-1}(-\infty, t]$ and $M^t = sd^{-1}[t, \infty)$. To deformation retract M_t to $|D_i| = |\mathrm{Del}_t(B)|$ we use $D \colon M_t \times [0,1] \to M_t$, which is the identity on $|D_i|$ and otherwise maps a point $x = (1-\lambda)y + \lambda z'$ to $D(x,s) = (1-s)x + sy$, for every $s \in [0,1]$. Symmetrically, we deformation retract M^t to $|V_i| = |\mathrm{Vor}^t(B)|$. We formally state the implications.

▶ **Theorem 5.4** (Crushing). *Let $B \subseteq \mathbb{R}^d \times \mathbb{R}$ have an injective and locally finite projection to $\mathbb{R}^d$ and $t \in \mathbb{R}$ be non-critical. Then $|\mathrm{Del}_t(B)| \simeq M_t$ and $|\mathrm{Vor}^t(B)| \simeq M^t$.*

In words, the channel can be split into halves, each half can be decomposed into line segments called fibers, and by retracting the fibers, we glue the boundaries of $|\mathrm{Del}_t(B)|$ and $|\mathrm{Vor}^t(B)|$ without altering the homotopy type, which is that of $\mathbb{R}^d$ or, after compactification, that of $\mathbb{S}^d$. Hence, Alexander duality applies, so we get $\beta_{q-1}(\mathrm{Vor}^t(B)) = \beta_p(\mathrm{Del}_t(B))$ for all dimensions $p + q = d$, except when $p = 0$ or $q = 0$ in which case the two ranks differ by 1. Recalling the parallel change of the two complexes discussed above, we now conclude that we see the birth of a p-dimensional homology class in $\mathrm{Vor}^t(B)$ iff we see the birth of a $(q-1)$-dimensional homology class in $\mathrm{Del}_t(B)$ at the same threshold, and similarly for the death of such classes.

6 Discussion

Motivated by challenges caused by data in non-general position, this paper explores the continuous and discrete functions that define Voronoi tessellations and Delaunay mosaics. Beyond the concrete results formulated as lemmas and theorems, we mention the generalization of key concepts in discrete Morse theory as one of the main contributions of this paper. In the process of gaining new insights into an old subject, we encountered questions we have not been able to answer:

- The piecewise linear $sd\colon \mathbb{R}^d \to \mathbb{R}$ can be defined for sets $B, C \subseteq \mathbb{R}^d \times \mathbb{R}$ that do not satisfy the polar relationship assumed in this paper. What are its properties, and what additional features does it enjoy when B and C are polar, as assumed in this paper?
- In $\mathbb{R}^3$, the union of balls is a popular model of a molecule [13], albeit in practice easier to compute and easier to display PL surfaces are preferred. These do generally not have the homotopy type of the boundary of the union of balls. The level set of $sd\colon \mathbb{R}^3 \to \mathbb{R}$ suggests itself as an easy to use yet topologically correct alternative. What are its combinatorial and geometric properties, and how fast can they be computed?
- We prove in this paper that the channel deformation retracts to the Voronoi complex as well as the complementing Delaunay complex. Can the same result be obtained with discrete methods, for example by collapsing the steps of the discrete versions of sd?

The discrete functions defined in this paper gives rise to a one-parameter family of complementing complexes. It would be interesting to connect these families to applications, such as the study of Raleigh–Bénard convection with its family of bi-partitions of space [12].

References

1 U. Bauer and H. Edelsbrunner. The Morse theory of Čech and Delaunay complexes. *Trans. Amer. Math. Soc.*, **369** (2017), 3741–3762.

2 M.K. Chari. On discrete Morse functions and combinatorial decompositions. *Discrete Math.* **217** (2000), 101–113.

3 B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk* **7** (1934), 793–800.

4 H. Edelsbrunner. The union of balls and its dual shape. *Discrete Comput. Geom.* **13** (1995), 415–440.

5 H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. *Discrete and Computational Geometry. The Goodman–Pollack Festschrift*, 379–404, eds.: B. Aronov, S. Basu, J. Pack and M. Sharir, Springer-Verlag, 2003.

6 H. Edelsbrunner, D.G. Kirkpatrick and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory* **IT-29** (1983), 551–559.

7 H. Edelsbrunner and P. Koehl. The geometry of biomolecular solvation. In *Discrete and Computational Geometry*, 243–275, eds. J.E. Goodman, J. Pach and E. Welzl, MSRI Publ. **52**, Cambridge Univ. Press, England, 2005.

8 H. Edelsbrunner and E.P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics* **13** (1994), 43–72.

9 R. Forman. Morse theory for cell complexes. *Adv. Math.* **134** (1998), 90–145.

10 R. Freij. Equivariant discrete Morse theory. *Discrete Math.* **309** (2009), 3821–3829.

11 J.D. Gardiner, J. Behnsen and C.A. Brassey. Alpha shapes: determining 3D shape complexity across morphologically diverse structures. *BMC Evolutionary Biology* **18** (2018), article 184.

12 K. Krishan, H. Kurtuldu, M. Schatz, M. Gameiro, K. Mischaikow and S. Madruga. Homology and symmetry breaking in Rayleigh–Bénard convection: experiments and simulations. *Phys. Fluids* **19** (2007), 117105.

13 B. Lee and F.M. Richards. The interpretation of protein structures: estimation of static accessibility. *J. Mol. Biol.* **55** (1971), 379–400.

14 M. Maddah and C.G.L. Cao. Application of the alpha shape method to visualize and analyze surgical motions. *Surgical Science* **08** (2017), 464–480.

15 J. Milnor. *Morse Theory.* Princeton Univ. Press, Princeton, New Jersey, 1963.

16 J.R. Munkres. *Elements of Algebraic Topology.* Perseus, Cambridge, Massachusetts, 1984.

17 D. Pedoe. *Geometry: a Comprehensive Course.* Dover, New York, 1988.

18 E.A. Ramos and B. Sadri. Geometric and topological guarantees for the Wrap reconstruction algorithm. *In* “Proc. 18th Ann. ACM-SIAM Sympos. Discrete Alg., 2007”, 1086–1095.

19 G. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième Mémoire: Recherches sur les parallélloèdres primitifs. *J. Reine Angew. Math.* **134** (1908), 198–287.

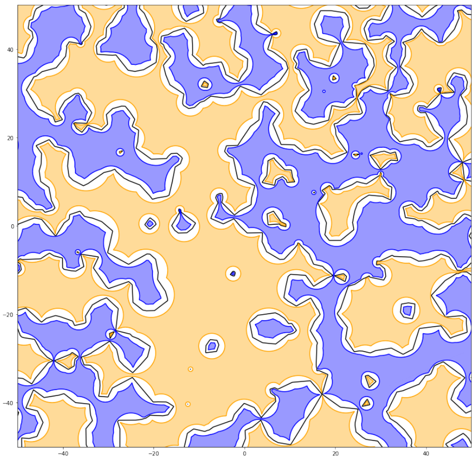

(a) Preimages of continuous functions.

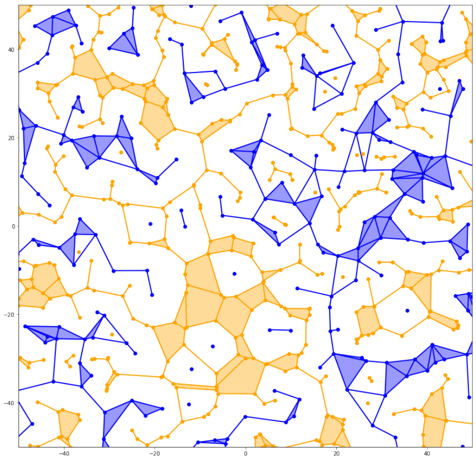

(b) Preimages of discrete functions.

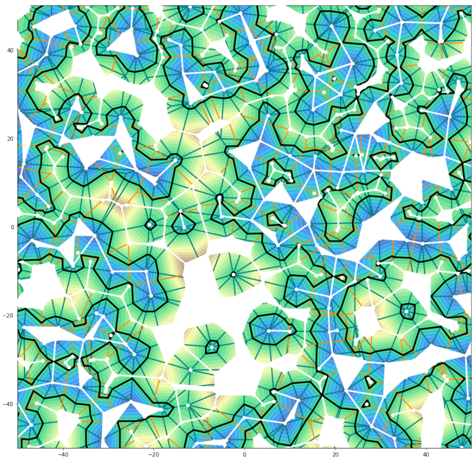

(c) Channel split by level set.

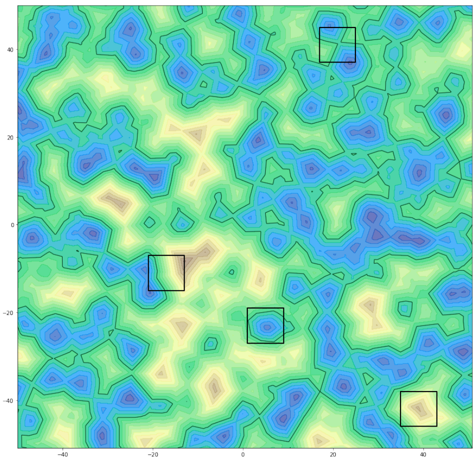

(d) Levels of piecewise linear function.

Figure 9: Pictures of the same decomposition of the plane into “land” and “water”. All geometric structures are for the same value of t: (a) sub-, super-, and level sets of three continuous functions; (b) sub- and superlevel sets of the discrete functions on the Voronoi tessellation and the Delaunay mosaic; (c) channel divided by level set of piecewise linear function; (d) level sets of piecewise linear functions, with square boxes marking the neighborhoods of a non-critical point, a minimum, a saddle, and a maximum.