

Multi-cover persistence and Delaunay mosaics

by

Georg Osang

January, 2021

*A thesis presented to the
Graduate School
of the
Institute of Science and Technology Austria, Klosterneuburg, Austria
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy*



Institute of Science and Technology

The thesis of Georg Osang, titled *Multi-cover persistence and Delaunay mosaics*, is approved by:

Supervisor: Herbert Edelsbrunner, IST Austria, Klosterneuburg, Austria

Signature: _____

Committee Member: Uli Wagner, IST Austria, Klosterneuburg, Austria

Signature: _____

Committee Member: Mohammad Saadatfar, Australian National University, Canberra, Australia

Signature: _____

Defense Chair: Scott Waitukaitis, IST Austria, Klosterneuburg, Austria, IST Austria, Klosterneuburg, Austria

Signature: _____

signed page is on file

© Georg Osang, January, 2021

This work is licensed under a Creative Commons Attribution 4.0 International License.



IST Austria Thesis, ISSN: 2663-337X

I hereby declare that this thesis is my own work and that it does not contain other people's work without this being so stated; this thesis does not contain my previous work without this being stated, and the bibliography contains all the literature that I used in writing the dissertation.

I declare that this is a true copy of my thesis, including any final revisions, as approved by my thesis committee, and that this thesis has not been submitted for a higher degree to any other university or institution.

I certify that any republication of materials presented in this thesis has been approved by the relevant publishers and co-authors.

Signature: _____

Georg Osang

January, 2021

signed page is on file

Abstract

In this thesis we study persistence of multi-covers of Euclidean balls and the geometric structures underlying their computation, in particular Delaunay mosaics and Voronoi tessellations.

The k -fold cover for some discrete input point set consists of the space where at least k balls of radius r around the input points overlap. Persistence is a notion that captures, in some sense, the topology of the shape underlying the input. While persistence is usually computed for the union of balls, the k -fold cover is of interest as it captures local density, and thus might approximate the shape of the input better if the input data is noisy. To compute persistence of these k -fold covers, we need a discretization that is provided by higher-order Delaunay mosaics.

We present and implement a simple and efficient algorithm for the computation of higher-order Delaunay mosaics, and use it to give experimental results for their combinatorial properties. The algorithm makes use of a new geometric structure, the rhomboid tiling. It contains the higher-order Delaunay mosaics as slices, and by introducing a filtration function on the tiling, we also obtain higher-order α -shapes as slices. These allow us to compute persistence of the multi-covers for varying radius r ; the computation for varying k is less straight-forward and involves the rhomboid tiling directly.

We apply our algorithms to experimental sphere packings to shed light on their structural properties. Finally, inspired by periodic structures in packings and materials, we propose and implement an algorithm for periodic Delaunay triangulations to be integrated into the Computational Geometry Algorithms Library (CGAL), and discuss the implications on persistence for periodic data sets.

Acknowledgments

First and foremost, I would like to thank Herbert, my supervisor, for guiding and supporting me, for giving me a lot of freedom and yet always being available when I needed help. I would also like to thank Uli and Mohammad for being on my committee, and Mohammad in particular for inviting me to Canberra to get a better understanding of his experiments and work in his field. I would like to thank Monique for welcoming me to collaborate on periodic Delaunay triangulations, and Mael for getting deep into the CGAL code and guiding me to make the project happen. I am also grateful to Michael Kerber for inviting me to Graz to discuss 2-parameter persistence, and, together with René Corbet and Michael Lesnick, collaborating on the topic. Thank you also to Vitaliy Kurlin for the opportunity to learn more about crystallographic applications of my work at the Materials Innovation Factory in Liverpool.

I would also like to thank my research group for the company throughout the years, and in particular Hubert Wagner for introducing me to the group and being my main point of contact in the early days.

I will refrain from individual mentions of family, friends and colleagues here. You know who you are.

About the Author

Georg Osang completed a double BSc in mathematics and informatics in 2012 at the Karlsruhe Institute of Technology. He was awarded a MMath in Combinatorics and Optimization from the University of Waterloo in 2014. He spent part of 2014 at the African Institute of Mathematical Sciences in Ghana as a tutor before joining IST Austria as a PhD student. Throughout his PhD he has stayed in contact with the maths education community in Africa via maths camps and collaborating with African Maths Initiatives, and has completed three internships at Google.

List of Publications

Publications which are part of the thesis:

1. H. Edelsbrunner, G. Osang. *The Multi-cover Persistence of Euclidean Balls*. Proceedings of the 34th International Symposium on Computational Geometry, 2018; journal version to appear in Discrete & Computational Geometry.
2. H. Edelsbrunner, A. Nikitenko, G. Osang. *A Step in the Delaunay Mosaic of Order k* . To appear in Journal of Geometry.
3. H. Edelsbrunner, G. Osang, M. Saadatfar. *Topological signatures and stability of hexagonal close packing and Barlow stackings*. Submitted to Soft Matter.
4. H. Edelsbrunner, G. Osang. *A Simple Algorithm for Computing Higher-order Delaunay Mosaics*. Submitted to Algorithmica.
5. G. Osang, M. Rouxel-Labbé, M. Teillaud. *Generalizing CGAL Periodic Delaunay Triangulations*. Proceedings of the 28th Annual European Symposium on Algorithms, 2020.
6. R. Corbet, M. Kerber, M. Lesnick, G. Osang. *Computing the multicover bifiltration*. Submitted to Symposium on Computational Geometry, 2021.

The content of (1) is distributed across Chapters 3 to 5, and includes the introduction of the rhomboid tiling, the computation of persistence in depth and the property that the rhomboid filtration function is generalized discrete Morse. The remainder of Chapter 3 is based on (4) and the remainder of Chapter 5 is based on (2). The results from (6) are summarized in Section 4.5. Chapter 6 is based on (3) and Chapter 7 is based on (5).

Publications not related to the thesis:

7. K. Chatterjee, G. Osang. *Pushdown Reachability with Constant Treewidth*. Information Processing Letters, 2017.
8. J. Alwen, P. Gaži, C. Kamath, K. Klein, G. Osang, K. Pietrzak, L. Reyzin, M. Rolínek, M. Rybár. *On the Memory-Hardness of Data-Independent Password-Hashing Functions*. Proceedings of the 2018 on Asia Conference on Computer and Communications Security.

Table of Contents

Abstract	v
Acknowledgments	vi
About the Author	vii
List of Publications	viii
1 Introduction	1
2 Background	8
2.1 Cell complexes	9
2.2 Voronoi tessellations	10
2.3 Hyperplane arrangements	11
2.4 Delaunay mosaics	13
2.5 Weighted notions	15
2.6 Algorithms and complexity	17
3 Computing higher-order Delaunay mosaics	18
3.1 Rhomboid Tiling	20
3.2 Combinatorial Properties	25
3.2.1 Predicting Cells	26
3.2.2 Identifying Vertices	28
3.3 Algorithm for higher-order Delaunay mosaics	29

3.4	Experimental Results	33
3.5	Extensions	38
4	Multi-cover persistence	41
4.1	Homotopy	43
4.2	Computing the Delaunay filtration	46
4.2.1	Optimizing the computation of constrained spheres	47
4.3	Homology and Persistence	50
4.3.1	Persistence in scale	53
4.4	Persistence in depth	55
4.4.1	Rhomboid filtration	55
4.4.2	Computation	59
4.5	2-parameter persistence	65
5	Discrete Morse theory	71
5.1	Topology of rhomboid steps	72
5.1.1	Computing intervals	75
5.2	Topology of Delaunay steps	78
5.3	Discussion	83
6	Analysis of sphere packings	84
6.1	Persistent features in FCC and HCP	88
6.2	Analysis of experimental sphere packings	93
6.2.1	Different densities.	93
6.2.2	Frequency measures.	94
6.2.3	Stability of FCC and HCP	95
6.3	Conclusions	96
7	Periodic Delaunay triangulations	100
7.1	Preliminaries	103







7.2	Algorithm	105
7.2.1	Phase 1	107
7.2.2	Transition	114
7.2.3	Phase 2	116
7.3	Experimental results	117
7.4	Discussion	119
7.4.1	Extension to weighted points	121
7.4.2	Extension to order- k Delaunay mosaics	123
7.4.3	Persistence of periodic point sets	124
	Bibliography	127

1 Introduction

Persistent homology [35; 74; 33], or persistence for short, is a notion used to characterize and quantify topological features in data. It captures changes in topology (or more specifically, homology) in a space that changes in accordance with some scale parameter. One specific setting is when the input data set is a discrete set of points. Given an input point set $P \subset \mathbb{R}^d$, the space in question is the union of balls of radius r around the points of P . Its topology changes as the parameter r increases from 0 to infinity, and these changes are captured by a *persistence diagram*. The persistence diagram is a 2-dimensional visual descriptor of the data set quantifying the life-span, called *persistence*, of its topological features, including the information when they emerge and disappear. For example, if the input P is thought of as a random sample of some underlying shape, then persistence captures topological features of the shape, see Figure 1.1. However persistence of the union of balls also has applications to other kinds of discrete point data sets, such as atom configurations of materials [50; 59] or sphere packings [76].

One crucial property of persistence diagrams is *stability* [22]: Small perturbations in the input imply small changes in the persistence diagram. However, persistence of the union of balls is not robust to noise, see Figure 1.3. The introduction of a single new data point can cause the persistence diagram to change significantly.

In this thesis we will extend the notion of persistence of the union of balls to take into account local density fluctuations. These fluctuations can be large — and the task may be the identification of regions with a prescribed density profile — or they can be small — and the goal may be to pick up subtle variations. For example, we may want to quantify local defects in lattice configurations or describe long-range differences between similar configurations, such as the *face-centered cubic* (FCC) lattice and the

(a) $r = 4$. Many small components.(b) $r = 9$. A moon-shaped loop emerges.(c) $r = 20$. The loop is about to close up.(d) $r = 25$. Star-shaped components merge.(e) $r = 46$. A short-lived loop appears.

(f) Persistence diagram.

Figure 1.1: The union of balls filtration and the corresponding persistence diagram. For example, the moon-shaped topological loop appears shortly before (b) at radius $r = 8$ and disappears shortly after (c) at $r = 22$, giving a point at $(8, 22)$ in the persistence diagram.

hexagonal close-packed (HCP) configuration in \mathbb{R}^3 . While both give densest sphere packings in \mathbb{R}^3 , physical particle systems prefer to settle in the FCC configuration [85; 41; 72]. The reason for this preference is not well understood. Our quantification of the long-range effects of density differences discriminates between the two configurations and in this way sheds light on this phenomenon.

To achieve this, we use a generalization of the union of balls which is the *k-fold cover*. It consist of those points in the underlying space \mathbb{R}^d that are contained in at least k balls of radius r around the points of P , see Figure 1.2.




Figure 1.2: Balls (pale pink) around six points in the plane; the 2-fold cover (dark pink) is where at least two balls overlap.

As we can see in Figure 1.3, the k -fold cover for appropriately chosen k can capture the underlying shape of a given point sample better than the union of balls if the input point set is noisy. Analogous to persistence of the union of balls, persistence of the k -fold cover, or *multi-cover persistence*, captures the changes in topology of the k -fold cover. The most natural change to the k -fold cover is incurred when varying the radius r of the balls, which we call *persistence in scale*. However, perhaps slightly less intuitively, one can also track the changes of the k -fold cover as the radius r stays fixed and instead the parameter k is changed, which we call *persistence in depth*. We will solve the computation of both persistence in scale and in depth, and furthermore discuss persistence in both parameters simultaneously.

Computation The well known Voronoi tessellations [88] decompose the union of balls into convex pieces. This gives rise to subcomplexes of their dual Delaunay triangulations,

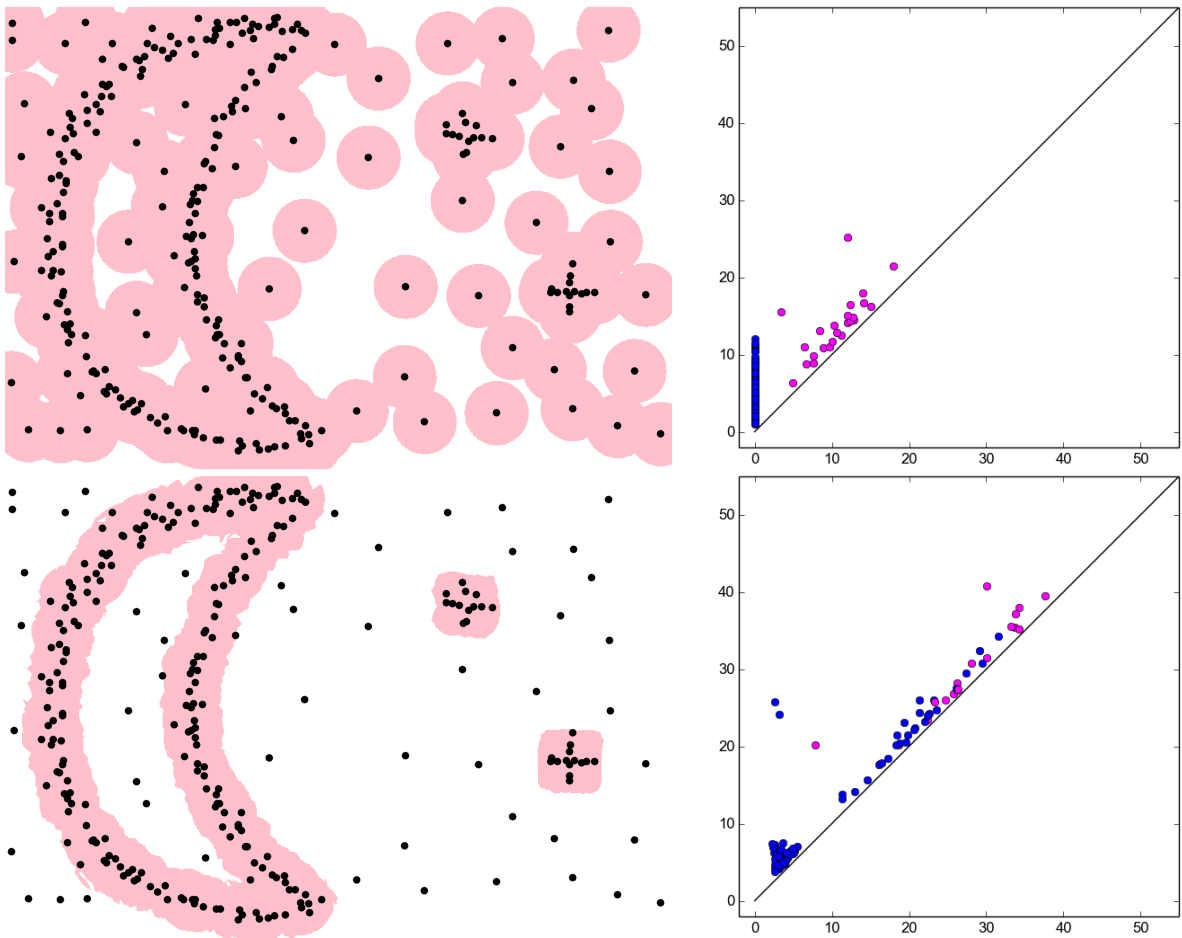


Figure 1.3: The union of balls (top) and 5-fold cover (bottom) of radius $r = 10$ for a noisy point set (left). On the right, the corresponding persistence diagrams. While the features from the persistence diagram in Figure 1.1f do not clearly stand out from the noise anymore in the top diagram, they are retained in the bottom diagram.

called α -shapes [34]. They provide a discretization of the union of balls while maintaining its topology. Similarly, in order to compute the topology of the k -fold cover, we need a discrete representation of it. Order- k Voronoi tessellations are a generalization of the classic Voronoi tessellations and were introduced by [77] as a data structure for fast k closest point queries. Similarly they decompose the k -fold covers and give rise to order- k α -shapes [55], which are subcomplexes of their dual order- k Delaunay mosaics [7] and capture the topology of the k -fold covers.

While algorithms for computing order- k Voronoi tessellations and their dual Delaunay mosaics have been proposed a long time ago, to our knowledge no efficient implementations have existed until now. We develop and implement a simple but efficient algorithm for computing order- k Delaunay mosaics, and have made it open-source [69; 70]. The algorithm

is very combinatorial in flavor. The only geometric operations are computing weighted Delaunay triangulations and means of points, which is made robust through the use of the Computational Geometry Algorithms Library, CGAL [81], and its exact arithmetics. Apart from performing k closest point searches and computing order- k α -shapes and multi-cover persistence (see below), another potential application of our algorithm is the computation of the distance-to-measure introduced in [20], which can be approximated using order- k Voronoi tessellations [44].

We use this opportunity to experimentally explore the complexity of order- k Delaunay mosaics. While tight worst-case bounds are known for the size of the first k Delaunay mosaics taken together [21], the size of individual order- k Delaunay mosaics is less understood. We observe that, like for Delaunay triangulation, for many practical point sets these worst-case bounds are not achieved, and relate the size of the order- k Delaunay mosaics to known input-specific bounds for order-1 Delaunay triangulations. We also investigate other combinatorial properties and non-asymptotic behavior of the complexity of the mosaics when k is relatively large compared to the input size.

Our algorithm uses a new geometric structure, which we call the *rhomboid tiling*, to obtain the order- k Delaunay mosaics. To compute persistence in scale, we introduce a radius function on the order- k Delaunay mosaics to give us the aforementioned subcomplexes as sublevel sets of this function. We give a recipe for computing this function, and with this persistence can be computed using standard techniques.

To compute persistence in depth, several innovative adaptations of the standard approach to persistence are needed. The main challenge is the combinatorial difference of the Delaunay mosaics from one value of k to the next. Here we introduce a radius function on the rhomboid tiling, and use the rhomboid tiling to get a zig-zag filtration that links consecutive order- k Delaunay mosaics and whose persistence diagram is the same as that of the filtration of multi-covers in depth.

Recent results [24] use a nerve construction to get a filtration in two parameters, k and the radius r , representing the multi-covers to get their persistent homology in both parameters simultaneously. In light of these results, we establish a link between the nerve construction and the rhomboid tiling, suggesting that the rhomboid tiling does in fact admit a similar 2-parameter filtration, without the need for a zig-zag filtration to compute persistence in depth. Our work is also related to the study of multi-covers based on Čech

complexes in [78]. While the relation between the different Čech complexes is simpler than that between the Delaunay mosaics, their explosive growth for increasing radius leads to algorithms with prohibitively long running time.

Finally, we investigate the radius function on the rhomboid tiling and the order- k Delaunay mosaics from a perspective of discrete Morse theory. We show that the level sets of the rhomboid tiling are intervals, and we use this insight to optimize the computation of the radius function on the rhomboid tiling, and by proxy also the order- k Delaunay mosaics. While the structure of the level sets on the order- k Delaunay mosaics is more complicated, we show that we can classify them into critical and non-critical steps with predictable impact on the topology.

Applications and extensions We apply these new tools and implementations to analyze sphere packings. Our main aim is to get a better understanding of why the FCC structure seems to occur preferredly over HCP in experimental sphere packings. We find that the FCC and the HCP configurations have the same persistence diagram for $k = 1, 2, 3$ but different persistence diagrams already for $k = 4$. With this insight we look at the persistence diagrams of the 4-fold cover in scale for sphere packings from physical experiments and molecular simulations. We define a measure to quantify the occurrence of FCC and HCP patterns, and apply it to packings of different densities. Unlike commonly used measures in the field, it has the aforementioned stability property inherited from persistence diagrams, i.e. small perturbations in the input can only cause small changes in the measure. Investigating time series data of FCC and HCP packings under external forces, we also observe from the persistence diagrams that FCC seems to be more robust to such forces.

Motivated by the applications of persistent homology to periodic structures, such as the regular sphere packings we just discussed or crystallographic materials like zeolites, we develop and implement an algorithm for Delaunay triangulations of periodic point sets to be integrated into the CGAL library. This extends the functionality of the library that is currently limited to periodic point sets with periodicity of the translational group generated by the standard basis vectors, whose fundamental domain is a cube. This algorithm combines ideas from both the existing implementation and classic results on periodic Delaunay triangulations. We discuss potential further extensions to order- k

Delaunay mosaics and how such periodic triangulations and mosaics lend themselves to the computation of persistent homology.

Outline We will first cover some background on order- k Delaunay mosaics and Voronoi tessellations (Chapter 2). In Chapter 3 we will then introduce the rhomboid tiling, and use it to shed some light on the combinatorics of order- k Delaunay mosaics. With this insight we present our algorithm to compute order- k Delaunay mosaics and the experimental investigation of their properties. Chapter 4 first covers background on the topological notions of homotopy and homology as well as persistence. We then present algorithms on how to compute the radius functions on the rhomboid tiling and the order- k Delaunay mosaics, followed by algorithms for persistence of multi-covers in scale and in depth. We then investigate these radius functions for their properties related to discrete Morse theory (Chapter 5). We apply our algorithm for persistence of multi-covers to analyze experimental sphere packings (Chapter 6), and close with Chapter 7 covering lattices, periodic point sets and our algorithm for computing periodic Delaunay triangulations.

2 Background

Order- k Voronoi tessellations are a generalization of the well known Voronoi tessellations [88]. Similar to the ordinary (order-1) Voronoi tessellation, the order- k Voronoi tessellation has a natural dual [7], the order- k Delaunay mosaic. Both will play a central role in this thesis, and we formally introduce them in this chapter, alongside related geometric structures and results. For a more in-depth exposition we refer to [31]. Because different flavors of cell complexes exist in the literature, we will also formally define the notions of cell complexes as we use them. Topological concepts will be introduced in later chapters when the need arises.

2.1 Cell complexes

A *hyperplane* h is a $(d - 1)$ -dimensional affine subspace of \mathbb{R}^d . It divides \mathbb{R}^d into two closed halves, which we refer to as *half-spaces*. A *polyhedron* is an intersection of a finite number of hyperplanes and half-spaces, and is thus by definition convex.

A *polyhedral complex* is a collection K of a finite number of distinct polyhedra, which we refer to as *cells*, that fulfills the following properties:

- the intersection of any two cells of K is another cell of K ,
- each point in $\bigcup_{\sigma \in K} \sigma$ is in the (relative) interior of a unique cell of K .

For a cell σ we denote its dimension as $\dim \sigma$. We call its e -dimensional cells *e-cells*, and we call 0-cells *vertices*. For a cell $\sigma \in K$, we call the cells $\tau \subseteq \sigma$ its *faces* and the cells $\tau \supseteq \sigma$ its *cofaces*. We call them *proper* faces and cofaces if $\sigma \neq \tau$, and we call them *facets*

and *cofacets* if $|\dim \sigma - \dim \tau| = 1$. A cell is a *maximal cell* if it does not have any proper cofaces.

A *polytope* σ is a convex hull of some finite set of points P that is in convex position. A polytope is thus convex, and we call P its *vertices*. We also say that σ is *spanned by* P . Any polytope is a polyhedron, and we call a polyhedral complex whose cells are polytopes a *regular complex*. We call a polytope σ whose number of vertices is $\dim \sigma + 1$ a *simplex*. A *geometric simplicial complex* K is a regular complex whose cells are simplices and further for each $\sigma = \text{conv } Q \in K$ also $\text{conv } Q'$ for $Q' \subseteq Q$ is in K .

Two polyhedral complexes K and K^* in \mathbb{R}^d are *dual* if there is a bijection between their cells such that for two dual cells σ and σ^* , $\dim \sigma + \dim \sigma^* = d$, and τ is a facet (or cofacet) of σ if and only if its dual cell τ^* is a cofacet (or facet) of σ^* , respectively.

2.2 Voronoi tessellations

Given a finite set of points $P \subseteq \mathbb{R}^d$, the *Voronoi domain* of $Q \subseteq P$ is $\text{dom}(Q) := \{x \in \mathbb{R}^d : \|x - q\| \leq \|x - p\|, \forall q \in Q, \forall p \in P \setminus Q\}$. Its *order* is $k = \#Q$, where $\#Q$ denotes the cardinality of Q . It is the subset of \mathbb{R}^d for which Q are the k closest points of P . For each positive integer k , the collection of Voronoi domains of order k , together with their non-empty intersections $\{\sigma_Q := \bigcap_{Q \in \mathcal{Q}} \text{dom}(Q), \sigma_Q \neq \emptyset \text{ where } \mathcal{Q} \subseteq \{Q \subseteq P : \#Q = k\}\}$ form a polyhedral complex which we refer to as the *order- k Voronoi tessellation*, denoted as $\text{Vor}_k(P)$, see Figure 2.1. The more commonly known Voronoi tessellations from [88] are the special case for $k = 1$.

The perpendicular bisector $b_{p,q}$ between two points p and q is the hyperplane that separates \mathbb{R}^d into the half-space $b_{p,q}^-$ of points closer to p , i.e. $b_{p,q}^- := \{x \in \mathbb{R}^d : \|x - p\| \leq \|x - q\|\}$, and the half-space $b_{p,q}^+$ of points closer to q . This gives us a different point of view on $\text{dom}(Q)$ as the intersection of all half-spaces $b_{q,p}^-$ with $q \in Q$ and $p \in P \setminus Q$. Similarly, the lower-dimensional cells of $\text{Vor}_k(P)$ are intersections of such half-spaces and bisectors, showing that $\text{Vor}_k(P)$ is indeed a polyhedral complex.

For completeness, we also introduce degree- k Voronoi tessellations [37]. For a point $p \in P$, the *k -th Brillouin zone* is the set of points $x \in \mathbb{R}^d$ for which p is a k -th closest point. To make this precise, a point p is a k -th closest point from x if there are at most $k - 1$ points $q \in P$ with $\|x - q\| < \|x - p\|$, but at least k points $q \in P$ with $\|x - q\| \leq \|x - p\|$.




Figure 2.1: The order-2 (solid lines) and order-1 (dotted lines) Voronoi tessellations of the point set P . The degree-2 Voronoi tessellation, $\text{Vor}_{1.5}(P)$, is the superposition of $\text{Vor}_1(P)$ and $\text{Vor}_2(P)$. In pink, the second Brillouin zone of the red point.

For $k = 1$ this coincides with the Voronoi domain of p , however for $k \geq 2$ the k -th Brillouin zones are not (convex) polyhedra themselves. Rather, they are composed of smaller polyhedra each of which corresponds to a unique set of $k - 1$ points of P that are strictly closer to the interior of the polyhedron than p . The collection of these polyhedra, for all points p , form the maximal cells of the *degree- k Voronoi tessellation*, which we denote as $\text{Vor}_{k-\frac{1}{2}}(P)$ for reasons to be seen later. Figure 2.1 illustrates these concepts. It refines the order- k Voronoi tessellation by decomposing its domains into these maximal cells. Similarly, the degree- k tessellation refines the order- $(k - 1)$ tessellation, and indeed $\text{Vor}_{k-\frac{1}{2}}(P)$ is the superposition of $\text{Vor}_k(P)$ and $\text{Vor}_{k-1}(P)$; see Figure 4.5b.

2.3 Hyperplane arrangements

Let $\mathcal{P}: \mathbb{R}^d \rightarrow \mathbb{R}$ be the function defined as $\mathcal{P}(x) = \frac{1}{2}\|x\|^2$. Its graph is a paraboloid which we also denote as \mathcal{P} in a slight abuse of notation. Given a finite set of points $P \subseteq \mathbb{R}^d$, for each point $p \in P$ we write $h_p: \mathbb{R}^d \rightarrow \mathbb{R}$ for the affine map defined by $h_p(x) = \langle x, p \rangle - \frac{1}{2}\|p\|^2 = \frac{1}{2}(\|x\|^2 - \|x - p\|^2)$. The graph of h_p is the hyperplane in \mathbb{R}^{d+1} that is tangent to \mathcal{P} at the point $(p, \mathcal{P}(p)) \in \mathbb{R}^{d+1}$, and we similarly denote it as h_p . We use h_p^+ to denote the half-space of points on or above h_p , and h_p^- for the half-space of points on or below h_p .

The collection of these hyperplanes h_p for $p \in P$ decomposes \mathbb{R}^{d+1} into convex cells,




Figure 2.2: The arrangement $\text{Arr}(P)$ for $P = \{a, b, c, d, e\} \subset \mathbb{R}^1$. The topmost chamber contains the paraboloid \mathcal{P} . The chamber below h_b and h_c (pink) projects to the domain $\text{dom}(\{b, c\})$ (red).

which we call the *hyperplane arrangement* of P , denoted $\text{Arr}(P)$; see Figure 2.2. The *cells* in the arrangement are intersections of these hyperplanes and half-spaces, and their interiors partition \mathbb{R}^{d+1} , justifying that we call it a polyhedral complex. More formally, a cell γ in the arrangement is uniquely described by an ordered three-partition $(P_{\text{above}}(\gamma), P_{\text{contains}}(\gamma), P_{\text{below}}(\gamma))$ with $P = P_{\text{above}}(\gamma) \sqcup P_{\text{contains}}(\gamma) \sqcup P_{\text{below}}(\gamma)$, and γ is the common intersection of the hyperplanes h_p^- for $p \in P_{\text{above}}(\gamma)$, the hyperplanes h_p for $p \in P_{\text{contains}}(\gamma)$ and the hyperplanes h_p^+ for $p \in P_{\text{below}}(\gamma)$. If for each point $y = (x, z) \in \mathbb{R}^d \times \mathbb{R}$ we define

$$P_{\text{above}}(y) := \{p \in P : z < h_p(x)\}, \quad (2.1)$$

$$P_{\text{contains}}(y) := \{p \in P : z = h_p(x)\}, \quad (2.2)$$

$$P_{\text{below}}(y) := \{p \in P : z > h_p(x)\}, \quad (2.3)$$

i.e. the set of points $p \in P$ whose hyperplanes h_p are above, contain or are below the point y respectively, then the interior of γ consists of the points y that satisfy $P_{\text{above}}(y) = P_{\text{above}}(\gamma)$, $P_{\text{contains}}(y) = P_{\text{contains}}(\gamma)$ and $P_{\text{below}}(y) = P_{\text{below}}(\gamma)$. Note that not all three-partitions of P correspond to arrangement cells.

We say a point set $P \subset \mathbb{R}^d$ is in *general position* if for $1 \leq e \leq d - 1$, no $e + 2$ points

of P lie on the same $(e - 1)$ -sphere. If P is in general position, the dimension of any cell γ is $e = d + 1 - \#P_{\text{contains}}(\gamma)$. We call $(d + 1)$ -dimensional cells of $\text{Arr}(P)$ *chambers*; they satisfy $P_{\text{contains}}(\gamma) = \emptyset$. The *depth* of a chamber γ is $\#P_{\text{above}}(\gamma)$ or, equivalently, the number of hyperplanes that are above this chamber.

Let $\text{proj}: \mathbb{R}^{d+1} \rightarrow \mathbb{R}^d$ be the map that projects a point $(x, z) \in \mathbb{R}^d \times \mathbb{R}$ to x . Then the following relationship between $\text{Arr}(P)$ and $\text{Vor}_k(P)$ holds [37].

Theorem 2.1 (Projections of chambers [37]). *In the arrangement $\text{Arr}(P)$ of a finite point set $P \subseteq \mathbb{R}^d$, the chamber γ defined by the three-partition $(Q, \emptyset, P \setminus Q)$ projects to $\text{proj}(\gamma) = \text{dom}(Q)$, the Voronoi domain of Q .*

In particular, the domains of $\text{Vor}_k(P)$ are the projections of the chambers of $\text{Arr}(P)$ at depth k .

Figure 2.2 shows an example. To provide some intuition for this result, recall that the Voronoi domain of some $Q \subseteq P$ is the intersection of half-spaces $b_{q,p}^-$ with $q \in Q$ and $p \in P \setminus Q$. Now the projection of the intersection of h_q and h_p is the bisector $b_{q,p}$, and the intersection of h_q^- and h_p^+ projects to $b_{q,p}^-$. As the chamber defined by the three-partition $(Q, \emptyset, P \setminus Q)$ is the intersection of h_q^- for $q \in Q$ and h_p^+ for $p \in P \setminus Q$, we see that its projection is $\text{dom}(Q)$.

A similar result holds for degree- k Voronoi tessellations: Each domain of the degree- k Voronoi tessellation is the projection of the intersection of a depth- k and a depth- $(k - 1)$ chamber. More precisely, the k -th level of $\text{Arr}(P)$ consists of those cells γ whose ordered three-partition $(P_{\text{above}}(\gamma), P_{\text{contains}}(\gamma), P_{\text{below}}(\gamma))$ satisfies $\#P_{\text{above}}(\gamma) \leq k - 1$ and $\#P_{\text{above}}(\gamma) + \#P_{\text{contains}}(\gamma) \geq k$. Then the cells of $\text{Vor}_{k-\frac{1}{2}}(P)$ are the vertical projections of the cells of the k -th level of $\text{Arr}(P)$.

2.4 Delaunay mosaics

A $(d - 1)$ -sphere is called an *empty sphere* if it contains no points from P in the interior of the ball bounded by the sphere (which to simplify language we also refer to as “interior of the sphere”). Assuming a point set P in general position, the (first-order) *Delaunay triangulation* of P is the geometric simplicial complex with vertex set P whose d -simplices are spanned by those sets of $d + 1$ vertices whose circumspheres are empty spheres. It was

shown in [37] that it is the projection of the boundary of special convex polytope in \mathbb{R}^{d+1} . This idea was generalized in [8] to yield a generalization of Delaunay triangulations, and we now outline this generalization.

Let $P \subseteq \mathbb{R}^d$ be a finite point set not necessarily in general position. We define the *lift* of $x \in \mathbb{R}^d$ onto the paraboloid \mathcal{P} as the point $\text{lift}(x) = (x, \mathcal{P}(x)) \in \mathbb{R}^{d+1}$. For each k -tuple $Q \subseteq P$, we take the sum of their lifts, $\sum_{q \in Q} \text{lift}(q)$, and define the *order- k Delaunay mosaic*, denoted $\text{Del}_k(P)$, as the vertical projection of the lower faces of the convex hull $\text{conv} \{ \sum_{q \in Q} \text{lift}(q) : Q \subseteq P, \#Q = k \}$ of these sums. Instead of sums, it is common to use the barycenters $\frac{1}{k} \sum_{q \in Q} \text{lift}(q)$, which yields the same complex scaled by a factor of $\frac{1}{k}$, and we will use this version in most of our figures. Figure 2.3 shows an example for $k = 2$. Independent of the choice of geometric embedding, for each cell $\sigma \in \text{Del}_k(P)$ we define the




Figure 2.3: Superposition of the order-2 Voronoi tessellation (in *black*) and the order-2 Delaunay mosaic (in *blue*) of a set of six points in the plane. Each domain of the tessellation corresponds to two of these six points, and the corresponding vertex of the mosaic is the average of these two points.

combinatorial vertex set $V(\sigma)$ as the collection of those sets Q whose sums (or barycenters) are geometric vertices of σ . We let $V(\text{Del}_k(P))$ denote the combinatorial vertex set of the whole complex, i.e. the union of $V(\sigma)$ for all cells $\sigma \in \text{Del}_k(P)$. The order- k Delaunay mosaic is a regular complex.

If P is in general position, the order-1 Delaunay mosaic coincides with the Delaunay triangulation. To see this, consider a facet σ of the lower convex hull of the lifts of P . Its vertex set is $\{\text{lift}(q) : q \in Q\}$ for some $Q \subseteq P, \#Q = d + 1$. Its supporting hyperplane h , i.e. the affine hull of its vertices, intersects the paraboloid \mathcal{P} in an ellipsoid. The projection of this ellipsoid is the circumsphere S of Q . Points on \mathcal{P} above h project to points outside

S while points on \mathcal{P} below h project to points inside S . However, as σ is a lower convex hull facet, none of the lifts of P are below h , and thus S is an empty sphere.

As we will see in Section 3.1, in dimension $d \geq 3$, the order- k Delaunay mosaic for $k \geq 2$ is not necessarily simplicial even if the points are in general position. The following relationship holds between the order- k Delaunay mosaic and the order- k Voronoi tessellation [8], also see Figure 2.3:

Theorem 2.2 (Duality [8]). $\text{Del}_k(P)$ is dual to $\text{Vor}_k(P)$.

2.5 Weighted notions

The notions of order- k Voronoi tessellation, hyperplane arrangement, order- k Delaunay mosaics and the results relating them to each other can be generalized to *weighted point sets*. Let P be a finite set of points \mathbb{R}^d , and $\text{wt}: P \rightarrow \mathbb{R}$ a function assigning real weights to the points.

The *power distance* of a point $x \in \mathbb{R}^d$ from a weighted point $p \in P$ with weight $\text{wt}(p)$ is $\text{dist}_p(x) = \|x - p\|^2 - \text{wt}(p)$. Then for $Q \subseteq P$, the corresponding *weighted Voronoi domain* is the set of points x that satisfy $\text{dist}_q(x) \leq \text{dist}_p(x)$ for all $q \in Q$ and $p \in P \setminus Q$. For each non-negative integer k , the *weighted order- k Voronoi tessellation*, is the collection of Voronoi domains for sets Q that satisfy $\#Q = k$. Setting $\text{wt}(p) = 0$ for all points $p \in P$, we get the unweighted situation as a special case. Note that while the perpendicular bisector between two weighted points might contain both points on the same side, it is still a hyperplane orthogonal to the straight line connecting the two points. This implies that the weighted order- k Voronoi tessellation is still a polyhedral complex, but also means that unlike in the unweighted case, some domains of the order-1 Voronoi tessellation may be empty.

To define the analog of the hyperplane arrangement for weighted points, let $h_p^{\text{wt}}: \mathbb{R}^d \rightarrow \mathbb{R}$ be defined by mapping $x \in \mathbb{R}^d$ to $h_p^{\text{wt}}(x) = h_p(x) + \frac{1}{2}\text{wt}(p)$. The graph of h_p^{wt} is parallel to h_p , the hyperplane tangent to \mathcal{P} , but it is shifted up by $\frac{1}{2}\text{wt}(p)$. The collection of these hyperplanes defines a hyperplane arrangement as before, and for these weighted definitions Theorem 2.1 holds as before. One difference that will be of importance later is that while all hyperplanes are incident to the top-most chamber in the unweighted arrangement, this is not necessarily the case in the arrangement for weighted points.

Order- k Delaunay mosaics similarly generalize to weighted point sets by merely adjusting the definition of the lift to $\text{lift}^{\text{wt}}(p) = (p, \mathcal{P}(p) - \frac{1}{2}\text{wt}(p))$, which is $\text{lift}(p)$ moved down by $\text{wt}(p)$. With this definition, the weighted order- k Delaunay mosaic is the projection of the lower convex hull of the barycenters of the lifts of P , and the duality from Theorem 2.2 still holds.

Also note that the unweighted order- k Delaunay mosaic for a point set P can be realized as a weighted first-order Delaunay mosaic for an appropriately chosen weighted point set P' . Choosing for each $Q \subseteq P$ with $\#Q = k$ the barycenter of Q as a new point $p' \in P'$, and setting its weight to $\text{wt}(p') = \|\frac{1}{k} \sum_{q \in Q} q\|^2 - \frac{1}{k} \sum_{q \in Q} \|q\|^2$, then $\text{lift}^{\text{wt}}(p')$ is the same as $\frac{1}{k} \sum_{q \in Q} \text{lift}(q)$, the barycenter of the lifts of Q . Thus the weighted first-order Delaunay mosaic of P' is the same as $\text{Del}_k(P)$.

We remark that while in the unweighted first-order Delaunay mosaic each point of P is a vertex of the mosaic, this is not the case anymore in the weighted setting. Intuitively, if a point $p \in P$ is surrounded by other points of P with much higher weight, then $\text{dom}(\{p\})$ can be empty, and dually p will not appear in $\text{Del}_1(P)$.

In order to generalize the notion of circumspheres and empty spheres that define cells in the Delaunay triangulations, we need to define what it means for a weighted point to be inside, on or outside a sphere. Let now $S = S_{r^2}(x)$ be the $(d-1)$ -dimensional *sphere* with center $x \in \mathbb{R}^d$ and squared radius $r^2 \in \mathbb{R}$. For $r^2 > 0$ this is an ordinary sphere, for $r^2 = 0$ it is a point, and for $r^2 < 0$ it is what we call an *imaginary sphere*. A weighted point p then is considered to be inside the sphere if $\text{dist}_p(x) < r^2$, on the sphere if $\text{dist}_p(x) = r^2$, and outside the sphere if $\text{dist}_p(x) > r^2$. The condition for a point to be on the sphere can be rewritten as $\|x - p\|^2 = r^2 + \text{wt}(p)$, which we geometrically interpret as having two spheres, $S_{r^2}(x)$ and $S_{\text{wt}(p)}(p)$, that intersect at a right angle. Then an empty sphere is a sphere that does not contain any weighted points of P inside it, and a circumsphere of $d+1$ weighted points is the sphere that has these points on it. Weighted circumspheres and their circumcenters are also called *orthospheres* and *orthocenters*, respectively. As for unweighted point sets, a weighted point set P is in *general position* if for $1 \leq e \leq d-1$, no $e+2$ points of P lie on the same $(e-1)$ -sphere.

In Chapters 3, 4 and 5 we will formulate all our statements in terms of unweighted points. However most of our definitions and results generalize to weighted point sets by substituting the above definitions for weighted points, and we will explicitly mention if a

result from these chapters does not.

2.6 Algorithms and complexity

In the plane, the number of domains in the order- k Voronoi tessellation or, equivalently, the number of vertices in the order- k Delaunay mosaic is $\Theta(k(n - k))$ where n is the number of input points; see [58; 77]. For dimensions $d \geq 3$, this number can vary significantly depending on the way the input points are distributed. The upper bound of $O(n^{\lfloor \frac{d+1}{2} \rfloor} k^{\lceil \frac{d+1}{2} \rceil})$ on the total size of the first k higher-order Delaunay mosaics is tight [21], while the lower bound of $\Omega(k^d n)$ is only conjectured [66]. For individual order- k Delaunay mosaics, the complexity is poorly understood. The problem is closely related to the $(d + 1)$ -dimensional k -set problem. Specifically, the points in $P \subseteq \mathbb{R}^d$ can be mapped to equally many points in \mathbb{R}^{d+1} such that the order- k domains in \mathbb{R}^d correspond to k -sets in \mathbb{R}^{d+1} , see e.g. [21].

The first algorithm to compute order- k Voronoi tessellations and Delaunay mosaics in the plane was described by Lee in [58]. The algorithm computes the Voronoi tessellations one by one, in increasing order and in time $O(k^2 n \log n)$. Mulmuley [66] extended this algorithm beyond two dimensions, computing the first k levels in the $(d + 1)$ -dimensional hyperplane arrangement, which implicitly yields the order- k Voronoi tessellations and Delaunay mosaics in time $O(s \log n + k^d n^2)$, in which s denotes the output size. Mulmuley [67] later described another algorithm, which instead adds hyperplanes one by one, and runs in time $O(n^{\lfloor \frac{d+1}{2} \rfloor} k^{\lceil \frac{d+1}{2} \rceil})$ for $d \geq 3$, which equals the worst-case output size. For $d = 2$, the expected runtime is $O(k^2 n \log \frac{n}{k})$. Another incremental algorithm with similar complexity for $d \geq 3$ has been described by Agarwal *et al.* [1].

To our knowledge, none of these algorithms have been implemented.

3 Computing higher-order Delaunay mosaics

In this chapter, we describe a new algorithm for computing order- k Delaunay mosaics in Euclidean space of dimension d that stands out in its simplicity. It selects the vertices of the order- k mosaic from incrementally constructed lower-order mosaics, using a new geometric construction in \mathbb{R}^{d+1} , the *rhomboid tiling*, that the order- k Delaunay mosaics are horizontal slices of. To construct the order- k mosaic from its vertices, it uses an algorithm for weighted first-order Delaunay mosaics as a black-box. Beyond this black-box, all operations of the algorithm are purely combinatorial, which facilitates easy implementation. It thus benefits from highly optimized existing implementations for weighted Delaunay triangulations and, if desired, can build upon their use of exact arithmetic. Its complexity depends on the complexity of the algorithm used for weighted Delaunay mosaics. Assuming it is linear in its output size, then the complexity of our algorithm is also linear in its output size. We implement this algorithm and run it on various point sets, shedding light on the size and other properties of order- k Delaunay mosaics. In particular, we compare the total size of the first k Delaunay mosaics of random point sets with the (tight) worst-case upper bound, and we study the size of individual order- k Delaunay mosaics, for which no tight bounds are known in general. As far as we are aware, no such experimental investigations have been performed in the past, possibly due to absence of a practical algorithm. We make open-source implementations of our algorithm available [69; 70].

We will introduce the rhomboid tiling in Section 3.1. We will explore the combinatorial properties of this tiling and, by proxy, the properties of order- k Delaunay mosaics in Section 3.2. Using these results, we explain our algorithm in Section 3.3. We present

experimental results obtained with the two implementations of this algorithm in Section 3.4. We close this chapter with a discussion of possible extensions and optimizations in Section 3.5.

3.1 Rhomboid Tiling

Given a finite set of points in \mathbb{R}^d , we are interested in the collection of Delaunay mosaics of all orders. Assuming the set is in general position, there exists a rhomboid tiling in \mathbb{R}^{d+1} such that the Delaunay mosaics are horizontal slices of the tiling. This section introduces the tiling and proves the relation to Delaunay mosaics.

Rhomboid tiling. Let $P \subseteq \mathbb{R}^d$ be finite and in general position. Every $(d-1)$ -dimensional sphere, S , in \mathbb{R}^d partitions P into the points *inside*, *on*, and *outside* S . We call this the *ordered three-partition* of P defined by S , and denote it as $P = \text{In}(S) \sqcup \text{On}(S) \sqcup \text{Out}(S)$. By assumption of general position, we have $0 \leq \#\text{On}(S) \leq d+1$, but there are no a priori upper bounds on the sizes of the other two sets.

We map each ordered three-partition defined by a $(d-1)$ -sphere, S , to a parallelepiped in \mathbb{R}^{d+1} , which we call the *rhomboid* of S , denoted $\text{rho}(S)$. To define it, we write $y_p = (p, -1) \in \mathbb{R}^{d+1}$, for every $p \in P$, and $y_Q = \sum_{q \in Q} y_q$ for every $Q \subseteq P$. The $(d+1)$ -st coordinate of y_Q is therefore $-\#Q$, and we call $\#Q$ the *depth* of the point. With this notation, $\text{rho}(S) = \text{conv}\{y_Q : \text{In}(S) \subseteq Q \subseteq \text{In}(S) \cup \text{On}(S)\}$. Equivalently, $\text{rho}(S)$ is the rhomboid spanned by the vectors y_p , with $p \in \text{On}(S)$, and translated along $y_{\text{In}(S)}$. Its dimension is the number of spanning vectors, $\#\text{On}(S)$.

While in general many spheres S yield the same rhomboid $\text{rho}(S)$, by definition a rhomboid is uniquely determined by the three-partition $\text{In}(S) \sqcup \text{On}(S) \sqcup \text{Out}(S)$, and for a rhomboid $\rho = \text{rho}(S)$ we define $P_{\text{in}}(\rho) := \text{In}(S)$, $P_{\text{on}}(\rho) := \text{On}(S)$ and $P_{\text{out}}(\rho) := \text{Out}(S)$. Combinatorially, each geometric vertex y_Q of ρ corresponds to a subset Q of P , and we define the *combinatorial vertex set* of ρ as

$$V(\rho) = \{P_{\text{in}}(\rho) \subseteq Q \subseteq P_{\text{in}}(\rho) \cup P_{\text{on}}(\rho)\}. \quad (3.1)$$

We refer to $P_{\text{in}}(\rho) \in V(\rho)$ as the (combinatorial) *anchor vertex* of ρ .

Observe that every face of $\text{rho}(S)$ is again the rhomboid defined by a sphere. To see this, we note that for every ordered partition of the points on S into three sets, $\text{On}(S) = O_{in} \sqcup O_{on} \sqcup O_{out}$, there is a sphere S' with $\text{In}(S') = \text{In}(S) \cup O_{in}$, $\text{On}(S') = O_{on}$, and $\text{Out}(S') = \text{Out}(S) \cup O_{out}$. There are $3^{\#\text{On}(S)}$ such ordered partitions, and each corresponds to a face of $\text{rho}(S)$.


By definition, the *rhomboid tiling* of P , denoted $\text{Rho}(P)$, is the collection of all rhomboids defined by spheres; see Figure 3.1a. As suggested by the figure, the ordered three partition $(\emptyset, \emptyset, P)$ is mapped to the origin of \mathbb{R}^{d+1} . Let H_k be the hyperplane at depth k that is orthogonal to the $(d+1)$ -st axis, i.e. $H_k := \{(x, z) \in \mathbb{R}^d \times \mathbb{R} : z = -k\}$. We claim the following properties.

Theorem 3.1 (Rhomboid Tiling). *Let $P \subseteq \mathbb{R}^d$ be locally finite and in general position. Then*


1. $\text{Rho}(P)$ is dual to the hyperplane arrangement $\text{Arr}(P)$;
2. $\text{Rho}(P)$ is the projection of the boundary of a zonotope in \mathbb{R}^{d+2} ;
3. the horizontal slice of $\text{Rho}(P)$ at depth k is the order- k Delaunay mosaic of P , i.e. each cell $\sigma \in \text{Del}_k(P)$ is $\sigma = \rho \cap H_k$ for some $\rho \in \text{Rho}(P)$.

Figure 3.1 illustrates the duality of Claim 1 in Theorem 3.1. Note that Claim 2 implies that the rhomboid tiling is a *geometric realization* of the dual of the arrangement in \mathbb{R}^{d+1} , that is: its rhomboids intersect in common faces but not otherwise. This also means that $\text{Rho}(P)$ is a regular complex. The remainder of this section proves the three claims.

Proof of Claim 1: hyperplane arrangement. Recall from Section 2.3 that $\text{Arr}(P)$ is the decomposition of \mathbb{R}^{d+1} into cells by hyperplanes h_p for $p \in P$, and each cell γ is uniquely described by the three-partition $(P_{\text{above}}(\gamma), P_{\text{contains}}(\gamma), P_{\text{below}}(\gamma))$ of points p for which h_p is above, contains, and is below γ respectively. Importantly, there is a bijection between the cells of $\text{Arr}(P)$ and the rhomboids in $\text{Rho}(P)$. To see this, map a point (x, z) in the interior of a cell γ to the sphere S with center x and squared radius $r^2 = \max\{0, \|x\|^2 - 2z\}$. Using the definition of $h_p(x) = \frac{1}{2}(\|x\|^2 - \|x - p\|^2)$, we observe that $\text{In}(S) = P_{\text{above}}(\gamma)$, $\text{On}(S) = P_{\text{contains}}(\gamma)$, and $\text{Out}(S) = P_{\text{below}}(\gamma)$. We can reverse the map, and while this will not reach the points with $\|x\|^2 - 2z < 0$, these points all belong to



(a) The rhomboid tiling of the five points. The highlighted 2-rhomboid ρ defined by $P_{in}(\rho) = \{c\}$ and $P_{on}(\rho) = \{b, d\}$ is the convex hull of the points $y_c, y_{\{b,c\}}, y_{\{c,d\}},$ and $y_{\{b,c,d\}}$. We simplify the labels here and e.g. write bcd instead of $y_{\{b,c,d\}}$. The horizontal line at depth k intersects the tiling in the order- k Delaunay mosaic.



(b) The dual hyperplane arrangement $\text{Arr}(P)$. Following the dotted lines connecting the points of P on the horizontal axis to the paraboloid, \mathcal{P} , we find the corresponding tangent hyperplanes h_p .

Figure 3.1: Rhomboid tiling, $\text{Rho}(P)$, and hyperplane arrangement, $\text{Arr}(P)$, for a 1-dimensional point set, $P = \{a, b, c, d, e\}$. Notice how for each p , the vector $y_p = (p, -1)$ in $\text{Rho}(P)$ from the origin to the vertex labelled p is orthogonal to the hyperplane h_p in $\text{Arr}(P)$. For each $0 \leq j \leq 2$, the red j -dimensional rhomboid from $\text{Rho}(P)$ is dual to the red $(2 - j)$ -dimensional cell of $\text{Arr}(P)$.

the chamber of the ordered three-partition $(\emptyset, \emptyset, P)$. This establishes the bijection between the cells and the rhomboids. This bijection reverses dimensions and preserves incidences, which justifies that we call it a duality between the rhomboid tiling and the hyperplane arrangement. This completes the proof of Claim 1 in Theorem 3.1.

Proof of Claim 2: zonotope. A *zonotope* [31, Section 1.7] is a special convex polytope, namely one obtained by taking the Minkowski sum of finitely many line segments. The zonotope of interest is constructed from the line segments that connect the origin to the points $v_p = (p, -1, \|p\|^2/2) \in \mathbb{R}^{d+2}$, with $p \in P$. Note that these line segments project to the vectors $y_p = (p, -1)$ used to build the rhomboid tiling. By construction, y_p is normal to the graph of h_p , which is the zero set of $F_p: \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ defined by $F_p(y) = \langle y, y_p \rangle - \|p\|^2/2$; see Figure 3.1. Adding a $(d+2)$ -nd coordinate, w , we introduce $G_p: \mathbb{R}^{d+2} \rightarrow \mathbb{R}$ defined by $G_p(y, w) = \langle y, y_p \rangle + w\|p\|^2/2$. Its zero-set is normal to v_p , the restriction of $G_p^{-1}(0)$ to $w = -1$ is the zero-set of F_p , and $G_p(0) = 0$. In other words, if we identify \mathbb{R}^{d+1} with the hyperplane $w = -1$ in \mathbb{R}^{d+2} , then the zero-sets of the G_p intersect \mathbb{R}^{d+1} in $\text{Arr}(P)$ and they all pass through the origin in \mathbb{R}^{d+2} .

By construction, the thus defined zonotope is dual to the arrangement of hyperplanes $G_p^{-1}(0)$ for $p \in P$. Therefore, the antipodal face pairs of the zonotope correspond dually to the cells of $\text{Arr}(P)$, provided we interpret the arrangement projectively, which means we combine antipodal pairs of unbounded cells; see also [31, Section 1.7]. We get a more direct dual correspondence by projecting the bottom side of the boundary of the zonotope to \mathbb{R}^{d+1} . By choice of the line segments, the vertices on this side project vertically to the vertices of $\text{Rho}(P)$, and since both are dual to $\text{Arr}(P)$, we conclude that $\text{Rho}(P)$ is the projection of this side of the zonotope. This completes the proof of Claim 2 in Theorem 3.1.

We remark that the relationship between the hyperplane arrangement defined by $G_p^{-1}(0)$ and the zonotope defined by the vectors v_p for $p \in P$ is a special instance of a more general idea by Coxeter introduced in the context of studying zonotopes [25]: For a given zonotope, Coxeter defines the so-called *second projective diagram* which is a hyperplane arrangement and has the property that it is dual to the zonotope.


Proof of Claim 3: Delaunay mosaics. By Theorem 2.1, for each Voronoi domain, there is a chamber in $\text{Arr}(P)$ that projects vertically to the domain. In particular, chambers at depth k project to domains of $\text{Vor}_k(P)$. By Claim 1, these chambers correspond to

the vertices of the rhomboid tiling at depth k . Since $\text{Rho}(P)$ is dual to $\text{Arr}(P)$, we get the dual of the Voronoi tessellation by taking the slice that is the intersection of H_k and $\text{Rho}(P)$. However by Theorem 2.2 the dual of the order- k Voronoi tessellation is precisely the order- k Delaunay mosaic. This completes the proof of Claim 3 in Theorem 3.1.


We see that the cells of $\text{Del}_k(P)$ are special slices of the rhomboids. More formally, for every $\sigma \in \text{Del}_k(P)$ there is a unique lowest-dimensional rhomboid $\rho \in \text{Rho}(P)$ such that $\sigma = \rho \cap H_k$. We denote this rhomboid as $\rho(\sigma)$, and call it the *rhomboid of σ* . Note that the qualification “lowest-dimensional” is required for uniqueness as vertices of $\text{Del}_k(P)$ are slices of both vertices and edges of $\text{Rho}(P)$. Equivalently, $\rho(\sigma)$ is the unique rhomboid whose interior intersected with H_k is the interior of σ . For vertices we have $\dim \sigma = \dim \rho = 0$, and for all higher-dimensional cells we have $\dim \sigma = \dim \rho - 1 \geq 1$.

Combinatorially, each rhomboid is a cube and, again combinatorially, each cell of $\text{Del}_k(P)$ is a slice orthogonal to the cube diagonal that passes through a non-empty set of the vertices. For the $(d+1)$ -cube, there are $d+2$ such slices, which we index from top to bottom by the *generation* $0 \leq g \leq d+1$. The g -th slice passes through $\binom{d+1}{g}$ vertices, so we have a vertex at generations $g = 0, d+1$, a d -simplex at generations $g = 1, d$, and some other d -dimensional polytope at generations $2 \leq g \leq d-1$. In $d+1 = 3$ dimensions, we have a vertex, a triangle, another triangle, and another vertex, see Figures 3.2a and 3.3; but already in $d+1 = 4$ dimensions, the middle slice is not a simplex; see Figure 3.2b. To describe these slices in general, let U_{d+1} be the $d+1$ unit coordinate vectors. The g -th slice is the convex hull of the points $\sum_{u \in Q} u$ with $Q \subseteq U_{d+1}, \#Q = g$, in which the empty sum is $(0, 0) \in \mathbb{R}^d \times \mathbb{R}$, by convention. To get an intuition, it might be easier to divide the sums by g , in which case the g -th slice is the convex hull of the barycenters of the $(g-1)$ -faces of the standard d -simplex.

Similar to the Delaunay mosaic, the *half-integer slice*, i.e. intersection of $\text{Rho}(P)$ with the hyperplane $H_{k-\frac{1}{2}}$ at depth $k - \frac{1}{2}$, is a regular complex in \mathbb{R}^d . We denote this complex as $\text{Del}_{k-\frac{1}{2}}(P)$. It is dual to the degree- k Voronoi tessellation, which motivates its earlier notation as $\text{Vor}_{k-\frac{1}{2}}(P)$.



(a) A 3-dimensional rhomboid ρ with $P_{in}(\rho) = \{p, q\}$ and $P_{on}(\rho) = \{a, b, c\}$. First-generation slice in red and second-generation slice in blue.



(b) Slices of a 4-dimensional rhomboid defined by $P_{in}(\rho) = \emptyset$ and $P_{on}(\rho) = \{a, b, c, d\}$. The non-trivial slices are a tetrahedron at generation $g = 1$, an octahedron at generation $g = 2$, and another tetrahedron at generation $g = 3$.

Figure 3.2: Combinatorics of 3- and 4-dimensional rhomboid slices. The first generation slice is always a simplex. The g -th generation slice is the convex hull of the barycenters of the $(g - 1)$ -faces of the first generation slice.

3.2 Combinatorial Properties

Recall that by our definition, the order- k Delaunay mosaic is the projection of the lower boundary cells of a convex polytope in \mathbb{R}^{d+1} . Specifically, for each subset $Q \subseteq P$ of k points, we take the barycenters (or sums) of the lifts of the points in Q , and the polytope in question is the convex hull of these barycenters. This directly gives us an recipe to compute order- k Delaunay mosaics, using an existing algorithm for convex hulls or, equivalently, weighted first-order Delaunay mosaics. However by itself, this approach does not scale well with k since there are $\binom{\#P}{k}$ such barycenters. Most barycenters, however, are irrelevant as they do not contribute to the lower faces of the convex hull. If we could, somehow, identify the relevant barycenters without wasting time on the irrelevant ones, this procedure would efficiently construct the cells of the order- k Delaunay mosaic by computing the weighted first-order Delaunay mosaic. We will see how this can be done in Section 3.2.2.

In $d \geq 3$ dimension, not all cells of $\text{Del}_k(P)$ are simplicial, even if the points in P are

in general position. The cells carry important information, which for some applications is essential and cannot be easily recovered from a triangulation. This poses an additional challenge because most algorithms for computing convex hulls or weighted first-order Delaunay mosaics return a triangulated version of the correct mosaic. As explained in the following section, we address this issue by predicting the cells from their corresponding rhomboids.

3.2.1 Predicting Cells

Given a cell σ in the order- k Delaunay mosaic, let $P_{in}(\sigma) := \bigcap V(\sigma)$, $P(\sigma) := \bigcup V(\sigma)$, $P_{on}(\sigma) := P(\sigma) \setminus P_{in}(\sigma)$ and $P_{out}(\sigma) := P \setminus P(\sigma)$. Then the following lemma identifies the rhomboid $\rho(\sigma)$ of σ , i.e. the rhomboid that σ is a slice of; see Figure 3.3 for an illustration. Clearly, $V(\sigma) \subseteq V(\rho)$.

Lemma 3.2. *Let $\sigma \in \text{Del}_k(P)$ be a e -dimensional cell, and $\rho = \rho(\sigma)$ be the rhomboid of σ . Then $P_{in}(\rho) = P_{in}(\sigma)$, $P_{on}(\rho) = P_{on}(\sigma)$, $P_{out}(\rho) = P_{out}(\sigma)$, and the generation of σ is $k - \#P_{in}(\sigma)$.*

Proof. Note that if σ is a vertex, ρ is a vertex and the statement is trivial. So assume σ is not a vertex. Then the dimension of ρ is $e + 1$, and the generation g of σ is between 1 and e . Recall that $V(\rho) = \{P_{in}(\rho) \subseteq Q \subseteq P_{in}(\rho) \cup P_{on}(\rho)\}$, in which $P_{in}(\rho)$ and $P_{on}(\rho)$ are disjoint. Since the depth of a vertex is determined by its cardinality, and the vertices of a slice are by definition all at the same depth, the vertices of the generation- g slice all satisfy $\#Q - \#P_{in}(\rho) = g$. Notice that the intersection of all g -subsets of $P_{on}(\rho)$ is empty, which implies that $P_{in}(\sigma)$, the intersection of the combinatorial vertices of the slice, is $P_{in}(\rho)$. Furthermore, $\bigcup V(\sigma) = \bigcup V(\rho)$ for every slice σ of ρ with generation $g \geq 1, g \leq e$. The union of all g -subsets of $P_{on}(\rho)$ is $P_{on}(\rho)$ itself, and thus $P_{on}(\sigma) = \bigcup V(\rho) \setminus P_{in}(\rho) = P_{on}(\rho)$. Finally, the generation of σ is the difference in depth of the anchor vertex, $P_{in}(\rho)$, and the slice defining σ . The depth of σ is k and the depth of $P_{in}(\rho)$ is its cardinality, which completes the proof. \square

Note that because all combinatorial vertices of a cell $\sigma \in \text{Del}_k(P)$ have the same cardinality, $\#P_{on}(\sigma) \neq 1$. While for 1-dimensional rhomboids ρ , i.e. edges, we can have $\#P_{on}(\rho) = 1$, this does not contradict our lemma because 1-dimensional rhomboids are not rhomboids of any cell σ .

We briefly discuss the implications for the dual order- k Voronoi tessellation. For a d -cell $\sigma \in \text{Del}_k(P)$, the dual Voronoi vertex is located at the circumcenter of $P_{on}(\sigma)$. As each $(d+1)$ -rhomboid ρ has d consecutive slices σ with the same $P_{on}(\sigma) = P_{on}(\rho)$, each order- k Voronoi vertex appears in d consecutive higher-order Voronoi tessellations (in the same location). More generally, each e -cell of $\text{Vor}_k(P)$ for $e \leq d-1$ appears in $d-e$ consecutive higher-order Voronoi tessellations.




Figure 3.3: First-, second-, and third-order Delaunay mosaics of the set $P = \{a, b, c, d, e\}$ in \mathbb{R}^2 as slices of the 3-dimensional rhomboid tiling. For clarity, only two of the rhomboids are shown, with their first-generation slices in *red* and second-generation slices in *dark blue*. The rhomboids on the *left* and *right* are defined by $P_{in} = \{b\}, P_{on} = \{a, c, d\}$ and $P_{in} = \emptyset, P_{on} = \{c, d, e\}$, respectively.

If all of our top-dimensional order- k Delaunay cells are triangulated—e.g. due to being the output of a weighted first-order Delaunay algorithm—we cannot directly apply Lemma 3.2. Assume τ is a simplex that is part of the triangulation of a cell σ , and $\rho = \rho(\sigma)$ is the rhomboid of σ . In general, we cannot obtain the information about ρ from τ alone, at least if $d \geq 4$. We can, however, still identify whether τ is a first-generation slice of ρ and thus in fact is equal to σ . Using Lemma 3.2, we can then obtain ρ .

Lemma 3.3. *Let $P \subset \mathbb{R}^d$ be a finite point set. A d -simplex, τ , in a triangulation of $\text{Del}_k(P)$*

is a first-generation d -cell of $\text{Del}_k(P)$ if and only if the intersection of its combinatorial vertices is of size $k - 1$.

Proof. Let σ be the d -cell in $\text{Del}_k(P)$ that contains τ in its triangulation, and assume σ is a generation- g slice of ρ . From Lemma 3.2, we know that $P_{in}(\rho) \subseteq v$ for all $v \in V(\sigma)$, and $\#P_{in}(\sigma) = k - g$. The remaining g points in every v are from $P_{on}(\rho)$. We have $V(\tau) \subseteq V(\sigma)$ with $\#V(\tau) = d + 1$. So for τ to consist of vertices whose common intersection is of size $k - 1$, there need to be $d + 1$ distinct g -subsets of $P_{on}(\rho)$ that all have $g - 1$ points in common. However, as $\#P_{on}(\rho) = d + 1$, this is not possible unless $g = 1$. \square

3.2.2 Identifying Vertices

Given a triangulation of the order- k Delaunay mosaic, we just saw how to identify its first-generation cells. From these, we can obtain the corresponding rhomboids and their higher-generation slices. We shall now prove that if we have triangulations of the order- j Delaunay mosaics, for all $j < k$, we can assemble the complete vertex set of the order- k Delaunay mosaic by taking slices at depth k obtained from first-generation cells at lower depths. We note that this only holds in the unweighted setting.

The proof of this result employs the hyperplane arrangement which the rhomboid tiling is dual to by Theorem 3.1. Recall that it was constructed from hyperplanes h_p which are tangent to the paraboloid \mathcal{P} for $p \in P$. These hyperplanes decompose \mathbb{R}^{d+1} into cells, the $(d + 1)$ -dimensional ones we call chambers.

We call a chamber γ a *bowl* if only one of its facets bounds it from above or, equivalently, if there is only one chamber at the next lower depth that shares a facet with γ . We call the hyperplane that contains this facet the *lid* of the bowl.

Lemma 3.4. *A hyperplane that is a lid of a bowl at depth 1 is not a lid of any other chambers.*

Proof. Let γ be a bowl at arbitrary depth, and let h be its lid. Every other hyperplane that contains a facet of γ bounds γ from below. The top facet of γ is the only part of h that is above all of these hyperplanes; that is: all other parts of h are below at least one of the other hyperplanes. This implies that every other bowl with lid h has at least one other hyperplane above it, and is thus of depth at least 2.

Now assume γ is at depth 1. If there were another bowl γ' with lid h , then the above argument would yield that all other bowls are at depth at least 2, contradicting our assumption on γ . Thus γ has to be the unique bowl with lid h . \square

With this lemma, we are ready to state and prove the main combinatorial insight that motivates our algorithm. In a nutshell, it says that the first-generation cells form *clusters* none of whose vertices are exclusive to the cluster, also see Figure 3.12. In \mathbb{R}^2 , this in particular implies that these clusters have outer-planar 1-skeletons.

Theorem 3.5. *Let $P \subseteq \mathbb{R}^d$ be a finite, unweighted point set and $k \geq 2$. Then every vertex in $\text{Del}_k(P)$ is vertex of some d -cell of generation $g \geq 2$.*

Proof. In the unweighted setting, each hyperplane is tangent to the paraboloid and contains a facet of the unique depth-0 chamber. Thus, each hyperplane is the lid to a chamber at depth 1. As this is true for every hyperplane, all chambers of depth 2 or higher have no lids by Lemma 3.4. This means that any chamber of depth at least 2 has at least two upper facets. Because the upper boundary is connected, there are two upper facets that meet in a $(d-1)$ -face, the dual rhomboid of this face has dimension 2, and its bottom vertex is dual to the chamber. Thus we can obtain this vertex, v , knowing the other three vertices of the 2-rhomboid.

Any 2-rhomboid is a face of some $(d+1)$ -dimensional rhomboid, ρ , which thus contains v at generation at least 2, i.e. v has depth at least $\#P_{in}(\rho) + 2$. Knowing $P_{in}(\rho)$ and $P_{on}(\rho)$, we obtain this vertex via Equation (3.1). \square

This result generally does not hold in the weighted setting. If not all hyperplanes of the arrangement are incident to the top chamber, then there can be bowls at depth greater than 1, whose dual vertices our algorithm does not find. Figure 3.4 shows an example of such a weighted configuration.

3.3 Algorithm for higher-order Delaunay mosaics

We outline our algorithm in this section; its correctness follows from the results of the previous sections. We compute the Delaunay mosaics one by one in sequence of increasing order. For $\text{Del}_1(P)$, the vertex set is the set P of input points. Whenever we have the




Figure 3.4: The hyperplane arrangement (left) and rhomboid tiling (right) of a weighted point set $P = \{a, b, c, d\}$. Bowls in $\text{Arr}(P)$ are marked in pink, and the dual vertices in $\text{Rho}(P)$ in red. Notice how we cannot obtain the red vertices as second-generation vertices of 2-rhomboids. In the order-2 Delaunay mosaic, the vertices ad , ac and ab as well as the vertices ab and bc both form clusters that have vertices exclusive to the cluster.

vertex set of $\text{Del}_j(P)$, we compute its (triangulated) d -cells using an off-the-shelf algorithm for weighted Delaunay triangulations. We use Lemma 3.3 to identify the first-generation d -cells, while discarding all other cells. From each first-generation cell, we obtain and save the higher-generation d -cells and vertices defined by the same rhomboid. These will appear in Delaunay mosaics of higher orders. By Theorem 3.5, once we have processed all $\text{Del}_j(P)$ for $j < k$, we will have obtained the complete vertex set of $\text{Del}_k(P)$ in the process, thus allowing our algorithm to continue until we have all Delaunay mosaics up to the desired order.

Algorithm 1 is a more formal write-up of the above outline, and Figure 3.3 visualizes the process. A dimension-agnostic `python` implementation and a 2- and 3-dimensional `C++` implementation using CGAL are available at [69; 70]. Note that from the first-generation d -cells σ of the order- k Delaunay mosaic we can reconstruct the $(d+1)$ -rhomboid $\rho = \rho(\sigma)$ that σ is a slice of by computing the anchor vertex $P_{in}(\rho) = P_{in}(\sigma)$ and $P_{on}(\rho) = P_{on}(\sigma)$. The implementation from [69] explicitly stores the vertices and d -cells of each order- k Delaunay mosaic. The implementation from [70] instead only stores these $(d+1)$ -rhomboids, which take up less space, and obtains the mosaics $\text{Del}_k(P)$ and $\text{Del}_{k-\frac{1}{2}}(P)$ via methods that reconstruct them on demand.

To get a handle on the runtime of the algorithm, we consider the two steps used to compute the order- k Delaunay mosaic after finishing the construction of the first $k-1$

Algorithm 1 computes the order- k Delaunay mosaic of a finite set of (unweighted) points, $P \subseteq \mathbb{R}^d$. We represent each d -cell of $\text{Del}_k(P)$ by the collection of its combinatorial vertices, stressing that these collections are sets and thus contain every combinatorial vertex only once. Duplicity is avoided by checking before adding. The locations of the combinatorial vertices are the barycenters of their points, and the cell is the convex hull of these locations.

While the software for computing the weighted Delaunay mosaic may return all cells triangulated, our algorithm outputs the (non-triangulated) cells of the order- k Delaunay mosaic. We recall that in $d \geq 3$ dimensions such non-simplicial cells appear generically for $k \geq 2$.

```

 $V(\text{Del}_1(P)) := P$ 
for  $j$  from 1 to  $k$  do
  // Compute the location and weight of each combinatorial vertex
  for all  $v \in V(\text{Del}_j(P))$  do
     $\text{loc}(v) := \frac{1}{j} \sum_{p \in v} p$ 
     $\text{wt}(v) := \left\| \frac{1}{j} \sum_{p \in v} p \right\|^2 - \frac{1}{j} \sum_{p \in v} \|p\|^2$ 
  end for
  // Get the (triangulated) cells of the order- $j$  Delaunay mosaic
   $D := \text{weightedDelaunay}(\text{loc}, \text{wt})$ 
  // Infer vertices and higher-generation cells of later Delaunay mosaics
  for all  $d$ -simplices  $\sigma$  in  $D$  do
    // Check whether the generation of  $\sigma$  is 1 via Lemma 3.3
    // We already obtained higher-generation cells of  $\text{Del}_j(P)$  earlier.
    if  $\#\cap V(\sigma) = j - 1$  then
      Add  $\sigma$  to  $\text{Del}_j(P)$ 
      // Get  $P_{in}(\rho)$  and  $P_{on}(\rho)$  via Lemma 3.2
       $P_{in}(\rho) := \cap V(\sigma)$ 
       $P_{on}(\rho) := \cup V(\sigma) \setminus P_{in}(\rho)$ 
      for  $g$  from 2 to  $d$  do
        // Get the generation- $g$  cell,  $\sigma'$ , of the rhomboid of  $\sigma$ , via Equation (3.1)
         $V(\sigma') := \{P_{in}(\rho) \cup Q : Q \in P_{on}(\rho), \#Q = g\}$ 
        Add all  $v \in V(\sigma')$  to  $V(\text{Del}_{j+g-1}(P))$ 
        Add  $\sigma'$  to  $\text{Del}_{j+g-1}(P)$ 
      end for
    end if
  end for
end for
return  $V(\text{Del}_k(P)), \text{Del}_k(P)$ 

```

mosaics. The first step is geometric and invokes the black-box algorithm to construct the weighted Delaunay mosaics from which we get vertices and cells of (unweighted) higher-order Delaunay mosaics. The runtime of this step depends on the runtime of the black box algorithm, which many cases is output-dependent. The second step is combinatorial and determines, for each output simplex from the first step, whether it is first generation, in which case it is a genuine cell of the mosaic. Assuming constant dimension, d , identifying whether an order- $(k - 1)$ cell is of first generation and, in this case, obtaining the higher-generation cells takes time $O(k)$. Thus, for a given k , the combinatorial step takes time $O(kC_k)$, in which C_k is the number of d -cells of $\text{Del}_k(P)$. With each vertex being represented as a k -tuple of points, this is linear in the output size, assuming we store each cell naively as a set of its vertices. If the runtime of each black-box invocation were linear in the output size, the total runtime for producing the first k higher-order Delaunay mosaics would thus be linear in the output size as well. In practice, it is more efficient to store a cell as a set of pointers or indices to its vertices, only requiring space $O(kV_k + C_k)$, with V_k denoting the number of vertices of $\text{Del}_k(P)$. Using this representation, the combinatorial step is not linear in the output size unless the number of cells of $\text{Del}_k(P)$ is linear in the number of vertices.

As the complexity of our algorithm is related to the output size, we recall the upper bound of $O(n^{\lfloor \frac{d+1}{2} \rfloor} k^{\lceil \frac{d+1}{2} \rceil})$ on the total size of the first k higher-order Delaunay mosaics [21]. At the same time, there exist point sets where this total size is only $\Theta(k^d n)$ [66]. We will also see in Proposition 4.8 from the next chapter that the complexity of the rhomboid tiling is the same regardless of the input, and is $\Theta(n^{d+1})$. As each $(e + 1)$ dimensional rhomboid yields e slices which are e -cells of $\text{Del}_k(P)$ for e consecutive values of k , this also means that the sum of the sizes of the order- k Delaunay mosaics is fixed. Thus for those point sets where the size of the order- k Delaunay mosaics grows faster for small k , we would expect the size of the mosaics to be smaller for large k closer to $n = \#P$. The next section will provide some experimental results confirming this expectation, and giving some insights on the size of individual order- k Delaunay mosaics which are not well understood.

3.4 Experimental Results

In 2 dimensions, the number of cells in the (order-1) Delaunay mosaic is always linear in the number of input points, while in $d \geq 3$ dimensions, the size of the mosaic depends on the input set itself—and not just its cardinality—and ranges from $\Omega(n)$ to $O(n^{\lceil d/2 \rceil})$ [61]. The asymptotic worst case is realized by points located on the moment curve, (t, t^2, \dots, t^d) with $t \in \mathbb{R}$, while e.g. uniformly sampled points within a sphere have expected linear size [30], as do uniformly sampled points on a convex polytope in \mathbb{R}^3 [43]. Under appropriate sampling conditions for points on a smooth surface, the size of the mosaic is $O(n \log n)$ [6].

Size in 3 dimensions. To shed light on the size range of order- k Delaunay mosaics, we compute them for a few 3-dimensional point sets relevant to these bounds. Note that for order- k Delaunay mosaics the number of vertices varies as well. Figure 3.5 shows the numbers of vertices and 3-dimensional cells for all higher-order Delaunay mosaics of four sets of size $n = 200$ each: points on the `moment curve`, points sampled on the `torus` (with major radius 1 and minor radius 0.5 obtained by uniformly sampling the angles of its parametrization), points uniformly sampled inside the `unit ball`, and a point set in convex position forming a `polytope` (obtained by uniformly sampling points inside a ball and randomly choosing 200 vertices of the convex hull).

The plots of vertex numbers and cell numbers generally resemble each other, with roughly three times as many cells as vertices. Other than in Figure 3.5, we therefore omit the information about the vertices and show only the plots for the cells. The `moment curve` and `polytope` sets are both in convex position. Nevertheless, the size of the mosaic for the moment curve grows large faster for small k , and reaches its peak at $k \approx n/3$, while for the polytope the peak is at $k \approx n/2$. Notice how a faster rise also goes along with an earlier decay. This is a consequence of the fact that the total size of all order- k Delaunay mosaics together—or, equivalently of the rhomboid tiling—only depends on the input size, n , and not on the relative position of the input points, see Proposition 4.8.

Size increase for small order. Looking more closely at the growth for small k relative to the input size, we observe that the `polytope` and `unit ball` exhibit linear growth while the size of the mosaic seems to grow quadratically for the `moment curve`, see Figure 3.6. This is consistent with the bounds on first-order Delaunay mosaics mentioned earlier. For

the **torus**, the size seems to grow slightly superlinearly, which is again consistent with the $O(n \log n)$ bound for smooth surfaces mentioned above.

Variance. To probe whether the graphs from Figure 3.6 are representative, we investigate the variance in number of cells for the **polytope** and the **unit ball**. As shown in Figure 3.7, the variance is particularly small for the **polytope**, and it is considerably larger of the **unit ball**. Curiously, the variance dips at $k = n/2$.

Generations. We also investigate the distribution of cells of different generations. All point sets exhibit a pattern similar to that in Figure 3.8, with the fraction of first-generation cells decreasing and the fraction of d -th-generation cells increasing as the order grows. The change is most prominent for small and large k , while the fractions remain almost constant in the range $k \approx n/2$, provided n is significantly larger than the dimension d .

Curse of dimensionality. Like many geometric structures, order- k Delaunay mosaics are subject to the dimensionality curse. Figure 3.9 shows how the size of order- k Delaunay mosaics behaves for point sets in different dimensions.

Vertex degrees. Order- k Delaunay mosaics exhibit an interesting distribution of vertex degrees for random point sets; see Figure 3.10. The distribution looks like the sum of two distributions—with the second one only covering values 2 modulo 3—and is exhibited for all k except very small and very large ones. We do not know the reason for vertices being frequently incident to $5, 8, 11, \dots$ d -cells, but suspect these numbers correspond to geometric configurations of cells of different generations, such as three octahedra sharing a common vertex with two tetrahedra.

Clusters. First-generation cells of any order- k Delaunay mosaic come in clusters connected by shared facets, see Figure 3.12 for an example. We investigate the distribution of their sizes, leaving the discussion of their potential algorithmic significance for Section 3.5.

Figure 3.11 shows cluster size distributions for different orders. For very small k , the distribution depends on how the points are sampled, while for all other k , the cluster sizes seem to follow an exponential distribution. The decay rate increases with k and seems




Figure 3.5: Number of vertices (left) and 3-dimensional cells (right) in the order- k Delaunay mosaics for four point sets of size 200, each in \mathbb{R}^3

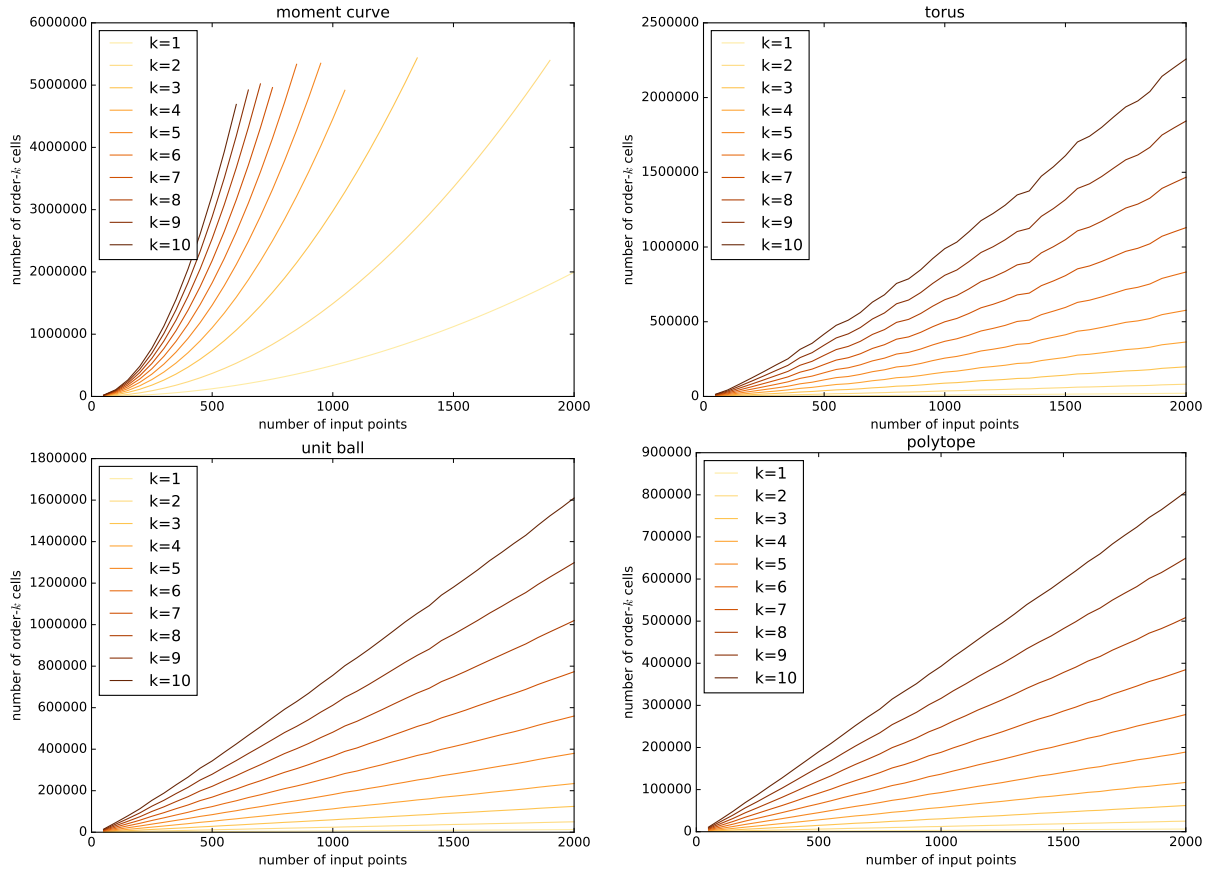


Figure 3.6: Number of cells in the order- k Delaunay mosaics for small k in relation to the input size, for various 3-dimensional point sets.




Figure 3.7: Variance of the number of 3-dimensional cells in the order- k Delaunay mosaics of randomly sampled points in convex position (*left*) and in a unit ball (*right*). The statistics of each plot are obtained from 30 sets of 50 points each. In *black*: the mean; in *red*: the range of one standard deviation around the mean; in *grey*: the range between the minimum and maximum.




Figure 3.8: Fraction of cells of each generation in the order- k Delaunay mosaic, for 50 random points in the unit 3-ball (*left*) and 20 random points in the unit 5-ball (*right*).




Figure 3.9: Number of d -cells in the order- k Delaunay mosaic for 20 points (*left*) and 50 points (*right*) randomly sampled in the unit ball for different dimensions d .




Figure 3.10: Vertex degree distribution in the order-50 Delaunay mosaic for 100 points sampled in the unit ball (*left*) and on the torus (*right*).

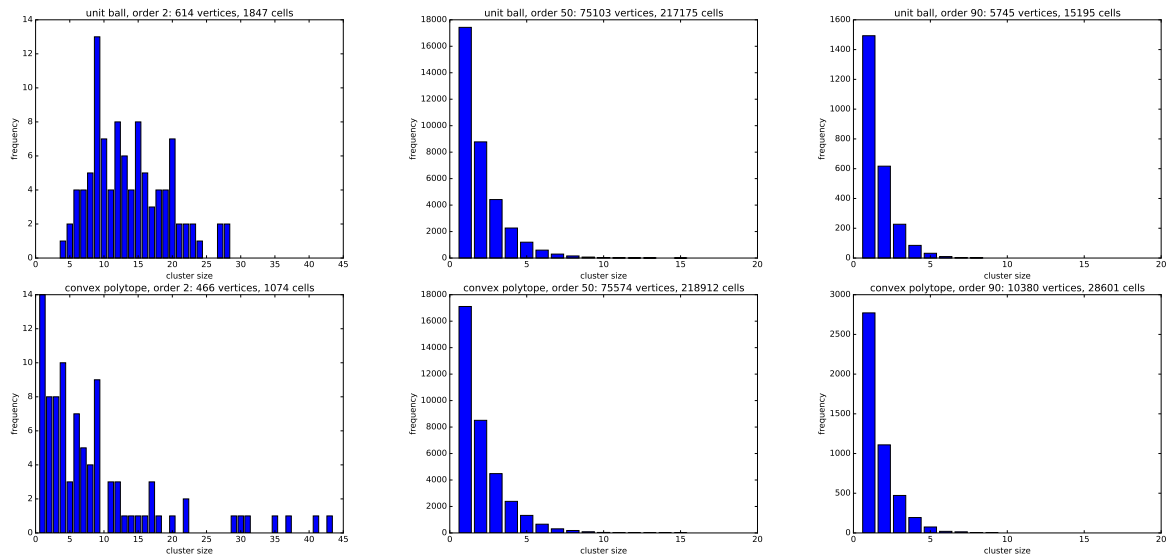


Figure 3.11: From *left to right*: distribution of cluster sizes in Delaunay mosaics of order 2, 50, and 90, for 100 random points in the unit ball (*top row*) and on a polytope (*bottom row*).




Figure 3.12: The order-2 Delaunay mosaic of 10 points (black) in the plane. First generation cells (red) form clusters (which are not necessarily convex); second generation cells in blue.

to be linked to the fraction of first-generation cells. It culminates in all clusters being singletons for $k = n - 3$. For $k > n - 3$, there are no more first-generation cells.

3.5 Extensions

In this chapter we presented a simple algorithm for computing order- k Delaunay mosaics of unweighted point sets in Euclidean space. It employs a new geometric structure, the rhomboid tiling, and we provided open-source implementations for this algorithm. The remainder of this section discusses possible extensions and optimizations of our algorithm.

Weighted setting. Our algorithm generalizes to points with real weights, but not easily. The main challenge is the extraction of the vertices of the order- k mosaic from lower-order mosaics. This extraction relies on Theorem 3.5, which does not hold for weighted points. Indeed, a crucial assumption in this theorem is that every lifted hyperplane is incident to the depth-0 chamber of the arrangement, and this property is generally violated for weighted points. This is the same assumption used in the prior dimension-agnostic algorithms [1; 66; 67]. For sets of weighted points that satisfy this assumption, our algorithm and these prior algorithms still work. To overcome this limitation, we would need a way to detect

all bowls in the arrangement, because they correspond to the vertices in the Delaunay mosaics our algorithm is not able to find. Identifying these vertices is an independent problem, and any solution to it can be combined with our algorithm. Once we know these vertices and add them to the appropriate mosaics, our algorithm works as before.

Clusters of cells. As mentioned in Section 3.4, first-generation cells in the order- k Delaunay mosaic are organized in clusters, also see Figure 3.12. To formally define them, consider the graph whose nodes are the cells and whose arcs are the shared facets (i.e. the 1-skeleton of the order- k Voronoi tessellation). A *cluster* is a connected component in the subgraph induced by the first-generation cells. It is not difficult to see that two such cells belong to a common cluster if and only if the corresponding rhomboids have the same anchor vertex. Let ρ be one of these rhomboids and recall that the anchor vertex is $P_{in}(\rho)$, which in this case is a collection of $k - 1$ points of P . Each combinatorial vertex of any cell in the cluster contains these $k - 1$ points, plus one additional point, which differentiates between these vertices. This means that the cluster is combinatorially equivalent to a subcomplex of the order-1 Delaunay mosaic of these additional points.

With this insight, in our algorithm we could replace the weighted Delaunay mosaic of the entire vertex set by multiple instances of unweighted Delaunay mosaics, namely one per cluster. This alternative strategy avoids the need to compute averages of points at the cost of extra book-keeping to group the vertex set of $\text{Del}_k(P)$ into clusters. We mention that in \mathbb{R}^2 , the structure of each cluster satisfies the requirements that allow for the construction in time linear in the number of points [2].

Exact arithmetic. The CGAL software library [81], which we employ in our C++ implementation, supports exact arithmetic by distinguishing between *exact constructions* and *exact predicates*. The latter are geometric tests with a **true** or **false** answer, such as whether or not a given point lies on a given sphere. By itself, the CGAL algorithm for weighted Delaunay triangulations requires exact predicates but no exact constructions. Our algorithm, on the other hand, computes averages of collections of input points, which are the locations of the vertices of the mosaic. This is an exact construction and indeed the only one needed to run our algorithm with exact arithmetic. In practice, exact constructions are a significant overhead with noticeable impact on the runtime, which would be nice to

avoid. One such possibility is to switch to computing clusters of first-generation cells as described above.

4 Multi-cover persistence

A *filtration* is a (not necessarily countable) sequence of increasing topological spaces or regular complexes linked by inclusions. Given a finite $P \subseteq \mathbb{R}^d$ and a radius $r \geq 0$, the k -fold covers give rise to two filtrations — one in *scale* obtained by fixing k and increasing r , and the other in *depth* obtained by fixing r and decreasing k — and we compute the persistence diagrams of both. While standard methods suffice for the filtration in scale, we need novel geometric and topological concepts for the filtration in depth due to the combinatorial difference of the Delaunay mosaics from one value of k to the next. In particular, we make use of the rhomboid tiling in \mathbb{R}^{d+1} whose horizontal integer slices are the order- k Delaunay mosaics of P , to connect consecutive Delaunay mosaics and construct a zigzag module that is isomorphic to the persistence module of the multi-covers. We get the persistence diagram using the algorithm in [13; 14].

After formally introducing the k -fold covers, in Section 4.1 we recall the notion of homotopy and introduce a radius function \mathcal{R}_k on the order- k Delaunay mosaic whose sublevel sets generalize the notion of α -shapes from $k = 1$ to orders $k \geq 1$ [34; 55]. We then proceed to give a recipe to explicitly compute the radius function (Section 4.2). In Section 4.3 we formally introduce persistence and show how to compute persistence of the multi-covers in scale. To compute persistence in depth, we first introduce a radius function on the rhomboid tiling and then construct a zig-zag sequence of complexes that yields the same persistence as the multi-covers (Section 4.4). We close the chapter with a discussion of 2-parameter persistence (Section 4.5).

k -fold cover. Let $P \subseteq \mathbb{R}^d$ be finite. Given a radius $r \geq 0$, the k -fold cover of P and r ,

denoted by $\text{Cover}_k(P, r)$, consists of all points $x \in \mathbb{R}^d$ for which there are k or more points $p \in P$ with $\|p - x\| \leq r$, or, in other words, the points $x \in \mathbb{R}^d$ that are covered by at least k of the balls of radius r centered at the points $p \in P$. A different viewpoint is to consider the k -th distance function $f_k: \mathbb{R}^d \rightarrow \mathbb{R}$ that maps each point x to the distance to the k -th closest point $p \in P$. Then the covers are $\text{Cover}_k(P, r) = f_k^{-1}(-\infty, r]$, the *sublevel sets* of f_k for $r \in \mathbb{R}$. We have

$$\text{Cover}_k(P, r) \subseteq \text{Cover}_k(P, s), \quad (4.1)$$

$$\text{Cover}_k(P, r) \subseteq \text{Cover}_\ell(P, r), \quad (4.2)$$

whenever $r \leq s$ and $\ell \leq k$, also see Figure 4.1. Thus these covers are a filtration for fixed k and increasing radius r , which we call the *filtration in scale*, but also for fixed r and decreasing k , which we call the *filtration in depth*. In fact, as the covers $\text{Cover}_k(P, r)$ are parametrized by two parameters, r and k , with inclusions with respect to either parameter, they are what we call a *bifiltration*.

We are interested in computing the persistent homology of the multi-cover filtrations, both in the direction of increasing radius and in the direction of decreasing order, which we refer to as *persistence in scale* and *persistence in depth* respectively. To do so, we represent the covers by complexes, namely by subcomplexes of the Delaunay mosaics. Varying the radius, we get a filtration of subcomplexes of the order- k Delaunay mosaic, and the persistent homology can be computed with standard methods; see e.g. [33, Chapter VII] for background. We will outline the details in the first few sections of this chapter. On the other hand, varying the order, we get subcomplexes of different Delaunay mosaics which do not form a filtration, and we need a novel algorithm to compute persistent homology (Section 4.4). While there is no notion of a persistence diagram for a bifiltration, one can still talk about *2-parameter persistence* (Section 4.5) which, in a sense, subsumes persistence in scale and depth.

In all settings, we exploit the rhomboid tiling to shed light on the filtration of multi-covers. In Section 4.4.1 we introduce a filtration on the rhomboid tiling which allows us to obtain sequences of subcomplexes that aid in the computation of persistent homology of the multi-covers. First, however, we have to cover some definitions on homotopy, homology and persistence to give concrete descriptions of our problem statements.

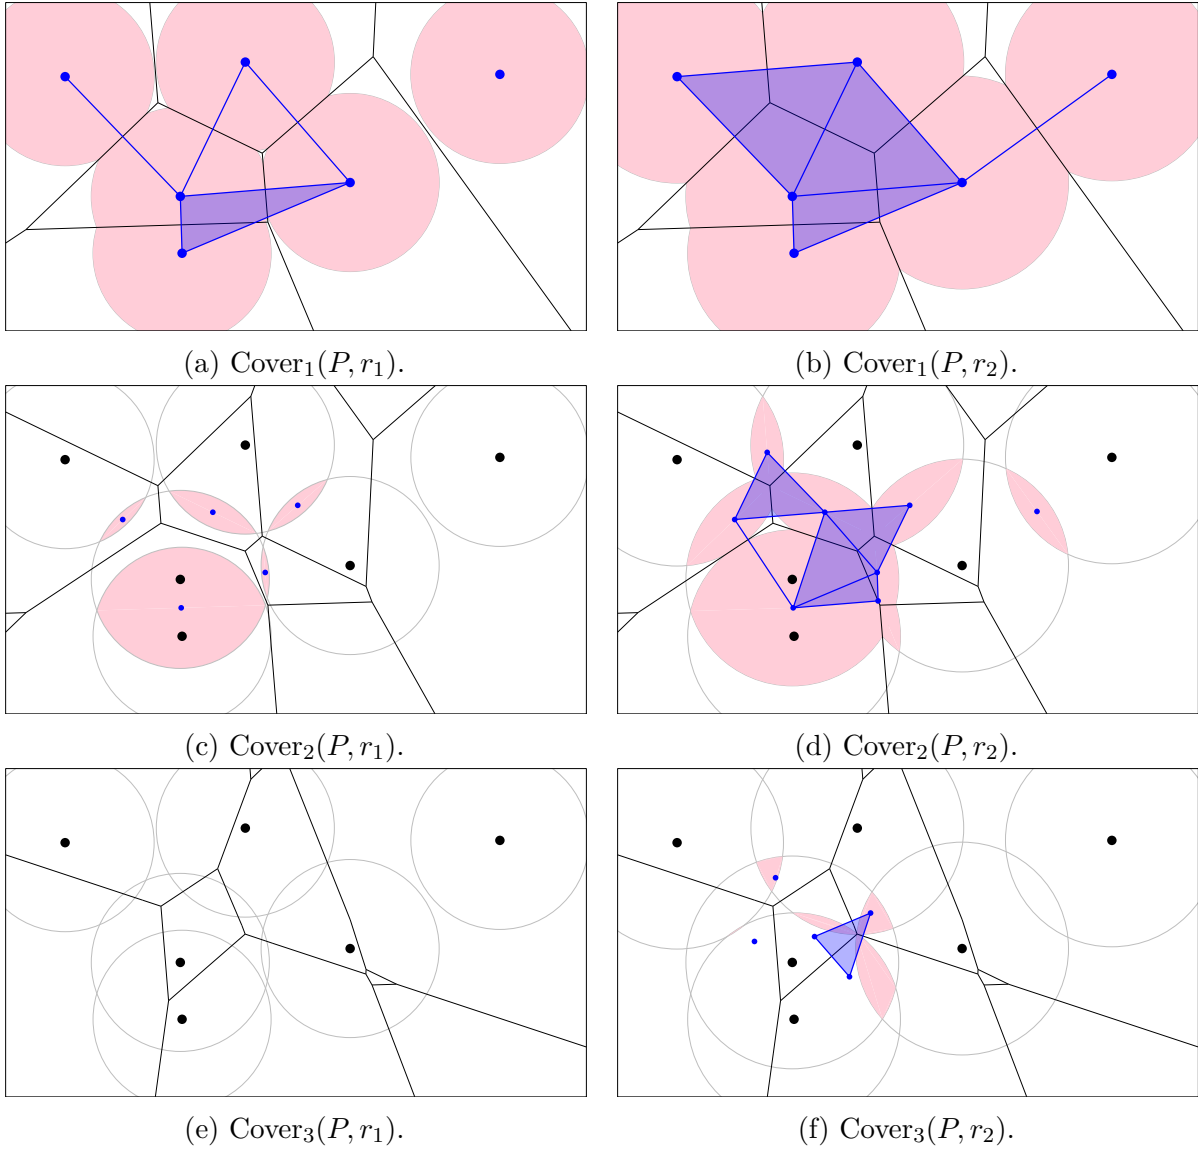


Figure 4.1: The k -fold covers (pink) with their corresponding order- k Voronoi decompositions (black) and homotopy equivalent order- k Delaunay subcomplexes (blue) for $k = 1, 2, 3$ and radii $r_1 < r_2$.

4.1 Homotopy

Before we can relate the multi-covers and order- k Delaunay complexes, we recall basic definitions and classic results from homotopy theory, e.g. see [33, Section III] or [48]. Let \mathbb{X} and \mathbb{Y} be two topological spaces, and $f, g: \mathbb{X} \rightarrow \mathbb{Y}$ continuous maps between them. A *homotopy* between f and g is a continuous map $H: \mathbb{X} \times [0, 1] \rightarrow \mathbb{Y}$ such that $H(x, 0) = f(x)$ and $H(x, 1) = g(x)$ for all $x \in \mathbb{X}$. If such a homotopy exists, we call f and g *homotopic* and write $f \sim g$. If there are two maps $f: \mathbb{X} \rightarrow \mathbb{Y}$ and $g: \mathbb{Y} \rightarrow \mathbb{X}$ such that $g \circ f \sim \text{id}_{\mathbb{X}}$ and $f \circ g \sim \text{id}_{\mathbb{Y}}$, we call \mathbb{X} and \mathbb{Y} *homotopy equivalent* or say they have the same *homotopy*.

type. In that case we call the maps f and g *homotopy equivalences*. A space is *contractible* if it is homotopy equivalent to a single point. For two spaces $\mathbb{Y} \subseteq \mathbb{X}$ and a continuous map $f: \mathbb{X} \rightarrow \mathbb{Y}$ with $f(y) = y$ for $y \in \mathbb{Y}$, we call f a *deformation retraction* if $f \sim \text{id}_{\mathbb{X}}$, and in that case we call \mathbb{Y} a *deformation retract* of \mathbb{X} .

An *abstract simplicial complex* K is a collection of sets such that if $\sigma \in K$ and $\tau \subseteq \sigma$, then also $\tau \in K$. The sets $\sigma \in K$ are its (*abstract*) *simplices*. The dimension of a simplex is $\dim \sigma = \#\sigma - 1$ and the dimension of K is $\max\{\dim \sigma: \sigma \in K\}$. We call $V(K) := \bigcup_{\sigma \in K} \sigma$ the vertex set of K .

Given a map $\phi: V(K) \rightarrow \mathbb{R}^d$, if the collection of $\text{conv}\{\phi(v): v \in V(\sigma)\}$ for all $\sigma \in K$ is a geometric simplicial complex, we call it a *geometric realization* of K in \mathbb{R}^d . By the Geometric Realization Theorem [33, Section III.1], any d -dimensional simplicial complex has a geometric realization in \mathbb{R}^{2d+1} . Furthermore all geometric realizations have the same homotopy type. Thus it makes sense to talk of the homotopy type of an abstract simplicial complex.

Let \mathcal{S} be a finite collection of sets $Q \in \mathbb{R}^d$. Then the *nerve* of \mathcal{S} is an abstract simplicial complex defined as

$$\text{Nerve } \mathcal{S} := \{Q \subseteq \mathcal{S}: \bigcap_{Q \in Q} Q \neq \emptyset\},$$

i.e. the collection of subsets of \mathcal{S} that have non-empty common intersection. If the sets are furthermore convex and closed, we have the following theorem.

Theorem 4.1 (Nerve theorem [60]). *Let \mathcal{S} be a finite collection of closed, convex sets $Q \in \mathbb{R}^d$. Then $\text{Nerve } \mathcal{S}$ and $\bigcup_{Q \in \mathcal{S}} Q$ are homotopy equivalent.*

As a first application of the Nerve Theorem, we will decompose the k -fold cover into convex pieces to obtain a simplicial complex of the same homotopy type. This complex is closely related to $\text{Del}_k(P)$ and will later aid us in computing multi-cover persistence.

The order- k Voronoi tessellation decomposes the k -fold cover into convex sets. To see this, let $\#Q = k$ and define the *restricted Voronoi domain* $\text{dom}(Q, r) = \text{dom}(Q) \cap \text{Cover}_k(Q, r)$, which is an intersection of convex sets and therefore convex. We write $\text{Vor}_k(P, r)$ for the collection of domains $\text{dom}(Q, r)$ with $\#Q = k$, and since $\text{dom}(Q, r) = \text{dom}(Q) \cap \text{Cover}_k(P, r)$, we refer to this as the *Voronoi decomposition* of $\text{Cover}_k(P, r)$.

Figure 4.1 provides an example. Since $\text{dom}(Q, r) \subseteq \text{dom}(Q)$, the dual of this decomposition is a subcomplex of the order- k Delaunay mosaic, which we call the *restricted order- k Delaunay complex* and denote $\text{Del}_k(P, r) \subseteq \text{Del}_k(P)$. The $\text{Del}_k(P, r)$ give a filtration for increasing radius r , and we define the *radius function* on the order- k Delaunay mosaic as $\mathcal{R}_k: \text{Del}_k(P) \rightarrow \mathbb{R}$ which assigns each cell σ the smallest r such that $\sigma \in \text{Del}_k(P, r)$. $\mathcal{R}_k(\sigma)$ is called the *radius value* or *filtration value* of σ . Note that \mathcal{R}_k is *monotonic*: $\mathcal{R}_k(\tau) \leq \mathcal{R}_k(\sigma)$ whenever τ is a face of σ . Monotonic functions on a regular complex K have the property that their sublevel sets are subcomplexes of K ; in this case the sublevel sets are the restricted Delaunay complexes: $\mathcal{R}_k^{-1}(-\infty, r] = \text{Del}_k(P, r)$.

The Nerve Theorem implies that the nerve of the non-empty $\text{dom}(Q, r)$ with $\#Q = k$, which we will call the *Delaunay Nerve* and denote as $\text{DelN}_k(P, r)$, is homotopy equivalent to $\text{Cover}_k(P, r)$. As for the order- k Delaunay mosaic, we define a radius function $\mathcal{RN}_k: \text{DelN}_k(P) \rightarrow \mathbb{R}$ which assigns each cell σ the smallest r such that $\sigma \in \text{DelN}_k(P, r)$. However, $\text{DelN}_k(P, r)$ is not the same as $\text{Del}_k(P, r)$: While the nerve is always a simplicial complex, $\text{Del}_k(P, r)$ is not, as observed in Section 3.1. The two complexes are closely related however. Their vertices correspond to the non-empty domains $\text{dom}(Q, r)$ with $Q \subseteq P, \#Q = k$, and like for the order- k Delaunay mosaic we define the *combinatorial vertex set* of the nerve, $V(\text{DelN}_k(P, r))$, as the collection of those subsets Q . The maximal cells of both complexes correspond to vertices $v \in \text{Vor}_k(P, r)$, and while a maximal cell σ' of $\text{DelN}_k(P, r)$ is a simplex spanned by the subsets Q whose $\text{dom}(Q, r)$ are incident to v , the corresponding maximal cell σ in $\text{Del}_k(P, r)$ is the convex hull of the geometric embeddings $\frac{1}{k} \sum_{q \in Q} q$ of those Q , and thus $V(\sigma) = V(\sigma')$. Finally, for both complexes, the intersection of two cells σ and τ is the cell spanned by the vertex set $V(\sigma) \cap V(\tau)$. This is sufficient to show that both complexes are, in fact, homotopy equivalent, by the following Lemma.

Lemma 4.2. *Assume we have two regular complexes K and K' with $V(K) = V(K')$, and a bijection between maximal cells such that corresponding $\sigma \in K$ and $\sigma' \in K'$ fulfill $V(\sigma) = V(\sigma')$. Furthermore assume that for any set S of maximal cells K and the set S' of corresponding maximal cells of K' , it holds that $V(\bigcap_{\sigma \in S} \sigma) = V(\bigcap_{\sigma' \in S'} \sigma')$. Then K and K' are homotopy equivalent.*

Proof. Let M and M' be the collection of maximal cells of K and K' respectively. Then the bijection between M and M' is a bijection on the vertices of Nerve M and Nerve M' ,

which extends to a bijection on the cells of Nerve M and Nerve M' due to the assumption about intersections of maximal cells. This means the two nerves only differ in a relabeling of their vertices, and thus also have to be homotopy equivalent. As the set of maximal cells of a regular complex is a closed, convex cover of the complex itself, K and K' are homotopy equivalent to this nerve by the Nerve theorem, and by transitivity homotopy equivalent to each other. \square

As a direct consequence, we get the homotopy equivalence of the k -fold cover and the restricted Delaunay mosaic.

Corollary 4.3 (Almost Nerve). *Let $P \subseteq \mathbb{R}^d$ be locally finite and in general position. For every integer $k \geq 1$ and real $r \geq 0$, $\text{Del}_k(P, r)$ and $\text{Cover}_k(P, r)$ have the same homotopy type.*

In the next section we will address the problem of computing the radius value of each cell of $\text{Del}_k(P)$ so we know for which r it appears in $\text{Del}_k(P, r)$.

4.2 Computing the Delaunay filtration

Recall from Section 3.2.1 that for a cell σ in the order- k Delaunay mosaic, we defined $P_{in}(\sigma) := \bigcap V(\sigma)$, $P(\sigma) := \bigcup V(\sigma)$, $P_{on}(\sigma) := P(\sigma) \setminus P_{in}(\sigma)$ and $P_{out}(\sigma) := P \setminus P(\sigma)$. For a cell σ of $\text{Del}_k(P)$ and a sphere S , we call S σ -constrained if it contains $P_{in}(\sigma)$ (in its interior or on its surface), has $P_{on}(\sigma)$ on its surface and no other points from P in its interior. Let $S_{min}(\sigma)$ be the smallest σ -constrained sphere of σ . Then the following lemma gives us a recipe for computing the radius value of any cell in the order- k Delaunay mosaic.

Lemma 4.4. *The radius value of an order- k Delaunay cell σ is the radius of $S_{min}(\sigma)$.*

Proof. By definition, we need to find the smallest radius r for which the intersection of $\text{dom}(Q, r)$ for $Q \in V(\sigma)$ is non-empty. Writing $P(\sigma) = P_{in}(\sigma) \cup P_{on}(\sigma) = \bigcup_{Q \in V(\sigma)} Q$, this is equivalent to the condition

$$\left(\bigcap_{p \in P(\sigma)} B_r(p) \right) \cap \left(\bigcap_{Q \in V(\sigma)} \text{dom}(Q) \right) \neq \emptyset.$$

So we need to find a ball $B_r(x)$ with minimal radius r that contains all points of $P(\sigma)$ and whose center, x , is in the intersection of the domains $\text{dom}(Q)$ for $Q \in V(\sigma)$. Let $x \in \mathbb{R}^d$ be such a point. From having to be in the intersection of the order- k Voronoi domains, we get the following constraints:

(1) Within P , all points from $P(\sigma)$ have to be among the k closest points to x . (Note that there can possibly be more than k of these if some points are tied for k -th closest). This also implies that all points not in $P(\sigma)$ are at least as far away as those from $P(\sigma)$.

(2) For a point p in $P_{on}(\sigma)$ there is in particular a $Q \in V(\sigma)$ that does not contain p . With $x \in \text{dom}(Q)$, this means that p is at least as far away from x as all the points from Q , and so it can only be tied for k -th closest. Therefore x must satisfy that all points p from $P_{on}(\sigma)$ have the same distance from x .

Taking these two constraints and adding the criterion that r should be minimal, this yields that $B_r(x)$ is the smallest enclosing ball of $P(\sigma)$ that has the points $P_{on}(\sigma)$ on its boundary and all other points from P outside of its interior, which is exactly the definition of $S_{min}(\sigma)$. \square

Algorithmically, $S_{min}(\sigma)$ can be computed using a generalization of Welzl's algorithm [91] for computing smallest enclosing spheres. This algorithm works in any dimension, and by design it already works with additional constraints for points that are required to be on the surface of the sphere. Algorithm 2 extends this algorithm to work with exclusion constraints, facilitating the computation of the radius values from the previous lemma. The computation of circumspheres, as required by the algorithm, is done by solving a system of linear equations.

4.2.1 Optimizing the computation of constrained spheres

When computing $S_{min}(\sigma)$ for some cell σ , notice how in the definition of $S_{min}(\sigma)$ we have one constraint for each point of P : Each point of P is required to be either inside, on, or outside the sphere. Therefore the previous algorithm takes $O(\#P)$ to compute $S_{min}(\sigma)$, regardless of the order k for which σ is a cell of $\text{Del}_k(P)$. In particular, for small order k , $P(\sigma)$ is small, while $P_{out}(\sigma)$ consists of most of the points of P . In the following lemma, we show that only a small subset of the constraints of $P_{out}(\sigma)$ need to be checked, allowing

Algorithm 2 computes the smallest enclosing sphere of P_{in} having P_{on} on its surface and not containing P_{out} in its interior.

```

function SMALLESTENCLOSINGSPHERE( $P_{in}, P_{on}, P_{out}$ )
  if  $\#P_{on} = d + 1$  then
     $CS := \text{CIRCUMSPHERE}(P_{on})$ 
    if not all  $p \in P_{in}$  inside  $CS$  or not all  $p \in P_{out}$  outside  $CS$  then
      return "No valid sphere exists."
    end if
  else if  $P_{out} \neq \emptyset$  then
    Pick  $p \in P_{out}$ 
     $CS := \text{SMALLESTENCLOSINGSPHERE}(P_{in}, P_{on}, P_{out} \setminus \{p\})$ 
    if  $p$  inside  $CS$  then
       $CS := \text{SMALLESTENCLOSINGSPHERE}(P_{in}, P_{on} \cup \{p\}, P_{out})$ 
    end if
  else if  $P_{in} \neq \emptyset$  then
    Pick  $p \in P_{in}$ 
     $CS := \text{SMALLESTENCLOSINGSPHERE}(P_{in} \setminus \{p\}, P_{on}, P_{out})$ 
    if  $p$  outside  $CS$  then
       $CS := \text{SMALLESTENCLOSINGSPHERE}(P_{in}, P_{on} \cup \{p\}, P_{out})$ 
    end if
  else
     $CS := \text{CIRCUMSPHERE}(P_{on})$ 
  end if
  return  $CS$ 
end function

```

to compute $S_{min}(\sigma)$ in $O(k)$. We remark that in Section 5.1.1 we will optimize this further using insights about the rhomboid tiling to yield a $O(1)$ algorithm.

Lemma 4.5. *$S_{min}(\sigma)$ is the smallest enclosing sphere of $P_{in}(\sigma)$ that has $P_{on}(\sigma)$ on its surface and does not contain any points of $P(\tau) \setminus P(\sigma)$ in its interior where τ is any co-facet of σ .*

Proof. Let x be the center of $S_{min}(\sigma)$. We claim that x must be in the Voronoi cell $\sigma^* := \bigcap_{Q \in V(\sigma)} \text{dom}(Q)$ which is dual to σ . From this claim the result follows, as by definition any point y in σ^* is the center of a σ -constrained sphere: all points $p \in P_{on}(\sigma)$ have the same distance from y , while points $p \in P_{in}(\sigma)$ are closer and points $p \in P_{out}(\sigma)$ are further away.

To prove this claim, note that the constraints from $P_{on}(\sigma)$ define an affine subspace $R_{on} \subseteq \mathbb{R}^d$ that must contain x . It is of dimension $d_{on} := d + 1 - \#P_{on}(\sigma)$, unless $P_{on}(\sigma) = \emptyset$ in which case $d_{on} = d$. Note that $\#P_{on}(\sigma) \neq 1$ and thus $d_{on} = d$ if and only if $P_{on}(\sigma) = \emptyset$. The Voronoi cell σ^* is contained in R_{on} and is also of dimension d_{on} . Thus if $d_{on} = 0$, our claim trivially holds. So assume $d_{on} \neq 0$, and $x \notin \sigma^*$. Then there is a facet τ^* of σ^* that separates the interior of σ^* from x , in the sense that the $(d_{on} - 1)$ -dimensional subspace R_{τ} containing τ^* divides R_{on} into two halves, one containing the interior of σ^* and the other containing x .

If $d_{on} = d$, then σ is a vertex v and $\sigma^* = \text{dom}(v)$. Then $P_{on}(\tau) = \{p, q\}$ for some point $p \in v$ and some point $q \in v'$ for some Voronoi domain $\text{dom}(v')$ adjacent to σ^* . Furthermore R_{τ} is the perpendicular bisector between p and q . Now being on the opposite side of R_{τ} than σ^* , $q \in P_{out}(\sigma)$ is closer to x than $p \in P_{in}(\sigma)$, a contradiction.

If $d_{on} \neq d$, then due to τ^* being in R_{on} and its dimension being $d_{on} - 1$, we know that $P_{on}(\tau) = P_{on}(\sigma) \cup \{p\}$ for some point $p \in P \setminus P_{on}(\sigma)$. R_{τ} is the set of points equidistant from the points of $P_{on}(\tau)$, thus for points $y \in R_{on}$ on one side of R_{τ} we have $\|y - p\| < \|y - q\|$ while on the other side we have $\|y - p\| > \|y - q\|$ for all $q \in P_{on}(\sigma)$.

Case 1: $p \in P_{in}(\sigma)$. Then for points $y \in \sigma^*$ we have $\|y - p\| < \|y - q\|$ for all $q \in P_{on}(\sigma)$, and thus $\|x - p\| > \|x - q\|$ for all $y \in P_{on}(\sigma)$, a contradiction.

Case 2: $p \in P_{out}(\sigma)$. Then for points $y \in \sigma^*$ we have $\|y - p\| > \|y - q\|$ for all $q \in P_{on}(\sigma)$, and thus $\|x - p\| < \|x - q\|$ for all $y \in P_{on}(\sigma)$. However we also have $p \in P(\tau)$ and thus $p \in P(\tau) \setminus P(\sigma)$, a contradiction.

4.3 Homology and Persistence

We recall essential definitions and results related to homology [48] and persistent homology [33]. Let K be a regular complex. An e -chain is a formal sum $\sum a_i \sigma_i$ where the σ_i are e -cells of K and a_i are coefficients from a field or ring; for our purposes we will use $a_i \in \mathbb{Z}_2$. Two chains $\sum a_i \sigma_i$ and $\sum b_i \sigma_i$ can be added to yield a chain $\sum (a_i + b_i) \sigma_i$, and the e -chains together with this addition form a *group of e -chains* $\mathbf{C}_e = \mathbf{C}_e(K)$. We define the *boundary map* $\partial_e: \mathbf{C}_e \rightarrow \mathbf{C}_{e-1}$ that maps each e -cell σ to $\partial_e(\sigma)$, the formal sum of its facets, and by extension maps a chain $\sum a_i \sigma_i$ to the sum $\sum a_i \partial_e(\sigma_i)$ of the boundaries of its cells. This map is a group homomorphism. The collection of boundary maps of a complex K , for all dimensions e , can be encoded by a *boundary matrix* B , which is a matrix with coefficients in \mathbb{Z}_2 whose rows and columns are indexed by the cells of K , and whose entry $B_{\tau, \sigma} = 1$ iff τ is a facet of σ . If we encode a chain as a \mathbb{Z}_2 -vector v with non-zero entries for those cells that are part of the chain, then Bv is the boundary of v .

An e -cycle is an e -chain whose boundary is 0, i.e. the empty sum. The set of e -cycles forms a group of e -cycles $\mathbf{Z}_e = \mathbf{Z}_e(K) := \ker \partial_e$. A e -boundary is a e -chain that is the boundary of some $(e + 1)$ -chain. The set of e -boundaries forms a group of e -boundaries $\mathbf{B}_e = \mathbf{B}_e(K) := \text{im } \partial_{e+1}$. The boundary map has the property that for every $(e + 1)$ -chain c it holds that $\partial_e(\partial_{e+1}(c)) = 0$. In other words, every e -boundary is a e -cycle, and in fact \mathbf{B}_e is a subgroup of \mathbf{Z}_e . This motivates the definition of the *e -th homology group* as the quotient $\mathbf{H}_e(K) = \mathbf{Z}_e(K)/\mathbf{B}_e(K)$. Intuitively, the e -th homology group consists of equivalence classes of e -cycles with the cycles in each class differing by e -boundaries, and e -cycles in the same homology class are called *homologous*, see Figure 4.2. Note that in our setting with chain coefficients in \mathbb{Z}_2 , the e -th homology group is in fact a vector space. We will write $\mathbf{H}_*(K)$ for the homology groups of all dimensions e .

Singular homology generalizes this theory from regular complexes to topological spaces. The cells that a regular complex consists of are replaced by continuous maps from the standard e -dimensional simplex into the topological space, called singular maps. Chains are formal sums of such maps. We omit the details here, but remark that it thus also makes sense to talk about homology (and persistence) for topological spaces. When applying




Figure 4.2: The blue 1-cycle c is homologous to the red 1-cycle c' because they differ in a boundary: $c = c' + \partial_2(\sigma)$.

homology to regular complexes or topological spaces linked by continuous maps, they are turned into homology groups that are linked by group homomorphisms. In this sense, homology can be viewed as a *functor*, and we will make occasional use of this terminology.

One important property is that if two topological spaces are homotopy equivalent, then their e -th homology groups are isomorphic, for any dimension e . The reverse is however not true. Thus homology gives us a coarser view on the topology of spaces than homotopy. While the question of the homotopy type of a cell complex is generally undecidable [80], homology is attractive because it is computable in polynomial time for chain coefficients in \mathbb{Z}_2 . Specifically, it comes down to computing the Smith normal form of the boundary matrix.

Let K be a regular complex, and $f: K \rightarrow \mathbb{R}$ a monotonic function. Let $a_1 < a_2 < \dots < a_m$ be the distinct function values of f , and let $K_i := f^{-1}(-\infty, a_i]$ be the sublevel sets of f . We get a sequence of complexes $\emptyset =: K_0 \subset K_1 \subset \dots \subset K_m = K$ linked by inclusions, which is the filtration of f . The inclusion map from K_i to K_j for $i \leq j$ induces a homomorphism $\phi_e^{i,j}: \mathbf{H}_e(K_i) \rightarrow \mathbf{H}_e(K_j)$ for each dimension e , and the filtration thus corresponds to a sequence of homology groups,

$$0 = \mathbf{H}_e(K_0) \rightarrow \mathbf{H}_e(K_1) \rightarrow \dots \rightarrow \mathbf{H}_e(K_m) = \mathbf{H}_e(K).$$

We also refer to such a sequence of homology groups with homomorphisms from left to right as a *persistence module*. We define the e -th *persistent homology group* as $\mathbf{H}_e^{i,j} := \text{im } \phi_e^{i,j}$ for $0 \leq i \leq j \leq m$. Intuitively, it consists of those homology classes of K_i that are still alive at K_j : $\mathbf{H}_e^{i,j} = \mathbf{Z}_e(K_i) / (\mathbf{B}_e(K_j) \cap \mathbf{Z}_e(K_i))$. We say a e -dimensional homology class $\gamma \in \mathbf{H}_e(K_i)$ is *born at* K_i if $\gamma \notin \mathbf{H}_e^{i-1,i}$, and we call a_i its *birth time*. We say it *dies entering* K_j if, while

going from K_{j-1} to K_j via $\phi^{j-1,j}$, it merges with a older class, i.e. $\phi_e^{i,j-1}(\gamma) \notin \mathbf{H}_e^{i-1,j-1}$ but $\phi_e^{i,j}(\gamma) \in \mathbf{H}_e^{i-1,j}$. In that case we call a_j its *death time*. Some homology classes may not have a death time; in that case we set its death time to $+\infty$. We call $a_j - a_i$ its *persistence* and the pair (a_i, a_j) its *persistence pair*. The multiplicity of a persistence pair is the number of independent e -dimensional homology classes born at K_i and dying entering K_j . Then the e -dimensional *persistence diagram*, $\text{Dgm}_e(f)$, is the multi-set of all persistence pairs with their respective multiplicities, and can be visualized in $\mathbb{R} \times (\mathbb{R} \cup \{+\infty\})$. For convenience, we will sometimes visualize all e -dimensional persistence diagrams in one single diagram $\text{Dgm}(f)$, in which every persistence pair is annotated with its dimension e .

Algorithmically, persistence is computed using the *boundary matrix reduction algorithm* [33, Chapter VII]. Given the boundary matrix of a complex K , indexed by the cells of K in order of their filtration value, a reduction of its columns akin to Gaussian elimination produces a simplified matrix from which the information about persistence pairs can be extracted. Thus this algorithm runs in time cubic in the number of cells n . This algorithm has been implemented and optimized [10; 82; 65]. It is widely used and in practice runs in almost linear time for many practical inputs.

The notion of persistence is not limited to filtrations, i.e. increasing sequences of regular complexes. We can also work with a zig-zag sequence of complexes that are connected by inclusions in either direction, e.g. $K_0 \subseteq K_1 \supseteq K_2 \subseteq \dots \supseteq K_m$. Algorithms for persistence in this setting have been implemented and their runtime is also $O(n^3)$ [14]. Algorithms also exist more generally for homology groups linked by zig-zag sequences of homomorphisms [13]. However in this case the running time is $O(n^4)$, and we would have to explicitly compute the homology groups first. Thus our aim will be to develop algorithms that only involve inclusion sequences of complexes.

In practice, such as in the case of the k -fold cover filtration, we might want to compute persistence for a sequence of topological spaces rather than regular complexes. While singular homology provides a definition, it is usually impractical to compute due to the infinite number of singular maps into the topological space. Even if we find a sequence of complexes each with the same homology as the corresponding topological space, as in the case of $\text{Cover}_k(P, r)$ and $\text{Del}_k(P, r)$, we need need to ensure that the resulting persistence diagrams are the same as well. The *Persistence Equivalence Theorem* [94] gives us a sufficient condition.

Consider two persistence modules $\mathbf{H}_e^{(i)}$ and $\widetilde{\mathbf{H}}_e^{(i)}$ linked by homomorphisms, and furthermore homomorphisms $\phi_i: \mathbf{H}_e^{(i)} \rightarrow \widetilde{\mathbf{H}}_e^{(i)}$ linking each pair of homology groups, as follows:

$$\begin{array}{ccccccc}
 \widetilde{\mathbf{H}}_e^{(0)} & \longrightarrow & \widetilde{\mathbf{H}}_e^{(1)} & \longrightarrow & \dots & \longrightarrow & \widetilde{\mathbf{H}}_e^{(m)} \\
 \uparrow \phi_0 & & \uparrow \phi_1 & & & & \uparrow \phi_m \\
 \mathbf{H}_e^{(0)} & \longrightarrow & \mathbf{H}_e^{(1)} & \longrightarrow & \dots & \longrightarrow & \mathbf{H}_e^{(m)}
 \end{array}$$

These could for example be the homology groups of sublevel sets of two different filtered complexes or spaces. If this diagram commutes and the ϕ_i are isomorphisms, then we call the two modules *isomorphic*. The *Persistence Equivalence Theorem* states that if two persistence modules are isomorphic, then the e -th persistence diagrams of the modules are the same. Note that this result also holds more generally for sequences of homology groups that are linked by zig-zag sequences of homomorphisms.

4.3.1 Persistence in scale

Recall that $\text{Cover}_k(P, r) \subseteq \text{Cover}_k(P, s)$, whenever $r \leq s$, and similarly $\text{Del}_k(P, r) \subseteq \text{Del}_k(P, s)$, giving us the k -fold cover filtration and the order- k Delaunay filtrations in scale. Furthermore $\text{Cover}_k(P, r)$ is homotopy equivalent to $\text{Del}_k(P, r)$ by Lemma 4.3. This implies that their homology groups are isomorphic, however a-priori we do not have map that would provide the isomorphism between their homology groups, nor a map from $\text{Cover}_k(P, r)$ to $\text{Del}_k(P, r)$ or vice versa. Neither $\text{Cover}_k(P, r) \subseteq \text{Del}_k(P, r)$ nor $\text{Cover}_k(P, r) \supseteq \text{Del}_k(P, r)$, regardless of the scaling of we use in the definition of the order- k Delaunay mosaic. However we can find a map embedding $\text{Del}_k(P, r)$ into $\text{Cover}_k(P, r)$, as follows. We first construct the barycentric subdivision, $\text{Sd Del}_k(P, r)$, which is a simplicial complex. Each vertex $u \in \text{Sd Del}_k(P, r)$ represents a j -cell in $\text{Del}_k(P, r)$, which is dual to a $(d - j)$ -dimensional Voronoi polyhedron, and we map u to the center of mass of the intersection of this polyhedron with the k -fold cover. By construction, this intersection is non-empty and convex, so it contains the center of mass in its interior. After mapping all vertices, we map the other simplices of $\text{Sd Del}_k(P, r)$ by piecewise linear interpolation; see Figure 4.3.

The resulting map is injective and provides a homeomorphism from $\text{Del}_k(P, r)$ to




Figure 4.3: Subdivision of $\text{Del}_2(P, r)$ with its embedding in $\text{Cover}_2(P, r)$.

$\text{Sd Del}_k(P, r)$, which is embedded in $\text{Cover}_k(P, r)$. In fact, $\text{Sd Del}_k(P, r)$ is a deformation retract of $\text{Cover}_k(P, r)$: an explicit deformation retraction can be constructed in a similar vein to the deformation retraction from $\text{Cover}_1(P, r)$ to $\text{Del}_1(P, r)$ that was given in [32]. This implies that the map from $\text{Del}_k(P, r)$ to $\text{Cover}_k(P, r)$ we just constructed is a homotopy equivalence.

With these maps to link $\text{Del}_k(P, r)$ and $\text{Cover}_k(P, r)$, the following diagram commutes, with *he* indicating the homotopy equivalences:

$$\begin{array}{ccc} \text{Cover}_k(P, r) & \xrightarrow{\subseteq} & \text{Cover}_k(P, s) \\ \uparrow \text{he} & & \uparrow \text{he} \\ \text{Del}_k(P, r) & \xrightarrow{\subseteq} & \text{Del}_k(P, s). \end{array}$$

Therefore this diagram commutes after applying the homology functor, and the resulting vertical maps are isomorphisms between the corresponding homology groups. The Persistence Equivalence Theorem implies that the persistence diagrams of the k -fold cover filtration and the order- k Delaunay filtration for fixed k and varying r are the same.

With this, we have the necessary tools to provide an algorithm to compute persistence of k -fold covers for varying radii. Algorithm 1 computes $\text{Del}_k(P)$. Lemma 4.5 provides the recipe for computing the radii of the cells of $\text{Del}_k(P)$, which is implemented by Algorithm 2. Thus we obtain the sublevel set filtration of $\text{Del}_k(P)$, whose persistence module is isomorphic to the persistence of $\text{Cover}_k(P, r)$ for varying radius r . Finally, the persistence diagram is obtained from the filtration via the boundary matrix reduction algorithm.

We implementated this algorithm in 2 and 3 dimensions in `python` [69] and `C++` [70], using the PHAT library [10] for boundary matrix reduction.

4.4 Persistence in depth

Recall that there is not only an inclusion relation between k -fold covers for varying radii, but also for varying orders k . In this section we will present results allowing us to compute persistence of k -fold covers for varying k , which we refer to as *persistence in depth*. While according to Lemma 4.3 we can use $\text{Del}_k(P, r)$ as a proxy for $\text{Cover}_k(P, r)$, there are no straightforward inclusions or maps between the order- k Delaunay complexes $\text{Del}_k(P, r)$ for different k . Thus a novel approach is needed in order to compute persistence in depth. We will make use of the rhomboid tiling, and begin by introducing a radius function on it.

4.4.1 Rhomboid filtration

To shed additional light on the subcomplexes of the Delaunay mosaics, we introduce a discrete function on the collection of rhomboids discussed in Section 3.1. Calling it the *squared radius function*, $\mathcal{R}^2: \text{Rho}(P) \rightarrow \mathbb{R}$, we define it by remembering that each j -dimensional rhomboid, $\rho \in \text{Rho}(P)$, corresponds to a $(d + 1 - j)$ -dimensional cell, $\rho^* \in \text{Arr}(P)$. Decomposing a point of the cell into its first d coordinates and its $(d + 1)$ -st coordinate, we write $y = (x, z) \in \mathbb{R}^d \times \mathbb{R}$, and we define $\mathfrak{r}(y) = \|x\|^2 - 2z$. With this notation, we define the squared radius function by mapping ρ to the minimum value of any point in its dual cell:

$$\mathcal{R}^2(\rho) = \min_{y \in \rho^*} \mathfrak{r}(y). \quad (4.3)$$

By convention, the value of the vertex that corresponds to the ordered three-partition $P = (\emptyset, \emptyset, P)$ is $\mathcal{R}^2(0) = -\infty$. To obtain a geometric interpretation of this construction, recall the paraboloid \mathcal{P} in \mathbb{R}^{d+1} which is defined by the equation $z = \frac{1}{2}\|x\|^2$. We introduce $\mathcal{P}_t(x): \mathbb{R}^d \rightarrow \mathbb{R}$ defined by $\mathcal{P}_t(x) = \frac{1}{2}(\|x\|^2 - t)$, and also use \mathcal{P}_t to refer to the graph of this function. This graph \mathcal{P}_t is the original paraboloid \mathcal{P} dropped down vertically by a distance $\frac{t}{2}$. With this notation, $\mathcal{R}^2(\rho)$ is the minimum t such that \mathcal{P}_t has a non-empty intersection with ρ^* , see Figure 4.4. This implicitly defines the *radius function* \mathcal{R} , and

while strictly speaking \mathcal{R} can take imaginary values, in the unweighted setting this is only of concern for the rhomboid that is the origin. In the weighted setting however this can apply to arbitrary rhomboids, and thus it is better to talk about the squared radius function \mathcal{R}^2 . Similarly, as in the weighted setting it makes sense to talk about spheres with negative square radius, it may be preferable to also talk about k -fold covers of balls of a given square radius and the square radius function \mathcal{R}_k^2 for Delaunay mosaics, in order to maintain generality.




Figure 4.4: The rhomboid radius function \mathcal{R}^2 in terms of the paraboloid \mathcal{P}_t sweeping the arrangement.

Clearly, \mathcal{R}^2 is *monotonic*, that is: $\mathcal{R}^2(\rho) \leq \mathcal{R}^2(\varrho)$ if ρ is a face of ϱ . Indeed, if ρ is a face of ϱ , then ϱ^* is a face of ρ^* , which implies that the paraboloid touches ρ^* at the same time or before it touches ϱ^* when dropped. It follows that the sublevel sets of the radius function are subcomplexes of the rhomboid tiling.

Relation to Delaunay radius function. Recall that the order- k Delaunay mosaic of P is the horizontal slice of the rhomboid tiling at depth k . In other words, every cell of $\text{Del}_k(P)$ is the horizontal slice of a rhomboid, and we used $\rho := \rho(\sigma)$ to denote the unique lowest-dimensional rhomboid that σ is a slice of.

It turns out that the previously defined radius function on the order- k Delaunay mosaic, $\mathcal{R}_k: \text{Del}_k(P) \rightarrow \mathbb{R}$, coincides with the restriction of \mathcal{R} to the horizontal slice, as we will

now show.

Lemma 4.6 (Delaunay Radius Function). *Let $P \subseteq \mathbb{R}^d$ be finite and in general position. For every integer $k \geq 1$ and every cell $\sigma \in \text{Del}_k(P)$, we have $\mathcal{R}_k^2(\sigma) = \mathcal{R}^2(\rho(\sigma))$.*

Proof. The paraboloid \mathcal{P}_t intersects \mathbb{R}^d in the sphere with squared radius t . More generally, the paraboloid intersects every d -plane tangent to $\mathcal{P} = \mathcal{P}_0$ in an ellipsoid whose vertical projection to \mathbb{R}^d is a sphere with squared radius t . Dropping the paraboloid vertically thus translates into growing balls simultaneously and uniformly centered at the points in P . By definition, $\mathcal{R}^2(\rho)$ is the value t_0 of t for which the paraboloid touches the dual cell, $\rho^* \in \text{Arr}(P)$, for the first time. More formally, the set of points $q \in \rho^*$ that lie on or above the graph of \mathcal{P}_t is empty for all $t < t_0$ and non-empty for all $t \geq t_0$.

Let $\rho := \rho(\sigma)$; let σ^* be the vertical projection of ρ^* to \mathbb{R}^d , and assume it is a polyhedron in some Voronoi tessellation of P . It belongs to $\text{Vor}_k(P)$ iff its dual cell, σ , belongs to $\text{Del}_k(P)$. In that case $\sigma^* \cap \text{Cover}_k(P, r)$ is empty for all $r < r_0$ and non-empty for all $r \geq r_0$, in which $r_0^2 = t_0 = \mathcal{R}^2(\rho)$. By definition, σ belongs to $\text{Del}_k(P, r)$ iff this intersection is non-empty, which implies $\mathcal{R}_k^2(\sigma) = \mathcal{R}^2(\rho)$, as required. \square

The equivalence of the previously defined order- k Delaunay radius function and the rhomboid radius function restricted to slices allows us to give an alternative characterization of the rhomboid radius function, akin to the one from Lemma 4.4 for order- k Delaunay mosaics.

In the same vein as for order- k Delaunay cells, we say a sphere is ρ -constrained if $\text{rho}(S) = \rho$, and denote the smallest such sphere as $S_{\min}(\rho)$. In other words, $S_{\min}(\rho)$ is the smallest enclosing sphere of $P_{\text{in}}(\rho)$ that has $P_{\text{on}}(\rho)$ on its surface and no other points from P in its interior.

Now if σ is an order- k Delaunay cell and $\rho := \rho(\sigma)$ the rhomboid of σ , then by Lemma 4.6 ρ and σ have the same radius value. Then it follows from Lemma 4.4 that the radius value of ρ is the radius of $S_{\min}(\sigma)$, which is the same as the radius of $S_{\min}(\rho)$ because $P_{\text{in}}(\rho) = P_{\text{in}}(\sigma)$, $P_{\text{on}}(\rho) = P_{\text{on}}(\sigma)$ and $P_{\text{out}}(\rho) = P_{\text{out}}(\sigma)$. While not all rhomboids ρ have such a slice σ , in particular any 1-dimensional rhomboids, the observation nevertheless generalizes:

Lemma 4.7. $\mathcal{R}^2(\rho)$ equals the squared radius of $S_{\min}(\rho)$.

Proof. Recall the correspondence between points in the hyperplane arrangement and spheres that was used in the proof of Claim 1 from Theorem 3.1. A point $y = (x, z)$ below the paraboloid \mathcal{P} was mapped to the sphere S with center x and squared radius $\|x\|^2 - 2z$. Crucially, we observed that if γ is the unique arrangement cell in whose interior y is contained in, then $\text{In}(S) = P_{\text{above}}(\gamma)$, $\text{On}(S) = P_{\text{contains}}(\gamma)$, and $\text{Out}(S) = P_{\text{below}}(\gamma)$.

Now let $t = \mathcal{R}^2(\rho)$, and let r^2 be the squared radius of $S_{\min}(\rho)$. Then by definition t is the smallest value for which \mathcal{P}_t contains a point $y \in \rho^*$. The aforementioned map maps y to a ρ -constrained sphere, thus $r^2 \leq t$. Conversely, when reversing this map, $S_{\min}(\rho)$ is mapped to a point $y \in \rho^*$. As t was the smallest value for which \mathcal{P}_t touches ρ^* , we have $t \leq r^2$. Thus the squared radius of $S_{\min}(\rho)$ equals $\mathcal{R}^2(\rho)$. \square

Complexity of the rhomboid tiling. Recall that, intuitively, the radius value of a rhomboid ρ is the first point in time when the paraboloid, as we drop it down through the hyperplane arrangement, encounters the dual cell ρ^* . This idea gives us convenient means to count the number of rhomboids in the rhomboid tiling, giving the result below. While this result has been long known, we still prove it here for completeness (also see e.g. [31, Section 1.2]).

Proposition 4.8. *Let P be a set of n points in general position in \mathbb{R}^d , and let $\Gamma_i^{d+1}(n)$ be the number of i -faces in $\text{Arr}(P)$. Then the number of j -dimensional rhomboids is $\Gamma_{d+1-j}^{d+1}(n) \leq \frac{(n+1)^{d+1}}{j!(d+1-j)!}$ and does not depend on P . The total number of rhomboids is at most $\frac{2^{d+1}}{(d+1)!}(n+1)^{d+1}$.*

Proof. The number of rhomboids or, equivalently, the cells in the dual hyperplane arrangement, is maximized when the n hyperplanes are in general position, and then they depend only on n and d . Observe first that for every $0 \leq i \leq d+1$, there are $\binom{n}{d+1-i}$ i -planes, each the common intersection of $d+1-i$ hyperplanes. The paraboloid \mathcal{P}_t used in the definition of \mathcal{R}^2 sweeps out the arrangement and encounters a new chamber whenever it first intersects one of the i -planes, for $0 \leq i \leq d+1$. This means that there is one chamber for each plane, which implies that the number of chambers in the arrangement, and equivalently the number of vertices in the rhomboid tiling, is

$$\Gamma_{d+1}^{d+1}(n) = \binom{n}{d+1} + \binom{n}{d} + \dots + \binom{n}{0} \leq \frac{(n+1)^{d+1}}{(d+1)!}. \quad (4.4)$$

The inequality on the right-hand side in (4.4) is easy to prove, by induction or otherwise. To count the i -cells in the arrangement, we observe that each i -plane carries an arrangement of $n - (d + 1 - i)$ $(i - 1)$ -planes. We get the number of (i -dimensional) chambers in this arrangement from (4.4), and multiplying with the number of i -planes, we get the number of i -cells:

$$\begin{aligned}\Gamma_i^{d+1}(n) &= \binom{n}{d+1-i} \Gamma_i^i(n - d - 1 + i) \\ &\leq \frac{n^{d+1-i}}{(d+1-i)!} \frac{(n+1)^i}{i!} \leq \frac{(n+1)^{d+1}}{(d+1-i)! i!}.\end{aligned}\tag{4.5}$$

Writing $j = d - i$, we get a $(j + 1)$ -rhomboid in $\text{Rho}(P)$ for every i -cell in the arrangement. In other words, (4.5) counts the $(j + 1)$ -rhomboids in the rhomboid tiling.


$$\#\text{cells} = \sum_{i=0}^{d+1} \Gamma_i^{d+1}(n) \leq \sum_{i=0}^{d+1} \frac{(n+1)^{d+1}}{(d+1-i)! i!} = \frac{2^{d+1}}{(d+1)!} (n+1)^{d+1}.\tag{4.6}$$

□


4.4.2 Computation

In this section, we develop an algorithm that computes the persistence of the nested sequence of multi-covers in depth (4.2). We follow the usual strategy of substituting a complex for each cover, but there are complications; for example, the vertex sets of consecutive Delaunay mosaics are different, see Figures 4.5a and 4.5c. We represent $\text{Cover}_k(P, r)$ by $\text{Del}_k(P, r)$ and we introduce additional complexes between contiguous Delaunay mosaics to realize the inclusion between the covers.


Half-integer slices. There are generally no convenient maps connecting $\text{Del}_k(P)$ with $\text{Del}_{k-1}(P)$. To finesse this difficulty, we use the horizontal half-integer slices $\text{Del}_\ell(P)$ of the rhomboid tiling at depth $\ell = k - \frac{1}{2}$ for $k \geq 1$. Recall that $\text{Del}_\ell(P)$ is dual to degree- k Voronoi tessellation $\text{Vor}_\ell(P)$, which is a refinement of the order- k Voronoi tessellation and thus also decomposes $\text{Cover}_k(P, r)$ into convex pieces. We refer to the decomposition of $\text{Cover}_k(P, r)$ by $\text{Vor}_\ell(P)$ as $\text{Vor}_\ell(P, r)$, see Figure 4.5b.



(a) The black order-2 Voronoi tessellation decomposes the 2-fold cover (stronger pink) into convex pieces. The corresponding subcomplex of the dual order-2 Delaunay mosaic is superimposed in blue.



(b) The 2-fold cover (stronger pink) with the 3-fold cover highlighted in pale red. In black, the degree-3 Voronoi tessellation, which is the superposition of the order-2 and order-3 Voronoi tessellations. It decomposes the 3-fold cover into convex pieces. In blue, the dual of this decomposition, $D_{2.5} := \text{Del}_{2.5}(P, r)$. The additional green cells are part of $E_{2.5} \supseteq D_{2.5}$.



(c) Like Figure 4.5a, but with 3-fold cover, order-3 Voronoi tessellation and dual $\text{Del}_3(P, r)$.

Figure 4.5: Delaunay mosaics and Voronoi decompositions of the k -fold cover for different k , for six points P in the plane with a pink ball of radius r centered at each.

There are natural piecewise linear maps from $\text{Del}_\ell(P)$ to $\text{Del}_k(P)$ and to $\text{Del}_{k-1}(P)$:

$$\cdots \rightarrow \text{Del}_{k-1}(P) \leftarrow \text{Del}_\ell(P) \rightarrow \text{Del}_k(P) \leftarrow \cdots$$

Specifically, we get $\text{Del}_\ell(P) \rightarrow \text{Del}_k(P)$ by mapping the vertices dual to the regions decomposing $\text{dom}(Q) \in \text{Vor}_k(P)$ to the vertex dual to $\text{dom}(Q)$. Symmetrically, we get $\text{Del}_\ell(P) \rightarrow \text{Del}_{k-1}(P)$. However, because such maps lead to complications in the persistence algorithm, we will construct a sequence of complexes linked by inclusions instead.

We use the horizontal slabs of the rhomboid tiling to connect the mosaics via inclusions. To formally define them, write H_ℓ^k for the points in \mathbb{R}^{d+1} that lie on or between H_ℓ and H_k . We define *slab mosaics* as intersections of such slabs with the rhomboid tiling. Analogous to $\text{Del}_k(P, r)$, we also define radius-dependent subcomplexes of these slab mosaics, as well as of half-integer mosaics:

$$\text{Del}_\ell^k(P, r) = \{\rho \cap H_\ell^k : \mathcal{R}^2(\rho) \leq r^2\}, \quad (4.7)$$

$$\text{Del}_\ell(P, r) = \{\rho \cap H_\ell : \mathcal{R}^2(\rho) \leq r^2\}, \quad (4.8)$$

$$\text{Del}_{k-1}^\ell(P, r) = \{\rho \cap H_{k-1}^\ell : \mathcal{R}^2(\rho) \leq r^2\}. \quad (4.9)$$

To simplify the notation, we fix r and write $D_k = \text{Del}_k(P, r)$, $C_k = \text{Cover}_k(P, r)$, etc. The half-integer Delaunay mosaic includes in both slab mosaics, D_k includes in the first, and D_{k-1} includes in the second, giving the following zig-zag sequence of inclusions:

$$\cdots \supseteq D_{k-1} \subseteq D_{k-1}^\ell \supseteq D_\ell \subseteq D_\ell^k \supseteq D_k \subseteq \cdots$$

Figure 4.6 provides an example. We note an important difference between the two slabs: $\text{Vor}_\ell(P, r)$ and $\text{Vor}_k(P, r)$ are different convex subdivisions of the same space, C_k , which implies that D_ℓ and D_k have the same homotopy type. Indeed, this is also the homotopy type of D_ℓ^k , and there are natural deformation retractions to D_ℓ and D_k . In contrast, D_{k-1} and D_ℓ have generally different homotopy types, and there is a deformation retraction from D_{k-1}^ℓ to D_{k-1} but not necessarily to D_ℓ ; see again Figure 4.6. To remedy this deficiency, we introduce mosaics that contain D_ℓ and D_{k-1}^ℓ as subcomplexes. To construct them, we recall that $\text{Vor}_\ell(P)$ is a refinement of $\text{Vor}_{k-1}(P)$, which implies that the polyhedra of




Figure 4.6: A sublevel set of the 3-dimensional rhomboid tiling whose slices are the mosaics of Figure 4.5. From *top* to *bottom*: D_2 in dark blue, $D_{2.5}$ in purple, and D_3 in dark red, with slabs connecting adjacent slices.

$\text{Vor}_\ell(P)$ intersect C_{k-1} in convex sets. We let E_ℓ be the dual of this convex decomposition of the $(k-1)$ -fold cover. Since $C_k \subseteq C_{k-1}$, we indeed have $D_\ell \subseteq E_\ell$; see Figure 4.5b. Furthermore, we let E_{k-1}^ℓ be the maximal subcomplex of $\text{Rho}(P) \cap H_{k-1}^\ell$ whose boundary complexes at depths $k-1$ and ℓ are D_{k-1} and E_ℓ . Clearly, D_{k-1}^ℓ is a subcomplex of E_{k-1}^ℓ , and because D_{k-1} and E_ℓ are deformation retracts of E_{k-1}^ℓ , these three mosaics have the same homotopy type. We will use these relations shortly in the computation of the persistence diagram of the filtration of multi-covers in depth (4.2).

Connecting the spaces. To prepare the construction of the persistence and zigzag modules, we connect the multi-covers and the corresponding Delaunay mosaics with maps. Fixing $r \geq 0$ and setting $\ell = k - \frac{1}{2}$, as before, we consider the following diagram in which identities and homotopy equivalences are marked as such:

$$\begin{array}{ccccccccccc}
 \longrightarrow & C_k & \xrightarrow{\text{id}} & C_k & \xrightarrow{\text{id}} & C_k & \xrightarrow{\text{id}} & C_k & \xrightarrow{\text{id}} & C_k & \longrightarrow & C_{k-1} & \xrightarrow{\text{id}} \\
 & \uparrow \text{he} & & \uparrow \text{he} & & \uparrow \text{he} & & \uparrow \text{he} & & \uparrow \text{he} & & \uparrow \text{he} & \\
 \longrightarrow & E_{\ell+1} & \xrightarrow{\text{he}} & E_k^{\ell+1} & \xleftarrow{\text{he}} & D_k & \xrightarrow{\text{he}} & D_\ell^k & \xleftarrow{\text{he}} & D_\ell & \longrightarrow & E_\ell & \xrightarrow{\text{he}}
 \end{array}$$

The top row stretches out the filtration by writing each multi-cover five times and connecting the copies with the identity. The remaining maps in this row are inclusions. The bottom

row contains the slice mosaics at integer and half-integer depths, and connects them with inclusions, using slab mosaics as intermediaries. As argued above, the first five mosaics all have the same homotopy type, and the inclusion maps between them are homotopy equivalences.

As the vertical map from D_k to C_k , we use the homotopy equivalence we constructed in Section 4.3.1, which intuitively embeds the barycentric subdivision $\text{Sd } D_k$ into C_k . Recall that D_ℓ is dual to $\text{Vor}_\ell(P, r)$, which is another convex decomposition of the k -fold cover. We therefore get the vertical map from D_ℓ to C_k the same way, first constructing $\text{Sd } D_\ell$ and second mapping the vertices to centers of mass. This is again a homotopy equivalence. Similarly, E_ℓ is dual to the convex decomposition of C_{k-1} with $\text{Vor}_\ell(P)$. As before, we get the vertical map by sending the vertices of E_ℓ to centers of mass, but we distinguish between two cases. If a polyhedron of $\text{Vor}_\ell(P)$ has a non-empty intersection with C_k , we send the corresponding vertex of $\text{Sd } E_\ell$ to the center of mass of this intersection. If, however, the intersection with C_k is empty but the intersection with C_{k-1} is non-empty, then we send the vertex to the center of mass of the latter. This ensures that the geometric embedding of $\text{Sd } D_\ell$ is contained in the geometric embedding of $\text{Sd } E_\ell$.

To finally map the slab mosaics, we first deformation retract them to slice mosaics and then map them reusing the barycentric subdivisions. Here we make arbitrary choices, mapping $E_k^{\ell+1}$ to $E_{\ell+1}$ to C_k and mapping D_ℓ^k to D_k to C_k . Note that all vertical maps are homotopy equivalences, as marked in the above diagram.

Modules. Applying the homology functor for a fixed coefficient field, we map all multi-covers and mosaics to vector spaces and all maps to homomorphisms (linear maps) between them. We denote the persistence module consisting of the top row of vector spaces with homomorphisms from left to right as $\text{MC}(r)$. The bottom row of vector spaces are connected by homomorphisms going from left to right or from right to left. This kind of structure is referred to as a *zigzag module*, and we denote it $\text{ZZ}(r)$. The advantage of the zigzag over the persistence module is that its maps are induced by inclusions between complexes, which lend themselves to computations. Our goal, however, is to compute the persistence diagram of $\text{MC}(r)$, and we do this by using $\text{ZZ}(r)$ as a proxy. The following result is therefore essential.

Lemma 4.9 (Isomorphism of Modules). *Let $P \subseteq \mathbb{R}^d$ be locally finite and in general*

position. Then the persistence diagrams of $\text{MC}(r)$ and of $\text{ZZ}(r)$ are the same for every $r \geq 0$.

Proof. Write C_k, D_k, E_k for the vector spaces obtained by applying the homology functor to C_k, D_k, E_k , etc. The goal is to show that the diagram of multi-covers and mosaics maps to a diagram of vector spaces in which all squares commute and most maps are isomorphisms:

$$\begin{array}{ccccccccccc}
 & \longrightarrow & C_k & \xrightarrow{\text{iso}} & C_k & \xrightarrow{\text{iso}} & C_k & \xrightarrow{\text{iso}} & C_k & \xrightarrow{\text{iso}} & C_k & \longrightarrow & C_{k-1} & \xrightarrow{\text{iso}} & \longrightarrow \\
 & & \uparrow \text{iso} & & \uparrow \text{iso} & & \uparrow \text{iso} & & \uparrow \text{iso} & & \uparrow \text{iso} & & \uparrow \text{iso} & & \\
 & \longrightarrow & E_{\ell+1} & \xrightarrow{\text{iso}} & E_k^{\ell+1} & \xleftarrow{\text{iso}} & D_k & \xrightarrow{\text{iso}} & D_\ell^k & \xleftarrow{\text{iso}} & D_\ell & \longrightarrow & E_\ell & \xrightarrow{\text{iso}} & \longrightarrow
 \end{array}$$

To prove commutativity, we consider the five squares shown in the above diagram. The first square commutes already before applying the homology functor, and so does the third square. Similarly, the fifth square commutes because the image of $\text{Sd } D_\ell$ in C_k includes in the image of $\text{Sd } E_\ell$ in C_{k-1} .

The second and fourth squares do not commute before applying the homology functor, but we argue they do after applying the functor. The two cases are similar, so we focus on the fourth square. Recall that $\text{Vor}_k(P, r)$ and $\text{Vor}_\ell(P, r)$ are two convex decompositions of the same space, which is C_k , and that $\text{Vor}_\ell(P, r)$ is a refinement of $\text{Vor}_k(P, r)$. D_k and D_ℓ are dual to these decompositions, with one or more vertices of D_ℓ corresponding to every one vertex of D_k . When we map D_ℓ to D_ℓ^k to C_k , the full subcomplex with these vertices is first contracted to the single vertex by the deformation retraction from D_ℓ^k to D_k , and second it is mapped to the center of mass of the corresponding domain in $\text{Vor}_k(P, r)$. In contrast, when we map D_ℓ to C_k directly, all these vertices map to different points in C_k , but all these points lie in the interior of the same domain in $\text{Vor}_k(P, r)$. Indeed, the full subcomplex with these vertices is dual to a convex decomposition of this domain and therefore contractible. It follows that the fourth square of homomorphisms commutes. Similarly, the second square commutes, and therefore all squares commute.

Isomorphisms are reversible, so we can draw them from left to right in the bottom row of the diagram. The result are two parallel persistence modules whose vector spaces are connected by isomorphisms. The Persistence Equivalence Theorem of [33, page 159] implies that the two modules have the same persistence diagram. \square

Algorithm and running time. We compute the persistence diagram of the filtration of multi-covers in depth (4.2) using the zigzag algorithm generically described in [13] and explained in detail for inclusion maps in [14]. Its worst-case running time is cubic in the input size, which is the total number of cells in the mosaics.

To count the cells, we assume a finite number of points in \mathbb{R}^d , $n = \#P$. Recall Proposition 4.8, which states that the number of j -rhomboids in $\text{Rho}(P)$ is $\Gamma_{d-j}^{d+1}(n) \leq \frac{(n+1)^{d+1}}{(d+1-j)! j!}$. All cells of our zig-zag filtration are horizontal slices or horizontal slabs of rhomboids in \mathbb{R}^{d+1} .

In particular, we have $\Gamma_{d+1}^{d+1}(n)$ vertices in the tiling. For $0 \leq j \leq d$, the interior of every $(j+1)$ -rhomboid has a non-empty intersection with $2j+1$ hyperplanes H_ℓ , in which 2ℓ is an integer. The $(j+1)$ -rhomboid thus contributes $2j+1$ j -cells to the Delaunay mosaics and $2j+2$ $(j+1)$ -prisms to the slab mosaics. Taking the sum over all dimensions, we get the total number of cells in the mosaics used in the construction of the zigzag module:

$$\begin{aligned} \#\text{cells} &= \Gamma_{d+1}^{d+1}(n) + \sum_{j=0}^d (4j+3) \Gamma_{d-j}^{d+1}(n) \\ &\leq \frac{(n+1)^{d+1}}{(d+1)!} + \sum_{i=0}^d 4(d+1-i) \frac{(n+1)^{d+1}}{(d+1-i)! i!} \\ &\leq \frac{(n+1)^{d+1}}{(d+1)!} + 4(n+1)^{d+1} \sum_{i=0}^d \frac{1}{(d-i)! i!} \leq 9(n+1)^{d+1}. \end{aligned} \quad (4.10)$$

Taking the third power, we get an upper bound for the worst-case running time of the algorithm and thus the main result of this section.

Theorem 4.10 (Multi-cover Persistence). *Let P be a set of n points in general position in \mathbb{R}^d . For every radius $r \geq 0$, the persistence diagram of the filtration of multi-covers with radius r can be computed in worst-case time $O(n^{3d+3})$.*

4.5 2-parameter persistence

Two-parameter persistence was introduced in [15], providing a 2-dimensional analog of persistence modules, called bipersistence modules. While the notion of persistence diagram does not generalize to bipersistence modules, one can explore the persistence diagrams of any of its 1-dimensional slices, a process aided by e.g. the software library

RIVET [83]. Furthermore other invariants of bipersistence modules have been proposed, such as multiparameter persistence landscapes [87].

For the multi-covers, we have inclusions $\text{Cover}_k(P, r) \subseteq \text{Cover}_\ell(P, s)$ whenever $r \leq s$ and $\ell \leq k$. Thus they form a bifiltration in two parameters, $r \in \mathbb{R}$ and $k \in \mathbb{Z}$ with $1 \leq k \leq \#P$, as illustrated below for $r_1 < r_2 < r_3$ and $k = 1, 2, 3$:

$$\begin{array}{ccccccc}
 & \vdots & & \vdots & & \vdots & \\
 & \downarrow & & \downarrow & & \downarrow & \\
 \dots & \hookrightarrow & \text{Cover}_3(P, r_1) & \hookrightarrow & \text{Cover}_3(P, r_2) & \hookrightarrow & \text{Cover}_3(P, r_3) \hookrightarrow \dots \\
 & \downarrow & & \downarrow & & \downarrow & \\
 \dots & \hookrightarrow & \text{Cover}_2(P, r_1) & \hookrightarrow & \text{Cover}_2(P, r_2) & \hookrightarrow & \text{Cover}_2(P, r_3) \hookrightarrow \dots \\
 & \downarrow & & \downarrow & & \downarrow & \\
 \dots & \hookrightarrow & \text{Cover}_1(P, r_1) & \hookrightarrow & \text{Cover}_1(P, r_2) & \hookrightarrow & \text{Cover}_1(P, r_3) \hookrightarrow \dots
 \end{array}$$

By applying the homology functor \mathbf{H}_* to this diagram, we get a *bipersistence module* with two parameters r and k . Note that it contains the persistence modules in scale in horizontal direction and in depth in vertical direction, but allows us to look at both aspects at once.

As in the 1-dimensional case, we call two bipersistence modules with parameters r and k isomorphic if there are point-wise isomorphisms between their homology groups for each r and k , and these maps commute with the maps of the bipersistence modules. If two bipersistence modules are isomorphic, then their aforementioned invariants are the same, in particular the persistence diagrams of any of their corresponding 1-dimensional submodules.

In a recent result, it was shown [24] that there is a simplicial complex admitting a bifiltration whose bipersistence module is isomorphic to the multi-cover bipersistence module. Let $n := \#P$, $1 \leq k \leq n$ and $r \in \mathbb{R}$. Extending the Delaunay Nerve, which we recall as $\text{DelN}_k(P, r) := \text{Nerve} \{ \text{dom}(Q, r) : Q \subseteq P, \#Q = k \}$ the authors define

$$\text{DelN}_k^{k+1}(P, r) := \text{Nerve} \{ \text{dom}(Q, r) : Q \subseteq P, \#Q = k \text{ or } \#Q = k + 1 \}$$

and construct a new complex by stacking these together, as follows:

$$\text{DelN}_k^\ell(P, r) := \bigcup_{j=k}^{\ell-1} \text{DelN}_j^{j+1}(P, r)$$

Then the main result is the following:

Theorem 4.11 ([24]). *The bipersistence modules $\mathbf{H}_*(\text{DelN}_k^n(P, r))$ and $\mathbf{H}_*(\text{Cover}_k(P, r))$ in parameters $r \in \mathbb{R}$ and $k \in \mathbb{Z}$ with $1 \leq k \leq n$ are isomorphic.*

One of the key insights of the proof is that $\text{DelN}_k^n(P, r)$ has the same homology as $\text{DelN}_k(P, r)$. An explicit algorithm with implementation for computing these complexes $\text{DelN}_k^n(P, r)$ in 2 dimensions is provided by the authors. By relating these complexes to the rhomboid tiling, we will now show that our algorithms and implementations for the rhomboid tiling and its radius function can be used to compute these complexes in arbitrary dimensions.

Slab complexes. Recall that we defined $H_k := \{(x_1, \dots, x_{d+1}) \in \mathbb{R}^{d+1} : x_{d+1} = k\}$ as the horizontal hyperplane at depth k , and H_k^ℓ as the points in \mathbb{R}^{d+1} that lie on or between H_k and H_ℓ . We observed in Theorem 3.1 that $\text{Del}_k(P, r)$ is the (cell-wise) intersection of $\text{Rho}(P, r)$ with H_k . In a similar vein, we define $\text{Del}_k^{k+1}(P, r)$ as the (cell-wise) intersection of $\text{Rho}(P, r)$ with H_k^{k+1} . Inspired by [24] and in full analogy to $\text{DelN}_k^\ell(P, r)$, for integers $1 \leq k < \ell \leq n$ we stack these complexes to define

$$\text{Del}_k^\ell(P, r) := \bigcup_{j=k}^{\ell-1} \text{Del}_j^{j+1}(P, r).$$

Figure 4.7 shows an example. Note that this is not the same as the intersection of H_k^ℓ with $\text{Rho}(P)$, because in $\text{Del}_k^\ell(P, r)$ the rhomboids are cut along all planes at integer depth. As before, we write $\text{DelN}_k^\ell(P)$ and $\text{Del}_k^\ell(P)$ for the corresponding complexes for infinite radius r , and call the smallest r such that a cell appears in $\text{DelN}_k^\ell(P, r)$ or $\text{Del}_k^\ell(P, r)$ respectively its *radius value*.

Note that for a simplex $S \in \text{DelN}_k^{k+1}(P)$ all combinatorial vertices correspond to subsets of P , and we can define $P_{in}(S) = \bigcap_{v \in S} v$, $P_{on}(S) = \bigcup_{v \in S} v \setminus P_{in}(S)$, $P_{out}(S)$, S -constrained sphere and $S_{min}(S)$ like we do for rhomboids or Delaunay cells. Then we have the following relationship between $\text{DelN}_k^{k+1}(P, r)$ and $\text{Del}_k^{k+1}(P, r)$ which naturally translates to the same relationship between $\text{DelN}_k^\ell(P, r)$ and $\text{Del}_k^\ell(P, r)$ for integers $k \leq \ell$:




Figure 4.7: At the top, the point set $P = \{a, b, c, d, e\} \subset \mathbb{R}^1$ with balls of radius r around them. Below in solid blue, the sublevel set $\text{Rho}(P, r)$ of the rhomboid tiling with the rhomboids sliced apart to give the stacking of the slab mosaics. Note how here $\text{Del}_k^n(P, r)$ with $n = 5$ is homotopy equivalent to $\text{Cover}_k(P, r)$ above.

Lemma 4.12. 1. For every cell σ from $\text{Del}_k^{k+1}(P)$, there is a simplex S in $\text{DelN}_k^{k+1}(P)$ with $S := V(\sigma)$. S and σ have the same radius value.

2. Conversely, for every simplex S from $\text{DelN}_k^{k+1}(P)$ there is a unique lowest-dimensional cell σ in $\text{Del}_k^{k+1}(P)$ whose vertex set is a superset of the vertex set of S . This cell σ is the intersection of the slab H_k^{k+1} with the rhomboid ρ defined by $P_{in}(\rho) = P_{in}(S)$, $P_{on}(\rho) = P_{on}(S)$ and $P_{out}(\rho) = P_{out}(S)$. S and σ have the same radius value.

3. In particular, there is a bijection between the maximal cells of $\text{Del}_k^{k+1}(P, r)$ and $\text{DelN}_k^{k+1}(P, r)$.

Proof. (1) Let $S := V(\sigma)$ be the set of combinatorial vertices of σ . Because $V(\sigma) \subseteq V(\rho)$ for some rhomboid ρ , $S_{min}(\sigma)$ exists, and let x be its center. To show that S is a simplex in $\text{DelN}_k^{k+1}(P)$, it is sufficient to show the domains $\text{dom}(Q)$ intersect for all $Q \in S$; in particular we will show that $x \in \text{dom}(Q)$ for all $Q \in S$. As any $Q \in S$ has $\#Q = k$ or $\#Q = k + 1$ and is the union of $P_{in}(\sigma)$ with some subset of $P_{on}(\sigma)$, Q is a set of k or $k + 1$ closest points for x . Thus $x \in \text{dom}(Q)$ for all $Q \in S$. The fact that S and σ have the

same radius value follows from (2).

(2) S is a set of subsets $Q \subseteq P$ such that all $\text{dom}(Q)$ for $Q \in S$ have a non-empty intersection. As furthermore $\#Q = k$ or $\#Q = k + 1$ for each Q , it follows that the chambers in the hyperplane arrangement $\text{Arr}(P)$ corresponding to the sets Q also intersect.

Recall that the chamber corresponding to Q is the intersection of h_q^- for $q \in Q$ and h_p^+ for $p \in P \setminus Q$. Then the intersection of these chambers for all $Q \in S$ yields a cell that is the intersection of h_p^- for $p \in P_{in}(S)$, h_p for $p \in P_{on}(S)$, and h_p^+ for $p \in P_{out}(S)$. This cell ρ^* is dual to the rhomboid $\rho \in \text{Rho}(P)$ with $P_{in}(S) = P_{in}(\rho)$, $P_{on}(S) = P_{on}(\rho)$ and $P_{out}(S) = P_{out}(\rho)$. In particular S is a subset of the vertices of ρ due to duality. Due to $\#Q = k$ or $\#Q = k + 1$ for $Q \in S$ we further have that S is a subset of $V(\sigma)$ for $\sigma := \rho \cap H_k^{k+1}$.

If S (and thus σ) consist of a single vertex, it is easy to see that they have the same radius value. If not, we know that $P_{on}(S) \neq \emptyset$. Now the projection δ^* of ρ^* is exactly the set of points $x \in \mathbb{R}^d$ for which $\|p - x\| =: r(x)$ is the same for all $p \in P_{on}(S)$, and $\|p - x\| \leq r(x)$ for $p \in P_{in}(S)$ and $\|p - x\| \geq r(x)$ for $p \in P_{out}(S)$. Thus any S -constrained sphere must have its center in δ^* . Let x be a point in δ^* that minimizes $r := r(x)$. Then r is the radius value of ρ and thus σ . However, r is also the smallest radius such that $\text{dom}(Q, r), Q \in S$ intersect: All balls $B_r(p)$ for $p \in P_{in}(S) \cup P_{on}(S)$, and thus $p \in Q$ for all $Q \in S$, contain x . For any smaller radius $r' < r$, the balls $B_{r'}(p)$ for $p \in P_{on}(S)$ (and $p \in P_{out}(S)$) will not intersect δ^* , leaving at most $|P_{in}(S)| < \max\{\#Q : Q \in S\}$ balls that intersect δ^* . In particular, $\text{dom}(Q, r')$ does not intersect δ^* for those Q of size $\max\{\#Q : Q \in S\}$.

(3) is a direct consequence of (1) and (2).

□

This allows us to compute $\text{DelN}_k^n(P, r)$ straight-forwardly from $\text{Del}_k^n(P, r)$. Therefore our dimension-agnostic implementation [70] can be employed to compute the simplicial bifiltration of $\text{DelN}_k^n(P, r)$ whose bipersistence module is isomorphic to the bipersistence module of the k -fold covers.

In fact, our lemma implies that there is a bijection of the maximal cells of $\text{DelN}_k^\ell(P, r)$ and $\text{Del}_k^\ell(P, r)$, and the conditions of Lemma 4.2 apply. This means that the two complexes

are homotopy equivalent, and in particular that $\text{Del}_k^n(P, r)$ and $\text{Cover}_k(P, r)$ have the same homology, see Figure 4.7.

This insight sparked a collaboration with Corbet, Kerber and Lesnick, the original authors of [24], to extend their manuscript by relating the bipersistence modules of their nerve construction $\text{DelN}_k^n(P, r)$ and the sliced rhomboid tiling $\text{Del}_k^n(P, r)$. Using a persistent version of the nerve theorem, the main new result is that the bipersistence modules of these two are in fact isomorphic. This in particular implies that the bipersistence modules of the k -fold covers and $\text{Del}_k^n(P, r)$ are isomorphic as well, and therefore the sliced rhomboid tiling can be used directly to compute 2-parameter persistence of k -fold covers. This is attractive as $\text{Del}_k^n(P)$ has a smaller number of cells than $\text{DelN}_k^n(P)$, in particular in higher dimensions.

Going one step further, it turns out that there is no need to slice the rhomboid tiling at all. Let the depth of a rhomboid be the minimal depth of its vertices. By thus defining a depth function on all cells of $\text{Rho}(P)$ rather than only the vertices, together with the radius function we obtain a bifiltration on the rhomboid tiling directly. Let $\text{Rho}(P, r)_{\geq k}$ be the sublevel set with respect to r and the superlevel set with respect to k . Then the corresponding bipersistence module is isomorphic to the bipersistence module of $\text{Del}_k^n(P, r)$: The main insight for this result is that $\text{Rho}(P, r)_{\geq k}$ is a deformation retraction of the corresponding $\text{Del}_k^n(P, r)$ for any k and r . But this means that $\text{Rho}(P, r)_{\geq k}$ is homotopy equivalent to $\text{Cover}_k(P, r)$, and the bipersistence modules of $\text{Rho}(P)$ and $\text{Cover}_k(P, r)$ are isomorphic as well. Thus we obtain the means of computing 2-parameter persistence of the multicovers on the rhomboid tiling directly, which is even smaller than $\text{Del}_k^n(P)$ and simplifies the algorithm further.

5 Discrete Morse theory

In this chapter we look at the radius function on the rhomboid tiling and the order- k Delaunay mosaic from a viewpoint of discrete Morse theory [40]. We show that the squared radius function \mathcal{R}^2 on $\text{Rho}(P)$ is a generalized discrete Morse function [42] which furthermore has the special property that all its intervals have a vertex as a lower bound (Section 5.1). Based on this result, in Section 5.1.1 we improve on the algorithm from the previous chapter that computes these radius functions.

On the other hand, $\mathcal{R}_k^2: \text{Del}_k(P) \rightarrow \mathbb{R}$ neither satisfies the requirements of a discrete Morse function [40] nor the slightly weaker requirements of a generalized discrete Morse function [42]. Nevertheless we prove that it behaves similar to a discrete Morse function, so that its increments can be meaningfully classified into critical and non-critical steps with predictable impact on the homotopy type (Section 5.2). A step, intuitively, is a connected component of the level set of the radius function. Noting that for all cells σ in a given step, $S := S_{\min}(\sigma)$ and in particular $P_{on} := P_{on}(\sigma)$ are the same, our main result of Section 5.2 is the following classification of the topology types of the steps:

- We call the configuration of cells that defines a step *self-centered* if the center of S is contained in the simplex spanned by P_{on} . 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 \mathcal{R}_k^2 .
- We call the configuration *altruistic* if the center of S is not contained in the simplex spanned by P_{on} . Adding the cells in the step preserves the homotopy type, so we refer to it as a *non-critical step* of \mathcal{R}_k^2 .

Discrete Morse Theory. To state our results more formally, we need to give a few definitions from discrete Morse theory, also see e.g. [33] for reference. Let K be a regular 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 $\dim \tau = \dim \sigma + 1$. Let $f: K \rightarrow \mathbb{R}$ be a monotonic function, i.e. $f(\sigma) \leq f(\tau)$ whenever $\sigma \subseteq \tau$. The *level set* for a value $r \in \mathbb{R}$ is the set of cells $f^{-1}(r) \subseteq K$. 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 $\sigma = v$; if σ is a proper face of v then the interval is *non-singular*. A monotonic function $f: K \rightarrow \mathbb{R}$ is *generalized discrete Morse* if every step is an interval; see [42]. For comparison, f is *discrete Morse* if every step is an interval of size 1 or 2; see [40] 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.

Indeed, consider two contiguous sublevel sets that differ by a level set: $f^{-1}[-\infty, r] \setminus f^{-1}[-\infty, r) = f^{-1}(r)$. If this difference is a non-singular interval, then the two sublevel sets have the same homotopy type, while if the difference is a singular interval, then they have different homotopy types. In Section 5.1 we will prove that the squared radius function \mathcal{R}^2 on the rhomboid tiling is a generalized discrete Morse function with the additional property that every sublevel set is contractible.

5.1 Topology of rhomboid steps

Recall how we defined the squared radius function on the rhomboid tiling, $\mathcal{R}^2: \text{Rho}(P) \rightarrow \mathbb{R}$. By continuously dropping down the paraboloid \mathcal{P} , eventually it intersects every cell in the dual hyperplane $\text{Arr}(P)$. With \mathcal{P}_t being \mathcal{P} shifted downwards by $\frac{t}{2}$, the squared radius value of a rhomboid ρ is thus the smallest t such that \mathcal{P}_t intersects its dual cell in the arrangement, ρ^* .

Assuming general position, the sequence in which the paraboloid encounters the cells in $\text{Arr}(P)$ follows a few simple rules. For example, when the paraboloid encounters a vertex in $\text{Arr}(P)$, 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. Thus all of these faces, or more precisely their dual rhomboids, have the same radius value. The following lemma generalizes and formalizes this observation.

Proposition 5.1 (Generalized Discrete Morse). *Let $P \subseteq \mathbb{R}^d$ be locally finite and in general position. Then $\mathcal{R}^2: \text{Rho}(P) \rightarrow \mathbb{R}$ is a generalized discrete Morse function. Furthermore, all intervals in the implied partition have a vertex as a lower bound, and there is only one singular interval, which contains the vertex at the origin.*


Proof. The vertex at the origin corresponds to the three-partition $(\emptyset, \emptyset, P)$, has radius $\mathcal{R}^2(0) = -\infty$, and forms a singular interval. Every other interval is defined by a point $y \in \mathbb{R}^{d+1}$ at which the dropping paraboloid first touches a cell of the arrangement. There is one such point on every plane that is the common intersection of hyperplanes forming the arrangement. By general position, all these points are different. Let y belong to an i -plane, which is common to $j = d + 1 - i$ hyperplanes. It belongs to the interior of an i -cell, which is common to 2^j chambers. Exactly one of these chambers has not already been touched before the i -cell. The paraboloid touches this chamber at the same point y and similarly every cell that is a face of this chamber and contains the i -cell as a face. The corresponding rhomboids form an interval of the radius function, with an upper bound of dimension j and a lower bound of dimension 0. We have $1 \leq j \leq d + 1$, which implies that the interval is not singular.

To show that \mathcal{R}^2 is a generalized discrete Morse function, we still need to make sure that intervals in the same level set are *separated*, by which we mean that no simplex of one interval is face of a simplex in the other interval. By assumption of general position, there is only one level set that contains more than one interval, namely $(\mathcal{R}^2)^{-1}(0)$. All its intervals are of the form $[y_p, 0y_p]$, in which p is a point in P , the origin $0 \in \mathbb{R}^{d+1}$ corresponds to the three-partition $(\emptyset, \emptyset, P)$, and $0y_p$ is the edge that connects 0 with y_p . While these edges all share 0, no two also share the other endpoint. It follows that these intervals are components of the Hasse diagram of the level set, as required. \square

Note that \mathcal{R}^2 is rather special because it has a limited collection of interval types. This is best seen by constructing $\text{Rho}(P)$ 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 \mathcal{R}^2 is contractible.

Interval vertices. Proposition 5.1 says that the paraboloid always enters a chamber together with a subset of its faces while sweeping. 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 now, the vertex is not necessarily the lowest vertex of the rhomboid, also see Figure 5.1.



(a) $\text{Rho}(P)$, with colors indicating filtration values of Delaunay cells σ which simultaneously indicate filtration value of the rhomboids $\rho(\sigma)$. Here r_Q denotes the circum-radius of the set of points Q .

(b) $\text{Arr}(P)$ in red, with the shifted paraboloid \mathcal{P}_t in blue. Dashed lines are pairwise intersections of hyperplanes.

(c) Top view of (b). The blue disks around a, b, c are the parts of the hyperplanes h_a, h_b, h_c that are above the paraboloid. Dotted circles: paraboloid-hyperplane intersections.

Figure 5.1: Rhomboid tiling and the dual arrangement for a planar point set $P = \{a, b, c\}$ as laid out in (c). The paraboloid is intersecting the bottom-most chamber (below h_a, h_b, h_c), but it is still above the intersection point of the three hyperplanes. Thus the bottom vertex in (a) has smaller radius value than the 3-dimensional rhomboid and does not form an interval with the rhomboid.

Letting $\rho \in \text{Rho}(P)$, 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 \mathcal{R}^2 . Proposition 5.1 implies that $\lambda = \text{last}(\rho)$ for every interval $[\lambda, \rho]$ of $\mathcal{R}^2: \text{Rho}(P) \rightarrow \mathbb{R}$. 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 5.2 (Last not Top Vertex). *Let P 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 \mathcal{R}^2 .*

Indeed, the chamber in $\text{Arr}(P)$ 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 \mathcal{R}^2 and therefore does not belong to the interval. In the following we will look at this result from a different angle to obtain a means for computing the vertex that a rhomboid forms an interval with.

5.1.1 Computing intervals

The alternative characterization of the rhomboid radius function from Lemma 4.7 provides another viewpoint for Lemma 5.2. To identify the vertex v that a rhomboid ρ forms an interval with, we need to identify its vertex with the same radius value. By Lemma 4.7 this means the radii of $S_{\min}(\rho)$ and $S_{\min}(v)$ have to be the same, and it is not hard to see that the spheres $S_{\min}(\rho)$ and $S_{\min}(v)$ are in fact the same. As $P_{\text{on}}(v) = \emptyset$ for any vertex v , the sphere achieving the radius value of v is defined solely by inclusions and exclusion constraints. Therefore all constraints of ρ that require points of $P_{\text{on}}(\rho)$ to be on the sphere need to be converted to inclusion and exclusion constraints without affecting the resulting sphere. We know that such constraints exist because the lower bound of the interval is a vertex. This observation gives rise to the following lemma.

Lemma 5.3. *Let ρ be a rhomboid that is an upper bound of an interval. Let $P_I \subseteq P_{\text{on}}(\rho)$ such that the smallest enclosing sphere S of P_I that excludes $P_{\text{on}}(\rho) \setminus P_I$ is the same as the circumsphere of $P_{\text{on}}(\rho)$. Then ρ forms an interval with the vertex $v = P_{\text{in}}(\rho) \cup P_I$.*

Proof. As ρ is an upper bound of an interval, its sphere $S_{\min}(\rho)$ is only supported by $P_{\text{on}}(\rho)$. If there were another point $p \in P_{\text{in}}(\rho)$ (or $p \in P_{\text{out}}(\rho)$) on the surface of this sphere, then the rhomboid ϱ with $P_{\text{on}}(\varrho) = P_{\text{on}}(\rho) \cup \{p\}$ and $P_{\text{in}}(\varrho) = P_{\text{in}}(\rho) \setminus \{p\}$ (or $P_{\text{out}}(\varrho) = P_{\text{out}}(\rho) \setminus \{p\}$) would be a higher dimensional rhomboid with the same sphere $S_{\min}(\varrho) = S_{\min}(\rho)$, contradicting that ρ is an upper bound of an interval.

As $S_{\min}(\rho)$ is only supported by $P_{\text{on}}(\rho)$, this means that $S_{\min}(\rho)$ is the same as the circumsphere of $P_{\text{on}}(\rho)$, which by our assumption is the same as S . Now the inclusion and exclusion constraints of S are part of the constraint set for $S_{\min}(v)$, but because $S = S_{\min}(\rho)$ it does in fact fulfill all the constraints of $S_{\min}(v)$. Thus $S_{\min}(v) = S = S_{\min}(\rho)$, proving that they are in the same interval. \square

Note that for unweighted point sets, this yields a stronger version of Lemma 5.2: If $\#P_I \leq 1$, the resulting sphere has squared radius 0, or $-\infty$ in the case of the origin. However the only rhomboids with radius value 0 are the vertices adjacent and edges incident to the origin. Thus for all other rhomboids ρ that are upper bounds, $\text{last}(\rho)$ cannot be a vertex at depth ≤ 1 within ρ , i.e. $\text{last}(\rho)$ is neither $\text{top}(\rho)$ nor a vertex adjacent to $\text{top}(\rho)$. This furthermore implies that apart from the edges incident to the origin, no other edges are upper bounds of any interval. These conclusions do not hold for weighted point sets.

Algorithm. Assume ρ is an e -rhomboid that is an upper bound of an interval. Let S be the circumsphere of $P_{on}(\rho)$. For each point in $p \in P_{on}(\rho)$ we need to decide whether to impose an inclusion or exclusion constraint on it. Let S_p be the circumsphere of $P_{on}(\rho) \setminus \{p\}$. If p is outside of S_p , then imposing an exclusion constraint for p would yield S_p rather than S , thus we add p to P_I in order to impose an inclusion constraint for it. Similarly, if p is inside of S_p , we have to impose an exclusion constraint for p and thus do not add it to P_I , see Figure 5.2.




Figure 5.2: Checking which constraint to impose on p_3 for a rhomboid ρ with $P_{on}(\rho) = \{p_1, p_2, p_3\}$. The circumsphere of $\{p_1, p_2\}$ (dashed) contains p_3 in its interior, therefore we need to impose an exclusion constraint on p_3 to get back the circumsphere of $P_{on}(\rho)$.

It remains to identify whether a rhomboid is an upper bound of an interval. While a-priori this is hard to tell for an individual rhomboid, it becomes straightforward if we want to compute all intervals in the rhomboid tiling. We know that all $(e + 1)$ -rhomboids are upper bounds of intervals. After marking all rhomboids that are contained in such intervals, we know that all remaining unmarked e -rhomboids are upper bounds of intervals. Thus by processing the rhomboids in decreasing dimension, all unmarked rhomboids we encounter are upper bounds. Algorithm 3 describes the resulting algorithm for computing the radius function on $\text{Rho}(P)$.

Algorithm 3 computes the squared radius function \mathcal{R}^2 for a given rhomboid tiling $\text{Rho}(P)$. Rhomboids come in intervals with the same radius value. Instead of computing the radius value for each rhomboid individually as the radius of a constrained sphere, this algorithm computes the radius value for each interval as the radius of a circumsphere.

```

function COMPUTERADII( $\text{Rho}(P)$ )
  for  $e$  from  $d + 1$  down to 0 do
    for all  $e$ -rhomboids  $\rho$  in  $\text{Rho}(P)$  do
      if  $\mathcal{R}^2(\rho)$  is not defined then
         $CS := \text{CIRCUMSPHERE}(P_{on}(\rho))$ 
         $\mathcal{R}^2(\rho) :=$  squared radius of  $CS$ 
         $P_I = \emptyset$ 
        for all  $p$  in  $P_{on}(\rho)$  do
           $CS := \text{CIRCUMSPHERE}(P_{on}(\rho) \setminus p)$ 
          if  $p$  outside  $CS$  then
            Add  $p$  to  $P_I$ 
          end if
        end for
         $P_O := P_{on}(\rho) \setminus P_I$ 
        //  $\rho$  forms an interval with lower bound  $v := P_{in}(\rho) \cup P_I$ 
        // Now iterate over all rhomboids in the interval
        for all subsets  $P_S$  of  $P_{on}(\rho)$  do
          Let  $\varrho$  be the rhomboid with
             $P_{in}(\varrho) = P_{in}(\rho) \cup (P_S \cap P_I)$ ,
             $P_{on}(\varrho) = P_{on}(\rho) \setminus P_S$ , and
             $P_{out}(\varrho) = P_{out}(\rho) \cup (P_S \cap P_O)$ 
           $\mathcal{R}^2(\varrho) := \mathcal{R}^2(\rho)$ 
        end for
      end if
    end for
  end for
  return  $\mathcal{R}^2$ 
end function

```

5.2 Topology of Delaunay steps

We are interested in the partition of $\text{Del}_k(P)$ into the steps of \mathcal{R}_k^2 . To this end, let σ and τ be two cells in $\text{Del}_k(P)$ and note that $\sigma \subseteq \tau$ iff $\rho(\sigma) \subseteq \rho(\tau)$. It follows that each step of \mathcal{R}_k^2 is the horizontal slice of a step of \mathcal{R}^2 .

Assuming $P \subseteq \mathbb{R}^d$ is locally finite and in general position, the radius function of the order-1 Delaunay mosaic is known to be a generalized discrete Morse function [9]. This property does not generalize to higher order. Nevertheless, we will show that we can still classify the steps of \mathcal{R}_k^2 into critical and non-critical types such that each critical step changes the homotopy type of the sublevel set in a predictable manner, and every non-critical step maintains the homotopy type of the sublevel set. The main result of this section is that critical and non-critical steps of \mathcal{R}_k^2 can be distinguished by whether $\text{last}(\rho)$ is equal to or different from $\text{btm}(\rho)$, with ρ the smallest rhomboid in $\text{Rho}(P)$ that contains the corresponding step of \mathcal{R}^2 . We begin with an enumeration of the types.

Step types. By Proposition 5.1, every step of \mathcal{R}^2 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 $\dim \rho = e + 1 \geq 0$, the rhomboid has vertices at $e + 2$ depth values, and letting k be the depth value of $\text{btm}(\rho)$, these values are $k - g$ for $0 \leq g \leq e + 1$. By Lemma 5.2, λ can assume only $e + 1$ of these depth values. If $\lambda \neq \text{btm}(\rho)$, then H_{k-g} has a non-empty intersection with the interior of at least one rhomboid in the interval for $1 \leq g \leq e$, and if $\lambda = \text{btm}(\rho)$, then there is one more, namely for $0 \leq g \leq e$. In total, we count $e^2 + e + 1$ possible types of slices; see Figure 5.3 for an illustration of the types for $e = 2$. Some of these types are symmetric. We refer to the $e + 1$ slices in case $\lambda = \text{btm}(\rho)$ as *self-centered* and the e^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 $P_{on}(\rho)$ contain the center of $S_{min}(\rho)$. 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.

Topology type. Letting $A \subseteq \text{Rho}(P)$ be a step of \mathcal{R}^2 , we write $H_k \cap A \subseteq \text{Del}_k(P)$ for the corresponding step of \mathcal{R}_k^2 . 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 $\chi(H_k \cap A) = \sum_{\sigma \in H_k \cap A} (-1)^{\dim \sigma}$, which may or may not be zero. We write $|A|$ for the union of interiors of the rhomboids in




Figure 5.3: 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. *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 correspond non-critical steps consist in both cases of a triangle with one edge, and a triangle with two edges and the shared vertex.

A , and $H_k \cap |A|$ for its slice at depth k . Let $\mathbb{H}^e \subseteq \mathbb{R}^e$ be the set of points with non-negative first coordinate, and note that $\chi(\mathbb{H}^e) = 0$ for all $e \geq 1$. 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}^e , for some value of e , while every self-centered configuration has non-zero Euler characteristic.

Theorem 5.4 (Topology of a Step). *Let P be a finite set of points in general position in \mathbb{R}^d , let $A = [\lambda, \rho]$ be a step of \mathcal{R}^2 , set $e + 1 = \dim \rho$, write k for the depth of $\text{btm}(\rho)$, and recall that $H_{k-g} \cap A$ is a step of \mathcal{R}_{k-g} .²*

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

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 H_{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 $1 \leq g \leq e$, H_{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 Figure 5.3 for the three cases that occur for $e = 2$, and see Figure 5.4 for three of the four cases that occur for $e = 3$. Therefore, $\chi(H_{k-g} \cap A) = \sum_{q=g+1}^{e+1} (-1)^q \binom{e+1}{q}$, in which the binomial coefficient is the number of q -dimensional faces of a $(e+1)$ -dimensional rhomboid that share a common vertex, namely λ . This sum evaluates to $(-1)^{g+1} \binom{e}{g}$, which is non-zero, as claimed.




Figure 5.4: 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.

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 e -dimensional convex polytope. Call the preimage of its (relative) boundary the *silhouette* of ρ , and note that it is a $(e-1)$ -dimensional topological sphere that contains all vertices of ρ other than λ and λ' ; see Figure 5.3. 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, H_k and $H_{k-(e+1)}$ both have empty *intersection* with the interiors of all rhomboids in $[\lambda, \rho]$, as claimed. We thus assume $1 \leq g \leq e$ for the remainder of this proof. At depth $k-g$, the horizontal hyperplane intersects ρ in a convex polytope of dimension e , 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 H_{k-g} , we have to cross the intersection of H_{k-g} with the silhouette, which we will prove is a topological $(e-2)$ -sphere. We conclude that an open $(e-1)$ -ball of the boundary belongs to $H_{k-g} \cap |A|$, and the complementary closed $(e-1)$ -ball does not belong to $H_{k-g} \cap |A|$. It follows that the slice of the interval has the topology type of \mathbb{H}^e , as claimed. The middle and right panels of Figure 5.3 illustrate the two altruistic

configurations for $e = 2$, and Figure 5.5 illustrates the nine altruistic configurations for $e = 3$. 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 $e = 3$.

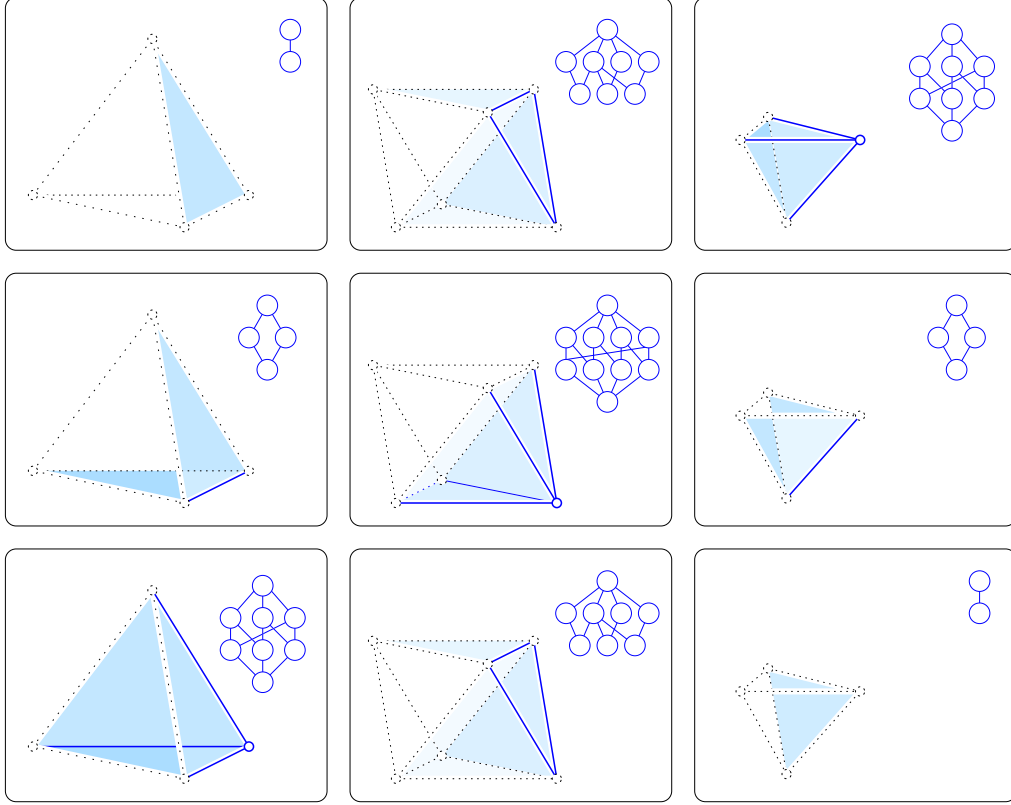


Figure 5.5: The 9 altruistic configurations in \mathbb{R}^3 . Compared to $\text{btm}(\rho)$ at depth k , the vertex $\lambda = \text{last}(\rho)$ has depth $k - 3$, $k - 2$, and $k - 1$ in the *top*, *middle*, and *bottom* row. Similarly, the slice is at depth $k - 3$, $k - 2$, and $k - 1$ in the *left*, *middle*, and *right* column.

We return to the intersection of H_{k-g} with the silhouette and the claim that this intersection is a topological sphere of dimension $e - 2$. For $e = 1$, ρ is a convex quadrangle, its silhouette consists of two vertices, $\text{top}(\rho)$ and $\text{btm}(\rho)$, and H_{k-1} passes through the other two vertices thus intersecting the silhouette in the empty set — the (-1) -sphere — as claimed. Assuming $e \geq 2$, we denote the silhouette by \mathbb{S} , we recall that it is a $(e - 1)$ -sphere, and we write $w: \mathbb{S} \rightarrow \mathbb{R}$ for the depth function on the silhouette. Its extreme values are $w(\text{top}(\rho)) = k - (e + 1)$ and $w(\text{btm}(\rho)) = k$, and $H_{k-g} \cap \mathbb{S} = w^{-1}(k - g)$. To prove that this level set is a $(e - 2)$ -sphere, it suffices to show that w has only two critical points, namely the minimum at $\text{top}(\rho)$ and the maximum at $\text{btm}(\rho)$. The case $e = 2$ 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 H_{k-g} meets the silhouette in two points — a 0-sphere — as claimed.

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

Consequences. We now turn Theorem 5.4 into a statement about the filtration of order- k Delaunay mosaics. Let $r_0^2 < r_1^2 < \dots$ be the sorted values of \mathcal{R}_k^2 and write $K_\ell = (\mathcal{R}_k^2)^{-1}[-\infty, r_\ell^2] \subseteq \text{Del}_k(P)$ for every $\ell \geq 0$. Assuming P 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, so to achieve this, we can move to the weighted setting and perturb the point set by assigning small weights. While it is not necessary, we simplify the following statement by using this stronger notion of general position.

Corollary 5.5 (Filtration of Order- k Delaunay Mosaics). *Let P 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 \mathcal{R}_k^2 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 \mathcal{R}_k^2 , then K_u and K_v have the same homotopy type.*

This corollary of Theorem 5.4 is a direct extension of a theorem about discrete Morse functions in [40]. Other results in this theory can be similarly extended.

5.3 Discussion

The main result of this chapter is a topological characterization of the incremental steps of the radius function on both the rhomboid tiling and the order- k Delaunay mosaic of a finite set of points in Euclidean space. For the rhomboid tiling, we show that except for the step consisting of the origin, all steps are intervals and thus the sublevel sets of the rhomboid tiling are contractable; we also give a recipe for computing these intervals. For the order- k Delaunay mosaic we provide a characterization of critical and non-critical steps. With this insight, we gain a topological interpretation of the probabilistic analysis of the order- k Delaunay radius function for a stationary Poisson point process [36]. 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 [9], the radius function in the order- k case is neither discrete Morse nor generalized discrete Morse [40; 42]. Since the function nevertheless behaves similar to a Morse function, it may be considered a geometrically motivated further extension of the framework; see also [54, 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 finite point sets that are not necessarily in general position.

6 Analysis of sphere packings

For a given d -dimensional metric space \mathbb{X} , a sphere packing is a collection of $(d - 1)$ -dimensional spheres (or equivalently d -dimensional balls) that are internally disjoint. Its packing density is the fraction of the containing space that is covered by those spheres. We will only consider packings of spheres with equal radius, also called *mono-disperse* packings. In 2-dimensional Euclidean space it is a long known result, see e.g. [84], that the optimal packing density is $\pi/\sqrt{12} \approx 91\%$, and it is uniquely achieved by the hexagonal packing, see Figure 6.1. In practice the 3-dimensional case is of more interest due to its




Figure 6.1: The hexagonal sphere packing in the plane.

applications in crystallography and material sciences. In his seminal work, Hales proved that $\pi/\sqrt{18} \approx 74\%$ is the maximum packing density that can be achieved in 3 dimensions [45]. This density is achieved by stacking layers of 2-dimensional hexagonal packings on top of each other. Such packings are called *Barlow stackings*. In the simplest of these structures, the layers are arranged as $ABAB\dots$, which corresponds to the *hexagonally close-packed* (HCP) structure, and $ABCABC\dots$ for the *face-centered cubic* (FCC) structure [49; 72], also see the first and last rows of Figure 6.6. More complex examples can be found in

structures of metal compounds where the anions are arranged in a Barlow stacking, with the cations occupying the holes of a HCP packing. [90; 64]. Such highly regular packings are also called *crystalline packings*.

In the physical world, when identical hard spheres are randomly packed together, they naturally form a disordered structure that fills 64% of the space. Although the packing structure is disordered, the limiting packing density of 64% is highly reproducible. This phenomenon was extensively studied by J. D. Bernal in the 1950s and 1960s and the limiting density of 64% is known as the Bernal density or φ_{Bernal} . Numerous experimental and numerical studies have extended Bernal's seminal findings and reported that a stable configuration of frictional mono-disperse spheres can exist at densities ranging from $\varphi = 0.55$ to $\varphi = 0.64$ [11; 5; 52]. It is possible to break through φ_{Bernal} by intensely vibrating [76] or cyclically shearing [73] of the packing structure. This forces the density of the packing to increase, up to the maximum of $\varphi_{max} \approx 74\%$. Beyond Bernal's density, crystalline clusters inevitably appear in mono-disperse sphere packings, see Figure 6.2. A natural question to ask is: How are the spheres packed together locally in the disordered regime and what happens to these local structures during the transition through φ_{Bernal} all the way to φ_{max} ?




Figure 6.2: 3D rendering of the experimental packing. A cut-out section shows the bulk crystallisation and partially crystallised regions

Due to the abundance and importance of such local motifs in physical systems, many techniques have been invented for their quantification, each with strengths and shortcomings. Bond orientational order parameters [64; 53] are widely used in condensed matter physics to characterise local crystalline structures. They quantify local structure within a short range of a given sphere, and depend on the choice of neighbourhood of that sphere, of which various notions exist [46]. Crucially, regardless of the neighbourhood notion, small perturbations in the sphere packing can cause the set of neighbours of a sphere to change, causing a discontinuity in the bond orientational order parameters [63].

Addressing this lack of continuity, sphere packings have been analyzed with persistent homology [76]. In this setting persistence is computed for the union of balls of radius r around the sphere centers of a packing. As the radius r increases from 0 to infinity, persistent homology tracks the emergence (births) and disappearance (deaths) of topological features of this union of balls. In the 3-dimensional setting, these topological features are *connected components* (or the *gaps* between them), *closed loops* (or the *tunnels* they form), and *closed surfaces* (or the hollow *voids* they surround). The resulting persistence diagram has the desirable properties that it is invariant under isometries (rotations, translations, reflections) of the input, and that small perturbations in the input cannot lead to large changes in the persistence diagram, which is called *stability* [22].

One shortcoming of persistent homology of the union of balls is that FCC and HCP packings yield the same persistence diagrams, and thus cannot be distinguished. Numerous studies have shown that in experimental sphere packings, FCC is the preferred structure [85; 41; 72] despite the fact that both FCC and HCP structures have an identical packing density of $\approx 74\%$. To shed light on this phenomenon without forfeiting the stability property of persistence diagrams, we use persistence of the k -fold cover in scale, or *k -cover persistence* for short, as introduced in Section 4.3.1, to analyse the sphere packings.

For a finite point set $P \subseteq \mathbb{R}^d$, recall that the k -fold cover, $\text{Cover}_k(P, r)$, is the subset of \mathbb{R}^d that is covered by at least k (closed) balls of radius r around the points P . The 1-fold cover as a special case is simply the union of the balls. For an input point set P , using persistence we obtain a persistence diagram $\text{Dgm}(f_k)$ as output. The notation is motivated by the fact that the k -fold covers are sublevel sets of the k -th distance function f_k for the points P . Just like the union of balls, the k -fold cover grows as we increase the radius of the balls, and its topology changes; see Figure 6.3 for a 2-dimensional example


with $k = 3$. The computed persistence diagram, $\text{Dgm}(f_3)$, of the point set P is shown in Figure 6.4, showing the birth and death times of various topological features in the 3-fold cover.

In the rest of this chapter, we analyze the geometric structure of FCC and HCP packings (Section 6.1) to show that with k -cover persistence already for $k = 4$, FCC and HCP have distinct persistence diagrams, as seen in Figure 6.5. We then use these insights to evaluate the 4-cover persistence diagrams of two experimental sphere packing data sets (Section 6.2), both courtesy of Mohammad Saadatfar. The first is a digital reconstruction of a real-world packing of 170000 spherical acrylic beads in a cylindrical container, see Figure 6.2, with local regions of different packing densities. The second is a time series obtained through a molecular dynamics simulation to assess the structural changes of FCC and HCP packings under external forces, see Figure 6.9.


6.1 Persistent features in FCC and HCP

In order to understand the persistence diagrams of our experimental datasets, we first have to analyze how geometric structures in FCC and HCP packings yield persistence pairs in $\text{Dgm}(f_4)$. The first and last rows of Figure 6.6 show schematic sketches of the FCC and HCP packings respectively. In particular we need to understand tetrahedral and octahedral cavities and their adjacency relations. Such cavities are formed by touching spheres whose center points span regular tetrahedra or octahedra. In our analysis, we will focus on $\text{Dgm}_0(f_4)$, the persistence diagram restricted to 0-dimensional topological features (i.e. connected components), as they are easier to interpret geometrically than the higher-dimensional features.


Assuming spheres of diameter 1, a component of the 4-fold cover emerges at the center of a tetrahedral cavity at radius $\sqrt{6}/4$. Within an octahedron, the 4-fold cover emerges at radius $\sqrt{2}/2$ at its center. (Both of these radii correspond to the circumradii of the tetrahedra and octahedra.) The radius for when these components merge with other components depends on a slightly larger neighbourhood. If two tetrahedra are face-adjacent, a configuration only present in HCP, the components merge earlier than if they are only edge-adjacent (as present in both FCC and HCP, but with different multiplicities).




(a) Nowhere do three or more balls overlap yet, thus the 3-fold cover is empty.




(b) Three connected components (F, G and H) of the 3-fold cover have emerged at the centers of the equilateral triangles.




(c) Another component I has emerged in the center of the square.



(d) The component G has merged into component F, while two new components have appeared.



(e) Component I has merged into component H, in the process creating a small hollow space, the topological loop K, between them.



(f) All components have merged back into component F, making the topology of the 3-fold cover trivial.

Figure 6.3: Development of the 3-fold cover for increasing radii, with the birth and death of topological features (components, loops). For each point, the ball of radius r is drawn in semi-transparent pink. The 3-fold cover is where at least three of these balls overlap, and is drawn in solid red.




Figure 6.4: 3-cover persistence diagram of the example point set from Figure 6.3, representing the topological changes over time. Each point is annotated with the label of the topological feature present in Figure 6.3. Features J and J' have the same times of birth and death, and are represented by point J which thus has multiplicity 2. Feature F is not present as its time of death is at infinity.

Table 6.1 summarizes the features we see in $\text{Dgm}_0(f_4)$ of the FCC and HCP packings. While zero-persistence features are usually not considered as features, we include the one stemming from octahedra for completeness as they will become non-zero persistence features once these octahedra get deformed, e.g. in Figures 6.6, 6.7 and 6.10.

Table 6.1: Summary of features in $\text{Dgm}_0(f_4)$ for FCC and HCP packings. Note that face-adjacent tetrahedra are only present in HCP packings, but not in FCC packings

Pair in $\text{Dgm}_0(f_4)$	Decimal	Feature
$(\sqrt{6}/4, \sqrt{6}/3)$	(0.61, 0.82)	A two face-adjacent tetrahedra
$(\sqrt{6}/4, \sqrt{3}/2)$	(0.61, 0.87)	B two edge-adjacent tetrahedra
$(\sqrt{2}/2, \sqrt{3}/2)$	(0.71, 0.87)	C octahedron-tetrahedron adjacency
$(\sqrt{2}/2, \sqrt{2}/2)$	(0.71, 0.71)	D octahedron (zero-persistence)

We will now give more detailed geometric explanations for the features in Table 6.1. Figure 6.3 can provide some intuition if equilateral triangles are used as an intuitive proxy for tetrahedra, and the square as proxy for an octahedron.

With each tetrahedral cavity, a component of the 4-fold cover emerges at radius $\sqrt{6}/4$, which is the circumradius of such a tetrahedron. The radius for which this component merges with other components depends on the neighbourhood of this tetrahedron. If it is face-adjacent, as present in the HCP packing, then it merges with that tetrahedron's component at radius $\sqrt{6}/3$, which is the height of a tetrahedron or equivalently the radius when a ball centered around one of the tetrahedral vertices touches the opposite face. This explains the pair $(\sqrt{6}/4, \sqrt{6}/3)$ in the persistence diagram of the HCP packing. If instead

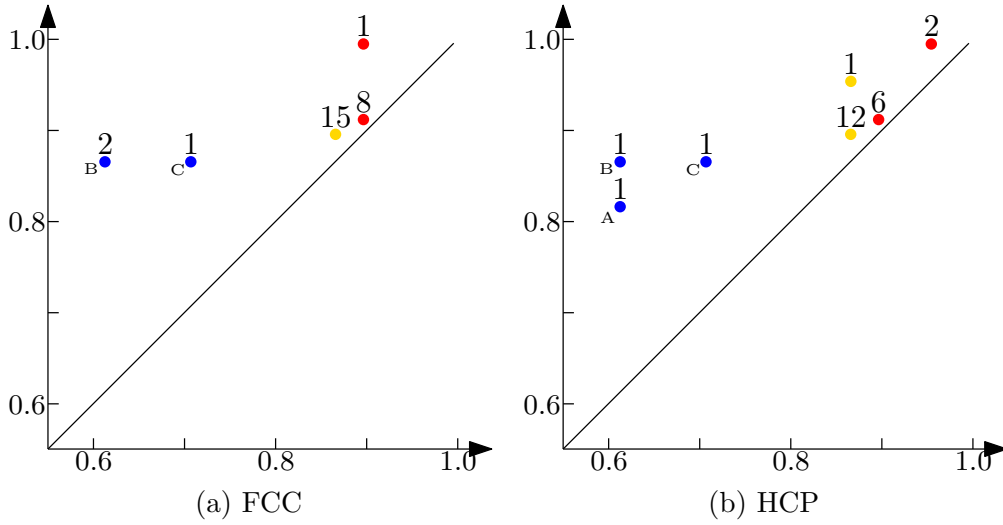


Figure 6.5: 4-cover persistence diagrams of FCC and HCP packings. 0-dimensional topological features (components) are shown in blue, 1-dimensional features (loops) in yellow, 2-dimensional ones (voids) in red. Each persistence pair is annotated above with a relative multiplicity, indicating how many topological features with the given birth and death coordinate exist per sphere. Furthermore each 0-dimensional persistence pair is annotated (in brackets) with the feature from Table 6.1 it corresponds to.

another tetrahedron is adjacent via an edge at the angle like in FCC and HCP, their components only merge at radius $\sqrt{3}/2$, which is the distance from a tetrahedral vertex to any of its opposite edges, equaling the height of a triangular face. This explains the pair $(\sqrt{6}/4, \sqrt{3}/2)$ in the persistence diagram. In FCC packings this configuration is twice as frequent as in HCP packings, and thus appears with multiplicity 2 there. Finally, both configurations contain octahedra that share faces with tetrahedra. The circumradius of an octahedron is $\sqrt{2}/2$. Thus within octahedra, the 4-fold cover only emerges at radius $\sqrt{2}/2$ at the circumcenter of the octahedron. It merges with the component of a face-adjacent tetrahedron at radius $\sqrt{3}/2$, which is half the distance between the vertex unique to the tetrahedron (i.e. not shared with the octahedron) to any of the vertices unique to the octahedron (i.e. not shared with the tetrahedron). This distance equals the height of a regular triangle, $\sqrt{3}/2$, as each of the non-shared tetrahedral faces is co-planar with the incident octahedral face. This explains the persistence pair at $(\sqrt{2}/2, \sqrt{3}/2)$. As there are 6 rather than merely 4 points at the same distance from the circumcenter of an octahedron, when the octahedron is slightly deformed it is possible that multiple components arise at a radius close to $\sqrt{2}/2$ which merge together very quickly. So while in a perfect packing these components have 0-persistence, in experimental packings we will often also see persistence pairs close to the diagonal at $(\sqrt{2}/2, \sqrt{2}/2)$, stemming from this phenomenon. These are

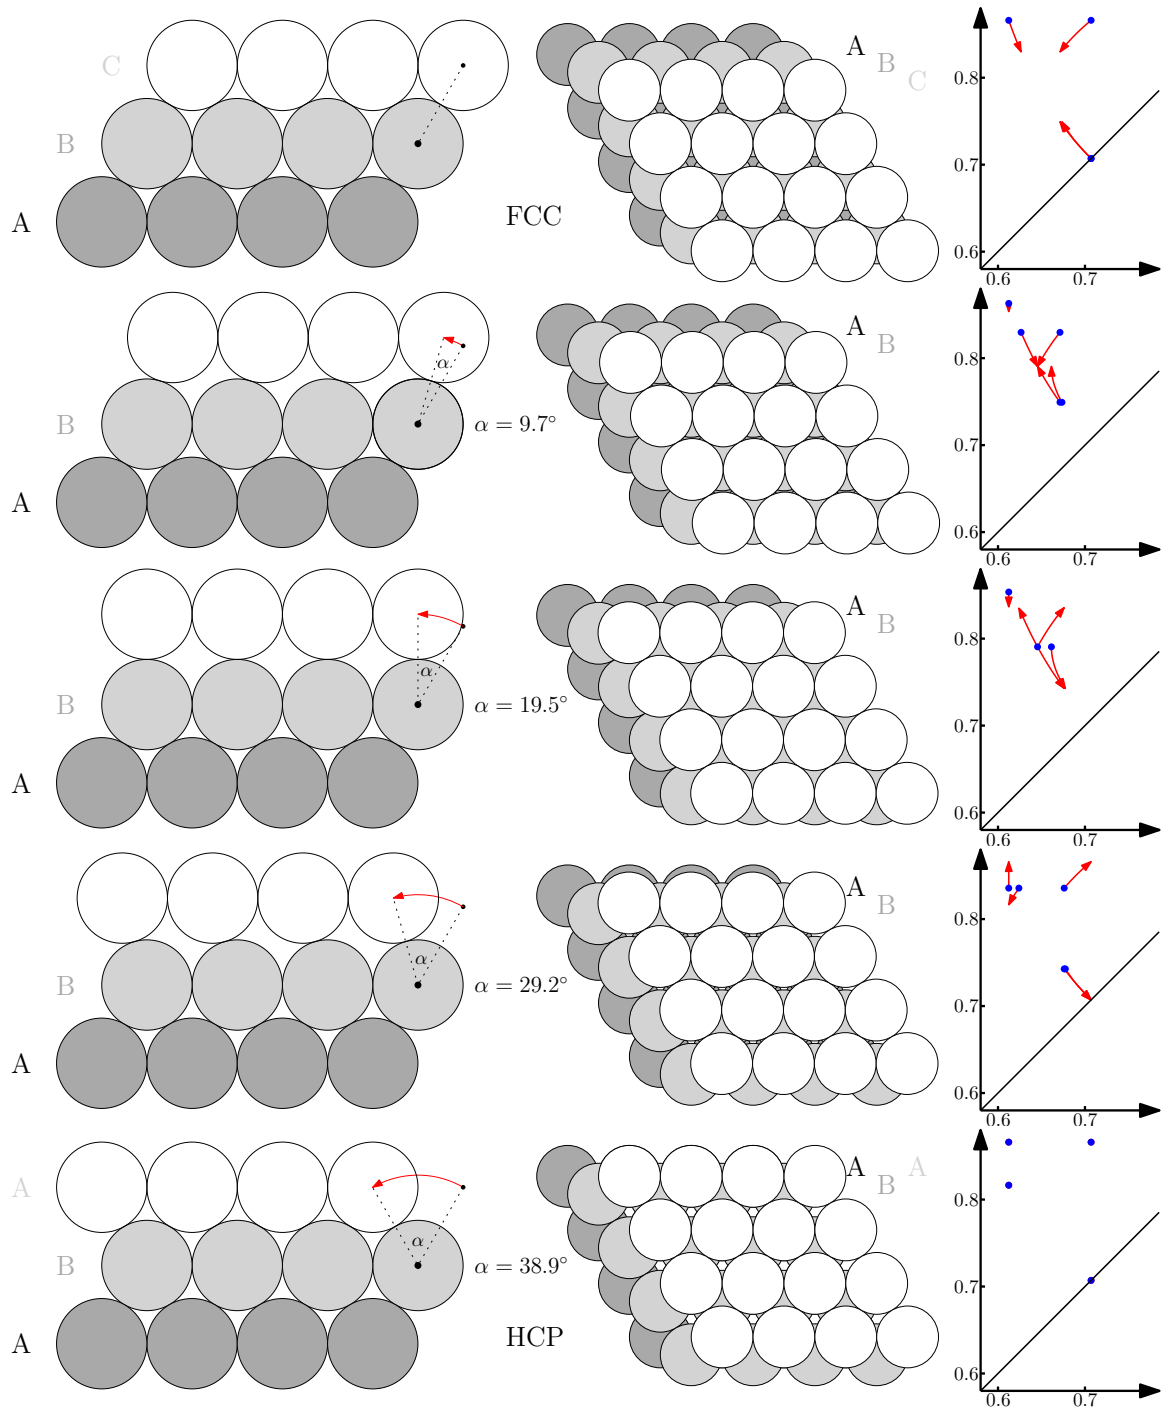


Figure 6.6: Deformation of an FCC packing (first row) into an HCP packing (last row) with 3 intermediate configurations. For each row, the first column shows a schematic sketch of the configuration with the corresponding shifting angle; the second column shows a projection from above; the third column shows the corresponding $\text{Dgm}_0(f_4)$ (blue) with red curves indicating the continuous change in $\text{Dgm}_0(f_4)$ when the configuration is continuously deformed into the next row's configuration.

present whenever octahedra are present, and thus appear in both slightly deformed FCC and HCP packings.

Deformations. We investigate the changes in $\text{Dgm}_0(f_4)$ as an FCC packing is deformed into an HCP packing by moving a layer of spheres over a saddle point. Figure 6.6 provides an illustration of the process and the resulting changes in $\text{Dgm}_0(f_4)$. Note that when a packing is continuously deformed, the persistence diagram changes continuously as well due to its stability property, and thus we see continuous curves in $\text{Dgm}_0(f_4)$.

6.2 Analysis of experimental sphere packings

Using 4-cover persistence and the insights about the persistence diagrams of FCC and HCP packings, we analyse the two aforementioned data sets. We first compare persistence diagrams for different packing densities, and then define a measure of frequency of FCC and HCP using 4-cover persistence to quantify these differences with specific focus on FCC- and HCP-like structures. Using simulation data, we then investigate how FCC and HCP packings are affected when shearing forces are applied.

Note that due to the high total number of spheres and thus persistence pairs, the diagrams in this section are visualized as heatmaps, each pixel representing the number of persistence pairs in its value range.

6.2.1 Different densities.

We investigate the relationship between packing densities and persistence diagrams for experimental sphere packing data sets. The main data set was obtained by putting approximately 170000 acrylic beads into a cylindrical container, subjecting the container to vibrations, and reconstructing the locations of the sphere centers via X-ray tomography and 3D image analysis. Figure 6.2 visualizes the complete data set. In order to get data sets for different packing densities, we analyze smaller regions of 4100–4600 spheres each.

Figure 6.7 shows $\text{Dgm}_0(f_4)$ for a few of these regions with different packing densities. In these experimental packings, we see persistence pairs in the locations from Table 6.1 characteristic of FCC and HCP structures, indicating the presence of these geometric

configurations in our packings. We furthermore see persistence pairs along the deformation curves from Figure 6.6, suggesting that these deformed configurations are present at various stages.

As the packing density increases, we see persistence pairs typical for FCC and HCP appearing in higher frequencies, indicating that crystalline domains appear in the packing. Interestingly, the persistence pair that is unique to HCP seems to be much less frequent than the two pairs shared with FCC. Exploiting these characteristic signatures in the persistence diagram, we will estimate the ratio of FCC and HCP present in a given packing in the following section.

6.2.2 Frequency measures.

Recall from Figure 6.5 that the persistence pairs $(0.61, 0.82)$, $(0.61, 0.87)$ and $(0.71, 0.87)$ appear in ratio 1:1:1 per sphere in HCP packings, while they appear in ratio 0:2:1 in FCC packings. We can “count” the number of persistence pairs close to these characteristic persistence pairs, to be explained in more detail later. Introducing two variables f_{HCP} and f_{FCC} to denote a frequency measure of HCP and FCC respectively, we find the least-squares solution to an over-defined linear system of 3 equations to estimate these frequencies of HCP and FCC and thus the ratio between HCP and FCC. The equation system is as follows:

$$\begin{aligned} f_{HCP} &= m_{face-tetra} \\ f_{HCP} + 2f_{FCC} &= m_{edge-tetra} \\ f_{HCP} + f_{FCC} &= m_{octa-tetra} \end{aligned}$$

where $m_{face-tetra}$, $m_{edge-tetra}$ and $m_{octa-tetra}$ denotes the number of persistence pairs “close” to $(0.61, 0.82)$, $(0.61, 0.87)$ and $(0.71, 0.87)$ respectively. Instead of directly counting how many persistence pairs are within a certain radius of each of these features, we rather weigh the contributions such that closer pairs contribute more to the count. Specifically, for each pair we compute the distance between the pair and the feature, and weigh it according to a Gaussian probability density function $g_\sigma(x)$ that is normalized such that $g_\sigma(0) = 1$, so that pairs that coincide exactly with the feature contribute 1. As standard deviation for the Gaussian we choose $\sigma = 0.005$, so that pairs within distance 0.005 still contribute approximately 0.6 to the count, while pairs at distance greater than 0.015 contribute less than 0.01. In particular, with this choice we ensure not to capture pairs belonging to

other features, as the two closest features have a distance of approximately 0.05. For the distance, we use L_∞ -distance which is the standard distance used in persistence diagrams. This choice of counting pairs ensures that the stability property of persistence diagrams translates to stability of our measure, meaning that small changes in the sphere packings imply small changes in estimated HCP and FCC frequency.

Figure 6.8 shows the values of our measure for HCP and FCC frequency as well as their ratio for different packing densities. Each data point is for an experimental packing of approximately 4000 spheres. We see similar results as observed using previous measures [47], in particular that FCC occurs more frequently in more crystalline packings than HCP.

Notice that our method will also distinguish more complex Barlow stackings, as long as the ratios between FCC-type layer groups (i.e. 3 consecutive distinct layers, ABC) and HCP-type layer groups (i.e. 3 consecutive layers with the first and third coinciding, ABA) in the stackings are different. For example, Sm or Mo_2S_3 with ABABCBCAC stacking [71; 90], Ti_4S_5 with ABABCBABAC stacking and Fe_3S_4 with ABCBCABABCAC stacking [90; 64] have ratios between FCC-type and HCP-type layer groups of 2:1, 3:2 and 1:1 respectively, and thus they exhibit different $\text{Dgm}_0(f_4)$ and our measure would indicate these different FCC/HCP ratios. For Barlow stackings with the same ratios, we expect $\text{Dgm}_0(f_k)$ for sufficiently large k to differ.

6.2.3 Stability of FCC and HCP

As seen in Figure 6.8, HCP structures appear less commonly in experimental packings than FCC structures. We investigate this phenomenon using a molecular dynamics simulation which allows us to dynamically track the “melting” of a crystalline structure from FCC/HCP to disorder. Figure 6.9 shows the beginning and end state of the simulation. A shearing force is applied to the packing of 6000 spheres by shifting the top layer of spheres horizontally. In each time step, forces between the spheres and the resulting displacements are computed using the Hertz–Mindlin contact model [47].

Using the 4-cover persistence diagrams, we analyse the stability of HCP and FCC patterns from a topological perspective as these motifs progress from order towards disorder. Figure 6.10 shows the temporal evolution of $\text{Dgm}_0(f_4)$ of the FCC and HCP packings at corresponding time steps. While both structures exhibit deformation patterns similar to

the ones shown in Figure 6.6, they are much more pronounced in the HCP melting. While the FCC packing stays stable in its (slightly perturbed) crystalline pattern, the HCP packing deteriorates into a disordered packing looking similar to Figure 6.7a. The birth and death coordinates of the FCC/HCP typical features are lower than they theoretically should be as shown in Figure 6.5, and they slightly decrease over time in the simulation. However this phenomenon is a side effect of the simulation allowing slight overlap of spheres under strong forces.

6.3 Conclusions

We applied persistence of k -fold covers in scale to investigate the occurrence of FCC- and HCP-like structures in sphere packings. This extends previous work [76] that used the special case for $k = 1$ which is unable to distinguish these two structures, but also maintains the stability under perturbations of the input, unlike other conventional parameters.

We defined a persistence based measure to estimate the relative occurrences of local FCC and HCP structures, and confirmed previous observations that FCC appears at higher frequencies in experimental packings than HCP, especially for higher packing densities. By applying 4-cover persistence to time series data sets of FCC and HCP packings under external forces, we observed that FCC has a more stable structure than HCP.

We would like to remark that the potential value of our framework is not limited to distinguishing FCC and HCP patterns in sphere packings. Persistence diagrams can be used in a much broader context, e.g. to compare different structures, to look for local structures with known signatures in their persistence diagram, or as a preprocessing step in conjunction with machine learning [56; 59]. Previous applications of persistence to discrete point sets include predicting physical properties of zeolites [59], or describing amorphous materials such as glass and distinguishing them from their liquid phase [68]. Like for the special case of $k = 1$ (union of balls), k -cover persistence is applicable wherever data is represented as a discrete point set, but it captures a wider range of features.

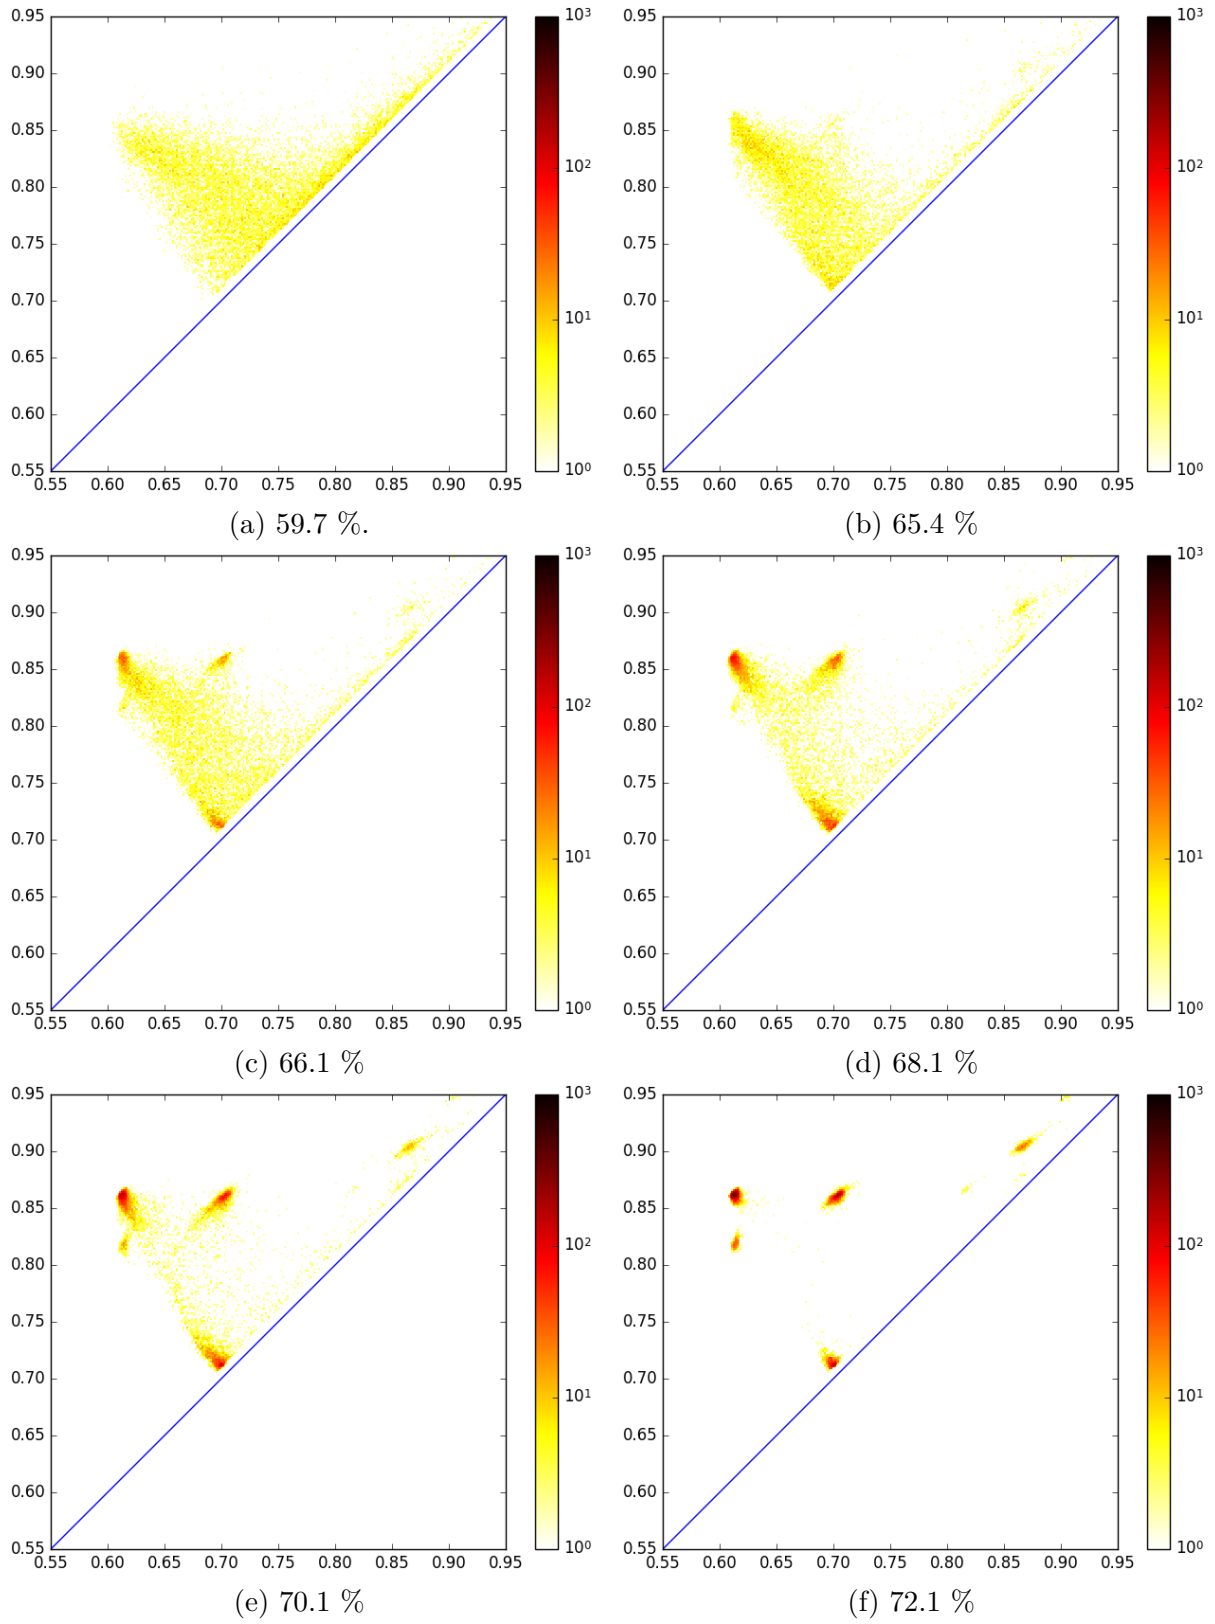


Figure 6.7: The 4-cover persistence diagram of dimension 0, $\text{Dgm}_0(f_4)$, for experimental sphere packings of various packing densities, each consisting of between 4100 and 4600 spheres.




Figure 6.8: Frequency measure of HCP- and FCC-like structures as well as ratio of the two in experimental packings of different densities, according to our measure.




Figure 6.9: Molecular dynamics simulations of the melting process from a regular HCP packing (left) to a disordered system (right) by inducing shear via moving the top layer.

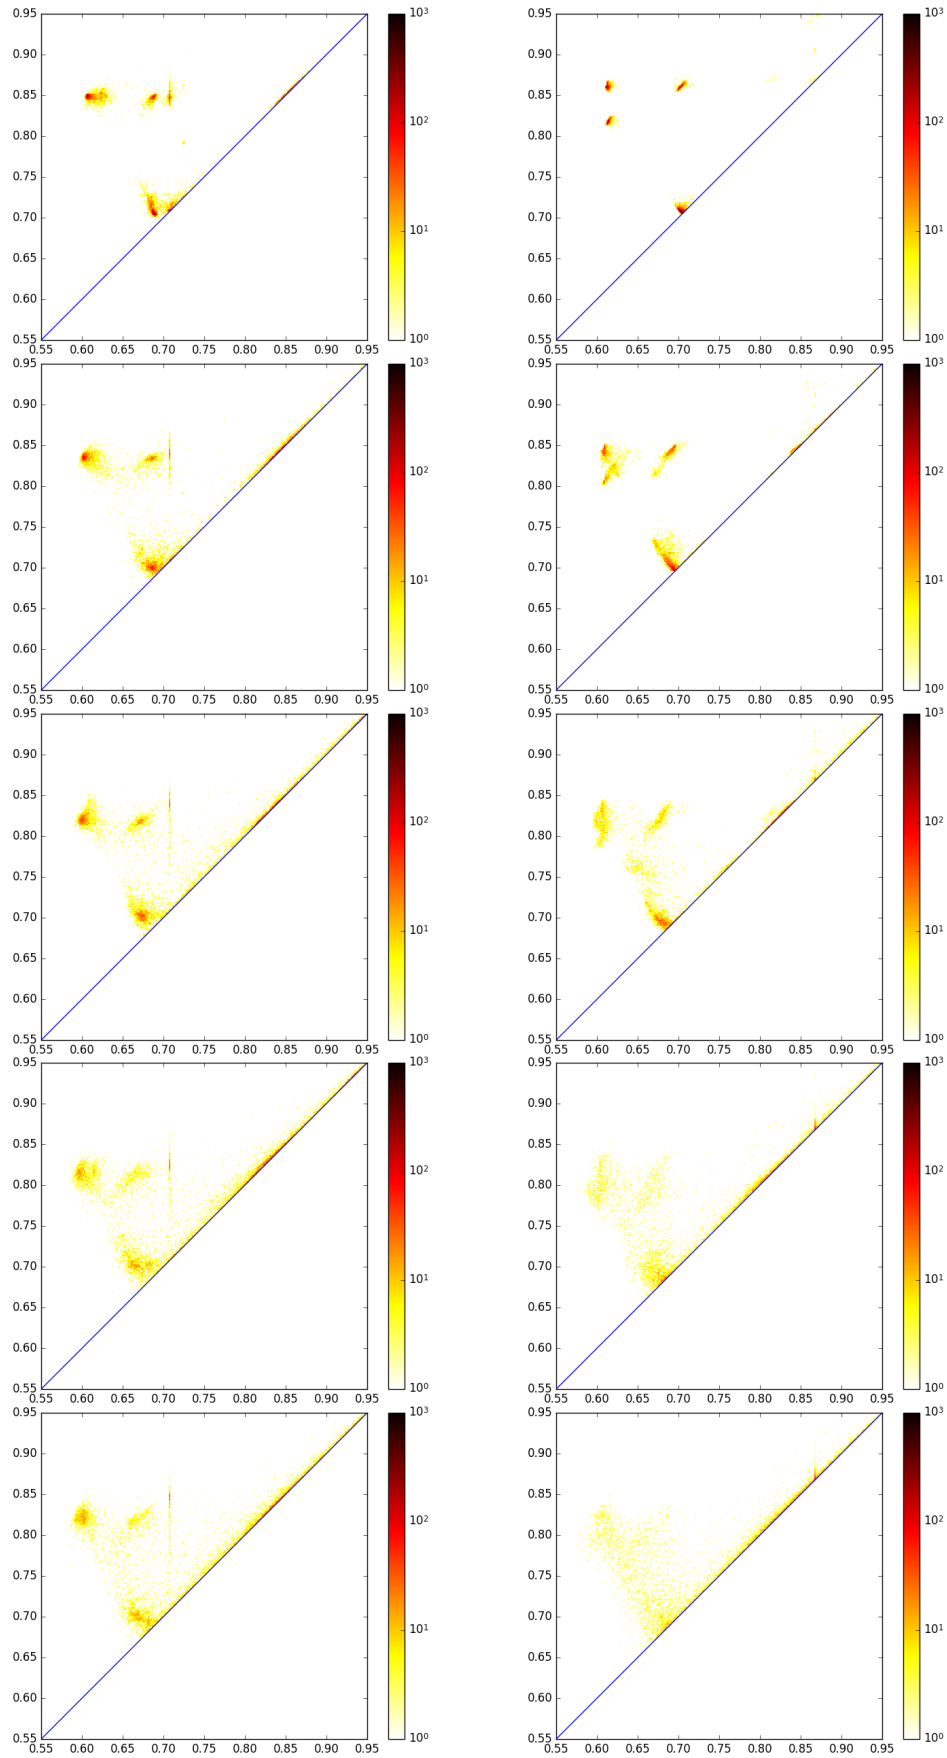


Figure 6.10: $\text{Dgm}_0(f_4)$ for simulated shearing of FCC (left) and HCP (right) packings at time steps 3, 11, 21, 33 and 49 of the simulation.

7 Periodic Delaunay triangulations

While nowadays many applications use Delaunay triangulations of finite point sets in Euclidean space, Delaunay originally introduced the notion in the context of infinite point sets with translational periodicity [26]. Such periodic point sets are abundant in fields such as crystallography and material sciences, and their communities would benefit from software that computes Delaunay triangulations of infinite periodic point sets in Euclidean space. Specifically, given a d -dimensional lattice and its associated translation group, the orbits of a given finite point set in \mathbb{R}^d with respect to this translation group define a periodic point set. Our aim is to compute a finite representation of the periodic Delaunay triangulation of such a periodic point set, specifically a projection of the triangulation onto the flat d -torus that is the quotient space of \mathbb{R}^d under the action of the translation group.

The first algorithm for this problem was already outlined in 1997 [29], yet to our knowledge no robust and efficient implementations for this problem exist to date. The Voro++ library [75] is focused on crystallographic applications in 3 dimensions, however is limited to orthogonal lattices. Zeo++ [92] extends its functionality to arbitrary 3-dimensional lattices. Both libraries compute the Voronoi cell of a given input point as an intersection of half-spaces, by searching for other points around it that have influence on its Voronoi cell. No easy access to the combinatorics of the Delaunay triangulation is provided, and the computations are done using floating point arithmetics. The CGAL library [81] provides packages for periodic Euclidean Delaunay triangulations in 2D and 3D, which currently are limited to the integer lattice, referred to as the square and cubic setting in 2 and 3 dimensions, respectively [57; 18; 17]. We propose an addition to CGAL to extend this functionality to arbitrary lattices.

The algorithm by Dolbilin and Huson [29] creates 3^d copies of each input point

and computes their finite Delaunay triangulation, from which a representation of the periodic Delaunay triangulation is extracted. The CGAL algorithm [16; 19] computes the triangulation in a finitely-sheeted covering space of the d -torus. It is based on the classical incremental algorithm by Bowyer and Watson [12; 89], and requires that the triangulation be a simplicial complex at any given time. Let us quickly recall that a triangulation is a simplicial complex, or is *simplicial* for short, if each of its simplices consists of a set of distinct vertices, and the intersection of any two simplices is either empty or a simplex. Operating directly on the d -torus does not guarantee this, see Figure 7.1a. Thus, in the cubic setting, a 3^d -sheeted cover is used (Figure 7.1b) until sufficiently many points have been inserted to guarantee that the triangulation in the 1-sheeted cover will be a simplicial complex. Unfortunately, a 3^d -sheeted cover is not sufficient for more general periodic point sets: As the 3^d copies of each point have to be inserted iteratively into the 3^d -sheeted cover, simpliciality can be violated in the intermediate stages of point insertion, see Figure 7.1c. While there always exists a finitely-sheeted covering space [19] that ensures simpliciality, the number of sheets might be prohibitively large. Thus we propose a different approach.

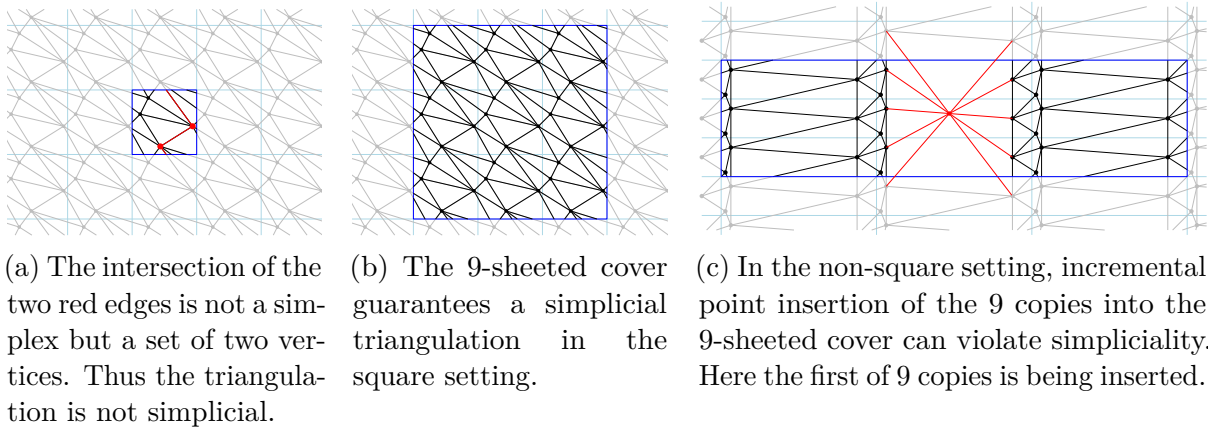


Figure 7.1: Representation of the projections of the periodic Delaunay triangulation into the 1-sheeted (left) and 9-sheeted cover (middle, right) of the 2-torus.

Overview After formally defining the problem in Section 7.1, we propose an algorithm (Section 7.2) for periodic Delaunay triangulations that combines two different approaches and consists of two phases, both of which use Bowyer-Watson’s algorithm. While for 2-dimensional periodic spaces algorithms based on flips circumvent the simpliciality requirement [27], and the idea of flips generalizes to higher dimensions in Euclidean space [38], we stick to Bowyer-Watson’s algorithm as its generalization to 3 (and higher) dimensions

is straightforward for both Euclidean and periodic point sets. Furthermore it allows for an efficient, clean, and easily maintainable implementation. The first phase of our algorithm (Section 7.2.1) refines the algorithm by Dolbilin and Huson [29], and its implementation details are based on some new results. It uses 3^d copies of each input point, regardless of the lattice, and computes a finite Euclidean Delaunay triangulation on this point set, from which a representation of the Delaunay triangulation on the d -torus is obtained. Once a simpliciality criterion is met, our algorithm switches to the second phase (Section 7.2.3), which conceptually follows the CGAL implementation of the cubic case [16]. It operates directly on the d -torus, maintaining only one copy of each input point, and thus provides better insertion running times than phase 1. A first version of our open-source implementation in 2D and 3D is available on [github](https://github.com/MaelRL/cgal/blob/Generative_Periodic_Triangulations/README.md).¹ We are working on integration into CGAL. Experiments (Section 7.3) show similar performances as the CGAL implementation restricted to cubic lattices [17]. We close with a discussion of future extensions in Section 7.4.

7.1 Preliminaries

Let us recall various notions [39; 23] that are employed throughout the algorithm. Let $B = \{b_1, b_2, \dots, b_d\}$ be a basis of \mathbb{R}^d . The point set $\Lambda := \{\sum_{i=1}^d z_i b_i : z_i \in \mathbb{Z}\}$ is called a *lattice*, and B is its *lattice basis*. The lattice Λ is associated with the translation group Γ consisting of the translations $\phi_\lambda : \mathbb{R}^d \rightarrow \mathbb{R}^d$ mapping the origin to λ , for each $\lambda \in \Lambda$. The group Γ acts on \mathbb{R}^d and each of its translations maps Λ onto itself. We denote the length of the shortest non-zero lattice vector as $\text{sv}(\Lambda)$. For a given lattice basis B , we call $B \cup \{b_0\}$ with $b_0 = -\sum_{i=1}^d b_i$ its *superbase*. A superbase is *obtuse* if for every pair b_i and b_j , $\langle b_i, b_j \rangle \leq 0$. A basis is *reduced* if its superbase is obtuse [39, Definition 4.4]. This notion is defined in such a way that we can easily compute $\text{sv}(\Lambda)$ and Dirichlet domains. The *Dirichlet domain* of a lattice point $\lambda \in \Lambda$ is the region of λ in the Voronoi tessellation of Λ , or more formally $\text{dom}(\lambda, \Lambda) := \{x \in \mathbb{R}^d : \|x - \lambda\| \leq \|x - \nu\| \forall \nu \in \Lambda\}$. It is a convex polytope, and we call the lattice Λ *generic* if each vertex of $\text{dom}(0, \Lambda)$ is incident to the Dirichlet domains of exactly d other lattice points. For 2-dimensional generic lattices the Dirichlet domains are hexagons, for 3-dimensional generic lattices they are combinatorially

¹ https://github.com/MaelRL/cgal/blob/Generative_Periodic_Triangulations/README.md

equivalent to truncated octahedra.

For a lattice vector λ , $\text{str}(\lambda) = \{x \in \mathbb{R}^d : -0.5 \leq \frac{\langle x, \lambda \rangle}{\langle \lambda, \lambda \rangle} < 0.5\}$ is an infinite half-open strip that contains the subspace orthogonal to λ through the origin. Then $\text{dom}(0, \Lambda)$ is the closure of the intersection of these strips for all non-zero lattice vectors. However, as $\text{dom}(0, \Lambda)$ only has a finite number of facets, there must be a finite subset of strips whose closed intersection yields $\text{dom}(0, \Lambda)$. Let \mathcal{V} be the minimal set of lattice vectors (together with their negatives) such that the closure of $\bigcap_{v \in \mathcal{V}} \text{str}(v)$ is $\text{dom}(0, \Lambda)$. The vectors in \mathcal{V} are commonly called *Voronoi-relevant vectors*. Each of them is a normal vector of a facet of the Dirichlet domain of 0, and thus their number is bounded by $2(2^d - 1)$ [39, Theorem 3.6]. Let $\mathcal{V}^+ \sqcup \mathcal{V}^-$ be a partition of \mathcal{V} such that if $v \in \mathcal{V}^+$, then $-v \in \mathcal{V}^-$, and vice versa. For a fixed choice of \mathcal{V}^+ (and implicitly \mathcal{V}^-), we define the *canonical* domain $\text{dom}^\square(0, \Lambda) := \bigcap_{v \in \mathcal{V}^+} \text{str}(v)$ (Figure 7.2). Its closure is $\text{dom}(0, \Lambda)$. We denote the images of $\text{dom}^\square(0, \Lambda)$ under Γ as $\text{dom}^\square(\lambda, \Lambda) := \phi_\lambda(\text{dom}^\square(0, \Lambda))$ for $\lambda \in \Lambda$, and note that similarly the closure of $\text{dom}^\square(\lambda, \Lambda)$ is $\text{dom}(\lambda, \Lambda)$. With $k\Lambda$ for $k \in \mathbb{Z}$ referring to the lattice with basis $\{kb_1, \dots, kb_d\}$, we note that $\text{dom}^\square(0, k\Lambda)$ is $\text{dom}^\square(0, \Lambda)$ scaled by a factor of k .

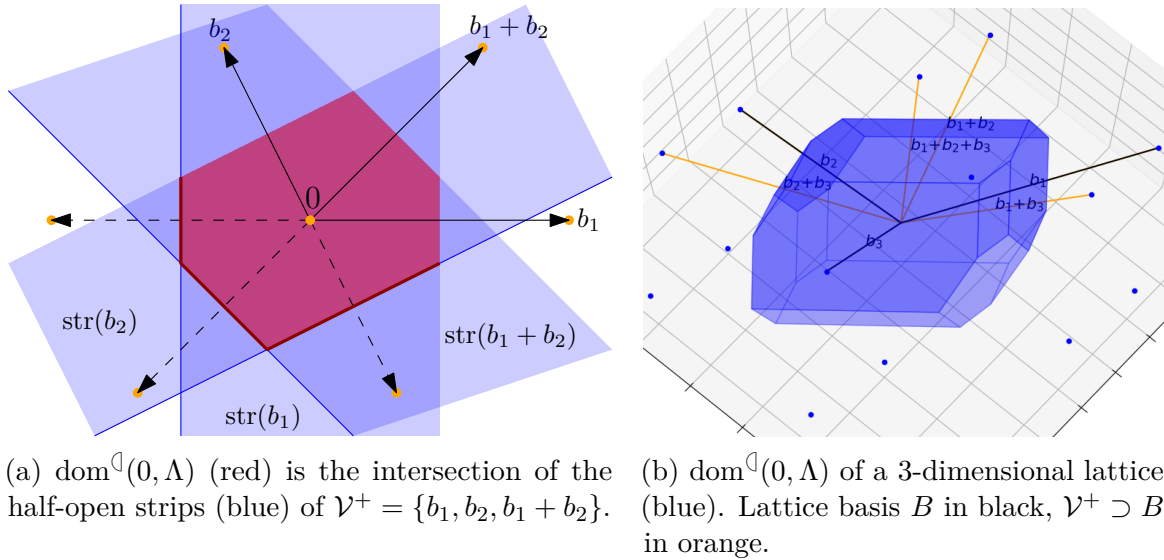


Figure 7.2: Canonical domains and Voronoi relevant vectors for lattices in 2D and 3D.

For $d \leq 3$, if we have a reduced lattice basis B with its superbase B_{sup} , then \mathcal{V} is a subset of $\{\sum_{b \in S} b : S \subset B_{\text{sup}}, S \neq \emptyset, S \neq B_{\text{sup}}\}$ [23, Theorems 3 and 8], with equality if the lattice is generic. We choose $\mathcal{V}^+ := \{\sum_{b \in S} b : S \subseteq B, S \neq \emptyset\} \cap \mathcal{V}$. This choice of \mathcal{V}^+ ensures the property that the translated domains $\text{dom}^\square(\lambda, \Lambda)$ for $\lambda \in \Lambda$ form a partition of \mathbb{R}^d . Note that $\text{sv}(\Lambda)$ can be obtained as the length of the shortest vector in \mathcal{V} .

Delaunay triangulations Recall that the Delaunay triangulation $\text{Del}(P)$ of an input point set $P \subset \mathbb{R}^d$ is a collection of simplices up to dimension d whose vertex set is P and each d -simplex (which in this chapter we simply refer to as *cells*) corresponds to a set of $d + 1$ points whose circumsphere does not contain any other points of P in its interior. While we previously assumed the input point set to be finite and in general position, the definition generalizes to infinite point sets as long as they are discrete, i.e. every ball of finite radius only contains a finite number of points of P . We can overcome the assumption of general position by using symbolic perturbation, see below.

Given a lattice Λ and a finite set of points P , we get the *periodic point set* $\Gamma P := \{\phi_\lambda(p) : p \in P \text{ and } \lambda \in \Lambda\}$ consisting of the elements of the orbits of P under Γ . ΓP is globally invariant under Γ , i.e. each translation of Γ maps ΓP onto itself. Then $\text{Del}(\Gamma P)$ is the *periodic Delaunay triangulation* of the infinite point set ΓP . Note that we can ignore degeneracies in ΓP by using the symbolic perturbation provided by CGAL [28]: it is translation-invariant, so, degeneracies are triangulated in a consistent way, which ensures that the computed $\text{Del}(\Gamma P)$ is actually invariant under Γ . The orbit space \mathbb{R}^d/Γ is a *flat torus*, and we denote its projection map as $\pi : \mathbb{R}^d \rightarrow \mathbb{R}^d/\Gamma$. The *torus triangulation* $\text{Del}(\Gamma P)/\Gamma$ is the projection of $\text{Del}(\Gamma P)$ into \mathbb{R}^d/Γ . Using $\text{dom}^\square(0, \Lambda)$ as a geometric representation of the torus, we can use $P_0 := \Gamma P \cap \text{dom}^\square(0, \Lambda)$ as *canonical* representatives of the vertex set of $\text{Del}(\Gamma P)/\Gamma$. While $\text{Del}(\Gamma P)/\Gamma$ gives us a finite representation of $\text{Del}(\Gamma P)$, unlike $\text{Del}(\Gamma P)$ it is not necessarily simplicial (see Figure 7.1a).

7.2 Algorithm

The input to our algorithm is a lattice basis B' for Λ and a set of points P defining the periodic point set. The output is an object representing $\text{Del}(\Gamma P)/\Gamma$. This object provides a uniform interface that, regardless of the internal state of our algorithm, allows the user to access the properties of Λ as well as the torus triangulation $\text{Del}(\Gamma P)/\Gamma$. For Λ this includes the reduced basis $B = \{b_1, \dots, b_d\}$. For $\text{Del}(\Gamma P)/\Gamma$ this includes the set P_0 of canonical representatives for its vertex set, and its cells. Cells are not solely defined by their vertex set, but have additional geometric information attached. Specifically, each vertex of a cell is represented as a point p from P_0 with an associated offset $o = (o_1, \dots, o_d)$, which is an integer vector. The geometric location of the vertex then is $p + \sum_{i=1}^d o_i b_i$. In alignment

with other CGAL triangulations, we also provide access to simplices of lower dimensions, represented with associated vertex offsets akin to cells, as well as neighborhood relations. These include querying a cell for its adjacent cells, or querying a vertex for its incident cells or lower dimensional simplices. While many steps generalize, we restrict our focus to 2- and 3-dimensional triangulations which are the most widely used.

Internally, our algorithm operates in two phases, which use two different data structures: The first phase maintains a finite Euclidean Delaunay triangulation while the second phase maintains a triangulation of $\text{Del}(\Gamma P)/\Gamma$.

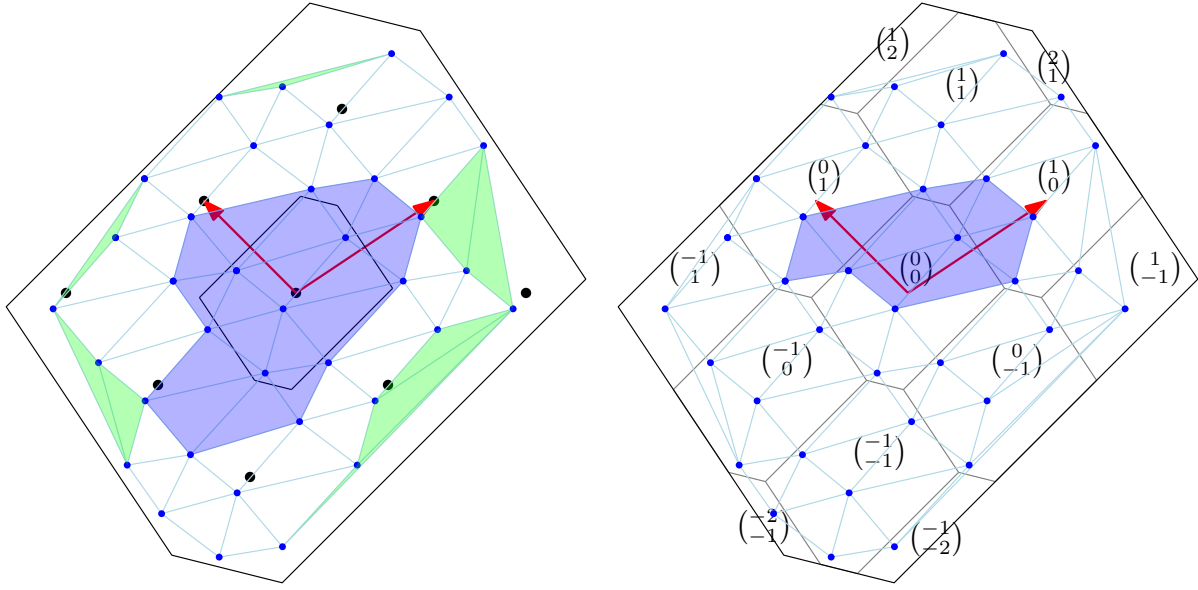
Let us now recall the result that is crucial to the first phase of our algorithm. For a point set P , let $P_3 := \text{dom}^\square(0, 3\Lambda) \cap \Gamma P$, i.e. all periodic copies of P that lie within the Dirichlet domain of 0 scaled by a factor of 3. We call a cell of $\text{Del}(P_3)$ a *periodic cell* if it is also a cell of the periodic triangulation $\text{Del}(\Gamma P)$ (see Figure 7.3a).

Proposition 7.1 ([29, Lemma 3.4]). *Given a point set P , each cell of $\text{Del}(P_3)$ that has at least one vertex in $\text{dom}^\square(0, \Lambda)$ is a periodic cell. Furthermore the set of these cells contains at least one periodic copy of each cell of $\text{Del}(\Gamma P)/\Gamma$.*

After some preprocessing that essentially consists in computing the canonical domain, the first phase internally maintains $\text{Del}(P_3)$ using the CGAL packages for Euclidean Delaunay triangulations [51; 93]. In Section 7.2.1 we develop a systematic way of computing P_3 and obtaining the interface for $\text{Del}(\Gamma P)/\Gamma$ from the internal data structure $\text{Del}(P_3)$.

Once we can guarantee that $\text{Del}(\Gamma P)/\Gamma$ is simplicial and will remain so for any future point insertions, we switch to the second phase. Simpliciality guarantees that the Bowyer-Watson algorithm can be used directly on $\text{Del}(\Gamma P)/\Gamma$. For this purpose we leverage the CGAL machinery for periodic triangulations from the cubic setting [57; 17] and enhance its underlying data structures to work for the generic setting. As we keep only one representative for each vertex, inserting points in this phase is more efficient than in the first phase.

The remainder of this section describes the two phases and the transition between them.



(a) Cells in blue are guaranteed to be periodic cells by Proposition 7.1. Green: non-periodic cells. Black vertices: Λ . (b) The set of canonical cells in blue. Each copy of $\text{dom}(0, \Lambda)$ is labeled with its offset.

Figure 7.3: In blue, the points P_3 with their Delaunay triangulation $\text{Del}(P_3)$. The large hexagon is $\text{dom}(0, 3\Lambda)$, while the smaller ones are $\text{dom}(0, \Lambda)$ and on the right also its periodic copies. The reduced lattice basis is drawn in red.

7.2.1 Phase 1

In a preprocessing step, we first compute the reduced lattice basis B , which allows us to compute \mathcal{V} , the face normals of the canonical domain. We use \mathcal{V} to represent $\text{dom}^\square(0, \Lambda)$ and check for containment of a point within $\text{dom}^\square(0, \Lambda)$, which helps obtaining the canonical copy of each input point. The main part of phase 1 maintains the Euclidean Delaunay triangulation of P_3 . For each new point to be inserted, we first find its canonical copy, then we compute its periodic copies that are contained within $\text{dom}^\square(0, 3\Lambda)$ using Lemmas 7.2 and 7.3, which are then inserted into $\text{Del}(P_3)$. To provide user access to the cells of $\text{Del}(\Gamma P)/\Gamma$, we define a notion of canonical cell in $\text{Del}(P_3)$ to get a representative for each cell of $\text{Del}(\Gamma P)/\Gamma$.

Lattice reduction

Many lattice related problems, such as the shortest non-zero vector problem (SVP), are believed to be hard in general [86; 4]. However in low dimensions, we can use two classical iterative algorithms to solve lattice reduction. Let B' be a lattice basis and $B'_{\text{sup}} = \{b'_0, \dots, b'_d\}$ its superbase, as defined in Section 7.1. Define $c_{ij} := \langle b'_i, b'_j \rangle$.

2-dimensional reduction A 2-dimensional lattice basis is Lagrange-reduced [39, Section 4.2] if $0 \leq 2|c_{12}| \leq c_{11} \leq c_{22}$. We can negate b'_2 if necessary so that $c_{12} \leq 0$. Then $\{b'_1, b'_2, b'_0 := -b'_1 - b'_2\}$ form an obtuse superbase, because $c_{01} = \langle b'_1, -b'_1 - b'_2 \rangle = -c_{11} - c_{12} \leq 0$ due to $2|c_{12}| \leq c_{11}$, and similarly $c_{02} \leq 0$. If a basis is not Lagrange-reduced with $2|c_{12}| > c_{11}$, we exchange b'_2 for a shorter vector that forms a basis with b'_1 : Let $b''_2 := b'_2 - sb'_1$, with $s = 1$ if $c_{12} > 0$ and $s = -1$ otherwise. Vector b''_2 is shorter than b'_2 since $c'_{22} = c_{22} + c_{11} - 2sc_{12}$. Since there are only finitely many pairs of lattice vectors within a ball of any given radius, a finite number of applications of this procedure will yield a Lagrange-reduced basis.

3-dimensional reduction We outline Selling's algorithm [39, Section 4.4], an iterative algorithm that obtains an obtuse superbase in 3 dimensions. In each step of the algorithm, if there is a $c_{ij} > 0$, the algorithm returns a new superbase B''_{sup} defined via $b''_i := -b'_i$, $b''_j := b'_j$, $b''_h := b'_h + b'_i$ and $b''_k := b'_k + b'_i$ where h and k are the remaining two indices different from i and j . For any basis B , let $\sigma(B) := \sum_{b \in B_{\text{sup}}} \|b\|^2$. Notice that in each step, $\sigma(B'') = \sigma(B') - 2c_{ij}$. In particular, $\sigma(B'') < \sigma(B')$. This fact, together with the fact that there are only finitely many lattice points in a ball of radius $\sqrt{\sigma(B')}$ (and thus only finitely many quadruplets of vectors whose square magnitudes sum up to at most $\sigma(B')$) guarantees termination of the algorithm.

Canonical points

For each point of P , finding its periodic copy that lies in $\text{dom}^{\mathbb{Q}}(0, \Lambda)$ is equivalent to solving the closest vector problem (CVP), i.e. given $p \in P$, determining the lattice point which is closest to p . For arbitrary dimensions this problem is known to be NP-hard [86]. For the exact version of CVP, various iterative algorithms have been described [3; 79; 62]. As we are only operating in 2 and 3 dimensions, any of them would suffice for us in practice in terms of running time, and we will describe the algorithm by Sommer et al [79] due to its simplicity.

For a real number r , define $\text{round}(r)$ to be the closest integer to r . If this integer is not unique, then it is the one with the smallest absolute value. This definition ensures convergence in cases where p is on the boundary of a Dirichlet domain of the lattice. Recall that \mathcal{V}^+ is a set of normals of the facets of the canonical domain, and can be obtained from

a reduced basis B (see Section 7.1). We first sort the vectors of \mathcal{V}^+ by their magnitude. As $\text{dom}^\square(0, \Lambda)$ is the intersection of the strips $\text{str}(v)$ for $v \in \mathcal{V}^+$, we need to find the periodic copy of p that is in all these $\text{str}(v)$. We loop through the vectors v of \mathcal{V}^+ and for each of them perform the following operations: Compute $c := \frac{\langle p, v \rangle}{\langle v, v \rangle}$. If $-0.5 \leq c < 0.5$, then already $p \in \text{str}(v)$. If not, then we subtract $\text{round}(c) \cdot v$ from p . Note that after such a step, it is guaranteed that $p \in \text{str}(v)$. However after modifying p for a longer vector v , it might happen that p is moved outside of $\text{str}(v')$ for some shorter vector v' again. Therefore we need to repeatedly loop through \mathcal{V}^+ and perform this operation until $p \in \text{str}(v)$ for all $v \in \mathcal{V}^+$.

Notice that in each step where p is modified, the magnitude of p strictly decreases. Therefore this algorithm terminates in finite time, as the number of lattice vectors within a ball of given radius $\|p\|$ is finite.

Computing P_3

For a given point $p \in P_0$, each of its periodic copies $q \in \Gamma p$ can be obtained as $\phi_\lambda(p)$ for some translation $\phi_\lambda \in \Gamma$. The corresponding lattice point λ can be uniquely written as $\sum_{i=1}^d o_i b_i$ for some $o \in \mathbb{Z}^d$ and we write $\lambda(o) := \lambda$. We then call $o =: o(q)$ the *offset* of q and $\phi_{\lambda(o)}$ the *translation* associated with o . For each point $p \in P_0$ we need to determine the 3^d offsets for which $\phi_{\lambda(o)}(p)$ is within $\text{dom}^\square(0, 3\Lambda)$. Fortunately, these offsets have to come from a fixed set of offsets that only depends on the lattice basis. Notice that if $\phi_{\lambda(o)}(\text{dom}^\square(0, \Lambda))$ does not intersect $\text{dom}^\square(0, 3\Lambda)$, then $\phi_{\lambda(o)}(p)$ cannot be in $\text{dom}^\square(0, 3\Lambda)$ for $p \in P_0$. Thus we only need to check those offsets for which $\phi_{\lambda(o)}(\text{dom}^\square(0, \Lambda)) \cap \text{dom}^\square(0, 3\Lambda) \neq \emptyset$.

Lemma 7.2. *There are at most $4^d - 2^d + 1$ translates of $\text{dom}^\square(0, \Lambda)$ that have non-empty intersection with $\text{dom}^\square(0, 3\Lambda)$.*

Proof. If $\text{dom}^\square(\lambda, \Lambda)$ intersects $\text{dom}^\square(0, 3\Lambda)$ for some $\lambda \in \Lambda$, then λ must be inside $\text{dom}^\square(0, 4\Lambda)$. There are 4^d lattice points within $\text{dom}^\square(0, 4\Lambda)$. For each $v \in \mathcal{V}^+$, the lattice vectors $-2v$ and $2v$ are on the boundary of $\text{dom}^\square(0, 4\Lambda)$, however only $-2v$ is one of those 4^d points within $\text{dom}^\square(0, 4\Lambda)$. For these lattice points $-2v$ on the boundary however the intersection between $\text{dom}^\square(-2v, \Lambda)$ and $\text{dom}^\square(0, 3\Lambda)$ is empty. Thus only $4^d - \#\mathcal{V}^+ = 4^d - (2^d - 1)$ translates of $\text{dom}^\square(0, \Lambda)$ intersect $\text{dom}^\square(0, 3\Lambda)$. \square

In general, we can find this set of offsets via a breadth-first search: We define a graph on Λ with each lattice point λ connected to $\lambda + v$ for $v \in \mathcal{V}$. We start the search at lattice point 0 and terminate once we have found $4^d - 2^d + 1$ offsets, or in the case of non-generic lattices once all 4^d lattice points inside $\text{dom}^\square(0, 4\Lambda)$ have been reached (with the ones on the boundary being discarded).

In 2 dimensions, there is in fact a fixed set S of 13 offsets such that for any lattice, $\text{dom}^\square(\lambda(o), \Lambda)$ only intersects $\text{dom}^\square(0, 3\Lambda)$ when $o \in S$ (Figure 7.3b).

Lemma 7.3. *Given a 2-dimensional lattice and an offset o , if $o \notin \{(\emptyset, \emptyset), (-1, -1), (\emptyset, 1), (1, \emptyset), (-1, \emptyset), (\emptyset, -1), (1, 1), (-1, -2), (1, 2), (-2, -1), (2, 1), (-1, 1), (1, -1)\}$, then $\text{dom}^\square(\lambda(o), \Lambda) \cap \text{dom}^\square(0, 3\Lambda) = \emptyset$.*

Proof. Each $\text{dom}^\square(\lambda, \Lambda)$ adjacent to $\text{dom}^\square(0, \Lambda)$ has all but one of its facets in the interior of $\text{dom}^\square(0, 3\Lambda)$. Except for those adjacent via one of those facets, all $\text{dom}^\square(\lambda, \Lambda)$ at graph distance 2 from 0 intersect $\text{dom}^\square(0, 3\Lambda)$. The ones that don't are exactly the 6 domains $\text{dom}^\square(2v, \Lambda)$ for $v \in \mathcal{V}$. As there are 19 domains at distance at most 2, we have found 13 domains that intersect $\text{dom}^\square(0, 3\Lambda)$. As this is the maximum possible number by Lemma 7.2, we have found all of them. \square

In 3 dimensions, the same argument yields a set of 51 fixed offsets. As the total number of domains that intersect $\text{dom}^\square(0, 3\Lambda)$ is 57 by Lemma 7.2 however, we need to find the remaining 6 domains. They do not always have the same offsets, however their offsets come from a fixed set of 24 offsets.

Lemma 7.4. *Let $O_3 := \{(3, 2, 1), (2, 1, -1), (3, 1, 2), (2, -1, 1), (1, -1, -2), (1, -2, -1), (2, 3, 1), (1, 2, -1), (1, 3, 2), (-1, 2, 1), (-1, 1, -2), (-2, 1, -1), (2, 1, 3), (1, -1, 2), (1, 2, 3), (-1, 1, 2), (-1, -2, 1), (-2, -1, 1), (-1, -2, -3), (-1, -3, -2), (-2, -1, -3), (-3, -1, -2), (-2, -3, -1), (-3, -2, -1)\}$*

and $O_{\leq 2} := \{(\emptyset, \emptyset, \emptyset), (-1, -1, -1), (\emptyset, \emptyset, 1), (\emptyset, 1, \emptyset), (1, \emptyset, \emptyset), (1, 1, \emptyset), (1, \emptyset, 1), (\emptyset, -1, -1), (\emptyset, 1, 1), (-1, \emptyset, -1), (-1, -1, \emptyset), (1, 1, 1), (\emptyset, \emptyset, -1), (\emptyset, -1, \emptyset), (-1, \emptyset, \emptyset), (1, 2, \emptyset), (1, \emptyset, 2), (-1, -2, -2), (2, 1, \emptyset), (2, \emptyset, 1), (1, -1, -1), (\emptyset, 1, 2), (-2, -1, -2), (\emptyset, 2, 1), (-1, 1, -1), (-2, -2, -1), (-1, -1, 1), (1, 1, 2), (-1, -1, -2),$

$(1, 2, 1), (\emptyset, 1, -1), (-1, -2, -1), (\emptyset, -1, 1), (2, 1, 1), (1, \emptyset, -1),$
 $(1, -1, \emptyset), (-2, -1, -1), (-1, \emptyset, 1), (-1, 1, \emptyset), (1, 2, 2), (-1, \emptyset, -2),$
 $(-1, -2, \emptyset), (2, 1, 2), (\emptyset, -1, -2), (2, 2, 1), (1, 1, -1), (\emptyset, -2, -1),$
 $(1, -1, 1), (-2, -1, \emptyset), (-2, \emptyset, -1), (-1, 1, 1)\}.$

Given a 3-dimensional lattice and an offset o , there is a subset S of size 6 of O_3 such that if $o \notin O_{\leq 2} \cup S$, then $\text{dom}^{\mathbb{Q}}(\lambda(o), \Lambda) \cap \text{dom}^{\mathbb{Q}}(0, 3\Lambda) = \emptyset$.

Proof. As mentioned in the proof of Theorem 3.5 from [39], $\frac{v}{2}$ is on the boundary of $\text{dom}^{\mathbb{Q}}(0, \Lambda)$ for every Voronoi-relevant vector v .

Claim 1: The orthogonal projection of v into the 1-dimensional subspace spanned by another Voronoi relevant vector $w \in \mathcal{V}^+$ has magnitude less than w .

Proof of Claim 1: All facets of $\text{dom}^{\mathbb{Q}}(0, \Lambda)$ are contained in $\text{str}(w)$ for all $w \in \mathcal{V}^+$, so because $\frac{v}{2}$ is on the boundary of $\text{dom}^{\mathbb{Q}}(0, \Lambda)$, it is contained in $\text{str}(w)$. By definition of $\text{str}(w)$ this implies $|\langle \frac{1}{2}v, w \rangle / \langle w, w \rangle| \leq \frac{1}{2}$, from which it follows that $|\langle v, w \rangle| \leq \langle w, w \rangle$.

Claim 2: Domains $\text{dom}^{\mathbb{Q}}(\nu, \Lambda)$ of graph-distance $(k - 2)$ from some $\lambda = kv$ for an integer k and $v \in \mathcal{V}$ cannot intersect $\text{dom}^{\mathbb{Q}}(0, 3\Lambda)$.

Proof of Claim 2: We have $\nu = kv + v_1 + \dots + v_{k-2}$ for some $v_i \in \mathcal{V}$. Then $\langle w, \nu \rangle = k\langle v, v \rangle + \langle v, v_1 \rangle + \dots + \langle v, v_{k-2} \rangle \geq k\langle v, v \rangle - (k - 2)\langle v, v \rangle = 2\langle v, v \rangle$ due to Claim 1. This means that ν is not in the interior of $\text{str}(4v)$ and thus the interior of the 4-scaled domain. Therefore $\text{dom}^{\mathbb{Q}}(\nu, \Lambda) \cap \text{dom}^{\mathbb{Q}}(0, 3\Lambda) = \emptyset$.

Proof of Lemma: Let $B_{\text{sup}} = \{b_0, b_1, b_2, b_3\}$. Note that because $b_0 = -(b_1 + b_2 + b_3)$, every lattice point λ can be written as a non-negative integer combination of three of these extended basis vectors, i.e. $\lambda = c_1a + c_2b + c_3c$ with $c_i \in \mathbb{Z}, c_1 \geq c_2 \geq c_3 \geq 0$ and $a, b, c \in B_{\text{sup}}$. This representation is unique up to permutation of basis vectors with the same coefficient.

Enumerating all $\text{dom}^{\mathbb{Q}}(\lambda, \Lambda)$ at graph distance 3 from $\text{dom}^{\mathbb{Q}}(0, \Lambda)$, we get lattice points

λ with non-negative integer combinations of the following types:

$$\begin{aligned}
& 3a \\
& 3a + b = 3a + b \\
& 3a + b + c = 3a + (b + c) \\
& 3a + 2b = 3(a + b) - b \\
& 3a + 2b + c \\
& 3a + 2b + 2c = 3(a + b + c) - (b + c) \\
& 3a + 3b \\
& 3a + 3b + c = 3(a + b) + c \\
& 3a + 3b + 2c = 3(a + b + c) - c \\
& 3a + 3b + 3c
\end{aligned}$$

All of these, except for type $3a + 2b + c$, can be written as $3v$ or $3v + w$ for some Voronoi relevant vectors $v, w \in \mathcal{V}$. This means that they are, or are adjacent to, a domain centered at $3v$, and from Claim 2 it follows that they cannot intersect $\text{dom}^\square(0, 3\Lambda)$. A similar argument shows that none of the domains at graph distance 4 from $\text{dom}^\square(0, \Lambda)$ can intersect $\text{dom}^\square(0, 3\Lambda)$, and therefore also none at higher graph distance.

Now the 51 offsets from $O_{\leq 2}$ are those that the argument from Lemma 7.3 yields, while O_3 contains the offsets corresponding to type $3a + 2b + c$. As by Lemma 7.2 there are 57 intersecting domains, only 6 of the offsets from O_3 can correspond to intersecting domains. \square

Extracting representative cells

Recall that our interface specifies access to the cells (and lower-dimensional simplices) of $\text{Del}(\Gamma P)/\Gamma$. Thus for each of its cells we have to provide one representative from $\text{Del}(P_3)$. For a cell σ of either $\text{Del}(\Gamma P)$ or $\text{Del}(P_3)$, with $V(\sigma) = \{p_1, \dots, p_d\}$, we define its offset as the vector $o(\sigma) = \min_{p \in V(\sigma)} \{o(p)\}$ where the minimum is taken lexicographically. Note that this definition differs from [16, Convention 3.3.1] where the coordinate-wise minimum is taken. If the offset of a cell is the 0-vector, we call it a *canonical cell*. Note that due to our definition a canonical cell always has a vertex with offset 0. Therefore by Proposition 7.1 a

canonical cell of $\text{Del}(P_3)$ is also periodic, see Figure 7.3b. For $\sigma \in \text{Del}(\Gamma P)$, the translated cell $\sigma - \sum_{i=1}^d o_i b_i$ is called its *canonical representative*. This means that for each class of cells in $\text{Del}(\Gamma P)$ (or equivalently each cell of $\text{Del}(\Gamma P)/\Gamma$), there is a unique canonical representative in $\text{Del}(P_3)$. Therefore we can get a set of representative cells by iterating over the cells of $\text{Del}(P_3)$ and selecting those that are canonical.

Neighborhood relations

Recall that our interface dictates that we provide neighborhood relations for the vertices and cells of $\text{Del}(\Gamma P)/\Gamma$. These vertices and cells are represented by the canonical vertices and cells of $\text{Del}(P_3)$, whose neighbors in $\text{Del}(P_3)$ (which we have access to) may differ from the neighbors in $\text{Del}(\Gamma P)/\Gamma$ (which we want to find). We outline how to get the neighbors of a cell of $\text{Del}(\Gamma P)/\Gamma$ from $\text{Del}(P_3)$, and note that other neighborhood relations work in a conceptually similar way. Note that we do not store these relations explicitly, but we compute them upon request and may cache them for future access.

Consider a canonical cell σ of $\text{Del}(P_3)$ and a neighboring cell σ_n . If σ_n is canonical, it is the neighbor of σ in $\text{Del}(\Gamma P)$. If σ_n is not canonical but has a vertex in $\text{dom}^\square(0, \Lambda)$, then it is periodic and we return the canonical representative of this cell. However, it is possible that all vertices of σ_n are outside $\text{dom}^\square(0, \Lambda)$, and thus σ_n might not be a periodic cell at all. In that case, we need to consider the facet τ separating σ and its neighbor σ_n . As τ is a facet of a canonical cell, it is also a periodic facet (i.e. a facet of $\text{Del}(\Gamma P)$). We compute τ 's offset $o := o(\tau)$. Then $\tau' := \phi_{\lambda(-o)}(\tau)$ is the canonical representative of τ and $\sigma' := \phi_{\lambda(-o)}(\sigma)$ as well as its neighbor σ'_n across τ' are periodic cells as they share τ' which has a vertex in $\text{dom}^\square(0, \Lambda)$. If σ'_n is canonical, then it is the canonical representative of the neighbor of σ in $\text{Del}(\Gamma P)$; if not, then its canonical representative is. Figure 7.4 illustrates this process.

In practice, we store the canonical representative of each vertex of P_3 . To obtain the canonical representative of a cell (or facet) σ , we need to choose a vertex p whose offset is minimal (lexicographically) among its vertices. This ensures that the periodic copy p' of p in the canonical version of the cell is inside $\text{dom}^\square(0, \Lambda)$. Then one of the cells (or facets) incident to p' in $\text{Del}(P_3)$ is the canonical representative of σ .




Figure 7.4: Finding the canonical neighbor of σ across the edge τ , which is σ''_n , the canonical representative of σ'_n . In blue, the points P_3 with their Delaunay triangulation $\text{Del}(P_3)$. The large hexagon is $\text{dom}(0, 3\Lambda)$, while the smaller ones are $\text{dom}(0, \Lambda)$ and its periodic copies labeled with their respective offsets. Pale blue cells are canonical cells.

7.2.2 Transition

Phase 2 is more efficient than phase 1 as it directly operates on $\text{Del}(\Gamma P)/\Gamma$, however we cannot use it from the start as the Bowyer-Watson algorithm [12; 89] comes with some constraints. The Bowyer-Watson algorithm is an incremental algorithm inserting points one by one. For each new point p , it determines the *conflict zone*, which is the set of cells whose circumsphere contains p . All these cells are removed, and all boundary facets of the resulting hole are connected to p to fill in the hole with new cells. The algorithm requires this hole to be a topological d -ball, which is not always guaranteed for $\text{Del}(\Gamma P)/\Gamma$. However the following criterion is a sufficient condition for the Bowyer-Watson algorithm to work [19].

Lemma 7.5 ([19, Criterion 3.11]). *If for every cell in $\text{Del}(\Gamma P)/\Gamma$ the circumradius is smaller than $\frac{1}{4}\text{sv}(\Lambda)$, then $\text{Del}(\Gamma P')/\Gamma$ is simplicial for every $P' \supseteq P$.*

If this criterion is fulfilled, we can safely switch to phase 2. To detect at which point in phase 1 the criterion is fulfilled, we maintain a set S_{big} of *big* cells, which are canonical cells whose circumradius is larger than or equal to $\frac{1}{4}\text{sv}(\Lambda)$. We update S_{big} during each point insertion.

Assume we wish to insert a new point p into the periodic triangulation of some point set P . On a high level, we need to remove the big cells in the conflict zone of p from S_{big} , and then add the newly created big cells to S_{big} . In practice, we have already computed the triangulation $\text{Del}(P_3)$ and have to insert the periodic copies of p that are within $\text{dom}^{\square}(0, 3\Lambda)$, i.e. $\Gamma x \cap \text{dom}^{\square}(0, 3\Lambda)$, into $\text{Del}(P_3)$. We first detect the conflict zone of p_0 , the canonical copy of p . For each big cell in the conflict zone we check if it has a canonical copy in S_{big} , and remove that copy from S_{big} in that case. While the conflict zone may contain non-periodic cells, Lemma 7.6 guarantees that we capture a representative of each cell from $\text{Del}(\Gamma P)$ that is in conflict with p . Next we insert all the copies $\Gamma x \cap \text{dom}^{\square}(0, 3\Lambda)$ into $\text{Del}(P_3)$. Finally, for each big cell incident to p_0 in the resulting triangulation, we add its canonical copy to S_{big} . Note that all these cells have a canonical copy by Proposition 7.1 as they have a vertex inside $\text{dom}^{\square}(0, \Lambda)$.

Lemma 7.6. *Let C be the conflict zone of some point $p \in \text{dom}^{\square}(0, \Lambda)$ with respect to $\text{Del}(\Gamma P)$. Then the conflict zone of p with respect to $\text{Del}(P_3)$ contains at least one periodic copy of each cell from C .*

Proof. First observe that if a cell of $\text{Del}(\Gamma P)$ has its circumcenter at the origin, then all its vertices must be within $\text{dom}(0, \Lambda)$, because if the circumsphere contains a point outside $\text{dom}(0, \Lambda)$, then it also contains its canonical copy. Via translation it follows that if a cell has its circumcenter in $\text{dom}^{\square}(0, 2\Lambda)$, then its vertices are in $\text{dom}^{\square}(0, 3\Lambda)$ and thus it is a cell of $\text{Del}(P_3)$. So assume we have a cell σ in C whose circumcenter c is not within $\text{dom}^{\square}(0, 2\Lambda)$. Then there is a facet of $\text{dom}^{\square}(0, 2\Lambda)$ with respect to which c is outside. Let f be its face normal. Then $p + f$ is also contained in the circumsphere of σ . Reversely, $\sigma - f$ is a cell whose circumsphere contains p , and furthermore its circumcenter is closer to 0 than c , see Figure 7.5. As there are only finitely many periodic copies of c within a given distance from 0, after applying this process a finite number of times we eventually obtain a cell whose circumsphere contains p and whose circumcenter is in $\text{dom}^{\square}(0, 2\Lambda)$. This cell is a periodic copy of σ , is contained in $\text{Del}(P_3)$ and thus also part of the conflict zone of p with respect to $\text{Del}(P_3)$. \square

Once S_{big} is empty, the criterion of Lemma 7.5 is fulfilled, and we can internally convert our triangulation from $\text{Del}(P_3)$ to $\text{Del}(\Gamma P)/\Gamma$, which is maintained in phase 2. We initialize the periodic triangulation data structure with the set of canonical vertices P_0




Figure 7.5: The red point p is in the conflict zone of the blue cell σ which is not a cell of $\text{Del}(P_3)$ and whose circumcenter c is outside $\text{dom}^Q(0, 2\Lambda)$. However the light blue cell $\sigma - f$ is a periodic copy of σ that is in $\text{Del}(P_3)$ and p is in its conflict zone.

and canonical cells obtained from $\text{Del}(P_3)$, as well as the adjacency and incidence relations outlined earlier.

7.2.3 Phase 2

Phase 2 operates directly on the torus triangulation $\text{Del}(\Gamma P)/\Gamma$. Thus it maintains only one copy of each cell and vertex and point insertion is faster than in phase 1. The data structure it uses to represent $\text{Del}(\Gamma P)/\Gamma$ is akin to the one used in CGAL for the cubic case, and closely resembles the interface we defined for our algorithm: Only P_0 is stored as vertex set, and each cell is represented by its vertices, with a vertex encoded as a pair (p, o) of a point $p \in P_0$ and an offset so that its geometric location is $\phi_{\lambda(o)}(p)$. While in the cubic case each offset coordinate either takes the value 0 or 1 and offsets can be encoded in d bits, in our more general setting offsets can take any of the lattice-specific values from Lemma 7.2. Unlike in phase 1, most neighborhood relations for $\text{Del}(\Gamma P)/\Gamma$ are already stored explicitly in the data structure: For each cell its adjacent cells are stored, and for each vertex one incident cell is stored. The remaining neighborhood relations required by our interface are obtained implicitly from the explicitly stored ones. With each point insertion, all stored neighborhood relations have to be updated accordingly.

Point insertion To insert a new point p into $\text{Del}(\Gamma P)/\Gamma$ using the Bowyer-Watson algorithm, first we compute the canonical copy p_0 . Then we locate the cell containing p_0 , via a traversal starting from an arbitrary cell. The conflict zone is computed via a search

starting in the cell containing the point p_0 . Whenever we are traversing cells, care has to be taken to maintain the correct offset of the affected cells relative to p_0 , and similarly when creating new cells to fill in the hole left by the deleted conflict zone. For the cubic case the details are described in [16, Section 3.3], and we omit the technical adjustments needed to make these steps work in the more general case.

7.3 Experimental results

Points until transition We experimentally evaluated the number of points required until the criterion of Lemma 7.5 is fulfilled and the transition to phase 2 occurs. Lemma 7.5 requires all of \mathbb{R}^d/Γ to be covered by the balls of radius $\frac{1}{4}\text{sv}(\Lambda)$ around P_0 . As the volume of the torus equals the volume of the Dirichlet domain of 0, denoted as $\text{vol}(\text{dom}(0, \Lambda))$, we expect the number of points until switching to phase 2 to be roughly proportional to $\text{vol}(\text{dom}(0, \Lambda))/\text{sv}(\Lambda)^d$, assuming the points are sampled uniformly at random.

To investigate this in 2 dimensions, we parametrize a 2-dimensional space of lattices. We call the parameters the elongation ℓ and the skew s . The basis of the lattice with elongation ℓ and skew s is $b_1 = (\ell, 0)$ and $b_2 = (s \cdot \ell/2, 1)$. With $\ell \geq 1$ and s between 0 and 1 we can parametrize all 2-dimensional lattices up to symmetry and scaling. Note that the skew affects $\text{sv}(\Lambda)$ but not $\text{vol}(\text{dom}(0, \Lambda))$, while the elongation is proportional to $\text{vol}(\text{dom}(0, \Lambda))$ but does not affect the $\text{sv}(\Lambda)$. Fixing the skew at 0 and varying the elongation (Figure 7.6a), we see that the number of random points needed until the phase switch appears to be proportional to the elongation. The same applies to the number of points until the resulting triangulation is simplicial for the first time. Figure 7.6b shows the same statistics for lattices of fixed area but varying skew. In addition we plot the inverse of $\text{sv}(\Lambda)^2$ for comparison, and observe that it behaves similarly albeit not entirely proportionally.

In 3-dimensions, a parametrization of lattice up to symmetry and scaling needs 5 parameters, and thus an analysis like in 2D is impractical. Instead we investigate the number of points until the switch to phase 2 relative to $\text{vol}(\text{dom}(0, \Lambda))/\text{sv}(\Lambda)^3$ which we might expect to be proportional. Figure 7.7 shows the result for 182 lattice whose basis vectors have randomly sampled direction and a random magnitude between 1 and 100.

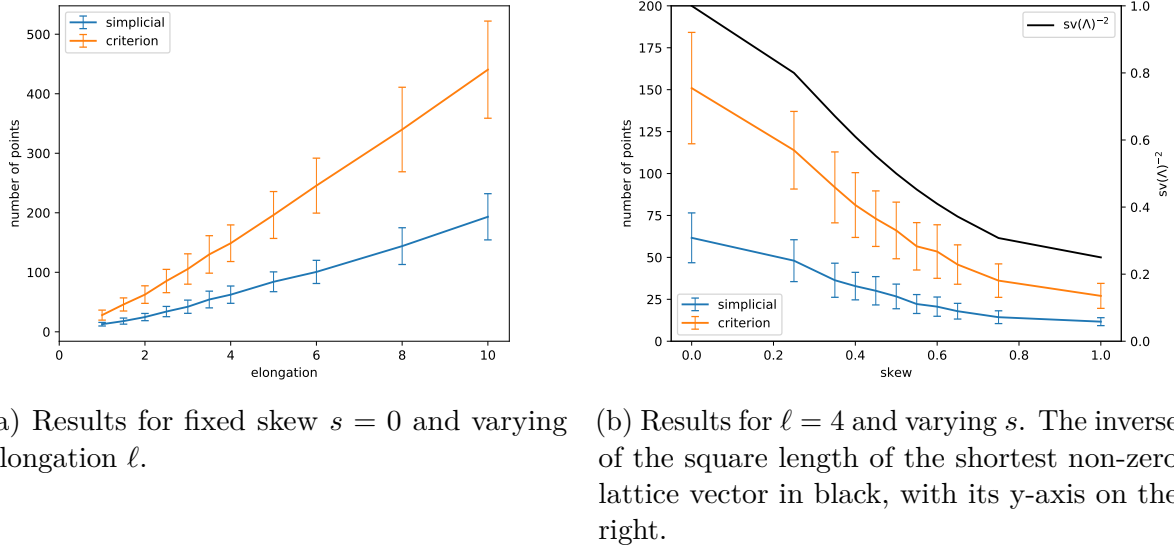


Figure 7.6: For different lattices, the number of points inserted into a periodic triangulation until (blue) $\text{Del}(\Gamma P)/\Gamma$ is simplicial for the first time, and (orange) $\text{Del}(\Gamma P)/\Gamma$ fulfills Lemma 7.5. Each data point is the mean of 200 trials, and the bars represent the standard deviation.

Running times We evaluate the running times of our algorithm for different 3-dimensional lattices, and compare them to the existing CGAL implementations. The data points collected are from Delaunay triangulations of 10^k uniformly sampled random points for k up to 7. Each data point is an average of 300 trials (10 trials for 10^k points). The experiments were conducted on a laptop running Fedora 30 64-bits, with two 6-core Intel(R) i9-8950HK CPU clocked at 2.90GHz, and with 32GB of RAM. The CGAL kernel used was `CGAL::Exact_predicates_inexact_constructions_kernel` and CPU time was measured using CGAL's timer tools. The code was compiled using clang 8.0.0 with compilation flags `-O3` and `-DNDEBUG`.

Table 7.1 shows a comparison between the CGAL implementation of Euclidean Delaunay triangulations [51] (with random points uniformly sampled in the unit cube), periodic Delaunay triangulations in the cubic setting [17] and our algorithm for the cubic lattice. Table 7.2 shows the running time of our algorithm for various lattices. For each lattice, we also measured the average number of points until the switch to phase 2. As our algorithm and the one for the cubic setting are based on the same code base when operating directly on the torus triangulation, their runtimes are comparable for large point sets. It should be noted that when similar experiments were run in 2010 to compare the Euclidean and cubic periodic algorithms [16, Section 3.6.2], both were performing comparably. Since then,




Figure 7.7: Phase switching in 3D: Each point in the plot represents a lattice, with $\text{vol}(\text{dom}(0, \Lambda))/\text{sv}(\Lambda)^3$ on the x -axis and the number of points inserted until the switch to phase 2 occurred on the y -axis. Each y -value is obtained as the average over 200 point sets, each distributed uniformly at random.

Euclidean Delaunay triangulations in CGAL have seen notable optimizations which were not applied to periodic triangulations. This also explains why our algorithm is faster than the cubic periodic algorithm for point sets where phase 1 takes up a significant portion of running time, as internally we use a Euclidean rather than periodic triangulation.

7.4 Discussion

Software distribution We aim to provide a CGAL package for periodic Delaunay triangulations for arbitrary lattices in 2 and 3 dimensions. The current state of our implementation is available online (see Footnote 1). Both the 2- and 3-dimensional implementations have been integrated into the CGAL codebase and only require some additional refactoring and optimizing to adhere to CGAL’s quality standards. Automated tests and documentation still have to be produced. An extension to weighted Delaunay triangulations, referred to as regular triangulations in CGAL, is planned for the future.

Implementation	Euclidean [51]	Cubic [17]	Lattice
10^0	0.0000	0.0001	0.0002
10^1	0.0000	0.0160	0.0033
10^2	0.0004	0.1848	0.0461
10^3	0.0049	0.5957	0.0858
10^4	0.0487	0.9591	0.3832
10^5	0.5679	4.8119	4.5153
10^6	6.5974	93.9327	95.1447
10^7	59.5152	2314.3618	2317.7867

Table 7.1: Running time (in seconds) of various Delaunay triangulation algorithms on random point sets of different sizes. Our algorithm (“Lattice”) is evaluated for the cubic lattice.

Lattice	Cubic	FCC	Λ_1	Λ_2	Λ_3	Λ_4
$\frac{\text{vol}(\text{dom}(0, \Lambda))}{\text{sv}(\Lambda)^3}$	1.0	0.71	0.96	1.80	12.5	346.41
n_{switch}	141	94	134	286	2519	89950
10^0	0.0002	0.0001	0.0005	0.0007	0.0003	0.0004
10^1	0.0033	0.0026	0.0048	0.0067	0.0044	0.0035
10^2	0.0461	0.0287	0.0602	0.1525	0.0460	0.0380
10^3	0.0858	0.0446	0.1042	0.3265	0.9812	0.6372
10^4	0.3832	0.1642	0.4375	0.5767	4.6602	16.5759
10^5	4.5153	2.7868	5.9151	4.4110	10.8139	362.7956
10^6	95.1447	51.5945	90.4999	93.0278	58.1568	517.9715
10^7	2317.7867	1215.2648	2394.9533	2385.1200	1799.4515	2983.0943

Table 7.2: Running times (in seconds) of our algorithm for the cubic lattice, FCC (face-centered cubic) lattice, and four other lattices. For each lattice we also record the $\text{vol}(\text{dom}(0, \Lambda))/\text{sv}(\Lambda)^3$ ratio and the average number of points until the switch to phase 2 (n_{switch}).

Dummy points For best performance, transition from phase 1 to phase 2 should occur as soon as possible. Intuitively, this happens when the input point distribution does not have large gaps. We can achieve this by first inserting additional *dummy points*, which are removed at the end (if possible) [19]. In practice, we can choose these points from a sufficiently fine hexagonal lattice (or body-centered cubic in 3 dimensions), such that the open spheres of radius $\frac{1}{4}\text{sv}(\Lambda)$ around the points cover the entire torus.

7.4.1 Extension to weighted points

Phase 1 of our algorithm readily generalizes to weighted point sets. In particular, Proposition 7.1 still holds for weighted point sets. We provide a proof below. Phase 2 only works for weighted points under additional restrictions because $\text{Del}(\Gamma P)/\Gamma$ can not be guaranteed to remain simplicial, in particular after inserting points with large weights. Such point insertions that break simpliciality can be prevented by requiring points to be inserted in decreasing order of weight, or like in the cubic implementation in CGAL by severely restricting the range of weights a point can have. Thus an implementation is subject to a tradeoff between flexibility (phase 1) and performance (phase 2).

Proposition 7.7 (Generalization of Lemma 3.2–3.4 from [29]). *Given a set P of representatives for our point set, let $P_3 := \text{dom}^\square(0, 3\Lambda) \cap \Gamma P$, i.e. all periodic copies of these points that lie within the Dirichlet domain of 0 scaled by a factor of 3. Let T_0 be those d -cells of $\text{Del}(P_3)$ that have at least one vertex in $\text{dom}^\square(0, \Lambda)$. Then the cells of T_0 are all part of the triangulation of ΓP . Furthermore, these cells contain at least one representative of each class of d -cells from $\text{Del}(\Gamma P)$.*

Proof. We will prove the proposition in 3 steps.

Claim 1: Assume one of our points is $p_0 = 0$ (with arbitrary weight). Then the orthocenter c_T of the vertices of any d -cell of $\text{Del}(\Gamma P)$ which has p_0 as a vertex is within $\text{dom}(0, \Lambda)$.

Proof of Claim 1: Assume not. Then there is a face F_i of $\text{dom}(0, \Lambda)$ such that F_i separates p_0 from c_T . Let f_i be the corresponding translation lattice vector orthogonal to F_i . Then $p_0 + f_i$ is strictly closer to c_T than p_0 because it has the same weight as p_0 . This is a contradiction to the empty-sphere property.

Claim 2: For a d -cell of $\text{Del}(\Gamma P)$ containing p_0 as a vertex, all vertices are in $\text{dom}(0, 2\Lambda)$.

Proof of Claim 2: Assume some vertex p is not. Consider the perpendicular bisector $b_{p-f_i, p}$ between p and $p - f_i$ where f_i is the face normal vector orthogonal to the face F_i of $\text{dom}(0, \Lambda)$ with respect to which p is outside of $\text{dom}(0, 2\Lambda)$. Note that the perpendicular bisector between two points of the same weight, as in this case, is the same as the unweighted perpendicular bisector. The bisector $b_{p-f_i, p}$ is separated from p_0 and the orthocenter c_T by F_i , the face of $\text{dom}(0, \Lambda)$ that is parallel to $b_{p-f_i, p}$. In particular, it follows from Claim 1 that c_T is in the interior of $b_{p-f_i, p}^-$, i.e. on the $p - f_i$ side of $b_{p-f_i, p}$. So $p - f_i$ is closer to c_T than p and thus the empty-sphere property is violated. See Figure 7.8 for an illustration.




Figure 7.8: A cell (dashed) with a vertex p that is outside of $\text{dom}(0, 2\Lambda)$. The dotted line is the perpendicular bisector between p and $p - f_i$.

Claim 3: All d -cells $\sigma \in \text{Del}(\Gamma P)$ having a vertex in $\text{dom}^\square(0, \Lambda)$ are entirely contained in $\text{dom}^\square(0, 3\Lambda)$.

Proof of Claim 3: Let p be a vertex of σ that is within $\text{dom}^\square(0, \Lambda)$. Shift the entire point set ΓP and its triangulation by the vector $-p$ so that p now coincides with 0. From Claim 2 it follows that the other vertices of the shifted cell are within $\text{dom}(0, 2\Lambda)$. Adding the vector $+p$ to these shifted vertices we get back the vertices of σ , but because $p \in \text{dom}^\square(0, \Lambda)$ we also know now that these vertices are within $\text{dom}^\square(0, 3\Lambda)$, as $\text{dom}^\square(0, 3\Lambda)$ is the Minkowski sum of $\text{dom}^\square(0, \Lambda)$ and $\text{dom}(0, 2\Lambda)$.

Proof of Proposition: Every d -cell of $\text{Del}(\Gamma P)$ that has a vertex in $\text{dom}^\square(0, \Lambda)$ has all its vertices in P_3 . Furthermore, because it fulfils the empty-sphere property in ΓP , then it also fulfils this property in $P_3 \subset \Gamma P$, and thus is present in $\text{Del}(P_3)$. As every

point from P has a representative in $\text{dom}^\square(0, \Lambda)$, also each d -cell from $\text{Del}(\Gamma P)$ has a representative in T_0 , proving the statement.

□

7.4.2 Extension to order- k Delaunay mosaics

In this section we will show that a result similar to Proposition 7.1 also holds for order- k Delaunay mosaics, giving us a recipe to compute them for periodic point sets.

In the following we will assume barycenters for the geometric embeddings of vertices of the order- k Delaunay mosaic. For a given periodic point set ΓP and an order k , let $r := r_{\Gamma P, k}$ be the radius of the largest sphere containing at most $k - 1$ points of ΓP in its interior. Let D be the Minkowski sum of $B_{2r}(0)$ and $\text{dom}^\square(0, \Lambda)$. Then the following result holds:

Proposition 7.8. *The d -cells of $\text{Del}_k(\Gamma P \cap D)$ that have at least one vertex in $\text{dom}^\square(0, \Lambda)$ are cells of $\text{Del}_k(\Gamma P)$. They contain at least one representative of each translational equivalence class of cells of $\text{Del}_k(\Gamma P)$.*

Proof. **Claim 1:** For a d -cell $\sigma \in \text{Del}_k(\Gamma P)$ that has a (geometric) vertex at the origin, all points of $P(\sigma)$ are inside $B_{2r}(0)$.

Proof of Claim 1: Note that $S_{\min}(\sigma)$ contains all points of $P(\sigma)$. As every geometric vertex of σ is a barycenter of subsets of $P(\sigma)$, $S_{\min}(\sigma)$ contains all geometric vertices of σ and thus in particular 0. Furthermore, $S_{\min}(\sigma)$ may not contain any points of ΓP other than $P_{\text{in}}(\sigma)$ in its interior. As $\#P_{\text{in}}(\sigma) \leq k - 1$, due to the definition of r the radius of $S_{\min}(\sigma)$ is at most r . Thus $S_{\min}(\sigma)$ is contained inside $B_{2r}(0)$, and therefore also all points from $P(\sigma)$ are.

Claim 2: For a d -cell $\sigma \in \text{Del}_k(\Gamma P)$ that has a (geometric) vertex p inside $\text{dom}^\square(0, \Lambda)$, all points of $P(\sigma)$ lie inside D .

Proof of Claim 2: This follows directly in the same vein as in Proposition 7.7, by translating the point set ΓP by the vector $-p$ and applying Claim 1.

Proof of Proposition: Let $\sigma \in \text{Del}_k(\Gamma P)$ be a d -cell that has a (geometric) vertex inside $\text{dom}^\square(0, \Lambda)$. The sphere $S_{\min}(\sigma)$ is the circumsphere of $P_{\text{on}}(\sigma)$ and contains $P_{\text{in}}(\sigma)$ in its interior, all of which are points within D . Therefore σ is also a d -cell in $\text{Del}_k(\Gamma P \cap D)$.

As every vertex class of $\text{Del}_k(\Gamma P)$ has a representative in $\text{dom}^{\mathbb{Q}}(0, \Lambda)$, also each class of d -cells of $\text{Del}_k(\Gamma P)$ has a representative in $\text{Del}_k(\Gamma P \cap D)$, proving the statement. \square

This result leaves us with the need for bounds on $r_{\Gamma P, k}$. Asymptotically, for a fixed point set ΓP , we would expect this radius to be $O(\sqrt{k})$, as for evenly distributed point sets the number of points in a ball of radius r is proportional to r^2 . However in practice we need exact bounds. These could be in terms of the lattice, regardless of the periodic point set, or e.g. in terms of the largest empty circumsphere of ΓP , which we can extract from the order-1 Delaunay triangulation that Algorithm 1 will have to compute along the way anyway. We leave these questions open as directions for further research.

Rather than applying Algorithm 1 directly to a sufficiently large subset of ΓP to get $\text{Del}_k(\Gamma P)$, a different approach could be to reconstruct $\text{Del}_j(\Gamma P)$ for all intermediate orders $j < k$: As each order- k Delaunay mosaic is computed as a weighted first-order Delaunay triangulation, we know from Proposition 7.7 that it is sufficient to triangulate $\text{dom}(0, 3\Lambda) \cap \Gamma P$. Thus by maintaining the combinatorial vertex set of each class of d -cells of $\text{Del}_j(\Gamma P)$, with offsets associated to each point of a vertex, we can likely reconstruct the vertex set of $\text{Del}_{j+1}(\Gamma P)$ in a similar way to Algorithm 1. The details would still have to be worked out.

7.4.3 Persistence of periodic point sets

For both the first order Delaunay triangulation and higher order Delaunay mosaics, the radius function can be computed as before. While a-priori there are infinitely many points imposing exclusion constraints on the sphere $S_{\min}(\sigma)$ for a cell σ , Lemma 4.5 tells us that only a small finite subset of these has to be checked. If we obtained $\text{Del}_k(\Gamma P)$ directly from a sufficiently large finite subset $\Gamma P \cap D$, like in Proposition 7.8, we can even use the optimization from Algorithm 3 to obtain radius values directly from the rhomboid tiling, as for cells σ that are periodic all points imposing constraints on $S_{\min}(\sigma)$ are in $\Gamma P \cap D$.

Due to the infinite nature, we cannot directly compute persistence in scale on the infinite order- k Delaunay filtration. We can however compute persistence on the torus using $\text{Del}(\Gamma P)/\Gamma$ and its analog $\text{Del}_k(\Gamma P)/\Gamma$. We would like to remark however that persistence on the torus has a few unexpected properties. It captures the inherent homology classes of the torus, giving multiple persistence pairs with infinite death value. Furthermore, when

computing persistence of ΓP on the torus $\mathbb{R}^d/m\Gamma$ instead of \mathbb{R}^d/Γ for some integer m , the multiplicities of persistence pairs behave in non-straightfoward ways; in particular, a persistence pair with multiplicity 1 on \mathbb{R}^d/Γ does not necessarily have multiplicity m^d on $\mathbb{R}^d/m\Gamma$, and in fact persistence pairs can emerge in the diagram for $\mathbb{R}^d/m\Gamma$ that are not present in the diagram for \mathbb{R}^d/Γ .

The situation for multi-cover persistence in depth on the torus is less clear, and open questions remain for future research: Does our algorithm for persistence in depth generalize to points on the torus? In particular, what does the rhomboid tiling for points on the torus look like, and does the rhomboid tiling of a periodic point set have a finite representation, like $\text{Del}_k(\Gamma P)$ has a finite representation on the quotient space \mathbb{R}^d/Γ ? Or alternatively, can we use the nerve construction from [24] also for point sets on the torus to get a simplicial bifiltration with the same persistence as the k -fold covers?

Bibliography

- [1] P.K. Agarwal, M. De Berg, J. Matousek, and O. Schwarzkopf. Constructing levels in arrangements and higher order Voronoi diagrams. *SIAM J. Comput.*, 27(3):654–667, 1998.
- [2] A. Aggarwal, L.J. Guibas, J. Saxe, and P.W. Shor. A linear-time algorithm for computing the Voronoi diagram of a convex polygon. *Discrete Comput. Geom.*, 4:591–604, 1989.
- [3] E. Agrell, T. Eriksson, A. Vardy, and K. Zeger. Closest point search in lattices. *IEEE Trans. Inform. Theory*, 48(8):2201–2214, 2002.
- [4] M. Ajtai. The shortest vector problem in L2 is NP-hard for randomized reductions. In *Proc. 30th Ann. ACM Symp. Theory Comput.*, pages 10–19, 1998.
- [5] T. Aste, M. Saadatfar, and T.J. Senden. Geometrical structure of disordered sphere packings. *Phys. Rev. E*, 71:061302, Jun 2005.
- [6] D. Attali, J.-D. Boissonnat, and A. Lieutier. Complexity of the Delaunay triangulation of points on surfaces the smooth case. In *Proc. 19th Ann. Symp. Comput. Geom.*, pages 201–210, 2003.
- [7] F. Aurenhammer. A new duality result concerning Voronoi diagrams. *Discrete Comput. Geom.*, 5(3):243–254, 1990.
- [8] F. Aurenhammer and O. Schwarzkopf. A simple on-line randomized incremental algorithm for computing higher order Voronoi diagrams. *Int. J. Comput. Geom. Appl.*, 2(04):363–381, 1992.
- [9] U. Bauer and H. Edelsbrunner. The Morse theory of Čech and Delaunay complexes. *Trans. Amer. Math. Soc.*, 369(369):3741–3762, 2017.

- [10] U. Bauer, M. Kerber, J. Reininghaus, and H. Wagner. PHAT—persistent homology algorithms toolbox. *J. Symb. Comput.*, 78:76–90, 2017.
- [11] J.D. Bernal. A geometrical approach to the structure of liquids. *Nature*, 183:141 – 147, 1959.
- [12] A. Bowyer. Computing Dirichlet tessellations. *The Computer J.*, 24(2):162–166, 1981.
- [13] G. Carlsson and V. de Silva. Zigzag persistence. *Found. Comput. Math.*, 10(4):367–405, 2010.
- [14] G. Carlsson, V. De Silva, and D. Morozov. Zigzag persistent homology and real-valued functions. In *Proc. 25th Ann. Symp. Comput. Geom.*, pages 247–256. ACM, 2009.
- [15] G. Carlsson and A. Zomorodian. The theory of multidimensional persistence. *Discrete Comput. Geom.*, 42(1):71–93, 2009.
- [16] M. Caroli. *Triangulating Point Sets in Orbit Spaces*. PhD thesis, Université Nice Sophia Antipolis, 2010.
- [17] M. Caroli, A. Pellé, M. Rouxel-Labbé, and M. Teillaud. 3D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.0.2 edition, 2020.
- [18] M. Caroli and M. Teillaud. 3D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 3.5 edition, 2009.
- [19] M. Caroli and M. Teillaud. Delaunay triangulations of closed Euclidean d -orbifolds. *Discrete Comput. Geom.*, 55(4):827–853, 2016.
- [20] F. Chazal, D. Cohen-Steiner, and Q. Mérigot. Geometric inference for probability measures. *Found. Comput. Math.*, 11(6):733–751, 2011.
- [21] K.L. Clarkson and P.W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4(1):387–421, 1989.
- [22] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. In *Proc. 21st Ann. Symp. Comput. Geom.*, SCG ’05, pages 263–271, New York, NY, USA, 2005. Association for Computing Machinery.

- [23] J.H. Conway and N.J.A. Sloane. Low-dimensional lattices VI: Voronoi reduction of three-dimensional lattices. *Proc. R. Soc. Lond. A*, pages 55–68, 1992.
- [24] R. Corbet, M. Kerber, and M. Lesnick. Computing the multicover bifiltration. Manuscript, 2019.
- [25] H. S. M. Coxeter. The classification of zonohedra by means of projective diagrams. *J. Math. Pures Appl.*, 41(9):137–156, 1962.
- [26] B. Delaunay. Sur la sphère vide. A la memoire de Georges Voronoï. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskikh i Estestvennykh Nauk*, 7:793–800, 1934.
- [27] V. Despré, J.-M. Schlenker, and M. Teillaud. Flipping geometric triangulations on hyperbolic surfaces. In *Proc. 36th Ann. Symp. Comput. Geom.*, 2020. To appear. Preliminary version: <https://hal.inria.fr/hal-02400219>.
- [28] O. Devillers and M. Teillaud. Perturbations for Delaunay and weighted Delaunay 3D triangulations. *Comput. Geom.: Theory & App.*, 44:160–168, 2011.
- [29] N. Dolbilin and D. Huson. Periodic Delone tilings. *Periodica Mathematica Hungarica*, 34:1-2:57–64, 1997.
- [30] R.A. Dwyer. Higher-dimensional Voronoi diagrams in linear expected time. *Discrete Comput. Geom.*, 6(3):343–367, 1991.
- [31] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag, Heidelberg, Germany, 1987.
- [32] H. Edelsbrunner. The union of balls and its dual shape. *Discret. Comput. Geom.*, 13:415–440, 1995.
- [33] H. Edelsbrunner and J.L. Harer. *Computational Topology. An Introduction*. American Mathematical Society, Providence, RI, 2010.
- [34] H. Edelsbrunner, D. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory*, 29(4):551–559, 1983.
- [35] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. In *Proc. 41st Ann. IEEE Symp. Found. Comput. Sci.*, pages 454–463, 2000.

- [36] H. Edelsbrunner and A. Nikitenko. Poisson–Delaunay mosaics of order k . *Discrete Comput. Geom.*, 62(4):865–878, 2019.
- [37] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete Comput. Geom.*, 1(1):25–44, 1986.
- [38] H. Edelsbrunner and N.R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica*, 15(3):223–241, 1996.
- [39] P. Engel. *Geometric crystallography: an axiomatic introduction to crystallography*. Springer, Dordrecht, 1986.
- [40] R. Forman. Morse theory for cell complexes. *Adv. Math.*, 134(1):90–145, 1998.
- [41] N. Francois, M. Saadatfar, R. Cruikshank, and A. Sheppard. Geometrical frustration in amorphous and partially crystallized packings of spheres. *Phys. Rev. Lett.*, 111(14):148001, 2013.
- [42] R. Freij. Equivariant discrete Morse theory. *Discrete Math.*, 309(12):3821–3829, 2009.
- [43] M.J. Golin and H.-S. Na. On the average complexity of 3d-Voronoi diagrams of random points on convex polytopes. *Comput. Geom.*, 25(3):197–231, 2003.
- [44] L.J. Guibas, D. Morozov, and Q. Mérigot. Witnessed k -distance. *Discrete Comput. Geom.*, 49(1):22–45, 2013.
- [45] T.C. Hales. A proof of the Kepler conjecture. *Ann. Math.*, 162(3):1065–1185, 2005.
- [46] M. Hanifpour, N. Francois, S.M. Vaez Allaei, T. Senden, and M. Saadatfar. Mechanical characterization of partially crystallized sphere packings. *Phys. Rev. Lett.*, 113(14):148001, 2014.
- [47] M. Hanifpour, N. Francois, V. Robins, A. Kingston, S.M. Vaez Allaei, and M. Saadatfar. Structural and mechanical features of the order-disorder transition in experimental hard-sphere packings. *Phys. Rev. E*, 91:062202, 2015.
- [48] A. Hatcher. *Algebraic Topology*. Cambridge Univ. Press, Cambridge, 2000.

- [49] S. Heitkam, W. Drenckhan, and J. Fröhlich. Packing spheres tightly: influence of mechanical stability on close-packed sphere structures. *Phys. Rev. Lett.*, 108(14):148302, 2012.
- [50] Y. Hiraoka, T. Nakamura, A. Hirata, E.G. Escolar, K. Matsue, and Y. Nishiura. Hierarchical structures of amorphous solids characterized by persistent homology. *Proc. Nat. Acad. Sci. USA*, 113(26):7035–7040, 2016.
- [51] C. Jamin, S. Pion, and M. Teillaud. 3D triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.0.2 edition, 2020.
- [52] M. Jerkins, M. Schröter, H.L. Swinney, T.J. Senden, M. Saadatfar, and T. Aste. Onset of mechanical stability in random packings of frictional spheres. *Phys. Rev. Lett.*, 101:018301, 2008.
- [53] B.A. Klumov, Y. Jin, and H.A. Makse. Structural properties of dense hard sphere packings. *J. Phys. Chem. B*, 118(36):10761–10766, 2014.
- [54] D. Kozlov. *Combinatorial Algebraic Topology*. Springer, 2008.
- [55] D. Krasnoshchekov and V. Polishchuk. Order- k α -hulls and α -shapes. *Inform. Process. Lett.*, 114(1-2):76–83, 2014.
- [56] A.S. Krishnapriyan, M. Haranczyk, and D. Morozov. Topological descriptors help predict guest adsorption in nanoporous materials. *arXiv:2001.05972*, 2020.
- [57] N. Kruithof. 2D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.4 edition, 2014.
- [58] D.-T. Lee. On k -nearest neighbor Voronoi diagrams in the plane. *IEEE Trans. Comput.*, 31(6):478–487, 1982.
- [59] Y. Lee, S.D. Barthel, P. Dlotko, S.M. Moosavi, K. Hess, and B. Smit. High-throughput screening approach for nanoporous materials genome using topological data analysis: application to zeolites. *J. Chem. Theory Comput.*, 14(8):4427–4437, 2018.
- [60] J. Leray. Sur la forme des espaces topologiques et sur les points fixes des représentations. *J. Math. Pures Appl. (9)*, 24:95–167, 1945.

- [61] P. McMullen. The maximum numbers of faces of a convex polytope. *Mathematika*, 17(2):179–184, 1970.
- [62] D. Micciancio and P. Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. *SIAM J. Comput.*, 42(3):1364–1391, 2013.
- [63] W. Mickel, S.C. Kapfer, G.E. Schröder-Turk, and K. Mecke. Shortcomings of the bond orientational order parameters for the analysis of disordered particulate matter. *J. Chem. Phys.*, 138(4):044501, 2013.
- [64] T.M. Middlemas, F.H. Stillinger, and S. Torquato. Hyperuniformity order metric of barlow packings. *Phys. Rev. E*, 99(2):022111, 2019.
- [65] D. Morozov. The Dionysus library. <https://mrzv.org/software/dionysus2/>, 2020.
- [66] K. Mulmuley. Output sensitive construction of levels and Voronoi diagrams in \mathbb{R}^d of order 1 to k . In *Proc. 22nd Ann. ACM Symp. Theory Comput.*, pages 322–330. ACM, 1990.
- [67] K. Mulmuley. On levels in arrangements and Voronoi diagrams. *Discrete Comput. Geom.*, 6(3):307–338, 1991.
- [68] T. Nakamura, Y. Hiraoka, A. Hirata, E.G. Escolar, K. Matsue, and Y. Nishiura. Description of medium-range order in amorphous structures by persistent homology. *arXiv:1501.03611*, 2015.
- [69] G. Osang. Higher order Delaunay mosaics in C++ and python. <https://github.com/geoo89/orderkdelaunay>, 2019.
- [70] G. Osang. Rhomboid tiling and order- k Delaunay mosaics in C++. <https://github.com/geoo89/rhomboidtiling>, 2020.
- [71] A.W. Overhauser. Crystal structure of lithium at 4.2 K. *Phys. Rev. Lett.*, 53(1):64, 1984.
- [72] F. Pacheco-Vázquez, G.A. Caballero-Robledo, and J.C. Ruiz-Suárez. Superheating in granular matter. *Phys. Rev. Lett.*, 102(17):170601, 2009.

- [73] F. Rietz, C. Radin, H.L. Swinney, and M. Schröter. Nucleation in sheared granular matter. *Phys. Rev. Lett.*, 120(5):055701, 2018.
- [74] V. Robins and K. Turner. Principal component analysis of persistent homology rank functions with case studies of spatial point patterns, sphere packing and colloids. *Physica D: Nonlinear Phenomena*, 334:99–117, 2016.
- [75] C. Rycroft. Voropp: A three-dimensional Voronoi cell library in C++. Technical report, Lawrence Berkeley Nat. Lab. (LBNL), Berkeley, CA (United States), 2009.
- [76] M. Saadatfar, H. Takeuchi, V. Robins, N. Francois, and Y. Hiraoka. Pore configuration landscape of granular crystallization. *Nature Comm.*, 8(1):1–11, 2017.
- [77] M.I. Shamos and D. Hoey. Closest-point problems. *Proc. 16th Ann. IEEE Symp. Found. Comput. Sci.*, pages 151–162, 1975.
- [78] D.R. Sheehy. A multi-cover nerve for geometric inference. In *Proc. Canadian Conf. Comput. Geom.*, 2012.
- [79] N. Sommer, M. Feder, and O. Shalvi. Finding the closest lattice point by iterative slicing. *SIAM J. Discrete Math.*, 23(2):715–731, 2009.
- [80] M. Tancer. Recognition of collapsible complexes is NP-complete. *Discrete Comput. Geom.*, 55(1):21–38, 2016.
- [81] The CGAL Project. *CGAL User and Reference Manual*. CGAL Editorial Board, 5.0 edition, 2019.
- [82] The GUDHI Project. *GUDHI User and Reference Manual*. GUDHI Editorial Board, 3.2.0 edition, 2020.
- [83] The RIVET Developers. Rivet. <http://rivet.online>, 2018.
- [84] L. Fejes Tóth. *Lagerungen in der Ebene, auf der Kugel und im Raum*, volume 65. Springer-Verlag, 1954.
- [85] A. van der Net, W. Drenckhan, D. Weaire, and S. Hutzler. The crystal structure of bubbles in the wet foam limit. *Soft Matter*, 2(2):129–134, 2006.

- [86] P. van Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. Technical report, Mathematische Instituut, Uni. Amsterdam Report, 04 1981.
- [87] O. Vipond. Multiparameter persistence landscapes. *J. Machine Learning Res.*, 21(61):1–38, 2020.
- [88] G. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. premier mémoire. sur quelques propriétés des formes quadratiques positives parfaites. *J. f. reine u. angew. Math.*, 1908(133):97–102, 1908.
- [89] D.F. Watson. Computing the n -dimensional Delaunay tessellation with application to Voronoi polytopes. *The Computer J.*, 24(2):167–172, 1981.
- [90] A.F. Wells. Structural inorganic chemistry. Oxford University Press, 1984.
- [91] E. Welzl. Smallest enclosing disks (balls and ellipsoids). In *Results and New Trends in Comp. Sci.*, pages 359–370. Springer, 1991.
- [92] T.F. Willems, C.H. Rycroft, M. Kazi, J.C. Meza, and M. Haranczyk. Algorithms and tools for high-throughput geometry-based analysis of crystalline porous materials. *Microporous and Mesoporous Materials*, 149(1):134–141, 2012.
- [93] M. Yvinec. 2D triangulation. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.0.2 edition, 2020.
- [94] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, 33(2):249–274, 2005.