Persistence and Morse Theory for Discrete Geometric Structures

by

Sebastiano Cultrera di Montesano

March, 2024

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

Committee in charge:

Francesco Locatello, Chair
Herbert Edelsbrunner
Kathryn Hess Bellwald
Yasuaki Hiraoka

The thesis of Sebastiano Cultrera di Montesano, titled *Persistence and Morse Theory for Discrete Geometric Structures*, is approved by:

Supervisor: Herbert Edelsbrunner, ISTA, Klosterneuburg, Austria

Signature: _____

Committee Member: Kathryn Hess Bellwald, EPFL, Lausanne, Switzerland

Signature: _____

Committee Member: Yasuaki Hiraoka, Kyoto University Institute for Advanced Study, Kyoto, Japan

Signature: _____

Defense Chair: Francesco Locatello, ISTA, Klosterneuburg, Austria

Signature: _____

Signed page is on file

ISTA 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: _____

Sebastiano Cultrera di Montesano
March, 2024

Signed page is on file

Abstract

Point sets, geometric networks, and arrangements of hyperplanes are fundamental objects in discrete geometry that have captivated mathematicians for centuries, if not millennia. This thesis seeks to cast new light on these structures by illustrating specific instances where a topological perspective, specifically through discrete Morse theory and persistent homology, provides valuable insights.

At first glance, the topology of these geometric objects might seem uneventful: point sets essentially lack of topology, arrangements of hyperplanes are a decomposition of $\mathbb{R}^d$, which is a contractible space, and the topology of a network primarily involves the enumeration of connected components and cycles within the network. However, beneath this apparent simplicity, there lies an array of intriguing structures, a small subset of which will be uncovered in this thesis.

Focused on three case studies, each addressing one of the mentioned objects, this work will showcase connections that intertwine topology with diverse fields such as combinatorial geometry, algorithms and data structures, and emerging applications like spatial biology.

Acknowledgements

I would like to express my deepest gratitude to my advisor, Herbert Edelsbrunner, for his invaluable guidance and teachings throughout my academic journey. His approach of allowing me to explore, make mistakes, and learn from them has been instrumental in my growth. Even more remarkable is his skill in balancing this guidance while keeping a light-hearted perspective. I also appreciate his efforts in assembling an exceptional team of individuals whose impact on my research during my PhD has been profound.

Special thanks to my coauthors Morteza, Ranita, and Ondra. Collaborating closely with you during various stages of my PhD has been not only immensely enjoyable but also incredibly educational. In fact, the research undertaken over the past 5 years, which this thesis builds on, should really be considered a group effort rather than an individual one. This thesis endeavors to weave our findings into a cohesive narrative.

Our collaboration began shortly after I joined the group, which had expanded significantly following Herbert's receipt of two major grants in 2018. My initial project involved delving into a discrete Morse theory book [101] alongside Ranita and Morteza, who were also new to the group. This joint study, combining our diverse backgrounds in combinatorics (Morteza), computer science (Ranita), and topology (myself), proved to be a fertile ground for raising new research questions. Several of these questions, tackled during the isolation of COVID times, laid the foundations for our ongoing collaborative work. A few years into the game, Ondra joined the team, initially drawn by our applied project focusing on developing a computational pipeline for the analysis of cellular arrangements in histology slides. His instrumental role quickly became evident, leading us to the development of the topological theory of chromatic alpha complexes, a central theme of this thesis.

My heartfelt thanks go to my family for their support during challenging times and for their genuine interest in my work. I am even more grateful to Marti for her constant support in all aspects of life, for believing in my abilities and for helping me learn how to present my work more effectively.

I extend my gratitude to the team at Owkin, especially Elodie and Paul, for their patience and trust. A special acknowledgment to Thomas for fostering a vibrant culture within the company and allowing me to be part of it. This experience significantly shaped my life, affirming my passion for life sciences and the biotech industry, which is the path I intend to pursue next.

I am thankful to Caroline Uhler for the invitation to spend time at the Schmidt Center at the Broad Institute of MIT and Harvard. My experience in Kyoto, thanks to Yasu Hiraoka, was unforgettable and contributed significantly to the completion of this thesis. I am also indebted to Kathryn Hess for her unwavering support throughout my PhD and for being an example of inspiring leader in academia.

Lastly, I express my fondness for the fussball table at ISTA, a constant presence in my life there. Its absence will certainly be felt.

Funding. The research undertaken over past 5 years would have not been possible without the research grants that were awarded to my supervisor. In particular, the various projects that this thesis builds on have received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant no. 788183, from the Wittgenstein Prize, Austrian Science Fund (FWF), grant no. Z 342-N31, and from the DFG Collaborative Research Center TRR 109, 'Discretization in Geometry and Dynamics', Austrian Science Fund (FWF), grant no. I 02979-N35.

About the Author

Sebastiano Cultrera di Montesano completed a BSc in Mathematics at the University of Glasgow and an MSc in Pure Mathematics at the University of Cambridge. He joined ISTA in September 2018. His research interests lie in applied and computational topology and his work has been published in leading journals and conferences of the field. In 2021, Sebastiano received the best student presentation award at the Symposium of Computational Geometry.

During his doctoral studies, Sebastiano conducted research at Owkin, a French biotech scale-up focused on personalizing patient treatment through computational methods. This experience sparked Sebastiano's interest in applying topological methods to biology, leading to an invitation from the Broad Institute of MIT and Harvard, where he further explored this research direction.

Throughout his PhD, Sebastiano engaged with diverse audiences beyond the mathematical community, particularly biologists and medical doctors. He presented his research at several internationally renowned biomedical institutions, including the Broad Institute, the European Institute of Oncology in Milan, the Policlinico Gemelli in Rome, the Center for Molecular Medicine in Vienna, and the biotech company Owkin.

List of Collaborators and Publications

Publications which are part of this thesis:

1. R. Biswas, S. Cultrera di Montesano, O. Draganov, H. Edelsbrunner, and M. Saghafian. On the size of chromatic Delaunay mosaics. arXiv:2212.03121 (2022).
2. R. Biswas, S. Cultrera di Montesano, H. Edelsbrunner and M. Saghafian. Depth in arrangements: Dehn–Sommerville–Euler relations with applications. To appear in *J Appl. and Comput. Topology* (2023).
3. R. Biswas, S. Cultrera di Montesano, H. Edelsbrunner and M. Saghafian. Geometric characterization of the persistence of 1D maps. *J Appl. and Comput. Topology* (2023), 1–19.
4. S. Cultrera di Montesano, O. Draganov, H. Edelsbrunner, and M. Saghafian. Chromatic alpha complexes. arXiv:2212.03128 (2024).
5. S. Cultrera di Montesano, H. Edelsbrunner, M. Henzinger and L. Ost. Dynamically Maintaining the Persistent Homology of Time Series. *In* "Proc. 35th Ann. ACM-SIAM Sympos. Discrete Alg. 2024", 243–295.

Chapter 3 is based (2). Chapter 4 is based on (3) and (5). Chapter 5 is based on (1) and (4).

Table of Contents

List of Figures

List of Tables

CHAPTER 1

Introduction

The shapes that exist in the natural world are fundamental building blocks dictating the structure, function, and behaviour of living organisms. Consider the intricate folds and convolutions of the human brain, which house a multitude of cognitive processes and define the uniqueness of human intelligence. In the realm of biology, understanding the shapes of proteins and their complex interactions is crucial to unlocking the mysteries of cellular functions and discovering potential targets for therapeutic intervention. Furthermore, the striking patterns of leaves in plants not only contribute to their visual allure but also reveal essential adaptations to environmental conditions, enabling them to thrive in various habitats. From the formation of organ structures to the evolution of species, the profound influence of shapes permeates the very fabric of life, driving scientific inquiries across disciplines.

Exploring shapes has long captivated mathematicians, sparking the development of fresh mathematical concepts and principles. In ancient times, the study of geometry emerged as humanity's first attempt to formalize and understand the shapes that surrounded them. The Greeks, notably Euclid, provided a rigorous foundation for geometric principles, introducing axiomatic reasoning and the notion of proofs [67]. However, as mathematical thinking advanced, the focus shifted from measuring shapes to understanding their intrinsic properties. A notable such example was observed by Leonhard Euler, who noticed a relation between the number of vertices, the number of edges and the number of faces in a bounded convex polyhedron. Indeed, he showed that their alternating sum always adds up to 2. This is often considered as the starting point of the mathematical field of *topology*: it is a global statement that does not depend on the precise geometry of the shape. It took more than a century to find a satisfying framework that includes Euler's original observation as a special case. In fact, Euler's formula was generalized to arbitrary d-dimensional polytopes by Schläfli (1852) [100] but the first correct proof was given 50 years later by Poincaré [90].

The first result that we will prove in this thesis is reminiscent, at least in spirit (and certainly not in impact), to Euler's formula from the 18th century. Specifically, we show that there is a system of independent equations that control the combinatorial structure of an *arrangement of hyperplanes* in Euclidean space. The latter is a decomposition of $\mathbb{R}^d$ into connected open cells of dimensions $0, 1, \cdots, d$. Just like in Euler's case such equations are given in terms of alternating sums. Our result fits into a well established paradigm: using topological methods to study questions in discrete geometry; see survey articles [13, 111] and books [81]. The latter reference illustrates the most striking example of such paradigm by showcasing the applications of the Borsuk-Ulam theorem – a celebrated topological result – in combinatorial geometry. The

theorem states that for every continuous function from a d-sphere to d-dimensional Euclidean space, there exist a pair of antipodal points (opposite points) on the sphere that map to the same point in Euclidean space. Perhaps surprisingly at first, the theorem finds use in problems related to fair division and graph coloring, providing elegant solutions that bridge the abstract realm of topology with challenging combinatorial problems.

Besides being interesting in their own right, arrangements of hyperplanes have served as a unifying structure for many problems in discrete and computational geometry [106, Chapter 28]. In the preface to his book *Algorithms in Combinatorial Geometry* [42], Edelsbrunner writes: "[...] arrangements of hyperplanes are at the very heart of computational geometry and this is my belief now more than ever". One of the main reasons for Edelsbrunner's claim, which is of particular importance for this thesis, is the *duality* between finite sets of points and finite sets of hyperplanes. In fact, in the same book [42, Chapter 12] the author describes a collection of problems stated for point configurations and solved by operating on their corresponding dual arrangements.

For the specific problem treated in Chapter 3, this duality allows us to reformulate questions about splitting finite point sets as questions about a discrete function defined on a corresponding arrangement. To avoid the case analysis needed to distinguish bounded and unbounded cells, we work with arrangements of great-spheres on $\mathbb{S}^d$ rather than of hyperplanes in $\mathbb{R}^d$. Assuming non-vertical great-spheres (which do not pass through the north-pole and the south-pole) the *depth function* maps every cell of the arrangement to the number of great-spheres that separate the cell from the north-pole.

Aspects of this function have been studied in the past, such as the maximum number of chambers (top-dimensional cells) at a given depth, which relates to counting *k-sets* in a set of n points, which are subsets of k points that can be separated from the remaining $n-k$ points by a straight line; see e.g. [53]. Giving tight bounds on the number of k-sets is still open, with substantial gaps between the current best upper and lower bounds in all dimensions larger than or equal to 2. We propose to focus on the topological aspects of the depth function, in particular the occurrence of critical cells of different types. In the top dimension, we have a chamber containing the north-pole (a minimum at depth 0), a chamber containing the south-pole (a maximum at depth n), and otherwise only non-critical chambers connecting the minimum to the maximum. There is nothing much topological to learn from such a *bi-polar* depth function, but its restrictions to common intersections of great-spheres display a richer topology, which can be studied with methods from discrete Morse theory [55] and persistent homology [43]. The core result presented in Chapter 3 is a system of Dehn–Sommerville type relations for level sets of the depth function. We will also show how this system can be used to rederive known cell-counting formulas for order-k Voronoi tessellations in $\mathbb{R}^2$ [73] and $\mathbb{R}^3$ [10]. These tessellations are partitions $\mathbb{R}^d$ on the basis of the first k closest sites (without distinguishing order among them). Can the persistent homology on the restrictions of the depth function to one-dimensional subarrangements be used to improve the current best upper bound for the number of k-sets in $\mathbb{R}^2$?

Motivated by the question above, the Chapter 4 of the thesis delves into the study of continuous functions defined on compact one-dimensional spaces such as the the unit circle, the unit interval, and more general geometric networks. Such maps are ubiquitous and arise in developmental biology (e.g. rythmic gene expression [34]), physiology (e.g. heart-rate [63]), and numerous other fields. We focus on the persistence homology of these functions, with the goal of using the persistent diagrams as practical bookkeeping devices for relevant combinatorial quantities. It will become apparent in Chapter 4 that the persistent homology of functions

on one dimensional spaces exhibits strikingly combinatorial characteristics. We will show how this distinctive feature allows for a geometric characterization of persistent pairs, enabling the derivation of elementary proofs for theorems related to the symmetry of persistence diagrams. In particular, our analysis identifies branching points and endpoints within networks as the sole sources of asymmetry.

Another implication of the geometric characterization of persistent pairs that we will derive in Chapter 4 is a relation between the variation and the total persistence of maps on networks. The variation of a real-valued map quantifies the total amount of local change in the map. According to the Koksma–Hlawka inequality, the error of a numerical integration is bounded from above by the variation of the map times the discrepancy of the points at which the map is evaluated [68, 70]. For a 1-dimensional, compact, and piecewise differentiable map, the variation is the integral of the absolute derivative. It is also the total persistence of the map, as we will prove in Chapter 4. The variation is thus a numerical summary of the more detailed information expressed in the persistence diagram. Not unlike the Fourier transform, this diagram decomposes the variation into different scales.

Furthermore, we will show how the structural results proven in this chapter have algorithmic implications. Specifically, they serve as the foundations for fast algorithms to maintain a collection of sorted lists together with their persistence diagrams. In fact, our structural results have culminated in the development of a new data structure designed to efficiently update and manage persistence information in the specific setting of functions defined on intervals (e.g. time series). In essence, the combinatorial nature of persistence in the 1-dimensional setting not only enhances our understanding of persistence itself but also unlocks new avenues for practical applications across disciplines where the analysis of time series data holds substantial significance.

The relationship between topology and combinatorics constitutes a dynamic interplay rather than a one-way street, where both disciplines mutually inform and enrich one another. While topological methods serve as powerful tools for solving intricate combinatorial problems, the converse is also true—topology can be effectively approached through a combinatorial lens. It's interesting to observe that the combinatorial approach to topology actually emerged before the use of topological methods in combinatorial problems. Historically, *combinatorial topology*, a predecessor of algebraic topology, characterized an era where topological invariants of spaces, like the Betti numbers, were perceived as derivable from combinatorial decompositions—such as decompositions into simplicial complexes.

In contemporary domains such as Topological Data Analysis [43], this reciprocal interdependence thrives, highlighting the indispensable reliance on discrete combinatorial structures for computing topological descriptors, notably persistent homology. The ability to govern the scale of these combinatorial structures stands pivotal, enabling efficient computational methodologies. Chapter 5 of this thesis will illuminate this interdependence by introducing a novel and efficient combinatorial structure tailored to investigate spatial interactions within point set data.

The investigation is driven by recent impactful technological advances in biology. Up until recently, extracting genetic information from a cell, including gene expression, necessitated tissue dissociation, leading to the loss of individual cell spatial data. However, the emergence of cutting-edge technologies, such as multiplex imaging and spatial transcriptomics, now allows direct observation of cellular phenotypes while retaining their original spatial context within the tissue microenvironment. These breakthroughs present new avenues for unraveling molecular organization, deciphering intra-tumor heterogeneity, and identifying biomarkers indicative of

responses to existing treatments. They not only open new directions in biological research but also create a pressing need for methodologies that effectively characterize the intricate geometry underlying cellular arrangements.

The theory of alpha complexes introduced by Edelsbrunner and collaborators [44, 48] is a compelling starting point to achieve this goal since it was originally designed to quantitatively describe the spatial configuration of a point set. The idea behind the mathematical construction is to grow discs around each data point with an increasingly larger radius and track the evolution of the homological features (i.e. the connected components and the cycles) as the radius grows. Persistent homology allows not only to identify radii at which homological features appear and disappear, but also pair these events to quantify how long each feature persists. The persistence of each homological feature is compactly stored in a "bookkeeping device" called the persistent diagram [43]. Motivated by the complex spatial interactions between multiple cell types in the tumor microenvironment, we extend the theory of alpha complexes to the setting where the input points are labeled according to the cellular phenotype of the cell they represent.

The necessary step needed to describe not only the shape of each point set separately, but also how such shapes relate to each other, is to study the inclusion map between the union of disks of one color (say red) into the union of disks of multiple colors (say red and blue) as their radii grow. For example, if a cycle is present in the union of red disks at a certain radius, it might or might not also be present in the union of the red and blue disks. If it is, such cycle will be in the image of the induced map in homology, while if it is not, it will be in the kernel. This observation led us to engineer a collection of homological quantifiers that describe what happens to the persistent homology of one set when we also consider points belonging to another set. The persistent homology of chromatic sets was introduced in [25], where we proved various relations between the homological quantifiers involved.

In [9], we showed that the continuous structures induced by the union of disks can be efficiently computed using a new carefully designed combinatorial representation, which we named the *chromatic alpha complex*. This structure uses an extra dimension for each color beyond the first to capture the interaction between colors. Counter-balancing the increase in dimension, we showed that the combinatorial size of the complex is moderate for small number of colors [9]. For example, we give linear bounds on the expected size for randomly colored points in two dimensions. The compelling mathematical properties of the algorithm, such as its stability with respect to small perturbations of the point set and its computational efficiency, encourage the use of chromatic alpha complexes in biomedical settings.

Collectively, these projects exemplify the versatility of topological methods and their profound connection with questions in combinatorics and geometry, as well as algorithms and data structures. The final chapter will suggest that topological thinking can transcend the traditional confines of mathematics and computer science, and serve as a powerful tool for the computational scientist.

Outline. Chapter 2 will provide the necessary background in discrete and combinatorial geometry that the rest of the thesis builds on. This includes introducing Voronoi tessellations, Delaunay mosaics and their relation with arrangements of hyperplanes. The remainder of the thesis will illustrate three case studies, each focused on a specific discrete geometric structure, where we gained insights by studying these objects from a topological perspective. In Chapter 3, we will delve into the study of arrangements of hyperplanes and derive formulas connecting the number of cells across various dimensions. In Chapter 4, we will explore the persistent

homology of functions on geometric networks, and provide a geometric characterisation of the persistent pairs in this setting. Chapter 5 will focus on the generalization of the alpha complex construction to scenarios where points have a label associated with them. We will describe the framework, illustrate its properties and prove combinatorial bounds for its size.

Disclaimer. The order in which the chapters are presented is mostly chronological, illustrating the author's personal journey through his PhD studies. The narrative style not only captures the evolution of his understanding but also reflects his changing tastes, marking a shift from theoretical to applied focus towards the thesis's conclusion. An effort has been made to weave the topics into a cohesive story, yet deviations in the narrative are inevitable. These deviations are not mere digressions but represent unexpected discoveries — moments of serendipity that are intrinsic to the nature of mathematical exploration. The thesis builds heavily on five research papers [9, 25, 11, 12, 27] that the author worked on during his PhD. If you have already read them carefully, you should probably stop here as not much else is in this thesis besides speculations at the end of each chapter.

CHAPTER 2

Discrete and Combinatorial Geometry

In this chapter, we examine foundational concepts in discrete and combinatorial geometry, essential for the subsequent sections of this thesis. For a thorough and detailed exploration of this field, readers are encouraged to consult [42].

We will introduce convex polytopes, Voronoi tessellations, and Delaunay mosaics, while highlighting how arrangements of hyperplanes serve as a cohesive framework, linking these concepts together. This chapter examines the structure and size of these objects, which is an essential prerequisite for determining what is the most efficient way to compute them, as well as determining how long it takes to do so.

2.1 Convex Polytopes

The study of convex polytopes in Euclidean space of two and three dimensions is one of the oldest branches of mathematics. In fact, they were already studied by Plato in the fourth century BC, who hypothesized in one of his dialogues, the *Timaeus*, that the classical elements were made of five regular solids: the tetrahedron, cube, octahedron, dodecahedron, and icosahedron. In a survey article on convex polytopes [65], Grünbaum and Shephard remarked that there were three developments which foreshadowed the modern theory of convex polytopes.

1. The publication of Euclid's Elements [67] and the five Platonic solids. In modern terms, these are the regular 3-polytopes, where regular means that the faces are congruent (identical in shape and size) regular polygons (all angles congruent and all edges congruent), and the same number of faces meet at each vertex.
2. Euler's Theorem, already mentioned in the introduction, which states that $v - e + f = 2$ holds for any 3-dimensional polytope, where v, e and f denote the number of vertices, edges and facets, respectively.
3. The discovery of polytopes in dimensions greater or equal to four by Schläfli [100].

Perhaps surprisingly, many of the interesting properties of polytopes have been discovered comparatively recently. After introducing some terminology, we will recall some of these properties and we recommend [110] for a comprehensive introduction to convex polytopes.

Key definitions. A set X is *convex* if it has the property that for any pair of points $x, y \in \mathbb{R}^d$, the line segment connecting x to y lies entirely in X. Note that the intersection of convex sets in also convex. 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 finitely many half-spaces, and is therefore convex. A (convex) *polytope* is a bounded polyhedron. Every convex polytope can alternatively be constructed by taking the *convex hull* of its vertices [110, Sect. 1.1]. This is the intersection of all convex sets that contain all vertices of the polytope. The *dimension* of a polytope is the dimension of the smallest Euclidean space which contains it. See Figure 2.1 for three examples of convex polytopes in $\mathbb{R}^3$.

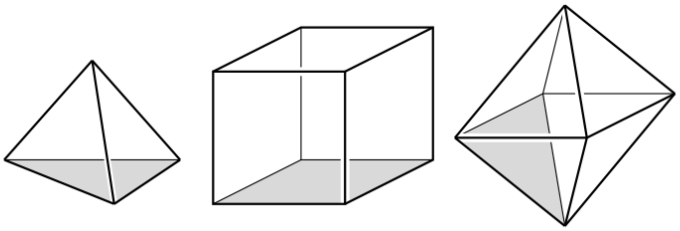

Figure 2.1: Examples of 3-dimensional polytopes: a tetrahedron, a cube and an octahedron.

A hyperplane h supports a closed bounded convex set X if $h \cap X \neq 0$, and X lies in one of the two closed half-spaces bounded by h. If h supports X, then $h \cap X$ is called a *face* of X. In the case of polytopes, each face is itself a polytope. Every point in the boundary of a polytope P lies on some supporting hyperplane of P, and so belongs to some face of P. To summarize, for a d-polytope P the following properties hold.

1. The faces of P are polytopes.
2. P possesses faces of every dimension $0, 1, \cdots, d-1$. A 0-face is called a *vertex*, a 1-face is called an *edge*, and a $(d-1)$-face is called a facet of P. If σ is a face of τ, we say that τ is a *coface* of σ.
3. Every face of a face of P is also a face of P.

The convex hull of $d+1$ affinely independent points is a d-polytope known as a *d-simplex* (a 3-simplex is often called a tetrahedron). If all the proper faces of a d-polytope P are simplices, then P is called a simplicial polytope.

Two polytopes P and P^* in $\mathbb{R}^d$ are *dual* if there is a bijection between their faces such that for two dual faces 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. A dual of a simplicial polytope is called a *simple polytope*. Note that the dual polytope of a cube is an octahedron, while a tetrahedron is dual to itself.

The *face lattice* of a polytope P is the set of all faces of P, partially ordered by inclusion. Two polytopes have isomorphic face lattices if there exists a one-to-one inclusion-preserving mapping between them. The combinatorial theory of polytopes may be regarded as a study of the face-lattices; it is concerned with combinatorial equivalence classes of polytopes rather than with polytopes themselves.

Face Counting. The f-*vector* of a convex polytope is given by $(f_0, \cdots, f_{d-1})$, where f_i enumerates the number of i-dimensional faces in the d-dimensional polytope. Since the problem of characterising the face-lattices of polytopes has proved too difficult, it is natural to attempt the apparently simpler one of characterising the f-vectors of polytopes. The theorem below takes a step in that direction.

Theorem 2.1.1 (Euler–Poincaré Relation)**.** *The f-vector of a polytope satisfies the **Euler–Poincaré relation**:*

$$f_0 - f_1 + f_2 - \cdots + (-1)^{d-1} f_{d-1} = 1 - (-1)^d$$

equivalently,

$$\sum_{i=-1}^{d} (-1)^i f_i = 0,$$

where f_{-1} denotes the number of empty faces ($= 1$) and $f_d = 1$ counts the entire polytope.

Proof. See [42, Theorem 6.8]. □

In other words, the f-vectors of all polytopes P lie on a certain hyperplane in Euclidian space.

Dehn-Sommerville Equations. For simplicial (and simple) polytopes it turns out that other remarkable equations – besides the Euler-Poincaré formula – hold among the number of i-faces. These equations were discovered by Dehn for $d = 4, 5$ in 1905 [32] and by Sommerville in the general case in 1927 [103]. For $d = 3$, every edge belongs to two facets and every facet has three edges. It follows that $2f_1 = 3f_2$. Together with Euler's formula $f_0 - f_1 + f_2 = 2$, we see that $f_1 = 3f_0 - 6$ and $f_2 = 2f_0 - 4$, namely, that the number of vertices of a simplicial 3-polytope determines its number of edges and faces, these being linear functions of the number of vertices. More generally, the following relations for simple polytopes hold.

Theorem 2.1.2 (Dehn-Sommerville Equations)**.** *Let $P \subseteq \mathbb{R}^d$ be a simple polytope. Then*

$$\sum_{i=k}^{d} (-1)^i \binom{i}{k} f_i = \sum_{i=d-k}^{d} (-1)^{d-i} \binom{i}{d-k} f_i,$$

for all $0 \leq k \leq d$.

Proof. See [42, Theorem 6.10]. □

Note that when $k = 0$, this becomes Euler's formula. We refer to [64, Section 9.2] for a more in depth introduction to the Dehn–Sommerville relations for convex polytopes.

The Upper Bound Theorem and Cyclic Polytopes. For a polytope of dimension d with n vertices, what is the maximal number of i-faces it can possess? This question is not only important from a theoretical point of view but also from a computational point of view because of its implications for algorithms in combinatorial optimization and in computational geometry. The answer to the above problem is that there is a class of polytopes called *cyclic polytopes* such that the cyclic d-polytope, $C_d(n)$, has the maximum number of i-faces among all d-polytopes with n vertices. This result stated by Motzkin in 1957 became known as the

upper bound conjecture until it was proved by McMullen in 1970 [82] . It is now known as the *Upper Bound Theorem*.

For fixed positive integers d and n the *cyclic polytope* $C_d(n)$ is the convex hull of n distinct points on the moment curve $(t, t^2, \cdots, t^d)$. Figure 2.2, which is borrowed from [106, Chapter 15], shows (a projection of) a 3-dimensional cyclic polytope.

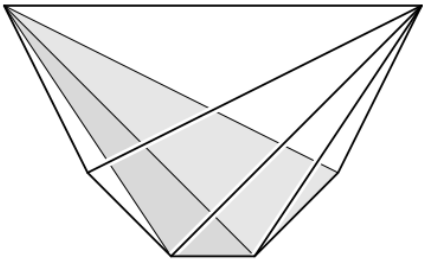

Figure 2.2: A 3-dimensional cyclic polytope with 6 vertices

2.2 Voronoi Tessellations, Delaunay Mosaics and Alpha Complexes

Given some number of points in the plane, their *Voronoi tessellation* divides the plane according to the nearest-neighbor rule: each point is associated with the region of the plane closest to it; see Figure 2.3. The concept of Voronoi tessellation has independently emerged, and

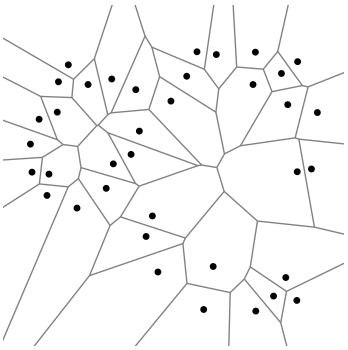

Figure 2.3: A Voronoi tessellation in the plane.

proven useful, in various fields of science. Different names particular to the respective field have been used, such as *Wigner-Seitz zones* in chemistry and physics, *domains of action in crystallography*, and *Thiessen polygons* in meteorology and geography. The mathematicians Dirichlet [38] and Voronoi [107] were the first to formally introduce this concept, even though Descartes had drawings of Voronoi tessellations about a century earlier in his book on the principles of philosophy [31]. His illustrations show a decomposition of space into convex regions, each consisting of matter revolving around one of the fixed stars [6]. Dirichlet and Voronoi used it for the study of quadratic forms; the sites are integer lattice points, and influence of each site is measured by the Euclidean distance. See [42, 6, 56] for standard reference in the subject.

Key definitions. Letting $A \subseteq \mathbb{R}^d$ be a finite set of points, the *Voronoi domain* of $a \in A$, denoted $\mathrm{dom}(a, A)$, is the set of points $x \in \mathbb{R}^d$ that satisfy $\|x - a\| \leq \|x - b\|$ for all $b \in A$. Observe that $\mathrm{dom}(a, A)$ is the intersection of finitely many closed half-spaces and therefore a closed convex polyhedron. The *Voronoi tessellation* of A, denoted $\mathrm{Vor}(A)$, is the collection of Voronoi domains defined by the points in A. These domains cover $\mathbb{R}^d$ while their interiors are pairwise disjoint. Nevertheless, a collection of these polyhedra may overlap in a shared face, which we refer to as a *Voronoi cell*. For a generic set, A, the dimension of a Voronoi cell is determined by the number of Voronoi domains that contain it.

Definition 2.2.1 (Generic Point Set)**.** *We call a point set, $A \subseteq \mathbb{R}^d$,* generic *if, for every $1 \leq p \leq d$, no $p + 2$ points in A lie on a common $(p - 1)$-dimensional sphere in $\mathbb{R}^d$.*

Then, indeed, the common intersection of any $p + 1$ Voronoi domains is either empty or a convex polyhedron of dimension $d - p$. Note that our notion of genericity allows for more than $p + 1$ points on a p-dimensional affine subspace.

Writing $n = \#A$, it is clear that $\mathrm{Vor}(A)$ has precisely n d-cells. For $d = 2$, this implies that there are at most $3n$ edges and at most $2n$ vertices. More generally for n points in $\mathbb{R}^d$, the Voronoi tessellation has $O(n^{\lceil d/2 \rceil})$ cells. While this bound is tight, the number of cells depends on the relative position of the points and is much smaller for many sets. For example, the Voronoi tessellation of n points chosen uniformly at random inside the unit cube in a constant-dimensional Euclidean space has only $O(n)$ cells in expectation; see e.g. [41].

The *Delaunay mosaic* of $A \subseteq \mathbb{R}^d$, denoted $\mathrm{Del}(A)$, is the dual of the Voronoi tessellation of A. To be specific, consider a p-cell of $\mathrm{Vor}(A)$, and observe that it is the common intersection of $m \geq d - p + 1$ Voronoi domains. Assuming this collection of domains is maximal, and writing $a_1, a_2, \ldots, a_m$ for the points in A that generate them, we call the convex hull of the a_i the *dual Delaunay cell* of the Voronoi p-cell. Its dimension is $q = d - p$. The Delaunay mosaic of A is the collection of Delaunay cells dual to cells of $\mathrm{Vor}(A)$.

We note that $\mathrm{Del}(A)$ is a polyhedral complex; that is: it consists of closed polyhedral cells such that the boundary of each cell is the union of lower-dimensional cells in the complex. Similarly, the collection of cells of $\mathrm{Vor}(A)$ is a polyhedral complex, but note that $\mathrm{Vor}(A)$ is, by definition, only the collection of Voronoi domains, which is not a complex.

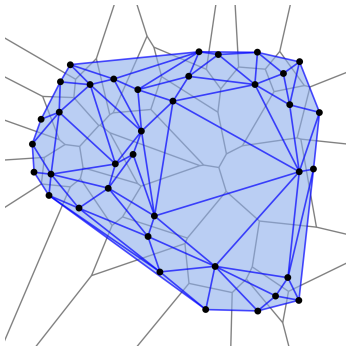

Figure 2.4: A Delaunay mosaic in the plane.

Call a $(d - 1)$-dimensional sphere *empty* of points in A if no point in A is enclosed by the sphere. The points may lie on the sphere or outside the sphere, but they are not allowed to lie inside the sphere. It is not difficult to see that the convex hull of m points in A is a cell in

$\mathrm{Del}(A)$ iff these m points lie on an empty $(d-1)$-sphere, while all other points in A lie strictly outside this sphere. Indeed, the center of such an empty sphere is a point in the interior of the dual Voronoi cell, and the Voronoi domains generated by the m points all share the cell.

When $A \subseteq \mathbb{R}^d$ is *generic*, all cells in $\mathrm{Del}(A)$ are simplices, so $\mathrm{Del}(A)$ is a simplicial complex in $\mathbb{R}^d$. Correspondingly, every p-cell of $\mathrm{Vor}(A)$ is the common intersection of exactly $d-p+1$ Voronoi domains, so the common intersection of any $d+2$ Voronoi domains is necessarily empty. This is what we call a *simple* decomposition of $\mathbb{R}^d$. In this case, the Delaunay mosaic is isomorphic to the *nerve* of the Voronoi tessellation, which consists of all collections of domains in $\mathrm{Vor}(A)$ that have a non-empty common intersection. The assumption that A be generic often simplifies matters, and it can be simulated computationally [47] to avoid cumbersome special cases.

Write A_r for the set of points at distance at most r from at least one data point, and note that the Voronoi tessellation decomposes A_r into convex sets, each the intersection of a round ball with a convex polyhedron, see left panel of Figure 2.5. Taking the dual of this decomposition, we get the *alpha complex* for radius r, denoted $\mathrm{Alf}_r(A)$, which we observe is a subcomplex of the Delaunay mosaic; see the right panel of Figure 2.5. The *alpha shape* for radius r is the part of $\mathbb{R}^2$ covered by the simplices of $\mathrm{Alf}_r(A)$. Since the Delaunay mosaic is generically a simplicial complex, so is $\mathrm{Alf}_r(A)$ for every r. When r increases, $\mathrm{Alf}_r(A)$ stays constant or gains new simplices. It follows that for each simplex there is a threshold beyond which the simplex belongs to the alpha complex. Write $\mathrm{Rad}\colon \mathrm{Del}(A) \to \mathbb{R}$ for the function that maps each simplex to this threshold, refer to Rad as the *radius function* on the Delaunay mosaic, and observe that the alpha complexes are its sublevel sets: $\mathrm{Alf}_r(A) = \mathrm{Rad}^{-1}[0,r]$ for every $0 \leq r \leq \infty$. The radius function on the Delaunay mosaic was first introduced in [44], along with its sublevel sets. Three-dimensional alpha shapes have found ample applications in shape modeling [48, 59, 78] and in the analysis of biomolecules [45].

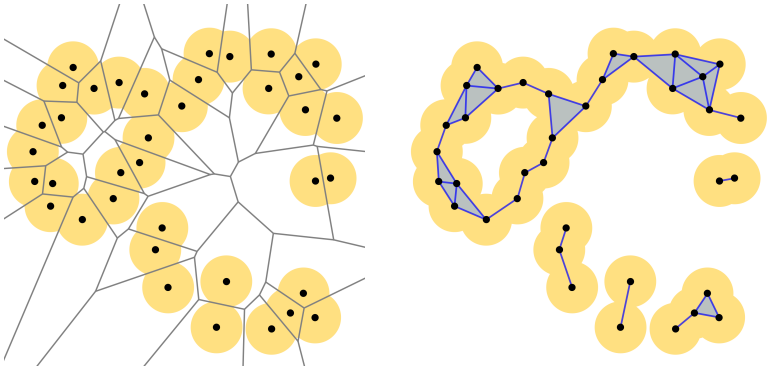

Figure 2.5: *Left*: the Voronoi decomposition of the union of balls centered at the data points. *Right*: the dual alpha complex.

Given simplices $\alpha \subseteq \gamma$ in a simplicial complex K, write $[\alpha, \gamma]$ for the simplices β that satisfy $\alpha \subseteq \beta \subseteq \gamma$; that is: $[\alpha, \gamma]$ is an *interval* in the face poset of K. Given a monotonic function $f\colon K \to \mathbb{R}$, an *interval of* f is an interval on which f is constant, and it is *maximal* if it is not contained in a larger interval of f.

Definition 2.2.2 (Generalized Discrete Morse Function)**.** *A monotonic function on a simplicial complex,* $f\colon K \to \mathbb{R}$*, is* generalized discrete Morse *if the maximal intervals of* f *partition* K*.*

Equivalently, f is generalized discrete Morse if every level set, $K_t = f^{-1}(t)$, is the disjoint union of maximal intervals. Discrete Morse theory was introduced by Forman [55]. In parallel, Edelsbrunner developed the wrap algorithm in an industrial setting, asking for the connection to discrete Morse theory, which was later established by Bauer and Edelsbrunner [8].

Theorem 2.2.1 ([8, Corollary 4.6])**.** *Let $A \subseteq \mathbb{R}^d$ be finite and generic. Then $Rad\colon Del(A) \to \mathbb{R}$ is a generalized discrete Morse function.*

The significance of this result is that we can construct each alpha complex by adding one interval at a time. If this interval consists of two or more simplices, then the addition does not affect the homotopy type of the complex. Indeed, the complex before is a deformation retract of the complex after the addition of the interval. On the other hand, if the interval consists of a single simplex, then the addition of this simplex changes the homology of the complex in a controlled manner.

Lifting. Voronoi tessellations and Delaunay mosaics correspond to certain convex polytopes in one dimension higher. To see this, consider a finite set of points A in the Euclidean space $\mathbb{R}^d$ and *lift* them onto a paraboloid living in $\mathbb{R}^{d+1}$. The Delaunay triangulation of A is the projection of the downward-facing faces of the convex hull of the set of lifted points. The connection has first been studied by Brown [15] and then refined by Edelsbrunner and Seidel [50]. We illustrate the concept precisely in two dimensions.

Let $P = \{(x_1, x_2, x_3) | x_3 = x_1^2 + x_2^2\}$ denote the paraboloid depicted in Figure 2.6. For each point $a = (a_1, a_2)$ in the plane, let $a' = (a_1, a_2, a_1^2 + a_2^2)$ denote its lifted image onto P.

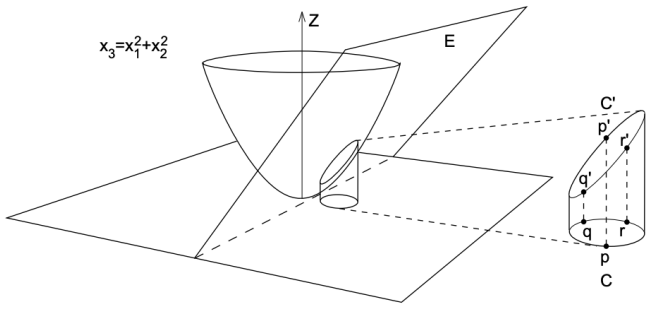

Figure 2.6: Lifting circles onto the paraboloid

Lemma 2.2.1 (Lifted Circles)**.** *Let C be a circle in the plane. Then C' is the intersection of the paraboloid P with a plane.*

Proof. Suppose that C is given by the equation

$$r^2 = (x_1 - c_1)^2 + (x_2 - c_2)^2 = x_1^2 + x_2^2 - 2x_1c_1 - 2x_2c_2 + c_1^2 + c_2^2.$$

By substituting $x_1^2 + x_2^2 = x_3$ we obtain

$$x_3 - 2x_1c_1 - 2x_2c_2 + c_1^2 + c_2^2 - r^2 = 0$$

for the points of C'. This equation defines a plane in 3-dimensions. □

This lemma has an interesting consequence. To make it precise, consider the lower boundary of the convex hull of the lifted points, by which we mean the part of the boundary of the convex hull which is visible from the (x_1, x_2)-plane. Then the following lemma holds.

Lemma 2.2.2 (Delaunay and Convex Hulls)**.** *Each cell of the Delaunay mosaic of A is the vertical projection onto the (x_1, x_2)-plane of a face in the lower boundary of the convex hull of A'.*

Similarly, Voronoi tessellations can be obtained as projection of certain hyperplanes in one dimension up. For $a \in \mathbb{R}^2$ let H_a denote the plane of tangency to the paraboloid P at the lifted point a'. This is the graph of the affine map $h_a : \mathbb{R}^2 \to \mathbb{R}$ defined by $h_a(x) = 2\langle x, a\rangle - \|a\|^2$. Note that h_a encodes the squared Euclidean distance from a; namely $\|x - a\|^2 = \|x\|^2 - h_a(x)$. For any point $b \in A$ the (vertical) distance between the lifted point b' on the paraboloid and the tangent plane at a is

$$\begin{aligned} b' - h_a(b) &= \|b\|^2 - (2\langle b, a\rangle - \|a\|^2) \\ &= \|a - b\|^2. \end{aligned}$$

In words, given points a and b, wherever the tangent plane at a is higher than the tangent plane at b, we are closer to a than to b. We use h_a^+ to denote the half-space of points on or above h_a, and h_a^- for the half-space of points on or below h_a.

By taking the tangent planes at all the lifted points in the set, we obtain an arrangement of planes that decomposes $\mathbb{R}^3$ into convex cells. Roughly speaking, looking at this arrangement "from above", yields the Voronoi tessellation of A. To make this precise, we introduce the *upper envelope* of a plane arrangement, which is the point-wise maximum of the affine functions defining it.

Lemma 2.2.3 (Voronoi and Plane Arrangement)**.** *Each cell of the Voronoi tessellation of A is the vertical projection onto the (x_1, x_2)-plane of a face of the upper envelope of the corresponding plane arrangement.*

Note that the lemma holds in higher dimensions too. Namely, the Voronoi tessellation of points in $\mathbb{R}^d$ can be obtain from projecting the faces of the upper envelope of the corresponding plane arrangement in $\mathbb{R}^{d+1}$.

Algorithms. The deep connection mentioned above has algorithmic consequences, as any $(d+1)$-dimensional convex hull algorithm can be used to compute a d-dimensional Delaunay mosaic and Voronoi tessellation. For completeness, we include a Table 2.7 from [106, Chapter 27], showing the various algorithms that can be used to compute Delaunay mosaic for points in Euclidean space; in fact the divide-and-conquer, incremental, and gift-wrapping algorithms are specialized convex hull algorithms. Running times are given both for worst-case inputs, and for inputs chosen uniformly at random inside a sphere, with expectation taken over input distribution.

The flipping algorithm starts with an arbitrary triangulation and flips the edges according to a local condition which ensures that we obtain the Delaunay triangulation. The plane sweep algorithm computes a planar Delaunay triangulation using a horizontal line that sweeps upward across the plane. The planar divide-and-conquer algorithm uses a splitting line to partition the point set into two equal halves, recursively computes the Delaunay triangulation of each half, and then merges the two subtriangulations in linear time. The incremental algorithm adds sites one by one, updating the Delaunay triangulation after each addition.

ALGORITHM	DIM	WORST CASE	UNIFORM
Flipping	2	$O(n^2)$	
Plane sweep	2	$O(n \log n)$	
Divide-and-conquer	2	$O(n \log n)$	$O(n)$
Randomized incremental	2	$O(n \log n)$	
Randomized incremental	≥ 3	$O(n^{\lceil d/2 \rceil})$	$O(n \log n)$
Gift-wrapping	≥ 2	$O(n^{\lceil d/2 \rceil + 1})$	$O(n)$

Figure 2.7: Algorithms Delaunay mosaic for points in Euclidean space

2.3 Higher-order Complexes

Higher-order Voronoi tessellations are natural and useful generalizations of classical Voronoi tessellations. Given a set A of n point sites in d-space, and an integer k between 1 and $n-1$, the order-k Voronoi tessellation of A, $\mathrm{V}or_k(A)$, partitions the space into regions such that each point within a fixed region has the same k closest sites. $\mathrm{V}or_1(A)$ is just the classical Voronoi tessellation of A. The regions of $\mathrm{V}or_k(A)$ are convex polyhedra, as they arise as the intersection of halfspaces bounded by bisecting hyperplanes of the sites. A subset M of k sites in A has a non-empty region in $\mathrm{V}or_k(A)$ if there is a sphere that encloses M but no site in $A \setminus M$. In fact, the region of M in $\mathrm{V}or_k(A)$ just is the set of centers of all such spheres. Two differences to the classical Voronoi tessellation are apparent. A region need not contain its defining sites, and the bisector of two sites may contribute more than one facet. See the left panel of Figure 2.8, which illustrates a planar order-1 and order-2 Voronoi tessellation superimposed.

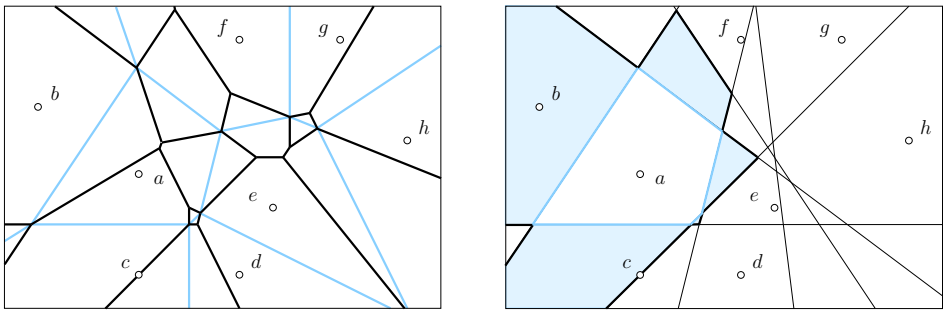

Figure 2.8: *Left panel:* starting with the *blue* (order-1) Voronoi tessellation of the points, we construct the order-2 Voronoi tessellation by dividing up the order-1 regions with *solid black* lines and merging them across the *blue* lines. *Right panel:* the bisectors of a and all other points divide the plane into the Brillouin zones of a. The highlighted second Brillouin zone is where a expands from the order-1 to the order-2 Voronoi tessellation; compare with *left panel*.

While Voronoi tessellations go back more than 100 years to the seminal work of Voronoi [107] or earlier, higher order Voronoi tessellations have been introduced only recently, by Shamos and Hoey [97] in computational geometry and by Gabor Fejes Toth [54] in discrete geometry. Particularly important for this thesis is the incremental algorithm of Lee [73], which also serves as inductive counting argument and establishes that the order-k Voronoi tessellation of n points in $\mathbb{R}^2$ has $\Theta(kn)$ vertices, edges, and regions. This implies that the first k higher order Voronoi tessellations have size $\Theta(k^2 n)$. The latter bound was extended to $\Theta(k^{\lceil \frac{d+1}{2} \rceil} n^{\lfloor \frac{d+1}{2} \rfloor})$ in

$\mathbb{R}^d$ by Clarkson and Shor [20]. Indeed, it is easy to give tight bounds on the total size, over all orders $1 \leq k \leq n-1$, but there are no good bounds known for individual orders beyond 2 dimensions.

To illustrate the difficulties, we mention that the size of the (order-1) Voronoi tessellation of n points in $\mathbb{R}^3$ depends not only on n but also on how the points are distributed in space. If the points are uniformly distributed within the unit cube, the expected size is $\Theta(n)$, but if the points are placed on the moment curve, then the size is $\Theta(n^2)$. On the other hand, the total size, over all orders, depends only on n and is therefore the same for both sets. This suggests that for large values of k, the uniformly distributed points have larger Voronoi tessellations than the points on the moment curve, and this has been experimentally quantified in [49].

Key definitions. Let A be a finite set of points in $\mathbb{R}^d$ and write $n = \#A$ for the cardinality. For any subset $Q \subseteq A$, the *region* of Q is the set of points in $\mathbb{R}^d$ that are at least as close to the points in Q as to the points not in Q. Each such region is a d-dimensional convex polyhedron, and the common intersection of any collection of regions, each defined by the same number of points, is either empty or a face common to all of them. We follow [73, 97] in defining the *order-k Voronoi tessellation* of A, denoted $\mathrm{Vor}_k(A)$, as the polyhedral complex whose cells are the regions defined by subsets Q of size k together with all their faces; see Figure 2.8, left panel. By definition, the order-0 tessellation consists of a single region, which is the entire $\mathbb{R}^d$.

The set of points in $\mathbb{R}^d$ for which $a \in A$ is the k-th nearest is the *k-th Brillouin zone* of a. As illustrated in the right panel in Figure 2.8, this set consists of a number of regions in the arrangement formed by the bisectors of a and the other points in A. The first Brillouin zone is a convex polyhedron, and each of the other zones has the homotopy type of a sphere. Furthermore, the union of the first k zones is star-convex, with a in the kernel; see [54]. Importantly, for $k \geq 2$, every region in the k-th Brillouin zone is a d-dimensional convex polytope whose boundary can be partitioned into the *near boundary*, which is visible from a, the *far boundary*, which is not visible from a, and the *silhouette*, which separates the near and far boundaries. By convexity, the silhouette is homeomorphic to a $(d-2)$-sphere that splits the boundary into two pieces, each homeomorphic to an open $(d-1)$-ball.

Lemma 2.2.3 extends to the setting of higher-order Voronoi tessellations. Consider the collection of d-dimensional planes in $\mathbb{R}^{d+1}$ obtained by mapping each point $a \in A$ to the affine function $h_a \colon \mathbb{R}^d \to \mathbb{R}$ defined by $h_a(x) = 2\langle x, a\rangle - \|a\|^2$. Note that h_a encodes the squared Euclidean distance from a: $\|x-a\|^2 = \|x\|^2 - h_a(x)$. The graph of h_a is a (non-vertical) d-plane in $\mathbb{R}^{d+1}$. The collection of d-planes decomposes $\mathbb{R}^{d+1}$ into convex cells of dimension $0 \leq i \leq d+1$, referred to as the *arrangement* of d-planes. We call the $(d+1)$-cells *chambers*, and the d-cells *facets*. For $1 \leq k \leq n$, the *k-th level* of the arrangement is the set of points $(x, y) \in \mathbb{R}^d \times \mathbb{R}$ such that $h_a(x) < y$ for at most $k-1$ affine maps and $h_a(x) > y$ for at most $n-k$ affine maps. The *k-th belt* is the set of points between the k-th level and the $(k+1)$-st level.

Lemma 2.3.1 ([50, From Arrangement to Tessellation])**.** *Let A be a set of n points in $\mathbb{R}^d$, let $0 \leq k \leq n$, and recall that A defines an arrangement of n non-vertical d-planes in $\mathbb{R}^{d+1}$.*

- *There is a bijection between the regions of $\mathrm{Vor}_k(A)$ and the chambers of the k-th belt such that each region is the vertical projection of the corresponding chamber.*

- *The k-th Brillouin zone of $a \in A$ is the vertical projection of the k-th level intersected with the d-plane defined by a.*

2.4 Arrangements

In his book *Lectures on Discrete Geometry*, Matoušek starts the chapter on arrangements [80, Chapter 6] with the following quote: "Arrangements of lines in the plane and their higher-dimensional generalization, arrangements of hyperplanes in $\mathbb{R}^d$, are a basic geometric structure whose significance is comparable to that of convex polytopes. In fact, arrangements and convex polytopes are quite closely related: A cell in a hyperplane arrangement is a convex polyhedron, and conversely, each hyperplane arrangement in $\mathbb{R}^d$ corresponds canonically to a convex polytope in $\mathbb{R}^{d+1}$ of a special type, the so-called *zonotope*". While this thesis does not delve into this connection further, it's worth noting, especially when considered alongside Lemma 2.3.1, that arrangements serve as a cohesive structure unifying the geometric objects discussed in this chapter.

Line arrangements. Consider a finite set of n lines in the Euclidean plane. The lines induce a decomposition of $\mathbb{R}^2$ into connected open cells, which we call an *arrangement of lines* or *line arrangement*. A 0-dimensional cell (a vertex) is the intersection point of two lines in L; a 1-dimensional cell (an edge) is a maximal connected portion of a line in L that is not intersected by any other line in L; and a 2-dimensional cell (a face) is a maximal connected region of $\mathbb{R}^2$ not intersected by any line in L. An arrangement of lines is *simple* if no three lines go through a point, and no two lines are parallel. The *complexity* of L is its number of vertices, edges and faces. In a simple arrangement of n lines, every pair of lines intersects in a vertex. We therefore have $\binom{n}{2}$ vertices. Furthermore, each line is subdivided into n edges, giving n^2 edges in total. The number of faces can be deduced using the Euler characteristic, which is 1 for the entire plane. Hence $\#f = 1 - \#v + \#e = 1 - \binom{n}{2} + n^2 = 1 + n(n+1)/2$.

The *k-level* of a line arrangement is the polygonal chain formed by the edges that have exactly $k-1$ other lines directly below them. Finding matching upper and lower bounds for the complexity of a k-level remains a major open problem in discrete geometry. It is often referred as the dual version of the two-dimensional k-set problem. We summarize the known results for the case $k = n/2$ (with n even) in the paragraph below. We state them in the primal version using the notion of halving line.

Halving lines. Assuming n is an even positive integer, a *halving line* of a set of n points in $\mathbb{R}^2$ has $n/2$ points on each side. To determine the maximum number of halving lines n points can have is one of the most vexing combinatorial questions in discrete geometry. The earliest non-trivial results were a lower bound of $\Omega(n \log n)$ and an upper bound of $O(n\sqrt{n})$; see [52, 53, 76]. The upper bound was improved to $O(n\sqrt{n}/\log^* n)$ in [85]. The current best bounds are $\Omega(n \exp(c\sqrt{\log n}))$, for some positive constant c, due to Geza Tóth [105], and $O(n^{4/3})$ due to Tamal Dey [35]. For sets with minimum distance 1 and diameter $O(\sqrt{n})$, the upper bound can be further improved to $O(n^{7/6})$; see [51, Theorem 2]. In contrast to counting halving lines, bounding the number of lines that separate at most k points from the rest much easier and a tight upper bound of kn for $k < n/2$ can be found in [2]; see also [62, 109].

Spherical arrangements. Analogously to the two dimensional case, an *arrangement of hyperplanes* in $\mathbb{R}^d$ is a decomposition of $\mathbb{R}^d$ into connected open cells of dimensions $0, 1, \cdots, d$ [106, Chapter 28]. To finesse the inconvenience of unbounded cells, it is sometimes convenient to work with arrangements of $(d-1)$-dimensional great-sphere instead of hyperplanes; and consider the arrangement formed by these great-spheres in $\mathbb{S}^d$. Besides having only bounded cells, the great-sphere arrangement is centrally symmetric and thus has two antipodal cells for each bounded cell and each pair of diametrically opposite unbounded cells in the hyperplane arrangement. See Figure 2.9. Two points in $\mathbb{S}^d$ are distinguished: the *north-pole* at the very

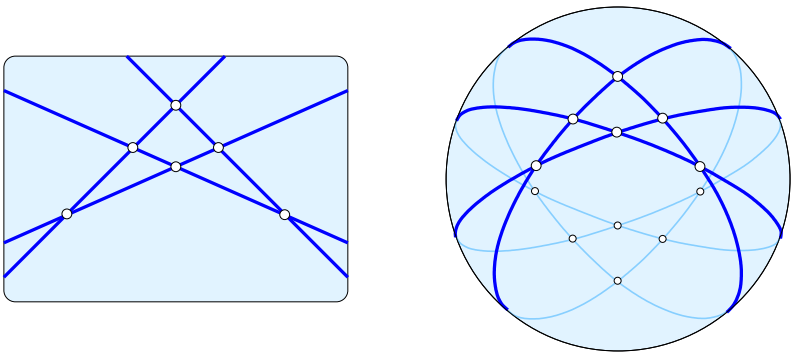

Figure 2.9: An arrangement of four lines in $\mathbb{R}^2$ on the *left* and the corresponding arrangement of four great-circles in $\mathbb{S}^2$ on the *right*.

top and the *south-pole* at the very bottom of the sphere . We assume that none of the great-spheres passes through the two poles. Letting σ be a great-sphere in $\mathbb{S}^d$, we write σ^- for the closed *lower hemisphere* bounded by σ, which contains the south-pole, and we write σ^+ for the closed *upper hemisphere*, which contains the north-pole.

Letting A be the collection of great-spheres, each *cell* in the *arrangement* corresponds to a tri-partition, $A = A^- \sqcup A^0 \sqcup A^+$, such that the cell is the common intersection of the lower hemispheres σ^-, with $\sigma \in A^-$, the great-spheres σ, with $\sigma \in A^0$, and the upper hemispheres σ^+, with $\sigma \in A^+$.

We write $\mathcal{A}$ for the arrangement defined by A, we refer to a cell of dimension p as a p-*cell*, and for $p = 0,\ 1,\ 2,\ d-1,\ d$, we call it a *vertex*, *edge*, *polygon*, *facet*, *chamber*, respectively. The *faces* of a cell are the cells contained in it, which includes the cell itself.

The intersection of great-spheres is again a great-sphere, albeit of a smaller dimension. To avoid any confusion, we will explicitly mention the dimension if it is less than $d-1$. We call the arrangement *simple* if all great-spheres avoid the two poles and the common intersection of any $d-p$ great-spheres is a p-dimensional great-sphere in $\mathbb{S}^d$. This implies that any d great-spheres intersect in a pair of antipodal points, and any $d+1$ or more great-spheres have an empty common intersection. For each $0 \leq p \leq d$, we write $C^p = C^p(\mathcal{A})$ for the number of p-cells in the arrangement, and $C^p(n,d)$ for the maximum over all arrangements of n great-spheres in $\mathbb{S}^d$. Importantly, the number of cells is maximized if the arrangement is simple, and in this case it depends on the number of great-spheres but not on the great-spheres themselves. A fundamental question in the study of arrangements is how complex a certain arrangement (or portion of it) can be. Answering this question is often a prerequisite to the analysis of algorithms to construct arrangements.

Proposition 2.4.1 (Number of Cells)**.** *Any simple arrangement of $n \geq d$ great-spheres in $\mathbb{S}^d$ has $C^p(n,d) = 2\left[\binom{d}{p}\binom{n}{d} + \binom{d-2}{p-2}\binom{n}{d-2} + \ldots + \binom{d-2i}{p-2i}\binom{n}{d-2i}\right]$ p-cells, in which $i = \lfloor p/2 \rfloor$.*

The formula for the number of p-cells is not new and can be derived from similar formulas for arrangements in d-dimensional real projective space [64, Section 18.1] or in d-dimensional Euclidean space [42, Section 1.2].

CHAPTER 3

Arrangements of Hyperplanes

The core result of this chapter, which is based on [11], is a system of Dehn–Sommerville type relations for level sets of the depth function on spherical arrangements. This is different but related to the more direct generalization of the Dehn–Sommerville relations to levels in arrangements proved by Linhart, Yao and Phillip [74]. Similar to their classic relatives and the generalization in [74], our relations are based on double-counting, but instead counting cells, we take sums of topological indicators. To state the relations, let $\mathcal{A}$ be an arrangement of n great-spheres in $\mathbb{S}^d$, and write $C_k^p(\mathcal{A})$ for the number of p-cells at depth k in $\mathcal{A}$. For each p-cell, consider the alternating sum of its faces at the same depth, and write $E_k^p(\mathcal{A})$ for the sum of such alternating sums over all p-cells at depth k. If $\mathcal{A}$ is simple, then we have a system of linear relations for $0 \leq p \leq d$ and $0 \leq k \leq n - d + p$:

$$\sum_{i=0}^{p}(-1)^i \tbinom{d-i}{d-p} E_k^i(\mathcal{A}) = C_k^p(\mathcal{A}) = \sum_{i=0}^{p} \tbinom{d-i}{d-p} E_{k+i-p}^i(\mathcal{A}), \tag{3.1}$$

which we refer to as *Dehn–Sommerville–Euler relations*. The system has applications to *cyclic polytopes*—which are convex hulls of finitely many points on the moment curve—and the broader class of *neighborly polytopes*—which are characterized by the property that every $(q-1)$-simplex spanned by $q \leq d/2$ vertices is a face of the polytope. A celebrated result in the field is the Upper Bound Theorem proved by McMullen [82], which states that every cyclic polytope has at least as many faces of any dimension as the convex hull of any other set of n points in $\mathbb{R}^d$. All cyclic polytopes with n vertices in $\mathbb{R}^d$ have isomorphic face complexes with a structure that is simple enough to allow for counting the faces, and expressions for these numbers can be found in textbooks, such as [110]. In contrast, neighborly polytopes with n vertices in $\mathbb{R}^d$ can have non-isomorphic face complexes, but they still have the same number of faces in every dimension. Within our framework, the structural simplicity is expressed by having bi-polar restrictions of the depth function to the intersection of any $q \leq d/2$ great-spheres. We call an arrangement in $\mathbb{S}^d$ that has this property a *neighborly arrangement*. Writing $p = d - q$ and counting only the cells of the subarrangement, $\mathcal{B}$, in the intersection of the q great-spheres, straightforward topological arguments imply

$$E_k^p(\mathcal{B}) = \begin{cases} 1 & \text{for } k = 0, \\ 0 & \text{for } 1 \leq k \leq n + p - d - 1, \\ (-1)^p & \text{for } k = n + p - d. \end{cases} \tag{3.2}$$

Together with the Dehn–Sommerville–Euler relations in (3.1), this implies expressions in n, d, p, and k for the number of p-faces, for *every* $0 \leq p \leq d$, and thus generalizes the result for

convex polytopes to levels in neighborly arrangements. Surprisingly, the neighborly property not only determines the number of faces of the convex hull but in fact of every level of the corresponding dual arrangement. The special case of cyclic polytopes, in which the hyperplanes are dual to points on the moment curve, has been solved in [3].

3.1 Depth Function

Given a set A of n great-spheres in $\mathbb{S}^d$, none passing through the two poles, we define the *depth* of a point $x \in \mathbb{S}^d$ as the number of great-spheres $\sigma \in A$ with $x \in \sigma^- \setminus \sigma$. In words, the depth of the point is the number of great-spheres that cross the shortest arc connecting x to the north-pole. If x and y are two interior points of the same cell, then they have the same depth. Recalling that $\mathcal{A}$ is the arrangement defined by A, we introduce the *depth function*, $\theta\colon \mathcal{A} \to [0, n]$, which we define by mapping each cell to the depth of its interior points. Depending on the situation, we think of θ as a discrete function on the arrangement or a piecewise constant function on $\mathbb{S}^d$, namely constant in the interior of every cell in $\mathcal{A}$.

Let c be a p-cell in $\mathcal{A}$, with corresponding tri-partition $A^- \sqcup A^0 \sqcup A^+$. The depth of every interior point $x \in c$ is $\theta(x) = \theta(c) = \#A^-$, and if the arrangement is simple, then $p = d - \#A^0$. Let $b \subseteq c$ be a face of dimension $i \leq p$, with corresponding tri-partition $B^- \sqcup B^0 \sqcup B^+$. We have $B^- \subseteq A^-$, $A^0 \subseteq B^0$, $B^+ \subseteq A^+$, and if the arrangement is simple, we also have $i = d - \#B^0$. Given the depth of c, this implies the following bounds on the depth of b:

Lemma 3.1.1 (Depth of Face)**.** *Let $\mathcal{A}$ be a simple arrangement of great-spheres in $\mathbb{S}^d$. For every i-face, b, of a p-cell, c, we have $\max\{0, \theta(c) + i - p\} \leq \theta(b) \leq \theta(c)$, and both bounds on the depth of b are tight.*

Proof. Since the arrangement is simple, we have $\#B^- \geq \#A^- - [\#B^0 - \#A^0] = \#A^- + i - p$, which implies the first inequality. The second inequality follows from $\#B^- \leq \#A^-$, which holds for general and not necessarily simple arrangements.

To prove the second inequality is tight, we show the existence of a p-cell that shares b with c and has the same depth as b. To this end, consider the tri-partition $(B^+ \cup X) \sqcup (B^0 \setminus X) \sqcup B^-$, in which $X \subseteq B^0$ has cardinality $p - i$. The cell defined by this tri-partition is non-empty because it contains b as a face. Furthermore, this cell has dimension p and the same depth as b. The proof that the first inequality is tight is symmetric and omitted. □

To relate this concept to the prior literature, we mention that [42, Chapter 3] introduces the *k-th level* of an arrangement of n non-vertical hyperplanes in d dimensions as the points $x \in \mathbb{R}^d$ below fewer than k and above fewer than $n - k$ of the hyperplanes. In other words, the k-th level consists of all facets at depth $k - 1$ and all their faces. Assuming the arrangement is simple, Lemma 3.1.1 implies that a p-cell belongs to the k-th level iff its depth is between $k - d + p$ and $k - 1$.

3.2 Sublevel Sets

Assume that $\mathcal{A}$ has at least one vertex, which in the simple case is implied by $n \geq d$. For $0 \leq k \leq n$, we write $\mathcal{A}_k = \theta^{-1}[0, k]$ for the *sublevel set* of θ at k. It consists of all cells in $\mathcal{A}$ whose depth is k or less. Recall that θ is *monotonic*, by which we mean that the depth of

every cell is at least as large as the depth of any of its faces. It follows that $\mathcal{A}_k$ is a complex, with well defined *Euler characteristic*:

$$\chi(\mathcal{A}_k) = \sum_{c \in \mathcal{A}_k} (-1)^{\dim c}. \tag{3.3}$$

The right-hand side of (3.3) explains how the Euler characteristic changes from $\mathcal{A}_{k-1}$ to $\mathcal{A}_k$, namely by adding the alternating sum of all cells at depth k. By Lemma 3.1.1, every cell at depth k is a face of a chamber at depth k. We can therefore construct $\mathcal{A}_k$ from $\mathcal{A}_{k-1}$ by adding all chambers at depth k together with their faces at the same depth. This motivates the following two definitions.

Definition 3.2.1 (Relative Euler and Depth Characteristic)**.** *For a cell $c \in \mathcal{A}$, let $F = F(c)$ be the complex of faces, which includes c, and let $F_0 \subseteq F$ be a subcomplex. The* relative Euler characteristic *of the pair of complexes is $\chi(F, F_0) = \sum_{b \in F \setminus F_0} (-1)^{\dim b}$. If F_0 is the set of faces $b \subseteq c$ with $\theta(b) < \theta(c)$, denoted $U = U(c)$, we call $\varepsilon(c) = \chi(F, U)$ the* depth characteristic *of c, and we call c* critical *for θ if $\varepsilon(c) \neq 0$.*

For example, if all faces have the same depth as c, then the depth characteristic of c is $\varepsilon(c) = \chi(F, \emptyset) = 1$, and if all proper faces have depth strictly less than c, then the depth characteristic of c is $\varepsilon(c) = \chi(F, F \setminus \{c\}) = (-1)^{\dim c}$. In both cases, c is critical.

Lemma 3.2.1 (Relative and Absolute Euler Characteristic)**.** *Let $F = F(c)$ be the face complex of a cell, c, in an arrangement, and let $F_0 \subseteq F$ be a subcomplex. Then the relative Euler characteristic of the pair is $\chi(F, F_0) = 1 - \chi(F_0)$.*

Proof. By definition, $\chi(F, F_0) + \chi(F_0)$ is the sum of $(-1)^{\dim b}$ over all cells $b \in F \setminus F_0$ as well as all $b \in F_0$, and therefore over all $b \in F$. Hence, this sum is $\chi(F)$, which is equal to 1 because c is closed and convex. The claimed equation follows. □

We write $C_k^p = C_k^p(\mathcal{A})$ for the number of p-cells at depth k, and $E_k^p = E_k^p(\mathcal{A}) = \sum_c \varepsilon(c)$ for the sum of depth characteristics over all p-cells at depth k. To see the motivation behind taking sums of depth characteristics, consider the subcomplex of cells at depth at most k in a p-dimensional subarrangement of the d-dimensional arrangement. It is pure p-dimensional, by which we mean that every cell in this subcomplex is a face of a p-cell. Furthermore, the Euler characteristic of this pure complex is the sum of depth characteristics of its p-cells. In other words, we can construct the subarrangement by adding its p-cells in the order of non-decreasing depth. Whenever we add a p-cell, c, we also add the yet missing faces, and we know that $\varepsilon(c)$ is the increment to the Euler characteristic of the subcomplex. Hence, E_k^p is the increment to the total Euler characteristic of the subcomplexes in the p-dimensional subarrangements when we add the p-cells at depth k together with their yet missing faces.

3.3 Local Configurations

Most arguments in the subsequent technical sections accumulate local quantities, each counting faces or cofaces of a cell. In a simple arrangement, the coface structure depends only on the dimension, so we study it first.

3.3.1 Coface Structure

In the generic case, the local neighborhood of a vertex in an arrangement in $\mathbb{S}^d$ looks like that of the origin in the arrangement of the d coordinate planes in $\mathbb{R}^d$. Each of these $(d-1)$-planes bounds an open half-space in which the corresponding coordinate is strictly negative. Accordingly, we define the *depth* of a point $x \in \mathbb{R}^d$ as the number of negative coordinates, and the *depth* of a cell in the arrangement as the depth of its interior points. To study this arrangement, consider $[-1,1]^d \subseteq \mathbb{R}^d$ and let $S^p(d)$ be the number of q-sides of the d-cube, in which we write $q = d - p$. The dual correspondence provides an incidence reversing bijection between the p-cells of the arrangement and the q-sides of the cube. We label each side with

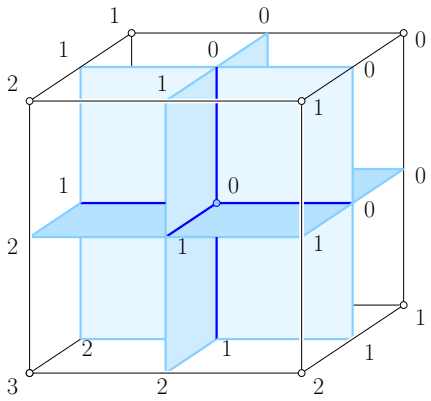

Figure 3.1: The neighborhood of the origin in $\mathbb{R}^3$ and the dual cube centered at the origin. The labels of the sides are the depths of the corresponding cells in the arrangement of coordinate planes.

the depth of the corresponding cell in the arrangement, and write $S^p_k(d)$ for the number of q-sides labeled k. As illustrated in Figure 3.1, this amounts to labeling $S^d_k(d) = \binom{d}{k}$ vertices with k, for $0 \le k \le d$, and labeling each side with the minimum label of its vertices. Note that the label of a q-side cannot exceed $d - q = p$.

Lemma 3.3.1 (Coface Structure of Vertex)**.** *Consider the arrangement defined by the d coordinate planes in $\mathbb{R}^d$.*

(i) For $0 \le k \le p \le d$, the number of p-cells at depth k is $S^p_k(d) = \binom{d-k}{d-p}\binom{d}{k}$.

(ii) There is one cell at depth d, namely the negative orthant, and for $0 \le k < d$, the alternating sum of cells at depth k vanishes; that is: $\sum_{p=k}^{d}(-1)^p S^p_k(d) = 0$.

Proof. The p-cells counted in (i) correspond to the q-sides with label k, in which $p + q = d$. To count these q-sides, we recall that the d-cube has $\binom{d}{k}$ vertices at depth k. For each such vertex, u, consider the largest side for which u is the vertex with minimum label. This largest side is a cube of dimension $d-k$, which contains $\binom{d-k}{q}$ q-sides incident to u. We thus get

$$S^p_k(d) = \binom{d-k}{q}\binom{d}{k} = \binom{d-k}{d-p}\binom{d}{k} \tag{3.4}$$

q-sides with label k, which proves (i).

To see (ii), consider a $(d-k)$-cube with label k. The alternating sum of sides with the same label is $\sum_{q=0}^{d-k}(-1)^q\binom{d-k}{q}$, which vanishes for $d-k>0$, and equals 1 for $d-k=0$. Likewise,

the sum of alternating sums over all $(d-k)$-sides with label k vanishes for $d-k>0$ and equals 1 for $k=d$. This implies (ii) by duality. □

It is easy to generalize Lemma 3.3.1 from a vertex to a cell of dimension $i \geq 0$. To see this geometrically, we slice the i-cell and its cofaces with a $(d-i)$-plane orthogonal to the i-cell. In this slice, the i-cell appears as a vertex, and each coface of dimension p appears as a $(p-i)$-cell.

Corollary 3.3.1 (Coface Structure of Cell)**.** *Consider the arrangement defined by the d coordinate planes in $\mathbb{R}^d$, and let c be an i-cell at depth $0 \leq \ell \leq i$.*

(i) For $0 \leq k-\ell \leq p-i \leq d-i$, the number of p-cells at depth k that contain c is $S^{p-i}_{k-\ell}(d-i) = \binom{d-i-k+\ell}{d-p}\binom{d-i}{k-\ell}$.

(ii) There is one cell at depth d, and for $\ell \leq k < d$, the alternating sum of cells at depth k that contain c vanishes; that is: $\sum_{p=k}^{d}(-1)^p S^{p-i}_{k-\ell}(d-i) = 0$.

3.3.2 Face Structure

The face structure of a cell in a simple arrangement is not quite as predictable as its coface structure. Nevertheless, we can say something about it. As before, we write $F = F(c)$ for the face complex of a cell, c, and we let $F_0 \subseteq F$ be a subcomplex. Furthermore, we write

$$\mathrm{X}(F, F_0) = \sum\nolimits_{b \in F \setminus F_0} (-1)^{\dim b} \chi(F(b), F_0 \cap F(b)) \tag{3.5}$$

for the alternating sum of relative Euler characteristics.

Lemma 3.3.2 (Face Structure of Cell)**.** *Let c be a cell in a simple arrangement of great-spheres in $\mathbb{S}^d$, and let $F_0 \subseteq F(c)$ be a subcomplex of the face complex of the cell. Then $X(F, F_0) = 1$ if $F_0 \neq F$ and $X(F, F_0) = 0$ if $F_0 = F$.*

Proof. If $F_0 = F$, then $\mathrm{X}(F, F_0)$ is a sum without terms, which is 0. We can therefore assume $F_0 \neq F$, which implies $c \in F \setminus F_0$. Fix a cell $a \in F \setminus F_0$ with dimension $i = \dim a$ less than or equal to $p = \dim c$. It contributes $(-1)^{i+j}$ for every j-cell $b \in F \setminus F_0$ that contains a as a face. The contribution of a to $\mathrm{X}(F, F_0)$ is therefore $(-1)^i \sum_{j=1}^{p}(-1)^j \binom{p-i}{j-i}$, which vanishes for all $i < p$ and is equal to 1 for $i = p$. Hence, the only non-zero contribution to $\mathrm{X}(F, F_0)$ is for $a = c$, which implies the claim. □

There is a symmetric form of the lemma, which we get by introducing the *codepth function*, $\vartheta \colon \mathcal{A} \to [0, n]$ defined by $\vartheta(x) = n - q - \theta(x)$, where q is the number of great-spheres that pass through x. Observe that $\vartheta(x)$ is the number of great-spheres that cross the shortest arc connecting x to the south-pole. We write $B^p_\ell(\mathcal{A})$ for the number of p-cells with codepth ℓ. If the arrangement is simple, then

$$B^p_\ell(\mathcal{A}) = C^p_k(\mathcal{A}), \quad \text{with } k + \ell + (d-p) = n. \tag{3.6}$$

Indeed, there are $d-p$ great-spheres that contain a p-cell, c, and if k great-spheres pass above c, then $\ell = n - (k + d - p)$ great-spheres pass below c. Recall that $\varepsilon(c) = \chi(F, U)$ is the depth characteristic, in which $F = F(c)$ is the face complex, and $U \subseteq F$ is the subcomplex of faces at depth strictly less than $\theta(c)$. Symmetrically, we call $\delta(c) = \chi(F, L)$ the *codepth*

characteristic of c, in which $F = F(c)$ as before, and $L \subseteq F$ is the subcomplex of faces at codepth strictly less than $\vartheta(c)$. In a simple arrangement, the two characteristics agree on even-dimensional cells, and they are the negative of each other for odd-dimensional cells.

Lemma 3.3.3 (Depth and Codepth Characteristics)**.** *For a p-cell, c, in a simple arrangement of great-spheres, we have $\delta(c) = (-1)^p \varepsilon(c)$.*

Proof. The boundary of c is a $(p-1)$-sphere, which is decomposed by the complex of proper faces of c. We write L for the proper faces with codepth strictly less than $\vartheta(c)$, and U for the proper faces with depth strictly less than $\theta(c)$. L and U exhaust the proper faces of c. More precisely, L and U partition the $(p-1)$-faces, and each of the two subcomplexes is the closure of its set of $(p-1)$-faces. Hence, $L \cap U$ consists of all $(p-2)$-faces shared by a $(p-1)$-face in L and another $(p-1)$-face in U, together with all faces of these $(p-2)$-faces. Since the arrangement is simple, the cells in $L \cap U$ decompose a $(p-2)$-manifold.

Case 1: p is odd. Then $L \cap U$ decomposes an odd-dimensional manifold. By Poincaré duality, $\chi(L \cap U) = 0$. The Euler characteristic of the boundary of c is 2, which implies $\chi(L) + \chi(U) - \chi(L \cap U) = \chi(L) + \chi(U) = 2$. By Lemma 3.2.1, $\varepsilon(c) = 1 - \chi(U)$ and therefore $\delta(c) = 1 - \chi(L) = 1 - [2 - \chi(U)] = -\varepsilon(c)$, as claimed.

Case 2: p is even. The boundary of c is an odd-dimensional sphere, so its Euler characteristic vanishes. By Alexander duality, $\chi(L) = \chi(U)$, and by Lemma 3.2.1, $\varepsilon(c) = 1 - \chi(U)$ and $\delta(c) = 1 - \chi(L)$, which implies $\delta(c) = \varepsilon(c)$, as claimed.

□

3.4 Relations

In this section, we prove linear relations for the cells at given depths. The relations are similar to the classic Dehn–Sommerville relations for convex polytopes, and we prove them the same way by straightforward double counting; see [64, Section 9.2]. We begin with the easy bi-polar case.

3.4.1 Bi-polar Depth Functions

We recall that the depth function on an arrangement of great-spheres is bi-polar if there is a chamber above all great-spheres. By construction, the arrangement and its depth function are antipodal, which implies that there is also a chamber below all great-spheres. With the great-spheres given in $\mathbb{S}^d$, the depth function on $\mathbb{S}^d$ is necessarily bi-polar, but its restrictions to subarrangements inside the common intersection of one or more great-spheres are not necessarily bi-polar.

Theorem 3.4.1 (Bi-polar Depth Functions)**.** *Let $\mathcal{A}$ be a simple arrangement of $n \geq d$ great-spheres in $\mathbb{S}^d$, let $\mathcal{B}$ be the p-dimensional subarrangement inside the intersection of $d-p$ of the great-spheres, and assume that the restriction of the depth function to $\mathcal{B}$ is bi-polar. Then*

$$E_k^p(\mathcal{B}) = \begin{cases} 1 & \text{for } k = 0, \\ 0 & \text{for } 1 \leq k \leq n-d+p-1, \\ (-1)^p & \text{for } k = n-d+p. \end{cases} \tag{3.7}$$

Proof. Let c_N be the (p-dimensional) chamber at depth 0 in $\mathcal{B}$, and let c_S be the antipodal chamber at depth $n-d+p$. We write $\mathbb{S}^p$ for the intersection of the $d-p$ great-spheres, fix a point $N \in \mathbb{S}^p$ inside the interior of c_N, and let $S \in \mathbb{S}^p$ in the interior of c_S be the antipodal point. We partition $\mathbb{S}^p \setminus \{N, S\}$ into open fibers, each half a great-circle connecting N to S. Along each fiber, the depth is non-decreasing. Consider the set of fibers that intersect a chamber $c \neq c_N, c_S$. They partition the boundary of c into the *upper boundary*, along which the fibers enter the chamber, the *lower boundary*, along which the fibers exit the chamber, and the *silhouette*, along which the fibers touch but do not enter the chamber. Since c is p-dimensional and spherically convex (the common intersection of closed hemispheres) this implies that the silhouette is a $(p-2)$-sphere, and the upper and lower boundaries are open $(p-1)$-balls. The depth characteristic of c is $(-1)^{p-1}$—for the open lower boundary—plus $(-1)^p$—for the chamber itself. It follows that the depth characteristic of c vanishes, and so does the depth characteristic of every other chamber, except for c_N and c_S. Because c_N has the same depth as its entire boundary, we have $\varepsilon(c_N) = 1$, and because c_S has larger depth than its entire boundary, we have $\varepsilon(c_S) = (-1)^p$. This implies (3.7). □

3.4.2 Alternating Sums of Depth Characteristics

In the general case, the restrictions of the depth function to subarrangements are not necessarily bi-polar. The depth characteristics may therefore violate (3.7), but they satisfy a system of linear relations, as we prove next.

Theorem 3.4.2 (Dehn–Sommerville–Euler for Levels)**.** *Let $\mathcal{A}$ be a simple arrangement of $n \geq d$ great-spheres in $\mathbb{S}^d$. Then for every dimension $0 \leq p \leq d$, we have*

$$\sum_{i=0}^{p} (-1)^i \tbinom{d-i}{p-i} E_k^i(\mathcal{A}) = C_k^p(\mathcal{A}) = \sum_{i=0}^{p} \tbinom{d-i}{p-i} E_{k+i-p}^i(\mathcal{A}) \quad \text{for } 0 \leq k \leq n-d+p. \tag{3.8}$$

Proof. Let c be a p-cell at depth k, let $F = F(c)$ be the face complex of c, and let $U \subseteq F$ be the subcomplex of faces at depth strictly less than k. Note that U does not contain c, so $U \neq F$, and Lemma 3.3.2 implies $\mathrm{X}(F, U) = 1$. Taking the sum over all p-cells at depth k thus gives the number of such p-cells, which is $C_k^p(\mathcal{A})$. By Corollary 3.3.1 (i), a single i-cell contributes to the alternating sums of $S_0^{p-i}(d-i) = \binom{d-i}{p-i}$ p-cells, which implies that the first sum in (3.8) is the total alternating sum of depth characteristics over all cells at depth k and dimension at most p. The second relation in (3.8) is the upside-down version of the first relation. Indeed, we can substitute codepth for depth and get the following relation using the notation of Section 3.3.2:

$$B_\ell^p(\mathcal{A}) = \sum_{i=0}^{p} (-1)^i \tbinom{d-i}{p-i} D_\ell^i(\mathcal{A}). \tag{3.9}$$

To translate this back in term of depth, we set $\ell = n - (k + d - p)$ so that a p-cell at codepth ℓ has depth $n - (\ell + d - p) = k$. Hence, $B_\ell^p(\mathcal{A}) = C_k^p(\mathcal{A})$. To write the Ds in terms of the Es, we multiply with $(-1)^i$ because of Lemma 3.3.3, and we change the index from $\ell = n - (k + d - p)$ to $k + i - p = n - (\ell + d - i)$ because of (3.6). This gives the right relation in (3.8). □

As an example consider the case $d = 2$. We get equations (3.10), (3.11), (3.12) by setting $p = 0, 1, 2$ in (3.8):

$$E_k^0 = C_k^0 = E_k^0, \tag{3.10}$$

$$2E_k^0 - E_k^1 = C_k^1 = 2E_{k-1}^0 + E_k^1, \tag{3.11}$$

$$E_k^0 - E_k^1 + E_k^2 = C_k^2 = E_{k-2}^0 + E_{k-1}^1 + E_k^2, \tag{3.12}$$

Equation (3.10) just says that the depth characteristic of every vertex is 1. (3.11) implies $E^1_k = E^0_k - E^0_{k-1}$, and (3.12) implies $E^1_k + E^1_{k-1} = E^0_k - E^0_{k-2}$, which follows from the relation implied by (3.11). Note that adding the depth characteristics of the edges gives a telescoping series, which implies $E^1_0 + E^1_1 + \ldots + E^1_k = E^0_k$.

3.4.3 Alternating Sums of Cells

For comparison, we state the more traditional version of the Dehn–Sommerville relations, which apply to cell complexes; see [83] and [74, Theorem 1]. It counts the p-cells at depth k, which together with all their faces form a cell complex. For each dimension $0 \leq i \leq p$, this includes all i-cells at depths $k + i - p$ to k.

Proposition 3.4.1 (Dehn–Sommerville for Levels)**.** *Let $\mathcal{A}$ be a simple arrangement of $n \geq d$ great-spheres in $\mathbb{S}^d$. For every dimension $0 \leq p \leq d$, we have*

$$C^p_k(\mathcal{A}) = \textstyle\sum_{i=0}^{p} (-1)^i \binom{d-i}{d-p} \sum_{j=0}^{p-i} \binom{p-i}{p-i-j} C^i_{k+i-p+j}(\mathcal{A}) \quad \text{for } 0 \leq k \leq n-d+p. \tag{3.13}$$

We get a non-trivial relation in (3.13) for $p = 1$, which asserts $C^1_k = dC^0_{k-1} + dC^0_k - C^1_k$. Indeed, twice the number of edges is the sum of vertex degrees. For $p = 2$, we get

$$C^2_k = \tbinom{d}{2} C^0_k - (d-1)C^1_k + C^2_k + (d-1)dC^0_{k-1} - (d-1)C^1_{k-1} + \tbinom{d}{2} C^0_{k-2}, \tag{3.14}$$

in which the polygons cancel and the rest is equivalent to the relation for $p = 1$. More generally, the term on left-hand side of (3.13) cancels whenever p is even.

3.5 Application to Higher-order Voronoi Tessellations

In this section, we give evidence for the unifying power of the system of Dehn–Sommerville–Euler relations by rederiving cell-counting formulas for higher-order Voronoi tessellations proved in [10, 73]. The difference forms of the relations are particularly convenient, which we present in dimensions 3 and 4.

3.5.1 Two Dimensions

Before discussing the 2-dimensional order-k Voronoi tessellations, we introduce the 3-dimensional difference relations implied by Theorems 3.4.1 and 3.4.2.

Corollary 3.5.1 (Difference Relations in $\mathbb{S}^3$)**.** *Let $\mathcal{A}$ be a simple arrangement of $n \geq 3$ great-spheres in $\mathbb{S}^3$. Then*

$$E^1_k(\mathcal{A}) = \tfrac{3}{2}[E^0_k(\mathcal{A}) - E^0_{k-1}(\mathcal{A})], \quad \text{for } 0 \leq k \leq n, \tag{3.15}$$

$$E^2_k(\mathcal{A}) = \tfrac{1}{3}[E^1_k(\mathcal{A}) - E^1_{k-1}(\mathcal{A})] + 2, \quad \text{for } 0 \leq k \leq n, \tag{3.16}$$

$$E^3_k(\mathcal{A}) = \begin{cases} 1 & \text{for } k = 0, \\ 0 & \text{for } 1 \leq k \leq n-1, \\ -1 & \text{for } k = n. \end{cases} \tag{3.17}$$

Proof. We get (3.15) by setting $d = 3$ and $p = 1$ in (3.8) and (3.17) by setting $d = p = 3$ in (3.7). To get (3.16), we begin by setting $d = p = 3$ in (3.8), which gives

$$E^2_k - E^2_{k-1} = [E^0_{k-3} - E^0_k] + [E^1_{k-2} + E^1_k] + 2E^3_k = \tfrac{1}{3}[E^1_k - 2E^1_{k-1} + E^1_{k-2}] + 2E^3_k, \tag{3.18}$$

in which we use $E^0_{k-3} - E^0_k = -\frac{2}{3}[E^1_k + E^1_{k-1} + E^1_{k-2}]$ implied by (3.15). Moving E^2_{k-1} to the right-hand side and substituting it recursively implies (3.16) because $\sum_{\ell=0}^{k} E^3_\ell = 1$ by (3.17). □

If every 2-dimensional subarrangement is bipolar, then each arrangement of $n-1$ great-circles inside a great-sphere has polygons of predictable depth characteristics, namely a minimum (with depth characteristic 1) at depth 0, a maximum (with depth characteristic 1) at depth $n-1$, and otherwise only non-critical polygons connecting the minimum to the maximum. Hence,

$$E^2_k(\mathcal{A}) = \begin{cases} n & \text{for } k = 0, n-1, \\ 0 & \text{for } 1 \le k \le n-2, \end{cases} \tag{3.19}$$

$$E^1_k(\mathcal{A}) = 3n - 6(k+1) \quad \text{for } 0 \le k \le n-2, \tag{3.20}$$

$$E^0_k(\mathcal{A}) = 2n(k+1) - 4\tbinom{k+2}{2} \quad \text{for } 0 \le k \le n-3, \tag{3.21}$$

in which we get (3.20) from (3.19) and (3.16), and we get (3.21) from (3.20) and (3.15). For values of k outside the given limits, the sums of Euler characteristics are zero.

As defined in [97], the *order-k Voronoi tessellation* of n points in $\mathbb{R}^2$ is a decomposition of the plane into closed convex regions such that any two points in a region share the same k nearest points in the given set; but see also [54]. It can be obtained by mapping each of the n points, $u = (u_1, u_2)$, to the plane $x_3 = u_1 x_1 + u_2 x_2 + \frac{1}{2}(u_1^2 + u_2^2)$, forming the arrangement of the n planes, and projecting the chambers at depth k to the regions of the tessellation. The boundaries of the regions are obtained by projecting the edges at depth $k-1$ and the vertices at depths $k-2$ and $k-1$. In 1982, Der-Tsai Lee counted the regions, edges, and vertices in these tessellations [73], and found that the numbers depend on n and k but barely on how the points are placed in the plane. Indeed, if we modify the setting slightly by turning the planes into great-spheres—as explained in Section 2.4—then a general position assumption suffices for these numbers to depend solely on n and k. Using Theorem 3.4.2 and the expressions for E^p_ℓ in the case of bipolar 2-dimensional subarrangements given in (3.21), (3.20), (3.19), (3.17), we get

$$C^0_{k-2} + C^0_{k-1} = E^0_{k-2} + E^0_{k-1} = 2(n-k)(2k-1) - 2k, \tag{3.22}$$

$$C^1_{k-1} = 3E^0_{k-2} + E^1_k = 3(n-k)(2k-1) - 3k, \tag{3.23}$$

$$C^3_k = E^0_{k-3} + E^0_{k-2} + E^0_{k-1} + E^0_k = (n-k)(2k-1) - k + 2 \tag{3.24}$$

for the number of vertices, edges, and regions. Modulo the difference between $\mathbb{R}^2$ and $\mathbb{S}^2$, these are the same expressions as in [73].

3.5.2 Three Dimensions

Before discussing the 3-dimensional order-k Voronoi tessellations, we introduce the 4-dimensional difference relations implied by Theorems 3.4.1 and 3.4.2.

Corollary 3.5.2 (Difference Relations in $\mathbb{S}^4$)**.** *Let $\mathcal{A}$ be a simple arrangement of $n \geq 4$ great-spheres in $\mathbb{S}^4$. Then*

$$E^1_k(\mathcal{A}) = 2[E^0_k(\mathcal{A}) - E^0_{k-1}(\mathcal{A})], \quad \text{for } 0 \leq k \leq n, \tag{3.25}$$

$$E^2_k(\mathcal{A}) = \tfrac{1}{2}[E^1_k(\mathcal{A}) - E^1_{k-1}(\mathcal{A})] + \textstyle\sum_{\ell=0}^{k} E^3_\ell, \quad \text{for } 0 \leq k \leq n, \tag{3.26}$$

$$E^4_k(\mathcal{A}) = \begin{cases} 1 & \text{for } k = 0, n, \\ 0 & \text{for } 1 \leq k \leq n-1. \end{cases} \tag{3.27}$$

Proof. We get (3.25) by setting $d = 4$ and $p = 1$ in (3.8), and we get (3.27) by setting $d = p = 4$ in (3.7). To get (3.26), we begin by setting $d = 4$ and $p = 3$ in (3.8), which gives

$$E^2_k - E^2_{k-1} = 2[E^0_{k-3} - E^0_k] + \tfrac{3}{2}[E^1_{k-2} + E^1_k] = \tfrac{1}{2}[E^1_k - 2E^1_{k-1} + E^1_{k-2}] + E^3_k, \tag{3.28}$$

in which we use $2[E^0_{k-3} - E^0_k] = -[E^1_k + E^1_{k-1} + E^1_{k-2}]$ implied by (3.25). Moving E^2_{k-1} to the right-hand side and substituting iteratively, we get (3.26). □

Note the absence of any relation for E^3_k. However, if we assume that all 3-dimensional subarrangements are bipolar, there is additional information about the facets and therefore also about the polygons:

$$E^3_k(\mathcal{A}) = \begin{cases} n & \text{for } k = 0, \\ 0 & \text{for } 1 \leq k \leq n-2, \\ -n & \text{for } k = n-1, \end{cases} \tag{3.29}$$

$$E^2_k(\mathcal{A}) = \tfrac{1}{2}[E^1_k(\mathcal{A}) - E^1_{k-1}(\mathcal{A})] + n, \quad \text{for } 0 \leq k \leq n-2, \tag{3.30}$$

in which we get (3.30) from (3.29) and (3.26).

By straightforward generalization from 2 to 3 dimensions, the *order-k Voronoi tessellation* of n points in $\mathbb{R}^3$ decomposes space into convex regions, each associated with the k nearest of the n points. In analogy to the 2-dimensional case, we map the points to 3-planes in $\mathbb{R}^4$—or to great-spheres in $\mathbb{S}^4$—so that the tessellation is the projection of a subset of the cells. Despite this similarity, the expressions for the number of cells of the 2-dimensional tessellations derived by Lee in 1982 [73] have been extended to 3 dimensions only recently. The main reason for such delay is that the number of cells do not only depend on n and k, but also on how the points are distributed in space. Indeed, compared to the 2-dimensional case, we have the same number of relations but one more variable. Specifically, we have relations (3.25), (3.30), (3.29), (3.27), and we count vertices, edges, polygons, and (3-dimensional) regions, which are obtained by projecting the $C^0_{k-1} + C^0_{k-2} + C^0_{k-3}$ vertices at depths $k-1, k-2, k-3$, the $C^1_{k-1} + C^1_{k-2}$ edges at depths $k-1, k-2$, the C^2_{k-1} polygons at depth $k-1$, and the C^4_k chambers at depth k. Using Theorem 3.4.2 and the four mentioned relations for bipolar 3-dimensional subarrangements, we get

$$E^0_{k-3} + E^0_{k-2} + E^0_{k-1} = \mathbf{E}^2_{k-3} + \mathbf{E}^2_{k-2} + \mathbf{E}^2_{k-1} - \tfrac{n}{2}[3k^2 - 3k + 2], \tag{3.31}$$

$$4E^0_{k-2} - E^1_{k-2} + 4E^0_{k-1} - E^1_{k-1} = 2\mathbf{E}^2_{k-2} + 4\mathbf{E}^2_{k-1} + 2\mathbf{E}^2_k - 2n[2k^2 - 2k + 1], \tag{3.32}$$

$$6E^0_{k-1} - 3E^1_{k-1} + E^2_{k-1} = \mathbf{E}^2_{k-2} + E^2_{k-1} - 3n[k^2 - k], \tag{3.33}$$

$$E^0_k - E^1_k + E^2_k - E^3_k + E^4_k = \mathbf{E}^2_{k-2} - \tfrac{n}{2}[k^2 - k + 2] \tag{3.34}$$

for the number of vertices, edges, polygons, and regions in the order-k Voronoi tessellation for $1 \leq k \leq n-1$, in which $\mathbf{E}^2_k = \sum_{m=0}^{k} \sum_{\ell=0}^{m} E^2_\ell$. To see that these are the same expressions as in [10], we note that $\mathbf{E}^2_k = N_{k+1}$ and $E^2_k = J_{k+1}$ in the notation of that paper.

3.6 Application to Neighborly Arrangements

Recall that an arrangement in $\mathbb{S}^d$ is neighborly if the great-spheres are dual to the vertices of a neighborly polytope. Equivalently, all subarrangements of dimension $p \geq d/2$ have bi-polar depth functions. We generalize the face-counting formulas for neighborly polytopes to the levels in neighborly arrangements. In particular, we show that the number of p-cells at depth k is a function of n, d, p, and k alone. For the special case of cyclic polytopes, this was proved before by Andrezejak and Welzl [3, Theorem 5.1], who also derived explicit formulas for the number of cells.

3.6.1 Equations in Matrix Form

We write $d = 2t - 1$ for odd d and $d = 2t$ for even d. Let $\mathcal{A}$ be a neighborly arrangement of n great-spheres in $\mathbb{S}^d$, so all subarrangements of dimension $t \leq p \leq d$ are bi-polar. By Theorem 3.4.1, the E_k^p are simple functions in n, d, p, and k, for all $t \leq p \leq d$. In addition, we get t independent relations for every k from Theorem 3.4.2. Specifically, for every odd p between 0 and d, we get a relation by equating the left-hand side of (3.1) with the right-hand side of (3.1). This gives what we call a *giant linear system* with variables E_k^0 to E_k^{t-1} for $0 \leq k \leq n$. To describe it, we introduce the $t \times t$ matrices M_d. For odd d, it is a straightforward configuration of binomial coefficients, which is however interrupted by -2s replacing $-\binom{2t-j}{2i-2} = -1$ in row i and column j whenever $2t - j = 2i - 2$:

$$M_{2t-1} = \begin{bmatrix} \binom{2t-1}{0} & -\binom{2t-2}{0} & \binom{2t-3}{0} & -\binom{2t-4}{0} & \dots & \pm\binom{t}{0} \\ \binom{2t-1}{2} & -\binom{2t-2}{2} & \binom{2t-3}{2} & -\binom{2t-4}{0} & \dots & \pm\binom{t}{2} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \binom{2t-1}{2t-4} & -\binom{2t-2}{2t-4} & \binom{2t-3}{2t-4} & -2 & \dots & 0 \\ \binom{2t-1}{2t-2} & -2 & 0 & 0 & \dots & 0 \end{bmatrix}. \tag{3.35}$$

These replacements will be important shortly. For even d, the matrix M_{2t} has the same number of entries, with $\binom{2t-j+1}{2i-1}$ in row i and column j replacing $\binom{2t-j}{2i-2}$ in M_{2t-1}. The -2s and 0s are the same in both matrices. In d dimensions, the giant system is given by a $t(n+1) \times t(n+1)$ matrix, with $n+1$ copies of M_d along the diagonal. All entries to the lower left of this diagonal of $t \times t$ blocks are zero, while there are sporadic non-zero entries to the upper right.

Lemma 3.6.1 (Invertible Blocks Imply Invertible Systems)**.** *For every $d \geq 1$, if M_d is invertible, then the giant system of linear relations in d dimensions is invertible.*

Proof. If M_d is invertible, then we can use row and column operations to turn M_d into an upper triangular matrix with non-zero entries along the diagonal. Applying the same operations to the giant matrix, we get a giant upper triangular matrix with non-zero entries along the entire diagonal. □

3.6.2 Everything Modulo 2

We prove the invertibility of M_{2t-1} by proving that its determinant is odd. Equivalently, we write P_{2t-1} for the matrix M_{2t-1} in which every entry is replaced by its parity, and we show that the mod 2 determinant of P_{2t-1} is 1. Before doing so, we show that the invertibility of

M_{2t-1} implies the invertibility of M_{2t}. Let N_{2t} be the matrix M_{2t} after dividing each column by the largest power of 2 that divides all its entries, and write P_{2t} for the matrix N_{2t} in which every entry is replaced by its parity.

Lemma 3.6.2 (Odd Imply Even Invertible Blocks)**.** $P_{2t} = P_{2t-1}$.

Proof. Recall that the entry in row i and column j is $\binom{2t-j}{2i-2}$ in M_{2t-1} and $\binom{2t-j+1}{2i-1}$ in M_{2t}, unless this entry is -2 or 0, in which case it is the same in the two matrices. Assuming the former case, the ratio of the two entries is $\binom{2t-j+1}{2i-1}/\binom{2t-j}{2i-2} = (2t-j+1)/(2i-1)$. Since $2i-1$ is odd, the largest power of 2 that divides $\binom{2t-j+1}{2i-1}$ is the largest power of 2 that divides $\binom{2t-j}{2i-2}$ times the largest power of 2 that divides $2t-j+1$. The latter is the same for all entries in a column. We thus divide column j in M_{2t} by the largest power of 2 that divides $2t-j+1$, which is 1 for all even j. The even columns of M_{2t} are the ones that contain the -2s, so after dividing, the parities of corresponding terms in M_{2t} and M_{2t-1} are the same. Equivalently, $P_{2t} = P_{2t-1}$. □

Henceforth, we focus on the odd case. We use a consequence of Kummer's Theorem [71] to get the parity version of M_{2t-1}:

Lemma 3.6.3 (Odd Binomial Coefficients)**.** *For all* $0 \le k \le n$, $\binom{n}{k}$ *is odd iff the binary representations of* n, k, *and* $n-k$ *satisfy* $n_2 = k_2 \ \ xor \ \ (n-k)_2$.

In words: the 1s in the binary representations of k and $n-k$ are at disjoint positions. It follows that the positions of the 1s in the binary representation of k are a subset of the positions of the 1s in the binary representation of n, and similarly for $n-k$ and n. A compelling visualization of Lemma 3.6.3 is the Pascal triangle in binary, whose 1s form the Sierpinski gasket as shown in Figure 3.2. To transform the Sierpinski gasket into a matrix that contains P_{2t-1}, for every

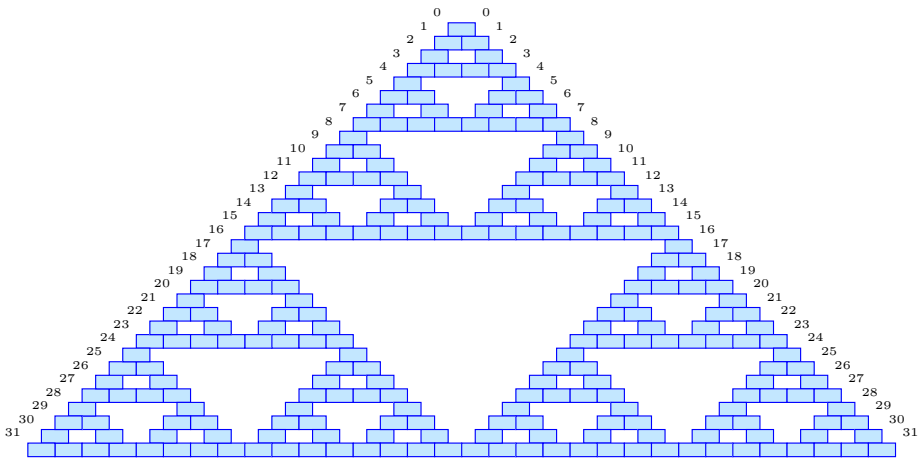

Figure 3.2: The Pascal triangle in modulo 2: the *blue* bricks are odd entries, and the *white* bricks (not shown) are even entries.

$t \ge 1$, we drop every other up-slope (whose label, given along the down-slope in Figure 3.2, is odd), we draw the remaining up-slopes as rows, and we draw the horizontal lines in the gasket as columns. Finally, we convert the last 1 in each row to a 0. These are the binomial coefficients that change from -1 to -2 in M_{2t-1}; see Figure 3.3.

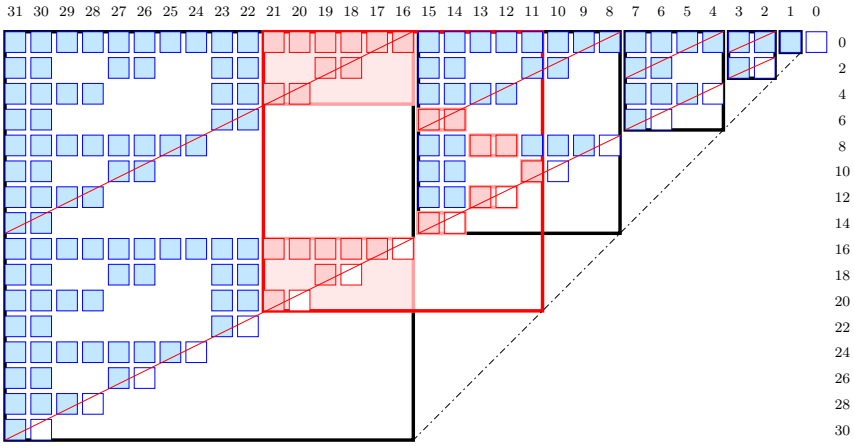

Figure 3.3: Each *blue* and *pink* square is a 1 in the matrix, and each *white* square is a 0 (only those originally equal to -2 are shown). The *bold black* frames mark the exponential blocks, the *bold red* frame marks the 11-th block, P_{21}, and the *pink* boxes inside the *red* frame mark the tops and bottoms of the NE- and SW-incursions that arise in its reduction.

3.6.3 Reducing Exponential Blocks

Observe that P_{2t-1} is the submatrix consisting of the rows labeled $2i$, for $0 \leq i \leq t-1$, and the columns labeled j, for $t \leq j \leq 2t-1$; see Figure 3.3. We call this the t*-th block*. For the time being, we focus on *exponential blocks*, for which t is a power of 2. Note the symmetry between the upper and lower halves of an exponential block: the bottom is a copy of the top, except that the last 1 in each row is turned into a 0. We use this property to reduce exponential blocks.

Reduction 3.6.1 (Exponential Block)**.** *Let P_{2t-1} be an exponential block, with $t = 2^n$, and write $s = 2^{n-1}$. We reduce P_{2t-1} in three steps:*

1. *For $0 \leq i \leq s-1$, add the row with label $2i + 2s$ to the row with label $2i$. Thereafter, we have a 1 in each row and each even column, and otherwise only 0s in the upper half of the exponential block.*
2. *Zero out the even columns in the lower half using the rows in the upper half. After consolidating the lower half by removing the even columns, which are all zero, we get an upper triangular matrix with 1s in the diagonal.*
3. *Reduce this upper triangular matrix to the $s \times s$ identity matrix. Adding the even columns back, we have a 1 in each row and each odd column, and otherwise only 0s in the lower half of the exponential block.*

Assuming $t = 2^n$, the above reduction algorithm turns P_{2t-1} into a $t \times t$ permutation matrix, whose determinant is of course 1. This is the parity of the determinant of M_{2t-1}, which is therefore non-zero. To extend this result to integers, t, that are not necessarily powers of 2, we need a few properties of an exponential block. Being a square matrix with $t = 2^n$ rows and columns, it decomposes into four quarters of $s = 2^{n-1}$ rows and columns each. By combining the NE- and NW-quarters, we get the *northern half* of the exponential block, and we draw the line from its bottom-left to top-right corners, calling it the *northern diagonal*; see Figure 3.3. Similarly, we merge the SE- and SW-quarters to get the *southern half* and draw the *southern*

diagonal from the bottom-left to top-right corner. Note that the southern half of P_{2t-1} is a copy of everything to the right of the northern half, namely the exponential blocks of size $1, 2, 4, \ldots, 2^{n-1}$ plus the 0s below and to the right of them.

An *NE-incursion* is a submatrix whose bottom-left corner lies on the southern diagonal and whose top-right corner is the top-right corner of the exponential block. As an example consider the rows labeled 0 to 20 and columns labeled 21 to 16, which is an NE-incursion of P_{31} in Figure 3.3. We decompose the NE-incursion into three rectangular matrices stacked on top of each other: the *top*, the *middle*, and the *bottom*, in which the top and bottom are twice as wide as they are high, and the middle fills the space in between. Importantly, the middle is zero, and the top and bottom combine to a square matrix whose structure is such that Reduction 3.6.1 can reduce it to the identity matrix.

Symmetrically, an *SW-incursion* is a submatrix whose top-right corner lies on the northern diagonal and whose bottom-left corner is the bottom-left corner of the exponential block. As an example consider the rows labeled 6 to 14 and columns labeled 15 to 14, which is an SW-incursion of P_{15} in Figure 3.3. As before, we decompose the SW-incursion into three rectangular matrices, in which the *top* and *bottom* are twice as wide as they are high, and the *middle* consists of the remaining rows in between. The top and bottom combine again to a square matrix that can be reduced to the identity matrix by Reduction 3.6.1. However, the middle is not necessarily zero. On the other hand, all entries to the right of the top but still within the exponential block are zero.

3.6.4 Reducing General Blocks

We thus have the necessary ingredients to reduce a not necessarily exponential block, P_{2t-1}. Assuming t is not a power of 2, let u be the power of 2 such that $u/2 < t < u$, and write $s = u/2$. The overlap of P_{2t-1} with P_{2u-1} is an NE-incursion of the latter.

Reduction 3.6.2 (NE-incursion)**.** *Let I be the overlap of P_{2t-1} and P_{2u-1}. We reduce I and zero out portions of P_{2t-1} outside I:*

1. *Combine the top and bottom of I and reduce it using Reduction 3.6.1.*
2. *Add back the middle, which we recall is 0.*
3. *Use the columns of the reduced I to zero out the rectangular regions of P_{2t-1} to the right of the top and bottom of I.*

Step 1 may contaminate the regions to the right of the bottom of I with non-zero entries, but Step 3 cleans up the contamination at the end. We are thus left with an un-reduced submatrix of size $(u-t) \times (u-t)$, which we denote P'_{2t-1}. It is a bottom-left submatrix but not necessarily an SW-incursion of P_{2s-1}. Assuming $s < 2(u-t)$, there is a largest SW-incursion of P_{2s-1} contained in P'_{2t-1}, which has the same number of rows as P'_{2t-1}.

Reduction 3.6.3 (SW-incursion)**.** *Assume $s < 2(u-t)$ and let J be the largest SW-incursion of P_{2s-1} contained in P'_{2t-1}. We reduce J as follows:*

1. *Combine the top and bottom of J and reduce it using Reduction 3.6.1.*

2. Add back the middle and zero it out using row operations.

We note that the regions of P'_{2t-1} to the right of the top and bottom of J are zero because J is an SW-incursion, and P'_{2t-1} is contained in P_{2s-1}. Step 1 preserves this property, so Step 2 can zero out the middle without contaminating the remaining un-reduced matrix of size $(s-u+t)\times(s-u+t)$, which we denote P''_{2t-1}.

It is also possible that $s \geq 2(u-t)$, in which case there is no non-empty SW-incursion of P_{2s-1} contained in P'_{2t-1}. We thus substitute the SW-quarter of P_{2s-1} for P_{2s-1}, or the SW-quarter of that SW-quarter, etc. This square matrix is a copy of the exponential block of the same size, so Reduction 3.6.3 still applies. Similarly, P''_{2t-1} is a copy of the $(s-u+t)$-th block. Since $s-u+t<t$, we can reduce it by induction. The correctness of the reduction algorithms implies

Lemma 3.6.4 (Blocks are Invertible)**.** *For every $d \geq 1$, M_d is invertible.*

Proof. For $d=2t-1$, Reductions 3.6.1, 3.6.2, 3.6.3 together with induction imply that P_{2t-1} can be reduced to the identity matrix. By Lemma 3.6.2 this is also the case for P_{2t}. Since P_d is the parity version of M_d, this implies that M_d is invertible. □

3.6.5 Number of Cells

The invertibility of the blocks implies the invertibility of the giant linear systems, which implies that the number of cells in the levels of neighborly arrangements are independent of the geometry of the great-spheres defining the arrangement.

Theorem 3.6.1 (Neighborly Arrangements)**.** *Let $\mathcal{A}$ be a neighborly arrangement of $n \geq d$ great-spheres in $\mathbb{S}^d$. Then the $E^p_k(\mathcal{A})$ and the $C^p_k(\mathcal{A})$ are functions of n, d, p, and k.*

Proof. By Lemma 3.6.4, the matrix M_d is invertible, which by Lemma 3.6.1 implies that the giant linear system created from Theorems 3.4.1 and 3.4.2 is invertible. Hence, the $E^p_k(\mathcal{A})$ of the d-dimensional arrangement are determined; that is: they are functions of n, d, p, and k, but not of the great-spheres defining the arrangement. By Theorem 3.4.2, the $C^p_k(\mathcal{A})$ are determined by the $E^p_k(\mathcal{A})$, so they are also functions of n, d, p, and k. □

As an example, consider a neighborly arrangement of n great-spheres in $\mathbb{S}^4$. All subarrangements of dimension 2, 3, and 4 have bi-polar depth functions, so we get the E^p_k for $p=2,3,4$ from Theorem 3.4.1, and we use Theorem 3.4.2 to get them for $p=0,1$:

$$E^0_k = \tfrac{1}{2}(k+1)n(n-k-3) \qquad \text{for} \qquad 0 \leq k \leq n-4, \tag{3.36}$$

$$E^1_k = n(n-2k-3) \qquad \text{for} \qquad 0 \leq k \leq n-3, \tag{3.37}$$

$$E^2_k = \tbinom{n}{2},\ 0,\ \tbinom{n}{2} \qquad \text{for } k=0,\ 1 \leq k \leq n-3,\ k=n-2, \tag{3.38}$$

$$E^3_k = n,\ 0,\ -n \qquad \text{for } k=0,\ 1 \leq k \leq n-2,\ k=n-1, \tag{3.39}$$

$$E^4_k = 1,\ 0,\ 1 \qquad \text{for } k=0,\ 1 \leq k \leq n-1,\ k=n. \tag{3.40}$$

Using the relations $C_k^0 = E_k^0$, $C_k^1 = 4E_k^0 - E_k^1$, etc., from Theorem 3.4.2, we get the number of cells with given depth:

$$C_k^0 = \tfrac{1}{2}(k+1)n(n-k-3) \qquad \text{for} \qquad 0 \le k \le n-4, \tag{3.41}$$

$$C_k^1 = n[n(2k+1) - 2k^2 - 6k - 3] \qquad \text{for} \qquad 0 \le k \le n-3, \tag{3.42}$$

$$C_k^2 = \tbinom{n}{2},\ 3nk(n-k-2),\ \tbinom{n}{2} \qquad \text{for } k=0,\ 1 \le k \le n-3,\ k = n-2, \tag{3.43}$$

$$C_k^3 = n,\ n[(2k-1)n - 2k^2 - 2k + 3],\ 6\tbinom{n}{2},\ 2\tbinom{n}{2},\ n \qquad \text{for } k=0,\ 1 \le k \le n-4,\ k=n-3,\ k=n-2,\ k=n-1, \tag{3.44}$$

$$C_k^4 = 1,\ \tfrac{1}{2}n[n(k-1) - k^2 + 3],\ n(n-3),\ \tbinom{n}{2},\ n,\ 1 \qquad \text{for } k=0,\ 1 \le k \le n-4,\ k=n-3, k=n-2,\ k=n-1,\ k=n. \tag{3.45}$$

3.7 Discussion and Open Problems

The main contribution of the paper that this chapter builds on [11] is the introduction of the discrete depth function as a topological framework to approach questions in discrete geometry, and the establishment of the system of Dehn–Sommerville–Euler relations for levels of this function. We have illustrated the use of this system by rederiving known cell-counting formulas for order-k Voronoi tessellations in $\mathbb{R}^2$ and $\mathbb{R}^3$, and by extending the classic face-counting formulas for neighborly polytopes to the levels in neighborly arrangements. This work suggests further research to deepen our understanding of the framework:

- Establish effective relations expressing the connections between the restrictions of the depth function to subarrangements.
- Relate the stability of the persistence diagrams of restrictions of the depth function to combinatorial questions in geometry.

While our framework sheds new light on well studied questions in discrete geometry, there is plenty of work that remains. The following questions are of particular interest:

- Give bounds on the topological quantities that arise in counting the regions of order-k Voronoi tessellations. As established in [10], the relevant quantity in $\mathbb{R}^3$ is the double sum of depth characteristics of the 2-dimensional cells (the polygons) in the corresponding arrangement of great-spheres in $\mathbb{S}^4$. How do these results extend beyond 3 dimensions?
- Generalize the results on neighborly arrangements to counting the k-sets of general sets of n points in $\mathbb{R}^d$. Specifically, use the framework of depth functions to improve the current best upper bounds on the maximum number of k-sets, which are $O(n^{4/3})$ in $\mathbb{R}^2$ [35], $O(n^{5/2})$ in $\mathbb{R}^3$ [99], and $O(n^{d-\varepsilon_d})$ for a small constant $\epsilon_d > 0$ in $\mathbb{R}^d$ [112].

CHAPTER 4

Geometric Networks

This chapter, which is based on [12, 27], draws its direct inspiration from the final open question posed at the conclusion of the previous section: can persistent homology of the depth map, defined for n points in $\mathbb{R}^2$, be leveraged to establish new bounds on the number of halving lines? Despite diligent efforts, we have yet to find a definitive answer. However, these endeavors have not been futile; they have sparked a more in-depth exploration into the persistence of functions on one-dimensional compact spaces, which we call *geometric networks*. This chapter delves into these discoveries and their subsequent algorithmic implications, which are discussed towards the chapter's end. Our exploration begins by defining the spaces under consideration and their *extended persistent homology.*

The core result in this chapter is a local characterization of the pairing of critical points in persistent homology for continuous functions defined on geometric networks. The characterization is formulated in terms of *windows*, each the product of a connected subset of the geometric network and the range of the function restricted to this subset.

4.1 Maps and Spaces

Let $f\colon \mathbb{S}^1 \to \mathbb{R}$ be a continuous map on the unit circle; see Figure 4.1. We call f *generic* if there is no non-empty open interval along which f is constant, so all critical points are isolated, and the values of these critical points are distinct. A *minimum* of such a map is a point $a \in \mathbb{S}^1$ for which there exists a neighborhood, $N(a) \subseteq \mathbb{S}^1$, such that $f(a) < f(x)$ for all $x \in N(a)$. Symmetrically, a *maximum* is a point $a \in \mathbb{S}^1$ such that $f(a) > f(x)$ for all $x \in N(a)$. The minima and maxima alternate in a trip around the circle, which implies that there are equally many of them. There is exactly one *global minimum*, a_0, and one *global maximum*, b_0, which satisfy $f(a_0) \leq f(x) \leq f(b_0)$ for all $x \in \mathbb{S}^1$. We note that the stability of the persistence diagram [21] makes the assumption of distinct critical values unnecessary, but we include it in the definition of genericity to simplify arguments at many places.

By a *geometric network* we mean the realization of an abstract graph in some Euclidean space: each vertex is mapped to a point, and each edge to a curve connecting the images of its vertices. We are not concerned with the details of the embedding, except that different vertices map to different points, and curves do not intersect except possibly at shared endpoints. For convenience, we restrict ourselves to finite graphs in which every vertex has degree 1 or 3. The constraint on the vertex degrees is not really a limitation since we can replace a degree-k

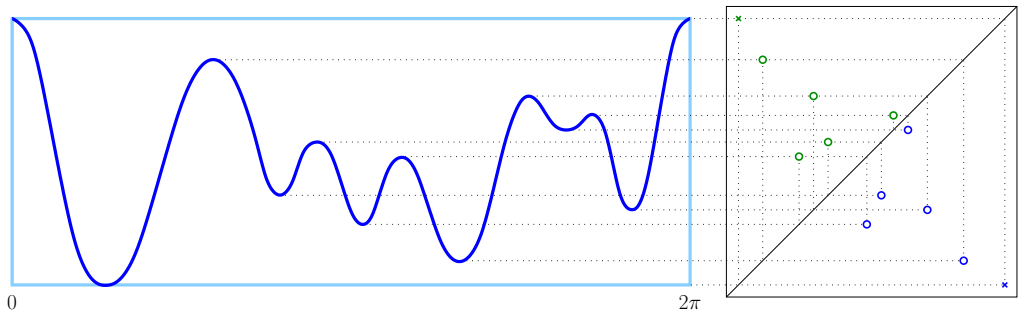

Figure 4.1: *Left:* the graph of a generic function on the circle with the global maximum at $0 = 2\pi$. The six minima alternate with the six maxima. *Right:* the persistence diagram of the map. The two points that correspond to the global min-max pair are marked by crosses, while all other points are marked by small circles.

vertex by a tree with $k-2$ vertices, all of degree 3, and if the edges in the tree approach zero length, we can recover the original topology in the limit. Similar substitutions can be used to model multi-edges and cycles. If a geometric network is connected and without cycle, we call it a *geometric tree*.

Besides minima and maxima, geometric networks contain two other types of critical points that play a central role in our analysis: the *endpoints* (degree-1 vertices) and *branching points* (degree-3 vertices). An endpoint has $\searrow$*-type* or $\nearrow$*-type* if its value is larger or smaller than the values of the points in a sufficiently small neighborhood, respectively. Similarly, we allow for two types of degree-3 vertices: *y-type* (one $\searrow$- and two $\nearrow$-type vertices glued to each other) and λ*-type* (one $\nearrow$- and two $\searrow$-type vertices glued to each other). We call the function $f\colon \mathbb{G} \to \mathbb{R}$ *generic* if all critical points are isolated and their values are distinct. A *critical value* of f is the value of a critical point; all other values are *non-critical*.

Call $f\colon \mathbb{G} \to \mathbb{R}$ *piecewise differentiable* if there is a decomposition of $\mathbb{G}$ into a finite number of curves such that the restriction of f to the interior of each curve is differentiable. For example, if $\mathbb{G}$ is a straight-line embedding of a finite graph, and f is the linear extension of its values at the vertices, then f is piecewise linear and therefore also piecewise differentiable. We define the *variation* of a piecewise differential map as the integral, over all interior points of the curves in the decomposition, of the absolute value of the derivative at these points.

We treat the circle separately and before considering more general geometric networks because it is the only connected 1-manifold among them.

4.2 Extended Persistent Homology

In a nutshell, persistent homology is the embodiment of the idea that features exist on many scale levels, and rather than preferring one scale over another, it quantifies the features in terms of the range of scales during which they appear. The concept of *persistent homology* has been introduced for components by Frosini and Landi [57] and for general homology groups by Robins [95] and independently by Edelsbrunner, Letscher, and Zomorodian [46]. The latter paper gives the first fast algorithm for persistence. A generalization of the notion of persistence to coefficient groups that are not fields can be found in [113], but will be not be treated in this thesis.

4.2.1 Homology with $\mathbb{Z}/2\mathbb{Z}$ Coefficients

To keep the algebraic discussion of homology as elementary as possible, we use modulo-2 arithmetic; that is: we construct homology groups with $\mathbb{Z}/2\mathbb{Z}$ coefficients. For compact 1-dimensional spaces, such groups are straightforward objects, so we can side-step the formal introduction of homology. For a more comprehensive treatment, we recommend standard texts in algebraic topology, for example Hatcher [66].

Given a map, $f\colon \mathbb{S}^1 \to \mathbb{R}$, the *sublevel set* at $t \in \mathbb{R}$ is $f_t = f^{-1}(-\infty, t]$, and the *superlevel set* is $f^t = f^{-1}[t, \infty)$. Let A_0 and B_0 be the values of the global minimum and the global maximum, respectively. For a non-critical value, t, we have the following three cases:

- $t < A_0$: $f_t = \emptyset$ and $f^t = \mathbb{S}^1$;
- $A_0 < t < B_0$: f_t consists of a positive number of connected components, each a closed arc with non-empty interior, and f^t consists of the same number of connected components of the same type;
- $t > B_0$: $f_t = \mathbb{S}^1$ and $f^t = \emptyset$.

We use *homology* to formally distinguish between these cases. In particular, the rank of $\mathsf{H}_0(f_t)$ is the number of connected components of the sublevel set, and the rank of $\mathsf{H}_1(f_t)$ is the number of cycles. Note that we have no cycle for $t < B_0$ and one cycle for $t > B_0$. Compare this with the *homology* of $\mathbb{S}^1$ *relative* to f^t, denoted $\mathsf{H}_i(\mathbb{S}^1, f^t)$, where we have $rank\,\mathsf{H}_0(\mathbb{S}^1, f^t) = rank\,\mathsf{H}_1(f_t) = 0$ for $t < B_0$ and $rank\,\mathsf{H}_0(\mathbb{S}^1, f^t) = rank\,\mathsf{H}_1(f_t) = 1$ for $t > B_0$. More interesting is the case $i = 1$, for which the relative homology group counts the open arcs in $\mathbb{S}^1 \setminus f^t$. By Lefschetz duality, the (absolute) homology groups and the relative homology groups are isomorphic: $\mathsf{H}_i(f_t) \simeq \mathsf{H}_{1-i}(\mathbb{S}^1, f^t)$, for $i = 0, 1$ and for all non-critical values, t of f. This is an elementary insight for the circle and is also true for higher-dimensional manifolds. It does not hold for more general spaces, not even for the unit interval. On the other hand, both homology and relative homology generalize and can be used to count connected components and cycles in geometric networks and the sub- and superlevel sets of maps on them.

4.2.2 Persistent Homology

Persistent homology arises when we keep track of sub- and superlevel sets while t changes continuously. We again take advantage of the relative simplicity provided by the restriction to compact 1-dimensional spaces and avoid the introduction of the concept in full generality. For more comprehensive background, we refer to [43]. Specifically, we use the framework that is referred to as *extended persistent homology*, which is constructed in two phases, first growing the sublevel set until it exhausts the space, and second doing the same with the superlevel set. We explain this for a generic map on the unit circle.

In *Phase One*, we increase t from $-\infty$ to ∞ and use $\mathsf{H}_0(f_t)$ and $\mathsf{H}_1(f_t)$ to do the book-keeping. A connected component is *born* when t passes the value of a minimum, and the component *dies* merging into another, older component when t passes the value of a maximum. There is one exception: when t passes B_0, then no component dies and instead a cycle is born. We pair

up the minimum, a, and the maximum, b, responsible for the birth and death of a component and represent the two events by the point $(f(a), f(b))$ in the plane.

In *Phase Two*, we decrease t from ∞ to $-\infty$ and use $\mathsf{H}_0(\mathbb{S}^1, f^t)$ and $\mathsf{H}_1(\mathbb{S}^1, f^t)$ to do the book-keeping. We enter Phase Two with a component born at $A_0 = f(a_0)$ and a cycle born at $B_0 = f(b_0)$, both of which did not yet die. The component dies in relative homology right at the beginning of Phase Two, when t passes B_0 (from the top going down), while the cycle lasts until the end, and dies when t passes A_0. This gives two pairs represented by the points (A_0, B_0) and (B_0, A_0). During Phase Two, a (relative 1-dimensional) cycle is born when t passes the value of a (non-global) maximum, and this cycle dies when t passes the value of a (non-global) minimum. Like in Phase One, we pair up the maximum, b, with the minimum, a, responsible for the birth and death of the cycle and represent the two events by the point $(f(b), f(a))$ in the plane.

4.2.3 Persistence Diagrams

The events during the two phases are recorded in the *persistence diagram* of f, denoted $\mathsf{D}gm(f)$, which is a multi-set of points, each marking the birth and death of a component or cycle; see Figure 4.1. We distinguish between three disjoint subdiagrams, $\mathsf{D}gm(f) = \mathsf{O}rd(f) \sqcup \mathsf{R}el(f) \sqcup \mathsf{E}ss(f)$, in which the *ordinary subdiagram* records the pairs in Phase One, the *relative subdiagram* records the pairs in Phase Two, and the *essential subdiagram* records the pairs that straddle the two phases. Whenever convenient, we list the dimension as a subscript, writing $\mathsf{D}gm_i(f)$ for the points that represent i-dimensional homology classes, and similarly for the subdiagrams. For a 1-dimensional map, we have $\mathsf{O}rd(f) = \mathsf{O}rd_0(f)$, $\mathsf{R}el(f) = \mathsf{R}el_1(f)$, but $\mathsf{E}ss(f) = \mathsf{E}ss_0(f) \sqcup \mathsf{E}ss_1(f)$. Recall that for a map on the unit circle, Lefschetz duality implies that the pairs in Phase One are the same as in Phase Two, only reversed. Similarly, for every pair straddling the two phases, there is also the reversed pair straddling the two phases. This implies that $\mathsf{D}gm(f)$ is symmetric across the main diagonal, with the caveat that a point $(f(a), f(b)) \in \mathsf{D}gm_i(f)$ maps to the point $(f(b), f(a)) \in \mathsf{D}gm_{1-i}(f)$; see Figure 4.1 and [22] for details. This property no longer holds for maps on non-manifold spaces, such as the unit interval, geometric trees, and general geometric networks. Nevertheless, the persistence diagram and its subdiagrams are useful book-keeping tools for the homology of the sub- and superlevel sets of maps on such more general spaces. Specifically, the ordinary subdiagram records the components of the sublevel set that are born and die during Phase One. The essential subdiagram records the homology of the geometric network, since its classes are born but do not die during Phase One. Finally, the relative subdiagram records the relative cycles in the network modulo the superlevel set.

For a point $(A, B) \in \mathsf{D}gm(f)$, we think of $|B - A|$ as the life-time or *persistence* of the corresponding component or cycle. Taking the sum, over all points in the multi-set, we get what we call the *total persistence* of f:

$$\|\mathsf{D}gm(f)\|_1 = \sum_{(A,B)\in \mathsf{D}gm(f)} |B - A|. \tag{4.1}$$

For a map on the unit circle, the global minimum and the global maximum contribute $2|B_0 - A_0|$ to this measure. Everything beyond that is due to wrinkles in the map and may be regarded as a measure of how interesting or noisy the map is.

An important property of persistence diagrams is their stability, which was first proved in [21]. Assuming f and g are generic maps on the same compact geometric network, this theorem

asserts that the bottleneck distance between $\mathrm{D}gm(f)$ and $\mathrm{D}gm(g)$ is bounded from above by $\|f-g\|_\infty$. It follows that every continuous map has a generic perturbation whose total persistence is arbitrarily close to that of the original map. It also implies that the restriction to maps whose critical points have distinct values is unnecessary while convenient.

4.3 The Circle Case

We consider generic maps on the unit circle and introduce the notion of a window to characterize the critical points paired by persistent homology. After establishing this connection, we get elementary proofs of fundamental properties of maps on the circle.

Let a be a minimum and b a maximum of a generic map, $f\colon \mathbb{S}^1 \to \mathbb{R}$, write $A = f(a)$, $B = f(b)$, and let $J = J(a,b)$ be the connected component of $f^{-1}[A,B]$ that contains both a and b. It may be a closed interval, the entire circle, or empty if no such component exists. We call $W(a,b) = J \times [A,B]$ the *frame* with *support* J *spanned* by a and b, and we say $W(a,b)$ *covers* the points $x \in J$. When J is an interval, a and b decompose it into three (closed) subintervals, which we read in a direction so that a precedes b: J_{in} before a, J_{mid} between a and b, and J_{out} after b. Correspondingly, we call $J_{\text{in}} \times [A,B]$, $J_{\text{mid}} \times [A,B]$, and $J_{out} \times [A,B]$ the *in-*, *mid-*, and *out-panels* of $W(a,b)$. We orient the in- and mid-panels away from the minimum, while we leave the out-panel without orientation; see Figure 4.2.

Definition 4.3.1 (Windows for Circles)**.** *Let a be a (non-global) minimum and b a (non-global) maximum, and assume that $J(a,b)$ is non-empty. We call $W(a,b)$, a* triple-panel window with simple wave *if the values at the endpoints of J_{in}, J_{mid}, J_{out} are B, A, B, A in this sequence.*

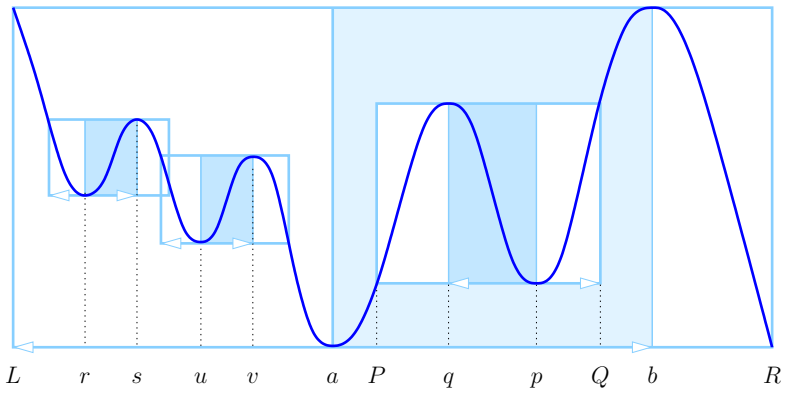

Figure 4.2: The triple-panel window with simple wave spanned by a and b. There are two children in the in-panel, spanned by r,s and u,v, there is one child in mid-panel, spanned by q,p, and there is no child in the out-panel. The triple-panel windows spanned by r,s and u,v overlap, while the corresponding double-panel windows are disjoint.

We will sometimes consider a *double-panel window*, which consists of the in-panel and the mid-panel. It contains the graph of the component in the sublevel set that grows from the minimum until it merges with another component at the corresponding maximum. We show that the windows with wave characterize the paired critical points, while noting that the global min-max pair is special and not subject to the following claim.

Theorem 4.3.1 (Characterization for Circles)**.** *Let $f\colon \mathbb{S}^1 \to \mathbb{R}$ be generic, let a be a (non-global) minimum with $f(a) = A$, and b a (non-global) maximum with $f(b) = B$. Then*

(A, B) and (B, A) are points in the ordinary and relative subdiagrams of $Dgm(f)$ iff the frame spanned by a and b is a triple-panel window with simple wave.

Proof. "$\Longleftarrow$". Let a, b span $W(a, b) = [L, R] \times [A, B]$, and assume that a is to the left of b, as in Figure 4.2. Consider the component of f_t that contains a as t increases from $-\infty$ to ∞. This component is born at $t = A$. Since $A \leq f(x) \leq B$ for all $L \leq x \leq b$, the component grows—occasionally by incorporating other, younger components—but never dies before t reaches B. At $t = B$, the component meets another component at b, and since $W(a, b)$ is a triple-panel window with wave, this other component is older. It follows that a, b are paired.

"$\Longrightarrow$". We suppose that a, b are paired. In other words, a component of f_t is born at $t = A$, and a remains the point with minimum value in this component until $t = B$, when the component merges with another, older component. Let $[L, b]$ and $[b, X]$ be the components right before merging. The graph of f restricted to $[L, b]$ describes the history of the component born at $t = A$, which implies that it is contained in $[L, b] \times [A, B]$. The other component is born earlier, so $[b, X]$ contains points that have the same value as a. Let R be the leftmost such point. By construction, the graph of f restricted to $[L, R]$ is contained in $[L, R] \times [A, B]$, which implies that $W(a, b)$ is a triple-panel window with simple wave. □

In addition to the points in the ordinary and relative subdiagrams—which are characterized by Theorem 4.3.1—D$gm(f)$ contains two more points, namely (A_0, B_0) and (B_0, A_0) in the essential subdiagram. With $A_0 < B_0$ the values of the global minimum and the global maximum, the first point represents the component and the second the cycle of the circle. There is no ambiguity which critical points of f are paired in persistent homology. Theorem 4.3.1 thus implies that for every minimum there is a unique maximum such that the corresponding frame is a window.

4.3.1 Nesting of Windows

As illustrated in Figure 4.2, two windows can be *nested* (one is a subset of the other), they can be *disjoint*, and they can *overlap*. We will see that any overlap is limited. We call $W(u, v)$ a *child* of $W(a, b)$, and $W(a, b)$ a *parent* of $W(u, v)$, if $W(u, v)$ is nested inside one of the panels of $W(a, b)$, and there is no other window nested between the two.

Lemma 4.3.1 (Nesting in Circle)**.** *Let $f : \mathbb{S}^1 \to \mathbb{R}$ be generic, and let $W(a, b)$ be a triple-panel window with simple wave and supports J_{in}, J_{mid}, J_{out} of its panels. If $W(u, v)$ is another triple-panel window and $u \in J_{in}, J_{mid}, J_{out}$, then $W(u, v)$ is nested inside the corresponding panel of $W(a, b)$.*

Proof. We first consider the mid-panel of $W(a, b)$, which we assume is oriented from left to right, so $a < b$. Moving from $x = a$ to $x = b$, we encounter an alternating sequence of minima and maxima, starting with a and ending with b. If a and b are the only critical points in this sequence, then the statement is vacuously true. Otherwise, let $a < p < b$ be the minimum with the smallest value, $f(p)$. There is at least one maximum to its left, and we let $a < q < p$ be the maximum with the largest value, $f(q)$; see Figure 4.2. Drawing a horizontal line from $(p, f(p))$ to the left, we intersect the graph of f in $(P, f(p))$, and drawing a horizontal line from $(q, f(q))$ to the right, we intersect the graph in $(Q, f(q))$. By construction, $a < P < q < p < Q < b$ as well as $f(p) \leq f(x) \leq f(q)$ for all $P \leq x \leq Q$. Hence, $W(p, q)$ is a triple-panel window with simple wave nested inside the mid-panel of $W(a, b)$. To continue,

we subdivide $[a, b]$ at q and p, and apply the same argument in each of the three sections to get a pairing of all critical points in the interior of $[a, b]$. Their frames are therefore triple-panel windows with simple waves nested inside mid-panel of $W(a, b)$. Repeating the argument for the in-panel and the out-panel, we obtain the desired claim. □

Recall that a double-panel window is obtained by dropping the out-panel. The double-panel windows can be nested or disjoint, but in contrast to the triple-panel windows, they cannot overlap. Indeed by Lemma 4.3.1, non-nested windows do not cover each other's critical points. It follows that the overlap is limited to the in-panel of one and the out-panel of the other window. Since we drop the out-panel, double-panel windows cannot overlap.

4.3.2 Consequences: Symmetry and Variation

We use the hierarchies of triple- and double-panel windows to prove two folklore results about real-valued maps on the circle. The first is a statement of symmetry that follows from Alexander duality. Given a multiset of points in $\mathbb{R}^2$, such as $\mathrm{D}gm(f)$, we write $\mathrm{D}gm^{\circ}(f)$ for the central reflection, which negates coordinates. Similarly, we write $\mathrm{D}gm^{R}(f)$ for the reflection across the major diagonal, which switches coordinates, and $\mathrm{D}gm^{r}(f)$ for the reflection across the minor diagonal, which negates and switches coordinates.

Corollary 4.3.1 (Strong Symmetry for Circles)**.** *Let $f\colon \mathbb{S}^1 \to \mathbb{R}$ be generic. Then $Dgm(f) = Dgm^{R}(f)$ and $Dgm(-f) = Dgm^{r}(f)$.*

Proof. A window with simple wave of f is also such a window of $-f$. Hence, $(A, B) \in \mathrm{O}rd(f)$ iff $(B, A) \in \mathrm{R}el(f)$. Recall also that $\mathrm{E}ss(f)$ consists of two points, (A_0, B_0) and (B_0, A_0), in which $A_0 = \min_x f(x)$ and $B_0 = \max_x f(x)$. This implies $\mathrm{D}gm(f) = \mathrm{D}gm^{R}(f)$.

To relate f with $-f$, note that both have the same critical points, except that minima switch with maxima. Since $W(a, b) = J \times [A, B]$ is a triple-panel window of f iff $W(b, a) = J \times [-B, -A]$ is a triple-panel window of $-f$, this implies that we get the diagram of $-f$ by negating and switching the coordinates; that is: $\mathrm{D}gm(-f) = \mathrm{D}gm^{r}(f)$. □

To state the second result, we recall that the *variation* of a 1-dimensional function is the total amount of climbing up and down. We claim that for $f\colon \mathbb{S}^1 \to \mathbb{R}$, this is the total persistence of f, which we recall is the sum of $|B - A|$ over all points $(A, B) \in \mathrm{D}gm(f)$.

Corollary 4.3.2 (Variation for Circles)**.** *Let $f\colon \mathbb{S}^1 \to \mathbb{R}$ be generic. Then the variation equals the total persistence of f: $Var(f) = \|Dgm(f)\|_1$.*

Proof. We use induction, considering the double-panel windows defined by min-max pairs of f in a sequence in which the children precede their parents. Observe that f restricted to the support of a double-panel window without children consists of two monotonic pieces. Its contribution to the variation of f is twice the height of the window, and so is its contribution to the total persistence. Indeed, the min-max pair corresponds to a point each in the ordinary and the relative subdiagrams, or it corresponds to two points in the essential subdiagram. After recording these contributions, we locally flatten f to remove the double-panel window and continue the inductive argument. □

The relation between the variation and the total persistence of a map on $\mathbb{S}^1$ expressed in Corollary 4.3.2 was known before. For example, it is used to measure to what extent a noisy cyclic map is periodic [34]. Its generalization to maps on geometric networks stated in Corollary 4.5.2 is however new.

4.4 The Tree Case

In this section, we consider networks without cycles, which if connected are trees. We begin with a single edge and continue with geometric trees whose interior vertices have degree 3.

4.4.1 Maps on the Interval

The simplest compact 1-dimensional space that is not a 1-manifold is a line segment, which we refer to as an *interval* and parametrize from 0 to 1. Recall that a map $f\colon [0,1] \to \mathbb{R}$ is *generic* if the minima, maxima, and endpoints are isolated and their values are distinct. Theorem 4.3.1 applies in the interior of the interval, but we need new kinds of windows that cover the endpoints. Let a be a minimum or $\nearrow$-type endpoint and b a maximum or $\searrow$-type endpoint of f, write $A = f(a)$ and $B = f(b)$, and recall that $J = J(a,b)$ is the component of $f^{-1}[A,B]$ that contains both a and b, with $J = \emptyset$ if no such component exists.

Definition 4.4.1 (Windows for Intervals)**.** *Let a be a (non-global) minimum or $\nearrow$-type endpoint, and b a (non-global) maximum. We call the non-empty frame, $W(a,b) = J \times [A,B]$, a* triple-panel window with short wave *if its in-, mid-, out-panels are delimited by $0 \leq a < b < x < 1$ or by $1 \geq a > b > x > 0$ such that $f(x) = A$.*

Observe that Definition 4.4.1 allows for the cases $a = 0$ and $a = 1$. As illustrated in Figure 4.3, a window with short wave covers exactly one endpoint of the interval, and this endpoint is either a or a $\searrow$-type endpoint. The case in which the window covers both endpoints is also possible but different and introduced in Definition 4.5.1. In contrast to windows with simple wave, windows with short wave do not come in symmetric pairs; that is: if $W(a,b)$ is a window with short wave of f, then $W(b,a)$ is not a window with short wave of $-f$.

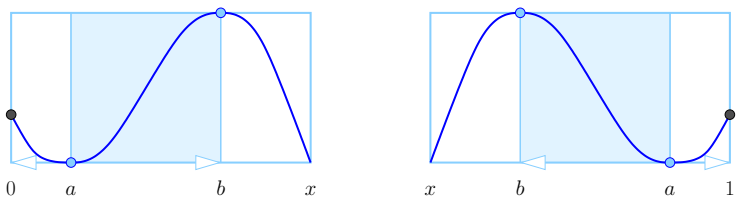

Figure 4.3: Two triple-panel windows with short wave, oriented from left to right on the *left* and from right to left on the *right*. Both cases may degenerate to zero-width in-panels. The *black* points correspond to endpoints of the interval. There are different ways how a frame can fail to be a window, one being that $f(x) > f(a)$.

Because of the asymmetry of windows with short wave, the extension of Theorem 4.3.1 to intervals requires a separate treatment of the ordinary and relative subdiagrams of $\mathrm{D}gm(f)$.

Theorem 4.4.1 (Characterization for Intervals)**.** *Let $f\colon [0,1] \to \mathbb{R}$ be generic, let a be a minimum or $\nearrow$-type endpoint, with $f(a) = A$, and let b be a maximum or $\searrow$-type endpoint, with $f(b) = B$. Then*

(i) $(A, B) \in Ord(f)$ *iff* $W(a, b)$ *is a triple-panel window with simple or short wave of* f*,*

(ii) $(B, A) \in Rel(f)$ *iff* $W(b, a)$ *is a triple-panel window with simple or short wave of* $-f$*.*

Proof. The pairs in (i) correspond to components of the sublevel set, which are counted by H_0, while the points in (ii) correspond to relative cycles, which are counted by H_1. The proof of (i) is almost verbatim the same as that of Theorem 4.3.1, and we omit the details.

Write $\mathbb{I} = [0, 1]$ and recall that $f^t = f^{-1}[t, \infty)$. To prove (ii), we explain the details of how $\mathsf{H}_0(f^t)$ and $\mathsf{H}_1(\mathbb{I}, f^t)$ are related. To this end, we decrease t from ∞ to $-\infty$ and show that the two groups change their ranks in parallel, with only one exception at $t = B_0$, the value of the global maximum, when $\mathsf{H}_0(f^t)$ goes from rank 0 to 1 while $\mathsf{H}_1(\mathbb{I}, f^t)$ remains at rank 0. For this purpose, we consider the long exact sequence of the pair $(\mathbb{I}, f^t)$. We recall that *exactness* means that the image of a map is the kernel of the next map in order along the sequence; see [43, Section IV.4] or [66, Section 2.1] for details. In the 1-dimensional case, all homology groups of dimension other than 0 and 1 are trivial, so the long exact sequence is rather short:

$$0 \to \mathsf{H}_1(f^t) \to \mathsf{H}_1(\mathbb{I}) \to \mathsf{H}_1(\mathbb{I}, f^t) \to \mathsf{H}_0(f^t) \to \mathsf{H}_0(\mathbb{I}) \to \mathsf{H}_0(\mathbb{I}, f^t) \to 0. \tag{4.2}$$

We have $rank\, \mathsf{H}_0(\mathbb{I}) = 1$ and $rank\, \mathsf{H}_1(\mathbb{I}) = rank\, \mathsf{H}_1(f^t) = 0$ for every t. There are only four possibly non-trivial groups, which we related to each other in a case analysis.

- For $t > B_0$, the only non-trivial groups are $\mathsf{H}_0(\mathbb{I})$ and $\mathsf{H}_0(\mathbb{I}, \emptyset)$, which both have rank 1. In particular, $\mathsf{H}_1(\mathbb{I}, f^t)$ and $\mathsf{H}_0(f^t)$ are both trivial and therefore isomorphic.
- For $t \leq B_0$, $\mathsf{H}_0(\mathbb{I}, f^t)$ is trivial, so by the exactness of (4.2), $rank\, \mathsf{H}_1(\mathbb{I}, f^t) = rank\, \mathsf{H}_0(f^t) - 1$.

To finish the argument, we remove the class born at $t = B_0$ from all groups $\mathsf{H}_0(f^t)$ to get two isomorphic persistence modules. It follows that the implied pairing of the critical values is the same, whether we track the components of f^t or the relative cycles of $(\mathbb{I}, f^t)$. Claim (ii) thus follows from (i). □

In addition to the points in the ordinary and relative subdiagrams—which are characterized by Theorem 4.4.1—$\mathsf{D}gm(f)$ contains one more point, namely (A_0, B_0) in the essential subdiagram. This point will be discussed in Section 4.5.

4.4.2 Maps on Geometric Trees

If we glue intervals at their endpoints without forming a cycle in the process, we get a *geometric tree*, $\mathbb{A} = (V, E)$, with *vertices*, V, and *edges*, E. We restrict ourselves to *degree-*3 *trees*, in which each vertex is an endpoint of either one or three edges. We call $f\colon \mathbb{A} \to \mathbb{R}$ *generic* if all critical points are isolated, they have distict values, and every degree-3 vertex is either a y-type or a λ-type vertex. It is tempting to consider $\nearrow$- and y-type vertices as minima and $\searrow$- and λ-type vertices as maxima, but note that components of sublevel sets are born at $\nearrow$-type but not at y-type vertices, and they die at λ-type but not at $\searrow$-type vertices.

Geometric trees introduce the topological phenomenon of branching, which requires yet another extension of the notion of window with wave. Let a be a minimum or $\nearrow$-type vertex, with $f(a) = A$, and b a maximum or λ-type vertex, with $f(b) = B$. Recall that $J = J(a, b)$ is the component of $f^{-1}[A, B]$ that contains both a and b, which is a geometric tree, and that a, b subdivide J into subtrees $J_{\text{in}}, J_{\text{mid}}, J_{out}$.

Definition 4.4.2 (Windows for Geometric Trees)**.** *We call a non-empty frame, $W(a, b) = J \times [A, B]$, a* triple-panel window with branching wave *if $f(x) > A$ for every point $x \neq a$ in $J_{in} \cup J_{mid}$, and $f(y) = A$ for at least one point $y \neq b$ in J_{out}.*

Note that the triple-panel windows with simple and short wave satisfy the conditions of Definition 4.4.2, but there are also others, as illustrated in Figure 4.4. We can now generalize Theorem 4.4.1 from intervals to geometric trees.

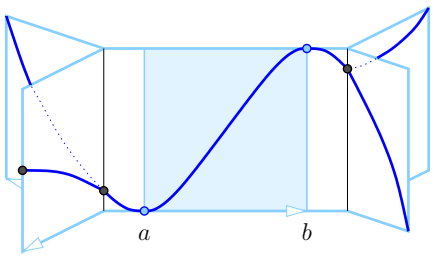

Figure 4.4: A triple-panel window with branching wave, $W(a, b)$. There is a branching point in the in-panel on the *left* and another in the out-panel on the *right*. Branching points and endpoints are marked in *black*. Note that $W(b, a)$ violates the conditions in Definition 4.4.2 for the negated map.

Theorem 4.4.2 (Characterization for Geometric Trees)**.** *Let $f : \mathbb{A} \to \mathbb{R}$ be a generic map on a compact geometric degree-3 tree, let a be a minimum, $\nearrow$-type, or y-type vertex, with $f(a) = A$, and let b be a maximum, λ-type, or $\searrow$-type vertex, with $f(b) = B$. Then*

(i) $(A, B) \in Ord(f)$ iff $W(a, b)$ is a triple-panel window with branching wave of f,

(ii) $(B, A) \in Rel(f)$ iff $W(b, a)$ is a triple-panel window with branching wave of $-f$.

Proof. The proof is almost verbatim the same as that of Theorem 4.4.1. We thus restrict ourselves to discussing what happens at the branching points as we monitor the sublevel set of f while t moves from $-\infty$ to ∞. A connected component is born at a minimum or a $\nearrow$-type vertex, a, and it dies at a maximum or a λ-type vertex, b. There is neither a birth nor a death when t passes the value of a $\searrow$-type vertex or a y-type vertex.

Assume b is a λ-type vertex. For $A < t < B$, we have $a \in f_t$ and $b \notin f_t$. At $t = B$, the component of the sublevel set that contains a merges with another, older component and therefore dies. Indeed, this other component approaches b from the other branch leading up to b, and after t passes B, the component extends along the branch leaving b upwards. □

Note that every vertex is paired only once: the $\nearrow$-type and λ-type vertices in Phase One, and the $\searrow$-type and y-type vertices in Phase Two. This is in contrast to the critical points in the

interior of the edges, which are paired twice. Indeed, according to Definition 4.4.2, $W(a, b)$ is not a window of f if a is a y-type vertex or b is a $\searrow$-type vertex. Symmetrically, $W(b, a)$ is not a window of $-f$ if b is a λ-type vertex or a is a $\nearrow$-type vertex. In addition to the points in the ordinary and relative subdiagrams—which are characterized by Theorem 4.4.2—D$gm(f)$ contains one point representing the one component, which is the entire geometric tree, in the essential subdiagram.

4.4.3 Consequences: Symmetry and Variation

For a map, f, on a geometric tree, the upside-down version of a window of f is not necessarily a window of $-f$. The strong symmetry statement in Corollary 4.3.1 thus fails to generalize and must be replaced by a weaker statement of symmetry. Recall that D$gm^{\circ}(f)$ and D$gm^{r}(f)$ are the reflections of D$gm(f)$ through the origin and across the minor diagonal.

Corollary 4.4.1 (Weak Symmetry for Geometric Trees)**.** *Let $f\colon \mathbb{A} \to \mathbb{R}$ be a generic map on a compact geometric tree. Then D$gm(-f)$ = O$rd^{\circ}(f) \sqcup$ R$el^{\circ}(f) \sqcup$ E$ss^{r}(f)$.*

Proof. Recall that D$gm(f) =$ O$rd(f) \sqcup$ R$el(f) \sqcup$ E$ss(f)$. By Theorem 4.4.2, the triple-panel windows with branching wave of f characterize O$rd(f)$, and those of $-f$ characterize R$el(f)$. For $-f$, we turn all windows upside-down, which switches and negates coordinates as well as switches the phases in which the windows are constructed. Hence, O$rd(-f) =$ R$el^{\circ}(f)$ and R$el(-f) =$ O$rd^{\circ}(f)$. There is only one point $(A_0, B_0) \in$ E$ss(f)$, in which A_0 and B_0 are the values of the global minimum and the global maximum of f. Similarly E$ss(-f)$ consists of a single point, $(-B_0, -A_0)$, which completes the proof. □

In contrast, Corollary 4.3.2 does generalize to geometric trees. However, the windows with non-simple wave complicate the proof of this generalization.

Corollary 4.4.2 (Variation for Geometric Trees)**.** *Let $f\colon \mathbb{A} \to \mathbb{R}$ be a generic map on a geometric tree. Then the variation equals the total persistence: $Var(f) = \|$D$gm(f)\|_1$.*

Proof. To formulate the proof strategy, we interpret each point $(A, B) \in$ D$gm(f)$ as the interval with endpoints A and B on the real line. We will show that for each non-critical value, $t \in \mathbb{R}$, the cardinality of $f^{-1}(t)$ is equal to the number of intervals in D$gm(f)$ that contain t. The claimed equation follows.

To begin, we add every minimum and maximum of f as a vertex to $\mathbb{A}$, so that f is monotonic on every edge of the thus subdivided geometric tree. We have six types of vertices, two each of degree 1, 2, and 3. We are interested in the change of the sublevel set and the superlevel set when t passes the value of a vertex:

- $\nearrow$-type endpoint: a component of f_t is born;
- $\searrow$-type endpoint: a cycle of $(\mathbb{A}, f^t)$ is born, unless the endpoint is the global maximum, in which case a component of f_t dies.
- minimum: a component of f_t is born and a cycle of $(\mathbb{A}, f^t)$ dies;
- maximum: a component of f_t dies, and a cycle of $(\mathbb{A}, f^t)$ is born, unless the maximum is the global maximum, in which case another component of f_t dies;

- y-type vertex: a cycle of $(\mathbb{A}, f^t)$ dies;
- λ-type vertex: a component of f_t dies.

We now increase t from $-\infty$ to ∞. The births and deaths of components correspond to start- and end-points of intervals, while the births and deaths of cycles correspond to end- and start-points of intervals, respectively. Accordingly, the number of intervals that contain t increases by 1 when t passes the value of a $\nearrow$-type endpoint or a y-type vertex, it decreases by 1 when t passes a $\searrow$-type endpoint or a λ-type vertex, it increases by 2 when t passes a minimum, and it decreases by 2 when t passes a maximum. The induction basis is provided by t smaller than the value of the global minimum, when there are no intervals that contain t and there are no points in $f^{-1}(t)$. The induction step is the observation that $\#f^{-1}(t)$ changes in the same way as the number of intervals that contain t, namely $\#f^{-1}(t)$ increases by 1 when t passes the value of a $\nearrow$-type endpoint or a y-vertex, etc. □

4.5 The Geometric Network Case

In this section, we take the step from maps on the unit circle and on geometric trees to maps on more general geometric networks. In contrast to a geometric tree, we do not assume that a geometric network is connected.

4.5.1 Stable Marriage

We call an element of $\mathsf{H}_1(\mathbb{G})$ a *cycle*, which by definition is an even degree and not necessarily connected subgraph of the network. We relate the global minima and maxima of the cycles in $\mathbb{G}$ to each other using the notion of a stable marriage. Let $f\colon \mathbb{G} \to \mathbb{R}$ be a generic map on a compact geometric network, and write $k = rank\, \mathsf{H}_1(\mathbb{G})$ for the rank of the cycle space. For $\Lambda \in \mathsf{H}_1(\mathbb{G})$, we introduce notation for the global minimum and maximum of f along Λ:

$$\mathrm{lo}(\Lambda) = \arg\min_{x\in\Lambda} f(x), \tag{4.3}$$
$$\mathrm{hi}(\Lambda) = \arg\max_{x\in\Lambda} f(x), \tag{4.4}$$

calling them the *low point* and the *high point* of the cycle. If cycles $\Lambda \neq \Lambda'$ have the same low point, then genericity implies the existence of a common arc that contains the shared low point in its interior. This arc does not belong to the sum, hence $f(\mathrm{lo}(\Lambda + \Lambda')) > f(\mathrm{lo}(\Lambda)) = f(\mathrm{lo}(\Lambda'))$. The symmetric inequality holds for cycles with shared high point. Write $\mathrm{Lo}(f)$ and $\mathrm{Hi}(f)$ for the collections of low and high points of all cycles. We begin by proving that both collections have cardinality k.

Lemma 4.5.1 (Low and High Points)**.** *Let $f\colon \mathbb{G} \to \mathbb{R}$ be a generic map on a compact geometric network. Then $\#Lo(f) = \#Hi(f) = rank\, \mathsf{H}_1(\mathbb{G})$.*

Proof. It suffices to prove that $\#\mathrm{Lo}(f)$ is equal to $k = rank\, \mathsf{H}_1(\mathbb{G})$ as the other equality is symmetric. Since $\mathsf{H}_1(\mathbb{G})$ is a vector space, every one of its bases consists of k cycles. Let $\Lambda_1, \Lambda_2, \ldots, \Lambda_k$ be a basis that maximizes $\sum_{i=1}^{k} f(\mathrm{lo}(\Lambda_i))$. We claim that their low points are distinct. Indeed, if $\mathrm{lo}(\Lambda_i) = \mathrm{lo}(\Lambda_j)$ with $i \neq j$, then $f(\mathrm{lo}(\Lambda_i + \Lambda_j)) > f(\mathrm{lo}(\Lambda_j))$ and we can substitute $\Lambda_i + \Lambda_j$ for Λ_j to get a new basis with larger sum of values. This contradiction implies $\mathrm{lo}(\Lambda_i) \neq \mathrm{lo}(\Lambda_j)$ whenever $i \neq j$ and therefore $\#\mathrm{Lo}(f) \geq k$.

To get $\#\mathrm{L}o(f) \leq k$, we observe that the low point of a sum of cycles in the basis (with distinct low points) is the lowest low point of these cycles and therefore one of the k low points we already observed exist. Thus, $\#\mathrm{L}o(f) = k$, as claimed. □

Since there are equally many low and high points, we can pair them up. Of particular interest is the solution to a *stable marriage* problem [58]. To formulate it, we call $b \in \mathrm{H}i(f)$ a *candidate* of $a \in \mathrm{L}o(f)$, and vice versa, if there exists a cycle, Λ, with $a = \mathrm{lo}(\Lambda)$ and $b = \mathrm{hi}(\Lambda)$. Among its candidates, a low point prefers high points with small function values, and a high point prefers low points with large function values. We write $\mathrm{hi}(a)$ and $\mathrm{lo}(b)$ for the *favorites* among their candidates and claim that everybody can be paired with its favorite.

Lemma 4.5.2 (Stable Marriage)**.** *Let $Lo(f)$ and $Hi(f)$ be the low and high points of a generic map on a compact geometric network, $f\colon \mathbb{G} \to \mathbb{R}$. Then $\mu\colon Lo(f) \to Hi(f)$ defined by $\mu(a) = hi(a)$ is a bijection, and it satisfies $\mu^{-1}(b) = lo(b)$.*

Proof. We show $b = \mathrm{hi}(a)$ iff $a = \mathrm{lo}(b)$, for all $a \in \mathrm{L}o(f)$ and $b \in \mathrm{H}i(f)$, which implies the claim. To reach a contradiction, suppose $b = \mathrm{hi}(a)$ but $a' = \mathrm{lo}(b)$ with $a' \neq a$. By definition of favorite, there exists a cycle, Λ, with $\mathrm{lo}(\Lambda) = a$ and $\mathrm{hi}(\Lambda) = b$. Hence, a is a candidate of b. However, since $a' \neq a$ is the favorite of b, this implies $f(a') > f(a)$. Let Λ' be the cycle with $\mathrm{lo}(\Lambda') = a'$ and $\mathrm{hi}(\Lambda') = b$. Then $\mathrm{lo}(\Lambda + \Lambda') = a$ and $f(\mathrm{hi}(\Lambda + \Lambda')) < f(b)$, which contradicts that b is the favorite of a. □

4.5.2 Maps on Geometric Networks

The components and cycles of $\mathbb{G}$ give rise to points in the 0- and 1-dimensional essential subdiagrams of $\mathrm{D}gm(f)$. They need new kinds of windows to be recognized. The more interesting case is that of a cycle. Let $a \in \mathrm{L}o(f)$, $b \in \mathrm{H}i(f)$, and recall the definition of $J = J(a, b)$. If a and b are candidates of each other, then $J \neq \emptyset$ as it contains at least the cycles whose low and high points are a and b. Even if a and b are not candidates of each other, $J \neq \emptyset$ is possible, but then it does not contain any cycle through the two points.

Definition 4.5.1 (Windows for Geometric Networks)**.** *Let $a \in \mathbb{G}$ be a minimum, $\nearrow$-type, or y-type vertex, with $f(a) = A$, and $b \in \mathbb{G}$ a maximum, $\searrow$-type, or λ-type vertex, with $f(b) = B$. Recall that $J = J(a, b)$ is the component of $f^{-1}[A, B]$ that contains both a and b, with $J = \emptyset$ if no such component exists.*

(i) $W(a, b) = J \times [A, B]$ *is a* window of component *if J is an entire component of $\mathbb{G}$.*

(ii) $W(a, b)$ *is a* window of cycle *if J contains a cycle that passes through a and b such that $J \setminus \{a, b\}$ is not connected.*

The window of cycle is illustrated in Figure 4.5: (a, A) and (b, B) lie on the lower and upper boundaries of the cylindrical strip. If $W(a, b)$ does not satisfy the conditions in Definition 4.5.1 (ii), then cutting the strip along vertical lines at a and b does not split it into two connected pieces. On the other hand, if $W(a, b)$ is a window of cycle, then the two cuts split the strip into two components. Note that a window with wave can neither be a window of component nor of cycle. On the other hand, it is possible that a window of component is also a window of cycle.

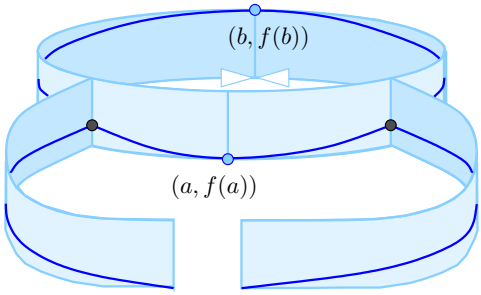

Figure 4.5: A window of cycle. If the two arms met at the ends, this would be a violation of the conditions in Definition 4.5.1 (ii) since cutting at a and b would not disconnect the strip.

The proof of Lemma 4.5.2 implies that $W(a, b)$ is a window of cycle iff a and b are each other's favorites. We show that this is also equivalent to being paired in persistent homology.

Theorem 4.5.1 (Characterization for Compact Geometric Networks)**.** *Let $f : \mathbb{G} \to \mathbb{R}$ be a generic map on a compact geometric network, let a be a minimum, $\nearrow$-type, or y-type vertex, with $A = f(a)$, and let b be a maximum, $\searrow$-type, or λ-type vertex, with $B = f(b)$. Then*

(i) $(A, B) \in Ess_0(f)$ iff $W(a, b)$ is a window of component,

(ii) $(B, A) \in Ess_1(f)$ iff $W(a, b)$ is a window of cycle.

Proof. (i) is obvious enough so we omit the proof. To see (ii), assume a and b are each other's favorites, and let Λ be a cycle whose low and high points are a and b. When $t \in \mathbb{R}$ reaches B in Phase One, Λ is born along with all cycles $\Lambda + \Lambda'$, in which Λ' is a cycle born before Λ. All these cycles die when t reaches A in Phase Two. Indeed, if Λ' died earlier, then $\Lambda + \Lambda'$ would become homologous to Λ, but since Λ is born after Λ', the sum of the two cycles cannot die yet. On the other hand, $\Lambda + \Lambda'$ dies at $t = A$ because it becomes homologous to Λ', which was born earlier. □

The characterization of points in the essential subdiagram of $Dgm(f)$ in Theorem 4.5.1 together with the characterization of the points in the ordinary and relative subdiagrams in Theorem 4.4.2 completes the proof of the Main Theorem stated in the Introduction.

4.5.3 Consequences: Symmetry and Variation

The weak symmetry assertion for geometric trees stated in Corollary 4.4.1 generalizes to geometric networks.

Corollary 4.5.1 (Weak Symmetry for Geometric Networks)**.** *Let $f : \mathbb{G} \to \mathbb{R}$ be a generic map on a compact geometric network. Then $Dgm(-f) = Ord^{\circ}(f) \sqcup Rel^{\circ}(f) \sqcup Ess^{r}(f)$.*

Proof. The argument for the triple-panel windows with wave is the same as in the proof of Corollary 4.4.1. Since geometric networks are not necessarily connected, we can have more than one window of component, which is different for geometric trees, which are connected. Nevertheless, the argument for such windows is the same as in the proof of Corollary 4.4.1.

It remains to argue about the cycles in the network. By Lemma 4.5.2, the cycles are represented by pairing their low and high points in a symmetric manner. Specifically, each low point is

paired with the lowest candidate high point, and because the candidate relation is symmetric, this is equivalent to pairing each high point with the highest candidate low point. Each such pair generated in Phase One corresponds to a point $(A, B) \in \mathrm{E}ss(f)$, and by symmetry to a point $(-B, -A) \in \mathrm{E}ss(-f)$, which completes the proof. □

The equality of the variation and the total persistence generalizes from circles and geometric trees to geometric networks. We can reuse the proof of Corollary 4.4.2, which we complement with an argument about cycles.

Corollary 4.5.2 (Variation for Geometric Networks)**.** *Let $f\colon \mathbb{G} \to \mathbb{R}$ be a generic map on a compact geometric network. Then the variation is the total persistence: $Var(f) = \|Dgm(f)\|_1$.*

Proof. We cut each cycle in $\mathbb{G}$ at its high point to obtain a geometric network, $\mathbb{G}'$, with one less cycle. Let $\eta\colon \mathbb{G}' \to \mathbb{G}$ be the surjection that reverses the cut, and let $g\colon \mathbb{G}' \to \mathbb{R}$ be defined by $g(x) = f(\eta(x))$. Since the maps are essentially the same, we have $\mathrm{V}ar(g) = \mathrm{V}ar(f)$.

To show that the total persistence remains the same, let Λ be a cycle in $\mathbb{G}$, $a = \mathrm{lo}(\Lambda)$ its low point, and $b = \mathrm{hi}(\Lambda)$ its high point. Assume that $W(a, b)$ is a window of cycle, so that $(A, B) \in \mathrm{E}ss_1(f)$, in which $A = f(a)$ and $B = f(b)$, as usual. The cut at b removes the cycle and thus the point (A, B) from the diagram. There is a second window, generated by b and another point $x \in \mathbb{G}$, whose corresponding point, (B, X), is removed from the diagram. In lieu of b, we get two $\nearrow$-type endpoints in $\mathbb{G}'$, which we denote b' and b''. By definition of η, we have $g(b') = g(b'') = B$. Since b' and b'' are endpoints, they are paired only once. By the local characterization of windows in Theorems 4.3.1, 4.4.1, 4.4.2, 4.5.1, all windows of f other than $W(a, b)$ and $W(b, x)$ are also windows of g. Hence b' and b'' can only be paired with a and x. We thus get two new points, (B, A) and (B, X) in $\mathrm{D}gm(g)$. Their persistence is the same as that of the two points they replace, so $\|\mathrm{D}gm(g)\|_1 = \|\mathrm{D}gm(f)\|_1$.

We now repeat the argument, cutting one cycle at the time, until we reach a collection of geometric trees. Corollary 4.4.2 implies that the variation is equal to the total persistence. Since both quantities did not change during the process, we thus established the equality also for compact geometric networks. □

4.6 Maintaining Persistence for Time Series data

Can we leverage the structural results of this chapter to *update* the persistent homology of one-dimensional data as the function, or the space, undergo a local change? The question was tackled in a joint effort with Herbert, as well as Monika Henzinger and her student Lara Ost, whose expertise is in the design of efficient algorithms and data structures. This subsection is taken from our joint paper [27] recently published at SODA and aims to summarize the key findings of the project. Instead of working with geometric networks, this section focuses on the simpler setting of functions on intervals, often referred to as time series; see Figure 4.6 for an example of such map.

Setting. We restrict ourselves to one-dimensional input data, i.e., a list of m points (or *items*) in an interval of $\mathbb{R}$ with each item i being assigned a *value* $f(i)$. Persistent homology has been applied to such time series data in multiple contexts, for example to heart-rate data

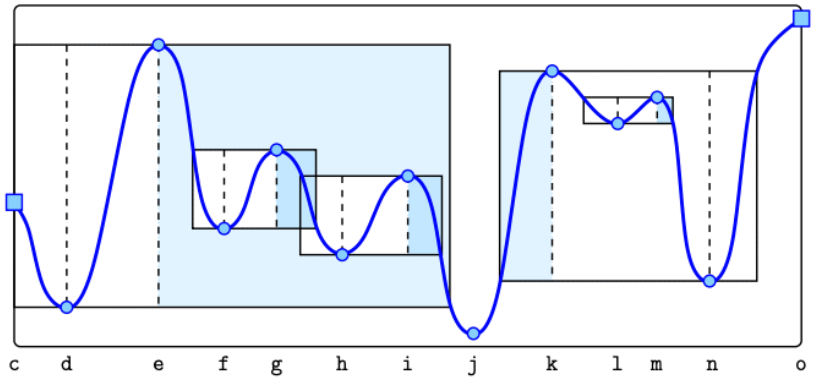

Figure 4.6: The graph of a generic map on a closed interval. All windows shown are with simple wave, except for the *leftmost* window, whose wave is short. The global window as well as the (tiny) windows caused by the hooks are not shown. The *light-blue shaded* out-panels are part of the triple- but not of the double-panel windows.

[19, 63], gene expression data [34, 88], and financial data [60]. The persistent homology of one-dimensional data can be derived from the merge tree [102], which records more detailed information about the structure of the persistence diagram, called the *history of the connected components in the filtration of sublevel sets*. Without recovering this history, the persistence information can be computed in $O(m)$ time [61]. To the best of our knowledge, the new $O(m)$ time algorithm that we introduced in our SODA paper [27] is the first linear-time algorithm that can also recover the history, which we store in the *augmented persistence diagram* of the filtration of sublevel sets. This diagram is the extended persistence diagram of [22], which we introduced in Section 4.2.3, together with a relation that encodes the merge tree. We draw an arrow between two points in the persistent diagram whenever the corresponding windows are nested, and there is no other window nested between the two. See Figure 4.7 for an example.

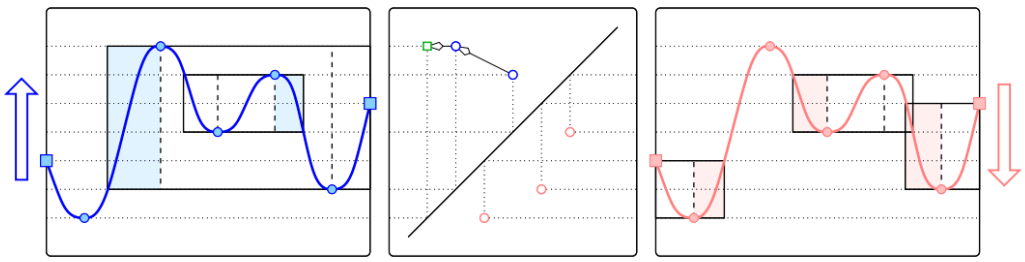

Figure 4.7: *Left:* a real-valued map on a closed interval, f, with three minima and two maxima. *Right:* the map $-f$ drawn upside-down. *Middle:* the augmented persistence diagram with two (*blue*) points in the ordinary subdiagram $\mathrm{Ord}(f)$ above the diagonal, three (*pink*) points in the relative subdiagram $\mathrm{Rel}(f)$ below the diagonal, and one (*green*) point in the essential subdiagram $\mathrm{Ess}(f)$.

As the data may change, it is an interesting question whether persistent homology can be maintained efficiently under update operations. The historically first such algorithm [24] takes time linear in the complex size per swap in the ordering of the simplices; see also [77, 89]. For

one-dimensional data this corresponds to a change of the value of an item, which reduces to a sequence of interchanges of f-values, each costing time linear in the size of the persistence diagram. More recently, [36] has shown that this can be strengthened to logarithmic time if the input complex is a graph. Our paper [27] is the first to maintain the persistent homology of dynamically changing one-dimensional input with a tailor-made data structure under a larger suite of update operations, which includes the *insertion* of a new item, the *deletion* of an item, the *adjustment* of the value of an item, the *cutting* of a list of items into two, and the *concatenation* of two lists into one.

The running time per operation is $O(\log n + k)$, in which n is the current number of critical items, and k is the number of changes to the augmented persistence diagram caused by the operation.

Our novel dynamic data structure is based on the characterization of the items in the persistence diagram through *windows*, as recently established in [12]. The main data structure is a *binary tree* ordered by position as well as value (see [108] for the introduction of such a binary tree), a *path-decomposition* of this tree dictated by persistent homology such that each path represents a window, and a final relaxation obtained by splitting each path into a *left trail* of nodes with right children on the path and a *right trail* of nodes with left children on the path. This split of each path into two trails is crucial for our results, and without it, the update time would have a linear dependence on the depth of the tree, which might be $\Theta(n)$.

Data Structures Overview. We propose a novel collection of data structures—some classic and some new—for maintaining nested windows. We show that using these data structures allows us to maintain the augmented persistence diagram to reflect the change from one map to another in the desired time bound. We maintain the following data structures:

(1) a *doubly-linked list* of all items, critical or not, ordered by their positions in the interval;

(2) two *binary search trees*, called *dictionaries*, one storing the minima and the other storing the maxima, both ordered by their positions in the interval;

(3) two *banana trees* (described next) representing the information in the augmented persistence diagram by storing the minima and maxima while reflecting their ordering by position as well as by function value.

Note that the minima and maxima are subsets of all items and are therefore represented in all of the above data structures. To reduce the special cases in the algorithms, we add two artificial items, called *hooks*, at the very beginning and the very end of the interval. They make sure that the formerly first and last items are proper minima or maxima, but we ignore them and the technicalities involved in this overview and give slightly simplified descriptions of the data structures and the algorithms.

Banana Trees. We introduce these trees in three stages. In the *first stage*, we organize all windows in a full binary tree, whose leaves are the minima and whose internal nodes are the maxima, such that (i) the in-order traversal of the tree visits the nodes in increasing order of their positions in the interval, and (ii) the nodes along any path from a leaf to the root are ordered by increasing function value. Such a tree always exists and it is unique. In the following, we do not distinguish between a node in the tree and the critical item it represents.

In the *second stage*, we add a *special root* labeled with value larger than the global maximum whose only child is the previous root. Call the resulting tree T and note that it has an equal number of leaves and internal nodes. We then form paths, each starting at a leaf, a, and ending at an internal node, b, such that a and b span a window, $W(a, b)$. Call this path $P(a, b)$. Based on structural properties of T, this leads to a partition of T into edge-disjoint (but not vertex-disjoint) paths; see the left drawing in Figure 4.8. The node b ending the path that starts at a is locally determined: it is the first internal node encountered while walking up from a, such that a and the descending leaf with minimum function value lie on different sides (in different subtrees) of this internal node. Hence, every maximum on the path from a to b spans a window that is immediately nested in $W(a, b)$. It follows that every maximum, b, belongs to two paths: $P(a, b)$ and $P(p, q)$ such that $W(a, b)$ is immediately nested in $W(p, q)$. This even holds for the root (of the full binary tree), which spans a path and also belongs to the path that ends at the special root. Given a map, f, the tree, T, and its partition into paths are unique. The strict dependence on f may force T to be unbalanced, and indeed have linear depth, so that efficient maintenance algorithms are challenging. This is why we need another modification.

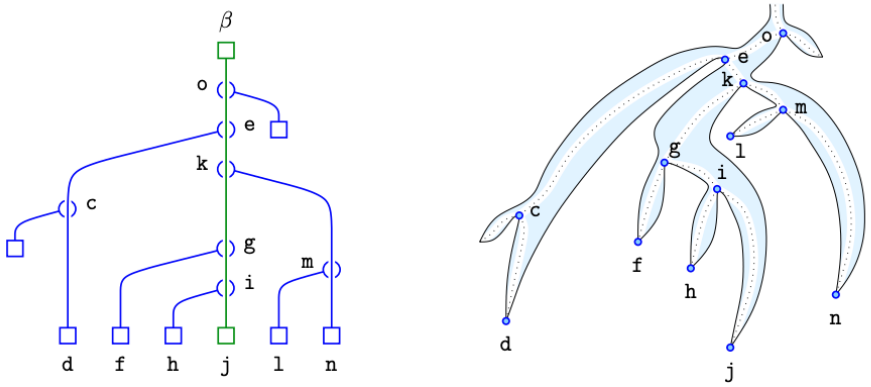

Figure 4.8: The path-decomposed binary tree associated to the map in Figure 4.6 with special root, β, on the *left*, and the corresponding banana tree on the *right*.

In the *third stage*, we split each path into two trails; see the right drawing in Figure 4.8. The *left trail* of $P(a, b)$ contains a and every maximum u on $P(a, b)$ with $u \leq a$, while the *right trail* contains a and every maximum u on $P(a, b)$ with $a \leq u$. A node v on $P(a, b)$ is the right child of its parent, u, in the binary tree iff u is on the left trail. Furthermore, v is the left child of u iff u is on the right trail. Thus, to which trail u belongs to can be decided based on local information at u. To simplify language, we also give a second name to the trails. If $a < b$, then $W(a, b)$ consists of the in-panel on the left, the mid-panel in the middle, and the out-panel on the right. In this case, u belongs to the left trail iff the window it spans is nested inside the in-panel, so we alternatively call the left trail the *in-trail*, and we alternatively call the right trail the *mid-trail*. If $b < a$, then the right trail is the in-trail and the left trail is the mid-trail. So what happened to the out-trail? It is indeed needed, but only if $W(a, b)$ is with simple wave. Such a window corresponds to $P(a, b)$ in the banana tree of f and to $P(b, a)$ in the banana tree of $-f$. The out-panel of $W(a, b)$ is the in-panel of $W(b, a)$, so we can get information about the out-panel from the in-panel of the same window in the other banana tree when we need it.

To speed up the algorithms, we maintain various pointers, such as between the occurrences of the same critical item in the banana trees, the dictionaries, and the doubly-linked list, for

example. Importantly, every internal node, b, stores a pointer to the descending leaf with minimum function value, $low(b)$, and every leaf, a, stores a pointer to the endpoint of its path, $dth(a)$. Observe that $q = dth(low(b))$ is the maximum that spans the window in which $W(a, b)$ is immediately nested, so this window can be obtained in constant time.

Construction. While the main focus of the paper [27] is the maintenance of the data structures through local updates, we also consider the construction from scratch. It is straightforward to derive the augmented persistence diagram during a single traversal of the banana trees in linear time, so the question we study is how fast this diagram can be constructed from a given one-dimensional input list. Assuming all non-critical items have been removed and we are given the remaining sequence of n critical items, there are standard algorithms that can be adapted to construct the banana trees in $O(n \log n)$ time. There is also an $O(n)$ time algorithm for computing the persistence diagram [61], but this algorithm does not extend to the augmented persistence diagram. To the best of our knowledge our algorithm is the first to construct the augmented persistence diagram in $O(n)$ time.

The main structure of the algorithm is a left-to-right scan of the data. We interpret the item i with value $c_i = f(i)$ as the point (i, c_i) in the plane and maintain a decreasing staircase such that all processed items are points on or below the staircase. Each step of the staircase corresponds to an unfinished banana. One of the difficulties is that before a banana is completed, we do not know whether it will be attached to a left or a right trail. We tentatively assume it will be attached to a left trail but are prepared to move the banana to the other side when this turns out to be necessary. When we process the next item, we may remove any number of steps, turning each into a finished banana, but we can add at most one new step. Since a step that is removed was added earlier, this proves that the algorithm runs in $O(1)$ amortized time per item, and therefore in $O(n)$ time altogether.

Local Maintenance. Given a list of m items with real function values, we consider the operations that *insert* a new item, *delete* an item, and *change the value* of an item. All three operations reduce to a sequence of *interchanges*—which can be between two maxima or between two minima—possibly preceded by an *anti-cancellation* or a *slide*, and possibly succeeded by a *cancellation* or a *slide*. In a slide, a minimum or maximum next to a non-critical item becomes non-critical, and the non-critical item becomes a minimum or maximum, respectively. Similarly in a cancellation, a minimum and a neighboring maximum simultaneously become non-critical. Here we will focus on the interchanges, because they are most common as well as most interesting, and on the anti-cancellations, because they pose an unexpected challenge.

Consider two maxima, b and q, of f, and assume $f(b) < f(q)$ before the interchange. To avoid confusing language, we write $g(b)$ and $g(q)$ for the values after the operation but assume that f and g agree on all items except for b. Furthermore, we assume that $g(b) > g(q)$ and that b and q are the only two items for which the ordering by f-value differs from the ordering by g-value. In many cases, the interchange of b and q does not affect the banana trees. Indeed, only if b is a child of q is it necessary to update the order of the two nodes. And even if b and q are consecutive maxima on a path, there is no structural change unless b and q also belong to a common trail. This is the main reason for splitting each path into two trails as explained in the third stage of the introduction of the banana trees: to avoid any cost to occur for interchanges that have no structural affect on the augmented persistence diagram. When b and q interchange while belonging to different trails, then the banana tree is oblivious to this change and requires no update. On the other hand, if b and q belong to the same trail,

then they swap positions along this trail, and there is a change of the augmented persistence diagram that pays for the time it takes to update the banana tree.

The interchange of two minima, a and p, is quite different because it does not affect the ordered binary tree at the first stage of the banana tree. However, the interchange affects the path-decomposition, so some of the bananas may have to be updated. To be specific, assume $f(a) > f(p)$ before and $g(a) < g(p)$ after the operation. As before, we also assume that f and g agree on all items except for a, and that a and p are the only two items for which the ordering by f-value differs from the ordering by g-value. Let b and q be the internal nodes so that $P(a, b)$ and $P(p, q)$ are paths in the decomposition of the banana tree of f. The interchange of a and p has no effect on the banana tree, unless b is a node on $P(p, q)$. If b lies on $P(p, q)$, then we extend $P(a, b)$ to $P(a, q)$ and we shorten $P(p, q)$ to $P(p, b)$. The nodes u on the path from b to the child of q change their pointer from $low(u) = p$ to $low(u) = a$. There can be arbitrarily many such nodes, but each change causes the adjustment of an arrow in the extended persistence diagram, to which it can be charged in the running time analysis.

This shows that it is possible to perform an interchange of two minima within the desired time bound, but it is not clear how to find them. Considering the scenario in which a decreases its value continuously, it may cause a sequence of interchanges with other minima, but since these minima are not sorted by function value, it is not clear how to find them, and how to ignore the ones without structural consequences. Here is where the relation between the banana trees of f and $-f$ becomes important. The minima of f are the maxima of $-f$, so the interchange of two minima in the banana tree of f corresponds to the interchange of two maxima in the banana tree of $-f$. We already know how to find the interchanges of two maxima and how to ignore the ones without structural consequences, so we use them to identify the interchanges of minima. More precisely, we prove that the interchange of the minima, a and p of f, affects the structure of the banana tree of f only if the interchange of the maxima, a and p of $-f$, affect the structure of the banana tree of $-f$. The converse does not hold, but the implication suffices since the interchange of the maxima of $-f$ leads to a change in the extended persistence diagram which can be charged for the interchange of the minima, which costs only $O(1)$ time if there are no structural adjustments.

Next, we sketch what happens in the remaining operations. A *slide* occurs if a maximum decreases its value so that it becomes non-critical, while a neighboring non-critical item becomes a maximum. However, if this neighboring item is a maximum, then both items become non-critical at the same time, in which case the operation is called a *cancellation*. There are also the symmetric operations in which a minimum increases its value and becomes non-critical, while a neighboring item changes from non-critical to minimum (a *slide*) or from maximum to non-critical (a *cancellation*). The corresponding updates are easily performed within the required time bounds.

A more delicate operation is the *anti-cancellation*, in which two neighboring non-critical items become critical at the same moment in time. Let a and b be these two items and assume a is a minimum and b is a maximum after the operation, so $g(a) < g(b)$ and, by assumption, $f(a') > f(a) > f(b) > f(b')$, in which a', a, b, b' are four consecutive items in the list. Hence, a and b are not present in the banana tree of f, but $P(a, b)$ is a path in the path-decomposed tree of g, so a, b form a minimal banana in the banana tree of g. All we need to do is find the correct place to attach this banana, but this turns out to be difficult. We first explain why it is difficult, then present an algorithm that works within the current data structure, and finally sketch a modification of the data structure that accelerates the anti-cancellation to

$O(\log n)$ time. While this leads to a speed-up in the worst case for this type of operation, it slows down other operations by a logarithmic factor.

We use mirrors to explain in what situation an anti-cancellation is difficult. Their representation in the banana tree is indirect: the maximum is the upper end of a mid-trail, and its mirror is the upper end of the matching in-trail. In the tree, the two upper ends are the same node, but if we traverse the trails of the banana trees in sequence, we visit them at different times, and these times correspond to the positions along the interval, which are different for the maximum and its mirror. Every maximum has at most one mirror, so adding all mirrors to the list increases it to less than double the original size. Nevertheless, it is easy to construct a case in which there is a long subsequence of mirrors, and these mirrors are consecutive with the exception of a and b appearing somewhere in their midst. In this situation, finding the correct attachment of $P(a, b)$ in the banana tree of g needs the mirrors immediately to the left and right of b. The data structure as described provides no fast search mechanisms among the mirrors, so we find the correct attachment by linear search starting from the first maximum to the left of a. Suppose we pass k mirrors before we find this attachment. Then we have a nested sequence of k windows, and $W(a, b)$ is nested inside all of them. Hence, there is one new immediate nesting pair, while the transitive closure of the nesting relation gains k new pairs. The cost of the anti-cancellation can be charged to the change of this transitive closure. In many cases, the change in the transitive closure will be comparable in size to that of the transitive reduction, but it can also be significantly larger, like in the case we just described. We thus entertain alternatives.

There is a modification of the data structure that allows for fast searches among subsequences of mirrors: for any consecutive min-max pair in the list, store the mirrors between them in a binary search tree. In practice, there will be many very small such trees, but it is possible that the mirrors accumulate and produce a few large trees. With this modification, an anti-cancellation can be performed in $O(\log n)$ time. Note however, that these binary search trees have to be maintained, which adds a factor of $O(\log n)$ to the time of every operation, in particular to every interchange, of which there can be many.

Topological Maintenance. We call an operation *topological* if it cuts the list of points into two, or it concatenates two lists into one. More challenging topological operations, such as gluing the two ends of a list to form a cyclic list, or gluing several lists to form a geometric tree or more complicated geometric network are feasible but have not yet been worked out in detail.

When we *cut* the list of data into two, we *split* the banana trees of the map and its negative into two each. We mention both banana trees since we need information from the other to be able to split one within the desired time bound. Splitting one banana tree is superficially similar to splitting a binary search tree, but more involved. Let $f\colon [1, m] \to \mathbb{R}$ be the map before the operation and $g\colon [0, \ell] \to \mathbb{R}$ and $h\colon [\ell + 1, m] \to \mathbb{R}$ the maps after the operation. We write $z = \ell + \frac{1}{2}$ for the position at which the time series is cut. Since the banana trees are determined by the maps, we need to understand the difference between the windows of f and those of g and h. Whether or not a frame of f is also a window depends solely on the restriction of f to the support of the frame. To decide about the future of a window of f, let $[x, y]$ be its support. This window is also a window of g if $y < z$, and it is also a window of h if $z < x$. The windows that need attention are the ones with $x < z < y$. A triple-panel window consists of three panels, so we distinguish between three cases:

- $W(a, b)$ suffers an *injury* if z cuts through the in-panel. Then a and b lie on the same side of z and they still span a window, albeit with short wave.
- $W(a, b)$ suffers a *fatality* if z cuts through the mid-panel. Then a and b lie on different sides of z and need to find new partnering critical items to span new windows.
- $W(a, b)$ suffers a *scare* if z cuts through the out-panel. These windows are difficult to find in the banana tree of f, but they are easy to find in the banana tree of $-f$, in which $W(b, a)$ suffers an injury.

The splitting of the banana tree proceeds in three steps: first, find the smallest banana that suffers an injury, fatality, or scare; second, find the remaining such bananas; and third, split the banana tree of f into the banana trees of g and h. We address all three steps and highlight the most interesting feature in each.

The first step is difficult because of the lack of an appropriate search mechanism in the banana tree. To explain this, we recall that the banana tree stores the mirrors implicitly, as the upper ends of the in-trails. Like in the case of an anti-cancellation, the challenging case is when mirrors accumulate and we have to locate z in their midst. As before, we locate z by linear search, scanning the mirrors in the order of decreasing function value. In contrast to the anti-cancellation, we can now charge the cost for the search to the changes in the extended persistence diagram. Indeed, every mirror we pass belongs to a window that experiences a scare. Such a window of f is neither a window of g nor of h, so its spanning critical items will re-pair and span new windows after the operation.

The second step traverses a path upward from the smallest affected banana we just identified. In each step of the traversal, we determine the corresponding window and push it onto the stacks of windows that experience an injury, fatality, or scare. A window that experiences an injury remains a window, but now with short instead of simple wave. Whether this window with short wave belongs to the banana tree of g or that of h depends on whether the spanning minimum is to the right or the left of the spanning maximum. A window that experiences a fatality falls apart, with one of the two spanning critical items in g and the other in h. Finally, a window that experiences a scare stops to be a window, in spite of having both critical points in g or in h. As mentioned earlier, such a window is difficult to find in the banana tree of f, but it is easy to find in the banana tree of $-f$, where it experiences an injury. We thus process both banana trees simultaneously, and distributed the windows or their corresponding bananas as needed. The upward traversal halts when we reach the *spine* of the tree, by which we mean the sequence of left children that start at the root, or the sequence of right children that start at the root. This is important because moving further until we reach the root is not necessary and can be costly because the spine can be arbitrarily long and the steps towards the root cannot be charged to changes in the extended persistence diagram.

The third step re-pairs the critical items of the windows that experience a fatality or scare. All new windows are with short wave, for else they would be windows of f before the splitting, which is a contradiction. Hence, their bananas belong to the spines of the banana trees of g and h. The bananas in the spines have the particularly simple structure that their critical items come in sequence. For the right spine of the banana tree of g, this sequence increases from right to left, for the left spine of the banana tree of h, the sequence increases from left to right, and both are consistent with the sequence in which we collect the corresponding windows in the second step. It follows that the available critical items can be paired up in sequence, which takes $O(1)$ time per item.

Corresponding to the concatenation of two lists, we pairwise *glue* the banana trees of the two maps. The operation is the inverse of splitting, so we omit further details. A final word about the cost paid by the changes in the augmented persistence diagram. How do we compare one diagram with two, which we get after splitting? To deal with this issue, we consider the maps g and h to be *one* map, namely a map on two intervals. Then we still have one augmented persistence diagram and the symmetric difference between the diagrams before and after the operation is well defined.

4.7 Discussion and Open Problems

Over the past year, Lara Ost has been implementing the data structure and its algorithms, and the first experiments seem to suggest that updates can be performed extremely fast. We will hopefully be able to share more concrete results soon. With such an implementation, we can deepen our understanding of the persistence of random lists, which in turn can be used as a baseline for our understanding of non-random time series, such as heart rate data. In scenarios where data is constantly being added, as is the case of heart rate data collected with smartwatches, the relevance of dynamic algorithms becomes evident. What is less clear is whether the information encoded in the persistence pairs represents a biologically relevant signal in such contexts.

The theoretical foundations for our dynamic algorithms are described in [12], where maps are defined on 1-dimensional domains, which include but go beyond intervals are described. Can the banana trees be extended to geometric trees (which allow for bifurcations) and geometric networks (which allow for bifurcations as well as loops) without deterioration of the asymptotic running time? Because of the unconventional representation of trees in terms of bananas (pairs of parallel trails), we finally ask whether there are other dynamic data structure questions that can benefit from this representation.

CHAPTER 5

Chromatic Point Sets

In this chapter, which builds upon [9, 25], we develop a novel topological approach to quantify spatial interactions between several point sets, differentiated by color. The objective is to formulate a mathematical language capable of addressing questions such as: "how, how often, and at what scale do blue points surround groups of red points?", or "are there cycles made out of blue, red, and green points that make essential use of all three colors?". We tackle these questions from a multi-scale homological perspective, with the goal of disentangling patterns such as the ones shown in Figure 5.1. Before delving into the mathematical aspects of the problem, let's recall how we came to think about it in the first place.

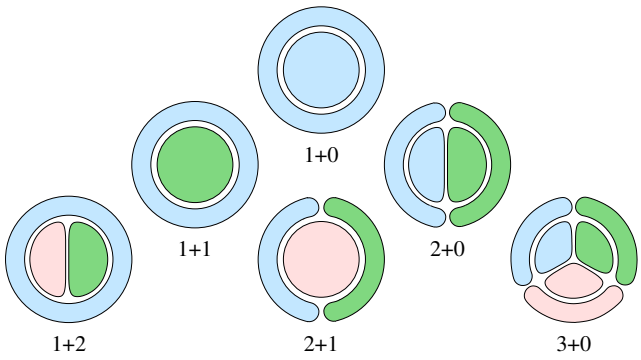

Figure 5.1: Mingling patterns distinguished by the number of colors needed to form a cycle and the number of additional colors needed to fill this cycle. The drawings are caricatures of similar patterns for cycles different from circles and fillings different from disks. The patterns are but a first attempt to differentiate types of interactions, and they are by no means precise or exhaustive. For example, two additional colors can fill a cycle in at least two different ways (see the pattern of type 1+2): in a collaboration as suggested in the drawing, or each individually, like two different patterns of type 1+1.

5.1 Motivation: Cellular Arrangements

Recent advances in technologies, including multiplexed imaging and spatial transcriptomics, have opened new avenues for observing both cellular location and phenotypes in native tissue environments, providing new opportunities for understanding the organizational principle

in biological systems [86]. This is particularly relevant in cancer research, where several recent studies have shown that the complex spatial arrangement formed by cancer cells and a collection of immune cells can provide mechanistic insights into disease progression and unveil biomarkers of response to existing treatments [104]. An example of such study is [96], where the authors demonstrated that the complex spatial arrangement of a specific type of immune cell, the tumor-infiltrating lymphocytes (TILs), is shown to be associated with disease progression and treatment response. However, the authors only considered the distribution of the lymphocytes, neglecting the distributions of other immune and cancer cells. Is their spatial arrangement clinically relevant? More subtly, instead of studying the patterns of the distribution of each cell type separately, would it be clinically useful to study the *interactions* between the distributions of cell types? Addressing this gap requires a mathematical model that can quantitatively describe the interactions between cell type distributions. Reading [96] while conducting research at Owkin motivated the author to think about how to approach this question topologically.

As an example, consider Figure 5.2, depicting a highly structured cellular arrangement in a healthy human lung. Such pattern, which is likely formed around a blood vessel, prompts the question: how can we quantitatively and precisely describe such system?

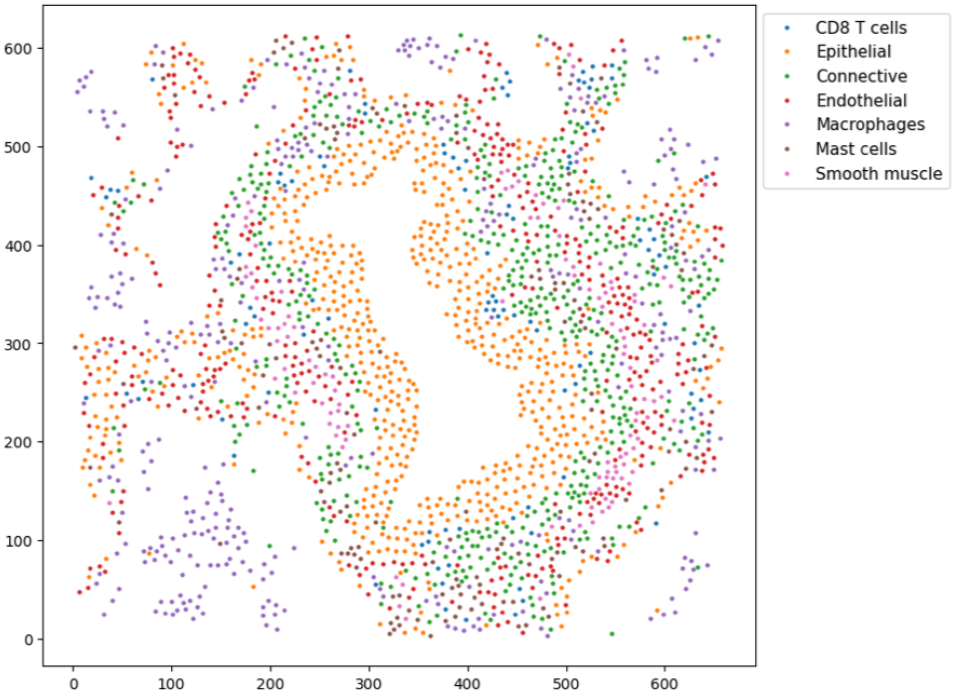

Figure 5.2: Two-dimensional cellular arrangement in a healthy human lung.

Persistent homology [46], when tailored to point set data, suggests itself as a natural candidate to achieve this goal. In fact, methods from topological data analysis have recently been applied in oncology, yielding critical insights in the treatment prognosis [30, 75, 37], tumor segmentation and diagnosis [91, 29, 28], disease classification [84, 5], and cellular architecture of tumor cells [98, 72, 7]. In this chapter, we focus on the latter application and we refer to [16] for a comprehensive survey on the different applications of TDA in oncology and [92] for the applications of topological methods in genomics and evolution.

It is worth noting that the rapid development of spatial transcriptomics is also shedding light on the complex relationship between a cell's microenvironment and its functional state [93].

Scientists in this field are starting to uncover how genetic changes influence cellular architecture and, and in parallel, how micro-environmental factors can alter cellular states, a concept known as plasticity. This leads to crucial questions about the causal relationships between intrinsic cellular properties and extrinsic ones (i.e. related to their environment), and the extent to which these factors influence metastatic tendencies of cancer cells. The author intends to focus on these research questions during his upcoming postdoctoral research.

Mathematical formulation and contributions To formally model such a system, we use points with an extra label, a *color*, to represent cells and their associated phenotypes. Methods from the field of topological data analysis [17], which are particularly suited for point cloud data, suggest themselves as a natural candidate to extract geometric features from such datasets. Specifically, when points have no labels, an appropriate filtration of the Delaunay mosaic, called the alpha filtration [48], can be used to quantitatively describe the spatial configuration of the point set. We seek for an analogous concept amenable to the setting in which points have an associated color. A solution for two colors was proposed by Reani and Bobrowski [94], which we generalize to arbitrarily many colors and whose structural and combinatorial properties we study. Given a locally finite set in $\mathbb{R}^d$ and a coloring with $s+1$ colors, this generalization places the points of different colors on $s+1$ parallel copies of $\mathbb{R}^d$, which intersect an orthogonal copy of $\mathbb{R}^s$ at the vertices of the standard s-simplex. This is a locally finite set in $\mathbb{R}^{d+s}$, and the *chromatic Delaunay mosaic* of the colored set in $\mathbb{R}^d$ is, by definition, the Delaunay mosaic of the set in $\mathbb{R}^{d+s}$. A similar set-up was used in [18] for the purpose of geometric morphing between $s+1$ shapes, so our work also sheds light on that proposal to construct a shape space.

The structural results we wish to highlight are as follows: (1) the chromatic Delaunay mosaic contains the chromatic Delaunay mosaic as well as the Delaunay mosaic of any subset of the $s+1$ colors as a subcomplex; in particular, it contains the Delaunay mosaic of each color individually and of all colors as subcomplexes; (2) the d-dimensional section of the colorful cells in the chromatic Delaunay mosaic (the cells that have at least one vertex of each color) is dual to the overlay of the $s+1$ mono-chromatic Voronoi tessellations.

Our combinatorial results help gauge the extent to which chromatic Delaunay mosaics can be used in applications. By the *size* of a mosaic we mean the number of cells, which we relate to the number of the points, denoted n. The dimension, d, and the number of colors, $s+1$, are assumed to be constants. We also consider locally finite but possibly infinite sets, namely Delone sets and Poisson point processes as examples of packed sets and random sets in $\mathbb{R}^d$, respectively. Here and later, we use the term 'packed' as a vague notion for locally finite sets of points that are, in a sense, d-dimensionally distributed. To facilitate the comparison with the results for finite sets, we count the points and cells within a sufficiently large ball centered at the origin.

As shown in Table 5.1, we have upper bounds for all three types of point sets assuming the colors are assigned at random. For any points and packed points, the bounds are formally stated in Theorems 5.3.1 and 5.3.2, and they are asymptotically tight. Not all bounds stated in Table 5.1 will be treated in this thesis – only the ones marked in bold will – and the interested reader should consult [9] for further details. Note the conspicuous absence of the number of colors in most bounds given, and this despite the fact that the chromatic Delaunay mosaic is a $(d+s)$-dimensional complex.

	chromatic Delaunay mosaic in $\mathbb{R}^{d+s}$			D. mosaic in $\mathbb{R}^{d+s}$
	any points	packed points	random points	(one color)
any colors	$n^{\min\{d,\lceil\frac{d+s}{2}\rceil\}}$ **Section 5.3**	$\min\{m^2, n^2\}$ in $\mathbb{R}^2$ $(^*)$ [9, Theor. 4.7]	?	$n^{\lceil\frac{d+s}{2}\rceil}$ [20]
random colors	$n^{\lceil d/2\rceil}$ **Theorem 5.3.1**	n **Theorem 5.3.2**	n [9, Theor. 5.1, 5.2]	

Table 5.1: Asymptotic size bounds for chromatic Delaunay mosaics of n points in $\mathbb{R}^d$ with $s+1$ colors. Constant factors are not shown. For the case of a packed set and any colors, we have a result only in $\mathbb{R}^2$ $(^*)$, in which m is the spread (the diameter divided by the minimum interpoint distance), which is at least a constant times $\sqrt{n}$. For comparison, we state the known maximum size of the (mono-chromatic) Delaunay mosaic of n points in $\mathbb{R}^{d+s}$ in the last column on the right [20].

5.2 Chromatic Complexes

The main concepts in this section are the chromatic Delaunay mosaics and Voronoi tessellations, which generalize the bi-chromatic construction in [94] to more than two colors.

5.2.1 Chromatic Delaunay Mosaics and Chromatic Voronoi Tessellation

Throughout this section, we let A be n points in $\mathbb{R}^d$, $\sigma = \{0, 1, \ldots, s\}$ a collection of colors, $\chi\colon A \to \sigma$ a coloring, and $A_j = \chi^{-1}(j)$ the subset of points with color j, for $0 \leq j \leq s$. We recall that the *standard s-simplex* is the convex hull of the $s+1$ unit coordinate vectors in $\mathbb{R}^{s+1}$. To map this simplex to s dimensions, we identify $\mathbb{R}^s$ with the s-plane defined by $x_1 + x_2 + \ldots + x_{s+1} = 1$ in $\mathbb{R}^{s+1}$ and parametrize it with the inherited $s+1$ *barycentric coordinates*. A subset of $t+1 \leq s+1$ unit coordinate vectors defines the standard t-simplex, which we map to $\mathbb{R}^t$ by parametrizing it with the $t+1$ barycentric coordinates inherited from $\mathbb{R}^{t+1}$. We are ready to construct the *chromatic Delaunay mosaic* of χ, denoted $\mathrm{D}el(\chi)$. We start by writing $\mathbb{R}^{s+d} = \mathbb{R}^s \times \mathbb{R}^d$, implying the explicit embeddings of $\mathbb{R}^s$ and $\mathbb{R}^d$ into $\mathbb{R}^{s+d}$, and then construct $\mathrm{D}el(\chi)$ in three steps:

Step 1: let $u_0, u_1, \ldots, u_s$ be the vertices of the standard s-simplex in $\mathbb{R}^s$;

Step 2: set $A' = A'_0 \sqcup A'_1 \sqcup \ldots \sqcup A'_s$, in which $A'_j = u_j + A_j \subseteq u_j + \mathbb{R}^d$, for each $0 \leq j \leq s$;

Step 3: construct $\mathrm{D}el(\chi) = \mathrm{D}el(A')$;

see Figure 5.3. As $\mathrm{D}el(\chi)$ is essentially a standard Delaunay mosaic, the paper can also be viewed as a study of Delaunay mosaics of point clouds restricted to a specific collection of affine spaces.

Similarly, we apply the construction to a subset of the colors, $\tau \subseteq \sigma$, and write $\mathrm{D}el(\chi|\tau)$, in which $\chi|\tau$ is our notation for the restriction of χ to $\chi^{-1}(\tau)$. This mosaic lives in $\mathbb{R}^{t+d}$, in which $t = 1 + \#\tau$. It is not difficult to see that $\mathrm{D}el(\chi|\tau)$ is a subcomplex of $\mathrm{D}el(\chi)$. To state this property formally, we call a cell in $\mathrm{D}el(\chi)$ *τ-colored* if the colors of its vertices belong to τ, and *τ-colorful* if it is τ-colored and has a vertex of every color in τ. Every cell is τ-colorful

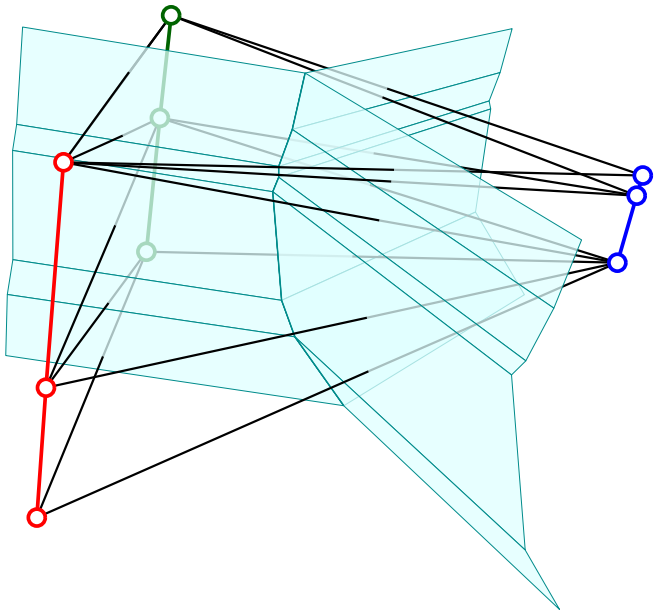

Figure 5.3: The chromatic Delaunay mosaic of three finite sets in $\mathbb{R}^1$ together with the stratification of space into membranes. The points of each set are placed on a copy of $\mathbb{R}^1$ orthogonal to the 2-plane that carries the standard triangle constructed in Step 1. The stratification consists of a 1-dimensional membrane geometrically located between the three lines, and three 2-dimensional membranes, one between any two of the lines.

for the smallest subset, $\tau \subseteq \sigma$, for which the cell is τ-colored. This implies that we get a partition of the cells into $2^{s+1} - 1$ classes. Note that the τ-colored cells form a subcomplex of $\mathrm{Del}(A)$, while the τ-colorful cells generally do not.

Lemma 5.2.1 (Sub-chromatic Delaunay Subcomplexes)**.** *Let $A \subseteq \mathbb{R}^d$ be finite, $\chi\colon A \to \sigma$ a coloring, and $\tau \subseteq \sigma$. Then the subcomplex of τ-colored cells in $\mathrm{Del}(\chi)$ is $\mathrm{Del}(\chi|\tau)$.*

Proof. Let H be a hyperplane in $\mathbb{R}^{d+s}$ that passes through all points with color in τ such that all other points in A' are contained in an open half-space bounded by H. The cells of $\mathrm{Del}(\chi|\tau)$ are characterized by the existence of an empty $(t+d-1)$-sphere in H that passes through the vertices of the cell and through no other points with color in τ. Since all points with color in $\sigma \setminus \tau$ lie in an open half-space bounded by H, we can extend this $(t+d-1)$-sphere to an empty $(d+s-1)$-sphere that passes through the same points. Hence, $\mathrm{Del}(\chi|\tau) \subseteq \mathrm{Del}(\chi)$, which implies the claim because $\mathrm{Del}(\chi|\tau)$ exhausts all τ-colored cells in $\mathrm{Del}(\chi)$. □

The *chromatic Voronoi tessellation* of $\chi\colon A \to \sigma$ is the Voronoi tessellation of $A' \subseteq \mathbb{R}^{d+s}$, and we write $\mathrm{Vor}(\chi) = \mathrm{Vor}(A')$. There is a bijection between the cells of $\mathrm{Vor}(\chi)$ and $\mathrm{Del}(\chi)$, denoted by mapping ν to $\nu^* \in \mathrm{Del}(\chi)$, such that $\dim \nu + \dim \nu^* = d + s$ and μ is a face of ν iff ν^* is a face of μ^*. We call ν *τ-colored* or *τ-colorful* if ν^* is τ-colored or τ-colorful, respectively. For each $\tau \subseteq \sigma$, we define the *τ-membrane* of χ as the union of the interiors of the τ-colorful cells of $\mathrm{Vor}(\chi)$, denoted $M(\tau)$. Since the interiors of the cells in $\mathrm{Vor}(\chi)$ partition $\mathbb{R}^{d+s}$, and the interior of each cell belongs to exactly one membrane, the membranes are pairwise disjoint and partition $\mathbb{R}^{d+s}$; see Figure 5.3.

Lemma 5.2.2 (Stratification into Membranes)**.** *Let $A \subseteq \mathbb{R}^d$ be finite, $\sigma = \{0, 1, \ldots, s\}$, and $\chi\colon A \to \sigma$ a coloring.*

1. *For each non-empty $\tau \subseteq \sigma$, $M(\tau)$ is a manifold homeomorphic to $\mathbb{R}^{s-t+d}$, with $t = \#\tau - 1$.*
2. *The collection of $M(\tau)$ forms a stratification of $\mathbb{R}^{d+s}$ with strata of dimension d to $d+s$, in which the p-stratum is the disjoint union of all $M(\tau)$ with $\#\tau = d+s-p+1$.*

Proof. We begin with $\tau = \sigma$. Let $w \in \mathbb{R}^d$ and consider $w + \mathbb{R}^s$, which is an s-plane parallel to $\mathbb{R}^s$ and therefore orthogonal to $\mathbb{R}^d$. By Pythagoras' theorem, the squared distance between points $x \in w + \mathbb{R}^s$ and $y \in \mathbb{R}^d$ is $\|x - w\|^2 + \|w - y\|^2$. Letting a be the point in A closest to x, this implies that a is the closest point in A to any point in $w + \mathbb{R}^s$. Similarly, if a'_j is the point in A'_j closest to x, then a'_j is the closest point in A'_j to any point in $w + \mathbb{R}^s$. There is a unique point $z(w) \in w + \mathbb{R}^s$ at equal distance to $a'_0, a'_1, \ldots, a'_s$. Hence, $z(w) \in M(\sigma)$ and it is indeed the only point of $w + \mathbb{R}^s$ in $M(\sigma)$. It follows that $M(\sigma)$ is the image of $z\colon \mathbb{R}^d \to \mathbb{R}^{d+s}$ defined by mapping w to $z(w)$. Note that z is continuous and its inverse is a projection, so $M(\sigma)$ is homeomorphic to $\mathbb{R}^d$. It is the stratum of the lowest dimension, d, in the claimed stratification.

To describe the remainder of the stratification, let $V(\sigma)$ be the Voronoi tessellation of $u_0, u_1, \ldots, u_s$ in $\mathbb{R}^s$. Since the u_j are the vertices of the standard s-simplex, this tessellation consists of a vertex at $0 \in \mathbb{R}^s$, $s+1$ half-lines emanating from 0, $\binom{s+1}{2}$ 2-dimensional wedges connecting the half-lines in pairs, etc. Returning to $w + \mathbb{R}^s$, we observe that it slices the stratification of $\mathbb{R}^{d+s}$ in a translate of this s-dimensional Voronoi tessellation, $z(w) + V(\sigma)$. Varying w over all points of $\mathbb{R}^d$, we get the claimed stratification of $\mathbb{R}^{d+s}$. □

How does $\mathrm{Del}(\chi)$ relate to $\mathrm{Del}(A)$? In the relatively straightforward simplicial case, $\mathrm{Del}(\chi)$ contains a subcomplex whose projection to $\mathbb{R}^d$ is $\mathrm{Del}(A)$; see Figure 5.4. In the general and therefore not necessarily simplicial case, we can for example have a convex quadrangle in $\mathrm{Del}(A)$ that is the projection of a tetrahedron in $\mathrm{Del}(\chi)$. We formulate the relationship that allows for this and similar cases in terms of the nerves of $\mathrm{Vor}(A)$ and $\mathrm{Vor}(\chi)$, which are possibly high-dimensional abstract simplicial complexes.

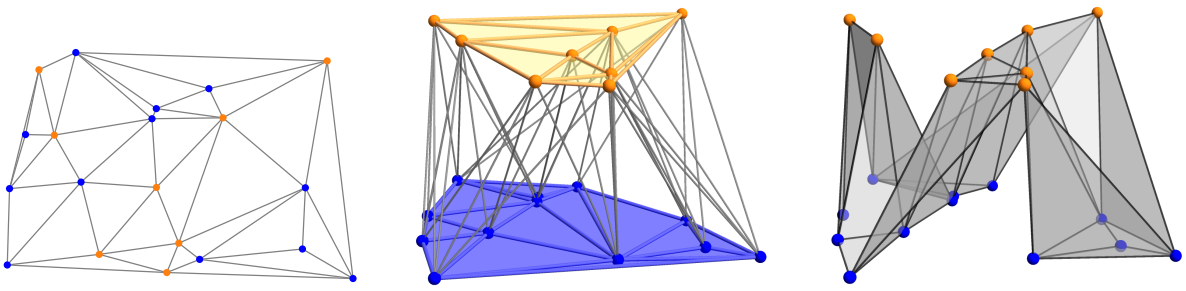

Figure 5.4: *Left:* the Delaunay mosaic of a bi-colored set in the plane, $\mathrm{Del}(A)$. *Middle:* the chromatic Delaunay mosaic, $\mathrm{Del}(\chi)$, with colorful triangles left unfilled for clarity. *Right:* the subcomplex of $\mathrm{Del}(\chi)$ that is isomorphic to $\mathrm{Del}(A)$.

Lemma 5.2.3 (Projection to Delaunay Mosaic)**.** *Let $A \subseteq \mathbb{R}^d$ be finite, $\sigma = \{0, 1, \ldots, s\}$, and $\chi\colon A \to \sigma$ a coloring. Then the nerve of the $(d+s)$-cells of $\mathrm{Vor}(\chi)$ in $\mathbb{R}^{d+s}$ has a subcomplex that projects to the nerve of the d-cells of $\mathrm{Vor}(A)$ in $\mathbb{R}^d$.*

Proof. Recall that $k+1$ points in A are the vertices of a cell in $\mathrm{Del}(A)$ iff there is an empty $(d-1)$-sphere, S, that passes through these $k+1$ points and through no other points of A. The nerve of the corresponding $k+1$ Voronoi d-cells is a k-simplex.

Following the construction of the chromatic Delaunay mosaic, we copy S to $u_j + S$ for each $j \in \sigma$. Let S' be the $(d+s-1)$-sphere in $\mathbb{R}^{d+s}$ whose intersection with $u_j + \mathbb{R}^d$ is $u_j + S$, for every $j \in \sigma$. It should be clear that S' exists: its center projected to $\mathbb{R}^s$ is the barycenter of the standard s-simplex and projected to $\mathbb{R}^d$ is the center of S. By construction, S' is empty and passes through the points $u_j + a$ with $a \in S$ and $\chi(a) = j$, and through no other points of A'. The nerve of the corresponding $(d+s)$-cells in $\mathrm{V}or(\chi)$ is again isomorphic to a k-simplex, and its projection to $\mathbb{R}^d$ is the k-simplex isomorphic to the nerve of the $k+1$ Voronoi d-cells we started with. The claim follows. □

5.2.2 Overlay of Mono-chromatic Voronoi Tessellations

Related to the strata are the overlays of tessellations. Given $A \subseteq \mathbb{R}^d$, $\sigma = \{0, 1, \ldots, s\}$, and $\chi\colon A \to \{0, 1, \ldots, s\}$, the *overlay* of the $s+1$ mono-chromatic Voronoi tessellations, denoted $\mathrm{V}or(A_j \mid j \in \sigma)$, is the decomposition of $\mathbb{R}^d$ obtained by drawing the Voronoi cells of dimension at most $d-1$ on top of each other; see Figure 5.5. More formally, each d-dimensional cell in the overlay is the common intersection of $s+1$ d-cells, one in each $\mathrm{V}or(A_j)$ for $j \in \sigma$, and the overlay consists of these d-dimensional cells and their faces. Even if the points in A are in general position, the overlay is not necessarily a simple decomposition of $\mathbb{R}^d$.

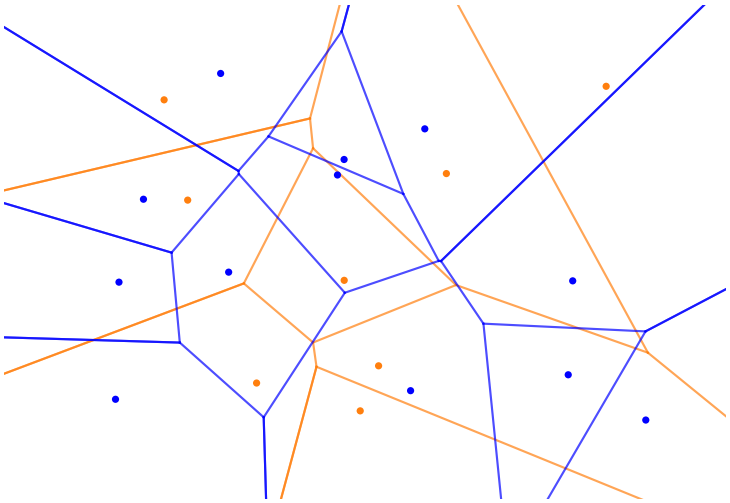

Figure 5.5: The overlay of a blue and a orange Voronoi tessellation in the plane. In the generic case, each of its vertices is either a vertex of a mono-chromatic tessellation, which has degree 3, or the crossing of two edges, which has degree 4.

Lemma 5.2.4 (Membranes and Overlays)**.** *Let $A \subseteq \mathbb{R}^d$ be finite, $\chi\colon A \to \{0, 1, \ldots, s\}$ a coloring, and $A_j = \chi^{-1}(j)$ for $0 \leq j \leq s$. For each $\tau \subseteq \sigma$, $Vor(A_j \mid j \in \tau)$ is the projection of the τ-membrane, $M(\tau)$, to $\mathbb{R}^d$.*

Proof. We begin with $\tau = \sigma$. By Lemma 5.2.2, $M(\sigma)$ is a manifold of dimension d, and in the proof of this lemma we learn that the orthogonal projection, $\pi\colon M(\sigma) \to \mathbb{R}^d$, is a homeomorphism. Indeed, π^{-1} is the restriction of $z\colon \mathbb{R}^d \to \mathbb{R}^{d+s}$ defined there. Since $M(\sigma)$ is decomposed into cells of $\mathrm{V}or(\chi)$, z is piecewise linear, so it suffices to prove that the linear pieces are the images of the cells in $\mathrm{V}or(A_j \mid j \in \sigma)$.

Let ν_j be a d-cell of $\mathrm{V}or(A_j)$ and write $a_j \in A_j$ for the point that generates ν_j, for $0 \leq j \leq s$. Assume that $\nu = \nu_0 \cap \nu_1 \cap \ldots \cap \nu_s$ has non-empty interior, so it is a d-cell of the overlay.

Correspondingly, the image of every point $x \in \nu$, $z(x) = \pi^{-1}(x)$, is equidistant from the points $u_j + a_j$, for $0 \leq j \leq s$. It follows that the image of ν is a subset of a linear piece in $M(\sigma)$. For every neighboring d-cell of ν in the overlay, we change one of the a_j, so their images belong to different linear pieces of $M(\sigma)$. This implies that the image of ν *is* a linear piece of $M(\sigma)$, as required.

To generalize, let $\tau \subseteq \sigma$ and use the above argument to conclude that $\mathrm{V}or(A_j \mid j \in \tau)$ is the projection of the τ-membrane to $\mathbb{R}^d$. Recall that $\mathrm{V}or(\chi|\tau)$ decomposes $\mathbb{R}^{t+d}$, and by Lemma 5.2.1, the extrusion of the τ-membrane in $\mathrm{V}or(\chi|\tau)$ along the remaining $s-t$ coordinate directions in $\mathbb{R}^{d+s}$ contains the τ-membrane in $\mathrm{V}or(\chi)$. Moreover, the projections of the two τ-membranes—one in $\mathrm{V}or(\chi|\tau)$ and the other in $\mathrm{V}or(\chi)$—to $\mathbb{R}^d$ are identical, which implies the claim. □

5.2.3 Chromatic Alpha Complexes

A direct analogy of the characterization of the Delaunay complex with empty spheres is the characterization of the chromatic Delaunay complex with what we call empty stacks. A *σ-stack* in $\mathbb{R}^d$ is a collection of $s+1$ concentric $(d-1)$-spheres, one for each color in σ; see Figure 5.6. We drop σ from the notation if it is clear from the context. The *radius* of the stack is the maximum radius of its spheres, and its *center* is the common center of the spheres. We label the spheres S_j, $j \in \sigma$, and say the stack is *empty* if S_j is empty of points in $A_j = \chi^{-1}(j)$, for each $j \in \sigma$. We say the stack *passes through* $\nu \subseteq A$ if S_j passes through all the points of $\nu \cap A_j$, for each color $j \in \sigma$.

Lemma 5.2.5. *Let $\chi\colon A \to \sigma$ be a chromatic point set in $\mathbb{R}^d$, write $A_j = \chi^{-1}(j)$, and let $\nu \subseteq A$ be a collection of points. Then $\nu \in Del(\chi)$ iff there exists an empty stack of spheres that passes through ν.*

Proof. Let $\nu_j = \nu \cap A_j$ be the j-colored points in ν, for each $j \in \sigma$. The existence of an empty sphere, S_j, with center x that passes through ν_j is equivalent to x being in the intersection of the corresponding Voronoi domains: $x \in \bigcap_{a \in \nu_j} \mathrm{d}om(a, A_j)$. Therefore, there exists an empty stack passing through ν centered at x iff $x \in \bigcap_{a \in \nu_j} \mathrm{d}om(a, A_j)$ for each $j \in \sigma$. This is the defining property of ν being in $\mathrm{D}el(\chi)$, namely that $\bigcap_{a \in \nu} \mathrm{d}om(a, A_{\chi(a)})$ is non-empty. □

Like in the mono-chromatic setting, we define chromatic alpha complexes as sublevel sets of the radius function defined on the chromatic Delaunay complex. We recall that the radius of a stack is the radius of its largest sphere.

Definition 5.2.1. *Let $\chi\colon A \to \sigma$ be a chromatic point set, and $Rad\colon Del(\chi) \to \mathbb{R}$ the* radius function *defined by mapping $\nu \in Del(\chi)$ to the radius of the smallest empty stack that passes through ν. The* chromatic alpha complex *of χ with radius $r \in \mathbb{R}$ is $Alf_r(\chi) = Rad^{-1}[0, r]$.*

Using the empty spheres and the empty stacks characterizations (Lemma 5.2.5), we can see a clear relation between alpha complexes and chromatic alpha complexes. If there exists an empty $(d-1)$-sphere, S, of radius r passing through $\nu \subseteq A$, then there also exists an empty stack of radius r passing through ν. Indeed, we can take $S_j = S$ for each $j \in \sigma$. Similarly, an empty sphere, S, that passes through points $\nu \subseteq A_j$ of the same color is itself an empty stack when we set S_i to be a sphere with zero radius for $i \neq j$. However, the same simplex can have a different radius in $\mathrm{D}el(A)$ and in $\mathrm{D}el(\chi)$: the smallest empty sphere can have strictly larger radius than the smallest empty stack passing through the same points; see Figure 5.7.

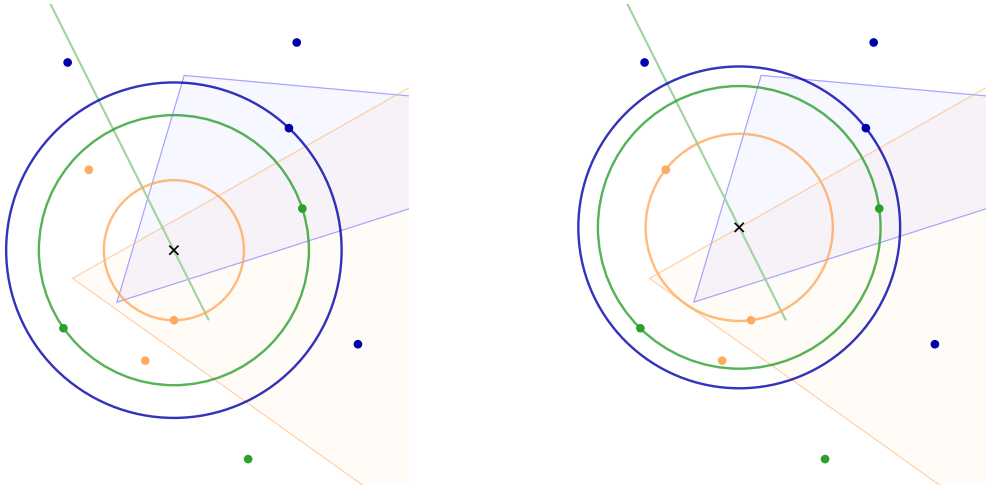

Figure 5.6: Two empty stacks in $\mathbb{R}^2$ that pass through one blue point, two green points, and one orange point forming a simplex $\nu \in \mathrm{D}el(\chi)$. (In fact, the stack on the *right* passes through *two* orange points, so it also passes through the *one* orange point that lies on the *left* orange circle.) The set of centers of all empty stacks that pass through these four points is the intersection of three Voronoi cells: a blue 2-cell, a green 1-cell, and an orange 2-cell. The *right panel* shows the smallest empty stack in this collection: its center lies on the boundary of the intersection of Voronoi cells, which is the reason why one of its circles passes through an extra point.

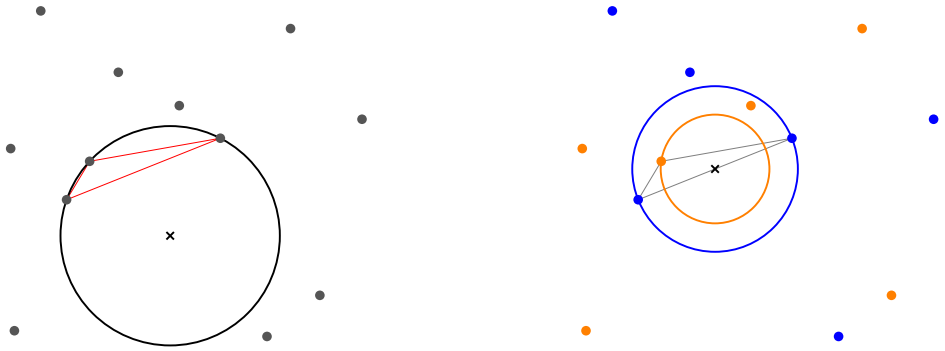

Figure 5.7: An obtuse triangle with two blue points and an orange point at the obtuse angle. On the *left:* the smallest empty sphere that passes through the three points. It has strictly larger radius than the smallest empty stack that passes through the three points, which is shown on the *right.* Therefore, the triangle belongs to both, the Delaunay complex and the chromatic Delaunay complex, but it has a different value in the two radius functions.

An important reason why the alpha complex is useful in the mono-chromatic setting is its correspondence to the union of balls growing from the input points. From the topological point of view, studying the growing union of balls is equivalent to studying the growing alpha complex. In the following, we draw an analogous connection for chromatic alpha complexes. One important distinction is that there is more structure to be preserved in the chromatic setting: not only the topological spaces themselves, but also how they are related to each other. For example, for a bi-chromatic point set as in Figure 5.8, we study the inclusion of the union of the blue disks into the union of all disks. We prove that we can equivalently study the inclusions of the blue alpha complex into the chromatic alpha complex.

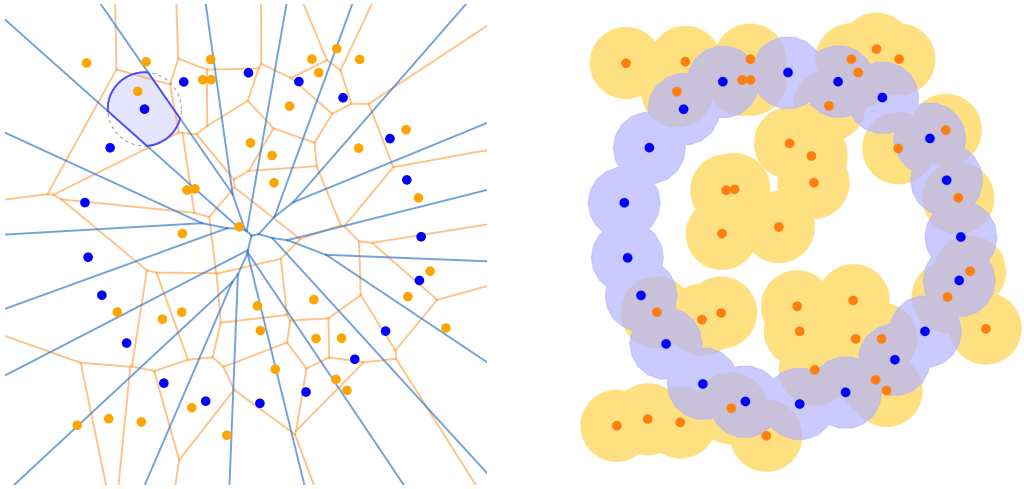

Figure 5.8: On the *left*: a chromatic set together with Voronoi tessellations of the blue and orange points overlaid, and one chromatic Voronoi ball highlighted. On the *right*: the union of blue disks and the union of orange disks; we study, e.g., the inclusion of the blue area into the union of all the disks.

For a point $a \in A$ in a chromatic set $\chi\colon A \to \sigma$, we define its *(chromatic) Voronoi ball* of radius r as the intersection of the ball of radius r with the Voronoi domain within its color class:

$$B'_r(a,\chi) = B'_r(a, A_{\chi(a)}) = B_r(a) \cap dom(a, A_{\chi(a)}). \tag{5.1}$$

Let $\nu \subseteq A$ be a set of points. Like in the proof of Lemma 5.2.5, we observe that x is the center of an empty stack of radius r passing through ν iff x is contained in the intersection of the Voronoi balls of radius r centered at the points in ν; that is: $x \in \bigcap_{a\in\nu} B'_r(a,\chi)$. This implies that $\mathrm{A}lf_r(\chi)$ is isomorphic to the nerve of all Voronoi balls $B'_r(a,\chi)$, $a \in A$. Since the union of the Voronoi balls is the same as the union of the balls, the Nerve Theorem yields the following:

Lemma 5.2.6. *$Alf_r(\chi) \simeq Alf_r(A)$, and both are homotopy equivalent to the union of balls, $\bigcup_{a\in A} B_r(a)$.*

We are currently studying the structural properties of the radius function $\mathrm{R}ad\colon \mathrm{D}el(\chi) \to \mathbb{R}$ as well as algorithms for its efficient computation. A second version of our paper [25] will contain these results, which informally spoil here. Under a suitable notion of general position, the chromatic radius function is generalized discrete Morse, extending the result of Bauer and Edelsbrunner in the monochromatic case 2.2.1.

5.3 Combinatorial Bounds

In this section, we are interested in the size of the chromatic Delaunay mosaic or, equivalently, of the overlays between the mono-chromatic Voronoi tessellations. We study extremal questions, in which we maximize over the point sets and their colorings, but we also consider random colorings. Refer to Table 5.1 for a summary of our results. We start by explaining how we get the bound in row 1, column 1.

We aim to give a bound on the size of any $\mathrm{D}el(\chi)$. Consider first the case in which there are equally many colors and dimensions: $d = s + 1$. For the lower bound, we assign each color to

about n/d points, and we place the points with color j in sequence on the j-th coordinate axis. This gives a constant times n^d colorful crossings. For the upper bound, we use the general bound on the number of simplices in a Delaunay mosaic of n (uncolored) points in $\mathbb{R}^{d+s}$ given in [20], which for $d+s=2d-1$ gives $O(n^d)$.

We get n^d as an upper bound also for the case in which there are more colors than dimensions: $d<s+1$. To see this, note that any crossing involves at most d colors, so multiplying the bound for $d=s+1$ with $\binom{s+1}{d}$, which is a constant, suffices.

This leaves the case $d>s+1$. Here we use the *moment curve*, which in $\mathbb{R}^d$ is the set of points $(t,t^2,\ldots,t^d)$, $t\in\mathbb{R}$. The upper bound is constant times $n^{\lceil (d+s)/2\rceil}$ again from the general case of n points in $\mathbb{R}^{d+s}$. For the lower bound, we assign each color to about $n/s+1$ points, place the points of the first color on the moment curve in $\mathbb{R}^{d-s}$ (which we assume is spanned by the first $d-s$ coordinate axes of $\mathbb{R}^d$), and the points of each other color on one of the s remaining coordinate axes. We get a constant times $n^{\lceil (d-s)/2\rceil}$ vertices in the Voronoi tessellation of the first color within $\mathbb{R}^{d-s}$. Each such vertex expands to the orthogonal s-plane in the Voronoi tessellation of the first color within $\mathbb{R}^d$, and this s-plane intersects the grid formed by the other s colors in a constant times n^s crossings. The total number of crossings is therefore a constant times $n^{\lceil (d-s)/2\rceil}\cdot n^s = n^{\lceil (d+s)/2\rceil}$, as required.

5.3.1 Few Spherical *k*-sets Imply Small Expected Overlays

Let A be a set of n points in $\mathbb{R}^d$. We call a subset of $k\leq n$ points a *spherical k-set* of A if there is a sphere that separates the k from the remaining $n-k$ points. Note that this differs from the classic notion of a k-set, for which there is a hyperplane that separates the k points of the k-set from the remaining $n-k$ points. In this section, we relate the number of spherical k-sets with the expected size of the overlay of mono-chromatic Voronoi tessellations for random colorings of A. Specifically, we prove the following lemma.

Lemma 5.3.1 (Spherical k-sets and Overlay)**.** *Let c,d,e be positive constants, and A a set of n points in $\mathbb{R}^d$ such that for every $1\leq k\leq n$, the number of spherical k-sets is $O(k^c n^e)$. Let furthermore $s\geq 0$ be a constant, let $\sigma=\{0,1,\ldots,s\}$, and write $A_j=\chi^{-1}(j)$, in which $\chi\colon A\to\sigma$ is a random coloring. Then the expected size of $Vor(A_j \mid j\in\sigma)$ is $O(n^e)$.*

Proof. We assume that the points in A are in general position and write $A_j=\chi^{-1}(j)$. Suppose we pick $s+1$ cells, one from each $\mathrm{V}or(A_j)$, and write i_j-1 for their co-dimensions. The common intersection of the $s+1$ cells is either empty or a cell of co-dimension $\sum_{j=0}^{s}(i_j-1)$. This is a vertex only if $\sum_{j=0}^{s}(i_j-1)=d$ or, equivalently, $\sum_{j=0}^{s} i_j = d+s+1$. To bound the expected size of the overlay, we bound the expected number of such vertices in the overlay, which is the sum of their probabilities to belong to the overlay. Below we argue that each spherical k-set can give rise to only a limited number of vertices, and we give a bound on the probability of any of them to appear in a random coloring.

Fix any $d+s+1$ points from A and a coloring $\chi\colon A\to\{0,1,\ldots,s\}$ such that every color is assigned to at least one of these $d+s+1$ points. Writing i_j for the number of points with color j, we have $\sum_{j=0}^{s} i_j = d+s+1$ and $1\leq i_j\leq d+1$ for each j. Let E_j be the set of points $y\in\mathbb{R}^d$ at equal distance to the i_j points with color j; it is a plane of co-dimension i_j-1. Since $\sum_{j=0}^{s}(i_j-1)=d$ and the $d+s+1$ points are in general position, the common intersection of the E_j is a point $x\in\mathbb{R}^d$. This point is a vertex of the overlay iff there is a stack of spheres, $S_0,S_1,\ldots,S_s$, with common center, x, such that S_j passes through the i_j

points with color j, and all other points in $A_j = \chi^{-1}(j)$ lie outside S_j. Suppose that S_0 is the largest of the $s+1$ spheres. Let k be the number of points on or inside S_0, note that this is a spherical k-set, and recall that there are at most $O(k^c n^e)$ such sets by assumption. Other than the $i_0 \leq d+1$ points on S_0, all points in the spherical k-set must have color different from 0. The probability of this is $s/(s+1)$ to the power $k - i_0 \geq k - d - 1$. The number of possible overlay vertices whose largest sphere of the corresponding stack of spheres separates the same spherical k-set is at most $\binom{k}{d+s+1}(s+1)^{d+s+1}$. This is the product of the number of subsets of size $d+s+1$ and the number of different colorings of such a set. Writing X for the number of vertices in the overlay, we thus get

$$\mathbb{E}[X] < \sum_{k=d+s+1}^{n} O(k^c n^e) \cdot \binom{k}{d+s+1}(s+1)^{d+s+1} \cdot \left(\frac{s}{s+1}\right)^{k-d-1} \tag{5.2}$$

$$< O(n^e) \cdot \sum_{k=0}^{\infty} \frac{(s+1)^{2d+s+2}}{s^{d+1}} \cdot k^{c+d+s+1} \cdot \left(\frac{s}{s+1}\right)^{k}. \tag{5.3}$$

The first factor within the latter sum is constant, the second is a constant degree polynomial, and the last factor is an exponential that vanishes as k goes to infinity. Because of the exponential decay, the sum converges to a constant that depends on c, d, and s but not on n. It follows that the number of vertices in the overlay is $O(n^e)$.

Observe that every vertex of the overlay belongs to only a constant number of cells of dimension 1 to d. Every such cell has at least one vertex, which implies that the number of cells of any dimension in the overlay is $O(n^e)$. □

As originally proved by Lee [73], the number of spherical k-sets of n points in $\mathbb{R}^2$ is less than $2kn$. The expected size of the overlay of the mono-chromatic Voronoi tessellations for a random coloring in $\mathbb{R}^2$ is therefore $O(n)$. To get a result for general dimensions, we note that the spherical k-sets in $\mathbb{R}^d$ correspond to (linear) k-sets in $\mathbb{R}^{d+1}$ via lifting to a paraboloid. For the latter, Clarkson and Shor [20] proved that the number of ℓ-sets, for $\ell = 1, 2, \ldots, k$, is $O(k^{\lceil (d+1)/2 \rceil} n^{\lfloor (d+1)/2 \rfloor})$. Lemma 5.3.1 thus implies the following theorem.

Theorem 5.3.1 (Overlay Size for Random Coloring)**.** *Let d and s be constants, let A be a set of n points in $\mathbb{R}^d$, let $\sigma = \{0, 1, \ldots, s\}$, and write $A_j = \chi^{-1}(j)$, in which $\chi\colon A \to \sigma$ is a random coloring. Then the expected number of cells in $Vor(A_j \mid j \in \sigma)$ is $O(n^{\lceil d/2 \rceil})$.*

This bound is asymptotically tight since even a single Voronoi tessellation of n points in $\mathbb{R}^d$ can have $\Omega(n^{\lceil d/2 \rceil})$ vertices, for example if the points are chosen on the moment curve in $\mathbb{R}^d$.

5.3.2 Delone Sets Have Small Expected Overlays

We start by showing that packed sets without big holes have few spherical k-sets. To formalize this claim, we recall that $A \subseteq \mathbb{R}^d$ is a *Delone set* if there are constants $0 < r < R < \infty$ such that every open ball of radius r contains at most one point of A, and every closed ball of radius R contains at least one point of A. The proof of the claim makes use of an extension of Voronoi tessellations to higher order. To define it, let $B \subseteq A$ and write $\mathrm{dom}(B, A)$ for the points $x \in \mathbb{R}^d$ that satisfy $\|x - b\| \leq \|x - a\|$ for all $b \in B$ and all $a \in A \setminus B$. Now fix an integer $k \geq 1$ and note that the cells $\mathrm{dom}(B, A)$ with $\#B = k$ cover the entire $\mathbb{R}^d$. The collection of the $\mathrm{dom}(B, A)$ with $\#B = k$ is referred to as the *order-k Voronoi tessellation* of A, denoted $\mathrm{Vor}_k(A)$. Note that $\mathrm{Vor}_1(A) = \mathrm{Vor}(A)$. Note also that there is a non-empty cell, $\mathrm{dom}(B, A)$, in $\mathrm{Vor}_k(A)$ iff B is a spherical k-set. For counting purposes, we say a spherical

k-set, $B \subseteq A$, *corresponds* to a point, $b \in A$, if there is a sphere that separates B from $A \setminus B$ and b is the point in A closest to the center of this sphere. Since the separating sphere is generally not unique, B may correspond to more than one point in A.

Lemma 5.3.2 (Bounded Correspondence)**.** *Let $A \subseteq \mathbb{R}^d$ be a Delone set. Then every point in A corresponds to at most $O(k^{d+1})$ spherical k-sets of A.*

Proof. Let x be a point in $\mathbb{R}^d$ and suppose that it lies in the interior of a d-cell of the order-k Voronoi tessellation of A. Assuming this cell is $\mathrm{dom}(B, A)$, then B is the unique spherical k-set that is separated from $A \setminus B$ by a sphere with center x. Letting t be the radius of one such sphere, we have $kr^d \leq (t+r)^d$ because the sphere with center x and radius $t+r$ encloses k disjoint open balls of radius r. Furthermore, $(t-R)^d \leq kR^d$ because the closed balls of radius R centered at the points of B cover the ball with center x and radius $t-R$. Hence

$$(\sqrt[d]{k}-1)r \leq t \leq (\sqrt[d]{k}+1)R. \tag{5.4}$$

Let $b \in A$ be a point with $x \in \mathrm{dom}(b, A)$. Since A is Delone, $\mathrm{dom}(b, A)$ is covered by the ball with center b and radius R. It follows that the sphere with center b and radius $(\sqrt[d]{k}+2)R$ encloses all spherical k-sets that correspond to b. The number of points in A enclosed by this sphere satisfies

$$\ell \leq \left[\left(\sqrt[d]{k}+2\right)R + r\right]^d / r^d, \tag{5.5}$$

which is $O(k)$ because r and R and therefore R/r are positive constants. For a finite set in $\mathbb{R}^d$, the number of ways it can be split into two by a sphere is less than the $(d+1)$-st power of its cardinality. Hence, there are at most $O(k^{d+1})$ spherical k-sets that correspond to b. □

Delone sets are necessarily infinite, so we let Ω be the closed ball with radius ω centered at the origin, and count a spherical k-set, $B \subseteq A$, only if there is a sphere that separates B from $A \setminus B$ whose center is in Ω.

Theorem 5.3.2 (Overlay Size for Delone Set)**.** *Let d and s be constants, let $A \subseteq \mathbb{R}^d$ be a Delone set, let $\sigma = \{0, 1, \ldots, s\}$, let $\chi\colon A \to \sigma$ be a random coloring, and let Ω be the ball of points at distance at most $\omega > R$ from the origin. Writing $n = \#(A \cap \Omega)$ and $A_j = \chi^{-1}(j)$, the expected number of cells in $Vor(A_j \mid j \in \sigma)$ that have at least one vertex in Ω is $O(n)$.*

Proof. Let $0 < r < R < \infty$ be constants for which A is Delone, and note that the number of points of A at distance at most $\omega + R$ from the origin is $O(n)$. Any spherical k-set that has a separating sphere with center in Ω corresponds to a point in this slightly larger ball, so Lemma 5.3.2 implies that the number of such spherical k-sets is $O(k^{d+1}n)$.

We count the vertices of the overlay using Lemma 5.3.1 but restricted to crossings inside Ω. We have $c = d+1$ and $e = 1$, so we get an expected number of $O(n)$ vertices in Ω. Assuming general position, every vertex belongs to only a constant number of cells, which implies the claimed bound on the number of cells with at least one vertex in Ω. □

A vertex of the overlay corresponds to an $(d+s)$-cell in the chromatic Delaunay mosaic whose circumcenter project to the vertex. Theorem 5.3.2 thus counts the cells in the chromatic Delaunay mosaic that are faces of $(d+s)$-cells whose circumcenters project into Ω.

5.4 Persistent Homology of Chromatic Alpha Complexes

In this section, we review the background needed to turn the chromatic alpha complexes into persistence diagrams, and we advocate the use of six such diagrams, which we refer to as a 6-pack. In addition, we discuss the relations between the diagrams in a 6-pack, and how to go beyond a 6-pack if the interest in in more than mutual interactions.

5.4.1 Kernel, Image, Cokernel Persistence

The goal of this subsection is to introduce the framework of persistent homology [43], together with its kernel, image, and cokernel generalizations [23]. We keep the formalism to a minimum by limiting ourselves to simplicial complexes and $\mathbb{Z}/2\mathbb{Z}$ coefficients.

Homology with $\mathbb{Z}/2\mathbb{Z}$ Coefficients

Loosely speaking, homology is an algebraic framework that defines and counts holes in a shape. Given a simplicial complex, K, a *p-chain* is a subset of p-simplices. The *sum* of two p-chains is the symmetric difference of the two sets: if a p-simplex belongs to both chains, the two copies erase each other, as $1 + 1 = 0$ in modulo-2 arithmetic. The *boundary* of a p-simplex is the set of $(p-1)$-dimensional faces, which is a $(p-1)$-chain. The p-chains with the sum operation form a group, $\mathsf{C}_p(K)$, and the *boundary operator*, $\partial_p \colon \mathsf{C}_p(K) \to \mathsf{C}_{p-1}(K)$, maps a p-chain to the sum of its simplices' boundaries. A *p-cycle* is a p-chain with empty boundary, a *filling* of this p-cycle is a $(p+1)$-chain whose boundary is the p-cycle, and a *p-boundary* is a p-cycle for which there exits a filling. The p-boundaries and the p-cycles form groups by themselves, and since every p-boundary is a p-cycle, and every p-cycle is a p-chain, we get three nested groups: $\mathsf{B}_p(K) \subseteq \mathsf{Z}_p(K) \subseteq \mathsf{C}_p(K)$. Two p-cycles are *homologous* if their sum has a filling or, equivalently, adding a p-boundary to one p-cycle gives the other p-cycle. Being homologous is an equivalence relation, whose equivalence classes are the elements of the *p-th homology group*: $\mathsf{H}_p(K) = \mathsf{Z}_p(K)/\mathsf{B}_p(K)$. All mentioned groups are vector spaces, so the *ranks* are their dimensions. Of particular importance is the *p-th Betti number* of K, which is the rank of the p-th homology group: $rank\, \mathsf{H}_p(K) = rank\, \mathsf{Z}_p(K) - rank\, \mathsf{B}_p(K)$.

Let L be a subcomplex of K. Relative homology describes the connectivity of the topological pair (K, L), which geometrically represents K with the subspace L identified as a single point. The chain groups are the quotients $C_p(K, L) = C_p(K) \big/ C_p(L)$. Cycle and boundary subgroups are defined as before, and their quotients are the *relative homology groups of the pair*, denoted $H_p(K, L)$. The homology groups and their relative cousins are related by the following long exact sequence:

$$\ldots \to \mathsf{H}_p(L) \to \mathsf{H}_p(K) \to \mathsf{H}_p(K, L) \to \mathsf{H}_{p-1}(L) \to \ldots \tag{5.6}$$

A well known basic property of long exact sequences is that the alternating sum of dimensions of the vector spaces vanishes.

Lemma 5.4.1. *Let $L_i \subseteq K_i$ be simplicial complexes. Then*

$$\textstyle\sum_{p \in \mathbb{Z}} (-1)^p [rank\, \mathsf{H}_p(L_i) - rank\, \mathsf{H}_p(K_i) + rank\, \mathsf{H}_p(K_i, L_i)] = 0. \tag{5.7}$$

Proof. By definition of exactness, the rank of each homology group in (5.6) can be written as the sum of two non-negative integers such that it shares one with the preceding group and the other with the succeeding group along the sequence. Since only finitely many groups have non-zero ranks, this implies that the alternating sum of ranks vanishes. □

Persistent Homology

In the following, let $f\colon K \to \mathbb{R}$ be monotonic, with values $r_1 < r_2 < \ldots < r_n$, and let $K_i = f^{-1}(-\infty, r_i]$ be its i-th sublevel set. Applying the p-th homology functor to $\emptyset = K_0 \subseteq K_1 \subseteq \ldots \subseteq K_n$, we get a sequence of vector spaces:

$$\mathsf{H}_p(K_0) \to \ldots \to \mathsf{H}_p(K_{i-1}) \to \mathsf{H}_p(K_i) \to \ldots \to \mathsf{H}_p(K_{j-1}) \to \mathsf{H}_p(K_j) \to \ldots \to \mathsf{H}_p(K_n). \tag{5.8}$$

There is one such sequence for each dimension, p. The inclusions $K_i \subseteq K_j$ induce maps $\mathsf{f}_i^j\colon \mathsf{H}_p(K_i) \to \mathsf{H}_p(K_j)$ for all $0 \leq i \leq j \leq n$. This sequence is called a *persistence module*. It can be written as a direct sum of indecomposable modules of the form $\ldots \to 0 \to \mathsf{k} \to \ldots \to \mathsf{k} \to 0 \to \ldots$, where $\mathsf{k} = \mathbb{Z}/2\mathbb{Z}$, all maps between these 1-dimensional vector spaces are identities, and all others are zero maps. Each indecomposable module has a concrete interpretation, namely a *birth* followed by a *death* of a homology class. Specifically, we have such an indecomposable module from position i to position $j-1$ if

- there is a class, $\gamma \in \mathsf{H}_p(K_i)$ that does not belong to the image of f_{i-1}^i, and
- $\mathsf{f}_i^{j-1}(\gamma)$ does not belong to the image of f_{i-1}^{j-1}, but $\mathsf{f}_i^j(\gamma)$ belongs to the image of f_{i-1}^j.

We say γ is *born* at K_i and *dies entering* K_j. We record this information with the point $(f(r_i), f(r_j))$, noting that the second coordinate is ∞ if the class is born but never dies. The resulting multi-set of points in the extended plane is the *p-th persistence diagram* of f, denoted $\mathsf{D}gm_p(f)$. Sometimes, we drop the index and write $\mathsf{D}gm(f)$ for the disjoint union of the $\mathsf{D}gm_p(f)$ over all dimensions, p.

If L is a subcomplex of K, we get a filtration, L_i, by restricting f to L. The inclusions of the pairs $(K_0, L_0) \subseteq (K_1, L_1) \subseteq \ldots \subseteq (K_n, L_n)$ give rise to a sequence of relative homology groups,

$$\mathsf{H}_p(K_0, L_0) \to \ldots \to \mathsf{H}_p(K_{i-1}, L_{i-1}) \to \mathsf{H}_p(K_i, L_i) \to \ldots \to \mathsf{H}_p(K_n, L_n). \tag{5.9}$$

Applying the above definitions to this sequence yields the *p-th relative persistent diagram*.

An important property of the persistence diagram is its stability. Specifically, the bottleneck distance between the diagrams of $f, g\colon K \to \mathbb{R}$ is bounded from above by the L_∞-distance between the two maps:

$$W_\infty(\mathsf{D}gm_p(f), \mathsf{D}gm_p(g)) \leq \|f - g\|_\infty; \tag{5.10}$$

see [21]. The *persistence* of a point in the persistence diagram is the vertical distance to the diagonal, $|f(r_j) - f(r_i)|$, and the *1-norm* is the sum of persistences of the points in the diagram, denoted $\|\mathsf{D}gm(f)\|_1$. To cope with points at infinity, we use a *cut-off*, C, which we effectively substitute for ∞. This gives finite 1-norms and preserves relationships implied by exact sequences, as expressed in Theorem 5.4.1.

Kernels, Images, and Cokernels

Let $L \subseteq K$ be simplicial complexes, $f_K \colon K \to \mathbb{R}$ monotonic, and $f_L \colon L \to \mathbb{R}$ the restriction of f_K to L. Taking sublevel sets, we get two parallel persistence modules and maps from one module to the other:

$$\begin{array}{ccccccccccc} \mathsf{H}_p(K_0) & \to & \ldots & \to & \mathsf{H}_p(K_i) & \to & \mathsf{H}_p(K_{i+1}) & \to & \ldots & \to & \mathsf{H}_p(K_n) \\ \uparrow & & \ldots & & \uparrow & & \uparrow & & \ldots & & \uparrow \\ \mathsf{H}_p(L_0) & \to & \ldots & \to & \mathsf{H}_p(L_i) & \to & \mathsf{H}_p(L_{i+1}) & \to & \ldots & \to & \mathsf{H}_p(L_n). \end{array} \tag{5.11}$$

Write $\kappa_i \colon \mathsf{H}_p(L_i) \to \mathsf{H}_p(K_i)$ for the vertical maps, which are induced by the inclusions $L_i \subseteq K_i$, for $0 \leq i \leq n$. These maps have kernels, images, and cokernels, which form persistence modules of their own:

$$\begin{array}{ccccccccccc} \mathrm{k}er_p\,\kappa_0 & \to & \ldots & \to & \mathrm{k}er_p\,\kappa_i & \to & \mathrm{k}er_p\,\kappa_{i+1} & \to & \ldots & \to & \mathrm{k}er_p\,\kappa_n, \\ \mathrm{i}m_p\,\kappa_0 & \to & \ldots & \to & \mathrm{i}m_p\,\kappa_i & \to & \mathrm{i}m_p\,\kappa_{i+1} & \to & \ldots & \to & \mathrm{i}m_p\,\kappa_n, \\ cok_p\,\kappa_0 & \to & \ldots & \to & cok_p\,\kappa_i & \to & cok_p\,\kappa_{i+1} & \to & \ldots & \to & cok_p\,\kappa_n. \end{array} \tag{5.12}$$

These persistence modules were introduced and studied in [23]. Following the notation in that paper, we write $\mathsf{D}gm(\mathrm{k}er\ f_L \to f_K)$, $\mathsf{D}gm(\mathrm{i}m\ f_L \to f_K)$, and $\mathsf{D}gm(cok\ f_L \to f_K)$ for the corresponding persistence diagrams. These diagrams are also stable under perturbations of the monotonic function, and they can be computed efficiently. We omit details and refer to [23], in particular for the matrix reduction algorithms, which we have implemented [39] to study the meaning of these derived persistence diagrams for chromatic point sets.

5.4.2 6-pack of Persistent Diagrams

The main new concept in this section is a collection of six related persistence diagrams, which we use to quantify the way different point sets mingle. We call this collection a 6-*pack*. A 6-pack can be defined for any pair of topological spaces $L \subseteq K$ with a filtration on K. We explain the construction on a concrete example: the blue circle on an orange background in Figure 5.9. Let $K = \mathrm{D}el(\chi)$ be the chromatic Delaunay complex of the portrayed chromatic set, and let $L \subseteq K$ be the blue subcomplex, consisting of those simplices in K that only have blue vertices. Let $f_K \colon K \to \mathbb{R}$ be the chromatic radius function, and write f_L and $f_{K,L}$ for the restrictions of f_K to L and $K \setminus L$. The radius function and its restrictions give rise to three persistence modules, and we get three additional persistence modules for the kernel, the image, and the cokernel of the map on homology induced by the inclusion $L \subseteq K$. The persistence diagrams in the 6-pack are arranged as in Table 5.2, in a manner that lends itself to comparing the information between them.

kernel: $\mathsf{D}gm(\mathrm{k}er\ f_L \to f_K)$	relative: $\mathsf{D}gm(f_{K,L})$	cokernel: $\mathsf{D}gm(cok\ f_L \to f_K)$
domain: $\mathsf{D}gm(f_L)$	image: $\mathsf{D}gm(\mathrm{i}m\ f_L \to f_K)$	codomain: $\mathsf{D}gm(f_K)$

Table 5.2: The arrangement of the persistence diagrams in the 6-pack for the pair $L \subseteq K$ in two rows and three columns. Read the six positions in a circle so that the diagram of the domain lies between those of the kernel and the image, the diagram of the image lies between those of the domain and the codomain, etc.

Figure 5.10 displays the 6-pack for the pair $L \subseteq K$ in the left panel of Figure 5.9. Not surprisingly, the circle of blue points gives rise to a persistent 1-cycle in L captured in the diagram of the domain. At the time of its birth, this 1-cycle includes into a non-trivial 1-cycle

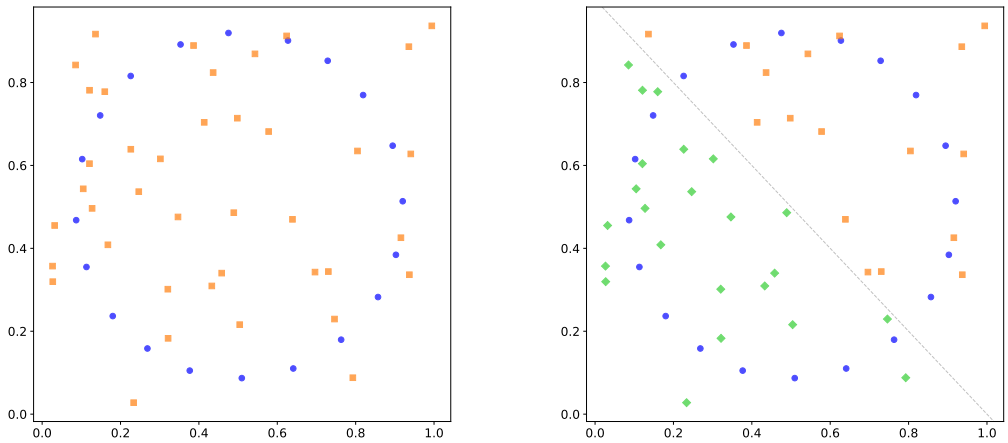

Figure 5.9: A bi-chromatic point set on the *left*, and a tri-chromatic point set on the *right*. The *dotted* line indicates the separation of *green* from *orange* points that form the background for the *blue* circle.

in K, so we get a point with the same birth-coordinate in the diagram of the image. When the circle is filled by orange disks, it becomes a trivial 1-cycle in K, which is marked by its death in the image and the simultaneous birth in the kernel. Eventually, the blue circle is filled by blue disks, so it dies in the domain and simultaneously in the kernel. To summarize, the point (a, c) in the diagram of the domain splits into two points, (a, b) in the diagram of the image, and (b, c) in the diagram of the kernel. While the split into two like this is a common phenomenon, not all points split in this manner; see the relations in the subsequence subsection. The point (b, c) can also be seen one dimension higher in the relative persisence diagram of the pair. Indeed, there is a non-bounding 2-cycle in the quotient space once the blue circle is filled by orange disks. Similarly, the point (a, b) can also be found in the diagram of the codomain. Both occurrences of (a, b) correspond to the 1-cycle representing the blue circle in homology, which explains why the point is missing in the diagram of the cokernel.

Note that other natural choices of L are the orange subcomplex or the union of the blue and orange subcomplexes, which is a choice that is symmetric with respect to the colors. We now revisit some of these observations in a more general setting, where the pair of topological spaces, $L \subseteq K$, is not necessarily formed by chromatic complexes.

5.4.3 Relations between Diagrams in a 6-pack

The inclusion of sublevel sets $L_i \subseteq K_i$ induces a map on homology $\kappa_i \colon \mathrm{H}(L_i) \to \mathrm{H}(K_i)$. This map has a component in each dimension, p, and we write $ker_p\, \kappa_i$, $im_p\, \kappa_i$, $cok_p\, \kappa_i$ for the kernel, image, cokernel of κ_i in dimension p.

Lemma 5.4.2. *Let $L_i \subseteq K_i$ be simplicial complexes and $\kappa_i \colon H(L_i) \to H(K_i)$ the induced map on homology. For each dimension, p, there are short exact sequences*

$$0 \to ker_p\, \kappa_i \to H_p(L_i) \to im_p\, \kappa_i \to 0, \tag{5.13}$$
$$0 \to im_p\, \kappa_i \to H_p(K_i) \to cok_p\, \kappa_i \to 0, \tag{5.14}$$
$$0 \to cok_p\, \kappa_i \to H_p(K_i, L_i) \to ker_{p-1}\, \kappa_i \to 0. \tag{5.15}$$

Proof. The first two exact sequences are obvious from the definitions and the isomorphism theorem. To see the third exact sequence, we recall the long exact sequence of the pair, see

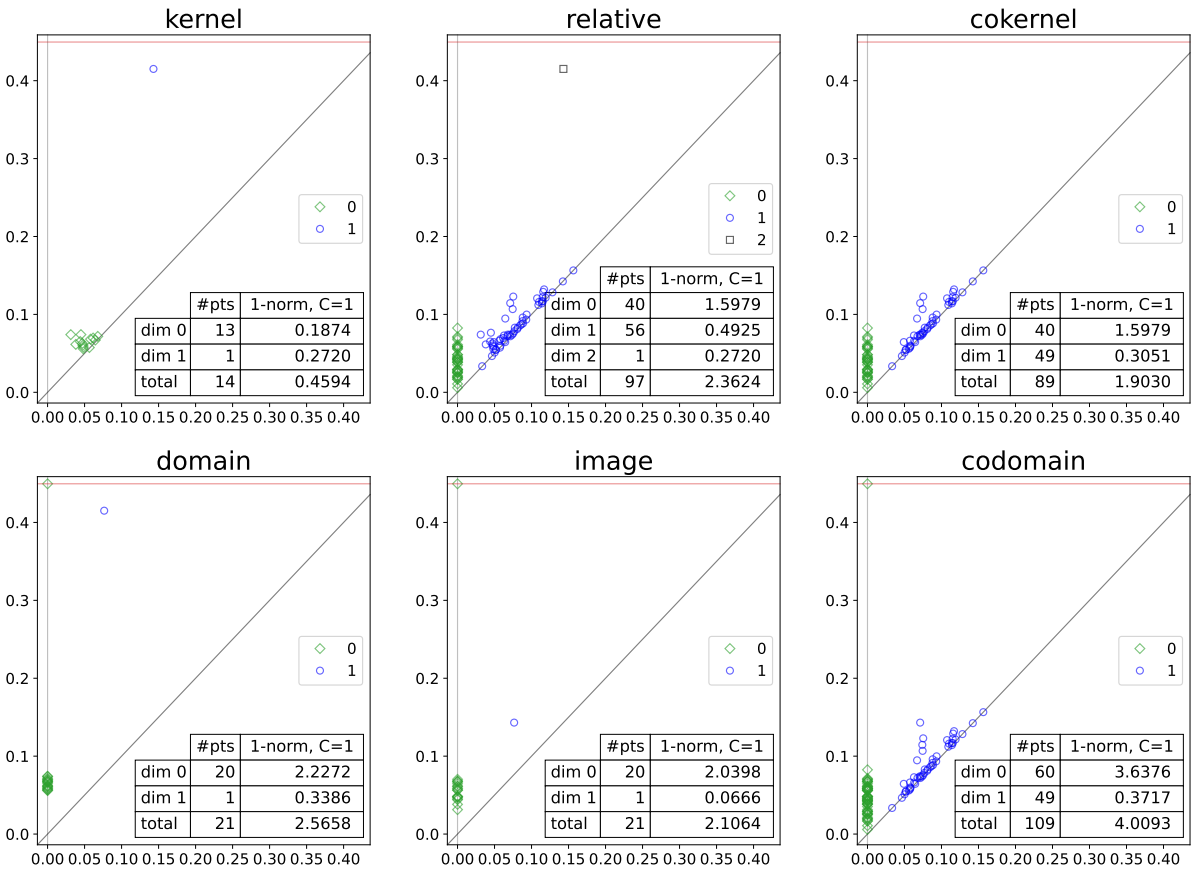

Figure 5.10: The 6-pack for the bi-chromatic point set in the *left panel* of Figure 5.9. The domain, L, is the blue subcomplex of the codomain, K, which is the 3-dimensional chromatic Delaunay mosaic of the *blue* and *orange* points.

Equation (5.6). Working with field coefficients, all homology groups are vector spaces and thus split. In particular, $\mathsf{H}_p(L_i) \simeq ker_p\,\kappa_i \oplus im_p\,\kappa_i$, in which $ker_p\,\kappa_i$ and $im_p\,\kappa_i$ are the images of the incoming and outgoing maps. We can therefore substitute $ker_p\,\kappa_i \to 0 \to im_p\,\kappa_i$ for $\mathsf{H}_p(L_i)$. By the same token, we substitute $im_p\,\kappa_i \to 0 \to cok_p\,\kappa_i$ for $\mathsf{H}_p(K_i)$, and we remove $0 \to im_p\,\kappa_i \to im_p\,\kappa_i$ to get

$$\ldots \to ker_p\,\kappa_i \to 0 \to cok_p\,\kappa_i \to \mathsf{H}_p(K_i, L_i) \to ker_{p-1}\,\kappa_i \to 0 \to cok_{p-1}\,\kappa_i \to \ldots, \quad (5.16)$$

which contains the required third short exact sequence. □

It follows that the ranks of $ker_p\,\kappa_i$ and $im_p\,\kappa_i$ add up to the rank of $\mathsf{H}_p(L_i)$, etc. This implies relations between the 1-norms of corresponding persistence diagrams.

Theorem 5.4.1. *Let $L \subseteq K$ be simplicial complexes, $f_K \colon K \to \mathbb{R}$ monotonic, and f_L, $f_{K,L}$ the restrictions of f_K to L and $K \setminus L$. For each dimension, p, and any fixed cut-off for the 1-norms, $C > 0$,*

$$\|Dgm_p(f_L)\|_1 = \|Dgm_p(ker\ f_L \to f_K)\|_1 + \|Dgm_p(im\ f_L \to f_K)\|_1, \quad (5.17)$$
$$\|Dgm_p(f_K)\|_1 = \|Dgm_p(im\ f_L \to f_K)\|_1 + \|Dgm_p(cok\ f_L \to f_K)\|_1, \quad (5.18)$$
$$\|Dgm_p(f_{K,L})\|_1 = \|Dgm_p(cok\ f_L \to f_K)\|_1 + \|Dgm_{p-1}(ker\ f_L \to f_K)\|_1. \quad (5.19)$$

Proof. We prove (5.17). We write $0 \leq r_1, r_2, \ldots, r_n$ for the values of f_K smaller than C. In addition, set $r_0 = -\infty$ and use the cut-off $r_{n+1} = C$ for the 1-norms. Letting $L_i = {f_L}^{-1}[0, r_i]$, note that $L_i = {f_L}^{-1}[0, r]$ for all $r_i \leq r < r_{i+1}$, and hence the ranks of the various groups are constant between two contiguous values. We can therefore write the 1-norm of $\mathrm{D}gm_p(f_L)$ as a sum of n contributions, and similar for the 1-norms of the kernel and image diagrams:

$$\|\mathrm{D}gm_p(f_L)\|_1 = \sum_{i=0}^{n}(r_{i+1} - r_i)\, rank\, \mathsf{H}_p(L_i), \tag{5.20}$$

$$\|\mathrm{D}gm_p(\ker f_L \to f_K)\|_1 = \sum_{i=0}^{n}(r_{i+1} - r_i)\, rank\, \ker_p \kappa_i, \tag{5.21}$$

$$\|\mathrm{D}gm_p(\mathrm{im}\, f_L \to f_K)\|_1 = \sum_{i=0}^{n}(r_{i+1} - r_i)\, rank\, \mathrm{im}_p \kappa_i. \tag{5.22}$$

We thus get (5.17) from (5.13). With the same argument applied to K_i, image, and cokernel, we get (5.18) from (5.14), and applied to (K_i, L_i), cokernel, and kernel, we get (5.19) from (5.15). □

We note that similar equations do not hold for the 0-norm, which counts the points in the diagrams. Putting the equations in Theorem 5.4.1 together yields a vanishing alternating sum:

$$\sum_{p\in\mathbb{Z}}(-1)^p \left[\|\mathrm{D}gm_p(f_L)\|_1 - \|\mathrm{D}gm_p(f_K)\|_1 + \|\mathrm{D}gm_p(f_{K,L})\|_1\right] = 0. \tag{5.23}$$

While there are relations between the diagrams in a 6-pack, no single diagram is necessarily determined by the others. Figure 5.11 shows one such example.

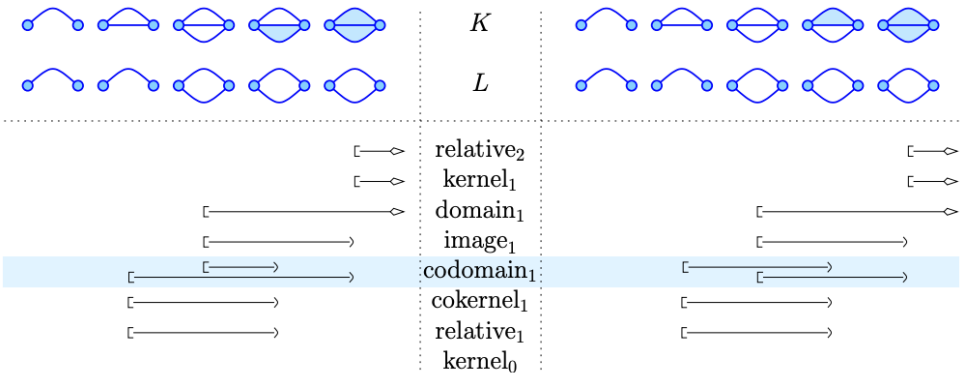

Figure 5.11: Example showing that five diagrams do not imply the sixth. The two filtrations differ by a single 2-dimensional cell added in the respective fourth steps of the filtrations. Correspondingly, five of the 1-dimensional persistence diagrams (shown as barcodes) are the same, while the highlighted diagrams of the codomain differ on the two sides.

Further relations among the diagrams in a 6-pack are suggested by the case-by-case analysis for simultaneous occurrence of births and deaths in various groups provided in [23]. For example, consider the triple $\ker \kappa_i$, $\mathsf{H}(L_i)$, $\mathrm{im}\, \kappa_i$. At a given radius, the rank of each group can change by at most one. The short exact sequence (5.13) reduces the twenty-six non-trivial combinations of changes down to only six. Out of those, [23] gives examples for five of them and shows that the sixth, death-nothing-birth, cannot occur because a death in $\ker \kappa_i$ always implies a death in $\mathsf{H}(L_i)$. This is an additional relation, which is not implied directly by (5.13). The same case is excluded for the triple in (5.14). Analogously, one can show that the case death-nothing-birth is excluded for the triple $cok_p\, \kappa_i$, $\mathsf{H}_p(K_i, L_i)$, $\ker_{p-1} \kappa_i$.

5.4.4 Relations between 6-packs of a Triplet

The framework described so far is amenable to a pair of complexes $L \subseteq K$, filtered by a monotonic function. This section addresses the next simplest case: when we have a sequence of three nested complexes, $M \subseteq L \subseteq K$, which gives rise to four long exact sequences:

$$\ldots \to \mathsf{H}_p(L) \to \mathsf{H}_p(K) \to \mathsf{H}_p(K,L) \to \mathsf{H}_{p-1}(L) \to \ldots, \tag{5.24}$$
$$\ldots \to \mathsf{H}_p(M) \to \mathsf{H}_p(K) \to \mathsf{H}_p(K,M) \to \mathsf{H}_{p-1}(M) \to \ldots, \tag{5.25}$$
$$\ldots \to \mathsf{H}_p(M) \to \mathsf{H}_p(L) \to \mathsf{H}_p(L,M) \to \mathsf{H}_{p-1}(M) \to \ldots, \tag{5.26}$$
$$\ldots \to \mathsf{H}_p(L,M) \to \mathsf{H}_p(K,M) \to \mathsf{H}_p(K,L) \to \mathsf{H}_{p-1}(L,M) \to \ldots. \tag{5.27}$$

To shed light on how they relate to each other, we draw them as sine-like curves, each directed

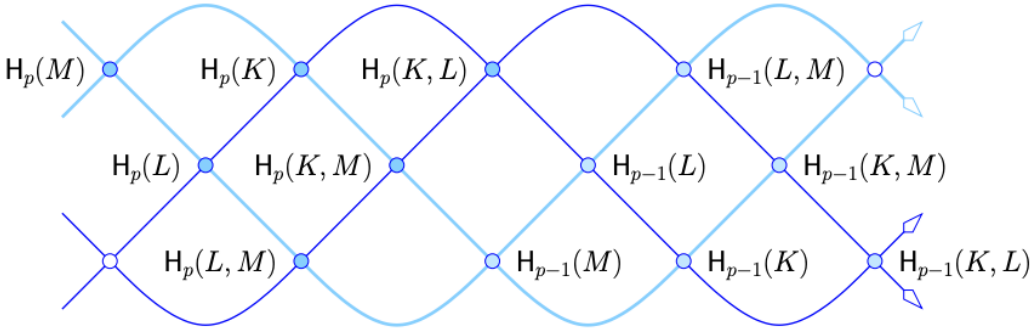

Figure 5.12: The four exact sequences for three complexes drawn along sine-like curves in the plane. After each half-period, the dimension of the homology group drops by one.

from left to right, with the homology groups sitting at the crossings between the curves; see Figure 5.12. Observe that the upper left triangular diagram commutes, which implies

$$ker\,[\mathsf{H}_p(M) \to \mathsf{H}_p(L)] \subseteq ker\,[\mathsf{H}_p(M) \to \mathsf{H}_p(K)], \tag{5.28}$$
$$im\,[\mathsf{H}_p(M) \to \mathsf{H}_p(K)] \subseteq im\,[\mathsf{H}_p(L) \to \mathsf{H}_p(K)] \tag{5.29}$$

for all dimensions p. Similar inclusions follow from the commutativity of the other regions in the arrangement of curves. The four inclusions that give rise to the sequences (5.24) to (5.27) yield four 6-packs, among which six diagrams appear twice, namely, $\mathsf{D}gm(f_K)$, $\mathsf{D}gm(f_L)$, $\mathsf{D}gm(f_M)$, $\mathsf{D}gm(f_{K,L})$, $\mathsf{D}gm(f_{K,M})$, $\mathsf{D}gm(f_{L,M})$. Therefore, we have eighteen unique diagrams, some of which are closely related.

5.4.5 A Tri-chromatic Case Study

While two colors give rise to interesting patterns, more colors do more so. With the increase in the number of colors, there is an explosive increase of configurations to study. We suggest looking at the relations between *k-chromatic subcomplexes* of $Del(\chi)$, which are the subcomplexes composed of all simplices with at most k colors. In this section, we focus on the tri-chromatic case, with colors $\sigma = \{0,1,2\}$. Let M be the 1-chromatic subcomplex, L the 2-chromatic subcomplex, and K the full 3-chromatic Delaunay complex. As before, $f_K \colon K \to \mathbb{R}$ is the chromatic squared radius function, and f_L, f_M, $f_{K,L}$, $f_{K,M}$, $f_{L,M}$ are its restrictions. A cycle can be formed by points of 1, 2, or 3 colors, and it can be filled by points of 0, 1, or 2 additional colors. Requiring that the sum of two numbers is at most 3, we get the

six mingling types sketched in Figure 5.1. Note that these six patterns are not independent. For example, the pattern 1+2 also gives rise to pattern 1+0, because the cycle gets filled by its own color eventually. However, different patterns corresponding to the same cycle will generally have different persistence, which quantifies which patterns is a better fit. Without a claim on completeness, we list where in the 6-pack one can find prominent cases of each of these six patterns.

Case 1+0: D$gm(f_M)$. The complex M is the disjoint union of the three mono-chromatic Delaunay mosaics. The diagram records the mono-chromatic cycles.

Case 2+0: D$gm(cok\ f_M \to f_L)$. The complex L contains all mono- and bi-chromatic cycles, and it shares the former with M. Therefore, we look at the cokernel to keep only the cycles that need two colors to be formed. A cycle dies either when it is filled by its own two colors, or when one of the two colors suffices to form a homologous cycle.

Case 3+0: D$gm(cok\ f_L \to f_K)$. As in the previous case, we look at the cokernel to capture cycles that are formed by all three colors, but not yet by any two.

Case 1+1: D$gm(\mathrm{k}er\ f_M \to f_L)$. A birth in this diagram occurs when a cycle formed by one color is filled by an additional one, and it persists until it is filled by its own color.

Case 2+1: D$gm(\mathrm{k}er\ f_{L,M} \to f_{K,M})$. The idea is similar to Case 1+1: we look at cycles formed by two colors that are filled when also using the third. Unlike the previous case, we consider the quotient spaces to filter out the mono-chromatic cycles.

Case 1+2: D$gm(cok\ f_{L,M} \to f_{K,M})$. Mono-chromatic p-cycles filled by the other colors appear in the pair (K, M) as $(p+1)$-cycles. Those that are filled by exactly one other color also appear in (L, M). We use the cokernel to filter them out.

We now look more closely at the concrete example displayed in the right panel of Figure 5.9: a circle of blue points with split background of green and orange points; compare with the mingling pattern 1+2. Focusing on this pattern, we search the 6-pack of the inclusion of the pairs $(L, M) \subseteq (K, M)$ in Figure 5.13. As suggested in Case 1+2 above, we expect a clear signal in the cokernel diagram, and indeed we see a single prominant point representing a 2-dimensional relative class. By construction, this class is born when the mono-chromatic 1-cycle is filled with two extra colors, and its persistence indicates how much longer it takes to fill the 1-cycle with just one extra color. Compare this with the even more prominent point in the diagram of the codomain, D$gm(f_{K,M})$. This point represents the same 1-cycle, but it expresses different information because it is not sensitive to whether the 1-cycle is filled by one or two additional colors.

It is interesting to interpret the two high persistence points in the diagram of the domain, which records classes in H(L_i, M_i). Since the background consists of two colors, it fills the blue 1-cycle with only one additional color twice, once with green and another time with orange. Both classes die at the same moment, namely when the blue cycle is filled by its own color.

5.5 Towards Defining the Topological Mingling Numbers

Just like alpha complexes can be used to quantify the *shape* of a point set, can chromatic alpha complexes to used to quantify the *mingling* of finite sets of points? This subsection

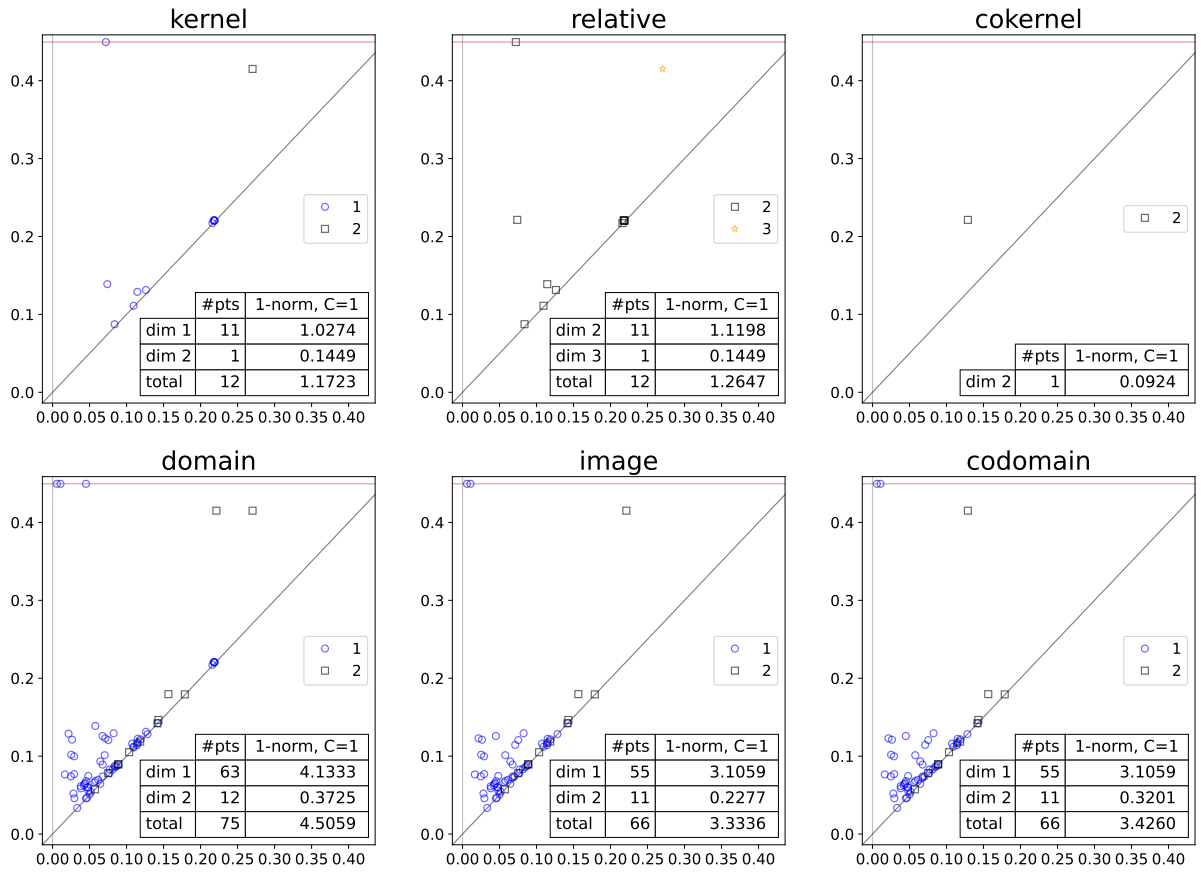

Figure 5.13: The 6-pack of $(L, M) \subseteq (K, M)$ for the data in the *right panel* of Figure 5.9. M, L, and K are the 1-, 2- and 3-chromatic subcomplexes of the chromatic Delaunay complex.

aims to collect the first thoughts of what is clearly a work in progress. We start by saying that this is not an easy task, for instance, which of the 3 coloring of the same point set would you say showcases a more mingled configuration in Figure 5.14?

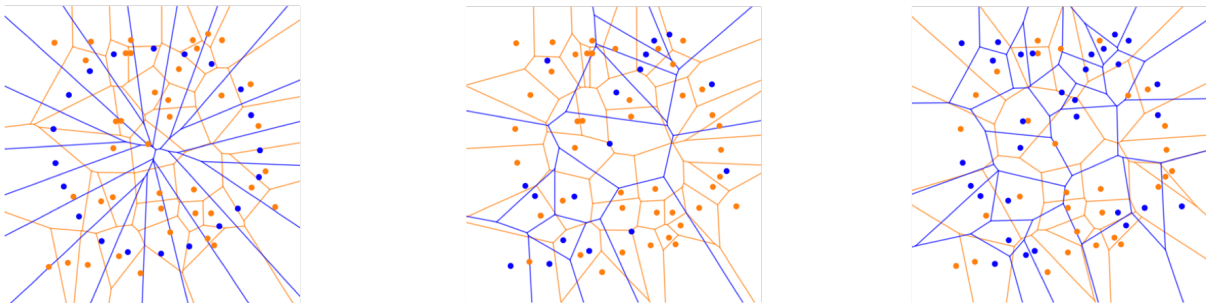

Figure 5.14: The overlay of the Voronoi tessellations for three coloring of the same point set.

When points of different colors are well-separated, say by a hyperplane, the kernel diagram is empty in every dimension. As we slowly *mix* the point sets, homology classes start migrating from the image diagram to the kernel diagram. Note that the sum of the 1-norm of the two diagrams is constant given a coloring of the point set. Can the share of homological information in the kernel diagram be used to obtain a set of quantifiers for how *mixed* or *mingled* two point sets are? Are those numbers really matching our intuition of what topological mixedness should be in all homological dimension? If yes, that would great as one could appeal to the

stability of the persistent diagrams with respect to small perturbations in the input point set, to obtain provably stable *topological mingling numbers*.

We have started investigating this question for two dimensional Euclidean bi-colored lattices (red and blue). As a first step, we are interested in understanding the 0-dimensional *image share*, which is the one norm of the image diagram, divided by the one norm of the domain diagram, where the latter is the union of the red and blue subcomplexes. Note that the information contained in the 0-dimensional persistent diagram is nothing more than a list of numbers representing the death radii of the connected components, and this is exactly the list of lengths of the edges in the *minimum spanning tree* of the point set divided by 2. As the image share depends on the coloring of the point set, a natural question arises: given a point set, how big can the sum of the lenghts of the blue and the red spanning trees get if we are allowed to chose the coloring?

To be more precise, let $A \subseteq \mathbb{R}^2$ be finite, write $\mathrm{M}ST(A)$ for the minimum spanning tree of A, and let $|\mathrm{M}ST(A)|$ be its Euclidean length. We note that the MST of A is generally not unique, but its length is.

Definition 5.5.1 (Minimum MST-ratio)**.** *Let $A \subseteq \mathbb{R}^d$ be finite and $\chi\colon A \to \{0, 1, \ldots, s\}$ a coloring. Hence, $A = A_1 \sqcup A_2 \sqcup \ldots \sqcup A_s$ is a partition of A, in which we write $A_j = \chi^{-1}(j)$. The* MST-ratio *of this partition is*

$$\mu(\chi) = \frac{|MST(A)|}{\sum_{j=0}^{s} |MST(A_j)|}, \tag{5.30}$$

and the minimum MST-ratio *of A is $\mu_s(A) = \min_\chi \mu(\chi)$, in which the minimum is taken over all partitions into $s+1$ subsets.*

We are interested in how small and how big the minimum MST-ratio can be. A manuscript is under development, where we prove tight bounds for the minimum MST-ratio for bi-colored lattices in $\mathbb{R}^2$: $\inf \min = 0.5$ and $\sup \min = 0.8$ [26].

Recently, Dumitrescu, Pach and Tóth studied the problem for bicolored points randomly sampled in the unit square in $\mathbb{R}^2$ [40]. In their setting, they maximize the inverse of our MST-ratio for a bipartition of the point set P into red and blue points, denoted R and B respectively. They showed that for a set of n random points uniformly distributed in $[0, 1]^2$, and any $\epsilon > 0$, the following (expected) ratio

$$\frac{|\mathrm{M}ST(R)| + |\mathrm{M}ST(B)|}{|\mathrm{M}ST(P)|},$$

in a maximizing partition is at least $\sqrt{2} - \epsilon$ with probability tending to 1 as the number of points tends to infinity.

It would be interesting to investigate the distribution of this number for random colorings (as opposed to only the maximizing one) in order to develop an intuition for the ratio's dependence on the coloring. Figure 5.15 shows a curious plot of the ratio studied by Pach and coauthors for 100 points randomly sampled from the square (repeated 1000 times). Each point is assigned the red color with probability 0.5. The mean is not much more than $\sqrt{2} = 1.414\cdots$, which is surprising as the coloring is not chosen to maximize the ratio, but randomly. We believe this not to be a coincidence and we are working on strengthening their result [40, Theorem 4].

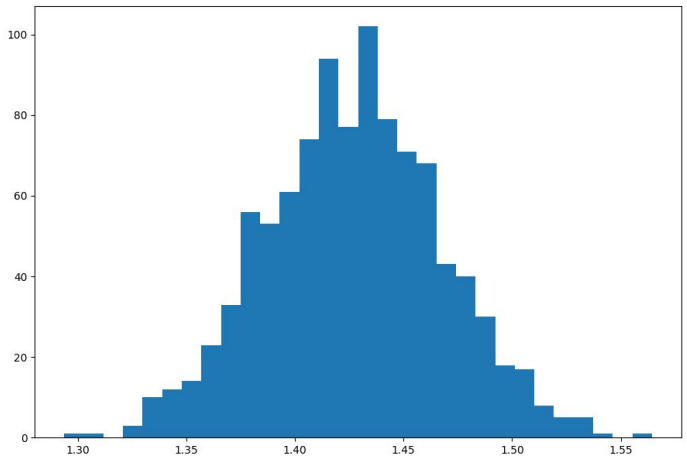

Figure 5.15: The MST ratio for 100 points randomly sampled from the square (repeated 1000 times). Each point is assigned color red with probability 0.5.

We plan to study the analogous ratio in higher dimensions and establish a better understanding of the average and extremal cases. This will serve as a baseline against which to compare real data, which often comes with a predefined coloring, such as the phenotype of a cell in our favorite biological application.

5.6 Discussion and Open Problems

The main contribution of the papers that this chapter builds on is the extension of the theory of alpha complexes to the setting where points are assigned a label. The work reported in this chapter suggests new directions of mathematical research aimed at solidifying our understanding of the chromatic setting. We list two possible directions.

- Develop a **chromatic variant of Forman's discrete Morse theory** [55]. Two concrete questions are the extension of the collapsibility of the Čech complex to the alpha complex proved in the mono-chromatic case [8] and the further collapse of $\mathrm{A}lf_r(\chi)$ to $\mathrm{A}lf_r(A)$.
- In many biological questions, the mingling between different populations of cells changes over time; see e.g. the study of cell segregation in early development [79] and an early topological approach in [69]. It would therefore be useful to **extend the vineyard algorithm** [24] to the chromatic setting introduced in this paper.

Bibliography

[1] P.K. Agarwal, B. Aronov, T.M. Chen and M. Sharir. On levels in arrangements of lines, segments, planes, and triangles. *Discrete Comput. Geom.* **19** (1998), 315–331.

[2] N. Alon and E. Győri. The number of semispaces of a finite set of points in the plane. *J. Combin. Theory Ser. A* **41** (1986), 154–157.

[3] A. Andrzejak and E. Welzl. In between k-sets, j-facets, and i-faces: (i, j)-partitions. *Discrete Comput. Geom.* **29** (2003), 105–131.

[4] C.R. Aragon and R.G. Seidel. Randomized search trees. *Algorithmica* **16** (1996), 464–497.

[5] J. Arsuaga, N. A. Baas, H. Daniel DeWoskin Mizuno, A. Pankov, C. Park Topological analysis of gene expression arrays identifies high risk molecular subtypes in breast cancer. *Applicable Algebra in Engineering, Communications and Comput.* **23** (2012), 3–15.

[6] F. Aurenhammer. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM Comput. Surv.*, **23**, (1991), 345–405.

[7] A. Aukerman, M. Carriere, C. Chen, K. Gardner, R. Rabadan, and R. Vanguri. Persistent homology based characterization of the breast cancer immune microenvironment: a feasibility study, *In* "Proc. 37th Ann. Sympos. Comput. Geom., 2020.", 183–206.

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

[9] R. Biswas, S. Cultrera di Montesano, O. Draganov, H. Edelsbrunner, and M. Saghafian. On the size of chromatic Delaunay mosaics. `arXiv:2212.03121` (2022).

[10] R. Biswas, S. Cultrera di Montesano, H. Edelsbrunner and M. Saghafian. Counting cells of order-k Voronoi tessellations in $\mathbb{R}^3$ with Morse theory. *In* "Proc. 37th Ann. Sympos. Comput. Geom., 2021", 16:1–16:15.

[11] R. Biswas, S. Cultrera di Montesano, H. Edelsbrunner and M. Saghafian. Depth in arrangements: Dehn–Sommerville–Euler relations with applications. Accepted in *Jour. of App. and Comput. Top.*, (2023).

[12] R. Biswas, S. Cultrera di Montesano, H. Edelsbrunner and M. Saghafian. Geometric characterization of the persistence of 1D maps. *J Appl. and Comput. Topology* (2023), 1–19.

[13] A. Björner. Topological methods. In *Handbook of Combinatorics*, eds. R. Graham, M. Grötschel and L. Lovász, 1819–1872, North-Holland, Amsterdam, 1995.

[14] K. Borsuk. On the imbedding of systems of compacta in simplicial complexes. *Fund. Math.* **35** (1948), 217–234.

[15] K. Q. Bwown. Voronoi diagrams from convex hulls. *Inf. Process. Lett.* **9** (1979), 223-228.

[16] A. Bukkuri, N. Andor, and I. K. Darcy. Applications of Topological Data Analysis in Oncology. *Frontiers in Artificial Intelligence* **4** (2021), 659037.

[17] G. Carlsson. Topology and data. *Bull. New Ser. Am. Math. Soc.* **46** (2009), 255–308.

[18] H.-L. Cheng, H. Edelsbrunner and P. Fu. Shape space from deformation. *Comput. Geom. Theory Appl.* **19** (2001), 191–204.

[19] Y.M. Chung, C.S. Hu, Y.L. Lo and H.T. Wu. A persistent homology approach to heart rate variability analysis with an application to sleep-walk classification. *Frontier Physiol.* **12** (2021), 637–684.

[20] K.L. Clarkson and P.W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.* **4** (1989), 387–421.

[21] D. Cohen-Steiner, H. Edelsbrunner and J. Harer. Stability of persistence diagrams. *Discrete Comput. Geom.* **37** (2007), 103–120.

[22] D. Cohen-Steiner, H. Edelsbrunner and J. Harer. Extending persistence using Poincaré and Lefschetz duality. *Found. Comput. Math.* **9** (2009), 79–103.

[23] D. Cohen-Steiner, H. Edelsbrunner. J. Harer and D. Morozov. Persistent homology for kernels, images, and cokernels. In "Proc. 20th Ann. ACM-SIAM Sympos. Discrete Alg., 2009", 1011–1020.

[24] D. Cohen-Steiner, H. Edelsbrunner and D. Morozov. Vines and vineyards by updating persistence in linear time. *In* "Proc. 22nd Ann. Sympos. Comput. Geom., 2006", 119-126.

[25] S. Cultrera di Montesano, O. Draganov, H. Edelsbrunner, and M. Saghafian. Chromatic alpha complexes. `arXiv:2212.03128` (2024).

[26] S. Cultrera di Montesano, O. Draganov, H. Edelsbrunner, and M. Saghafian. *Euclidean MST-ratio for Bi-colored Lattices.* In preparation.

[27] S. Cultrera di Montesano, H. Edelsbrunner, M. Henzinger and L. Ost. Dynamically Maintaining the Persistent Homology of Time Series. *In* "Proc. 35th Ann. ACM-SIAM Sympos. Discrete Alg. 2024", 243–295.

[28] D. R. Chittajallu, et al. Vectorized persistent homology representations for characterizing glandular architecture in histology images, *In* "IEEE 15th International Symposium on Biomedical Imaging 2018", 232–235.

[29] Y.-M. Chung, C.-S. Hu, A. Lawson, C. Smyth. Topological approaches to skin disease image analysis, *In* "IEEE International Conference on Big Data 2018", 100–105.

[30] L. Crawford, A. Monod, A. X. Chen, S. Mukherjee, R. Rabadán. Predicting clinical outcomes in glioblastoma: an application of topological and functional data analysis, *Journal of the American Statistical Association* (2020), 1139-1150.

[31] R. Descartes. *Principia Philosophiae.* Ludovicus Elzevirius, Amsterdam, 1644.

[32] M. Dehn. Die Eulersche Formel in Zusammenhang mit dem Inhalt in der nichtEuklidischen Geometrie. *Math. Ann.*, **61** (1905), 561–586.

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

[34] M.-L. Dequeant, et al. Comparison of pattern detection methods in microarray time series of the segmentation clock. *PLoS One* **3** (2008), e2856.

[35] T.K. Dey. Improved bounds for planar k-sets and related problems. *Discrete Comput. Geom.* **19** (1998), 373–382.

[36] T.K. Dey, T. Hou and S. Parsa. Revisiting graph persistence for updates and efficiency. *In* "Proc. Workshop Algor. Data Struct., 2023", 371-385.

[37] D. DeWoskin , J. Climent, I. Cruz-White, M. Vazquez, C. Park, J. Arsuaga. Applications of computational homology to the analysis of treatment response in breast cancer patients. *Topol. Appl.* **157** (2010), 157–164.

[38] G. L. Dirichlet. Über die Reduction der positiven quadratischen Formen mit drei unbestimmten ganzen Zahlen. *J. reine angew. Math.* **40**, (1850), 209–227.

[39] O. Draganov and M. Mahini. Chromatic-tda. `github.com/OnDraganov/chromatic-tda` (2023).

[40] A. Dumitrescu, J. Pach, G. Tóth. Two trees are better than one, `arXiv:2312.09916` (2024).

[41] R.A. Dwyer. Higher-dimensional Voronoi diagrams in linear expected time. *Discrete Comput. Geom.* **6** (1991), 343–367.

[42] H. Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer-Verlag, Heidelberg, 1987.

[43] H. Edelsbrunner and J.L. Harer. *Computational Topology. An Introduction.* American Mathematical Society, Providence, Rhode Island, 2010.

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

[45] H. Edelsbrunner and P. Koehl. The geometry of biomolecular solvation. In *Discrete Comput. Geom.* **52**, 2005, 243–275.

[46] H. Edelsbrunner, D. Letscher and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.* **28** (2002), 511–533.

[47] H. Edelsbrunner and E.P. Mücke. Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graphics* **9** (1990), 66–104.

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

[49] H. Edelsbrunner and G. Osang. A simple algorithm for computing higher-order Delaunay mosaics. *Algorithmica* **85** (2023), 277-295.

[50] H. Edelsbrunner, R. Seidel. Voronoi diagrams and arrangements. *Discrete Comput. Geom.* **1**, (1986), 25–44.

[51] H. Edelsbrunner, P. Valtr and E. Welzl. Cutting dense point sets in half. *Discrete Comput. Geom.* **17** (1997), 243–255.

[52] H. Edelsbrunner and E. Welzl. On the number of line separations of a finite set in the plane. *J. Combin. Theory Ser. A* **38** (1985), 15–29.

[53] P. Erdős, L. Lovász, A. Simmons and E.G. Straus. Dissection graphs of planar point sets. In *A Survey of Combinatorial Theory*, J.N. Srivastava, et al. (eds.), North-Holland, Amsterdam (1973), 139–149.

[54] G. Fejes Toth. Multiple packing and covering of the plane with circles. *Acta Math. Acad. Sci. Hung.* **27** (1976), 135–140.

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

[56] S.J. Fortune. Voronoi diagrams and Delaunay triangulations. *Computing in Euclidean Geometry*, 2nd edition, (1995), 225–265.

[57] P. Frosini and C. Landi. Size theory as a topological tool for computer vision. *Pattern Recognition and Image Analysis* **9** (1999), 596–603.

[58] D. Gale and L.S. Shapley. College admission and the stability of marriage. *Amer. Math. Monthly* **69** (1962), 9–14.

[59] J.D. Gardiner, J. Behnsen and C.A. Brassey. Alpha shapes: determining 3D shape complexity across morphologically diverse structures. *BMC Evolutionary Biology* **18** (2018), 1–16.

[60] M. Gidea and Y.A. Katz. Topological data analysis of financial time series: landscapes of crashes. *Phys. A: Stat. Mech. Appl.* **491** (2018), 820–834.

[61] M. Glisse. Fast persistent homology computation for functions on $\mathbb{R}$. arXiv:2301.04745v1 (2023).

[62] E. Goodman and R. Pollack. On the number of k-subsets of a set of n points in the plane. *J. Combin. Theory Ser. A* **36** (1984), 101–104.

[63] G. Graff, B. Graff, P. Pilarczyk, G. Jabloński, D. Gasecki and K. Narkiewicz. Persistent homology as a new method of the assessment of heart rate variability. *PLoS ONE* (2021), 0253851.

[64] B. Grünbaum. *Convex Polytopes.* Interscience, John Wiley & Sons, London, 1967.

[65] B. Grünbaum and G. C. Shephard. Convex polytopes, *Bull. London Math. Soc.* **1** (1969), 257–300.

[66] A. Hatcher. *Algebraic Topology.* Cambridge Univ. Press, Cambridge, England, 2002.

[67] T. L. Heath. *Euclid: The Thirteen Books of The Elements*, Dover, 1956.

[68] E. Hlawka. Funktionen von beschränkter Variation in der Theorie der Gleichverteilung. *Ann. Math. Pura Appl.* **54** (1961), 325–333.

[69] M. Kerber and H. Edelsbrunner. 3d kinetic alpha complexes and their implementation. In "Proc. Mtg. Algorithm Engin. Experiments. 2013", 70–77.

[70] J.F. Koksma. A general theorem from the theory of uniform distribution modulo 1. *Mathematica B (Zutphen)* **11** (1942), 7–11.

[71] E.E. Kummer. Über die Ergänzungssätze zu den allgemeinen Reciprocitätsgesetzen. *J. für die Reine und Angew. Math.* **44** (1852), 93–146.

[72] Lawson P, Sholl AB, Brown JQ, Fasy BT, Wenk C. Persistent Homology for the Quantitative Evaluation of Architectural Features in Prostate Cancer Histology. *Sci Rep.*, **9** (1):1139, 2019.

[73] D.-T. Lee. On k-nearest neighbor Voronoi diagrams in the plane. *IEEE Trans. Comput.* **31** (1982), 478–487.

[74] J. Linhart, Y. Yang and M. Philipp. Arrangements of hemispheres and halfspaces. *Discrete Math.* **223** (2000), 217–226.

[75] S. Lockwood, B. Krishnamoorthy. Topological features in cancer gene expression data. *Pac Symp Biocomput.*, ():108-19, 2015.

[76] L. Lovász. On the number of halving lines. *Ann. Univ. Sci. Budapest. Rolando Eőtvős Nomin. Sect. Math.* **14** (1971), 107–108.

[77] Y. Luo and B.J. Nelson. Accelerating iterated persistent homology computations with warm starts. arXiv:2108.05022 (2021).

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

[79] J.-L. Maître, et al. Adhesion functions in cell sorting by mechanically coupling the cortices of adhering cells. *Science* **338** (2012), 253–256.

[80] J. Matoušek. *Lectures on Discrete Geometry*, Springer, Berlin, 2002.

[81] J. Matoušek. *Using the Borsuk–Ulam Theorem. Lectures on Topological Methods in Combinatorics and Geometry.* Second edition, Universitext, Springer, Heidelberg, 2008.

[82] P. McMullen. The maximum numbers of faces of a convex polytope, *Mathematika* **17** (1970), 179–184.

[83] K. Mulmuley. Dehn–Sommerville relations, upper bound theorem, and levels in arrangements. *In* "Proc. 9th Ann. Sympos. Comput. Geom., 1993", 240–246.

[84] A. Oyama, Y. Hiraoka, I. Obayashi, Y. Saikawa, S. Furui, K. Shiraishi, S. Kumagai, T. Hayashi, J. Kotoku Hepatic tumor classification using texture and topology analysis of non-contrast-enhanced three-dimensional T1-weighted MR images with a radiomics approach. *Sci Rep.*; **9** (1):8764, 2019.

[85] J. Pach, W. Steiger and E. Szemerédi. An upper bound on the number of planar k-sets. *Discrete Comput. Geom.* **7** (1992), 109–123.

[86] G. Palla, D. S. Fischer, A. Regev and F. J. Theis. Spatial components of molecular tissue biology. *Nat Biotechnol* **40** (2022), 308–318.

[87] F. Pausinger and A.M. Svane. A Koksma–Hlawka inequality for general discrepancy systems. *J. Complexity* **31** (2015), 773–797.

[88] J.A. Perea, A. Deckard, S.B. Haase and J. Harer. SW1PerS: sliding windows and 1-persistence scoring; discovering periodicity in gene expression time series data. *BMC Bioinformatics* **16** (2015), 1–12.

[89] M. Piekenbrock and J.A. Perea. Move schedules: fast persistence computations in coarse dynamic settings. `arXiv:2104.12285v2` (2021).

[90] H. Poincaré. Complément à l'analysis situs. *Rendiconti del Circolo Matematico di Palermo* **13** (1899), 285–343.

[91] T. Qaiser, et al. Persistent Homology for Fast Tumor Segmentation in Whole Slide Histology Images. *Procedia Comput. Sci.* **90** (2016), 119–124.

[92] R. Rabadán and A. J. Blumberg. *Topological data analysis for genomics and evolution: topology in biology.* Cambridge University Press, 2019.

[93] S. Raghavan, et al. Microenvironment drives cell state, plasticity, and drug response in pancreatic cancer. *Cell* **184** (2021), 6119–6137.

[94] Y. Reani and O. Bobrowski. A coupled alpha complex. *Journal of Computational Geometry* **14** (2023), 221–256.

[95] V. Robins. Toward computing homology from finite approximations. *Topology Proceedings* **24** (1999), 503–532.

[96] J. Saltz, et al. Spatial organization and molecular correlation of tumor-infiltrating lymphocytes using deep learning on pathology images. *Cell Rep.* **23** (2018), 181–193.

[97] M.I. Shamos and D. Hoey. Closest-point problems. *In* "Proc. 16th Ann. IEEE Symp. Found. Comput. Sci., 1975", 151–162.

[98] N. Singh, H. D. Couture, J. S. Marron, C. Perou, M. Niethammer. Topological descriptors of histology images, *In* "Machine Learning in Medical Imaging 2014", 231–239.

[99] M. Sharir, S. Smorodinsky and G. Tardos. An improved bound for k-sets in three dimensions. *Discrete Comput. Geom.* **26** (2001), 195–204.

[100] L. Schläfli. Theorie der vielfachen Kontinuität. Zürich, Basel: Georg & Co. (1852), 167–387.

[101] N.A. Scoville. *Discrete Morse Theory*, American Mathematical Society, 2019.

[102] D. Smirnov and D. Morozov. Triplet merge trees. In *"Topological Methods in Data Analalysis and Visualization 2019"*, 1–17.

[103] D.M.Y. Sommerville. The relations connecting the angle sums and volume of a polytope in space of n dimensions. *Proceedings of the Royal Society Series A* **115** (1927), 103–119.

[104] M. Sorin, M. Rezanejad, E. Karimi, et al. Single-cell spatial landscapes of the lung tumour immune microenvironment. *Nature* **614** (2023), 548–554.

[105] G. Tóth. Point sets with many k-sets. *Discrete Comput. Geom.* **26** (2001), 187–194.

[106] C. D. Tóth, J. O'Rourke and J. E. Goodman., Handbook of Discrete and Computational Geometry, Boca Raton, FL, USA:CRC Press, 2017.

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

[108] J. Vuillemin. A unifying look at data structures. *Commun. ACM* **23** (1980), 229–239.

[109] E. Welzl. More on k-set of finite sets of points in the plane. *Discrete Comput. Geom.* **1** (1986), 95–100.

[110] G.M. Ziegler. *Lectures on Polytopes.* Springer-Verlag, New York, 1995.

[111] R.T. Živaljević. Topological methods in discrete geometry. In *Handbook of Discrete and Computational Geometry*, third edition, eds. J.E. Goodman, J. O'Rourke and C.D. Tóth, 551–580, CRC Press, Boca Raton, Florida, 2017.

[112] R.T. Živaljević and S.T. Vrećica. The colored Tverberg's problem and complexes of injective functions. *J. Combin. Theory Ser. A* **61** (1992), 309–318.

[113] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Comput. Geom.* **33** (2005), 249–274.