Computational Algebraic Topology
Lecture Notes

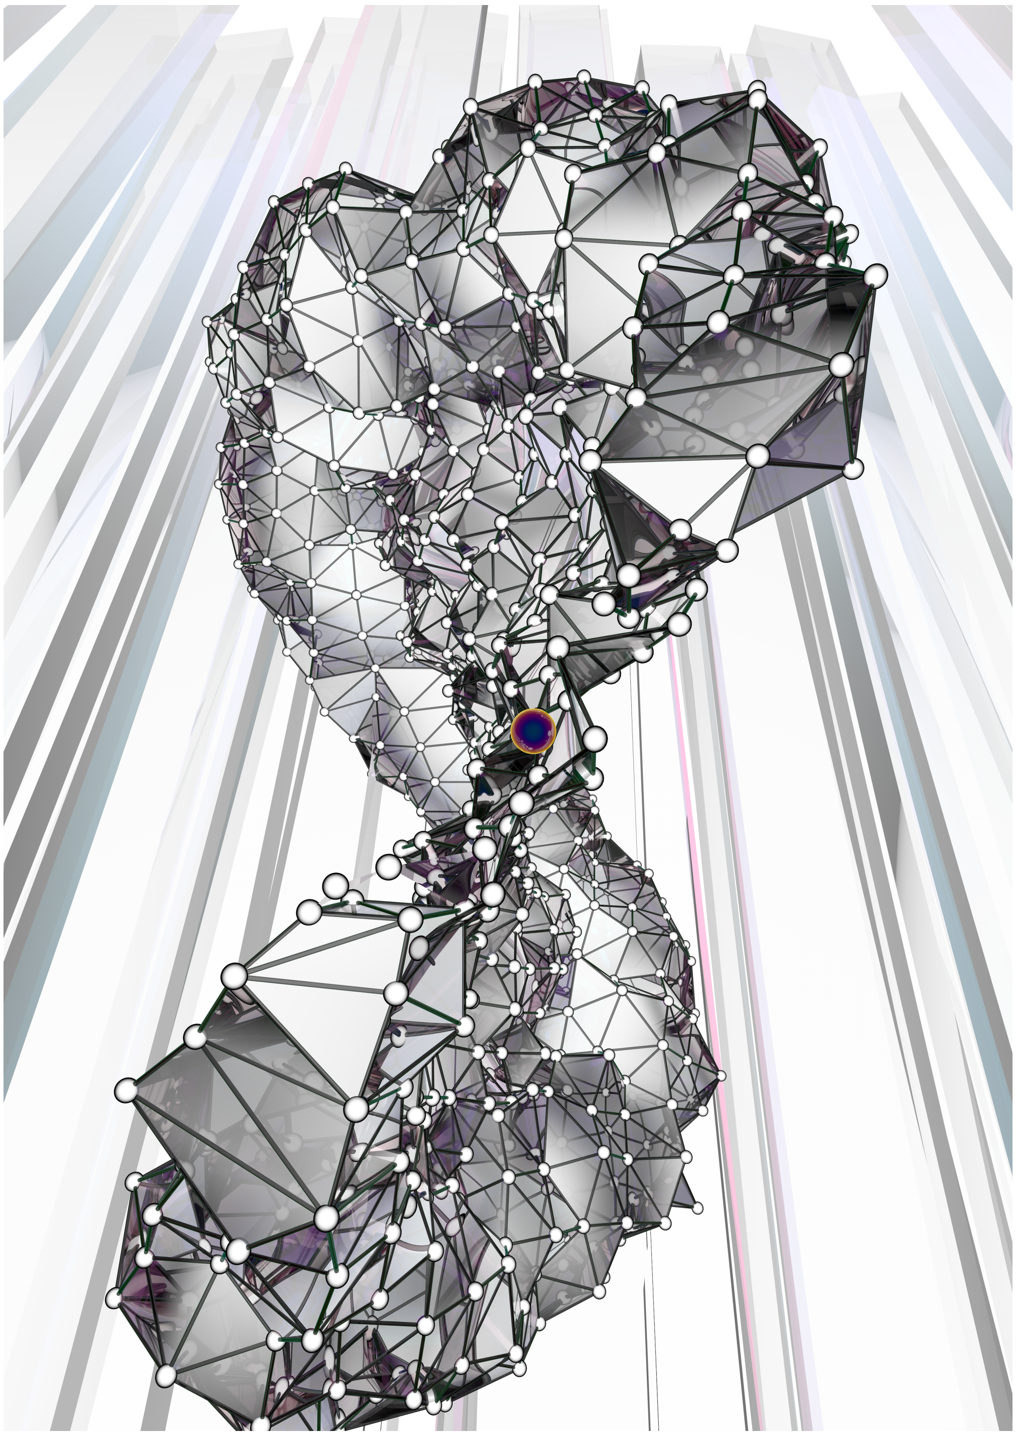

Vidit Nanda
with title illustrations by Robert Ghrist

CONTENTS

1. COMPLEXES

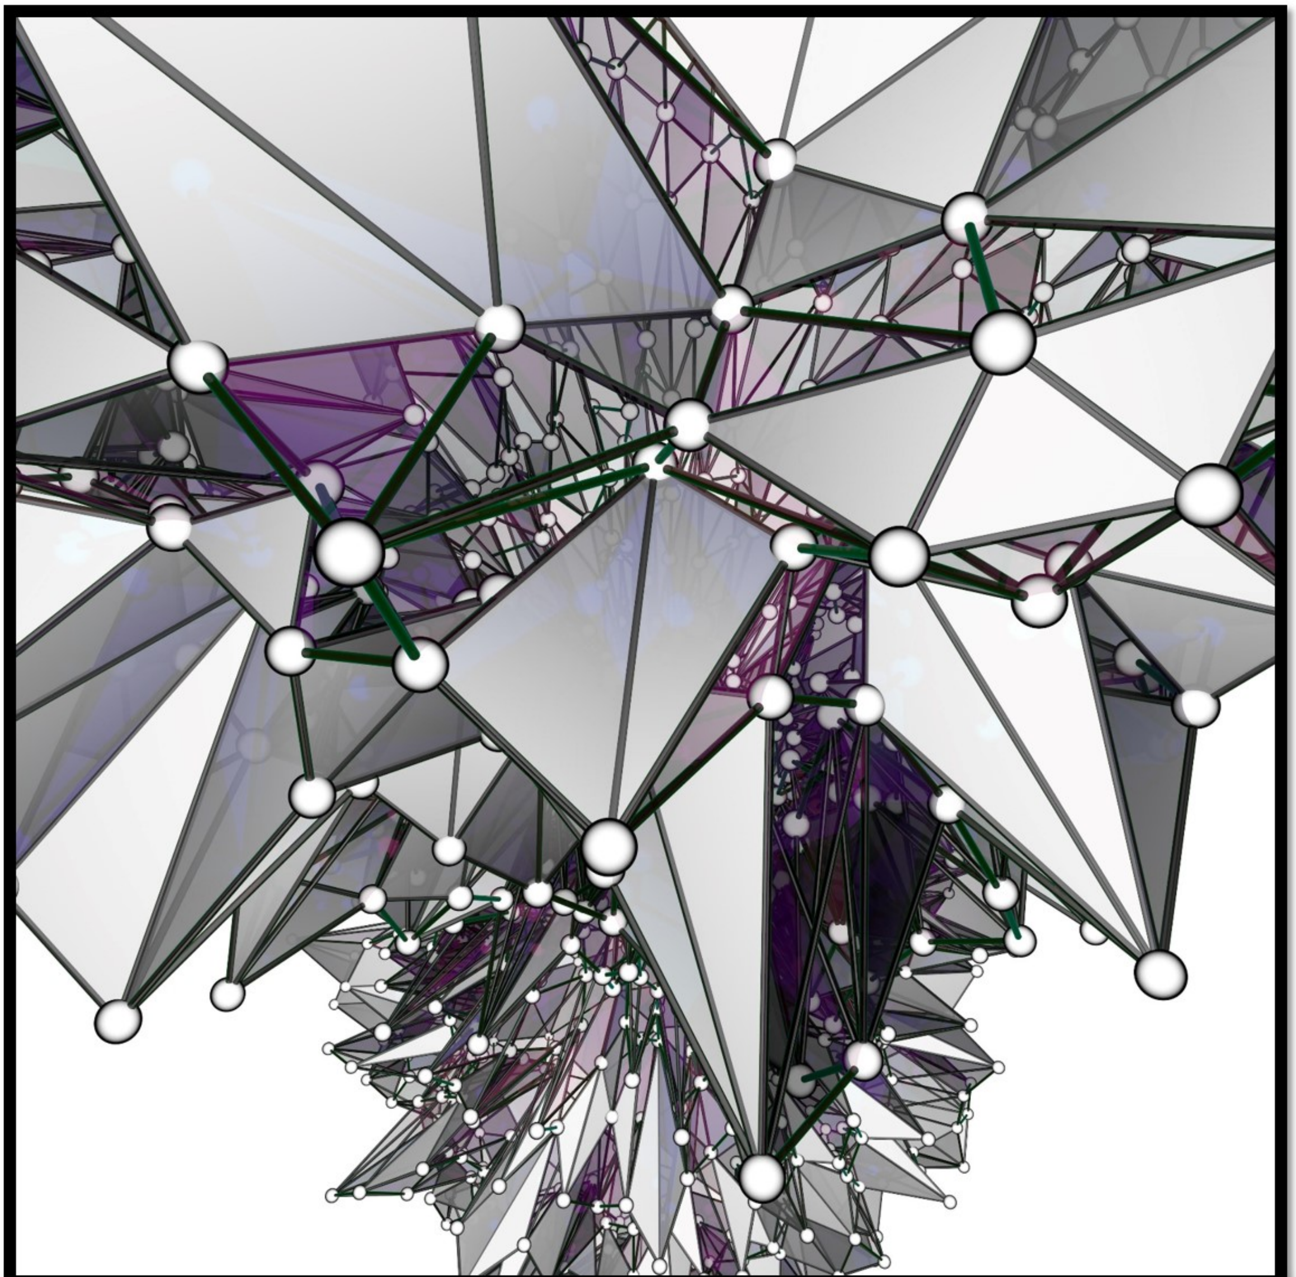

1.1 COMBINATORICS

Let V be a finite nonempty set whose elements we will call *vertices*.

DEFINITION 1.1. A **simplicial complex** on V is a collection K of nonempty subsets of V subject to two requirements:

- for each vertex v in V, the singleton $\{v\}$ is in K, and
- if τ is in K and $\sigma \subset \tau$ then σ must also be in K.

The nonempty subsets which lie in a simplicial complex K are called the **simplices** of K. The *dimension* of a simplex σ in K is defined to be

$$\dim \sigma = \#\sigma - 1,$$

where $\#\sigma$ denotes the cardinality of (or, the number of vertices contained in) σ. Thus, the singletons $\{v\}$ all lie in K and have dimension zero, all pairs $\{v, v'\}$ which happen to lie in K have dimension one, and so forth. The dimension of K itself is given by taking a maximum over constituent simplices, i.e.,

$$\dim K = \max\{\dim \sigma \mid \sigma \in K\}.$$

We will write K_i to denote the set of all i-dimensional simplices in K; the first requirement of Definition 1.1 guarantees that K_0 equals the vertex set V. The figure below contains cartoon depictions of a vertex set V with four elements, a simplicial complex K and a non-simplicial complex K' — the fact that the set $\{1, 2, 3\}$ is present in K' but the subset $\{1, 3\}$ is not disqualifies K' from being a simplicial complex.

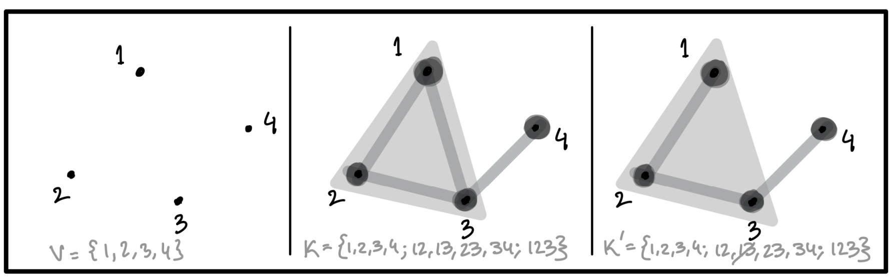

Here are some more exciting examples of simplicial complexes.

- **Graphs**: a (finite, undirected, simple) graph is a pair $G = (V, E)$ consisting of a finite set V (whose elements are called vertices as before) and a set $E \subset V \times V$ consisting of distinct vertex-pairs, usually called *edges*. Every graph automatically forms a one-dimensional simplicial complex K with $V = K_0$ and $E = K_1$.
- **Solid Simplices**: for each integer $k \geq 0$, the solid k-simplex is the simplicial complex $\Delta(k)$ defined on the vertex set $\{0, 1, \ldots, k\}$ whose simplices are *all* possible subsets of vertices.
- **Hollow simplices**: the hollow k-simplex (for each integer $k \geq 1$) is denoted $\partial\Delta(k)$ and defined exactly like a solid k-simplex, *except* that we discard the unique top-dimensional simplex $\{v_0, \ldots, v_k\}$. Thus, $\partial\Delta(k)$ has dimension $k - 1$.

The figure below illustrates a graph, a solid 2-simplex and a hollow 2-simplex respectively.

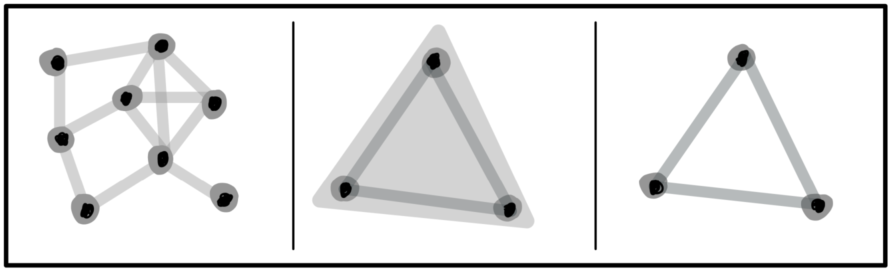

So far, the structure of a simplicial complex appears to be purely combinatorial — we are given a universal finite set V of vertices, and we may select any collection K of subsets of V provided that the two constraints of Definition 1.1 are satisfied. The first step towards expanding this perspective beyond combinatorics is to formally relate simplices with their subsets.

DEFINITION 1.2. Given two simplices σ and τ of a simplicial complex K, we say that σ is a **face** of τ, denoted $\sigma \leq \tau$, whenever every vertex of σ is also a vertex of τ.

Given a pair $\sigma \leq \tau$ of simplices of a simplicial complex K, we call the difference $\dim \tau - \dim \sigma$ the *codimension* of σ as a face of τ; note that the codimension is always a non-negative integer.

1.2 Subcomplexes, Closures and Filtrations

Knowledge of face relations between simplices allows us to define subsets of simplicial complexes which are simplicial complexes in their own right.

DEFINITION 1.3. Let K be a simplicial complex. A subset $L \subset K$ of simplices is called a **subcomplex** of K if it satisfies the following property: for each simplex τ in L, if σ is a face of τ in K, then σ also belongs to L.

In general, for a subcomplex $L \subset K$, we do *not* require every vertex of K to be a vertex of L.

EXAMPLE 1.4. Each hollow k-simplex $\partial\Delta(k)$ naturally forms a subcomplex of the corresponding solid k-simplex $\Delta(k)$; each vertex of a given simplicial complex is automatically a subcomplex.

If you are handed a collection K' of simplices in some simplicial complex K, it is often desirable to check how far K' is from being a subcomplex of K. The following notion is often helpful when performing such checks.

DEFINITION 1.5. The **closure** of a collection of simplices K' in a simplicial complex K is defined to be the smallest subcomplex $L \subset K$ satisfying $K' \subset L$.

Evidently, a nonempty subcollection $K' \subset K$ of simplices forms a subcomplex if and only if it equals its own closure. It should be noted that the closure of a given collection K' of simplices can be *much larger* than K'. The following exercise is highly recommended: if σ is a single k-dimensional simplex in a simplicial complex K, show that the closure of σ in K contains $2^k - 1$ simplices. Of particular interest to us here are ascending chains of subcomplexes.

DEFINITION 1.6. Let K be a simplicial complex; a **filtration** of K (of *length* n) is a nested sequence of subcomplexes of the form

$$\boldsymbol{F}_1 K \subset \boldsymbol{F}_2 K \subset \cdots \subset \boldsymbol{F}_{n-1} K \subset \boldsymbol{F}_n K = K.$$

In general, the dimensions of the intermediate $\mathbf{F}_i K$ are not constrained by i. On the other hand, in order to have a well-defined notion of length, we require $\mathbf{F}_i K \neq \mathbf{F}_{i+1} K$ for all i.

The figure below depicts a filtration of length four of the simplicial complex in the right-most panel; the things to check are that each panel contains a genuine simplicial complex, and that these simplicial complexes are getting strictly larger as we scan from left to right.

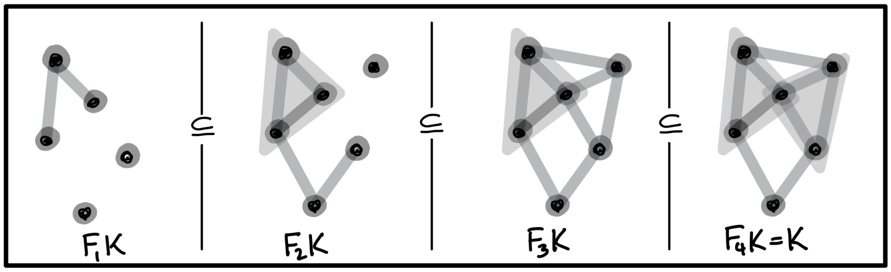

1.3 Geometric Realization

The **geometric simplex** spanned by a collection of points $\{x_0, x_1, \ldots, x_k\}$ in $\mathbb{R}^n$ is the closed subset of $\mathbb{R}^n$ given by

$$\left\{ \sum_{i=0}^{k} t_i x_i \;\middle|\; \text{where } t_i \geq 0 \text{ and } \sum_{i=0}^{k} t_i = 1 \right\}.$$

These points $\{x_0, \ldots, x_k\}$ are said to be *affinely independent* if the collection of vectors

$$\{(x_1 - x_0), (x_2 - x_0), \ldots, (x_k - x_0)\}$$

is linearly independent. There can, therefore, be at most $(n+1)$ affinely independent points in $\mathbb{R}^n$; the canonical example of such a set has x_0 as the origin while x_i for $0 < i \leq n$ is the standard basis vector with 1 in the i-th coordinate and zeros elsewhere.

Definition 1.7. Let $\phi : K_0 \to \mathbb{R}^n$ be any function that sends the vertices of K to points in $\mathbb{R}^n$. The **geometric realization** of K *with respect to* ϕ is the union

$$|K|_\phi = \bigcup_{\sigma \in K} |\sigma|_\phi,$$

where for each $\sigma = \{v_0, \ldots, v_k\}$ in K, the set $|\sigma|_\phi \subset \mathbb{R}^n$ is the geometric simplex spanned by the points $\{\phi(v_0), \ldots, \phi(v_k)\}$.

If we use a particularly degenerate $\phi : K_0 \to \mathbb{R}^n$, such as the map sending every vertex to the origin, then the topological space $|K|_\phi \subset \mathbb{R}^n$ might be quite uninteresting and bear no resemblance with K. We call $\phi : K_0 \to \mathbb{R}^n$ an **affine embedding** of K in $\mathbb{R}^n$ if ϕ is injective (i.e., it sends different vertices to different points) and if its image $\phi(K_0)$ is affinely independent. It turns out that the topology of $|K|_\phi$ is independent of the choice of ϕ provided that we stay within the realm of affine embeddings.

Proposition 1.8. *For any two affine embeddings $\phi, \psi : K_0 \to \mathbb{R}^n$, there is a homeomorphism $|K|_\phi \simeq |K|_\psi$ between the corresponding geometric realizations.*

Proof. Let $K_0 = \{v_0, \ldots, v_k\}$ be the vertex set of K; for each i in $\{1, \ldots k\}$ define the following sets of vectors in $\mathbb{R}^n$

$$x_i = \phi(v_i) - \phi(v_0) \quad \text{and} \quad y_i = \psi(v_i) - \psi(v_0).$$

Since the vectors $\{x_i\}$ and $\{y_i\}$ are linearly independent by our assumption on ϕ and ψ, they each span (possibly distinct) k-dimensional subspaces of $\mathbb{R}^n$. Thus, there is an invertible $n \times n$ matrix M sending x_i to y_i for each i, and this M maps $|K|_\phi$ to $|K|_\psi$ homeomorphically. $\square$

In light of the preceding result, we will usually write the geometric realization of a simplicial complex K as $|K|$, and omit any mention whatsoever of the affine embedding ϕ. It is often convenient to use the endpoints of standard basis vectors in $\mathbb{R}^n$ as targets of the vertices — this ensures, for instance, that every simplicial complex K has a geometric realization embeddable in $\mathbb{R}^n$ for $n = \#K_0$. The figure below depicts the geometric realizations of the solid simplices $\Delta(1)$ and $\Delta(2)$ with respect to this standard basis embedding.

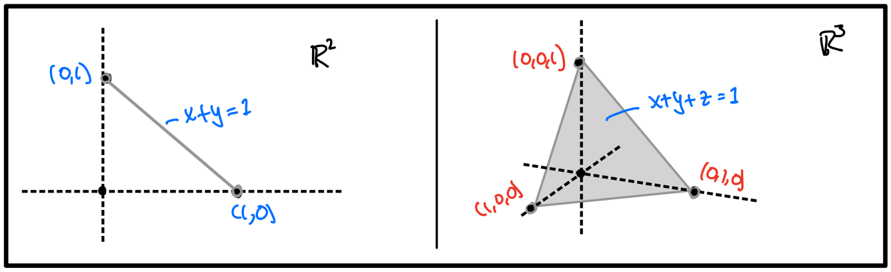

The geometric realization $|\Delta(k)|$ is homeomorphic to a k-dimensional disk while the realization of $\partial\Delta(k)$ is a homeomorphic to the $(k-1)$-dimensional sphere. Geometric realizations allow us to look beyond the combinatorial aspects of simplicial complexes and seek structure in the geometry and topology of their realizations. They also provide a rigorous justification for depicting simplices of dimension $0, 1, 2, 3, \ldots$ as points, lines, triangles, tetrahedra, and so forth.

1.4 Simplicial Maps

Let K and L be simplicial complexes.

Definition 1.9. A **simplicial map** $f : K \to L$ is an assignment $K_0 \to L_0$ of vertices to vertices which sends simplices to simplices. So for each simplex $\sigma = \{v_0, \ldots, v_k\}$ of K, the image $f(\sigma) = \{f(v_0), \ldots, f(v_k)\}$ must be a simplex of L.

It is important to note that f as defined above may not be injective, so in general we allow $f(v_i) = f(v_j)$ even when $v_i \neq v_j$. Thus, we only have an inequality $\dim f(\sigma) \leq \dim \sigma$.

Example 1.10. Whenever $L \subset K$ is a subcomplex, the **inclusion** map $K \hookrightarrow L$ sends each simplex of L to the same simplex in K. In the special case $L = K$, this inclusion is called the **identity map** of K. All such inclusion maps are injective by definition. At the other end of the spectrum, there is a unique surjective simplicial map $K \twoheadrightarrow \bullet$, where $\bullet$ denotes the trivial simplicial complex with only one vertex — so every simplex of K is sent to this single vertex!

One can compose simplicial maps in a straightforward way — given $f : K \to L$ and $g : L \to M$, the composite $g \circ f : K \to M$ sends each simplex σ of K to the simplex $g(f(\sigma))$ of L. We call the simplicial map $f : K \to L$ an **isomorphism** if there exists an inverse, i.e., a simplicial map $g : L \to K$ so that the composites $g \circ f$ and $f \circ g$ are the identity maps of K and L respectively. Simplicial maps induce honest continuous maps between geometric realizations, which behave as well as one might expected, as described in the following result.

Proposition 1.11. *For any simplicial map $f : K \to L$,*

(1) *there is an indued continuous function $|f| : |K| \to |L|$ between geometric realizations so that for each simplex σ in K, the geometric simplex $|f(\sigma)| \subset |L|$ is exactly the image under $|f|$ of the geometric simplex $|\sigma| \subset |K|$; and moreover,*
(2) *if we have a second simplicial map $g : L \to M$, then $|g \circ f|$ and $|g| \circ |f|$ coincide as continuous maps $|K| \to |M|$.*

The proof of both statements is a reasonable exercise once we explain how to construct $|f|$ from f. Let $\phi : K_0 \to \mathbb{R}^m$ and $\psi : L_0 \to \mathbb{R}^n$ be any affine embeddings. Now each point x in $|K| = |K|_\phi$ can be uniquely written as a linear combination $x = \sum_i t_i \cdot \phi(v_i)$ where v_i ranges over all the vertices of K and the t_i are non-negative real numbers satisfying $\sum_i t_i = 1$. The image $|f|(x)$ of this point in $|L| = |L|_\psi$ is then given by the formula

$$|f|(x) = \sum_i t_i \cdot \psi \circ f(v_i). \tag{1}$$

If you restrict this map to the realization of a single simplex $|\sigma|_\phi \subset |K|_\phi$, you will discover that $|f|$ is an honest linear map onto the realization of the image simplex $|f(\sigma)|_\psi \subset |L|_\psi$. For this reason, such continuous maps are called **piecewise-linear**, and their study forms a rich subject in its own right.

One natural question that you might ask is when two simplicial complexes K and L produce homeomorphic geometric realizations $|K|$ and $|L|$. It is a consequence of Proposition 1.8 that any simplicial isomorphism $f : K \to L$ induces a homeomorphism $|f|$ between $|K|$ and $|L|$ — but in general $|K|$ and $|L|$ can be homeomorphic even if there is no simplicial isomorphism relating K to L. We will describe examples of this phenomenon in the next section.

1.5 Barycentric Subdivision

Let K be a simplicial complex.

DEFINITION 1.12. The **barycentric subdivision** of K is a new simplicial complex **Sd** K defined as follows; for each dimension $i \geq 0$, the i-dimensional simplices are given by all sequences

$$\sigma_0 < \sigma_1 < \cdots < \sigma_{i-1} < \sigma_i$$

of (distinct) simplices in K ordered by the face relation.

This definition is liable to cause confusion until we see what barycentric subdivision looks like geometrically. The figures below depict (some) barycentric simplices within the geometric realizations of the solid simplices $\Delta(1)$ and $\Delta(2)$ as well as the hollow 3-simplex $\partial\Delta(3)$.

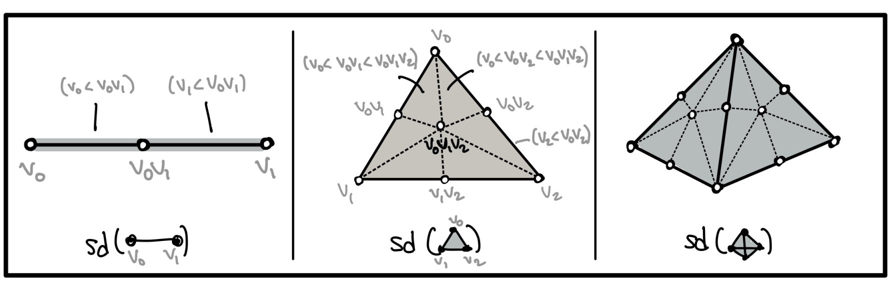

In light of these figures, it is clear that the geometric realizations $|K|$ and $|\mathbf{Sd}\, K|$ agree for every simplicial complex K; we record this not-too-surprising fact below.

PROPOSITION 1.13. *For any simplicial complex K, there is a homeomorphism between geometric realizations $|K|$ and $|\mathbf{Sd}\, K|$.*

You can check just by counting simplices across various dimensions that for non-trivial K there can be no simplicial isomorphism $K \to \mathbf{Sd}\, K$. Since $\mathbf{Sd}\, K$ is itself a simplicial complex, it can be further barycentrically subdivided. We refer to this *second* barycentric subdivision as $\mathbf{Sd}^2\, K = \mathbf{Sd}(\mathbf{Sd}\, K)$, and similarly define $\mathbf{Sd}^n\, K$ for all larger n. By Proposition 1.13, all the geometric realizations $|\,\mathbf{Sd}^n\, K|$ are homeomorphic regardless of $n \geq 1$, even though there are no simplicial isomorphisms which induce these homeomorphisms.

1.6 Filtrations from Data

By **data** here we mean a finite set of observations with a well-defined notion of *pairwise distance*, with the typical example being a finite collection of points in $\mathbb{R}^n$ equipped with the standard Euclidean distance. But in general such observations might not come with any embedding into Euclidean space. One common example is furnished by *dissimilarity matrices* — given a set of observations $O_1, \ldots, O_k$, one can often build a $k \times k$ symmetric matrix whose entry in the (i, j)-th position measures the difference between O_i and O_j. Here is a convenient mathematical framework which encompasses all notions of datasets that are relevant to us here.

Definition 1.14. A **metric space** (M, d) is a pair consisting of a set A and a function

$$d : M \times M \to \mathbb{R},$$

called the *metric*, satisfying four properties:

1. **identity**: $d(x, x) = 0$ for each x in M,
2. **positivity**: $d(x, y) > 0$ for each $x \neq y$ in M,
3. **symmetry**: $d(x, y) = d(y, x)$ for all x, y in M, and most importantly,
4. **triangle inequality**: $d(x, y) + d(y, z) \geq d(x, z)$ for all x, y, z in M.

When the metric is clear from context, we will denote the metric space simply by M; this happens, for instance, when M is a a subset of some Euclidean space $\mathbb{R}^n$. In this case, $d(x, y)$ is understood to be the Euclidean distance $\|x - y\|$ for all x and y in M. In fact, any subset $A \subset M$ of an ambient metric space (M, d) is automatically given the structure of a metric space in its own right, since we can simply restrict d to $A \times A$.

One fundamental idea behind topological data analysis is best viewed by considering the special case where M is a finite collection of points in the Euclidean space $\mathbb{R}^n$. For such *point clouds*, there is a well-defined notion of *thickening* by any scale $\epsilon > 0$ — namely, $M^{+\epsilon}$ is the union of ϵ-balls in $\mathbb{R}^n$ around the points of M. Various thickenings are illustrated in the figure below.

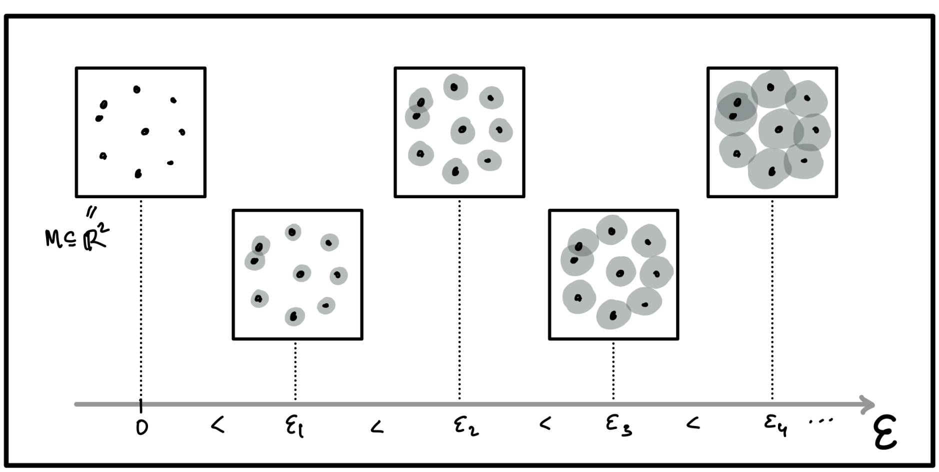

The hope is to better understand the geometric structure of M *across various scales*. One obstacle in this quest is that the union of balls $M^{+\epsilon}$ is a remarkably inconvenient object from the perspective of designing algorithms — for instance, if you were given a set of points $M \subset \mathbb{R}^2$ and a scale $\epsilon > 0$, how would you program a computer to determine whether or not $M^{+\epsilon}$ was connected? To address such questions, one replaces unions of balls by filtrations of simplicial complexes (which we encountered in Definition 1.6). There are two common choices of filtrations — **Vietoris-Rips** and **Čech**.[1]

DEFINITION 1.15. Let (M, d) be a finite metric space. The **Vietoris-Rips filtration** of M is an increasing family of simplicial complexes $\mathbf{VR}_\epsilon(M)$ indexed by the real numbers $\epsilon \geq 0$, defined as follows:

> a subset $\{x_0, x_1, \ldots, x_k\} \subset M$ forms a k-dimensional simplex in $\mathbf{VR}_\epsilon(M)$ if and only if the pairwise distances satisfy $d(x_i, x_j) \leq \epsilon$ for all i, j.

The astute reader may have noticed that we are indexing the simplicial complexes in this filtration by real numbers $\epsilon \geq 0$ rather than finite subsets of the form $\{1, 2, \ldots, n\}$ as demanded by Definition 1.6. The disparity between the two scenarios is artificial — since we have assumed that M is finite, there are only finitely many pairwise distances $d(x, y)$ encountered among the elements of M, so there are only finitely many ϵ values where new simplices are added to $\mathbf{VR}_\epsilon(M)$. Those who have not met Vietoris-Rips filtrations before can get better acquainted by verifying the following facts:

(1) the set $\mathbf{VR}_\epsilon(M)$ is a simplicial complex for each $\epsilon > 0$,
(2) the elements of M are vertices of each such $\mathbf{VR}_\epsilon(M)$, and
(3) for any pair $0 \leq \epsilon \leq \epsilon'$ of real numbers, $\mathbf{VR}_\epsilon(M)$ is a subcomplex of $\mathbf{VR}_{\epsilon'}(M)$.

We will see an example of a Vietoris-Rips filtration shortly; first let us examine the Čech alternative.

DEFINITION 1.16. Let M be a finite subset of a metric space (Z, d). The **Čech filtration** of M with respect to Z is the increasing family of simplicial complexes $\mathbf{C}_\epsilon$ indexed by $\epsilon \geq 0$ defined :

> a subset $\{x_0, x_1, \ldots, x_k\} \subset M$ forms a k-dimensional simplex in $\mathbf{C}_\epsilon(M)$ if and only if there exists some z in Z satisfying $d(z, x_i) \leq \epsilon$ for all i.

Although the larger metric space Z plays a starring role in deciding when a simplex lies inside $\mathbf{C}_\epsilon(M)$, it is customary to suppress it from the notation (in any case the typical scenario is $Z = \mathbb{R}^n$ with the Euclidean metric). This blatant dependence on Z is the biggest immediate difference between Čech filtrations and Vietoris-Rips filtrations — the Vietoris Rips filtration can be defined directly from knowledge of the metric on M whereas the Čech filtration can not.

To examine the key differences between these two filtrations, consider the three-element metric space (M, d) illustrated below.

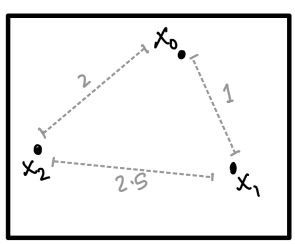

[1]This is pronounced "check".

The Vietoris Rips filtration of M at all scales $\epsilon \geq 0$ is given by the following lists of simplices:

$$\mathbf{VR}_\epsilon(M) = \begin{cases} \{x_0, x_1, x_2\} & 0 \leq \epsilon < 1 \\ \{x_0, x_1, x_2, x_0x_1\} & 1 \leq \epsilon < 2 \\ \{x_0, x_1, x_2, x_0x_1, x_0x_2\} & 2 \leq \epsilon < 2.5 \\ \{x_0, x_1, x_2, x_0x_1, x_0x_2, x_1x_2, x_0x_1x_2\} & \epsilon \geq 2.5 \end{cases}$$

It is crucial to note that the edge x_0x_2 and the 2-simplex $x_0x_1x_2$ enter the filtration at exactly the same scale, i.e., $\epsilon = 2.5$. Let us now contrast this with the Čech filtration for the same M, but now viewed as a subset of three points in the Euclidean plane $\mathbb{R}^2$. Here, the edge x_0x_2 and the 2-simplex $x_0x_1x_2$ will not appear simultaneously. Let $r > 0$ be the radius of the smallest ball which encloses all three points, like so:

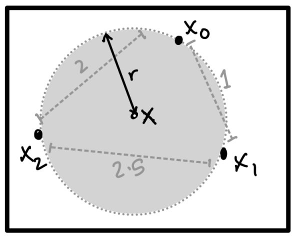

The Čech filtration of M as a subset of $\mathbb{R}^2$ is given by

$$\mathbf{C}_\epsilon(M) = \begin{cases} \{x_0, x_1, x_2\} & 0 \leq \epsilon < 0.5 \\ \{x_0, x_1, x_2, x_0x_1\} & 0.5 \leq \epsilon < 1 \\ \{x_0, x_1, x_2, x_0x_1, x_0x_2\} & 1 \leq \epsilon < 1.25 \\ \{x_0, x_1, x_2, x_0x_1, x_0x_2, x_1x_2\} & 1.25 \leq \epsilon < r \\ \{x_0, x_1, x_2, x_0x_1, x_0x_2, x_1x_2.x_0x_1x_2\} & \epsilon \geq r \end{cases}$$

Determining the radii of smallest enclosing balls (such as r above) is quite challenging algorithmically, which is why Vietoris-Rips filtrations are substantially easier to compute. On the other hand, the advantage of the Čech filtration is that it happens to be far more faithful to the underlying geometry of the union of balls $M^{+\epsilon}$ which we sought to approximate in the first place. For instance, given the union of ϵ-balls shown below, the Vietoris-Rips complex at scale 2ϵ is the solid 2-simplex (which fails to detect the hole) whereas the Čech filtration at scale ϵ equals the far more appropriate hollow 2-simplex.

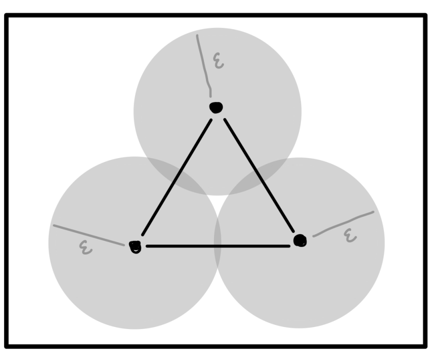

We will study this phenomenon much more carefully in the next Chapter.

1.7 Bonus: Local Geometry

The three notions introduced in this section (stars, links and cones) appear in the exercises of this Chapter and are invoked frequently in subsequent Chapters; but Theorem 1.20 below is not used anywhere else in this text.

Throughout this section, we fix a simplicial complex K as in Definition 1.1; our goal here is to describe the *neighborhood* of a given simplex σ in (the geometric realization of) K. The first step is to identify all the simplices which admit σ as a face.

Definition 1.17. The **open star** of σ in K is the collection of simplices

$$\mathbf{st}_K(\sigma) = \{\tau \text{ in } K \mid \sigma \leq \tau\}.$$

When the ambient simplicial complex K is clear from context (as it should be here), we simply denote the open star of each simplex σ by $\mathbf{st}(\sigma)$ rather than dragging K around in the subscript. The first panel below depicts (a part of) the geometric realization of a 2-dimensional simplicial complex; the open stars of the highlighted vertex v and edge e are shown in the next two panels (hollow vertices and dashed edges are *not* included).

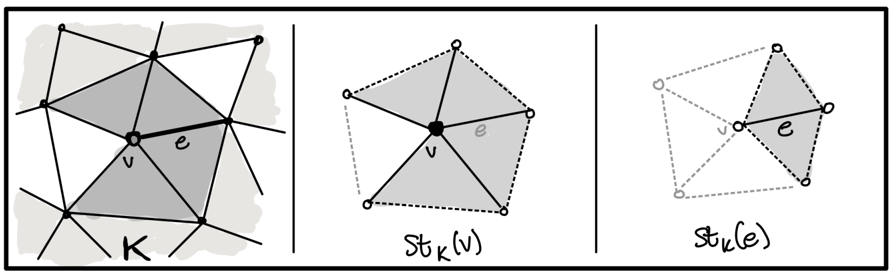

Clearly, the open star of σ describes a small simplicial neighborhood of σ in the geometric realization of K. Since $\mathbf{st}(\sigma)$ always contains σ, it is guaranteed to be non-empty — but as visible even in the simple examples drawn above, open stars are rarely subcomplexes of K since they tend to contain simplices without containing all of their faces. Passing to the closure of $\mathbf{st}(\sigma)$ as described in Definition 1.5 produces a bona fide subcomplex $\overline{\mathbf{st}}(\sigma) \subset K$, called the **closed star** of σ. Another useful subset of K that describes the local geometry of σ is called the **link**.

Definition 1.18. The **link** of σ in K is the collection $\mathbf{lk}_K(\sigma)$ of all simplices τ in K which simultaneously satisfy both $\tau \cup \sigma \in K$ and $\tau \cap \sigma = \varnothing$.

Unlike open stars, links of simplices in K might be empty (for example, the link of a top-dimensional simplex is always empty). But *if* the link of σ is non-empty, then it must be a subcomplex of K. Here are the links of the vertex v and edge e whose open stars we examined in the previous figure.

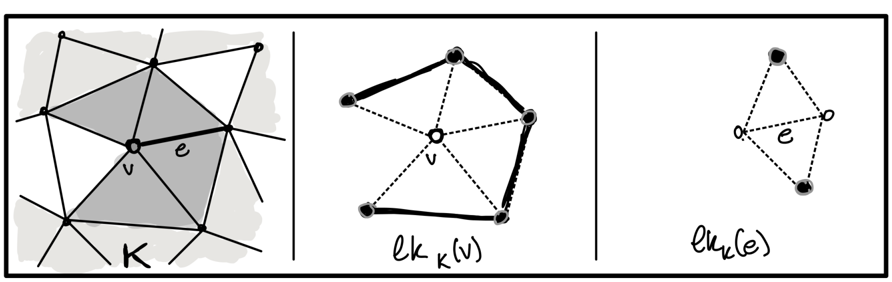

The final piece of the puzzle is the notion of a **cone** over a simplicial complex.

DEFINITION 1.19. The **cone** over K is a simplicial complex $\mathrm{Cone}(K)$ defined on the vertex set $K_0 \cup v_*$, where v_* is a new vertex not already present in K_0. For $d > 0$, a d-simplex of $\mathrm{Cone}(K)$ is either a d-simplex of K itself, or it is v_* adjoined with a $(d-1)$-simplex of K.

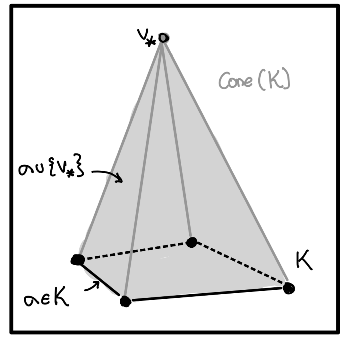

The reason this is called a cone becomes evident if one tries to draw the geometric realization of $\mathrm{Cone}(K)$ whenever $|K|$ is homeomorphic to a cicle, as shown here. Although we have banned empty simplicial complexes in Definition 1.1, it is convenient to adopt the convention that the cone over the empty set is just the one vertex v_*, i.e., the solid zero-simplex. It follows immediately from the definition of cones that $\dim \mathrm{Cone}(K)$ is always $1 + \dim K$.

The following result describes the smallest possible closed neighborhood of every d-dimensional simplex in an arbitrary simplicial complex: any such neighborhood decomposes into a product of the cone over the link of that simplex with the d-dimensional cube $[0,1]^d$.

THEOREM 1.20. *For any simplex σ in a simplicial complex K, there is a homeomorphism*

$$|\overline{\mathbf{st}}_K(\sigma)| \simeq |\mathrm{Cone}(\mathbf{lk}_K(\sigma))| \times [0,1]^{\dim \sigma}.$$

The left side here is the geometric realization of σ's closed star in K while the right side is a product of the geometric realization of σ's link with the closed unit cube in $\mathbb{R}^{\dim \sigma}$.

Here is an illustration of this product decomposition in the special case where σ is a 1-simplex that happens to be a face of three 2-simplices.

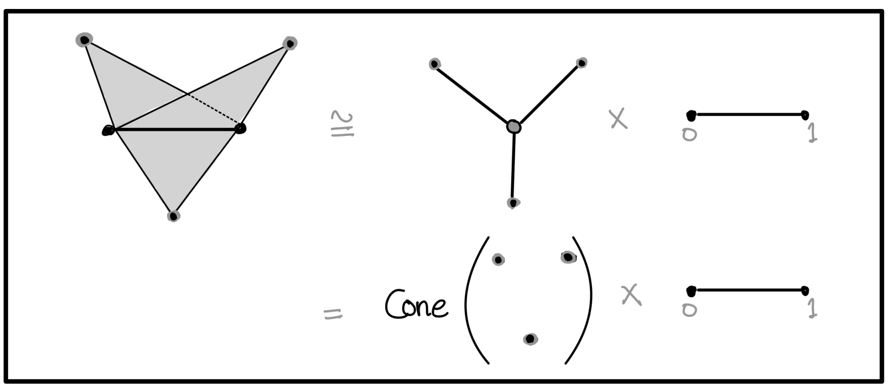

EXERCISES

EXERCISE 1.1. For each pair $i \leq k$ of non-negative integers, how many faces of codimension i does the solid k-simplex $\Delta(k)$ have?

EXERCISE 1.2. Show that the face relations between simplices in a finite simplicial complex satisfy the axioms of a *partially ordered set*.

EXERCISE 1.3. Show that the set of all subcomplexes of a finite simplicial complex K satisfy the axioms of a partially ordered set when ordered by containment $L \subset L'$.

EXERCISE 1.4. Either prove the following, or find a counterexample: if K is a simplicial complex and $L \subset K$ a subcomplex with $L \neq K$, then the complement $K - L$ is also a subcomplex of K.

EXERCISE 1.5. Assume that L, L' are two subcomplexes of a simplicial complex K with a nonempty intersection. Show that this intersection $L \cap L'$ is also a subcomplex of K.

EXERCISE 1.6. Let K be a k-dimensional simplicial complex, and for each dimension i in $\{0, 1, \ldots, k\}$ let n_i be the number of i-simplices in K. How many i-simplices does the barycentric subdivision $\mathbf{Sd}\, K$ have for each dimension i?

EXERCISE 1.7. Let M be a finite metric subspace of an ambient metric space (Z, d). Show, for each $\epsilon > 0$, that the associated Čech complex $\mathbf{C}_\epsilon(M)$ is a subcomplex of the Vietoris-Rips complex $\mathbf{VR}_{2\epsilon}(M)$. Then, show that – no matter what Z we had chosen – this $\mathbf{VR}_{2\epsilon}(M)$ is itself a subcomplex of $\mathbf{C}_{2\epsilon}(M)$.

EXERCISE 1.8. Consider any homeomorphism from $|\Delta(k)|$ to a closed k-dimensional disk for $k \geq 1$; where must this homeomorphism send the subspace $|\partial\Delta(k)|$?

EXERCISE 1.9. Let M be a finite subset of points in Euclidean space $\mathbb{R}^n$ (with its standard metric). As a function of n, can you find the *smallest* δ so that $\mathbf{VR}_\epsilon(M)$ is always a subcomplex of $\check{\mathbf{C}}_\delta(M)$? [Here the Čech complex has been constructed with respect to the ambient Euclidean space $\mathbb{R}^n$]

EXERCISE 1.10. If σ and τ are a pair of simplices in a simplicial complex K satisfying $\sigma \leq \tau$, show that $\mathbf{st}(\sigma) \supset \mathbf{st}(\tau)$ and $\mathbf{lk}(\sigma) \supset \mathbf{lk}(\tau)$.

EXERCISE 1.11. Show that if the link $\mathbf{lk}(\sigma)$ of a simplex σ in a simplicial complex K is nonempty, then $\mathbf{lk}(\sigma)$ is a subcomplex of K.

EXERCISE 1.12. Let σ be a simplex in a simplicial complex K. Show that a simplex τ lies in $\mathbf{lk}_K(\sigma)$ if and only if the following condition holds: the open stars of σ and τ have a non-trivial intersection *and* σ and τ have no common faces.

2. Homotopy

2.1 Basic Definitions

Given topological spaces X and Y, the set of all continuous functions from X to Y is typically quite large and complicated even in relatively simple cases (e.g., when both X and Y are the unit circle in $\mathbb{R}^2$). In order to study such functions, we are compelled to define interesting equivalence relations on them and restrict attention to equivalence classes. Among the deepest and most fruitful equivalence relations between functions $X \to Y$ is the notion of a *homotopy*.

Definition 2.1. Two continuous functions $F, G : X \to Y$ between topological spaces X and Y are **homotopic** if there is a third continuous function

$$\theta : X \times [0,1] \to Y$$

(called a **homotopy**) so that for all x in X, we have $\theta(x,0) = F(x)$ and $\theta(x,1) = G(x)$.

The requirement that θ also be continuous is absolutely essential here, since it is always possible to find discontinuous θ satisfying the requirements of this definition. Thus, the fundamental idea behind this definition is to put two functions in the same equivalence class whenever you can continuously interpolate from one to the other as a parameter $t \in [0,1]$ slides from 0 to 1. The picture below illustrates the homotopy equivalence of two maps F, G when X is a circle and Y is $\mathbb{R}^3$. These are homotopic if we can find a continuous θ from the cylinder $X \times [0,1]$ to $\mathbb{R}^3$ whose restriction to the lower boundary $X \times \{0\}$ coincides with F and restriction to the upper boundary $X \times \{1\}$ coincides with G.

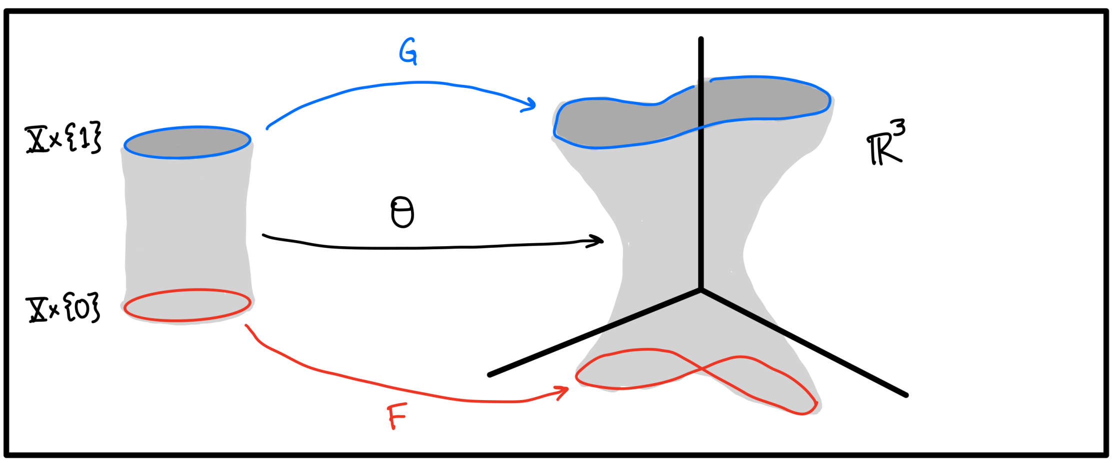

Homotopies between functions can be used in order to produce an equivalence relation on topological spaces as well.

Definition 2.2. Two topological spaces X and Y are **homotopy equivalent** if there are continuous maps $F : X \to Y$ and $G : Y \to X$ so that

(1) the composite $F \circ G$ is homotopic to the identity map on Y, while
(2) the composite $G \circ F$ is homotopic to the identity map on X.

A pair of continuous maps F and G satisfying the two conditions above are often called *homotopy inverses* of each other, although it is important to note that in general there is no uniqueness of such inverses — the set of homotopy inverses for a given F might contain several maps. Homotopy equivalence is a topological property that tends to be largely agnostic to metric information. The two panels below are designed to illustrate this phenomenon: in the first case, the 2-dimensional thickened figure-8 is homotopy equivalent to the thinner 1-dimensional figure-8 in its interior. But if we perturb this thinner curve ever so slightly to create a single loop, then homotopy equivalence no longer holds.

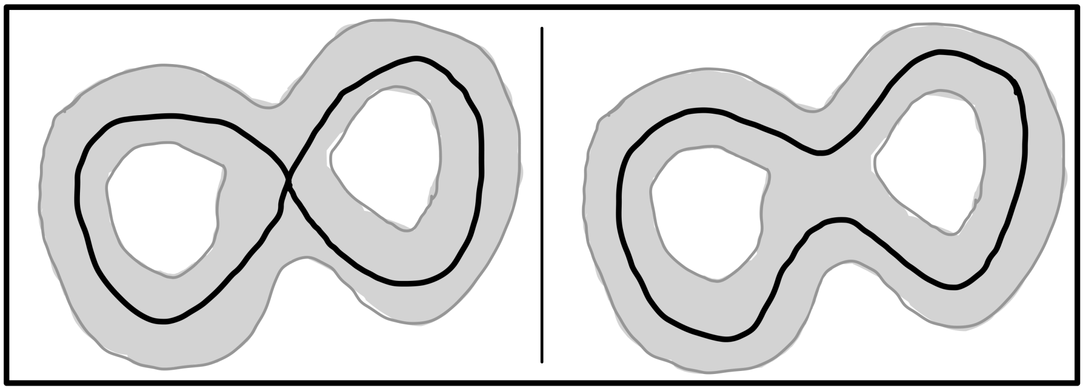

Two simplicial complexes K and L are said to be homotopy equivalent, or have the same *homotopy type*, whenever their geometric realizations $|K|$ and $|L|$ are homotopy equivalent in the sense of the definition above. It may not be immediately obvious that homotopy is an important equivalence relation between topological spaces — absorbing this fact takes time and experience. What should be clearer even at this early stage is that homotopy equivalence is far less rigid than homeomorphism: homeomorphic spaces are always homotopy equivalent, but the converse does not hold.

2.2 Contractible Spaces

The quest to study topological spaces up to homotopy equivalence has a natural starting point — we begin by asking which spaces are the least complicated from a homotopical perspective.

DEFINITION 2.3. A topological space X is **contractible** if it is homotopy equivalent to the one-point space.

You should check that X is contractible if and only if there exists some point $p \in X$ so that the identity map on X is homotopic to the constant map sending every point of X to p. In particular, the empty set $\varnothing$ is *not* contractible.

EXAMPLE 2.4. Here are several families of contractible simplicial complexes:

(1) **Solids**: for each $k \geq 0$ the solid k-simplex $\Delta(k)$ is contractible.
(2) **Cones**: the cone over *any* simplicial complex K (see Definition 1.19) is contractible.
(3) **Trees**: a tree is a connected graph with no cycles; these are all contractible.

We will prove the contractibility of these after developing some helpful machinery. For now, it is important to start building a mental database which contains as many contractible spaces as possible. The next few sections contain a suite of extremely powerful tools for detecting homotopy equivalence, and all of these tools rely in one way or another on your ability to recognize contractible spaces. The underlying reason for this dependence is the following vital result.

LEMMA 2.5. *Let X be a topological space and $k \geq 1$ an integer. If X is contractible, then any continuous map $F : |\partial\Delta(k)| \to X$ from the hollow k-simplex to X can be extended to a continuous map $F^+ : |\Delta(k)| \to X$ from the solid k-simplex.*

PROOF. Even the case $k = 1$ is quite insightful, so we will go over it carefully. Since $|\Delta(1)|$ is homeomorphic to the unit interval $[0,1]$ and $|\partial\Delta(1)|$ consists of the endoints $\{0,1\}$, we must show that X is *path-connected*, i.e., given any pair of points $F(0) = x_0$ and $F(1) = x_1$ in X, there is a continuous path in X from x_0 to x_1.

From Definition 2.3, we know that X is contractible if and only if there is some point $p \in X$ so that the identity map on X is homotopic to the constant map $X \to p$. Thus, there exists a homotopy $\theta : X \times [0,1] \to X$ satisfying $\theta(x,0) = x$ and $\theta(x,1) = p$ for all x in X. As we vary t from 0 to 1 for any given x in X, we obtain a continuous path $\theta(x,t)$ from x to our special point p — in particular both x_0 and x_1 admit paths to p. Thus, we can concatenate these two paths to get a path from x_0 to x_1 that passes through p; more explicitly, the desired extension $F^+ : [0,1] \to X$ is given in terms of θ by the piecewise-formula

$$F^+(t) = \begin{cases} \theta(x_0, 2t) & t \leq 1/2 \\ \theta(x_1, 2t-1) & t > 1/2. \end{cases}$$

This extension is continuous because at $t = 1/2$ both pieces are guaranteed to equal p. The following picture may help if the numerology of this formula is mysterious.

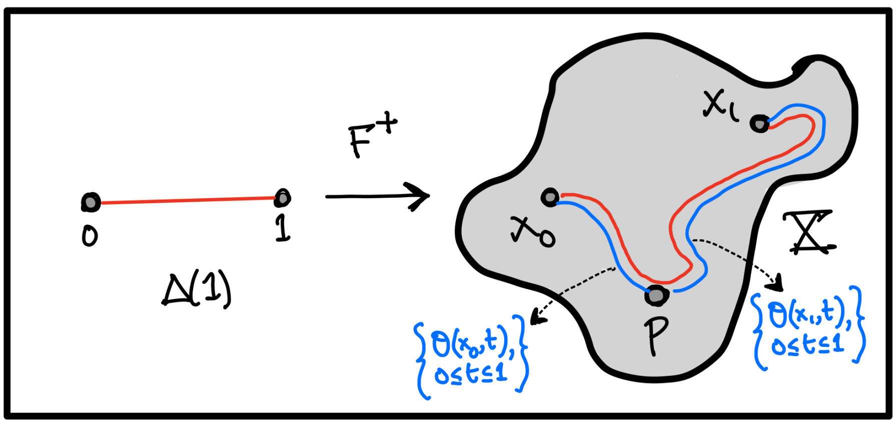

The argument for $k \geq 2$ is more technical and subscript-infested, but the basic principle remains the same — homotopies to constant maps allow us to "fill in" the F-images of hollow simplices to produce F^+-images of the corresponding solid simplices. $\square$

In the argument above, we used a homotopy θ to define an extension map without ever having an explicit formula for θ; this is quite typical because in general homotopies can get quite complicated even when relating simple maps between benign spaces. One refreshing exception to this unfortunate state of affairs is provided by the class of *straight-line homotopies*: given maps $f, g : X \to Y$ with $Y \subset \mathbb{R}^n$, one often attempts to use $\theta(x,t) = t \cdot f(x) + (1-t) \cdot g(x)$. Of course, there is no guarantee that the image of such a θ will actually lie in Y. Our next result highlights an important instance where this straight-line strategy succeeds.

PROPOSITION 2.6. *For each dimension $k \geq 0$, the solid k-dimensional simplex $\Delta(k)$ is contractible.*

PROOF. Let $\{x_0, \ldots, x_k\} \subset \mathbb{R}^n$ be any set of affinely independent points, so the geometric realization of $\Delta(k)$ is given (up to homeomorphism) by

$$|\Delta(k)| = \left\{ \sum_{i=0}^{k} t_i x_i \mid t_i \geq 0 \text{ and } \sum_{i=0}^{k} t_i = 1 \right\}.$$

Now consider the continuous map $\theta : |\Delta(k)| \times [0,1] \to \Delta(K)$ that sends each $x = \sum_{i=0}^{k} t_i x_i$ in $|\Delta(k)|$ and t in $[0,1]$ to the point

$$\theta(x,t) = [1 - t(1 - t_0)] \cdot x_0 + t \cdot \sum_{i=1}^{k} t_i x_i.$$

This formula prescribes a straight-line homotopy between the identity map (at $t = 1$) and the constant map (at $t = 0$) sending everything to x_0. Three routine verifications have been left as exercises: to complete the proof, one must show that $\theta(x,t)$ lies in $|\Delta(k)|$ for all t, that $\theta(x,0)$ is just the constant map to x_0, and that $\theta(x,1)$ is the identity map on $|\Delta(k)|$. □

Armed with knowledge of many contractible spaces, we are ready to explore a suite of homotopy equivalence detectors.

2.3 Carriers

Let K be a simplicial complex and X a topological space.

DEFINITION 2.7. A **carrier** C for K in X is an assignment of a subset $C(\sigma) \subset X$ to every simplex σ of K so that $C(\sigma) \subset C(\tau)$ holds whenever σ is a face of τ.

We say that C **carries** a continuous map $F : |K| \to X$ if for each simplex $\sigma \in K$ we have $F(|\sigma|) \subset C(\sigma)$. Similarly, we say that C carries a homotopy $\theta : |K| \times [0,1] \to Y$ if for each intermediate t in $[0,1]$ the map $\theta_t : |K| \to X$ given by

$$\theta_t(x) = \theta(x,t)$$

is carried by C in the sense described above. The next result is among the most powerful and widely-applicable tools for testing whether two maps $|K| \to X$ are homotopic.

LEMMA 2.8. **(The Carrier Lemma)** *Let C be a carrier for K in X. If the subset $C(\sigma) \subset X$ is contractible for each simplex $\sigma \in K$, then* **(a)** *there exists a continuous map $F : |K| \to X$ carried by C;* **(b)** *any two continuous maps $F, G : |K| \to X$ carried by C are homotopic; and* **(c)** *in fact, we can always choose a homotopy $\theta : |K| \times [0,1] \to X$ between F and G that is also carried by C.*

PROOF. Index the simplices of K as $\{\sigma_1, \sigma_2, \ldots, \sigma_m\}$ so that the faces of each simplex have lower indices than that simplex itself — this can be ensured for instance by indexing all the 0-dimensional simplices before all the 1-dimensional simplices, and so forth. There is now a filtration $\{\mathbf{S}_i K \mid 1 \leq i \leq m\}$ of K (see Definition 1.6) obtained by setting

$$\mathbf{S}_i K = \bigcup_{j \leq i} \sigma_j.$$

We will show **(b)** and **(c)** by induction on i; the argument for **(a)** is eerily similar and has been assigned as an exercise.

Base case: When $i = 1$ we must have a simplex σ_1 of minimum dimension, i.e., a vertex. By the hypotheses of this Theorem, the maps F and G send our vertex σ_1 to possibly distinct points (let's call them x_0 and x_1) in the contractible set $C(\sigma_1) \subset X$. The points x_0 and x_1 are evidently the image of a map $|\partial\Delta(1)| \to C(\sigma)$, so by Lemma 2.5 there is a path lying in $C(\sigma_1)$ from x_0 to x_1. This path prescribes a homotopy carried by C between the restrictions of F and G to $S_1 K = \sigma_1$.

Inductive step, part 1: Now let us assume that for some $i > 1$ the restrictions of F and G to $S_{i-1} K \subset K$ admit a homotopy $\theta : |\mathbf{S}_{i-1} K| \times [0,1] \to X$ carried by C. We must extend this θ continuously to the larger space $|\mathbf{S}_i K| \times [0,1]$; thus it suffices to define θ on the subset $|\sigma_i| \times [0,1]$, where σ_i is the unique simplex satisfying $S_i = S_{i-1} \cup \sigma_i$. Let $B \subset |\mathbf{S}_{i-1}|$ be the union of geometric realizations of all the faces $\tau \leq \sigma_i$ other than σ_i itself. Since all the $C(\tau)$ are subsets of $C(\sigma_i)$ by Definition 2.7, we note that the image $\theta(B \times [0,1])$ is entirely contained within $C(\sigma_i)$. Moreover, by the requirement that C carries F and G, both $F(|\sigma_i|)$ and $G(|\sigma_i|)$ also lie inside $C(\sigma_i)$.

Inductive step, part 2: The key observation here is as follows: writing $d = \dim \sigma_i$, the product $|\sigma_i| \times [0,1]$ is homeomorphic to $|\Delta(d)| \times [0,1]$, which in turn is homeomorphic to $|\Delta(d+1)|$. Consequently, the boundary[1] of $|\sigma_i| \times [0,1]$ is homeomorphic to the subset

$$|\partial\Delta(d+1)| \simeq \Big(|\partial\Delta(d)| \times [0,1]\Big) \cup \Big(|\Delta(d)| \times \{0,1\}\Big).$$

Here is a figure illustrating these spaces for $d = 2$:

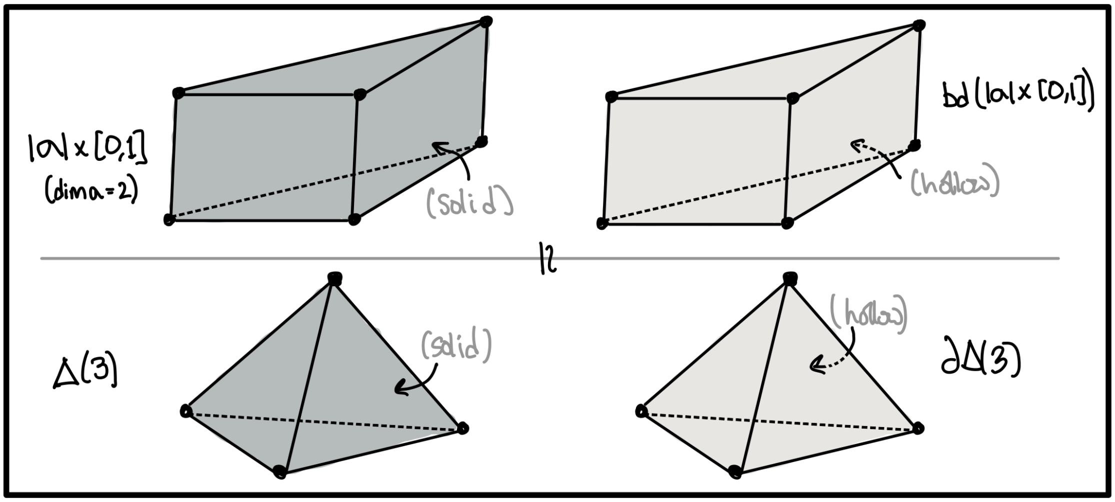

Now the first piece of this union $|\partial\Delta(d)| \times [0,1]$ is homeomorphic to $B \times [0,1]$ while the second piece is homeomorphic to two disjoint copies of $|\sigma_i|$. Our homotopy θ sends the first piece to $C(\sigma_i)$ by **part 1** of the inductive step. As for the second piece, we know that

$$\theta(|\sigma_i|, 0) = F(|\sigma_i|) \subset C(\sigma_i).$$

Here the equality follows from Definition 2.1 while the containment is a consequence of the assumption that C carries F. Similarly, we also have $\theta(|\sigma_i|, 1) = G(|\sigma_i|)| \subset C(\sigma_i)$. So up to homeomorphism, θ constitues a map from the entire boundary $\partial\Delta(d+1)$ to the contractible set $C(\sigma_i) \subset X$. Lemma 2.5 guarantees a continuous extension $\theta^+ : |\Delta(d+1)| \to C(\sigma_i)$, and using the homeomorphism $\Delta(d+1) \simeq |\sigma_i| \times [0,1]$ gives us the desired continuous extension of θ to $|\sigma_i| \times [0,1]$. ☐

The utility of the Carrier lemma in homotopically-oriented problems is difficult to overstate. Here is a simple consequence designed to work directly with simplicial maps. We say that two simplicial maps $f, g : K \to L$ are **contiguous** if for any simplex σ of K, the union $f(\sigma) \cup g(\sigma)$ is a simplex of L.

COROLLARY 2.9. *If $f, g : K \to L$ are contiguous, then they must be homotopic.*

PROOF. For each simplex σ in K, let $C(\sigma) \subset |L|$ be the geometric realization of the union-simplex $f(\sigma) \cup g(\sigma)$. This assignment C prescribes a carrier for K in $|L|$; clearly, C carries both f and g. And finally, since solid simplices are contractible by Proposition 2.6, the desired conclusion follows from Lemma 2.8 **(b)**. ☐

[1]Here we have used the fact that the boundary of a product $\mathrm{bd}(P \times Q)$ is the union $(P \times \mathrm{bd}\, Q) \cup (\mathrm{bd}\, P \times Q)$.

This result has satisfying and immediate applications: for instance, we can now easily show that $\mathrm{Cone}(K)$ is contractible for any simplicial complex K. Writing v_* for the additional vertex as in Definition 1.19, apply Corollary 2.9 to the case where f is the identity map on $\mathrm{Cone}(K)$ while g is the map sending every vertex to v_*.

2.4 Fibers

Let $f : K \to L$ be a simplicial map; for each simplex τ in L, the **fiber of f under τ** is the collection of simplices in K given by

$$\tau/f = \{\sigma \in K \mid f(\sigma) \leq \tau\}. \tag{2}$$

Each such fiber is a subcomplex of K; and moreover, τ/f is a subcomplex of τ'/f whenever $\tau \leq \tau'$ in L. We will use the Carrier lemma three times below to show that simplicial maps with contractible fibers induce homotopy equivalences — this forms a special case of a far more general result called *Quillen's Theorem A*.

THEOREM 2.10. **(Quillen's Fiber Theorem)** *Let $f : K \to L$ be a simplicial map. If the fiber τ/f is contractible for every simplex τ in L, then the induced continuous map $|f| : |K| \to |L|$ admits a homotopy inverse $G : |L| \to |K|$; and in particular, K and L are homotopy equivalent.*

PROOF. For each simplex τ of L, let $C(\tau) \subset |K|$ be the geometric realization of the fiber τ/F; this provides a carrier for L in $|K|$ with each $C(\tau)$ contractible, so by Lemma 2.8 **(a)** we know that there exists a continuous $G : |L| \to |K|$ satisfying $G(|\tau|) \subset C(\tau) = |\tau/f|$ for all τ in L. We will confirm that any such G is a homotopy inverse for $|f|$.

1. $|f| \circ G$ **is homotopic to the identity on** L**:** for each simplex τ in L, we have the containment

$$|f| \circ G(|\tau|) \subset |\tau|,$$

simply because $G(|\tau|)$ is contained in $|\tau/f|$. Therefore, the assignment $C_L(\tau) = |\tau|$ prescribes a carrier (for L in $|L|$) which carries both $|f| \circ G$ and the identity map on L. Since each $|\tau|$ is contractible by Proposition 2.6, we have from Lemma 2.8 **(b)** that $|f| \circ G$ is homotopic to the identity on L as desired.

2. $G \circ |f|$ **is homotopic to the identity on** K**:** for each simplex σ in K, we know from Proposition 1.11 that the $|f|$-image of $|\sigma|$ is exactly $|f(\sigma)| \subset |L|$. Recall that by our construction of G, we have the containment

$$G(|f(\sigma)|) \subset C(f(\sigma)) = |f(\sigma)/f|.$$

So if we define C_K to be the carrier for K in $|K|$ given by $C_K(\sigma) = |f(\sigma)/f|$, we know that C_K carries $G \circ |f|$. Note also that σ automatically lies in $f(\sigma)/f$ by (2), so C_K also carries the identity map on K. Since each $C_K(\sigma)$ is contractible by our assumption on the fibers of f, a final appeal to Lemma 2.8 **(b)** concludes the argument. $\square$

The strength of Quillen's fiber theorem lies in the fact that it allows us to conclude homotopy equivalence of simplicial complexes K and L given only a one-way simplicial map $f : K \to L$. As long as this f has contractible fibers, one is not required to painstakingly construct an explicit homotopy inverse $|L| \to |K|$.

2.5 Nerves

A finite **open cover** $U_\bullet$ of a topological space X is a collection of open subsets $U_\alpha \subset X$ (here α ranges over some finite index set A) satisfying

$$X = \bigcup_{\alpha \in A} U_\alpha.$$

By keeping track of how the different U_α intersect one another, we can build a simplicial complex on the vertex set A; the hope is to appropriately constrain the cover so that this simplicial complex is homotopy equivalent to X.

DEFINITION 2.11. The **nerve** $\mathbf{N}(U_\bullet)$ of an open cover $\{U_\alpha \mid \alpha \in A\}$ of a topological space X is the simplicial complex whose i-simplices are given by all subsets $\sigma \subset A$ of cardinality $(i+1)$ for which the intersection

$$\mathbf{Supp}(\sigma) := \bigcap_{\alpha \in \sigma} U_\alpha$$

is nonempty.

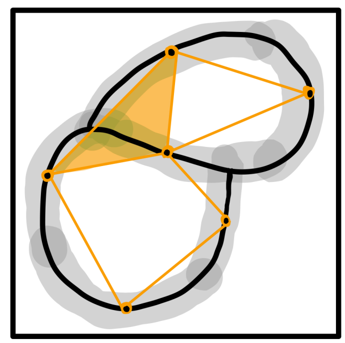

This intersection $\mathbf{Supp}(\sigma) \subset X$ is called the *support* of the simplex σ, and those encountering this notion for the first time should beware that $\sigma \leq \tau$ in $\mathbf{N}(U_\bullet)$ means $\mathbf{Supp}(\sigma) \supset \mathbf{Supp}(\tau)$ as subsets of X. In particular, the vertices of $|N(U_\bullet)|$ have larger supports than the edges which admit them as faces, and so on.

Having gone through the effort of finding an open cover $U_\bullet$ of a topological space X, one wonders to what extent the homotopy type of X is captured by the geometric realization $|N(U_\bullet)|$ of the associated nerve. The task appears absolutely hopeless at first glance — for instance, we could always choose $U_\bullet$ to consist of a single subset $U_1 = X$, in which case its nerve is just $\Delta(0)$ regardless of X. As with most of the other results described here, the key to solving this problem is contractibility. If we require all nonempty supports to be contractible subsets of X, then the following miracle occurs.

THEOREM 2.12. **(The Nerve theorem)** *Let $\{U_\alpha \mid \alpha \in A\}$ be a finite open cover of a topological space X. If each simplex $\sigma \in N(U_\bullet)$ has contractible support $\mathbf{Supp}(\sigma) \subset X$, then $|N(U_\bullet)|$ is homotopy equivalent to X.*

PROOF. Let $X(U_\bullet)$ be the subset of the product $X \times |N(U_\bullet)|$ containing all pairs (x,u) for which there is a simplex σ in $N(U_\bullet)$ satisfying both $x \in \mathbf{Supp}(\sigma)$ and $u \in |\sigma|$. There are natural projection maps from $X(U_\bullet)$ to both X and $|N(U_\bullet)|$:

$$\begin{tikzcd} & X(U_\bullet) \arrow[dl, two heads, "p"'] \arrow[dr, two heads, "q"] & \\ X & & {|N(U_\bullet)|} \end{tikzcd}$$

In particular, $p(x,u) = x$ and $q(x,u) = u$ for every (x,u) in $X(U_\bullet)$. Next, we show that for each point x in X and u in $|N(U_\bullet)|$, the fibers $p^{-1}(x)$ and $q^{-1}(u)$ are contractible subsets of $X(U_\bullet)$.

1. Fibers of p: For each point x in X, the fiber $p^{-1}(x) \subset X(U_\bullet)$ is homeomorphic to the set of all $u \in |N(U_\bullet)|$ lying in the realizations of simplices σ whose supports contain x. But all such σ must be faces of the single simplex σ_x in $N(U_\bullet)$ whose support is the intersection of *all*

U_α satisfying $x \in U_\alpha$. Thus, $p^{-1}(x)$ is homeomorphic to the geometric realization of $|\sigma_x|$, which must be contractible by Proposition 2.6.

2. Fibers of q: Given u in $|N(U_\bullet)|$, let $\sigma_u \in N(U_\bullet)$ be the unique simplex containing u in the interior of its realization. The fiber $q^{-1}(u) \subset X(U_\bullet)$ is homeomorphic to the support $\mathbf{Supp}(\sigma_u) \subset X$, which is contractible by assumption.

3. Finale: There is a variant of Theorem 2.10 which applies to a large class of continuous (not necessarily simplicial) maps between metric spaces (not necessarily simplicial complexes). In particular, this result implies that sufficiently well-behaved maps — such as our p and q — induce homotopy equivalences if their fibers over all *points* of their codomains are contractible[2]. An appeal to this modified fiber theorem establishes that X and $|N(U_\bullet)|$ are both homotopy equivalent to $X(U_\bullet)$ via p and q respectively, so the desired conclusion follows. □

There are at least three things to be noted about the Nerve theorem and its proof. First, it was really convenient to have a fiber theorem at our disposal — not only did we avoid having to build any homotopic inverses, but we even managed to avoid building a one-way map relating X to $|N(U_\bullet)|$. Second, the Nerve theorem gives us a mechanism for going back from topological spaces to simplicial complexes; in that sense, it constitutes a sort of converse to geometric realizations from Definition 1.7. And third, this theorem guarantees that Čech filtrations from Definition 1.16 accurately capture the homotopy type of the underlying union of balls at each scale.

COROLLARY 2.13. *Let $M \subset \mathbb{R}^n$ be a finite set of points. For each radius $\epsilon > 0$, the union $M^{+\epsilon} \subset \mathbb{R}^n$ of radius ϵ Euclidean balls around the points of M is homotopy equivalent to the geometric realization of the Čech complex $\mathbf{C}_\epsilon(M)$.*

PROOF. For each point x in M, let $B_\epsilon(x)$ be the open ball of radius ϵ around x. By definition of $M^{+\epsilon}$, we have

$$M^{+\epsilon} = \bigcup_{x \in M} B_\epsilon(x),$$

so the collection $\{B_\epsilon(x) \mid x \in M\}$ constitutes an open cover of $M^{+\epsilon}$. The Čech complex $\mathbf{C}_\epsilon(M)$ is precisely the nerve of this cover, so the desired conclusion follows from the Nerve theorem if we can show that nonempty intersections of Euclidean balls are contactible. Such intersections are always *convex* subsets of $\mathbb{R}^n$, and their contractibility will be established in one of the Exercises of this Chapter. □

There are no homotopical guarantees analogous to the above result which apply to the Vietoris-Rips filtration.

2.6 ELEMENTARY COLLAPSES

There is a simple combinatorial operation on simplicial complexes which allows us to find homotopy-equivalent subcomplexes by performing a series of moves; each such move removes two adjacent simplices $(\sigma < \tau)$ at a time, and has a very concrete and algorithmic flavor. For instance, one can show that $\Delta(2)$ is contractible simply by drawing the following diagram:

[2]For details, see the main result of S Smale's 1957 paper *A Vietoris Mapping Theorem for Homotopy*.

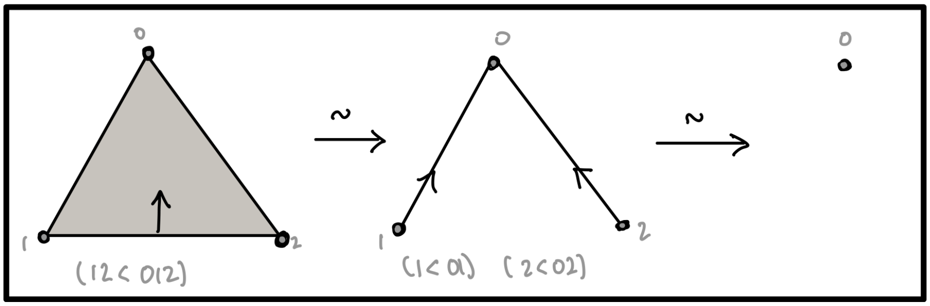

Our goal in this section is to describe these homotopy-preserving moves.

Let K be a simplicial complex. We call two distinct simplices $(\sigma < \tau)$ of K a **free face pair** if the open star of σ (see Definition 1.17) satisfies $\mathbf{st}_K(\sigma) = \{\sigma, \tau\}$. For such a pair we immediately have $\dim \tau = \dim \sigma + 1$; moreover, there can be no other simplices in K (besides σ and τ) which admit σ as a face.

PROPOSITION 2.14. *If $(\sigma < \tau)$ is a free face pair in K, then the collection $K' = K - \{\sigma, \tau\}$ forms a subcomplex of K, and in fact K' is homotopy equivalent to K.*

PROOF. Assume for the sake of contradiction that some simplex γ in K' is missing a face; such a γ would have to satisfy $\gamma > \sigma$ in K; this forces $\mathbf{st}_K(\sigma)$ to contain γ and violates our free face assumption. Thus, $K' \subset K$ is a subcomplex. To see the desired homotopy equivalence to K, consider the following figure:

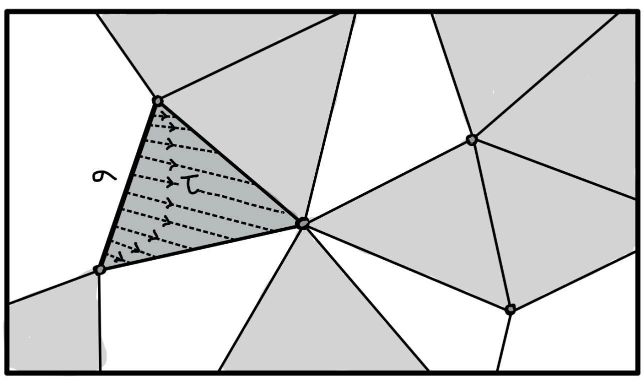

There is a map $r : |K| \to |K'|$ which is the identity away from $|\tau|$ and sends all points of $|\tau|$ along straight line segments to points in the union $\bigcup_{\sigma \neq \eta < \tau} |\eta|$ of realizations of all faces of τ except σ. This map serves as a homotopy inverse to the inclusion $i : K' \hookrightarrow K$; on the one hand, the composite $r \circ |i|$ is the identity map on $|K'|$. And on the other hand, these straight line segments generate a homotopy from $|i| \circ r$ to the identity map on K. □

The removal of a free face pair $(\sigma < \tau)$ from K is called an **elementary collapse**. These can be iterated, as shown in our diagrammatic reduction of $\Delta(2)$ to $\Delta(0)$ drawn above. One important point to note, visible already in the figure above, is that the subcomplex $K' = K - \{\sigma, \tau\}$ might contain free face pairs that were unavailable in K: when we remove the pair $(12 < 012)$ from $\Delta(2)$, the pairs $(1 < 01)$ and $(2 < 02)$ become free and can be safely removed in the second step. We say that K **collapses** onto a subcomplex L if there is a filtration (as in Definition 1.6) of the form

$$L = \boldsymbol{F}_1 K \subset \boldsymbol{F}_2 K \subset \cdots \subset \boldsymbol{F}_n K = K$$

where each $\boldsymbol{F}_i K$ is obtained by removing a single free-face pair from the subsequent $\boldsymbol{F}_{i+1} K$. By Proposition 2.14, all the $\boldsymbol{F}_i K$ are homotopy equivalent to each other in this case. Thanks to their simple combinatorial nature, elementary collapses can be algorithmically implemented on a computer.

2.7 BONUS: SIMPLICIAL APPROXIMATION

The contents of the section are not used elsewhere in this text; they have been included here because Theorem 2.15 described below is a foundational result in simplicial algebraic topology. It allows us to study homotopy classes functions between (geometric realizations of) simplicial complexes using only simplicial maps rather than arbitrary continuous ones.

Here is a fairly natural challenge in light of our quest to understand simplicial complexes up to homotopy equivalence.

> *Assume that $F : |K| \to |L|$ is a continuous map between the geometric realizations of two simplicial complexes K and L. Does there exist a simplicial map $f : K \to L$ so that $|f|$ is homotopic to F?*

Unfortunately, the answer to this question as stated is *no*. One way to see why (without doing any heavy computations) is to note that the set of all simplicial maps $K \to L$ is always finite, so it is unreasonable to expect simplicial maps to attain all possible homotopy types achievable by the (typically *very* infinite) set of continuous maps $|K| \to |L|$. The good news, however, is that the answer to our challenge becomes yes *if* we give ourselves the ability to barycentrically subdivide the domain K finitely many times (as described in Definition 1.12). The following result is called the **simplicial approximation theorem**.

THEOREM 2.15. *Let $F : |K| \to |L|$ be a continuous map between the geometric realizations of two simplicial complexes. There exists an integer $n \geq 0$ and a simplicial map $f : \mathbf{Sd}^n K \to L$ so that $|f|$ is homotopic to F.*

In general there is no known bound on how many barycentric subdivisions of K might be required to build the simplicial approximation f for a given F.

EXERCISES

EXERCISE 2.1. Prove that homotopy equivalence is an equivalence relation on the class of all topological spaces.

EXERCISE 2.2. Show that if two topological spaces X and Y are homeomorphic, then they must also be homotopy equivalent.

EXERCISE 2.3. Show that if Y is a contractible space, then then for any topological space X the product $X \times Y$ is homotopy equivalent to X.

EXERCISE 2.4. Show that if Y is contractible then any pair of maps $f, g : X \to Y$ are homotopic.

EXERCISE 2.5. A subset $P \subset \mathbb{R}^n$ is said to be *convex* if for every pair of points x, y in P the line segment $\{tx + (1-t)y \mid 0 \leq t \leq 1\}$ lies inside P. Show that every nonempty convex set is contractible.

EXERCISE 2.6. Show that the subspaces $X = \{0\}$ and $Y = \{0, 1\}$ of the real line $\mathbb{R}$ are not homotopy equivalent.

EXERCISE 2.7. Prove the assertion **(a)** from Lemma 2.8. [Hint: the filtration S_iK, the inductive strategy and Lemma 2.5 are all useful here].

EXERCISE 2.8. Consider two simplicial maps from $\partial\Delta(2)$ to the illustrated simplicial complex described as follows. The first one sends vertices $\{0, 1, 2\}$ to $\{a_0, a_1, a_2\}$ in order, while the second one sends the same vertices to $\{b_0, b_1, b_2\}$ respectively. Show that these two maps are homotopic, keeping in mind that the left-most edge in the figure is identified with the right-most edge. [Hint: first show that the unit square $[0, 1] \times [0, 1]$ is contractible by Proposition 2.6 plus Exercise 2.3, and then apply Lemma 2.8]

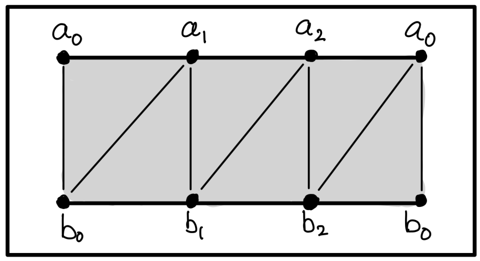

EXERCISE 2.9. Given a simplicial map $f : K \to L$, show that for each simplex τ in L the fiber τ/f as defined in (2) is a subcomplex of K; also show that τ/f is a subcomplex of τ'/f whenever $\tau \leq \tau'$ holds in L.

EXERCISE 2.10. Find the smallest open cover of the circle with contractible supports. What is the nerve of this cover?

EXERCISE 2.11. Find a cover of the circle containing at least two open sets which violates the hypotheses of the nerve lemma. What is the nerve of this bad cover?

EXERCISE 2.12. Show that trees (connected graphs with no cycles) are simple homotopy-equivalent to $\Delta(0)$, and hence contractible. [Hint: induction on edges plus Proposition 2.14].

EXERCISE 2.13. Use a suitable sequence of elementary collapses to show that the simplicial complex drawn in Exercise 2.8 is homotopy-equivalent to the subcomplex consisting of the simplices $\{b_0, b_1, b_2, b_0b_1, b_1b_2, b_0b_2\}$.

3. HOMOLOGY

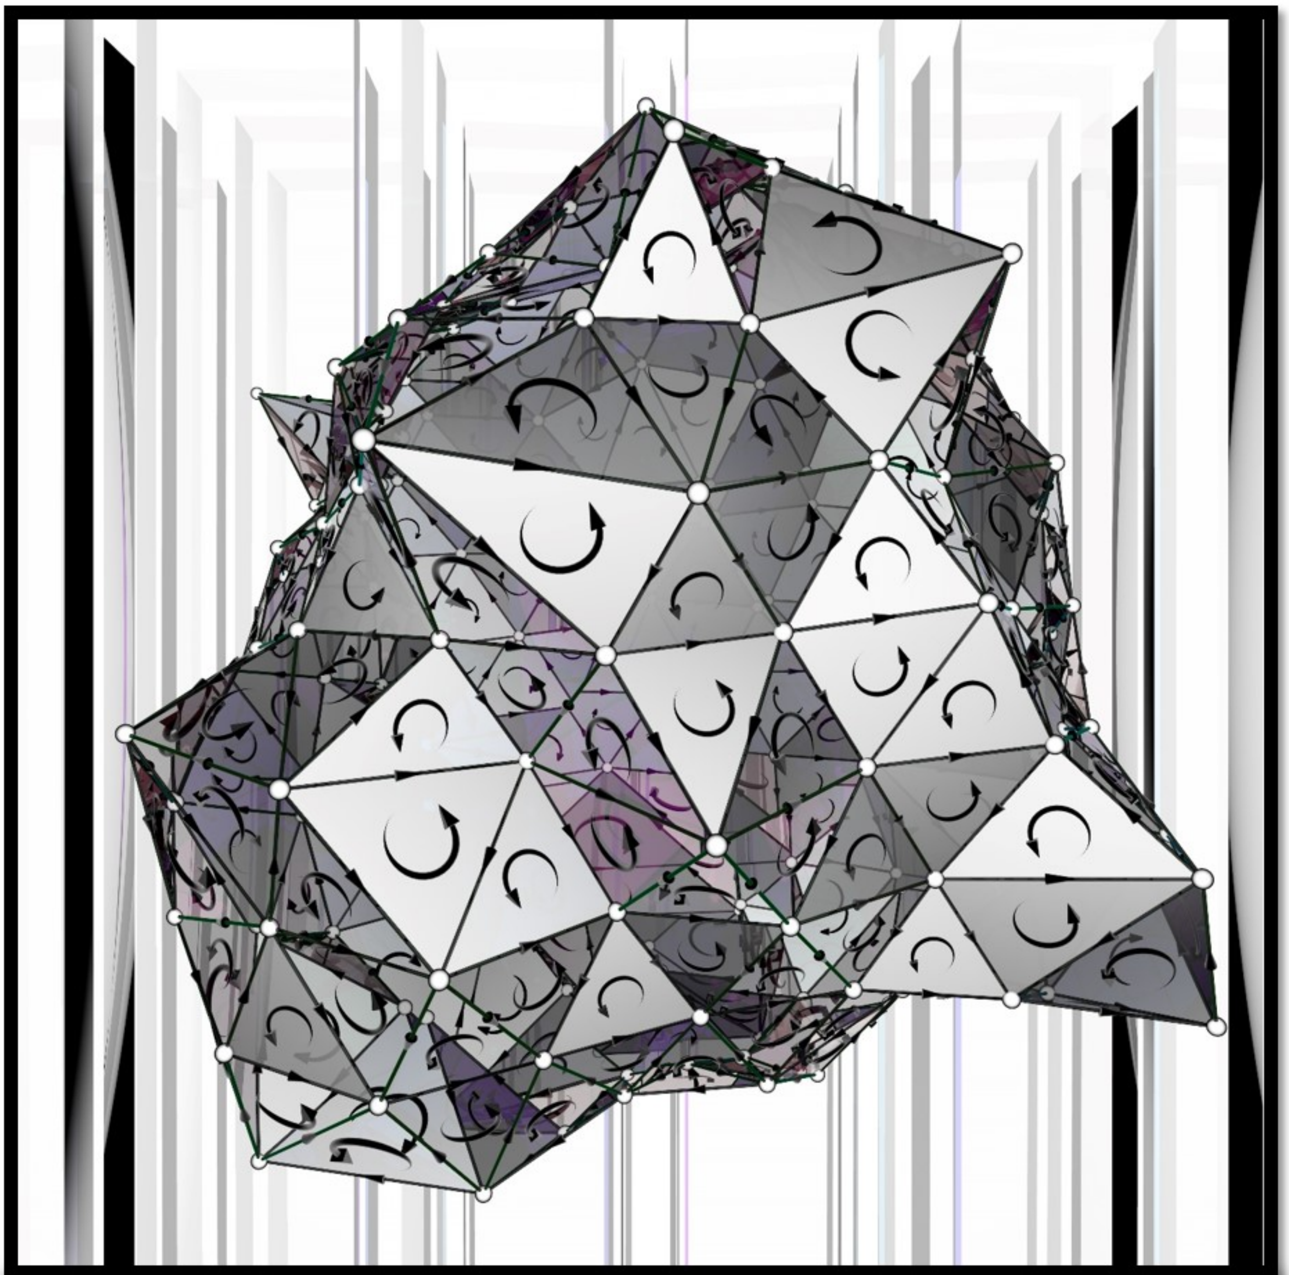

3.1 Euler Characteristic

Despite all the lemmas and theorems described in the previous Chapter, it remains difficult to explicitly compute – by hand or machine – whether two simplicial complexes are homotopy equivalent or not. As a consequence, we are forced to seek computable **homotopy invariants**; these are assignments $K \mapsto I(K)$ sending each simplicial complex K to some algebraic object $I(K)$ so that the following crucial property is satisfied. If K is homotopy equivalent to some other simplicial complex L, then $I(K)$ is equal (or at least isomorphic in a suitable sense) to $I(L)$. This invariance is necessarily a one-way street: we can not require $I(K) = I(L)$ to imply that K and L are homotopy equivalent, otherwise I-equivalence would be just as difficult to compute as homotopy equivalence. The oldest and simplest homotopy invariant is the **Euler characteristic**, defined as follows.

Definition 3.1. The **Euler characteristic** of a simplicial complex K is the integer $\chi(K) \in \mathbb{Z}$ given by the alternating sum of cardinalities

$$\chi(K) = \sum_{i=0}^{\dim K} (-1)^i \cdot \#K_i.$$

(Here $\#K_i$ indicates the number of all i-dimensional simplices in K).

Reducing an entire simplicial complex to a single integer might seem absurd at first glance, but this simple definition conceals several interesting mysteries. For one thing, it will turn out that χ is a homotopy invariant: if K and L are homotopy equivalent simplicial complexes, then $\chi(K) = \chi(L)$. This fact is by no means obvious, and even verifying it in the case where $L = \mathbf{Sd}\, K$ from the above definition appears painful. Algebraic topology was created to explain why so crude a summary of K should remain entirely unaffected by (combinatorially) enormous perturbations of K which lie within the same homotopy class. For instance, a simple computation reveals that the solid 0-simplex satisfies $\chi(\Delta(0)) = 1$, so by homotopy invariance we can immediately conclude that all contractible simplicial complexes (including $\Delta(k)$ for other choices of $k > 0$) also have $\chi = 1$.

Setting aside this mystery of homotopy invariance for the moment, one may wish to take a moment to marvel at how easily we can compute $\chi(K)$ for a single K. But even here, there are good reasons to tread with caution as described below.

Example 3.2. Let $S \subset \mathbb{N}$ denote the set of all **square-free** natural numbers — this consists of all those $n > 1$ which can be expressed as a product of *distinct* prime numbers. Thus, the first few numbers in S are $(2, 3, 5, 6 = 2 \times 3, \ldots)$; note that 4 is excluded because it equals 2^2.

For each $n \geq 1$, let $K(n)$ be the simplicial complex defined on the vertex set $V(n) = \{s \in S \mid s \leq n\}$ by the following rule: the k-simplices for $k > 0$ are all subsets $\{s_0, \ldots, s_k\} \subset V(n)$ of vertices so that each s_i divides the subsequent s_{i+1}. This is clearly a simplicial complex for each integer n, since the divisibility property is preserved when passing to subsets. Moreover, we have an inclusion $K(n) \subset K(n+1)$ for all n, so these simplicial complexes $K(n)$ constitute a filtration of unbounded length. The statement

$$\lim_{n \to \infty} \frac{|\chi(K(n))|}{n^{1/2+\epsilon}} = 0 \text{ for all } \epsilon > 0$$

is equivalent to the Riemann hypothesis. For more information on these simplicial complexes $K(n)$, see A. Bjorner's 2011 paper *A Cell Complex in Number Theory.*

The best way to see that χ is homotopy invariant is to recast it as a numerical reduction of a richer homotopy invariant; this richer invariant is called **homology**, and it forms the main theme of this Chapter.

3.2 Orientations and Boundaries

An **orientation** of a simplicial complex K is an injective function $o : K_0 \to \mathbb{N}$ which assigns unique natural numbers to vertices. The number assigned to each vertex will not be as important as the relative ordering of vertices induced by o, so we may as well require o to take values in the first $\#K_0$ natural numbers. Given an orientation of K, we will always write simplices as *ordered* subsets of K_0 — rather than writing each k-simplex as an unstructured set of vertices, we can uniquely write it as a tuple $(v_0, v_1, \ldots, v_k)$ inside $K_0 \times \cdots \times K_0$ satisfying

$$o(v_0) < o(v_1) < \cdots < o(v_k).$$

We call this ordered-tuple of vertices an *oriented simplex*.

DEFINITION 3.3. Let K be an oriented simplicial complex and let $\sigma = (v_0, \ldots, v_k)$ be an oriented k-simplex in K. For each i in $\{0, 1, \ldots, k\}$, the ***i*-th face** of σ is the $(k-1)$-dimensional simplex

$$\sigma_{-i} = (v_0, \ldots, v_{i-1}, \cancel{v_i}, v_{i+1}, \ldots, v_d)$$

obtained by removing the i-th vertex.

In the absence of an orientation, there is no coherent way to identify the i-the vertex of σ, so σ_{-i} is not a well-defined simplex — when being explicit about the choice of orientation, one may wish to write the i-th face of σ as σ^o_{-i}. Here are the ordered faces of the top-dimensional simplices of $\Delta(1)$ and $\Delta(2)$ if we assume the standard orientation on the vertices, i.e., $0 < 1 < 2$:

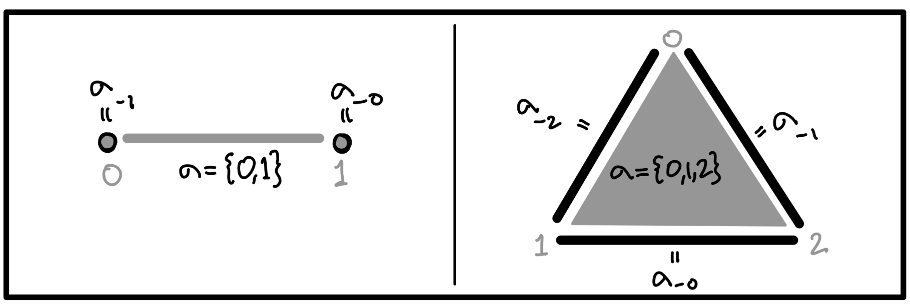

Oriented simplicial complexes form higher-dimensional generalizations of certain types of *directed graphs* — recall that a graph $G = (V, E)$ is directed if each edge comes with a preferred direction, usually indicated as an arrow from one vertex (the source) to the other (the target). The **incidence matrix** $I = I(G)$ of such a graph has the vertices V indexing its rows and edges E indexing its columns; the entry in row v and column e of I is given by the pleasant rule

$$I_{v,e} = \begin{cases} -1 & \text{if } u \text{ is the source of } e, \\ +1 & \text{if } u \text{ is the target of } e, \\ 0 & \text{ otherwise.} \end{cases}$$

Thus, each column of I contains exactly one $+1$ and one -1, with all other entries necessarily being zero. In contrast, the rows of I are far less structured; the number of ± 1 entries in each rows depends on the number of edges for which the associated vertex forms a source or target. If we label the vertices of an undirected graph by distinct natural numbers, we automatically induce a directed structure on the edges by forcing sources to have smaller values than targets.

A small example of a directed graph constructed via this vertex-labelling method is depicted below along with its adjacency matrix.

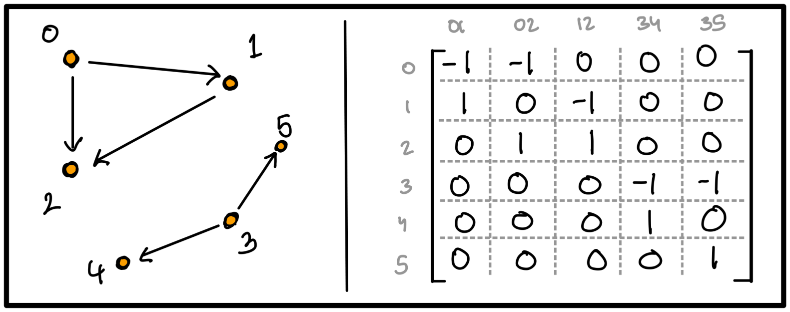

At first glance, the matrix I appears to simply be an alternate way to encode the structure of G in a manner that is specifically tailored to the needs of computers (or equivalently, algebraists). The advantages of this encoding become clearer when I is treated as a linear map — if we let $\mathbb{R}[V]$ and $\mathbb{R}[E]$ be the real vector spaces obtained by treating the vertices and edges of G respectively as orthonormal bases, then I prescribes a linear map $\mathbb{R}[E] \to \mathbb{R}[V]$ defined by the following action on every basis edge e. If e has source vertex u and target vertex v, then

$$I(e) = v - u.$$

Algebraic properties of $I : \mathbb{R}[E] \to \mathbb{R}[V]$ reflect geometric properties of the graph G — for instance, the number of connected components of G is $\#V - \text{rank}(I)$, and the number of undirected cycles in G is $\#E - \text{rank}(I)$.

With a simplicial complex, it becomes necessary to define an incidence matrix not only from 1-dimensional simplices to vertices as described above, but from k-dimensional simplices to $(k-1)$-dimensional simplices for *every* dimension k in $\{1, 2, \ldots, \dim K\}$. We therefore require a formula sending each k-simplex to a ± 1 linear combination of its codimension one faces.

DEFINITION 3.4. Let σ be a k-dimensional oriented simplex. The **algebraic boundary** of σ is the linear combination

$$\partial_k \sigma = \sum_{i=0}^{k} (-1)^i \sigma_{-i},$$

where σ_{-i} denotes the i-th face of σ as in Definition 3.3.

Since zero-dimensional simplices have no lower-dimensional faces, we require by convention that ∂_0 of every vertex is 0. The following figure illustrates algebraic boundaries for oriented simplices of dimension one and two:

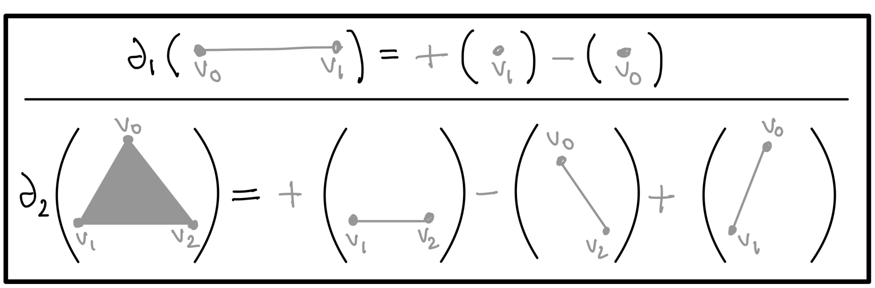

It will be convenient for the moment to not dwell too much on *where* these ± 1 coefficients are supposed live — we do not demand to know whether they are integers, rational numbers, real numbers, etc. For now we seek solace in the fact that ± 1 are defined in every unital ring (i.e., ring with a multiplicative identity). We will simply treat the algebraic boundary of each simplex

σ as a formal sum of its faces. The first purely algebraic miracle of this subject occurs when we try to compute boundaries of boundaries under the assumption that each ∂_k is a linear map.

PROPOSITION 3.5. *For any oriented simplex σ of dimension $k \geq 0$, we have*

$$\partial_{k-1} \circ \partial_k \sigma = 0.$$

PROOF. The proof of the full statement has been assigned as an exercise, but let's at least compute everything in the case $k = 2$. Consider an oriented 2-simplex $\sigma = (v_0, v_1, v_2)$ and note by Definition 3.4 that

$$\begin{aligned}\partial_2 \sigma &= (\sigma_{-0}) - (\sigma_{-1}) + (\sigma_{-2}) \\ &= (v_1, v_2) - (v_0, v_2) + (v_0, v_1).\end{aligned}$$

Assuming linearity of ∂_1, we have

$$\begin{aligned}\partial_1 \circ \partial_2 \sigma &= \partial_1(v_1, v_2) - \partial_1(v_0, v_2) + \partial_1(v_0, v_1) \\ &= [(v_2) - (v_1)] - [(v_2) - (v_0)] + [(v_1) - (v_0)].\end{aligned}$$

Now the desired conclusion follows by noticing that every vertex has appeared twice, but with opposite signs. □

Since ∂_0 is identically zero, the proposition above has non-trivial content only when $k \geq 2$; thus, this miraculous cancellation remains entirely un-witnessed in the realm of graphs and their incidence matrices.

3.3 CHAIN COMPLEXES

In order to take full advantage of Proposition 3.5, we must fix a *coefficient ring* to give precise meaning to the formal sums obtained when we take algebraic boundaries of simplices. The simplest choice, in terms of computation, is to work with a *field* $\mathbb{F}$ — typical choices include

- $\mathbb{F} = \mathbb{Q}$, the rational numbers;
- $\mathbb{F} = \mathbb{Z}/p$, integers modulo a prime number p, (often $p = 2$), and
- $\mathbb{F} = \mathbb{R}$, the real numbers.

The main advantage when using field coefficients is that we get to work with vector spaces and matrices, so all the standard machinery of linear algebra is at our disposal. With non-field coefficients (even the ring $\mathbb{Z}$ of integers), the algebraic objects at hand become considerably more intricate.

Let K be an oriented simplicial complex and $\mathbb{F}$ a field; both will remain fixed throughout this section.

DEFINITION 3.6. For each dimension $k \geq 0$, the k-th **chain group** of K is the vector space $\mathbf{C}_k(K)$ over $\mathbb{F}$ generated by treating the k-simplices of K as a basis.

Thus, every element γ in $\mathbf{C}_k(K)$ – which is called a k**-chain** of K – can be uniquely expressed as a linear combination of the form

$$\gamma = \sum_\sigma \gamma_\sigma \cdot \sigma,$$

where σ ranges over the k-simplices of K while the coefficients γ_σ are chosen from $\mathbb{F}$. Each k-simplex σ in K constitutes a basis vector in $\mathbf{C}_k(K)$, namely the chain γ whose coefficients are all zero except γ_σ, which equals the multiplicative identity $1 \in \mathbb{F}$. When k exceeds $\dim X$, there are no simplices to serve as basis elements, so $\mathbf{C}_k(X)$ is the trivial (i.e., zero-dimensional) vector space for all large k.

Having described chains as linear combinations of simplices, we are able to reinterpret the algebraic boundaries of Definition 3.4 as linear maps between consecutive chain groups.

DEFINITION 3.7. For each dimension $k \geq 0$, the k-th **boundary operator** of K is the $\mathbb{F}$-linear map $\partial_k^K : \mathbf{C}_k(K) \to \mathbf{C}_{k-1}(K)$ which sends each basis k-chain σ to the $(k-1)$-chain

$$\partial_k^K(\sigma) = \sum_{i=0}^{k} (-1)^i \cdot \sigma_{-i}.$$

In contrast to the formal sums of Definition 3.4, the $(-1)^i$ coefficients appearing in the boundary operator formula above do have a precise meaning — they are simply coefficients chosen from the field $\mathbb{F}$. Note that each σ_{-i} is a $(k-1)$-simplex of K, so the boundary $\partial_k^K(\sigma)$ of each k-simplex σ is automatically a $(k-1)$-chain as expected. To evaluate ∂_k^K on an arbitrary k-chain γ rather than a basis simplex, one simply exploits linearity:

$$\partial_k^K(\gamma) = \sum_{\sigma} \gamma_\sigma \cdot \partial_k^K(\sigma).$$

Here is an immediate consequence of Proposition 3.5

COROLLARY 3.8. *For every dimension $k \geq 0$, the composite*

$$\partial_k^K \circ \partial_{k+1}^K : \mathbf{C}_{k+1}(K) \to \mathbf{C}_{k-1}(K)$$

is the zero map. In other words, for each dimension k the image of ∂_{k+1}^K lies in the kernel of ∂_k^K.

Thus, we now have the ability to build (starting from an oriented simplicial complex K and a coefficient field $\mathbb{F}$) a sequence of finite-dimensional vector spaces connected by linear maps:

$$\cdots \xrightarrow{\partial_{k+1}^K} \mathbf{C}_k(K) \xrightarrow{\partial_k^K} \mathbf{C}_{k-1}(K) \xrightarrow{\partial_{k-1}^K} \cdots \xrightarrow{\partial_2^K} \mathbf{C}_1(K) \xrightarrow{\partial_1^K} \mathbf{C}_0(K) \longrightarrow 0$$

And moreover, this sequence has the magic property that whenever we compose two adjacent maps, the result is always zero. Such sequences play an enormous role in homology theory and beyond, so they have a special name.

DEFINITION 3.9. A **chain complex** $(C_\bullet, d_\bullet)$ over the field $\mathbb{F}$ is a collection of $\mathbb{F}$-vector spaces C_k (indexed by integers $k \geq 0$) and $\mathbb{F}$-linear maps $d_k : C_k \to C_{k-1}$ which satisfy the condition $d_k \circ d_{k+1} = 0$ for all k.

Chain complexes of the form $(\mathbf{C}_\bullet(K), \partial_\bullet^K)$ which arise from a simplicial complex K will be called *simplicial chain complexes* in order to distinguish them from the arbitrary chain complexes $(C_\bullet, d_\bullet)$ of Definition 3.9.

REMARK 3.10. Definition 3.9 works verbatim when we replace the field $\mathbb{F}$ with a commutative, unital ring. For many reasons, the most common choice of non-field coefficients is the ring of integers $\mathbb{Z}$. When working with $\mathbb{Z}$ coefficients, the chain groups C_k are abelian groups rather than vector spaces, and the boundary operators d_k are abelian group homomorphisms; the chain complex condition $d_k \circ d_{k+1} = 0$ makes sense in this context. But now one finds a stark difference between simplicial chain complexes and arbitrary ones: in a simplicial chain complex, each chain group C_k is always *free*, i.e., it has the form $\mathbb{Z}^n$ for some $n \geq 0$. On the other hand, arbitrary chain complexes over $\mathbb{Z}$ can have *torsion* in their chain groups, e.g., $C_k = \mathbb{Z} \oplus \mathbb{Z}/2$ is allowed. Torsion plays no role whatsoever in chain complexes (simplicial or otherwise) when we work with field coefficients.

3.4 Homology

Fix a chain complex $(C_\bullet, d_\bullet)$ of vector spaces and linear maps over some field $\mathbb{F}$:

$$\cdots \xrightarrow{d_{k+1}} C_k \xrightarrow{d_k} C_{k-1} \xrightarrow{d_{k-1}} \cdots \xrightarrow{d_2} C_1 \xrightarrow{d_1} C_0 \xrightarrow{d_0} 0$$

We recall from Definition 3.9 that $d_k \circ d_{k+1}$ is always the zero map from C_{k+1} to C_{k-1}, which means that the kernel $\ker d_k$ admits the image $\operatorname{img} d_{k+1}$ as a subspace for each $k \geq 0$.

DEFINITION 3.11. For each dimension $k \geq 0$, the k-th **homology group** of $(C_\bullet, d_\bullet)$ is defined to be the quotient vector space

$$\mathbf{H}_k(C_\bullet, d_\bullet) = {\ker d_k}/{\operatorname{img} d_{k+1}}$$

It is customary to refer to $\ker d_k \subset C_k$ as the subspace of k**-cycles** and to $\operatorname{img} d_{k+1}$ as the subspace of k**-boundaries**, so the mantra to chant is

Homology is cycles modulo boundaries.

The best way to become familiar with cycles and boundaries (and hence, with homology groups) is to try drawing them as subsets of simplicial complexes. One can avoid algebraic impediments when building geometric intuition by using $\mathbb{F} = \mathbb{Z}/2$ coefficients, so that it suffices to highlight which simplices have coefficient 1 in a given chain. Illustrated below are two 1-cycles γ and γ' in a triangulated annulus — the key point is that each vertex in sight must be a face of an even number of edges lying in γ (otherwise it will appear with a nonzero coefficient when we take the boundary $\partial_1\gamma$). The cycles γ and γ' represent the same element in the first homology group since they differ only by the boundary of the 2-simplex which has been shaded in the right panel.

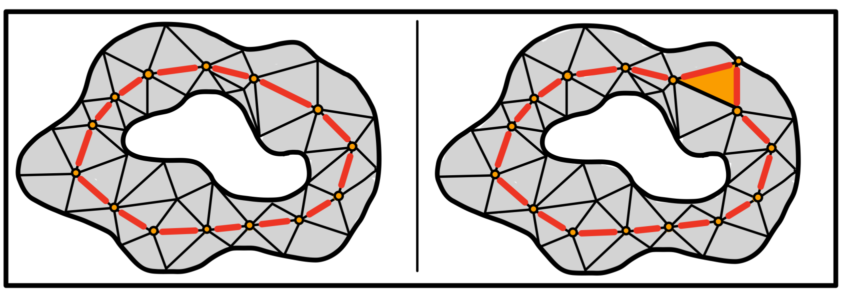

When $(C_\bullet, d_\bullet) = (\mathbf{C}_\bullet(K), \partial^K_\bullet)$ is the chain complex associated to a simplicial complex K, the associated homology groups are called the **simplicial homology groups** of K and denoted either $\mathbf{H}_k(K)$ or $\mathbf{H}_k(K; \mathbb{F})$ depending on how emphatically one is trying to showcase the choice of coefficient field. Simplicial homology groups are always finite-dimensional (since we require K to be finite), and for each $k \geq 0$ the dimension

$$\beta_k(K; \mathbb{F}) = \dim \mathbf{H}_k(K; \mathbb{F})$$

is called the k-th **Betti number** of K. As we have chosen to work with field coefficients, this single number completely determines $\mathbf{H}_k(K; \mathbb{F})$ up to isomorphism as a vector space, but it doesn't actually give us a basis of k-chains which generate $\mathbf{H}_k(K; \mathbb{F})$ as a vector space.

EXAMPLE 3.12. The Betti numbers of the solid 0-simplex over any field $\mathbb{F}$ are

$$\beta_k(\Delta_0) = \begin{cases} 1 & k = 0 \\ 0 & k > 0. \end{cases}$$

This can be seen by directly building the simplicial chain complex, which only admits a non-trivial chain group C_k for $k = 0$:

$$\cdots \to 0 \to 0 \to \mathbb{F} \to 0;$$

so all the homology groups are trivial except $\mathbf{H}_0$ which equals $\mathbb{F}$. On the other hand, the hollow 2-simplex has Betti numbers

$$\beta_k(\partial\Delta(2)) = \begin{cases} 1 & k \in \{0,1\} \\ 0 & k > 1. \end{cases}$$

This time the chain complex is non-trivial in dimensions 0 and 1:

$$\cdots \to 0 \to \mathbb{F}^3 \to \mathbb{F}^3 \to 0;$$

and the only non-trivial boundary map $d_1 : \mathbb{F}^3 \to \mathbb{F}^3$ has rank 2. Thus, its kernel has dimension 1 while its image has dimension 2, which means $\beta_0 = 3 - 2 = 1$ and $\beta_1 = 1 - 0 = 1$.

As suggested by these computations, the Betti numbers of K can be determined entirely by the ranks of boundary operators ∂_k^K which appear in the simplicial chain complex.

PROPOSITION 3.13. *Let K be a simplicial complex with K_k denoting the set of all k-dimensional simplices in K. Writing r_k for the rank of the boundary map $\partial_k^K : \mathbf{C}_k(K) \to \mathbf{C}_{k-1}(K)$ from Definition 3.7, for each dimension $k \geq 0$ we have*

$$\beta_k(K) = \#K_k - (r_k + r_{k+1})$$

PROOF. Since $\beta_k(K)$ is the dimension of the quotient $\ker d_k / \operatorname{img} d_{k+1}$, we have

$$\begin{aligned} \beta_k(K) &= \dim(\ker d_k) - \dim(\operatorname{img} d_{k+1}) \\ &= \dim(\ker d_k) - r_{k+1} \end{aligned}$$

By the **rank-nullity theorem** from linear algebra, we have $\dim(\ker d_k) = \dim \mathbf{C}_k(K) - r_k$, and from Definition 3.6 we know $\dim \mathbf{C}_k(K)$ is precisely the number of k-simplices in K. $\square$

In order to go beyond Betti numbers and actually extract *basis elements* for the vector spaces $\mathbf{H}_k(K;\mathbb{F})$, we require more serious linear algebra than the rank-nullity theorem. Before delving into algebraic manipulations, it will be helpful to have the following pictures firmly in mind. The first depicts a simplicial chain complex: for each dimension $k \geq 0$ we have a vector space generated by simplices of dimension k:

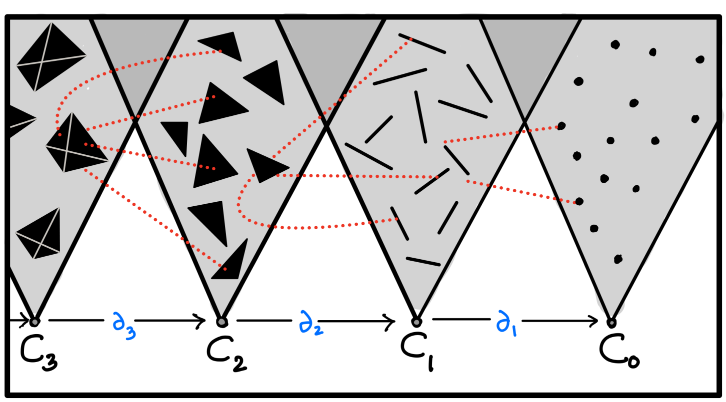

The dotted lines in this picture describe nonzero entries in matrix representations of the boundary operators: each basis 1-simplex has exactly two such nonzero entries under its column in ∂_1^K; these lie in the rows corresponding to the two 0-simplices which form its faces. Similarly, each basis 2-simplex has three nonzero entries in its ∂_2^K column, and so on. In this picture, the vector spaces $\mathbf{C}_k(K)$ have a very convenient and geometric description — basis elements are simplices σ, and these clearly form subspaces $|\sigma|$ of the geometric realization $|K|$. But the matrix representations of the boundary operators are a mess — there are lots of dotted lines flying around all over the place, linking each simplex to all of its faces. The key to computing the homology groups of K is to change the bases of all chain groups so that there is at most one incoming or one outgoing dotted line. This produces the following new picture of the same chain complex:

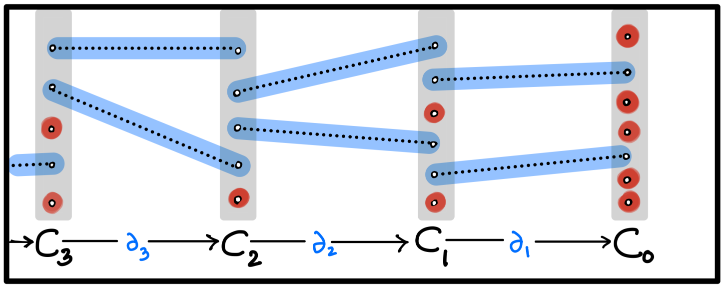

Now the basis elements which generate the chain groups do not have a convenient geometric description — they are weird linear combinations of simplices, and it is not clear how to attach a coherent geometric interpretation within $|K|$ to a linear combination $\gamma = 3\sigma - 5\tau$ of k-dimensional simplices. On the other hand, the boundary matrices have now become gloriously straightforward — there is a trichotomy for each basis element γ of $\mathbf{C}_k(K)$: either

(1) there is a single incoming dotted line to γ from a unique basis element of $\mathbf{C}_{k+1}(K)$, or
(2) there is a single outgoing dotted line from γ to a basis element of $\mathbf{C}_{k-1}(K)$, or
(3) γ remains entirely untouched by dotted lines.

In the first case, γ lies in the image of ∂^K_{k+1} while in the second case γ lies outside the kernel of ∂^K_k. The third case is the most interesting to us, since the basis elements which miss the dotted arrows entirely will simultaneously lie inside the kernel of ∂^K_k *and* outside the image $\operatorname{img} \partial^K_{k+1}$ — these are the desired generators of the homology group $\mathbf{H}_k(K)$.

In the next section, we will describe the algebraic operations which diagonalize the boundary matrices, thus turning the first chain complex picture into the second chain complex picture.

3.5 The Smith Decomposition

Fix a field $\mathbb{F}$ and consider a linear map $A : \mathbb{F}^m \to \mathbb{F}^n$ for some integers $m, n \geq 0$. We fix bases for the domain and codomain so that A has an explicit representation as an $n \times m$ matrix with each entry A_{ij} an element of $\mathbb{F}$.

THEOREM 3.14. **[Smith Normal Form]** *If $A : \mathbb{F}^m \to \mathbb{F}^n$ has rank r, then there exist invertible matrices $P : \mathbb{F}^n \to \mathbb{F}^n$ and $Q : \mathbb{F}^m \to \mathbb{F}^m$ satisfying the matrix equation*

$$D = PAQ,$$

where D is an $n \times m$ diagonal matrix whose entries are given by

$$D_{ij} = \begin{cases} 1 & \text{if } i = j \leq r, \\ 0 & \text{otherwise.} \end{cases}$$

(This matrix D is called the **Smith normal form** *of A.)*

It takes very little imagination (or computation) to produce the matrix D from A, since it only makes use of the numbers m, n and r — one simply creates an $n \times m$ matrix whose leading $r \times r$

minor is the identity matrix $\mathrm{Id}_{r\times r}$ and every other entry is zero. The information that we seek to extract from the *Smith decomposition* $D = PAQ$ is hidden in the invertible matrices P and Q: the first r columns of the inverse P^{-1} form a basis for $\operatorname{img} A \subset \mathbb{F}^n$ while the last $(m-r)$ columns of Q form a basis of $\ker A \subset \mathbb{F}^m$:

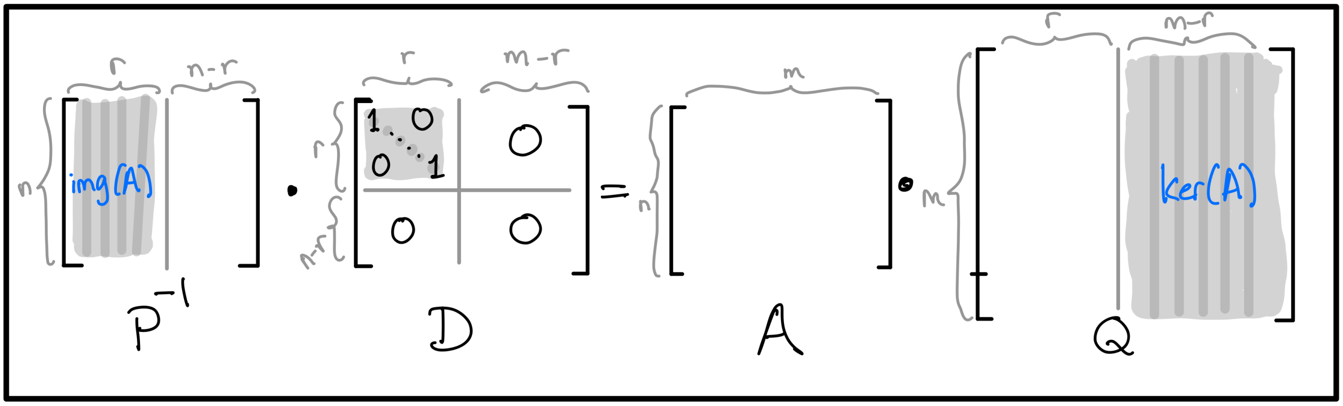

The good news is that computing these matrices P and Q from A requires nothing more sophisticated than the sorts of (hopefully familiar) row and column operations which one might use to put matrices in echelon form, namely: (1) add an $\mathbb{F}$-multiple of one row to another row, (2) scale a row by some nonzero element of $\mathbb{F}$, and (3) interchange two rows, plus the three corresponding operations for columns. You can compute everything at once by starting with the block matrix $\begin{bmatrix} \mathrm{Id}_{n\times n} & A \\ 0_{m\times n} & \mathrm{Id}_{m\times m}\end{bmatrix}$. As we perform row and column operations to diagonalize A, the identity matrices below and to the left will evolve accordingly, but the zero block remains unmolested. When A is fully diagonalized, we are conveniently left with $\begin{bmatrix} P & D \\ 0_{m\times n} & Q\end{bmatrix}$. Here are the first two moves of this computation for the matrix $A = \left[\begin{smallmatrix} 1 & 2 & 3 \\ 4 & 5 & 6\end{smallmatrix}\right]$, assuming that $\mathbb{F}$ is the field of rational numbers:

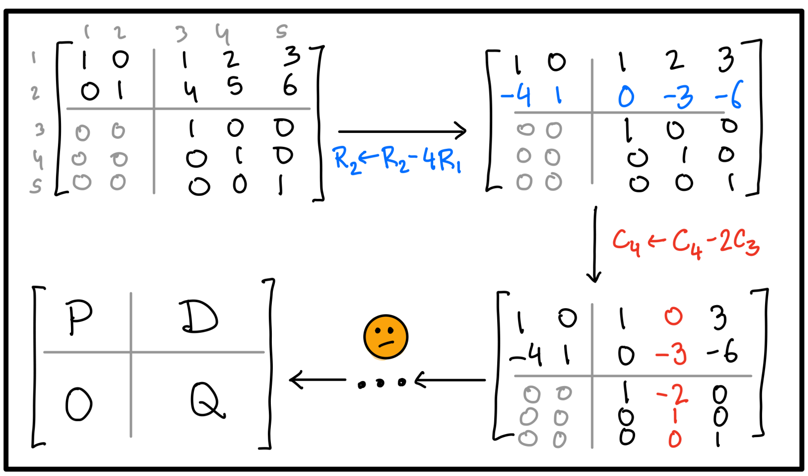

Let $(C_\bullet, d_\bullet)$ be a chain complex over $\mathbb{F}$ so that each C_k is finite-dimensional. Knowledge of the Smith decomposition of all the boundary operators d_k allows us to find basis vectors for the homology groups $\mathbf{H}_k(C_\bullet, d_\bullet)$ as follows.

PROPOSITION 3.15. *For each dimension $k \geq 0$, let $D_k = P_k d_k Q_k$ be the Smith decomposition of the boundary operator $d_k : C_k \to C_{k-1}$. For $n_k = \dim C_k$ and $r_k = \operatorname{rank} d_k$, let G_k be the matrix of size $(r_{k+1} + n_k - r_k) \times n_k$ given by the block decomposition*

$$G_k = [B_k \mid Z_k],$$

where B_k equals the first r_{k+1} columns of P_{k+1}^{-1} while Z_k equals the last $(n_k - r_k)$ columns of Q_k. If $E_k = [B'_k \mid Z'_k]$ is the reduced row echelon form of G_k obtained by performing row (but not column) operations over $\mathbb{F}$, then the columns of Z_k corresponding to the pivot columns of Z'_k form a basis for $\mathbf{H}_k(C_\bullet, d_\bullet)$.

PROOF. As discussed above, for each $k \geq 0$ the matrix B_k contains a basis for $\operatorname{img} d_{k+1}$, and so every column in the left block B'_k of E_k is guaranteed to have a pivot. Recall that $\operatorname{img} d_{k+1}$ is a subspace of $\ker d_k$ by Definition 3.9, and that the right block Z_k contains a basis for $\ker d_k$. Thus, there will be exactly $(n_k - r_k - r_{k+1})$ pivot columns in Z'_k once we are in row echelon form, and the corresponding columns of Z_k provide a basis for the quotient $\ker d_k / \operatorname{img} d_{k+1} = \mathbf{H}_k(C_\bullet, d_\bullet)$. $\square$

There is nothing unique about bases obtained via the procedure above — for one thing, we can always add vectors from $\operatorname{img} d_{k+1}$ to a basis vector to get a new basis vector. More seriously, we could take interesting linear combinations, e.g., replace basis vectors $\{\gamma_1, \gamma_2\}$ by $\{\gamma_1 + \gamma_2, \gamma_1 - \gamma_2\}$.

REMARK 3.16. Theorem 3.14 also holds when A is an integer matrix $\mathbb{Z}^m \to \mathbb{Z}^n$; in this case, both P and Q will be invertible integer matrices. But here the Smith normal form D has a more interesting structure; instead of a leading $r \times r$ identity matrix, we have a diagonal matrix populated by positive integers $a_1 \leq \cdots \leq a_r$, called the **invariant factors** of A. Each a_i divides the subsequent a_{i+1}, and whenever $a_i \neq 1$ we obtain torsion of the form $\mathbb{Z}/a_i$ in the integral homology groups.

We have not yet explained what homology groups have to do with the fact that the Euler characteristic is homotopy invariant. But we are able to take the first steps in the right direction: one of the exercises below asks you to confirm that $\chi(K)$ for a simplicial complex K can be completely recovered from its rational Betti numbers $\{\beta_k(K; \mathbb{Q}) \mid k \geq 0\}$. In the next Chapter, we will develop enough machinery to see that homology groups, and hence Betti numbers, are themseleves homotopy invariant.

EXERCISES

EXERCISE 3.1. What is the Euler characteristic $\chi(\partial\Delta(k))$ of the hollow k-simplex as a function of the dimension $k \geq 1$?

EXERCISE 3.2. For any sequence $a = (a_0, a_1, \ldots)$ of real numbers and simplicial complex K, define the real number $\chi_a(K)$ by

$$\chi_a(K) = \sum_{i=0}^{\dim K} a_i \cdot \# K_i.$$

Show that if $\chi_a(\Delta(k))$ is constant for all $k \geq 0$ then there exists some real number λ satisfying $a_i = \lambda \cdot (-1)^i$, so χ_a must be a scalar multiple of the Euler characteristic.

EXERCISE 3.3. Show that the Euler characteristic of a simplicial complex equals the alternating sum of its Betti numbers over $\mathbb{Q}$, i.e.,

$$\chi(K) = \sum_{i=0}^{\dim K} \beta_i(K; \mathbb{Q}).$$

[Hint: use Proposition 3.13]

EXERCISE 3.4. Prove Proposition 3.5 by extending the argument for $k = 2$ to arbitrary k. [Hint: evaluate $\partial_{k-1} \circ \partial_k \sigma$ as follows:

$$\partial_{k-1}\left(\sum_{i=0}^{k}(-1)^i\sigma_{-i}\right) = \sum_{i=0}^{k}(-1)^i \cdot \partial_{k-1}\sigma_{-i} = \sum_{i=0}^{k}(-1)^i \cdot \sum_{j=0}^{k-1}(-1)^j(\sigma_{-i})_{-j}.$$

Now decompose the double-sum into the parts where $j > i$ and $j < i$.]

EXERCISE 3.5. Let K be a one-dimensional oriented simplicial complex (i.e., a directed graph). Describe the simplicial chain complex of K in terms of its incidence matrix I.

EXERCISE 3.6. For K a one-dimensional oriented simplicial complex and coefficients in $\mathbb{F} = \mathbb{Z}/2$, show that $\ker \partial_1^K$ must consist entirely of cycles, i.e., paths which start and end at the same vertex.

EXERCISE 3.7. Let K be a simplicial complex and L a subcomplex so that $K - L$ only contains simplices of dimension k or above. Prove that $\beta_i(K) = \beta_i(L)$ for all $i < k - 1$.

EXERCISE 3.8. Write down the simplicial chain complex for the solid simplex $\Delta(2)$. Determine the ranks of the boundary operators and hence the Betti numbers of $\Delta(2)$.

EXERCISE 3.9. Compute the Smith decomposition $D = PAQ$ for the matrix $A = \left[\begin{smallmatrix} 1 & 2 & 1 \\ -2 & -4 & -2 \end{smallmatrix}\right]$ over $\mathbb{Q}$; use this to write down basis vectors for its kernel and image.

EXERCISE 3.10. Find bases for $\mathbf{H}_0$ and $\mathbf{H}_1$ of the hollow 2-simplex $\partial\Delta(2)$ with $\mathbb{F} = \mathbb{Z}/2$.

EXERCISE 3.11. A simplicial complex is *connected* if any two vertices u and v can be joined by a path of consecutive edges, i.e.,

$$u \leftrightarrow w_0 \leftrightarrow w_1 \leftrightarrow \cdots \leftrightarrow w_k \leftrightarrow v$$

Show that $\beta_0(K; \mathbb{Q}) = 1$ whenever K is connected.

4. SEQUENCES

4.1 Categories

When first meeting a new class of mathematical objects – such as topological spaces or abelian groups – it is natural to try and learn about the underlying structures (open sets, commutative multiplication laws) which characterize each object, and to carefully describe a reasonable notion of functions (continuous maps, group homomorphisms) which preserve that structure. In almost all cases of interest, it turns out that composing two such structure-preserving functions produces another such function. The following definition provides a convenient umbrella under which such (structure, function) pairs reside.

DEFINITION 4.1. A **category** $\mathscr{C}$ consists of

(1) a collection $\mathscr{C}_0$ whose elements are called *objects*,
(2) for every pair of objects x, y in $\mathscr{C}$ a set $\mathscr{C}(x,y)$ of *morphisms* from x to y, whose elements we denote $f, g, \ldots : x \to y$, and
(3) for each triple x, y, z of objects a *composition law* $\mathscr{C}(x,y) \times \mathscr{C}(y,z) \to \mathscr{C}(x,z)$ sending $f : x \to y$ and $g : y \to z$ to some $g \circ f : x \to z$,

subject to the **identity** and **associativity** axioms.

Our definition remains incomplete until we spell out these two axioms; here they are:

(1) for each x in $\mathscr{C}_0$, there is a distinguished *identity* morphism $1_x : x \to x$ satisfying both
$$g \circ 1_x = g \text{ and } 1_x \circ h = h$$
for any object y and morphisms $g : y \to x$ and $h : x \to y$;
(2) given any triple of morphisms of the form
$$x \xrightarrow{f} y \xrightarrow{g} z \xrightarrow{h} w,$$
the *associativity* condition $h \circ (g \circ f) = (h \circ g) \circ f$ holds.

Instances of (object, morphism) pairs in mathematics which satisfy these two axioms are ubiquitous — consider, for instance:

- the category **Set** of (sets, functions),
- the category **Grp** of (groups, group homomorphisms),
- its subcategory **AbGrp** of (abelian groups, abelian group homomorphisms),
- the category **SC** of (simplicial complexes, simplicial maps),
- the category $\mathbf{Vect}_{\mathbb{F}}$ of ($\mathbb{F}$-vector spaces, $\mathbb{F}$-linear maps) over a field $\mathbb{F}$, etc.

One can encode the associativity axiom in the form of a *commuting square*, like so:

$$\begin{tikzcd}
\mathscr{C}(x,y) \times \mathscr{C}(y,z) \times \mathscr{C}(z,w) \arrow[r, "(f,g,h)\mapsto(g\circ f,h)"] \arrow[d, "(f,g,h)\mapsto(f,h\circ g)"'] & \mathscr{C}(x,z) \times \mathscr{C}(z,w) \arrow[d, "(g\circ f,h)\mapsto h\circ(g\circ f)"] \\
\mathscr{C}(x,y) \times \mathscr{C}(y,w) \arrow[r, "(f,h\circ g)\mapsto(h\circ g)\circ f"'] & \mathscr{C}(x,w)
\end{tikzcd}$$

At first glance, this diagrammatic translation of $(h \circ g) \circ f = h \circ (g \circ f)$ might come across as an elaborate crime against brevity. There are, however, several compelling reasons to become familiar with the language of commuting diagrams — for one thing, there are many such diagrams in our immediate future. Another special feature of the categorical philosophy, besides this profusion of commuting diagrams, is that it can be turned inwards to reason about the theory of categories itself. Those under its influence naturally ask what key piece of structure must be preserved by functions which map one category $\mathscr{C}$ to another category $\mathscr{C}'$.

DEFINITION 4.2. A **functor** $\boldsymbol{F} : \mathscr{C} \to \mathscr{C}'$ assigns

(1) to each object x in $\mathscr{C}_0$ an object $\boldsymbol{F}x$ in $\mathscr{C}'_0$, and
(2) to each morphism $f : x \to y$ in $\mathscr{C}$ a morphism $\boldsymbol{F}f : \boldsymbol{F}x \to \boldsymbol{F}y$ in $\mathscr{C}'$,

subject to the following conditions:

(1) we have $\boldsymbol{F}1_x = 1_{\boldsymbol{F}x}$ for each x in $\mathscr{C}_0$, and
(2) for any pair of morphisms f in $\mathscr{C}(x, y)$ and g in $\mathscr{C}(y, z)$, we have
$$\boldsymbol{F}(g \circ f) = \boldsymbol{F}g \circ \boldsymbol{F}f$$
(Here the composition on the left takes place in $\mathscr{C}$ while the composition on the right takes place in $\mathscr{C}'$).

Thus, a functor $\mathscr{C} \to \mathscr{C}'$ sends $\mathscr{C}$-objects to $\mathscr{C}'$-objects and the corresponding $\mathscr{C}$-morphisms to $\mathscr{C}'$-morphisms in a manner that duly respects composition laws of both $\mathscr{C}$ and $\mathscr{C}'$. One of the exercises to this Chapter asks you to define the composite of two functors and hence construct the category **Cat** containing (categories, functors). We have been discussing categories and functors because of the next result, which catalogues one of the most important properties of simplicial homology (see Section 4).

THEOREM 4.3. *For each dimension $k \geq 0$, the assignment*
$$K \mapsto \mathbf{H}_k(K; \mathbb{F})$$
constitutes a functor from the category **SC** *of simplicial complexes and maps to the category* $\mathbf{Vect}_{\mathbb{F}}$ *of vector spaces over* $\mathbb{F}$.

We already know from Chapter 3 that every simplicial complex K can be assigned a vector space $\mathbf{H}_k(K; \mathbb{F})$ by first building the simplicial chain complex $(\mathbf{C}_\bullet(K), \partial^K)$ and then extracting the relevant quotient $\ker \partial_k^K / \operatorname{img} \partial_{k+1}^K$. So the new content of Theorem 4.3 lies entirely on the level of morphisms — we must first show that every simplicial map $f : K \to L$ induces a well-defined linear map $\mathbf{H}_k f : \mathbf{H}_k(K; \mathbb{F}) \to \mathbf{H}_k(L; \mathbb{F})$ of homology groups; and next, we have to confirm that given some other simplicial map $g : L \to M$, we have an equality
$$\mathbf{H}_k(g \circ f) = \mathbf{H}_k g \circ \mathbf{H}_k f$$
of linear maps $\mathbf{H}_k(K; \mathbb{F}) \to \mathbf{H}_k(M; \mathbb{F})$. These are our goals in the next two Sections.

4.2 CHAIN MAPS

Fix simplicial complexes K and L as well we as a simplicial map $f : K \to L$ and a coefficient field $\mathbb{F}$. We will continue to write $(\mathbf{C}_\bullet(K), \partial_\bullet^K)$ to indicate the simplicial chain complex of K (and similarly for L).

DEFINITION 4.4. For each dimension $k \geq 0$, let $\mathbf{C}_k f : \mathbf{C}_k(K) \to \mathbf{C}_k(L)$ be the $\mathbb{F}$-linear map between chain groups defined by the following action on each basis k-simplex σ of K:
$$\mathbf{C}_k f(\sigma) = \begin{cases} f(\sigma) & \text{if } \dim f(\sigma) = k \\ 0 & \text{otherwise.} \end{cases}$$

Perhaps the most interesting aspect of this definition is its piecewise nature — some simplices σ are faithfully mapped onto their images $f(\sigma)$ while others are sent to zero, depending on whether f is injective on their vertices or not. This is a necessary bit of book-keeping: we want to produce a map from k-chains to k-chains for each k, and the image $f(\sigma)$ of a k-simplex σ will not be a basis element of $\mathbf{C}_k(L)$ unless $\dim f(\sigma) = k$. Our next order of business is to see how the linear maps $\mathbf{C}_\bullet f$ interact with the two boundary operators $\partial_\bullet^K$ and $\partial_\bullet^L$. It turns out that the diagram below is a commuting square (in the category $\mathbf{Vect}_\mathbb{F}$) for each $k \geq 0$:

$$\begin{tikzcd}
\mathbf{C}_k(K) \arrow[r, "\partial_k^K"] \arrow[d, "\mathbf{C}_k f"'] & \mathbf{C}_{k-1}(K) \arrow[d, "\mathbf{C}_{k-1} f"] \\
\mathbf{C}_k(L) \arrow[r, "\partial_k^L"'] & \mathbf{C}_{k-1}(L)
\end{tikzcd}$$

PROPOSITION 4.5. *For each dimension $k \geq 0$, and k-simplex σ in K, we have an equality*

$$\partial_k^L \circ \mathbf{C}_k f(\sigma) = \mathbf{C}_{k-1} f \circ \partial_k^K(\sigma).$$

PROOF. Given Definition 4.4, the argument naturally decomposes into two cases.

Case 1: $\dim f(\sigma) = k$**.** Here $\mathbf{C}_k f(\sigma) = f(\sigma)$ and we have

$$\begin{aligned}
\partial_k^L \circ \mathbf{C}_k f(\sigma) &= \partial_k^L f(\sigma) && \text{since } \mathbf{C}_k f(\sigma) = f(\sigma), \\
&= \sum_{i=0}^{k} (-1)^i \cdot f(\sigma)_{-i} && \text{by Definition 3.7.}
\end{aligned}$$

On the other hand, $\partial_k^K(\sigma)$ equals $\sum_{i=0}^k (-1)^i \sigma_{-i}$, and since f is injective on the vertices of σ it must also be injective on the vertices of each face σ_{-i}. Thus, $\mathbf{C}_{k-1} f(\sigma_{-i}) = f(\sigma_{-i})$ for each i, and we have

$$\begin{aligned}
\mathbf{C}_{k-1} f \circ \partial_k^K(\sigma) &= \mathbf{C}_{k-1} f \left(\sum_{i=0}^{k} (-1)^i \cdot \sigma_{-i} \right) && \text{by Definition 3.7,} \\
&= \sum_{i=0}^{k} (-1)^i f(\sigma)_{-i} && \text{by Definition 4.4.}
\end{aligned}$$

Thus, $\partial_k^L \circ \mathbf{C}_k f(\sigma)$ and $\mathbf{C}_{k-1} f \circ \partial_k^K(\sigma)$ coincide in this case.

Case 2: $\dim f(\sigma) < k$**.** Here $\mathbf{C}_k f(\sigma)$ equals zero by definition, and hence so does its boundary in L. It suffices to show that the other composite $\mathbf{C}_{k-1} f \circ \partial_k^K(\sigma)$ is also zero. To this end, impose orientations o_K and o_L on K and L so that f is orientation-preserving, i.e, $o_K(v) < o_K(v')$ forces $o_L(f(v)) < o_L(f(v'))$. Writing σ as an o_K-oriented simplex $(v_0, \ldots, v_k)$, we must have $f(v_p) = f(v_{p+1})$ for some p in $\{0, \ldots, k-1\}$. Thus, f fails to be injective on the vertices of every face σ_{-i} of σ, with the possible exceptions of σ_{-p} and $\sigma_{-(p+1)}$. Now,

$$\begin{aligned}
\mathbf{C}_{k-1} f \circ \partial_k^K(\sigma) &= \mathbf{C}_{k-1} f \left(\sum_{i=0}^{k-1} (-1)^i \cdot \sigma_{-i} \right) && \text{by Definition 3.7} \\
&= \sum_{i=0}^{k-1} (-1)^i \mathbf{C}_{k-1} f(\sigma_{-i}) && \text{by linearity of } \mathbf{C}_k f \\
&= (-1)^p \left[f(\sigma_{-p}) - f(\sigma_{-(p+1)}) \right] && \text{by Definition 4.4.}
\end{aligned}$$

Here the last step follows from our observation that $\mathbf{C}_k f$ evaluates to zero on all other other faces of $f(\sigma)$ since f will not be injective on their vertices. But now, $f(\sigma_{-p})$ and $f(\sigma_{-p+1})$ are the same simplex in L, since both have vertices $(f(v_0),\ldots,f(v_p),f(v_{p+2}),\ldots,f(v_k))$. Thus, both $\partial_k^L \circ \mathbf{C}_k f(\sigma)$ and $\mathbf{C}_{k-1} f \circ \partial_k^K(\sigma)$ equal zero in this case. $\square$

As a consequence of this result, we are able to use the simplicial map $f : K \to L$ to produce a sequence of linear maps $\mathbf{C}_\bullet f : \mathbf{C}_\bullet K \to \mathbf{C}_\bullet L$ between the chain groups which form a ladder-shaped commuting diagram:

$$\begin{tikzcd}
\cdots \arrow[r, "\partial_{k+1}^K"] & \mathbf{C}_k(K) \arrow[r, "\partial_k^K"] \arrow[d, "\mathbf{C}_k f"] & \mathbf{C}_{k-1}(K) \arrow[r, "\partial_{k-1}^K"] \arrow[d, "\mathbf{C}_{k-1} f"] & \cdots \arrow[r, "\partial_2^K"] & \mathbf{C}_1(K) \arrow[r, "\partial_1^K"] \arrow[d, "\mathbf{C}_1 f"] & \mathbf{C}_0(K) \arrow[r, "0"] \arrow[d, "\mathbf{C}_0 f"] & 0 \arrow[d, equal] \\
\cdots \arrow[r, "\partial_{k+1}^L"'] & \mathbf{C}_k(L) \arrow[r, "\partial_k^L"'] & \mathbf{C}_{k-1}(L) \arrow[r, "\partial_{k-1}^L"'] & \cdots \arrow[r, "\partial_2^L"'] & \mathbf{C}_1(L) \arrow[r, "\partial_1^L"'] & \mathbf{C}_0(L) \arrow[r, "0"'] & 0
\end{tikzcd}$$

This is a standard example of a **chain map**, which can be used to relate arbitrary (i.e., not necessarily simplicial) chain complexes.

DEFINITION 4.6. A **chain map** $\phi_\bullet$ from $(C_\bullet, d_\bullet)$ to $(C'_\bullet, d'_\bullet)$ is defined to be a sequence of $\mathbb{F}$-linear maps $\{\phi_k : C_k \to C'_k \mid k \geq 0\}$ which satisfy

$$d'_k \circ \phi_k = \phi_{k-1} \circ d_k$$

for each $k \geq 0$.

Proposition 4.5 can be now be rephrased:

$$\text{simplicial maps } f : K \to L \text{ induce chain maps } \mathbf{C}_\bullet f : (\mathbf{C}_\bullet(K), \partial_\bullet^K) \to (\mathbf{C}_\bullet(L), \partial_\bullet^L).$$

It turns out that chain maps form the correct notion of morphisms in the category of chain complexes; their composition is not too difficult to define, and will be addressed by Exercise 4.3.

4.3 FUNCTORIALITY

To continue our proof of Theorem 4.3, we will use chain maps to construct maps of homology groups.

PROPOSITION 4.7. *Let $\phi_\bullet : (C_\bullet, d_\bullet) \to (C'_\bullet, d'_\bullet)$ be a chain map. For each dimension $k \geq 0$, there is a well-defined $\mathbb{F}$-linear map $\mathbf{H}_k\phi : \mathbf{H}_k(C_\bullet, d_\bullet) \to \mathbf{H}_k(C'_\bullet, d'_\bullet)$ induced by $\phi_\bullet$.*

PROOF. To induce a map of quotient vector spaces $\ker d_k / \operatorname{img} d_{k+1} \to \ker d'_k / \operatorname{img} d'_{k+1}$, it suffices to show that ϕ_k maps $\ker d_k$ to $\ker d'_k$ and $\operatorname{img} d_{k+1}$ to $\operatorname{img} d'_{k+1}$. First consider $\gamma \in C_k$ satisfying $d_k(\gamma) = 0$. Using Definition 4.6, we get

$$d'_k \circ \phi_k(\gamma) = \phi_{k-1} \circ d_k(\gamma) = 0,$$

so $\phi_k(\gamma)$ lies in $\ker d'_k$ as desired. Next, if $\alpha \in C_k$ lies in $\operatorname{img} d_{k+1}$, then we have $\alpha = d_{k+1}(\beta)$ for some β in C_{k+1}. Once again, Definition 4.6 gives us,

$$\phi_k(\alpha) = \phi_k \circ d_{k+1}(\beta) = d'_{k+1} \circ \phi_{k+1}(\zeta),$$

whence $\phi_k(\alpha)$ lies in $\operatorname{img} d_{k+1}$. Thus, for each γ in $\ker d_k$, our map $\mathbf{H}_k\phi$ sends $\gamma + \operatorname{img} d_{k+1}$ to $\phi(\gamma) + \operatorname{img} d'_{k+1}$, with a guarantee that $\phi(\gamma)$ lies in $\ker d'_k$. $\square$

Combining Proposition 4.5 with Proposition 4.7, we see that every simplicial map $f : K \to L$ indeed produces a well-defined linear map $\mathbf{H}_k(K;\mathbb{F}) \to \mathbf{H}_k(L;\mathbb{F})$ for each dimension $k \geq 0$. In order to avoid writing the monstrosity $\mathbf{H}_k\mathbf{C}f$ every time we want to mention this *induced map*, we will abbreviate it to $\mathbf{H}_k f : \mathbf{H}_k(K;\mathbb{F}) \to \mathbf{H}_k(L;\mathbb{F})$. It is not difficult to confirm that when f is the identity simplicial map $K \to K$, its induced map is the identity on $\mathbf{H}_k(K;\mathbb{F})$. The following result (which forms one of the exercises to this Chapter) takes a bit more work.

PROPOSITION 4.8. *Given chain maps $\phi_\bullet : (C_\bullet, d_\bullet) \to (C'_\bullet, d'_\bullet)$ and $\psi : (C'_\bullet, d'_\bullet) \to (C''_\bullet, d''_\bullet)$, we have*

$$\mathbf{H}_k(\psi \circ \phi) = \mathbf{H}_k\psi \circ \mathbf{H}_k\phi$$

for each dimension $k \geq 0$

Applying the above result to the special case where our chain maps are induced by simplicial maps (i.e., $\phi_\bullet = \mathbf{C}_\bullet f$ and $\psi_\bullet = \mathbf{C}_\bullet g$ for some $f : K \to L$ and $g : L \to M$) completes the proof of Theorem 4.3.

We now have the ability to study not just the homology groups of simplicial complexes but also linear maps of homology groups induced by simplicial maps (and more generally, chain maps); we will now examine various salient properties of such maps. A chain map $\phi_\bullet$ is called an isomorphism if each $\phi_k : C_k \to C'_k$ is an isomorphism from C_k to C'_k — for such a ϕ the induced maps $\mathbf{H}_\bullet\phi$ are also isomorphisms. But in general $\mathbf{H}_\bullet\phi$ can be an isomorphism even if $\phi_\bullet$ is not.

DEFINITION 4.9. A chain map $\phi_\bullet : (C_\bullet, d_\bullet) \to (C'_\bullet, d'_\bullet)$ is called a **quasi-isomorphism** if the induced map $\mathbf{H}_k\phi : \mathbf{H}_k(C_\bullet, d_\bullet) \to \mathbf{H}_k(C'_\bullet, d'_\bullet)$ is an isomorphism for each dimension $k \geq 0$.

In sharp contrast to testing whether a simplicial map induces homotopy equivalence or not, testing whether it induces a quasi-isomorphic chain map (and hence, isomorphisms of homology groups) is algorithmic and machine-computable.

REMARK 4.10. Consider a chain map $\phi : (C_\bullet, d_\bullet) \to (C'_\bullet, d'_\bullet)$; if the dimensions of C_k and C'_k are finite, then the computation of $\mathbf{H}_k\phi$ can be accomplished via the following linear algebraic procedure:

(1) Extract basis vectors B and B' for $\mathbf{H}_k(C_\bullet, d_\bullet)$ and $\mathbf{H}_k(C'_\bullet, d'_\bullet)$ via Proposition 3.15.
(2) For each basis chain b in B, write $\phi_k(b)$ as a linear combination of the basis chains of B':
$$\phi_k(b) = \sum_{b'} \alpha_{b,b'} \cdot b',$$
where each $\alpha_{b,b'}$ lies in the coefficient field $\mathbb{F}$. These α coefficients can be determined for all b at once by row-reducing the augmented matrix $[B'_k \mid \phi_k(B)]$.
(3) The coefficients $\{\alpha_{b,b'} \mid b \in B \text{ and } b' \in B'\}$ form a matrix $\mathbf{H}_k(C_\bullet, d_\bullet) \to \mathbf{H}_k(C'_\bullet, d'_\bullet)$; this matrix represents our linear map $\mathbf{H}_k\phi$ in terms of the bases B and B'.

Computability issues aside, induced maps on homology can be somewhat subtle.

EXAMPLE 4.11. The figure below illustrates two simplicial maps f, g from the hollow 2-simplex $\partial\Delta(2)$ to another simplicial complex K. The homology groups of K and $\partial\Delta(2)$ are isomorphic as rational vector spaces, i.e.,

$$\mathbf{H}_k(\partial\Delta(2);\mathbb{Q}) \simeq \mathbf{H}_k(K;\mathbb{Q}) = \begin{cases} \mathbb{Q} & \text{if } k \in \{0,1\}, \\ 0 & \text{otherwise .} \end{cases}$$

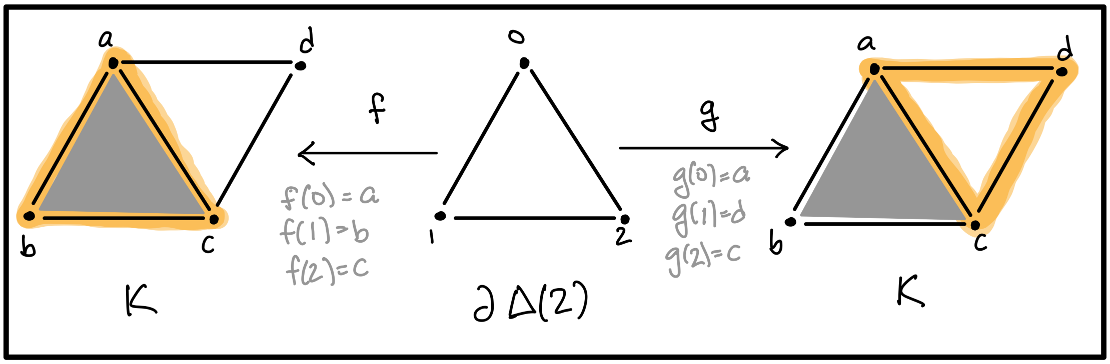

The chain map $\mathbf{C}_\bullet g$ is a quasi-isomorphism whereas $\mathbf{C}_\bullet f$ is not.

4.4 Chain Homotopy

There is a purely algebraic version of homotopy equivalence designed to work directly with chain complexes (rather than topological spaces). As usual, the first step is to define an equivalence relation between the set of all chain maps between a fixed pair of chain complexes.

Definition 4.12. A **chain homotopy** $\eta_\bullet$ between chain maps $\phi_\bullet, \psi_\bullet : (C_\bullet, d_\bullet) \to (C'_\bullet, d'_\bullet)$ is a collection of $\mathbb{F}$-linear maps $\eta_k : C_k \to C'_{k+1}$ which satisfy

$$\phi_k - \psi_k = \eta_{k-1} \circ d_k + d'_{k+1} \circ \eta_k$$

for each $k \geq 0$.

We write $\eta_\bullet : \phi_\bullet \Rightarrow \psi_\bullet$ to indicate that $\eta_\bullet$ is a chain homotopy as defined above; and the maps $\phi_\bullet$ and $\psi_\bullet$ are said to be **chain homotopic** whenever such an $\eta_\bullet$ exists. Chain homotopy is an equivalence relation on the set of all chain maps between a fixed pair of chain complexes.

Remark 4.13. It is important to note that the linear maps η_k are *not* required to satisfy any relations beyond the ones in the preceding defintion — in particular, they do not have to commute with d, d', ϕ or ψ. Even so, it is good to see how they fit within the commuting staircase diagrams that contain ϕ and ψ:

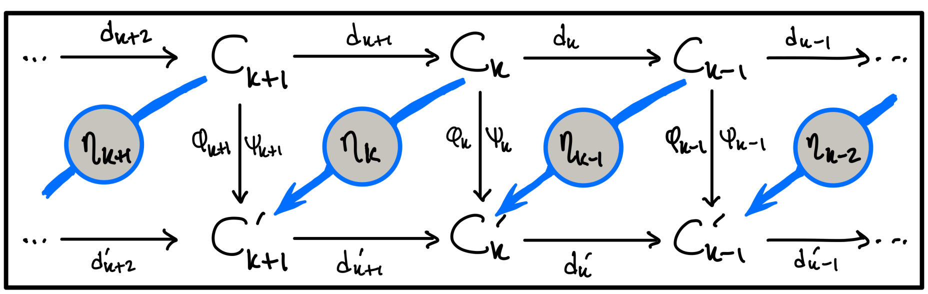

The following result highlights the utility of chain homotopy.

Proposition 4.14. *If $\phi_\bullet, \psi_\bullet : (C_\bullet, d_\bullet) \to (C'_\bullet, d'_\bullet)$ are chain homotopic, then their induced maps on homology coincide, i.e.,*

$$\mathbf{H}_k\phi = \mathbf{H}_k\psi$$

for each dimension $k \geq 0$.

Proof. Let $\eta_\bullet : \phi_\bullet \Rightarrow \psi_\bullet$ be a chain homotopy. For any chain $\gamma \in \ker d_k$, Definition 4.12 gives us

$$\phi_k(\gamma) - \psi_k(\gamma) = \eta_{k-1} \circ d_k(\gamma) + d'_{k+1} \circ \eta_k(\gamma).$$

But $d_k(\gamma) = 0$, so the first term on the right side disappears and the difference $\phi_k(\gamma) - \psi_k(\gamma)$ equals $d'_{k+1} \circ \eta_k(\gamma)$, which evidently lies in img d'_{k+1}. Thus, this difference is always a k-boundary, which is undetectable by homology. ☐

Chain homotopies are to chain maps what homotopies (as in Definition 2.1) are to continuous maps: they provide an indirect method for establishing that two chain complexes $(C_\bullet, d_\bullet)$ and $(C'_\bullet, d'_\bullet)$ are related by a quasi-isomorphism. The good news is that this method largely circumvents the tedious algebraic manipulations of Remark 4.10 and Proposition 3.15. But the bad news is that in order to avail of this method, we require not only a backwards chain map $\psi_\bullet : (C'_\bullet, d'_\bullet) \to (C_\bullet, d_\bullet)$ but also a pair of chain homotopies, described below.

DEFINITION 4.15. A pair of chain complexes is said to be **chain homotopy equivalent** if there are two chain maps

$$\phi_\bullet : (C_\bullet, d_\bullet) \to (C'_\bullet, d'_\bullet) \text{ and } \psi_\bullet : (C'_\bullet, d'_\bullet) \to (C_\bullet, d_\bullet)$$

along with chain homotopies

$$\eta_\bullet : 1_{(C_\bullet, d_\bullet)} \Rightarrow \psi_\bullet \circ \phi_\bullet \text{ and } \eta'_\bullet : \phi_\bullet \circ \psi_\bullet \Rightarrow 1_{(C'_\bullet, d'_\bullet)}.$$

Here $1_{(C_\bullet, d_\bullet)}$ is the identity chain map of $(C_\bullet, d_\bullet)$, etc.

It follows immediately from Proposition 4.14 that if two chain complexes are chain homotopy equivalent, then they must be related by quasi-isomorphisms and hence have isomorphic homology groups.

EXAMPLE 4.16. The cone Cone(K) over any simplicial complex K (see Definition 1.19) has homology groups isomorphic to those of $\Delta(0)$, namely:

$$\mathbf{H}_k(\text{Cone}(K); \mathbb{F}) = \begin{cases} \mathbb{F} & k = 0 \\ 0 & k > 0. \end{cases}$$

To see this, let $f : \text{Cone}(K) \to \Delta(0)$ be the simplicial map sending every vertex of the cone to the unique vertex 0, and let $g : \Delta(0) \to \text{Cone}(K)$ be the simplicial map sending 0 to the special vertex v_* which lies in Cone$(K) - K$. Now the composite $f \circ g$ equals the identity on $\Delta(0)$, and the other composite $g \circ f$ sends every vertex of Cone(K) to v_*. It remains to construct a chain homotopy from the identity chain map on $\mathbf{C}_\bullet(\text{Cone}(K))$ to the composite $\mathbf{C}_\bullet(g \circ f)$. This will be accomplished in one of the exercises to this Chapter.

4.5 THE SNAKE LEMMA

Our study of homotopy equivalence benefited greatly from a thorough analysis of contractible spaces, i.e., the spaces which have the simplest possible homotopy type. For analogous reasons, we ask which chain complexes have trivial homology.

DEFINITION 4.17. A sequence of vector spaces and linear maps

$$\cdots \xrightarrow{a_{k+2}} V_{k+1} \xrightarrow{a_{k+1}} V_k \xrightarrow{a_k} V_{k-1} \xrightarrow{a_{k-1}} \cdots$$

is said to be **exact at** k if ker a_k equals img a_{k+1} as subpaces of V_k. The entire sequence is called **exact** if it is exact at every $k \in \mathbb{N}$.

A casual glance at Definition 3.9 will confirm that every exact sequence is a chain complex, and another brief look at Definition 3.11 reveals that exact sequences are precisely those chain complexes whose homology group is trivial in every dimension $k \geq 0$. We call an exact sequence **short** if all but three of the V_i (let's say V_0, V_1 and V_2 without loss of generality) are required to be trivial. Short exactness relates to standard notions in linear algebra, for instance:

(1) $0 \to V_1 \to V_0$ is exact at $k = 1$, iff $V_1 \to V_0$ is injective,
(2) $V_2 \to V_1 \to 0$ is exact at $k = 1$ iff $V_2 \to V_1$ is surjective,
(3) $0 \to V_2 \to V_1 \to 0$ is exact at $k \in \{1, 2\}$ iff $V_2 \to V_1$ is an isomorphism, and
(4) $0 \to V_2 \to V_1 \to V_0 \to 0$ is exact iff $V_1 = V_0 \oplus V_2$.

The first three of these statements hold in broader contexts (i.e, we can replace the vector spaces by abelian groups) whereas the last one is specific to vector spaces. The definition of a short exact sequence also extends verbatim to chain complexes.

DEFINITION 4.18. A **short exact sequence of chain complexes** consists of three chain complexes and two chain maps arranged as follows:

$$(C_\bullet, d_\bullet) \xrightarrow{\phi_\bullet} (C'_\bullet, d'_\bullet) \xrightarrow{\psi_\bullet} (C''_\bullet, d''_\bullet),$$

with the additional requirement that for each $k \geq 0$ the chain groups

$$0 \longrightarrow C_k \xrightarrow{\phi_k} C'_k \xrightarrow{\psi_k} C''_k \longrightarrow 0$$

form a short exact sequence of $\mathbb{F}$-vector spaces.

The following lemma is by far the most important result in this Chapter; it forms the first of many miracles in the field of *homological algebra*.

LEMMA 4.19. **[The Snake lemma.]** *For each short exact sequence of chain complexes*

$$(C_\bullet, d_\bullet) \xrightarrow{\phi_\bullet} (C'_\bullet, d'_\bullet) \xrightarrow{\psi_\bullet} (C''_\bullet, d''_\bullet),$$

there exists a family of linear maps $D_k : \mathbf{H}_k(C''_\bullet, d''_\bullet) \to \mathbf{H}_{k-1}(C_\bullet, d_\bullet)$ which fit into an exact sequence of homology groups:

$$\cdots \xrightarrow{D_{k-1}} \mathbf{H}_k(C_\bullet, d_\bullet) \xrightarrow{\mathbf{H}_k\phi} \mathbf{H}_k(C'_\bullet, d'_\bullet) \xrightarrow{\mathbf{H}_k\psi} \mathbf{H}_k(C''_\bullet, d''_\bullet) \xrightarrow{D_k} \mathbf{H}_{k-1}(C_\bullet, d_\bullet) \xrightarrow{\mathbf{H}_{k-1}\phi} \cdots$$

The collection of linear maps $\{D_k \mid k \geq 1\}$ is called the **connecting homomorphism** of the given short exact sequence. The full proof of this lemma is a tedious affair, and tends to be far from enlightening. We will say just enough about it here to explain the serpentine etymology. To build D_k, one starts with the piece of the short exact sequence connecting dimensions k and $k-1$:

$$\begin{tikzcd}
0 \arrow[r] & C_k \arrow[r, hook, "\varphi_k"] \arrow[d, "d_k"] & C'_k \arrow[r, two heads, "\psi_k"] \arrow[d, "d'_k"] & C''_k \arrow[r] \arrow[d, "d''_k"] & 0 \\
0 \arrow[r] & C_{k-1} \arrow[r, hook, "\varphi_{k-1}"'] & C'_{k-1} \arrow[r, two heads, "\psi_{k-1}"'] & C''_{k-1} \arrow[r] & 0
\end{tikzcd}$$

Since both rows are exact by Definition 4.18, the ϕ maps are injective while the ψ maps are surjective. We'd like D_k to send elements of the homology group $\mathbf{H}_k(C''_\bullet, d''_\bullet)$ to elements of

$\mathbf{H}_{k-1}(C_\bullet, d_\bullet)$, so it makes sense to start with the upper-right corner of this diagram. There are four basic steps in the construction:

(1) Choose any γ lying in $\ker d''_k \subset C''_k$.
(2) By surjectivity, there is some β in C'_k satisfying $\psi_k(\beta) = \gamma$.
(3) Since ψ is a chain map, Definition 4.6 gives us

$$\psi_{k-1} \circ d'_k(\beta) = d''_k \circ \psi_k(\beta) = d''_k(\gamma) = 0;$$

thus, $d'_k(\beta)$ lies in $\ker \psi_{k-1}$.
(4) By exactness of the bottom row, this kernel equals the image of ϕ_{k-1}, so there is some α in C_{k-1} satisfying $\phi_{k-1}(\alpha) = d'_k(\beta)$.

One defines $D_k(\gamma) = \alpha$ as the desired map. The promised snake materializes when we trace the path taken in our short exact sequence $\gamma \mapsto \beta \mapsto d'_k\beta \mapsto \alpha$:

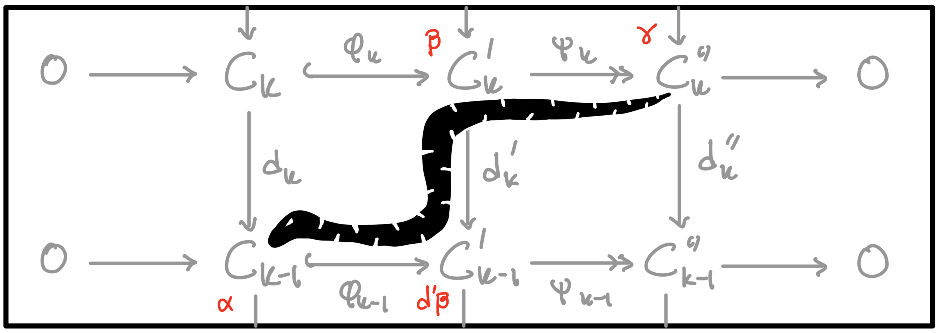

The argument is far from complete: one must show that D_k defines a well-defined map on homology independent of our choice of β, that α lies in $\ker d_{k-1}$, and that the sequence involving $\mathbf{H}_k\phi, \mathbf{H}_k\psi$ and D_k is exact. We will only perform the second check here:

$$\begin{aligned} \phi_{k-2} \circ d_{k-1}(\alpha) &= d'_{k-1} \circ \phi_{k-1}(\alpha) && \text{by Definition 4.6,} \\ &= d'_{k-1} \circ d'_k(\beta) && \text{since } \phi_{k-1}(\alpha) = d'_k(\beta), \\ &= 0 && \text{by Definition 3.9.} \end{aligned}$$

But ϕ_{k-2} is injective by exactness, so $d_{k-1}(\alpha) = 0$ as desired.

4.6 Pairs and Relative Homology

One of the first applications of Lemma 4.19 is the ability to relate the homology groups of a simplicial complex K, a subcomplex $L \subset K$ and the topological quotient $|K|/|L|$. This quotient does not form a simplicial complex in any natural way, but we are still able to define its homology by building an appropriate quotient chain complex as follows. Each chain group $\mathbf{C}_k(L)$ is a subspace of the corresponding $\mathbf{C}_k(K)$ by Definition 3.6. And since tje faces of every simplex in L themselves lie in L by the subcomplex property, the restriction of ∂_k^K to $\mathbf{C}_k(L)$ coincides with ∂_k^L by Definition 3.7. Thus, ∂_k^K induces a well-defined map of quotient spaces, which we denote

$$\partial_k^{K,L} : \mathbf{C}_k(K)/\mathbf{C}_k(L) \to \mathbf{C}_{k-1}(K)/\mathbf{C}_{k-1}(L).$$

Since $\partial_\bullet^K$ is a boundary operator, it follows that $\partial_k^{K,L} \circ \partial_{k+1}^{K,L} = 0$.

DEFINITION 4.20. Let $L \subset K$ be a pair of simplicial complexes; the **relative homology groups** $\mathbf{H}_k(K, L)$ are defined to be the homology groups of the chain complex defined as follows: its chain groups are

$$\mathbf{C}_k(K, L) = {}^{\mathbf{C}_k(K)}\!/\!{}_{\mathbf{C}_k(L)},$$

and the boundary operators are $\partial_k^{K,L} : \mathbf{C}_k(K, L) \to \mathbf{C}_{k-1}(K, L)$.

The Snake lemma enters the picture because whenever $L \subset K$ is a subcomplex, we have an apparent short exact sequence of chain complexes

$$(\mathbf{C}_\bullet(L), \partial_\bullet^L) \xhookrightarrow{\iota_\bullet} (\mathbf{C}_\bullet(K), \partial_\bullet^K) \xrightarrow{\pi_\bullet} (\mathbf{C}_\bullet(K, L), \partial_\bullet^{K,L}).$$

Here the chain map $\iota_\bullet$ is given by inclusions of subspaces while the chain map $\pi_\bullet$ is given by projections to quotient spaces. Applying Lemma 4.19 to this short exact sequence, we obtain a connecting homomorphism $D_k : \mathbf{H}_k(K, L) \to \mathbf{H}_{k-1}(L)$ and hence the following exact sequence relating homology groups.

DEFINITION 4.21. The **exact sequence of the pair** $L \subset K$ of simplicial complexes is given by

$$\cdots \xrightarrow{D_{k+1}} \mathbf{H}_k(L) \xrightarrow{\mathbf{H}_k\iota} \mathbf{H}_k(K) \xrightarrow{\mathbf{H}_k\pi} \mathbf{H}_k(K, L) \xrightarrow{D_k} \mathbf{H}_{k-1}(L) \xrightarrow{\mathbf{H}_{k-1}\iota} \cdots$$

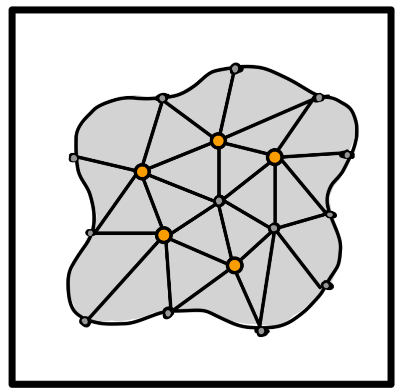

The exact sequence of a pair is a wonderful tool for computing relative homology groups $\mathbf{H}_\bullet(K, L)$ using prior knowledge of $\mathbf{H}_k(K)$ and $\mathbf{H}_k(L)$. Consider, for instance, the scenario where K is any simplicial complex whose realization $|K|$ is homeomorphic to the 2-dimensional disk, while the subcomplex $L \subset K$ consists of n interior vertices — the case $n = 5$ has been illustrated. Building the chain complex $\mathbf{C}_\bullet(K, L)$ which yields the relative homology is quite a chore, but the exact sequence of a pair works remarkably well. We know (or can compute, if asked) that the homology of K agrees with that of $\Delta(2)$, whereas L consists of n disjoint copies of $\Delta(0)$. Putting all this known information about K and L together, we have:

$$\mathbf{H}_k(K; \mathbb{F}) = \begin{cases} \mathbb{F} & k = 0 \\ 0 & k > 0 \end{cases} \quad \text{and} \quad \mathbf{H}_k(L; \mathbb{F}) = \begin{cases} \mathbb{F}^n & k = 0 \\ 0 & k > 0 \end{cases}.$$

All the non-trivial bits of the exact sequence of this pair concentrate in the lower dimensions — here is the relevant piece of the sequence:

$$\cdots \longrightarrow 0 \longrightarrow \mathbf{H}_1(K, L) \xrightarrow{D_0} \mathbf{H}_0(L) \xrightarrow{\mathbf{H}_0\iota} \mathbf{H}_0(K) \xrightarrow{\mathbf{H}_0\pi} \mathbf{H}_0(K, L) \longrightarrow 0$$

Everything depends on the rank of the map $\mathbf{H}_0\iota$ which is induced on 0-th homology by the inclusion of L into K. It is straightforward to check that this is not the zero map, and hence has rank 1. Now exactness of this sequence immediately forces the rank of D_0 to be $n - 1$ and the rank of $\mathbf{H}_0\pi$ to be 0. But D_0 is injective and $\mathbf{H}_0\pi$ is surjective (because of the leading and trailing 0's plus exactness), which gives

$$\mathbf{H}_k(K, L) = \begin{cases} \mathbb{F}^{n-1} & k = 1 \\ 0 & k \neq 1 \end{cases}.$$

REMARK 4.22. The relative homology of a pair $L \subset K$ generalizes ordinary simplicial homology of K if we allow ourselves the luxury of setting $L = \varnothing$; in this case the chain groups

$\mathbf{C}_\bullet(K)$ and $\mathbf{C}_\bullet(K,L)$ are equal, as are the boundary operators. On the other hand, the relative homology of a pair is further generalized by that of a **triple** $M \subset L \subset K$ of simplicial complexes. Here the short exact sequence of interest is

$$(\mathbf{C}_\bullet(L,M), \partial_\bullet^{L,M}) \xhookrightarrow{\iota_\bullet} (\mathbf{C}_\bullet(K,M), \partial_\bullet^{K,M}) \xrightarrow{\pi_\bullet} (\mathbf{C}_\bullet(K,L), \partial_\bullet^{K,L}),$$

Once again, the chain map $\iota_\bullet$ is an inclusion while the map π is a projection; the subcomplex property is necessary to get well-defined boundary operators of these relative chain complexes (as it was in Definition 4.20). The connecting homomorphisms $D_k : \mathbf{H}_k(K,L) \to \mathbf{H}_{k-1}(L,M)$ guaranteed by Lemma 4.19 fit into an exact sequence with $\mathbf{H}_\bullet\iota$ and $\mathbf{H}_\bullet\pi$.

4.7 The Mayer-Vietoris Sequence

A second enormously useful application of the Snake lemma is that it confers the ability to compute homology of a complicated simplicial complex K in terms of a decomposition into two (hopefully simpler) subcomplexes. Assume that L and M are subcomplexes of K satisfying $K = L \cup M$, and let's agree to write their intersection $L \cap M$ – which must also be a subcomplex of K – as I. There are now four chain complexes and four chain maps (all inclusions) to keep track of; these fit into the following diamond:

$$\begin{tikzcd}
 & (\mathbf{C}_\bullet(I), \partial_\bullet^I) \arrow[dl, hook] \arrow[dr, hook] & \\
(\mathbf{C}_\bullet(L), \partial_\bullet^L) \arrow[dr, hook] & & (\mathbf{C}_\bullet(M), \partial_\bullet^M) \arrow[dl, hook] \\
 & (\mathbf{C}_\bullet(K), \partial_\bullet^K) &
\end{tikzcd}$$

Both paths from the top to the bottom give the same chain map (the one which includes chains of I into chains of K); thus our diamond commutes in the category $\mathbf{Chain}_\mathbb{F}$. The crucial idea here is to generate a short exact sequence by combining the two chain complexes of the middle row into a single one.

The *direct sum* of $(\mathbf{C}_\bullet(L), \partial_\bullet^L)$ and $(\mathbf{C}_\bullet(M), \partial_\bullet^M)$ is the new chain complex defined as follows: in each dimension $k \geq 0$, it has

$$\text{chain groups } \mathbf{C}_k(L) \oplus \mathbf{C}_k(M) \text{ and boundary operator } \begin{bmatrix} \partial_k^L & 0 \\ 0 & \partial_k^M \end{bmatrix}.$$

The k-th homology group of this direct sum is $\mathbf{H}_k(L) \oplus \mathbf{H}_k(M)$. More interestingly, there is an injective chain map $\iota_k : \mathbf{C}_k(I) \to \mathbf{C}_k(M) \oplus \mathbf{C}_k(L)$ which sends every γ to the pair (γ, γ). There is also a second chain map $\pi_k : \mathbf{C}_k(M) \oplus \mathbf{C}_k(L) \to \mathbf{C}_k(K)$ that sends each pair (α, β) to the difference $(\beta - \alpha)$. This map $\pi_\bullet$ is evidently surjective because $K = L \cup M$; thus, we obtain

$$(\mathbf{C}_\bullet(I), \partial_\bullet^I) \xhookrightarrow{\iota_\bullet} (\mathbf{C}_\bullet(L), \partial_\bullet^L) \oplus (\mathbf{C}_\bullet(M), \partial_\bullet^M) \xrightarrow{\pi_\bullet} (\mathbf{C}_\bullet(K), \partial_\bullet^K).$$

This turns out to be a short exact sequence: note that $(\alpha, \beta) \in \mathbf{C}_k(L) \cap \mathbf{C}_k(M)$ lies in $\ker \pi_k$ if and only if $\alpha = \beta$. But this equality holds if and only if the chain α lies in the intersection $\mathbf{C}_k(I) = \mathbf{C}_k(L) \cap \mathbf{C}_k(M)$, whence (α, α) lies in $\operatorname{img} \iota_k$. Having obtained a short exact sequence of chain complexes, we appeal once more to the Snake lemma and obtain a connecting homomorphism $D_k : \mathbf{H}_k(K) \to \mathbf{H}_{k-1}(I)$.

DEFINITION 4.23. Let $K = L \cup M$ be a decomposition of the simplicial complex K into two subcomplexes L and M whose intersection is denoted I. The **Mayer-Vietoris exact sequence** associated to this partition is given by

$$\cdots \xrightarrow{D_{k-1}} \mathbf{H}_k(I) \xrightarrow{\mathbf{H}_k\iota} \mathbf{H}_k(L) \oplus \mathbf{H}_k(M) \xrightarrow{\mathbf{H}_k\pi} \mathbf{H}_k(K) \xrightarrow{D_k} \mathbf{H}_{k-1}(I) \xrightarrow{\mathbf{H}_{k-1}\iota} \cdots$$

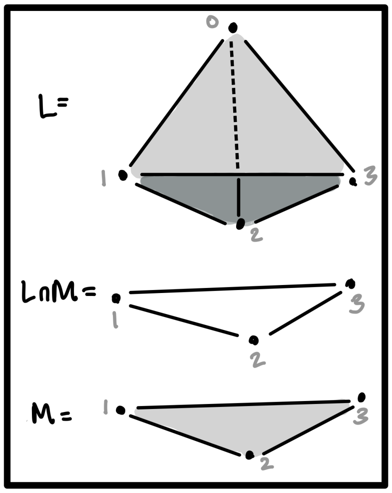

This exact sequence is particularly effective when combined with inductive arguments — we can use it to compute the i-th homology group of every hollow k-simplex $\partial\Delta(k)$ for $i > 1$. Consider the decomposition $\partial\Delta(k) = L \cup M$ where L is the closed star of the vertex 0 (see Definition 1.17) while M consists of the simplex $\{1, 2, \ldots, k\}$ along with all its faces. This decomposition is illustrated for $k = 3$ here. Note also that the intersection $L \cap M$ is the hollow simplex of one lower dimension, i.e., $\partial\Delta(k-1)$.

Now we claim that both L and M have the same homology as $\Delta(0)$. First note that L is clearly a cone over $\partial\Delta(k-1)$, so the conclusion follows from Example 4.16. And M is a solid k-simplex, which is a cone over a solid $(k-1)$-simplex, so once again Example 4.16 does the job. Consequently, the homology groups $\mathbf{H}_i(L)$ and $\mathbf{H}_i(M)$ are trivial for all $i > 0$, and hence so is their direct sum. So for each $i > 1$, we obtain the following snippet of the Mayer-Vietoris exact sequence:

$$0 \longrightarrow \mathbf{H}_i(\partial\Delta(k)) \xrightarrow{D_i} \mathbf{H}_{i-1}(\partial\Delta(k-1)) \longrightarrow 0$$

Exactness forces D_i to be an isomorphism for all $i > 1$, so it suffices to calculate the homology groups $\mathbf{H}_i(\partial\Delta(2); \mathbb{F})$ as a base case; we did this in Example 3.12, and can safely conclude that for $i > 0$ we have:

$$\mathbf{H}_i(\partial\Delta(k); \mathbb{F}) = \begin{cases} \mathbb{F} & i = k-1 \\ 0 & \text{otherwise.} \end{cases}$$

A separate (and somewhat easier) argument must be used to compute $\mathbf{H}_0(\partial\Delta(k))$.

4.8 BONUS: HOMOTOPY INVARIANCE

Theorem 4.24 below is vital from both a theoretical and practical perspective; its proof requires techniques which are outside our scope at the moment, but the ability to understand and apply it will be quite beneficial when working with homology.

As mentioned at the beginning of Chapter 3, the Euler characteristic inherits its homotopy invariance from homology.

THEOREM 4.24. *Let K and L be simplicial complexes. For any choice of coefficient field $\mathbb{F}$,*
(1) *if $f, g : K \to L$ are homotopic simplicial maps, then $\mathbf{H}_k f = \mathbf{H}_k g$ for every $k \geq 0$; and,*
(2) *if K and L are homotopy equivalent, then $\mathbf{H}_k(K)$ is isomorphic to $\mathbf{H}_k(L)$ for every $k \geq 0$.*

The second assertion follows from the first one if we use simplicial approximation (see Theorem 2.15). The basic idea is to start with topology and gradually descend to algebra: Assume that $\theta : |K| \times [0,1] \to |L|$ is a homotopy from $|f|$ to $|g|$. The first order of business is to build a simplicial complex homeomorphic to $|K| \times [0,1]$ — this is rendered difficult by the fact that in general the product of a simplex with $[0,1]$ is not itself a simplex in any natural way. Fortunately, such a product can be triangulated into a union of simplices, and putting these together produces a simplicial complex $P(K)$ whose realization is homeomorphic to $|K| \times [0,1]$. Using this homeomorphism, one approximates the homotopy θ as a simplicial map $\mathbf{Sd}^n\, P(K) \to L$ (where **Sd** stands for barycentric subdivision). This approximated version of θ then descends to a chain homotopy from $\mathbf{C}_\bullet f$ to $\mathbf{C}_\bullet g$. An appeal to Proposition 4.14 completes the argument.

Exercises

EXERCISE 4.1. Given functors $\boldsymbol{F} : \mathscr{C} \to \mathscr{C}'$ and $\boldsymbol{G} : \mathscr{C}' \to \mathscr{C}''$, define their composite $\boldsymbol{G} \circ \boldsymbol{F}$ and show that this is a functor $\mathscr{C} \to \mathscr{C}''$.

EXERCISE 4.2. Show that the collection of all simplicial complexes and simplicial maps satisfies the axioms of a category **SC**.

EXERCISE 4.3. Consider two chain maps $\phi : (C_\bullet, d_\bullet) \to (C'_\bullet, d'_\bullet)$ and $\psi : (C'_\bullet, d'_\bullet) \to (C''_\bullet, d''_\bullet)$. Show that the collection of maps $\psi_k \circ \phi_k : C_k \to C''_k$ prescribe a chain map from $(C_\bullet, d_\bullet)$ to $(C''_\bullet, d''_\bullet)$. Thus, chain maps are morphisms in the category $\mathbf{Chain}_{\mathbb{F}}$ of chain complexes over $\mathbb{F}$.

EXERCISE 4.4. Given simplicial maps $f : K \to L$ and $g : L \to M$, show that $\mathbf{C}_k(g \circ f)$ equals $\mathbf{C}_k(g) \circ \mathbf{C}_k(f)$. This shows that $\mathbf{C}$ is a functor from the category **SC** of Exercise 4.2 to the category $\mathbf{Chain}_{\mathbb{F}}$ of Exercise 4.3.

EXERCISE 4.5. Write down a proof of Proposition 4.8.

EXERCISE 4.6. Verify the assertions of Example 4.11.

EXERCISE 4.7. In the setting of Example 4.11, consider the simplicial map $h : K \to \partial\Delta(2)$ that sends vertex a to 0, vertex d to 1 and vertex c to 2. Show that $\mathbf{H}_k h$ is an inverse to $\mathbf{H}_k g$ for every $k \geq 1$. (Note that h and g themselves are not inverse to each other as chain maps!)

EXERCISE 4.8. Prove that chain homotopy is an equivalence relation on the set of all chain maps $(C_\bullet, d_\bullet) \to (C'_\bullet, d'_\bullet)$.

EXERCISE 4.9. Using $\mathbb{F} = \mathbb{Z}/2$ coefficients, complete the argument of Example 4.16 as follows. Define the linear maps $\eta_k : \mathbf{C}_k(\mathrm{Cone}(K)) \to \mathbf{C}_{k+1}(\mathrm{Cone}(K))$ that sends each basis k-simplex σ to

$$\eta_k(\sigma) = \begin{cases} \sigma \cup \{v_*\} & \sigma \in K \\ 0 & \sigma \in \mathrm{Cone}(K) - K. \end{cases}$$

Show that $\eta_\bullet$ prescribes a chain homotopy between the chain map $\phi_\bullet := \mathbf{C}_\bullet(g \circ f)$ and the identity chain map. [Hint: let d be the boundary operator for the simplicial chain complex of $\mathrm{Cone}(K)$. Over $\mathbb{F} = \mathbb{Z}/2$ it suffices to show that $\sigma + \phi(\sigma) = d \circ \eta(\sigma) + \eta \circ d(\sigma)$ for each simplex σ. Start with $\dim \sigma = 0$ and induct upwards along dimension.]

EXERCISE 4.10. For each $k \geq 1$, compute the relative homology group $\mathbf{H}_k(\Delta(k), \partial\Delta(k))$.

EXERCISE 4.11. Let K and L be simplicial complexes. Identify a vertex v of K with a vertex w of L to form a new simplicial complex $K \vee L$. Prove that $\mathbf{H}_k(K \vee L) = \mathbf{H}_k(K) \oplus \mathbf{H}_k(L)$ for all $k > 0$ [Hint: Mayer-Vietoris].

5. Cohomology

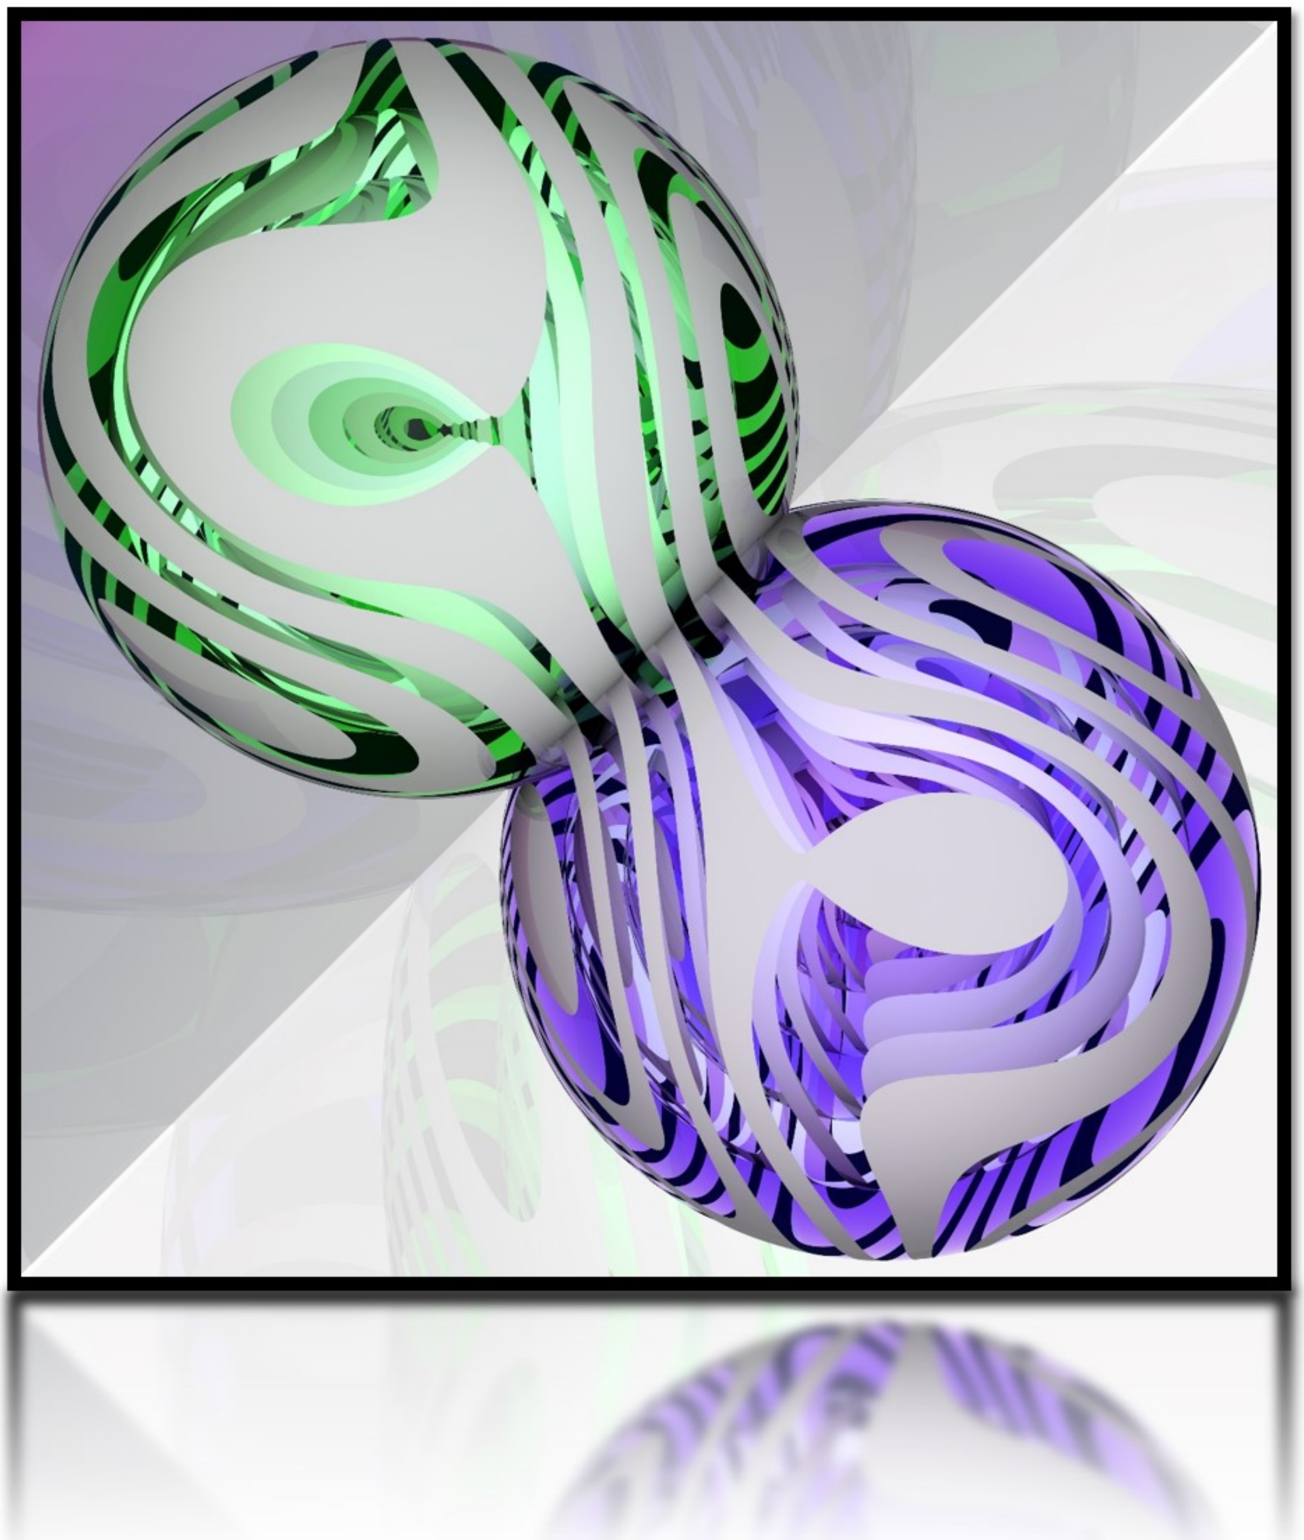

5.1 Cochain Complexes

The **dual** V^* of a $\mathbb{F}$-vector space consists of all linear maps $V \to \mathbb{F}$. It is not too painful to confirm that V^* is also a vector space over $\mathbb{F}$ — given linear maps p, q in V^* along with scalars α, β in $\mathbb{F}$, the linear combination $\alpha \cdot p + \beta \cdot q$ is evidently another linear map $V \to \mathbb{F}$ and hence constitutes an element of V^*. For finite-dimensional V one can describe the elements of V^* quite explicitly — every basis $\{e_1, \ldots, e_k\} \subset V$ has a corresponding *dual basis* $\{e_1^*, \ldots, e_k^*\} \subset V^*$ prescribed by the following action on the V-basis:

$$e_i^*(e_j) = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}.$$

Thus, we can transport any basis for V to a basis for V^* and express all the elements of V^* in terms of this dual basis.

Life gets considerably more interesting when one similarly attempts to dualize a linear map $A : V \to W$ of $\mathbb{F}$-vector spaces. Now A does not give us any straightforward way of sending V^*-elements to W^*-elements — every $p : V \to \mathbb{F}$ fits into an awkward zigzag with A:

$$W \xleftarrow{A} V \xrightarrow{p} \mathbb{F},$$

In sharp contrast, if we start with $q : W \to \mathbb{F}$, then there *is* an obvious map $V \to \mathbb{F}$:

$$V \xrightarrow{A} W \xrightarrow{q} \mathbb{F}.$$

Thus, for every $A : V \to W$ we get a **dual map** $A^* : W^* \to V^*$ which acts as $q \mapsto q \circ A$. Our goal here is to investigate some of the homological consequences of this dramatic reversal of domain and codomain that occurs when we dualize linear maps.

Let's start with a chain complex $(C_\bullet, d_\bullet)$ over $\mathbb{F}$

$$\cdots \xrightarrow{d_{k+1}} C_k \xrightarrow{d_k} C_{k-1} \xrightarrow{d_{k-1}} \cdots \xrightarrow{d_2} C_1 \xrightarrow{d_1} C_0 \longrightarrow 0$$

and dualize everything in sight:

$$\cdots \xleftarrow{d_{k+1}^*} C_k^* \xleftarrow{d_k^*} C_{k-1}^* \xleftarrow{d_{k-1}^*} \cdots \xleftarrow{d_2^*} C_1^* \xleftarrow{d_1^*} C_0 \longleftarrow 0$$

The important fact from out perspective is that even in this dualized form, adjacent maps compose to give zero; given any dimension $k \geq 0$ and linear map $\zeta : C_k \to \mathbb{F}$, we have $d_{k+2}^* \circ d_{k+1}^*(\zeta) = \zeta \circ d_{k+1} \circ d_{k+2}$, which must equal zero regardless of ζ by the defining property of a chain complex. If we write this dualized chain complex from left to right and shift the indexing of the dual boundary maps by 1, then we arrive at the following definition.

Definition 5.1. A **cochain complex** $(C^\bullet, d^\bullet)$ over $\mathbb{F}$ is a sequence of vector spaces and linear maps of the form

$$0 \longrightarrow C^0 \xrightarrow{d^0} C^1 \xrightarrow{d^1} C^2 \xrightarrow{d^2} \cdots \xrightarrow{d^{k-1}} C^k \xrightarrow{d^k} C^{k+1} \xrightarrow{d^{k+1}} \cdots$$

satisfying $d^{k-1} \circ d^k = 0$ for every $k \geq 1$.

Aside from the fact that the maps are going up the indexing rather than down, cochain complexes are not very different from the chain complexes of Definition 3.9. We call C^k the k-th *cochain group* and $d^k : C^k \to C^{k+1}$ the k-th *coboundary map* of $(C^\bullet, d^\bullet)$.

5.2 Cohomology

Let $(C^\bullet, d^\bullet)$ be a cochain complex over a field $\mathbb{F}$.

Definition 5.2. For each dimension $k \geq 0$, the k-th **cohomology group** of $(C^\bullet, d^\bullet)$ is the quotient vector space

$$\mathbf{H}^k(C^\bullet, d^\bullet) = {\ker d^k}/{\operatorname{img} d^{k-1}}$$

Elements of $\ker d^k$ are called *k-cocycles* while elements of $\operatorname{img} d^{k-1}$ are the *k-coboundaries*. To acquire geometric intuition for cohomology, we will retreat to the relative comfort of simplicial complexes.

Let K be a simplicial complex, so that each chain group $\mathbf{C}_k(K)$ is generated by treating the k-simplices as basis elements. Thus, each k-simplex σ in K corresponds to a distinguished cochain $\sigma^* : \mathbf{C}_k(K) \to \mathbb{F}$, defined by (linearity and) the following action on any given k-simplex τ:

$$\sigma^*(\tau) = \begin{cases} 1 & \tau = \sigma \\ 0 & \tau \neq \sigma \end{cases}$$

The collection of such cochains $\{\sigma^* : \mathbf{C}_k(K) \to \mathbb{F} \mid \dim(\sigma) = k\}$ forms a basis for the group of k-cochains of K. It is customary to write $\mathbf{C}^k(K)$ rather than the cumbersome $\mathbf{C}_k(K)^*$ to denote this *simplicial cochain group* of K — there is a long-standing convention in algebraic topology to index homology with subscripts and cohomology with superscripts.

The k-th *simplicial coboundary operator* is (unsurprisingly) denoted $\partial^k_K : \mathbf{C}^k(K) \to \mathbf{C}^{k+1}(K)$; by definition, this is the dual to the boundary operator $\partial^K_{k+1} : \mathbf{C}_{k+1}(K) \to \mathbf{C}_k(K)$, and hence satisfies $\partial^k_K(\sigma^*) = \sigma^* \circ \partial^K_{k+1}$ for each σ^* in $\mathbf{C}^k(K)$. It follows that for each general cochain ξ in $\mathbf{C}^k(K)$ and oriented $(k+1)$-simplex $\sigma = (v_0, \ldots, v_{k+1})$, we have the wonderfully convenient formula

$$\partial^k_K \xi(\sigma) = \sum_{i=0}^{k} (-1)^i \cdot \xi(\sigma_{-i}), \tag{3}$$

where σ_{-i} is the face of σ obtained by deleting the vertex v_i. Thus, with respect to our choice of basis elements, ∂^k_K is simply the transpose[1] of the boundary matrix ∂^K_{k+1} for each $k \geq 0$; we will discuss three advantages of adopting this perspective shortly. In any event, the k-th cohomology group of the simplicial cochain complex $(\mathbf{C}^\bullet(K), \partial^\bullet_K)$ is called the k-th **simplicial cohomology group** of K and denoted by the shorthand $\mathbf{H}^k(K; \mathbb{F})$ or simply $\mathbf{H}^k(K)$.

The first advantage of realizing that coboundary operators are transposes of boundary operators (with respect to our simplex-induced basis) is the ability to visualize low-dimensional simplicial cocycles at least over $\mathbb{F} = \mathbb{Z}/2$:

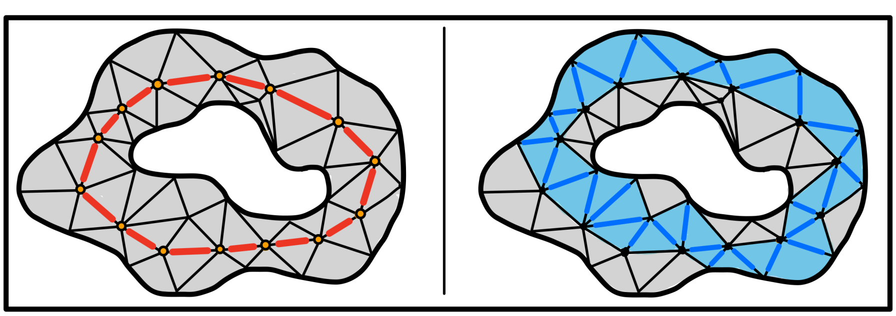

[1]When working with $\mathbb{F} = \mathbb{C}$ coefficients, this becomes a conjugate transpose.

To the left is a 1-cycle γ in a triangulated annulus, which we last saw when studying homology in Definition 3.11; and to the right we have a 1-cocycle ξ in the same annulus. All the edges being sent to 1 by ξ have been highlighted. On the left, every vertex had to be the face of an even number of edges in γ (otherwise the boundary $\partial_1(\gamma)$ would be nonzero). On the right, every 2-simplex must contain an even number of edges from ξ in its boundary (otherwise the coboundary $\partial^1(\xi)$ will be nonzero).

A second advantage is that we can also see in small examples when two cocycles lie in the same cohomology class; our 1-cocycle ξ represents the same cohomology class as new cocycle ξ' shown on the right, since they differ only by the coboundary of the highlighted vertex:

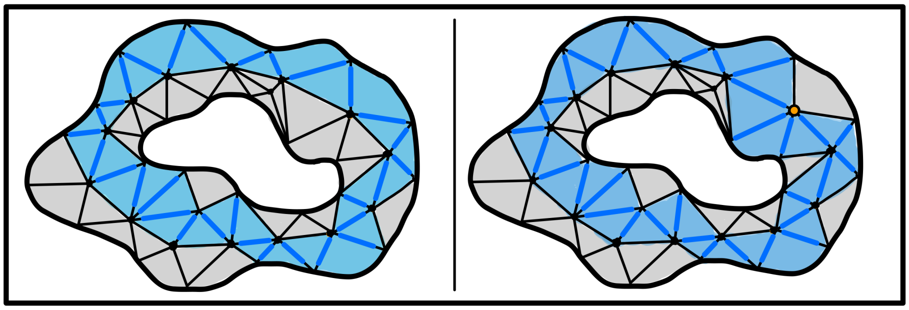

The third advantage of realizing that ∂_K^k is the transpose of ∂_{k+1}^K is the knowledge that they must have the same ranks as linear maps.

PROPOSITION 5.3. *Let* $(C_\bullet, d_\bullet)$ *be a chain complex over a field* $\mathbb{F}$ *so that* $\dim C_k$ *is finite for all* $k \geq 0$*, and let* $(C^\bullet, d^\bullet)$ *be its dual cochain complex. Then, we have*

$$\dim \mathbf{H}_k(C_\bullet, d_\bullet) = \dim \mathbf{H}^k(C^\bullet, d^\bullet)$$

in each dimension $k \geq 0$.

PROOF. This follows from the fact that $\dim C^k = \dim C_k$ and $\operatorname{rank} d^{k+1} = \operatorname{rank} d_k$ for all k:

$$\begin{aligned}
\dim \mathbf{H}_k(C_\bullet, d_\bullet) &= \dim\ker d_k - \dim\operatorname{img} d_{k+1} \\
&= (\dim C_k - \operatorname{rank} d_k) - \operatorname{rank} d_{k+1} \\
&= (\dim C^k - \operatorname{rank} d^{k-1}) - \operatorname{rank} d^k \\
&= (\dim C^k - \operatorname{rank} d^k) - \operatorname{rank} d^{k-1} \\
&= \dim\ker d^k - \dim\operatorname{img} d^{k-1} = \dim \mathbf{H}^k(C^\bullet, d^\bullet).
\end{aligned}$$

In particular, we have $\dim \mathbf{H}_k(K;\mathbb{F}) = \dim \mathbf{H}^k(K;\mathbb{F})$ for every simplicial complex K. $\square$

The machinery developed for homology in the previous two chapters is readily translatable to work for cohomology, with the standard caveat that duality will force various maps to point in the opposite direction. For instance, every simplicial map $f : K \to L$ induces *cochain maps* $\mathbf{C}^\bullet f : (\mathbf{C}^\bullet(L), \partial_L^\bullet) \to (\mathbf{C}^\bullet(K), \partial_K^\bullet)$, which in turn yield well-defined linear maps

$$\mathbf{H}^k f : \mathbf{H}^k(L;\mathbb{F}) \to \mathbf{H}^k(K;\mathbb{F})$$

between cohomology groups. There is an avatar of Proposition 3.15 which allows us to extract bases of all the cohomology groups using Smith decompositions of coboundary matrices. Similarly, one can define cochain homotopies, relative cohomology groups and Mayer-Vietoris sequences for cohomology. This is a worthy endeavour, strongly recommended for all those who are encountering cohomology for the first time. Instead of reinventing that wheel here, we will focus on those aspects of cohomology which are new and different.

5.3 The Cup Product

The remarkable benefit of cochains over chains is that they are functions taking values in a field $\mathbb{F}$, so we can multiply them with each other. Fix an oriented simplicial complex K, so that each k-simplex σ can be uniquely written as an increasing list of vertices $(v_0, \ldots, v_k)$. It will be convenient henceforth to write, for each i in $\{0, \ldots, k\}$ the i-th *front face* of σ is the i-dimensional simplex $\sigma_{\leq i} = (v_0, \ldots, v_i)$, and similarly the i-th *back face* of σ is the $(k-i)$-dimensional simplex $\sigma_{\geq i} = (v_i, \ldots, v_k)$.

DEFINITION 5.4. Let $\xi \in \mathbf{C}^k(K)$ and $\eta \in \mathbf{C}^\ell(K)$ be two simplicial cochains of K. Their **cup product** is a new cochain $\xi \smile \eta$ in $\mathbf{C}^{k+\ell}(K)$ defined by the following action on each $(k+\ell)$-dimensional simplex σ:

$$\xi \smile \eta(\sigma) = \xi(\sigma_{\leq k}) \cdot \eta(\sigma_{\geq k}).$$

(Here the multiplication on the right side takes place in the underlying field $\mathbb{F}$.)

Having obtained a new cochain $\xi \smile \eta$ by suitably multiplying ξ with η, we should lay to rest any curiosity regarding its coboundary.

PROPOSITION 5.5. *For any ξ in $\mathbf{C}^k(K;\mathbb{F})$ and η in $\mathbf{C}^\ell(K;\mathbb{F})$, we have*

$$\partial_K^{k+\ell}(\xi \smile \eta) = [\partial_K^k(\xi) \smile \eta] + (-1)^k \cdot [\xi \smile \partial_K^\ell(\eta)].$$

PROOF. Let τ be a $(k+\ell+1)$-dimensional oriented simplex with vertices $(v_0, \ldots, v_{k+\ell+1})$. We evaluate the two terms on the right side of the desired equality separately on τ. First,

$$\begin{aligned}
[\partial_K^k(\xi) \smile \eta](\tau) &= \partial_K^k(\xi)(\tau_{\leq k+1}) \cdot \eta(\tau_{\geq k+1}) && \text{by Definition 5.4,} \\
&= \left(\sum_{i=0}^{k+1} (-1)^i \cdot \xi((\tau_{\leq k+1})_{-i}) \cdot \eta(\tau_{\geq k+1}) \right) && \text{by (3).}
\end{aligned}$$

And similarly,

$$(-1)^k \cdot [\xi \smile \partial_K^\ell(\eta)](\tau) = \left(\sum_{j=0}^{\ell+1} (-1)^{k+j} \cdot \xi(\tau_{\leq k}) \cdot \eta((\tau_{\geq k})_{-j}) \right).$$

When we add these two expressions, the $i = k+1$ term of the first sum cancels the $j = 0$ term of the second; the terms which survive are exactly $\partial_K^{k+\ell}(\xi \smile \eta)(\tau)$ by (3). □

Using the above formula for the coboundary of $\xi \smile \eta$, one can confirm that the cup product of two cocycles is again a cocycle:

$$\begin{aligned}
\partial_K^{k+\ell}(\xi \smile \eta) &= [\partial_K^k(\xi) \smile \eta] + (-1)^k[\xi \smile \partial_K^\ell(\eta)] && \text{by Proposition 5.5,} \\
&= [0 \smile \eta] + (-1)^k[\xi \smile 0] && \text{since } \xi \text{ and } \eta \text{ are cocycles,} \\
&= 0 && \text{by (3).}
\end{aligned}$$

Now if $\xi = \partial_K^{k-1}(\xi')$ is a coboundary while and η is a cocycle as before, then their cup product is a coboundary:

$$\begin{aligned}
\xi \smile \eta &= \partial_K^{k-1}(\xi') \smile \eta \\
&= [\partial^{k-1}(\xi') \smile \eta] + (-1)^k \cdot [\xi' \smile \partial_K^\ell(\eta)] && \text{since } \partial_K^\ell(\eta) = 0 \\
&= \partial_K^{k+\ell}(\xi' \smile \eta). && \text{by Proposition 5.5.}
\end{aligned}$$

Similarly, if ξ is a cocycle and η a coboundary, then again their cup product is a coboundary. We have arrived at the following result.

PROPOSITION 5.6. *For each simplicial complex K and dimensions $k, \ell \geq 0$, the cup product map $\smile: \mathbf{C}^k(K;\mathbb{F}) \times \mathbf{C}^\ell(K;\mathbb{F}) \to \mathbf{C}^{k+\ell}(K;\mathbb{F})$ induces a well-defined bilinear map of cohomology groups.*

It is customary to use the same notation when describing the cup product on cohomology groups rather than cochains, i.e.,

$$\smile: \mathbf{H}^k(K;\mathbb{F}) \times \mathbf{H}^\ell(K;\mathbb{F}) \to \mathbf{H}^{k+\ell}(K;\mathbb{F}).$$

The direct sum $\bigoplus_{k\geq 0} \mathbf{H}^k(K;\mathbb{F})$ is evidently a vector space over $\mathbb{F}$; writing its elements as

$$\underline{\xi} = (\xi_1, \xi_2, \ldots, \xi_k, \ldots),$$

we say that $\underline{\xi}$ has *grade* k if all the ξ_i for $i \neq k$ are zero. The cup product gives us a bilinear multiplication law on this direct sum which is additive on grades, i.e., the cup product of a grade-k element with a grade-ℓ element is a grade-$(k+\ell)$ element. A graded $\mathbb{F}$-vector space equipped with such a graded bilinear multiplication is called a **graded algebra** over $\mathbb{F}$. While the direct sum of homology groups $\bigoplus_{k\geq 0} \mathbf{H}_k(K;\mathbb{F})$ also forms a graded vector space, there is no multiplication law analogous to the cup product. It is in this sense that cohomology is considered a richer algebraic invariant than homology, even though the dimensions of cohomology groups agree with those of homology groups when working over a field.

EXAMPLE 5.7. By a **torus** we mean any simplicial complex T whose geometric realization is homeomorphic to the product $\partial\Delta(2) \times \partial\Delta(2)$. Consider also the simplicial complex W obtained by first taking the disjoint union $\partial\Delta(3) \sqcup \partial\Delta(2) \sqcup \partial\Delta(2)$, and then identifying the vertices labelled $\{0\}$ of all three pieces to create a single connected simplicial complex. Now one can check that T and W have isomorphic homology groups over any field $\mathbb{F}$, namely

$$\mathbf{H}_k(T) = \mathbf{H}_k(W) = \begin{cases} \mathbb{F} & k \in \{0,2\} \\ \mathbb{F}^2 & k = 1 \\ 0 & k > 2 \end{cases}.$$

Let α and β denote any two 1-cycles which span $\mathbf{H}_1$ and examine their cup product $\alpha \smile \beta$. In T, this will be (a nonzero multiple of) the unique cycle generating $\mathbf{H}_2$, whereas in W this cup product will equal zero.

The cup product $\alpha \smile \beta$ in the torus is nontrivial for a viscerally geometric reason; one can choose α to be a cochain that runs along the equator while β runs along the meridian. Now there will be at least one 2-simplex whose 1-dimensional faces are *all* sent to nonzero elements of $\mathbb{F}$ by either α or β. We highlight such a 2-simplex for the illustrated α and β below:

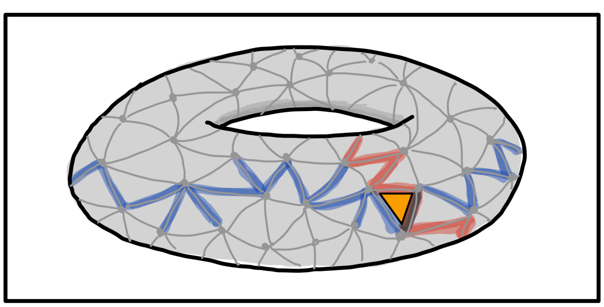

The miracle here is that no matter how much we perturb α and β within their respective cohomology classes, we will always have at least one such 2-simplex.

REMARK 5.8. There are no obstacles to defining cohomology with non-field coefficients, e.g., by using coefficients sourced from the ring of integers $\mathbb{Z}$. However, various subtleties arise from the fact that in general an abelian group G is not isomorphic to its *dual group* G^*; here G^* consists of all abelian group homomorphisms $G \to \mathbb{Z}$. In particular, G^* is blind to torsion in G and there is no analogue of Proposition 5.3 when using $\mathbb{Z}$ coefficients. Similarly, in this case the cup product prescribes the structure of a **graded ring** on the direct sum $\bigoplus_{k\geq 0} \mathbf{H}^k(K;\mathbb{Z})$ rather than a graded algebra.

5.4 THE CAP PRODUCT

There is a second (far stranger) product that mixes homology and cohomology. As before, let K be an oriented simplicial complex; each oriented k-simplex σ therefore has a front face $\sigma_{\leq i}$ and a back face $\sigma_{\geq i}$ for i in $\{0,\ldots,k\}$. Our new product arises from taking an i-cochain ξ for some $i \leq k$ and letting it act on σ by

$$\sigma \mapsto \xi(\sigma_{\leq i}) \cdot \sigma_{\geq i}.$$

That is, we evaluate σ on the front face of the appropriate dimension, and multiply the resulting scalar with the back face to produce a chain of dimension $(k-i)$. More formally, note that each k-chain γ in $\mathbf{C}_k(K)$ is uniquely expressible as a linear combination $\gamma = \sum_\sigma \gamma_\sigma \cdot \sigma$ where σ ranges over oriented k-simplices and each γ_σ is an element of the coefficient field $\mathbb{F}$.

DEFINITION 5.9. The **cap product** of an i-cochain ξ with a k-chain $\gamma = \sum_\sigma \gamma_\sigma \cdot \sigma$ is the new $(k-i)$-chain $\xi \frown \gamma$ defined by

$$\xi \frown \gamma = \sum_\sigma \gamma_\sigma \cdot \xi(\sigma_{\leq i}) \cdot \sigma_{\geq i}.$$

(For $i > k$ this sum is automatically zero).

The first thing to confirm about the cap product formula from the definition above is that the expression on the right side is a $(k-i)$-chain — each $\sigma_{\geq i}$ is a $(k-i)$-simplex obtained by deleting the first i vertices of the k-simplex σ, and the product $\gamma_\sigma \cdot \xi(\sigma_{\leq i})$ of two $\mathbb{F}$-elements clearly lies in $\mathbb{F}$. By definition, the cap product gives us bilinear maps $\mathbf{C}^i(K) \times \mathbf{C}_k(K) \to \mathbf{C}_{k-i}(K)$ for every pair of dimensions $i \leq k$. Since $\xi \frown \gamma$ is a chain, it has a boundary rather than a coboundary.

PROPOSITION 5.10. *For each ξ in $\mathbf{C}^i(K)$ and γ in $\mathbf{C}_k(K)$, we have*

$$\partial^K_{k-i}(\xi \frown \gamma) = (-1)^i \cdot [(\xi \frown \partial^K_k(\gamma)) - (\partial^i_K(\eta) \frown \gamma)]$$

The above result follows from a calculation which has a very similar structure to the one which we used when proving Proposition 5.5. This has been assigned as an exercise, unlike the the following corollary.

PROPOSITION 5.11. *For each simplicial complex K and dimensions $i \leq k$, the cap product $\frown$: $\mathbf{C}^i(K;\mathbb{F}) \times \mathbf{C}_k(K;\mathbb{F}) \to \mathbf{C}_{k-i}(K;\mathbb{F})$ induces a well-defined bilinear map of cohomology groups.*

PROOF. The desired result follows from the three claims described below, each of which is proved using the boundary formula from Proposition 5.10.

1. cocycle $\frown$ cycle is a cycle: if $\partial^i_K(\xi) = 0$ and $\partial^K_k(\gamma) = 0$, then we get

$$\partial^K_{k-i}(\xi \frown \gamma) = (-1)^i \cdot [(\xi \frown 0) - (0 \frown \gamma)] = 0,$$

as desired.

2. cocycle $\frown$ boundary is a boundary: if $\partial^i_K(\xi) = 0$ and β is any $(k+1)$-chain, then by Proposition 5.10, we have

$$\pm(\xi \frown \partial^K_{k+1}(\beta)) = \partial^K_k(\xi \frown \beta) \mp (\partial^i_K(\eta) \frown \beta);$$

now the second term on the right side vanishes because η is a cocycle. Thus, up to a choice of sign, $\xi \frown \partial^K_{k+1}(\beta)$ equals $\partial^K_k(\xi \frown \beta)$ which is evidently a boundary.

3. coboundary $\frown$ cycle is a boundary: this is very similar to the previous claim, and has therefore been assigned as an exercise. ☐

As with the cup product, it is standard to use the same notation for the cap product map on (co)homology groups as on (co)chain groups:

$$\frown : \mathbf{H}^i(K;\mathbb{F}) \times \mathbf{H}_k(K;\mathbb{F}) \to \mathbf{H}_{k-i}(K;\mathbb{F}).$$

The geometry of the cap product is all about *intersections*. If we choose a meridinal 1-cycle γ and an equatorial 1-cocycle ξ on a torus as drawn below, then there will necessarily be at least one edge with a nonzero coefficient in γ that is sent to a nonzero element of $\mathbb{F}$ by ξ; and the 0-chain $\xi \frown \gamma$ will have a nonzero coefficient on one of the two vertices lying in the boudnary of that edge:

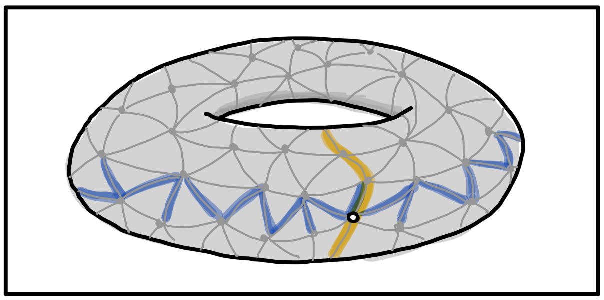

The power of the algebraic formulation of the cap product lies in the fact that the cycle $\xi \frown \gamma$ is well-defined on the level of homology even when ξ and γ are perturbed within their respective (co)homology classes.

5.5 Poincaré Duality

The cap product becomes extremely potent when applied to the study of *manifolds*. Throughout this section, we fix the following **assumption:**

> M is a simplicial complex whose geometric realization $|M|$ is a compact and connected n-dimensional manifold.

(The compactness requirement is overkill since we require simplicial complexes to be finite). The fact that every point on an n-manifold admits a local neighborhood homeomorphic to $\mathbb{R}^n$ forces every $(n-1)$-dimensional simplex of M to lie in the boundary of exactly two n-simplices.

Definition 5.12. We say that M is **orientable** over the field $\mathbb{F}$ if there exists a function

$$\omega : \{n - \text{simplices of } M\} \to \{\pm 1\}$$

assigning $\{\pm 1\} \subset \mathbb{F}$ to each top-dimensional simplex so that the chain $[M] = [M]_\omega$ given by

$$[M] = \sum_{\dim \sigma = n} \omega(\sigma) \cdot \sigma$$

is an n-cycle in $\mathbf{C}_n(M;\mathbb{F})$.

(It should be noted that any M satisfying our assumption above is automatically orientable in this sense over $\mathbb{F} = \mathbb{Z}/2$.) There is an unfortunate historical conflation of terminology here between orientability as defined above and the orderings of vertices which played a role in Definition 3.3. At any rate, if such a map ω exists then $[M]$ is called the **fundamental class** of M, and it generates all of $\mathbf{H}_n(M;\mathbb{F})$ which must necessarily be a one-dimensional vector space.

THEOREM 5.13. **[Poincaré duality.]** *Assume that M is a simplicial complex whose geometric realization is compact, connected and orientable over $\mathbb{F}$. For each i in $\{0, 1, \ldots, n\}$, the linear map*

$$\boldsymbol{D}_i : \mathbf{H}^i(M;\mathbb{F}) \to \mathbf{H}_{n-i}(M;\mathbb{F})$$

given by $\boldsymbol{D}_i(\xi) = \xi \frown [M]$ is an isomorphism of $\mathbb{F}$-vector spaces.

It is quite challenging to prove this result entirely within the realm of simplicial complexes, so we will not make any such attempts here. But it should be noted that Poincaré duality has strong consequences for even the simplest homotopy invariants of manifolds. Combining Theorem 5.13 with Proposition 5.3 produces the following suite of results for Euler characteristics and Betti numbers of manifolds (see Sections 1 and 4 of Chapter 3).

COROLLARY 5.14. *Let M be a simplicial complex satisfying the hypotheses of Theorem 5.13. The following assertions hold.*

(1) *The Betti numbers $\beta_0(M), \beta_1(M), \ldots, \beta_n(M)$ are palindromic, i.e., $\beta_k = \beta_{n-k}$ for all k.*
(2) *If n is odd, then the Euler characteristic $\chi(M)$ is zero.*
(3) *If $n = 2i$ is even, then the Euler characteristic $\chi(M;\mathbb{F})$ is odd if and only if the middle Betti number $\beta_i(M)$ is odd.*

PROOF. For the first assertion, note that

$$\begin{aligned}
\beta_k(M) &= \dim \mathbf{H}_k(\mathbb{F}) && \text{by definition,}\\
&= \dim \mathbf{H}^{n-k}(\mathbb{F}) && \text{by Theorem 5.13,}\\
&= \dim \mathbf{H}_{n-k}(\mathbb{F}) && \text{by Proposition 5.3,}\\
&= \beta_{n-k}(M) && \text{again by definition.}
\end{aligned}$$

The second assertion now follows from the first one by using (from Exercise 3.3) the fact that the Euler characteristic is the alternating sum of the Betti numbers:

$$\chi(M) = \sum_{k=0}^{n} (-1)^k \cdot \beta_k(M).$$

If n is odd, then β_k and β_{n-k} will appear with opposite signs and hence cancel. The third assertion follows from the same alternating sum — but for even $n = 2i$ all the β_k appear twice (with the same signs) except for β_i, which only appears once. Thus, the expression $\chi(M) \pm \beta_i(M)$ is always an even number. □

5.6 BONUS: THE KÜNNETH FORMULA

Let K and L be simplicial complexes. We have already lamented (in Section 8 of Chapter 4) that the product of simplicial complexes does not canonically have the structure of a simplicial complex. Even so, it is possible to find a simplicial complex P whose geometric realization is homeomorphic to $|K| \times |L|$, so it makes sense to define $\mathbf{H}_k(K \times L;\mathbb{F})$ as the usual homology groups of any such P, and similarly for cohomology groups. One naturally wonders how these product (co)homology groups of P relate to the (co)homology groups of the factors K and L.

Variants of the following result are called **Künneth formulas**.

THEOREM 5.15. *Let K and L be simplicial complexes and $\mathbb{F}$ a field. For each dimension $k \geq 0$ there is an isomorphism*

$$\mathbf{H}^k(K \times L; \mathbb{F}) \simeq \bigoplus_{i=0}^{k} \mathbf{H}^i(K; \mathbb{F}) \times \mathbf{H}^{k-i}(L; \mathbb{F}).$$

As a consequence of the Künneth formula and Proposition 5.3, one can compute Betti numbers of simplicial products via

$$\beta_k(K \times L) = \sum_{i=0}^{k} \beta_i(K) \cdot \beta_{k-i}(L).$$

There are several ways of proving Theorem 5.15; one strategy makes essential use of the cup product. Given a simplicial complex P whose realization is $|K| \times |L|$, assume that we have managed to simplicially approximate the natural projection maps from $K \times L$ to K and L, i.e.,

$$K \xleftarrow{\;\;f\;\;} P \xrightarrow{\;\;g\;\;} L$$

The goal now becomes to produce k-cochains in P from pairs of the form (ξ, η) where ξ is an i-cochain in K while η is a $(k-i)$-cochain in L. And the map which accomplishes this task is

$$(\xi, \eta) \mapsto \mathbf{C}^i f(\xi) \smile \mathbf{C}^{k-i} g(\eta).$$

EXERCISES

EXERCISE 5.1. Given a simplicial map $f : K \to L$, define the associated cochain maps $\mathbf{C}^k f : \mathbf{C}^k(L; \mathbb{F}) \to \mathbf{C}^k(K; \mathbb{F})$ and show that they commute with the coboundary operators (i.e., state and prove a cohomological version of Proposition 4.5).

EXERCISE 5.2. State a version of Definition 4.18 (short exact sequences) and Lemma 4.19 (the Snake lemma) that works for cochain complexes and cohomology.

EXERCISE 5.3. Show that the cup product is associative, i.e., for cochains ξ, η and ζ of a simplicial complex K, prove that

$$(\xi \smile \eta) \smile \zeta = \xi \smile (\eta \smile \zeta)$$

[Hint: by linearity, it suffices to evaluate both sides on a single simplex σ.]

EXERCISE 5.4. Let $f : K \to L$ be a simplicial map and consider a pair of cochains ξ in $\mathbf{C}^k(L)$ and η in $\mathbf{C}^\ell(L)$. Prove that $\mathbf{C}^k f(\xi) \smile \mathbf{C}^\ell f(\eta) = \mathbf{C}^{k+\ell} f(\xi \smile \eta)$. [Thus, we have $\mathbf{H}^k f(\xi) \smile \mathbf{H}^\ell f(\eta) = \mathbf{H}^{k+\ell}(\xi \smile \eta)$ whenever ξ and η lie in $\mathbf{H}^k(L)$ and $\mathbf{H}^\ell(L)$ respectively.]

EXERCISE 5.5. Prove Proposition 5.10.

EXERCISE 5.6. Prove the third claim of Proposition 5.11.

EXERCISE 5.7. Let $f : K \to L$ be a simplicial map. There is a diagram of $\mathbb{F}$-vector spaces, a part of which is shown below:

$$\begin{array}{ccccc} \mathbf{H}^i(K) & \times & \mathbf{H}_j(K) & \stackrel{\frown}{\longrightarrow} & \mathbf{H}_{j-i}(K) \\ \\ \mathbf{H}^i(L) & \times & \mathbf{H}_j(L) & \stackrel{\frown}{\longrightarrow} & \mathbf{H}_{j-i}(L) \end{array}$$

(1) draw three vertical arrows representing maps induced by f which connect the top row to the bottom row. What are the natural candidates for these maps?
(2) formulate an identity relating cap products and these three induced maps. You do not have to prove that this identity holds (but it is a good exercise to meditate on).

EXERCISE 5.8. Use the Künneth formula (Theorem 5.15) to find an expression for the k-th Betti number of the n-dimensional torus T^n obtained by taking the n-fold product of the hollow simplex $\partial\Delta(2)$.

6. Persistence

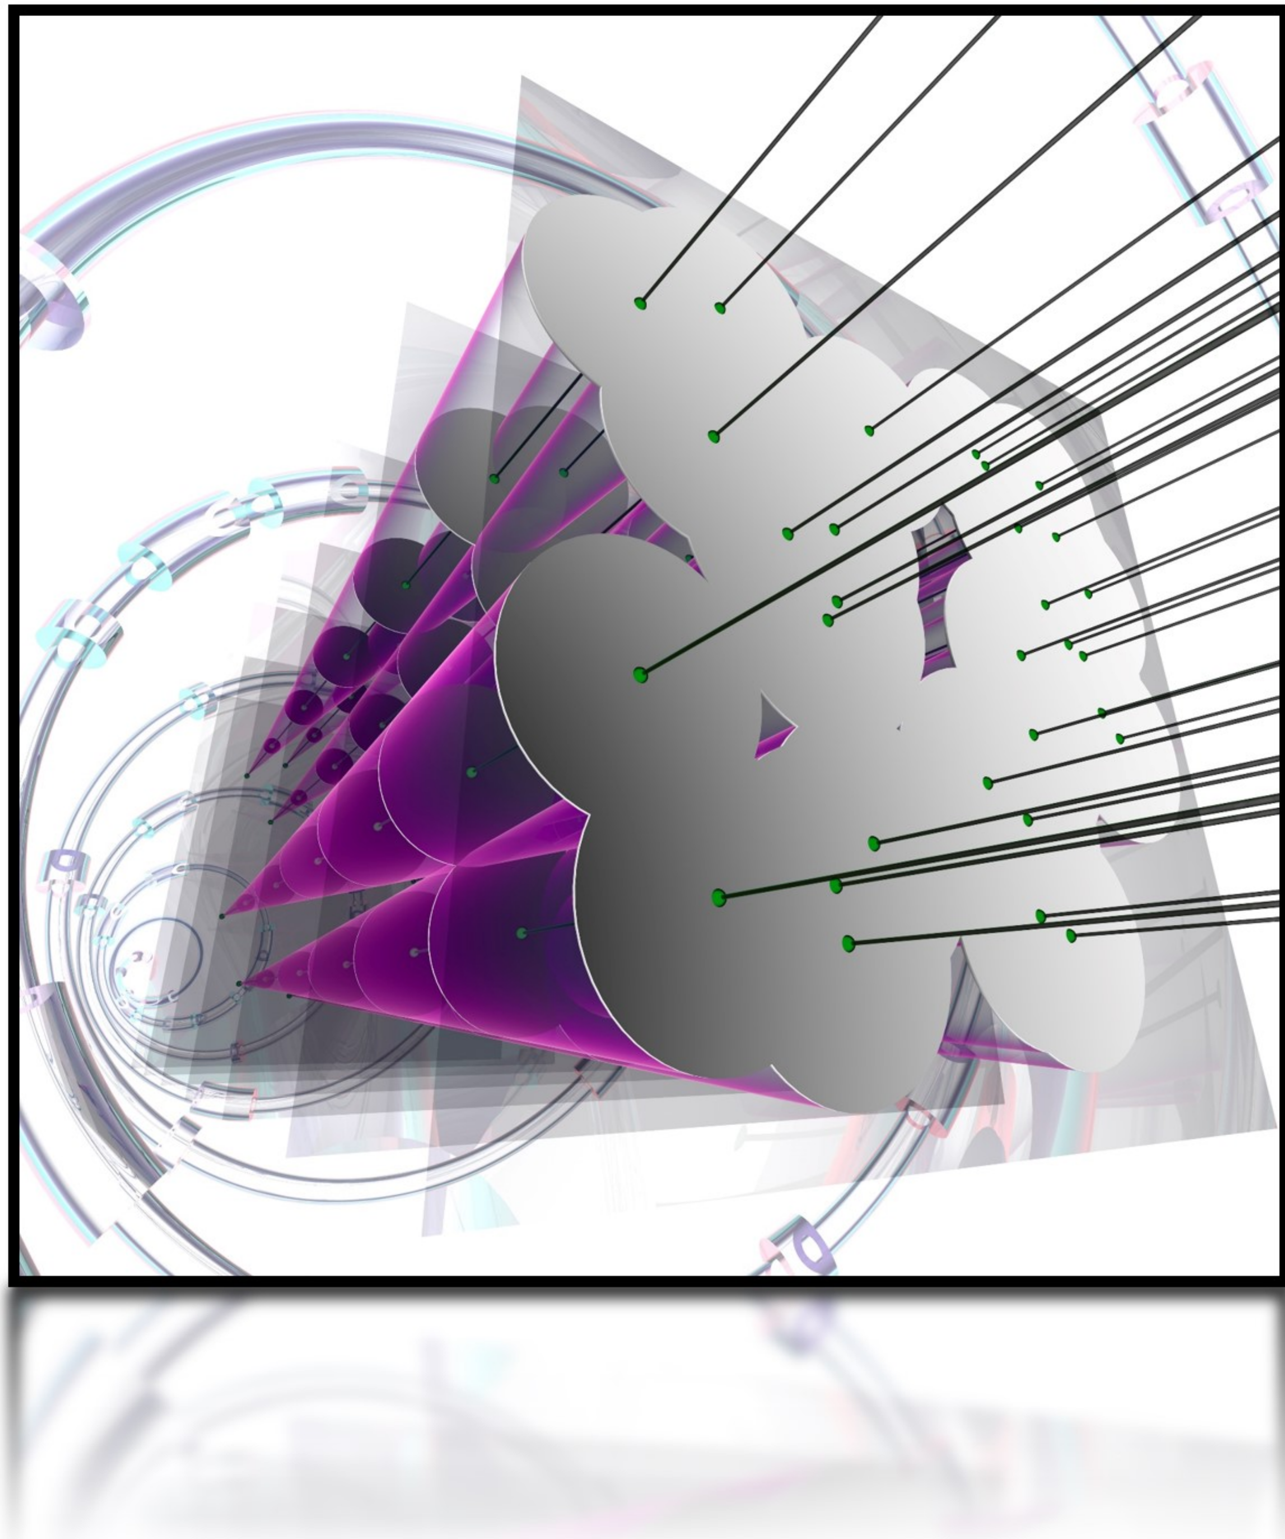

6.1 Persistent Homology

Consider a filtration $F_\bullet K$ of a simplicial complex K (as in Definition 1.6):

$$F_0K \subset F_1K \subset \cdots \subset F_{n-1}K \subset F_nK = K,$$

and denote the inclusion simplicial maps by $g_i : F_iK \hookrightarrow F_{i+1}K$. Here is one such filtration:

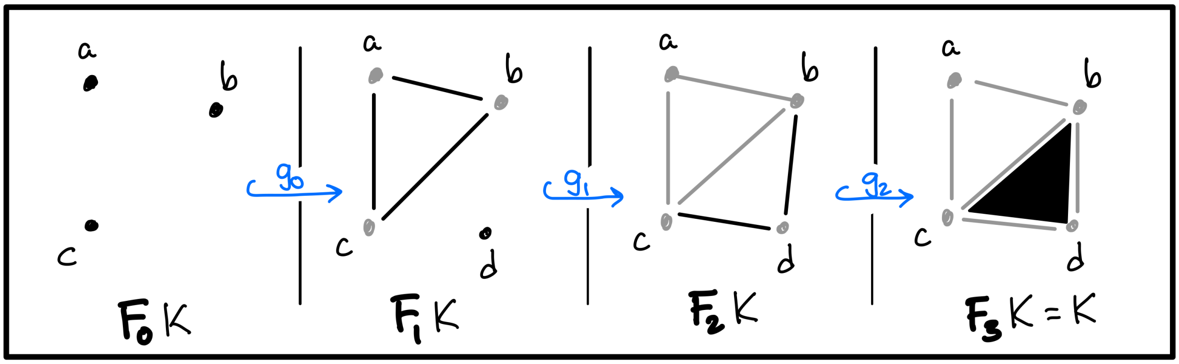

There are induced linear maps on homology $\mathbf{H}_kg_i : \mathbf{H}_k(F_iK) \to \mathbf{H}_k(F_{i+1}K)$ in every dimension $k \geq 0$ (see Sections 2 and 3 of Chapter 4). For a fixed k, these linear maps fit together into a sequence of vector spaces:

$$\mathbf{H}_k(F_0K) \xrightarrow{\mathbf{H}_kg_0} \mathbf{H}_k(F_1K) \xrightarrow{\mathbf{H}_kg_1} \cdots \xrightarrow{\mathbf{H}_kg_{n-2}} \mathbf{H}_k(F_{n-1}K) \xrightarrow{\mathbf{H}_kg_{n-1}} \mathbf{H}_k(F_nK).$$

There are several other induced maps on homology hiding in plain sight — for instance, we have said nothing about the inclusion $g_1 \circ g_0 : F_0K \hookrightarrow F_2K$. Fortunately for us, homology is functorial (see Proposition 4.8); so the missing map $\mathbf{H}_k(g_1 \circ g_0)$ is easily reconstructed by composing the available maps $\mathbf{H}_kg_1$ and $\mathbf{H}_kg_0$.

More generally, for any pair $i < j$ of filtration indices in $\{0, \ldots, n\}$, the map induced on homology by the inclusion $g_{i\to j} : F_iK \hookrightarrow F_jK$ is the composite $\mathbf{H}_k(F_iK) \to \mathbf{H}_k(F_jK)$ in our diagram of vector spaces, i.e.,

$$\mathbf{H}_kg_{i\to j} = \mathbf{H}_kg_{j-1} \circ \mathbf{H}_kg_{j-1} \circ \cdots \circ \mathbf{H}_kg_{i+1} \circ \mathbf{H}_kg_i.$$

Such maps $\mathbf{H}_kg_{i\to j}$ contain crucial information which allows us to coherently connect the k-th homology groups of *all* the subcomplexes which appear in the filtration $F_\bullet$ of K. The key point is that in order to extract this information, we must study sequences of vector spaces; thus, we are inexorably led to the following definition.

Definition 6.1. An $\mathbb{N}$-indexed **persistence module** over $\mathbb{F}$ is a sequence $(V_\bullet, a_\bullet)$ of $\mathbb{F}$-vector spaces V_k and linear maps a_k defined for $k \geq 0$ which fit into a diagram

$$V_0 \xrightarrow{a_0} V_1 \xrightarrow{a_1} V_2 \xrightarrow{a_2} \cdots \xrightarrow{a_{k-1}} V_k \xrightarrow{a_k} V_{k+1} \xrightarrow{a_{k+1}} \cdots$$

The maps $a_\bullet$ are not required to satisfy $a_k \circ a_{k-1} = 0$, so persistence modules need not be cochain complexes (compare Definition 5.1); conversely, every cochain complex is automatically a persistence module. In any event, for every pair $i \leq j$ in $\mathbb{N}$ we will write the composite map $a_{j-1} \circ a_{j-2} \circ \cdots \circ a_i$ via the shorthand $a_{i\to j} : V_i \to V_j$, with the implicit understanding that $a_{i\to i}$ is just the identity map on V_i.

Remark 6.2. We say that a persistence module $(V_\bullet, a_\bullet)$ is of **finite type** if $\dim V_i < \infty$ for all $i \geq 0$ and if the maps $a_i : V_i \to V_{i+1}$ are isomorphisms for all $i \gg 0$. Both these conditions are satisfied by persistence modules obtained from homology groups of filtered simplicial complexes.

We now turn to the objects of interest.

DEFINITION 6.3. For each pair $i \leq j$ of integers, the associated **persistent homology group** of a persistence module $(V_\bullet, a_\bullet)$ is the subspace of V_j given by

$$\mathbf{PH}_{i \to j}(V_\bullet, a_\bullet) = \operatorname{img}(a_{i \to j}).$$

It is not too difficult to check that $\mathbf{PH}_{i \to j}(V_\bullet, a_\bullet)$ is a subset of $\mathbf{PH}_{i' \to j}(V_\bullet, a_\bullet)$ whenever $i' \geq i$. We say that a vector v in V_i is **born** at filtration index i if v does not lie in img a_{i-1}; similarly, v is said to **die** at filtration index $j \geq i$ whenever j is the *smallest* number satisfying $a_{i \to j}(v) = 0$; by convention, the death index of v equals $+\infty$ no such j exists, i.e., if $a_{i \to j}(v)$ is nonzero for all $j \geq i$. The **persistence** of v is defined to be death minus birth, i.e., $(j - i)$.

REMARK 6.4. In the special case where our persistence module arises from taking the k-th homology groups of a filtered simplicial complex as described above, we will denote the persistent homology groups as $\mathbf{PH}_k g_{i \to j}(\boldsymbol{F}_\bullet K)$ for all $i \leq j$. The group $\mathbf{PH}_k g_{i \to j}(\boldsymbol{F}_\bullet K)$ consists of precisely those homology classes in $\mathbf{H}_k(\boldsymbol{F}_i K)$ which continue to generate nontrivial homology in the larger complex $\boldsymbol{F}_j K$ — geometrically, these are precisely those (equivalence classes of) k-cycles in $\boldsymbol{F}_i K$ which do not become k-boundaries in $\boldsymbol{F}_j K$. Writing ∂_k^i for the k-th boundary operator of each simplicial complex $\boldsymbol{F}_i K$, we have

$$\mathbf{PH}_k g_{i \to j}(\boldsymbol{F}_\bullet K) = \mathbf{H}_k g_{i \to j}(\ker \partial_k^i) / [\mathbf{H}_k g_{i \to j}(\ker \partial_k^i) \cap \operatorname{img} \partial_{k+1}^j].$$

And in particular, $\mathbf{PH}_k g_{i \to i}(\boldsymbol{F}_\bullet K)$ is just the k-th homology group of $\boldsymbol{F}_i K$.

The study of persistence modules is greatly facilitated by two miracles — an inherently algebraic **structure theorem** and a viscerally geometric **stability theorem**. The first of these allows us to represent every persistence module using the combinatorial data called its *barcode*. And the stability theorem asserts that the assignment of barcodes to modules is an isometry under certain natural metrics. We will describe the structure theorem in the next section

6.2 BARCODES

The quest to understand persistent homology groups begins, like many good quests, with the establishment of a categorical framework.

DEFINITION 6.5. A **morphism** between persistence modules $(V_\bullet, a_\bullet)$ and $(W_\bullet, b_\bullet)$ is a family of linear maps $\phi_k : V_k \to W_k$ which satisfy

$$b_i \circ \phi_i = \phi_{i+1} \circ a_i$$

for every $i \geq 0$

This definition amounts to requiring the commutativity of all squares in the following diagram of vector spaces:

$$\begin{tikzcd}
V_0 \arrow[r, "a_0"] \arrow[d, "\phi_0"'] & V_1 \arrow[r, "a_1"] \arrow[d, "\phi_1"'] & \cdots \arrow[r, "a_{k-1}"] & V_k \arrow[r, "a_k"] \arrow[d, "\phi_k"'] & V_{k+1} \arrow[r, "a_{k+1}"] \arrow[d, "\phi_{k+1}"'] & \cdots \\
W_0 \arrow[r, "b_0"'] & W_1 \arrow[r, "b_1"'] & \cdots \arrow[r, "b_{k-1}"'] & W_k \arrow[r, "b_k"'] & W_{k+1} \arrow[r, "b_{k+1}"'] & \cdots
\end{tikzcd}$$

The pair (persistence modules, their morphisms) forms a category in the sense of Definition 4.1. We call $\phi_\bullet : (V_\bullet, a_\bullet) \to (W_\bullet, b_\bullet)$ an *isomorphism* if every ϕ_i is an invertible linear map of vector spaces in the usual sense. If such an isomorphism exists, we write $(V_\bullet, a_\bullet) \simeq (W_\bullet, b_\bullet)$.

DEFINITION 6.6. The **direct sum** of two persistence modules $(V_\bullet, a_\bullet)$ and $(W_\bullet, b_\bullet)$ is a new persistence module $(V_\bullet \oplus W_\bullet, a_\bullet \oplus b_\bullet)$ defined as follows: its k-th vector space is the direct sum $V_i \oplus W_i$, while the linear map $a_i \oplus b_i$ has matrix representation $\begin{bmatrix} a_i & 0 \\ 0 & b_i \end{bmatrix}$.

Persistent homology groups of direct sums are direct sums of persistent homology groups (see Exercise 6.1 of this Chapter for a precise statement.) We say that a persistence module $(I_\bullet, c_\bullet)$ is **indecomposable** if it does not admit any interesting direct sum decompositions — in other words, anytime we have an isomorphism

$$(I_\bullet, c_\bullet) \simeq (V_\bullet, a_\bullet) \oplus (W_\bullet, b_\bullet),$$

of persistence modules, one of the factors on the right side will be isomorphic to $(I_\bullet, c_\bullet)$, while the other one will be zero everywhere. The following result highlights a particularly important class of indecomposable persistent modules.

PROPOSITION 6.7. *Let $(I_\bullet, c_\bullet)$ be a nonzero $\mathbb{N}$-indexed persistence module over a field $\mathbb{F}$. Assume that there exist indices $i \leq j$ with i in $\mathbb{N}$ and j in $\mathbb{N} \cup \{\infty\}$ so that*

$$\dim I_p = \begin{cases} 1 & i \leq p \leq j \\ 0 & \textit{otherwise} \end{cases}, \quad \textit{and} \quad \operatorname{rank}\,(c_p : I_p \to I_{p+1}) = \begin{cases} 1 & i \leq p < j \\ 0 & \textit{otherwise} \end{cases}.$$

Then, $(I_\bullet, c_\bullet)$ is indecomposable.

PROOF. Consider any direct sum decomposition $(I_\bullet, c_\bullet) \simeq (V_\bullet, a_\bullet) \oplus (W_\bullet, b_\bullet)$. For each p in $\{i, i+1, \ldots, j\}$ we have $\dim V_p + \dim W_p = \dim I_p = 1$; let's assume without loss of generality that $\dim V_i = 1$ and $\dim W_i = 0$. This forces the map b_i to be zero, and by Definition 6.5 we now have a commutative diagram which looks like:

$$\begin{tikzcd} \mathbb{F} \arrow[r, "\simeq"] \arrow[d, "\simeq"'] & \mathbb{F} \arrow[d, "\simeq"] \\ \mathbb{F} \oplus 0 \arrow[r, "\begin{bmatrix} a_i & 0 \\ 0 & 0 \end{bmatrix}"'] & V_{i+1} \oplus W_{i+1} \end{tikzcd}$$

with all arrows labelled $\simeq$ being vector space isomorphisms. It follows that a_i has rank one, $\dim V_{i+1} = 1$, and $\dim W_{i+1} = 0$. Continuing onwards by induction on i, we see that $(V_\bullet, a_\bullet)$ is isomorphic to $(I_\bullet, c_\bullet)$ while $(W_\bullet, b_\bullet)$ is trivial; thus, $(I_\bullet, c_\bullet)$ is indecomposable as desired. ☐

Up to isomorphism, every indecomposable module of the form described in the proposition above is completely characterized by knowledge of the pair of integers $i \leq j$ (allowing for the fact that j might equal ∞).

DEFINITION 6.8. For each pair $0 \leq i \leq j \leq \infty$ (with $i \neq \infty$), the $\mathbb{N}$-indexed **interval module** $(I_\bullet^{i,j}, c_\bullet^{i,j})$ over $\mathbb{F}$ is given by

$$I_p^{i,j} = \begin{cases} \mathbb{F} & i \leq p \leq j \\ 0 & \text{otherwise} \end{cases}, \quad \text{and} \quad c_p^{i,j} = \begin{cases} \mathrm{id}_{\mathbb{F}} & i \leq p < j \\ 0 & \text{otherwise} \end{cases}.$$

(Here $\mathrm{id}_{\mathbb{F}}$ denotes the identity map $\mathbb{F} \to \mathbb{F}$).

The first miracle of persistent homology is the following result, which allows us to uniquely express *any* $\mathbb{N}$-indexed persistence module of finite type as a direct sum of finitely many interval modules. Please do not panic (yet) if various terms in the proof appear intimidating — clarifying remarks and concrete computations will follow.

THEOREM 6.9. **[Structure Theorem]** *For any finite type $\mathbb{N}$-indexed persistence module $(V_\bullet, a_\bullet)$ over $\mathbb{F}$, there exists a set $\mathbf{Bar}(V_\bullet, a_\bullet)$ of integer pairs $0 \leq i \leq j \leq \infty$ (with $i \neq \infty$) and a function $\mu : \mathbf{Bar}(V_\bullet, a_\bullet) \to \mathbb{N}_{>0}$ to the nonzero natural numbers with the following property: there is a direct sum decomposition*

$$(V_\bullet, a_\bullet) \simeq \bigoplus_{[i,j]} (I_\bullet^{i,j}, c_\bullet^{i,j})^{\mu(i,j)}.$$

Here the indices $[i,j]$ range over elements of $\mathbf{Bar}(V_\bullet, a_\bullet)$. Moreover, this direct sum decomposition is unique (up to isomorphism of persistence modules).

PROOF. Since $(V_\bullet, a_\bullet)$ is of finite type, there is some $n \geq 0$ so that every $a_i : V_i \to V_{i+1}$ is an isomorphism for $i > n$. Consider the vector space $V = \bigoplus_0^n V_i$ and the linear map $t : V \to V$ sending each vector $v = (v_0, v_1, \ldots, v_n)$ to the shifted vector

$$t(v) = \big(0, a_0(v_0), a_1(v_1), \ldots, a_{n-1}(v_{n-1})\big).$$

This gives V the structure of a finitely generated $\mathbb{F}[t]$-module where $\mathbb{F}[t]$ is the polynomial ring over $\mathbb{F}$ in a single variable t. Since $\mathbb{F}[t]$ is a *principal ideal domain* whenever $\mathbb{F}$ is a field, every finitely generated $\mathbb{F}[t]$-module decomposes uniquely as a direct sum into two parts

$$V = F \oplus T,$$

where F is called *free* while T is *torsion*. Moreover, F is a direct sum of $\mathbb{F}[t]$-modules of the form $t^i \cdot \mathbb{F}[t]$ for some $i \geq 0$; each such free summand is isomorphic to an interval module of the form $(I_\bullet^{i,\infty}, c_\bullet^{i,\infty})$. Similarly, the torsion part T is a direct sum of modules of the form $t^i \cdot \mathbb{F}[t]/(t^j)$, i.e., a free module quotient by an ideal $(t^j) \lhd \mathbb{F}[t]$ with $0 \leq i < j$; each such summand is isomorphic to the interval module $(I_\bullet^{i,j}, c_\bullet^{i,j})$. These (free and torsion) interval modules might occur in the decomposition with any multiplicities ≥ 1, which are catalogued by the function μ. $\square$

While quite miraculous in its outcomes, this argument has two serious drawbacks arising from the fact that it invokes the classification of finitely generated $\mathbb{F}[t]$-modules. First, this proof strategy will not survive if we attempt something similar with $\mathbb{Z}[t]$-modules or even $\mathbb{F}[t_1, t_2]$-modules. Second, the *deus ex machina* nature of appealing to this classification renders life somewhat difficult for those who seek to understand the decomposition of $(V_\bullet, a_\bullet)$ on a more concrete level. There is no remedy for the first problem, but we can offer some solace to those afflicted by the second malady. The next Section contains a very concrete algorithm for computing interval-decompositions in the case of maximal interest to us, i.e., where $(V_\bullet, a_\bullet)$ arises from the homology groups of a filtered simplicial complex.

DEFINITION 6.10. For each ($\mathbb{N}$-indexed, finite type) persistence module $(V_\bullet, a_\bullet)$ over $\mathbb{F}$, the collection $\mathbf{Bar}(V_\bullet, a_\bullet)$ of intervals $[i,j]$ and their multiplicities $\mu(i,j) \geq 1$ (whose existence and uniqueness is guaranteed by Theorem 6.9) is called the **barcode** of $(V_\bullet, a_\bullet)$.

The content of Theorem 6.9 is that every finite type persistence module is uniquely determined up to isomorphism by the combinatorial data consisting of intervals $[i,j]$ in $\mathbf{Bar}(V_\bullet, a_\bullet)$ and their multiplicities $\mu(i,j)$. For brevity, we will denote multiplicities as superscripts, so $[1,4]^3$ means that the bar $[1,3]$ occurs with multiplicity $\mu(1,4) = 3$ in a given barcode.

6.3 Algorithm (for Filtrations)

Let $\boldsymbol{F}_\bullet K$ be a filtered simplicial complex

$$\boldsymbol{F}_0 K \subset \boldsymbol{F}_1 K \subset \cdots \subset \boldsymbol{F}_{n-1} K \subset \boldsymbol{F}_n K = K,$$

and for each simplex σ of K let $b(\sigma)$ denote the smallest index i in $\{0, \ldots, n\}$ for which σ lies in $\boldsymbol{F}_i K$. Since each $\boldsymbol{F}_i K$ forms a subcomplex of K, it follows that b is an order preserving map on the simplices of K, i.e., $\sigma \leq \tau$ in K implies $b(\sigma) \leq b(\tau)$. In more prosaic terms, a simplex can only enter the filtration at index i if all of its faces are already present. Writing $g_{i \to j}$ for the inclusion map $\boldsymbol{F}_i K \hookrightarrow \boldsymbol{F}_j K$ for $i \leq j$, here we will describe an efficient algorithm which computes *all* the persistent homology groups $\mathbf{PH}_k g_{i \to j}(\boldsymbol{F}_\bullet K)$ at once by exploiting Theorem 6.9.

0. Setup: Order the simplices of K as $\{\sigma_1, \sigma_2, \ldots, \sigma_N\}$ so that σ precedes τ in this ordering when either on of the following conditions holds:

- we have $b(\sigma) \leq b(\tau)$, or
- we have $b(\sigma) = b(\tau)$ and σ is a face of τ in K.

Aside from these two constraints, the simplices of K may be ordered arbitrarily.

1. Input: The input to the algorithm is an $N \times N$ matrix $\boldsymbol{D}$ described as follows. For each pair (p, q) in $\{1, \ldots, N\}^2$, the entry of $\boldsymbol{D}$ in the p-th row and q-th column is given by

$$\boldsymbol{D}_{pq} = \begin{cases} \pm 1 & \text{if } \sigma_p \leq \sigma_q \text{ with } \dim \sigma_q - \dim \sigma_p = 1 \\ 0 & \text{otherwise} \end{cases}.$$

Here the sign ± 1 depends on an ordering of K's vertices; in particular, this is the same sign as the one used in the algebraic boundary operator of Definition 3.4. We will indicate the q-th column of $\boldsymbol{D}$ by $\mathrm{col}(q)$ and write $\mathrm{low}(q)$ to indicate the largest p satisfying $\boldsymbol{D}_{pq} \neq 0$, with the explicit understanding that $\mathrm{low}(q) = 0$ whenever the $\mathrm{col}(q)$ is entirely zero.

2. Procedure: The entire routine can be described with only six lines of pseudocode.

01	**For** $q = 1$ to N
02	**Set** $p = \mathrm{low}(q)$
03	**While** some $r < q$ satisfies $\mathrm{low}(r) = p \neq 0$
04	**Add** $\left(-\boldsymbol{D}_{pq}/\boldsymbol{D}_{pr}\right) \cdot \mathrm{col}(r)$ to $\mathrm{col}(q)$
05	**End** While
06	**End** For

3. Output: This procedure modifies the matrix $\boldsymbol{D}$ to produce a new matrix $\boldsymbol{D}'$ — this matrix $\boldsymbol{D}'$ is related to $\boldsymbol{D}$ by a change of basis since we only used column operations. In particular, lines 03-05 attempt to incrementally zero out the q-th column of $\boldsymbol{D}$ by adding preceding columns whose lowest nonzero entry coincides with that of $\mathrm{col}(q)$. Thus, when the algorithm terminates, the p-th row of $\boldsymbol{D}'$ can be the lowest nonzero entry $\mathrm{low}(q)$ of at most one column q — if there is such a q, then the entry $\boldsymbol{D}'_{pq}$ is said to form a *pivot* in the output matrix $\boldsymbol{D}'$.

4. The Barcodes: For each $k \geq 0$, let's write $\mathbf{Bar}_k(\boldsymbol{F}_\bullet K)$ to indicate the barcode of the persistence module obtained by taking the k-th homology groups of $\boldsymbol{F}_\bullet K$. We can read off such barcodes (and hence determine these persistence modules thanks to Theorem 6.9) by traversing the columns of $\boldsymbol{D}'$ and applying this handy flow-chart:

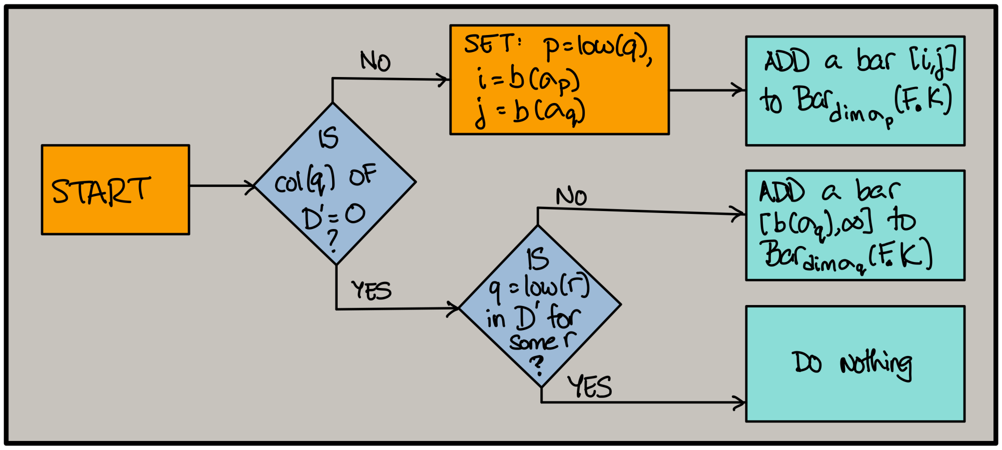

5. Example: When this algorithm is run on the filtration depicted in Section 1 (reproduced below), it will output the barcode $\{[0,\infty],[0,1]^2,[1,2]\}$ for 0-dimensional persistent homology and the barcode $\{[1,\infty],[2,3]\}$ for 1-dimensional persistent homology, perfectly capturing the evolution of connected components and loops at various stages in the filtration:

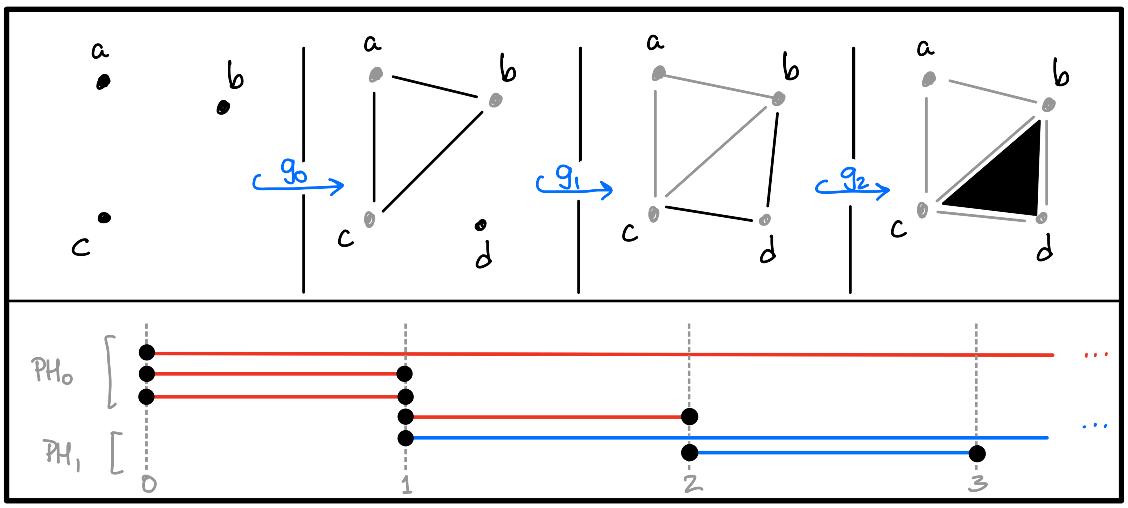

The starting point of the algorithm for this filtration is the following matrix $\boldsymbol{D}$ as described in the **Input** step above — all unlabelled entries are zero:

$$
\mathcal{D} = \begin{array}{c|ccc|cccc|cc|c}
 & & b=0 & & & & b=1 & & b=2 & & b=3 \\
 & a & b & c & d & ab & ac & bc & bd & cd & bcd \\
\hline
a & & & & & -1 & -1 & & & & \\
b & & & & & 1 & & -1 & -1 & & \\
c & & & & & & 1 & 1 & & -1 & \\
d & & & & & & & & 1 & 1 & \\
ab & & & & & & & & & & \\
ac & & & & & & & & & & \\
bc & & & & & & & & & & 1 \\
bd & & & & & & & & & & -1 \\
cd & & & & & & & & & & 1 \\
bcd & & & & & & & & & &
\end{array}
$$

No operations are performed on ab's column, so the 1 in that column (in b's row) serves as a pivot. This pivot will contribute one of the two $[0, 1]$ bars in the 0-dimensional barcode of this filtration. The first interesting column operation occurs when the 1 in ac's column is used to clear out the 1 in bc's column (both corresponding to c's row). This changes the lowest entry in bc's column to the -1 in b's row, and we then use our pivot 1 in ab's column to cancel this new lowest entry. This will completely clear out bc's column, and contribute the $[1, \infty]$ bar in to the 1-dimensional barcode.

REMARK 6.11. Even on the small example described above, it is difficult to carry out the entire algorithm by hand. Fortunately, there are several good software packages available for computing persistent homology of filtered simplicial complexes arising in practice. In particular, one can find many implementations of this algorithm which will compute barcodes of Vietoris-Rips filtrations built around finite metric spaces (see Definition 1.15).

6.4 INTERLEAVING DISTANCE

Having witnessed the algebraic miracle of Theorem 6.9, we now turn to the geometric miracle, which takes the form of a **stability** result. Roughly, the set of finite type persistence modules admits the structure of a metric space, as does the set of barcodes; and with respect to the two chosen metrics, the assignment of a barcode to a module is an isometry. Here we will describe the desired metric on persistence modules after suitably upgrading them (and their barcodes) to be indexed by real numbers rather than natural numbers.

DEFINITION 6.12. An $\mathbb{R}_+$**-indexed persistence module** over $\mathbb{F}$ is a pair $(V_\bullet, a_\bullet)$ consisting of an $\mathbb{F}$-vector space V_t for each real number $t \geq 0$ and a linear map $a_{s\leq t} : V_s \to V_t$ for each pair $s \leq t$ of non-negative real numbers; these maps must satisfy

(1) $a_{t\leq t}$ is the identity map on V_t for each $t \geq 0$, and
(2) $a_{s\leq t} \circ a_{r\leq s} = a_{r\leq t}$ for every triple $0 \leq r \leq s \leq t$ of real numbers.

Put more succintly, these new persistence modules are functors of the form $(\mathbb{R}_+, \leq) \to \mathbf{Vect}_{\mathbb{F}}$ (see Definition 4.2). Here $(\mathbb{R}_+, \leq)$ is the category whose objects are all non-negative real numbers, with a unique morphism $s \to t$ whenever $s \leq t$; and the codomain is the usual category of (vector spaces, linear maps) over $\mathbb{F}$.

These persistence modules are more general than the $\mathbb{N}$-indexed ones from Definition 6.1: we can always replace an $\mathbb{N}$-indexed $(V_\bullet, a_\bullet)$ by an equivalent $\mathbb{R}_+$-indexed $(V'_\bullet, a'_\bullet)$ by interpolation as follows. Writing $\lfloor t \rfloor$ for the largest integer smaller than each t in $\mathbb{R}_+$, define

$$V'_t = V_{\lfloor t \rfloor} \quad \text{and} \quad a'_{s\leq t} = a_{\lfloor s \rfloor \to \lfloor t \rfloor}. \tag{4}$$

Henceforth, by persistence module we will mean the $\mathbb{R}_+$-indexed version defined above. For numerous reasons, it will be extremely convenient to visualize these as a continuum of of vector spaces living along a semi-infinite line segment connected by linear maps going from left to right, like so:

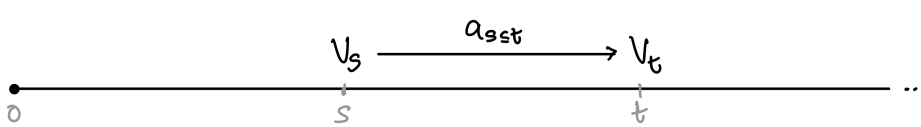

In order to guarantee barcodes for these new persistence modules a la Theorem 6.9, one must impose some finiteness constraints.

DEFINITION 6.13. A persistence module $(V_\bullet, a_\bullet)$ is called **tame** if two properties hold:

(1) the vector spaces V_t are finite-dimensional for all $t \geq 0$, and
(2) there are only finitely many $t \geq 0$, called *critical values*, for which the map $a_{t-\epsilon \leq t+\epsilon} : V_{t-\epsilon} \to V_{t+\epsilon}$ fails to be an isomorphism for arbitrarily small $\epsilon > 0$.

Tameness allows us to use Theorem 6.9 with impunity even with the $\mathbb{R}_+$-indexing — each tame persistence module $(V_\bullet, a_\bullet)$ can be reduced to a finite type $\mathbb{N}$-indexed persistence module $(V'_\bullet, a'_\bullet)$ as follows: let $0 \leq t_1 < t_2 < \cdots < t_n \leq \infty$ be the critical values of $(V_\bullet, a_\bullet)$ and set

$$V'_i = V_{t_i} \quad \text{and} \quad a'_i = a_{t_i \leq t_{i+1}}. \tag{5}$$

The barcode of $(V'_\bullet, a'_\bullet)$ can now be reinterpreted as the barcode of $(V_\bullet, a_\bullet)$ by sending each interval $[i, j]$ to the corresponding $[t_i, t_j]$. The *interval module* $(I_\bullet^{t_i,t_j}, c_\bullet^{t_i,t_j})$ supported on $[t_i, t_j]$ has the obvious definition:

$$I_t^{t_i,t_j} = \begin{cases} \mathbb{F} & t_i \leq t \leq t_j \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad c_{s\leq t}^{t_i,t_j} = \begin{cases} \mathrm{id}_{\mathbb{F}} & [s,t] \subset [t_i, t_j] \\ 0 & \text{otherwise} \end{cases}.$$

We have arrived at the following Corollary of Theorem 6.9; to fully appreciate its content, one should define (iso)morphisms and direct sums of tame persistence modules (as we did for their $\mathbb{N}$-indexed cousins).

COROLLARY 6.14. *For every tame persistence module $(V_\bullet, a_\bullet)$, there is a finite set $\mathbf{Bar}(V_\bullet, a_\bullet)$ of intervals of the form $[s,t] \subset \mathbb{R}_+$ (possibly with $t = \infty$) and a multiplicity $\mu : \mathbf{Bar}(V_\bullet, a_\bullet) \to \mathbb{N}_{>0}$ so that we have a unique direct sum decomposition into interval modules*

$$(V_\bullet, a_\bullet) \simeq \bigoplus_{[s,t]} (I_\bullet^{s,t}, c_\bullet^{s,t})^{\mu(s,t)},$$

with $[s,t]$ ranging over the intervals in $\mathbf{Bar}(V_\bullet, a_\bullet)$.

We now seek to measure distances between persistence modules. The following notion plays a central role.

DEFINITION 6.15. For each $\epsilon \geq 0$, an ϵ-**interleaving** between persistence modules $(V_\bullet, a_\bullet)$ and $(W_\bullet, b_\bullet)$ consists of two families of linear maps

$$\{\Phi_t : V_t \to W_{t+\epsilon} \mid t \geq 0\} \text{ and } \{\Psi_t : W_t \to V_{t+\epsilon} \mid t \geq 0\},$$

which satisfy four criteria. First, there are two *parallelogram relations*:

(1) for all $s \leq t$, we have $\Phi_t \circ a_{s\leq t} = b_{s+\epsilon \leq t+\epsilon} \circ \Phi_s$, and
(2) for all $s \leq t$, we have $\Psi_t \circ b_{s \leq t} = a_{s+\epsilon\leq t+\epsilon} \circ \Psi_s$.

And second, there are two *triangle relations*:

(1) for all t, we have $\Psi_{t+\epsilon} \circ \Phi_t = a_{t \leq t+2\epsilon}$, and
(2) for all t, we have $\Phi_{t+\epsilon} \circ \Psi_t = b_{t\leq t+2\epsilon}$.

These four criteria might appear opaque at a first reading; the best method of acquiring an intuitive grasp on interleavings is to draw the commutative diagrams implied by the parallelogram and triangle relations. This will require us to visualize both $V_\bullet$ and $W_\bullet$ along line segments as suggested before, so that the maps Φ_t and Ψ_t connect each point $t \geq 0$ on one of these lines to the point $t + \epsilon$ on the other. Here, for instance, is the commuting diagram which represents the first parallelogram relation:

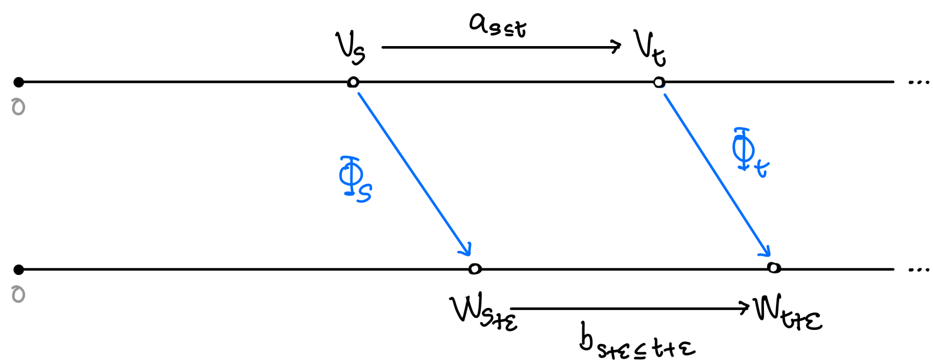

Of course, we have one such commuting diagram for *every* choice of $s \leq t$. Similarly, here is an illustration of the first triangle relation (there is one such commuting triangle for every t).

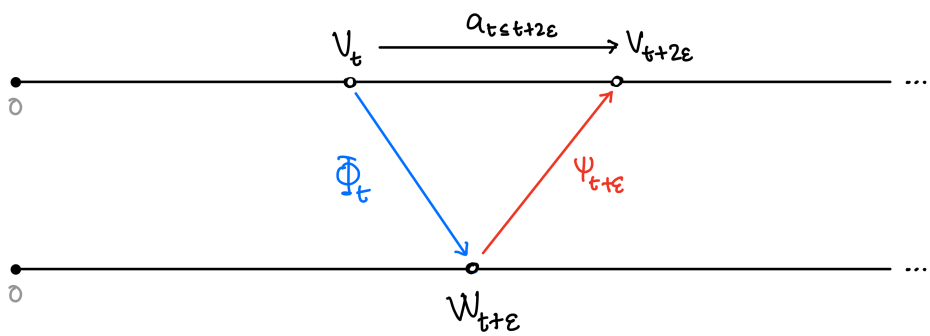

It might also be helpful to verify that 0-interleavings are isomorphisms of persistence modules — this is one of the Exercises. Finally, here is the promised metric on persistence modules.

DEFINITION 6.16. The **interleaving distance** $\boldsymbol{d}_{\mathrm{Int}}((V_\bullet, a_\bullet), (W_\bullet, b_\bullet))$ between persistence modules $(V_\bullet, a_\bullet)$ and $(W_\bullet, b_\bullet)$ is the infimum over all $\epsilon \geq 0$ for which there exists an ϵ-interleaving between them. If no such interleaving exists, then $\boldsymbol{d}_{\mathrm{Int}}(V_\bullet, W_\bullet) = \infty$.

6.5 THE STABILITY THEOREM

The barcodes $\mathbf{Bar}(V_\bullet, a_\bullet)$ whose existence and uniqueness is guaranteed by Corollary 6.14 for each tame persistence module $(V_\bullet, a_\bullet)$ are finite multi-sets of intervals $[s, t] \subset \mathbb{R}_+ \cup \infty$. Here by multi-set we simply mean that each interval $[s, t]$ might have several copies within the barcode, the precise number being given by the function $\mu(s, t)$. Our next goal is to impose a metric on the set of all such multi-sets of intervals.

DEFINITION 6.17. For $\epsilon \geq 0$, an **ϵ-matching** between two multi-sets B and B' of intervals is a bijection $\rho : B_0 \to B_0'$ between a pair of multi-subsets $B_0 \subset B$ and $B_0' \subset B'$ subject to the following constraints:

(1) Every $[s, t]$ in $(B - B_0) \cup (B' - B_0')$ has length $t - s \leq 2\epsilon$, and
(2) If $\rho[s, t] = [s', t']$ for some $[s, t]$ in B_0, then $|s - s'| \leq \epsilon \geq |t - t'|$.

Thus, if ρ is an ϵ-matching between multi-sets B and B', then it must pair all intervals of length exceeding 2ϵ of B with those of B'. And moreover, if ρ pairs $[s, t]$ with $[s', t']$, then we can obtain s' and t' by perturbing s and t respectively by no more than ϵ:

DEFINITION 6.18. The **bottleneck distance** $d_{\text{Bot}}(B, B')$ between multi-sets of intervals B and B' is the infimum over all $\epsilon \geq 0$ for which there exists an ϵ-matching between them.

Here is the geometric miracle of persistence modules.

THEOREM 6.19. **[Stability Theorem]** *For every pair* $(V_\bullet, a_\bullet)$ *and* $(W_\bullet, b_\bullet)$ *of tame persistence modules, we have*

$$d_{\text{Int}}((V_\bullet, a_\bullet), (W_\bullet, b_\bullet)) = d_{\text{Bot}}(\mathbf{Bar}(V_\bullet, a_\bullet), \mathbf{Bar}(W_\bullet, b_\bullet)).$$

Thus, the assignment of a barcode to a tame persistence module constitutes an isometry from the metric space of tame persistence modules (with interleaving distance) to the metric space of multi-sets of intervals (with bottleneck distance). All known proofs of the stability theorem are too technical to be included here[1]. The key advantage of the stability theorem is that it confers a certain geometric robustness to the following *topological data analysis* pipeline:

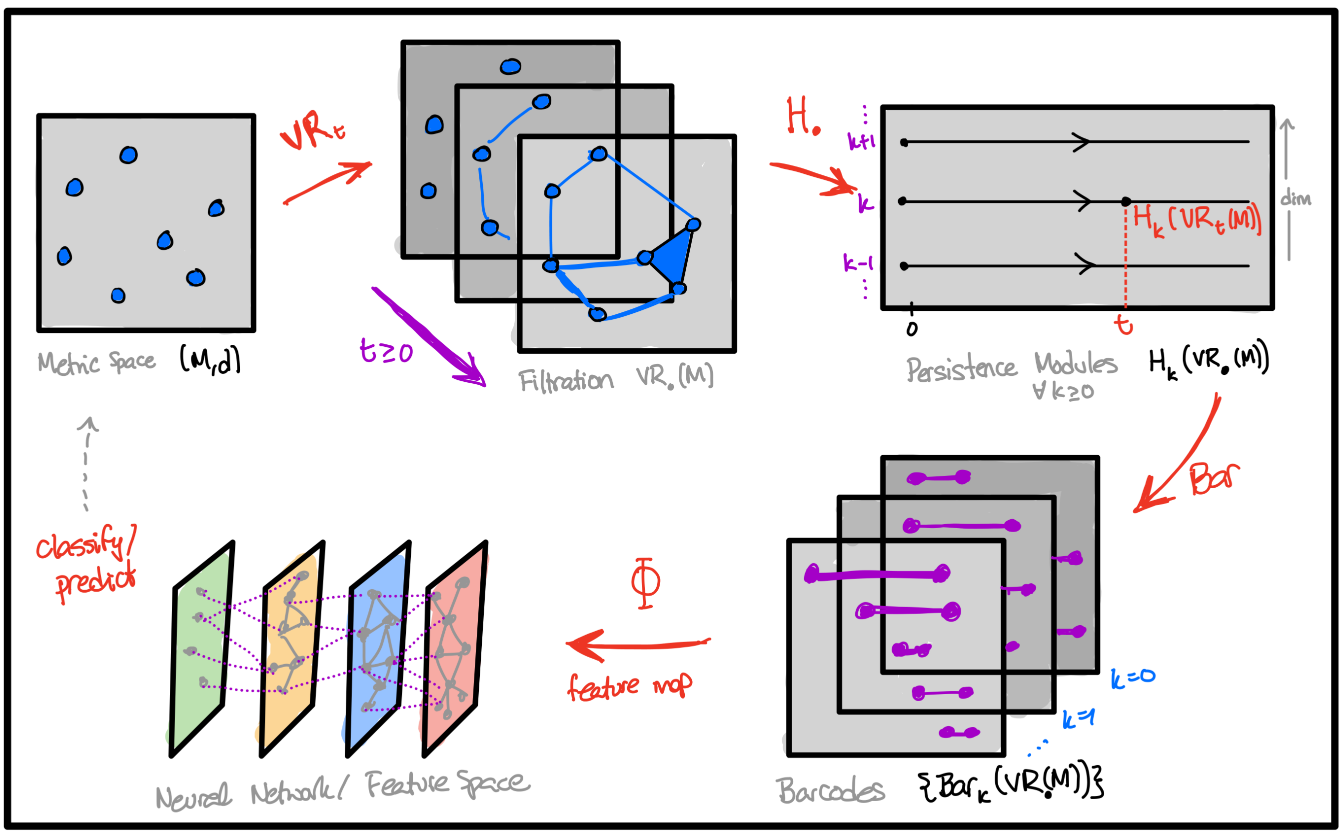

The first step describes the passage from a finite metric space to a filtered simplicial complex (as in Section 6 of Chapter 1). From there we compute persistent homology barcodes as described in Section 3 above. Since barcodes are combinatorial (rather than algebraic) objects, they can

[1]See Bauer and Lesnick's 2015 paper *Induced Matchings and the Algebraic Stability of Persistence Barcodes* for the most elementary proof known at present.

easily be vectorized and fed as input into neural networks or other statistical inference tools. The stability theorem enters the picture due to the following result.

PROPOSITION 6.20. *Let P and Q be two finite point-sets in $\mathbb{R}^n$ which are close in the following sense: there is some $\epsilon > 0$ so that*

(1) *there is a point of Q within distance ϵ of any point of P, and*
(2) *there is a point of P within distance ϵ of every point of Q.*

Then for each dimension $k \geq 0$, the k-th persistent homology modules of the Vietoris-Rips filtrations $\mathbf{VR}_\bullet(P)$ and $\mathbf{VR}_\bullet(Q)$ are 2ϵ-interleaved.

PROOF. Let $\alpha : P \to Q$ and $\beta : Q \to P$ be any pair of functions guaranteed by the ϵ-closeness of P and Q; thus, the Euclidean distance $\|p - \alpha(p)\|$ is no larger than ϵ for all p in P (and similarly for β). Now α induces simplicial maps $\{\alpha_t : \mathbf{VR}_t(P) \to \mathbf{VR}_{t+2\epsilon}(Q) \mid t \geq 0\}$ — to see why, note that if $\|p - p'\| \leq t$ then $\|\alpha(p) - \alpha(p')\| \leq t + 2\epsilon$ by the triangle inequality. Similarly, we get simplicial maps $\beta_t : \mathbf{VR}_t(Q) \to \mathbf{VR}_{t+2\epsilon}(P)$ for every $t \geq 0$. For each dimension $k \geq 0$, there are induced maps on homology $\mathbf{H}_k\alpha_t$ and $\mathbf{H}_k\beta_t$. We will now confirm that these induced maps $\mathbf{H}_k\alpha_t$ and $\mathbf{H}_k\beta_t$ satisfy the requirements of a 2ϵ-interleaving (Definition 6.15) between the persistence modules $\mathbf{PH}_k(\mathbf{VR}_\bullet P)$ and $\mathbf{PH}_k(\mathbf{VR}_\bullet Q)$.

1. Parallelogram Relations: For each $s \leq t$, let's denote the Vietoris-Rips inclusion maps as

$$i_{s\leq t} : \mathbf{VR}_s(P) \hookrightarrow \mathbf{VR}_t(P) \quad \text{and} \quad j_{s\leq t} : \mathbf{VR}_s(Q) \hookrightarrow \mathbf{VR}_t(Q).$$

By definition, we have $\alpha_t \circ i_{s\leq t} = j_{s+2\epsilon\leq t+2\epsilon} \circ \alpha_s$; now functoriality (i.e., Theorem 4.8) guarantees that the maps induced on k-th homology by α_s and α_t satisfy the parallelogram relation (see Definition 6.15).

$$\mathbf{H}_k\alpha_t \circ \mathbf{H}_k i_{s\leq t} = \mathbf{H}_k j_{s+2\epsilon\leq t+2\epsilon} \circ \mathbf{H}_k\alpha_s$$

An eerily similar argument establishes the parallelogram relation for $\mathbf{H}_k\beta_t$.

2. Triangle Relations: For each $t \geq 0$, note that the composite simplicial map

$$\beta_{t+2\epsilon} \circ \alpha_t : \mathbf{VR}_t(P) \to \mathbf{VR}_{t+4\epsilon}(P)$$

sends each vertex p to the vertex $p' = \beta \circ \alpha(p)$; by the triangle inequality we have $\|p - p'\| \leq 4\epsilon$. If $\sigma = (p_0, \ldots, p_m)$ is any m-simplex in $\mathbf{VR}_t(P)$, then the inclusion map $i_{t\leq t+4\epsilon}$ sends σ to σ, while the composite $\beta_{t+2\epsilon} \circ \alpha_t$ sends it to $\sigma' = (p'_0, \ldots, p'_m)$, with $p'_i = \beta \circ \alpha(p_i)$ for all i. It is easily confirmed that $\sigma \cup \sigma'$ is a simplex in $\mathbf{VR}_{t+4\epsilon}(P)$ by the triangle inequality. Thus, the simplicial maps $i_{t\leq t+4\epsilon}$ and $\beta_{t+2\epsilon} \circ \alpha_t$ are contiguous (in the sense of Corollary 2.9) and hence homotopic. By the homotopy invariance of homology (Theorem 4.24), their induced maps on homology coincide, and we obtain the desired triangle relation

$$\mathbf{H}_k\beta_{t+2\epsilon} \circ \mathbf{H}_k\alpha_t = \mathbf{H}_k i_{t\leq t+4\epsilon}.$$

A similar argument (with the roles of α and β interchanged) establishes the second triangle relation as well, and yields the desired result. □

As a consequence of the stability theorem, we see that for any $P, Q \subset \mathbb{R}^n$ satisfying the hypotheses of the above result, the k-th Vietoris-Rips persistent homology barcodes of P and Q must have the same number of sufficiently long bars, i.e., there is a bijection between bars of length $\geq 4\epsilon$ between the two barcodes in every homological dimension k. In this sense, the longer bars are stable to the sorts of perturbations which would replace P with Q. On the other hand, persistent homology is not stable to egregious outliers. In other words, if one obtains Q from P by adding just one point very far away from the existing points of P, then there is no relationship in general between the barcodes of P and those of Q.

EXERCISES

EXERCISE 6.1. Let $(V_\bullet, a_\bullet)$ and $(W_\bullet, b_\bullet)$ be $\mathbb{N}$-indexed persistence modules over a field $\mathbb{F}$. Show that for all $i \leq j$, there is an isomorphism

$$\mathbf{PH}_{i\to j}\big((V_\bullet, a_\bullet) \oplus (W_\bullet, b_\bullet)\big) \simeq \mathbf{PH}_{i\to j}(V_\bullet, a_\bullet) \oplus \mathbf{PH}_{i\to j}(W_\bullet, b_\bullet)$$

of persistent homology groups.

EXERCISE 6.2. Let $L \subset K$ be a two-step filtration of a simplicial complex K. Describe how to extract the dimension of the relative homology group $\mathbf{H}_k(K, L)$ for each $k \geq 0$ given the barcodes (with multiplicity) of this filtration.

EXERCISE 6.3. Let $\boldsymbol{F}_\bullet K$ be a filtration of a simplicial complex K. For each dimension $k \geq 0$ and filtration index i, describe how to compute the k-th Betti number of $\boldsymbol{F}_i K$ from the barcode $\mathbf{PH}_k(\boldsymbol{F}_\bullet K)$.

EXERCISE 6.4. Show that the interpolation of (4) produces an $\mathbb{R}_+$-indexed persistence module from an $\mathbb{N}$-indexed one.

EXERCISE 6.5. Describe a notion of morphisms which turn $\mathbb{R}_+$-indexed persistence modules into a category (if this is done correctly, the $\mathbb{N}$-indexed persistence modules will form a subcategory via (4)). What are the isomorphisms?

EXERCISE 6.6. Show that every $\mathbb{R}_+$-indexed interval module is tame.

EXERCISE 6.7. Show that sending a finite type $\mathbb{N}$-indexed persistence module $(V_\bullet, a_\bullet)$ to a tame $\mathbb{R}_+$-indexed one via (5), and then going back via (4), gives us $(V_\bullet, a_\bullet)$ back.

EXERCISE 6.8. Show that two ($\mathbb{R}_+$-indexed) persistence modules are isomorphic if and only if they admit a 0-interleaving.

EXERCISE 6.9. Draw commuting diagrams which represent the second parallelogram relation and the second triangle relation from Definition 6.15.

EXERCISE 6.10. Show that the interleaving distance satisfies the triangle inequality. [Hint: show that an ϵ-interleaving between $(U_\bullet, a_\bullet)$ and $(V_\bullet, b_\bullet)$ can always be combined with an ϵ'-interleaving between $(V_\bullet, b_\bullet)$ and $(W_\bullet, c_\bullet)$ to produce an $(\epsilon + \epsilon')$-interleaving between $(U_\bullet, a_\bullet)$ and $(W_\bullet, c_\bullet)$.]

EXERCISE 6.11. Let $a < a' < b < b'$ be four positive real numbers. What is the interleaving distance between the two $\mathbb{R}_+$-indexed interval modules $(I_\bullet^{a,b}, c_\bullet^{a,b})$ and $(I_\bullet^{a',b'}, c_\bullet^{a',b'})$?

EXERCISE 6.12. Show that the bottleneck distance satisfies the triangle inequality.

EXERCISE 6.13. State and prove a variant of Proposition 6.20 for Čech filtrations.

EXERCISE 6.14. Let $\mathscr{S}$ be a sheaf over a simplicial complex K and Σ an $\mathscr{S}$-compatible acyclic partial matching. Mimic the argument from Proposition 8.8 to show that the Morse complex of Σ with coefficients in $\mathscr{S}$ (see Definition 8.17) is a cochain complex.

7. SHEAVES

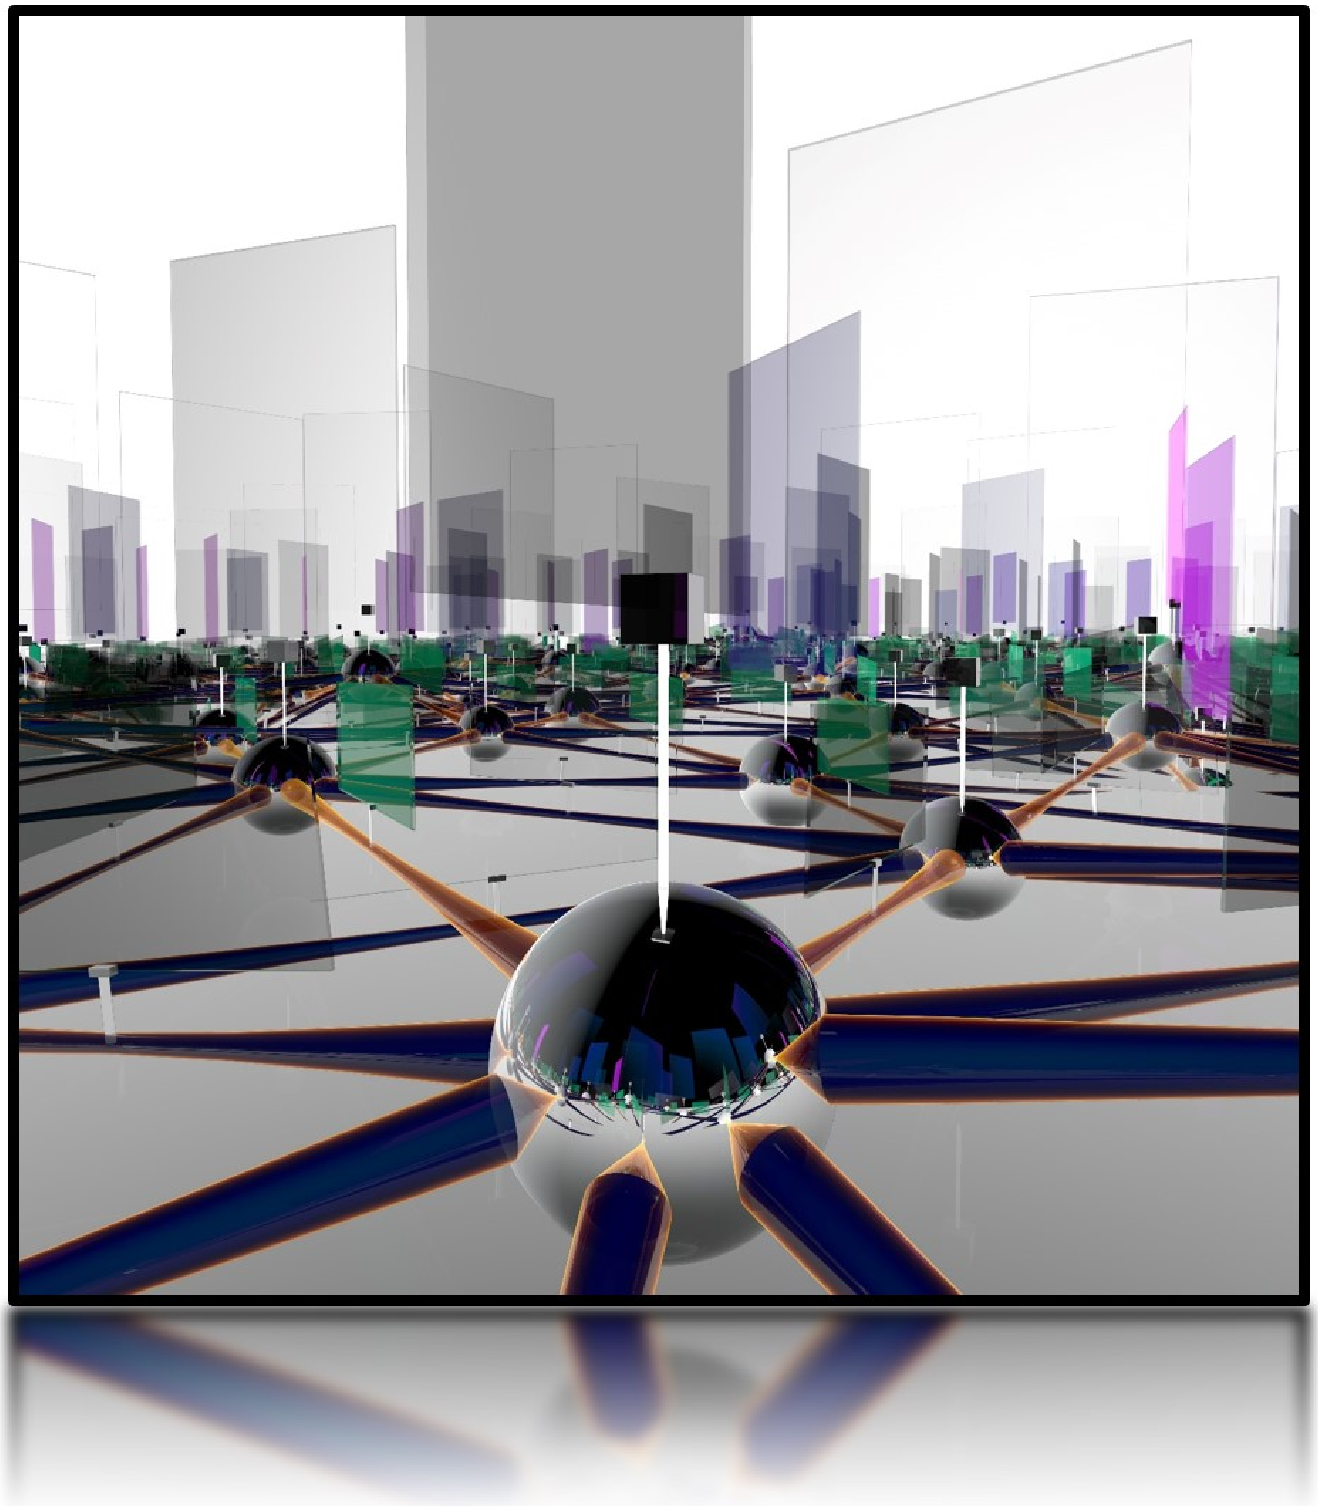

7.1 Fibers and Persistence

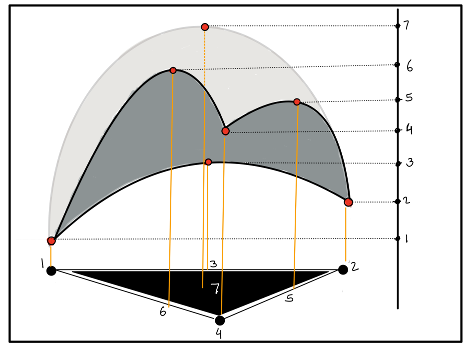

Let $f : K \to \mathbb{R}_+$ be a function that assigns a non-negative real number $f(\sigma)$ to every simplex σ of a simplicial complex K. We call f *monotone* if it satisfies $f(\sigma) \leq f(\tau)$ whenever σ is a face of τ in K. Specifying a monotone f is equivalent to imposing an $\mathbb{R}_+$-indexed filtration $\boldsymbol{F}_\bullet$ on K — to discover this filtration, one uses the rule $\boldsymbol{F}_t K = \{\sigma \in K \mid f(\sigma) \leq t\}$. We call $\boldsymbol{F}$ the **sublevelset filtration** of K with respect to f. Conversely, if we are given a filtration $\boldsymbol{F}_\bullet$ of K, then the corresponding monotone function $f : K \to \mathbb{R}_+$ is given by $f(\sigma) = \inf\{t \in \mathbb{R}_+ \mid \sigma \in \boldsymbol{F}_t K\}$. Thus, much of persistent homology (particularly its application to the study of filtered simplicial complexes) can be interpreted as the systematic analysis of homology groups associated to certain *fibers* of f — for each $t \in \mathbb{R}^+$, the fiber of interest is a subcomplex of K:

$$\{f \leq t\} := \{\sigma \in K \mid 0 \leq f(\sigma) \leq t\}$$

Thanks to the finiteness of K, taking the k-th homology of sublevelset filtrations always produces tame persistence modules (in the sense of Definition 6.13); thus these modules admit a barcode decomposition as guaranteed by Corollary 6.14. These barcodes satisfy two special properties: first, they allow us to combinatorially describe the homology of each fiber $\{f \leq t\}$ and the rank of the linear maps

$$\mathbf{H}_k\big(\{f \leq t\}\big) \to \mathbf{H}_k\big(\{f \leq s\}\big)$$

induced on k-th homology by inclusion of fibers for all pairs $t \leq s$. Second, if we have a another monotone function $f' : K \to \mathbb{R}$ that is ϵ-close to our f, i.e., if we have

$$|f(\sigma) - f'(\sigma)| < \epsilon \text{ for every } \sigma \text{ in } K,$$

then the barcodes for f' will be no more than ϵ-apart from those of f with respect to the bottleneck distance (see Definition 6.18 and Exercise 7.1). Thus, all intervals longer than 2ϵ in the barcode of f correspond to fiber homology classes that are stable with respect to ϵ-perturbations of f.

Card-carrying mathematicians will immediately wonder whether similar stability results can be obtained for maps $K \to X$ when X is more complicated than $\mathbb{R}_+$: *ars gratia artis*. Those with the ability to withstand this temptation to generalize might instead be compelled by more practical considerations. A monotone map $f : K \to \mathbb{R}_+$ associates a (real-valued) measurement to each simplex, and we are often interested in several such measurements $\{f_i : K \to \mathbb{R}_+ \mid 1 \leq i \leq n\}$ and wish to study (the homology of) their common sublevelsets $\bigcap_{i=1}^n \{f_i \leq t_i\}$ simultaneously. Thus, we may as well assign

$$\sigma \mapsto \big(f_1(\sigma), \ldots, f_n(\sigma)\big)$$

and study the fibers of this single vector-valued map $K \to \mathbb{R}_+^n$.

Even more interesting from a topological viewpoint is the scenario where the f_i associate angles in $[0, 2\pi)$ to simplices; in this case, we have a map $f : K \to \mathbb{T}^n$ to the n-torus (i.e., the product of n circles). Now it no longer makes sense to seek monotonicity or ask about fibers of the form $\{f_i \leq t_i\}$, since there is no natural partial order on points of the n-torus. On the other hand, we can certainly triangulate the torus so that f becomes a simplicial map and study the fiberwise homology of f over simplices (or subcomplexes) of $\mathbb{T}^n$. It is, therefore, in our interest to understand the (co)homology groups of fibers of simplicial maps $f : K \to L$. The optimal data structure which coherently organizes these fiber homology groups is called a **sheaf**.

7.2 Sheaves

Let L be a simplicial complex and $\mathbb{F}$ a field. We write $(L, \leq)$ to denote the poset of simplices in L ordered by the face relation.

Definition 7.1. A **sheaf** over L is a functor $\mathscr{S} : (L, \leq) \to \mathbf{Vect}_{\mathbb{F}}$. In other words, $\mathscr{S}$ assigns

(1) to each simplex τ of L an $\mathbb{F}$-vector space $\mathscr{S}(\tau)$ called the *stalk*, and
(2) to each $\tau \leq \tau'$ in L a linear map $\mathscr{S}(\tau \leq \tau') : \mathscr{S}(\tau) \to \mathscr{S}(\tau')$ called the *restriction map*,

subject to the usual (identity and associativity) categorical axioms:

(1) for every simplex τ in L, the map $\mathscr{S}(\tau \leq \tau)$ is the identity on $\mathscr{S}(\tau)$, and
(2) for every triple $\tau \leq \tau' \leq \tau''$ in L, we have $\mathscr{S}(\tau' \leq \tau'') \circ \mathscr{S}(\tau \leq \tau') = \mathscr{S}(\tau \leq \tau'')$.

We call L the **base space** of the sheaf $\mathscr{S}$. From a purely algebraic perspective, $\mathscr{S}$ is an arrangement of $\mathbb{F}$-vector spaces and linear maps parametrized by the simplices of L and their face relations. Alternately, one may view $\mathscr{S}$ as a gadget which weights these simplices and face relations by vector spaces and linear maps respectively. Although the stalks of a sheaf can vary drastically from simplex to simplex, the associativity requirement places severe constraints on restriction maps. For instance, both composite paths from $\mathscr{S}(\tau)$ to $\mathscr{S}(\tau'')$ in the accompanying figure must evaluate to $\mathscr{S}(\tau \leq \tau'')$. On the other hand, if L is one-dimensional then associativity holds automatically because there are no ascending triples $\tau < \tau' < \tau''$ of simplices.

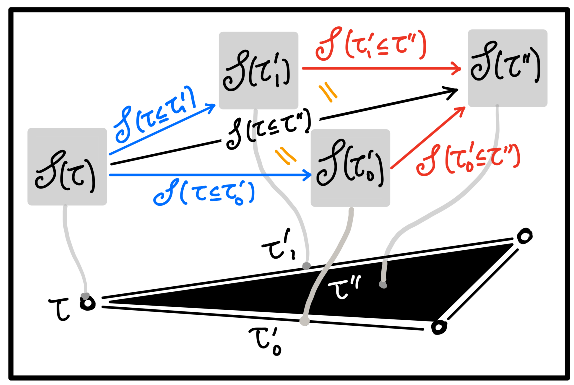

Example 7.2. Here are three examples of sheaves on a simplicial complex L, in increasing order of complexity.

(1) The **zero** sheaf $\underline{0}_L$, as suggested by its name, assigns the trivial (i.e., zero-dimensional) $\mathbb{F}$-vector space to every simplex. This forces all the restriction maps to also be zero.
(2) Given a simplex τ of L, the associated **skyscraper** sheaf $\underline{\mathrm{Sk}}_\tau$ over L assigns the trivial vector space to every simplex except τ, whose stalk is the one-dimensional vector space $\mathbb{F}$. The restriction map associated to $\tau \leq \tau$ is the identity, while all other restriction maps must be zero.
(3) The **constant** sheaf $\underline{\mathbb{F}}_L$ assigns the one-dimensional stalk $\mathbb{F}$ to every simplex of L and the identity restriction map $\mathbb{F} \to \mathbb{F}$ to every face relation in sight.

More interesting examples will become available later.

As mentioned in the previous Section, our main interest in sheaves comes from their remarkable ability to encode the homology groups of fibers of simplicial maps. Recall from (2) that the fiber of a simplicial map $f : K \to L$ under a simplex τ of L is the subcomplex of K given by

$$\tau/f = \{\sigma \in K \mid f(\sigma) \leq \tau\}.$$

And moreover, for any pair $\tau \leq \tau'$ in L there is an obvious inclusion of fibers $\tau/f \hookrightarrow \tau'/f$ because any σ in K satisfying $f(\sigma) \leq \tau$ automatically satisfies $f(\sigma) \leq \tau'$. Thus, fitting the homology groups $\mathbf{H}_k(\tau/f; \mathbb{F})$ into a sheaf over L becomes a matter of invoking the functoriality of homology with respect to inclusion maps.

PROPOSITION 7.3. *Let $f : K \to L$ be a simplicial map. For each dimension $k \geq 0$, the assignments*

$$\tau \mapsto \mathbf{H}_k(\tau/f), \text{ and}$$
$$(\tau \leq \tau') \mapsto \mathbf{H}_k(\tau/f \hookrightarrow \tau'/f)$$

constitute a sheaf over L, which we denote $\mathscr{F}_f^k$ and call the k-th **fiber homology sheaf** *of f.*

The proof is not complicated — for any triple of simplices $\tau \leq \tau' \leq \tau''$ in L, the inclusion $\tau/f \hookrightarrow \tau''/f$ factors through τ'/f; the identity and associativity axioms of Definition 7.1 are satisfied simply because homology is functorial. It should also be noted that in general some fiber τ/f might be empty, in which case we would have $\mathscr{F}_f^k(\tau) = \mathbf{H}_k(\tau/f) = 0$ for all k.

EXAMPLE 7.4. Fiber homology sheaves of the identity simplicial map $\mathrm{id} : L \to L$ are already familiar to us — for each simplex τ of L, the fiber τ/id is the subcomplex $\overline{\tau}$ consisting of the single simplex τ along with all of its faces. Each such fiber is contractible by Proposition 2.6, and hence has the homology of a point $\Delta(0)$. Consequently,

$$\mathscr{F}_{\mathrm{id}}^k(\tau) = \begin{cases} \mathbb{F} & k = 0 \\ 0 & k \neq 0 \end{cases}$$

Thus, $\mathscr{F}_{\mathrm{id}}^k$ is the zero sheaf $\underline{0}_L$ whenever $k > 0$. With a bit of effort, one can discover that the restriction maps of $\mathscr{F}_{\mathrm{id}}^0$ are all identities $\mathbb{F} \to \mathbb{F}$, and so $\mathscr{F}_{\mathrm{id}}^0$ is the constant sheaf $\underline{\mathbb{F}}_L$.

Those experiencing nostalgia for persistent homology have no cause for concern: every sheaf $\mathscr{S}$ is filled to the brim with persistence modules. Take any ascending sequence

$$\tau_0 \leq \tau_1 \leq \cdots \leq \tau_n$$

of simplices in the base space L, and note that the restriction maps produce a persistence module

$$\mathscr{S}(\tau_0) \xrightarrow{\mathscr{S}(\tau_0 \leq \tau_1)} \mathscr{S}(\tau_1) \xrightarrow{\mathscr{S}(\tau_1 \leq \tau_2)} \cdots \xrightarrow{\mathscr{S}(\tau_{n-1} \leq \tau_n)} \mathscr{S}(\tau_n).$$

It follows from the associativity axiom of Definition 7.1 that the number of intervals $[i, j]$ in the barcode of this persistence module must equal the rank of $\mathscr{S}(\tau_i \leq \tau_j)$.

7.3 SHEAF COHOMOLOGY

Taking the perspective of sheaves as *algebraic weights on simplices* seriously produces a suite of new cohomology theories for simplicial complexes. To define these sheaf-infused cohomology groups, we must first build a suitable cochain complex using the data of a sheaf; to this end, fix a sheaf $\mathscr{S}$ on a simplicial complex L.

DEFINITION 7.5. For each dimension $k \geq 0$, the vector space of **k-cochains of L with $\mathscr{S}$-coefficients** is the product

$$\mathbf{C}^k(L; \mathscr{S}) = \prod_{\dim \tau = k} \mathscr{S}(\tau)$$

of the stalks of $\mathscr{S}$ over all the k-dimensional simplices of L.

Depending on which sheaf $\mathscr{S}$ is being used as the *coefficient system* in the definition above, the cochain groups $\mathbf{C}^\bullet(L; \mathscr{S})$ might be quite different from the familiar simplicial cochain groups $\mathbf{C}^\bullet(L; \mathbb{F})$ of Definition 5.1 — for instance, when $\mathscr{S} = \underline{0}_L$, we obtain trivial cochain groups in all dimensions regardless of L. But for $\mathscr{S} = \underline{\mathbb{F}}_L$, we recover the usual simplicial cochain groups of

L. The key point is that while the constant sheaf identifies a unique one-dimensional subspace of $\mathbf{C}^k(L)$ with every k-simplex of L, using a different sheaf $\mathscr{S}$ allows us to upgrade the contribution of some simplices (by assigning them stalks of dimension > 1) and diminishing the contribution of others (by assigning them zero stalks).

Let's assume that the vertices of L are ordered so that each k-simplex τ has a well-defined i-th face τ_{-i} for i in $\{0, \ldots, k\}$ (see Definition 3.4). For each pair of simplices τ, τ' in L we write

$$[\tau : \tau'] := \begin{cases} +1 & \tau = \tau'_{-i} \text{ for } i \text{ even,} \\ -1 & \tau = \tau'_{-i} \text{ for } i \text{ odd,} \\ 0 & \text{otherwise.} \end{cases}$$

Thus, $[\tau : \tau'] \in \mathbb{F}$ is precisely the coefficient of τ' in the simplicial coboundary of τ, or equivalently, the coefficient of τ in the simplicial boundary of τ'.

DEFINITION 7.6. For each $k \geq 0$, the k-th **coboundary map** of L **with $\mathscr{S}$-coefficients** is the linear map

$$\partial^k_{\mathscr{S}} : \mathbf{C}^k(L; \mathscr{S}) \to \mathbf{C}^{k+1}(L; \mathscr{S})$$

defined via the following block-action: for each pair of simplices $\tau \leq \tau'$ with $\dim \tau = k$ and $\dim \tau' = k+1$, the $\mathscr{S}(\tau) \to \mathscr{S}(\tau')$ component of $\partial^k_{\mathscr{S}}$ is given by

$$\partial^k_{\mathscr{S}}|_{\tau,\tau'} = [\tau : \tau'] \cdot \mathscr{S}(\tau \leq \tau') \tag{6}$$

From a computational perspective, it often helps to view $\partial^k_{\mathscr{S}}$ as an enormous block-matrix whose columns are indexed by (stalks of) all the k-simplices in L and rows are indexed by (stalks of) all the $(k+1)$-simplices; the component $\partial^k_{\mathscr{S}}|_{\tau,\tau'}$ is the block in the column of τ and the row of τ':

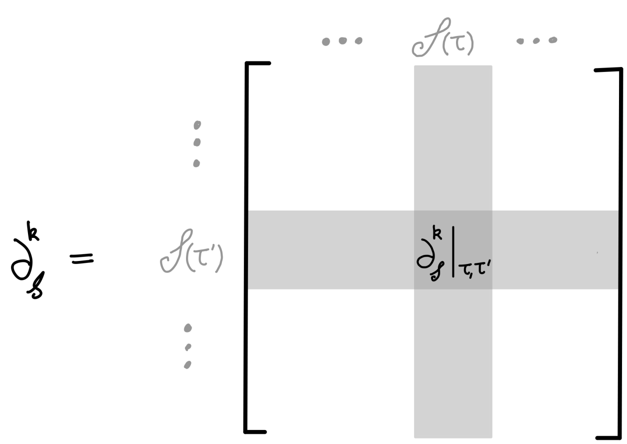

The expression (6) for $\partial^k_{\mathscr{S}}|_{\tau,\tau'}$ involves a restriction map, but note that it makes sense even when τ is not a face of τ': in this case, the scalar $[\tau : \tau']$ is zero, so the entire block is zero.

REMARK 7.7. If $\mathscr{S}$ is the constant sheaf $\underline{\mathbb{F}}_L$, then all the rows and columns have width one (since all the stalks are one-dimensional); and since the restriction maps in this case are all identities, the entry $\partial^k_{\mathscr{S}}|_{\tau,\tau'}$ lies in $\{0, \pm 1\}$ depending on whether or not τ is a face of τ'. Thus, both $\mathbf{C}^k(L; \mathscr{S})$ and $\partial^k_{\mathscr{S}}$ reduce to the familiar objects from Definition 5.1 when $\mathscr{S} = \underline{\mathbb{F}}_L$.

The harsh constraints placed on restriction maps of $\mathscr{S}$ by the associativity axiom of Definition 7.1 will now start yielding rich dividends. The following result establishes that the choice of terminology (cochains and coboundary operators) for the objects $\mathbf{C}^k(L; \mathscr{S})$ and $\partial^k_{\mathscr{S}}$ is apposite.

PROPOSITION 7.8. *The sequence*

$$0 \longrightarrow \mathbf{C}^0(L;\mathscr{S}) \xrightarrow{\partial^0_{\mathscr{S}}} \mathbf{C}^1(L;\mathscr{S}) \xrightarrow{\partial^1_{\mathscr{S}}} \cdots \xrightarrow{\partial^{k-1}_{\mathscr{S}}} \mathbf{C}^k(L;\mathscr{S}) \xrightarrow{\partial^k_{\mathscr{S}}} \mathbf{C}^{k+1}(L;\mathscr{S}) \xrightarrow{\partial^{k+1}_{\mathscr{S}}} \cdots$$

forms a cochain complex over $\mathbb{F}$. *In other words,* $\partial^k_{\mathscr{S}} \circ \partial^{k-1}_{\mathscr{S}}$ *equals zero for all* $k \geq 1$.

PROOF. It suffices to verify that the composite of two adjacent coboundary operators equals zero block-wise. Namely, for each $(k-1)$-simplex τ and $(k+1)$-simplex τ'' we will show that the $\mathscr{F}(\tau) \to \mathscr{F}(\tau'')$ block of this composite is the zero map, from which the desired conclusion immediately follows. For any vector v in $\mathscr{F}(\tau)$, we calculate

$$\begin{aligned}
\partial^k_{\mathscr{S}} \circ \partial^{k-1}_{\mathscr{S}}(v) &= \sum_{\dim \tau' = k} \partial^k_{\mathscr{S}}|_{\tau',\tau''} \circ \partial^{k-1}_{\mathscr{S}}|_{\tau,\tau'}(v) && \text{by Definition 7.6} \\
&= \sum_{\tau<\tau'<\tau''} \partial^k_{\mathscr{S}}|_{\tau',\tau''} \circ \partial^{k-1}_{\mathscr{S}}|_{\tau,\tau'}(v) && \text{eliminating zero terms} \\
&= \sum_{\tau<\tau'<\tau''} [\tau' : \tau''] \cdot [\tau : \tau'] \cdot \mathscr{S}(\tau' \leq \tau'') \circ \mathscr{S}(\tau \leq \tau')(v) && \text{by (6)} \\
&= \sum_{\tau<\tau'<\tau''} [\tau' : \tau''] \cdot [\tau : \tau'] \cdot \mathscr{S}(\tau \leq \tau'')(v) && \text{associativity axiom!} \\
&= \left(\sum_{\tau<\tau'<\tau''} [\tau' : \tau''] \cdot [\tau : \tau'] \right) \cdot \mathscr{S}(\tau \leq \tau'')(v) && \text{collecting scalars}
\end{aligned}$$

But now the scalar in parentheses is zero because it equals the coefficient of τ'' in the composite $\partial^k_L \circ \partial^{k-1}_L(\tau)$. Since our choice of v was arbitrary, the composite $\partial^k_{\mathscr{S}} \circ \partial^{k-1}_{\mathscr{S}}$ is identically zero as desired. □

Having produced a cochain complex from $\mathscr{S}$, we can safely define the associated cohomology groups in the usual fashion.

DEFINITION 7.9. For each dimension $k \geq 0$, the k-th **cohomology group of L with coefficients in $\mathscr{S}$** is the quotient vector space

$$\mathbf{H}^k(L;\mathscr{S}) = \ker \partial^k_{\mathscr{S}} / \operatorname{img} \partial^{k-1}_{\mathscr{S}}.$$

At the moment, this definition is simply a way of producing cohomology groups from sheaves. We know, based on the discussion above, that this *sheaf cohomology* agrees with standard cohomology whenever $\mathscr{S}$ is the constant sheaf $\underline{\mathbb{F}}_L$. It is challenging to visualize sheaf cohomology for more general choices of $\mathscr{S}$; but in the next Section, we will provide a topological interpretation for the simplest sheaf cohomology group $\mathbf{H}^0(L;\mathscr{S})$ for arbitrary $\mathscr{S}$.

7.4 The Étale Space and Sections

Let L be a simplicial complex and $\mathscr{S}$ a sheaf on L; both will remain fixed throughout this section. We recall that the geometric realization of every simplex τ in L is denoted $|\tau| \subset |L|$ (see Definition 1.7) and its open star (from Definition 1.17) is denoted $\mathbf{st}(\tau) \subset L$. The realization of this open star is

$$|\,\mathbf{st}(\tau)| = \bigcup_{\tau \leq \tau'} |\tau'|^\circ,$$

where $|\tau'|^\circ$ stands for the interior of $|\tau'|$ in $|L|$. For each $x \in |L|$ there is a unique simplex $\tau \in L$ with $x \in |\tau|^\circ$, which we will denote by τ_x throughout this section.

DEFINITION 7.10. The **étale space** of a sheaf $\mathscr{S}$ on L is the topological space $\mathbf{E}\mathscr{S}$ defined as follows. Its underlying set consists of pairs

$$\mathbf{E}\mathscr{S} = \{(x, v) \mid v \in \mathscr{S}(\tau_x)\}.$$

A basis for the topology is prescribed by open sets $U_{\tau,v} \subset \mathbf{E}\mathscr{S}$ indexed by pairs (τ, v) where $\tau \in L$ is a simplex and $v \in \mathscr{S}(\tau)$ is a vector lying in its stalk. Each such *basic open* set is:

$$U_{\tau,v} = \{(x, w) \mid \tau_x \geq \tau \text{ and } w = \mathscr{S}(\tau \leq \tau_x)(v)\}.$$

There is a natural projection $\pi_{\mathscr{S}} : \mathbf{E}\mathscr{S} \twoheadrightarrow |L|$ sending each (x, v) to x; this is called the **étale map** of $\mathscr{S}$ and it satisfies two strong properties. First, its restriction to each basic open $U_{\tau,v}$ is a homeomorphism onto $|\,\mathbf{st}(\tau)|$. And second, for each x in L we have

$$\pi_{\mathscr{S}}^{-1}(x) = \{x\} \times \mathscr{S}(\tau_x).$$

Thus, $\pi_{\mathscr{S}}^{-1}(x)$ has the structure of a vector space for each x in $|L|$. The étale space is is home to some very special subspaces; these can be discovered by attempting to find right-inverses for the affiliated étale map.

DEFINITION 7.11. Let $L' \subset L$ be any subcollection of simplices (which do not necessarily form a subcomplex). A **section of $\mathscr{S}$ over** L' is any continuous map $h : |L'| \to \mathbf{E}\mathscr{S}$ for which the composite $\pi_{\mathscr{S}} \circ h$ equals the identity map on $|L'|$. The set of all such sections is denoted $\Gamma(L'; \mathscr{S})$.

The case $L = L'$ is of particular interest — we call $\Gamma(L; \mathscr{S})$ the set of **global sections** of $\mathscr{S}$. Since any section h in $\Gamma(L', \mathscr{S})$ satisfies $\pi_{\mathscr{S}} \circ h = \text{id}$, it must at least send each point x of $|L'|$ to a vector $h(x)$ in the stalk $\mathscr{S}(\tau_x)$. Since h is also continuous, we can make two stronger claims.

PROPOSITION 7.12. *For any subcollection $L' \subset L$ of simplices,*

1. *each section h in $\Gamma(L'; \mathscr{S})$ is constant on $|\tau|^\circ$ for each τ in L'; moreover,*
2. *the set $\Gamma(L'; \mathscr{S})$ has the structure of a vector space.*

PROOF. Fix any simplex τ in L'. Since $\pi_{\mathscr{S}} \circ h$ is the identity, it follows that $h(|\,\mathbf{st}(\tau)|)$ is a subset of $\pi_{\mathscr{S}}^{-1}(|\,\mathbf{st}(\tau)|)$. By definition, there is a decomposition

$$\pi_{\mathscr{S}}^{-1}(|\,\mathbf{st}(\tau)|) \simeq \coprod_{v \in \mathscr{S}(\tau)} U_{\tau,v},$$

where each $U_{\tau,v}$ is a basic open set. Since h is continuous and $|\,\mathbf{st}(\tau)|$ is connected, there is a single v in $\mathscr{S}(\tau)$ so that $h(|\,\mathbf{st}(\tau)|) \subset U_{\tau,v}$. Thus, any two points x and x' in $|\tau|^\circ$ are sent by h to the same vector $\mathscr{S}(\tau \leq \tau)(v) = v$, which proves the first claim. Armed with this knowledge, we may as well view h as a function sending each simplex $\tau \in L'$ to a vector $h(\tau) \in \mathscr{S}(\tau)$. With this shift in perspective, the vector space structure on $\Gamma(L'; \mathscr{S})$ becomes obvious: for any pair of scalars α, β in $\mathbb{F}$ and sections h, g in $\Gamma(L'; \mathscr{S})$, we can form the linear combination $\alpha \cdot h + \beta \cdot g$ that sends each τ to the vector $\alpha \cdot h(\tau) + \beta \cdot g(\tau)$ in $\mathscr{S}(\tau)$. □

Writing sections as assignments of stalk-vectors to simplices of L' (rather than to points of $|L'|$) allows us to view them as finite objects. Implicit in the proof of the above result is the following observation, which establishes that sections correspond to choices of stalk-vectors that are **compatible** with respect to the restriction maps of $\mathscr{S}$.

COROLLARY 7.13. *If h is a section in $\Gamma(L', \mathscr{S})$, then for every pair of simplices $\tau \leq \tau'$ in L' we have the equality*

$$\mathscr{S}(\tau \leq \tau')(h(\tau)) = h(\tau').$$

We have been discussing sections of sheaves because they are intimately related to the sheaf cohomology groups from Definition 7.9.

THEOREM 7.14. *For any sheaf $\mathscr{S}$ over a simplicial complex L, there is a vector space isomorphism*

$$\mathbf{H}^0(L; \mathscr{S}) \simeq \Gamma(L; \mathscr{S})$$

between the zeroth cohomology groups of L with coefficients in $\mathscr{S}$ and the global sections of $\mathscr{S}$.

PROOF. Although this proof has been assigned as an exercise, we show the first step of the argument here as a (substantial) hint. The zeroth cohomology $\mathbf{H}^0(L; \mathscr{S})$ is precisely the kernel of the coboundary map $\partial^0_{\mathscr{S}}$, whose block structure has been described in Definition 7.6. The row-blocks are indexed by the 1-simplices, each of which contains exactly two vertices in its boundary. The row corresponding to a 1-simplex $\tau = (u_0, u_1)$ can only have nonzero blocks in the two columns corresponding to its vertices u_0 and u_1. Thus, a cochain v in $\mathbf{C}^0(L; \mathscr{S})$ lies in the kernel of this coboundary matrix if and only if its components $v_i \in \mathscr{S}(u_i)$ for i in $\{0, 1\}$ satisfy

$$\mathscr{S}(u_0 \leq \tau)(v_0) = \mathscr{S}(u_1 \leq \tau)(v_1).$$

This is the first step in showing that v constitutes a section. $\square$

REMARK 7.15. When defining sections of $\mathscr{S}$ over subsets of L, we only used the topology of $\mathbf{E}\mathscr{S}$ and properties of the map $\pi_{\mathscr{S}} : \mathbf{E}\mathscr{S} \twoheadrightarrow |L|$. In fact, one can completely recover $\mathscr{S}$ from its étale map: the stalk $\mathscr{S}(\tau)$ over each simplex τ of L is the vector space of sections $\Gamma(|\,\mathbf{st}(\tau)|; \mathscr{S})$ over its open star, and the restriction map associated to $\tau \leq \tau'$ is obtained by using the fact that every section $|\,\mathbf{st}(\tau)| \to \mathbf{E}\mathscr{S}$ restricts to a section over the smaller set $|\,\mathbf{st}(\tau')|$.

7.5 PUSHFORWARDS AND PULLBACKS

There is a natural way to define maps of sheaves over a fixed simplicial complex L.

DEFINITION 7.16. A **morphism of sheaves** $\Phi_\bullet : \mathscr{S} \to \mathscr{S}'$ over L consists of linear maps $\Phi_\tau : \mathscr{S}(\tau) \to \mathscr{S}'(\tau)$ indexed by simplices $\tau \in L$ so that the following diagram of vector spaces commutes for each $\tau \leq \tau'$:

$$\begin{tikzcd} \mathscr{S}(\tau) \arrow[r, "\Phi_\tau"] \arrow[d, "\mathscr{S}(\tau\leq\tau')"'] & \mathscr{S}'(\tau) \arrow[d, "\mathscr{S}'(\tau\leq\tau')"] \\ \mathscr{S}(\tau') \arrow[r, "\Phi_{\tau'}"'] & \mathscr{S}'(\tau'). \end{tikzcd}$$

These morphisms endow the set of all sheaves over L with the structure of a category, which we will denote by $\mathbf{Sh}(L)$. Sheaf morphisms induce well-defined maps on sheaf cohomology (this is an exercise to this Chapter).

Our goal here is to show how sheaves can be transported back and forth between a pair of simplicial complexes K and L by using a simplicial map $f : K \to L$. Surprisingly, the easier direction is backwards: we can construct a sheaf on K from a sheaf on L in a relatively straightforward manner.

DEFINITION 7.17. The **pullback** of a sheaf $\mathscr{S}$ over L across the simplicial map $f : K \to L$ is a new sheaf $f^*\mathscr{S}$ over K defined as follows. The stalk over every simplex σ in K is

$$f^*\mathscr{S}(\sigma) = \mathscr{S}(f(\sigma)),$$

while the restriction map for $\sigma \leq \sigma'$ is

$$f^*\mathscr{S}(\sigma \leq \sigma') = \mathscr{S}(f(\sigma) \leq f(\sigma'))$$

Transporting sheaves from K forwards to L along $f : K \to L$ is more intricate, because now the direction of f works against us. For each simplex τ of L, there might be a large collection of simplices in K which get mapped to (a co-face of) τ; we must somehow combine the $\mathscr{T}$-stalks over all these simplices in order to produce a sheaf over K. Here it helps to utilize the perspective from Remark 7.15 and define the desired sheaf in terms of its étale space.

DEFINITION 7.18. The **pushforward** of a sheaf $\mathscr{T}$ on K along a simplicial map $f : K \to L$ is a new sheaf $f_*\mathscr{T}$ on L whose étale space equals

$$\mathbf{E}f_*\mathscr{T} = \left\{ \big(|f|(x), v\big) \mid (x, v) \in \mathbf{E}\mathscr{T} \right\};$$

here $|f| : |K| \to |L|$ is the continuous map induced by f.

By our recipe for extracting sheaves from their étale spaces, it follows that the stalk $f_*\mathscr{T}(\tau)$ for each simplex τ of L is the vector space of sections $\Gamma(|f/\tau|; \mathscr{T})$, where f/τ is the *dual* fiber

$$f/\tau = \{\sigma \in K \mid f(\sigma) \geq \tau\}.$$

Although this dual fiber is not generally a subcomplex of K unlike τ/f, the space of $\mathscr{T}$'s sections over it is still well-defined.

REMARK 7.19. Pullbacks and pushforwards are functors between $\mathbf{Sh}(K)$ and $\mathbf{Sh}(L)$ — so, we can pull and push not only sheaves but also their morphisms. Moreover, they form a dual *adjoint pair* in the following sense. Given a simplicial map $f : K \to L$ along with sheaves $\mathscr{S} \in \mathbf{Sh}(L)$ and $\mathscr{T} \in \mathbf{Sh}(K)$, there is a bijection

$$\begin{bmatrix} \text{Morphisms} \\ f^*\mathscr{S} \to \mathscr{T} \\ \text{in } \mathbf{Sh}(K) \end{bmatrix} \simeq \begin{bmatrix} \text{Morphisms} \\ \mathscr{S} \to f_*\mathscr{T} \\ \text{in } \mathbf{Sh}(L) \end{bmatrix}$$

To prove this, one must first discover natural sheaf morphisms

$$\mathscr{S} \to f_*f^*\mathscr{S} \quad \text{and} \quad f^*f_*\mathscr{T} \to \mathscr{T}$$

in $\mathbf{Sh}(L)$ and $\mathbf{Sh}(K)$ respectively. The best way to become familiar with pushforwards and pullbacks is to find these morphisms on your own and use them to establish this bijection.

7.6 BONUS: COSHEAVES

Sheaves come with a cohomology theory because of the directions of their restriction maps, which point from low-dimensional simplices to high-dimensional ones. In order to produce an equal and opposite homology theory, one requires maps going in the other direction; this is achieved by reversing the partial order on the simplices of the base space L.

DEFINITION 7.20. A **cosheaf** over L is a functor $\mathscr{C} : (L, \geq) \to \mathbf{Vect}_{\mathbb{F}}$.

Thus, $\mathscr{C}$ assigns an $\mathbb{F}$-vector space $\mathscr{C}(\tau)$ (called the *costalk*) to each simplex τ of L; and it assigns a a linear map $\mathscr{S}(\tau \geq \tau') : \mathscr{S}(\tau) \to \mathscr{S}(\tau')$ (called the *extension map*) to each coface relation $\tau \geq \tau'$ in L. Moreover, we require the expected axioms to hold:

(1) the map $\mathscr{C}(\tau \geq \tau)$ is the identity on $\mathscr{C}(\tau)$, and
(2) the equality $\mathscr{C}(\tau' \geq \tau'') \circ \mathscr{C}(\tau \geq \tau') = \mathscr{C}(\tau \geq \tau'')$ holds for every triple of simplices $\tau \geq \tau' \geq \tau''$.

All of the constructions and results which have been described for sheaves in this Chapter also admit cosheafy analogues — for instance, every cosheaf $\mathscr{C}$ on L induces a chain complex

$$\cdots \xrightarrow{\partial^{\mathscr{C}}_{k+1}} \mathbf{C}_k(L;\mathscr{C}) \xrightarrow{\partial^{\mathscr{C}}_{k}} \mathbf{C}_{k-1}(L;\mathscr{C}) \xrightarrow{\partial^{\mathscr{C}}_{k-1}} \cdots \xrightarrow{\partial^{\mathscr{C}}_{2}} \mathbf{C}_1(L;\mathscr{C}) \xrightarrow{\partial^{\mathscr{C}}_{1}} \mathbf{C}_0(L;\mathscr{C}) \longrightarrow 0$$

which gives rise to the **homology of** L **with coefficients in** $\mathscr{C}$. Similarly, there are dual notions of étale spaces, pushforwards and pullbacks for cosheaves.

EXERCISES

EXERCISE 7.1. Given two monotone functions $f, f' : K \to \mathbb{R}$ on a simplicial complex K, assume there exists some $\epsilon > 0$ so that $|f(\sigma) - f'(\sigma)| < \epsilon$ holds for every simplex σ of K. Letting $\boldsymbol{F_\bullet}$ and $\boldsymbol{F'_\bullet}$ denote the sublevelset filtrations of K with respect to f and f' respectively, show that the barcodes of $\mathbf{H}_k(\boldsymbol{F_\bullet}K)$ and $\mathbf{H}_k(\boldsymbol{F'_\bullet}K)$ have bottleneck distance at most ϵ for every $k \geq 0$. [Hint: find an ϵ interleaving of the two persistence modules and use Theorem 6.19]

EXERCISE 7.2. Describe the stalks and restriction maps of the fiber homology sheaves $\mathscr{F}^k_f$ for $k \geq 0$ when f is the inclusion $\partial\Delta(k) \hookrightarrow \Delta(k)$.

EXERCISE 7.3. Let L be a simplicial complex and τ a simplex in L of dimension $k \geq 0$. What are the cohomology groups of L with coefficients in the skyscraper sheaf $\underline{\mathrm{Sk}}_\tau$?

EXERCISE 7.4. Let $f : \partial\Delta(2) \hookrightarrow \Delta(2)$ be the inclusion map and $\mathscr{F}^k_f$ the associated fiber homology sheaf for each $k \geq 0$. Compute the cohomology groups $\mathbf{H}^i(\Delta(2), \mathscr{F}^j_f)$ for all four pairs $0 \leq i, j \leq 1$.

EXERCISE 7.5. Find a sheaf $\mathscr{S}$ on a contractible simplicial complex L for which $\mathbf{H}^1(L;\mathscr{S})$ is nonzero.

EXERCISE 7.6. Show how Corollary 7.13 follows from the argument which was used to prove Proposition 7.12.

EXERCISE 7.7. Show that every morphism $\Phi : \mathscr{S} \to \mathscr{S}'$ of sheaves over a simplicial complex L induces well-defined linear maps $\mathbf{H}^k(L;\mathscr{S}) \to \mathbf{H}^k(L;\mathscr{S}')$ of sheaf cohomology groups.

EXERCISE 7.8. Show that the pullback $f^*\mathscr{S}$ of a sheaf over L across a simplicial map $f : K \to L$ is a sheaf over K

EXERCISE 7.9. Complete the proof of Theorem 7.14.

EXERCISE 7.10. Show that for every simplicial map $f : K \to L$ and each dimension $k \geq 0$, the assignment of fiberwise *cohomology* groups $\tau \mapsto \mathbf{H}^k(\tau/f)$ constitutes a cosheaf over L.

8. Gradients

8.1 Acyclic Partial Matchings

Let K be a simplicial complex. For any pair of simplices σ, τ in K, we write $\sigma \lhd \tau$ to indicate that σ is a codimension one face of τ, i.e., that $\sigma \leq \tau$ and $\dim \tau - \dim \sigma = 1$.

Definition 8.1. A **partial matching** on K is a collection $\Sigma = \{(\sigma_\bullet \lhd \tau_\bullet)\}$ of simplex-pairs in K subject to the following constraint: if a pair $(\sigma \lhd \tau)$ lies in Σ, then neither σ nor τ appear in any other pair of Σ.

More elaborately, a partial matching Σ consists of two disjoint subsets of simplices $S_\Sigma, T_\Sigma \subset K$ along with a bijection $\mu_\Sigma : S_\Sigma \xrightarrow{\sim} T_\Sigma$ so that $\sigma \lhd \mu(\sigma)$ holds for every σ in S_Σ. Crucially, we do not require $K = S_\Sigma \cup T_\Sigma$, so there might be simplices in K which remain untouched by the matching. These unmatched simplices lying in the complement $C_\Sigma := K - (S_\Sigma \cup T_\Sigma)$ are called **Σ-critical**. It should also be noted that none of the sets S_Σ, T_Σ and C_Σ are required by this definition to be subcomplexes of K.

Partial matchings are relevant to us because under certain assumptions (to be described in gory detail below), we can compute the homology groups of K using a chain complex whose chain groups are built using *only the critical simplices* of a partial matching. Thus, finding a good partial matching with very few critical simplices makes it possible to drastically reduce the algorithmic burden of computing homology groups. Before describing all this machinery, we will examine some examples (and non-examples) of partial matchings.

Example 8.2. Partial matchings are usually illustrated using arrows pointing from the smaller simplex σ to the larger simplex τ whenever $(\sigma \lhd \tau)$ lies in Σ. Consider the diagrams **I-IV** below:

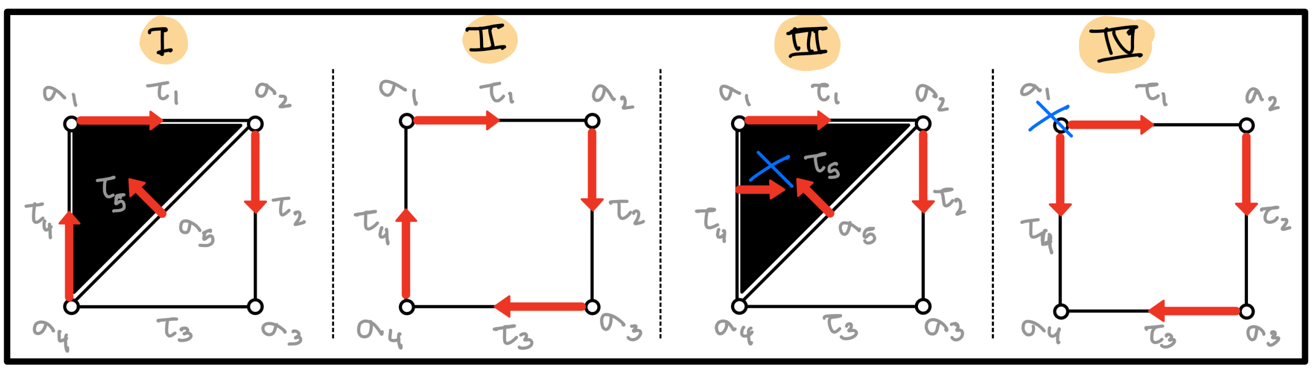

Both **I** and **II** constitute legal partial matchings — the elements of S_Σ are sources of arrows while the elements of T_Σ are targets. The simplices σ_3 and τ_3 in **I** remain untouched by arrows and are therefore critical (but note that **II** has no critical simplices). Neither **III** nor **IV** are partial matchings — in **III** there is a simplex with two incoming arrows whereas in **IV** there is a simplex with two outgoing ones.

Fix a partial matching Σ on K.

Definition 8.3. A **Σ-path** is a zigzag sequence of distinct simplices in K of the form

$$\rho = (\sigma_1 \lhd \tau_1 \rhd \sigma_2 \lhd \tau_2 \rhd \cdots \rhd \sigma_m \lhd \tau_m), \tag{7}$$

where $(\sigma_i \lhd \tau_i)$ lies in Σ for all i in $\{1, \ldots, m\}$. Such a path is **gradient** if either $m = 1$ or σ_1 is not a face of τ_m. We say that Σ is an **acyclic** partial matching if all of its paths are gradient.

Of the two legal partial matchings depicted in Example 8.2 above, only **I** is acyclic — the non-gradient paths in **II** can be discovered by starting at any vertex and following arrows until the loop is completed. Henceforth we will only consider acyclic partial matchings; our interest in this special subset is primarily motivated by the following result.

THEOREM 8.4. *Let Σ be an acyclic partial matching on a simplicial complex K, and let $\mathbb{F}$ be any coefficient field. There exists a chain complex (of $\mathbb{F}$-vector spaces)*

$$\cdots \xrightarrow{d^\Sigma_{k+1}} \mathbf{C}^\Sigma_k(K;\mathbb{F}) \xrightarrow{d^\Sigma_k} \mathbf{C}^\Sigma_{k-1}(K;\mathbb{F}) \xrightarrow{d^\Sigma_{k-1}} \cdots \xrightarrow{d^\Sigma_2} \mathbf{C}^\Sigma_1(K;\mathbb{F}) \xrightarrow{d^\Sigma_1} \mathbf{C}^\Sigma_0(K;\mathbb{F}) \longrightarrow 0$$

satisfying three properties:

(1) *each chain group $\mathbf{C}^\Sigma_k(K;\mathbb{F})$ is $\bigoplus_\alpha \mathbb{F}$, indexed by critical k-simplices $\alpha \in C_\Sigma$,*
(2) *the boundary operators d^Σ_k are explicitly determined by knowledge of Σ-paths, and*
(3) *the homology groups of $(\mathbf{C}^\Sigma_\bullet(K;\mathbb{F}), d^\Sigma_\bullet)$ are isomorphic to those of K.*

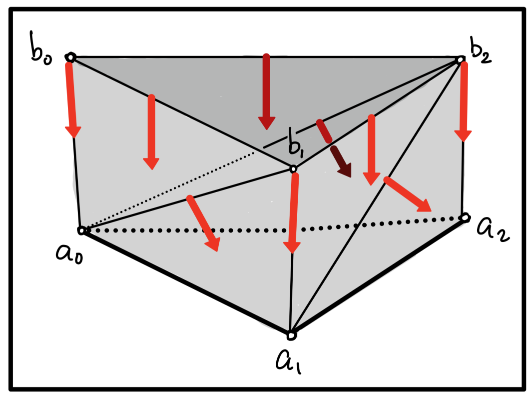

The next two Sections are devoted to the task of building the boundary operators $d^\Sigma_\bullet$ from Σ-paths and proving the isomorphism on homology as promised by properties (2) and (3) respectively. If the set of critical simplices $C_\Sigma \subset K$ forms a subcomplex of K, then the Theorem above can be proved without much difficulty. The illustration here contains one example of this easy case: the complex K is a triangulation of the cylinder $\partial\Delta(2) \times [0,1]$, and the critical simplices C_Σ consist of the base circle (spanned by the vertices a_0, a_1, a_2 and the three edges between them). In this case there is a sequence of elementary collapses (as in Proposition 2.14) from K to C_Σ. This establishes a homotopy equivalence, and hence the desired isomorphisms on homology by Theorem 4.24. Thus, our challenge in proving Theorem 8.4 stems from the fact that in general $C_\Sigma \subset K$ will not be a subcomplex.

REMARK 8.5. Acyclic partial matchings are combinatorial analogues of *gradient vector fields* from differential geometry, and the main idea behind the proof of Theorem 8.4 is to deform the original chain complex $(\mathbf{C}_\bullet(K), \partial^K_\bullet)$ to the smaller chain complex $(\mathbf{C}^\Sigma_\bullet(K), d^\Sigma_\bullet)$ by flowing down along the arrows of this combinatorial gradient vector field. As such, Theorem 8.4 forms the simplicial analogue of one of the main results from smooth *Morse theory*. For these historical reasons, $(\mathbf{C}^\Sigma_\bullet(K), d^\Sigma_\bullet)$ is called the **Morse chain complex** associated to Σ, and the study of acyclic partial matchings is called **discrete Morse theory**.

8.2 THE MORSE CHAIN COMPLEX

Let K be a simplicial complex with ordered vertices. Given any simplices σ and τ in K, let $[\tau : \sigma] \in \{0, \pm 1\}$ indicate the coefficient of σ in the boundary of τ (see Definition 3.4) — this number is nonzero if and only if $\sigma \lhd \tau$. Fix an acyclic partial matching Σ on K as in Definition 8.3. Here we will build the boundary operators $d^\Sigma_\bullet$ whose existence was promised in the statement of Theorem 8.4. The first step in this direction is to associate an algebraic contribution to each Σ-path.

DEFINITION 8.6. The **weight** $w(\rho) \in \{\pm 1\}$ of the Σ-path

$$\rho = (\sigma_1 \lhd \tau_1 \rhd \sigma_2 \lhd \tau_2 \rhd \cdots \rhd \sigma_m \lhd \tau_m),$$

is defined to be the product

$$w(\rho) = \frac{-1}{[\tau_1 : \sigma_1]} \cdot [\tau_2 : \sigma_1] \cdot \frac{-1}{[\tau_2 : \sigma_2]} \cdots [\tau_{m-1} : \sigma_m] \cdot \frac{-1}{[\tau_m : \sigma_m]}$$

One can equivalently collect numerators and denominators to express the weight of each Σ-path ρ as a single ratio

$$w(\rho) = (-1)^m \cdot \frac{\prod_{i=1}^{m-1} [\tau_i : \sigma_{i+1}]}{\prod_{i=1}^{m} [\tau_i : \sigma_i]},$$

but the un-collected version will be more convenient for our purposes.

Recall (from the statement of Theorem 8.4) that the vector space $\mathbf{C}_k^{\Sigma}(K)$ has as its basis the set of all k-dimensional Σ-critical simplices. We will define the desired linear maps from assertion (2) of Theorem 8.4 as matrices with respect to these chosen bases. And for each gradient path ρ as in (7), we indicate the first simplex σ_1 and last simplex τ_m by σ_ρ and τ_ρ respectively.

DEFINITION 8.7. For each dimension $k \geq 0$, the k-th **Morse boundary operator** is the linear map $d_k^{\Sigma} : \mathbf{C}_k^{\Sigma}(K) \to \mathbf{C}_{k-1}^{\Sigma}(K)$ given by the following matrix representation: its entry in the column of a critical k-simplex α and the row of a critical $(k-1)$-simplex ω is given by

$$[\alpha : \omega]_{\Sigma} = [\alpha : \omega] + \sum_{\rho} [\alpha : \sigma_\rho] \cdot w(\rho) \cdot [\tau_\rho : \omega], \tag{8}$$

where ρ ranges over all the Σ-paths.

There are three aspects of the formula (8) which might merit deeper consideration. First, the term $[\alpha : \omega]$ on the right side is precisely the entry in ω's column and α's row within the simplicial boundary matrix ∂_k^K — thus, the difference between this original entry and our new Σ-perturbed one is precisely the sum-over-paths term. Second, we don't have to sum over *all* the paths; the only paths that make a non-zero contribution are the ones which flow from α to ω like so:

$$\alpha \rhd \left(\sigma_1 \lhd \tau_1 \rhd \sigma_2 \lhd \tau_2 \rhd \cdots \rhd \sigma_m \lhd \tau_m\right) \rhd \omega.$$

And third, life gets much simpler when working over the field $\mathbb{F} = \mathbb{Z}/2$ because in this case each path connecting α to ω has weight 1; thus, it suffices to simply count the odd/even parity of the number of such connecting Σ-paths.

PROPOSITION 8.8. *The pair* $(\mathbf{C}_\bullet^{\Sigma}(K), d_\bullet^{\Sigma})$ *constitutes a chain complex.*

PROOF. It suffices by induction to show that the desired result holds when Σ consists of a single pair $(\sigma \lhd \tau)$ of simplices in K; thus the set of critical simplices is $C_\Sigma = K - \{\sigma, \tau\}$, and the only Σ-path is

$$\rho = (\sigma \lhd \tau).$$

To show that $d_\bullet^{\Sigma}$ is a boundary operator, we must establish that for each fixed $\alpha, \omega \in C_\Sigma$, the sum

$$B = \sum_{\xi} [\alpha : \xi]_{\Sigma} \cdot [\xi : \omega]_{\Sigma}$$

equals zero when indexed over all $\xi \in C_\Sigma$. Using the formula (8), the contribution of each ξ to this sum is the product

$$B_\xi = \left([\alpha:\xi] - \frac{[\alpha:\sigma]\cdot[\tau:\xi]}{[\tau:\sigma]}\right)\cdot\left([\xi:\omega] - \frac{[\xi:\sigma]\cdot[\tau:\omega]}{[\tau:\sigma]}\right).$$

The negated term in the first factor disappears whenever $\dim\xi \neq \dim\sigma$, and the negated term in the second factor disappears whenever $\dim\xi \neq \dim\tau$. Thus, only three of the four terms survive when we multiply these two factors:

$$B_\xi = [\alpha:\xi]\cdot[\xi:\omega] - \frac{[\alpha:\sigma]\cdot[\tau:\xi]\cdot[\xi:\omega]}{[\tau:\sigma]} - \frac{[\alpha:\xi]\cdot[\xi:\sigma]\cdot[\tau:\omega]}{[\tau:\sigma]}$$

Summing over $\xi \in C_\Sigma$, we have $B = \sum_\xi B_\xi$ given by

$$B = \sum_\xi[\alpha:\xi]\cdot[\xi:\omega] - \frac{[\alpha:\sigma]}{[\tau:\sigma]}\sum_\xi[\tau:\xi]\cdot[\xi:\omega] - \frac{[\tau:\omega]}{[\tau:\sigma]}\sum_\xi[\alpha:\xi]\cdot[\xi:\sigma].$$

It is now straightforward to check that $B = 0$ because $\partial^K_\bullet$ is a boundary operator on $\mathbf{C}_\bullet(K)$. In particular, the first sum evaluates to $-([\alpha:\sigma]\cdot[\sigma:\omega] + [\alpha:\tau]\cdot[\tau:\omega])$, while the second term evaluates to $[\alpha:\sigma]\cdot[\sigma:\omega]$ and the third term to $[\alpha:\tau]\cdot[\tau:\omega]$. ☐

As mentioned before, we call $(\mathbf{C}^\Sigma_\bullet(K), d^\Sigma_\bullet)$ the *Morse chain complex* associated to our acyclic partial matching Σ; although we have not yet shown that it has the same homology as $(\mathbf{C}_\bullet(K), \partial^K_\bullet)$, this is a good time to examine a few known cases and verify this assertion experimentally. One can build an acyclic partial matching on any simplicial complex by performing these two steps over and over until all simplices have been classified as matched or critical — initially, all simplices are unclassified:

(1) classify a simplex of lowest available dimension as critical; then,
(2) while there exist pairs $(\sigma \lhd \tau)$ of unclassified simplices so that σ is the only unclassified face of τ, classify $(\sigma \lhd \tau)$ as matched.

Although this process is not guaranteed to produce the largest acyclic partial matching (i.e., the one containing the fewest possible critical simplices), it is devastatingly effective in practice.

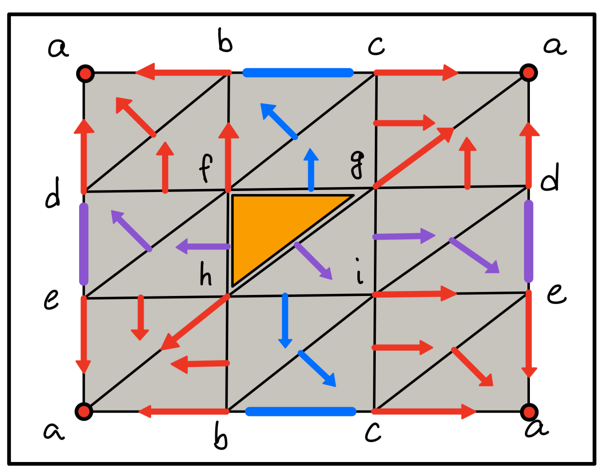

Illustrated here is the acyclic partial matching imposed by this simple two-step algorithm on the torus (note that the left and right edges of the figure have been identified, as have the top and bottom ones). In the first stage, one classifies the vertex a as critical; this creates various edges (such as ab, ad, etc.) with only one unclassified vertex in their boundaries — these produce the matchings indicated by red arrows. At the end of this process, all the vertices have been matched with edges, but there are several 2-simplices remaining with more than one unmatched edge in their boundaries. Next, we classify bc as critical and are allowed to make matchings indicated by the blue arrows. Next, we classify de as critical and make the purple matchings. Finally, only the simplex fgh remains unclassified, so it becomes critical. The critical simplices lie far away from each other, and do not form a subcomplex of the torus.

EXAMPLE 8.9. Let K be the triangulated torus and Σ the overlaid acyclic partial matching illustrated above. The Σ-critical simplices are $\{a, bc, de, fgh\}$, so the associated Morse chain complex has the form

$$\cdots \longrightarrow 0 \longrightarrow \mathbb{F} \xrightarrow{d_2^\Sigma} \mathbb{F}^2 \xrightarrow{d_1^\Sigma} \mathbb{F} \longrightarrow 0$$

To really determine the boundary operators using (8) for arbitrary $\mathbb{F}$, we would have to impose an ordering on the vertices and keep careful track of minus signs. Let's instead work over $\mathbb{Z}/2$ and count gradient paths — there are two from bc to a, namely:

$$bc \rhd (b \lhd ab) \rhd a \quad \text{and} \quad bc \rhd (c \lhd ac) \rhd a.$$

Since there is an even number of connecting gradient paths, the entry $d_1^\Sigma|_{bc,a}$ equals 0. Proceeding similarly, one can check (exercise!) that both d_1^Σ and d_2^Σ are zero maps, which makes it trivial to compute the homology of the torus.

8.3 THE EQUIVALENCE

Let Σ be an acyclic partial matching on a simplicial complex K. Our goal here is to complete the proof of Theorem 8.4 by showing establishing the following result.

PROPOSITION 8.10. *The Morse chain complex $(\mathbf{C}_\bullet^\Sigma(K), d_\bullet^\Sigma)$ of Proposition 8.8 is chain homotopy equivalent to the standard simplicial chain complex $(\mathbf{C}_\bullet(K), \partial_\bullet^K)$.*

In other words, we will describe two chain maps

$$\psi_\bullet : \mathbf{C}_\bullet(K) \to \mathbf{C}_\bullet^\Sigma(K) \quad \text{and} \quad \phi_\bullet : \mathbf{C}_\bullet^\Sigma(K) \to \mathbf{C}_\bullet(K)$$

along with a pair of chain homotopies relating $\phi_\bullet \circ \psi_\bullet$ and $\psi_\bullet \circ \phi_\bullet$ to the identity chain maps on $\mathbf{C}_\bullet(K)$ and $\mathbf{C}_\bullet^\Sigma(K)$ respectively. The best way to build $\psi_\bullet$ and $\phi_\bullet$ is by processing the simplex-pairs $(\sigma \lhd \tau)$ in Σ one at a time. Given this strategy, it is instructive to first examine the special case where Σ contains a single pair $(\sigma \lhd \tau)$.

Consider the entries (in the usual matrix representation) of $\partial_{\dim \tau}^K$ corresponding not only to our chosen pair $(\sigma \lhd \tau)$, but also two arbitrary simplices α and ω.

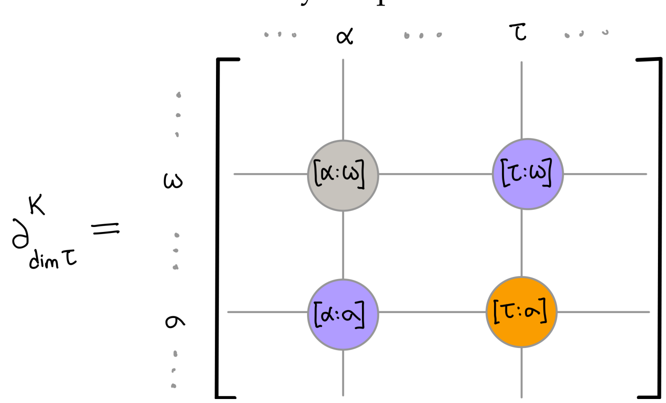

In order to algebraically disentangle σ and τ from the other simplices, we treat the ± 1 entry $[\tau : \sigma]$ as a pivot and seek to clear out all the other entries in both $\mathrm{Col}(\tau)$ and $\mathrm{Row}(\sigma)$. This requires performing row and column operations of the form

$$\mathrm{Row}(\omega) \leftarrow \mathrm{Row}(\omega) - \frac{[\tau : \omega]}{[\tau : \sigma]} \cdot \mathrm{Row}(\sigma) \quad \Big| \quad \mathrm{Col}(\alpha) \leftarrow \mathrm{Col}(\alpha) - \frac{[\alpha : \sigma]}{[\tau : \sigma]} \cdot \mathrm{Col}(\tau). \tag{9}$$

After these operations have been performed, the entry in α's column and ω's row equals

$$[\alpha : \omega] + [\alpha : \sigma] \cdot \frac{-1}{[\tau : \sigma]} \cdot [\tau : \omega], \tag{10}$$

which agrees with the expression for $[\alpha : \omega]_\Sigma$ from (8) because there is only one Σ-path $\sigma \lhd \tau$. More importantly, the row and column operations of (9) suggest the structure of the desired chain maps which take us from $\mathbf{C}_\bullet(K)$ to $\mathbf{C}^\Sigma_\bullet(K)$ and back. This allows us to prove Proposition 8.10 in the special case where Σ contains only one pair.

LEMMA 8.11. *Let Σ be an acyclic partial matching on K containing only one pair $(\sigma \lhd \tau)$. Then the simplicial chain complex $(\mathbf{C}_\bullet(K), \partial^K_\bullet)$ is chain homotopy equivalent to the Morse complex $(\mathbf{C}^\Sigma_\bullet(K), d^\Sigma_\bullet)$.*

PROOF. For each $k \geq 0$, define the linear maps $\psi_k : \mathbf{C}_k(K) \to \mathbf{C}^\Sigma_k(K)$ by the following matrix representation; for each pair of k-simplices (α, ω) in $K \times (K - \{\sigma, \tau\})$, the entry in α's column and ω's row is

$$\psi_k\big|_{\alpha,\omega} = \begin{cases} -\frac{[\tau:\omega]}{[\tau:\sigma]} & \alpha = \sigma \\ 1 & \alpha = \omega \neq \tau \\ 0 & \text{otherwise.} \end{cases} \tag{11}$$

Conversely, define the linear maps $\phi_k : \mathbf{C}^\Sigma_k(K) \to \mathbf{C}_k(K)$ by placing the following entry in the column of ω in $K - \{\sigma, \tau\}$ and the row of α in K:

$$\phi_k\big|_{\omega,\alpha} = \begin{cases} -\frac{[\omega:\sigma]}{[\tau:\sigma]} & \alpha = \tau \\ 1 & \omega = \alpha \neq \sigma \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

Checking that both $\psi_\bullet$ and $\phi_\bullet$ are chain maps has been relegated to two of the Exercises. To extract the chain homotopies, first note that $\psi_\bullet \circ \phi_\bullet$ equals the identity map on $\mathbf{C}^\Sigma_\bullet(K)$. Conversely, the composite $\phi_\bullet \circ \psi_\bullet$ is given by

$$\phi_k \circ \psi_k\big|_{\alpha,\alpha'} = \begin{cases} -\frac{[\tau:\alpha']}{[\tau:\sigma]} & \alpha = \tau \neq \alpha' \\ -\frac{[\alpha:\sigma]}{[\tau:\sigma]} & \alpha \neq \sigma = \alpha' \\ 1 & \alpha = \alpha' \\ 0 & \text{otherwise.} \end{cases}$$

One can now check that the linear maps $\theta_k : \mathbf{C}_k(K) \to \mathbf{C}_{k+1}(K)$ prescribed by

$$\theta\big|_{\alpha,\beta} = \begin{cases} \frac{1}{[\tau:\sigma]} & \alpha = \sigma \text{ and } \beta = \tau \\ 0 & \text{otherwise} \end{cases} \tag{13}$$

furnish the desired chain homotopy between $\phi_k \circ \psi_k$ and the identity chain map. □

The acyclicity of Σ plays an important role when attempting to iteratively apply Lemma 8.11 for the purposes of proving Proposition 8.10. Acyclicity guarantees that removing a single pair $(\sigma \lhd \tau) \in \Sigma$ from K does not alter the entry $[\tau' : \sigma']$ in the boundary matrix corresponding to another pair $(\sigma' \lhd \tau') \in \Sigma$. To see why, note from (10) that the difference between the old and new entries equals

$$\frac{[\tau' : \sigma] \cdot [\tau : \sigma']}{[\tau : \sigma]}.$$

Assuming that the numerator is nonzero, we are forced to conclude that the the Σ-path $\sigma \lhd \tau \rhd \sigma' \lhd \tau'$ is not gradient, which leads to the desired contradiction. As a consequence, the repeated application of Lemma 8.11 correctly converges to the Morse complex regardless of the order in which we remove the simplex-pairs lying in Σ.

8.4 FOR PERSISTENCE

The machinery of acyclic partial matchigns and Morse complexes is extremely flexible, and admits powerful generalizations. Here we will describe how to construct filtered Morse complexes for the purposes of simplifying the persistent homology computations which formed the focus of Chapter 6. Let $\boldsymbol{F}_\bullet K$ be a ($\mathbb{R}_+$-indexed) filtration of a simplicial complex K, and let $b : K \to \mathbb{R}_+$ be the associated monotone function $\sigma \mapsto \inf\{t \geq 0 \mid \sigma \in \boldsymbol{F}_t(K)\}$.

DEFINITION 8.12. An acyclic partial matching Σ on K is **$\boldsymbol{F}_\bullet$-compatible** if $b(\sigma) = b(\tau)$ holds for every pair of simplices $(\sigma \lhd \tau)$ in Σ.

This compatibility requirement forces Σ-paths to be decreasing with respect to b.

PROPOSITION 8.13. *Let Σ be an $\boldsymbol{F}_\bullet$-compatible acyclic partial matching on K. For any Σ-path*

$$\rho = \sigma_1 \lhd \tau_1 \rhd \cdots \rhd \sigma_m \lhd \tau_m,$$

we have $b(\sigma_i) \geq b(\sigma_j)$ for all $i \leq j$.

PROOF. For each $i \in \{1, \ldots, m\}$ we have an equality $b(\sigma_i) = b(\tau_i)$ by the $\boldsymbol{F}_\bullet$-compatibility of Σ and an inequality $b(\tau_i) \geq b(\sigma_{i+1})$ by the monotonicity of $b : K \to \mathbb{R}$. $\square$

This elementary observation has some wonderful consequences when it comes to simplifying computations of persistent homology. For each $t \in \mathbb{R}_+$, let $\Sigma_t \subset \Sigma$ be the restriction of Σ to (pairs which lie in) the subcomplex $\boldsymbol{F}_t K \subset K$, and let $(M^t_\bullet, d^t_\bullet)$ be shorthand for the affiliated Morse complex $(\mathbf{C}^{\Sigma_t}_\bullet(\boldsymbol{F}_t K), \partial^{\Sigma_t}_\bullet)$.

COROLLARY 8.14. *For each pair $0 \leq t \leq s$ of real numbers, there is an inclusion $(M^t_\bullet, d^t_\bullet) \hookrightarrow (M^s_\bullet, d^s_\bullet)$ of Morse chain complexes.*

PROOF. The critical simplices in $\boldsymbol{F}_t K$ remain critical in $\boldsymbol{F}_s K$, so M^t_k is naturally a subspace of M^s_k for all $k \geq 0$. Thus, it suffices to check that the Morse boundary operator d^s_k equals d^t_k when restricted to the subspace M^t_k. But this follows directly from the formula (8) — consider a Σ-critical k-simplex $\alpha \in \boldsymbol{F}_t K$, and a Σ-path of the form

$$\rho = (\sigma_1 \lhd \tau_1 \rhd \cdots \rhd \sigma_m \lhd \tau_m)$$

so that $\alpha \rhd \sigma_1$. By the monotonicity of b, we have $t \geq b(\alpha) \geq \sigma_1$. Now Proposition 8.13 guarantees that all subsequent Σ-paired simplices $\sigma_i \lhd \tau_i$ appearing in ρ must have b-values bounded above by t. In particular, adding new simplices from $(\boldsymbol{F}_s - \boldsymbol{F}_t)$ can not possibly change the Σ-paths over which we sum when evaluating the Morse boundary of α in M^s_k, whence $d^s_k(\alpha) = d^t_k(\alpha)$ as desired. $\square$

Having found a nested sequence of Morse complexes, one seeks to relate persistent homology groups of $\mathbf{H}_k(\boldsymbol{F}_\bullet K)$ to those of $\mathbf{H}_k(M^\bullet, d^\bullet)$. The basic idea, as one might expect, is to unite all the chain homotopy equivalences $\{\psi_t, \phi_t \mid t \geq 0\}$ promised by Proposition 8.10 between $\mathbf{C}_\bullet(\boldsymbol{F}_t K)$ and M^t for each $t \geq 0$ into a single equivalence relating the two persistence modules.

THEOREM 8.15. *For each dimension $k \geq 0$ and pair of real numbers $0 \leq t \leq s$, there are isomorphisms*
$$\mathbf{PH}_{t\to s}\mathbf{H}_k(\boldsymbol{F}_\bullet K) \simeq \mathbf{PH}_{t\to s}\mathbf{H}_k(M^\bullet, d^\bullet)$$
of persistent homology groups. Therefore, the barcodes of $\mathbf{H}_k(\boldsymbol{F}_\bullet K)$ and $\mathbf{H}_k(M^\bullet, d^\bullet)$ are equal.

PROOF. Enumerate all the simplex-pairs in Σ according to their b-values, i.e., write
$$\Sigma = \{(\sigma_1 \lhd \tau_1), (\sigma_2 \lhd \tau_2), \ldots, (\sigma_m \lhd \tau_m)\}$$
so that $b(\sigma_i) \leq b(\sigma_j)$ whenever $i \leq j$. Applying Lemma 8.11 to the Σ-pairs in this order, we obtain a family of chain homotopy equivalences indexed by $t \geq 0$:
$$\psi^t_\bullet : \mathbf{C}_\bullet(\boldsymbol{F}_t K) \to M^t_\bullet \quad \text{and} \quad \phi^t_\bullet : M^t_\bullet \to \mathbf{C}_\bullet(\boldsymbol{F}_t K)$$
which fit into a commuting diagram with the natural inclusion maps. Namely, for any pair of positive real numbers $t \leq s$ and dimension $k \geq 0$, the following diagrams of vector spaces commute:
$$\begin{tikzcd} \mathbf{C}_k(\boldsymbol{F}_t K) \arrow[r, hook] \arrow[d, "\psi^t_k"'] & \mathbf{C}_k(\boldsymbol{F}_s K) \arrow[d, "\psi^s_k"] \\ M^t_k \arrow[r, hook] & M^s_k \end{tikzcd} \qquad \begin{tikzcd} \mathbf{C}_k(\boldsymbol{F}_t K) \arrow[r, hook] & \mathbf{C}_k(\boldsymbol{F}_s K) \\ M^t_k \arrow[r, hook] \arrow[u, "\phi^t_k"] & M^s_k \arrow[u, "\phi^s_k"'] \end{tikzcd}$$
Since ψ^t and ϕ^t form two halves of a chain homotopy equivalence, they induce isomorphisms on k-th homology for all $k \geq 0$. Thus, we obtain a 0-interleaving between the two k-th homology persistence modules, which guarantees that all their persistent homology groups are isomorphic. □

From the perspective of using this result to simplify computations, it is important to note that large $\boldsymbol{F}_\bullet$-compatible partial matchings can only be found on filtrations where lots of simplices share the same b-values. Fortunately, this requirement is always satisfied by the Vietoris-Rips filtration. Consider a collection of points $P = \{p_0, \ldots, p_k\}$ so that the largest pairwise distance $d(p_i, p_j)$ equals $t' > 0$, corresponding to a single edge (p_i, p_j). Then the set of simplices born at this scale t' in $\mathbf{VR}_\bullet(P)$ include not only our edge, but also every other simplex containing this edge in its boundary.

8.5 FOR SHEAVES

Aside from the usual cognitive dissonance caused by reversing arrows when transitioning from homology to cohomology, there are not too many obstacles involved in using acyclic partial matchings to simplify sheaf cohomology computations. Let $\mathscr{S}$ be a sheaf (see Definition 7.1) on a simplicial complex K.

DEFINITION 8.16. An acyclic partial matching Σ on K is $\mathscr{S}$**-compatible** if the restriction map $\mathscr{S}(\sigma \leq \tau)$ is an isomorphism for every pair $(\sigma \lhd \tau)$ in Σ.

The weights of gradient paths from Definition 8.6 must now be upgraded from scalars to linear maps. It will be convenient, for simplices α, β in K, to define the *scaled restriction map*

$\mathscr{S}_{\alpha,\beta} : \mathscr{S}(\alpha) \to \mathscr{S}(\beta)$ as

$$\mathscr{S}_{\alpha,\beta} = [\beta : \alpha] \cdot \mathscr{S}(\alpha \leq \beta) = \begin{cases} +\mathscr{S}(\alpha \leq \beta) & \alpha = \beta_{-i} \text{ for even } i, \\ -\mathscr{S}(\alpha \leq \beta) & \alpha = \beta_{-i} \text{ for odd } i, \\ 0 & \text{otherwise.} \end{cases}$$

This linear map forms the block in α's column and β's row in the coboundary operator $\partial^{\mathscr{S}}_{\bullet}$ from Definition 7.6. For each Σ-path

$$\rho = (\sigma_1 \lhd \tau_1 \rhd \sigma_2 \lhd \tau_2 \rhd \cdots \rhd \sigma_m \lhd \tau_m);$$

define the $\mathscr{S}$**-weight** $w_{\mathscr{S}}(\rho)$ to be the composite linear map $\mathscr{S}(\tau_m) \to \mathscr{S}(\sigma_1)$ given by

$$(-1)^m \cdot \left[\mathscr{S}^{-1}_{\sigma_1,\tau_1} \circ \mathscr{S}_{\sigma_2,\tau_1} \circ \mathscr{S}^{-1}_{\sigma_2,\tau_2} \circ \cdots \circ \mathscr{S}^{-1}_{\sigma_m,\tau_m}\right].$$

Unsurprisingly, these $\mathscr{S}$-weights make an appearance when defining the Morse complex of Σ with $\mathscr{S}$-coefficients.

DEFINITION 8.17. Let $\mathscr{S}$ be a sheaf over the simplicial complex K and Σ an $\mathscr{S}$-compatible acyclic partial matching. The **Morse complex of Σ with coefficients in $\mathscr{S}$** is a cochain complex

$$\left(\mathbf{C}^{\bullet}_{\Sigma}(K;\mathscr{S}), \partial^{\bullet}_{\mathscr{S},\Sigma}\right)$$

defined as follows. For each dimension $k \geq 0$,

(1) the vector space $\mathbf{C}^k_{\Sigma}(K;\mathscr{S})$ equals the product of stalks $\prod_\alpha \mathscr{S}(\alpha)$ where α ranges over the k-dimensional Σ-critical simplices, and
(2) the linear map $\partial^k_{\mathscr{S},\Sigma} : \mathbf{C}^k_{\Sigma}(K;\mathscr{S}) \to \mathbf{C}^{k+1}_{\Sigma}(K;\mathscr{S})$ is represented by a block-matrix whose entry in α's column and ω's row equals

$$\partial^k_{\mathscr{S},\Sigma}\Big|_{\alpha,\omega} = \mathscr{S}_{\alpha,\omega} + \sum_{\rho} \mathscr{S}_{\sigma_\rho,\omega} \circ w_{\mathscr{S}}(\rho) \circ \mathscr{S}_{\alpha,\tau_\rho},$$

where ρ ranges over all the Σ-paths.

The fact that this definition actually produces a cochain complex follows from arguments analogous to the ones which we used in the proof of Proposition 8.8; the most significant difference is that unlike scalars of the form $[\alpha : \omega]$ used throughout that proof, the linear maps $\mathscr{S}_{\alpha,\omega}$ do not (necessarily) commute with each other.

Similarly, all the results of Section 3 admit direct generalizations to the sheafy context, with two caveats. First, we are working with cohomology rather than homology, so the boundary matrix is transposed. And second, we are working with an arbitrary sheaf, so the coboundary matrix is populated by block sub-matrices rather than scalar entries. For each $(\sigma \lhd \tau)$ in Σ, the motivating picture is provided by the usual matrix representation of the coboundary $\partial^{\dim \sigma}_{\mathscr{S}}$:

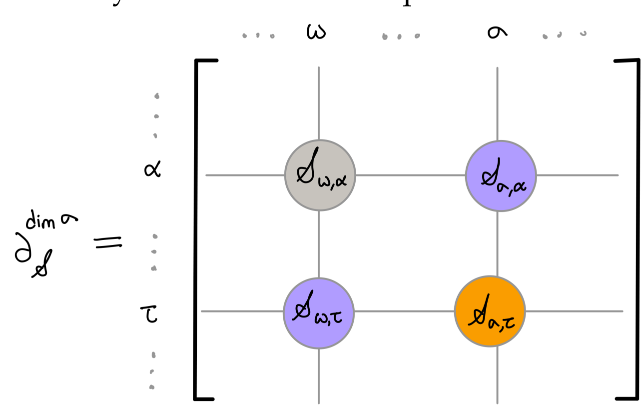

From this picture, one can discover the row and column operations that are required to turn the (invertible!) block $\mathscr{S}_{\sigma,\tau}$ into a pivot, and hence deduce the cochain homotopy equivalences which form counterparts of the maps ψ and ϕ from Lemma 8.11. Here is the aftermath.

THEOREM 8.18. *Let $\mathscr{S}$ be a sheaf on a simplicial complex K and let Σ be a $\mathscr{S}$-compatible acyclic partial matching on K. Then for each dimension $k \geq 0$, the sheaf cohomology group $\mathbf{H}^k(K;\mathscr{S})$ is isomorphic to the k-th cohomology group of the Morse cochain complex $(\mathbf{C}^\bullet_\Sigma(K;\mathscr{S}), \partial^\bullet_{\mathscr{S},\Sigma})$.*

The advantage of using the Morse complex in practice for computing sheaf cohomology is that it tends to be much smaller, since the cochain groups are built using stalks of the critical simplices (rather than all simplices). On the other hand, the compatibility requirement on Σ is quite severe — to find large acyclic partial matchings which happen to be compatible with a sheaf, we require the presence of many simplex-pairs $(\sigma \lhd \tau)$ for which the associated restriction map is invertible.

EXERCISES

EXERCISE 8.1. Let Σ be an acyclic partial matching on a simplicial complex K. Show that the Euler characteristic of K is given by

$$\chi(K) = \sum_{k=0}^{\dim K} (-1)^k \cdot m_k,$$

where m_k is the number of k-dimensional Σ-critical simplices.

EXERCISE 8.2. Write down all the gradient paths between critical simplices in Example 8.9 and confirm that the Morse chain complex has zero boundary operators over $\mathbb{Z}/2$.

EXERCISE 8.3. When not functioning as an occult symbol, the **Petersen graph** serves as the source of many counterexamples in graph theory.

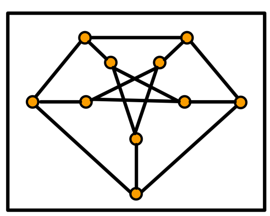

Impose an acyclic partial matching on this graph and use it to compute the homology groups over $\mathbb{Z}/2$ without performing any matrix operations.

EXERCISE 8.4. Show that the maps $\psi_\bullet : \mathbf{C}_\bullet(K) \to \mathbf{C}^\Sigma_\bullet(K)$ defined in (11) form a chain map .

EXERCISE 8.5. Show that the maps $\phi_\bullet : \mathbf{C}^\Sigma_\bullet(K) \to \mathbf{C}_\bullet(K)$ defined in (12) form a chain map.

EXERCISE 8.6. Show that the maps $\theta_k : \mathbf{C}_k(K) \to \mathbf{C}_{k+1}(K)$ from (13) serve as a chain homotopy between $\phi_\bullet \circ \psi_\bullet$ and the identity chain map on $\mathbf{C}_k(K)$.

EXERCISE 8.7. Verify that the two diagrams in the proof of Theorem 8.15 actually commute.

EXERCISE 8.8. State and prove a version of Lemma 8.11 in the context of a sheaf $\mathscr{S}$ on a simplicial complex K equipped with an $\mathscr{S}$-compatible acyclic partial matching Σ.