1. Simplicial Complexes

Here we introduce elementary concepts of algebraic topology indispensable for the subsequent chapters, most notably geometric and abstract simplicial complexes, homotopy, and homotopic equivalence of spaces.

Most of this material is usually covered in introductory courses on algebraic topology. But our presentation may deviate from others in details of notation and terminology, and it also includes some less commonly treated results. So even those fluent in algebraic topology may want to go through the chapter quickly.

The central notion for us is *simplicial complex*, which provides a link from combinatorics to topology. It can be viewed as a purely combinatorial object, namely, a hereditary set system. But it also describes a continuous object: a topological space. Many kinds of combinatorial objects—graphs, hypergraphs, partitions, and so on—can be associated with hereditary set systems, sometimes even in several natural ways. By viewing these hereditary set systems as simplicial complexes, we also assign topological spaces to the considered combinatorial objects. These spaces can be studied by methods of algebraic topology, and their topological properties are often related to combinatorial properties of the original object in interesting ways. Of course, creating simplicial complexes at every possible occasion is no panacea, but sometimes it does lead to meaningful results.

1.1 Topological Spaces

Although this may be unnecessary for most readers, we first review a few concepts from general topology. We begin with recalling the definition of a topological space, which is a mathematical structure capturing the notions of "nearness" and "continuity" on a very general level.

1.1.1 Definition. *A* **topological space** *is a pair* $(X, \mathcal{O})$*, where* X *is a (typically infinite) ground set and* $\mathcal{O} \subseteq 2^X$ *is a set system, whose members are called the* **open sets***, such that* $\emptyset \in \mathcal{O}$*,* $X \in \mathcal{O}$*, the intersection of finitely many open sets is an open set, and so is the union of an arbitrary collection of open sets.*

For example, the standard topology of the real line $\mathbb{R}$ or, more generally, of $\mathbb{R}^d$, is usually taught to freshmen: The open sets in $\mathbb{R}^d$ are defined, although one does not necessarily speak of topology. Namely, a set $U \subseteq \mathbb{R}^d$ is open exactly if for every point $\boldsymbol{x} \in U$ there exists an $\varepsilon > 0$ such that the ε-ball around $\boldsymbol{x}$ is contained in U. The same definition applies for any *metric space*, which for many readers may also be a notion more familiar and more intuitive than a topological space.

The theory dealing with topological spaces in general, *point-set topology* or *general topology*, often investigates fairly exotic examples. However, in our text, as well as in most of algebraic topology, one deals only with topological spaces that are subspaces of some $\mathbb{R}^d$, or at least can be identified with such subspaces.

What is a subspace? Let $(X, \mathcal{O})$ be a topological space. Every subset $Y \subseteq X$ defines a subspace, namely, the topological space $(Y, \{U \cap Y : U \in \mathcal{O}\})$.

For example, let $Y \subseteq \mathbb{R}^d$ be an arbitrary set. What are the open sets in the topology of the subspace defined by Y? They are exactly the intersections of open sets in $\mathbb{R}^d$ with Y; note that they need *not* be open as subsets of $\mathbb{R}^d$ (take Y as a closed segment in $\mathbb{R}^2$, for example).

Conventions. In the formulation of some topological definitions and theorems, it would be artificial to restrict our attention to subspaces of Euclidean spaces. But everywhere we assume that the considered spaces are (at least) *Hausdorff*, meaning that for every two distinct points $x, y \in X$ there are disjoint open sets U, V with $x \in U$ and $y \in V$.

Let us remark that if X is a set and the topology on X is understood, say if $X \subseteq \mathbb{R}^d$ and X is considered with the subspace topology, one usually does not mention the topology in the notation and writes "topological space X" even when formally X is only a set. We will also often say just "space" instead of "topological space."

Continuous maps. If $(X_1, \mathcal{O}_1)$ and $(X_2, \mathcal{O}_2)$ are topological spaces, a mapping $f: X_1 \to X_2$ is called *continuous* if preimages of open sets are open; i.e., $f^{-1}(V) \in \mathcal{O}_1$ for every $V \in \mathcal{O}_2$.

For mappings $\mathbb{R} \to \mathbb{R}$, for example, many readers may be accustomed to the "epsilon–delta" definition of continuity: For every $x \in \mathbb{R}$ and every $\varepsilon > 0$ there exists a $\delta > 0$ such that all points of the δ-neighborhood of x are mapped to the ε-neighborhood of $f(x)$. Or equivalently, for every sequence $x_1, x_2, x_3, \ldots$ converging to a limit a, we have $\lim_{n\to\infty} f(x_n) = f(a)$. Such readers may rest assured that these definitions of continuity are equivalent to the general one given above (for mappings $\mathbb{R} \to \mathbb{R}$, or more generally, for mappings between metric spaces). Or instead of resting, they may also want to prove it.

Convention: all maps are continuous. We implicitly assume that *all considered mappings between topological spaces are continuous*, although we do not always explicitly say so. More precisely, this applies to unspecified mappings in statements like, “Let $f: S^n \to \mathbb{R}^n$ be a mapping....” Sometimes, of course, after having constructed some mapping, we have to verify its continuity.

“The same” topological spaces: homeomorphism. As was remarked above, the topology of $\mathbb{R}^d$ is induced by the usual Euclidean metric, so why speak about topology? In the considerations of algebraic topology, the metric plays only an auxiliary role; often it is a convenient tool, but ultimately it is only the topology of a space that really matters. Two spaces that look metrically quite different can be topologically the same. An example is the real line $\mathbb{R}$ and the open interval $(0, 1)$.

The notion of “being the same” for topological spaces is similar to many other mathematical structures, such as groups, rings, and graphs. For most mathematical structures, one speaks about isomorphism, which is a bijective mapping preserving the considered structure (group or ring operations, graph edges, etc.). For topological spaces, the corresponding notion is traditionally called a *homeomorphism*.

1.1.2 Definition. *A* **homeomorphism** *of topological spaces* $(X_1, \mathcal{O}_1)$ *and* $(X_2, \mathcal{O}_2)$ *is a bijection* $\varphi: X_1 \to X_2$ *such that for every* $U \subseteq X_1$, $\varphi(U) \in \mathcal{O}_2$ *if and only if* $U \in \mathcal{O}_1$. *In other words, a bijection* $\varphi: X_1 \to X_2$ *is a homeomorphism if and only if both* φ *and* φ^{-1} *are continuous.*

(Warning: There are examples of continuous bijections for which the inverse mapping is not continuous, so both the continuity of φ and the continuity of φ^{-1} need checking in general.)

If X and Y are topological spaces and there is a homeomorphism $X \to Y$, we write $X \cong Y$ (read “X is homeomorphic to Y”).

Closure, boundary, interior. A set F in a topological space X is *closed* iff $X \setminus F$ is open. The *closure* of a set $Y \subseteq X$ in X, denoted by $\mathrm{cl}_X Y$, is the intersection of all closed sets in X containing Y (the subscript X is omitted if X is understood). For $Y \subseteq X = \mathbb{R}^d$, we have $\mathrm{cl}\, Y = \{\boldsymbol{x} \in \mathbb{R}^d: \mathrm{dist}(\boldsymbol{x}, Y) = 0\}$, where $\mathrm{dist}(\boldsymbol{x}, Y) := \inf\{\|\boldsymbol{x} - \boldsymbol{y}\| : \boldsymbol{y} \in Y\}$. The *boundary* of Y is $\partial Y := \{\mathrm{cl}\,(Y) \cap \mathrm{cl}\,(X \setminus Y)\}$ and the *interior* $\mathrm{int}\, Y := Y \setminus \partial Y$.

Compactness. We conclude this nano-course on general topology by recalling compactness. A space $X \subseteq \mathbb{R}^d$ is compact if and only if X is a closed and bounded set. (In general, a topological space X is compact if for every collection $\mathcal{U}$ of open sets with $\bigcup \mathcal{U} = X$, there exists a finite subcollection $\mathcal{U}_0 \subseteq \mathcal{U}$ with $\bigcup \mathcal{U}_0 = X$.) In a compact metric space, every infinite sequence has a convergent subsequence.

If X is a compact space and $f: X \to \mathbb{R}$ is a continuous real function, then f attains its minimum; that is, there is an $x \in X$ with $f(x) \le f(y)$ for all $y \in X$. Moreover, a continuous function on a compact metric space is *uniformly continuous*; that is, for every $\varepsilon > 0$ there is a $\delta > 0$ such that any two points at distance at most δ are mapped to points at distance at most ε.

Notes. Among many textbooks of topology, we mention Munkres [Mun00], which deals both with general topology and with elements of algebraic topology. A large menagerie of topological spaces is collected in [SS78].

Exercises

1. Verify the following homeomorphisms:
(a) $\mathbb{R} \cong (0,1) \cong (S^1 \setminus \{(0,1)\})$;
(b) $S^1 \cong \partial([0,1]^2)$.
2. (a) Let X and Y be topological spaces. Check that a mapping $f\colon X \to Y$ is continuous if and only if $f^{-1}(F)$ is closed for every closed set $F \subseteq Y$.
(b) Let X be covered by finitely many closed sets $A_1, A_2, \ldots, A_n$ (i.e., $X = A_1 \cup A_2 \cup \cdots \cup A_n$), and let $f\colon X \to Y$ be a mapping whose restriction to each A_i is continuous. Verify that f is continuous.

1.2 Homotopy Equivalence and Homotopy

In algebraic topology, two spaces are considered "the same" under an equivalence relation even coarser than homeomorphism. This notion is called homotopy equivalence. Similarly, continuous maps are classified into classes according to so-called homotopy.

Deformation retract. Before plunging into subtleties of homotopy equivalence, we introduce the perhaps more intuitive notion of deformation retract. The figure 8 below drawn by the thick line is a deformation retract of the gray area with two holes:

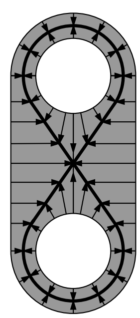

This means that the gray area can be continuously shrunk to the figure 8 while keeping the points of the 8 fixed. The motion is shown by arrows: Each point moves in the indicated direction at uniform speed until it hits the 8, where it stops. In general, if X is a space and $Y \subseteq X$ a subspace of it, a *deformation retraction* of X onto Y is a family $\{f_t\}_{t\in[0,1]}$ of continuous maps $f_t\colon X \to X$ (we can think of t as time), such that

- f_0 is the identity map on X,
- $f_t(y) = y$ for all $y \in Y$ and all $t \in [0,1]$ (Y remains stationary), and
- $f_1(X) = Y$.

Moreover, the mappings should depend continuously on t. That is, if we define the mapping $F: X \times [0,1]$ by $F(x,t) = f_t(x)$, this mapping should be continuous. Explicitly, this means that if we choose $x \in X$, $t \in [0,1]$, and an arbitrarily small neighborhood V of $F(x,t)$, there are $\delta > 0$ and a neighborhood U of x such that $F(x',t') \in V$ for all $x' \in U$ and all $t' \in (t+\delta, t-\delta)$. In most of the literature, a deformation retraction is formally viewed as the mapping F, rather than a family of maps; we will use both of these presentations interchangeably.

If a deformation retraction as above exists, Y is called a *deformation retract* of X.

The intuition for deformation retraction, that X can be continuously shrunk to Y, has to be used with some care. Namely, the shrinking motion has to take place *within* X. One can think of an "old" copy of X, which is solid and remains motionless, and a "new" elastic copy of X, which shrinks within the old copy. (It seems quite tempting, for X sitting in $\mathbb{R}^d$, to imagine a motion in the ambient space, rather than within X, but this is wrong.)

Homotopy equivalence. If Y is a deformation retract of X, then X and Y are homotopy equivalent. But obviously, being a deformation retract is not an equivalence relation. As the following picture illustrates, one space can have several rather different-looking deformation retracts:

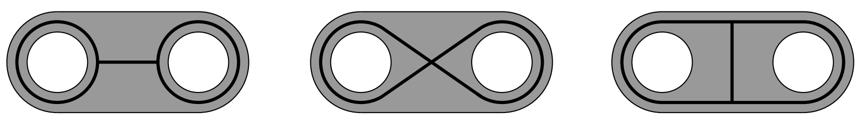

Homotopy equivalence can be introduced as follows: Spaces X and Y are homotopy equivalent, in symbols $X \simeq Y$, iff there exists a space Z such that both X and Y are deformation retracts of Z. For example, the three spaces drawn by the thick line are all homotopy equivalent.

The usual definition of homotopy equivalence is different; it is technically more convenient but perhaps less intuitive. To state it, we first need to introduce homotopy of maps.

1.2.1 Definition. *Two continuous maps $f, g: X \to Y$ are* **homotopic** *(written $f \sim g$) if there is a "continuous interpolation" between them; that is, a family $\{f_t\}_{t\in[0,1]}$ of maps $f_t: X \to Y$ depending continuously on t (i.e., the associated bivariate mapping $F(x,t) := f_t(x)$ is a continuous map $X \times [0,1] \to Y$, similar to deformation retraction above) such that $f_0 = f$ and $f_1 = g$.*

In particular, a map $X \to Y$ is called *nullhomotopic* if it is homotopic to a constant map that maps all of X to a single point $y_0 \in Y$ (so "nullhomotopic" is a misnomer; it would be more logical to say "constant-homotopic," but we stick to the traditional terminology). It is not hard to verify that "being homotopic" is an equivalence on the set of all continuous maps $X \to Y$.

1.2.2 Definition (Homotopy equivalence). *Two spaces X and Y are homotopy equivalent (or have the same* **homotopy type***) if there are continuous maps $f: X \to Y$ and $g: Y \to X$ such that the composition $f \circ g: Y \to Y$ is homotopic to the identity map id_Y and $g \circ f \sim \mathrm{id}_X$.*

The equivalence of this definition to the characterization above (homotopy equivalent spaces are deformation retracts of the same space) is nontrivial; see, e.g., [Hat01, Chapter 0].

A space that is homotopy equivalent to a single point is called *contractible*. Some spaces are "obviously" contractible, such as the ball B^d, but for others, contractibility is not easy to visualize. A beautiful example of this is "Bing's house"; see [Hat01, Chapter 0] for a nice presentation. It is tempting to think that a contractible space can always be deformation-retracted to a point, but this is false in general (it can happen that all points are forced to move during any contraction; see Exercise 7).

The task of determining whether two given spaces are homotopy equivalent is in general very difficult. Without a sophisticated technical apparatus, it is quite hard to prove even "obvious" facts such as that the circle S^1 is not contractible. But the spaces arising in many topological proofs of combinatorial or geometric theorems happen to be relatively simple, and often they turn out to be homotopy equivalent to a sphere.

Exercises

1. Show that the dumbbell 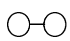 and the letter θ are homotopy equivalent, using Definition 1.2.2 (exhibit suitable mappings f and g).
2. Verify that if spaces X and Y are both deformation retracts of the same space Z, then X and Y are homotopy equivalent.
3. Take a 2-dimensional sphere (in $\mathbb{R}^3$) and connect the north and south poles by a segment, obtaining a space X. Let Y be a 2-dimensional sphere with a circle attached by one point to the north pole of the sphere. Show that $X \simeq Y$ (using both of the definitions of homotopy equivalence given in the text).
4. Consider two embeddings f and g of the circle S^1 into $\mathbb{R}^3$, where f just inserts the circle into $\mathbb{R}^3$ without changing its shape while g maps it to the trefoil knot:

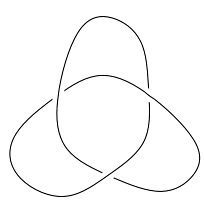

Are f and g homotopic? Substantiate your answer at least informally.

5. (a) Prove that homotopy is an equivalence relation on the set of all continuous maps $X \to Y$.
   (b) Prove that homotopy equivalence is indeed an equivalence relation on the class of all topological spaces (check transitivity).
6. (a) Prove that a space X is contractible if and only if for every space Y and every continuous map $f: X \to Y$, f is nullhomotopic.
   (b) Prove that a space X is contractible if and only if for every space Y and every continuous map $f: Y \to X$, f is nullhomotopic.
7.* The ***topologist's comb*** is the subspace $X := (R\times[0,1])\cup([0,1]\times\{0\})$ of $\mathbb{R}^2$, where R denotes the set of all rational numbers in the interval $[0,1]$. (Here $\mathbb{R}^2$ is taken with the usual topology and X has the subspace topology.) Let Y be made of countably many copies of X arranged in a zigzag fashion into a doubly infinite chain:

   ... 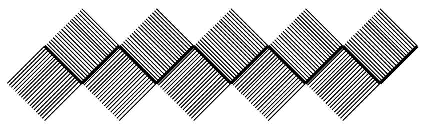 ...

   Show that Y is contractible.
   It can be proved that no point is a deformation retract of Y (you may want to try this as well). In $\mathbb{R}^3$, one can even construct a contractible compact Y with this property; see the exercises to Chapter 0 in Hatcher [Hat01].

1.3 Geometric Simplicial Complexes

Many topologically interesting subspaces of $\mathbb{R}^d$ can be described as simplicial complexes. This means that they are pasted together from simple building blocks, called simplices and including segments, triangles, and tetrahedra, in a way respecting simple rules. As we will see later, simplicial complexes have a purely combinatorial description, and they are particularly significant in the interplay of topology and combinatorics.

First we need to introduce affine independence and simplices.

1.3.1 Definition. *Let $\boldsymbol{v}_0, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$ be points in $\mathbb{R}^d$. We call them* **affinely dependent** *if there are real numbers $\alpha_0, \alpha_1, \ldots, \alpha_k$, not all of them 0, such that $\sum_{i=0}^k \alpha_i \boldsymbol{v}_i = \boldsymbol{0}$ and $\sum_{i=0}^k \alpha_i = 0$. Otherwise, $\boldsymbol{v}_0, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$ are called* **affinely independent**.

For $k = 2$, affine independence simply means $\boldsymbol{v}_0 \neq \boldsymbol{v}_1$; for $k = 3$ it means that $\boldsymbol{v}_0, \boldsymbol{v}_1, \boldsymbol{v}_2$ do not lie on a common line; for $k = 4$ it means that $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_3$ do not lie on a common plane; and so on.

Here are two further simple but useful characterizations of affine independence.

1.3.2 Lemma. *Both of the following conditions are equivalent to affine independence of points* $\boldsymbol{v}_0, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_k \in \mathbb{R}^d$*:*

- *The* k *vectors* $\boldsymbol{v}_1 - \boldsymbol{v}_0, \boldsymbol{v}_2 - \boldsymbol{v}_0, \ldots, \boldsymbol{v}_k - \boldsymbol{v}_0$ *are linearly independent.*
- *The* $(d+1)$*-dimensional vectors* $(1, \boldsymbol{v}_0), (1, \boldsymbol{v}_1), \ldots, (1, \boldsymbol{v}_k) \in \mathbb{R}^{d+1}$ *are linearly independent.*

We leave the easy proof as a warmup exercise. We also note that $d+1$ is the largest size of an affinely independent set of points in $\mathbb{R}^d$.

Simplices. Here are examples of simplices: a point, a line segment, a triangle, and a tetrahedron:

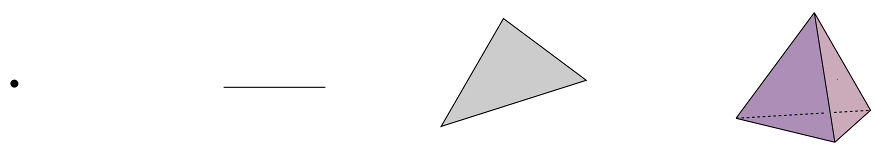

These examples have dimensions 0, 1, 2, and 3, respectively.

1.3.3 Definition. *A* **simplex** σ *is the convex hull of a finite affinely independent set* A *in* $\mathbb{R}^d$. *The points of* A *are called the* **vertices** *of* σ. *The* **dimension** *of* σ *is* $\dim \sigma := |A| - 1$. *Thus every* $\boldsymbol{k}$**-simplex** *(*k*-dimensional simplex) has* $k+1$ *vertices.*

1.3.4 Definition. *The convex hull of an arbitrary subset of vertices of a simplex* σ *is a* **face** *of* σ *(this is a special case of the definition of a face of a convex polytope). Thus every face is itself a simplex.*

The **relative interior** *of a simplex* σ *arises from* σ *by removing all faces of dimension smaller than* $\dim \sigma$.

For illustration, we count the faces of a triangle: the whole triangle, 3 edges, 3 vertices, and the empty set; altogether we have 8 faces.

Every simplex is a *disjoint* union of the relative interiors of its faces. Thus we get a (closed) triangle as a union of its relative interior (i.e., an open triangle), 3 open line segments (the edges without their endpoints), and 3 vertices.

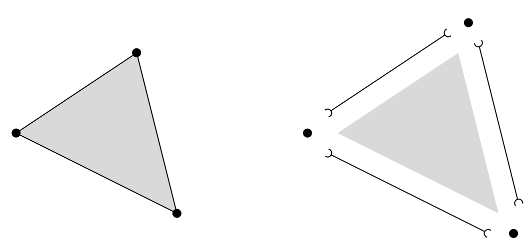

Here are the simple rules for putting simplices together to form a simplicial complex.

1.3.5 Definition. *A nonempty family* Δ *of simplices is a* **simplicial complex** *if the following two conditions hold:*

(1) *Each face of any simplex* $\sigma \in \Delta$ *is also a simplex of* Δ.
(2) *The intersection* $\sigma_1 \cap \sigma_2$ *of any two simplices* $\sigma_1, \sigma_2 \in \Delta$ *is a face of both* σ_1 *and* σ_2.

The union of all simplices in a simplicial complex Δ *is the* **polyhedron** *of* Δ *and is denoted by* $\|\Delta\|$.

The **dimension** *of a simplicial complex is the largest dimension of a simplex:* $\dim \Delta := \max\{\dim \sigma : \sigma \in \Delta\}$.

The **vertex set** *of* Δ, *denoted by* $V(\Delta)$, *is the union of the vertex sets of all simplices of* Δ.

In particular, note that every simplicial complex contains the empty set as a face (this is different from what appears in some other sources, such as [Mun84] and [Bjö95], where the empty face is excluded!).

The simplicial complex that consists only of the empty simplex is defined to have dimension -1. Zero-dimensional simplicial complexes are just configurations of points, while 1-dimensional simplicial complexes correspond to graphs (represented geometrically with straight edges that do not cross). The following picture shows one 2-dimensional simplicial complex in the plane and two cases of putting simplices together in ways forbidden by the definition of a simplicial complex:

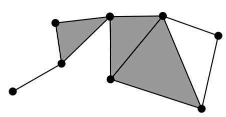

good

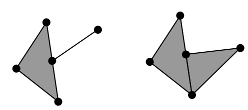

bad!!!

We are going to consider only *finite* simplicial complexes (with finitely many simplices). From the topological point of view, this is quite a restrictive assumption, since then the polyhedra are only compact spaces, and we cannot express, for example, the space $\mathbb{R}^d$ as the polyhedron of a finite simplicial complex. But finite simplicial complexes are sufficient for our combinatorial applications, and this assumption spares us some trouble (namely, of discussing too much point set topology).

Support. Just as in the case of a single simplex, the relative interiors of all simplices of a simplicial complex Δ form a partition of the polyhedron $\|\Delta\|$: For each point $\boldsymbol{x} \in \|\Delta\|$ there exists exactly one simplex $\sigma \in \Delta$ containing $\boldsymbol{x}$ in its relative interior. This simplex is denoted by $\operatorname{supp}(\boldsymbol{x})$ and called the *support* of the point $\boldsymbol{x}$.

It may seem obvious at this point that the set of all faces of a simplex forms a simplicial complex. Still, to be on the safe side, and for further use, we include a proof.

1.3.6 Lemma. *The set of all faces of a simplex is a simplicial complex.*

Proof. Let $V \subset \mathbb{R}^d$ be affinely independent and let $F, G \subseteq V$. It suffices to show that

$$\operatorname{conv}(F) \cap \operatorname{conv}(G) = \operatorname{conv}(F \cap G),$$

where $\operatorname{conv}(F) \cap \operatorname{conv}(G) \supseteq \operatorname{conv}(F \cap G)$ is trivial. We write $\boldsymbol{x} \in \operatorname{conv}(F) \cap \operatorname{conv}(G)$ as

$$\boldsymbol{x} = \sum_{\boldsymbol{u} \in F} \alpha_{\boldsymbol{u}} \boldsymbol{u} = \sum_{\boldsymbol{v} \in G} \beta_{\boldsymbol{v}} \boldsymbol{v},$$

with $\alpha_{\boldsymbol{u}}, \beta_{\boldsymbol{v}} \geq 0$ and $\sum_{\boldsymbol{u} \in F} \alpha_{\boldsymbol{u}} = 1 = \sum_{\boldsymbol{v} \in G} \beta_{\boldsymbol{v}}$. By subtracting we get

$$\sum_{\boldsymbol{u} \in F \setminus G} \alpha_{\boldsymbol{u}} \boldsymbol{u} - \sum_{\boldsymbol{v} \in G \setminus F} \beta_{\boldsymbol{v}} \boldsymbol{v} + \sum_{\boldsymbol{w} \in F \cap G} (\alpha_{\boldsymbol{w}} - \beta_{\boldsymbol{w}}) \boldsymbol{w} = \mathbf{0}.$$

The points in $F \cup G$ are affinely independent, and thus all coefficients on the left-hand side of this equation must be 0. In particular, $\alpha_{\boldsymbol{w}}, \beta_{\boldsymbol{w}}$ can be nonzero only for $\boldsymbol{w} \in F \cap G$, and thus $\boldsymbol{x} \in \operatorname{conv}(F \cap G)$.

A simplicial complex consisting of all faces of an arbitrary n-dimensional simplex (including the simplex itself) will be denoted by σ^n. Hence $\|\sigma^n\|$ is a (geometric) n-simplex.

The notion of subcomplex is defined as everyone would expect:

1.3.7 Definition. *A* **subcomplex** *of a simplicial complex Δ is a subset of Δ that is itself a simplicial complex (that is, it is closed under taking faces).*

An important example of a subcomplex is the *k-skeleton* of a simplicial complex Δ. It consists of all simplices of Δ of dimension at most k, and we denote it by $\Delta^{\leq k}$.

1.4 Triangulations

Let X be a topological space. A simplicial complex Δ such that $X \cong \|\Delta\|$, if one exists, is called a *triangulation* of X. We give a few examples.

The simplest triangulation of the sphere S^{n-1} is the boundary of an n-simplex, that is, the subcomplex of σ^n obtained by deleting the single n-dimensional simplex (but retaining all of its proper faces). Indeed, the boundary of an n-simplex is homeomorphic to S^{n-1}, as can be seen using the central projection:

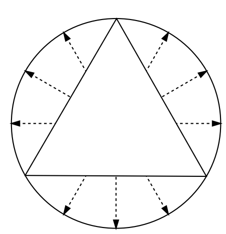

Other triangulations of spheres are obtained from convex polytopes. A convex polytope $P \subset \mathbb{R}^d$ is called *simplicial* if all of its proper faces, i.e., all faces except possibly for P itself, are simplices. For the familiar 3-dimensional convex polytopes, it means that all the 2-dimensional faces are triangles, as is the case for the regular octahedron or icosahedron. It can be shown without much difficulty that the set of all proper faces of any simplicial polytope P is a simplicial complex. Since the boundary ∂P is homeomorphic to S^{d-1} for every d-dimensional convex polytope P, we obtain various triangulations of the sphere in this way (although for $d > 3$, by far not all possible triangulations; see Section 5.6!).

Particularly nice and important symmetric triangulations of S^{d-1} are provided by crosspolytopes.

1.4.1 Definition. *The d-dimensional* **crosspolytope** *is the convex hull*

$$\operatorname{conv}\{\boldsymbol{e}_1, -\boldsymbol{e}_1, \ldots, \boldsymbol{e}_d, -\boldsymbol{e}_d\}$$

of the vectors of the standard orthonormal basis and their negatives:

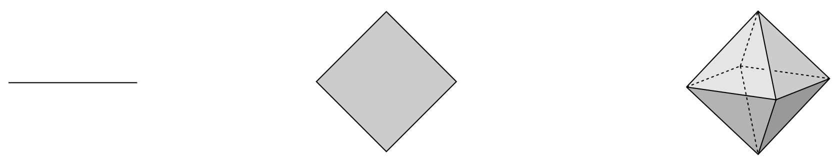

$d = 1$ $d = 2$ $d = 3$

Alternatively, it is the unit ball of the ℓ_1-norm: $\{\boldsymbol{x} \in \mathbb{R}^d : \|\boldsymbol{x}\|_1 \leq 1\}$.

It is not hard to show that a subset $F \subseteq \{\boldsymbol{e}_1, -\boldsymbol{e}_1, \ldots, \boldsymbol{e}_d, -\boldsymbol{e}_d\}$ forms the vertex set of a proper face of the crosspolytope if and only if there is no $i \in [d]$ with both $\boldsymbol{e}_i \in F$ and $-\boldsymbol{e}_i \in F$ (Exercise 2).

The next example is more sophisticated and surprising. Although we will not need it in the sequel, it is worth considering at least briefly.

1.4.2 Example (Cube triangulation). The cube $[0,1]^d$ can be triangulated as follows: Let S_d denote the set of all permutations of $[d]$, and for every $\pi \in S_d$, let $\sigma_\pi = \operatorname{conv}\{\mathbf{0}, \boldsymbol{e}_{\pi(1)}, \boldsymbol{e}_{\pi(1)} + \boldsymbol{e}_{\pi(2)}, \ldots, \boldsymbol{e}_{\pi(1)} + \cdots + \boldsymbol{e}_{\pi(d)}\}$. Each σ_π is a d-simplex, and all the σ_π together plus all of their faces form a triangulation of $[0,1]^d$. We leave the (somewhat laborious) verification as Exercise 4.

Another approach to triangulating the cube, involving a generally useful auxiliary construction, is outlined in Exercise 3.

Notes. To construct "suitable" triangulations of given geometric shapes is a major topic in many fields of applied mathematics, such as numerical analysis and computer aided design (CAD).

In contemporary algebraic topology, simplicial complexes are often considered old-fashioned. Spaces can usually be described much more economically if we allow for more general ways of gluing the basic building blocks together than is permitted in simplicial complexes. For example, the torus (also known, at least in the United States, as the surface of a doughnut) can be produced by a suitable gluing of the edges of a single square in $\mathbb{R}^3$,

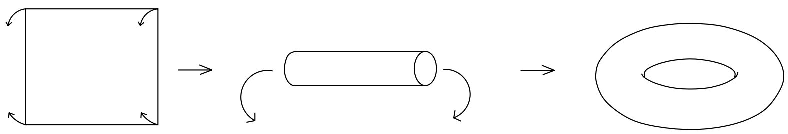

while a triangulation of the torus requires quite a number of simplices (Exercise 1). Moreover, there are quite "reasonable" spaces (4-dimensional manifolds) that cannot be triangulated at all, while they can be obtained using more general ways of gluing.

However, these more general ways of building spaces, most notably CW-complexes (discussed in Section 4.5), do not admit as direct a combinatorial interpretation as simplicial complexes do.

Exercises

1. Draw a triangulation of a torus. Use as few simplices as you can.
2. (a) Prove the claim about the faces of the crosspolytope below Definition 1.4.1 (use the definition of a polytope face mentioned in the Preliminaries).
   (b) Count the number of faces of each dimension.
3.* (Triangulation of a simplicial prism) Let σ be a simplex with vertices $\boldsymbol{v}_0, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_d$, and let $P = \sigma \times [0,1]$ be the $(d+1)$-dimensional "prism above σ."

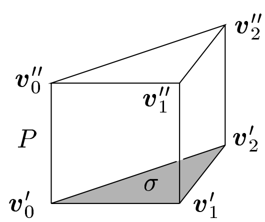

Let the vertices of P be $\boldsymbol{v}'_0, \boldsymbol{v}'_1, \ldots, \boldsymbol{v}'_d, \boldsymbol{v}''_0, \boldsymbol{v}''_1, \ldots, \boldsymbol{v}''_d$, where each $\boldsymbol{v}'_i$ is a bottom vertex and $\boldsymbol{v}''_i$ is the top vertex above it. For $i = 0, 1, 2, \ldots, d$, let σ_i be the simplex $\mathrm{conv}\{\boldsymbol{v}'_0, \boldsymbol{v}'_1, \ldots, \boldsymbol{v}'_i, \boldsymbol{v}''_i, , \boldsymbol{v}''_{i+1}, \ldots, \boldsymbol{v}''_d\}$ (we take the first $i+1$ of the bottom vertices and the last $d+1-i$ of the top vertices).
(a) Let $d = 2$; draw the simplices $\sigma_0, \sigma_1, \sigma_2$ and check that they triangulate P.
(b) Prove that $\sigma_0, \sigma_1, \ldots, \sigma_d$ are indeed $(d+1)$-dimensional simplices, they cover P, and they have disjoint interiors.

(c) Show that the σ_i and all of their faces form a simplicial complex.
(d) Let Δ be a simplicial complex with $\|\Delta\| \subseteq \mathbb{R}^d$. Describe how the above construction can be used to triangulate $\|\Delta\| \times [0,1]$. Explain how this construction, applied inductively, triangulates the cube $[0,1]^d$.
(e) Count the number of d-dimensional simplices in the inductive triangulation of the d-dimensional cube as in (d).

4.* This refers to the cube triangulation in Example 1.4.2.
(a) Check that each simplex σ_π is d-dimensional and can be written as $\sigma_\pi = \{\boldsymbol{x} \in [0,1]^d : x_{\pi(d)} \leq x_{\pi(d-1)} \leq \cdots \leq x_{\pi(1)}\}$. Conclude that $\bigcup_{\pi \in S_n} \sigma_\pi = [0,1]^d$.
(b) Let $\preceq$ be a *linear quasiordering* of $[d]$, i.e., a transitive relation in which every two numbers are comparable, $i \preceq j$ or $j \preceq i$ (but it may happen that both $i \preceq j$ and $j \preceq i$ even if $i \neq j$). Define $\sigma_{\preceq} := \{\boldsymbol{x} \in [0,1]^d : x_i \leq x_j \text{ whenever } i \preceq j\}$. Check that $\sigma_{\preceq}$ is a simplex, determine its dimension (in terms of $\preceq$), and describe its vertices.
(c) Show that the intersection $\sigma_{\preceq_1} \cap \sigma_{\preceq_2}$ is again of the form $\sigma_{\preceq}$ for a suitable linear quasiordering $\preceq$. How do we obtain $\preceq$ from $\preceq_1$ and $\preceq_2$?
(d) What are the faces of σ_π? Verify that the σ_π and their faces form a simplicial complex.
(e) Can this triangulation of the cube be obtained by the inductive procedure using Exercise 3(e)? Do we always obtain a triangulation as in Example 1.4.2 by that inductive procedure?
(f) Show that the copies of the triangulation in Example 1.4.2 translated by each integer vector in $\{0,1,\ldots,n-1\}^d$ form a triangulation of $[0,n]^d$.

1.5 Abstract Simplicial Complexes

We introduce a combinatorial object called an abstract simplicial complex. In order to distinguish it from the simplicial complex defined in Section 1.3, which is a geometric object, we will call the latter a *geometric simplicial complex*. However, this distinction will not be maintained for very long: Soon we will see that an abstract simplicial complex and a geometric simplicial complex are essentially two different descriptions of the same mathematical object. One can thus simply speak of a simplicial complex, and use both the combinatorial and geometric aspects as convenient.

1.5.1 Definition. *An* **abstract simplicial complex** *is a pair (V, K), where V is a set and $\mathsf{K} \subseteq 2^V$ is a hereditary system of subsets of V; that is, we require that $F \in \mathsf{K}$ and $G \subseteq F$ imply $G \in \mathsf{K}$ (in particular, $\varnothing \in \mathsf{K}$ whenever $\mathsf{K} \neq \varnothing$). The sets in K are called (abstract) simplices. Further, we define the* **dimension** $\dim(\mathsf{K}) := \max\{|F|-1 : F \in \mathsf{K}\}$.

Abstract simplicial complexes are denoted by sans-serif capital letters like $\mathsf{K}, \mathsf{L}, \mathsf{N}, \ldots$ in this book.

Usually we may assume that $V = \bigcup \mathsf{K}$; thus it suffices to write K instead of (V, K), where V is understood to equal $\bigcup \mathsf{K}$.

Each geometric simplicial complex Δ determines an abstract simplicial complex. The points of the abstract simplicial complex are all vertices of the simplices of Δ, so we set $V := V(\Delta)$, and the sets in the abstract simplicial complex are just the vertex sets of the simplices of Δ. The set system (V, K) obtained in this way is clearly an abstract simplicial complex. For example, for the geometric simplicial complex

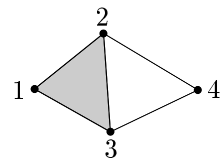

we have the abstract simplicial complex $\{\emptyset, \{1\}, \{2\}, \{3\}, \{4\}, \{1,2\}, \{1,3\}, \{2,3\}, \{2,4\}, \{3,4\}, \{1,2,3\}\}$.

In this situation we call Δ a *geometric realization* of K, and the polyhedron of Δ is also referred to as a *polyhedron of* K (soon we will see that a polyhedron of K is unique up to homeomorphism).

It is easy to see that any abstract simplicial complex (V, K) with V finite (which we always assume) has a geometric realization. Let $n := |V| - 1$ and let us identify V with the vertex set of an n-dimensional simplex σ^n. We define a subcomplex Δ of σ^n by $\Delta := \{\mathrm{conv}(F) : F \in \mathsf{K}\}$. This is a geometric simplicial complex, and its associated abstract simplicial complex is K. So every simplicial complex on $n+1$ vertices can be realized in $\mathbb{R}^n$ (later on, we will prove a much sharper result).

Simplicial mappings. Now we show that a geometric realization is unique up to homeomorphism. At this occasion we also introduce the very important notion of a simplicial mapping, which is a *combinatorial counterpart of a continuous mapping.*[1]

1.5.2 Definition. *Let* K *and* L *be two abstract simplicial complexes. A* **simplicial mapping** *of* K *into* L *is a mapping* $f\colon V(\mathsf{K}) \to V(\mathsf{L})$ *that maps simplices to simplices, i.e., such that* $f(F) \in \mathsf{L}$ *whenever* $F \in \mathsf{K}$.

A bijective simplicial mapping whose inverse mapping is also simplicial is called an **isomorphism** *of abstract simplicial complexes. The existence of an isomorphism of simplicial complexes* K *and* L *will be denoted by* $\mathsf{K} \cong \mathsf{L}$.

Isomorphic abstract simplicial complexes are thus "the same" set systems; they differ only in the names of the vertices. In the sequel, we will not usually distinguish among isomorphic simplicial complexes.

[1] In earlier days of algebraic topology, approximation of arbitrary continuous maps of spaces by simplicial maps of sufficiently fine triangulations was one of the main tools for converting topological statements into algebraic or combinatorial ones.

We also note that for an arbitrary simplicial mapping, a k-simplex in K can be mapped to a simplex of L of any dimension $\ell \leq k$.

With each simplicial mapping f of simplicial complexes we are going to associate a continuous mapping $\|f\|$ of their polyhedra. Namely, we extend f affinely on each simplex. To state this precisely, we first note that if $\sigma \subset \mathbb{R}^d$ is a k-simplex with vertices $\boldsymbol{v}_0, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$, then each point $\boldsymbol{x} \in \sigma$ can be uniquely written as a convex combination $\boldsymbol{x} = \sum_{i=0}^{k} \alpha_i \boldsymbol{v}_i$, where $\alpha_0, \ldots, \alpha_k \geq 0$ and $\sum_{i=0}^{k} \alpha_i = 1$. Indeed, at least one such convex combination exists because $\boldsymbol{x} \in \operatorname{conv}\{\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k\}$, and if there were two distinct convex combinations equal to $\boldsymbol{x}$, we would get a contradiction to the affine independence of $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k$ by subtracting them.

1.5.3 Definition. *Let Δ_1 and Δ_2 be geometric simplicial complexes, let K_1 and K_2 be their associated abstract simplicial complexes, and let $f\colon V(\mathsf{K}_1) \to V(\mathsf{K}_2)$ be a simplicial mapping of K_1 into K_2. We define the mapping*

$$\|f\|\colon \|\Delta_1\| \longrightarrow \|\Delta_2\|,$$

the **affine extension** *of f, by extending f affinely to the relative interiors of the simplices of Δ_1, as follows: If $\sigma = \operatorname{supp}(\boldsymbol{x}) \in \Delta_1$ is the support of $\boldsymbol{x}$, the vertices of σ are $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k$, and $\boldsymbol{x} = \sum_{i=0}^{k} \alpha_i \boldsymbol{v}_i$ with $\alpha_0, \ldots, \alpha_k \geq 0$ and $\sum_{i=0}^{k} \alpha_i = 1$, we put $\|f\|(\boldsymbol{x}) = \sum_{i=0}^{k} \alpha_i f(\boldsymbol{v}_i)$.*

First we note that the mapping $\|f\|$ is well-defined, because the set $\{f(\boldsymbol{v}_0), \ldots, f(\boldsymbol{v}_k)\}$ is always the vertex set of a simplex in Δ_2. With some more effort, one can check the following proposition, whose proof we omit.

1.5.4 Proposition. *For every simplicial mapping f as in Definition 1.5.3, $\|f\|$ is a continuous map $\|\Delta_1\| \to \|\Delta_2\|$. If f is injective, then $\|f\|$ is injective too, and if f is an isomorphism, then $\|f\|$ is a homeomorphism.*

In particular, this proposition shows that each (finite) abstract simplicial complex (V, K) defines a topological space uniquely up to homeomorphism.

Simplicial complexes: a connection between combinatorics and topology. We summarize the contents of the last three sections and add some remarks.

Every finite hereditary set system can be regarded as an abstract simplicial complex, and it specifies a topological space (the polyhedron of a geometric realization) up to homeomorphism. Simplicial maps of simplicial complexes yield continuous maps of the corresponding spaces.

Conversely, if a topological space admits a triangulation, it can be described purely combinatorially by an abstract simplicial complex. (This description is not unique.)

A continuous map, even between triangulated spaces, generally cannot be described by a simplicial map. On the other hand, there are theorems stating that under suitable conditions, a continuous map is homotopic to a simplicial map between sufficiently fine triangulations of the considered spaces, and it can be approximated by such simplicial maps with any prescribed precision; see [Mun84] or [Hat01]. We will not prove a general theorem of this kind (a *simplicial approximation theorem*), but we will encounter some special cases.

Convention. In the sequel, a simplicial complex will formally be understood as an abstract simplicial complex (i.e., it will be a set system as a mathematical object). But we will speak of a polyhedron $\|\mathsf{K}\|$ for an abstract simplicial complex K (which is well-defined up to homeomorphism in view of Proposition 1.5.4). We will even freely use topological notions such as "K is contractible" instead of "$\|K\|$ is contractible."

Exercises

1. The *chessboard complex* ♖$_{m,n}$ has the squares of the $m \times n$ chessboard as vertices, and simplices are all subsets of squares such that no two squares lie in the same row or column (so if we place rooks on these squares they do not threaten one another). Describe the "geometric shape" of $\|$♖$_{3,4}\|$.

1.6 Dimension of Geometric Realizations

Here is the promised sharper result about realizability of d-dimensional simplicial complexes.

1.6.1 Theorem (Geometric realization theorem). *Every finite d-dimensional simplicial complex K has a geometric realization in $\mathbb{R}^{2d+1}$.*

For $d = 1$, the theorem says that every graph can be represented in $\mathbb{R}^3$, with edges being straight segments. The dimension 3 is the smallest possible in general, since there are nonplanar graphs. A theorem of Van Kampen and Flores, which we will prove later (Theorem 5.1.1), shows that for every d there are d-dimensional simplicial complexes that cannot be realized in $\mathbb{R}^{2d}$, and so the dimension $2d+1$ in the geometric realization theorem is optimal for all d. Of course, this applies only in the worst case, since there are many d-dimensional simplicial complexes that can be realized in dimensions lower than $2d+1$ (say the d-simplex).

In the proof of Theorem 1.6.1, we use the following sufficient condition for a geometric realization.

1.6.2 Lemma. *If K is a simplicial complex and $f\colon V(\mathsf{K}) \to \mathbb{R}^d$ is an injective map such that $f(F \cup G)$ is affinely independent for all $F, G \in \mathsf{K}$, then the assignment*

$$F \longmapsto \sigma_F := \operatorname{conv}(f(F))$$

provides a geometric realization of K in $\mathbb{R}^d$.

Proof. If $f(F \cup G)$ is affinely independent, then σ_F and σ_G are two faces of the simplex with the vertex set $f(F \cup G)$. So $\sigma_F \cap \sigma_G = \sigma_{F \cap G}$, since the faces of a geometric simplex form a simplicial complex (Lemma 1.3.6).

A suitable placement of vertices can be defined using the moment curve. Later on, we will meet this useful curve several more times.

1.6.3 Definition. *The curve* $\{\gamma(t) : t \in \mathbb{R}\}$ *given by* $\gamma(t) := (t, t^2, \ldots, t^d)$ *is the* **moment curve** *in* $\mathbb{R}^d$.

The following lemma expresses a key property of the moment curve (any curve with this property would do in the sequel). It is a little stronger than needed here.

1.6.4 Lemma. *No hyperplane intersects the moment curve* γ *in* $\mathbb{R}^d$ *in more than* d *points. Consequently, every set of* $d+1$ *distinct points on* γ *is affinely independent. Moreover, if* γ *intersects a hyperplane* h *at* d *distinct points, then it crosses* h *from one side to the other at each intersection.*

Proof. A hyperplane h has an equation $a_1x_1 + a_2x_2 + \cdots + a_dx_d = b$ with $(a_1, \ldots, a_d) \neq \mathbf{0}$. If a point $\gamma(t)$ lies in h, then we have $a_1t + a_2t^2 + \cdots + a_dt^d = b$. This means that the values of t corresponding to intersections with h are the real roots of the nonzero polynomial $p(t) = (\sum_{i=1}^{d} a_it^i) - b$ of degree at most d. Such a $p(t)$ has at most d roots, and so there are no more than d intersections.

If there are d distinct intersections, then $p(t)$ has d distinct roots, which must all be simple. Therefore, $p(t)$ changes sign at each root, and this means that γ passes from one open half-space defined by h to the other at each intersection.

Proof of Theorem 1.6.1. We choose a map $f: V(\mathsf{K}) \to \mathbb{R}^{2d+1}$ such that the vertices of K are assigned distinct points on the moment curve in $\mathbb{R}^{2d+1}$. Then for $F, G \in \mathsf{K}$ we have $|F \cup G| \leq (d+1) + (d+1) = 2d+2$, and thus by Lemma 1.6.4 the corresponding points in $f(F \cup G)$ are affinely independent. Hence we are done by Lemma 1.6.2.

1.7 Simplicial Complexes and Posets

We recall that a *partially ordered set*, or *poset* for short, is a pair $(P, \preceq)$, where P is a set and $\preceq$ is a binary relation on P that is reflexive ($x \preceq x$), transitive ($x \preceq y$ and $y \preceq z$ imply $x \preceq z$), and weakly antisymmetric ($x \preceq y$ and $y \preceq x$ imply $x = y$). When the ordering relation $\preceq$ is understood, it is sometimes omitted from the notation, and we say only “a poset P.”

As we will see, there is a correspondence between (finite) simplicial complexes and (finite) posets. It is not quite one-to-one, but each poset is assigned a unique topological space, up to homeomorphism.

1.7.1 Definition. *The* **order complex** *of a poset P is the simplicial complex $\Delta(P)$, whose vertices are the elements of P and whose simplices are all chains (i.e., linearly ordered subsets, of the form $\{x_1, x_2, \ldots, x_k\}$, $x_1 \prec x_2 \prec \cdots \prec x_k$) in P.*

The **face poset** *of a simplicial complex K is the poset $P(\mathsf{K})$, which is the set of all* **nonempty** *simplices of K ordered by inclusion.*

For example, the simplicial complex

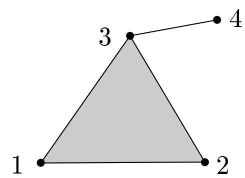

has the face poset

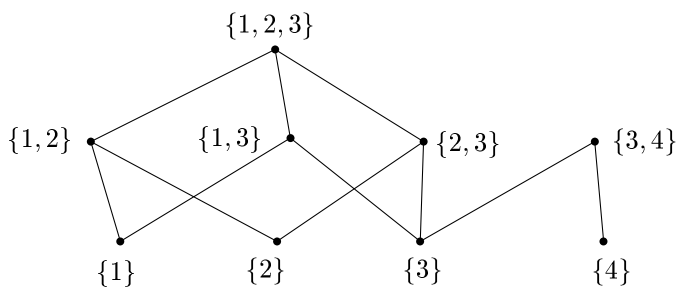

(this is the *Hasse diagram* of the poset, where each element is connected to its immediate predecessors and immediate successors, with the predecessors lying below it and the successors above it). Here is the order complex of this poset, together with a meadow saffron (also called autumn crocus; *Colchicum autumnale L.*) as an extra bonus:

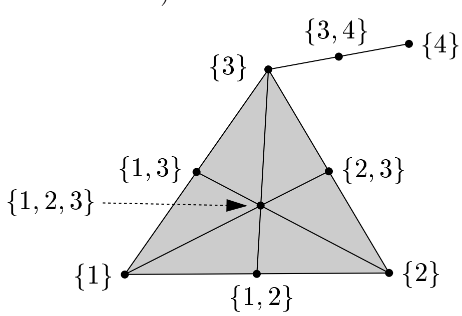

The operation we just did on the original simplicial complex, namely passing to the face poset and then to its order complex, is very important and has a name:

1.7.2 Definition. *For a simplicial complex* K, *the simplicial complex*

$$\mathrm{sd}(\mathsf{K}) := \Delta(P(\mathsf{K}))$$

is called the **(first) barycentric subdivision** *of* K.

More explicitly, the vertices of $\mathrm{sd}(\mathsf{K})$ are the nonempty simplices of K, and the simplices of $\mathrm{sd}(\mathsf{K})$ are chains of simplices of K ordered by inclusion.

Given a geometric realization of K, we can place the vertex of $\mathrm{sd}(\mathsf{K})$ corresponding to a simplex σ at the center of gravity (barycenter) of σ, as we did in the above picture. It turns out that, as the picture suggests, $\|\mathrm{sd}(\mathsf{K})\|$ is always (canonically) homeomorphic to $\|\mathsf{K}\|$. It suffices to prove this for K being (the simplicial complex of) a simplex; we leave this to the reader's diligence.

In algebraic topology, mainly in the earlier days, iterated barycentric subdivision was used for constructing arbitrarily fine triangulations of a given polyhedron. In the applications in this book, we will mainly encounter barycentric subdivision in its combinatorial meaning, in connection with posets.

Monotone maps and simplicial maps. Let $(P_1, \preceq_1)$ and $(P_2, \preceq_2)$ be posets. A mapping $f\colon P_1 \to P_2$ is called *monotone* if $x \preceq_1 y$ implies $f(x) \preceq_2 f(y)$. We have the following simple but useful result.

1.7.3 Proposition. *Every monotone mapping* $f\colon P_1 \to P_2$ *between posets is also a simplicial mapping* $V(\Delta(P_1)) \to V(\Delta(P_2))$ *between their order complexes.*

We again leave the very easy verification to the reader.

1.7.4 Corollary. *Let* K_1 *and* K_2 *be simplicial complexes. Consider an arbitrary mapping* f *that assigns to each simplex* $F \in \mathsf{K}_1$ *a simplex* $f(F) \in \mathsf{K}_2$ *(*f *is not necessarily induced by a mapping of vertices!), and suppose that if* $F' \subseteq F$, *then also* $f(F') \subseteq f(F)$. *Then* f *can be regarded as a simplicial mapping of* $\mathrm{sd}(\mathsf{K}_1)$ *into* $\mathrm{sd}(\mathsf{K}_2)$, *and so it induces a continuous map* $\|f\|\colon \|\mathsf{K}_1\| \to \|\mathsf{K}_2\|$.

Notes. In many books and papers, $\mathrm{sd}(\mathsf{K})$ is denoted by K', and sometimes it is called the *derived* of K.

If we iterate the barycentric subdivision of a geometric simplicial complex sufficiently many times, the diameter of all simplices decreases below any prescribed threshold (Exercise 3). This is a standard way of producing arbitrarily fine triangulations. However, it is not very suitable for algorithmic applications where the number and shape of simplices are important.

The order complex $\Delta(P)$ is an instance of a more general construction of a *classifying space*; see, e.g., [Hat01, Chapter 2].

Let us mention a result somewhat similar to the geometric realization theorem (Theorem 1.6.1), which provides an upper bound on the dimension necessary for embedding a given simplicial complex. First we recall the notion of *Dushnik–Miller dimension* (or *order dimension*) of a poset. As is easy to check, if $(P, \preceq)$ is a finite poset, there exist linear orderings $\leq_1, \leq_2, \ldots, \leq_k$ such that $x \preceq y$ iff $x \leq_i y$ for all $i \in [k]$. In other words, $\preceq = \bigcap_{i=1}^{k} \leq_i$ (here $\leq_i$ stands for the ith linear ordering considered as a binary relation on P, that is, a subset of $P \times P$). The smallest possible k for such a representation of $\preceq$ by linear orderings is the Dushnik–Miller dimension $\dim(P, \preceq)$. Ossona de Mendez [Oss99] proved, using Scarf's construction, that every finite simplicial complex K can be geometrically realized in $\mathbb{R}^{d-1}$ with $d := \dim(P(\mathsf{K}))$. For a proof, let $\leq_1, \ldots, \leq_d$ be linear orderings of K witnessing $\dim(P(\mathsf{K})) = d$. We restrict the orderings $\leq_i$ to the set $V := V(\mathsf{K})$ (the vertices are also simplices of K), and we let φ_i be the injective map $V \to [n]$, $n = |V|$, that is monotone with respect to $\leq_i$ (that is, $u <_i v$ iff $\varphi_i(u) < \varphi_i(v)$ for every $u, v \in V$). We define $f_0 \colon V \to \mathbb{R}^d$ by $f_0(v) = ((d+1)^{\varphi_1(v)}, (d+1)^{\varphi_2(v)}, \ldots, (d+1)^{\varphi_d(v)})$, and finally, we let $f(v)$ be the projection of $f_0(v)$ from $\mathbf{0}$ on the hyperplane $\sum_{i=1}^{d} x_i = 1$. Then it can be shown that f satisfies the condition of Lemma 1.6.2 and thus provides a realization of K in $\mathbb{R}^{d-1}$.

A converse of this theorem is known for $d = 3$: If we regard a graph G as a 1-dimensional simplicial complex, then the dimension of the face poset is at most 3 if and only if G is planar [Sch89]; also see [BT93], [BT97], [Fel01] for related results.

Exercises

1.* Prove that a simplex is homeomorphic to its barycentric subdivision (a rigorous proof takes some work!).
2. Prove Proposition 1.7.3 and Corollary 1.7.4.
3.* (a) Prove that the diameter of an arbitrary simplex σ is equal to the distance between some two vertices of σ.
(b) Prove that for every n and $\delta > 0$ there exists k such that if σ^n is any n-dimensional simplex of diameter 1, then all simplices of $\mathrm{sd}^k(\sigma^n)$ (barycentric subdivision iterated k times) have diameter at most δ. Does k have to depend on n?