Chapter 7

Basics of Combinatorial Topology

7.1 Simplicial and Polyhedral Complexes

In order to study and manipulate complex shapes it is convenient to discretize these shapes and to view them as the union of simple building blocks glued together in a "clean fashion".

The building blocks should be simple geometric objects, for example, points, lines segments, triangles, tehrahedra and more generally simplices, or even convex polytopes.

Definition 7.1.1 Let $\mathcal{E}$ be any normed affine space, say $\mathcal{E} = \mathbb{E}^m$ with its usual Euclidean norm. Given any $n+1$ affinely independent points $a_0, \ldots, a_n$ in $\mathcal{E}$, the *n-simplex (or simplex) σ defined by $a_0, \ldots, a_n$* is the convex hull of the points $a_0, \ldots, a_n$, that is, the set of all convex combinations $\lambda_0 a_0 + \cdots + \lambda_n a_n$, where $\lambda_0 + \cdots + \lambda_n = 1$ and $\lambda_i \geq 0$ for all i, $0 \leq i \leq n$.

We call n the *dimension* of the n-simplex σ, and the points $a_0, \ldots, a_n$ are the *vertices* of σ.

Given any subset $\{a_{i_0}, \ldots, a_{i_k}\}$ of $\{a_0, \ldots, a_n\}$ (where $0 \leq k \leq n$), the k-simplex generated by $a_{i_0}, \ldots, a_{i_k}$ is called a *k-face* or simply a *face* of σ.

A face s of σ is a *proper face* if $s \neq \sigma$ (we agree that the empty set is a face of any simplex). For any vertex a_i, the face generated by $a_0, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n$ (i.e., omitting a_i) is called the *face opposite* a_i.

Every face that is an $(n-1)$-simplex is called a *boundary face* or *facet*. The union of the boundary faces is the *boundary of* σ, denoted by $\partial\sigma$, and the complement of $\partial\sigma$ in σ is the *interior* $\mathrm{Int}\,\sigma = \sigma - \partial\sigma$ of σ. The interior $\mathrm{Int}\,\sigma$ of σ is sometimes called an *open simplex*.

It should be noted that for a 0-simplex consisting of a single point $\{a_0\}$, $\partial\{a_0\} = \emptyset$, and $\mathrm{Int}\,\{a_0\} = \{a_0\}$.

Of course, a 0-simplex is a single point, a 1-simplex is the line segment (a_0, a_1), a 2-simplex is a triangle (a_0, a_1, a_2) (with its interior), and a 3-simplex is a tetrahedron (a_0, a_1, a_2, a_3) (with its interior).

The inclusion relation between any two faces σ and τ of some simplex, s, is written $\sigma \preceq \tau$.

Clearly, a point x belongs to the boundary $\partial\sigma$ of σ iff at least one of its barycentric coordinates $(\lambda_0, \ldots, \lambda_n)$ is zero, and a point x belongs to the interior $\mathrm{Int}\,\sigma$ of σ iff all of its barycentric coordinates $(\lambda_0, \ldots, \lambda_n)$ are positive, i.e., $\lambda_i > 0$ for all i, $0 \leq i \leq n$.

Then, for every $x \in \sigma$, there is a unique face s such that $x \in \mathrm{Int}\,s$, the face generated by those points a_i for which $\lambda_i > 0$, where $(\lambda_0, \ldots, \lambda_n)$ are the barycentric coordinates of x.

A simplex σ is convex, arcwise connected, compact, and closed. The interior $\operatorname{Int}\sigma$ of a simplex is convex, arcwise connected, open, and σ is the closure of $\operatorname{Int}\sigma$.

We now put simplices together to form more complex shapes. The intuition behind the next definition is that the building blocks should be "glued cleanly".

Definition 7.1.2 A *simplicial complex in* $\mathbb{E}^m$ (for short, a *complex* in $\mathbb{E}^m$) is a set K consisting of a (finite or infinite) set of simplices in $\mathbb{E}^m$ satisfying the following conditions:

(1) Every face of a simplex in K also belongs to K.

(2) For any two simplices σ_1 and σ_2 in K, if $\sigma_1 \cap \sigma_2 \neq \emptyset$, then $\sigma_1 \cap \sigma_2$ is a common face of both σ_1 and σ_2.

Every k-simplex, $\sigma \in K$, is called a *k-face* (or *face*) of K. A 0-face $\{v\}$ is called a *vertex* and a 1-face is called an *edge*. The *dimension* of the simplicial complex K is the maximum of the dimensions of all simplices in K.

If $\dim K = d$, then every face of dimension d is called a *cell* and every face of dimension $d-1$ is called a *facet*.

Condition (2) guarantees that the various simplices forming a complex intersect nicely. It is easily shown that the following condition is equivalent to condition (2):

(2′) For any two distinct simplices σ_1, σ_2,
$\mathrm{Int}\,\sigma_1 \cap \mathrm{Int}\,\sigma_2 = \emptyset$.

Remarks:

1. A simplicial complex, K, is a combinatorial object, namely, a *set* of simplices satisfying certain conditions but not a subset of $\mathbb{E}^m$. However, every complex, K, yields a subset of $\mathbb{E}^m$ called the geometric realization of K and denoted $|K|$. This object will be defined shortly and should not be confused with the complex.

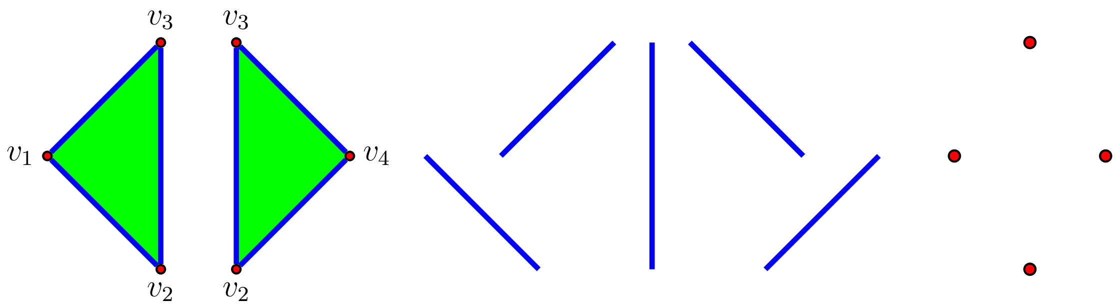

Figure 7.1: A set of simplices forming a complex

Figure 7.1 illustrates this aspect of the definition of a complex. For clarity, the two triangles (2-simplices) are drawn as disjoint objects even though they share the common edge, (v_2, v_3) (a 1-simplex) and similarly for the edges that meet at some common vertex.

2. Some authors define a *facet* of a complex, K, of dimension d to be a d-simplex in K, as opposed to a $(d-1)$-simplex, as we did. This practice is not consistent with the notion of facet of a polyhedron and this is why we prefer the terminology *cell* for the d-simplices in K.

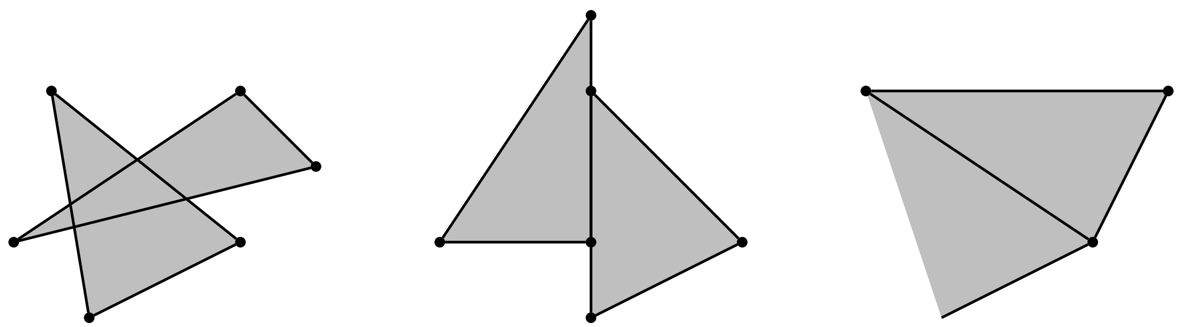

Figure 7.2: Collections of simplices not forming a complex

3. It is important to note that in order for a complex, K, of dimension d to be realized in $\mathbb{E}^m$, the dimension of the "ambient space", m, must be big enough. For example, there are 2-complexes that can't be realized in $\mathbb{E}^3$ or even in $\mathbb{E}^4$. There has to be enough room in order for condition (2) to be satisfied. It is not hard to prove that $m = 2d+1$ is always sufficient. Sometimes, $2d$ works, for example in the case of surfaces (where $d = 2$).

Some collections of simplices violating some of the conditions of Definition 7.1.2 are shown in Figure 7.2.

On the left, the intersection of the two 2-simplices is neither an edge nor a vertex of either triangle.

In the middle case, two simplices meet along an edge which is not an edge of either triangle.

On the right, there is a missing edge and a missing vertex.

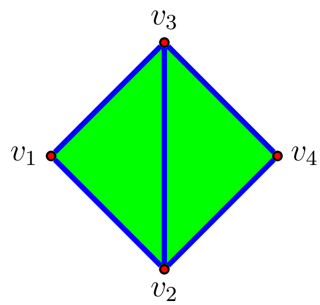

Figure 7.3: The geometric realization of the complex of Figure 7.1

The union $|K|$ of all the simplices in K is a subset of $\mathbb{E}^m$. We can define a topology on $|K|$ by defining a subset F of $|K|$ to be closed iff $F \cap \sigma$ is closed in σ for every face $\sigma \in K$.

It is immediately verified that the axioms of a topological space are indeed satisfied.

The resulting topological space $|K|$ is called the *geometric realization of* K.

The geometric realization of the complex from Figure 7.1 is show in Figure 7.1.

Some "legal" simplicial complexes are shown in Figure 7.4.

Obviously, $|\sigma| = \sigma$ for every simplex, σ. Also, note that distinct complexes may have the same geometric realization. In fact, all the complexes obtained by subdividing the simplices of a given complex yield the same geometric realization.

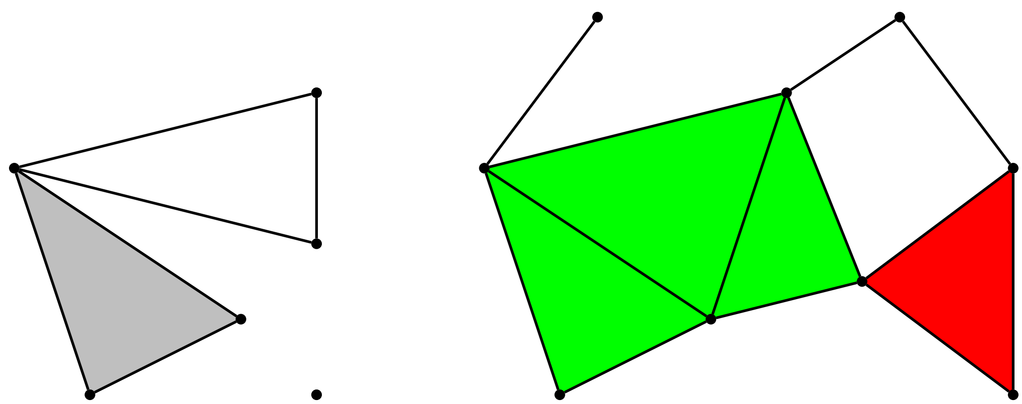

Figure 7.4: Examples of simplicial complexes

A *polytope* is the geometric realization of some simplicial complex. A polytope of dimension 1 is usually called a *polygon*, and a polytope of dimension 2 is usually called a *polyhedron*.

When K consists of infinitely many simplices we usually require that K be *locally finite*, which means that every vertex belongs to finitely many faces. If K is locally finite, then its geometric realization, $|K|$, is locally compact.

In the sequel, we will consider only finite simplicial complexes, that is, complexes K consisting of a finite number of simplices.

In this case, the topology of $|K|$ defined above is identical to the topology induced from $\mathbb{E}^m$. For any simplex σ in K, $\mathrm{Int}\,\sigma$ coincides with the interior $\overset{\circ}{\sigma}$ of σ in the topological sense, and $\partial\sigma$ coincides with the boundary of σ in the topological sense.

Definition 7.1.3 Given any complex, K_2, a subset $K_1 \subseteq K_2$ of K_2 is a *subcomplex* of K_2 iff it is also a complex. For any complex, K, of dimension d, for any i with $0 \leq i \leq d$, the subset

$$K^{(i)} = \{\sigma \in K \mid \dim \sigma \leq i\}$$

is called the *i-skeleton* of K. Clearly, $K^{(i)}$ is a subcomplex of K. We also let

$$K^i = \{\sigma \in K \mid \dim \sigma = i\}.$$

Observe that K^0 is the set of vertices of K and K^i is not a complex.

A simplicial complex, K_1 is a *subdivision* of a complex K_2 iff $|K_1| = |K_2|$ and if every face of K_1 is a subset of some face of K_2.

A complex K of dimension d is *pure* (or *homogeneous*) iff every face of K is a face of some d-simplex of K (i.e., some cell of K). A complex is *connected* iff $|K|$ is connected.

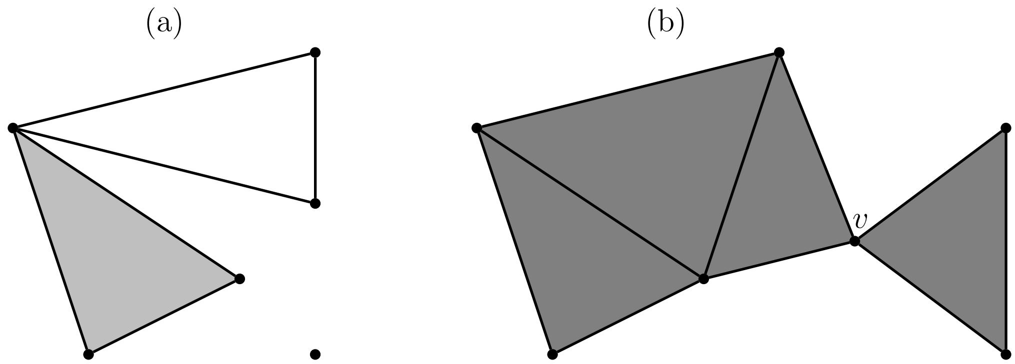

Figure 7.5: (a) A complex that is not pure. (b) A pure complex

It is is easy to see that a complex is connected iff its 1-skeleton is connected.

The intuition behind the notion of a pure complex, K, of dimension d is that a pure complex is the result of gluing pieces all having the same dimension, namely, d-simplices.

For example, in Figure 7.5, the complex on the left is not pure but the complex on the right is pure of dimension 2.

Most of the shapes that we will be interested in are well approximated by pure complexes, in particular, surfaces or solids.

However, pure complexes may still have undesirable "singularities" such as the vertex, v, in Figure 7.5(b).

The notion of link of a vertex provides a technical way to deal with singularities.

Definition 7.1.4 Let K be any complex and let σ be any face of K. The *star*, $\mathrm{St}(\sigma)$ (or if we need to be very precise, $\mathrm{St}(\sigma, K)$), of σ is the subcomplex of K consisting of all faces, τ, containing σ and of all faces of τ, *i.e.*,

$$\mathrm{St}(\sigma) = \{s \in K \mid (\exists \tau \in K)(\sigma \preceq \tau \quad \text{and} \quad s \preceq \tau)\}.$$

The *link*, $\mathrm{Lk}(\sigma)$ (or $\mathrm{Lk}(\sigma, K)$) of σ is the subcomplex of K consisting of all faces in $\mathrm{St}(\sigma)$ that do not intersect σ, i.e.,

$$\mathrm{Lk}(\sigma) = \{\tau \in K \mid \tau \in \mathrm{St}(\sigma) \quad \text{and} \quad \sigma \cap \tau = \emptyset\}.$$

To simplify notation, if $\sigma = \{v\}$ is a vertex we write $\mathrm{St}(v)$ for $\mathrm{St}(\{v\})$ and $\mathrm{Lk}(v)$ for $\mathrm{Lk}(\{v\})$.

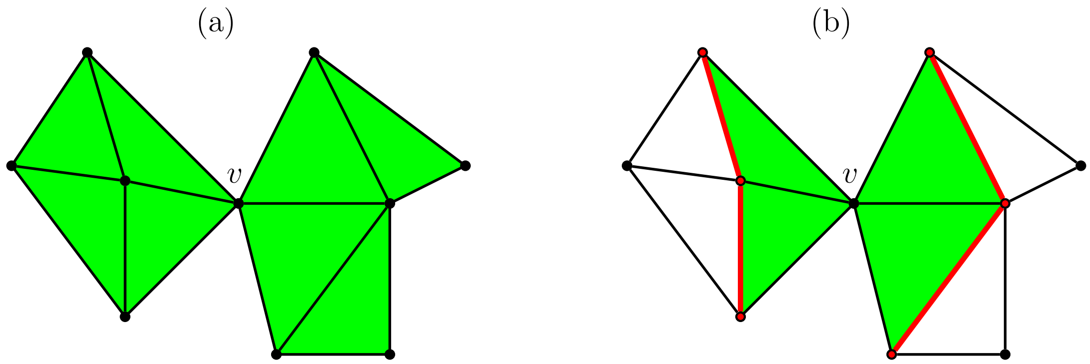

Figure 7.6: (a) A complex. (b) Star and Link of v

Figure 7.6 shows:

(a) A complex (on the left).

(b) The star of the vertex v, indicated in gray and the link of v, showed as thicker lines.

If K is pure and of dimension d, then $\mathrm{St}(\sigma)$ is also pure of dimension d and if $\dim \sigma = k$, then $\mathrm{Lk}(\sigma)$ is pure of dimension $d - k - 1$.

For technical reasons, following Munkres [**?**], besides defining the complex, $\mathrm{St}(\sigma)$, it is useful to introduce the *open star* of σ, denoted $\mathrm{st}(\sigma)$, defined as the subspace of $|K|$ consisting of the union of the interiors, $\mathrm{Int}(\tau) = \tau - \partial\,\tau$, of all the faces, τ, containing, σ.

According to this definition, the open star of σ is not a complex but instead a subset of $|K|$.

Note that

$$\overline{\mathrm{st}(\sigma)} = |\mathrm{St}(\sigma)|,$$

that is, the closure of $\mathrm{st}(\sigma)$ is the geometric realization of the complex $\mathrm{St}(\sigma)$.

Then, $\mathrm{lk}(\sigma) = |\mathrm{Lk}(\sigma)|$ is the union of the simplices in $\mathrm{St}(\sigma)$ that are disjoint from σ.

If σ is a vertex, v, we have

$$\mathrm{lk}(v) = \overline{\mathrm{st}(v)} - \mathrm{st}(v).$$

However, beware that if σ is not a vertex, then $\mathrm{lk}(\sigma)$ is properly contained in $\overline{\mathrm{st}(\sigma)} - \mathrm{st}(\sigma)$!

One of the nice properties of the open star, $\mathrm{st}(\sigma)$, of σ is that it is open. This follows from the fact that the open star, $\mathrm{st}(v)$, of a vertex, v is open.

Furthermore, for every point, $a \in |K|$, there is a unique smallest simplex, σ, so that $a \in \mathrm{Int}(\sigma) = \sigma - \partial\,\sigma$.

As a consequence, for any k-face, σ, of K, if $\sigma = (v_0, \ldots, v_k)$, then

$$\mathrm{st}(\sigma) = \mathrm{st}(v_0) \cap \cdots \cap \mathrm{st}(v_k).$$

Consequently, $\mathrm{st}(\sigma)$ is open and path connected.

Unfortunately, the "nice" equation

$$\mathrm{St}(\sigma) = \mathrm{St}(v_0) \cap \cdots \cap \mathrm{St}(v_k)$$

is false! (and anagolously for $\mathrm{Lk}(\sigma)$.)

For a counter-example, consider the boundary of a tetrahedron with one face removed.

Recall that in $\mathbb{E}^d$, the *(open) unit ball*, B^d, is defined by

$$B^d = \{x \in \mathbb{E}^d \mid \|x\| < 1\},$$

the *closed unit ball*, $\overline{B}^d$, is defined by

$$\overline{B}^d = \{x \in \mathbb{E}^d \mid \|x\| \leq 1\},$$

and the $(d-1)$-*sphere*, S^{d-1}, by

$$S^{d-1} = \{x \in \mathbb{E}^d \mid \|x\| = 1\}.$$

Obviously, S^{d-1} is the boundary of $\overline{B}^d$ (and B^d).

Definition 7.1.5 Let K be a pure complex of dimension d and let σ be any k-face of K, with $0 \leq k \leq d-1$. We say that σ is *nonsingular* iff the geometric realization, $\mathrm{lk}(\sigma)$, of the link of σ is homeomorphic to either S^{d-k-1} or to $\overline{B}^{d-k-1}$; this is written as $\mathrm{lk}(\sigma) \approx S^{d-k-1}$ or $\mathrm{lk}(\sigma) \approx \overline{B}^{d-k-1}$, where $\approx$ means homeomorphic.

In Figure 7.6, note that the link of v is not homeomorphic to S^1 or B^1, so v is singular.

It will also be useful to express $\mathrm{St}(v)$ in terms of $\mathrm{Lk}(v)$, where v is a vertex, and for this, we define the notion of cone.

Definition 7.1.6 Given any complex, K, in $\mathbb{E}^n$, if $\dim K = d < n$, for any point, $v \in \mathbb{E}^n$, such that v does not belong to the affine hull of $|K|$, the *cone on K with vertex v*, denoted, $v * K$, is the complex consisting of all simplices of the form $(v, a_0, \ldots, a_k)$ and their faces, where $(a_0, \ldots, a_k)$ is any k-face of K. If $K = \emptyset$, we set $v * K = v$.

It is not hard to check that $v * K$ is indeed a complex of dimension $d + 1$ containing K as a subcomplex.

Proposition 7.1.7 *For any complex, K, of dimension d and any vertex, $v \in K$, we have*

$$\mathrm{St}(v) = v * \mathrm{Lk}(v).$$

More generally, for any face, σ, of K, we have

$$\overline{\mathrm{st}(\sigma)} = |\mathrm{St}(\sigma)| \approx \sigma \times |v * \mathrm{Lk}(\sigma)|,$$

for every $v \in \sigma$ and

$$\overline{\mathrm{st}(\sigma)} - \mathrm{st}(\sigma) = \partial\,\sigma \times |v * \mathrm{Lk}(\sigma)|,$$

for every $v \in \partial\,\sigma$.

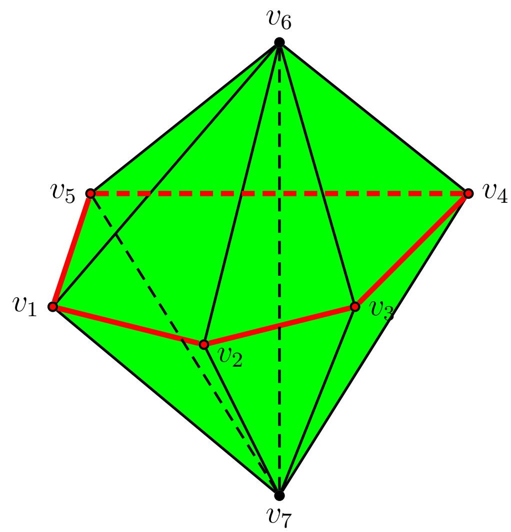

Figure 7.7: More examples of links and stars

Figure 7.7 shows a 3-dimensional complex. The link of the edge (v_6, v_7) is the pentagon $P = (v_1, v_2, v_3, v_4, v_5) \approx S^1$. The link of the vertex v_7 is the cone $v_6 * P \approx B^2$. The link of (v_1, v_2) is $(v_6, v_7) \approx B^1$ and the link of v_1 is the union of the triangles (v_2, v_6, v_7) and (v_5, v_6, v_7), which is homeomorphic to B^2.

Remark: Unfortunately, the word "cone" is overloaded. It might have been better to use the term *pyramid* as some authors do (including Ziegler).

Given a pure complex, it is necessary to distinguish between two kinds of faces.

Definition 7.1.8 Let K be any pure complex of dimension d. A k-face, σ, of K is a *boundary* or *external* face iff it belongs to a single cell (i.e., a d-simplex) of K and otherwise it is called an *internal* face ($0 \leq k \leq d-1$). The *boundary* of K, denoted $\mathrm{bd}(K)$, is the subcomplex of K consisting of all boundary facets of K together with their faces.

It is clear by definition that $\mathrm{bd}(K)$ is a pure complex of dimension $d-1$.

Even if K is connected, $\mathrm{bd}(K)$ is not connected, in general.

For example, if K is a 2-complex in the plane, the boundary of K usually consists of several simple closed polygons (i.e, 1 dimensional complexes homeomorphic to the circle, S^1).

Proposition 7.1.9 *Let K be any pure complex of dimension d. For any k-face, σ, of K the boundary complex,* $\mathrm{bd}(\mathrm{Lk}(\sigma))$, *is nonempty iff σ is a boundary face of K ($0 \leq k \leq d-2$). Furthermore,*

$$\mathrm{Lk}_{\mathrm{bd}(K)}(\sigma) = \mathrm{bd}(\mathrm{Lk}(\sigma))$$

for every face, σ, of $\mathrm{bd}(K)$, *where* $\mathrm{Lk}_{\mathrm{bd}(K)}(\sigma)$ *denotes the link of σ in* $\mathrm{bd}(K)$.

Proposition 7.1.9 shows that if every face of K is nonsingular, then the link of every internal face is a sphere whereas the link of every external face is a ball.

Proposition 7.1.10 *Let K be any pure complex of dimension d. If every vertex of K is nonsingular, then* $\mathrm{st}(\sigma) \approx B^d$ *for every k-face, σ, of K ($1 \leq k \leq d-1$).*

Here are more useful propositions about pure complexes without singularities.

Proposition 7.1.11 *Let K be any pure complex of dimension d. If every vertex of K is nonsingular, then for every point, $a \in |K|$, there is an open subset, $U \subseteq |K|$, containing a such that $U \approx B^d$ or $U \approx B^d \cap \mathbb{H}^d$, where $\mathbb{H}^d = \{(x_1, \ldots, x_d) \in \mathbb{R}^d \mid x_d \geq 0\}$.*

Proposition 7.1.12 *Let K be any pure complex of dimension d. If every facet of K is nonsingular, then every facet of K, is contained in at most two cells (d-simplices).*

Proposition 7.1.13 *Let K be any pure and connected complex of dimension d. If every face of K is nonsingular, then for every pair of cells (d-simplices), σ and σ', there is a sequence of cells, $\sigma_0, \ldots, \sigma_p$, with $\sigma_0 = \sigma$ and $\sigma_p = \sigma'$, and such that σ_i and σ_{i+1} have a common facet, for $i = 0, \ldots, p-1$.*

Proposition 7.1.14 *Let K be any pure complex of dimension d. If every facet of K is nonsingular, then the boundary,* $\mathrm{bd}(K)$, *of K is a pure complex of dimension $d-1$ with an empty boundary. Furthermore, if every face of K is nonsingular, then every face of* $\mathrm{bd}(K)$ *is also nonsingular.*

The building blocks of simplicial complexes, namely, simplicies, are in some sense mathematically ideal. However, in practice, it may be desirable to use a more flexible set of building blocks.

We can indeed do this and use convex polytopes as our building blocks.

Definition 7.1.15 A *polyhedral complex in* $\mathbb{E}^m$ (for short, a *complex* in $\mathbb{E}^m$) is a set, K, consisting of a (finite or infinite) set of convex polytopes in $\mathbb{E}^m$ satisfying the following conditions:

(1) Every face of a polytope in K also belongs to K.

(2) For any two polytopes σ_1 and σ_2 in K, if $\sigma_1 \cap \sigma_2 \neq \emptyset$, then $\sigma_1 \cap \sigma_2$ is a common face of both σ_1 and σ_2.

Every polytope, $\sigma \in K$, of dimension k, is called a *k-face* (or *face*) of K. A 0-face $\{v\}$ is called a *vertex* and a 1-face is called an *edge*. The *dimension* of the polyhedral complex K is the maximum of the dimensions of all polytopes in K. If $\dim K = d$, then every face of dimension d is called a *cell* and every face of dimension $d - 1$ is called a *facet*.

Every Polytope, P, yields two natural polyhedral complexes:

(i) The polyhedral complex, $\mathcal{K}(P)$, consisting of P together with all of its faces. This complex has a single cell, namely, P itself.

(ii) The *boundary complex*, $\mathcal{K}(\partial P)$, consisting of all faces of P other than P itself. The cells of $\mathcal{K}(\partial P)$ are the facets of P.

The notions of k-skeleton and pureness are defined just as in the simplicial case.

The notions of star and link are defined for polyhedral complexes just as they are defined for simplicial complexes except that the word “face” now means face of a polytope.

Now, by Theorem 6.3.1, every polytope, σ, is the convex hull of its vertices. Let $\mathrm{vert}(\sigma)$ denote the set of vertices of σ.

We have the following crucial observation: Given any poyhedral complex, K, for every point, $x \in |K|$, there is a *unique* polytope, $\sigma_x \in K$, such that
$x \in \mathrm{Int}(\sigma_x) = \sigma_x - \partial\, \sigma_x$.

Now, just as in the simplicial case, the open star, $\mathrm{st}(\sigma)$, of σ is given by

$$\mathrm{st}(\sigma) = \bigcap_{v \in \mathrm{vert}(\sigma)} \mathrm{st}(v).$$

and $\mathrm{st}(\sigma)$ is open in $|K|$.

The next proposition is another result that seems quite obvious, yet a rigorous proof is more involved that we might think.

This proposition states that a convex polytope can always be cut up into simplices, that is, it can be subdivided into a simplicial complex.

In other words, every convex polytope can be triangulated. This implies that simplicial complexes are as general as polyhedral complexes.

Proposition 7.1.16 *Every convex d-polytope, P, can be subdivided into a simplicial complex without adding any new vertices, i.e., every convex polytope can be triangulated.*

With all this preparation, it is now quite natural to define combinatorial manifolds.

7.2 Combinatorial and Topological Manifolds

The notion of pure complex without singular faces turns out to be a very good "discrete" approximation of the notion of (topological) manifold because it of its highly to computational nature.

Definition 7.2.1 A *combinatorial d-manifold* is any space, X, homeomorphic to the geometric realization, $|K| \subseteq \mathbb{E}^n$, of some pure (simplicial or polyhedral) complex, K, of dimension d whose faces are all nonsingular. If the link of every k-face of K is homeomorphic to the sphere S^{d-k-1}, we say that X is a combinatorial manifold *without boundary*, else it is a combinatorial manifold *with boundary*.

Other authors use the term *triangulation* for what we call a computational manifold.

It is easy to see that the connected components of a combinatorial 1-manifold are either simple closed polygons or simple chains (simple, means that the interiors of distinct edges are disjoint).

A combinatorial 2-manifold which is connected is also called a *combinatorial surface* (with or without boundary). Proposition 7.1.14 immediately yields the following result:

Proposition 7.2.2 *If X is a combinatorial d-manifold with boundary, then* $\mathrm{bd}(X)$ *is a combinatorial* $(d-1)$*-manifold without boundary.*

Now, because we are assuming that X sits in some Euclidean space, $\mathbb{E}^n$, the space X is Hausdorff and second-countable.

(Recall that a topological space is second-countable iff there is a countable family of open sets of X, $\{U_i\}_{i\geq 0}$, such that every open subset of X is the union of open sets from this family.)

Since it is desirable to have a good match between manifolds and combinatorial manifolds, we are led to the following definition:

Recall that

$$\mathbb{H}^d = \{(x_1, \ldots, x_d) \in \mathbb{R}^d \mid x_d \geq 0\}.$$

Definition 7.2.3 For any $d \geq 1$, a *(topological) d-manifold with boundary* is a second-countable, topological Hausdorff space M, together with an open cover, $(U_i)_{i\in I}$, of open sets in M and a family, $(\varphi_i)_{i\in I}$, of homeomorphisms, $\varphi_i \colon U_i \to \Omega_i$, where each Ω_i is some open subset of $\mathbb{H}^d$ in the subset topology.

Each pair (U, φ) is called a *coordinate system*, or *chart*, of M, each homeomorphism $\varphi_i \colon U_i \to \Omega_i$ is called a *coordinate map*, and its inverse $\varphi_i^{-1} \colon \Omega_i \to U_i$ is called a *parameterization* of U_i. The family $(U_i, \varphi_i)_{i\in I}$ is often called an *atlas* for M.

A *(topological) bordered surface* is a connected 2-manifold with boundary. If for every homeomorphism, $\varphi_i \colon U_i \rightarrow \Omega_i$, the open set $\Omega_i \subseteq \mathbb{H}^d$ is actually an open set in $\mathbb{R}^d$ (which means that $x_d > 0$ for every $(x_1, \ldots, x_d) \in \Omega_i$), then we say that M is a *d-manifold*.

Note that a d-manifold is also a d-manifold with boundary.

Letting $\partial\mathbb{H}^d = \mathbb{R}^{d-1} \times \{0\}$, it can be shown using homology, that if some coordinate map, φ, defined on p maps p into $\partial\mathbb{H}^d$, then every coordinate map, ψ, defined on p maps p into $\partial\mathbb{H}^d$.

Thus, M is the disjoint union of two sets ∂M and $\mathrm{Int}\, M$, where ∂M is the subset consisting of all points $p \in M$ that are mapped by some (in fact, all) coordinate map, φ, defined on p into $\partial\mathbb{H}^d$, and where $\mathrm{Int}\, M = M - \partial M$.

The set ∂M is called the *boundary* of M, and the set $\mathrm{Int}\, M$ is called the *interior* of M, even though this terminology clashes with some prior topological definitions.

A good example of a bordered surface is the Möbius strip. The boundary of the Möbius strip is a circle.

The boundary ∂M of M may be empty, but $\operatorname{Int} M$ is nonempty. Also, it can be shown using homology, that the integer d is unique.

It is clear that $\operatorname{Int} M$ is open, and an d-manifold, and that ∂M is closed.

It is easy to see that ∂M is an $(d-1)$-manifold.

Proposition 7.2.4 *Every combinatorial d-manifold is a d-manifold with boundary.*

Proof. This is an immediate consequence of Proposition 7.1.11. $\square$

Is the converse of Proposition 7.2.4 true?

It turns out that answer is yes for $d = 1, 2, 3$ but **no** for $d \geq 4$. This is not hard to prove for $d = 1$.

For $d = 2$ and $d = 3$, this is quite hard to prove; among other things, it is necessary to prove that triangulations exist and this is very technical.

For $d \geq 4$, not every manifold can be triangulated (in fact, this is undecidable!).

What if we assume that M is a triangulated manifold, which means that $M \approx |K|$, for some pure d-dimensional complex, K?

Surprinsingly, for $d \geq 5$, there are triangulated manifolds whose links are not spherical (i.e., not homeomorphic to $\overline{B}^{d-k-1}$ or S^{d-k-1}).

Fortunately, we will only have to deal with $d = 2, 3$!

Another issue that must be addressed is orientability.

Assume that fix a total ordering of the vertices of a complex, K. Let $\sigma = (v_0, \ldots, v_k)$ be any simplex.

Recall that every permutation (of $\{0, \ldots, k\}$) is a product of *transpositions*, where a transposition swaps two distinct elements, say i and j, and leaves every other element fixed.

Furthermore, for any permutation, π, the parity of the number of transpositions needed to obtain π only depends on π and it called the *signature* of π.

We say that two permutations are equivalent iff they have the same signature. Consequently, there are two equivalence classes of permutations: Those of even signature and those of odd signature.

Then, an *orientation* of σ is the choice of one of the two equivalence classes of permutations of its vertices. If σ has been given an orientation, then we denote by $-\sigma$ the result of assigning the other orientation to it (we call it the *opposite orientation*).

For example, $(0,1,2)$ has the two orientation classes:

$$\{(0,1,2),(1,2,0),(2,0,1)\}$$

and

$$\{(2,1,0),(1,0,2),(0,2,1)\}.$$

Definition 7.2.5 Let $X \approx |K|$ be a combinatorial d-manifold. We say that X is *orientable* if it is possible to assign an orientation to all of its cells (d-simplices) so that whenever two cells σ_1 and σ_2 have a common facet, σ, the two orientations induced by σ_1 and σ_2 on σ are opposite. A combinatorial d-manifold together with a specific orientation of its cells is called an *oriented manifold*. If X is not orientable we say that it is *non-orientable*.

There are non-orientable (combinatorial) surfaces, for example, the Möbius strip which can be realized in $\mathbb{E}^3$. The Möbius strip is a surface with boundary, its boundary being a circle.

There are also non-orientable (combinatorial) surfaces such as the Klein bottle or the projective plane but they can only be realized in $\mathbb{E}^4$ (in $\mathbb{E}^3$, they must have singularities such as self-intersection).

We will only be dealing with orientable manifolds, and most of the time, surfaces.

One of the most important invariants of combinatorial (and topological) manifolds is their *Euler characteristic*.

In the next chapter, we prove a famous formula due to Poincaré giving the Euler characteristic of a convex polytope. For this, we will introduce a technique of independent interest called *shelling*.