3

Simplicial Topology

We present a brief overview of the theory of homology and homotopy for simplicial complexes and quotients of simplicial complexes. We also list some of the most important classes of simplicial complexes such as contractible and shellable complexes.

In Section 3.1, we consider simplicial homology theory, stating the main definitions and presenting the important Mayer-Vietoris exact sequence. In Section 3.2, we proceed with relative homology and present the long exact sequence for pairs of simplicial complexes. We also state the main result about Alexander duality. Section 3.3 provides the basic definitions from simplicial homotopy theory. In Section 3.4, we discuss acyclic, contractible, collapsible, and nonevasive complexes. We will need some results about quotient complexes, most notably the Contractible Subcomplex Lemma; we present these results in Section 3.5. Section 3.6 is devoted to Cohen-Macaulay, constructible, shellable, and vertex-decomposable complexes. We proceed with balls and spheres in Section 3.7 and conclude the chapter in Section 3.8 with a few comments about the well-known Stanley-Reisner correspondence between simplicial complexes and monomial rings and ideals.

3.1 Simplicial Homology

We review the basic concepts of simplicial homology. Simplicial homology is well-known to coincide with the restriction of singular or cellular homology to simplicial complexes; see Munkres [101, §34, §39].

Throughout this section, let $\mathbb{F}$ be a field or $\mathbb{Z}$, the ring of integers.

Chain Groups

Let Δ be a simplicial complex. For $d \geq -1$, let $\tilde{C}_d(\Delta; \mathbb{F})$ be the free $\mathbb{F}$-module with one basis element, denoted as $[s_1] \wedge \cdots \wedge [s_{d+1}]$, for each d-dimensional face $\{s_1, \ldots, s_{d+1}\}$ of Δ. This means that the rank of $\tilde{C}_d(\Delta; \mathbb{F})$ equals the number of faces of Δ of dimension d. By convention, we set $\tilde{C}_d(\Delta; \mathbb{F})$ equal to 0 for $d < -1$ and for $d > \dim \Delta$. For any permutation $\pi \in \mathfrak{S}_{[d+1]}$ and any face $\sigma = \{s_1, \ldots, s_{d+1}\}$, we define

$$[s_{\pi(1)}] \wedge [s_{\pi(2)}] \wedge \cdots \wedge [s_{\pi(d+1)}] = \mathrm{sgn}(\pi) \cdot [s_1] \wedge [s_2] \wedge \cdots \wedge [s_{d+1}]. \tag{3.1}$$

We will find it convenient to write $[\sigma] = [s_1] \wedge [s_2] \wedge \ldots \wedge [s_{d+1}]$, implicitly assuming that we have a fixed linear order on the 0-cells in Δ. Whenever σ and τ are disjoint faces such that $\sigma \cup \tau \in \Delta$, we define $[\sigma] \wedge [\tau]$ in the natural manner. Note that $[\emptyset] \wedge z = z$ for all z.

Boundary Map

The *boundary map* $\partial_d : \tilde{C}_d(\Delta; \mathbb{F}) \to \tilde{C}_{d-1}(\Delta; \mathbb{F})$ is the homomorphism defined by

$$\partial_d([s_1] \wedge \ldots \wedge [s_{d+1}]) = \sum_{i=1}^{d+1} (-1)^{i-1} [s_1] \wedge \ldots \wedge [s_{i-1}] \wedge [s_{i+1}] \wedge \ldots \wedge [s_{d+1}].$$

One easily checks that this definition is consistent with (3.1). Combining all ∂_d, we obtain an operator ∂ on the direct sum $\tilde{C}(\Delta; \mathbb{F})$ of all $\tilde{C}_d(\Delta; \mathbb{F})$. It is well-known and easy to see that $\partial^2 = 0$. This means that the pair $(\tilde{C}(\Delta; \mathbb{F}), \partial)$ forms a (graded) *chain complex*.

Let Δ_1 and Δ_2 be complexes on disjoint sets of 0-cells. Given any elements $c_1 \in \tilde{C}_{d_1}(\Delta_1; \mathbb{F})$ and $c_2 \in \tilde{C}_{d_2}(\Delta_2; \mathbb{F})$, the element $c_1 \wedge c_2 \in \tilde{C}_{d_1+d_2+1}(\Delta_1 * \Delta_2; \mathbb{F})$ satisfies the following identity:

$$\partial(c_1 \wedge c_2) = \partial(c_1) \wedge c_2 + (-1)^{d_1+1} c_1 \wedge \partial(c_2). \tag{3.2}$$

Homology

For the chain complex $(\tilde{C}(\Delta; \mathbb{F}), \partial)$ on the simplicial complex Δ, we refer to elements in $\partial^{-1}(\{0\})$ as *cycles* and elements in $\partial(\tilde{C}(\Delta; \mathbb{F}))$ as *boundaries*. Define the d^{th} *reduced homology group* of Δ with coefficients in $\mathbb{F}$ as the quotient $\mathbb{F}$-module

$$\tilde{H}_d(\Delta; \mathbb{F}) := \partial_d^{-1}(\{0\})/\partial_{d+1}(\tilde{C}_{d+1}(\Delta; \mathbb{F})) = \ker \partial_d / \mathrm{im}\, \partial_{d+1}.$$

Defining $\tilde{C}_{-1}(\Delta; \mathbb{F})$ to be zero, we obtain *unreduced* homology groups, denoted $H_d(\Delta; \mathbb{F})$ ("H" instead of "$\tilde{H}$"). We will be mainly concerned with reduced homology.

Just to give a simple example, we note that $\tilde{H}_d(\Delta; \mathbb{F}) = 0$ for all d whenever $\Delta = \mathsf{Cone}_x(\Sigma)$ for some Σ. Namely, we may write any element c in $\tilde{C}(\Delta; \mathbb{F})$ as $c = [x] \wedge c_1 + c_2$, where c_1 and c_2 are elements in $\tilde{C}(\Sigma; \mathbb{F})$. If c is a cycle, then $\partial(c_2) = -c_1$, which implies that $\partial([x] \wedge c_2) = c$; hence every cycle is a boundary.

Theorem 3.1 (see Munkres [101, Th. 25.1]). *For any pair of simplicial complexes Δ and Γ, we have the Mayer-Vietoris long exact sequence*

$$\begin{array}{cccccc} \cdots & \longrightarrow & \tilde{H}_{d+1}(\Delta;\mathbb{F}) \oplus \tilde{H}_{d+1}(\Gamma;\mathbb{F}) & \longrightarrow & \tilde{H}_{d+1}(\Delta \cup \Gamma;\mathbb{F}) & \\ \longrightarrow \; \tilde{H}_d(\Delta \cap \Gamma;\mathbb{F}) & \longrightarrow & \tilde{H}_d(\Delta;\mathbb{F}) \oplus \tilde{H}_d(\Gamma;\mathbb{F}) & \longrightarrow & \tilde{H}_d(\Delta \cup \Gamma;\mathbb{F}) & \\ \longrightarrow \tilde{H}_{d-1}(\Delta \cap \Gamma;\mathbb{F}) & \longrightarrow & \tilde{H}_{d-1}(\Delta;\mathbb{F}) \oplus \tilde{H}_{d-1}(\Gamma;\mathbb{F}) & \longrightarrow & \cdots & \end{array}$$

$\square$

Corollary 3.2. *Let Δ and Γ be simplicial complexes. Then the wedge $\Delta \vee \Gamma$ with respect to any identified* 0*-cells $x \in \Delta$ and $y \in \Gamma$ satisfies*

$$\tilde{H}_d(\Delta \vee \Gamma;\mathbb{F}) \cong \tilde{H}_d(\Delta;\mathbb{F}) \oplus \tilde{H}_d(\Gamma;\mathbb{F}).$$

for all $d \geq -1$.

Proof. We have that $\Delta \cap \Gamma = \{\emptyset, x\}$, which implies that $\tilde{H}_d(\Delta \cap \Gamma;\mathbb{F}) = 0$ for all d. By the Mayer-Vietoris sequence (Theorem 3.1), we are done. $\square$

Remark. Throughout this book, whenever we discuss the homology of a simplicial complex, we are referring to the reduced $\mathbb{Z}$-homology unless otherwise specified.

3.2 Relative Homology

Let $\Delta \subset \Gamma$ be two simplicial complexes. We refer to the family $\Gamma \setminus \Delta$ as a *quotient complex* and denote it as Γ/Δ. We define the *relative* chain complex of Γ/Δ in the following manner: Define the d^{th} chain group $\tilde{C}_d(\Gamma/\Delta;\mathbb{F})$ as the quotient group $\tilde{C}_d(\Gamma;\mathbb{F})/\tilde{C}_d(\Delta;\mathbb{F})$. This means that $\tilde{C}_d(\Gamma/\Delta;\mathbb{F})$ is a free $\mathbb{F}$-module with one generator $[\sigma]$ for each face $\sigma \in \Gamma \setminus \Delta$ of dimension d. Since the boundary map on $\tilde{C}_d(\Gamma;\mathbb{F})$ maps elements in $\tilde{C}_d(\Delta;\mathbb{F})$ to elements in $\tilde{C}_{d-1}(\Delta;\mathbb{F})$, this boundary map induces a boundary map $\partial_d : \tilde{C}_d(\Gamma/\Delta;\mathbb{F}) \to \tilde{C}_{d-1}(\Gamma/\Delta;\mathbb{F})$. If Δ is the void complex, then we obtain the ordinary chain complex of Γ.

Define the d^{th} *relative homology group* of Δ with coefficients in $\mathbb{F}$ as the quotient $\mathbb{F}$-module

$$\tilde{H}_d(\Gamma/\Delta;\mathbb{F}) := \partial_d^{-1}(\{0\})/\partial_{d+1}(\tilde{C}_d(\Delta/\Gamma;\mathbb{F})) = \ker \partial_d / \mathrm{im}\, \partial_{d+1}.$$

It is clear that this definition depends only on $\Gamma \setminus \Delta$. Specifically, we may replace Γ and Δ with any Γ' and Δ' such that $\Gamma' \setminus \Delta' = \Gamma \setminus \Delta$ without affecting the chain complex structure.

Note that the traditional notation is $\tilde{H}_d(\Gamma, \Delta;\mathbb{F})$ rather than the more streamlined $\tilde{H}_d(\Gamma/\Delta;\mathbb{F})$ that we have chosen.

Theorem 3.3 (see Munkres [101, Th. 23.3]). *For any pair of simplicial complexes $\Delta \subset \Gamma$, we have the following long exact sequence for the pair (Γ, Δ):*

$$\begin{array}{ccccccc} & \cdots & \longrightarrow & \tilde{H}_{d+1}(\Gamma;\mathbb{F}) & \longrightarrow & \tilde{H}_{d+1}(\Gamma/\Delta;\mathbb{F}) \\ \xrightarrow{f} & \tilde{H}_d(\Delta;\mathbb{F}) & \longrightarrow & \tilde{H}_d(\Gamma;\mathbb{F}) & \longrightarrow & \tilde{H}_d(\Gamma/\Delta;\mathbb{F}) \\ \xrightarrow{f} & \tilde{H}_{d-1}(\Delta;\mathbb{F}) & \longrightarrow & \tilde{H}_{d-1}(\Gamma;\mathbb{F}) & \longrightarrow & \cdots \end{array} \tag{3.3}$$

The map f is induced by the boundary operator ∂ in the chain complex of Γ. The other maps are defined in the natural manner. $\square$

A simple observation is that the relative homology of the pair (Γ, Δ) coincides with the simplicial homology of $\Gamma \cup \mathsf{Cone}(\Delta)$; consider the long exact sequence for the pair $(\Gamma \cup \mathsf{Cone}(\Delta), \mathsf{Cone}(\Delta))$ and observe that $\mathsf{Cone}(\Delta)$ has vanishing reduced homology in all dimensions.

Let $\sigma \in \Gamma$ and write $\Delta = \mathrm{fdel}_\Gamma(\sigma)$. It is immediate from the definition that

$$\tilde{H}_d(\Gamma/\Delta;\mathbb{F}) \cong \tilde{H}_{d-|\sigma|}(\mathrm{lk}_\Gamma(\sigma);\mathbb{F}).$$

By Theorem 3.3, we thus have a long exact sequence relating Γ and the link and face deletion of Γ with respect to σ. We will use this fact in Section 5.2.1 when we examine semi-nonevasive and semi-collapsible complexes.

In situations where there is no torsion, the homology of the Alexander dual of a complex is easy to compute via relative homology:

Theorem 3.4. *Let $\mathbb{F}$ be a field or $\mathbb{Z}$ and let Δ be a simplicial complex on a nonempty set X with $\mathbb{F}$-free homology. Then*

$$\tilde{H}_d(\Delta;\mathbb{F}) \cong \tilde{H}_{|X|-d-3}(\Delta^*_X;\mathbb{F}). \tag{3.4}$$

Proof. By Theorem 3.3, $\tilde{H}_d(\Delta;\mathbb{F}) \cong \tilde{H}_{d+1}(2^X/\Delta;\mathbb{F})$ for all d. Almost by definition, we have that $\tilde{H}_{d+1}(2^X/\Delta;\mathbb{F}) \cong \tilde{H}^{|X|-d-3}(\Delta^*_X;\mathbb{F})$, where $\tilde{H}^i(\Delta^*_X;\mathbb{F})$ denotes the i^{th} *cohomology group*; see Munkres [101]. Applying duality between homology and cohomology for complexes with free homology (see Munkres [101, Th. 45.8]), we obtain the desired result. $\square$

We cannot drop the condition that the homology be free; see Munkres [101].

3.3 Homotopy Theory

A *pointed space* is a topological space X together with a *base point* $x_0 \in X$. Let X and Y be pointed spaces with base points $x_0 \in X$ and $y_0 \in Y$. A *(pointed) map* from X to Y is a continuous function $f : X \to Y$ such that

$f(x_0) = y_0$. Let I be the interval $[0,1] = \{x \in \mathbb{R} : 0 \leq x \leq 1\}$. For maps $f, g : X \to Y$, a *homotopy* from f to g is a continuous function $F : I \times X \to Y$ such that $F_t(x_0) := F(t, x_0) = y_0$ for all $t \in I$ and such that $F_0(x) = f(x)$ and $F_1(x) = g(x)$ for all $x \in X$. We say that f and g are *homotopic* if such a homotopy exists.

X and Y are *homotopy equivalent*, denoted $X \simeq Y$, if there exist maps $f : X \to Y$ and $h : Y \to X$ such that $h \circ f : X \to X$ is homotopic to the identity map on X and $f \circ h : Y \to Y$ is homotopic to the identity map on Y. We will sometimes express this as saying that X has the *homotopy type* of Y. The choice of base point makes a difference only if the space is not path-connected. As almost all our spaces turn out to be path-connected, we will suppress the notion of base point from now on.

Lemma 3.5. *Let Y be a topological space and let X be a subspace. Suppose that there is a homotopy $F : I \times Y \to Y$ such that F_0 is the identity, the restriction of F_1 to X is the identity, and $F_1(Y) = X$. Then X and Y are homotopy equivalent.*

Proof. Define $f : Y \to X$ by $f(y) = F_1(y)$ and $g : X \to Y$ by $g(x) = F_0(x) = x$. We obtain that $f \circ g$ is the identity on X and that $g \circ f = F_1$. Since F_1 is homotopic to the identity F_0 on Y, we are done. □

Let Δ be a nonvoid abstract simplicial complex on a set X, say $X = [n]$. By some abuse of notation, we define the *topological realization* of Δ as any topological space homeomorphic to the following space $\|\Delta\|$: Let $\mathsf{e}_1, \ldots, \mathsf{e}_n$ be an orthonormal basis for Euclidean space $\mathbb{R}^n$. For a face σ, let $\|\sigma\|$ denote the set

$$\left\{ \sum_{x \in \sigma} \lambda_x \mathsf{e}_x : \sum_{x \in \sigma} \lambda_x = 1, \lambda_x > 0 \text{ for all } x \in \sigma \right\}. \tag{3.5}$$

Define $\|\Delta\|$ as the union $\bigcup_{\sigma \in \Delta} \|\sigma\|$; this is a disjoint union. Note that $\|2^\sigma\| = \bigcup_{\tau \subseteq \sigma} \|\tau\|$; this is the convex hull of the set $\{\mathsf{e}_x : x \in \sigma\}$. Also note that $\|\{x\}\| = \{\mathsf{e}_x\}$. We refer to $\|\Delta\|$ as the *canonical realization* of Δ

Let Δ and Γ be defined on two disjoint vertex sets X and Y. One easily checks that the canonical realization of the join $\Delta * \Gamma$ is the set

$$\{\lambda x + (1 - \lambda) y : x \in \|\Delta\|, y \in \|\Gamma\|, \lambda \in [0,1]\}.$$

The join operation preserves homeomorphisms and homotopies:

Lemma 3.6. *If $\|\Delta_1\| \cong \|\Delta_2\|$ and $\|\Gamma_1\| \cong \|\Gamma_2\|$, then $\|\Delta_1 * \Gamma_1\| \cong \|\Delta_2 * \Gamma_2\|$. If $\|\Delta_1\| \simeq \|\Delta_2\|$ and $\|\Gamma_1\| \simeq \|\Gamma_2\|$, then $\|\Delta_1 * \Gamma_1\| \simeq \|\Delta_2 * \Gamma_2\|$.*

Proof. Given homeomorphisms $f : \|\Delta_1\| \to \|\Delta_2\|$ and $g : \|\Gamma_1\| \to \|\Gamma_2\|$, a homeomorphism $h : \|\Delta_1 * \Gamma_1\| \to \|\Delta_2 * \Gamma_2\|$ is given by $h(\lambda x + (1 - \lambda) y) = \lambda f(x) + (1 - \lambda) g(y)$ for each $x \in \|\Delta_1\|$, $y \in \|\Gamma_1\|$, and $\lambda \in [0,1]$. This is well-defined, because we may extract λx from $\lambda x + (1 - \lambda) y$ by restricting to the coordinates corresponding to the elements in X, and we may extract λ from λx by summing the coordinates of λx.

In the same manner, one easily establishes the statement about homotopy equivalence. □

We say that an abstract simplicial complex Δ is homotopy equivalent to a pointed space X if the topological realization of Δ is homotopy equivalent to X. More generally, whenever we discuss topological properties of an abstract simplicial complex Δ, we are referring to its topological realization.

The void complex $\emptyset$ is by convention homotopy equivalent to a point (i.e., a 0-simplex).

We will frequently use the following well-known facts without reference; see Munkres [101] for details.

- Two simplicial complexes with the same homotopy type have the same homology (the converse is not true in general).
- The homotopy type of a wedge of two simplicial complexes Δ and Γ with respect to given identified 0-cells $x \in \Delta$ and $y \in \Gamma$ does not depend on the choice of x and y as long as each of Δ and Γ is connected.
- Any simplicial complex is homeomorphic to its first barycentric subdivision.

Occasionally, we will need to consider cell complexes. For a vector $\mathbf{x} = (x_1, \ldots, x_n)$, write $\|\mathbf{x}\| = \sqrt{x_1^2 + \ldots + x_n^2}$. The unit *$n$-ball* B^n is the set $\{\mathbf{x} = (x_1, \ldots, x_n) : \|\mathbf{x}\| \leq 1\}$ in $\mathbb{R}^n$. The unit *$(n-1)$-sphere* S^{n-1} is the boundary $\{\mathbf{x} = (x_1, \ldots, x_n) : \|\mathbf{x}\| = 1\}$ of B^n. By convention, B^0 is a point and S^{-1} is the empty set. Int $B^n = B^n \setminus S^{n-1}$ is the unit *open n-ball.* A topological space D is an *open n-cell* if D is homeomorphic to an open n-ball.

A Hausdorff topological space X is a finite *cell complex* if the following conditions are satisfied [101, §38]:

- X is the disjoint union of a finite number of open cells $\{D_i : i \in I\}$.
- For each open cell D_i, there is a continuous map
$$\varphi_i : B^{n_i} \to X$$
($n_i = \dim D_i$) such that the restriction of φ_i to Int B^{n_i} defines a homeomorphism to D_i and such that $\varphi_i(S^{n_i-1})$ is contained in the $(n_i - 1)$-skeleton of X (the union of all open cells D_j of dimension at most $n_i - 1$).
- A set C is closed in X if and only if $C \cap \overline{D}_i$ is closed in $\overline{D}_i$ for each cell D_i, where $\overline{D}_i = \varphi_i(B^{n_i})$.

The topological realization of a nonvoid simplicial complex Δ is a cell complex; for every face σ of Δ of dimension $d \geq 0$, the set $\|\sigma\|$ is homeomorphic to an open d-cell and the boundary of $\|\sigma\|$ is contained in the $(d-1)$-skeleton of $\|\Delta\|$. A simplicial complex is a *regular* cell complex, meaning that each map φ_i defines a homeomorphism to its image and $\varphi_i(S^{n_i-1})$ is equal to a union of smaller cells. We refer to Hatcher [59] or Munkres [101] for a more detailed exposition on cell complexes.

Some results in this book about simplicial complexes generalize to larger classes of cell complexes, but we will not state these generalizations unless we really need them.

We obtain a *wedge* of topological spaces $Y_1, \ldots, Y_r$ by taking the disjoint union of the spaces, choosing points $y_i \in Y_i$, and identifying the points $y_1, \ldots, y_r$. We may interpret a wedge X of spheres as a cell complex; the identified point y is a 0-cell and the space $X \setminus \{y\}$ is a disconnected space in which each component is a cell in X. Many simplicial complexes in this book are homotopy equivalent to such wedges of spheres.

3.4 Contractible Complexes and Their Relatives

We define the classes of acyclic, contractible, collapsible, and nonevasive complexes. In this book, we are particularly interested in the latter two classes, which we will generalize in Chapter 5.

3.4.1 Acyclic and k-acyclic Complexes

Let $\mathbb{F}$ be a field or $\mathbb{Z}$. A simplicial complex Δ is *acyclic* over $\mathbb{F}$ or $\mathbb{F}$-*acyclic* if Δ has no reduced homology over $\mathbb{F}$. By the universal coefficient theorem [59, Th. 3A.3], a complex Δ is $\mathbb{Z}$-acyclic if and only if Δ is $\mathbb{F}$-acyclic for each field $\mathbb{F}$. However, for any field $\mathbb{F}$, there exist $\mathbb{F}$-acyclic complexes that are not $\mathbb{Z}$-acyclic. For example, any triangulation of the real projective plane (e.g., the one in Figure 5.3 in Section 5.2.1) is $\mathbb{F}$-acyclic whenever $\mathbb{F}$ is a field of odd or zero characteristic but not $\mathbb{Z}_2$-acyclic or $\mathbb{Z}$-acyclic.

A complex Δ is k-*acyclic* over $\mathbb{F}$ if the homology group $\tilde{H}_d(\Delta;\mathbb{F})$ vanishes for $d \le k$. If a complex Δ is k-acyclic over $\mathbb{Z}$, then Δ is k-acyclic over $\mathbb{F}$ for every field, but the converse is again false for $k \ge 1$.

Proposition 3.7. *Let $d_1, d_2 \ge 0$. If Δ is (d_1-1)-acyclic over $\mathbb{F}$ and Γ is (d_2-1)-acyclic over $\mathbb{F}$, then $\Delta * \Gamma$ is (d_1+d_2)-acyclic over $\mathbb{F}$.*

Proof. Throughout this proof, c_i and $\hat{c}_i$ denote elements in $\tilde{C}_i(\Delta;\mathbb{F})$ and c'_j denotes an element in $\tilde{C}_j(\Gamma;\mathbb{F})$. Let $a \le d_1 + d_2$ and let z be a nonzero cycle in $\tilde{C}_a(\Delta * \Gamma;\mathbb{F})$. We can write

$$z = \sum_{i=r}^{s} c_i \wedge c'_{a-i-1} \tag{3.6}$$

for some $r \le s$, where the first term and the last term are both nonzero. It is clear that c_r and c'_{a-s-1} are cycles. Since $a \le d_1 + d_2$ and $s \ge r$, we cannot simultaneously have that $r \ge d_1$ and $a - s - 1 \ge d_2$. By symmetry, we may assume that $r \leq d_1 - 1$; hence there is an element $\hat{c}_{r+1}$ such that $\partial(\hat{c}_{r+1}) = c_r$. Consider the element $\hat{z} = \partial(\hat{c}_{r+1} \wedge c'_{a-r-1}) = c_r \wedge c'_{a-r-1} \pm \hat{c}_{r+1} \wedge \partial(c'_{a-r-1})$. If $r = s$, then $\hat{z} = z$; hence z is a boundary. Otherwise, $z - \hat{z}$ is a sum as in (3.6) but from $r + 1$ to s. By induction on $s - r$, $z - \hat{z}$ is a boundary, which concludes the proof. □

3.4.2 Contractible and k-connected Complexes

A simplicial complex Δ is *contractible* if Δ is homotopy equivalent to a single point. A contractible complex Δ is acyclic over $\mathbb{Z}$, but the converse is not necessarily true unless Δ is simply connected; the famous Poincaré homology 3-sphere [106] is one example. For $k \geq 0$, a topological space X is *k-connected* if the following holds for all $d \in [0, k]$:

- Every continuous map $f : S^d \to X$ has a continuous extension $g : B^{d+1} \to X$.

By convention, X is (-1)-connected if and only if X is nonempty. Note that X is 0-connected if and only if X is path-connected. One typically refers to 1-connected complexes as *simply connected*. The *connectivity degree* of X is the largest integer k such that X is k-connected ($+\infty$ if X is k-connected for all k). Increasing the connectivity degree by one, we obtain the *shifted* connectivity degree; this value is the smallest integer k such that X is not k-connected. In many situations, the shifted connectivity degree coincides with the smallest integer d such that the homology in dimension d is nonvanishing:

Theorem 3.8 (see Hatcher [59, Th. 4.32]). *For $k \geq 1$, a simplicial complex Δ is k-connected if and only if Δ is k-acyclic over $\mathbb{Z}$ and simply connected. Δ is contractible if and only if Δ is acyclic over $\mathbb{Z}$ and simply connected. For $k \in \{-1, 0\}$, a complex Δ is k-connected if and only if Δ is k-acyclic.* □

Corollary 3.9. *For $k \geq 0$, if Δ_1 and Δ_2 are k-connected and $\Delta_1 \cap \Delta_2$ is $(k-1)$-connected, then $\Delta_1 \cup \Delta_2$ is k-connected.*

Proof. The corollary is clear if $k = 0$. Assume that $k \geq 1$. By the Mayer-Vietoris exact sequence (Theorem 3.1), $\Delta_1 \cup \Delta_2$ has no homology below dimension k. Now, Δ_1 and Δ_2 are simply connected, whereas $\Delta_1 \cap \Delta_2$ is path-connected. As a consequence, $\Delta_1 \cup \Delta_2$ is simply connected by the van Kampen theorem (see Hatcher [59, Th. 1.20]). Thus we are done by Theorem 3.8. □

Corollary 3.10. *If Δ is a k-connected subcomplex of Γ and the dimension of each face of $\Gamma \setminus \Delta$ is at least $k + 1$, then Γ is k-connected.*

Proof. We are done if $\Delta = \Gamma$. Otherwise, let σ be a maximal face of $\Gamma \setminus \Delta$; by assumption, $\dim \sigma > k$. By induction, $\Gamma \setminus \{\sigma\}$ is k-connected. Now, 2^σ is k-connected, whereas $\partial 2^\sigma$ is $(k-1)$-connected. Since $\Gamma = (\Gamma \setminus \{\sigma\}) \cup 2^\sigma$ and $\partial 2^\sigma = (\Gamma \setminus \{\sigma\}) \cap 2^\sigma$, Corollary 3.9 yields that Γ is k-connected. □

Theorem 3.11. *If Δ is connected and* $\dim \Gamma \geq 0$ *(i.e., Γ is (-1)-connected), then $\Delta * \Gamma$ is simply connected.*

Proof. If $\Gamma = \{\emptyset, \{x\}\}$, then $\Delta * \Gamma$ is a cone and hence simply connected. Otherwise, let x be a 0-cell in Γ and write $\Gamma_1 = \mathrm{del}_\Gamma(x)$ and $\Gamma_2 = \mathsf{Cone}_x(\mathrm{lk}_\Gamma(x))$. It is clear that $\Gamma = \Gamma_1 \cup \Gamma_2$ and that $\Delta * (\Gamma_1 \cap \Gamma_2)$ is connected. By induction, $\Delta * \Gamma_1$ and $\Delta * \Gamma_2$ are simply connected; each of Γ_1 and Γ_2 is (-1)-connected. By Corollary 3.9, it follows that $\Delta * \Gamma$ is simply connected. □

Corollary 3.12. *Let $d_1, d_2 \geq 0$. If Δ is $(d_1 - 1)$-acyclic over $\mathbb{Z}$ and Γ is $(d_2 - 1)$-acyclic over $\mathbb{Z}$, then $\Delta * \Gamma$ is $(d_1 + d_2)$-connected.*

Proof. The corollary is clearly true for $d_1 = d_2 = 0$. Assume that $d_1 + d_2 \geq 1$. By Proposition 3.7, $\Delta * \Gamma$ is $(d_1 + d_2)$-acyclic. Theorem 3.11 yields that $\Delta * \Gamma$ is simply connected; hence we are done by Theorem 3.8. □

Theorem 3.13. *Let $d \geq 0$. If Δ is $(d-1)$-connected and* $\dim \Delta \leq d$*, then Δ is homotopy equivalent to a wedge of spheres of dimension d.*

Proof. The theorem is trivial for $d = 0$. If $d = 1$, then Δ is a connected graph, which is homotopy equivalent to a wedge of circles. Otherwise, Δ is simply connected and $(d-1)$-acyclic by Theorem 3.8. As a consequence, all homology of Δ is concentrated in dimension d. Since $\dim \Delta \leq d$, this homology must be torsion-free and hence of the form $\mathbb{Z}^r$ for some $r \geq 0$. By the homology version of Whitehead's theorem (see Hatcher [59, Prop. 4C.1]), this implies that Δ is homotopy equivalent to a cell complex consisting of r cells of dimension d and one 0-cell, hence a wedge of r spheres of dimension d. □

3.4.3 Collapsible Complexes

Recall that a complex is collapsible if the complex is void or can be collapsed to a point $\{\emptyset, \{v\}\}$. Collapsible complexes are contractible, but not all contractible complexes are collapsible; the dunce hat [150] is one example. One may characterize collapsible complexes in the following manner:

Definition 3.14. We define the class of *collapsible* simplicial complexes recursively as follows:

(i) The void complex $\emptyset$ and any 0-simplex $\{\emptyset, \{v\}\}$ are collapsible.
(ii) If Δ contains a nonempty face σ such that the face-deletion $\mathrm{fdel}_\Delta(\sigma)$ and the link $\mathrm{lk}_\Delta(\sigma)$ are collapsible, then Δ is collapsible.

We discuss further properties of collapsible complexes in Section 5.4.

3.4.4 Nonevasive Complexes

To obtain the class of nonevasive complexes, we use Definition 3.14 with the restriction that the face σ in (ii) must be a 0-cell:

Definition 3.15. We define the class of *nonevasive* simplicial complexes recursively as follows:

(i) The void complex $\emptyset$ and any 0-simplex $\{\emptyset, \{v\}\}$ are nonevasive.
(ii) If Δ contains a 0-cell x such that $\mathrm{del}_\Delta(x)$ and $\mathrm{lk}_\Delta(x)$ are nonevasive, then Δ is nonevasive.

For example, cones are nonevasive. A complex is *evasive* if it is not nonevasive. We explain this terminology in Chapter 5. As Kahn, Saks, and Sturtevant [78] observed, nonevasive complexes are collapsible. The converse is not true in general; in Proposition 5.13, we present a counterexample due to Björner. We discuss further properties of nonevasive complexes in Section 5.4.

3.5 Quotient Complexes

Let $X \subseteq Y$ be two topological spaces such that X is nonempty. Let p be an isolated point not in Y. One defines the *quotient space* Y/X as the set $(Y \setminus X) \cup \{p\}$ equipped with the topology induced by the map $\alpha : Y \to (Y \setminus X) \cup \{p\}$ defined by

$$\alpha(x) = \begin{cases} x \text{ if } x \in Y \setminus X; \\ p \text{ if } x \in X. \end{cases}$$

That is, M is open in Y/X if and only if $\alpha^{-1}(M)$ is open in Y. By convention, we set $Y/\emptyset$ equal to the union of Y and a discrete point $\{p\}$ not in Y.

Let $\Delta \subseteq \Gamma$ be simplicial complexes such that Δ is nonvoid. We define the *topological realization* of the quotient complex Γ/Δ to be any space homeomorphic to $\|\Gamma\|/\|\Delta\|$. One easily checks directly from the definition that $\|\Gamma\|/\|\Delta\|$ is homeomorphic to $\|\Gamma'\|/\|\Delta'\|$ whenever $\Gamma \setminus \Delta = \Gamma' \setminus \Delta'$. Note that $\|\Gamma\|/\|\{\emptyset\}\| = \|\Gamma\| \cup \{p\}$, because $\|\{\emptyset\}\| = \emptyset$.

One may interpret the space $\|\Gamma\|/\|\Delta\|$ as a cell complex. Specifically, we have one cell $\|\sigma\|$ for each face $\sigma \in \Gamma \setminus \Delta$ plus one additional 0-cell $\{p\}$ corresponding to Δ. The boundary of $\|2^\sigma\|$ is the same as in $\|\Gamma\|$ except that we identify all points in $\|\partial 2^\sigma\| \cap \|\Delta\|$ with p.

Whenever we talk about the topology of Γ/Δ, we are referring to the space $\|\Gamma\|/\|\Delta\|$. The following lemma is known as the *Contractible Subcomplex Lemma*.

Lemma 3.16 (see Hatcher [59, Prop. 0.17]). *Let Γ and Δ be simplicial complexes such that Δ is a contractible subcomplex of Γ. Then Γ/Δ and Γ are homotopy equivalent.*

Proof. Let E be the set of 0-cells in Γ. It is well-known and easy to prove that there is a homeomorphism from $\|\Gamma\|$ to $\|\mathrm{sd}(\Gamma)\|$ such that restriction to $\|\Delta\|$ is a homeomorphism to $\|\mathrm{sd}(\Delta)\|$. In particular, $\|\Gamma\|/\|\Delta\|$ and $\|\mathrm{sd}(\Gamma)\|/\|\mathrm{sd}(\Delta)\|$ are homeomorphic. As a consequence, we may assume without loss of generality that Δ coincides with the induced subcomplex of Γ on some set $E_0 \subset E$ of 0-cells; thus $\Delta = \Gamma \cap 2^{E_0}$. Let $\Delta^\perp$ be the induced subcomplex on the set $E \setminus E_0$.

Let $F : I \times \|\Delta\| \to \|\Delta\|$ be a homotopy from the identity to a constant function; $F_0(x) = x$ and $F_1(x) = y$ for some $y \in \|\Delta\|$. Each element x in $\|\Gamma\|$ has a unique representation $x = \lambda q + (1-\lambda)r$, where $q \in \|\Delta\|$, $r \in \|\Delta^\perp\|$, and $\lambda \in I$. Define $G : I \times \|\Gamma\| \to \|\Gamma\|$ to be the homotopy given by

$$G_t(\lambda q + (1-\lambda)r) = \begin{cases} (1+t)\lambda q + (1-(1+t)\lambda)r & \text{if } \lambda \le 1/(1+t); \\ F_{(t+1)\lambda-1}(q) & \text{if } \lambda \ge 1/(1+t) \end{cases}$$

for all relevant $q \in \|\Delta\|$ and $r \in \|\Delta^\perp\|$. This is indeed a homotopy, because $\lambda = 1/(t+1)$ yields the same result q in both formulas.

We have that G_t induces a homotopy $\tilde{G}_t : \|\Gamma\|/\|\Delta\| \to \|\Gamma\|/\|\Delta\|$. Moreover, G_1 induces a continuous map $\hat{G}_1 : \|\Gamma\|/\|\Delta\| \to \|\Gamma\|$; G_1 maps the entirety of $\|\Delta\|$ to $F_1(\|\Delta\|) = \{y\}$. Define $\alpha : \|\Gamma\| \to \|\Gamma\|/\|\Delta\|$ to be the projection map. Now, $\hat{G}_1 \circ \alpha = G_1$, which is homotopic to the identity G_0. Moreover, $\alpha \circ \hat{G}_1 = \tilde{G}_1$, which is homotopic to the identity $\tilde{G}_0$; hence we are done. □

Corollary 3.17. *Let Γ be a simplicial complex and let Δ be a subcomplex of Γ. Let Σ be a complex on a 0-cell set disjoint from the 0-cell set of Γ such that $\Sigma * \Delta$ is contractible. Then Γ/Δ is homotopy equivalent to $\Gamma \cup (\Sigma * \Delta)$.*

Proof. Since $\Gamma/\Delta = (\Gamma \cup (\Sigma * \Delta))/(\Sigma * \Delta)$, the Contractible Subcomplex Lemma 3.16 implies the desired result. □

Lemma 3.18. *Let Γ be a contractible simplicial complex and let Δ be a subcomplex of Γ. Then Γ/Δ is homotopy equivalent to the suspension of Δ. Moreover, $\tilde{H}_{i+1}(\Gamma/\Delta;\mathbb{F}) = \tilde{H}_i(\Delta;\mathbb{F})$ for $i \ge -1$.*

Proof. Let x and y be two 0-cells not in Γ. By Corollary 3.17, Γ/Δ is homotopy equivalent to $\Gamma \cup \mathsf{Cone}_x(\Delta)$. Since Γ is contractible, the Contractible Subcomplex Lemma 3.16 implies that $\Gamma \cup \mathsf{Cone}_x(\Delta)$ is homotopy equivalent to $(\Gamma \cup \mathsf{Cone}_x(\Delta))/\Gamma$ and hence to $\mathsf{Cone}_x(\Delta)/\Delta$. Another application of Corollary 3.17 yields that $\mathsf{Cone}_x(\Delta)/\Delta$ is homotopy equivalent to $\mathsf{Cone}_x(\Delta) \cup \mathsf{Cone}_y(\Delta) = \mathrm{Susp}_{x,y}(\Delta)$, which concludes the proof. For the last claim, use the long exact sequence in Theorem 3.3. □

In this context, it might be worth stating the following fact about suspensions.

Lemma 3.19 (Björner and Welker [16, Lemma 2.5]). *If $\Delta \simeq \bigvee_{i\in I} S^{d_i}$, then* $\mathrm{Susp}(\Delta) \simeq \bigvee_{i\in I} S^{d_i+1}$. □

The converse is not true. For example, the suspension of a d-dimensional complex with homology only in top dimension $d \geq 1$ is simply connected by Theorem 3.11 and hence homotopy equivalent to a wedge of spheres by Theorem 3.13.

The following lemma is a special case of a much more general result about homotopy type being preserved under join.

Lemma 3.20. *Let Δ be a simplicial complex and let Γ and Γ' be quotient complexes. If $\Gamma \simeq \Gamma'$, then $\Delta * \Gamma \simeq \Delta * \Gamma'$.*

Proof. Write $\Gamma = \Gamma_1/\Gamma_0$, where Γ_1 and Γ_0 are simplicial complexes. By Corollary 3.17,

$$\Delta * \Gamma = \frac{\Delta * \Gamma_1}{\Delta * \Gamma_0} \simeq (\Delta * \Gamma_1) \cup \mathsf{Cone}_x(\Delta * \Gamma_0) = \Delta * (\Gamma_1 \cup \mathsf{Cone}_x(\Gamma_0)).$$

By Corollary 3.17 and Lemma 3.6, we obtain that the homotopy type of $\Delta * \Gamma$ is uniquely determined by the homotopy type of each of Δ and Γ, which concludes the proof. □

3.6 Shellable Complexes and Their Relatives

We define the classes of Cohen-Macaulay, constructible, shellable, and vertex-decomposable complexes along with nonpure versions. In Section 3.6.5, we present some basic topological results about these complexes. For our purposes, the class of vertex-decomposable complexes is by far the most important. See Section 6.3 for some specific results related to this class.

3.6.1 Cohen-Macaulay Complexes

Definition 3.21. Let Δ be a pure simplicial complex. Δ is *homotopically Cohen-Macaulay* (CM) if $\mathrm{lk}_\Delta(\sigma)$ is $(\dim \mathrm{lk}_\Delta(\sigma) - 1)$-connected for each σ in Δ. Let $\mathbb{F}$ be a field or $\mathbb{Z}$. Δ is *Cohen-Macaulay over* $\mathbb{F}$ (denoted as $CM/\mathbb{F}$) if $\mathrm{lk}_\Delta(\sigma)$ is $(\dim \mathrm{lk}_\Delta(\sigma) - 1)$-acyclic for each σ in Δ.

By Theorem 3.13, $\mathrm{lk}_\Delta(\sigma)$ is $(\dim \mathrm{lk}_\Delta(\sigma) - 1)$-connected if and only if $\mathrm{lk}_\Delta(\sigma)$ is homotopy equivalent to a wedge of spheres of dimension $\dim \mathrm{lk}_\Delta(\sigma)$. See Section 3.8 for the ring-theoretic motivation of Definition 3.21.

Define the *homotopical depth* of a complex Δ as the largest integer k such that the k-skeleton of Δ is homotopically CM. Define the *depth over* $\mathbb{F}$ of Δ as the largest integer k such that the k-skeleton of Δ is $CM/\mathbb{F}$. Equivalently, the depth over $\mathbb{F}$ equals

$$\min\{m : \tilde{H}_{m-|\sigma|}(\mathrm{lk}_\Delta(\sigma), \mathbb{F}) \neq 0 \text{ for some } \sigma \in \Delta\}.$$

This is closely related to the ring-theoretic concept of depth; see Section 3.8.

Define the *pure d-skeleton* $\Delta^{[d]}$ of Δ as the subcomplex of Δ generated by all d-dimensional faces of Δ. Stanley [132] extended the concept of Cohen-Macaulayness to nonpure complexes:

Definition 3.22. A simplicial complex Δ is *sequentially homotopy-CM* if the pure d-skeleton $\Delta^{[d]}$ is homotopically CM for every $d \geq 0$. Let $\mathbb{F}$ be a field or $\mathbb{Z}$. Δ is *sequentially $CM/\mathbb{F}$* if the pure d-skeleton $\Delta^{[d]}$ is $CM/\mathbb{F}$ for every $d \geq 0$.

3.6.2 Constructible Complexes

Definition 3.23. We define the class of *constructible* simplicial complexes recursively as follows:

(i) Every simplex (including $\emptyset$ and $\{\emptyset\}$) is constructible.
(ii) If Δ_1 and Δ_2 are constructible complexes of dimension d and $\Delta_1 \cap \Delta_2$ is a constructible complex of dimension $d-1$, then $\Delta_1 \cup \Delta_2$ is constructible.

Hochster [63] introduced constructible complexes.

Let us extend the concept of constructibility to nonpure complexes. For a simplicial complex Δ, define $\mathcal{F}(\Delta)$ to be the family of maximal faces of Δ.

Definition 3.24. We define the class of *semipure constructible* simplicial complexes recursively as follows:

(i) Every simplex (including $\emptyset$ and $\{\emptyset\}$) is semipure constructible.
(ii) Suppose that Δ_1, Δ_2, and $\Gamma = \Delta_1 \cap \Delta_2$ are semipure constructible complexes such that the following conditions are satisfied:
 (a) $\mathcal{F}(\Delta_1 \cup \Delta_2)$ is the disjoint union of $\mathcal{F}(\Delta_1)$ and $\mathcal{F}(\Delta_2)$.
 (b) Every member of $\mathcal{F}(\Gamma)$ is a maximal face of either $\Delta_1 \setminus \mathcal{F}(\Delta_1)$ or $\Delta_2 \setminus \mathcal{F}(\Delta_2)$ (possibly of both).
 Then $\Delta_1 \cup \Delta_2$ is semipure constructible.

Expressed in terms of pure skeletons, condition (b) is equivalent to the identity

$$\Delta_1{}^{[d]} \cap \Delta_2{}^{[d]} = \Gamma^{[d-1]} \cup \Gamma^{[d]}$$

for each d.

One may refer to semipure constructible complexes that are not pure as *nonpure constructible.*

3.6.3 Shellable Complexes

The class of shellable complexes is arguably the most well-studied class of Cohen-Macaulay complexes. Indeed, proving shellability is in many situations the most efficient way of establishing Cohen-Macaulayness; see Björner and Wachs [12] for just one of many examples. In this respect, this book constitutes an exception, as our proofs of Cohen-Macaulayness typically go via vertex-decomposability (see Section 3.6.4). Therefore, we confine ourselves to presenting basic definitions and refer the interested reader to Björner [9] for more information and further references.

Definition 3.25. We define the class of *shellable* simplicial complexes recursively as follows:

(i) Every simplex (including $\emptyset$ and $\{\emptyset\}$) is shellable.
(ii) If Δ is pure and contains a nonempty face σ – a *shedding face* – such that $\mathrm{fdel}_\Delta(\sigma)$ and $\mathrm{lk}_\Delta(\sigma)$ are shellable, then Δ is also shellable.

This way of defining shellability is easily seen to be equivalent to more conventional approaches; see Provan and Billera [108].

We say that a lifted complex $\Sigma = \Delta * \{\rho\}$ (see Section 2.3.7) is shellable if the underlying simplicial complex Δ is shellable. A sequence $(\Sigma_1, \ldots, \Sigma_r = \Sigma)$ is a *shelling* of Σ if each Σ_i is a pure lifted complex over ρ of dimension $\dim \Sigma$ such that $\Sigma_i \setminus \Sigma_{i-1}$ has a unique maximal face τ_i and a unique minimal face σ_i for each $i \in [1, r]$; $\Sigma_0 = \emptyset$. The i^{th} *shelling pair* is the pair (σ_i, τ_i). Note that $\sigma_1 = \rho$.

Let Δ be a lifted complex over ρ. The recursive procedure in (ii) of Definition 3.25 gives rise to a shelling of Δ. Specifically, assume inductively that we have shellings $(\Delta_1, \ldots, \Delta_q)$ of $\mathrm{fdel}_\Delta(\sigma)$ and $(\Delta_{q+1}, \ldots, \Delta_r)$ of $\Delta(\sigma, \emptyset)$ (we lift the link $\mathrm{lk}_\Delta(\sigma)$). If $\Delta = \mathrm{lk}_\Delta(\sigma) * 2^\sigma$, then $(\Delta_{q+1} * 2^\sigma, \ldots, \Delta_r * 2^\sigma)$ is a shelling of Δ. Otherwise, $(\Delta_1, \ldots, \Delta_q, \Delta_{q+1}, \ldots, \Delta_r)$ is a shelling of Δ; the unique minimal element in $\Delta_{q+1} \setminus \Delta_q$ is σ.

Conversely, it is easy to prove that Δ admits a shelling if and only if Δ is shellable in terms of Definition 3.25; use the last minimal face σ_r as the first shedding face.

Björner and Wachs [13] extended shellability to complexes that are not necessarily pure:

Definition 3.26. We define the class of *semipure shellable* simplicial complexes recursively as follows:

(i) Every simplex (including $\emptyset$ and $\{\emptyset\}$) is semipure shellable.
(ii) If Δ contains a nonempty face σ – a *shedding face* – such that $\mathrm{fdel}_\Delta(\sigma)$ and $\mathrm{lk}_\Delta(\sigma)$ are semipure shellable and such that every maximal face of $\mathrm{fdel}_\Delta(\sigma)$ is a maximal face of Δ, then Δ is also semipure shellable.

To see that Definition 3.26 is equivalent to the original definition [13, Def. 2.1], adapt the proof of Björner and Wachs [14, Th. 11.3]. One may refer to semipure shellable complexes that are not pure as *nonpure shellable*.

3.6.4 Vertex-Decomposable Complexes

Definition 3.27. We define the class of *vertex-decomposable* (VD) simplicial complexes recursively as follows:

(i) Every simplex (including $\emptyset$ and $\{\emptyset\}$) is VD.
(ii) If Δ is pure and contains a 0-cell x – a *shedding vertex* – such that $\mathrm{del}_\Delta(x)$ and $\mathrm{lk}_\Delta(x)$ are VD, then Δ is also VD.

Vertex-decomposable complexes were introduced by Provan and Billera [108].

As for shellability, one readily extends vertex-decomposability to lifted complexes. An alternative approach to vertex-decomposability including lifted complexes is as follows:

Definition 3.28. We define the class of VD lifted complexes recursively as follows.

(i) Every simplex (including $\emptyset$ and $\{\emptyset\}$) is VD.
(ii) If Δ contains a 0-cell v such that $\Delta(v, \emptyset)$ and $\Delta(\emptyset, v)$ are VD of the same dimension, then Δ is also VD.
(iii) If Δ is a cone over a VD complex Δ', then Δ is also VD.
(iv) If $\Delta = \Sigma * \{\sigma\}$ and Σ is VD, then Δ is also VD.

The restriction of this definition to simplicial complexes is easily seen to be equivalent to the original Definition 3.27.

Just as for shellability, Björner and Wachs [13] extended the concept of vertex-decomposability to nonpure complexes:

Definition 3.29. We define the class of *semipure* VD simplicial complexes recursively as follows:

(i) Every simplex (including $\emptyset$ and $\{\emptyset\}$) is semipure VD.
(ii) If Δ contains a 0-cell x – a *shedding vertex* – such that $\mathrm{del}_\Delta(x)$ and $\mathrm{lk}_\Delta(x)$ are semipure VD and such that every maximal face of $\mathrm{del}_\Delta(x)$ is a maximal face of Δ, then Δ is also semipure VD.

One may refer to semipure VD complexes that are not pure as *nonpure* VD.

3.6.5 Topological Properties and Relations Between Different Classes

Theorem 3.30. *The properties of being CM, sequentially CM, constructible, semipure constructible, shellable, semipure shellable, VD, and semipure VD are all closed under taking link and join.*

Proof. The properties being closed under taking link is straightforward to prove in all cases. The $CM/\mathbb{F}$ and sequentially $CM/\mathbb{F}$ properties are closed under taking join, because the join of a (d_1-1)-acyclic complex and a (d_2-1)-acyclic complex is (d_1+d_2)-acyclic by Proposition 3.7. By Corollary 3.12, the homotopically CM and sequentially homotopy-CM properties are also closed under taking join.

For the remaining properties, use a simple induction argument, decomposing with respect to the first complex in the join and keeping the other complex fixed. In each case, the base case is a join of two simplices, which is again a simplex.

For example, suppose that $\Delta = \Delta_1 \cup \Delta_2$ and Δ' are semipure constructible and that Δ_1 and Δ_2 satisfy the properties in (ii) in Definition 3.24. By induction, $\Delta_1 * \Delta'$, $\Delta_2 * \Delta'$, and $(\Delta_1 \cap \Delta_2) * \Delta'$ are all semipure constructible. Moreover, one easily checks that conditions (a) and (b) in Definition 3.24 hold for $\Delta_1 * \Delta'$ and $\Delta_2 * \Delta'$. It hence follows that $(\Delta_1 \cup \Delta_2) * \Delta'$ is semipure constructible as desired.

The treatment of the other properties is equally straightforward. □

Proposition 3.31. *The following properties hold for any pure simplicial complex Δ:*

(i) *Δ is sequentially CM if and only if Δ is CM in the sense of Definition* 3.21.
(ii) *Δ is semipure constructible if and only if Δ is constructible in the sense of Definition 3.23.*
(iii) *Δ is semipure shellable if and only if Δ is shellable in the sense of Definition 3.25.*
(iv) *Δ is semipure VD if and only if Δ is VD in the sense of Definition 3.27.*

Proof. (i) This is obvious.

(ii) Constructible complexes are easily seen to be semipure constructible. The other direction is obvious if Δ satisfies (i) in Definition 3.24. Suppose that $\Delta = \Delta_1 \cup \Delta_2$ and that the conditions in (ii) are satisfied. Condition (a) yields that Δ_1 and Δ_2 are pure, whereas condition (b) yields that their intersection is pure of dimension one less than Δ_1 and Δ_2. By induction, all these three complexes are constructible, which implies that the same is true for Δ.

(iii) It is clear that Δ is shellable if Δ is semipure shellable. The other direction is immediate, except that we need to check the case that we have a shedding face σ in Definition 3.25 such that the dimension of $\mathrm{fdel}_\Delta(\sigma)$ is strictly smaller than that of Δ. This implies that $\Delta = 2^\sigma * \mathrm{lk}_\Delta(\sigma)$. Namely, Δ is generated by the maximal faces of the lifted complex $\Delta(\sigma, \emptyset)$. By Theorem 3.30, semipure shellability is preserved under join, which implies that Δ is semipure shellable as desired.

(iv) This is proved in exactly the same manner as (iii). □

Lemma 3.32. *Let Δ_1 and Δ_2 be homotopically CM of dimension d such that the $(d-1)$-skeleton of $\Delta_1 \cap \Delta_2$ is homotopically CM. Then $\Delta_1 \cup \Delta_2$ is homotopically CM.*

Proof. Note that $\mathrm{lk}_{\Delta_1 \cup \Delta_2}(\sigma) = \mathrm{lk}_{\Delta_1}(\sigma) \cup \mathrm{lk}_{\Delta_2}(\sigma)$ and analogously for the intersection. As a consequence, the lemma follows immediately from Corollary 3.9. □

Theorem 3.33. *We have the following implications:*

$$VD \Longrightarrow \textit{Shellable} \Longrightarrow \textit{Constructible} \Longrightarrow \textit{Homotopically CM}.$$

The analogous implications hold for the semipure variants.

Proof. By Proposition 3.31, it suffices to prove that the implications hold for the semipure variants.

Semipure $VD \Longrightarrow$ Semipure shellable. Trivial.

Semipure shellable $\Longrightarrow$ Semipure constructible. The theorem is obvious if Δ is a simplex. Otherwise, let σ be a shedding face as in (ii) in Definition 3.26. Write $\Delta_1 = \mathrm{fdel}_\Delta(\sigma)$ and $\Delta_2 = 2^\sigma * \mathrm{lk}_\Delta(\sigma)$; it is clear that $\Delta = \Delta_1 \cup \Delta_2$. Now, Δ_1 is semipure shellable by assumption. Moreover, by Theorem 3.30, $\Delta_2 = 2^\sigma * \mathrm{lk}_\Delta(\sigma)$ is semipure shellable. Finally, the intersection $\Delta_1 \cap \Delta_2$ equals $\partial 2^\sigma * \mathrm{lk}_\Delta(\sigma)$. The boundary of a simplex is well-known to be shellable and hence semipure shellable; hence $\Delta_1 \cap \Delta_2$ is semipure shellable. Induction yields that all these complexes are semipure constructible.

It remains to prove that conditions (a) and (b) in Definition 3.24 are satisfied. Condition (a) follows immediately from Definition 3.26. To prove condition (b), consider a maximal face $(\sigma - x) \cup \tau$ of $\partial 2^\sigma * \mathrm{lk}_\Delta(\sigma)$; $\tau \in \mathcal{F}(\mathrm{lk}_\Delta(\sigma))$. One easily checks that the only maximal face of $\Delta_2 = 2^\sigma * \mathrm{lk}_\Delta(\sigma)$ containing this face is $\sigma \cup \tau$; thus we are done.

Semipure constructible $\Longrightarrow$ Sequentially homotopy-CM. This is obvious if Δ satisfies (i) in Definition 3.24. Suppose that $\Delta = \Delta_1 \cup \Delta_2$ and that the conditions in (ii) are satisfied. We need to prove that the pure d-skeleton $\Delta^{[d]}$ is CM for every $d \geq 0$.

By induction, each of ${\Delta_1}^{[d]}$ and ${\Delta_2}^{[d]}$ is CM. Moreover, by construction, ${\Delta_1}^{[d]} \cap {\Delta_2}^{[d]} = \Gamma^{[d-1]} \cup \Gamma^{[d]}$, where $\Gamma = \Delta_1 \cup \Delta_2$. By induction, each of $\Gamma^{[d-1]}$ and $\Gamma^{[d]}$ is CM. Their intersection equals the $(d-1)$-skeleton of $\Gamma^{[d]}$ and is hence CM. As a consequence, Lemma 3.32 yields that the $(d-1)$-skeleton of ${\Delta_1}^{[d]} \cap {\Delta_2}^{[d]}$ is CM. Another application of the same lemma yields that $\Delta^{[d]} = {\Delta_1}^{[d]} \cup {\Delta_2}^{[d]}$ is CM, which concludes the proof. $\square$

All implications in Theorem 3.33 turn out to be strict; see Proposition 5.13, Proposition 5.14, and Björner [9, §11.10]. We also have the following implications for any field $\mathbb{F}$:

$$\text{homotopically } CM \Longrightarrow CM/\mathbb{Z} \Longrightarrow CM/\mathbb{F}.$$

These implications are valid also for sequentially CM complexes. Again, all implications are strict.

Corollary 3.34. *Let Δ be a pure complex of dimension d. If Δ is VD, shellable, constructible, or homotopically CM, then the homotopical depth of Δ is equal to d.* $\square$

By the following result due to Björner, Wachs, and Welker, sequentially CM complexes have a nice topological structure.

Theorem 3.35 (Björner et al. [15]). *If Δ is sequentially homotopy-CM, then Δ is homotopy equivalent to a wedge of spheres. Moreover, there is no sphere of dimension d in this wedge unless there are maximal faces of Δ of dimension d.* $\square$

Björner and Wachs [13, Th. 4.1] earlier proved Theorem 3.35 in the special case that Δ is semipure shellable.

Let us prove the homology version of Theorem 3.35.

Proposition 3.36 (Wachs [144]). *Assume that Δ is sequentially $CM/\mathbb{Z}$. Then the homology of Δ is torsion-free. Moreover, there is no homology in dimension d unless there are maximal faces of Δ of dimension d. Indeed, the homomorphism $\iota^* : \tilde{H}_d(\Delta \setminus \mathcal{F}(\Delta); \mathbb{Z}) \to \tilde{H}_d(\Delta; \mathbb{Z})$ induced by the inclusion map is zero for all d.*

Proof. We can write ι^* as a composition

$$\tilde{H}_d(\Delta \setminus \mathcal{F}(\Delta); \mathbb{Z}) \to \tilde{H}_d(\Delta^{[d+1]}; \mathbb{Z}) \to \tilde{H}_d(\Delta; \mathbb{Z})$$

of maps induced by inclusion maps; every d-dimensional face of $\Delta \setminus \mathcal{F}(\Delta)$ is contained in a $(d+1)$-dimensional face of Δ. Since $\Delta^{[d+1]}$ is $CM/\mathbb{Z}$, we have that $\tilde{H}_d(\Delta^{[d+1]}; \mathbb{Z}) = 0$ and hence that $\iota^* = 0$ as desired. By the long exact sequence for the pair $(\Delta, \Delta \setminus \mathcal{F}(\Delta))$, it follows that the homology of Δ is torsion-free. $\square$

3.7 Balls and Spheres

We summarize some well-known properties of balls and spheres. Such objects do not play a central part in this book, but they are of some interest in the analysis of the homology of certain complexes. Specifically, in some situations, one may interpret the homology in terms of fundamental cycles of spheres; see Chapters 19 and 20 for the most notable examples.

A simplicial complex Δ is a *d-ball* if there is a homeomorphism $\|\Delta\| \to B^d$. Δ is a *d-sphere* if there is a homeomorphism $\|\Delta\| \to S^d$. For example, the full d-simplex is a d-ball, whereas the boundary of a d-simplex is a $(d-1)$-sphere. We define the *boundary* $\partial\Delta$ of a d-ball Δ as the pure $(d-1)$-dimensional complex with the property that σ is a maximal face of $\partial\Delta$ if and only if σ is contained in exactly one maximal face of Δ.

In general, balls and spheres are not as nice as one may suspect. For example, they are not necessarily homotopically CM; see Björner [9, §11.10]. However, the balls and spheres to be considered in this book are indeed nice.

If a simplicial complex Δ is homeomorphic to a d-dimensional sphere, then $\tilde{H}_d(\Delta; \mathbb{Z})$ is generated by a cycle z that is unique up to sign. We refer to z as the *fundamental cycle* of Δ.

A set $P \subset \mathbb{R}^n$ is a *convex polytope* if P is the convex hull of a finite set P_0 of points in convex position; no point p in P_0 is in the convex hull of $P_0 \setminus \{p\}$. P is homeomorphic to a d-ball for some $d \leq n$; hence P has a well-defined boundary ∂P, which is homeomorphic to a $(d-1)$-sphere. A simplicial complex Δ with 0-cell set Δ_0 is the *boundary complex* of P if there is a bijection $\varphi : \Delta_0 \to P_0$ such that a point $\sum_{i \in \Delta_0} \lambda_i \varphi(i)$ ($\sum_i \lambda_i = 1$, $\lambda_i \geq 0$)

belongs to ∂P if and only if $\{i : \lambda_i > 0\}$ is a face of Δ. We refer to such a complex Δ as *polytopal*. If Δ is the boundary complex of a polytope, then Δ is a shellable sphere; see Bruggesser and Mani [24].

3.8 Stanley-Reisner Rings

We conclude this chapter with a few words about Stanley-Reisner rings. We will only occasionally discuss such rings and include this section merely for completeness.

Let Δ be a simplicial complex on the set Y and let $\mathbb{F}$ be a field. Let R denote the commutative polynomial ring $\mathbb{F}[x_i : i \in Y]$. For each set $\sigma \subseteq Y$, identify σ with the monomial $x_\sigma = \prod_{i \in \sigma} x_i$. Define $I(\Delta)$ to be the monomial ideal in R generated by the (minimal) nonfaces of Δ. This means that a monomial m belongs to $I(\Delta)$ if and only if x_σ divides m for some $\sigma \notin \Delta$. The *Stanley-Reisner ring* or *face ring* of Δ is $R(\Delta) = R/I(\Delta)$.

Let Δ be a $(d-1)$-dimensional simplicial complex with depth $p-1$ over $\mathbb{F}$. Some well-known properties of the Stanley-Reisner ring $R(\Delta)$ are as follows; see a textbook on commutative algebra [26, 43] for ring-theoretic definitions.

- The Krull dimension of $R(\Delta)$ is equal to d.
- The depth of $R(\Delta)$ is equal to p.
- $R(\Delta)$ is a Cohen-Macaulay ring if and only if Δ is $CM/\mathbb{F}$.
- The multiplicity of $R(\Delta)$ is equal to the number of d-dimensional faces of Δ.

We refer the reader to Section 6.2, Reisner [113], and Stanley [132] for details and further references.

Note that this correspondence provides some motivation for examining the topology of simplicial complexes; the problems of counting faces of maximal dimension and examining Cohen-Macaulay properties of a simplicial complex indeed have very natural ring-theoretic counterparts. Moreover, under favorable circumstances, it is possible to combine the theory of Stanley and Reisner with Gröbner basis theory to obtain important information about rings that are not necessarily Stanley-Reisner rings. One example is the work on determinantal ideals by Herzog and Trung [62]; see Section 1.1.6 for more information.