Homology without simplices

Michael Weiss
Mathematisches Institut, Universität Münster
Winter 2013/2014

CHAPTER 1

Homotopy

1.1. The homotopy relation

Let X and Y be topological spaces. (If you are not sufficiently familiar with topological spaces, you should assume that X and Y are metric spaces.) Let f and g be continuous maps from X to Y. Let $[0,1]$ be the unit interval with the standard topology, a subspace of $\mathbb{R}$.

Definition 1.1.1. A *homotopy* from f to g is a continuous map

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

such that $h(x,0) = f(x)$ and $h(x,1) = g(x)$ for all $x \in X$. If such a homotopy exists, we say that f and g are *homotopic*, and write $f \simeq g$. We also sometimes write $h : f \simeq g$ to indicate that h is a homotopy from the map f to the map g.

Remark 1.1.2. If you made the assumption that X and Y are metric spaces, then you should use the product metric on $X \times [0,1]$ and $Y \times [0,1]$, so that for example

$$d((x_1,t_1),(x_2,t_2)) := \max\{d(x_1,x_2), |t_1 - t_2|\}$$

for $x_1, x_2 \in X$ and $t_1, t_2 \in [0,1]$. If you were happy with the assumption that X and Y are "just" topological spaces, then you need to know the definition of *product of two topological spaces* in order to make sense of $X \times [0,1]$ and $Y \times [0,1]$.

Remark 1.1.3. A homotopy $h : X \times [0,1] \to Y$ from $f : X \to Y$ to $g : X \to Y$ can be seen as a "family" of continuous maps

$$h_t : X \to Y\ ;\ h_t(x) = h(x,t)$$

such that $h_0 = f$ and $h_1 = g$. The important thing is that h_t depends continuously on $t \in [0,1]$.

Example 1.1.4. Let $f : \mathbb{R}^n \to \mathbb{R}^n$ be the identity map. Let $g : \mathbb{R}^n \to \mathbb{R}^n$ be the map such that $g(x) = 0 \in \mathbb{R}^n$ for all $x \in \mathbb{R}^n$. Then f and g are homotopic. The map $h : \mathbb{R}^n \times [0,1]$ defined by $h(x,t) = tx$ is a homotopy from f to g.

Example 1.1.5. Let $f : S^1 \to S^1$ be the identity map, so that $f(z) = z$. Let $g : S^1 \to S^1$ be the antipodal map, $g(z) = -z$. Then f and g are homotopic. Using complex number notation, we can define a homotopy by $h(z,t) = e^{\pi i t} z$.

Example 1.1.6. Let $f : S^2 \to S^2$ be the identity map, so that $f(z) = z$. Let $g : S^2 \to S^2$ be the antipodal map, $g(z) = -z$. Then f and g are *not* homotopic. We will prove this later in the course.

Example 1.1.7. Let $f : S^1 \to S^1$ be the identity map, so that $f(z) = z$. Let $g : S^1 \to S^1$ be the constant map with value 1. Then f and g are *not* homotopic. We will prove this quite soon.

PROPOSITION 1.1.8. *"Homotopic" is an equivalence relation on the set of continuous maps from X to Y.*

PROOF. Reflexive: For every continuous map $f : X \to Y$ define the *constant homotopy* $h : X \times [0,1] \to Y$ by $h(x,t) = f(x)$.
Symmetric: Given a homotopy $h : X \times [0,1] \to Y$ from a map $f : X \to Y$ to a map $g : X \to Y$, define the *reverse homotopy* $\bar{h} : X \times [0,1] \to Y$ by $\bar{h}(x,t) = h(x, 1-t)$. Then $\bar{h}$ is a homotopy from g to f.
Transitive: Given continuous maps $e, f, g : X \to Y$, a homotopy h from e to f and a homotopy k from f to g, define the *concatenation homotopy* $k * h$ as follows:

$$(x,t) \mapsto \begin{cases} h(x,2t) & \text{if } 0 \leqslant t \leqslant 1/2 \\ k(x,2t-1) & \text{if } 1/2 \leqslant t \leqslant 1 \,. \end{cases}$$

Then $k * h$ is a homotopy from e to g. □

DEFINITION 1.1.9. The equivalence classes of the above relation "homotopic" are called *homotopy classes*. The homotopy class of a map $f : X \to Y$ is often denoted by $[f]$. The set of homotopy classes of maps from X to Y is often denoted by $[X, Y]$.

PROPOSITION 1.1.10. *Let X, Y and Z be topological spaces. Let $f : X \to Y$ and $g : X \to Y$ and $u : Y \to Z$ and $v : Y \to Z$ be continuous maps. If f is homotopic to g and u is homotopic to v, then $u \circ f : X \to Z$ is homotopic to $v \circ g : X \to Z$.*

PROOF. Let $h : X \times [0,1] \to Y$ be a homotopy from f to g and let $w : Y \times [0,1] \to Z$ be a homotopy from u to v. Then $u \circ h$ is a homotopy from $u \circ f$ to $u \circ g$ and the map $X \times [0,1] \to Z$ given by $(x,t) \mapsto w(g(x),t)$ is a homotopy from $u \circ g$ to $v \circ g$. Because the homotopy relation is transitive, it follows that $u \circ f \simeq v \circ g$. □

DEFINITION 1.1.11. Let X and Y be topological spaces. A (continuous) map $f : X \to Y$ is a *homotopy equivalence* if there exists a map $g : Y \to X$ such that $g \circ f \simeq \mathrm{id}_X$ and $f \circ g \simeq \mathrm{id}_Y$.
We say that X is *homotopy equivalent* to Y if there exists a map $f : X \to Y$ which is a homotopy equivalence.

DEFINITION 1.1.12. If a topological space X is homotopy equivalent to a point, then we say that X is *contractible*. This amounts to saying that the identity map $X \to X$ is homotopic to a constant map from X to X.

EXAMPLE 1.1.13. $\mathbb{R}^m$ is contractible, for any $m \geq 0$.

EXAMPLE 1.1.14. $\mathbb{R}^m \smallsetminus \{0\}$ is homotopy equivalent to S^{m-1}.

EXAMPLE 1.1.15. The general linear group of $\mathbb{R}^m$ is homotopy equivalent to the orthogonal group $O(m)$. The Gram-Schmidt orthonormalisation process leads to an easy proof of that.

1.2. Homotopy classes of maps from the circle to itself

Let $p \colon \mathbb{R} \to S^1$ be the (continuous) map given in complex notation by $p(t) = \exp(2\pi i t)$ and in real notation by $p(t) = (\cos(2\pi t), \sin(2\pi t))$. In the first formula we think of S^1 as a subset of $\mathbb{C}$ and in the second formula we think of S^1 as a subset of $\mathbb{R}^2$.
Note that p is surjective and $p(t+1) = p(t)$ for all $t \in \mathbb{R}$. We are going to use p to understand the homotopy classification of continuous maps from S^1 to S^1. The main lemma is as follows.

LEMMA 1.2.1. *Let* $\gamma\colon [0,1] \to S^1$ *be continuous, and* $a \in \mathbb{R}$ *such that* $p(a) = \gamma(0)$*. Then there exists a unique continuous map* $\tilde{\gamma}\colon [0,1] \to \mathbb{R}$ *such that* $\gamma = p \circ \tilde{\gamma}$ *and* $\tilde{\gamma}(0) = a$*.*

PROOF. The map γ is *uniformly continuous* since $[0,1]$ is compact. It follows that there exists a positive integer n such that $d(\gamma(x), \gamma(y)) < 1/100$ whenever $|x-y| \leq 1/n$. Here d denotes the standard (euclidean) metric on S^1 as a subset of $\mathbb{R}^2$. We choose such an n and write

$$[0,1] = \bigcup_{k=1}^{n} [t_{k-1}, t_k]$$

where $t_k = k/n$. We try to define $\tilde{\gamma}$ on $[0, t_k]$ by induction on k. For the induction beginning we need to define $\tilde{\gamma}$ on $[0, t_1]$ where $t_1 = 1/n$. Let $U \subset S^1$ be the open ball of radius $1/100$ with center $\gamma(0)$. (Note that *open ball* is a metric space concept.) Then $\gamma([0,t_1]) \subset U$. Therefore, in defining $\tilde{\gamma}$ on $[0,t_1]$, we need to ensure that $\tilde{\gamma}([0,t_1])$ is contained in $p^{-1}(U)$. Now $p^{-1}(U) \subset \mathbb{R}$ is a disjoint union of open intervals which are mapped homeomorphically to U under p. One of these, call it V_a , contains a, since $p(a) = \gamma(0) \in U$. The others are translates of the form $\ell + V_a$ where $\ell \in \mathbb{Z}$. Since $[0,t_1]$ is connected, its image under $\tilde{\gamma}$ will also be connected, whatever $\tilde{\gamma}$ is, and so it must be contained entirely in exactly one of the intervals $\ell + V_a$. Since we want $\tilde{\gamma}(0) = a$, we must have $\ell = 0$, that is, image of $\tilde{\gamma}$ contained in V_a. Since the map p restricts to a homeomorphism from V_a to U, we must have $\tilde{\gamma} = q\gamma$ where q is the inverse of the homeomorphism from V_a to U. This formula determines the map $\tilde{\gamma}$ on $[0,t_1]$.
The induction steps are like the induction beginning. In the next step we define $\tilde{\gamma}$ on $[t_1, t_2]$, using a "new" a which is $\tilde{\gamma}(t_1)$ and a "new" U which is the open ball of radius $1/100$ with center $\gamma(t_1)$. □

Now let $g\colon S^1 \to S^1$ be any continuous map. We want to associate with it an integer, the degree of g. Choose $a \in \mathbb{R}$ such that $p(a) = g(1)$. Let $\gamma = g \circ p$ on $[0,1]$; this is a map from $[0,1]$ to S^1. Construct $\tilde{\gamma}$ as in the lemma. We have $p\tilde{\gamma}(1) = \gamma(1) = \gamma(0) = p\tilde{\gamma}(0)$, which implies $\tilde{\gamma}(1) = \tilde{\gamma}(0) + \ell$ for some $\ell \in \mathbb{Z}$.

DEFINITION 1.2.2. This ℓ is the degree of g, denoted $\deg(g)$.

It looks as if this might depend on our choice of a with $p(a) = g(1)$. But if we make another choice then we only replace a by $m + a$ for some $m \in \mathbb{Z}$, and we only replace $\tilde{\gamma}$ by $m + \tilde{\gamma}$. Therefore our calculation of $\deg(g)$ leads to the same result.

Remark. Suppose that $g\colon S^1 \to S^1$ is a continuous map which is close to the constant map $z \mapsto 1 \in S^1$ (complex notation). To be more precise, assume $d(g(z), 1) < 1/1000$ for all $z \in S^1$. Then $\deg(g) = 0$.
The verification is mechanical. Define $\gamma\colon [0,1] \to S^1$ by $\gamma(t) = g(p(t))$. Let $V \subset \mathbb{R}$ be the open interval from $-1/100$ to $1/100$. The map p restricts to a homeomorphism from V to $p(V) \subset S^1$, with inverse $q\colon p(V) \to V$. Put $\tilde{\gamma} = q \circ \gamma$, which makes sense because the image of γ is contained in $p(V)$ by our assumption. Then $p \circ \tilde{\gamma} = \gamma$ as required. Now the image of $\tilde{\gamma}$ is contained in V and therefore

$$|\deg(g)| = |\tilde{\gamma}(1) - \tilde{\gamma}(0)| \leq 2/100$$

and so $\deg(g) = 0$.

Remark. Suppose that $f, g\colon S^1 \to S^1$ are continuous maps. Let $w\colon S^1 \to S^1$ be defined by $w(z) = f(z) \cdot g(z)$ (using the multiplication in $S^1 \subset \mathbb{C}$). Then $\deg(w) = \deg(f) + \deg(g)$.
The verification is also mechanical. Define $\varphi, \gamma, \omega\colon [0,1] \to S^1$ by $\varphi(t) = f(p(t))$, $\gamma(t) =$

$g(p(t))$ and $\omega(t) = w(p(t))$. Construct $\tilde{\varphi}\colon [0,1] \to \mathbb{R}$ and $\tilde{\gamma}\colon [0,1] \to \mathbb{R}$ as in lemma 1.2.1. Put $\tilde{\omega} := \tilde{\varphi} + \tilde{\gamma}$. Then $p \circ \tilde{\omega} = \omega$, so

$$\deg(w) = \tilde{\omega}(1) - \tilde{\omega}(0) = \cdots = \deg(f) + \deg(g).$$

LEMMA 1.2.3. *If $f, g\colon S^1 \to S^1$ are continuous maps which are homotopic, $f \sim g$, then they have the same degree.*

PROOF. Let $h\colon S^1 \times [0,1] \to S^1$ be a homotopy from f to g. As usual let $h_t\colon S^1 \to S^1$ be the map defined by $h_t(z) = h(z,t)$, for fixed $t \in [0,1]$. For fixed $t \in [0,1]$ we can find $\delta > 0$ such that $d(h_t(z), h_s(z)) < 1/1000$ for all $z \in S^1$ and all s which satisfy $|s-t| < \delta$. Therefore $h_s(z) = g_s(z) \cdot h_t(z)$ for such s, where $g_s\colon S^1 \to S^1$ is a map which satisfies $d(g_s(z), 1) < 1/1000$ for all $z \in S^1$. Therefore $\deg(g_s) = 0$ by the remarks above and so $\deg(h_s) = \deg(g_s) + \deg(h_t) = \deg(h_t)$.
We have now shown that the the map $[0,1] \to \mathbb{Z}$ given by $t \mapsto \deg(h_t)$ is locally constant (equivalently, *continuous* as a map of metric spaces) and so it is constant (since $[0,1]$ is connected). In particular $\deg(f) = \deg(h_0) = \deg(h_1) = \deg(g)$. $\square$

LEMMA 1.2.4. *If $f, g\colon S^1 \to S^1$ are continuous maps which have the same degree, then they are homotopic.*

PROOF. Certainly f is homotopic to a map which takes 1 to 1 and g is homotopic to a map which takes 1 to 1 (using complex notation, $1 \in S^1 \subset \mathbb{C}$). Therefore we can assume without loss of generality that $f(1) = 1$ and $g(1) = 1$.
Let $\varphi\colon [0,1] \to S^1$ and $\gamma\colon [0,1] \to S^1$ be defined by $\varphi(t) = f(p(t))$ and $\gamma(t) = g(p(t))$. Construct $\tilde{\varphi}$ and $\tilde{\gamma}$ as in the lemma, using $a = 0$ in both cases, so that $\tilde{\varphi}(0) = 0 = \tilde{\gamma}(0)$. Then

$$\tilde{\varphi}(1) = \deg(f) = \deg(g) = \tilde{\gamma}(1).$$

Note that f can be recovered from $\tilde{\varphi}$ as follows. For $z \in S^1$ choose $t \in [0,1]$ such that $p(t) = z$. Then $f(z) = f(p(t)) = \varphi(t) = p\tilde{\varphi}(t)$. If $z = 1 \in S^1$, we can choose $t = 0$ or $t = 1$, but this ambiguity does not matter since $p\tilde{\varphi}(1) = p\tilde{\varphi}(0)$. Similarly, g can be recovered from $\tilde{\gamma}$. Therefore we can show that f is homotopic to g by showing that $\tilde{\varphi}$ is homotopic to $\tilde{\gamma}$ *with endpoints fixed.* In other words we need a continuous

$$H\colon [0,1] \times [0,1] \to \mathbb{R}$$

where $H(s,0) = \tilde{\varphi}(s)$, $H(s,1) = \tilde{\gamma}(s)$ and $H(0,t) = 0$ for all $t \in [0,1]$ and $H(1,t) = \tilde{\varphi}(1) = \tilde{\gamma}(1)$ for all $t \in [0,1]$. This is easy to do: let $H(s,t) = (1-t)\tilde{\varphi}(s) + t\tilde{\gamma}(s)$. $\square$

Summarizing, we have shown that the degree function gives us a well defined map from $[S^1, S^1]$ to $\mathbb{Z}$, and moreover, that this map is injective. It is not hard to show that this map is also surjective! Namely, for arbitrary $\ell \in \mathbb{Z}$ the map $f\colon S^1 \to S^1$ given by $f(z) = z^\ell$ (complex notation) has $\deg(f) = \ell$. (Verify this.)

COROLLARY 1.2.5. *The degree function is a bijection from $[S^1, S^1]$ to $\mathbb{Z}$.* $\square$

CHAPTER 2

Fiber bundles and fibrations

2.1. Fiber bundles and bundle charts

DEFINITION 2.1.1. Let $p\colon E \to B$ be a continuous map between topological spaces and let $x \in B$. The subspace $p^{-1}(\{x\})$ is sometimes called the *fiber* of p over x.

DEFINITION 2.1.2. Let $p\colon E \to B$ be a continuous map between topological spaces. We say that p is a *fiber bundle* if for every $x \in B$ there exist an open neighborhood U of x in B, a topological space F and a homeomorphism $h\colon p^{-1}(U) \to U \times F$ such that h followed by projection to U agrees with p.

Note that h restricts to a homeomorphism from the fiber of f over x to $\{x\} \times F$. Therefore F must be homeomorphic to the fiber of p over x.

Terminology. Often E is called the *total space* of the fiber bundle and B is called the *base space*. A homeomorphism $h\colon p^{-1}(U) \to U \times F$ as in the definition is called a *bundle chart*. A fiber bundle $p\colon E \to B$ whose fibers are discrete spaces (intuitively, just sets) is also called a *covering space*. (A *discrete space* is a topological space $(X, \mathcal{O})$ in which $\mathcal{O}$ is the entire power set of X.)
Here is an easy way to make a fiber bundle with base space B. Choose a topological space F, put $E = B \times F$ and let $p\colon E \to B$ be the projection to the first factor. Such a fiber bundle is considered unexciting and is therefore called *trivial*. Slightly more generally, a fiber bundle $p\colon E \to B$ is *trivial* if there exist a topological space F and a homeomorphism $h\colon E \to B \times F$ such that h followed by the projection $B \times F \to B$ agrees with p. Equivalently, the bundle is trivial if it admits a bundle chart $h\colon p^{-1}(U) \to U \times F$ where U is all of B.
Two fiber bundles $p_0\colon E_0 \to B$ and $p_1\colon E_1 \to B$ with the same base space B are considered *isomorphic* if there exists a homeomorphism $g\colon E_0 \to E_1$ such that $p_1 \circ g = p_0$. In that case g is an *isomorphism* of fiber bundles.

According to the definition above a fiber bundle is a *map*, but the expression is often used informally for a space rather than a map (the total space of the fiber bundle).

PROPOSITION 2.1.3. *Let $p\colon E \to B$ be a fiber bundle where B is a connected space. Let $x_0, y_0 \in B$. Then the fibers of p over x_0 and y_0, respectively, are homeomorphic.*

PROOF. For every $x \in B$ choose an open neighborhood U_x of x, a space F_x and a bundle chart $h_x\colon p^{-1}(U_x) \to U_x \times F_x$. The open sets U_x for all $x \in B$ form an open cover of B. We make an equivalence relation R on the set B in the following manner: xRy means that there exist elements

$$x_0, x_1, \dots, x_k \in B$$

such that $x_0 = x$, $x_k = y$ and $U_{x_{j-1}} \cap U_{x_j} \neq \emptyset$ for $j = 1, \dots, k$. Clearly xRy implies that F_x is homeomorphic to F_y. Therefore it suffices to show that R has only one equivalence class. Each equivalence class is open, for if $x \in B$ belongs to such an equivalence class, then U_x is contained in the equivalence class. Each equivalence class is closed, since its complement is open, being the union of the other equivalence classes. Since B is connected, this means that there can only be one equivalence class. □

EXAMPLE 2.1.4. One example of a fiber bundle is $p\colon \mathbb{R} \to S^1$, where $p(t) = \exp(2\pi i t)$. We saw this in section 1. To show that it is a fiber bundle, select some $z \in S^1$ and some $t \in \mathbb{R}$ such that $p(t) = z$. Let $V =]t-\delta, t+\delta[$ where δ is a positive real number, not greater than $1/2$. Then p restricts to a homeomorphism from $V \subset \mathbb{R}$ to an open neighborhood $U = p(V)$ of z in S^1; let $q\colon U \to V$ be the inverse homeomorphism. Now $p^{-1}(U)$ is the disjoint union of the translates $\ell + V$, where $\ell \in \mathbb{Z}$. This amounts to saying that

$$g\colon U \times \mathbb{Z} \to p^{-1}(U)$$

given by $(y, m) \mapsto m + q(y)$ is a homeomorphism. The inverse h of g is then a bundle chart. Moreover $\mathbb{Z}$ plays the role of a discrete space. Therefore this fiber bundle is a covering space. It is not a trivial fiber bundle because the total space, $\mathbb{R}$, is not homeomorphic to $S^1 \times \mathbb{Z}$.

EXAMPLE 2.1.5. The Möbius strip leads to another popular example of a fiber bundle. Let $E \subset S^1 \times \mathbb{C}$ consist of all pairs (z, w) where $w^2 = c^2 z$ for some $c \in \mathbb{R}$. This is a (non-compact) implementation of the Möbius strip. There is a projection

$$q\colon E \to S^1$$

given by $q(z, w) = z$. Let us look at the fibers of q. For fixed $z \in S^1$, the fiber of q over z is identified with the space of all $w \in \mathbb{C}$ such that $w^2 = c^2 z$ for some real c. This is equivalent to $w = c\sqrt{z}$ where $\sqrt{z}$ is one of the two roots of z in $\mathbb{C}$. In other words, w belongs to the one-dimensional linear *real* subspace of $\mathbb{C}$ spanned by the two square roots of z. In particular, each fiber of q is homeomorphic to $\mathbb{R}$. The fact that all fibers are homeomorphic to each other should be taken as an indication (though not a proof) that q is a fiber bundle. The full proof is left as an exercise, along with another exercise which is slightly harder: show that this fiber bundle is not trivial.

In preparation for the next example I would like to recall the concept of *one-point compactification.* Let $X = (X, \mathcal{O})$ be a locally compact topological space. (That is to say, X is a Hausdorff space in which every element $x \in X$ has a compact neighborhood.) Let $X^c = (X^c, \mathcal{U})$ be the topological space defined as follows. As a set, X^c is the disjoint union of X and a singleton (set with one element, which in this case we call ∞). The topology $\mathcal{U}$ on X^c is defined as follows. A subset V of X^c belongs to $\mathcal{U}$ if and only if

- either $\infty \notin V$ and $V \in \mathcal{O}$;
- or $\infty \in V$ and $X^c \setminus V$ is a *compact* subset of X.

Then X^c is compact Hausdorff and the inclusion $u\colon X \to X^c$ determines a homeomorphism of X with $u(X) = X^c \setminus \{\infty\}$. The space X^c is called the *one-point compactification* of X. The notation X^c is not standard; instead people often write $X \cup \infty$ and the like. The one-point compactification can be characterized by various good properties; see books on point set topology. For use later on let's note the following, which is clear from the definition of the topology on X^c. Let $Y = (Y, \mathcal{W})$ be any topological space. A map $g\colon Y \to X^c$ is continuous if and only if the following hold:

- $g^{-1}(X)$ is open in Y
- the map from $g^{-1}(X)$ to X obtained by restricting g is continuous

- for every compact subset K of X, the preimage $g^{-1}(K)$ is a closed subset of Y (that is, its complement is an element of $\mathcal{W}$).

EXAMPLE 2.1.6. A famous example of a fiber bundle which is also a crucial example in homotopy theory is the Hopf map from S^3 to S^2, so named after its inventor Heinz Hopf. (Date of invention: around 1930.) Let's begin with the observation that S^2 is homeomorphic to the one-point compactification $\mathbb{C}\cup\infty$ of $\mathbb{C}$. (The standard homeomorphism from S^2 to $\mathbb{C}\cup\infty$ is called *stereographic projection.*) We use this and therefore describe the Hopf map as a map

$$p\colon S^3 \to \mathbb{C}\cup\infty.$$

Also we like to think of S^3 as the unit sphere in $\mathbb{C}^2$. So elements of S^3 are pairs (z,w) where $z,w\in\mathbb{C}$ and $|z|^2+|w|^2=1$. To such a pair we associate

$$p(z,w)=z/w$$

using complex division. This is the Hopf map. Note that in cases where $w=0$, we must have $z\neq 0$ as $|z|^2=|z|^2+|w|^2=1$; therefore z/w can be understood and must be understood as $\infty\in\mathbb{C}\cup\infty$ in such cases. In the remaining cases, $z/w\in\mathbb{C}$.
Again, let us look at the fibers of p before we try anything more ambitious. Let $s\in\mathbb{C}\cup\infty$. If $s=\infty$, the preimage of $\{s\}$ under p consists of all $(z,w)\in S^3$ where $w=0$. This is a circle. If $s\notin\{0,\infty\}$, the preimage of $\{s\}$ under p consists of all $(z,w)\in S^3$ where $w\neq 0$ and $z/w=s$. So this is the intersection of $S^3\subset\mathbb{C}^2$ with the one-dimensional complex linear subspace $\{(z,w)\mid z=sw\}\subset\mathbb{C}^2$. It is also a circle! Therefore all the fibers of p are homeomorphic to the same thing, S^1. We take this as an indication (though not a proof) that p is a fiber bundle.
Now we *show* that p is a fiber bundle. First let $U=\mathbb{C}$, which we view as an open subset of $\mathbb{C}\cup\infty$. Then

$$p^{-1}(U)=\{(z,w)\in S^3\subset\mathbb{C}^2 \mid w\neq 0\}\,.$$

A homeomorphism h from there to $U\times S^1=\mathbb{C}\times S^1$ is given by

$$(z,w)\mapsto (z/w, w/|w|).$$

This has the properties that we require from a bundle chart: the first coordinate of $h(z,w)$ is $z/w=p(z,w)$. (The formula $g(y,z)=(yz,z)/\|(yz,z)\|$ defines a homeomorphism g inverse to h.) Next we try $V=(\mathbb{C}\cup\infty)\smallsetminus\{0\}$, again an open subset of $\mathbb{C}\cup\infty$. We have the following commutative diagram

$$\begin{array}{ccc}
S^3 & \xrightarrow{\alpha} & S^3 \\
\downarrow{\scriptstyle p} & & \downarrow{\scriptstyle p} \\
\mathbb{C}\cup\infty & \xrightarrow{\zeta} & \mathbb{C}\cup\infty
\end{array}$$

where $\alpha(z,w)=(w,z)$ and $\zeta(s)=s^{-1}$. (This amounts to saying that $p\circ\alpha=\zeta\circ p$.) Therefore the composition

$$p^{-1}(V)\xrightarrow{\alpha} p^{-1}(U)\xrightarrow{h} U\times S^1 \xrightarrow{(s,w)\mapsto(s^{-1},w)} V\times S^1$$

has the properties required of a bundle chart. Since $U\cup V$ is all of $\mathbb{C}\cup\infty$, we have produced enough charts to know that p is a fiber bundle. □

2.2. Restricting fiber bundles

Let $p\colon E \to B$ be a fiber bundle. Let A be a subset of B. Put $E_{|A} = p^{-1}(A)$. This is a subset of E. We want to regard A as a subspace of B (with the subspace topology) and $E_{|A}$ as a subspace of E.

PROPOSITION 2.2.1. *The map* $p_A\colon E_{|A} \to A$ *obtained by restricting* p *is also a fiber bundle.*

PROOF. Let $x \in A$. Choose a bundle chart $h\colon p^{-1}(U) \to U \times F$ for p such that $x \in U$. Let $V = U \cap A$, an open neighborhood of x in A. By restricting h we obtain a bundle chart $h_A\colon p^{-1}(V) \to V \times F$ for p_A. □

Remark. In this proof it is important to remember that a bundle chart as above is not just *any* homeomorphism $h\colon p^{-1}(U) \to U \times F$. There is a condition: for every $y \in p^{-1}(U)$ the U-coordinate of $h(y) \in U \times F$ must be equal to $p(y)$. The following informal point of view is recommended: A bundle chart $h\colon p^{-1}(U) \to U \times F$ for p is just a way to specify, simultaneously and continuously, homeomorphisms h_x from the fibers of p over elements $x \in U$ to F. Explicitly, h determines the h_x and the h_x determine h by means of the equation

$$h(y) = (x, h_x(y)) \in U \times F$$

when $y \in p^{-1}(x)$, that is, $x = p(y)$.

Let $p\colon E \to B$ be any fiber bundle. Then B can be covered by open subsets U_i such that $E_{|U_i}$ is a trivial fiber bundle. This is true by definition: choose the U_i together with bundle charts $h_i\colon p^{-1}(U_i) \to U_i \times F_i$. Rename $p^{-1}(U_i) = E_{|U_i}$ if you must. Then each h_i is a bundle isomorphism of $p_{|U_i}\colon E_{|U_i} \to U_i$ with a trivial fiber bundle $U_i \times F_i \to U_i$. There are cases where we can say more. One such case merits a detailed discussion because it takes us back to the concept of homotopy.

LEMMA 2.2.2. *Let* B *be any space and let* $q\colon E \to B \times [0,1]$ *be a fiber bundle. Then* B *admits a covering by open subsets* U_i *such that*

$$q_{|U_i \times [0,1]}\colon E_{|U_i \times [0,1]} \longrightarrow U_i \times [0,1]$$

is a trivial fiber bundle.

PROOF. We fix $x_0 \in B$ for this proof. We try to construct an open neighborhood U of $\{x_0\}$ in B such that $q_{|U\times[0,1]}\colon E_{|U\times[0,1]} \longrightarrow U \times [0,1]$ is a trivial fiber bundle. This is enough.

To minimize bureaucracy let us set it up as a proof by *analytic induction*. So let J be the set of all $t \in [0,1]$ for which there exist an open $U' \subset B$ and an open subset U'' of $[0,1]$ which is also an interval containing 0, such that $x_0 \in U'$ and $t \in U''$ and such that $q_{|U'\times U''}$ is a trivial fiber bundle. The following should be clear.

- J is an open subset of $[0,1]$.
- J is nonempty since $0 \in J$.
- If $t \in J$ then $[0,t] \subset J$; hence J is an interval.

If $1 \in J$, then we are happy. So we assume $1 \notin J$ for a contradiction. Then $J = [0,\sigma[$ for some σ where $0 < \sigma \le 1$. Since q is a fiber bundle, the point (x_0, σ) admits an open neighborhood V in $B \times [0,1]$ with a bundle chart $g\colon q^{-1}(V) \to V \times F_V$. Without loss of generality V has the form $V' \times V''$ where $V' \subset B$ is an open neighborhood of x_0 in B and V'' is an interval which is also an open neighborhood of σ in $[0,1]$. There exists $r < \sigma$ such that $V'' \supset [r,\sigma]$. Then $r \in J$ and so there exists $W = W' \times W''$ open in $B \times [0,1]$

with a bundle chart $h\colon q^{-1}(W) \to U \times F_W$ such that $x_0 \in W'$ and $W'' = [0, \tau[$ where $\tau > r$. Without loss of generality, $W' = V'$. Now $W'' \cup V''$ is an open subset of $[0,1]$ which is an interval (since $r \in W'' \cap V''$). It contains both 0 and σ. Now let $U' = V'$ and $U'' = W'' \cup V''$. If we can show that $q_{|U' \times U''}$ is a trivial fiber bundle, then the proof is complete because $U' \times U''$ contains $\{x_0\} \times [0, \sigma]$, which implies that $\sigma \in J$, which is the contradiction that we need. Indeed we can make a bundle chart

$$k\colon q^{-1}(U' \times U'') \to (U' \times U'') \times F_W$$

as follows. For $(x,t) \in U' \times U''$ with $t \le r$ we take $k_{(x,t)} = h_{(x,t)}$. For $(x,t) \in U' \times U''$ with $t \ge r$ we take

$$k_{(x,t)} = h_{(x,r)} \circ g_{(x,r)}^{-1} \circ g_{(x,t)}\,.$$

Decoding: $h_{(x,t)}$ is a homeomorphism from the fiber of q over $(x,t) \in W \subset B \times [0,1]$ to F_W. Similarly $g_{(x,t)}$ is a homeomorphism from the fiber of q over $(x,t) \in V \subset B \times [0,1]$ to F_V. Also note that

$$h_{(x,r)} \circ g_{(x,r)}^{-1}$$

is a homeomorphism from F_V to F_W, depending on $x \in V_1 = W_1 \subset B$. □

2.3. Pullbacks of fiber bundles

Let $p\colon E \to B$ be a fiber bundle. Let $g\colon X \to B$ be any continuous map of topological spaces.

Definition 2.3.1. The pullback of $p\colon E \to B$ along g is the space

$$g^*E := \{\,(x,y) \in X \times E \mid g(x) = p(y)\}.$$

It is regarded as a subspace of $X \times E$ with the subspace topology.

Lemma 2.3.2. *The projection $g^*E \to X$ given by $(x,y) \mapsto x$ is a fiber bundle.*

Proof. First of all it is helpful to write down the obvious maps that we have in a commutative diagram:

$$\begin{tikzcd} g^*E \arrow[r, "r"] \arrow[d, "q"] & E \arrow[d, "p"] \\ X \arrow[r, "g"] & B \end{tikzcd}$$

Here q and r are the projections given by $(x,y) \mapsto x$ and $(x,y) \mapsto y$. *Commutative* means that the two compositions taking us from g^*E to B agree. Suppose that we have an open set $V \subset B$ and a bundle chart

$$h\colon p^{-1}(V) \xrightarrow{\ \cong\ } V \times F.$$

Now $U := g^{-1}(V)$ is open in X. Also $q^{-1}(U)$ is an open subset of g^*E and we describe elements of that as pairs (x,y) where $x \in U$ and $y \in E$, with $g(x) = p(y)$. We make a homeomorphism

$$q^{-1}(U) \to U \times F$$

by the formula $(x,y) \mapsto (x, h_{g(x)}(y)) = (x, h_{p(y)}(y))$. It is a homeomorphism because the inverse is given by

$$(x,z) \mapsto (x, (h_{g(x)})^{-1}(z))$$

for $x \in U$ and $z \in F$, so that $(g(x), z) \in V \times F$. Its is also clearly a bundle chart. In this way, every bundle chart

$$h\colon p^{-1}(V) \xrightarrow{\ \cong\ } V \times F$$

for $p\colon E \to B$ determines a bundle chart

$$q^{-1}(U) \xrightarrow{\cong} U \times F$$

with the same F, where U is the preimage of V under g. Since $p\colon E \to B$ is a fiber bundle, we have many such bundle charts $p^{-1}(V_j) \to V_j \times F_j$ such that the union of the V_j is all of B. Then the union of the corresponding U_j is all of X, and we have bundle charts $q^{-1}(U_j) \to U_j \times F_j$. This proves that q is a fiber bundle. □

This proof was too long and above all too formal. Reasoning in a less formal way, one should start by noticing that the fiber of q over $z \in X$ is essentially the same (and certainly homeomorphic) to the fiber of p over $g(z) \in B$. Namely,

$$q^{-1}(z) = \{(x,y) \in X \times E \mid g(x) = p(y),\ x = z\} \;=\; \{z\} \times p^{-1}(\{g(z)\})\ .$$

Now recall once again that a bundle chart $h\colon p^{-1}(U) \to U \times F$ for p is just a way to specify, simultaneously and continuously, homeomorphisms h_x from the fibers of p over elements $x \in U$ to F. If we have such a bundle chart for p, then for any $z \in g^{-1}(U)$ we get a homeomorphism from the fiber of q over z, which "is" the fiber of p over $g(z)$, to F. And so, by letting z run through $g^{-1}(U)$, we get a bundle chart for q.

Example 2.3.3. Restriction of fiber bundles is a special case of pullback, up to isomorphism of fiber bundles. More precisely, suppose that $p\colon E \to B$ is a fiber bundle and let $A \subset B$ be a subspace, with inclusion $g\colon A \to B$. Then there is an isomorphism of fiber bundles from $p_A\colon E_{|A} \to A$ to the pullback $g^*E \to A$. This takes $y \in E_{|A}$ to the pair $(p(y), y) \in g^*E \subset A \times E$.

2.4. Homotopy invariance of pullbacks of fiber bundles

Theorem 2.4.1. *Let $p\colon E \to B$ be a fiber bundle. Let $f, g\colon X \to B$ be continuous maps, where X is a compact Hausdorff space. If f is homotopic to g, then the fiber bundles $f^*E \to X$ and $g^*E \to X$ are isomorphic.*

Remark 2.4.2. The compactness assumption on X is unnecessarily strong; *paracompact* is enough. But paracompactness is also a more difficult concept than compactness. Therefore we shall prove the theorem as stated, and leave a discussion of improvements for later.

Remark 2.4.3. Let X be a compact Hausdorff space and let $U_0, U_1, \dots, U_n$ be open subsets of X such that the union of the U_i is all of X. *Then there exist continuous functions*

$$\varphi_0, \varphi_1, \dots, \varphi_n\colon X \to [0,1]$$

such that $\sum_{j=0}^{n} \varphi_j \equiv 1$ *and such that* supp(φ_j)*, the support of* φ_j *, is contained in* U_j *for* $j = 0, 1, \dots, n$. Here supp(φ_j) is the closure in X of the open set

$$\{x \in X \mid \varphi_j(x) > 0\}.$$

A collection of functions $\varphi_0, \varphi_1, \dots, \varphi_n$ with the stated properties is called a *partition of unity subordinate to the open cover of* X *given by* $U_0, \dots, U_n$. For readers who are not aware of this existence statement, here is a reduction (by induction) to something which they might be aware of.

First of all, if X is a compact Hausdorff space, then it is a *normal* space. This means, in addition to the Hausdorff property, that any two disjoint closed subsets of X admit disjoint open neighborhoods. Next, for any normal space X we have the *Tietze-Urysohn extension lemma.* This says that if A_0 and A_1 are disjoint closed subsets of X, then there is a continuous function $\psi\colon X \to [0,1]$ such that $\psi(x) = 1$ for all $x \in A_1$ and $\psi(x) = 0$ for all $x \in A_0$. Now suppose that a normal space X is the union of two open subsets U_0 and U_1. Because X is normal, we can find an open subset $V_0 \subset U_0$ such that the closure of V_0 in X is contained in U_0 and the union of V_0 and U_1 is still X. Repeating this, we can also find an open subset $V_1 \subset U_1$ such that the closure of V_1 in X is contained in U_1 and the union of V_1 and V_0 is still X. Let $A_0 = X \smallsetminus V_0$ and $A_1 = X \smallsetminus V_1$. Then A_0 and A_1 are disjoint closed subsets of X, and so by Tietze-Urysohn there is a continuous function $\psi\colon X \to [0,1]$ such that $\psi(x) = 1$ for all $x \in A_1$ and $\psi(x) = 0$ for all $x \in A_0$. This means that supp(ψ) is contained in the closure of $X \smallsetminus A_0 = V_0$, which is contained in U_0. We take $\varphi_1 = \psi$ and $\varphi_0 = 1 - \psi$. Since $1 - \psi$ is zero on A_1, its support is contained in the closure of V_1, which is contained in U_1. This establishes the induction beginning (case $n = 1$).

For the induction step, suppose that we have an open cover of X given by $U_0, \dots, U_n$ where $n \geq 2$. By inductive assumption we can find a partition of unity subordinate to the cover $U_0 \cup U_1, U_2, \dots, U_n$ and by the induction beginning, another partition of unity subordinate to $U_0, U_1 \cup U_2 \cup \cdots U_n$. Call the functions in the first partition of unity $\varphi_{01}, \varphi_2, \dots, \varphi_n$ and those in the second ψ_0, ψ_1, we see that the functions $\psi_0\varphi_{01}, \psi_1\varphi_{01}, \varphi_2, \dots, \varphi_n$ form a partition of unity subordinate to the cover by $U_0, \dots, U_n$. □

Proof of theorem 2.4.1. Let $h\colon X \times [0,1] \to B$ be a homotopy from f to g, so that $h_0 = f$ and $h_1 = g$. Then $h^*E \to X \times [0,1]$ is a fiber bundle. We give this a new name, say $q\colon L \to X \times [0,1]$. Let ι_0 and ι_1 be the maps from X to $X \times [0,1]$ given by $\iota_0(x) = (x, 0)$ and $\iota_1(x) = (x, 1)$. It is not hard to verify that the fiber bundle $f^*E \to X$ is isomorphic to $\iota_0^*L \to X$ and $g^*E \to X$ is isomorphic to $\iota_1^*L \to X$. Therefore all we need to prove is the following.

Let $q\colon L \to X \times [0,1]$ be a fiber bundle, where X is compact Hausdorff. Then the fiber bundles $\iota_0^*L \to X$ and $\iota_1^*L \to X$ obtained from q by pullback along ι_0 and ι_1 are isomorphic. To make this even more explicit: given the fiber bundle $q\colon L \to X \times [0,1]$, we need to produce a homeomorphism from $L_{|X\times\{0\}}$ to $L_{|X\times\{1\}}$ which fits into a commutative diagram

$$\begin{tikzcd}
L_{|X\times\{0\}} \arrow[r, "\text{our homeom.}"] \arrow[d, "\text{res. of } q"'] & L_{|X\times\{1\}} \arrow[d, "\text{res. of } q"] \\
X \times \{0\} \arrow[r, "(x,0)\mapsto(x,1)"] & X \times \{1\}
\end{tikzcd}$$

Here $L_{|K}$ means $q^{-1}(K)$, for any $K \subset X \times [0,1]$.

By a lemma proved last week (lecture notes week 2), we can find a covering of X by open subsets U_i such that that $q_{U_i\times[0,1]}\colon L_{|U_i\times[0,1]} \to U_i \times [0,1]$ is a trivial bundle, for each i. Since X is compact, finitely many of these U_i suffice, and we can assume that their names are $U_1, \dots, U_n$. Let $\varphi_1, \dots, \varphi_n$ be continuous functions from X to $[0,1]$ making up a partition of unity subordinate to the open covering of X by $U_1, \dots, U_n$. For $j = 0, 1, 2, \dots, n$ let $\nu_j = \sum_{k=1}^{j} \varphi_k$ and let $\Gamma_j \subset X \times [0,1]$ be the graph of ν_j. Note that Γ_0 is $X \times \{0\}$ and Γ_n is $X \times \{1\}$. It suffices therefore to produce a homeomorphism

$e_j\colon L_{|\Gamma_{j-1}} \to L_{|\Gamma_j}$ which fits into a commutative diagram

$$\begin{tikzcd} L_{|\Gamma_{j-1}} \arrow[rr, "e_j"] \arrow[d, "\text{res. of } q"'] & & L_{|\Gamma_j} \arrow[d, "\text{res. of } q"] \\ \Gamma_{j-1} \arrow[rr, "(x,\nu_{j-1}(x)) \mapsto (x,\nu_j(x))"] & & \Gamma_j \end{tikzcd}$$

(for $j = 1, 2, \dots, n$). Since $q_{U_j \times [0,1]}\colon L_{|U_j \times [0,1]} \to U_j \times [0,1]$ is a trivial fiber bundle, we have a single bundle chart for it, a homeomorphism

$$g\colon L_{|U_j \times [0,1]} \longrightarrow (U_i \times [0,1]) \times F$$

with the additional good property that we require of bundle charts. Fix j now and write $L = L' \cup L''$ where L' consists of the $y \in L$ for which $q(y) = (x,t)$ with $x \notin \operatorname{supp}(\varphi_j)$, and L'' consists of the $y \in L$ for which $q(y) = (x,t)$ with $x \in U_j$. Both L' and L'' are open subsets of L. Now we make our homeomorphism $e = e_j$ as follows. By inspection, $L_{|\Gamma_{j-1}} \cap L' = L_{|\Gamma_j} \cap L'$, and we take e to be the identity on $L_{|\Gamma_{j-1}} \cap L'$. By restricting the bundle chart g, we have a homeomorphism $L_{|\Gamma_{j-1}} \cap L'' \to U_j \times F$; more precisely, a homeomorphism from $L_{|\Gamma_{j-1}} \cap L''$ to $(\Gamma_{j-1} \cap U_j \times [0,1]) \times F$. By the same reasoning, we have a homeomorphism $L_{|\Gamma_{j-1}} \cap L'' \to U_j \times F$; more precisely, a homeomorphism from $L_{|\Gamma_{j-1}} \cap L''$ to $(\Gamma_j \cap U_j \times [0,1]) \times F$. Therefore we have a preferred homeomorphism from $L_{|\Gamma_{j-1}} \cap L''$ to $L_{|\Gamma_j} \cap L''$, and we use that as the definition of e on $L_{|\Gamma_{j-1}} \cap L''$. By inspection, the two definitions of e which we have on the overlap $L_{|\Gamma_{j-1}} \cap L' \cap L''$ agree, so e is well defined. □

Corollary 2.4.4. *Let $p\colon E \to B$ be a fiber bundle where B is compact Hausdorff and contractible. Then p is a trivial fiber bundle.*

Proof. By the contractibility assumption, the identity map $f\colon B \to B$ is homotopic to a constant map $g\colon B \to B$. By the theorem, the fiber bundles $f^*E \to B$ and $g^*E \to B$ are isomorphic. But clearly $f^*E \to B$ is isomorphic to the original fiber bundle $p\colon E \to B$. And clearly $g^*E \to B$ is a trivial fiber bundle. □

Corollary 2.4.5. *Let $q\colon E \to B \times [0,1]$ be a fiber bundle, where B is compact Hausdorff. Suppose that the restricted bundle*

$$q_{B \times \{0\}}\colon E_{|B \times \{0\}} \to B \times \{0\}$$

admits a section, i.e., there exists a continuous map $s\colon B \times \{0\} \to E_{|B \times \{0\}}$ such that $q \circ s$ is the identity on $B \times \{0\}$. Then $q\colon E \to B \times [0,1]$ admits a section $\bar{s}\colon B \times [0,1] \to E$ which agrees with s on $B \times \{0\}$.

Proof. Let $f, g\colon B \times [0,1] \to B \times [0,1]$ be defined by $f(x,t) = (x,t)$ and $g(x,t) = (x,0)$. These maps are clearly homotopic. Therefore the fiber bundles $f^*E \to B \times [0,1]$ and $g^*E \to B \times [0,1]$ are isomorphic fiber bundles. Now $f^*E \to B \times [0,1]$ is clearly isomorphic to the original fiber bundle

$$q\colon E \to B \times \{0,1\}$$

and $g^*E \to B \times [0,1]$ is clearly isomorphic to the fiber bundle

$$E_{|B \times \{0\}} \times [0,1] \to B \times [0,1]$$

given by $(y,t) \mapsto (q(y),t)$ for $y \in E_{|B \times \{0\}}$, that is, $y \in E$ with $q(y) = (x,0)$ for some $x \in B$. Therefore we may say that there is a homeomorphism $h\colon E_{|B \times \{0\}} \times [0,1] \to E$

which is over $B \times [0,1]$, in other words, which satisfies

$$(q \circ h)(y,t) = (q(y),t)$$

for all $y \in E_{|B\times\{0\}}$ and $t \in [0,1]$. Without loss of generality, h satisfies the additional condition $h(y,0) = y$ for all $y \in E_{|B\times\{0\}}$. (In any case we have a homeomorphism $u\colon E_{|B\times\{0\}} \to E_{|B\times\{0\}}$ defined by $u(y) = h(y,0)$. If it is not the identity, use the homeomorphism $(y,t) \mapsto h(u^{-1}(y),t)$ instead of $(y,t) \mapsto h(y,t)$.) Now define $\bar{s}$ by $\bar{s}(x,t) = h(s(x),t)$ for $x \in B$ and $t \in [0,1]$. $\square$

2.5. The homotopy lifting property

DEFINITION 2.5.1. A continuous map $p\colon E \to B$ between topological spaces is said to have the *homotopy lifting property* (HLP) if the following holds. Given any space X and continuous maps $f\colon X \to E$ and $h\colon X\times[0,1] \to B$ such that $h(x,0) = p(f(x))$ for all $x \in X$, there exists a continuous map $H\colon X\times[0,1] \to E$ such that $p\circ H = h$ and $H(x,0) = f(x)$ for all $x \in X$. A map with the HLP can be called a *fibration* (sometimes *Hurewicz fibration*).

It is customary to summarize the HLP in a commutative diagram with a dotted arrow:

$$\begin{tikzcd} X \arrow[r, "f"] \arrow[d, "x\mapsto(x,0)"'] & E \arrow[d, "p"] \\ X\times[0,1] \arrow[r, "h"] \arrow[ur, dotted, "H"] & B \end{tikzcd}$$

Indeed, the HLP for the map p means that once we have the data in the outer commutative square, then the dotted arrow labeled H can be found, making both triangles commutative. More associated customs: we think of h as a homotopy between maps h_0 and h_1 from X to B, and we think of $f\colon X \to E$ as a *lift* of the map h_0 , which is just a way of saying that $p \circ f = h_0$.

More generally, or less generally depending on point of view, we say that $p\colon E \to B$ satisfies the HLP for a class of spaces $\mathcal{Q}$ if the dotted arrow in the above diagram can always be supplied when the space X belongs to that class $\mathcal{Q}$.

PROPOSITION 2.5.2. *Let $p\colon E \to B$ be a fiber bundle. Then p has the HLP for compact Hausdorff spaces.*

PROOF. Suppose that we have the data X, f and h as in the above diagram, but we are still trying to construct or find the diagonal arrow H. We are assuming that X is compact Hausdorff. The pullback of p along h is a fiber bundle $h^*E \to X\times[0,1]$. The restricted fiber bundle

$$(h^*E)_{|X\times\{0\}} \to X\times\{0\}$$

has a continuous section s given essentially by f, and if we say it very carefully, by the formula

$$(x,0) \mapsto ((x,0),f(x)) \in h^*E \subset (X\times[0,1]\,)\times E\ .$$

The section s extends to a continuous section $\bar{s}$ of $h^*E \to X\times[0,1]$ by corollary 2.4.5. Now we can define $H := r\circ\bar{s}$, where r is the standard projection from h^*E to E. $\square$

EXAMPLE 2.5.3. Let $p\colon S^3 \to S^2$ be the Hopf fiber bundle. Assume if possible that p is nullhomotopic; we shall try to deduce something absurd from that. So let

$$h\colon S^3 \times [0,1] \to S^2$$

be a nullhomotopy for p. Then $h_0 = p$ and h_1 is a constant map. Applying the HLP in the situation

$$\begin{tikzcd} S^3 \arrow[r, "\mathrm{id}"] \arrow[d, "x \mapsto (x,0)"'] & S^3 \arrow[d, "p"] \\ S^3 \times [0,1] \arrow[r, "h"] \arrow[ur, dotted, "H"] & S^2 \end{tikzcd}$$

we deduce the existence of $H\colon S^3 \times [0,1] \to S^3$, a homotopy from the identity map $H_0 = \mathrm{id}\colon S^3 \to S^3$ to a map $H_1\colon S^3 \to S^3$ with the property that $p \circ H_1$ is constant. Since p itself is certainly not constant, this means that H_1 is not surjective. If H_1 is not surjective, it is nullhomotopic. (A non-surjective map from any space to a sphere is nullhomotopic; that's an exercise.) Consequently $\mathrm{id}\colon S^3 \to S^3$ is also nullhomotopic, being homotopic to H_1. This means that S^3 is contractible.
Is that absurd enough? We shall prove later in the course that S^3 is not contractible. Until then, what we have just shown can safely be stated like this: *if S^3 is not contractible, then the Hopf map $p\colon S^3 \to S^2$ is not nullhomotopic.* (I found this argument in Dugundji's book on topology. Hopf used rather different ideas to show that p is not nullhomotopic.)

Let $p\colon E \to B$ be a fibration (for a class of spaces $\mathcal{Q}$) and let $f\colon X \to B$ be any continous map between topological spaces. We define the pullback f^*E by the usual formula,

$$f^*E = \{(x,y) \in X \times E \mid f(x) = p(y)\ \}.$$

LEMMA 2.5.4. *The projection $f^*E \to X$ is also a fibration for the class of spaces $\mathcal{Q}$.*

The proof is an exercise. □

In example 2.5.3, the HLP was used for something resembling a computation with homotopy classes of maps. Let us try to formalize this, as an attempt to get hold of some algebra in homotopy theory. So let $p\colon E \to B$ be a continuous map which has the HLP for a class of topological spaces $\mathcal{Q}$. Let $f\colon X \to B$ be any continuous map of topological spaces. Now we have a commutative square

$$\begin{tikzcd} f^*E \arrow[r, "q_2"'] \arrow[d, "q_1"] & E \arrow[d, "p"] \\ X \arrow[r, "f"] & B \end{tikzcd}$$

where q_1 and q_2 are the projections. Take any space W in the class $\mathcal{Q}$. There is then a commutative diagram of sets and maps

$$\begin{tikzcd} {[W, f^*E]} \arrow[r] \arrow[d] & {[W, E]} \arrow[d] \\ {[W, X]} \arrow[r] & {[W, B]} \end{tikzcd}$$

PROPOSITION 2.5.5. *The above diagram of sets of homotopy classes is "half exact" in the following sense: given* $a \in [W, X]$ *and* $b \in [W, E]$ *with the same image in* $[W, B]$*, there exists* $c \in [W, f^*E]$ *which is taken to* a *and* b *by the appropriate maps in the diagram.*

PROOF. Represent a by a map $\alpha\colon W \to X$, and b by some map $\beta\colon W \to E$. By assumption, $f \circ \alpha$ is homotopic to $p \circ \beta$. Let $h = (h_t)_{t\in[0,1]}$ be a homotopy, so that $h_0 = p \circ \beta$ and $h_1 = f \circ \alpha$, and $h_t\colon W \to B$ for $t \in [0, 1]$. By the HLP for p, there exists a homotopy $H\colon W \times [0, 1] \to E$ such that $p \circ H = h$ and $H_0 = \beta$. Then H_1 is homotopic to $H_0 = \beta$, and $p \circ H_1 = f \circ \alpha$. Therefore the formula $w \mapsto (\alpha(w), H_1(w))$ defines a map $W \to f^*E$. The homotopy class c of that is the solution to our problem. □

Looking back, we can say that example 2.5.3 is an application of proposition 2.5.5 with $p\colon E \to B$ equal to the Hopf fibration and f equal to the inclusion of a point (and $\mathcal{Q}$ equal to the class of compact Hausdorff spaces, say). We made some unusual choices: $W = E$ and $b = [\mathrm{id}] \in [W, E]$.

2.6. Remarks on paracompactness and fiber bundles

Quoting from many books on point set topology: a topological space $X = (X, \mathcal{O})$ is *paracompact* if it is Hausdorff and every open cover $(U_i)_{i\in\Lambda}$ of X admits a locally finite refinement $(V_j)_{j\in\Psi}$.

There is a fair amount of open cover terminology in that definition. In this formulation, we take the view that an open cover of X is a *family*, i.e., a map from a set to $\mathcal{O}$ (with a special property). This is slightly different from the equally reasonable view that an open cover of X is a subset of $\mathcal{O}$ (with a special property), and it justifies the use of round brackets as in $(U_i)_{i\in\Lambda}$, as opposed to curly brackets. Here the map in question is from Λ to $\mathcal{O}$. There is an understanding that $(V_j)_{j\in\Psi}$ is also an open cover of X, but Ψ need not coincide with Λ. *Refinement* means that for every $j \in \Psi$ there exists $i \in \Lambda$ such that $V_j \subset U_i$. *Locally finite* means that every $x \in X$ admits an open neighborhood W in X such that the set $\{j \in \Psi \mid W \cap V_j \neq \emptyset\}$ is a finite subset of Ψ.

It is wonderfully easy to get confused about the meaning of paracompactness. There is a strong similarity with the concept of compactness, and it is obvious that *compact* (together with Hausdorff) implies paracompact, but it is worth emphasizing the differences. Namely, where compactness has something to do with open covers and *sub*-covers, the definition of paracompactness uses the notion of *refinement* of one open cover by another open cover. We require that every V_j is *contained* in some U_i ; we do not require that every V_j is *equal* to some U_i. And *locally finite* does not just mean that for every $x \in X$ the set $\{j \in \Psi \mid x \in V_j\}$ is a finite subset of Ψ. It means more.

For some people, the Hausdorff condition is not part of *paracompact*, but for me, it is.

An important theorem: every metrizable space is paracompact. This is due to A.H. Stone who, as a Wikipedia page reminds me, is not identical with Marshall Stone of the Stone-Weierstrass theorem and the Stone-Čech compactification. The proof is not very complicated, but you should look it up in a book on point-set topology which is not too ancient, because it was complicated in the A.H. Stone version.

Another theorem which is very important for us: in a paracompact space X, every open cover $(U_i)_{i\in\Lambda}$ admits a subordinate partition of unity. In other words there exist continuous functions $\varphi_i\colon X \to [0, 1]$, for $i \in \Lambda$, such that

- every $x \in X$ admits an open neighborhood W in X for which the set
$$\{\, i \in \Lambda \mid W \cap \operatorname{supp}(\varphi_i) \neq \emptyset \}$$
is finite;
- $\sum_{i\in\Lambda} \varphi_i \equiv 1$;
- $\operatorname{supp}(\varphi_i) \subset U_i$.

The second condition is meaningful if we assume that the first condition holds. (Then, for every $x \in X$, there are only finitely many nonzero summands in $\sum_{i\in\Lambda} \varphi_i(x)$. The first condition also ensures that for any subset $\Xi \subset \Lambda$, the sum $\sum_{i\in\Xi} \varphi_i$ is a continuous function on X.)
The proof of this theorem (existence of subordinate partition of unity for any open cover of a paracompact space) is again not very difficult, and boils down mostly to showing that paracompact spaces are *normal*. Namely, in a normal space, locally finite open covers admit subordinate partitions of unity, and this is easy.

Many of the results about fiber bundles in this chapter rely on partitions of unity, and to ensure their existence, we typically assumed compactness here and there. But now it emerges that paracompactness is enough.
Specifically, in theorem 2.4.1 it is enough to assume that X is paracompact. In corollary 2.4.4 it is enough to assume that B is paracompact (and contractible). In corollary 2.4.5 it is enough to assume that B is paracompact. In proposition 2.5.2 we have the stronger conclusion that p has the HLP for paracompact spaces.

Proof of variant of thm. 2.4.1. Here we assume only that X is paracompact (previously we assumed that it was compact). By analogy with the case of compact X, we can easily reduce to the following statement. *Let* $q\colon L \to X \times [0,1]$ *be a fiber bundle, where* X *is paracompact. Then the fiber bundles* $\iota_0^* L \to X$ *and* $\iota_1^* L \to X$ *obtained from* q *by pullback along* ι_0 *and* ι_1 *are isomorphic.* And to make this more explicit: given the fiber bundle $q\colon L \to X \times [0,1]$, we need to produce a homeomorphism h from $L_{|X\times\{0\}}$ to $L_{|X\times\{1\}}$ which fits into a commutative diagram

$$\begin{tikzcd}
L_{|X\times\{0\}} \arrow[r, "h"] \arrow[d, "\text{res. of } q"'] & L_{|X\times\{1\}} \arrow[d, "\text{res. of } q"] \\
X \times \{0\} \arrow[r, "(x,0)\mapsto(x,1)"] & X \times \{1\}
\end{tikzcd}$$

By a lemma proved in lecture notes week 2, we can find an open cover $(U_i)_{i\in\Lambda}$ of X such that that $q_{U_i\times[0,1]}\colon L_{|U_i\times[0,1]} \to U_i \times [0,1]$ is a trivial bundle, for each $i \in \Lambda$. Let $(\varphi_i)_{i\in\Lambda}$ be a partition of unity subordinate to $(U_i)_{i\in\Lambda}$. So $\varphi_i\colon X \to [0,1]$ is a continuous function with $\operatorname{supp}(\varphi_i) \subset U_i$, and $\sum_i \varphi_i \equiv 1$. Every $x \in X$ admits a neighborhood W in X such that the set
$$\{i \in \Lambda \mid \operatorname{supp}(\varphi_i) \cap W \neq \emptyset\}$$
is finite.
Now choose a total ordering on the set Λ. (A total ordering on Λ is a relation $\leq$ on Λ which is transitive and reflexive, and has the additional property that for any distinct $i, j \in \Lambda$, precisely one of $i \leq j$ or $j \leq i$ holds. We need to assume something here to get such an ordering: for example the Axiom of Choice in set theory is equivalent to the

Well-Ordering Principle, which states that every set can be well-ordered. A well-ordering is also a total ordering.) Given $x \in X$, choose an open neighborhood W of x such that the set of $i \in \Lambda$ having $\mathrm{supp}(\varphi_i) \cap W \neq \emptyset$ is finite; say it has n elements. We list these elements in their order (provided by the total ordering on Λ which we selected):

$$i_1 \leq i_2 \leq i_3 \leq \cdots i_n \ .$$

The functions $\varphi_{i_1}, \varphi_{i_2}, \ldots, \varphi_{i_n}$ (restricted to W) make up a partition of unity on W which is subordinate to the covering by open subsets $W \cap U_{i_1}, W \cap U_{i_2}, \ldots W \cap U_{i_n}$. Now we can proceed exactly as in the proof of theorem 2.4.1 to produce (in n steps) a homeomorphism h_W which makes the following diagram commute:

$$\begin{tikzcd} L_{|W\times\{0\}} \arrow[r, "h_W"] \arrow[d, "\text{res. of } q"'] & L_{|W\times\{1\}} \arrow[d, "\text{res. of } q"] \\ W \times \{0\} \arrow[r, "(x,0)\mapsto(x,1)"] & W \times \{1\} \end{tikzcd}$$

Finally we can regard W or x as variables. If we choose, for every $x \in X$, an open neighborhood W_x with properties like W above, then the W_x for all $x \in X$ constitute an open cover of X. For each W_x we get a homeomorphism h_{W_x} as above. These homeomorphisms agree with each other wherever this is meaningful, and so define together a homeomorphism $h \colon L_{|X\times\{0\}} \to L_{|X\times\{1\}}$ with the property that we require. $\square$

CHAPTER 3

Presheaves and sheaves on topological spaces

3.1. Presheaves and sheaves

DEFINITION 3.1.1. A *presheaf* on a topological space X is a rule $\mathcal{F}$ which to every open subset U of X assigns a set $\mathcal{F}(U)$, and to every pair of nested open sets $U \subset V \subset X$ a map

$$\mathrm{res}_{V,U} \colon \mathcal{F}(V) \to \mathcal{F}(U)$$

which satisfies the following conditions.

- For open sets $U \subset V \subset W$ in X we have $\mathrm{res}_{V,U} \circ \mathrm{res}_{W,V} = \mathrm{res}_{W,U}$ (an equality of maps from $\mathcal{F}(W)$ to $\mathcal{F}(U)$).
- $\mathrm{res}_{V,V} = \mathrm{id} \colon \mathcal{F}(V) \to \mathcal{F}(V)$ for every open V in X.

EXAMPLE 3.1.2. An important and obvious example for us is the following. Fix X as above and let Y be another topological space. For open U in X let $\mathcal{F}(U)$ be the *set* of all continuous maps from U to Y. Note that we make no attempt here to define a topology on $\mathcal{F}(U)$; we just take it as a set. For open sets $U \subset V \subset X$ there is an obvious restriction map $\mathcal{F}(V) \to \mathcal{F}(U)$. That is, a continuous map from V to Y determines by restriction a continuous map from U to Y. The conditions for a presheaf are clearly satisfied.

EXAMPLE 3.1.3. Let $p \colon Y \to X$ be any continuous map. We can use this to make a presheaf $\mathcal{F}$ on X as follows. For an open set U in X, let $\mathcal{F}(U)$ be the set of continuous maps $g \colon U \to Y$ such that $p \circ g = \mathrm{id}_U$. For open sets $U \subset V \subset X$ let $\mathrm{res}_{V,U} \colon \mathcal{F}(V) \to \mathcal{F}(U)$ be given by restriction in the usual sense. Namely, if $f \in \mathcal{F}(V)$, then $f \colon V \to Y$ is a continuous map which satisfies $p \circ f = \mathrm{id}_V$, and so the restriction $f_{|U}$ is a continuous map $U \to Y$ which satisfies $p \circ f_{|U} = \mathrm{id}_U$.

EXAMPLE 3.1.4. Suppose that X happens to be a differentiable (smooth) manifold (in which case it is also a topological space). For open U in X, let $\mathcal{F}(U)$ be the set of smooth functions from U to $\mathbb{R}$. For open subsets $U \subset V \subset X$, let $\mathrm{res}_{V,U} \colon \mathcal{F}(V) \to \mathcal{F}(U)$ be given by restriction in the usual sense. The conditions for a presheaf are clearly satisfied by $\mathcal{F}$.

EXAMPLE 3.1.5. Given a topological space X and a set S, define $\mathcal{F}(U) = S$ for every open U in X. For open sets $U \subset V \subset X$, let $\mathrm{res}_{V,U} \colon \mathcal{F}(V) \to \mathcal{F}(U)$ be the identity map of S. The conditions for a presheaf are clearly satisfied.

EXAMPLE 3.1.6. Fix X as above and let Y be another topological space. For open U in X put $\mathcal{F}(U) = [U, Y]$, the set of homotopy classes of continuous maps from U to Y. For open sets $U \subset V \subset X$ there is an obvious restriction map $\mathcal{F}(V) \to \mathcal{F}(U)$. That is, a homotopy class of continuous maps from V to Y determines by restriction a homotopy class of continuous maps from U to Y. The conditions for a presheaf are clearly satisfied. This example looks as if it might become very important in this course, since it connects presheaves and the concept of homotopy. But it will not become very important except as a source of homework problems and counterexamples.

EXAMPLE 3.1.7. Fix X as above and let Y be another topological space. For an open subset U of X let $\mathcal{F}(U)$ be the set of *formal linear combinations* (with integer coefficients) of continuous maps from U to Y. So an element of $\mathcal{F}(U)$ might look like $5f - 3g + 9h$ where f, g and h are continuous maps from U to Y. We do not insist that f, g, h in this expression are distinct, but if for example f and g are equal, then we take the view that $5f - 3g + 9h$ and $2f + 9h$ define the same element of $\mathcal{F}(U)$. This remark is important when we define the restriction map

$$\mathrm{res}_{V,U}\colon \mathcal{F}(V) \to \mathcal{F}(U)$$

This is of course determined by restriction of continuous maps. So for example, if

$$3a - 6b + 10c - d$$

is an element of $\mathcal{F}(V)$, and here we may as well assume that the continuous maps $a, b, c, d\colon V \to Y$ are distinct (because we can simplify the expression if not), then $\mathrm{res}_{V,U}$ takes that element to $3(a_{|U}) - 6(b_{|U}) + 10(c_{|U}) - d_{|U} \in \mathcal{F}(U)$. And here we can not assume that the continuous maps $a_{|U}, b_{|U}, c_{|U}, d_{|U}\colon U \to Y$ are all distinct. In any case the conditions for a presheaf are clearly satisfied.

This example looks silly and unimportant, but it is not silly and it will become very important in this course. Let's also note that there are more grown-up ways to describe $\mathcal{F}(U)$ for this presheaf $\mathcal{F}$. Instead of saying *the set of formal linear combinations with integer coefficients of continuous maps from U to Y*, we can say: the free abelian group generated by the set of continuous maps from U to Y. Or we can say: the free $\mathbb{Z}$-module generated by the set of continuous maps from U to Y. (See also subsection 4.4 for some clarifications.)

With a view to the next definition we introduce some practical notation. Let X be a space, let $\mathcal{F}$ be a presheaf on X, and suppose that U, V are open subsets of X such that $U \subset V$. Then we have the restriction map $\mathrm{res}_{V,U}\colon \mathcal{F}(V) \to \mathcal{F}(U)$. Let $s \in \mathcal{F}(V)$. Instead of writing $\mathrm{res}_{V,U}(s) \in \mathcal{F}(U)$, we sometimes write $s_{|U} \in \mathcal{F}(U)$.

DEFINITION 3.1.8. A presheaf $\mathcal{F}$ on a topological space X is called a *sheaf* on X if it has the following additional properties. For every collection of open subsets $(W_i)_{i\in\Lambda}$ of X, and every collection

$$(s_i \in \mathcal{F}(W_i))_{i\in\Lambda}$$

with the property $s_{i|W_i\cap W_j} = s_{j|W_i\cap W_j} \in \mathcal{F}(W_i \cap W_j)$, there exists a unique

$$s \in \mathcal{F}(\bigcup_{i\in\Lambda} W_i)$$

such that $s_{|W_i} = s_i$ for all $i \in \Lambda$. In particular, $\mathcal{F}(\emptyset)$ has exactly one element.

In a slightly more wordy formulation: if we have elements $s_i \in \mathcal{F}(W_i)$ for all $i \in \Lambda$, and we have agreement of s_i and s_j on $W_i \cap W_j$ for all $i, j \in \Lambda$, then there is a unique $s \in \mathcal{F}(\bigcup_i W_i)$ which agrees with s_i on each W_i.

To silence a particularly nagging and persistent type of critic, including the critic within myself, let me explain in detail why this implies that $\mathcal{F}(\emptyset)$ has exactly one element. Put $\Lambda = \emptyset$. For each $i \in \Lambda$, select an open subset W_i. (Easy, because there is no $i \in \Lambda$.) For each $i \in \Lambda$, select an element $s_i \in \mathcal{F}(W_i)$. (Easy.) Verify that, for each i and j in Λ, we have $s_{i|U_i\cap U_j} = s_{j|U_i\cap U_j}$. (Easy.) Conclude that there exists a *unique*

$$s \in \mathcal{F}(\bigcup_{i\in\Lambda} W_i)$$

such that $s_{|W_i} = s_i$ for every $i \in \Lambda$. Now note that $\bigcup_{i\in\Lambda} W_i = \emptyset$ and verify that *every* $t \in \mathcal{F}(\emptyset)$ satisfies the condition $t_{|W_i} = s_i$ for every $i \in \Lambda$. (Easy.) Therefore *every* element t of $\mathcal{F}(\emptyset)$ must be equal to that distinguished element s which we have already spotted.

Obviously it is now our duty to scan the list of the examples above and decide for each of these presheaves $\mathcal{F}$ whether it is a sheaf. It is a good idea to ask first in each case whether $\mathcal{F}(\emptyset)$ has exactly one element. If that is not the case, then it is not a sheaf. It looks like a mean reason to refuse sheaf status to a presheaf. But often when $\mathcal{F}(\emptyset)$ does not have exactly one element, the presheaf $\mathcal{F}$ turns out to have other properties which prevent us from promoting it to sheaf status by simply redefining $\mathcal{F}(\emptyset)$. — The following lemma is also a good tool in testing for the sheaf property.

LEMMA 3.1.9. *Let $\mathcal{F}$ be a sheaf on X and let $(W_i)_{i\in\Lambda}$ be a collection of pairwise disjoint open subsets of X. Then the formula $s \mapsto (s_{|W_i})_{i\in\Lambda}$ determines a bijection*

$$\mathcal{F}(\bigcup_{i\in\Lambda} W_i) \longrightarrow \prod_{i\in\Lambda} \mathcal{F}(W_i).$$

PROOF. Take an element in $\prod_{i\in\Lambda} \mathcal{F}(W_i)$ and denote it by $(s_i)_{i\in\Lambda}$, so that s_i is an element of $\mathcal{F}(W_i)$. Since $W_i \cap W_j = \emptyset$ and $\mathcal{F}(\emptyset)$ has exactly one element, the matching condition

$$s_{i|W_i\cap W_j} = s_{j|W_i\cap W_j}$$

is vacuously satisfied for all $i, j \in \Lambda$. Hence by the sheaf property, there is a unique element s in $\mathcal{F}(\bigcup_{i\in\Lambda} W_i)$ such that $s_{|W_i} = s_i$ for all $i \in \Lambda$. This means precisely that $s \mapsto (s_{|W_i})_{i\in\Lambda}$ is a bijection. (The surjectivity is expressed in *there is* and the injectivity in the word *unique*.) □

Discussion of example 3.1.2. This is a sheaf. What is being said here is that if we have open $W_i \subset X$ for each $i \in \Lambda$, and continuous maps $f_i \colon W_i \to Y$ for each i such that f_i and f_j agree on $W_i \cap W_j$ for all $i, j \in \Lambda$, then we have a unique continuous map f from $\bigcup W_i$ to Y which agrees with f_i on W_i for each $i \in \Lambda$.

Discussion of example 3.1.3. This is a sheaf. We can reason as in the case of example 3.1.2.

Discussion of example 3.1.4. This is a sheaf. What is being said here is that if X is a smooth manifold, and we have open $W_i \subset X$ for each $i \in \Lambda$, and smooth functions $f_i \colon W_i \to \mathbb{R}$ for each i such that f_i and f_j agree on $W_i \cap W_j$ for all $i, j \in \Lambda$, then we have a unique smooth $f \colon \bigcup W_i \to Y$ which agrees with f_i on W_i for each $i \in \Lambda$. An interesting aspect of this example is that, in contrast to examples 3.1.2 and 3.1.3, it seems to express something which is not part of the world of topological spaces, something "differentiable". So I am suggesting that the notion of *smooth manifold* could be redefined along the following lines: a smooth manifold is a topological Hausdorff space X together with a sheaf $\mathcal{F}$ ... which we would call the sheaf of *smooth* functions (on open subsets of X) and which would presumably have to be a subsheaf (notion yet to be defined) of the sheaf of *continuous* functions on open subsets of X. That would be an alternative to defining smooth manifolds using charts and atlases. Of course this has been noticed and has been done by the ancients, but I am digressing.

Discussion of example 3.1.5. Here we have to make a case distinction. If S has exactly one element, then this presheaf $\mathcal{F}$ is a sheaf, and the verification is easy. If S has more than one element, or is empty, then $\mathcal{F}$ is not a sheaf because $\mathcal{F}(\emptyset)$ does not have exactly one element.

Can we fix this by redefining $\mathcal{F}(\emptyset)$ to have exactly one element? Let us try. So let $\mathcal{G}$ be the presheaf on X defined by $\mathcal{G}(U) = S$ when U is nonempty, and $\mathcal{G}(\emptyset) = \{*\}$, a set with a single element $*$. It is a presheaf as follows: for open subsets $U \subset V$ of X we let $\mathrm{res}_{V,U}\colon \mathcal{G}(V) \to \mathcal{G}(U)$ be the identity map of S if $U \neq \emptyset$; otherwise it is the unique map of sets from $\mathcal{G}(V)$ to $\{*\}$.
Is this presheaf $\mathcal{G}$ a sheaf? The answer depends a little on X, and on S. Suppose that X has disjoint open nonempty subsets U_1 and U_2. By lemma 3.1.9, the diagonal map from $S = \mathcal{G}(U_1 \cup U_2)$ to $S \times S = \mathcal{G}(U_1) \times \mathcal{G}(U_2)$ is bijective. We have a problem with that if S has more than one element. The case where S has exactly one element was excluded, so only the possibility $S = \emptyset$ remains. And indeed, if S is empty, we don't have a problem: $\mathcal{G}$ is a sheaf. Also, if X does not have any disjoint nonempty open subsets U_1 and U_2, we don't have a problem: $\mathcal{G}$ is a sheaf, no matter what S is.

Discussion of example 3.1.6. In general, this is not a sheaf, although it responds nicely to the two standard tests. (One standard test is to ask: what is $\mathcal{F}(\emptyset)$? Here we get the set of homotopy classes of maps from $\emptyset$ to Y, and that set has exactly one element, as it should have if $\mathcal{F}$ were a sheaf. The other standard test comes from lemma 3.1.9. If $(W_i)_{i\in\Lambda}$ is a collection of disjoint open subsets of X, then

$$\mathcal{F}(\textstyle\bigcup_i W_i) = [\bigcup_i W_i, Y]$$

which is in bijection with $\prod_{i\in\Lambda}[W_i, Y]$ by composition with the inclusions $W_j \to \bigcup_{i\in\Lambda} W_i$ for each $j \in \Lambda$.) For a counterexample, let $X = Y = S^1$. In X we have the open sets U_1 and U_2 where $U_1 = S^1 - \{1\}$ and $U_2 = S^1 \smallsetminus \{-1\}$, using complex number notation. Since U_1 and U_2 are contractible and Y is path connected, both $\mathcal{F}(U_1)$ and $\mathcal{F}(U_2)$ have exactly one element. Since $U_1 \cap U_2$ is the disjoint union of two contractible open sets V_1 and V_2, we get

$$\mathcal{F}(U_1 \cap U_2) = \mathcal{F}(V_1 \cup V_2)$$

which is in bijection with $\mathcal{F}(V_1) \times \mathcal{F}(V_2)$, which again has exactly one element. If $\mathcal{F}$ were a sheaf, it would follow from these little calculations that $\mathcal{F}(U_1 \cup U_2)$ has exactly one element. But $\mathcal{F}(U_1 \cup U_2) = \mathcal{F}(X) = [X, Y] = [S^1, S^1]$, and we know that this has infinitely many elements.

Discussion of example 3.1.7. This is obviously not a sheaf because $\mathcal{F}(\emptyset)$ has more than one element. Indeed, there is exactly one continuous map from $\emptyset$ to Y. So $\mathcal{F}(\emptyset)$ is the free $\mathbb{Z}$-module one one generator, which means that it is isomorphic to $\mathbb{Z}$.
It might seem pointless to look for further reasons to deny sheaf status to $\mathcal{F}$. It is like kicking somebody who is already down. Nevertheless, because this is an important example, it will be instructive for us to know more about it, and we could argue that by showing interest we are showing some patience and kindness. Also, there is a new aspect here: the sets $\mathcal{F}(U)$ always always carry the structure of abelian groups alias $\mathbb{Z}$-modules, and the maps $\mathrm{res}_{V,U}$ are always homomorphisms.
Suppose that $X = \{1, 2, 3, 4, 5, 6\}$ with the discrete topology (every subset of X is declared to be open). Let $Y = \{a, b\}$, a set with two elements, also with the discrete topology. We note that X is the disjoint union of six open subsets U_i , where $i = 1, 2, 3, 4, 5, 6$ and $U_i = \{i\}$. We have $\mathcal{F}(U_i) = \mathbb{Z} \oplus \mathbb{Z} = \mathbb{Z}^2$ (free $\mathbb{Z}$-module on two generators) because each U_i has exactly two continuous maps to Y. We have $\mathcal{F}(\bigcup_i U_i) = \mathcal{F}(X) = \mathbb{Z}^{64}$ (free $\mathbb{Z}$-module on 64 generators) because there are 64 continuous maps from X to Y. It follows that the map

$$\mathcal{F}(\textstyle\bigcup_i U_i) \longrightarrow \prod_{i=1}^{6} \mathcal{F}(U_i)$$

of lemma 3.1.9 (which in the present circumstances is a $\mathbb{Z}$-module homomorphism) cannot be bijective, because that would make it a $\mathbb{Z}$-module isomorphism between $\mathbb{Z}^{64}$ and $\mathbb{Z}^{12}$. (For an abstract interpretation of what is happening, the notion of *tensor product* is useful. Namely, $\mathcal{F}(\bigcup_i U_i) \cong \mathbb{Z}^{64}$ is isomorphic to the tensor product

$$\mathcal{F}(U_1) \otimes \mathcal{F}(U_2) \otimes \cdots \otimes \mathcal{F}(U_6).$$

It is unsurprising that this is *not* isomorphic to the product $\prod_{i=1}^{6} \mathcal{F}(U_i)$. So it emerges that $\mathcal{F}$ fails to have the sheaf property because it has another respectable property.)
Next, re-define X and Y in such a way that X and Y are two topological spaces related by a covering map $p\colon Y \to X$ with finite fibers. In other words, p is a fiber bundle whose fibers are finite sets (viewed as topological spaces with the discrete topology). For simplicity, suppose also that X is connected. Choose an open covering $(W_j)_{j\in\Lambda}$ of X such that p admits a bundle chart over W_j for each j:

$$h_j\colon p^{-1}(W_j) \to W_j \times F$$

where F is a finite set (with the discrete topology). There is no loss of generality in asking for the same F in all cases, independent of j, because X is connected.[1] For $j \in \Lambda$ and $z \in F$ there is a continuous map $\sigma_{j,z}\colon W_j \to Y$ given by $\sigma_{j,z}(x) = h_j^{-1}(x, z)$ for $x \in W_j$. Define

$$s_j = \sum_{z\in F} \sigma_{j,z}\,.$$

This is a formal linear combination of continuous maps from W_j to Y which has meaning as an element $\mathcal{F}(W_j)$. So we can write $s_j \in \mathcal{F}(W_j)$. The matching condition

$$s_{i|W_i\cap W_j} = s_{j|W_i\cap W_j}$$

is satisfied. However it seems to be hard or impossible to produce $s \in \mathcal{F}(X) = \mathcal{F}(\bigcup_j W_j)$ such that $s_{|W_i} = s_i$ for all $i \in \Lambda$. This indicates another violation of the sheaf property. (Unfortunately, showing that in many cases such an s does not exist is also hard; we may return to this when we are wiser.)

3.2. Categories, functors and natural transformations

The concept of a *category* and the related notions *functor* and *natural transformation* emerged in the middle of the 20th century (Eilenberg-MacLane, 1945) and were immediately used to re-organize algebraic topology (Eilenberg-Steenrod, 1952). Later these notions became very important in many other branches of mathematics, especially algebraic geometry. Category theory has many definitions of great depth, I think, but in the foundations very few theorems and fewer proofs of any depth. Among those who love difficult proofs, it has a reputation of shallowness, boring-ness; for many of the theorizers who appreciate good definitions, it is an ever-ongoing revelation. Young mathematicians tend to like it better than old mathematicians ... probably because it helps them to see some order in a multitude of mathematical facts.

DEFINITION 3.2.1. A *category* $\mathcal{C}$ consists of a class $\mathrm{Ob}(\mathcal{C})$ whose elements are called the *objects of* $\mathcal{C}$ and the following additional data.

- For any two objects c and d of $\mathcal{C}$, a set $\mathrm{mor}_{\mathcal{C}}(c, d)$ whose elements are called the *morphisms* from c to d.

[1] Lecture notes week 2, prop. 2.3.

- For any object c in $\mathcal{C}$, a distinguished element $\mathrm{id}_c \in \mathrm{mor}_{\mathcal{C}}(c,c)$, called the *identity morphism* of c.
- For any three objects b,c,d of $\mathcal{C}$, a map from $\mathrm{mor}_{\mathcal{C}}(c,d) \times \mathrm{mor}_{\mathcal{C}}(b,c)$ to $\mathrm{mor}_{\mathcal{C}}(b,d)$ called *composition* and denoted by $(f,g) \mapsto f \circ g$.

These data are subject to certain conditions, namely:

- Composition of morphisms is associative.
- The identity morphisms act as two-sided neutral elements for the composition.

The associativity condition, written out in detail, means that

$$(f \circ g) \circ h = f \circ (g \circ h)$$

whenever a,b,c,d are objects of $\mathcal{C}$ and $f \in \mathrm{mor}_{\mathcal{C}}(c,d)$, $g \in \mathrm{mor}_{\mathcal{C}}(b,c)$, $h \in \mathrm{mor}_{\mathcal{C}}(a,b)$. The condition on identity morphisms means that $f \circ \mathrm{id}_c = f = \mathrm{id}_d \circ f$ whenever c and d are objects in $\mathcal{C}$ and $f \in \mathrm{mor}_{\mathcal{C}}(c,d)$. Saying that $\mathrm{Ob}(\mathcal{C})$ is a *class*, rather than a *set*, is a subterfuge to avoid problems which are likely to arise if, for example, we talk about *the set of all sets* (Russell's paradox). If the object class is a set, which sometimes happens, we speak of a *small category*.

Notation: we shall often write $\mathrm{mor}(c,d)$ instead of $\mathrm{mor}_{\mathcal{C}}(c,d)$ if it is obvious that the category in question is $\mathcal{C}$. Morphisms are often denoted by arrows, as in $f\colon c \to d$ when $f \in \mathrm{mor}(c,d)$. It is customary to say in such a case that c is the *source* or *domain* of f, and d is the *target* or *codomain* of f.

A morphism $f\colon c \to d$ in a category $\mathcal{C}$ is said to be an *isomorphism* if there exists a morphism $g\colon d \to c$ in $\mathcal{C}$ such that $g \circ f = \mathrm{id}_c \in \mathrm{mor}_{\mathcal{C}}(c,c)$ and $f \circ g = \mathrm{id}_d \in \mathrm{mor}_{\mathcal{C}}(d,d)$.

Example 3.2.2. The prototype is $\mathcal{S}\mathrm{ets}$, the category of sets. The objects of that are the sets. For two sets S and T, the set of morphisms $\mathrm{mor}(S,T)$ is the set of all maps from S to T. Composition is composition of maps as we know it and the identity morphisms are the identity maps as we know them.

Another very important example for us is $\mathcal{T}\mathrm{op}$, the category of topological spaces. The objects are the topological spaces. For topological spaces $X = (X, \mathcal{O}_X)$ and $Y = (Y, \mathcal{O}_Y)$, the set of morphisms $\mathrm{mor}(X,Y)$ is the set of continuous maps from X to Y. Composition is composition of continuous maps as we know it and the identity morphisms are the identity maps as we know them.

Another very important example for us is $\mathcal{H}\mathrm{o}\mathcal{T}\mathrm{op}$, the homotopy category of topological spaces. The objects are the topological spaces, as in $\mathcal{T}\mathrm{op}$. But the set of morphisms from $X = (X, \mathcal{O}_X)$ to $Y = (Y, \mathcal{O}_Y)$ is $[X,Y]$, the set of *homotopy classes* of continuous maps from X to Y. Composition $\circ$ is defined by the formula

$$[f] \circ [g] = [f \circ g]$$

for $[f] \in [Y,Z]$ and $[g] \in [X,Y]$. Here $f\colon Y \to Z$ and $g\colon X \to Y$ are continuous maps representing certain elements of $[Y,Z]$ and $[X,Y]$, and $f \circ g\colon X \to Z$ is their composition. There is an issue of well-defined-ness here, but fortunately we settled this long ago in chapter 1. As a result, associativity of composition is not in doubt because it is a consequence of associativity of composition in $\mathcal{T}\mathrm{op}$. The identity morphisms in $\mathcal{H}\mathrm{o}\mathcal{T}\mathrm{op}$ are the homotopy classes of the identity maps.

Another popular example is $\mathcal{G}\mathrm{roups}$, the category of groups. The objects are the groups. For groups G and H, the set of morphisms $\mathrm{mor}(G,H)$ is the set of group homomorphisms from G to H. Composition of morphisms is composition of group homomorphisms.

The definition of a category as above permits some examples which are rather strange.

One type of strange example which will be important for us soon is as follows. Let $(P, \le)$ be a partially ordered set, alias poset. That is to say, P is a set and $\le$ is a relation on P which is transitive ($x \le y$ and $y \le z$ forces $x \le z$), reflexive ($x \le x$ holds for all x) and antisymmetric (in the sense that $x \le y$ and $y \le x$ together implies $x = y$). We turn this setup into a small category (nameless) such that the object set is P. We decree that, for $x, y \in P$, the set $\mathrm{mor}(x, y)$ shall be empty if x is not $\le y$, and shall contain exactly one element, denoted $*$, if $x \le y$. Composition

$$\circ : \mathrm{mor}(y, z) \times \mathrm{mor}(x, y) \longrightarrow \mathrm{mor}(x, z)$$

is defined as follows. If y is not $\le z$, then $\mathrm{mor}(y, z)$ is empty and so $\mathrm{mor}(y, z) \times \mathrm{mor}(x, y)$ is empty, too. There is exactly one map from the empty set to $\mathrm{mor}(x, z)$ and we take that. If x is not $\le y$, then $\mathrm{mor}(y, z) \times \mathrm{mor}(x, y)$ is empty, and we have only one choice for our composition map, and we take that. The remaining case is the one where $x \le y$ and $y \le z$. Then $x \le z$ by transitivity. Therefore $\mathrm{mor}(y, z) \times \mathrm{mor}(x, y)$ has exactly one element, but more importantly, $\mathrm{mor}(x, z)$ has also exactly one element. Therefore there is exactly one map from $\mathrm{mor}(y, z) \times \mathrm{mor}(x, y)$ to $\mathrm{mor}(x, z)$ and we take that.
Another type of strange example (less important for us but still instructive) can be constructed by starting with a specific group G, with multiplication map $\mu \colon G \times G \to G$. From that we construct a small category (nameless) whose object set has exactly one element, denoted $*$. We let $\mathrm{mor}(*, *) = G$. The composition map

$$\mathrm{mor}(*, *) \times \mathrm{mor}(*, *) \to \mathrm{mor}(*, *)$$

now has to be a map from $G \times G$ to G, and for that we choose μ, the multiplication of G. Since μ has an associativity property, composition of morphisms is associative. For the identity morphism $\mathrm{id}_* \in \mathrm{mor}(*, *)$ we take the neutral element of G.
There are also some easy ways to make new categories out of old ones. One important example: let $\mathcal{C}$ be any category. We make a new category $\mathcal{C}^{\mathrm{op}}$, the *opposite* category of $\mathcal{C}$. It has the same objects as $\mathcal{C}$, but we let

$$\mathrm{mor}_{\mathcal{C}^{\mathrm{op}}}(c, d) := \mathrm{mor}_{\mathcal{C}}(d, c)$$

when c and d are objects of $\mathcal{C}$, or equivalently, objects of $\mathcal{C}^{\mathrm{op}}$. The identity morphism of an object c in $\mathcal{C}^{\mathrm{op}}$ is the identity morphism of c in $\mathcal{C}$. Composition

$$\mathrm{mor}_{\mathcal{C}^{\mathrm{op}}}(c, d) \times \mathrm{mor}_{\mathcal{C}^{\mathrm{op}}}(b, c) \longrightarrow \mathrm{mor}_{\mathcal{C}^{\mathrm{op}}}(b, d)$$

is defined by noting $\mathrm{mor}_{\mathcal{C}^{\mathrm{op}}}(c, d) \times \mathrm{mor}_{\mathcal{C}^{\mathrm{op}}}(b, c) = \mathrm{mor}_{\mathcal{C}}(d, c) \times \mathrm{mor}_{\mathcal{C}}(c, b)$ and going from there to $\mathrm{mor}_{\mathcal{C}}(c, b) \times \mathrm{mor}_{\mathcal{C}}(d, c)$ by an obvious bijection, and from there to $\mathrm{mor}_{\mathcal{C}}(d, b) = \mathrm{mor}_{\mathcal{C}^{\mathrm{op}}}(b, d)$ using composition of morphisms in the category $\mathcal{C}$.

It turns out that there is something like a *category of all categories.* Let us not try to make that very precise because there are some small difficulties and complications in that. In any case there is a concept of morphism between categories, and the name of that is *functor.*

DEFINITION 3.2.3. A *functor* from a category $\mathcal{C}$ to a category $\mathcal{D}$ is a rule F which to every object c of $\mathcal{C}$ assigns an object $F(c)$ of $\mathcal{D}$, and to every morphism $g \colon b \to c$ in $\mathcal{C}$ a morphism $F(g) \colon F(b) \to F(c)$ in $\mathcal{D}$, subject to the following conditions.

- For any object c in $\mathcal{C}$ with identity morphism id_c, we have $F(\mathrm{id}_c) = \mathrm{id}_{F(c)}$.
- Whenever a, b, c are objects in $\mathcal{C}$ and $h \in \mathrm{mor}_{\mathcal{C}}(a, b)$, $g \in \mathrm{mor}_{\mathcal{C}}(b, c)$, we have $F(g \circ h) = F(g) \circ F(h)$ in $\mathrm{mor}_{\mathcal{D}}(F(a), F(c))$.

EXAMPLE 3.2.4. A functor F from the category $\mathcal{T}op$ to the category $\mathcal{S}ets$ can be defined as follows. For a topological space X let $F(X)$ be the set of path components of X. A continuous map $g\colon X \to Y$ determines a map $F(g)\colon F(X) \to F(Y)$ like this: $F(g)$ applied to a path component C of X is the unique path component of Y which contains $g(C)$.
Fix a positive integer n. Let $\mathcal{R}ings$ be the category of rings and ring homomorphisms. (For me, a ring does not have to be commutative, but it should have distinguished elements 0 and 1 and in this example I require $0 \neq 1$.) A functor F from $\mathcal{R}ings$ to $\mathcal{G}roups$ can be defined by $F(R) = \mathrm{GL}_n(R)$, where $\mathrm{GL}_n(R)$ is the group of invertible $n \times n$ matrices with entries in R. A ring homomorphism $g\colon R_1 \to R_2$ determines a group homomorphism $F(g)$ from $F(R_1)$ to $F(R_2)$. Namely, in an invertible $n \times n$-matrix with entries in R_1 , apply g to each entry to obtain an invertible $n \times n$-matrix with entries in R_2.
Let G be a group which comes with an action on a set S. In example 3.2.2 we constructed from G a category with one object $*$ and $\mathrm{mor}(*,*) = G$. A functor F from that category to $\mathcal{S}ets$ can now be defined by $F(*) = S$, and $F(g) =$ translation by g, for $g \in \mathrm{mor}(*,*) = G$. More precisely, to $g \in G = \mathrm{mor}(*,*)$ we associate the map $F(g)$ from $S = F(*)$ to $S = F(*)$ given by $x \mapsto g \cdot x$ (which has a meaning because we are assuming an action of G on S).
Let $\mathcal{C}$ be any category and let x be any object of $\mathcal{C}$. A functor F_x from $\mathcal{C}$ to $\mathcal{S}ets$ can be defined as follows. Let $F_x(c) = \mathrm{mor}_{\mathcal{C}}(x, c)$. For a morphism $g\colon c \to d$ in $\mathcal{C}$ define $F_x(g)\colon F_x(c) \to F_x(d)$ by $F_x(g)(h) = g \circ h$. In more detail, we are assuming $h \in F_x(c) = \mathrm{mor}_{\mathcal{C}}(x, c)$ and $g \in \mathrm{mor}_{\mathcal{C}}(c, d)$, so that $g \circ h \in \mathrm{mor}_{\mathcal{C}}(x, d) = F_x(d)$.

The functors of definition 3.2.3 are also called *covariant functors* for more precision. There is a related concept of *contravariant functor*. A contravariant functor from $\mathcal{C}$ to $\mathcal{D}$ is simply a (covariant) functor from $\mathcal{C}^{\mathrm{op}}$ to $\mathcal{D}$ (see example 3.2.2). If we write this out, it looks like this. A contravariant functor F from $\mathcal{C}$ to $\mathcal{D}$ is a rule which to every object c of $\mathcal{C}$ assigns an object $F(c)$ of $\mathcal{D}$, and to every morphism $g\colon c \to d$ in $\mathcal{C}$ a morphism $F(g)\colon F(d) \to F(c)$; note that the source of $F(g)$ is $F(d)$, and the target is $F(c)$. And so on.

EXAMPLE 3.2.5. Let $\mathcal{C}$ be any category and let x be any object of $\mathcal{C}$. A contravariant functor F^x from $\mathcal{C}$ to $\mathcal{S}ets$ can be defined as follows. Let $F^x(c) = \mathrm{mor}_{\mathcal{C}}(c, x)$. For a morphism $g\colon c \to d$ in $\mathcal{C}$ define

$$F^x(g)\colon F^x(d) \to F^x(c)$$

by $F^x(g)(h) = h \circ g$. In more detail, we are assuming $h \in F^x(d) = \mathrm{mor}_{\mathcal{C}}(d, x)$ and $g \in \mathrm{mor}_{\mathcal{C}}(c, d)$, so that $h \circ g \in \mathrm{mor}_{\mathcal{C}}(c, x) = F^x(c)$.
There is a contravariant functor P from $\mathcal{S}ets$ to $\mathcal{S}ets$ given by $P(S) =$ power set of S, for a set S. In more detail, a morphism $g\colon S \to T$ in $\mathcal{S}ets$ determines a map $P(g)\colon P(T) \to P(S)$ by "preimage". That is, $P(g)$ applied to a subset U of T is $g^{-1}(U)$, a subset of S. (You may have noticed that this example of a contravariant functor is not very different from a special case of the preceding one; we will return to this point later.)
Next, let $\mathcal{M}an$ be the category of smooth manifolds. The objects are the smooth manifolds (of any dimension). The morphisms from a smooth manifold M to a smooth manifold N are the smooth maps from M to N. For any fixed integer $k \geq 0$ the rule which assigns to a smooth manifold M the real vector space $\Omega^k(M)$ of smooth differential k-forms is a contravariant functor from $\mathcal{M}an$ to the category $\mathcal{V}ect$ of real vector spaces (with linear maps as morphisms). Namely, a smooth map $f\colon M \to N$ determines a linear map $f^*\colon \Omega^k(N) \to \Omega^k(M)$. (You must have seen the details if you know anything about differential forms.)

A presheaf $\mathcal{F}$ on a topological space X is nothing but a contravariant functor from the poset of open subsets of X to $\mathcal{S}$ets. In more detail, write $\mathcal{O}$ for the topology on X, the set of open subsets of X. We can regard $\mathcal{O}$ as a partially ordered set (poset) in the following way: for $U, V \in \mathcal{O}$ we decree that $U \le V$ if and only if $U \subset V$. A partially ordered set is a small category, as explained in example 3.2.2; therefore $\mathcal{O}$ is (the object set of) a small category. For $U, V \in \mathcal{O}$, there is exactly one morphism from U to V if $U \subset V$, and none if U is not contained in V. To that one morphism (if $U \subset V$) the presheaf $\mathcal{F}$ assigns a map $\mathrm{res}_{V,U} \colon \mathcal{F}(V) \to \mathcal{F}(U)$. The conditions on $\mathcal{F}$ in definition 3.1.1 are special cases of the conditions on a contravariant functor.

The story does not end there. The functors from a category $\mathcal{C}$ to a category $\mathcal{D}$ also form something like a category. There is a concept of morphism between functors from $\mathcal{C}$ to $\mathcal{D}$, and the name of that is *natural transformation.*

Definition 3.2.6. Let F and G be functors, both from a category $\mathcal{C}$ to a category $\mathcal{D}$. A *natural transformation* from F to G is a rule ν which for every object c in $\mathcal{C}$ selects a morphism $\nu_c \colon F(c) \to G(c)$ in $\mathcal{D}$, subject to the following condition. Whenever $u \colon c \to d$ is a morphism in $\mathcal{C}$, the square of morphisms

$$\begin{tikzcd} F(c) \arrow[r, "\nu_c"] \arrow[d, "F(u)"] & G(c) \arrow[d, "G(u)"] \\ F(d) \arrow[r, "\nu_d"] & G(d) \end{tikzcd}$$

in $\mathcal{D}$ commutes; that is, the equation $G(u) \circ \nu_c = \nu_d \circ F(u)$ holds in $\mathrm{mor}_{\mathcal{D}}(F(c), G(d))$.

Example 3.2.7. MacLane (in his book *Categories for the working mathematician*) gives the following pretty example. For a fixed integer $n \ge 1$ the rule which to a ring R assigns the group $\mathrm{GL}_n(R)$ can be viewed as a functor GL_n from the category of rings to the category of groups, as was shown earlier. There we allowed non-commutative rings, but here we need commutative rings, so we shall view GL_n as a functor from the category c$\mathcal{R}$ings of commutative rings to $\mathcal{G}$roups. Note that $\mathrm{GL}_1(R)$ is essentially the group of units of the ring R. The group homomorphisms

$$\det \colon \mathrm{GL}_n(R) \to \mathrm{GL}_1(R)$$

(one for every commutative ring R) make up a natural transformation from the functor $\mathrm{GL}_n \colon \mathrm{c}\mathcal{R}\mathrm{ings} \to \mathcal{G}\mathrm{roups}$ to the functor $\mathrm{GL}_1 \colon \mathrm{c}\mathcal{R}\mathrm{ings} \to \mathcal{G}\mathrm{roups}$.
Returning to smooth manifolds and differential forms: we saw that for any fixed $k \ge 0$ the assignment $M \mapsto \Omega^k(M)$ can be viewed as a contravariant functor from $\mathcal{M}$an to $\mathcal{V}$ect. The exterior derivative maps

$$d \colon \Omega^k(M) \longrightarrow \Omega^{k+1}(M)$$

(one for each object M of $\mathcal{M}$an) make up a natural transformation from the contravariant functor Ω^k to the contravariant functor Ω^{k+1}.

Notation: let F and G be functors from $\mathcal{C}$ to $\mathcal{D}$. Sometimes we describe a natural transformation ν from F to G by a strong arrow, as in $\nu \colon F \Rightarrow G$.
Remark: one reason for being a little cautious in saying *category of categories* etc. is that the functors from one big category (such as $\mathcal{T}$op for example) to another big category (such as $\mathcal{S}$ets for example) do not obviously form a set. Of course, some people would not exercise that kind of caution and would instead say that the definition of category as

given in 3.2.1 is not bold enough. In any case, it must be permitted to say *the category of small categories.*

3.3. The category of presheaves on a space

Let $X = (X, \mathcal{O})$ be a topological space. We have seen that a presheaf $\mathcal{F}$ on X is the same thing as a contravariant functor from the poset $\mathcal{O}$ (partially ordered by inclusion, and then viewed as a category) to $\mathcal{S}\mathrm{ets}$. Therefore it is not surprising that we define a *morphism* from a presheaf $\mathcal{F}$ on X to a presheaf $\mathcal{G}$ on X to be a natural transformation between contravariant functors from $\mathcal{O}$ to $\mathcal{S}\mathrm{ets}$. Writing this out in detail, we obtain the following definition.

DEFINITION 3.3.1. Let $\mathcal{F}$ and $\mathcal{G}$ be presheaves on the topological space X. A *morphism* or *map* of presheaves from $\mathcal{F}$ to $\mathcal{G}$ is a rule which for every open set U in X selects a map $\lambda_U \colon \mathcal{F}(U) \to \mathcal{G}(U)$, subject to the following condition. Whenever U and V are open subsets of X and $U \subset V$, the diagram

$$\begin{tikzcd} \mathcal{F}(U) \arrow[r, "\lambda_U"] & \mathcal{G}(U) \\ \mathcal{F}(V) \arrow[u, "\mathrm{res}_{V,U}"] \arrow[r, "\lambda_V"] & \mathcal{G}(V) \arrow[u, "\mathrm{res}_{V,U}"'] \end{tikzcd}$$

in $\mathcal{S}\mathrm{ets}$ commutes; that is, the maps $\mathrm{res}_{V,U} \circ \lambda_V$ and $\lambda_U \circ \mathrm{res}_{V,U}$ from $\mathcal{F}(V)$ to $\mathcal{G}(U)$ agree.

With this definition of morphism, it is clear that there is a category of presheaves on X. It is a small category.

EXAMPLE 3.3.2. Let X be a topological space. Let $\mathcal{F}$ be the presheaf on X such that $\mathcal{F}(U)$, for open $U \subset X$, is the set of continuous maps from U to $\mathbb{R}$, and such that $\mathrm{res}_{V,U} \colon \mathcal{F}(V) \to \mathcal{F}(U)$ is given by restriction of functions. Let $\mathcal{G}$ be the presheaf on X such that $\mathcal{G}(U)$, for open $U \subset X$, is the set of all open subsets of X which are contained in U. More precisely $\mathcal{G}$ is a presheaf because in the situation $U \subset V$ we define

$$\mathrm{res}_{V,U} \colon \mathcal{G}(V) \to \mathcal{G}(U)$$

by $W \mapsto W \cap U$ for an open subset W of X contained in V. (Then $W \cap U$ is an open subset of X contained in U.) A morphism α from presheaf $\mathcal{F}$ to presheaf $\mathcal{G}$ is defined by

$$\alpha_U(g) = g^{-1}(\,]0,\infty[\,)$$

for $g \in \mathcal{F}(U)$. In a more wordy formulation: to an element g of $\mathcal{F}(U)$, alias continuous function $g \colon U \to \mathbb{R}$, the morphism $\alpha \colon \mathcal{F} \to \mathcal{G}$ assigns an element of $\mathcal{G}(U)$, alias open set of X contained in U, by taking the preimage of $]0,\infty[$ under g.

3.4. (Appendix): Abelian group vocabulary

It is customary to describe the binary operation in an abelian group by a $+$ sign, if there is no danger of confusion. Thus, if A is an abelian group and $a, b \in A$, we like to write $a + b$ instead of ab or $a \cdot b$; also $-b$ instead of b^{-1} and 0 instead of 1 for the neutral element.

The expression *abelian group* is synonymous with $\mathbb{Z}$*-module*. The name $\mathbb{Z}$*-module* is a reminder that there is some interaction between the ring $\mathbb{Z}$ and the elements of any abelian group A. This looks a lot like the multiplication of vectors by scalars in a vector space. Namely, let A be an abelian group (written with $+$ etc.), let a be an element of A and $z \in \mathbb{Z}$. Then we can define

$$z \cdot a \in A$$

as follows: if $z \geq 0$ we mean $a + a + \cdots + a$ (there are z summands in the sum); if $z \leq 0$ then we know already what $(-z) \cdot a$ means and $z \cdot a$ should be the inverse, $z \cdot a = -((-z) \cdot a)$. This "scalar multiplication" has an associativity property:

$$(wz) \cdot a = w \cdot (z \cdot a)$$

and also two distributivity properties, $(w+z) \cdot a = w \cdot a + z \cdot a$ as well as $z \cdot (a+b) = z \cdot a + z \cdot b$. Furthermore, $1 \cdot a = a$ for all $a \in A$ and $z \cdot 0 = 0$ for all $z \in \mathbb{Z}$. We might feel tempted to say that A is a vector space over the field $\mathbb{Z}$, but there is the objection that $\mathbb{Z}$ is not a field.

(Of course there is a more general concept of R-*module*, where R can be any ring. An R-module is an abelian group A with a map $R \times A \to A$ which we write in the form $(r, a) \mapsto r \cdot a$. That map is subject to many conditions, such as $(rs) \cdot a = r \cdot (s \cdot a)$ and $r \cdot (a + b) = r \cdot a + r \cdot b$, for all $r \in R$ and $a, b \in A$, and a few more. Look it up in any algebra book.)

Definition 3.4.1. Let S be a set. The *free abelian group* generated by S is the set A_S of all functions $f \colon S \to \mathbb{Z}$ such that $\{s \in S \mid f(s) \neq 0\}$ is a *finite* subset of S. It is an abelian group by pointwise addition; that is, for $f, g \in A_S$ we define $f + g \in A_S$ by $(f + g)(s) = f(s) + g(s) \in \mathbb{Z}$.

Notation. Elements of the free abelian group A_S generated by S can also be thought of as *formal linear combinations*, with integer coefficients, of elements of S. In other words, we may write

$$\sum_{s \in S} a_s \cdot s$$

where $a_s \in \mathbb{Z}$ for all $s \in \mathbb{Z}$, and we mean the function $f \in A_S$ such that $f(s) = a_s$ for all $s \in S$. Now it is important to insist that the sum have only finitely many (nonzero) summands, $a_s \neq 0$ for only finitely many $s \in S$.
My notation A_S for the free abelian group generated by S is meant to be temporary. I can't think of any convincing standard notation for it.

An important property of the free abelian group generated by S. The group A_S comes with a distinguished map $u \colon S \to A_S$ so that $u(s)$ is the function from S to $\mathbb{Z}$ taking s to 1 and all other elements of S to 0. Together, the abelian group A_S and the map (of sets) $u \colon S \to A_S$ have the following property. *Given any abelian group* B *and map* $v \colon S \to B$, *there exists a unique homomorphism of abelian groups* $q_v \colon A_S \to B$ *such that* $q_v \circ u = v$. Diagrammatic statement:

$$\begin{tikzcd} S \arrow[r, "u"] \arrow[dr, "v"'] & A_S \arrow[d, dotted, "q_v"] \\ & B \end{tikzcd}$$

The proof is easy. Every element a of A_S can be written uniquely in the form

$$\sum_{s \in S} a_s \cdot u(s)$$

with $a_s \in \mathbb{Z}$, with only finitely many nonzero a_s. Therefore

$$q_v(a) = q_v\Big(\sum_{s\in S} a_s \cdot u(s)\Big) = \sum_{s\in S} q_v\big(a_s \cdot u(s)\big) = \sum_{s\in S} a_s \cdot q_v(u(s)) = \sum_{s\in S} a_s \cdot v(s) \ .$$

(The following complaint can be made: *Just a minute ago you said that we can write elements* a *of* A_S *in the form* $\sum_{s\in S} a_s \cdot s$ *, but now it is* $\sum_{s\in S} a_s \cdot u(s)$*, or what?* The complaint is justified: $\sum_{s\in S} a_s \cdot s$ is a short and imprecise form of $\sum_{s\in S} a_s \cdot u(s)$.)

4.5. (Appendix): Preview

If our main interest is in understanding notions like homotopy and classifying topological spaces up to homotopy equivalence, why should we learn something about presheaves and sheaves? In this appendix I try to give an answer, very much from the point of view of category theory.

Summarizing the experience of the first few weeks in category language, we might agree on the following. In the category $\mathcal{T}\mathrm{op}$ of topological spaces (and continuous maps), we introduced the homotopy relation $\simeq$ on morphisms. This led to a new category $\mathcal{H}\mathrm{o}\mathcal{T}\mathrm{op}$ with the same objects as $\mathcal{T}\mathrm{op}$, where a morphism from X to Y is a homotopy class of continuous maps from X to Y. We made some attempts to understand sets of homotopy classes $[X,Y] = \mathrm{mor}_{\mathcal{H}\mathrm{o}\mathcal{T}\mathrm{op}}(X,Y)$ in some cases; for example we understood $[S^1,S^1]$ and we showed that $[S^3,S^2]$ has more than one element. A vague impression of computability may have taken hold, but nothing very systematic emerged.

Here is a very simple-minded attempt to make things easier by introducing some algebra into topology. We can make a new category $\mathbb{Z}\mathcal{T}\mathrm{op}$ where the objects are still the topological spaces and where the set of morphisms from X to Y is the *free abelian group* generated by the set of continuous maps from X to Y. In other words, a morphism from X to Y in $\mathbb{Z}\mathcal{T}\mathrm{op}$ is a formal linear combination (with integer coefficents) of continuous maps from X to Y, such as $4f - 3g + 7u + 1v$, where $f, g, u, v\colon X \to Y$ are continuous maps. Note that *formal* is *formal*; we make no attempt to simplify such expressions, except by allowing $4f - 3g + 7u + 1v = 4f + 4u + 1v$ if we happen to know that $g = u$, and the like. How do we compose morphisms in $\mathbb{Z}\mathcal{T}\mathrm{op}$? We use the composition of morphisms in $\mathcal{T}\mathrm{op}$ and enforce a distributive law, so we say for example that the composition of the morphism $4f - 3g + 7u$ from X to Y with the morphism $-2p + 5q$ from Y to Z is

$$-8(p \circ f) + 6(p \circ g) - 14(p \circ u) + 20(q \circ f) - 15(q \circ g) + 35(q \circ u),$$

a morphism from X to Z. In many ways $\mathbb{Z}\mathcal{T}\mathrm{op}$ is a fine category, and perhaps better than $\mathcal{T}\mathrm{op}$; the morphism sets are abelian groups and composition of morphisms

$$\mathrm{mor}_{\mathbb{Z}\mathcal{T}\mathrm{op}}(Y,Z) \times \mathrm{mor}_{\mathbb{Z}\mathcal{T}\mathrm{op}}(X,Y) \longrightarrow \mathrm{mor}_{\mathbb{Z}\mathcal{T}\mathrm{op}}(X,Z)$$

is bilinear. That is, post-composition with a fixed element of $\mathrm{mor}_{\mathbb{Z}\mathcal{T}\mathrm{op}}(Y,Z)$ gives a homomorphism of abelian groups, from $\mathrm{mor}_{\mathbb{Z}\mathcal{T}\mathrm{op}}(X,Y)$ to $\mathrm{mor}_{\mathbb{Z}\mathcal{T}\mathrm{op}}(X,Z)$, and pre-composition with a fixed element of $\mathrm{mor}_{\mathbb{Z}\mathcal{T}\mathrm{op}}(X,Y)$ gives a homomorphism of abelian groups from $\mathrm{mor}_{\mathbb{Z}\mathcal{T}\mathrm{op}}(Y,Z)$ to $\mathrm{mor}_{\mathbb{Z}\mathcal{T}\mathrm{op}}(X,Z)$. We can relate $\mathcal{T}\mathrm{op}$ to $\mathbb{Z}\mathcal{T}\mathrm{op}$ by a functor

$$\mathcal{T}\mathrm{op} \to \mathbb{Z}\mathcal{T}\mathrm{op}$$

which takes any object to the same object, and a continuous map $f\colon X \to Y$ to the formal linear combination $1f$. And yet, it is hard to believe that any of this will give us new insights into anything.

But let us try to raise a well-formulated objection. We have lost something in replacing $\mathcal{T}\mathrm{op}$ by $\mathbb{Z}\mathcal{T}\mathrm{op}$: the sheaf property. More precisely, we know that we can construct a continuous map $f\colon X \to Y$ by specifying an open cover $(U_i)_{i\in\Lambda}$ of X, and for each i a continuous map $f_i\colon U_\imath \to Y$, in such a way that

$$f_{i|U_i\cap U_j} = f_{j|U_i\cap U_j}$$

for all $i, j \in \Lambda$. (Then there is a unique continuous map $f\colon X \to Y$ such that $f_{|U_i} = f_i$ for all $i \in \Lambda$.) We could take the view that this is a property of $\mathcal{T}\mathrm{op}$ which is important to us, one that we don't want to sacrifice when we experiment with modifications of $\mathcal{T}\mathrm{op}$. But as we have seen, the sheaf property fails in so many ways in $\mathbb{Z}\mathcal{T}\mathrm{op}$; see example 3.1.7 and the elaborate discussion of that example. I propose that we regard that as the one great weakness of $\mathbb{Z}\mathcal{T}\mathrm{op}$.

Fortunately, in sheaf theory there is a fundamental construction called *sheafification* by which the sheaf property is enforced. In the following chapters we will apply that construction to $\mathbb{Z}\mathcal{T}\mathrm{op}$ to restore the sheaf property. When that is done, we can once again speak of homotopies and homotopy classes, and it will turn out that we have a very manageable situation.

CHAPTER 4

Sheafification

4.1. The stalks of a presheaf

Let $\mathcal{F}$ a presheaf on a topological space X. Fix $z \in X$. There are situations where we want to understand the behavior of $\mathcal{F}$ near z, that is to say, in small neighborhoods of z. Then it is a good idea to work with pairs (U, s) where U is an open neighborhood of z and s is an element of $\mathcal{F}(U)$. Two such pairs (U, s) and (V, t) are considered to be *germ-equivalent* if there exists an open neighborhood W of z such that $W \subset U \cap V$ and $s_{|W} = t_{|W}$ in $\mathcal{F}(W)$. It is easy to show that germ equivalence is indeed an equivalence relation.

DEFINITION 4.1.1. The set of equivalence classes is called the *stalk* of $\mathcal{F}$ at z and denoted by $\mathcal{F}_z$. The elements of $\mathcal{F}_z$ are often called *germs* (at z, of something ... depending on the meaning of $\mathcal{F}$).

EXAMPLE 4.1.2. Let X and Y be topological spaces. Let $\mathcal{F}$ be the sheaf on X where $\mathcal{F}(U)$, for open $U \subset X$, is the set of continuous maps from U to Y. For $z \in X$, an element of the stalk $\mathcal{F}_z$ is called a *germ of continuous maps from* (X, z) *to* Y.

EXAMPLE 4.1.3. Fix a continuous map $p\colon Y \to X$. Let $\mathcal{F}$ be the sheaf on X where $\mathcal{F}(U)$ is the set of continuous maps $s\colon U \to Y$ such that $p \circ s$ is the inclusion $U \to X$. An element of $\mathcal{F}(U)$ can be called a *continuous section of* p *over* U. For $z \in X$, an element of $\mathcal{F}_z$ can be called a *germ at* z *of continuous sections of* $p\colon X \to Y$.

EXAMPLE 4.1.4. Let X be the union of the two coordinate axes in $\mathbb{R}^2$. For open U in X, let $\mathcal{G}(U)$ be the set of connected components of $X \smallsetminus U$. For open subsets U, V of X such that $U \subset V$, define

$$\mathrm{res}_{V,U}\colon \mathcal{G}(V) \to \mathcal{G}(U)$$

by saying that $\mathrm{res}_{V,U}(C)$ is the unique connected component of $X \smallsetminus U$ which contains C (where C can be any connected component of $X \smallsetminus V$). These definitions make $\mathcal{G}$ into a presheaf on X. For $z \in X$, what can we say about the stalk $\mathcal{G}_z$? If z is the origin, $z = (0, 0)$, then $\mathcal{G}_z$ has four elements. In all other cases $\mathcal{G}_z$ has two elements. (Despite that, for any $z \in X$ and any open neighborhood V of z in X, there exists an open neighborhood W of z in X such that $W \subset V$ and $\mathcal{G}(W)$ has more than 1000 elements.)

Now let $\alpha\colon \mathcal{F} \to \mathcal{G}$ be a map (morphism) of sheaves on X. Again fix $z \in X$. Then every pair (U, s), where U is an open neighborhood of z and $s \in \mathcal{F}(s)$, determines another pair $(U, \alpha(s))$ where U is still an open neighborhood of z and $\alpha(s) \in \mathcal{G}(U)$. The assignment $(U, s) \mapsto (U, \alpha(s))$ is compatible with germ equivalence. That is, if V is another open neighborhood of z in X, and $t \in \mathcal{F}(V)$, and (U, s) is germ equivalent to (V, t), then $(U, \alpha(s))$ is germ equivalent to $(V, \alpha(t))$. In short, α determines a map of sets from $\mathcal{F}_z$ to $\mathcal{G}_z$ which takes the equivalence class (the germ) of (U, s) to the equivalence class (the germ) of $(U, \alpha(s))$. In category jargon: the assignment

$$\mathcal{F} \mapsto \mathcal{F}_z$$

is a *functor* from $\mathcal{P}\mathrm{re}\mathcal{S}\mathrm{h}(X)$, the category of presheaves on X, to $\mathcal{S}\mathrm{ets}$.

When a presheaf $\mathcal{F}$ on X is a sheaf, the stalks $\mathcal{F}_z$ carry a lot of information about $\mathcal{F}$. The following theorem illustrates that.

THEOREM 4.1.5. *Let $\beta\colon \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves on X. Suppose that for every element z of X, the map of stalks $\mathcal{F}_z \to \mathcal{G}_z$ determined by β is a bijection. Then β is an isomorphism.*

PROOF. The claim that β is an isomorphism means, abstractly speaking, that there exists a morphism $\gamma\colon \mathcal{G} \to \mathcal{F}$ of sheaves such that $\beta \circ \gamma$ is the identity on $\mathcal{G}$ and $\gamma \circ \beta$ is the identity on $\mathcal{F}$. In more down-to-earth language it means simply that $\beta_U\colon \mathcal{F}(U) \to \mathcal{G}(U)$ is a bijection for every open U in X, so this is what we have to show. To ease notation, we write $\beta\colon \mathcal{F}(U) \to \mathcal{G}(U)$.
We fix U, an open subset of X. First we want to show that $\beta\colon \mathcal{F}(U) \to \mathcal{F}(G)$ is *injective.* For that we set up a commutative square of sets and maps:

$$\begin{tikzcd} \prod_{z \in U} \mathcal{F}_z \arrow[r, "\beta"] & \prod_{z \in U} \mathcal{G}_z \\ \mathcal{F}(U) \arrow[u] \arrow[r, "\beta"] & \mathcal{G}(U) \arrow[u] \end{tikzcd}$$

The left-hand vertical arrow is obtained by noting that each $s \in \mathcal{F}(U)$ determines a pair (U, s) representing an element of $\mathcal{F}_z$, for each $z \in U$. The right-hand vertical arrow is similar. We show that the left-hand vertical arrow is injective. Suppose that $s, t \in \mathcal{F}(U)$ have the same image in $\prod_{z \in U} \mathcal{F}_z$. It follows that every $z \in U$ admits a neighborhood W_z in U such that $s_{|W_z} = t_{|W_z}$. Selecting such a W_z for every $z \in U$, we have an open cover

$$(W_z)_{z \in U}$$

of U. Since $s_{|W_z} = t_{|W_z}$ for each of the open sets W_z in the cover, the sheaf property for $\mathcal{F}$ implies that $s = t$. Hence the left-hand vertical arrow in our square is injective, and so is the right-hand arrow by the same argument. But the top horizontal arrow is bijective by our assumption. Therefore $\beta\colon \mathcal{F}(U) \to \mathcal{F}(G)$ is injective.
Next we show that $\beta\colon \mathcal{F}(U) \to \mathcal{F}(G)$ is *surjective.* We can use the same commutative square that we used to prove injectivity. An element $s \in \mathcal{G}(U)$ determines an element of $\prod_{z \in U} \mathcal{G}_z$ (right-hand vertical arrow) which comes from an element of $\prod_{z \in U} \mathcal{F}_z$ because the top horizontal arrow is bijective. So for each $z \in U$ we can find an element of $\mathcal{F}_z$ which under β is mapped to the germ of s at z (an element of $\mathcal{G}_z$). In terms of representatives of germs, this means that for each $z \in U$ we can find an open neighborhood V_z of z in U and an element $t_z \in \mathcal{F}(V_z)$ such that $\beta(t_z) = s_{|V_z} \in \mathcal{G}(V_z)$. Selecting such a V_z for every $z \in U$, we have an open cover

$$(V_z)_{z \in U}$$

of U and we have $t_z \in \mathcal{F}(V_z)$. Can we use the sheaf property of $\mathcal{F}$ to produce $t \in \mathcal{F}(U)$ such that $t_{|V_z} = t_z$ for all $z \in U$? We need to verify the matching condition,

$$t_{z|V_z \cap V_y} = t_{y|V_z \cap V_y} \in \mathcal{F}(V_z \cap V_y)$$

whenever $y, z \in U$. By the injectivity of $\beta\colon \mathcal{F}(V_z \cap V_y) \to \mathcal{G}(V_z \cap V_y)$, which we have established, it is enough to show

$$\beta(t_z)_{|V_z \cap V_y} = \beta(t_y)_{|V_z \cap V_y} \in \mathcal{G}(V_z \cap V_y).$$

This clearly holds as $\beta(t_z) = s_{|V_z}$ by construction, so that both sides of the equation agree with $s_{|V_z \cap V_y}$. So we obtain $t \in \mathcal{F}(U)$ such that $t_{|V_z} = t_z$ for all $z \in U$. Now it is easy to show that $\beta(t) = s$. Indeed we have $\beta(t)_{|V_z} = s_{|V_z}$ by construction, for all open sets V_z in the covering $(V_z)_{z \in U}$ of U, so the sheaf property of $\mathcal{F}$ implies $\beta(t) = s$. Since $s \in \mathcal{G}(U)$ was arbitrary, this means that $\beta\colon \mathcal{F}(U) \to \mathcal{G}(U)$ is surjective. □

4.2. Sheafification of a presheaf

PROPOSITION 4.2.1. *Let $\mathcal{F}$ be a presheaf on a topological space X. There is a sheaf $\Phi\mathcal{F}$ on X and there is a morphism $\eta\colon \mathcal{F} \to \Phi\mathcal{F}$ of presheaves such that, for every $z \in X$, the map of stalks $\mathcal{F}_z \to (\Phi\mathcal{F})_z$ determined by η is bijective.*

PROOF. Let U be an open subset of X. We are going to define $(\Phi\mathcal{F})(U)$ as a subset of the product

$$\prod_{z \in U} \mathcal{F}_z \, .$$

Think of an element of that product as a function s which for every $z \in U$ selects an element $s(z) \in \mathcal{F}_z$. The function s qualifies as an element of $(\Phi\mathcal{F})(U)$ if and only if it satisfies the following *coherence condition.* For every $y \in U$ there is an open neighborhood W of y in U and there is $t \in \mathcal{F}(W)$ such that the pair (W, t) simultaneously represents the germs $s(z) \in \mathcal{F}_z$ for all $z \in W$.

From the definition, it is clear that there are restriction maps

$$\mathrm{res}_{V,U}\colon (\Phi\mathcal{F})(V) \to (\Phi\mathcal{F})(U)$$

whenever U, V are open in X and $U \subset V$. Namely, a function s which selects an element $s(z) \in \mathcal{F}_z$ for every $z \in V$ determines by restriction a function $s_{|U}$ which selects an element $s(z) \in \mathcal{F}_z$ for every $z \in U$. The coherence condition is satisfied by $s_{|U}$ if it is satisfied by s. With these restriction maps, $\Phi\mathcal{F}$ is a presheaf. Furthermore, it is straightforward to see that $\Phi\mathcal{F}$ satisfies the sheaf condition. Indeed, suppose that $(V_i)_{i \in \Lambda}$ is a collection of open subsets of X, and suppose that elements $s_i \in (\Phi\mathcal{F})(V_i)$ have been selected, one for each $i \in \Lambda$, such that the matching condition

$$s_{i|V_i \cap V_j} = s_{j|V_i \cap V_j}$$

is satisfied for all $i, j \in \Lambda$. Then clearly we get a function s on $V = \bigcup_i V_i$ which for every $z \in V$ selects $s(z) \in \mathcal{F}_z$ by declaring, unambiguously,

$$s(z) := s_i(z)$$

for any i such that $z \in V_i$. The coherence condition is satisfied because it is satisfied by each s_i .

The morphism of presheaves $\eta\colon \mathcal{F} \to \Phi\mathcal{F}$ is defined in the following mechanical way. Given $t \in \mathcal{F}(U)$, we need to say what $\eta(t) \in (\Phi\mathcal{F})(U)$ should be. It is the function which to $z \in U$ assigns the element of $\mathcal{F}_z$ represented by the pair (U, t), that is to say, the germ of (U, t) at z.

Last not least, we need to show that for any $z \in X$ the map $\mathcal{F}_z \to (\Phi\mathcal{F})_z$ determined by η is a bijection. We fix z. *Injectivity*: we consider elements a and b of $\mathcal{F}_z$ represented by pairs (U_a, s_a) and (U_b, s_b) respectively, where U_a, U_b are neighborhoods of z and

$s_a \in \mathcal{F}(U_a)$, $s_b \in \mathcal{F}(U_b)$. Suppose that a and b are taken to the same element $t \in (\Phi\mathcal{F})_z$ by η. Then in particular $t(z) \in \mathcal{F}_z$ is the germ at z of s_a, and also the germ at z of s_b, so the germs of s_a and s_b (elements of $\mathcal{F}_z$) are equal. *Surjectivity*: let an element of $(\Phi\mathcal{F})_z$ be represented by a pair (U, t) where U is an open neighborhood of z in X and $t \in (\Phi\mathcal{F})(U)$. By the coherence condition, there exists an open neighborhood W of z in U and there exists $s \in \mathcal{F}(W)$ such that $t_{|W}$ is the function which to $y \in W$ assigns the germ at y of (W, s), an element of $\mathcal{F}_y$. But this means that the map of stalks $\mathcal{F}_z \to (\Phi\mathcal{F})_z$ determined by the morphism η takes the element of $\mathcal{F}_z$ represented by (W, s) to the element of $(\Phi\mathcal{F})_z$ represented by (U, t). □

Example 4.2.2. Let T be any set. Let $\mathcal{F}$ be the constant presheaf on X given by $\mathcal{F}(U) = T$ for all open subsets U of X (and $\mathrm{res}_{V,U} : \mathcal{F}(V) \to \mathcal{F}(U)$ is id_T). What does the sheaf $\Phi\mathcal{F}$ look like? This question has quite an interesting answer. Let's keep a cool head and approach it mechanically. For any $z \in X$, the stalk $\mathcal{F}_z$ can be identified with T. This is easy. Let U be an open subset of X. The elements of $(\Phi\mathcal{F})(U)$ are functions s which for every $z \in U$ select an element $s(z) \in \mathcal{F}_z = T$, subject to a coherence condition. So the elements of $(\Phi\mathcal{F})(U)$ are maps s from U to T subject to a coherence condition. What is the coherence condition? The condition is that s must be locally constant, i.e., every $z \in U$ admits an open neighborhood W in U such that $s_{|W}$ is constant. So the elements of $(\Phi\mathcal{F})(U)$ are the locally constant maps s from U to T. A locally constant map s from U to T is the same thing as a continuous map s from U to T, if we agree that T is equipped with the discrete topology (every subset of T is declared to be open). Summing up, $(\Phi\mathcal{F})(U)$ is the set of continuous functions from U to T. We can say that $\Phi\mathcal{F}$ is the sheaf of continuous functions (from open subsets of X) to T.
To appreciate the beauty of this answer, take a space X which is a little out of the ordinary; for example, $\mathbb{Q}$ with the standard topology inherited from $\mathbb{R}$, or the Cantor set (a subset of $\mathbb{R}$). For T, any set with more than one element is an interesting choice. (What happens if T has exactly one element? What happens if $T = \emptyset$?)

There are a few things of a general nature to be said about proposition 4.2.1 — not difficult, not surprising, but important. The construction Φ is a functor; we can view it as a functor from the category $\mathcal{P}\mathrm{re}\mathcal{S}\mathrm{h}(X)$ to itself. This means in particular that any morphism of presheaves $\alpha : \mathcal{F} \to \mathcal{G}$ on X determines a morphism

$$\Phi\alpha : \Phi\mathcal{F} \to \Phi\mathcal{G} .$$

Namely, for $s \in \Phi\mathcal{F}(V)$ we define $t = (\Phi\alpha)(s) \in \Phi\mathcal{G}(V)$ in such a way that $t(z) \in \mathcal{G}_z$ is the image of $s(z) \in \mathcal{F}_z$ under the map $\mathcal{F}_z \to \mathcal{G}_z$ induced by α. (It is easy to verify that t satisfies the coherence condition.)
Furthermore η is a natural transformation from the identity functor id on $\mathcal{P}\mathrm{re}\mathcal{S}\mathrm{h}(X)$ to the functor $\Phi : \mathcal{P}\mathrm{re}\mathcal{S}\mathrm{h}(X) \to \mathcal{P}\mathrm{re}\mathcal{S}\mathrm{h}(X)$. This means that, for a morphism of presheaves $\alpha : \mathcal{F} \to \mathcal{G}$ on X as above, the diagram

$$\begin{tikzcd} \mathcal{F} \arrow[r, "\alpha"] \arrow[d, "\eta"] & \mathcal{G} \arrow[d, "\eta"] \\ \Phi\mathcal{F} \arrow[r, "\Phi\alpha"] & \Phi\mathcal{G} \end{tikzcd}$$

in $\mathcal{P}\mathrm{re}\mathcal{S}\mathrm{h}(X)$ is commutative. That is also easily verified.
There is one more thing of a general nature which must be mentioned. Let $\mathcal{F}$ be any

presheaf on X. What happens if we apply the functor Φ to the morphism $\eta_{\mathcal{F}}\colon \mathcal{F} \to \Phi\mathcal{F}$? The result is obviously a morphism of sheaves

$$\Phi(\eta_{\mathcal{F}})\colon \Phi\mathcal{F} \to \Phi(\Phi\mathcal{F}).$$

It is an *isomorphism* of sheaves. The verification is easy using theorem 4.1.5.

The sheaf $\Phi\mathcal{F}$ is *the sheafification* (or the *associated sheaf*) of the presheaf $\mathcal{F}$; also Φ may be called the sheafification functor, or the associated sheaf functor.

COROLLARY 4.2.3. *Let* $\beta\colon \mathcal{F} \to \mathcal{G}$ *be any morphism of presheaves on* X. *If* $\mathcal{G}$ *is a sheaf, then* β *has a unique factorization* $\beta = \beta_1 \circ \eta_{\mathcal{F}}$ *where* $\eta_{\mathcal{F}}\colon \mathcal{F} \to \Phi\mathcal{F}$ *is the morphism of proposition 4.2.1:*

$$\begin{tikzcd}
\mathcal{F} \arrow[r, "\beta"] \arrow[d, "\eta_{\mathcal{F}}"] & \mathcal{G} \\
\Phi\mathcal{F} \arrow[ur, dotted, "\beta_1"'] &
\end{tikzcd}$$

PROOF. Apply Φ and η to $\mathcal{F}$, $\mathcal{G}$ and β to obtain a commutative diagram

$$\begin{tikzcd}
\mathcal{F} \arrow[r, "\beta"] \arrow[d, "\eta_{\mathcal{F}}"] & \mathcal{G} \arrow[d, "\eta_{\mathcal{G}}"] \\
\Phi\mathcal{F} \arrow[r, "\Phi\beta"] & \Phi\mathcal{G}
\end{tikzcd}$$

By proposition 4.2.1, the vertical arrows determine bijections from $\mathcal{F}_z$ to $(\Phi\mathcal{F})_z$ and from $\mathcal{G}_z$ to $(\Phi\mathcal{G})_z$ for every $z \in X$. Both $\mathcal{G}$ and $\Phi\mathcal{G}$ are sheaves, so theorem 4.1.5 applies and we may deduce that the right-hand vertical arrow is an isomorphism of sheaves on X. Let $\lambda\colon \Phi\mathcal{G} \to \mathcal{G}$ be an inverse for that isomorphism. The factorization problem has a solution, $\beta_1 = \lambda \circ \Phi\beta$.
To see that the solution is unique, apply Φ and η to the commutative diagram

$$\begin{tikzcd}
\mathcal{F} \arrow[r, "\beta"] \arrow[d, "\eta_{\mathcal{F}}"] & \mathcal{G} \\
\Phi\mathcal{F} \arrow[ur, dotted] &
\end{tikzcd}$$

in $\mathcal{P}re\mathcal{S}h(X)$. The result is a commutative diagram in $\mathcal{P}re\mathcal{S}h(X)$ in the shape of a prism:

$$\begin{tikzcd}
\mathcal{F} \arrow[rr, "\beta"] \arrow[d, "\eta_{\mathcal{F}}"] \arrow[drr] & & \mathcal{G} \arrow[drr, "\eta_{\mathcal{G}}"] & & \\
\Phi\mathcal{F} \arrow[urr, dotted] \arrow[drr] & & \Phi\mathcal{F} \arrow[rr, "\Phi\beta"] \arrow[d, "\Phi(\eta_{\mathcal{F}})"] & & \Phi\mathcal{G} \\
& & \Phi(\Phi\mathcal{F}) \arrow[urr, dotted] & &
\end{tikzcd}$$

Here the arrow labeled $\Phi(\eta_{\mathcal{F}})$ is an isomorphism of sheaves, as noted above under things of a general nature. This makes the lower dotted arrow unique. But the arrow labeled $\eta_{\mathcal{G}}$

is also an isomorphism by theorem 4.1.5 and the property of $\eta_{\mathcal{G}}$ stated in proposition 4.2.1. This ensures that the upper dotted arrow is determined by the lower dotted arrow. $\square$

4.3. Mapping cycles

Let X and Y be topological spaces. One of the first examples of a sheaf that we saw was the sheaf $\mathcal{F}$ on X such that

$$\mathcal{F}(U) = \text{set of continuous maps from } U \text{ to } Y$$

etc., for open U in X. From that we constructed a presheaf $\mathcal{G}$ on X such that that

$$\mathcal{G}(U) = \text{free abelian group generated by } \mathcal{F}(U)$$

etc., for open U in X. In other words, $\mathcal{G}(U)$ is the set of formal linear combinations (with coefficients in $\mathbb{Z}$) of continuous functions from X to Y. It turned out that $\mathcal{G}$ is never a sheaf, and for many reasons. The stalk $\mathcal{G}_z$ at $z \in X$ can be described (after some unraveling) as the set of formal linear combinations, with integer coefficients, of germs of continuous maps from (X, z) to Y. (Recall that *germ of continuous maps from* (X, z) *to* Y means an equivalence class of pairs (U, f) where U is an open neighborhood of z in X and $f\colon U \to Y$ is continuous.) Of course, we ask what $\mathcal{G}_z$ is because it feeds into the construction of $\Phi\mathcal{G}$, the sheafification of $\mathcal{G}$. It is permitted and even exciting to evaluate $\Phi\mathcal{G}$ on X, since X is an open subset of X.

Definition 4.3.1. An element of $(\Phi\mathcal{G})(X)$ will be called a *mapping cycle* from X to Y.

So what is a mapping cycle from X to Y ?

First answer. A mapping cycle from X to Y is a function s which for every $z \in X$ selects $s(z)$, a formal linear combination with integer coefficients of germs[1] of continuous maps from (X, z) to Y. There is a coherence condition to be satisfied: it must be possible to cover X by open sets W_i such that all values $s(z)$, where $z \in W_i$, can be simultaneously represented by one formal linear combination

$$\sum_j b_{ij} f_{ij}$$

where $f_{ij}\colon W_i \to Y$ are continuous maps and the b_{ij} are integers.

Second answer. A mapping cycle from X to Y can be specified (described, constructed) by choosing an open cover $(U_i)_{i\in\Lambda}$ of X and for every $i \in \Lambda$ a formal linear combination s_i with integer coefficients of continuous maps[2] from U_i to Y. There is a matching condition to be satisfied[3]: for any $i, j \in \Lambda$ and any $x \in U_i \cap U_j$, there should exist an open neighborhood W of x in $U_i \cap U_j$ such that $s_{i|W} = s_{j|W}$.

(The second answer is in some ways less satisfactory than the first because it does not say explicitly what a mapping cycle *is*, only how we can construct mapping cycles. But it can indeed be useful when we need to construct mapping cycles.)

Some of the "counter"examples which we saw previously now serve as illustrations of the concept of mapping cycle.

[1] Grown-up formulation: selects an element in the free abelian group generated by the set of germs ...

[2] Grown-up formulation: for every $i \in \Lambda$ an element s_i in the free abelian group generated by the set of continuous maps ...

[3] Did you expect to see the condition $s_{i|U_i\cap U_j} = s_{j|U_i\cap U_j}$? Sheaf theory dictates a weaker condition!

EXAMPLE 4.3.2. If S is a set with 6 elements and T is a set with 2 elements, both to be viewed as topological spaces with the discrete topology, then the abelian group of mapping cycles from S to T is isomorphic to $\mathbb{Z}^{12} \cong \prod_{i=1}^{6}(\mathbb{Z} \oplus \mathbb{Z})$. Do not confuse with $\mathbb{Z}/12$.

EXAMPLE 4.3.3. Let X and Y be two topological spaces related by a covering map

$$p\colon Y \to X$$

with finite fibers. In other words, p is a fiber bundle whose fibers are finite sets (viewed as topological spaces with the discrete topology). For simplicity, suppose also that X is connected. Choose an open covering $(W_j)_{j\in\Lambda}$ of X such that p admits a bundle chart over W_j for each j:

$$h_j\colon p^{-1}(W_j) \to W_j \times F$$

where F is a finite set (with the discrete topology). For $j \in \Lambda$ and $z \in F$ there is a continuous map $\sigma_{j,z}\colon W_j \to Y$ given by $\sigma_{j,z}(x) = h_j^{-1}(x,z)$ for $x \in W_j$. Define

$$s_j = \sum_{z\in F} \sigma_{j,z}\,.$$

This is a formal linear combination of continuous maps from W_j to Y. Clearly

$$s_{i|W_i\cap W_j} = s_{j|W_i\cap W_j}$$

(yes, this is more than we require). Therefore, by *"second answer"*, we have specified a mapping cycle from X to Y (which agrees with s_j on W_j).

EXAMPLE 4.3.4. Let X and Y be topological spaces. Suppose that $X = V_1 \cup V_2$ where V_1 and V_2 are open subsets of X. Let continuous maps $f, g\colon V_1 \to Y$ be given such that

$$f_{|V_1\cap V_2} = g_{|V_1\cap V_2}\,.$$

Then it makes (some) sense to view the formal linear combination $f - g = 1\cdot f + (-1)\cdot g$ as a mapping cycle from X to Y. How? We have the open cover of X consisting of V_1 and V_2 , and we specify $s_1 = f - g$ (a mapping cycle from V_1 to Y), and $s_2 = 0$ (a mapping cycle from V_2 to Y). Then $s_{1|V_1\cap V_2} = 0 = s_{2|V_1\cap V_2}$. So the matching condition is satisfied, and so by *"second answer"* we have specified a mapping cycle from X to Y.

Mapping cycles are complicated entities, but I hope that readers having survived the excursion into sheaf theory remain sufficiently intoxicated to find the definition obvious and unavoidable. With that, the excursion into sheaf theory is over (for now). Next we shall try to develop a comfortable relationship with mapping cycles. Here is a list of some of their good uses and properties.

(1) Every continuous map from X to Y determines a mapping cycle from X to Y.
(2) The mapping cycles from X to Y form an abelian group.
(3) A mapping cycle from X to Y can be composed with a (continuous) map from Y to Z to give a mapping cycle from X to Z. A mapping cycle from Y to Z can be composed with a (continuous) map from X to Y to give a mapping cycle from X to Z. But more remarkably, a mapping cycle from X to Y can be composed with a mapping cycle from Y to Z to give a mapping cycle from X to Z.
(4) Composition of mapping cycles is bilinear.
(5) Mapping cycles satisfy a sheaf property: if $(U_i)_{i\in\Lambda}$ is an open covering of X and $s_i\colon U_i \to Y$ is a mapping cycle, for each $i \in \Lambda$, such that

$$s_{i|U_i\cap U_j} = s_{j|U_i\cap U_j}$$

for all $i, j \in \Lambda$, then there is a unique mapping cycle s from X to Y such that $s_{|U_i} = s_i$ for all $i \in \Lambda$.

(6) There is exactly one mapping cycle from X to $\emptyset$. And there is exactly one mapping cycle from $\emptyset$ to Y, for any space Y.

(7) Mapping cycles from a topological disjoint union $X_1 \coprod X_2$ to Y are in bijection with pairs (s_1, s_2) where s_i is a mapping cycle from X_i to Y for $i = 1, 2$. Mapping cycles from X to a topological disjoint union $Y_1 \coprod Y_2$ are in bijection with pairs (s_1, s_2) where s_i is a mapping cycle from X to Y_i for $i = 1, 2$.

Some comments on that.
(1) A continuous map $f\colon X \to Y$ determines a mapping cycle $s = s_f$ where $s(z)$ is the germ of f at z. Interesting observation: the map $f \mapsto s_f$ from the set of continuous maps from X to Y to the set of mapping cycles from X to Y is injective.
(2) Obvious.
(3) Given a mapping cycle s from X to Y and a continuous map $g\colon Y \to Z$ we get a mapping cycle $g \circ s$ from X to Z by $x \mapsto \sum b_j(g \circ f_j)$ when $x \in X$ and $s(x) = \sum b_j f_j$. Given a mapping cycle s from Y to Z and a continuous map $g\colon X \to Y$ we get a mapping cycle $s \circ g$ from X to Z by $x \mapsto \sum b_j(f_j \circ g)$ when $x \in X$ and $s(x) = \sum b_j f_j$. Given a mapping cycle s from X to Y and a mapping cycle t from Y to Z we get a mapping cycle $t \circ s$ from X to Z by the formula

$$x \mapsto \sum (b_j c_{ij})(f_{ij} \circ g_j)$$

when $x \in X$ and $s(x) = \sum_j b_j g_j$ and $t(g_j(x)) = \sum_i c_{ij} f_{ij}$. (The notation is not fantastically precise or logical; in any case b_j, c_{ij} etc. are meant to be integers while f_{ij} , g_j etc. are meant to be germs of continuous functions. Note that f_{ij} in the displayed formula is a germ at $g_j(x)$, while g_j is a germ at x.)
(4) Should be clear from the last formula in the comment on (3).
(5) By construction.
(6) Mapping cycles from $\emptyset$ to Y: there is exactly one by construction. A mapping cycle s from X to $\emptyset$ is a function which for each $x \in X$ selects a formal linear combination of germs of continuous maps from (X, x) to $\emptyset$, etc.; since there no such germs, the only possible formal linear combination is the zero linear combination. This does satisfy the coherence condition.
(7) By construction and by inspection.

In category language, we can say that there is a category $\mathcal{A}\mathcal{T}\mathrm{op}$ whose objects are the topological spaces and where a morphism from space X to space Y is a mapping cycle from X to Y. There is an "inclusion" functor

$$\mathcal{T}\mathrm{op} \to \mathcal{A}\mathcal{T}\mathrm{op}$$

taking every object X to the same object X, and taking a morphism $f\colon X \to Y$ (continuous map) to the corresponding mapping cycle as explained in (1). For each X and Y, the set $\mathrm{mor}_{\mathcal{A}\mathcal{T}\mathrm{op}}(X, Y)$ is equipped with the structure of an abelian group. Composition of morphisms is bilinear. There is a zero object X in $\mathcal{A}\mathcal{T}\mathrm{op}$, i.e., an object with the property that $\mathrm{mor}_{\mathcal{A}\mathcal{T}\mathrm{op}}(X, Y)$ has exactly one element and $\mathrm{mor}_{\mathcal{A}\mathcal{T}\mathrm{op}}(Y, X)$ has exactly one element for arbitrary Y. Indeed, $X = \emptyset$ is a zero object in $\mathcal{A}\mathcal{T}\mathrm{op}$. The property expressed in (7) can also be formulated in category language, but we must postpone it because the vocabulary for that has not been introduced so far. In all, we can say that $\mathcal{A}\mathcal{T}\mathrm{op}$ is an *additive category*.

CHAPTER 5

Homotopies in $\mathcal{A}\mathcal{T}\mathrm{op}$

5.1. The homotopy relation

DEFINITION 5.1.1. Let X and Y be topological spaces. We call two mapping cycles f and g from X to Y *homotopic* if there exists a mapping cycle h from $X \times [0,1]$ to Y such that $f = h \circ \iota_0$ and $g = h \circ \iota_0$. Here $\iota_0, \iota_1 \colon X \to X \times [0,1]$ are defined by $\iota_0(x) = (x, 0)$ and $\iota_1(x) = (x, 1)$ as usual. Such a mapping cycle h is a *homotopy* from f to g.

Remark. In that definition, $X \times [0,1]$ still means the product of X and $[0,1]$ in $\mathcal{T}\mathrm{op}$. We saw some evidence suggesting that in $\mathcal{A}\mathcal{T}\mathrm{op}$ this does not have the properties that we might expect from a product (in a category sense).

LEMMA 5.1.2. *"Homotopic" is an equivalence relation on the set of mapping cycles from* X *to* Y. The set of equivalence classes will be denoted by $[[X, Y]]$ and the equivalence class of a mapping cycle f will be denoted by $[[f]]$.

PROOF. Reflexivity and symmetry are fairly obvious. Transitivity is more interesting. (I am indebted to S. Mahanta for the following pretty argument.) Let h be a homotopy from e to f and k a homotopy from f to g, where e, f and g are mapping cycles from X to Y. We can agree that it suffices to produce a mapping cycle ℓ from $X \times [0,2]$ to Y such that ℓ restricted to $X \times \{0\}$ agrees with e and ℓ restricted to $X \times \{1\}$ agrees with g. Let

$$u \colon X \times [0,2] \longrightarrow X \times [0,1], \qquad v \colon X \times [0,2] \longrightarrow X \times [0,1], \qquad p \colon X \times [0,2] \to X$$

be the continuous maps given by $u(x,t) \mapsto (x, \min\{t, 1\})$, $v(x,t) = (x, \max\{t, 1\})$ and $p(x,t) = x$. Put

$$\ell := u^*h + v^*k - p^*f.$$

For that we can also write $\ell = (h \circ u) + (k \circ v) - (f \circ p)$. □

PROPOSITION 5.1.3. *The set* $[[X, Y]]$ *is an abelian group.*

PROOF. This amounts to observing that the homotopy relation is compatible with addition of mapping cycles. In other words, if f is homotopic to g and u is homotopic to v, where f, g, u, v are mapping cycles from X to Y, then $f + u$ is homotopic to $g + v$. Indeed, if h is a homotopy from f to g and k is a homotopy from u to v, then $h + k$ is a homotopy from $f + u$ to $g + v$. □

LEMMA 5.1.4. *A composition map* $[[Y, Z]] \times [[X, Y]] \to [[X, Z]]$ *can be defined by* $([[f]], [[g]]) \mapsto [[f \circ g]]$. *Composition is bilinear, i.e., for fixed* $[[g]]$ *the map* $[[f]] \mapsto [[f \circ g]]$ *is a homomorphism of abelian groups and for fixed* $[[f]]$ *the map* $[[g]] \mapsto [[f \circ g]]$ *is a homomorphism of abelian groups.* □

As a result there is a homotopy category $\mathcal{H}\mathrm{o}\mathcal{A}\mathcal{T}\mathrm{op}$ whose objects are (still) the topological spaces, while the set of morphisms from X to Y is $[[X, Y]]$.

5.2. First calculations

Write $\star$ for a singleton, alias one-point space.

PROPOSITION 5.2.1. *For any space* X *the abelian group* $[[X,\star]]$ *is isomorphic to the set of continuous (=locally constant) functions from* X *to* $\mathbb{Z}$*, where* $\mathbb{Z}$ *has the discrete topology.*

PROOF. We learned in example 4.2.2 that the set of mapping cycles from X to $\star$ is identified with the set of continuous functions from X to $\mathbb{Z}$. (It is $(\Phi\mathcal{G})(X)$ where $\Phi\mathcal{G}$ is the sheaf associated to the constant presheaf $\mathcal{G}$ which has $\mathcal{G}(U) = \mathbb{Z}$ for all open $U \subset X$.) Similarly, the set of mapping cycles from $X \times [0,1]$ to $\star$ is identified with the set of continuous functions from $X \times [0,1]$ to $\mathbb{Z}$. But a continuous function h from $X \times [0,1]$ to $\mathbb{Z}$ is constant on $\{x\} \times [0,1]$ for each $x \in X$, and so will have the form $h(x,t) = g(x)$ for a unique continuous $g \colon X \to \mathbb{Z}$. It follows that the homotopy relation on the set of mapping cycles from X to $\star$ is trivial, i.e., two mapping cycles from X to $*$ are homotopic only if they are equal. $\square$

EXAMPLE 5.2.2. Take $X = \mathbb{Q}$, a subspace of $\mathbb{R}$ with the standard topology. The group $[[\mathbb{Q},\star]]$ is uncountable because the set of continuous maps from $\mathbb{Q}$ to $\mathbb{Z}$ is uncountable.

LEMMA 5.2.3. *For a path-connected (non-empty) space* Y *the abelian group* $[[\star, Y]]$ *is isomorphic to* $\mathbb{Z}$*.*

PROOF. Fix some point $z \in Y$. A mapping cycle from $\star$ to Y is the same thing as a formal linear combination of points in Y, say $\sum_j b_j y_j$ where $b_j \in \mathbb{Z}$ and $y_j \in Y$. In the abelian group $[[\star, Y]]$ we have

$$[[\Sigma_j b_j y_j]] = \Sigma_j b_j [[y_j]] = (\Sigma_j b_j)[[z]].$$

(Here $[[y_j]]$ for example denotes the homotopy class of the mapping cycle determined by the continuous map $\star \to Y$ which has image $\{y_j\}$. As that continuous map is homotopic to the map $\star \to Y$ which has image $\{z\}$, we obtain $[[y_j]] = [[z]]$.) Therefore $[[\star, Y]]$ is cyclic, generated by the element $[[z]]$. To see that it is infinite cyclic we use the homomorphism

$$[[\star, Y]] \to [[\star,\star]]$$

given by composition with the continuous map $Y \to \star$. Now $[[\star,\star]]$ is infinite cyclic by proposition 5.2.1. It is also clear that the homomorphism just above takes $[[z]]$ to the generator of $[[\star,\star]]$, the class of the identity mapping cycle. Hence it must be an isomorphism and so $[[\star, Y]]$ is infinite cyclic. $\square$

COROLLARY 5.2.4. *For any space* Y *the abelian group* $[[\star, Y]]$ *is isomorphic to the free abelian group generated by the set of path components of* Y*.*

PROOF. The abelian group of mapping cycles from $\star$ to Y is simply the free abelian group A generated by the underlying set of Y. Write this as a direct sum $\bigoplus_{\lambda\in\Lambda} A_\lambda$ where Λ is an indexing set for the path components Y_λ of Y and A_λ is the free abelian group generated by the underlying set of Y_λ. Now fix some λ. *Claim*: If $f \in A$ is homotopic to $g \in A$, by a mapping cycle $h\colon [0,1] \to Y$, then the coordinate of f in A_λ is homotopic to the coordinate of g in A_λ , by a mapping cycle $[0,1] \to Y_\lambda$. To see this, cover the interval $[0,1]$ by finitely many open subsets U_i such that $h_{|U_i}$ can be represented by a formal linear combination of continuous maps from U_i to Y. This is possible by the coherence condition on h. Choose a subdivision

$$0 = t_0 < t_1 < \cdots t_{N-1} < t_N = 1$$

of $[0,1]$ such that for each of the the subintervals $[t_r, t_{r+1}]$, where $r = 0,1,\ldots,N-1$, there exists U_i containing it. Let $h_{t_r} \in A$ be obtained by restricting h to t_r. Then $h_{t_0} = f$ and $h_{t_N} = g$, so it suffices to show that the λ-coordinate of h_{t_r} is homotopic to the λ-coordinate of $h_{t_{r+1}}$, for $r = 0,1,\ldots,N-1$. But $[t_r, t_{r+1}]$ is contained in some U_i and so there is a formal linear combination

$$\sum_j b_j u_j$$

where $b_j \in \mathbb{Z}$ and the u_j are continuous maps from $[t_r, t_{r+1}]$ to Y, and $\sum_j b_j u_j$ restricts to h_{t_r} on t_r and to $h_{t_{r+1}}$ on t_{r+1}. Each u_j maps to only one path component of Y; in the formal linear combination $\sum_j b_j u_j$, select the terms $b_j u_j$ where u_j is a map to Y_λ and discard the others. Then the selected linear sub-combination is a homotopy from the λ-component of h_{t_r} to the λ-component of $h_{t_{r+1}}$. This proves the claim.
Therefore $[[\star, Y]]$ is the direct sum of the $[[\star, Y_\lambda]]$. By the lemma above, each $[[\star, Y_\lambda]]$ is isomorphic to $\mathbb{Z}$. $\square$

PROPOSITION 5.2.5. *For topological spaces X and Y where X is a topological disjoint union $X_1 \amalg X_2$, there is an isomorphism*

$$[[X,Y]] \longrightarrow [[X_1,Y]] \times [[X_2,Y]]\ ;\ [[f]] \mapsto ([[f_{|X_1}]], [[f_{|X_2}]])\ .$$

For topological spaces X and Y where Y is a topological disjoint union $Y_1 \amalg Y_2$, there is an isomorphism

$$[[X,Y_1]] \oplus [[X,Y_2]] \longrightarrow [[X,Y]]\ ; [[f]] \oplus [[g]] \mapsto [[j_1 \circ f + j_2 \circ g]]$$

where $j_1 \colon Y_1 \to Y$ and $j_2 \colon Y_2 \to Y$ are the inclusions.

PROOF. First statement: the set $\mathrm{mor}_{\mathcal{A}\mathcal{T}\mathrm{op}}(X,Y)$ of mapping cycles breaks up as a product $\mathrm{mor}_{\mathcal{A}\mathcal{T}\mathrm{op}}(X_1,Y) \times \mathrm{mor}_{\mathcal{A}\mathcal{T}\mathrm{op}}(X_2,Y)$ by restriction to X_1 and X_2, and a similar statement holds for the set $\mathrm{mor}_{\mathcal{A}\mathcal{T}\mathrm{op}}(X \times [0,1], Y)$. Second statement: the set $\mathrm{mor}_{\mathcal{A}\mathcal{T}\mathrm{op}}(X,Y)$ of mapping cycles breaks up as a direct sum $\mathrm{mor}_{\mathcal{A}\mathcal{T}\mathrm{op}}(X,Y_1) \times \mathrm{mor}_{\mathcal{A}\mathcal{T}\mathrm{op}}(X,Y_2)$, and a similar statement holds for $\mathrm{mor}_{\mathcal{A}\mathcal{T}\mathrm{op}}(X \times [0,1], Y)$. $\square$

PROPOSITION 5.2.6. *For any topological space X we have*

$$[[\emptyset, X]] = 0 = [[X, \emptyset]]\ .$$

PROOF. The abelian group of mapping cycles from X to $\emptyset$ is a trivial group and the abelian group of mapping cycles from $\emptyset$ to X is a trivial group. $\square$

5.3. Homology and cohomology: the definitions

DEFINITION 5.3.1. For $n \geq 0$, the n-th *homology group* of a topological space X is the abelian group

$$H_n(X) := [[S^n, X]]/[[\star, X]]\ .$$

The n-th *cohomology group* of X is the abelian group

$$H^n(X) := [[X, S^n]]/[[X, \star]].$$

Comments. There is an understanding here that $[[\star, X]]$ is a subgroup of $[[S^n, X]]$. How? By pre-composing mapping cycles from $\star$ to X with the unique continuous map from S^n to $\star$, we obtain a (well defined) homomorphism $[[\star, X]] \to [[S^n, X]]$. Conversely, by pre-composing mapping cycles from S^n to X with a selected continuous map from $\star$ to S^n, inclusion of the base point, we obtain a homomorphism $[[S^n, X]] \to [[\star, X]]$. The composition $[[\star, X]] \to [[S^n, X]] \to [[\star, X]]$ is the identity on $[[\star, X]]$. So we can say that $[[\star, X]]$

is a direct summand of $[[S^n, X]]$. We remove it, suppress it etc., when we form $H_n(X)$. Similarly, by post-composing mapping cycles from X to S^n with the unique continuous map $S^n \to \star$, we obtain a homomorphism $[[X, S^n]] \to [[X, \star]]$. Conversely, by post-composing mapping cycles from X to $\star$ with a selected continuous map $\star \to S^n$, inclusion of the base point, we obtain a homomorphism $[[X, \star]] \to [[X, S^n]]$. The composition $[[X, \star]] \to [[X, S^n]] \to [[X, \star]]$ is the identity on $[[X, \star]]$. Therefore $[[X, \star]]$ is a direct summand of $[[X, S^n]]$. We remove it, suppress it etc., when we form $H^n(X)$.

You will be unsurprised to hear that H_n is a functor from $\mathcal{T}$op to the category of abelian groups. We can also say that it is a functor from $\mathcal{A}\mathcal{T}$op to abelian groups. Both statements are obvious from the definition. Equally clear from the definition, but important to keep in mind: if $f, g\colon X \to Y$ are homotopic maps, then the induced homomorphisms f_* and g_* from $H_n(X)$ to $H_n(Y)$ are the same. (Therefore we might say that H_n is a functor from $\mathcal{H}$o$\mathcal{T}$op to the category of abelian groups. Indeed it is a functor from $\mathcal{H}$o$\mathcal{A}\mathcal{T}$op to abelian groups ...)
Similarly H^n is a contravariant functor from $\mathcal{T}$op (or from $\mathcal{A}\mathcal{T}$op, or from $\mathcal{H}$o$\mathcal{T}$op, or from $\mathcal{H}$o$\mathcal{A}\mathcal{T}$op) to the category of abelian groups.

So far we have few tools available for computing $H_n(X)$ and $H^n(X)$ in general. But in the cases $n = 0$, arbitrary X, we are ready for it, and in the case where n is arbitrary and $X = \star$ we are also ready for it.

Example 5.3.2. Take $n = 0$ and X arbitrary. Then $H_0(X) = [[S^0, X]]/[[\star, X]]$. For S^0 we can write $\star \amalg \star$ (disjoint union of two copies of $\star$), and using the first part of proposition 5.2.5, we get $[[S^0, X]] \cong [[\star, X]] \times [[\star, X]]$. Therefore $H_0(X) \cong [[\star, X]]$. Using corollary 5.2.4, it follows that $H_0(X)$ is identified with the free abelian group generated by the set of path components of X. For example, if X is path connected, then $H_0(X)$ is isomorphic to $\mathbb{Z}$.
By a very similar calculation, $H^0(X)$ is isomorphic to $[[X, \star]]$. Using proposition 5.2.1, we then obtain that $H^0(X)$ is isomorphic to the abelian group of continuous maps from X to $\mathbb{Z}$. For example, if X is connected, then $H^0(X)$ is isomorphic to $\mathbb{Z}$.

Example 5.3.3. Take n arbitrary and $X = \star$. Now $H_n(\star) = [[S^n, \star]]/[[\star, \star]]$. Using proposition 5.2.1, we find $[[S^n, \star]] \cong \mathbb{Z}$ when $n > 0$ and $[[S^0, \star]] \cong \mathbb{Z} \oplus \mathbb{Z}$; also $[[\star, \star]] = \mathbb{Z}$. By an easy calculation, the quotient $[[S^n, \star]]/[[\star, \star]]$ is therefore 0 when $n > 0$, and isomorphic to $\mathbb{Z}$ when $n = 0$. So we have:

$$H_n(\star) \quad \cong \quad \begin{cases} \mathbb{Z} & \text{if } n = 0 \\ 0 & \text{if } n > 0 \end{cases}$$

Similarly, $H^n(\star) = [[\star, S^n]]/[[\star, \star]]$. Using corollary 5.2.4 this time, we find that $[[\star, S^n]] \cong \mathbb{Z}$ when $n > 0$ and $[[\star, S^0]] \cong \mathbb{Z} \oplus \mathbb{Z}$. By an easy calculation, the quotient $[[\star, S^n]]/[[\star, \star]]$ is therefore 0 when $n > 0$, and isomorphic to $\mathbb{Z}$ when $n = 0$. Therefore:

$$H^n(\star) \quad \cong \quad \begin{cases} \mathbb{Z} & \text{if } n = 0 \\ 0 & \text{if } n > 0 \end{cases}$$

CHAPTER 6

The homotopy decomposition theorem and the Mayer-Vietoris sequence

6.1. The homotopy decomposition theorem

Notation for the following theorem and the corollary: X and Y are topological spaces, V and W are open subsets of Y such that $V \cup W = Y$, and C is a closed subset of X. We assume that X is *paracompact.*

THEOREM 6.1.1. *Let* $\gamma : X \times [0,1] \to Y$ *be a mapping cycle which restricts to zero on an open neighborhood of* $X \times \{0\}$*. Then there exists a decomposition*

$$\gamma = \gamma^V + \gamma^W \ ,$$

where $\gamma^V \colon X \times [0,1] \to V$ *and* $\gamma^W \colon X \times [0,1] \to W$ *are mapping cycles, both zero on an open neighborhood of* $X \times \{0\}$*. If* γ *is zero on some neighborhood of* $C \times [0,1]$*, then it can be arranged that* γ_V *and* γ_W *are zero on a neighborhood of* $C \times [0,1]$*.*

The proof of this is hard. We postpone it.

COROLLARY 6.1.2. *Let* $a \in [[X,V]]$ *and* $b \in [[X,W]]$ *be such that the images of* a *and* b *in* $[[X,Y]]$ *agree. Then there exists* $c \in [[X, V \cap W]]$ *whose image in* $[[X,V]]$ *is* a *and whose image in* $[[X,W]]$ *is* b*.*

PROOF. Let α be a mapping cycle which represents a and let β be a mapping cycle which represents b. Choose a mapping cycle $\gamma \colon X \times [0,1] \to Y$ which is a homotopy from 0 to $\beta - \alpha$. It is easy to arrange this in such a way that γ is zero on a neighborhood of $X \times \{0\}$. Use the theorem to obtain a decomposition $\gamma = \gamma^V + \gamma^W$. Let γ_1^V and γ_1^W be the restrictions of γ^V and γ^W to $X \times \{1\}$. Then α and $\alpha + \gamma_1^V$ are homotopic as mapping cycles $X \to V$, by the homotopy $\alpha \circ p + \gamma^V$, where p is the projection $X \times [0,1] \to X$. Similarly $\beta = \alpha + \gamma_1^V + \gamma_1^W$ and $\alpha + \gamma_1^V$ are homotopic as mapping cycles $X \to W$. Finally, $\alpha + \gamma_1^V = \beta - \gamma_1^W$ lands in $V \cap W$ by construction. So $c = [[\alpha + \gamma_1^V]]$ is a solution. □

REMARK 6.1.3. The corollary is in a formal way very reminiscent of proposition 2.5.5. However the assumptions there were somewhat different. Instead of a union-intersection square of spaces serving as targets, we had a pullback square and a fibration condition. We can ask whether that was necessary or appropriate. Does corollary 6.1.2 have a more direct analogue in $\mathcal{H}o\mathcal{T}op$? In other words, given spaces X and $Y = V \cup W$ as in corollary 6.1.2, and elements $a \in [X,V]$ and $b \in [X,W]$ such that the images of a and b in $[X,Y]$ agree, does there exist $c \in [X, V \cap W]$ whose image in $[X,V]$ is a and whose image in $[X,W]$ is b ? Interestingly the answer is no in general. A relatively easy counterexample (easier for you if you know the concept *fundamental group*) can be constructed as follows. Let $p, q \in \mathbb{R}^2$, $p = (0,1)$ and $q = (0,-1)$. Let $Y = \mathbb{R}^2 \smallsetminus \{q\}$, $V = \mathbb{R}^2 \smallsetminus \{p,q\}$ and W the open upper half-plane. Then $V \cap W = W \smallsetminus \{p\}$. For X take S^1. It is rather easy to invent $a \in [X, V]$ which maps to the zero element in $[X, Y]$, but which does not come from $[X, V \cap W]$. Therefore if we take $b \in [X, W]$ to be the class of the constant map, we have a "situation". Picture of a map in the homotopy class a:

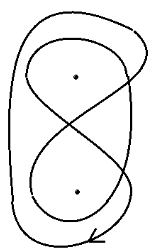

There are also deeper counterexamples where $X = S^n$ for some $n > 1$. For those we need to work harder.

6.2. The Mayer-Vietoris sequence in homology

A sequence of abelian groups $(A_n)_{n\in\mathbb{Z}}$ together with homomorphisms

$$f_n \colon A_n \to A_{n-1}$$

for all $n \in \mathbb{Z}$ is called an *exact sequence of abelian groups* if the kernel of f_n is equal to the image of f_{n+1} , for all $n \in \mathbb{Z}$. More generally, we sometimes have to deal with diagrams of abelian groups and homomorphisms in the shape of a string

$$A_n \to A_{n-1} \to A_{n-2} \to \cdots \to A_{n-k} \ .$$

Such a diagram is *exact* if the kernel of each homomorphism in the string is equal to the image of the preceding one, if there is a preceding one.

DEFINITION 6.2.1. *Alternative definition of homology*: Write $I = [0, 1]$. For a space Y, and $n \geq 0$, re-define $H_n(Y)$ as the abelian group of homotopy classes of mapping cycles $I^n \to Y$ which vanish on some open neighborhood of ∂I^n.

Comment. In this definition, we regard two mapping cycles $I^n \to Y$ which vanish on some neighborhood of ∂I^n as *homotopic* if they are related by a homotopy $I^n \times I \to Y$ which vanishes on some neighborhood of $\partial I^n \times I$. Such a homotopy will be called (informally) a *homotopy relative to* ∂I^n or a *homotopy rel* ∂I^n.
To relate the old definition of $H_n(Y)$ to the new one, we make a few observations. Given a mapping cycle $\alpha \colon I^n \to Y$ which vanishes on some neighborhood of ∂I^n, we immediately obtain a mapping cycle from the quotient $I^n/\partial I^n$ to Y. To view this as a mapping cycle $\beta \colon S^n \to Y$, we pretend that $S^n = \mathbb{R}^n \cup \infty$ (one-point compactification of $\mathbb{R}^n$) and specify a homeomorphism $u \colon I^n/\partial I^n \to \mathbb{R}^n \cup \infty$ taking base point to base point. (Note that $I^n/\partial I^n$ has a preferred base point, the point represented by all elements of ∂I^n.) We are specific enough if we say that u is smooth and orientation preserving on $I^n \smallsetminus \partial I^n$ (i.e., the Jacobian determinant is everywhere positive). Conversely, given a mapping cycle $\beta \colon S^n \to Y$ representing an element of $H_n(Y)$ according to the old definition, we may subtract a suitable constant to arrange that β is zero when restricted to the base point of S^n. We can also assume that β is zero on a neighborhood of the base point; if not,

compose with a continuous map $S^n \to S^n$ which is homotopic to the identity and takes a neighborhood of the base point to the base point. Then $\beta \circ u$ is a mapping cycle $I^n/\partial \to Y$ which can also be viewed as a mapping cycle $I^n \to Y$ vanishing on a neighborhood of ∂I^n.

DEFINITION 6.2.2. Suppose that Y comes with two open subspaces V and W such that $V \cup W = Y$. The *boundary homomorphism*

$$\partial : H_n(Y) \to H_{n-1}(V \cap W)$$

is defined as follows, using the alternative definition of H_n. Let $x \in H_n(Y)$ be represented by a mapping cycle $\gamma : I^n \to Y$ which is zero on an open neighborhood of ∂I^n. Think of γ as a homotopy, $\gamma : I^{n-1} \times I \to Y$. Choose a decomposition $\gamma = \gamma^V + \gamma^W$ as in theorem 6.1.1. The theorem guarantees that γ^V and γ^W can be arranged to vanish on a neighborhood of $\partial I^{n-1} \times I$. Let $\partial(x)$ be the class of the mapping cycle

$$\gamma_1^W : I^{n-1} \to V \cap W ,$$

composition of γ^W with the map $\iota_1 : I^{n-1} \to I^{n-1} \times I$ defined by $\iota_1(x) = (x, 1)$. (Then γ_1^W vanishes on some open neighborhood of ∂I^{n-1}.)

We must show that this is well defined. There were two choices involved: the choice of representative γ, and the choice of decomposition $\gamma = \gamma^V + \gamma^W$. For the moment, keep γ fixed, and let us see what happens if we try another decomposition of γ. Any other decomposition will have the form

$$(\gamma^V + \eta) + (\gamma^W - \eta)$$

where $\eta : I^{n-1} \times I \to V \cap W$ is a mapping cycle which vanishes on an open neighborhood of $\partial I^{n-1} \times I$ and on an open neighborhood of $I^{n-1} \times \{0\}$. We need to show that $\gamma_1^W - \eta_1$ is homotopic (rel boundary of I^{n-1}) to γ_1^W. But η_1 is homotopic to zero by the homotopy η.
Next we worry about the choice of representative γ. Let φ be another representative of the same class x, and let $\lambda : I \times I^n \to Y$ be a homotopy from φ to γ. (Writing the factor I on the left might help us to avoid confusion.) We can think of λ as a homotopy in a different way:

$$(I \times I^{n-1}) \times I \longrightarrow Y.$$

Then we can apply the homotopy decomposition theorem and choose a decomposition $\lambda = \lambda^V + \lambda^W$ where λ^V and λ^W vanish on a neighborhood of $I \times \partial I^{n-1} \times I$. We then find that λ_1^W is a mapping cycle from $X = I \times I^{n-1}$ to $V \cap W$ which we may regard as a homotopy (now with parameters written on the left). The homotopy is between γ_1^W and φ_1^W, provided the decompositions $\gamma = \gamma^V + \gamma^W$ and $\varphi = \varphi^V + \varphi^W$ are the ones obtained by restricting the decomposition $\lambda = \lambda^V + \lambda^W$. $\square$

The boundary homomorphisms ∂ can be used to make a sequence of abelian groups and homomorphisms

$$\begin{tikzcd}
\cdots \arrow[r] & H_{n+1}(Y) \arrow[d, "\partial"] & & \\
& H_n(V \cap W) \arrow[r] & H_n(V) \oplus H_n(W) \arrow[r] & H_n(Y) \arrow[d, "\partial"] \\
& & & H_{n-1}(V \cap W) \arrow[r] & \cdots
\end{tikzcd}$$

where $n \in \mathbb{Z}$. (Set $H_n(X) = 0$ for $n < 0$ and any space X. The unlabelled homomorphisms in the sequence are as follows: $H_n(V) \oplus H_n(W) \to H_n(Y)$ is $j_{V*} + j_{W*}$, the sum of the two maps given by composition with the inclusions $j_V \colon V \to Y$ and $j_W \colon W \to Y$, and $H_n(V \cap W) \to H_n(V) \oplus H_n(W)$ is $(e_{V*}, -e_{W*})$, where e_{V*} and e_{W*} are given by composition with the inclusions $e_V \colon V \cap W \to V$ and $e_W \colon V \cap W \to W$.) The sequence is called the homology *Mayer-Vietoris* sequence of Y and V, W.

THEOREM 6.2.3. *The homology Mayer-Vietoris sequence of Y and V, W is exact.*[1]

Terminology for the proof. Let X and Q be topological spaces and let $h \colon X \times I \to Q$ be a map or mapping cycle (which we think of as a homotopy). Let $p \colon X \times I \to X$ be the projection and let $\iota_0, \iota_1 \colon X \to X \times I$ be the maps given by $x \mapsto (x, 0)$ and $x \mapsto (x, 1)$, respectively. We say that h is *stationary* near $X \times \{0, 1\}$ if there exist open neighborhoods U_0 and U_1 of $X \times \{0\}$ and $X \times \{1\}$, respectively, in $X \times I$ such that h agrees with $h \circ \iota_0 \circ p$ on U_0 and with $h \circ \iota_1 \circ p$ on U_1.

PROOF. (i) Exactness of the pieces $H_n(V \cap W) \to H_n(V) \oplus H_n(W) \to H_n(Y)$ follows from corollary 6.1.2, for all $n \in \mathbb{Z}$. (It is more convenient to use the standard definition of H_n at this point.) More precisely, we have exactness of

$$[[S^n, V \cap W]] \to [[S^n, V]] \oplus [[S^n, W]] \to [[S^n, Y]]$$

by corollary 6.1.2, and we have exactness of

$$[[\star, V \cap W]] \to [[\star, V]] \oplus [[\star, W]] \to [[\star, Y]]$$

by corollary 6.1.2. Note also that $[[\star, V]] \oplus [[\star, W]] \to [[\star, Y]]$ is surjective. Then it follows easily that

$$\frac{[[S^n, V \cap W]]}{[[\star, V \cap W]]} \to \frac{[[S^n, V]] \oplus [[S^n, W]]}{[[\star, V]] \oplus [[\star, W]]} \to \frac{[[S^n, Y]]}{[[\star, Y]]}$$

is exact.

(ii) Next we look at pieces of the form

$$H_n(V) \oplus H_n(W) \longrightarrow H_n(Y) \xrightarrow{\partial} H_{n-1}(V \cap W) .$$

The cases $n < 0$ are trivial. In the case $n = 0$, the claim is that the homomorphism $H_0(V) \oplus H_0(W) \to H_0(Y)$ is surjective. This is a pleasant exercise. Now assume $n > 0$. It is clear from the definition of ∂ that the composition of the two homomorphisms is zero. Suppose then that $[\gamma] \in H_n(Y)$ is in the kernel of ∂, where $\gamma : I^n \to Y$ vanishes on a neighborhood of ∂I^n. We must show that $[\gamma]$ is in the image of $H_n(V) \oplus H_n(W) \to H_n(Y)$. As above, we think of γ as a homotopy, $I^{n-1} \times I \to Y$, which we decompose, $\gamma = \gamma^V + \gamma^W$ as in theorem 6.1.1, where γ^V and γ^W vanish on a neighborhood of $\partial I^{n-1} \times I$. We can also arrange that the homotopies γ^V and γ^W are stationary near $I^{n-1} \times \{0, 1\}$. The assumption $\partial[\gamma] = 0$ then means that the zero map

$$I^{n-1} \to V \cap W$$

is homotopic to γ_1^W by a homotopy $\lambda : I^{n-1} \times I \to V \cap W$ which vanishes on a neighborhood of $\partial I^{n-1} \times I$. We can arrange that λ is stationary near $I^{n-1} \times \{0, 1\}$. Then $\gamma^V + \lambda$ and $\gamma^W - \lambda$ are mapping cycles from $I^{n-1} \times I = I^n$ to V and W, respectively. Both vanish on a neighborhood of ∂I^n. Hence they represent elements in $H_n(V)$ and $H_n(W)$ whose

[1] If you wish, view this as a sequence of abelian groups and homomorphisms indexed by the integers, by setting for example $A_{3n} = H_n(Y)$ for $n \geq 0$, $A_{3n+1} = H_n(V) \oplus H_n(W)$ for $n \geq 0$, $A_{3n+2} = H_n(V \cap W)$ for $n \geq 0$, and $A_m = 0$ for all $m \leq 0$.

images in $H_n(Y)$ add up to $[\gamma]$.
(iii) We show that the composition

$$H_{n+1}(Y) \xrightarrow{\partial} H_n(V\cap W) \longrightarrow H_n(V) \oplus H_n(W) \text{ .}$$

is zero. We can assume $n \geq 0$. Represent an element in $H_n(Y)$ by a mapping cycle $\gamma : I^n \times I \to Y$, vanishing on a neighborhood of the entire boundary; decompose as usual, and obtain $\partial[\gamma] = [\gamma_1^W]$. Now $\gamma_1^W = -\gamma_1^V$ viewed as a mapping cycle $I^n \to V$ is homotopic to zero by the homotopy $-\gamma^V$ vanishing on a neighborhood of $\partial I^{n-1} \times I$. Therefore $\partial[\gamma]$ maps to zero in $H_n(V)$. A similar calculation shows that it maps to zero in $H_n(W)$.
(iv) Finally let $\varphi : I^n \to V \cap W$ be a mapping cycle which vanishes on a neighborhood of ∂I^n, and suppose that $[\varphi] \in H_n(V \cap W)$ is in the kernel of the homomorphism $H_n(V \cap W) \to H_n(V) \oplus H_n(W)$. Choose a homotopy $\gamma^V : I^n \times I \to V$ from zero to $-\varphi$, and another homotopy $\gamma^W : I^n \times I \to W$ from zero to φ, both vanishing on a neighborhood of $\partial I^n \times I$, and both stationary near $I^n \times \{0, 1\}$. Then $\gamma := \gamma^V + \gamma^W$ vanishes on the entire boundary of $I^n \times I$, hence represents a class $[\gamma] \in H_{n+1}(Y)$. It is clear that $\partial[\gamma] = [\varphi]$. $\square$

Remark 6.2.4. The Mayer-Vietoris sequence has a naturality property. The statement is complicated. Suppose that Y and Y' are topological spaces, $g\colon Y \to Y'$ is a continuous map, $Y = V \cup W$ where V and W are open subsets, $Y' = V' \cup W'$ where V" and W' are open subsets, $g(V) \subset V'$ and $g(W) \subset W'$. Then the Mayer-Vietoris sequences for Y, V, W and Y', V', W' can be arranged in a ladder-shaped diagram

$$\begin{tikzcd}
\vdots \arrow[d] & \vdots \arrow[d] \\
H_{n+1}(Y) \arrow[r, "g_*"] \arrow[d, "\partial"] & H_{n+1}(Y') \arrow[d, "\partial"] \\
H_n(V\cap W) \arrow[r, "g_*"] \arrow[d] & H_n(V'\cap W') \arrow[d] \\
H_n(V) \oplus H_n(W) \arrow[r, "g_*"] \arrow[d] & H_n(V') \oplus H_n(W') \arrow[d] \\
H_n(Y) \arrow[r, "g_*"] \arrow[d, "\partial"] & H_n(Y') \arrow[d, "\partial"] \\
H_{n-1}(V\cap W) \arrow[r, "g_*"] \arrow[d] & H_{n-1}(V'\cap W') \arrow[d] \\
\vdots & \vdots
\end{tikzcd}$$

This diagram is *commutative*; that is the naturality statement. The proof is not complicated (it is by inspection).
Often this can be usefully combined with the following observation: if, in the Mayer-Vietoris sequence for Y and V, W we interchange the roles (order) of V and W, then the

homomorphisms ∂ and $H_n(V \cap W) \to H_n(V) \oplus H_n(W)$ change sign. To be more precise, we set up a diagram

$$\begin{tikzcd} H_{n+1}(Y) \arrow[r, "="] \arrow[d, "\partial"] & H_{n+1}(Y) \arrow[d, "\partial"] \\ H_n(V \cap W) \arrow[r, "="] \arrow[d] & H_n(W \cap V) \arrow[d] \\ H_n(V) \oplus H_n(W) \arrow[r, "\cong"] & H_n(W) \oplus H_n(V) \end{tikzcd}$$

where the columns are bits from the Mayer-Vietoris sequence of Y, V, W and Y, W, V, respectively. The diagram is *not* (always) commutative; instead each of the small squares in it commutes up to a factor (-1). The proof is by inspection.

CHAPTER 7

Homology of spheres and applications

7.1. Homology of spheres

PROPOSITION 7.1.1. *The homology groups of* S^1 *are* $H_0(S^1) \cong \mathbb{Z}$, $H_1(S^1) \cong \mathbb{Z}$ *and* $H_k(S^1) = 0$ *for all* $k \neq 0, 1$.

PROOF. Choose two distinct points p and q in S^1. Let $V \subset S^1$ be the complement of p and let $W \subset S^1$ be the complement of q. Then $V \cup W = S^1$. Clearly V is homotopy equivalent to a point, W is homotopy equivalent to a point and $V \cap W$ is homotopy equivalent to a discrete space with two points. Therefore $H_k(V) \cong H_k(W) \cong \mathbb{Z}$ for $k = 0$ and $H_k(V) \cong H_k(W) = 0$ for all $k \neq 0$. Similarly $H_k(V \cap W) \cong \mathbb{Z} \oplus \mathbb{Z}$ for $k = 0$ and $H_k(V \cap W) = 0$ for all $k \neq 0$. The exactness of the Mayer-Vietoris sequence associated with the open covering of S^1 by V and W implies immediately that $H_k(S^1) = 0$ for $k \neq 0, 1$. The part of the Mayer-Vietoris sequence which remains interesting after this observation is

$$0 \longrightarrow H_1(S^1) \xrightarrow{\partial} \mathbb{Z} \oplus \mathbb{Z} \longrightarrow \mathbb{Z} \oplus \mathbb{Z} \longrightarrow H_0(S^1) \longrightarrow 0$$

Since S^1 is path-connected, the group $H_0(S^1)$ is isomorphic to $\mathbb{Z}$. The homomorphism from $\mathbb{Z} \oplus \mathbb{Z}$ to $H_0(S^1)$ is onto by exactness, so its kernel is isomorphic to $\mathbb{Z}$. Hence the image of the homomorphism $\mathbb{Z} \oplus \mathbb{Z} \to \mathbb{Z} \oplus \mathbb{Z}$ is isomorphic to $\mathbb{Z}$, so its kernel is again isomorphic to $\mathbb{Z}$. Now exactness at $H_1(S^1)$ leads to the conclusion that $H_1(S^1) \cong \mathbb{Z}$. □

THEOREM 7.1.2. *The homology groups of* S^n *(for* $n > 0$*) are*

$$H_k(S^n) \cong \begin{cases} \mathbb{Z} & \text{if } k = n \\ \mathbb{Z} & \text{if } k = 0 \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. We proceed by induction on n. The induction beginning is the case $n = 1$ which we have already dealt with separately in proposition 7.1.1. For the induction step, suppose that $n > 1$. We use the Mayer-Vietoris sequence for S^n and the open covering $\{V, W\}$ with $V = S^n \smallsetminus \{p\}$ and $W = S^n \smallsetminus \{q\}$ where $p, q \in S^n$ are the north and south pole, respectively. We will also use the homotopy invariance of homology. This gives us

$$H_k(V) \cong H_k(W) \cong \begin{cases} \mathbb{Z} & \text{if } k = 0 \\ 0 & \text{otherwise} \end{cases}$$

because V and W are homotopy equivalent to a point. Also we get

$$H_k(V \cap W) \cong \begin{cases} \mathbb{Z} & \text{if } k = n-1 \\ \mathbb{Z} & \text{if } k = 0 \\ 0 & \text{otherwise.} \end{cases}$$

by the induction hypothesis, since $V \cap W$ is homotopy equivalent to S^{n-1}. Furthermore it is clear what the inclusion maps $V \cap W \to V$ and $V \cap W \to W$ induce in homology:

an isomorphism in H_0 and (necessarily) the zero map in H_k for all $k \neq 0$. Thus the homomorphism

$$H_k(V \cap W) \longrightarrow H_k(V) \oplus H_k(W)$$

from the Mayer-Vietoris sequence takes the form

$$\mathbb{Z} \longrightarrow \mathbb{Z} \oplus \mathbb{Z}$$

when $k = 0$, and

$$\mathbb{Z} \longrightarrow 0$$

when $k = n - 1$. In all other cases, its source and target are both zero. Therefore the exactness of the Mayer-Vietoris sequence implies that $H_0(S^n)$ and $H_n(S^n)$ are isomorphic to $\mathbb{Z}$, while $H_k(S^n) = 0$ for all other $k \in \mathbb{Z}$. □

THEOREM 7.1.3. *Let* $f\colon S^n \to S^n$ *be the antipodal map. The induced homomorphism* $f_*\colon H_n(S^n) \to H_n(S^n)$ *is multiplication by* $(-1)^{n+1}$.

PROOF. We proceed by induction again. For the induction beginning, we take $n = 1$. The antipodal map $f : S^1 \to S^1$ is homotopic to the identity, so that $f^* : H_1(S^1) \to H_1(S^1)$ has to be the identity, too. For the induction step, we use the setup and notation from the previous proof. Exactness of the Mayer-Vietoris sequence for S^n and the open covering $\{V, W\}$ shows that

$$\partial : H_n(S^n) \longrightarrow H_{n-1}(V \cap W)$$

is an isomorphism. The diagram

$$\begin{array}{ccc} H_n(S^n) & \xrightarrow{\partial} & H_{n-1}(V \cap W) \\ {\scriptstyle f_*}\big\downarrow & & {\scriptstyle f_*}\big\downarrow \\ H_n(S^n) & \xrightarrow{\partial} & H_{n-1}(W \cap V) \end{array}$$

is meaningful because f takes $V \cap W$ to $V \cap W = W \cap V$. But the diagram is not commutative (i.e., it is not true that $f_* \circ \partial$ equals $\partial \circ f_*$). The reason is that f interchanges V and W, and it does matter in the Mayer-Vietoris sequence which of the two comes first. Therefore we have instead

$$f_* \circ \partial = -\partial \circ f_*$$

in the above square. By the inductive hypothesis, the f_* in the left-hand column of the square is multiplication by $(-1)^n$, and therefore the f^* in the right-hand column of the square must be multiplication by $(-1)^{n+1}$. □

7.2. The usual applications

THEOREM 7.2.1. (Brouwer's fixed point theorem). *Let* $f : D^n \to D^n$ *be a continuous map, where* $n \geq 1$. *Then* f *has a fixed point, i.e., there exists* $y \in D^n$ *such that* $f(y) = y$.

PROOF. Suppose for a contradiction that f does not have a fixed point. For $x \in D^n$, let $g(x)$ be the point where the ray (half-line) from $f(x)$ to x intersects the boundary S^{n-1} of the disk D^n. Then g is a smooth map from D^n to S^{n-1}, and we have $g|S^{n-1} = \mathrm{id}_{S^{n-1}}$. Summarizing, we have

$$S^{n-1} \xrightarrow{j} D^n \xrightarrow{g} S^{n-1}$$

where j is the inclusion, $g \circ j = \mathrm{id}_{S^{n-1}}$. Therefore we get

$$H_{n-1}(S^{n-1}) \xrightarrow{j_*} H_{n-1}(D^n) \xrightarrow{g_*} H_{n-1}(S^{n-1})$$

where $g_* j_* = \mathrm{id}$. Thus the abelian group $H_{n-1}(S^{n-1})$ is isomorphic to a direct summand of $H_{n-1}(D^n)$. But from our calculations above, we know that this is not true. If $n > 1$ we have $H_{n-1}(D^n) = 0$ while $H_{n-1}(S^{n-1})$ is not trivial. If $n = 1$ we have $H_{n-1}(D^n) \cong \mathbb{Z}$ while $H_{n-1}(S^{n-1}) \cong \mathbb{Z} \oplus \mathbb{Z}$. □

Let $f\colon S^n \to S^n$ be any continuous map, $n > 0$. The induced homomorphism f_* from $H_n(S^n)$ to $H_n(S^n)$ is multiplication by some number $n_f \in \mathbb{Z}$, since $H_n(S^n)$ is isomorphic to $\mathbb{Z}$.

DEFINITION 7.2.2. The number n_f is the *degree* of f.

Remark. The degree n_f of $f\colon S^n \to S^n$ is clearly an invariant of the homotopy class of f.

Remark. In the case $n = 1$, the definition of degree as given just above agrees with the definition of degree given in section 1. See exercises.

EXAMPLE 7.2.3. According to theorem 7.1.3, the degree of the antipodal map $S^n \to S^n$ is $(-1)^{n+1}$.

PROPOSITION 7.2.4. *Let $f\colon S^n \to S^n$ be a continuous map. If $f(x) \neq x$ for all $x \in S^n$, then f is homotopic to the antipodal map, and so has degree $(-1)^{n+1}$. If $f(x) \neq -x$ for all $x \in S^n$, then f is homotopic to the identity map, and so has degree 1.*

PROOF. Let $g\colon S^n \to S^n$ be the antipodal map, $g(x) = -x$ for all x. Assuming that $f(x) \neq x$ for all x, we show that f is homotopic to g. We think of S^n as the unit sphere in $\mathbb{R}^{n+1}$, with the usual notion of distance. We can make a homotopy $(h_t : S^n \to S^n)_{t\in[0,1]}$ from f to g by "sliding" along the unique minimal geodesic arc from $f(x)$ to $g(x)$, for every $x \in S^n$. In other words, $h_t(x) \in S^n$ is situated $t \cdot 100$ percent of the way from $f(x)$ to $g(x)$ along the minimal geodesic arc from $f(x)$ to $g(x)$. (The important thing here is that $f(x)$ and $g(x)$ are not antipodes of each other, by our assumptions. Therefore that minimal geodesic arc is unique.)
Next, assume $f(x) \neq -x$ for all $x \in S^n$. Then, for every x, there is a unique minimal geodesic from x to $f(x)$, and we can use that to make a homotopy from the identity map to f. □

COROLLARY 7.2.5. (Hairy ball theorem). *Let ξ be a tangent vector field* (explanations follow) *on S^n. If $\xi(z) \neq 0$ for every $z \in S^n$, then n is odd.*

Comments. A tangent vector field on $S^n \subset \mathbb{R}^{n+1}$ can be defined as a continuous map ξ from S^n to the vector space $\mathbb{R}^{n+1}$ such that $\xi(x)$ is perpendicular to (the position vector of) x, for every $x \in S^n$. We say that vectors in $\mathbb{R}^{n+1}$ which are perpendicular to $x \in S^n$ are *tangent* to S^n at x because they are the velocity vectors of smooth curves in $S^n \subset \mathbb{R}^n$ as the pass through x.

PROOF. Define $f : S^n \to S^n$ by $f(x) = \xi(x)/\|\xi(x)\|$. Then $f(x) \neq x$ and $f(x) \neq -x$ for all $x \in S^n$, since $f(x)$ is always perpendicular to x. Therefore f is homotopic to the antipodal map, and also homotopic to the identity. It follows that the antipodal map is homotopic to the identity. Therefore n is odd by theorem 7.1.3. □

REMARK 7.2.6. Theorem 7.1.3 has an easy generalization which says that the degree of the map $f\colon S^n \to S^n$ given by

$$(x_1, x_2, \ldots, x_{n+1}) \mapsto (x_1, \ldots, x_k, -x_{k+1}, \ldots, -x_{n+1})$$

is $(-1)^{n+1-k}$. Here we assume $n \geq 1$ as usual. The proof can be given by induction on $n+1-k$. The induction step is now routine, but the induction beginning must cover all cases where $n=1$. This leaves the three possibilities $k=0,1,2$. One of these gives the identity map $S^1 \to S^1$, and another gives the antipodal map $S^1 \to S^1$ which is homotopic to the identity. The interesting case which remains is the map $f\colon S^1 \to S^1$ given by $f(x_1, x_2) = (x_1, -x_2)$. We need to show that it has degree -1, in the sense of definition 7.2.2. One way to do this is to use the following diagram

$$\begin{tikzcd} H_1(S^1) \arrow[r, "f_*"] \arrow[d, "\partial"] & H_1(S^1) \arrow[d, "\partial"] \\ H_0(V \cap W) \arrow[r, "f_*"] & H_0(W \cap V) \end{tikzcd}$$

where $V = S^1 \smallsetminus \{(0,1)\}$ and $W = S^1 \smallsetminus \{(0,-1)\}$. We know from the previous chapter that it commutes up to a factor (-1). In the lower row, we have the identity homomorphism $\mathbb{Z} \oplus \mathbb{Z} \to \mathbb{Z} \oplus \mathbb{Z}$. The vertical arrows are injective (seen earlier in the proof of proposition 7.1.1). Therefore the upper horizontal arrow is multiplication by -1.

To state this result in a more satisfying manner, let us note that the orthogonal group $O(n+1)$ (the group of orthogonal $(n+1) \times (n+1)$-matrices with real entries) is a topological group which has two path components. The two path components are the preimages of $+1$ and -1 under the homomorphism

$$\det\colon O(n+1) \to \{-1, +1\}.$$

Let $f\colon S^n \to S^n$ be given by $f(z) = Az$ for some $A \in O(n+1)$. Because $\deg(f)$ depends only on the homotopy class of f, it follows that $\deg(f)$ depends only on the path component of A in $O(n+1)$, and hence only on $\det(A)$. What we have just shown means that $\deg(f)$ is *equal* to $\det(A)$.

CHAPTER 8

Proving the homotopy decomposition theorem

8.1. Reductions

Here we reduce the proof of the homotopy decomposition theorem to the following lemmas.

LEMMA 8.1.1. *Let* Z *be a paracompact topological space,* Y *any topological space. Let* $\beta\colon Z\times[0,1]\to Y$ *be a mapping cycle. Write* $\iota_0,\iota_1\colon Z\to Z\times[0,1]$ *for the maps given by* $\iota_0(z)=(z,0)$ *and* $\iota_1(z)=(z,1)$. *If there exists a decomposition*

$$\beta\circ\iota_0=\beta_0^V+\beta_0^W$$

where β_0^V *and* β_0^W *are mapping cycles from* Z *to* V *and* W*, respectively, then there exists a decomposition* $\beta\circ\iota_1=\beta_1^V+\beta_1^W$.

LEMMA 8.1.2. *In the situation of lemma 8.1.1, every element of* Z *has an open neighborhood* U *such that the restriction* $\beta_{U\times[0,1]}$ *of* β *to* $U\times[0,1]$ *admits a decomposition*

$$\beta_{U\times[0,1]}=\beta^V_{U\times[0,1]}+\beta^W_{U\times[0,1]}$$

where $\beta^V_{U\times[0,1]}$ *and* $\beta^W_{U\times[0,1]}$ *are mapping cycles from* $U\times[0,1]$ *to* V *and* W*, respectively.*

SHOWING THAT LEMMA 8.1.2 IMPLIES LEMMA 8.1.1. In the situation of lemma 8.1.1, choose an open cover $(U_k)_{k\in\Lambda}$ such that the restriction $\beta_{[k]}$ of β to $U_k\times[0,1]$ admits a decomposition

$$\beta_{[k]}=\beta^V_{[k]}+\beta^W_{[k]}\ .$$

Such an open cover exists by lemma 8.1.2. Since Z is paracompact, there is no loss of generality in assuming that the open cover is locally finite. Moreover, there exists a partition of unity $(\varphi_k)_{k\in\Lambda}$ subordinate to the cover $(U_k)_{k\in\Lambda}$. Choose a total ordering of Λ. If Λ is finite, we can proceed as follows. We may assume that Λ is $\{1,2,3,\dots,m\}$ for some m, with the standard ordering. For $k\in\{0,1,\dots,m\}$ let

$$f_k\colon Z\to Z\times[0,1]$$

be the function $z\mapsto(z,\sum_{\ell=1}^k\varphi_\ell)$. Then $f_0=\iota_0$ and $f_m=\iota_1$ in the notation of lemma 8.1.1. By induction on k we define a decomposition

$$\beta\circ f_k=(\beta\circ f_k)^V+(\beta\circ f_k)^W\ .$$

For $k=0$ this decomposition (of $\beta\circ f_0=\beta\circ\iota_0$) is already given to us. If we have constructed the decomposition for $\beta\circ f_{k-1}$, where $0<k\le m$, we define it for $\beta\circ f_k$ in such a way that

$$(\beta\circ f_k)^V=(\beta\circ f_{k-1})^V+\beta^V_{[k]}\circ f_k-\beta^V_{[k]}\circ f_{k-1}$$

on $U_k\subset Z$ and $(\beta\circ f_k)^V=(\beta\circ f_{k-1})^V$ outside the support of φ_k. Similarly, define

$$(\beta\circ f_k)^W=(\beta\circ f_{k-1})^W+\beta^W_{[k]}\circ f_k-\beta^W_{[k]}\circ f_{k-1}$$

on U_k and $(\beta\circ f_k)^W = (\beta\circ f_{k-1})^W$ outside the support of φ_k. Then on U_k we have

$$(\beta\circ f_k)^V + (\beta\circ f_k)^W = \beta\circ f_{k-1} + \beta\circ f_k - \beta\circ f_{k-1} = \beta\circ f_k$$

and outside the support of φ_k we have

$$(\beta\circ f_k)^V + (\beta\circ f_k)^W = (\beta\circ f_{k-1})^V + (\beta\circ f_{k-1})^W = \beta\circ f_{k-1} = \beta\circ f_k\ .$$

Therefore $(\beta\circ f_k)^V + (\beta\circ f_k)^W = \beta\circ f_k$ as required. The case $k = m$ is the decomposition of $\beta\circ\iota_1 = \beta\circ f_m$ that we are after.

If Λ is not finite, we can proceed as follows. Choose $z\in Z$ and an open neighborhood Q of z in Z such that the set

$$J = \{\, k\in\Lambda \mid Q\cap U_k \neq \emptyset \,\}$$

is finite. Now J is a finite set with a total ordering, and the φ_j where $j\in J$ constitute a partition of unity for Q, subordinate to the open cover $(U_k\cap Q)_{k\in J}$ of Q. Use this as above to find a decomposition of $\beta\circ\iota_1$, restricted to Q, into summands which are mapping cycles from Q to V and W, respectively. Do this for every z and open neighborhood Q. The decompositions obtained match on overlaps, and so define a decomposition of $\beta\circ\iota_1$ of the required sort. □

Showing that lemma 8.1.1 implies the homotopy decomposition theorem. Given X, Y and a mapping cycle $\gamma\colon X\times[0,1]\to Y$, we look for a decomposition $\gamma = \gamma^V + \gamma^W$ where γ^V and γ^W are mapping cycles from $X\times[0,1]$ to V and W, respectively. There is an additional condition to be satisfied. Namely, γ is zero on an open neighborhood U of $(X\times\{0\})\cup(C\times[0,1])$ in $X\times[0,1]$, and we want γ^V,γ^W to be zero on some (perhaps smaller) open neighborhood U' of $(X\times\{0\})\cup(C\times[0,1])$ in $X\times[0,1]$.

Put $Z = X\times[0,1]$. Since X was assumed to be paracompact, Z is also paracompact; it is a general topology fact that the product of a paracompact space with a compact Hausdorff space is paracompact. We have a map

$$h\colon Z\times[0,1]\to Z$$

defined by $h((x,s),t)) = (x,st)$ for $(x,t)\in X\times[0,1] = Z$ and $t\in[0,1]$. Now $\beta := \gamma\circ h$ is a mapping cycle from $Z\times[0,1]$ to Y. In the notation of lemma 8.1.1, we have

$$\beta\circ\iota_1 = \gamma, \qquad \beta\circ\iota_0 \equiv 0\ .$$

There exists a decomposition $\beta_0 = \beta_0^V + \beta_0^W$ because we can take $\beta_0^V\equiv 0$ and $\beta_0^W\equiv 0$. Therefore, by lemma 8.1.1, there exists a decomposition $\beta\circ\iota_1 = \beta_1^V + \beta_1^W$, and we can write that in the form

$$\gamma = \beta_1^V + \beta_1^W\ .$$

This is a decomposition of the kind that we are looking for. Unfortunately there is no reason to expect that β_1^V, β_1^W are zero on $(X\times\{0\})\cup(C\times[0,1])$, or on a neighborhood of that in $X\times[0,1]$.

But it is easy to construct a continuous map $\psi\colon X\times[0,1]\to X\times[0,1]$ such that $\psi(X\times[0,1])$ is contained in the open set U specified above, and such that ψ agrees with the identity on some open neighborhood U' of $(X\times\{0\})\cup(C\times[0,1])$ in $X\times[0,1]$. Then obviously $U'\subset U$. Now let

$$\gamma^V = \beta_1^V - (\beta_1^V\circ\psi), \qquad \gamma^W = \beta_1^W - (\beta_1^W\circ\psi).$$

Then $\gamma^V + \gamma^W = (\beta_1^V + \beta_1^W) - (\beta_1^V + \beta_1^W)\circ\psi = \gamma - \gamma\circ\psi$. Furthermore $\gamma\circ\psi$ is zero because γ is zero on U and the image of ψ is contained in U. So $\gamma^V + \gamma^W = \gamma$. Also γ^V and γ^W are zero on U' by construction, since ψ agrees with the identity on U'. □

8.2. Local homotopy decomposition

Proof of lemma 8.1.2. Call an open subset P of $Z \times [0,1]$ *good* if the mapping cycle $\beta_{|P}$ from P to Y can be written as the sum of a mapping cycle from P to V and a mapping cycle from P to W. The goal is to show that every $z \in Z$ has an open neighborhood U such that $U \times [0,1]$ is good.
The proof is based on two observations.

- Every element of $Z \times [0,1]$ admits a good open neighborhood.
- If U is open in Z and A, B are open subsets of $[0,1]$ which are also intervals, and if $U \times A$ and $U \times B$ are both good, then $U \times (A \cup B)$ is good.

To prove the first observation, fix $(z,t) \in Z \times [0,1]$ and choose an open neighborhood Q of that in $Z \times [0,1]$ such that $\beta_{|Q}$ can be written as a formal linear combination, with coefficients in $\mathbb{Z}$, of continuous maps from Q to Y. Such a Q exists by the definition of *mapping cycle*. Making Q smaller if necessary, we can arrange that each of the (finitely many) continuous maps which appear in that formal linear combination is either a map from Q to V or a map from Q to W. It follows immediately that Q is good.
In proving the second observation, we can easily reduce to a situation where $A \cap B$ contains an element t_0, where $0 < t_0 < 1$, and $A \cup B$ is the union of $A \cap [0, t_0]$ and $B \cap [t_0, 1]$. Choose a continuous map $\psi \colon B \to B \cap A$ such that $\psi(s) = s$ for all $s \in B \cap [0, t_0]$. Since $P := U \times A$ is good by assumption, we can write

$$\beta_{|P} = \beta^{V,P} + \beta^{W,P}$$

where the summands in the right-hand side are mapping cycles from P to V and from P to W, respectively. Similarly, letting $Q := U \times B$ we can write

$$\beta_{|Q} = \beta^{V,Q} + \beta^{W,Q}\,.$$

Let $\varphi \colon Q \to P \cap Q$ be given by $\varphi(z,t) = (z, \psi(t))$. Define $\beta^{V,P\cup Q}$, a mapping cycle from $P \cup Q$ to V, as follows:

$$\beta^{V,P\cup Q} = \begin{cases} \beta^{V,P} & \text{on} \quad P \cap (U \times [0, t_0[\,) \\ \beta^{V,Q} - (\beta^{V,Q} \circ \varphi) + (\beta^{V,P} \circ \varphi) & \text{on} \quad Q\,. \end{cases}$$

This is well defined because the two formulas agree on the intersection of Q and $U \times [0, t_0[$, where φ agrees with the identity. Similarly, define $\beta^{W,P\cup Q}$, a mapping cycle from $P \cup Q$ to W, as follows:

$$\beta^{W,P\cup Q} = \begin{cases} \beta^{W,P} & \text{on} \quad P \cap (U \times [0, t_0[\,) \\ \beta^{W,Q} - (\beta^{W,Q} \circ \varphi) + (\beta^{W,P} \circ \varphi) & \text{on} \quad Q\,. \end{cases}$$

An easy calculation shows that $\beta^{V,P\cup Q} + \beta^{W,P\cup Q} = \beta_{|P\cup Q}$. Therefore $P \cup Q = U \times (A \cup B)$ is good. The second observation is established.
Now fix $z_0 \in Z$. By the first of the observations, it is possible to choose for each $t \in [0,1]$ a good open neighborhood Q_t of (z_0, t) in $Z \times [0,1]$. By a little exercise, there exists an open neighborhood U of z_0 in Z and a small number $\delta = 1/n$ (where n is a positive integer) such that each of the open sets

$$U \times [0, 2\delta[\,, \quad U \times]1\delta, 3\delta[\,, \quad U \times]2\delta, 4\delta[\,, \quad \dots,$$

$$U \times]1 - 3\delta, 1 - 1\delta[\,, \quad U \times]1 - 2\delta, 1]$$

in $Z \times [0,1]$ is contained in Q_t for some $t \in [0,1]$. Therefore these open sets

$$U \times [0, 2\delta[\,, U \times]1\delta, 3\delta[\,, \dots$$

are also good. By the second of the two observations, applied $(n-2)$ times, their union, which is $U \times [0,1]$, is also good. $\square$

8.3. Relationship with fiber bundles

The proof of the homotopy decomposition theorem as given above has many surprising similarities with proofs in section 3 related to fiber bundles (theorem 3.4, corollaries 3.7 and 3.8., and improvements in section 3.4). I cannot resist the temptation to explain these similarities now, after the proof.

Let E and B be topological spaces and let $p\colon E \to B$ be a fiber bundle. We need to be a little more precise by requiring that $p\colon E \to B$ be a fiber bundle *with fiber* F, for a fixed topological space F. This is supposed to mean that every fiber of p is homeomorphic to F in some way. (We learned in section 2 that every fiber bundle over a path connected space is a fiber bundle with fiber F, for some F.) With this situation we can associate two presheaves $\mathcal{T}$ and $\mathcal{H}_F$ on B.

- For an open set U in B, let $\mathcal{H}_F(U)$ be the group of homeomorphisms h from $U \times F$ to $U \times F$ respecting the projection to U.
- For an open set U in B let $\mathcal{T}(U)$ be the set of trivializations of the fiber bundle $E_{|U} \to U$, that is, the set of all homeomorphisms $p^{-1} \to U \times F$ respecting the projections to U.
- An inclusion of open sets $U_0 \hookrightarrow U_1$ in B induces maps

$$\mathcal{H}_F(U_1) \to \mathcal{H}_F(U_0), \qquad \mathcal{T}(U_1) \to s\mathsf{T}(U_0)$$

  by restriction of homeomorphisms.

In fact it is clear that $\mathcal{T}$ and $\mathcal{H}_F$ are sheaves. Clearly $\mathcal{H}_F$ is a sheaf of groups, that is, each set $\mathcal{H}_F(U)$ comes with a group structure and the restriction maps $\mathcal{H}_F(U_1) \to \mathcal{H}_F(U_0)$ are group homomorphisms. By contrast $\mathcal{T}$ is not a sheaf of groups in any obvious way. But there is an *action* of the group $\mathcal{H}_F(U)$ on the set $\mathcal{T}(U)$ given by

$$(h, g) \mapsto h \circ g$$

(composition of homeomorphisms, where $h \in \mathcal{H}_F(U)$ and $g \in \mathcal{T}(U)$). This is compatible with restriction maps (reader, make this precise). Moreover:

(1) for any $g \in \mathcal{T}(U)$, the map $\mathcal{H}_F(U) \to \mathcal{T}(U)$ given by $h \mapsto h \circ g$ is a bijection;
(2) every $z \in B$ has an open neighborhood U such that $\mathcal{T}(U) \neq \emptyset$.

(Of course, despite (1), it can happen that $\mathcal{T}(U)$ is empty for some open subsets U of B, for example, $U = B$.) The proof of (1) is easy and by inspection; (2) holds by the definition of *fiber bundle*. There are words and expressions to describe this situation: we can say that $\mathcal{H}_F$ is a sheaf of groups on B and $\mathcal{T}$ is an $\mathcal{H}_F$*-torsor*.

This reasoning shows that a fiber bundle on B with fiber F determines an $\mathcal{H}_F$-torsor on B. It is also true (and useful, and not very hard to prove, though it will not be explained here) that the process can be reversed: every $\mathcal{H}_F$-torsor on B determines a fiber bundle with fiber F on B. So it transpires that section 3 about fiber bundles could alternatively have been written in the language of sheaves (of sets or groups) and torsors. Note that we are often interested in questions like this one: is $\mathcal{T}(B)$ nonempty? This amounts to asking whether the fiber bundle p is a trivial fiber bundle.

REMARK 8.3.1. For the sake of honesty it should be pointed out that $\mathcal{H}_F$ is a sheaf on *all topological spaces* simultaneously, and this would become important if we really wanted to rewrite section 3 in sheaf language. In more detail:

- We can view $\mathcal{H}_F$ as a contravariant functor from topological spaces to groups. Indeed, for a topological space X let $\mathcal{H}_F(X)$ be the group of homeomorphisms from $X \times F$ to $X \times F$ respecting the projection to X. A continuous map $X_0 \to X_1$ induces a map $\mathcal{H}_F(X_1) \to \mathcal{H}_F(X_0)$ which is a group homomorphism.
- If we evaluate this functor only on open subsets of a fixed space X, and on inclusion maps $U_0 \to U_1$ of open subsets of X, then the resulting presheaf on X is in fact a sheaf on X.

There are also words and expressions for this; to keep it short, I will just say that $\mathcal{H}_F$ is a *big sheaf.*

Now try to forget fiber bundles for a while. We return to the homotopy decomposition theorem. Assume that $Y = V \cup W$ as in the homotopy decomposition theorem. Let Z be any topological space and fix α, a mapping cycle from Z to Y. We introduce two presheaves $\mathcal{F}$ and $\mathcal{G}$ on Z.

- For an open set U in Z, let $\mathcal{G}(U)$ be the abelian group of mapping cycles from U to $V \cap W$.
- For open U in Z let $\mathcal{F}(U)$ be the set of mapping cycles β from U to V such that $\alpha_{|U} - \beta$ is a mapping cycle from U to W. To put it differently: an element β of $\mathcal{F}(U)$ is, or amounts to, a sum decomposition
$$\alpha_{|U} = \beta + (\alpha_{|U} - \beta)$$
where the two summands β and $\alpha_{|U} - \beta$ are mapping cycles from U to V and from U to W, respectively.
- An inclusion of open sets $U_0 \hookrightarrow U_1$ in Z induces maps
$$\mathcal{G}(U_1) \to \mathcal{G}(U_0), \qquad \mathcal{F}(U_1) \to sF(U_0)$$
by restriction of mapping cycles.

It is easy to see that $\mathcal{F}$ and $\mathcal{G}$ are *sheaves*, and $\mathcal{G}$ is even a sheaf of abelian groups on Z. By contrast $\mathcal{F}$ is not in an obvious way a sheaf of abelian groups. But there is an *action* of the group $\mathcal{G}(U)$ on the set $\mathcal{F}(U)$ given by
$$(\lambda, \beta) \mapsto \lambda + \beta\,.$$
(In this formula, $\lambda \in \mathcal{G}(U)$ and $\beta \in \mathcal{F}(U)$; then $\lambda + \beta$ can be viewed as a mapping cycle from U to V and it turns out to be an element of $\mathcal{F}(U)$.) Moreover:

(1) for any $\beta \in \mathcal{F}(U)$, the map $\mathcal{G}(U) \to \mathcal{F}(U)$ given by $\lambda \mapsto \lambda + \beta$ is a bijection;
(2) every $z \in Z$ has an open neighborhood U such that $\mathcal{F}(U) \neq \emptyset$.

(Of course it is quite possible, despite (1), that $\mathcal{F}(U)$ is empty for some open subsets U of Z, for example, $U = Z$.) The proof of (1) is easy and by inspection; the proof of (2) was given in a special case earlier, but it can be repeated. Choose a neighborhood U of z such that $\alpha_{|U}$ can be represented by a formal linear combination, with integer coefficients, of continuous maps from U to Y. Making U smaller if necessary, we can assume that each of the (finitely many) continuous maps which appear in that formal linear combination is either a map from U to V or a map from U to W. Then it is clear that $\alpha_{|U}$ can be written as a sum of two mapping cycles, one from U to V and the other from U to W. So $\mathcal{F}(U)$ is nonempty.

So we see that $\mathcal{G}$ is a sheaf of abelian groups on Z and $\mathcal{F}$ is a $\mathcal{G}$-torsor. Again we are interested in questions like this one: is $\mathcal{F}(Z)$ nonempty? This is equivalent to asking whether our fixed mapping cycle α from Z to Y can be written as a sum of two mapping cycles, one from Z to V and one from Z to W. And again, for the sake of honesty, it should be noted that $\mathcal{G}$ is a big sheaf of abelian groups. (If we wanted to rewrite the proof of the homotopy decomposition theorem in sheaf and torsor language, that would have to be used.)

CHAPTER 9

Combinatorial description of some spaces

9.1. Vertex schemes and simplicial complexes

DEFINITION 9.1.1. A *vertex scheme* consists of a set V and a subset $\mathcal{S}$ of the power set $\mathcal{P}(V)$, subject to the following conditions: every $T \in \mathcal{S}$ is finite and nonempty, every subset of V which has exactly one element belongs to $\mathcal{S}$, and if T' is a nonempty subset of some $T \in \mathcal{S}$, then $T' \in \mathcal{S}$.
The elements of V are called *vertices* (singular: *vertex*) of the vertex scheme. The elements of $\mathcal{S}$ are called *distinguished subsets* of V.

EXAMPLE 9.1.2. The following are examples of vertex schemes:

(i) Let $V = \{1, 2, 3, \ldots, 10\}$. Define $\mathcal{S} \subset \mathcal{P}(V)$ so that the elements of $\mathcal{S}$ are the following subsets of V: all the singletons, that is to say $\{1\}, \{2\}, \ldots, \{10\}$, and $\{1, 2\}$, $\{2, 3\}$, ..., $\{9, 10\}$ as well as $\{10, 1\}$.
(ii) Let $V = \{1, 2, 3, 4\}$ and define $\mathcal{S} \subset \mathcal{P}(V)$ so that the elements of $\mathcal{S}$ are exactly the subsets of V which are nonempty and not equal to V.
(iii) Let V be any set and define $\mathcal{S}$ so that the elements of $\mathcal{S}$ are exactly the nonempty finite subsets of V.
(iv) Take a regular icosahedron. Let V be the set of its vertices (which has 12 elements). Define $\mathcal{S} \subset \mathcal{P}(V)$ in such a way that the elements of $\mathcal{S}$ are all singletons, all doubletons which are connected by an edge, and all tripletons which make up a triangular face of the icosahedron. (There are twenty such tripletons, which is supposed to explain the name *icosahedron*.)

The *simplicial complex* determined by a vertex scheme $(V, \mathcal{S})$ is a topological space $X = |V|_{\mathcal{S}}$. We describe it first as a set. An element of X is a function $f \colon V \to [0, 1]$ such that

$$\sum_{v \in V} f(v) = 1$$

and the set $\{v \in V \mid f(v) > 0\}$ is an element of $\mathcal{S}$.
It should be clear that X is the union of certain subsets $\Delta(T)$, where $T \in \mathcal{S}$. Namely, $\Delta(T)$ consists of all the functions $f \colon V \to [0, 1]$ for which $\sum_{v \in V} f(v) = 1$ and $f(v) = 0$ if $v \notin T$. The subsets $\Delta(T)$ of X are not always disjoint. Instead we have $\Delta(T) \cap \Delta(T') = \Delta(T \cap T')$ if $T \cap T'$ is nonempty; also, if $T \subset T'$ then $\Delta(T) \subset \Delta(T')$.
The subsets $\Delta(T)$ of X, for $T \in \mathcal{S}$, come equipped with a preferred topology. Namely, $\Delta(T)$ is (identified with) a subset of a finite dimensional real vector space, the vector space of all functions from T to $\mathbb{R}$, and as such gets a subspace topology. (For example, $\Delta(T)$ is a single point if T has one element; it is homeomorphic to an edge or closed interval if T has two elements; it looks like a compact triangle if T has three elements; etc. We say that $\Delta(T)$ is a *simplex* of dimension m if T has cardinality $m + 1$.) These topologies are compatible in the following sense: if $T \subset T'$, then the inclusion $\Delta(T) \to \Delta(T')$ makes a

homeomorphism of $\Delta(T)$ with a subspace of $\Delta(T')$.
We decree that a subset W of X shall be *open* if and only if $W \cap \Delta(T)$ is open in $\Delta(T)$, for every T in $\mathcal{S}$. Equivalently, and perhaps more usefully: a map g from X to another topological space Y is continuous if and only if the restriction of g to $\Delta(T)$ is a continuous from $\Delta(T)$ to Y, for every $T \in \mathcal{S}$.

EXAMPLE 9.1.3. The simplicial complex associated to the vertex scheme (i) in example 9.1.2 is homeomorphic to S^1. In (ii) and (iv) of example 9.1.2, the associated simplicial complex is homeomorphic to S^2.

EXAMPLE 9.1.4. The simplicial complex associated to the vertex scheme $(V, \mathcal{S})$ where $V = \{1,2,3,4,5,6,7,8\}$ and

$$\mathcal{S} = \left\{ \begin{array}{l} \{1\},\{2\},\{3\},\{4\},\{5\},\{6\},\{7\},\{8\},\{1,3\},\{2,3\},\{3,4\}, \\ \{3,5\},\{3,6\},\{4,5\},\{5,6\},\{5,7\},\{7,8\},\{3,4,5\},\{3,5,6\} \end{array} \right\}$$

looks like this:

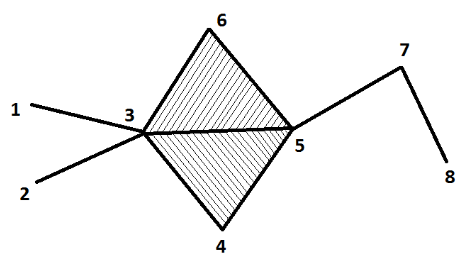

LEMMA 9.1.5. *The simplicial complex* $X = |V|_{\mathcal{S}}$ *associated with a vertex scheme* $(V, \mathcal{S})$ *is a Hausdorff space.*

PROOF. Let f and g be distinct elements of X. Keep in mind that f and g are functions from V to $[0,1]$. Choose $v_0 \in V$ such that $f(v_0) \neq g(v_0)$. Let $\varepsilon = |f(v_0) - g(v_0)|$. Let U_f be the set of all $h \in X$ such that $|h(v_0) - f(v_0)| < \varepsilon/2$. Let U_g be the set of all $h \in X$ such that $|h(v_0) - g(v_0)| < \varepsilon/2$. From the definition of the topology on X, the sets U_f and U_g are open. They are also disjoint, for if $h \in U_f \cap U_g$ then $|f(v_0) - g(v_0)| \leq |f(v_0) - h(v_0)| + |h(v_0) - g(v_0)| < \varepsilon$, contradiction. Therefore f and g have disjoint neighborhoods in X. □

LEMMA 9.1.6. *Let* $(V, \mathcal{S})$ *be a vertex scheme and* $(W, \mathcal{T})$ *a vertex sub-scheme, that is,* $W \subset V$ *and* $\mathcal{T} \subset \mathcal{S} \cap \mathcal{P}(W)$. *Then the evident map* $\iota\colon |W|_{\mathcal{T}} \to |V|_{\mathcal{S}}$ *is a closed, continuous and injective map and therefore a homeomorphism onto its image.*

PROOF. The map ι is obtained by viewing functions from W to $[0,1]$ as functions from V to $[0,1]$ by defining the values on elements of $V \smallsetminus W$ to be 0. A subset A of $|V|_{\mathcal{S}}$ is closed if and only if $A \cap \Delta(T)$ is closed for the standard topology on $\Delta(T)$, for every $T \in \mathcal{S}$. Therefore, if A is a closed subset of $|V|_{\mathcal{S}}$, then $\iota^{-1}(A)$ is a closed subset of $|W|_{\mathcal{T}}$; and if C is a closed subset of $|W|_{\mathcal{S}}$, then $\iota(C)$ is closed in $|V|_{\mathcal{S}}$. □

REMARK 9.1.7. The notion of a simplicial complex is old. Related vocabulary comes in many dialects. I have taken the expression *vertex scheme* from Dold's book *Lectures on* *algebraic topology* with only a small change (for me, $\emptyset \notin \mathcal{S}$). It is in my opinion a good choice of words, but the traditional expression for that appears to be *abstract simplicial complex*. Most authors agree that a *simplicial complex* (non-abstract) is a topological space with additional data. For me, a simplicial complex is a space of the form $|V|_{\mathcal{S}}$ for some vertex scheme $(V, \mathcal{S})$; other authors prefer to write, in so many formulations, that a simplicial complex is a topological space X together with a homeomorphism $|V|_{\mathcal{S}} \to X$, for some vertex scheme $(V, \mathcal{S})$.

9.2. Semi-simplicial sets and their geometric realizations

Semi-simplicial sets are closely related to vertex schemes. A semi-simplicial set has a *geometric realization*, which is a topological space; this is similar to the way in which a vertex scheme determines a simplicial complex.

DEFINITION 9.2.1. A semi-simplicial set Y consists of a sequence of sets

$$(Y_0, Y_1, Y_2, Y_3, ...)$$

(each Y_k is a set) and, for each injective order-preserving map

$$f\colon \{0, 1, 2, \dots, k\} \longrightarrow \{0, 1, 2, \dots, \ell\}$$

where $k, \ell \geq 0$, a map $f^*\colon Y_\ell \to Y_k$. The maps f^* are called *face operators* and they are subject to conditions:

- if f is the identity map from $\{0, 1, 2, \dots, k\}$ to $\{0, 1, 2, \dots, k\}$ then f^* is the identity map from Y_k to Y_k.
- $(g \circ f)^* = f^* \circ g^*$ when $g \circ f$ is defined (so $f\colon \{0, 1, \dots, k\} \to \{0, 1, \dots, \ell\}$ and $g\colon \{0, 1, \dots, \ell\} \to \{0, 1, \dots, m\}$).

Elements of Y_k are often called k*-simplices* of Y. If $x \in Y_k$ has the form $f^*(y)$ for some $y \in Y_\ell$, then we may say that x is a *face* of y corresponding to face operator f^*.

REMARK 9.2.2. The definition of a semi-simplicial set can be reformulated in category language as follows. There is a category $\mathcal{C}$ whose objects are the sets $[n] = \{0, 1, \dots, n\}$, where n can be any non-negative integer. A morphism in $\mathcal{C}$ from $[m]$ to $[n]$ is an order-preserving injective map from the set $[m]$ to the set $[n]$. Composition of morphisms is, by definition, composition of such order-preserving injective maps.
A semi-simplicial set is a contravariant functor Y from $\mathcal{C}$ to the category of sets. We like to write Y_n when we ought to write $Y([n])$. We like to write $f^*\colon Y_n \to Y_m$ when we ought to write $Y(f)\colon Y([n]) \to Y([m])$, for a morphism $f\colon [m] \to [n]$ in $\mathcal{C}$.
Nota bene: if you wish to define (invent) a semi-simplicial set Y, you need to invent sets $Y_0, Y_1, Y_2, \dots$ (one set Y_n for each integer $n \geq 0$) *and* you need to invent maps $f^*\colon Y_n \to Y_m$, one for each order-preserving injective map $f\colon [m] \to [n]$. *Then* you need to convince yourself that $(g \circ f)^* = f^* \circ g^*$ whenever $f\colon [k] \to [\ell]$ and $g\colon [\ell] \to [m]$ are order-preserving injective maps.

EXAMPLE 9.2.3. Let $(V, \mathcal{S})$ be a vertex scheme as in the preceding (sub)section. Choose a total ordering of V. From these data we can make a semi-simplicial set Y as follows.

- Y_n is the set of all order-preserving injective maps β from $\{0, 1, \dots, n\}$ such that $\mathrm{im}(\beta) \in \mathcal{S}$. Note that for each $T \in \mathcal{S}$ of cardinality $n+1$, there is exactly one such β.
- For an order-preserving injective $f\colon \{0, 1, \dots, m\} \to \{0, 1, \dots, n\}$ and $\beta \in Y_n$, define $f^*(\beta) = \beta \circ f \in Y_m$.

In order to warm up for geometric realization, we introduce a (covariant) functor from the category $\mathcal{C}$ in remark 9.2.2 to the category of topological spaces. On objects, the functor is given by

$$\{0,1,2,\dots,m\} \mapsto \Delta^m$$

where Δ^m is the space of functions u from $\{0,1,\dots,m\}$ to $\mathbb{R}$ which satisfy the condition $\sum_{j=0}^m u(j) = 1$. (As usual we view this as a subspace of the finite-dimensional real vector space of all functions from $\{0,1,\dots,n\}$ to $\mathbb{R}$. It is often convenient to think of $u \in \Delta^n$ as a vector, $(u_0, u_1, \dots, u_m)$, where all coordinates are ≥ 0 and their sum is 1.) Here is a picture of Δ^2 as a subspace of $\mathbb{R}^3$ (with basis vectors e_0, e_1, e_2):

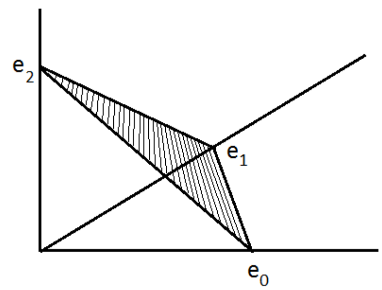

For a morphism f, meaning an order-preserving injective map

$$f\colon \{0,1,2,\dots,m\} \longrightarrow \{0,1,2,\dots,n\},$$

we want to see an induced map

$$f_*\colon \Delta^m \to \Delta^n.$$

This is easy: for $u = (u_0, u_1, \dots, u_m) \in \Delta^m$ we define

$$f_*(u) = v = (v_0, v_1, \dots, v_n) \in \Delta^n$$

where $v_j = u_i$ if $j = f(i)$ and $v_j = 0$ if $j \notin \mathrm{im}(f)$.
(Keep the following conventions in mind. For a covariant functor G from a category $\mathcal{A}$ to a category $\mathcal{B}$, and a morphism $f\colon x \to y$ in $\mathcal{A}$, we often write $f_*\colon G(x) \to G(y)$ instead of $G(f)\colon G(x) \to G(y)$. For a contravariant functor G from a category $\mathcal{A}$ to a category $\mathcal{B}$, and a morphism $f\colon x \to y$ in $\mathcal{A}$, we often write $f^*\colon G(y) \to G(x)$ instead of $G(f)\colon G(y) \to G(x)$.)

The geometric realization $|Y|$ of a semi-simplicial set Y is a topological space defined as follows. Our goal is to have, for each $n \geq 0$ and $y \in Y_n$, a preferred continuous map

$$c_y\colon \Delta^n \to |Y|$$

(the *characteristic map* associated with the simplex $y \in Y_n$). These maps should match in the sense that whenever we have an injective order-preserving

$$f\colon \{0,1,\dots,m\} \to \{0,1,\dots,n\}$$

and $y \in Y_n$, so that $f^*y \in Y_m$, then the diagram

$$\begin{tikzcd} \Delta^n \arrow[r, "c_y"] & {|Y|} \\ \Delta^m \arrow[u, "f_*"] \arrow[r, "c_{f^*y}"] & {|Y|} \arrow[u, "="'] \end{tikzcd}$$

is commutative. There is a "most efficient" way to achieve this. As a set, let $|Y|$ be the set of all symbols $\bar{c}_y(u)$ where $y \in Y_n$ for some $n \geq 0$ and $u \in \Delta^n$, modulo the relations[1]

$$\bar{c}_y(f_*(u)) \sim \bar{c}_{f^*y}(u)$$

(notation and assumptions as in that diagram). This ensures that we have maps c_y from Δ^n to $|Y|$, for each $y \in Y_n$, given in the best tautological manner by

$$c_y(u) := \text{ equivalence class of } \bar{c}_y(u) \text{ .}$$

Also, those little squares which we wanted to be commutative are now commutative because we enforced it. Finally, we say that a subset U of $|Y|$ shall be *open* (definition coming) if and only if $c_y^{-1}(U)$ is open in Δ^n for each characteristic map $c_y \colon \Delta^n \to |Y|$.

A slightly different way (shorter but possibly less intelligible) to say the same thing is as follows:

$$|Y| := \left(\coprod_{n \geq 0} Y_n \times \Delta^n \right) \Big/ \sim$$

where $\sim$ is a certain equivalence relation on $\coprod_n Y_n \times \Delta^n$. It is the smallest equivalence relation which has $(y, f_*(u))$ equivalent to (f^*y, u) whenever $f \colon \{0, 1, \ldots, m\} \to \{0, 1, \ldots, n\}$ is injective order-preserving and $y \in Y_n$, $u \in \Delta^m$. Note that, where it says $Y_n \times \Delta^n$, the set Y_n is regarded as a topological space with the discrete topology, so that $Y_n \times \Delta^n$ has meaning; we could also have written $\coprod_{y \in Y_n} \Delta^n$ instead of $Y_n \times \Delta^n$.
This new formula for $|Y|$ emphasizes the fact that $|Y|$ is a *quotient space* of a topological disjoint union of many standard simplices Δ^n (one simplex for every pair (n, y) where $y \in Y_n$). Go ye forth and look up *quotient space* or *identification topology* in your favorite book on point set topology.— To match the second description of $|Y|$ with the first one, let the element of $|Y|$ represented by $(y, u) \in Y_n \times \Delta^n$ in the second description correspond to the element which we called $c_y(u)$ in the first description of $|Y|$.

EXAMPLE 9.2.4. Fix an integer $n \geq 0$. We might like to invent a semi-simplicial set

$$Y = \underline{\Delta}^n$$

such that $|Y|$ is homeomorphic to Δ^n. The easiest way to achieve that is as follows. Define Y_k to be the set of all order-preserving injective maps from $\{0, 1, \ldots, k\}$ to $\{0, 1, \ldots, n\}$. So Y_k has $\binom{n+1}{k+1}$ elements (which implies $Y_k = \emptyset$ if $k > n$). For an injective order-preserving map

$$g \colon \{0, 1, \ldots, k\} \to \{0, 1, \ldots, \ell\},$$

define the face operator $g^* \colon Y_\ell \to Y_k$ by $g^*(f) = f \circ g$. This makes sense because $f \in Y_\ell$ is an order-preserving injective map from $\{0, 1, \ldots, \ell\}$ to $\{0, 1, \ldots, n\}$. There is a unique

[1] *Modulo the relations* is short for the following process: form the smallest equivalence relation on the set of all those symbols $\bar{c}_y(u)$ which contains the stated relation. Then pass to the set of equivalence classes for that equivalence relation. That set of equivalence classes is $|Y|$.

element $y \in Y_n$, corresponding to the identity map of $\{0, 1, \ldots, n\}$. It is an exercise to verify that the characteristic map $c_y \colon \Delta^n \to |Y|$ is a homeomorphism.

EXAMPLE 9.2.5. Up to relabeling there is a unique semi-simplicial set Y such that Y_0 has exactly one element, Y_1 has exactly one element, and $Y_n = \emptyset$ for $n > 1$. Then $|Y|$ is homeomorphic to S^1. More precisely, let $z \in Y_1$ be the unique element; then the characteristic map

$$c_z \colon \Delta^1 \longrightarrow |Y|$$

is an identification map. (Translation: it is surjective and a subset of the target is open in the target if and only if its preimage is open in the source.) The only identification taking place is $c_z(a) = c_z(b)$, where a and b are the two boundary points of Δ^1.

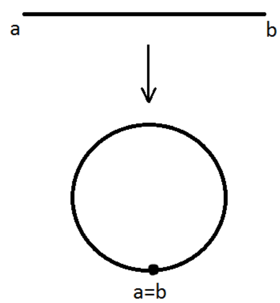

9.3. Technical remarks concerning the geometric realization

Let Y be a semi-simplicial set. We reformulate the definition of the geometric realization $|Y|$ once again.
From the semi-simplicial set Y, we make a category $\mathcal{C}_Y$ as follows. An object is a pair (n, z) where n is a non-negative integer and $z \in Y_n$. A morphism from (m, y) to (n, z) is, by definition, an order-preserving injective map $g \colon \{0, 1, 2, \ldots, m\}$ to $\{0, 1, 2, \ldots, n\}$ which has the property $g^*(z) = y$ (where $g^* \colon Y_n \to Y_m$ is the face operator determined by g).
We define a covariant functor F_Y from $\mathcal{C}_Y$ to the category of topological spaces as follows. The definition of F_Y on objects is simply

$$F_Y(n, z) = \Delta^n$$

where Δ^n is the standard n-simplex. (Recall that this is the space of all functions u from $\{0, 1, \ldots, n\}$ to $[0, 1]$ which satisfy $\sum_j u(j) = 1$, viewed as a subspace of the real vector space of all functions from $\{), 1, \ldots, n\}$ to $\mathbb{R}$.) If we have a morphism from (m, y) to (n, z) given by an an order-preserving injective map $g \colon \{0, 1, 2, \ldots, m\}$ to $\{0, 1, 2, \ldots, n\}$, then we define

$$F_Y(f) = g_* \colon \Delta^m \to \Delta^n,$$

that is to say, $F_Y(f)(u_1, \ldots, u_m) = (v_1, \ldots, v_n)$ where $v_i = u_j$ if $i = g(j)$ and $v_i = 0$ if i is not of the form $g(j)$. Note that I have written u_i instead of $u(i)$ etc. ; strictly speaking $u(i)$ is correct because we said that u is a function from $\{0, 1, \ldots, m\}$ to $[0, 1]$.

Now the definition of $|Y|$ can be recast as follows:

$$|Y| = \left(\coprod_{(n,z)} F_Y(n,z) \right) \Big/ \sim$$

where $\sim$ is the equivalence relation generated by

$$F_Y(m,y) \ \ni \ (u_1,\dots,u_m) \sim F_Y(g)(u_1,\dots,u_m) \in F_Y(n,z)$$

whenever g is a morphism from (m,y) to (n,z); in other words g is an order-preserving injective map from $\{0,1,2,\dots,m\}$ to $\{0,1,2,\dots,n\}$ which has $g^*(z) = y$. It may look as if the formula defines $|Y|$ only as a set, but we want to view it as a formula defining a topology on $|Y|$ as well, namely, the *quotient topology*. Therefore, a subset of $|Y|$ is considered to be *open* (definition) if and only if its preimage in $\coprod_{(n,z)} F_Y(n,z)$ is open.

Warning: do not read these $2\frac{1}{2}$ *lines unless you are somewhat familiar with category theory.* You will notice that $|Y|$ has been defined to be the direct limit (also called colimit) of the functor F_Y.

EXAMPLE 9.3.1. Let $(V,\mathcal{S})$ be a vertex scheme, choose a total ordering on V, and let Y be the associated semi-simplicial set, as in example 9.2.3. We are going to show that the geometric realization $|Y|$ is homeomorphic to the simplicial complex $|V|_{\mathcal{S}}$.
An element of Y_n is an order-preserving injective map from $\{0,1,\dots,n\}$ to V. This is determined by its image T, a distinguished subset of V (where *distinguished* means that $T \in \mathcal{S}$). So we can pretend that Y_n is simply the set of all distinguished subsets of V that have exactly $n+1$ elements. Furthermore, if $T' \in Y_m$ and $T \in Y_n$, then there exists at most one morphism from T' to T in the category $\mathcal{C}_Y$. It exists if and only if $T' \subset T$. Therefore we may say that $\mathcal{C}_Y$ is the category whose objects are the distinguished subsets $T, T', \dots$ of V, with exactly one morphism from T' to T if $T' \subset T$, and no morphism from T' to T otherwise. In this description, the functor F_Y is given on objects by

$$F_Y(T) = \Delta(T)$$

where $\Delta(T)$ replaces Δ^n (assuming that T has exactly $n+1$ elements) and means: the space of functions u from T to $[0,1]$ that satisfy $\sum_{j\in T} u(j) = 1$. For $T' \subset T$ we have exactly one morphism from T' to T, and the induced map $F_Y(T') = \Delta(T') \to \Delta(T) = F_Y(T)$ is given by $u \mapsto v$ where $v(t) = u(t)$ if $t \in S'$ and $v(t) = 0$ if $t \in S \smallsetminus S'$. Therefore

$$|Y| = \Big(\coprod_{T\in\mathcal{S}} \Delta(T) \Big) \Big/ \sim$$

where the equivalence relation is generated by $u \in \Delta(T') \sim v \in \Delta(T)$ if $T' \subset T$ and $v(t) = u(t)$ for $t \in T'$, $v(t) = 0$ for $t \in T \smallsetminus T'$.
There is a map of sets

$$\coprod_{T\in\mathcal{S}} \Delta(T) \longrightarrow |V|_{\mathcal{S}}$$

which is equal to the inclusion $\Delta(T) \to |V|_{\mathcal{S}}$ on each $\Delta(T)$. That map clearly determines a *bijective* map

$$|Y| = \Big(\coprod_{T\in\mathcal{S}} \Delta(T) \Big) \Big/ \sim \qquad \longrightarrow \qquad |V|_{\mathcal{S}} \ .$$

By our definition of the topology on $|V|_{\mathcal{S}}$, a subset of $|V|_{\mathcal{S}}$ is open if and only if its preimage in $\coprod_{T\in\mathcal{S}} \Delta(S)$ is open; and by our definition of the topology in $|Y|$, that happens if and only if its preimage in $|Y|$ is open. So that bijective map from $|Y|$ to $|V|_{\mathcal{S}}$ is a homeomorphism.

LEMMA 9.3.2. *Let* Y *be any semi-simplicial set. For every element* a *of* $|Y|$ *there exist unique* $m \geq 0$ *and* $(z,w) \in Y_m \times \Delta^m$ *such that* $a = c_z(w)$ *and* w *is in the "interior" of* Δ^m*, that is, the coordinates* $w_0, w_1, \ldots, w_m$ *are all strictly positive.*
Furthermore, if $a = c_x(u)$ *for some* $(x,u) \in Y_k \times \Delta^k$*, then there is a unique order-preserving injective* $f\colon \{0,1,\ldots,m\} \to \{0,1,2,\ldots,k\}$ *such that* $f^*(x) = z$ *and* $f_*(w) = u$*, for the above-mentioned* $(z,w) \in Y_m \times \Delta^m$ *with* $w_0, w_1, \ldots, w_m > 0$.

PROOF. Let us call such a pair (z,w) with $a = c_z(w)$ a *reduced presentation* of a; the condition is that all coordinates of w must be positive. More generally we say that (x,u) is a *presentation* of a if $(x,u) \in Y_k \times \Delta^k$ for some $k \geq 0$ and $a = c_x(u)$. First we show that a admits a reduced presentation and then we show uniqueness.
We know that $a = c_x(u)$ for some $(x,u) \in Y_k \times \Delta^k$. Some of the coordinates $u_0, \ldots, u_k$ can be zero (not all, since their sum is 1). Suppose that $m+1$ of them are nonzero. Let $f\colon \{0,1,\ldots,m\} \to \{0,1,\ldots,k\}$ be the unique order-preserving map such that $u_{f(j)} \neq 0$ for $j = 0,1,2,\ldots,m$. Then $a = c_z(w)$ where $z = f^*(x)$ and $w \in \Delta^m$ with coordinates $w_j = u_{f(j)}$. (Note that $f_*(w) = u$.) So (z,w) is a reduced presentation of a.
We have also shown that any presentation (x,u) of a (whether reduced or not) determines a reduced presentation. Namely, there exist unique m, f and $w \in \Delta^m$ such that $v = f_*(w)$ for some $w \in \Delta^m$ with all $w_i > 0$; then $(f^*(x), w)$ is a reduced presentation of a.
It remains to show that if a has two presentations, say $(x,u) \in Y_k \times \Delta^k$ and $(y,v) \in Y_\ell \times \Delta^\ell$, then they determine the *same* reduced representation of a. If indeed $a = c_x(u) = c_y(v)$ then $\bar{c}_x(u)$ and $\bar{c}_y(v)$ are equivalent, and so (recalling how that equivalence relation was defined) we find that there is no loss of generality in assuming that $x = g^*(y)$ and $v = g_*(u)$ for some order-preserving injective $g\colon \{0,1,\ldots,k\} \to \{0,1,\ldots,\ell\}$. Now determine the unique m and order-preserving injective $f\colon \{0,1,\ldots,m\} \to \{0,1,\ldots,k\}$ such that $u = f_*(w)$ where $w \in \Delta^m$ and all $w_i > 0$. Then we also have $v = g_*(u) = g_*(f_*(w)) = (g \circ f)_*(w)$ and it follows that we get the same reduced presentation, $(f^*(x), w) = ((g \circ f)^*(y), w)$, in both cases. □

COROLLARY 9.3.3. *The space* $|Y|$ *is a Hausdorff space.*

PROOF. For $a \in Y$ with reduced presentation $(z,w) \in Y_m \times \Delta^m$ and $\varepsilon > 0$, define $N(a,\varepsilon) \subset |Y|$ as follows. It consists of all $b \in |Y|$ with reduced presentation $(x,u) \in Y_k \times \Delta^k$ such that there exists an order-preserving injective $f\colon \{0,1,\ldots,m\} \to \{0,1,\ldots,k\}$ for which $f^*(x) = z$ and $f_*(w)$ is ε-close to u, that is, the maximum of the numbers $|w_{f(j)} - u_j|$ is $< \varepsilon$. From the definitions, $N(a,\varepsilon)$ is open in $|Y|$; so it is a neighborhood of a.
Let $a' \in |Y|$ be another element, with reduced presentation $(y,v) \in Y_n \times \Delta^n$. We assume $a \neq a'$ and proceed to show that $N(a',\varepsilon) \cap N(a,\varepsilon) = \emptyset$ if ε is small enough. More precisely, we take ε to be less than half the minimum of the coordinates of v and w; and if it should happen that $m = n$ and $y = z$, then we know $v, w \in \Delta^m$ but $v \neq w$, and we take ε to be less than half the maximum of the $|v_j - w_j|$ as well. Now suppose for a contradiction that $b \in N(a,\varepsilon) \cap N(a',\varepsilon)$ and that b has reduced presentation $(x,u) \in Y_k \times \Delta^k$. Then there exist order-preserving injective $f\colon \{0,1,\ldots,m\} \to \{0,1,\ldots,k\}$ and $g\colon \{0,1,\ldots,n\} \to \{0,1,\ldots,k\}$ such that $f^*(x) = z$, $g^*(x) = y$ and $f_*(w), g_*(v)$ are both ε-close to u in Δ^k. Then $f_*(w)$ is 2ε-close to $g_*(v)$ in Δ^k, and now we can deduce that $m = n$ and $f = g$. (Otherwise there is some $j \in \{0,1,\ldots,k\}$ which is in the image of g but not in the image of f, or vice versa, and then the j-th coordinate of $g_*(w)$ differs by more than 2ε from the j-th coordinate of $f_*(v)$.) Therefore $z = f^*(x) = g^*(x) = y$ and so a has reduced presentation (z,w) while a' has reduced presentation (z,v), with

$v, w \in \Delta^m$ and the same $z \in Y_m$. It follows that v and w are already 2ε-close in Δ^m. This contradicts our choice of ε. □

REMARK 9.3.4. In the proof above, and in a similar proof in the previous section, arguments involving distances make an appearance, suggesting that we have a metrizable situation. To explain what is going on let me return to the situation of a vertex scheme $(V, \mathcal{S})$ with simplicial complex $|V|_{\mathcal{S}}$, which is easier to understand. A metric on the set $|V|_{\mathcal{S}}$ can be introduced for example by $d(f,g) = (\sum_v (f(v)-g(v))^2)^{1/2}$ or $d(f,g) = \sum_v |f(v)-g(v)|$. Here we insist/remember that elements of $|V|_{\mathcal{S}}$ are functions $f, g, \ldots : V \to [0,1]$ subject to some conditions. The sums in the formulas for $d(f,g)$ are finite, even though V might not be a finite set. It is not hard to show that the two formulas for $d(f,g)$, although different as metrics, determine the same topology. However the topology on $|V|_{\mathcal{S}}$ that we have previously decreed (let me call it the *weak* topology) is not in all cases the same as that metric topology. Every subset of $|V|_{\mathcal{S}}$ which is open in the metric topology is also open in the weak topology. But the weak topology can have more open sets. (We reasoned that the weak topology is Hausdorff because it has all the open sets that the metric topology has, and perhaps a few more, and the metric topology is certainly Hausdorff.) In the case where V is finite, weak topology and metric topology on $|V|_{\mathcal{S}}$ coincide. (Exercise.)

9.4. A more economical but less conceptual definition of *semi-simplicial set*

Every injective order-preserving map from $[k] = \{0, 1, \ldots, k\}$ to $[\ell] = \{0, 1, \ldots, \ell\}$ is a composition of $\ell - k$ injective order preserving maps

$$[m-1] \longrightarrow [m]$$

where $k < m \leq \ell$. It is easy to list the injective order-preserving maps from $[m-1]$ to $[m]$; there is one such map f_i for every $i \in [m]$, characterized by the property that the image of f_i is

$$[m] \smallsetminus \{i\}.$$

(This f_i really depends on two parameters, m and i. Perhaps we ought to write $f_{m,i}$, but it is often practical to suppress the m subscript.) We have the important relations

$$f_i f_j = f_j f_{i-1} \quad \text{if } j < i \tag{9.4.1}$$

(You are allowed to read this from left to right or from right to left! It is therefore a formal consequence that $f_i f_j = f_{j+1} f_i$ when $j \geq i$.) These *generators and relations* suffice to describe the category $\mathcal{C}$ (lecture notes week 11) whose objects are the sets $[k] = \{0, 1, \ldots, k\}$ for $k \geq 0$ and whose morphisms are the order-preserving injective maps between those sets. In other words, the structure of $\mathcal{C}$ as a category is pinned down if we say that it has objects $[k]$ for $k \geq 0$ and that, for every $k > 0$ and $i \in \{0, 1, \ldots, k\}$, there are certain morphisms $f_i \colon [k-1] \to [k]$ which, under composition when it is applicable, satisfy the relations (9.4.1). Prove it!
Consequently a semi-simplicial set Y, which is a contravariant functor from $\mathcal{C}$ to spaces, can also be described as a sequence of sets $Y_0, Y_1, Y_2, \ldots$ and maps

$$d_i \colon Y_k \to Y_{k-1}$$

which are subject to the relations

$$d_j d_i = d_{i-1} d_j \quad \text{if } j < i \tag{9.4.2}$$

Here $d_i \colon Y_k \to Y_{k-1}$ denotes the map induced by $f_i \colon [k-1] \to [k]$, whenever $0 \leq i \leq k$. Because of contravariance, we have had to reverse the order of composition in translating relations (9.4.1) to obtain relations (9.4.2).

CHAPTER 10

Chain complexes and their homology groups

10.1. Chain complexes

Definition 10.1.1. A *chain complex* C consists a sequence of abelian groups $C_0, C_1, C_2, \ldots$ and homomorphisms $\partial_n \colon C_n \to C_{n-1}$ for $n > 0$ such that

$$\partial_n \circ \partial_{n-1} = 0$$

for all $n > 1$. We say that C_n is the n-th chain group of the chain complex C. The homomorphisms ∂_n are called *differentials* or (sometimes) *boundary operators*.

Think of a chain complex as a diagram

$$C_0 \xleftarrow{\partial_1} C_1 \xleftarrow{\partial_2} C_2 \xleftarrow{\partial_3} C_3 \xleftarrow{\partial_4} C_4 \xleftarrow{\partial_5} \cdots$$

where the composition of any two consecutive arrows is zero. It is very common to drop the subscript n in ∂_n. So a more standard picture of a chain complex looks like

$$C_0 \xleftarrow{\partial} C_1 \xleftarrow{\partial} C_2 \xleftarrow{\partial} C_3 \xleftarrow{\partial} C_4 \xleftarrow{\partial} \cdots,$$

and we just write $\partial\partial = 0$ instead of writing $\partial_n \circ \partial_{n-1} = 0$ for all $n > 1$.

Unsurprisingly, chain complexes are the objects of a category. The morphisms in that category are called *chain maps*. A chain map from a chain complex C to a chain complex D is a sequence of homomorphisms

$$f_n \colon C_n \to D_n$$

(for $n \geq 0$) making the diagram

(10.1.2)
$$\begin{tikzcd}
C_0 & C_1 \arrow[l, "\partial"'] \arrow[d, "f_1"] & C_2 \arrow[l, "\partial"'] \arrow[d, "f_2"] & C_3 \arrow[l, "\partial"'] \arrow[d, "f_3"] & C_4 \arrow[l, "\partial"'] \arrow[d, "f_4"] & \cdots \arrow[l, "\partial"'] \\
D_0 \arrow[from=u, "f_0"] & D_1 \arrow[l, "\partial"'] & D_2 \arrow[l, "\partial"'] & D_3 \arrow[l, "\partial"'] & D_4 \arrow[l, "\partial"'] & \cdots \arrow[l, "\partial"']
\end{tikzcd}$$

commutative; in other words $\partial \circ f_n = f_{n-1} \circ \partial$ for all $n > 0$. The preferred shorthand notation for such a morphism is $f \colon C \to D$.

Remark 10.1.3. Some would say that what has been defined above is a chain complex *graded over the non-negative integers*. There are also chain complexes graded over the integers, which look like

$$\cdots \xleftarrow{\partial_{-2}} C_{-2} \xleftarrow{\partial_{-1}} C_{-1} \xleftarrow{\partial_0} C_0 \xleftarrow{\partial_1} C_1 \xleftarrow{\partial_2} C_2 \xleftarrow{\partial_3} C_3 \xleftarrow{\partial_4} C_4 \xleftarrow{\partial_5} \cdots$$

If a chain complex C comes along being graded over the non-negative integers, then it is often a good idea to view it as a chain complex graded over the integers by setting $C_n = 0$ for $n < 0$.

DEFINITION 10.1.4. For $n \geq 0$, the n*-th homology group* $H_n(C)$ of a chain complex C is the (group-theoretic) quotient

$$H_n(C) = \frac{\ker[\partial_n \colon C_n \to C_{n-1}]}{\operatorname{im}[\partial_{n+1} \colon C_{n+1} \to C_n]} .$$

Note that the kernel of ∂_n and the image of ∂_{n+1} are both subgroups of C_n, and the kernel of ∂_n contains the image of ∂_{n+1}.

When $n = 0$, we need to make sense of $\ker[\partial_0 \colon C_0 \to C_{-1}]$. In agreement with the convention that chain complexes graded over the non-negative integers can be viewed as chain complexes graded over the integers, we take the view that this is all of C_0 and so

$$H_0(C) = \frac{C_0}{\operatorname{im}[\partial_1 \colon C_1 \to C_0]}$$

REMARK 10.1.5. For fixed n the rule $C \mapsto H_n(C)$ is a (covariant) functor from the category of chain complexes to the category of abelian groups. Let $f \colon C \to D$ be a chain map, consisting of homomorphisms $f_n \colon C_n \to D_n$ for $n \geq 0$. Then f_n takes $\ker[\partial_n \colon C_n \to C_{n-1}]$ to $\ker[\partial_n \colon D_n \to D_{n-1}]$, and takes $\operatorname{im}[\partial_{n+1} \colon C_{n+1} \to C_n]$ to $\operatorname{im}[\partial_{n+1} \colon C_{n+1} \to C_n]$, and so determines a homomorphism

$$H_n(C) \to H_n(D) .$$

10.2. The combinatorial chain complex of a semi-simplicial set

A semi-simplicial set Y determines a chain complex $C(Y)$, the *combinatorial chain complex of* Y, in the following way.

The chain group $C(Y)_n$ is defined to be the direct sum of copies of $\mathbb{Z}$, one copy for each $z \in Y_n$. We can write

$$C(Y)_n = \bigoplus_{z \in Y_n} \mathbb{Z} .$$

(It is also customary to say that $C(Y)_n$ is the *free abelian group generated by the set* Y_n.) If we agree to denote the element "1" in the summand corresponding to $z \in Y_n$ by $\langle z \rangle$, then we can describe elements of $C(Y)_n$ as linear combinations

$$\sum_{z \in Y_n} a_z \cdot \langle z \rangle$$

where the coefficients a_z are integers (and the sum is understood to be finite, that is, $a_z \neq 0$ for only finitely many $z \in Y_n$.) The differential or boundary operator

$$\partial_n \colon C(Y)_n \longrightarrow C(Y)_{n-1}$$

is defined by

$$\langle z \rangle \quad \mapsto \quad \sum_{j=0}^{n} (-1)^j \langle d_j z \rangle \in C(Y)_{n-1}$$

where $d_j \colon Y_n \to Y_{n-1}$ is the face operator discussed previously, corresponding to the unique monotone injective map from $\{0, 1, \ldots, n-1\}$ to $\{0, 1, \ldots, n\}$ which has image $\{0, 1, \ldots, n\} \smallsetminus \{j\}$.

Now we need to show that $\partial_{n-1}\partial_n = 0$ for all $n > 1$. This is a straightforward calculation based on the relations (9.4.2).

$$\partial_{n-1}(\partial_n(\langle z \rangle)) = \cdots = \sum_{j=0}^{n-1} \sum_{i=0}^{n} (-1)^{i+j} \langle d_j d_i z \rangle$$

The double sum can be split into two parts, one part comprising the terms $(-1)^{i+j}\langle d_j d_i z\rangle$ where $i < j$ and the other comprising the terms $(-1)^{i+j}\langle d_j d_i z\rangle$ where $i \geq j$. Each summand $(-1)^{k+\ell}\langle d_\ell d_k z\rangle$ in the first part part cancels exactly one in the other part, $(-1)^{\ell+k-1}\langle d_{k-1} d_\ell z\rangle$, where we are using (9.4.2).

EXAMPLE 10.2.1. The projective plane $\mathbb{R}P^2$ can be described as $|Y|$ for a semi-simplicial set Y. To construct this we start with a simplicial complex or vertex scheme $(V, \mathcal{S})$ describing the upper hemisphere of S^2. Picture:

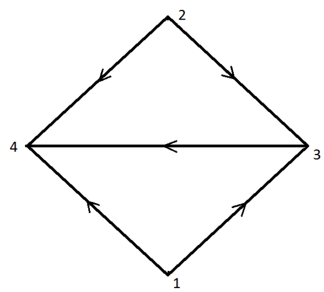

Therefore $V = \{1,2,3,4\}$ and

$$\mathcal{S} = \Big\{\ \{1,3,4\},\{2,3,4\},\{1,3\},\{1,4\},\{2,3\},\{2,4\},\{3,4\},\{1\},\{2\},\{3\},\{4\}\ \Big\}.$$

The vertex set V is already ordered and we can immediately pass to a semi-simplicial set X where

$$\begin{aligned} X_2 &= \{(134),(234)\} \\ X_1 &= \{(13),(14),(23),(24),(34)\} \\ X_0 &= \{(1),(2),(3),(4)\}. \end{aligned}$$

(To clarify the improvised notation: X_2 has two elements, X_1 has five elements and X_0 has four elements.) For $n > 2$ we set $X_n = \emptyset$. The operators d_i are defined by omitting the digit in position i (but we label the positions with integers from 0 upwards), so that for example

$$d_0(134) = (34),\ d_1(134) = (14),\ d_2(134) = (13),\ d_0(13) = (3).$$

By a certain proposition we have $|X| = |V|_{\mathcal{S}}$, which we think of as the upper hemisphere of S^2, but now we want to identify opposite points on the boundary (=equator). In the semi-simplicial set code, this means that we enforce

$$(14) \sim (23), \qquad (13) \sim (24), \qquad (1) \sim (2), \qquad (3) \sim (4).$$

(NB: it seems to me that I had to think fairly hard to get the numbering of vertices right, so that by making these identifications we do in fact identify opposite points on the equator when we pass to geometric realizations.) In this way we get a new semi-simplicial set Y where

$$\begin{aligned} Y_2 &= \{(134),(234)\} \\ Y_1 &= \{(13)=(24),(14)=(23),(34)\} \\ Y_0 &= \{(1)=(2),(3)=(4)\}. \end{aligned}$$

(To clarify the very improvised notation: Y_2 has two elements, Y_1 has three elements and Y_0 has two elements.) For $n > 2$ we still have $Y_n \geq 0$. The operators d_i are defined by omitting the digit in position i (but we label the positions with integers from 0 upwards). Now we are ready to follow instructions above to make the chain complex $C(Y)$:

$$\mathbb{Z}^2 \leftarrow \mathbb{Z}^3 \leftarrow \mathbb{Z}^2 \leftarrow 0 \leftarrow 0 \leftarrow \cdots$$

The boundary operators can be described as matrices with integer entries: a 2×3 matrix for ∂_1 and a 3×2 matrix for ∂_2. (For the matrix descriptions we need and we have *ordered* bases: so the columns for example in the 2×3 matrix are labeled with the three elements of Y_1, in the order in which they are listed above.) The differential ∂_1 is given by

$$(13) \mapsto (3)-(1), \quad (14) \mapsto (4)-(1)=(3)-(1), \quad (34) \mapsto (4)-(3)=0$$

which in matrix form is

$$\begin{bmatrix} -1 & -1 & 0 \\ 1 & 1 & 0 \end{bmatrix}$$

The differential ∂_2 is given by

$$(134) \mapsto (34)-(14)+(13), \qquad (234) \mapsto (34)-(24)+(23)=(34)-(13)+(14)$$

which in matrix form is

$$\begin{bmatrix} 1 & -1 \\ -1 & 1 \\ 1 & 1 \end{bmatrix}$$

Note that the product of the two matrices (in the correct order) *is* zero, confirming that $\partial_1\partial_2 = 0$, as it should be. It is easy to see using the matrix description that the image of ∂_1 consists of all elements in $\mathbb{Z}^2$ which have coordinate sum equal to 0, and it follows immediately that

$$H_0(C(Y)) = \operatorname{coker}[\partial_1] \cong \mathbb{Z} .$$

Determining $H_1(C(Y))$ is not straightforward. The kernel of ∂_1 consists of all elements of $\mathbb{Z}^3$ which are perpendicular to the row vector $\begin{bmatrix} 1 & 1 & 0 \end{bmatrix}$, and so $\ker[\partial_1]$ has a complement in $\mathbb{Z}^3$, the subgroup A of $\mathbb{Z}^3$ spanned by the element

$$\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} .$$

Therefore

$$H_1(C(Y) = \frac{\ker[\partial_1]}{\operatorname{im}[\partial_2]} \cong \frac{\mathbb{Z}^3/A}{\operatorname{im}[p\circ\partial_2]}$$

where $p\colon \mathbb{Z}^3 \to \mathbb{Z}^3/A$ is the projection. We can identify $\mathbb{Z}^3/A$ with $\mathbb{Z}^2$ in the obvious manner. Then we can describe $p\circ\partial_2$ as a homomorphism from $\mathbb{Z}^2$ to $\mathbb{Z}^2$ by a square matrix

$$\begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix}$$

obtained by deleting the top row in the matrix description of ∂_2 . This 2×2 matrix has determinant -2 and so the cokernel of the homomorphism $\mathbb{Z}^2 \to \mathbb{Z}^2$ which it describes has order $|-2| = 2$. It must be cyclic of order 2. Similar reasoning shows that ∂_2 is injective. Indeed $p\circ\partial_2$ is injective, because you can use Kramer's rule and the nonzero determinant to recover elements in the source from their values in the target. Therefore

$$H_1(C(Y)) \cong \mathbb{Z}/2, \qquad H_2(C(Y)) = 0$$

and clearly $H_n(C(Y)) = 0$ for all $n > 2$ as well.
Let's not fail to observe that the groups $H_n(C(Y))$ coincide with the groups $H_n(\mathbb{R}P^2) = H_n(|Y|)$, for every $n \geq 0$. This is not an accident, as we shall see in the next section.

CHAPTER 11

Homology of the geometric realization

11.1. A formula

THEOREM 11.1.1. *For every semi-simplicial set* Y *and every* $n \geq 0$*, there is an isomorphism of the* n*-th homology group of the chain complex* $\mathsf{C}(\mathsf{Y})$ *with the* n*-th homology group of* $|\mathsf{Y}|$*:*

$$H_n(C(Y)) \cong H_n(|Y|) \ .$$

This can be stated in a slightly more precise way: ... *for every* $n \geq 0$ *there is a natural isomorphism*

$$H_n(C(Y)) \xrightarrow{\cong} H_n(|Y|).$$

For this stronger form of the statement, we need to be aware that semi-simplicial sets form a category and that the rule $Y \mapsto C(Y)$ is a functor. How do semi-simplicial sets form a category? One of several definitions of *semi-simplicial set* said that such a thing is a contravariant functor from a certain category $\mathcal{C}$ (with objects $[k] = \{0, 1, \dots, k\}$ for $k \geq 0$ and with injective order-preserving maps as morphisms) to the category of sets. Therefore it is very reasonable to say that a morphism from one simplicial set X to another, Y, is a natural transformation $X \Rightarrow Y$ between such contravariant functors. In detail, this means that a morphism $\alpha\colon X \to Y$ is given by a sequence $(\alpha_k)_{k \geq 0}$ where each $\alpha_k\colon X_k \to Y_k$ is a map (of sets). The maps α_k are together subject to a strong condition: for each order-preserving injective map $g\colon [k] \longrightarrow [\ell]$ the diagram

$$\begin{array}{ccc} X_\ell & \xrightarrow{\alpha_\ell} & Y_\ell \\ \downarrow{\scriptstyle g^*} & & \downarrow{\scriptstyle g^*} \\ X_k & \xrightarrow{\alpha_k} & Y_k \end{array}$$

is commutative, $g^*\alpha_\ell = \alpha_k g^*$. (It is enough to verify this when $k = \ell - 1$, in which case the custom is to write d_i for g^*, where i is the element of $[\ell]$ not in the image of g. Then the condition becomes $d_i\alpha_\ell = \alpha_{\ell-1}d_i\colon X_\ell \to Y_{\ell-1}$.)
How is the rule $Y \mapsto C(Y)$ a functor from the category of semi-simplicial sets to the category of chain complexes? This is now clear: a morphism $\alpha\colon X \to Y$ of simplicial sets defines a chain map $C(\alpha)\colon C(X) \to C(Y)$ given by homomorphisms

$$C(\alpha_k)\colon C(X)_k \longrightarrow C(Y)_k$$

defined by $\langle x \rangle \mapsto \langle \alpha_k(x) \rangle$, where $\langle x \rangle$ is the generator corresponding to an element $x \in X_k$. The condition for a chain map, $\partial_k \circ C(\alpha_k) = C(\alpha_{k-1}) \circ \partial_k$, is satisfied because ∂_k is defined in terms of the d_i , which α respects.

11.2. Skeletons

The naturality part of theorem 11.1.1 is important because it will help us to prove the theorem. The proof occupies most of this chapter. We are going to proceed by induction on *skeletons*.

DEFINITION 11.2.1. The n-*skeleton* of a semi-simplicial set Y is the semi-simplicial set $Y^{\leq n}$ defined by

$$Y_k^{\leq n} = \begin{cases} Y_k & \text{if } k \leq n \\ \emptyset & \text{if } k > n \end{cases}$$

and face operators $Y_\ell^{\leq n} \to Y_k^{\leq n}$ defined like the corresponding ones in Y if $k, \ell \leq n$. (Thus $Y^{\leq n}$ is a semi-simplicial *sub*set of Y. For induction purposes it is useful to allow $n = -1$; we define $Y_k^{\leq -1} = \emptyset$ for all k.)

LEMMA 11.2.2. Let X be a semi-simplicial subset of Y, so that $X_n \subset Y_n$ for all $n \geq 0$ and the face operators $f^*\colon Y_\ell \to Y_k$ take X_ℓ to X_k. Then the map $|X| \to |Y|$ induced by the inclusion $X \to Y$ has closed image and is a homeomorphism onto its image (so that it can be viewed as the inclusion of a closed subspace). In particular, $|Y^{\leq n}|$ can be identified with a closed subspace of $|Y|$.

PROOF. It follows from lemma 9.2 for example that $|X| \to |Y|$ is injective. The image is a closed subspace of $|Y|$ by the definition of the topology in $|Y|$. (This is easier to see if we reformulate that definition as follows: *a subset A of $|Y|$ is closed if its preimage under each of the characteristic maps $c_y\colon \Delta^k \to |Y|$ is closed in Δ^k, where $y \in Y_k$.* When $A = |X|$, the preimage of A under c_y is the union of some faces of Δ^k. This is obviously a closed subset of Δ^k.) By the same reasoning, applied to $|X|$ and to $|Y|$, we see that a subset of $|X|$ is closed if and only if its image in $|Y|$ is closed. □

LEMMA 11.2.3. *Let X and Y be spaces, X compact Hausdorff. For any mapping cycle α from X to Y, there exists a compact subspace $K \subset Y$ such that α factors through K.*

PROOF. Choose a finite open cover $(U_i)_{i=1,2,\dots,k}$ of X such that α restricted to any U_i can be written as a formal linear combination, with integer coefficients, of (finitely many) continuous maps: $\sum_j a_{ij} f_{ij}$ where $a_{ij} \in \mathbb{Z}$ and the $f_{ij}\colon U_i \to Y$ are continuous maps. Choose another finite open cover $(V_i)_{i=1,2,\dots,k}$ of X such that the closure $\bar{V}_i$ of V_i in X is contained in U_i. (This is possible because X is compact Hausdorff.) Let $K \subset Y$ be the union of the finitely many compact sets $f_{ij}(\bar{V}_i)$. □

LEMMA 11.2.4. *Let Y be a semi-simplicial set. For any compact subset K of $|Y|$, there exists a finite semi-simplicial subset X of Y such that $K \subset |X| \subset |Y|$.*

PROOF. Suppose that K is not contained in any subspace of the form $|X|$, where X is a finite semi-simplicial subset of Y. Then it must be possible to choose elements $z_j \in K$ for $j = 1, 2, 3, 4, \dots$ (infinitely many) such that z_j has reduced presentation of the form $c_{y_j}(u)$, where the $y_j \in Y_{n_j}$ are all *distinct* (and $c_{y_j}\colon \Delta^{n_j} \to |Y|$ is the characteristic map associated with y_j, and we are assuming that $u \in \Delta^{n_j}$ does not belong to the boundary). Let

$$W_i = |Y| \smallsetminus \{z_i, z_{i+1}, \dots\}\ .$$

By construction and by the definition of the topology in $|Y|$, the sets W_i are open in $|Y|$. Clearly $W_i \subset W_{i+1}$ and $\bigcup_i W_i = |Y|$. Therefore the union of all the sets $W_i \cap K$ is all of K. We have found an open covering of K which does not have a finite subcover; contradiction. □

COROLLARY 11.2.5. *For every element* u *of* $H_k(|Y|)$*, there exists a finite semi-simplicial subset* X *of* Y *such that* u *is in the image of the inclusion-induced homomorphism*

$$H_k(|X|) \to H_k(|Y|).$$

If X *and* X' *are finite semi-simplicial subsets of* Y *and* $v \in H_k(|X|)$*,* $w \in H_k(|X'|)$ *have the same image in* $H_k(|Y|)$*, then there exists another finite semi-simplicial subset* X'' *of* Y *such that* $X \subset X''$*,* $X' \subset X''$ *and* v, w *have the same image in* $H_k(|X''|)$*.*

PROOF. Given $u \in H_k(|Y|)$, represent it by a mapping cycle from S^k to $|Y|$. By lemmas 11.2.3 and 11.2.4, that mapping cycle factors through $|X|$ for some finite semi-simplicial subset X of Y. This proves the first part. Now suppose that $v \in H_k(|X|)$ and $w \in H_k(|X'|)$ have the same image in $H_k(|Y|)$, for finite semi-simplicial subsets X and X' of Y. Recall that $H_k(|X|)$ can be defined as the cokernel of $[[\star, |X|]] \to [[S^k, |X|]]$ or alternatively as the kernel of $[[S^k, |X|]] \to [[\star, |X|]]$, where the first homomorphism is induced by the projection $S^k \to \star$ and the other is induced by the inclusion of the base point in the sphere, $\star \to S^k$. Here the second definition is more useful, so we represent v by a mapping cycle $\alpha\colon S^k \to |X|$ such that the composition of α with $\star \to S^k$ is homotopic to 0 as a mapping cycle. In the same way, we represent w by a mapping cycle $\beta\colon S^k \to |X'|$ such that the composition of β with $\star \to S^k$ is homotopic to 0 as a mapping cycle. Since v and w have the same image in $H_k(|Y|)$, there exists a suitable homotopy, i.e., a mapping cycle $\gamma\colon S^k \times [0,1] \to |Y|$ which restricts to α on $S^k \times \{0\} \cong S^k$ and to β on $S^k \times \{1\} \cong S^k$. By lemma 11.2.3 and 11.2.4, that mapping cycle γ factors through $|X''|$ for some finite semi-simplicial subset X'' of Y, and we can enlarge X if necessary to ensure $X'' \supset X$ and $X'' \supset X'$. Then clearly v and w have the same image in $H_k(|X''|)$. □

LEMMA 11.2.6. *For every element* u *of* $H_k(|Y|)$*, there exists* $n \geq 0$ *such that* u *is in the image of the inclusion-induced homomorphism from* $H_k(|Y^{\leq n}|)$ *to* $H_k(|Y|)$*. Moreover, if* $v \in H_k(|Y^{\leq m}|)$*,* $w \in H_k(|Y^{\leq n}|)$ *have the same image in* $H_k(|Y|)$*, then there exists* $\ell \geq m, n$ *such that* v *and* w *have the same image already in* $H_k(|Y^{\leq \ell}|)$*.*

This can be deduced from corollary 11.2.5, or proved in the same way. — Let's now ask how $H_*(|Y^{\leq n-1}|)$ is related to $H_*(|Y^{\leq n}|)$.

LEMMA 11.2.7. *The homomorphism*

$$H_k(|Y^{\leq n-1}|) \to H_k(|Y^{\leq n}|)$$

induced by the inclusion of $|Y^{\leq n-1}|$ *in* $|Y^{\leq n}|$ *is an isomorphism for* $k < n-1$*, whereas* $H_k(|Y^{\leq n-1}|) = 0$ *for* $k > n-1$*. For* $k = n-1$ *this homomorphism is part of an exact sequence*

$$0 \longrightarrow H_n(|Y^{\leq n}|) \xrightarrow{\gamma_n} \bigoplus_{x \in Y_n} \mathbb{Z} \xrightarrow{\beta_n} H_{n-1}(|Y^{\leq n-1}|) \longrightarrow H_{n-1}(|Y^{\leq n}|) \longrightarrow 0\,.$$

PROOF. We can assume $n \geq 1$. Let $W = |Y^{\leq n}| \smallsetminus |Y^{\leq n-1}|$. In other words, W consists of all points which can be written in the form $c_y(u)$ for some $y \in Y_n$ (with characteristic map $c_y\colon \Delta^n \to |Y^{\leq n}|$) and some $u \in \Delta^n$ whose barycentric coordinates are all nonzero: $u_0, \dots, u_n > 0$. Let V be the subset of $|Y^{\leq n}|$ obtained by taking out all points of the form $c_y(b)$ where $y \in Y_n$ and $b \in \Delta^n$ is the barycenter, that is, $b_0 = b_1 = \cdots = b_n = 1/(n+1)$. Then V and W are open subsets and $V \cup W = |Y^{\leq n}|$. Clearly W is homeomorphic to a disjoint union of copies of $\mathbb{R}^n$, one copy for each $y \in Y_n$. Therefore clearly $V \cap W$ (viewed as a subspace of W, if you wish) is homeomorphic to a disjoint union of copies of $\mathbb{R}^n \smallsetminus \{0\}$, and consequently homotopy equivalent to a disjoint union of copies of S^{n_1}, one copy for each $y \in Y_n$. This means that we know the homology groups of W and of $V \cap W$.

Regarding V, we show that the inclusion $\iota\colon |Y^{\leq n-1}| \to V$ is a homotopy equivalence. A map in the opposite direction, $r\colon V \to |Y^{\leq n-1}|$, is given by $r(z) = z$ if $r \notin V \cap W$ and

$$r(c_y(u)) \mapsto c_y(\rho(u))$$

where we assume $y \in Y_n$ and $u \in \Delta^n$ not equal to the barycenter b , and $\rho(u)$ is the point in the boundary of Δ^n where the straight line through b and u meets the boundary. Then the composition $r \circ \iota$ is the identity on $|Y^{\leq n-1}|$, and $\iota \circ r$ is homotopic to the identity on V.

Now we are in a good position to understand the Mayer-Vietoris sequence:

$$\cdots \to H_k(V \cap W) \to H_k(V) \oplus H_k(W) \to H_k(V \cup W) \to H_{k-1}(V \cap W) \to \cdots$$

Since $V \cap W$ is homotopy equivalent to a disjoint union of copies of S^{n-1}, the homology groups $H_k(V \cap W)$ are nonzero only for $k = n-1$ and $k = 0$. The homology groups $H_k(W)$ are nonzero only for $k = 0$. It is routine to make the following deduction from the exactness of the Mayer-Vietoris sequence:

The inclusion $V \to V \cup W$ induces an isomorphism

$$H_k(V) \to H_k(V \cup W)$$

when $k < n-1$. (This is trivially true when $n = 1$. In the cases $n > 1$, the inclusion $V \cap W \to W$ induces an isomorphism from $H_0(V \cap W)$ to $H_0(W)$, and this should be used, too.)

Exactness of the Mayer-Vietoris sequence also permits us to show by induction on n that

$H_k(|Y^{\leq n}|) = H_k(V \cup W)$ is zero for $k > n$.

Therefore, if $n > 1$, the interesting part of the Mayer-Vietoris sequence (where k is close to n) is an exact sequence

$$0 \to H_n(V \cup W) \to \bigoplus_{x \in Y_n} \mathbb{Z} \to H_{n-1}(V) \to H_{n-1}(V \cup W) \to 0$$

where the 0 on the right is justified as the kernel of the (injective) homomorphism $H_{n-2}(V \cap W) \to H_{n-2}(V) \oplus H_{n-2}(W)$. If $n = 1$ we have instead an exact sequence

$$0 \to H_1(V \cup W) \to \bigoplus_{x \in Y_n} (\mathbb{Z} \oplus \mathbb{Z}) \to H_0(V) \oplus \bigoplus_{x \in Y_n} \mathbb{Z} \to H_0(V \cup W) \to 0\,.$$

Here the composite homomorphism

$$\bigoplus\nolimits_x (\mathbb{Z} \oplus \mathbb{Z}) \longrightarrow H_0(V) \oplus \bigoplus\nolimits_x \mathbb{Z} \xrightarrow{\text{proj.}} \bigoplus\nolimits_x \mathbb{Z}$$

is onto by inspection and its kernel is the antidiagonal $\bigoplus_x \mathbb{Z}_a$, where $\mathbb{Z}_a \subset \mathbb{Z} \oplus \mathbb{Z}$ consists of all pairs of integers of the form $(r, -r)$. Therefore we can remove some terms and obtain an exact sequence

$$0 \to H_1(V \cup W) \to \bigoplus_{x \in Y_n} \mathbb{Z} \to H_0(V) \to H_0(V \cup W) \to 0$$

in the case $n = 1$, too. Finally, using that $V \cup W = |Y^{\leq n}|$ and $V \simeq |Y^{\leq n-1}|$, we have the exact sequence that we wanted. $\square$

COROLLARY 11.2.8. *The inclusion* $|Y^{\leq n}| \to |Y|$ *induces an isomorphism in* H_k *for* $k < n$, *and a surjection for* $k = n$.

PROOF. We have a sequence of inclusion-induced homomorphisms

$$H_k(|Y^{\leq n}|) \to H_k(|Y^{\leq n+1}|) \to H_k(|Y^{\leq n+2}|) \to \cdots$$

If $k < n$, all homomorphisms in the sequence are isomorphisms by lemma 11.2.7. Combine this with lemma 11.2.6 to deduce that $H_k(|Y^{\leq n}|) \to H_k(|Y|)$ is an isomorphism. If $k = n$, the first homomorphism in the sequence is surjective, the others are isomorphisms. Again combine his with lemma 11.2.6 to deduce that $H_k(|Y^{\leq n}|) \to H_k(|Y|)$ is surjective.

COROLLARY 11.2.9. $H_k(Y) \cong H_k(D(Y))$ *where* $D(Y)$ *is the chain complex defined as follows:*

$$D(Y)_n = C(Y)_n = \bigoplus_{x \in Y_n} \mathbb{Z}$$

for $n \geq 0$ *and the differential* $D(Y)_n \to D(Y)_{n-1}$ *is* $\gamma_{n-1} \circ \beta_n$ *with notation as in lemma 11.2.7.*

(We take the view that γ_0 is defined and is an isomorphism from $H_0(|Y^{\leq 0}|)$ to $\bigoplus_{y \in Y_0} \mathbb{Z}$. This is in good agreement with lemma 11.2.7. Note also that

$$(\gamma_{n-1} \circ \beta_n) \circ (\gamma_n \circ \beta_{n+1}) = 0$$

because already $\beta_n \circ \gamma_n$ is zero, as can be seen in lemma 11.2.7. Therefore $D(Y)$ is indeed a chain complex.)

PROOF. From lemma 11.2.7 we obtain

$$H_n(|Y|) \cong H_n(|Y^{\leq n+1}|) \cong \frac{H_n(|Y^{\leq n}|)}{\mathrm{im}(\beta_{n+1})} \cong^a \frac{\ker(\beta_n)}{\mathrm{im}(\gamma_n \circ \beta_{n+1})} =^b \frac{\ker(\gamma_{n-1} \circ \beta_n)}{\mathrm{im}(\gamma_n \circ \beta_{n+1})}$$

if $n > 0$. The isomorphism with superscript a is obtained by using the injective homomorphism γ_n to identify $H_n(|Y^{\leq n}|)$ with the subgroup $\ker(\beta_n)$ of $\bigoplus_{y \in Y_n} \mathbb{Z}$. The equality with superscript b is based on the observation that γ_{n-1} is injective. For $n = 0$ we get

$$H_0(|Y|) \cong H_0(|Y^{\leq 1}|) \cong \frac{H_0(|Y^{\leq 0}|)}{\mathrm{im}(\beta_1)} \cong^a \frac{\bigoplus_{y \in Y_0} \mathbb{Z}}{\mathrm{im}(\gamma_0 \circ \beta_1)}$$

which is again exactly what we want. □

11.3. Naturality considerations

To finish the proof of theorem 11.1.1 we need to show that the chain complexes $D(Y)$ and $C(Y)$ are the same, or at least isomorphic. We already have $C(Y)_n = D(Y)_n$ by construction. It would be wonderful to know that the boundary homomorphism

$$\gamma_{n-1} \circ \beta_n \colon D(Y)_n \to D(Y)_{n-1}$$

agrees with the boundary homomorphism $C(Y)_n \to C(Y)_{n-1}$ from the definition of the chain complex $C(Y)$.

LEMMA 11.3.1. *These two boundary homomorphisms agree up to a sign; in other words*

$$(\gamma_{n-1} \circ \beta_n)\langle x \rangle = \pm \sum_{i=0}^{n} (-1)^i \langle d_i x \rangle$$

for $x \in Y_n$ *with corresponding basis element* $\langle x \rangle \in D(Y)_n$*. The sign* $\pm$ *depends on* n*, but not on* Y *or* $x \in Y_n$*.*

PROOF. We begin with the important observation that the construction of $\gamma_{n-1} \circ \beta_n$ was *natural*. To be more explicit, a morphism $f\colon X \to Y$ of semi-simplicial sets determines (for any fixed $n \geq 0$) a homomorphism from $D(X)_n$ to $D(Y)_n$ given by $\langle x \rangle \mapsto \langle f(x) \rangle$ for $x \in X_n$, and the resulting diagram

$$\begin{tikzcd} D(X)_n \arrow[r, "\gamma_{n-1}\circ\beta_n"] \arrow[d] & D(X)_{n-1} \arrow[d] \\ D(Y)_n \arrow[r, "\gamma_{n-1}\circ\beta_n"] & D(Y)_{n-1} \end{tikzcd} \tag{11.3.2}$$

is *commutative*. (So $Y \mapsto D(Y)_n$ and $Y \mapsto D(Y)_{n-1}$ are functors, rather obviously, and now $\gamma_{n-1} \circ \beta_n$ is claimed to be a natural transformation from one to the other.) The reason for this is that

- $f\colon X \to Y$ determines a map $|f^{\leq k}|\colon |X^{\leq k}| \to |Y^{\leq k}|$ for every k;
- the preimages under the map $|f^{\leq k}|$ of certain open subsets $V = V_{Y,k}$ and $W = W_{Y,k}$ of $|Y^{\leq k}|$ which we used to set up a Mayer-Vietoris sequence and so to construct β_k and γ_k are precisely $V_{X,k}$ and $W_{X,k}$, open subsets of $|X^{\leq k}|$.

Therefore naturality of the Mayer-Vietoris sequence (remark 6.2.4) applies, and we get the commutativity of (11.3.2). — It follows almost immediately that it suffices to prove our formula

$$(\gamma_{n-1} \circ \beta_n)\langle x \rangle = \pm \sum_{i=0}^{n} (-1)^i \langle d_i x \rangle$$

in the very special case where $Y = \underline{\Delta}^n$ as in example 9.2.4, and x is the unique element in Y_n. Indeed, if Z is any other semi-simplicial set and $x' \in Z_n$ and we wish to know what $\gamma_{n-1} \circ \beta_n$ does to $\langle x' \rangle$, then we observe that there is precisely one morphism

$$Y = \underline{\Delta}^n \longrightarrow Z$$

which takes the unique $x \in Y_n$ to $x' \in Z_n$. So we know what $\gamma_{n-1} \circ \beta_n$ does to $\langle x' \rangle$ if we know what it does to $\langle x \rangle$, by the commutativity of (11.3.2).
Having made the observation, we proceed by induction on n. The case $n = 1$ is the induction start. It is not completely trivial. There are exactly two distinct morphisms $\underline{\Delta}^0 \to \underline{\Delta}^1$ of semi-simplicial sets. We know that they induce the same homomorphism $H_0(|\underline{\Delta}^0|) \to H_0(|\underline{\Delta}^1|)$. It follows with lemma 11.2.7 that they induce the same homomorphism

$$H_0(D(\underline{\Delta}^0)) \longrightarrow H_0(D(\underline{\Delta}^1)).$$

Therefore $\langle d_0 x \rangle - \langle d_1 x \rangle$ represents zero in $H_0(D(\underline{\Delta}^1))$, and so it must be in the image of

$$\gamma_0 \circ \beta_1\colon\ D(\underline{\Delta}^1)_1 \cong \mathbb{Z} \ \longrightarrow\ D(\underline{\Delta}^1))_0 = \mathbb{Z} \oplus \mathbb{Z}.$$

That can only happen if the generator $\langle x \rangle$ of $D(\underline{\Delta}^1)_1 = \mathbb{Z}$ is taken to

$$\pm(\,\langle d_0 x \rangle - \langle d_1 x \rangle\,)$$

by $\gamma_0 \circ \beta_1$. This takes care of the case $n = 1$.
For the induction step we assume $n > 1$. The inductive assumption tells us what

$$\gamma_{n-2} \circ \beta_{n-1}\colon D(\underline{\Delta}^n)_{n-1} \to D(\underline{\Delta}^n)_{n-2}$$

is, up to sign. It follows by direct computation that the element

$$\sum_{i=0}^{n}(-1)^i\langle d_i x\rangle \ \in D(\underline{\Delta}^n)_{n-1}$$

is in the kernel of $\gamma_{n-2}\circ\beta_{n-1}$. Therefore it is in the image of $\gamma_{n-1}\circ\beta_n$, since $H_{n-1}(D(\underline{\Delta}^n))\cong H_{n-1}(|\underline{\Delta}^n|)=0$ by lemma 11.2.7. But this can only happen if the generator $\langle x\rangle$ of $D(\underline{\Delta}^n)_n=\mathbb{Z}$ is taken to $\pm(\sum_{i=0}^n(-1)^i\langle d_i x\rangle)$ by $\gamma_{n-1}\circ\beta_n$. □

Proof of theorem 11.1.1. Write ∂_n for the differentials in $C(Y)$, and ∂'_n for the differentials in $D(Y)$. By lemma 11.3.1, we have $\partial'_n=a_n\cdot\partial_n$ where $a_n\in\{-1,+1\}$. This is meaningful because $D(Y)_n=C(Y)_n$ for all n. An isomorphism u of chain complexes from $C(Y)$ to $D(Y)$ can be defined by $u_0=\mathrm{id}\colon C(Y)_0\to D(Y)_0$ and

$$u_n=a_n a_{n-1}\cdots a_2 a_1\cdot\mathrm{id}\colon C(Y)_n\to D(Y)_n$$

for $n>0$. □

Remark. The undetermined sign in lemma 11.3.1 is slightly annoying. It can be determined with more work (to me, more annoying). Without a doubt it is also possible to re-set some basic definitions in such a way that the undetermined sign turns out to be always $+$. There are a few places where we had a choice of sign: notably in defining the boundary operator of the Mayer-Vietoris sequence, but also in choosing the order of V and W in the proof of lemma 11.2.7.

11.4. Semi-simplicial sets and the foundations of homology theory

There is a construction (a functor) which to a topological space X associates a semi-simplicial set SX. This is quite important in the more standard treatments of homology theory, even though in some of those standard treatments it does not appear explicitly.

Definition 11.4.1. The semi-simplicial set SX has

$$SX_n = \text{ set of continuous maps from } \Delta^n \text{ to } X$$

for $n\geq 0$. Face operators are given by composition. More precisely, if

$$f\colon\{0,1,\dots,m\}\to\{0,1,\dots,n\}$$

is monotone injective, and $\sigma\in SX_n$, meaning that $\sigma\colon\Delta^n\to X$ is a continuous map, then we let $f^*\sigma:=\sigma\circ f_*\in SX_m$, where $f_*\colon\Delta^m\to\Delta^n$ is the “linear” map determined by f.

In most cases SX is huge. For example, if $X=S^1$, then clearly each of the sets SX_n is uncountable. But for every X there is a comparison map

$$\kappa\colon|SX|\longrightarrow X\ .$$

It is defined in such a way that for every $\sigma\in SX_n$, the composition of characteristic map $c_\sigma\colon\Delta_n\to|SX|$ with $\kappa\colon|SX|\to X$ agrees with the map $\sigma\colon\Delta^n\to X$ itself. In many cases the map $\kappa\colon|SX|\to X$ is a homotopy equivalence. (It is not a homotopy equivalence in all cases because there are topological spaces which are not at all homotopy equivalent to the geometric realization of any semi-simplicial set. In fact this is the only source of trouble. If X *is* homotopy equivalent to the geometric realization of some semi-simplicial set, then $\kappa\colon|SX|\to X$ is a homotopy equivalence. But the proof will not be given here.)
Even though $\kappa\colon|SX|\to X$ is not a homotopy equivalence in all cases, it can be shown that

the homomorphism $H_n(|SX|) \to H_n(X)$ induced by κ is always an isomorphism, for all $n \geq 0$. (That proof will not be given here either.) Therefore

$$H_n(X) \cong H_n(|SX|) \cong H_n(C(SX))$$

using theorem 11.1.1. The right-hand expression, $H_n(C(SX))$, is the definition of the n-th homology group of X given in many standard treatments. This means that chain complexes and their homology are prominent from the very beginning in those treatments. The standard simplices and the maps $f_*\colon \Delta^m \to \Delta^n$ are also prominent from the start. Semi-simplicial sets need not make an explicit appearance, because it is easy to describe the chain complex $C(SX)$ is without explaining what a semi-simplicial set is.